Systems & Control Transactions

Volume 5

Proceedings of the 36th

European Symposium on Computer Aided Process Engineering

ESCAPE 36

PSE PRESS

Solomon F. Brown
Maria Papathanasiou
David Bogle
Joan Cordiner
Michael Fairweather

Systems & Control Transactions

Vol 5

Proceedings of the 36th European Symposium on Computer Aided Process Engineering

ESCAPE 36

June 21–24, 2026
Sheffield, United Kingdom

PSE PRESS

PSEcommunity.org

PSE Press • Hamilton • Notre Dame

Systems & Control Transactions, Volume 5
PSE Press
Hamilton, Ontario, Canada

Individual articles copyright © 2026 by the authors and licensed to PSEcommunity.org and PSE Press.
Remaining text © 2026 by the PSE Press. All rights reserved.
Cover photo credit: Photo 119650316 © Scanrail | Dreamstime.com

This book and all individual articles are released under the Creative Commons CC-BY-SA 4.0 License
Credit must be given to creator and adaptations must be shared under the same terms.
See https://creativecommons.org/licenses/by-sa/4.0/ for license details.

Living Archive for Process Systems Engineering (LAPSE) Archive ID: LAPSE:2026.0200
Digital Object Identifier (DOI): 10.69997/sct.161019

Library and Archives Canada Catalogue
A catalogue record and copy of this book is available from Library and Archives Canada

Archival Permissions
This book, as well as the individual articles within, may be deposited in digital research archives, personal archives, or institutional archives.

Electronic Book
ISBN (Volume 5) 978-1-7779403-5-5
ISSN (Series): 2818-4734

10 9 8 7 6 5 4 3

The information presented in this book was provided by individual contributors in the form of new scientific research contributions, data, and findings that were believed to be reliable and provided in good faith. Some, but not all, content has undergone peer review as marked. Neither PSE Press, PSEcommunity.org, the University of Sheffield, the EFCE CAPE Working Party, the CACHE Foundation, Inc., the editors, nor authors guarantee the accuracy or completeness of any information provided, nor shall they be responsible for errors, omissions, incomplete information, or damages arising from its use. The information within does not consist of engineering or professional services.

Typeset in Inter

Suggested Citation for Individual Articles
Student MC, Supervisor QR. Title of the article. *Systems & Control Transactions* 5:startpage-endpage (2026). https://doi.org/10.69997/sct.restofarticledoi

Suggested Citation for Entire Proceedings
Brown SF (ed). Proceedings of the 36th European Symposium on Computer Aided Process Engineering, Sheffield, United Kingdom, June 21-24 (2026). ISBN: 978-1-7779403-5-5

Scientific Editors
Solomon F. Brown (University of Sheffield), Maria Papathanasiou (Imperial College London), David Bogle (University College London), Joan Cordiner (University of Sheffield), Michael Fairweather (University of Leeds)

PSE Press Production Team
Thomas A. Adams II (NTNU), Mina Naeini (Jacobs Consultancy), Noelia Gudiño, Daniel R. Lewin (Technion)

Table of Contents

PREFACE

Introduction ..... **xxiii**
Solomon F. Brown

Peer Review Policy ..... **xxiv**

International Scientific Committee ..... **xxvi**

PEER REVIEWED CONFERENCE PROCEEDINGS

Section 1: CAPE in Circular Economy

Circular Zero Liquid Discharge Systems with Renewable Energy Integration: A Technoeconomic Assessment **2**
Fatima Mansour, Sabla Y. Alnouri, Sabah Solim, Ali Al-Sharshani, and Dhabia Al-Mohannadi

Development of a Novel Microwave-assisted Process that Converts Mixed Plastic Waste to Olefins and Aromatics **11**
Aseel Al-Sakkaf, Chunlin Luo, Yuxin Wang, and Srinivas Palanki

Process-Informed Design of Electrochemical Cells for Urea Production: A Techno-Economic and Systems Engineering Approach **18**
Zhimian Hao, Shilong Fu, Chengtian Cui, Ruud Kortlever, Ruud van Ommen, and Ana Somoza-Tornos

Computed-Aided Design of an Intensified Process for the Sustainable Production of Biodiesel from Waste Cooking **25**
Tania G. Salgado-Rodríguez, Fernando I. Gómez-Castro, and Nelly Ramírez-Corona

Integrated Multiproduct Facility for the Green Production of Chemicals and Food from Apples **31**
Vanessa Villazón-León, Carlos Sanz, Adrián Bonilla-Petriciolet, and Mariano Martín

Hybrid Modeling of a Sewage-Sludge Gasifier using Flowsheet Simulation and Machine Learning **37**
Malte Lutz, William Würpel, Fabian E. Habicht, Burcu Aker, and Jan C. Schöneberger

Lifetime-Adjusted LCA of Biochemical and Thermochemical Circular Plastic Pathways **45**
Alexandra Krestnikova and Gonzalo Guillén-Gosálbez

Optimization of Circular Supply Chains for Electric Vehicle Batteries **52**
Kaapo Kopra and Iiro Harjunkoski

Integration of carbon dioxide capture in a wine effluent biorefinery through the use of deep eutectic solvents . **59**
Carlos E. Guzmán Martínez, Valeria Caltzontzin Rabell, Sergio I. Martínez Guido, Salvador Hernández, and Claudia Gutiérrez Antonio

Techno-economic feasibility of gallium recovery from semiconductor wastewater **68**
Kilian Kozerke, Thomas Eberius, and Nathanial J. Cooper

Integrated solvent and process design with technoeconomic and lifecycle assessment for solvent-based recycling of end-of-life vehicle plastics **77**
Riccardo Standish, Jian Yin, Jakob Burger, Mirjana Minceva, Hannah Mangold, Christian Holtze, Markus Schoerner, Bernhard von Vacano, Amparo Galindo, George Jackson, and Claire S. Adjiman

(article retracted) 87

Advancing Circularity in Biopharma: Leveraging Industrial Symbiosis for Resource Efficiency 94
Miriam Sarkis, Andrea Bernardi, Cleo Kontoravdi, and Maria M. Papathanasiou

Beyond Decarbonization: Quantifying Circularity in Energy System Planning 100
Javiera Vergara-Zambrano and Styliani Avraamidou

Hybrid Modelling of Segmented Flow Extraction Process for Digital Twin Development in Critical Metals Recovery 108
Arun Pankajakshan, Konstantinos Katsoulas, Malik Olasinde, Cong Chao, Eric S. Fraga, Panagiota Angeli, and Federico Galvanin

Life Cycle Modeling towards Regional Symbiosis for Valorizing Mixed-Lignocellulosic Biomass from Agriculture and Forestry 117
Yasunori Kikuchi, Nobuhide Takahashi, and Satoshi Ohara

Design and Assessment of Regional Symbiosis: A Case Study of Plant-oil Production in Japan 125
Yukito Watanabe, Yuichiro Kanematsu, Hiro Tabata, Heng Yi Teah, Kousuke Hiromori, and Yasunori Kikuchi

OpenAD-lib: Open-Source Framework for Uncertainty-Aware Anaerobic Digestion Digital Twins 131
Benaissa Dekhici, Rohit Murali, and Michael Short

Achieving Net-Zero Emissions in Industrial Parks Through Optimized Symbiotic Exchanges and Carbon Capture Utilization 143
Ricardo N. Dias, Fátima N. Serralha, Carla I.C. Pinheiro

(article retracted) 152

Techno-Economic Assessment and Optimisation of Self-Sufficient Biomethane Systems for Regional Decarbonisation 161
Meshkat Dolat, Benaissa Dekhici, and Michael Short

Dynamic material flow analysis of iridium circularity in proton exchange membrane water electrolysers in Japan 169
Shu Yamaki, Ayumi Yamaki, Heng Yi Teah, Yuichiro Kanematsu, and Yasunori Kikuchi

Value-Based Assessment for Strategic Selection and Optimization of POME Valorization Pathways 175
Yiwei Gao, Heng Yi Teah, Yuichiro Kanematsu, and Yasunori Kikuchi

Chemical Additives in Plastics: Understanding the Reactions, Fate, and Releases during Pyrolysis 183
Ronald Borja-Roman, Andres Castellar-Freile, John D. Chea, Monica Rodriguez Morris, Gerardo J. Ruiz-Mercado, and Kirti M. Yenkie

Banana peel as a Source of Polyphenols: Perspectives of Sustainability 192
Kajal Chaudhary and Soumyajit SEN GUPTA

A Whole Systems Thinking Model Towards Optimal Decarbonization Strategies for China's Cement Sector 198
Yushu Wang, Wenli Du, Minglei Yang, and Vassilis M. Charitopoulos

A Decision-Support Framework for Process Design of Sustainable Aviation Fuel Production via Integrated Biorefineries 205
Vibhu Baibhav, Daniel Florez-Orrego, and Francois Maréchal

Modeling standardized industrial profiles for the optimization of eco-industrial parks 215
Dilafruz Mavlanova, Marianne Boix, Rachid Ouaret, and Stephane Negny

Understanding Environmental Impacts of Lithium-Ion Battery Recycling 222

Marco Vaccari, Leonardo Tognotti, and Monica Puccini

Accelerating Design of Chemical Recycling of Plastic Waste through Digitalization: A Bubbling Fluidized Bed Reactor Case Study 229
Stefano Iannello, Vassilis M. Charitopoulos, and Massimiliano Materazzi

Principal Component Analysis (PCA) for Evaluation of Fatty Acid Monoalkyl Ester (FAME) Quality towards Sustainable Biodiesel Production from Indonesian Microalgae Strains 236
Dea Nilam Mustika and Dimitrios I. Gerogiorgis

Section 2: CAPE in Clean Energy Systems

A MIBLP model for a Northern European negative-emission hydrogen supply chain with CCS in the North Sea 244
Matthias Maier, Sungho Shin, Simon Roussanaly, and Thomas A. Adams II

Integration of exergy and economic optimization for green hydrogen and power co-generation based on sorbent-enhanced biogas reforming with CO2 capture 252
Arthur-Maximilian Báthori and Calin-Cristian Cormos

Green Hydrogen Supply Chain Design Towards Social Sustainability: A Case Study in Brazil 259
Leonardo Santana, Fernando Pessoa, and Ana Barbosa-Póvoa

Sustainable Design of an Integrated Seawater-Based Green Hydrogen Production Process 267
Antonio Torres-Ayala, Eduardo Sánchez-Ramírez, Marcelino Carrera-Rodríguez, and Juan Gabriel Segovia-Hernández

Net Carbon Balance (NCB): a Better Way to Evaluate and Optimize Carbon Capture Technologies 274
André F. Young, Aline R. Eckstein, Leonardo D. Ferreira, Vítor M. Sermoud, and Ingrid A. de Oliveira

Multi-Scale Design for Clean Energy Systems: Industrial Electrification and Flexible Operation of Ammonia Synthesis 282
Nicholas N. Kalamaris and Christos T. Maravelias

Optimizing Renewable Energy Storage Systems to Accelerate Sustainable Data Center Deployment 289
Matthew J. Palys and Prodromos Daoutidis

Temporal aggregation bias in model-based Direct Air Capture performance under weather variability 297
Eleni Chalasti, Gbemi Oluleye, Maria M. Papathanasiou, and Ronny Pini

Dynamic Optimization of an Adsorption Heat Storage to satisfy the Heat Demand of a House 306
Alix Untrau, Lorenz T. Biegler, and Sabine Sochard

Direct iron reduction system analysis with mixture of hydrogen feed.....314
Marwa Mortadi, Carl Haikarainen, Guanwei Zhou, and Henrik Saxén

Optimization of Site-wide Heat-Integrated Utility Systems with Heat Pumps using MILP..... 320
Thorben Hochhaus, Marcus Grünewald, and Julia Riese

Investigating the Effects of Heat Ingress and Tank Motion on the Ullage Space of a Partially Filled Liquid Hydrogen Tank Using CFD 328
Anna Pakarinen and Anders Brink

Development of a methodology for heat pump-based heat integration in batch processes 334
Johannes Wloch, Marcus Grünewald, and Julia Riese

Pareto-Optimal Pathways for Refinery Decarbonization: Retrofit of Small Modular Nuclear Reactors..... 342
Aditya S Khatu, Sampriti Chattopadhyay, and Ana I Torres

Exploring the Thermal Coupling of Solid Oxide Electrolysis and Ammonia Synthesis: A Plantwide Energy Integration Assessment 352
Alessandra Macchi, Federica Longobardi, Régis Anghilante, Andrea Isella, and Davide Manca

Techno-Economic Assessment of Decarbonization Pathways for Methanol and Formaldehyde Production: A Superstructure Optimization Approach 360
Rafailia Mitraki, Muhammad Salman, and Grégoire Léonard

Integrated Operating Strategies and Parameter Optimization for PEM Electrolyzers in Power-to-X Energy Systems 369
Luka Bornemann, Yifan Wang, and Martin Kaltschmitt

Evaluating the Potential of Sustainable Aviation Fuel for Decarbonization of the Aviation Sector: An Agent-based Model 378
Geeta Joshi, Tejeswi Ramprasad, Harmandeep Singh, Narayanan Rajaraman, Vikrant Urade, Arnoud Higler, and Rajagopalan Srinivasan

A Data-Driven Optimization Framework for the Design and Operation of Adaptive and Resilient Energy Supply Chain Networks under Uncertainty 386
Halil Iseri, Funda Iseri, Mahmoud El-Halwagi, Eleftherios Iakovou, and Efstratios Pistikopoulos

Powering AI Beyond the Grid: Optimal allocation and Behind the Meter Investment Portfolios for Data Centers396
Mohamed Abdelhady, Eleftherios Iakovou, and Efstratios N. Pistikopoulos

Evaluating the potential of e-fuels for decarbonizing European truck transport: A techno-economic and life cycle approach 404
Marion Andritz, Severin Sendlhofer, Rafailia Mitraki, Grégoire Léonard, and Christoph Markowitsch

Integration of computer aided design and emerging technology development based on a series of scale-up demonstration tests; Case study of thermal energy storage 413
Shoma Fujii and Yasunori Kikuchi

Optimizing Heat Storage Integration for Solar Thermal Systems in Industrial Process Heat Networks 419
Håvard Falch, Henri Tande Espen, and Rahul Anantharaman

Energy planning towards absolute environmental sustainability: identifying key demand-side sufficiency levers to stay within planetary boundaries using sensitivity analysis tool 424
Nicolas Ghuys, Diederik Coppitters, Anne van den oever, Mahdi Kchaou, Hervé Jeanmart, and Francesco Contino

Assessing the potential of vehicle-to-grid (V2G) systems using dynamic simulation and life cycle assessment 432
Ayumi Yamaki and Yasunori Kikuchi

Techno-economic assessment of green ammonia plants with multi-scale capacity 440
Ruitao Sun and Jie Li

Terawatts for Petabytes: Exploring the impact of AI data centres on Europe's net zero goals 446
Mohammad Hemmati and Vassilis M. Charitopoulos

Simulation of Methanol Production from Biogas: Impact of Feedstock Composition and Stoichiometric Number Adjustment 454
Muhammad Zulkefal, Magne Hillestad, Truls Gundersen, and Bjørn Austbø

Strategic Design of CO2-Reuse Pathways for Sustainable Aviation Fuel: A Game-Theoretic Techno-Economic Analysis 462
Andrés I. Cárdenas, Víctor A. Soria, and Ana I. Torres

Feasibility of Integrating Sugarcane-Derived Biogas into the Allam–Fetvedt Cycle for BECCS Power Generation 470
João Pedro B. Vasconcelos, Miguel D. Carvalho, and Jean F. Leal Silva

Optimising Waste-to-Energy Power Generation in Trinidad and Tobago 477
Sherard Sadeek, Thomas Hart, Arun Mangra, Daniel Chernick, Andrew Ross, Dhurjati Prasad Chakrabarti, Maria M. Papathanasiou, and Keeran Ward

Harnessing waste heat in the optimal operation of power-to-X energy systems using detailed process models 483
Yifan Wang, Luka Bornemann, and Niklas von der Assen

Comparative Life Cycle Assessment of Electrochemical and Conventional Regeneration Pathways in KOH-Based Direct Air Capture Systems 492
Georgia Ioanna Prokopou, Zoi Drakopoulou, Dominik Bongartz, and Alexander Mitsos

Comparative Techno-economic and Environmental Evaluation of Single-Step vs. Dual-Step CO2-to-Methanol Processes using Multiobjective Optimization 500
Biswarup Mondal, Johannes Leipold, and Achim Kienle

A Multi-Objective Optimization and Superstructure-Based Decision-Support Tool for Regional Low-Carbon Hydrogen Roadmaps: Methodology and Application to a region of Spain 507
Silvia Moreno, Alejandro Aragón-García, Ángel L. Villanueva-Perales, Bernabé Alonso-Fariñas, and Pedro Haro

Electrified refineries in the Power Flow Network 514
Sampriti Chattopadhyay, Ana I. Torres, Ignacio E. Grossmann, and Saif R Kazi

Techno-economic Analysis of Alternatives for Carbon Capture and Utilization and Green Ammonia Production from a Cement Plant Flue Gas 526
Miguel A. Pedro, Ana S. Amorim, and Henrique A. Matos

CFD-based optimal design of a portable and stackable alkaline water electrolyser for hydrogen production ... 534
Akepogu Venkateshwarlu, Gianluca Li-Puma, and Brahim Benyahia

Modeling and experimental validation of a flat-conduit dense-phase receiver for concentrated solar power ... 543
Mustapha Hamdan, Malak Hamdan, Bogdan Dorneanu, and Harvey Arellano-Garcia

Section 3: CAPEing with Uncertain Futures

The Value of Multi-Stage Stochastic Programming in Power Grid Capacity Expansion Planning 550
Sergio Bugosen, Tomas Valencia, Jean-Paul Watson, Chrysanthos E. Gounaris, and Carl D. Laird

Designing in an Unpredictable World: Novel Methods for Uncertainty Characterization, Quantification, and Optimization in Process Engineering 558
Diederik Coppitters, Antoine Laterre, Mahdi Kchaou, Kevin Verleysen, Panagiotis Tsirikoglou, Jerome Stock, Matthias Weigold, Konstantinos Kyprianidis, Ward De Paepe, and Francesco Contino

Enhancing Interpretability of Stochastic Programming Solutions: A Multiparametric Approach 567
Parth Brahmbhatt and Styliani Avraamidou

Exploring Robust Early-Stage Decisions in Energy Transitions Using Near-Optimal Pathways and Multi-Armed Bandits 574
Mahdi Kchaou, Diederik Coppitters, and Francesco Contino

Optimization-based Design, Simulation and Data-Driven Learning for Resilient Manufacturing Systems 583
Miriam Sarkis and Efstratios Pistikopoulos

A Machine Learning Implementation for Fermentation Quality Prediction in Wine Manufacturing 592

Matthew A.J. Hill and Dimitrios I. Gerogiorgis

Global Optimization of Robust AC OPF **602**
Yuhui Yin and Vassilis M. Charitopoulos

Section 4: Pharmaceutical and Biotechnological Systems

Uncertainty-Aware Model Validation Framework for Pharmaceutical Process Development **611**
Kensaku Matsunami, Yash Barhate, and Zoltan K. Nagy

Comparison of Centralised and Decentralised Pharmaceutical Manufacturing Paradigms: An Agent-Based Simulation Study **618**
Farshid Babaei, Mohammad Salehian, David Robins, Cameron J. Brown, Daniel Markl, Alastair J. Florence, and Solomon Brown

Developing predictive models for batch cooling crystallization of APIs with limited data availability **625**
Mauro Davanzo, Emanuele Tomba, Enrico Carlassare, Riccardo Motterle, Massimiliano Barolo, Zoltan K. Nagy, and Fabrizio Bezzo

An in silico/in vitro approach for uncertainty-aware hybrid models for template-induced protein crystallisation systems **631**
Daniele Pessina, Jerry Y. Y. Heng, and Maria M. Papathanasiou

Capturing mixing effects on aggregation kinetics of monoclonal antibodies during viral inactivation **640**
T. Marella, F. Cenci, P. Thompson, M. Muhieddine, and F. Bezzo

A Generative AI Approach to Inverse Design for Continuous Pharmaceutical Manufacturing **648**
Consuelo Del Pilar Vega-Zambrano and Vassilis M. Charitopoulos

Development of Symbolic Regression-Based ATR-FTIR Calibration Models **655**
Fernando A. R. D. Lima, Inga S. Nordhus, Marcellus G. F. de Moraes, M. Enis Leblebici, Argimiro R. Secchi, Mauricio B. de Souza Jr., and Idelfonso Nogueira

Beyond Solid-Phase: Comparative Assessment of Liquid Phase Oligonucleotide Synthesis with Single- and Dual-Stage Diafiltration **664**
Alberto Saccardo, Rachel Ha, Zoe Fang, and Benoît Chachuat

In silico solvent selection for green and cost-effective pregabalin crystallisation **672**
Matthew Blair and Dimitrios I. Gerogiorgis

Molecular Similarity Coefficient in Chemical Design and Analysis **681**
Youquan Xu, Zhijiang Shao, Abdulelah S. Alshehri, Mansour S. Alhoshan, and Anjan K. Tula

Pareto Front Guided Sampling for Efficient Bioprocess Experimentation **692**
Stricker Samuel, Lucas Francisco dos Santos, Claus Wirnsperger, Alessandro Butté, Antonio del Rio Chanona, Mehmet Mercangöz, and Gonzalo Guillén Gosálbez

Sensitivity-Based Comparison of Resource Competition Models for Optogenetic Gene Circuit Design **702**
Pratham Kapavarapu, Satyajeet S. Bhonsale, Simen Akkermans, and Jan F.M. Van Impe

Genome to Production: A Multiscale Model for Bioprocess Design **709**
Rajiv Kailasanathan, Mohammad Reza Boskabadi, Abhishek Sivaram, and Seyed Soheil Mansouri

Modelling & optimization of recombinant protein production in a microbial cultivation with tunable induction **714**
Philipp Pably, I Gede Eka Perdana Putra, Gerd Seibold, Jakob K. Huusom, and Julian Kager

Automatic kLa determination in stirred tank reactors by model-based design of experiments **721**
Ana Helena V. Caetano, Krist V. Gernaey, and Julian Kager

Decarbonizing API Manufacturing: Conceptual Design and Scale-up Analysis of Continuous-Flow Electrosynthesis for Ibuprofen Production **727**
Tuse Asrav, Merlin Alvarado-Morales, and Gürkan Sin

A Multi-Level Hybrid EKF–Machine Learning Soft Sensor for Robust Bioprocess Monitoring **733**
Mohammad Reza Boskabadi, Rajiv Kailasanathan, Luis Ricardez-Sandoval, and Seyed Soheil Mansouri

Understanding the Impact of Ribbon Splitting on Tablet Properties Using a Hybrid Mechanistic–Machine Learning Framework **742**
Shumaiya Ferdoush, Mohammad Shahab, Xinle Zhang, Jayden A. Pierce, Emma Jeffries, Adaugo Ufomba, Zoltan K. Nagy, Gintaras V. Reklaitis, and Marcial Gonzalez

Physics-Informed Neural Networks for NIR Spectroscopy Analysis of Pharmaceutical Tablet Properties **751**
Xinle Zhang, Shumaiya Furdoush, Marcial Gonzalez, and Gintaras V. Reklaitis

Probabilistic design spaces from small DoEs – A boundary-focused workflow using quantile surrogates **756**
Tobias Overgaard, Emmanouil Papadakis, Maria-Ona Bertran, and Maria M. Papathanasiou

(article retracted) **763**

Combined PBM-PBPK Modeling for Optimized Integrated Oral Solid Dosage Form and Dosing Strategy Design **769**
Meng-Hua Yang, Francesco Rossi, Gintaras V. Reklaitis, and Zoltan K. Nagy

Experiments & Modelling of Batch Fermentation of Fusarium venenatum on Glucose-Fructose Mixtures **776**
Tom Vinestock and Miao Guo

Towards Digital Threads for FAIR, Trustworthy, and Human-Centric Bioprocess Development **783**
Jonas M. Karsten, Ernesto C. Martínez, and Mariano N. Cruz Bournazou

Safe and Sustainable by Design Pharmaceuticals through Combined Computer-Aided Retrosynthesis, Techno-Economic Analysis, and Life Cycle Assessment **791**
Shang Gao and Brahim Benyahia

Section 5: Modelling and Simulation

An Open-Source IDAES Framework for Simulating Inductively Heated Adsorption Processes **802**
Sudip Sharma and Thomas A. Adams II

Development of a Predictive Model for Microbial Growth under Variable Conditions Using a Multilayer Perceptron Neural Network: Application to Candida guilliermondii **810**
Jazmín Cortez-González, Juan Gabriel Segovia-Hernández, Salvador Hernández, Varinia López-Ramírez, Arturo Hernández-Aguirre, and Rodolfo Murrieta-Dueñas

Dynamic optimization of glucose feed in cell cultivation for monoclonal antibody production process design balancing productivity and impurity generation **816**
Kosuke Nemoto, Yuki Yoshiyama, Mizuki Morisasa, Junshin Iwabuchi, Yusuke Hayashi, Sara Badr, and Hirokazu Sugiyama

Development of ANN-based models for dye removal through electrochemical advanced oxidation techniques **822**
Zaira J. Mosqueda-Huerta, Oscar D. Lara-Montaño, Juan Manuel Peralta-Hernández, and Fernando I. Gómez-Castro

Model Screening and Identifiability Analysis of Methanol Synthesis Kinetics: Information-Guided Evaluation of Operating Conditions **828**
Eblagh Ahmad, Biasin Alberto, Nardi Luca, and Federico Galvanin

A Process Modeling Approach for Water and Energy Optimization in Geologic Hydrogen Extraction 838
Caroline Kaitano and Thokozani Majozi

Optimal Simulation of an Electrodialysis Reactor for the Desalination and Regeneration of Multi-Ionic Wastewater 846
Vicent Ayala-Andreu, Miguel A. Montiel, Vicente Montiel, and Juan A. Labarta

A Data-Efficient Symbolic Regression Framework for Automated Interpretable Bioprocess Modelling 855
Luca Riezzo, Alexander Rogers, Harry Kay, and Dongda Zhang

Modelling Pressure Effects in Boiling Brazed Aluminum Heat Exchangers: A Software Comparison 861
Hamza Karim, Rim Khodr, Rodolphe Sardeing, and Gaëtan Becker

From Drift to Adaptation to the failed ML model: Transfer Learning in Industrial MLOps 870
Waqar Muhammad Ashraf, Talha Ansar, Fahad Ahmed, Jawad Hussain, Muhammad Mujtaba Abbas, and Vivek Dua

Design of a Chemical Heat Pump based on Methylcyclohexane, Toluene and Hydrogen 879
Rajalakshmi Krishnadoss, Félix Le Bot, and Thomas A. Adams II

An Integrated Process of Multi Effect Distillation Based Desalination with Renewable Energies: Evaluation of Power Generation Efficiency and Freshwater Production Cost 886
Mohammed Adam, Mudhar A. Al-Obaidi, and I. M. Mujtaba

Dynamic Modelling of Renewable-driven CO2 Methanation using Recurrent Neural Networks 895
M. Andrea Pappagallo, Diego A. Romero Lombo, Mattia Vallerio, and Emanuele Moioli

Chemical Language Transformers for the Inverse Design of Novel Surfactants 903
Alexander W. Rogers, Ruediger Zillmer, Amanda Lane, Adam Kowalski, and Dongda Zhang

A Modeling Framework Integrating Data Trends and Reference Information for Predicting Temperature-Dependent Thermophysical Properties 910
Shuai Zhang, Abdulelah S. Alshehri, Mansour S. Alhoshan, and Anjan Tula

Exploring Molecular Pretraining and Mechanism-Aware Modeling for Reaction Yield Prediction 919
Yongkyu Lee, Won Bo Lee, and Lauren Ye Seol Lee

Synergistic integration of direct air capture in bioenergy systems 927
Nor Syuriaty Jaafar, Norhuda Abdul Manaf, Noor Fatina Emelin Nor Fadzil, and Nilay Shah

An Adaptive Framework for Robust Energy Forecasting under Concept Drift and Feature Uncertainty 934
Francesco Marcato, Alessio Santecchia, Manuel Ruivo de Oliveira, Francesco Silvestri, and Rafael Castro-Amoedo

Analysis and comparison of technologies for the regeneration of a capture solution in DAC absorption systems 943
Grazia Leonzio and Nilay Shah

From Plastic Waste to Platform Chemicals: Aspen Plus Modeling of Polystyrene Conversion Through Hydrothermal Processing into Value-added Chemicals 950
MohammadSina HajiHashemi, Corinna Schulze-Netzer, and Thomas A. Adams II

Uncertainty Prioritisation for Water-Energy-Food-Land Nexus Optimisation 958
Md Shamsul Alam, I. David L. Bogle, and Vivek Dua

Data Transformation Techniques and its Influence in Hybrid Model Performance 964
Juan Federico Herrera-Ruiz, Carlos Eduardo Robles-Rodriguez, Cesar Arturo Aceves-Lara, Javier Fontalvo, and Oscar Andrés Prado-Rubio

Simulation and analysis of carbon capture process using piperazine for large scale biomass-fired power plant972
Shengyuan Huang, Olajide Otitoju, Yao Zhang, and Meihong Wang

Differentiable Programming for Cyclic Adsorption Processes 978
Alex Glover, Maria M. Papathanasiou, and Ronny Pini

Bayesian Optimization Framework for Agrochemical Formulation Design 986
Yipei Zhao, Robin Wesley, and Joan Cordiner

Variational Bayesian Neural Networks for Modelling and Uncertainty Quantification in Bioprocessing..... 992
George Spencer, Harini Narayanan, Claus Wirnsperger, Alessandro Butté, Cleo Kontoravdi, and Maria M. Papathanasiou

Multi-objective simulation-based optimisation of pharmaceutical process systems 1001
Artemis Tsochatzidi, Francesca Cenci, Magdalini Aroniada, and Lazaros G. Papageorgiou

Modeling, Simulation, and Optimization of an Anion Exchange Membrane Cell for Ammonia Electrolysis 1007
Laureen E. Hernández Lefrán, Leonardo A. Cáceres Avilez, Antonio E. Bresciani, Claudio A. Oller do Nascimento, and Rita M. Brito Alves

A Coarse-Graining Algorithm for Complex Chemical Reaction Networks 1014
Yi Tao and Tong Qiu

Automated Construction of Bayesian Networks of Chemical Process for Dynamic Risk Assessment 1022
Kai Yin, Hao Kang, and Jinsong Zhao

CO2 Conversion: Three-Dimensional Modelling of Gas Diffusion Electrodes 1030
Cristina González-Fernández, Camilo Peralta, Jose Antonio Abarca, Esther Santos, Guillermo Díaz-Sainz, and Ángel Irabien

(article retracted) 1036

Multiscale Modeling of PHBV Production: Explicit Polymerization Modeling and Improved Prediction of Chain Length Distributions 1043
Stefan Hempfling, Rudolph Kok, Stefanie Duvigneau, Achim Kienle, and Robert Dürr

Evaluating Extrapolation of Modular Hybrid Process Models for Pilot-Scale Batch Separation Processes..... 1050
Søren Villumsen, Jakob K. Huusom, Xiaodong Liang, and Jens Abildskov

Exploiting Input-Space Separation in Kolmogorov–Arnold Networks to Prevent Catastrophic Forgetting in Industrial NIR Systems 1058
Imam M. Iqbal, Isabell Viedt, and Leon Urbas

Renewables to X: Micro-Reactor Pathways towards Methanol and Dimethyl Ether Production 1064
David T. Hren and Andreja Nemet

Dynamic modeling of fouling development during dead-end filtration of dusty superheated steam 1073
Felipe de Oliveira, Wijtze Nijhuis, Marcel Meinders, and Edwin Zondervan

Robust Design of Transient Flow Experiments for the Identification of Kinetic Models in Flow Reactor Systems with Catalyst Deactivation 1080
Jinwen Cui and Federico Galvanin

Process modelling and multi-objective optimisation of solid sorbent-based direct air capture 1088
Toluleke E. Akinola and Meihong Wang

Uncertainty Quantification of Stochastic Gene Expression 1095
Francisca Pizarro Galleguillos, Satyajeet S. Bhonsale, and Jan F.M. Van Impe

Practical Identifiability and Optimal Experiment Design for Hybrid Cybernetic Models: An E.Coli Case Study . 1102
Stylianos Floros, Satyajeet S. Bhonsale, Simen Akkermans, and Jan F.M. Van Impe

Semi-Supervised Generative Augmentation Improves Surfactant Surface Tension Prediction from Limited Experimental Data 1111
Gabriela C. Theis Marchan, Kyle Territo, and Jose A. Romagnoli

A CFD Analysis of Dielectric Fluid Performance in Thermal Management of Li-ion Cells 1119
Margarita G. Correa-Ibarra, Jorge A. Alfaro-Ayala, Jose de J. Ramirez-Minguela, Zeferino Gamiño-Arroyo, and Agustin R. Uribe-Ramirez

Simulation of Fixed-Bed Reactor System for Combined Ca–Cu Chemical Looping with Integrated Combustion and CO2 Capture 1127
Levente Biró, Norbert-Botond Mihály, and Ana-Maria Cormos

Comparison of Various Hydrogen Flux Trajectories in a Catalytic Membrane Reactor Operating Dehydrogenation of Ethylbenzene to Styrene..... 1135
Nabeel S. Abo-Ghander

A Neural Model of Pinch-Based Multicomponent Distillation for Applications in Flowsheet Synthesis..... 1145
Alexander B. Wolf, Mirko Skiborowski, and Jakob Burger

Physics-informed Graph Neural Networks to Predict Thermodynamically Consistent Activity Coefficients in Multicomponent Mixtures 1153
Lifeng Zhang, Benoît Chachuat, and Claire S. Adjiman

SMILE: Smell Maximisation In Low-cost Eau de parfum 1160
Flora Esposito, Ulderico Di Caprio, Mattia Collu, Raffaele Graziano, Vincenzo Guida, Hasan Sildir, Idelfonso B.R. Nogueira, Florence Vermeire, M. Enis Leblebici

Dynamic Operation of a Haber-Bosch Loop with Quench-Cooled Converter for Power-to-Ammonia Systems 1167
José M. Pires, Diogo A. C. Narciso, Carla I. C. Pinheiro

A Symbolic Regression-based approach for Modeling Fouling Resistance in Heat Exchangers..... 1175
Fernando A. R. D. Lima, Antonioni B. Campos, Bruna Carla G. de Assis, Livia Pereira L. Costa, Fabio S. Liporace, Mauricio B. de Souza Jr., and Argimiro R. Secchi

Energy Integration Via Heat Pump in a Simulated Fluidized TSA Column for CO2 Capture from Biomass-Derived Flue Gases..... 1183
Eduardo S. Funcia, Yuri S. Beleli, Enrique Vilarrasa-Garcia, Marcelo M. Seckler, José L. Paiva, and Galo A. C. Le Roux

Towards White-box Environmental and Economic Process Optimization: Tailoring Modelling Approaches to Multi-scale Simulations..... 1189
Thomas Hietala, Sonja Herres-Pawlis

Designing MgCl2-Based Ethanol Dehydration Systems: A Multi-Objective Approach with Open-Loop Controllability 1197
Josué J. Herrera Velázquez, J. Rafael Alcántara Avila, Salvador Hernández, and Julián Cabrera Ruiz

Addressing Matrix Effects Through A Physical Prior-Informed Calibration Model For Quantitative Analysis ...1205
Onur C. Boy, Ulderico Di Caprio, Idelfonso Nogueira, and M. Enis Leblebici

Multi-Objective CAPE Simulation of Agro-Industrial Systems Integrating High-Yield Sugarcane and the Inversion Process..... 1212
Satoshi Ohara, Yoshifumi Terajima, Hiro Tabata, and Yasunori Kikuchi

Towards the Decarbonization of a Conventional Ammonia Plant by the Gradual Incorporation of Green Hydrogen and Air-Separated Nitrogen **1219**
João Fortunato, Diogo A. C. Narciso, and Henrique A. Matos

Hybrid Modeling of Wastewater Treatment Dynamics Using Hammerstein-Wiener Structures **1226**
Arne Tirez, Niels Stevens, Dominik Bongartz, and José Matias Assumpcao

Re-parametrisation of NRTL model for C1+ organics and alcohols recovery from aqueous phase in pyrolysis oil production **1234**
Matteo Gilardi, Filippo Bisotti, Trung Trinh, and Bernd Wittgens

A Framework based on Population Balance Modeling for Predicting Li–O2 Battery Discharge and Life Cycle Behavior **1242**
Nadia G. Khouri, Jean F. Leal Silva, Letícia M. S. Barros, Viktor O. C. Concha, and Rubens Maciel Filho

A Computational Framework for Simulation and Energy Evaluation of Sustainable Biodiesel Production Routes **1249**
Ian B. B. Batata, Emílio E X. Guimarães Filho, Victor H. S. Ramos, Maria R. Wolf Maciel, Nadia G. Khouri, and Rubens Maciel Filho

(article retracted) **1256**

Enhancing Pharmaceutical Supply Chain Robustness via Simulated Annealing **1262**
Nelson Chibeles-Martins, Maria A. Monge, and Tânia Pinto-Varela

Assessing the Impact of Thermodynamic and other Modelling Choices in MEA-based CO2 Capture Simulations **1269**
Hassan Khaled Hassan Baabbad, Alberto Fernández, Fèlix Llovell, and Carlos Pozo

Experimental and Kinetic Study of Iron Oxide Reduction in a Fixed Bed Reactor using a Dynamic Shrinking Core Model **1277**
Emiliano Salucci, Antonio D'Angelo., Vincenzo Russo, Henrik Grénman, and Henrik Saxén

A General Framework for Model Recognition in Chemical Reactor Systems Using Artificial Neural Networks Classifiers **1285**
Emmanuel Agunloye, Asterios Gavriilidis, and Federico Galvanin

Predicting Flow Regimes in a Wiped Film Evaporator Using the Volume of Fluid Method **1292**
Gonçalo V.L. Pardal and Fernando P. Bernardo

Modelling of fouling dynamics in a falling-film evaporator **1298**
Johanne L. Christensen, Lukas S. Theisen, Kevin Feldmann, and Jakob K. Huusom

Modeling Slug Flow Dynamics in Offshore Wells using Universal Differential Equations **1306**
Gustavo Luís Rodrigues Caldas, Giovani Gerevini, Fábio C. Diehl, Idelfonso B. R. Nogueira, Maurício B.de Souza Jr., and Argimiro R. Secchi

Estimation of Thermodynamic Properties for Cellulosic Biomass-Derived Compounds: Application to Heat and Work Balances in Process Simulation **1318**
Anthony D. Anastasi, Cornelius M. Masuku, Praveen Ravikumar, Shishir P.S. Chundawat, and Diane Hildebrandt

Municipal Solid Waste Valorization into Chemical Solvents for Industrial Symbiosis: Techno-Economic and Environmental Assessment **1328**
Oktay Boztas, Daniel A. Flórez-Orrego, Meire E. G. R. Domingos, and François Maréchal

Energy Baseline Surrogates for Modular Reactors from Generated Recipe-Based Process Data **1336**
Shreyas Parbat, Greeshmanth Rajanala, Isabell Viedt, and Leon Urbas

Evaluation of dual pressure low-temperature distillation for LNG Production in CO2-rich fields 1345
Victor S. V. Mercado, Dirceu Noriler, Laura Plazas Tovar, Radin Suhaib Salihuddin, Amiza Bt Surmi, Fadhli Hadana Rahman, and Jean F. Leal Silva

Multi-scale Metabolic Modeling and Simulation 1353
Peter E. Carstensen, Teddy Groves, Lars K. Nielsen, Ulrich Krühne, Krist V. Gernaey, and John B. Jørgensen

Enhancing Parameter Identifiability in Capacitive Deionization: A Model-Based Design of Experiments Approach 1360
Yuxuan Yang and Federico Glavanin

Process Intensification for LNG Purification: Modeling CO2 Separation in a Rotating Packed Bed.....1366
Alexander A. Zerwas, Bruna L. V. Maia, Wilson Santos Neto, Radin Suhaib Salihuddin, Amiza Bt Surmi, Fadhli Hadana Rahman, Jean F. Leal Silva, and Dirceu Noriler

Desing and optimization of a multi-objective plant to obtain the best furfural derivates 1374
Melanie Coronel Muñoz, Brenda Huerta Rosas, Eduardo Sánchez Ramírez, Juan Gabriel Segovia Hernández, and Juan José Quiroz Ramírez

Multi-Level Optimization of Crane Scheduling 1383
Sophia Onyshkevych, Bianca Springub, and Christos Galanopoulos

Modelling of carbon dioxide methanation in radial flow reactor..... 1391
Salvatore Capasso, Vincenzo Russo, and Henrik Grénman

Dynamic Modeling of a Biomass Fluidized-Bed Gasifier 1398
Jefferson D. C. Araujo, Fréderic Marias, and Sabine Sochard-Reneaume

Comprehensive Framework for Model Discovery and Discrimination Based on Symbolic Regression and Structural Identifiability - Application to a Partially Observed Chemical Reaction System 1406
Xuming Yuan and Brahim Benyahia

SEMPRE-BIO project: comparison of three innovative scaled up and optimised technologies for biomethane production and its purification.....1414
Filippo Bisotti, Matteo Gilardi, and Bernd Wittgens

Modeling and Simulation of Nitrogen Generation by Pressure Swing Adsorption for Power-to-Ammonia 1422
Marcus J. Schytt, Lorenz T. Biegler, and John B. Jørgensen

Energy recovery from process purges: steam turbine integration and operation optimisation in biogas upgrading within SEMPRE-BIO project.....1431
Filippo Bisotti, Matteo Gilardi, and Bernd Wittgens

ProcessSimulator.jl: A Symbolic-Numeric Open-Source Framework for Process Simulation in Julia Language1439
Vinicius V. Santana, Christopher V. Rackauckas, and Idelfonso Nogueira

Nanoparticle Nucleation and Growth Model Exploration with Perturbative Analysis..... 1445
Stephen T. King, Antonios Armaou, Themis Matsoukas, Griffin A. Canning, and Robert M. Rioux

Sustainable and Optimized Biorefinery Design for the Production of High-Value Catechol Derivatives from Lignin 1455
Alden Paul Rangel-López, Eduardo Sánchez-Ramírez, Maricruz Juarez-García, Jesús Manuel Núñez-López, and Juan Gabriel Segovia-Hernández

Integrated Data-Driven Optimisation of LNG Hot Section for Energy Efficiency and Decarbonization.....1461
Aisha Al-Hammadi, Dr Tareq Al-Ansari, Dr Ahmed AlNouss, and Abdul Aziz Shaikh

Modeling and Optimization of Sonochemical Reactors through simulations 1467

Nikolaos I. Vittas and Antonios Armaou

MCSGP dynamic simulation for peptides separation using Aspen Chromatography 1476
Ivan Chóez-Guaranda, Emmanuel Appiah-Danquah, Bogdan Dorneanu, and Harvey Arellano-García

Multisectorial Energy Integration of Low-Temperature Brewery Process, Manufacturing Industry and District Heating Network 1483
Pullah Bhatnagar, Daniel Florez-Orrego, Oktay Boztas, Meire Ribeiro Domingos, Manuele Margni, and Francois Marechal

Process-Intensified Oscillatory Opposed-Jet Mixers: Mixing Quantification and Operational Guidelines 1494
Sofia P. Brandão, Ricardo J. Santos, Madalena M. Dias, José C. Lopes, and Margarida S. C. A. Brito

Section 6: Concepts, Methods and Tools

A Multimodal Framework Integrating Procedural Texts and Visual Perception for Laboratory Safety Monitoring 1503
Shuo Xu and Jinsong Zhao

Optimal Stopping of Batch Processes with Stochastic Dynamics – A Study of When to Act under Uncertainty 1513
Rafif S. Ramadhan, Luca Grebe, Maximilian Maschmeier, Johannes Pastoors, and Eike Cramer

A Multi-objective Experimental Design Framework Leveraging Hybrid Modelling and Gaussian Process Optimization.....1520
Michael Aku, Solomon Gajere Bawa, Ye Seol Lee, and Federico Galvanin

A Universal Framework for Automated Reaction Network Identification and Interpretable Rate Model Generation 1529
Harry Kay, Alexander Rogers, and Dongda Zhang

Unveiling Reaction Patterns in Thermal and Catalytic Biomass Pyrolysis Using PCA and Multivariate Analysis1539
Martín Rodríguez-Fragoso, Sandro González-Arias, Octavio Elizalde-Solis, and Edgar Ramírez-Jiménez

Optimisation of Synthetic Natural Gas Production via Direct Air Capture and Utilisation using Reduced Models under a Novel Trust-Region Funnel Method 1551
Gul Hameed, Tao Chen, Antonio del Rio Chanona, Lorenz T Biegler, and Michael Short

Control-Guided Reinforcement Learning for Cooperative Energy Management.....1558
Isabela Fons Moreno-Palancas, Raquel Salcedo Díaz, Rubén Ruiz Femenía, José A. Caballero, and Antonio del Río Chanona

New tools, new thinking: Biomimetic Process Design through Parametric Modelling and Simulation.....1565
Alix Saury, Thibaut Houette, Pierre-Emmanuel Fayemi, Jean-Matthieu Cousin, Jérôme Fortin, and Arnaud Dujany

A Comparative Analysis of Sequential Active Learning Approaches: Statistical Design of Experiments versus Bayesian Optimisation.....1573
Daniel V. Batista and Marco S. Reis

Advancing Industrial Fermentation across scales: Model Development, Cost Analysis, and Predictive Control 1582
Marc Lemperle, Pedram Ramin, Julian Kager, Benny Cassells, Stuart Stocks, and Krist V. Gernaey

Predicting Ecotoxicity (HC50) Values Using Symbolic Regression for Transparent Life Cycle Assessment 1592
Abdulhakeem Ahmed, Nitya Kasera, and Ana I. Torres

Adaptive soft sensor to estimate alite fraction in clinker production through quasi-ensemble PLS modelling . 1601
Mihnea Stefan, Wilson R. Leal da Silva, Fabrizio Bezzo, and Pierantonio Facco

Semantic PEA Datasheets for digitalised modular plant documentation **1608**
Sascha Lamm, Sebastian Tecl, Ingo Dietrich, Sissy Sommer, Markus Heinbücher, and Peter F. Pelz

libDIPS: An Open-Source Platform for Global Optimization of Hierarchical Optimization Problems **1617**
Adrian W. Lipow, Daniel Jungen, Aron Zingler, Hatim Djelassi, and Alexander Mitsos

Model verification and Uncertainty Quantification methods using the CCSI simulation model for CO2 capture **1626**
Jessica V. Scheffer, Serena Delgado, Olivier Authier, Valentin Loubière, Franchine Ni, Christophe Castel, and Jean-Marc Commenge

An End-to-End Pure Component Property Prediction Framework Based on a Hierarchical Molecular Fragmentation Method **1634**
Jianfeng Jiao and Jie Li

Global Optimization of a Hydrodealkylation Flowsheet through Spatial Decomposition with SNoGloDe **1643**
Madeline Leppla, Georgia Stinchfield, Norman Tran, and Carl D. Laird

Superstructure Framework for Feasibility and Flexibility Analysis Methods in Modular Plant Design **1650**
Julian Pamperin, Jonathan Mädler, Amy Koch, Isabel Viedth, and Leon Urbas

Process Flowsheet Synthesis via Quantum Reinforcement Learning with Improved Scalability **1659**
Austin Braniff, Fengqi You, and Yuhe Tian

A Unified Multi-Scale TCN Framework for Batch Manufacturing Soft Sensing and Monitoring **1666**
Yee Hung Hong and Zhao Jinsong

Optimizing the Solubility of Organic Molecules in Mixed Solvents Using Bayesian Optimization and Multicomponent Directed-Message Passing Neural Networks **1679**
Simona Buzzi, Ulderico Di Caprio, Dominik Bongartz, and Florence Vermeire

Coupling Analytical Derivatives with Adjoint Automatic Differentiation in a Modular Process Simulator **1687**
Andrés Piña-Martinez and Jean-Marc Commenge

Hand-crafted Feature Fusion for Deep Learning-Based Instance Segmentation in Microfluidics **1696**
Wenle Xu, Lin Sheng, Qichen Shang, Mengqi Liu, Tong Qiu, Kai Wang, and Guangsheng Luo

Methodology to assess the integrity of Water and Energy Integration Systems (WEIS) models using the ThermWatt computational tool **1705**
Miguel Castro Oliveira, Rita Castro Oliveira, and Henrique A. Matos

Targeted Olfactory Molecule Generation for Vanilla Scents Using Generative Flow Networks **1713**
Bruno C. L. Rodrigues, Paul J. Groening, Laura Sisson, Mumin Enis Leblebici, and Idelfonso B. R. Nogueira

A Strategy for Limiting the Effects of Nonconvexities in Mixed-Integer Nonlinear Programming Reformulation of Nonconvex Generalized Disjunctive Programs **1721**
Miloš Bogataj, Chiara Železnik, and Zdravko Kravanja

Beyond Tennessee Eastman: Benchmarking Deep Anomaly Detection on Real-World Pilot-Scale Continuous Distillation Data **1728**
Fabian Hartung, Aparna Muraleedharan, Marius Kloft, and Jakob Burger

From P&ID Drawings to Process Graphs: A Multimodal Language Model Approach **1737**
Baikai Zhu, Samuel Duong, Javal Vyas, and Mehmet Mercangöz

Task-Conditioned Hierarchical Representations for Controllable AI-Assisted Process Synthesis **1746**
Ali Tarik Karagoz, Omar Alqusair, and Jie Li

Recommendation System for Prediction of Adsorption Properties using Kernelized Probabilistic Matrix Factorization **1754**

Gnaneshwar Sampathirao, Sasidhar Gumma, and Nabil Magbool Jan

Optimizing MIP-Heuristics: Generic Formulation and Code **1761**
Sophie Hildebrandt, Meik Franke, Edwin Zondervan, and Guido Sand

A Large Language Model Enhanced Fault Diagnosis Framework for Chemical Processes **1769**
Jingkang Liang and Gürkan Sin

A Unified Python/JAX Framework for Thermodynamic Modeling, Nonlinear Solvers, and DAE Solution of Hydrocarbon Systems **1776**
Carlos C. Sanz and Galo Le Roux

Digital Twin Supported FAIR Electronic Lab Notebooks for Simulated Experiments **1786**
Amy Koch, Isabell Viedt, and Leon Urbas

Nonconvex Robust Optimization for Process Design with Artificial Neural Networks Embedded **1793**
Diego Izquierdo González, Basit Adeogun, Yuhui Yin, and Vassilis M. Charitopoulos

Using Active Learning to Efficiently Calibrate Foundation Models on Raman Spectra in Upstream Bioprocess Fermentations **1801**
Christoph Lange, Ernesto Martínez, Peter Neubauer, and Mariano Nicolas Cruz Bournazou

Utilizing Machine Learning for Phenomena-based Synthesis of Intensified Process Flowsheets **1809**
Omar Alqusair and Jie Li

Separation of Concern Capabilities of Information Model Candidates for Modular Plant System Engineering Lines **1817**
Tobias Kock, Isabell Viedt, Amy Koch, and Leon Urbas

A Framework for Flexible Start/Stop Operation of Electrified Chemical Processes **1825**
Samuel Mercer and Michael Baldea

PhoSim V.0 - Towards A Digital Twin for an Industrial Wet-Process Phosphoric Acid Production **1833**
Ilias Bouchkira, Sanae Elmisaoui, and Abderrazak Latifi

MatStudio: A Human-in-the-Loop Framework for Microstructure Segmentation with SAM-Guided Refinement **1841**
Yao Xue, Yanhu Wang, and Antonios Armaou

Hybrid Multi-Task Learning for Sustainability-Aware Pharmaceutical Molecular Design **1847**
Yiming Ma, Shang Gao, and Brahim Benyahia

Machine Learning-Assisted Multi-PAT Data Fusion for Physics Consistent Crystallization Monitoring **1856**
Yiming Ma, Xuming Yuan, and Brahim Benyahia

Analysis of Ultrasound-Assisted Transesterification for Sustainable Biodiesel Production via Inline Raman spectroscopy **1866**
Ilias Bouchkira and Adel Mhamdi

Evaluating and adapting modelling strategies for data-driven prediction of solvent effects on reaction barriers **1876**
Daeun Shin, Lingfeng Gui, Jonggeol Na, Won Bo Lee, and Lauren Ye Seol Lee

GlycoPy: An Equation-Oriented and Object-Oriented Python Framework for Process Modeling, Optimization and Optimal Control **1884**
Yingjie Ma, Jing Guo, and Richard D. Braatz

Automated workflow for the configuration of modular plants and HAZOP analysis by utilizing DEXPI P&ID **1891**
Mathias Schmitz, Janis Weber, and Norbert Kockmann

(article retracted) 1899

Accelerating Efficient Dimethyl Ether Synthesis through Machine Learning-Based Process Optimization..... 1908
Mitra Jafari, Jefferson Santos da Silva, Wilson Sousa Mercês Neto, Lucas Fonseca Couto, Bogdan Dorneanu, Karen Valverde Pontes, and Harvey Arellano-Garcia

Decomposition of MINLP Formulations in Process Family Design using Progressive Hedging 1916
Ali Asger, Bernard Knueven, and Carl Laird

Physics Constrained Machine Learning for Modeling and Optimization of Chemical Process Systems 1925
Rahul Golder, Bimol Nath Roy, and M. M. Faruque Hasan

Section 7: Process Design, Scheduling and Optimisation

Effect of the feed composition on the performance of a double-dividing wall distillation column.....1933
Carlos E. Guzmán-Martínez, Araceli G. Romero-Izquierdo, Claudia Gutiérrez-Antonio, Salvador Hernández, Massimiliano Errico, and Fernando I. Gómez-Castro

System-Level CO2 Allocation under Supply Constraints in Industrial Clusters 1943
Razan Sawaly, Ahmad Abushaikha, and Tareq Al-ansari

Virtual Plant–Model Pair as a Step Towards Real-Time Optimization of a Simulated Moving Bed System..... 1950
Guilherme C. Amaral, Alexandre F. P. Ferreira, Ana M. Ribeiro, Idelfonso B. R. Nogueira, and Diogo Rodrigues

Rolling-Horizon Scheduling for Dynamic Market-Driven Operation of an Air Separation Plant.....1958
Kieran McKenzie and Christopher L. E. Swartz

Deep Kernel Learning with Kolmogorov–Arnold Networks for Bayesian Optimization 1967
Zhongtao Shang, Zhihong Yuan, Lifeng Zhang, and Yiyang Dai

Machine Learning and Adaptive Sampling Powered Feasible Path Algorithm for Black-box Optimization 1974
Zixuan Zhang, Xiaowei Song, Jiaming Li, Yujiao Zeng, Yaling Nie, Min Zhu, Dongyun Lu, Yibo Zhang, Xin Xiao, and Jie Li

Work and Heat Exchanger Networks as a General Energy-Integration Strategy for Chemical Processes 1986
José A. Caballero, Zinet Mekidiche-Martínez, and Juan A. Labarta

Designing a Load-Flexible Renewable Ammonia Plant for Variable Green Hydrogen Supply.....1992
Niklas Groll and Gürkan Sin

A framework for dynamic rescheduling under disruptions and resource constraints..... 2001
David Robins, Farshid Babaei, Joan Cordiner, and Solomon F. Brown

Optimization of Large-Scale Lycopene Production from Tomato Waste: A Comparative Study of Different Processing Technologies.....2008
Nereyda V. Hernández-Camacho, Fernando I. Gómez-Castro, Mariano Martín, Ehecatl A. del Rio-Chanona, and Oscar D. Lara-Montaño

Auxiliary flexibility in an integrated green steel plant participating in Day-ahead and Intra-day electricity markets 2015
Santeri Vaara and Iiro Harjunkoski

Development of a process modeling library for the design and optimization of beverage production plants... 2023
Valentin Becher, Christian Prommesberger, Ulrike Paap, Anna Afanasev, Anna Bechtold, and Jörg Zacharias

Discrete multi-criteria optimisation of a modular heterogeneous electrolysis system 2029
Hannes Lange, Lukas Furtner, Michael Große, Isabell Viedt, and Leon Urbas

Optimizing Flexible Operation of Grid-Connected Electrolyzers: Storage Capacity as the Key to Economic Viability **2036**
Julian Pamperin, Hannes Lange, Michael Große, and Leon Urbas

Comparative Techno-Economic Assessment of Hybrid-Green Ammonia Layouts for Available-to-Date Decarbonization of the Fertilizer Industry **2045**
Andrea Isella and Davide Manca

Optimizing Steam flux for Energy efficiency in Ammonia Recovery during Sodium carbonate production **2051**
Ediane S. Alves, Mohamad A. Chahine, Denis Guillaume, and Julien Gornay

Structural Constraints with the P-graph Framework: Application to an Ammonia Synthesis Process **2058**
Darrick Hillaby, Andrés Piña Martinez, Jean-François Portha, and Laurent Falk

High Performance Heat Pumps Using Tailored Refrigerants **2066**
Finlay M. Sandham, Andrew Muumbo, Kenneth Mathew, Sarthak Sinha, and Smitha Gopinath

Generative AI for the optimal design of seawater desalination processes **2074**
Valentin ZARLENGA, Antonio ROCHA AZEVEDO, Alvaro MARTINEZ-TRIANA, and Thibaut NEVEUX

Assessing the Impact of Solvent Recycling in Cooling Crystallization using Computer-Aided Molecular and Process Design **2082**
Gaurav Seth, Saman Naseri Boroujeni, Shubhani Paliwal, Amparo Galindo, George Jackson, and Claire S. Adjiman

Reinforcement Learning-driven Process Intensification Synthesis – Design and Optimization of Reaction/Separation Systems **2091**
Dylan Nice, Daniel Wenck Ribeiro, Kristina Savitskaya, Rahul Bindlish, Efstratios N. Pistikopoulos, and Yuhe Tian

A Graph Reinforcement Learning Framework for Batch Process Scheduling in State-Task Networks **2099**
Syu-Ning Johnn, Victor-Alexandru Darvariu, and Vassilis M. Charitopoulos

Two-stage stochastic programming optimization of natural gas pipeline network under cost and carbon emission reduction **2107**
Huiyu Hao and Bohong Wang

Optimization-based design of distillation processes with embedded pressure drop and HETP correlations **2115**
Sina Bertram, Jonas Schnurr, and Mirko Skiborowski

Data-Driven Multi-Objective Optimization of Energy, Environmental, and Economic Performances in Manufacturing with Physics-Consistent Deep Learning **2123**
Hyeonrok Choi, Jaewook Lee, Won Yang, and Seong-il Kim

Joint Optimization of Feedstock Procurement and Production Planning in AD: A Deep Learning-Integrated Stochastic Programming Framework **2131**
Ruosi Zhang and Michael Short

A Method for Uniquely Determining Robust Operating Conditions in Simulated Moving Bed Chromatography **2140**
Kensuke Suzuki, Tomoyuki Yajima, and Yoshiaki Kawajiri

Design of Fluid Distribution Devices Using Topology Optimization **2145**
Osamu Tonomura, Shunya Doi, Naohiro Akashi, and Ken-ichiro Sotowa

An Extended Superstructure Formulation for Non-Isobaric Flowsheet Synthesis **2152**
Harrison A. Fraser, Smitha Gopinath, Jan Sefcik, George Jackson, Amparo Galindo, and Claire S. Adjiman

Superstructure Optimization of CCUS Value Chain: Case Study on Sohar Freezone in Oman **2161**
Shaima Al-Salmani and Muhammad Abdul Qyyum

Design Optimization of Shell-and-Tube Heat Exchangers Under Operational Uncertainty: A Comparative Study Across Three Paradigms **2170**
Fernando Israel Gómez-Castro, Sergio Iván Martínez-Guido, Claudia Gutiérrez-Antonio, and Oscar Daniel Lara-Montaño

Reactor network synthesis of enzymatic cascades using superstructure optimization **2176**
Swastik Chandra, Leandros Paschalidis, Siv Kinau, and Mirko Skiborowski

Designing Multi-Objective Optimization Models for Vaccine Supply Chains: Economic, Environmental, and Social Trade-offs in the COVID-19 Context **2183**
Jonathan Jair Cuevas Lopez, Sofía De-León Almaraz, Alberto A. Aguilar Lasserre, and Catherine Azzaro-Pantel

Multiperiod optimisation of a European CCS supply chain under capture-cost uncertainty. **2191**
José A. Álvarez-Menchero, Rubén Ruiz-Femenia, Raquel Salcedo-Díaz, and José A. Caballero

Optimal Operation of an Alkaline Electrolyzer in an Industrial Setting Using Effective Linearization Techniques **2201**
Jonaed Bin Mustafa Kamal, Loukas Kyriakidis, Saskia Bublitz, Bogdan Dorneanu, and Harvey Arellano-Garcia

Set-based Formulations for the State Task Network Scheduling Problem **2209**
David A. Liñán, Georgia Stinchfield, Carl D. Laird, and Jan Kronqvist

Ammonia as Fuel for Gas Turbines – The Impact of Heat Integrated Partial Decomposition **2218**
Julian Straus, John C. Morud, and Elettra Vantaggiato

Research on Dynamic Scheduling of Multi-line Polyolefin Production Based on Deep Reinforcement Learning **2226**
Zhineng Tao, Tong Qiu, Zhenzhi Gong, Fenglian Dong, Zhi Wei, and Yunlong Guan

Multi-Objective Optimisation of Pressure Swing Adsorption Systems via Symbolic Regression **2234**
Carine Menezes REBELLO, Amilton Barbosa BOTELHO JUNIOR, Anderson Rapello DOS SANTOS, and Idelfonso B. R. NOGUEIRA

GPU-Accelerated Nonlinear Multi-Period AC Optimal Power Flow for Large-Scale Power–Hydrogen Systems **2243**
Geunseo Song, Dirk Lauinger, Sungho Shin, and Jonggeol Na

Foundation Model-Guided Optimization of Chemical Reaction Spaces for Autonomous Experimentation **2252**
Youhyun Kim and Jonggeol Na

Techno-Economic Optimization of Electrified Airports as Collaborative Energy Hubs **2261**
Mohammadreza Babaei, Stavros Vouros, Konstantinos Kyprianidis, and John D. Hedengren

Particle Swarm Optimization for simultaneous design and optimization of heat pumps considering Mixed Integer problems **2272**
Beatriz C. da Silva, Ana M. Ribeiro, Alírio E. Rodrigues, Alexandre F.P. Ferreira, Diogo Rodrigues, and Idelfonso B.R. Nogueira

Techno-economic analysis of hydrogen refueling station with on-site production from a novel blue H2 and N2 production system **2280**
Adrian R. Irhamna and George M. Bollas

Distributed low-carbon hydrogen for freight corridors: siting hydrogen refueling station with onsite production on New England highways **2288**
Adrian R. Irhamna, Burcu Beykal, and George M. Bollas

Superstructure Modelling of Membrane Systems for the Optimization and Flexible Design of Post-combustion Carbon Capture Processes **2297**

Stefania Bempeli and Marina Micari

Multi-scenario Optimization of Groundwater-Sourced Water Production Networks With Daily Well Shutdown Requirements **2303**
Pedro H. CALLIL-SOARES, René P. SCHNEIDER, and Galo A. CARRILLO LE ROUX

Towards the Resilient Design of Power-to-Ammonia Systems via Linear Optimization Tools **2311**
José M. Pires, João Fortunato, Henrique A. Matos, and Diogo A. C. Narciso

Transfer Learning–Enhanced Deep Probabilistic Surrogates for Scalable Multi-Fidelity Bayesian Optimisation in Process Design **2320**
Jaewook Lee, Ethan Errington, and Miao Guo

Optimal Biogas Utilization Planning in a Pig Farm Under Sustainability Indicators **2327**
Jaime David Ponce-Rocha, Martín Picón-Núñez, César Ramírez-Márquez, José María Ponce-Ortega, and Ricardo Morales-Rodriguez

Process design for the recovery of valuable organic compounds from pyrolysis oil aqueous phase **2334**
Matteo Gilardi, Filippo Bisotti, Stefan Schmidt, Rune Myrstad, Camilla Otterlei, Trung Trinh, and Bernd Wittgens

Intensified liquid-liquid process design for critical metals extraction from e-waste **2340**
Konstantinos Katsoulas, Arun Pankajakshan, Malik Olasinde, Cong Chao, Federico Galvanin, Panagiota Angeli, and Eric S. Fraga

(article retracted) **2349**

A novel decomposition-based approach to solve heterogeneous capacitated vehicle routing problems **2355**
Vakil Vamsi Krishna, Mangesh Kapadi, Pankaj Verma, and Shamik Misra

(article retracted) **2362**

Simulation-Optimization vs. MILP Approaches for Real-Time Scheduling of Multiproduct Batch Plants **2370**
Engelbert Pasieka and Sebastian Engell

Design and Optimization of Supply Chain for Citrus Biorefineries: A Regional Approach for Waste Valorization in Brazil **2377**
Marilia G. L. Cavenaghi, Larissa T. Bruschi, and Moises T. dos Santos

Logistics Management of Agri-Industrial Waste for Energy Valorization in Uruguay **2386**
Milena Lagarmilla, Ivan Guchin, Mauro Gambetta, Darío Huelmo, Adrián Ferrari, and Soledad Gutierrez

Section 8: Process Control and Operation

Design and Control of Heat Pump Assisted Distillation Processes for Flexible E-methanol Production **2393**
Lucas A.T. Poker, Marija Saric, Jan Wilco Dijkstra, Vladimir Dikic, and Anton A. Kiss

Long-Cycle Operation for Residue Hydrotreating Processes with Bayesian Optimization **2400**
Pengcheng Zhu, Han Wang, Gang Chen, Bo Chen, Fei Zhao, and Xi Chen

Decentralized Causal Monitoring in High-Dimensional Systems: Revealing the Topological Drivers behind Fault Detection Performance **2406**
Rodrigo Paredes and Marco S. Reis

Control Structure Design of Novel Microwave-Catalyzed Process for Simultaneous Production of Ammonia and Ethylene **2418**
Md Mizanur Rahman, Omar Almaraz, Snehitha Baddam, Jianli Hu, and Srinivas Palanki

Towards Safety-Intelligent Cyber-Physical Systems: A Real-time Monitoring and Control Framework **2425**

Zhane Ann Tizon, Yuanxing Liu, Sahithi Srijana Akundi, Austin Braniff, Beatriz Dantas, Yuhe Tian, Faisal I. Khan, and Efstratios N. Pistikopoulos

Extremum seeking control by perturb and observe applied to dividing wall column pilot 2433
Ivar J. Halvorsen, Bart M. A. Bergers, Giovanni Merlo, Leontine I.M. Aarnoudse, Mark A.M. Haring, and Sigurd Skogestad

Managing Renewable Energy Uncertainty in Green Hydrogen Production Systems 2441
Matteo Lea Casagrande, Andrea Isella, and Davide Manca

Forecasting Time-to-Cyclic Steady State in Periodic Bioprocesses via a Multi-Feature k-Nearest Neighbours Framework 2449
Yasser Algoufily, Foteini Michalopoulou, Maria M. Papathanasiou, and Mehmet Mercangöz

Enhancing Control in Chemical Processes using Reinforcement from Human Feedback 2457
Hilde Gerold, Dean Bradner, and Sergio Lucia

Open-Source Optimization Algorithm for the Simulated Moving Bed Process using CasADi 2466
João Nunes, Ana M. Ribeiro, Alexandre Ferreira, Diogo Rodrigues, and Alírio E. Rodrigues

Connecting the Dots: A Graph-based Approach for Unsupervised Learning and Adaptive Process Monitoring with LLM-assisted Fault Diagnosis 2473
Kyle Territo and Jose Romagnoli

A Hybrid Data-Driven Approach for the Optimization of an Industrial Alkylation Unit 2481
Rastislav Fáber, Karol Lubušký, and Radoslav Paulen

Energy Management of a Renewable-Powered Alkaline Electrolyzer System: A Comparative Study of Nonlinear Optimization Methods 2488
Loukas Kyriakidis, Jonaed Bin Mustafa Kamal, Saskia Bublitz, Bogdan Dorneanu, and Harvey Arellano-Garcia

Real-Time Estimation and Optimal Control of Supersaturation in Sugar Crystallization using Model-based Soft Sensor 2497
Ananya Lohani, Adam Fedor, Július Kurucz, and Radoslav Paulen

Data Reconciliation for Inventory Monitoring in a Petrol Refinery 2505
Jakub Gaborcík, Karol Lubušký, and Radoslav Paulen

Advanced Process Control Structures for Energy-Efficient Downstream Processing in HMF Biorefineries 2511
Norbert B. Mihály, Miruna Prodan, Vasile M. Cristea, and Anton A. Kiss

CMLM: A Cascade of Machine Learning Models to detect and diagnose the performance of model predictive controllers 2518
Elizabeth V. Melo, Argimiro R. Secchi, and Maurício B. de Souza Jr.

Utilization of Additional Equipment Information for Drift Diagnosis in Chemical Plants 2527
Linda Eydam, Julius Lorenz, and Leon Urbas

Reinforcement Learning Supervisory Control with Fuzzy-Logic Reward for Multistage Gas Compression 2534
José R. Torraca Neto, Sergio A. C. Giraldo, Mario C. M. M. Campos, Gustavo L. R. Caldas, Bruno D. O. Capron, and Argimiro R. Secchi

Relating Loss Geometry to Empirical Generalization in Recurrent Neural Net Surrogates: Three Tanks Case Study 2542
Ricardo M. Roxas II and Karl Ezra Pilario

Active-Constraint Regions and Power Distribution in Multi-Stack PEM Water Electrolysis Systems 2551

Marius Fredriksen and Johannes Jäschke

Exploiting the line pack potential of gaseous CO2 pipelines 2558
Archana Kumaraswamy and Johannes Jäschke

Causal Discovery for the Spatial Autoregressive Model: Application to Defect Analysis in the Plastic Injection Molding Process 2565
Ryosuke Tanaka and Koichi Fujiwara

Industrial batch process online fault detection using machine learning..... 2572
Oliver Pennington, Adam Wilson, Carolina Cruz, and Dongda Zhang

Enhancing plasma etching efficiency via physics-based modeling and machine learning 2578
Eneri Boniakou, Yao Xue, Tzannis Vasileiadis, Sotiris Mouchtouris, Katerina Oikonomou, Chloi Zormpa, Antonios Armaou, Vassilios Constantoudis, Evangelos Gogolides, and George Kokkoris

Identification and Self-optimization of Robust Nominal Operating Ranges Using Proximal Policy Optimization2587
Ashish Yewale and Brahim Benyahia

Hybrid Physics-Informed Neural Networks for Thermal Process Identification and Control..... 2594
Sahar Hemmati, Mohammadreza Babaei, and John Hedengren

Section 9: CAPE in Education, Knowledge Transfer and Entrepreneurship

Benchmarking generative AI on fermentation knowledge 2600
Fiammetta Caccavale, Ulrich Krühne, Krist V. Gernaey, and Carina L. Gargalo

Assessing Workflow Automation Platforms in Engineering Education: Towards an Ethical, Technical, and Pedagogical Framework..... 2607
Daniela Galatro, Stuart Grey, and Sourojeet Chakraborty

Artificial Intelligence (AI) Usage in an Undergraduate Chemical Engineering Course: Strengths, Pitfalls, and Future Insights..... 2613
Sourojeet Chakraborty, Stuart Grey, and Daniela Galatro

LLM-Based Intelligent Data Extraction System for Industrial Equipment..... 2622
Zean Chen, Kaicheng Song, Lingyu Zhu, Anjan Kumar Tula, and Xi Chen

The Imperial College Integrated Design Project 2631
Paul S. Fennell, Klaus Hellgardt, and Daniel R. Lewin

Generative AI in Process Design Instruction: A Survey of Students and Faculty 2638
Daniel R. Lewin, Thomas A. Adams II, Dominik Bongartz, Seyed Soheil Mansouri, and Edwin Zondervan

Understanding Student's Preferences for Computational Tools in Chemical Engineering Assessment 2646
Sakiru Badmos

Empirical survey among experts on the relevance of various criteria for optimizing modular electrolysis systems 2652
Hannes Lange, Lucien Beisswenger, Daniel Erdmann, Isabell Viedt, and Leon Urbas

Mapping "Digital Chemical Engineering" in the UK: A Sector-Level Audit of IChemE MEng Curricula 2658
E. Routoula, M. Mohammad Zadeh, M. Granollers Mesa, M. Malekshahian, D. Dikicioglu, M. Pollock, and M. Zandi

Enhancing Robotics and Automation Education Through the Development of Simulation Tool for Material Synthesis 2666
Hsuan Chang, Adedayo Ogunnoiki, and Solomon Gajere Bawa

A Techno-economic Analysis of Simulated Wind Farms 2674

Isaac N. James, Laura Edwards, and Dhurjati Chakrabarti

An Engineering Clinic-Based Approach to Teaching Process Design and Modeling: Bridging Theory and Practice 2680

Barnabas Gao, Thien An Pham, Amarelys Rios, Corbin Tinker, Saugat Bhandari, Robert Hesketh, C. Stewart Slater, and Kirti M. Yenkie

A pedagogical framework for sustainability learning : the case of Industrial Ecology 2689

Marianne Boix, Sydney Thomas, Lea van der Werf, and Ludovic Montastruc

Closing the Digital Gap: A Scaffolded Pathway for Developing Digitalisation Skills in Undergraduate Chemical Engineering Curricula 2696

E Routoula, J Bestenlehner, and M Zandi

INDEX

Author Index ..... Index 1

Keyword Index ..... Index 9

Introduction

Solomon Brown, Maria Papathanasiou, David Bogle, Joan Cordiner, Michael Fairweather (eds.)

The 36th European Symposium on Computer Aided Process Engineering (ESCAPE|36) was held in the UK from Sunday 21 June to Wednesday 24 June 2026 at the University of Sheffield. The organisation was led by members of the University's School of Chemical, Materials and Biological Engineering, in collaboration with the Institution of Chemical Engineers (IChemE) and the Grantham Centre for Sustainable Futures.

Since it began, ESCAPE has become a prestigious platform for the process systems engineering community to connect with experts, share innovative ideas, and explore latest advancements. Spread across four days, researchers, industry professionals, and students, will enjoy a wide range of presentations, posters, plenary speakers, workshops, panel discussions, networking opportunities, and social events.

In an era defined by climate challenges, resource constraints, and rapid technological advancement, the need for sustainable practices has never been more urgent or complex. The main theme of this year's conference, Resilient Sustainability through CAPE, reflects the shift toward integrating resilience through design and the critical role computer-aided process engineering has when developing innovative tools and adaptive, efficient systems that can withstand uncertainty while minimising environmental impact. This conference brings together leading minds to explore how advanced modelling, simulation, optimisation, and data-driven decision-making can drive sustainable transformation across sectors.

ESCAPE|36 received 687 abstract submissions, of which 654 were invited for a 5-to-8 page full paper or a 2-page short paper after a first round review. Upon rigorous re-review of all papers submitted, 331 conference papers have been published in full open access by PSE Press as a volume of the *Systems & Control Transactions* series. Another 162 contributions being made available as short papers separately.

The scientific contributions were organised into oral sessions, poster presentations, workshops and demonstrations, within nine thematic areas:

1. CAPE in Circular Economy
2. CAPE in Clean Energy Systems
3. CAPEing with Uncertain Futures
4. Pharmaceutical and Biotechnological Systems
5. Modelling and Simulation
6. Concepts, Methods and Tools
7. Process Design, Scheduling and Optimisation
8. Process Control and Operation
9. CAPE in Education, Knowledge Transfer and Entrepreneurship

We would like to extend our sincere thanks to all authors, reviewers, topic coordinators, plenary and keynote speakers, and members of the Scientific and Organizing Committees. We gratefully acknowledge the University of Sheffield and the IChemE for their support and express our heartfelt appreciation to our numerous sponsors. This edition marks the second volume with PSE Press as the official publisher of the ESCAPE proceedings (full papers volume) – continuing this important collaboration in promoting high-quality, open access dissemination of PSE research.

We hope this volume captures the scientific excellence, collaborative spirit, and forward-looking vision that defined ESCAPE|36 and continue to serve as a valuable reference for researchers and practitioners in the years ahead.

ESCAPE|36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

Peer Review Policy

TERMINOLOGY

- **Conference Paper:** This is a peer-reviewed research paper describing an **original** research contribution and published in the conference proceedings.

APPLICATION

- All conference papers are subject to peer review.

PUBLICATIONS

- All conference papers which have passed peer review and have been accepted by the conference chairs will be published in the conference proceedings.

PEER REVIEW PROCESS

Initial Screening

The conference chairs will screen the initial submissions (one-page abstracts) of contributed papers. At their option, chairs may choose to reject these submissions for any reason, but the primary reasons for rejection at this stage are (a) subject does not match conference themes and topics; (b) in appropriate, unprofessional, unethical, or dishonest material; (c) duplicate submissions; (d) abandoned submissions; (e) multiple submissions from same presenter; (f) obvious poor quality; (g) reasonable suspicion of machine-generated text; (h) need to limit acceptances because of conference space constraints. Chairs may also transfer papers between sessions at this stage.

Conference chairs will then invite authors of the one-page abstracts that have passed the initial screening to submit a conference paper for peer review.

Conference Paper Peer Review

Conference chairs act as editors. They will assign peer reviewers to assess the quality of the conference papers and manage the peer review process for each paper. Peer reviewers are generally selected from the International Scientific Committee who are scientists, engineers, or researchers with technical expertise in the conference topic. Peer reviewers must be considered technical experts in their field, hold a PhD, and have published peer reviewed scientific works previously. Each work must be reviewed by a minimum of two peer-reviewers.

A “single-blind” peer review system is used, in which the peer reviewers have access to the identities of the authors, but the authors are never given the identities of the peer reviewers. This is consistent with most journal peer-review procedures in our field.

Peer reviewers will be given a set of quality criteria that may include scientific and technical quality, quality of writing, originality and novelty, appropriateness for the conference topic and theme, interest to the community, and other factors of merit. Peer reviewers are expected to be rigorous and critical in their technical assessments and adhere to the highest standards in the field in order to ensure high quality. Peer reviewers are asked to comment on their assigned conference papers and issue a recommendation to the conference chairs based on this quality criteria. Reviewers may recommend one of the following:

- **Accept Paper.** No technical changes are necessary before publication. Only typographical, spelling, grammar, or other minor changes are necessary which do not require technical review.
- **Accept Paper with Minor Revisions.** Some minor technical issues need to be addressed either through technical changes to the manuscript or through rebuttal to reviewer comments. The reviewer does not believe the issues are significant enough to require additional technical review by the reviewers. The reviewer has provided enough commentary such that the editor can decide if the minor issues have been addressed in a future revision.
- **Reject Without Reconsideration.** The technical issues are so significant that it is unlikely that an acceptable

manuscript could be produced by the deadline; or, the manuscript is out of scope, inappropriate, or computer-generated.

Peer reviewers may be asked to provide numerical scores or rankings, as well as provide written comments intended for the chairs and/or the authors. Peer reviewers are encouraged to provide specific and constructive feedback that will aid the authors in improving the work and provide advice to the chairs.

The conference chairs are responsible for making the final decision on each paper and may require several rounds of author changes if necessary to meet quality standards. Conference chairs are not required to follow the recommendations of the peer reviewers in making these decisions. It is possible that conference chairs will choose to not accept some papers that still pass through rigorous technical peer review, especially when limited by available space or when papers do not sufficiently promote conference objectives. Conference chairs may also transfer papers between sessions as desired.

CONFLICTS OF INTEREST AND ETHICAL GUIDELINES

The chairs and peer reviewers must ensure they do not have a conflict of interest that may bias their decisions, such as ensuring that authors are at “arm’s length” and have no financial conflicts of interest. Peer reviewers and chairs must disclose if they have such a conflict of interest with a specific submission, and if so, a different chair and/or peer reviewer should be assigned to handle that submission.

To help determine conflicts of interest, and for all other ethical guidelines, we use the Systems and Control Transactions ethical guidelines described at: https://psecommunity.org/contributor-guidelines

International Scientific Committee

Antonis Kokossis	National Technical University of Athens
Daniel Lewin	Technion - Israel Institute of Technology
Dominik Bongartz	KU Leuven
Grégoire Léonard	University of Liège, Liège, Belgium
Gürkan Sin	Technical University of Denmark
Ludovic Montastruc	INP Toulouse
Miroslav Fikar	Slovak University of Technology in Bratislava
Seyed Soheil Mansouri	Technical University of Denmark
Solomon Brown	University of Sheffield
Vasileios Charitopoulos	University College London
Abderrazak Latifi	Université de Lorraine
Alessandro Di Pretoro	Laboratoire de Génie Chimique de Toulouse
Alexander Mitsos	RWTH Aachen University
Ana Barbosa-Povoa	Instituto Superior Técnico
Andrea Bernardi	Imperial College London
Andrea Galeazzi	Imperial College London
Andrzej Kraslawski	LUT University
Antoine Rouxhet	University of Liège
Antonio Del Rio	Imperial College London
Antonio Sánchez García	Universidad de Salamanca
Artur Schweidtmann	Delft University of Technology
Athanasios Papadopoulos	Centre for Research and Technology Hellas
Bogdan Dorneanu	Brandenburg University of Technology Cottbus-Senftenberg
Borja Hernández	Universidad Rey Juan Carlos
Brahim Benyahia	Loughborough University
Catinca Secuianu	Norwegian University of Science and Technology
Cristhian Molina Fernández	University of Liège
Daniel Alexander Florez Orrego	Ecole Polytechnique Fédérale de Lausanne
Daniel Cristiu	Imperial College London
David Bogle	University College London
Deenesh K Babi	Novo Nordisk
Dimitrios Gerogiorgis	University of Edinburgh
Eike Cramer	University College London
Eleni Routoula	University of Sheffield
Emilia Kondili	University of West Attica
Emmanouil Papadakis	Novo Nordisk
Eoin Syron	University College Dublin
Erdal Aydın	Koç University
Erik Resendiz Mora	University of Sheffield
Ezio Nicola D'addario	AIDIC Sezione Centro
Farshid Babaei	University of Sheffield
Fabrizio Bezzo	University of Padova
Fernando Martins	University of Porto
Ferrasse Jean-Henry	Aix Marseille University
Filip Logist	BASF Antwerpen NV
Filippo Bisotti	SINTEF Industry
Flavio Manenti	Politecnico di Milano
Florence Vermeire	KU Leuven
Francois Marechal	Ecole Polytechnique Fédérale de Lausanne
Gendebien Samuel	University of Liège
Grazia Leonzio	University of Cagliari
Henrique Matos	Instituto Superior Técnico

Hirokazu Sugiyama	The University of Tokyo
Iqbal Mujtaba	University of Bradford
Jan Verstraete	IFP Energies Nouvelles
Joan Cordiner	University of Sheffield
Johannes Jäschke	Norwegian University of Science and Technology
Laura Helleckes	Imperial College London
Lidija Čuček	University of Maribor, Maribor, Slovenia
Luigi Piero Di Bonito	University of Naples Federico II
Manuel Rodrigo	University of Castilla la Mancha
Marco Vaccari	University of Pisa
Maria Papathanasiou	Imperial College London
Maria-Ona Bertran	Novo Nordisk
Marianne Boix	Université de Toulouse
Mariano Martin	Universidad de Salamanca
Marija Tasic	University of Nis
Massimiliano M Villone	University of Naples Federico II
Mathew Wilkes	University of Sheffield
Matteo Gazzani	Utrecht University
Michael Fairweather	University of Leeds
Michael Georgiadis	Aristotle University of Thessaloniki
Mo Zandi	University of Sheffield
Monika Bakosova	Slovak University of Technology in Bratislava
Monika Dokl	University of Maribor, Maribor, Slovenia
Muhammad Salman	University of Liège
Mumin Enis Leblebici	KU Leuven
Niki Triantafyllou	Imperial College London
Olga Arsenyeva	University Paderborn
Paulina Quintanilla	University College London
Philippe Nimmegeers	University of Antwerp
Pullah Bhatnagar	Ecole Polytechnique Fédérale de Lausanne
Radoslav Paulen	Slovak University of Technology in Bratislava
Ranjith Chiplunkar	Imperial College London
Ronny Pini	Imperial College London
Sabla Alnouri	Qatar University
Satyajeet Bhonsale	KU Leuven
Sebastian Engell	TU Dortmund
Sergio Lucia	TU Dortmund University, Dortmund, Germany
Sigurd Skogestad	Norwegian University of Science and Technology
Simon Harvey	Chalmers University of Technology
Smitha Gopinath	University of Sheffield
Somnath Nandi	Savitribai Phule Pune University
Sönke Bröcker	Evonik Operations GmbH
Stavros Papadokonstantakis	Technische Universität Wien
Stefan Radl	TU Graz
Stratos Pistikopoulos	Texas A&M
Theodoros Damartzis	Aristotle University of Thessaloniki
Thomas Adams	Norwegian University of Science and Technology
Ulderico Di Caprio	KU Leuven
Vasile Mircea Cristea	Babes-Bolyai University
Vibhu Baibhav	Ecole Polytechnique Fédérale de Lausanne

Section 1

CAPE in Circular Economy

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Circular Zero Liquid Discharge Systems with Renewable Energy Integration: A Technoeconomic Assessment

Fatima Mansour[a], Sabla Y. Alnouri[b], Sabah Solim[c], Ali Al-Sharshani[c], and Dhabia Al-Mohannadi[d*]

[a] American University of Beirut, Department of Chemical Engineering, Beirut, Lebanon
[b] Qatar University, Gas Processing Centre, College of Engineering, Doha, Qatar
[c] Qatar Shell Research and Technology Center, Doha, Qatar
[d] Hamad Bin Khalifa University, College of Engineering, Doha, Qatar
* Corresponding Author: dalmohannadi@hbku.edu.qa

ABSTRACT

The transition toward circular economy principles in water treatment requires advanced process systems engineering tools to evaluate the trade-offs between environmental sustainability and economic viability, particularly for energy-intensive Zero Liquid Discharge (ZLD) systems. While classic ZLD systems treat concentrated brine as waste, circular ZLD (CZLD) systems incorporate salt recovery technologies that generate marketable salt product. This study presents a comprehensive technoeconomic assessment framework for CZLD systems integrated with renewable energy. The framework is developed to evaluate different CZLD configurations that generate saleable sodium chloride. The assessment methodology integrates solar photovoltaic systems with increasing capacities (100-1400 kW) to analyze renewable energy penetration and energy storage requirements. The renewable energy integration model incorporates hierarchical energy dispatch algorithms prioritizing direct solar utilization, battery storage, and grid backup systems. Solar PV integration demonstrates renewable fractions ranging from 15-100%, with corresponding emissions reductions of up to 72% compared to grid-powered baselines. The technoeconomic analysis reveals a critical trade-off: while renewable integration substantially reduces carbon emissions, it significantly increases operational costs due to higher solar costs compared to subsidized grid electricity. However, configurations with substantial salt revenue generation maintain economic viability even under complete renewable operation. Beyond configuration screening, this work makes three specific contributions to process systems engineering: (i) a hierarchical energy-dispatch model for intermittent solar integration in energy-intensive separations; (ii) a revenue-function formulation that embeds salt valorization into ZLD process economics, shifting the design objective from cost minimization to net revenue maximization; and (iii) systematic identification of the thermodynamic basis for membrane–thermal hybrid superiority through analysis of osmotic pressure versus latent heat energy barriers across 26 configurations. These contributions establish a reproducible framework applicable to brine management in resource-constrained industrial settings.

Keywords: circular water system, zero liquid discharge, zero liquid, resource recovery

INTRODUCTION

Global water stress affects over 2 billion people, with industrial sectors consuming approximately 20% of freshwater resources while generating increasingly saline effluents. Traditional treatment paradigms achieving 70-90% recovery leave substantial volumes requiring disposal, creating environmental liabilities and regulatory challenges. Zero liquid discharge technologies present compelling alternatives by maximizing water recovery beyond 99%, yet implementation barriers persist due to energy requirements reaching 20-80 kWh/m³ and treatment costs spanning 2-15 $/m³ [1].

Recent technological advances enable transformation of ZLD systems through circular economy integration. Salt crystallization and purification technologies

convert waste brines into marketable products, fundamentally altering system economics. Simultaneous renewable energy deployment addresses sustainability concerns, particularly in regions with abundant solar resources coinciding with water scarcity [2]. As a result, ZLD technology adoption has accelerated over the past decade, driven by regulatory tightening and increasing water scarcity.

ZLD process design methodologies have evolved from single-technology studies toward multi-stage hybrid systems. [3] establish that membrane-based ZLD is thermodynamically more efficient than thermal alternatives for feeds below 150, 000 mg/L TDS, while [4] quantify the theoretical minimum energy for ZLD at 3–5 kWh/m³, establishing benchmarks against which practical configurations should be compared. Renewable energy integration with desalination has been extensively reviewed by [2], yet the specific challenge of solar coupling with thermal-intensive ZLD operations remains understudied. [5] show that solar-integrated MED systems operating with thermal energy storage achieve substantially higher renewable fractions than daytime-only configurations, underscoring the importance of storage design in solar-thermal desalination. Salt valorization in ZLD systems is an emerging field with significant economic implications. [6] conduct a techno-economic assessment comparing two ZLD configurations, one producing mixed solid salt and one producing high-purity NaCl, and found that while pure NaCl recovery incurs 1.4× higher energy consumption, it yields 1.52× greater profitability, demonstrating the economic premium of product-grade salt separation. At the ion separation stage enabling such purity, [7] demonstrate that monovalent-selective nanofiltration can preferentially pass monovalent ions while rejecting divalent species, providing a pre-crystallization fractionation pathway. However, systematic frameworks for optimizing these integrated systems remain underdeveloped.

The present work addresses this gap through three distinct contributions to the process systems engineering field. First, a configuration generation algorithm is developed that applies thermodynamic feasibility constraints (including maximum inlet TDS thresholds and process sequencing rules) as an a priori screening step across viable configuration design space, analogous to superstructure-based process synthesis methods but adapted for rapid feasibility assessment. Second, a hierarchical renewable energy dispatch model is formulated that prioritizes direct solar utilization, battery storage, and grid supplementation sequentially, extending established industrial microgrid energy management heuristics to the intermittent solar energy profiles. Third, a multi-criteria performance framework is introduced that applies parallel coordinate visualization to map a six-dimensional performance space (spanning energy consumption, treatment cost, salt revenue, net revenue, carbon emissions, and renewable fraction) across configurations, enabling systematic identification of Pareto-optimal design regions and partitioning of the configuration space into context-sensitive application archetypes. Together, these contributions reframe CZLD design from a single-technology selection problem into a systematic process synthesis and integration challenge, directly advancing circular water treatment systems.

METHODOLOGY

System configuration development

The technoeconomic framework evaluates CZLD configurations through systematic technology combination. Process units span four categories: pretreatment (chemical precipitation, nanofiltration), concentration (reverse osmosis, electrodialysis), thermal processing (multi-effect distillation, multi-stage flash, brine concentrators), and crystallization systems for sodium chloride recovery. Each configuration undergoes mass and energy balance calculations to determine feasibility boundaries.

Feed water characteristics assume 35, 000 mg/L total dissolved solids, representative of industrial brines. Technology sequencing is subject to a technical viability check ensuring that inlet feed concentration does not exceed maximum allowable limit. Sequential staging ensures progressive concentration while respecting operational limits.

Solar integration model

Renewable energy integration employs dynamic allocation algorithms matching solar generation profiles with process demands. The model incorporates direct solar utilization during generation periods, battery storage for load shifting (assumed to be five times the cost premium), and grid supplementation for unmet demands

Solar availability assumes 13.75 hours/day peak generation (Qatar conditions) [8], with photovoltaic costs at 0.06 $/kWh [9] versus 0.035 $/kWh grid electricity [10]. Energy allocation prioritizes direct consumption, followed by stored energy deployment during non-generation periods.

Economic and environmental assessment

Economic evaluation incorporates capital expenditure, operational costs, and revenue streams from recovered salt products. Salt pricing correlates with purity levels: high-grade NaCl (>95%) commands 0.15-0.25 $/kg, medium-grade (90-95%) achieves 0.05-0.15 $/kg, while mixed salts yield 0.02-0.08 $/kg [11]. Net treatment costs account for salt revenues, providing comprehensive economic assessment.

Environmental metrics quantify carbon emissions based on energy sources. Grid electricity generated from natural gas assumes 0.2 kg CO_2/kWh emission factor [12], while solar generation via photovoltaics produces lifecycle emissions of 0.048 kg CO_2/kWh [13].

RESULTS AND ANALYSIS

Energy-economic performance trade-offs

Comprehensive evaluation reveals significant performance variations across CZLD configurations. Energy consumption spans 16.1-89.2 kWh/m³, with membrane-dominant systems achieving highest efficiency while thermal processes maximize salt quality. Figure 1 illustrates the energy-economic landscape, identifying optimal configurations balancing competing objectives.

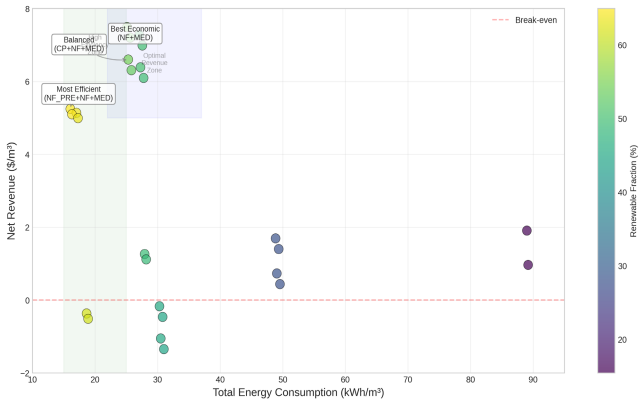

Figure 1. CZLD configuration performance showing energy-economic trade-offs with renewable fraction.

Analysis identifies three distinct performance clusters: (i) high-efficiency membrane systems (16-20 kWh/m³) with moderate revenues (5-6 $/m³), (ii) balanced membrane-thermal hybrids (25-30 kWh/m³) maximizing net revenue (6-7.5 $/m³), and (iii) thermal-intensive configurations (50-90 kWh/m³) with variable economics depending on salt quality. The optimal configuration (NF + MED + Crystallizer + Polish) generates 7.49 $/m³ net revenue at 25.1 kWh/m³ consumption, demonstrating synergies between selective separation and thermal concentration.

Detailed examination of the evaluated configurations reveals systematic performance patterns based on technology combinations. Configurations employing chemical precipitation pretreatment (CP) demonstrate 8-12% higher energy consumption compared to membrane-only pretreatment due to mixing and flocculation requirements. Electrodialysis-based concentration stages (ED) consume 18-22 kWh/m³ for brine concentration, positioning them between reverse osmosis and thermal technologies. The relationship between water recovery rate and specific energy consumption exhibits nonlinear behavior across configuration types.

Table 1 provides a summary of key performance metrics for selected configurations.

Table 1: Performance metrics for selected high-revenue CZLD configurations under base conditions.

Configuration	Energy (kWh/m³)	Total Cost ($/m³)	Salt Rev. ($/m³)	Net Rev. ($/m³)	Emissions (kg CO_2)
NF + MED + Crystallizer	25.1	1.48	8.97	7.49	2.50
RO + MED + Crystallizer	25.6	1.77	8.97	7.20	2.55
CP + NF + MED + Crystallizer	25.3	2.36	8.96	6.60	2.53
NF_PRE + NF + MED + Crystallizer	16.1	3.72	8.97	5.25	1.61

Solar integrations and emissions reduction

Photovoltaic integration demonstrates substantial environmental benefits across all configurations. Renewable fractions vary from 15% for energy-intensive thermal systems to over 70% for more efficient configurations. Figures 2 and 3 compare emissions performance between grid-powered and solar-integrated scenarios, quantifying achievable reductions through renewable deployment.

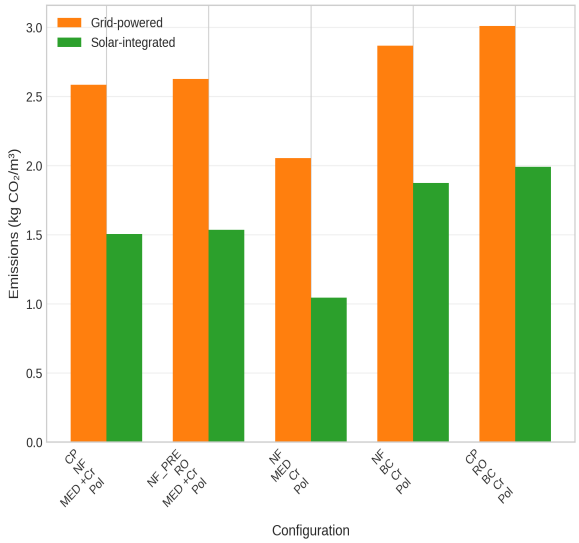

Figure 2. Emission comparison between grid powered and solar integrated configurations.

The selected configurations represent distinct operational archetypes: highly efficient membrane dominant systems (NF_PRE + NF + MED at 16.1 kWh/m³), optimal hybrid configurations (NF + MED + Crystallizer at 25.1 kWh/m³, CP + NF + MED at 25.3 kWh/m³), and more

energy-intensive alternatives with brine concentrators and crystallizers. Emissions reductions range from 1.0 to 1.55 kg CO_2/m^3 across the displayed configurations, representing 38-50% decreases relative to grid-only baselines. Notably, the most efficient achieves the lowest absolute emissions in both scenarios (2.58 kg CO_2/m^3 grid, 1.52 kg CO_2/m^3 solar) but demonstrates smaller absolute reduction (1.06 kg CO_2/m^3) compared to higher-energy configurations with emission performance as follows: 3.01 kg CO_2/m^3 grid, 2.01 kg CO_2/m^3 solar, 1.00 kg CO_2/m^3 reduction. This pattern reflects the dual influence of baseline energy consumption and renewable fraction achieved—lower-energy systems have less emissions to reduce, while higher-energy systems can achieve larger absolute reductions despite lower renewable penetration percentages.

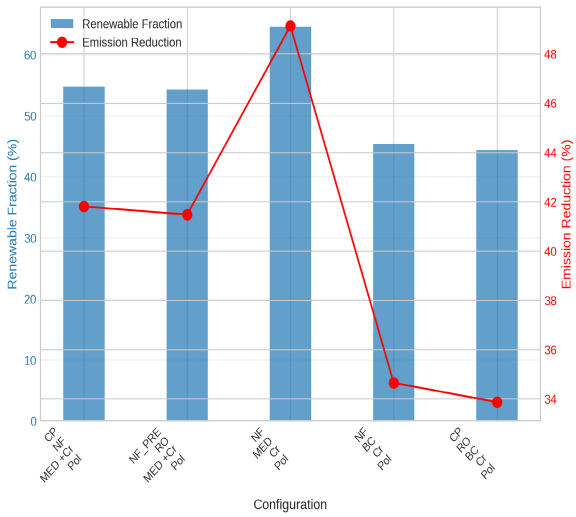

Figure 3. Solar integration performance metrics across selected configurations.

Emission reductions correlate strongly with process energy profiles and renewable fraction achieved. Membrane-based configurations demonstrate highest renewable integration potential (50-65%) due to consistent power demands matching solar generation profiles. Thermal processes, despite higher absolute consumption, benefit from load shifting through strategic battery deployment, achieving 40-55% renewable operation. Economic analysis reveals solar premium of 71% over grid electricity offset by emission reductions exceeding 40% for optimized configurations.

Figure 3 indicates that renewable integration potential is non-uniform with membrane dominant configurations achieving renewable fractions of 54-58%, while more energy-intensive systems demonstrate lower penetration at 44-46% despite identical available solar capacity. The emission reduction trend line demonstrates a striking nonlinear relationship: configurations achieving 54-58% renewable fraction deliver 42-49% emissions reductions, exhibiting near-proportional correlation, while the drop to 44-46% renewable fraction corresponds to disproportionately lower 34-35% emissions reductions. This nonlinearity reflects the combined influence of reduced renewable displacement and the carbon intensity differential between grid and solar lifecycle emissions.

Salt recovery economics

Salt valorization transforms CZLD economics from cost-minimization to revenue-maximization paradigms. Crystallization with polishing generates high-purity NaCl (>95%) commanding an average premium price of 0.25 $/kg, while mixed salt byproducts contribute additional revenue at 0.05 $/kg. Figure 4 illustrates the relationship between salt revenue generation and treatment costs across configuration types. Blue dots represent configurations using chemical precipitation as pretreatment. The orange dots represent configurations using nanofiltration as pretreatment. Green dots are configurations with no pretreatment (direct processing).

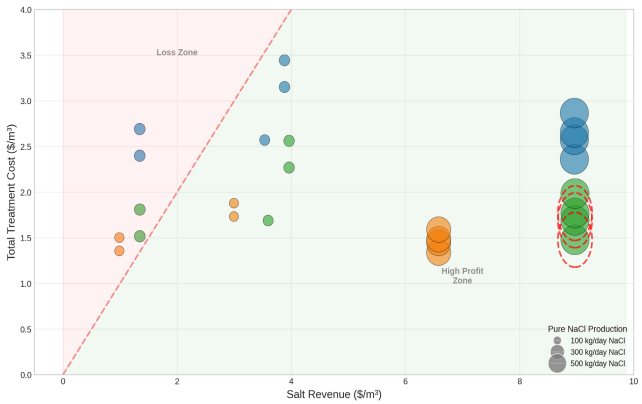

Figure 4. Salt recovery economics showing revenue potential versus treatment costs.

The diagonal clustering of configurations reveals a strong positive correlation between salt recovery capacity and overall economic viability, with configurations generating >7 $/m³ salt revenue consistently achieving positive net revenues of 5-7.5 $/m³. The scatter pattern demonstrates three distinct groupings: (i) low-revenue membrane-only systems (0-2 $/m³ salt revenue) operating at net costs of 2-4 $/m³, highlighting the economic burden of treatment without resource recovery; (ii) intermediate crystallization systems (4-6 $/m³ salt revenue) approaching break-even economics; and (iii) optimized crystallization configurations with polishing (7-9 $/m³ salt revenue) achieving substantial net revenues. The slope of the configuration clusters indicates that each additional dollar of salt revenue translates to approximately 0.85-0.90 $/m³ improvement in net revenue, with the slight offset representing the marginal energy and capital costs associated with crystallization and

purification processes.

Thus, salt revenue can significantly offset treatment costs for optimal configurations. Thermal concentration coupled with crystallization maximizes product quality, generating 468 kg/day pure NaCl and 2, 954 kg/day mixed salts for 1, 000 m^3/day treatment capacity. Market analysis confirms regional demand exceeding 500, 000 tons/year in Qatar alone [14], meaning that the implementation of CZLD systems could eliminate salt imports while advancing circular economy objectives through waste-to-resource transformation.

Thermodynamic efficiency analysis

While energy consumption provides a practical basis for comparing configurations, it conflates work and heat inputs of fundamentally different thermodynamic quality, and does not situate results relative to the theoretical limits of separation. To address this, the present results are contextualized using a second-law framework. [15] establish the minimum specific exergy of separation for a saline feed at approximately 1.06 kWh/m^3 at seawater salinity, representing the thermodynamic lower bound irrespective of technology pathway. Against this benchmark, the best-performing configuration in the present study (NF + MED + crystallizer, 25.1 kWh/m^3) operates at a second-law efficiency of approximately 4.2%, consistent with published values for industrial-scale MED systems, which typically range from 3–6% [16]. Membrane-dominant configurations in this study approach 6–7% second-law efficiency, reflecting the lower irreversibility of pressure-driven versus thermally-driven separation. The observed spread of 4–7% across the 26 configurations confirms that substantial thermodynamic improvement potential remains, primarily through heat recovery within thermal stages and pressure energy recovery in membrane modules. It should be noted that the present analysis treats electrical and thermal energy inputs on a kWh basis without explicit exergy weighting; full exergy accounting, including the chemical exergy of recovered NaCl product streams, would further differentiate configurations by the thermodynamic value of outputs and is reserved for future work.

Sensitivity analysis

Salt price sensitivity is the dominant source of economic uncertainty. For the optimal configuration (NF + MED + NaCl Cr + NaCl Polish), all recovered salt is high-purity NaCl (3, 588 kg/day at 100 m^3/day feed), and the base case assumes the upper bound of the high-grade price range at 0.25 $/kg. This represents an optimistic but commercially justified assumption given that the NaCl Polish step delivers a >95% purity product. Applying the lower bound of the high-grade price range (0.15 $/kg), a 40% price reduction, cuts salt revenue from 8.97 to 5.38 $/$m^3$, reducing net revenue at 100 kW solar penetration from 7.32 to 3.73 $/$m^3$. This remains profitable but represents a substantial compression of economic margin. At the midpoint price of 0.20 $/kg, net revenue settles at approximately 5.05 $/$m^3$, which constitutes a reasonable central estimate. Configurations relying on brine crystallizers, which produce mixed salts at 0.02–0.08 $/kg, are already near or below breakeven at base case pricing and carry no meaningful upside from product purity improvements, making them economically dominated regardless of price scenario. The qualitative ranking of configurations is therefore preserved across the full high-grade price range, though the absolute magnitude of net revenue is sensitive to realized market price.

Energy cost sensitivity is of secondary magnitude but varies with solar penetration level. At 100 kW solar capacity, the lowest penetration modeled, energy cost for the optimal configuration is 0.62 $/$m^3$, representing approximately 38% of total treatment cost of 1.65 $/$m^3$. A 50% increase in the grid tariff from 0.035 to 0.053 $/kWh raises treatment cost by approximately 0.31 $/$m^3$, reducing net revenue from 7.32 to ~7.01 $/$m^3$ at the base case salt price. However, at the lower bound salt price of 0.15 $/kg, this same energy cost increase reduces net revenue further from 3.73 to 3.42 $/$m^3$, illustrating that energy and salt price sensitivities interact: configurations operating at compressed salt revenue margins are disproportionately exposed to energy cost increases. At higher solar penetration levels, grid supplementation constitutes a smaller share of total energy supply and this sensitivity diminishes accordingly. Process efficiency parameters, including membrane rejection rates, thermal process performance ratios, and energy conversion efficiencies, govern energy consumption and water recovery rather than salt product purity, and their influence on net revenue is therefore, for the purposes of this analysis, bounded by the energy cost sensitivity established above. The qualitative configuration ranking, membrane-thermal hybrids with high-purity crystallization outperforming purely thermal or purely membrane alternatives, does not change across the price perturbations examined. A full simulation-based parametric sensitivity analysis across the complete solar penetration range, including systematic quantification of process efficiency uncertainty, is reserved for the future work.

Multi-criteria assessment

Optimal configuration selection requires balancing multiple competing objectives including energy efficiency, economic performance, environmental impact, and operational reliability. Figure 5 illustrates the performance trends of key metrics across different configurations (indicated by the pink dots). The parallel axes display normalized performance metrics spanning energy consumption (kWh/m^3), total treatment cost ($/$m^3$), salt revenue generation ($/$m^3$), net revenue ($/m^3), carbon

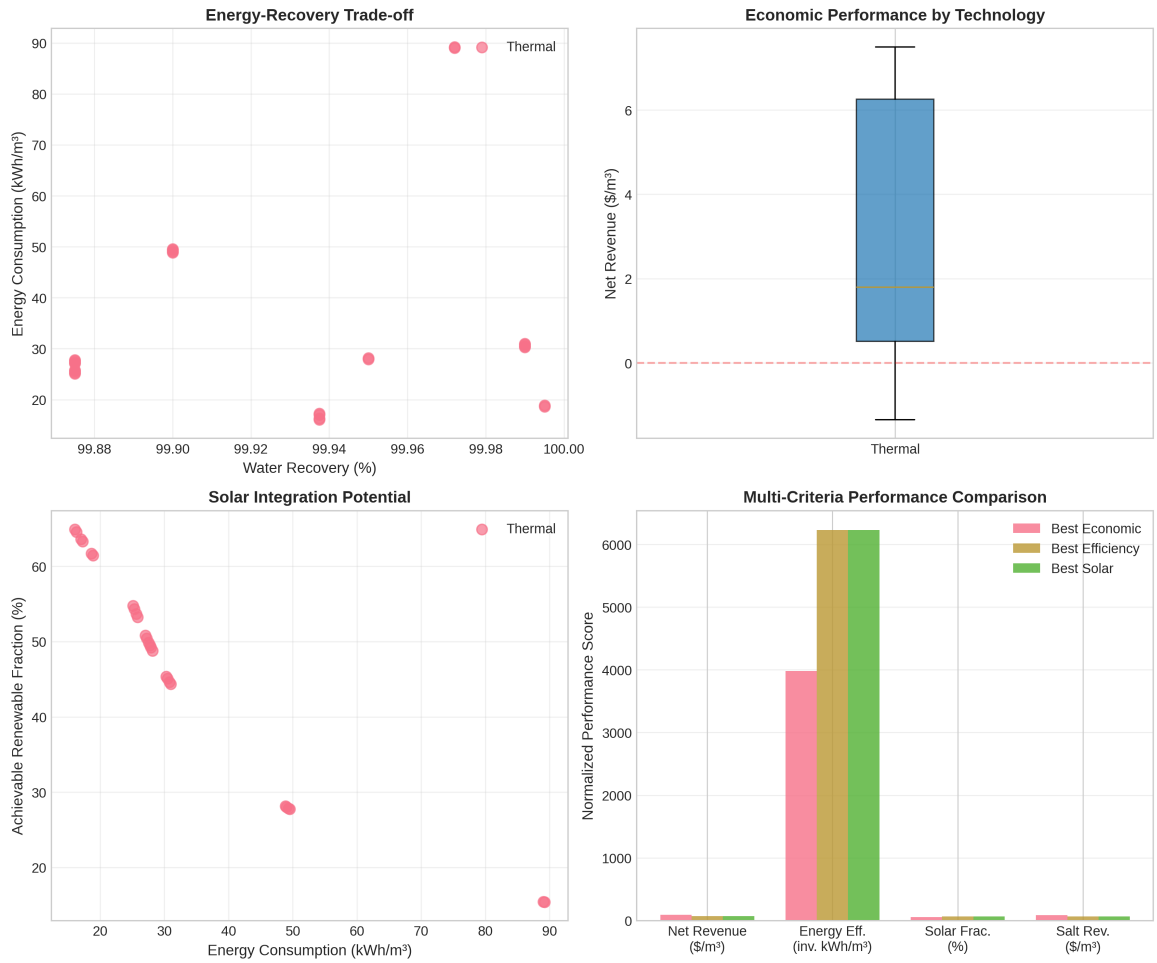

Figure 5. Multi-criteria performance matrix showing (a) energy-recovery trade-offs, (b) economic performance distribution, (c) solar integration potential, and (d) normalized comparison of top configurations.

emissions (kg CO_2/m^3), and renewable energy fraction (%). Each pink line represents a distinct configuration, with line patterns revealing the inherent trade-offs and synergies across objectives. Configurations achieving low energy consumption (left side, 16-25 kWh/m³) typically exhibit moderate salt revenues (6-8 $/m³) and higher renewable fractions (50-65%), while energy-intensive configurations (70-90 kWh/m³) demonstrate variable economic performance depending on crystallization integration.

For revenue maximization, NF + MED + Crystallizer configuration dominates with 7.49 $/m³ net revenue. Energy minimization favors NF_PRE + NF + MED at 16.1 kWh/m³ consumption. Balanced performance emerges from CP + NF + MED configuration, achieving 6.60 $/m³ revenue at 25.3 kWh/m³ with 54% renewable fraction.

The visualization reveals no single dominant configuration across all criteria, confirming the necessity of context-specific optimization based on regional priorities. Notably, the optimal configurations identified in earlier analyses (NF + MED + Crystallizer, CP + NF + MED + Crystallizer) appear as lines maintaining favorable positions across multiple axes, demonstrating balanced performance. The convergence of lines at high salt revenue values (8-9 $/m³) indicates that crystallization-equipped systems achieve similar revenue generation regardless of energy pathway, while divergence in the renewable fraction axis reflects fundamental differences in the energy consumption between membrane-based and thermal-intensive processing routes.

DISCUSSION

The findings of this study highlight several fundamental insights that advance the understanding of Circular Zero Liquid Discharge (CZLD) systems and their integration with renewable energy. The diverse performance observed across configurations underscores that CZLD is not a singular technology but a design space shaped by thermodynamic constraints, process integration strategies, and market economics. This variability reaffirms that there is no “one-size-fits-all” approach; instead,

optimal configurations emerge from careful matching of feed characteristics, separation objectives, and local energy-market conditions.

The technical analysis of the configurations reveals critical technical insights into the fundamental trade-offs governing circular water treatment systems. A key insight is that all configurations achieve similar high-water recovery (>99.5%), but energy consumption varies widely (16-90 kWh/m³). The stark variation in energy consumption directly correlates with the thermodynamic barriers each technology must overcome. Membrane processes operating near osmotic pressure limits require 16-30 kWh/m³ primarily for pressurization, while thermal systems consuming 50-90 kWh/m³ must supply latent heat of vaporization to achieve phase change. The optimal performers (NF + MED + Crystallizer at 25.1 kWh/m³ generating 7.49 $/m³ revenue) leverage synergistic coupling where nanofiltration selectively removes divalent ions at lower energy penalties than reverse osmosis, producing a monovalent-rich stream ideal for MED thermal concentration and subsequent crystallization into high-purity (>95%) NaCl.

Thus, one of the clearest outcomes is the dominance of membrane–thermal hybrid configurations. This matches well-documented industry trends showing that hybridization can reduce specific energy consumption by up to 40–60% relative to purely thermal ZLD systems. From a process integration perspective, the superior performance of hybrid configurations stems from strategic placement of separation barriers such as using membrane processes for initial volume reduction (50-70% water removal at 2-6 kWh/m³) before engaging energy-intensive thermal systems (70-99.9% recovery at 15-70 kWh/m³) minimizes the volume subjected to high-energy treatment, explaining why pure thermal systems (BC, BCr alone at 80-90 kWh/m³) prove economically inferior despite achieving similar recovery rates. To further contextualize these energy differences within a rigorous thermodynamic framework, the observed energy consumption values are benchmarked against the theoretical minimum exergy of separation for the present feed conditions. This technical understanding indicates that future optimization should focus on intermediate precipitation steps to remove sparingly soluble salts before thermal stages, potentially reducing energy consumption by an additional 20-30%. The benchmarking of these energy values against the theoretical minimum exergy of separation for the present feed conditions confirms that all configurations operate at second-law efficiencies of 4.8–7.5%, consistent with published values for industrial-scale membrane-thermal hybrid systems and indicative of substantial remaining thermodynamic improvement potential.

The strong revenue contribution from salt recovery highlights a paradigm shift in ZLD economics. Historically, ZLD has been framed as a compliance-driven, cost-intensive environmental obligation. Introducing resource recovery transforms the system into a productive industrial asset capable of offsetting treatment costs or even achieving positive net revenues. The results indicate that for regions with high salt demand, such as Gulf countries, where industrial salt consumption can exceed several hundred thousand tons annually, CZLD plants have the potential to establish localized salt supply chains, reducing reliance on imports and contributing to national resource security. This reframing supports broader circular economy strategies and diversification of industrial inputs. The sensitivity analysis further confirms that while absolute net revenue is sensitive to realized NaCl market price, the qualitative dominance of high-purity crystallization configurations is preserved across the full reported price range, supporting the robustness of these economic conclusions.

Renewable energy integration further enriches the system design space by coupling water, waste, and energy flows. The analysis shows that solar integration can achieve renewable fractions ranging from 15% to 70%, depending largely on the configuration's load profile. Membrane-based systems, with relatively uniform and lower power requirements, exhibit strong alignment with solar generation availability, leading to higher renewable penetration and smoother load matching. Thermal systems, however, exhibit high-power, nonuniform demand patterns requiring substantial battery buffering to reach similar renewable fractions. This highlights a key operational constraint: the degree to which CZLD systems can decarbonize is not only a function of total energy consumption but also of energy temporal dynamics.

Furthermore, the nonlinear relationship between renewable fraction and emissions reduction reveals the influence of system-scale inefficiencies such as battery cycling losses and inverter inefficiencies. Thus, while renewable energy offers meaningful environmental gains (up to 72% reduction in emissions relative to grid-only operation), these benefits diminish for high-energy thermal systems unless paired with optimized storage strategies. Future research should explore alternative load-management tools such as thermal energy storage, flexible scheduling, or hybrid renewable portfolios that combine solar with other intermittent resources.

The multi-criteria assessment emphasizes the inherent trade-offs in CZLD design. Maximum revenue configurations do not minimize energy use, and lowest-energy systems do not maximize salt recovery or renewable fraction. Decision-making must therefore be context-sensitive, weighing environmental priorities, capital availability, operational risk tolerance, and market conditions. The results suggest three broad application archetypes:

- Energy-Minimizing Designs for regions with limited

power capacity or high energy prices;

- Revenue-Maximizing Designs for regions with high salt demand or favorable salt-market conditions;
- Balanced Hybrid Designs suitable for industrial clusters seeking cost-effective, decarbonized water recovery solutions.

From a systems perspective, CZLD facilities hold strategic potential in industrial symbiosis frameworks, where recovered water, salts, and thermal energy streams can directly feed neighboring industries such as chlor-alkali plants, cooling towers, or district cooling systems. This extends the value chain beyond individual plants into regional industrial ecosystems, supporting broader sustainability and economic diversification goals.

While this study provides a robust basis for CZLD feasibility, real-world deployment requires addressing practical considerations such as brine variability, operational fouling, crystallizer control strategies, and long-term battery degradation. Pilot-scale validation remains an essential next step, especially for hybrid configurations that rely on tight coupling of membrane and thermal processes. Moreover, policy levers such as carbon pricing, renewable energy incentives, or circular-economy mandates could significantly shift the economic balance in favor of CZLD, accelerating adoption in water-stressed regions.

Overall, the transition from conventional ZLD to CZLD reframes industrial brine from an environmental liability to a recoverable resource, directly aligning with circular economy principles. By generating high-purity salts, reducing freshwater dependency, and integrating low-carbon solar energy, CZLD systems close material and energy loops that traditionally remained linear. Moreover, the recovered salt streams can supply local industrial users, enabling industrial symbiosis and reducing reliance on imported materials. These findings demonstrate that CZLD is not simply a treatment technology but a circular value-creation platform that enables economic, environmental, and resource-efficiency gains across the water-energy-materials nexus, if implemented strategically.

CONCLUSION

This research establishes comprehensive technoeconomic assessment frameworks for solar-integrated CZLD systems, demonstrating economic viability through strategic salt valorization. Configuration analyses indicate that membrane-thermal hybrid systems are optimal solutions, generating net revenues exceeding 7 $/m³ while achieving >99% water recovery. Solar integration reduces emissions with achievable renewable fractions of 15-65% depending on process configuration.

Key findings include: (i) multi-effect distillation with crystallization maximizes economic performance at 25-30 kWh/m³ consumption and (ii) high-purity salt recovery profitably offsets treatment costs transforming waste management to resource generation.

Results provide quantitative guidance for industrial practitioners and policymakers advancing circular economy principles in water management. Future research should explore dynamic optimization under uncertainty, pilot-scale validation of promising configurations, and integration with broader industrial symbiosis networks. The framework establishes pathways for sustainable water management in resource-constrained environments through technology integration and business model transformation.

ACKNOWLEDGEMENTS

Please acknowledge your funders in this section, with grant or funding numbers as appropriate. Note that most funding agencies require this.

REFERENCES

1. Panagopoulos A, Giannika V. Decarbonized and circular brine management/valorization for water & valuable resource recovery via minimal/zero liquid discharge (MLD/ZLD) strategies. Journal of Environmental Management 324:116239 (2022). https://doi.org/10.1016/j.jenvman.2022.116239
2. Alawad SM, Mansour RB, Al-Sulaiman FA, Rehman S. Renewable energy systems for water desalination applications: a comprehensive review. Energy Conversion and Management 286:117035 (2023). https://doi.org/10.1016/j.enconman.2023.117035
3. Tong T, Elimelech M. The global rise of zero liquid discharge for wastewater management: drivers, technologies, and future directions. Environ. Sci. Technol. 50:6846-6855 (2016). https://doi.org/10.1021/acs.est.6b01000
4. Davenport DM, Deshmukh A, Werber JR, Elimelech M. High-pressure reverse osmosis for energy-efficient hypersaline brine desalination: current status, design considerations, and research needs. Environ. Sci. Technol. Lett. 5:467-475 (2018). https://doi.org/10.1021/acs.estlett.8b00274
5. Muftah AK, Zili-Ghedira L, Abugderah MM, Hassen W, Becheikh N, Alshammari BM, Kolsi L. Sustainable water production: solar energy integration in multi-effect desalination plants. Water 17:647 (2025). https://doi.org/10.3390/w17050647
6. Panagopoulos A. Beneficiation of saline effluents from seawater desalination plants: fostering the zero liquid discharge (ZLD) approach - a techno-

economic evaluation. Journal of Environmental Chemical Engineering 9:105338 (2021). https://doi.org/10.1016/j.jece.2021.105338
7. Nativ P, Leifman O, Lahav O, Epsztein R. Desalinated brackish water with improved mineral composition using monovalent-selective nanofiltration followed by reverse osmosis. Desalination 520:115364 (2021). https://doi.org/10.1016/j.desal.2021.115364
8. Banibaqash A, Hunaiti Z, Abbod M. An analytical feasibility study for solar panel installation in qatar based on generated to consumed electrical energy indicator. Energies 15:9270 (2022). https://doi.org/10.3390/en15249270
9. IRENA, I.R.E.A., RENEWABLE POWER GENERATION COSTS IN 2023. 2024.
10. Kahrama. *Tariff rates*. [cited 2025 May 20]; Available from: https://www.km.qa/CustomerService/Pages/TariffRates.aspx.
11. Tridge. *Global Salt price*. 2025; Available from: https://dir.tridge.com/prices/salt.
12. US EPA. *Emission factors for greenhouse gas inventories*. 2021 [cited 2025 May 20]; Available from: https://www.epa.gov/sites/default/files/2022-01/documents/emission-factors_jan_2022.pdf.
13. Schlömer, S., et al. Annex III: Technology-specific cost and performance parameters. . in In: Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Edenhofer, O., R. Pichs-Madruga, Y. Sokona, E. Farahani, S. Kadner, K. Seyboth, A. Adler, I. Baum, S. Brunner, P. Eickemeier, B. Kriemann, J. Savolainen, S. Schlömer, C. von Stechow, T. Zwickel and J.C. Minx (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA. 2014.
14. TrendEconomy. Qatar | Imports and Exports | World | Salt, pure sodium chloride, sea water | Value (US$) and Value Growth, YoY (%) | 2012 - 2023. 2023 [cited 2025 July]; Available from: https://trendeconomy.com/data/h2/Qatar/2501
15. Mistry KH, McGovern RK, Thiel GP, Summers EK, Zubair SM, Lienhard JH V. Entropy generation analysis of desalination technologies. Entropy 13:1829-1864 (2011). https://doi.org/10.3390/e13101829
16. Shahzad MW, Burhan M, Ang L, Ng KC. Energy-water-environment nexus underpinning future desalination sustainability. *Desalination* 413:52-64. https://doi.org/10.1016/j.desal.2018.03.009

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Development of a Novel Microwave-assisted Process that Converts Mixed Plastic Waste to Olefins and Aromatics

Aseel Al-Sakkaf[a], Chunlin Luo[a], Yuxin Wang[a], and Srinivas Palanki[a]*

[a]Department of Chemical and Biomedical Engineering, West Virginia University, Morgantown, WV 26506, USA
* Corresponding Author: srinivas.palanki@mail.wvu.edu

ABSTRACT

Plastic waste represents an abundant and underutilized resource that can be converted into valuable products through microwave-assisted pyrolysis. In this research, a novel microwave-assisted processing plant that converts mixed plastic waste to olefins and aromatics is developed and simulated on Aspen Plus (v.14) guided by laboratory-scale experimental data. The experimental results show that at a bulk temperature of 400°C and ambient pressure, approximately 95% of a solid waste plastic feed comprised of equal portions of polypropylene and polyethylene is converted to gases, with nearly two-thirds of the resulting effluent gas composed of olefins. Simulation results show that 2889.1 kg/h propylene, 2088.0 kg/h ethylene and 96.3 kg/h aromatics (benzene and toluene) are produced as main products from 8000 kg/h of mixed plastic feed. High-purity propane and ethane streams were also recovered and sold as byproducts. A technoeconomic analysis is subsequently conducted, revealing that the proposed novel plant achieves a positive net present value of USD 61.6 million at a processing capacity of 8000 kg/h of mixed plastic waste feed. The levelized cost of propylene production is also calculated to be USD 481.1 per ton, estimated while keeping the price of the other coproducts fixed at their average market selling price. Such results confirm the economic feasibility of the microwave-assisted novel plant at the investigated scale.

Keywords: Process Design, Technoeconomic Analysis, Microwave-assisted Heating, Plastic Waste Pyrolysis

INTRODUCTION

Olefins, such as propylene and ethylene, are important feedstocks, used primarily in the manufacture of polymers such as polyethylene and polypropylene. The global production of propylene is projected to reach 191 million metric ton (Mt) by 2030, while ethylene is expected to exceed 407 Mt by the same year [1, 2]. Conventional routes used to produce olefins include steam cracking (SC), which is an energy intensive process that cracks large hydrocarbon fractions, such as naphtha, into smaller molecules, at elevated temperatures of around 800°C [3]. The heat required for the process is typically supplied by the combustion of fossil-based fuels, and industrial steam crackers with gas-fired furnaces are reported to emit globally about 366 Mt of carbon dioxide (CO_2) annually [3]. Statistics show that around 65% of the global propylene production is derived from SC, but the production of propylene from steam crackers is highly dependent on several factors, including the composition of the feedstock [4]. However, with the shifting to ethane-based feed, the amount of the propylene produced through this pathway as a byproduct is declining rapidly, which led to a propylene gap that needs to be bridged by other technologies, such as propane dehydrogenation (PDH) using the OLEFLEX by UOP Honeywell [5]. However, both naphtha cracking as well as PDH utilize fossil fuels and conventional heating methods, which are considered inefficient from an energy point of view and produce significant quantities of greenhouse gases. For these reasons there is current interest in producing light olefins from alternative feedstocks that involve waste plastics and explore alternative heating strategies that are energy efficient to improve process sustainability.

Several researchers have investigated plastic waste upcycling using microwave-assisted heating to convert plastic waste into high-value products, such as olefins

and aromatics, attempting to increase recycling rates and promoting a circular plastic economy [6, 7]. Microwave-assisted heating has emerged as a transformative and advantageous electrified heating method compared to conventional thermal heating routes. Unlike conventional heating methods, microwaves heat the material uniformly, and this volumetric and selective heating enables faster processing times, lowers energy losses hence increases efficiency, and allows precise control over temperature gradients [8].

In this research, a novel microwave-assisted process to convert mixed plastic waste into olefins and aromatics is designed. First, experimental data is obtained in a laboratory-scale microwave reactor that converts a mixture of polypropylene and polyethylene to olefins and aromatics. Then, this data is utilized to develop and simulate the novel process on Aspen Plus (v.14) that produces mainly olefins and aromatics at an industry scale. The objective of this research is to investigate the economic feasibility of this proposed novel scheme. Heat integration tools are implemented in the designed plant to reduce the required external hot and cold utilities. A detailed technoeconomic analysis (TEA) study is performed on Aspen Process Economic Analyzer (APEA) to evaluate important economic indicators, including the net present value (NPV) and levelized cost of propylene production (LCoP) to assess the economic feasibility of the microwave-assisted novel process.

MICROWAVE-ASSISTED PLASTIC PYROLYSIS

In the laboratory, the plastic waste feed to the microwave reactor consists of equal portions of polypropylene and polyethylene obtained from real-world post-consumer plastic waste, each accounting for 50% of the total feed. Because the process is catalyzed, a 4 wt.% Ru/α-Fe_2O_3 catalyst was synthesized via the wetness impregnation technique as described in a previous related publication [7]. The plastic particles, after being crushed and sieved to small particle size, were mixed with the metal-oxide catalyst at a 1:1 mass ratio before being charged into a quartz tube reactor. The microwave-assisted pyrolysis of the plastic waste mixture at lab-scale was conducted in a microwave system which utilizes a Sairem microwave control system (Sairem Corp., Peachtree Corners, GA, USA) having a 2.45 GHz solid-state generator, which can provide up to 1 kW of maximum forward power. The system used is also equipped with a TE10 mono-mode cavity which serves as an applicator [6, 7]. The temperature of the reacting bed was measured using an infrared pyrometer, which regulated incident microwave power, and the external surface temperatures were independently monitored using an infrared thermometer (Micro Epsilon America, IL, USA). Prior to microwave irradiation, the system, being the plastic/catalyst mixture, was purged with nitrogen at 50 mL/min for a period of 30 min. The mixed plastic waste was then converted under microwave-assisted irradiation, and the gaseous effluent of the system was collected in sampling gas bag, with nitrogen as the carrier gas. The gaseous stream was analyzed in a four-channel Inficon Fusion micro-gas chromatograph (Inficon, NY, USA) and the gas phase product yields were quantified and reported. The system is equipped with the following analytical columns: Rt-Molsieve 5A (O_2, N_2, H_2, CO, CH_4), Rt-U-Bond (CO_2, C_2 hydrocarbons), Alumina-Na_2SO_4 (C_3–C_4 hydrocarbons), and Rxi-1ms (aromatic species such as benzene and toluene) [6, 7].

Table 1 summarizes the gas-phase compositions of the lab-scale microwave reactor outlet. It was observed that 94.8% of the solid plastic waste feed was converted to gas phase at a bulk temperature of 400°C and ambient pressure. The outlet consists of significant fractions of propylene and ethylene, which make up almost two-thirds of the gas products, while aromatics components (i.e., benzene and toluene) comprise 1.2% of the gaseous stream. These experimental data are utilized to develop a novel process flowsheet that produces polymer-grade (PG) propylene and ethylene which can be used as key feedstocks to produce polypropylene, polypropylene, and other high-value polyolefins. The aromatics stream is recovered as a mixed stream of benzene and toluene. Since alkanes, namely propane and ethane, make up 9% of the gaseous stream and are readily obtained as part of the PG propylene and PG ethylene purification process, respectively, they are also recovered to be sold as by-products since they can be used in various industrial applications.

DEVELOPMENT OF SIMULATION MODEL

Figure 1 presents an illustration of the novel microwave-assisted process for mixed plastic waste conversion. In the process flow diagram (PFD) that is developed for this study, the feed is specified to be 8000 kg/h of mixed plastic waste, comprising equal portions of polypropylene and polyethylene. Initially, the feed is directed to two single-screw extruders for feed processing before it enters the microwave reactor. The extruder configuration is representative of standard industrial units, featuring a 150 mm diameter and a 5 m length. Each extruder is capable of a processing capacity of 4000 kg/h [9].

Table 1: Gas Composition of Microwave Reactor Outlet.

Component	Gas composition (wt.%)
Hydrogen	1.4
Methane	10.2
Ethylene	25.9
Ethane	6.2
Acetylene	0.3
Propane	2.8
Propylene	36.4
Butane and butene(s)	15.6
Benzene	1.0
Toluene	0.2
Total	100
Gas Percent	94.8%

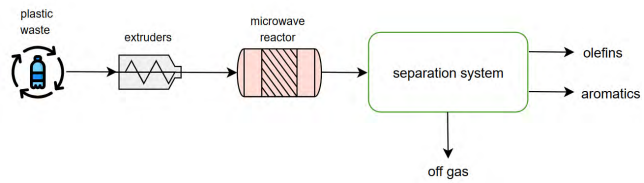

Figure 1. Illustration of microwave-assisted conversion of mixed plastic waste to olefins and aromatics.

The feed plastic mixture is pyrolyzed in the microwave reactor, and it is assumed in this study that the large-scale reactor will produce the same product stream as that observed in the lab-scale reactor. The separation train is modeled in Aspen Plus for the purification and recovery of the product streams as summarized in Figure 2. The hot reactor effluent, exiting at 400°C and ambient pressure, is first routed to a heat exchanger, where it is cooled down and its temperature is reduced before it enters the train of separation. The resulting cooled stream undergoes multi-stage compression, increasing the pressure from near-atmospheric conditions to 15 bar, before being introduced to the first distillation column. This column separates the heavy hydrocarbon fractions (S1)

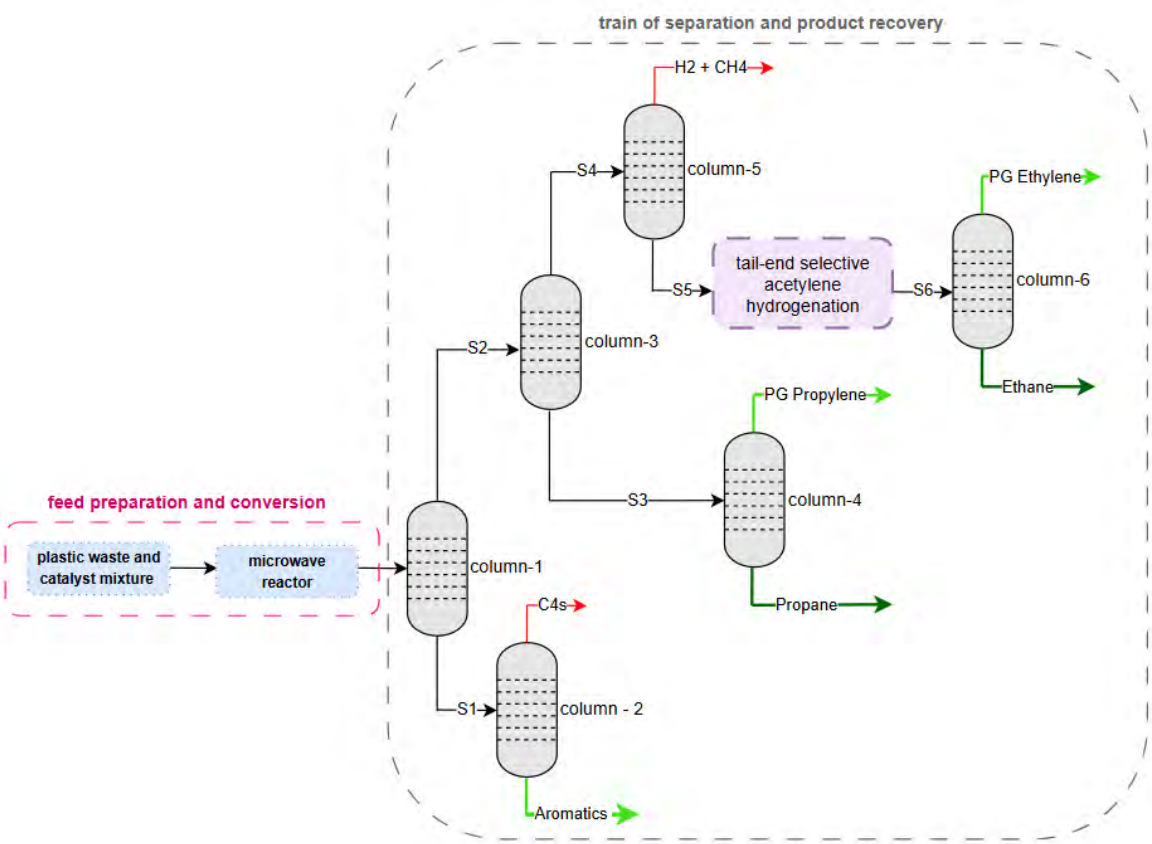

Figure 2. Train of seperation and recovered streams of the simulated novel microwave-assisted process. The main product streams are labeled *PG Propylene*, *PG Ethylene*, and *Aromatics*. The two byproduct streams are labeled

from the lighter components (S2). The heavy components stream (S1) is directed to another distillation column that effectively separates the aromatics mixture of benzene and toluene as a bottom product, which is composed of 14.4 mol% toluene and 84.9 mol% benzene at a rate of 96.3 kg/h. The top product stream of the distillation column consists primarily of a mixture of C_4 hydrocarbons and is utilized as an on-site fuel source. The stream which contains the propylene, propane, and lighter components (S2) is sent to another compressor to increase its pressure to 30 bar and then enters a cryogenic column which separates the propylene/propane mixture (S3) from the lighter fractions (S4). The propylene/propane mixture (S3) undergoes heat exchange and is expanded to 14 bar, which is a typical operating pressure of the propylene/propane splitter, before being redirected to the column for separating the two components. Because this separation is energy-intensive, the column operating pressure is selected to avoid the need for cryogenic cooling, thereby reducing the overall cost of the column. The column efficiently separates PG propylene, with a high purity of 99.85 mol% as the top product stream, and 96.7 mol% propane as a bottom product stream. The flowrates of the PG propylene and propane streams are 2889.1 and 222.7 kg/h, respectively.

In parallel, the pressure of the light ends stream (S4) is reduced to 25 bar before it is fed into a second cryogenic column to remove the C2-rich mixture (S5) from the mixture of hydrogen and methane. Similar to the C4-rich stream, the stream composed of predominantly hydrogen and methane is also used as on-site fuel. Because the C2-rich stream (S5) contains 1.01 mol% acetylene, it must be removed before ethylene purification, and a tail-end selective acetylene hydrogenation process is implemented for this purpose. The C2-rich stream (S5) has a pressure of 25 bar, which lies in the range of the reported operating pressures for acetylene hydrogenation [10]. The stream (S5) subsequently mixes with hydrogen to enable acetylene conversion via selective hydrogenation. The heated mixture enters the selective hydrogenation system, where the acetylene is converted across two beds: a lead bed and a guard bed, with interstitial heat exchangers installed to limit temperature rise which result due to the highly exothermic nature of the process [10]. The effluent stream of the hydrogenation system (S6) is cooled down before entering the ethylene/ethane splitter. The final column separates PG ethylene with 99.91 mol% purity containing < 1 ppm of acetylene, satisfying the recommended threshold for PG ethylene. The bottom product of this column is 97.9 mol% ethane, which is sold as a byproduct alongside the propane recovered as the bottom stream of the propylene/propane splitter. These byproducts can be used downstream, where they are further cracked into high-value olefins. The flowrates of the PG ethylene and ethane streams amount to 2088.0 and 509.5 kg/h, respectively. Figure 2 shows the schematic of the entire plant, including the detailed train of separation, which enables product recovery and purification.

COLUMNS COMPOSITION PROFILES

In Aspen Plus, the specifications of the distillation columns, such as the reflux ratio, feed stage location, and required number of equilibrium stages are first estimated using the DSTWU shortcut columns. The parameters obtained from the DSTWU columns are then used as inputs for the more rigorous RadFrac distillation column [9]. The propylene and ethylene recovered in the process need to meet specific purity requirements to qualify as PG olefins. Therefore, the propylene/propane and ethylene/ethane columns are optimized carefully, by selecting the optimal stage number and reflux ratio to achieve the required minimum purity. Figure 3 shows the composition profile of the propylene/propane splitter plotted as a function of stage number in the simulated plant. From the plot, it can be observed that the vapor mole fraction of propylene at the topmost stage, stage 1, is 0.9985 and the liquid mole fraction of propane at stage 130 is 0.9671.

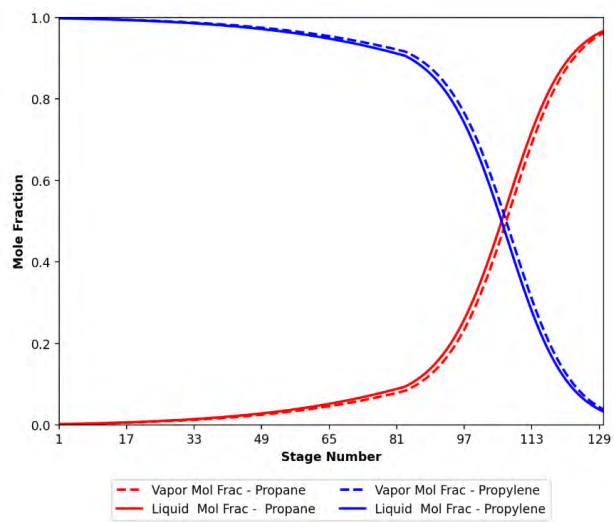

Figure 3. Composition profile in the propylene/propane splitter.

Figure 4 presents the molar composition profile in the ethylene/ethane splitter as a function of stage number. Polymer-grade ethylene is separated as the top product stream, and the vapor mole fraction of ethylene in stage 1 is 0.9991. Conversely, the liquid mole fraction of ethane in the bottommost stage, namely stage 82, is 0.9787. Moreover, in both distillation columns, partial-vapor condensers are employed since recovery of the

distillate in the liquid phase is not required. This configuration results in energy savings and a more cost-effective separation process.

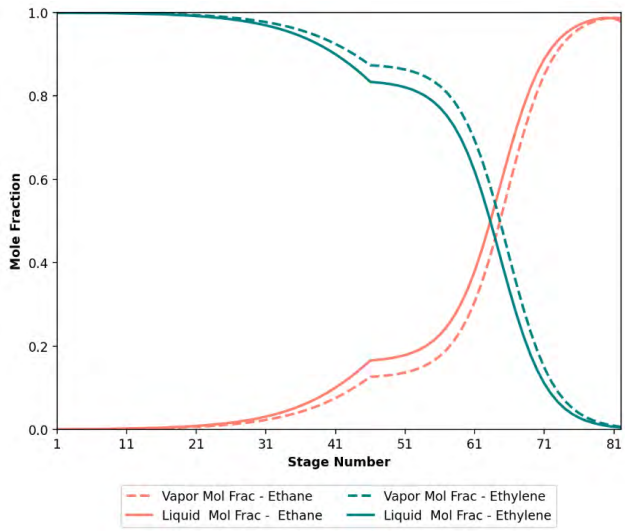

Figure 4. Composition profile in the ethylene/ethane splitter.

TECHNOECONOMIC ANALYSIS

The designed distillation train successfully achieved the required separations, yielding three main product streams and two byproduct streams. The economic evaluation of the simulated novel scheme is then conducted on APEA to assess the economic feasibility of the microwave-assisted novel plant. Since microwave-assisted processes are still in the developmental stage, a contingency factor of 26% is used for the economic study and the economic life of the project is taken to be 15 years [9]. In addition, a tax rate of 21% is assumed, and the raw material, products, and utility escalations are taken to be 1% per year. The cost of the microwave reactor is estimated using the numbering up approach and entered into APEA as a quoted item. Similarly, the costs of the single-screw extruders are also entered as quoted items in the software based on industry prices [9]. However, the costs of the remaining equipment are obtained by mapping the units from the developed PFD in Aspen Plus to APEA, where the software automatically performs equipment sizing and cost estimation. The cost of electricity in the economic evaluation is taken to be USD 0.06 per kWh as per the recommendation of the U.S. Department of Energy (DOE), while the cost of the other utilities is obtained from previous published TEA studies [9]. Using these economic assumptions, the TEA is carried out, and the total capital investment of the simulated novel plant is determined to be USD 48.1 million. Meanwhile, the total operating expenditure (OPEX) amounted to USD 18.8 million per year, with the raw material cost calculated to be USD 8.8 million per year. Since three main product streams are recovered and sold alongside two byproducts, the total product sales amount to USD 38.0 million per year.

Because microwave-assisted heating relies heavily on electricity, the total utility cost for the plant is high and is calculated to be USD 5.6 million per year. To investigate the contribution of each utility to the overall utility cost, Figure 5 presents a pie chart illustrating the distribution of costs among the five utilities used. As observed from the plot, electricity ranks first, representing just over half of the total utility cost, demonstrating the high electricity demand of the microwave-assisted novel process. The high contribution of electricity to the total utility cost highlights its role as the dominant factor among all utilities. The electricity cost is then followed by approximately 31% for refrigeration, which is primarily required for the cryogenic columns used in the separation of the light fractions and the purification and recovery of the PG ethylene stream. Low pressure steam (LP steam) comprises approximately 16% of the utility cost and is primarily used in the reboilers of the distillation columns as well as providing the necessary heat in the heat exchangers throughout the plant. Cooling water contributes minimally to the total utility cost, accounting for roughly 1% of the cost, while high pressure steam (HP steam) represents merely < 1% of the utility cost. It is rather crucial to investigate the utility cost distribution to identify the major cost drivers and prioritize opportunities for efficiency improvements.

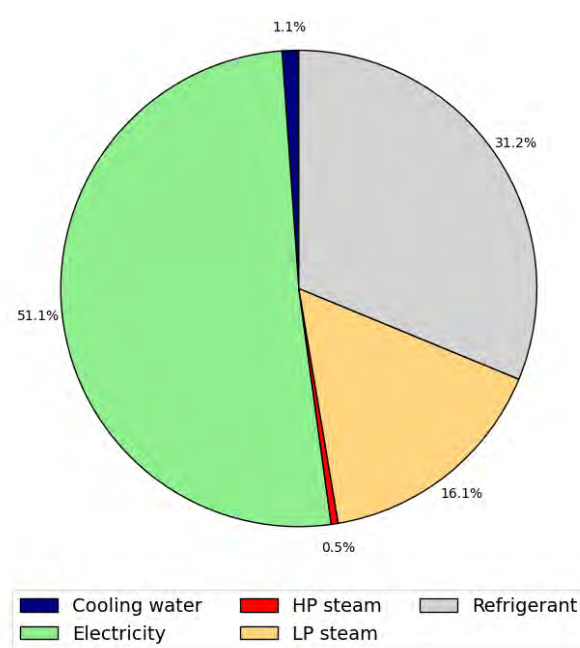

Figure 5. Total utility cost breakdown.

The NPV of the designed plant totals to USD 61.6 million, confirming the economic potential of the proposed novel plant for mixed plastic waste conversion at the investigated scale. Since propylene exhibits the highest yield among all product streams, the levelized cost of production is attributed to and calculated for propylene alone. In this case, revenues from the remaining four product streams are treated as credits that offset the total production cost. The LCoP is calculated using a discounted cash flow (DCF) analysis, assuming a 10% internal rate of return (IRR) over the 15-year plant lifespan, with the condition that the NPV at the estimated LCoP value amounts to zero. Thus, the LCoP is estimated to be USD 481.1 per ton, and the detailed TEA results of the proposed novel microwave-assisted simulated plant are summarised in Table 2.

Table 2: Economic Summary for the Simulated Novel Process.

Parameter and unit	Value
Total capital investment (million USD)	48.1
Operating cost (million USD/year)	18.8
Raw material cost (million USD/year)	8.8
Utility cost (million USD/year)	5.6
Product sales (million USD/year)	38.0
NPV (million USD)	61.6
LCoP (USD/ton)	481.1

Since the OPEX has a substantial impact on the plant's economics, a sensitivity analysis is performed on the raw material and electricity costs, as they are the main cost drivers, to demonstrate the effect of changing these costs on the OPEX. Figure 6 presents the effect of changing the cost of the electricity and raw material on the OPEX. From this figure, the following can be observed: when the electricity cost is increased by 50%, the OPEX increases by 8.2% and conversely, a 50% reduction in the electricity cost leads to a consequent reduction of 8.2% in the OPEX. In comparison, changes in the raw material cost appears to have a more pronounced effect, where a 50% increase in the raw material cost results in a 25.3% rise in the OPEX, while a 50% decrease leads to a corresponding 25.3% reduction in the OPEX.

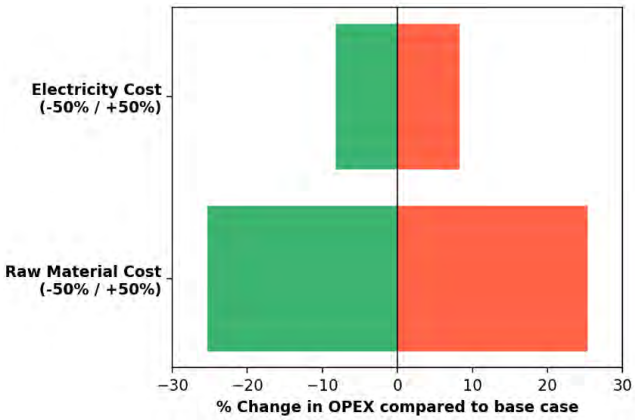

Figure 6. Effect of changing the cost of electricity and raw material on the OPEX.

CONCLUSIONS

The novel process presented in this work utilizes microwave-assisted heating to convert mixed plastic waste of polyethylene and polypropylene into olefins and aromatics at a bulk temperature of 400°C and atmospheric pressure. Experimental results from the conversion of mixed plastic waste to valuable products in a lab-scale microwave reactor are used to develop the flowsheet in Aspen Plus. The simulated separation train enabled the recovery of three main streams of high-value products, namely PG propylene, PG ethylene, as well as a mixed aromatics stream of benzene and toluene. In addition, propane and ethane are also recovered as high-purity by-product streams and can be sold to downstream cracking plants for further conversion into higher-value products. Economic evaluation confirmed the potential of the simulated novel plant at the investigated scale, with a positive NPV of USD 61.6 million in the economic lifetime and a LCoP of USD 481.1 per ton. Such results demonstrate the technoeconomic feasibility of producing valuable products from mixed plastic waste through microwave-assisted heating, highlighting the potential for efficient plastic waste recycling and the opportunity to contribute to a circular plastics economy.

PE and PP are the two most common plastics in use today. An initial 50/50 mixture of PP and PE was considered in this research to account for these two common plastics and a technoeconomic analysis was conducted to see if this process is economically feasible. However, mixed plastic waste could have other components, or the ratio of PE and PP could be different. Ideally, one would like a flexible plant that could take many different ratios and types of plastics and convert them to higher-value olefins profitably. Alternatively, one would have to take a mixed plastic waste and add either PE or PP to it so that the ratio of PE and PP is in a range that makes the production of olefins from the feed profitable, despite the extra cost of adjusting the feed composition. The

methods developed here provide the framework for this profitability analysis and is the subject of our current research.

ACKNOWLEDGEMENTS

This study is supported by the United States Department of Energy (DOE), Office of Science, under grant no. DE-SC0025622.

REFERENCES

1. Zhang Z, Tian J, Wu X, Surin I, Pérez-Ramírez J, Hemberger P, Bodi A. Unraveling radical and oxygenate routes in the oxidative dehydrogenation of propane over boron nitride. J. Am. Chem. Soc. 145:7910-7917 (2023). https://doi.org/10.1021/jacs.2c12970
2. Ethylene Global Industry Report 2024-2030: US to Dominate, Polyethylene to Account for Highest Growth Segment - ResearchAndMarkets.com [Internet]. [cited 2026 Jan 17]. Available from: https://www.businesswire.com/news/home/20240513898212/en/Ethylene-Global-Industry-Report-2024-2030-US-to-Dominate-Polyethylene-to-Account-for-Highest-Growth-Segment---ResearchAndMarkets.com
3. Tijani MEH, Zondag H, Van Delft Y. Review of electric cracking of hydrocarbons. ACS Sustainable Chem. Eng. 10:16070-16089 (2022). https://doi.org/10.1021/acssuschemeng.2c03427
4. Akah A, Al-Ghrami M. Maximizing propylene production via FCC technology. Appl Petrochem Res 5:377-392 (2015). https://doi.org/10.1007/s13203-015-0104-3
5. Agarwal A, Sengupta D, El-Halwagi M. Sustainable process design approach for on-purpose propylene production and intensification. ACS Sustainable Chem. Eng. 6:2407-2421 (2018). https://doi.org/10.1021/acssuschemeng.7b03854
6. Wang Y, Biddle T, Jiang C, Luong T, Chen R, Brown S, Jie X, Hu J. Microwave-driven upcycling of single-use plastics using zeolite catalyst. Chemical Engineering Journal 465:142918 (2023). https://doi.org/10.1016/j.cej.2023.142918
7. Luo C, Liu S, Li S, Robinson B, Wu T, Hu J, Wang Y. Boosting upcycling polypropylene waste to propylene via microwave catalysis. Applied Catalysis B: Environment and Energy 381:125834 (2026). https://doi.org/10.1016/j.apcatb.2025.125834
8. Deng Y, Bai X, Abdelsayed V, Shekhawat D, Muley PD, Karpe S, Mevawala C, Bhattacharyya D, Robinson B, Caiola A, Powell JB, van Bavel AP, Hu J, Veser G. Microwave-assisted conversion of methane over h-(fe)-zsm-5: evidence for formation of hot metal sites. Chemical Engineering Journal 420:129670 (2021). https://doi.org/10.1016/j.cej.2021.129670
9. Al-Sakkaf A, Luo C, Wang Y, Palanki S. Technoeconomic analysis of a microwave-assisted novel process that converts polypropylene plastic waste to propylene. Ind. Eng. Chem. Res. 64:23568-23581 (2025). https://doi.org/10.1021/acs.iecr.5c03596
10. Dehghani Khold O, Parhoudeh M, Rahimpour MR, Raeissi S. A new configuration in the tail-end acetylene hydrogenation reactor to enhance catalyst lifetime and performance. Journal of the Taiwan Institute of Chemical Engineers 65:8-21 (2016). https://doi.org/10.1016/j.jtice.2016.04.027

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Process-Informed Design of Electrochemical Cells for Urea Production: A Techno-Economic and Systems Engineering Approach

Zhimian Hao[a], Shilong Fu[b], Chengtian Cui[ac], Ruud Kortlever[b], Ruud van Ommen[a*], and Ana Somoza-Tornos[a*]

[a] Department of Chemical Engineering, Delft University of Technology, Delft, The Netherlands
[b] Department of Process and Energy, Delft University of Technology, Delft, The Netherlands
[c] Process and Systems Engineering Laboratory, Åbo Akademi University, Turku, 20500, Finland
* Corresponding Author: A.SomozaTornos@tudelft.nl, J.R.vanOmmen@tudelft.nl

ABSTRACT

Conventional urea production is a centralized and fossil-intensive process associated with significant greenhouse-gas (GHG) emissions and limited flexibility for deep decarbonization. As an alternative, the Integrated COnversion of NItrate and Carbonate steams (ICONIC) project is developing innovative electrochemical urea (e-urea), via the co-electroreduction of nitrogen and carbon sources using renewable power. While recent research advances in electrocatalysis have demonstrated promising Faradaic efficiencies (FE) toward urea, the design of electrochemical systems involves inherent trade-offs between key performance indicators (KPIs) such as current density, cell voltage, and FE. Crucially, the implications of electrolyzer-level performance on plant-level economics and environmental impacts remain poorly understood. To address this gap, we integrate process modelling with techno-economic and life-cycle assessment (TEA–LCA) to evaluate the trade-offs of KPIs from a process systems perspective. Lab-scale electrolyzer KPIs are transformed into process-level environmental/economic metric, using Python-based electrolyzer modelling and Aspen Plus downstream simulations. Fossil urea is used as a benchmark, and production scenarios are evaluated for wind-powered operation in the Netherlands and solar-powered operation in Spain. The results identify quantitative performance targets for electrolyzer design and demonstrate that urea recovery in the downstream are decisive bottlenecks. This work reveals the essentials for expanding the system boundary beyond the electrolyzer cell, as to identify R&D prioritization and deployment strategies during early-stage technology development and lays the groundwork for scalable and low-carbon urea production.

Keywords: Multiscale Modelling, Carbon Dioxide Sequestration, Life Cycle Analysis, Technoeconomic Analysis, Process Design, urea electrosynthesis

1 INTRODUCTION

Nitrogen mismanagement presents a dual global challenge. On the one hand, excess nitrogen species released into aquatic environments cause coastal eutrophication. On the other hand, the agriculture sector remains heavily dependent on nitrogen fertilizers, such as urea, and faces challenges regarding supply security. For example, the European Union (EU) needs to import approximately 6 Mt urea annually from Asia, Africa and other regions [1].

Conventional urea production relies on centralized Haber–Bosch ammonia synthesis followed by thermal urea formation by Bosch-Meiser process. This pathway is fossil-dependent, energy-intensive and poorly compatible with variable renewable electricity, resulting in substantial CO_2 emissions and limited potential for deep decarbonization [2, 3].

To address these challenges, the Integrated COnversion of NItrate and Carbonate steams (ICONIC) project positions electrochemical urea (e-urea) as an alternative [4]. By coupling renewable electricity with

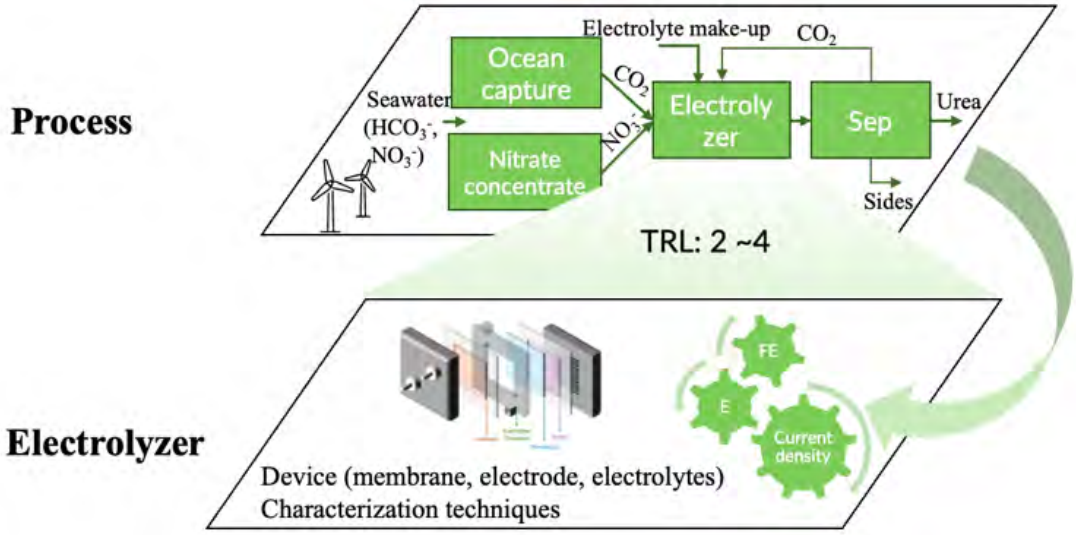

Figure 1. Process-informed KPI framework for electrochemical cells.

electrochemical C–N bond formation from nitrate and CO_2, ICONIC envisions a distributed, electricity-driven pathway in which e-urea directly substitutes fossil urea in existing markets.

ICONIC's e-urea is currently at technology readiness level (TRL) 2-4 and will soon face the *Valley of Death* challenge from R&D to commercialization. In recent years, research on e-urea has expanded rapidly [5], with reported FE approaching 75% in 2023 [3] and 90% in 2025 [6] but all in low current density with low yield of urea. However, most research scopes are limited to 'point view' for the development of electrocatalysts, membranes, and electrolyzer concepts for cell-level performance metrics [5], with limited consideration of system-level integration.

At higher TRLs, where demonstration activities dominate, fundamental changes to process concepts become increasingly risky, costly, and difficult to implement [7]. Conversely, at low to intermediate TRLs, system-level redesign remains feasible and can substantially influence long-term cost and environmental performance. This makes early-stage, process-informed analysis particularly valuable on cost and risk management.

Specifically, e-urea synthesis faces a critical gap: electrolyzer KPIs are rarely linked quantitatively to plant-level cost, GHG emissions, and deployment choices such as separation strategy and electricity mix (distinguished by location [8]). Without such linkage, it is difficult to define realistic performance targets that guide fundamental research toward commercial relevance.

This study addresses the following research question: What electrolyzer KPI targets (FE, current density, and electrolyzer cost) are required for ICONIC e-urea to become economically and environmentally competitive with fossil urea? Process models are employed to systematically propagate electrolyzer performance to plant-level metrics. Following this, this work aims to identify key technology bottlenecks at an early stage and provide guidance for R&D and scale-up decisions.

2 METHODOLOGY

This study adopts a process systems engineering (PSE) methodology to target the KPIs of electrolyser for ICONIC e-urea production. A steady-state process model is developed to represent the full value chain, including CO_2 capture, electrochemical C–N conversion, and downstream product recovery. The electrolyzer is parameterised using experimental data as the baseline operating point. Future technology development is explored through parametric variation of these KPIs to reflect potential improvements in catalysts, membranes, and cell design.

Mass and energy balances by process simulations are used to quantify system-level electricity demand and electrolyzer sizing. These results are subsequently integrated into a techno-economic assessment to evaluate levelized e-urea cost and identify cost and energy drivers. By systematically mapping electrolyser-level KPIs to process-level performance, the framework derives process-informed performance targets for early-stage technology development.

2.1. Process description

The system boundary extends beyond the electrolyzer to capture the full cradle-to-gate e-urea process. The ICONIC concept comprises:

- Nitrate separation from aqueous streams;

- CO_2 capture from ocean carbonate sources;
- C–N co-electrolysis in a BPM-based flow cell;
- Downstream purification to obtain solid urea.

CO_2 is supplied via ocean capture, which exploits the ocean carbonate buffer system as a large, naturally equilibrated carbon reservoir. Seawater is electrochemically acidified to shift dissolved inorganic carbon (HCO_3^-/CO_3^{2-}) toward dissolved CO_2, which is subsequently released and separated as a concentrated CO_2 stream. The treated stream is neutralised prior to discharge. The energy consumption for ocean capture is reported at 11.3 GJ/t-CO_2 [9].

The downstream purification consists of steam stripping for ammonia recovery, followed by evaporation for water removal, crystallisation of urea, and solid–liquid separation by filtration as to obtain the solid urea product. This flowsheet reflects mature industrial practice and enables systematic evaluation of energy consumption to purify the urea from the liquid phase following the electrolyzer.

Process modelling is used to link electrolyzer performance to system-level economic and environmental impacts. Process model covers all modules except the nitrate pre-treatment, where energy consumption is assumed as negligible while achieving a nitrate concentration of 0.1M NO_3^-, which is commonly used for the C-N electrolyzer operation [5].

2.2 Electrolyzer Modelling

The electrolyzer is modelled in Python using mass and energy balances for a three-compartment BPM flow-cell (cathode, anode, membrane). The electrolyser model is informed by experimental research conducted at TU Delft using a small-scale flow-cell electrolyser for the co-electrolysis of CO_2 and nitrate to urea. Current experimental result (FE, current density, and cell voltage) is used as the initial operating point for the process model. To represent future scale-up and technology maturation, improved electrolyser performance is systematically explored through parametric variation of these KPIs, reflecting plausible advancements in catalyst activity, membrane performance, and cell design. Scale-up is represented via modular replication of identical cells, assuming preservation of cell-level performance.

Given inputs include the urea production rate, nitrate and CO_2 conversion, FE, and current density; the model outputs are the required electrode area, energy balance (cell voltage), mass balance (outlet stream compositions).

2.3 Downstream Process Simulation

Downstream separation is simulated in Aspen Plus to recover urea from the dilute aqueous product stream produced by electrochemical C–N co-electrolysis. As no established industrial separation route exists for urea recovery at such low concentrations, representative downstream configurations are designed as engineering case studies to systematically evaluate energy and cost associated with urea recovery. The base case consists of steam stripping for NH_3 recovery, followed by multi-effect evaporation, urea crystallisation, and solid–liquid separation by filtration, providing a conservative thermally driven reference.

As to reduce the energy consumption, alternative process configurations incorporate pre-concentration (e.g. reverse osmosis, RO) and heat-pump (HP) assisted evaporation [10, 11].

2.4 Techno-Economic and Life-Cycle Assessment

The functional unit is 1 kg of urea on a cradle-to-gate, covering raw material supply, co-electrolysis and urea recovery. Fossil urea is used as a benchmark (0.38 $/kg-urea; 1.38 kg CO_2/kg-urea). ICONIC targets a production cost at 1 €/kg-urea.

Electrolyzer CAPEX and electricity price can be two important drivers for e-chemicals [12, 13]. ICONIC electrolyzer CAPEX for C/N co-electrolysis for e-urea is assumed to range from 5, 000 $/m² (base) to 2, 500 $/m² (optimistic), reflecting inheritance from both CO_2 reduction reaction (CO2RR) at 900–1800 $/m² [13] and nitrogen reduction reaction (NRR) technologies (4000–7500 $/m² by 2050) [12].

Electricity cost and carbon intensity are strongly location-dependent and therefore explicitly considered through the selection of two representative production locations. The Netherlands (NL) is chosen to represent a system dominated by offshore wind power, characterised by relatively low-carbon but moderate-cost electricity. Spain (ES) is selected to represent a solar-dominated system with lower electricity costs but a higher carbon intensity. Both base and optimistic scenarios are evaluated for each location to systematically assess the combined effects of electricity price and electrolyzer cost on the economic performance of e-urea production (Table 1).

Table 1. Scenarios of locations and electrolyzer costs for e-urea production.

Scenarios	Energy cost ($/kWh)	Carbon intensity (gCO_2/kWh)	Electrolyzer cost ($/m²)
NL base (wind)	0.07	12	5000
ES base (solar)	0.04	33	5000
NL optimistic (wind)	0.03	12	2500
ES optimistic (solar)	0.01	33	2500

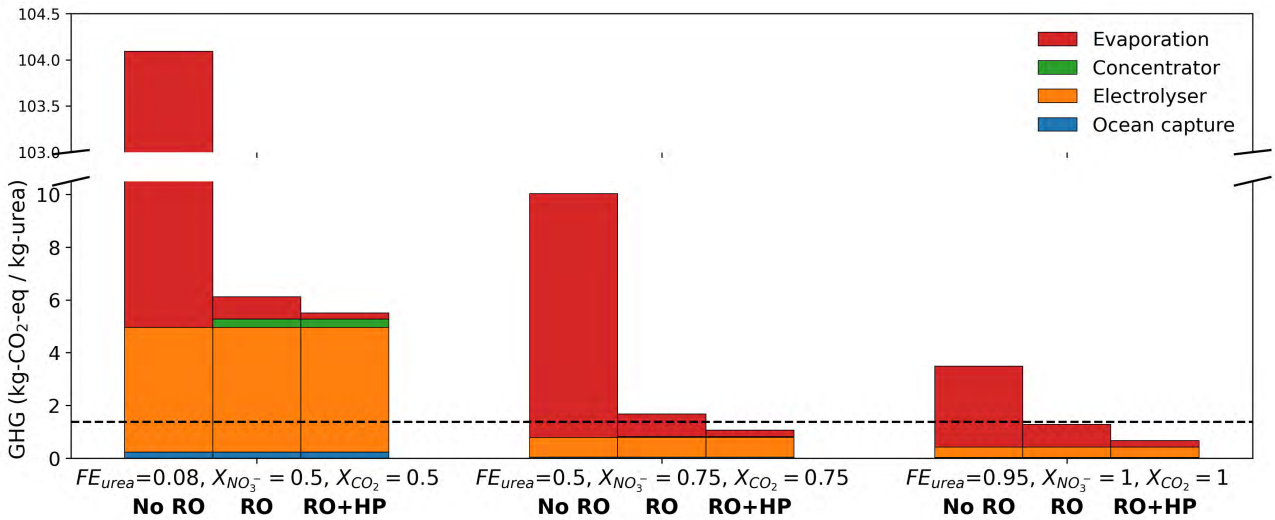

Figure 2. GHG breakdowns with the improvement of electrolyzer and different downstream design options. FE_urea is the FE to urea; X_NO3- and X_CO2 is the conversion.

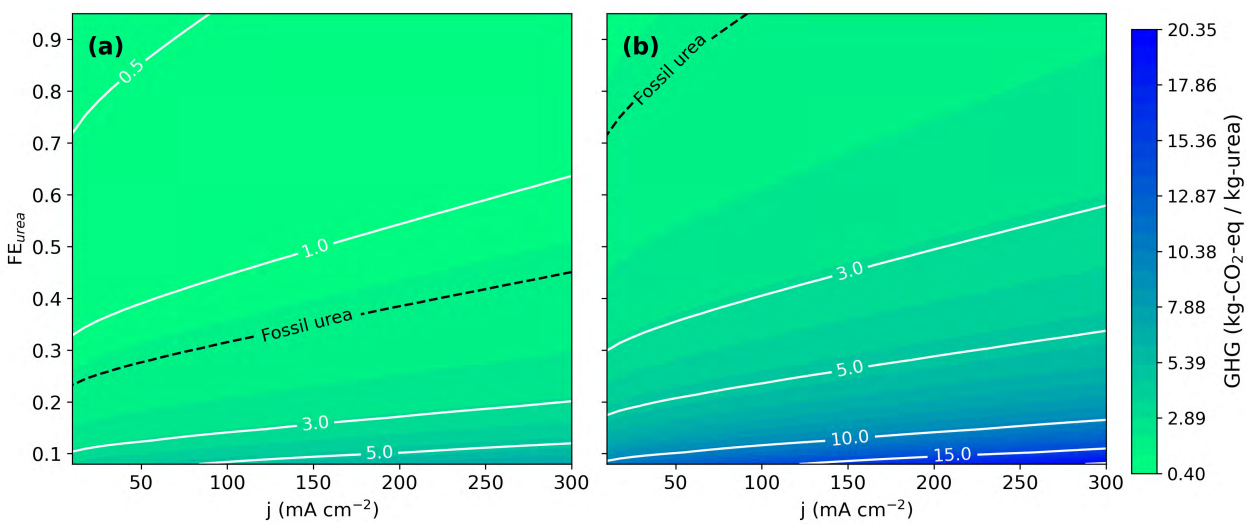

Figure 3. Targeting electrolyzer cell KPIs by comparing GHG emissions of e-urea to fossil urea: (a) wind-powered operation in the Netherlands and (b) solar-powered operation in Spain. RO pre-concentration and HP-assisted evaporation are included. Complete nitrate and CO_2 conversion is assumed.

3 RESULTS AND DISCUSSIONS

3.1 GHG Emissions

Figure 2 presents the GHG breakdowns for e-urea production by ocean capture for CO_2, electrolyzer, urea concentrator, downstream separation under the improving electrolyzer performances. The dashed line indicates the GHG intensity of fossil urea for reference.

For each electrolyzer case, three downstream separation configurations are considered:

- No RO, where urea recovery mainly relies on thermal evaporation.
- RO, where reverse osmosis (RO) is deployed to concentrate urea prior to evaporation.
- RO+HP, where RO is combined with heat-pump-assisted evaporation.

Notably, downstream evaporation dominates GHG emissions, primarily due to the extremely dilute urea concentration (0.02–0.3 wt%) exiting the electrolyzer. Even with near-ideal electrolyzer performance (the right 'no RO' in Figure 2) and the benefits of renewable electricity, the intensive evaporation energy alone can result in more GHG emissions than fossil urea. The results show that, without pre-concentration, downstream purification dominates total GHG emissions, whereas RO and RO+HP

substantially reduce the thermal energy demand and associated emissions.

Introducing RO concentration can reduce thermal energy demand by up to ~90%. HP-assisted evaporation can further cut thermal duty by additionally 90%. RO and HP substantially expand the feasible operating window for electrolyzer KPIs.

Comparing to the GHG of fossil urea, Figure 3 presents sensitivity analysis for the feasible operating window for electrolyzer. Lower current density and higher FE to urea are environmentally preferred. Under wind power in NL, e-urea achieves lower GHG emissions than fossil urea at FE > 0.4. In ES, higher FE (>0.7) and moderate current densities (and j < 100 mA/cm²) are required due to the higher carbon intensity of solar electricity than wind (**Table 1**).

3.2 Techno-Economic Assessment

Figure 4 illustrates the effect of electrolyzer current density on e-urea production cost. Increasing current density reduces electrolyser CAPEX due to the reduction of required electrolyzer area, while simultaneously increases electricity-related OPEX. As a result, the total production cost exhibits a shallow minimum at intermediate current densities (≈50–100 mA·cm⁻²). The dashed lines indicate the fossil urea benchmark and the ICONIC target cost for comparison.

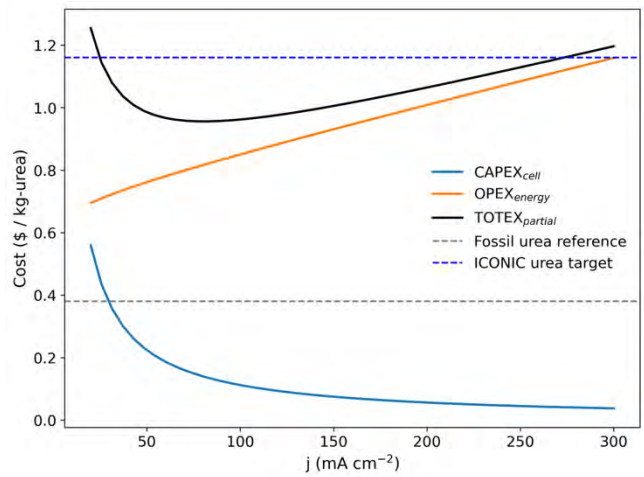

Figure 4. Cost of urea with respect to current density.

Figure 4 shows e-urea is more than twice costly than fossil urea, while a proper operating window ranging from 30 to 270 $mA \cdot cm^{-2}$can still achieve the ICONIC target at 1 €/kg-urea.

As shown in Figure 5, sensitivity analysis shows a

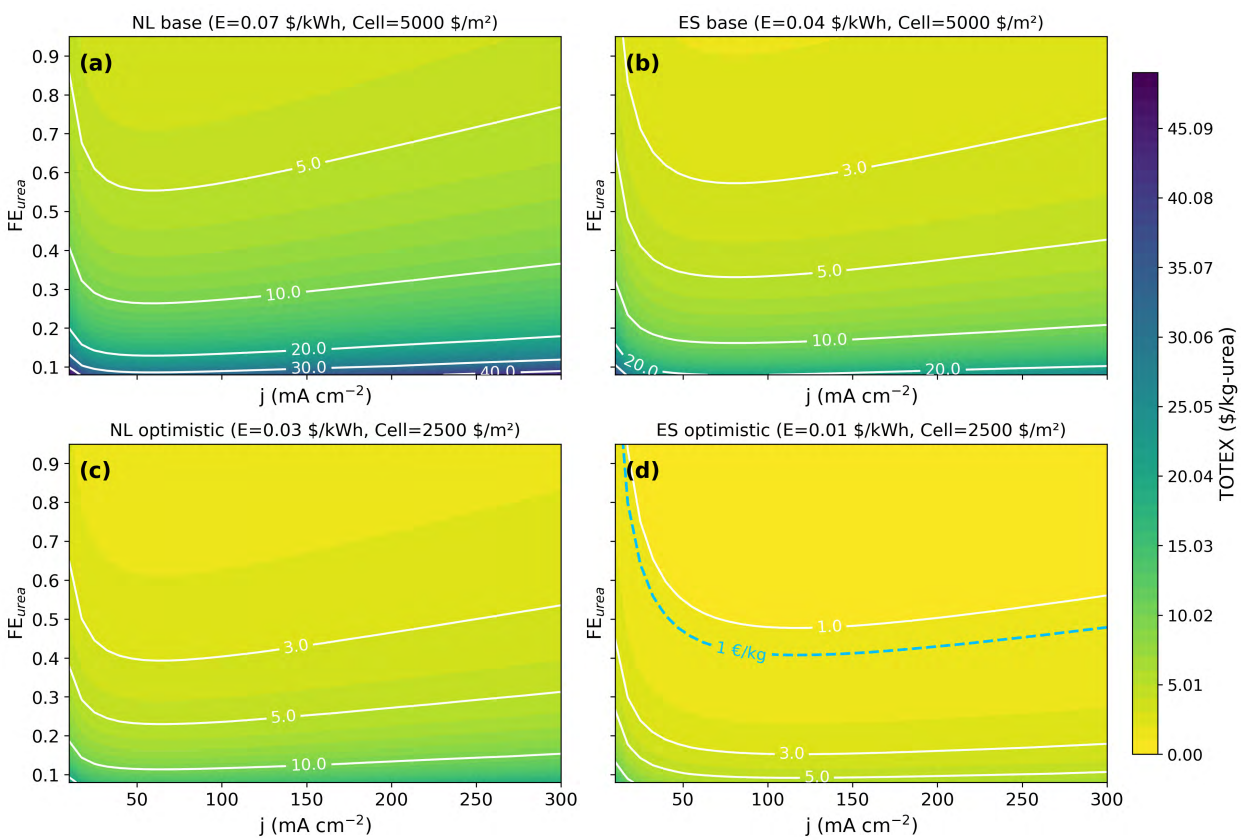

Figure 5. Targeting electrolyzer cell KPIs by the production cost of e-urea. (a) NL – base; (b) ES – base; (c) NL – optimistic; (d) ES – optimistic. RO pre-concentration and HP-assisted evaporation are included. Complete nitrate and CO_2 conversion is assumed.

higher FE always reduces cost and a proper current density range: 50–100 mA/cm² can further reduce the cost. In the base scenario, neither NL nor ES achieves cost parity with fossil urea. However, under optimistic ES conditions (FE > 0.4, j > 50 mA/cm²), ICONIC e-urea approaches (~1 $/kg-urea), compared to the ICONIC target at 1 €/kg-urea.

CONCLUSION CONCLUSIONS AND FUTURE OUTLOOK

This study demonstrates that it is challenging for e-urea to environmentally and economically compete with fossil urea, if the electrolyzer alone is considered**.** Key insights for the electrolysis include:

- Faradaic efficiency: higher FE is always preferred, as it lowers specific energy (kWh/kg-urea), reduces side-product formation, and shrinks downstream duty.
- Current density: Moderate targets (~50–100 mA·cm^{-2}) are preferred, as it balances CAPEX and OPEX.
- Electrolyzer cost: Achieving competitiveness requires CAPEX well below 5, 000 $/m², with a medium-term target of ~2, 500 $/m².

Notably, downstream (upstream) processes of electrolyzer matters. The dominant bottleneck is dilute urea from electrolyzer outlet for the downstream, which is directly related to the low NO_3^- concentration at 0.1 M in the electrolyzer inlet for the upstream (the concentration of NO_3^- in coastal water is even much lower, so upstream for NO_3^- separation also matters). Novel energy-efficient recovery and heat integration can significantly expand the feasible design space for electrolyzer KPIs.

As per location, manufacturing in the Netherlands results in lower GHG emissions but remains economically less viable than production in Spain.

Future research should shift beyond catalysts and membranes for electrolyzer toward process-level innovations, including electrolyzer configurations that directly concentrate urea and advanced liquid-phase urea recovery technologies. By linking lab-scale KPIs to system-level performance, this TEA–LCA framework provides actionable guidance for scaling up e-urea production in early stage and enhance the resilience of the fertilizer supply chain.

ACKNOWLEDGEMENTS

Authors gratefully acknowledge Prof. Tony Kiss and Dr. Tamara Jankovic for the insightful discussions on the design options for downstream separation from diluted solution, and Qing Li for the discussions on heat pump-assisted distillation. This work received the financial support from European Union: ICONIC (101115204) project.

USE OF AI TOOLS

Authors used ChatGPT-5 to improve readability and language of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

REFERENCES

1. European Commission, EU Fertiliser Market analysis, (n.d.). https://agriculture.ec.europa.eu/document/download/7834f9ad-c7ca-444a-bdd4-1e7bf209eb5f_en?filename=fertilisers-mo-2023-11-24-presentation_en.pdf (accessed October 1, 2025).
2. Mao C, Byun J, MacLeod HW, Maravelias CT, Ozin GA. Green urea production for sustainable agriculture. Joule 8:1224-1238 (2024). https://doi.org/10.1016/j.joule.2024.02.021
3. Luo Y, Xie K, Ou P, Lavallais C, Peng T, Chen Z, Zhang Z, Wang N, Li XY, Grigioni I, Liu B, Sinton D, Dunn JB, Sargent EH. Selective electrochemical synthesis of urea from nitrate and CO2 via relay catalysis on hybrid catalysts. Nat Catal 6:939-948 (2023). https://doi.org/10.1038/s41929-023-01020-4
4. ICONIC, Integrated COnversion of NItrate and Carbonate steams, (n.d.). https://iconicproject.eu (accessed May 1, 2025).
5. Zhou Y, Xing G, Sun C, Chen Z, Li S, Yang R, Chen M, Zhang P, Feng C, Abudula A, Guan G. Electrosynthesis of urea from carbon dioxide and waste nitrates: history, recent progress, and future prospects. Adv Funct Materials 36: (2025). https://doi.org/10.1002/adfm.202515635
6. Gong C, Peng Y, Xu M, Wei X, Sheng G, Liu J, Wu X, Han X, Dai F, Dong J, Chen Z, Zhu Y, Ye W, Cui Y. Selective electrocatalytic synthesis of urea using entangled iron porphyrins in covalent organic frameworks. Nat. Synth 4:720-729 (2025). https://doi.org/10.1038/s44160-025-00742-6
7. J. Weber, WEC Technology Readiness and Performance Matrix – finding the best research technology development trajectory, (n.d.).
8. Hao Z, Hamad WY, Yaseneva P. Understanding the environmental impacts of large-scale cellulose nanocrystals production: case studies in regions dependent on renewable and fossil fuel energy sources. Chemical Engineering Journal 478:147160 (2023). https://doi.org/10.1016/j.cej.2023.147160
9. de Lannoy CF, Eisaman MD, Jose A, Karnitz SD,

DeVaul RW, Hannun K, Rivest JLB. Indirect ocean capture of atmospheric CO2: part I. prototype of a negative emissions technology. International Journal of Greenhouse Gas Control 70:243-253 (2018). https://doi.org/10.1016/j.ijggc.2017.10.007

10. Cui C, Qi M, Zhang X, Sun J, Li Q, Kiss AA, Wong DSH, Masuku CM, Lee M. Electrification of distillation for decarbonization: an overview and perspective. Renewable and Sustainable Energy Reviews 199:114522 (2024). https://doi.org/10.1016/j.rser.2024.114522
11. Dal Mas R, Carta A, Somoza-Tornos A, Kiss AA. Coupling CO2 electrolysis and downstream processing via heat pump-based waste heat recovery. Computers & Chemical Engineering 204:109330 (2026). https://doi.org/10.1016/j.compchemeng.2025.109330
12. Izelaar B, Ramdin M, Vlierboom A, Pérez-Fortes M, van der Slikke D, Sajeev Kumar A, de Jong W, Mulder FM, Kortlever R. Techno-economic assessment of different small-scale electrochemical nh₃ production plants. Energy Environ. Sci. 17:7983-7998 (2024). https://doi.org/10.1039/d4ee03299c
13. Jouny M, Luc W, Jiao F. General techno-economic analysis of co₂ electrolysis systems. Ind. Eng. Chem. Res. 57:2165-2177 (2018). https://doi.org/10.1021/acs.iecr.7b03514

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Computed-Aided Design of an Intensified Process for the Sustainable Production of Biodiesel from Waste Cooking

Tania G. Salgado-Rodríguez[a], Fernando I. Gómez-Castro[a]*, Nelly Ramírez-Corona[b]

[a] Universidad de Guanajuato, Campus Guanajuato, División de Ciencia Naturales y Exactas, Departamento de Ingeniería Química, Guanajuato, Guanajuato, México.
[b] Universidad de las Américas Puebla, Departamento de Alimentos, Ingeniería Química y Ambiental, Puebla, Puebla, México.
* Corresponding Author: fgomez@ugto.mx.

ABSTRACT

The utilization of low-quality vegetable oils as raw materials helps to reduce the production costs of biodiesel. Waste cooking oils are examples of this type of raw material, having a high content of free fatty acids. While biodiesel is a sustainable alternative to fossil fuels, conventional production methods face challenges due to low reaction rates and high energy demands. This study investigates pathways for biodiesel production from waste cooking oil using process intensification technologies in combination with biowaste-derived heterogeneous catalyst. Conventional and intensified (reactive distillation-based) processes are compared using Aspen Plus simulations as an analysis tool, focusing on the production of biodiesel from waste cooking oil (WCO) using CaO as a catalyst. According to the results, the conventional method at 60°C and 6:1 methanol-oil ratio achieves 95% conversion but suffers from high methanol use and long reaction times (65 min). The intensified process at 65–70°C reduces energy needs by more than 29%, requiring shorter residence time while enhancing efficiency. Overall, it demonstrates that intensified and sustainable biodiesel production can reduce environmental footprints and support the transition to cleaner energy systems. This work emphasizes the transition from traditional processing routes to innovative, intensified technologies for sustainable fuel production within the framework of the circular economy.

Keywords: reactive distillation, process design, biodiesel, waste cooking oil

INTRODUCTION

Recently, the global energy demand climbed 63% between 2010 and 2019, reaching 14, 406 Mtoe due to population expansion and modern lifestyle changes. The International Energy Agency (IEA) projects a further 20% increase by 2050. Thus, the transition to renewable energy sources, particularly biofuels, has become more urgent than ever [1]. Biofuels (bioethanol, biodiesel, biogas, etcetera) produced from different organic feedstocks such as food, municipal, and agricultural waste might provide an alternative solution to overcome the global challenges faced by the excessive use of fossil fuels, including economy, reservoir depletion, and pollution [2].

Biodiesel is a very attractive option to overcome the energy crisis, since it is derived from renewable feedstocks, which may lead to more sustainable production schemes. Moreover, its use as energy source may provide substantial reductions in greenhouse gas emissions [3].

Lipidic waste such as waste cooking oils (WCOs), and slaughterhouse waste are potential feedstocks for biodiesel production. The waste cooking oil is inexpensive and widely available due to its numerous sources (restaurants, food processing industries, fast food outlets, and households), which makes it a potential alternative as raw material. The utilization of waste cooking oil to produce biofuels is helpful to deal with the energy crisis, food safety, and environmental problems because of the renewable, non-toxic, and clean combustion alternative of such derivatives [2].

An industrial-scale biodiesel production facility usually employs alkali catalysts for transesterification. However, the use of alkaline catalyst may lead to undesired side reactions for raw materials with high composition of free fatty acids. In these cases, esterification can convert

FFA into fatty acid methyl esters (FAMEs) using an acid catalyst such as sulfuric acid as a pretreatment step [5]. Therefore, an additional process is required when changing the feedstock from refined vegetable oil to WCO to decrease FFA in the raw material before transesterification.

A heterogeneous catalyst derived from biowaste, specifically eggshells, has been used to improve the environmental and operational performance of biodiesel production. Environmental assessment has indicated that CaO-based heterogeneous catalyst significantly reduces impacts on human health and ecosystems while maintaining high biodiesel yields. Compared to conventional catalysts, biowaste-derived CaO provides a sustainable alternative, balancing high conversion efficiency with reduced environmental impacts and aligning with circular economy principles. [6]

Process intensification seeks to transform the industry through integrated design, which aims at reducing the mass and heat transfer resistances while overcoming thermodynamic limitations through integrated designs and operation. In retches enchainment years, PI has attracted considerable academic interest as a potential mean for process improvement, to meet the increasing demands for sustainable production [7]. A clear example of intensified equipment is given by the reactive distillation columns.

The concept of reactive distillation combines these two important functions for the enhancement of the process. The separation of the product from the reaction mixture may not need a separate distillation step, which saves energy for heating and material.

When considering the characteristics of the transesterification reaction, a multifunctional reactor combining chemical reaction and distillation in a single column is a promising way. The simultaneous separation and reaction inside the column allow the equilibrium-limited reaction to take place more completely and reduce the separation cost of unreacted reactants [8].

The design of the reactive distillation focuses on the identification of the characteristics and operating conditions of the separation systems, which are significantly more complex than those involved in ordinary distillation. However, its advantages may compensate this additional complexity, considerably enhancing the performance of the process and its economic and environmental indicators.

Although production of biodiesel is technically feasible, strategies must be sought to make it cost-competitive with fossil diesel, implementing strategies that make more efficient its production. One of the key alternatives that allows for reducing production costs is the use of low-cost raw materials, as well as the implementation of intensified technologies [2]. Therefore, the objective of this research is to evaluate the implementation of intensified technologies to improve process efficiency and reduce energy impact in the production of biodiesel with heterogeneous catalysis.

CASE STUDY

The biodiesel production processes studied in this work use waste cooking oil as raw material. The designs of the processes have been obtained by simulation in the commercial software Aspen Plus. The first steps in developing the process simulation were defining components and selecting appropriate thermodynamic models. In addition, unit operations and operating conditions, input conditions and plant capacity must be all selected and specified.

Based on the aforementioned information, a process feeding of 2900 kg/h of waste cooking oil has been proposed; this is an estimated amount generated annually in Mexico. [9] Table 1 presents the lipidic profile employed to represent the oil. Likewise, numerous studies have reported the interaction parameters between glycerol, methanol, and water. Some authors have studied the behavior of liquid-liquid equilibria, which are fundamental to the design of biodiesel purification processes, comparing experimental data with calculations provided by thermodynamic models. In the study by Tomassetti et al. [10], who measured liquid-liquid equilibrium, the UNIQUAC model accurately follows the LLE of biodiesel + glycerol and biodiesel + water, while the UNIFAC model can predict values with most of the experimental data for the FAMEs + glycerol system. The use of the UNIFAC L-L model is also reaffirmed by Souza et al. [11]. In the study of Pasha et al., the non-random two-liquid (NRTL) model was selected as the property package due to the presence of polar compounds and the non-ideal nature of the transesterification reaction system, with methyl esters being of primary interest [12]. NRTL has been chosen for most stages of the process studied in this work. UNIFAC-LL is used for phase separation equipment in the presence of ethyl esters with glycerol and methanol.

Table 1: Lipid composition of waste cooking oil.[13]

Fatty acids (6% wt)	Triglicerides (94% wt)
Palmitc acid (12.3% wt)	Trypalmitin (12.3% wt)
Oleic acid (24% wt)	Triolein (24% wt)
Linoleic acid (56.6% wt)	Trilinolein (56.6% wt)
Linoleic acid (7.1% wt)	Trilinolenin (7.1% wt)

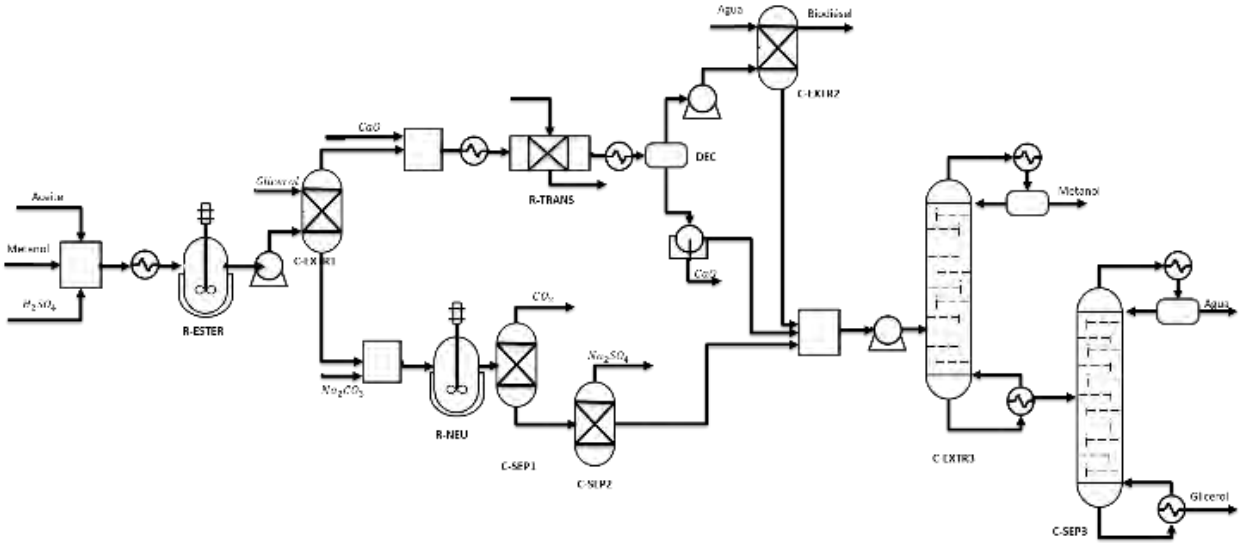

Figure 1. General conventional process treatment.

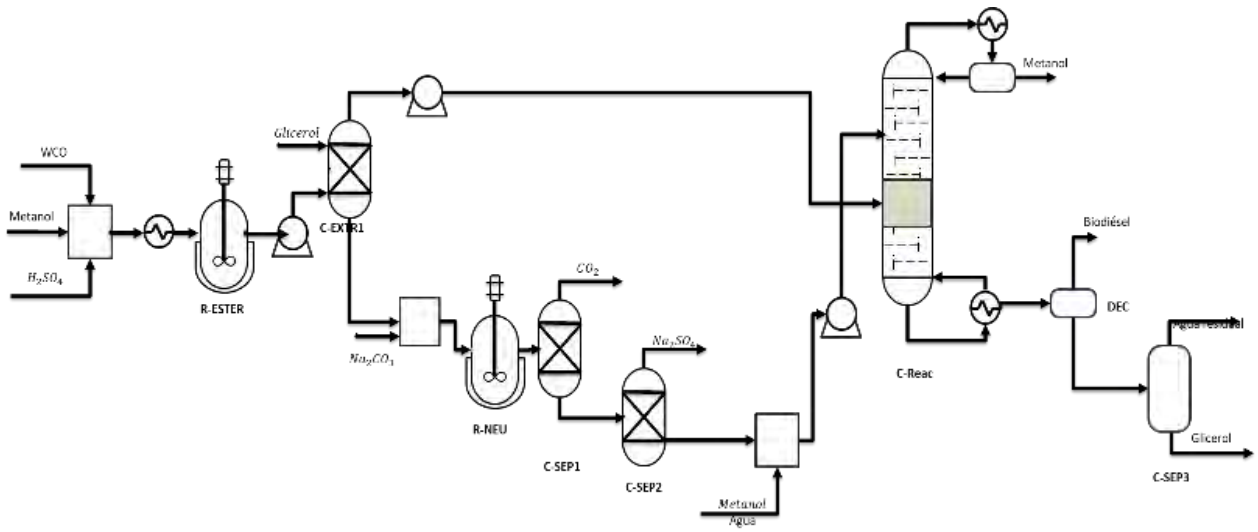

Figure 2. General intensification process treatment.

Conventional process

Figure 1 shows the conventional process proposed to produced biodiesel. The waste cooking oil pretreatment stage is implemented; this step is catalyzed by a sulfuric acid. At the outlet of the esterification reactor, the triglycerides together with the excess alcohol and the catalyst. The mixture passes to an extraction column where the triglycerides are separated from the alcohol and the catalyst, so it can move on to the neutralization stage and recover part of the alcohol. Downstream of the esterification reactor, there is a neutralization reactor, where the catalyst is neutralized, and the resulting salts are separated. The second reactive zone is where the transesterification of triglycerides to methyl esters take place. A plug flow reactor is used, where water is used as cooling fluid. With the help of a decanter, the biodiesel phase and the aqueous phase are separated; this last phase includes methanol, glycerol and water. The oily phase goes into an extraction column where the biodiesel is washed and separated from the impurities to finally obtain the refined biofuel. On the other hand, the aqueous phase is directed to two distillation columns where most of the methanol and glycerol present in the stream are recovered. Sensitivity analysis has been applied to the process equipment with available degrees of freedom, mainly the reactors and distillation columns, to obtain the best values operating variables, aiming to keeps costs down and maintain product quality. The goal was to maintain 99 % purity for both the glycerol recovery and the biodiesel.

Intensified process

Figure 2 shows the proposed intensified process. The initial step is the same of the conventional process, where the esterification reaction followed by the neutralization of the resulting mixture take place. Upon completion, the triglycerides exiting the extraction column are fed directly into the reactive distillation column. A fraction of the methanol originally allocated for esterification is redirected to facilitate the transesterification stage. The core of the process occurs within a reactive distillation column designed, a sensitivity analysis was performed to determine the operating parameters, using the minimum thermal load as the primary criterion while maintaining the required purity in the condenser, where the excess methanol is recovered. It was decided to maintain a reflux ratio variation ranging from 1.5 to 3.5, which are average values commonly found in the literature [14]. The design considered 18 equilibrium stages and a feed composition of 6:1, six parts methanol to one-part triglycerides. The reaction zone is concentrated within two specific reactive stages, the stages 7 to 9, containing a solid calcium oxide (CaO) catalyst. The sensitivity analysis was conducted by evaluating combinations of the number of stages and the feed tray locations. The column yields two distinct output streams, the top stream recovering the methanol, where this alcohol is collected as distillate for potential recycling into the system. The bottom stream consists of a mixture of fatty acid methyl esters (biodiesel), glycerol, and a residual percentage of unreacted triglycerides.

Final purification is achieved through a two-step separation sequence. First, the bottoms stream is sent to a decanter for primary biodiesel recovery, where two phases contained in the mixture, organic and inorganic, are separated. The operating conditions were 1 atm of pressure and 30°C of temperature. Subsequently, the remaining mixture is processed in a flash tank where the glycerol is recovered. This equipment was chosen with a view to obtaining the highest flow in the bottom stream, with the lowest thermal load in the column and recovering most of the glycerol. In addition to the difference in volatility of the components.

RESULTS

The production of biodiesel through the transesterification of waste cooking oil with methanol in a reactive distillation has been investigated in this work. Once the initial oil passes through the esterification reactor, the initial free fatty acid content goes from 6% to less than 0.9% by weight.

The simulations result for biodiesel production via conventional and intensified processes are presented in Table 2 and 3. The key findings highlight the advantages of process intensification in terms of material efficiency, energy, and product yield.

Table 2: Reactant feed streams

Materials (kg/h)	CP	IP
Methanol	1342	992
CaO	90	90
H_2SO_4	29	29
Na_2CO_3	28.5	28.5
Glycerol	288.2	288.25
Water	3500	0

Table 3: Products and by-products generated in simulations for biodiesel production.

Materials (kg/h)	CP	Purity (%)	IP	Purity (%)
CO_2	13	100	13	100
Na_2SO_4	26.82	100	26.82	100
Biodiesel	2887	99.9	2906.15	99.9
Glycerol	288.2	99.95	288.25	99.8
Waste water	3170	-	35	-

Table 4: Operating conditions of key equipment in the biodiesel production process.

Equipment	Variable	CP	IP
Esterification reactor	Temperature (°C)	70	
	Molar ratio of alcohol to oil	6:1	
	Conversion (%wt)	92	
	Free fatty acids (%wt)		
Equipment	**Variable**	**CP**	**IP**
Transesterification reactor	Temperature (°C)	60	-
	Molar ratio of alcohol to oil	6:1	-
	Conversion (%wt)	98	-
Equipment	**Variable**	**CP**	**IP**
Reactive	Molar ratio of	-	6:1

distillation column	alcohol to oil		
	Reactive stages	-	7-9
	Temperature in reactive stages	-	75.5-80
	Conversion (% wt)	-	99
	Number of stages	-	16
	Distillate flow (kg/h)	-	660
	Methanol purity (% mass)	-	99
Equipment	**Variable**	**CP**	**IP**
Biodiesel separation	Water flow (kg/h)	3500	0
	Temperature (°C)	30	30
	Biodiesel purity (%mol)	99.9	99.9
	Biodiesel flow (kg/h)	2866	2906.15
Equipment	**Variable**	**CP**	**IP**
Methanol Column	Number of plates	18	-
	Feed plate	6	-
	Flow rate (kg/h)	3, 675.71	-
	Methanol purity (%mass)	99	-
Equipment	**Variable**	**CP**	**IP**
Glycerol column	Number of plates	6	1
	Feed plate	2	1
	Flow rate (kg/h)	3, 182	414
	Methanol purity (%mass)	99	99

Table 5: Energy requirements of the process.

Equipment	CP (kW)	IP (kW)
Esterification reactor	1.9	1.9
Transesterification reactor	848.6	-
Reactive distillation column	-	1, 736.13
Biodiesel separation	792.5	157.4
Methanol column	1650.2	-
Glycerol column	15, 801.3	51.5
Total	19, 094.5	1946.9

When comparing the two processes, improvement in feedstock utilization yield is observed, increasing from 94 % with a conventional plug flow reactor (PFR) to 95 % with the reactive distillation. Furthermore, the transesterification reaction produces glycerol, a commercially valuable by product, yielding 288.25 kg/ hr in the intensified process and 288.19 kg/ hr in the conventional route. Since the conventional process generates a high-volume entering the first distillation column, it consequently requires significantly higher energy for purification. Regarding methanol consumption, the intensified process uses substantially less, which translates not only into raw material savings but also a direct reduction in condenser duty. While the reduction in the number of stages influenced the holdup, increasing the reaction rate. The most significant improvement in seen in the reboiler energy savings, around 66%. This is due to the heat integration by the transesterification reaction within the reactive column, where the heat is utilized directly for phase separation, unlike the sequence used in conventional distillation. Consequently, this results in an overall energy savings of 29% when comparing the complete intensified process to the conventional one.

The conventional process requires cooling services for the PFR, and water for biodiesel washing and methanol separation. The intensified process reduces wastewater generation. In the reactive distillation column, the methanol is recovered, simplifying the separation of the final product and by-products. This reduces energy costs, as there is no need for extraction equipment for biodiesel washing in the intensified scheme.

CONCLUSIONS

The design of two biodiesel production processes has been carried out in a simulation environment, allowing the analysis for the complete process to be carried out. Based on the results, the most convenient configuration for biodiesel production employs an intensified system. This configuration can produce high biodiesel yields by incorporating a pretreatment stage that significantly reduces free fatty acids, as well as the choice of favorable operating conditions.

The intensified process offered better performance in terms of biodiesel yield and energy consumption compared with the conventional one. The results demonstrated the benefits of adopting intensified processes in the transition toward more sustainable fuel production systems.

Ultimately, this study establishes that intensified biodiesel production, when integrated with renewable feedstocks and catalysts, plays a critical role in the global transition to low-carbon fuels. Advancing these technologies toward large-scale implementation will depend on coordinated progress in engineering efficiency, economic viability, and regulatory support. Such an integration ensures that biodiesel productions align with the broader goals of resource scarcity reduction and the establishment of sustainable energy systems.

ACKNOWLEDGEMENTS

The authors acknowledge the support provided by Universidad de Guanajuato. The authors also thank to the Secretaría de Ciencia, Humanidades, Tecnología e Innovación (SECIHTI, México) for the financial support provided through a Master´s a fellowship awarded to Tania Gisell Salgado-Rodriguez.

AUTHOR IDENTIFIERS

Author ORCIDs:
Salgado-Rodríguez TG: 0009-0009-7823-8709
Gómez-Castro FI: 0000-0003-4906-063X
Ramírez-Corona N: 0000-0003-3823-733X

REFERENCES

1. International Energy Agency (2021). Global primary energy demand growth by scenario 2019–2030. Available from: https://www.iea.org/data-and-statistics/charts/global-primary-energy-demand-growth-by-scenario-2019-2030., Access March 2025.
2. Sharma P, Usman M, Salama ES, Redina M, Thakur N, Li X. Evaluation of various waste cooking oils for biodiesel production: a comprehensive analysis of feedstock. Waste Management 136:219-229 (2021). https://doi.org/10.1016/j.wasman.2021.10.022
3. Petchsoongsakul N, Ngaosuwan K, Kiatkittipong W, Wongsawaeng D, Assabumrungrat S. Different water removal methods for facilitating biodiesel production from low-cost waste cooking oil containing high water content in hybridized reactive distillation. Renewable Energy 162:1906-1918 (2020). https://doi.org/10.1016/j.renene.2020.09.115
4. Zahan K, Kano M. Biodiesel production from palm oil, its by-products, and mill effluent: a review. Energies 11:2132 (2018). https://doi.org/10.3390/en11082132
5. Che F, Sarantopoulos I, Tsoutsos T, Gekas V. Exploring a promising feedstock for biodiesel production in mediterranean countries: a study on free fatty acid esterification of olive pomace oil. Biomass and Bioenergy 36:427-431 (2012). https://doi.org/10.1016/j.biombioe.2011.10.005
6. Motevali A, Hosseingholilou B, Ghafari A, Ghobadian B, Fayyazi E, Khoshnevisan B. Sustainable biodiesel production through process intensification: a comparative life cycle assessment. Energy Conversion and Management 350:120884 (2026). https://doi.org/10.1016/j.enconman.2025.120884
7. Segovia-Hernández JG, Hernández S, Bonilla Petriciolet A. Reactive distillation: a review of optimal design using deterministic and stochastic techniques. Chemical Engineering and Processing: Process Intensification 97:134-143 (2015). https://doi.org/10.1016/j.cep.2015.09.004
8. Simasatitkul L, Siricharnsakunchai P, Patcharavorachot Y, Assabumrungrat S, Arpornwichanop A. Reactive distillation for biodiesel production from soybean oil. Korean J. Chem. Eng. 28:649-655 (2011). https://doi.org/10.1007/s11814-010-0440-z
9. Aguilar-Aguilar FA, Mena-Cervantes VY, Hernández-Altamirano R. Analysis of public policies and resources for biodiesel production in méxico. Biomass and Bioenergy 196:107762 (2025). https://doi.org/10.1016/j.biombioe.2025.107762
10. Tomassetti S, Di Nicola G, Polonara F, Moglie M, Santori G. Liquid-liquid equilibria of biodiesel + glycerol and biodiesel + water binary systems. Fluid Phase Equilibria 595:114431 (2025). https://doi.org/10.1016/j.fluid.2025.114431
11. Souza MF, Hirata GF, Batista EAC. Evaluation of kinetics and thermodynamic parameters for simulation of palm oil biodiesel production. Fluid Phase Equilibria 525:112792 (2020). https://doi.org/10.1016/j.fluid.2020.112792
12. Pasha MK, Dai L, Liu D, Guo M, Du W. An overview to process design, simulation and sustainability evaluation of biodiesel production. Biotechnol Biofuels 14: (2021). https://doi.org/10.1186/s13068-021-01977-z
13. Gasca-González, R., Gómez-Castro, F. I., & Romero-Izquierdo, A. G. (2018). Design of a low-cost process for the production of biodiesel using waste oil as raw material. *Computer Aided Chemical Engineering*, *43*, 1529-1534. https://doi.org/10.1016/B978-0-444-64235-6.50267-9
14. Gonçalves PC, Monteiro LPC, Santos LS. Multi-objective optimization of a biodiesel production process using process simulation. Journal of Cleaner Production 270:122322 (2020). https://doi.org/10.1016/j.jclepro.2020.122322

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Integrated Multiproduct Facility for the Green Production of Chemicals and Food from Apples

Vanessa Villazón-León[a], Carlos Sanz[b], Adrián Bonilla-Petriciolet[a], Mariano Martín[b]*

[a] Instituto Tecnológico de Aguascalientes, Departamento de Ingeniería Química, Aguascalientes, Mexico
[b] Universidad de Salamanca, Departamento de Ingeniería Química, Salamanca, Spain
* Corresponding Author: mariano.m3@usal.es

ABSTRACT

An integrated multiproduct facility for the valorization of apple pomace in the green production of high value-added chemicals such as phenolic compounds and pectin, bioethanol, as well as apple juice, was optimally designed. The process includes green extraction technologies relying on subcritical water extraction and ethanol produced on-site through fermentation of residues. Two scenarios were evaluated: one based on purchasing the ethanol and another with on-site ethanol production. The units of the process were modeled using first principles. The process superstructure was formulated as a mixed-integer nonlinear programming problem, solved by decomposition. Investment and production costs of both alternatives were similar, with unit production requirements ranging from 1.09 €/L to 1.13 €/L. The discounted payback periods were 6.7 years and 6.4 years for the on-site ethanol production and purchased ethanol scenarios, respectively, while the internal rates of return were 36.4 and 37.8% for both processes, respectively. Both scenarios showed economic profitability with higher returns under the purchased ethanol scenario, but the on-site production scenario allowed waste reduction, self-sufficiency and ethanol reuse. A green extraction process based on subcritical water for the pectin extraction and the solvent recycling avoided the use of conventional acids and promoted the circular economy. The findings of this project demonstrated the potential of apple pomace as a valuable feedstock for integrated biorefinery applications in the agri-food sector.

Keywords: Integrated facility, Apples, Mathematical optimization, Apple pomace, Value-added products

INTRODUCTION

Food waste represents an important challenge leading to environmental problems due to the high volumes of residues generated (59 million tons alone in Europe in 2024) and the reduction of the supply chain efficiency [1]. Food waste valorization is an opportunity to mitigate this problem. Within this context, the beverage and juice processing industry not only generates large amounts of residues but also discards products with high value-added that could be recovered and reused. Fruit byproducts contain pectin, phenolic compounds, and fermentable sugars that can be transformed into value-added products and bioenergy contributing to the circular economy [2]. Orange peels spent coffee grounds, apple pomace, and olive oil residues have been used to produce value-added products [3].

The extraction process (i.e., the use of solvents, supercritical fluids or steam) to produce value-added products is a critical step. To evaluate and scale up the process, techno-economic analysis of the process using simulation-based approaches for comparing technologies is a promising tool [4]. Optimization-based design enables simultaneous synthesis, and mass and energy integration analyses. This allows the reduction of the need for external chemicals and utilities within the process. Polyphenols, oils, essential oils and pectin are the main value-added products extracted from food wastes. Furthermore, green solvents and end-use chemicals have been produced from these biomasses. This framework has been extended to the apple residues valorization.

The extraction of apple juice produces high volumes of apple pomace that can be valorized. The optimization of the process is carried out by developing a design of experiments for determining the optimal operating conditions. Apple pomace contains pectin, antioxidants and

polyphenols. Therefore, the development of cleaner extraction technologies such as subcritical water, microwave-assisted, supercritical CO_2, and ultrasound-assisted contribute to improve process efficiency and reduce solvent consumption.

Different approaches using commercial simulation software (i.e., SuperPro or Aspen Plus) to produce polyphenols, and to compare the biogas, pectin and bio-oil production have been reported in the literature. Vukušić et al. [5] have developed a pilot plant biorefinery analysis for producing apple juice, pellets, ethanol and pectin from apples. Okoro et al. [6] compared the production of lactic acid, xylitol and hydroxymethyl furfural using an optimization approach (NSGA-II algorithm within Aspen Plus). However, an integration of the facility has not been conducted. Commercial simulation software has limitations to model different biomass processing stages. Furthermore, their databases lack chemicals and properties needed to perform accurate modeling of processes.

In this study, an integrated facility was optimized to produce apple juice, high value-added products (e.g., pectin, phenolic compounds, and ethanol used as a solvent in extraction processes) based on circular economy concepts and using a mathematical optimization model. An economic analysis was carried out to assess economic feasibility of the facility under two scenarios: one producing ethanol on-site and another purchasing it externally. This framework promotes more sustainable strategies within the agri-food sector and contributes to maximizing economic benefits and efficiency.

PROCESS DESCRIPTION

The extraction processes for the pectin and phenolic compounds recovery require high volumes of solvents, which can be covered by the full utilization of apple residues for producing green ethanol by fermentation. The process has been divided into four sections as shown in Figure 1.

The apple juice production involves washing, milling, and pressing the fruit using pectinase enzyme to achieve higher yields in juice extraction. A clarification process is carried out to hydrolyze the pectin (responsible for cloudiness) with a conversion of 0.99 using 0.05 kg of immobilized pectinase per liter of juice for 6 hours at 30 °C. A stirred tank reactor operating for 30 min at ambient temperature in the presence of bentonite and gelatin enables the coagulation of protein complexes, tannins and suspended solids. Pasteurization is carried out using a continuous high-temperature short-time pasteurizer at 85 °C before bottling.

The apple residues generated during juice production are fed to a dryer (DR-01), followed by a mill (M-02) and pre-treating unit using ethanol for 20 min at 50 °C, ensuring the phenolic compounds extraction. The solid content is sent to the section of pectin extraction containing 30% moisture by weight. A flash separation unit enables the ethanol recovery before being reused. This unit operates at 90 °C. A molecular sieve (T-01), which operates at 70 °C, isolates phenolic compounds with 99.9% purity.

The pectin extraction is carried out in a stirred reactor (TK-04) through a subcritical water extraction process at 150 °C and 4.6 bar, with a yield of 0.17 g of pectin per gram of pomace. In the TK-04, cellulose and hemicellulose hydrolysis to fermentable sugars is achieved, with yields of 50 and 70%, respectively [7]. The pectin recovery is carried out employing ethanol at 95% purity in a stirred reactor (TK-05) operating for 20 min at 30 °C. The pectin is extracted as a gel and washed using ethanol to remove impurities. The solid pectin containing 10% moisture is isolated from the liquid phase using a press. Finally, a dryer (DR-02) reduces pectin humidity to 5%, before milling and storage.

The ethanol-rich streams from the extraction processes units are sent to the distillation column (DC-01) as a liquid at saturation conditions, for the ethanol recovery at 95% purity. The fermentation unit (R-01) is fed with the distillation column bottoms and operates at 38 °C for 24 h with 0.7 g of inoculum per liter of solution [8]. The main substrates assumed in the fermentation were xylose and glucose, with conversion rates of 0.83 and 0.955, respectively. The reactions involved in this process are shown below according to the data from the NREL database.

$$C_6H_{12}O_6 \longrightarrow 2C_2H_5OH + 2CO_2 \quad (1)$$

$$3C_5H_{10}O_5 \longrightarrow 5C_2H_5OH + 5CO_2 \quad (2)$$

$$C_6H_{12}O_6 + 1.2NH_3 \longrightarrow 6CH_{0.56}O_{0.5}N_{0.2} + 2.4H_2O + 0.3O_2 \quad (3)$$

$$C_5H_{10}O_5 + NH_3 \longrightarrow 5CH_{0.56}O_{0.5}N_{0.2} + 2H_2O + 0.25O_2 \quad (4)$$

A molecular sieve was used for the bioethanol purification up to 95% weight. The bed of the molecular sieve is regenerated with air at 70 °C and continuous operation is enabled by two parallel beds. The stream from the molecular sieve is mixed with the top of the distillation column and is recycled to the process, whereas the remaining ethanol is purified to meet fuel-grade purity requirements.

The scenario that considers purchased ethanol, shares the first three stages (juice production, phenolic compounds extraction and pectin recovery) and differs in the last stage, where only ethanol is recovered at 95% purity using a distillation column and the distillate is mixed with make-up ethanol and recycled to meet the plant's ethanol demand.

MODELING AND SOLUTION PROCEDURE

The units were modeled using first principles (i.e.,

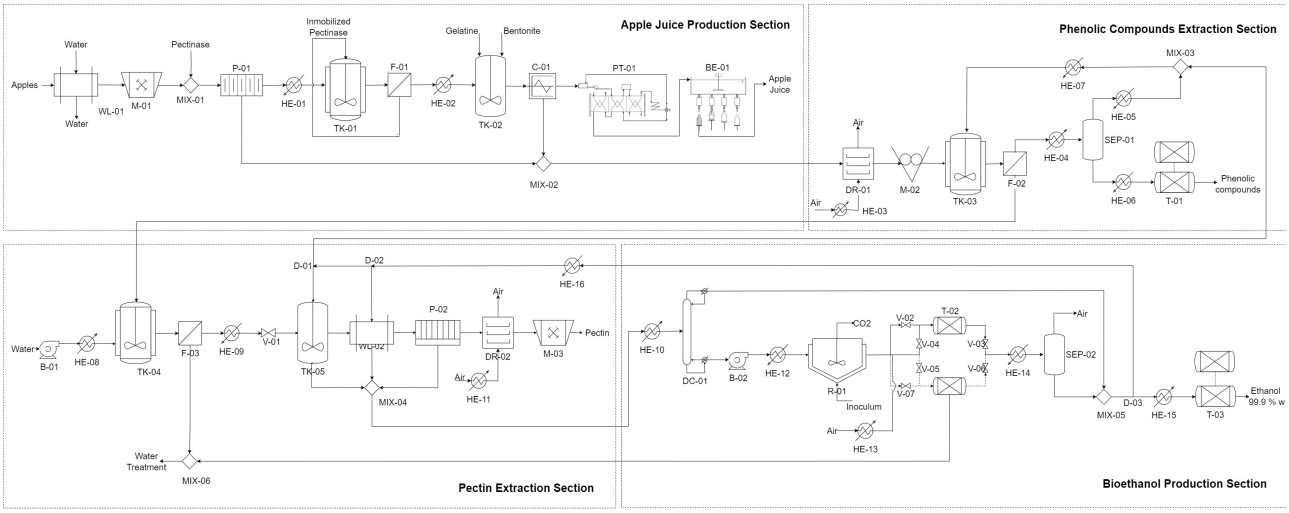

Figure 1. Flow diagram of the integrated facility divided by sections.

energy and mass balances), experimental kinetics, process yields, vapor-liquid equilibria and power consumption of equipment. The distillation column was modeled employing the CHEMCAD process simulator. The flash separator units were modeled employing vapor-liquid equilibrium relationships and the Antoine equation for the vapor pressure calculations. The problem was solved by maximizing a profit function comprising the product revenues, the cost of raw materials and utilities as follows:

$$Z = \sum_{i=Products} C_i \bullet flow(i) - \sum_{r=Raw\ materials} C_r \bullet flow(r) - \sum_{k=Mechanical\ equipment} C_e \bullet W(k) - \sum C_{Wa} \bullet Wa_{utility} - \sum C_{Steam} \bullet Steam_{utility} \quad (5)$$

where C_{Steam} is the cost of process steam, C_{Wa} is the water cost, C_e corresponds to electricity cost, W is the power consumption of mechanical equipment, C_r denotes the raw materials costs, and C_i represents the selling prices of the products obtained.

The superstructure of alternatives was represented by a mixed integer nonlinear programming (MINLP) model. Two scenarios (i.e., one producing ethanol on-site, and another purchasing it) were modeled by two nonlinear programming (NLP) formulations written in GAMS. The scenario involving the purchasing ethanol production is composed of 3, 854 equations and 4, 812 variables, while the scenario involving ethanol production on-site is composed of 4, 763 equations and 5, 979 variables. The problems were solved using a multistart optimization framework, selecting CONOPT as the preferred solver. After modelling all equipment, an economic evaluation was carried out.

RESULTS

The facility capacity was assumed to be 2000 kg/h of apples which corresponds to 40% of the production of the annual production in the region of Castilla y León, Spain [9]. The facility operates 24 h a day for 334 days in a year, considering a month for maintenance. A scaling study was carried out to evaluate economic feasibility.

Mass and Energy Balances

The energy and mass balances results for both scenarios are reported in Table 1. The needs for raw materials are equal in both cases, as they are associated with the production of apple juice. Ethanol is produced on-site under the first scenario, resulting in an additional amount of solvent generated (i.e., 16.61 L per m^3 of apple juice). However, in the purchased ethanol scenario, a total of 134 L of ethanol is needed to be acquired. Producing ethanol contributes to the circular economy principles, avoids dependence on external resources and promotes sustainable waste management practices.

Apple juice accounted for the largest production volume, exceeding 10, 000 m^3 per year. However, its selling price is low. A total of 3 t per year of phenolic compounds were produced, but they are high-priced, accounting for a substantial economic profit. Pectin was produced by 148 t per year, and its value is relatively high.

The first scenario requires 120.04 kWh of electricity per m^3 of apple juice, twice the energy requirements of the second scenario (67.58 kWh per m^3 of apple juice). This difference is due to the high energy consumption of the fermentation and ethanol recovery units. Cooling water and ammonia needs are similar in both scenarios. A

reduction in steam requirements in the ethanol purification process under the first scenario was achieved using the molecular sieve. Ethanol losses in this purification unit remained minimal (1%).

This study proposes the use of subcritical water for the pectin extraction, avoiding the need for sulfuric acid or citric acid and neutralizing agents compared to previous works [10]. This sustainable alternative also reduces operational risks without compromising process yields.

Table 1: Energy and material balances results for the study cases.

	Produced on-site ethanol	Purchased ethanol
Apples (t/year)	16, 032	16, 032
Water (m^3/m^3 juice)	9.05	9.05
Gelatin (kg/m^3 juice)	0.05	0.05
Bentonite (kg/m^3 juice)	0.18	0.18
Pectinase (kg/m^3 juice)	0.22	0.22
Inoculum (kg/m^3 juice)	3.97	
Ethanol consumed (L/m^3 juice)	-	134
Ammonia (kg/m^3 juice)	0.31	-
Apple juice (m^3/year)	10, 895.11	10, 895.11
Phenolic compounds (kg/m^3 juice)	0.29	0.29
Pectin (kg/m^3 juice)	13.67	13.67
Ethanol produced (L/m^3 juice)	16.61	
Power consumption (kWh/m^3 juice)	120.04	67.58
Steam (kg/m^3 juice)	20, 634.10	20, 631.50
Cooling water (m^3/m^3 juice)	780.57	785.08
Ammonia (kg/m^3 juice)	75.21	75.21

Economic Evaluation

An economic evaluation was conducted to compare the profitability of producing ethanol as a solvent on-site with purchasing it from the market. Investment and production costs were calculated, and a sensitivity analysis were performed to estimate several economic indicators. These indicators were evaluated by varying the ethanol market price. In addition, scale-up and scale down studies were carried out to assess the effect of production capacity.

Investment costs were estimated using the factorial method described by Towler and Sinnott [11] and employing correlations reported in the literature to estimate the equipment costs [12]. Figure 2 shows that investment costs are largely dominated by the ethanol recovery and production section, which represents around 50% of the total investment in both scenarios. This is because the distillation column and heat exchangers are expensive units, meaning that avoiding on-site ethanol production does not significantly reduce the investment. In the remaining stages, phenolic compound recovery involves higher investment costs than juice production, which partly explains the limited industrial adoption of waste valorization strategies without economic incentives.

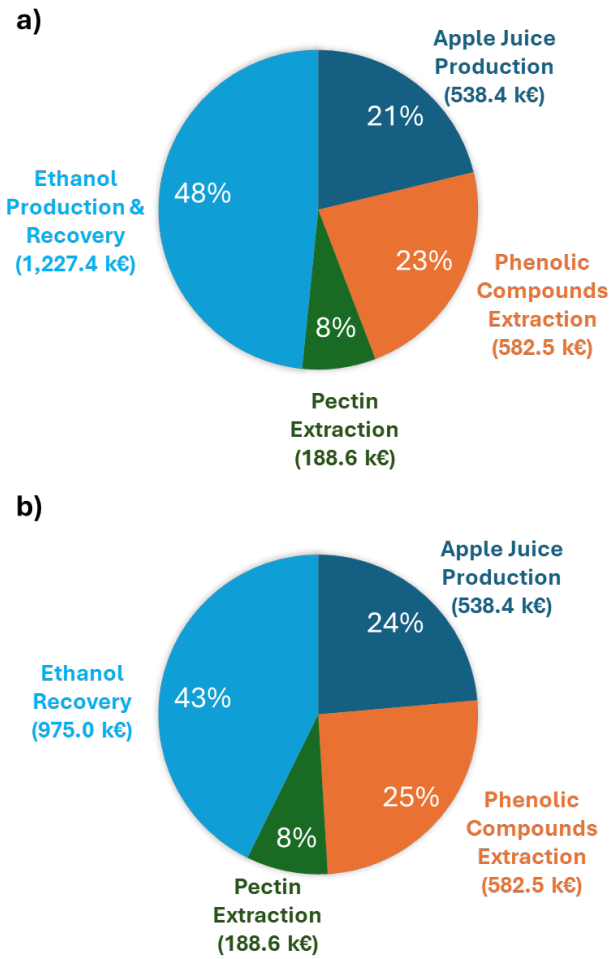

Figure 2. Investment cost distribution by stage for the on-site ethanol production scenario (a) and purchased ethanol scenario (b).

Production costs are classified into variable and fixed costs and were estimated following the procedure described in Towler and Sinnott [11]. Figure 3 shows that raw materials and utilities account for more than 70% of the total operating costs and therefore represent the highest expenditure in both scenarios. Purchasing ethanol slightly reduces utility consumption but increases raw material costs, resulting in only minor overall cost differences between both alternatives. Overall, both processes are economically viable, with unit production costs between 1.09 and 1.13 €/L, well below the market price of apple juice in Spain (1.5 €/L). In this way, due to

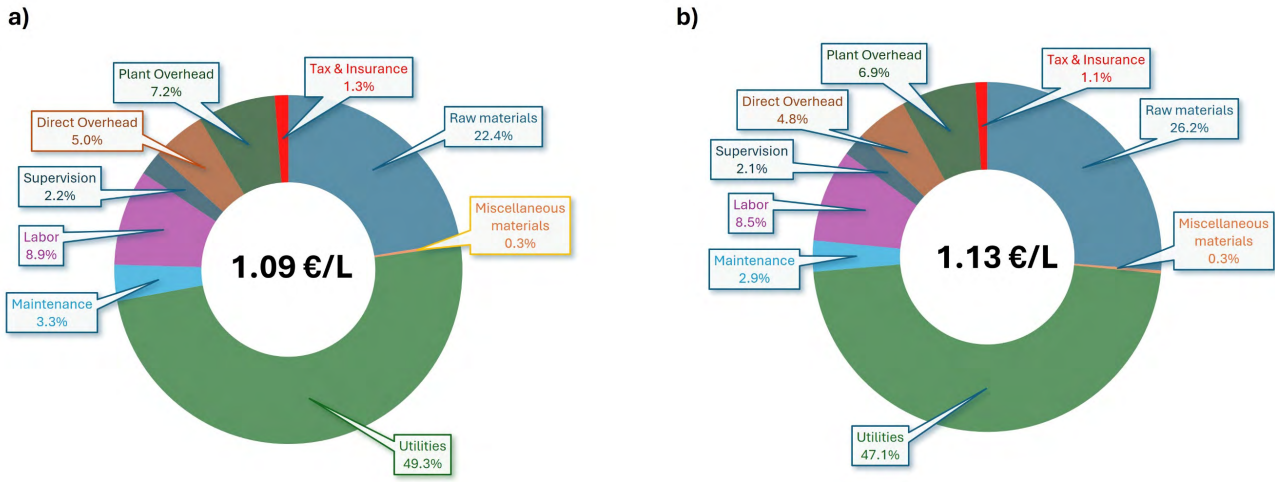

Figure 3. Breakdown of the production costs for the (a) ethanol produced on-site and (b) purchased ethanol (0.5 €/kg).

the difficulty in identifying the most attractive scenario, a sensitivity analysis is carried out using the internal rate of return (IRR) and the discounted payback period (DPP) indicators while considering several ethanol market prices (0, 25-1 €/kg). The results show that at low ethanol prices, purchasing ethanol is economically preferable, as it leads to higher internal rates of return and shorter payback periods. However, as ethanol prices and production capacity increase, on-site ethanol production becomes more competitive, lowering unit production costs and offsetting the higher investment.

The scaling results showed that purchasing ethanol is always cheaper from an investment point of view, whereas production costs are slightly higher in the ethanol on-site production scenario, regardless of production capacity. Nevertheless, these results are consistent with the conclusions drawn from the sensitivity analysis.

CONCLUSIONS

The optimization of an integrated facility to produce juice, phenolic compounds, pectin, and ethanol as a green solvent from apples was conducted. The scenario with on-site ethanol production promotes circular economy concepts and reduces the external dependence on solvents for the extraction processes. Both scenarios showed similar profitability despite ethanol production consumes substantial electricity in fermentation and ethanol recovery stages. However, the plant configuration is determined by the production or waste processing capacity. In general, the scenario based on purchasing ethanol from the market remains economically favorable. However, at large plant capacities and when ethanol prices increase above 0.5 €/kg, on-site ethanol production becomes a competitive alternative. Under these conditions, producing ethanol in situ not only mitigates exposure to external market fluctuations but also enhances process sustainability. In addition, the subcritical water extraction process avoids the need for acids and neutralizing agents. This study highlights the potential of apple pomace as a valuable feedstock for the development of sustainable and economically viable valorization pathways, supporting the transition to resource-efficient and sustainable agri-food systems.

ACKNOWLEDGEMENTS

The authors acknowledge the JCyL research project SA085P24.

AUTHOR IDENTIFIERS

Author ORCIDs:
Villazón-León V: 0000-0003-3820-5776
Sanz C: 0009-0008-9293-1780
Bonilla-Petriciolet A: 0000-0002-0197-3539
Martin M: 0000-0001-8554-4813

REFERENCES

1. Eurostat. Food waste and food waste prevention. European Union. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Food_waste_and_food_waste_prevention_-_estimates (accessed 2025-09-20).
2. B?aszczyk A, Sady S, Pacho?ek B, Jakubowska D, Grzybowska-Brzezi?ska M, Krzywonos M, Popek S. Sustainable management strategies for fruit processing byproducts for biorefineries: a review.

Sustainability 16:1717 (2024). https://doi.org/10.3390/su16051717
3. Martín M, El-Halwagi M. Biorefinery process design for the integrated valorization of biomass. NATO Science for Peace and Security Series C: Environmental Security :257-282 (2025). https://doi.org/10.1007/978-94-024-2345-7_16
4. Fonseca-Pérez RM, Almena A, Ramírez-Márquez C, Bonilla-Petriciolet A, Martín M. Techno-economic and environmental comparison of processes for the production of grape oil. Journal of Cleaner Production 441:141041 (2024). https://doi.org/10.1016/j.jclepro.2024.141041
5. Vukuši? JL, Millenautzki T, Cieplik R, Obst V, Saaid AM, Clavijo L, Zlatanovic S, Hof J, Mösche M, Barbe S. Reshaping apple juice production into a zero discharge biorefinery process. Waste Biomass Valor 12:3617-3627 (2020). https://doi.org/10.1007/s12649-020-01245-5
6. Okoro OV, Nie L, Hobbi P, Shavandi A. Valorization of waste apple pomace for production of platform biochemicals: a multi-objective optimization study. Waste Biomass Valor 12:6887-6901 (2021). https://doi.org/10.1007/s12649-021-01487-x
7. Pourbafrani M, Forgács G, Horváth IS, Niklasson C, Taherzadeh MJ. Production of biofuels, limonene and pectin from citrus wastes. Bioresource Technology 101:4246-4250 (2010). https://doi.org/10.1016/j.biortech.2010.01.077
8. Salakkam A, Phukoetphim N, Laopaiboon P, Laopaiboon L. Mathematical modeling of bioethanol production from sweet sorghum juice under high gravity fermentation: applicability of monod-based, logistic, modified gompertz and weibull models. Electronic Journal of Biotechnology 64:18-26 (2023). https://doi.org/10.1016/j.ejbt.2023.03.004
9. Junta de Castilla y León. Avances de superficies y producciones Agricultura y Ganadería. https://agriculturaganaderia.jcyl.es/web/es/estadistica-informacion-agraria/avances-superficies-producciones.html (accessed 2025-09-05).
10. Sanz C, Martín M. Developing a circular economy around jam production wastes. Ind. Eng. Chem. Res. 64:4991-5002 (2025). https://doi.org/10.1021/acs.iecr.4c03885
11. Towler G, Sinnott R. CHEMICAL ENGINEERING DESIGN: PRINCIPLES, PRACTICE AND ECONOMICS OF PLANT AND PROCESS DESIGN. Elsevier (2021).
12. Almena A, Martín M. Technoeconomic analysis of the production of epichlorohydrin from glycerol. Ind. Eng. Chem. Res. 55:3226-3238 (2015). https://doi.org/10.1021/acs.iecr.5b02555

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Hybrid Modeling of a Sewage-Sludge Gasifier using Flowsheet Simulation and Machine Learning

Malte Lutz[a1], William Würpel[b1], Fabian E. Habicht[a], Burcu Aker[c] and Jan C. Schöneberger[d*]

[a] Shit2Power GmbH, Berlin, Berlin, Germany
[b] DBFZ Deutsches Biomasseforschungszentrum gemeinnützige GmbH, Leipzig, Saxony, Germany
[c] Datacor, Inc., Florham Park, New Jersey, USA
[d] Berliner Hochschule für Technik, Department 8, Berlin, Berlin, Germany
* Corresponding Author: Jan.Schoeneberger@bht-berlin.de
[1] These authors contributed equally to this work.

ABSTRACT

This work presents a hybrid modelling approach for a downdraft sewage sludge gasifier within the Shit2Power (S2P) process. The gasifier is represented in CHEMCAD NXT by a series of four standard reactors that combine stoichiometric and equilibrium models with a data-driven correction step to account for deviations from ideal Gibbs equilibrium. Reaction conversions in the correction reactor are fitted to experimental synthesis gas compositions reported by Werle (2014) [7] for 30 operating points with varying equivalence ratios and reactor inlet air temperatures. The calibrated hybrid reactor model is evaluated against these data and shows conservative agreement for the combustible gas components of the synthesis gas. To overcome the limitations of linear interpolation between fitted operating points, several machine learning approaches are evaluated to predict the reaction conversions, and boosted neural networks are selected as a compromise between prediction accuracy and smooth behaviour in the operating window. After implementation in the CHEMCAD flowsheet, sensitivity studies with respect to equivalence ratio and reactor inlet air temperature show good agreement between predicted and fitted gas compositions, considering the limited data available. This demonstrates that the combination of flowsheet simulation and machine learning provides a promising framework for modelling sewage sludge gasification in the S2P process, but that significantly more experimental and training data are needed for practical application.

Keywords: gasification, sewage sludge, flowsheet simulation, hybrid model, machine learning, data-driven

1 INTRODUCTION

Across Europe, wastewater treatment facilities are adapting sludge-management strategies in response to evolving regulatory requirements and circular-economy objectives. The revised EU Urban Wastewater Treatment Directive (UWWTD) [1], in force since 1 January 2025, tightens requirements for nutrient removal and introduces quaternary treatment for micropollutant removal, while setting an energy-neutrality target for larger plants by 2045. These measures increase operational complexity and energy demand, strengthening the need for robust and compliant sludge treatment pathways.

In parallel, phosphorus has been classified as a Critical Raw Material at EU level and several countries are moving toward binding recovery obligations. This reinforces the strategic relevance of recovering phosphorus from sewage sludge, particularly via thermal routes that concentrate phosphorus in solid residues and facilitate downstream separation.

Together, these developments accelerate the shift away from land application and mixed-incineration toward mono-thermal, resource-oriented solutions. For many operators - especially small to medium plants - decentralized on-site concepts are favorable, as they reduce transport needs, mitigate dependency on constrained external capacities, and allow solutions tailored to local sludge characteristics.

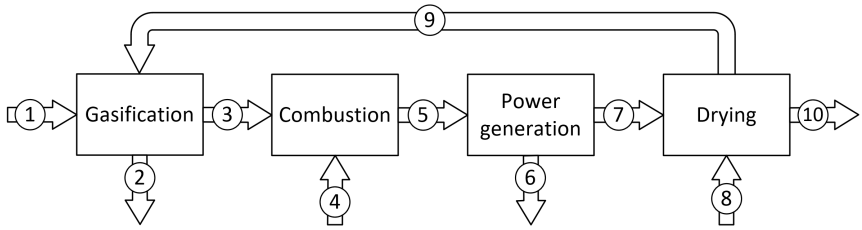

Figure 1. Block diagram of the Shit2Power process with the streams (1) air, (2) ash, (3) synthesis gas, (4) air, (5) hot flue gas, (6) electrical energy, (7) flue gas, (8) dewatered sewage sludge, (9) dried sludge, and (10) exhaust gas.

1.1 Shit2Power process

Shit2Power GmbH is developing a modular sludge-to-energy technology designed for decentralized deployment and plug-and-play integration at wastewater treatment plants. As shown in Figure 1, the system follows a four-module setup - sludge drying, gasification, synthesis gas combustion, and power generation - implemented as an energy-integrated process.

Dewatered sludge is first dried to a moisture content to less than 15 % using internally recovered heat. The dried feed is then converted in a dedicated gasifier under sub-stoichiometric air supply, producing a combustible synthesis gas (primarily H_2, CO, CH_4, CO_2 and N_2) and a solid ash fraction. The synthesis gas is subsequently oxidized in a purpose-designed combustion chamber, and the released thermal energy is converted to electricity in an ORC-based power generation unit. A defined part of the heat is recirculated to the dryer to enable self-sustaining operation, while surplus heat can be made available for site-specific external use.

1.2 Process simulation

In order to predict optimal operation regimes and to develop process control strategies, the Shit2Power (S2P) process is simulated using the CHEMCAD process simulation software from Datacor Inc. While the majority of the process steps can be adequately represented by physical models, no standard approach is available to model the complex reactions that govern the gasification process. To address this, the present study employs a hybrid approach that integrates first-principal models with data-driven models.

2 MODEL DEVELOPPEMENT

This section provides a comprehensive overview of the physical and chemical principles relevant to the gasification process. In addition, it offers a detailed description of the structure and sequence of the flowsheet simulation. The concept of the hybrid reactor model is derived, and the selected machine learning approach is described.

2.1 Gasification chemistry

Gasification is a thermo-chemical conversion process that transforms carbonaceous solids into a combustible synthesis gas under sub-stoichiometric oxygen supply. The primary components of the gas are hydrogen (H_2), carbon monoxide (CO), methane (CH_4), carbon dioxide (CO_2), and nitrogen (N_2). Additionally, this process leads to the formation of water vapor (H_2O) and other minor species, including tars (C_xH_x), ammonia (NH_3), hydrogen sulfide (H_2S), and hydrogen chloride (HCl). These species are of significance for the subsequent stages of gas cleaning and material selection because they can reduce the technical lifetime considerably. [2]

2.1.1 Reactor type and key reactions

In the present work, a downdraft fixed-bed gasifier is considered, which has been found to be suitable for a decentralized application with a thermal power well below 1 MW in prior investigations [3]. The dried sludge is introduced at the upper end of the reactor and moves downward through the bed, while air is injected into the reactor. The resulting gas stream follows the same direction. [2]

Four distinct reaction zones can be identified. In the drying zone (i), residual moisture is evaporated at temperatures between 100 °C and 200 °C in an oxygen-free environment. In the pyrolysis zone (ii), thermal decomposition occurs between approximately 150 °C and 500 °C mostly in the absence of oxygen. This process converts the sludge into char (C), gases (CO, CO_2, H_2, H_2O, CH_4), tars and ash. In the oxidation zone (iii), exothermic reactions with oxygen occur at temperatures of up to 1000 °C, resulting in the formation of carbon dioxide (CO_2), carbon monoxide (CO), and water (H_2O). These reactions provide the necessary heat for the endothermic reactions in the pyrolysis zone and the final stage of the gasification process, known as the reduction zone (iv). Here, a series of endothermic reactions lead to temperatures ranging from 500 °C to 700 °C. These reactions include the Boudouard, water–gas, hydrogasification, and water–gas shift reactions. Together, these reactions increase the calorific value of the gas by forming H_2, CO, and CH_4. [2, 3, 4]

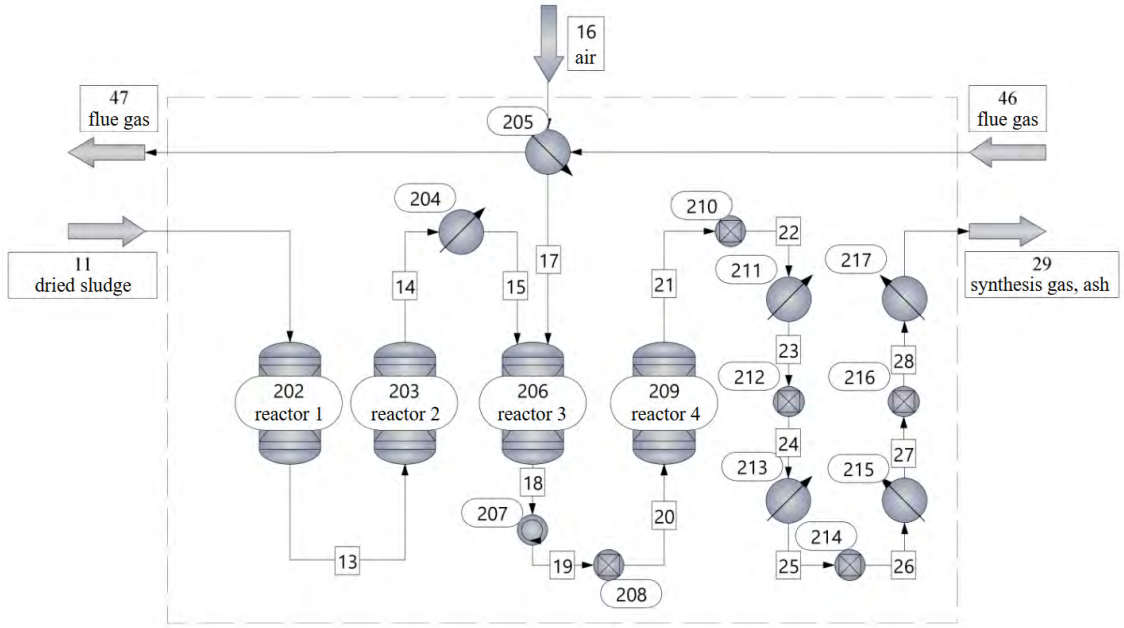

Figure 2. Image of the hybrid reactor model in the CHEMCAD simulation.

2.1.2 Operating parameters

Key factors influencing the gasification process in the described reactor configuration are the equivalence ratio (ER), the temperature of the supplied air, and the sludge properties - dry-matter content, lower heating value, and chemical composition. Within these variables, equivalence ratio and supply air temperature during operation can be influenced by the operator. The ER is defined as the ratio of the actual air supply to the air supply required for complete stoichiometric combustion. In the gasification of sewage sludge, the equivalence ratio is typically in the range of 0.15 to 0.35 [5, 6].

2.2 Flowsheet model

The flowsheet simulation is executed with CHEMCAD NXT (version 1.1.3.18169, Datacor Inc.), which employs a block-oriented representation of the plant, using unit operations to depict the processes. For each unit operation, mass and energy balances are solved using thermodynamic property models and, where applicable, equilibrium calculations. The tool holds a comprehensive library of thermodynamic models and pure component and mixture property data. The resulting process model is solved numerically using a sequential calculation approach.

2.2.1 Sewage sludge component

The sewage sludge, which serves as the primary input for the S2P process, is modeled as a mixture of water, inert ash (implemented as SiO2), and water- and ash-free (waf) sewage sludge. In the process of creating the waf sewage sludge component in CHEMCAD, the lower heating value and the elemental composition (carbon, hydrogen, oxygen, nitrogen, sulfur, chlorine) are taken into account.

For all simulations conducted as part of this study, the sewage sludge SS1 from Werle (2014) [7] is used.

2.2.2 Hybrid reactor model

Due to the highly complex reaction process involved in thermochemical gasification, it is not possible to rigorously model all reactions. Therefore, a hybrid reactor model that integrates a physical and a data-driven approach is developed.

As shown in Figure 2, a series of four CHEMCAD standard reactors is employed. The feed streams are dried sewage sludge (11) and preheated gasification air (16). The outputs of this process consist of a mixture of synthesis gas, ash and other by-products (29), which are separated in a subsequent simulation step.

The initial reactor employed is a stoichiometric reactor (202), in which the sewage sludge is completely broken down into its elementary components. These act as reactants for the following three reactors.

The second reactor is an equilibrium reactor (203). Here, relevant side reactions take place, producing the following by-products: ammonia, hydrogen sulfide, and hydrochloric acid. In the remaining two reactors of the hybrid reactor model, these substances are considered inert. [8]

In reactor (206), adiabatic minimization of Gibbs free energy results in a product gas containing the following main components: hydrogen, carbon monoxide, carbon dioxide, methane, and nitrogen. Tars are neglected. The quantity of oxygen present is governed

through the settings of a feedback controller (208). The airflow (16) is preheated using a heat exchanger (205). The heat required for this process is extracted from the combustion chambers flue gas flow (46) .

However, complete minimization of the free Gibbs energy during gasification is not achieved in reality. Consequently, the deviations in the synthesis gas composition must be adjusted to experimental results. To do this, the fourth equilibrium reactor (209) is employed, embodying the data-driven modeling approach. Five reverse reactions (equation 1 – 5) are incorporated in this reactor, which cause the products of the Gibbs reactor to react back.

$$\text{R1: } 2\,CO \xrightarrow{X_{R1}} C + CO_2 \quad (1)$$

$$\text{R2: } CH_4 \xrightarrow{X_{R2}} C + 2\,H_2 \quad (2)$$

$$\text{R3: } CO + H_2 \xrightarrow{X_{R3}} C + H_2O \quad (3)$$

$$\text{R4: } 3\,H_2O + CO \xrightarrow{X_{R4}} CH_4 + H_2O \quad (4)$$

$$\text{R5: } H_2 + CO_2 \xrightarrow{X_{R5}} CO + H_2O \quad (5)$$

Reactors 1, 2, and 4 are calculated isothermally, with each heat duty subsequently being taken into account in the energy balance of the synthesis gas stream via four heat exchangers (211, 213, 215, 217) and controller blocks (208, 210, 212, 214, 216). Heat exchanger (204) is inactive. It can be used to introduce the temperature at which the Gibbs equilibrium is calculated as an additional degree of freedom for the data reconciliation.

2.2.3 Data Reconciliation

In the calibration of the reactor model, the reaction conversions X_{Ri} in reactor 4 (209) (reactions 1 - 5) are adjusted to align the simulated synthesis gas concentrations with the experimental measurements. In addition to the combustible components - hydrogen, methane, and carbon monoxide - the non-combustible gas components nitrogen and carbon dioxide are also taken into account at this point. Werle (2014) [7] provides experimental data for 30 operating points as a function of the gasification air temperature (T = 25, 100, 150, 200, and 250 °C) and the equivalence ratio (ER = 0.12, 0.14, 0.16, 0.18, 0.23, and 0.27). The conversion rates are fitted using CHEMCAD's proprietary data reconciliation tool. In this optimization task multiple restarts are used in order to overcome local minima.

The results of the data reconciliation can be found in Fig. 6 and Fig. 7, where calculated and experimental compositions are compared. The resulting errors are shown in Table 1. In general, the simulated molar fractions of the primary synthesis gas components are lower than the experimental values. The model underestimates the content of the calorific gases, thereby providing conservative predictions. The errors are smallest at the boundaries of the investigated equivalence ratio range and largest in the central region. The most significant discrepancies emerge for the CO_2 molar fraction, exhibiting a relative error of 50 % (absolute error of 7.6 mol%). However, the impact on the overall process assessment is mitigated by the negligible contribution of CO_2 to the syngas lower heating value. The authors are aware that the amount of data used is very limited and that for a more robust model more data points are required.

Table 1: Errors between experimental [7] and fitted synthesis gas concentrations.

Parameter	Rel. error, avg. [%]	Absolute error [mol%]		
		avg.	min.	max.
y_{H2}	6	0.4	< 0.01	1.6
y_{CO}	29	7.7	2.8	13.4
y_{CO2}	50	7.6	3.5	10.5
y_{CH4}	6	0.1	< 0.01	0.2
y_{N2}	29	15.0	7.5	21.4

2.3 Machine learning approach

Only a limited number of 30 operating points are available from the literature data [7]. Therefore, interpolation is necessary to investigate intermediate operating points. After conducting initial tests with linear interpolation, this study focuses on implementing and evaluating a machine learning (ML) approach. This approach has several advantages over linear interpolation:

- No requirement for a structured data grid
- Simplified data handling and improved usability
- Easier integration of additional sewage sludge parameters
- Representation of nonlinear relationships
- Potential for extrapolation, given a sufficient data basis

2.3.1 Training and test data

The training data for the machine learning model consists of the inputs equivalence ratio and reactor inlet air temperature, and the outputs are the conversion rates X_{R1} to X_{R5} of reactions R1 to R5 (see eq. 1 - 5). The conversion rates were obtained by calibrating the reactor model to the synthesis gas compositions that were experimentally determined (see section 2.2.3).

In the interest of expanding the functionality of the reactor model, it is also possible to introduce additional inputs, provided that a suitable data basis is available. For instance, if data from other sewage sludges is available, the sewage sludge composition and the lower heating value can be considered.

The resulting dataset consists of 30 samples structured in an evaluable CSV file. The sample size is very

limited for training a neural network. Consequently, the present study places its primary emphasis on the methodology of hybrid modelling.

2.3.2 Model screening

The commercial machine learning software JMP Pro was used for model selection. Following the importation of the training data, a model comparison of the widely used machine learning methodologies (Neural Boosted, Boosted Tree, Bootstrap Forest, Support Vector Machines, Fit Least Squares, Fit Stepwise, Generalized Regression Lasso, K Nearest Neighbors) was conducted automatically. For each of these algorithms, JMP Pro generated five individual models with ER and reactor inlet air temperature as inputs and a chemical reaction conversion (X_{Ri}) as the output.

2.3.3 Integration in CHEMCAD

The utilization of a VBA function module facilitates the implementation of the machine learning model into the flowsheet simulation. The so-called VBA unit is designed to execute VBA code during the simulation process, thereby allowing it to access the inputs for the machine learning model from the simulation and subsequently send the outputs back to CHEMCAD. The machine learning model itself can be implemented in two ways: either directly as VBA code or via a Python interface, controlled from the VBA unit.

3 RESULTS AND DISCUSSION

This section shows the result of the selection of the most suitable ML model. Subsequently, the performance of the hybrid reactor model is examined and discussed.

3.1 Data analysis and model selection

In order to conduct a meaningful comparison of the machine learning approaches, Figure 3 shows the R^2 values of the best-performing algorithms, which are Boosted Tree (BT), boosted neural network (NB), and Support Vector Machine (SVM). For each approach, the five different models were considered; together, they provide predictions for all five outputs.

It is evident that there are substantial issues with all models in terms of their capacity to accurately predict the outputs X_{R2} and X_{R5}. The R^2 values for X_{R2} range between 0.0833 (FLS) and 0.2927 (BT), and for X_{R5} between 0.1918 (NB) and 0.5816 (BT). These values are very low. This finding suggests the absence of dependencies between the corresponding training values, indicating a data issue. The remaining outputs, X_{R1}, X_{R3}, and X_{R4}, can be predicted with a reasonable degree of accuracy. The boosted tree shows the best performance, with an average R^2 of 0.7302, followed by the boosted neural network with an average R^2 of 0.609.

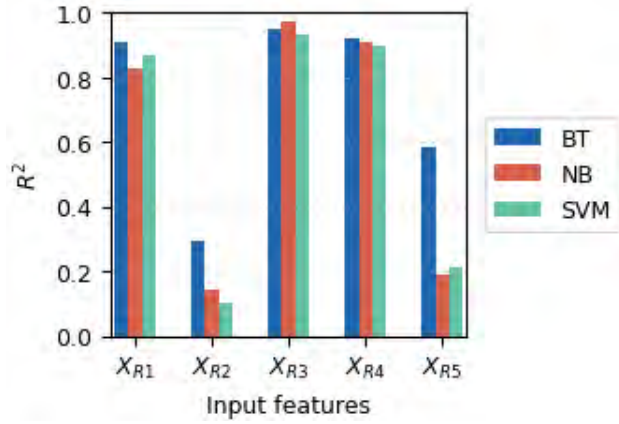

Figure 3. Comparison of R^2 for the predictions of X_{R1} to X_{R5}; for each value one of the following machine learning approaches was used: boosted tree (BT), boosted neural network (NB), and support vector machine (SVM).

In order to evaluate the suitability of the two best machine learning approaches, Figure 4 presents X_{R2} as a function of ER and the reactor inlet air temperature (T). In contrast to the smooth curve of the boosted neural network, the boosted tree exhibits pronounced discontinuities. The X_{R2} value shows sudden rises or falls around the points of the training data. This behaviour indicates that the use of the boosted tree is not suitable for the reactor model, as it does not permit reliable interpolation between the fitted chemical conversions X_{R1} to X_{R5}.

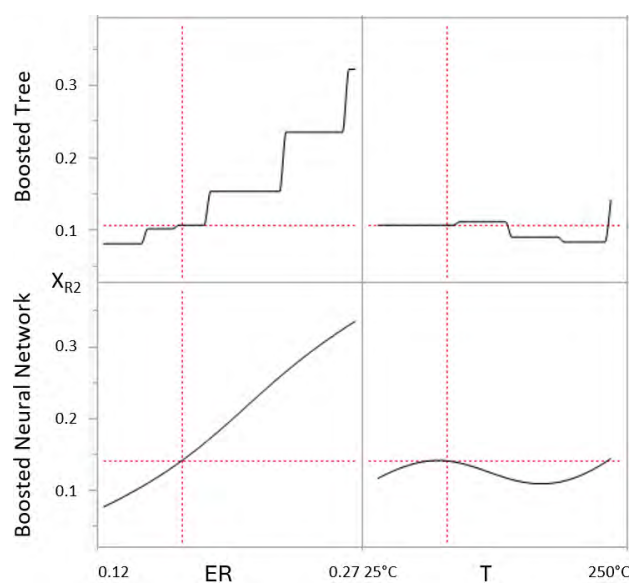

Figure 4. Comparison of the X_{R2} curve for a boosted tree and a boosted neural network as a function of ER and the reactor supply air temperature.

Consequently, boosted neural networks was selected for implementation in this study. The number of layers and neurons is documented in Table 2, along with the R^2 values for training and validation. The validation

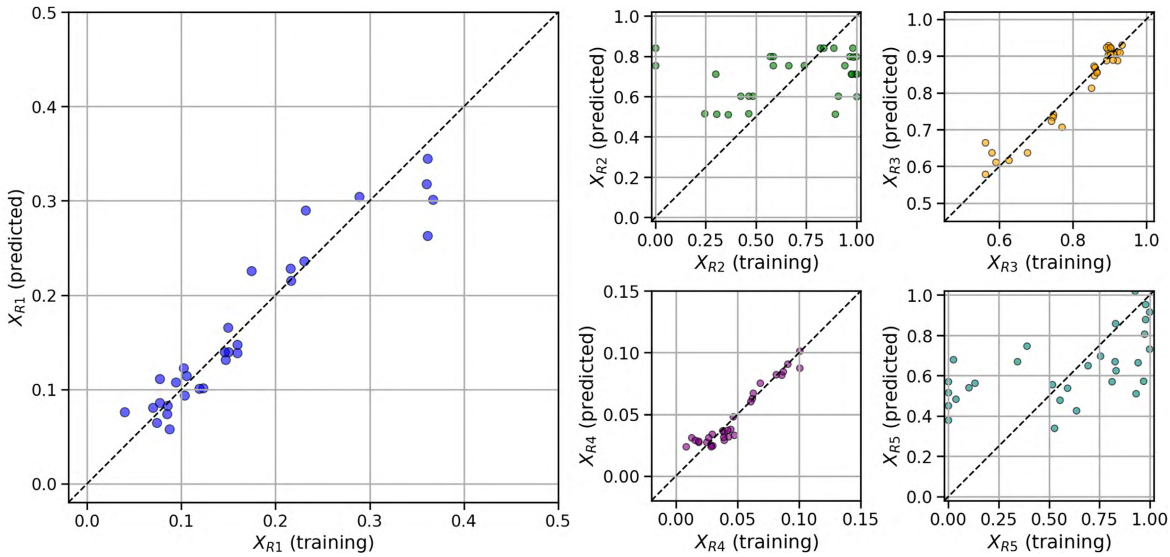

Figure 5. Comparison of X_{R1} to X_{R5} training data with the prediction values of the boosted neural networks.

was performed using random holdback, which did not indicate any overfitting for the evaluated models. The hyperbolic tangent (tanh) function was utilised as the activation function. Furthermore, Figure 5 compares the chemical conversions X_{R1} to X_{R5} from the training data with the predicted values of the boosted neural networks.

As already indicated by the data in Table 2, the conversions X_{R2} and X_{R5} cannot be predicted accurately (see Figure 5). It is notable, that these conversions range from 0 to 1, while the others scatter in a much smaller range. In a future work this difference in range will be analyzed more in detail.

Table 2: Documentation of the architecture and performance of the boosted neural networks X_{R1} to X_{R5} containing the number of layers, neurons, and R2 of the training and validation.

Model	Number of		R^2	
	layer	neurons	train.	valid.
X_{R1}	3	60	0.8259	0.9049
X_{R2}	3	6	0.1458	0.1083
X_{R3}	3	60	0.9710	0.9028
X_{R4}	3	60	0.9098	0.8177
X_{R5}	3	21	0.1918	0.2323

3.2 Performance of the hybrid reactor model

After the implementation of the five boosted neural networks in the CHEMCAD flowsheet, a series of sensitivity studies was performed. The ER and the reactor inlet air temperature were varied, and the composition of the synthesis gas was logged. The results are shown in Figures 6 and 7, with the left-hand side showing examples of the curves at the highest (250 °C) and lowest (25 °C) temperatures over the variable ER and the right-hand side displays the curves as a function of the reactor inlet air temperature and an ER of 0.12 and 0.27. In order to facilitate comparison between these graphs, both the values as determined through experimentation [7], and the composition based on the fitted chemical conversions, which act as training data for the neural network, are displayed. The dashed lines in the figure represent the composition of the synthesis gas, which can be simulated by using linear interpolation between the fitted data points. It is evident that the predicted concentration curves align closely with the fitted concentrations, particularly in consideration of the limited dataset. Furthermore, an error calculation was performed between the fitted and predicted concentrations for all 30 operating points, not solely the visualized ones. The results of this calculation are presented in Table 3.

Table 3: Average errors between the fitted and predicted concentrations of H_2, CO, CH_4, CO_2, and N_2.

	H_2	CO	CH_4	CO_2	N_2
Rel. error [%]	11.4	4.7	8.4	5.7	1.6
Abs. error [mol%]	0.72	0.84	0.10	0.43	1.12

The data shows that the absolute errors, despite the relatively large relative errors of up to 11.4% (H_2), are between 0.10 mol% (CH_4) and 1.12 mol%. Given the limited data available, these errors are considered minor, and they suggest that the machine learning-optimized hybrid reactor model can plausibly simulate the synthesis gas composition.

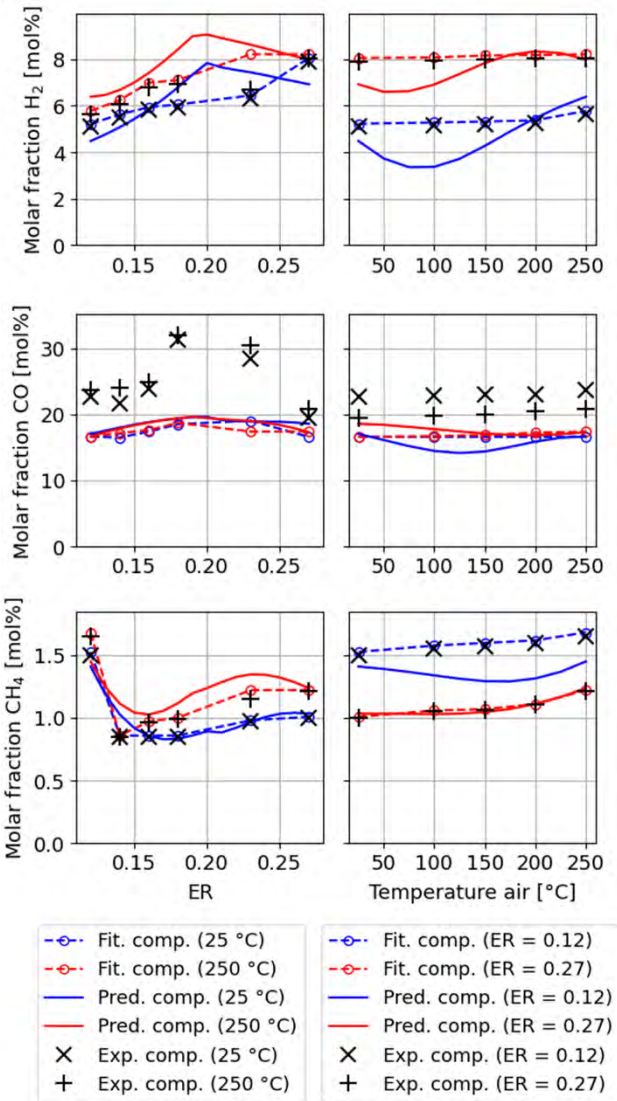

Figure 6. Concentration curve of the oxidizable gases H_2, CO, and CH_4 as a function of ER (left) and reactor inlet air temperature (right); comparison of experimental, fitted, and predicted concentrations.

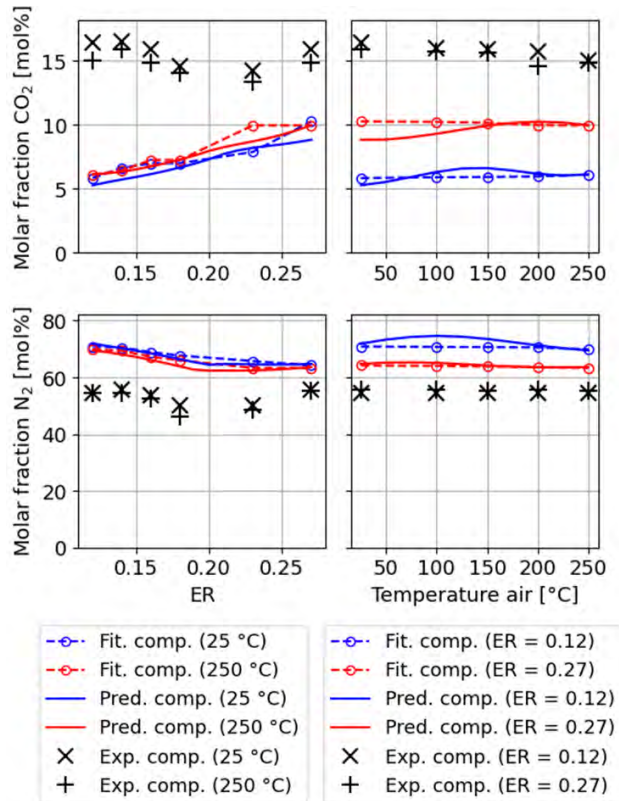

Figure 7. Concentration curve of the non-oxidizable gases CO_2 and N_2 as a function of ER (left) and reactor inlet air temperature (right); comparison of experimental, fitted, and predicted concentrations.

4 CONCLUSIONS

The present study develops and evaluates a hybrid modelling approach for a downdraft sewage sludge gasifier within the Shit2Power process. A four-reactor concept is implemented in CHEMCAD, in which a sequence of standard reactor models is combined with a data-driven correction step to account for the deviation from ideal Gibbs equilibrium. The correction is based on reaction conversions fitted to synthesis gas compositions [7]. This enables the model to reproduce the behavior of real sewage sludge gasifiers with conservative estimates of the energy content of the resulting syngas.

On this basis, several machine learning approaches were systematically compared regarding their suitability to predict the reaction conversions in the hybrid reactor model. Although boosted trees achieved the highest average coefficient of determination (R^2), they exhibited pronounced discontinuities and were therefore deemed unsuitable for interpolation between experimental operating points. Boosted neural networks were ultimately selected as the most appropriate compromise between prediction accuracy and smoothness, providing satisfactory representations for three of the five reaction conversions, given the small training dataset of 30 samples. After integration into the CHEMCAD flowsheet simulation, the resulting hybrid reactor model reproduces the fitted synthesis gas compositions with good agreement, while maintaining the overall mass and energy balances.

Overall, the results show that the combination of rigorous flowsheet simulation and machine learning provides a promising framework for modelling complex thermochemical conversion processes such as sewage sludge gasification. Within the Shit2Power concept, the hybrid reactor model provides a robust basis for a digital representation of the gasification reactor and enables future work on process optimization, advanced control, and evaluation of varying sewage sludge characteristics.

DIGITAL SUPPLEMENTARY MATERIAL

The VBA code, which includes the interface between the Python machine learning model and the CHEMCAD simulation, can be viewed in the Living Archive for Process Systems Engineering under ID LAPSE:2026.0012.

ACKNOWLEDGEMENTS

We would like to express our gratitude to the German Federal Ministry for Economic Affairs and Climate Action and the European Social Fund, which provided financial support for the research conducted for this study through the exist Gründungsstipendium.

AUTHOR IDENTIFIERS

Author ORCIDs:
Lutz M: 0009-0002-5782-568X
Habicht FE: 0000-0002-7379-4636
Schöneberger JC: 0009-0008-8296-4810

REFERENCES

1. European Parliament, Council of the European Union. Directive (EU) 2024/3019 of the European Parliament and of the Council of 27 November 2024 concerning urban wastewater treatment (recast). (2024) https://eur-lex.europa.eu/eli/dir/2024/3019/oj/eng
2. Kaltschmitt M, Hofbauer H, Lenz V. Energie aus Biomasse: Thermo-chemische Konversion. Springer Vieweg (2021).
3. Basu P. Pyrolysis and torrefaction. Biomass Gasification Design Handbook :65-96 (2010). https://doi.org/10.1016/b978-0-12-374988-8.00003-9
4. Rösch C, Wintzer D. Vergasung und Pyrolyse von Biomasse. 2. Sachstandsbericht zum Monitoring "Nachwachsende Rohstoffe". Büro für Technikfolgen-Abschätzung beim Deutschen Bundestag (TAB) (1997)
5. Manyà JJ, Sánchez JL, Gonzalo A, Arauzo J. Air gasification of dried sewage sludge in a fluidized bed: effect of the operating conditions and in-bed use of alumina. Energy Fuels 19:629-636 (2005). https://doi.org/10.1021/ef0497614
6. de Andrés JM, Vedrenne M, Brambilla M, Rodríguez E. Modeling and model performance evaluation of sewage sludge gasification in fluidized-bed gasifiers using aspen plus. Journal of the Air & Waste Management Association 69:23-33 (2018). https://doi.org/10.1080/10962247.2018.1500404
7. Werle S. Impact of feedstock properties and operating conditions on sewage sludge gasification in a fixed bed gasifier. Waste Manag Res 32:954-960 (2014). https://doi.org/10.1177/0734242x14535654
8. Di Fraia S, Massarotti N, Uddin MR, Vanoli L. Conversion of sewage sludge to combined heat and power: modeling and optimization. Smart Energy 5:100061 (2022). https://doi.org/10.1016/j.segy.2021.100061

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Lifetime-Adjusted LCA of Biochemical and Thermochemical Circular Plastic Pathways

Alexandra Krestnikova[a], Gonzalo Guillén-Gosálbez[a]*

a Institute for Chemical and Bioengineering, Department of Chemistry and Applied Biosciences, ETH Zürich, 8093 Zürich, Switzerland
* Corresponding Author: gonzalo.guillen.gosalbez@chem.ethz.ch

ABSTRACT

The transition from a linear, fossil-based polymer economy to a circular bio-economy is critical for mitigating resource depletion and greenhouse gas emissions. This study provides a rigorous comparison of two biomass-to-plastic pathways: a biochemical route (PLA via enzymatic hydrolysis) and a thermochemical route (bio-PE via gasification and MTO). Based on Aspen Plus simulations and a "lifetime-adjusted" lifecycle assessment framework, we evaluate the environmental performance of these routes in the transition from linear to circular systems. Unlike standard "cut-off" methods, the lifetime-adjusted model accounts for virgin make-up and molecular retention across multiple recycling cycles. Results indicate that at current 15% recycling rates, PLA exhibits the lowest global warming potential due to significant biogenic carbon sequestration. However, as recycling rates reach 75%, process efficiency becomes the dominant factor; the precise biochemical recycling of PLA continues to outperform mechanical HDPE recycling, whereas the energy-intensive thermochemical bio-PE route loses its competitive advantage. Our sensitivity analysis reveals that LCA modeling choices significantly diverge at intermediate recycling rates, resulting in different values for environmental impact. Ultimately, while bioplastics serve as a vital agent for CO_2 sequestration during the transition to circularity, long-term sustainability necessitates a shift toward "recyclable-by-design" materials.

Keywords: Aspen Plus, Biomass, Environment, Polymers, Life Cycle Analysis, Circular Economy

INTRODUCTION

Currently, polymer production is mostly based on fossil feedstocks, leading to burdens, such as pollution, resource depletion and green house gas emissions [1]. Nonetheless, production volumes continue to rise, especially driven by the plastic packaging sector, which made up over 30% of the global plastic demand in 2019 [2], where recycled plastics remain the exception, with recycling rates reaching 15% at most [3]. The sector's reliance on linear manufacturing and single-use packaging threatens to intensify its environmental crisis [1], highlighting the need for a circular plastics economy, leveraging various renewable technologies [4].

Biomass, especially agricultural residues such as wheat straw, emerges as a promising renewable feedstock for CO_2 sequestration, and has multiple processing pathways: biochemical and thermochemical. Polylactic acid (PLA) is currently leading the global bioplastic production [5] as a versatile packaging material and follows a biochemical processing path, using hydrolysis and subsequent fermentation to produce lactic acid [6]. Meanwhile bio-polyethylene (bio-PE) is a "drop-in" plastic, that can be produced thermochemically via biomass gasification [7] and methanol-to-olefins (MTO) technology [8]. However, bioplastics alone will not solve the environmental crisis in the plastics sector and can contribute to added strain on land use [4].

A possible solution for a more sustainable polymer economy is combining chemical recycling, with renewable feedstocks, which saves virgin resources and facilitates the shift away from fossil-dependence. While bioplastics are currently mostly landfilled, they have great potential for chemical depolymerization, conserving their polymer quality, as opposed to mechanical recycling. PLA can be recycled using enzymatic hydrolysis, which is a very precise, single step biochemical process [9]. Similar to its production, bio-PE can be recycled

thermochemically by polymer gasification and MTO [10, 11]. While previous studies have evaluated these pathways in isolation, a rigorous comparison of these technologies is missing to our best knowledge.

Furthermore, standard lifecycle assessment (LCA) "cut-off" methods often fail to capture the long-term value of molecular retention and do not account for the impact of virgin make-up. This leads to comparisons of isolated processes that neglect the synergies between chemical recycling and the use of renewable makeup, where the latter can bring additional benefits relative to the standard consumption of fossil make up in circular routes. This assumption fails to accurately quantify the true impact of plastics based on the recycling rate, the virgin make-up and the end-of-life treatment of the losses. In contrast, this paper uses a "lifetime- adjusted" framework, following the work of Nicholson et al. [12], to evaluate the global warming potential (GWP) of these technologies as they transition from linear to circular systems and assess the difference between the cut-off model and lifetime-adjusted model for impact calculation.

PROCESS DESCRIPTION & SIMULATION

Biochemical Pathway

The simulations are carried out in Aspen Plus V15. ELECNRTL is used as the property model, and missing model parameter values are taken from NREL's Physical Property Database for Biofuels Components [13] or estimated with the UNIFAC method. The considered substances are modeled as conventional substances, with the exception of biomass, enzymes, PLA, and calcium lactate, which are modeled as solids. The reactors are modeled with RSTOIC units, except for the polymerization of the lactide, where a batch is used. The flashes and distillation columns are modeled using FLASH2 and RADFRAC respectively. The process flowsheets and inventories can be found in the SI.

Production begins with steam explosion of wheat straw (160°C, 0.62 MPa) [14] to hydrolyze hemicellulose into xylose, followed by organosolv pretreatment (170°C, 65% ethanol, 0.5% H_2SO_4) to solubilize lignin [15]. The remaining cellulose is enzymatically hydrolyzed (50°C) to glucose [16].

These sugars are fermented (37°C, pH 7) into lactic acid (LA) with 90% yield and 99.5% optical purity [17]. Neutralization with lime forms calcium lactate, which is acidified with sulfuric acid to liberate LA and remove gypsum [18, 19]. Purification utilizes reactive distillation via esterification and hydrolysis, achieving commercial concentrations of 85-88% [20].

PLA polymerization follows the continuous ring-opening polymerization (ROP) route. LA is dehydrated (130°C, 0.5 atm) to a prepolymer, then thermally depolymerized (180°C, 0.1 atm) to cyclic lactide. Purified lactide undergoes ROP in the melt using a tin octanoate catalyst (180°C, 4h) with 95% conversion [21, 22]. Residual monomers are devolatilized and recycled.

Enzymatic recycling employs PLA-depolymerases to degrade the polymer back to monomers [23]. PLA is hydrolyzed in water (70°C, 230 g/L solids) [24], followed by neutralization and acidification as described for LA production from biomass conversion above. Finally, reverse osmosis recovers LA at target purity, closing the material loop [25].

Thermochemical Pathway

For the thermochemical benchmark, existing work and inventories were used, as shown in the SI. The process is based on high-TRL biomass gasification [7] followed by a MTO [8, 10] route to produce bio-PE. Circularity is achieved via polymer gasification, returning the waste plastic to syngas [11]. This route is characterized by higher energy intensity and a higher number of steps but benefits from seamless integration into existing infrastructure.

METHODOLOGY

Lifetime-adjusted LCA

The LCA follows ISO 14040/14044 standards [26] using Ecoinvent 3.10 [27] as the background system and process simulation outputs as foreground data. Calculations are performed in Brightway2 [28], utilizing the IPCC 2021 climate change method for GWP over a time horizon of 100 years and including hydrogen and biogenic carbon.

Beyond the standard cut-off approach, we implement a lifetime-adjusted impact framework following Nicholson et al. [12]. This approach calculates the number of material lifetimes L based on the process yield of polymer recycling (a), the recycling rate (b), and the fraction of recycled monomer in the new material (c), which make up the total fraction of recyclate Φ, following equation (1). The polymer recycling yield is process-specific and taken from the respective inventories of enzymatic hydrolysis (a = 95%) and polymer gasification (a = 81%). The recycling rate describes how much plastic is collected for recycling. The fraction of recycled monomer in the new material was assumed to be one for both PLA and bio-PE, as the chemical recycling process produces polymers of "virgin-like" quality, such that no other virgin material is required.

$$\frac{ab}{c} = \Phi \tag{1}$$

It is assumed that anything not recycled ($1-\Phi$) ends up in waste streams, that ultimately are treated by municipal waste incineration. The number of lifetimes L is used to scale the products environmental impact x_L, ensuring each product shares the burden of virgin production x_v and recycling x_r across all cycles n in the

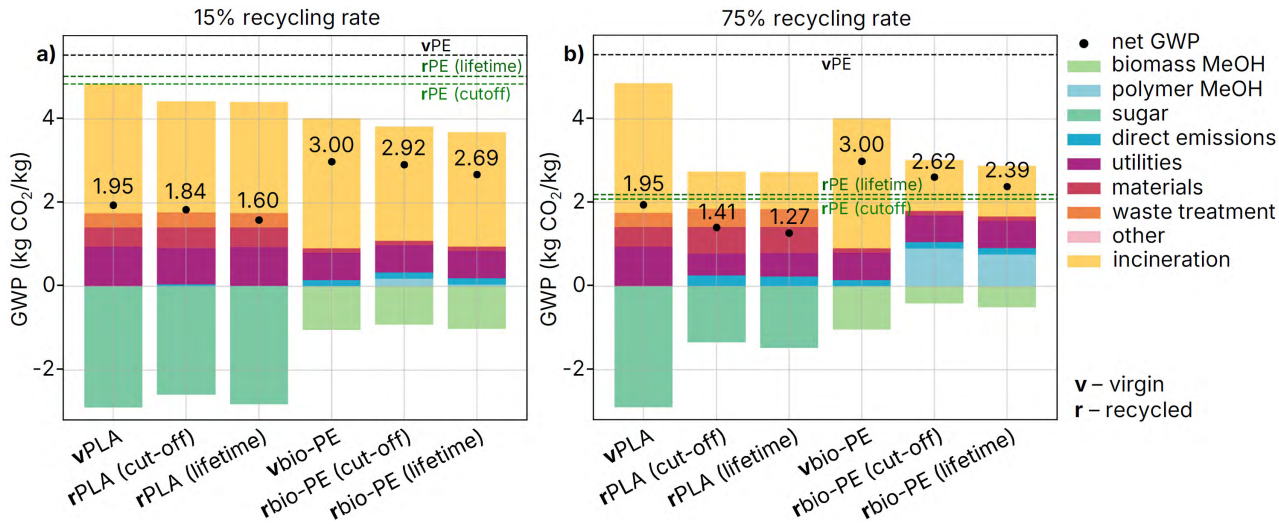

Figure 1. a) Breakdown of average GWP impacts for virgin and recycled PLA, bio-PE, and fossil-PE across different life-cycle calculation models (cut-off vs. lifetime-adjusted) at 15% recycling rate. **b)** Breakdown of average GWP impacts for virgin and recycled PLA, bio-PE, and fossil-PE across different life-cycle calculation models (cut-off vs. lifetime-adjusted) at 75% recycling rate.

production and recycling cascade as can be seen in equations (2) and (3).

$$L = \sum_{n=0}^{\infty}\left(\frac{ab}{c}\right)^n = \frac{c}{c-ab}\text{, for } \frac{ab}{c} < 1 \quad (2)$$

$$x_L = \frac{x_v + (L-1)x_r}{L} \quad (3)$$

One limitation of the lifetime calculation, however, is that the amount of monomer recovered from chemical recycling (ab) must be less than the amount used to resynthesize the equivalent amount of polymer product (c), as seen in equation (2). Hence parameter c is also fixed at one for mechanically recycled conventional PE.

While many studies compare the cut-off impacts of a conventional material with its recycled counterpart, this approach assumes that all plastic is recycled. With our current recycling rates of just 15%, our study applies an averaged impact of circular plastics accounting for a virgin make-up, as shown in equation (4).

$$x_{average} = \Phi\, x_r + (1-\Phi)x_v \quad (4)$$

The model further accounts for end-of-life treatment of losses, distinguishing between material entering the loop and virgin makeup. Equations (5) and (6) present the contribution of the end-of-life treatment x_{eol} to the overall impacts in both the cut-off and lifetime adjusted models. As Nicholson's model only assumes the same life-time adjusted impact over all levels in the cascade for the material entering the loop, the fraction that is virgin but then lost, needs to be accounted separately. For an extended derivation, please consult the SI.

$$x_{cut-off} = \Phi\, x_r + (1-\Phi)x_v + (1-\Phi)x_{eol} \quad (5)$$

$$x_{lifetime} = \Phi x_L + (1-\Phi)\Phi x_L + (1-\Phi)^2 x_v + (1-\Phi)x_{eol} \quad (6)$$

By comparing both models and benchmarking against linear impacts of fossil high-density polyethylene (HDPE) and mechanically recycled HDPE across varying recycling rates, we highlight how different calculation methods can steer policy decisions. The inventories for HDPE from naphtha cracking, henceforth called conventional PE, and from mechanical recycling (recycled conventional PE) were taken from Ecoinvent 3.10. While the use phase is excluded due to lack of data, the comparative results across plastic types remain robust.

RESULTS & DISCUSSION

Impacts at current state of recycling

Figure 1 a) shows the breakdown of the averaged GWP impacts for the biochemical pathway, represented by virgin and recycled PLA, as well as the impacts from the thermochemical production path for virgin and recycled bio-PE at current recycling rates of 15%. In the case of both virgin polymers, waste incineration has the biggest contribution to the total GWP. However, PLA is able to offset most of this impact with the production of the sugars, which uses biomass, resulting in a negative contribution. The negative contribution stems from the fact that biomass uses CO_2 from the atmosphere during its growth, such that this take-up is subtracted from the total process and end-of-life emissions and can be viewed as a carbon credit. On the other hand, bio-PE from gasification also uses biomass and even to higher extents, but the production of methanol involves many more steps and much more energy, resulting in a smaller offset.

Nonetheless, both bioplastics are outperforming conventional virgin and even recycled PE at these lower recycling rates in terms of lower environmental burden.

At the lower recycling rate, the contribution of virgin material and the losses to be incinerated are quite high, explaining the small relative improvement of the bioplastics when they are recycled with respect to the virgin material. All the same, recycling bioplastics presents an environmental benefit, as the amount of material getting incinerated decreases in both cases, bio- and thermochemical, of recycling.

Comparing the two impact calculation methods, one observes that for both bioplastics applying the lifetime-adjusted framework results in 13% and 8% lower GWP for recycled PLA and bio-PE respectively. This stems from accounting for the first cycle of the virgin material that has negative carbon credits and hence decreases the impact of the plastics, even when recycled. Similarly, the averaged impacts of conventional PE in the recycling scenario differ depending on the calculation method. GWP of conventional PE when recycled is underestimated, as the contribution of the virgin material is disregarded.

Impacts at higher recycling rates

Following the lifecycle analysis of bioplastics at low recycling rates, we investigate how these impacts change with an increased rate of collection. Figure 1 b) shows the impact breakdown for both bioplastics, PLA and bio-PE, in the linear and circular scenarios, when the recycling rate increases to 75%.

As expected, when more plastic is recycled the GWP decreases for all types of plastic in the circular scenario. Now the contribution of the virgin material decreases and hence the carbon credits decrease, as less biomass is used. At the same time, much less plastic is lost and incinerated, such that the burden from the end-of-life treatment decreases as well, resulting in lower impacts. These end-of-life impacts differ between PLA and bio-PE as their monomer yields, and hence the fraction of recyclate is different. In this scenario, recycling plastics becomes much more beneficial, however it becomes evident, that while recycling PLA via enzymatic hydrolysis still outperforms the mechanical recycling of conventional PE, the gasification of bio-PE does not yield any environmental benefits anymore. This is due to the lower monomer yield and higher energy consumption from the gasification process, so as recycling rates increase, the more precise and less energy intensive biochemical pathway offers better performance than the thermochemical path.

One also observes that the difference in impacts for the recycled material between the two recycling rates is even more significant for conventional PE than for the bioplastics, because of the big difference between virgin and recycled impacts.

Further, the differences in GWP of the recycled materials from the two calculation methods are not as pronounced anymore. This is also explained by the decreasing impact of the virgin material, suggesting that at higher recycling rates the choice of calculation method becomes less relevant. At 75% recycling rate the difference is still significant, however moving towards a circular future this might not be the case anymore, warranting a sensitivity analysis.

Sensitivity to recycling rates

Given the decreasing influence of the impact calculation model, a sensitivity analysis with respect to the recycling rate was performed for PLA and conventional PE using both models, as shown in Figure 2. The GWP of the virgin plastics that are being incinerated remains unchanged across different recycling rates as they never enter the material loop and are hence independent of the rate of recycling.

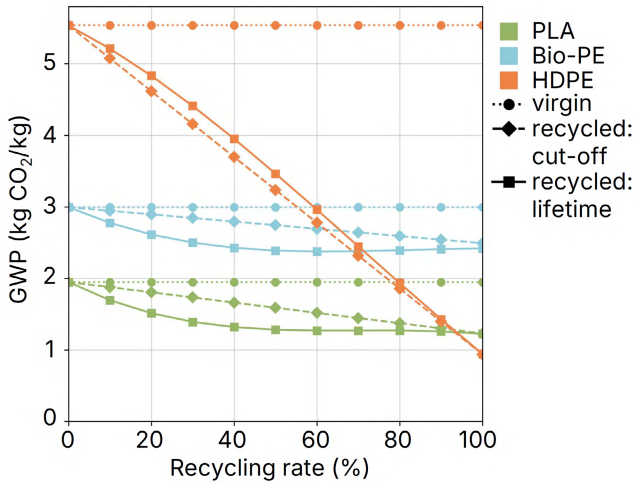

Figure 2. Sensitivity analysis of GWP impacts of virgin conventional PE (orange), virgin bio-PE (blue) and virgin PLA (green), as well as impacts when recycled, calculated with the cut-off (diamond) and the lifetime-adjusted (square) methods, across all possible recycling rates.

Modeling choice

Focusing on the scenarios where the plastics are recycled, we notice that naturally, at 0% recycling rate all impacts for the recycled plastic are equal to the impact of the respective virgin plastic regardless of the model, as no recycling means that all material is ultimately incinerated, corresponding to a linear economic scenario. As the recycling rates increase, the impacts of the more circular scenarios decrease and a difference between the impact calculation models becomes noticeable. This divergence between the two models increases, as the cut-off model follows a linear trend as per equation (5), while

the lifetime-adjusted model traces a non-linear path, due to the consideration of the virgin material in the lifetime-adjusted impacts as per equations (3) and (6). The disparity starts decreasing after a recycling rate of about 40%, as with a higher fraction of the recyclate the influence of the virgin material and the end-of-life treatment decreases and the source of the carbon becomes less relevant. At 100% collection rate there is virtually no virgin material and no losses, except to account for the deficit during the depolymerization process. This means that the source of carbon is now irrelevant and the limiting factor becomes the process efficiency and not the feedstock, such that the impacts of both models converge at the cradle-to-gate impact of recycling the material x_r. These observations suggest that for a certain range of recycling rates, the usually applied cut-off method may overestimate impacts for bioplastics and underestimate them for fossil plastics within a circular economy.

Technology choice

Comparing the biopolymers and conventional PE, the feedstock plays a major role at low recycling rates, so that bioplastic whether virgin or recycled, has much lower GWP than conventional PE derived from cracking naphtha, due to the carbon credits from biomass consumption. At higher recycling rates, however, the difference quickly diminishes. As the difference between virgin and purely recycled impacts of conventional PE is so high, the impact decreases more steeply than for PLA or bio-PE. The trends between PLA and bio-PE are very similar with the difference that bio-PEs impacts are higher across all recycling rates. Nonetheless, recycling rates of nearly 80% are needed for mechanically recycled conventional PE to beat virgin PLA in terms of GWP, and nearly 95% of plastic need to be collected for recycled conventional PE to have lower environmental impacts than the biochemically recycled PLA, as the impact of the virgin material can be neglected and the process itself becomes the main constraint. This suggests that ultimately, in a perfectly circular economy, the source of carbon would be virtually irrelevant and bioplastic would offer no environmental advantage over fossil plastic when both are fully recycled, a very ideal unlikely case. However, with current recycling rates of 15% we are far away from this goal, so as long as there is plastic loss the carbon credits from biomass derived plastics offer a necessary buffer to mitigate long-term environmental damage.

While this work has focused on environmental impacts in terms of global warming potential, there are other indicators that are relevant in sustainability assessments. Previous studies show that bioplastics such as PLA can have higher impacts in other categories like human toxicity and eutrophication, due to the agricultural phase and the use of fertilizers [29, 30]. This in turn highlights the importance of plastic recycling to mitigate the impacts from virgin resource extraction. Though there are many studies, their differing methodologies result in a big variation of insights, warranting further analysis under a harmonized framework [30].

CONCLUSION

A lifecycle assessment comparing the biochemical and thermochemical pathways for biomass valorization to PLA and bio-PE respectively was performed. Two impact calculation models (cut-off and lifetime adjusted) were used to evaluate the average GWP impact when considering the virgin make-up in a circular economy. The analysis showcased the environmental advantage of both processing pathways at current recycling rates of 15%, with PLA having the lowest impact, due to its highly-targeted biomass conversion and energy economy. At higher recycling rates (75%), impacts for the circular materials decrease, such that biomass and polymer gasification to produce PE cannot compete anymore, owing to the lower depolymerization yield and higher energy consumption during methanol production. Meanwhile PLA recycled via enzymatic hydrolysis remains superior to mechanically recycled conventional PE in terms of GWP.

The lifetime-adjusted framework offers a more accurate estimation of life cycle impacts in a circular economy, as it accounts for the first lifecycle of the virgin material and therefore attributes the corresponding carbon credit or penalty to the averaged impact. However, towards maximum recycling rates the models converge, as the impact of the virgin material becomes negligible and the process impact becomes the main limitation, such that mechanically recycled PE shows a lower GWP than recycled PLA.

At high recycling rates the source of carbon becomes irrelevant, making bioplastics less appealing for decarbonizing the plastic sector. However, we are very far from reaching the goal of these high recycling rates and continue to leak plastic from the closed material loop. While this remains the case, bioplastics offer the necessary carbon offset and facilitate the transition to our circular future, by mitigating damage caused by excess CO_2 emissions already now.

Nonetheless this study has shown that in the long-term depending solely on bioplastics is not enough, as aside from the heavy burden on land use, recycling can offer better environmental benefits. This means the research community's focus should move towards the development of materials that are recyclable-by-design and processes that facilitate that beyond identifying the most suitable renewable feedstock.

DIGITAL SUPPLEMENTARY MATERIAL

Digital supplementary material is provided at

http://PSEcommunity.org/LAPSE under the LAPSE ID LAPSE:2026.0007, including more detailed process descriptions, flowsheets and inventory data, as well as the full explanation and derivation of the lifetime-adjusted framework.

ACKNOWLEDGEMENTS

This publication was created as part of NCCR Catalysis (grant number 225147), a National Centre of Competence in Research funded by the Swiss National Science Foundation.

AUTHOR IDENTIFIERS

Author ORCIDs:
Krestnikova A: 0009-0007-7681-1976
Guillén-Gosálbez G: 0000-0001-6074-8473

REFERENCES

1. Grace R. Closing the circle: reshaping how products are conceived & made. Plastics Engineering 73:8-11 (2017). https://doi.org/10.1002/j.1941-9635.2017.tb01670.x
2. Statista. Distribution of plastic consumption worldwide in 2019, by application. https://www.statista.com/statistics/1002055/plastic-consumption-share-worldwide-by-application/
3. Plastics Europe. Plastics the Fast Facts 2025. https://plasticseurope.org/knowledge-hub/plastics-the-fast-facts-2025/
4. Bachmann M, Zibunas C, Hartmann J, Tulus V, Suh S, Guillén-Gosálbez G, Bardow A. Towards circular plastics within planetary boundaries. Nat Sustain 6:599-610 (2023). https://doi.org/10.1038/s41893-022-01054-9
5. Statista. Distribution of the production capacities of bioplastics worldwide in 2024 with a forecast to 2029, by material type. https://www.statista.com/statistics/678775/production-capacity-distribution-of-bioplastics-worldwide-by-material/
6. Vink ETH, Davies S. Life cycle inventory and impact assessment data for 2014 ingeo ™ polylactide production. Industrial Biotechnology 11:167-180 (2015). https://doi.org/10.1089/ind.2015.0003
7. Liu Y, Li G, Chen Z, Shen Y, Zhang H, Wang S, Qi J, Zhu Z, Wang Y, Gao J. Comprehensive analysis of environmental impacts and energy consumption of biomass-to-methanol and coal-to-methanol via life cycle assessment. Energy 204:117961 (2020). https://doi.org/10.1016/j.energy.2020.117961
8. Ioannou I, D'Angelo SC, Martín AJ, Pérez?Ramírez J, Guillén?Gosálbez G. Hybridization of fossil? and co₂?based routes for ethylene production using renewable energy. ChemSusChem 13:6370-6380 (2020). https://doi.org/10.1002/cssc.202001312
9. Nosrati-Ghods N, Petersen AM, Bosman CE, Myburgh MW, Favaro L, Viljoen-Bloom M, Görgens JF. Techno-economic assessment of waste poly(lactic acid) enzymatic hydrolysis and recycling for poly(lactic acid) production. Biochemical Engineering Journal 215:109615 (2025). https://doi.org/10.1016/j.bej.2024.109615
10. Ioannou I, Javaloyes-Antón J, Caballero JA, Guillén-Gosálbez G. Economic and environmental performance of an integrated co₂ refinery. ACS Sustainable Chem. Eng. 11:1949-1961 (2023). https://doi.org/10.1021/acssuschemeng.2c06724
11. Salah C, Istrate R, Bjørn A, Tulus V, Pérez-Ramírez J, Guillén-Gosálbez G. Environmental benefits of circular ethylene production from polymer waste. ACS Sustainable Chem. Eng. 12:13897-13906 (2024). https://doi.org/10.1021/acssuschemeng.4c04241
12. Nicholson SR, Rorrer JE, Singh A, Konev MO, Rorrer NA, Carpenter AC, Jacobsen AJ, Román-Leshkov Y, Beckham GT. The critical role of process analysis in chemical recycling and upcycling of waste plastics. Annu. Rev. Chem. Biomol. Eng. 13:301-324 (2022). https://doi.org/10.1146/annurev-chembioeng-100521-085846
13. Wooley RJ, Putsche V. Development of an ASPEN PLUS Physical Property Database for Biofuels Components. Tech. rep. National Renewable Energy Laboratory. (1996)
14. Liu ZH, Chen HZ. Xylose production from corn stover biomass by steam explosion combined with enzymatic digestibility. Bioresource Technology 193:345-356 (2015). https://doi.org/10.1016/j.biortech.2015.06.114
15. Martino DC, Colodette JL, Chandra R, Saddler J. Steam explosion pretreatment used to remove hemicellulose to enhance the production of a eucalyptus organosolv dissolving pulp. Wood Sci Technol 51:557-569 (2017). https://doi.org/10.1007/s00226-016-0889-y
16. Sulzenbacher D, Atzmüller D, Hawe F, Richter M, Cristobal-Sarramian A, Zwirzitz A, Optimization of steam explosion parameters for improved biotechnological use of wheat straw. *Biomass Convers. Biorefin.*, 13, 1035–1046. (2023)
17. Zhao J, Xu L, Wang Y, Zhao X, Wang J, Garza E, Manow R, Zhou S. Homofermentative production of optically pure l-lactic acid from xylose by

genetically engineered escherichia coli b. Microb Cell Fact 12: (2013). https://doi.org/10.1186/1475-2859-12-57
18. Gezae Daful A, Görgens JF. Techno-economic analysis and environmental impact assessment of lignocellulosic lactic acid production. Chemical Engineering Science 162:53-65 (2017). https://doi.org/10.1016/j.ces.2016.12.054
19. Bressanin JM, Sampaio ILM, Geraldo VC, Klein BC, Chagas MF, Bonomi A, Filho RM, Cavalett O. Techno-economic and environmental assessment of polylactic acid production integrated with the sugarcane value chain. Sustainable Production and Consumption 34:244-256 (2022). https://doi.org/10.1016/j.spc.2022.09.009
20. Su CY, Yu CC, Chien IL, Ward JD. Plant-wide economic comparison of lactic acid recovery processes by reactive distillation with different alcohols. Ind. Eng. Chem. Res. 52:11070-11083 (2013). https://doi.org/10.1021/ie303192x
21. Gruber P, O'Brien M. Polylactides "natureworks® PLA". Biopolymers Online : (2005). https://doi.org/10.1002/3527600035.bpol4008
22. Gruber P. CONTINUOUS PROCESS FOR THE MANUFACTURE OF LACTIDE AND LACTIDE POLYMERS. US Patent 6, 326, 458 B1. (2001)
23. Li F, Zhan Y, Liu D, Liu D, Xia H, Chen S. Production of poly(l-lactide)-degrading enzyme by amycolatopsis orientalis ssp. orientalis and its catalytic ability in biological recycling of poly(l-lactide). J Polym Environ 21:1143-1149 (2013). https://doi.org/10.1007/s10924-013-0590-2
24. Cao X, Lee HJ, Yun HS, Koo YM. Solubilities of calcium and zinc lactate in water and water-ethanol mixture. Korean J. Chem. Eng. 18:133-135 (2001). https://doi.org/10.1007/bf02707210
25. Li Y, Shahbazi A, Williams K, Wan C. Separate and concentrate lactic acid using combination of nanofiltration and reverse osmosis membranes. Appl Biochem Biotechnol 147:1-9 (2007). https://doi.org/10.1007/s12010-007-8047-5
26. International Standards Organization (ISO) ISO 14044: Environmental Management—Life Cycle Assessment—Principles and Framework, 2006. https://www.iso.org/standard/37456.html
27. Wernet G, Bauer C, Steubing B, Reinhard J, Moreno-Ruiz E, Weidema B. The ecoinvent database version 3 (part I): overview and methodology. Int J Life Cycle Assess 21:1218-1230 (2016). https://doi.org/10.1007/s11367-016-1087-8
28. Mutel C. Brightway: an open source framework for life cycle assessment. JOSS 2:236 (2017). https://doi.org/10.21105/joss.00236
29. Fonseca A, Ramalho E, Gouveia A, Figueiredo F, Nunes J. Life cycle assessment of PLA products: a systematic literature review. Sustainability 15:12470 (2023). https://doi.org/10.3390/su151612470
30. Walker S, Rothman R. Life cycle assessment of bio-based and fossil-based plastic: a review. Journal of Cleaner Production 261:121158 (2020). https://doi.org/10.1016/j.jclepro.2020.121158

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimization of Circular Supply Chains for Electric Vehicle Batteries

Kaapo Kopra[a] and Iiro Harjunkoski[a*]

[a] Aalto University, Department of Chemical and Metallurgical Engineering, Espoo, Finland
* Corresponding Author: iiro.harjunkoski@aalto.fi

ABSTRACT

The increasing popularity of electric vehicles (EVs) leads to an expected rise in the quantity of end-of-life lithium-ion batteries (LIBs) that require efficient management. This paper presents a State Task Network (STN) based optimization model to analyze and optimize the supply chain for LIBs, allowing for the selection of optimal processing routes, facility locations, capacities and re-integration of recovered materials, as well as analyzing the possible trade-offs between different end-of-life management strategies. Based on available data from the literature, the model is demonstrated with the LIB supply chain considering both primary production and different end-of-life strategies for spent LIBs (recycling and reuse). The case study reveals that mechanical pretreatment followed by hydrometallurgical recycling is the optimal pathway and it outperforms the linear supply chain in both costs and emissions. The cost optimal solution opts for more centralized collection and disassembly, whereas when minimizing the total greenhouse gas emissions, the optimal collection and disassembly center locations are decentralized. The model can aid in the design of circular supply chains to select and locate the optimal processes and to examine how the network should evolve over time.

Keywords: Batteries, GAMS, Optimization, Supply Chain, Circular Economy

INTRODUCTION

The recent energy transition to renewable energy sources has resulted in increased demand for metals such as lithium, cobalt, nickel and manganese that are needed to produce lithium-ion batteries (LIBs) used mainly in electric vehicles (EVs). At the same time, the natural mineral deposits are reaching critical concentrations of these valuable metals and thus increasing the environmental impacts associated with metals extraction via the conventional primary production routes [1]. As only a limited fraction of LIBs are currently recycled [2], a wider adoption of recycling processes and advancements in circular economy (CE) could help alleviate future supply-demand disparities for battery materials, reduce environmental burdens, and lower the production costs by decreasing the required amount of primary production of battery metals.

The concept of circular economy (CE) has been gaining popularity in businesses, government and the research community [3]. In contrast with linear production systems, the aim of circular economy is that materials never become waste. It is based on the following principles: to eliminate waste and pollution, to circulate products and materials and to regenerate nature [4]. Aligned with the principles of CE, end-of-life (EOL) EV LIBs can be managed through several alternatives to disposal: reuse in second-life applications, remanufacturing suitable battery cells, and recycling to recover valuable materials from EOL LIBs. The general structure of the circular LIB system is presented in Figure 1.

A vast amount of research has focused on designing recycling processes for LIBs to recover valuable materials and facilitate closed-loop recycling of battery materials. Currently, mechanical pretreatment processes, hydrometallurgical processes and pyrometallurgical processes are conducted on an industrial scale, whereas emerging direct recycling processes are currently on a pilot or laboratory scale level and thus lack industrial maturity [5].

- Mechanical pre-treatment: involves discharging, crushing and separation stages for the LIBs to

obtain black mass powder which contains valuable metals, mainly from the cathode active material.

- Hydrometallurgical recycling: the black mass is leached, and the battery metals are selectively separated and purified.
- Pyrometallurgical recycling: the spent LIBs are fed into a high-temperature process to produce a metal alloy as well as a slag phase as a byproduct.

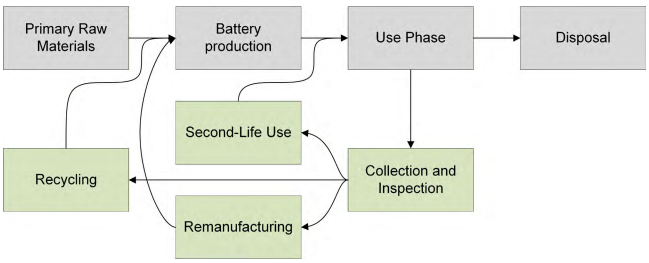

Figure 1. Schematic representation of the circular LIB system.

METHODOLOGY

The need for systematic methods to assess and optimize circular economy networks has been recognized in the process systems engineering literature [3]. Recent examples include the design and optimization of circular economy networks for polyethylene terephthalate (PET) [6] and the optimization of circular food supply chains [7].

Ahmed et al. [6] proposed a State Task Network (STN) modeling approach for the optimal design of circular economy networks and applied it to a case study on polyethylene terephthalate (PET) networks with a one-year assessment period. Baratsas et al. [7] presented a systems engineering framework for the optimization of circular food supply chains using a Resource Task Network (RTN) representation. Finally, Iakovou et al. [8] presented an RTN-based approach for the optimal design of reverse logistics networks for photovoltaic recycling.

The multi-period modeling of the supply chain for batteries is crucial, especially given the growing availability of EOL materials. By modeling the entire supply chain, we can examine how different end-of-life management strategies can be efficiently integrated with the primary production processes.

Model formulation

The optimization model is based on the State Task Network representation first introduced by [9] for short-term scheduling of batch processes. The State Task Network representation allows for flexible representation of production processes where materials are represented as states that are transformed in processes referred to as tasks. Thus, the framework allows for the preparation of a superstructure, which includes all the possible production processes that can be used to produce the required products in the supply chain. As the STN formulation inherently allows for material circulation in the network, a complete supply chain network can be represented which includes both primary production processes and all possible processes to treat end-of-life batteries, as well as the possible reintegration routes in the supply chain.

The goal is to be able to present the holistic supply chain model for lithium-ion batteries that allows the investigation of optimal EOL LIB treatment options and how the recover materials can be integrated back into the different stages of the battery production supply chain.

Sets, Parameters and Variables

The nomenclature used in the optimization model is summarized in Tables 1-3.

Table 1: Sets used in the optimization model.

$c \in C$	Capacity selection
$d \in D$	Delay time periods
$i \in I$	Processing tasks
$l \in L$	Locations
$s \in S$	Materials
$t \in T$	Time periods
$S^F \subset S$	Subset for final products
$S^V \subset S$	Subset for primary raw materials

Table 2: Parameters used in the optimization model.

$Cap_{i,c}^{Max}$	Maximum capacity for processing task
$Cap_{i,c}^{Min}$	Minimum capacity for processing task
$CostP_{s,t}$	Cost of purchased materials
$CostS_{s,t}$	Cost of sold/sink materials
$CTrans$	Transportation cost
$Dem_{l,s,t}$	Material demand
$Dist_{l,l'}$	Distance between locations
E_i	Emissions of processing tasks
$EPur_s$	Material procurement emissions
$ETrans$	Transportation emissions
$FC_{i,c,t}$	Fixed costs of processing tasks
$Sink_{s,t}^{Max}$	Maximum sink for materials
$Supply_{s,t}^{Max}$	Maximum supply of materials
$VC_{i,t}$	Variable costs of processing tasks
$\rho_{s,i}$	Conversion of materials in processing tasks
$\rho_{s,i}^{+}$	Generation of material s in i, $(\rho_{s,i} > 0)$
$\rho_{s,i}^{-}$	Consumption of material s in i, $(\rho_{s,i} < 0)$
$\alpha_{i,d}$	Fraction of outputs from process i released after delay d

Table 3: Variables used in the optimization model.

$B_{i,l,t}$	Flow through processing task i at location l during time period t.
$P_{l,s,t}$	Purchasing of material s in location l during time period t.
$S_{l,s,t}$	Selling/sink of material s in location l during time period t.
$Tr_{l,l',s,t}$	Transportation of material s from location l to location l' during time period t.
$Y_{i,l,t,c}$	1 if processing task i with capacity level c is selected at location l during time period t, 0 otherwise.

Objective Functions

To evaluate the potential trade-offs between economic performance, environmental impacts, and supply chain circularity, three objective functions are considered: supply chain costs, greenhouse gas emissions generated and the consumption of primary raw materials.

First, the total supply chain costs are minimized that consists of fixed and variable processing costs, transportation costs, and material purchasing costs and revenues generated from selling materials.

$$\min \sum_i \sum_l \sum_t B_{i,l,t} \cdot VC_{i,t} + \sum_i \sum_l \sum_t \sum_c Y_{i,l,t,c} \cdot FC_{i,c,t} + \sum_l \sum_{l'} \sum_s \sum_t CTrans \cdot Tr_{l,l',s,t} \cdot \mathrm{Dist}_{l,l'} + \sum_l \sum_s \sum_t (P_{l,s,t} \cdot CostP_{s,t} - \mathrm{S}_{l,s,t} \cdot CostS_{s,t}) \quad (1)$$

Second, the environmental sustainability of the supply chain is addressed by minimizing the total emissions that consists of process emissions, transportation emissions and emissions related to material procurement.

$$\min \sum_i \sum_l \sum_t B_{i,l,t} \cdot E_i + \sum_l \sum_{l'} \sum_s \sum_t ETrans \cdot Tr_{l,l',s,t} \cdot Dist_{l,l'} + \sum_l \sum_s \sum_t P_{l,s,t} \cdot EPur_s \quad (2)$$

Finally, as proposed in [6], the total purchasing of primary raw materials is minimized and is considered as a proxy for material circularity in the supply chain.

$$\min \sum_l \sum_{s^v} \sum_t P_{l,s,t} \quad (3)$$

Constraints

Net flow through each processing task cannot exceed the maximum production capacity at the selected capacity level.

$$B_{i,l,t} \leq \sum_C Y_{i,l,t,c} \cdot Cap_{i,c}^{Max} \ \forall i, l, t \quad (4)$$

Net flow through each processing tasks must exceed the minimum production capacity level.

$$B_{i,l,t} \geq \sum_C Y_{i,l,t,c} \cdot Cap_{i,c}^{Min} \ \forall i, l, t \quad (5)$$

Material purchases in the supply chain cannot exceed the maximum supply of materials.

$$\sum_l P_{l,s,t} \leq Supply_{s,t}^{Max} \ \forall s, t \quad (6)$$

Material sales in the supply chain cannot exceed the maximum allowed sink of materials.

$$\sum_l S_{l,s,t} \leq Sink_{s,t}^{Max} \ \forall s, t \quad (7)$$

The demand for the final products must be satisfied in the supply chain. In a circular system, the final products are consumed by the demand zones.

$$\sum_i -B_{i,l,t} \cdot \rho_{s,i}^{-} = Dem_{l,s,t} \ \forall l, s^F, t \quad (8)$$

The mass balance at all locations and all time periods is formulated based on material generation or consumption, transportation and the purchasing or selling of materials. As in the case of electric vehicle batteries, the use phase introduces a time delay before the material becomes available for end-of-life processes. Accordingly, the material generation term is modified to model this delay as a fractional release, where each task with a delay generates materials in predetermined fractions.

$$\sum_i \sum_d B_{i,l,t-d} \cdot \rho_{s,i}^{+} \cdot \alpha_{i,d} + P_{l,s,t} + \sum_{l'} Tr_{l',l,s,t} = \sum_i -B_{i,l,t} \cdot \rho_{s,i}^{-} + S_{l,s,t} + \sum_{l'} Tr_{l,l',s,t} \ \forall s, l, t \quad (9)$$

For each processing task and each time period, at most one capacity level can be selected.

$$\sum_c Y_{i,l,t,c} \leq 1 \ \forall i, l, t \quad (10)$$

The selected production capacity of the processing tasks at time t cannot be lower than at previous time period t-1.

$$\sum_c Y_{i,l,t,c} \cdot Cap_{i,c}^{Max} \geq \sum_c Y_{i,l,t-1,c} \cdot Cap_{i,c}^{Max} \ \forall i, l, t \quad (11)$$

Thus, the objective functions defined in Eqs. (1-3) are used to study trade-offs in circular supply chain design, and the resulting MILP formulation can be expressed as:

$$\min Z$$
$$s.t. Eq. (4-11) \quad (12)$$
$$B_{i,l,t}, P_{l,s,t}, S_{l,s,t}, Tr_{l,l',s,t} \geq 0, Y_{i,l,t,c} \in \{0,1\}$$

By constraining the material flows through selected echelons and by fixing the location of existing facility locations in the supply chain, the model facilitates "what if" analyses and enables the investigation of different scenarios.

CASE STUDY

To illustrate the proposed optimization model for circular supply chain design, a representative case study of NMC622 LIBs used in EVs is considered. The case study investigates the high-value metals present in the cathode active material of the LIBs: refined lithium, cobalt, nickel and manganese, while other materials in the lithium-ion battery cells and packs are considered to be inert. Model parameters are chosen to reflect typical

economic, environmental and process conditions that are reported in the literature ([10-12]). The case study considers 10 uniformly randomly generated locations within a 1000 x 1000 km grid, each associated with identical demand for LIBs with an annual demand growth rate of 1.5%. The supply chain is optimized for the time period of 30 years, assuming that no end-of-life LIBs are available at the initial time period. The main assumptions of the case study are summarized in Table 4, and Figure 2 presents the candidate locations and locations fixed in the linear reference scenario.

Table 4: The main assumptions used in the case study.

Parameter	Value
Cell chemistry	NMC622
Cell specific energy	0.218 kWh/kg
CAM to Cell ratio	0.324
Cell to pack ratio	0.6
Lifespan of EV batteries	10 years
Lifespan of reused batteries	4 years

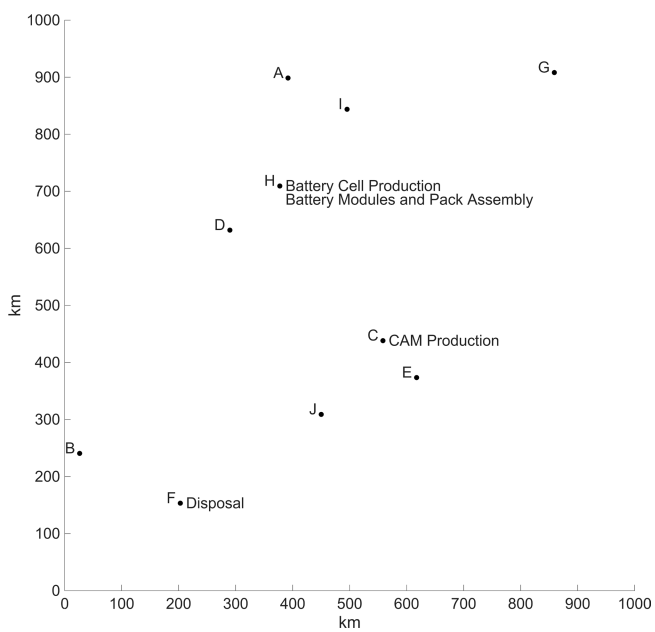

Figure 2. Representation of the case study showing fixed battery production sites (C, F, H) and candidate locations.

In the case study, the production of LIB packs begins with the production of cathode active material (CAM) which consumes refined lithium, cobalt, nickel and manganese to produce the CAM. The CAM is then supplied to the cell production phase, which manufactures battery cells that are then assembled into battery modules and finally into battery packs in the assembly process. From the assembly, the produced battery packs are transported to demand zones where the battery packs are used in electric vehicles. For the case study, a fixed lifetime of 10 years is assumed before the batteries become available for end-of-life treatment processes. Figure 3 illustrates the superstructure of the studied network.

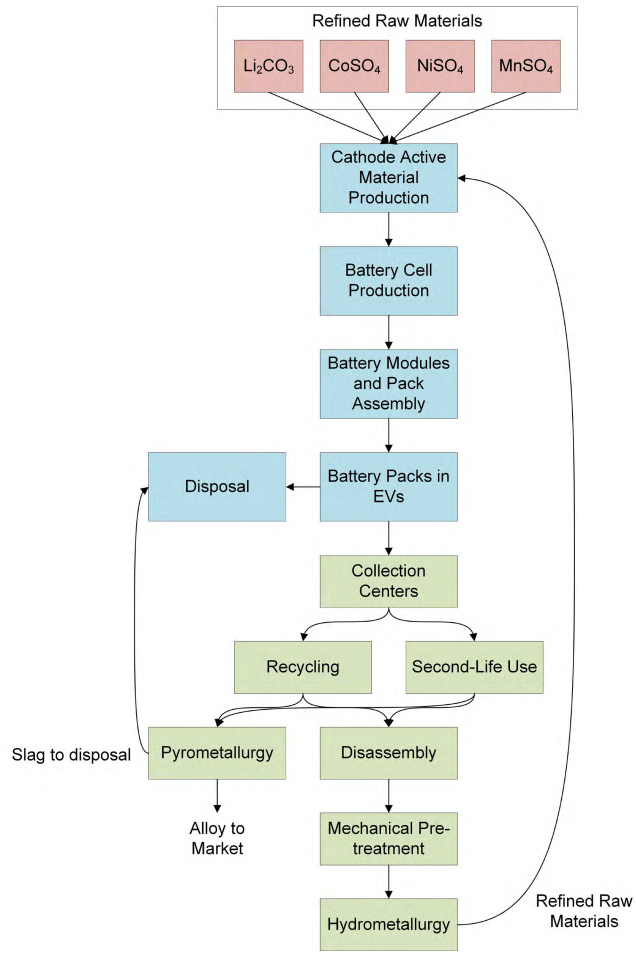

Figure 3. Superstructure representation of the circular LIB network, focusing on the high-value refined battery metals present in cathode active material (CAM).

In the linear reference scenario, all EOL LIB packs are transported to the disposal facility located at F. In the recycling scenario, EOL LIB packs can be collected at collection and sorting centers, and either processed directly in pyrometallurgical recycling facilities or first disassembled into battery cells, followed by mechanical pretreatment to obtain black mass. The black mass can be subsequently treated in hydrometallurgical recycling facilities to recover refined battery materials that can be used in the CAM production. Finally, in the circular scenario, the EOL battery packs can be processed using the same recycling pathways, or a maximum of 20% can be sent to location F to be reused as stationary energy storage. Here we assume a fixed lifetime of 4 years as energy storage after which the LIBs are treated in downstream recycling processes. The analyzed scenarios in the case

study are summarized in Table 5.

Table 5: Studied scenarios in the case study.

Scenario	Description
Linear	Linear benchmark supply chain, where EOL LIBs are transported to disposal.
Recycling	EOL LIBs can be processed in recycling processes.
Circular	EOL LIBs can be processed in recycling processes and a fraction reused in stationary energy storage applications.

RESULTS AND DISCUSSION

The optimization model is implemented in GAMS and solved using CPLEX solver version 22.1.1. The evaluated scenarios (Table 5) were run with all three objective functions (Equations 1-3) to assess the trade-offs in the optimal supply chain design. Further multi-objective optimization and pareto analysis was not conducted at this stage.

Cost optimal solution

Compared with the linear supply chain, the recycling scenario reduced the total cost by 4.4% in the time periods studied. The optimal strategy was to use hydrometallurgical recycling for spent LIBs as it allows for cost savings due to smaller quantity of required raw material purchases. The cost-optimal supply chain uses centralized collection and disassembly and, at the final time period, requires three collection centers, three disassembly centers and a mechanical pretreatment center and a hydrometallurgical facility, as shown in Figure 4. Mechanical pretreatment and hydrometallurgical recycling are located at the same location as the CAM production to avoid transportation of the refined battery metals.

In the circular scenario, when the possibility of second-life use of LIBs was allowed, the optimal solution included using a fraction of the LIBs as energy storage in the allowed location F to generate more revenue, resulting in total supply chain cost decrease of 10.8% compared to the linear supply chain.

Emission optimal solution

When the total supply chain emissions were minimized, the hydrometallurgical processing route was also selected, to reduce the supply chain emissions from raw material procurement. However, in contrast with the cost optimal strategy, the supply chain shifts towards a decentralized collection and disassembly to avoid emissions generated via transportation. The optimal supply chain for the recycling scenario consisted of collection and disassembly facilities at all locations, three mechanical pretreatment facilities at locations C, F, and H and a hydrometallurgical recycling facility at location C, as presented in Figure 5. In the circular scenario, the model did not allocate any collected batteries to secondary use, as no environmental benefit was assigned to reuse.

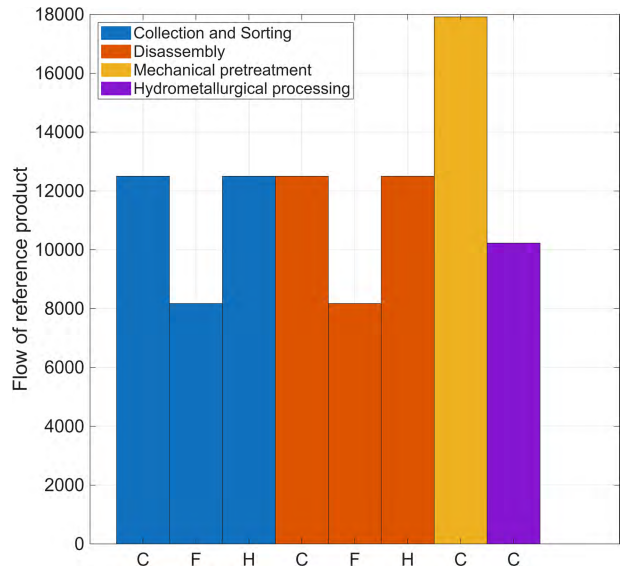

Figure 4. Cost optimal capacities at the final period for the recycling scenario.

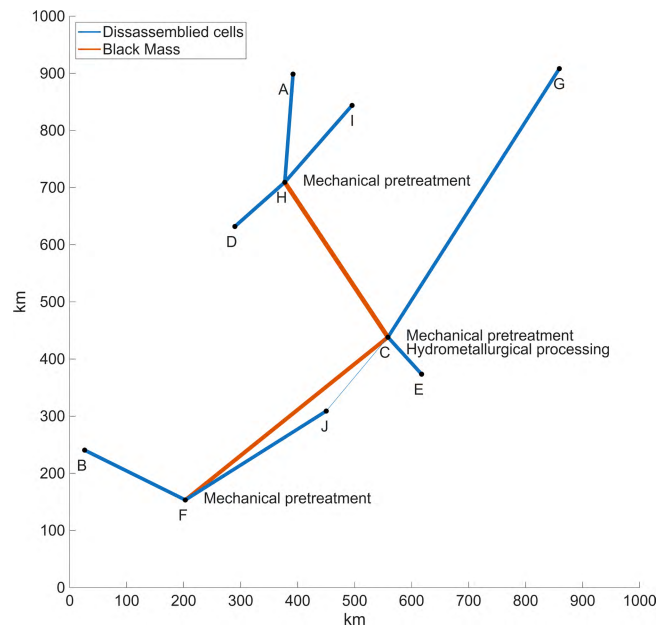

Figure 5. GHG-optimal decentralized supply chain for LIB recycling.

Raw material optimal solution

For the studied generic pyrometallurgical and hydrometallurgical processes, the optimal solution for minimizing primary raw material procurement was to utilize the hydrometallurgical recycling route in all scenarios where recycling was permitted by the model constraints. By recycling EOL LIBs, the supply chain reduced its total refined battery materials consumption by 55% over the

studied time periods.

CONCLUSIONS AND FUTURE WORK

The significant rise in the popularity of electric vehicles results in the increasing demand for lithium-ion batteries, consequently increasing the amount of EOL LIBs, which require efficient management. Thus, an optimization model is proposed as a tool to design and optimize the circular supply chain network for LIBs. The State Task Network based optimization framework allows for the optimal design of facility locations, technologies, capacities and the reintegration of recovered materials back into the different stages of the supply chain.

The results from the illustrative case study tested with literature data showed that both in terms of costs and emissions, the hydrometallurgical process is preferred. With the circular scenario, the optimal strategy was to utilize a fraction of the collected LIBs as stationary energy storage to generate more revenue. The case study revealed that the cost optimal solution prefers centralized collection, while the GHG optimal network prefers decentralized collection and mechanical pretreatment for LIBs. The preliminary results align well with the identified spoke-and-hub network recognized to reduce costs in LIB recycling [13].

Future work could involve developing a more realistic case study as well as studying different cathode active materials used in EVs. Also, considering the full bill of materials enables the optimization and assessment of the circularity of electrolytes and other materials present in LIBs.

AUTHOR IDENTIFIERS

Author ORCIDs:
Kopra: 0009-0005-5428-1818
Harjunkoski: 0000-0003-0428-0751

REFERENCES

1. Velázquez-Martínez O, Valio J, Santasalo-Aarnio A, Reuter M, Serna-Guerrero R. A critical review of lithium-ion battery recycling processes from a circular economy perspective. Batteries 5:68 (2019). https://doi.org/10.3390/batteries5040068
2. Machala ML, Chen X, Bunke SP, Forbes G, Yegizbay A, de Chalendar JA, Azevedo IL, Benson S, Tarpeh WA. Life cycle comparison of industrial-scale lithium-ion battery recycling and mining supply chains. Nat Commun 16: (2025). https://doi.org/10.1038/s41467-025-56063-x
3. Avraamidou S, Baratsas SG, Tian Y, Pistikopoulos EN. Circular economy - a challenge and an opportunity for process systems engineering. Computers & Chemical Engineering 133:106629 (2020). https://doi.org/10.1016/j.compchemeng.2019.106629
4. "Circular economy introduction." Accessed: Nov. 05, 2025. [Online]. Available: https://www.ellenmacarthurfoundation.org/topics/circular-economy-introduction/overview
5. Latini D, Vaccari M, Lagnoni M, Orefice M, Mathieux F, Huisman J, Tognotti L, Bertei A. A comprehensive review and classification of unit operations with assessment of outputs quality in lithium-ion battery recycling. Journal of Power Sources 546:231979 (2022). https://doi.org/10.1016/j.jpowsour.2022.231979
6. Ahmed A, Nair A, Torres AI. Design and optimization of circular economy networks—a case study of PET. Computers & Chemical Engineering 200:109164 (2025). https://doi.org/10.1016/j.compchemeng.2025.109164
7. Baratsas SG, Pistikopoulos EN, Avraamidou S. A systems engineering framework for the optimization of food supply chains under circular economy considerations. Science of The Total Environment 794:148726 (2021). https://doi.org/10.1016/j.scitotenv.2021.148726
8. Iakovou E, Pistikopoulos EN, Walzberg J, Iseri F, Iseri H, Chrisandina NJ, Vedant S, Nkoutche C. Next-generation reverse logistics networks of photovoltaic recycling: perspectives and challenges. Solar Energy 271:112329 (2024). https://doi.org/10.1016/j.solener.2024.112329
9. Kondili E, Pantelides CC, Sargent RWH. A general algorithm for short-term scheduling of batch operations—i. MILP formulation. Computers & Chemical Engineering 17:211-227 (1993). https://doi.org/10.1016/0098-1354(93)80015-f
10. Gutsch M, Leker J. Costs, carbon footprint, and environmental impacts of lithium-ion batteries – from cathode active material synthesis to cell manufacturing and recycling. Applied Energy 353:122132 (2024). https://doi.org/10.1016/j.apenergy.2023.122132
11. Q. Dai, J. Spangenberger, S. Ahmed, L. Gaines, J. Kelly, and M. Wang, "EverBatt: A Closed-loop Battery Recycling Cost and Environmental Impacts Model, " ANL--19/16, 1530874, 153050, Apr. 2019. doi: 10.2172/1530874.
12. Rosenberg S, Glöser-Chahoud S, Huster S, Schultmann F. A dynamic network design model with capacity expansions for eol traction battery recycling – a case study of an OEM in germany. Waste Management 160:12-22 (2023). https://doi.org/10.1016/j.wasman.2023.01.029

13. Soldan Cattani N, Weinert C, Mussehl V, Frieges M, Kampker A. Economic and structural challenges of lithium-ion battery recycling in europe: a stakeholder-based assessment. Waste Management 205:114962 (2025). https://doi.org/10.1016/j.wasman.2025.114962

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Integration of carbon dioxide capture in a wine effluent biorefinery through the use of deep eutectic solvents

Carlos E. Guzmán Martínez[a], Valeria Caltzontzin Rabell[a], Sergio I. Martínez Guido[a], Salvador Hernández[b] and Claudia Gutiérrez Antonio[a]*

[a] Universidad Autónoma de Querétaro, Facultad de Ingeniería, Querétaro, Querétaro, Mexico
[b] Universidad de Guanajuato, Departamento de Ingeniería Química, Guanajuato, Guanajuato, Mexico
* Corresponding Author: claudia.gutierrez@uaq.mx

ABSTRACT

The wine industry generates large volumes of organic effluents, whose inadequate management poses significant environmental challenges but also offers opportunities for resource recovery. In this work, an integrated biorefinery scheme for the valorization of winery effluents is proposed and evaluated through steady-state simulation in Aspen Plus®. The biorefinery converts winery wastewater into a portfolio of value-added chemicals and biofuels, including levulinic acid, propylene glycol, formic acid, light gases, naphtha, sustainable aviation fuel, green diesel, and bioethanol, while enabling water recovery and carbon dioxide management. Two alternative CO_2 capture routes are analyzed and compared: a conventional CaO-based carbonation–calcination process and an innovative absorption system using deep eutectic solvents (DES), specifically choline chloride–urea. Technical performance is assessed through chemical oxygen demand (COD) removal, recovery, conversion, yield, and product mass ratios. Economic feasibility is evaluated using profit-based indicators, while environmental performance is quantified through CO_2-equivalent emissions associated with utility consumption. Results show that the proposed biorefinery achieves a COD removal efficiency of 99.99%, producing treated water compliant with Mexican regulations for internal reuse. The DES-based configuration reduces raw material costs by 48.86%, enables hydrogen recovery as an additional valuable product, and increases overall profit by 2.86% compared to the CaO-based scheme. Although the relative reduction in total CO_2 emissions is modest (≈0.5%), the DES configuration achieves an absolute annual reduction of 2, 198 t CO_2. Overall, the results demonstrate that integrating DES-based CO_2 capture into winery effluent biorefineries enhances economic performance and supports circular economy principles through waste valorization, water reuse, and emissions mitigation.

Keywords: Deep Eutectic Solvents, biorefinery, wine effluents, techno-economic analysis

INTRODUCTION

The wine industry plays an essential role in the global economy, with sales ranging between €25.8 and €31.4 billion [1]. Mexico, ranked 32nd among the world's wine producers, contributed approximately 400, 000 hectoliters in 2022 [2]. Moreover, the wine industry stands out due to the considerable volumes of organic effluents generated during pretreatment, fermentation, clarification, and bottling operations. On average, 2–4 liters of effluent are generated per liter of wine produced, underscoring the need for sustainable management strategies. In this context, waste valorization offers a significant opportunity [3].

The wine effluents contain high concentrations of carbohydrates, alcohols, and organic matter. Because of their rich organic and mineral content, winery effluents are often reused for irrigation or discharged into water bodies and surrounding environments. However, if not properly treated, their high organic load and seasonal variability pose serious environmental risks. The impacts extend beyond aquatic systems to terrestrial ones. On land, effluents can cause salinization, acidification, structural soil damage, and unpleasant odors, while reducing

mineral availability, microbial diversity, and biological productivity. In aquatic systems, they can lead to turbidity, eutrophication, oxygen depletion, decreased photosynthesis, and toxic effects on aquatic life [3]. Nonetheless, their organic content also makes these effluents a promising feedstock for biorefinery systems. Due to this, several strategies have been analyzed to treat and valorize effluents derived from the wine industry. Among the analyzed processes, the valorization of wine effluents through biorefinery schemes is highlighted.

A biorefinery integrates a set of processes designed to convert wine effluents into value-added products, chemicals, biofuels, and bioenergy. In this processing scheme, the inclusion of carbon capture technologies aligns with circular-economy principles, promoting closed-loop systems that minimize waste and emissions. Among the available CO_2-capture approaches, absorption with chemical solvents is one of the most efficient and technologically mature. However, conventional amine-based solvents such as monoethanolamine (MEA) have significant drawbacks, including high regeneration energy requirements, solvent degradation, and environmental toxicity. As an alternative, deep eutectic solvents (DES) have emerged as a new generation of environmentally friendly solvents with outstanding physicochemical properties.

Various DES formulations have been explored, such as choline chloride/ethanolamine, choline chloride/L-arginine/glycerol, and choline chloride/1, 2-propanediol. Notably, the combination of choline chloride and urea forms a liquid solvent known as choline chloride:urea (ChCl:U) or reline, which demonstrates a high capacity for selectively absorbing CO_2 from gas mixtures [4].

Therefore, the research aim is to design and evaluate a conceptual process that integrates deep eutectic solvents-based CO_2 absorption with a biorefinery scheme. This scheme treats winery effluents to produce levulinic acid, sustainable aviation fuel (SAF), naphtha, light gas, formic acid, green diesel, propylene glycol, dipropylene glycol, and bioethanol. Additionally, in this biorefinery, water can be recovered to quality levels permissible under Mexican regulations.

METHODOLOGY

In general, the procedure implemented to develop the proposed biorefinery is summarized in the block flowchart shown in Figure 1.

First, the case study is defined in terms of the effluent composition and its mass flow rate. Subsequently, the target products are established, together with the unit operations required for their recovery and/or production; this stage corresponds to the conceptual process design.

To satisfy the CO_2 capture requirement within the proposed biorefinery, two alternative capture routes are developed. The first route integrates CO_2 capture using deep eutectic solvents (DESs), whereas the second route is based on CaO as the capturing agent. Each route is incorporated into the overall process flowsheet to meet the CO_2 capture demand of the system. Consequently, these alternatives define two different biorefinery schemes, each one characterized by its corresponding CO_2 capture technology.

The kinetic and thermodynamic information reported in the literature is compiled and used to support the process modeling. Afterwards, the manipulated variables for each unit operation are identified, and the complete biorefinery is modeled and simulated under steady-state conditions using Aspen Plus® V15. Finally, technical, economic, and environmental assessments are carried out to evaluate the overall performance of the proposed biorefinery.

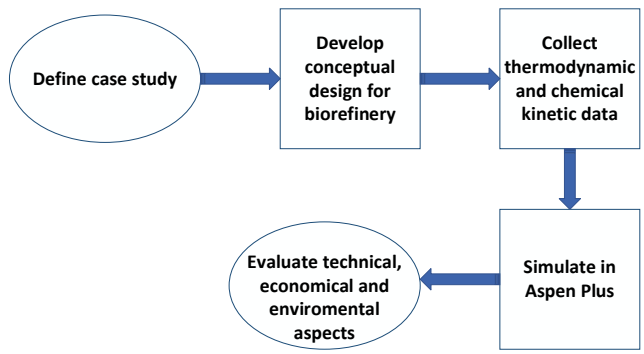

Figure 1. Methodological strategy

The technical evaluation aims to quantify the capacity of the biorefinery to remove contaminants from the effluent. This assessment is performed using the Chemical Oxygen Demand (COD) as a key indicator, calculated directly in Aspen Plus. Specifically, the technical performance is quantified through the relative reduction between the initial COD of the wine effluent and the COD of the treated water. In addition, several technical indicators are computed to characterize the overall process performance, namely: component recovery (rec), conversion (cv), yield (Yd), and product mass ratio (pmr), as defined by Equations (1)–(4) [5].

$$rec = \frac{oms}{fms} \tag{1}$$

$$cv = \frac{\dot{r}_{in} - \dot{r}_{out}}{\dot{r}_{in}} \tag{2}$$

$$yd = \frac{pmf}{tpm} \tag{3}$$

$$pmr = \frac{mp}{rmm} \tag{4}$$

Where: rec = component recuperation, oms = output mass flow (kg/h), fms = feed mass flow (kg/h) cv =conversion, r_{in} = reactant input to process (mol/h), r_{out} = reactant output to process (mol/h), Yd = Yield, pmf = moles of desired product formed (mol), tpm = theoretical product moles (mol), pmr = product mass ratio, mp = product mass flow (kg/h), rmm = raw material mass flow (kg/h)

The economic evaluation is based on the estimation of the process profit (PFT), defined as the difference between incomes and outcomes. Incomes correspond to product sales, whereas outcomes include raw material costs, total utility costs, and equipment-related costs. Equipment costs are estimated using the Guthrie method [6]. Since equipment investment represents a time-independent cost, an annualization factor (k) is incorporated to express this contribution on an annual basis [6]. The mathematical framework for the economic assessment is described by Equations (5)–(12). The parameters required to solve Equations (5)–(12), including Guthrie constants, are taken from Tourton et al. [6]. Utility prices were obtained from the Aspen Plus® V15 database. It is important to note that Guthrie parameters were originally developed in 2001; therefore, equipment costs were updated to 2025 values using the Marshall and Swift index and Equation (10) [6].

Finally, the environmental evaluation focuses on the environmental burden associated with utility consumption. The main indicator employed is the total CO_2-equivalent emissions, defined by Equation (13) and based on US-EPA-RULE-E9-5711. This parameter is also calculated using Aspen Plus.

$$PFT = (Ah)(Psale - RMC - UC_T) - (k)(EC) \tag{5}$$

$$Psale = \sum_{1}^{j}\left((mp_j)(pp_j)\right) \tag{6}$$

$$RMC = \sum_{1}^{g}\left((mp_g)(pp_g)\right) \tag{7}$$

$$UC_T = \sum_{1}^{i}(UC_i) \tag{8}$$

$$k = \frac{y(1+y)^h}{(1+y)^h - 1} \tag{9}$$

$$EC = C_2 = C_1\left(\frac{I_2}{I_1}\right) \tag{10}$$

$$C_{BM} = C_p^o(B_1 + B_2F_MF_P) \tag{11}$$

$$log_{10}C_p^o = \left(K_1 + K_2log_{10}(A) + K_3\left(log_{10}(A)\right)^2\right) \tag{12}$$

$$TCO_2e = \sum_{1}^{l}(CO_2e_l) \tag{13}$$

Where: PFT = profit (USD/year), Ah = hours per year (8766 h/year), PSALE = Product sales (USD/h), RMC = Raw material cost (USD/h), UCT = Total utilities cost (USD/h), k = annualization factor, EC = Equipment cost (USD), mp = product mass flow (kg/h), rmm = raw material mass flow (kg/h), pp = product price (USD/kg), prmm = raw material mass price (USD/kg), UC = Utility cost, y = interest rate (10 %), h = Life period of each technology (20 years), C_p^o= purchased cost for condition (USD), A = the capacity or size parameter for the equipment, (K_1, K_2, K_3, B_1, B_2) = Guthrie constants, FM = Material factor, FP = Pressure factor, CBM = bare module cost (USD), C_1 = Purchased cost to base time (USD), C_2 = Purchased cost to time when cost is desired (USD), I_1 = Cost index to base time, I_2 = Cost index to time when cost is desired, TCO_2e= Total CO_2 equivalent (ton CO_2 equivalent/year), CO_2e_l= CO_2 equivalent per utility (ton CO_2 equivalent/year). Subscripts: j = product stream, g = raw material stream, i = utility.

RESULTS

As previously mentioned, Mexico reported a wine production of 400, 000 hectoliters in 2022 [2]. According to Buitrón et al. [3], the wine industry generates between 0.2 and 4 L of wastewater per liter of wine produced. For the definition of the study case, the average value within this range was adopted, resulting in 2.1 L of wastewater per liter of wine. Therefore, the annual volume of wine effluent released in Mexico was estimated at 8.4 × 10^7 L·year^{-1}, which corresponds to 9, 589.04 L·h^{-1}. For modeling purposes, this value was rounded and, assuming the effluent density is close to that of water (as water is the predominant component), the study case was defined as a mass flow rate of 10, 000 kg·h^{-1}.

According to Buitrón et. al. [3] winery effluent is composed of several components; however, the most abundant are water, solids, ethanol, and glucose. On the other hand, De Iseppi et al. [7] report that the solid fraction is constituted by 25 wt% lees, 25 wt% lignin, 25 wt% cellulose, and 25 wt% hemicellulose, mainly. It is important to mention that wine lees are dead yeast that are deposited and discarded at the clarification step of wine production. Table 1 shows the effluent characterization.

Regarding the thermodynamic model, it is not possible to use only one for the entire biorefinery; this is because the process uses several components, which have different characteristics such as organic or inorganic nature, in different amounts; therefore, thermodynamics models for each section are defined by Carlson algorithm [8]. Those used in the entire process are NRTL, BK10, and PENG-ROBINSON.

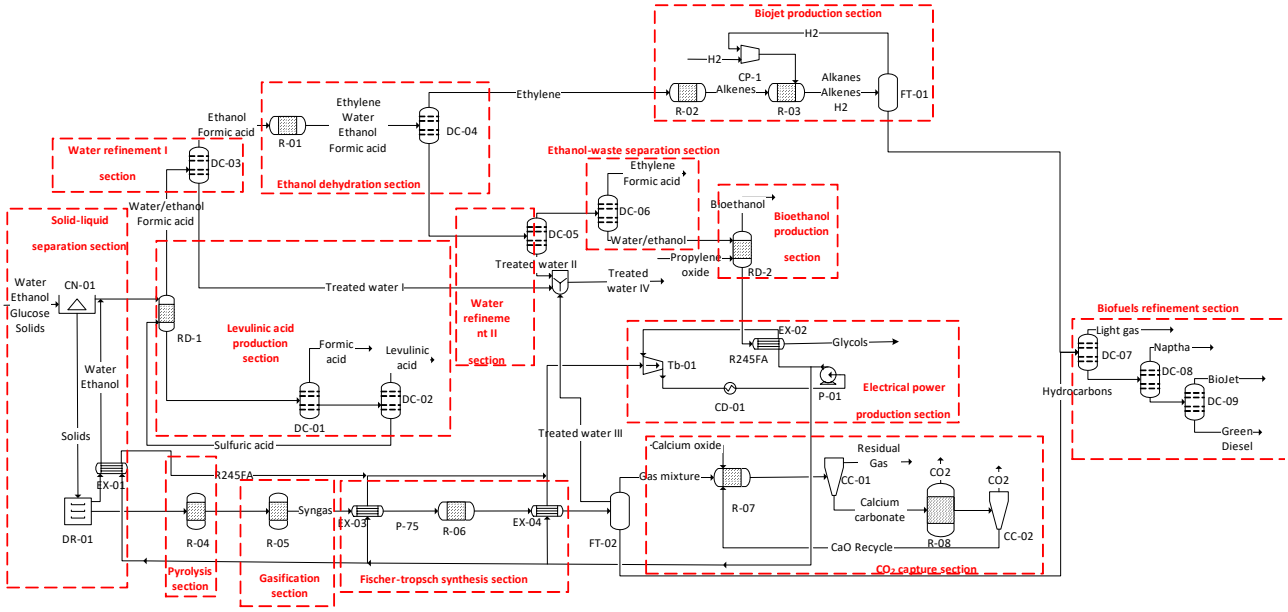

Figure 2. Biorefinery process flowsheet for the CaO-based CO_2 capture configuration.

Table 1: Winery effluent characterization

Winery effluent	Value
Mass flow (kg/h)	10, 000
Water (wt%)	71.76
Ethanol (wt%)	7.39
Glucose (wt%)	5.26
Solids (wt%)	15.59
DQO (PPM)	210, 052

The overall biorefinery scheme is presented in Figure 2. For clarity, the flowsheet is divided into several process sections. Figure 2 represents both the CaO-based and DES-based configurations; however, the only difference between them lies in the CO_2 capture section. For the DES configuration, the corresponding CO_2 capture section is shown separately in Figure 3, where CO_2 is captured using DESs [9].

The biorefinery process (Figure 2) begins with the separation of suspended solids from the bulk effluent. This step is performed using a centrifuge, and the recovered solids are routed to an independent processing pathway. The clarified liquid stream, mainly composed of glucose, ethanol, and water, is subsequently fed to a reactive distillation column (RD-01). In this unit, glucose is first converted to 5-hydroxymethylfurfural (5-HMF) and then further transformed into levulinic acid and formic acid. These reactions correspond to R1 and R2, respectively, while the kinetic parameters employed in the model are summarized in Table 2. The reaction system is catalyzed by H_2SO_4. Finally, the resulting acid mixture is purified using two conventional distillation columns (DC-01 and DC-02), enabling catalyst recovery and recycling.

$$C_6H_{12}O_6 \rightarrow C_6H_6O_3 + 3H_2O \quad \text{(R1)}$$

$$C_6H_6O_3 + 2H_2O \rightarrow CH_2O_2 + C_5H_8O_3 \quad \text{(R2)}$$

Table 2: Kinetic parameters for levulinic acid production [10]

Reaction	**Reaction R1**	**Reaction R2**
Preexponential factor (s^{-1})	2.09×10^7	27, 261.3
Activation energy (kJ/mol)	86.33	56.95

An overhead stream containing water, ethanol, and formic acid is withdrawn from RD-01. This stream is then fed to a third distillation column (DC-03), where ethanol and formic acid are recovered as the distillate, while treated water is obtained as the bottom product. The ethanol-rich stream is subsequently sent to reactor R-01, where ethanol is dehydrated to produce ethylene and water.

The ethylene produced in R-01 is purified in distillation column DC-04. The overhead stream, rich in ethylene, is sent to reactor R-02, where the polymerization step is carried out. The bottom stream from DC-04, containing heavier compounds and impurities, is directed to DC-05 for further treatment.

In R-02, the reaction yields a hydrocarbon mixture mainly composed of alkanes. This stream is subsequently hydrogenated to ensure full saturation of the products. Unreacted hydrogen is recovered in a flash tank (FT-01) and recycled, while the resulting hydrocarbon stream is

routed to the biofuels upgrading section.

Regarding the waste stream from DC-04, the water content is separated as the bottom product of DC-05, generating a second treated-water stream. The distillate from DC-05 is further processed in DC-06. In this column, ethylene and formic acid are recovered as the overhead product, whereas an ethanol–water mixture is obtained as the bottom stream.

The ethanol–water mixture is subsequently used as feedstock for bioethanol production. Ethanol dehydration is performed in a reactive distillation unit, where water reacts with propylene oxide to form propylene glycol. In addition, a side reaction between propylene oxide and propylene glycol leads to the formation of dipropylene glycol. These reactions correspond to R3 and R4, respectively, and the kinetic parameters employed are summarized in Table 3.

$$C_3H_6O + H_2O \rightarrow C_3H_8O_2 \quad \text{(R3)}$$

$$C_3H_6O + C_3H_8O_2 \rightarrow C_6H_{14}O_3 \quad \text{(R4)}$$

Table 3: Bioethanol production. Reactive dehydration parameters [11]

Reaction	**Reaction R3**	**Reaction R4**
Preexponential factor (s^{-1})	$9.6x10^7$	75, 362
Activation energy (kJ/mol)	$9.6x10^6$	82, 899

The solid fraction separated from the effluent is dried in DR-01. The generated vapor is condensed prior to being recycled to RD-01.

The dried solids are subsequently subjected to pyrolysis in reactor R-04, yielding bio-oil and biochar. These products are then routed to the gasification section (R-05), where they are converted into syngas. The produced syngas is upgraded via a Fischer–Tropsch (FT) synthesis in reactor R-06 [12], generating a broad hydrocarbon distribution ranging from C1 to C30. However, water and CO_2 are also formed as by-products. Therefore, the multiphase effluent from R-06 is sent to a flash tank (FT-02), which separates the stream into three phases. The aqueous phase is recovered as the third treated-water stream in the overall process. The hydrocarbon phase is directed to the biofuels refinement section, while the gas phase, mainly composed of CO_2, along with methane and hydrogen, is routed to the CO_2 capture section.

In the CaO-based capture configuration, the gas stream is fed to reactor R-07, where a selective carbonation reaction takes place. In this step, calcium oxide reacts with CO_2 to form calcium carbonate ($CaCO_3$), as described by reaction R5. The kinetic model employs an activation energy of 0.3699 kJ/mol and a pre-exponential factor of 10.2 s^{-1} [13].

$$CaO + CO_2 \rightarrow CaCO_3 \quad \text{(R5)}$$

After CO_2 capture, the remaining gas stream, rich in methane and hydrogen, is recovered as a potential fuel product using cyclone CC-01. In parallel, the calcium carbonate particles are sent to reactor R-08, where calcination is performed to release the previously captured CO_2 and regenerate CaO. The resulting stream exiting R-08 is routed to a second cyclone (CC-02) to separate residual CO_2 from the regenerated solid. Finally, the recovered CaO is recycled back to reactor R-07.

The biofuels upgrading section consists of three distillation columns. Light gases are recovered as the overhead stream from DC-07, naphtha is obtained as the distillate from DC-08, and sustainable aviation fuel is recovered as the overhead stream from DC-09. The bottom product of DC-09 corresponds to green diesel.

Electrical power generation is achieved through an organic Rankine cycle (ORC) employing R245fa as the working fluid. The ORC recovers waste heat from different sections of the biorefinery through heat exchangers, including streams associated with DR-01, R-05, R-06, and RD-02.

In addition to the CaO-based configuration, a DES-based CO_2 capture alternative is assessed. The only difference between both biorefinery schemes lies in the CO_2 capture section. In the DES configuration, the gas mixture from FT-02 is fed to the bottom of an absorption column (AC-01), while a deep eutectic solvent (DES) is introduced from the top. The DES consists of choline chloride:urea and is used in a 1:1 DES:water ratio. The solvent selectively absorbs carbon-containing compounds, enabling hydrogen recovery as the overhead product, which is considered a valuable co-product.

The CO_2-loaded (exhausted) solvent stream (Ex-DES) is subsequently sent to flash tanks FL-03 and FL-04 for solvent regeneration and CO_2 release. The overhead stream from FL-04 corresponds to a residual gas containing approximately 80 mol% CO_2. The process flowsheet of the DES-based CO_2 capture configuration is presented in Figure 3.

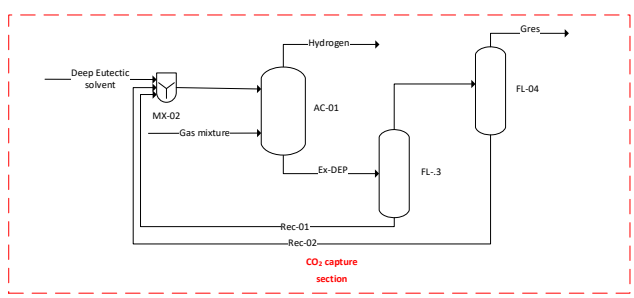

Figure 3. Process flowsheet for the DES-based CO_2 capture configuration.

The product slate obtained for each biorefinery configuration is summarized in Table 4. The reported water product corresponds to the total amount of treated water, obtained as the sum of all treated-water streams generated throughout the process. The resulting treated-water stream exhibits a COD of 19.1 ppm.

As shown in Table 4, the overall product yields are identical for both biorefinery configurations; the only differences are associated with the waste and hydrogen streams. In both the CaO-based and DES-based schemes, the waste stream is mainly composed of (i) the ethylene–formic acid stream generated in the ethanol/waste separation section, in addition to (ii) a residual gas stream and (iii) a CO_2-rich stream.

In both configurations, treated water is the dominant product, accounting for 70.5% of the total output. This is followed by sustainable aviation fuel, levulinic acid, naphtha, and light gases, contributing 4.13%, 2.20%, 2.10%, and 1.53%, respectively. The primary difference between the two schemes lies in the relative proportions of waste and hydrogen streams. In the CaO-based configuration, waste streams account for 17.21% of the total products, whereas in the DES-based configuration this fraction is reduced to 16.20%. As a result, hydrogen emerges as an additional valuable product in the DES-based scheme, representing 1.16% of the overall product distribution.

Table 4: Products obtained from the different schemes.

Product (kg/h)	**CaO**	**DES**
Scheme	Figure 2	Figure 3
Water	7,094.84	7,094.84
Levulinic acid	221.57	221.57
Light gases	153.62	153.62
Naphtha	211	211
Sustainable aviation fuel	415	415
Green diesel	74.8	74.8
Bioethanol	6.93	6.93
Di/Propylene glycol	49.3	49.3
Formic acid	96.4	96.4
Waste	1,730.26	1,631.49
Hydrogen	-----	117.01

The chemical oxygen demand (COD) removal capacity is defined as the relative reduction between the initial COD of the winery effluent and the final COD of the treated-water product. Based on this definition, the proposed biorefinery achieves a COD reduction of 99.99%, fully complying with the maximum COD limits established by Mexican regulations for water intended for human use (COD < 70 ppm), as specified in NOM-001-SEMARNAT-2021 and NOM-127-SSA1-2021. Nevertheless, even if the treated water were to reach excellent quality (e.g., COD ≈ 0 ppm), NOM-127-SSA1-2021 does not allow its direct use for drinking purposes. Instead, the treated-water stream is suitable for internal reuse, such as recycled water within vineyard operations.

Regarding technical performance indicators, recovery is calculated exclusively for water and ethanol, as these are considered the relevant products of the biorefinery. Water recovery reaches 98.87% (Eq. (1)), based on the water content in the winery effluent as the input and the total treated water produced by the biorefinery as the output. In terms of sectional contributions, 84.73% of the recovered water is obtained from treated-water stream I, 13.56% from treated-water stream II, and the remaining fraction is generated within the Fischer–Tropsch section.

Ethanol recovery is relatively low (0.95%), as ethanol is produced solely as a bioethanol stream. It is important to emphasize that this outcome does not indicate an inefficient biorefinery configuration. Rather, it reflects the intentional process design, where ethanol contained in the winery effluent is primarily valorized as a feedstock for sustainable aviation fuel production pathway. Thus, only a small fraction of ethanol is available for recovery.

Table 5: Technical indicators

Technical parameter	**CaO**	**DES**
Scheme	Figure 2	Figure 3
Recuperation		
Water	98.87%	98.87%
Ethanol	0.95%	0.95%
Conversion		
Glucose	100%	100%
Solids	100%	100%
Ethanol	99.05%	99.05%
Calcium oxide	99.86%	-----
Yield		
Levulinic acid	65.19%	65.19%
Glycol	75.57%	75.57%
Ethylene	98.8%	98.8%
Calcium carbonate	99.99%	-----
Production mass ratio		
Light gases	0.0154	0.0154
Naphtha	0.0211	0.0211
Green diesel	0.0075	0.0075
Sustainable aviation fuel	0.0415	0.0415
Levulinic acid	0.0221	0.0221
Bioethanol	0.0007	0.0007
Water	0.7095	0.7095
Glycols	0.0049	0.0049
Formic acid	0.0088	0.0088
Wastes	0.173	0.1631

The product ratios reported in Table 5 provide valuable input for subsequent systems-level analyses, including supply chain development. These ratios can be directly incorporated into mathematical models to represent process outputs, thereby avoiding the introduction

of complex nonlinearities that typically arise when embedding detailed process simulations into optimization frameworks.

The main economic indicators of the proposed biorefinery are summarized in Table 6. A comparative assessment of the economic indicators for the CaO- and DES-based configurations reveals a significant reduction of 48.86% in raw material costs when CaO is replaced by DES. This reduction is primarily attributed to the internal generation of hydrogen within the CO_2 capture section, which supplies the hydrogen required for alkene hydrogenation in the alcohol-to-jet pathway. Specifically, hydrogen is produced in unit AC-01. In contrast, in the CaO-based configuration, hydrogen must be purchased as a raw material, accounting for 50.16% of the total raw material cost. The DES-based scheme also exhibits a 2.45% increase in product sales relative to the CaO-based configuration. This improvement is mainly associated with hydrogen recovery as a valuable product. In the CaO scheme, hydrogen is effectively wasted due to the lack of purification capability. In the DES-based process, hydrogen becomes the third most abundant product, following levulinic acid and sustainable aviation fuel, contributing 91.22% and 3.39% of the product distribution, respectively.

Table 6: Economic indicators

	CaO	**DES**
Scheme	Figure 2	Figure 3
Raw material (USD/y)	$1,138,265.10	$582,134.16
Products (USD/y)	$255,625,939.62	$261,888,720.64
Utilities (USD/y)	$21,848,553.72	$21,901,820.16
Equipment cost (USD)	$50,954,841.13	$53,374,306.62
Annualized equipment cost (USD/y)	$5,985,136.52	$6,269,326.03
Total annual cost (USD/y)	$28,971,887.65	$28,753,280.34
Profit (USD/y)	$226,654,024.68	$233,135,440.3

Despite these advantages, utility and equipment costs increase by 0.24% and 4.75%, respectively, in the DES-based configuration compared to the CaO-based scheme. The higher equipment cost is mainly driven by the inclusion of a compressor required to feed CO_2 to unit AC-01 and a heat exchanger used to pretreat the DES prior to absorption, operating at 20 atm and 15 °C. In the DES-based scheme, the compressor and heat exchanger represent 2.67% and 2.09% of the total equipment cost, respectively. By comparison, all equipment associated with CO_2 capture in the CaO-based scheme accounts for only 0.57% of the total equipment cost.

Regarding utility costs, the CO_2 capture section in the CaO-based configuration requires $32.97 USD h^{-1}, while the DES-based configuration requires $38.76 USD h^{-1}. These values correspond to 1.32% and 1.55% of the total utility costs of each process, respectively, explaining the observed 0.24% increase. In the CaO-based scheme, utility consumption is mainly associated with reactors R-07 and R-08, whereas in the DES-based scheme it includes the compressor feeding AC-01, the heat exchanger, and flash tanks FL-03 and FL-04.

Overall, the total annual cost is reduced by 0.75% in the DES-based configuration relative to the CaO-based scheme. As a result, the difference between revenues and operating costs leads to a 2.86% increase in profit for the DES-based process.

Table 7 summarizes the selected environmental indicators used for CO_2 quantification. In terms of CO_2-equivalent emissions associated with utilities, the DES-based scheme achieves a reduction of 0.48% compared to the CaO-based configuration. Although the CaO-based scheme releases a substantial amount of CO_2 due to $CaCO_3$ calcination during CaO regeneration, the DES-based scheme is unable to significantly reduce overall emissions. This limitation is mainly attributed to the large flowrate of exhausted DES (12.17 t h^{-1}) and the energy-intensive desorption step, which requires heating and handling a high mass of solvent, resulting in CO_2 emissions comparable to those of the CaO-based process.

Table 7: Environmental indicators

	CaO	**DES**
Scheme	Figure 2	Figure 3
CO_2e utilities (Ton CO_2/y)	431,843.59	429,766.75
CO_2 streams (Ton CO_2/y)	12,015.37	11,894.15
Total	443,858.96	441,660.90

The CO_2 streams correspond to the desorption stages in which both CaO- and DES-based systems release the captured CO_2. The DES-based configuration exhibits a slightly lower CO_2 release, with a reduction of 1.01% compared to the CaO-based scheme. This behavior is attributed to partial CO_2 retention within the DES matrix, whereas in the CaO-based process CO_2 is almost entirely released during regeneration.

Although the relative difference in total CO_2 emissions between the CaO- and DES-based configurations is modest (approximately 0.5%), the absolute impact is significant. The DES-based process enables an annual reduction of 2, 198 t CO_2, highlighting its environmental

relevance despite the small percentage difference.

CONCLUSIONS

This study demonstrates the technical and economic feasibility of producing value-added chemicals and biofuels from winery effluents through the modelling and simulation of an integrated biorefinery scheme. The proposed process enables the recovery of levulinic acid, propylene glycol, formic acid, and biofuels, including light gases, naphtha, sustainable aviation fuel, green diesel, and bioethanol, transforming an agro-industrial waste stream into a diversified portfolio of marketable products.

From a circular economy perspective, the proposed biorefinery closes material and energy loops by simultaneously valorizing organic matter, recovering water for internal reuse, and separating CO_2 emissions from hydrogen mixtures through the incorporation of deep eutectic solvents (DES) as a capture and utilization strategy. This approach reduces dependency on external raw materials, lowers total annual costs, and enhances profitability by integrating hydrogen as an additional value-added product. Moreover, the DES-based configuration contributes to emissions mitigation, reinforcing the environmental performance of the system while maintaining economic competitiveness. For the selected case study, in Baja California, Mexico, the biorefinery achieves attractive economic indicators and complies with national environmental and water quality regulations. In particular, a COD removal efficiency of 99.99% allows the recovery of treated water suitable for internal reuse, in accordance with NOM-001-SEMARNAT-2021 and NOM-127-SSA1-2021.

Overall, the proposed scheme exemplifies a circular and sustainable pathway for winery effluent management; this proposal allows shifting from waste treatment to resource recovery, and highlighting the role of integrated biorefineries as enabling technologies for the transition toward low-carbon and circular agro-industrial systems.

ACKNOWLEDGEMENTS

The authors acknowledge the support provided by Secretaría de Ciencia, Humanidades, Tecnología e Innovación (SECIHTI, México) for the financial support granted to develop the post-doctoral research (grant no. 2397403) of C. E. Guzmán-Martínez.

AUTHOR IDENTIFIERS

Author ORCIDs:
Guzmán-Martínez C.E: 0000-0002-3404-2650
Caltzontzin-Rabell V: 0000-0002-5739-0957
Martínez-Guido S.I: 0000-0002-0198-9338
Hernández S: 0000-0002-4598-3392
Gutiérrez-Antonio C: 0000-0002-7557-2471

REFERENCES

1. Mader AE, Holtman GA, Welz PJ. Treatment wetlands and phyto-technologies for remediation of winery effluent: challenges and opportunities. Science of The Total Environment 807:150544 (2022). https://doi.org/10.1016/j.scitotenv.2021.150544
2. https://www.oiv.int/sites/default/files/documents/OIV_State_of_the_world_Vine_and_Wine_sector_in_2022_2.pdf〉(accessed February 25).
3. Buitrón G, Martínez-Valdez FJ, Ojeda F. Biogas production from a highly organic loaded winery effluent through a two-stage process. Bioenerg. Res. 12:714-721 (2019). https://doi.org/10.1007/s12155-019-09984-7
4. Romero-García AG, Ramírez-Márquez C, Sánchez-Ramírez E, Ponce-Ortega JM, González-Campos JB, De Blasio C, Segovia-Hernández JG. Implementation of the deep eutectic solvent, choline urea chloride (1:2), to evaluate the sustainability of its application during CO2 capture. Process Integr Optim Sustain 8:741-758 (2023). https://doi.org/10.1007/s41660-023-00383-2
5. Felder, R. M., Rousseau, R. W., & Bullard, L. G. (2020). Elementary principles of chemical processes. John Wiley & Sons.
6. Turton, R., Bailie, R. C., Whiting, W. B., & Shaeiwitz, J. A. (2008). Analysis, synthesis and design of chemical processes. Pearson Education.
7. De Iseppi A, Lomolino G, Marangon M, Curioni A. Current and future strategies for wine yeast lees valorization. Food Research International 137:109352 (2020). https://doi.org/10.1016/j.foodres.2020.109352
8. Carlson, E. C.: Don't Gamble with Physical Properties for Simulations. Chemical Engineering Progress, 35–46 (1996)
9. Romero-García AG, Ramírez-Corona N, Sánchez-Ramírez E, Alcocer-García H, De Blasio C, Segovia-Hernández JG. Sustainability assessment in the CO2 capture process: multi-objective optimization. Chemical Engineering and Processing - Process Intensification 182:109207 (2022). https://doi.org/10.1016/j.cep.2022.109207
10. Solis-Sanchez JL, Alcocer-Garcia H, Sanchez-Ramirez E, Segovia-Hernandez JG. Innovative reactive distillation process for levulinic acid production and purification. Chemical Engineering Research and Design 183:28-40 (2022). https://doi.org/10.1016/j.cherd.2022.04.041
11. Guzmán-Martínez, C. E., Castro-Montoya, A. J., &

Nápoles-Rivera, F. (2019). Economic and environmental comparison of bioethanol dehydration processes via simulation: reactive distillation, reactor–separator process and azeotropic distillation. Clean Technologies and Environmental Policy, 21, 2061-2071.
12. Guzmán-Martínez CE, Caltzontzin-Rabell V, Martínez-Guido SI, Gutiérrez-Antonio C. Valorization of effluents from the wine industry through a biorefinery scheme to obtain sustainable aviation fuel, levulinic acid, water, and value-added compounds. Biochemical Engineering Journal 221:109798 (2025). https://doi.org/10.1016/j.bej.2025.109798
13. Kinoshita C. Production of hydrogen from bio-oil using cao as a CO2 sorbent. International Journal of Hydrogen Energy : (2003). https://doi.org/10.1016/s0360-3199(02)00203-3

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Techno-economic feasibility of gallium recovery from semiconductor wastewater

Kilian Kozerke[ab]*, Thomas Eberius[b], Nathanial J. Cooper[a]

[a] University of Cambridge, Department of Engineering, Cambridge, UK
[b] RWTH Aachen University, Faculty of Mechanical Engineering, Aachen, Germany
* Corresponding Author: kilian.kozerke@rwth-aachen.de, kozerke.work@gmail.com

ABSTRACT

Gallium is a critical material with increasing demand driven by compound semiconductors such as gallium nitride (GaN) and gallium arsenide (GaAs) used in power electronics and optoelectronics and a highly concentrated supply chain, with 98 % of refined production occurring in China. While recycling remains limited, GaAs semiconductor manufacturing generates wastewater that can contain gallium concentrations ranging from 1–35 mg·L^{-1}, representing an underutilized secondary resource. This study evaluates the technical and economic feasibility of recovering gallium from GaAs semiconductor wastewater across an input range of 10–100 m^3·d^{-1} using process modelling and a techno-economic analysis comparing two alternative separation routes: ion exchange (IX) and solvent extraction (SX). Using a real-world industrial wastewater composition, IX achieves a higher overall recovery than single-stage SX (80.40 % vs. 62.54 %), which translates into consistently lower levelized costs of gallium. The resulting levelized cost decreases from 563 to 144 €·kg^{-1} for IX and from 701 to 161 €·kg^{-1} for SX. The results demonstrate that separation step efficiency is the dominant determinant of economic performance and show that gallium recovery from semiconductor wastewater can be economically viable at fab-relevant scale. The analysis provides a proof-of-concept for integrating gallium recovery into semiconductor wastewater management and highlights its potential contribution to resource security and circular material supply.

Keywords: Gallium recovery, GaAs semiconductor wastewater, techno-economic analysis, ion exchange, solvent extraction, critical raw material, circular economy

INTRODUCTION

The compound semiconductor industry is a key enabler of renewable energy and low-carbon technologies through materials such as gallium arsenide (GaAs) and gallium nitride (GaN), which enable high-efficiency power electronics and optoelectronic devices [1]. Semiconductor applications include photovoltaics, power conversion in wind and solar technologies and the electrification of transport [1].

As demand increases, gallium has become a critical material with highly concentrated supply chains, as 98% of refined production occurs in China [2]. As a companion metal primarily derived from bauxite processing and aluminum refining, independent scale-up is challenging. Global demand is projected to exceed 3000 t by 2050 [3], yet total potential capacity in 2024 was only 1100 t and primary production output was around 760 t.

Recent export restrictions have exposed the fragility of this dependency, with gallium prices rising from below 185 € kg^{-1} in 2020 to more than 460 € kg^{-1} in 2023 [4]. These developments highlight the urgency of developing alternative and secondary gallium sources to enhance supply resilience and reduce exposure to geopolitical and regulatory disruptions.

Recycling of gallium remains negligible and is largely confined to production scrap, while end-of-life recovery is virtually absent [1]. Industrial wastewater from semiconductor manufacturing represents a promising yet underexplored secondary source of gallium. In particular, GaAs wastewater from wafer manufacturing has shown gallium concentrations of up to 35 mg L^{-1} [5]. Recent material flow analyses indicate that dominant gallium losses in both wafer- and device-oriented compound semiconductor manufacturing arise from the same unit

operations, particularly substrate processing and back-grinding, implying comparable gallium-bearing wastewater streams across these sectors [6]. A material flow analysis identified wafer and device production as major loss hotspots, accounting for 19 ± 5 % and 26 ± 5 % of gallium losses within the production system, respectively, and highlights wastewater treatment as a key leverage point for reducing these losses, revealing a substantial untapped recovery potential of both economic and environmental relevance [7].

To address this gap, this study compares two separation technologies: ion exchange (IX, e.g., chelating resins) and solvent extraction (SX, e.g., organophosphorus extractants). IX is the current state-of-the-art method and is already applied on a large industrial scale to recover gallium from Bayer® liquor, an alkaline process stream generated during alumina production, owing to its high selectivity and operational stability [8]. However, while IX dominates industrial gallium recovery from alkaline process streams, solvent extraction remains chemically well suited for acidic, arsenic-bearing solutions and is frequently reported in the literature on GaAs-related aqueous process solutions. For such systems, including scrap-derived leachates and simulated process solutions representative of GaAs manufacturing, both IX and SX have been tested mainly at laboratory scale, achieving high gallium recovery yields of up to 99 % and successful arsenic removal [9, 10].

Despite this demonstrated technical potential, neither technology has been implemented at scale for fab wastewaters, largely due to dilute gallium concentrations and complex co-contaminant matrices that undermine economic viability [5]. At the same time, recent laboratory studies report high selectivity and recovery for GaAs-representative acidic solutions, and emerging downstream conversion concepts offer additional enabling options; however, their implications for full-scale process design, cost drivers, and break-even conditions remain insufficiently quantified [9, 11, 12].

This work develops a system-level process model and techno-economic analysis to translate reported separation performance into scaled process designs and quantify cost drivers, break-even conditions, and design trade-offs across IX and SX. The results provide actionable design guidance for fab integration of resource recovery and water reuse and support policymakers and commodity strategists with quantitative evidence on circular supply options that can reduce dependence on highly concentrated primary gallium supply chains.

MATERIALS AND METHODS

Framework for the case study

This work develops and applies a techno-economic analysis (TEA) to evaluate the economic viability of gallium recovery from GaAs semiconductor wastewater. This section defines the case study, system boundaries, and modelling framework used to quantitatively compare two alternative recovery routes—ion exchange and solvent extraction.

First, process models, which include mass and energy balances to track gallium and key co-contaminants, are developed for each recovery train, identifying the major pretreatment, separation, and conversion steps for the two alternative separation technologies considered. The process models calculate gallium production rates, energy and additive demands, equipment quantities and sizes, and material replacement requirements, which are subsequently used as inputs to the techno-economic analysis. Next, a system-level TEA model calculates the levelized cost of gallium (LCOGa).

Key technical and economic inputs (feed composition, operating setpoints, performance parameters, costs and lifetime assumptions) are summarized in Tables 1–2 and applied consistently across both routes to ensure comparability. The TEA reports LCOGa and its decomposition into capital, operating, replacement, and labour cost contributions.

The case study is defined using measured wastewater compositions reported for real GaAs wafer manufacturing effluents. All economic parameters (electricity price, labour cost, and industrial cost indices) are referenced to a German industrial context, reflecting a representative European semiconductor manufacturing location. Gallium and arsenic concentrations of 34.6 mg L^{-1} and 35.3 mg L^{-1}, respectively, together with an acidic pH of 3.8, were taken from semiconductor wastewater characterized by Jain et al. and are used to define the representative feed composition in this work [5].

The wastewater is assumed to be available at continuous flow rates typical for industrial semiconductor fabrication facilities, reflecting segregated process wastewaters from shared unit operations such as back-grinding and etching that occur across both wafer and device manufacturing [6]. A throughput range of 10–100 m^3 d^{-1} was selected based on reported industrial effluent volumes and is applied consistently for equipment sizing, and cost correlations [13]. Additional aggregate water quality parameters, including total dissolved solids (TDS) and total suspended solids (TSS), were adopted from GaAs semiconductor wastewater treatment studies reported by Li et al. [14]. While Table 1 reports the key parameters used to define the case study, the process modelling accounts for the full wastewater composition reported in the literature.

As previously noted, this study examines IX and SX as the two alternative separation technologies for gallium recovery from GaAs semiconductor wastewater. In this work, the IX route is modelled using a chelating ion exchange resin (Diaion CR11) [12], while the SX route

employs an organophosphorus extractant (Cyanex 272) diluted in kerosene [9], consistent with experimental studies on gallium recovery from acidic GaAs-derived process solutions.

Table 1: Key feed characteristics and operating conditions defining the GaAs semiconductor wastewater case study

Parameter	Unit	Value	Source
Volume flow	$m^3\,d^{-1}$	10–100	[13]
Temperature	°C	25	[13]
pH		3.8	[5]
Gallium	$mg\,L^{-1}$	34.6	[5]
Arsenic	$mg\,L^{-1}$	35.3	[5]
Total dissolved solids	$mg\,L^{-1}$	1500	[14]
Total suspended solids	$mg\,L^{-1}$	30	[14]

These two separation technologies are evaluated as alternative recovery routes under identical wastewater feed definitions and system boundaries, enabling a consistent techno-economic comparison. The selected wastewater dataset therefore serves as a well-documented blueprint case for system-level modelling of gallium recovery from semiconductor wastewater.

This study targets the recovery of gallium with a purity of 99.99 % (4N). This grade is the standard commercial form of commodity gallium and is subsequently refined to high-purity (6N to 7N) for semiconductor device fabrication. Using this framework, the levelized cost of 4N gallium production is compared against current market prices, which were approximately 423 € kg^{-1} in 2024 [4].

Techno-economic analysis

The objective function to be minimized is the LCOGa, defined as the total annual cost of the system TC divided by the total annual production of 4N gallium, $m_{\mathrm{Ga,ann}}$ ($kg \cdot a^{-1}$):

$$LCOGa = \frac{TC}{m_{Ga,ann}} \tag{1}$$

The total annual gallium production is a function of the wastewater feed flow rate (Q_{feed}), the initial gallium concentration (c_{Ga}), the plant availability factor (y), and the overall recovery rate (TR), which is calculated from the process models developed in this work:

$$m_{Ga,ann} = Q_{feed} \times c_{Ga} \times TR \times y \times 365\ days \tag{2}$$

The total annualized cost of the gallium recovery system is defined as:

$$TC = AF \times CapEx + OpEx + Rep + LC \tag{3}$$

where $CapEx$ represents the capital expenditure required to install the gallium recovery process train, AF is the annualization factor, $OpEx$ denotes the annual operating expenses, Rep accounts for annualized replacement costs of materials and equipment with finite lifetimes, LC represents labour costs associated with plant operation.

The formula for the annualization factor is dependent on the annual discount rate, r, and the project lifetime in years, n:

$$AF = \frac{r(1+r)^n}{(1+r)^n - 1} \tag{4}$$

Capital expenditures are calculated by summing purchased equipment costs across all process steps and applying step-specific Lang factors $f_{Lang,s}$, to convert purchased costs to installed costs (accounting for installation, piping, instrumentation, for example):

$$CapEx = \sum_s (f_{Lang,s} \times \sum C_{equip,s}) \tag{5}$$

Here, $\sum C_{\mathrm{equip},s}$ is the total purchased cost of all equipment items assigned to process step s, and is the corresponding step-specific Lang factor.

Operating expenditures comprise electricity costs, chemical/additive consumption, and operations & maintenance (O&M) costs. Electricity costs are calculated from the annual electricity demand E_{ann} and the electricity price EP, while chemical costs depend on the annual consumption $m_{k,\,\mathrm{ann}}$ and unit price AP_k of each chemical or additive k. O&M costs are approximated as a fixed fraction OM of installed $CapEx$:

$$OpEx = E_{ann} \times EP + \sum_k (m_{k,ann} \times AP_k) + OM \times CapEx \tag{6}$$

Annual replacement costs Rep account for equipment and materials with finite lifetimes. They are calculated by annualizing the purchased cost of each replaceable item $C_{\mathrm{equip},i}$ over its expected lifetime τ_i:

$$Rep = \sum_i \frac{C_{equip,i}}{\tau_i} \tag{7}$$

The annual labour cost contribution, LC, is calculated as:

$$LC = C_{FTE} \times f_{FTE} \tag{8}$$

where C_{FTE} is the annual fully loaded cost of one full-time equivalent (FTE) employee, including salary, social charges, overhead, and benefits, and f_{FTE} is the fraction of a full-time equivalent required for operation of the gallium recovery process.

Values for the key parameters used in the analysis can be found in Table 2. It is worth noting that performance figures for the RO98pHt membrane, IX resin, SX, and electrowinning (EW) are sourced from experimental studies.

Table 2: Key parameters for the techno-economic analysis and the values associated with them.

Parameter symbol	Description	Value	Source
n	Plant lifetime	20	[15]
y	Availability factor	0.95	[15]
r	Discount factor	0.08	[16]
OM	Annual rate cost of operations and maintenance	0.02	[16]
EP	Electricity price [€ kWh^{-1}]	0.1648	[17]
C_{FTE}	Annual labour cost [€ a^{-1}]	69, 164	[18]
f_{FTE}	Fraction of an FTE	0.25	[19]
f_{Lang}	Lang factor	1.4 – 4.9	[20–22]
AP_{HCL}	Cost of hydrochloric acid feedstock [€ kg^{-1}]	0.15	[23]
AP_{H2SO4}	Cost of sulfuric acid feedstock [€ kg^{-1}]	0.13	[23]
AP_{NaOH}	Cost of sodium hydroxide feedstock [€ kg^{-1}]	0.3	[24]
C_{IX}	IX resin cost [€ kg^{-1}]	138.91	[25]
τ_{IX}	IX resin lifetime [y]	5	[26]
C_{SX}	SX extractant cost [€ kg^{-1}]	10	[27]
f_{makeup}	SX make up rate [% a^{-1}]	15	[28]
SEC_{EW}	EW energy [kWh kg^{-1} Ga]	9.05	[11]
$C_{RO98pHT}$	RO98pHt membrane cost [€ pc^{-1}]	1, 088	[29]

RESULTS AND DISCUSSION

Process trains and models

Process trains for each separation technology, IX and SX, were developed to recover 4N gallium from GaAs semiconductor wastewater. As shown schematically in Figure 1, both trains share an identical pretreatment and concentration front-end and a shared downstream conversion block; they differ only in the separation step. This structure enables a controlled comparison in which route-level differences in recovery and cost are dominated by separation performance and its implications for downstream conditioning.

The process begins with filtration to reduce suspended solids and mitigate downstream fouling while preserving dissolved species, including Ga [20]. Reverse osmosis (RO) is then used as a concentrator to reduce the flow volume and increase gallium concentration, improving the feasibility of selective recovery at comparatively small, fab-relevant module throughputs. The RO design uses a commercial spiral-wound membrane configuration (RO98pHt), which has demonstrated high rejection of multivalent ions and dissolved metals in acidic, metal-bearing waters [30]. In the process model, RO is implemented as a volumetric split with component-specific partitioning of gallium (and other tracked species) between permeate and concentrate. The permeate stream is treated as a reuse or discharge stream depending on site-specific integration, while the concentrate is routed to chemical conditioning [31].

The RO concentrate is adjusted to pH 2 using hydrochloric acid to ensure a stable operating window compatible with both separation routes and to avoid premature hydroxide precipitation. This setpoint follows reported operating conditions for gallium recovery in acidic solutions and provides a consistent basis for comparing IX and SX under the same feed chemistry. [9, 12, 23]

In the IX process train, Ga^{3+} is selectively complexed on a chelating resin and subsequently recovered through an acid wash. This step yields a gallium-rich product stream, while the treated aqueous stream leaving the column constitutes a gallium-depleted waste stream. Diaion CR-11 is used as the baseline resin and it is regenerated using 0.1 M H_2SO_4, consistent with reported gallium recovery from GaAs-derived feeds and supported by manufacturer specifications and literature data [12].

In the SX process train, gallium is transferred from the aqueous phase into an organic phase consisting of Cyanex 272 diluted in kerosene via pH-dependent complex formation, and subsequently stripped into an aqueous recovery solution using 0.5 M HCl, consistent with experimental studies on simulated GaAs waste etching solutions and hydrometallurgical practice [9]. Solvent make-up is explicitly included in the model to account for extractant losses and operability constraints relevant to industrial scale-up. The IX and SX recovery streams differ in acid chemistry (sulfate versus chloride), which propagates to downstream neutralization requirements and associated reagent demand.

After separation and regeneration, the gallium-rich recovery streams are not directly suitable for electrowinning due to their low gallium concentration, acidic chemistry, and residual matrix components, which would lead to severe hydrogen evolution and poor current efficiency. To enable efficient electroreduction, gallium is therefore first precipitated as gallium hydroxide by alkalinization using sodium hydroxide, effecting a solid–liquid separation that concentrates gallium while removing the bulk aqueous matrix and associated impurities [8].

The precipitate is then subjected to alkaline leaching to generate a gallium-rich electrolyte suitable for electroreduction [11]. This step is essential because direct

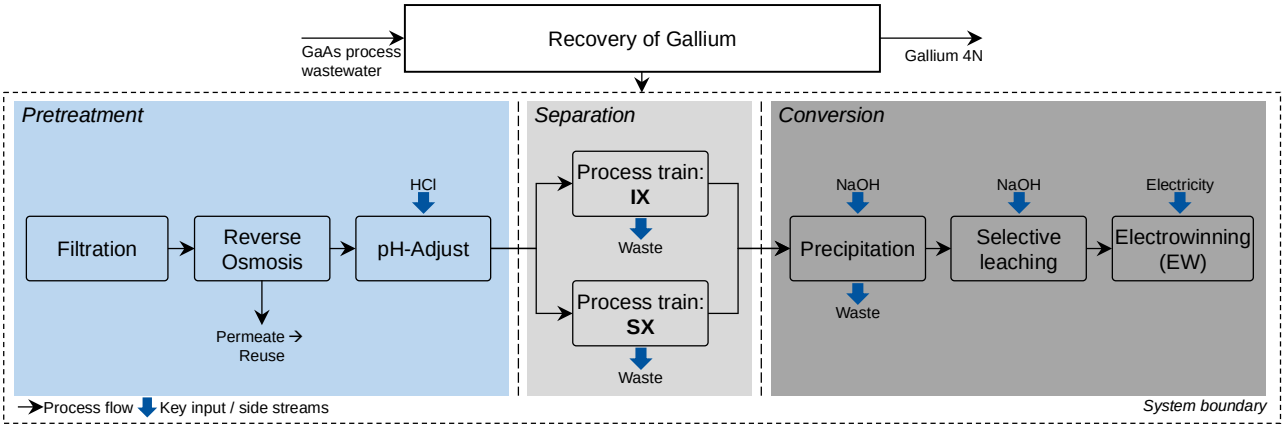

Figure 1: Schematic diagram of the processing trains developed and analyzed for recovery of 4N gallium from GaAs process wastewater, structured into pretreatment, separation, and conversion blocks. Pretreatment comprises filtration followed by reverse osmosis (RO), operated as a volumetric split to concentrate gallium in the retentate while generating a permeate stream assumed available for reuse. The RO concentrate is chemically conditioned by pH adjustment to establish an acidic matrix suitable for selective gallium separation.
Two alternative separation technologies are considered: (i) ion exchange (IX), in which gallium is selectively adsorbed on a chelating resin and recovered by acid elution; and (ii) solvent extraction (SX), in which gallium is transferred into an organic phase and subsequently stripped into an aqueous recovery stream. In both cases, the gallium-rich stream is routed to a shared conversion train comprising precipitation, alkaline leaching, and electrowinning to produce 4N gallium using a titanium cathode. The dashed box indicates the system boundary. Black arrows denote the main process flow, while blue arrows indicate key chemical and utility inputs as well as qualitative side streams.

electrowinning from acidic recovery solutions would be kinetically unfavorable and dominated by hydrogen evolution.

Final recovery is achieved by electrowinning using a titanium cathode, following the cyclone electrowinning concept reported by Xu et al [11]. In this process, forced electrolyte circulation enhances mass transfer and suppresses concentration polarization, enabling stable deposition of metallic gallium from alkaline solutions. The deposited metallic gallium is mechanically separated, washed, and vacuum-dried to remove residual electrolyte and moisture, yielding 4N gallium as the final product, which represents the standard commercial grade and serves as feedstock for subsequent high-purity upgrading. [11]

Table 3: The fraction of gallium recovered using each process train for both technologies.

Process train	Recovery fraction
IX	80.4 %
SX	62.54 %

Overall gallium recovery for each train is computed from the stepwise recovery fractions across RO, separation, precipitation, leaching, and electrowinning, with the corresponding recovery values reported in Table 3. Because both routes share the same front-end and downstream conversion, differences in overall recovery are dominated by separation performance: IX achieves a single-pass recovery of 96.9 %, whereas single-stage SX is limited to 77.3 %, resulting in overall recoveries of 80.40 % and 62.54 %, respectively. In the baseline, the SX route is modelled as a single extraction/stripping stage following Chen et al., such that higher recovery would require additional stages and increased solvent inventory, implying a cost–recovery trade-off [9]. In contrast, the IX basis follows Cheng et al. and provides higher effective single-pass recovery due to strong Ga^{3+} affinity and efficient elution, thereby reducing separation-stage losses [12].

Techno-economic performance

The process models assume an unpressurized feed, with all unit operations sized explicitly using flow-dependent design equations and operating setpoints. Given the comparatively small target throughputs considered in this study (baseline feed flow rates of 10–100 $m^3\ d^{-1}$), a bottom-up costing approach is adopted based on explicit equipment sizing and vendor or literature pricing [32]. Generic large-plant cost correlations are not applied, as they are sparse or non-representative at this scale [20].

Waste and side streams (e.g., spent regenerant, RO permeate, spent electrolyte) are quantified for mass-balance closure, disposal costs for these streams are not charged to the recovery system in the baseline TEA. This

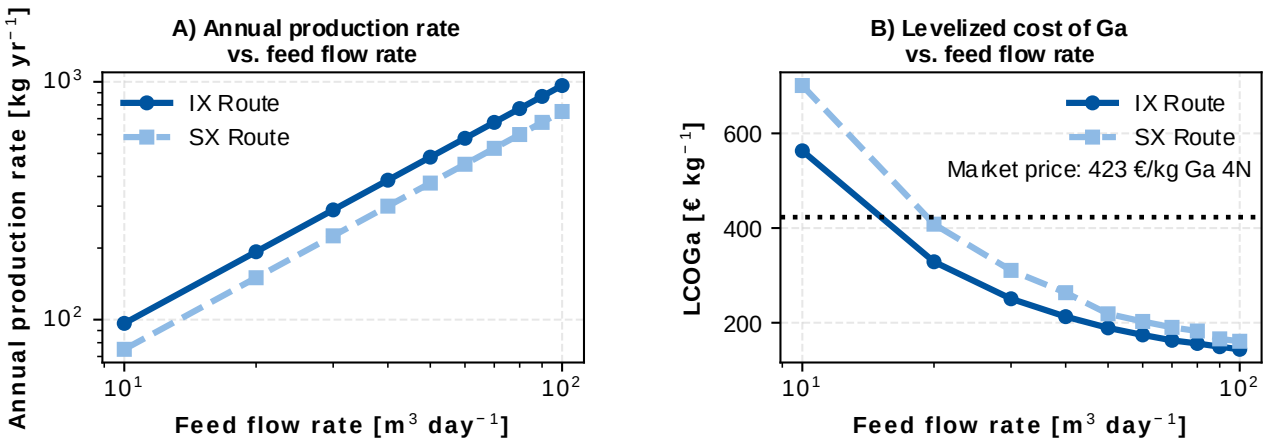

Figure 2: Baseline techno-economic results for gallium recovery from GaAs semiconductor wastewater comparing ion exchange (IX) and solvent extraction (SX) over feed flow rates of Q = 10–100 m³ d⁻¹. (A) Annual production rate (kg yr⁻¹) versus feed flow rate (log–log), showing near-linear scaling with throughput at fixed feed composition and route-dependent yield differences arising from different overall recovery. (B) Levelized cost of gallium (€/kg 4N Ga) versus feed flow rate (log-x), including a reference market price for 4N gallium (423 €/kg).

framing represents an incremental "bolt-on" recovery module at a fab site with existing wastewater treatment infrastructure; stand-alone deployment would require explicit treatment/disposal costing and may shift the economic results [33].

The production rate shown in Figure 2A scales linearly with the feed flow rate. The factors determining annual production in these designs are primarily the total recovery and conversion yields across the process train (Table 3), which depend on separation performance (recovery and selectivity), reaction yields, and operational setpoints, rather than on the absolute volume of feed processed. Accordingly, the model applies fixed recovery fractions for each route in the baseline case, resulting in proportional scaling of annual production with Q. At the baseline composition (c_{Ga} = 34.6 mg L⁻¹), the IX route produces 96 kg per year at Q = 10 m³ per day and 965 kg per year at Q = 100 m³ per day, whereas the SX route produces 75 kg per year and 750 kg per year at the same throughputs. This difference in yield is not merely a production metric: because the LCOGa is normalized per kilogram of product, a lower total recovery reduces the denominator (kilograms of gallium produced per cubic meter processed) and therefore increases the cost in euros per kilogram, even if the processing cost per cubic meter were unchanged.

The levelized cost trends are shown in Figure 2B. LCOGa decreases steeply with throughput for both routes, falling from 563 to 144 € kg⁻¹ for IX and from 701 to 161 € kg⁻¹ for SX when scaling Q from 10 to 100 m³ per day. The IX route is consistently lower in cost because it combines a higher overall recovery with a more favourable operating-cost structure at scale. This effect is particularly evident at Q = 100 m³ d⁻¹, where operating expenses account for a larger fraction of total cost for SX, approximately 43%, compared to approximately 24% for IX, as shown in Figure 3. This is consistent with the additional consumables and operational burdens associated with the SX route in the baseline model.

To assess economic viability, the LCOGa curves are compared against a representative market price for 4N gallium of 423 € kg⁻¹. Both routes reach economic viability within the evaluated throughput range. Interpolation of the base-case LCOGa curves yields break-even throughput thresholds of approximately Q* = 14.2 m³ d⁻¹ for IX and Q* = 19.3 m³ d⁻¹ for SX under baseline assumptions. These thresholds provide a practical screening criterion for deployment: below Q* the recovery module is not viable at baseline performance, whereas above Q* it becomes viable without requiring optimistic price assumptions.

Figure 3 decomposes LCOGa into CapEx, OpEx, replacement, and labour to better understand the major cost drivers of the LCOGa. At Q = 10 m³·d⁻¹, labour and annualized CapEx dominate for both routes (~30 % each), consistent with a small-plant regime where fixed and semi-fixed costs are spread over limited production. At Q = 100 m³·d⁻¹, labour becomes less significant (~12–13 %), while CapEx and OpEx dominate, explaining the strong decline in €·kg⁻¹ with scale.

In this bottom-up model, most variable consumptions (electricity, reagents, membrane area, and separation materials) scale approximately with treated volume, so returns to scale arise mainly from dilution of fixed/semi-fixed costs rather than improved specific consumptions. Parts of the installed equipment are also

discretely sized (e.g., standard pumps/ancillaries), such that selected items can cover a wide operating window and contribute similar installed costs at Q = 10 and 100 $m^3 \cdot d^{-1}$, producing stepwise CapEx behaviour. Consequently, within the two-train comparison, the most effective levers to reduce LCOGa are higher overall recovery TR (dominated by separation) and reductions in labour intensity and replacement burden.

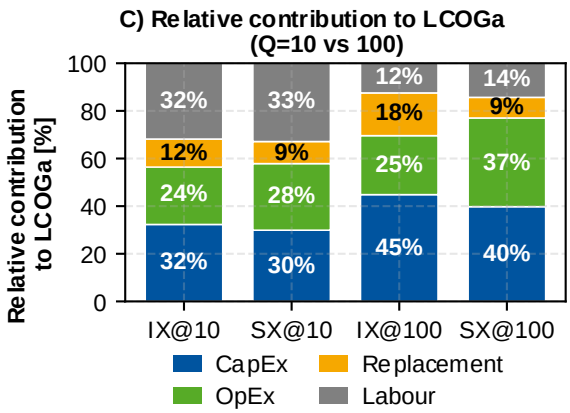

Figure 3: (C) Relative contribution of aggregated cost centres to the levelized cost of gallium (LCOGa) at two representative scales (Q = 10 and 100 m^3 d^{-1}) for both routes

Deployment implications follow directly from the controlled IX vs. SX comparison. First, throughput and feed concentration jointly determine viability, as production scales with Q and c_{Ga} at fixed recovery (Eq. 2). Second, within the two-train comparison the dominant technical lever is overall recovery TR, driven primarily by the separation step; higher SX recovery would require additional stages and solvent inventory, introducing a clear cost–recovery trade-off. Third, the strong cost decline with scale indicates two regimes: at low feed flow rates, fixed and semi-fixed costs (labour, annualised equipment) dominate, whereas at higher feed flow rates consumables and energy become increasingly important.

CONCLUSION

This study assessed the techno-economic feasibility of recovering 4N gallium from GaAs semiconductor wastewater (c_{Ga} = 34.6 $mg \cdot L^{-1}$, pH = 3.8) across a throughput range of 10–100 $m^3 \cdot d^{-1}$. Under identical system boundaries with shared pretreatment and downstream conversion, ion exchange (IX) achieves higher overall recovery than single-stage solvent extraction (SX) (80.40 % vs. 62.54 %), resulting in consistently lower LCOGa across the evaluated range (144–563 $€ \cdot kg^{-1}$ for IX vs. 161–701 $€ \cdot kg^{-1}$ for SX). This performance gap is driven primarily by the markedly higher single-pass separation efficiency of IX (96.9% vs. 77.3% for SX), which dominates recovery losses and propagates directly into total plant recovery and cost. Compared against a representative 4N gallium market price of 423 $€ \cdot kg^{-1}$, the corresponding break-even throughputs are $Q^* \approx 14.2$ $m^3 \cdot d^{-1}$ for IX and 19.3 $m^3 \cdot d^{-1}$ for SX, providing a transferable proof-of-concept blueprint for bolt-on integration of gallium recovery into semiconductor fabrication facilities.

Future work should quantify environmental performance and investigate the cost–recovery trade-offs of multi-stage SX configurations to overcome single-stage limitations. In parallel, sensitivity analyses should be used to test the robustness and transferability of these findings across different wastewater flow rates, gallium concentrations, and semiconductor wastewater chemistries, and to explore their implications for wider deployment of gallium recovery across the semiconductor industry.

DIGITAL SUPPLEMENTARY MATERIAL

The Python code and input files required to reproduce the results of this study are available at: https://github.com/Kilian-Kozerke/escape36-gallium-tea

REFERENCES

1. Løvik AN, Restrepo E, Müller DB. The global anthropogenic gallium system: determinants of demand, supply and efficiency improvements. Environ. Sci. Technol. 49:5704-5712 (2015). https://doi.org/10.1021/acs.est.5b00320
2. Jaskula B Mineral commodity summaries 2024.
3. Song H, Wang C, Sen B, Liu G. China factor: exploring the byproduct and host metal dynamics for gallium–aluminum in a global green transition. Environ. Sci. Technol. 56:2699-2708 (2022). https://doi.org/10.1021/acs.est.1c04784
4. Jaskula B Mineral commodity summaries 2025.
5. Jain R, Fan S, Kaden P, Tsushima S, Foerstendorf H, Barthen R, Lehmann F, Pollmann K. Recovery of gallium from wafer fabrication industry wastewaters by desferrioxamine b and e using reversed-phase chromatography approach. Water Research 158:203-212 (2019). https://doi.org/10.1016/j.watres.2019.04.005
6. Han Z, Liu Q, Ouyang X, Song H, Gao T, Liu Y, Wen B, Dai T. Tracking two decades of global gallium stocks and flows: a dynamic material flow analysis. Resources, Conservation and Recycling 202:107391 (2024). https://doi.org/10.1016/j.resconrec.2023.107391
7. Ueberschaar M, Otto SJ, Rotter VS. Challenges for critical raw material recovery from WEEE – the case study of gallium. Waste Management 60:534-545

(2017). https://doi.org/10.1016/j.wasman.2016.12.035
8. Luo H, Huang TY, Wu X, Zhao F. Life cycle assessment for primary gallium production at industrial-scale. Int J Life Cycle Assess 30:1545-1559 (2025). https://doi.org/10.1007/s11367-025-02492-1
9. Chen WS, Tien KW, Wang LP, Lee CH, Chung YF. Recovery of gallium from simulated gaas waste etching solutions by solvent extraction. Sustainability 12:1765 (2020). https://doi.org/10.3390/su12051765
10. Nguyen TH, Lee MS. A review on separation of gallium and indium from leach liquors by solvent extraction and ion exchange. Mineral Processing and Extractive Metallurgy Review 40:278-291 (2018). https://doi.org/10.1080/08827508.2018.1538987
11. XU ZP, LIU ZW, GUO XY, LI D, ZOU MJ, TIAN QH. Gallium recovery by cyclone electrowinning from alkaline electrolyte with titanium cathode. Transactions of Nonferrous Metals Society of China 34:1027-1036 (2024). https://doi.org/10.1016/s1003-6326(23)66451-9
12. Cheng TH, Liu CJ, Tsai TY, Shen YH. A process for the recovery of gallium from gallium arsenide scrap. Processes 7:921 (2019). https://doi.org/10.3390/pr7120921
13. Huang C, Pan JR, Lee M, Yen S. Treatment of high?level arsenic?containing wastewater by fluidized bed crystallization process. J of Chemical Tech & Biotech 82:289-294 (2007). https://doi.org/10.1002/jctb.1671
14. Li A, Luo J, Xu W, Pan T. Arsenic removal of high-arsenic wastewater from gallium arsenide semiconductor production by enhanced two-stage treatment. Desalination and Water Treatment :1-8 (2014). https://doi.org/10.1080/19443994.2014.927795
15. Kettani M, Bandelier P. Techno-economic assessment of solar energy coupling with large-scale desalination plant: the case of morocco. Desalination 494:114627 (2020). https://doi.org/10.1016/j.desal.2020.114627
16. Helal AM, Al-Malek SA, Al-Katheeri ES. Economic feasibility of alternative designs of a PV-RO desalination unit for remote areas in the united arab emirates. Desalination 221:1-16 (2008). https://doi.org/10.1016/j.desal.2007.01.064
17. SMARD | industriestrompreise. SMARD | Industriestrompreise. [Accessed: 22 Jan. 2026] https://www.smard.de/page/home/topic-article/211972/215546/industriestrompreise.
18. Chemiewirtschaft in zahlen - 2024.
19. Pryce D, Kapelan Z, Memon FA. Economic evaluation of a small wastewater treatment plant under different design and operation scenarios by life cycle costing. Development Engineering 7:100103 (2022). https://doi.org/10.1016/j.deveng.2022.100103
20. Peters MS et al. Plant design and economics for chemical engineers. (2004).
21. Huang X, Guida S, Jefferson B, Soares A. Economic evaluation of ion-exchange processes for nutrient removal and recovery from municipal wastewater. npj Clean Water 3: (2020). https://doi.org/10.1038/s41545-020-0054-x
22. Malek A, Hawlader MNA, Ho JC. Design and economics of RO seawater desalination. Desalination 105:245-261 (1996). https://doi.org/10.1016/0011-9164(96)00081-1
23. Amusat OO, Dudchenko AV, Atia AA, Bartholomew T. Cost-optimal selection of ph control for mineral scaling prevention in high recovery reverse osmosis desalination. Systems and Control Transactions 3:253-260 (2024). https://doi.org/10.69997/sct.143335
24. Mike Sodium hydroxide price index.
25. Diaion™ cr11(na), chelating resin, highly porous type, 0.5 mmole/ml on poly(styrene-divinylbenzene) 1 kg. Diaion™ CR11(Na), chelating resin, highly porous type, 0.5 mmole/ml on poly(styrene-divinylbenzene) 1 kg. [Accessed: 22 Jan. 2026] https://www.thermofisher.com/order/catalog/product/046544.A1.
26. AldexChemical-resin-service-life-guide.
27. Buy cyanex 272(83411-71-6) from luyunjia chemistry xiamen limited. Buy Cyanex 272(83411-71-6) from LUYUNJIA CHEMISTRY XIAMEN LIMITED. [Accessed: 22 Jan. 2026] https://www.echemi.com/produce/pr2503031015-cyanex-27283411-71-6.html.
28. Kaminski MD and I (United S Argonne National Lab. Evaluation of extractant-coated magnetic microparticles for the recovery of hazardous metals from waste solution.
29. Membrane & hsn code 8421 exports from denmark. Membrane & HSN Code 8421 Exports from Denmark. [Accessed: 22 Jan. 2026] https://www.volza.com/p/membrane/export/export-from-denmark/hsn-code-8421/.
30. Ambiado K, Bustos C, Schwarz A, Bórquez R. Membrane technology applied to acid mine drainage from copper mining. Water Science and Technology 75:705-715 (2016). https://doi.org/10.2166/wst.2016.556
31. Mendret J, Azais A, Favier T, Brosillon S. Urban wastewater reuse using a coupling between nanofiltration and ozonation: techno-economic

assessment. Chemical Engineering Research and Design 145:19-28 (2019). https://doi.org/10.1016/j.cherd.2019.02.034
32. Mohajerani S, Kumar A, Oni AO. A techno-economic assessment of gas-to-liquid and coal-to-liquid plants through the development of scale factors. Energy 150:681-693 (2018). https://doi.org/10.1016/j.energy.2018.03.005
33. Hess JC Chip production's ecological footprint: mapping climate and environmental impact.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Integrated solvent and process design with technoeconomic and lifecycle assessment for solvent-based recycling of end-of-life vehicle plastics

Riccardo Standish[a], Jian Yin[b], Jakob Burger[c], Mirjana Minceva[b], Hannah Mangold[d], Christian Holtze[d], Markus Schoerner[d], Bernhard von Vacano[d], Amparo Galindo[a], George Jackson[a], Claire S. Adjiman[a*]

[a] Department of Chemical Engineering, Sargent Centre for Process Systems Engineering, Imperial College London, London SW7 2AZ UK
[b] Technical University of Munich, TUM School of Life Sciences, Freising, 85354, Germany
[c] Technical University of Munich, Campus Straubing for Biotechnology and Sustainability, Straubing, 94315, Germany
[d] BASF SE, Carl-Bosch-Strasse 38, 67056 Ludwigshafen am Rhein, Germany
* Corresponding Author: c.adjiman@imperial.ac.uk.

ABSTRACT

The accumulation of automotive plastic waste poses a growing environmental threat; while recycling has the potential to address this, its use remains limited by the complexity of the materials used in vehicle components. Specifically, the presence of mixtures of polypropylene (PP), polyethylene (PE), and polyoxymethylene (POM) in the materials makes mechanical recycling challenging due to the difficulty of separation. To address the inefficiency of current end-of-life management, we present a systematic computational framework integrating computer-aided molecular and process design (CAMPD) with technoeconomic assessment (TEA) and life cycle analysis (LCA) to design a solvent-based recycling process capable of producing near-virgin quality resins. This framework involves utilizing the SAFT-γ Mie equation of state to predict thermodynamic properties and employing nonlinear programming (NLP) to perform process optimization. From an evaluation of 875 solvent candidates, we identify 72 feasible solvent combinations, amongst which cymene and cyclohexanone is the most promising solvent pair to separate polyolefins selectively and dissolve POM, enabling reduced dissolution temperatures and minimized unit capital costs. TEA and LCA indicate that the proposed process is highly competitive, as all feasible designs yield a minimum selling price (MSP) below the market price of virgin polymers and achieve global warming potentials (GWP) significantly lower than those of virgin polymer production. Sensitivity analysis confirms the robustness of the design, showing that both economic viability and environmental competitiveness relative to virgin production are maintained, even under conservative solvent loss scenarios of up to 5%. These findings suggest that solvent-based recycling offers a commercially viable pathway to meet emerging EU circularity mandates for the automotive sector.

Keywords: Solvent-based plastic recycling, Process optimization, SAFT-γ Mie, Technoeconomic analysis, Lifecycle Assessment, Solvent design

INTRODUCTION

Automotive plastic waste has emerged as a critical environmental concern in Europe. With a fleet of approximately 285.6 million vehicles, roughly 6.5 million reach their end of life annually [1]. The automotive sector is a major consumer of polymers, with more than 2.2 million tons of plastics being used in vehicles introduced to the EU market each year. Key engineering polymers in this sector include polypropylene (PP), polyethylene (PE), and polyoxymethylene (POM) [2]. Despite this significant material footprint, current end-of-life management is inefficient; less than 19% of automotive plastic waste is recycled, while the vast majority is directed to incineration or landfill [2].

Current mechanical recycling approaches often

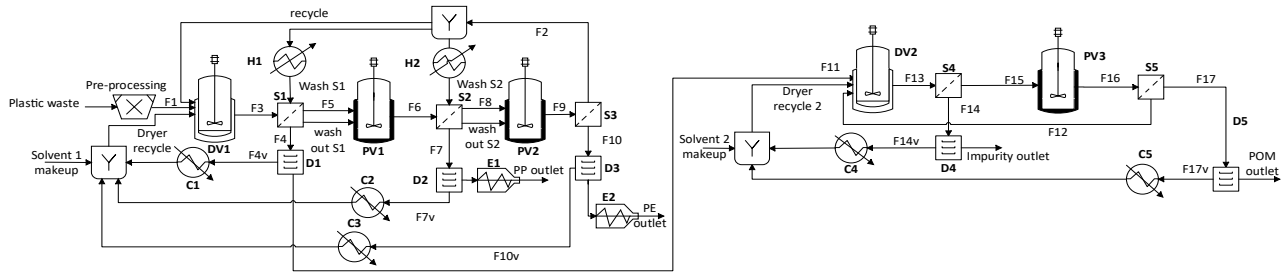

Figure 1: Process flow diagram of the plastic waste recycling process considered in the current work.

yield downgraded materials with compromised performance due to polymer chain scission, contamination from additives, and the presence of glass fibers [3]. These limitations severely restrict the reuse of recycled plastics in high-performance automotive applications. The new EU circularity rules, however, mandate at least 20% recycled plastic in new vehicle types within six years [1], creating an urgent need for technologies capable of producing recycled resins comparable in quality to virgin materials.

In this context, solvent-based recycling offers a promising solution. By selectively dissolving target polymers, this method can effectively remove contaminants and precipitate near-virgin quality materials [4]. This technology is particularly suitable for PE and PP, whose inert backbones make chemical depolymerization impractical [5]. However, the design of solvent-based polymer recycling processes is complex, requiring the identification of solvent systems that can achieve high selectivity and purity while maintaining economic viability [6].

Recent computational efforts have sought to streamline the design of these processes. Notably, Ikegwu et al. [7] developed a fast computational framework for designing solvent-based recycling systems, specifically targeting the recovery of multilayer packaging films via the solvent-targeted recovery and precipitation (STRAP) methodology. Their work demonstrated the potential of computational screening to identify solvents for distinct polymer layers. However, their analysis also highlighted the significant difficulty of separating chemically similar polyolefins; they reported finding very few viable solvent candidates capable of effectively separating PE from PP under fixed processing conditions.

This scarcity of obvious candidates underscores the need for more integrated approaches to molecular and process design. While broad screening is effective for identifying solvents with high solubility, successfully separating PE, PP, and POM requires further optimization of dissolution and precipitation operating conditions. Rigorous optimization of process temperatures and recycle rates may render a solvent inviable for the multicomponent mixture, even if it shows low or high solubility for the single solute [8].

Building on the challenges identified in the literature, we move beyond standard screening to systematically optimize the process conditions for 875 solvent combinations (one solvent for each dissolution step) through a novel computer-aided molecular and process design (CAMPD) framework specific to the ternary separation of automotive PE, PP, and POM waste.

We propose a process flowsheet comprising sequential dissolution/precipitation steps with solvent recycle. To identify the optimal design, we employ an enumeration strategy coupled with rigorous nonlinear programming (NLP). For each solvent combination, we solve a dedicated NLP problem to convergence using a model of the process that combines mass and energy balances with the predictive SAFT-γ Mie [9, 10] equation of state to evaluate the thermodynamic properties of the multicomponent mixtures. This allows us to identify specific temperature and concentration ranges where separation is thermodynamically and economically feasible; screening methods might overlook some of these designs as they lack the flexibility to adapt operating points to each solvent/solvent mixture.

We integrate the CAMPD framework with detailed technoeconomic assessment (TEA) and life cycle analysis (LCA) to provide a holistic evaluation of the trade-offs between total annualized cost (TAC) and environmental metrics, including global warming potential (GWP), cumulative energy demand (CED), and the ReCiPe and Eco-indicator 99 (EI99) endpoint indicators. Furthermore, we determine the minimum selling price (MSP) and internal rate of return (IRR) to assess the financial viability of the process. Additionally, we perform sensitivity analyses to quantify the impact of the solvent loss rate, a highly influential and uncertain parameter, on the MSP and GWP.

The paper is structured as follows: the Design Methodology section covers the problem definition, process flowsheet, CAMPD formulation, and TEA and LCA methods. The Results and Discussion section analyzes the optimization results in conjunction with economic and environmental assessments. The paper closes with final conclusions.

DESIGN METHODOLOGY

Problem definition

The process configuration, shown in Figure 1, entails

the sequential separation and recovery of 3 polymer product streams, PE, PP and POM. A 3000 ton/year plastic waste stream made up of 45 wt% PE, 45 wt% PP and 10 wt% POM is pre-processed and introduced into dissolution vessel DV1, where PE and PP are selectively dissolved in Solvent 1 to separate them from POM and insoluble impurities. The slurry is passed to filtration unit S1, where the POM and impurities are removed. A portion of the lean solvent recycle stream, *F2*, is utilized to wash the filter cake, thereby recovering the majority of the liquid PE and PP and improving polymer purity. The resulting filter cake is dried in D1 to recover Solvent 1 within the POM matrix; this solvent is subsequently condensed in C1 and recycled to DV1.

The filtrate from S1, stream *F5*, and the S1 wash outlet are fed to the precipitation vessel PV1. Here, cooling is applied to selectively precipitate the PP. The solid PP is recovered in filter S2, where the filter cake is washed with a fraction of the lean solvent, stream *F2*, to recover residual liquid PE. The cake from S2, stream *F7*, is dried in D2 for solvent recovery (condensed in C2 and recycled) and is then extruded to produce near-virgin PP.

The filtrate from S2, stream *F8,* combined with the S2 wash outlet, enters precipitation vessel PV2, where cooling is applied to precipitate PE and any remaining liquid PP. Filtration is employed in S3 to produce a PE-rich cake, *F10*. This cake is dried in D3 to recover solvent (condensed in C3 and recycled), producing a PE outlet stream that is extruded in E2 to form near-virgin PE.

For the final recovery stage, the POM-containing stream, *F11*, is fed to dissolution vessel DV2 for POM to be dissolved in Solvent 2, separating it from the remaining insoluble impurities. These impurities are removed in filter S4 and fed to dryer D4 for solvent recovery (condensed in C4 and recycled to DV2). The filtrate from S4, stream *F15*, is directed to a precipitation stage where cooling is applied to recover the solid POM. In unit S5, the solid POM is separated by filtration, and the cake is dried – enabling the recovery of Solvent 2 – before being extruded in E3 to produce near-virgin POM.

Modeling Framework and Optimization constraints

The models of each process unit, including the relevant material and energy balances, are implemented in gPROMS Process 2025.2.0. Thermodynamic properties, including phase equilibria and enthalpy calculations, are evaluated using the SAFT-γ Mie equation of state implementation within gProperties 2.10, utilizing group contribution parameters developed by Haslam et al. [11] and Paliwal et al. [12].

Process constraints

To ensure the formation of a processable solution and to prevent gel formation, upper bounds are imposed on the total polymer mass fractions in the dissolution outlet streams *F3* and *F13* [13]:

$$\sum_{pol\in\{PP,PE,POM\}} w^L_{F3,pol} \le 0.3, \tag{1}$$

$$\sum_{pol\in\{PP,PE,POM\}} w^L_{F13,pol} \le 0.3. \tag{2}$$

In DV1, the operating temperature, T_{F3}, is determined by imposing solid-liquid equilibrium (SLE) equation on PP [13]. To ensure selective dissolution, where PP dissolves while PE remains in the liquid phase and POM remains solid, we enforce specific lower and upper bounds on T_{F3}:

$$T^{\mathrm{SLE}}_{PE,F3} \le T_{F3} \le T^{\mathrm{SLE}}_{POM,F3}, \tag{3}$$

where $T^{\mathrm{SLE}}_{PE,F3}$ and $T^{\mathrm{SLE}}_{POM,F3}$ are the SLE temperatures of PE and POM at $\boldsymbol{x}^L_{F3}$, the liquid composition of stream *F3* (with $T^{\mathrm{SLE}}_{POM,F3}$ computed at $x^L_{F3,\mathrm{POM}} = 10^{-20}$).
Additionally, to guarantee that the polymer-solvent mixture forms a single stable liquid phase in DV1, a logical stability function $\tilde{S}_l(T_3, P_3, \boldsymbol{x}^L_3)$ is constrained to 1. A value of 1 indicates that a temperature-pressure flash calculation at the specified conditions detects only one stable (liquid) phase, while 0 indicates the presence of another fluid phase (a vapor or a second liquid). A similar function $\tilde{S}_l(T_{13}, P_{13}, \boldsymbol{x}^L_{13})$ is constrained to 1 to ensure the polymers and solvent form a single stable liquid phase in DV2.

In PV1, the operating temperature, T_{F6}, is also obtained by imposing SLE equation for PP. To prevent the co-precipitation of PE, a lower limit on T_{F6} is enforced:

$$T_{F6} \ge T^{\mathrm{SLE}}_{PE,F6}, \tag{4}$$

where $T^{SLE}_{PE,F6}$ is the SLE temperature of PE at $\boldsymbol{x}^L_{F6}$, the molar liquid composition of stream *F6*. In PV2, the design objective is to fully recover PE and any residual PP. Consequently, the operating temperature (T_{F9}) must be sufficiently low to ensure full precipitation but must remain above the solvent's freezing point to avoid crystallization:

$$T_{F9} \le T^{\mathrm{SLE}}_{PE,F9}; \quad T_{F9} \le T^{\mathrm{SLE}}_{PP,F9}; \quad T^{\mathrm{fus}}_{sol} \le T_{F9} \tag{5, 6, 7}$$

where $T^{\mathrm{SLE}}_{PE,F9}$ and $T^{\mathrm{SLE}}_{PP,F9}$ are computed for a trace mole fraction of 10^{-20} and T^{fus}_{sol} is the solvent temperature of fusion.

Finally, because the polymers are thermally sensitive, strict upper limits must be applied to the temperatures of dryers D1, D2, D3, and D5. These temperatures must remain below the polymer temperature of fusion (T^{fus}_{pol}) and below or equal to the solvent bubble point (T^{b}_{sol}) to prevent overpressurization:

$$T_{F4} \le T^{\mathrm{fus}}_{POM} - 5; \; T_{F4} \le T^{\mathrm{b}}_{sol}; \tag{8, 9}$$

$$T_{F7} \le T^{\mathrm{fus}}_{PP} - 5; \; T_{F7} \le T^{\mathrm{b}}_{sol}; \tag{10, 11}$$

$$T_{F10} \leq T_{PE}^{\text{fus}} - 5;\ T_{F10} \leq T_{sol}^{\text{b}} \quad (12, 13)$$

$$T_{F14} \leq T_{sol}^{\text{b}}; \quad (14)$$

$$T_{F17} \leq T_{POM}^{\text{fus}} - 5;\ T_{F17} \leq T_{sol}^{\text{b}}; \quad (15, 16)$$

where T_{F4}, T_{F7} and T_{F10} correspond to the temperatures of dryers D1, D2 and D3, respectively. The temperature of dryer D4, T_{F14}, can be above the temperature of fusion of the polymers as the filter cake from S1, *F4*, consists primarily of insoluble impurities. A safety margin of 5 K has been applied to the lower bounds.

Optimization variables and specifications

The optimization problem is constrained by strict product quality specifications, which require a minimum recovery of 95% for each polymer component (PP, PE, and POM) relative to the initial feed stream (*F1*), along with a minimum product purity of 99% by weight. The set of decision variables includes the identities of solvents 1 and 2, the weight fractions of PP in the dissolution outlet (*F3*) and precipitation outlet (*F6*), denoted as $w_{F3,PP,liquid}$ and $w_{F6,PP,liquid}$, respectively, as well as the operating temperatures of the dryers $(T_{F4}, T_{F7}, T_{F10}, T_{F17})$ and the split fractions of the pure solvent stream (*F2*) utilized for washing the filter cakes in units S1 (WF_{S1}) and S2 (WF_{S2}).

CAMPD formulation

The process model and thermodynamic constraints are integrated within a CAMPD framework. This is in principle a mixed integer problem, but it is tackled here via enumeration of the integer variables and solution of the resulting NLP problem for each combination. The solvent candidates are selected from two predefined lists, $\mathcal{L}_1$ (for Polyolefins) and $\mathcal{L}_2$ (for POM). Each solvent l_1 and l_2 is characterized by a vector of functional group numbers, $\boldsymbol{n}_{l_1}$ and $\boldsymbol{n}_{l_2}$, using the sets G_1 and G_2 of SAFT-γ Mie groups (see Table 1). The solvents selected include commonly used in solvent-based recycling processes and solvents with promising RED values according to Hansen Solubility Parameters (HSPs) [14]. Two binary variables, y_{l_1} and y_{l_2}, are introduced to denote whether the solvent candidates l_1 and l_2 are selected and are defined such that only one solvent from each list can be chosen:

$$\sum_{l_1 \in \mathcal{L}_1} y_{l_1} = 1; \qquad \sum_{l_2 \in \mathcal{L}_2} y_{l_2} = 1. \quad (17, 18)$$

Assignment constraints are used to set the values of the SAFT-γ Mie group numbers. The number of groups of type g_i in solvent l_i is given by

$$n_{g_i} = \sum\nolimits_{l_i \in \mathcal{L}_i} y_{l_i} n_{g_i,l_i}. \quad (19, 20)$$

The solution strategy is implemented via a Python script that enumerates the solvent combinations and passes them to gPROMS via gO:RUN, where the NLP is solved using NLPSQP.

The problem is formulated using the sum of the scaled TAC (economic metric), $\widehat{\text{TAC}}$, and scaled ReCiPe Endpoint indicator (environmental metric), $\widehat{\text{ReC1Pe}}$, as the objective function. The two scaled metrics are given equal weight. Normalization is performed using scalarization-based multiobjective optimization (cf. SI).

Table 1: Solvent lists and SAFT-γ Mie groups involved in the two solvent lists.

List	Solvents	SAFT-γ Mie groups
$\mathcal{L}_1$	Benzene, butane, cumene, cyclohexane, cyclohexanone, cymene, decalin, n-decane, dibutoxymethane, dimethylcyclohexane, n-dodecane, diisopropylbenzene, ethyl acetate, ethylbenzene, hexane, isopropylxylene, MEK, mesitylene, methylcyclohexane, n-octane, n-pentane, propane, toluene, n-undecane, xylene.	$G_1 = \{CH_3, CH_2, cCH, cCH_2, cCO, mO, aCH, aCCH, aCCH_3, aCCH_2, COO, CH_3CO\}$
$\mathcal{L}_2$	$\mathcal{L}_1$ + acetone, benzyl alcohol, butyl acetate, CPME, 1-decanol, ethyl isobutyrate, ethylene glycol, 1-octanol and thymol.	$G_2 = G_1 \cup \{CH_3COCH_3, aCOH, OH, CH_2OH, CH\}$

TEA and LCA methodology

The TEA analysis is performed using the methodologies established by Towler and Sinnott [15] to compute the TAC, MSP, and IRR. The TAC is calculated as

$$\text{TAC} = \text{OPEX} + \frac{[i(1+i)^n]}{[(1+i)^n - 1]} \times \text{CAPEX}; \quad (21)$$

where i represents the discount rate, set at 20% [16], and n denotes the plant life, assumed to be 20 years [17]. The operating expenditure (OPEX) is divided into variable and fixed categories. Variable costs include steam generated from natural gas, grid electricity, cooling water supplied via cooling towers, solvent makeup, and waste incineration. Fixed costs comprise pre-treatment, salaries, maintenance, benefits, overheads, land rent, tax, and insurance. Further details on the TEA methodology are provided in the SI.

The LCA is conducted on a cradle-to-gate basis, accounting for all impacts associated with production up to the factory gate. The analysis is performed using the openLCA 2.3.1 software with background data from the Ecoinvent v3.6 database [18]. The life cycle inventory (LCI) is generated using the process model developed in gPROMS, supplemented by external sources for specific inputs. Data for the GWP, EI99, and CED of solvent

Table 2: Best 5 and worst 5 Solvent Pairs (SP) ranked from lowest to highest objective value (obj=$\widehat{\text{TAC}}$+ $\widehat{\text{ReCiPe}}$.)

Label	Solvent 1 + Solvent 2	TAC/M$/y	ReCiPe/-	T_{F3}/K	T_{F13}/K	$F_{F3,sol}$ /kton/y	$F_{F13,sol}$ /kton/y
SP1	cymene+cyclohexanone	2.90	10.63	389.67	415.43	10.09	5.57
SP2	ethylbenzene+cyclohexanone	2.90	10.75	395.99	415.41	10.25	5.75
SP3	cymene+n-decane	2.91	10.73	389.96	419.48	9.68	3.85
SP4	cymene+mesitylene	2.90	10.82	389.02	427.36	11.11	5.29
SP5	cymene+n-undecane	2.91	10.75	390.18	419.90	9.39	3.91
SP68	dibutoxymethane+1-octanol	3.27	11.70	400.07	426.51	11.72	4.54
SP69	n-undecane+1-octanol	3.32	11.39	391.01	426.50	11.37	4.59
SP70	n-undecane+1-decanol	3.33	11.38	391.28	426.49	10.92	4.46
SP71	n-dodecane+1-octanol	3.35	11.42	393.45	426.52	10.77	4.33
SP72	n-dodecane+1-decanol	3.35	11.41	393.40	426.51	10.85	4.21

production is obtained from Finechem [19, 20]. The ReCiPe values for solvent production are estimated using the correlation developed by Dieterich et al. [21] between EI99 and ReCiPe endpoint indicators for solvent production. Further details on the LCA methodology are provided in the SI.

RESULTS AND DISCUSSION

Optimization results

Through the optimization framework, we systematically evaluate 875 potential solvent combinations, out of which 72 are identified as feasible candidates satisfying all thermodynamic stability and process constraints. Among these, the combination of cymene+cyclohexanone emerges as the best-performing candidate, minimizing the performance objective within the investigated search space. This configuration results in a 13.43% reduction in TAC and 6.84% lower ReCiPe compared to n-dodecane+1-decanol, the lowest-ranked design. Lower OPEX and environmental burdens are achieved due to lower operating temperatures in the dissolution vessels (T_{F3}=389.47 K and T_{F13}=415.43 K) compared to competing designs.

The top-performing candidates (SP1–SP5) all utilize the aromatic solvents cymene and ethylbenzene for the polyolefin dissolution/precipitation loop, with alkanes, cyclics, and aromatics favored in the POM separation section. Conversely, the poorest performing combinations (SP68–SP72) employ dibutoxymethane and alkanes for the PE/PP process, while relying on alcohols for POM dissolution. These inferior candidates require operating temperatures in DV2 near the polymer's fusion point (see SI), resulting in higher thermal energy demand.

Regarding solvent use, the top 5 SPs use similar amounts of solvent to the bottom 5 SPs when considering both the PE/PP and POM processes. The differences in TAC can be largely attributed to the cost of the filter units, as shown in Figure 2, which is highly dependent on solvent choice. Filter cost is influenced not only by the solvent flow rate but also by its viscosity [22]. Further details on the filter costing methodology and viscosity calculations are provided in the SI.

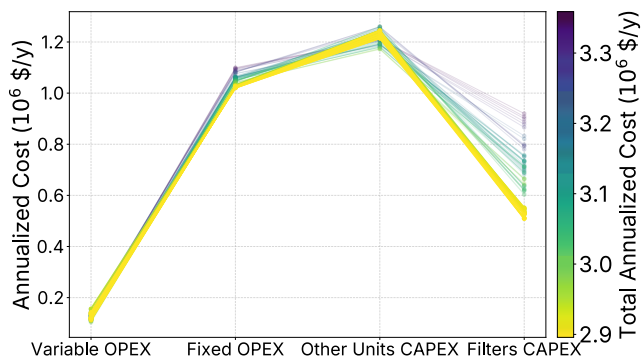

Figure 2. Parallel-coordinates plot illustrating the differences in variable and fixed OPEX, other unit, and filter CAPEX. Each line corresponds to one of the 72 SPs, colour-coded by TAC.

To evaluate the trade-offs between conflicting design criteria, we show a parallel coordinates plot (Figure 3). The analysis reveals strong consistency across the environmental metrics; solvents that perform well in GWP also tend to rank highly in CED, EI99, and ReCiPe, with only slight overlap in ranking lines observed between EI99 and ReCiPe for the poorest performers. This suggests that minimizing a single environmental indicator is generally a reliable proxy for overall sustainability in this process.

However, we observe a notable discrepancy between the economic metric (TAC) and these environmental indicators. As the plant is small-scale, we calculate that CAPEX accounts for approximately 60% of the TAC for most solvents (see SI). Consequently, if the economic metric is chosen as the sole objective, designs involving

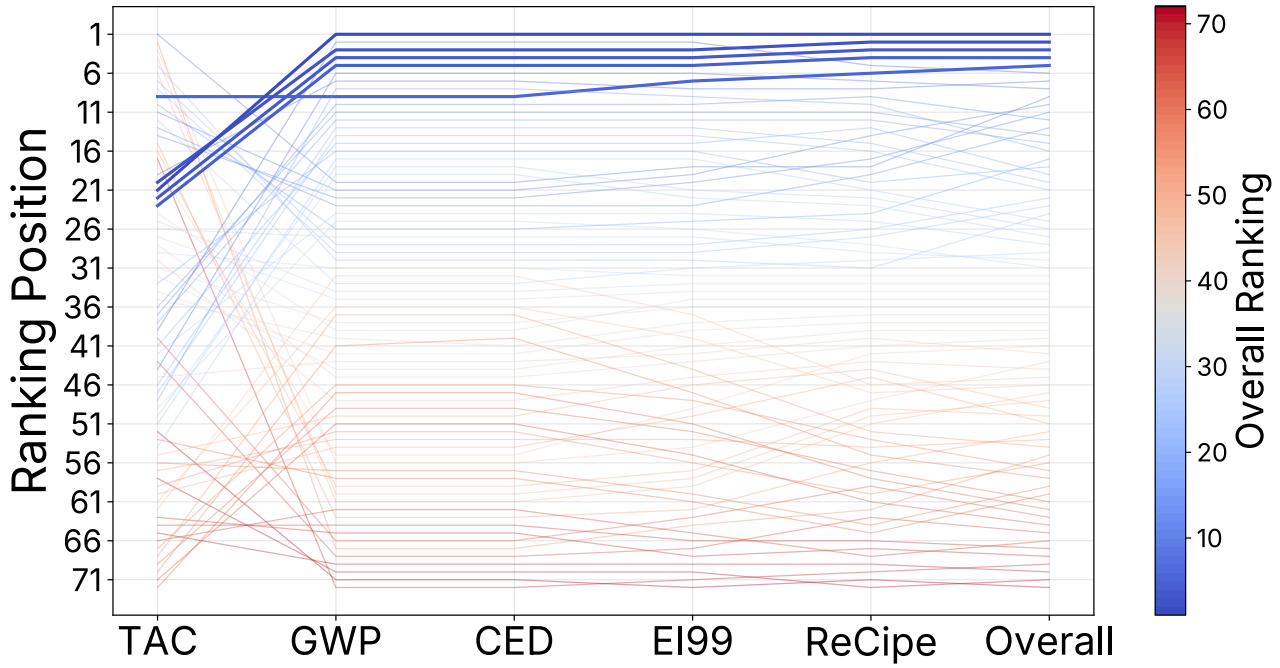

Figure 3. Parallel-coordinates plot displaying the ranking consistency of the 72 feasible solvent candidates across economic (TAC) and environmental (GWP, CED, EI99, ReCiPe) performance metrics. The vertical axes represent the ranking position (1 = best, 72 = worst). Blue lines indicate top-tier candidates, and red lines indicate lower-ranked candidates based on the objective function.

lower solvent inventories are favored to minimize vessel sizes, even if those solvents possess high specific heat capacities and are thus associated with significant energy demand. In contrast, we find that environmental performance is predominantly dependent on utility use (steam and electricity). Therefore, environmental metrics favor solvents that operate at lower temperatures, possess lower heat capacities and enthalpies of vaporization, regardless of the equipment size required. These trade-offs highlight the importance of considering both economic and environmental performance metrics to identify optimal designs.

TEA and LCA results

In Figure 4, we present the unified techno-economic and environmental assessment (MSP vs. GWP) for all 72 feasible designs. The results demonstrate that the proposed solvent-based recycling process is highly competitive against virgin polymer production. Specifically, we observe that all evaluated designs yield an MSP below the current European market price for virgin polymers. The highest MSP calculated is $1.34/kg, which remains comfortably below the market floor of $1.50/kg for virgin polypropylene [23].

The environmental analysis confirms that all 72 designs achieve a GWP lower than the CO_2 emissions associated with virgin PP and LDPE production [24] (comparative data for virgin POM was unavailable). We identify a cluster of optimal designs near the lower-left boundary of the plot, representing the set of non-dominated solutions that offer both minimal cost and minimal environmental impact. As expected, higher IRRs are observed for designs with lower MSPs. Crucially, we find that all designs exceed the standard hurdle rate of 20% required to cover the cost of capital, and the majority surpass the 18-23% threshold typically required to justify investment in novel chemical technologies [25], indicating strong potential for commercial viability. Although only one weight vector was used in the objective function, the figure shows that markedly different performance can be achieved and points to the value of multi-objective optimization to explore the space of solutions further.

We evaluate the robustness of the process by varying the solvent loss rate from the baseline value of 0.002% [26] up to 5%. In Figure 5 the impact of these losses on the breakdown of GWP contributors is shown. In this analysis, the contributions from solvent production and waste disposal are scaled to reflect the increased makeup requirements and incineration of waste solvent, while holding other contributions fixed. As can be seen, at a 5% loss rate, the solvent and disposal contributions become approximately equal; however, they remain lower than the contributions from power, heating and pre-processing. Overall, even at this worst-case loss rate

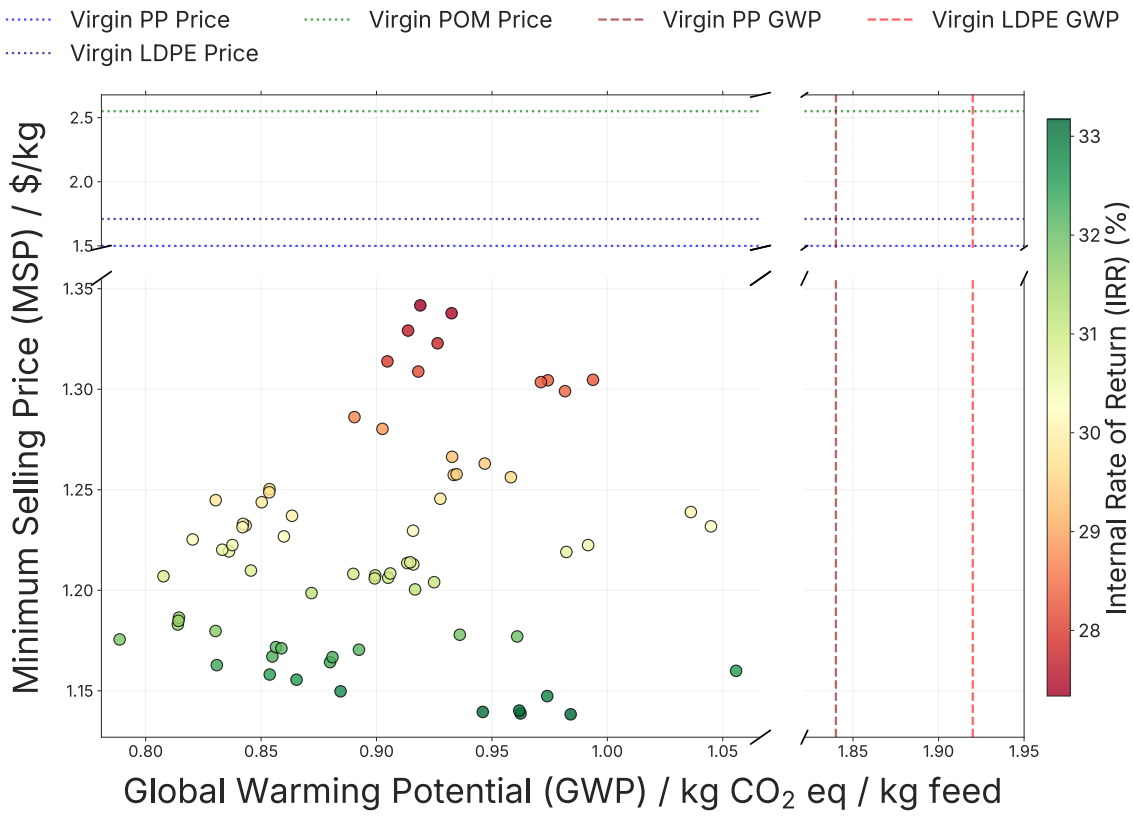

Figure 4. Scatter plot of MSP versus GWP for the 72 feasible solvent combinations. The color scale indicates the IRR.

(5%), the total system GWP remains below the carbon intensity of virgin PE and PP production.

In Figure 6, we evaluate the economic sensitivity of the process to solvent loss rates ranging from 0.002% to 5%. Variations in this parameter directly influence variable OPEX, increasing solvent makeup requirements and associated incineration costs. The data indicates that at 5% solvent loss, the MSP for the majority of design configurations exceeds the market price of virgin PP ($1.50/kg). Notably, however, solvent pair SP2 (ethylbenzene + cyclohexanone) exhibits significant economic resilience, achieving an MSP of $1.49/kg, thereby marginally retaining profitability relative to virgin PP even under maximum loss conditions.

Moreover, all evaluated designs remain substantially more cost-effective than virgin POM ($2.55/kg), while a significant proportion maintain competitiveness against virgin LDPE ($1.71/kg). Given that the weighted average market value of the polymers is approximately $1.70/kg, we confirm that several designs preserve economic viability across the full spectrum of operational efficiencies. When coupled with anticipated fiscal incentives for plastic recycling, these findings suggest the process is poised to retain robust competitiveness under diverse operational scenarios.

CONCLUSIONS

In this work, we have developed and applied an integrated CAMPD-TEA-LCA framework to design a solvent-based recycling plant for a three-component automotive plastic waste (PP/PE/POM). We have evaluated 875 solvent combinations and identified 72 feasible candidates that satisfy process and thermodynamic constraints, capable of separating chemically similar polyolefins from engineering plastics.

Our analysis highlights a notable tension between economic and environmental design drivers. We have found that while economic performance (TAC) is dominated by CAPEX (approx. 60%) and favors solvents that minimize equipment volume and solvent viscosity, environmental metrics are driven by utility consumption and favor solvents with lower heat capacities and enthalpies of vaporization. We have identified cymene + cyclohexanone as the best-performing solvent pair, effectively balancing these trade-offs by leveraging lower operating temperatures, solvent circulation rates, and transport properties to minimize both thermal energy demand and

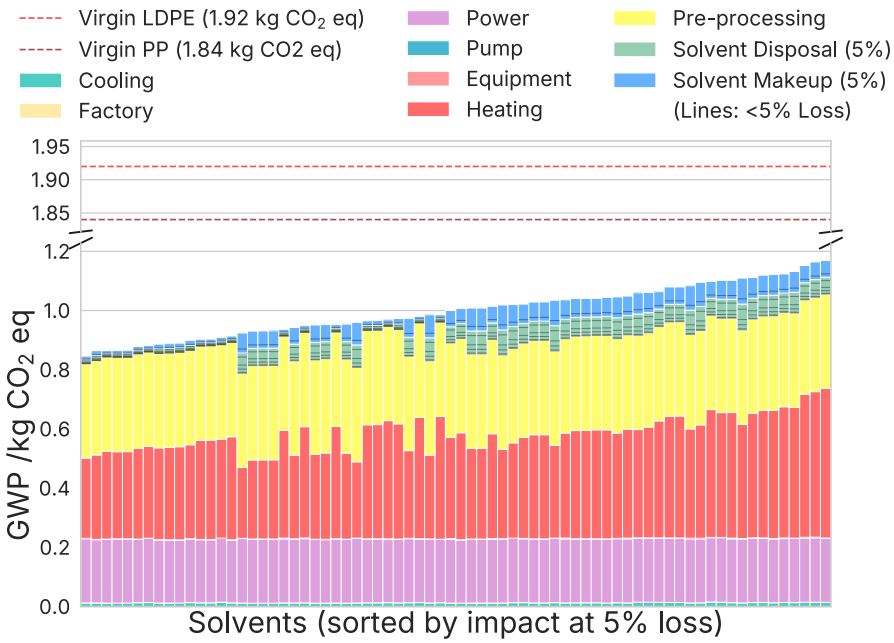

Figure 5. Sensitivity analysis of GWP to solvent loss rates ranging from 0.002% to 5%. The stacked bars illustrate the changing contributions of utilities, solvent makeup, and disposal, while the black lines indicate the total GWP values.

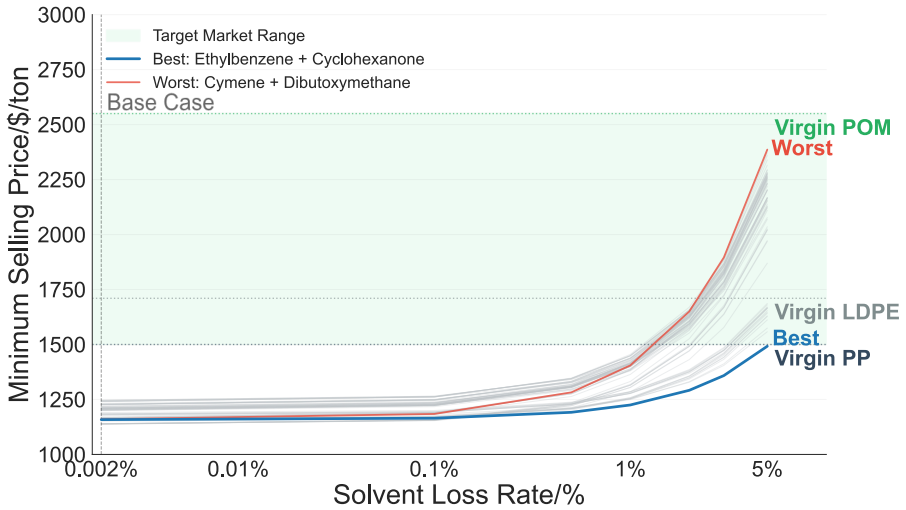

Figure 6. Sensitivity analysis of MSP to solvent loss rates. The calculated MSPs are compared against market price benchmarks for virgin PP, LDPE, and POM to assess economic competitiveness under varying operating conditions.

equipment size.

Techno-economic analysis has indicated that the process is financially attractive, with IRRs consistently exceeding the 20% hurdle rate for novel technologies. Furthermore, the environmental benefits are substantial, with carbon footprints remaining below those of virgin production even under worst-case scenarios of solvent loss (5%). We conclude that solvent-based recycling represents a robust, scalable solution for closing the loop on automotive plastics, providing a necessary technological enabler for the industry to meet the 20% recycled content mandate for new vehicles.

In future work, we plan to formulate a CAMPD problem to construct optimal solvent structures from functional groups via mixed-integer nonlinear programming, allowing for a significantly broader search space to be explored. Furthermore, we will incorporate heat integration to minimize energy demands and conduct

experiments to validate the simulated unit operations and conduct multi-objective optimization.

DIGITAL SUPPLEMENTARY MATERIAL

The supplementary material can be consulted at http://PSEcommunity.org/LAPSE:2026.0029.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge financial support from the Joint Academy of Doctoral Sciences and the EPSRC DTP grant (Grant Ref: EP/W524323/1).

REFERENCES

1. European Parliament. Press Release: New EU rules on automotive sector (2025). Link.
2. Baldassarre B, Maury T, Tazi N, Mathieux F, Sala S. Increasing plastic circularity in the automotive sector: supply chain analysis and policy options from the european union (EU). Resources, Conservation and Recycling 218:108216 (2025). https://doi.org/10.1016/j.resconrec.2025.108216
3. Schyns ZOG, Shaver MP. Mechanical recycling of packaging plastics: a review. Macromol. Rapid Commun. 42: (2020). https://doi.org/10.1002/marc.202000415
4. BASF. Solvent-based recycling pilot. Link
5. Xu Z, Sanchez-Rivera K, Granger C, Zhou P, Munguia-Lopez AC, Ikegwu UM, Avraamidou S, Zavala VM, Van Lehn RC, Bar-Ziv E, De Meester S, Huber GW. Solvent-based plastic recycling technologies. Nat Chem Eng 2:407-423 (2025). https://doi.org/10.1038/s44286-025-00247-1
6. Munguía-López AC, Sánchez-Zarco XG, Owusu-Boateng A. Process systems engineering approaches for sustainable plastics management. Ind. Eng. Chem. Res. 64:13519-13535 (2025). https://doi.org/10.1021/acs.iecr.5c00658
7. Ikegwu UM, Zhou P, Van Lehn RC, Zavala VM, Munguía-López AC. A fast computational framework for the design of solvent-based plastic recycling processes. Computers & Chemical Engineering 199:109148 (2025). https://doi.org/10.1016/j.compchemeng.2025.109148
8. Muhieddine MH, Viswanath SK, Armstrong A, Galindo A, Adjiman CS. Model-based solvent selection for the synthesis and crystallisation of pharmaceutical compounds. Chemical Engineering Science 264:118125 (2022). https://doi.org/10.1016/j.ces.2022.118125
9. Papaioannou V, Lafitte T, Avendaño C, Adjiman CS, Jackson G, Müller EA, Galindo A. Group contribution methodology based on the statistical associating fluid theory for heteronuclear molecules formed from mie segments. The Journal of Chemical Physics 140: (2014). https://doi.org/10.1063/1.4851455
10. Dufal S, Papaioannou V, Sadeqzadeh M, Pogiatzis T, Chremos A, Adjiman CS, Jackson G, Galindo A. Prediction of thermodynamic properties and phase behavior of fluids and mixtures with the saft-? mie group-contribution equation of state. J. Chem. Eng. Data 59:3272-3288 (2014). https://doi.org/10.1021/je500248h
11. Haslam AJ, et al. Expanding the applications of the SAFT-γ Mie group-contribution equation of state: Prediction of thermodynamic properties and phase behavior of mixtures. J Chem Eng Data 65:5862–5890 (2020) 10.1021/acs.jced.0c00746
12. Paliwal S, Perdomo FA, Bernet T, Alyazidi A, Haslam AJ, Jackson G, Galindo A. Extending the SAFT- ? mie group-contribution approach to model cyclic ketones: application to monoterpenoids. Molecular Physics 123: (2025). https://doi.org/10.1080/00268976.2025.2493320
13. Standish R, Yin J, Minceva M, Burger J, Galindo A, Jackson G, Adjiman CS. Optimizing the selection of solvents for the dissolution and precipitation of polyethylene. Systems and Control Transactions 4:485-490 (2025). https://doi.org/10.69997/sct.143217
14. Hansen, C. M. Hansen solubility parameters: a user's handbook; CRC press, 2007.
15. Towler G, Sinnott R. Economic evaluation of projects. Chemical Engineering Design :305-337 (2022). https://doi.org/10.1016/b978-0-12-821179-3.00009-1
16. Chaudhari US, Kulas DG, Peralta A, Hossain T, Johnson AT, Hartley DS, Handler RM, Reck BK, Thompson VS, Watkins DW, Shonnard DR. Solvent based dissolution–precipitation of waste polyethylene terephthalate: economic and environmental performance metrics. RSC Sustainability 1:1849-1860 (2023). https://doi.org/10.1039/d3su00231d
17. Naviroj, P.; Treacy, J.; Urffer, C.Chemical Recycling of Plastics by Dissolution; University of Pennsylvania, 2019. Link
18. Burhan S, Srocka M, Ciroth A, Recanati F. ecoinvent v.3.6 in openLCA. (2019)
19. Wernet G, Hellweg S, Fischer U, Papadokonstantakis S, Hungerbühler K. Molecular-structure-based models of chemical inventories using neural networks. Environ. Sci. Technol. 42:6717-6722 (2008). https://doi.org/10.1021/es7022362
20. Wernet G, Papadokonstantakis S, Hellweg S,

Hungerbühler K. Bridging data gaps in environmental assessments: modeling impacts of fine and basic chemical production. Green Chem. 11:1826 (2009). https://doi.org/10.1039/b905558d
21. Dieterich, P., 2017. Short-Cut Models for Environmental Impact Metrics of Chemical Production Master's Thesis. ETH Zurich, Switzerland.
22. Tarleton ES, Wakeman RJ. Process design for continuous separations. Solid/Liquid Separation :329-382 (2007). https://doi.org/10.1016/b978-185617421-3/50007-9
23. Businessanalytiq. https://businessanalytiq.com/
24. Uekert T, Singh A, DesVeaux JS, Ghosh T, Bhatt A, Yadav G, Afzal S, Walzberg J, Knauer KM, Nicholson SR, Beckham GT, Carpenter AC. Technical, economic, and environmental comparison of closed-loop recycling technologies for common plastics. ACS Sustainable Chem. Eng. 11:965-978 (2023). https://doi.org/10.1021/acssuschemeng.2c05497
25. McKinsey & Company. Chemical innovation: An investment for the ages. Link
26. Plastic recycling system using solvent-targeted recovery and precipitation (STRAP). US Patent Application 20240059858. (2024)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Advancing Circularity in Biopharma: Leveraging Industrial Symbiosis for Resource Efficiency

Miriam Sarkis[a], Andrea Bernardi[bc], Cleo Kontoravdi[bc], and Maria M. Papathanasiou[bc]*

[a] Texas A&M Energy Institute, Texas A&M University, College Station, 77843, TX, United States
[b] The Sargent Centre for Process Systems Engineering, Imperial College London, London, SW72AZ, United Kingdom
[c] Department of Chemical Engineering, Imperial College London, London, SW72AZ, United Kingdom
* Corresponding Author: maria.papathanasiou11@imperial.ac.uk

ABSTRACT

The biopharmaceutical sector has traditionally focused on cost-efficient process design and capacity planning to meet rising demand. Recently, sustainability pressures have increased, driving efforts to reduce the environmental footprint of manufacturing and supply chains; however, strict quality and sterilization requirements can limit the implementation of fully circular resource-use strategies. In this space, adopting an industrial-cluster systems view could unlock opportunities to improve sustainability of industrial clusters through coordinated material and energy exchange, supporting resource efficiency at cluster level and still meet sector-specific quality/sterilization requirements. In this work, we present life cycle assessment (LCA)-based comparative analyses to investigate the potential of industrial symbiosis within monoclonal antibody (mAb) manufacturing, whereby LCA process models are based on comprehensive techno-economic analyses that quantify resource inputs and waste streams. The presented results discuss impact reduction pathways including integrating renewable wind-based electricity, low-grade waste heat use, heat recovery through incineration and mechanical recycling of single-use plastics. The compound reductions of industrial symbiosis and energy source decarbonization are estimated, corresponding to 75% reductions in GWP potential, 70% in fossil resource scarcity and 60% in water use, thus highlighting the benefit of adopting a large-scale system view rather than siloed approaches.

Keywords: Life Cycle Analysis, Modelling, SimaPro, Environment, Industrial Symbiosis

INTRODUCTION

With sustainability becoming a growing priority in the pharmaceutical sector, manufacturers are pressured to quantify and mitigate footprints where possible, whilst continuing to meet product demands at global scale. Pharmaceutical manufacturing is typically water- and energy-intensive and generate significant wastewater flows, particularly in the case of multi-use (MU) stainless steel technologies. Aiming for improved flexibility, scalability and shorter set-up timelines, the sector is moving towards single-use (SU) based technologies, which in turn result in less water-related impacts, but significant increase of solid waste from the use and disposal of consumables such as plastics bioreactor, storage, mixing bags, tubes and filters.

Life cycle assessment (LCA) is an established systematic tool for comparative assessments and standardization of footprint quantification and related uncertainty, to identify more sustainable process designs. LCA[1, 2] studies comparing the MU vs SU alternatives have highlighted that footprints are largely process/product-dependent, thus requiring a case-by-case analysis to support decision-making.

LCA effectively helps pinpoint the critical resources consumed and waste flows that impact the environment through the quantification of direct and indirect emissions. Additionally, it offers the capability of assessing through what-if analyses the impact of sourcing and waste treatment decisions, given its large database of options across geographical regions, production and treatment processes. From an industrial symbiosis point of view, this translates into a systematic way to compare resource exchange pathways across industrial sectors and facilities and compute overall footprints of the integrated system.

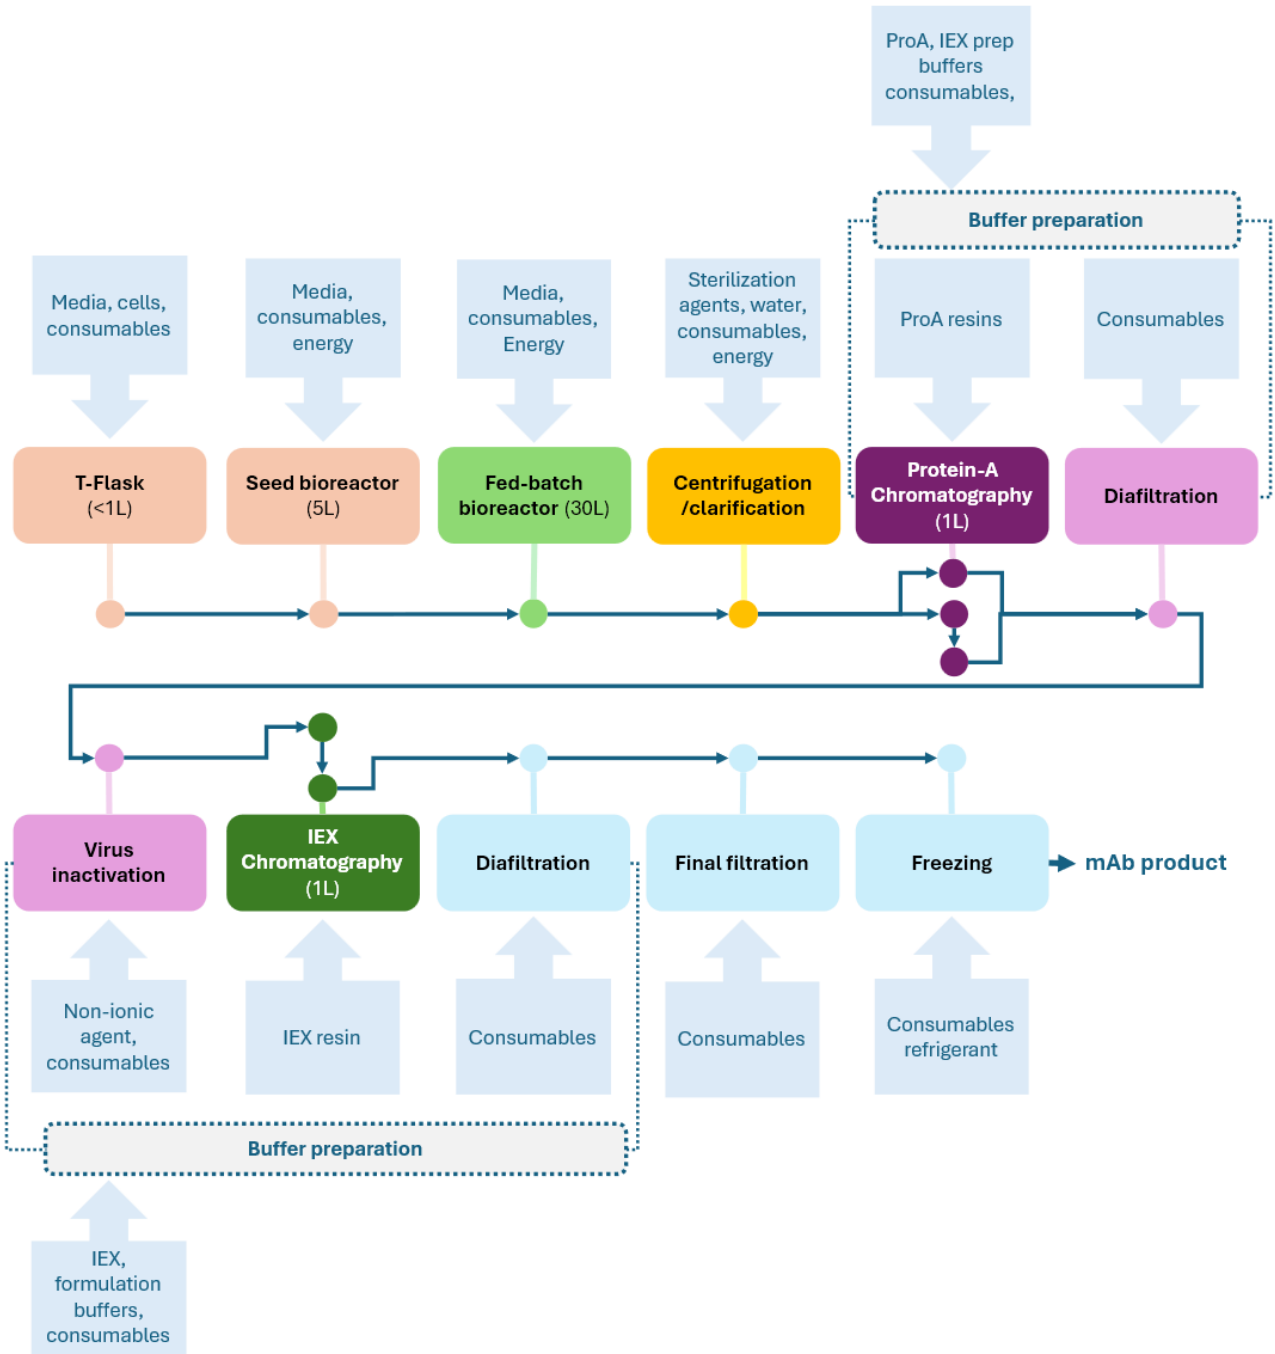

Figure 1: mAB process flowsheet illustrating equipment volumes, unit sequence and resource inputs.

Hence, in this work footprint reduction sourcing and treatment strategies are investigated, including (i) low-grade heat use in manufacturing, (ii) plastics waste incineration, (iii) incineration with energy recovery, (iv) mechanical recycling, with the aim of gaining an understanding promising pathways towards reducing impacts in biopharmaceutical manufacturing processes, specifically monoclonal antibody (mAb) production.

Effectively, the pathway-specific data presented in this work can serve as inputs within optimization-based tools to be developed for the support systematic decision-making with respect to optimal resource exchange across sectors and cost-footprint trade-off analyses in industrial planning problems.

METHODS

The framework presented in this work represents an

extension of previously published work[3], where candidate mAb process designs are generated through superstructure-based optimization and assessed via techno-economic and LCA. Here, we present further footprint analyses considering alternative sourcing and waste treatment options within the LCA analysis tool.

Process flowsheet

We consider the purity-constrained process design for mAb production[3]. Techno-economic modelling in SuperPro Designer (Intelligen), enabled the simulation of the 30-L scale platform. Figure 1 illustrates the process steps within the process sections including (i) inoculum preparation, (ii) cell culture, (iii) primary recovery, (iv) Protein A chromatography, (v) viral inactivation, (vi) IEX chromatography, (vii) filling, (viii) buffer prep and hold, which supplies buffer-based purification/separation steps. Additionally, sequential and parallel equipment units in parallel are highlighted. For reference, detailed process description and the resulting input/output materials for the LCA process models are provided in Sarkis et al. (2025)[3].

Life cycle analysis models

Environmental performance of the mAb process is assessed using LCA, in accordance with the ISO 14040/14044 standards[4].

Goal and scope

The aim of the LCA is to determine the environmental footprint of producing mAb, hence the functional unit corresponds to 1 gram of mAb produced by a 30L pilot scale process. A cradle-to-gate LCA is performed, thus including raw material production and activities up until the mAb product is sold. As a result, the disposal of single-use plastic equipment is considered as this is part of the activities within the production process. It is worth noting that plastics consumption is estimated based on the material composition of the single-use bags, thus not considering packaging.

Life cycle inventories

Materials and energy flows for the mAb production are obtained from the techno-economic models and form the foreground life cycle inventory (LCI) data, which is presented in Appendix C of Sarkis et al. (2025) [3]. and linked to the background LCI data, within EcoInvent v3.10, through Simapro[5]. The limited data availability for some resources used in the biopharmaceutical production process required a set of assumptions, including the selection of proxies for certain chemical components, including ethylene vinyl alcohol (EVOH), sodium hydrophospate and sodium dihydrophosphate. Additionally, stoichiometric modelling is integrated for TRIS-based solvents. Critically, plastics-based resources are integrated in the LCA model using mass compositions of each consumable, which therefore results in the LCA model not accounting for the production process of the consumables and the related indirect environmental impacts.

End-of-life alternatives for the plastic waste comprise landfill, incineration, and mechanical recycling. Landfill and incineration are available in Ecoinvent for polypropylene and for mixed plastic waste, while two new processes have been created for the mechanical recycling of Polyethylene and Nylon-6, considering the LCIs reported in Ruggeri et al.[6] and Tonsi et al.[7] respectively.

Life-cycle impact assessment

To illustrate the methodology in this work we focus on three midpoint indicators: global warming potential (GWP), fossil fuel scarcity, and water use. GWP and fossil fuel scarcity are estimated using the ReCiPe2016[8] life-cycle impact assessment (LCIA) methodology, while the impact of water use is characterised in terms of water scarcity using the latest consensus-based LCIA methodology AWARE[9].

Uncertainty quantification

Uncertainty in the LCA results related to the background LCI data is quantified through Monte Carlo sampling. Specifically, uncertainty of LCI within Ecoinvent is defined based on a pedigree matrix, which gives a score between 1 and 5, with 1 representing the highest quality, to five criteria: completeness, reliability, temporal correlation, geographic correlation, and technological correlation. These scores are translated into log-normal distributions of the LCIs.

In this work, the uncertainty analysis was performed using the tool included into Simapro, which uses Monte Carlo method to sample considering 1000 runs for each scenario. The uncertainty was characterised for GWP and fossil resource scarcity, while it has not been quantified for the water use as it is artificially inflated by the sampling method. Specifically, in background processes such as hydropower, where large amount of water enters and exit the system in equal amount the net water consumption is zero. However, these flows are sampled independently and may lead to large imbalances that are not physically meaningful.

Scenario analysis

In this work, the LCA models comparing different sourcing and waste treatment strategies are developed to investigate the potential for industrial symbiosis for biopharmaceutical manufacturing. Figure 2 provides a representation of the scenarios considered, which are described below:

a. **Business as usual (BAU):** plastics are purchased from their respective markets, heat requirement is met with industrial steam production, electricity is provided by the UK national grid, waste plastics are incinerated.

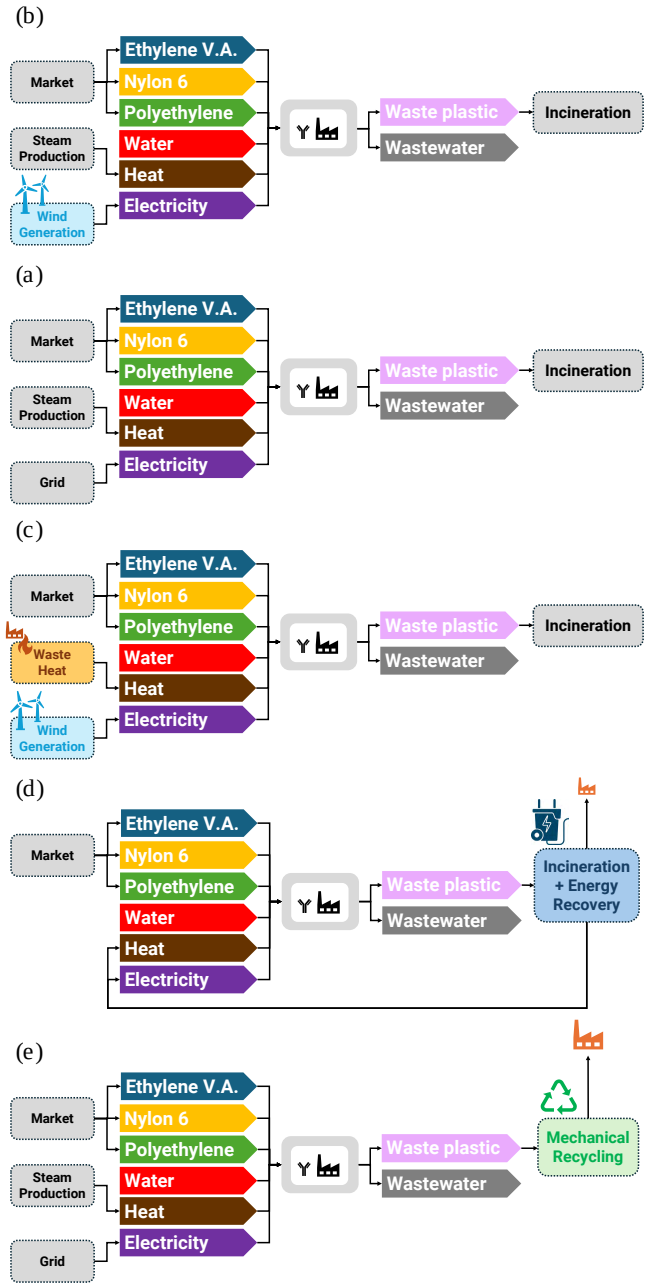

Figure 2. Footprint mitigation and symbiotic routes considered: (a) business as usual, (b) Scenario 1, (c) Scenario 2, (d) Scenario 3, (e) Scenario 4.

b. **Scenario 1**: electricity is generated through renewables, in this case an onshore wind energy
c. **Scenario 2**: electricity is provided by onshore wind energy and heat is derived from low-grade waste heat sources.
d. **Scenario 3:** waste plastics are incinerated in an energy from waste (EfW) plant, producing heat and electricity, either sold to other plants or re-used within the mAb manufacturing process.
e. **Scenario 4:** mechanical recycling of Ethylene and Nylon, producing lower grade plastics that can purchased by other industries.

RESULTS AND DISCUSSION

Figure 3 shows the environmental impacts associated with producing 1 g of mAb in a 30 L pilot-scale process. Global warming potential (GWP), fossil resource scarcity, and water use are presented in Figures 3a, 3b, and 3c, respectively. Stacked bar charts highlight the main environmental hotspots, while net impacts are indicated by white squares. Uncertainty bars for GWP and fossil resource scarcity represent the 5% and 95% percentiles obtained from Monte Carlo simulations. In the business-as-usual case (Figure 3a) the environmental hotspots for GWP correspond to plastics use (~50%), comprising EVA, Nylon and Polyethylene, the use of electricity and heat (~25%) and the end-of-life incineration of the plastics (~20%). Thus, switching to wind-based energy as captured in Scenario 1 only provides reductions up to (~10%) of the overall GWP footprint, whilst plastics-related impacts remain significant. This is even more evident with respect to fossil fuel scarcity, where in the BAU case plastics-disposal is not a contributor to resource scarcity itself, whilst plastics use is even more critical (66%). Here, switching to wind-based energy reduces the impact by 11% (Figure 3b).

To fully decarbonize the energy requirements of the plant, Scenario 2 also introduces the switch to waste-heat sources from other industries within the industrial cluster, thus effectively corresponding to a symbiotic exchange. In this scenario, the impact from heat generation is removed, resulting in GWP reductions of ~22% (Figure 3a) and fossil resource scarcity of 23% (Figure 3b). Whilst energy requirements are tackled through the wind-energy and low-grade heat exchange proposed in Scenario 2, waste plastics account for the remainder of GWP and the fossil fuel scarcity impacts.

Hence, in Scenario 3 the option of recovering energy from the incineration process of plastics is investigated. This corresponds to the case where the energy recovered could be used (i) internally to feed into the heating and electricity requirements of the process or (ii) sold to industrial partners as part of the symbiotic system. As can be seen in Figure 3a, the resulting reduction is displayed in terms of credits gained (EfW – Heat, EfW – Electricity) and corresponds to nearly 15% of the overall GWP. It is worth mentioning that these credits are calculated based on current state of the UK grid. If grids become greener, the share of renewables increase, the credits would decrease. Hence, this translates in a long-term uncertainty in the benefits of Scenario 3, as these credits are highly dependent on the carbon intensity of the grid and will likely decrease with time. The case for fossil fuel scarcity is similar (Figure 3b), with overall impact reduction (15%) achieved through credits gained for

avoided energy production, quantified through EfW – Heat and Electricity.

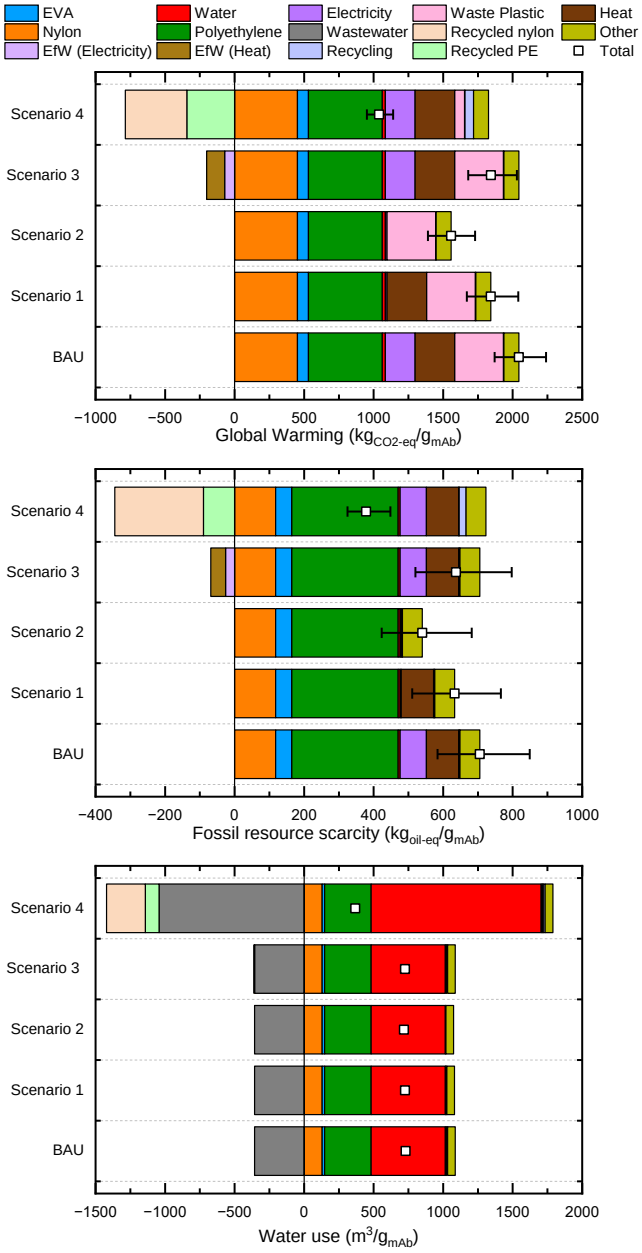

Figure 3. (a) Global warming potential, (b) Fossil fuel scarcity, (c) water depletion for the production of 1 g_{mAb}. Error bars represent the 5% and 95% percentile.

Finally, Scenario 4 considers mechanical recycling of the consumables used. The process includes shredding, water separation, washing, extrusion, and pellet production, and results in downcycled plastic material that does not meet the quality standards required for biopharmaceutical manufacturing. Therefore, the recovered plastic is assumed to be sold to other industries, generating credits for avoided virgin plastic production within the symbiotic system.

Impacts associated with mechanical recycling—mainly driven by electricity, heat, and water requirements—are shown in lavender in Figure 3. As shown in Figure 3a and 3b, the additional contributions to overall GWP and fossil resource scarcity are small compared to the electricity and heat impacts of the mAb production process. In contrast, credits from avoided plastic production are substantial, leading to ~52% reductions in overall GWP and ~47% reductions in overall fossil resource scarcity.

In terms of water use, mechanical recycling requires a large amount of water, nearly doubling direct water demand. However, after wastewater treatment, most of this water is released back to the environment. In the base case, indirect water demand associated with plastic production accounts for ~40% of total water demand. As a result, credits from avoided plastic production led to an overall ~50% reduction in net water use within the symbiotic system.

The results suggest that mechanical recycling is more promising than energy from waste in a symbiotic system. Additionally, combining mechanical recycling with wind electricity and waste heat utilization, could result in ~75% reduction in GWP and fossil fuel scarcity, and ~60% reduction in water use compared to the base case. However, it is important to note that the feasibility and efficiency of mechanical recycling depend on if and how well plastic consumables can be separated into monomaterial streams. In this work we are considering a 20 and 30% material loss for Polyethylene and Nylon[6, 7], but future work should aim at quantifying consumables-specific efficiency.

Chemical recycling techniques, such as pyrolysis, gasification, plasma-based decomposition could produce virgin-like materials, thus reducing the reliance of the biopharmaceutical industry on fossil-based products, and process unsorted mixed plastic waste that could not be mechanically recycled[10]. Although promising, low TRL level and high dependency on feedstock composition require consumable-specific modelling to develop even a preliminary LCA model of chemical recycling alternative processes and will be addressed in future research effort (Figure 4).

Overall, for the scenarios considered, this work demonstrates the benefits of adopting proactive strategies for footprint mitigation, primarily through industrial symbiosis to improve integrated system sustainability rather than adopting a siloed approach.

This motivates future research avenues in studying feasible resource exchanges through optimization-based system-system integration, with challenges emerging from potentially asynchronous operational cycles, resource transport and storage constraints.

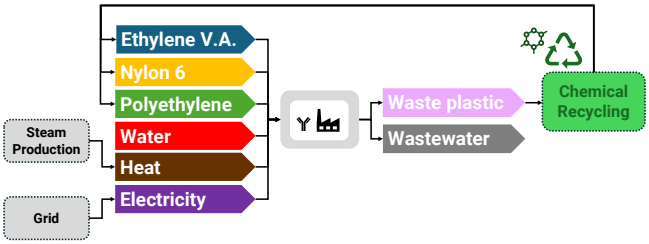

Figure 4. Footprint mitigation route through chemical recycling.

CONCLUSION

Industrial symbiosis offers opportunities to improve sustainability of industrial clusters through collective efforts and resource exchange. Through the systematic analysis presented in this study, its potential within biopharmaceutical industry has been explored, with a focus on a case study from mAb manufacturing relying on single-use and multi-use equipment. These result in emissions from both plastics consumption and disposal and energy and water use. To mitigate these, life cycle assessment models are presented to quantify footprint breakdown and uncertainty in footprint estimation and compare alternative supplier and waste treatment options representing industrial symbiosis solutions. The findings demonstrate how amongst the scenario considered, mechanical recycling is a promising option as it directly tackles the impact from waste plastics disposal, through instead the recovery of lower-grade plastics flows thus resulting in avoided emissions from low-grade plastics production elsewhere. Results further demonstrate that a combination of plastics recycling, low-grade heat supply and renewable electricity greatly reduces overall impacts, suggesting that adopting industrial cluster-wide system approach could greatly reduce footprints through smart resource exchanges.

REFERENCES

1. Pietrzykowski M, Flanagan W, Pizzi V, Brown A, Sinclair A, Monge M. An environmental life cycle assessment comparison of single-use and conventional process technology for the production of monoclonal antibodies. Journal of Cleaner Production 41:150-162 (2013). https://doi.org/10.1016/j.jclepro.2012.09.048
2. Budzinski K, Constable D, D'Aquila D, Smith P, Madabhushi SR, Whiting A, Costelloe T, Collins M. Streamlined life cycle assessment of single use technologies in biopharmaceutical manufacture. New Biotechnology 68:28-36 (2022). https://doi.org/10.1016/j.nbt.2022.01.002
3. Sarkis M, Monteiro M, Bernardi A, Chiplunkar R, Kontoravdi C, Papathanasiou MM. Sustainable by design: a first attempt on bioprocessing. Computers & Chemical Engineering 205:109454 (2026). https://doi.org/10.1016/j.compchemeng.2025.109454
4. ISO, ISO 14040:2006(En), Environmental Management — Life Cycle Assessment — Principles and Framework: Technical report, *International Organisation for Standardization* (2006)
5. Wernet G, Bauer C, Steubing B, Reinhard J, Moreno-Ruiz E, Weidema B. The ecoinvent database version 3 (part I): overview and methodology. Int J Life Cycle Assess 21:1218-1230 (2016). https://doi.org/10.1007/s11367-016-1087-8
6. Ruggeri M, Savastano M, Rizzi A, Veronesi A, Vinci G. Life cycle assessment of post-industrial plastics mechanical recycling. a case study from italy. Science of The Total Environment 998:180246 (2025). https://doi.org/10.1016/j.scitotenv.2025.180246
7. Tonsi G, Caccia M, Maesani C, Simon Ostan P, Ortenzi MA, Alini S, Pirola C. Polyamide recycling by selective dissolution approach: life cycle assessment study and environmental impacts comparison with different recycling technologies. Polymer Engineering & Sci 65:765-782 (2025). https://doi.org/10.1002/pen.27041
8. Huijbregts MAJ, Steinmann ZJN, Elshout PMF, Stam G, Verones F, Vieira M, Zijp M, Hollander A, van Zelm R. Recipe2016: a harmonised life cycle impact assessment method at midpoint and endpoint level. Int J Life Cycle Assess 22:138-147 (2016). https://doi.org/10.1007/s11367-016-1246-y
9. Boulay AM, Bare J, Benini L, Berger M, Lathuillière MJ, Manzardo A, Margini M, Motoshita M, Núñez M, Pastor AV, Ridoutt B, Oki T, Worbe S, Pfister S, The WULCA consensus characterization model for water scarcity footprints: assessing impacts of water consumption based on available water remaining (AWARE), *Int. J. Life Cycle Assess.*, 23: 368-378 (2018) https://doi.org/10.1007/s1137-017-1333-8
10. Schade A, Melzer M, Zimmermann S, Schwarz T, Stoewe K, Kuhn H. Plastic waste recycling?a chemical recycling perspective. ACS Sustainable Chem. Eng. 12:12270-12288 (2024). https://doi.org/10.1021/acssuschemeng.4c02551

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Beyond Decarbonization: Quantifying Circularity in Energy System Planning

Javiera Vergara-Zambrano[a], Styliani Avraamidou[a]*

[a] University of Wisconsin-Madison, Department of Chemical and Biological Engineering, Madison, Wisconsin, USA
* Corresponding Author: avraamidou@wisc.edu

ABSTRACT

While the transition from traditional energy sources to renewable energy is necessary to reduce greenhouse gas (GHG) emissions, it introduces new challenges related to material use, both in quantity and type, potentially leading to resource scarcity, biodiversity loss, and waste accumulation. Therefore, incorporating circular economy (CE) principles into the design and planning of energy systems becomes essential. Despite the growing recognition of circularity, current assessments in energy systems focus on economic performance and GHG emissions. In this work, we propose a metric for quantifying circularity of energy systems based on the CE assessment framework MICRON, addressing the gap between CE metrics and energy systems planning. The framework is adapted to energy systems by accounting for the specific characteristics of energy technologies and by incorporating metrics associated with critical material use, scarcity, and durability. Its applicability is demonstrated through a case study of energy system planning at the University of Wisconsin-Madison, considering a grid-connected system with solar, wind, and lithium-ion battery technologies. Results show that wind-only portfolios achieve higher overall circularity scores than solar-only and hybrid systems, reflecting the higher efficiency and availability of wind energy. Hybrid systems exhibit higher durability and more efficient material use by avoiding system oversizing. Regarding decarbonization strategies, reducing grid reliance and associated emissions does not necessarily improve circularity, as energy storage is required to ensure reliability. Storage systems increase material demand, the share of critical materials, and replacement frequency. Finally, a sensitivity analysis was performed, highlighting that end-of-life recovery is a key factor influencing circularity.

Keywords: Renewable and Sustainable Energy, Energy Systems, Circular Economy, Energy Planning

INTRODUCTION

The increasing energy demand, along with the environmental impacts of traditional energy sources, highlight the need for an energy transition. While moving from fossil fuels to renewable energy is essential to reduce greenhouse gas (GHG) emissions, focusing solely on emissions can be insufficient and overlook other environmental sustainability aspects. A sustainable energy transition must simultaneously consider, amongst others, resource efficiency, material use, land footprint, waste generation, energy security, and socioeconomic impacts, in addition to GHG emissions [1]. However, current energy transition policies and decarbonization strategies mostly focus on the closure of fossil fuel power plants and the deployment of new infrastructure, often without accounting for material demand, the increased reliance on critical materials, and the associated waste generation [2].

Clean energy technologies, including solar photovoltaics (PV), wind turbines, and power storage systems, require significant mineral resources and introduce new challenges related to resource scarcity, biodiversity loss, land use changes, and waste generation at the end of their lifetime, a concern that is amplified by their relatively short lifetimes compared to fossil fuel plants [3, 4]. For instance, a 100 MW wind power plant requires nearly 67, 850 tons of materials, including composites and polymers that are difficult to recycle [5]. Similarly, the projected 1.75 TW of installed PV capacity in the United States by

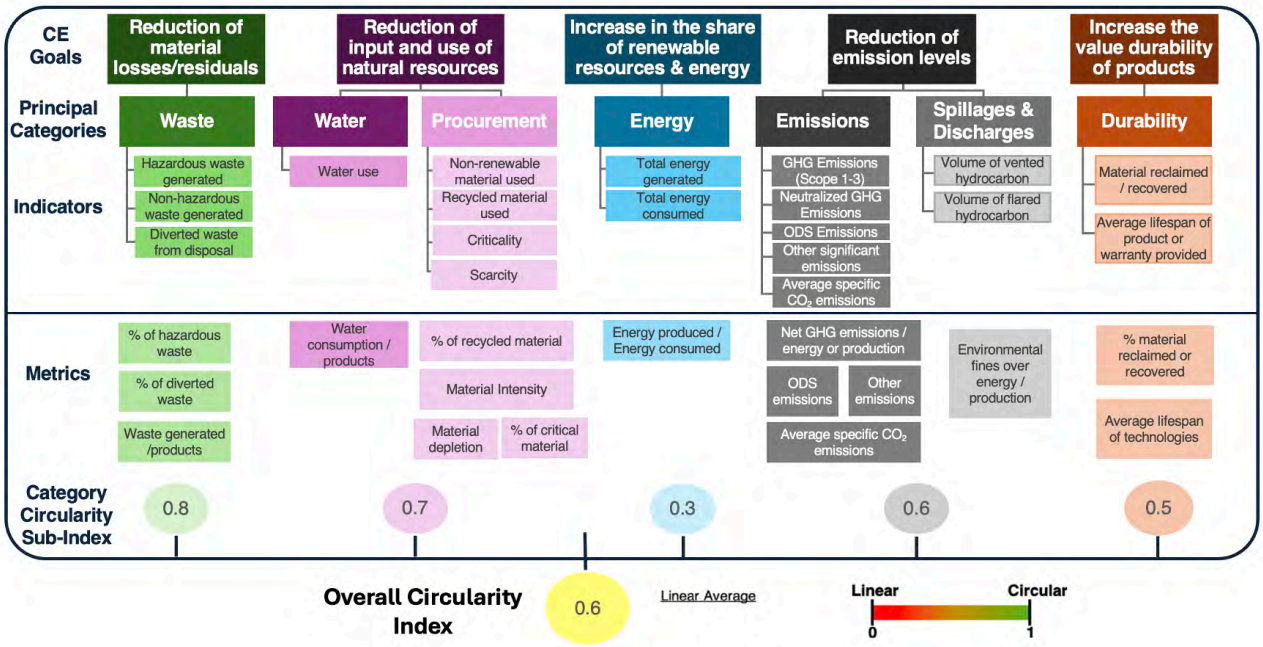

Figure 1. Schematic of MICRON CE assestment framework for energy systems, with arbitrary values on category-based circularity sub-indices.

2050 will demand nearly 97 million tons of virgin materials and generate over 8 million tons of lifecycle waste [6]. Therefore, integrating circular economy (CE) principles into the design and planning of energy systems becomes essential to reduce material demand and carbon intensity, and address end-of-life (EoL) management challenges. Without a circular approach, the energy transition risks replacing carbon emissions challenges with new environmental burdens, such as material scarcity and waste accumulation [7].

Circular Economy Metrics for Energy Systems

Progress toward CE can be measured using indicators defined at different scales, including product (micro), system (meso), and economy-wide (macro) levels. The Material Circularity Indicator (MCI) [8] evaluates material inputs and outputs at the product level, while the circularity index [9] accounts for recovered EoL material and its energy requirements at a system level. The global circularity metric [10] is an economy-wide indicator that measures the percentage of materials reused or cycled back into the global economy relative to the total material inputs. The Circular Economic Value (CEV) [11] is a system-level indicator addressing the reduction of virgin materials and waste generation. Material Flow Analysis (MFA) [12] provides a flow-based system-level framework to track material stocks and flows across processes and sectors, and recycling rates [13] are process- or material-specific metrics commonly used to evaluate EoL management performance.

When applied to energy systems, these CE metrics exhibit several limitations. Most indicators focus on mass flows and do not capture broader system interactions or the dynamic characteristics of energy infrastructures. Energy systems are often treated as inputs or waste streams rather than as systems requiring dedicated circularity metrics. Common energy-related indicators include energy productivity and energy consumption [14]. Moreover, despite increasing recognition of the need for circularity, current energy systems assessments focus on economic performance and GHG emissions reduction. For renewable energy technologies, operational emissions are negligible, and the sustainability concerns arise from the upstream and downstream stages, i.e., material extraction, and EoL management [15].

In response to these shortcomings, recent studies have introduced CE metrics tailored to specific energy sectors. Serrano-Arevalo et al. [16] integrated three CE indicators into an energy planning optimization problem. The metrics proposed were relative material impact, absolute material impact, and energy self-sufficiency [14], accounting for the material requirements and analyzing the producer-consumer relationship. The results showed that incorporating CE indicators led to significant changes in optimal system configurations, yielding additional reductions of 12% in CO_2-eq emissions compared to approaches based solely on emissions minimization. These findings demonstrate that circularity-oriented metrics can guide the selection of technology

combinations that reduce non-renewable resource use and environmental burdens. Similarly, Xin et al. [15] introduced a framework to evaluate the circularity of EV battery systems, including 15 indicators across three categories: material flow, end-of-life, and energy. The findings show that the material outflow significantly influences the circularity score, highlighting the need to improve collection and recycling rates of equipment at end-of-life.

Despite these advances, to our knowledge, no integrative assessment exists to holistically evaluate circularity performance across energy infrastructures. A circularity metric for energy systems should evaluate CE goals across all circularity dimensions, accounting for material scarcity and criticality, as well as EoL management, including downstream effects [16]. To address this gap, this study proposes an integrated metric to assess and quantify CE in energy systems based on the MICRON CE framework [17], accounting for the specific characteristics of energy technologies.

METHODS

The MICRON CE framework

The MICRON (MIcro CirculaR ecOnomy iNdex) framework, developed by Baratsas et.al. [17], is a quantitative CE assessment tool based on the Global Reporting Initiative (GRI) standards at a micro-level. The framework defines circularity scores from 0 to 1, where 0 represents a fully linear system and 1 a perfect circular system. MICRON was designed around five CE goals: i) reduction of material losses, ii) reduction of input and use of natural resources, iii) increase in the share of renewable resources, iv) reduction of emissions and pollutants, and v) increase the durability of products. Based on these goals, seven principal categories (waste, water, procurement, energy, emissions, spillages, and durability) are defined, each associated with assessment indicators and quantitative metrics. These metrics are used to derive the Category-based Circularity Index. The linear average of the category-based indices constitutes the Overall Circularity Index. MICRON was originally developed for the automotive, energy and utilities, manufacturing and service sector, and has also been applied to the recycling of food packaging and transportation fuel systems [18-19].

Adapting MICRON for energy systems

In this work, we have adapted the MICRON framework for its application to energy systems, accounting for the specific characteristics of energy technologies. Figure 1 shows the schematic of the MICRON framework for energy systems, indicating the principal categories and their associated metrics. The main contributions of this work are: i) the addition of three indicators under the procurement principal category to capture material-related challenges of renewable energy, ii) the adaptation of the durability metric to account for systems that have technologies with different lifetimes and replacement frequencies, and iii) the demonstration of the framework's applicability through a case study of energy system planning at the University of Wisconsin-Madison. The complete set of metrics considered is included in Table 1.

Table 1: CE metrics for energy systems.

Principal Category	Metric
Waste	**1a**. Waste Generated [ton]/Energy Produced [MWh] **1b.** % Hazardous Waste over Total Waste Generated **1c.** % Diverted Waste over Total Waste Generated
Water	**2a.** Water use [L]/Energy Produced [MWh]
Procurement	**3a.** % Recycled Material Used **3b.** Material intensity [ton/MWh] **3c.** % Critical Material Used **3d.** Abiotic Depletion Potential [Sb-eq]
Energy	**4a.** % Renewable Energy over Total Energy Produced **4b.** Energy Return of Investment (EROI) [-]
GHG Emissions	**5a.** Net Total Emissions [kg CO2-eq]/Energy Produced [MWh] **5b.** Ozone Depleting Substances (ODS) [metric tons CFC-11]/Energy Produced [MWh] **5c.** NOX, SOX Emissions [ton Nox+SOx]/Energy Produced [MWh]
Durability	**6a**. % Material to be Reclaimed/Recover **6b.** Average Lifetime of the System [years] **6c.** Average Remaining Lifetime [years]

Material-related procurement metrics

The deployment of new infrastructure implies an intense use of resources, both in terms of quantity and type. We have included material intensity, criticality, and scarcity as metrics in the procurement category to account for this characteristic.

1) Material intensity: It is defined as the amount of material required by the system relative to the energy produced.

$$SMI = \frac{M_{syst}^{tot}}{E_{syst}^{tot}} \quad (1)$$

Figure 2. Schematic of the CE assestment for energy systems using MICRON.

Where M_{syst}^{tot} is the mass of material required by the system and E_{syst}^{tot} is the energy produced by the installed technologies.

2) Criticality: Captures the percentage of critical material used relative to the system's total material requirements.

$$CR = \frac{M_{syst}^{crit}}{M_{syst}^{tot}} = \frac{\sum_{y\in Y}\sum_{t\in T}\sum_{i\in I_t} Exp_{t,y}\cdot m_{t,i}\cdot x_i}{\sum_{y\in Y}\sum_{t\in T}\sum_{i\in I_t} Exp_{t,y}\cdot m_{t,i}} \quad (2)$$

where M_{syst}^{crit} is the amount of critical material required by the system, $Exp_{t,y}$ is the capacity of technology t installed in year y, $m_{t,i}$ is the mass of material i per unit of technology t and x_i is a binary parameter equal to 1 if the material i is classified as critical and 0 otherwise. The classification of critical materials follows the assessments developed by the U.S. Department of Energy (DOE) and the European Union [16-17].

3) Material Scarcity: It is evaluated using the abiotic depletion potential (ADP) factor, which measures the reduction of finite, non-living natural resources due to human extraction. The ADP of a resource is defined as the ratio of its annual production to the square of its ultimate reserve, normalized by the same ratio for a reference resource [18].

$$SR = \sum_{y\in Y}\sum_{t\in T}\sum_{i\in I} Exp_{t,y}\cdot m_{t,i}\cdot ADP_i^{\prime} \quad (3)$$

where $ADP_i^{\prime}$ is the abiotic depletion potential of resource/material i.

Durability metrics

For energy systems, durability must be addressed at the portfolio level rather than by comparing single technologies. A high-durability system implies it has technologies that have long lifetimes, lower replacement frequency and low material turnover, whereas a low durability system involves frequent replacement and short lifetime technologies. In this study, the durability of an energy system is measured using the average lifetime of the system, the average remaining lifetime, and the remaining capacity. In this way, both the lifetime of the system and the potential use of the system beyond the modeling horizon are considered.

The Average Lifetime of the System (ALS) is defined as

$$ALS = \frac{\sum_{y\in Y}\sum_{t\in T} Exp_{t,y}\cdot L_t}{\sum_{y\in Y}\sum_{t\in T} Exp_{t,y}} \quad (4)$$

where L_t is the lifetime of technology t.

On the other hand, the average remaining lifetime (ARL) is defined as

$$ARL = \frac{\sum_{y\in Y}\sum_{t\in T} Exp_{t,y}\cdot RL_{t,y_f}}{\sum_{y\in Y}\sum_{t\in T} Exp_{t,y}} \quad (5)$$

where $RL_{t,y}$ is the remaining lifetime at the end the of planning horizon (year y_f) defined as $\max(0, L_t - (y_f - y))$.

Finally, the remaining capacity (RC) at the end of the modeling horizon is calculated as

$$SR = \sum_{t\in T} Exp_{t,y}\cdot RL_{t,y_f} \quad (6)$$

Technology lifetime determines the replacement and end-of-life decisions, influencing the three metrics. This allows the framework to account for energy portfolios with technologies with different lifetimes and the installation year decisions.

Other categories

The remaining categories, including water, waste, energy, emissions and spillages, are adopted from the original MICRON framework, and adapted for energy systems by normalizing with respect to the energy produced.

CASE STUDY

The framework is applied to a case study that considers the energy transition at the University of Wisconsin-Madison to evaluate how decarbonization strategies affect circularity. The analysis focuses on the planning of the electricity system at a campus scale over a 25-year horizon (2025-2050). A cost-minimizing energy planning model developed in [19] is used to determine capacity expansion and energy dispatch decisions. The model represents a grid connected system with solar, wind, and lithium-ion battery storage technologies. Grid reliance is included to account for the GHG emissions operational targets and is defined as the share of the total electricity

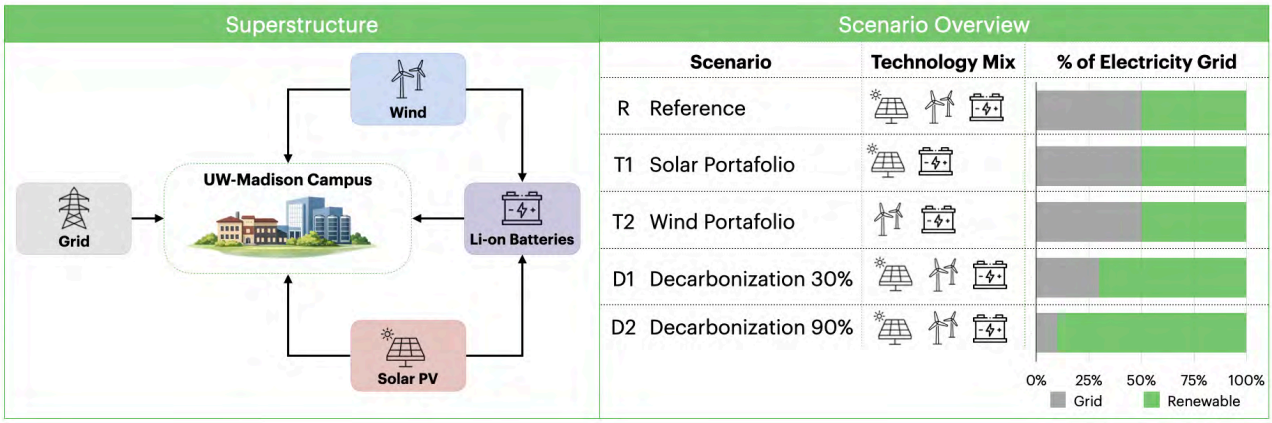

Figure 3. Summary of the case study including superstructure and scenarios studied.

Table 2: Summary of results for different technology portfolios.

Scenario	Technology Portfolio	% Grid Electricity	Storage Use	Total Costs	Scope 2 Emissions	Scope 3 Emissions	Total GHG Emissions	Circularity Overall
				MM USD	ton CO_2-eq	ton CO_2-eq	ton CO_2-eq	Index
R	Solar-Wind	50%	No	1,378	829	168	997	0.61
T1	Solar	50%	No	6,252	791	950	1,742	0.39
T2	Wind	50%	No	1,391	801	949	949	0.66
D1	Solar-Wind	30%	Yes	1,629	539	259	798	0.56
D2	Solar-Wind	10%	Yes	2,892	273	580	853	0.55

demand supplied by the grid. Circularity is quantified using the MICRON framework as a post-processing step. The overall modeling and assessment framework is illustrated in Figure 2.

The scenarios evaluated and the superstructure of the problem are summarized in Figure 3. A reference configuration is defined and used to compare the impact of technology portfolio (Scenarios T1-T2) and increasing decarbonization stringency through reduced grid reliance (Scenarios D1-D2). All scenarios assume identical electricity demand, technology performance parameters, and cost assumptions. As an initial approach, no recycling of technologies is considered, i.e., all waste is sent to landfill. Later, the effect of recycling is studied in the sensitivity analysis.

RESULTS & DISCUSSION

Table 2 summarizes the system's costs, GHG emissions and the overall circularity index for the scenarios analyzed. The reference scenario shows a circularity index of 0.61 with a cumulative capacity of solar PV of 96 MW and 116 MW for wind. Compared to other portfolios, the solar only case (scenario T1) exhibits the lowest circularity (0.30) and the highest GHG emissions. Meeting the demand with only solar energy required approximately 1, 580 MW of capacity, which substantially increases the material requirements and indirect GHG (Scope 3) emissions. Additionally, as the system does not include energy storage system, there is a significant curtailment of excess energy generation, reducing the system performance and penalizing the circularity score. Conversely, the only wind portfolio achieves the highest circularity (0.66) among the three configurations. The higher efficiency of wind turbines and greater availability of the resource, improve the water and procurement, energy and emissions sub-indices, translating in a higher overall circularity index. These results indicates that resource availability and efficiency are as important as material composition for circularity.

Analyzing the effect of grid reliance, we can see that lower grid reliance decreases operational emissions (Scope 2), but it can reduce overall circularity. This reduction is driven by two factors. First, reducing grid reliance requires the addition of energy storage, represented here by lithium batteries, which increases the share of critical materials and scarcity. Second, achieving higher self-sufficiency systems requires larger deployment of renewable energy, raising total material demand and associated indirect (Scope 3) emissions. Overall, the results show that reducing operational emissions (scope 2) that not always reduce overall emissions. Therefore,

the reduction in operational emissions can be downside by the increased impact of material extraction and waste generation. It is important to notice that the systems studied do not have Scope 1 emission for Scope 1 as it does not include any technology with fuel consumption.

Figure 4 shows the circularity subindex values for each scenario. Differences are primarily driven by the material-related categories. In terms of subindices, the reference scenario is relative balanced, with the highest score for Durability. This performance is explained by the hybrid option, which reduces oversizing of the system and storage while maintaining reliability. Scenario T1 has the lowest scores in Water & Procurement and Energy, due to the oversized solar PV system and energy curtailment. Scenarios D1 and D2 show low Durability scores due to the relatively short lifetime of batteries which requires more frequent replacement compared to wind turbines or PV panels. Consequently, improvements in Emissions sub-index are offset by effects of deployment and EoL, reducing overall circularity. These findings suggest that decarbonization strategies focusing solely on emission reduction may shift into a material burden. All scenarios analyzed have similar Waste sub-index scores, as no recycling of technologies is considered.

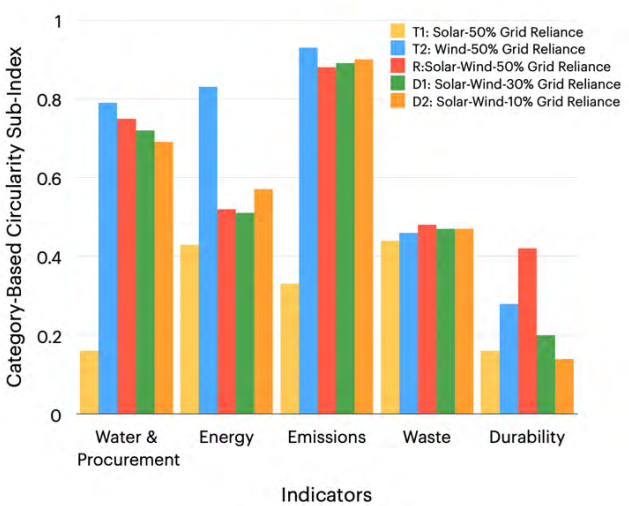

Figure 4. Circularity Sub-Index values for different scenarios analyzed.

Sensitivity Analysis

A single parameter sensitivity analysis was performed to identify the key drivers of circularity. Each parameter is varied individually around the reference scenario (R) while all other parameters are kept constant. The parameters analyzed included the share of critical material, material scarcity, material intensity, % of recycled material used, and % of material recovered at the EoL. Figure 5 shows the corresponding effect, a ± 25% variation was applied to each parameter. Results indicate that the percentage of material recovered has the largest influence on circularity; in contrast, no change was seen for amount of recycled material used during manufacturing. This highlights the importance of EoL management of renewable technology to achieve a sustainable energy transition. On the other hand, small changes were seen for the percentage of critical material and scarcity, suggesting that it is important to focus on the type of material used and not only on the amount.

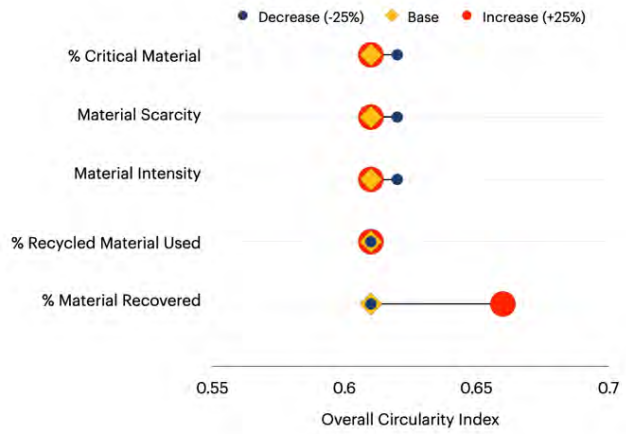

Figure 5. Single-parameter sensitivity analysis on the overall circularity index.

CONCLUSIONS

This study proposed an integrated metric to assess and quantify circularity in energy systems based on the MICRON CE framework. The framework was adapted to account for the specific characteristics of energy technologies by incorporating metrics related to critical material, scarcity, material intensity and durability. The proposed metric was applied to evaluate circularity in an electricity planning problem at the University of Wisconsin-Madison.

While the scenarios studied showed similar GHG emissions, they differ in material use, durability and the overall circularity index. This demonstrates that GHG emissions alone are insufficient to characterize sustainability in the design and planning of energy systems. The results also show that lowering grid reliance, hence reducing operational emissions, does not necessarily improve circularity. The reduction of operational emissions can be offset by higher material requirements and increased waste generation.

The percentage of material recovered influences improvements in circularity. Therefore, future work will focus on evaluating the circularity of different recycling pathways for energy systems and integrating the circularity metric as a design criterion in an optimization planning problem, enabling the systematic design of energy systems that balance decarbonization and circularity objectives.

ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation under Grant No. CMMI-2328160, the Wisconsin Alumni Research Foundation, and the Department of Chemical & Biological Engineering at the University of Wisconsin-Madison.

AUTHOR IDENTIFIERS

Author ORCIDs:
Vergara-Zambrano J: 0000-0001-6899-6607
Avraamidou S: 0000-0002-9334-9951

REFERENCES

1. Wang Q, Wang X, Li R. Energy transition and environmental sustainability: the interplay with natural resource rents and trade openness. Humanit Soc Sci Commun 12: (2025). https://doi.org/10.1057/s41599-025-05521-4
2. Gielen, D. Critical minerals for the energy transition. International Renewable Energy Agency (2021)
3. Avraamidou S, Baratsas SG, Tian Y, Pistikopoulos EN. Circular economy - a challenge and an opportunity for process systems engineering. Computers & Chemical Engineering 133:106629 (2020). https://doi.org/10.1016/j.compchemeng.2019.106629
4. Renaldi V. Sustainable energy transition: circular economy and sustainable financing for environmental, social and governance (ESG) practices. Social and Environmental Accountability Journal 44:266-269 (2024). https://doi.org/10.1080/0969160x.2024.2401212
5. Mali S, Garrett P. Life cycle assessment of electricity production from an onshore V150-4.2 MW wind plant (Technical report). Vestas (2022)
6. Mirletz H, Ovaitt S, Sridhar S, Barnes TM. Circular economy priorities for photovoltaics in the energy transition. PLoS ONE 17:e0274351 (2022). https://doi.org/10.1371/journal.pone.0274351
7. Avraamidou S, Torres AI. Circular Economy: Definitions, Challenges and Opportunities. FOCAPO (2023)
8. Ellen MacArthur Foundation. Circularity Indicators: An Approach to Measure Circularity. Methodology. (2015)
9. Cullen JM. Circular economy: theoretical benchmark or perpetual motion machine?. J of Industrial Ecology 21:483-486 (2017). https://doi.org/10.1111/jiec.12599
10. De Wit M, Hoogzaad J, Ramkumar S, Friedl H, Douma A. The Circularity Gap Report. Circle Economy (2018)
11. Fogarassy C, Horvath B, Kovacs A, Szoke L, Takacs-Gyorgy K. A circular evaluation tool for sustainable event management – An Olympic case study. APH 14: (2018). https://doi.org/10.12700/aph.14.7.2017.7.10
12. Kalmykova Y, Sadagopan M, Rosado L. Circular economy – from review of theories and practices to development of implementation tools. Resources, Conservation and Recycling 135:190-201 (2018). https://doi.org/10.1016/j.resconrec.2017.10.034
13. Fellner J, Lederer J. Recycling rate – the only practical metric for a circular economy?. Waste Management 113:319-320 (2020). https://doi.org/10.1016/j.wasman.2020.06.013
14. López-de-Lérida-Elías C, Díaz-Alvarado F, Palma-Behnke R. Beyond emission reductions: circular economy indicators for the energy transition of the power sector. Ind. Eng. Chem. Res. 64:9338-9348 (2025). https://doi.org/10.1021/acs.iecr.5c00415
15. Weisser D. A guide to life-cycle greenhouse gas (GHG) emissions from electric supply technologies. Energy 32:1543-1559 (2007). https://doi.org/10.1016/j.energy.2007.01.008
16. Serrano-Arévalo TI, Díaz-Alvarado FA, Tovar-Facio J, Ponce-Ortega JM. The impact of circular economy indicators in the optimal planning of energy systems. Sustainable Production and Consumption 44:234-249 (2024). https://doi.org/10.1016/j.spc.2023.12.018
17. Bin X, Thakur J. Circular economy metrics for batteries: enhancing sustainability in energy storage systems. Sustainable Production and Consumption 55:470-485 (2025). https://doi.org/10.1016/j.spc.2025.02.014
18. Corona B, Shen L, Reike D, Rosales Carreón J, Worrell E. Towards sustainable development through the circular economy—a review and critical assessment on current circularity metrics. Resources, Conservation and Recycling 151:104498 (2019). https://doi.org/10.1016/j.resconrec.2019.104498
19. Baratsas SG, Pistikopoulos EN, Avraamidou S. A quantitative and holistic circular economy assessment framework at the micro level. Computers & Chemical Engineering 160:107697 (2022). https://doi.org/10.1016/j.compchemeng.2022.107697
20. Munoz-Briones P, del Carmen Munguía-López A, Rivera KS, Avraamidou S. Integrated decision-making approach for the simultaneous design of food packaging and waste management technologies to achieve a circular economy. Computers & Chemical Engineering 202:109269

(2025). https://doi.org/10.1016/j.compchemeng.2025.109269
21. Martinez JK, Brionez-Munoz PA, Avraamidou S Circular Economy Assessment of Viable for Fuels for Mobility. Comput. Aided Chem. Eng. 53:2335-2340 (2024) https://doi.org/10.1016/B978-0-443-28824-1.50390-2
22. U.S. Department of Energy. Critical Materials Assessment. Tech. Report (2023)
23. European Commission. Critical Raw Materials Resilience: Charting a Path towards greater Security and Sustainability. Tech. Report (2020)
24. van Oers L, Guinée JB, Heijungs R. Abiotic resource depletion potentials (adps) for elements revisited—updating ultimate reserve estimates and introducing time series for production data. Int J Life Cycle Assess 25:294-308 (2019). https://doi.org/10.1007/s11367-019-01683-x
25. Vergara-Zambrano J, Brahmbhatt P, Avraamidou S. A multi-scale optimization framework for energy transition planning in urban areas: insights from a university campus case study. Sustainable Energy, Grids and Networks 44:101996 (2025). https://doi.org/10.1016/j.segan.2025.101996

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Hybrid Modelling of Segmented Flow Extraction Process for Digital Twin Development in Critical Metals Recovery

Arun Pankajakshan, Konstantinos Katsoulas, Malik Olasinde, Cong Chao, Eric S. Fraga, Panagiota Angeli and Federico Galvanin*

Department of Chemical Engineering, University College London, Torrington Place, WC1E 7JE, London, United Kingdom
* Corresponding Author: f.galvanin@ucl.ac.uk.

ABSTRACT

Critical metals are indispensable in renewable, low-carbon, and hydrogen technologies due to their unique catalytic and electrochemical properties. They are primarily sourced through mining, which is associated with significant environmental impacts and geopolitical risks due to the uneven global distribution of ore deposits. As a result, efficient recovery of these metals from secondary sources such as electronic waste has become increasingly important. In this context, liquid-liquid extraction (LLE) has emerged as a promising separation technique due to its high selectivity and scalability. The development of intensified, continuous-flow LLE in small channels offers further advantages in terms of mass transfer efficiency, solvent utilization, and process sustainability, making it an attractive approach for the recovery of critical metals. A flow pattern known as segmented flow further enhances mass transfer in LLE in small channels. This work presents a hybrid modelling approach for developing a predictive model of a segmented flow LLE process, intended for digital twin implementation in critical metals recovery. Within the hybrid modelling framework, mass transfer is modelled using a lumped approach, which allows to treat mass transfer independently from flow hydrodynamics. Further, hydrodynamic and mass transfer models are developed in parallel using Gaussian process (GP)-based active learning (AL) and model-based design of experiments (MBDoE), respectively. Prior knowledge of flow regimes is used in developing the hydrodynamic model. The method was tested in an *in silico* case study and shown to efficiently develop reliable models for segmented flow extraction in small channels.

Keywords: hybrid modelling, active learning, model-based design of experiments, critical metals, extraction, segmented flow

INTRODUCTION

Critical metals are essential to the functioning of modern technological systems, with applications ranging from lithium-ion batteries in smartphones and laptops to rare earth elements (REEs) used in permanent magnets for wind turbines and electric vehicles. These vital metals are used in many sectors including communications, renewable energy, transport and life sciences (Department for Business & Trade, 2026). The primary sources of critical metals are economically viable ore deposits, whose distribution across the globe is greatly uneven. In addition, the process of mining to obtain critical metals from their ore deposits is a highly resource-intensive process and can cause enormous damage to the environment. As a result, the supply of critical metals is increasingly vulnerable due to geopolitical, environmental and economic constraints. This risk in supply chain has pushed the recovery of critical metals from secondary sources such as electronic waste (e-waste) (Sun et al., 2017).

The recovery of critical metals from e-waste usually proceeds through pyrometallurgical or hydrometallurgical routes. Hydrometallurgical routes are preferred due to their flexibility, high recovery rate and sustainable nature. Hydrometallurgical processes essentially involve dissolving e-waste in acid solutions followed by selective separation of critical metals via liquid-liquid extraction (LLE). LLE involves contacting two immiscible liquids to selectively transfer solutes between them. Intensified, continuous flow LLE in small channels offers further

advantages such as enhanced mass transfer, less material consumption, precise control of operating conditions, real-time monitoring and ease of automation and digitization. Based on the way in which the immiscible liquids are contacted, in small channels, LLE can proceed through different flow patterns. One of the possible flow patterns is segmented flow which is a two-phase flow pattern in which one immiscible liquid forms discrete segments within a continuous phase of the other liquid (Garstecki et al., 2006; Mac Giolla Eain et al., 2013). Segmented flow is particularly attractive due its regular hydrodynamics and enhanced mass transfer, supporting the development of predictive models and advanced process control. Consequently, implementing digital twins of segmented flow LLE in small channels can accelerate the development of critical metals recovery from e-waste. A digital twin is a software model (Ljung, 1999) of a physical system which interacts with the physical system continuously and acts as a reliable virtual entity of the real world described by the physical system. Compared to the traditional simulation models, digital twins are adaptive and responsive real time simulations enabled via fast computation, advanced sensing and automation. Digital twins are becoming increasingly important to automate decision making related to productivity, cost, energy consumption, and sustainability for complicated systems. The core of such digital twins are mathematical models of physical systems, which are instrumental for simulation and prediction.

The aim of this study is in developing a digital twin model of a segmented flow LLE system for recovering critical metals from e-waste. The geometry of segmented flow pattern and the resulting hydrodynamic characteristics strongly influence the mass transfer between the two liquid phases. Consequently, modelling segmented flow requires to develop both hydrodynamic and mass transfer models.

Hydrodynamic modelling of segmented flow in small channels has traditionally relied on a combination of mechanistic and semi-empirical approaches. Mechanistic models, based on computational fluid dynamics (CFD) simulations, offer detailed insight but are computationally expensive. Semi-empirical models, on the other hand, use fitted correlations for flow parameters and are typically calibrated for specific fluid systems and operating conditions. They often fail to reliably predict system behaviour under new operating conditions which limit their application as digital twins. In this work we combine two powerful modelling techniques – probabilistic modelling via Gaussian processes (GPs) and model-based design of experiments (MBDoE) (Geremia et al., 2026) methods to develop a hybrid modelling framework for the segmented flow. In the proposed framework, hydrodynamic models are developed using GP and active learning (AL) (Lewis & Catlett, 1994), while MBDoE methods are employed to precisely estimate the parameters of the mass transfer model. The framework is tested on a simulated case study and is further shown to be effective in developing hybrid models of segmented flow LLE processes.

2. BACKGROUND AND METHODOLOGY

As shown in Fig. 1, segmented flow in small channels is characterised by repeated volumes of one phase (dispersed phase) in a continuous volume formed by the other phase (continuous phase). The dispersed phase is called plug and the continuous phase volume between two plugs is called slug. A single plug and slug constitute a unit cell. The size of plug, slug and unit cell are characterised by their lengths L_p, L_s and L_u, respectively. Moreover, plugs are assumed to be cylindrical in shape with two hemispherical caps attached at both ends.

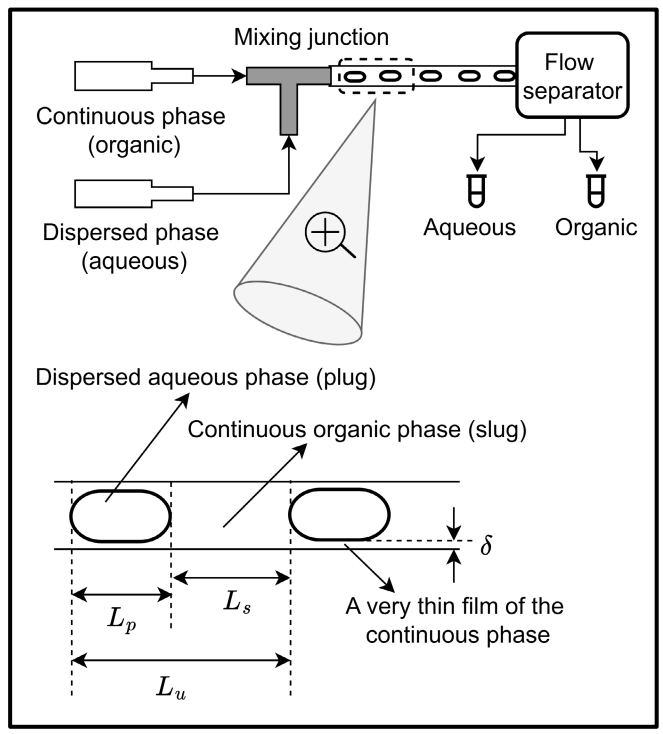

Figure 1. Schematic showing a segmented flow contactor set up, the segmented flow pattern and hydrodynamic characteristics of the segmented flow pattern.

In this work we consider a segmented flow LLE process in which the mass transfer by extraction happens from the dispersed aqueous phase to the continuous organic phase. Typically, the rate of mass transfer is affected by three main factors – the velocity and mixing patterns within each of the phases, the concentration driving forces and the interfacial area available for mass transfer. These three factors are influenced by several variables. Some of them for example, the design of the

mixing junctions (T-shaped junction or Y shaped junction, rectangular or circular cross section), the channel diameter and length and the properties of the contacted liquids are fixed for a specific system. While some variables such as the mixture velocity, the flow rate fraction of any of the two phases or the initial concentration of metals in the aqueous phase can be varied to study their effects on the plug-slug geometries as well as on the rate of mass transfer.

2.1 Fundamentals of modelling segmented flow extraction in small channels

Mathematical models of segmented flow LLE in small channels consist of a hydrodynamic model, a model for the pressure drop and a mass transfer model. Important hydrodynamic parameters related to segmented flow LLE are: a) primary flow geometries such as L_p, L_s and L_u, b) film thickness (δ), c) the dispersed phase holdup (ε_d) and d) the specific interfacial area (a). As shown in Fig. 1, a thin film of the continuous phase separates the slugs from the channel wall. The dispersed phase holdup is the volume fraction of the dispersed phase within one unit cell of the segmented flow:

$$\varepsilon_{\mathrm{d}} = \frac{\text{volume of dispersed phase in one unit cell}}{\text{total volume of the unit cell}} \tag{1}$$

The specific interfacial area is the area available for mass transfer per volume of the unit cell, and is defined as:

$$a = \frac{\text{total surface area of the dispersed phase}}{\text{total volume of the unit cell}} \tag{2}$$

The hydrodynamic modelling involves developing mathematical relationships between the inputs or the operating conditions to the primary flow geometries as well as to the film thickness. Typically, the relationship between inputs and film thickness is modelled using physics-based correlations such as Bretherton-type scaling laws, often extended to account for fluid properties and operating conditions (Mac Giolla Eain et al., 2013). Such correlations are often functions of dimensionless numbers such as the capillary number and the Weber number, depending on whether the flow is in the visco-capillary regime or in the visco-inertial regime (Mac Giolla Eain et al., 2013). Given L_p, L_s, L_u and δ, and assuming a plug geometry consisting of a cylindrical core with two hemispherical caps, expressions for ε_d and a can be derived using simple geometric relations. The rate of mass transfer between the dispersed phase and the continuous phase depends on the driving force, i.e., the metal concentration difference between the two phases, and the interfacial area available for mass transfer. Assuming an ideal plug flow behaviour (producing plugs of same dimension with no axial mixing) and negligible concentration gradient along the radial direction, the mass transfer along the channel length can be described by:

$$u_{\mathrm{mix}} \frac{\mathrm{d}C_{\mathrm{aq}}}{\mathrm{d}z} = -\frac{K_{\mathrm{aq}} a}{\varepsilon_{\mathrm{d}}} \left(C_{\mathrm{aq}}(z) - C_{\mathrm{aq}}^{\mathrm{eq}}\right) \tag{3}$$

In Eq. (3), $u_{mix}, K_{aq}, C_{aq}(z)$ and C_{aq}^{eq} are the mixture velocity, aqueous phase mass transfer coefficient, concentration of metal in the aqueous phase at a distance z from the channel inlet and the equilibrium concentration of metal, respectively. The mixture velocity is defined as the ratio of the total volumetric flow rate (Q_T) and the cross-sectional area of the channel with internal diameter d_{int}. The total volumetric flow rate is defined as the sum of volumetric flow rates of the continuous phase (Q_c) and the dispersed phase (Q_d). The equilibrium concentration is defined as:

$$C_{\mathrm{aq}}^{\mathrm{eq}} = \frac{C_{\mathrm{aq}}^{\mathrm{in}}}{D(Q_{\mathrm{d}}/Q_{\mathrm{c}})+1} \tag{4}$$

In Eq. (4), D, C_{aq}^{in}, Q_d and Q_c are the distribution coefficient, initial concentration of metal in the aqueous phase, volumetric flow rate of the dispersed phase and the volumetric flow rate of the continuous phase, respectively. The distribution coefficient is usually evaluated from batch equilibrium experiments. Pressure drop is a key parameter in designing segmented flow contactors as it provides crucial information regarding energy consumption. However, it becomes more significant in designing scale-out segmented flow systems (Garciadiego Ortega et al., 2018) where multiple channels are involved. For this reason, in this study, a pressure drop model is not considered.

2.2 A hybrid modelling framework for segmented flow extraction in small channels

The mass transfer model in the form provided in Eq. (3) imposes two challenges. The first is its direct dependency on hydrodynamic terms such as the specific interfacial area and the dispersed phase holdup. This prevents the independent development (estimation and validation) of the mass transfer model as the predictive quality of the mass transfer model strongly depends on the accurate evaluation of ε_{d} and a using the hydrodynamic model. Secondly, the aqueous phase mass transfer coefficient K_{aq} in Eq. (3) relates the mass flux through the interface to the concentration driving force. Consequently, when the interfacial area changes (i.e., the plug shape changes), for example due to variations in the flow rate, the value of mass transfer coefficient changes. For this reason, it is a local property that depends on the specific hydrodynamic conditions and cannot be uniquely estimated as a single model parameter over a range of operating conditions. As a solution to both the challenges, we lump the mass transfer coefficient, the specific interfacial area and the dispersed phase holdup into a single parameter k' given by:

$$k' = K_{\mathrm{aq}} a / \varepsilon_{\mathrm{d}} \tag{5}$$

The lumped parameter k' physically signifies mass transfer per unit interfacial area per unit volume of the dispersed phase. This parameter can be treated as constant (Garciadiego-Ortega et al., 2020) across a wide range of operating conditions that affect the primary plug-slug geometries, and subsequently the interfacial area and holdup. However, the parameter k' can be affected by the mixture velocity, which promotes internal circulation within the plugs and slugs and enhances mass transfer even without changing the plug geometry. Therefore, analogously to the Arrhenius equation describing the temperature dependency of reaction rate constants, we employ a simple scaling rule to account for velocity dependency of k', given by:

$$k' = k'_0\left(\frac{u_{\mathrm{mix}}}{u_{\mathrm{mix}}^{\mathrm{ref}}}\right)^p \tag{6}$$

In Eq. (6), k'_0 is the lumped mass transfer coefficient corresponding to a reference mixture velocity u_{mix}^{ref}. The exponent p in Eq. (6) represents the combined sensitivity of the effective mass transfer rate to mixture velocity, accounting for velocity-induced changes in interfacial area, phase holdup, and local mass transfer. By substituting Eq. (5) and (6) in Eq. (3), the mass transfer model can be modified as:

$$u_{\mathrm{mix}}\frac{\mathrm{d}C_{\mathrm{aq}}}{\mathrm{d}z} = -k'_0\left(\frac{u_{\mathrm{mix}}}{u_{\mathrm{mix}}^{\mathrm{ref}}}\right)^p\left(C_{\mathrm{aq}}(z) - C_{\mathrm{aq}}^{\mathrm{eq}}\right) \tag{7}$$

Now, Eq. (7) contains two parameters k'_0 and p that can be uniquely estimated so that Eq. (7) holds true for a wide range of operating conditions. This approach allows the mass transfer model to be developed independently of the hydrodynamic model. We suggest the application of closed loop experimentation via optimal feedback loop (with minimum data points) to develop both hydrodynamic and mass transfer model. To this end, we employ GP modelling via AL for the development of hydrodynamic model and MBDoE for the mass transfer model development. Details of these methods are given below.

2.2.1 Gaussian processes and active learning for hydrodynamic model development

Consider a small channel segmented flow LLE system in which the aqueous phase forms the plug and the organic phase forms the slug. Let the system is affected by two inputs – mixture velocity and the continuous phase flow rate fraction (ϕ_c), which is defined as the ratio of continuous phase volumetric flow rate over the total volumetric flow rate. Now, an experiment i is identified by the input vector $\mathbf{u}_i = [u_{i1}, u_{i2}]$ and the outputs - $\mathbf{y}_{h_i} = [y_{h_{i1}}, y_{h_{i2}}]$ and y_{m_i}. Here, u_{i1} and u_{i2} respectively denote the values of mixture velocity and continuous phase flow rate fraction that define experiment i. Similarly, $y_{h_{i1}}, y_{h_{i2}}$ and y_{m_i} respectively denote the observed values of plug length [m], unit cell length [m] and concentration of metal [mol/L] in the dispersed aqueous phase at the outlet of the channel of length L_z. It is assumed that the observations $y_{h_{i1}}, y_{h_{i2}}$ and y_{m_i} are corrupted by independent and identically distributed Gaussian noise with zero mean and standard deviations of σ_1, σ_2 and σ_c, respectively.

Consider a dataset from n experiments consisting of inputs $\mathbf{u}_1, \dots, \mathbf{u}_{\mathrm{n}}$ and outputs $\mathbf{y}_{h_1}, \dots, \mathbf{y}_{h_n}$ and $y_{m_1}, \dots, y_{m_n}$. The relationship between $\mathbf{u}$ and $\mathbf{y}_h$, which is the relationship between the inputs and the primary plug-slug geometries is usually modelled as empirical correlations validated using polynomial regression. Disadvantages of this approach are the lack of proper quantification of the uncertainty in model predictions and the lack of accurate predictions under new operating conditions, which were not used in calibration. As a better approach that guarantees accurate estimation of prediction uncertainties, and that captures non-linear system behaviour, we model the relationship between $\mathbf{u}$ and $\mathbf{y}_h$ using a multiple-output GP model (Bonilla et al., 2007).

In a multiple-output GP modelling framework, several related output variables are modelled together by assuming each output having an underlying hidden function called the latent function. Mathematically, the observed value of output l at input $\mathbf{u}_i$ can be represented as:

$$y_{h_{il}} = f_l(\mathbf{u}_i) + \epsilon_l, \epsilon_l \sim \mathcal{N}(0, \sigma_l^2) \tag{8}$$

In Eq. (8), f_l is the latent function for output l and ϵ_l is the error term. In the GP framework, the latent functions are treated as random functions drawn from a GP, meaning that for any set of inputs $\mathbf{u}_1, \dots, \mathbf{u}_n$, the latent function values follow a multivariate Gaussian distribution.

$$[f_{11}, \dots, f_{\mathrm{n}1}, f_{12}, \dots, f_{\mathrm{n}2}; \mathbf{u}_1, \dots, \mathbf{u}_{\mathrm{n}}] \sim \mathcal{N}(\mathbf{m}, \mathbf{K}) \tag{9}$$

In Eq. (9), $\mathbf{m}$ and $\mathbf{K}$ represent the mean vector and the covariance matrix, respectively. The covariance between any two latent function values is given by:

$$Cov\left(f_l(\mathbf{u}), f_j(\mathbf{u}')\right) = \kappa_{lj}^f \kappa_u(\mathbf{u}, \mathbf{u}') \tag{10}$$

The RHS of Eq. (10) is called a kernel function. Essentially, it assigns covariance between the latent function values based on the covariance or similarities observed among inputs and outputs. According to Eq. (10), similarity between inputs is captured by $\kappa_u(\mathbf{u}, \mathbf{u}')$ whereas similarity between outputs is captured by κ^f. We employed the intrinsic coregionalization model (ICM) kernel proposed by Bonilla et al., (2007) for the product $\kappa_{lj}^f \kappa_u(\mathbf{u}, \mathbf{u}')$, with the radial basis function (RBF) kernel for $\kappa_u(\mathbf{u}, \mathbf{u}')$.

The GP model training involves estimating the optimal values of hyperparameters of the kernel function that minimises a loss function. We used the negative marginal log-likelihood as the loss function. Once trained with few points, the GP model can be used to predict the plug and unit length for any unseen input data point. Along with these predictions, the model also provides an estimate of

uncertainty associated with the predictions. This uncertainty will be largest in regions less explored, where the model is least confident. Therefore, selecting new experiments in the largest uncertainty region and retraining the model with those new data points helps to reduce the model prediction uncertainty over the experimental iterations. This iterative approach of training the model with informative data points is known as active learning (AL) (Lewis & Catlett, 1994). Mathematically, AL is an optimisation problem of the form:

$$\max_{\mathbf{u}_{\mathrm{lb}} \le \mathbf{u} \le \mathbf{u}_{\mathrm{ub}}} \psi\{\mathrm{Var}[f(\mathbf{u})]\} \tag{11}$$

where, $Var[f(\mathbf{u})]$ is the posterior predictive variance of the GP model at conditions defined by the vector $\mathbf{u}$ and ψ is a scalar function of the variance vector. We have chosen ψ as the ℓ^2-norm of the variance vector. The AL procedure was applied iteratively to develop a reliable GP model that maps the mixture velocity and continuous phase flow rate fraction to the plug and unit lengths.

2.2.2 Model-based design of experiments for mass transfer model development

The development of mass transfer model presented in Eq. (7) essentially requires precise estimation of the two model parameters lumped in the N_θ (here, 2) dimensional vector $\theta = [k_0', p]$. Let us revisit the dataset formed from n experiments. The mass transfer data consists of the set of inputs $\mathbf{u}_1, \dots, \mathbf{u}_\mathrm{n}$ and outputs $y_{m_1}, \dots, y_{m_n}$. Each input vector corresponds to an experimental condition specified by the values of mixture velocity and continuous phase flow rate fraction, while the output variables denote the metal concentration in the dispersed phase at the channel outlet. With the availability of such a dataset, a parameter estimation can be performed to obtain the best values of the model parameters that minimises the discrepancy between predicted and observed concentration values. We employed maximum likelihood method of parameter estimation to obtain the optimal parameter estimates $\hat{\theta}$.

The uncertainty in estimation of $\hat{\theta}$ can be quantified as the parameter covariance matrix $\mathbf{V}_\theta$, which in turn can be approximated as the inverse of observed Fisher information matrix $\mathbf{H}_\theta$ (Geremia et al., 2026).

$$\mathbf{H}_\theta = \left[\mathbf{V}_\theta^0\right]^{-1} + \sum_{\mathrm{i}=1}^{n} \mathbf{S}_\mathrm{i}^\mathrm{T} (\sigma_\mathrm{c}^2)^{-1} \mathbf{S}_\mathrm{i} \tag{12}$$

In Eq. (12), $\mathbf{S}_\mathrm{i}$ represents the parameter sensitivity matrix with dimensions equal to number of output variables times the number of model parameters. Since, we have the case with only one output variable, $\mathbf{S}_\mathrm{i}$ is a row vector whose elements represent the first derivatives of the output variable with respect to the model parameters, evaluated at the estimated parameter values and at the condition investigated in experiment i. The parameter covariance matrix $\mathbf{V}_\theta$ is evaluated as the inverse of $\mathbf{H}_\theta$. The statistical precision in which the parameters are estimated can be evaluated using a t-test. From the parameter covariance matrix, the test statistic t-value for the t-test can be computed for each parameter estimate as:

$$t_\mathrm{i} = \frac{\hat{\theta}_\mathrm{i}}{t_{N-N_\theta}(1-\alpha/2)\sqrt{\mathbf{V}_{\theta_\mathrm{ii}}}} \tag{13}$$

When the computed t-value for an estimate is greater than a reference t-value, $t_{ref} = t_{N-N_\theta}(1-\alpha/2)$, that parameter is identified as precisely estimated, otherwise the test fails and new experiments need to be designed to improve the statistical precision. This can be achieved by solving the MBDoE optimisation problem, which suggests new condition that minimises a scalar function ξ of the expected parameter covariance matrix $\hat{\mathbf{V}}_\theta$.

$$\min_{\mathbf{u}_{\mathrm{lb}} \le \mathbf{u} \le \mathbf{u}_{\mathrm{ub}}} \xi\left[\hat{\mathbf{V}}_\theta(\mathbf{u})\right] \tag{14}$$

Popular choices of ξ are the alphabetic design criteria with A-, D- and E-designs, respectively minimising the trace, determinant and maximum eigenvalue of the expected covariance matrix. Like AL, MBDoE can be also performed iteratively in closed loop until all the model parameters are precisely estimated.

2.3 Case study

A simulation study was designed to test the proposed hybrid modelling approach. For the simulation study, the extraction of uranium U(VI) by tributyl phosphate (TBP) (100 %) (Garciadiego-Ortega et al., 2020) in a channel of internal diameter 2 mm was selected as the case study. In their study, Garciadiego-Ortega et al. (2020) have developed the models provided in Eq. (15) and (16) to predict the dimensionless plug and slug lengths. They also validated these models over the range of operating conditions provided in Table 1.

Table 1: Experimental design space in terms of variable bounds.

Input	Bounds
Mixture velocity (cm/s)	[1.062, 4.24]
Continuous phase flow rate fraction (–)	[0.2, 0.5]

In this study, the models with the coefficients in Eq. 15 and 16 are assumed as the "true models" representing the hydrodynamic component of the segmented flow LLE system. The design space provided in Table 1 was chosen to be the design space in this simulation study.

$$\ln\left(\frac{L_\mathrm{p}}{d}\right) = -2.56 - 0.872\ln\left(\frac{Q_\mathrm{c}}{Q_\mathrm{T}}\right) - 0.36\ln\left(\frac{\mu_\mathrm{d}}{\mu_\mathrm{c}}\right) - 0.280\ln(\mathrm{Ca_c}) + 0.099\ln\left(\frac{\mathrm{Re_c}}{\mathrm{Ca_c}}\right) \tag{15}$$

$$\ln\left(1 - \frac{L_\mathrm{p}}{L_\mathrm{u}}\right) = -0.82 + 1.181\ln\left(\frac{Q_\mathrm{c}}{Q_\mathrm{T}}\right) - 0.099\ln\left(\frac{\mu_\mathrm{d}}{\mu_\mathrm{c}}\right) -$$

$$0.130 \ln(\mathrm{Ca_c}) \quad (16)$$

In their study, Garciadiego-Ortega et al. (2020) have used the following correlation proposed by Mac Giolla Eain et al., (2013) to calculate the film thickness:

$$\frac{2\delta}{d_{\mathrm{int}}} = 0.35(\mathrm{Ca}_c)^{0.354}(\mathrm{We}_c)^{0.097} \quad (17)$$

According to Mac Giolla Eain et al., (2013), Eq. (17) is only valid within a specific range of operating conditions defined by the dimensionless numbers listed in Table 2. However, a subregion of the chosen design space in Table 1 resulted in values of the dimensionless numbers outside the feasible bounds provided in Table 2.

Table 2: Dimensionless numbers and their ranges over which the film thickness correlation in Eq. (17) is valid.

Dimensionless number	Description	Desired range
Capillary (Ca)	$\mathrm{Ca}_c = \mu_c u_{mix}/\gamma$	0.002 – 0.119
Weber (We)	$\mathrm{We}_c = \rho_c u_{mix}^2 d/\gamma$	0.047 – 0.697
Reynolds (Re)	$\mathrm{Re}_c = \rho_c u_{mix} d/\mu_c$	14.46 – 100.96

Therefore, in the simulation study, the AL was solved as a constrained optimization problem with the constraint $0.047 \leq \mathrm{We} \leq 0.697$. Based on the work by Garciadiego-Ortega et al., (2020), the true values of the parameters of the lumped mass transfer model - k_0' and p were chosen as 0.055 and 0.8, respectively with $u_{mix}^{ref} = 1.062$ cm/s. For *in silico* data generation, it is assumed that the plug length, unit length and concentration data are corrupted by independent and identically distributed Gaussian noise with zero mean and standard deviations of 0.0004 m, 0.0004 m and 0.001 mol/L, respectively. The physical properties of the fluids such as viscosity (μ), density (ρ) and interfacial tension (γ) required to calculate Eq. (15-17) and the dimensionless numbers in Table 2 as well as the value of the distribution coefficient D were obtained from Garciadiego-Ortega et al., (2020). The channel length L_z was assumed as 0.5 m.

RESULTS AND DISCUSSION

Initial experiments to calibrate both the hydrodynamic and mass transfer models were generated using Latin hypercube sampling (LHS) applied in the design space provided in Table 1. Four initial experiments were designed using the LHS method. The outputs (plug length, unit length and metal concentration in the dispersed phase at the channel outlet) corresponding to these experiments were obtained by simulating the true models and adding noise to them. The initial hydrodynamic dataset was used to train the GP model, and the initial mass transfer dataset was used to estimate parameters of the lumped mass transfer model. The results of parameter estimation of the mass transfer model using the initial dataset (i.e., at the end of 4 DoE experiments) are shown in Table 3.

Table 3: Parameter estimation results after DoE and AL. The reference *t*-values after DoE and AL are 2.91 and 1.71 respectively.

Parameter	Estimate	*t*-value	Evaluation
k_0'	0.060 ±0.03	**1.98**	End of DoE
p	0.704 ±0.49	**1.43**	
k_0'	0.058 ±0.005	10.39	End of AL
p	0.734 ±0.08	8.35	

As shown in Table 3, the parameter estimates obtained from the initial DoE exhibited relatively large uncertainty. The *t*-values of the parameter estimates were smaller than the reference *t*-value (2.91). At this stage, one can perform either AL, MBDoE or a multi-objective optimization using AL and MBDoE to design next set of experiments for developing the segmented flow model. In this work, we focussed on the development of hydrodynamic model. Thus, after the DoE, a campaign of AL experiments, consisting of 20 new experiments were designed iteratively to reduce the prediction uncertainty of the GP model. In this simulated experimental campaign, AL was solved as a constrained optimization problem with the constraint equation $0.047 \leq \mathrm{We} \leq 0.697$.

The experiments designed in the AL campaign along with the initial LHS experiments are shown in Fig. 2. As shown in the Fig. 2, in the constrained AL problem, a low velocity region appeared to be infeasible as it resulted in Re and We below their validated range. A physical significance of this observation is that the correlation for film thickness provided in Eq. (17) applies to the visco-inertial regions, where the velocity or inertia force start to have an influence on the film thickness. However, at low velocity, this influence is negligible with capillary and viscous forces dominating (Mac Giolla Eain et al., 2013). Hence, even the small weighting of the Weber number is not appropriate in calculating film thickness at low velocities. As a result, this low velocity region proved to be the infeasible region in the constrained AL problem. It can be also inferred from Fig. 2 that the constrained AL problem always resulted in feasible solutions, however, very close to the feasible boundary. This is because AL picks new experiments in the less explored regions.

The experiment numbers showing the exact iteration of experiments are provided as labels in Fig. 2. It is also understood from Fig. 2 that the AL campaign was successful in exploring the design space. Further evidence of the good performance of the AL campaign in reducing the prediction uncertainty of the GP model is

provided in Fig. 3.

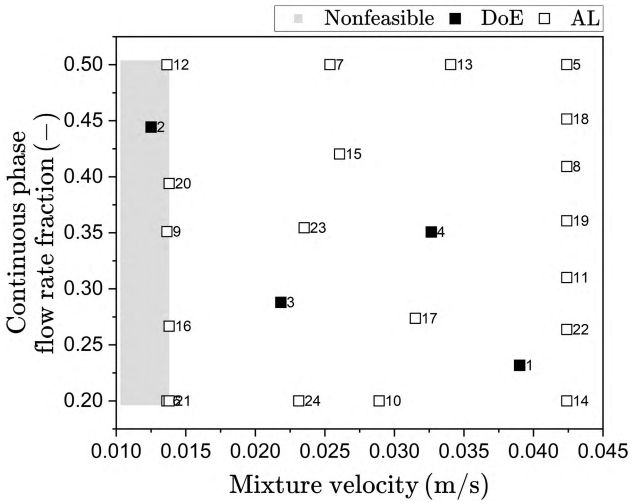

Figure 2. The simulated experimental campaign showing the DoE (filled squares) and AL (open squares) experiments. Labels adjacent to the scatter points denote the experiment number. The nonfeasible design space is shown using grey scatter points.

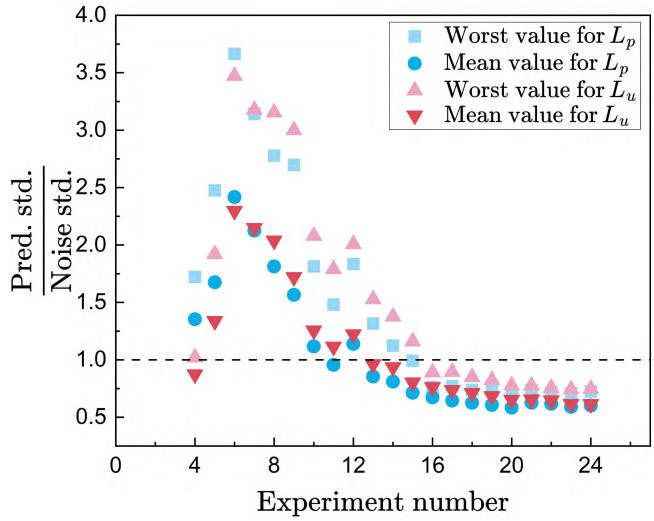

Figure 3. Convergence of the AL method in terms of the ratio of standard deviation of prediction to that of noise.

As shown in Fig. 3, for each output (L_p and L_u), the ratio of GP predictive standard deviation to the corresponding observation noise standard deviation was used as a performance indicator of the AL method. Specifically, the worst and mean values of this ratio evaluated

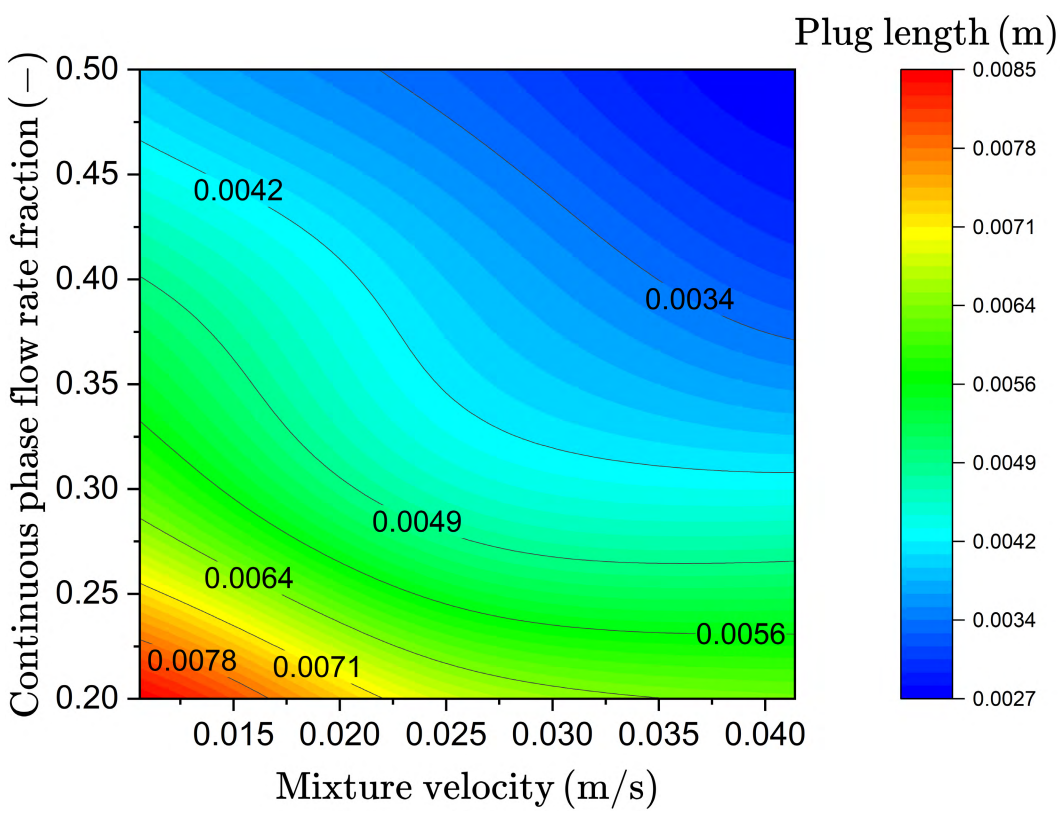

Figure 4. Prediction of plug length by the GP model developed using constrained AL method. The numbers on the contour lines represent the predicted values of plug length.

posteriori (after each design and execution) across the design space were used as the performance indicators. As expected, the ratio decreased gradually as new informative training point was picked in each iteration of the AL method. Ideally, one can stop performing the AL experimental campaign when the ratio of the standard deviations drops below 1 (marked by the dashed line in Fig. 3), indicating that the uncertainty in predicting L_p and L_u by the GP model can be fully explained by the uncertainty in the random nature of measurements.

Besides the prediction uncertainty, the prediction behaviour of the final GP model (after 20 AL experiments) developed through the constrained AL method is illustrated in Fig. 4.

As shown in Fig. 4, the plug length increases with decreasing mixture velocity and continuous phase flow rate fraction. A decrease in the continuous phase flow fraction increases the proportion of the dispersed phase in the channel, resulting in larger plugs. Also, as the mixture velocity decreases, the drag force which helps detach the plugs at the inlet decreases, again producing longer plugs (Garstecki et al., 2006). The behaviour of the GP model is quite appealing as its predictions of the plug length align well with known physical phenomena affecting plug formation and its size.

CONCLUSION

A hybrid modelling approach has been proposed to develop a digital twin model of a segmented flow LLE process to recover critical metals. In the proposed method, a lumped mass transfer model was developed to decouple the mass transfer and hydrodynamic models. Furthermore, the proposed approach enables the closed-loop development of both hydrodynamic and mass transfer models via GP-based AL and MBDoE, respectively. When prior constraints regarding flow regimes or hydrodynamic characteristics are available, this knowledge can be incorporated to guide the development of the GP model describing hydrodynamic features such as plug and slug lengths. Although the approach decouples the hydrodynamic and mass transfer components of the segmented flow model, it allows both models to be developed simultaneously through multi-objective optimization of AL and MBDoE. The approach was tested on an *in silico* case study and proved to be very efficient in developing a combined data-driven and knowledge-based model of segmented flow LLE in small channels.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the funding received from the UK Engineering and Physical Sciences Research Council (EPSRC) through the grant EP/Z533798/1, which supported this work.

AUTHOR IDENTIFIERS

Author ORCIDs:
Pankajakshan A: 0000-0002-3365-2668
Katsoulas K: 0009-0004-9833-0127
Olasinde M: 0000-0001-6001-1404
Fraga ES: 0000-0002-5732-6082
Angeli P: 0000-0001-9575-4259
Galvanin F: 0000-0002-3296-1581

NOMENCLATURE

a	Specific interfacial area
C_{aq}	Concentration of metal in the aq. phase
D	Distribution coefficient
d_{int}	Channel internal diameter
f	Latent function
K_{aq}	Aqueous phase mass transfer coefficient
k'	Lumped mass transfer coefficient
L_p, L_s, L_u	Plug, slug and unit cell length
L_z	Channel length
Q	Volumetric flow rate
u	Input variable
u_{mix}	Mixture velocity
y	Output or response variable
z	Axial coordinate
δ	Film thickness
ε_d	Dispersed phase holdup
κ^f	Kernel function for cross-covariance
κ_u	Kernel function for input covariance
σ	Standard deviation
$\mathbf{H}_\theta$	Observed Fisher information matrix
$\mathbf{K}$	Covariance matrix of the GP kernel
$\mathbf{m}$	Mean vector of the GP kernel
$\mathbf{S}$	Parameter sensitivity matrix
$\mathbf{V}_\theta$	Parameter covariance matrix
θ	Model parameter vector

REFERENCES

1. Williams O. A switched gaussian process for estimating disparity and segmentation in binocular stereo. Advances in Neural Information Processing Systems 19 :1497-1504 (2007). https://doi.org/10.7551/mitpress/7503.003.0192
2. Department for Business & Trade. (2026, January 23). *Vision 2035: Critical Minerals Strategy*. https://www.gov.uk/government/publications/uk-critical-minerals-strategy/vision-2035-critical-minerals-strategy
3. Garciadiego Ortega E, Tsaoulidis D, Angeli P. Predictive model for the scale-out of small channel

two-phase flow contactors. Chemical Engineering Journal 351:589-602 (2018). https://doi.org/10.1016/j.cej.2018.06.020
4. Garciadiego-Ortega E, Tsaoulidis D, Pineda M, Fraga ES, Angeli P. Hydrodynamics and mass transfer in segmented flow small channel contactors for uranium extraction. Chemical Engineering and Processing - Process Intensification 153:107921 (2020). https://doi.org/10.1016/j.cep.2020.107921
5. Garstecki P, Fuerstman MJ, Stone HA, Whitesides GM. Formation of droplets and bubbles in a microfluidic t-junction—scaling and mechanism of break-up. Lab Chip 6:437 (2006). https://doi.org/10.1039/b510841a
6. Geremia M, Macchietto S, Bezzo F. A review on model-based design of experiments for parameter precision – open challenges, trends and future perspectives. Chemical Engineering Science 319:122347 (2026). https://doi.org/10.1016/j.ces.2025.122347
7. Lewis DD, Catlett J. Heterogeneous uncertainty sampling for supervised learning. Machine Learning Proceedings 1994 :148-156 (1994). https://doi.org/10.1016/b978-1-55860-335-6.50026-x
8. Ljung, L. (1999). *System Identification: Theory for the User* (T. Kailath, Ed.; Second Edition). Prentice Hall PTR.
9. Mac Giolla Eain M, Egan V, Punch J. Film thickness measurements in liquid–liquid slug flow regimes. International Journal of Heat and Fluid Flow 44:515-523 (2013). https://doi.org/10.1016/j.ijheatfluidflow.2013.08.009
10. Sun Z, Cao H, Xiao Y, Sietsma J, Jin W, Agterhuis H, Yang Y. Toward sustainability for recovery of critical metals from electronic waste: the hydrochemistry processes. ACS Sustainable Chem. Eng. 5:21-40 (2016). https://doi.org/10.1021/acssuschemeng.6b00841

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Life Cycle Modeling towards Regional Symbiosis for Valorizing Mixed-Lignocellulosic Biomass from Agriculture and Forestry

Yasunori Kikuchi[abcd]*, Nobuhide Takahashi[e], Satoshi Ohara[bdf]

[a] the University of Tokyo, Institute for Future Initiatives, Bunkyo-ku, Tokyo, Japan
[b] the University of Tokyo, Department of Chemical System Engineering, Bunkyo-ku, Tokyo, Japan
[c] the University of Tokyo, Presidential Endowed Chair for "Platinum Society", Bunkyo-ku, Tokyo, Japan
[d] the University of Tokyo, UTokyo LCA Center for Future Strategy, Meguro-ku, Tokyo, Japan
[e] Faculty of Textile Science and Technology, Shinshu University, Ueda, Japan
[f] the University of Tokyo, Research Center for Advanced Science and Technology, Tokyo, Japan
* Corresponding Author: ykikuchi@ifi.u-tokyo.ac.jp.

ABSTRACT

Regional deployment of bioenergy and bio-based products is often constrained by the seasonality, heterogeneity, and dispersed availability of lignocellulosic biomass. This work demonstrates a computer-aided process engineering (CAPE) workflow that integrates experimental characterization, process modeling, and life cycle assessment (LCA) to support regional symbiosis design using mixed feedstocks from agriculture and forestry. A case study is developed for Tanegashima, a remote Japanese island where unused woody residues and sugarcane bagasse are locally available but temporally mismatched. Torrefaction is modeled in an autothermal configuration: char is the main product, while torrefaction gas and condensables are recovered for internal heat supply and any excess is treated as an energy coproduct. Laboratory measurements (220–400°C, 20°C interval) provide temperature-dependent yields of char, tar, aqueous condensate, and gas, alongside ultimate analysis and heating values of solids and representative gas composition. These data are translated into a mass–energy inventory for LCA foreground modeling up to pelletization. The functional unit is 1 t of wet biomass (40 wt% moisture) processed on-island. Results show that increasing temperature decreases solid yield but increases charification and heating value. Autothermal operation becomes feasible above approximately 340°C under a mid-level drying-energy assumption, and the threshold is highly sensitive to drying performance. Across assessed feedstocks and temperatures, life-cycle GHG emissions are negative when credits are given for substituting fossil coal and kerosene-equivalent fuels. The proposed workflow enables composition- and seasonality-aware screening of regional biomass portfolios under consistent inventory assumptions.

Keywords: torrefaction, mixed lignocellulosic biomass, autothermal operation, CAPE, life cycle assessment

INTRODUCTION

Utilizing lignocellulosic biomass from forestry and agriculture can contribute to decarbonization and energy security, particularly in regions that rely on imported fossil fuels. In Japan, expanding the use of domestic woody biomass is discussed not only as a renewable energy measure but also as a route to revitalize forestry and local economies [1]. Realizing these benefits requires long-term alignment between forest management, labor availability, and roundwood demand; recent modeling work explicitly treats this as a "forest transition" problem that must be planned over multi-decade horizons [2]. At regional scale, the feasibility of bioenergy systems is therefore shaped as much by supply-side constraints as by conversion technology performance [3, 4].

Alongside forestry resources, agricultural residues can provide concentrated feedstocks with established

collection points. Sugarcane bagasse is produced at sugar mills during the crushing season, and its use has been studied as part of coordinated agricultural–industrial systems design [5] and as a lever for reducing greenhouse gas (GHG) emissions in sugar/ethanol production systems [6]. However, residue supply is seasonal and storage is not trivial: large bagasse stockpiles can pose risks related to dust emissions, leachate, and spontaneous heating [7]. In island contexts such as Tanegashima, where agricultural and forestry residues coexist but are temporally mismatched, combining feedstocks can help stabilize plant operation while improving local resource circulation [3, 4, 8].

Thermochemical pretreatment technologies are attractive for heterogeneous biomass because they can stabilize the feedstock and improve fuel quality. Torrefaction is an oxygen-deficient thermal treatment typically carried out over 200–300°C, although practical systems may operate into higher temperatures where mild pyrolysis becomes relevant. Reviews report that torrefaction reduces moisture and oxygen content, increases energy density, improves grindability, and enhances hydrophobicity, which can reduce biological degradation during storage and facilitate pelletization and co-firing [9]. Techno-economic assessments further highlight that process performance strongly depends on reactor configuration and heat integration, and that overall efficiency hinges on how moisture removal and volatile utilization are handled [10].

From a life-cycle perspective, torrefaction can improve energy efficiency per delivered megajoule but can also add process energy demand and associated impacts. In particular, the energy required for drying and water removal is frequently a dominant contributor and uncertainty driver in torrefied pellet life-cycle assessments [11, 12]. Autothermal operation—supplying process heat primarily from internally generated torrefaction gas and condensables ("torgas" and related volatiles)—is often proposed as a design target because it can reduce or eliminate external fuel demand. Sensitivity analyses indicate that the extent of by-product utilization is critical for both emissions and economics [12].

Computer-aided process engineering (CAPE) tools can support these decisions by generating consistent mass–energy inventories for alternative pathway options and by enabling composition-aware screening. Recent work demonstrated simulation-based design for regional biomass thermochemical conversion systems, integrating process modeling and LCA for improved environmental and socio-economic performance [13]. Building on this foundation, the present study focuses on autothermal dry torrefaction as a key pretreatment step for mixed lignocellulosic biomass and illustrates how experimental yield data can be translated into foreground inventories suitable for LCA and for future superstructure-based regional design.

The objective of this paper is to demonstrate a CAPE-supported workflow for designing regionally appropriate lignocellulosic resource utilization systems using Tanegashima as an illustrative case. Specifically, we (i) characterize temperature-dependent thermal decomposition behavior and product yields for wood chips and bagasse (220–400°C); (ii) formulate an autothermal torrefaction process model that converts experimental yields into a consistent mass–energy inventory up to pelletization; and (iii) evaluate life-cycle GHG emissions (LC-GHG) across feedstocks, temperatures, and operating assumptions, highlighting the conditions under which autothermal operation is feasible.

MATERIALS AND METHODS

Introduction to case study: Tanegashima

Tanegashima is a remote island in Kagoshima Prefecture, Japan, characterized by an islanded energy system and a regional economy where agriculture and forestry coexist. Prior studies have analyzed industrial symbiosis centered on a regional combined heat and power (CHP) plant using local resources [8] and have shown that designing the appropriate plant scale requires careful matching between sustainably available woody biomass supply and local heat/electricity demand [3]. A co-learning case study further emphasized that aligning technology choices with local constraints (resource availability, stakeholder preferences, and infrastructure) is essential for regional transformation [4].

In addition to forestry residues, Tanegashima hosts a sugar-mill industry that produces sugarcane bagasse as a coproduct. Coordinated integration of agricultural and industrial processes has been studied for sugarcane-derived production systems, highlighting how cultivar choice and mill operation influence material and energy flows [5]. Bagasse supply is concentrated in the crushing season, while forestry residues can be available more continuously but depend on long-term forest management and logistics. The coexistence of these two feedstocks motivates a mixed-biomass utilization strategy to stabilize conversion facilities and improve regional resilience.

Thermal decomposition experiments and yield characterization

Thermal decomposition behavior was characterized under controlled inert conditions to provide process-relevant foreground data. Isothermal experiments were conducted using a thermogravimetric analyzer (TGA-51) with nitrogen carrier gas at 100 mL/min. Samples were heated to target temperatures between 220 and 400°C (20°C interval) and held isothermally for 120 min to observe mass-yield trajectories [14, 15]. These

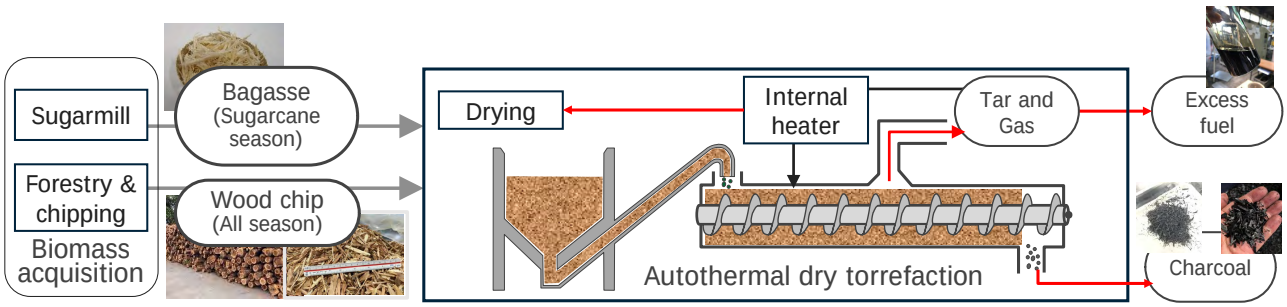

Figure 1. Life cycle boundary for autothermal dry torrefaction of sugarcane bagasse and wood chip as regional lignocellulosic biomass.

Table 1: Settings of parameters for process modeling.

Parameter	Baseline	Range / cases	Source / note
Functional unit basis	1 t wet biomass	—	40 wt% moisture (wet basis).
Cradle-to-gate inventory	Wood chip		Japan average of Sugi (Cryptomeria japonica) [17]
Cradle-to-gate inventory	Bagasse		Tanegashima average [6]
Drying energy factor (water removal)	6 MJ/kg-water	3 / 6 / 9 MJ/kg-water	Drying sensitivity per torrefied pellet LCA [11].
Torrefaction temperature	Variable	220–400°C (20°C interval)	Yield functions fitted to experiments (Section 2.2).
Heat recovery efficiency	0.85	0.7–0.9	Aggregate efficiency for by-product to useful heat.
Electricity: grinding + pelletization	Literature-based	Updated when measured	Foreground parameter; informed by pellet/co-firing LCAs [11, 17].
Impact category	LC-GHG (GWP100)	—	IPCC AR4 factors [19].
Background data	IDEA v3.5.1	—	Japan LCI database (industry standard).

measurements provide a consistent basis for mapping temperature-dependent solid yield and for supporting yield-function fitting.

In addition, bench-scale carbonization experiments in an electric tubular furnace were performed for Japanese cedar wood chips and sugarcane bagasse sourced from Tanegashima to quantify product distributions. For each temperature, mass yields of solid char, tar, aqueous condensate (wood-vinegar-like fraction), and non-condensable gas were determined. Ultimate analysis (C/H/O) and HHV/LHV were obtained for raw feedstock and char, and representative gas composition (CO_2/CO/CH_4) was measured by gas chromatography (GC). These data allow estimation of by-product energy available for internal heat supply.

A mass–energy balance was constructed using measured compositions. For solids, carbon, hydrogen, and oxygen contents were directly measured. For the gas phase, the measured species were used to estimate the combustible fraction and the lower heating value of the gas mixture. Liquid products were separated into tar and aqueous fractions; because comprehensive speciation was outside the scope of this proceeding, their aggregate energy content was inferred from the overall energy balance using measured HHV/LHV of the feed and char and the estimated gas energy.

Although these experiments were performed at laboratory scale with inert carrier gas and differ from industrial reactors in heat and mass transfer, the measured yields and compositions are valuable anchors for process modeling. In the process model (Section 2.3), we treat experimental yields at 30 mins as primary inputs and represent scale-up-relevant uncertainties (e.g., drying energy and heat losses) through scenario ranges consistent with LCA literature [11].

Autothermal dry torrefaction process modeling

Figure 1 illustrates the process boundary for the modeled torrefaction system. The target configuration is autothermal dry torrefaction, in which process heat is supplied predominantly by combusting internally generated torrefaction gas and condensable by-products. The

model includes biomass acquisition and preprocessing (chipping/handling and transport), drying, torrefaction (220–400°C), solid cooling, grinding, and pelletization. Char is the main product, while by-products are used for internal heat demand and any surplus is treated as an energy coproduct.

Experimental temperature-dependent yields (char, tar, aqueous, and gas) are mapped to a steady-state mass flow model for a reference feed input. **Table 1** summarizes key modeling parameters and scenario settings (moisture content, drying energy per kg evaporated water, electricity demand for size reduction and pelletization, and heat recovery efficiency). Drying is represented by reducing feed moisture from the functional-unit basis (40 wt% wet basis) to a target moisture suitable for torrefaction. The required drying heat is calculated from the latent heat of vaporization and an "effective drying energy" factor that implicitly captures sensible heating, exhaust losses, and heat recovery performance; this factor is varied in three cases to reflect uncertainty (Section 2.4).

Autothermal feasibility is assessed by comparing recoverable by-product energy with process thermal demand. Let Q_demand be the sum of drying heat, reactor sensible/latent duties, and heat losses. Let Q_byprod be the usable heat from by-product combustion after accounting for combustion efficiency and heat recovery (Table 1). Autothermal operation is considered feasible if Q_byprod ≥ Q_demand (electricity excluded). This criterion provides a transparent link between measured yield distributions (which determine Q_byprod) and design/operation choices (which determine Q_demand).

Electricity use is included for material handling, grinding, and pelletization. When site-specific measurements are unavailable, electricity intensities are adopted from torrefied pellet and co-firing LCA studies and treated as scenario inputs [11, 16, 17]. This approach keeps the process model modular: measured thermochemical yields drive thermal balances, while mechanical-electric demands can be updated as site-specific data become available.

Life cycle assessment (LCA)

The goal of the LCA is to quantify life-cycle greenhouse gas emissions associated with producing charcoal via autothermal dry torrefaction on Tanegashima and to compare outcomes across feedstocks (wood chip, bagasse, and mixed operation) and operating temperatures. The system boundary (**Figure 1**) covers biomass acquisition and preprocessing, drying, torrefaction, pelletization, and the use of products that substitute fossil fuels. The functional unit is the processing of 1 t of wet biomass (40 wt% moisture) delivered to the torrefaction system.

Foreground models for forestry operations and logistics are aligned with prior Tanegashima studies that evaluated woody biomass CHP and regional symbiosis under local supply constraints [3, 8]. For sugarcane bagasse, upstream burdens and allocation within the sugar-mill system follow Japanese case studies on sugarcane-derived production and GHG reduction options [5, 6]. Background processes for fuels, electricity, transport, and materials are taken from IDEA, and island-specific electricity assumptions can be applied following earlier Tanegashima LCA practice [8].

LC-GHG is calculated as the sum of upstream and on-site emissions minus substitution credits. Substitution is modeled on an energy basis: charcoal displaces general coal, while any externally available by-product energy displaces kerosene-equivalent fuel. This consequential framing is used to highlight mitigation potential in an island context where imported fossil fuels dominate. Sensitivity analysis is performed by varying the effective drying energy factor (3/6/9 MJ/kg-water) following the approach of torrefied pellet LCAs that explicitly tested drying assumptions [11].

RESULTS AND DISCUSSION

Process modeling results: yield and fuel quality trends

For both Japanese cedar wood chips and sugarcane bagasse, experimental results show consistent qualitative trends with increasing temperature. Solid yield decreases as temperature increases due to progressive devolatilization. In parallel, the degree of charification increases: carbon fraction rises, oxygen fraction falls, and the solid heating value increases. These patterns are consistent with the broader torrefaction and mild pyrolysis literature [9, 10].

Comparing feedstocks at the same temperature, wood chips generally exhibit higher solid yield, while bagasse tends to reach a higher charification degree (higher C fraction and lower O/C) at comparable severity. Such differences can be attributed to compositional and ash effects that influence devolatilization pathways. This observation reinforces the value of feedstock-specific (and ultimately composition-based) yield models for reliable process screening.

By-product streams (gas and condensables) increase with severity and underpin autothermal feasibility. Measured gas composition indicates that combustible species (CO, CH_4) become more prominent at higher temperatures, raising gas LHV and improving suitability for internal heat generation. When mapped into the process model, this shift increases Q_byprod, while decreasing solid yield reduces the energy retained in the main product. The trade-off between solid yield and autothermal heat availability therefore defines an operational "sweet spot" that depends strongly on drying and heat recovery assumptions.

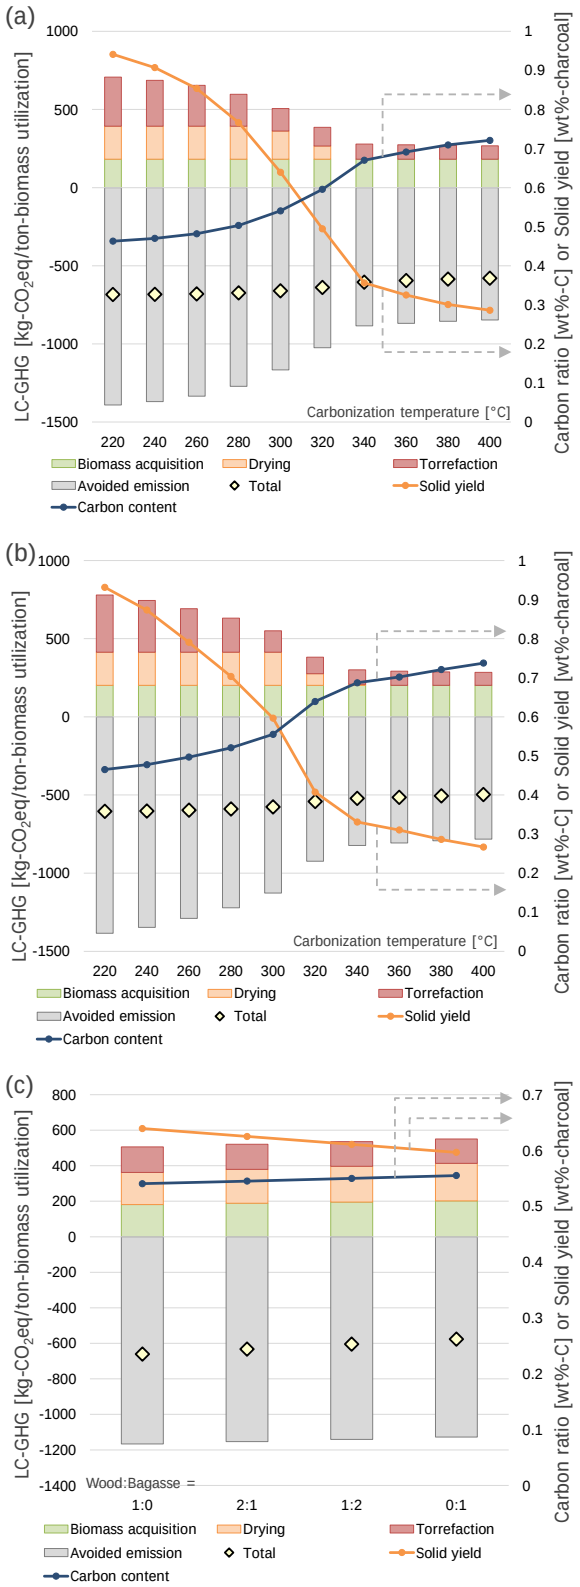

Figure 2. LC-GHG results (a) wood chip, (b) sugarcane bagasse, (c) combination of wood chip and bagasse.

Life cycle GHG results

Figure 2 summarizes LC-GHG results for (a) wood chips only, (b) bagasse only, and (c) combined operation.

For each case, the bar chart shows total LC-GHG per functional unit and the line plots indicate solid yield and degree of charification. The wood-only and bagasse-only cases show similar LC-GHG magnitudes across temperatures, suggesting that upstream differences are partially offset by differences in yield and product energy content when credits are applied.

A key operational result is the temperature threshold for autothermal feasibility. Under the baseline (middle) drying-energy case, autothermal operation—meeting all process heat demand with internally generated by-products (excluding electricity)—becomes feasible above approximately 340°C. Below this range, by-product energy is insufficient to cover drying and reactor duties, implying an external heat requirement. This threshold shifts with drying assumptions, consistent with the strong sensitivity to drying energy reported in prior torrefied pellet LCAs [11].

Across all assessed temperatures and feedstocks, total LC-GHG is negative when displacement credits are applied. The dominant driver is avoided emissions from coal substitution by charcoal on an energy basis. This result highlights the mitigation potential of converting locally available biomass into a storable, transportable solid fuel in an island context. It also clarifies the modeling dependence on the chosen use and substitution assumptions, motivating transparent reporting and scenario testing when LCAs are used for decision support.

Discussion: implications for CAPE-supported regional design

The experimental mass–energy balances imply that, under baseline drying assumptions, autothermal feasibility is not obtained within the conventional torrefaction band (220–300°C). This does not necessarily mean that autothermal torrefaction is infeasible in practice. Instead, it emphasizes that drying and heat integration dominate the energy balance and that laboratory yields must be interpreted together with realistic process-scale drying performance and heat recovery. Adams et al. [11] explicitly varied drying energy in three cases and showed that drying assumptions can substantially change comparative outcomes; similarly, McNamee et al. highlighted that using torgas for process heat is a major lever for emission savings [12]. For Tanegashima, practical strategies such as air-drying prior to mechanical drying, using waste heat, or co-locating with heat sources could lower effective drying energy and shift the autothermal threshold.

Beyond autothermal operation, the by-product streams expand the design space for regional symbiosis. If process heat is supplied by an additional renewable source, gas and condensables need not be fully consumed internally and could be recovered as energy or upgraded products. In the Tanegashima context, demonstration activities have explored the use of torrefaction-derived fuels in practical energy applications; for example, finely ground char slurried with diesel (biomass-oil-mixture, BOM) was combusted in a commercial diesel engine to generate electricity [20]. While BOM fuel requires blending with fossil diesel and is therefore not a fully defossilized route, such demonstrations illustrate the value of flexible product portfolios and the need to evaluate multiple utilization options alongside LCA results.

Feedstock seasonality and storage are practical constraints. Bagasse is produced seasonally, and stockpiling can lead to dust, leachate, and spontaneous heating risks if unmanaged [7]. Torrefaction offers stabilization by reducing moisture and increasing hydrophobicity, enabling longer-term storage and more flexible use. For woody residues, mobilization depends on sustained forest management and logistics, which must be evaluated over long time horizons [2, 3]. A mixed-feedstock strategy can therefore improve plant utilization factors while sharing logistics and infrastructure across sectors, consistent with regional symbiosis concepts observed in Tanegashima case studies [4, 8].

From a CAPE standpoint, the main contribution is the translation of experimentally verified yields into consistent life-cycle inventories suitable for superstructure-based screening. By expressing conversion performance as temperature- and feedstock-dependent functions, the model enables rapid evaluation of alternative operating points, feedstock mixes, and product allocations under consistent assumptions. While this proceeding focuses on LC-GHG, the same inventories can be extended to additional impact categories, uncertainty propagation, and multi-objective trade-offs, supporting future CAPE-driven regional planning.

CONCLUSION

This paper demonstrated a CAPE-supported workflow for designing and evaluating regional lignocellulosic biomass utilization based on autothermal dry torrefaction, using Tanegashima as a case study and focusing on Japanese cedar wood chips and sugarcane bagasse. Laboratory experiments (220–400°C) provided temperature-dependent yields and fuel-quality indicators for char and by-products, which were translated into a process inventory up to pelletization.

Process modeling and LCA showed a clear trade-off between solid yield and charification with increasing temperature. Under a baseline drying-energy assumption, autothermal operation becomes feasible above approximately 340°C, and the threshold is highly sensitive to drying performance. With substitution credits for coal and kerosene-equivalent fuels, LC-GHG results are negative across assessed feedstocks and temperatures, indicating substantial mitigation potential in an island context.

The results underline that drying energy and heat integration are decisive factors for both autothermal feasibility and LC-GHG outcomes. Future work will refine site-specific drying and heat-recovery assumptions, integrate the torrefaction model into a broader regional superstructure, and evaluate multi-product portfolios and seasonal feedstock allocation strategies with combined environmental and socio-economic objectives.

ACKNOWLEDGEMENTS

This work was supported by MEXT/JSPS KAKENHI (Grant Number JP23K11521), and JST COI-NEXT JPMJPF2003. Activities of the Presidential Endowed Chair for "Platinum Society" at the University of Tokyo are supported by Mitsui Fudosan Corporation, Sekisui House, Ltd., the East Japan Railway Company, and Toyota Tsusho Corporation.

AUTHOR IDENTIFIERS

Author ORCIDs:
Kikuchi Y: 0000-0003-0002-8233
Takahashi N: 0000-0001-9157-5958
Ohara S: 0000-0002-5273-8969

REFERENCES

1. Ministry of Agriculture, Forestry and Fisheries (MAFF). The situation surrounding the use of biomass. https://www.maff.go.jp/e/policies/env/attach/pdf/biomass-6.pdf
2. Han Z, Oosawa K, Kanematsu Y, Teah HY, Kikuchi Y. Simulating forest transition for sustainable forestry planning in japan: a backcasting model of roundwood supply and carbon stocks with operational and regional considerations. Resources, Conservation & Recycling Advances 29:200314 (2026). https://doi.org/10.1016/j.rcradv.2026.200314
3. Kanematsu Y, Oosawa K, Okubo T, Kikuchi Y. Designing the scale of a woody biomass CHP considering local forestry reformation: a case study of tanegashima, japan. Applied Energy 198:160-172 (2017). https://doi.org/10.1016/j.apenergy.2017.04.021
4. Kikuchi Y, Nakai M, Kanematsu Y, Oosawa K, Okubo T, Oshita Y, Fukushima Y. Application of technology assessments to co-learning for regional transformation: a case study of biomass energy systems in tanegashima. Sustain Sci 15:1473-1494 (2020). https://doi.org/10.1007/s11625-020-00801-1
5. Ouchida K, Kanematsu Y, Fukushima Y, Ohara S, Sugimoto A, Hattori T, Terajima Y, Okubo T, Kikuchi Y. Coordinated integration of agricultural and industrial processes: a case study of sugarcane-derived production. Process Integr Optim Sustain 7:1191-1209 (2023). https://doi.org/10.1007/s41660-023-00337-8
6. OHARA S, KIKUCHI Y, OUCHIDA K, SUGIMOTO A, HATTORI T, YASUHARA T, FUKUSHIMA Y. Reduction of greenhouse gas emissions in the introduction of inversion system to produce sugar and ethanol from sugarcane. Journal of Life Cycle Assessment, Japan 15:86-100 (2019). https://doi.org/10.3370/lca.15.86
7. Rainey TJ, O'Hara IM, Mann AP, Bakir CH, Plaza F. Effect of depithing on the safety and environmental aspects of bagasse stockpiling. Process Safety and Environmental Protection 91:378-385 (2013). https://doi.org/10.1016/j.psep.2012.09.002
8. Kikuchi Y, Kanematsu Y, Ugo M, Hamada Y, Okubo T. Industrial symbiosis centered on a regional cogeneration power plant utilizing available local resources: a case study of tanegashima. J of Industrial Ecology 20:276-288 (2015). https://doi.org/10.1111/jiec.12347
9. van der Stelt MJC, Gerhauser H, Kiel JHA, Ptasinski KJ. Biomass upgrading by torrefaction for the production of biofuels: a review. Biomass and Bioenergy 35:3748-3762 (2011). https://doi.org/10.1016/j.biombioe.2011.06.023
10. Batidzirai B, Mignot APR, Schakel WB, Junginger HM, Faaij APC. Biomass torrefaction technology: techno-economic status and future prospects. Energy 62:196-214 (2013). https://doi.org/10.1016/j.energy.2013.09.035
11. Adams PWR, Shirley JEJ, McManus MC. Comparative cradle-to-gate life cycle assessment of wood pellet production with torrefaction. Applied Energy 138:367-380 (2015). https://doi.org/10.1016/j.apenergy.2014.11.002
12. McNamee P, Adams PWR, McManus MC, Dooley B, Darvell LI, Williams A, Jones JM. An assessment of the torrefaction of north american pine and life cycle greenhouse gas emissions. Energy Conversion and Management 113:177-188 (2016). https://doi.org/10.1016/j.enconman.2016.01.006
13. Corradini LL, Heiho A, Kanematsu Y, Shimono R, Ohara S, Kikuchi Y. Simulation-based design of regional biomass thermochemical conversion system for improved environmental and socio-economic performance. Comput Aided Chem Eng 52:2361-2366 (2023) https://doi.org/10.1016/B978-0-443-15274-0.50375-9
14. Matsumoto Y, Shimada I, Osada M, Fukunaga H, Takahashi N. Prediction of changes in mass yield

and elemental composition of Sugi (Cryptomeria japonica) in the course of pyrolysis based on pyrolysis behavior of the main components of biomass. 35th Annual Meeting of The Japan Institute of Energy (2025).
15. Segawa H, Sasaki T, Shimada I, Osada M, Takahashi N, Fukunaga H. Effect of pore-forming agent on concentration overpotential of carbon-free pt electrodes for PEFC. Meet. Abstr. MA2024-02:4700-4700 (2024). https://doi.org/10.1149/ma2024-02674700mtgabs
16. Tsalidis GA, Joshi Y, Korevaar G, de Jong W. Life cycle assessment of direct co-firing of torrefied and/or pelletised woody biomass with coal in the netherlands. Journal of Cleaner Production 81:168-177 (2014). https://doi.org/10.1016/j.jclepro.2014.06.049
17. Cho K, Lee Y. Life cycle assessment of torrefied residual biomass co-firing in coal-fired power plants: aspects of carbon dioxide emission. Energies 17:6165 (2024). https://doi.org/10.3390/en17236165
18. Nakano K, Shibahara N, Nakai T, Shintani K, Komata H, Iwaoka M, Hattori N. Greenhouse gas emissions from round wood production in japan. Journal of Cleaner Production 170:1654-1664 (2018). https://doi.org/10.1016/j.jclepro.2016.10.024
19. Intergovernmental Panel on Climate Change (IPCC). Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the IPCC. Cambridge University Press (2007).
20. IFI-UTokyo (Institute for Future Initiatives, the University of Tokyo). Co-JUNKAN Project – Successfully demonstrated the combustion of a novel biofuel derived from Tanegashima's biomass in a commercial diesel engine at JAXA's Okinawa Tracking and Communications Station to generate electricity. https://ifi.u-tokyo.ac.jp/en/news/13330/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Design and Assessment of Regional Symbiosis: A Case Study of Plant-oil Production in Japan

Yukito Watanabe[a], Yuichiro Kanematsu[b], Hiro Tabata[b], Heng Yi Teah[b], Kousuke Hiromori[c], and Yasunori Kikuchi[abd]*

[a] Department of Chemical System Engineering, The University of Tokyo, Bunkyo-ku, Tokyo, Japan
[b] Presidential Endowed Chair for "Platinum Society", The University of Tokyo, Bunkyo-ku, Tokyo, Japan
[c] Department of Chemical Engineering, Tohoku University, Aoba-ku, Sendai, Japan
[d] Institute for Future Initiatives, The University of Tokyo, Bunkyo-ku, Tokyo, Japan
* Corresponding Author: ykikuchi@ifi.u-tokyo.ac.jp

ABSTRACT

This study conducted a life cycle assessment to assess and design regional symbiosis at plant-oil production. These industries face challenges including dependence on fossil fuels and the generation of underutilized by-products, while effective regional symbiosis requires the selection of diverse regional unused resources and assessment based on process models that consider future technological prospects. Mathematical models for plant-oil production were developed using industrial data from literature to calculate inventory data. The case study showed that introducing woody biomass combined heat and power reduced GHG emissions by 8% in the Cradle-to-Grave system boundary, while recycling technology for soap stock using Kolbe electrolysis achieved a 3% reduction. Regional analysis indicated that 33 prefectures in Japan could meet woody biomass demand through sustainable forestry management, potentially reducing GHG emissions in Japan by approximately 0.041%. These results suggest that regional symbiosis can enhance resource-use efficiency, reduce environmental impacts, and contribute to sustainable regions.

Keywords: Life Cycle Analysis, Modelling, Biomass, Energy, Materials

INTRODUCTION

Japan has set a governmental target to achieve net-zero greenhouse gas emissions by 2050[1]. In Japan, annual greenhouse gas emission reduction targets have been steadily achieved; however, further reduction of environmental impacts will require mitigation efforts not only in large-scale industries such as power generation, petrochemicals, and steel production but also in small- and medium-scale industries operating in local regions.

Plant-based production systems, which are widely distributed across local regions of Japan, utilize plant-derived resources as feedstocks and produce plant-based products through processes such as pressing and refining. However, these processes remain highly dependent on fossil-derived energy and often generate large amounts of unused by-products, which represents a challenge for improving environmental impacts [2]. To address this issue, regional symbiosis has been proposed, in which neighboring industries mutually utilize underutilized by-products, energy, and other resources [3]. Life cycle assessment (LCA) has been widely applied to quantitatively evaluate the environmental impacts of such systems [4]. Nevertheless, previous studies have focused on limited industrial processes [5], and evaluation of technological options based on diverse regionally available resources remains insufficient.

The objective of this study is to evaluate and design regional symbiosis enabled by the introduction of new technologies into plant-oil production from a life cycle perspective. To assess the effects of technological implementation, we develop process models for both the existing plant-oil production process and newly introduced technologies and conduct life cycle assessment (LCA). Regional symbiosis can enable the utilization of underutilized regional resources that are difficult to exploit within individual industries, thereby improving resource efficiency and reducing dependence on fossil resources. Such systems have the potential to reduce regional environmental impacts while enhancing the

sustainability and socio-economic viability of local production systems.

In this study, plant-oil production is selected as the target industry. Plant-oil production facilities are widely distributed throughout Japan and already possess infrastructure for handling regional biomass resources, indicating their potential to serve as hubs for regional symbiosis. Previous studies have examined regional symbiosis using LCA, including the conceptual design of energy systems centered on sugar factories [6] and the design of energy hubs for pulp and paper mills [7]. However, evaluations based on process models that consider the selection of diverse regional resources, and the future potential of technological options remain limited. Consequently, design challenges and the environmental benefits associated with technological implementation for regional symbiosis in plant-oil production have not yet been fully clarified. As illustrated in **Figure 1**, this study assumes a regional symbiosis framework in which the introduction of new technologies utilizing regional resources into existing plant-oil production systems promotes resource circulation within the region.

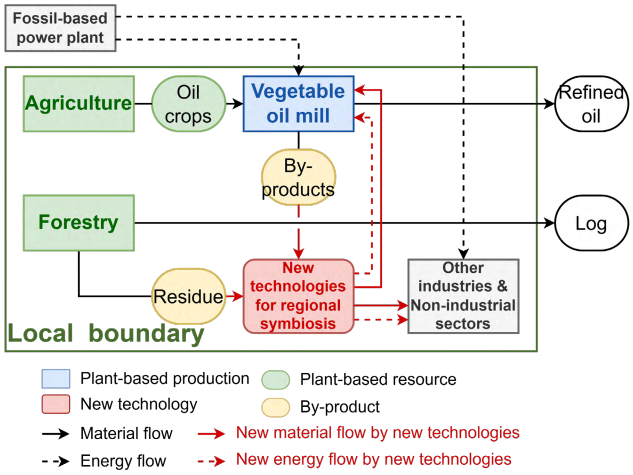

Figure 1. Future visions of regional symbiosis at plant-oil production

METHOD

Life Cycle Assessment

Figure 2 illustrates the system boundary of the LCA case study, which assumes the introduction of a woody biomass combined heat and power (CHP) system using a steam turbine [8] and a resource recovery technology for neutralization residues soap stocks based on the Kolbe electrolysis reaction [9] into conventional plant-oil production. In the LCA, the functional unit was defined as the conversion of 1 kg of oil crops into the refined oil and a series of by-products. The system boundary was set as cradle-to-grave (excluding the use phase). The assessment indicator was greenhouse gas (GHG) emissions.

Regarding the end-of-life (EoL) of the main product (food), processes such as digestion, wastewater treatment, and sludge incineration are assumed; however, for simplicity, all carbon contained in the product was assumed to be uniformly released as CO_2. In contrast, plant resources were assumed to absorb CO_2 during cultivation, and by-products were regarded as substituting for conventional production and were therefore accounted for as negative emissions.

The production process and the newly introduced technologies were treated as foreground processes, and models were developed to enable the calculation of mass and energy balances required for the LCA. The production process model was formulated as a mathematical model capable of reproducing processes of different scales and operating conditions through parameter adjustment, based on empirical data extracted from the literature and interviews with companies. All other processes were treated as background processes, and data were obtained from IDEA ver. 3.5.1, the Japanese background LCA database [10].

Process Modeling of Plant-oil Production

In the plant-oil production process, oil crops are separated into crude oil and oil meal through pressing and extraction. The crude oil then undergoes a degumming process, in which warm water is added to hydrate and remove hydratable phospholipids as gums. Subsequently, phosphoric acid is added to convert non-hydratable phospholipids remaining after water degumming into forms that can be more easily entrained into soapstock in the subsequent neutralization step. In the neutralization process, alkaline agents are added to remove free fatty acids as soapstock. In the bleaching process, pigments and other impurities are removed by adsorption using activated bleaching earth, and in the deodorization process, volatile and odorous compounds are removed by high-temperature steam stripping, resulting in the production of edible refined oil. The oil meal, on the other hand, undergoes a drying process and is obtained as a by-product used for animal feed [11, 12].

A mathematical model was developed based on actual industrial data extracted from the literature, including parameters such as chemical input ratios, separation ratios, and energy intensities. In addition, through interviews with companies, the model was refined to capture the specific characteristics of the industrial processes.

Process Modeling of Emerging Technologies

Woody Biomass CHP

Woody biomass CHP is an energy supply technology that uses wood chips as fuel. In this study, the supplied wood chips were assumed to be produced from underutilized woody resources such as thinned wood and waste wood, and a steam turbine-based CHP system

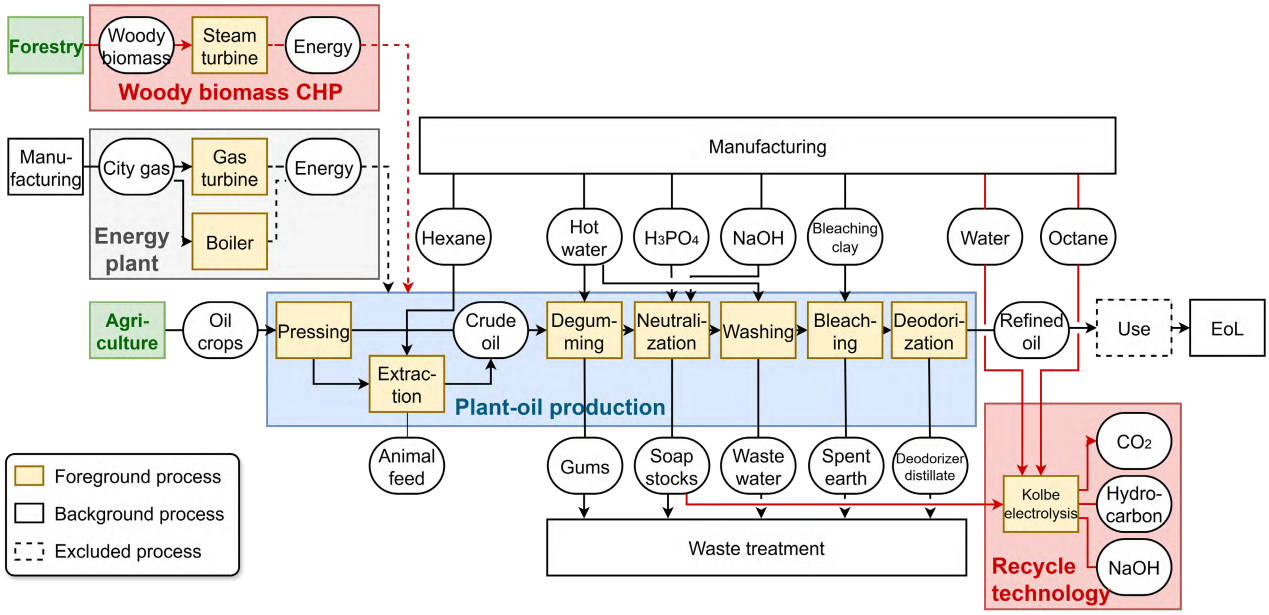

Figure 2. The system boundary for the cradle-to-grave assessment (excluding the refined oil use phase) of plant oil production and technology implementation. The system boundary includes, as conventional processes, oil crop cultivation, plant oil production processes, additive manufacturing, waste treatment, fossil-based energy supply, and EoL processes. In addition, as newly introduced technologies, the system includes woody biomass combined heat and power (CHP) and a recycling technology for neutralization residues using Kolbe electrolysis.
Yellow rectangles indicate foreground processes, while white rectangles represent background processes. The dotted rectangles denote processes excluded from the system boundary. Rounded rectangles indicate avoided processes, which represent conventional production substituted by the system.

was considered. The system was assumed to operate under a heat-led mode to meet the thermal demand of the process. When surplus electricity was generated, it was assumed to be exported to the grid, whereas when electricity was insufficient, it was assumed to be supplied from the grid. The fuel was assumed to be woody chips with a moisture content of 11%, produced from coniferous forests through harvesting, drying, and chipping processes. The electricity generation efficiency of the steam turbine was calculated as a function [13] of the power output, and the required amount of woody biomass was determined based on the energy demand of the process.

Recycle Technology for Neutralization Residue

Soapstock, a residue generated in the neutralization process, poses a disposal challenge for many edible oil producers, and there is a strong demand for technologies that can both reduce environmental burdens and convert this residue into value-added products. At present, a laboratory-scale technology has been developed in which soapstock is diluted with water and subjected to electrolysis followed by separation, enabling the production of sodium hydroxide and fuel-grade hydrocarbons [9]. The overall reaction of the process is expressed by Eq. (1), yielding a 48 wt% aqueous sodium hydroxide solution and long-chain hydrocarbons for fuel, while carbon dioxide and hydrogen are purged as by-products.

$$2RCOONa + 2H_2O \rightarrow RR + 2NaOH + H_2 + 2CO_2$$

Based on experimental data and empirical data from other industrial processes, a conceptual design of the separation system including recycle streams was developed for industrial-scale implementation [14]. Material and energy balances were calculated using a process simulator. For thermal energy analysis, pinch analysis was applied [15], and the Grand composite curve was constructed. The minimum heating requirement derived from the analysis was assumed to represent the additional heat demand associated with the introduction of the resource recovery technology.

RESULTS AND DISCUSSIONS

Process Flow Analysis

Using the mathematical model, the material and energy balances of the plant oil production system were calculated as shown in **Figure 3**. The material flow analysis indicates that approximately 32% of the oil crops are converted into refined oil, while about 54% are converted into animal feed. The flow rate of soapstock, which serves as the feedstock for the resource recovery technology based on Kolbe electrolysis, corresponds to

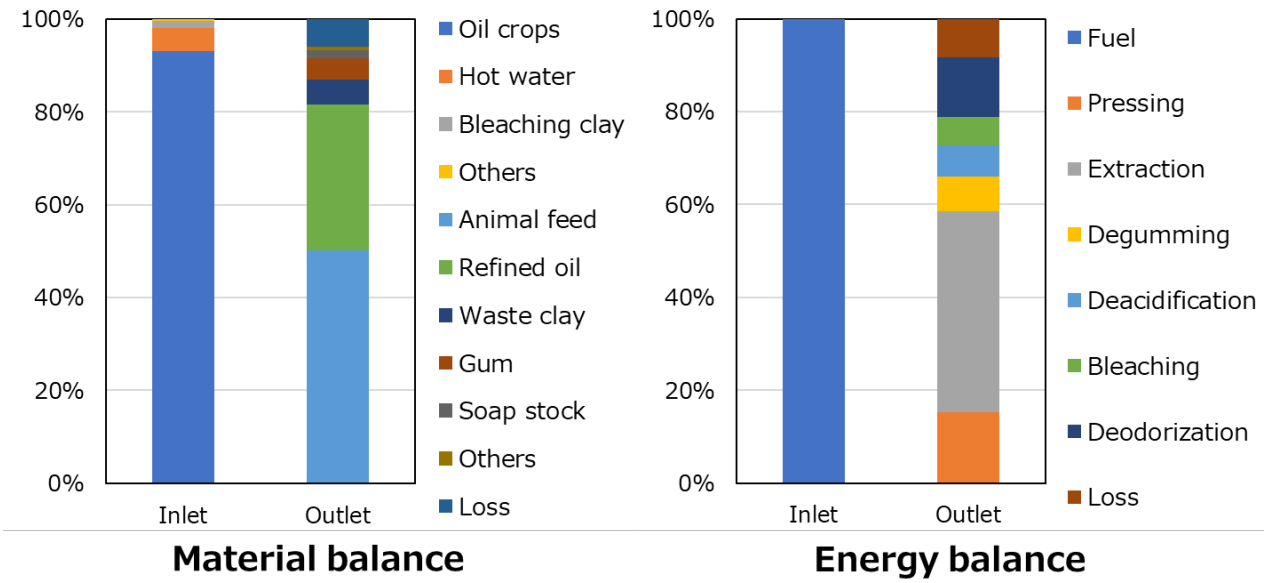

Figure 3. Material (left) and energy (right) balance calculated from flow analysis of plant-oil production based on mathematical model

approximately 1.8% of the oil crop input.

The energy flow analysis reveals that the overall heat-to-power ratio of the plant oil production process is approximately 3.4, indicating a heat-intensive system. A large share of the thermal energy demand arises in the extraction process, with particularly high heat consumption in the drying process for converting oilseed meal into animal feed. Regarding temperature levels of the thermal demand, the deodorization tower operates at approximately 210–240 °C, whereas most of the other processes require relatively low temperatures of around 70–90 °C, at which thermal degradation of plant oil is minimal.

Life Cycle Assessment

Figure 4 presents the GHG emission results for plant oil production with and without technology implementation. The conventional process results in approximately 1.32 kg-CO_2 eq per kg of oil crop. The introduction of woody biomass CHP reduces GHG emissions by 8%, the resource recovery technology by 3%, and the combined implementation of both technologies by 12%.

The reduction achieved by woody biomass CHP can be attributed to the decarbonization of the fuel supply, as the GHG emissions from biomass combustion are offset by CO_2 uptake during the growth of the woody biomass. In addition, because the CHP system is assumed to operate in a heat-led mode, surplus electricity that cannot be fully utilized within the plant oil production process is exported to the grid and substitutes for grid electricity, which is another major contributor to the emission reduction. For the resource recovery technology, soapstock is converted into sodium hydroxide and long-chain hydrocarbons for fuel. These products are assumed to substitute for their conventional counterparts, and the resulting substitution effects are reflected as avoided emissions.

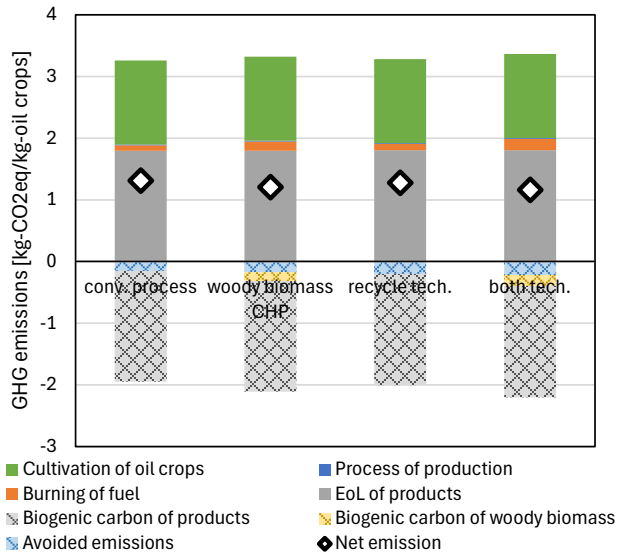

Figure 4. GHG emissions of plant-oil production and introduction of technology

Regional Symbiosis with Forestry in Japan

Next, regional symbiosis between forestry and plant oil production was examined by comparing woody biomass availability and demand on a prefectural basis. The demand was defined as the woody biomass required if all existing plant oil factories in Japan were assumed to adopt the same process as the plant oil production model

and to introduce heat-led woody biomass CHP systems. The demand was calculated based on the plant oil production model and national statistics [16]. The supply, on the other hand, was defined as the required harvesting volume to equalize forest age-class distributions under a sustainable long-term forestry plan [17], assuming that woody biomass is supplied within the same prefecture where plant oil production facilities are located. As shown in **Figure 5**, in regions where the demand-to-supply ratio is less than or equal to 1, excluding Ibaraki, Osaka, and Kagawa etc. prefectures, the woody biomass demand does not exceed the required harvesting volume. This suggests that the utilization of woody biomass CHP can generate stable long-term demand for woody biomass as fuel and may contribute to promoting harvesting and replanting cycles based on long-term forestry plans.

Furthermore, because the CHP systems are assumed to operate in a heat-led mode, surplus electricity is generated from woody biomass CHP depending on the scale of the plant oil production facilities, as not all electricity can be fully consumed within the plants. When this surplus electricity is assumed to be supplied within the local region, it was found that, in all prefectures, the existing electricity demand substantially exceeds the amount of surplus electricity. This indicates that GHG emissions can be reduced by supplying woody biomass-derived energy to both the energy demand of plant oil production and a portion of regional electricity demand, thereby substituting for existing fossil-based energy.

Based on these assumptions, regional symbiosis of plant oil production could achieve an annual GHG emission reduction of approximately 4.19×10^5 t-CO_2 eq nationwide, which corresponds to about 0.041% of Japan's annual GHG emissions [18].

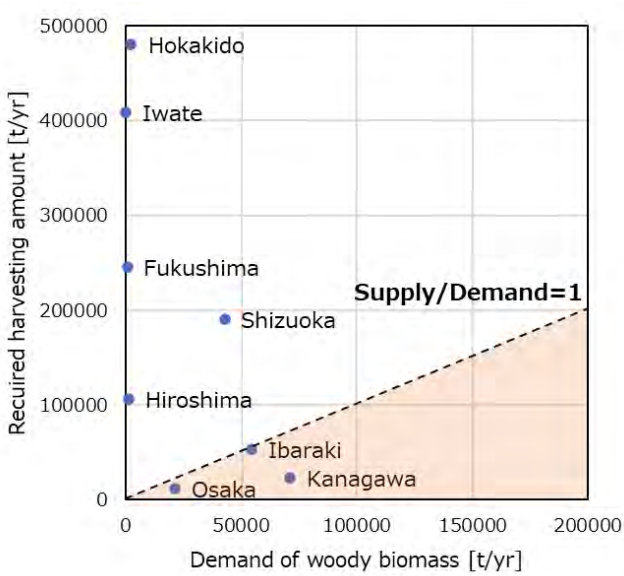

Figure 5. Comparison between woody biomass demand and required harvesting volume in Japan

CONCLUSIONS

In this study, as a case study of regional symbiosis at plant-oil production, mathematical models of a plant oil production process and newly developed technologies were developed to consider their future potential, regional resource supply–demand balance, and geographical constraints. Using these models, LCA was conducted assuming the introduction of woody biomass CHP and a resource recovery technology for neutralization residues in plant oil production. The evaluation confirmed that regional symbiosis of plant-oil production can provide benefits such as environmental impact reduction.

Plant oil production is energy-intensive and is often located in rural areas. In Japan, where woody biomass from forestry is abundantly available, coupling plant oil production with woody biomass CHP represents a feasible pathway for regional symbiosis and decarbonization.

The recycle technology using Kolbe electrolysis process can effectively convert neutralization residues into valuable products; however, the total amount of such residues is relatively small compared to the quantity of refined oil, as indicated by the process model. Consequently, its overall contribution to emission reduction is limited, and its feasibility is expected to depend on the scale and conditions of individual production sites.

ACKNOWLEDGEMENTS

This study was conducted with support from the JST COI-NEXT program (JPMJPF2003) and the Environment Research and Technology Development Fund (Grant Number JPMEERF20243M01) of the Environmental Restoration and Conservation Agency of Japan. The University of Tokyo's "Platinum Society" Endowed Chair operates with contributions from Mitsui Fudosan Co., Ltd., Sekisui House, Ltd., East Japan Railway Company, and Toyota Tsusho Corporation.

REFERENCES

1. Ministry of Environment, Annual Report on the Environment in Japan 2025, (2025)
2. Conteratto C, Artuzo FD, Benedetti Santos OI, Talamini E. Biorefinery: a comprehensive concept for the sociotechnical transition toward bioeconomy. Renewable and Sustainable Energy Reviews 151:111527 (2021). https://doi.org/10.1016/j.rser.2021.111527
3. Huang B, Zhuge Y, Rameezdeen R, Xing K, Liu Y. Optimizing resource efficiency in regional industrial symbiosis: integrated modelling and scenario analysis. Journal of Cleaner Production 489:144744 (2025). https://doi.org/10.1016/j.jclepro.2025.144744
4. Aissani L, Lacassagne A, Bahers J, Féon S. Life

cycle assessment of industrial symbiosis: a critical review of relevant reference scenarios. J of Industrial Ecology 23:972-985 (2019). https://doi.org/10.1111/jiec.12842
5. Neves A, Godina R, Azevedo SG, Matias JCO. A comprehensive review of industrial symbiosis. Journal of Cleaner Production 247:119113 (2020). https://doi.org/10.1016/j.jclepro.2019.119113
6. Kikuchi Y, Nakai M, Kanematsu Y, Oosawa K, Okubo T, Oshita Y, Fukushima Y. Application of technology assessments to co-learning for regional transformation: a case study of biomass energy systems in tanegashima. Sustain Sci 15:1473-1494 (2020). https://doi.org/10.1007/s11625-020-00801-1
7. Yamaki A, Fujii S, Kanematsu Y, Kikuchi Y. Life cycle greenhouse gas emissions of cogeneration energy hubs at japanese paper mills with thermal energy storage. Energy 270:126886 (2023). https://doi.org/10.1016/j.energy.2023.126886
8. Kanematsu Y, Oosawa K, Okubo T, Kikuchi Y. Designing the scale of a woody biomass CHP considering local forestry reformation: a case study of tanegashima, japan. Applied Energy 198:160-172 (2017). https://doi.org/10.1016/j.apenergy.2017.04.021
9. Hiromori K, Konno Y, Katagami K, Takahashi A, Shibasaki-Kitakawa N. Kolbe electrolysis for recycling non-aqueous fatty acid salts. Journal of Chemical Engineering of Japan 57: (2024). https://doi.org/10.1080/00219592.2024.2332621
10. National Institute of Advanced Industrial Science and Technology (AIST), Research Institute of Science for Safety and Sustainability, IDEA Lab., AIST-IDEA ver. 3.5.1, (2025)
11. Goto N., Revised New Edition: Knowledge of Fats and Oils Products (in Japanese), Saiwaishobo, (2018)
12. Fujita J., Japan, Chemical Process Compendium (in Japanese), Tokyo Kagaku Dojin, 959–967, (1970)
13. Kikuchi Y, Kanematsu Y, Sato R, Nakagaki T. Distributed cogeneration of power and heat within an energy management strategy for mitigating fossil fuel consumption. J of Industrial Ecology 20:289-303 (2015). https://doi.org/10.1111/jiec.12374
14. Hiromori K., Investigating the applicability of Kolbe electrolysis for sustainable soapstock recycling, AOCS Annual Meeting & Expo, (2025)
15. FUCHINO T, MURAKI M. Constitution of a plant complex with heat recovery.. Sekiyu Gakkaishi 40:474-481 (1997). https://doi.org/10.1627/jpi1958.40.474
16. Ministry of Agriculture, Forestry and Fisheries, Oilseed Production Statistics (Final Report for FY2024) (in Japanese), (2025)
17. Han Z, Oosawa K, Kanematsu Y, Teah HY, Kikuchi Y. Simulating forest transition for sustainable forestry planning in japan: a backcasting model of roundwood supply and carbon stocks with operational and regional considerations. Resources, Conservation & Recycling Advances 29:200314 (2026). https://doi.org/10.1016/j.rcradv.2026.200314
18. Takahashi K, Shino J. Greenhouse gas emissions and loan markets in japan. Environmental Technology Innovation and ESG Investment :71-90 (2024). https://doi.org/10.1007/978-981-99-9768-8_4

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

OpenAD-lib: Open-Source Framework for Uncertainty-Aware Anaerobic Digestion Digital Twins

Benaissa Dekhici[ab], Rohit Murali[a], and Michael Short[ab]*

[a] School of Chemistry and Chemical Engineering, University of Surrey, Guildford, GU2 7XH, Surrey, United Kingdom
[b] Supergen Bioenergy Impact Hub, Energy and Bioproducts Research Institute, UK
* Corresponding Author: m.short@surrey.ac.uk.

ABSTRACT

This paper presents OpenAD-lib, an open-source Python framework for anaerobic digestion (AD) digital twins, unifying mechanistic models, machine learning (ML) surrogates, and model predictive control (MPC) within a modular ecosystem. OpenAD-lib addresses the critical fragmentation in AD digitalisation by bridging mechanistic and data-driven paradigms under explicit uncertainty. By integrating uncertainty-aware feedstock characterisation with robust process control, the platform enables the transition from isolated research tools to fully integrated digital twins, delivering economic and environmental value in AD systems.

Keywords: Digitalisation, Bioenegry, Anaerobic digestion, Digital twins, Uncertainty quantification, Model Predictive Control, Machine learning, Open-source framework

INTRODUCTION

Anaerobic digestion (AD) is critical to the UK's energy transition and circular economy, producing renewable biogas and nutrient-rich digestates that enhance soil health and carbon sequestration [5]. However, most AD facilities lack production-ready computational tools integrating mechanistic modelling with data-driven methods, resulting in passive monitoring, sub-optimal biogas recovery, and underutilised operational potential.

Since the introduction of Anaerobic Digestion Model No. 1 (ADM1) in 2002 [2], mechanistic AD modelling has advanced substantially, with ADM1 establishing the international reference standard for capturing complex substrate heterogeneity, microbial dynamics, and inhibitory mechanisms across 32 state variables [1]. However, four interconnected limitations impede sector digitalisation. First, the AD modelling landscape remains fragmented [5]. No unified platform integrates mechanistic models, ML surrogates, and control strategies. Practitioners assemble tools ad-hoc with divergent input/output conventions, and the absence of standardised feedstock descriptors across six documented characterisation methodologies for the ADM1 inputs [13] necessitates resource-intensive site-specific calibration, preventing model transfer between facilities. Second, existing tools deliver deterministic predictions without quantifying uncertainty or confidence bounds. Real-time operational decisions including biogas commitments and control setpoints require probabilistic guarantees rather than nominal estimates. Third, mechanistic and data-driven components operate in isolation. Recent advances including Long Short-Term Memory networks (LSTM) surrogates [9], Physics-Informed Machine Learning (PIML)[12], and multi-task Gaussian processes [4] (MTGP) demonstrate that hybrid approaches achieve both computational efficiency and mechanistic fidelity, yet no integrated commercial or open-source platform exists, forcing practitioners to choose between high-fidelity simulation or fast data-driven methods [5]. Finally, scattered AD models [4, 7, 9, 12] across disconnected repositories lack unified access, standardised comparison frameworks, and industrial-grade implementation practices required for operational deployment.

OpenAD-lib, developed within the Supergen Bioenergy Impact Hub (SBIH) [11] is an open-source Python framework that bridges mechanistic understanding with data-driven approaches under explicit uncertainty. The framework integrates feedstock characterisation, mechanistic models, ML surrogates, and control strategies in seamless workflows (unified platform), well-documented, comprehensively tested, and designed for industrial extensibility (production-ready), propagates uncertainty explicitly from feedstock characterisation

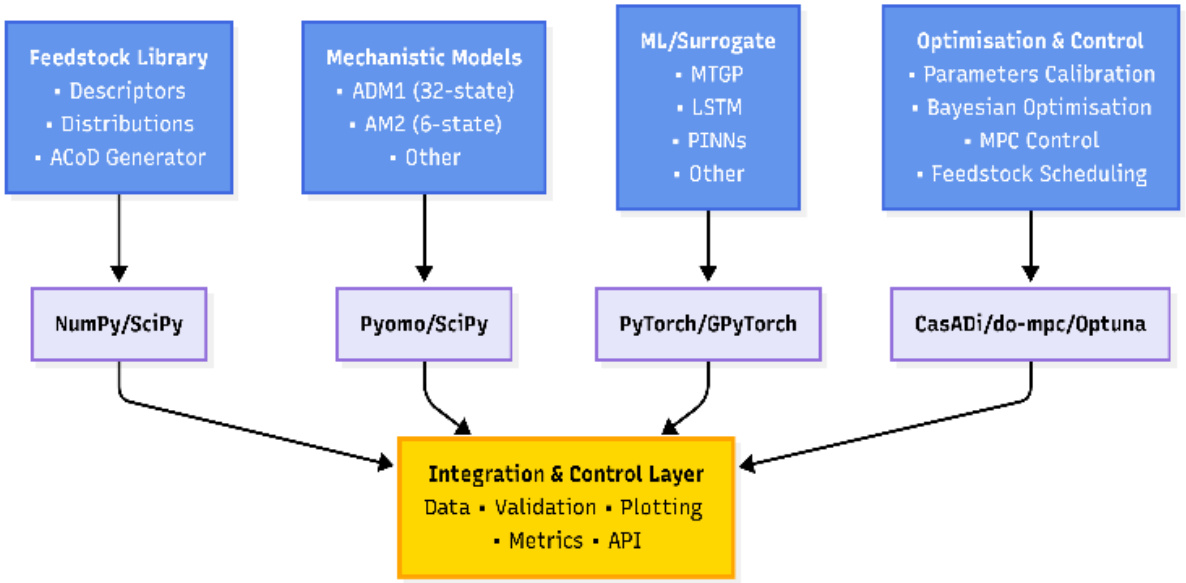

Figure 1. OpenAD-lib four-layer modular architecture.

through mechanistic simulation, surrogate approximation, and control optimisation to enable quantified confidence in operational decisions (uncertainty-aware by design), and enables users to select components (ADM1-only, surrogate-only, or hybrid approaches) matching their computational and accuracy requirements, avoiding unnecessary complexity while enabling gradual adoption (modular composability). Beyond AD, OpenAD-lib is designed for extensibility to complementary bioenergy pathways (gasification, pyrolysis) aligned with SBIH's partnership ecosystem.

The paper is organised as follows: Section 2 presents the software architecture, Section 3 describes key features and implementation, Section 4 demonstrates a case study, and Section 5 provides conclusions and future work.

SOFTWARE ARCHITECTURE

OpenAD-lib implements a four-layer modular architecture that enables uncertainty propagation from feedstock characterisation through process prediction to optimisation and control. Figure 1 presents the layered technology stack of OpenAD-lib, demonstrating how the framework leverages Python's scientific computing ecosystem to integrate each layer within a cohesive software architecture. This layered design philosophy reflects a systems level approach to digitalisation of AD processes, enabling users to employ components independently or in combination to accommodate diverse use cases ranging from offline feasibility analysis to real-time adaptive control:

- **Layer 1 -Feedstock Characterisation Library:** Encodes statistical representations of feedstock properties.
- **Layer 2 -Mechanistic Modelling:** Wraps standard solvers for ADM1 (32-states) and AM2 (reduced-order only 6-states) simulations.
- **Layer 3 -ML/Surrogate Models:** Provides pre-configured PyTorch modules for fast inference.

Layer 4- Optimisation and Control: Integrates optimization solvers for calibration and MPC.

The modular interconnection of these layers is orchestrated through a unified *Integration & Control Layer* that manages data flow and uncertainty propagation across all operational modes. The OpenAD-lib can be accecced in Github via the following link: https://github.com/benmola/OpenAD-lib

IMPLEMENTATION AND KEY FEATURES

Feedstock Characteristics Library

Reliable characterisation of feedstock properties is fundamental to predictive accuracy in AD modelling [6]. OpenAD-lib addresses this through a structured feedstock library that captures the variability inherent in real-world substrates and translates readily available measurements into the required inputs for the mechanistic models. The library integrates comprehensive representations of diverse feedstock types commonly encountered in typical UK AD sites.

Table 1: OpenAD-lib Feedstock Library: Data Coverage Summary

Feedstock Category	Included Feedstocks
Agricultural/ Energy Crops	Maize Silage, Wholecrop Cereals, Grass, Fodder Beet
Agro-industrial	Rice Bran, Corn Gluten/Rapemeal, Crimped Wheat
Food Processing	Lactose (Whey), Apple Pomace, Potatoes
Manure & Residues	Chicken Litter, Animal Manure

Table 2: Standardised descriptor set for uncertainty-aware feedstock characterisation.

Symbol	Descriptor	Unit
TS	Total Solids	Mass % or kg m^{-3}
VS	Volatile Solids	g kg^{-1} TS or kg m^{-3}
COD_t	Total COD	kg COD m^{-3}
$COD_{s,p}$	Soluble / Particulate COD	kg COD m^{-3}
BMP	Biochemical Methane Potential	NmL CH_4 kg^{-1} VS
f_{bio}	Biochemical Composition	g kg^{-1} or kg m^{-3}
$C:N$	Carbon-to-Nitrogen Ratio	Molar ratio
TAN	Total Ammonia Nitrogen	g N m^{-3}
V_{vfa}	Volatile Fatty Acids	kg m^{-3}

Feedstock Characterisation and Uncertainty Representation

The OpenAD-lib feedstock library consists of 12+ substrates consitsting of agricultural residues, energy crops, agro-industrial by-products and food waste streams as shown in Table 1. To capture the inherent variability of biological substrates, OpenAD-lib encodes feedstock properties using the standardised descriptor set detailed in Table 2. Unlike deterministic approaches, the framework models these inputs as probabilistic distributions ($x_{feedstock} \sim \mathcal{P}(\theta)$), parametrised via Maximum Likelihood Estimation (MLE) using data from multi-facility laboratory measurements, literature meta-analysis, and expert elicitation. To ensure physical consistency during ensemble generation, the library automatically assigns Beta distributions to bounded fractions (e.g., biodegradability coefficients) and Log-normal or Gamma distributions to strictly positive quantities (e.g., TS, VS, COD), thereby preventing non-physical sampling while facilitating the propagation of measurement uncertainty into downstream process control.

Mechanistic Modelling

The mechanistic modelling layer simulates AD dynamics using coupled ordinary differential equation (ODE) systems [1]. OpenAD-lib adopts a multi-model approach, offering both comprehensive and reduced-order formulations to balance process fidelity with computational efficiency for varying application requirements. The complex nonlinear dynamics are represented by a general state-space mass balance equation [1]:

$$\dot{x} = D(x_{in} - x) + Kr(x) \quad (2)$$

where $x \in R^n$ is the state vector of concentrations were is n is the number of the model state variables. D is the dilution rate (scalar or diagonal matrix, $\frac{Q_{in}}{V}$ where Q_{in} is the flow rate and V is the reactor volume). x_{in} is the vector of influent concentrations. K is the stoichiometric matrix containing yield coefficients. $r(x)$ is the vector of kinetic reaction rates (Monod/inhibition functions). The mechanistic layer provides Python wrappers for the ADM1 and the reduced-order AM2 model. The engine uses SciPy/Numpy for standard simulation and Pyomo/CasADi for symbolic gradient-based optimization tasks. OpenAD-lib implements the ADM1 as the primary high-fidelity representation. The implementation follows [2, 7] comprising a system of 32 ($x \in R^{32}$) coupled ODEs that describe seven biochemical pathways. ADM1 offers exceptional detail for hypothesis testing and steady-state design [2]. However, the requirement to estimate ~100 parameters and the computational stiffness of the 32-state system make it challenging for real-time control applications or scenarios requiring thousands of iterative simulations [1, 3]. To address the control-oriented limitations of ADM1, OpenAD-lib incorporates the Anaerobic Model for Control (AMOCO), hereafter referred to as the AM2 model [3]. AM2 aggregates the complex dynamics of ADM1 into a two-step biological process (acidogenesis and methanogenesis) governed by two lumped biomass populations. The AM2 dynamics follow the same general mass balance structure as in Eq (2), but with only 6 state variables ($x \in R^6$). A major barrier in applying mechanistic models is the complex mapping required to convert routine physicochemical measurements (e.g., TS, VS, COD) into the state variables required by ADM1/AM2. To resolve this, OpenAD-lib provides a standardised framework that mathematically transforms available feedstock data into stoichiometrically consistent ADM1 input fractions, ensuring mass conservation (COD, Nitrogen) across the model interface [6]. The framework also supports other widely-adopted characterisation protocols depending on available data fidelity [13]. For facilities operating with substrate mixtures, the library implements a mixing module that aggregates individual feedstock descriptors. Once

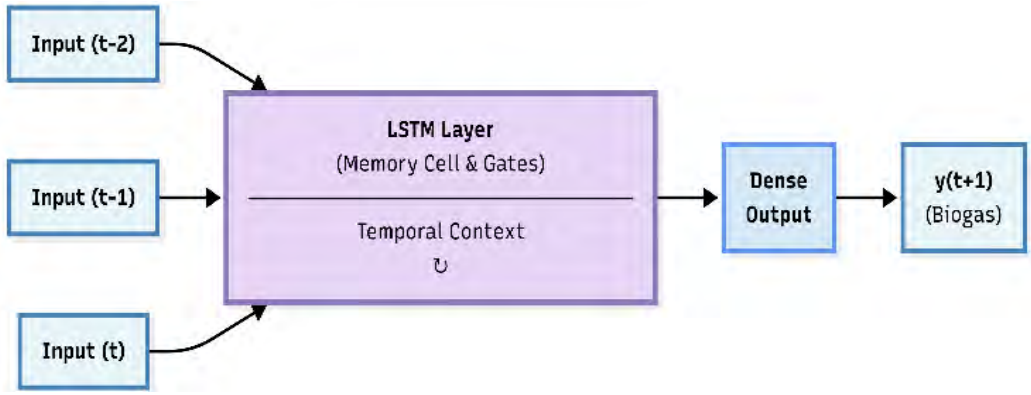

Figure 2. LSTM Architecture (Temporal Forecasting) [9].

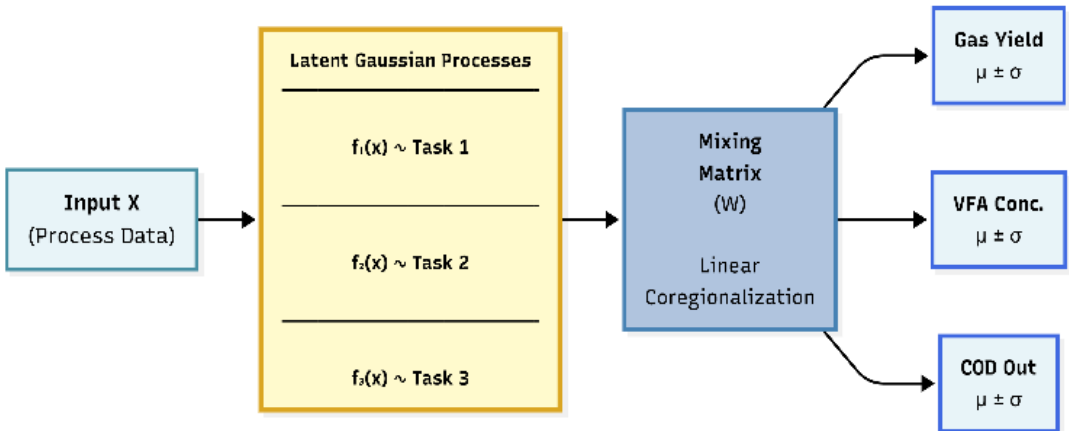

Figure 3. MTGP Architecture (Multi-OUTPUT & Uncertainty) [4].

characterised, the composite influent state vector $x_{mixture}$ is computed via weighted averaging based on volumetric or mass flow ratios (w_i) :

$$x_{mixture} = \sum_{i=1}^{N} w_i x_i \quad \text{with} \quad \sum_i w_i = 1 \tag{3}$$

This formulation preserves the mass balance of all state variables while naturally propagating the uncertainty from individual substrate characterisations to the final blended influent. By providing a single, standardised interface to multiple established characterisation protocols, OpenAD-lib removes the historical barrier whereby practitioners implemented ad-hoc, undocumented conversions from measurements to ADM1 inputs, often introducing unquantified error and reducing model transparency [6, 13].

ML/Surrogate Models

The ML surrogate approximation problem is formulated as learning an efficient mapping:

$$\hat{y} = \mathcal{S}(x; \phi) \tag{4}$$

where $\mathcal{S}$ denotes the learned surrogate function parameterised by ϕ (e.g. neural network weights, Gaussian process hyperparameters), and the ML surrogate is trained to approximate outputs $\hat{y}$ (e.g. cumulative methane/biogas production, VFA concentration trajectory, digestate composition) given inputs x (real system data such as feedstock composition and feeding, operating conditions). However, ML surrogates require hyperparameter tuning of the network and it remains a "black box" approach. OpenAD-lib serves as a repository for ML surrogates tailored to AD and biogas production forecasting. Currently, the framework implements two validated models: LSTM [9] and MTGP [4], both developed by the authors and natively integrated into the library. Future versions will progressively incorporate additional surrogate architectures from peer-reviewed studies in the AD and bioenergy domain. The LSTM (Figure 2) utilises a many-to-one architecture, ingesting a sliding window of historical operational data (e.g., 30 days) to explicitly model the retention time lag, predicting future yield based on past feeding trends [9]. In contrast, the MTGP (Figure 3) utilises a non-recurrent structure designed for robust state estimation. It maps current

process inputs to multiple correlated outputs (Gas, Stability, Effluent) through a set of shared latent Gaussian functions. Crucially, the MTGP output layer includes a covariance term, generating confidence intervals (shaded regions) that quantify the uncertainty of the prediction [4]. To address the computational stiffness of ODE solvers, OpenAD-lib provides native PyTorch modules that share the same API as the mechanistic models. To synergise the interpretability of mechanistic descriptions with the fitting power of deep learning, the library implements a Physics-Informed Machine Learning (PIML) framework. As illustrated in Figure 4, this hybrid approach integrates the "Mechanistic Model" (ADM1 or AM2) directly into the learning loop of the "ML Model." By embedding governing ODEs as physical regularisation terms or using mechanistic simulations to generate synthetic pre-training data, the framework constrains the neural network to biologically plausible solutions [12]. This interaction creates an "Improved Model" that balances the high fidelity of mass-balance equations with the computational speed of data- driven inference. The implementation supports "grey-box" configurations where the ML component compensates for unmodelled phenomena (e.g., temperature shocks) while the physical component guides extrapolation in data-sparse regimes, preventing the prediction of chemically impossible states (e.g., negative concentrations) outside the training domain.

Optimisation and Control

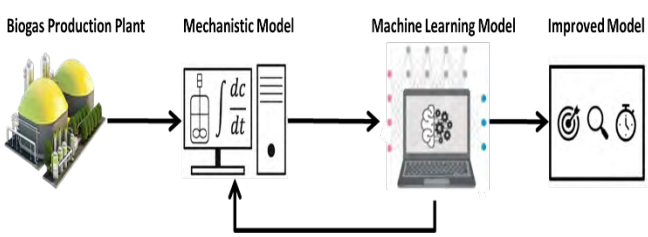

Figure 4. Physics-Informed Machine Learning workflow.

For real-time process optimisation and control, OpenAD-lib integrates a modular control layer designed to be model-agnostic. This architecture allows users to deploy any calibrated model, whether a mechanistic system (AM2, ADM1) or a data-driven surrogate (LSTM, MTGP) as the internal prediction engine (f_{model}) without altering the control logic.

Automated Calibration and Tuning

Prior to deployment, the library leverages the *Optuna* framework [8], a state-of-the-art Bayesian optimisation library to automate parameter estimation. For mechanistic models, this involves identifying kinetic constants (μ_{max}, K_s), while for ML surrogates, it tunes architectural hyperparameters (e.g., learning rate, hidden layers). The library utilises the Tree-structured Parzen Estimator (TPE), a Bayesian optimisation algorithm that constructs probabilistic density estimates of promising versus unpromising parameter regions [8]. TPE is chosen for its ability to converge in significantly fewer iterations than random search [8]. The objective is to minimise the deviation between simulated outputs and experimental time-series data:

$$J_{calib}(\Theta) = \sum_{t=1}^{T} ||y_{meas}(t) - y_{sim}(t, \Theta)||^2 \quad (4)$$

where θ is the vector of unknown kinetic parameters. This automated routine ensures that the internal model accurately reflects the biological reality of the specific digester before control authority is granted.

Model Predictive Control (MPC) Formulation

The control layer wraps the *do-mpc* library [10] to solve nonlinear Optimal Control Problems (OCP). The controller discretises the process dynamics using orthogonal collocation on finite elements, ensuring numerical stability even for stiff biological systems (e.g., rapid pH changes). At each sampling instant k, the solver uses the calibrated prediction engine (f_{model}) to compute a sequence of optimal control inputs u (e.g., dilution rate, feedstock blend) over a finite horizon N_p:

$$\underset{u}{\text{MIN}}\, J_{MPC} = \sum_{k=0}^{N_p} \left(l(x_k, u_k)\right) + V\left(x_{N_p}\right) \quad (5)$$

Subject to:

- **Model Dynamics:** $x_{k+1} = f_{model}(x_k, u_k)$ (Mechanistic ODEs or ML Inference)
- **Path Constraints**: $g(x_k) \leq 0$ (e.g., $pH > 6, VFA < 2000\ mg/L$)
- **Input Bounds:** $u_{min} \leq u_k \leq u_{max}$

where $l(x_k, u_k)$ is the running cost (penalising deviation from setpoints or maximising yield) and $V\left(x_{N_p}\right)$ represents the terminal cost, which penalises the final state deviation to ensure infinite-horizon stability. The optimisation problem is transcribed into a Non-Linear Programming (NLP) problem via CasADi and solved using IPOPT implemented within the do-mpc library [10]. The implementation natively supports two configurable control strategies: i) Yield Maximisation: the stage cost is defined to maximise energy recovery ($J = -Q_{\text{gas}}$), driving the system to the highest loading rate possible within stability constraints. ii) Setpoint Tracking: the controller minimises deviations from a reference trajectory ($J = ||y - y_{ref}||^2$), useful for maintaining effluent quality standards (e.g., COD limits).

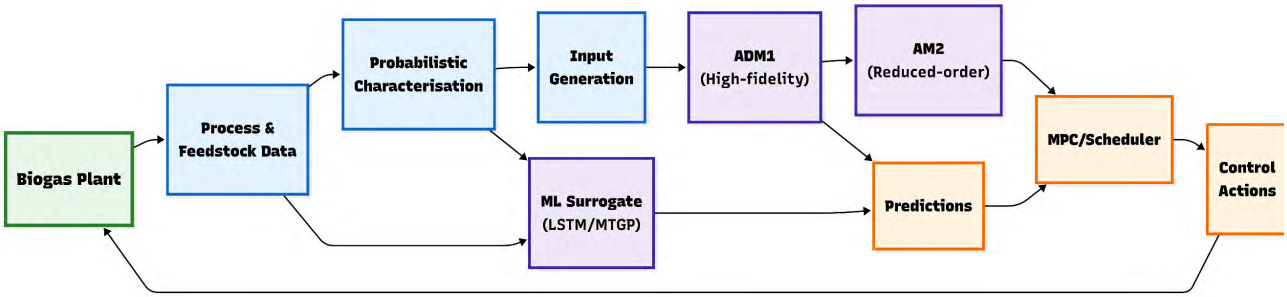

Figure 5: OpenAD-lib framework architecture showing the integrated workflow from feedstock characterization through mechanistic models, ML surrogates to control with closed-loop feedback.

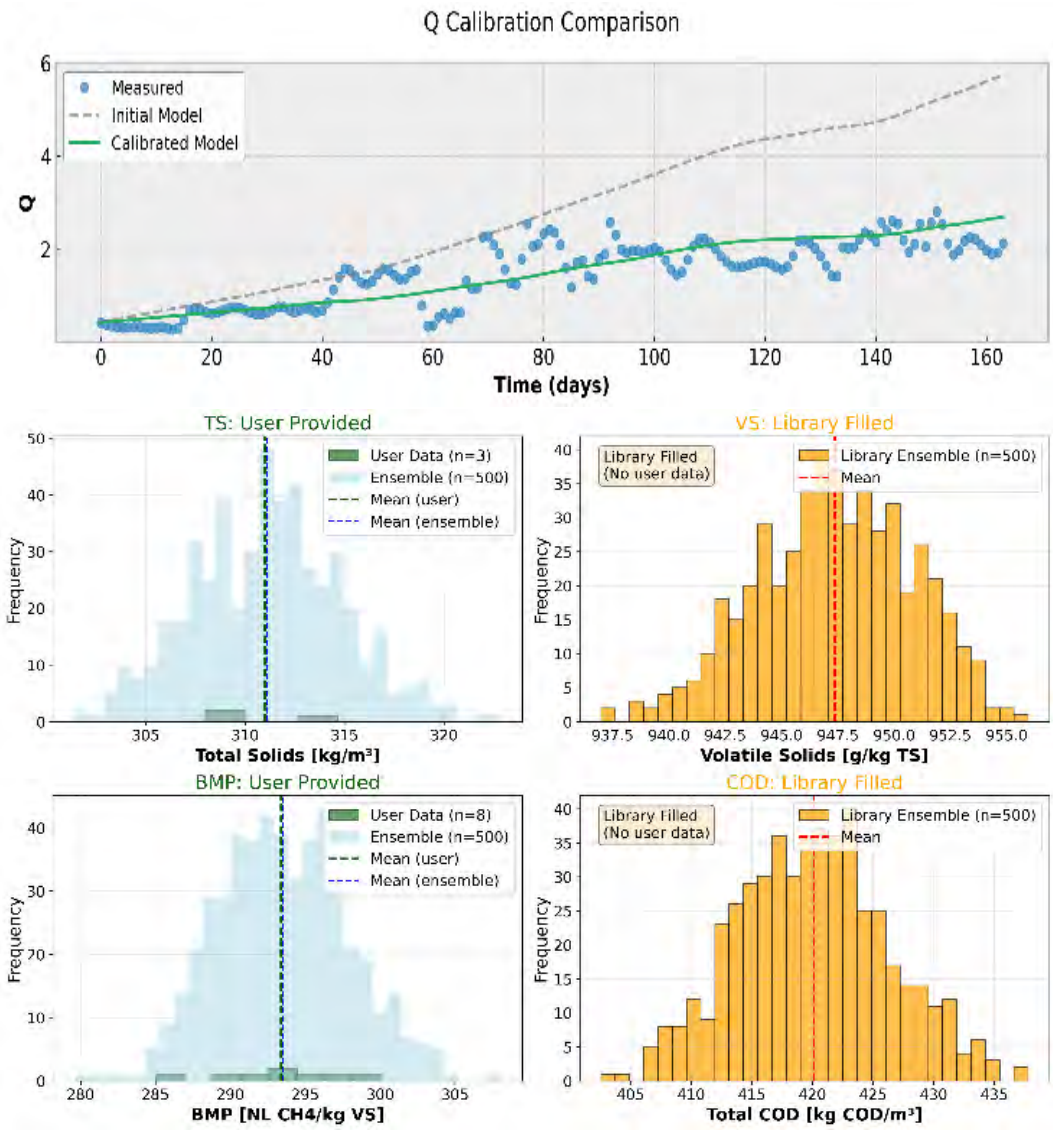

Figure 6. Probabilistic Feedstock Characterisation Sampling.

CASE STUDY

We present a comprehensive case study of a co-digestion biogas plant to showcase the 'plug-and-play' nature of the OpenAD-lib ecosystem. This digital twin integrates the entire process pipeline, transforming sparse feedstock inputs into optimised control trajectories. To

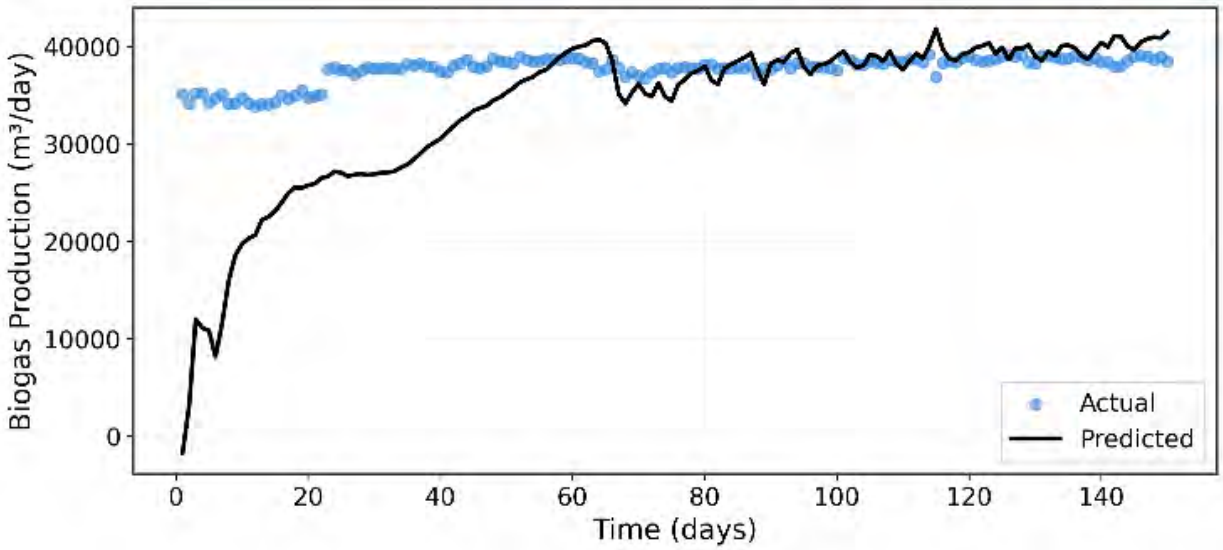

Figure 7. ADM1 mechanistic simulation validation.

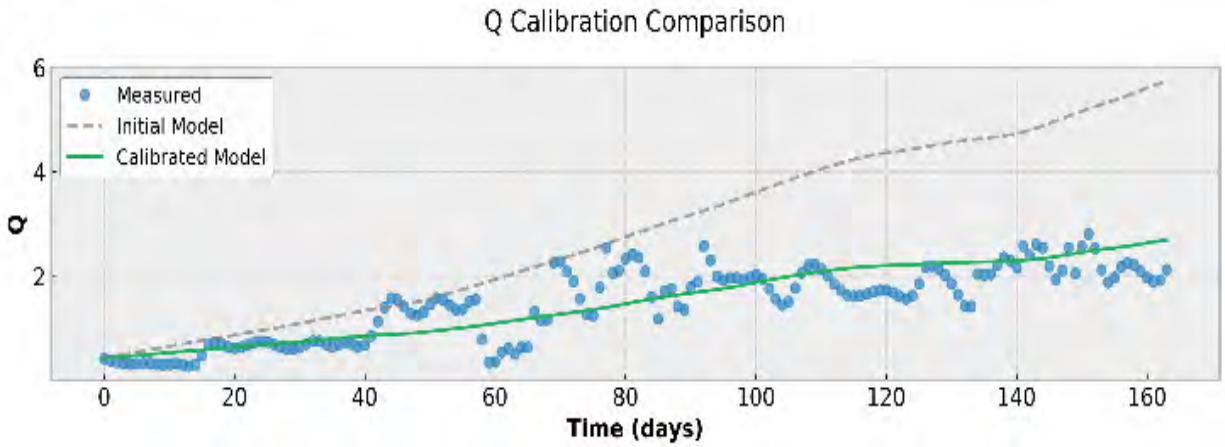

Figure 8. Tune simplified model to plant data with optimised parameters, before/after comparison.

validate the framework across varying degrees of information availability, the case study leverages a multi-source data strategy: it combines real-world sensor readings with synthetic data generated to simulate rare operational events, alongside standard datasets adopted from the original manuscripts of the ADM1/AM2 and MTGP /LSTM models [2, 3, 4, 9]. This hybrid approach ensures the digital twin is tested against both theoretical baselines and realistic operational noise. Figure 5 outlines the complete architecture and data flow of this integrated workflow.

Phase 0: Probabilistic Feedstock Characterisation

This phase Handles sparse user measurements with library uncertainty filling. Figure 6 shows the probabilistic feedstock characterisation with sparse user measurements (3-8 samples) combined with library-filled distributions for TS, VS, BMP, and COD parameters.

```
import openad_lib as openad
# Initialize feedstock library
lib = openad.FeedstockLibrary()
# User provides LIMITED measurements (real-
istic scenario)
user_maize_data = {
 'ts': [310, 315, 308], # Only 3 TS measure-
ments from lab
 'bmp': [285, 300, 293, 290, 295, 292, 298,
294] # 8 BMP measurements
# Missing: vs, cod_total, proteins, lipids
}
# Library fills missing parameters automati-
cally
maize_prob = lib.get_probabilistic("Maize",
user_data=user_maize_data)
```

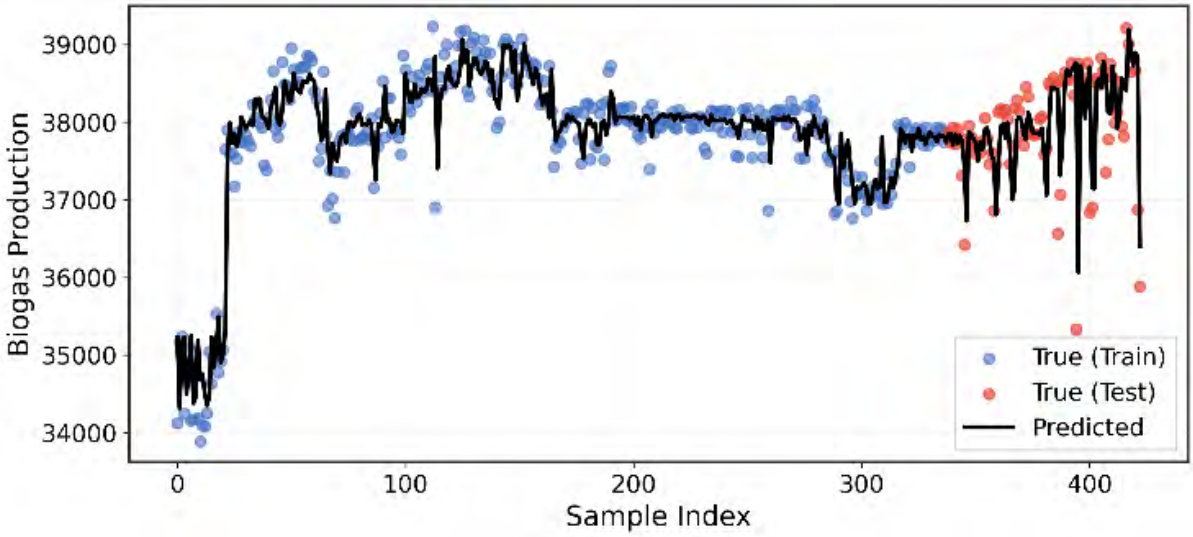

Figure 9. LSTM biogas production predictions.

```
# Generate ensemble for uncertainty propaga-
tion
maize_ensemble = maize_prob.sample(n=500,
random_state=42)
```

Phase 1: ADM1 Mechanistic Baseline

After generating the ADM1 model inputs from the probabilistic characterisation we establish the physics-based understanding by running ADM1 simulations and compare the results of the measured and simulated outputs as shown in Figure 7.

```
# Generate influent from feedstock ratios
influent_df = openad.acod.generate_influ-
ent_data(feedstock_data)
# Run ADM1 simulation (32 states)
adm1_model = openad.ADM1Model()
results = adm1_model.simulate(influent_df)
# Evaluate performance
metrics = openad.utils.metrics.compute_met-
rics(y_true, y_pred)
```

Phase 2: AM2 Parameters Calibration

In this step the AM2 model calibration using Optuna is achieved. Figure 8 show a comparison of pre-calibration predictions of biogas production showing poor fit versus post-calibration performance after 30 optimisation trials.

```
# Calibrate with Optuna (30 trials)
calibrator = openad.AM2Calibrator(am2_model)
best_params = calibrator.calibrate(
 params_to_tune=['m1', 'K1', 'm2', 'Ki',
'K2'], n_trials=30)
```

Figure 8. Tune simplified model to plant data with optimised parameters, before/after comparison.

Phase 3A: LSTM Time Series Prediction

LSTM time-series prediction for biogas production with 30-day lookback window. Figure 9 show the results with training phase (blue, days 0-60) demonstrates model learning on historical feeding patterns. Test phase (orange, days 60-90) validates generalisation capability. LSTM architecture: input dimension = 6 (operational + feedstock features), hidden dimension = 24, training epochs = 30 as shown in the code below.

```
# Create LSTM model
lstm = openad.LSTMModel(input_dim=6, hid-
den_dim=24)
# Prepare time series with lags
X, y, dataset = lstm.prepare_time_se-
ries_data(data, features, target, n_in=1)
# Train and evaluate
lstm.train(X_train, y_train, epochs=30)
metrics = lstm.evaluate(X_test, y_test)
```

Phase 3B: Multi-Task GP with Uncertainty

Next, multi-output prediction with confidence intervals using the MTGP of three correlated outputs with uncertainty quantification is shown. Figure 10 presentes the results of (a) soluble COD (SCOD_out) predictions with 95% confidence intervals (shaded region) achieving 94% empirical coverage; (b) VFA concentration predictions capturing process dynamics; (c) biogas production rate with mean ± 2σ bounds.

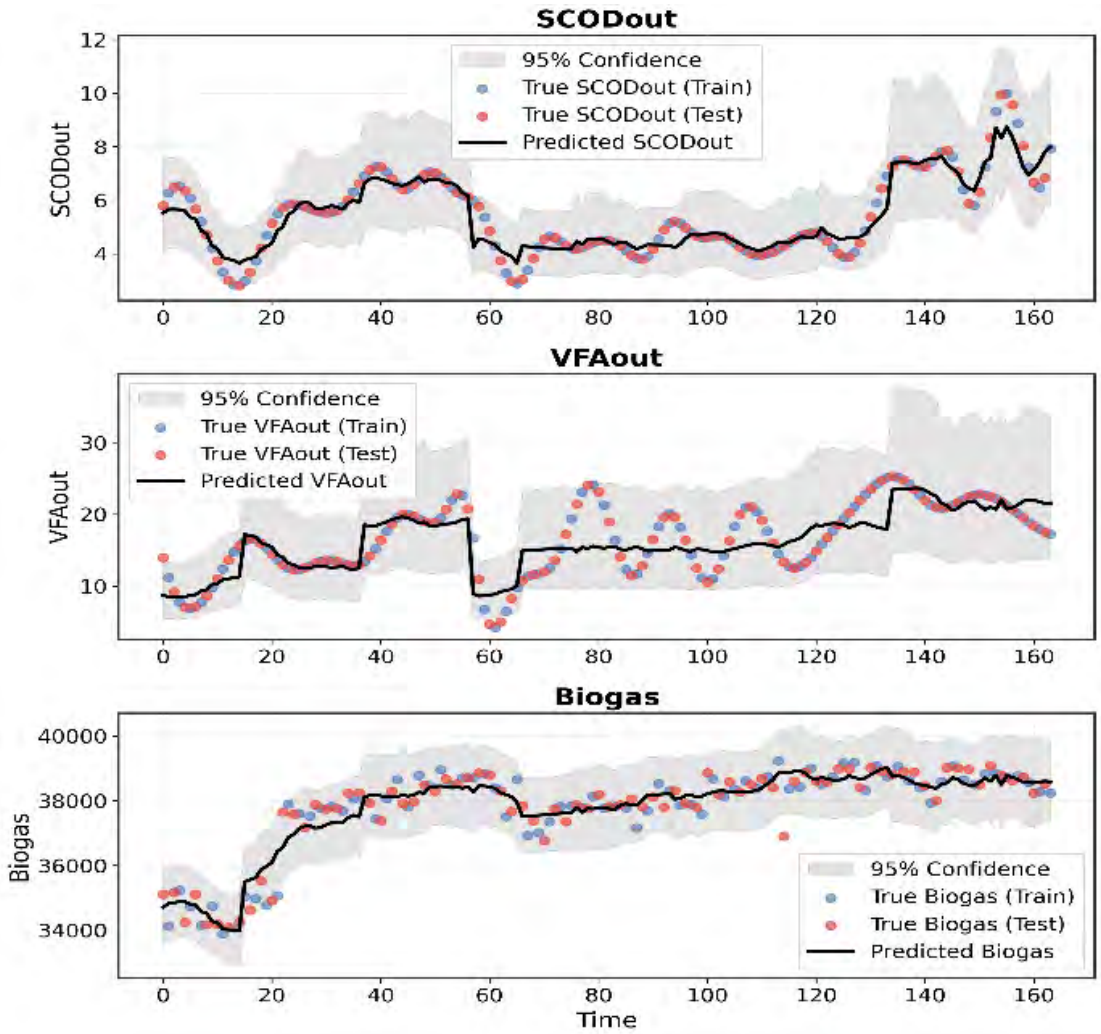

Figure 10. MTGP OF SCOD, VFA, and biogas with 95% confidence intervals.

```
# Initialize MTGP (3 outputs: SCOD, VFA, Bi-
ogas)
mtgp = openad.MultitaskGP(num_tasks=3,
num_latents=3)
# Train and predict with uncertainty
mtgp.train(X_train, Y_train, epochs=300)
mean, lower, upper = mtgp.predict(X_test,
return_std=True)
```

Phase 4A: MPC Biogas Maximisation

For control, MPC for maximum biogas production. Tow-panel trajectory showing (Figure 11): (a) optimal dilution rate (control input) computed at each MPC step subject to process constraints; (b) biogas production increase from 0.001 L/d (baseline) to 0.03 L/d under optimal feeding strategy. MPC configuration: prediction horizon = 10 steps, control update frequency = 1 day, with AM2 internal model.

```
# Setup MPC controller
mpc = openad.AM2MPC(params)
mpc.setup_controller(
 horizon=10,
 objective_type='maximize_biogas',
 D_max=0.5
)
# Run closed-loop simulation
for k in range(30):
 u_opt, y_next =
mpc.run_step(S1in_val=S1in_k, pH_val=pH_k)
```

Phase 4B: MPC VFA Tracking

MPC for VFA concentration setpoint tracking to maintain process stability. Figure 12 demonstrates: (a) VFA concentration regulated at 2.0 mmol/L target setpoint (dashed line) with tight tracking despite

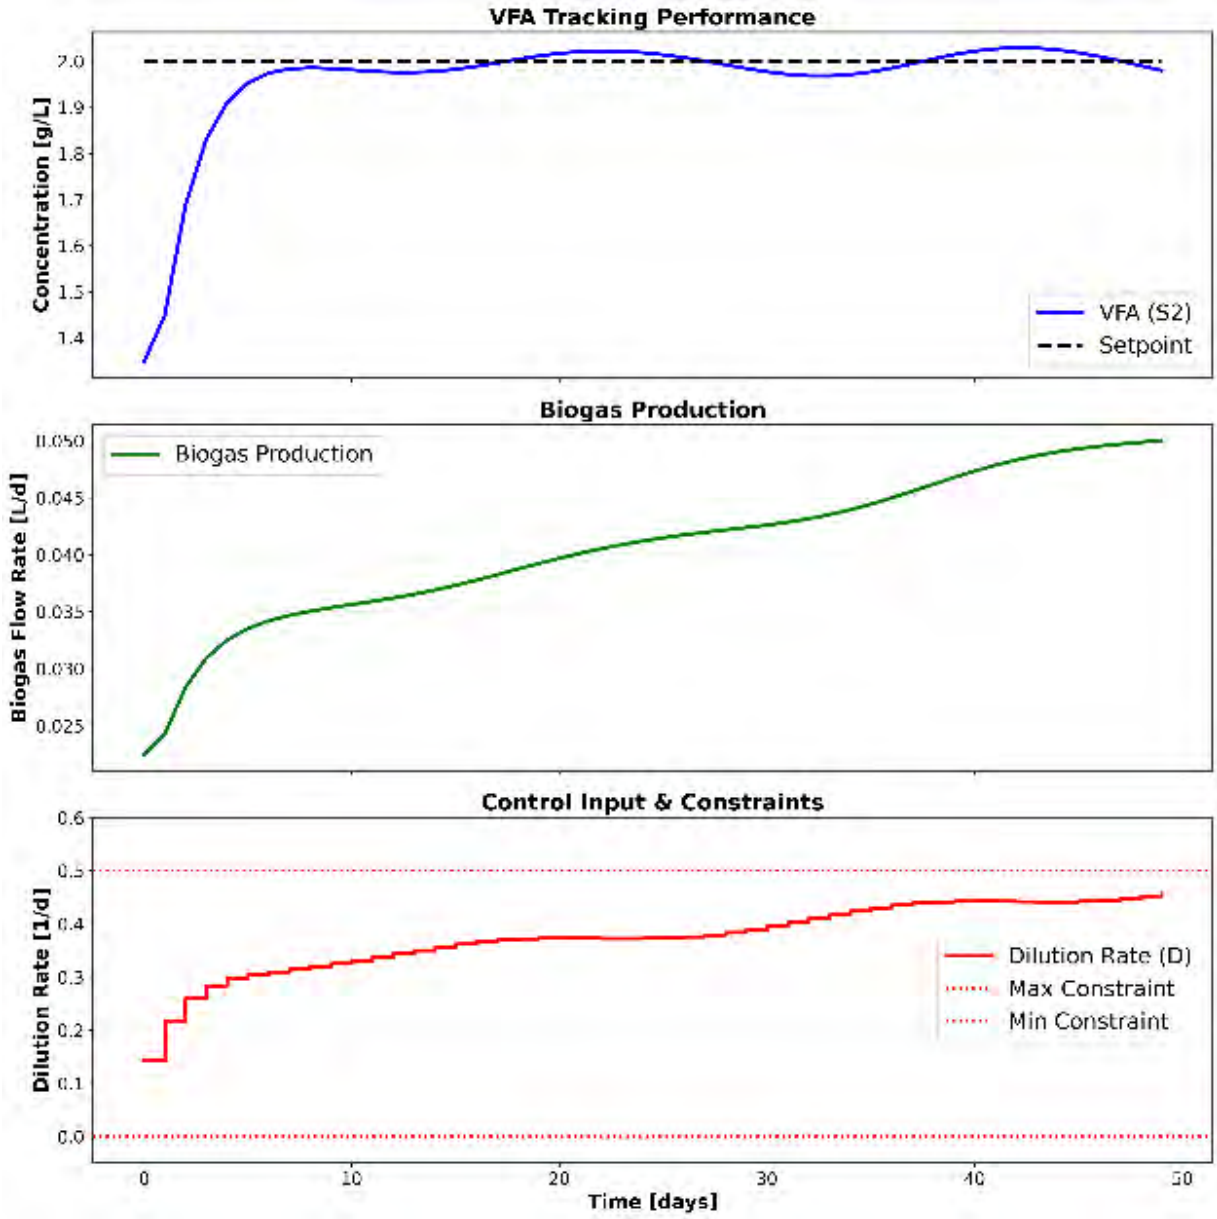

Figure 11. MPC for maximum biogas production.

feedstock composition disturbances; (b) control effort (dilution rate adjustments) showing smooth manipulated variable trajectory without excessive actuator action.

```
# Configure VFA tracking
mpc.setup_controller(
objective_type='track_vfa',
```

Figure 12. MPC for VFA concentration setpoint tracking to maintain process stability.

```
vfa_setpoint=2.0
)
# Run tracking simulation
for k in range(30):
  u_opt, y_next = mpc.run_step(...)
```

CONCLUSION

OpenAD-lib addresses the digital gap in AD industry by unifying mechanistic understanding, data-driven methods, and uncertainty quantification. By providing a production-ready, open-source ecosystem, it enables the development of robust strategies for real-world AD application. Future iterations will extend to bioenergy

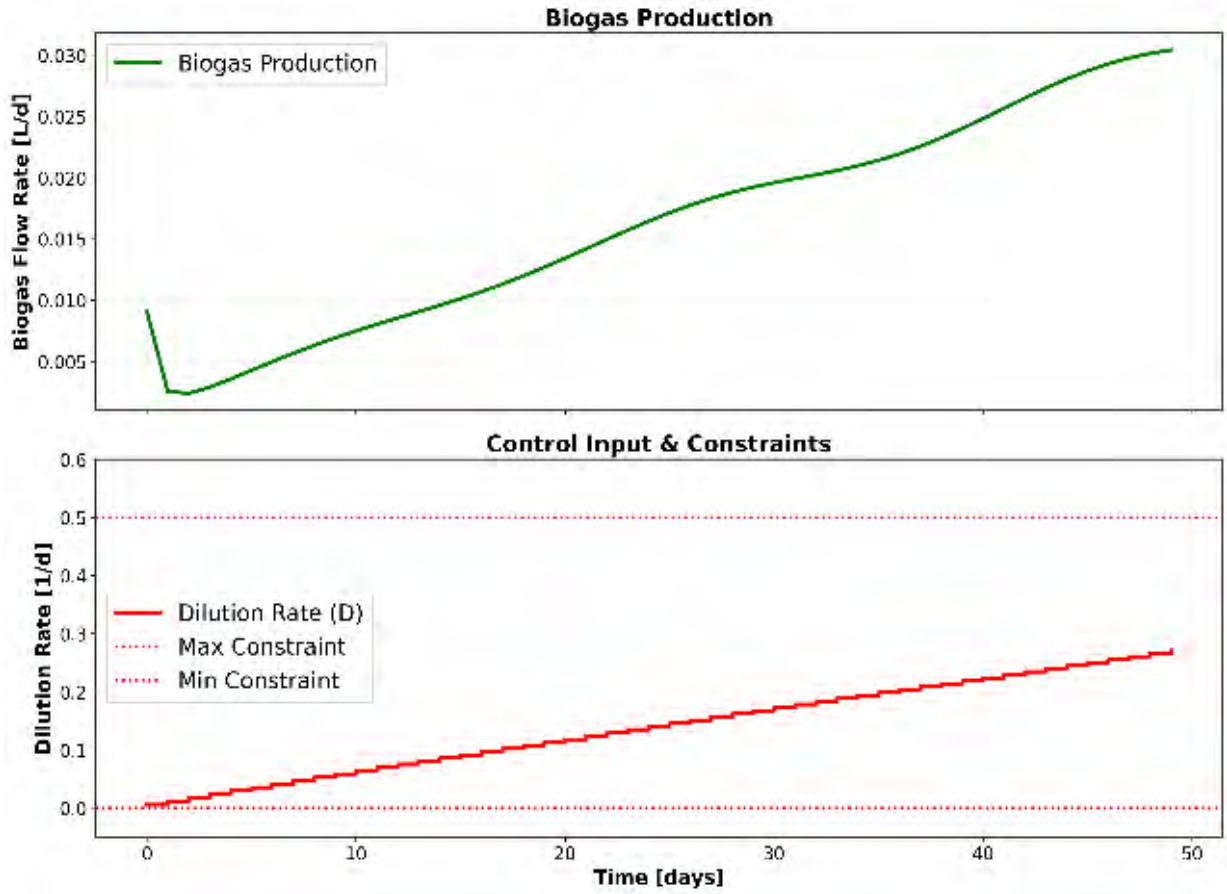

Figure 12. MPC for VFA concentration setpoint tracking to maintain process stability.

technologies (such as gasification and pyrolysis) demonstrating the platform's role in the wider bioenergy transition.

ACKNOWLEDGEMENTS

The authors thank both UKRI grants: the Supergen Bioenergy Impact Hub, EP/Y016300/1 and AI for Net Zero, EP/Y005600/1, for their support.

AUTHOR IDENTIFIERS

Author ORCIDs:
Benaissa Dekhici: 0000-0002-2660-1029
Rohit Murali: 0009-0001-5515-7721
Michael Short: 0000-0003-3227-1844

REFERENCES

1. Dekhici B. Data-driven modeling, order reduction and control of anaerobic digestion processes. Ph.D. thesis, University of Tlemcen (2024)
2. Batstone DJ, Keller J, Angelidaki I, Kalyuzhnyi SV, Pavlostathis SG, Rozzi A, Sanders WTM, Siegrist H, Vavilin VA. The IWA anaerobic digestion model no 1 (ADM1). Water Science and Technology 45:65-73 (2002). https://doi.org/10.2166/wst.2002.0292
3. Bernard O, Hadj?Sadok Z, Dochain D, Genovesi A, Steyer J. Dynamical model development and parameter identification for an anaerobic wastewater treatment process. Biotech & Bioengineering 75:424-438 (2001). https://doi.org/10.1002/bit.10036
4. Dekhici B, Short M. Data-driven modelling of biogas production using multi-task gaussian processes. Systems and Control Transactions 4:26-32 (2025). https://doi.org/10.69997/sct.121877
5. Murali R, Bywater A, Dolat M, Dekhici B, Zarei M, Hilton L, Sadhukhan J, Zhang D, Short M. Anaerobic digestion site-wide optimisation and decision-making: an industrial perspective and review. Renewable and Sustainable Energy Reviews 226:116402 (2026). https://doi.org/10.1016/j.rser.2025.116402

6. Arnell M, Astals S, Åmand L, Batstone DJ, Jensen PD, Jeppsson U. Modelling anaerobic co-digestion in benchmark simulation model no. 2: parameter estimation, substrate characterisation and plant-wide integration. Water Research 98:138-146 (2016). https://doi.org/10.1016/j.watres.2016.03.070
7. Sadrimajd P, Mannion P, Howley E, Lens PNL. PyADM1: a Python implementation of Anaerobic Digestion Model No. 1. bioRxiv (2021) https://doi.org/10.1101/2021.03.03.433746
8. Akiba T, Sano S, Yanase T, Ohta T, Koyama M. Optuna: A Next-generation Hyperparameter Optimization Framework. arXiv:1907.10902 (2019) https://arxiv.org/abs/1907.10902
9. Murali R, Dekhici B, Chen T, Zhang D, Short M. Mechanistic and data-driven models for predicting biogas production in anaerobic digestion processes. Systems and Control Transactions 4:388-393 (2025). https://doi.org/10.69997/sct.176459
10. Fiedler F, Karg B, Lüken L, Brandner D, Heinlein M, Brabender F, Lucia S. Do-mpc: towards FAIR nonlinear and robust model predictive control. Control Engineering Practice 140:105676 (2023). https://doi.org/10.1016/j.conengprac.2023.105676
11. Supergen Bioenergy Impact Hub.
12. https://www.supergen-bioenergy.net/research/digitalisation-and-ai-for-ad/
13. Shaw KM, Poh PE, Ho YK, Chen ZY, Chew IML. Modeling the anaerobic digestion of palm oil mill effluent via physics-informed deep learning. Chemical Engineering Journal 485:149826 (2024). https://doi.org/10.1016/j.cej.2024.149826
14. Mayowa F. Oladele, George M. Bollas, Leveraging Digital Twin Modeling for Anaerobic Digesters using Anaerobic Digestion Model No. 1 (ADM1) and Neural Network within the Pyomo Framework, Computer Aided Chemical Engineering, Elsevier, Volume 53, 2024. https://doi.org/10.1016/B978-0-443-28824-1.50191-5
15. Ficara E, Hassam S, Allegrini A, Leva A, Malpei F, Ferretti G. Anaerobic digestion models: a comparative study. IFAC Proceedings Volumes 45:1052-1057 (2012). https://doi.org/10.3182/20120215-3-at-3016.00186

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Achieving Net-Zero Emissions in Industrial Parks Through Optimized Symbiotic Exchanges and Carbon Capture Utilization

Ricardo N. Dias[ab*], Fátima N. Serralha[b], Carla I.C. Pinheiro[a]

[a] Centro de Química Estrutural, IMS, Departamento de Engenharia Química, Instituto Superior Técnico/Universidade de Lisboa, Av. Rovisco Pais 1, 1049-001 Lisboa, Portugal
[b] Escola Superior de Tecnologia do Barreiro, Instituto Politécnico de Setúbal, 2839-001 Lavradio, Portugal
* Corresponding Author: ricardo.n.dias@tecnico.ulisboa.pt.

ABSTRACT

The integration of industrial symbiosis (IS) and carbon capture utilization (CCU) is recognized as a key strategy for achieving net-zero emissions in industrial parks (IPs). However, the optimization of these processes in combination remains an undeveloped research area. This work presents a multi-objective optimization framework implemented in Pyomo and linked with OpenLCA to simultaneously maximize material exchanges and minimize operational costs while evaluating CCU deployment. Applied to Portugal's largest industrial park, the model identified 26 feasible symbiotic exchange routes involving 14 enterprises and 7 potential CCU technologies. Maximum material exchange optimization yields 3,042,107 ton/year across 26 symbiotic routes and 7 CCU units, achieving 89.8 % reduction in climate change impact (from 13.5 to 1.76 million CO_2eq/year); Cost minimization achieves 2,223,298 ton/year with 27 % fewer exchanges, delivering 87.0 % environmental reduction and net revenue of €65.5M/year via carbon tax avoidance. The proposed trade-off configuration deploys 2 CCU units (both for methanol production), establishes 12 symbiotic routes, exchanges 2,925,556 ton/year, and generates €65.5M/year net benefit. Decision variables encompassed inter-enterprise flows, binary decisions to select exchanges, and CCU unit selection informed by technology readiness level (TRL). The epsilon-constraint method generated a Pareto front revealing the cost-benefit trade-offs of the symbiotic network. IS with CCU integration can reduce IP climate impact by 87-90 % while generating substantial economic returns. Methanol emerges as the critical CCU pathway for replacing fossil-imported feedstocks. This framework provides decision-makers with actionable optimization strategies for achieving near net-zero industrial parks by 2050, not only through CCU but also through IS.

Keywords: carbon capture, Pyomo, life cycle assessment, optimization

INTRODUCTION

Industrial Parks (IPs) face a fundamental sustainability challenge, optimizing economic benefits often contradicts minimizing environmental impacts. While Industrial Symbiosis (IS) aims to perform well on those two dimensions through resource sharing, achiev–ing genuine economic and environmental impact, optimization remains an unsolved challenge. To complicate things, there is always a social dimension to be considered. Through IS, it is possible to reduce in–dustrial waste, promoting the Circular Economy (CE) and extending the life cycle of several materials [1,2].

The evidence is compelling: industrial parks generate waste volumes nearly 18 times those of municipal waste, creating pollution hotspots and ecosystem degradation [3]. Yet established symbiosis networks demonstrate that material, energy, and water exchanges can simultaneously reduce environmental footprints, cut operating costs, and generate economic growth [4,5]. The question is not whether industrial symbiosis can deliver sustainability benefits—empirical evidence from Kalundborg confirms it does—but rather how to systematically design industrial park networks that reconcile

economic performance, environmental, and social well-being without subordinating local firm viability to global system efficiency or compromising one sustainability dimension to advance another [5].

The optimization of IP has been the focus of numerous studies, encompassing the identification of potential symbiotic pathways [6] and the identification of redundancy and resilience in these potential exchanges [7]. The integration and optimization of energy and water resources for IPs have also been a subject of extensive research. However, optimizing symbiotic pathways while accounting for the introduction of new facilities and their dynamics remains an active research field. A solid decision-making process is essential for proposing the introduction of process-specific alternations [8] and even more so for sounder still modifications at the IP level.

In this work, we propose a novel approach to IS in IPs. This approach is based on a model that aims to maximize material exchanges while minimizing the costs of implementation of such strategies. Furthermore, the model evaluates the implementation of a carbon capture and utilization (CCU) unit into the IP, based on some proven CCU technologies.

METHODOLOGY

The methodology presented in this section focuses on an optimization model that maximizes material exchanges between enterprises while minimizing the costs for the IP. A sustainability assessment is also carried to quantify the environmental gains of the symbiotic routes (OpenLCA is used to this end), enabling a tri-way assessment of the proposed IP configuration, total mass exchanges ($Total_{EM}$), total costs ($Total_{Cost}$), and environmental accountability [9].

$$maximize\ Z = Total_{EM} \tag{1}$$

$$minimize\ Z = Total_{Cost} = Operational_{cost} + CAPEX \tag{2}$$

Subject to:

The presented objective functions are subject to a set of constraints here presented:

The $Total_{EM}$ is determined by summing the flow of all material exchanges between enterprises (EM), between enterprises and proposed CCU units (EM_{CCU}), and the flow from CCU units to enterprises (EM_{CCUR}). Each exchanged material variable is subject to specific sets, namely, the donor enterprise (n) or CCU unit (n''), the material category (j) whether it is a residue, raw material, coproduct, or product. Furthermore, the material code (i), and the receiving enterprise or CCU unit. For clarity, all EM variables are indexed according to these sets.

$$Total_{EM} = \sum_n \sum_j \sum_i \sum_{n'} EM_{n,j,i,n'} + \sum_n \sum_j \sum_i \sum_{n''} EM_{CCU_{n,j,i,n''}} + \sum_{n''} \sum_j \sum_i \sum_n EM_{CCUR_{n'',j,i,n}} \tag{3}$$

Capital Expenditure

The CAPEX is estimated by a combination of the costs of introducing a new CCU unit into the IP C_{CCU} and the costs of establishing a new exchange route, C_{newEM}.

$$CAPEX = C_{CCU} \cdot Y_{CCU} + C_{newEM} \tag{4}$$

Gross operational costs from the exchanges

The operation costs, $Operational_{cost}$, are obtained as a sum of the current market cost of each material, $C_{i,j}$, multiplied by the total respective material exchanged. This cost varies by material category; for example, a material that is residue will have a different cost compared to its product value.

$$Operational_{cost} = \sum_n \sum_j \sum_i \sum_{n'} EM_{n,j,i,n'} \cdot C_{i,j} + \sum_n \sum_j \sum_i \sum_{n''} EM_{CCU_{n,j,i,n''}} \cdot C_{i,j} + \sum_{n''} \sum_j \sum_i \sum_n EM_{CCUR_{n'',j,i,n}} \cdot C_{i,j} \tag{5}$$

ε - Constraint

This constraint ensures that multiple points between the maximum and minimum $Total_{EM}$ obtained are covered in the solving of the multiobjective problem.

$$Total_{EM} \geq \varepsilon, \varepsilon \in \{Total_{EM}^{min}, Total_{EM}^{max}\} \tag{6}$$

Upper bound limited by availability

The following set of constraints ensures that the decision variables are within the available resources in each of the considered enterprises or CCU units. This data comes from parameters, real data, $RD_{n,j,i}$, and CCU data $RD_{CCU_{n'',j,i}}$.

$$\sum_n EM_{n,j,i,n'} \leq RD_{n,j,i} \cdot Y_{symb_{n,j,i,n'}}, \forall n \in \{1, \dots, N\}, j \in \{0, \dots, J\}, i \in \{1, \dots, I\} \tag{7}$$

Upper bound limited by reception capacity

$$\sum_{n'} EM_{n,j,i,n'} \leq RD_{n,0,i} \cdot Y_{symb_{n,j,i,n'}}, \forall n \in \{1, \dots, N\}, j \in \{0, \dots, J\}, i \in \{1, \dots, I\} \tag{8}$$

Lower bound of the flow

$$EM_{n,j,i,n'} \geq MinEM_{n,j,i} \cdot Ysymb_{n,j,i,n}, \forall n \in \{1, \dots, N\}, \quad i \in \{1, \dots, I\} \tag{9}$$

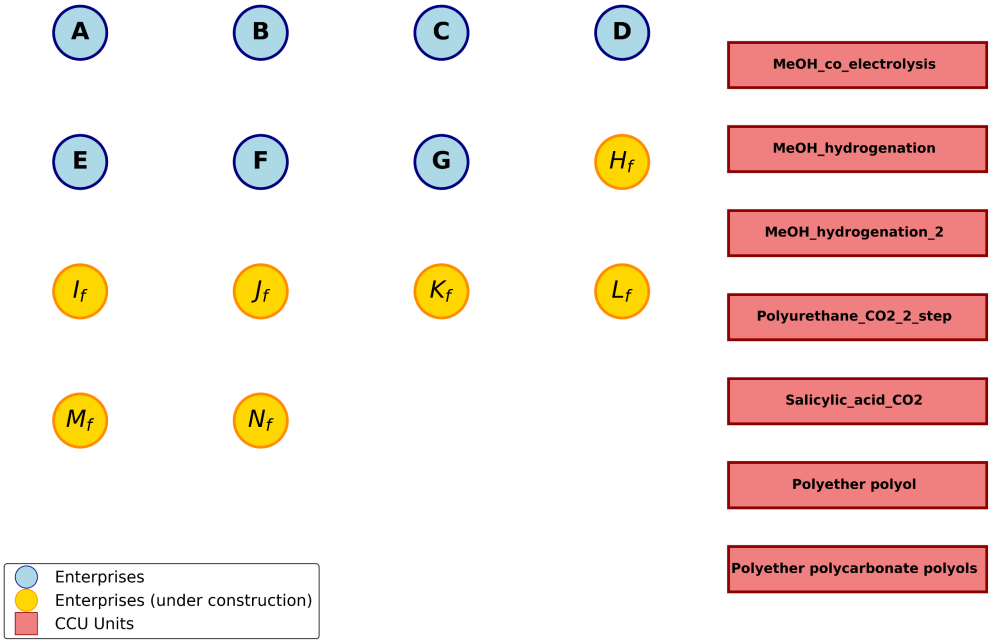

Figure 1. Enterprises considered in the IP and the possible CCU units to be activated.

$$MinEM_{n,j,i} = min\{RD_{n,j,i}\sum Y_{symb_{n,j,i,n'}}, RD_{n,0,i}\sum Y_{symb_{n,j,i,n'}}\}, \quad \forall n \in \{1,\dots,N\}, j \in \{0,\dots,J\}, i \in \{1,\dots,I\}, \wedge n \neq n' \tag{10}$$

Constraints for integrating CCU technologies into the IP

$$\sum_{n''} EM_{CCU_{n,j,i,n''}} \leq RD_{n,j,i} \cdot Y_{CCU_{n''}}, \forall n \in \{1,\dots,N\}, j \in \{0,\dots,J\}, i \in \{1,\dots,I\} \tag{11}$$

$$\sum_{n} EM_{CCU_{n,0,i,n''}} \leq RD_{CCU_{n'',0,i}} \cdot Y_{CCU_{n''}}, \forall n'' \in \{1,\dots,N\}, i \in \{1,\dots,I\} \tag{12}$$

$$\sum_{n''} Y_{CCU_{n''}} \leq 7 \tag{13}$$

$$\sum_{n} EM_{CCUR_{n'',j,i,n}} \leq RD_{CCU_{n'',j,i}} \cdot Y_{CCU_{n''}}, \forall n'' \in \{1,\dots,N\}, j \in \{0,\dots,J\}, i \in \{1,\dots,I\} \tag{14}$$

$$\sum_{n''} EM_{CCUR_{n'',0,i,n}} \leq RD_{n,0,i} \cdot Y_{CCU_{n''}}, \forall n \in \{1,\dots,N\}, \quad i \in \{1,\dots,I\} \tag{15}$$

CCU mass balance

$$\sum_{n}\sum_{j} EM_{n,j,i,n'} + \sum_{c}\sum_{j} EM_{CCUR_{n,j,i,c}} \leq RD_{n',0,i}, \quad \forall\{n \neq n', j \neq 0\} \tag{16}$$

$$Y_{n,j,i,n'}\begin{cases} 1 \vee 0 \; if \; RD_{n',0,i} \wedge RD_{n',0,i} > 0 \\ 0 \; otherwise \end{cases} \tag{17}$$

Sustainability assessment

OpenLCA is used to obtain climate change characterization factors [10]. They are then used to quantify the environmental gains that can be obtained with the symbiotic exchanges calculated. The characterization factors are considered in eq. 18 in parameter CC_i, in CO_{2eq}.

$$\sum_{n}\sum_{j}\sum_{i}\sum_{n'} EM_{n,j,i,n'} \bullet CC_i + \sum_{n}\sum_{j}\sum_{i}\sum_{n''} EM_{CCU_{n,j,i,n''}} \bullet CC_i + \sum_{n''}\sum_{j}\sum_{i}\sum_{n} EM_{CCUR_{n'',j,i,n}} \bullet CC_i \leq 0 \tag{18}$$

Key assumptions

To implement the presented model, some key assumptions were made and are here listed:

- The costs of each material were determined form current market prices for Europe.
- The current EU ETS (European Union European Trading System) price for CO_2 was used to monetize the CO_2 emissions. As according to EU

ETS the free allowances are intended to be phased out until 2034. Therefore, a worst-case scenario, in which enterprises must pay 100 % of its emissions was considered.

- A cooperative approach to symbiotic routes was considered, as if economically viable the partners are expected to establish such exchanges.
- The presented results in this work are case-study specific, even though the Framework is generic and can be applied to any IP.

CASE STUDY

The methodology previously presented has been implemented in a case study, the largest industrial park in Portugal. Only publicly available data was assessed by the model. The IP features a large oil refinery, a petrochemical plant, and several other smaller enterprises. The list of the considered enterprises is presented in **Table 1**. It must be noted that enterprises from H_f to N_f are not yet built or operating in the IP, these have already received the necessary licensing approvals required for commissioning. Therefore, these might start operations in the near future, J_f and K_f, for example, are already under construction and are L_f is still to begin construction.

As stated in the Introduction, the incorporation of CCU technologies into the IP was considered. The data on CCU technologies were derived from existing literature and are presented in **Table 2**. A range of products is under consideration, and all technologies under review have currently a TRL of 9. This indicates that these technologies have already been proven and deployed at an industrial scale. Methanol, a chemical compound classified as a building block, facilitates the direct substitution for its fossil-based counterpart in various chemical industries [11]. In addition, methanol has been studied as a hydrogen carrier due to its superior transport properties compared with pure hydrogen [12]. In this case study, methanol assumes particular significance to the IP. Currently, there are currently imports of methanol to the IP, with local production of e-methanol some or indeed all of the imported methanol could be displaced. While salicylic acid is predominantly recognized for its pharmaceutical applications, it is noteworthy that this chemical compound utilizes CO_2 in its production process, thereby contributing to potential carbon conversion. The polyurethane and polyether polyols under consideration are exemplars of products of added value that can enable the utilization of CO_2 in their synthesis [13]. The primary applications of these materials are in construction, where they are predominantly left untouched for long periods, thereby promoting carbon storage in the products themselves.

Table 1: Industries considered in the IP and the respective code considered.

Enterprise	Working field
A	Oil refining
B	Petrochemicals
C	Fibrous resins
D	Light gases production
E	Biodiesel
F	Diesel from slops
G	Terephthalic acid production
H_f	Batteries
I_f	Hydrogen
J_f	Biodiesel and SAF
K_f	Polypropylene and Polyethylene
L_f	Methanol
M_f	Hydrogen
N_f	Ammonia

Table 2: CCU technologies considered and their respective CAPEX and CO_2 capture capacity.

Product	CO_2 captured (ton/year)	CAPEX (M€)	Ref.
MeOH, co-electrol.	7056	100	[11]
MeOH, hydrogen	352000	72.92	[12]
MeOH, hydrogen	644000	200	[17]
Salicilic acid	468	100	[18]
Polyurethane	20000	100	[13]
Polyether polyol	23000	383	[19]
Polyether polycarbonate polyols	1000	15	[20]

The considered CCU technologies were gathered considering state of the art peer-reviewed papers which concluded that these are the current CO_2 utilization pathways with TRL 9 [14–16].

RESULTS

The results obtained of the optimization model, applied to the case study data, are presented here. First, the results for maximizing the material exchanged and minimizing the total cost are presented, second the multiobjective optimization results and a Pareto front are presented and discussed.

The model identified a total of 26 possible exchange routes to be established between existing IP enterprises. This derives from the $Y_{symb_{n,j,i,n'}}$ variable, which enables the establishment of a symbiotic route, which can then be optimized by the model. Furthermore, the introduction of CCU units into the IP will open several possible symbiotic routes, from residues to products, that might displace the current import of raw materials by locally produced ones. Figure 1 presents a graph of all possible

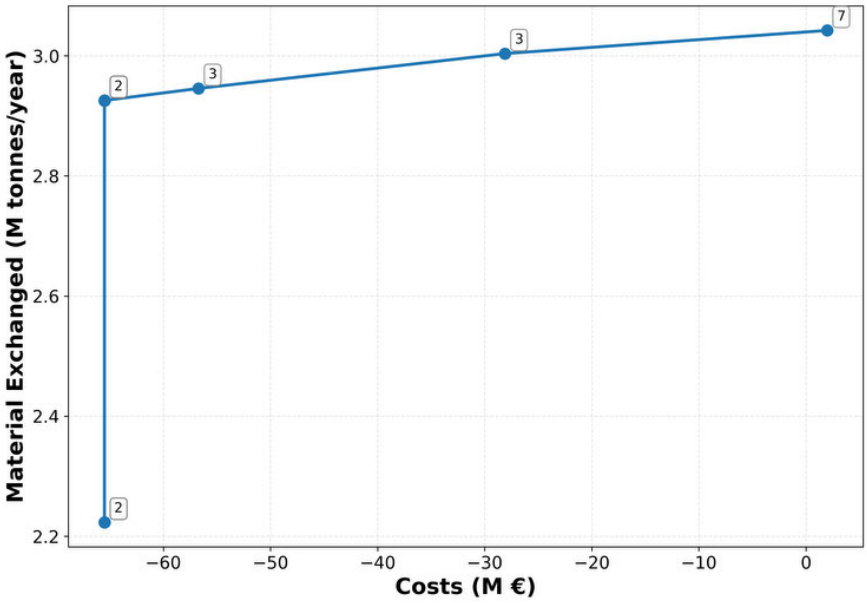

Figure 2: Pareto front between costs and material exchange. The number of CCU units is presented alongside each point.

symbiotic routes that can be established, with the introduction of a CCU unit highlighted to better illustrate the opportunities it can offer to the IP in this case study.

Material exchange maximization

First, the model was run to optimize the IP configuration and symbiotic routes to maximize material exchange between enterprises. This results in a total of 3,042,107 ton of material to be exchanged per year. This means that several residues, products, and co-products can be exchanged within the IP, leading to reduced environmental impacts. Mainly when residues can be directly or partially used in other processes, or when fossil-derived raw materials can be displaced by green products produced locally from captured CO_2. This is the case for methanol, which can be produced from CO_2 captured locally and is currently being imported to the IP. Seven CCU units were determined to be viable pathways for utilizing the CO_2 produced in the IP.

Total cost minimization

Secondly, the model was used to optimize total cost, resulting in a total material exchange of 2,223,298 ton/year. This value is 27 % lower than the obtained total material exchanged when optimized for material exchanges. The model identified the possibility of introducing only two CCU units, as these are expensive to build (significant CAPEX) and might not lead to further IS beyond CO_2, without proper incentive, might be hard to justify.

Multiobjective problem solving

As stated in the methodology section, to solve the multiobjective problem, an epsilon-constraint approach was used. The model was run to minimize total costs considering that the total material exchange should be greater than 2,223,298 plus an epsilon value, which lead to 20 values between the 2,223,298 and 3,042,107. The Pareto front of the non-dominated solutions can be found in **Figure 2**, where the obtained Pareto front is presented between the two objective functions. It is possible to identify that maximizing total material exchange also increases costs, requiring construction of 7 CCU units. However, this will also increase revenue as new products will be produced under the IP. It must be noted that the L_f unit, to which will be built in the upcoming years, be produce enough methanol to fulfil the current methanol requirements of the IP., This way, the introduction of a CCU unit to produce methanol from CO_2 will bring redundancy to the IP, in terms of methanol production. Yet, with the right incentives [4], and the increasing prices on CO_2 emissions [21], this might go from a redundant methanol production unit to a key player in the CO_2 valorization market. Furthermore, methanol is a building block chemical that can be used to produce several other products. It was also found that the costs show minimal variation among the first few points of the Pareto front. This is because only 2 CCU units are deployed in these first few points.

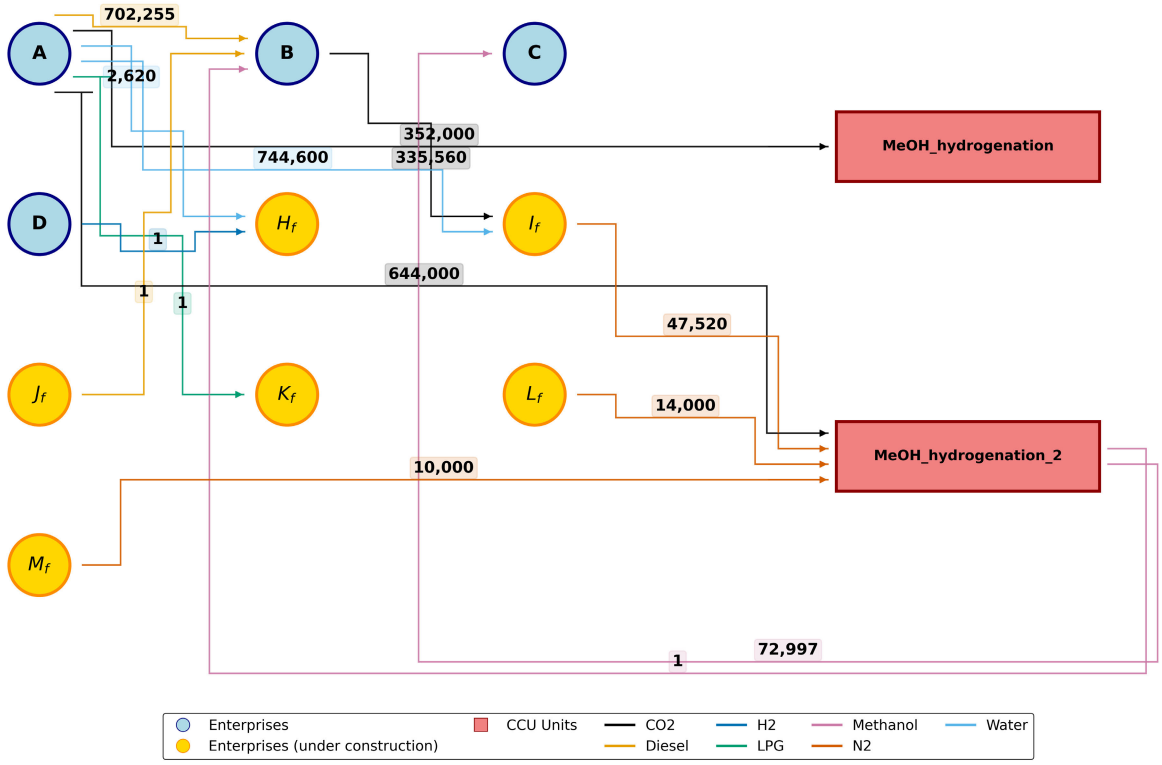

Figure 3. Proposed IP configuration, a trade-off between material exchanges and costs.

Table 3: Main results for the extreme pareto points, business as usual and proposed configuration of the IP with the introduction of IS and CCU units.

Iteration	T. EM (Mt/year)	T. Costs (M€/year)	Nr. connections	CCU units
BAU	-	-	-	-
Min Costs	2,223	-65,503	10	2
Max EM	3,042	1,974	26	7
Proposed config	2,925	-65,503	14	2

The IP configuration with the introduction of 2 CCU units is determined to be the most effective approach and is presented in **Figure 3**, as it enables the reduction of the IP's climate change impacts through carbon capture and valorization. Furthermore, the introduction of symbiotic routes, where methanol from fossil fuels is replaced with locally produced methanol from CO_2. Even when minimizing the implementation costs, by missing out on several exchanges, the environmental gains are significant, with 87.4 % reduction in climate change.

Table 4: Climate change variation with the introduction of IS.

Iteration	Climate change (t CO_{2eq})	Change (%)
BAU	13,505,894	-
Min Costs	1,759,455	-87.0
Max EM	1,383,357	-89.8
Proposed config	1,707,930	-87.4

Environmental gains

The total climate change impact was also calculated to quantify the environmental gains of establishing new symbiotic routes. These routes are particularly beneficial in that they redirect residues away from waste treatment plants (such as landfills or specialized treatment) toward symbiotic valorization. As presented in **Table 4**, the BAU (business-as-usual) configuration—defined as the conventional park setup without symbiotic exchanges yields 13,505,894 ton CO_{2eq}. Conversely, all symbiosis-enabled configurations resulting in a reduction of this impact to between 1.38 and 1.76 million ton CO_{2eq}. Specifically, the Minimum Costs and Proposed Configuration scenarios achieved comparable climate change values (1,759,455 and 1,707,930 t CO_{2eq}, respectively), corresponding to reductions of 87.0 % and 87.4 % relative to BAU. In

contrast, the Maximum EM scenario yielded the lowest impact (1,383,357 t CO_{2eq}), representing an 89.8 % reduction.

SENSITIVITY ANALYSIS

To assess the robustness of the multi-objective optimization, a two-part sensitivity analysis was carried out. First, a deterministic scenario analysis was applied to CCU processing efficiency and material market prices using factors (κ) of 70 %, 85 %, 115 %, and 130 % to the baseline values. These parameters act uniformly on all exchanges and are well suited to assess the effects of CCU efficiency and market price in the overall results. To assess the enterprise-level feedstock availability, a deterministic approach that would consider a constant facto applied to all flows would not have been realistic, so a Monte Carlo (MC) analysis was performed, with 100 samples. In this case, each donor enterprise received an independent random perturbation drawn from a uniform distribution U [1−δ, 1+δ], with δ ∈ {10 %, 20 %}, to reflect heterogeneous supply conditions across enterprises. For each deterministic or MC scenario, the full ε-constraint routine was run (as presented in the Methodology section), and the resulting Pareto front was evaluated through the calculation of a hypervolume indicator (HV), which provides a single scalar measure of front quality. Robustness of individual Pareto solutions was quantified as the fraction of MC runs in which each baseline Pareto point remained non-dominated.

The main results for the sensitivity analysis are presented in **Table 5**. The deterministic analysis was assessed using a one at a time approximated first-order sensitivity index [22], which shows that the Pareto front is considerably more sensitive to market price than to CCU efficiency (approximate first-order indices $S_1 \approx 1.64$ and $S_1 \approx 0.35$, respectively). In the MC analysis (with 100 samples), the ±10 % feedstock perturbation led to a mean HV reduction of 0.29 % relative to the baseline, concurrently a coefficient of variation of 0.85 %, indicating that supply fluctuations of up to 10 % have a negligible impact on the trade-off structure. The ± 20 % perturbations resulted in a mean HV loss of 1.61 % (and a coefficient of variation of 3.01 %), with larger spread (as it would be expected) but with still a modest degradation of the mean HV of around 1.3 % difference between the 10 and 20 % MC runs.

A Mann–Whitney U test (a non-parametric statistical test used to compare two independent distributions without assuming normality [23]) was applied to compare the HV distributions between the ±10 % and ±20 % bands. This test confirmed that the two HV distributions are statistically different ($p = 0.0069$), yet the magnitude of this difference remains small from an engineering perspective. Importantly, the minimum-cost exchange level (total exchanged material lower bound) was unchanged across all MC runs, meaning that every solution on the baseline Pareto front remained non-dominated in all simulated supply scenarios.

Table 5: Sensitivity analysis results for key model parameters using deterministic and Monte Carlo (MC) methods. The perturbation κ is applied as a factor in the deterministic method. For MC analysis, parameters were sampled uniformly within the bands.

Parameter	Method	κ Range	HV (%)	CV (%)	S_1
CCU efficiency	Deterministic	0.70–1.30	—	—	0.35
Market price	Deterministic	0.70–1.30	—	—	1.64
		MC bands			
Feedstock availability	MC	± 10 %	0.29	0.85	—
Feedstock availability	MC	± 20 %	1.61	3.01	—

MC – Monte Carlo, HV - hypervolume, CV – coefficient of variance, S_1 - first order index.

CONCLUSIONS

The concept of industrial symbiosis has the potential to simultaneously promote economic growth and environmental gain. In this study, a multi-objective optimization model is presented and solved to assess an industrial park. The analysis revealed that the optimization of IP for material exchanges resulted in enhanced environmental benefits. Conversely, cost minimization led to a decrease in the volume of exchanges.

The findings indicate that the maximum annual exchange capacity is 3,042,107 ton/year, thereby reducing overall IP climate change by 89.4 %. This objective is realized through the implementation of 26 symbiotic routes and the introduction of 7 CCU units into the IP. Optimizing for costs has been shown to result in 87.0 % reduction in the climate change impact of the overall IP, while concurrently achieving a net return (65.5 M €/year), primarily through the avoidance of carbon tax payments on avoided CO_2 emissions. The proposed IP configuration enables the introduction of 2 CCU units, a total of 12 symbiotic routes, 1,707,930 ton/year of total material exchanged, and a net gain of 65.5 M €/year. Additionally, it can be concluded that, while promoting IP integration, it is feasible to reduce both the economic costs of residues

and associated with environmental impacts. The findings of this study lay the foundation for further exploration into specific symbiotic pathways. These pathways have the potential to not only reduce current disposal costs but also to enhance circularity.

The sensitivity analysis demonstrated that the model is robust to parameter uncertainty: the Pareto front showed limited sensitivity to CCU efficiency (S_1 = 0.35) but higher sensitivity to market prices (S_1 = 1.64), while stochastic perturbations of up to ± 20 % in feedstock availability led to a modest mean HV reduction of only 1.61 %, with all baseline Pareto solutions remaining non-dominated across all MC scenarios.

Subsequent research endeavors should concentrate on time discretization and deployment assessment, to facilitate a deeper understanding of IP dynamics and how IP actors interact.

ACKNOWLEDGEMENTS

The authors of this work are grateful for the Fundação para a Ciência e a Tecnologia, I.P./MCTES through national funds PIDDAC - CQE UID/00100/2025 and UID/PRR/100/2025. And IMS the Associated Laboratory funded by Project LA/P/0056/2020 (DOI 10.54499/LA/P/0056/2020). We are also grateful for the Fundação para a Ciência e a Tecnologia PhD grant PRT/BD/154409/2023 (DOI 10.54499/PRT/BD/154409/2023), under which this work was developed.

AUTHOR IDENTIFIERS

Author ORCIDs:
Dias RN: 0000-0002-6559-727X
Pinheiro CIC: 0000-0002-2568-3974
Serralha FN: 0000-0002-6731-3832

NOMENCLATURE

Sets and indices

n, n' - Index for enterprises in the industrial park

n' - Index for CCU units

i – Index for material code

j - Index for material category (residue, raw material, coproduct, product)

Parameters

$RD_{n,j,i}$ – Available quantity of material i in category j at enterprise n

$RD_{n,0,i}$ – Reception capacity for material i at enterprise n

$RD_{CCU_{n'',j,i}}$ – Capacity of CCU unit n'' for material i and category j

$C_{i,j}$ – Unit cost associated with material i in category j

C_{CCU} – Investment cost of a CCU unit

C_{newEM} – Investment cost associated with establishing new exchange routes

$MinEM_{n,j,i}$ – Minimum flow allowed for an exchange of material i and category j from enterprise n

CC_i – Climate change characterization factor of material i (CO_2-eq per unit of material)

$Total_{EM}^{\min}$, $Total_{EM}^{\max}$ – Minimum and maximum values of total exchanged material used in the epsilon-constraint

ε – Epsilon value enforcing the lower bound on $Total_{EM}$ in the multiobjective formulation

κ – Deterministic multiplication factor

δ – Monte Carlo band

Continuous Variables

$EM_{n,j,i,n'}$ - Material flow exchanged from enterprise n to enterprise n' for material i and category j

$EM_{CCU_{n,j,i,n''}}$ - Material flow from enterprise n to CCU unit n'' for material i and category j

$EM_{CCUR_{n'',j,i,n}}$ - Material flow from CCU unit n'' back to enterprise n for material i and category j

$Total_{EM}$ - Total material exchanged in the industrial park

$Total_{Cost}$ - Total annual cost of the industrial park configuration

Binary Variables

$Y_{Symb_{n,j,i,n'}}$ - Binary variable for symbiotic exchange between enterprises n and n' is established for material i and category j

$Y_{CCU_{n''}}$ - Binary variable for the activation of CCU units n''

REFERENCES

1. Lyu Y, Liu Y, Guo Y, Sang J, Tian J, Chen L. Review of green development of chinese industrial parks. Energy Strategy Reviews 42:100867 (2022). https://doi.org/10.1016/j.esr.2022.100867
2. Dias RN, Filipe RM, Matos HA. Sustainability analysis of a solar-driven calcium looping plant for thermochemical energy storage. Journal of Cleaner Production 429:139551 (2023). https://doi.org/10.1016/j.jclepro.2023.139551
3. United Nations Development Programme. Circular Parks Initiative. Accessed January 15, 2026. https://www.undp.org/chemicals-waste/projects/circular-parks-initiative

4. Henriques J, Ferrão P, Castro R, Azevedo J. Industrial symbiosis: a sectoral analysis on enablers and barriers. Sustainability 13:1723 (2021). https://doi.org/10.3390/su13041723
5. Paché G. Kalundborg industrial symbiosis: circular strategy in the light of mutualism. JSIS 19: (2024). https://doi.org/10.33423/jsis.v19i3.7330
6. Patricio J, Kalmykova Y, Rosado L, Cohen J, Westin A, Gil J. Method for identifying industrial symbiosis opportunities. Resources, Conservation and Recycling 185:106437 (2022). https://doi.org/10.1016/j.resconrec.2022.106437
7. Fraccascia L, Yazan DM, Albino V, Zijm H. The role of redundancy in industrial symbiotic business development: a theoretical framework explored by agent-based simulation. International Journal of Production Economics 221:107471 (2020). https://doi.org/10.1016/j.ijpe.2019.08.006
8. Dias RN, Filipe RM, Matos HA. Decision-making based on sustainability analysis using GREENSCOPE. Clean Techn Environ Policy 26:755-770 (2023). https://doi.org/10.1007/s10098-023-02647-4
9. Dias RN, Serralha MF, Pinheiro CIC. Modeling industrial symbiosis to reduce CO ₂ emissions in a portuguese industrial park. Ind. Eng. Chem. Res. 65:2836-2847 (2026). https://doi.org/10.1021/acs.iecr.5c03145
10. OpenLCA. Why we started the development of openLCA. 2005. Accessed July 5, 2023. https://www.openlca.org/the-idea/
11. Rivarolo M, Bellotti D, Magistri L, Massardo AF. Feasibility study of methanol production from different renewable sources and thermo-economic analysis. International Journal of Hydrogen Energy 41:2105-2116 (2016). https://doi.org/10.1016/j.ijhydene.2015.12.128
12. Lonis F, Tola V, Cau G. Renewable methanol production and use through reversible solid oxide cells and recycled CO2 hydrogenation. Fuel 246:500-515 (2019). https://doi.org/10.1016/j.fuel.2019.02.108
13. von der Assen N, Bardow A. Life cycle assessment of polyols for polyurethane production using co₂ as feedstock: insights from an industrial case study. Green Chem. 16:3272-3280 (2014). https://doi.org/10.1039/c4gc00513a
14. Chauvy R, Meunier N, Thomas D, De Weireld G. Selecting emerging CO2 utilization products for short- to mid-term deployment. Applied Energy 236:662-680 (2019). https://doi.org/10.1016/j.apenergy.2018.11.096
15. Dziejarski B, Krzy?y?ska R, Andersson K. Current status of carbon capture, utilization, and storage technologies in the global economy: a survey of technical assessment. Fuel 342:127776 (2023). https://doi.org/10.1016/j.fuel.2023.127776
16. Leung DYC, Caramanna G, Maroto-Valer MM. An overview of current status of carbon dioxide capture and storage technologies. Renewable and Sustainable Energy Reviews 39:426-443 (2014). https://doi.org/10.1016/j.rser.2014.07.093
17. Anicic B, Trop P, Goricanec D. Comparison between two methods of methanol production from carbon dioxide. Energy 77:279-289 (2014). https://doi.org/10.1016/j.energy.2014.09.069
18. Pérez-Fortes M, Schöneberger JC, Boulamanti A, Tzimas E. Methanol synthesis using captured CO2 as raw material: techno-economic and environmental assessment. Applied Energy 161:718-732 (2016). https://doi.org/10.1016/j.apenergy.2015.07.067
19. Fernández-Dacosta C, van der Spek M, Hung CR, Oregionni GD, Skagestad R, Parihar P, Gokak DT, Strømman AH, Ramirez A. Prospective techno-economic and environmental assessment of carbon capture at a refinery and CO2 utilisation in polyol synthesis. Journal of CO2 Utilization 21:405-422 (2017). https://doi.org/10.1016/j.jcou.2017.08.005
20. Hartmann J. New “Dream Production” CO2-to-polyols plant dedicated at Dormagen / Vision for the future is to dispense with petrochemical route. Plasteurope. 2016. Accessed July 11, 2024. https://www.plasteurope.com/news/detail.asp?id=234342
21. Mbatha S, Everson RC, Musyoka NM, Langmi HW, Lanzini A, Brilman W. Power-to-methanol process: a review of electrolysis, methanol catalysts, kinetics, reactor designs and modelling, process integration, optimisation, and techno-economics. Sustainable Energy Fuels 5:3490-3569 (2021). https://doi.org/10.1039/d1se00635e
22. Glen G, Isaacs K. Estimating sobol sensitivity indices using correlations. Environmental Modelling & Software 37:157-166 (2012). https://doi.org/10.1016/j.envsoft.2012.03.014
23. Ahmad Z, Nguyen TK, Rai A, Kim JM. Industrial fluid pipeline leak detection and localization based on a multiscale mann-whitney test and acoustic emission event tracking. Mechanical Systems and Signal Processing 189:110067 (2023). https://doi.org/10.1016/j.ymssp.2022.110067

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Techno-Economic Assessment and Optimisation of Self-Sufficient Biomethane Systems for Regional Decarbonisation

Meshkat Dolat[a], Benaissa Dekhici[a], and Michael Short[ab]*

[a] School of Chemistry and Chemical Engineering, University of Surrey, Guildford GU2 7XH, UK
[b] Institute for Sustainability, University of Surrey, Guildford GU2 7XH, UK
* Corresponding Author: m.short@surrey.ac.uk.

ABSTRACT

Existing gas network infrastructure are important national energy assets, transporting mostly fossil-derived natural gas to end-users. Biomethane, methane derived from anaerobic digestion (AD) of organic matter, presents a potential route to replace fossil fuels with home-grown renewable gas. Combined with carbon capture and storage (CCS) of the CO_2 in the biogas potentially results in carbon negative energy. This work seeks to understand the feasibility of operating a part of the gas network isolated from the main natural gas network fully on biomethane in Scotland. We present an integrated techno-economic optimisation framework for designing self-sufficient biomethane islands, applied to the Inverness network. The model, implemented as a nonlinear program (NLP), maximises annual net profit from biomethane sales and Green Gas Support Scheme (GGSS) tariffs subject to practical constraints such as GGSS-compliance of ≥50 % waste-derived biomethane, seasonal supply, land/scale, demand balancing with centralised liquefied natural gas (LNG) storage, and a life-cycle global warming potential (GWP) metric. Three archetypes are analysed: Type A (crop-dominated, manure co-digestion), Type B (food/industrial wastes, grass/manure support), and Type C (distillery residues + grass/manure). In Inverness, feasible solutions include: Type A (2 large ~92, 000 m^3 digestion plants at ~23 ha/site producing 97.39 $Mm^3\,y^{-1}$ of gas; net revenue £13.4 M y^{-1}; GWP ~42.2 $ktCO_2e/y$), Type B (1 ~97, 000 m^3 plant at ~24 ha, producing 48.69 $Mm^3\,y^{-1}$ gas; net revenue £9.4 m y^{-1}; GWP ~48.0 $ktCO_2e/y$), and Type C (2 large ~110, 000 m^3 plants at ~27 ha, producing 97.4 $Mm^3\,y^{-1}$ of gas; net revenue £13.0 m y^{-1}; GWP ~47.2 $ktCO_2e/y$). Type B is most profitable per unit capacity due to gate-fee feedstocks but carries higher GWP (mostly from grass-silage cultivation). The model balances a combination of dynamic feeding of different recipes with using a centralised LNG storage to buffers seasonal deficits and maximise asset utilisation; optional CO_2 liquefaction (~87.7 kt y^{-1} per large site at ~151 kWh t^{-1}) enables near/net-negative operation under low-carbon power. Our results find that the business model is feasible for Inverness and highlight the value of systems thinking and the need for policy reform (particularly lifting the 250 GWh y^{-1} cap for GGSS and rewarding carbon intensity rather than just waste-derived methane) to unlock larger, efficient, low-emission regional systems.

Keywords: Alternative Fuels, Modelling and Simulations, Biofuels, Renewable and Sustainable Energy, Technoeconomic Analysis

INTRODUCTION

Anaerobic digestion (AD) has become an increasingly important technology in the transition to low-carbon and circular energy systems, owing to its capacity to stabilise organic wastes, recover renewable biomethane, and reduce greenhouse-gas emissions from agriculture and food supply chains. The broader environmental and policy context for this transition is underscored by national targets for Net Zero, 2050 for the UK and 2045 for Scotland, which position renewable gas as a critical enabler of decarbonisation across heat, transport, and

industrial applications. The UK Biomass Strategy targets 30-40 TWh of biomethane by 2050, up from 6-8TWh [1], with the National Energy System Operator (NESO) suggesting 64 TWh by 2050[2]. The national gas infrastructure is an important national energy asset that risks decommissioning should fossil fuel natural gas decline as electrification of heating and transport dominate. Within these contexts, we seek to assess the feasibility, both technically and economically, of operating regional gas networks entirely on locally produced biomethane, integrating waste management, renewable energy conversion, and seasonal storage into a unified system design.

From a modelling perspective, AD has traditionally been described using mechanistic kinetic models, including simplified first-order approaches and more detailed models such as the ADM1. First-order kinetics remain widely applied in strategic decision-support settings due to their tractability and ability to capture degradation dynamics across co-digestion systems [3], but still need to assure control of methane yield, inhibition, and stability, for co-digestion by considering C/N ratios, solids loading, retention time and synergy effects. In our previous work [4], we developed a mixed-integer linear programming scheduling model to optimise feed scheduling for a single AD plant to optimise both electricity and biomethane gas-to-grid for maximising profit, considering co-digestion recipes, dynamic pricing, and gas storage. The model used a co-digestion model for understanding the gas yields from blending different feedstocks from [3] Additional literature indicates growing interest in hybrid frameworks that combine mechanistic understanding with optimisation and control, including the incorporation of price forecasting, nonlinear programming (NLP), and model predictive control (MPC) to manage temperature, substrate scheduling, and economic performance in AD systems [5].

At system level, optimisation frameworks have increasingly integrated logistics, crop availability, plant sizing, price signals, and life cycle assessments (LCA) [6], [7]. Despite this progress, relatively few studies integrate logistics, co-digestion modelling, economic optimisation, environmental performance, and policy constraints within a single modelling framework [8]. This gap is particularly relevant when systems must meet regulatory requirements such as the Green Gas Support Scheme (GGSS), which mandates that at least 50 % of biomethane be derived from eligible waste sources. This constraint strongly influences feedstock blending, reactor sizing, and plant economics due to the subsidy scheme. In practical deployment, additional challenges arise from seasonal demand variability, land availability, feedstock availability and transport distances, and the need for flexible gas-balancing infrastructure such as centralised bio-LNG storage, all of which directly affect the feasibility of establishing a self-sufficient regional gas system.

To address these challenges, working with gas networks and industrial partners, this work develops a unified NLP optimisation framework that integrates feedstock logistics, co-digestion performance, methane-release dynamics, LNG storage requirements, and life-cycle greenhouse gas emissions into a coherent decision-support tool to optimise the design and economics of AD sites that could realistically deliver flexible green gas to a part of the gas network. The framework is applied to the Inverness region as a detailed case study considering real feedstock availabilities and potential site locations, enabling the systematic evaluation of multiple plant archetypes under shared technical, economic, and policy constraints. The resulting analysis offers insight into the techno-economic viability of regional biomethane systems, the role of seasonal storage [9], and the implications of current UK policy structures for large-scale renewable-gas deployment.

METHODS

We formulate an NLP in Pyomo [10] solved with IPOPT [11]. The model couples: (i) feedstock logistics with seasonality and a 50-mile catchment from identified potential site locations; (ii) co-digestion performance via first-order decay over solids retention time (SRT) and empirical blending correlations (C:N, biodegradability); (iii) demand balancing with potential for centralised bio-LNG storage on the network; (iv) GGSS constraints (≥50 % waste share; tiering of tariffs based on gas amounts); and (v) GWP covering cultivation, transport and parasitic energy. A scale-factor adjusts digestion volume/land for lower-energy wastes and GGSS compliance; excessive scaling is penalised in the objective.

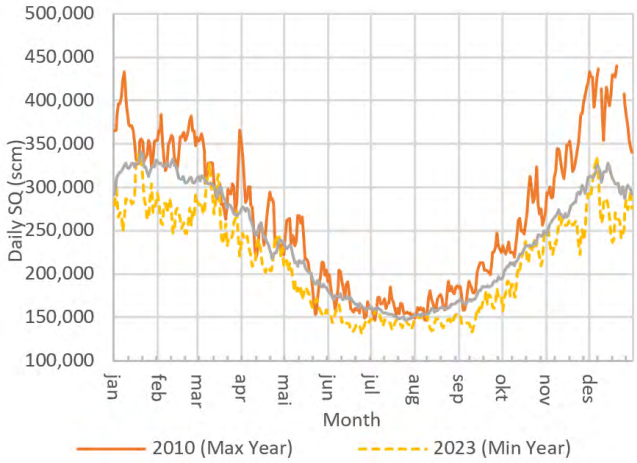

Figure 1. Inverness daily demand profiles (min/avg/max; 2010 as conservative 1-in-20 winter). Values are derived from the historic Scotland LDZ dataset.

Mathematical Formulation

The NLP combines economic, operational, and policy-driven constraints to optimise the site design and feeding patterns to meet regional gas demands at maximum profit. The mathematical model is discretised in time for weekly demand for strategic scenario analysis and sufficiently detailed to capture the dynamic methane-release behaviour, seasonal balancing requirements, and regulatory conditions (particularly the ≥50 % waste-derived biomethane threshold specified by the GGSS).

The optimisation seeks to maximise the annual net profit of the biomethane island. Revenue is derived from biomethane sales to the gas grid and, where applicable, GGSS tariff payments; costs include feedstock procurement, transport logistics, annualised capital expenditure, operational expenditures, and scale-related penalties applied when plant size significantly exceeds benchmark industrial references. This leads to the following general objective:

$$\max_{\mathbf{x}} \Pi = \left(P_{\text{gas}} Q_{\text{bm}} + T_{\text{GGSS}} Q_{\text{bm}}^{\text{elig}}\right) - \left(C_{\text{feed}} + C_{\text{OPEX}} + C_{\text{CAPEX}}^{\text{ann}} + C_{\text{transport}} + \phi_{\text{scale}}\right) \quad (1)$$

where Q_{bm} represents total biomethane production, $Q_{\text{bm}}^{\text{elig}}$ is the portion eligible for GGSS support, and $\phi_{\text{scale}}$is the nonlinear scaling penalty used to discourage unrealistic reactor upscaling beyond benchmark industrial AD site references.

A defining feature of the system is compliance with the GGSS requirement that at least 50 % of the biomethane energy content must originate from eligible wastes. Because each feedstock carries a distinct methane-yield potential, this constraint is expressed on an energy basis, ensuring that compliance cannot be met merely by mass-based blending:

$$\frac{\sum_{i \in \mathcal{W}} E_i Q_i}{\sum_{i \in \mathcal{F}} E_i Q_i} \geq 0.50 \quad (2)$$

where E_i denotes the methane energy yield per tonne of substrate i, Q_i its annual quantity, and $\mathcal{W}$ is the set of GGSS-eligible waste feedstocks from all feeds $\mathcal{F}$. This constraint is highly influential in determining plant scale, especially for crop-dominated (Type A) systems, and is responsible for driving the 35–62 % volume-expansion factors observed across Inverness scenarios.

To represent the time-dependent degradation behaviour of substrates, the model incorporates first-order kinetic decay, a widely used simplification in strategic AD planning [3]. This allows gas production to be traced over the effective SRT, enabling weekly matching between production and district gas demand:

$$M_{i,t} = \sum_{\tau=0}^{t} Q_{i,\tau}\, Y_i (1 - e^{-k_i(t-\tau)}), t \in \{1, \dots, 52\} \quad (3)$$

where $M_{i,t}$ is methane released at week t from substrate i, Y_i is its specific methane yield, and k_i a first-order rate constant. This reflects the dynamic gas-release effects, which are particularly important when high-moisture residues such as pot ale dominated blends.

The Inverness network exhibits pronounced winter peaks, requiring that surplus summer biomethane be liquefied and stored as bio-LNG for later reinjection. The required LNG capacity is defined as the maximum cumulative surplus over the annual horizon:

$$V_{\text{LNG}}^{\text{req}} = \max_{t \in [1,365]} \left(\sum_{\tau=1}^{t} Q_{\text{bm},\tau} - D_\tau\right) \cdot \rho_{\text{LNG}} \quad (4)$$

where D_t is district demand and ρ_{LNG} the gas-to-LNG volumetric conversion factor. This constraint is central to identifying feasible island-mode operation, ensuring that biomethane production and storage remain sufficient to meet peak-demand shortfalls.

CASE STUDY

The case study in this work is a real gas network in Inverness, Scotland, with data obtained from our industrial partners. We use historical regional gas demands to define the demand profiles (see Figure 1), with the 2010 demand profile used as a conservative 1-in-20 winter standard to ensure the system is designed for robustness. We use an existing industrial site as a large-scale reference, with a 57 MW class site (~500 GWh y^{-1} historically with a footprint of ~17 ha). For "real demand-scale" testing, the GGSS 250 GWh cap is relaxed while maintaining the ≥50 % waste rule; tiering is applied to eligible volumes. Available feedstocks are obtained from an existing national feedstock database from industrial partners, with their characteristics and availabilities shown in Table 1. GWP figures for crops are derived from [12]. GWP numbers provided in the results aggregates cultivation (dominant for grass silage), transport (feeds + digestate), and site parasitic load (CHP electricity, auxiliary use). CH_4 leakage is set to 0 % in the base cases but has been noted at ~3–4 % in empirical studies.

Plant Archetypes

Three representative plant archetypes were defined to capture the main categories of feedstocks available within the 50-mile Inverness sourcing radius, obtained from industrial partners and to reflect typical UK AD deployment models. These archetypes differ in substrate composition, yield characteristics, infrastructure requirements, and their ease of compliance with the GGSS ≥50 % waste-derived biomethane rule.

Energy-Crop-Dominated Systems (Type A): These plants are based primarily on energy crops (wholecrop and cereal silages), co-digested with grass silage and cattle manure to ensure stable nutrient profiles. Energy crops provide relatively high and predictable methane yields, supporting flexible scheduling to match seasonal demand. However, because cultivated crops are not GGSS-eligible wastes, Type A systems must incorporate

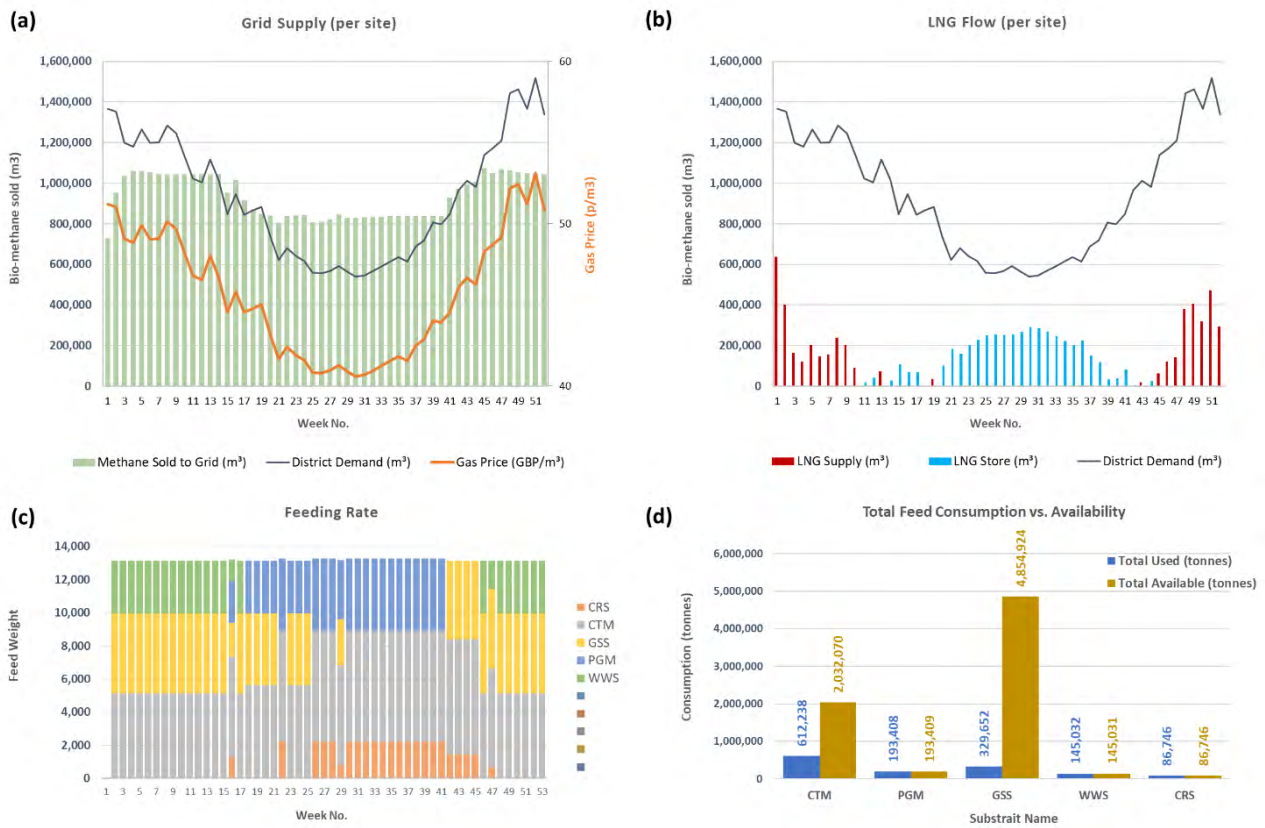

Figure 2. Operational performance – Type A plant (Inverness): (a) production vs demand; (b) LNG schedule; (c) blend ratios; (d) feed availability/use.

substantial manure volumes to meet the ≥50 % waste threshold, which drives the required 35 % scale increase (digester volume factor 1.35 in comparison to our reference industrial sites) observed in the Inverness configuration.

Food- and Industrial-Waste-Based Systems (Type B): These utilise household food waste (HFW) and industrial food waste (IFW), both of which count entirely toward the GGSS waste fraction. These substrates typically have low or negative net cost due to gate fees, enabling strong economic performance. However, their compositional variability necessitates co-digestion with grass silage and manure to maintain buffering capacity and prevent acidification. Feedstock availability limits Inverness to a single large-scale Type B plant, requiring 43 % larger digesters above benchmarks.

Distillery-Residue-Based Systems (Type C): These configurations are built around distillery by-products, particularly draff and pot ale, both abundant and fully waste-classified within the Inverness region. These substrates offer favourable co-digestion synergies with grass silage and manure, but raw pot ale's low solids content significantly reduces energy density. Consequently, Type C requires the largest scale increase of 62 % above the benchmark digester sizing, leading to higher digestion volumes and land requirements compared with Types A and B.

RESULTS

Type A: Crop-Dominated Configuration

The crop-based Type A scenario makes extensive use of energy crops (wholecrop and cereal silages), grass silage, and regionally abundant cattle manures. The optimisation determines that two large-scale AD plants constitute the maximum feasible configuration for Inverness under this scenario, with each site requiring a 35 % scale increase above the benchmark to comply GGSS requirements. This results in a digestion volume of approximately 92, 000 m³ per plant, with an associated land footprint of roughly 23 ha per site. The resulting biomethane output reaches 97.39 million m³ year⁻¹, of which 48.69 million m³ year⁻¹ is derived from waste feedstocks, meeting the GGSS threshold exactly. Table 2 summarises the techno-economic performance of the sites with annual sales of £86.2 million, feedstock costs of £38.9 million, annualised CAPEX of £19.1 million, and OPEX of £14.9 million yielding a net annual revenue of £13.4 million across the two-plant configuration to meet demands.

Figure 2a (Inverness Type A performance) shows

Table 1: Feedstocks for Type-A and Type-B plants in the selected area. Note EBMP is experimental biomethane potential and TBMP is theoretical biomethane potential. k is the rate constant, C/N is carbon to nitrogen ratio and GWP is the global warming potential associated with growing 1 tonne of the feedstock. Type-C feedstocks not shown to save space as these were not shown to be a part of the optimal configurations.

Substrate	Abbr.	Available weight	Cost	Total Solids	Volatile Solids	EBMP	TBMP	C/N	k	GWP
		Tonnes	£/t	kg_{TS} / t_{feed}	kg_{VS} / t_{TS}	m^3CH_4 / t_{VS}	m^3CH_4 / t_{VS}	-	$week^{-1}$	-
Cattle Manure	CTM	2,032,070	15	250	800	230	380	20	0.7	0
Pig Manure	PGM	193,409	10	250	800	250	410	14	1	0
Grass Silage	GSS	4,854,924	35	350	900	370	440	25	0.9	0.39
Wholecrop Wheat Silage	WWS	145,031	70	350	870	390	450	32	1	0.31
Cereal Straw	CRS	86,746	70	850	850	220	380	75	0.4	0
Household Food Wast	HFW	77,703	15	240	920	500	600	15	1.8	0
Industrial Food Waste	IFW	42,000	20	240	920	600	650	20	1.5	0

that production tracks seasonal demand through deliberate modulation of feedstock blends. Higher-yield crops are favoured during winter months, whereas lower-energy materials dominate during periods of suppressed demand. This allows the plants to maximise the digester volume usage and still hit the GGSS requirements across the year, while surpluses are still generated to maximise profit and still hit overall gas demands via gas storage. This flexible scheduling produces a yearly surplus of approximately 4.8 million m^3 $plant^{-1}$, which is redirected to the centralised LNG facility. The associated GWP for Type A totals ~42, 200 tCO_2e $year^{-1}$, driven predominantly by grass-silage cultivation (33, 400 tCO_2e) and transport impacts (8, 800 tCO_2e).

Type B: HFW and IFW Configuration

Type B shifts the focus to HFW and IFW, supported by grass silage and manure to maintain process stability. Due to limited volumes of food waste within the Inverness catchment, the optimisation identifies only one feasible large-scale site, which would need to be partnered with a Type A site to deliver the full region's gas demand. This plant requires a 43 % volume expansion over the industrial benchmark sizing, resulting in a total digestion volume of approximately 97, 000 m^3 and a land area of 24 ha.

Annual biomethane production in this configuration is 48.69 million m^3, with half (24.35 million m^3) derived from eligible waste streams. Despite being only half the scale of the Type A two-plant scenario, Type B shows strong economic performance, benefiting from the gate-fee revenues associated with food waste. As shown in Table 3, annual sales revenue reaches £43 million, with feedstock procurement costing only £16.2 million, significantly lower than in Type A. When combined with annualised CAPEX (£9.8 million) and OPEX (£7.6 million), this yields a net annual revenue of £9.4 million for a single facility.

Unlike Type A, the food-waste-dominated system exhibits less flexibility in modulating feedstock blends, resulting in a greater reliance on LNG storage to buffer seasonal mismatches. The Type B plant generates a larger relative surplus, approximately 6.4 million m^3 $year^{-1}$, equivalent to 13 % of total production. However, Type B also carries the highest GWP among the scenarios (48, 000 tCO_2e $year^{-1}$), again dominated by emissions from grass-silage cultivation (~38, 700 tCO_2e) and transportation (9, 300 tCO_2e).

Type C: Distillery Residue Configuration

The Type C scenario evaluates draff and pot-ale feedstocks from local distillery operations, supplemented by grass silage and manure. Feedstock availability supports the deployment of two large-scale plants, but low methane-potential substrates (particularly raw pot-ale) necessitate the largest scale expansion of the three scenarios: a 62 % increase (scale factor 1.62), giving a digestion volume of approximately 110, 000 m^3 per plant and a land requirement of 27 ha per site.

Combined annual biomethane production is 97.4 million m^3, with 48.7 million m^3 from eligible wastes. Total sales revenue reaches £85.8 million, with feedstock costs (£36 million), annualised CAPEX (£21 million), and OPEX (£16 million) resulting in a net revenue of £13 million $year^{-1}$, broadly comparable to Type A. Production under this scenario remains relatively stable throughout the year due to the low degree of flexibility offered by distillery residues. The system accordingly generates a larger annual surplus of ~7.2 million m^3 (15 %), which must be stored as LNG to satisfy winter deficits.

Table 2: Optimisation results- Estimated performance and economic metrics for Type-A 57MW AD plants in Inverness

Number of benchmark sites required	MW	2 x 57
Max. feasible number of sites	MW	2 x 57
Total estimated land required	ha	2 x 23
Performance/Economic Metric	**Unit**	**Total Sites Value**
Total biomethane Produced	m^3/yr	97,385,400
Waste-derived Volume	m^3/yr	48,692,700
Waste-derived Share	%	50%
Total Sell	£m/yr	86.2
Total Feed Cost	£m/yr	38.9
Total CapEx	£m	202.4
Total OpEx	£m/yr	14.9
Annualised CapEx	£m/yr	19.1
Net Revenue	£m/yr	13.4
Volume Scale Factor	-	1.35
Total Digestion Vol. per plant	m^3	92,000

Table 3: Performance, cost, and LNG implications of the selected plant configuration for Inverness biomethane island

Selected configuration		(1xA) + (1xB)
Total land required	ha	(23 ha) + (24 ha)
Total annual costs	£m/yr	69.9
Total net revenue	£m/yr	16.1
Total GWP	tCO_2e	90,200
Total LNG storage	m^3LNG	18,500
Cent. LNG facility total capital	£m	67
Cent. LNG facility Opex	£m	2.0

The Type C configuration exhibits a total GWP of approximately 47, 200 tCO_2e $year^{-1}$, with the majority (78 %) attributable to grass-silage cultivation and the remainder to logistics. Notably, the exclusion of pot-ale syrup, a higher-solids, higher-methane-potential variant, suggests that future scenarios including concentrated distillery co-products could significantly reduce both plant size and GWP.

Optional CO_2 liquefaction (per large site)

For each large-scale Inverness AD facility, the biogenic CO_2 stream from upgrading can be processed via an optional on-site liquefaction unit. Our model estimates an annual CO_2 flow of ~87, 700 t CO_2 $year^{-1}$ per site, requiring ~£61 million in capital expenditure and ~£0.71 million $year^{-1}$ in O&M costs, with electricity demand of 151 kWh t^{-1} CO_2. These energy requirements are treated as a parasitic load on the CHP system rather than external grid purchases, ensuring consistency with the broader techno-economic system boundary. Although optional, CO_2 liquefaction has the potential to shift AD systems toward near- or net-negative emissions, particularly when the captured CO_2 can be diverted for sequestration or utilised in local industrial markets.

The feasibility of implementing CO_2 liquefaction, however, depends heavily on future policy incentives and the evolution of CO_2 commodity pricing. The current GGSS framework does not reward CO_2 capture, meaning the investment must presently be justified solely on environmental grounds, as current UK CO_2 market prices do not provide a profitable investment case [13]. Shifting policy from a strict ≥50 % waste-content rule toward a carbon-intensity-based support mechanism would materially improve the economic case for CO_2 capture and valorisation. As emerging low-carbon markets develop, particularly in Scotland's food, beverage, and synthetic-fuel sectors, liquefied CO_2 could represent a meaningful future revenue stream, strengthening the long-term viability of regional biomethane systems.

DISCUSSION

The optimisation results for Inverness reveal that system feasibility depends on the interplay between feedstock availability, GGSS constraints, and seasonal balancing needs. The most profitable solution identified in the report is a hybrid configuration comprising one Type A and one Type B plant based on feedstock availability. This combination leverages the high methane yield and operational stability of crop-supported co-digestion (Type A) together with the strong economic performance of food-waste-based digestion (Type B), achieving £16.1 million year^{-1} in net revenue and requiring 18, 500 m^3 of centralised LNG storage to meet winter peak demand. The details of the optimal system is given in Table 3.

This hybrid arrangement also moderates the limitations of the individual systems. By combining the two plant types, the system achieves a more favourable balance of economic efficiency and environmental performance, while reducing sensitivity to feedstock variability. Seasonal surplus remains critical across all scenarios, necessitating dedicated centralised LNG infrastructure.

CONCLUSION

This work applied a unified techno-economic optimisation framework to assess the feasibility of an Inverness biomethane island. We showed that it is feasible in this region to meet all current gas demands within current policy and feedstock availability restrictions with a novel arrangements of technologies and feeding patterns, while still providing an attractive business case. Future work should examine hybrid designs that combine the economic strengths of Type B with the lower emissions and operational flexibility of crop- or distillery-based systems, as well as the potential of higher-value co-products such as pot-ale syrup to reduce required plant scale. Policy implications are significant: the 250 GWh y^{-1} GGSS cap and binary ≥50 % waste rule limit optimal system design that do not encourage flexible operation. Shifting toward carbon-intensity-based incentives, alongside support for CO_2 capture, would better align economic and environmental outcomes and accelerate deployment.

ACKNOWLEDGEMENTS

Biomethane Islands (NIA2_SGN0072) is funded via the Network Innovation Allowance. We acknowledge project partners SGN, Green Giraffe B.V., NNFCC Ltd and Axis PED Ltd. for data and review. We also acknowledge funding from Engineering and Physical Research Council and Biotechnology and Biological Sciences Research Council via the Supergen Bioenergy Hub (EP/Y016300/1) and AI for Net Zero (EP/Y005600/1) projects.

REFERENCES

1. D. for Energy Security and N. Zero, "Biomass Strategy 2023, " 2023. [Online]. Available: www.gov.uk/official-documents
2. Cadent, "Accelerating biomethane in the UK, " 2025. Accessed: Feb. 01, 2026. [Online]. Available: https://cadentgas.com/getContentAsset/4ba7a0cd-4b18-478e-938c-f109400b4c73/1edc10b3-193a-4a87-9cfc-cbb68531e06b/Cadent_Accelerating-Biomethane-in-the-UK.pdf?language=en
3. Moretta F, Goracci A, Manenti F, Bozzano G. Data-driven model for feedstock blending optimization of anaerobic co-digestion by BMP maximization. Journal of Cleaner Production 375:134140 (2022). https://doi.org/10.1016/j.jclepro.2022.134140
4. Dolat M, Murali R, Zarei M, Zhang R, Pincam T, Liu YQ, Sadhukhan J, Bywater A, Short M. Dynamic feed scheduling for optimised anaerobic digestion: an optimisation approach for better decision-making to enhance revenue and environmental benefits. Digital Chemical Engineering 13:100191 (2024). https://doi.org/10.1016/j.dche.2024.100191
5. Zarei M, Dolat M, Murali R, Zhu M, Pennington O, Zhang D, Short M. Real-time dynamic optimisation for sustainable biogas production through anaerobic co-digestion with hybrid models. Systems and Control Transactions 4:2423-2428 (2025). https://doi.org/10.69997/sct.130144
6. Mayerle SF, Neiva de Figueiredo J. Designing optimal supply chains for anaerobic bio-digestion/energy generation complexes with distributed small farm feedstock sourcing. Renewable Energy 90:46-54 (2016). https://doi.org/10.1016/j.renene.2015.12.022
7. Ó Céileachair D, O'Shea R, Murphy JD, Wall DM. The effect of seasonal biomass availability and energy demand on the operation of an on-farm biomethane plant. Journal of Cleaner Production 368:133129 (2022). https://doi.org/10.1016/j.jclepro.2022.133129
8. Murali R, Bywater A, Dolat M, Dekhici B, Zarei M, Hilton L, Sadhukhan J, Zhang D, Short M. Anaerobic digestion site-wide optimisation and decision-making: an industrial perspective and review. Renewable and Sustainable Energy Reviews 226:116402 (2026). https://doi.org/10.1016/j.rser.2025.116402
9. He F, Short M, Chen Q, Liu L. Integrating multi-timescale energy storage into net-zero electricity systems under evolving technologies and policy environment: insights from the united kingdom's case study. Energy Conversion and Management 352:121076 (2026). https://doi.org/10.1016/j.enconman.2026.121076
10. Hart WE, Watson JP, Woodruff DL. Pyomo: modeling and solving mathematical programs in

python. Math. Prog. Comp. 3:219-260 (2011). https://doi.org/10.1007/s12532-011-0026-8
11. Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math. Program. 106:25-57 (2005). https://doi.org/10.1007/s10107-004-0559-y
12. R. Zhang *et al.*, "Novel Life Cycle GHG Formulations of Anaerobic Digestion Systems Aligned with Policy, " 2024, doi: 10.2139/ssrn.4837715.
13. Zhang D, Li D, Bywater A, Short M, Sadhukhan J. Carbon credits monetary value for anaerobic digestion systems and energy policy implication in the UK. The Innovation Energy 2:100066 (2025). https://doi.org/10.59717/j.xinn-energy.2024.100066

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Dynamic material flow analysis of iridium circularity in proton exchange membrane water electrolysers in Japan

Shu Yamaki[a], Ayumi Yamaki[b], Heng Yi Teah[c], Yuichiro Kanematsu[c], and Yasunori Kikuchi[abc]*

[a] The University of Tokyo, Department of Chemical System Engineering, Tokyo, Japan
[b] The University of Tokyo, Institute for Future Initiatives, Tokyo, Japan
[c] The University of Tokyo, Presidential Endowed Chair for "Platinum Society" Organization for Interdisciplinary Research Project, Tokyo, Japan
* Corresponding Author: ykikuchi@ifi.u-tokyo.ac.jp

ABSTRACT

Achieving and sustaining net-zero greenhouse-gas emissions will require the long-term deployment of green hydrogen. Proton exchange membrane water electrolysers (PEMWEs) are attractive for variable renewable electricity (VRE) because of their fast dynamic response; however, they rely on iridium (Ir) anode catalysts, and Ir supply is severely constrained. Here, a Japan-specific dynamic material flow analysis (DMFA) model is developed for 2025–2100 to quantify Ir circularity in PEMWE deployment under a backcasting-oriented hydrogen production pathway. The model tracks Ir in anode catalysts only and represents: (i) Ir demand for new capacity additions and replacements, (ii) end-of-life (EoL) outflows governed by a Weibull lifetime distribution, and (iii) closed-loop recycling characterised by an overall recycling rate across collection, separation/pre-processing, and refining. Sensitivity analyses show that long-term primary Ir requirements are governed by the coupled effects of catalyst durability and recycling performance. Although lowering Ir loading reduces initial Ir intensity, accelerated degradation can increase replacement frequency and cumulative primary Ir demand unless recycling rate improves concurrently. Importantly, conclusions depend on the analysis horizon: Ir loading dominates cumulative demand up to 2050, whereas recycling becomes increasingly decisive toward 2100. These results highlight that Ir-thrifting catalyst designs must be evaluated jointly with durability under dynamic operation and recyclability to ensure scalable and continuous PEMWE deployment.

Keywords: Hydrogen, Proton Exchange Membrane Water Electrolyser, Material Flow Analysis, Iridium

INTRODUCTION

Green hydrogen is expected to play a central role in decarbonizing industrial sectors where deep emissions reductions are difficult without major production cuts. Water electrolysis is a key pathway for green hydrogen production. Among electrolyser technologies, proton exchange membrane water electrolysers (PEMWEs) are attractive because their fast transient response is compatible with variable renewable electricity (VRE) such as photovoltaics [1]. However, state-of-the-art PEMWEs rely on iridium-based catalysts at the anode to sustain oxygen evolution in strongly acidic environments [2]. Iridium (Ir) is among the scarcest metals, with global annual production on the order of only about 8 t [3]. This scarcity has motivated extensive research into reducing Ir usage through lower catalyst loadings [4].

A critical challenge is that Ir-thrifting strategies can compromise durability. Several studies report accelerated degradation at low catalyst loading and under dynamic operation [5]. From a systems perspective, faster degradation shortens electrolyser lifetime, increases replacement frequency, and enlarges the end-of-life (EoL) flow requiring recycling. These effects can raise the risk of Ir dissipation along the recycling chain and ultimately undermine circularity. Therefore, evaluating Ir circularity requires a framework that links catalyst loading, degradation-driven lifetime, and recycling rather than treating them independently.

Dynamic material flow analysis (DMFA) provides a time-resolved approach for quantifying primary demand, in-use stock, and EoL flows of critical materials, thereby

informing recycling infrastructure planning and technology targets. Existing Ir-focused DMFA studies remain limited both globally [6–8] and at the country level [9], and many studies emphasize time horizons up to 2050, the net-zero target year for many countries. Net-zero, however, is not an endpoint: maintaining a net-zero state requires continued infrastructure renewal and resource management beyond 2050.

In this study, a Japan-specific DMFA model is developed for Ir in PEMWEs through 2100 and quantifies how catalyst properties (Ir loading and degradation rate) and recycling performance jointly govern long-term primary Ir requirements. Japan is a relevant case due to its industrial scale, potential demand for green hydrogen, and notable role in Ir-consuming products (e.g., spark plugs, crucibles), which together position Japan as an important node in the global Ir cycle. The analysis also highlights durability and recycling priorities that connect catalyst/material chemistry to long-term material circularity.

MATERIALS AND METHODS

Scenario and scope

This study adopts a Japan-focused decarbonisation scenario [10] in which only photovoltaic-derived VRE is available as the electricity supply, and PEMWEs are selected because of their operational responsiveness. The model period is 2025–2100, and commercial introduction of PEMWEs is assumed to begin in 2030. The system boundary includes only Ir in the PEMWE anode catalyst.

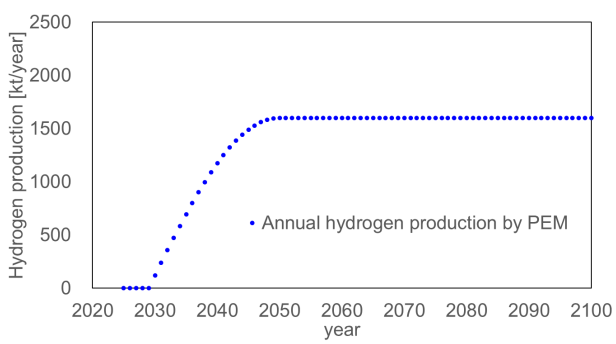

Figure 1. Scenario-based trajectory of annual hydrogen production in Japan from 2030 to 2100.

The objective is to estimate the primary Ir requirement needed to achieve and sustain the specified hydrogen production pathway. To isolate the effects of catalyst loading, durability, and recycling performance, technological progress is not explicitly represented. On the demand side, Japan's ongoing population decline and the expected contraction of domestic manufacturing motivate a simplified trajectory in which PEMWE-based hydrogen production increases to 1.6 Mt/year by 2050 and remains constant thereafter (**Figure 1**). This PEMWE contribution (1.6 Mt/year) is less than 10% of Japan's 2050 hydrogen target of 20 Mt.

DMFA model

Annual Ir demand for PEMWE deployment in year t, Ir_t^d [g], is calculated from (i) the additional hydrogen production required between years $t-1$ and t, and (ii) Ir associated with end-of-life PEMWE equipment in year $t-1$. Let H_t [kg] be the annual target hydrogen production, EI_t [kWh/kg-H_2] the electrical energy intensity of PEM water electrolysis, MI_t [g/kW] the Ir material intensity of new equipment (hereafter referred to as Ir loading), and h [h/year] the annual operating hours. The required additional capacity ΔC_t is:

$$\Delta C_t = \max(0, H_t - H_{t-1}) EI_t / h \tag{1}$$

The corresponding Ir demand is:

$$Ir_t^d = \Delta C_t MI_t + Ir_{t-1}^w \tag{2}$$

where Ir_{t-1}^w [g] denotes Ir contained in waste (retired) PEMWE equipment. Recycled Ir, Ir_t^r [g], is computed using an overall recycling rate RR [–], defined as the yield across all processes required to return Ir to a state suitable for reuse in PEMWE manufacturing (collection, separation, and refining):

$$Ir_t^r = Ir_{t-1}^w RR \tag{3}$$

Primary Ir demand is then:

$$Ir_t^p = Ir_t^d - Ir_t^r \tag{4}$$

Recycled Ir is prioritized for PEMWE manufacturing, and any remaining demand is met by primary Ir.

End-of-life outflow and lifetime distribution

The EoL outflow Ir_t^w is determined by applying a lifetime distribution to historical inflows of Ir embedded in newly manufactured and replacement PEMWE anode catalysts. Specifically, Ir_t^w is computed as the convolution of past inflows with the discrete retirement probability implied by a Weibull distribution. The Weibull shape parameter is fixed (**Table 1**), and the scale parameter is calibrated so that the resulting lifetime distribution is consistent with the lifetime L defined below, enabling time-dependent replacement and recycling needs to be represented within a consistent stock–flow framework.

Lifetime modeling and parameter settings

Previous DMFA studies mainly assume a uniform lifetime [6–8]. Here, lifetime is linked to electrochemical durability via the voltage degradation rate. End-of-life is defined using the US Department of Energy (DOE) criterion of a 10% voltage increase at fixed current density relative to beginning-of-life [12]. If the average degradation rate is γ [%/kh], the corresponding lifetime L [h] is:

$$L = \frac{10}{\gamma} \times 1000 \quad (5)$$

For example, $\gamma = 0.25$ [%/kh], the US DOE "2022 Status" value, corresponds to $L = 40,000$ h. This definition connects catalyst durability to replacement frequency, thereby linking catalyst chemistry/structure to long-term Ir circularity.

Because the scenario adopted in this analysis assumes photovoltaic-derived VRE as the power source, the average operating time was constrained to 2920 h/year (8 h/day). With respect to operational variability under VRE, it was assumed to have no effect on the hydrogen production stage; during operating hours, hydrogen production was assumed to proceed at a constant rate.

Table 1: Parameter settings for DMFA calculation.

Fixed Parameters	Value
Weibull shape parameter	5
End of life voltage loss [%]	10
Average operation hour [h/year]	2920
Electrical efficiency [kWh/kg-H_2]	51

RESULTS AND DISCUSSION

DMFA and sensitivity analysis results

Figure 2 illustrates the time evolution of total Ir demand, primary Ir demand, and recycled Ir supply for an illustrative parameter set intended to reduce primary Ir dependence: Ir loading 0.05 g/kW (below the US DOE 2026 target), recycling rate 0.9 and degradation rate 0.25 %/kh corresponding to a 40, 000 h lifetime. Under these assumptions, the cumulative primary Ir demand by 2100 is limited to approximately 2.0 t, a level that could be secured by a single country.

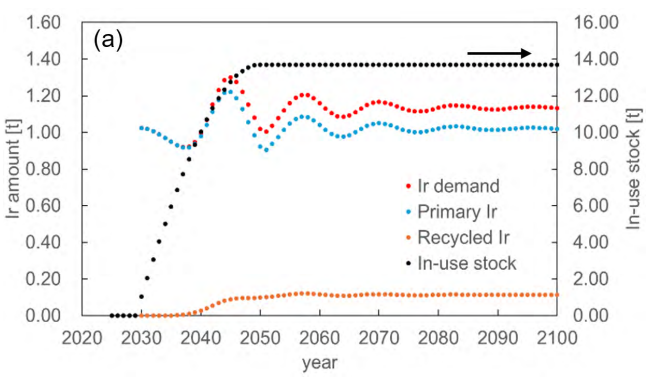

Figure 2. DMFA results (Ir loading = 0.05 g/kW, recycling rate = 0.9, degradation rate = 0.25 %/kh).

Utilizing this DMFA model, **Figure 3** presents the results of a sensitivity analysis in which Ir loading is held constant (0.05 g/kW) and the degradation rate and recycling rate are varied as input parameters, with cumulative primary Ir demand in 2100 as the output. Because global annual Ir production is approximately 8 t, it would be challenging for Japan alone to consume several tens of tons of Ir. **Figure 3** indicates that high recycling alone cannot compensate for rapid degradation, and conversely that durability gains deliver limited benefits when recycling losses are large. These results motivate Ir-reduction strategies that improve both durability and recycling performance rather than optimising lever in isolation.

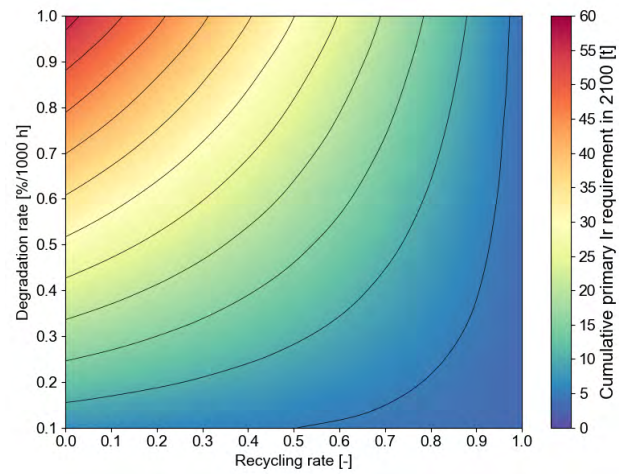

Figure 3. Results of sensitivity analysis on cumulative primary Ir demand in 2100 using recycling rate and degradation rate as parameters. Ir loading is fixed at 0.05 g/kW.

Relationship between Ir loading and degradation rate

Several studies report that reducing Ir content can worsen degradation [5], implying that Ir loading and degradation rate should not be treated as independent. Accordingly, the relationship between Ir loading and degradation rate is parameterized using literature data: a linear function is derived based on Ir loading and degradation-rate data reported in the US DOE "2022 Status" [12], together with additional literature data [13] obtained at the same current density (2 A/cm²).

Figure 4 and **Figure 5** show sensitivity results using Ir loading and recycling rate as input variables, with degradation rate determined from the parameterized Ir loading and degradation rate relationship. A non-monotonic response emerges: lowering Ir loading reduces Ir intensity per unit capacity, but it can also shorten lifetime by accelerating degradation, thereby increasing replacement frequency. Consequently, cumulative primary Ir demand can increase in a low-loading region when accelerated replacement outweighs the benefit of reduced initial Ir intensity. This is a critical systems-level implication of durability–loading trade-offs: Ir-thrifting strategies can

backfire if durability degradation dominates.

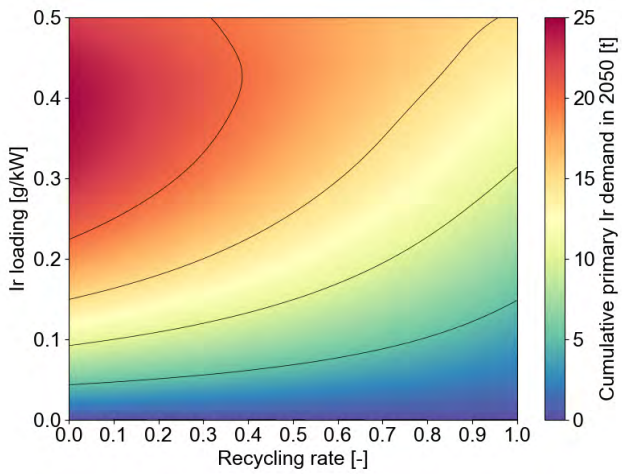

Figure 4. Results of sensitivity analysis on cumulative primary Ir demand in 2050 using recycling rate and Ir loading as parameters.

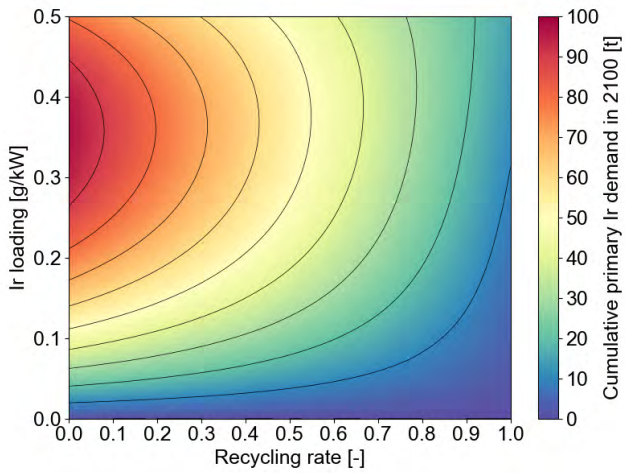

Figure 5. Results of sensitivity analysis on cumulative primary Ir demand in 2100 using recycling rate and Ir loading as parameters.

The analysis horizon strongly affects the apparent priority of interventions. Up to 2050, reducing Ir loading has a larger effect than increasing recycling rate, because the system is still building in-use stock and cumulative replacements are limited. By 2100, however, the system is dominated by replacement and recycling loops rather than initial stock build-out, and improving recycling rate becomes more decisive. Therefore, DMFA conclusions intended to inform carbon-neutrality planning should state the terminal year explicitly.

Figure 6 illustrates the mechanism behind the increase in cumulative primary Ir demand observed in **Figure 4** and **5** despite reduced Ir loading. At a low recycling rate (0.1), comparing a higher loading case (0.50 g/kW) with a lower loading case (0.35 g/kW) shows that the higher-loading case has a larger in-use stock but benefits from a longer lifetime and therefore a lower EoL frequency. In contrast, the lower-loading case experiences more frequent replacements, which increases annual Ir demand over time and elevates cumulative primary Ir demand. Accordingly, catalyst development should target low Ir loading while maintaining durability under dynamic operation, and electrode designs should explicitly consider recyclability and separability to increase the overall recycling rate.

Limitations of this study

First, the relationship between Ir loading and degradation rate depends on catalyst state, distribution, and electrode architecture, and it is not necessarily linear. Due to limited long-duration durability data under realistic dynamic operation, this study assumed a proportional relationship derived from available datasets. More comprehensive durability testing under VRE-relevant operating profiles is required.

Second, although the analysis extends to 2100, parameters are assumed constant over time to isolate effects, and long-term technological progress is not explicitly represented. This may overestimate Ir requirements if

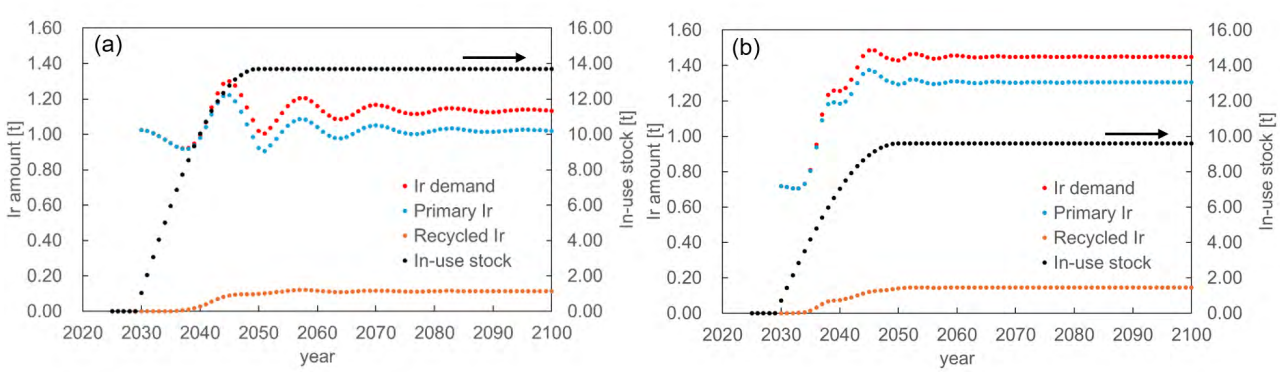

Figure 6: DMFA results considering the relationship between Ir loading and degradation (recycling rate = 0.1) (a) Ir loading = 0.5 g/kW; (b) Ir loading = 0.35 g/kW.

future innovations improve efficiency, durability, or recycling. Conversely, the scenario assumes that PEMWE-supplied hydrogen demand remains constant after 2050, which may underestimate Ir requirements if demand continues to grow.

Third, the system boundary includes only Ir in the anode catalyst and excludes other Ir uses and potential competition across sectors. Including cross-sectoral coupling could further constrain available Ir and alter feasible deployment pathways.

CONCLUSIONS

This study developed a Japan-specific DMFA model to quantify long-term Ir circularity in PEMWE deployment through 2100, explicitly linking catalyst loading, degradation-driven lifetime, and the overall recycling rate. The results show that primary Ir dependence is governed by the coupled effects of durability and recycling: lowering Ir loading reduces initial Ir intensity, but it can increase cumulative primary Ir demand when it accelerates degradation and replacement frequency, particularly under low recycling rate. The time horizon further conditions these conclusions, with Ir loading effects dominating up to 2050 and recycling performance becoming increasingly decisive toward 2100.

These findings indicate that Ir-thrifting strategies must be pursued alongside (i) durability improvements under dynamic operation and (ii) improvements in the overall recycling rate across separation/pre-processing and refining to ensure scalable and sustainable PEMWE-based hydrogen deployment. In practice, robust Ir circularity requires integrated catalyst and electrode design for both stability and recyclability, together with strategic planning of recycling infrastructure.

ACKNOWLEDGEMENTS

This research contained the achievements of JST COI-NEXT (JPMJPF2003). This research contained the achievements of the JST-MIRAI Program Grant Number JPMJMI19C7, and JSPS KAKENHI (Grant Number JP21H03660). Activities of the Presidential Endowed Chair for “Platinum Society” at the University of Tokyo are supported by Mitsui Fudosan Corporation, Sekisui House, Ltd., the East Japan Railway Company, and Toyota Tsusho Corporation.

AUTHOR IDENTIFIERS

Author ORCIDs:
Yamaki S: 0009-0004-8594-0962
Yamaki A: 0000-0003-4162-7806
Kikuchi Y: 0000-0003-0002-8233
Heng YT: 0000-0003-2923-4938
Kanematsu Y: 0000-0001-9978-7795

REFERENCES

1. Shiva Kumar S, Lim H. An overview of water electrolysis technologies for green hydrogen production. *Energy Reports.* 8:13793–13813 (2022). http://dx.doi.org/10.1016/j.egyr.2022.10.127
2. Zhang K, Ni X, Liang X, Zou X. Iridium management strategies for scalable proton exchange membrane water electrolysis. *Chem Catalysis.* (2025). http://dx.doi.org/10.1016/j.checat.2025.101570
3. JohnsonMatthey. *PGM market report May 2025.*
4. Sezer N, Bayhan S, Fesli U, Sanfilippo A. A comprehensive review of the state-of-the-art of proton exchange membrane water electrolysis. *Materials Science for Energy Technologies.* 8:44–65 (2025). http://dx.doi.org/10.1016/j.mset.2024.07.006
5. Alia SM, Stariha S, Borup RL. Electrolyzer Durability at Low Catalyst Loading and with Dynamic Operation. *Journal of The Electrochemical Society.* 166(15):F1164–F1172 (2019). http://dx.doi.org/10.1149/2.0231915jes
6. Clapp M, Zalitis CM, Ryan M. Perspectives on current and future iridium demand and iridium oxide catalysts for PEM water electrolysis. *Catalysis Today.* 420 (2023). http://dx.doi.org/10.1016/j.cattod.2023.114140
7. Minke C, Suermann M, Bensmann B, Hanke-Rauschenbach R. Is iridium demand a potential bottleneck in the realization of large-scale PEM water electrolysis? *International Journal of Hydrogen Energy.* 46(46):23581–23590 (2021). http://dx.doi.org/10.1016/j.ijhydene.2021.04.174
8. Riedmayer R, Paren BA, Schofield L, Shao-Horn Y, Mallapragada D. Proton Exchange Membrane Electrolysis Performance Targets for Achieving 2050 Expansion Goals Constrained by Iridium Supply. *Energy & Fuels.* 37(12):8614–8623 (2023). http://dx.doi.org/10.1021/acs.energyfuels.3c01473
9. Kiemel S, Smolinka T, Lehner F, Full J, Sauer A, Miehe R. Critical materials for water electrolysers at the example of the energy transition in Germany. *International Journal of Energy Research.* 45(7):9914–9935 (2021). http://dx.doi.org/10.1002/er.6487
10. CGC. *Net Zero Japan 2050 -Summary for Business Leaders- Interim Report on Decarbonization Scenarios for 2050* 2023.
11. Lallana M, Torrubia J, Valero A. Metals for energy & digital transition in Spain: Demand, recycling and sufficiency alternatives. *Resources, Conservation and Recycling.* 205 (2024). http://dx.doi.org/10.1016/j.resconrec.2024.107597

12. U.S.DOE. *Technical Targets for Proton Exchange Membrane Electrolysis* 2022.
13. Arshad M, Bashir A, Hassan H, et al. Recent advances in catalysts and interface engineering for high-performance proton exchange membrane water electrolyzers. *Materials Science and Engineering: R: Reports.* 167 (2026). http://dx.doi.org/10.1016/j.mser.2025.101124

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Value-Based Assessment for Strategic Selection and Optimization of POME Valorization Pathways

Yiwei Gao[a], Heng Yi Teah[b], Yuichiro Kanematsu[b], Yasunori Kikuchi[abc]*

[a] Department of Chemical System Engineering, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan
[b] Presidential Endowed Chair for "Platinum Society", The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan
[c] Institute for Future Initiatives, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8654, Japan
* Corresponding Author: ykikuchi@ifi.u-tokyo.ac.jp

ABSTRACT

Palm oil mill effluent (POME) represents a major environmental burden in the palm oil industry while offering opportunities for resource recovery. This study develops and applies a value-based assessment framework to examine how technological choice influences the integrated environmental–economic performance of POME valorization. Biomethane production and bio-hydrogen production are selected as representative mature and emerging technologies, respectively. Life-cycle environmental performance is quantified using greenhouse gas (GHG) emissions midpoint indicator and natural resources endpoint indicator, reflecting broader environmental damages. A techno-economic assessment is performed to show the economic performance. In addition to conventional return of investment (ROI), the benefits of mitigating environmental impacts are accounted using the return of value (ROV) methodology. The results indicate that the attractiveness of POME valorization pathways depends strongly on how environmental gains are represented. Midpoint-based valuation shows the climate mitigation potential of the biomethane and bio-hydrogen, whereas endpoint-based valuation reveals higher environmental impacts for the energy-intensive hydrogen route, favoring the mature biomethane option in terms of integrated sustainability performance. This contrast demonstrates that technological maturity is reflected not only in cost structure and process stability but also in the balance between avoided impacts and newly introduced environmental burdens. Overall, the proposed value-based assessment framework enables environmental impacts and economic performance to be evaluated on a common value scale, providing a consistent basis for sustainability-oriented decision-making in POME valorization and supporting strategic planning for wastewater management and bio-economy development in the palm oil sector.

Keywords: Life Cycle Assessment, Wastewater valorization, Palm oil mill effluent, Bioenergy pathways, Value-based assessment

INTRODUCTION

The palm oil industry plays a pivotal role in the global vegetable oil market, accounting for approximately 33% of global oils and fats production and nearly 40% of international edible oil trade [1]. However, the large-scale processing of fresh fruit bunches (FFB) into crude palm oil is highly water-intensive, particularly during sterilization and clarification stages, resulting in the generation of substantial quantities of POME. On average, approximately 0.1 tons of effluent per ton of produced crude palm oil is generated as POME. [2]. POME is characterized by high chemical oxygen demand (COD), biological oxygen demand (BOD), oil and grease content, and acidic pH [3]. If inadequately managed, conventional open-pond treatment systems can lead to severe water pollution and significant methane emissions, contributing to climate change [4]. To address these challenges, POME valorization has emerged as an important strategy to transform wastewater from an environmental burden into a valuable resource. Importantly, the POME valorization technologies span a wide range of technology readiness levels (TRLs). On one end, mature technologies have already been deployed at industrial scale, delivering

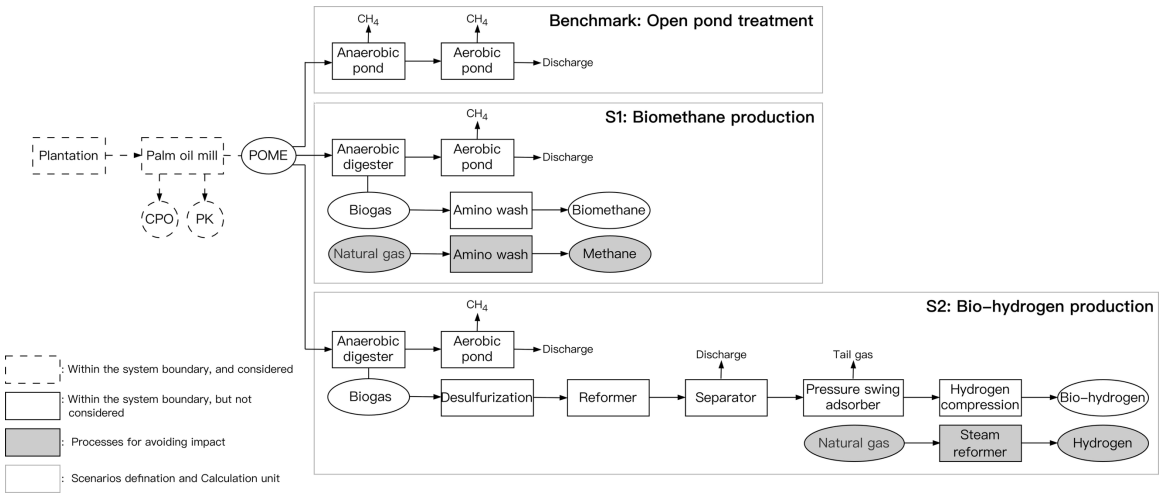

Figure 1: Cradle-to-gate system boundary of POME valorization pathways for benchmark and two valorization

near-term environmental benefits and relatively stable economic returns. On the other end, emerging technologies remain at pilot or laboratory stages, offering potentially higher long-term value but accompanied by greater technical uncertainty, higher costs, and less certain environmental performance. This coexistence of mature and emerging valorization options poses a critical challenge for decision-makers: how to systematically compare technologies with fundamentally different maturity levels on a consistent and transparent basis. Most existing studies evaluate single POME valorization pathways using either environmental indicators, such as life cycle greenhouse gas (GHG) emissions, or economic metrics, such as production cost and profitability, or evaluate these two aspects individually [5, 6]. However, these single-dimensional assessments fail to explicitly integrate environmental impacts with economic performance into a single level to enable holistic comparison across heterogeneous technologies.

This study aims to first characterize the life-cycle performance of POME valorization by examining two representative pathways corresponding to mature and emerging technologies. Building on this characterization, a value-based assessment framework is applied in which environmental impacts are represented in monetary value and added to the economic indicator. The objective is to elucidate how technological choice influences the integrated environmental–economic performance of POME valorization and to provide a quantitative basis for strategic technology selection in sustainable wastewater management and bio-economy development in the palm oil sector.

METHOD

Scenario Design and Settings

To reflect the technological heterogeneity of POME valorization options, this study selects two representative pathways in different levels of technological maturity. Biomethane production via anaerobic digestion is chosen as the mature technology because it has been widely implemented at industrial scale [7] and its relatively stable technical performance and production costs. In contrast, bio-hydrogen production from POME via steam reforming is selected as the emerging technology. Although hydrogen is associated with higher market value and strategic relevance for future low-carbon energy systems [8], this pathway remains at pilot or laboratory scale, leading to higher technical uncertainty and production costs. These two pathways are therefore used as representative cases to examine how differences in technological maturity and value potential influence the integrated environmental–economic performance of POME valorization.

The two scenarios are: S1, POME-to-biomethane with biogas capture and purification, and S2, POME-to-bio-hydrogen via steam reforming. A conventional open-pond treatment without methane removal is used as a benchmark for comparison.

Life Cycle Assessment

This life cycle assessment (LCA) aims to evaluate environmental impact of POME treatment and valorization pathways. The process modeling is based on a mass and energy balance framework, in which each unit operation is represented by literature-based input–output data and linked to corresponding datasets in ecoinvent v3.11. The functional unit (FU) is defined as 1 t of POME utilization. The system boundary is illustrated in **Figure 1** and focuses on POME treatment and down-stream valorization processes. Upstream plantation and palm oil mill (POM) operations are not considered, and POME is considered as the sole feedstock entering the system. The

benchmark includes processes of anaerobic and aerobic ponds, and the treated effluent is discharged, while methane generated during treatment is emitted to the atmosphere without removal. In S1, POME is treated by an anaerobic digester (AD) followed by aerobic ponds. The biogas produced from anaerobic digestion is purified through an amino washing process to produce biomethane, which is assumed to substitute fossil-derived methane. In S2, the system boundary is further extended to include biogas desulfurization, steam reforming, gas separation, pressure swing adsorption, and hydrogen compression for bio-hydrogen production. The produced bio-hydrogen is assumed to substitute hydrogen derived from fossil-based methane.

Inventory datas for material, energy and process shown in **Table 1** are primarily derived from literature sources [5, 8]. These data are scaled to the defined FU using proportional coefficients, according to:

$$Flow_{FU} = Flow_{ref} \times \frac{FU}{FU_{ref}} \quad (1)$$

where $Flow_{FU}$ is the inventory flow per FU, and $Flow_{ref}$ is the reference flow reported in the literature for a given FU FU_{ref}. All inventory flows are scaled linearly to the functional unit based on reported reference data, assuming steady-state operation and constant conversion efficiencies.

For the benchmark, the generated biogas is considered as direct emissions. For S1, inventory data include water and electricity input, and the purification process. A biomethane purification efficiency of 97% is assumed, while the remaining fraction is released as losses and accounted for in the emission calculation. For S2, additional material and energy inputs include monoethanolamine (MEA) for gas treatment and metal oxides (iron oxide, copper oxide, and nickel) as catalysts for reforming and downstream conversion. In this study, catalyst recovery and re-use are not considered. Steam consumption includes low-, medium-, and high-pressure steam, which are represented by an average energy intensity of 2.80 MJ/kg steam to calculate both steam demand and its associated environmental impacts. Tail gas is also generated from the pressure swing adsorption unit. The composition consists of 94.66 % CO_2, 2.00 % H_2S, 1.25 % water, 0.93 % NH_3, 0.40 % CH_4, 0.40 % C_2H_7NO, 0.21 % H_2, 0.11 % $C_{13}H_{25}O_7N_3S$, and 0.04 % CO [8]. It is assumed to be directly released, and only CO_2 (587 kg) and CH_4 (2.48 kg) are included in the GHGs calculation. Wastewater and digestate from two valorization pathways are estimated based on mass balance. For the fossil-derived alternatives, cradle-to-gate environmental impacts are adopted from ecoinvent v3.11 database, which includes natural gas production, amine washing, and steam reforming.

Table 1: Inventory data of each scenario.

S1	Material/Energy	Unit	Value
Input	POME	t	1
	Water	m^3	0.01
	Electricity	kWh	0.03
	Biogas-to-biomethane purification (amine washing)	m^3	20.79
Output	Biomethane	m^3	20.17
	Digestate	t	0.11
	Wastewater	t	0.89
S2			
Input	POME	t	1
	Monoethanolamine	kg	0.03
	Iron oxide	kg	0.05
	Copper oxide	kg	0.08
	Nickel	kg	0.06
	Water	kg	17.50
	Steam	kg	570.43
	Heat	MJ	1693.50
Output	Hydrogen	kg	6.08
	Tail gas	kg	620.21
	Wastewater	t	0.59

Direct methane emissions from each stage are obtained by multiplying the corresponding treatment organic load by the emission factors for facultative and other treatment and for discharge to the water environment [4]. Chemical oxygen demand concentrations [5, 9] are assumed at three treatment stages: from palm oil mill to anaerobic digesters, from anaerobic digesters to aerobic digesters, and from aerobic digesters to final discharge. For the benchmark, where methane is not recovered, biogas produced in the anaerobic stage is assumed to be directly released to the atmosphere. Therefore, total methane emissions are calculated as the sum of emissions from all three treatment stages, resulting in 14.2 kg CH_4 per FU and representing a major source of GHG emissions. For S1, methane generated in the anaerobic stage is assumed to be captured as biogas. Direct emissions in S1 include those from the aerobic stages and the residual biogas released during purification, which is assumed to be 3% of the produced biogas, resulting in 1.7 kg CH_4. For S2, direct emissions include methane released during the mill-to-anaerobic digester stage and emissions associated with tail gas from hydrogen production. It is assumed that, apart from the above, no methane is released from other processes.

Two impact assessment methods are applied based

on the ReCiPe 2016 framework. First, the ReCiPe 2016 midpoint (H) method is used with a focus on global warming potential (GWP100), expressed as kg CO_2-eq, to quantify the climate change impact associated with POME treatment and conversion. Second, the ReCiPe 2016 endpoint (H) method is applied with a focus on the Natural Resources damage category, expressed in monetary units (USD), to represent resource depletion impacts. Background inventory data are all taken from ecoinvent 3.11, using datasets representative of Malaysia, one of the major palm oil–producing countries. These two results are subsequently used as two alternative representations of environmental gain in the value-based assessment, providing complementary perspectives on climate-related and resource-related impacts.

Techno-economic assessment

Techno-economic assessment (TEA) is performed, which parameters are shown in **Table 2**. Production costs consist of capital costs and operating costs. Capital costs are annualized directly based on resource data [8, 10]. It should be noted that the techno-economic parameters differ significantly between the two pathways due to their different technology maturity levels. The objective of this study is not to normalize these differences, but to evaluate their influence on integrated environmental–economic performance under consistent functional unit and system boundary assumptions.

Table 2: Techno-economic analysis assumptions.

	Economic parameters	Unit	Value
S1	Lifetime	Year	15
	Annualized factor		0.096
	Input/output rate		0.65
	Total CAPEX	USD	155000
	Annual operating time	h/year	4380
	Annual OPEX	USD/year	25600
	Water selling price	USD/t	0.79
	Electricity selling price	USD/kWh	0.09
	Biomethane selling price	USD/m^3	0.55
S2	H_2 production	t/year	963.34
	Lifetime	Year	20
	Annualized factor		0.1
	Hydrogen output	kg	6.08
	Total CAPEX	USD	121080920
	Annual OPEX	USD/year	52717984.74
	Hydrogen selling price	USD/kg	19.68

The annualized capital cost is calculated as:

$$Annualized\ CAPEX = CAPEX_{tot} \times AF \tag{2}$$

where $CAPEX_{tot}$ is the total capital cost of the plant, and AF is the annualized capital cost factor derived from the plant lifetime.

For S1, the capital cost allocated to the FU is then obtained based on the annual operating time and the output or input rate of the system:

$$CAPEX_{FU-S1} = \frac{Annualized\ CAPEX}{R \times t_{op}} \tag{3}$$

where $CAPEX_{FU-S1}$ is the capital cost per FU for S1, R is the output rate of the system (per hour), and t_{op} is the annual operating time (h·yr^{-1}).

For S2, the capital cost allocated to the FU is based on annual hydrogen production and hydrogen production per FU:

$$CAPEX_{FU-S2} = \frac{Annualized\ CAPEX \times\ H_2\ output}{Annual\ H_2\ production} \tag{4}$$

where $CAPEX_{FU-S2}$ is the capital cost per FU for S2, $H_2\ output$ is amount of hydrogen produced per FU, and $Annual\ H_2\ production$ is annual production of hydrogen.

For S1, operating costs are calculated based on the annual operating cost and the cost of input materials per FU:

$$OPEX_{FU-S1} = \frac{OPEX_{ann}}{R \times t_{op}} + \sum(Input_i \times Price_i) \tag{5}$$

where $OPEX_{FU}$ is the operating cost per FU, $OPEX_{ann}$ is the annual operating cost provided in the resource data, R is the output rate of the system (per hour), and t_{op} is the annual operating time (h·yr^{-1}). $Input_i$ is the quantity of input i required per FU, and $Price_i$ is the corresponding unit price.

For S2, operating costs are calculated based on the annual operating cost, annual hydrogen production and hydrogen production per FU:

$$OPEX_{FU-S2} = \frac{Annualized\ OPEX \times\ H_2\ output}{Annual\ H_2\ production} \tag{6}$$

where $OPEX_{FU-S2}$ is the operating cost per FU for S2, $H_2\ output$ is amount of hydrogen produced per FU, and $Annual\ H_2\ production$ is annual production of hydrogen. And cost of input materials and energy in S2 are included in $OPEX_{ann}$.

Product revenues are calculated as the market values of POME-derived products:

$$Total\ revenue = \sum(Product_j \times Price_j) \tag{7}$$

where $Product_j$ is the output quantity of product j per FU, and $Price_j$ is the corresponding market price.

The net economic value (NEV) is evaluated by subtracting total revenues (TR) with the sum of capital and operating costs per FU.

The benchmark is not included in the cost calculation, because this study focuses exclusively on the costs associated with valorizing an existing POME treatment facility.

Value-based assessment

Unlike conventional investment analysis that considers only ROI, value-based assessment also incorporates environmental gains to represent the sustainability value of impact mitigation in a unified metric.

Firstly, the return of investment (ROI) is used to evaluate the conventional economic performance of each scenario. ROI is defined as:

$$ROI = \frac{Financial\ gain - Financial\ loss}{Total\ investment}$$

where $Financial\ gain$ represents the revenue obtained from product sales and $Financial\ loss$ includes capital and operating costs.

ROI is applied to quantify the economic attractiveness of each pathway without considering environmental externalities and serves as a reference to the integrated sustainability performance indicated by ROV.

The integrated sustainability performance is quantified using the ROV [11]:

$$ROV = \frac{Value\ gain - Value\ loss}{Total\ investment} \tag{8}$$

where $Value\ gain$ includes both financial and environmental gains from investment, while $Value\ loss$ represents the total cost, including both $CAPEX$ and $OPEX$.

Two approaches are applied to quantify environmental gains. Under the midpoint-based approach, environmental gain is estimated by multiplying the avoided GHG emissions by the projected Malaysian carbon price:

$$Enivironmental\ gain_{GHG} = \Delta GHG \times P_{carbon} \tag{9}$$

where $\Delta GHG$is the reduction from benchmark and including the avoided fossil-derived products in GWP100 and $P_{carbon}$is the carbon price (RM 35–45/t CO_2-eq) [12], we use the average 40 RM/t CO_2-eq, or 10 USD/t CO_2-eq.

Under the endpoint approach, environmental gain directly uses the Natural Resources damage category expressed in monetary units (USD). This value is derived from the ReCiPe2016 endpoint (H) method, which converts resource depletion into economic terms based on the marginal increase in future extraction costs caused by current resource use, thereby representing the damage to natural capital in value units.

The economic gain is calculated as NEV from product sales, while total investment includes $OPEX_{FU}$ and $CAPEX_{FU}$. The ROV indicator therefore integrates economic profitability and environmental externalities into a unified value metric, enabling direct comparison between POME valorization technologies.

RESULTS AND DISCUSSION

The environmental impacts of the two scenarios are evaluated using two complementary approaches. **Figure 2a** presents the GHG emission results, with the conventional treatment used as a benchmark. Both S1 and S2 achieve net GHG reductions, which are 392.57 and 158.99 kg CO_2-eq, respectively. This indicates that diverting POME from direct discharge to valorization pathways can substantially mitigate climate impacts. For S1, the dominant contribution is the avoided GHG emissions associated with fossil-derived energy substitution through biomethane utilization, totally 447.38 kg CO_2-eq. This negative contribution outweighs the positive emissions arising from electricity use, biogas purification, and material inputs, resulting in a net GHG reduction. Direct methane emissions are relatively low because methane generated in the anaerobic stage is assumed to be captured and upgraded, with only minor losses during purification. In S2, avoided GHG emissions is 513.97 kg CO_2-eq, but this benefit is partially offset by process emissions. Steam and heat consumption for reforming and gas separation represent the largest positive contributors, followed by tail gas release. As a result, although S2 still achieves net GHG reduction compared with the benchmark, its total GHG benefit is smaller than that of S1. Overall, the GHG results highlight contrasting mitigation mechanisms: S1 mainly reduces emissions through methane capture and fossil energy substitution, whereas in S2 the avoided emissions from hydrogen substitution are largely offset by substantial process emissions. Therefore, biomethane is more effective than bio-hydrogen in reducing GHG emissions under current conditions.

Figure 2b shows the endpoint results for natural resources in terms of avoided damage, incurred damage, and net damage for the two valorization scenarios. For S1, the net damage is -6.17 USD, showing that the avoided damage associated with fossil-derived substitution exceeds the resource damage generated during processing. This result indicates that biomethane production provides a small net benefit in the natural resources indicator, as the resource savings from displaced fossil energy outweigh the additional resource use required for treatment and purification. In S1, the avoided damage associated with fossil-derived energy substitution slightly exceeds the resource damage generated during processing, resulting in a small negative net damage. This indicates that biomethane production leads to a marginal net benefit in the natural resources category, as the resource savings from displaced fossil energy offset the additional resource use required for treatment and upgrading. In contrast, the net damage of S2 is 12.41 USD, exhibiting a positive net damage. Although hydrogen substitution leads to a larger avoided damage than in S1, this benefit is more than offset by the substantial resource damage

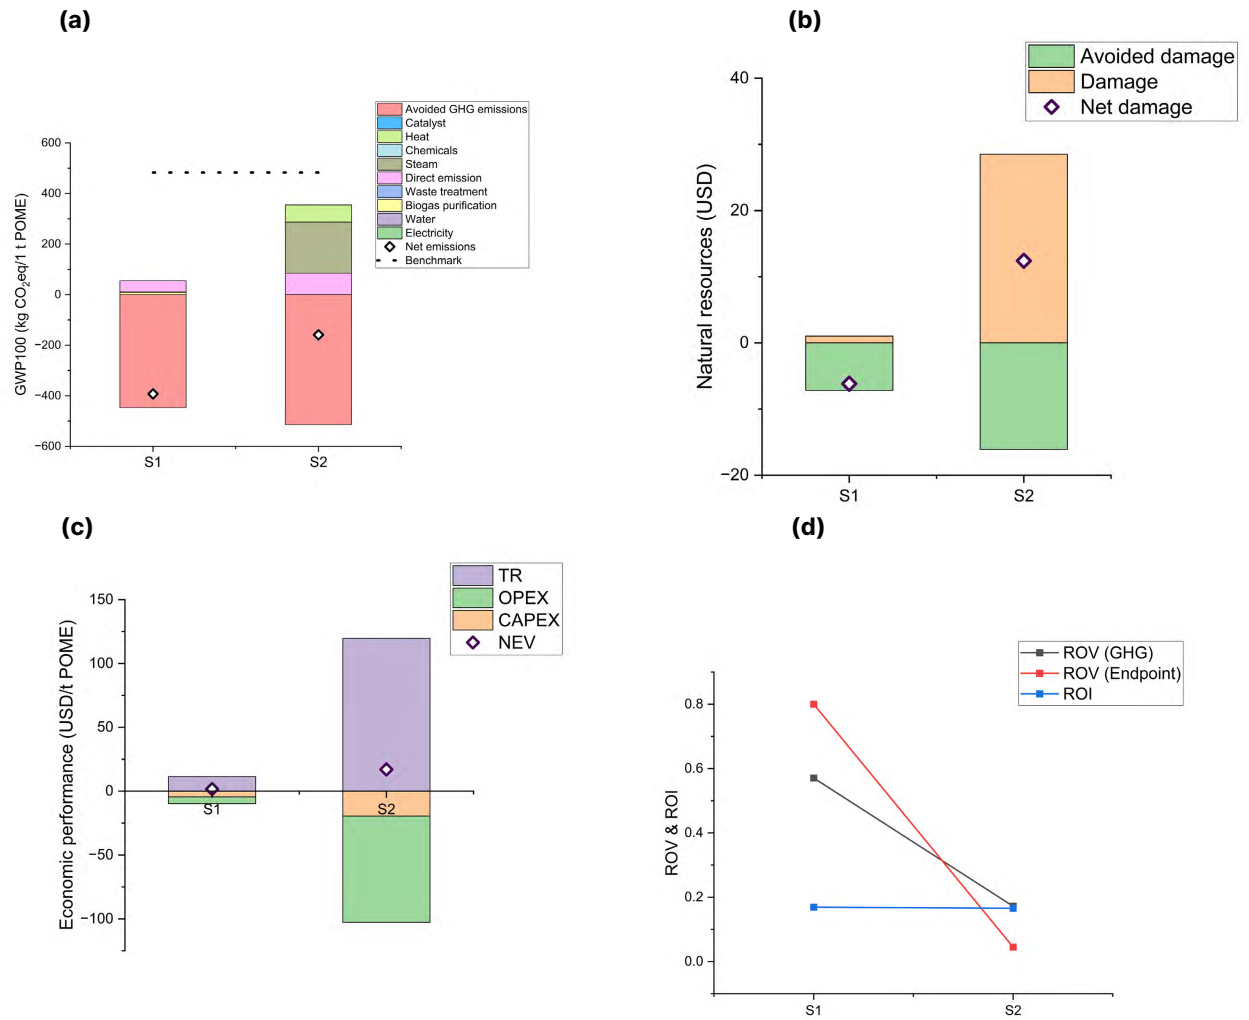

Figure 2. Integrated results of POME management scenarios (B–S2): (2a) GHG emissions (GWP100), (2b) endpoint indicator (Natural resources), (2c) economic performance (CAPEX, OPEX, TR, and NEV), and (2d) ROV based on GHG- and endpoint-derived environmental gains and financial gains and ROI only based on financial gains.

caused by conversion processes. Steam contributes significantly to the positive damage term. As a result, S2 shows an overall increase in natural resource impacts.

Overall, the environmental assessment indicates that biomethane production delivers more favorable environmental performance because the bio-hydrogen production is much more energy-intensive.

Figure 2c illustrates the economic performance of the two valorization scenarios in terms of capital expenditure (CAPEX), operating expenditure (OPEX), TR, and NEV. In S1, both CAPEX and OPEX are relatively low, reflecting the technological maturity of the biomethane pathway and its moderate energy and material requirements. The resulting TR from biomethane production slightly exceeds the total production cost, leading to a relatively low NEV. This indicates that biomethane production can approximately balance costs and revenues under the assumed conditions. In contrast, S2 exhibits substantially higher production costs, particularly in OPEX, which is mainly driven by intensive steam, heat, and chemical inputs required for hydrogen production. Although S2 achieves a much higher TR due to the higher market value of hydrogen, this revenue is largely offset by the elevated operating costs. As a result, the NEV of S2 remains positive but is only moderately higher than that of S1. Overall, the economic results suggest that S1 represents a low-risk and cost-stable option, whereas S2 offers higher revenue potential at the expense of significantly increased cost intensity.

Figure 2d compares the performance of the two valorization scenarios using ROI and ROV. In terms of conventional ROI, both scenarios exhibit similar profitability, with ROI values of 17% for S1 and 18% for S2. This indicates that biomethane and bio-hydrogen pathways provide comparable financial returns under the assumed market and cost conditions. When environmental gains

are incorporated using the midpoint impact approach (GHG-only), a clear differentiation emerges. The ROV_{GHG} of S1 reaches 0.57, substantially higher than the value of 0.17 obtained for S2. This result reflects the stronger climate mitigation effect of biomethane production through effective methane capture and fossil energy substitution, whereas the larger process-related emissions in S2 reduce its net GHG-based value gain. A larger contrast is observed when environmental gains are evaluated using the endpoint-based approach. The $ROV_{endpoint}$ of S1 reaches 0.8, while that of S2 drops to 0.04. This indicates that, once broader environmental damage categories are considered, the integrated sustainability performance of biomethane production is markedly superior to that of bio-hydrogen production. The intensive energy and material requirements of the hydrogen pathway substantially increase its endpoint damages, which are not fully compensated by its higher economic return. Overall, the value-based assessment demonstrates that scenario ranking is highly sensitive to how environmental gains are represented, with S1 consistently favored when both climate-related and damage-oriented indicators are considered.

CONCLUSIONS

This study applied a value-based assessment framework to examine how technological choice influences the integrated environmental–economic performance of POME valorization. By selecting biomethane production and bio-hydrogen production as representative mature and emerging technologies, the analysis first characterized their life-cycle environmental and economic profiles and then integrated these dimensions using the ROV.

The ROV results demonstrate that the attractiveness of POME valorization pathways depends strongly on how environmental gains are represented. When environmental performance is expressed through midpoint indicators focusing on GHG emissions, the bio-hydrogen pathway can appear competitive due to its high product value and credited climate benefits. In contrast, when endpoint indicator is applied to capture a broader environmental burden on natural resources, the biomethane pathway shows a superior integrated performance. This difference indicates that technological maturity is reflected not only in process reliability and cost structure, but also in the balance between avoided impacts and newly introduced environmental burdens. From a methodological perspective, value-based assessment offers a means to integrate LCA and TEA by expressing environmental impacts in value terms, but its interpretation is constrained by the assumptions underlying environmental valuation. GHG-based valuation depends on carbon price scenarios that presuppose the existence of policy instruments such as carbon taxes or markets, while endpoint-based resource valuation reflects theoretical damage costs rather than actual economic transactions. Several limitations should be acknowledged. In the economic assessment, labor employment and associated wage costs were not explicitly considered, which may lead to an underestimation of total production costs. In addition, catalyst recovery and reuse were not included, and tail gas was assumed to be directly emitted, which may overestimate environmental burdens for advanced configurations.

Overall, the proposed framework provides a quantitative basis for comparing mature and emerging POME valorization technologies on a common value scale. It supports strategic technology selection by revealing how technological maturity shapes integrated sustainability performance and offers a structured approach for guiding sustainable wastewater management and bio-economy development in the palm oil sector.

ACKNOWLEDGEMENTS

This work contains the results of projects supported by the JST COI-NEXT JPMJPF2003, and JSPS KAKENHI (JP22KK0169). The activities of the Presidential Endowed Chair of the “Platinum Society” at the University of Tokyo are supported by Mitsui Fudosan Corporation, Sekisui House, Ltd., East Japan Railway Company, and Toyota Tsusho Corporation.

REFERENCES

1. Murphy DJ, Goggin K, Paterson RRM. Oil palm in the 2020s and beyond: challenges and solutions. CABI Agric Biosci 2: (2021). https://doi.org/10.1186/s43170-021-00058-3
2. Mohd Yusof MAB, Chan YJ, Chong CH, Chew CL. Effects of operational processes and equipment in palm oil mills on characteristics of raw palm oil mill effluent (POME): a comparative study of four mills. Cleaner Waste Systems 5:100101 (2023). https://doi.org/10.1016/j.clwas.2023.100101
3. Soo PL, Bashir MJK, Wong LP. Recent advancements in the treatment of palm oil mill effluent (POME) using anaerobic biofilm reactors: challenges and future perspectives. Journal of Environmental Management 320:115750 (2022). https://doi.org/10.1016/j.jenvman.2022.115750
4. Intergovernmental Panel on Climate Change (IPCC), 2019 Refinement to the 2006 IPCC Guidelines for National Greenhouse Gas Inventories, Volume 5: Waste, IPCC (2019) https://www.ipcc-nggip.iges.or.jp/public/2019rf/
5. Aziz NIHA, Hanafiah MM. Life cycle analysis of biogas production from anaerobic digestion of palm oil mill effluent. Renewable Energy 145:847-

857 (2020).
https://doi.org/10.1016/j.renene.2019.06.084
6. Raketh, M., Kongjan, P., Sompong, O., Mamimin, C., Jariyaboon, R., & Promnuan, K., Life cycle assessment (LCA) and economic analysis of two-stage anaerobic process of co-digesting palm oil mill effluent (POME) with concentrated latex wastewater (CLW) for biogas production, Process Saf. Environ. Prot., 192, 450–459 (2024) https://doi.org/10.1016/j.psep.2024.10.024
7. Kan KW, Chan YJ, Tiong TJ, Lim JW. Maximizing biogas yield from palm oil mill effluent (POME) through advanced simulation and optimisation techniques on an industrial scale. Chemical Engineering Science 285:119644 (2024). https://doi.org/10.1016/j.ces.2023.119644
8. Wee ANCH, Erison AE, Edward Anyek EH, Pakpahan GR, Lim JR, Tiong ANT. Techno-economic assessment of hydrogen production via steam reforming of palm oil mill effluent. Sustainable Energy Technologies and Assessments 53:102575 (2022). https://doi.org/10.1016/j.seta.2022.102575
9. Chan YJ, Chong MF, Law CL. Biological treatment of anaerobically digested palm oil mill effluent (POME) using a lab-scale sequencing batch reactor (SBR). Journal of Environmental Management 91:1738-1746 (2010). https://doi.org/10.1016/j.jenvman.2010.03.021
10. Tan QT, Rajakal JP, Lim CH, Hassim MH, Ng DKS. Toward carbon neutrality: systematic approach to decarbonize palm oil value chain. Ind. Eng. Chem. Res. 63:1903-1925 (2024). https://doi.org/10.1021/acs.iecr.3c02405
11. Ng, D.K.S. Value-Based Sustainable Chemical Engineering Design Framework. Process Integr Optim Sustain (2025) https://doi-org.utokyo.idm.oclc.org/10.1007/s41660-025-00556-1
12. Medina, A. F. (2025, October 15). *Malaysia's 2026 budget: What it means for foreign investors*. ASEAN Briefing. https://www.aseanbriefing.com/news/malaysias-2026-budget-what-it-means-for-foreign-investors/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Chemical Additives in Plastics: Understanding the Reactions, Fate, and Releases during Pyrolysis

Ronald Borja-Roman[a], Andres Castellar-Freile[a], John D. Chea[b], Monica Rodriguez Morris[c], Gerardo J. Ruiz-Mercado[de], and Kirti M. Yenkie[a]*

[a] Rowan University, Department of Chemical Engineering, Glassboro, New Jersey, U.S.
[b] Bright Path Laboratories, Scottsdale, Arizona, U.S.
[c] U.S. Environmental Protection Agency, Oak Ridge Institute for Science and Education, hosted by the Office of Research & Development, Cincinnati, Ohio, U.S.
[d] U.S. Environmental Protection Agency, Office of Research & Development, Cincinnati, Ohio, U.S.
[e] Universidad del Atlántico, Chemical Engineering Graduate Program, Puerto Colombia, Atlántico, Colombia.
* Corresponding Author: yenkie@rowan.edu.

ABSTRACT

Plastic pyrolysis is widely promoted as a techno-economic industrial scale recycling strategy. Nevertheless, the fate and reactivity of plastic chemical additives during pyrolysis are mostly overlooked in product quality and environmental release assessments. Here, we present an integrated modeling framework to elucidate the role of additives in plastic pyrolysis and evaluate the implications of their transformation products and environmental releases. Using high-density polyethylene (HDPE) as a case study, chemical additives of concern are selected based on occurrence, concentration data, and potential risk to human health and the environment. Bond dissociation energies are predicted using a machine learning model to identify dominant radical species formed under pyrolytic conditions. These additive-derived radicals are incorporated into an automatic chemical reaction mechanism generator that constructs kinetic models composed of elementary chemical reaction steps. These kinetic models are simulated using kinetic Monte Carlo (kMC) methods to predict product distributions and yields. The results show that common additives readily form stabilized alkyl and aryl radicals at energies accessible during pyrolysis, enabling their active participation in polymer degradation pathways. These interactions influence product formation and may contribute to the generation of environmentally relevant by-products. Overall, this study provides a mechanistic and risk-informed perspective on plastic pyrolysis, emphasizing the importance of explicitly accounting for additive chemistry in the development of safer and more sustainable chemical recycling technologies.**Disclaimer**: The views expressed in this work are those of the authors and do not necessarily represent the views or policies of the EPA.

Keywords: Plastic Recycling, Machine Learning, Reaction Engineering, Stochastic Simulations, Environment

INTRODUCTION

Plastic pollution remains a critical global challenge due to rapidly increasing waste generation and the limited effectiveness of existing recycling technologies. Approximately 2.01 billion metric tons of municipal solid waste (MSW) are generated annually worldwide, a figure projected to rise to 3.4 billion metric tons by 2050 [1]. Plastics constitute a substantial fraction of this end-of-life (EoL) material stream, accounting for roughly 7–12% by weight in the United States [2]. Despite ongoing recovery efforts, only 9–10% of plastic MSW is recycled globally, while 12–19% is incinerated, and the majority is landfilled or released into the environment, leading to persistent accumulation and adverse environmental and human health impacts [3].

In this context, pyrolysis has emerged as a promising thermochemical upcycling strategy for converting plastic waste into fuels, monomers, and chemical feedstocks under oxygen-free conditions. However, plastics are chemically complex materials that contain not only polymer matrices but also a wide range of intentionally

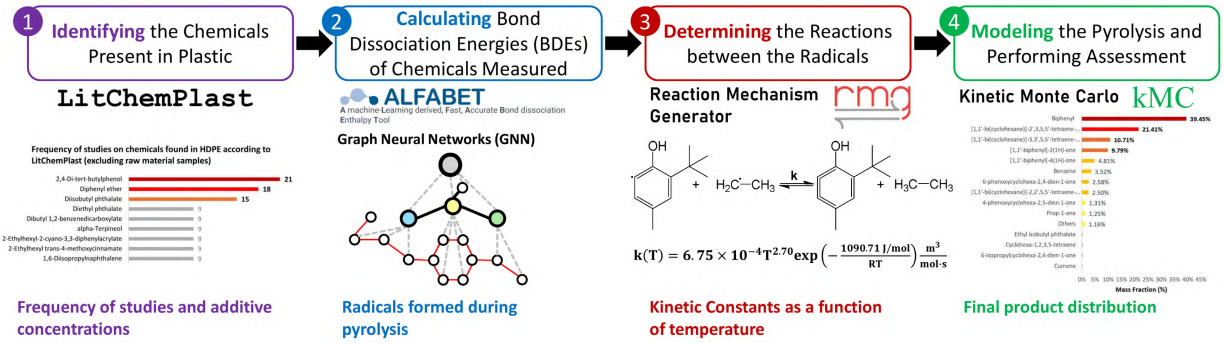

Figure 1. Proposed framework comprising the LitChemPlast database, the machine learning model ALFABET, the Reaction Mechanism Generator (RMG), and kinetic Monte Carlo (kMC) simulations to estimate the reactions, fate, and releases of chemical additives in plastic pyrolysis.

and non-intentionally added substances, including residual monomers, additives, and processing aids. Data from the LitChemPlast database show that over 3, 500 chemicals have been detected in plastics across 372 studies, with more than 16, 000 substances linked to plastic products, including over 4, 200 classified as hazardous [4]. Despite this complexity, most pyrolysis studies focus on polymer degradation, largely overlooking the fate and reactivity of additive-derived chemicals.

Mechanistic modeling approaches, including kinetic Monte Carlo (kMC) simulations, have been widely applied to describe polymer degradation and predict product distributions based on molecular-scale reaction kinetics [5]. However, polymer–additive interactions remain poorly characterized due to limited experimental data and incomplete reaction mechanisms, reducing the predictive capability of conventional models [6]. Machine learning (ML) offers a powerful means to address these limitations by enabling systematic identification of chemically plausible reaction pathways and reactive intermediates [7].

In this work, we integrate ML-based tools with kMC simulations to explicitly account for additive chemistry during plastic pyrolysis. Representative additives are selected based on occurrence and concentration data from the LitChemPlast database. Bond dissociation energies are predicted using ALFABET to identify dominant radical species formed under pyrolytic conditions [8], and their subsequent interactions are explored using the reaction mechanism generator (RMG) [9]. The resulting reaction networks are simulated using kMC to predict product distributions and yields and used to assess the environmental fate and impacts of pyrolysis-derived products. The framework is demonstrated through a case study of high-density polyethylene (HDPE) pyrolysis that incorporates three frequently reported additives (2, 4-di-tert-butylphenol, diphenyl ether, and diisobutyl phthalate).

Methodology

This work introduces an integrated framework to predict product yields and the formation of chemicals of concern during plastic pyrolysis. Machine learning (ML) tools, coupled with kinetic Monte Carlo (kMC) simulations, capture additive-derived reaction pathways often neglected in conventional analyses. The framework identifies representative plastic-associated chemicals, predicts dominant radicals using bond dissociation energy (BDE) analysis, and constructs reaction networks that are used to determine the pyrolysis degradation products. Figure 1 outlines the workflow from chemical identification and radical generation to kMC-based product prediction.

Identifying the Chemicals Present in Plastic

The LitChemPlast database is used to identify and quantify chemicals present in plastics. It compiles and harmonizes chemo-analytical measurements across multiple stages of the plastic life cycle, addressing the limited availability of publicly accessible data on plastic chemical composition. The database includes studies published between 1978 and 2021, identified through systematic searches via Web of Science and SciFinder followed by a two-stage screening process. Extracted data include bibliographic information, sample characteristics (product categories and polymer types), analytical methods, and measured chemicals, with concentrations standardized and reported in parts per million (ppm) [4].

In this study, the comprehensive state-of-the-art open database of chemical additives in plastics, LitChemPlast, is used to identify chemicals of concern with the highest reported occurrence and concentrations across different plastic types and applications. This information enables the selection of representative compounds that are most likely to participate in chemical reactions during the pyrolysis process and therefore influence reaction pathways and product formation.

Calculating Bond Dissociation Energies (BDEs) of Chemicals Measured in Plastics

To predict radical formation during pyrolysis, bonds most susceptible to thermal cleavage must be identified. Bond dissociation energies (BDEs) quantify bond strength and are used to determine bond-breaking under pyrolytic conditions. A BDE is the energy required for homolytic cleavage of a bond (A–B) into radicals ($A\bullet$ and $B\bullet$) [10]. As thermodynamic parameters, BDEs inform reactivity, kinetics, and accessible pathways in radical-driven systems. Bond dissociation is shown in Equation (1), and BDEs are calculated from the enthalpies of formation ($\Delta_f H$) of the parent molecule and the resulting radicals (Equation (2)).

$$A - B \rightarrow A\bullet + B\bullet \tag{1}$$

$$BDE = \Delta_f H(A\bullet) + \Delta_f H(B\bullet) - \Delta_f H(A - B) \tag{2}$$

Experimental determination of BDEs is challenging and restricted to a few techniques, including radical kinetics, photoionization mass spectrometry, and the acidity–electron affinity thermochemical cycle. These methods require high-precision measurements, advanced instrumentation, and strict control, limiting their applicability, particularly for complex or transient species [11]. Quantum chemical approaches, such as density functional theory (DFT), provide accurate BDE estimates but are computationally intensive and sensitive to the level of theory options chosen.

To address these limitations, we employ ALFABET, a machine learning (ML)-based model that predicts homolytic BDEs of organic molecules with near-DFT accuracy at sub-second computational cost. ALFABET is trained on over 290, 000 DFT-calculated BDEs using a graph neural network that operates directly on molecular structures derived from simplified molecular input line entry system (SMILES) strings, eliminating manual feature engineering. The model achieves mean absolute errors of 2.4 kcal mol^{-1} for in-distribution molecules and 3.4 kcal mol^{-1} for out-of-distribution molecules, compared with DFT values of 2.1 kcal mol^{-1} and 2.5 kcal mol^{-1}, respectively. This accuracy loss is compensated by a significant decrease in the CPU time per molecule, where DFT methods require 31.4 h while ALFABET 1 ms [8].

In this work, ALFABET is used to estimate the BDEs of compounds selected from the LitChemPlast database. For each compound, the lowest BDE is used to identify the most probable radical species formed during pyrolysis, which are subsequently incorporated into the reaction network to predict the content of additive-derived chemicals in pyrolysis products and environmental releases.

Determining the Reactions between the Radicals Formed During Plastic Pyrolysis

The reaction network describing interactions between additive-derived radicals is constructed using the kinetic search module of the RMG. Starting from the radical species identified by the BDE analysis, RMG systematically enumerates chemically plausible elementary reactions and assigns temperature-dependent rate coefficients to each reaction. Pressure-independent kinetics are described using a modified Arrhenius expression as shown in equation (3).

$$k(T) = A\left(\frac{T}{T_0}\right)^n \exp\left(-\frac{E_0 + \alpha \Delta H_{\mathrm{rxn}}}{RT}\right) \tag{3}$$

Where A is the pre-exponential factor, n is the temperature exponent, T_0 is a reference temperature, E_0 is the activation energy for a thermoneutral reaction, ΔH_{rxn} is the reaction enthalpy, α is the Evans-Polanyi coefficient, R is the universal gas constant, and T is the system temperature. The Evans-Polanyi corrections are not applied in this work; therefore, α and ΔH_{rxn} are set to zero, and the reported activation energy corresponds directly to E_0.

When assigning kinetic parameters, RMG uses a hierarchical strategy prioritizing data reliability. Parameters are taken from user-provided seed mechanisms, followed by curated reaction libraries, matched training reactions, exact rule-based matches, and, when necessary, estimated average rates. When multiple sources exist within a category, parameters are ranked by level of theory or experimental evidence, with experimental data and high-level quantum calculations prioritized over lower-level estimates [9].

Using this structured approach, RMG generates a consistent and chemically informed reaction network describing radical–radical interactions during pyrolysis. The resulting elementary reactions and associated kinetic parameters are subsequently used as inputs to the kinetic Monte Carlo (kMC) simulations to predict product formation, concentrations, and yields.

Modeling Pyrolysis using Stochastic Simulations

The kMC method is a stochastic simulation framework based on the stochastic simulation algorithm (SSA) originally developed by Gillespie. It enables the time-resolved simulation of chemically reacting systems by explicitly considering the probabilistic nature of individual reaction events [5]. As shown in Figure 2, implementing kMC requires kinetic and system-specific inputs, including reaction-specific pre-exponential factors and activation energies, as well as operating conditions such as temperature, system volume, and the initial populations of all chemical species. The kinetic inputs are obtained from the RMG results, and the process conditions are established based on previous plastic pyrolysis studies.

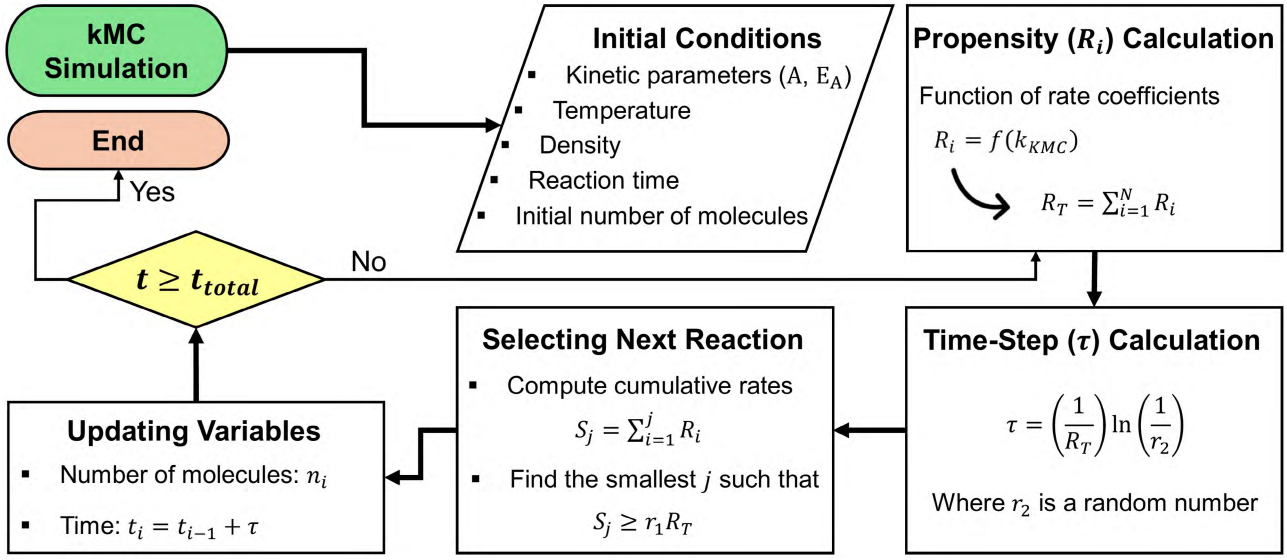

Figure 2. Flowchart of the kinetic Monte Carlo (kMC) framework used for reaction network simulations.

From these inputs, stochastic rate constants are derived from the corresponding macroscopic rate coefficients, accounting for reaction order and molecularity. Specifically, distinct expressions are used for first-order unimolecular reactions, second-order bimolecular reactions between identical species, and second-order bimolecular reactions between different species, as shown from equations (4) to (6), respectively, where $k_a^{\exp}, k_{aa}^{\exp}$, and $k_{ab}^{\exp}$ denote the macroscopic rate coefficients, N_A is Avogadro's number, and V is the control volume [12]. These stochastic rate constants are used to calculate the reaction propensities for each elementary reaction channel. Equation (7) shows the total propensity of the system (R_T), which is obtained by summing up the propensities of all reactions (R_i).

$$k^{MC} = k_a^{\exp} \tag{4}$$

$$k^{MC} = \frac{2k_{aa}^{\exp}}{V \cdot N_A} \tag{5}$$

$$k^{MC} = \frac{k_{ab}^{\exp}}{V \cdot N_A} \tag{6}$$

$$R_T = \sum_{i=1}^{N} R_i \tag{7}$$

The time to the next reaction event is sampled by generating a uniform random number $r_2 \in (0,1)$, which determines the stochastic time increment. A second random number $r_1 \in (0,1)$is then used to select the reaction event by identifying the first reaction whose cumulative propensity exceeds $r_1 \cdot R_T$. Following reaction selection, the populations of reactant and product species are updated according to the reaction stoichiometry. This procedure is repeated, allowing the reaction network to evolve stochastically over time while rigorously preserving the correct statistical distribution of reaction events.

Following the kMC simulations, time-resolved product distributions and final yields are obtained, and high-concentration products are prioritized for downstream analysis due to their relevance for environmental exposure and risk evaluation [13]. This study introduces a fully integrated, mechanistic framework that explicitly links additive-specific reaction chemistry to pyrolysis product formation and environmental releases, which is an aspect that has not been systematically addressed in prior pyrolysis modeling efforts. Whereas most existing models are constrained to polymer-only degradation due to the severe lack of kinetic data for additive-derived reactions, the present approach overcomes this limitation by combining ML-based reactivity prediction, automated reaction mechanism generation, and stochastic kinetic modeling. By doing so, this work enables, for the first time, predictive exploration of additive-driven degradation pathways, improves the reliability of pyrolysis outcome predictions, and supports early identification of environmentally relevant and potentially hazardous degradation products.

RESULTS AND DISCUSSION

This section presents the validation of the proposed methodology through a case study on plastic pyrolysis. The most frequently reported chemicals associated with the selected plastic are identified, and their BDEs are used to determine the dominant radical species formed during thermal degradation. The interactions between additive-derived radicals are then analyzed to determine key reaction pathways and their influence on product

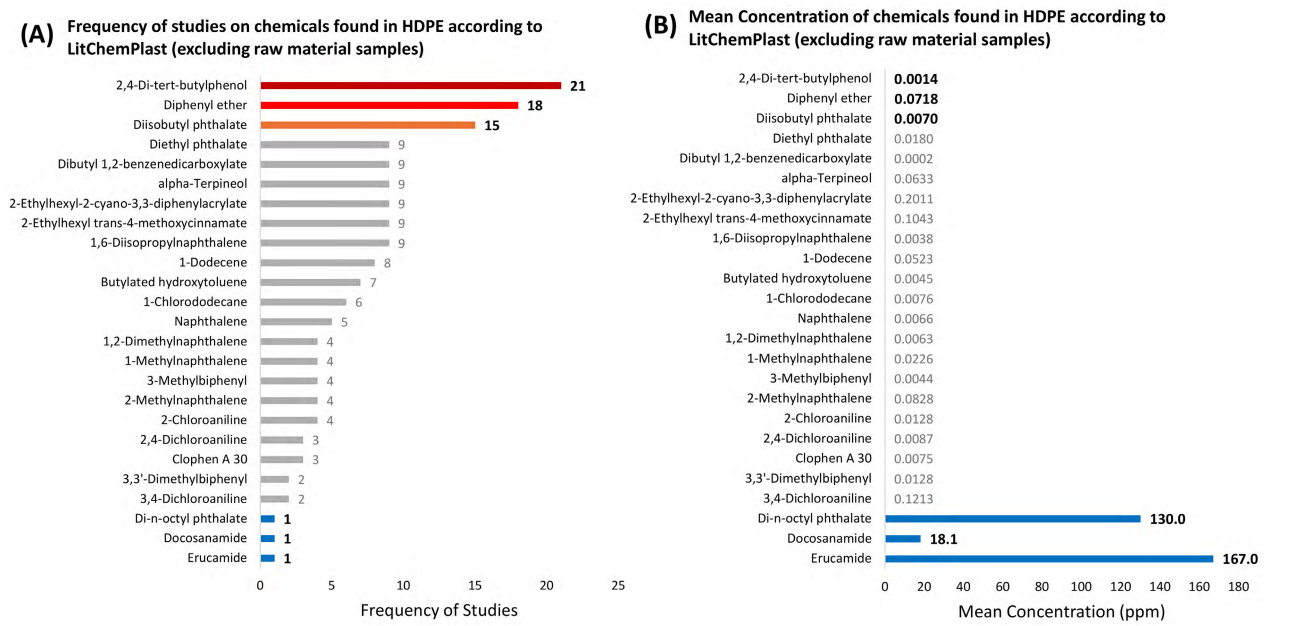

Figure 3. A) Frequency of studies and **B)** Mean concentration in parts per million (ppm) of additives found in HPDE according to the LitChemPlast database, excluding raw material samples.

formation. Based on these reactions, the resulting product distribution is predicted, and the environmental implications and fate of the released chemicals are subsequently assessed.

Case Study: Pyrolysis HDPE with Chemical Additives

The case study used to validate the proposed methodology is based on the pyrolysis of HDPE, considering the chemical additives identified in different HDPE samples as reported in the LitChemPlast database. Figure 3 shows the frequency of studies reporting chemicals detected in HDPE and their corresponding mean concentrations, as compiled from the LitChemPlast database [4].

As shown in Figure 3A, the chemicals most frequently reported in HDPE samples are 2, 4-di-tert-butylphenol, diphenyl ether, and diisobutyl phthalate, identified in 21, 18, and 15 samples, respectively. Additive selection was primarily based on occurrence frequency across the LitChemPlast dataset to ensure that the case study compounds are broadly representative of HDPE samples.

In contrast, di-*n*-octyl phthalate, docosanamide, and erucamide were each detected in only a single sample but exhibited the highest concentrations (Figure 3B). While such high-concentration additives represent interesting case studies, their limited occurrence restricts statistical representativeness. Nevertheless, when present in specific EoL streams, their high concentration may significantly influence the emission risk profile: erucamide undergoes autocatalytic degradation below 90°C generating nitrile species and alkyl radicals [14], while di-n-octyl phthalate decomposes via ester bond cleavage to yield oxygenated aromatics [15]. These chemicals were excluded from this case study; however, the proposed framework is generalizable and can readily be applied to such scenarios.

Table 1: Chemical structure, description, SMILES string, and primary hazards of the most measured chemical compounds in HDPE samples.

Information	**2, 4-di-tert-butylphenol**	**Diphenyl ether**	**Diisobutyl Phthalate**
Chemical Structure			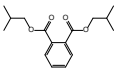
Description	Antioxidant [16]	Degradation product [17]	Plasticizer[18]
SMILES String	CC(C)(C)c1ccc(O)c(C(C)(C)C)c1	c1ccc(Oc2ccccc2)cc1	CC(C)COC(=O)c1ccccc1C(=O)OCC(C)C
Primary Hazards	Corrosive, Irritant, Health Hazard, Environmental [16]	Irritant, Environmental hazard [17]	Health Hazard [18]

Consequently, this work focuses on the three most frequently reported chemicals in HDPE (excluding raw material samples), as identified in the LitChemPlast database: 2, 4-di-tert-butylphenol, diphenyl ether, and

diisobutyl phthalate. Notably, all three compounds are associated with hazard classifications, including potential carcinogenicity and aquatic toxicity [16] [17] [18]. Table 1 summarizes their key characteristics, including their functional roles in HDPE and SMILES representations, which are subsequently used to compute the corresponding BDEs.

Computing BDEs of Chemicals Measured

The calculated BDEs for the three compounds selected in the case study are summarized in Table 2. For each compound, the two lowest BDE values are reported to provide an overview of the energetically most accessible bond-cleavage pathways.

Table 2: The two lowest BDE values obtained with ALFABET for 2, 4-di-tert-butylphenol, diphenyl ether, and diisobutyl phthalate.

Bond	2, 4-di-tert-butylphenol	Diphenyl ether	Diisobutyl Phthalate
#1	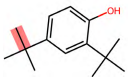 Bond Type: C-C *BDE(ML):* 76.1kcal/mol	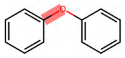 Bond Type: C-O *BDE(ML):* 84.2kcal/mol	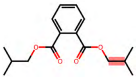 Bond Type: C-C *BDE(ML):* 87.1kcal/mol
#2	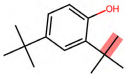 Bond Type: C-C *BDE(ML):* 76.7kcal/mol	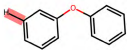 Bond Type: C-H *BDE(ML):* 110.9kcal/mol	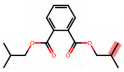 Bond Type: C-C *BDE(ML):* 89.0kcal/mol

Across all compounds, the lowest BDEs range from approximately 76.1 to 87.1 kcal mol^{-1}, indicating that radical formation from these additives is thermodynamically accessible under typical pyrolysis temperatures (≥ 400-600 °C), where thermal energies are sufficient to activate bond cleavage within this range.

For 2, 4-di-tert-butylphenol, the lowest BDEs (76.1 and 76.7 kcal mol^{-1}) correspond to C-C bond cleavage within the tert-butyl substituents. These values are approximately 10 kcal mol^{-1} lower than the phenolic O-H bond dissociation energy (85.9 kcal mol^{-1}) and nearly 20 kcal mol^{-1} lower than aromatic C-C bond cleavage (96.5 kcal mol^{-1}). This significant energetic preference indicates that pyrolysis is initiated predominantly through the formation of tert-butyl-derived alkyl radicals rather than phenoxy or aryl radicals.

In the case of diphenyl ether, the weakest bond is the aryl-O bond, with a BDE of 84.2 kcal mol^{-1}. This value is significantly lower than the aromatic C-H BDEs (110.9-112.0 kcal mol^{-1}), indicating that homolytic cleavage of the ether bond is the primary pathway for radical formation. The resulting phenyl and phenoxy radicals are stabilized by resonance and are known to participate in secondary reactions that lead to aromatic growth. Therefore, for the formation of alkyl radicals observed in 2, 4-di-tert-butylphenol, diphenyl ether predominantly contributes to aromatic radical species, which may affect the formation of polycyclic aromatic hydrocarbons and char during pyrolysis.

For diisobutyl phthalate, the lowest BDEs (87.1 and 89.0 kcal mol^{-1}) are associated with C-C bond cleavage in the isobutyl side chains. These values are 8-10 kcal mol^{-1} lower than C-H bond dissociation energies in the ester backbone (95.9-96.7 kcal mol^{-1}), indicating that radical formation is dominated by the release of isobutyl radicals rather than direct fragmentation of the aromatic or ester part.

Overall, the BDE analysis demonstrates that additive-derived radicals are thermodynamically favored to form under pyrolytic conditions, and their nature strongly depends on the additive structure. These results provide a quantitative basis for including additive-derived radicals in the reaction network and support their potential role in modifying polymer degradation pathways, product distributions, and the formation of environmentally relevant by-products.

Determining the Reactions between the Radicals Formed During the Plastic Pyrolysis

Based on the BDE analysis obtained using ALFABET, the radical species considered in the pyrolysis simulations are summarized in Table 3. For each additive, the lowest BDE value is selected as the most probable bond-cleavage pathway, resulting in two corresponding radical species that are incorporated into the reaction network. These additive-derived radicals constitute the initial reactive pool used for reaction network generation with the RMG.

To determine the interactions among radical species, binary interactions between all species are assumed. This assumption is widely used in kinetic modeling of pyrolysis and combustion systems, where bimolecular reactions dominate radical propagation, abstraction, and termination pathways. In contrast, higher-order reactions are statistically less probable under typical thermal decomposition conditions due to the low likelihood of simultaneous multi-body collisions. Consistent with this framework, the RMG constructs reaction networks primarily from unimolecular and bimolecular reaction steps based on established kinetic theory [19]. Considering this, Figure 4 summarizes the reaction space generated

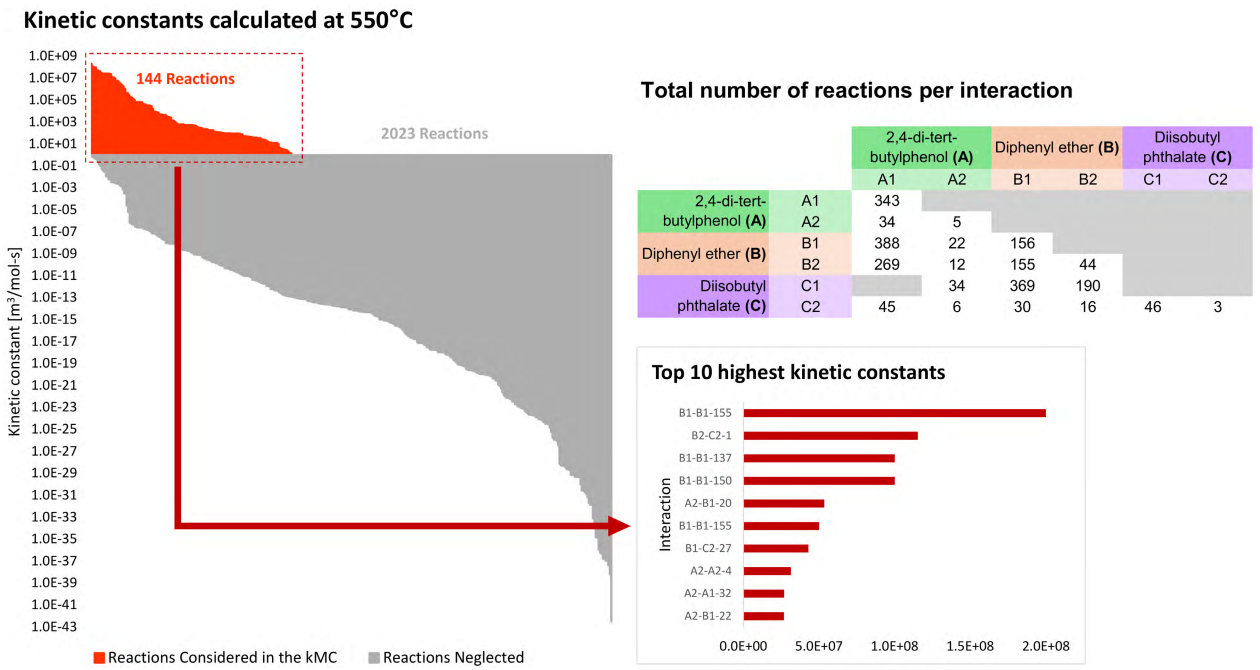

Figure 4. Total number of reactions generated for each radical-radical interaction, along with the corresponding kinetic parameters obtained from the Reaction Mechanism Generator (RMG). Of the 2,167 reactions identified, 144 are selected for inclusion in the kinetic Monte Carlo (kMC) simulations.

by RMG for interactions between additive-derived radicals and the subset retained for kMC simulations. 2, 167 elementary reactions are identified across all radical–radical interactions; however, only 144 reactions are selected based on their kinetic relevance at the pyrolysis temperature (550 °C). These retained reactions exhibit rate constant values several orders of magnitude higher than the remainder of the network, indicating that a relatively small subset of reactions dominates system kinetics.

The distribution of kinetic constants spans more than 40 orders of magnitude, with the fastest reactions primarily involving diphenyl ether-derived radicals (B-type), which account for the highest number of reaction pathways and the largest rate constants. The highest-rate reactions are predominantly radical–radical recombination or hydrogen abstraction events, consistent with expected high-temperature pyrolysis chemistry.

The reduction of the reaction network relies on the assumption that reactions with negligible rate constants do not significantly contribute to product formation within the simulated time scale. This filtering strategy preserves the chemically and kinetically dominant pathways while maintaining computational tractability, enabling efficient and physically meaningful kMC simulations of product formation.

Table 3: Radical species considered for generating the reaction network using the RMG.

Original Compound	Radical ID	Chemical Structure	SMILES
2, 4-di-tert-butylphenol	A1	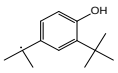	[C](C)(C)C1=CC(=C(C=C1)O)C(C)(C)C
2, 4-di-tert-butylphenol	A2	$\cdot CH_3$	[CH3]
Diphenyl ether	B1		C1=CC=C(C=C1)[O]
Diphenyl ether	B2		[c]1ccccc1
Diisobutyl phthalate	C1	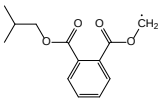	CC(C)COC(=O)C1=CC=CC=C1C(=O)O[CH2]C[CH]C
Diisobutyl phthalate	C2		

Modeling the Reaction Network using Stochastic Simulations

Figure 5 presents the product distribution predicted by the kMC simulations, showing the final mass fractions

of the compounds.

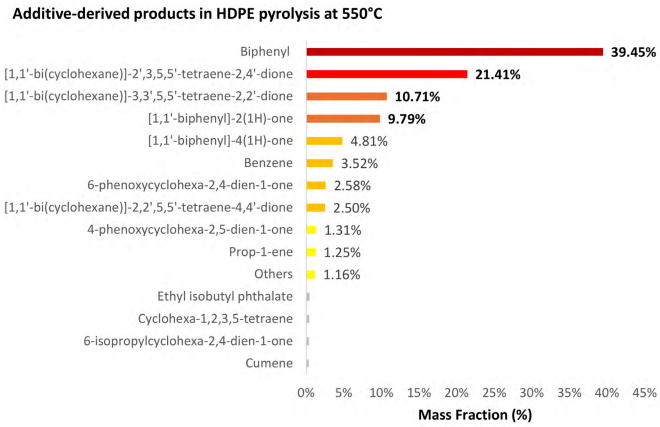

Figure 5. Product distribution predicted by the kMC simulation, considering just the chemicals derived from additives.

The kMC results show that product formation is dominated by a small number of aromatic and partially hydrogenated cyclic compounds, indicating that kinetically favored pathways govern system behavior. Biphenyl accounts for nearly 40 wt% of the additive-derived products, followed by phenolated ketones and cyclohexadienedione derivatives. These findings are consistent with the high stability and reactivity of aromatic radicals identified in the BDE and RMG analyses, especially those originating from diphenyl ether, which exhibits both high reaction connectivity and elevated rate constants.

The prominence of aromatic products suggests that additives can redirect polymer degradation toward aromatic-rich pathways, rather than acting as inert spectators during pyrolysis. This behavior contrasts with conventional polymer-only pyrolysis models, which typically predict product distributions dominated by aliphatic hydrocarbons. The formation of stabilized aromatic species early in the reaction network also explains their persistence at high concentrations in the final product mixture.

From an environmental perspective, several of the dominant products raise concerns. Biphenyl is persistent, hydrophobic, and toxic to aquatic organisms, with documented bioaccumulation potential [20]. Benzene, although produced at lower yields, is a well-established human carcinogen and a regulated air pollutant [21]. Phenolated ketones and cyclohexadienedione derivatives may serve as precursors to polycyclic aromatic hydrocarbons under secondary thermal or environmental conditions, contributing to mutagenic and carcinogenic risks. These results highlight that additive chemistry can significantly influence not only product yields but also the environmental risk profile of pyrolysis outputs.

Overall, these results show that plastic additives actively influence pyrolysis chemistry by serving as radical precursors that redirect degradation pathways and product formation. By integrating additive occurrence data with ML-based bond dissociation analysis, automated reaction mechanism generation, and kinetic Monte Carlo simulations, this study establishes a mechanistic link between additive chemistry, product distributions, and potential environmental impacts. Neglecting additive-derived reactions leads to systematic underprediction of aromatic and environmentally relevant by-products, limiting the reliability of polymer-only pyrolysis models. The proposed framework enables predictive assessment of pyrolysis outcomes and early identification of potentially hazardous degradation products, addressing key challenges associated with scarce kinetic data for additive reactions and supporting safer, more sustainable plastic recycling technologies.

CONCLUSIONS

The framework established in this study offers a risk-informed perspective on plastic pyrolysis by shifting the analytical focus from isolated polymer matrices to consider the complex chemical interactions of additives within EoL streams. By leveraging the ALFABET model to predict bond dissociation energies with near-DFT accuracy at sub-second speeds, we have enabled a systematic, scalable assessment of additives present in commercial plastics. As thermal recycling strategies like pyrolysis are promoted globally, future kinetic and environmental assessments must move beyond polymer-only models to incorporate the additive-derived reaction pathways identified here. Ultimately, the generalization of this framework into a computational tool will be crucial for characterizing other polymers and complex mixtures, ensuring the production of safe, high-value chemical feedstock from plastic EoL material.

ACKNOWLEDGEMENTS

The authors acknowledge the Department of Chemical Engineering at Rowan University and the U.S. Environmental Protection Agency.

REFERENCES

1. Babaremu KO, Okoya SA, Hughes E, Tijani B, Teidi D, Akpan A, Igwe J, Karera S, Oyinlola M, Akinlabi ET. Sustainable plastic waste management in a circular economy. Heliyon 8:e09984 (2022). https://doi.org/10.1016/j.heliyon.2022.e09984
2. United States Environmental Protection Agency (EPA), “Guide to the Facts and Figures Report about Materials, Waste and Recycling, ” (2025).
3. . Global plastics outlook. OECD Publishing (2022). https://doi.org/10.1787/de747aef-en

4. Wiesinger H, Shalin A, Huang X, Siegrist A, Plinke N, Hellweg S, Wang Z. <tt>litchemplast</tt>: an open database of chemicals measured in plastics. Environ. Sci. Technol. Lett. 11:1147-1160 (2024). https://doi.org/10.1021/acs.estlett.4c00355
5. Gillespie DT. Stochastic simulation of chemical kinetics. Annu. Rev. Phys. Chem. 58:35-55 (2007). https://doi.org/10.1146/annurev.physchem.58.032806.104637
6. Radhakrishnan H, Mohammed AABA, Coffman I, Bai X. Influence of functional additives, fillers, and pigments on thermal and catalytic pyrolysis of polyethylene for waste plastic upcycling. Green Chem. 27:5861-5882 (2025). https://doi.org/10.1039/d5gc00688k
7. Gao W, Luo S, Coley CW. Generative AI for navigating synthesizable chemical space. Proc. Natl. Acad. Sci. U.S.A. 122: (2025). https://doi.org/10.1073/pnas.2415665122
8. P. C. St. John, Y. Guan, Y. Kim, S. Kim, and R. S. Paton, "Prediction of organic homolytic bond dissociation enthalpies at near chemical accuracy with sub-second computational cost, " *(*2025)
9. M. Liu *et al.*, "Reaction Mechanism Generator v3.0: Advances in Automatic Mechanism Generation, " *J. Chem. Inf. Model.*, vol. 61, no. 6, (2021)
10. Zhao L, Zhi M, Frenking G. The strength of a chemical bond. Int J of Quantum Chemistry 122: (2021). https://doi.org/10.1002/qua.26773
11. Blanksby SJ, Ellison GB. Bond dissociation energies of organic molecules. Acc. Chem. Res. 36:255-263 (2003). https://doi.org/10.1021/ar020230d
12. Pires da Mata Costa L, Brandão ALT, Pinto JC. Modeling of polystyrene degradation using kinetic monte carlo. Journal of Analytical and Applied Pyrolysis 167:105683 (2022). https://doi.org/10.1016/j.jaap.2022.105683
13. O. US EPA, "TSCA Section 5(b)(4) Concern List." Accessed: Mar. 24, (2026).
14. Peloso CW, O'Connor MJ, Bigger SW, Scheirs J. Characterising the degradation of the polymer slip additive erucamide in the presence of inorganic antiblock agents. Polymer Degradation and Stability 62:285-290 (1998). https://doi.org/10.1016/s0141-3910(98)00009-3
15. Saido K, Taguchi H, Yada S, Ishihara Y, Kuroki T, Ryu IJ, Chung SY. Thermal decomposition products of phthalates with poly(vinyl chloride) and their mutagenicity. Macromol. Res. 11:178-182 (2003). https://doi.org/10.1007/bf03218349
16. National Center for Biotechnology Information, "PubChem Compound Summary for CID 7311, 2, 4-Di-tert-butylphenol, " National Center for Biotechnology Information, United States, 2026. Accessed: (2026).
17. National Center for Biotechnology Information, "PubChem Compound Summary for CID 7583, Diphenyl Ether, " National Center for Biotechnology Information, United States, (2026).
18. National Center for Biotechnology Information, "PubChem Compound Summary for CID 6782, Diisobutyl phthalate, " National Center for Biotechnology Information, United States, (2026).
19. Gao CW, Allen JW, Green WH, West RH. Reaction mechanism generator: automatic construction of chemical kinetic mechanisms. Computer Physics Communications 203:212-225 (2016). https://doi.org/10.1016/j.cpc.2016.02.013
20. National Center for Biotechnology Information, "PubChem Compound Summary for CID 7095, Biphenyl, " National Center for Biotechnology Information, United States, (2026)
21. National Center for Biotechnology Information, "PubChem Compound Summary for CID 241, Benzene, " National Center for Biotechnology Information, United States. (2026)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Banana peel as a Source of Polyphenols: Perspectives of Sustainability

Kajal Chaudhary[a], and Soumyajit SEN GUPTA[a]*

[a] Department of Chemical Engineering, Indian Institute of Technology (Indian School of Mines) Dhanbad, India
* Corresponding Author: soumyajit@iitism.ac.in.

ABSTRACT

The case of banana peel as a rich source of pharmaceutically important polyphenols is well-known. Our research focuses on techno-economic and environmental impact assessment of sourcing polyphenols from banana peel. Firstly, phenolic acids and flavonoids, the two major classes of polyphenols, are extracted from dried banana peel powder. The extraction process is carried out with different solvents with varying properties and different extraction techniques with differences in their working mechanism. These extraction protocols were compared on the basis of techno-economics. Distilled water was the most favored solvent in that regard. On the other hand, from the perspective of carbon dioxide emission potential, ethanol stood out as the best. Recovery of extraction solvent was considered and its impact, on techno-economics and emission potential of the protocols was analyzed. The exact nature of the impact was a strong function of the specific protocol and the target polyphenols.

Keywords: Extraction, Food & Agricultural Processes, Batch Process, Polyphenols, Techno-economics, Emission potential

INTRODUCTION

Valorization of agro-food waste is key to waste management in a sustainable manner. While such waste products, so commonly available world-wide, has great value as fuel and fertilizer, they are also useful source of important chemicals, many of which have critical pharmaceutical applications owing to their anti-oxidant properties [1, 2]. Among pharmaceutically active intra-cellular metabolites of plants, polyphenols are most abundant of which phenolic acids and flavonoids are the major subclasses [3]. Accordingly, in this work, the aim is to evaluate various extraction strategies for obtaining phenolic acids and flavonoids from banana peel.

While there exists rich literature on evaluating performance of different extraction techniques, solvents as well as other factors (such as time, temperature, parameters of extraction agents such as microwave power, ultrasound power etc.) for polyphenols extraction from various natural sources [4, 5], studies on analyzing their sustainability perspectives are very rare. Thus, in this work, we aim to bridge this gap by performing analysis based on techno-economic and environmental impact for production of phenolics and flavonoids from banana peel.

EXTRACTION PROTOCOLS

Banana peels, freshly obtained from the whole fruit was cut into pieces and dried in microwave oven to remove moisture (to around 85 %) which otherwise might interfere with solvent extraction efficiency. Our study involved optimizing critical extraction decisions in terms of techno-economics and environmental impact. With this motivation, the microwave-dried banana peel powder (BPP) was subjected to solvent extraction using five different solvents – distilled water (DW), methanol (Me), ethanol (Et), acetone (Ac) and acetic acid (AA). In terms of extraction solvent, the chosen solvents varied in terms of the key properties relevant for leaching – such as polarity (which impacts the type of polyphenols a solvent would extract), surface tension, viscosity, vapor pressure (these properties of solvent determine how good a solvent would perform in presence of ultrasound), dielectric properties (which impact suitability of solvents in microwave).

The extraction process was carried out with five different methods. Only these two factors which would fundamentally alter the extraction performance were varied.

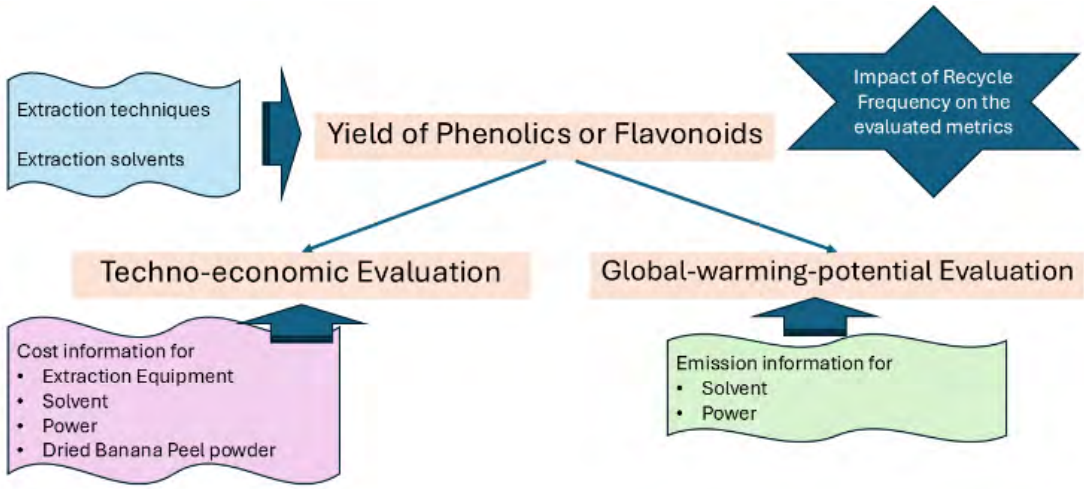

Figure 1. Scope of the Present Work and its Correlation with the Yield of Polyphenols and the Relevant Input Parameters.

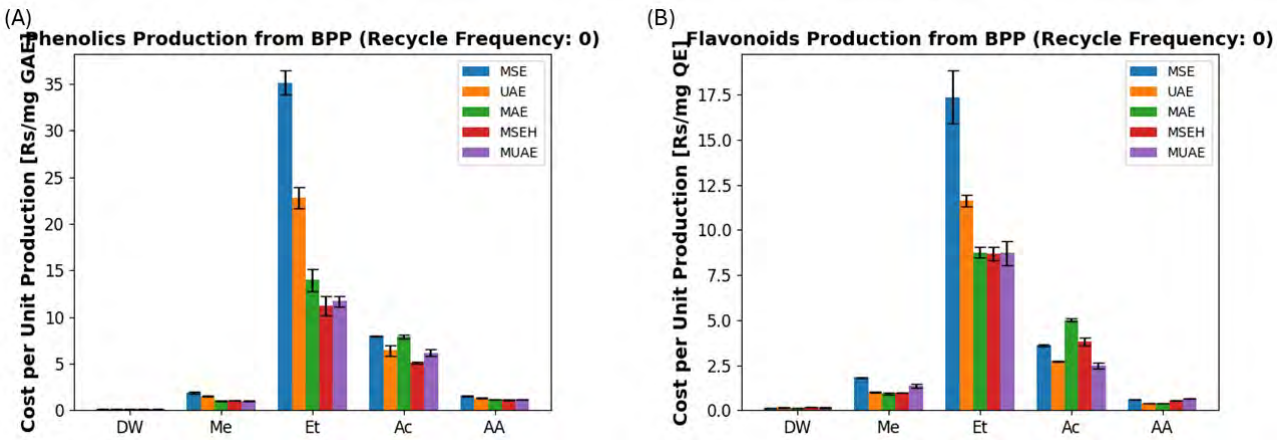

Figure 2. Variation in Cost per Unit Production of either type of Polyphenols for all the different Extraction Protocols: (A): Phenolics and (B): Flavonoids

Other potential factors like temperature, solvent amount, US power etc. were not independently varied. Among the five extraction techniques, for Magnetic Stirred-based Extraction (MSE), a simple leaching was conducted at room temperature (around 25^0 C). Ultrasound and microwave have been reported to enhance the yield by helping rupture the cell walls of BPP [6, 7], thereby allowing the solvents an easy access to dissolve them. Thus, two of our techniques (UAE and MAE) employed these to incorporate these impacts. While temperature of extraction was maintained at 25^0 C for UAE, microwave also increased the temperature of medium. The extraction temperature was maintained at 50^0 C for MAE. Thus, we also had an extraction technique (Magnetic Stirred-based Extraction with Heating, i.e., MSEH) that would ensure 50^{0}C operation with conventional heating – for a fair comparison between two modes of heating. Finally, for the case of MUAE operated at 50^0 C, both the ultrasound and microwave effects acted in parallel. Extraction time was 30 times for all the scenarios and solvent to solid ratio was 10 mL/g. Total Phenolics Content (TPC) and Total Flavonoids Content (TFC) were computed with UV-Visible Spectrophotometer. While the lab-scale experiments were conducted using 50 mL solvent, for computations in this work, a scaling of 10 times was considered since 500 mL was maximum capacity of equipment with which we performed MAE and MUAE. Accordingly, for the scaled scenario, electricity demand was considered to be also 10 times (as against lab scale one) to ensure identical profile for stirring, US, MW, heating to keep efficiency of extraction same.

METRIC FOR COMPARING VARIOUS EXTRACTION DECISION-ALTERNATIVES

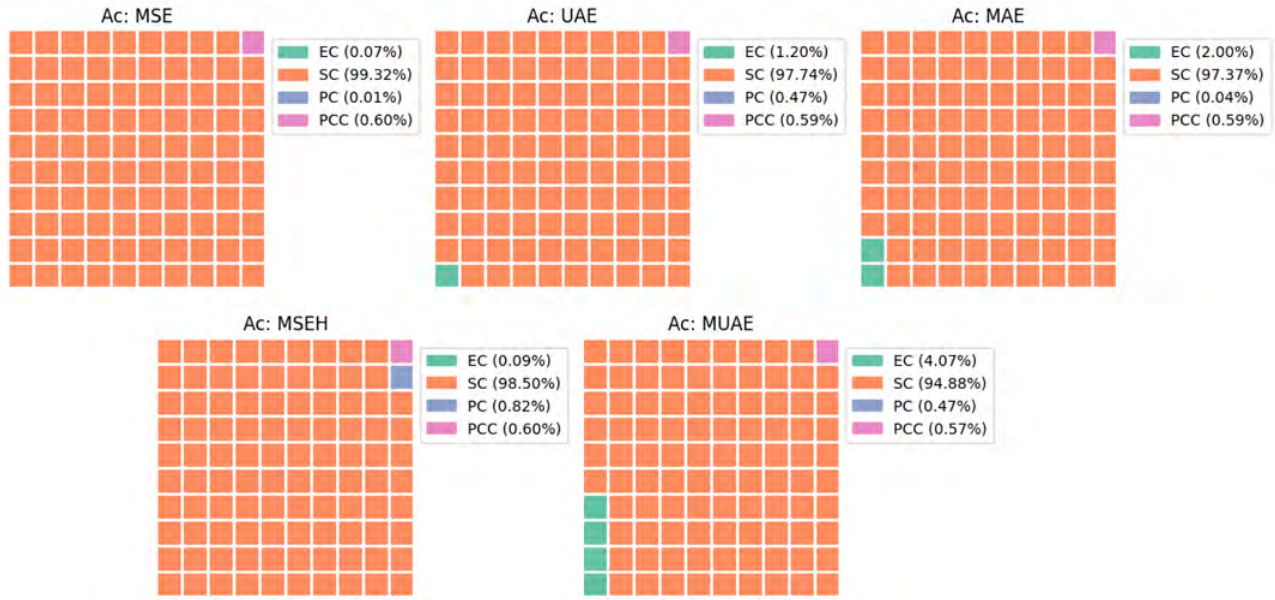

Figure 3. Relative Contribution of various cost components for varioud extraction techniques with acetne as solvent (for the use of fresh solvent, i.e., recycle frequency = 0)

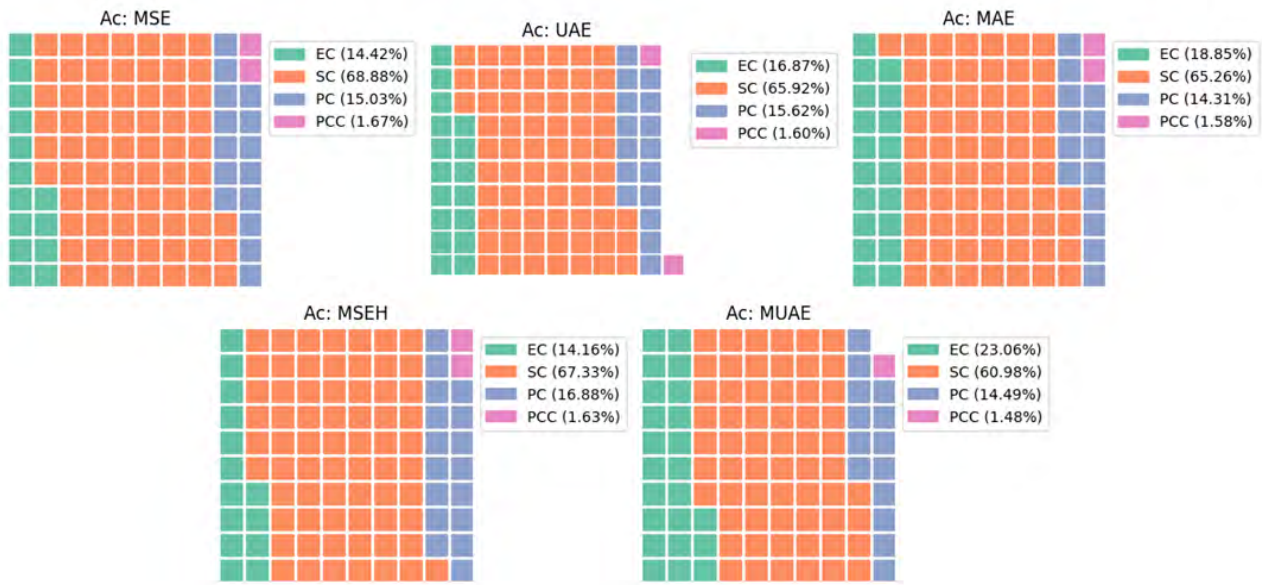

Figure 4. Relative Contribution of various cost components for varioud extraction techniques with acetne as solvent (for recycle frequency = 5)

In this section, we present and discuss about the basis on which techno-economic analysis was performed. We define Techno-Economic Metric (TEM) as a function of total cost incurred over a year for unit production of phenolics or flavonoids. Total cost was assumed to have four components – equipment cost (EC) required to procure extraction equipment as per the specific technique, solvent cost (SC) required to purchase necessary amount of a specific solvent, power cost (PC) necessary to meet the demand of electricity for stirring, heating, introducing ultrasound and microwave, as the case may be, and BPP cost (BPPC) incurred to procure the necessary amount of raw material. Among the four components, EC, being a capital cost, was annualized as per capital recovery factor. 300 days of operation was assumed over one year and 16 batches, over one day with each batch,

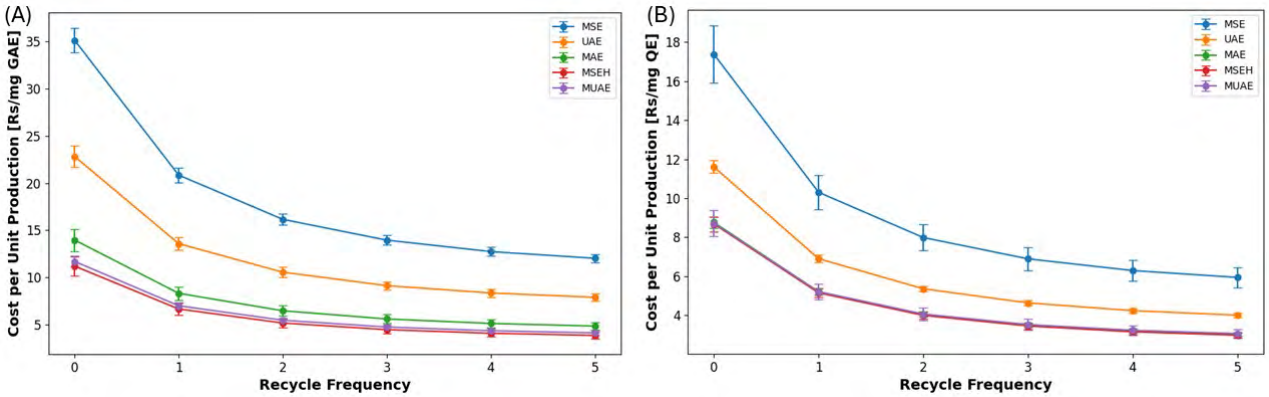

Figure 5. Variation in Cost per Unit Production of Polyphenols with Increase in Recycle Frequency: (A): Phenolics, (B): Flavonoids, considering ethanol as an example solvent

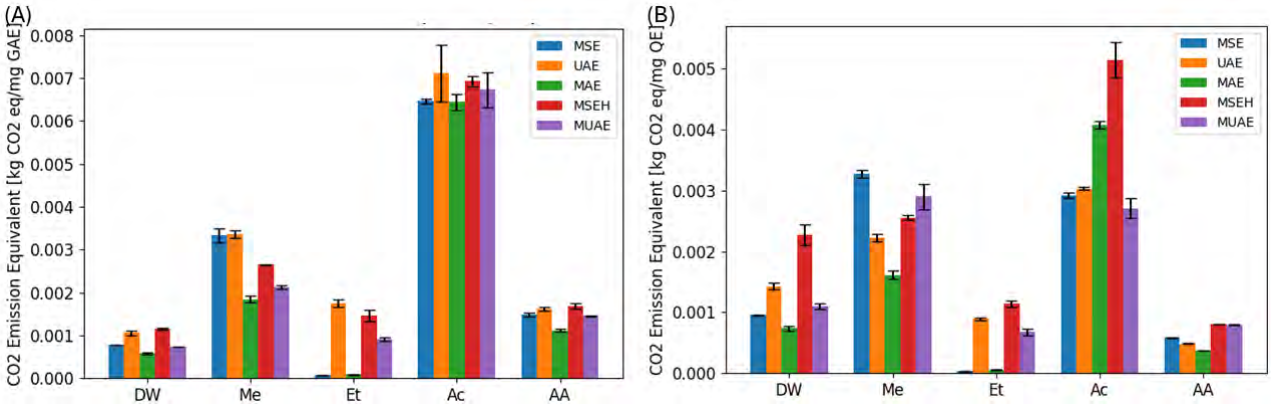

Figure 6. Variation in Emission per Unit Production of either type of Polyphenols for all the different Extraction Protocols: (A): Phenolics and (B): Flavonoids, using fresh solvent (n: 0).

employing 500 mL of a specific extraction solvent at a relative ratio of 10 ml/g solid.

Instances of solvent recovery was also considered. For such cases, the extra terms for cost for recovery scenario involved the purchase of a rotary evaporator, vacuum pump and recirculating chiller as well as the electricity needed for the recovery operation, while cost for fresh solvent got reduced. Detailed scheduling of extraction and recovery operation was ignored since labor cost was ignored.

Figure 1 presents a schematic of the scope of this study. Extraction yield of polyphenols is a strong function of the extraction protocols, i.e., technique and solvent. In conjunction with yield and relevant inputs related to cost of various contributing factors, we obtain cost of production (CP) per unit of a specific type of polyphenols, and that is considered as a metric of evaluation. Similarly, the CO2 emission potential for production of the organic solvents (as per standard values in literature [8-10] in addition with the electricity required during extraction, in conjunction with yield provide us a mean to compare the protocols as per emission from production (EP) for unit production of a specific type of polyphenols over the same operation period, as stated above. For DW, the emission was based on power demand in the distillation equipment. DW was not recovered as it would always be available in any lab facility at minimal cost.

It is to be noted that the recycle frequency plays an important role in our study. We consider six scenarios in this regard – recycle frequency of zero defining use of fresh solvent daily. For recycle frequency of 'n", the scenario assumes that extraction starts with fresh solvent on day 1 and then, for n consecutive days, solvents are continued to be recovered after each extraction operation. After n days, again the process is repeated with one day of operation with fresh solvent and n days again, with recovered solvent. Yield is assumed to be dropping similar to that reported in literature [11] by a factor of 0.9 for each recovery frequency. The trade-off that all these factors bear on CP and EP per unit production is important and discussed in due course.

RESULTS AND DISCUSSIONS

Based on yield, it was noticed that for phenolics, DW

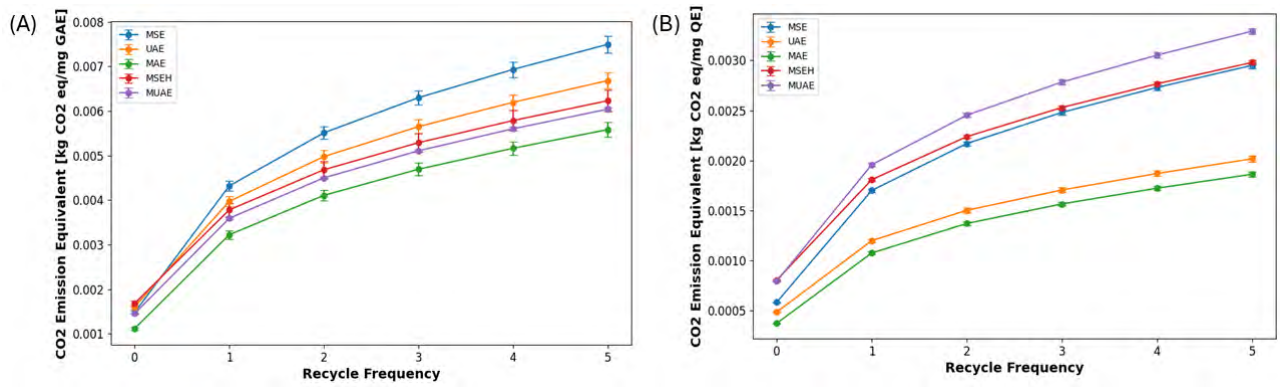

Figure 7. Variation in Emission per Unit Production of Polyphenols with Increase in Recycle Frequency: (A): Phenolics, (B): Flavonoids, considering acetic acid as an example solvent.

with MUAE was the most efficient, while for flavonoids, it was AA, with MAE. Variation in yield with extraction protocols closely followed prospective impact of most critical solvent properties, as relevant for a specific extraction technique. As stated earlier, the major focus of this present work is comparative evaluation of five extraction techniques and those many solvents from the perspectives of sustainability.

Techno-economic Analysis

We firstly focus on the findings from a techno-economic evaluation. The trade-off between yield obtained by a certain protocol and the associated cost incurred is studied here. This study is important since without this, one would not know whether a protocol with low yield with a low cost is preferred over one with high yield at high cost.

From Figure 2, we notice that DW was the most favoured solvent for all the extraction techniques from the perspective of achieving a high enough yield of either type of polyphenols at low enough expenditure. Ethanol, being the most expensive solvent among all (for the AR/ACS grade, in India), was the least favoured. Among the extraction techniques, the relative superiority varied with the choice of solvent as well as the target polyphenol and there was no clear trend identical for all the solvents for either of the polyphenols. Figure 3 presents the relative contribution of the four cost components – for acetone, as an example. We notice that for all the techniques, solvent cost was the most dominant cost components. The findings were same (with difference only in % contribution values) for the other solvents. These % contribution were same for both the types of polyphenols since the cost for utilizing 500 mL solvent in extraction over one year as stated before would be same for either type of polyphenols.

Impact of Recovery of Solvent

As the frequency of recycle increases, fresh solvent amount drops – this, its cost and emission would also drop. However, recycle of solvent also requires equipment for solvent recovery – thereby increasing equipment cost and power cost (for power, operating those extra equipment).

With the yield dropping and cost also majorly dropping (as solvent contribution was so dominant on process economics for extraction with fresh solvent), the net impact can be seen in Figure 5, wherein we noticed a drop in cost as n increases for all the scenarios. In terms of relative cost contribution, equipment and power now had more foothold as expected (Figure 4 as against Figure 3, for acetone as an example).

Emission Potential of the Extraction Protocols

We also evaluated the EP for unit amount of polyphenols as shown in Figure 6 for fresh solvent. Like CP, EP was also, as expected, a strong function of extraction protocol. However, percentage change in EP was not that much significant across the techniques for any solvent, for majority of the instances. Among all the scenario, acetone's outlook was the poorest for phenolics and also for MAE and MSEH for flavonoids extraction. Ethanol with MSE was the most superior among all options.

Figure 7 reflects relative change in EP with increase in recycle frequency for acetic acid, as an example solvent. Even though emission due to solvent production dropped (for use of lower amount of fresh solvent), that due to power demand for solvent recovery increased. Moreover, with increase in recycle frequency, production of polyphenols also took a hot. As a result, we notice an increase in EP with recycle frequency for all the extraction techniques. A similar trend was observed for other solvents, as well.

CONCLUSIONS

Comparative evaluation of extraction protocols

(technique and solvent) from the critical perspectives of techno-economics and emission potential was the major focus of this study. We found that the relative superiority of extraction protocols as per either of these perspectives differed as per choice of target polyphenols. While this work is an important step towards sustainability assessment for phytochemical extraction from agro-food waste, consideration of other environmental impact assessment indicators would be an important next step. Moreover, pareto-optimality of this two-objective (cost and environmental impact) would enable us to have a sounder decision-making for extraction protocol. Moreover, while this work has highlighted important findings for polyphenols extraction from a given quantity of banana peel (since volume of solution per batch was same), it would also be very informative to explore the optimal decisions for a certain demand profile of polyphenols.

ACKNOWLEDGEMENTS

The authors acknowledge the facilities and support at IIT(ISM) Dhanbad and specifically, those in Department of Chemical Engineering in connection with this work.

REFERENCES

1. González-Montelongo R, Gloria Lobo M, González M. Antioxidant activity in banana peel extracts: testing extraction conditions and related bioactive compounds. Food Chemistry 119:1030-1039 (2010). https://doi.org/10.1016/j.foodchem.2009.08.012
2. Ozdemir M, Gungor V, Melikoglu M, Aydiner C. Solvent selection and effect of extraction conditions on ultrasound-assisted extraction of phenolic compounds from galangal (alpinia officinarum). Journal of Applied Research on Medicinal and Aromatic Plants 38:100525 (2024). https://doi.org/10.1016/j.jarmap.2023.100525
3. Rezaie M, Farhoosh R, Iranshahi M, Sharif A, Golmohamadzadeh S. Ultrasonic-assisted extraction of antioxidative compounds from bene (pistacia atlantica subsp. mutica) hull using various solvents of different physicochemical properties. Food Chemistry 173:577-583 (2015). https://doi.org/10.1016/j.foodchem.2014.10.081
4. Tai N, Linh MN, Thuy NM. Optimization of extraction conditions of phytochemical compounds in “xiem” banana peel powder using response surface methodology. J App Biol Biotech : (2021). https://doi.org/10.7324/jabb.2021.9607
5. Vu HT, Scarlett CJ, Vuong QV. Maximising recovery of phenolic compounds and antioxidant properties from banana peel using microwave assisted extraction and water. J Food Sci Technol 56:1360-1370 (2019). https://doi.org/10.1007/s13197-019-03610-2
6. Martínez-Ramos T, Benedito-Fort J, Watson NJ, Ruiz-López II, Che-Galicia G, Corona-Jiménez E. Effect of solvent composition and its interaction with ultrasonic energy on the ultrasound-assisted extraction of phenolic compounds from mango peels (mangifera indica L.). Food and Bioproducts Processing 122:41-54 (2020). https://doi.org/10.1016/j.fbp.2020.03.011
7. Mao Y, Robinson J, Binner E. Understanding heat and mass transfer processes during microwave-assisted and conventional solvent extraction. Chemical Engineering Science 233:116418 (2021). https://doi.org/10.1016/j.ces.2020.116418
8. Juan D. Medrano-García, Rubén Ruiz-Femenia, Jose A. Caballero, Revisiting Classic Acetic Acid Synthesis: Optimal Hydrogen Consumption and Carbon Dioxide Utilization, Editor(s): Anton A. Kiss, Edwin Zondervan, Richard Lakerveld, Leyla Özkan. Computer Aided Chemical Engineering, Elsevier, Volume 46: 145-150 (2019)
9. Methanol Institute, https://methanol.org/wp-content/uploads/2022/01/CARBON-FOOTPRINT-OF-METHANOL-PAPER_1-31-22.pdf
10. The Innovation Platform ISSUE 16. https://www.pyroco2.eu/wp-content/uploads/2023/11/Unlocking-the-green-potential-Transforming-industrial-CO2-into-acetone-with-carbon-capture-and-utilisation-technology.pdf
11. Zhou P, Wang X, Liu P, Huang J, Wang C, Pan M, Kuang Z. Enhanced phenolic compounds extraction from morus alba L. leaves by deep eutectic solvents combined with ultrasonic-assisted extraction. Industrial Crops and Products 120:147-154 (2018). https://doi.org/10.1016/j.indcrop.2018.04.071

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Whole Systems Thinking Model Towards Optimal Decarbonization Strategies for China's Cement Sector

Yushu Wang[ab], Wenli Du[a], Minglei Yang[a], and Vassilis M. Charitopoulos[b]*

[a] Key Laboratory of Smart Manufacturing in Energy Chemical Process, Ministry of Education, East China University of Science and Technology, Shanghai, 200237, China
[b] Department of Chemical Engineering, The Sargent Centre for Process Systems Engineering, University College London (UCL), Torrington Place, WC1E 7JE, UK
* Corresponding Author: v.charitopoulos@ucl.ac.uk.

ABSTRACT

China's cement industry accounts for over half of global production and contributes 8% of global CO_2 emissions, making its decarbonization critical for achieving climate targets. While carbon capture and storage (CCS) and carbon capture and utilization (CCU) are essential deep decarbonization technologies, existing research has not adequately addressed the regional and temporal variations needed for optimal pathway selection across China's diverse provinces. This study develops a comprehensive whole-systems optimization model to design provincial-scale decarbonization pathways for China's cement industry from 2025 to 2060. The model reveals significant spatial and temporal heterogeneity in optimal technology combinations. Before 2050, traditional cement processes integrated with CCS (TCP-CCS) represent the dominant bridging technology for low-carbon transition. However, reaching carbon neutrality by 2060 necessitates an eventual shift toward widespread deployment of novel chemical processes combined with green hydrogen (NCP-H2) as the ultimate decarbonization pathway. Notably, regional natural gas prices significantly affect technology feasibility and deployment timing. The optimized transition pathway reduces unit clinker costs by approximately 58% compared to 2025 levels while successfully meeting the 2060 carbon neutrality target. This research framework provides quantitative decision support for policymakers to design cost-effective, province-specific decarbonization strategies that align with local resource endowments and economic conditions, ultimately advancing deep decarbonization across China's industrial sector.

Keywords: Cement decarbonization, Whole systems thinking, Supply chains, CCS, Green hydrogen, MILP

INTRODUCTION

In the face of global warming, countries worldwide are pursuing economic development while striving for environmental sustainability [1-2]. The cement industry, as a cornerstone of the built environment, accounts for a significant percentage of global carbon emissions, making it a critical area for decarbonization [3]. China, the world's largest cement producer [4], plays a pivotal role in achieving both domestic and global cement industry carbon neutrality. Given that 58% of cement industry emissions originate from calcination processes [5], deep decarbonization cannot be achieved solely through alternative energy sources or equipment improvements. Even emerging methane reforming technology for carbonate decarbonization requires integration with renewable energy technologies such as green hydrogen and carbon capture and utilization (CCU) to achieve net-zero emissions [6-7]. Additionally, carbon capture and storage (CCS) pathways represent another important pathway for deep decarbonization [8]. A comprehensive comparison of economic and environmental performance for these technologies, along with their applicable conditions, is urgently needed.

While numerous studies have investigated carbon reduction methods for the cement industry across China [9], South Korea [10], Europe [11], and globally [12]. However, most studies only evaluate and compare different processes under single or uniform conditions, failing to consider how regional variations might affect technology

performance rankings and optimal application scenarios. Few studies have addressed regional heterogeneity through refined spatial modeling, with notable examples including comparative analyses across European countries. [13]. For China, significant resource disparities exist across regions, with varying geological resources, renewable energy potential, and cement demand [14]. A one-size-fits-all approach has clear limitations; each province needs to adopt appropriate decarbonization methods based on local conditions to meet the national 2060 carbon neutrality vision. However, research comparing green hydrogen and CCS technologies across different Chinese provinces remains limited.

Furthermore, given the long timeframe to 2060, considering dynamic changes in decarbonization technologies over time is imperative [15-16]. Most studies focus on technology deployment endpoints rather than how decarbonization pathways will evolve throughout the entire transition period, including specific deployment timing and sequencing [17]. These long-term, large-scale technological transformations will significantly impact overall economic and environmental performance, necessitating whole systems thinking models.

This study develops a comprehensive optimization model for China's provincial cement industry low-carbon transformation. It integrates regional heterogeneity including renewable energy resources, fossil fuel prices, electricity costs, CO_2 geological storage conditions, and cement demand variations, while considering inter-provincial transportation of materials. Under constraints regarding cement demand, carbon emission targets, and technological capacity limitations, the model minimizes the total discounted net cost of system to provide tailored transformation solutions for different regions and identify advantageous scenarios for different technologies. Through multiperiod spatially explicit optimization, the model considers annual capacity changes, production adjustments, and inventory management, providing optimal deployment timing and sequencing for low-carbon technologies. This research framework offers detailed transformation pathway planning for cement enterprises and provides policymakers with optimal solutions and policy recommendations for achieving carbon neutrality.

MATHEMATICAL FORMULATION

Objective Function

The model minimizes the total discounted net cost (TDNC) of system from 2025 to 2060:

$$\min \mathrm{TDNC} = \sum_t \mathrm{DF}_t \times \mathrm{TC}_t \tag{1}$$

$$\mathrm{DF}_t = \frac{1}{(1+\phi)^{t-1}} \quad \forall t \tag{2}$$

where DF_t the discount factor for year t (assuming 5% annual discount rate), and TC_t is the total cost in year t.

Total cost comprises capital costs (cement processes and energy facilities), fixed operating costs, and variable operating costs including material procurement, transportation, electricity/hydrogen trading, carbon tax, CO_2 storage, and inventory holding, retirement costs and offset by methanol revenue.

$$\begin{aligned}\mathrm{TC}_t = {} & \mathrm{C}_t^{cap,cem} + \mathrm{C}_t^{cap,ene} + \mathrm{C}_t^{fix,cem} + \mathrm{C}_t^{fix,ene} + \mathrm{C}_t^{mat} + \mathrm{C}_t^{inv} \\ & + \mathrm{C}_t^{trans} + \mathrm{C}_t^{ele} + \mathrm{C}_t^{H_2} + \mathrm{C}_t^{Ctax} + \mathrm{C}_t^{sto} + \mathrm{C}_t^{ret} - R_t^{MeOH} \quad \forall t\end{aligned} \tag{3}$$

Material and Energy Balance Constraints

Material flows are constrained to adjacent provinces, with transportation between non-adjacent regions achieved through sequential routing via intermediate provinces.

$$\begin{aligned}&\mathrm{Trans}_{r,r',t}^{Coal} = \mathrm{Trans}_{r,r',t}^{CO_2} = \mathrm{Trans}_{r,r',t}^{cem} = 0 \\ &\forall t, (r,r') \notin ADJ, r = r'\end{aligned} \tag{4}$$

where *ADJ* denotes the set of adjacent province pairs.

The cement inventory balance ensures production plus net imports equals demand plus net exports, maintaining dynamic equilibrium across different years. Similar balance equations govern coal and CO_2 flows across provinces, enabling inter-provincial resource allocation and transportation optimization.

$$\begin{aligned}I_{r,t}^{cem} = {} & I_{r,t-1}^{cem} + \sum_i \mathrm{Pro}_{i,r,t} + \sum_{r' \in ADJ(r)} Trans_{r',r,t}^{cem} \\ & - \mathrm{dem}_{r,t} - \sum_{r' \in ADJ(r)} Trans_{r,r',t}^{cem} \quad \forall r,t\end{aligned} \tag{5}$$

Electric power balance:

$$\begin{aligned}&\sum_{j \in \{\mathrm{Wind,PV}\}} \left(\omega_j^{\mathrm{Ele}} \times \mathrm{Cap}_{j,r,t}^{ene} \times \mathrm{h}\right) \\ &+ Ele_{r,t}^{buy} = \sum_i \left(\alpha_{i,\mathrm{Ele}} \times \mathrm{Pro}_{i,r,t}\right) \\ &+ \left(\omega_{\mathrm{EC}}^{\mathrm{Ele}} \times \mathrm{Cap}_{EC,r,t}^{ene} \times \mathrm{h}\right) + Ele_{r,t}^{sell} \quad \forall \mathrm{r,t}\end{aligned} \tag{6}$$

Hydrogen balance:

$$\begin{aligned}&\omega_{\mathrm{EC}}^{\mathrm{H_2}} \times \mathrm{Cap}_{EC,r,t}^{ene} \times \mathrm{h} + H_{r,t}^{buy} = \\ &\sum_i \left(\alpha_i^{\mathrm{H_2}} \times \mathrm{Pro}_{i,r,t}\right) + H_{r,t}^{sell} \quad \forall \mathrm{r,t}\end{aligned} \tag{7}$$

where h is 8760 hours. ω represents the electricity and hydrogen production coefficients per unit capacity of energy facilities (MWh/MW or ton H_2/MW). α represents the electricity and hydrogen consumption coefficients for cement production process i (MWh/ton clinker or ton H_2/ton clinker).

Capacity Evolution Constraints

The evolution of cement production capacity over time is calculated in Equations (8) and (9). These equations account for new investments, automatic lifetime retirements, and early decommissioning, ensuring a

dynamic balance of capacity.

$$\begin{aligned}&\mathrm{Cap}_{i,r,t}^{cem}=\mathrm{Cap}_{i,r,t-1}^{cem}+\mathrm{Cap}_{i,r,t}^{cem,new}\\&-\mathrm{Cap}_{i,r,t-te}^{cem,new}-\mathrm{Cap}_{i,r,t}^{cem,dec}\quad\forall i,r,t>1\end{aligned}\tag{8}$$

$$\mathrm{Cap}_{i,r,t}^{cem}=\mathrm{cap}_{i,r,0}^{cem}\quad\forall i,r,t=1\tag{9}$$

Equations (10) establish lower and upper bounds on new capacity additions, linking the continuous variable to the binary investment indicator.

$$\begin{aligned}&\mathrm{cap}_i^{Min}\times\mathrm{Build}_{i,r,t}\le\mathrm{Cap}_{i,r,t}^{cem,new}\\&\le\mathrm{cap}_i^{Max}\times\mathrm{Build}_{i,r,t}\quad\forall i,r,t\end{aligned}\tag{10}$$

$$\mathrm{cap}_i^{Max}=\lambda\times cap_r^{cem,totalMax}\tag{11}$$

where $\mathrm{Build}_{i,r,t}$ is a binary variable indicating whether process i is built in province r and year t, $cap_r^{cem,totalMax}$ represents the maximum allowable total cement capacity in province r, determined by historical data, cap_i^{Min} is the minimum economic capacity for process i, and cap_i^{Max} are the maximum economic capacity, calculated in Equation (11) as a fraction λ of the total provincial limit.

Several constraints are imposed to limit capacity and production. Equation (12) constrains the aggregate cement production capacity across all technologies within each province. Equation (13) ensures that early decommissioning does not exceed existing capacity from the previous period. Equation (14) constrains clinker production to not exceed the installed capacity.

$$\sum_i\mathrm{Cap}_{i,r,t}^{cem}\le cap_r^{cem,totalMax}\quad\forall r,t\tag{12}$$

$$\mathrm{Cap}_{i,r,t}^{cem,dec}\le\mathrm{Cap}_{i,r,t-1}^{cem}\quad\forall i,r,t>1\tag{13}$$

$$\mathrm{Pro}_{i,r,t}\le\mathrm{Cap}_{i,r,t}^{cem}\quad\forall i,r,t\tag{14}$$

Similar power generation capacity constraints are employed, which for the sake of brevity are not explicitly introduced herein.

CO_2 Management Constraints

CO_2 balance ensures sources (process emissions and inter-provincial imports) equal sinks (storage and exports).

$$\begin{aligned}&\sum_{i\in\{TCP-CCS,NCP-CCS\}}\left(\beta_i^{cem}\times\mathrm{Pro}_{i,r,t}\right)+\sum_{r'\in ADJ(r)}\mathrm{Trans}_{r',r,t}^{CO_2}\\&=\mathrm{CS}_{r,t}+\sum_{r'\in ADJ(r)}\mathrm{Trans}_{r,r',t}^{CO_2}\quad\forall r,t\end{aligned}\tag{15}$$

where β_i is the emission factor of process i for clinker production.

The system emissions must not exceed carbon cap targets aligned with China's carbon neutrality pathway in each year t.

$$\mathrm{Emi}_t^{dir}=\sum_i\sum_r\beta_i^{cem}\times\mathrm{Pro}_{i,r,t}\quad\forall t\tag{16}$$

$$\mathrm{Emi}_t^{dir}\le tce_t^{Cap}\quad\forall t\tag{17}$$

RESULTS

Baseline Scenario Model Results

The MILP model was implemented in GAMS 48.1.0 on a desktop computer equipped with a 13th Gen Intel(R) Core (TM) i9-13900K CPU @ 3.00 GHz and 128 GB RAM, and was solved by the Gurobi 12.0.1 solver to obtain the globally optimal solution.

Based on data from various literature sources and institutions such as International Energy Statistics (IES) [18], China's cement clinker demand is projected to decline annually from approximately 1, 233 million tons (Mt) in 2025 to 563 Mt in 2060. Figure 1a demonstrates that the model's total cement clinker installed capacity remains relatively smooth and closely tracks total demand throughout the planning horizon. Furthermore, actual production levels closely approximate installed capacity, indicating high-capacity utilization rates and optimal model performance. The divergence between production and demand curves represents the model's strategic allocation of cement inventory.

Figures 1b and 1c present the total and unit clinker production costs for each year during the planning period (2025-2060), respectively. The black and red lines in Figures 1b and 1c represent pre-discount and post-discount cost values, respectively. While the black line shows an overall upward trend, the red line exhibits a downward trend. The red line is better highlighting both the advantages and necessity for enterprises to transition to low-carbon production. Although carbon tax prices increase annually, enterprises can reduce costs by adopting low-carbon processes. The red line in Figure 1c shows that unit clinker production cost decreases from 540 CNY/ton to 226 CNY/ton (comparing smooth curve values while ignoring peaks). Notably, the red line peaks around 2025, 2040, and 2050, indicating these years represent peak periods for decommissioning old processes and constructing new processes facilities. Figure 1e further proves this observation through its depiction of economic cost composition. This suggests China's cement industry will undergo three distinct low-carbon transformation periods by 2060.

As illustrated in Figure 1d, the low-carbon transformation of cement clinker production can be divided into three stages: from 2025 to 2050, traditional cement processes (TCP) will significantly decline until complete phase-out; from 2030 to 2050, traditional cement processes with carbon capture and storage systems (TCP-CCS) will serve as the primary decarbonization pathway; from 2040 onwards, novel cement processes with green

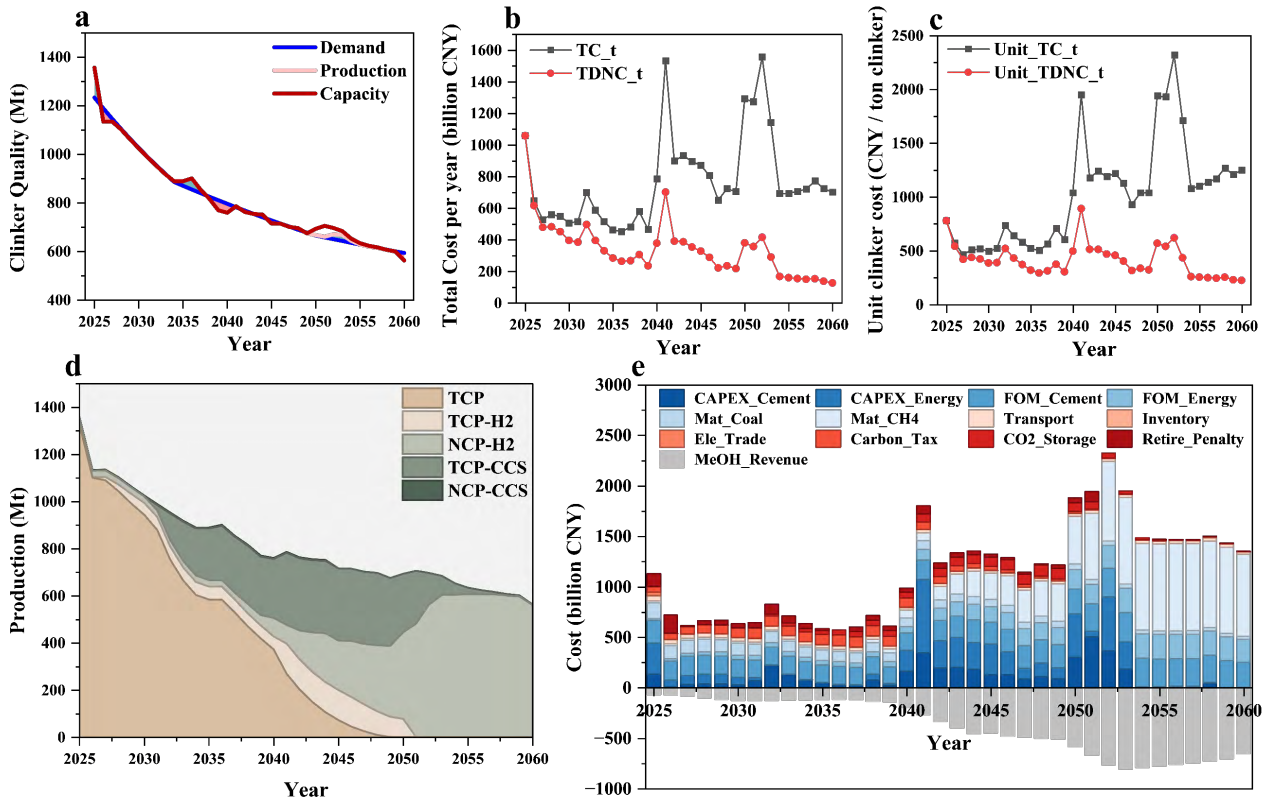

Figure 1. (a) Temporal evolution of cement clinker demand, production, and installed capacity. (b, c) Total cost and unit clinker cost trajectories over time, where the red line represents total discounted net cost and the black line represents pre-discount total cost. (d) Optimal cement decarbonization pathways from 2025 to 2060, where TCP denotes traditional cement process, NCP represents novel chemical process via methane reforming, TCP-H2 and NCP-H2 indicate integration with green hydrogen technology, and TCP-CCS and NCP-CCS represent integration with carbon capture and storage systems. (e) Economic cost composition across different years.

hydrogen (NCP-H2) will gradually demonstrate superior decarbonation advantages, becoming the dominant technology after 2050 and ultimately capturing nearly the entire cement clinker market. This progression occurs because NCP-H2 possesses stronger deep decarbonization capabilities compared to other processes, and its net-zero emission capacity makes it the optimal choice for meeting China's 2060 carbon neutrality target.

Meanwhile, Figure 1e also reveals that methanol revenue gradually increases after 2040, reaching a peak in 2053. After 2053, methanol revenue declines as cement clinker demand and process capacity decreases. Additionally, natural gas costs account for a substantial proportion of total costs because the NCP-H2 process requires natural gas as a raw material, meaning natural gas prices will significantly influence cement clinker production costs during this period.

Regional Technology Deployment Pathways

Figure 2a illustrates the temporal evolution of low-carbon cement processes across different Chinese provinces. Figure 2b displays natural gas prices across different provinces, highlighting how regional price disparities result in different economic advantages for NCP-H2 and influence its adoption. The evolution of low-carbon cement processes in each province can be broadly categorized into three patterns:

Firstly, in provinces where natural gas prices remain below 5, 050 CNY per ton (represented by yellow bars), the NCP-H2 process demonstrates superior economics compared to TCP-CCS. These provinces show greater willingness to adopt NCP-H2 for low-carbon transformation, with provinces having lower natural gas prices transitioning from traditional cement processes (TCP) to NCP-H2 earlier in the planning horizon.

Secondly, in provinces where natural gas prices range from 5, 050 to 5, 500 CNY per ton (represented by light green bars), the economic advantage of NCP-H2 over TCP-CCS becomes marginal. Consequently, these provinces primarily utilize TCP-CCS for decarbonization before 2050, then transition to NCP-H2 in the post-2050 period to meet increasingly stringent carbon neutrality

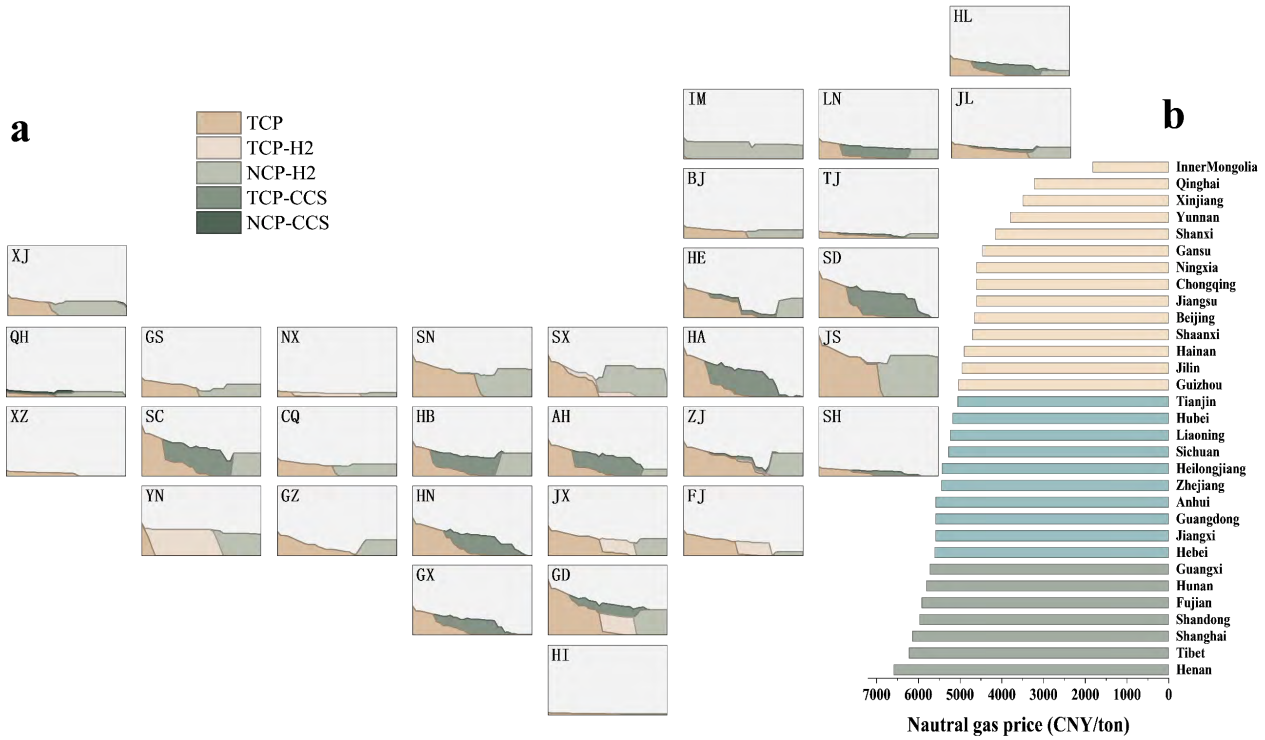

Figure 2. (a) Provincial cement decarbonization pathway trajectories across China (excluding Hong Kong, Macao, and Taiwan), with all subplots sharing identical x-axis (years 2025-2060) and y-axis (cement clinker production) scales. (b) Natural gas prices across Chinese provinces, categorized into three tiers (low, medium, high) represented by different colors. Provinces across different price tiers demonstrate markedly divergent decarbonization pathway trajectories.

requirements.

For provinces where natural gas prices exceed 5, 500 CNY per ton (represented by dark green bars), NCP-H2 loses its economic competitiveness entirely. These provinces rely exclusively on TCP-CCS throughout their decarbonization pathway. When local TCP-CCS clinker production capacity cannot satisfy demand, additional clinker is sourced from neighboring provinces through inter-provincial trade.

Several provinces exhibit unique decarbonization patterns worth noting. Yunnan (YN), benefiting from favorable solar radiation conditions conducive to photovoltaic development and green hydrogen production. Simultaneously, it faces substantial clinker demand that makes premature TCP retirement economically prohibitive. Therefore, it adopts TCP-H2 as a transitional process, offering a balanced compromise between cost and environmental performance. Meanwhile, Zhejiang (ZJ), Jiangxi (JX), and Fujian (FJ), constrained by insufficient geological capacity for CO_2 storage, demonstrate stronger preference for hydrogen-based processes (TCP-H2 or NCP-H2) rather than CCS-dependent processes.

CONCLUSIONS

This study presents a comprehensive whole-system optimization model for China's cement industry decarbonization, incorporating spatial heterogeneity across provinces, cement technologies, renewable energy integration, and CCS. The model provides quantitative insights into optimal technology deployment pathways, identifying when, where, and which technologies should be adopted to achieve carbon neutrality cost-effectively.

The optimization results reveal the principles for selecting low-carbon transformation processes in China's cement industry. First, the low-carbon transformation of China's cement industry requires a gradual transition. Before 2050, the traditional cement process combined with CCS (TCP-CCS) is an important low-carbon transition process. However, to achieve the carbon neutrality target by 2060, the adoption of a novel methane reforming and green hydrogen process (NCP-H2) is the ultimate solution, as it is the only process with the potential for net-zero or even negative emissions. It is worth noting, natural gas prices are a significant factor influencing the NCP-H2 process and determining whether it can surpass TCP-CCS technology and be adopted by different

provinces.

The modeling framework developed in this study provides a template applicable to other hard-to-abate industrial sectors (steel, chemicals) facing similar decarbonization challenges, contributing to broader climate mitigation strategies in China and globally.

Future work should extend this framework to incorporate demand-side uncertainties through two-stage stochastic programming, integrate life cycle assessment across multiple environmental impact categories, and couple with techno-economic analysis for comprehensive sustainability evaluation.

ACKNOWLEDGEMENTS

This work was supported by the National Key Research and Development Program of China (2022YFB3305900) and the Engineering & Physical Sciences Research Council (EPSRC) under grants EP/W003317/1 and EP/V051008/1.

REFERENCES

1. Bounitsis GL, Charitopoulos VM. The value of ammonia towards integrated power and heat system decarbonisation. Sustainable Energy Fuels 8:2914-2940 (2024). https://doi.org/10.1039/d4se00449c
2. Charitopoulos VM, Fajardy M, Chyong CK, Reiner DM. The impact of 100% electrification of domestic heat in great britain. iScience 26:108239 (2023). https://doi.org/10.1016/j.isci.2023.108239
3. Wang Y, Wang R, Tanaka K, Ciais P, Penuelas J, Balkanski Y, Sardans J, Hauglustaine D, Liu W, Xing X, Li J, Xu S, Xiong Y, Yang R, Cao J, Chen J, Wang L, Tang X, Zhang R. Accelerating the energy transition towards photovoltaic and wind in china. Nature 619:761-767 (2023). https://doi.org/10.1038/s41586-023-06180-8
4. Dinga CD, Wen Z. China's green deal: can china's cement industry achieve carbon neutral emissions by 2060?. Renewable and Sustainable Energy Reviews 155:111931 (2022). https://doi.org/10.1016/j.rser.2021.111931
5. Ke J, Zheng N, Fridley D, Price L, Zhou N. Potential energy savings and CO2 emissions reduction of china's cement industry. Energy Policy 45:739-751 (2012). https://doi.org/10.1016/j.enpol.2012.03.036
6. Varnier L, d'Amore F, Clausen K, Melitos G, de Groot B, Bezzo F. Combined electrification and carbon capture for low-carbon cement: techno-economic assessment of different designs. Journal of Cleaner Production 498:145029 (2025). https://doi.org/10.1016/j.jclepro.2025.145029
7. Wang Y, Yang M, Charitopoulos VM, Zhou M, Du W. Multi?objective bilevel operational optimization of low?carbon cement industries with integrated energy systems. AIChE Journal 72: (2025). https://doi.org/10.1002/aic.70101
8. Wang Y, Yang M, Shen F, Zhou M, Du W. Conceptual design and life-cycle environmental and economic assessment of low-carbon cement manufacturing processes. Journal of Cleaner Production 471:143349 (2024). https://doi.org/10.1016/j.jclepro.2024.143349
9. Ren M, Ma T, Fang C, Liu X, Guo C, Zhang S, Zhou Z, Zhu Y, Dai H, Huang C. Negative emission technology is key to decarbonizing china's cement industry. Applied Energy 329:120254 (2023). https://doi.org/10.1016/j.apenergy.2022.120254
10. Lee H, McJeon H, Yu S, Liu Y, Kim H, Eom J. Decarbonization pathways for korea's industrial sector towards its 2050 carbon neutrality goal. Journal of Cleaner Production 476:143749 (2024). https://doi.org/10.1016/j.jclepro.2024.143749
11. Cavalett O, Watanabe MDB, Voldsund M, Roussanaly S, Cherubini F. Paving the way for sustainable decarbonization of the european cement industry. Nat Sustain 7:568-580 (2024). https://doi.org/10.1038/s41893-024-01320-y
12. Cheng D, Reiner DM, Yang F, Cui C, Meng J, Shan Y, Liu Y, Tao S, Guan D. Projecting future carbon emissions from cement production in developing countries. Nat Commun 14: (2023). https://doi.org/10.1038/s41467-023-43660-x
13. Strunge T, Küng L, Sunny N, Shah N, Renforth P, Van der Spek M. Finding least-cost net-zero CO _{2e} strategies for the european cement industry using geospatial techno-economic modelling. RSC Sustainability 2:3054-3076 (2024). https://doi.org/10.1039/d4su00373j
14. Wang Y, Wen Z, Xu M, Chen J, He P. Plant-level green transformation strategy in china's cement industry: considering energy conservation and emission reduction co-benefits. Journal of Cleaner Production 467:142945 (2024). https://doi.org/10.1016/j.jclepro.2024.142945
15. Wu X, Meng J, Liang X, Sun L, Coffman D, Kontoleon A, Guan D. Technological pathways for cost-effective steel decarbonization. Nature 647:93-101 (2025). https://doi.org/10.1038/s41586-025-09658-9
16. Jeong Y, Zhang Q. Transition toward a decarbonized steel industry: a supply chain analysis for the united states. Ind. Eng. Chem. Res. 64:16263-16280 (2025). https://doi.org/10.1021/acs.iecr.5c01586
17. Wang Y, Wen Z, Xu M, Doh Dinga C. Long-term transformation in china's steel sector for carbon capture and storage technology deployment. Nat

Commun 16: (2025). https://doi.org/10.1038/s41467-025-59205-3
18. Liao S, Wang D, Xia C, Tang J. China's provincial process CO2 emissions from cement production during 1993–2019. Sci Data 9: (2022). https://doi.org/10.1038/s41597-022-01270-0

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Decision-Support Framework for Process Design of Sustainable Aviation Fuel Production via Integrated Biorefineries

Vibhu Baibhav[a]*, Daniel Florez-Orrego[a], Francois Maréchal[a]

[a] Industrial Process and Energy Systems Engineering Laboratory, École Polytechnique Fédérale de Lausanne (EPFL), Valais-Wallis, Sion, Switzerland.
* Corresponding Author: vibhu.baibhav@epfl.ch.

ABSTRACT

Sustainable aviation fuel (SAF) is a key pathway for mitigating greenhouse gas emissions in aviation, yet its large-scale deployment is constrained by high energy demand and production costs. Among available conversion routes, the hydroprocessed esters and fatty acids (HEFA) pathway is the most commercially mature, but it requires substantial hydrogen input and high-temperature heat, affecting both economic and environmental performance. This study presents a decision-support framework for SAF production via integrated biorefineries, using rapeseed oil extraction coupled with HEFA conversion as a case study. Detailed process simulations are combined with energy integration and mixed-integer linear programming optimization to enable system-level analysis. Material integration strategies include internal hydrogen generation via glycerol steam reforming and valorization of rapeseed meal through gasification for syngas production. Heat integration considers cross-process heat recovery and the application of high-temperature heat pumps to reduce external energy demand. The results demonstrate that coordinated optimization of heat and material flows significantly improves energy efficiency and reduces reliance on fossil utilities. The framework provides a practical tool for identifying economically and environmentally robust SAF production pathways by integrating feedstock conversion, co-product utilization, and utility optimization. These findings support strategic decision-making for next-generation biorefineries and strengthen circular economy principles in sustainable aviation fuel production.

Keywords: Sustainable Aviation Fuel (SAF), HEFA, Biorefinery, Process Integration, MILP Optimization

INTRODUCTION

The aviation sector plays a critical role in enabling economic development and global connectivity, yet it is also a significant contributor to climate change [1]. While aviation accounts for approximately 2.5% of global carbon dioxide emissions, its contribution to total radiative forcing is estimated to be closer to 4% due to the altitude at which emissions are released [2]. Decarbonizing aviation is therefore both urgent and particularly challenging, given the sector's reliance on high-energy-density liquid fuels, stringent fuel specifications, and long aircraft lifetimes.

Sustainable Aviation Fuel (SAF) is widely recognized as the most viable near- to mid-term solution for reducing aviation emissions. SAFs are drop-in fuels with physicochemical properties similar to conventional jet fuel, allowing the use of existing aircraft fleets and fuel infrastructure. In contrast, non-drop-in alternatives such as hydrogen would require entirely new aircraft designs and energy supply chains [3]. Despite this, large-scale SAF deployment
remains constrained by high production costs, significant fossil energy consumption, and limitations associated with feedstock availability and logistics.

Current SAF studies are typically focused on fuel yield and product quality, while energy and utility systems are often neglected. As a result, SAF pathways often exhibit high demands for steam, electricity, and hydrogen, frequently supplied by fossil-based utilities, undermining both economic competitiveness and environmental performance.

Among the available SAF pathways, Hydroprocessing of Esters and Fatty Acids (HEFA) is the most mature and commercially deployed technology, accounting for more than 99% of current global SAF production from lipid-based biomass [4]. In the HEFA process, acylglycerols and free fatty acids are upgraded into drop-in hydrocarbons suitable for jet fuel applications. Triglycerides, which constitute the dominant fraction of lipid-based feedstocks, require extensive upgrading involving deoxygenation and saturation reactions. Oxygen removal proceeds through hydrodeoxygenation, decarbonylation, and decarboxylation pathways, yielding linear hydrocarbons typically in the C8–C18 carbon range [5]. Subsequent hydrocracking and isomerization steps are required to meet jet fuel cold-flow specifications, followed by product separation into jet fuel and co-products such as diesel, naphtha, and propane.

Hydrogen availability is a critical determinant of both the yield and carbon efficiency of the HEFA process. Higher hydrogen addition promotes hydrodeoxygenation and improves carbon retention by reducing CO and CO_2 formation, but it also substantially increases hydrogen demand. Consequently, the environmental benefits of HEFA-based SAF are strongly dependent on the hydrogen source. The use of fossil-derived hydrogen significantly limits achievable greenhouse gas reductions, particularly given the scale of hydrogen consumption required for jet fuel production. While on-site hydrogen production via water electrolysis is a potential solution, it is associated with high capital costs and challenges related to the intermittency of renewable electricity supply. An alternative approach involves producing hydrogen from biomass-derived intermediates, such as glycerol generated during triglyceride hydrolysis, enabling partial internal hydrogen supply and improved process integration [6–7].

Waste cooking oil (WCO) and animal fats are often considered preferred HEFA feedstocks due to their low associated life-cycle emissions and classification as waste-derived resources. However, large-scale deployment is constrained by limited and geographically dispersed supply, collection logistics, and variability in feedstock quality. As a result, waste-derived lipids alone are insufficient to meet projected SAF demand, particularly in Europe. To ensure stable and continuous SAF production, vegetable oils can therefore serve as complementary feedstocks when waste lipid supply is insufficient, as the generic triglyceride to jetfuel conversion pathway remains applicable with comparatively less feedstock pretreatment requirements.

Rapeseed is the most widely cultivated oilseed crop in Europe, driven primarily by demand for rapeseed oil [8], and benefits from an established agricultural and processing infrastructure. Rapeseed oil extraction also generates rapeseed meal as a major co-product [9]. Due to the presence of anti-nutritional compounds, rapeseed meal is currently utilized primarily as animal feed and is not suitable for direct human consumption [8, 10]. While this co-product is traditionally treated as a low-value output in fuel-oriented system assessments, it represents a significant source of biogenic carbon and energy.

Integrating rapeseed oil extraction with downstream HEFA upgrading enables a biorefinery configuration in which both the primary product and co-products are considered explicitly within a unified energy and material system. Such integration creates additional degrees of freedom for process and utility optimization that are not accessible in stand-alone fuel conversion pathways. In particular, coupling feedstock processing, fuel upgrading, and internal hydrogen generation allows systematic assessment of co-product utilization and waste heat recovery within a single modeling framework.

Accordingly, this study develops detailed ASPEN Plus process models for both HEFA-based SAF production and rapeseed oil extraction to enable consistent evaluation of mass and energy flows, utility requirements, and integration opportunities across the full value chain. The integrated modeling framework provides the basis for systematic analysis of heat integration, hydrogen sourcing strategies, and co-product valorization pathways, with the objective of improving overall energy efficiency, reducing dependence on fossil-based utilities, and enhancing carbon utilization within SAF production systems.

METHODOLOGY

This study adopts a system-level approach to evaluate an integrated biorefinery configuration for sustainable aviation fuel (SAF) production. The methodology is structured to first describe the individual process models and subsequently assess their integration at the energy and material level within a unified biorefinery framework.

Jet fuel production process model

An ASPEN Plus process model of the HEFA-based jet fuel production pathway was developed to evaluate mass and energy flows, hydrogen demand, and product yields. The model includes triglyceride hydrolysis, hydrogen production via glycerol steam reforming, hydrotreating, hydrocracking and isomerization, and final product separation.

Triglyceride hydrolysis was modeled to produce free fatty acids and glycerol. Reaction parameters were adapted from Iliuta et al. [11]. Figure 1 illustrates the ASPEN Plus flowsheet of the triglyceride hydrolysis section, including pressure reduction, power recovery, and phase separation into free fatty acids and glycerol. The hydrolysis product stream was expanded from 50 bar to 20 bar and subsequently to 1 bar. Power recovery from flashed

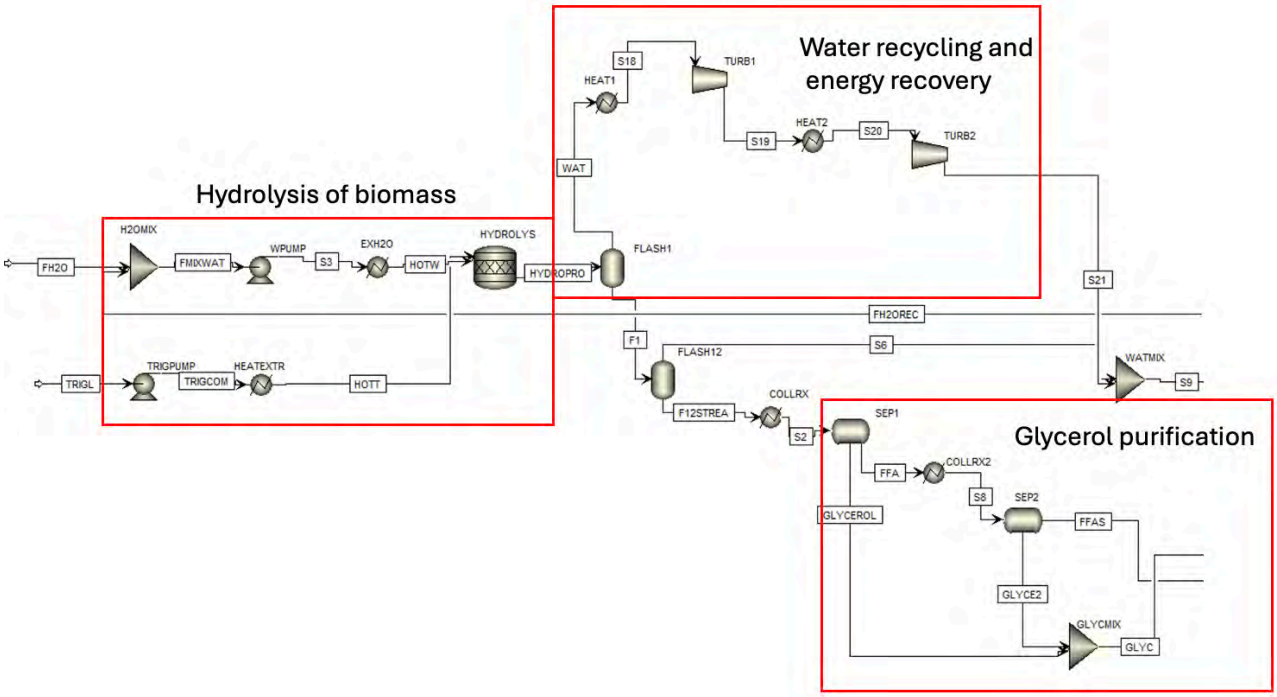

Figure 1. Process flowsheet of the triglyceride hydrolysis section, illustrating the production of free fatty acids and glycerol, pressure reduction stages, power recovery from flashed water vapor, and gravity-based phase separation.

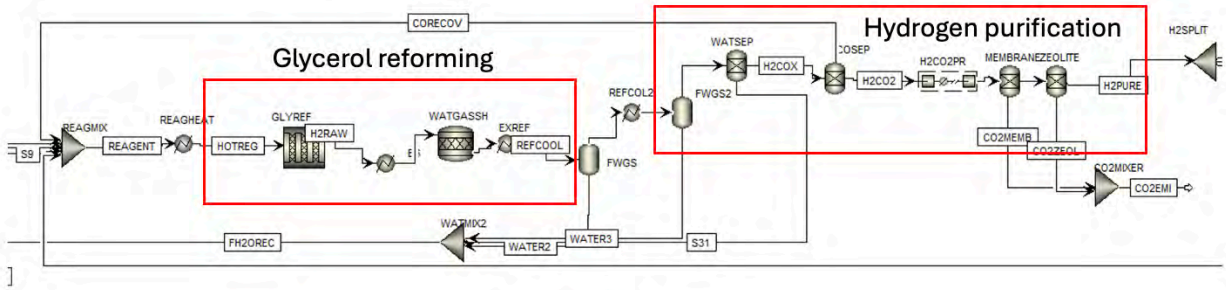

Figure 2. Process flowsheet of the glycerol steam reforming and downstream gas separation units used to produce hydrogen for deoxygenation, ismoerization and cracking steps.

water vapor at 20 bar was modeled using a two-stage turbine. Gravity-based separation of free fatty acids and glycerol was represented using decanter units, following an approach similar to biodiesel process models in Baibhav et al. [12].

The separated glycerol stream was sent to steam reforming for hydrogen production. Glycerol steam reforming was modeled using kinetic parameters obtained from Cheng et al. [13]. Figure 2 presents the ASPEN Plus model of the glycerol steam reforming section and the gas separation units. Downstream gas separation units were included to produce high-purity hydrogen suitable for use in subsequent upgrading reactions. The hydrogen produced via this pathway was supplied internally to the HEFA upgrading section.

Hydrotreating reactions were modeled considering the simultaneous occurrence of hydrodeoxygenation, decarbonylation, and decarboxylation pathways. These reactions convert saturated fatty acids into alkanes while producing water, carbon monoxide, or carbon dioxide, depending on the dominant pathway. The coexistence of these reactions reflects realistic operating conditions and catalyst behavior.

Hydrocracking and isomerization were modeled using RStoic reactors operating at 250 °C and 30 bar, with conversion levels and product yields adopted from Elias et al. [14]. Final product separation was performed using distillation columns, yielding propane, naphtha, diesel, and jet fuel fractions. The resulting jet fuel stream, along with co-products, was assessed in terms of mass and

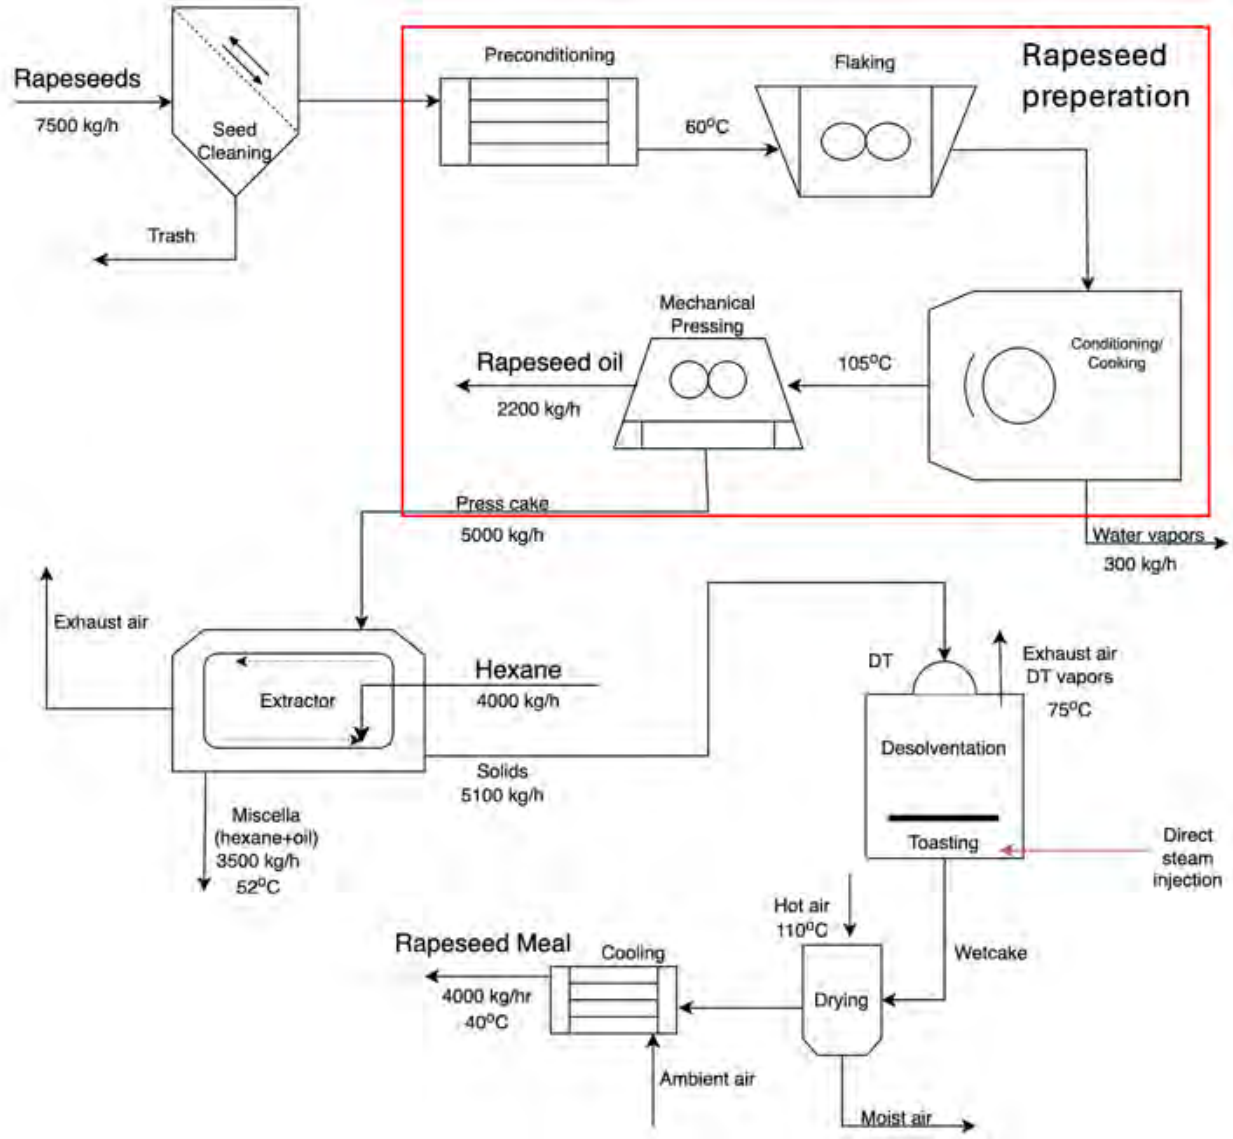

Figure 3. Process flow diagrams of the rapeseed preparation and oil extraction process. **(a)** Rapeseed preparation, crushing, and extraction process along with meal desolventization steps.

energy balances, providing the basis for subsequent process integration and system-level evaluation.

Rapeseed oil production process model

Rapeseed oil production was modeled as an upstream unit operation to supply the lipid feedstock to the HEFA-based jetfuel production process. Unlike soybean oil extraction, which relies exclusively on solvent extraction, rapeseed oil production combines mechanical pressing and solvent-based extraction due to the higher oil content of rapeseed (approximately 40% w/w, compared to ~20% w/w for soybeans). A soybean crushing and oil extraction process previously developed in ASPEN Plus by Granjo et al. [15] was adapted in this study using rapeseed-specific operating conditions and parameters reported by Carré Patrick [16].

Figure 3a shows the process flow diagram of the rapeseed oil production system modeled in ASPEN Plus in this study. Rapeseeds are first cleaned to remove impurities and are then preheated to approximately 60 °C prior to flaking. The flaked seeds are subsequently conditioned at around 100 °C to reduce moisture content and improve oil release and are then fed to a mechanical press for primary oil extraction. Approximately 70% of the total oil content is recovered at this stage, while the remaining solid fraction (press cake) is directed to a solvent extraction unit for further oil recovery.

Hexane, the most widely used industrial solvent for oil extraction [9], is employed in this model to extract the remaining oil from the press cake. The solvent flow rate

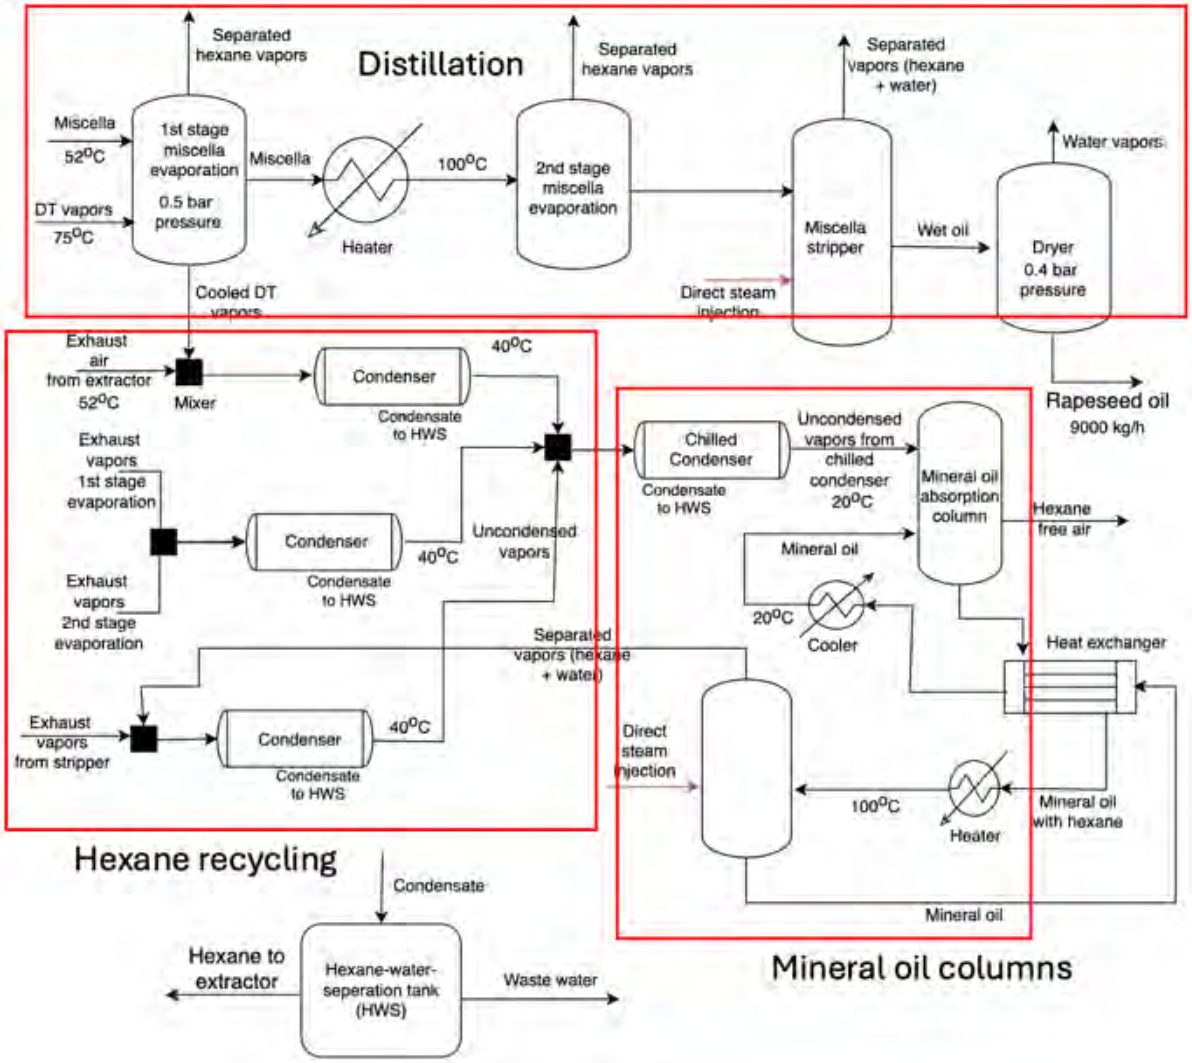

Figure 3 Continued. Process flow diagrams of the rapeseed preparation and oil extraction process. **(b)** Miscella distillation and hexane condensation along with mineral oil absorption steps.

corresponds to approximately 55% (w/w) of the incoming seed flow. In the extractor, hexane recovers the remaining 30% of the oil content and exits the unit via three outlet streams: extracted rapeseed cake, miscella (hexane–oil mixture), and exhaust air. In addition to miscella distillation, hexane is also recovered from the solid cake and exhaust air streams and recycled within the system. Overall, more than 99.9% of the added hexane is recovered [16-17], resulting in negligible solvent losses and limited contribution to operating costs and emissions.

Figure 3b presents the miscella distillation and hexane recovery section of the process. The miscella is treated through a two-stage evaporation system followed by a stripper, where direct steam injection is used to remove residual hexane. The crude oil is subsequently dried to remove moisture introduced during the stripping step and is sent for the refining process. The exhaust air from the extraction process is treated via a mineral oil absorption system to ensure hexane emissions remain below environmental regulatory limits [18]. Efficient absorption requires cooling the gas stream to approximately 20 °C, which is represented in the model by a refrigeration cycle producing chilled water. Absorbed hexane is recovered by direct steam injection into the heated mineral oil and recycled back to the extraction unit.

The solid press cake leaving the extractor, which is saturated with hexane, is sent to the desolventization–toasting–drying–cooling (DTDC) section to produce hexane-free rapeseed meal suitable for animal feed. The desolventization and toasting steps involve heating the meal to approximately 100 °C with direct steam injection, resulting in a DT exhaust stream containing hexane and water vapor. This exhaust is combined with extractor exhaust vapors and routed to the hexane recovery system.

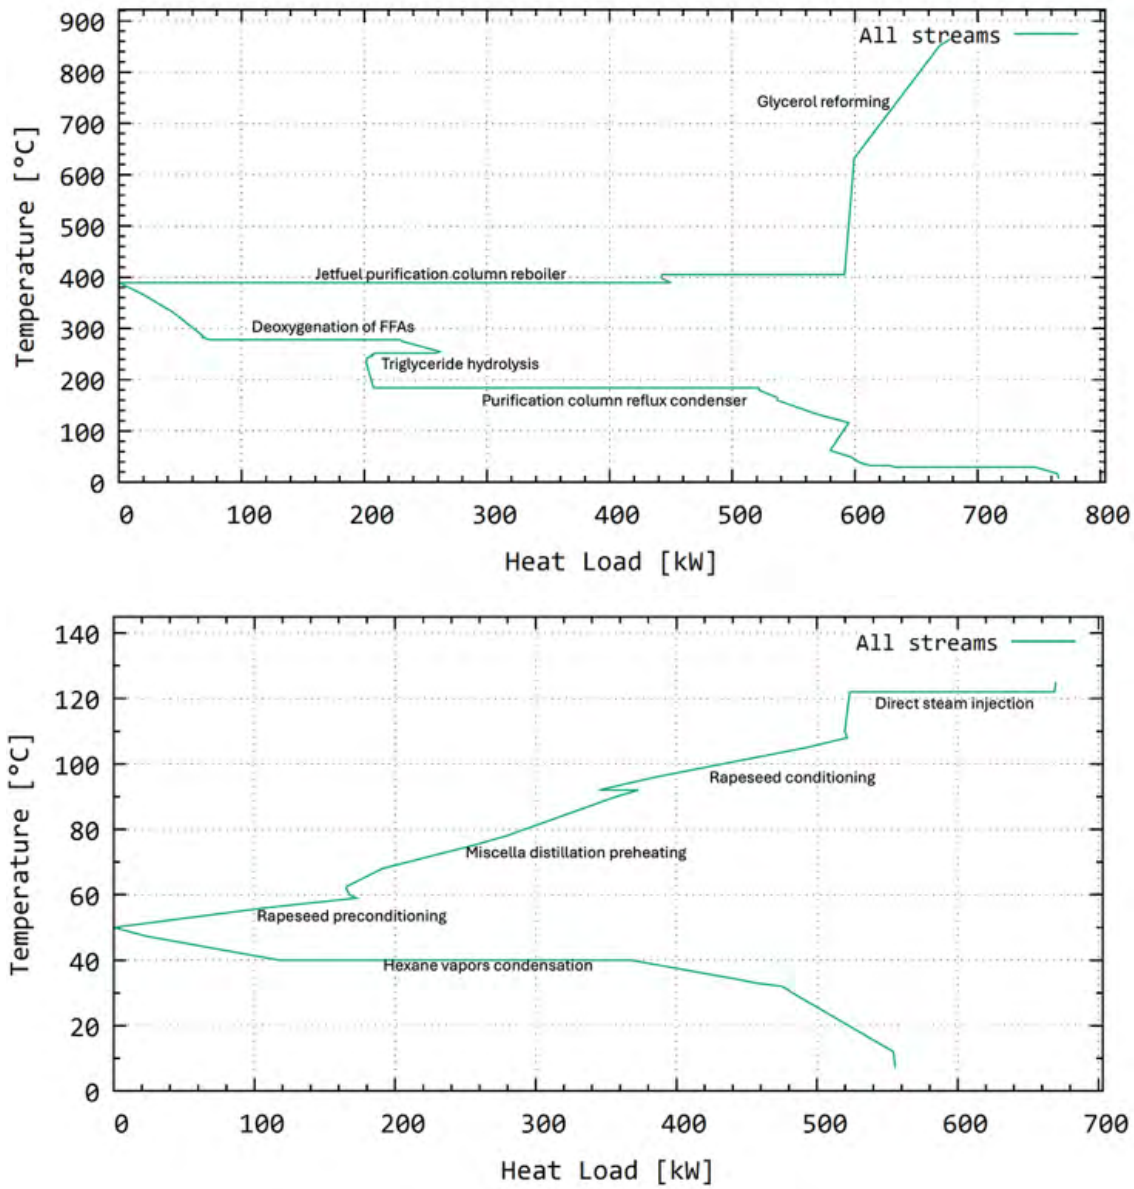

Figure 4. Grand composite curves of the jetfuel production and oil extraction process. (a, top) HEFA-based jet fuel production process (1 t h^{-1} capacity) with glycerol reforming for hydrogen production. Process pinch point temperature of 389 ˚C shows waste heat recovery potential. (b, bottom) Rapeseed oil extraction process (7.5 t h^{-1} crushing capacity) with a low-temperature pinch point at 50 °C.

Due to steam injection, the meal exiting the DT unit has a high moisture content. It therefore undergoes drying using hot air at approximately 105 °C, followed by cooling to around 40 °C using ambient air.

All recovered hexane streams were condensed and sent to a gravity–based hexane–water separation unit, with the recycled solvent returned to the extractor and the separated water directed to wastewater treatment. The extracted rapeseed oil undergoes subsequent refining, while rapeseed meal, produced as a major co-product, was explicitly considered within the integrated biorefinery framework for further valorization.

Process integration and biorefinery

The mass and heat streams obtained from the ASPEN Plus process models were exported to OSMOSE, a system-level optimization platform developed at IPESE, EPFL [19–20], to evaluate energy integration and utility system design. Heat integration was first performed separately for the HEFA-based jet fuel production process

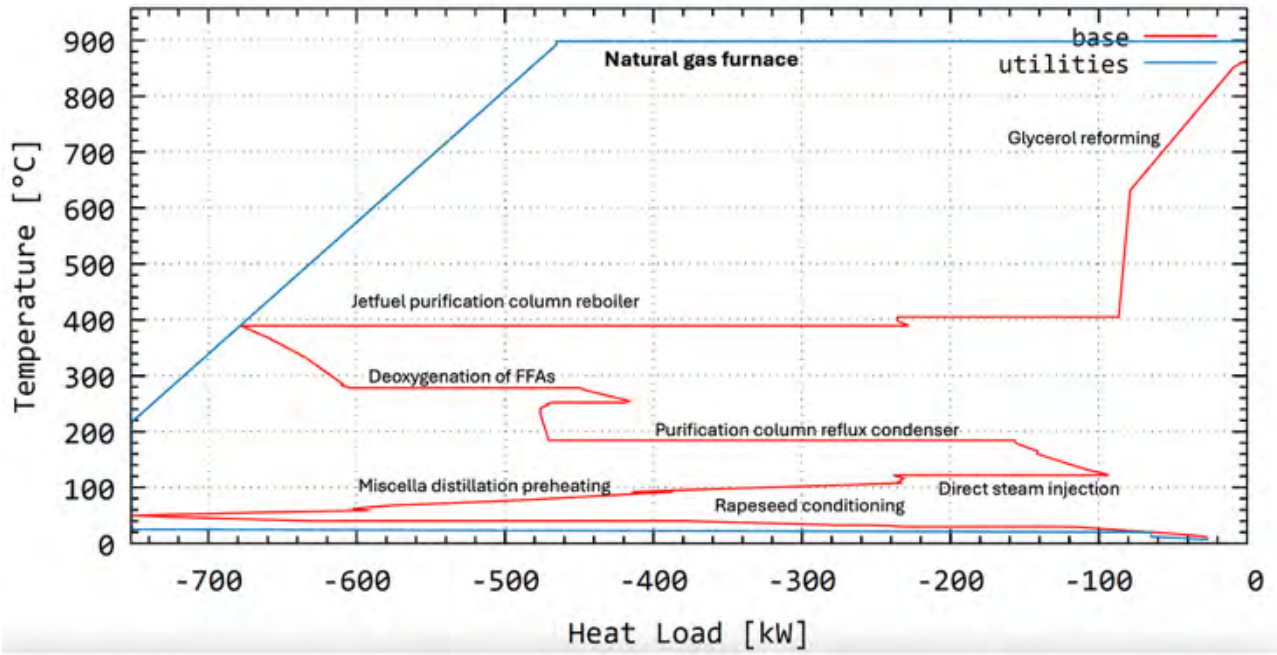

Figure 5. Integrated composite curve of the combined HEFA and rapeseed oil extraction processes. Thermal coupling between subsystems results in approximately 89% reduction in external heating requirement compared to stand-alone production.

and the rapeseed oil extraction process, and subsequently for the c ombined biorefinery configuration to assess potential symbiotic benefits.

The integration framework includes energy conversion technologies such as boilers, cooling water systems, refrigeration cycles, and heat pumps. A mixed-integer linear programming (MILP) formulation was implemented to determine the optimal utility configuration and operating strategy that minimizes total annualized system cost, accounting for both capital and operational expenditures while satisfying process integration constraints. Environmental constraints were incorporated to evaluate decarbonization pathways, including Scope 1 emissions associated with biogenic CO_2 release and natural gas consumption for process heating, as well as Scope 2 emissions linked to electricity use. In addition to heat integration, material integration between subsystems was evaluated, with emphasis on internally generated hydrogen and the valorization of rapeseed meal. Hydrogen produced via glycerol steam reforming was directly supplied to the HEFA upgrading section, reducing external hydrogen requirements. Rapeseed meal was assessed as an alternative energy and carbon carrier through steam gasification, and its use for syngas production was compared against its conventional use as animal feed.

The gasification process was modeled in ASPEN Plus as a sequence of drying, pyrolysis, reduction, and combustion reactions, following the methodology reported by Florez et al. [21]. The proximate composition of rapeseed meal was assumed to be 5.7% moisture, 13.5% fixed carbon, 74% volatile matter, and 6.8% ash, consistent with literature values [22–23]. Mass yields in the pyrolysis step and the resulting syngas composition were calculated using empirical correlations derived from the referenced model [21], providing the basis for integration within the overall biorefinery configuration. The resulting optimization framework enables a systematic assessment of cost–emission trade-offs while simultaneously evaluating heat and material integration options, allowing identification of biorefinery configurations that enhance energy efficiency, reduce external resource dependence, and improve carbon utilization within the SAF production system.

RESULTS AND DISCUSSION

The described modeling framework was applied to evaluate the energetic performance of the rapeseed-based HEFA biorefinery. The analysis assumes a production capacity of 1 ton per hour of jet fuel and a rapeseed crushing capacity of 7.5 tons per hour for oil extraction. Results are presented sequentially, beginning with the individual subsystems, followed by their integration.

Stand-alone process energy analysis

The grand composite curves (GCC) of the HEFA and rapeseed oil extraction processes are shown in Figure 4. The HEFA process (Figure 4a) has minimum heating

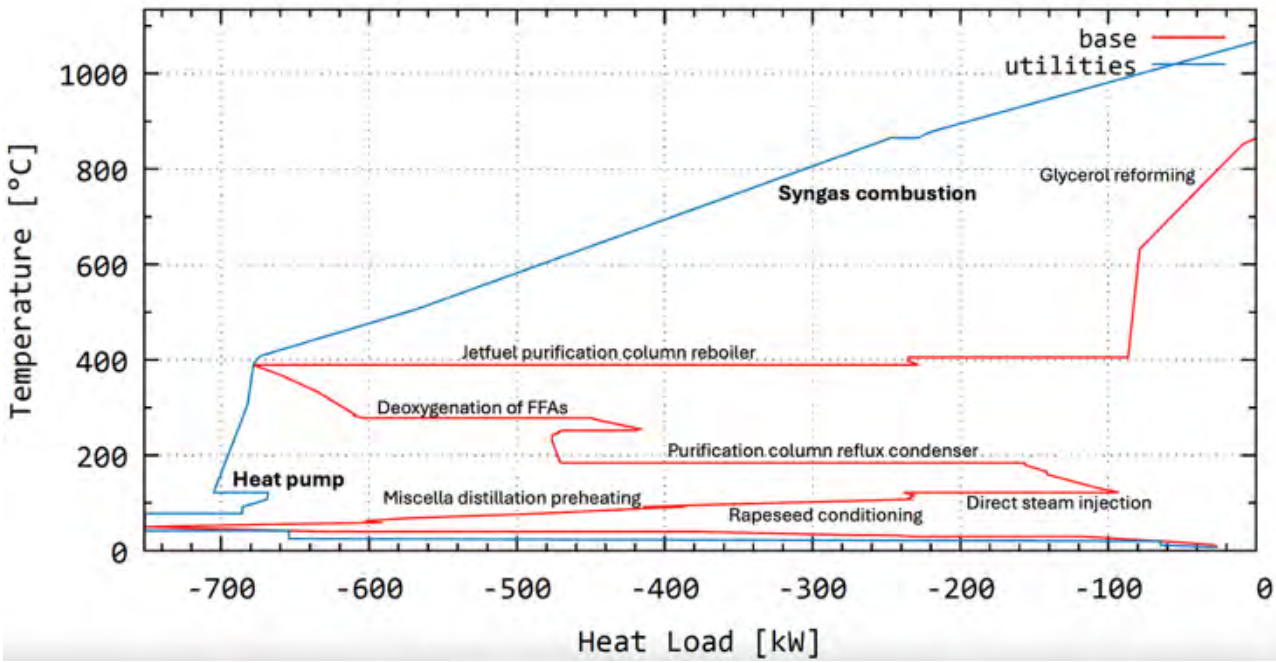

Figure 6. Integrated composite curve of the biorefinery configuration including HEFA production, rapeseed oil extraction, and rapeseed meal gasification. Conversion of 6.2% of the available rapeseed meal to syngas enables full thermal self-sufficiency of the integrated system.

requirements of 677 kWh for 1 ton of jetfuel production. Glycerol steam reforming, operating at 827 °C, supplies the hydrogen required for deoxygenation, isomerization, and hydrocracking. Recycling the diesel fraction from the distillation column back to the hydrocracking reactor increases overall jet fuel yield to 73%, corresponding to approximately 1.37 t h^{-1} of biomass feedstock conversion. The pinch point temperature of 389 °C reflects the external heating requirements at elevated temperature levels with significant thermal demand primarily driven by reforming reactions and the distillation columns.

In contrast, the rapeseed oil extraction process shows a pinch point at 50 °C (Figure 4b). At a crushing capacity of 7.5 t h^{-1}, the process exhibits the same minimum heating requirement of 677 kW as the HEFA subsystem. This operation yields 2.2 t h^{-1} of press oil, 0.9 t h^{-1} of hexane-extracted oil, and 4 t h^{-1} of rapeseed meal as a co-product. The dominant heat duties arise from seed conditioning, hexane recovery, and steam generation at 120 °C for direct injection.

The observed temperature-level mismatch between the two subsystems suggests strong potential for thermal symbiosis, with high-temperature waste heat from the HEFA process potentially satisfying heating requirements of the extraction section.

Heat Integration of the Combined Biorefinery

Figure 5 shows the integrated composite curve of the combined jetfuel and oil extraction process. In this configuration, a natural gas furnace and cooling tower satisfy the residual external heating and cooling demands after internal heat recovery. A chilled-water-based refrigeration cycle is required to meet sub-ambient cooling demands associated with the mineral oil absorption process in the extraction section. The optimized furnace capacity is 746 kW, corresponding to an approximately 89% reduction in external heating demand compared to stand-alone operation of the rapeseed oil extraction process. The reduction in fossil-based heating demand directly lowers operational energy requirements and contributes to improved economic and environmental performance of the integrated HEFA biorefinery.

To further reduce external utility demand, integration with rapeseed meal gasification was evaluated. Figure 6 shows the integrated composite curve with the gasification subsystem included in the biorefinery configuration. The results indicate that only 6.2% of the available rapeseed meal must be converted to syngas to fully satisfy the remaining process heating demand. A two-stage heat pump upgrades low-temperature heat and generates steam at 120 °C required for the gasification process. Under this configuration, the biorefinery achieves full thermal self-sufficiency while preserving the majority of rapeseed meal for its conventional market application as animal feed. The combined heat and material integration strategy, therefore, eliminates fossil-based heating requirements and significantly improves the overall energy autonomy of the HEFA-based jet fuel

production system. Such system-level integration is therefore critical for enhancing the competitiveness of sustainable aviation fuel production pathways.

Outlook and Extended Integration Opportunities

Beyond the discussed configurations, additional valorization pathways may further enhance system circularity and resource efficiency. Rapeseed straw, because of its relatively high calorific value [17], could be considered for direct combustion or co-combustion with other biomass resources to provide supplementary thermal energy. In addition, ash produced from the rapeseed meal gasification process, which is rich in phosphorus and contains significant quantities of potassium, calcium, and magnesium [23], presents potential for agricultural reuse as fertilizer. The application of biochar and nutrient-rich ash to rapeseed cultivation could strengthen the circular-economy framework of the HEFA-based jet fuel production system.

Since only a small fraction of rapeseed meal is required to achieve thermal self-sufficiency, the remaining high-temperature heat from syngas combustion could be further utilized for power generation through a Rankine cycle, potentially covering a significant share of the plant's electricity demand. Future work may also explore the integration of Power-to-Liquid (PtL) pathways for the utilization of biogenic CO_2 streams originating from the HEFA decarboxylation process. Potential strategies include seasonal or underground CO_2 storage in salt caverns [24], coupled with electrolysis-based hydrogen production during periods of excess renewable electricity. Such extensions could further enhance carbon efficiency and address the current limitations in the broader adoption of sustainable aviation fuels.

CONCLUSION

This study presented a system-level evaluation of an integrated rapeseed-based biorefinery for sustainable aviation fuel production via the HEFA pathway. Standalone analysis of the individual subsystems revealed complementary thermal characteristics, with the HEFA process exhibiting a high pinch point temperature and the rapeseed oil extraction process showing a low-temperature pinch. This facilitated substantial cross-process heat recovery when the systems were integrated. Coupling the HEFA and oil extraction processes reduced external heating demand by up to 89% compared to operating the processes independently. Further integration through rapeseed meal gasification demonstrated that only a small fraction of the available biomass byproduct needs to be valorized to achieve thermal self-sufficiency of the combined biorefinery. Also, internally generated hydrogen from glycerol steam reforming reduced reliance on external hydrogen supply, enhancing overall resource efficiency.

These findings underscore the importance of simultaneously evaluating heat and material integration within a unified optimization framework. By leveraging inter-process symbiosis and flexible biomass allocation, the proposed configuration substantially lowers fossil utility dependence while maintaining co-product valorization pathways. System-level integration strategies of this kind are essential for improving the economic and environmental performance of sustainable aviation fuel production and for enabling its broader, large-scale deployment.

REFERENCES

1. Becken S, Mackey B, Lee DS. Implications of preferential access to land and clean energy for sustainable aviation fuels. Science of The Total Environment 886:163883 (2023). https://doi.org/10.1016/j.scitotenv.2023.163883
2. Ritchie, H. (2024). What share of global CO_2 emissions comes from aviation? *Our World in Data*. https://ourworldindata.org/global-aviation-emissions
3. Voß S, Bube S, Kaltschmitt M. Aviation fuel production pathways from lignocellulosic biomass via alcohol intermediates – a technical analysis. Fuel Communications 17:100093 (2023). https://doi.org/10.1016/j.jfueco.2023.100093
4. Maghzian A, (John) Saddler J. Which processes to making biojet/sustainable aviation fuel (SAF) are likely to supplement the predominant, hefa/lipid-to-biojet production route?. Bioresource Technology 441:133631 (2026). https://doi.org/10.1016/j.biortech.2025.133631
5. Rahmawati Z, Santoso L, McCue A, Azua Jamari NL, Ninglasari SY, Gunawan T, Fansuri H. Selectivity of reaction pathways for green diesel production towards biojet fuel applications. RSC Adv. 13:13698-13714 (2023). https://doi.org/10.1039/d3ra02281a
6. Sad ME, Duarte HA, Vignatti C, Padró CL, Apesteguía CR. Steam reforming of glycerol: hydrogen production optimization. International Journal of Hydrogen Energy 40:6097-6106 (2015). https://doi.org/10.1016/j.ijhydene.2015.03.043
7. Luo N, Zhao X, Cao F, Xiao T, Fang D. Thermodynamic study on hydrogen generation from different glycerol reforming processes. Energy Fuels 21:3505-3512 (2007). https://doi.org/10.1021/ef070066g
8. Di Lena G, Sanchez del Pulgar J, Lucarini M, Durazzo A, Ondrejí?ková P, Oancea F, Frincu RM, Aguzzi A, Ferrari Nicoli S, Casini I, Gabrielli P,

Caproni R, ?erve? I, Lombardi-Boccia G. Valorization potentials of rapeseed meal in a biorefinery perspective: focus on nutritional and bioactive components. Molecules 26:6787 (2021). https://doi.org/10.3390/molecules26226787
9. Dworakowska S, Bogda? D, Bednarz S. Production of biodiesel from rapeseed oil. Proceedings of The 1st World Sustainability Forum : (2011). https://doi.org/10.3390/wsf-00631
10. Wongsirichot P, Gonzalez-Miquel M, Winterburn J. Rapeseed meal biorefining: fractionation, valorization and integration approaches. Biocatalysis and Agricultural Biotechnology 62:103460 (2024). https://doi.org/10.1016/j.bcab.2024.103460
11. Iliuta I, Radfarnia HR, Iliuta MC. Hydrogen production by sorption?enhanced steam glycerol reforming: sorption kinetics and reactor simulation. AIChE Journal 59:2105-2118 (2012). https://doi.org/10.1002/aic.13979
12. Baibhav V, Orrego DF, Bhatnagar P, Mar?chal F. Process integration and waste valorization for sustainable biodiesel production toward a transportation sector energy transition. Systems and Control Transactions 4:869-875 (2025). https://doi.org/10.69997/sct.147707
13. Cheng CK, Foo SY, Adesina AA. Glycerol steam reforming over bimetallic co?ni/al₂o₃. Ind. Eng. Chem. Res. 49:10804-10817 (2010). https://doi.org/10.1021/ie100462t
14. Martinez-Hernandez E, Ramírez-Verduzco LF, Amezcua-Allieri MA, Aburto J. Process simulation and techno-economic analysis of bio-jet fuel and green diesel production — minimum selling prices. Chemical Engineering Research and Design 146:60-70 (2019). https://doi.org/10.1016/j.cherd.2019.03.042
15. Granjo JFO, Duarte BPM, Oliveira NMC. Integrated production of biodiesel in a soybean biorefinery: modeling, simulation and economical assessment. Energy 129:273-291 (2017). https://doi.org/10.1016/j.energy.2017.03.167
16. Carré P. Recovery of low temperature heat in oil mills. OCL 19:324-331 (2012). https://doi.org/10.1051/ocl.2012.0482
17. Suchocki T. Energy utilization of rapeseed biomass in europe: a review of current and innovative applications. Energies 17:6177 (2024). https://doi.org/10.3390/en17236177
18. Anderson International Corp, 2023 https://www.andersonintl.com/understanding-hexane-extraction-of-vegetable-oils/
19. Flórez-Orrego D, Ribeiro Domingos M., Maréchal F. (Sept 2022). A systematic framework for the multi-time integration of industrial complexes and urban systems, CPOTE 2022, Warsaw, Poland
20. Dardor D, Flórez-Orrego D, Terrier C, Ribeiro Domingos ME, Platteau C, da Silva JC, Lopez M, Maréchal F. ROSMOSE: a web-based decision support tool for the design and optimization of industrial and urban energy systems. Energy 304:132182 (2024). https://doi.org/10.1016/j.energy.2024.132182
21. Flórez-Orrego D, Maréchal F, de Oliveira Junior S. Comparative exergy and economic assessment of fossil and biomass-based routes for ammonia production. Energy Conversion and Management 194:22-36 (2019). https://doi.org/10.1016/j.enconman.2019.04.072
22. Angili TS, Grzesik K. Unlocking the circular economy potential: techno-economic analysis of rapeseed meal valorisation through pyrolysis. Journal of Water and Land Development :122-132 (2025). https://doi.org/10.24425/jwld.2025.154256
23. Jerzak W, Gao N, Kalemba-Rec I, Magdziarz A. Catalytic intermediate pyrolysis of post-extraction rapeseed meal by reusing ZSM-5 and zeolite y catalysts. Catalysis Today 404:63-77 (2022). https://doi.org/10.1016/j.cattod.2021.10.023
24. Khavari F, Choi Y. Profit-driven scheduling of CO2 storage and reuse in salt caverns for solar-to-x applications: a two-stage optimization approach. Solar Energy 300:113855 (2025). https://doi.org/10.1016/j.solener.2025.113855

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Modeling standardized industrial profiles for the optimization of eco-industrial parks

Dilafruz Mavlanova[a], Marianne Boix[a*], Rachid Ouaret[a], and Stephane Negny[a]

[a] Laboratoire de Génie Chimique, UMR 5503, Toulouse INP, CNRS, Université Toulouse 3, 4 Allée Emile Monso, 31432 Toulouse Cedex 4, France
* Corresponding Author: Marianne.boix@ensiacet.fr

ABSTRACT

The ecological transition demands innovative frameworks to reduce industrial resource consumption and environmental impacts. Industrial ecology, particularly through Eco-Industrial Parks (EIPs), provides a promising pathway by enabling exchanges of materials, energy, and water between firms. However, the deployment of EIPs is limited by the lack of standardized industrial profiles and transferable modeling approaches. This study develops a generic framework for representing industrial actors as standardized input–output black-box models, consolidating data on resource consumption, energy demand, by-products, and waste streams. These profiles are structured into a harmonized database to support resource-exchange analysis and scalable optimization across diverse contexts. Complementary mappings of processes and resources, as well as energy and heat demand profiles, enhance the feasibility of identifying synergies such as heat cascading and material reuse. The framework is designed to integrate with optimization approaches to evaluate trade-offs between economic and environmental objectives. The framework highlights the benefits of generalizing industrial profiles for robust and transferable EIP modeling, while also identifying gaps in the integration of social indicators, uncertainty analysis, and governance aspects. This work contributes a methodological foundation for supporting circular, low-carbon industrial systems and provides practical tools for the design and deployment of sustainable industrial zones.

Keywords: industrial ecology, eco-industrial park, resource exchanges, optimization

INTRODUCTION

The ecological transition requires innovative strategies to reduce resource consumption and mitigate industrial environmental impacts. Industrial ecology, and particularly Eco-Industrial Parks (EIPs), has emerged as a promising framework for promoting circular resource use through exchanges of energy, water, and materials between companies. While several studies have demonstrated the economic feasibility of industrial symbiosis, there remains a lack of integration of environmental performance and model flexibility across various industrial contexts.

Industrial ecology (IE) provides the theoretical foundation for rethinking industrial systems as ecosystems in which resources circulate rather than follow a linear path. Chertow (2000) emphasized that the central principle of IE is the closing of loops that allows transforming waste from one process into a valuable input for another. This idea is operationalized in EIPs, where firms engage in industrial symbiosis through exchanges of materials, energy, water, and by-products to achieve both environmental and economic benefits [1].

The Kalundborg symbiosis in Denmark is the most widely cited and pioneering example of this approach. Over several decades, companies in Kalundborg developed a network of exchanges, such as the transfer of steam, surplus heat, gypsum, and wastewaterthat reduced costs, minimized environmental impacts, and demonstrated the feasibility of symbiosis in practice. This case illustrates how industrial systems can evolve organically towards resource efficiency and collaborative sustainability.

Recent works extend EIP thinking toward Hubs for Circularity (H4C), which integrate industrial, urban, and rural exchanges to reach near-zero emissions and enhance local welfare. This broader framing highlights the importance of considering not only intra-industrial

synergies but also urban heat recovery, agricultural waste valorization, and community-level benefits [2].

In Europe, Italy has pioneered regional EIP initiatives, particularly in Emilia-Romagna, where governance and institutional support are critical drivers of symbiotic networks. Similarly, in Spain, co-located industrial symbiosis networks have contributed to multiple UN Sustainable Development Goals (SDGs), including water savings, emission reductions, and job creation [3].

Despite successful examples, EIP design and operation are complex due to conflicting objectives, local versus global optima, and operational uncertainties. Kuznetsova and Zio propose a methodological framework treating EIPs as complex engineered systems comparable to smart grids, requiring multi-level optimization and robustness to uncertainties [4]. Mathematical programming, particularly mixed-integer linear programming (MILP), is widely applied to model optimal material and energy exchanges, balancing costs, emissions, and resource use.

Recent advances include:

- Multi-objective MILP models optimizing economic, environmental, and social performance in desert EIPs, such as the Ulan Buh Demonstration Park in Inner Mongolia.
- Stochastic and robust optimization frameworks to account for uncertainty in demand, waste availability, and energy prices within IS networks and H4Cs.
- Leader–slave optimization structures that reconcile global EIP objectives with local firm interests.

Such formulations increasingly integrate multi-resource networks (water, energy, materials, gases), though joint modeling across all domains remains limited. The literature demonstrates the evolution from classical EIP design to complex, multi-objective, and multi-scale frameworks. Key insights include: EIPs can deliver significant environmental and economic benefits but require systematic modeling under uncertainty. Simulation and visualization methods bridge the gap between technical models and stakeholder decision-making.

Despite progress, gaps remain in integrated multi-resource modeling, incorporation of social objectives (e.g., job creation, equity), and operational-level optimization beyond strategic design. This work builds upon these foundations by developing a generic framework for modeling network components in industrial resource exchanges, integrating black-box process representations with systemic indicators (e.g., MFA, Sankey diagrams) to improve adaptability across contexts.

PROBLEM STATEMENT

The primary objective of this study is to support the deployment of industrial ecology by developing standardized industrial profiles that can be used as building blocks in a generic optimization framework for eco-industrial parks. The work involves reviewing existing literature and case studies, defining representative profiles for selected industrial sectors, and characterizing their typical processes in terms of resource consumption, energy demand, by-products, and waste streams.

These profiles are modeled using a black-box approach to identify potential resource exchanges and synergies between companies. Finally, the profiles are integrated into an adaptable optimization framework and applied to a case study, with the aim of evaluating both economic and environmental performances.

The purpose of this formulation is not to provide a site-specific optimal solution, but to demonstrate how standardized profiles reduce modeling effort and enable scalable optimization across heterogeneous industrial contexts. This research work aims to answer the following questions:

- What are the defining features of industrial profiles in terms of resource usage?
- How can these profiles be modeled in a way that supports generic optimization?
- What is the impact of using standardized profiles on the scalability and adaptability of eco-industrial models?

METHODOLOGY

The methodology outlines the approach adopted to develop a generalized framework for modeling resource exchanges within eco-industrial parks. Within the scope of the French industry decarbonization program, this research addresses the need for a standardized database capable of integrating diverse industrial profiles across the country. The procedure involved structuring representative profiles into a harmonized format and designing a network model to support resource exchange analysis. This approach provides the foundation for evaluating eco-industrial park topologies in a way that is both generalizable and adaptable to local contexts.

This study is defined as quantitative and model-based, drawing on the frameworks of industrial ecology and optimization modeling. Industrial ecology provides the systemic perspective for analyzing material, energy, and water exchanges within industrial ecosystems [5], [7]. Within this methodological framework, EIPs operationalize industrial symbiosis by linking firms through complementary resource flows and assessing their performance with eco-efficiency indicators [8].

Case studies such as the Ulsan Eco-Industrial Park

Table 1. Representative enterprise types and their primary input–output flows.

Category	Facility / Plant	Inputs	Outputs
Energy Utility	Combined Heat & Power (CHP) Plant	Fuel, Water	Electricity, Steam, Waste Heat
	Biomass Power Plant	Biomass	Electricity, Heat, Biochar
	Waste-to-Energy Plant	Municipal Waste	Electricity, Heat, Ash
	Hydrogen Production Plant	Electricity, Water	Hydrogen, Oxygen
Wastewater TP	Municipal Wastewater Treatment Plant	Wastewater, Electricity	Treated Water, Biogas, Sludge
Heavy Industry	Cement Factory	Limestone, Slag, Energy	Cement, Clinker, Heat
	Steel Manufacturing Plant	Iron Ore, Coal, Water	Steel, Slag, Waste Heat
	Glass Manufacturing Plant	Silica, Soda Ash, Energy	Glass Products, Heat
	Brick or Tile Factory	Clay, Ash	Bricks, Tiles
Chemical Industry	Chemical Processing Plant	Feedstocks, Steam, Water	Chemicals, Heat, Wastewater
	Plastic Recycling Facility	Plastic Waste, Energy	Recycled Plastic
	Paint or Coating Factory	Solvents, Chemicals	Paint, Coating
	Fertilizer Plant	Ammonia, CO_2	Fertilizers, CO_2 Reuse
Agri-Food	Food Processing Plant	Raw Food, Water	Packaged Food, Organic Waste
	Brewery or Winery	Grains, Water	Beer/Wine, Spent Grains
	Slaughterhouse	Livestock	Animal Products, Fat Waste
	Dairy Plant	Milk, Steam, Water	Milk Products, Whey
Light Manufacturing	Textile Factory	Cotton, Dyes, Water	Fabrics, Dye Wastewater
	Paper Mill	Pulp, Water	Paper, Sludge
	Wood Workshop	Wood	Sawdust, Wood Products
	Electronics Assembly Plant	Electronic Components	Electronics, E-waste
Other industries	Urban Greenhouse / Vertical Farm	CO_2, Heat, Water	Vegetables, CO_2 Utilization
	Composting Facility	Food Waste, Green Waste	Compost
	Algae Production Facility	Wastewater, CO_2	Algae, Biomass
	Biogas Plant (Anaerobic Digester)	Organic Waste	Biogas, Digestate
Small Business	Bakery	Flour, Water, Electricity	Bread, Waste Heat
	Laundromat	Water, Electricity	Hot Water Discharge
	Café / Restaurant	Water, Ingredients, Energy	Food Waste, Wastewater

demonstrate how "designed" industrial symbiosis networks can evolve from research and development initiatives into practical business models, enabling systematic evaluation of exchanges across diverse enterprises [9].

Optimization modeling has been widely applied to EIPs to evaluate synergies and design efficient exchange networks [10]. Methods such as mixed-integer linear programming (MILP) and superstructure optimization enable systematic assessment of trade-offs between economic and environmental objectives, while metaheuristic approaches such as ant colony optimization (ACO) and agent-based modeling (ABM) have been proposed to handle the complexity of large-scale networks.

The objectives of this research are therefore fourfold:

- To develop generalized industrial profiles that abstract enterprises into standardized input–output black-box models;
- To model resource exchanges in EIPs through network superstructures that capture all feasible synergies;
- To design a harmonized dataset to support optimization across diverse industrial contexts;
- To evaluate the adaptability of this framework through application to a case study.

MAPPING OF INDUSTRIAL PROFILES

Industrial profiles and enterprises types

A substantial body of literature exists on material flow analysis (MFA) [1], circular economy implementation, and the optimization of industrial symbiosis [6]. These studies demonstrate the potential of resource exchanges within eco-industrial parks (EIPs) and provide valuable methodologies for flow quantification and network optimization. However, most of these works remain

Table 2. Process to Resources mapping

Enterprise	Process	Resources Required
Steel Manufacturing	Blast Furnace	Iron ore, Coal, Coke, Oxygen, Heat
	Basic Oxygen Furnace	Oxygen, Scrap Metal, Heat
	Rolling	Electricity, Lubricants
	Waste Heat Recovery	Waste Gas, Heat Exchanger
	Slag Handling	Electricity, Water
Cement Factory	Raw Material Crushing	Limestone, Electricity
	Preheating	Waste Heat, Fuel
	Clinker Production	Fuel, Heat, Electricity
	Grinding & Packing	Electricity, Gypsum
Food Processing	Cleaning	Water, Electricity
	Cooking	Steam, Electricity, Ingredients
	Pasteurization	Steam, Water, Electricity
	Packaging	Plastic, Packaging Gas

Table 3. Resources -> Process mapping

Resource Input	Used In Process	Industry Examples
Steam	Sterilization	Food & Beverage Pharmaceutical
	Boiling	Brewery, Food Processing Food processing
	Evaporation	Sugar Refinery Dairy
	Drying	Paper Mill, Textile, Chemical
	Heating reactors	Chemical Manufacturing, Petrochemical
	Steam Stripping	Petrochemical, Refining
	Textile finishing	Textile, Apparel
	Steam cleaning	Dairy, Brewery
	Pasteurization	Dairy, Beverage
	Cooking	Canning, Packaged Food
	Distillation	Refinery, Spirits Production
	Preheating	Steel, Chemical

context-specific, focusing on individual case studies or sectoral analyses, rather than proposing a generalized and reusable dataset of industrial profiles.

Currently, there is no consolidated datasheet that categorizes enterprises into standardized groups with generalized input/output flow data. This gap makes it difficult to design transferable models that can efficiently simulate exchanges between companies in different locations. Existing datasets are often fragmented (case-specific "stories" of companies or sectoral MFA studies), but they are not harmonized into a universal resource-exchange typology.

Industries in this study were first classified into representative categories (e.g., power plants, chemical industries, steel plants, bakeries, wastewater treatment facilities). The Enterprise Types sheet of the database defines the scope of participants, their primary inputs (fuels, raw materials, water), and outputs (electricity, steam, waste heat, by-products). This provides the foundation for a generic library of industrial actors in an EIP.

The Input-Output sheets consolidate resource balances by identifying primary inputs, main outputs, and secondary outputs such as waste heat, CO,ÇÇ, or solid residues. This abstraction ensures that industries with very different processes can be compared under the same modeling structure.

A datasheet (Table 1) has been created for each industry profile, providing a general view and supporting

Table 4. Energy and heat demand profiles of industries

Enterprise Name	Process	Electricity Demand (Typical)	Heat/Steam Demand (Typical)	Estimated Temperature Range (°C)
Steel Manufacturing Plant	Raw material handling	Very High	High	900–1800
	Blast furnace	Very High	High	900–1800
	BOF	Very High	High	900–1800
	Rolling & finishing	Very High	High	900–1800
	Waste heat recovery	Very High	High	900–1800
Cement Factory	Crushing	High	Very High	1200–1500
	Preheating	High	Very High	1200–1500
	Clinker formation	High	Very High	1200–1500
	Grinding	High	Very High	1200–1500
	Packing	High	Very High	1200–1500
Chemical Processing Plant	Reactor operation	High	High	100–300
	Separation	High	High	100–300
	Distillation	High	High	100–300
	Storage	High	High	100–300
	Cooling	High	High	100–300

the construction of a generic model. The sheets are structured around the following characteristics:

- Enterprise type (energy, chemical industry, heavy industry, etc.);
- Primary and secondary inputs and outputs of enterprises.

Process-resource mapping

Two complementary sheets: Process→Resources and Resources→Process—map industrial processes to the resources they require and produce (Table 2, 3). For example, in steelmaking, the blast furnace consumes iron ore, coke, and oxygen, while emitting CO_2 and slags. This mapping ensures that exchanges can be matched based on technical feasibility.

ENERGY AND HEAT PROFILES

To systematically evaluate exchange opportunities, the energy and heat requirements of industrial processes are characterized in terms of both demand level and temperature range. The Electricity–Heat Demand profile provides a comparative mapping of industries such as steel manufacturing, power generation, and food processing, indicating whether their electricity and heat consumption is low, medium, or very high.

In addition, typical operating temperature ranges are included to capture the quality of heat required or released. This parameterization is crucial for identifying feasible synergies: for instance, excess low-temperature steam from one process may be unsuitable for another process requiring high-grade heat, but could be matched to district heating or drying applications.

By grouping industries according to their energy demand intensity and thermal characteristics, the framework facilitates heat cascading strategies and improves the robustness of optimization models for eco-industrial parks.

OPTIMIZATION AND MODELLING STRATEGIES

The standardized industrial profiles developed in this study are designed to be adaptable to a range of optimization and modeling approaches. This creates opportunities to strengthen the design and deployment of Eco-Industrial Parks (EIPs) by linking resource-exchange data to advanced decision-support tools.

Mathematical optimization methods, such as Mixed-Integer Linear Programming (MILP), are particularly well-suited for systematically identifying exchanges that minimize cost or emissions while respecting process constraints. By structuring industries as black-box profiles, the framework ensures compatibility with such models and facilitates scalable, scenario-based evaluation.

Based on previous works [6] and to demonstrate the compatibility of the proposed standardized industrial profiles with process-systems engineering approaches, this section outlines a generic optimization formulation for eco-industrial park (EIP) design based on these

profiles. The objective is not to solve a fully instantiated case-specific problem, but to illustrate how the standardized profiles can be directly embedded into a superstructure-based optimization framework.

Each industrial actor $i \in I$ is represented as a standardized black-box profile defined by its resource inputs, outputs, by-products, and energy demands. These profiles constitute the building blocks of a superstructure that enumerates all feasible resource exchanges between actors within the EIP.

Let R denote the set of resources (e.g. water, steam, heat, materials), and x_{ij}^r the continuous decision variable representing the flow of resource $r \in R$ exchanged from actor i to actor j. Binary variables y_{ij}^r indicate whether an exchange link is activated.

A generic cost–emission optimization problem can be formulated as:

$$\min \left(\sum_{i,j,r} C_{ij}^r x_{ij}^r \, , \; \sum_{i,j,r} EF^r x_{ij}^r \right)$$

subject to:

- Mass and energy balance constraints

$$\sum_j x_{ji}^r + S_i^r = \sum_j x_{ij}^r + D_i^r \forall i, r$$

where S_i^r and D_i^r denote standardized supply and demand parameters derived from the industrial profiles.

- Capacity constraints

$$0 \le x_{ij}^r \le \bar{X}_{ij}^r y_{ij}^r$$

- Technical compatibility constraints

$$x_{ij}^r = 0$$

$$\text{if profiles } i \text{ and } j \text{ are incompatible for resource } r$$

The standardized profiles provide harmonized bounds $\bar{X}_{ij}^r$, emission factors EF^r, and unit exchange costs C_{ij}^r, enabling the formulation to be instantiated across different territorial contexts without redefining the model structure.

This formulation illustrates how standardized profiles reduce model-building effort, facilitate scalability, and support comparative scenario analysis, while remaining fully compatible with classical MILP and multi-objective optimization approaches used in EIP design.

The standardized industrial profiles can be applied to generate and analyze generic resource-exchange scenarios representative of typical eco-industrial park configurations. Rather than focusing on a single site-specific case, the objective is to demonstrate how the profile-based approach supports systematic identification of symbiosis opportunities.

Based on the compiled profiles, industries are currently classified according to their dominant resource roles (producer, consumer, or prosumer) and thermal characteristics. This enabled the identification of recurrent exchange patterns, including:

- heat and steam cascading between high-temperature and low-temperature processes,
- by-product valorization chains across manufacturing and recycling activities,
- water reuse loops involving industrial and utility actors.

For instance, standardized profiles for energy-intensive processes (e.g. power generation, steelmaking) are matched with medium- and low-temperature demand profiles (e.g. food processing, district heating), revealing generic heat recovery configurations applicable across multiple industrial zones. Similarly, material by-products identified in standardized profiles enabled the construction of cascading exchange chains connecting primary production, transformation, and recycling actors.

These scenarios can be directly encoded as candidate exchange links in the optimization superstructure, reducing the search space while preserving technical realism. Compared to purely case-specific modeling, the standardized profile approach enables rapid scenario generation, cross-site comparison, and transferability of symbiosis configurations.

Looking forward, the main opportunity lies in combining these approaches: using standardized profiles as building blocks for multi-objective optimization, heuristic exploration, and dynamic modeling. This integration will not only improve technical robustness but also support decision-making by balancing economic, environmental, and social objectives. While the present framework relies on static categorizations of energy and heat demand (e.g., low, medium, high), this simplification corresponds to an aggregated representation aimed at enabling generic and transferable modeling across industrial sectors. However, real industrial systems often exhibit significant temporal variability due to batch operations, shift-based production, or seasonal fluctuations.

The proposed formulation can be naturally extended to account for such dynamics by introducing a multi-period MILP structure. In this case, the set of decision variables and parameters (e.g., resource flows, demands, and availabilities) can be indexed over discrete time intervals (e.g., hourly, daily, or seasonal periods). This allows capturing temporal mismatches between resource supply and demand, such as excess heat production during off-peak hours and consumption during peak periods.

Additional components such as storage units (e.g., thermal storage, tanks) and time-dependent operational constraints can also be integrated to ensure feasibility of exchanges over time. Importantly, the standardized industrial profiles developed in this work remain fully

compatible with this extension, as their input–output structure can be enriched with time-dependent parameters without modifying the underlying modeling framework.

This perspective opens the way toward more realistic and operational eco-industrial park designs, combining the scalability of standardized profiles with the accuracy of dynamic optimization.

CONCLUSION

This study proposes a process-based framework to support the design and optimization of Eco-Industrial Parks (EIPs) through the development of generalized industrial profiles. By abstracting enterprises into standardized black-box models defined by their inputs and outputs, the framework makes it possible to compare diverse sectors under a common structure and to systematically identify potential exchanges of materials, energy, and water.

The results highlight that structuring industrial data in this way improves both the robustness and scalability of optimization models. It also enhances transferability across different contexts, making the framework suitable for regional and national applications in the deployment of industrial ecology. Furthermore, the use of systemic indicators such as Sankey diagrams and material flow analysis strengthens the capacity to communicate technical results to stakeholders, supporting collaborative decision-making.

While the present work has focused primarily on economic and environmental objectives, future developments should integrate social and governance dimensions, life cycle assessment (LCA), and explicit modeling of uncertainty. In combination with advanced optimization and simulation strategies, the proposed framework can thus serve as a practical tool for advancing circular, low-carbon industrial systems and supporting territorial decarbonization policies.

ACKNOWLEDGEMENTS

This work was supported by a government grant managed by the Agence Nationale de la Recherche under the France 2030 program (PEPR SPLEEN), under the reference ANR-22-PESP-0004.

AUTHOR IDENTIFIERS

Author ORCIDs:
Marianne Boix: 000-0003-0993-3395
Rachid Ouaret: 0000-0002-0502-2108
Stéphane Negny: 0000-0001-7533-1857

REFERENCES

1. Chertow MR. Industrial symbiosis: literature and taxonomy. Annu. Rev. Energy. Environ. 25:313-337 (2000). https://doi.org/10.1146/annurev.energy.25.1.313
2. HBrunner P, Rechberger H., 2004. Practical Handbook of Material Flow Analysis. The International Journal of Life cycle assessment. 9, 5. DOI:10.1007/BF02979426
3. Neves A, Godina R, Azevedo SG, Matias JCO. A comprehensive review of industrial symbiosis. Journal of Cleaner Production 247:119113 (2020). https://doi.org/10.1016/j.jclepro.2019.119113
4. Kuznetsova E, Zio E, Farel R. A methodological framework for eco-industrial park design and optimization. Journal of Cleaner Production 126:308-324 (2016). https://doi.org/10.1016/j.jclepro.2016.03.025
5. Yu H, Da L, Li Y, Chen Y, Geng Q, Jia Z, Zhang Y, Li J, Gao C. Industrial symbiosis promoting material exchanges in ulan buh demonstration eco-industrial park: a multi-objective MILP model. Journal of Cleaner Production 414:137578 (2023). https://doi.org/10.1016/j.jclepro.2023.137578
6. Boix M, Montastruc L, Pibouleau L, Azzaro-Pantel C, Domenech S. Industrial water management by multiobjective optimization: from individual to collective solution through eco-industrial parks. Journal of Cleaner Production 22:85-97 (2012). https://doi.org/10.1016/j.jclepro.2011.09.011
7. Ghali MR, Frayret JM, Ahabchane C. Agent-based model of self-organized industrial symbiosis. Journal of Cleaner Production 161:452-465 (2017). https://doi.org/10.1016/j.jclepro.2017.05.128
8. Lange KPH, Korevaar G, Oskam IF, Nikolic I, Herder PM. Agent-based modelling and simulation for circular business model experimentation. Resources, Conservation & Recycling Advances 12:200055 (2021). https://doi.org/10.1016/j.rcradv.2021.200055
9. Saghafi Z, Roshandel R. Agent-based simulation for technology implementation in an energy-based industrial symbiosis network. Resources, Conservation & Recycling Advances 21:200201 (2024). https://doi.org/10.1016/j.rcradv.2023.200201
10. UNIDO Eco-Industrial Parks - Tools.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Understanding Environmental Impacts of Lithium-Ion Battery Recycling

Marco Vaccari*, Leonardo Tognotti, Monica Puccini

University of Pisa, Department of Civil and Industrial Engineering, Largo Lucio Lazzarino 2, 56126 Pisa (Italy)
* Corresponding Author: marco.vaccari@unipi.it

ABSTRACT

The increasing deployment of lithium-ion batteries (LIBs) requires effective recycling strategies to reduce environmental impacts and dependence on critical raw materials. In this study, a comparative life cycle assessment (LCA) of two LIB recycling routes, a pyrometallurgical process (Pyro) and a hydrometallurgical process with co-precipitation (Hydro), was performed using a Python-based process modeling framework. The LCA was carried out using an attributional approach, with impacts referred to 1 kg of spent LIBs treated at the recycling facility inlet, considering a representative mix of battery formats and cathode chemistries. Results showed that, when normalized per kilogram of treated batteries, the Hydro route is more impactful than the Pyro one, particularly in terms of global warming potential. The Pyro process does not enable direct cathode regeneration but allows the recovery of high-purity metal salts, whereas the Hydro route enables the production of re-formed NMC-111 cathode material that exhibits lower environmental impacts than virgin cathode production. Overall, the study highlights how rigorous process modeling integrated with LCA enables a detailed and non-trivial assessment of LIB recycling sustainability, revealing trade-offs that are not evident from simplified or purely qualitative analyses.

Keywords: Life Cycle Analysis, Modelling and Simulations, Batteries, Process Operations, SimaPro

INTRODUCTION

The rapid growth of electric mobility has led to a sharp increase in the deployment of lithium-ion batteries (LIBs). As a consequence, the management of spent LIBs has become a critical environmental and resource-security issue [1]. Spent LIBs contain significant amounts of valuable and critical raw materials (CRMs), including lithium (Li), cobalt (Co) and nickel (Ni), whose extraction is associated with high environmental burdens and supply risks. Recycling is expected to play a central role in mitigating these challenges by reducing reliance on primary resources and supporting circular economy strategies along the battery value chain [2]. However, the environmental performance of recycling technologies strongly depends on process design and operational conditions.

Among the recycling technologies proposed for LIBs, pyrometallurgical and hydrometallurgical routes represent the most mature industrial options [3]. Pyrometallurgical routes are characterized by high robustness and flexibility with respect to battery chemistry and format but typically exhibit high energy demand and limited selectivity toward Li recovery [4]. Hydrometallurgical routes, in contrast, rely on leaching and selective separation steps enabling higher recovery efficiencies for CRMs, at the cost of increased process complexity and reagent consumption [5].

Life Cycle Assessment (LCA) has therefore been widely adopted as a decision-support tool to evaluate the environmental implications of LIB recycling technologies. Several studies have shown that recycling does not automatically result in lower environmental impacts compared to primary material production, and that LCA outcomes are highly sensitive to methodological choices such as system boundaries and functional unit definition [6]. Comparisons based on the mass of treated battery waste may lead to different conclusions than those based on the quantity or quality of recovered materials, highlighting the importance of transparent and consistent assumptions [7].

In recent years, process modeling has emerged as a key element for improving the robustness of LCA studies

applied to battery recycling systems. Detailed mass and energy balances derived from process simulations enable the construction of consistent inventories and allow environmental results to be interpreted in direct relation to specific unit operations and material inputs [8]. This approach is particularly relevant for complex and evolving recycling technologies, where industrial-scale data are often limited and scenario-based analyses are required to explore potential optimization pathways [9].

Within this framework, the present study provides a comparative LCA of two LIB recycling routes: a pyrometallurgical process (Pyro) and a hydrometallurgical process with co-precipitation (Hydro) based on a unified process modeling approach. Material and energy balances obtained from a Python-based simulation model were implemented in SimaPro to assess the environmental impacts associated with each recycling strategy. By integrating process-level modeling and LCA, the study aims to clarify the main environmental trade-offs between the two technologies and to identify the process features that most strongly influence their sustainability performance.

RECYCLING ROUTES AND PROCESS MODELING BACKGROUND

The two alternative recycling routes for LIBs considered in this study, Pyro and Hydro, were modeled in detail to quantify material and energy flows and to identify the main process features influencing their environmental performance [9]. A conceptual overview of the process workflow adopted in this work is provided in Figure 1.

Pyrometallurgical recycling route (Pyro)

The Pyro route is based on a high-temperature smelting process, in which spent LIBs are treated in a furnace to recover a metallic alloy mainly composed of Ni, Co and copper. During smelting, organic components and graphite are oxidized and provide part of the thermal input, while Li, aluminum and other elements are largely transferred to the slag phase, making them not directly recovered as a valuable product in this route. The metallic alloy produced in the smelter is subsequently processed through a series of hydrometallurgical refining steps aimed at separating and purifying individual metal salts. These downstream refining steps allow the recovery of Ni and Co in the form of high-purity metal salts, which represent the main valuable outputs of the Pyro route. These operations involve leaching, precipitation and washing stages, requiring the use of alkaline reagents, in particular sodium hydroxide (NaOH), as well as additional water and energy inputs. Together with the high thermal energy demand of the smelting stage, these inputs contribute significantly to the overall material and energy balances of the Pyro route. The main advantage of this recycling strategy lay in its robustness and flexibility with respect to battery chemistry and format, as well as in its relatively simple process structure. However, this comes at the expense of high energy consumption and limited selectivity for Li recovery [3].

Hydrometallurgical recycling route with co-precipitation (Hydro)

The Hydro route is instead based on a sequence of mechanical treatments, leaching and selective separation steps. After mechanical preprocessing, the black mass is subjected to acidic leaching to solubilize valuable metals. Impurities were removed through controlled precipitation, followed by a co-precipitation stage aimed at recovering a mixed metal precursor suitable for the synthesis of NMC-111 cathode material. In this route, the main valuable output is a mixed-metal precursor whose composition is

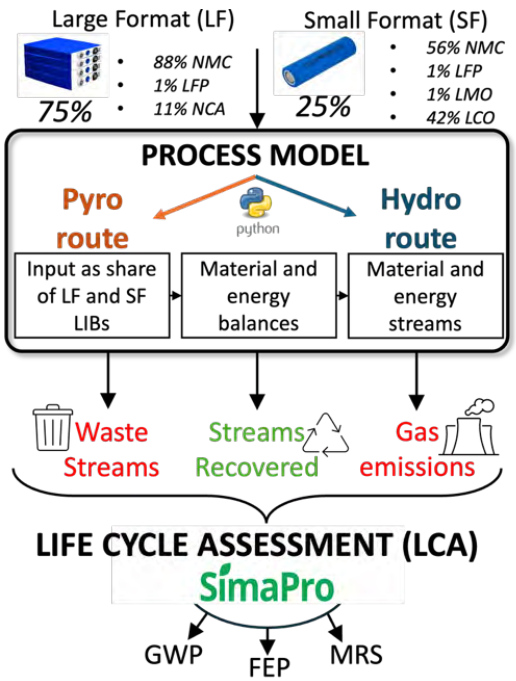

Figure 1. Conceptual framework for the LCA of LIB recycling: from Python-based LCI to results.

adjusted to match the target NMC-111 stoichiometry, requiring the controlled addition of Ni, Co and Mn salts during co-precipitation. Compared to the Pyro route, the Hydro process involved a higher number of unit operations and a more complex flowsheet. In addition to inorganic reagents, the process required the use of organic solvents, such as acetonitrile (ACN), in specific separation steps. These solvents are partially recovered and recycled, but their use contributes to auxiliary material consumption and to the generation of solvent-containing waste streams. The Hydro route allows a higher degree of selectivity in the recovery of CRMs, including Li. At the

same time, the increased complexity of the process results in additional material, water and energy inputs that are explicitly accounted for in the process model [3].

Process modeling framework

Both recycling routes are simulated using a Python-based process modeling framework previously developed and described in detail in [9]. The model represents the main unit operations of each recycling route and calculates the corresponding material and energy balances under steady-state conditions. Model parameters are derived from literature data and industrial references, and the resulting balances are used as primary data for the subsequent LCA. To facilitate the integration of the process simulations into the LCA framework, individual unit operations were aggregated into process blocks representing the main sections of each recycling route. This aggregation allowed a transparent implementation of the recycling processes in SimaPro while preserving the link between process modeling assumptions and environmental results. The resulting system boundaries and block structure adopted for the LCA are illustrated in Figure 2.

LIFE CYCLE ASSESSMENT METHODOLOGY

Goal and scope definition

The LCA is performed to compare the environmental performance of the two LIB recycling routes, consistently with the process modeling framework developed in previous work. As evidenced in Figure 2, the analysis focuses exclusively on the recycling stage, with the aim of identifying and comparing the environmental burdens associated with the two alternative technological solutions. The functional unit defined is 1 kg of spent LIBs at the inlet of the recycling facility. This choice is deliberately adopted to ensure a consistent comparison between the two recycling routes, which differ significantly in terms of process configuration, material selectivity and recovery efficiency, and overall output products. Normalizing all results to the same amount of input battery waste allows the environmental impacts of the two processes to be directly compared without introducing additional allocation assumptions. A complementary comparison based on the mass of recovered NMC-111 cathode material was also carried out as a supporting analysis for the Hydro route only. In this case, the environmental impacts associated with the recycled cathode material are compared with those of a reference NMC-111 production process available in the SimaPro database. This comparison is not intended to define an alternative functional unit, but rather to provide additional context for the interpretation of the results related to material recovery.

System boundaries and life cycle inventory

The system boundaries of the LCA include the recycling processes from the reception of spent LIBs at the recycling facility to the generation of recovered material streams and waste emissions. Battery use, collection and transport to the recycling facility are excluded from the analysis. The system boundaries and the aggregation of unit operations into process blocks adopted in the LCA model are illustrated in Figure 2. Life cycle inventory (LCI) data are derived from a Python-based process simulation model, which provided detailed material and energy balances for both recycling routes [9]. The process model quantified the associated input and output streams and the resulting mass and energy balances are subsequently implemented in SimaPro to build the LCI. An attributional approach was adopted for the LCA, using the Ecoinvent 3 database (cut-off by classification) as background data source. Recycling facilities are assumed to be located in Europe, in line with the industrial context considered in the process simulations. Consequently, European electricity mixes and region-specific energy carriers were used to model energy inputs. This assumption is particularly relevant for energy-intensive operations and has a significant influence on the overall environmental performance of the recycling routes.

Life cycle impact assessment and impact categories

Life cycle impact assessment (LCIA) is carried out using the ReCiPe 2016 Midpoint (Hierarchist) method implemented in SimaPro. The method allows the transformation of inventory flows into a set of midpoint environmental impact indicators, providing a comprehensive assessment of the potential environmental burdens associated with the recycling processes. Although the full set of midpoint impact categories available in the selected method was calculated, only a subset of impact categories is reported and discussed in detail for clarity and relevance. The selection of the impact categories reported in Figure 3 reflects an interpretative and application-oriented choice, rather than compliance with product category rules or environmental product declaration criteria. In particular, Global Warming Potential (GWP) is included due to its direct relevance to European climate policies and regulations [2]. Mineral Resource Scarcity (MRS) is selected to reflect the primary objective of LIB recycling, namely the recovery of CRMs and the reduction of virgin resource extraction. Freshwater Eutrophication Potential (FEP) is considered to account for water-related impacts, which are particularly relevant for hydrometallurgical processes due to reagent use and wastewater generation. Finally, Ozone Depletion Potential (ODP) is included as an indicator of environmentally harmful emissions not strongly dependent on local conditions, in contrast to several ecotoxicity-related categories. This selection

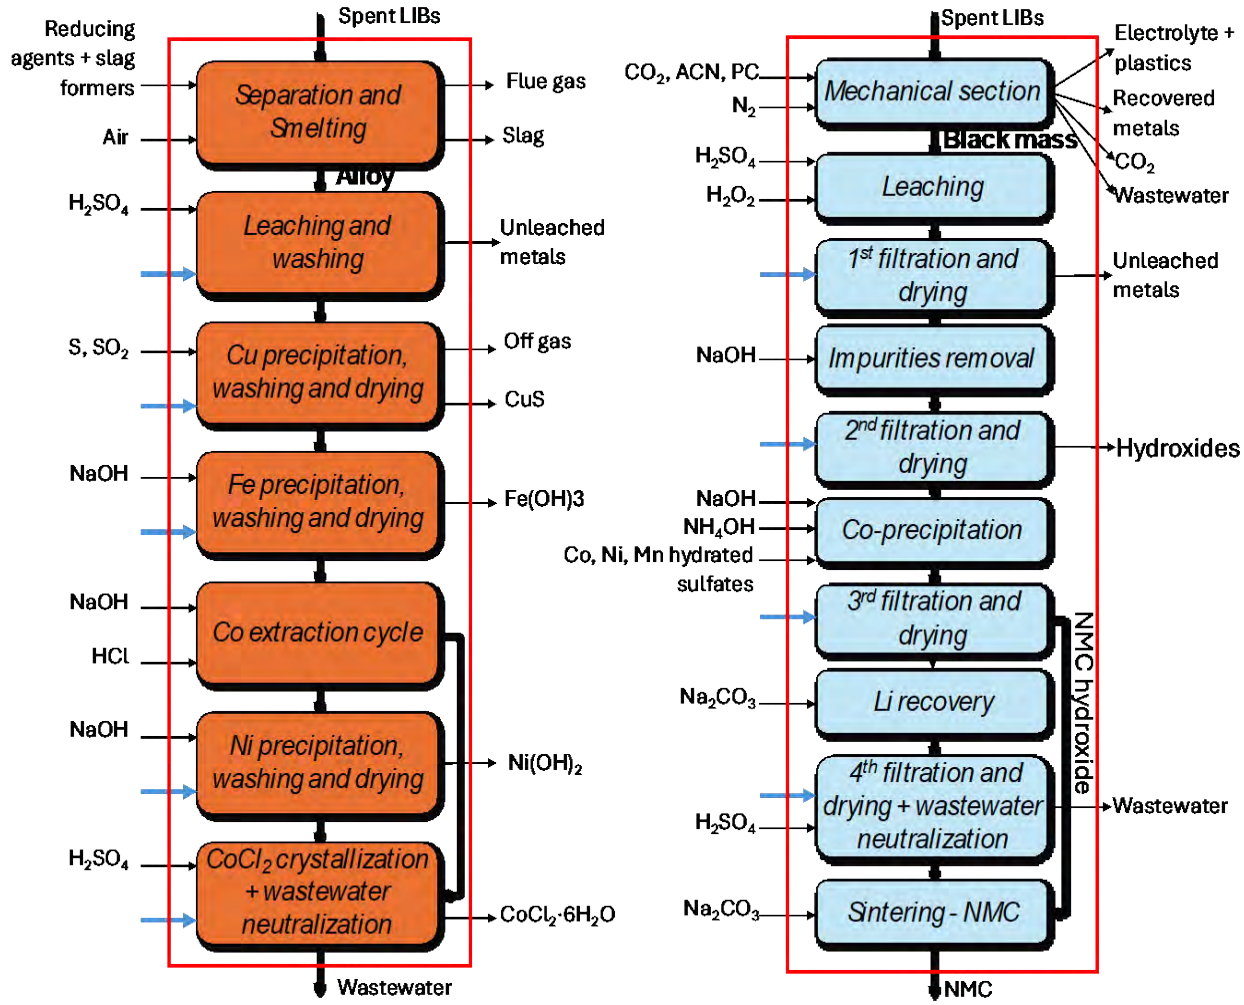

Figure 2. System boundaries for LCA analysis of the Pyro route (on the left) and of the Hydro route (on the right) highlighting the main material streams entering and exiting the blocks. The sequences of operations is taken from [9] and simplified/grouped according to SimaPro implementation. The blue arrows represent inlet washing water, while energy streams and transport input are not represented even if included in the analysis

enabled a clear and physically meaningful comparison between the two routes while avoiding redundancy and potential misinterpretation of highly location-dependent categories. This methodological framework enables a consistent interpretation of the contribution analysis and the quantitative comparisons presented in the following section.

RESULTS AND DISCUSSION

LCA results per kg of spent LIBs

Figure 3 reports a normalized contribution analysis of the Pyro and Hydro recycling routes. Impacts are normalized within each process and category to identify the most impactful unit operations within each route. As illustrated in Figure 1, the process feed consisted of a representative mix of LF and SF LIBs, characterized by different cathode chemistries, to reflect realistic industrial recycling conditions and to avoid results biased toward a single battery type. For all the impact categories considered, the two processes exhibited markedly different environmental profiles, reflecting both their technological characteristics and the associated material and energy inputs. For the Pyro route, Figure 3 clearly highlights the smelting stage as the dominant contributor to GWP. This result is consistent with the high-temperature nature of the process and the associated fuel consumption. In addition, the iron precipitation stage represents a relevant contribution in all displayed impact categories, reflecting the intensive use of alkaline reagents, particularly NaOH, required for impurity removal in downstream hydrometallurgical refining steps. Other unit operations

show comparatively minor contributions, indicating that environmental impacts in the Pyro route are strongly concentrated in a limited number of process blocks. In the Hydro route, the contribution analysis reveals a different impact distribution. The mechanical separation section and the co-precipitation stage emerge as the most impactful unit operations across the selected categories. The mechanical separation section is mainly associated with electricity consumption and solvent-related operations, including electrolyte separation based on ACN, which constitutes a major contributor to GWP. In contrast, the co-precipitation stage is dominated by the consumption of chemical reagents necessary to control metal ratios and achieve the desired cathode stoichiometry. Overall, the contribution analysis highlights a clear distinction between the main impact drivers of the two dominant unit operations in the Hydro route. Mechanical separation stages are primarily associated with energy-related contributions, mainly electricity consumption and solvent handling, whereas the co-precipitation stage is driven by material-related inputs linked to reagent consumption. This separation between energy- and material-driven contributions provides a process-level interpretation of the impact distribution observed in Figure 3.

While Figure 3 provides insight into the internal structure of environmental impacts for each recycling route, a quantitative comparison between Pyro and Hydro is reported in Table 1. When impacts are normalized to 1 kg of spent LIBs treated, the Hydro route exhibits a higher GWP than the Pyro route (3.97 vs. 3.61 kg CO_2 eq), mainly due to solvent use, reagent consumption and wastewater treatment requirements. A similar trend is observed for FEP, where the Hydro route shows markedly higher values, reflecting the intensive use of water and aqueous process streams. Consistently, the Pyro route shows MRS values that are approximately two orders of magnitude lower than those of the Hydro route. Although Li and Al are lost to the slag phase during smelting, the pyrometallurgical process recovers Ni and Co as high-purity salts, thereby reducing the demand for virgin CRMs. In contrast, the higher MRS observed for the Hydro route is mainly associated with the intentional use of Co, Ni and Mn sulphates to adjust the stoichiometry of the re-formed NMC-111 cathode, rather than with inefficiencies in metal recovery.

An important outcome of the combined analysis of Figure 3 and Table 1 concerns the potential for process optimization. In the Pyro route, the dominant contributions originate from the smelting stage, which is intrinsically linked to high-temperature operation and therefore difficult to mitigate without major process redesign. In contrast, the Hydro route exhibits impact hotspots associated with specific unit operations and material choices, such as solvent-based separations and reagent-intensive co-precipitation. This suggests that targeted modifications, including alternative solvents, improved solvent recovery or process simplification, could significantly reduce the environmental impacts of the Hydro route without compromising its material recovery performance.

Table 1: Quantitative comparison of Hydro and Pyro for 1 kg of spent LIBs treated

Impact category	Unit	Hydro	Pyro
GWP	kg CO_2eq	3.974	3.611
ODP	kg CFC11eq	1.63E-06	1.47E-06
FEP	kg Peq	1.77E-01	1.23E-03
MRS	kg Cueq	1.69E-01	8.25E-03

Table 2: Quantitative comparison of recycled and virgin NMC-111 cathode production (1 kg basis).

Impact category	Unit	Hydro (recycled NMC-111)	Virgin NMC-111
GWP	kg CO_2eq	20.522	26.116
ODP	kg CFC11eq	8.42E-06	1.96E-05
FEP	kg Peq	9.11E-01	3.595
MRS	kg Cueq	8.74E-01	5.633

Supporting comparison with virgin NMC-111 cathode production

To further contextualize the environmental performance of the Hydro recycling route, a supporting comparison was carried out between the production of recycled NMC-111 cathode material and a virgin NMC-111 cathode production process available in the SimaPro database. This comparison was not intended to introduce a second functional unit, but rather to provide an additional reference for interpreting the potential benefits of cathode material recovery. The virgin NMC-111 production process was selected from the Ecoinvent 3 database (cut-off by classification) and represents an industrial-scale synthesis of NMC-111 cathode material based on the use of primary raw materials. The process includes the upstream extraction and refining of Ni, manganese (Mn) and Co compounds, followed by precursor synthesis and high-temperature calcination to produce the final cathode material. Energy inputs and emissions were modeled using region-specific background processes consistent with the European context adopted in this study. Table 2 reports the quantitative comparison between the Hydro recycling route and the virgin NMC-111 production process for selected impact categories. The results show that, when normalized per kilo of production of NMC-111 from recycled materials the Hydro route leads to lower impacts with respect to virgin cathode production for the chosen indicators. This outcome reflects the avoidance of upstream mining and refining

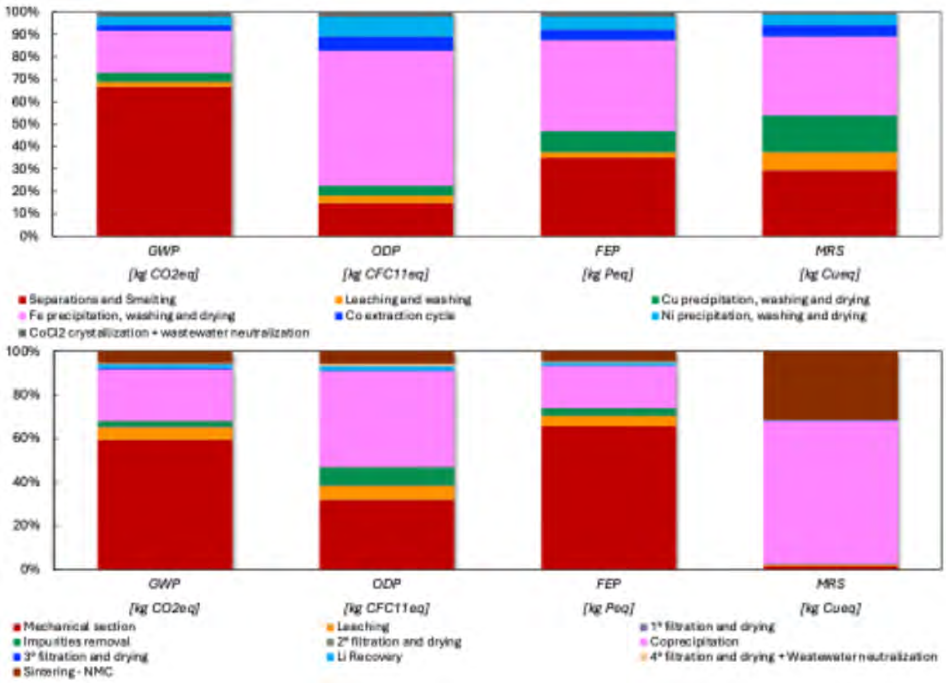

Figure 3. Normalized contribution analysis for Pyro (top) and Hydro (bottom) recycling routes, highlighting the relative contribution of the main unit operations, grouped as in Figure 2, to the selected impact categories.

stages for CRMs, which are among the most impact-intensive phases of virgin cathode production. These results support the interpretation that the environmental performance of Hydro recycling should not be assessed solely based on impacts per kilogram of treated waste, but also in relation to the quality and functionality of the recovered products. When evaluated from this perspective, Hydro recycling shows a clear potential to reduce the environmental burden associated with cathode material production, particularly in scenarios where process optimization measures are implemented.

Taken together, the results reported in Table 1 and Table 2 show that the environmental performance of Hydro recycling strongly depends on the selected reference basis. When impacts are expressed per kilogram of treated battery waste, the Hydro route appears more impactful due to its higher process complexity and material intensity. However, when the same process is evaluated in terms of its main functional output, namely the production of NMC-111 cathode material, Hydro recycling exhibits a clear environmental advantage over conventional production from primary raw materials.

CONCLUSIONS

This work presented a comparative LCA of two LIB recycling routes, namely a pyrometallurgical process and a hydrometallurgical process with co-precipitation, based on detailed process modeling and consistent inventory data. The analysis was carried out using a single functional unit defined as 1 kg of spent LIBs at the recycling facility inlet, considering a representative mix of battery formats and chemistries. The results showed that the two recycling routes exhibit markedly different environmental profiles. The comparison highlighted that some impact-driving elements of the pyrometallurgical process are intrinsically linked to the technology and therefore difficult to mitigate. In contrast, the hydrometallurgical route offers a greater potential for environmental optimization, for instance through alternative solvent choices, improved solvent recovery or process simplification, without necessarily compromising material recovery performance. In addition, when compared to virgin NMC-111 cathode production, the hydrometallurgical recycling route can achieve lower environmental burdens, emphasizing the importance of evaluating recycling processes not only in terms of waste treatment efficiency, but also in relation to the quality and functionality of the recovered products. Overall, the study confirms that the environmental sustainability of LIB recycling strongly depends on process design choices, system boundaries and functional unit selection. The integration of process modeling and LCA proved to be a valuable approach for identifying critical operations and supporting informed optimization strategies.

Further improvements in data quality and the inclusion of industrial-scale operational data would allow more robust conclusions and support informed decision-making in the development of sustainable battery recycling technologies.

AUTHOR IDENTIFIERS

Author ORCIDs:
Vaccari M: 0000-0001-6044-4541
Tognotti L: 0000-0001-5039-515X
Puccini M: 0000-0001-7039-0604

REFERENCES

1. Battery Report 2024. Global battery manufacturing, recycling trends and end-of-life projections. https://www.batteryreport2024.com
2. European Commission. Regulation (EU) 2023/1542 of the European Parliament and of the Council concerning batteries and waste batteries (2023).
3. Latini D, Vaccari M, Lagnoni M, Orefice M, Mathieux F, Huisman J, Tognotti L, Bertei A. A comprehensive review and classification of unit operations with assessment of outputs quality in lithium-ion battery recycling. Journal of Power Sources 546:231979 (2022). https://doi.org/10.1016/j.jpowsour.2022.231979
4. Gaines L. Lithium-ion battery recycling processes: research towards a sustainable course. Sustainable Materials and Technologies 17:e00068 (2018). https://doi.org/10.1016/j.susmat.2018.e00068
5. Harper G, Sommerville R, Kendrick E, Driscoll L, Slater P, Stolkin R, Walton A, Christensen P, Heidrich O, Lambert S, Abbott A, Ryder K, Gaines L, Anderson P. Recycling lithium-ion batteries from electric vehicles. Nature 575:75-86 (2019). https://doi.org/10.1038/s41586-019-1682-5
6. Peters JF. Best practices for life cycle assessment of batteries. Nat Sustain 6:614-616 (2023). https://doi.org/10.1038/s41893-023-01067-y
7. Ciez RE, Whitacre JF. Examining different recycling processes for lithium-ion batteries. Nat Sustain 2:148-156 (2019). https://doi.org/10.1038/s41893-019-0222-5
8. Perocillo YK, Pirard E, Léonard A. Process simulation-based LCA: li-ion battery recycling case study. Int J Life Cycle Assess 30:2494-2518 (2025). https://doi.org/10.1007/s11367-025-02478-z
9. Vaccari M, Parlanti F, Manni FM, Orefice M, Mathieux F, Pannocchia G, Tognotti L, Bertei A. Assessing performance in lithium-ion batteries recycling processes: a quantitative modeling perspective. Resources, Conservation and Recycling 206:107643 (2024). https://doi.org/10.1016/j.resconrec.2024.107643

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Accelerating Design of Chemical Recycling of Plastic Waste through Digitalization: A Bubbling Fluidized Bed Reactor Case Study

Stefano Iannello*, Vassilis M. Charitopoulos, Massimiliano Materazzi

University College London, Department of Chemical Engineering, Torrington Place, London WC1E 7JE, UK
* Corresponding Author: stefano.iannello.18@ucl.ac.uk.

ABSTRACT

The reliable identification of feasible and optimal operating conditions is a key challenge in the design and optimization of thermochemical conversion processes, where kinetics, limited data availability, and strict physical constraints coexist. In this work, a novel data-driven strategy based on Physics-Informed Neural Networks (PINNs) is proposed to explore the operability space of a bubbling fluidized bed (BFB) plastic pyrolysis process. The approach integrates mechanistic knowledge through explicit mass balance constraints with data-driven learning, enabling accurate prediction of and feasibility boundaries. An adaptive sampling framework is employed to iteratively augment the training dataset. The trained PINN surrogate is then used to predict feasible regions and perform constrained optimization aimed at minimizing tar production, which is one of the most problematic byproducts in plastic pyrolysis processes. Beyond classical optimality, a robustness-oriented uncertainty quantification methodology is introduced, combining local perturbations with feasibility filtering to assess the reliability of the optimal solution. Quantitative robustness metrics, including relative objective uncertainty, feasibility retention ratio, and an overall robustness index, are introduced and used to track convergence and select the most reliable operating point. Results show that the proposed framework efficiently converges to accurate feasibility regions and identifies operating conditions that balance optimal performance with robustness against uncertainty. This work provides a methodology for physics-driven process optimization and operability analysis, with potential applicability to a wide range of complex energy and chemical systems.

Keywords: Plastics Recycling, Circular Economy, Surrogate Modelling, Pyrolysis, Data-driven Operability, Physics-Informed Neural Networks

INTRODUCTION

The rampant surge in plastic waste generation and lack of sustainable management practises have posed a substantial threat to the ecosystem and human health in the last two decades [1], and plastic waste production is projected to triple by 2050 reaching 1, 200 million metric tons [2]. Chemical recycling through advanced thermochemical technologies (ATTs), such as pyrolysis [3, 4], represent a promising pathway to tackle the plastic waste issue, allowing us to produce a broad spectrum of chemical feedstock for the manufacturing sector [5] and promoting the transition to a true circular economy [6, 7]. Plastics pyrolysis involves feedstock thermal decomposition at temperatures in the range of approximately 400 – 800 °C and in complete absence of oxidizing agents. This process promotes the formation of both gas and liquid products, whose relative amount depends on the reactor operating conditions. The pyrolysis gas is mainly composed of hydrogen and light hydrocarbons. After cleaning and refining, this gas can be used as chemical feedstock by the manufacturing industry. The liquid product, also known as pyrolysis oil, can be upgraded and used to produce olefins, paraffins and substitute fuels. To achieve these results, bubbling fluidized bed reactors (BFBR) are among the best reactor design for ATTs, due to their favourable operating flexibility, and enhanced heat and mass transfer [8]. However, when processing plastic waste, they are subject to hydrody-

namic instabilities due to the plastic melting and unwanted byproducts formation, such as tar and waxes, that can lead to products quality deviations and even early plant shutdown if not properly monitored during operation. This challenge raises the urgency of developing reliable modelling approaches to control BFBR for chemical recycling, de-risking their implementation at industrial scale and improving their performance. While physics-based or mechanistic models are successfully used in the design stage of industrial chemical reactors, their integration for predicting product quality and controlling operations comes at significant computing costs. The opposite approach involves using fully data-driven models, which can ultimately lead to similar issues as they completely lack physical understanding of the process. This work proposes a combination of established process-specific semi-empirical approaches and physics-informed machine learning algorithms to optimize BFBR for chemical recycling of plastic waste. The objective of this work is to find the optimal operating conditions that minimize the yield of tar produced during pyrolysis and elucidate design space trade-offs [9, 10, 11].

A mechanistic model of a fluidized bed reactor for plastic waste pyrolysis is employed, based on kinetics semi-empirical correlations obtained from the literature [12]. This model predicts the yield of the main pyrolysis products, including volatile matter, tar and char, which represent unreacted carbon. Although this model is relatively simple, in general directly embedding pure mechanistic models within an optimisation framework is computationally prohibiting and this motivates the use of surrogate models [13, 14]. Using adaptive sampling, we develop a Physics-Informed Neural Network (PINN) as a surrogate [15, 16], by minimizing simultaneously data loss and physics loss. In this work, the physics loss enforces the conservation of mass between inlet and outlet of the BFBR. The use of neural networks for feasible region identification has already been proven successful, as highlighted by Metta et al. (2020) [17]. Finally, we use the trained PINN within an optimization framework to quantify maximum operability limits [18] and identify the optimal operating conditions that minimize the tar yield.

METHODOLOGY

Mechanistic model of plastic pyrolysis in BFB

A mechanistic model $\mathcal{M}$ for plastic pyrolysis in a bubbling fluidized bed reactor is employed to describe the relationship between reactor operating conditions and product yields. The model is based on semi-empirical correlations obtained by Vitale et. al (2024) [12], capturing the dependence of temperature on the yield of volatiles, tar and char during pyrolysis of polypropylene, as follows:

$$\boldsymbol{y}(F,T) = \boldsymbol{A}T + \boldsymbol{B} \tag{1}$$

Where $\boldsymbol{y}$ is the vector of volatiles, tar and char mass flow rates $\dot{m}$, and the hydrogen-to-methane mass ratio h in the pyrolysis gas (H_2/CH_4), while $\boldsymbol{A}$ and $\boldsymbol{B}$ are vectors of coefficients obtained from experimental data fitting. The parameter h is a key indicator of gas quality and it is considered in this work as one of the criteria to determine the feasible operability region.

Although the mechanistic model employed in this work consists of linear semi-empirical correlations that are computationally inexpensive, the primary objective of this study is to develop a general algorithmic strategy for feasibility discovery through adaptive sampling and robustness-driven design using PINN.

The reactor system inputs are the plastic feed flow rate F and operating temperature T, which also constitute the decision variables of the optimization problem. The BFB is considered to operate with nitrogen at twice the minimum fluidization condition. The set of operating conditions $\mathcal{D}$ of the BFB is given by the temperature range of 650 – 850 °C and the plastic feed mass flow rate range of 20 – 30 kg/h. The temperature limits are chosen in order to remain consistent with the experimental conditions used by the authors, while the flow rates are chosen to vary in a reasonable range for medium size BFB pyrolizers.

PINN surrogate model

A PINN is developed to act as a surrogate for the mechanistic BFB model. The PINN maps the reactor inputs F and T to the vector of outputs including volatiles, tar and char mass flow rates, and H_2/CH_4. The neural network consists of a fully connected feedforward architecture with 1 hidden layer of 20 neurons and hyperbolic tangent activation functions. Several authors observed that a single hidden layer is a good function approximator, provided sufficient number of hidden neurons [17]. The PINN is trained by minimizing a composite loss function including three contributions, as shown in the equation below: a data loss term, a physics loss term, and an L2 regularization term to avoid overfitting. The loss terms are calculated using a MSE metric.

$$\mathcal{L}(\theta) = \frac{1}{N}\sum_{i=1}^{N}\|\hat{y}_i - y_i\| + \frac{1}{N}\sum_{i=1}^{N}\left\|\hat{\dot{m}}_{out,i} - F_i\right\| + \lambda_{L2}\sum_k \theta_k \tag{2}$$

Where N are the collocation points in the domain, $\hat{y}$ is the PINN prediction, y is the output of the mechanistic model, λ_{L2} is the regularization weight, θ_k is the kth trainable parameter of the PINN, and $\hat{\dot{m}}_{out}$ and F are the overall mass flowrates at outlet and inlet of the BFB, respectively.

The physics loss enforces mass conservation by constraining the predicted outlet flow rate to match the inlet feed flow rate in real physical units. This constraint

embeds first-principles knowledge into the learning process and improves extrapolation behavior in sparsely sampled regions of the operating space. Input and output variables are standardized prior to training to ensure numerical stability.

PINN-driven feasibility analysis

An initial training dataset of 100 points is generated using Latin Hypercube Sampling (LHS) over the predefined set of operating conditions for feed flow rate and reactor temperature. The corresponding mechanistic model outputs are used to initialize the PINN training. This initial step is summarized by the following equations:

$$\{x_i\}_{i=1}^{N_0} \sim LHS(\mathcal{D}) \tag{3}$$

$$y_i = \mathcal{M}(x_i) \tag{4}$$

To progressively improve surrogate accuracy in regions of interest, an adaptive sampling strategy is employed. At each adaptive iteration, the PINN is retrained using all available data, after which a dense pool of candidate points is generated within the operating conditions set using LHS. An acquisition function is defined and evaluated on this pool to identify new informative samples. The acquisition is chosen to encourage the PINN to learn regions with high physics residuals, where the surrogate violates mass conservation. The candidate point maximizing this acquisition function is selected, the corresponding output is evaluated using the mechanistic model and added to the training dataset. This approach encourages the PINN to learn regions of the space according to physics constrains. This is the main novelty aspect of the methodology developed in this work. The adaptive sampling algorithm at each iteration can be summarized by the following steps:

1. Train PINN

$$\theta = arg \min_{\theta} \mathcal{L}(\theta) \tag{5}$$

2. Generate a dense pool of candidate points

$$\{\tilde{\mathbf{x}}_j\}_{j=1}^{N_c} \sim LHS(\mathcal{D}) \tag{6}$$

3. Evaluate surrogate and physics residual

$$r_j = \left|\hat{m}_{out}(\tilde{\mathbf{x}}_j) - \tilde{F}_j\right| \tag{7}$$

4. Evaluate acquisition function

$$a(\tilde{\mathbf{x}}_j) = \frac{r_j}{max(r)} \tag{8}$$

5. Select new sample

$$x_{new} = arg \max_{\tilde{\mathbf{x}}_j} a(\tilde{\mathbf{x}}_j) \tag{9}$$

6. Update training dataset

$$\mathcal{D}_{k+1} = \mathcal{D}_k \cup \{x_{new}, \mathcal{M}(x_{new})\} \tag{10}$$

Subsequently, the process feasibility region is determined through a combination of physics and operational constraints, as follows:

$$\Omega = \left\{(F,T) \in \mathcal{D} \;\middle|\; \begin{matrix} r(F,T) \le \varepsilon \\ h_{min} \le \hat{h}(F,T) \le h_{max} \end{matrix}\right\} \tag{11}$$

The physics residual associated with mass conservation must remain below a prescribed tolerance. Moreover, the H_2/CH_4 is constrained to lie within a specified range, reflecting pyrolysis gas requirements for downstream utilization. The values of ε, h_{min} and h_{max} used to run the algorithm in this work are 1e-2, 0.4 and 0.6, respectively.

Optimization and sensitivity analysis

The trained PINN surrogate is embedded within a constrained optimization framework to identify operating conditions that minimize tar production while satisfying feasibility constraints. The optimal operating point is selected as the feasible point with minimum tar within the feasible region, as follows:

$$x^* = arg \min_{(F,T)\in\Omega} \hat{m}_{tar}(x) \tag{12}$$

To quantify local uncertainty, a local perturbation analysis is performed on the optimum point. The perturbations were generated at each adaptive iteration from a normal distribution with zero mean, as follows:

$$\delta_p \sim \mathcal{N}(0, \sigma) \tag{13}$$

$$x_p^* = x^* + \delta_p \tag{14}$$

Where x_p^* is the pth perturbated optimum point. Each perturbated sample is evaluated using the PINN and it is retained if satisfies the same feasibility constraints shown in Eq. 11. If N_δ and N_{feas} are the number of perturbated samples and feasible perturbated samples for the tar mass flow rate, respectively, it is possible to define the following metrics:

$$FRR = \frac{N_{feas}}{N_\delta} \tag{15}$$

$$\sigma_{tar} = \frac{\Delta\hat{m}_{tar}}{\hat{m}_{tar*}} \tag{16}$$

Where FRR is the feasibility retention ratio, σ_{tar} is the relative uncertainty of the optimal tar value, and $\Delta\hat{m}_{tar}$ and $\hat{m}_{tar}*$ are the 95% confidence interval (CI) for the outlet tar mass flow rate and the optimum tar mass flow rate, respectively. A high value of FRR indicates that small perturbations around the optimum remain feasible with high probability, while σ_{tar} captures the sensitivity of the objective value to local perturbations around the optimum. Finally, an overall robustness index can be calculated as:

$$RI = (1 - \sigma_{tar})\, FRR \tag{17}$$

This metric simultaneously rewards small tar uncertainty and high retention of feasible perturbations. The best feasibility region and optimum are chosen to be those with the highest value of RI.

A block diagram of the entire algorithm is shown in Figure 1.

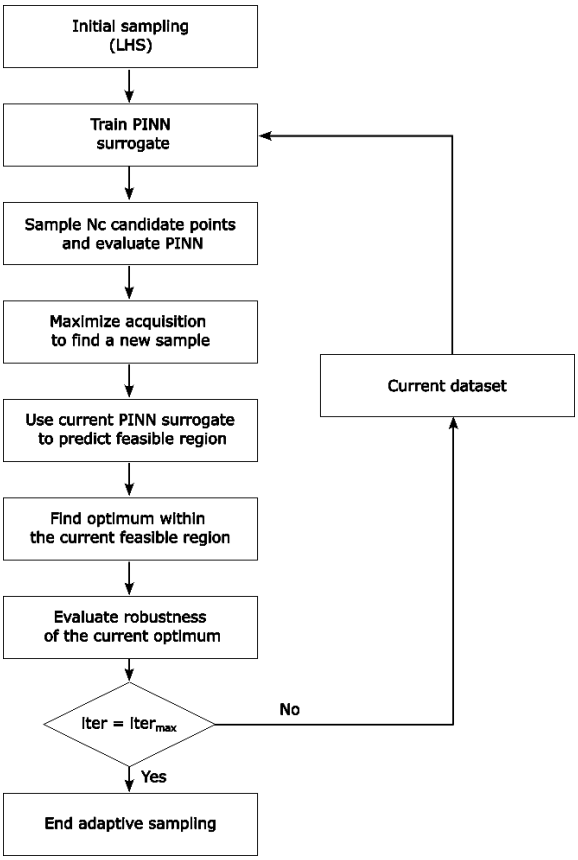

Figure 1. Adaptive sampling algorithm.

Table 1 presents the values of relevant parameters used to run the algorithm.

Table 1: Parameters used in this study.

Parameter	Description	Value/Name
F_{min}, F_{max}	Bounds of plastic mass feed flow rate	20 - 30 kg/h
T_{min}, T_{max}	Bounds of BFB temperature	650 - 850 kg/h
N_0	Number of initial design points (LHS)	100
N_c	Number of samples in the dense pool of candidate points (LHS)	5000
iter	Maximum number of adaptive sampling iterations	200
N_{layers}	Number of PINN hidden layers	1
$N_{neurons}$	Number of PINN hidden neurons	20
η	Learning rate for Adam optimizer	1e-3
N_{epochs}	Number of PINN training epochs	7000
ε	Mass balance residual tolerance	1e-2
h_{min}, h_{max}	Bounds of H_2/CH_4 ratio	0.4 – 0.6
N_δ	Number of perturbations per iteration	5000
σ	Standard deviation for perturbation	0.3

RESULTS AND DISCUSSION

Feasible region identification and optimum

Figure 2 compares the feasible operating regions and tar flow rate iso-contours obtained at two different stages of the adaptive sampling procedure: the last iteration (200 adaptive points) and the iteration corresponding to the maximum robustness index (191 adaptive points).

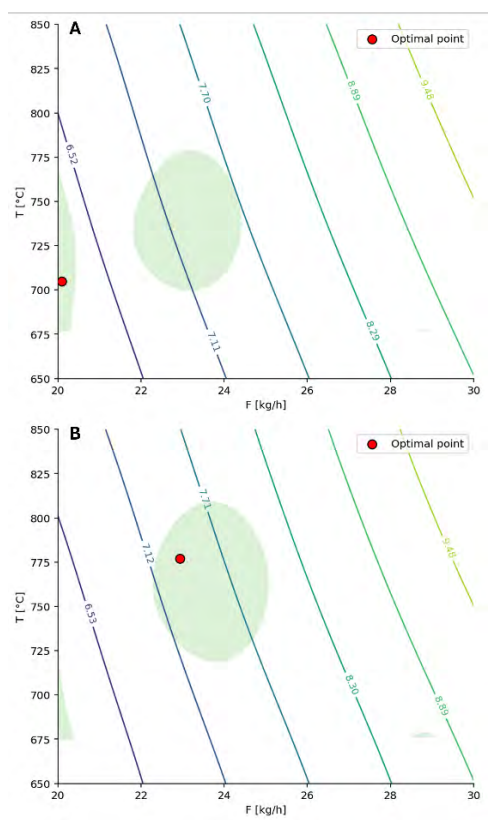

Figure 2. Feasible operating region with tar iso-contours and optimal point after 200 (A) and 191 adaptive samples (B).

Although both feasibility regions are obtained using the same PINN architecture and feasibility criteria, they reveal fundamentally different behaviors in terms of operability structure, and optimum point location. At the last iteration (Figure 2-A, iter = 200), the feasible region appears more fragmented, with greater disconnected areas. The shaded feasibility domain is narrow in the temperature direction. In contrast, at the robustness-optimal iteration (Figure 2-B, iter = 191), the feasible region shows smaller disconnected area at low feed flowrate values when compared with Figure 2-A. The smoother and more compact geometry suggests that the PINN has captured the feasibility constraints more consistently, resulting in a better-defined operating window. A key difference between the two figures lies in the location of the optimal operating point relative to the feasible region. In Figure 2-A, the optimal point lies very close to the feasibility boundary, particularly near the lower temperature and feed flow rate limits. In this case, even small uncertainties in the model predictions or minor disturbances in operating conditions could lead to constraint violation. However, in Figure 2-B, the optimal point is located farther from the boundary of the feasible region. This interior positioning provides a substantial operability margin, making the solution more resilient to uncertainty in process operation.

Perturbation and robustness analysis

Table 2 compares the robustness index RI and optimum for the two different adaptive sample numbers (191 and 200).

Table 2. Optimal solution and *RI* at the two iterations investigated.

Iteration	RI	Optimal values
191	0.538	$\hat{F}^* = 22.93\ kg/h$ $\hat{T}^* = 777.14$ °C $\hat{\dot{m}}_{tar}^{\ *} = 7.37$ kg/h $\hat{h}^* = 0.545$ $r = 2.96e-3$
200	0.157	$\hat{F}^* = 20.10\ kg/h$ $\hat{T}^* = 704.75$ °C $\hat{\dot{m}}_{tar}^{\ *} = 6.17$ kg/h $\hat{h}^* = 0.437$ $r = 3.13e-3$

The discussion in the previous section agrees with the values obtained for the RI. The robustness index integrates objective certainty and feasibility retention into a single metric, and its values are 0.157 and 0.538 at iteration 200 and 191, respectively. Figure 3 compares the uncertainty boundaries obtained by perturbing the optimal operating point identified at the iteration with the highest robustness index and at the final adaptive iteration. Each one represents feasible operating points generated by local perturbations around the nominal optimum and filtered by feasibility constraints as explained in the Methodology section, providing a direct visualization of local operability robustness.

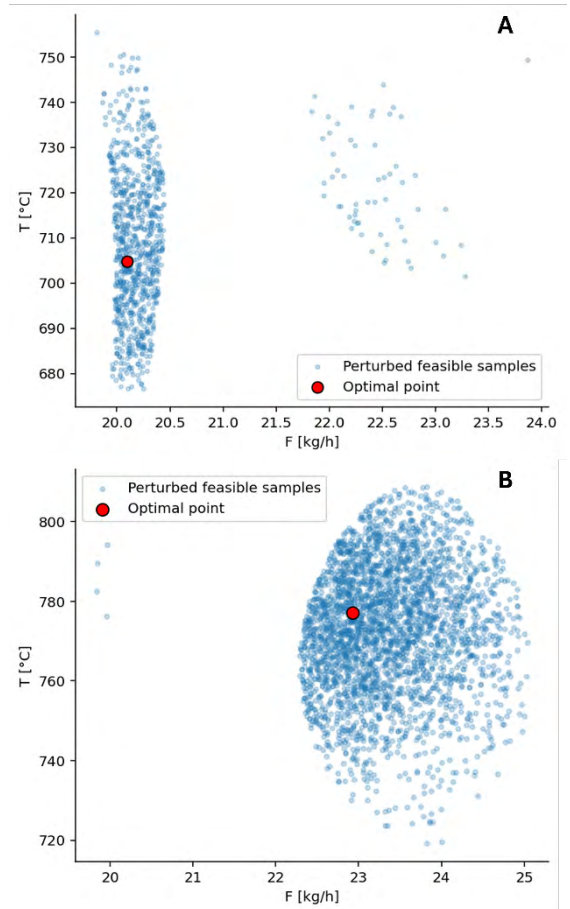

Figure 3. Optimum and feasible perturbed points after 200 (A) and 191 adaptive samples (B).

For the best-robustness iteration (Figure 3-B), feasible perturbations are more densely distributed around the optimum, indicating that small variations in feed rate and temperature may not readily violate feasibility constraints. In contrast, the final iteration exhibits a markedly different perturbation pattern (Figure 3-A). The feasible samples form a narrow, elongated region, particularly constrained in the feed rate dimension. While feasibility is retained under some perturbations, the admissible region collapses rapidly along one coordinate, suggesting that the optimal point lies close to a constraint. This explains the lower robustness index observed at the final iteration. Despite further adaptive refinement of the surrogate, this indicates increased sensitivity of feasibility to small input variations and highlights the risk of operating near constraint boundaries, even when the objective value appears to be favorable.

These results emphasize that continued adaptive sampling does not necessarily improve operational robustness but can lead to over-refinement around feasibility boundaries. While later iterations may sharpen these boundaries or marginally improve objective values, they can also drive the optimizer toward regions of higher sensitivity. Whereas the iteration with the highest robustness index achieves the best trade-off between performance and reliability.

CONCLUSIONS

This work presents a novel physics-informed strategy based on PINNs to explore the design space and predict the feasibility region for a BFB pyrolysis process. Optimization and sensitivity analysis are also performed to find the optimum operating point, as well as investigate its robustness. By embedding fundamental mass balance constraints directly into the learning process, the proposed PINN framework achieves physically consistent predictions while maintaining flexibility in data-driven modeling. The introduction of an adaptive sampling scheme enables efficient exploration of the input space, progressively refining the surrogate model in regions of high violation of the mass balance constraint. This results in accurate operability maps that capture complex feasible regions without requiring exhaustive mechanistic simulations. A key contribution of this study is the explicit integration of robustness analysis into the optimization workflow. Local perturbations around candidate optima are used to quantify uncertainty in both operating variables and objective values, while feasibility filtering ensures that robustness is evaluated under strict process constraints. The combined use of relative objective uncertainty and feasibility retention ratio allows the definition of a robustness index that discriminates between purely optimal solutions and those that are reliably operable under uncertainty. The comparison between early and late adaptive iterations highlights that continued sampling does not necessarily improve robustness, highlighting the importance of robustness-aware criteria. The results show that the most robust operating point may emerge before full convergence of the feasible region, offering practical insights for reducing computational effort while maintaining reliability.

Overall, the proposed methodology provides a framework for feasibility prediction, constrained optimization, and robustness assessment in complex thermochemical systems. Its modular structure and physics-informed features make it readily extendable to higher-dimensional design spaces, additional constraints, and alternative reactor configurations. As this study is a proof of concept based on a semi-empirical literature model, its industrial applicability is currently limited. Future works should be directed towards validation using experimental and plant data.

ACKNOWLEDGEMENTS

The authors would like to acknowledge EPSRC and UK Research and Innovation (UKRI) for funding this research project. Grant numbers: UKRI787, EP/V051008/1, EP/W003317/1.

AUTHOR IDENTIFIERS

Author ORCIDs:
Iannello S: 0000-0001-7894-398X
Charitopoulos V.M.: 0000-0001-9051-917X
Materazzi M: 0000-0001-6350-2559

REFERENCES

1. OECD (2022), Global Plastics Outlook. https://www.oecd.org/en/publications/2022
2. Badejo O, Hernández B, Ierapetritou M. Design of the transition for sustainable plastic waste management under carbon policies in the united states. Ind. Eng. Chem. Res. 64:5476-5492 (2025). https://doi.org/10.1021/acs.iecr.4c03942
3. Crîstiu D, d'Amore F, Bezzo F. Economic and environmental optimisation of mixed plastic waste supply chains in northern italy comparing incineration and pyrolysis technologies. Computers & Chemical Engineering 180:108503 (2024). https://doi.org/10.1016/j.compchemeng.2023.108503
4. García RC, Erickson ED, Martín M, Zavala VM. Evaluating the potential of plastic waste upcycling using thermochemical technologies: a case study in spain. Computers & Chemical Engineering 206:109524 (2026). https://doi.org/10.1016/j.compchemeng.2025.109524
5. Chari S, Sebastiani A, Paulillo A, Materazzi M. The environmental performance of mixed plastic waste gasification with carbon capture and storage to produce hydrogen in the U.K.. ACS Sustainable Chem. Eng. 11:3248-3259 (2023). https://doi.org/10.1021/acssuschemeng.2c05978
6. Avraamidou S, Baratsas SG, Tian Y, Pistikopoulos EN. Circular economy - a challenge and an opportunity for process systems engineering. Computers & Chemical Engineering 133:106629 (2020). https://doi.org/10.1016/j.compchemeng.2019.106629
7. Nixon KD, Schyns ZOG, Luo Y, Ierapetritou MG, Vlachos DG, Korley LTJ, Epps, III TH. Analyses of circular solutions for advanced plastics waste

recycling. Nat Chem Eng 1:615-626 (2024). https://doi.org/10.1038/s44286-024-00121-6
8. Iannello S, Morrin S, Materazzi M. Fluidised bed reactors for the thermochemical conversion of biomass and waste. KONA 37:114-131 (2020). https://doi.org/10.14356/kona.2020016
9. Pistikopoulos EN, Tian Y, Bindlish R. Operability and control in process intensification and modular design: challenges and opportunities. AIChE Journal 67: (2021). https://doi.org/10.1002/aic.17204
10. Sachio S, Kontoravdi C, Papathanasiou MM. A model-based approach towards accelerated process development: a case study on chromatography. Chemical Engineering Research and Design 197:800-820 (2023). https://doi.org/10.1016/j.cherd.2023.08.016
11. Tian H, Ierapetritou MG. A surrogate?based framework for feasibility?driven optimization of expensive simulations. AIChE Journal 70: (2024). https://doi.org/10.1002/aic.18364
12. Vitale A, Di Carlo A, Foscolo PU, Papa AA. Kinetic model implementation of fluidized bed devolatilization. Energies 17:3154 (2024). https://doi.org/10.3390/en17133154
13. Scelfo G, Charitopoulos VM, Vicari F, Tamburini A, Bogle DI, Micale G, Cipollina A. Surrogate-based optimization analysis and scale-up of an ultra-concentrated brine pilot treatment chain for sustainable mineral recovery. Desalination 621:119682 (2026). https://doi.org/10.1016/j.desal.2025.119682
14. Misener R, Biegler L. Formulating data-driven surrogate models for process optimization. Computers & Chemical Engineering 179:108411 (2023). https://doi.org/10.1016/j.compchemeng.2023.108411
15. Cai S, Mao Z, Wang Z, Yin M, Karniadakis GE. Physics-informed neural networks (pinns) for fluid mechanics: a review. Acta Mech. Sin. 37:1727-1738 (2022). https://doi.org/10.1007/s10409-021-01148-1
16. Iannello S, Friso A, Galvanin F, Materazzi M. A hybrid physics–machine learning approach for modeling plastic–bed interactions during fluidized bed pyrolysis. Energy Fuels 39:4549-4564 (2025). https://doi.org/10.1021/acs.energyfuels.4c05870
17. Metta N, Ramachandran R, Ierapetritou M. A novel adaptive sampling based methodology for feasible region identification of compute intensive models using artificial neural network. AIChE Journal 67: (2020). https://doi.org/10.1002/aic.17095
18. Meng Q, Bogle D, Charitopoulos VM. Probabilistic design space exploration and optimization via bayesian approach for a fluid bed drying process. European Journal of Pharmaceutical Sciences 210:107116 (2025). https://doi.org/10.1016/j.ejps.2025.107116

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Principal Component Analysis (PCA) for Evaluation of Fatty Acid Monoalkyl Ester (FAME) Quality towards Sustainable Biodiesel Production from Indonesian Microalgae Strains

Dea Nilam Mustika and Dimitrios I. Gerogiorgis*

School of Engineering, University of Edinburgh, Edinburgh, EH9 3FB, UK
*Corresponding Author: D.Gerogiorgis@ed.ac.uk

ABSTRACT

Biodiesel production from sustainably cultivated plant sources holds extremely high promise globally [6]. Microalgae have been intensely explored as a next-generation source for transport fuel production [9], as they combine attractive characteristics: rapid growth, high lipid content, and environmental benefits. Nevertheless, technical challenges abound regarding the feedstock potential, cultivation process, and its fatty acid mono-alkyl ester (FAME) properties. Performance evaluations for specific microalgal strains [2, 4, 10, 15] are thus of particular interest. The case of Indonesia is particularly significant due to the country's large size, population, biodiversity of terrestrial and marine plant species, and the variety of microalgae that can be harvested and used on an industrial scale for biodiesel production, especially in different media and cultivation methods. Many strains, such as *Botryococcus braunii*, *Chlorella* sp., *Chlamydomonas* sp., and *Nannochloropsis* sp., have a high identified potential [1, 5]. This paper presents the implementation of an established statistical methodology (Principal Component Analysis, PCA) to evaluate the potential of some nitrogen sources for sustainable biodiesel production from Indonesian microalgae strains.

Keywords: Biodiesel, Indonesia, Microalgae, Nitrogen, Principal Component Analysis

INTRODUCTION

Biodiesel is an alternative fuel supply for diesel engines composed of FAME derived from renewable feedstock. Biodiesel has received much scientific attention since it offers comparable fuel characteristics [3, 5]. It is generally produced from a feedstock of vegetable oils, animal fats, and any renewable stock through a chemical reaction called transesterification [1]. Compared to the conservative fuels, biodiesel is more manageable, storable, and transportable due to its biodegradability, high flash point, improved cetane number, and reduced exhaust emission [3, 4]. Nonetheless, the production of biodiesel from second-generation feedstock has led to major conversation concerning the competition between food and fuel. Conversely, the third generation of biodiesel derived from microalgae does not vie with food crops for arable land and water resources.

Microalgae like *Botryococcus braunii*, *Chlorella* sp., *Chlamydomonas* sp., and *Nannochloropsis* sp., which are the most potential strains that are typically found in Indonesian freshwater, marine water, and brackish water systems, are one of the most promising feedstocks for biodiesel production due to their strong adaptability to a wide range of conditions and their characteristics of rapid growth, high oil content, and especially sequestration of atmospheric CO_2 (**Figure 1**) [2, 4, 9]. Microalgae are predicted to produce biodiesel with 15–300 times more oil and 10 times more per unit area of land than a typical terrestrial oleaginous crop used in the first and second generations of biodiesel [2, 9, 10, 14].

Figure 1. Geographical Representation of Microalgae Biodiversity Research Across Indonesia [3].

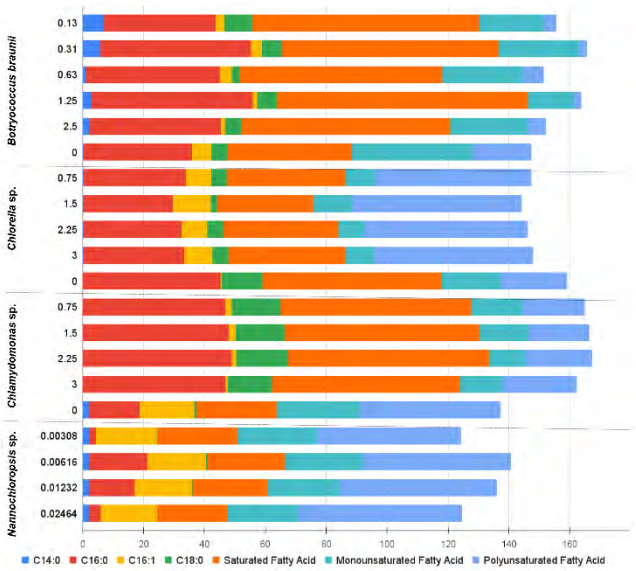

Figure 2. Fatty acid profiles across nitrogen concentration on *Botryococcus braunii, Chlorella* sp., *Chlamydomonas* sp., *Nannochloropsis* sp. [7, 16, 17].

Considering microalgae's potential and impact on the environment, numerous investigations have explored efforts to develop biodiesel production research, such as the different nitrogen sources' application impacts in microalgae cultivation [2, 4, 10, 15]. Nitrogen (N) is an essential core for microalgae's growth; it is a precipitant in all the structural and functional proteins such as chlorophylls, peptides, enzymes, energy transfer molecules, and genetic materials in algal cells. Nitrogen concentration in a culture medium strongly affects both cellular biochemical compositions and growth rate in microalgae, such as decreasing cell growth rate and raising the lipid or carbohydrate content, as well as reducing protein synthesis. It is shown in **Figure 2** that different nitrogen sources and concentrations significantly influenced the FAME of all four microalgal strains, although the proportions and trend of the response were highly dependent on the strain [7, 16, 17]. Changes were observed not only in the total fatty accumulation but also in the relative distribution of saturated (SFA), monounsaturated (MUFA), and polyunsaturated fatty acids (PUFA). These results demonstrate that nitrogen restriction systematically influences lipid metabolism; nevertheless, the consequent FAME profiles are specified by variation in species strategies, which underscores the value of aligning nitrogen use with the microalgae selection to achieve ideal lipid compositions.

In terms of using variations of microalgae, there are several nitrogen sources that can be used in the cultivation stage of biodiesel production, such as KNO_3, $NaNO_3$, $CO(NH_2)_2$, NH_4HCO_3, $(NH_4)_2CO_3$, NH_4NO_3, NH_4Cl, and $(NH_4)_2SO_4$. Different ways to choose nitrogen sources have different effects on the biodiesel production performance; for example, ammonium (NH_4^+) and nitrate (NO_3^-) with either limitation or starvation conditions have been reported as the critical culturing conditions affecting lipid accumulation in many microalgae [15].

The research on variable nitrogen supplementation in microalgae cultivation has increasingly understood how different nitrogen sources might influence biodiesel production and analyzed it based on the FAME properties they have to find out the biodiesel productivity and its overall lipid quality.

The study stated that nitrogen does not simply affect biomass accumulation but also drives substantial shifts in the FAME profile, which then ultimately determines the fuel performance, stability, and suitability for the large-scale application. Therefore, to accurately characterize these shifts and to identify which nitrogen sources offer the most favorable outcomes for sustainable biodiesel generation, a thorough study of statistical processing methodology using principal component analysis (PCA) is needed to pinpoint the most promising potential nitrogen inputs for sustainable biodiesel production from diverse microalgae strains [3].

Various PCA models concerning FAME characteristics are prevalent in **Figure 3**. This study investigated four strains: *Botryococcus braunii*, *Chlorella* sp., *Chlamydomonas* sp., and *Nannochloropsis* sp., which are the most promising for high lipid and oil production in Indonesia, based on several studies about Advanced Biodiesel Production, as shown in **Table 1**. PCA works as a multivariate statistical method that combines the information from several variables and reduces them in order to focus on the essential features and information needed, the indicators with maximum variables [6].

METHODOLOGY

The PCA models we use are scree plot and biplot models. The scree plot model visualizes how much information each PC contributes to the dataset and plots it on the x-axis and the explained variance ratio on the y-axis; it helps us determine how many components should be retained from the graph. The scree plot in **Figure 4** showed us the amount of variances each principal component (PC) explains across all four microalgae strains. It is clear that PC1 and PC2 together explained most of the variance, which is sufficient for interpreting 85.9% of the overall data. The PCA biplot serves as a scatter plot to visualize the relationship between the zoriginal data and principal components.

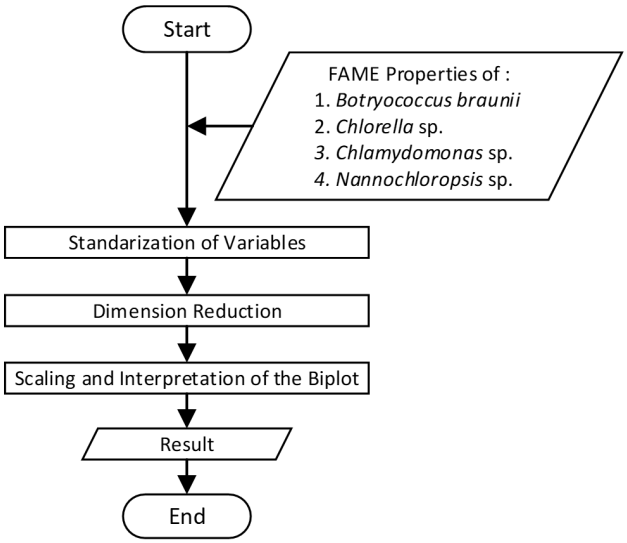

Figure 3. PCA Workflow on Different Nutrient Sources' Effects on Biodiesel Production from Indonesian Micro–algae Strains [6].

The primary advantage of these approaches is the identification of the most effective nitrogen source for strain cultivation in biodiesel production. As shown in **Figure 5**, the PCA requires several steps to get the suitable and compatible variables for making the biplot: standardization of the variables, dimension reduction, scaling and interpreting the biplot, optional de-emphasizing of cases or variables in the biplot, and optional adding of supplementary variables to the biplot. Biplot in **Figure 5** demonstrates that microalgal species are clearly different from each other based on the FAME composition. PC1 displays the condition of lipid saturation very clearly, while PC2 illustrates the length of the chain and how it changes with monounsaturation. These distinctions show that each strain has its own technique of producing lipids and adjusting to its surroundings.

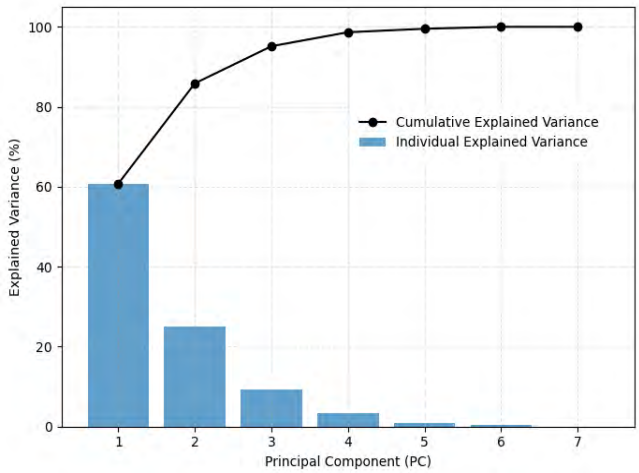

Figure 4. PCA scree plot highlighting Nitrogen concentration relationships between *Botryococcus braunii, Chlorella* sp., *Chlamydomonas* sp., and *Nanno–chloropsis* sp.

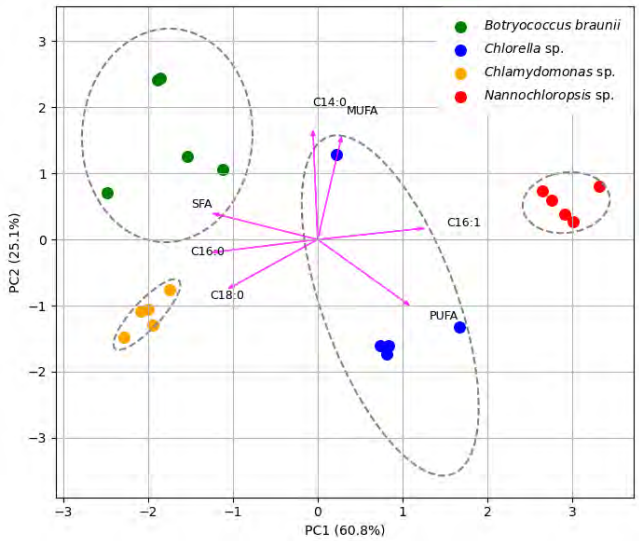

Figure 5. PCA biplot illustrating the relationships among *Botryococcus braunii*, *Chlorella* sp., *Chlamydomonas* sp., and Nannochloropsis sp. under different nitrogen concentrations, based on their FAME profiles.

This research focuses on the impact of nitrogen concentration and its linearity to the type of nitrogen used on Indonesian microalgae strains for bioethanol production. Therefore, the workflow we use reached the analytical point with the dataset we use from **Table 1** using the same scree plot and biplot method. In **Figure 6**, we use a similar scree plot to understand how much information we can gain from each dataset; we observe that PC1 explains the most variance, approximately 34.9%, indicating it captures the dominant pattern in the data. On the other hand, PC2 and PC3 make a big difference and contribute substantially with 26.1% and 19%. The explained variance therefore drops more gradually

after PC3 until it's almost negligible. There is a clear elbow around PC3 and PC4, which means the first two PCs are sufficient to represent most of the meaningful structure in the data.

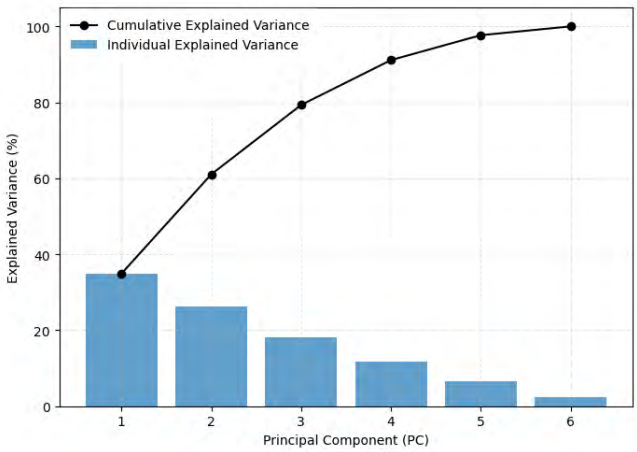

Figure 6. PCA scree plot showing the nitrogen source relationships among *Botryococcus braunii*, *Chlorella* sp., *Chlamydomonas* sp., and *Nannochloropsis* sp.

The FAME data from each strain included different FAME properties, including lauric acid (C12:0), tridecanoic acid (C13:0), tetradecanoic acid (C14:0), myristoleic acid (C14:1), trans-myristoleic acid (C14:1T), cis-myristoleic acid (C14:1C), pentadecanoic acid (C15:0), methyl pentadecanoic (C15:1), palmitic acid (C16:0), palmitoleic acid (C16:1), hexadecadienoic acid (C16:2), heptadecanoic acid (C17:0), heptadecenoic acid (C17:1), stearic acid (C18:0), elaidic acid (C18:1n9t), and oleic acid (C18:1n9c), as shown in **Table 1**. The main focus of this research is 4 different selected variables: lauric acid, tetradecanoic acid, palmitic acid, palmitoleic acid, stearic acid, and oleic acid. The lauric acid and tetradecanoic acid are medium-chain saturated fatty acids that increase biodiesel's cetane number and its combustion quality. Palmitic and stearic acid are long-chain saturated fatty acids that strongly increase the cetane number but significantly worsen the cold-flow behavior, such as higher cloud point, pour point, and Cold Filter Plugging Point (CFPP). Palmitoleic acid enhances cold-flow properties relative to saturated fats, while oleic acid is regarded as the optimal biodiesel component due to its strong balance between good cetane number, excellent cold-flow characteristics, and huge oxidative stability [11].

RESULTS & DISCUSSION

Using the scree plot and biplot, we have visualized clear patterns in how four microalgae strains responded to different nitrogen sources. The PCA results showed clear separation patterns driven by both strain differences and the nitrogen source effect in **Figure 6** and **Figure 7.** The first principal component (PC1) shows that all four microalgae strains formed distinct clusters, indicating that the strains' identity is the dominant factor influencing the fatty acid composition. *Chlamydomonas* sp. was placed at the upper-right region of the score plot and reflected a slightly different behavior from the other strains, while the *Nannochloropsis* sp. formed a separate cluster in the upper-middle area. *Chlorella* sp. shows a tighter grouping in the lower-left area, which means that the lipid profile is more stable across different levels of nitrogen. On the other hand, *Botryococcus braunii* shows the most significant area along PC1, especially when $CO(NH_2)_2$ is present. The spread in this case means strain metabolism changes a lot depending on the nitrogen supply employed in the study.

The second main component (PC2) from **Table 1** data shows how different nitrogen sources are making any impact. Nitrogen sources like NH_4NO_3 and NH_4Cl made the PC2 trend in Chlamydomonas sp. increase significantly, showing how much they affect this strain's lipid metabolism. On the other hand, $NaNO_3$ circumstances tended to diminish PC2 values in many species. **Figure 7**'s biplot analysis showed that saturated fatty acids like C12:0, C13:0, C16:0, and C18:0 were the main factors that caused separation along PC1. This supports their significance in telling species apart. At the same time, fatty acids like C17:1 and C15:0 had greater loadings on PC2, which shows that they are sensitive to changes in nitrogen.

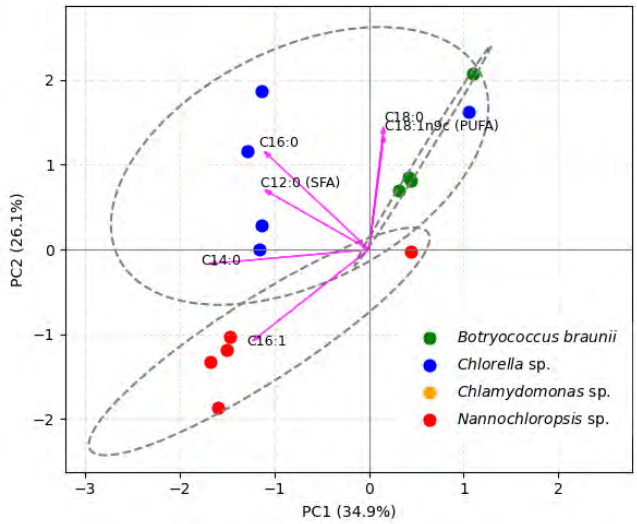

Figure 7. PCA biplot demonstrating the correlation among *Botryococcus braunii*, *Chlorella* sp., *Chlamy–domonas Botryococcus braunii*, *Chlorella* sp., *Chlamy–domonas* sp., and *Nannochloropsis* sp. under C12:0, C14:0, C16:1, C18:0, and C18:1N9C FAME properties.

The discussion in **Figure 6** and **Figure 7** analyzes all

Table 1. FAME profiles of selected Indonesian microalgae strains cultivated with various nitrogen sources. *[a] C12:0, [b] C13:0, [c] C14:0, [d] C14:1. [e] C14:1t, [f] C14:1c, [g] C15:0, [h] C15:1, [i] C16:0, [j] C16:1, [k] C17:0, [l] C17:1, [m] C18:0, [n] C18:1n9t, and [o] C18:1n9c.

Strains	Nitrogen Source	a*	b*	c*	d*	e*	f*	g*	h*
Botryococcus braunii	KNO_3	6.96	6.71	1.41	1.92			0.5	0.24
	$NaNO_3$	6.87	5.3	1.45	3.31			0.53	0.41
	$CO(NH_2)_2$	6.18	5.97	1.24	2.42			0.42	1.14
	NH_4HCO_3	4.56	3.44	1.81	3.37			0.67	0.62
Chlorella sp.	NaNO3	14.3		5					
	$CO(NH_2)_2$	3.5		2.6					
	$(NH_4)2CO_3$	21.4		4.2					
	KNO_3	24.4		4.5					
Chlamydomonas sp.	NH_4NO_3	23.1		4.9					
	NH_4NO_3	0.31		0.4	0.05				
	NH_4Cl	0.29		0.32	0.05				
	KNO_3	0.29		0.4	0.07				
Nannochloropsis sp.	NH_4HCO_3	0.2		5			0.9		
	$(NH_4)2SO_4$			5.6			0.9		
	$NaNO_3$			3.8			0.7		
	NH_4Cl	0.5		4.4		3.6			
	NH_4NO_3	0		5.1			0.6		

Strains	Nitrogen Source	i*	j*	k*	l*	m*	n*	o*	Ref.
Botryococcus braunii	KNO_3	43.25	4.17	3.62	0.12	0.8	0.01	17.8	[4]
	$NaNO_3$	43.48	6.54	0.03	0.03	1.02	0.25	16.64	
	$CO(NH_2)_2$	32.88	4.78	9.93		13.23	0.17	10.3	
	NH_4HCO_3	39.44	8.77	8.44		5.81	0.14	11.18	
Chlorella sp.	NaNO3	41	7.8			7.6		16.8	[15]
	$CO(NH_2)_2$	26.3	4.8	6.5				36.4	
	$(NH_4)2CO_3$	39.7	7.6			.4.4		11.5	
	KNO_3	28	7.5					7	
Chlamydomonas sp.	NH_4NO_3	27.6	6			3.1		3.7	[2]
	NH_4NO_3	10.2	2.91			1.89		4.95	
	NH_4Cl	9.7	3.6			1.6		4.61	
	KNO_3	12.1	3.76			2.64		5.76	
Nannochloropsis sp.	NH_4HCO_3	38.2	32	0.2	0.3	1.7		9.6	[14]
	$(NH_4)2SO_4$	26	29.7	0.6	1	0.2		6.1	
	$NaNO_3$	33.8	0.7	0.4		2.3		9.6	
	NH_4Cl	34.8	29.2	0.2	0.3	2.2		6.8	
	NH_4NO_3	32.2	30.7		0.5	1.6		6.9	

four strains with the specific chosen FAME properties, such C12:0, C14:0, C16:0, C16:1, C18:0, SFA, and PUFA. These FAME properties are explicitly showing us the lipid quality and having the majority of the data from all 4 strains analyzed. The scree plot in **Figure 6** shows that the first two principal components (PC1 and PC2) account for 61.0% of total variances, which means these two components sufficiently summarize the major variability in Table 1 and justify the use of PC1 and PC2 plots for interpretation. The score plot above demonstrated that PC1 primarily separated all four strains' variables based on the inherent species-level differences. Chlamydomonas sp. formed a distinct cluster on the positive side of PC1, while *Nannochloropsis* sp. was very close to the negative side. The gap between the strains shows that there are significant differences in the fatty

acid properties. In contrast, *Botryococcus braunii* showed minimal dispersion, which means its fatty acid profile was relatively stable across the nitrogen treatments. On the other hand, *Chlorella* sp. indicated a steadier movement and metabolic response to different variants of nitrogen.

The biplot in **Figure 7** showed us further insight into the variables after integrating both sample positioning and contribution of fatty acids. C18:1n9c and C18:0 aligned strongly with the positive PC1 axis, indicating that these components were elevated particularly in *Chlamydomonas* sp. However, *Nannochloropsis* sp. seemed to be related with negative PC1 values, which were linked to greater amounts of shorter-chain fatty acids such C14:0 and C16:1. Along PC2, saturated fatty acids such as C12:0, C16:0, and partially C18:0, were positively aligned, reflecting the upward distribution of *Chlorella* sp., which means that these nitrogen sources caused more saturated lipids in this strain.

The combined PCA value shows that strains identity is the main thing that causes variance (PC1), though PC2 may reflect variation associated with the nitrogren treatments we used, as evidenced by the differential movement of all four strains along this axis. The results also point out the significant uniqneness of biochemical responses of each strain, *Chlamydomonas* sp. is characterized by longer-chain fatty acids, *Nannochloropsis* sp. by shorter-chain and monounsaturated fatty acids, and *Chlorella* sp. shows a deep sensitivity to nitrogen form through heightened saturated fatty acid accumulation. This multivariate method clearly shows how the amount of nitrogen changes the lipid makeup in different species.

CONCLUSIONS

PCA on this research effectively shows the kind of strain and nitrogen source affect FAME in Indonesian microalgae relevant to the biodiesel production. While PC1 stipulated that essential strain-level differences were the major causes of variance, PC2 captured additional variability that may be associated with different nitrogen source effects on lipid metabolism. *Botryococcus braunii* exhibited minimal dispersion which indicates a steady fatty acid profile across nitrogen treatments. *Chlamydomonas* sp. was characterized by longer-chain fatty acids, *Nannochloropsis* sp. by shorter-chain and monounsaturated fatty acids, and *Chlorella* sp. shows sensitivity to nitrogen form through heightened saturated fatty acid accumulation. These strain-specific lipid characteristics are relevant to biodiesel quality parameters such as oxidative stability and fuel performance, it concluded that all four Indonesian strains have great potential to be used in biodiesel through an appropriate nitrogen treatment.

REFERENCES

1. Aletheia SP, Syauqi A, Kelvin , Khaira K, Rafi MM. Techno-economic analysis of biodiesel and bioethanol production from chlorella sp. algae biomass. E3S Web Conf. 503:02004 (2024). https://doi.org/10.1051/e3sconf/202450302004
2. Amin NF, et al. Effect of Some Nitrogen Sources on Growth and Lipid of Microalgae *Chlorella* sp. For Biodiesel Production. *Journal of Applied Sciences Research* 9(8): 4845-4855 (2013).
3. Andriani F, Hendrayanti D, Yasman Y. Exploring microalgae diversity in indonesia: harnessing potential for lead bioremediation. BES 8:26-42 (2024). https://doi.org/10.22236/jbes/13225
4. Ashokkumar V, et al. Optimization and characterization of biodiesel production from microalgae *Botryococcus* grown at semi-continous system. *Energy Conversion and Management* 88: 936-946 (2014). http://dx.doi.org/10.1016/j.enconman.2014.09.019
5. Deshmukh S, Kumar R, Bala K. Microalgae biodiesel: a review on oil extraction, fatty acid composition, properties and effect on engine performance and emissions. Fuel Processing Technology 191:232-247 (2019). https://doi.org/10.1016/j.fuproc.2019.03.013
6. Greenacre M, et al. Principal Component Analysis. *Economic Working Series* 1856. Nature Reviews Methods Primers 2(1) (2022). https://doi.org/2010.1038/s43586-022-00184-w
7. Huang X, Huang Z, Wen W, Yan J. Effects of nitrogen supplementation of the culture medium on the growth, total lipid content and fatty acid profiles of three microalgae (tetraselmis subcordiformis, nannochloropsis oculata and pavlova viridis). J Appl Phycol 25:129-137 (2012). https://doi.org/10.1007/s10811-012-9846-9
8. Huber GW, Iborra S, Corma A. Synthesis of transportation fuels from biomass: chemistry, catalysts, and engineering. Chem. Rev. 106:4044-4098 (2006). https://doi.org/10.1021/cr068360d
9. . World energy outlook 2022. OECD (2022). https://doi.org/10.1787/3a469970-en
10. Jahirul MI, Rasul MG, Brown RJ, Senadeera W, Hosen MA, Haque R, Saha SC, Mahlia TMI. Investigation of correlation between chemical composition and properties of biodiesel using principal component analysis (PCA) and artificial neural network (ANN). Renewable Energy 168:632-646 (2021). https://doi.org/10.1016/j.renene.2020.12.078
11. Kasmin H, Lazim AM, Awang R. Determination of fatty acid composition and quality characteristics of oils from palm fruits using solvent extraction. AIP

Conference Proceedings 1678:050002 (2015). https://doi.org/10.1063/1.4931281
12. Lee THW, Show PL, Ong HC, Ling TC, Lan JCW, Chang JS. Microalgae isolation and cultivation technology for mass production. Handbook of Biorefinery Research and Technology :1-29 (2022). https://doi.org/10.1007/978-94-007-6724-9_10-1
13. Lertwattanasakul N, Rodrussamee N, Kumakiri I, Pattanakittivorakul S, Yamada M. Potential of thermo-tolerant microorganisms for production of cellulosic bioethanol. Handbook of Biorefinery Research and Technology: Production of Biofuels and Biochemicals :33-62 (2024). https://doi.org/10.1007/978-981-97-7586-6_35
14. Mata TM, Martins AA, Caetano NS. Microalgae for biodiesel production and other applications: a review. Renewable and Sustainable Energy Reviews 14:217-232 (2010). https://doi.org/10.1016/j.rser.2009.07.020
15. An M, Wang Y, Liu F, Qi X, Zheng Z, Ye N, Sun C, Miao J. Biomass, nutrient uptake and fatty acid composition of chlamydomonas sp. ICE-L in response to different nitrogen sources. Acta Oceanol. Sin. 36:105-110 (2017). https://doi.org/10.1007/s13131-017-0984-4
16. Ruangsomboon S. Effects of different media and nitrogen sources and levels on growth and lipid of green microalgae *Botryococcus braunii* KMITL and its biodiesel properties based on fatty acid composition. *Bioresource Technology* 191: 377-384 (2015). http://dx.doi.org/10.1016/j.biortech.2015.01.091
17. Sivaramakrishnan R, Incharoensakdi A. Enhancement of total lipid yield by nitrogen, carbon, and iron supplementation in isolated microalgae. Journal of Phycology 53:855-868 (2017). https://doi.org/10.1111/jpy.12549

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Section 2

CAPE in Clean Energy Systems

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A MIBLP model for a Northern European negative-emission hydrogen supply chain with CCS in the North Sea

Matthias Maier[a], Sungho Shin[b], Simon Roussanaly[c], and Thomas A. Adams II[a]*

[a] Norwegian University of Science and Technology, Trondheim, Norway
[b] Massachusetts Institute of Technology (MIT), Department of Chemical Engineering, Cambridge, USA
[c] SINTEF Energy Research, Trondheim, Norway
* Corresponding Author: thomas.a.adams@ntnu.no

ABSTRACT

Hydrogen from biomass gasification combined with carbon capture and storage (CCS) can lead to negative emissions and support Europe's energy transition. This study presents a mixed-integer bilinear optimization model for the cost-optimal design of a Northern European hydrogen supply chain with integrated CCS, focusing on exports from Norway to Germany and CO_2 sequestration in Norway. The model is formulated as a superstructure problem and implemented in Pyomo, considering multiple locations for infrastructure nodes and transport options for hydrogen, wood chips, and CO_2. The results show that shipping wood chips and CO_2 is generally more cost-effective than shipping compressed hydrogen. Supply chain costs range from 35–55 NOK/kg H_2, and net-negative emissions (scope 1 and scope 2) are achieved at CO_2 capture rates above approximately 30%.

Keywords: Superstructure optimization, mixed-integer bilinear problem, hydrogen supply chain, carbon capture and storage, negative emissions, hydrogen

INTRODUCTION

The European Union is dependent on imports of fossil fuels today. Hydrogen could be an alternative fuel to natural gas, since it can be produced with low carbon emissions and does not produce CO_2 upon combustion. In 2022, the European Commission presented the REPowerEU plan that aims to import 10 million tonnes of renewable hydrogen into the EU by 2030 [2]. On the other end, Norway could become a large export nation for hydrogen by establishing a sustainable and resilient hydrogen supply chain between Norway and the EU. Moreover, Norway could develop international supply chains for carbon capture and storage in the near future. In August 2025, Northern Lights JV started operations of their storage facility in Kollsnes, Norway with plans to expand to store 5 million tonnes of CO_2 per year [3]. If hydrogen is produced from biomass gasification and the generated CO_2 is captured and permanently sequestered underground, this could even lead to negative emissions. In this study, we calculated the overall costs and emissions and investigated the structure of a cost-optimal supply chain for the export of hydrogen from Norway to Germany combined with carbon capture and storage (CCS) in Kollsnes, Norway.

METHODOLOGY

We developed an optimization model using the python-based optimization framework Pyomo to find the cost-optimal supply chain for export of biomass gasification-based hydrogen from Norway to mainland Europe given a final hydrogen demand and a rate of CO_2 to be captured and sequestered. In previous work, we developed a supply chain model for hydrogen export from Norway to Wilhelmshaven, Germany [1], but it did not consider the possibility of CCS. The model allows for hydrogen production via biomass gasification in various Norwegian locations and in the import terminal in Germany. At the same time, both compressed hydrogen and wood chips can be shipped from Norway to Germany. In this work, we extended the existing model to include new

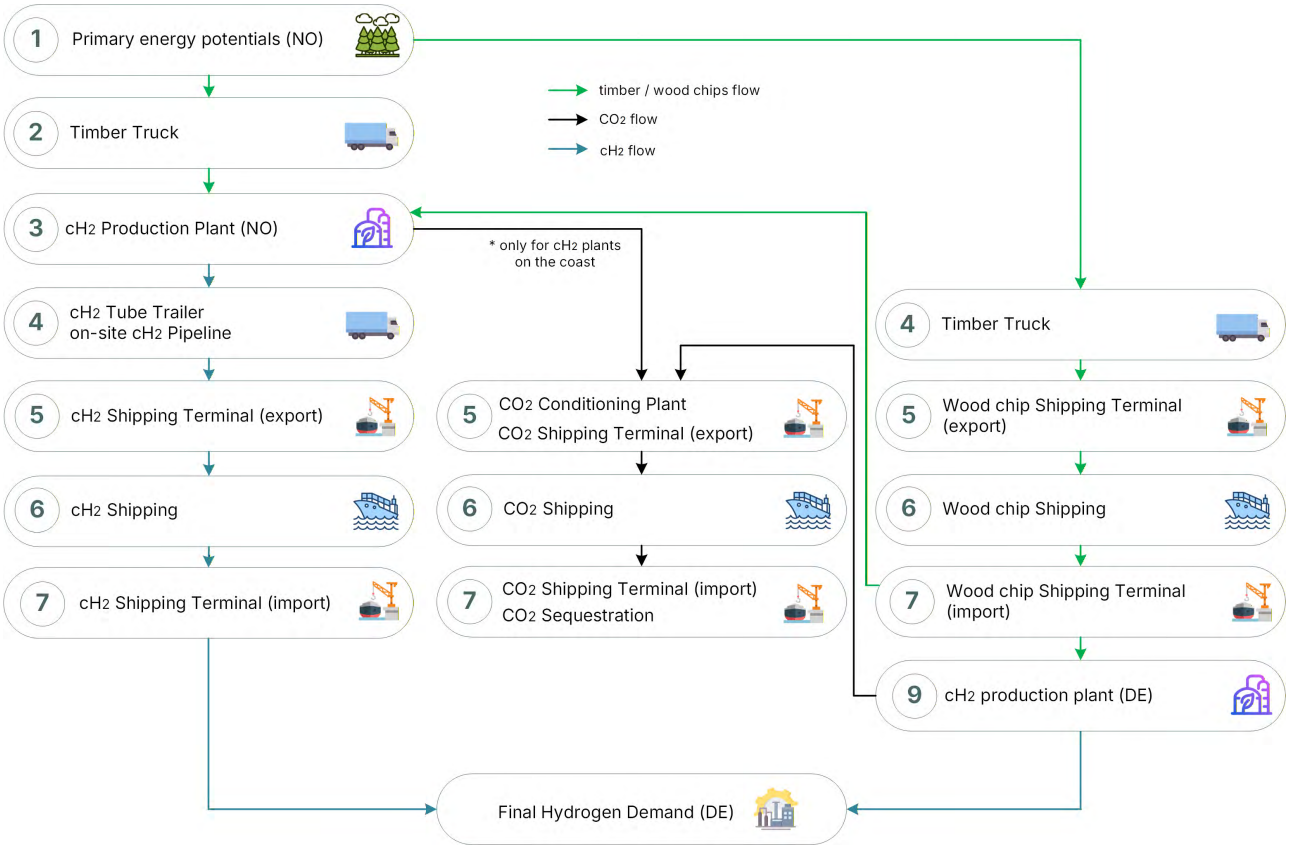

Figure 1. Superstructure design (modified from [1])

commercial CCS options. Since the previous study [1] used 2024 as the base year for currency value, we followed that system and defined NOK2024 as the default economic currency, using an exchange rate of 11.63 NOK/EUR. As for supply chain emissions, we only consider emissions from utilities. Applied emission factors are stated in the supplementary material. The model is categorized as a mixed integer bilinear problem (MIBLP) and was solved to its global optimum using the commercial solver Gurobi [4].

Superstructure design

Figure 1 shows the design of the superstructure graph, which defines the underlying logic of the model. Possible energy conversion routes are defined in the nodes and possible transport routes are defined in the edges. The superstructure graph (see **Figure 1**) is structured in a hierarchical flow-oriented manner, where energy carriers flow from top to bottom. Odd-numbered levels contain nodes, while even-numbered levels contain the edges connecting those nodes (see **Table 1**). At the first level, the model includes renewable biomass potentials at 1, 000 geographically specific locations randomly distributed across Norway (see [1]). At the second level, timber can be transported to hydrogen production facilities within Norway, which must be constructed by the model. At level three, hydrogen is produced through gasification of wood chips. The model permits the construction of hydrogen production plants at 40 randomly distributed geographic locations in Norway.

At level four, (a) timber may be transported from forest resources to existing timber shipping terminals and/or (b) compressed hydrogen may be transported from hydrogen production plants to hydrogen shipping terminals that must also be built. Hydrogen transport can occur either by compressed hydrogen tube trailer or by hydrogen pipeline; the latter is restricted to local transport (i.e., distances under 5 km).

At level five, the model allows construction of up to 20 compressed hydrogen shipping terminals distributed along the Norwegian coast. In addition, shipping terminals for liquefied CO_2 exports may be built at this level at locations where hydrogen production plants are constructed and CO_2 from gasification is captured (see **Figure 2**). At level six, (a) hydrogen can be shipped from Norwegian export terminals to hydrogen import terminals in mainland Europe, (b) wood chips can be shipped from export terminals to import terminals, and/or (c) CO_2 can

be shipped from CO_2 export terminals to CO_2 import terminals. These import terminals must be developed at level seven. For hydrogen, the model considers an import terminal in Wilhelmshaven, Germany. For wood chips, import terminals are considered in Kollsnes, Norway, and Wilhelmshaven, Germany. For CO_2, the model includes an import terminal in Kollsnes, Norway.

At level nine, a hydrogen production plant may be constructed at the wood-chip import terminal in Wilhelmshaven, Germany. Finally, the total amount of hydrogen imported plus the hydrogen produced in Wilhelmshaven must equal the value specified as an input parameter to the model.

The novelty of the work is a) the extension of the hydrogen supply chain model with the possibility of CCS and b) an improvement of the compressed hydrogen (cH_2) shipping model. In [1], the terminal times at the compressed hydrogen export terminals and the compressed hydrogen import terminals were pre-calculated using the relations reported in [5] and then implemented as a parameter in the model. Thereby, all terminals receive the same loading and unloading times. While this logic is a good approximation for simple, single-route supply chains, we found that a more distributed supply chain, as is investigated in this study, requires a more thorough approach. Thus, we implemented the terminal times for the cH_2 export terminals and cH_2 import terminals as decision variables in the model.

Table 1: Indices sets for nodes (modified from [1])

Set	Description
I_1	Timber production sites in Norway
I_3	Biomass gasification hubs in Norway
I_{5,cH_2}	Shipping terminals for export of cH_2
I_{5,CO_2}	Shipping terminals for export of CO_2
$I_{5,CO_2,NO}$	Subset of I_{5,CO_2} with terminals in Norway
$I_{5,wc}$	Shipping terminals for export of wood chips
I_{7,cH_2}	Shipping terminals for import of cH_2
$I_{7,wc}$	Shipping terminals for import of wood chips
I_{7,CO_2}	Shipping terminals for import of CO_2

Optimization objective

The objective function is stated below, where θ is the set of meta parameters and X are the decision variables in the supply chain graph. The model receives a set of meta parameters which contain data for nodes and edges upon model creation. For every run, the desired hydrogen supply and the desired amount of CO_2 to be sequestered must be declared.

$$p_{\text{supply chain}}(\theta, m^{cH_2,\text{demand}}, m^{CO_2\text{ sequestered}}):$$
$$\min_X f_{cost}(X, \theta, m^{cH_2,\text{demand}}, m^{CO_2\text{ sequestered}})$$

$$\begin{aligned}
f_{cost}(X, \theta, m^{cH_2,\text{demand}}, m^{CO_2\text{ sequestered}}) = {} & \underbrace{\sum_{i\in I_1} c_i^{\text{timber 01}}}_{\text{Timber production}} \\
& + \underbrace{c^{\text{timber truck}}}_{\text{Timber truck transport}} + \underbrace{\sum_{i\in I_3} c_i^{\text{gasification 03}}}_{\text{Gasification hub (NO)}} + \underbrace{c^{cH_2\text{ tube trailer 04}}}_{cH_2\text{ tube trailer}} \\
& + \underbrace{\sum_{i\in I_{5,cH_2}} c_i^{cH_2\text{ terminal 05}}}_{cH_2\text{ shipping terminal (export)}} + \underbrace{\sum_{i\in I_{5,CO_2}} c_i^{CO_2\text{ terminal 05}}}_{CO_2\text{ shipping terminal (export)}} \\
& + \underbrace{\sum_{i\in I_{5,wc}} c_i^{\text{wc terminal 05}}}_{\text{wc shipping terminal (export)}} \\
& + \underbrace{\sum_{i\in I_{5,cH_2}} \sum_{j\in I_{7,cH_2}} c_{ij}^{cH_2\text{ shipping,var 06}} + c^{cH_2\text{ shipping,fixed 06}}}_{cH_2\text{ shipping}} \\
& + \underbrace{\sum_{i\in I_{5,CO_2}} \sum_{j\in I_{7,CO_2}} c_{ij}^{CO_2\text{ shipping,var 06}} + c^{CO_2\text{ shipping,fixed 06}}}_{CO_2\text{ shipping}} \\
& + \underbrace{\sum_{i\in I_{5,wc}} \sum_{j\in I_{7,wc}} c_{ij}^{\text{wc shipping,var 06}} + c^{\text{wc shipping,fixed 06}}}_{\text{Wood chip shipping}} \\
& + \underbrace{\sum_{i\in I_{7,cH_2}} c_i^{cH_2\text{ terminal cost 07}}}_{cH_2\text{ shipping terminal (import)}} + \underbrace{\sum_{i\in I_{7,CO_2}} c_i^{CO_2\text{ terminal 07}}}_{CO_2\text{ shipping terminal (import)}} \\
& + \underbrace{\sum_{i\in I_{7,wc}} c_i^{\text{wc terminal 07}}}_{\text{wc shipping terminal (import)}} + \underbrace{\sum_{i\in I_{7,wc}} c_i^{\text{gasification 09}}}_{\text{Gasification hub (DE)}}
\end{aligned}$$

MODEL DESCRIPTION

Below, the amendment of nodes for CO_2 shipping terminals as well as CO_2 shipping edges is presented. For the modification of the cH_2 shipping structure, we direct the reader to the supplementary material.

CO_2 export terminals

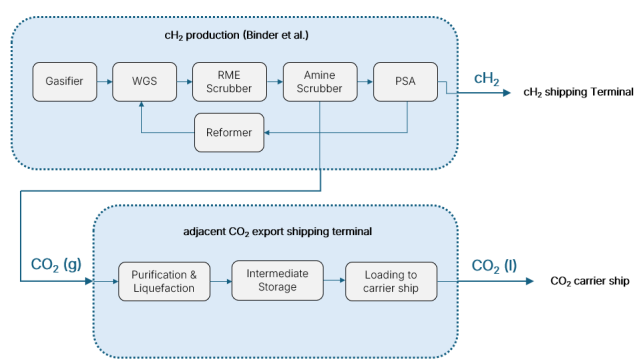

Figure 2. Flow chart of CO_2 export terminal

The superstructure nodes for shipping terminals for export of CO_2 are located at 20 evenly spread-out locations along the Norwegian coast as well as one node in Kollsnes and one node in Wilhelmshaven. At each CO_2 export terminal, a hydrogen production plant must be built, and the CO_2 throughput of the terminal can only be as

large as the CO_2 captured from that hydrogen production plant. We do not consider inland transport of liquified CO_2 from other hydrogen production plants to the terminal. Thus, the CO_2 export terminals only handle locally captured CO_2.

The processes diagram of the CO_2 export terminal is shown in **Figure 2**. Since the design of the hydrogen production plants already contains an amine scrubber prior to PSA for removal of CO_2 (see Binder et al. [6]), the objective of the CO_2 export terminal is a) to purify and liquefy the captured CO_2 stream, b) to store the liquefied CO_2 in an intermediate storage tank and c) to load the liquid CO_2 on incoming ships. As for Kollsnes, the terminal serves mathematically both as an export and an import terminal, since the model allows for hydrogen production combined with CO_2 capture and on-site sequestration. In this special case, the shipping of CO_2 is omitted.

Specific investment costs for CO_2 conditioning as a function on system size were obtained from Deng et al. [7]. Fixed operational costs for CO_2 conditioning as a function of system size were obtained from Bjerketvedt et al. [8]. Specific investment costs for the shipping terminal as well as utility consumptions were calculated using the iCCS tool developed by SINTEF Energy Research [9].

In the optimization model, the shipping terminals for the export of CO_2 are described by the variables for expansion size of terminal i $m_i^{CO_2\ terminal\ 05}$ [kt CO_2 p.a.], expansion decision of terminal i $z_i^{CO_2\ terminal\ 05}$ [-], fixed annual cost at terminal i $c_i^{CO_2\ terminal,fixed\ 05}$ [kNOK2024 p.a.], total annual cost at terminal i $c_i^{CO_2\ terminal\ 05}$ [kNOK2024 p.a.], and annual CO_2 emissions from operations at CO_2 export terminals $e^{CO_2\ terminal\ 05}$ [tCO_{2eq} p.a.]. The CO_2 export shipping terminal i can only process as much CO_2 as is created and captured at the biomass gasification plant i. The captured CO_2 at biomass gasification plants is directly related to its expansion size $m_i^{gasification}$ [kt cH_2 p.a.] via the conversion factor $e^{gasification,direct\ CO_2}$ [$kgCO_2/kgH_2$].

$$m_i^{CO_2\ terminal\ 05} \leq m_i^{gasification} \bullet e^{gasification,direct\ CO_2}\ \forall i \in I_{5,CO_2}$$

The expansion size of each terminal i is also linked to the outflow of the terminal i. All of the CO_2 which is liquified must be exported by ship.

$$m_i^{CO_2\ terminal\ 05} = \sum_{j \in I_{7,CO_2}} m_{ij}^{CO_2\ shipping\ 06}\ \forall i \in I_{5,CO_2}$$

The CO_2 emissions from operations at the terminals arise from electricity consumption $m^{CO_2\ terminal,el\ cons\ 05}$ [MWh/kt CO_2] given the emission factor $e^{el,NO}$ [tCO_{2eq}/MWh] for terminals in Norway or the emission factor $e^{el,DE}$ [tCO_{2eq}/MWh] for terminals in Germany and are linked to the expansion size of the terminals.

$$e^{CO_2\ terminal\ 05} = \sum_{i \in I_{5,CO_2,NO}} m_i^{CO_2\ terminal\ 05} \bullet m^{CO_2\ terminal,el\ cons\ 05} \bullet e^{el,NO} + \sum_{i \in I_{5,CO_2,DE}} m_i^{CO_2\ terminal\ 05} \bullet m^{CO_2\ terminal,el\ cons\ 05} \bullet e^{el,DE}$$

The total annual cost at each terminal i includes the fixed annual costs, electricity costs given the electricity price at terminal i c_i^{el} [kNOK2024/MWh] and an assumed 1 EUR/tonne CO_2 harbour cost.

$$c_i^{CO_2\ terminal\ 05} = c_i^{CO_2\ terminal,fixed\ 05} \bullet z_i^{CO_2\ terminal\ 05} + m_i^{CO_2\ terminal\ 05} \bullet m^{CO_2\ terminal,el\ cons\ 05} \bullet c_i^{el} + 11.63\ [NOK/t] \bullet m_i^{CO_2\ terminal\ 05}\ \forall i \in I_{5,CO_2}$$

The fixed annual costs include the investment and fixed operational costs for CO_2 conditioning, cost for buffer storage (i.e., 1.25 times the CO_2 ship size) and shipping terminal construction costs. The cost parameter is in general a nonlinear function of the expansion size and was linearized via a piecewise linear regression. In the model, a piecewise linear constraint was implemented to link the fixed annual costs to the terminal expansion size for all terminals $i \in I_{5,CO_2}$.

$$c_i^{CO_2\ terminal,fixed\ 05} = a_p \bullet m_i^{CO_2\ terminal\ 05} + b_p\ \forall i \in I_{5,CO_2}$$

where a_p is the specific cost factor and b_p is the offset for the piece $p_{lower} \leq m_i^{CO_2\ terminal\ 05} \leq p_{upper}$. The pieces of the cost function are defined by their breakpoints p_{lower} and p_{upper}. The values for a_p, b_p and the breakpoints can be found in **Table 2**.

Table 2: Piecewise cost function CO_2 export terminals

Parameter	Values
a_p [kNOK2024/kt]	67.7, 60.9
b_p [kNOK2024/a]	8612.6, 16917.4
p [kt/a]	0, 1217.5, 2500

CO_2 import terminal

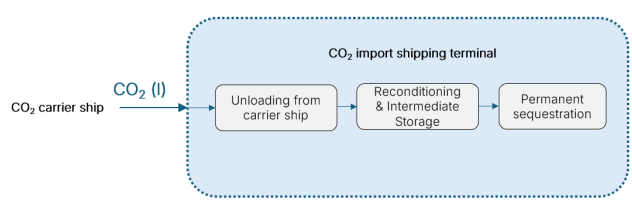

Figure 3. Flow chart of CO_2 import terminal

The superstructure node for the shipping terminal for the import of CO_2 is located in Kollsnes, Norway. The maximum amount of CO_2 to be sequestered from operations investigated in this study is 2.2 Mt/a CO_2,

corresponding to a theoretical 700 MW hydrogen supply chain with full CO_2 capture. This implies an annual wood chip consumption of 1.1 Mt p.a. for hydrogen production.

The processes diagram of the CO_2 export terminal is shown in **Figure 3**. The objective of the CO_2 import terminal is a) to unload the liquid CO_2 from the ship, b) to recondition and intermediately store the CO_2 in storage tanks and c) to sequester the CO_2 in a permanent underground reservoir. Data for investment costs for the terminal, utility consumption and a cost estimate of CO_2 sequestration was calculated with the iCCS tool developed by SINTEF Energy Research [9]. Fixed operational costs were taken from Bjerketvedt et al. [8].

In the optimization model, the shipping terminals for the import of CO_2 are described by the variables for expansion size of terminal i $m_i^{CO_2\ terminal\ 07}$ [kt CO_2 p.a.], total annual cost at terminal i $c_i^{CO_2\ terminal\ 07}$ [kNOK2024 p.a.], and annual CO_2 emissions from operations at CO_2 import terminals $e^{CO_2\ terminal\ 07}$ [tCO_{2eq} p.a.]. The expansion size of each terminal i is linked to the inflow of CO_2 to terminal i via liquid CO_2 shipping.

$$m_i^{CO_2\ terminal\ 07} = \sum_{j \in I_{5,CO_2}} m_{ji}^{CO_2\ shipping\ 06} \quad \forall i \in I_{7,CO_2}$$

The CO_2 emissions from operations at the terminals arise from electricity consumption $m^{CO_2\ terminal,el\ cons\ 07}$ [MWh/kt CO_2] and are linked to the expansion size of the terminals.

$$e^{CO_2\ terminal\ 07} = \sum_{i \in I_{7,CO_2}} m_i^{CO_2\ terminal\ 07} \bullet m^{CO_2\ terminal,el\ cons\ 07} \bullet e^{el,NO}$$

The total annual costs consider investment costs, fixed operational cost, electricity costs as well as an assumed 1 EUR/tonne CO_2 harbour costs and an assumed 19.7 EUR/tonne CO_2 sequestration costs.

$$c_i^{CO_2\ terminal\ 07} = (0.3 + 0.51 + 1 + 19.7)\ [EUR/t] \bullet 11.63\ [NOK/EUR] \bullet m_i^{CO_2\ terminal\ 05} + m_i^{CO_2\ terminal\ 07} \bullet m^{CO_2\ terminal,el\ cons\ 07} \bullet c_i^{el}\ \forall i \in I_{7,CO_2}$$

CO_2 carrier ship

Since the Northern Lights project [3] is based on a 15 bar supply chain design, we adopted this assumption. The design of the ships which carry liquid CO_2 from the on-shore terminal to the off-shore sequestration site follows the design parameters outlined in the Northern Lights concept report [10], using a payload of 7500 m^3, ship speed of 14 knots, and operating pressure of 15 bar. The specific fuel consumption and investment costs were calculated based on data from Roussanaly at al. [9]. Fixed operational costs were assumed to be 5% of the ship's investment costs.

The CO_2 shipping edges connect the nodes for CO_2 shipping terminals (export) $i \in I_{5,CO_2}$ with the nodes for CO_2 shipping terminals (import) $i \in I_{7,CO_2}$. The edges are described by the variables for CO_2 shipping amount $m_{ij}^{CO_2\ shipping\ 06} \mid i \in I_{5,CO_2}, j \in I_{7,CO_2}$ [kt CO_2 p.a.], number of annual trips on the route ij $n_{ij}^{CO_2\ shipping\ trips\ 06} \mid i \in I_{5,CO_2}, j \in I_{7,CO_2}$ [-], variable cost on the route ij $c_{ij}^{CO_2\ shipping,var\ 06} \mid i \in I_{5,CO_2}, j \in I_{7,CO_2}$ [kNOK2024 p.a.], number of CO_2 ships in the fleet $n^{CO_2\ shipping\ ships\ 06}$ [-], fixed annual costs of the fleet $c^{CO_2\ shipping,fixed\ 06}$ [kNOK2024 p.a.], and annual CO_2 emissions from CO_2 shipping $e^{CO_2\ shipping\ 06}$ [tCO_{2eq} p.a.].

The number of ships in the fleet must be greater than the required number of ships (float) in the fleet. The number of ships required in the fleet is calculated as the sum of ships required on each route. We assume that a ship can sail different routes throughout the year. The number of ships required on route ij is calculated as the number of trips required on the route ij divided by number of trips a ship can do per year on route ij. The number of trips required on the route ij is defined as the shipping amount $m_{ij}^{CO_2\ shipping\ 06}$ divided by the payload of the ship $m^{CO_2\ ship,payload}$ [tonnes CO_2]. The number of trips a ship can do per year on route ij is calculated as the maximum annual ship utilization $t^{CO_2\ ship,annual}$ [h/a] divided by the trip time. The trip time is calculated as the time for sailing on route ij from the export terminal to the import terminal and back to the export terminal given the distance $d_{ij}^{CO_2\ shipping\ 06}$ [km] and a speed of $v^{CO_2\ ship}$ [km/day] plus the time spent at terminal $t^{CO_2\ shipping,terminal}$ [h/trip]. The time spent at each terminal is assumed to be 12h. The shipping distances $d_{ij}^{CO_2\ shipping\ 06}$ [km] were pre-calculated using the methodology presented in [1].

$$n^{CO_2\ shipping\ ships\ 06} \geq \sum_{i \in I_{5,CO_2}} \sum_{j \in I_{7,CO_2}} m_{ij}^{CO_2\ shipping\ 06} / (m^{CO_2\ ship,payload} \bullet t^{CO_2\ ship,annual}) \bullet 1000\ [t/kt] \bullet (d_{ij}^{CO_2\ shipping\ 06} \bullet 2 \bullet 24\ [h/day] / v^{cH_2\ ship} + t^{CO_2\ shipping,terminal})$$

The number of trips conducted on route ij (integer) must be greater than or equal to the number of trips required on route ij.

$$n_{ij}^{CO_2\ shipping\ trips\ 06} \geq m_{ij}^{CO_2\ shipping\ 06} / m^{CO_2\ ship,payload} \bullet 1000\ [t/kt]\ \forall i \in I_{5,CO_2}, \forall j \in I_{7,CO_2}$$

The variable cost on route ij is calculated as the product of the number of trips conducted on route ij and the specific trip cost on route ij $c_{ij}^{CO_2\ shipping,trip\ 06}$ [kNOK2024/trip]. The specific trip cost is calculated as the product of specific fuel consumption during sailing $m^{CO_2\ shipping,fuel}$ [g/tCO_2/km] and the bunkering cost for ship fuel. We assume that the ship runs on LNG and consumes 6.87 g/tCO_2/km fuel [9] with a bunkering cost c^{LNG}

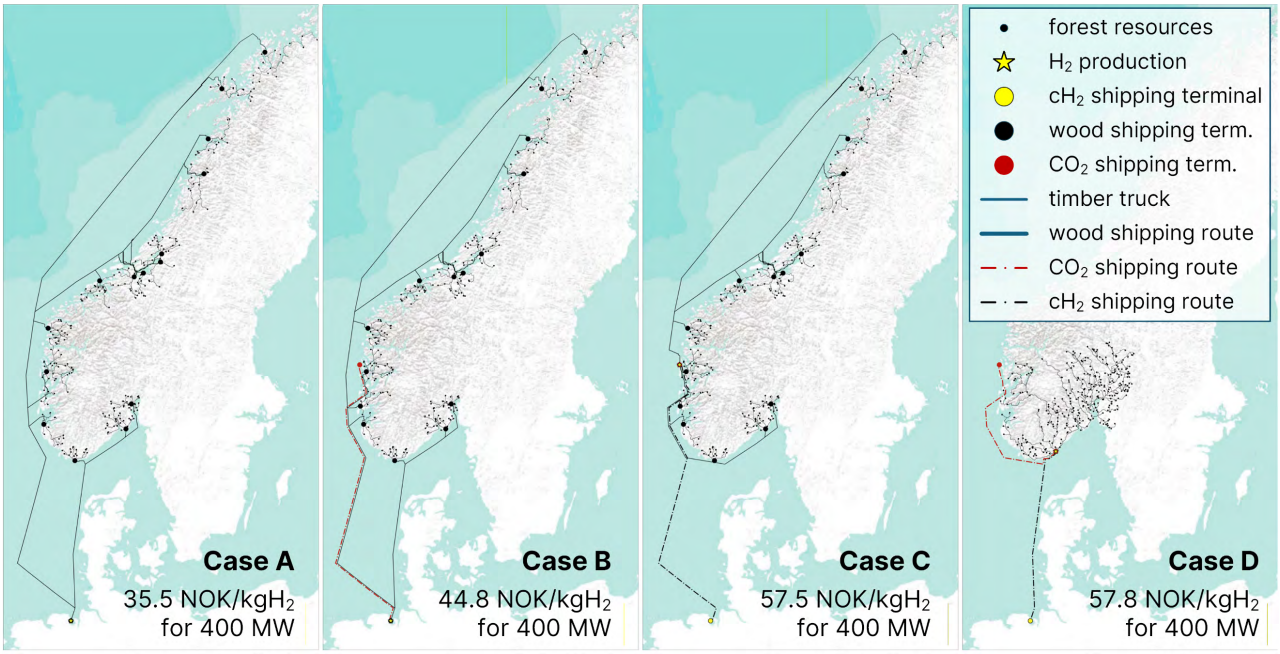

Figure 4. Comparison of supply chain structure results for different cases.

of 7.6 kNOK2024/tonne [11].

$$c_{ij}^{CO_2\ shipping,var\ 06} = n_{ij}^{CO_2\ shipping\ trips\ 06} \bullet c_{ij}^{CO_2\ shipping,trip\ 06}\ \forall i \in I_{5,CO_2}, \forall j \in I_{7,CO_2}$$

$$c_{ij}^{CO_2\ shipping,trip\ 06} = m^{CO_2\ shipping,fuel\ 06} \bullet d_{ij}^{CO_2\ shipping\ 06} \bullet m^{CO_2\ ship,payload} \bullet c^{LNG}/1000/1000\ \forall i \in I_{5,CO_2}, \forall j \in I_{7,CO_2}$$

The fixed annual costs of the fleet are calculated as the product of the number of ships in the fleet times the fixed annual costs per CO_2 ship $c^{CO_2\ shipping,fixed,ship\ 06}$ [kNOK2024/a]. This cost factor includes investment costs, staff costs, insurance and maintenance costs and is based on a previous study [5].

$$c^{CO_2\ shipping,fixed\ 06} = n^{CO_2\ shipping\ ships\ 06} \bullet c^{CO_2\ shipping,fixed,ship\ 06}$$

The annual CO_2 emissions from CO_2 shipping are calculated as the product of annual fuel consumption on all CO_2 shipping routes times the emission factor for LNG fuel e^{LNG} [kgCO_{2eq}/kg].

$$e^{CO_2\ shipping\ 06} = \sum_{i \in I_{5,CO_2}} \sum_{j \in I_{7,CO_2}} m^{CO_2\ shipping,fuel\ 06} \bullet d_{ij}^{CO_2\ shipping\ 06} \bullet m^{CO_2\ ship,payload}/1000/1000 \bullet n_{ij}^{CO_2\ shipping\ trips\ 06} \bullet e^{LNG}$$

RESULTS

Case comparison

Figure 4 shows the optimal supply chain design for a final hydrogen demand of 400 MW and different system constraints. The following cases were investigated:

- **Case A:** The model was run with no additional constraints. The solution is the cost-optimal supply chain without CCS.
- **Case B:** The model was run with the additional constraint that all of the direct CO_2 generated at hydrogen production must be captured and sequestered in Kollsnes.
- **Case C:** The model was run with the additional constraints that a) all of the direct CO_2 generated at hydrogen production must be captured and sequestered in Kollsnes and b) hydrogen must be produced in Norway and exported via compressed hydrogen shipping.
- **Case D:** The model was run with the additional constraints that a) all of the direct CO_2 generated at hydrogen production must be captured and sequestered in Kollsnes and b) hydrogen must be produced in Norway and exported via compressed hydrogen shipping and c) the cH_2 export shipping terminal in the very south of Norway (i.e., Grimstad) must be built out.

It can be seen that the most cost-effective supply chain design is shipping wood chips from various locations along the Norwegian coast to Germany and producing hydrogen there. In Case B, the optimizer chooses to build out a hydrogen production plant at the wood chip

import terminal in Germany and to ship the captured and liquefied CO_2 from Germany to Kollsnes, omitting cH_2 shipping. In Case C, hydrogen production in Norway and export of compressed hydrogen is mandated. Unlike the other cases, shipping of CO_2 is omitted by building a hydrogen production plant in Kollsnes, where the captured CO_2 is stored directly. Wood chips are shipped from the entire Norwegian coast to Kollsnes to serve as feedstock for hydrogen production. In that way, the CO_2 shipping supply chain is replaced by a cH_2 shipping supply chain, which increases the overall supply chain costs by 12.6 NOK/kg H_2 or 28%. This clearly shows that establishing a supply chain for shipping of liquid CO_2 is economically superior to establishing a supply chain for compressed hydrogen shipping. In Case D, by moving the hydrogen production to the south of Norway, the length of the cH_2 shipping route is decreased, reducing the cH_2 shipping costs. On the other hand, shipping of CO_2 from Grimstad to Kollsnes is required. However, this solution is only insignificantly more expensive than Case C, showing that more than one practical solution with very similar hydrogen supply costs exists.

Case B, which is the cost optimal supply chain for hydrogen export combined with carbon capture, has supply chain costs of 44.8 NOK per kgH_2 produced and delivered, with corresponding cradle-to-product net emissions of -7.6 $kgCO_{2eq}/kgH_2$ (scope 1 plus scope 2). Given these values, we calculated the cost of CO_2 (CCA) avoided using

$$\mathrm{CCA} = \frac{c_{\mathrm{cH_2,B}} - c_{\mathrm{ref}}}{e_{\mathrm{ref}} - e_{\mathrm{cH_2,B}}},$$

where $c_{cH_2,B}$ and $e_{cH_2,B}$ describe the specific costs and emissions case B, while c_{ref} and e_{ref} describe the specific costs and emissions for the reference system, which was defined as natural gas combustion. Herein, we assumed specific emissions for the reference system of 201 $kgCO_2/MWh$ (LHV) [12] and approximated the natural gas price (LHV) from the natural gas price (HHV) $\mathrm{p_{ng,HHV}}$ using a factor of 0.9. Assuming a natural gas price of 20 €/MWh (LHV), the cost of CO_2 avoided is 218 €/tonne CO_2. Assuming a natural gas price of 60 €/MWh (LHV), the cost of CO_2 avoided is 114 €/tonne CO_2

Sensitivity analysis

Figure 5 shows the sensitivity analysis of the supply chain costs and emissions in the cost-optimal supply chain given different final hydrogen demands and different CO_2 sequestration rates. It can be seen that the estimated supply chain costs lay between 35 and 55 NOK/kg H_2. While increasing CO_2 sequestration rates lead to higher hydrogen supply costs, increasing hydrogen demand leads to lower hydrogen supply costs. Especially for small hydrogen demands lower than 150 MW and high CO_2 sequestration rates, the hydrogen supply costs are very high.

The specific emissions are only a function of CO_2 sequestration rate and are not significantly affected by the final hydrogen demand. Herein, all direct and indirect supply chains emissions were summed up, and the sequestered CO_2 was counted as negative emissions. A capture rate above 30% is required in order to achieve a negative emission hydrogen supply chain.

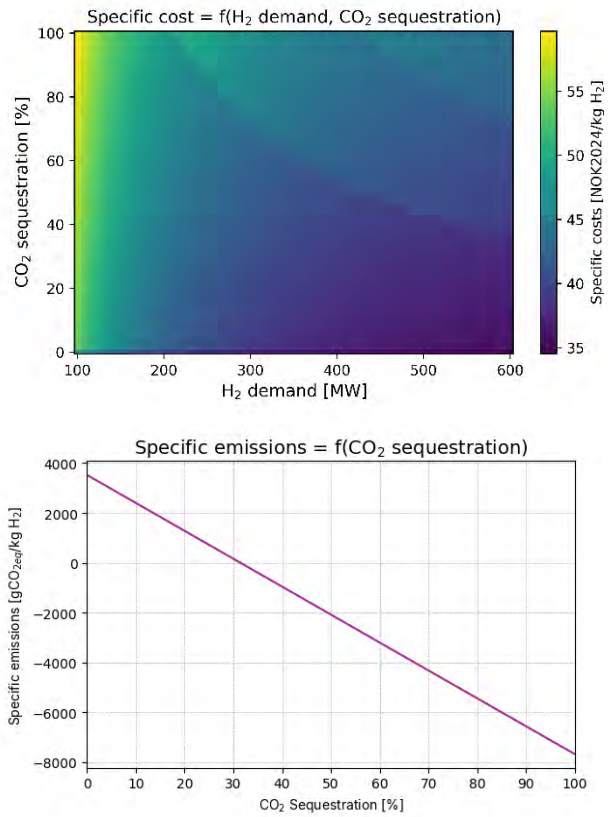

Figure 5. Sensitivity analysis (upper: specific supply chain costs, lower: specific supply chain emissions)

CONCLUSION

This work demonstrates that negative-emission hydrogen supply chains based on biomass gasification combined with CCS can be designed in a cost-effective manner in the Northern European context. Moreover, we note that our analysis does not account for potential additional revenues that could be received from private and public financial incentives from achieving net-negative emissions. We found that even though shipping of CO_2 could be omitted by producing hydrogen in Kollsnes and sequestering CO_2 at the same location while shipping compressed hydrogen to Germany, it is significantly more cost-effective to produce the hydrogen in Germany and ship the CO_2 to Kollsnes.

While low cost of CO_2 avoided can be achieved with

the supply chains investigated in this study, supply chains focusing on different energy carriers might lead to even better CCAs. Future research should investigate the possibilities of integrating alternative energy carriers such as ammonia or methanol. Moreover, the calculation of the size of the intermediate cH_2 storage is based on average ship arrival times and the sizing for buffer storage for CO_2 is based on conventional assumptions. While these assumptions allow the model to be in bilinear space, the validity of these assumptions should be confirmed in future work.

DIGITAL SUPPLEMENTARY MATERIAL

The source code for the model can be found in https://psecommunity.org/LAPSE:2026.0004. The supplementary material can be found at https://psecommunity.org/LAPSE:2026.0008

ACKNOWLEDGEMENTS

This project was funded by the NTNU-MIT Energy Research Programme (project NMERP-C1).

AUTHOR IDENTIFIERS

Author ORCIDs:
Maier, Matthias: 0009-0008-7835-1261
Shin, Sungho: 0000-0002-9889-3278
Roussanaly, Simon: 0000-0002-4757-2829
Adams, Thomas A., II: 0000-0002-9871-9851

REFERENCES

1. Maier M, Shin S, Adams TA II, Export of bioenergy from Norway – Hydrogen or wood chips? *Energy Conversion and Management: X* 30:101887 (2026). https://doi.org/10.1016/j.ecmx.2026.101887
2. Directorate-General for Energy, The European Commission. Hydrogen. (2025) Available from: https://energy.ec.europa.eu/topics/eus-energy-system/hydrogen_en
3. Northern Lights. Northern Lights JV has successfully stored first CO2. August 25 (2025). Available from: https://norlights.com/news/northern-lights-jv-has-successfully-stored-first-co%e2%82%82/
4. Gurobi Optimization, LLC. Gurobi 13.0 Documentation. https://docs.gurobi.com/current/ (Accessed January 2026)
5. Maier M, Adams TA II. Models for a hydrogen supply chain in the Nordics with compressed hydrogen tube trailers and compressed hydrogen shipping. Chemical Product and Process Modeling (2026). https://doi.org/10.1515/cppm-2026-0027
6. Kraussler M, Binder M, Schindler P, Hofbauer H. Hydrogen production within a polygeneration concept based on dual fluidized bed biomass steam gasification. Biomass and Bioenergy 111:320-329 (2018). https://doi.org/10.1016/j.biombioe.2016.12.008
7. Deng H, Roussanaly S, Skaugen G. Techno-economic analyses of CO2 liquefaction: impact of product pressure and impurities. International Journal of Refrigeration 103:301-315 (2019). https://doi.org/10.1016/j.ijrefrig.2019.04.011
8. Bjerketvedt VS, Tomasgard A, Roussanaly S. Deploying a shipping infrastructure to enable carbon capture and storage from norwegian industries. Journal of Cleaner Production 333:129586 (2022). https://doi.org/10.1016/j.jclepro.2021.129586
9. Roussanaly S, Deng H, Skaugen G, Gundersen T. At what pressure shall CO2 be transported by ship? an in-depth cost comparison of 7 and 15 barg shipping. Energies 14:5635 (2021). https://doi.org/10.3390/en14185635
10. Equinor. Northern Lights Project Concept report. RE-PM673-00001 (2019).
11. Ship & Bunker. Rotterdam Bunker Prices LNG. 2025; Available from: https://shipandbunker.com/prices/emea/nwe/nl-rtm-rotterdam#LNG
12. BAFA. Informationsblatt CO2-Faktoren. 2023, Bundesamt für Wirtschaft und Ausfuhrkontrolle. Available from: https://www.bafa.de/SharedDocs/Downloads/DE/Energie/eew_infoblatt_co2_faktoren_2024.html

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Integration of exergy and economic optimization for green hydrogen and power co-generation based on sorbent-enhanced biogas reforming with CO_2 capture

Arthur-Maximilian Báthori[a]*, Calin-Cristian Cormos[a]

[a] Babes-Bolyai University, Faculty of Chemistry and Chemical Engineering, Chemical Engineering Department, 11 Arany Janos, Postal code: 400028, Cluj-Napoca, Romania
* Corresponding Author: arthur.bathori@ubbcluj.ro.

ABSTRACT

In the urgent effort to reduce greenhouse gas (GHG) emissions in the industrial sector, biogas-derived green hydrogen and power co-generation represents a promising solution. Biogas, a renewable and carbon-neutral resource, provides a flexible feedstock for decentralized energy systems, particularly in regions with well-developed agricultural or waste biomass infrastructure. This approach allows the deployment of cost-efficient systems aligned with climate targets and industrial decarbonization roadmaps. Compared to steam methane reforming (SMR), sorbent-enhanced SMR (SE-SMR) with integrated calcium looping (CaL) CO_2 capture reduces process emissions while enhancing hydrogen yield. This study investigates the economic and exergy-based implications of partially splitting hydrogen from a SE-SMR-CaL system producing 50, 000 Nm^3/h of H_2 from desulfurized biogas. Following heat integration using the PINCH methodology, an electrically self-sufficient base case was established. Economic and exergy analyses were conducted in Excel, with cash flow allocation based on exergy contributions. An automated optimization routine identified competitive levelized costs of electricity (LCOE) and hydrogen (LCOH). Results show that increasing the hydrogen split to power generation maintains nearly constant LCOE (<1% variation) while reducing overall exergy efficiency. Partial hydrogen splitting achieves LCOE below 30.5 €/MWh at the expense of increased LCOH, highlighting a trade-off between electricity and hydrogen economics in flexible, multi-vector systems. The workflow demonstrates that incorporating electricity generation into a sorbent-enhanced hydrogen production system using air combustion is technically and economically feasible.

Keywords: Green hydrogen and power, Sorbent-enhanced biogas reforming, Exergy analysis, Modelling, simulation and optimization, Techno-economic assessment

INTRODUCTION

Low carbon technologies and carbon neutrality are key considerations in discussions of a future in which humanity lives in a sustainable manner with respect to available resources. The main challenged addressed in the context of sustainability is anthropogenically induced climate change, which has led to the implementation of a number of significant and ambitious policies worldwide [1]. At the core of these policies is the carbon element, primarily in the form of carbon dioxide (CO_2), a greenhouse gas emitted through various anthropogenic activities, including energy production, the bulk chemical industry, and transportation. By mitigating (or offsetting) CO_2 emissions by choosing renewable raw materials and/or electricity sources, short- to mid-term solutions can be developed that benefit from the use of existing infrastructure [2]. In this context, carbon capture, utilization and storage (CCUS) technologies emerge as carbon sinks by either completely sequestering (via storage) or off-setting (via utilization) the CO_2 footprint of a given technology. However, the transition to carbon-free energy vectors, such as hydrogen and ammonia, represents a long-term solution [3] that requires substantially greater research, development and deployment efforts than previously discussed approaches to increase the

technology readiness level (TRL) and accelerate industrial deployment.

In this context [4], the present work evaluates the techno-economic and exergy-based implications of partially splitting hydrogen (H_2) from a sorbent-enhanced steam methane reforming (SE-SMR) plant, coupled with calcium looping technology (CaL) [5] with a design capacity of 50, 000 Nm3/h of high-purity H_2 for electricity generation. In line with the considerations discussed above, desulfurized biogas is employed as the raw material for hydrogen production [6], offering a sustainable energy pathway for the deployment of an innovative, cost-efficient, and decarbonized system that aligns with climate targets and industrial decarbonization roadmaps. By employing process modelling and simulation using ChemCAD, techno-economic and exergy analyses using Excel, and automation using Visual Basic for Applications (VBA), a robust framework was developed for evaluating multi-product renewable energy systems. In particular, an optimal range for the hydrogen split fraction was identified, within which competitive levelized costs of electricity (LCOH) and hydrogen (LCOH) can be achieved under realistic economic bounds.

PLANT CONFIGURATION, DESIGN ASSUMPTIONS, EXERGY AND ECONOMIC ANALYSIS METHODOLOGY

The plant (Figure 1) is designed to produce 50, 000 Nm^3/h of H_2 using desulfurized biogas, while also providing inherent carbon sequestration through carbonation-calcination cycles. The resulting hydrogen-enriched gas stream from the reformer is further purified via pressure-swing adsorption (PSA). A flexible fraction of pure H_2 is diverted to generate electricity through air combustion for high-pressure (HP) steam generation in a Rankine cycle, while the remainder is compressed and stored for downstream utilization (e.g., chemical applications). The SE-SMR-CaL process employs a dual circulating fluidized bed system, enabling continuous operation with high CO_2 capture efficiency and stable hydrogen production. The process produces high-purity hydrogen and CO_2 streams, with the latter being further processed to meet sequestration or utilization criteria, primarily through water removal using triethylene glycol (TEG) dehydration. In this study, the net power output considers the electricity required for oxygen production from the air separation unit (ASU) and for compressing pure H_2, CO_2, and streams within the TEG unit. After heat integration of the available hot streams using the PINCH methodology to maximize heat recovery (Figure 2), an electrically self-sufficient base case was established.

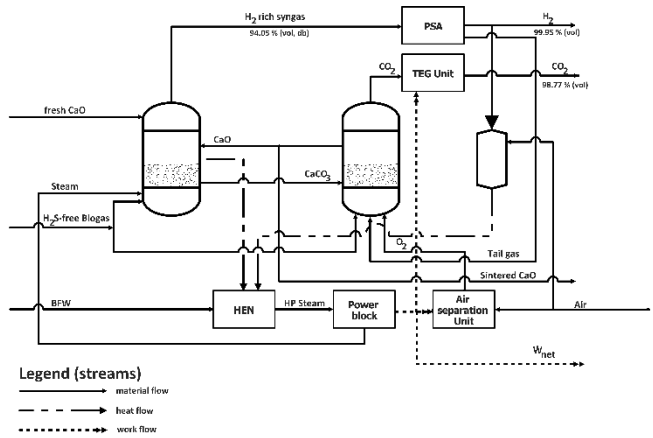

Figure 1. Schematic design of the SE-SMR-CaL plant.

The chemically reactive system within the reformer is described by the following reactions:

$$r_1\ CH_4 + H_2O \leftrightarrow CO + 3H_2\ \Delta H_1 = 206.1\ kJ/mol \quad (1)$$

$$r_2\ CO + H_2O \leftrightarrow CO_2 + H_2\ \ \Delta H_2 = -41.1\ kJ/mol \quad (2)$$

$$r_3\ CH_4 + 2H_2O \leftrightarrow CO_2 + 4H_2\ \Delta H_3 = 165.0\ kJ/mol \quad (3)$$

$$r_4\ CaO + CO_2 \leftrightarrow CaCO_3\ \ \ \Delta H_4 = -178.0\ kJ/mol \quad (4)$$

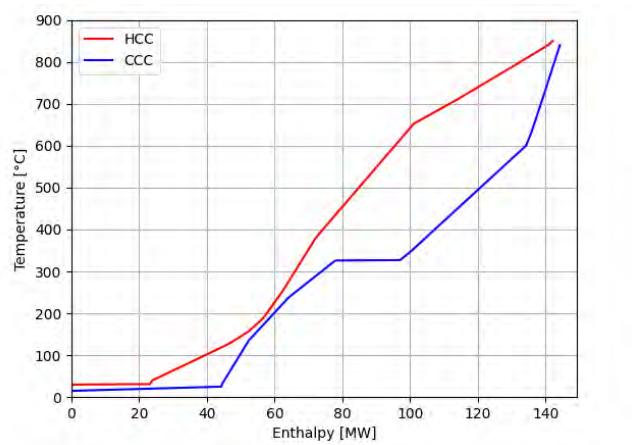

Figure 2. Thermal integration analysis of the SE-SMR-CaL plant.

The reformer was modelled as a kinetic reactor, with the following reaction rates and additional expressions implemented using VBA within ChemCAD [7, 8]:

$$r_1 = \frac{\frac{k_1}{p_{H_2}^{2,5}} \cdot \left(p_{CH_4} \cdot p_{H_2O} - \frac{p_{H_2}^3 \cdot p_{CO}}{K_1}\right)}{\left(1 + K_{CO} \cdot p_{CO} + K_{H_2} \cdot p_{H_2} + K_{CH_4} \cdot p_{CH_4} + \frac{K_{H_2O} \cdot p_{H_2O}}{p_{H_2}}\right)^2} \quad (5)$$

$$r_2 = \frac{\frac{k_2}{p_{H_2}} \cdot \left(p_{CO} \cdot p_{H_2O} - \frac{p_{H_2} \cdot p_{CO_2}}{K_2}\right)}{\left(1 + K_{CO} \cdot p_{CO} + K_{H_2} \cdot p_{H_2} + K_{CH_4} \cdot p_{CH_4} + \frac{K_{H_2O} \cdot p_{H_2O}}{p_{H_2}}\right)^2} \quad (6)$$

$$r_3 = \frac{\frac{k_3}{p_{H_2}^{3.5}} \cdot \left(p_{CH_4} \cdot p_{H_2O}^2 - \frac{p_{H_2}^4 \cdot p_{CO_2}}{K_3}\right)}{\left(1 + K_{CO} \cdot p_{CO} + K_{H_2} \cdot p_{H_2} + K_{CH_4} \cdot p_{CH_4} + \frac{K_{H_2O} \cdot p_{H_2O}}{p_{H_2}}\right)^2} \quad (7)$$

$$r_4 = \frac{3k_4}{\bar{R_p}} \bullet (1 - \bar{X}_{CaO})^{\frac{2}{3}} \bullet \left(\frac{p_{CO_2} - p_{CO_2,eq}}{RT}\right)^{0.66} \quad (8)$$

$$p_{CO_2,eq} = 1.01325 * 10^{\left(\frac{-8308}{T} + 7.079\right)} \quad (9)$$

$$k_i = k_i^0 \bullet e^{-\frac{E_i}{RT}} \quad (10)$$

$$K_j = K_J^0 \bullet e^{-\frac{\Delta H_J}{RT}} \quad (11)$$

The kinetic and thermodynamic parameters employed in expressions (5)-(11) are summarized in Table 1, while general modelling assumptions of the SE-SMR plant are summarized in Table 2.

Table 1: Kinetic and thermodynamic parameters employed for reformer modelling.

Symbol	Value	Unit
k_1^0	$4.225 \bullet 10^{15}$	$mol\ bar^{0.5}\ kg^{-1} cat\ s^{-1}$
k_2^0	$1.995 \bullet 10^{6}$	$mol\ kg^{-1} cat\ s^{-1}\ bar^{-1}$
k_3^0	$1.020 \bullet 10^{15}$	$mol\ kg^{-1} cat\ s^{-1}$
k_4^0	3.05	$m\ s^{-1}$
E_1	240.1	$kJ\ mol^{-1}$
E_2	67.13	$kJ\ mol^{-1}$
E_3	243.9	$kJ\ mol^{-1}$
E_4	32.6	$kJ\ mol^{-1}$
K_C^0	$5.8 \bullet 10^{-4}$	$bar^{-0.5}$
K_H^0	$1.6 \bullet 10^{-2}$	$bar^{-0.5}$
K_{CO}^0	$8.23 \bullet 10^{-5}$	bar^{-1}
$K_{H_2}^0$	$6.12 \bullet 10^{-9}$	bar^{-1}
$K_{CH_4}^0$	$6.65 \bullet 10^{-4}$	bar^{-1}
$K_{H_2O}^0$	$1.77 \bullet 10^{5}$	–
ΔH_C	-42.00	$kJ\ mol^{-1}$
ΔH_H	-16.00	$kJ\ mol^{-1}$
ΔH_{CO}	-70.65	$kJ\ mol^{-1}$
ΔH_{H_2}	-82.90	$kJ\ mol^{-1}$
ΔH_{CH_4}	-32.28	$kJ\ mol^{-1}$
ΔH_{H_2O}	88.68	$kJ\ mol^{-1}$
$\bar{X}_{CaO}$	0.72	–
$\bar{R_p}$	125	μm

Table 2: Primary design assumptions of the SE-SMR-CaL plant.

Plant subsystem	Design specification
Fuel characteristics	Biogas purity (vol. %): 59.75% CH_4, 40.00% CO_2, 0.20% N_2, 0.05% O_2 Lower heating value (LHV): 17.58 MJ/kg
Air Separation Unit (ASU)	Oxygen purity (vol. %): 95% O_2, 3% Ar, 2% N_2 ASU power consumption: 200 kWh/t O_2
Reforming reactor	Reactor temperature & pressure: 650°C & 1.25 bar Reactor model: kinetic reactor Thermal operation mode: isothermal Pressure drop: 0.12 bar Sorbent make-up: 5%
Calcination reactor	Reactor temperature & pressure: 850°C & 1 bar Reactor model: Equilibrium reactor $CaCO_3$ decomposition rate: 98% Pressure drop: 0.15 bar
H_2 purification and compressing	Purification unit: Pressure Swing Adsorption (PSA) H_2 recovery yield: 85% H_2 purity: min. 99.95% (vol.) Final compressing pressure: 60 bar Compressor efficiency: 80%
CO_2 drying and compressing	Drying unit: TEG (Triethylene glycol) CO_2 composition (vol. %): >95% CO_2, <2, 000 ppm CO, <250 ppm H_2O, <100 ppm H_2S, <4% non-condensable gases (N_2, H_2, Ar) Delivery pressure: 120 bar
Steam cycle	HP steam parameters: 600°C & 120 bar LP steam parameters: 152°C & 1.25 bar Condensation pressure (steam turbine): 0.05 bar Steam expander efficiency: 92%
Heat exchanger network (HEN)	ΔT_{min}: 10°C Pressure drop: 2 - 4% of inlet pressure
Thermodynamic packages	Biogas processing unit: SRK Gas compression units: BWRS CO_2 drying unit: TEG Dehydration

Exergy efficiencies (ε) are defined according to the fuel (F)-product (P) paradigm [9], accounting for the physical and chemical contributions of inlet and outlet streams within a given control volume, defined such that the boundary temperature equals the environmental temperature. In this study, pre-treated biogas, oxygen from the ASU, required boiler feed water (BFW) for steam reforming, and fresh CaO are considered fuels, while the purified H_2, captured and processed CO_2, and net electricity output are treated as products. Exergy losses are attributed to CO_2 emissions from the TEG unit, removed sintered sorbent stream, and process water discharged as condensate streams.

$$\varepsilon = \frac{\dot{E}_P}{\dot{E}_F} = 1 - \frac{\dot{E}_D + \dot{E}_L}{\dot{E}_F} \tag{12}$$

$$e^{PH} = (h - h_0) - T_0 * (s - s_0) \tag{13}$$

$$\overline{e}^{CH} = \sum_1^k x_k * \overline{e}_k^{CH} + \overline{R} T_0 * \sum_1^k x_k * ln\,(x_k) \tag{14}$$

$$\dot{E} = \dot{m} * \left(e^{PH} + e^{CH}\right) \tag{15}$$

To facilitate the computation of chemical exergy using ChemCAD's thermodynamic packages, a custom unit operation was developed to duplicate a stream with a specified composition at a given temperature and pressure and bring it to the restricted dead state (RDS) by flashing it at T_0 = 25 °C and p_0 = 1 bar. All RDS streams are subsequently mixed to verify the mass balance of the overall plant.

The economic analysis focuses around the estimation of the following relevant economic performance indicators:

Capital expenditure of a process unit (C_E) with capacity Q is calculated using the cost correlation approach [10], which considers base cost of the evaluated subsystem (C_B), its corresponding base production size (Q_B) and a technological maturity exponent (M):

$$C_E = C_B * \left(\frac{Q}{Q_B}\right)^M \tag{16}$$

The reference data (C_B, Q_B and M) were extracted from the literature and the outdated values were updated to the present-day conditions by usage of Chemical Engineering Plant Cost Index (CEPCI):

$$\frac{Cost_{y1}}{Cost_{y2}} = \frac{CEPCI_{y1}}{CEPCI_{y2}} \tag{17}$$

Operational and maintenance (O&M) costs account for both fixed costs (e.g., maintenance, labour, and administrative expenses) and variable costs (e.g., feedstock chemicals, catalyst replacement, cold utilities, and costs associated with CO_2 transport and storage) over the operational lifetime of the plant.

The levelized cost of hydrogen (LCOH) was calculated as the sum of the annualized capital cost ($CC_{Annualized}$), assuming straight-line depreciation, and the operational and maintenance (O&M) costs, divided by the annual hydrogen thermal output ($Q_{H2,\,Annual}$):

$$LCOH = \frac{CC_{Annualized} + O\&M}{Q_{H_{2,Annual}}} \tag{17}$$

The Levelized cost of electricity (LCOE) was computed in a similar manner than levelized cost of hydrogen (LCOH), with additional consideration for the correlation factors that account for the combined energy output (hydrogen and net power) of the plant.

Varying the hydrogen split fraction decreases the required capacity for the hydrogen compression system while increasing the power block capacity, impacting capital expenditures and, consequently, levelized costs. The optimization focuses on determining the net present value (NPV) of the overall plant, accounting for changes in the cost correlation method and allocating the cash flow factor based on the exergy values of the produced hydrogen and electricity.

Excel Solver is used to determine the levelized costs of electricity (LCOE) and hydrogen (LCOH) that yield an NPV of zero, subject to non-negativity constraints:

Find LCOE, LCOH such that: NPV(LCOE, LCOH) = 0, subject to LCOE ≥ 0, LCOH ≥ 0, NPV ≥ 0.

Automation is achieved by employing the sensitivity study tool and data maps within ChemCAD to extract the relevant thermodynamic data of the main material streams (specific enthalpy, specific entropy, stream composition, stream mass rate) and unit operation parameters (e.g., compressor power, divider split fraction), in combination with a VBA subroutine that links each sensitivity study run to an instance of Excel Solver:

```
Sub RunExcelSolverFromChemCAD()
 Dim xlApp As Object
 Dim xlWB As Object
 Set  xlApp  =  CreateObject("Excel.Applica-
tion")
 xlApp.Visible = False
 Set        xlWB        =        xlApp.Work-
books.Open("path_to_solver")
 xlApp.RUN "RunSolverAuto"
 xlWB.Close SaveChanges:=True
 xlApp.Quit
 Set xlWB = Nothing
 Set xlApp = Nothing
End Sub
```

RESULTS AND DISCUSSION

The mathematical model of the plant was validated primarily by analyzing the dry volumetric composition of the simulated reformer gaseous output with respect to available experimental data [12, 13] (Table 3).

Table 3: Validation of the kinetic reformer model.

Parameter	Simulation	Experimental
% CH_4	2.15	2-2.5
% $CO + CO_2 + N_2$	3.8	3-4
% H_2	94.05	94-94.5
CH_4 conversion	90.77	90

As shown, the simulation data are in range and close agreement to the experimental data, with relative errors for hydrogen composition and methane conversion below 0.5%. This underlines the reliability of the model and the robust behaviour of the kinetically modelled reformer. A more detailed model validation is presented in a previous work of the authors [11] which serves as the basis of the current optimization study, including sorbent deactivation [14] as a function of the number of looping cycles and several arguments for the sorbent make-up fraction and selected mean sorbent conversion within the reformer.

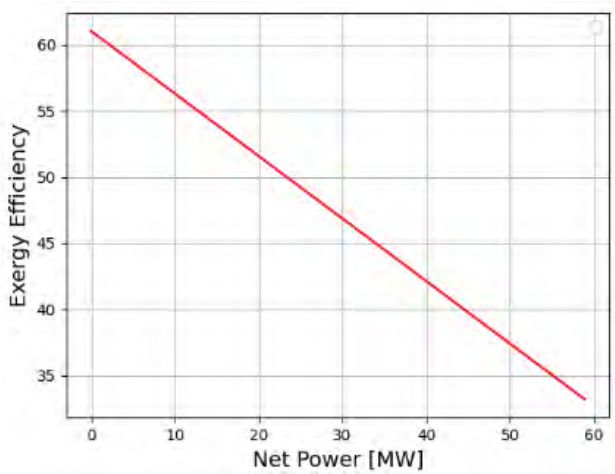

Figure 3. Plant exergy efficiency variation with hydrogen split fraction.

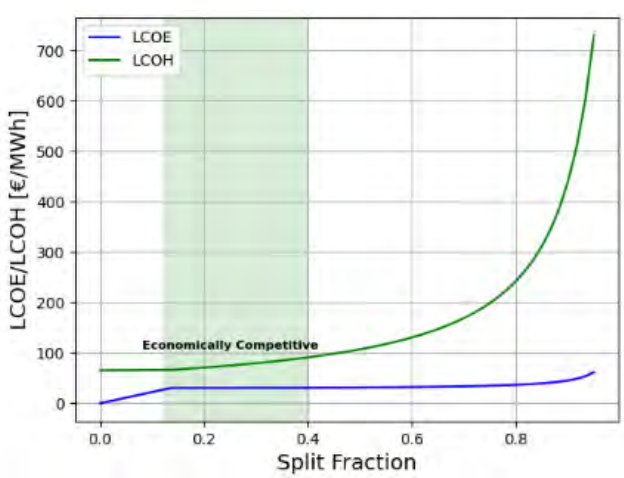

Figure 4. LCOE and LCOH dependence of hydrogen split fraction in the first sensitivity analysis.

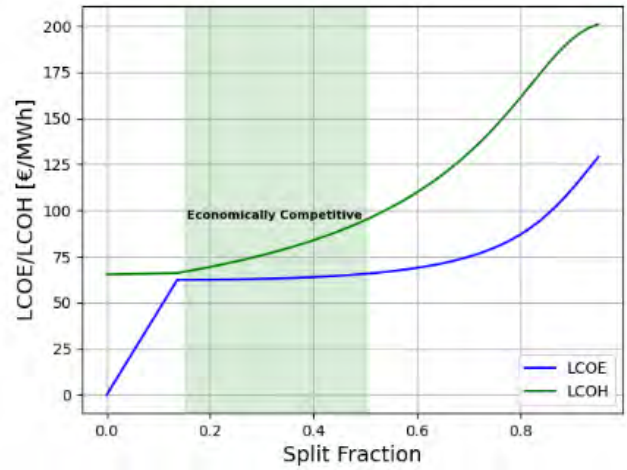

Figure 5. LCOE and LCOH dependence of hydrogen split fraction in the second sensitivity analysis.

Once the steady-state output of the reformer was validated, two successive sensitivity analyses were conducted in ChemCAD, employing the previously mentioned data map and Excel solver configuration. Both runs indicated that as the hydrogen split fraction for electricity generation increases up to about 0.40, the levelized cost of electricity (LCOE) remains essentially constant, varying by less than 1%, while the overall system exergy efficiency declines. This highlights an important compromise between electricity generation and exergy efficiency. The observed linear decrease in exergy efficiency (Figure 3) is expected, as the high chemical exergy of the pure hydrogen stream is inefficiently utilized to generate steam and, subsequently, electricity through combustion, an inherent exergy sink, and a Rankine cycle, operating at constant efficiency across all split fractions.

The optimization results reveal two distinct operational regimes. In the first case (Figure 4), a plateau in electricity cost is observed where LCOE remains below 30.5 €/MWh, while the LCOH steadily increases from 80 to 100 €/MWh in a linear manner. Beyond this split-fraction range, the LCOH rises exponentially up to 700 €/MWh, highlighting a threshold where hydrogen production becomes economically prohibitive. In this regime, the economic burden is primarily borne by hydrogen production, as electricity costs remain stable up to a split fraction of 0.85.

In the second optimization study (Figure 5), the solver was initialized to previously obtained values [11], with both LCOE and LCOH starting around 60-70 €/MWh. In this case, LCOE remains approximately constant across a wider range of hydrogen split fractions (up to about 0.55), while LCOH increases gradually to 100 €/MWh and, at high split fractions, only reaches about 200 €/MWh. These findings indicate that while affordable electricity prices can be maintained across a broad operational range, hydrogen costs are highly

sensitive to the split fraction beyond certain thresholds. In essence, the system preserves economically competitive electricity prices across a wide operating range by progressively transferring the cost burden to hydrogen production as the split fraction increases.

The results further indicate an important trade-off between electricity and hydrogen economics. Allowing for a moderate increase in electricity costs (approximately 50% compared to the first operational regime) enables hydrogen production costs to remain stable over a wider range of hydrogen split fractions. This defines a flexible operating window, extending up to a split fraction of 0.55, within which the plant can simultaneously supply moderately priced electricity and economically competitive hydrogen. The application of exergy-based allocation to the cash flow factor enhances the interpretation of these results by revealing the relative efficiency and economic contribution of each product, thereby supporting more robust design and investment decisions for integrated multi-vector energy systems.

From a global market perspective, the LCOE achieved in the most competitive scenarios (< 30.5 €/MWh) is below typical world averages for newly commissioned renewable generation such as solar PV (about 40–60 €/MWh [15]) and onshore wind (about 42 €/MWh [16]). Even the higher-cost case (about 60–70 €/MWh) remain within the range of real-world LCOE values, indicating that the proposed system is economically feasible and competitive in multiple energy markets.

CONCLUSIONS

This study shows that partial hydrogen splitting can be optimized to balance electricity affordability and hydrogen production costs, highlighting a clear trade-off between electricity and hydrogen economics. Two distinct operational cases were identified, showing that low-cost electricity can be maintained across a range of split fractions while hydrogen costs remain economically competitive. As a key finding, despite the overall variability in LCOH, operating at a hydrogen split fraction between 0.4 and 0.5 provides a reliable and economically competitive range for hydrogen production.

The proposed methodology provides a reproducible and adaptable framework for evaluating multi-product renewable energy systems, which can be extended to alternative hydrogen-to-power pathways, including fuel cells, supporting the development of sustainable and flexible hybrid energy strategies.

ACKNOWLEDGEMENTS

This work was supported by a grant of the Ministry of Education and Research, CCCDI – UEFISCDI, project number: PN-IV-P8-8.1-PRE-HE-ORG-2024-0228, within PNCDI IV.

REFERENCES

1. United Nations Framework Convention on Climate Change. Paris agreement. https://unfccc.int/sites/default/files/english_paris_agreement.pdf
2. Davoodi S, Al-Shargabi M, Wood DA, Rukavishnikov VS, Minaev KM. Review of technological progress in carbon dioxide capture, storage, and utilization. Gas Science and Engineering 117:205070 (2023). https://doi.org/10.1016/j.jgsce.2023.205070
3. Flores F, Feijoo F, DeStephano P, Herc L, Pfeifer A, Dui? N. Assessment of the impacts of renewable energy variability in long-term decarbonization strategies. Applied Energy 368:123464 (2024). https://doi.org/10.1016/j.apenergy.2024.123464
4. Masoudi Soltani S, Lahiri A, Bahzad H, Clough P, Gorbounov M, Yan Y. Sorption-enhanced steam methane reforming for combined CO2 capture and hydrogen production: a state-of-the-art review. Carbon Capture Science & Technology 1:100003 (2021). https://doi.org/10.1016/j.ccst.2021.100003
5. Arias B, Alvarez Criado Y, Méndez A, Marqués P, Finca I, Abanades JC. Pilot testing of calcium looping at TRL7 with co₂ capture efficiencies toward 99%. Energy Fuels 38:14757-14764 (2024). https://doi.org/10.1021/acs.energyfuels.4c02472
6. Fagerström Anton, Al Seadi Teodorita, Rasi Saija, Briseid Tormod. The role of anaerobic digestion and biogas in the circular economy. IEA Bioenergy 2019.
7. Xu J, Froment GF. Methane steam reforming, methanation and water?gas shift: I. intrinsic kinetics. AIChE Journal 35:88-96 (2004). https://doi.org/10.1002/aic.690350109
8. Johnsen K, Grace JR, Elnashaie SSEH, Kolbeinsen L, Eriksen D. Modeling of sorption-enhanced steam reforming in a dual fluidized bubbling bed reactor. Ind. Eng. Chem. Res. 45:4133-4144 (2006). https://doi.org/10.1021/ie0511736
9. Bejan A, Tsatsaronis G, Moran M. Thermal design and optimization. first edition, John Wiley & Sons Hoboken, USA (1995)
10. Smith R. Chemical process design and integration. second edition, Wiley, Hoboken, USA (2016)
11. Báthori AM, Cormos CC. Green hydrogen production by intensified biogas reforming via sorption-enhanced technology: an energy, exergy, economic and environmental (4E) analysis. International Journal of Hydrogen Energy 155:150363 (2025).

https://doi.org/10.1016/j.ijhydene.2025.150363
12. Hyper - Bulk hydrogen production by sorbent enhanced steam reforming. https://hyperh2.co.uk/
13. García-Lario AL, Aznar M, Martinez I, Grasa GS, Murillo R. Experimental study of the application of a nio/nial2o4 catalyst and a cao-based synthetic sorbent on the sorption enhanced reforming process. International Journal of Hydrogen Energy 40:219-232 (2015). https://doi.org/10.1016/j.ijhydene.2014.10.033
14. Greco-Coppi M, Strohle J, Epple B. A carbonator model for CO_2 capture based on results from pilot tests. Part II: deactivation and reaction model. *Chem Eng J* 508:159041 (2025) https://doi.org/10.1016/j.apenergy.2016.12.153
15. Lai CS, McCulloch MD. Levelized cost of electricity for solar photovoltaic and electrical energy storage. Applied Energy 190:191-203 (2017). https://doi.org/10.1016/j.apenergy.2016.12.153
16. Stehly T, Duffy P, Hernando DM, Cost of Wind Energy Review: 2024 Edition. National Renewable Energy Laboratory 2024

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Green Hydrogen Supply Chain Design Towards Social Sustainability: A Case Study in Brazil

Leonardo Santana[a]*, Fernando Pessoa[a], and Ana Barbosa-Póvoa[b]

[a] SENAI CIMATEC University, Department of Computational Modeling, Salvador, Bahia, Brazil
[b] University of Lisbon, Instituto Superior Técnico, Lisbon, Portugal
* Corresponding Author: leosantana049@gmail.com

ABSTRACT

When designing and planning Green Hydrogen Supply Chains (GHSCs), sustainability considerations are increasingly recognized as essential, particularly in light of decarbonization goals and climate policy targets. Existing research has largely focused on economic and environmental however, social sustainability aspects remain significantly underexplored. This work aims to develop a mathematical programming model to design a GHSC, considering simultaneously economic and social aspects. Solar PV, wind power, and PPA (wind) as energy sources are integrated, while transportation options include the construction of new pipelines, compared to the use of existing highways for trucks carrying liquefied or compressed hydrogen to deliver hydrogen to an oil refinery. The model is applied to a case study conducted in the Brazilian state of Bahia, where different social indicators will be explored, characterizing the case study context while allowing generalization to other contexts. Results allow us to establish a trade-off analysis between economic and social concerns, offering valuable insights for both policymakers and companies, supporting strategies that strike a balance between economic competitiveness and regional social development.

Keywords: Supply Chain, Hydrogen, Optimization, Social Sustainability, Brazil case study

INTRODUCTION

Green hydrogen produced without CO_2 emissions, namely electrolytic hydrogen, has emerged as a key enabler for the decarbonization of hard-to-abate sectors [1]. It is widely recognized in literature as a promising solution to support the achievement of the United Nations Sustainable Development Goals (SDGs). However, despite its potential, the large-scale deployment of green hydrogen still requires policy support and subsidies to overcome cost barriers and enable market uptake [2]. In this context, national hydrogen strategies and roadmaps have been proposed and continuously updated to facilitate the development and deployment of green hydrogen supply chains (GHSCs) [3].

Nevertheless, the design and planning of GHSCs must go beyond purely techno-economic considerations and explicitly account for sustainability metrics. Although several studies have proposed mathematical programming models for GHSC design, relatively few have critically examined how environmental, economic, and social dimensions are jointly addressed within these frameworks [4]. In particular, a significant gap persists in the assessment of social sustainability, as social aspects remain largely underrepresented in optimization-based formulations, often treated qualitatively or excluded altogether [4, 5].

This gap is especially concerning given that the development of green hydrogen is explicitly framed as a pathway toward sustainable and inclusive energy transitions. Beyond reducing greenhouse gas emissions, such transitions must also proactively address potential social risks, including neocolonial patterns of resource extraction, uneven distribution of economic benefits, exclusion of local communities, water scarcity, and land-use conflicts. Recent studies emphasize that the integration of energy justice principles is essential for the equitable and socially robust development of GHSCs [6].

Accordingly, socio-economic factors, encompassing income levels, employment generation, fiscal contributions, subsidies, and broader social development indicators such as the Human Development Index, have

emerged as critical elements in sustainability assessments [7]. Despite their relevance, these dimensions are rarely incorporated explicitly into optimization-based HSC design models, remaining largely disconnected from strategic planning and investment decisions [4].

In the context of a developing country such as Brazil, the inclusion of a social sustainability assessment is particularly important. While Brazil seeks to position itself as a competitive producer of low-cost green hydrogen, decision-making processes must also account for regional inequalities and development needs [8]. These objectives are often conflicting, highlighting the need for integrated modeling approaches. To address this challenge, this study proposes a bi-objective mixed-integer linear programming (MILP) model for the strategic design of a GHSC in Brazil, explicitly integrating economic performance and location-dependent socio-economic indicators into the decision-making process.

LITERATURE REVIEW

Kim and Moon (2008) were the first to address a social criterion for HSC design using MILP, by considering cost and safety in the design of hydrogen infrastructure [9]. Since then, the models normally use safety to access the social aspects of the HSC, optimizing risk-related factors. This risk-based interpretation of the social dimension has become the dominant perspective in the literature [4].

Almaraz et al. (2013) was one of the precursors to the proposal of sustainable HSC, highlighting the importance of considering the triple bottom line in sustainability in the strategic stage of any new design to ensure system efficiency. However, the social aspects are restricted to the risks linked to the activities of each stage of the HSC [10].

Fazli-Khalaf et al. (2020) implemented in their modeling the maximization of social responsibility, considering among their criteria the generation of employment opportunities, reduction of health damage, and promotion of sustainable and balanced development between regions. Despite this broader interpretation of social sustainability, the adopted indicators remain highly aggregated and are not explicitly linked to regional socio-economic conditions or location-specific vulnerability [11].

In the Brazilian context, recent MILP-based studies on HSC optimization, such as da Silva et al. (2024), Santana et al (2025), and Camelo et al. (2025), focus primarily on economic and technical performance, without explicitly incorporating social or socio-economic criteria [12-14]. This highlights a persistent gap in the integration of location-dependent social indicators into strategic HSC design models.

To address this gap, the present study incorporates four socio-economic indicators into a bi-objective MILP framework, explicitly linking regional socio-economic conditions to strategic supply chain design decisions.

PROBLEM STATEMENT

As a deterministic capacitated facility location problem, this work aims to design optimal investments that consider both economic and social aspects in the GHSC to meet the demand of Bahia's oil refinery, as shown in the superstructure in Figure 1.

Given:

(i) A predefined hydrogen demand cluster j (representing the refinery).

(ii) Possible facility locations i (12 locations in Bahia).

(iii)Scenarios involving energy sources e (solar, wind, or PPA), production technology p (PEM), and transportation methods l (liquid tanker, tube trailer, or pipelines).

Determine the optimal:

(i) Hydrogen production plant location.

(ii)The transportation method for the production plant to the demand cluster.

Objective: Minimize the Levelized Cost of hydrogen (LCOH) while maximizing social impact (SOCimp) for different social indexes.

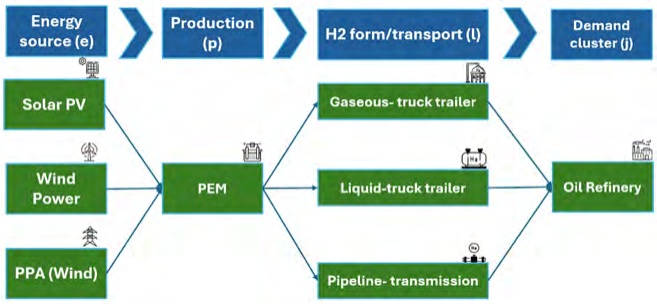

Figure 1. Superstructure of the Green Hydrogen Supply Chain Model.

The demand cluster requires 200 MW (29 kt/year) of hydrogen, according to a previous study [14]. The geographic disposition of the candidate site productions and the demand cluster is shown in Figure 2. Energy sources include solar PV, Wind power, and purchase power agreements of wind farms (PPA-wind). A hub-based configuration is adopted, in which a single production site is selected to supply the entire demand cluster. This assumption reflects the economies of scale and allows isolating location-dependent trade-offs between cost and social performance, which is essential for assessing the influence of regional socio-economic indicators in strategic decision-making.

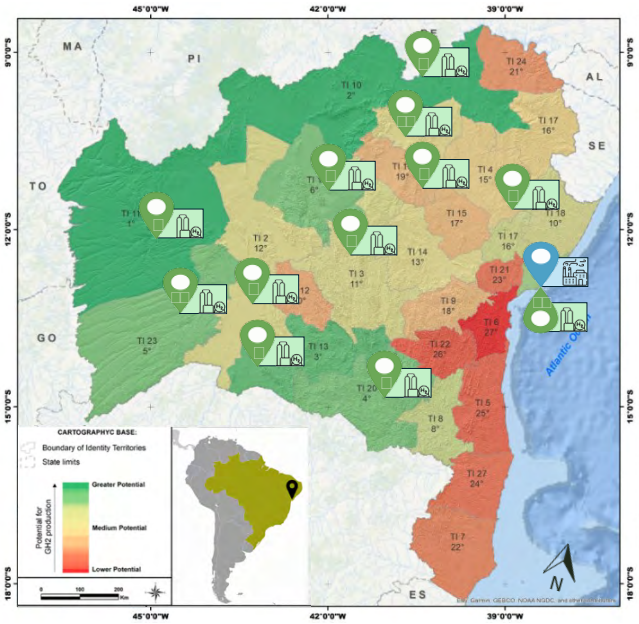

Figure 2. Geographic Location of Candidate Production Sites (green) and Demand Cluster (blue) in Bahia, Brazil.

METHODOLOGY

The study follows three main steps. First, a bi-objective MILP model is formulated to represent strategic decisions on hydrogen production site selection and transport mode selection, minimizing economic costs and maximizing social impact. Second, the model is parametrized using local socio-economic data from the Green Hydrogen Atlas of Bahia and complementary statistical sources [15]. At the same time, technical and cost coefficients are obtained from peer-reviewed literature. Finally, the optimization problem is implemented in Python and solved using the Gurobi solver, applying the ε-constraint method to generate discrete Pareto fronts.

OPTIMIZATION FRAMEWORK

Mathematical modeling

A basic mathematical formulation of this work is based on previous models by Santana et al. (2025), was developed just to test the social component of such networks. The social objective is inspired by the work of Mota et al. (2018) [16, 17].

Sets:

i: Production site location

j: Demand cluster location

e: Energy source type

p: Production technology type

l: Transport mode

Parameters:

$a^{p,l}$: Literature parameter for production technology p and energy source e, or transport mode l.

$b^{p,l}$: Literature parameter for production technology p and energy source e, or transport mode l.

$c^{p,l}$: Literature parameter for production technology p and energy source e, or transport mode l.

d^{l}: Literature parameter for transport mode l.

EP^{e}: Electricity price of energy source e

L^{l}_{ij}: Distance between production plant i and demand cluster j per transport mode l.

$FTE^{e,p,l}$: Full time employment per energy source e, or production technology p, or transport mode l.

SV^{si}: Regional social index value for index s

$SV^{s,ref}$: Reference social index value for index s

μ^{s}_{i}: Regional factors based on social indexes s

Variables:

x_i: Plant capacity of production plant i (MW)

Y_i: Binary variable, 1 if the site i is activated, 0 otherwise.

T_{li}: Binary variable, 1 if the transportation mode l is activated at location i, 0 otherwise.

In this specific case study, as only one production location may be selected among all candidates, variable x_i becomes $x_i = 200Y_i$.

Equations:

The model considers that the flow between production site and demand cluster is equal to the plant capacity x_i, thus the levelized cost of hydrogen for producing at site i and transport to demand cluster j, are:

$$LCOH^{p}(x_i) = \sum_{i=1}^{n}(a^{p}x_i + b^{p}EP^{e}Y_i + c^{p}Y_i) \quad (1)$$

$$LCOH^{l}(x_i) = \sum_{i=1}^{n}(a^{l}L^{l}_{ij}T_i + b^{l}T_i + c^{l}x_i + d^{l}T_i) \quad (2)$$

Regional social factor definition:

$$\mu^{s}_{i} = \frac{SV^{s,ref}}{SV^{si}} \quad (3)$$

Social impact of installing a production plant in a location i:

$$SOCimp_i = \mu^{s}_{i}x_i(FTE^{e} + FTE^{p} + FTE^{l}L^{l}_{ij}) \quad (4)$$

Constraints:

Hydrogen production from all selected plants must equal the supply chain demand:

$$\sum_i x_i = D \quad (5)$$

Table 1: Regression Parameters for Levelized Cost of Hydrogen (LCOH) Calculation.

	a	b	c	d	Ref
p (PEM)	-0.0249	13.3604	2.1039	0	[17]
l (Pipeline)	0.2176	-7.105x10^-15	0	0	[12]
l (Liq truck)	0.0006	0.1203	-0.0078	3.5535	[18]
l (Gas truck)	0.0031	0.3231	0	0	[18]

Table 2: Employment Factors (FTE) for Different Supply Chain Echelons [19].

	Employment factor (FTE)	Unit
e (Wind)	0.4	FTE/MW
e (Solar)	0.12	FTE/MW
e (PPA-Wind)	0.0	FTE/MW
p (PEM)	0.9	FTE/MW
l (Pipeline)	0.04	FTE/km
l (Liq truck)	2	FTE/(GW*km)
l (Gas truck)	0.3	FTE/(GW*km)

Table 3: Candidate Production Sites: Characteristics, Distances, and Socio-Economic Indicators [15;20-21].

Site (i)	Location	Energy source (e)	Transport mode distance (km)		GDP (kBRL per inhabitant)	IFDM	SoVI	Salary (BRL)
			Pipe	Road				
1	Irecê	Wind	407.4	443	10.011	0.534	0.436	1652
2	BJ Lapa	Solar	553.35	743	10.095	0.533	0.467	1652
3	Seabra	Wind	364.35	430	10.466	0.546	0.459	1790
4	Serrinha	Wind	130.2	149	9.348	0.536	0.465	1652
5	Juazeiro	Wind	438.9	472	15.551	0.525	0.427	1790
6	Barreiras	Solar	735	828	37.867	0.593	0.352	1937
7	Guanambi	Wind	508.2	640	13.946	0.608	0.379	1652
8	Jacobina	Wind	283.5	305	10.842	0.513	0.447	1790
9	VConquista	PPA	360.15	481	15.124	0.550	0.396	1790
10	SM Vitoria	Solar	644.7	832	18.674	0.547	0.405	1790
11	S Bonfim	Wind	319.2	349	12.213	0.532	0.434	1937
12	SF Conde	PPA	0	0	30.8183	0.653	0.3536	2522
		Reference value (Bahia's average)			*19.54*	*0.5564*	*0.419*	*1865*

Hydrogen production is an integer variable:

$$x_i \in \mathbb{Z}^+, \forall i \quad (6)$$

Hydrogen production is constrained by its maximum capacity and is dependent on whether the plant is selected:

$$x_i \leq D\, Y_i, \forall i \quad (7)$$

Binary decision variables Y_i and T_{li} indicates whether a specific production site and transport mode is selected:

$$Y_i \in \{0,1\}, \forall i \quad (8)$$

$$T_{li} \in \{0,1\}, \forall i \quad (9)$$

Only one transportation mode may be selected if a location is active:

$$\sum_i t_{il} = Y_i, \forall i \quad (10)$$

Objective functions:

Minimize the Supply chain's LCOH:

$$Min\ LCOH = LCOH^p(x_i) + LCOH^l(x_i) \quad (12)$$

Maximize the Supply chain's social impact:

$$Max\ SOCimp\ (x_i) = \sum_{i=1}^{n} x_i \mu_i^s (FTE^e + FTE^p + FTE^l L_{ij}^l) \quad (13)$$

Multiobjective optimization strategy:

The bi-objective problem is solved using the ε-constraint method. The economic objective is minimized to obtain the minimum economic cost and the corresponding social impact. Then, the social objective is maximized to obtain the maximum attainable social impact. So, a set

of ε values is generated between these two extremes, and the economic objective is minimized subject to:

$$SOCimp \geq \varepsilon_k \quad (14)$$

In this study, a discrete Pareto frontier is constructed using five ε values uniformly distributed between the economic and social extremes.

Model parametrization

The literature parameters: $a^{p,l}$; $b^{p,l}$; $c^{p,l}$; d^{l}, are displayed in table 1. The production parameters are derived from Santana et al. (2024), which quantified the impact of various variables on the levelized cost of hydrogen (LCOH) for different plant capacity scales [17]. This parameter set incorporates costs associated with electricity, water pre-treatment, PEM electrolyzer system, initial compression, and gaseous storage buffering.

Based on the published data, a bilinear function was regressed to estimate the LCOH as a function of two key variables: plant capacity (MW) and electricity price (US$/MWh). This formulation captures the main influencing factors identified in the reference study. In the present model, the electricity price is an exogenous variable, fixed and defined specifically for each location i.

The parameters for transportation and subsequent conditioning costs (gaseous or liquefied form) were sourced from the meta-analysis by Santos (2024) [18]. This work provides a comprehensive review of hydrogen production, storage, and transportation technologies, detailing their cost and greenhouse gas emission profiles. Given that initial compression costs are already internalized within the production parameters, the transportation and conditioning costs applied here refer to subsequent logistical steps. Originally, non-linear or scale-incompatible parameters were adapted and linearized where necessary to align with the scope and structure of the current model.

The employment parameters were sourced from Ganter et al. (2024), who synthesized the job creation potential for each echelon of the hydrogen supply chain. Their synthesis was based on a combination of data from academic publications, industry sources, and anonymous industry surveys and is presented in table 2 [19]. For PPA-wind is assumed to have no Job creation once the energy production isn't in the same region of hydrogen's.

The selection of potential plant locations i was based on the results of the Green Hydrogen Atlas of Bahia study [15]. This study applied an Analytic Hierarchy Process (AHP)-based multicriteria analysis to rank the most suitable regions within the Brazilian state of Bahia for green hydrogen production. The selection of the energy source type e for each locality i was performed by integrating data from the Atlas with information from the Brazilian National Electric System Operator [20].

Due to budget constraints, a single plant installation is to be considered in this case. Also, the technology to be installed is a single one, and it will be an electrolyser.

The distances between production locations i and demand points j were calculated as follows: for pipeline transport, straight-line distances were obtained via Google Earth, with a 5% markup applied to account for potential route deviations. For truck transport (liquid or gaseous), road distances were acquired using Google Maps, and the electricity price is based on literature data, being $0.05 per kWh for solar PV; $0.03 per kWh for onshore wind power and assume as premised by Mughal and Sajid, 2025 the electricity cost of PPA-Wind equal to solar PV ($0.05 per kWh) [21, 22].

The data for social indices were either directly obtained or calculated from the values presented in the Green Hydrogen Atlas for the Municipal Development Index (IFDM), Social Vulnerability Index (SoVI), and Gross Domestic Product (GDP) [15]. Income data, however, were sourced directly from the Superintendency of Economic and Social Studies of Bahia (SEI) [23].

Table 3 presents the selected plant locations, distances, and social indices.

RESULTS AND DISCUSSION

The results highlight a trade-off between economic efficiency and social sustainability in the design of a GHSC in Bahia. Using an ε-constraint multi-objective MILP formulation with a single-sourcing structure, the resulting Pareto fronts are inherently discrete, with each non-dominated solution corresponding to a specific combination of production site, energy source, and transport mode.

Table 4 presents the trade-off between levelized cost (LCOH) and social impact when the Municipal GDP is used as the social indicator. The Pareto-efficient solutions range from the cost-optimal site (4-Serrinha) to the socially optimal one (6-Barreiras), illustrating the additional cost associated with investing in municipalities with lower economic activity to boost local development.

Table 4: Pareto-Optimal Solutions for the GDP-based Social Indicator.

Site (i)	Energy source (e)	Transport mode (l)	LCOH ($/kg)	Social impact (GDP)
4	Wind	Gas truck	1.92	128
5	Wind	Liq truck	3.54	357
6	Solar	Liq Truck	6.42	1037

For the Municipal Development Index (IFDM), Table 5 summarizes the Pareto-optimal solutions obtained when maximizing its social impact. The sequence of selected sites demonstrates how the model shifts

investment towards locations with greater infrastructural and social needs as the social constraint is tightened, with a corresponding increase in LCOH.

Table 5: Pareto-Optimal Solutions for the IFDM-based Social Indicator.

Site (i)	Energy source (e)	Transport mode (l)	LCOH ($/kg)	Social impact (IFDM)
4	Wind	Gas truck	1.92	259
8	Wind	Liq truck	3.44	352
3	Wind	Liq Truck	3.51	424
7	Wind	Liq truck	3.64	564
6	Solar	Liq Truck	6.42	571

In turn, Table 6 details the compromise solutions identified when the optimization incorporates the minimization of Social Vulnerability (SoVI). The results show a progression towards sites with higher social fragility, underscoring the potential role of green hydrogen projects in addressing regional vulnerabilities, although at a higher production cost.

Table 6: Pareto-Optimal Solutions for the SoVI-based Social Indicator.

Site (i)	Energy source (e)	Transport mode (l)	LCOH ($/kg)	Social impact (SoVI)
4	Wind	Gas truck	1.92	298
8	Wind	Liq truck	3.44	408
3	Wind	Liq truck	3.51	473
2	Solar	Liq truck	6.37	559

Finally, the Pareto-optimal configurations for the social indicator based on the local Average Salary are presented in Table 7. The solutions reveal a trade-off where selecting sites with lower average wages (aiming to improve local income) leads to higher supply chain costs, highlighting a key equity-efficiency tension in strategic planning.

Table 7: Pareto-Optimal Solutions for the Salary-based Social Indicator.

Site (i)	Energy source (e)	Transport mode (l)	LCOH ($/kg)	Social impact (Salary)
4	Wind	Gas truck	1.92	238
8	Wind	Liq truck	3.44	366
11	Wind	Liq truck	3.46	415
6	Solar	Liq truck	6.42	555

Across all social indicators, the economically optimal solution consistently selects Serrinha (site 4), supplied by wind power and connected to the refinery via compressed hydrogen trucks (Gas truck), achieving the lowest LCOH of 1.92 USD/kgH_2. This configuration benefits from short transport distances and low electricity costs but yields the lowest social impact. These results indicate that, in the case of Bahia, cost-minimizing strategies tend to concentrate investments in already consolidated regions, offering limited contributions to broader regional social development goals.

As social requirements become more restrictive, the model shifts the production location and systematically adopts liquefied hydrogen transport (Liq truck), despite its higher cost, due to its significantly higher FTE factor per distance. Intermediate Pareto-efficient solutions, such as Juazeiro (site 5, Table 4), Jacobina (site 8, Tables 5, 6, 7), and Seabra (site 3, Tables 5 and 6), present LCOH values between 3.4 and 3.6 USD/kgH_2, while delivering substantially higher social impacts.

On the other hand, the socially optimal solutions for indicators like GDP, IFDM, and Salary are consistently associated with Barreiras (site 6, powered by solar PV and relying on Liq truck), achieving the highest social impact scores (Tables 4, 5, 7). However, these configurations present the highest LCOH, approximately 6.4 USD/kgH_2, reflecting longer transport distances and higher electricity costs. This outcome reveals a structural tension: regions with the greatest potential for social benefits are often those facing logistical and energy cost disadvantages, limiting their competitiveness under purely market-driven conditions.

Although each social indicator emphasizes different dimensions of local development, consistent qualitative patterns emerge. GDP-based results favor municipalities with lower economic activity, while IFDM and Salary indicators shift solutions toward regions with greater development needs. The Social Vulnerability Index (SoVI) further reorients the Pareto frontier toward socially fragile municipalities, such as Bom Jesus da Lapa (site 2, Table 6), even at the expense of higher hydrogen production costs.

From a policy perspective, the results suggest that achieving socially inclusive GHSCs requires targeted interventions, such as regional incentives, subsidies, or infrastructure investments in inland regions. Without such measures, cost-driven solutions are likely to prevail, limiting the social benefits of hydrogen deployment. For supply chain planners and investors, the intermediate Pareto-efficient solutions detailed in the tables offer viable compromise strategies that balance economic competitiveness with measurable social contributions, aligning with emerging ESG-oriented investment frameworks.

CONCLUSION

This work proposed a bi-objective MILP framework

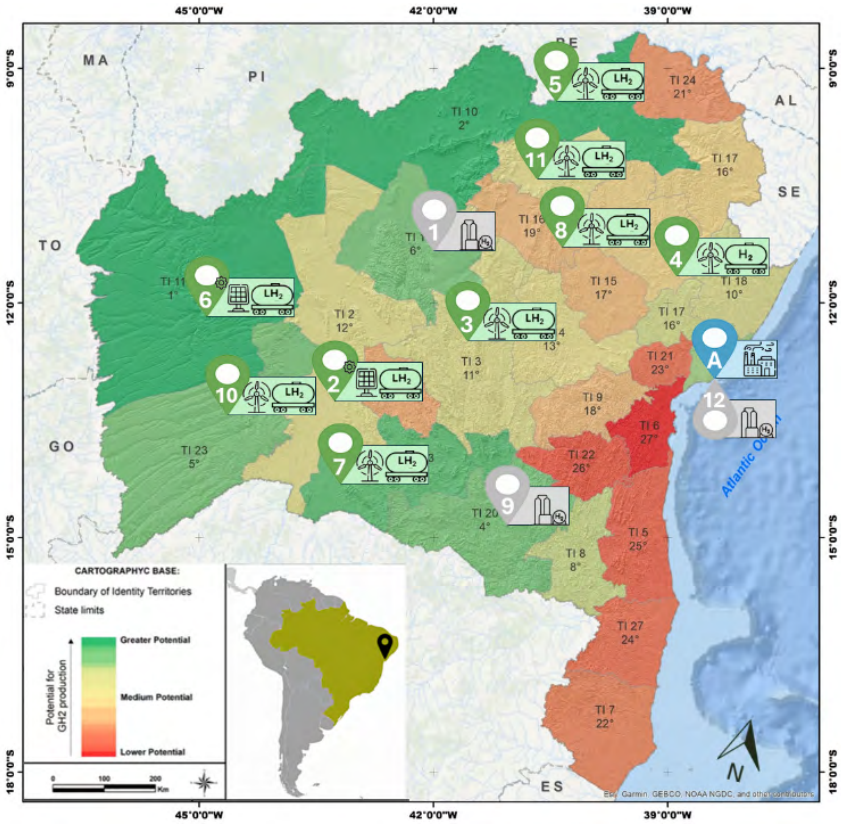

Figure 3. Snapshot of the Pareto-optimal site selections for green hydrogen production in Bahia.

for the strategic design of GHSCs that explicitly integrates economic performance and location-dependent social sustainability indicators. By incorporating regional socio-economic metrics into the optimization process, the model enables a transparent assessment of trade-offs between cost efficiency and social impact. The case study in Bahia, Brazil, demonstrates that cost-minimizing solutions tend to favor already developed regions, while socially oriented solutions shift investments toward inland municipalities with greater development needs, at the expense of higher hydrogen costs. The resulting Pareto fronts provide actionable insights for policymakers and investors by identifying compromise solutions that balance competitiveness and social contribution. Although the analysis is restricted to a single production hub, the proposed framework is flexible and can be extended to multi-facility, multi-period, and uncertainty-aware formulations, supporting broader applications in sustainable hydrogen planning. The model proposed is a basic model that will be further developed to include further detailed characteristics of the GHSC, such as fugitive emissions and multiple demand clusters and technologies for hydrogen production.

ACKNOWLEDGEMENTS

The authors acknowledge Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES), SENAI-CIMATEC University, Instituto Superior Técnico de Lisboa, for technical and financial support, and to CEGIST (FCT - Foundation for Science and Technology, I.P., under the projects UIDB/00097/2025 and UIDP/00097/2025 (CEGIST).

AUTHOR IDENTIFIERS

Author ORCIDs:
Santana L.: 0000-0001-9087-8594
Pessoa F.: 0000-0003-3667-3875
Barbosa-Póvoa A.: 0000-0001-6594-9653

REFERENCES

1. Mansilha C, Barbosa-Póvoa A, Tarelho L, Fonseca A. A comprehensive review of green hydrogen production technologies: current status, challenges, research trends and future directions. Renewable and Sustainable Energy Reviews 225:116119 (2026). https://doi.org/10.1016/j.rser.2025.116119
2. Martins FP, De-León Almaraz S, Botelho Junior AB, Azzaro-Pantel C, Parikh P. Hydrogen and the sustainable development goals: synergies and trade-offs. Renewable and Sustainable Energy

Reviews 204:114796 (2024). https://doi.org/10.1016/j.rser.2024.114796
3. Sgarbossa F, Arena S, Tang O, Peron M. Renewable hydrogen supply chains: a planning matrix and an agenda for future research. International Journal of Production Economics 255:108674 (2023). https://doi.org/10.1016/j.ijpe.2022.108674
4. Oliveira VESF, do Carmo BBT, da Costa Fontes FF. Hydrogen supply chains: modeling challenges and sustainability gaps. Applied Energy 401:126827 (2025). https://doi.org/10.1016/j.apenergy.2025.126827
5. De-León Almaraz S, Azzaro-Pantel C. Integration of techno-economic, environmental, and social perspectives in the hydrogen economy. Accelerating the Transition to a Hydrogen Economy :3-32 (2025). https://doi.org/10.1016/b978-0-443-24002-7.00015-x
6. Patonia A. Green hydrogen and its unspoken challenges for energy justice. Applied Energy 377:124674 (2025). https://doi.org/10.1016/j.apenergy.2024.124674
7. De-León Almaraz S, Kocsis T, Azzaro-Pantel C, Szántó ZO. Identifying social aspects related to the hydrogen economy: review, synthesis, and research perspectives. International Journal of Hydrogen Energy 49:601-618 (2024). https://doi.org/10.1016/j.ijhydene.2023.10.043
8. Bisognin Garlet T, de Souza Savian F, Duarte Ribeiro JL, Mairesse Siluk JC. Unlocking brazil's green hydrogen potential: overcoming barriers and formulating strategies to this promising sector. International Journal of Hydrogen Energy 49:553-570 (2024). https://doi.org/10.1016/j.ijhydene.2023.09.227
9. Kim J, Moon I. Strategic design of hydrogen infrastructure considering cost and safety using multiobjective optimization. International Journal of Hydrogen Energy 33:5887-5896 (2008). https://doi.org/10.1016/j.ijhydene.2008.07.028
10. De-León Almaraz S, Azzaro-Pantel C, Montastruc L, Pibouleau L, Senties OB. Assessment of mono and multi-objective optimization to design a hydrogen supply chain. International Journal of Hydrogen Energy 38:14121-14145 (2013). https://doi.org/10.1016/j.ijhydene.2013.07.059
11. Fazli-Khalaf M, Naderi B, Mohammadi M, Pishvaee MS. Design of a sustainable and reliable hydrogen supply chain network under mixed uncertainties: a case study. International Journal of Hydrogen Energy 45:34503-34531 (2020). https://doi.org/10.1016/j.ijhydene.2020.05.276
12. Nascimento da Silva G, Lantz F, Rochedo PRR, Szklo A. Developing and applying the hydrogen economics and infrastructure optimization model (HERA). International Journal of Hydrogen Energy 61:1170-1186 (2024). https://doi.org/10.1016/j.ijhydene.2024.02.015
13. Camelo MM, de Andrade CF, Prata BA. A mixed-integer linear programming model for optimizing green hydrogen supply chain networks. International Journal of Hydrogen Energy 118:134-145 (2025). https://doi.org/10.1016/j.ijhydene.2025.02.138
14. Santana LOS, Santos GS, Pessoa FLP, Barbosa-P?voa AP. The green hydrogen supply chain in the brazilian state of bahia: a deterministic approach. Systems and Control Transactions 4:2038-2043 (2025). https://doi.org/10.69997/sct.185907
15. Mota B, Gomes MI, Carvalho A, Barbosa-Povoa AP. Sustainable supply chains: an integrated modeling approach under uncertainty. Omega 77:32-57 (2018). https://doi.org/10.1016/j.omega.2017.05.006
16. Santos, G. Modelo para Tomada de Decisão das Rotas de Produção do Hidrogênio. Universidade SENAI CIMATEC. 2024. 10.13140/RG.2.2.22624.80642
17. Santana L, Santos G, Santos A, Marinho C, Bispo A, Villardi H, Pessoa F. Evaluating the economic influence of water sources on green hydrogen production: a cost analysis approach. International Journal of Hydrogen Energy 89:353-363 (2024). https://doi.org/10.1016/j.ijhydene.2024.09.274
18. Ganter A, Lonergan KE, Büchi HM, Sansavini G. Shifting to low-carbon hydrogen production supports job creation but does not guarantee a just transition. One Earth 7:1981-1993 (2024). https://doi.org/10.1016/j.oneear.2024.10.009
19. SENAI-CIMATEC. Atlas H2V Bahia. (2023).
20. ONS.https://www.ons.org.br/Paginas/resultados-da-operacao/historico-da-operacao/fator-capacidade.aspx
21. Our world in data. https://ourworldindata.org/grapher/levelized-cost-of-energy?country=~BRA
22. Mughal SUR, Sajid MB. Energy-economic-environment (3E) analysis of green hydrogen production: comparing power purchase agreement and onsite renewable power. International Journal of Hydrogen Energy 175:151467 (2025). https://doi.org/10.1016/j.ijhydene.2025.151467
23. SEI. https://www.ba.gov.br/sei/bahia

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Sustainable Design of an Integrated Seawater-Based Green Hydrogen Production Process

Antonio Torres-Ayala[a], Eduardo Sánchez-Ramírez[a], Marcelino Carrera-Rodríguez[b] and Juan Gabriel Segovia-Hernández[s]*

[a] Universidad de Guanajuato, Department of Chemical Engineering, Guanajuato, Gto., México
[b] Instituto Politécnico Nacional, Unidad Profesional Interdisciplinaria de Ingeniería Campus Guanajuato, Silao de la Victoria, Gto., México
* Corresponding Author: gsegovia@ugto.mx.

ABSTRACT

Green hydrogen constitutes a strategic energy vector for achieving the Sustainable Development Goals (SDGs 7, 9, 12, and 13) due to its high energy density, flexibility for renewable energy storage, and direct emission-free operation. However, its production critically depends on the supply of high-purity water, which is unsustainable in the context of a projected 40% global water deficit by 2030. Given that more than 97% of available water is saline, integrating desalination processes with electrolysis constitutes an essential strategy for transitioning toward circular economy models in water resource management. This work presents the conceptual design, detailed modeling, and optimization of an integrated process for the sustainable production of green hydrogen from saline water. The system couples a desalination technology (Solar Distillation) with two electrolysis technologies (AEL and SOEC), modeled through physicochemical, electrochemical, and thermodynamic principles. The objective is to determine technological configurations, materials, and operating conditions that maximize the energy efficiency and minimize the LCOH, contributing to the development of viable routes within the energy transition and circular economy. The results shows that the viability of a integrated Seawater Green Hydrogen system with a detail model is possible. In contrast to black-box models, this model yielded a detailed, geometrically accurate simulation that captures phenomenological effects. It enables the development of processes that contribute to the energy transition, reduce freshwater consumption, and assess their economic viability. The sensitivity analysis revealed promising solutions, with LCOH values ranging from 4.22 to 7.10 USD/kg and overall energy efficiencies between 67% and 82%.

Keywords: Hydrogen, Energy, Renewable and Sustainable Energy, Process Design, Optimization, Sensitivity Analysis, Desalinisation

1. INTRODUCTION

Green hydrogen, produced through electrolysis powered by renewable energy, represents a key technology for decarbonizing hard-to-abate industrial and energy sectors, in accordance with SDGs 7 (Affordable and Clean Energy) and 13 (Climate Action) [1, 2]. However, its production depends on the supply of ultrapure water, which poses a critical challenge in the context of water stress and the need to transition toward systems that promote circular water use.

By 2030, the European Union aims to reach an annual hydrogen production of 10 million metric tons. From a stoichiometric standpoint, this target demands 89 million metric tons of water, excluding supplementary requirements like cooling systems [3]. For comparison, the average domestic water consumption in nations like Germany and the United Kingdom is currently 120 and 140 liters per capita per day (lcd), respectively [4]. Based on the current EU population (~450 million) [5], the total domestic water use amounts to 19, 700 million metric tons annually. Therefore, scaling up hydrogen production would introduce an additional water demand of 0.45 % per year. Although this increase might not pose a systemic threat to Europe's total water reserves, it constitutes a substantial volume. This demand becomes

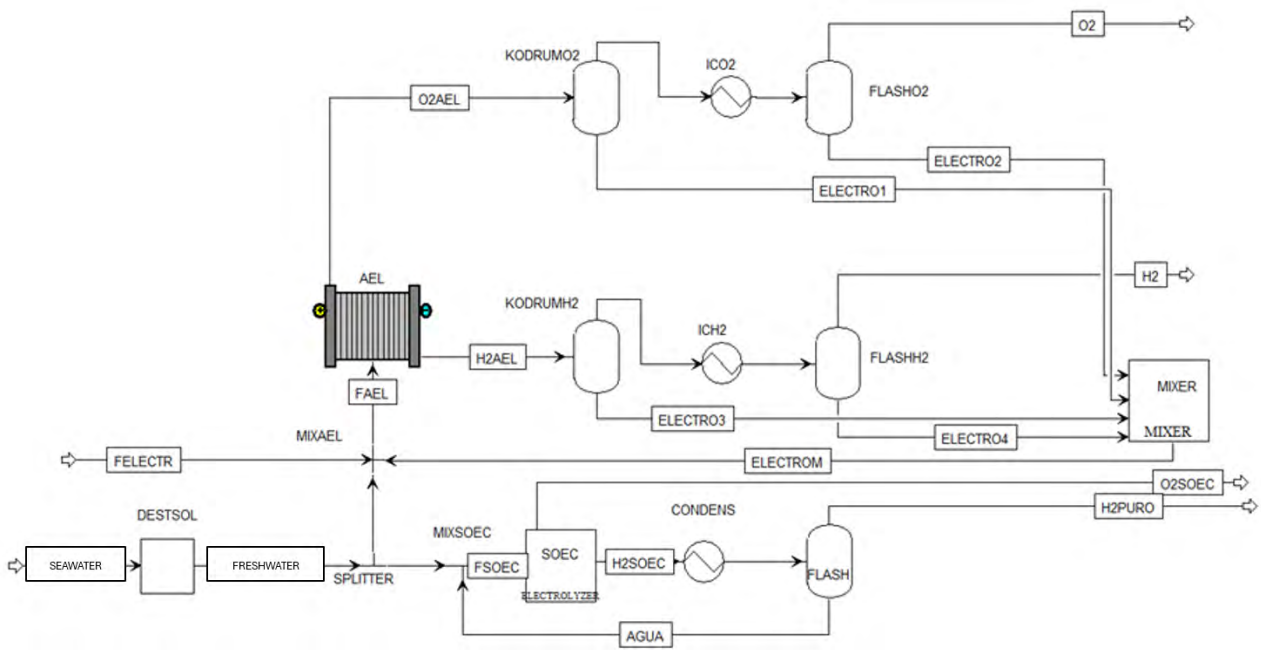

Figure 1. Hydrogen production process flowsheet.

particularly challenging in decentralized, water-stressed regions with limited freshwater availability, making the integration of seawater desalination a strategic necessity for sustainable deployment in these regions.

Integrating desalination technologies with advanced electrolysis systems avoids the use of potable water and improves the sustainability of the energy-hydrogen–water cycle [2, 6].

While desalination and electrolysis are mature technologies individually, their thermal and mass integration presents non-trivial challenges. Most technoeconomic studies rely on "black-box" models that fix the efficiency of the electrolyzer regardless of operating conditions. However, the specific energy consumption of an Alkaline Electrolyzer (AEL) is sensitive to the void fraction generated by bubbles at high current densities, a phenomenon that cannot be captured without a phenomenological approach. Furthermore, traditional experimental optimization is time-consuming and capital-intensive, limiting the ability to screen novel electrode materials or geometric configurations.

This analysis forms the basis for future multi-objective optimization steps, addresses the geometric gaps by developing a rigorous simulation framework that couples electrochemical kinetics with a detailed thermodynamic property package (Aspen Plus) and a desalinisation system. Unlike standard models, this tool allows for the "virtual prototyping" of the cell simultaneously varying geometric parameters (gap, area) and electrolyte concentrations, using both empirical correlation and fundamental equations-based optimization of the coupled water-hydrogen nexus before physical piloting.

2. METHODOLOGY

The process simulation strategy relies on a hybrid architecture. The mass and energy balances, along with the electrochemical equations, are solved in a Python environment, which dynamically queries Aspen Plus V14 to retrieve physicochemical properties using the [e-NRTL / Pitzer] method.

The process consists of a coupled desalination-electrolysis system, beginning with a seawater desalination unit. Subsequently, a splitter separates the water supply into an AEL line and a SOEC line, as shown in Figure 1. After each electrolyzer produces hydrogen, the separation units purify and recirculate the water and electrolyte.

2.1 Electrochemical Modeling

Mathematical models were developed in an open-source language for Alkaline (AEL) electrolysis technology and Solid Oxide technology, incorporating mass and energy balances, electrochemical models for overpotentials, ionic transport, and equations for thermodynamic behavior under different materials, and operating conditions.

The cell voltage (V_{cell}) is the primary performance indicator in this type of analysis, and is calculated by summing the reversible potential and the overpotentials, as defined in Eq. (1) reported by Abdin et al. (2017) [7] report in their model:

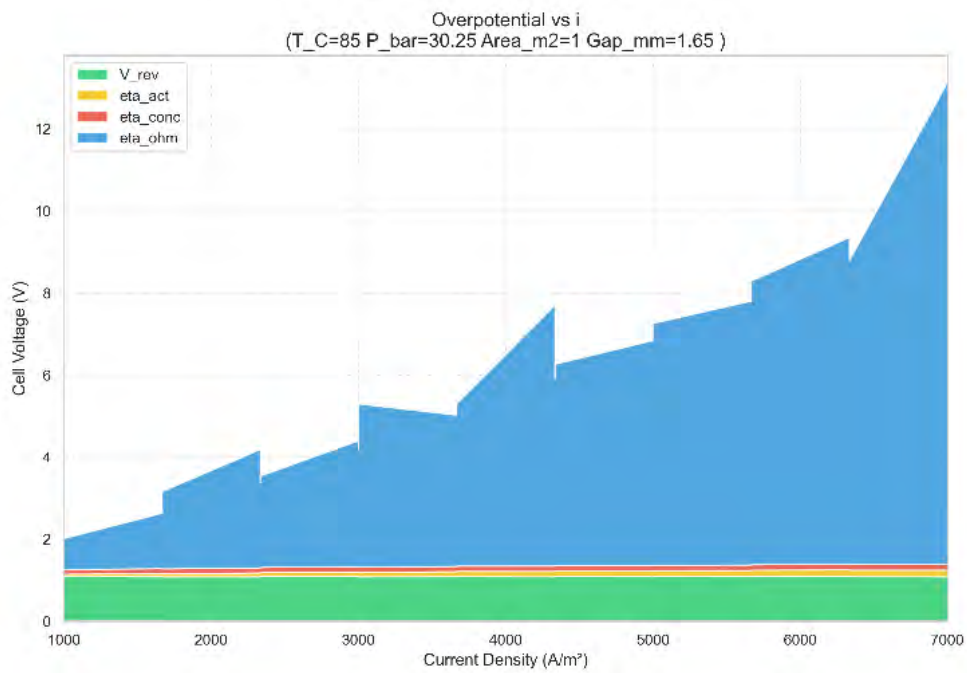

Figure 2. Contribution of overpotentials to the total cell voltage for Ni-Mo and Ni-Co electrodes.

$$V_{cell} = V_{oc} + \eta_{act} + \eta_{ohm} + \eta_{conc} \quad (1)$$

The overpotential includes in a explicit way the effect of physic, geometric and transport variables such as the ohmic overpotential (η_{ohm}) in Eq. (2) or the activation overpotential (η_{act}) in Eq. (3) shows:

$$\eta_{ohm} = I_{cell}\left(\frac{\delta_{anode}\rho_{ref}^{anode}}{A_{anode}(1-\varepsilon_{anode})^{1.5}} + \frac{\delta_{cathode}\rho_{ref}^{cathode}}{A_{cathode}(1-\varepsilon_{cathode})^{1.5}} + \frac{1}{\kappa_{electrolyte}A_{anode}}\left[(l_{anode-membrane} - \beta_{\text{anode}}) + \left(\frac{\beta_{anode}}{(1-\theta_{anode})^{m_{brug}}}\right)\right] + \frac{1}{\kappa_{electrolyte}A_{cathode}}\left[(l_{cathode-membrane} - \beta_{\text{cathode}}) + \left(\frac{\beta_{cathode}}{(1-\theta_{cathode})^{m_{brug}}}\right)\right] + \frac{\delta_{membrane}\tau_{membrane}^2}{\kappa_{electrolyte}\omega_{membrane}\varepsilon_{membrane}A_{membrane}}\right) \quad (2)$$

$$\eta_{act} = \frac{RT_{cell}}{a_{cathode}F_c} ln\left[\left(\frac{i}{2i_{exch}^{cathode}(1-\theta_{cathode})^{m_{brug}}}\right) + \sqrt{\left(\frac{i}{2i_{exch}^{cathode}(1-\theta_{cathode})^{m_{brug}}}\right)^2 + 1}\right] + \frac{RT_{cell}}{a_{anode}F_c} ln\left[\left(\frac{i}{2i_{exch}^{anode}(1-\theta_{anode})^{m_{brug}}}\right) + \sqrt{\left(\frac{i}{2i_{exch}^{anode}(1-\theta_{anode})^{m_{brug}}}\right)^2 + 1}\right] \quad (3)$$

The (η_{ohm}) accounts for the contribution of electrode gap (l), the bubble coverage (θ) effect on electrolyte conductivity (third and fourth term of sum) and membrane electric resistance (fifth term of sum) to name some variables. This rigorous treatment allows the model to predict the sharp voltage increase at high current densities due to bubble accumulation, a critical factor for sizing the circulation pumps and optimizing the operating pressure and current density. The production of Hydrogen was estimated with Eq. (4):

$$\dot{N}_{H_2}^{cathode,gen} = \frac{I_{cell}}{2F_c} \quad (4)$$

2.2 Desalinisation system

The desalination subsystem was modeled to ensure the water requirements of electrolysis were met, allowing for future evaluation of their energy integration under water efficiency criteria. This part of the process was developed with a Solar Still technology proposed by Abdel-Rehim & Lasheen (2005) [8].

This model allows estimating the necessary surface area of the solar still to supply water production, taking into account heat transfer under specific climatic conditions (solar resource, wind speed and ambient temperature), as well as energy and mass balances.

2.3 Separation units

The separation train (flash units and heat exchangers) was modeled using phase equilibrium equations and energy balances, allowing the electrolyte to be recovered and recirculated, which improves the overall efficiency of the process.

2.4 Sensitivity Analysis

A sensitivity analysis was performed using Monte

Table 1. Variables and ranges explored in the sensitivity analysis.

Variable	Lower Limit	Upper Limit	Units
Split	0	100	%
Distilled water recovery fraction	35	55	%
Temperature step size for flash stages	5.0	30.0	°C
AEL recirculation ratio	20.0	60.0	-
Electrolyte fraction in AEL cathode	40	60	%
AEL operating temperature	35	95	°C
AEL operating pressure	1	40	Bar
AEL current density	1000	6000	A/m²
Electrolyte concentration	20	30	wt%
Cathode-membrane gap	0.0002	0.006	m
Anode-membrane gap	0.0002	0.006	m
AEL cathode thickness	0.0005	0.003	m
AEL anode thickness	0.0005	0.003	m
SOEC water consumption fraction	60	85	%
SOEC operating temperature	700	900	°C
SOEC condenser temperature	25	60	°C
SOEC flash operating pressure	1	5	Bar
SOEC current density	2000	10000	A/m²
SOEC electrolyte thickness	5e-6	50e-6	m
SOEC cathode thickness	250e-6	1000e-6	m
SOEC anode thickness	20e-6	100e-6	m

Table 2. List of material in the sensitivity analysis.

Component	Material 1	Material 2	Material 3
AEL cathode material	Ni-Mo alloy	Classic Ni-Al alloy (Ni Raney)	SS 316L
AEL anode material	Ni-Fe alloy	Ni-Co spinel	SS 316L

Carlo simulations complemented by machine learning techniques (Random Forest) to identify the variables with the greatest impact on LCOH and energy efficiency. The variables explored and the evaluated ranges are detailed in Table 1.

Material Screening

Furthermore, the model's versatility was utilized to screen different electrocatalyst combinations by adjusting the kinetic parameters (e.g., exchange current density). Three materials were compared on each electrode, as shown in Table 2.

3. RESULTS AND DISCUSSION

The integrated model was simulated across a wide range of operating conditions, as described in the methodology section, enabling various types of analysis.

3.1. Voltage Breakdown Analysis

Understanding the distribution of energy losses is vital for stack design. Figure 2 presents the breakdown of the cell voltage at representative conditions (85°C, 30 bar).

The dominance of ohmic losses at high current densities highlights the necessity of zero-gap configurations.

It is observed that at low current densities, activation losses dominate due to the catalytic barrier. However, as production increases, the ohmic component becomes the governing factor, driven by bubble accumulation and electrolyte resistance. This validates the importance of operating at high pressures to compress gas bubbles and reduce the void fraction effect.

3.2. Sensitivity and Operating Maps

Sensitivity analysis allowed for the robust identification of the design and operating variables with the greatest influence on the overall system performance. Specifically, it was observed that operating temperature and current density are the dominant factors in energy consumption due to their direct impact on overpotentials and, consequently, on cell voltage, as shown in Figures 3, 4 and 5.

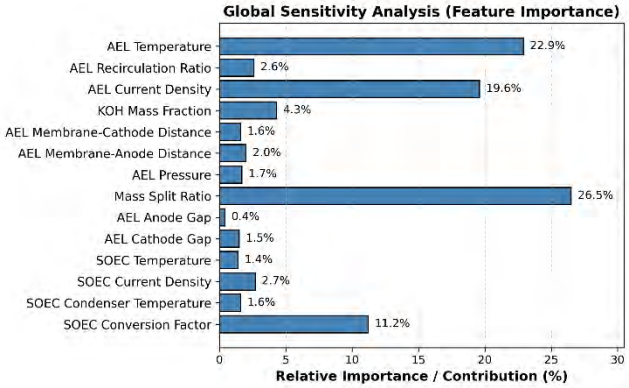

Figure 3. Relative importance of each variable on LCOH.

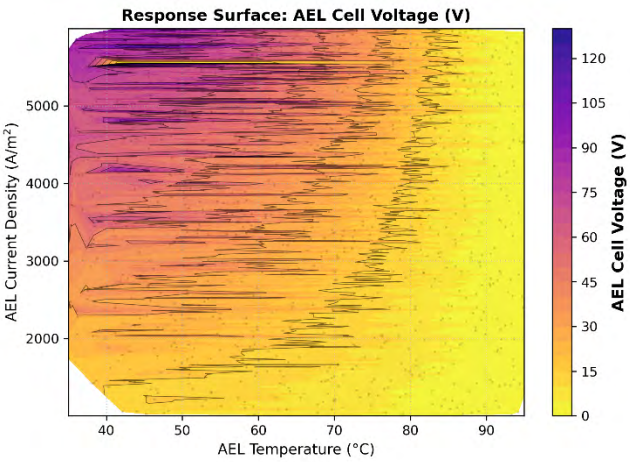

Figure 4. Cell voltage response surface.

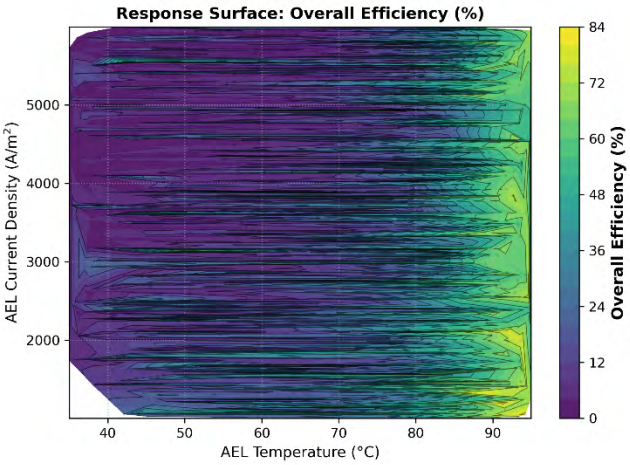

Figure 5. Impact of T of AEL and I of AEI on Overall Efficiency.

High-temperature operation minimizes voltage but requires careful thermal management to prevent electrolyte boiling. The results indicate that increasing temperature from 60°C to 90°C results in a voltage reduction of approximately 150-200 mV at high currents, which translates to a significant decrease in energy consumption.

There are zones where the temperature can be lower and maintain good efficiencies at high currents as shown in Figure 5, but in general, the best results are found at high temperatures.

3.3. Geometric Optimization

The impact of the geometric variables showed high sensitivity and highly dispersed values even among the most promising simulations, highlighting the need to measure the impact of these variables when there is a trade-off between configurations with lower ohmic resistances and those with lower manufacturing and material costs, when applying a multi-objective optimization approach.

The inter-electrode gap is a critical design parameter. Figure 6 illustrates the sensitivity of the voltage to gap variations.

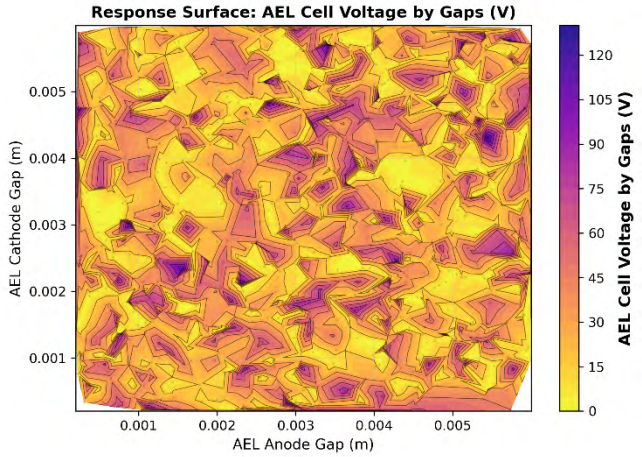

Figure 6. Impact of electrode gap distance on cell voltage.

Reducing the gap is the most effective strategy to minimize ohmic resistance in theory but the analysis suggests that there are multiple gap combinations that minimize the cell voltage and is lower when cathode have a greater gap than the anode.

3.4. Integrated Process Performance

The three best plants obtained by plotting the Pareto front shown in Figure 7 correspond to designs with a low LCOH (<5 USD/kg) and high overall efficiency (>65%). These points share an operating temperature of the AEL around to 90 °C, although they are quite heterogeneous in the other variables, including the Split system, which has values of 90%, 66%, and 95% at these points.

These results demonstrate that the integration of the desalination and electrolysis system is technically feasible and allow for the delimitation of feasible operating regions as a starting point for more aggressive optimization algorithms that consider environmental indicators.

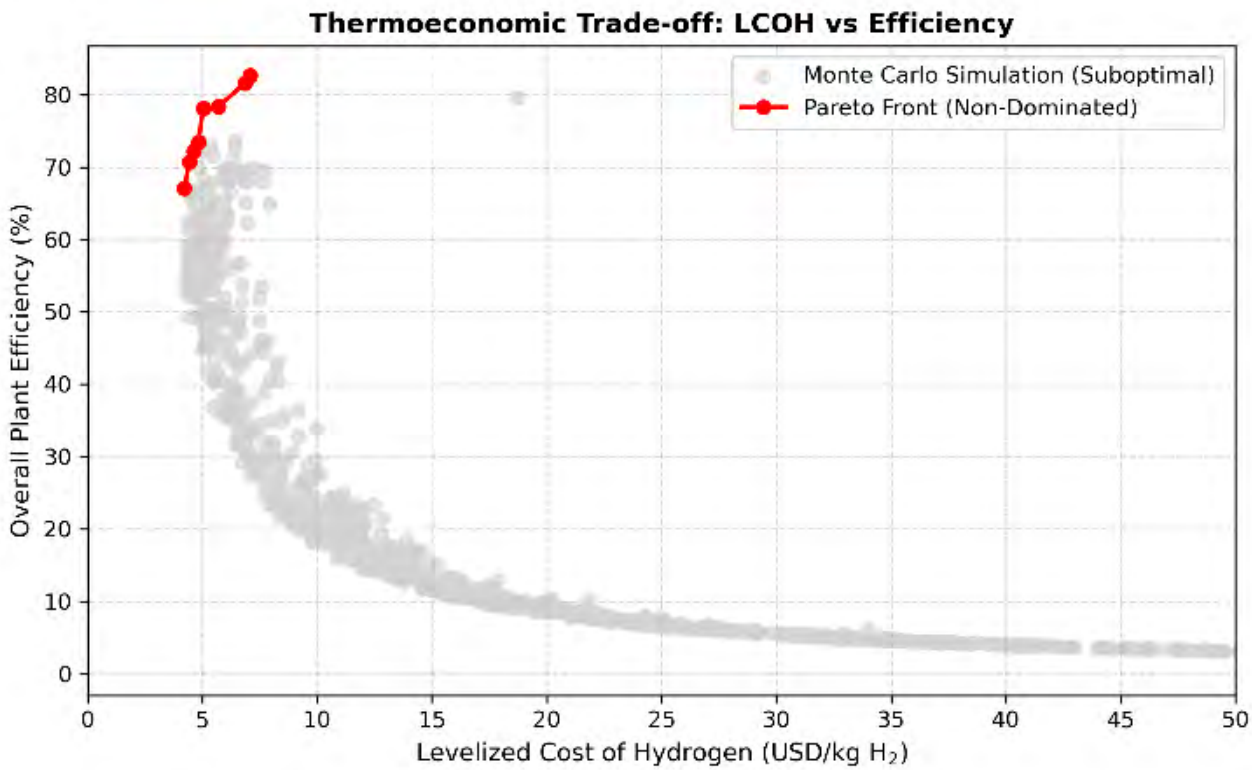

Figure 7. LCOH vs Global Energy Efficiency.

From an economic perspective, the results indicate that it is possible to achieve substantial reductions in LCOH through the appropriate selection of operating variables and materials. Operating regions are identified where the trade-off between energy efficiency and investment costs is optimal, validating the importance of sensitivity analysis as a tool for reducing the search space in complex optimization problems.

4. CONCLUSIONS

Taken together, the results not only demonstrate the viability of the integrated system but also provide key information on the most promising technological pathways. Furthermore, they highlight the importance of simultaneously considering energy, water, and economic aspects in the design of sustainable processes, aligning with the principles of process intensification and the circular economy.

The integration of desalination technologies with advanced electrolysis systems represents a robust strategy for ensuring the sustainable production of green hydrogen, reducing dependence on freshwater and strengthening the transition to circular models in resource management. This approach allows for the decoupling of hydrogen production from the use of conventional water resources, strengthening the resilience of energy systems and promoting circular water use schemes.

The development and application of rigorous mathematical models accurately capture the complexity of the phenomena involved, providing a solid foundation for decision-making in the early design stages. Through sensitivity analysis, it was possible to identify the variables most influential on system performance, facilitating the implementation of more efficient and targeted optimization strategies.

The results obtained confirm that the appropriate selection of operating conditions, materials, and geometric configurations can lead to significant reductions in the levelized cost of hydrogen and substantial improvements in energy efficiency. In this sense, the work lays the groundwork for the development of multi-objective optimization methodologies that integrate economic, environmental, and safety criteria.

Finally, this study directly contributes to the advancement of technological solutions aligned with the 2030 Agenda by proposing a green hydrogen production scheme with a smaller water footprint, greater efficiency, and high potential for industrial scalability, thus consolidating its role within the global energy transition.

ACKNOWLEDGEMENTS

Authors acknowledge the support provided by ADRESS Group, SECIHTI and Universidad de Guanajuato.

AUTHOR IDENTIFIERS

Author ORCIDs:
Antonio Torres-Ayala: 0009-0005-8865-8943
Eduardo Sánchez-Ramírez: 0000-0002-4326-4837
Marcelino Carrera-Rodríguez: 0000-0003-1988-2077
Juan Gabriel Segovia-Hernández: 0000-0003-4175-9809

REFERENCES

1. Werner and Yannick, "Modeling Electrolyzers and Gas Networks for Integration with Power Systems" *DTU Wind and Energy Systems, (*2023) https://doi.org/10.11581/DTU.00000301
2. Y. Zheng, "Power-to-Hydrogen Systems: Modelling, Operation and Techno-economic Analysis, " *DTU Wind and Energy Systems*, (2023) https://doi.org/10.11581/DTU.00000270
3. Hydrogen Europe. (2022). Hydrogen production & water consumption. https://hydrogeneurope.eu/wp-content/uploads/2022/02/Hydrogen-production-water-consumption_fin.pdf
4. Chenoweth, J., Lopez-Aviles, A., Morse, S., & Druckman, A. *Water And Wellbeing: Is There A Link?* World Water Congress (2015), Edinburgh Scotland 1. Global challenges for water governance. IWRA. https://iwra.org/proceedings/index.php?page=286&eventid=5&abstract_id=2865
5. Eurostat. Population on 1 January. (2024). https://ec.europa.eu/eurostat/
6. I. Ivanova, "Hydrogen's Potential in Desalination, " hydrogenera.eu. https://hydrogenera.eu/tpost/6ebz9l9ii1-hydrogens-potential-in-desalination-prov
7. Abdin Z, Webb CJ, Gray EM. Modelling and simulation of an alkaline electrolyser cell. Energy 138:316-331 (2017). https://doi.org/10.1016/j.energy.2017.07.053
8. Abdel-Rehim ZS, Lasheen A. Improving the performance of solar desalination systems. Renewable Energy 30:1955-1971 (2005). https://doi.org/10.1016/j.renene.2005.01.008

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Net Carbon Balance (NCB): a Better Way to Evaluate and Optimize Carbon Capture Technologies

André F. Young[a]*, Aline R. Eckstein[b], Leonardo D. Ferreira[b], Vítor M. Sermoud[b], and Ingrid A. de Oliveira[b]

[a] Federal Fluminense University, Department of Chemical and Petroleum Engineering, Niterói, Rio de Janeiro, Brazil
[b] Federal University of Rio de Janeiro, School of Chemistry, Rio de Janeiro, Rio de Janeiro, Brazil
* Corresponding Author: ayoung@id.uff.br.

ABSTRACT

The objective of this paper is to present a single equation format for quantifying the net carbon balance (NCB) in the evaluation of CO_2 capture technologies, and to discuss the benefits of this approach. The equation must take into account indirect emissions, especially the contributions from utility generation systems (heating, cooling and electricity), making use of efficiency values and emission factors. The idea is to synthesize, in a single expression, the quantification of the environmental footprint of a technology, in a practical way so that it could be used as an efficient metric in technical evaluation studies, or as objective function/constraint in optimization problems. It also facilitates demonstrating the relationship between capture efficiency and environmental performance, as well as the contribution of each term to total emissions, and to compare different technologies in terms of time, location and available energy sources. To illustrate the application of the proposed NCB formulation, two processes were compared in two different scenarios: post-combustion and pre-combustion capture with MEA absorption, in a Base Case formulated with literature premises and an NCB Zero Case meaning carbon neutrality. A process simulator – Aspen HYSYS – was used to solve the mass and energy balances associated with each process, providing the necessary information for NCB calculation, whose parameters were obtained from the technical literature and were discussed in relation to their importance to the proposed analysis.

Keywords: Carbon Dioxide, Process Design, Modelling and Simulations, Energy Efficiency, Environment, Emissions

INTRODUCTION

In recent years, studies assessing the technical and economic feasibility of carbon capture technologies have become increasingly common, meeting industrial and government demand for emissions mitigation and energy transition solutions. Among these studies, optimization works can be identified. However, most studies focus solely on maximizing carbon capture efficiency, occasionally taking into account energy consumption, but not the associated emissions. Therefore, the environmental impact reduction associated with these solutions becomes questionable.

For a more realistic technology assessment, it is necessary to consider indirect emissions related to energy generation in the carbon capture technologies (not to mention fugitive emissions), which can be very significant. Therefore, in this work, we propose that technologies should be evaluated and optimized in terms of their Net Carbon Balance (NCB), defined, for example, according to Equation 1:

$$NCB = f_{e,h} \cdot \sum \frac{q_h}{\eta_h} + f_{e,c} \cdot \sum q_c + f_{e,e} \cdot \sum \frac{q_e}{\eta_e} - w_{CO2,cap} \quad (1)$$

where: q_h is the thermal energy consumed for heating, in MJ per year; q_c is the thermal energy removed by cooling, in MJ per year; q_e is the electric energy balance between what is consumed by pumps, blowers and compressors, and what is eventually recovered with turbines, in MJ per year; η_h is the efficiency associated with heating utilities generation (e.g, steam generation in a boiler); η_e is the efficiency associated with electricity consumption (i.e., conversion between electricity and shaft work); $f_{e,h}$ is an

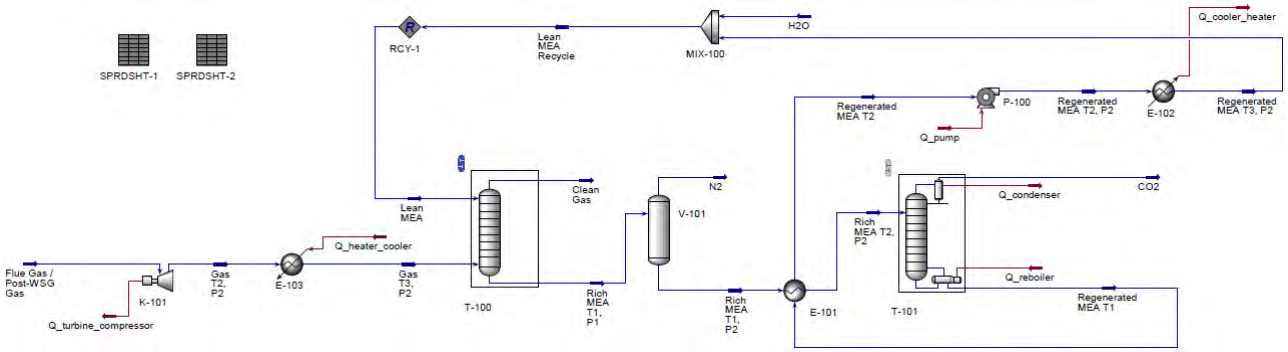

Figure 1. Flowsheet of the carbon capture process simulated in this work: MEA absorption.

emission factor relative to heating utilities utilization (e.g., fuel combustion for steam generation in a boiler), in Mtons of CO_{2e} per MJ; $f_{e,\,c}$ is an emission factor relative to cooling utilities pumping and regeneration (e.g, electricity consumption in the fan of a water refrigeration tower), in Mtons of CO_{2e} per MJ; $f_{e,\,e}$ is an emission factor relative to electricity utilization (e.g., a regional emission factor, depending on the energy matrix), in Mtons of CO_{2e} per MJ; and $w_{CO2,\,cap}$ is the amount of CO_2 captured, in Mtons per year.

The emission factors, as well as typical efficiency values, can easily be obtained from official technical reports and scientific literature. The amount of heat and electricity consumed in a carbon capture technology can be obtained through process simulation, along with a good estimation of its CO_2 capture potential. Defined this way, negative values for NCB indicate a real emissions mitigation scenario, while positive values indicate that emissions are not being effectively reduced, which compromises the environmental meaning of the proposed solution.

In this work, to illustrate the use of NCB in evaluating carbon capture technologies, two processes were compared: carbon capture from a typical post-combustion stream (flue gas); and carbon capture from a typical pre-combustion stream (hydrogen production process by biomass gasification, after water-gas shift), both using MEA absorption. These streams differ in terms of flow rate, temperature, pressure and composition, especially the CO_2 content. The processes were simulated in Aspen HYSYS, from where the mass and energy information required for NCB calculation was obtained.

METHODOLOGY

Emission Sources

To exemplify the use of the NCB formalism as a metric of evaluation of different carbon capture technologies, we selected two well-established processes: MEA absorption, applied to a post-combustion CO_2 emission source and to a pre-combustion CO_2 emission source.

As a good representative of a post-combustion source, the flue gas stream from the study by Moghaddam et al. [1] was chosen. The authors performed an optimization study of the MEA (monoethanolamine) absorption process based on process simulation and response surface methodology, and it was also the basis for the simulations performed in this work. This article was selected because it is a recent process simulation study, and the reported flue gas stream has a low CO_2 concentration, which is not only typical of post-combustion sources – in contrast to pre-combustion sources –, but also interesting for the analysis proposed in this work, about capture efficiency, energy demand, etc. Also, as a simplification, contaminants were not simulated, which is welcome because they would hinder the simulation and not contribute to the discussions that were intended to be made.

For the selection of a pre-combustion source, we performed a brief literature review on hydrogen production processes from lignocellulosic biomass, a topic of interest for our research group. In these processes, a residual biomass is usually gasified, generating syngas (mainly composed by H_2 and CO, but also containing water and contaminants, such as CH_4, CO_2, H_2S and particulates). The syngas stream then follows though some cleaning stages and undergo water-gas shift (WGS) for the production of more H_2 and CO_2. The CO_2 must be removed, and that is where the carbon capture technologies can be applied. The search was restricted to process simulation studies published in the last five years, and that presented a high level of detail (temperature, pressure, flow rate and composition) on the desired process stream. The thing is that the stream composition, especially the CO_2 content, can vary greatly depending on the feedstock processed, the gasification conditions and the subsequent steps, being usually higher, though, than that found in post-combustion sources.

Three articles called our attention: Meramo-Hurtado et al. [2] simulated the gasification of a mixture of rice and cassava waste. The produced syngas followed through the mentioned processing stages up to the CO_2

removal step, where a membrane technology was evaluated. Shamsi et al. [3] presented a steam reforming process aimed at removing the tars generated in the gasification of lignocellulosic biomass, increasing H_2 production and improving energy efficiency compared to the conventional process. The process was simulated with a pressure swing adsorption (PSA) unit, followed by MEA-PZ (monoethanolamine-piperazine) absorption. Finally, Cook and Hagen [4] simulated hydrogen production from different mixes of woody biomass, in a process fashion that also combined gasification and reforming. The removal of CO_2 was done with PSA. Comparing the process streams right after WGS in these three studies, that from Shamsi et al.'s study was the one that we considered the best representative of such streams and was the one selected for this study.

Table 1 presents the two emission sources used as examples in the present work, which can be considered typical post-combustion and pre-combustion sources. As can be seen, the post-combustion source (flue gas) is available at a higher temperature and lower pressure than the pre-combustion source (post-WGS syngas). Considering the conditioning steps required before absorption, these different conditions would typically imply different demands in terms of heating/cooling, and input/output of electric energy for pressurization/depressurization. All of that is accounted for in the NCB equation proposed in this work. In addition, the mole fraction of CO_2 is ten times higher in the pre-combustion source, which makes the absorption process easier, consequently demanding less energy and implying lower indirect emissions, as will be shown.

Table 1: CO_2 emission sources evaluated in this work.

Stream	Post-Comb. [1]	Pre-Comb. [3]
Temp. (ºC)	100.00	25.00
Pressure (bar)	1.00	12.00
Flow rate (kg/hr)	273654.06	27411.22
Mole fraction		
CO_2	0.0232	0.2740
H_2O	0.1019	0.0030
N_2	0.8749	0.3130
H_2	0.0000	0.4100

An interesting fact – which also helps to justify the choice for the selected references [1, 3] – is that the flow rate of the post-combustion stream is around 10 times higher than the flow rate of the pre-combustion stream, making the flow rates of captured CO_2 of the same order.

Process Simulations

The process simulations were performed in the Aspen HYSYS® software, utilizing the Acid Gas – Chemical Solvents thermodynamic package, which employs the Electrolyte NRTL (Non-Random Two-Liquid) activity coefficient model for the liquid phase and the Peng-Robinson equation of state for the vapor phase. In all the cases, the solvent used for CO_2 absorption was an aqueous solution of MEA at 30% by mass.

Process description

The process flowsheet is illustrated in Figure 1. In the literature, the temperature range in the feed stream of the absorption column is generally cited as being in the range of 40 °C to 60 °C. In this work, the temperature was set at 40 °C. To ensure the gas stream from the emission source met the temperature and pressure conditions of the absorber inlet, a turbine or compressor (K-101) and a heater or cooler (E-103) were positioned upstream of the absorption column. The solvent stream entering at the top of the column contains a small amount of CO_2, as it is a recycle stream. The solvent flow rate was adjusted so that the capture ratio, defined as the mass of MEA used per mass of CO_2 captured, met the specification (see the definition of the Base Case scenario in the next section).

The absorption column was specified to operate at 3 bar in the post-combustion scenarios and 2 bar in the pre-combustion scenarios, with plastic Mellapak™ 250X packing. The number of stages was determined as the minimum required for the desired robustness (avoiding non-convergence problems), resulting in 80 theoretical stages for the post-combustion scenarios and 30 theoretical stages for the pre-combustion scenarios. The columns were sized by the simulation software, but the size of the columns was not subject of analysis in this work (of course, for a rigorous techno-economic evaluation, alternative designs should be assessed; for example, a very high column could be partitioned in two columns).

After absorption, the CO_2-rich solvent exiting the bottom passes through a flash vessel (V-101), used to remove traces of nitrogen present and to adjust the pressure before the regeneration column (T-101). Next, the stream passes through an integration heat exchanger (E-101), which raises its temperature. The inlet temperature to the regeneration column was defined based on the established heat integration, aiming to recover most of the heat available in the bottom product stream.

In the regeneration column, reflux ratio and reboiler duty were adjusted so that the heat ratio, defined as the heat supplied per mole of CO_2 captured, met the specification (see the definition of the Base Case scenario in the next section). The column was specified to operate at 1 bar, with total reflux condenser. Once again, the number of stages was determined as the minimum required for the desired robustness, resulting in 75 theoretical stages for the post-combustion scenarios and 40 theoretical stages for the pre-combustion scenarios.

After regeneration, the solvent passes through the

Table 2: Stream properties, Base Case, post-combustion process.

	Gas T3, P2	Clean Gas	Lean MEA	Rich MEA T2, P2	Regenerated MEA T1	CO2
Temperature (ºC)	40.00	39.14	40.00	100.00	102.10	40.16
Pressure (bar)	3.00	3.00	3.00	1.00	1.00	1.00
Flow rate (kmol/hr)	9212.43	8947.61	16000.00	16264.52	16024.71	239.81
Mole fraction						
CO_2	0.0252	0.0010	0.0208	0.0342	0.0208	0.9243
MEA	0.0000	0.0000	0.1101	0.1083	0.1099	0.0000
H_2O	0.0252	0.0212	0.8691	0.8575	0.8692	0.0752
N_2	0.9497	0.9777	0.0000	0.0000	0.0000	0.0005
H_2	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

heat integration exchanger again, where it is cooled, then follows to a pump (P-100) for pressure adjustment and subsequently to a cooler (E-102) for temperature reduction. The final step to enable solvent recycling is the make-up of water and MEA (MIX-100), compensating for eventual losses. The solvent then returns to the top of the absorption column, closing the cycle.

Simulated scenarios

Two scenarios were simulated for both the post-combustion and pre-combustion processes: a Base Case scenario and an NCB Zero scenario, which consists in configurations for which NCB equals zero.

The Base Case was defined by establishing three well-accepted procedural premises for carbon capture by MEA absorption: capture efficiency of at least 90% in relation to the amount of CO_2 present in the absorption column inlet stream; capture ratio of up to 10 kgMEA/kgCO_2; heat ratio of approximately 200 kJ/molCO_2 in the regeneration column. In general, to meet these criteria, the following variables were manipulated: the solvent flow rate at the inlet of the absorption column, the reflux ratio and the reboiler duty of the regeneration column. Although a formal optimization study is not the objective of this work, optimizations performed informally by the authors in terms of NCB demonstrated that the values presented by most of the analyzed parameters were sufficiently close to their optimal values. The stream properties for the Base Case scenarios of the post-combustion and pre-combustion processes can be seen in Tables 2 and 4, respectively.

The NCB Zero scenario was achieved by means of manual modifications of the manipulated variables. The reason for bringing the NCB Zero Case into this analysis was to identify the system's equilibrium point in terms of carbon neutrality, in order to demonstrate the potential of the proposed carbon balance formulation to correlate process performance parameters with the actual carbon footprint of the evaluated technologies. The stream properties for the Base Case scenarios of the post-combustion and pre-combustion processes can be seen in Tables 5 and 6, respectively.

Carbon Balance Analysis

All the parameter values necessary for solving Equation 1 are presented in Table 3, with the appropriate references. It is worth mentioning that the value adopted for the emission factor related to heating utilities utilization ($f_{e,h}$) is compatible with the use of natural gas as fuel, while the emission factor for electric energy utilization ($f_{e,e}$) corresponds to the value published by the Brazilian Ministry of Science, Technology and Innovation (MCTI) for 2023.

Table 3: Parameters for NCB calculation.

Parameter	Value	Ref.
Steam boiler efficiency (η_h)	0.90	[5]
Pump efficiency ($\eta_{e,pump}$)	0.75	[6]
Turb./Comp. efficiency ($\eta_{e,t/c}$)	0.85	[7]
Fan efficiency ($\eta_{e,ref}$)	0.90	[8]
Heating emission factor ($f_{e,h}$)[a]	5.59×10^{-5}	[9]
Electricity emission factor ($f_{e,e}$)[a]	1.07×10^{-5}	[10]
Water specific heat (Cp_{H2O})[b]	4.18	[11]
Water density (ρ_{H2O})[c]	1000	[11]
Water temp. variation (ΔT)[d]	10	[12]

[a] Emission factors given in kgCO_2/kJ.
[b] kJ/kg.ºC.
[c] kg/m³.
[d] ºC.

The total thermal energy consumed in the heating stages was calculated as the sum of the duties of the regeneration column's (T-101) reboiler and E-103, when it is a heater. The total thermal energy removed in the cooling stages was calculated as the sum of duties of the regeneration column's condenser, the cooler (E-102) and E-103, when it is a cooler. All these duties can be obtained directly from the process simulator.

The total electric energy consumed or generated in the CO_2 absorption process was obtained from the sum

Table 4: Stream properties, Base Case, pre-combustion process.

	Gas T3, P2	Clean Gas	Lean MEA	Rich MEA T2, P2	Regenerated MEA T1	CO_2
Temperature (ºC)	40.00	40.30	40.00	80.00	101.70	32.80
Pressure (bar)	2.00	2.00	2.00	1.00	1.00	1.00
Flow rate (kmol/hr)	1263.00	959.40	13002.30	13305.70	12955.80	349.90
Mole fraction						
CO_2	0.2739	0.0148	0.0250	0.0494	0.0251	0.9493
MEA	0.0000	0.0000	0.1115	0.1090	0.1119	0.0000
H_2O	0.0030	0.0334	0.8634	0.8416	0.8630	0.0504
N_2	0.3130	0.4120	0.0000	0.0000	0.0000	0.0001
H_2	0.4101	0.5397	0.0000	0.0000	0.0000	0.0002

of the energy contributions of the pump (P-100) and the turbine or compressor (K-101), whose power ratings can be obtained directly from the process simulator. It is important to note that the turbine reduces net electricity consumption, as it generates energy. Therefore, its contribution must be considered with a negative sign in the summation, representing an energy gain rather than a cost. The possibility of harnessing energy from pressurized streams is a characteristic of pre-combustion streams, which can make a difference in certain cases.

The electric energy consumption related to the cooling water plant can be calculated according to Equation 2 [8]:

$$E_{ref} = \frac{F_{tower} \cdot 2.432 \times 10^{-4}}{\eta_{ref}} \quad (2)$$

where: E_{ref} is the heat exchanged in the water refrigeration tower, in GJ/hr; F_{tower} is the flow rate of water in the refrigeration tower, in m^3/hr; and η_{ref} is the efficiency of the refrigeration tower fan.

The water flow rate in the refrigeration tower, in turn, can be calculated according to Equation 3:

$$F_{tower} = \frac{\sum q_c}{Cp_{H2O} \cdot \Delta T \cdot \rho_{H2O}} \quad (3)$$

where: Cp_{H2O} is the specific heat of water at constant pressure, in kJ/kg.ºC; ΔT is the temperature variation of the cooling water in the tower, in °C; and ρ_{H2O} is the density of water, in kg/m³.

The cold utility contribution to NCB can then be calculated by multiplying E_{ref} by the electricity emission factor ($f_{e,e}$), subsequently correcting it to kJ/hr as shown in Equation 4:

$$f_{e,c} \cdot \sum q_c = f_{e,e} \cdot E_{ref} \times 10^6 \quad (4)$$

As formulated, NCB should ideally be negative, and the lower its value, the better the environmental performance of the process, in absolute terms of CO_2 emission abatement. However, in order to compare different processes, especially CO_2 capture technologies applied to emission streams with very different flow rates and compositions, or solutions obtained for the same capture technology that exhibit different capture efficiencies, it is convenient to define a Net Carbon Ratio (NCR):

$$NCR = \frac{NCB}{w_{CO2,cap}} \quad (5)$$

In this work, it allows a fairer comparison of the results for the Base Case scenarios.

RESULTS AND DISCUSSION

Base Case

The stream properties for the Base Case scenarios of the post-combustion and pre-combustion processes are presented in Tables 2 and 4, respectively, and the results for the main process parameters can be seen in Table 7. Comparatively, the post-combustion scenario required a higher MEA-to-CO_2 capture ratio to achieve an equivalent performance, which is expected due to the much lower CO_2 content in the emission source (leading to thermodynamic and mass transfer limitations). Besides that, the need for more solvent has direct impact on the heat ratio in the reboiler of the regeneration column, since more solvent must be evaporated.

This result can be translated in terms of NCB (and NCR): both processes can be considered effective because they contribute to a net abatement of CO_2, as indicated by the negative sign of NCB. However, the pre-combustion capture process is more effective than the post-combustion process, since it allows twice the relative abatement (NCR of -0.62) of the post-combustion process (NCR of -0.29). In other words, 71% of the captured CO_2 were offset by indirect emissions in the post-combustion process, while only 38% were offset in the pre-combustion process. It is important to emphasize that these processes were not optimized and are being presented in this paper only to illustrate the methodology.

Table 5: Stream properties, NCB Zero Case, post-combustion process.

	Gas T3, P2	Clean Gas	Lean MEA	Rich MEA T2, P2	Regenerated MEA T1	CO2
Temperature (ºC)	40.00	39.15	40.00	100.00	103.70	68.73
Pressure (bar)	3.00	3.00	3.00	1.00	1.00	1.00
Flow rate (kmol/hr)	9212.43	8929.63	16250.00	16532.51	16206.37	326.14
Mole fraction						
CO_2	0.0252	0.0004	0.0343	0.0475	0.0343	0.7027
MEA	0.0000	0.0000	0.1645	0.1617	0.1649	0.0000
H_2O	0.0252	0.0199	0.8012	0.7908	0.8007	0.2970
N_2	0.9497	0.9797	0.0000	0.0000	0.0000	0.0004
H_2	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 6: Stream properties, NCB Zero Case, pre-combustion process.

	Gas T3, P2	Clean Gas	Lean MEA	Rich MEA T2, P2	Regenerated MEA T1	CO2
Temperature (ºC)	40.00	40.30	40.00	80.00	100.30	94.80
Pressure (bar)	2.00	2.00	2.00	1.00	1.00	1.00
Flow rate (kmol/hr)	1263.00	952.54	10414.13	10724.48	8629.03	2095.45
Mole fraction						
CO_2	0.2739	0.0085	0.0180	0.0490	0.0217	0.1612
MEA	0.0000	0.0000	0.1178	0.1144	0.1422	0.0000
H_2O	0.0030	0.0328	0.8642	0.8366	0.8361	0.8387
N_2	0.3130	0.4150	0.0000	0.0000	0.0000	0.0000
H_2	0.4101	0.5437	0.0000	0.0000	0.0000	0.0000

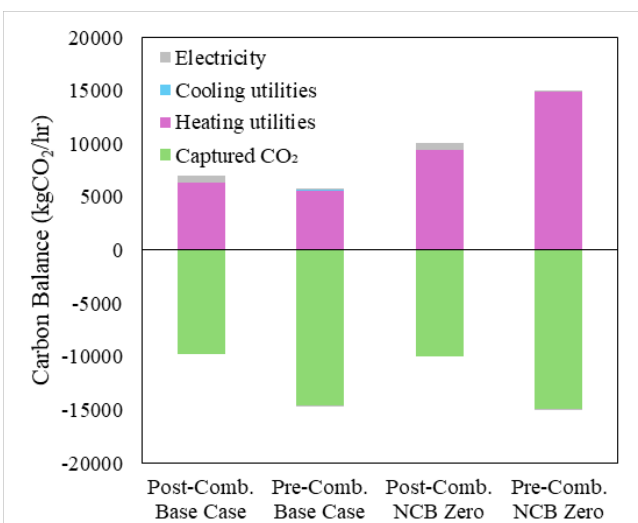

Figure 2. NCB breakdown.

It is also interesting to look at the relative contributions of each energy type to the net carbon balance, which can be done with Figure 2. The reboiler duty corresponds to 99% of the energy demand for heating in the post-combustion process, and 95% in the pre-combustion process. As a consequence of the use of natural gas for vapor generation in the boiler, it accounts for around 90% and 95% of the total indirect CO_2 emissions, respectively. It means that the environmental effectiveness of these carbon capture processes could be greatly enhanced with the use of lower emission energy sources – eventually renewable energy – for heating, which would reduce the value of $f_{e,h}$. This information is not new, but the advantage of using the proposed NCB formulation in this analysis is that it can be directly inferred from the equation, which synthesizes, in the emission factors, the results of sustainability analyses performed elsewhere.

Electricity consumption accounts for less than 10% of the indirect emissions in the post-combustion MEA absorption process, because the inlet stream must only be pressurized from 1 to 3 bar. In the pre-combustion process, its contribution is even negative, because inlet stream depressurization from 12 to 2 bar can be used to generate electricity (there is an implicit assumption here, that this electricity would be useful elsewhere in the plant, since it is more than what is needed in the carbon capture process itself). In this regard, two comments must be done: first, the relative importance of this term can be higher for technologies that require significant pressure increases, such as membrane separation, for example; second, it strongly depends on the environmental footprint of electricity generation, translated in $f_{e,e}$ (the Brazilian electric energy matrix, used as a basis for

Table 7: Result comparison.

Indicators	Post-Comb. Base Case	Pre-Comb. Base Case	Post-Comb. NCB Zero	Pre-Comb. NCB Zero
CO_2 recovery (%)	95.58	96.02	98.82	97.66
Capture ratio (kgMEA/kgCO_2)	10.99	6.07	16.23	5.04
Heat ratio (kJ/molCO_2)	460.18	259.60	663.10	696.79
Process Parameters				
CO_2 flow rate in inlet gas (kmol/hr)	231.91	345.94	231.91	345.94
CO_2 flow rate in stripping outlet (kmol/hr)	221.66	332.16	229.17	337.83
Energy consumed in the reboiler (T-101) (kJ/hr)	1.02E+08	8.62E+07	1.52E+08	2.35E+08
Energy consumed in the heater (E-102) (kJ/hr)	6.11E+05	-	-	-
Energy consumed in the heater (E-103) (kJ/hr)	-	4.48E+06	-	4.48E+06
Total thermal energy consumed for heating (kJ/hr)	1.03E+08	9.07E+07	1.52E+08	2.40E+08
Energy removed in the condenser (T-101) (kJ/hr)	9.35E+07	3.26E+07	1.21E+08	1.12E+08
Energy removed in the cooler (E-102) (kJ/hr)	-	5.17E+07	1.78E+07	4.18E+07
Energy removed in the cooler (E-103) (kJ/hr)	1.05E+08	-	1.05E+08	-
Total thermal energy removed for cooling (kJ/hr)	1.98E+08	8.43E+07	2.43E+08	1.54E+08
Energy consumed in the pump (P-100) (kJ/hr)	9.55E+04	3.97E+04	1.06E+05	2.82E+04
Energy consumed in the compressor (K-101) (kJ/hr)	5.25E+07	-	5.25E+07	-
Energy removed in the turbine (K-101) (kJ/hr)	-	3.78E+06	-	3.78E+06
Total electric energy consumed/removed (kJ/hr)	5.26E+07	3.74E+06	5.26E+07	3.75E+06
Make-up water consumption (kmol/hr)	22.20	48.00	38.86	1786.00
Make-up MEA consumption (kmol/hr)	0.12	-	0.18	-
Carbon Balance				
NCB (kgCO_2/hr)	-2801.51	-9015.79	~0	~0
NCR	-0.2872	-0.6167	-	-

this analysis, is much cleaner than the ones in US or China, for example). It shows that the proposed NCB formulation is an efficient way to compare different locations for carbon capture technologies too. The moment in time can also be evaluated, since national and international agencies are constantly updating these emission factors.

NCB Zero

The stream properties for the NCB Zero scenarios of the post-combustion and pre-combustion processes are presented in Table 5 and 6, respectively, and the results for the main process parameters can be seen in Table 7. Since the Base Case scenarios for both processes were effective in mitigating net CO_2 emissions (negative values of NCB), the equilibrium point would occur under more demanding conditions, for example for obtaining higher CO_2 recoveries, as can be seen in the tables. It would be achieved, in MEA absorption systems, with higher heat ratios in the reboiler of the regeneration column.

As a consequence of the increased energy consumption for heating, more fuel must be burned, resulting in higher indirect emissions. In this context, 100% of the CO_2 recovered in the capture process is emitted in the same plant, nullifying the environmental benefits of carbon capture. This is an easy way to demonstrate that higher CO_2 recovery rates do not necessarily result in lower net emissions, and that indirect emissions must always be taken into account in the evaluation of carbon capture technologies.

The single equation format proposed for NCB, therefore, presents great potential not only for use as an objective function, but also as a constraint in optimization problems, particularly when other environmental or economic aspects are being considered. It is important to point out that this format is not unchangeable: other terms, such as fugitive emissions, could be added, and the emission factors could be recalculated in a dedicated work through the formulation of inventories (life cycle assessment). It is convenient, however, to have a simple

and straightforward way to quantify the net carbon balance. This is the contribution of this paper to the technical and scientific literature.

CONCLUSION

In this paper, we propose a single equation format for NCB calculation, accounting for direct and indirect CO_2 emissions associated with the evaluated technologies. It can be easily adapted, however, to account for fugitive emissions, as well as emissions throughout the entire life cycle of the analyzed processes or products. For example, it is worth noting that the proposed equation only considers operating carbon emissions, but does not consider emissions related to building the capture plant itself. It could be accounted for, if the user wishes.

With this formulation, it was possible to easily demonstrate that higher capture efficiencies do not necessarily correspond to maximum emission mitigation scenarios. Indeed, lower emissions are associated with lower energy consumption, which will depend on the type of stream (composition, condition) being treated, and, in this context, the CO_2 content in the emission source is a relevant parameter. However, better than simply accounting for energy consumption is directly assessing the associated emissions, since emission factors are readily available and can be used to evaluate different technologies in terms of time, location and available energy sources.

In next works, we plan to show how the NCB formalism can be applied in optimization studies, serving as a successful, quantitative, non-subjective, objective function and/or constraint.

ACKNOWLEDGEMENTS

The authors are grateful for the financial support received from the National Council for Scientific and Technological Development (CNPq), Brazil (Grant Nº. 301625/2025-5).

AUTHOR IDENTIFIERS

Author ORCIDs:
Young AF: 0000-0002-3022-6006
de Oliveira IA: 0000-0002-2527-656X
Sermoud VM: 0000-0001-8643-6507

REFERENCES

1. Moghaddam AH, Esfandyari M, Sakhaeinia H. Optimization of amine-based carbon capture: simulation and energy efficiency analysis of absorption section. Results in Engineering 24:103574 (2024). https://doi.org/10.1016/j.rineng.2024.103574
2. Meramo-Hurtado SI, Puello P, Cabarcas A. Process analysis of hydrogen production via biomass gasification under computer-aided safety and environmental assessments. ACS Omega 5:19667-19681 (2020). https://doi.org/10.1021/acsomega.0c02344
3. Shamsi M, Obaid AA, Farokhi S, Bayat A. A novel process simulation model for hydrogen production via reforming of biomass gasification tar. International Journal of Hydrogen Energy 47:772-781 (2022). https://doi.org/10.1016/j.ijhydene.2021.10.055
4. Cook B, Hagen C. Techno-economic analysis of biomass gasification for hydrogen production in three us-based case studies. International Journal of Hydrogen Energy 49:202-218 (2024). https://doi.org/10.1016/j.ijhydene.2023.07.219
5. Turton R, Shaeiwitz J, Bhattacharyya D, Whiting W. Analysis, synthesis, and design of chemical processes. Pearson Education (2018)
6. Seider WD, Lewin DR, Seader JD, Widagdo S, Gani R, Ng KM. Product and process design principles: synthesis, analysis, and evaluation. John Wiley & Sons (2017)
7. Ipieca. https://www.ipieca.org/resources/energy-efficiencysolutions/power-recovery-turbines-2013
8. Caxiano IN, Junqueira PG, Mangili PV, Prata DM. Eco-efficiency analysis and intensification of the acetic acid purification process. Chemical Engineering and Processing - Process Intensification 147:107784 (2020). https://doi.org/10.1016/j.cep.2019.107784
9. Campbell L, Toolen J, Grubert D, Napp G. Compendium of greenhouse gas emissions methodologies for the natural gas and oil industry. American Petroleum Institute (2021)
10. MCTI – Brazilian Ministry of Science, Technology and Innovation. https://www.gov.br/mcti/pt-br/acompanhe-o-mcti/noticias/2024/02/fator-de-emissao-de-co2-na-geracao-de-energia-eletrica-no-brasil-em-2023-e-o-menorem-12-anos
11. Haynes WM. CRC Handbook of Chemistry and Physics. CRC Press (2016)
12. Volkers B-D. Carbon dioxide for calcite scale control in cooling water systems. University of Groningen, Netherlands (2016)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Multi-Scale Design for Clean Energy Systems: Industrial Electrification and Flexible Operation of Ammonia Synthesis

Nicholas N. Kalamaris[a], Christos T. Maravelias[ab]*

[a] Princeton University, Department of Chemical Engineering, Princeton, New Jersey, United States of America
[b] Andlinger Center for Energy and the Environment, Princeton, New Jersey, United States of America
* Corresponding Author: maravelias@princeton.edu.

ABSTRACT

Flexible, electrified systems for chemical and energy production are promising alternatives to traditional, hydrocarbon-based processes. Flexible systems have the potential to reduce costs and emissions, but the interconnection between design and operation makes these systems challenging to implement. We use an operation-informed design framework to model a flexible, electrified ammonia synthesis system. We examine the levelized cost and carbon intensity of ammonia in response to different grid emissions (0-420 kg/MWh). We find levelized costs from 700-1200 \$/ton-NH_3 and observe non-monotonicity in carbon-intensity with respect to grid emissions. We rationalize this trend as a design transition from large, grid-reliant systems to smaller, flexible designs that are grid independent. We then study how synergies in demand and unit-operation flexibility can lower both the price and carbon-intensity of ammonia production. We find that for seasonal, or yearly demand (rather than hourly), a fully flexible Haber Bosch process can achieve 20% lower costs and reduce its carbon-intensity by ~70-100%. We highlight major challenge in decarbonization, but also the importance flexibility plays in reducing carbon-intensity. Together, these analyses demonstrate that flexibility helps electrified industrial systems achieve financial and environmental goals.

Keywords: Multiscale Modelling, Energy Systems, Process Design

INTRODUCTION

Global manufacturing relies on chemical and energy systems to interconvert matter and energy. Fundamental production processes for energy, commodity chemicals, and construction material are the basis of the international supply chain. Fossil fuels serve as cheap precursors for these processes, but petroleum resources dwindle as material and energy demand increase. The future security of manufacturing requires a switch to sustainable production.

Flexible, electrified systems are attractive alternatives to traditional processes because they can selectively access intermittent renewable power, avoid operation during energy price spikes, and participate in demand response programs. Despite these benefits, flexible systems struggle to penetrate the chemical sector due to challenges in design and operation. Many unit operations have limited flexibility, requiring mass and energy storage devices to buffer their exposure to variability. The interplay between storage, flexible, and inflexible units creates a challenging problem; the optimal design depends on the holistic operation of the system. Newly developed frameworks for operation-informed design can be implemented to understand how chemical and energy systems respond to multi-scale variability [1]. These frameworks can also be used to identify transitions in design and operation for different conditions such as energy availability, grid renewability, and demand side flexibility.

To identify and address systemic challenges in decarbonization, we implement an operation-informed design framework to perform a case study on electrified ammonia synthesis. Previous studies have examined the levelized cost of ammonia (LCOA) for electrified systems [2, 3, 4]. However, these studies do not fully examine the

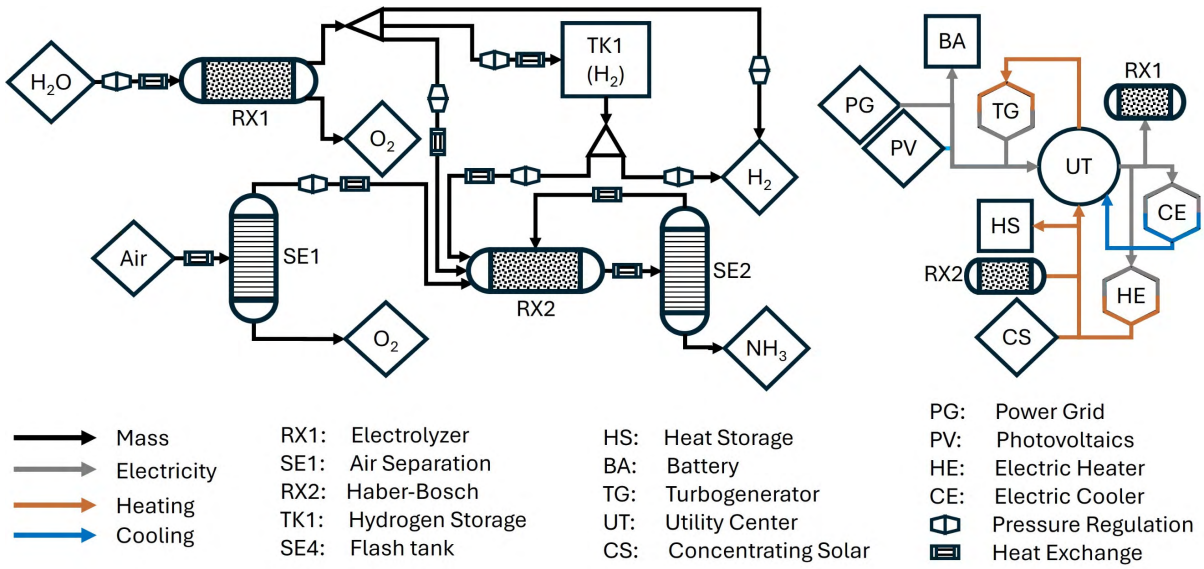

Figure 1. Superstructure for flexible, electrifed, ammonia sythesis. The system is connected to the power grid, and can also generate and interconvert power from solar irradiance onsite. Mass and energy storage allow for flexible operation in response to variability in energy price, availability, and associated emissions.

interplay between optimal design and operation. They are either simulation based (and require a pre-supposed design) or focus on operation under uncertainty, and sacrifice detail in terms of the system design and unit models. Recently, studies have suggested that process flexibility can positively impact the environmental and economic outcomes of ammonia synthesis [5, 6]. However, they utilize coarse grain system models and focus broadly on ammonia synthesis across different locations and modes of production. More detailed models exist in the context of process control, but do not optimize design [7].

The operation-informed design framework we use captures non-linearities in design (economies of scale) and operation (thermodynamics) in a single mathematical program. We use this to study the impact of grid emissions on the environmental and economic outcomes of this system. We then understand these outcomes by examining transitions in optimal design and operation. We also analyze the synergy between unit-operation and demand flexibility to identify conditions that can reduce the cost and carbon intensity of electrified ammonia production.

METHODS

Superstructure

We consider a system superstructure for flexible, electrified ammonia synthesis (Figure 1). The system can utilize grid electricity ($\leq$ 1 GW) or generate power onsite using photovoltaics and/or concentrating solar units ($\leq$ 10 km^2 total panel area). We utilize an emission penalty of 300 \$/ton-CO_2 for grid related emissions [8]. Heating, cooling, and electricity can be interconverted using power conversion units (heaters, coolers, turbogenerators). Hydrogen is generated via water electrolysis, and nitrogen is purified using cryogenic distillation of air. Hydrogen and nitrogen are converted to ammonia according to the Haber-Bosch reaction. The system can install hydrogen storage ($\leq$ 1500 m^3), as well as electricity and heat storage ($\leq$ 500 MWh each). All units are assumed to be flexible except for the Haber-Bosch reactor, which has limited flexibility.

Time Structure

The time-structure for the system is a 30-year time horizon with six epochs, each corresponding to five-years. An epoch is represented by a one-year operating period composed of four seasons. Each season contains representative days chosen from 3-years of power-grid data (Figure 2). The representative period accounts for hourly variability in electricity price, grid emissions, and solar irradiance using data from the CAISO grid [9, 10, 11].

Representative days for a season are generated using an IP clustering algorithm which groups individual days around a representative [1]. The algorithm minimizes the distance between all days and their representative at a fixed cost per representative. The time points in a representative day are weighted by the number of days they represent; representatives for larger clusters are weighted more heavily than smaller clusters (which typically represent atypical days). The representative days are hierarchy concatenated to form a season and subsequently a year, which is used as the

operating period for the epoch. We allow for discontinuity in process operation between, but not within, seasons. We do not consider retrofitting in this analysis, and the operation across epochs is repeated.

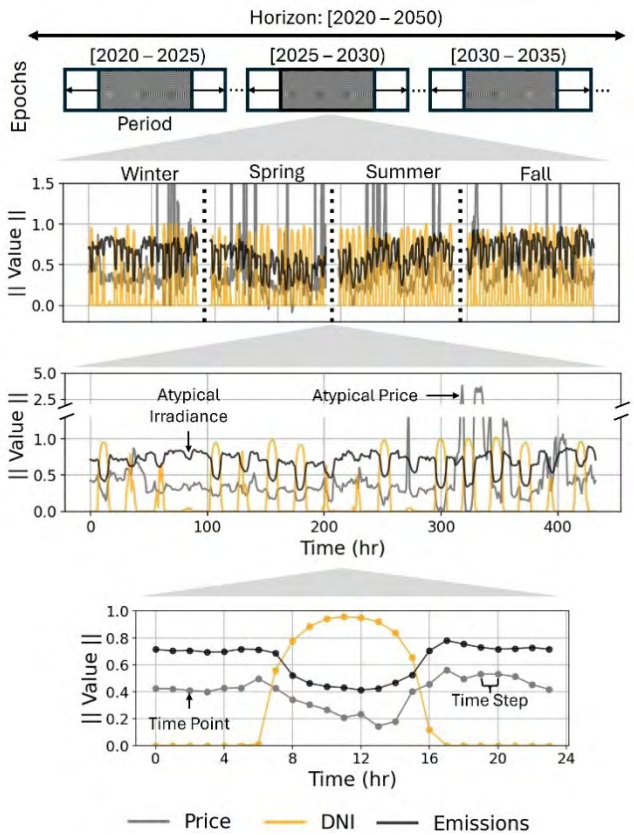

Figure 2. A 30 year horizon is partitioned into six epochs of five years. Epochs corresponds to a representative year with four seasons. Seasons are composed of representative days at an hourly resolution.

Mathematical Program

We utilize an operation-informed design framework to formulate a mathematical program for the system. This framework includes unit models for conversion reactors, separators, heaters, coolers, expanders, compressors, as well as mass and energy storage. Each unit has constraints enforcing mass (eq. 1) and energy (eq. 2) balances, as well as thermodynamic relationships (eq. 3) as described in [1].

$$\begin{aligned} \min \quad & C_{NH3} \\ \text{s.t.} \quad & \mathbf{m}(\mathbf{x}_t, \mathbf{x}_{t-1}) = 0 && \forall\, t \in \mathbf{T}\backslash\mathbf{T^0} && (1) \\ & \mathbf{e}(\mathbf{x}_t, \mathbf{x}_{t-1}) = 0 && \forall\, t \in \mathbf{T}\backslash\mathbf{T^0} && (2) \\ & \mathbf{h}(\mathbf{x}_t) = 0 && \forall\, t \in \mathbf{T} && (3) \\ & \mathbf{l}(\mathbf{y}) \le \mathbf{o}(\mathbf{x}_t) \le \mathbf{u}(\mathbf{y}) && \forall\, t \in \mathbf{T} && (4) \\ & \mathrm{NPV} = 0 && && (5) \end{aligned}$$

Here, $\mathbf{x}_t$ are the operating variables at time t and $\mathbf{y}$ are design variables. Limits on operation $\mathbf{o}(\mathbf{x}_t)$ are set relative to functions of the design: $\mathbf{l}(\mathbf{y}), \mathbf{u}(\mathbf{y})$ (eq. 4). $\mathbf{T}$ is the set of time points in the time structure, and $\mathbf{T^0}$ is the first time point in a season. The cost for unit operations is calculated based on their capacity, and the total capital expenditure is amortized over the process horizon. The system generates revenue by selling ammonia, and has operating expenses related to maintenance and labor, electricity purchase, and emissions penalties. There is a 21% corporate tax on profit. We find the system LCOA by minimizing the selling price of ammonia (C_{NH3}) subject to a net present value of zero (eq. 5).

First, we demonstrate how the optimal system LCOA, design, and operation change with respect to different levels of grid emissions. We consider representative profiles for the power grid with average emissions from 0-420 kg/MWh. The ammonia production capacity is a design variable between 5 and 100 tons/hr. We assume that there is no demand flexibility, such that the system maintains production at its nominal capacity. For this analysis we use a time structure containing 16 representative days for a period (Figure SI.1a).

Then, to understand how flexibility impacts the LCOA of the system, we fix ammonia production to 45 tons/hr and examine different combinations of demand and process flexibility. We consider systems with hourly demand (eq. 6), seasonal demand (eq. 7), and yearly demand (eq. 8). Average grid emissions are 220 kg/MWh. The sensitivity analysis utilizes 16 representative days for a period. We then analyze design and operation of flexible and inflexible systems using a 77-day time structure (Figure SI.1b).

$$M_{t,NH3} = \delta \;\forall\, t \in \mathbf{T} \tag{6}$$

$$\sum_{t \in \mathbf{T}_s^S} \omega_t M_{t,NH3} = \frac{1}{4}\delta \;\forall\, s \in \mathbf{S} \tag{7}$$

$$\sum_{t \in \mathbf{T}_p^P} \omega_t M_{t,NH3} = \delta \;\forall\, p \in \mathbf{P} \tag{8}$$

$\mathbf{P}$ the set of periods, $\mathbf{S}$ is the set of seasons, $\mathbf{T}_s^S$ is the set of time points for a season, and $\mathbf{T}_p^P$ is the set of time points for a period. $M_{t,NH3}$ is the production of ammonia at time t, δ is the ammonia demand, and ω_t is the weight for a time point (weights sum to one over $\mathbf{T}_p^P$). For each demand type, we quantify the LCOA for different levels (0-100%) of hourly flexibility from the nominal design of the Haber-Bosch reactor. We then examine how design and operational changes improve the LCOA.

The model is solved using CONOPT. Computations were performed on the Princeton Della computing cluster, using 2.8 GHz Intel Cascade Lake CPUs. Additional details can be found in the supplemental information.

RESULTS

Grid Renewability Analysis

We find a range of LCOA from 700-1200 $/ton over a

range of grid emissions from 0-420 kg-CO_2/MWh. We identify that the trend in LCOA increases linearly from 0-140 kg-CO_2/MWh before exhibiting a non-linear response to changes in grid emissions (figure 3a).

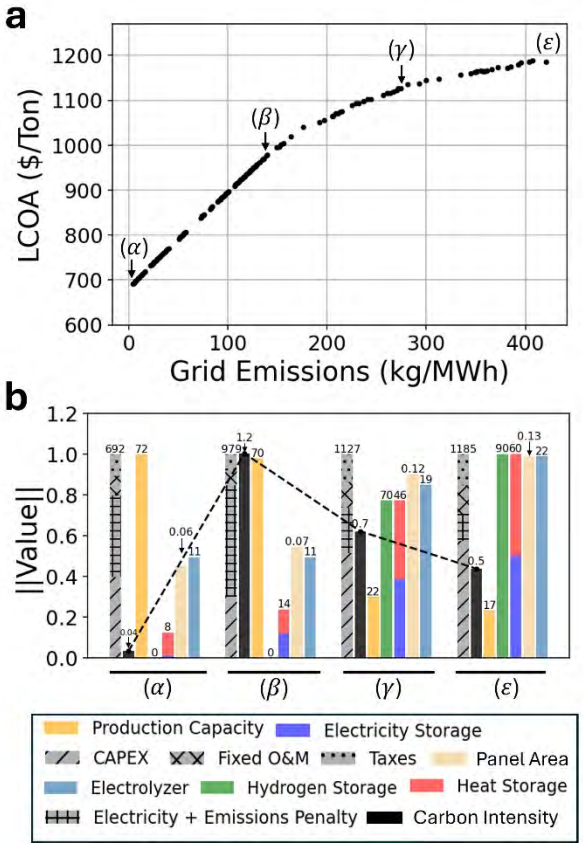

Figure 3. The levelized cost of ammonia for different average grid emissions (a). The contributions to the LCOA are shown in grey, along with production capacity in orange (ton-NH_3/hr), and carbon intensity in black (ton-CO_2/ton-NH_3). Relative to production capacity, mass storage (m^3), energy storage (MWh), solar facility panel area (km^2), and electyrolyzer capacity (MW) are also shown (b). Values have been normalized to lie between zero and one for the purpose of visualization; actual values appear above each bar.

This change is rationalized as a transition point in design and operation when examining instances $\alpha, \beta, \gamma, \varepsilon$. From 0-140 kg-$CO_2$/MWh, the production capacity remains constant, but the fraction of expenses due to emissions rises steadily. In this range, the process has a high production capacity but relies on the grid to provide power during the night, when onsite renewables are inactive. At ~140 kg-CO_2/MWh, operating costs dominate capital expenditure, and it becomes optimal to redesign the process: production capacity decreases, and the proportional level of mass and energy storage increase. This allows the process to operate more flexibly. Increased storage reduces power usage from the grid during levels of low solar irradiance. In turn, operating costs drop but capital costs increase disproportionately due to decreased economies of scale and overdesign (Figure 3b). Specifically, as emissions rise from 140 to 420 kg-CO_2/MWh, the relative electrolyzer capacity and solar facility size double. Overdesign allows these systems to operate independently from the grid but drastically increases capital costs.

The interplay between capital and operating expenditure causes non-monotonic behavior in the carbon-intensity of the system. Systems operating in low or high emission grids produce ammonia at a lower carbon intensity than grids with intermediate emissions. This presents a major challenge to ammonia decarbonization through electrification. Moderate improvements in grid renewability may not improve system level emissions. Carbon intensive grids must decarbonize significantly for green ammonia production to be competitive. Unfortunately, changes to the grid are often slow and incremental. However, these results suggest that increasing system flexibility is the key to achieving lower emissions and prices.

System Flexibility Analysis

Storage is not the only sources of flexibility; it can also be achieved by unit operations or the demand side.

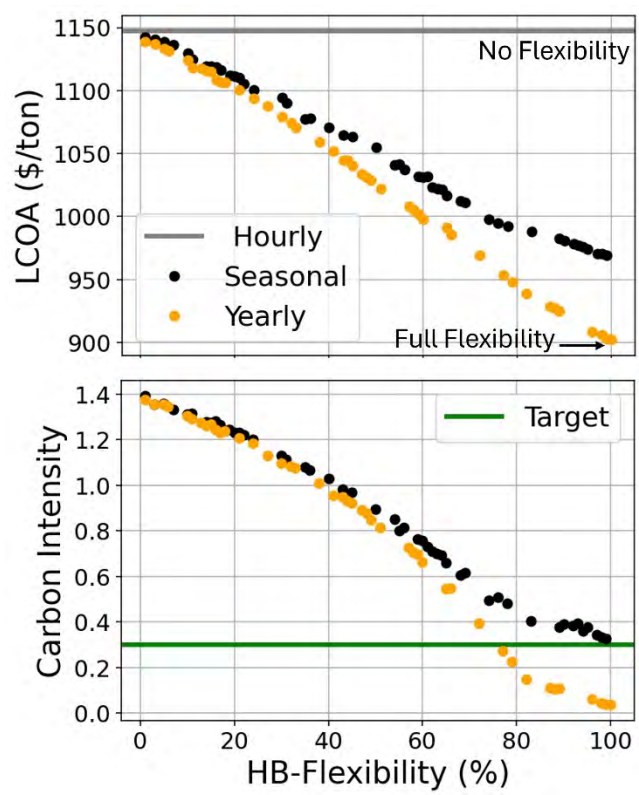

Figure 4. Levelized cost of ammonia and carbon intensity (ton-CO2/ton-NH_3) for different levels of Haber-Bosch (HB) flexibility with hourly, seasonal or yearly demand.

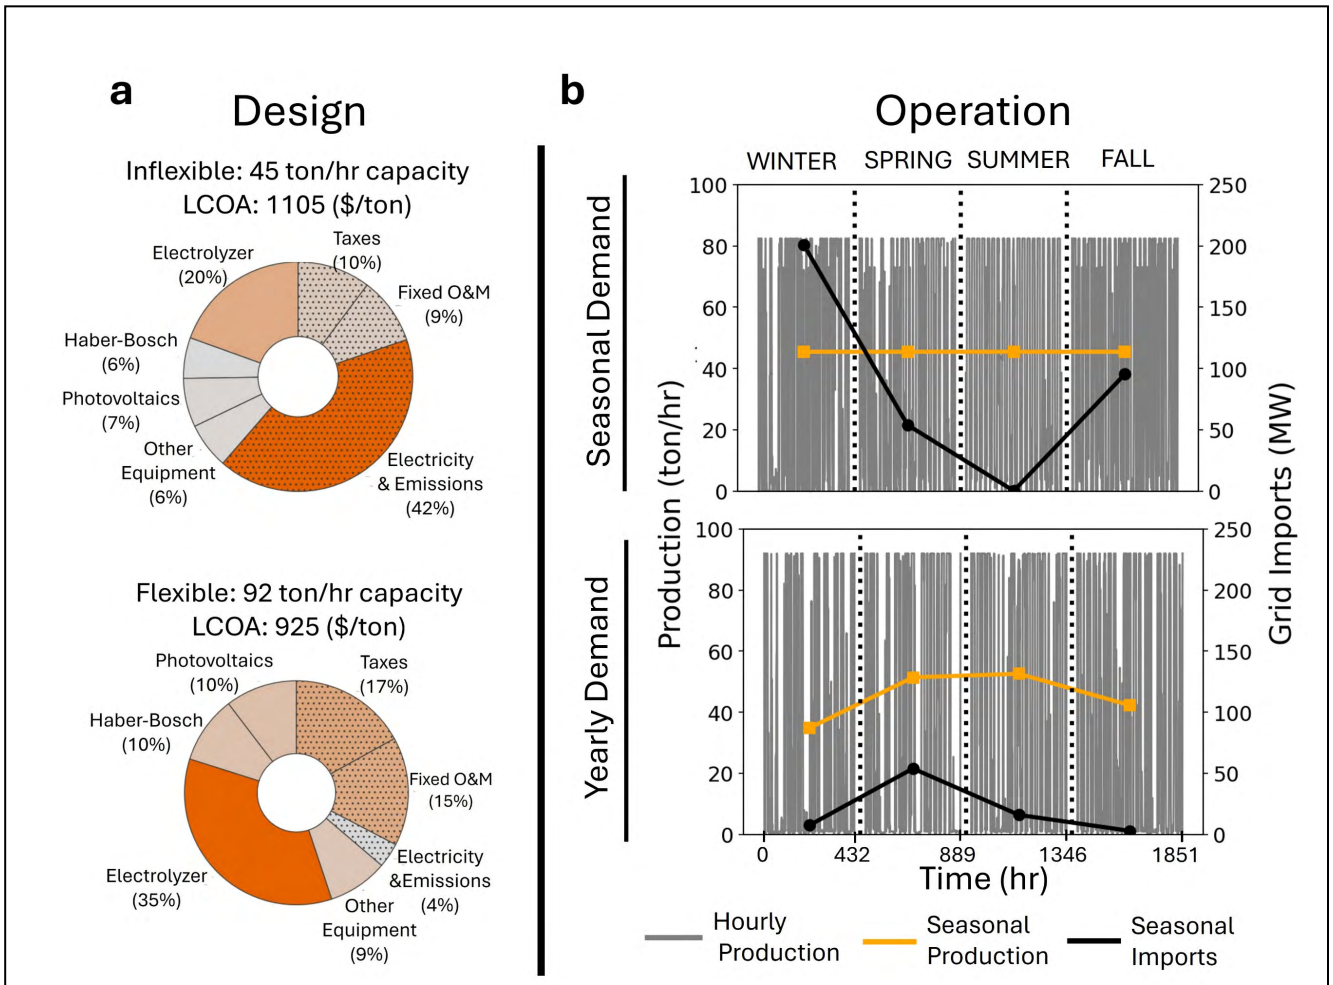

Figure 5. Contributions to the LCOA for an inflexible and flexible system (a). Hourly production of ammonia, as well as the seasonal average production and electricity imports from the power grid (b).

We find that synergies between demand and unit operation flexibility drastically reduce ammonia price and emissions. When considering seasonal or yearly demand, Haber-Bosch flexibility can reduce the LCOA by almost 20% (Figure 4a) without requiring inordinate storage capacity. Increased flexibility in the Haber-Bosch process can also move electrified systems towards, and even below, the target carbon intensity for green ammonia (0.3 ton-CO_2/ton-NH_3); unit flexibility, when coupled with seasonal or yearly demand flexibility, reduces the systems emissions by almost 70% and 100%, respectively.

When examining the LCOA of inflexible and flexible processes, we see that a high level of flexibility results in an islanded process that purchases very little electricity from the power grid. The flexible design is dominated by capital expenditure and is designed at approximately twice the nominal capacity of the inflexible design (Figure 5a). Examining the operation, it is observed that flexible processes produce twice as much material per hour during the day when compared to the inflexible designs but forgo operation during hours of low solar irradiance. Flexible systems benefit from increased economies of scale, and because they do not rely on the grid for power, they also are insensitive to spikes in electricity price and emissions, which often coincide with hour of low solar irradiation.

While the difference in the LCOA and carbon intensity between systems with seasonal and yearly demand is negligible at low to moderate levels of unit flexibility, the difference becomes pronounced at high levels of unit flexibility. Systems with seasonal demand have a different relationship with the power grid than those with yearly demand. Irradiance is low in the winter and fall, so systems with seasonal demand meet the demand requirement by supplementing their onsite power production with grid electricity. For yearly demand, higher production during the spring and summer make up for decreased production during the winter and fall. These seasonal changes in operation allow systems with yearly demand flexibility to achieve almost complete utility independence (Figure 5b); electricity is purchased sparingly from the grid during the spring and summer, and only because very low prices and emissions are achieved during the day.

CONCLUSIONS

Using a framework for operation-informed design, we have modeled a flexible, electrified ammonia synthesis system. We use this model to quantify how the levelized cost of ammonia changes for different grid emissions and understand the transitions in design and operation associated with these different grids. We also find non-monotonicity in the carbon intensity of this system due to transitions in production capacity. At low grid emissions, large, grid-reliant processes are optimal due to economies of scale. These designs remain optimal even as grid emissions rise, causing a sharp increase in carbon intensity. At emission rates of approximately 140 kg-CO_2/MWh, the operating expenditure for electricity and emissions penalties justify a significant downscaling of production capacity. Smaller systems can better leverage onsite energy and mass storage, operating with increased flexibility and reduced grid-reliance. It is important to note that these systems do utilize more land relative to production when compared to grid-reliant systems. Considering the ecological impact of land usage was outside the scope of this work, but in regions with limited land availbility, or sensitive ecosystems, solar power may not be a viable option physically or environmentally. The trends identified in this analysis highlight the importance of operation-informed design, and the role of flexibility in achieving ammonia decarbonization.

We also identify synergies between demand and unit-operation flexibility. We find that improved Haber-Bosch flexibility can lower LCOA by 20% and significantly reduce carbon-intensity for systems with seasonal or yearly demand. Highly flexible systems are designed at much higher capacities than inflexible systems but only operate during hour of peak solar irradiation. Flexible systems are not dependent on the grid, even during hours of low irradiance. Systems with fully flexible unit operations and seasonal demand can achieve a carbon-intensity of approximately 0.3 ton-CO_2/ton-NH_3 (a target for green ammonia). With yearly demand, these systems can achieve nearly zero emissions by adjusting operation to match seasonal changes in energy availbility.

The results demonstrate the importance of flexibility in achieving economic and environmental targets. They also illustrate synergy between multiple forms of flexibility: unit-level, demand-side, storage. Because optimal flexible systems can have drastically different designs from their inflexible counterparts (and even from one other), it is important to evaluate all contributions to flexibility. Operation-informed design frameworks are particularly powerful for these analyses. While there are many challenges to decarbonization through industrial electrification, studying optimal design and operation together uncovers nuanced trends in these systems, and allows us to make informed and intelligent decision about the transition to sustainability.

DIGITAL SUPPLEMENTARY MATERIAL

Supplementary Text S1: Detailed system specifications, parametric data, and additional figures can be found at LAPSE:2026.0018.

ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2444107. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. The authors would also like to thank Siemens for additional support.

AUTHOR IDENTIFIERS

Author ORCIDs:
Kalamaris NN: 0000-0002-3886-1979
Maravelias CT: 0000-0002-4929-1748

REFERENCES

1. Kalamaris NN, Maravelias CT. A general optimization framework for designing chemical & energy systems subject to multi-scale temporal variability. Computers & Chemical Engineering 203:109315 (2025). https://doi.org/10.1016/j.compchemeng.2025.109315
2. Pistolesi C, Giaconia A, Bassano C, De Falco M. Flexible green ammonia production: impact of process design on the levelized cost of ammonia. Fuels 6:39 (2025). https://doi.org/10.3390/fuels6020039
3. Nayak-Luke R, Bañares-Alcántara R, Wilkinson I. “green” ammonia: impact of renewable energy intermittency on plant sizing and levelized cost of ammonia. Ind. Eng. Chem. Res. 57:14607-14616 (2018). https://doi.org/10.1021/acs.iecr.8b02447
4. Pozo D, Sauma E, Bolado-Lavín R. Levelized cost analysis for renewable ammonia production in chile. Energy 335:137554 (2025). https://doi.org/10.1016/j.energy.2025.137554
5. Smith C, Torrente-Murciano L. Cost efficiency versus energy utilization in green ammonia production from intermittent renewable energy. Nat Chem Eng 2:261-272 (2025). https://doi.org/10.1038/s44286-025-00207-9
6. Palys MJ, Moot T, Parvathikar S, Zhang K, Kavale CA, Gupta JK, Daoutidis P. Techno-economic

optimization of flexible distributed green ammonia production. International Journal of Hydrogen Energy 206:153230 (2026). https://doi.org/10.1016/j.ijhydene.2025.153230
7. Kong B, Zhang Q, Daoutidis P. Nonlinear model predictive control of flexible ammonia production. Control Engineering Practice 148:105946 (2024). https://doi.org/10.1016/j.conengprac.2024.105946
8. Moore FC, Drupp MA, Rising J, Dietz S, Rudik I, Wagner G. Synthesis of evidence yields high social cost of carbon due to structural model variation and uncertainties. *PNAS* 121:e2410733121 (2025) https://doi.org/10.1073/pnas.2410733121
9. CAISO. https://www.caiso.com/todays-outlook/prices
10. Singularity. https://singularity.energy/open-grid-emissions
11. NREL. https://nsrdb.nrel.gov/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimizing Renewable Energy Storage Systems to Accelerate Sustainable Data Center Deployment

Matthew J. Palys and Prodromos Daoutidis*

Department of Chemical Engineering & Materials Science, University of Minnesota Twin Cities, Minneapolis, MN 55455, USA
* Corresponding Author: daout001@umn.edu

ABSTRACT

Behind-the-meter generation from variable renewable energy is a potential pathway for new data centers to obtain power more quickly and more sustainably than interconnecting to existing electrical grids. Energy storage is needed to accommodate the variability of wind and solar energy across multiple timescales. Hydrogen from electrolysis and ammonia made from this hydrogen can be used as fuel for dispatchable power generation while offering lower $/MWh storage costs than batteries. In this work, we analyze the economics of using hydrogen, and/or ammonia along with batteries in hybrid energy storage systems to enable data centers to be powered by 100% renewables. We perform this analysis using an optimization model for the selection, sizing, and coordinated hourly operation of constituent energy storage technologies toward minimizing the levelized cost of energy (LCOE). The model uses an hourly resolution scheduling horizon of five years to account for hourly, seasonal, and interannual renewable variability. We use this model to specifically analyze the economics of to 5 MW, 50 MW, and 500 MW sustainable data center energy systems in three relevant American locations. We find that including ammonia energy storage in hybrid systems with batteries and hydrogen enables lower LCOE across all scales and locations by reducing energy storage capital investments despite deploying an order-of-magnitude more storage capacity while also enabling less investment in renewable generation. Ammonia is used for seasonal and interannual energy storage and its inclusion in these hybrid storage systems is most advantageous for larger-scale data centers in locations with high wind potential.

Keywords: Optimization, Energy Storage, Renewable and Sustainable Energy, Technoeconomic Analysis, Alternative Fuels

INTRODUCTION

In the United States, the development of renewable generation projects outpaces the addition of the new transmission infrastructure needed to maximize the use of this renewable power within existing electrical grids [1]. In fact, the average wait time for grid interconnection requests for U.S.-based generation projects is five years [2]. The wait time of this interconnection queue is similar for large industrial customers who wish to purchase power from the grid. As a result, new industrial power users who wish to more quickly begin operations prefer “behind-the-meter” power supply wherein they are directly connected to a generation source. Data centers are an increasingly important example of these industrial consumers. At present, they are being rapidly deployed to facilitate the growing adoption of artificial intelligence. By 2030, data centers are projected to be responsible for at least 10% of total U.S. energy consumption [3, 4]. Using behind-the-meter renewables to provide power for data centers provides an opportunity to ensure that this growing customer segment is sustainable from an energy perspective while also lessening the impact of interconnection constraints on the continued development of renewable generation projects.

Energy storage technologies are a necessary aspect of meeting data center power demands, which are generally constant over time, using variable renewable energy. This variability manifests over multiple timescales from within a day to between consecutive years. Batteries are the most well-developed energy storage technology and are well-suited for managing short-term variability, but their deployment is limited by high energy storage costs for long-term supply-demand imbalances. As

such, previous research on energy storage systems for data centers has considered the use of dispatchable fossil-based generation to limit the required size of battery installations [5, 6]. Hydrogen obtained from water via renewable-powered electrolysis can act as dispatchable fuel source without incurring any carbon dioxide emissions upon reconversion to electrical power. The economics of hydrogen energy storage have previously been investigated in the context of data centers [7, 8]. In [7], hydrogen storage was used in a grid-connected data center to facilitate demand response. In [8], hydrogen storage was coupled with bifacial PV in Arctic locations to enable 100% renewable operation. To this end, over 1, 000 metric tons (mt) of hydrogen storage was needed, which may not be suitable for all locations. Furthermore, analysis in [8] did not consider interannual solar generation variability, the management of which could require even more energy storage capacity in a behind-the-meter, 100% renewables system.

Converting hydrogen to ammonia by combining it with nitrogen separated from air can facilitate higher energy storage capacities than hydrogen itself. Ammonia energy storage has been of particular interest over the past decade for high capacity, long duration energy storage due to higher achievable storage volumes as well as lower storage costs [9]. As with hydrogen, ammonia can be used for dispatchable power generation via fuel cells or combustion [10, 11]. To our knowledge, the economics of ammonia energy storage in the context of data centers have not been previously analyzed. We have previously investigated the economics of ammonia energy storage for 100% renewable energy storage in 1 MW residential microgrids U.S. throughout the continental U.S. [12] and combined heat and power systems in remote U.S. locations [13]. Ammonia energy storage economics have also been analyzed in facilitating the energy transition of different provinces in Spain [14], enabling exclusively solar power to meet 4 MW residential power demand in Florence, Italy [15], and using ammonia for both energy storage and export in Can Dao, an archipelago off Vietnam's Southern Coast [16].

These previous efforts employed quantitative modeling approaches to determine the best deployment of ammonia energy storage to manage renewable variability. Specifically, these works used optimization models wherein technology selection, sizing, and operational scheduling decisions are made simultaneously toward lowest cost. This is the state-of-the-art approach for modeling renewables-based chemical production and utilization [17, 18]. In the context of ammonia energy storage, this optimization approach has been implemented using mixed integer linear programming (MILP) models solved via branch-and-cut algorithms [12-14, 16] or using software based on hierarchical operational rules for scheduling [15].

Scheduling decisions in these optimization models are typically made over a yearlong horizon with hourly resolution, including in [12-16]. This allows optimal energy storage system design to account for renewable variability across hourly to seasonal timescales. A representative year which agglomerates multiple years of historical or forecasted weather data is usually selected for this purpose. However, previous work on ammonia energy storage has not considered interannual variability in renewable generation. In general, renewable potential can vary considerably over consecutive years and its consideration in the design of renewable energy systems can increase reliability [19]. Interannual variability has only been previously considered in the economic analysis of renewable ammonia production [20]. We hypothesize that ammonia's low storage cost can facilitate economical interannual energy storage.

In this work, we investigate the economics of using hydrogen and/or ammonia along with batteries in hybrid energy storage systems to facilitate meeting data center power demands using only renewable energy. We specifically formulate a combined design and scheduling model which employs multiple years of historical, hourly resolution time series data to capture interannual renewable variability as well as hourly and seasonal variability which have been considered in previous work. We use this model to analyze the optimal economics of sustainable data center energy systems at 5 MW, 50 MW, and 500 MW power demand scales in three American metropolitan areas: Northern Virginia (NoVa), Dallas-Fort Worth (DFW), and Las Vegas. The goal of this work is to determine the merit of hydrogen and ammonia energy storage toward providing economical and sustainable energy for data centers, the most rapidly growing segment of industrial power consumption.

CASE STUDY DESCRIPTION

Battery-H_2-NH_3 Energy Storage

The energy storage system we analyze in this work must meet data center power demand using energy generated exclusively from wind and solar resources and energy storage (see Figure 1). Energy can be stored directly in its electrical form using battery storage. It can also be used to produce hydrogen via water electrolysis which can be stored as gas at 200 bar and used for dispatchable power generation. Electrical energy can also be used to separate nitrogen from air and combine it with hydrogen to make ammonia, the synthesis of which is also electrically powered. Ammonia can be stored as liquid and used for dispatchable power generation.

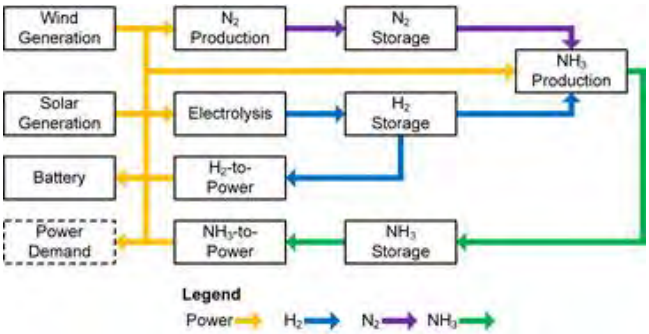

Figure 1. Conceptual diagram of 100% renewable energy system with battery, hydrogen, and ammonia energy storage pathways.

Combined Optimal Sizing & Scheduling Model

The optimization model we use minimizes the levelized cost of energy (LCOE) of meeting data center power demands using behind-the-meter renewables and hybrid energy storage. The LCOE includes capital investments (CapEx) of all technologies in Figure 1, annualized by a capital recovery factor (CRF), as well as annual operating costs (OpEx):

$$LCOE\left(\frac{\$}{MWh}\right) = \frac{CapEx \times CRF + OpEx}{Demand} \quad (1)$$

The model simultaneously optimizes the selection, sizing, and operational scheduling of all constituent technologies in Figure 1. Optimal scheduling decisions for battery state-of-charge and charge/discharge rates, chemical production rates and storage inventories, and dispatchable power generation rates are made across a multi-year horizon of hourly resolution renewable availability time series data.

Hourly power balance constraints ensure that renewable generation plus battery discharge plus dispatchable hydrogen- or ammonia-fueled power generation meet or exceed constant data center power demand plus battery charge plus chemical production power demands. Hourly mass balance equations for hydrogen and ammonia govern the accumulation of their storage inventory between consecutive hours. This storage inventory is constrained to be cyclical between the first and last hour of the multi-year operating horizon. Chemical production and dispatchable power generation technologies are constrained to operating limits defined by their respective operational flexibility. This model is formulated as a MILP and solved using CPLEX in GAMS. Please see our previous work [12] for more technical details, including model equations.

U.S. Data Centers Case Study Data

We perform a case study considering data centers with power demand scales of 5 MW, 50 MW, and 500 MW in NoVa, Dallas-Fort Worth, and Las Vegas. NoVa has highest current data center capacity in United States. DFW and Las Vegas respectively have high wind and solar potential, and data center capacity is rapidly growing in both metropolitan areas. We optimize the LCOE of meeting data center power demands via (i) battery, (ii) battery-hydrogen, and (iii) battery-hydrogen-ammonia energy storage systems using five years of historical renewable generation data from 2019 to 2024 in each location. The annual average wind and solar capacity factors in NoVa, DFW, and Las Vegas are respectively 39%, 45%, and 27% and 16%, 18%, and 22% [21].

Table 1: Capital and operating cost coefficients for renewable generation, battery charge/discharge capacity and electrolysis at different power demand scales. Data for all technologies except for electrolysis are from [22] while electrolysis data are from [23].

Technology at Demand Scale (MW)	CapEx ($MM/MW)	OpEx (%CapEx)
Wind (5)	2.04	1.86
Wind (>50)	1.41	2.08
PV (5)	1.49	1.06
PV (>50)	1.19	1.51
Battery - Power (5)	0.88	2.41
Battery - Power (>50)	0.43	2.15
Electrolysis (5)	1.99	2.5
Electrolysis (50)	1.37	2.5
Electrolysis (500)	0.94	2.5
ICE Genset (5)	1.27	2.0
CCGT (>50)	1.43	2.2

The capital and operating costs of energy system technologies depend on scale (see Table 1). Battery energy capacity (in MWh) has a capital cost of $0.18/MWh and $0.26/MWh respectively at the 5 MW and >50 MW demand scales [22]. Batteries have a charge efficiency of 90% and a discharge efficiency of 95%. PEM electrolysis is considered in this work for its wide operating range and fast dynamics. Its energy efficiency is assumed to be 51 MWh/mtH_2 [23]. High pressure gaseous hydrogen storage has a capital cost of $0.9MM/mt$H_2$. Reconversion of hydrogen and ammonia to electrical power occurs through combustion in internal combustion engine gensets at the 5 MW scale and combined cycle gas turbine (CCGT) plants at 50 MW and 500 MW scales. ICE genset and CCGT efficiencies are respectively 35% and 60% of LHV. ICE gensets have 90% turndown (i.e., can operate between 10% and 100% of installed capacity) while CCGTs have 60% turndown. The capital costs of nitrogen and ammonia production are subject to conventional power law cost scaling (see Table 2). Each technology has an annual operating cost of 2% of capital cost. Nitrogen and ammonia storage have capital costs of $\$5\times10^{-2}MM/mtN_2$ and $\$1.5\times10^{-3}MM/mtNH_3$ respectively [24].

Total energy system capital cost is annualized by a

Table 2: Capital cost correlations and performance data for nitrogen and ammonia production technologies [24].

	Capacity (mt/y)	CapEx (MM$)	Scaling Exponent	Efficiency (MWh/mt)	Turndown (%)
Pressure swing adsorption (PSA)	10,000	6.6	0.75	0.25	90
Cryogenic distillation	197,500	29.5	0.55	0.18	50
Haber-Bosch synthesis	237,500	89	0.65	0.45	90

Table 3: Optimal LCOE, renewable generation and energy storage capacity, and renewable curtailment for 50 MW data centers in Northern Virginia. Hydrogen and ammonia storage capacities in MWh are based on LHV.

Storage Pathway	Battery	Battery-H_2	Battery-H_2-NH_3
LCOE ($/MWh)	328	185	140
Wind (MW)	192	198	124
PV (MW)	556	130	104
Battery (MWh)	2,115	405	441
H_2 Storage (MWh)	0	8,382	255
NH_3 Storage (MWh)	0	0	43,340
Wind Curtailment (%)	62	33	10
PV Curtailment (%)	73	22	17

CRF of 0.073 y^{-1}. This corresponds to a project lifetime of 30 years with a 6.04% weighted average cost of capital (WACC). This is the average WACC for green and sustainable energy projects in the United States [25].

RESULTS AND DISCUSSION

Incorporating hydrogen and ammonia as lower cost energy storage media into hybrid renewable energy storage systems for sustainable data centers results in meaningfully lower LCOE than using battery storage alone (see Table 3). Combining hydrogen storage with battery storage results in a 44% LCOE reduction. This is driven by considerably lower solar generation and battery costs. Battery energy storage is too expensive to include enough capacity to store close to all generated MWh of renewable energy specifically due to variability over longer timescales. The battery-only energy storage system thus overbuilds solar generation to enable more demand to be met via direct generation and respectively curtails 62% and 73% of wind and solar over its modeled operational lifetime. Nonetheless, more than 2 GWh of battery storage is still required. In contrast, the battery-hydrogen system installs around 4x more total energy storage capacity despite 5x less battery capacity than the battery system to facilitate more effective renewable utilization. The optimal solar generation capacity for the battery-hydrogen system is 77% less than the battery-only system while optimal wind and solar curtailment are 33% and 22% respectively. Overall, the hybrid battery-hydrogen system still uses batteries for short-term storage and hydrogen for longer-term storage.

Using ammonia in a hybrid battery-hydrogen-ammonia energy system further reduces the LCOE by 24%. The same general trends as described in the previous paragraph are more pronounced. Ammonia's low storage cost results in approximately 20x more total optimal storage capacity than the battery-only system as well as approximately 5x more total storage capacity than the battery-hydrogen system despite 97% less hydrogen storage. Consequently, renewable energy generation capacities and curtailment are both reduced relative to the energy systems where energy storage technologies have higher $/MWh storage costs. Comparing battery-hydrogen and battery-hydrogen-ammonia systems specifically, we note that incorporating ammonia allows for considerably more effective use of wind energy with 37% less wind capacity installed and wind curtailment reduced from 33% to 10%. Wind potential in NoVa exhibits more seasonal and interannual variability than solar potential, and ammonia is more cost effective at balancing supply-demand mismatches that occur over longer timescales.

Interestingly, the optimal hybrid battery-hydrogen-ammonia energy system does not install any hydrogen-to-power capacity. In this system, the installed hydrogen storage only acts as a buffer between electrolysis and the less flexible Haber-Bosch synthesis. This occurs because, for the combustion technologies considered in this work, the roundtrip efficiency of the hydrogen storage pathway is only 39% compared to 32% for the ammonia storage pathway, but hydrogen storage cost is 100x higher on a $/MWh basis. Ultimately, the hybrid battery-hydrogen-ammonia system uses battery storage for short-term cycling and ammonia for longer-duration storage with hydrogen only acting as an intermediate. Specifically, batteries are cycled on short timescales to take advantage of their high roundtrip efficiency. Figure 2

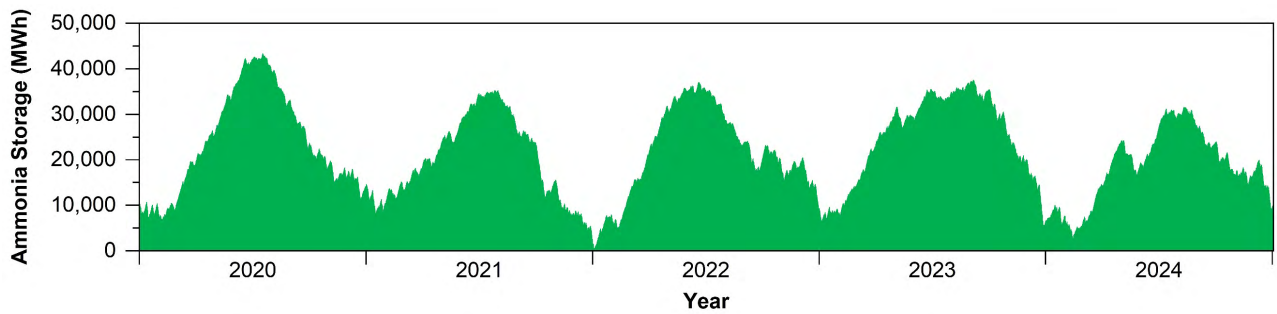

Figure 2. Optimal ammonia energy storage schedule for 50 MW data center using battery-hydrogen-ammonia storage pathway in Northern Virginia. Ammonia storage capacity in MWh based on LHV.

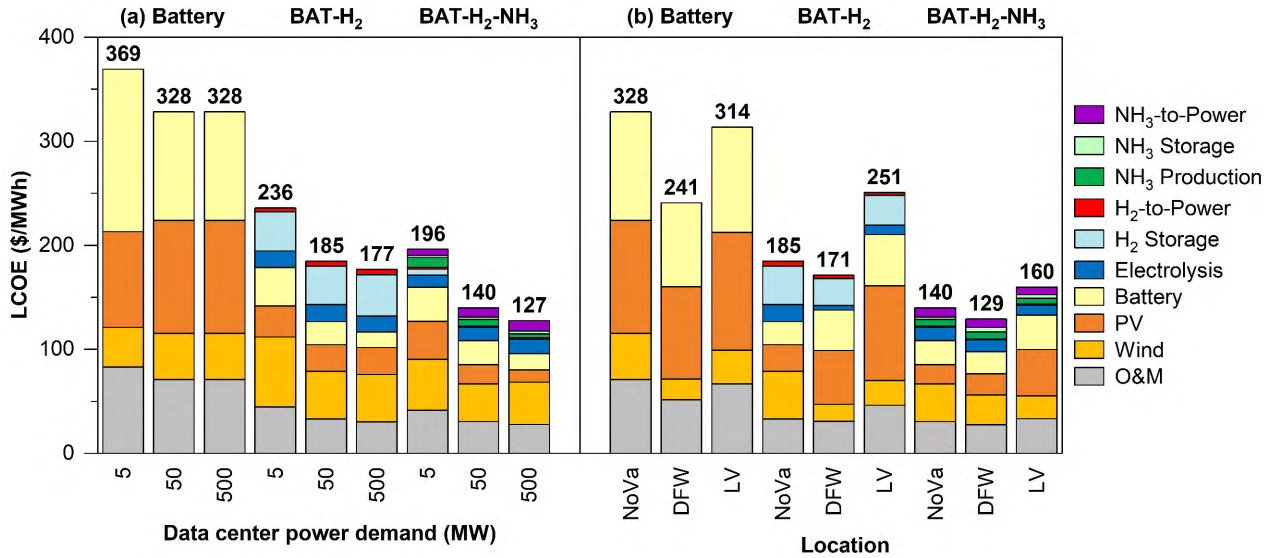

Figure 3: Breakdown of optimal LCOE for (a) 5 MW, 50 MW, and 500 MW data centers in Northern Virginia and (b) 50 MW data centers in Northern Virginia, Dallas-Fort Worth, and Las Vegas. Nitrogen production and storage as well as Haber-Bosch synthesis are included in "NH_3 Production".

shows the optimal ammonia energy storage schedule over for optimal operation from 2019 to 2024. It is evident that ammonia is accumulated through the first half of each year before being consumed to accommodate for seasonally low wind generation in second half of each year. From an interannual standpoint, the ammonia inventory is depleted by the end of 2021 to account for its lowest wind capacity factor while the highest net ammonia accumulation occurs in 2022 because its capacity factor is the highest.

Analysis of Power Demand Scale

Sustainable data center using 100% renewables and energy storage exhibit clear economies of scale (see Figure 3a). The LCOE for the 5 MW demand scale is considerably higher than for 50 MW and 500 MW demand scales regardless of the energy storage pathways under consideration. At this small scale, all technologies from wind and solar generation through to ammonia synthesis have higher capital intensity. The loss in economies of scale for nitrogen and ammonia production is particularly impactful. Data center energy systems at the 5 MW demand scale benefit least from the incorporation of ammonia relative to LCOE reductions at 50 MW and 500 MW demand scales. Additionally, the hybrid battery-hydrogen-ammonia storage system at the 5 MW demand scale uses both hydrogen- and ammonia-to-power combustion technologies, unlike the larger systems which only use ammonia for dispatchable power generation. This occurs to reduce the capacity and accompanied higher capital intensity of small-scale ammonia synthesis despite hydrogen having a higher storage cost.

The economic outlook of powering 50 MW and 500 MW data centers with renewables is similar. Renewable generation and batteries exhibit little to no cost savings with this scaleup while electrolysis does exhibit some economies of scale. This is evident in comparing the battery and battery-hydrogen LCOEs which are reduced by

less than $10/MWh with the 10x scaleup from 50 MW to 500 MW. Ammonia energy storage is well-suited for large centers on the order of 500 MW. Incorporating ammonia at this scale results in deeper cost reductions relative to smaller scales. Ammonia is also important for energy storage at hyperscale data centers from an optimal capacity standpoint. The optimal battery and hydrogen storage capacities are 21, 000 MWh and 3, 200 mtH_2 respectively when these are the standalone storage pathways for 500 MW NoVa data centers. These capacities are not realistic relative to current deployments, requiring 7x more battery capacity [26] and 10x more hydrogen storage capacity [27] than the largest facilities for each currently operational in the United States. Conversely, the hybrid battery-hydrogen-ammonia data center requires only 2, 800 MWh battery and 100 mtH_2 because long-duration, high-capacity storage can be facilitated by 105, 000 $mtNH_3$ ammonia storage.

Large-scale data centers are also best suited to ammonia energy storage from a safety standpoint. The ammonia synthesis and storage capacities these data centers would require are similar to those already used in the ammonia fertilizer supply chain. In contrast, using ammonia for energy storage across many small data centers would increase the number of facilities at which ammonia must be stored and handled and would thus increase the associated risk.

Analysis of Data Center Location

Hybrid battery-hydrogen and even more so hybrid battery-hydrogen-ammonia storage offers the same qualitative advantages across the three studied locations. Storage pathways with lower storage costs than batteries enable more complete utilization of lower optimal renewable generation capacities, which results in lower generation cost contributions to the LCOE, especially for solar (see Figure 3b). Battery storage cost contributions are also meaningfully less when storage can occur in the form of hydrogen and/or ammonia. When the ammonia storage pathway is considered, the optimal energy storage systems for 50 MW demand data centers in DFW and Las Vegas only use ammonia for dispatchable power generation as was previously discussed for the optimal hybrid battery-hydrogen-ammonia NoVa energy storage system.

The location of data centers does have a quantitative effect of the economics of sustainable data centers for each storage pathway (see Figure 3b). For battery-only storage systems, NoVa has the highest LCOE. Both DFW and Las Vegas have higher solar potential and less seasonal solar variation, making these locations better suited to batteries alone. Hydrogen and even more so ammonia energy storage enables better use of wind generation. This results in Las Vegas having the highest LCOE for these hybrid systems because its wind potential is lower than NoVa or DFW and must install considerably more solar generation as a result. Conversely, DFW has better wind and solar potential than NoVa and considerably better wind potential than Las Vegas. This results in the DFW system having the lowest LCOE regardless of the storage pathways under consideration because these systems have lower total renewable generation and total energy storage costs. Nonetheless, hybrid battery-hydrogen-ammonia energy storage results in a NoVa LCOE only 9% higher than the DFW LCOE. In comparison, the NoVa battery LCOE is 36% higher than the DFW battery LCOE. Overall, ammonia storage makes sustainable data centers more economically competitive in more locations.

CONCLUSIONS

In this work, we analyze the economics of using hydrogen and/or ammonia along with batteries in hybrid energy storage systems to meet data center power demands using 100% behind-the-meter renewable energy. We perform this analysis through the lens of an economic optimization model which minimizes the levelized cost of energy (LCOE) by simultaneously making technology selection, sizing, and operational scheduling decisions. Optimal scheduling occurs for a five-year operating horizon with hourly resolution based on time series data for historical renewable availability to ensure that optimal selection and sizing decisions account for hourly, seasonal, and interannual variability. We use this model to determine the optimal LCOE of using 100% renewable and hybrid energy storage systems for data centers with 5 MW, 50 MW, and 500 MW power demand in Northern Virginia (NoVa), Dallas-Fort Worth (DFW), and Las Vegas

We find that hybrid battery-hydrogen and battery-hydrogen-ammonia energy systems provide lower LCOE than using batteries alone for all scales and locations. These lower storage cost pathways enable larger storage capacity so that more energy can be economically stored and used rather than curtailed, resulting in less investment in renewable generation. Hybrid battery-hydrogen-ammonia systems provide the lowest LCOE and ammonia provides at least 95% of total energy storage capacity in these hybrid systems due to its low storage cost. Ammonia energy storage is best suited to large-scale data centers due to the economies of scale for nitrogen and ammonia production. It is also best suited to locations with high wind potential, the seasonal and interannual variability of which can be economically balanced with data center power demands via ammonia storage.

Future work will compare the viability of ammonia energy storage with conventional fossil fuel technologies as well as small modular nuclear reactors in the context of data centers across different U.S. locations. This future work will include techno-economic analysis of these

energy storage pathways under a consistent set of assumptions. We will also compare the emissions and water intensity of these different energy storage pathways through a life cycle assessment framework.

ACKNOWLEDGEMENTS

This work was funded in part by the Advanced Research Projects Agency-Energy (ARPA-E), US Department of Energy, United States of America, under Award Number DE-AR0001479. The views and opin-ions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

AUTHOR IDENTIFIERS

Author ORCIDs:
Palys MJ: 0000-0002-3738-6846
Daoutidis P: 0000-0003-2459-6817

REFERENCES

1. Gorman W, Kemp JM, Rand J, Seel J, Wiser R, Manderlink N, Kahrl F, Porter K, Cotton W. Grid connection barriers to renewable energy deployment in the united states. Joule 9:101791 (2025). https://doi.org/10.1016/j.joule.2024.11.008
2. Rand J, Strauss R, Gorman W, Joachim S, Kemp J, Seongeun J, Robson D, & Wiser R. Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection As of the End of 2022. Lawrence Berkeley National Laboratory (2023) https://doi.org/10.2172/1969977
3. Shehabi A, Smith SJ, Hubbard A, Newkirk A, Lei N, Siddik MAB, Holecek B, Koomey J, Masanet E, & Sartor D. 2024 United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory (2024) eta.lbl.gov/publications/2024-lbnl-data-center-energy-usage-report
4. Electric Power Research Institute. Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption (2024) epri.com/research/products/3002028905
5. He W, Xu Q, Liu S, Wang T, Wang F, Wu X, Wang Y, Li H. Analysis on data center power supply system based on multiple renewable power configurations and multi-objective optimization. Renewable Energy 222:119865 (2024). https://doi.org/10.1016/j.renene.2023.119865
6. Ren X, Wang J, Hu X, Sun Z, Zhao Q, Chong D, Xue K, Yan J. A novel demand response-based distributed multi-energy system optimal operation framework for data centers. Energy and Buildings 305:113886 (2024). https://doi.org/10.1016/j.enbuild.2024.113886
7. Long X, Li Y, Li Y, Ge L, Gooi HB, Chung C, Zeng Z. Collaborative response of data center coupled with hydrogen storage system for renewable energy absorption. IEEE Trans. Sustain. Energy 15:986-1000 (2024). https://doi.org/10.1109/tste.2023.3321591
8. Temiz M, Dincer I. A unique bifacial PV and hydrogen-based cleaner energy system with heat recovery for data centers. Applied Thermal Engineering 206:118102 (2022). https://doi.org/10.1016/j.applthermaleng.2022.118102
9. Sepulveda NA, Jenkins JD, Edington A, Mallapragada DS, Lester RK. The design space for long-duration energy storage in decarbonized power systems. Nat Energy 6:506-516 (2021). https://doi.org/10.1038/s41560-021-00796-8
10. Cesaro Z, Ives M, Nayak-Luke R, Mason M, Bañares-Alcántara R. Ammonia to power: forecasting the levelized cost of electricity from green ammonia in large-scale power plants. Applied Energy 282:116009 (2021). https://doi.org/10.1016/j.apenergy.2020.116009
11. Sánchez A, Castellano E, Martín M, Vega P. Evaluating ammonia as green fuel for power generation: a thermo-chemical perspective. Applied Energy 293:116956 (2021). https://doi.org/10.1016/j.apenergy.2021.116956
12. Palys MJ, Daoutidis P. Using hydrogen and ammonia for renewable energy storage: a geographically comprehensive techno-economic study. Computers & Chemical Engineering 136:106785 (2020). https://doi.org/10.1016/j.compchemeng.2020.106785
13. Palys MJ, Mitrai I, Daoutidis P. Renewable hydrogen and ammonia for combined heat and power systems in remote locations: optimal design and scheduling. Optim Control Appl Methods 44:719-738 (2021). https://doi.org/10.1002/oca.2793
14. Sánchez A, Zhang Q, Martín M, Vega P. Towards a new renewable power system using energy storage: an economic and social analysis. Energy Conversion and Management 252:115056 (2022). https://doi.org/10.1016/j.enconman.2021.115056
15. Veltroni V, Calabrese M, Carcasci C, Pagliai M. Techno-economic assessment of power-to-ammonia-to-power for long-term energy storage. Energy Conversion and Management 348:120621 (2026). https://doi.org/10.1016/j.enconman.2025.120621
16. Su K, Wang L, Jiang D, Lin J, Yang C, Wu F, Liu D, Wu F, Zheng C, Gao X. Towards zero-carbon

power supply: development of a multi-energy coupled offshore green energy production island. Sustainable Energy Technologies and Assessments 85:104791 (2026). https://doi.org/10.1016/j.seta.2025.104791
17. Schack D, Rihko-Struckmann L, Sundmacher K. Linear programming approach for structure optimization of renewable-to-chemicals (r2chem) production networks. Ind. Eng. Chem. Res. 57:9889-9902 (2018). https://doi.org/10.1021/acs.iecr.7b05305
18. Palys MJ, Daoutidis P. Renewable ammonia for islanded energy storage. Current Opinion in Green and Sustainable Chemistry 48:100946 (2024). https://doi.org/10.1016/j.cogsc.2024.100946
19. Javed MS, Jurasz J, Guezgouz M, Canales FA, Ruggles TH, Ma T. Impact of multi-annual renewable energy variability on the optimal sizing of off-grid systems. Renewable and Sustainable Energy Reviews 183:113514 (2023). https://doi.org/10.1016/j.rser.2023.113514
20. Salmon N & Bañares-Alcántara R (2022). Importance of interannual renewable energy variation in the design of green ammonia plants. In *Comput. Aided Chem. Eng.* 49: 757-762 (2022). https://doi.org/10.1016/B978-0-323-85159-6.50126-3
21. NREL. National solar radiation database (2024) https://nsrdb.nrel.gov/data-viewer
22. NREL. Annual Technology Baseline (2025) https://atb.nrel.gov/electricity/2024b/about
23. EPRI. Electrolysis techno-economic analysis – CapEx rate (2024) https://apps.epri.com/lcri-electrolysis-tea/en/capex-rates.html
24. Palys MJ, Moot T, Parvathikar S, Zhang K, Kavale CA, Gupta JK, Daoutidis P. Techno-economic optimization of flexible distributed green ammonia production. International Journal of Hydrogen Energy 206:153230 (2026). https://doi.org/10.1016/j.ijhydene.2025.153230
25. Damodaran A. Cost of Equity and Capital (US) (2026) https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/wacc.html
26. Thomas M. Cleanview: Battery storage projects in the United States (2024) https://cleanview.co/battery-storage-projects/us
27. Swanger A. World's largest liquid hydrogen tank nearing completion. *Cold Facts* 38(2): 20220004276 (2022). https://ntrs.nasa.gov/citations/20220004276

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Temporal aggregation bias in model-based Direct Air Capture performance under weather variability

Eleni Chalasti[abc], Gbemi Oluleye[bc], Maria M. Papathanasiou[ab], Ronny Pini[ab]*

[a] Imperial College London, Department of Chemical Engineering, London, United Kingdom
[b] Imperial College London, Sargent Centre for Process Systems Engineering, London, United Kingdom
[c] Imperial College London, Grantham Institute - Climate Change and the Environment, London, United Kingdom
* Corresponding Author: rpini@imperial.ac.uk.

ABSTRACT

Direct Air Capture (DAC) is a negative emissions technology whose performance is inherently linked to ambient conditions, which directly affect its primary feed stream (air). A common simplification in DAC model simulations is the use of fixed weather conditions, which can bias the predicted performance under weather variability. In response, this study quantifies the impact of local meteorological variability and temporal weather aggregation on the performance of DAC units. Building on a previously developed and validated 1D mechanistic model of a fixed-bed Steam-assisted Temperature Vacuum Swing Adsorption (S-TVSA) DAC process, we simulate its operation using weather data from the Met Office station at Buchan (UK), near the Saint Fergus terminal - a strategic hub for Carbon Capture and Storage (CCS) activities in Scotland. A two-branch methodological framework is developed combining optimization and forward simulations. Operating conditions are optimized using a multi-objective genetic algorithm (NSGA-II) to maximize productivity (Pr) and minimize specific equivalent work (Weq) at two temporal resolutions. Furthermore, daily weather inputs are aggregated on a monthly and yearly scale to assess the impact of data resolution on model predictions and real operational gains. Results show that temporal weather aggregation to yearly averages biases DAC key performance indicators, overestimating Pr by up to 5% while underestimating Weq by up to 31%, relative to performance based on daily weather variations. Moreover, optimization strategies that explicitly account for monthly weather variability present monthly gains, by increasing Pr by up to 10%. Yet, these monthly gains do not necessarily translate into significant operational performance benefits at the annual scale when daily weather data is propagated in the process model.

Keywords: Carbon Capture, Process Design, Adsorption, Dynamic Modelling, Genetic Algorithm, United Kingdom, Industrial Clusters, Direct Air Capture, Temporal Weather Aggregation

INTRODUCTION

In light of the ongoing need to pursue mitigation pathways towards climate change, consistent with the Paris Agreement targets, Direct Air Capture (DAC) has emerged as a carbon dioxide (CO_2) removal (CDR) option capable of delivering negative emissions by counterbalancing legacy emissions and hard-to-abate residual industrial CO_2 that cannot be eliminated through conventional abatement measures [1]. DAC is a novel yet technically complex CDR technology during which atmospheric CO_2 is artificially removed through an engineered air contactor [2]. Unlike point-source capture technologies, DAC targets the dilute CO_2 in ambient air and is intrinsically sensitive to local weather conditions [3]. Furthermore, although DAC systems are claimed to be location-independent, their operation depends on the availability and integration of utilities, land and energy sources. To avoid the need for building new DAC-specific infrastructure and to reduce their energy consumption, the integration of DAC systems within existing industrial clusters is a beneficial engineering solution. In fact, co-locating DAC units near existing industrial sites can enhance the overall efficiency of the system thanks to the availability of suitable, existing infrastructure, available waste utilities (heat, cooling water) and low likelihood of

environmental constraints [1].

Despite ongoing, promising momentum, the widespread adoption of DAC as a viable CDR at large-scale is currently constrained by several interrelated, unresolved technical complexities that include the integration of resources within existing industrial activity [1], [3], high cost of capture [4], [5], and performance uncertainties due to operational variability under real-world conditions [3], [4], [6].

Dynamic process simulations have highlighted the sensitivity of DAC performance to ambient weather, including spatiotemporal variations in energy demand driven by different temperature (T) and relative humidity (RH) profiles across different locations [4], [7]. To the best of the authors' knowledge, only few studies have quantified the impact of fine weather variability and temporal aggregation on technical key performance indicators (KPIs) [3], [6], [8], employing optimization approaches ranging from mixed-integer linear programming (MILP) to data-driven surrogate-based branch-and-bound (DDSBB) and Bayesian optimization. Most DAC studies report performance under a limited set of representative conditions or rely on temporally averaged weather inputs, without constructing year-long dynamic performance profiles that resolve intra-annual variability. Consequently, the impact of variable ambient conditions on deviations from design-point operation - particularly near constraint limits and along energy-productivity trade-offs - remains insufficiently quantified. In addition, DAC is often assessed in isolation from region-specific deployment constraints such as industrial geography, local climate, and evolving CO_2 transport-and-storage (T&S) infrastructure.

To address this gap, we take an initial step toward a decision-support framework for DAC siting using the UK as a case study. We evaluate a Steam-assisted Temperature-Vacuum Swing Adsorption (S-TVSA) fixed-bed process for DAC at an industrial UK location and develop dynamic performance profiles under off-design ambient conditions. We compare predictions obtained using different weather resolutions to quantify temporal aggregation bias, and we optimize operating variables using NSGA-II to maximize productivity (Pr) and minimize specific equivalent work (Weq) under capture-grade specifications. Finally, we quantify the value of higher optimization resolution by comparing a single yearly optimization against twelve month-specific optimizations accounting for intra-annual variability.

METHODOLOGY

Temperature-Vacuum-Swing Adsorption Cycle

Previous work focused on the development of a high-fidelity mechanistic model in MATLAB to simulate and optimize the performance of an S-TVSA DAC unit, which constitutes the modeling basis for this work [2]. The system separates CO_2 from ambient air containing 400 ppm CO_2, nitrogen (N_2), and water vapor (H_2O) at atmospheric pressure (1 bar) and local ambient temperature, using the benchmark adsorbent APDES-NFC. It comprises five sequential steps: (i) adsorption, (ii) blowdown, (iii) heating, (iv) desorption, and (v) pressurization, as shown in Figure 1. Each step operates under a distinct set of conditions, ranging from low pressure (P) and temperature (P_L, T_L) to high pressure and temperature (P_H, T_H). During adsorption, CO_2 is selectively captured while H_2O co-adsorbs on the sorbent's surface. During blowdown, pressure is reduced to vacuum to remove residual N_2. Steam-assisted heating under vacuum promotes the desorption step, producing a CO_2/H_2O stream which is separated downstream by condensation. At the last step, the bed is then repressurized to initiate the next cycle using air and cooling water.

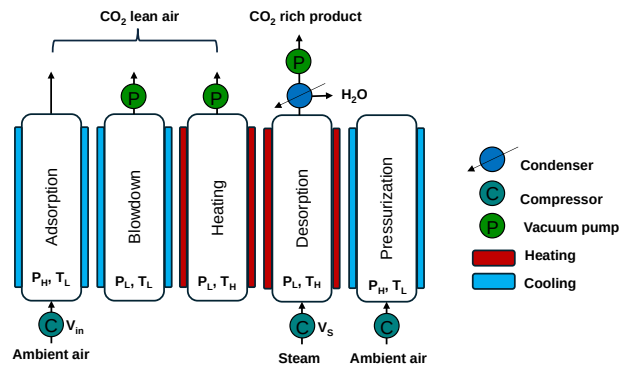

Figure 1. Schematic representation of the 5-step S-TVSA cycle considered in this study.

The high-fidelity, one dimensional (1D) fixed bed model describes competitive CO_2/H_2O adsorption equilibria as functions of pressure, temperature, and gas composition using thermodynamically consistent isotherm models (GAB for H_2O and modified Toth for CO_2 [9]), assuming negligible N_2 adsorption. Adsorption dynamics are captured by coupled partial differential equations for mass, energy, and momentum balances with axial dispersion. Intraparticle mass transfer is captured using a Linear Driving Force (LDF) approximation. The cycle is simulated until cyclic steady state is reached, defined by <0.5% variation in state variables between consecutive cycles at all axial locations.

Multi-objective optimization

The multi-objective optimization problem to minimize Weq and maximize Pr, subject to a product purity constraint (s.t. Pu ≥ 95%) is formulated as shown below:

$$\min(-Pr + \varphi, W_{eq} + \varphi) \quad (1)$$

$$s.t\ \theta_L \le \theta \le \theta_U$$

$$\varphi = 0.25 \times [\max(0, 95 - Pu)]^2$$

where θ is the vector of process operating variables, bounded between a lower (θ_L) and upper (θ_U) value as shown in Table 1. The equivalent work term represents the total energy duty of the system, combining the thermal (heating, steam purge) and electrical energy requirements (compressor, vacuum pump, condenser) into a single metric [2].

Table 1. Upper and lower bounds applied to the operating variables θ for optimization of the S-TVSA process.

	t_{ads} (s)	t_{heat} (s)	t_{des} (s)	P_L (Pa)	T_H (K)	v_s (m/s)	v_{in} (m/s)
θ_L	1,000	500	1,000	500	363	0.0015	0.005
θ_U	15,000	1,500	40,000	5,000	373	0.005	0.01

Building on previous work, the purity constraint is represented as a quadratic loss penalty function with a drawdown coefficient c = 0.25 [2]. The decision variables that are optimized include the adsorption time (t_{ads}), the heating time (t_{heat}), the desorption time (t_{des}), the vacuum pressure (P_L), the steam temperature (T_H), the steam velocity (v_s), and the air feed velocity (v_{in}). The duration of the blowdown (t_{bd}) and the pressurization (t_{pres}) steps is 30s [2].

Industrial Clusters – Location case-study selection

To enable a cluster-specific assessment and the construction of a year-long dynamic performance profile under realistic ambient conditions, a representative site within the Scottish Industrial Cluster was selected. The Scottish Cluster is particularly relevant for large-scale DAC deployment due to its proximity to existing and planned CO_2 transport and storage infrastructure [10]. The area of Saint Fergus was chosen as a reference location, serving as a proxy for cluster-level deployment and enabling a detailed, high-resolution assessment of weather-driven performance variability. Year-long ambient meteorological data for 2024 (T_d, RH_d) were obtained on a daily resolution from the Buchan station, the nearest UK Met Office station to the selected reference site [11].

Effect of weather variability and temporal aggregation

Dynamic Performance Profiles

Previous DAC process-modelling studies have typically employed constant ambient conditions to evaluate system performance. Although computationally efficient, this assumption neglects the intrinsic temporal variability of atmospheric T and RH that directly define DAC inlet conditions and may therefore bias predicted KPIs. To systematically evaluate the impact of weather resolution on DAC performance, we develop a two-branch methodological framework combining optimization and forward simulations under different temporal representations of ambient conditions. By developing dynamic performance profiles, we can quantify how real operation at a fixed operating set of conditions (fixed operating strategy) responds to different temporal resolutions of weather.

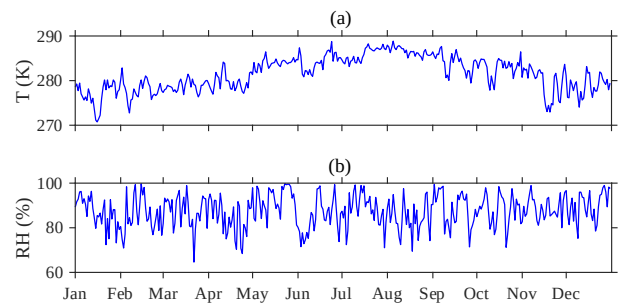

Figure 2. Daily ambient conditions at Buchan in 2024: (a) temperature T_d (K) and (b) relative humidity RH_d (%). Over the year, T_d ranged from 270.8-288.9 K and RH_d from 64.6-99.8%. Annual mean values were $\bar{T}_y = 281.6$ K and $\overline{RH}_y = 87.4\%$.

Furthermore, we can evaluate whether temporal aggregation of weather data yields values that are consistent with the fixed monthly and yearly inputs commonly assumed in DAC process simulations, and quantify any discrepancies introduced by this simplification. Daily weather data (T_d, RH_d) for the chosen location, with $d \in \{1, \dots, D\}$ where D are the number of days in 2024 (366), are obtained directly from the data source. The first step of the workflow involves running a multi-objective optimization under yearly averaged $(\bar{T}_y, \overline{RH}_y)$ conditions. From the resulting Pareto set, as shown in Figure 3, the maximum-productivity point $(\mathrm{Pr}^*_{y,max})$ is selected to define the fixed operating strategy for this workflow. Next, this fixed vector of decision variables, θ^*_y, that describe optimal operating conditions at $\mathrm{Pr}^*_{y,max}$ is propagated into forward simulations in the DAC model using the full sequence of daily weather inputs (T_d, RH_d). This procedure yields time-dependent daily trajectories of KPIs, enabling a direct correlation between ambient weather inputs and DAC performance. Daily weather inputs are averaged on a monthly and yearly basis as shown below:

$$\bar{T}_{m,i} = \frac{1}{D_{m,i}} \sum_{d=1}^{D_{m,i}} T_d \tag{2}$$

$$\overline{\mathrm{RH}}_{m,i} = \frac{1}{D_{m,i}} \sum_{d=1}^{D_{m,i}} RH_d \quad (3)$$

$$\bar{T}_y = \frac{1}{\mathrm{D}} \sum_{d=1}^{\mathrm{D}} T_d \quad (4)$$

$$\overline{\mathrm{RH}}_y = \frac{1}{\mathrm{D}} \sum_{d=1}^{\mathrm{D}} RH_d \quad (5)$$

$D_{m,i}$ denotes the number of days in month i.

Forward Simulations

To assess the impact of weather data temporal resolution on performance predictions, the daily KPI outputs are aggregated into monthly and yearly values.

$$\overline{KPI}_{m,i} = \frac{1}{D_{m,i}} \sum_{d=1}^{D_{m,i}} KPI_d(T_d, RH_d, \theta_y^*) \quad (6)$$

$$\overline{KPI}_y = \frac{1}{\mathrm{D}} \sum_{d=1}^{\mathrm{D}} KPI_d(T_d, RH_d, \theta_y^*) \quad (7)$$

The resulting averaged KPIs are then compared against KPIs obtained from simulations that use fixed monthly $(KPI_{m,i}(\bar{T}_{m,i}, \overline{RH}_{m,i}, \theta_y^*))$ and fixed yearly $(KPI_y^*(\bar{T}_y, \overline{RH}_y, \theta_y^*))$ ambient conditions by testing the following null hypotheses H_0.

$$H_0: \overline{KPI}_{m,i} - KPI_{m,i}(\bar{T}_{m,i}, \overline{RH}_{m,i}, \theta_y^*) = 0 \quad (8)$$

$$H_0: \overline{KPI}_y - KPI_y^*(\bar{T}_y, \overline{RH}_y, \theta_y^*) = 0 \quad (9)$$

In addition, we can also estimate yearly average KPIs using weighted fixed-month simulations:

$$\bar{\overline{KPI}}_y = \sum_{i=1}^{12} \frac{D_{m,i}}{D} KPI_{m,i}(\bar{T}_{m,i}, \overline{RH}_{m,i}, \theta_y^*) \quad (10)$$

This methodology allows us to (i) determine whether the temporal averaging implied in fixed-weather modelling is valid and (ii) to quantify the operational penalties that arise when a DAC unit optimized for nominal yearly conditions is exposed to real atmospheric variability.

Process Optimization

The second branch of the workflow will allow us to quantify the effect of temporal resolution in process improvements of DAC optimal performance. We solve the optimization problem using NSGA-II under two temporal representations of ambient conditions: yearly constant conditions as described above $(\bar{T}_y, \overline{RH}_y)$ and monthly-constant conditions $(\bar{T}_{m,i}, \overline{RH}_{m,i})$, yielding 12 different sets of decision variables $(\theta_{m,i}^*)$. For each optimization scenario, two representative Pareto solutions are extracted: the maximum Pr $(\mathrm{Pr}_{y,max}^*)$ and the minimum Weq points $(\mathrm{Weq}_{y,min}^*)$. For the monthly optimization case, annual mean KPIs are computed by averaging weighted monthly optima:

$$\bar{\overline{\mathrm{KPI}}}_y^* = \sum_{i=1}^{12} \frac{D_{m,i}}{D} KPI_{m,i}^*(\bar{T}_{m,i}, \overline{RH}_{m,i}, \theta_{m,i}^*) \quad (11)$$

This branch of the workflow tests whether increasing the temporal resolution of the optimization inputs leads to improved predicted annual-mean optimal performance relative to a single yearly-constant optimization. To this end, the following null hypothesis is tested:

$$H_0: KPI_y^*(\bar{T}_y, \overline{RH}_y, \theta_y^*) - \bar{\overline{\mathrm{KPI}}}_y^*(\bar{T}_{m,i}, \overline{RH}_{m,i}, \theta_{m,i}^*) = 0 \quad (12)$$

Furthermore, to quantify the month-specific value of re-optimization, we compare the performance obtained when operating each month using the annually optimal decision vector (θ_y^*) against the performance achieved under month-optimal settings $(\theta_{m,i}^*)$, both evaluated under the same month-averaged ambient conditions $(\bar{T}_{m,i}, \overline{RH}_{m,i})$. This test isolates the operational benefit attributable solely to adjusting operating parameters at a monthly timescale, relative to maintaining fixed yearly operation:

$$H_0: KPI_{m,i}(\bar{T}_{m,i}, \overline{RH}_{m,i}, \theta_y^*) - KPI^*{}_{m,i}(\bar{T}_{m,i}, \overline{RH}_{m,i}, \theta_{m,i}^*) = 0 \quad (13)$$

Aggregated on an annual scale, we also evaluate:

$$H_0: \bar{\overline{KPI}}_y - \bar{\overline{KPI}}_y^* = 0 \quad (14)$$

Finally, to evaluate whether gains identified under aggregated ambient conditions translate into improved annual operation under fully time-resolved weather, we test the following null hypothesis:

$$H_0: \overline{KPI}_y - \overline{KPI_y^*} = 0 \quad (15)$$

, where we define:

$$\overline{KPI_y^*} = \frac{1}{\mathrm{D}} \sum_{d=1}^{\mathrm{D}} KPI_d(T_d, RH_d, \theta_{m,i}^*) \quad (16)$$

RESULTS

Dynamic Performance Profiles

The DAC unit was optimized at nominal weather conditions that correspond to the yearly average. From the resulting Pareto Front shown in Figure 3, the combination of operating variables corresponding to the maximum productivity was chosen to create our fixed operating strategy.

The simulated DAC productivity exhibits intra-annual variability (Figure 4b), ranging from 5.6 to 6.2 kg/m^3/day, with lower values occurring under combinations of high RH and medium-to-high ambient T. Although modest in magnitude, a non-uniform productivity profile can propagate across the CO_2 removal value chain. At plant level, it can reduce capacity utilization and

can bias sizing when designs rely on annual-average performance. At cluster scale, time-varying CO_2 output may undermine steady-supply assumptions for downstream compression, transport, and storage infrastructure, lowering utilization of shared assets. Even without explicitly modeling the Carbon Capture & Storage (CCS) chain in this study, these results suggest that sustained productivity fluctuations can materially affect integration, cost estimates, and deployment robustness. Energy consumption (Figure 4c) is strongly affected, displaying peak-to-trough variations of up to 76% over the year - an effect that is obscured when steady-state or temporally averaged conditions are assumed. It is interesting to notice that high RH values cause a simultaneous increase in energy consumption and reduction in productivity, consistent with CO_2-H_2O co-adsorption: at very humid conditions, H_2O adsorption dominates and water-loaded sorption sites require additional regeneration, increasing energy requirements while reducing effective CO_2 throughput. Finally, purity shows a clearer intra-annual trend, decreasing during summer months (Figure 4a). This is consistent with the decreasing adsorption capacity and sorbent's reduced selectivity at higher temperatures, which gives more space for H_2O to compete by occupying and accumulating on the void spaces, yielding a mixture which is less rich in CO_2. Although the purity constraint of 95% was met for this combination of ambient conditions and marginal peak-to-through variability (1%) was observed, variable product purity can entail the risk of increasing downstream processing costs to maintain the stream within spec, with significant impacts on the downstream CCS chain.

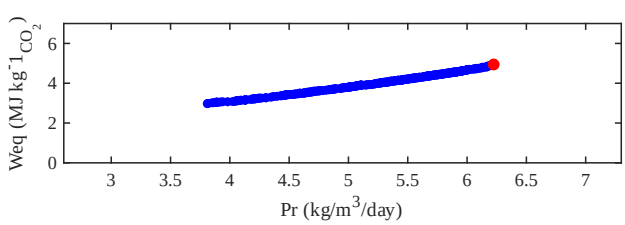

Figure 3. Pareto Front resulting from the multi-objective optimization problem of maximizing productivity (Pr) and minimizing equivalent work (Weq) using NSGA-II.

Temporal Weather Averaging

Figure 5 compares normalized DAC performance indicators obtained using different temporal representations of ambient conditions while maintaining a fixed operating strategy (θ_y^*). KPIs calculated using yearly average conditions are included as a reference (dashed lines) to reflect the common simplifying assumption adopted in much of the existing DAC modeling literature. The results demonstrate that the use of temporally averaged weather inputs leads to systematically biased performance predictions. When monthly- (green bins) or yearly-averaged (dashed line) T and RH values are employed, KPIs can be either over- or under-estimated relative to simulations driven by daily resolved meteorological data (purple bins). System productivity is consistently overestimated, while energy requirements are significantly underestimated when constant yearly ambient conditions are assumed.

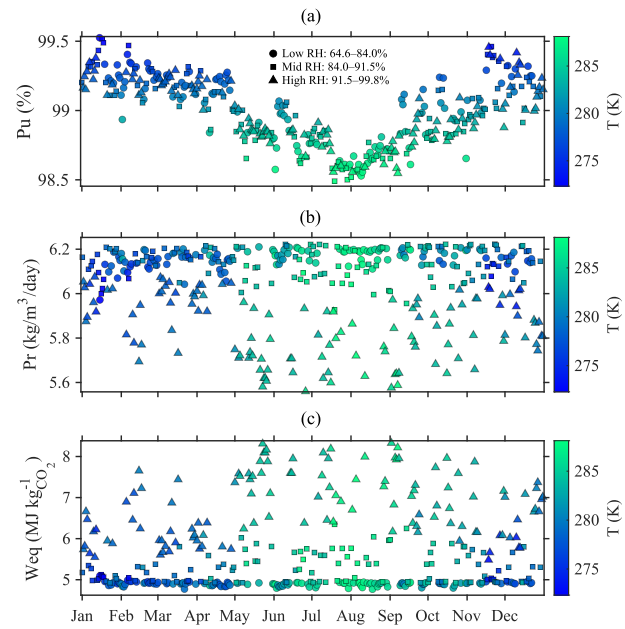

Figure 4. Dynamic annual profiles of (a) purity (Pu), (b) productivity (Pr), and (c) equivalent work (Weq) for the optimized DAC unit operating at St. Fergus.

May corresponds to the maximum deviation from the yearly nominal operating conditions, with productivity overpredicted by up to 5% and energy demand underestimated by approximately 31% when compared to the daily-resolved cases. Despite these discrepancies, CO_2 purity is shown to be comparatively insensitive to temporal aggregation, with deviations remaining below 1% across all cases. This indicates that purity constraints are robust to weather averaging for the operating regime considered. Other critical KPIs exhibit much stronger nonlinear responses to ambient variability. Electrical energy consumption is underestimated by up to 195% when yearly fixed temporal representations are used underscoring the extent to which averaging suppresses high-energy operating periods. On an annual scale (purple bins), the same trends are observed compared to the yearly-averaged simulations (dashed line). This discrepancy highlights that the annual mean of a nonlinear KPI does not correspond to the indicator evaluated at annual-mean ambient conditions. Hence, we show that:

$$\overline{KPI}_y \neq KPI_y^*(\bar{T}_y, \overline{RH}_y, \theta_y^*)$$

A similar effect is observed at the monthly level. When KPIs obtained from daily-resolved simulations (light purple bins) are averaged over each month and compared to simulations performed using fixed monthly-average ambient conditions (green bins), systematic deviations persist. Specifically, while purity remains largely unaffected, productivity is overestimated by up to 2.5% and energy consumption is underestimated by up to 11% depending on the month. Hence, we conclude that:

$$\overline{KPI}_{m,i} \neq KPI_{m,i}(\bar{T}_{m,i}, \overline{RH}_{m,i}, \theta_y^*)$$

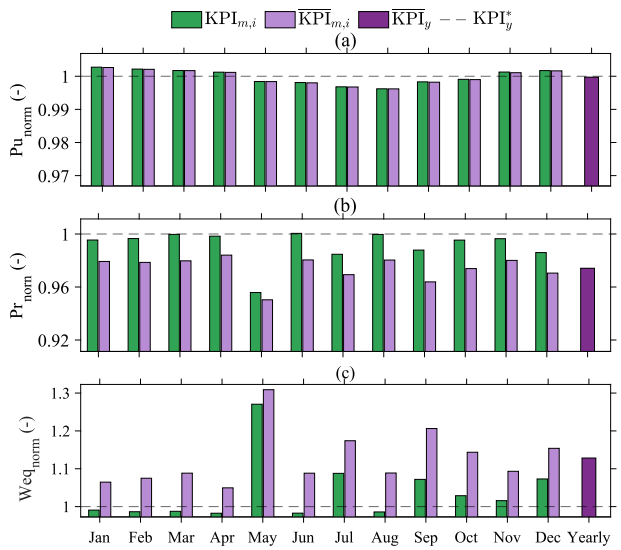

Figure 5. Comparison of normalized (a) purity (Pu_{norm}), (b) productivity (Pr_{norm}), and (c) equivalent work (Weq_{norm}) obtained at different temporal resolutions. All results are obtained using the same operating variables θ_y^* and are normalized by the nominal case that produces KPI_y^*.

These deviations arise because short-term excursions to unfavorable temperature and humidity conditions - particularly those associated with high energy penalties - are smoothed out when average inputs are used. These findings demonstrate that temporal aggregation masks the true operational variability of DAC systems by suppressing extreme but impactful weather conditions. While low-resolution weather data may be sufficient for predicting product quality, it fails to accurately capture energy requirements and the system's productivity. This has important implications for techno-economic assessments, utility sizing, and system integration studies, where underestimating peak and cumulative energy demand may lead to overly optimistic cost projections and misinformed process design.

Optimization at temporal resolutions

Building on the forward-simulation analysis, we move to the optimization framework which enables detangling (i) the benefit of re-optimizing operating conditions at a monthly timescale, (ii) the month- and yearly specific improvement achievable relative to fixed annual operation, and (iii) the extent to which apparent gains under aggregated weather inputs translate into improved performance under fully time-resolved (dynamic) ambient conditions. These hypotheses provide a structured basis to determine whether increasing temporal resolution yields meaningful operational improvements, or whether it primarily increases computational burden without material gains in annual performance. As seen in Table 2, monthly re-optimization ($\overline{\overline{\text{KPI}}}_y^*$) does not yield consistent annual-mean performance gains relative to a single yearly-constant optimization (KPI_y^*). At $Pr_{y,max}^*$, monthly optimization produces negligible change in productivity but increases equivalent work from 4.95 to 5.16 (+4.35%), indicating a deterioration in energy performance. In contrast, at $Weq_{y,min}^*$, monthly optimization substantially increases productivity (+10.13%) but at the expense of increased energy consumption (+15.09%). Overall, these results indicate that increasing temporal resolution in the optimization inputs shifts the optimal trade-off between productivity and energy consumption rather than delivering uniform improvements across the Pareto front. Hence, we conclude that:

$$KPI_y^*(\bar{T}_y, \overline{RH}_y, \theta_y^*) \neq \overline{\overline{\text{KPI}}}_y^*(\bar{T}_{m,i}, \overline{RH}_{m,i}, \theta_{m,i}^*)$$

Table 2. Annual-mean KPI comparison for yearly-constant and monthly-constant optimization (max-Pr, min-W_{eq}). The relative KPI changes are calculated with respect to the yearly optimization strategy that produces KPI_y^* at each operating point. $\Delta\text{KPI} > 0$ signifies an increase in the measured KPI.

Operating point	Strategy	Pr	Weq	$\Delta\,\text{Pr}\,\%$	$\Delta Weq\,\%$
$Pr_{y,max}^*$	KPI_y^*	6.22	4.95	-	-
	$\overline{\overline{\text{KPI}}}_y^*$	6.20	5.16	-0.37	4.35
$Weq_{y,min}^*$	KPI_y^*	3.81	2.98	-	-
	$\overline{\overline{\text{KPI}}}_y^*$	4.20	3.43	10.13	15.09

Furthermore, as shown in Figure 6, by investigating the month-by-month value of re-optimization at $Pr_{y,max}^*$ we observe that monthly re-optimization yields higher productivity in most months, with $\%\Delta Pr \geq 0$ in 9 out of 12 months. The largest productivity increase occurs in January (+4.75%), while the most pronounced productivity penalty occurs in November (-6.98%), indicating that re-optimization does not consistently preserve the productivity-extreme solution across the full seasonal

envelope of monthly mean conditions. The energy response at $Pr^*_{y,max}$ is mixed and season dependent. The equivalent work decreases (i.e., $\%\Delta W_{eq} \leq 0$) in 8 out of 12 months, reaching a maximum reduction of -11.68% in April and -7.56% in May, but increases in the remaining months, with the largest penalty occurring in October (+7.79%). Consequently, monthly re-optimization at the productivity-driven operating point does not yield uniform co-improvement but instead induces month-specific trade-offs where productivity gains may be accompanied by higher energy consumption (e.g., January, October, December), and energy reductions may coincide with reduced productivity (notably November). At the minimum-energy operating point, monthly re-optimization produces a consistent productivity improvement in every month, with $\%\Delta \mathrm{Pr} > 0$ for 12 out of 12 months. The magnitude of the productivity uplift is substantial, ranging from +1.48% (March) up to +23.02% (November), demonstrating that the fixed yearly operating strategy is systematically conservative in terms of throughput when applied under month-specific ambient conditions.

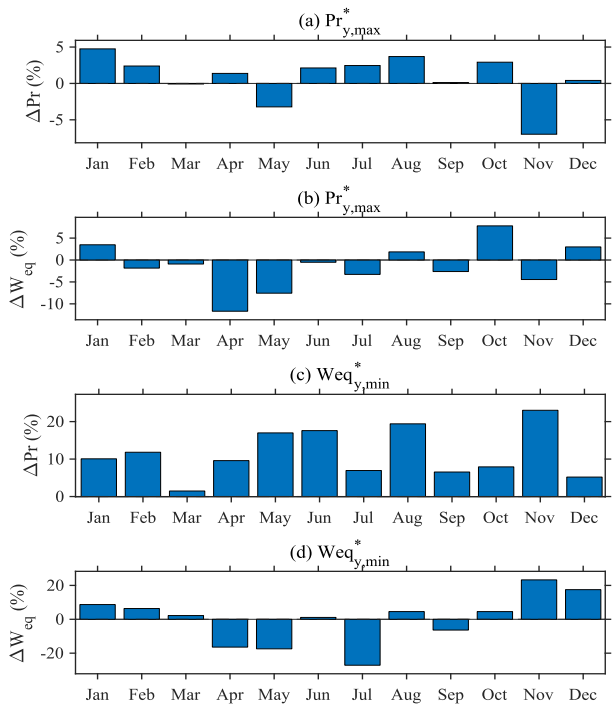

Figure 6. Monthly percentage change in productivity and equivalent work when applying monthly re-optimization relative to fixed yearly-optimized operation evaluated at $Pr^*_{y,\max}$ (a/b) and at $W^*_{\mathrm{eq},y,\min}$ (c/d).

In contrast, the energy effect remains highly seasonal. Equivalent work is reduced (i.e., $\%\Delta W_{eq} \leq 0$) in only 4 out of 12 months - specifically April (-16.42%), May (-17.44%), July (-27.12%), and September (-6.35%) - but increases in the remaining months, with pronounced penalties in November (+23.27%) and December (+17.50%). This indicates that month-specific re-optimization does not consistently minimise energy month-by-month but rather relocates the operating point along the Pareto surface in a way that frequently favours productivity increases, occasionally at the expense of higher energy demand. When averaged on an annual scale, at $Weq^*_{y,min}$, montlhy re-optimization provides a clear annual productivity increase (~11%) with a modest reduction in Weq (~2%). A similar decrease in Weq is observed at $Pr^*_{y,max}$ coupled with a marginal increase in Pr (< 1%). Hence, we conclude that:

$$KPI_{m,i}(\bar{T}_{m,i}, \overline{RH}_{m,i}, \theta^*_y) \neq KPI^*{}_{m,i}(\bar{T}_{m,i}, \overline{RH}_{m,i}, \theta^*_{m,i})$$

$$\overline{KPI}_y \neq \overline{\mathrm{KPI}}^*_y$$

The above analysis serves as a tool to depict that finer data can improve the predicted KPIs and hence show that literature-assumptions of constant weather prevents the system from reaching its maximum potential. However, these gains are achieved under the assumption of monthly fixed weather inputs. The real operational gains should be evaluated when daily weather data is propagated along with the optimal values of operating variables. Under time-resolved weather propagation, monthly scheduling provides no meaningful benefit at $Pr^*_{y,max}$, yielding negligible productivity change (+0.03%) while increasing equivalent work (+1.72%). In contrast, at $W^*_{\mathrm{eq},y,\min}$, monthly scheduling delivers a substantial productivity gain (+10.24%) but no corresponding energy reduction, with equivalent work remaining effectively unchanged (+0.18%). Overall, the 'real gains' analysis shows that improvements inferred under aggregated ambient conditions do not necessarily translate into energy savings under intra-month variability, while productivity gains persist primarily in the energy-minimization regime. Hence, we conclude that:

$$\overline{KPI}_y \neq \overline{\mathrm{KPI}}^*_y$$

CONCLUSION

To conclude, this work leverages on a previously developed high-fidelity, mechanistic, model to assess the impact of local weather variability and temporal aggregation on model predictions and real operational gains for an S-TVSA DAC unit. We show that DAC performance responds non-linearly to temperature and relative humidity. Consequently, KPIs evaluated at monthly or yearly average conditions do not reproduce true annual performance under time-resolved operation. Productivity is consistently overestimated, while electrical energy demand and equivalent work are substantially underestimated when yearly average weather inputs are used, with important implications for techno-economic

assessment, plant sizing, and system integration.

The study further shows that increasing the temporal resolution of the optimization inputs from yearly-mean to monthly-mean ambient conditions does not deliver uniform improvements in annual performance but instead repositions the optimal operating strategy along the productivity-energy trade-off surface. Month-specific optimizations can yield localized performance peaks (e.g., productivity uplift in selected months or reduced equivalent work) that can survive aggregation and deliver an improved performance on an annual scale. Importantly, operating at the energy minimization regime can offer substantial productivity gains. However, this work highlights that the value of these strategies must be assessed through dynamic evaluation: when day-to-day weather variability is propagated through the process model, the apparent benefits inferred under aggregated conditions are attenuated, and meaningful operational improvements remain confined to specific operating regimes. Collectively, these findings indicate that weather should be treated as a primary operational input for DAC design and scheduling - rather than a secondary disturbance to be averaged out if robust, region-specific deployment near industrial clusters is to be achieved. Although demonstrated here for the UK, the proposed workflow is directly transferable to other regions through substitution of site-specific meteorological datasets and deployment constraints, enabling consistent cross-regional benchmarking of DAC performance and operating strategies.

ACKNOWLEDGEMENTS

This work was funded by Graham Taylor, an alumnus of Imperial College London who left a generous gift in his will to fund research into energy policy.

AUTHOR IDENTIFIERS

Author ORCIDs:
Eleni Chalasti: 0009-0002-5835-8331
Gbemi Oluleye: 0000-0001-7753-0315
Maria M. Papathanasiou: 0000-0002-8886-0624
Ronny Pini: 0000-0002-9443-3573

REFERENCES

1. C. Science and C. Change Committee, “Assessing the Feasibility for Large-scale DACCS Deployment in the UK Climate Change Committee Assessing the Feasibility for Large-scale DACCS Deployment in the UK, ” Feb. 2025. Accessed: Oct. 02, 2025. [Online]. Available: https://www.theccc.org.uk/wp-content/uploads/2025/02/Assessing-the-Feasibility-for-Large-scale-DACCS-Deployment-in-the-UK-2.pdf
2. Ward A, Papathanasiou MM, Pini R. The impact of design and operational parameters on the optimal performance of direct air capture units using solid sorbents. Adsorption 30:1829-1848 (2024). https://doi.org/10.1007/s10450-024-00526-y
3. Cai X, Coletti MA, Sholl DS, Allen-Dumas MR. Assessing impacts of atmospheric conditions on efficiency and siting of large-scale direct air capture facilities. JACS Au 4:1883-1891 (2024). https://doi.org/10.1021/jacsau.4c00082
4. Sendi M, Bui M, Mac Dowell N, Fennell P. Geospatial analysis of regional climate impacts to accelerate cost-efficient direct air capture deployment. One Earth 5:1153-1164 (2022). https://doi.org/10.1016/j.oneear.2022.09.003
5. A. Luukkonen and J. Elfving, “Optimizing direct air capture: evaluating the impact of process parameters on productivity, energy requirement, and cost, ” 2024. [Online]. Available: https://ssrn.com/abstract=5014185
6. Jung H, Kim K, Jeong J, Jamal A, Koh DY, Lee JH. Exploring the impact of hourly variability of air condition on the efficiency of direct air capture. Chemical Engineering Journal 508:160840 (2025). https://doi.org/10.1016/j.cej.2025.160840
7. Balasubramaniam BM, Thierry PT, Lethier S, Pugnet V, Llewellyn P, Rajendran A. Process-performance of solid sorbents for direct air capture (DAC) of CO2 in optimized temperature-vacuum swing adsorption (TVSA) cycles. Chemical Engineering Journal 485:149568 (2024). https://doi.org/10.1016/j.cej.2024.149568
8. Wiegner JF, Grimm A, Weimann L, Gazzani M. Optimal design and operation of solid sorbent direct air capture processes at varying ambient conditions. Ind. Eng. Chem. Res. 61:12649-12667 (2022). https://doi.org/10.1021/acs.iecr.2c00681
9. Stampi-Bombelli V, van der Spek M, Mazzotti M. Analysis of direct capture of $${\hbox {CO}}_{2}$$ from ambient air via steam-assisted temperature–vacuum swing adsorption. Adsorption 26:1183-1197 (2020). https://doi.org/10.1007/s10450-020-00249-w
10. S. Cluster, “Ready to Deliver Industrial Decarbonisation.” [Online]. Available: www.theccc.org.uk/publication/net-zero-the-uks-contribution-to-stopping-global-warming
11. UK Met Office, “ MIDAS Open: UK daily weather observation data, v202507, ” CEDA Archive. Accessed: Jan. 08, 2026. [Online]. Available: https://data.ceda.ac.uk/badc/ukmo-midas-open/data/uk-daily-weather-obs/dataset-version-202507/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Dynamic Optimization of an Adsorption Heat Storage to satisfy the Heat Demand of a House

Alix Untrau[a]*, Lorenz T. Biegler[b], and Sabine Sochard[a]

[a] Universite de Pau et des Pays de l'Adour, E2S UPPA, LaTEP, Pau, France
[b] Carnegie Mellon University, Chemical Engineering Department, Pittsburgh, Pennsylvania, USA
* Corresponding Author: alix.untrau@univ-pau.fr.

ABSTRACT

This study presents the modeling and operation optimization of an adsorption heat storage to improve the supply of renewable heat to a house. The system configuration is an open system with water being carried by an air flow and adsorbed on zeolite 13X beads in a packed bed. A numerical model is developed based on mass and energy balances, using a Langmuir adsorption isotherm and a Linear Driving Force (LDF) mass transfer equation. The model is implemented in Pyomo and solved with the NLP solver IPOPT. A sensitivity analysis on the discretization parameters is performed to choose a good compromise between accuracy and computational time. The chosen model is then validated against experimental data from the literature, with a mean absolute percentage error less than 5%. The dynamic optimization of the operation of the system to satisfy a heat demand is then performed. The trajectory for the inlet fluid velocity is optimized in several heat demand scenarios. The results show that this numerical framework is able to follow realistic heat demands for a house even in the case of large peaks corresponding to domestic hot water production, with quadratic errors between the demand and supplied power less than 1%.

Keywords: Adsorption, Energy Storage, Modelling and Simulations, Optimization, Pyomo

INTRODUCTION

Heat production through renewable energy such as solar energy is a key driver of the energy transition. Due to the intermittency of solar irradiation, an energy storage solution is required to achieve a continuous heat delivery. Thermal energy storage technologies are developed based on three different principles: sensible, latent and thermochemical heat storage. The later includes adsorption heat storage based on the exothermic adsorption of a molecule on a porous structure. In this case, the heat is stored as an adsorption potential allowing long storage duration without heat losses. Moreover, this type of storage has a high energy density, up to 10 times larger than a sensible storage and 2 to 3 times larger than a latent storage [1]. Adding salts to the porous structure increases the thermal effects due to the exothermicity of the hydration reaction. However, such thermochemical systems suffer from poor cycling capacity [1] and are thus excluded from this study.

Adsorption heat storage is still under development and research efforts are required in several areas: adsorbent materials, adsorbate choice, storage structure, process optimization [2]. Some prototypes, such as in [3], have been built to assess the adsorption storage feasibility. However, it was not tested in realistic house heating demand scenarios and the operating conditions were not adjusted accordingly. The energy density of experimental prototypes is often much lower than the theoretical energy density of the adsorbent/adsorbate couple [4], highlighting the need for research on system design and operation. In particular, Lefebvre et al. emphasized the need to adjust the operating parameters in order to reach high energy density values [1].

Optimizing a system's design or operation requires a reliable and accurate model that faithfully represents its behavior. Numerous studies have developed a numerical model for an adsorption heat storage from detailed CFD models to simplified 1D models, such as in [5]. However, the models built might not be compatible with optimization studies, where a good trade-off between accuracy and computational time is essential.

Most optimization studies on adsorption heat storage focus on the design of the system ([6], [7]). Zamengo

et al., used a multi-objective methodology to assist the design of a packed bed thermochemical storage [6]. Optimal values for the adsorbent porosity and thickness are determined while maximizing heat storage density, storage rate and exergy efficiency. The conclusion suggests that modular thin beds are more efficient. Anilkumar et al. developed a 3D dynamic mathematical model using analytical equations and applied a multi-objective algorithm to design an efficient adsorption system [7]. Dimensions for the reactor tubes and the fins added to enhance heat transfer are obtained and realistic systems are simulated based on the recommendations. Those studies did not consider the optimization of the operation of the system.

Wang et al. studied the configuration planning of an adsorption heat storage for different applications in a house [8]. Charging and discharging steps are considered and decision variables include the inlet parameters: inlet air temperature, flow rate and humidity. However, constant values are obtained for each heating scenario.

In this study, a physical 1D numerical model for an adsorption heat storage is developed and validated with experimental data from [3]. The dynamic model is implemented in an optimization framework. Trajectories for the inlet conditions (air velocity) are determined in order to satisfy a heating demand from a house. The dynamic behavior of the system is further analyzed.

The remaining parts of the paper are organized as follows: first, the adsorption heat storage is described. Then, the numerical model is presented along with a sensitivity analysis on the discretization parameters and the validation of the model. After that, the optimization problem is formulated and the results from case studies are explained. Finally, a discussion presents the limitations and future work before a conclusion on this study.

SYSTEM DESCRIPTION

This study considers the supply of renewable heat from solar collectors to a house. An adsorption storage is added to tackle the intermittency of solar energy and reduce the need of a back-up heat production. The storage is the main focus of this work. The system considered is an open system, using air as a carrier and water as the adsorbate. This choice keeps the complexity of the system and the investment cost lower than other configurations. The adsorbent is zeolite 13X in a packed bed.

The two modes of operation are schematized in Figure 1. When excess solar heat is available, it can be used to desorb water from the zeolite, and thus charge the heat storage. This step requires high temperatures, in the range of [120;180]°C which can be achieved with suitable solar collector technologies such as evacuated tube collectors. The choice of the collectors and the detailed analysis of the desorption step is outside the scope of this study. Once the storage is charged, i.e. all the water has been removed, it can be kept in this state for extended periods without any heat losses. When heat is needed in the house but cannot be directly supplied by the solar collectors, humid air is sent to the adsorption storage. This air can be the interior air of the house or the outside air if its humidity and temperature allow for a sufficient temperature lift. The water contained in the air is adsorbed on the zeolite, releasing heat which is collected and carried by the dried air flow. The heat can then be used for space heating or domestic hot water production in the house.

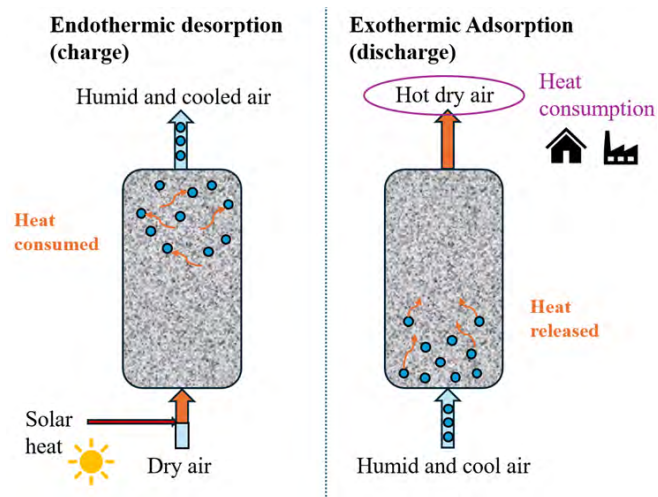

Figure 1. Schematic of the adsorption heat storage utilization in a house

NUMERICAL MODEL

Model equations

The model of the adsorption heat storage is a 1D model considering only spatial variations along the height of the packed bed. Since the system is open, a constant pressure equal to the atmospheric pressure is considered. There are 7 variables in the model, described in the nomenclature, each depending on the spatial coordinate z and time t, which are omitted for improved readability: $\rho_w, \rho_{da}, \bar{q_w}, q_{eq}, \rho, T, u$.

The model is based on mass and energy balances:

- Mass balance on the dry air, assuming a constant dry air flow

$$\frac{\partial(u\rho_{da})}{\partial z} = 0 \quad (1)$$

- Mass balance on water, neglecting axial dispersion

$$\frac{\partial\rho_w}{\partial t} = -\frac{\partial(u\rho_w)}{\partial z} - \frac{\rho_b}{\varepsilon_b}\frac{\partial\overline{q_w}}{\partial t} \quad (2)$$

- Mixture composition

$$\rho = \rho_w + \rho_{da} \quad (3)$$

- Energy balance on the whole system, assuming a

homogeneous temperature between the solid and the fluid

$$\left[\rho C_{v_f}\frac{\varepsilon_b}{1-\varepsilon_b}+\rho_s(1-\varepsilon_s)C_{p_s}+\rho_s(1-\varepsilon_s)\bar{q_w}C_{p_w}\right]\frac{\partial T}{\partial t}=C_{p_f}\frac{\varepsilon_b}{1-\varepsilon_b}\frac{\partial(u\rho T)}{\partial z}-\frac{\Delta H_{ads}\rho_s(1-\varepsilon_s)}{M_w}\frac{\partial\bar{q_w}}{\partial t} \quad (4)$$

- Langmuir isotherm

$$q_{eq}=\frac{Kq_mP_w}{1+KP_w} \quad (5)$$

$$K=K_0\exp\left(-\frac{\Delta H_{ads}}{RT}\right) \quad (6)$$

-Linear Driving Force (LDF)

$$\frac{\partial\bar{q_w}}{\partial t}=K_{LDF}\left(q_{eq}-\bar{q_w}\right) \quad (7)$$

-Ideal gas law for the fluid phase

$$\rho=\frac{PM}{RT} \quad (8)$$

The model uses a simple adsorption isotherm, Langmuir, with parameters values as follows: $K_0=4.06\mathrm{e}-12\ Pa^{-1}$ and $q_m=0.228\ kg_w/kg_s$. The LDF model is used to describe the mass transfer in the adsorbent. The lumped parameter K_{LDF} takes into account the external diffusion and internal diffusion in the particle and adsorption on the zeolite crystal. A constant K_{LDF} value of $0.0025\ s^{-1}$ was chosen based on the average value computed with the equations detailed in [5].

Initial conditions for the temperature, water concentration in fluid and adsorbed water concentration along the height of the column are required. Inlet conditions for the fluid temperature, water concentration and velocity are also fixed.

Resolution method

The model was implemented in the Pyomo framework [9]. Pyomo.dae was used to discretize the partial differential equations. The finite difference method was chosen for both time and space discretization, using a fully implicit Backward Euler scheme (first order, constant step). The discretized equations system forms a NonLinear Programming (NLP) problem solved with the IPOPT solver [10]. For a simulation, the number of variables matches the number of constraints, and a dummy objective is added for the resolution.

The model comprises a large number of variables and nonlinear equations and has to represent a steep adsorption front travelling through the storage tank. This makes the problem stiff and difficult to solve. An initialisation strategy was therefore developed as shown in Figure 2. The model was first solved one time step at a time, reducing drastically the complexity of the problem. The results at the end of a time step are used as initial conditions for the next, replacing the implicit scheme for time discretization by an explicit formulation. The results for every time steps are then used together to initialize the simulation over the complete time horizon. The results from the complete simulation are then used to initialize the dynamic optimization, which will be described later in the article.

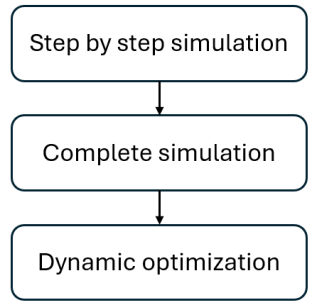

Figure 2. Initialisation strategy

Sensitivity on discretization

A sensitivity analysis on the discretization both in time and space was performed, considering zero degree of freedom, i.e. in a simulation configuration. The inlet conditions are fixed to constant values: $u_{in}=0.3\ m/s, T_{in}=293\ K, \rho_{w_{in}}=0.012\ kg_w/m_f^3$ from [3]. Initially, the storage is at 298K and does not contain water.

Figure 3 presents the temperature profiles obtained at the outlet of the column during adsorption simulations with three different time steps: 50s, 100s and 200s and a fixed spatial discretization of 20 points. The temperature quickly rises due to the exothermic adsorption and stays at a high temperature during several hours. Once the column is saturated and adsorption can no longer take place, the temperature decreases back to the inlet temperature. The profiles obtained with the three discretization time steps are following similar trends, reaching the same adsorption temperature, with only small differences during the transitions. A finer discretization allows to represent steeper transitions at the beginning and at the end of the adsorption phenomenon. Given the small differences between the results, a time step of 100s was chosen, as a good compromise between accuracy and computational time due to the number of variables.

Figure 4 presents the results of adsorption simulations with various numbers of spatial discretization points and a time step of 100s. There are some discrepancies between the profiles. First, the adsorption temperature reached varies by 2°C between 10 and 100 discretization points. Besides, the steepness of the temperature profiles at the beginning and at the end of the adsorption phase is reduced when the number of discretization points is reduced. This can lead to an inaccurate estimation of the adsorption duration. In particular, 10 discretization points does not allow to represent the curvature at the end of the adsorption phase, while all the other plots are superimposed in this zone of the graph.

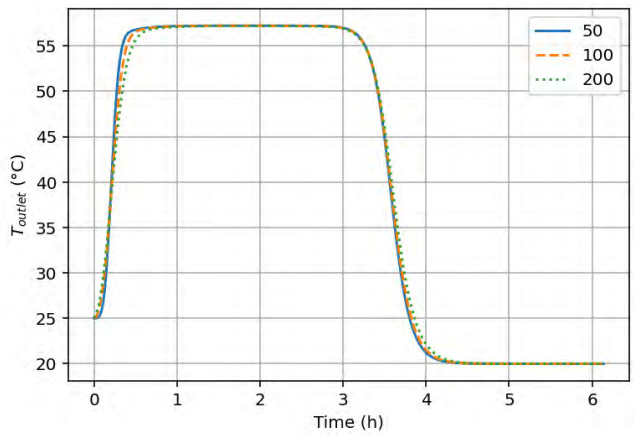

Figure 3. Sensitivity analysis on the time step duration

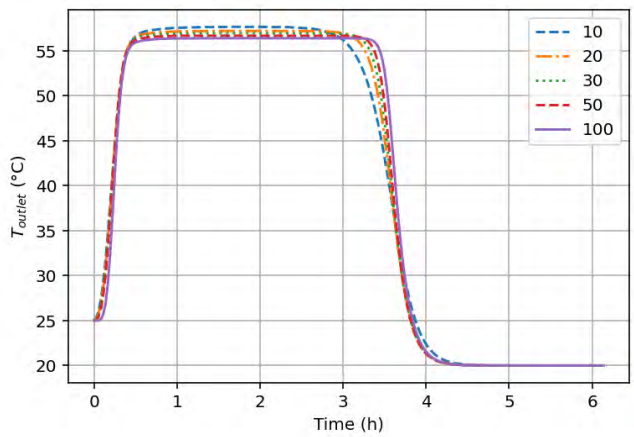

Figure 4. Sensitivity analysis on the number of spatial discretization points

Increasing the number of discretization points thus improves the accuracy of the model resolution but at the cost of a larger problem which will be more difficult to use in optimization studies. Table 1 presents the number of variables for each discretization tested. Since most variables depends on the spatial coordinate z, the number of variables increases greatly with the number of discretization points. In order to keep the complexity of the problem as low as possible for optimization while keeping a good accuracy, 20 points were chosen.

Table 1: Number of variables in the problem for each spatial discretization

Number of discretization points	Number of variables
10	33267
20	64317
30	95367
50	157467
100	312717

Validation

The model with the chosen discretization of 20 points in the spatial dimension and a time step of 100s was validated against experimental data from the work of Tatsidjodoung et al. [3]. The configuration from their prototype was reproduced in the model for all design and operation parameters, as well as initial and inlet conditions. The cylindrical bed considered has a volume of 144L for a height of 20cm. In the experimental bench, air diffusers ensure a uniform distribution of the flow over the entire section of the bed, justifying the use of a 1D model despite the large diameter of 72cm.

The experiment chosen for the validation is their "test 12" where an adsorption step is performed for two adsorption beds in series. The results at the outlet of the first bed (M1 in [3]) are considered for the validation because of the availability of the corresponding data.

The temperature profiles at the outlet of the bed are plotted in Figure 5. The shapes of the profiles are similar, even in the transition regions when temperature stabilizes at the beginning and the end of the adsorption phase. This ensures that the model predicts an accurate adsorption duration. The region when the temperature starts to decrease, after about 3 hours of heat supply, is not well reproduced by the model. There seems to be a break in the experimental curve at 3.5 hours without physical justification. Therefore, this part of the curve cannot invalidate the numerical model. The overall mean absolute percentage error between the two curves is 4.46% with an adjusted value of ΔH_{ads} equal to -56398 J/mol in the model, which lies in the range of experimental measurements of [45-75] kJ/mol [11]. The error below 5% ensures the validation of the model.

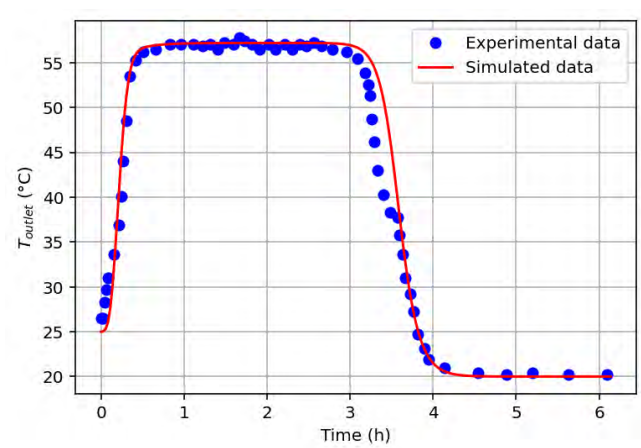

Figure 5. Comparison of the experimental [3] and simulated outlet temperature profiles

OPTIMIZATION

Formulation of the optimization problem

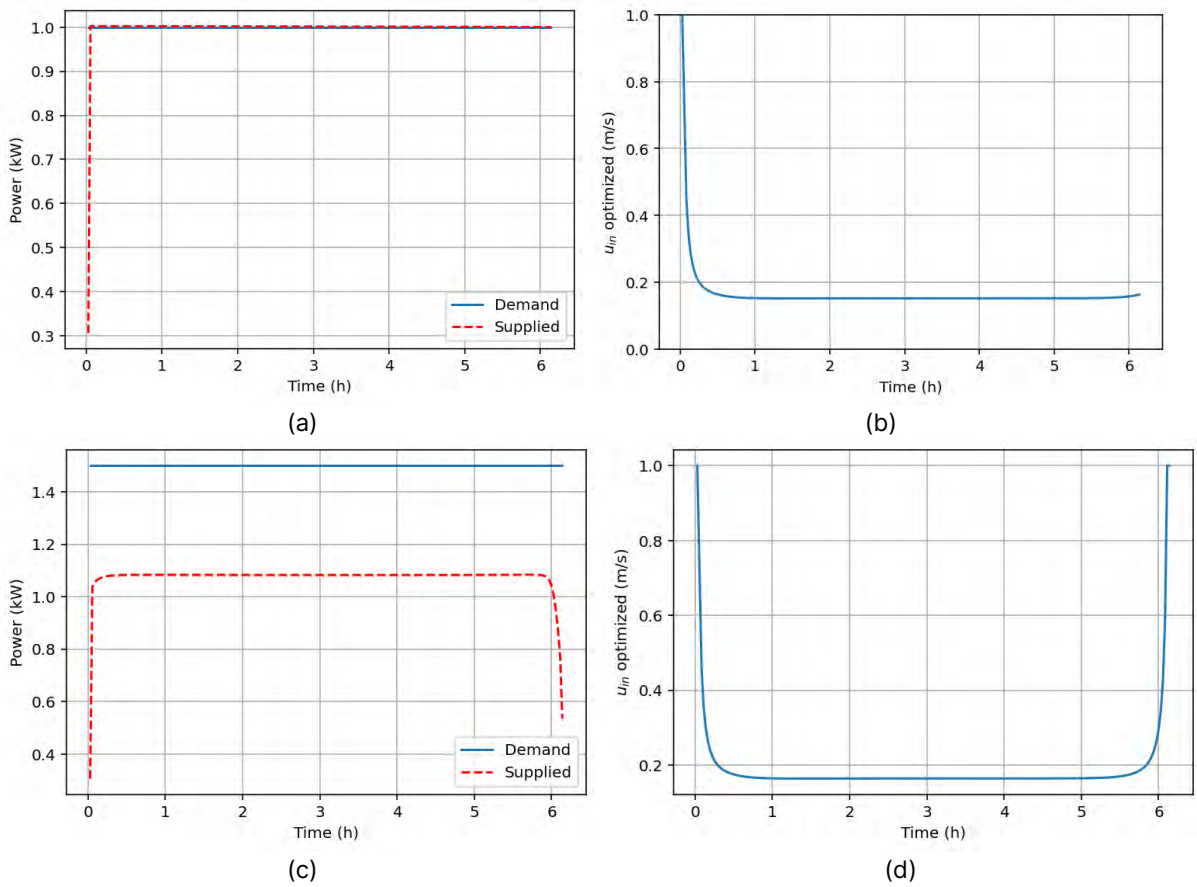

Figure 6: Various power demand scenarios and the obtained power supplied and the optimal velocity profiles – constant heat demand

The validated numerical model described previously is then used to perform dynamic optimization of the operation of the heat storage. The design of the system is fixed. The dimensions used are the same as in the validation part, taken from the experimental bench in [3].
The objective function is to minimize the squared difference between the heat power demand and the heat power supplied by the storage at each time step (quadratic error):

$$\min_{u_{in}(t_i)} \sum_{t_0}^{t_f} \left(\dot{Q}_{demand}(t_i) - \dot{Q}_{supplied}(t_i) \right)^2 \quad (8)$$

The optimization variables are the inlet velocity $u_{in}(t_i)$ at each time step. The velocity is bounded in the interval [0;1] m/s, to ensure that it is lower than the fluidization velocity. The other inlet conditions, namely the temperature and water concentration, remain fixed to the previously mentioned values because in an open system there are generally imposed by the ambient conditions. The only controllable variable is thus the velocity, through the manipulation of the fan.

Results

To assess the capability of the storage system to satisfy the heat demand of a house, several demand scenarios have been created. These scenarios are realistic heat demands of a room. The heat demand of a whole house is larger and would require several modules of the studied adsorption heat storage. To simplify, only one module is studied here.

Figure 6 shows the results for two distinct scenarios considering a constant heat demand, corresponding to the heating during the night or the middle of the day, when no big change occurs in the demand. Figure 6 (a) shows a heat demand of 1kW, that is almost perfectly matched by the heat supplied, with an objective value (sum of squared errors) of 0.91. The only discrepancy is at the beginning, because the adsorption heat delivery at the outlet of the column is not instantaneous. The corresponding optimal inlet velocity profile is shown in Figure 6 (b). At the initial time, the velocity is as high as possible to start the heat delivery as soon as possible and reduce the difference between the heat demand and the heat supply in the first minutes. Once the heat delivery has started, the velocity goes quickly to its constant value ensuring the required power.

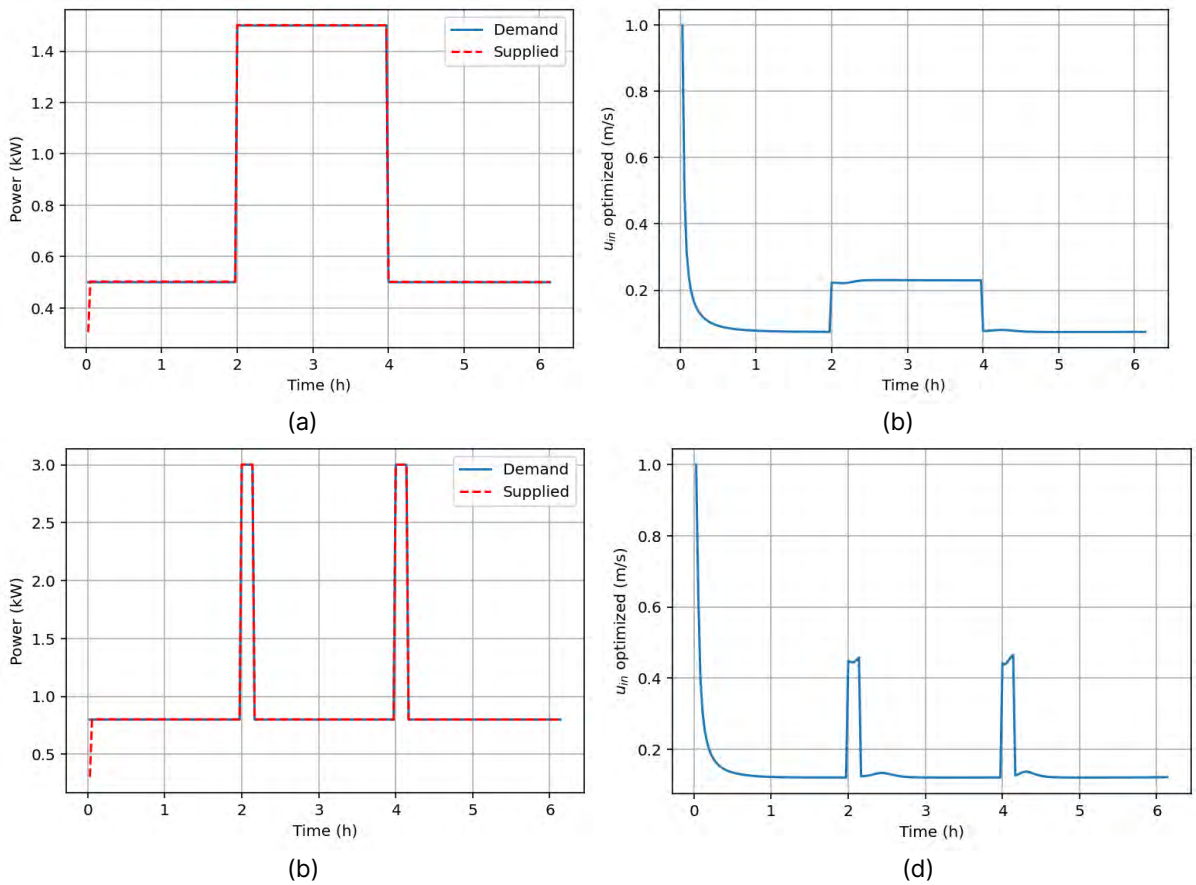

Figure 7: Various power demand scenarios and the obtained power supplied and the optimal velocity profiles – variable heat demand

In Figure 6 (c), a similar scenario with a greater heat demand of 1.5kW is shown. In this case, the adsorption storage is not able to satisfy the heat demand, the heat supplied is lower than the demand and the objective is still high with a value of 42.8. Indeed, the capacity of the storage is exceeded, the zeolite cannot adsorb more water and saturates. The optimal solution is here to use the adsorption storage for preheating over the whole time period. A back-up heater will then be necessary to satisfy the heat demand. The optimal inlet velocity for this scenario is shown in Figure 6 (d). As previously, the velocity is high at the very beginning, to start the heat delivery as quickly as possible. It then stabilizes to a value ensuring a stable heat delivery over the whole time horizon. Towards the end of the time period, the storage tank is nearly saturated and not much heat can still be extracted. The velocity increases again to try to match the heat demand, but the supplied power still decreases. This case shows the capacity of the methodology to adjust the inlet velocity to satisfy a heat demand. When the heat demand is too high and the system will not be able to satisfy it over the whole time horizon, the methodology adjusts the inlet velocity to supply the maximum power during the whole time period. In another approach, the volume of the heat storage unit could be optimized to ensure that the heat demand can always be satisfied.

Additional scenarios with variable heat demands were also tested and are shown in Figure 7.

Figure 7 (a) shows a small baseline heat demand of 0.5kW and a larger power of 2 kW for 2 hours. This could be the heat demand of a room with a peak during the hours 6-8pm when people come back home and need more thermal comfort. In this case, the adsorption storage is able to track the heat demand efficiently, with an objective value of 0.06, adjusting the value of the inlet velocity as shown in Figure 7 (b).

Finally, a scenario representing domestic hot water needs has been tested. A baseline heat demand of 0.8kW is chosen and two peaks at 3kW, each lasting 10 minutes, are added. Again, the storage system is able to track the variable heat demand and adjust the fluid velocity, with an objective value of 0.06. Satisfying the real power of domestic hot water production in a house would again require several modules of storage operating in parallel.

Discussion

The design of the system was fixed in this study, following the prototype dimensions from [3], who tested two adsorption beds in series. Several beds are required to satisfy the heat demand of a whole house. However, this study focused on the dynamic behavior rather than the sizing of the system. Thus, heat demands for a single room were tested according to the size of the adsorption heat storage modelled. Future work should focus on the optimization of the design of the system according to the real heat demand of a house.

This study shows the capability of an adsorption heat storage to follow a heat demand of a house. The dynamic behavior is even able to track fast changes, such as the ones required for domestic hot water production. However, the velocity required to achieve those fast changes is high, leading to high operational costs due to the electricity consumption of the fan. In this study, the water concentration at the inlet is fixed because it is not controllable in most open configuration adsorption systems. Nevertheless, adding a humidifier to increase the inlet humidity of the air flow could help reach higher supplied power values without using as much electricity. Moreover, this added degree of freedom could help achieve a more efficient operation of the storage through dynamic optimization.

Another approach would be to change the system for a closed configuration allowing more control over the inlet variables.

CONCLUSION

In this work, a numerical model for a packed bed adsorption heat storage was developed. The model is based on mass and energy balances and uses a Langmuir isotherm to describe the adsorption phenomenon and a LDF model for the mass transfer.

A sensitivity analysis was conducted to choose the discretization parameters offering a good compromise between accuracy and model complexity. A time step of 100s and 20 spatial discretization points were chosen. The model was then validated against experimental data from the literature. The value of ΔH_{ads} in the model was adjusted within the range of experimental data to the value of - 56398 J/mol, in order to reduce the modeling error to less than 5%. Dynamic optimization was then performed on the chosen design. Optimal trajectories for the inlet velocity were computed in order to match a heating demand of part of a house. The methodology was able to track constant and variable heat demands with sums of quadratic errors less than 1%. When the heat demand is too high to be satisfied, the methodology plans the optimal pre-heating.

The results confirmed the potential of the dynamic optimization methodology for the future operation of a real adsorption heat storage unit. Future work will focus on the optimization of added degrees of freedom, such as the inlet humidity, as well as the simultaneous optimization of the design and operation of the system.

NOMENCLATURE

Symbol	Description	Unit
ε	Porosity	[-]
ρ_i	Species mass concentration in fluid	$[kg_i/m_f^3]$
ρ	Density	$[kg/m^3]$
ΔH_{ads}	Heat of adsorption	$[J/mol]$
C_p	Specific heat capacity at constant pressure	$[J/(kg.K)]$
C_v	Specific heat capacity at constant volume	$[J/(kg.K)]$
K	Langmuir equilibrium constant	$[Pa^{-1}]$
K_{LDF}	Linear Driving Force constant	$[s^{-1}]$
M	Molar mass	$[g/mol]$
$\bar{q_w}$	Average mass concentration of water in the adsorbed phase	$[kg_w/kg_s]$
q_{eq}	Mass concentration of water in the adsorbed phase in equilibrium with the fluid phase	$[kg_w/kg_s]$
$\dot{Q}$	Heating power	$[W]$
R	Ideal gas constant	$[J/(mol.K)]$
P	Pressure	$[Pa]$
T	Temperature	$[K]$
u	Intersticial velocity	$[m/s]$
Subscripts	**Description**	
b	Bed	
da	Dry air	
s	Solid	
w	Water	

ACKNOWLEDGEMENTS

The authors would like to acknowledge Dr. Sakshi Naik for her valuable support and insightful contributions to the development of the Pyomo code.

AUTHOR IDENTIFIERS

Author ORCIDs:
Untrau A: 0000-0002-6855-0768
Biegler L. T: 0000-0003-3875-4441
Sochard S: 0000-0002-8955-1575

REFERENCES

1. Lefebvre D, Tezel FH. A review of energy storage technologies with a focus on adsorption thermal energy storage processes for heating applications. Renewable and Sustainable Energy Reviews 67:116-125 (2017). https://doi.org/10.1016/j.rser.2016.08.019
2. Yan T, Xu XK, Wang Z, Zhang H, Pan WG, Wang LW. Review on solid-gas sorption heat storage: materials, processes and system optimization. Journal of Energy Storage 100:113589 (2024). https://doi.org/10.1016/j.est.2024.113589
3. Tatsidjodoung P, Le Pierrès N, Heintz J, Lagre D, Luo L, Durier F. Experimental and numerical investigations of a zeolite 13x/water reactor for solar heat storage in buildings. Energy Conversion and Management 108:488-500 (2016). https://doi.org/10.1016/j.enconman.2015.11.011
4. Yu N, Wang RZ, Wang LW. Sorption thermal storage for solar energy. Progress in Energy and Combustion Science 39:489-514 (2013). https://doi.org/10.1016/j.pecs.2013.05.004
5. Ferreira S, Sochard S, Serra S, Marias F, Reneaume JM. Modelling of an adsorption heat storage system and study of operating and design conditions. Processes 9:1885 (2021). https://doi.org/10.3390/pr9111885
6. Zamengo M, Wu S, Yoshida R, Morikawa J. Multi-objective optimization for assisting the design of fixed-type packed bed reactors for chemical heat storage. Applied Thermal Engineering 218:119327 (2023). https://doi.org/10.1016/j.applthermaleng.2022.119327
7. Anilkumar S, Chandrakala B, Kumar EA. Design optimization of thermochemical reactor for 1 TR adsorption cooling systems. Applied Thermal Engineering 270:126186 (2025). https://doi.org/10.1016/j.applthermaleng.2025.126186
8. Wang Y, Xu Z, Liu S, Shen Y, Ji W, Chen T, Li Y. Multi-objective optimization of thermochemical energy storage systems with configuration planning for different applications. Applied Thermal Engineering 274:126578 (2025). https://doi.org/10.1016/j.applthermaleng.2025.126578
9. Hart WE, Watson JP, Woodruff DL. Pyomo: modeling and solving mathematical programs in python. Math. Prog. Comp. 3:219-260 (2011). https://doi.org/10.1007/s12532-011-0026-8
10. Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math. Program. 106:25-57 (2005). https://doi.org/10.1007/s10107-004-0559-y
11. Kim KM, Oh HT, Lim SJ, Ho K, Park Y, Lee CH. Adsorption equilibria of water vapor on zeolite 3A, zeolite 13X, and dealuminated y zeolite. J. Chem. Eng. Data 61:1547-1554 (2016). https://doi.org/10.1021/acs.jced.5b00927

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Direct iron reduction system analysis with mixture of hydrogen feed

Marwa Mortadi*, Carl Haikarainen, Guanwei Zhou and Henrik Saxén

Process and Systems Engineering Laboratory, Åbo Akademi University, Turku, Finland
* Corresponding Author: marwa.mortadi@abo.fi.

ABSTRACT

The present work explores computationally the potential of hybrid operation of a direct reduction shaft furnace with a feed gas mixture of methane and hydrogen, considering the operation of the associated gas handling and conversion units (reformers, condensers, heat exchangers, compressors, heaters, etc.). The state of the system is dependent on both the performance of the shaft furnace, which is sensitive to the feed gas composition and quantity, and the reformer, which converts the methane and the recycled top gas from the shaft furnace into feed gas. This coupling gives rise to nonlinear behavior of the system with respect to changes in the boundary conditions, which justifies a model-based approach for holistic analysis. The system is modelled in Aspen Plus to simulate the operation of an industrial-scale plant using a detailed model of the process units. Three configurations of the system are evaluated: (i) co-feeding of hydrogen with methane, (ii) feeding hydrogen directly to the shaft furnace, and (iii) splitting the hydrogen feed between the two injection points. Further, the effect of H_2 share in the feed was evaluated in terms of energy consumption, CO_2 emissions and energy cost with low- and relatively high-carbon electricity. The results show that increasing the share of hydrogen in the fresh feed leads to lower CO_2 emissions, considering low-carbon electricity.

Keywords: Hydrogen, Direct reduction, Steam reforming, CO2 emissions, Energy cost

INTRODUCTION

The iron- and steelmaking industry is the foremost source of industrial CO_2 emissions with a share of 27%, emitting 2.6 Gt/a in 2022 [1], [2]. The sector is particularly energy intensive accounting for 20% of global industrial energy demand, relying mostly on fossil sources. The global production of steel has almost doubled within the last two decades, creating an urgent need to accelerate the decarbonization of this sector [3]. Various pathways have been investigated for a green transition of iron and steel industries, relying on process efficiency and switching to alternative feedstocks and fuels with lower carbon footprints. A strong concept that has been explored is the replacement of blast furnaces, which rely heavily on coke and coal, by direct reduction shaft furnaces (SF) operating on hydrogen-rich gas mixtures. A full transition to hydrogen-based iron ore reduction is a promising pathway to limit the carbon footprint. However, there are still challenges that hinder its adoption by steelmakers, notably the endothermic nature of the reduction reactions, the lower heat capacity of the gas, and, foremost, the high price of green hydrogen. Recent studies have indicated that it may be beneficial to include some carbon monoxide in the feed gas due to its positive impact on the heat supply and less endothermic reduction reactions [4].

The MIDREX process is one primarily gas-based technology for direct reduction of iron ore, representing about 80% of SF-based global steel production [5]. To support the transition toward H_2-based reduction, the MIDREX-flex process offers flexible operation using both natural gas (NG) and H_2, allowing a smooth and controllable transition for steelmakers [6]. It also enables the retrofitting of existing plants rather than requiring a complete redesign. However, the process behavior during this transition is not clearly known as most of the research works evaluate a complete transition to H_2.

The current study evaluates computationally the operation of an iron ore direct reduction process with top gas recycling by gas handling and conversion units

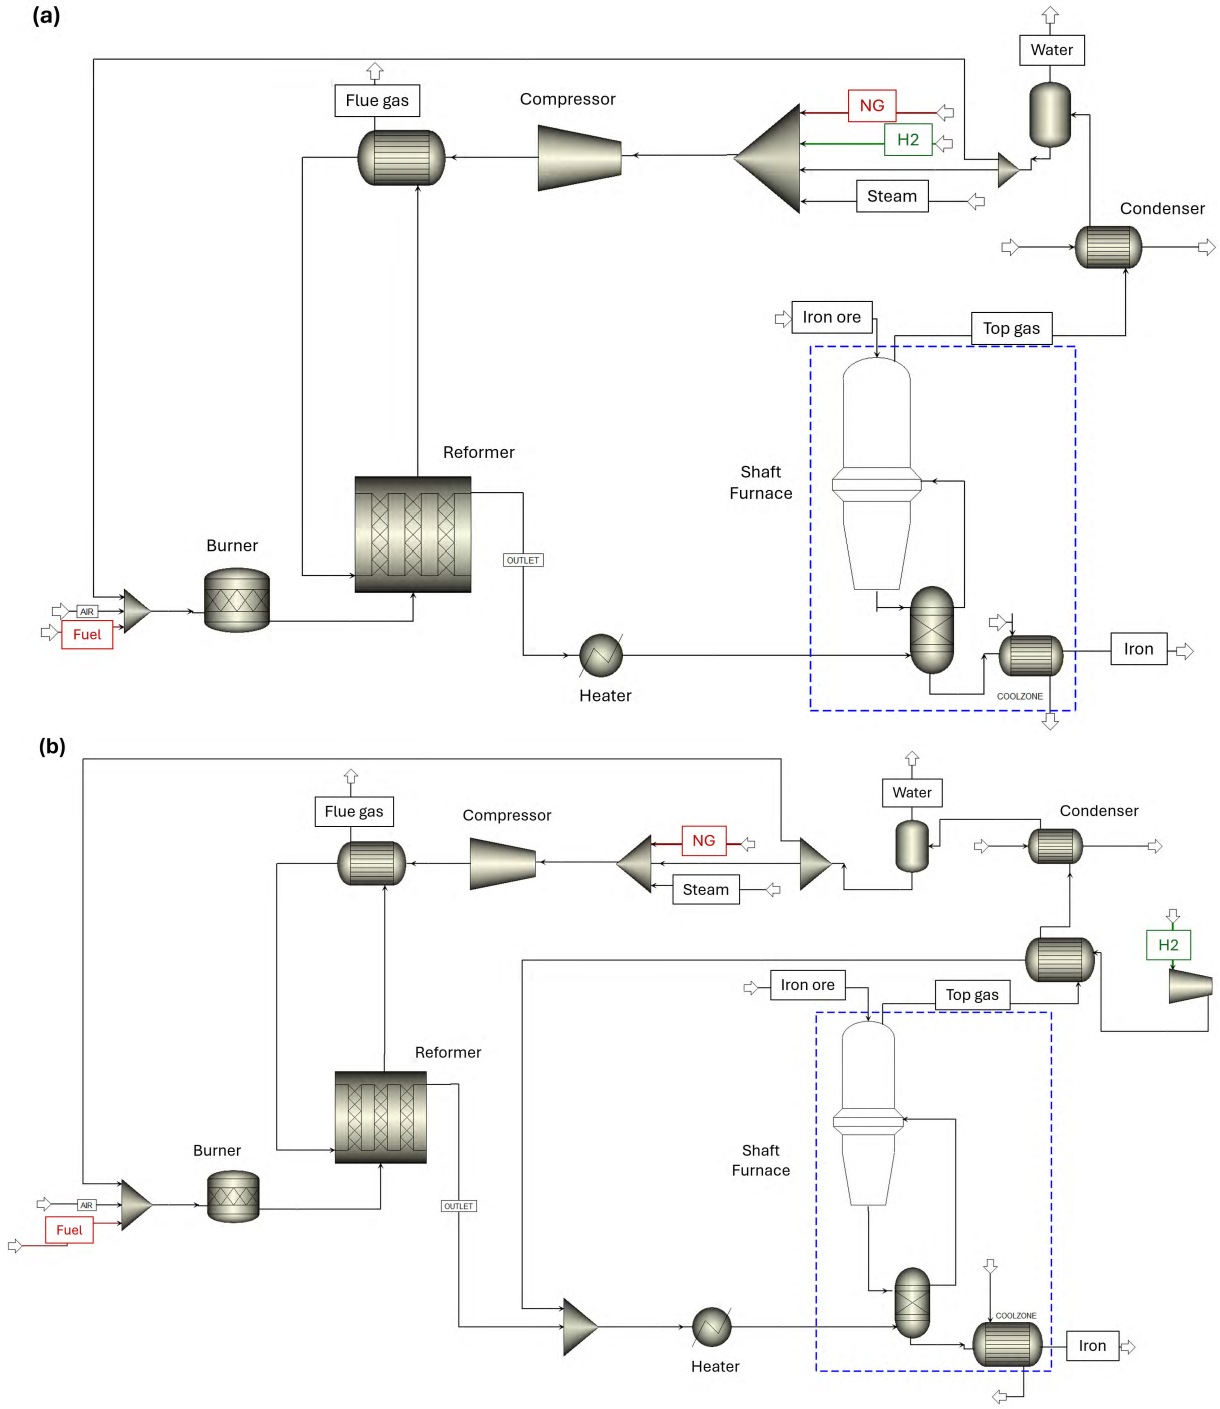

Figure 1. Process flowsheet of the direct reduction process.

(reformers, condensers, heat exchangers, compressors, heaters, etc.). The process operation is studied under varied shares of H_2 in the feed with different possible injection points in the system. Fresh gas and fuel demands, energy consumption and costs are analysed along with the total CO_2 emissions. The study provides guidelines for possible future operation states of iron direct reduction.

METHODOLOGY

Process description

The current work considers a shaft furnace

Table 1: Process and fuel gas consumption for the evaluated cases.

Description	**Case 1** H_2 injection at reformer upstream	**Case 2** H_2 injection at SF level	**Case 3** H_2 injection split at reformer upstream and SF inlet	**Case 4** Only H_2 based reduction	**Case 5** Only CH_4 based reduction
H_2 feed (kNm^3/h)		27.11		80.69	—
CH_4 feed (kNm^3/h)			22.09		31.38
Fuel consumption (kNm^3/h)		56.03		—	56.03

employing rich gas mixtures for the iron ore reduction. The process operates with both H_2 and NG, where H_2 represents 55 mol% of the combined feed. The shaft furnace top gas is recycled back to the system after water removal and mixing with the fresh feed. The process gas is compressed and preheated before entering a steam methane reformer (SMR) for converting CH_4 into a rich mixture of H_2 and CO, in the presence of recycled top gas. An electric heater is included for further heating of the syngas. It is important to note that 70% of the SF top gas is recycled as process gas, while the other part is fed to the reformer burner along with fuel gas (NG). **Figure 1 (a)** illustrates the process layout of the base case scenario, where both H_2 and NG are fed to the reformer.

The paper investigates two other cases with different injection points of H_2. The second case assesses the scenario where H_2 is supplied directly at the inlet of the SF after being preheated by the top gas, as represented by **Figure 1 (b)**. In the third case, the H_2 feed is split so that half of it is fed with the NG and the other half is supplied at the SF inlet. Fully H_2-based and CH_4-based direct reduction scenarios are considered as reference cases, corresponding to cases 4 and 5, respectively. **Table 1** summarizes the cases with H_2, CH_4, and fuel consumptions.

It is worth mentioning that the feed of H_2 and CH_4 is kept constant for the three first cases, while for cases 4 and 5 it is varied to achieve a production of 100 t_{DRI}/h with a metallization rate of 0.93, which is the target production for all the cases.

Model framework

Reformer model

The system is modeled using Aspen Plus software (V14). First, the model of the SMR is a plug flow reactor with co-current thermal fluid to simulate an industrial reformer with multi-tubular reactor geometry. The reactor catalytic tubes have a length of 7.9 m, with an internal diameter of 0.2 m [7], and a heat transfer coefficient estimated as 750 W/m^2 K [8]. The Peng-Robinson thermodynamic model is used for mass and energy balances of the system, while the pressure drop throughout the reformer is estimated using the Ergun equation.

The main chemical reactions occurring are the steam methane reforming reaction (Eq. (1)) and the water-gas-shift reaction (Eq. (2)):

$$CH_4 + H_2O \rightleftharpoons CO + 3H_2 \; \Delta H_{298\,K} = 206.1\; kJ/mol \quad (1)$$

$$CO + H_2O \rightleftharpoons CO_2 + H_2 \;\; \Delta H_{298\,K} = -41.15\; kJ/mol \quad (2)$$

The reactions kinetics over nickel-alumina catalysts have been widely investigated in literature [9], [10], [11]. The Langmuir-Hinshelwood-Hougen-Watson (LHHW) model is employed to determine the reaction rates, as an accurate model that considers surface reactions making it suitable for modelling catalyst-based processes [12], such as methane reforming.

Shaft furnace model

A one-dimensional open source steady-state model, developed by Aspen Technology [13], is used to simulate the SF. The model was implemented in Aspen Custom Modeler (ACM), modeling a counter-current flow of descending iron ore pellets and ascending reducing gases. An RGibbs reactor is integrated to simulate the water-gas-shift and methane reforming reactions over an iron oxide catalyst, and a heat exchange is employed for cooling down the metallized iron to 100 °C. The key reactions for the direct reduction can be expressed as follows:

$$\mathrm{Fe_2O_3 + 3H_2 \longrightarrow 2Fe + 3H_2O} \quad (3)$$

$$\mathrm{Fe_2O_3 + 3CO \longrightarrow 2Fe + 3CO_2} \quad (4)$$

The evaluation of energy consumption considers the required compression power, heating duty, and H_2 production via electrolysis. Total carbon dioxide emissions are investigated by accounting the emissions linked to process electricity consumption in addition to the

direct emissions with flue gases. Finnish and EU average energy and emissions contexts are used for the analysis.

SIMULATION RESULTS

Effect of hydrogen feed point

To find optimal operation points of the direct reduction process, the possible injection points of hydrogen are evaluated as described in **Table 1**. The system behavior is analysed in terms of the metallization rate of direct reduced iron (DRI), energy consumption, and total CO_2 emissions.

Table 2 lists the energy consumption of the process in terms of electricity consumption (EC) and primary energy consumption (PEC). Comparing the three injection points of H_2 (Cases 1—3), the first scenario, where hydrogen is co-fed with methane, enables the least energy consumption. This is because less heating is required at the inlet of the SF to maintain syngas temperature at 940 °C. The process gas exits the reformer with relatively high temperature (around 850 °C). Thus, the heater needs to provide a limited temperature increase, whereas the alternative injection options involve cooler gas entering the heater and therefore increasing the heating duty. On the other hand, it is observed that H_2-based direct reduction requires higher energy input, most of it in the form of electricity, with a large share for powering the electrolyser (83%).

Table 2: Electricity and primary energy consumption of the system in different scenarios.

	PEC (MW)	EC (MW)
Case 1	549.8	160.9
Case 2	553.3	163.8
Case 3	551.6	162.4
Case 4	647.1	471.5
Case 5	406.9	32.5

With regards to the syngas composition, it is found that the injection point of H_2 does not have a significant effect on the molar composition. In the current steady-state model, the feed of H_2 and CH_4 is similar for the three cases, in addition to the fuel consumption that permits a CH_4 conversion around 0.9. Thus, the overall mass balance of the system remains nearly unchanged.

Figure 2 represents the total CO_2 emissions of the system considering (a) Finnish and (b) EU average contexts. The grid emission factors used for the study are 29.5 and 187 gCO_2/kWh, respectively, based on the statistics of 2024 ([14], [15]). With regards to the different scenarios (cf. **Table 1**), direct emissions are the major contributor to total CO_2 emissions, considering the current share of H_2 in the system. Further, H_2-based direct reduction enables total emissions abatement in the order of 81%, compared with methane-based reduction process, in the Finnish context. However, the emissions of hydrogen-based system rose to surpass the conventional system (case 5) by 12% in the average EU context. Therefore, the hybrid and hydrogen-based processes are environmentally attractive only when electricity has lower carbon footprint and a higher part of renewables.

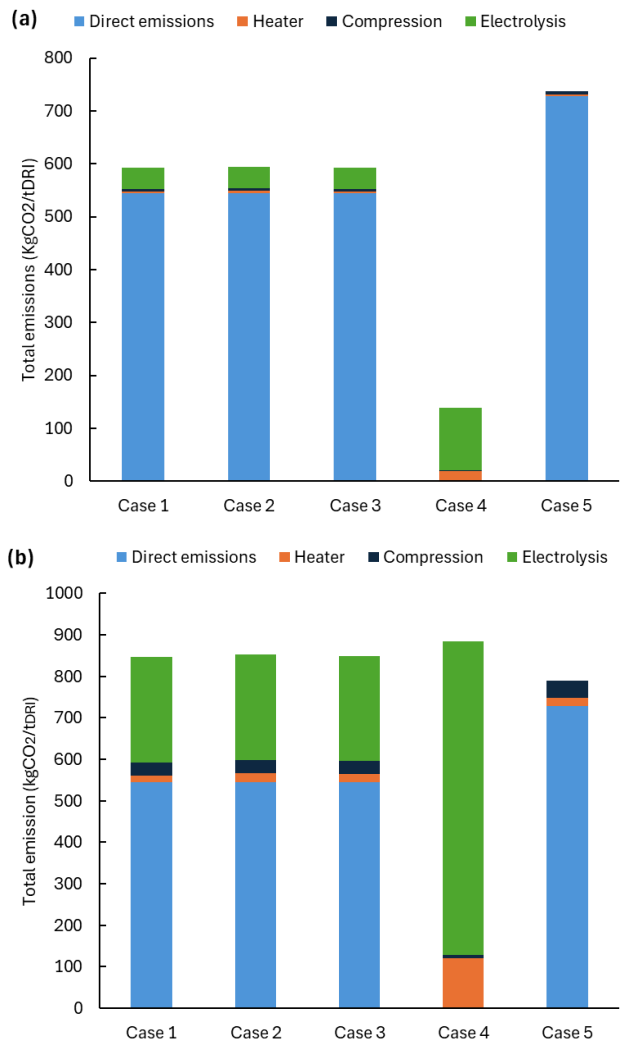

Figure 2. Total emissions of the system under different scenarios for (a) Finnish and (b) EU average contexts.

The analysis of energy cost shows that the reduction route combining H_2 and CH_4 results in a 5% cost reduction, compared with CH_4 based DRI process, while the fully hydrogen-based route allows 32% of energy cost reduction, in the Finnish scenario (**Table 3**). In contrast, for the average EU context, the energy cost was increased by 1.6 to 3 folds, comparing with the conventional system, due to higher electricity consumption.

Table 3: Energy cost under different scenarios (€/t_{DRI})

	Finland	EU average
Case 1	300.8	348
Case 2	302.2	352.1
Case 3	301.5	350.1
Case 4	217.4	652.2
Case 5	317.7	212.5

Effect of hydrogen share in the feed

To evaluate the process behavior under different shares of H_2, **Figure 3** illustrates the variation of energy consumption and CO_2 emissions. It is observed that as hydrogen share increases in the fresh feed, the energy consumption increases as well, because of increased electricity demand. Total emissions of the system decrease with higher shares of hydrogen, caused mainly by lower direct emissions. Note that the Finnish context is considered in this part of the study.

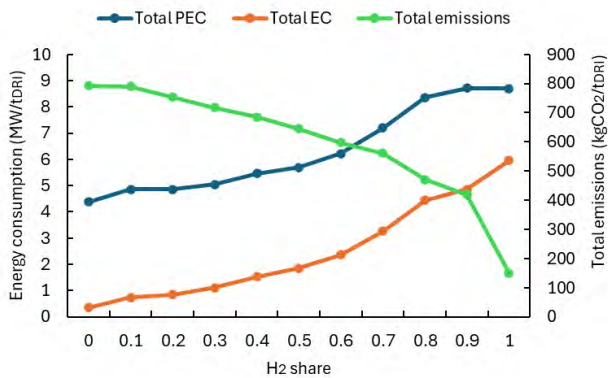

Figure 3. Effect of H_2 share on energy consumption and CO_2 emissions.

Figure 4 shows the variation of energy cost with the hydrogen share in the feed. Following the increase of electricity consumption, the electricity cost rises linearly with the H_2 part. At low to moderate H_2 shares (up to 0.7), the cost of NG, used as reducing agent and fuel, is the primary driver for the energy cost. In these cases the process relies significantly on NG to achieve the desired iron ore reduction rate while the cost of NG in Finland is relatively high [16]. Therefore, the transition to hydrogen-based direct reduction is economically wise in regions where electricity is cheaper. Total energy cost remains approximately constant over a wide range of hydrogen shares, as the savings in NG cost are offset by the rising electricity demand.

It is worth mentioning that more in-depth economic analysis is required to assess the feasibility of the green transition, accounting for investment costs in addition to operational costs.

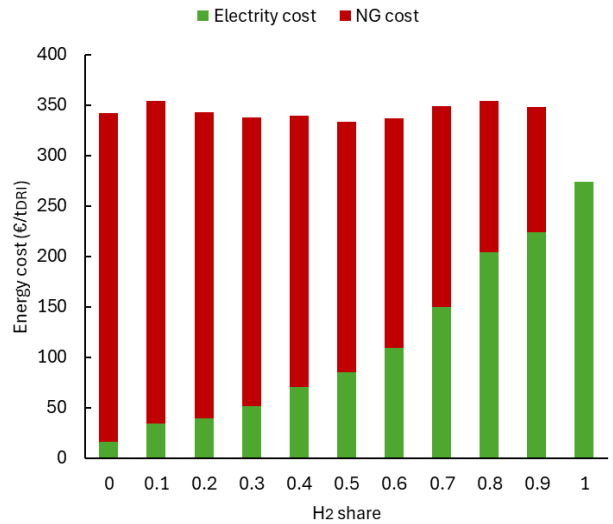

Figure 4. Effect of H_2 share on energy cost.

CONCLUSION

The present work investigates a process for direct reduction of iron using a mixture of methane and hydrogen. The system includes steam methane reforming and a direct reduction shaft furnace in addition to heat exchangers, a condenser, and a heater, which are modeled using Aspen Plus software. The effect of having different hydrogen injection points in the process was evaluated, along with the variation of H_2 share in the fresh gas feed. Energy consumption, total emissions, and energy cost are the key performance indicators employed to evaluate the process.

The results outline that the feeding point of hydrogen in the system has a negligible impact under steady state mode. However, co-feeding hydrogen with methane is favored for its relatively lower energy consumption. Evaluating the total CO_2 emissions shows that a hybrid route for direct reduction can reduce 25% of the emissions compared with fully methane-based reduction. Further, the carbon intensity of electricity is a critical determinant of overall system emissions; up to 97% of CO_2 can be avoided when low-carbon electricity is available, whereas a high-carbon electricity supply can limit these benefits and even lead to higher emissions than the conventional reference case. On the other hand, increasing the hydrogen part in the fresh feed leads to lower total emissions when low-carbon electricity is available. From an energy cost perspective, it is found that hydrogen-based direct reduction is economically wise. However, a more in-depth economic study is necessary to identify the cost effectiveness of the system.

ACKNOWLEDGEMENTS

The authors are grateful for the support from Business Finland and companies within the Towards Fossil-free Steel 2 project (decision number 5561/31/2023) and for the strategic support from the Research Council of Finland for university research profiling.

AUTHOR IDENTIFIERS

Author ORCIDs:
Mortadi M: 0000-0002-0379-0242

REFERENCES

1. International Energy Agency, "Industry - Energy system." Accessed: Sep. 26, 2025. [Online]. Available: https://www.iea.org/energy-system/industry
2. International Energy Agency, "World Energy Outlook 2024".
3. World steel association, "World steel in figures, " 2024. [Online]. Available: https://worldsteel.org/
4. Zhai Y, Shao L, Zhao C, Zhang X, Saxén H. A numerical study of the thermochemical and aerodynamic states of h₂?intensive direct reduction shaft furnace. steel research int. 96: (2024). https://doi.org/10.1002/srin.202400567
5. Sheng K, Wang X, Si F, Zhou Y, Liu Z, Hua H, Wang X, Duan Y. Rational capacity investment for renewable hydrogen-based steelmaking systems: a multi-stage expansion planning strategy. Applied Energy 372:123746 (2024). https://doi.org/10.1016/j.apenergy.2024.123746
6. "MIDREX Flex | Midrex Technologies." [Online]. Available: https://www.midrex.com/midrex-process/midrex-flex/
7. Toan Vu T, Seo J, Kim E, Ryoo SG, Park BC, Song D. Techno-economic analysis of integrated MIDREX process with CO2 capture and storage: evaluating sustainability and viability for iron production. Process Safety and Environmental Protection 189:1314-1322 (2024). https://doi.org/10.1016/j.psep.2024.07.005
8. Cui C, Vo DN, Zhao Y, Qi M, Xia M, Ramkrishna D, Masuku CM. Rigorous development and comparison of multi-dimensional reactor models encompassing the catalyst domains for steam methane reforming. Chemical Engineering Journal 496:153581 (2024). https://doi.org/10.1016/j.cej.2024.153581
9. Alhumaizi K, Ajbar A, Soliman M. Modelling the complex interactions between reformer and reduction furnace in a midrex?based iron plant. Can J Chem Eng 90:1120-1141 (2011). https://doi.org/10.1002/cjce.20596
10. Jun HJ, Park MJ, Baek SC, Bae JW, Ha KS, Jun KW. Kinetics modeling for the mixed reforming of methane over ni-ceo2/mgal2o4 catalyst. Journal of Natural Gas Chemistry 20:9-17 (2011). https://doi.org/10.1016/s1003-9953(10)60148-x
11. Snoeck JW, Froment GF, Fowles M. Steam/co₂reforming of methane. carbon formation and gasification on catalysts with various potassium contents. Ind. Eng. Chem. Res. 41:3548-3556 (2002). https://doi.org/10.1021/ie010665p
12. Trinca A, Patrizi D, Verdone N, Bassano C, Vilardi G. Toward green steel: modeling and environmental economic analysis of iron direct reduction with different reducing gases. Journal of Cleaner Production 427:139081 (2023). https://doi.org/10.1016/j.jclepro.2023.139081
13. "AspenTech." [Online]. Available: https://www.aspentech.com/en
14. "Statistical databases | Statistics Finland." [Online]. Available: https://stat.fi/tup/tilastotietokannat/index_en.html
15. "European Environment Agency EEA." [Online]. Available: https://www.eea.europa.eu/en
16. "Eurostat." [Online]. Available: https://ec.europa.eu/eurostat/en/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimization of Site-wide Heat-Integrated Utility Systems with Heat Pumps using MILP

Thorben Hochhaus[a*], Marcus Grünewald[a], and Julia Riese[b]

[a] Ruhr University Bochum, Laboratory of Fluid Separations, Bochum, Germany
[b] Paderborn University, Fluid Process Engineering, Paderborn, Germany
* Corresponding Author: thorben.hochhaus@fluidvt.rub.de.

ABSTRACT

The reduction of CO_2-emissions in the chemical industry is essential to meet European climate targets. Particularly, the reliance on fossil fuels for process heat supply is a key factor for CO_2-emissions. Electrically driven compression heat pumps are a promising option to reduce fossil fuel consumption by upgrading low-temperature waste heat to a higher temperature level, provided that low-carbon electricity is available. However, the integration of heat pumps into chemical utility systems remains a challenge due to economic constraints and the high complexity associated with site-wide heat integration and retrofit of existing structures. This work presents a mixed-integer linear programming (MILP) approach for the optimization of utility systems with integrated heat pumps. To address computational complexity, candidate utility temperature levels are pre-selected, and feasible heat pump coefficients of performance (COP) are precomputed. The framework is applied to both greenfield and retrofit scenarios for a synthetic case study consisting of 400 process streams. In the greenfield scenario, optimal utility temperature levels and heat pump integration configurations are identified. For the retrofit scenario, temperature levels of an existing utility system are modified to reduce total annual costs (TAC). Additionally, sensitivity analysis is conducted to assess the influence of key economic and environmental parameters. The presented case studies demonstrate short solution times, highlighting the suitability of the proposed framework for screening studies and systematic sensitivity analyses in early-stage design and retrofit applications.

Keywords: Energy Systems, Optimization, Renewable and Sustainable Energy

INTRODUCTION

The CO_2-emissions of the chemical industry must be reduced significantly to meet Europe's climate goals of zero net-emissions by 2050 [1]. A large proportion of the process heat demand is still met by burning fossil fuels such as natural gas which results in the emission of CO_2 [2]. Therefore, an effort should be made to reduce the share of process heat supplied from fossil fuels. This can be achieved by improving the efficiency of chemical processes or using alternative, low-carbon sources of heat supply. Compression heat pumps offer a possibility of providing heat with low CO_2-emissions when powered by electricity with a low CO_2-emission factor [3]. The use of formerly unused waste heat which is lifted to a higher temperature level by the heat pump can substitute process heat provided by fossil fuels.

Despite the advantages of heat pumps for providing process heat with low CO_2-emissions, heat pump integration into chemical processes is still limited by mayor challenges. Among these are economic feasibility and the inherent complexity of the systems involved. [4] In the context of chemical sites, complexity arises, on the one hand, from the large number of potential heat sources and sinks that must be considered for heat integration. On the other hand, both heat integration and the integration of heat pumps increase system complexity by introducing intra- and inter-process dependencies. The integration task then includes the appropriate sizing of heat pumps and identifying optimal integration points that allow independent operation of different considered processes.

In chemical sites, heat is usually provided to the processes via different utility levels such as steam and hot

water at different temperatures. Waste heat from the processes is either used to supply heat to the heating utilities or collected via cooling utilities such as cooling water. The temperature levels of these utilities are often based on historical design decisions. Over time, the design may have become suboptimal due to changes in process conditions or configurations as well as evolving energy trends such as energy price developments and alternative sources of heat supply.

The scope of this contribution is heat pump integration between different temperature levels of the site-wide utility system. This allows the utilization of surplus heat from the lower level and electricity to supply heat to the higher temperature utility level to substitute fossil-based heat. While the integration of heat pumps between different utility levels is limited in efficiency compared to the integration in selected process heat sources and sinks with lower temperature lifts, it offers several advantages. Centralized heat pump integration requires only one power line, it mitigates space limitations associated with integration close to the processes and the utility system may act as a buffer in case of operational fluctuations. Increasing the efficiency of centrally integrated heat pumps may require changing the temperature levels of the existing utility system to optimize waste heat usage and resulting temperature lifts for heat pumps which will be discussed in this work.

LITERATURE REVIEW

The concept of pinch analysis, introduced by Linnhoff and Hindmarsh in 1983, represents a foundational approach to process heat integration and enables the identification of minimum external energy requirements [5]. Using the problem table algorithm, feasible heat transfer from higher to lower temperature levels is ensured through heat cascading [6]. The methodology was later extended to chemical sites by the concept of total site heat integration which allows targeting and optimization of utility levels as well as cogeneration potentials [7].

Building on these concepts, mathematical optimization approaches have since been developed to address increasingly complex heat integration and utility design problems in process plants and chemical sites. Papoulias and Grossmann proposed MILP optimization based on the transshipment model to minimize utility consumption in a heat recovery network by splitting the temperature range into temperature intervals and solving interval energy balances [8]. Bagajewicz and Barbaro introduced a non-linear targeting approach that optimizes heat pump integration in total sites with continuous decision variables for heat pump temperature levels [9]. Oluleye et al. incorporated heat pump integration in the optimization of existing site cogeneration systems by means of MILP modeling [10]. They focused on different heat upgrading technologies and optimal operational parameters to minimize emissions and total annual costs. Wallerand et al. presented a superstructure-based MINLP optimization of heat pumps into industrial processes [11]. They used a decomposition strategy to solve the complex non-convex optimization task by iteratively solving non-linearities in a master level and technology selection by means of integer variables in a slave level.

The combinatorial complexity of heat transfer between process streams, utility levels and technologies such as cogeneration and heat pumps poses significant challenges for large-scale optimization problems. This may lead to convergence issues for large scale problems which is why many studies employ several simplifications to reduce the solution space. Simplifying utility system optimization is often done by applying simplified equipment models [9, 10, 12], discretizing or using fixed utility candidate temperature levels often from heuristic preselection [13, 14, 15] or linearization of non-linear constraints such as cost functions [16]. While this might exclude the global optimum from the solution space, applying simplified models allow faster solution times which can be valuable for gaining an understanding of the system, optimization under uncertainty or scenario and parameter studies [17].

Heat pump integration into utility systems of chemical sites has been subject to several studies as it provides a valuable option for waste heat utilization and reduction of external heating services [4, 9, 10]. Liew and Walmsley propose a total site targeting method that incorporates open cycle vapor recompression heat pumps with the goal of decreasing boiler energy demand to supply high quality steam [18]. Miah et al. proposed a framework for identifying possible heat integration at different scales incorporating heat pumps to upgrade waste heat and substitute external heating [19].

While previous optimization-based approaches address heat pump integration in utility systems, they often rely on fixed or finely discretized or continuous candidate temperature levels, which can be computationally expensive for large-scale systems. This limits their applicability for extensive parameter studies such as future price scenarios. Based on these limitations, this contribution proposes a mixed-integer linear programming framework for integrating heat pumps into utility systems of chemical plants. Computational complexity is reduced by a structured precomputation of candidate utility temperature levels and corresponding heat pump coefficients of performance, enabling short solution times. The approach is applicable to large-scale systems with several hundred heat sources and sinks and is well suited for screening studies under varying economic and operational conditions.

METHODOLOGY

Problem formulation

In chemical sites, heating and cooling services for cold and hot process streams are commonly supplied by site-wide utility systems with different temperature levels, often realized by steam at different pressure levels or water at different temperatures. This allows providing heat to cold streams with a temperature level higher than the cold stream target temperature plus the required minimum temperature difference in heat exchangers. The same principle applies to hot streams which can be cooled down by supplying heat to a utility level with a lower temperature than the process target temperature minus the required minimum temperature difference in heat exchangers. This way, heat surplus from hot process streams can substitute external heat supply to the utility level to supply the cold stream heat demand. The degree of heat integration in existing plants can often be improved by optimizing the utilisation of surplus heat to possibly maintain high quality heat in terms of exergy. While this might increase required heat exchanger area, it may further decrease external heat demand. The same principle applies to heat sinks, where the optimal temperature level for supplying heat to the process stream is not always selected. This aspect becomes particularly relevant when heat pumps are used to supply heat to utility systems. The COP of heat pumps decreases with higher temperature differences between heat source and heat sink and they are limited to sink temperatures of around 200°C [20, 21]. Therefore, process heat sources and heat sinks should be connected to optimal temperature levels to increase the potential for heat pump integration in utility systems of chemical sites.

This can become a challenging task in chemical sites as often several hundred hot and cold streams are connected to the utility system as they cannot be integrated in the processes themselves. The MILP formulation proposed in this work considers the optimization of utility systems. Simple constraint modifications in the MILP allow the analysis of retrofit scenarios as well as the design of new utility systems. The user must supply duties and target temperatures of hot and cold process streams. The approach considers central closed cycle heat pump integration between different temperature levels of the utility system and connections between process streams and utility levels. Optimal utility level temperatures, external heating or cooling requirements for every utility level, optimal connections between streams and utility levels as well as heat pump design and integration are selected during the optimization.

Candidate temperature level reduction

To enable fast MILP solution times, the number of candidate utility temperature levels must be reduced to reduce combinatorial complexity. In this approach, candidate temperature levels are selected by identifying the highest and lowest feasible temperature levels for hot and cold streams, respectively, that still allow reaching the target temperatures. These levels correspond to the temperatures at which kinks occur in the hot and cold composite curves. While this enables fast calculation, solutions requiring intermediate utility levels for cascaded heat pump configurations are excluded, which is discussed later as a model limitation.

Heat pump modeling and COP precomputation

As the presented approach should serve as a screening support in decision-making, a simple heat pump model that incorporates thermodynamic principles is applied. The COP is modeled as a function of the Carnot COP by including source (T_{source}) and sink (T_{sink}) temperatures, required temperature differences in heat exchangers ($\Delta T_{min,source}$ and $\Delta T_{min,sink}$ for source and sink respecively) and compressor efficiency (η_C) as shown in equation (1).

$$COP = \eta_C \cdot \frac{T_{sink}+\Delta T_{min,sink}}{T_{sink}+\Delta T_{min,sink}-(T_{source}-\Delta T_{min,source})} \quad (1)$$

The maximum temperature lift was assumed to be 60 K for single stage heat pumps with a maximum sink temperature of 200°C [20]. As the candidate temperature levels are predefined, COP values for all feasible combinations of source and sink utility level were also calculated to avoid nonlinearities in the optimization model. The heat pump must further follow the energy balance shown in equation (2) with electricity power demand P_{el}, heat demand at the source temperature $\dot{Q}_{source}$ and heat supply at the sink temperature $\dot{Q}_{sink}$ and the COP relation between heat supply and electricity power demand as shown in equation (3).

$$\dot{Q}_{sink} = \dot{Q}_{source} + P_{el} \quad (2)$$

$$COP = \frac{\dot{Q}_{sink}}{P_{el}} \quad (3)$$

Another reduction of model complexity was achieved by eliminating thermodynamically infeasible matches between process streams and utility levels and only keeping feasible matches as sets to be considered during optimization.

Objective function and constraints

The optimization model is set up as a MILP which minimizes the objective function c subject to sets of linear equality and inequality constraints, represented by g and f, respectively. The decision variables include binary variables x and continuous variables y.

$$\min_{x,y} c(x,y) \quad (4)$$

$$\text{s.t. } g(x, y) = 0$$
$$f(x, y) \leq 0$$
$$x \in \{0, 1\}^m, y \in \mathbb{R}^n$$

The MILP was solved in Python using Gurobi Optimizer [22] on a desktop computer equipped with an Intel® Core™ i7-12700 CPU @ 2.10 GHz (12 cores, 20 threads). The objective function can be chosen to be TAC, operational expenditure (OPEX), capital expenditure (CAPEX) or CO_2-emissions. In this work, only heat pump investment costs were considered as we focused on retrofit scenarios. We assume electricity driven heat pump compressors and natural gas boilers as existing external heating equipment. Detailed cost calculations and all assumed cost parameters are given in the supplementary material.

Binary variables control the existence of components such as heat pumps and utility temperature levels and the connection between process streams and utility levels. We further define that every process stream must be connected to exactly one utility level.

To accurately model the utility system, energy balances were set up for every utility level. As depicted in Figure 1 of the supplementary material, each utility level can supply heat to cold process streams and receive heat from hot process streams. If that results in a heat deficit, heat must be supplied to this utility level either by an external heat source (e.g. natural gas boiler), a heat pump or through heat cascading from a higher temperature level. If it results in a heat surplus, the utility level must reject heat either via external cooling, to a heat pump or through heat cascading to a lower utility level.

Model limitations

As this modeling approach is designed for screening purposes in utility system optimization, it focuses on fast solution times. In the following, we briefly discuss several simplifications that must be considered when interpreting the results. The first simplification arises from the candidate temperature level reduction. By only allowing temperature levels based on process stream data, heat pump integration might be limited e.g. in cases with gaps between candidate utility levels that exceed the maximum allowable temperature lift of a heat pump. Another limitation comes from the shortcut heat pump model which neglects detailed design of heat pump components as well as working fluid selection.

Moreover, no investment cost for new utility levels is considered as this would require tracking the length of required piping which introduces nonlinearities. While this might be sufficient in retrofit scenarios, where only temperature levels of the existing utility system are adjusted slightly, this becomes important in scenarios where new utility levels are added. For screening purposes, fixed investment costs can be assumed for each candidate levels to address this simplification. The same applies to heat exchangers, where no additional costs were considered, as heat exchangers were assumed to exist in retrofit scenarios. It should be checked whether the existing heat exchanger area is sufficient for the proposed utility level and process stream connection. Fixed cost parameters for additional heat exchanger area requirements could be provided as input values for every process stream with predefined utility levels. For heat pumps, fixed specific investment costs were considered (see supplementary material). As effects of scale exist for heat pumps [20], this could be accounted for by including linear estimations of cost correlations from the literature. We further assume no temperature glide in heat exchangers on the utility side. While this Is accurate for utility levels with latent heat supply and demand via condensation and evaporation, it neglects sensible heat transfer e.g. in cooling water. Finally, no stream splitting is allowed, which neglects potentials for higher energy efficiency especially with process streams with a high temperature difference between starting and target temperature.

CASE STUDY

The synthetic case study discussed in this work consists of 200 hot streams that require cooling and 200 hot streams that require heating. Inlet and outlet temperatures and required heat duties are given as input data, and the respective distributions as well as grand composite curve are presented in the supplementary material. The minimum temperature difference for heat exchangers for every stream of 10 K is provided for every stream. Heat cascading between utility levels is allowed for all scenarios in this work. Sensitivity analysis and parameter studies conducted on the case study are presented in the results and discussion section. The scenarios include analysis of optimal temperature levels for a new utility system as well as retrofit scenarios on an existing utility system.

RESULTS AND DISCUSSION

In the following, the results for several application scenarios for the proposed utility system optimization approach are presented and discussed.

Greenfield scenario

The first scenario is the design of a new utility system for the synthetic case study. This includes the selection of optimal temperature levels and the design of heat pumps. The maximum number of utility levels was increased stepwise from 2 to 15, and the resulting impact on TAC and the number of required heat pumps was evaluated for each configuration. All other parameters were

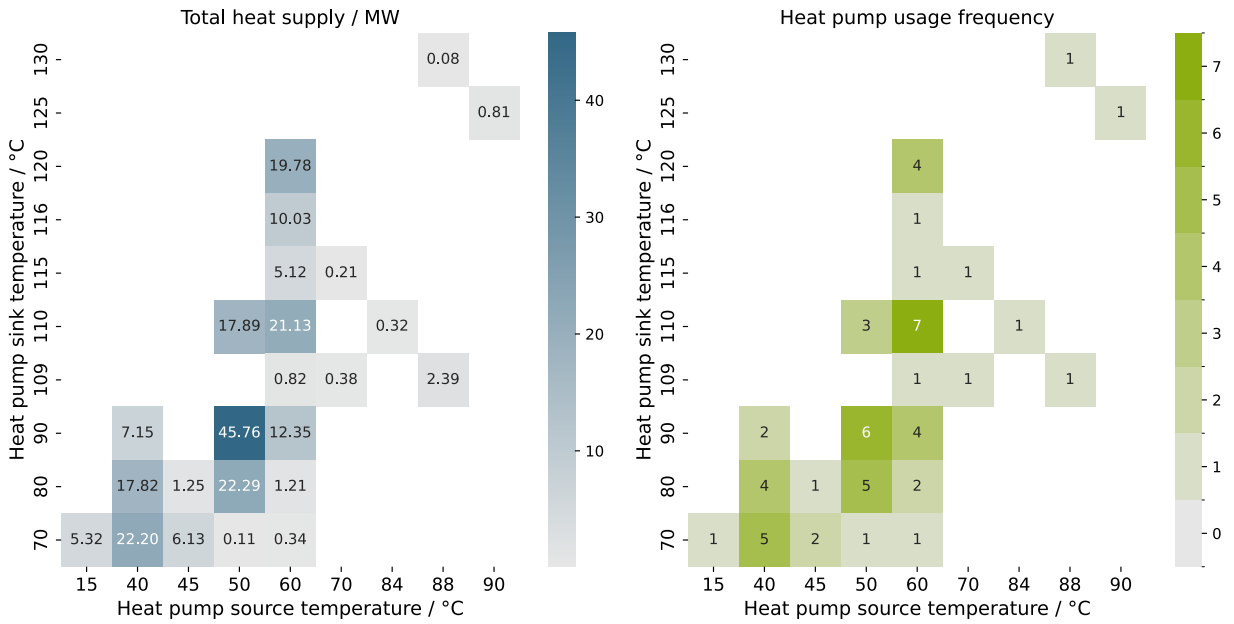

Figure 2: Total heat supply of heat pumps and usage frequency for different configurations of maximum number of utility levels.

kept constant. Candidate temperature level reduction resulted in 156 possible utility temperature which results in an average solver time of 8.5 seconds per configuration. As shown in Figure 1, an increasing number of utility levels results in lower TAC as this allows heat integration and heat pump integration becomes more feasible.

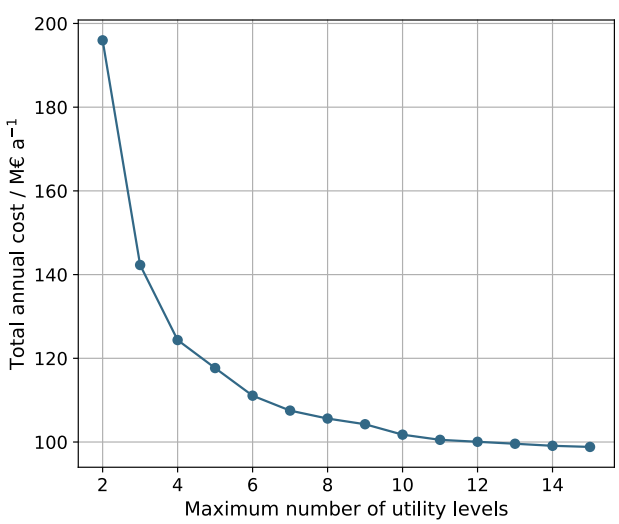

Figure 1. Total annual costs for different allowable utility level configurations.

In Figure 2, selected heat pump configurations for different maximum number of utility levels are shown. In the utility level configurations studied, heat pumps with a source and sink temperature of 60 °C and 110 °C were selected most frequently. The analysis further reveals the total heat supply provided by every configuration. Comparison with the selection frequency reveal configurations with a high individual substitution potential. Although the heat pump configuration with a source and sink temperature of 60 °C and 116 °C is only selected in one utility configuration it provides a high substitution potential of over 10 MW. Through this analysis, relevant heat pump configurations can be assessed quickly. It further supports decision-making in selecting an appropriate number of utility levels by clearly illustrating the trade-off between TAC and system complexity.

Moreover, the moderate computation times enable the evaluation of multiple cost assumptions, facilitating informed decisions and serving as a solid basis for subsequent detailed analysis.

Retrofit scenario 1

In the following, a retrofit scenario for the synthetic case study is discussed. The assumed utility system has temperature levels at 13 °C, 25 °C, 80 °C, 143 °C, 214 °C, 285 °C and 1600 °C. In a first analysis, the potential for changing the temperature of a single utility level is analyzed. The temperature may be adjusted by 10 K. Resulting TAC savings including heat pump investment costs are compared to the existing utility system structure without heat pump integration.

The lowest original temperature level at 13 °C is modified to 15 °C as this is sufficient to supply cooling to the stream with the lowest outlet temperature but it does not affect the potential for heat pump integration. No

modification could be identified for the levels 80 °C, 214 °C, 285 °C and 1600 °C to improve heat pump integration. TAC savings of 0.28 % can be achieved for the levels 13 °C, 80 °C, 214 °C, 285 °C and 1600 °C by integrating the same heat pump between the levels 25 °C and 80 °C. Only the modification of the utility level at 25 °C and 143 °C improve the potential for heat pump integration and therefore results in TAC savings of 0.36 % and 1.08 % respectively. To achieve this reduction in TAC, the level at 25 °C must be increased to 35 °C and the level at 143 °C must be decreased to 136 °C.

Retrofit scenario 2

In a second retrofit scenario, the same existing utility levels were assumed but two temperature levels could be modified simultaneously. Again, TAC compared to the existing utility system structure without heat pump integration were assessed and included annualized heat pump investment cost and operational cost such as cost for natural gas, electricity and emission certificates as described in the supplementary material. The TAC savings achieved by integrating heat pumps and modifying combinations of utility levels are shown in Figure 3.

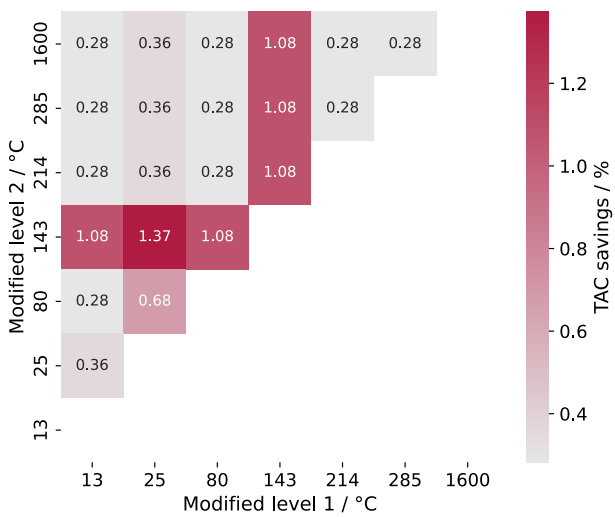

Figure 3. Effect of modifying two utility level temperatures on TAC savings.

While for some combinations, the same savings as in the first retrofit scenario can be seen as no additional heat pump integration is feasible, the combined modification of temperature levels 25 °C and 80 °C leads to TAC savings of 0.68 %. For this, the levels must be adjusted to 35 °C and 90 °C. By modifying level temperatures of 25 °C to 35 °C and 143 °C to 134 °C, higher TAC savings compared to the existing utility structure without heat pumps of 1.37 % can be achieved. Notably, when modifying the 25 °C and 143 °C levels simultaneously, the 143 °C level is adjusted to 134 °C instead of 136 °C, as observed in retrofit scenario 1 where only a single level was modified at a time.

As both retrofit scenarios only allow temperature adjustments in a range, the number of integer variables for candidate utility levels decreases significantly, resulting in average solver times of 0.2 seconds per configuration which underlines the potential for heat pump integration screening studies of this approach.

Retrofit scenario 3

In cases, where temperature levels of the utility system should not be modified, the complexity of the optimization problem is greatly reduced to only integrating and sizing of heat pumps as well as optimizing connections of process streams to the utility system to increase heat integration. One possible application for this is discussed in the last retrofit scenario. In this scenario, the synthetic case study is used with an existing utility system with the modified temperatures that led to the highest TAC savings in the second retrofit scenario (15 °C, 35 °C, 80 °C, 134 °C, 214 °C, 285 °C, 1600 °C). Sensitivity analysis is used to analyze the effect of different parameters on TAC. The results for varying electricity cost, emission certificate cost, specific heat pump investment cost as well as electricity emission factor are shown in Figure 4 where 0 % variation represents the base value.

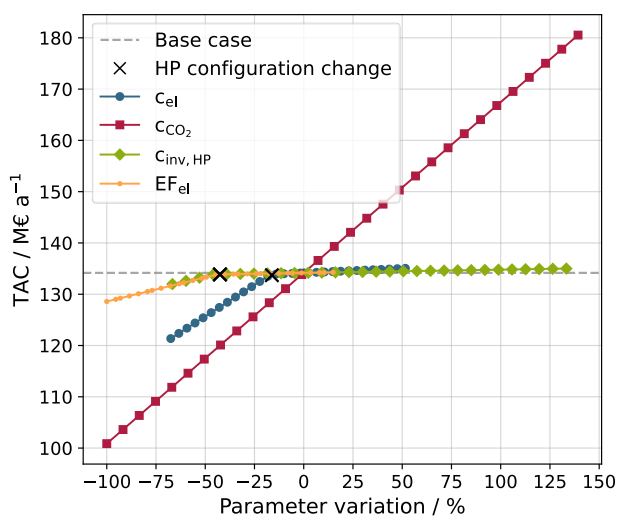

Figure 4. Sensitivity of total annual cost (TAC) to selected model parameters.

Due to the significant reduction of model complexity, the required computational time was around 0.1 seconds per parameter variation.

It is evident that CO_2 certificate costs for emissions associated with heat supplied by burning natural gas and electricity consumption in heat pumps have the highest impact on TAC. At high temperature levels, technical limitations prevent heat pumps from fully replacing external

heat demand, leading to residual emissions. With increasing emission certificate costs in the future, this underlines the importance of defossilization to stay economically competitive. For increases relative to the base value, the remaining parameters lead to only minor rises in TAC. For negative variations, a tipping point, marked as a change in heat pump configuration in Figure 4, can be observed for electricity cost, specific heat pump investment cost and the electricity emission factor. In the base configuration, a single heat pump is used to upgrade heat from the 35 °C utility level to 80 °C. For all cases located to the left of the markers, a second heat pump becomes feasible, enabling a further upgrade from 80 °C to 134 °C. Compared to the base values, a reduction of approximately 45 % in both the specific heat pump investment cost and the electricity emission factor is required for the second heat pump to become feasible, while a reduction of about 16 % in the electricity price is sufficient. In cases with two heat pumps, the total heat supplied by the heat pumps increases significantly, which explains the steeper gradient observed in the sensitivity analysis compared to cases where only one heat pump is feasible. Conducting sensitivity analysis can thus be an important decision support in combination with the utility system optimization in identifying critical model parameters or evaluating future cost and emission scenarios.

CONCLUSION AND OUTLOOK

In this work, a MILP framework for utility system optimization including heat pump integration was proposed. Model complexity was reduced by preselecting relevant utility level candidates, precomputation of feasible heat pump COP as well as elimination of thermodynamically infeasible matches between process streams and utility levels. While the proposed simplifications may exclude certain intermediate optima, they lead to a substantial reduction in computational effort. The resulting short solution times, as demonstrated in the presented case studies, make the approach well suited for screening studies and sensitivity analyses. The optimization framework can be applied to various optimization tasks, including greenfield design and the retrofit of existing utility systems. This was demonstrated with a greenfield design of a utility system, analyzing the impact of allowed utility levels on TAC and heat pump integration. Three retrofit scenarios were analyzed in which selected temperature levels of an existing utility system were modified, and promising adjustments were identified. In addition, the effects of different model parameters on TAC and heat pump integration were identified. Overall, the demonstrated computational performance supports the use of the model in early-stage design and retrofit studies, where multiple design alternatives and parameter variations must be evaluated efficiently. Future studies should further address the simplifications discussed in the model limitations section, for example by incorporating utility investment costs or additional heat exchanger costs, particularly for applications requiring a more detailed economic assessment. In addition, the influence of the candidate temperature level reduction on solution quality should be systematically assessed and compared to alternative discretization strategies, to quantify the trade-off between computational efficiency and optimality. Overall, the proposed framework provides a robust and computationally efficient basis for the systematic evaluation of heat pump integration options in complex utility system optimization problems.

DIGITAL SUPPLEMENTARY MATERIAL

Detailed description of the model, parameters for different case studies and additional figures. LAPSE:2026.0015

ACKNOWLEDGEMENTS

This work was funded by the Federal Ministry for Economic Affairs and Energy (BMWE) within the project WinPro (project funding reference number 03EN2115A).

AUTHOR IDENTIFIERS

Author ORCIDs:
Hochhaus T: 0000-0002-2566-6656
Grünewald G: 0000-0002-4007-8270
Riese J: 0000-0002-3053-0534

REFERENCES

1. United Nations, Paris Agreement. United Nations Framework Convention on Climate Change (2015) https://unfccc.int/sites/default/files/english_paris_agreement.pdf
2. Wolf S. Integration von Wärmepumpen in industrielle Produktionssysteme. Universität Stuttgart (2017) http://dx.doi.org/10.18419/opus-9593
3. Kiss AA, Smith R. Rethinking energy use in distillation processes for a more sustainable chemical industry. Energy 203:117788 (2020). https://doi.org/10.1016/j.energy.2020.117788
4. Schlosser F, Jesper M, Vogelsang J, Walmsley TG, Arpagaus C, Hesselbach J. Large-scale heat pumps: applications, performance, economic feasibility and industrial integration. Renewable and Sustainable Energy Reviews 133:110219 (2020). https://doi.org/10.1016/j.rser.2020.110219
5. Linnhoff B, Hindmarsh E. The pinch design method for heat exchanger networks. Chemical

Engineering Science 38:745-763 (1983). https://doi.org/10.1016/0009-2509(83)80185-7
6. Linnhoff B, Flower JR. Synthesis of heat exchanger networks: I. systematic generation of energy optimal networks. AIChE Journal 24:633-642 (2004). https://doi.org/10.1002/aic.690240411
7. Dhole VR, Linnhoff B. Total site targets for fuel, co-generation, emissions, and cooling. Computers & Chemical Engineering 17:S101-S109 (1993). https://doi.org/10.1016/0098-1354(93)80214-8
8. Papoulias SA, Grossmann IE. A structural optimization approach in process synthesis—ii. Computers & Chemical Engineering 7:707-721 (1983). https://doi.org/10.1016/0098-1354(83)85023-6
9. Bagajewicz MJ, Barbaro AF. On the use of heat pumps in total site heat integration. Computers & Chemical Engineering 27:1707-1719 (2003). https://doi.org/10.1016/s0098-1354(03)00149-2
10. Oluleye G, Jobson M, Smith R. Process integration of waste heat upgrading technologies. Process Safety and Environmental Protection 103:315-333 (2016). https://doi.org/10.1016/j.psep.2016.02.003
11. Wallerand AS, Kermani M, Kantor I, Maréchal F. Optimal heat pump integration in industrial processes. Applied Energy 219:68-92 (2018). https://doi.org/10.1016/j.apenergy.2018.02.114
12. Shang Z, Kokossis A. A transhipment model for the optimisation of steam levels of total site utility system for multiperiod operation. Computers & Chemical Engineering 28:1673-1688 (2004). https://doi.org/10.1016/j.compchemeng.2004.01.010
13. Colmenares TR, Seider WD. Synthesis of utility systems integrated with chemical processes. Ind. Eng. Chem. Res. 28:84-93 (2002). https://doi.org/10.1021/ie00085a015
14. Varbanov PS, Doyle S, Smith R. Modelling and optimization of utility systems. Chemical Engineering Research and Design 82:561-578 (2004). https://doi.org/10.1205/026387604323142603
15. Dardor D, Flórez-Orrego D, Terrier C, Ribeiro Domingos ME, Platteau C, da Silva JC, Lopez M, Maréchal F. ROSMOSE: a web-based decision support tool for the design and optimization of industrial and urban energy systems. Energy 304:132182 (2024). https://doi.org/10.1016/j.energy.2024.132182
16. Papoulias SA, Grossmann IE. A structural optimization approach in process synthesis—i. Computers & Chemical Engineering 7:695-706 (1983). https://doi.org/10.1016/0098-1354(83)85022-4
17. Kantor I, Robineau JL, Bütün H, Maréchal F. A mixed-integer linear programming formulation for optimizing multi-scale material and energy integration. Front. Energy Res. 8: (2020). https://doi.org/10.3389/fenrg.2020.00049
18. Liew PY, Walmsley TG. Heat pump integration for total site waste heat recovery. Chem. Eng. Trans. 52:817-822 (2016) https://doi.org/10.3303/CET1652137
19. Miah JH, Griffiths A, McNeill R, Poonaji I, Martin R, Leiser A, Morse S, Yang A, Sadhukhan J. Maximising the recovery of low grade heat: an integrated heat integration framework incorporating heat pump intervention for simple and complex factories. Applied Energy 160:172-184 (2015). https://doi.org/10.1016/j.apenergy.2015.09.032
20. Klute S, Budt M, van Beek M, Doetsch C. Steam generating heat pumps – overview, classification, economics, and basic modeling principles. Energy Conversion and Management 299:117882 (2024). https://doi.org/10.1016/j.enconman.2023.117882
21. Vieren E, Demeester T, Beyne W, Magni C, Abedini H, Arpagaus C, Bertsch S, Arteconi A, De Paepe M, Lecompte S. The potential of vapor compression heat pumps supplying process heat between 100 and 200 °C in the chemical industry. Energies 16:6473 (2023). https://doi.org/10.3390/en16186473
22. Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual. (2025) https://www.gurobi.com

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Investigating the Effects of Heat Ingress and Tank Motion on the Ullage Space of a Partially Filled Liquid Hydrogen Tank Using CFD

Anna Pakarinen[a]*, and Anders Brink[a]

[a] Åbo Akademi University, Faculty of Science and Engineering, Turku, Finland
* Corresponding Author: anna.pakarinen@abo.fi.

ABSTRACT

Cryogenic fuel tanks used in ships are continuously subjected to heat ingress and motions which affect the thermal behavior of the fluid inside the tanks. In this study, the ullage space of a liquid hydrogen (LH2) tank subjected to heat ingress and periodic rolling motion is analyzed using Computational Fluid Dynamics (CFD). A two-dimensional transient model utilizing the dynamic mesh approach is created to represent the ullage space of a partially filled LH2 tank. Three cases are studied where the properties of the thermal insulation of the tank model are varied, resulting in a different heat ingress for each case. A low-frequency motion is applied to the model domain, which induces mixing of temperature layers and cooling due to vapor contact with wetted walls. After 60 s of tank motion, most mixing is observed in the case with the smallest heat ingress, whereas in the cases with larger heat ingress and, consequently, larger thermal and density gradients, separation into a warmer and colder vapor region can be observed. The model developed in this work can be used to gain insight into how liquid hydrogen tanks used on ships are affected by heat ingress and ship motions.

Keywords: Computational Fluid Dynamics, Hydrogen, Heat Ingress, Tank Motion, Cryogenic Fuel Tank

INTRODUCTION

Hydrogen is considered a potential fuel in the shipping industry to mitigate greenhouse gas emissions from the propulsion system. Like natural gas, hydrogen can be liquefied to increase its energy density and stored in cryogenic tanks; this requires cooling down the hydrogen to temperatures around 20 K. Even with efficient thermal insulation, LH2 tanks are affected by heat ingress leading to boil-off and thermal stratification [1, 2]. Ship motions, on the other hand, can cause mixing and destratification of the fluid, which affects e.g. the tank pressure. Understanding fluid behavior under such conditions is crucial for safe operation and design of the fuel tanks.

Various studies have been conducted on fluid behavior in cryogenic tanks. Kang et al. [3] used Finite Element Analysis to investigate thermal stratification in liquid cryogen tanks and the effects of thermal aspect ratio. S B et al. [4] conducted experiments on thermal stratification and pressurization in a partially filled liquid nitrogen (LN2) tank and developed a CFD model to predict fluid behavior in the tank. Jeon et al. [5] performed experiments on the thermal behavior and generation of boil-off gas in a Type C LN2 tank and conducted numerical simulations using CFD.

Furthermore, Jeon et al. [6] conducted experiments on the Type C LN2 tank subjected to rolling motions and presented a multiphase CFD model using the volume of fluid (VOF) method and a phase change model to predict the effects of liquid sloshing in the tank, while Liu et al. [7] implemented the VOF method to investigate the effects of intermittent external excitations on the thermal performance in a liquid oxygen tank.

Most of these studies implemented two-phase models and either simulated the thermal stratification phenomenon in stationary tanks or the response to motions in tanks with initially uniform fluid temperatures. However, not many studies have been conducted on the effects of tank motions on an initially thermally stratified ullage space of a cryogenic fuel tank through

computational methods.

In this study, CFD modelling with dynamic mesh is used to simulate the fluid behavior in the ullage space of an LH2 tank subjected to external heat input and rolling motion. The vapor phase is modelled exclusively for detailed prediction of thermal behavior in the ullage while limiting the required mesh size and avoiding additional computational costs from including the liquid phase. The resulting model is used to demonstrate how mixing affects the thermal behavior of the vapor phase in an LH2 tank with different heat ingress and temperature gradients. Furthermore, the modelling approach in the present work is validated against experimental cases and datasets sourced from previous studies. Thermal stratification in the ullage is validated against experiments in a stationary LN2 tank [8], while the effect of tank motion is validated against experiments performed in a pressurized hot water-water vapor tank [9].

METHODOLOGY

Model LH2 Tank Concept and System Boundary Conditions

A Type C cryogenic storage tank with double-walled vacuum insulation filled with perlite powder is used as the tank concept and visualized in Figure 1. The design choices made here are assumptions and estimates based on literature findings for LH2 tanks [1, 2, 10].

The tank has an inner length and diameter of 1500 mm and 500 mm, respectively. Stainless steel of the grade 316L is used as wall material for both the inner and outer wall, which both have a thickness of 3.5 mm, while the thickness of the annular space between the walls is 100 mm. The liquid filling level of the tank is 50%, with the pressure inside the tank at 4 bar and liquid temperature as well as initial gas temperature at 26 K, while the temperature of the outer wall is set at 288 K. In the CFD simulations, the tank is modelled using a two-dimensional representation.

In total, three cases are considered in this work, where the thermal boundary conditions are varied for each case by adjusting the thermal conductivity of the insulation. Ratnakar et al. [10] investigated the effective thermal conductivity of various insulation materials for cryogenic storage tanks and considered different vacuum levels in their study. The conductivity values used for the tank insulation in this work are taken within the ranges presented in the study for perlite powder vacuum insulation. Here, the thermal conductivity of the insulation is assumed to be varied by adjusting the vacuum level in the annular space.

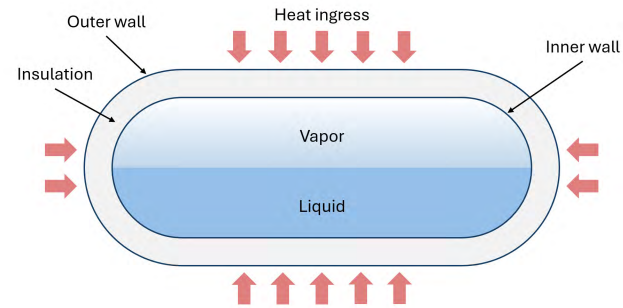

Figure 1. Design concept of the model LH2 tank.

The thermal conductivity of the walls and insulation, as well as the total thermal conductivity for the entire wall structure in each case are summarized in Table 1. The heat transfer rates into the liquid zone and the boil-off rates are given in Table 2. These are calculated based on the heat conduction through the walls to the liquid zone with inner wall temperature of 26 K.

Table 1: Thermal conductivity of the inner and outer walls (k_w), insulation (k_{ins}) and the total thermal conductivity (k_{tot}).

	Case 1	**Case 2**	**Case 3**
k_w [W/(m²•K)]	15	15	15
k_{ins} [W/(m²•K)]	0.0005	0.001	0.002
k_{tot} [W/(m²•K)]	0.000535	0.00107	0.00214

Table 2: Heat transfer rate into liquid (Q_l) and boil-off rate at the liquid surface (m_{evap}).

	Case 1	**Case 2**	**Case 3**
Q_l [W]	2.3	4.7	9.4
m_{evap} [kg/s]	$5.2 \cdot 10^{-6}$	$1.0 \cdot 10^{-5}$	$2.1 \cdot 10^{-5}$

For the tank motion simulations, a periodic rolling motion is applied to the model domain to simulate the motions of a moving ship. In each case, the amplitude is 10° and the period is 12 s. The liquid surface remains flat throughout the motion. This can be seen as an approximate representation of a tank with floating baffles to mitigate sloshing. Furthermore, the tank motion causes wetting of wall surfaces, where wetted surfaces are assumed to have the same temperature as the gas-liquid interface (26 K). Figure 2 shows the ullage at different stages of the tank motion and the direction of gravity perpendicular to the liquid surface. This approach for modelling the ullage and gravity during tank motion allows for simpler implementation of the dynamic mesh while maintaining realistic domain geometry and gravity conditions.

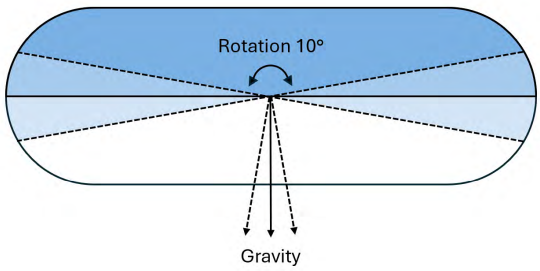

Figure 2. Ullage position at different points of the tank motion.

Numerical Model

CFD analyses of the three cases are conducted using the Ansys Fluent 2024 R1 software. A two-dimensional transient model with dynamic meshing is created to represent the vapor phase inside the tank. The simulations are performed in two parts: Part 1: the formation of the thermal gradient in the ullage due to heat ingress and natural convection, and Part 2: the tank motion and its effect on the thermal gradient is investigated.

Part 1: Thermal Gradient

The fluid domain consists of hydrogen gas surrounded by the tank walls and the liquid surface at the bottom. For the first part of the simulations, the domain is not moving, and the liquid surface is modelled as a wall with a mass source term to represent boil-off. A pressure outlet of the size 10 mm is set at the middle of the top wall to maintain the tank pressure close to 4 bar. The boundary conditions for the walls and liquid surface for each case are set as described in the previous section for computing the thermal gradient.

In the first part, the thermal gradient in the ullage for the static tank is computed. This is conducted by first calculating the thermal gradients for steady state conditions for a low Rayleigh value by reducing gravity and, subsequently, switching to a transient model and incrementally increasing gravity to normal conditions. This approach is to prevent the solution from diverging. The simulations are run until the temperature gradients become stable, using a time step of 0.01 s.

Part 2: Tank Motions

Results from the first part of the simulations are then used as the initial conditions for the tank motion simulations. The dynamic mesh method with user-defined functions for the mesh motion is applied to represent the motion and change in shape of the vapor domain over time, see Figure 2. The direction of gravity is also set to follow the motion so that it is always perpendicular to the liquid surface. For the tank motion simulations, the domain is closed. The mass source term for the liquid surface is set to zero, which means that the vapor mass will remain constant. Furthermore, the wetted parts of the walls are set at the same temperature as the liquid surface. The time step used here is 0.005 s.

Solver Settings

The effects of heat ingress and buoyancy are considered by enabling the energy equation and using the ideal gas law for fluid density to solve the thermal gradients. The governing equations used in the model are continuity, momentum and energy conservation. Species transport is also used to account for the mass source term at the liquid surface, and the fluid flow is modelled using the k-ω SST turbulence model with near-wall treatment. Furthermore, the pressure-based transient solver with PRESTO! pressure discretization scheme and PISO algorithm are used for the solver settings. The convergence conditions for the residuals are set as 10^{-6} for energy and 10^{-5} for continuity and transport equations.

METHOD VALIDATION

Before modelling the LH2 cases, the model was validated against experimental data from the literature. Firstly, the model for developing the thermal gradient in the static ullage was validated using experiments reported by Singh et al. [8], where they measured the vertical temperature gradient inside an LN2 tank at atmospheric conditions with liquid filling levels of 57%, 69% and 77%, with heat ingress of 0.27 W, 0.26 W and 0.29 W into the vapor phase, respectively, as well as a heat ingress of 2.03 W into the liquid for all cases. The tank used in the study consisted of an inner and outer vessel made of aluminum sheets with vacuum multi-layer insulation in between. The inner height of the tank was 360 mm, and the diameter was 300 mm [8]. In the present work, two-dimensional meshes with a mesh size of 5 mm were used to model the vapor phase of each validation case. In Figure 3, results from the simulations are compared to the experimental data from the study by Singh et al. [8].

Secondly, the tank motion model was applied to one of the experiments conducted in the study by Grotle and Æsøy [9]. In their study, they measured the thermal response in a pressurized hot water tank subjected to rolling motions using an amplitude of 3° and various frequencies. The tank used in their study was an insulated cylindrical pressure tank with spherical tank heads and supplied with a heating element. The length of the cylindrical part was 0.89 m, the diameter was 0.35 m, and the tank had a liquid filling level of 50%. Before the tank motion tests, the liquid was heated to 393 K, most of the air was removed and the tank pressure was increased to 4.0254 bar [9]. The tank motion model with the dynamic mesh used in the present work was applied to the hot water tank experiment case with a frequency of 0.13 s^{-1}, where the liquid surface remains flat [9]. A two-dimensional mesh with a mesh size of 5 mm was used to model

the vapor phase of the hot water tank case. In Figure 4, the simulation results using the tank motion model are compared to the experimental data of the pressure reduction in the hot water tank.

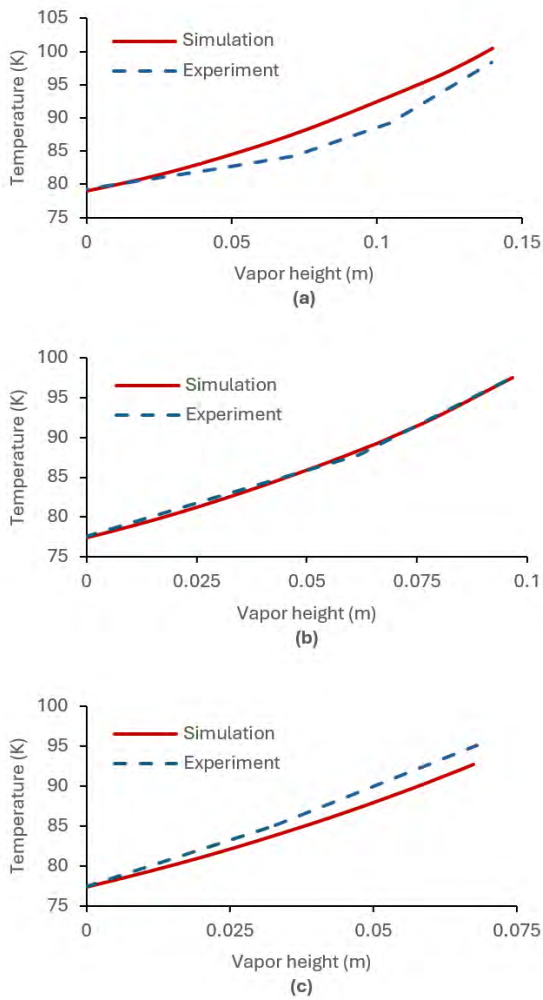

Figure 3. Simulation results for the vertical thermal gradient validated against experimental data for liquid filling levels of (a) 57%, (B) 69% and (c) 77% [8].

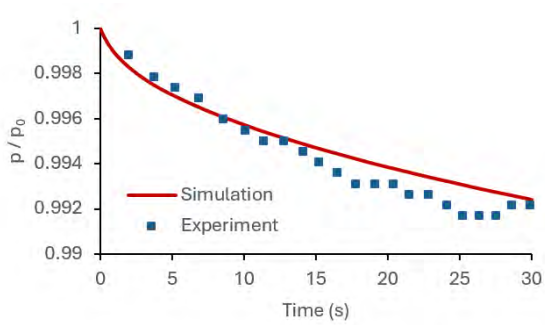

Figure 4. Simulation results for the tank motion model validated against experimental data [9]. Results are for tank pressure over time.

RESULTS FOR LH2 TANK CASES

Mesh Sensitivity Analysis

A mesh sensitivity analysis was conducted to determine the required mesh size for the LH2 tank ullage model. Three different mesh sizes were compared against each other: cell sizes of 2.5 mm, 5 mm and 10 mm with corresponding minimum cell sizes of 0.5 mm, 1 mm and 2 mm around the pressure outlet. The heat transfer rates at the tank walls and the vertical temperature gradients for the different mesh sizes were investigated for Case 2, see Table 1 and 2. The difference between the heat transfer rates was less than 1% between all cases and Figure 5 reveals that the temperature gradients were almost identical. A mesh size of 5 mm is chosen for this study to balance accuracy and computational costs.

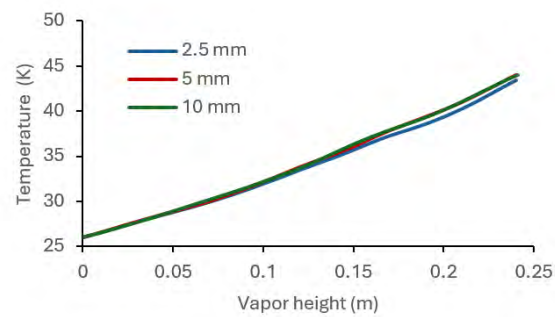

Figure 5. Vertical temperature gradient in LH2 tank ullage computed for mesh sizes 2.5 mm, 5 mm and 10 mm.

LH2 Tank Ullage Simulations

The initial temperature gradients for Case 1, 2 and 3 at stationary conditions were established using the method described earlier. In each case, the transient run was 120 s in total, at which point the gradients had reached a steady state. The temperature contours of the ullage for Case 1, 2 and 3 with the initial thermal gradients are shown in Figure 6(a1), (b1) and (c1), respectively. The time is marked as 0 s to indicate the initial state for the tank motion runs. Furthermore, the corresponding vertical thermal gradients for each case are shown in Figure 7(a). The sizes of the thermal gradients are 10.9 K, 18.7 K and 29.6 K for Case 1, 2 and 3, respectively. These are dependent on the total thermal conductivity values of the tank walls and the associated heat ingress into the tank.

Once the initial thermal gradients were computed, the dynamic mesh method was applied and the tank motion cases were computed for a total run time of 60 s, or five complete oscillations of 12 s periods. The temperature distribution in the ullage for each case at the start and the end of the moving tank runs are shown in Figure 6(a2), (b2) and (c2), and the corresponding vertical temperature gradients are shown in Figure 7(b).

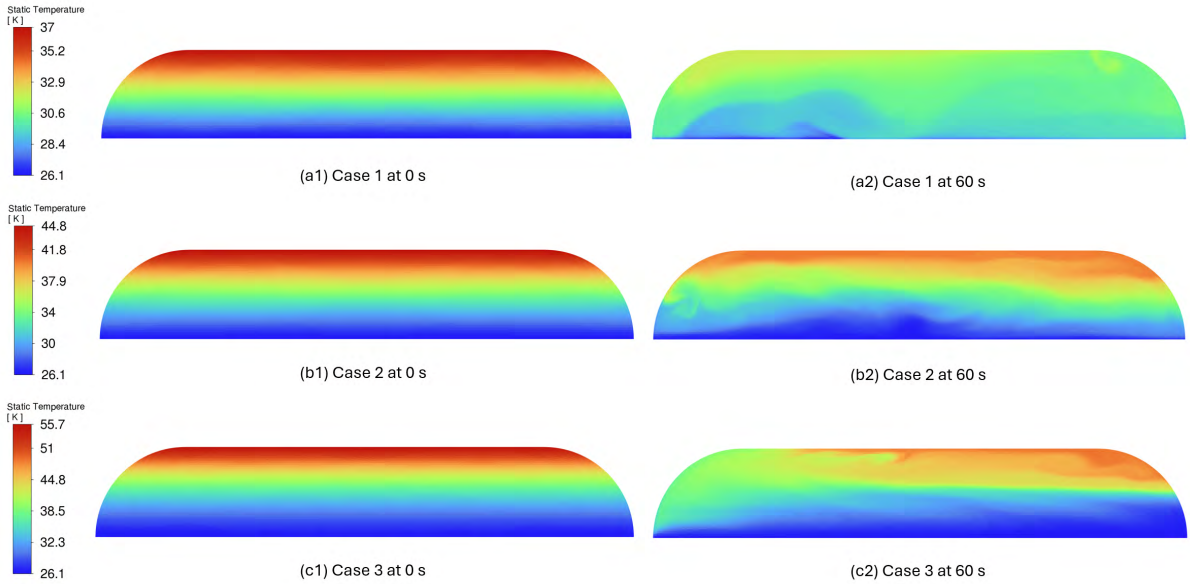

Figure 6. Temperature profile of the ullage at stationary conditions ((a1), (b1) and (c1)), and after 60 s of tank motion ((a2), (b2) and (c2)).

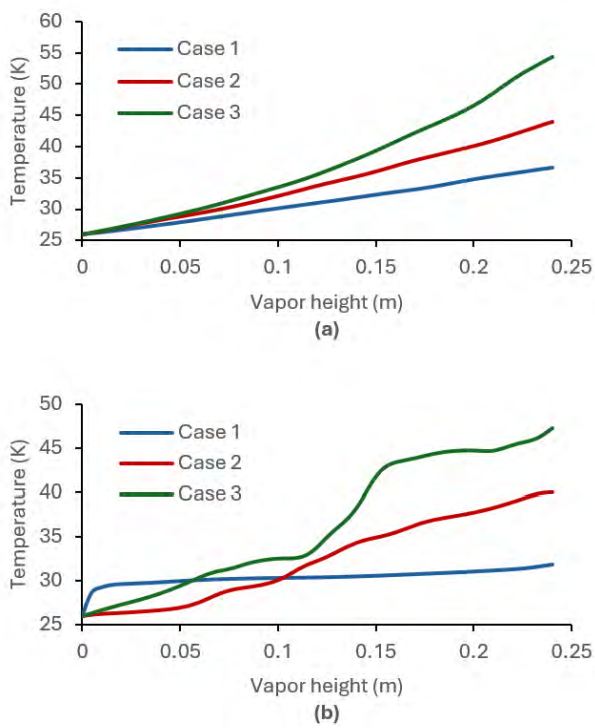

Figure 7. Vertical thermal gradient in the ullage for Case 1, 2, and 3 with (a) initial thermal stratification and (b) after 60 s of tank motion.

After 60 s of tank motion, there is a significant reduction in overall temperature and some mixing of temperature layers in all cases. Most mixing can be observed for Case 1, which also has the smallest initial temperature gradient. This is expected behavior, as a smaller temperature difference also means a smaller difference in density between fluid layers, which enhances mixing. For Case 3, on the other hand, the thermal gradient can be seen separating into two fluid layers around the middle of the domain: the denser cold zone at the bottom and the warmer zone with lower density on top. Case 2 also shows separation between fluid layers, although not as clearly as Case 3 and it still has moderate mixing between cold and warm zones.

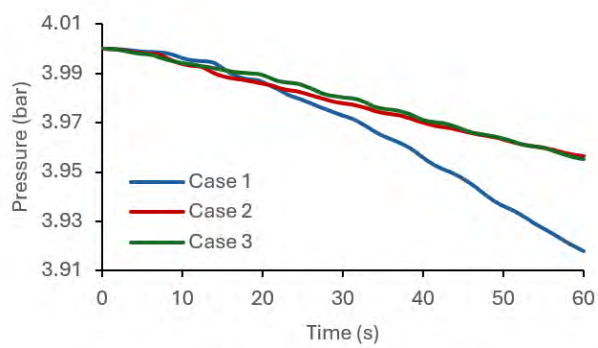

Figure 8. Tank pressure over time for Case 1, 2, and 3.

Furthermore, the vapor cooling also causes a reduction in the domain pressure, as shown in Figure 8. Case 1 has a significantly larger pressure reduction compared to Case 2 and 3, which have a similar pressure evolution to each other over time. The mixing and destratification of fluid layers in Case 1 allows warmer fluid to come into

contact with wetted walls and the liquid surface, resulting in more efficient heat transfer from the vapor to the liquid, and consequently, a larger pressure reduction. Both in Case 2 and 3, the fluid separation into two density layers hinders the mixing of the fluid and the subsequent heat transfer from vapor to liquid, slowing down the pressure reduction rate.

CONCLUSIONS

A numerical investigation on the effects of heat ingress and ship motions on the thermal behavior in the ullage of an LH2 tank was conducted using CFD modelling with dynamic meshing. Three cases were studied, where the thermal conductivity value of the tank insulation was varied and the subsequent response to the thermal gradient and destratification effect due to rolling motion was analyzed. Varying the thermal conductivity value of the insulation affected the size of the temperature gradient. It was observed that for the case with the smallest temperature gradient, the tank motion resulted in destratification and an almost uniform temperature profile in the ullage within five complete oscillations. In the other two cases, the larger difference in density between fluid layers caused separation into cold and warm fluid regions. The results of this work demonstrate the effects of heat ingress and ship motion on the ullage space of an LH2 fuel tank by varying the properties of the thermal insulation. The developed model can be used for further investigations of similar effects in cryogenic tanks.

ACKNOWLEDGEMENTS

This work was supported by the Flexible Clean Propulsion Technologies project (ref. 8324/31/2023), which is co-funded by Business Finland.

AUTHOR IDENTIFIERS

Author ORCIDs:
Pakarinen A: 0009-0005-5825-0066
Brink A: 0000-0002-3236-9692

REFERENCES

1. Ustolin F, Campari A, Taccani R. An extensive review of liquid hydrogen in transportation with focus on the maritime sector. JMSE 10:1222 (2022). https://doi.org/10.3390/jmse10091222
2. Aziz M. Liquid hydrogen: a review on liquefaction, storage, transportation, and safety. Energies 14:5917 (2021). https://doi.org/10.3390/en14185917
3. Kang M, Kim J, You H, Chang D. Experimental investigation of thermal stratification in cryogenic tanks. Experimental Thermal and Fluid Science 96:371-382 (2018). https://doi.org/10.1016/j.expthermflusci.2017.12.017
4. S B V, Bhowmick S, Kuzhiveli BT. Experimental and numerical investigation of stratification and self pressurization in a high pressure liquid nitrogen storage tank. Energy Sources, Part A: Recovery, Utilization, and Environmental Effects 44:2580-2594 (2019). https://doi.org/10.1080/15567036.2019.1651424
5. Jeon GM, Park JC, Kim JW, Lee YB, Kim DS, Kang DE, Lee SB, Lee SW, Ryu MC. Experimental and numerical investigation of change in boil-off gas and thermodynamic characteristics according to filling ratio in a c-type cryogenic liquid fuel tank. Energy 255:124530 (2022). https://doi.org/10.1016/j.energy.2022.124530
6. Jeon GM, Jeong SM, Park JC. Experimental and numerical investigation of the influences of sloshing motion on the change in boil-off gas/boil-off rate in a cryogenic liquid tank. Ocean Engineering 298:117173 (2024). https://doi.org/10.1016/j.oceaneng.2024.117173
7. Liu Z, Cai H, Feng Y, Liu Y, Li Y. Thermodynamic characteristic in a cryogenic storage tank under an intermittent sloshing excitation. International Journal of Hydrogen Energy 45:12082-12094 (2020). https://doi.org/10.1016/j.ijhydene.2020.02.134
8. Singh SK, Sharma A, Atrey MD. Experimental investigation of cryogen storage in isobaric and non-isobaric conditions. Appl Therm Eng 277:127044 (2025) https://doi.org/10.1016/j.applthermaleng.2025.127044
9. Grotle EL, Æsøy V. Dynamic modelling of the thermal response enhanced by sloshing in marine LNG fuel tanks. Applied Thermal Engineering 135:512-520 (2018). https://doi.org/10.1016/j.applthermaleng.2018.02.086
10. Ratnakar RR, Sun Z, Balakotaiah V. Effective thermal conductivity of insulation materials for cryogenic LH2 storage tanks: a review. International Journal of Hydrogen Energy 48:7770-7793 (2023). https://doi.org/10.1016/j.ijhydene.2022.11.130

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Development of a methodology for heat pump-based heat integration in batch processes

Johannes Wloch[a]*, Marcus Grünewald[a] and Julia Riese[b]

[a] Ruhr University Bochum, Laboratory of Fluid Separations, Bochum, Germany
[b] Paderborn University, Fluid Process Engineering, Paderborn, Germany
* Corresponding Author: wloch@fluidvt.rub.de.

ABSTRACT

Heat pumps offer the possibility of reducing CO_2-emissions in the chemical industry. However, the integration of heat pumps, especially in non-continuous processes, faces several challenges. Energy storage facilitates a way to enhance heat integration by providing a continuous supply of heat flows. By doing so, the question arises as to whether this implementation should be applied to the process or to the utility level. At the process level, there is usually more freedom, as one is not bound by the existing temperature levels of the utility system, which are mostly difficult to retrofit. Therefore, this study presents an approach that generates heat integration concepts at the process level based on two different criteria. These criteria influence which process streams are grouped for a storage implementation and therefore influence the heat integration. The aim is to maintain the heat flows as continuous as possible by integrated heat storages. Finally, the possible heat integration concept is evaluated in terms of energy efficiency by a know method for continuous process streams, here the pinch analysis.

Keywords: Energy Systems, Batch Systems, Energy Storage, Optimization, Energy Efficiency, Heat Integration

INTRODUCTION

The CO_2-emissions of the chemical industry must be reduced significantly to meet Europe's climate goals of zero net-emissions by 2045 [1]. A large proportion of the heat demand in process industry is still met by burning fossil fuels such as natural gas which results in significant amount of CO_2-emissions to the atmosphere [2]. Therefore, the supply of process heat from fossil fuels should be reduced. This can be achieved by improving the efficiency of chemical processes or using alternative, low-carbon sources of heat supply. Compression heat pumps offer a possibility of providing heat with low CO_2-emissions when powered by electricity with a low CO_2-emission factor [3]. The use of formerly unused waste heat which is lifted to a higher temperature level by the heat pump can substitute process heat provided by fossil fuels.

Despite the advantages of heat pumps for providing process heat with low CO_2-emissions, heat pump integration into chemical processes is still limited by major challenges. Especially non-continuous processes with fluctuating waste heat supply and process heat demand add an additional challenge for heat pump integration [4]. Furthermore, a heat pump can be integrated at either utility or process level. A heat pump integration at utility level can have a negative effect on the heat pump due to a higher temperature rise, resulting in a lower coefficient of performance (COP). The reason for that is, that the temperature lift depends on the temperatures of the utilities. An adjustment of the temperatures of the utilities is only possible to a limited extend, because in a chemical site other plants would be influenced by this adjustment as well. Therefore, a heat pump integration at process level can be reasonable. Because of the fluctuating nature of the heat flows in non-continuous processes, energy storages offer the possibility of reducing fluctuations and they should be considered as additional possible heat integration concepts. The integration of energy storages would increase the operation time of heat pumps as well [5].

Within the scope of this study, the energetic influence of the number and size of thermal energy storages is analyzed, particularly under the assumption that heat flows are not stored in an ideal manner from an energetic standpoint. Although energetically optimal storage, as

applied in the time average model (TAM), for example, can serve as a reference or benchmark, it may result in heat flows with similar temperature levels being stored in separate energy storages. The question arises to what extent the number of storage units can be reduced and according to which principles heat flows should be aggregated.

To address this issue, the present study provides a concise overview of various methods for heat integration in non-continuous processes with fluctuating waste heat supply at the process level. Furthermore, a systematic procedure is introduced from which different heat integration concepts can be derived based on two distinct criteria that specify which process streams are stored together. The influence of these criteria-particularly with respect to the aggregation and storage of heat flows is analyzed. Finally, one selected variant of these heat integration concepts is evaluated in terms of its energy efficiency.

LITERATURE REVIEW

Different types of methods were developed to receive heat integration concepts. One the on hand there are graphical methods like the pinch analysis. This is a method, which determines the theoretical maximal internal heat exchange based on the cold and hot process streams [5]. This method assumes that alle streams are continuously available and have a constant start and end temperature. Based on the pinch analysis, the time average model (TAM) was developed for non-continuous processes. The result of the TAM is a maximal threshold of possible internal heat exchange. To use this method, the non-continuous process is overdriven into a continuous representative process, by averaging the duties over the process time. A pinch analysis can now be performed using the average heat flows. The results of the TAM are used to design thermal energy storages conceptually. [6] In addition to TAM, there is also the time slice model (TSM). This is a conceptual method which, in contrast to the TAM approach, attempts to take into account the temporal influence on heat integration. To do this, the batch process is divided into time slices so that continuous behavior occurs in these slices. The external utility demand can now be determined in each interval, by using the pinch analysis for each slice. The sum of the external utility over the entire observation period also represents a minimum threshold value. [6] Other approaches to determining heat integration concepts for batch processes use mathematical optimization. A hybrid approach, which is a combination of a graphical and an optimization part takes time slices into account and also considers the integration of a heat pump. A grand composite curve is generated for each time slice. Based on this, the heat pump is dimensioned once so that surplus heat can be stored on the source or sink side by means of an energy storage and released again in times of deficit. [7] Of course, recurring patterns for the slices must first be found. At the same time, the complexity of possible heat transfer networks increases with an increased number of slices. Walden et al. also consider mathematical optimization as an approach for heat pump integration. The grand composite curve for each hour is used as input and then a heat pump is integrated with the help of a mathematical optimization problem. The integration is carried out for different heat outputs: Either for averaged values (TAM), for each hour an optimization, or an optimization over the entire period under consideration. Using this optimization approach it was possible to reduce the heat pump size and thus the operating costs. Thermal storages were not taken into account in this work. [8] Prendl et al. use a linearization of the operating pattern of the heat pump. At the same time, storage and the heat exchange network are taken into account so that a minimum TAC can be achieved. The linearization of MINLP problems has the advantage that MILP problems are not dependent on start values and still achieve a global optimum; and MILP problems enable lower computing times [9, 11]. Becker et al. also use a mathematical optimization problem, which can be used to optimally integrate a heat pump [10]. In addition to dimensioning the relevant equipment for heat integration, the extent to which the scheduling of non-continuous processes affects heat integration was also investigated. Optimization problems were solved for this purpose [12, 13]. Fluctuating heat supply occurs not only in the process industry, but also in the residential sector. Renaldi et al. develop a mathematical optimization model, which considers the design and the operation für a heating system within a heat pump and a thermal energy storage [14]. Furthermore, cluster algorithms were also used, which, for example, aim to identify recurring operating conditions. Wickramasinghe et al. use the k-means algorithm to group temperature time series for commercial buildings [15]. When using heat pump systems, operating data was also used to determine representative load and temperature profiles. For this purpose, the patterns were extracted using dynamic time warping [16].

PROBLEM STATEMENT

We consider a case study involving hot and cold process streams for which heat integration is investigated. Heat integration at the utility level is not the first choice due to unfavorable temperature levels, so we have started with heat integration at process level. To this end, a potential study was carried out in advance to identify low-hanging fruits, i.e., the process streams that are interesting for direct heat integration. The influence of simultaneous availability, temperature difference, and

the ratio of heating supply to heating demand were taken into account in the potential study. The process streams identified by the preliminary potential study as suitable for direct heat integration are not considered in the following problem formulation, thereby reducing the number of process streams and slightly simplifying the problem. The remaining process flows are now to be integrated at the process level. Because there are many possible combinations, an optimization problem that considers heat pumps and storage systems simultaneously was ruled out.

METHODOLOGY

The method aims to determine the impact of the number of storage units on the heat integration concept. The purpose of the storage units is to achieve steady-state conditions, thereby facilitating subsequent, conventional heat integration. For example, heat integration for steady-state conditions can be carried out using well-known analysis methods such as pinch analysis. The question arises which streams should be stored together. To address this question, the method introduces two criteria, each grouping the process streams according to a different objective. These two criteria each have an influence on the grouping of process streams and thus also on the size of the thermal energy storage. The structure of this method is shown in Figure 1. The method starts with a pre-analysis to check the process data for consistency and group them into hot and cold streams. Afterwards it is possible to choose between two different criteria according to which the hot and cold streams are grouped. An energy storage is designed for each of these groups so that stationary conditions are achieved and a following-up pinch analysis is possible. The following section discusses the criteria and storage dimensioning in more detail. The entire methodology is implemented in Python.

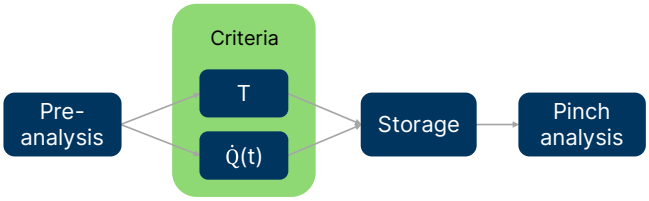

Figure 1. Structure of the method.

Criteria 1: Temperature

This first of the two criteria follows the approach of clustering process streams into groups so that the process streams within a group have similar temperatures. The k-means cluster algorithm is used for grouping [17]. Data points, in our case temperatures, should be grouped so that they are as similar as possible within a cluster and as different as possible between clusters. To use the algorithm, one needs to specify how many clusters the data will be grouped into. The benefit of k-means is that it is easy to implement. This criterion actively influences the temperature prevailing in the energy storage.

Criteria 2: Temporal heat duty

The second criterion aims to ensure that, if several energy storages are available, the process flows are grouped in such a way that they fluctuate less over the time horizon. This has the advantage that the thermal energy storage can be smaller in size, as it needs to buffer less thermal energy. To determine how the streams are grouped, a small optimization problem is solved. This is a MINLP problem, as the objective function takes variance into account. The objective function minimizes the total variance of all groups over the entire time horizon. At the same time, the constraints must be met, e.g., that each stream may only be assigned to one cluster, and that each cluster should be filled with at least one stream. In the optimization problem, the number of clusters is predefined and varied iteratively within an identical cluster range to maintain comparability with the first criterion. The detailed calculation is provided in the supplementary material. The optimization problem was set up and solved in Gurobi [18].

Storage

Once the groups have been selected, a storage dimension can be determined based on the selected process streams. A two-layer storage is selected as thermal energy storage, as this allows start and end temperatures to be specified and then used for pinch analysis. The temperature level within the layers are determined in such a way that, for example, in the case of hot streams that need to be cooled, the temperature in the top layer corresponds to the minimum temperature to which it must be cooled minus a temperature difference for heat transfer. Then for the energy storage, a temperature difference of 10 K is assumed for the temperature of the second layer. This value naturally influences the required storage capacity and can of course be adjusted. The output duty of the energy storage is calculated from the sum of the total energy flowing into the storage device. For a specified storage medium, in this case water, the respective heat capacity results in a mass balance that is accumulated over time and then set to a start value so that the storage content is never negative. Figure 2 shows the structure of the storage system. The specific calculation can be found in the supplementary material. The fluctuating process flows are shown on the right-hand side and are grouped together for a storage. In this case, these heat flows are cooled by the storage and the heat flow on the left-hand side remains constant.

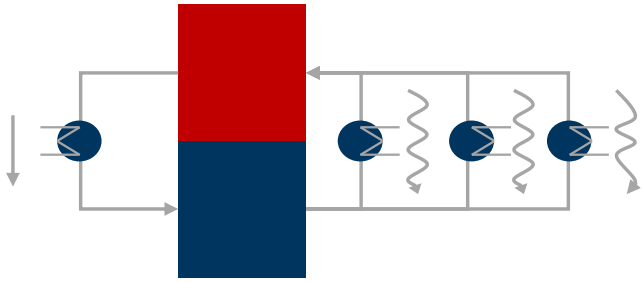

Figure 2. Schematic structure of the energy storage.

Pinch Analysis

The pinch analysis is then performed for a different number of clusters for the hot and cold streams respectively. A minimum temperature difference of 10 K is assumed in each case. The result of the pinch analysis ultimately serves, for example, as the basis for a percentage saving of the total heating requirement. Based on this a final heat integration concept can be derived.

CASE STUDY

The case study involves a batch process comprising 13 cold streams and 23 hot streams, with a temperature range extending from 5 °C to 144 °C. Due to the absence of a periodic operating pattern, the heat duties do not follow a consistent or recurring behavior, thereby complicating heat integration. The dataset used for analysis is provided at 5-minute intervals over a continuous period of 7 days and captures process- and apparatus-level heat duties for each of the 36 streams. This case study originates from real plant data that was anonymized accordingly.

RESULTS

The results of the two different criteria are presented below. The subsequent results, such as storage dimensioning and pinch analysis, are only presented for one criterion. All results refer only to the cold process streams.

Criterion 1:

The first criterion considers the final temperatures to which the streams are heated or cooled. Figure 3 shows the distribution of the different temperatures depending on the number of clusters. The range in which the clusters were varied extends from one cluster to seven different clusters. This is because a maximum of seven different temperatures occur. This does not necessarily mean that there are only seven streams, as streams can also be heated to the same temperature. It is noticeable that the three temperatures around 90 °C are grouped together for most of the clusters tested. The different colors of the temperature data points indicate the cluster and therefore the storage, to which each temperature belongs. The number of clusters corresponds to the number of available storages. For instance, if two clusters are defined, two storages (S1 and S2) are used. In this example, two distinct colors represent the clustering: all streams with temperatures below 100 °C are assigned to one cluster, while those above 100 °C are assigned to another. Consequently, all streams below 100 °C (yellow markers) are grouped together and stored in S2, whereas the remaining streams-requiring temperatures above 100 °C (purple markers)-are grouped into the other cluster and assigned to storage S1. This criterion is used to assign process streams to storage based on the number of available storages, with the objective of minimizing temperature differences among process streams within the same storage. In general, an increasing number of available storages leads to a reduction in temperature variability within each storage.

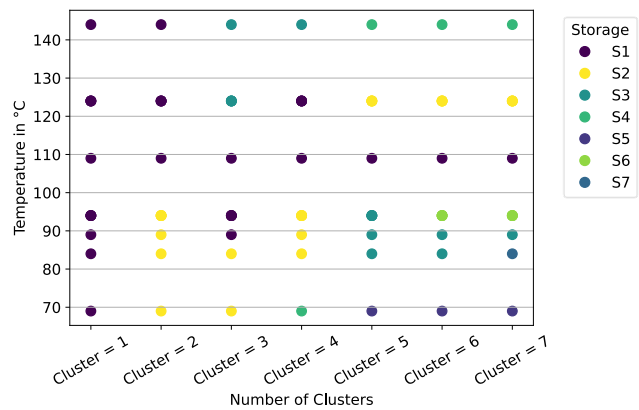

Figure 3. Cold Streams Clusterdistribution for criterion 1.

Criterion 2:

As with criterion 1, the process streams were grouped into a predefined number of clusters. The number of clusters was varied over the same range as for criterion 1. Figure 4 illustrates the composition of a single cluster when the total number of clusters is four. The clusters were chosen such that the variance in the sum of their values was as small as possible. For example, stream W-12 complements the other two streams (W-3 and W-15) very well from approximately hour 80 to hour 130. Nevertheless, there may also be overlaps between the individual streams. The average duty of the streams of this cluster is approximately 380 kW. The presence of overlapping streams arises from the lack of alternative combinations that would result in a lower total variance for the specified number of storages (and corresponding clusters), which is four in this example. The resulting stream distribution is strongly dependent by the individual heat duty profiles.

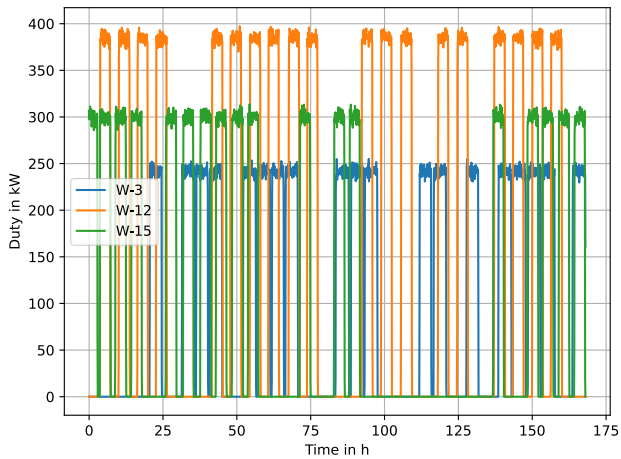

Figure 4. Composition of Cluster 1 for criterion 2 for cold streams.

Comparison of the criteria

As previously mentioned, the cluster criteria have a direct influence on energy storage. Figure 5 shows the influence of the two cluster criteria on the total storage mass and the average maximum temperature difference of the flows belonging to a single thermal storage. It is noticeable that an increase in the number of clusters and therefore also in the number of storages leads to an increase in the total mass of all storages used for both criteria. If only one storage is used, the total storage mass for both criteria is approximately 900 kg. If considering only one cluster, the criterion is irrelevant, as all process flows must be included in one cluster and therefore one storage. If the number of clusters and thus also the number of storage tanks is increased, the total mass of all storages used increases. For criterion 2 (dashed red line), the total mass is always below the solid red line. This is reasonable, as criterion 2 selects the clusters in such a way that the total mass is minimized. The difference between the two red curves also increases with an increase in the number of clusters, as there are now many more possible combinations, meaning that the solutions can deviate more strongly from each other. The average maximum temperature difference of the process flows assigned to each cluster shows a different trend. The temperature difference is determined as follows. For example, for two possible clusters, the maximum temperature difference between the process streams flowing into the storage units is 35 °C and 25 °C. The average of these two temperatures is 30° C. As the number of clusters increases, the maximum temperature difference between the storages decreases. This is reasonable, as more clusters or storage units lead to a finer distribution of the temperature. The solid blue line reaches a temperature difference of zero for seven clusters, as there are seven different temperature levels for the cold process streams, and if the corresponding number of storages are used, only one temperature level flows into each storage. The fact that the two blue curves for one cluster have the same total temperature difference is due to the same reason as for the red curve.

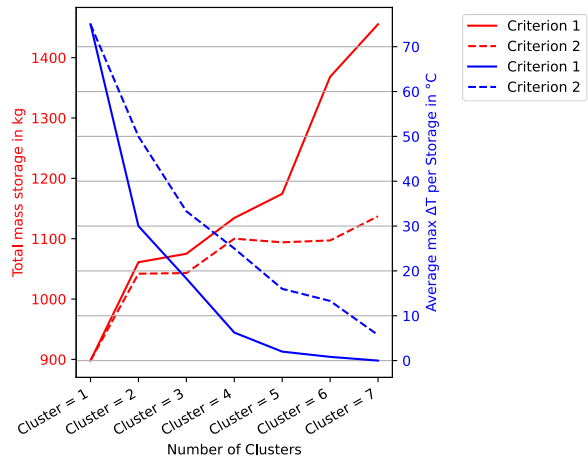

Figure 5. Influence of the number of clusters on the total storage mass and on the total deltaT depending on the cluster criterion.

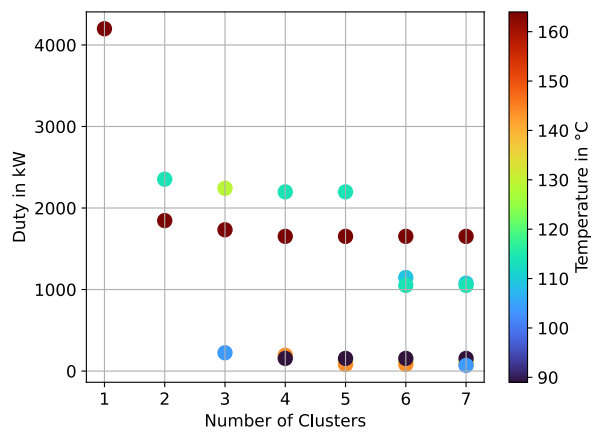

Figure 6. Resulting storage temperatures and duties for cold streams exemplified for criterion 1.

Storage for Criterion 1:

The following section examines the extent to which the number of clusters affects relevant storage sizes and therefore the continuous heat duty. This is illustrated using cluster criterion 1 as an example. Figure 6 shows both the influence on the heat duty, which is then stationary, and on the temperatures inside the energy storage. Since these are storages for heating of cold streams, the temperature of the hotter layer was selected. Basically, the sum of the outgoing duty across the different cluster or storage numbers must be identical. This is because, for all cluster cases, the total energy required to heat the streams over a constant observation period must be

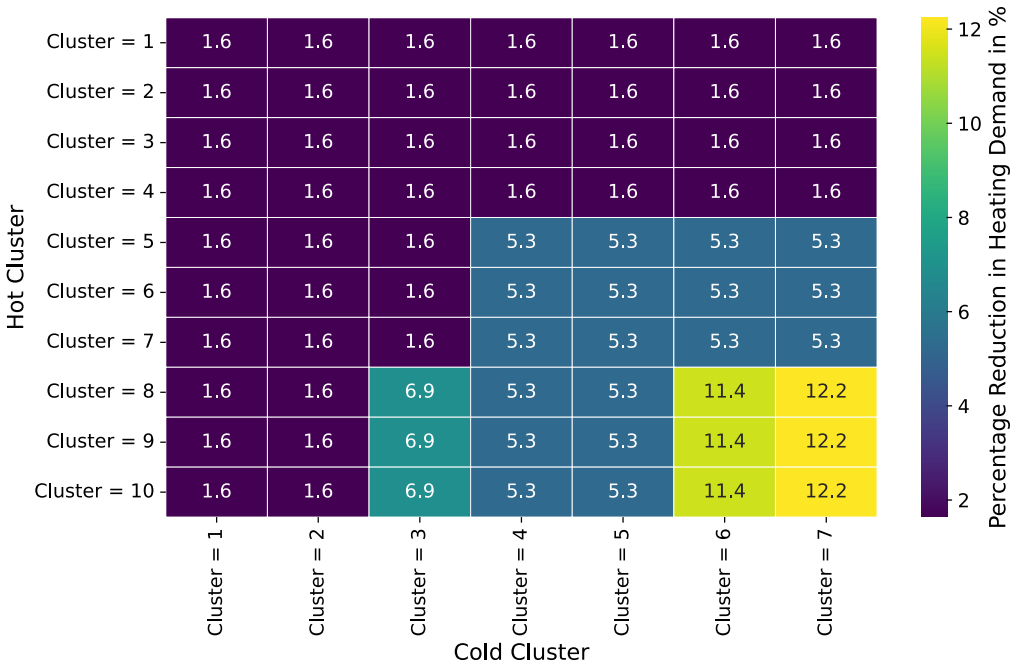

Figure 7. Clustermatrix for cold and hot streams.

provided. In addition, it can be seen that the temperature levels used for heating run towards a limit value based on the number of storage. For example, the 169 °C requires a storage with a heating capacity of over 4000 kW. The heating capacity required for this temperature decreases when the number of clusters or storage is increased, as the process streams can now be distributed more finely in terms of temperature distribution (criterion 1), and thus fewer process streams are heated with this temperature level.

Pinch Analysis for Criterion 1:

Once the relevant storage variables have been determined and the heat flows are in steady-state, a pinch analysis can be performed based on this information. The pinch analysis is performed iteratively for all clusters of hot and cold flows. The results can be seen in Figure 7. This shows the percentage savings in heating duty compared to the case where no heat is integrated. If there are more clusters and thus more storage tanks, greater savings can be achieved. This is because the heating and cooling requirements are at more suitable temperature levels, enabling heat integration without a heat pump. The corresponding grand composite curve for the combination of 9 clusters for the hot process flows and 7 clusters for cold process flows is shown in Figure 8. On the one hand, the pocket shows where heat integration can take place. It is now up to the user to decide whether they also want to integrate storage outside the pocket in order to provide stationary conditions outside of possible heat integration. This can be particularly interesting if, for example, a heat pump can be integrated. In this case, a heat pump could be suitable for the two temperature levels of 40 °C and 120 °C.

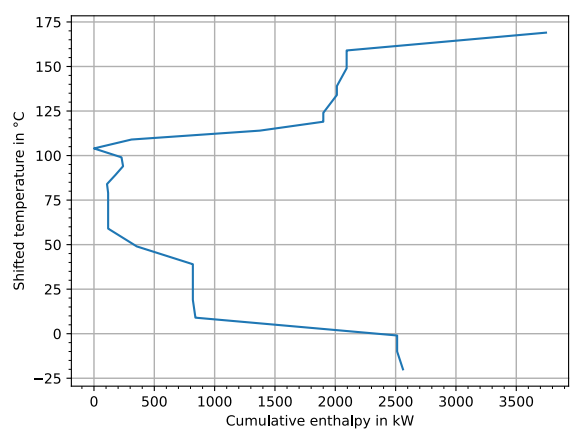

Figure 8. Grand Composite Curve for 9 clusters of hot streams and 7 clusters of cold streams.

CONCLUSION

In this study, we analyzed the influence of the number and size of thermal energy storages, by introducing two different grouping criteria for heat integration of batch process streams at the process level. It was shown that both criteria fulfil the purpose. The first criterion employs k-means clustering to group process streams with

similar temperatures and assign them to the same storage tank. As a result, streams directed to a single storage exhibit comparable temperatures. However, the required storage mass may increase due to mismatches in the heat duty of the grouped process streams. With the second criterion, it was exactly the opposite. To this end, a small optimization problem was solved, which minimizes the variance of the flows belonging to an energy storage. Accordingly, the total mass of all storage tanks was smaller than that of the storage tanks grouped according to temperature. The use of energy storage aimed to compensate for fluctuations in heat flows in batch processes, thereby establishing steady-state conditions that enable the application of established heat integration methods. An energy assessment was carried out using pinch analysis for a storage combination of hot and cold flows. It was shown that internal heat exchange is possible. The remaining heating demand can then be covered by a heat pump. From our point of view, it is important that further possible evaluation criteria are found. Perhaps one that allows a weighting between temperature focus and storage size focus. In addition, it makes sense to identify not just two, but also more layer within an energy storage as integration options. This can be an interesting approach, especially when streams are grouped according to temperature. Finally, it should be added that this is not the optimal heat integration solution. This would involve a larger mathematical optimization problem. Instead, this approach aims to identify promising heat integration concepts at an early stage. These solutions can also be used as starting values for optimization.

DIGITAL SUPPLEMENTARY MATERIAL

The specific calculation of the storage model and the optimization model can be found in the appendix. LAPSE:2026.0028

ACKNOWLEDGEMENTS

This work was funded bei the Federal Ministery for Econimic Affairs and Energy (BMWE) within the project WinPro (project funding reference number 03EN2115A).

AUTHOR IDENTIFIERS

Author ORCIDs:
Wloch J: 0009-0005-8498-5904
Grünewald M: 0000-0002-4007-8270
Riese J: 0000-0002-3053-0534

REFERENCES

1. United Nations, Paris Agreement. United Nations Framework Convention on Climate Change (2015) https://unfccc.int/sites/default/files/english_paris_agreement.pdf)
2. Wolf S. Integration von Wärmepumpen in industrielle Produktionssysteme. Universität Stuttgart (2017) http://dx.doi.org/10.18419/opus-9593
3. Kiss AA, Smith R. Rethinking energy use in distillation processes for a more sustainable chemical industry. Energy 203:117788 (2020). https://doi.org/10.1016/j.energy.2020.117788
4. Agner R, Ong BHY, Stampfli JA, Krummenacher P, Wellig B. A Graphical Method for Combined Heat Pump and Indirect Heat Recovery Integration. Energies (2022) https://www.mdpi.com/1996-1073/15/8/2829#metrics
5. Kemp IC. Key concepts of pinch analysis. Pinch Analysis and Process Integration :15-40 (2007). https://doi.org/10.1016/b978-075068260-2.50007-9
6. Fernández I, Renedo CJ, Pérez SF, Ortiz A, Mañana M. A review: energy recovery in batch processes. Renewable and Sustainable Energy Reviews 16:2260-2277 (2012). https://doi.org/10.1016/j.rser.2012.01.017
7. Stampfli JA, Atkins MJ, Olsen DG, Walmsley MRW, Wellig B. Practical heat pump and storage integration into non-continuous processes: a hybrid approach utilizing insight based and nonlinear programming techniques. Energy 182:236-253 (2019). https://doi.org/10.1016/j.energy.2019.05.218
8. Walden JVM, Wellig B, Stathopoulos P. Heat pump integration in non-continuous industrial processes by dynamic pinch analysis targeting. Applied Energy 352:121933 (2023). https://doi.org/10.1016/j.apenergy.2023.121933
9. Prendl L, Schenzel K, Hofmann R. Simultaneous integration of heat pumps and different thermal energy storages into a tightened multi-period MILP HENS superstructure formulation for industrial applications. Computers & Chemical Engineering 147:107237 (2021). https://doi.org/10.1016/j.compchemeng.2021.107237
10. Becker H, Maréchal F. Energy integration of industrial sites with heat exchange restrictions. Computers & Chemical Engineering 37:104-118 (2012). https://doi.org/10.1016/j.compchemeng.2011.09.014
11. Beck A, Hofmann R. A novel approach for linearization of a MINLP stage-wise superstructure formulation. Computers & Chemical Engineering 112:17-26 (2018).

https://doi.org/10.1016/j.compchemeng.2018.01.010
12. Hegyháti M, Friedler F. Combinatorial algorithms of the s-graph framework for batch scheduling. Ind. Eng. Chem. Res. 50:5169-5174 (2011). https://doi.org/10.1021/ie101800x
13. Seid ER, Majozi T. Heat integration in multipurpose batch plants using a robust scheduling framework. Energy 71:302-320 (2014). https://doi.org/10.1016/j.energy.2014.04.058
14. Renaldi R, Kiprakis A, Friedrich D. An optimisation framework for thermal energy storage integration in a residential heat pump heating system. Applied Energy 186:520-529 (2017). https://doi.org/10.1016/j.apenergy.2016.02.067
15. Wickramasinghe A, Muthukumarana S, Loewen D, Schaubroeck M. Temperature clusters in commercial buildings using k-means and time series clustering. Energy Inform 5: (2022). https://doi.org/10.1186/s42162-022-00186-8
16. Zhang S, Zhang L, Zhang X. Clustering based on dynamic time warping to extract typical daily patterns from long-term operation data of a ground source heat pump system. Energy 249:123767 (2022). https://doi.org/10.1016/j.energy.2022.123767
17. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay É. Scikit-learn: Machine Learning in Python. JMLR (2011).
18. Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual. (2024) https://www.gurobi.com

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Pareto-Optimal Pathways for Refinery Decarbonization: Retrofit of Small Modular Nuclear Reactors

Aditya S Khatu[a†], Sampriti Chattopadhyay[a†], Ana I Torres[a*]

[a] Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh 15213, PA, USA
* Corresponding author: aitorres@cmu.edu
† authors contributed equally

ABSTRACT

Refineries are major sources of direct CO_2 emissions, primarily from steam generation, fluid catalytic cracking, and hydrogen production. This study develops a superstructure optimization framework to evaluate the economic and environmental viability of retrofitting existing refineries with small modular nuclear reactors (SMRs) for cogeneration of heat and electricity. A multi-period mixed-integer quadratically constrained program is formulated, simultaneously minimizing the present cost of retrofitting and CO_2 emissions over the time horizon. This problem is solved to generate a Pareto frontier via the ε-constraint method. Two cases are analyzed for a medium-scale refinery, considering 1) inflexible operation under average annual electricity prices and 2) flexible operation under hourly prices with the possibility of installation of storage devices. Compared to a benchmark without SMRs in the superstructure, allowing their installation leads to reduced costs at lower or comparable emission levels. The results show that SMRs are primarily used for high-pressure steam generation. Flexible operation and the inclusion of thermal energy storage further reduce costs. Overall, SMRs appear in multiple non-dominated solutions, highlighting their potential as a cost-effective refinery decarbonization strategy.

Keywords: Electricity & Electrical Devices, Energy Storage, Nuclear, Process Design, Refining

INTRODUCTION

Reducing emissions from energy-intensive industries is essential for meeting climate targets. Oil refineries are among the largest industrial sources of direct CO_2 emissions, primarily due to fossil-fired steam generation, hydrogen production from steam methane reforming, and large point sources such as fluid catalytic crackers [1]. Electrification and carbon capture can lower emissions, but their effectiveness depends on electricity prices, availability of renewable energy, natural gas prices, and operational flexibility. This has increased interest in exploring other low-carbon sources of heat and power that can be deployed directly at industrial sites.

Small modular nuclear reactors (SMRs) are among the promising options for decarbonizing the chemical industry and refineries. Unlike conventional gigawatt-scale nuclear plants, SMRs typically produce up to about 300 MWe per unit and are designed for modular construction and incremental deployment [2]. Inherent safety features enable simplified plant designs and greater flexibility, particularly for deployment in remote or infrastructure-limited industrial sites. This increased adaptability, along with the ability to cogenerate electricity and steam, makes them especially attractive for chemical plants and refineries, where continuous demand for high-quality utilities exists. Despite these advantages, several challenges remain. High first-of-a-kind capital costs, financing risks, and construction uncertainties pose significant barriers to deployment. Additional challenges include waste-management considerations and other regulatory requirements, which make smaller units expensive to operate. Addressing these issues is essential for realizing the benefits of new technologies [3, 4].

As shown in Fig. 1, the nuclear reactor consists of a reactor vessel, which houses the reactor core, coolant, moderator and control rods. The core contains fuel assemblies where nuclear fission reactions generate heat, surrounded by a moderator (light water or graphite) that slows neutrons to sustain the chain reaction. Control rods regulate the reaction rate by absorbing neutrons. Heat is

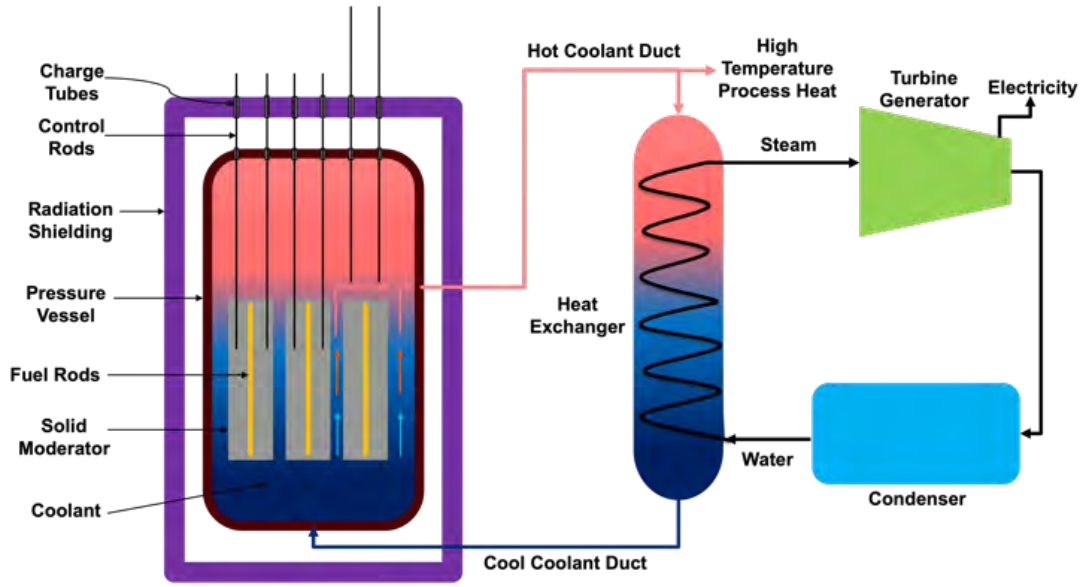

Fig. 1. Schematic showing the structure of a nuclear reactor

removed by the coolant and transferred to steam generators or heat exchangers. Nuclear reactors are classified by coolant and moderator type. Table 1 summarizes the types illustrating the trade-offs between maturity and achievable process temperatures.

Table 1. Achievable outlet temperatures for different small modular reactor types [11]

Reactor type	Coolant/moderator	Outlet Temperature (in centigrade)	Maturity
iPWR	Light water/ Light water	300-330	High
MSR	Molten fluoride or chloride salt/ graphite	600-750	Low
HTGR	Helium/ Graphite	750-950	Low

Existing techno-economic assessments indicate that construction time, financing conditions, and assumptions regarding modularization dominate the uncertainty in levelized cost of electricity estimates for SMRs [5, 6]. Learning-curve analyses, such as the review by Mignacca et al. [3], suggest that if modularization is achieved at scale, each doubling of installed capacity could lead to cost reductions of approximately 5–10% through "economies of number, " arising from the deployment of many identical units rather than traditional economies of scale. Reactor-level optimization studies investigate design trade-offs related to core configuration and power conversion systems [7, 8]. System-level studies examining SMRs in power system applications and industrial applications highlight the value of cogeneration and flexible operation [9–11].

This paper extends our previous studies [12] by evaluating the potential for deploying small modular reactors to achieve economically viable decarbonization of refineries. The superstructure, shown in Fig. 2, includes both conventional and low-carbon pathways for hydrogen production and steam generation. SMRs enable additional decarbonized pathways for high-pressure steam (HPS) and hydrogen (H_2) generation through the co-production of low-carbon electricity and heat. SMR-generated electricity can be used to power proton exchange membrane and alkaline electrolyzers for hydrogen production and electric boilers for steam generation, while reactor heat can be directly utilized to generate HPS. Other pathways for decarbonizing hydrogen production include steam methane reforming with pre- or post-combustion carbon capture and electrolysis using electricity purchased from the grid. Additional pathways for high-pressure steam (HPS) generation include natural gas boilers with post-combustion carbon capture and electric boilers powered by grid electricity. Three storage options are considered: electricity storage in lithium-ion batteries, hydrogen storage in salt caverns, and SMR-generated heat stored in molten nitrate salt, which is referred to as thermal energy storage (TES) in the rest of the paper.

MATHEMATICAL MODELING

Small Modular Reactor

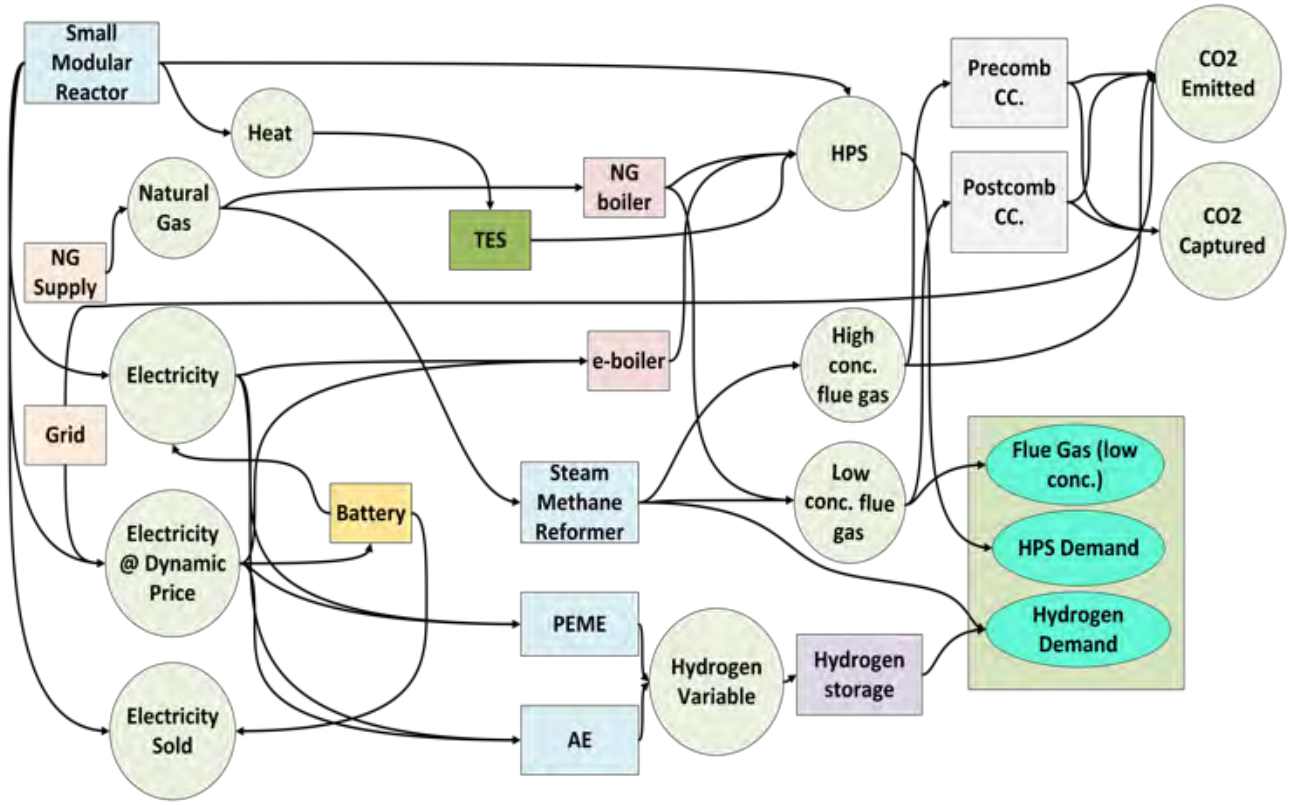

Fig. 2: Superstructure considered in this study

We begin by describing the constraints associated with design variables, including decisions on module selection and construction timelines. The stated equations are valid $g \in G, n \in N, yr \in YR$ (where G is the set of SMR reactor technology types considered (e.g., $\mathrm{G} = (\mathrm{iPWR}, \mathrm{HTGR}), N$ is the set of candidate module counts (i.e., number of SMR modules that can be installed), $N = \{1, 2, 3, 4, 5, 6\}$, and YR is the set of planning years over the optimization horizon, $YR = \{1, 2, \ldots, 19\}\}$), unless specified otherwise. Equation (1) ensures that no more than one reactor type *g* is installed in any given year, *yr*, as determined by the binary variable $s_{g,yr}$. As shown in equation (2), the module construction decisions $m_{g,n,yr}$ depend on this selection; a module of type g can be constructed in a year yr when the corresponding reactor type is selected. The availability indicator $a_{g,n,yr}$ indicates whether a given module is already constructed, evaluated as a cumulative sum of $m_{g,n,yr}$ over time, as shown in equation (3), each module becomes available in the year it is built and remains available thereafter. According to equation (4), the installed capacity in the year $yr, Q_{g,n,yr}$, is given by the sum of the prior year's capacity $Q_{g,n,yr-1}$and the expansion $QE_{g,n,yr}$undertaken during that year. Annual expansion $QE_{g,yr}$ is bounded between technology-specific lower and upper limits, conditional on the module construction decision binary $m_{g,n,yr}$, as shown in equation (5).

$$\sum_{g \in G} s_{g,yr} \leq 1 \tag{1}$$

$$m_{g,n,yr} \leq s_{g,yr}) \tag{2}$$

$$a_{g,n,yr} = \sum_{tt=1}^{yr} m_{g,n,tt} \tag{3}$$

$$Q_{g,n,yr-1} + QE_{g,n,yr} = Q_{g,n,yr} \tag{4}$$

$$L\, m_{g,n,yr} \leq QE_{g,n,yr} \leq U\, m_{g,n,yr} \tag{5}$$

Now we present the constraints showing operational restrictions. The stated equations are valid $g \in G, n \in N, hr \in HR, d \in D, yr \in YR$(where HR is the set of hourly time steps within a representative day $HR = \{1, 2, \ldots, 24\}$, and D is the set of representative days used to capture intra-year operating variability $D = \{1, 2, 3\}$ for three representative days per year), unless specified otherwise. A module n can be committed in hour hr of representative day d and year yr, denoted by binary $o_{g,n,hr,d,yr}$, only if it is already built by that time, denoted by $a_{g,n,yr}$, as shown in equation (6). Electricity production $e_{g,n,hr,d,yr}$ is limited by the installed nameplate capacity $Q_{g,n,yr}$ of each module, equation (7). The reactor's core output $r_{g,n,hr,d,yr}$ scaled by heat-transfer and turbine efficiencies η_{HT_g} and η_{TE_g}, expressed in electrical units, cannot exceed $Q_{g,n,yr}$ when the unit is committed, as enforced in equation (8). Ramp-

rate limits on the core output, bound hour-to-hour variation in $r_{g,n,hr,d,yr}$ by RR_g/η_{TE_g}, as modeled by equation (9). During implementation this constraint is decomposed into two inequality constraints for each possible value of the absolute function.

$$o_{g,n,hr,d,yr} \leq a_{g,n,yr} \tag{6}$$

$$e_{g,n,hr,d,yr} \leq Q_{g,n,yr}, \tag{7}$$

$$r_{g,n,hr,d,yr}\,\eta_{HT_g}\,\eta_{TE_g} \leq Q_{g,n,yr}\,o_{g,n,hr,d,yr} \tag{8}$$

$$|r_{g,n,hr-1,d,yr} - r_{g,n,hr,d,yr}| \leq \frac{RR_g}{\eta_{TE_g}} \tag{9}$$

Startup and shutdown events are tracked by binaries $\nu_{g,n,hr,d,yr}$ and $w_{g,n,hr,d,yr}$, which relate to changes in commitment status $o_{g,n,hr,d,yr}$ across consecutive hours through equation (10). Equation (11) prevents simultaneous startup and shutdown in the same hour. A module is committed to electricity generation, denoted by binary $oe_{g,n,hr,d,yr}$, only when the module is online, as shown by equation (12). Equation (13) shows that the electricity generated $e_{g,n,hr,d,yr}$ is zero whenever the unit is not committed for generation and is capped by $Q_{g,yr}$ otherwise. Minimum downtime is enforced via Equation (14), which requires the module to remain offline for at least MDT_g hours after a decommitment before a new startup is permitted.

$$\nu_{g,n,hr,d,yr} = o_{g,n,hr,d,yr} - o_{g,n,hr-1,d,yr} + w_{g,n,hr,d,yr} \tag{10}$$

$$\upsilon_{g,n,hr,d,yr} + w_{g,n,hr,d,yr} \leq 1 \tag{11}$$

$$oe_{g,n,hr,d,yr} \leq o_{g,n,hr,d,yr} \tag{12}$$

$$e_{g,n,hr,d,yr} \leq Q_{g,n,yr}\,oe_{g,n,hr}, \tag{13}$$

$$\sum_{i=1}^{MDT_g}\left(1 - o_{g,n,hr-i,d,yr}\right) \geq \nu_{g,n,hr,d,yr}\,MDT_g, \tag{14}$$

Finally, we outline the financial constraints. Operating costs (OPEX) include startup costs, $SC_g\,\nu_{g,n,hr,d,yr}$, variable costs proportional to overall reactor output ($VC_g\,r_{g,n,hr,d,yr}$), and fixed annual costs, $FC_g\,Q_{g,n,yr}$ as shown by Equation (15). The per unit startup and variable cost components are multiplied by the weights of the representative days, defined as the number of days each representative day denotes, and are then summed to determine the total. Capital expenditure (CAPEX) incurred for building SMR modules each year is a function of the annual capacity expansion and specific capital cost $\delta_{g,yr}$, as shown in Equation (16).

$$OPEX_{SMR,yr} = \sum_{d\in D} w_d \sum_{hr\in HR}\sum_{g\in G}\sum_{n\in N}\left(SC_g\,\nu_{g,n,hr,d,yr} + VC_g\,r_{g,n,hr,d,yr}\right) + FC_g\,Q_{g,n,yr}$$

$$\forall yr \in YR \tag{15}$$

$$CAPEX_{SMR,yr} = \sum_{g\in G}\sum_{n\in N} QE_{g,n,yr}\;\delta_{g,yr}\;\forall \mathrm{yr} \in \mathrm{YR} \tag{16}$$

Thermal energy storage (TES)

The TES system is defined for each reactor type g and module n. We begin with the equations relating design variables associated with TES. Capacity expansion in TES, $QE_{TES,g,yr}$ is bounded and activated by a binary installation decision $k_{g,n,yr}$, as shown in equation (17). Equation (18) shows that installed TES capacity is calculated as a function of annual capacity expansions analogous to SMR capacity. Equations (19)–(24) define the operational constraints for TES. The state of charge, representing the amount of heat stored $l_{g,n,hr,d,yr}$, is bounded by the installed thermal energy storage capacity $Q_{\mathrm{TES},g,n,yr}$, as specified in equation (19). The maximum charging $c_{g,n,hr,d,yr}$and discharging $d_{g,n,hr,d,yr}$ rates are limited to a fraction of the total installed storage capacity, where the parameter EP represents the minimum charge/discharge time. The evolution of the state of charge is governed by the energy balance in equation (22), which accounts for net heat inflows and charging and discharging efficiencies. Equations (23) and (24) impose cyclic continuity across representative days. These constraints are valid given the use of a chronology-preserving agglomerative hierarchical clustering (AHC) method, following Tso et al. [13], to derive representative days.

$$L\,k_{g,n,yr} \leq QE_{TES,g,n,yr} \leq U\,k_{g,n,yr} \tag{17}$$

$$Q_{TES,g,n,yr-1} + QE_{TES,g,n,yr} = Q_{TES,g,n,yr} \tag{18}$$

$$l_{g,n,hr,d,yr} \leq Q_{\mathrm{TES},g,n,yr} \tag{19}$$

$$c_{g,n,hr,d,yr} \leq \frac{Q_{\mathrm{TES},g,yr}}{EP} \tag{20}$$

$$d_{g,n,hr,d,yr} \leq \frac{Q_{\mathrm{TES},g,n,yr}}{EP} \tag{21}$$

$$l_{g,n,hr,d,yr} = l_{g,n,hr-1,d,yr} + c_{g,n,hr,d,yr}\,\eta_{\mathrm{Charging}} - \frac{d_{g,n,hr,d,yr}}{\eta_{\mathrm{Discharging}}} \tag{22}$$

$$l_{g,n,1,d,yr} = l_{g,n,24,d,yr} + c_{g,n,24,d,yr}\,\eta_{\mathrm{Charging}} - \frac{d_{g,n,24,d,yr}}{\eta_{\mathrm{Discharging}}} \tag{23}$$

$$l_{g,n,1,d+1,yr} = l_{g,n,24,d,yr} + c_{g,n,24,d,yr}\,\eta_{\mathrm{Charging}} - \frac{d_{g,n,24,d,yr}}{\eta_{\mathrm{Discharging}}} \tag{24}$$

The annual capital expenditure (CAPEX) resulting from the installation of TES is a function of the annual capacity expansion and the capital cost per unit, as shown in Equation (25). For this study, it is assumed that TES incurs no additional operating expenditure.

$$\mathrm{CAPEX}_{\mathrm{TES},yr} = \sum_{g\in G}\sum_{n\in N}\delta_{\mathrm{TES},g,yr}\;QE_{\mathrm{TES},g,n,yr}\;\forall yr \in YR \tag{25}$$

Constraints modelling the rest of the superstructure

We implement an extension of our previous formulation in Karthikeyan et al. [12], while accounting for flexible operations and the possibility of installation of storage units. The constraints keeping track of the installed capacities of each technology j, in each year, yr ($Q_{j,yr}$), after capacity expansion and decommissioning decisions are included from our previous work (Equations (7) – (10) in [12]). These constraints are now also applicable to the battery and H_2 storage units.

The stated operating and physical constraints in the superstructure are applicable to $hr \in HR, d \in D$, and $yr \in YR$, unless otherwise stated. Operating constraints ensure that the hourly mass flow rates of key components (k'_j) and power flow rates ($F_{k'_j,j,hr,d,yr}$/ $P_{k,j,hr,d,yr}$) to unit j remains within the installed capacity limits and above a minimum turndown fraction χ_j of the installed capacity, respectively. We implement an extension to equations (11)-(12) in [12] applicable to mass and power flows at each hour of each representative day.

Constraint (26) ensures that each conversion technology operates above a minimum average annual utilization factor, ϕ_j, preventing under-use of capital-intensive assets.

$$\phi_j Q_{j,yr} \leq \sum_{d\in D} w_d \sum_{hr\in HR} \sum_{k\in In(j)} P_{k,j,hr,d,yr} \; \forall j \in \{\text{e-boiler, PEME, AE}\} \; \forall yr \in YR \quad (26)$$

For hydrogen storage units, the hourly charging and discharging rates are constrained by equation (27), the installed storage capacity, and the charge/discharge duration $\tau_{j,yr}$, which represents the minimum number of hours required to fully charge or discharge the storage at its rated power.

$$F_{H_2 variable\ or\ H_2\ demand, H_2\ storage, hr,d,yr} \leq \frac{Q_{H_2 storage,yr}}{\tau_{H_2 storage,yr}} \quad (27)$$

The hydrogen stored in the system, $stor_{H_2,H_2 storage,hr,d,yr}$, at any time, is limited by the installed storage capacity, as shown in equation (28).

$$stor_{H_2,H_2 storage,hr,d,yr} \leq Q_{H_2 storage,yr} \quad (28)$$

Similarly, for battery storage units, the hourly charging and discharging power are constrained symmetrically by the installed battery capacity and the charge/discharge duration $\tau_{j,yr}$, as shown in equations (29) and (30).

$$P_{elec.or\ elec.at\ dynamic\ price, Battery,hr,d,yr} \leq \frac{Q_{Battery,yr}}{\tau_{Battery,yr}} \quad (29)$$

Equation (30) shows that the state of charge of the battery, $stor_{elec.,Battery,hr,d,yr}$, is bounded by its installed energy capacity.

$$stor_{elec.,Battery,hr,d,yr} \leq Q_{Battery,yr} \quad (30)$$

The proposed formulation also specifies physical balances across the task nodes (technology units) and the resource nodes (components) across the superstructure. The task balances are defined by hourly extensions to equations (1)-(2) and (4)-(6) in [12], which relate the inlet and outlet component mass/energy flow rates, ($F_{k,j,hr,d,yr}$/$P_{k,j,hr,d,yr}$) to the flow of a key component $F_{k'_j,j,hr,d,yr}$ using simplified yield-based models for each hour of each representative day.

The total hydrogen and high-pressure steam (HPS) produced across all pathways and aggregated at each hour of each representative day must satisfy the refinery's overall hydrogen and HPS demand. CO_2 capture and emissions downstream of carbon capture units are assumed proportional to inlet CO_2 flows, governed by the capture efficiency of the respective CC technology.

Hourly mass and energy balances are also imposed across resource nodes at each hour, extending equation (3) in [12].

The mass balance on hydrogen storage, as shown in equation (33) tracks the evolution of stored hydrogen, $stor_{H_2,j,hr,d,yr}$, as a function of the net flow into the storage unit at each hour.

$$stor_{H_2,j,hr+1,d,yr} = stor_{H_2,j,hr,d,yr} + \left(\sum_{k\in In(j)} F_{k,j,hr,d,yr} - \sum_{l\in Out(j)} F_{l,j,hr,d,yr} \right) \cdot \Delta t \quad \forall\, j \in \{H_2\ storage\} \quad (33)$$

Equation (34) enforces continuity between the storage level at the last hour of each day and the first hour of the same representative day, valid due to chronology-preserving AHC.

$$stor_{H_2,j,1,d,yr} = stor_{H_2,j,24,d,yr} + \left(\sum_{k\in In(j)} F_{k,j,24,d,yr} - \sum_{l\in Out(j)} F_{l,j,24,d,yr} \right) \cdot \Delta t \quad \forall\, j \in \{H_2\ storage\} \quad (34)$$

Equation (35) links the storage level at the last hour of each day to that at the first hour of the following representative day, valid due to chronology-preserving AHC.

$$stor_{H_2,j,1,d+1,yr} = stor_{H_2,j,24,d,yr} + \left(\sum_{k\in In(j)} F_{k,j,24,d,yr} - \sum_{l\in Out(j)} F_{l,j,24,d,yr} \right) \cdot \Delta t \quad \forall\, j \in \{H_2\ storage\} \quad (35)$$

The intra-day state-of-charge dynamics of battery energy storage between consecutive hours are modeled in a manner similar to equation (33) in equation (36), accounting for both charging and discharging efficiencies.

$$stor_{elec.,j,hr+1,d,yr} = stor_{elec.,j,hr,d,yr}$$

$$+\left(\sum_{k\in In(j)} P_{k,j,hr,d,yr}\cdot\eta^{+}-\sum_{l\in Out(j)} P_{l,j,hr,d,yr}\right)/\eta^{-}\,]\cdot\Delta t$$

$$\forall\, j\in\{battery\} \tag{36}$$

Equations (37) and (38) establish continuity between the state of charge in batteries on consecutive days, analogous to equations (34) and (35).

$$stor_{elec.,j,1,d,yr} = stor_{elec.,j,24,d,yr}$$

$$+\left(\sum_{k\in In(j)} P_{k,j,24,d,yr}\cdot\eta^{+}-\sum_{l\in Out(j)} P_{l,j,24,d,yr}\right)/\eta^{-}\,]\cdot\Delta t$$

$$\forall\, j\in\{battery\} \tag{37}$$

$$stor_{elec.,j,1,d+1\,yr} = stor_{elec.,j,24,d,yr}$$

$$+\left(\sum_{k\in In(j)} P_{k,j,24,d,yr}\cdot\eta^{+}-\sum_{l\in Out(j)} P_{l,j,24,d,yr}\right)/\eta^{-}\,]\cdot\Delta t$$

$$\forall\, j\in\{battery\} \tag{38}$$

Equations relating variables from SMR, TES modules and the rest of the superstructure

The electricity balance around the SMR is enforced in equation (39). At each hour and representative day, the total electrical output of all SMR types and modules, $\sum_g\sum_n e_{g,n,hr,d,yr}$, must be fully allocated to: (i) electricity supplied directly to internal equipment, $P_{\text{elec., SMR},hr,d,yr}$; (ii) electricity flow to charge batteries $P_{\text{elec. at dynamic price, SMR},hr,d,yr}$, and (iii) electricity sold to the grid at the dynamic selling price, $P_{elec.sold,SMR,hr,d,yr}$

$$\sum_{g\in G}\sum_{n\in N} e_{g,n,hr,d,yr} = P_{\text{elec., SMR},hr,d,yr} + P_{\text{elec. sold, SMR},hr,d,yr} + P_{\text{elec. at dynamic price, SMR},hr,d,yr} \tag{39}$$

Equation (40) states that for each reactor type and module, the total energy produced at the core of the nuclear reactor discounted by the losses due to heat transfer, $r_{g,n,hr,d,yr}\,\eta_{HT_g}$, is split consistently between three competing uses: direct onsite steam production $h_{g,n,hr,d,yr}$, TES charging $\frac{c_{g,n,hr,d,yr}}{\eta^{+}}$, and electricity generation via the turbine, represented by $e_{g,n,hr,d,yr}/\eta_{TE_g}$.

$$h_{g,n,hr,d,yr} + \frac{c_{g,n,hr,d,yr}}{\eta^{+}} + \frac{e_{g,n,hr,d,yr}}{\eta_{TE_g}} = r_{g,n,hr,d,yr}\,\eta_{HT_g} \tag{40}$$

Finally, equation (41) closes the thermal balance between the SMR–TES system and the refinery steam network. The sum, over all SMR types and modules, of direct reactor heat used for steam production, scaled by the overall heat-transfer efficiency η_{HT}, $h_{g,n,hr,d,yr}\eta_{HT}$, and the energy generated by discharging the TES, $d_{g,n,hr,d,yr}$, scaled by the discharge efficiency, scaled by the overall heat-transfer efficiency η^{-} is equated to the effective high-pressure steam stream $F_{\text{HPS, SMR},hr,d,yr}$ multiplied by the enthalpy of vaporization ΔH_{vap}.

$$\sum_{g\in G}\sum_{n\in N}\left(h_{g,n,hr,d,yr}\eta_{HT} + d_{g,n,hr,d,yr}\eta^{-}\right) = F_{\text{HPS, SMR},hr,d,yr}\,H_{vap} \tag{41}$$

Finally, the annual operational expenditure accounts for costs related to natural gas and electricity procurement, operations of carbon capture equipment, storage system operation, carbon tax incurred on residual CO_2 emissions and operations of the SMR, shown in equation (42). The contribution of each representative day is weighted by its occurrence frequency w_d.

$$\begin{aligned}\text{OPEX}_{\text{yr}} = {} & \Delta t\cdot\sum_{d\in D} w_d \sum_{hr\in HR}\sum_{\{j\in Out(NG)\}} F_{NG,j,hr,d,yr}\cdot \$_{NG,yr}\\ & P_{elec.at\ dynamic\ price,grid,hr,d,yr}\cdot \$_{elec.,hr,d,yr}\\ & + \sum_{j\in\{CC\ units\}} F_{CO_2Captured,j,hr,d,yr}\cdot\left(\frac{\text{NG}_{\text{consumed}}}{\text{unit mass CO}_2}\$_{NG,yr} + \$_{rest\ of\ OPEX,j,yr}\right)\\ & + \sum_{j\in H_2Storage} F_{H_2,j,hr,d,yr}\cdot \$_{H_2\ Storage,yr}\\ & +\left(P_{elec.at\ dynamic\ price,hr,d,yr} + P_{elec,Battery,hr,d,yr}\right)\cdot \$_{Battery,yr}\\ & \sum_{j\in EmissionSources} F_{CO_2Released,j,hr,d,yr}\cdot \$_{CO_2tax,yr}\\ & + OPEX_{SMR,yr}\quad \forall yr\in YR\end{aligned} \tag{42}$$

Annual revenue from SMR's electricity sales is computed as the weighted sum of electricity sold at a fraction of the buying price (discounting for grid congestion effects), as shown in equation (43).

$$Revenue_{yr} = \Delta t\cdot\sum_{d\in D} w_d\left(\sum_{hr\in HR}\left(P_{elec.sold,SMR,hr,d,yr} + P_{elec.sold,battery,hr,d,yr}\right)\cdot \$_{electricity,hr,d,yr}\ f_{selling}\right)\forall yr\in \text{YR} \tag{43}$$

Capital expenditure for process capacity expansion are modeled in Equation (44) using a fixed-plus-variable cost formulation. As in [12], the first term, $\alpha_{j,yr}$ denotes the fixed cost incurred for technology j in year yr, while the second term, $\beta_{j,yr}QE_{j,yr}$ accounts for the variable cost. Both components are assumed to be paid throughout the construction period $dur_{j,}$. The CAPEX associated with the SMR and TES units are added separately.

$$\begin{aligned}CAPEX_{yr} = {} & \sum_{j\in J\setminus\{SMR,TES\}}\left(\sum_{h=0}^{dur(j)-1}\frac{\alpha_{j,yr+h}y_{yr+h}}{dur(j)} + \sum_{h=0}^{dur(j)-1}\frac{\beta_{j,yr+h}\,QE_{j,yr+h}}{dur(j)}\right) + CAPEX_{SMR,yr}\\ & + CAPEX_{TES,yr}\quad \forall yr\in YR\end{aligned} \tag{44}$$

The first objective minimizes the net present cost as the decarbonization strategy, accounting for annual

capital costs, operating costs, and revenues over the planning horizon and appropriately discounting them, as shown in equation (45).

$$Z_1 = \min \sum_{yr \in YR} \frac{CAPEX_{yr} + OPEX_{yr} - Revenue_{SMR,yr}}{(1+int)^{yr}} \quad (45)$$

The second objective minimizes the overall CO_2 equations over the time horizon, as shown in equation (46).

$$Z_2 = \sum_{yr \in YR} \Delta t \cdot \sum_{d \in D} w_d \sum_{hr \in HR} \sum_{j \in In(CO_2)} F_{CO_2,j,hr,d,yr} \quad (46)$$

RESULTS

The results are presented as Pareto fronts (obtained using ε-constraint method) showing the trade-off between present cost and cumulative CO_2 emissions under different policy and market scenarios. Data for the refinery system follows Chattopadhyay et al. [1], and those for the SMR have been obtained from Vanatta et al. [11]. We used five years of hourly price data from the Houston hub of the ERCOT electricity market [14]. Detailed assumptions on battery and hydrogen storage costs will be reported in the full-length paper. Each point on the front represents a non-dominated solution: cost cannot be reduced without increasing emissions, and emissions cannot be reduced without raising costs. Comparing fronts across (i) cases with and without carbon taxes, and (ii) models with and without flexible operation and storage, reveals economically attractive decarbonization pathways. Case (1) considers no operational flexibility and uses average annual prices over a 19-year horizon. This resulted in a MIQCP in Pyomo 6.6.1 comprising 4,674 constraints (1,881 equalities,) and 4,504 variables (1,539 binaries), which was solved with BARON 24.10.10, with a MIP gap tolerance of 0.01%. In contrast, the case considering hourly varying electricity prices and the possibility of installation of storage, modeled using three representative days per year, resulted in a much larger MIQCP with 303,222 constraints (119,833 equalities) and 251,598 variables (66,519 binaries), which was solved with Gurobi 10.0.1 with a MIP gap tolerance of 0.01%. Generation of the Pareto frontier, constructed from 20 discrete points, required 4.5 and 43.6 minutes of CPU time for the inflexible and flexible cases, respectively.

Figure 3 compares case 1 and case 2 scenarios with the benchmark frontiers previously reported in our work [12], in which refinery decarbonization options did not include SMRs. Across the entire emissions range, the inclusion of the possibility of installing SMRs reduces the overall retrofitting costs, pushing the Pareto front downward. The possibility of flexible operations further reduces costs. The leftmost point on the Pareto frontier yielding minimum emissions has also moved leftward, due to a reduced value of minimum attainable feasible scope 1 and 2 emissions. Overall, CO_2 taxes raise retrofitting costs across scenarios, and the non-dominated solutions are compressed towards lower emission values. Beyond a certain threshold, increasing the ε value no longer increases CO_2 emissions, as the ε-constraint becomes inactive. Staying below the allowed ε limit is financially beneficial due to carbon tax penalties for exceeding it.

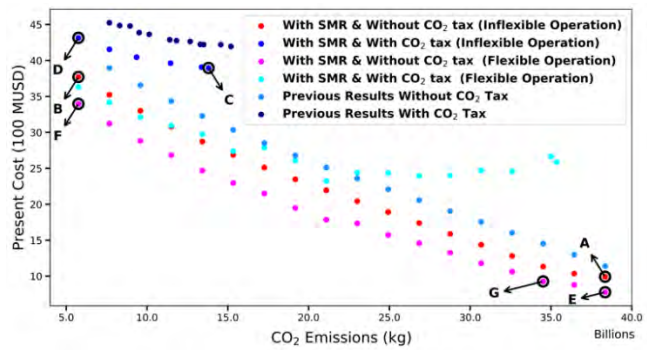

Fig. 3: Comparison of Pareto fronts obtained under all cases. Previous results correspond to those reported in [12].

Case (1): Inflexible operations assuming average annual electricity Prices (No Storage)

Fig. 4 shows the solutions obtained for points (A) and (B) labeled in Fig. 3. When considering average annual electricity prices without CO_2 taxes, the lowest-cost solutions on the rightmost part of the frontier rely on conventional operations during the initial years. In this period, hydrogen is supplied entirely by the steam methane reformer; no carbon capture units are installed, and high-pressure steam (HPS) is primarily produced by the natural gas boiler. The SMR mainly functions as a power unit, with electricity sales at 0.8× the purchasing price generating revenue to offset installation costs. In the later years, as natural gas prices are expected to rise, SMRs are partially used for HPS generation. Moving left along the frontier as the ε-constraint is tightened, HPS production increasingly comes from the heat generated by SMRs earlier, while hydrogen continues to be produced via reforming, and electrolyzers remain unused. Both pre-combustion and post-combustion carbon capture units are installed: post-combustion units treat dilute flue gas streams, whereas pre-combustion capture is applied to the higher-concentration reformer gas.

Under carbon taxes implemented at rates, the same as those discussed in Chattopadhyay et al [1], capture is implemented earlier as compared to the high emission solution in the case without taxes, and HPS generation is from SMR-generated heat earlier in the planning horizon

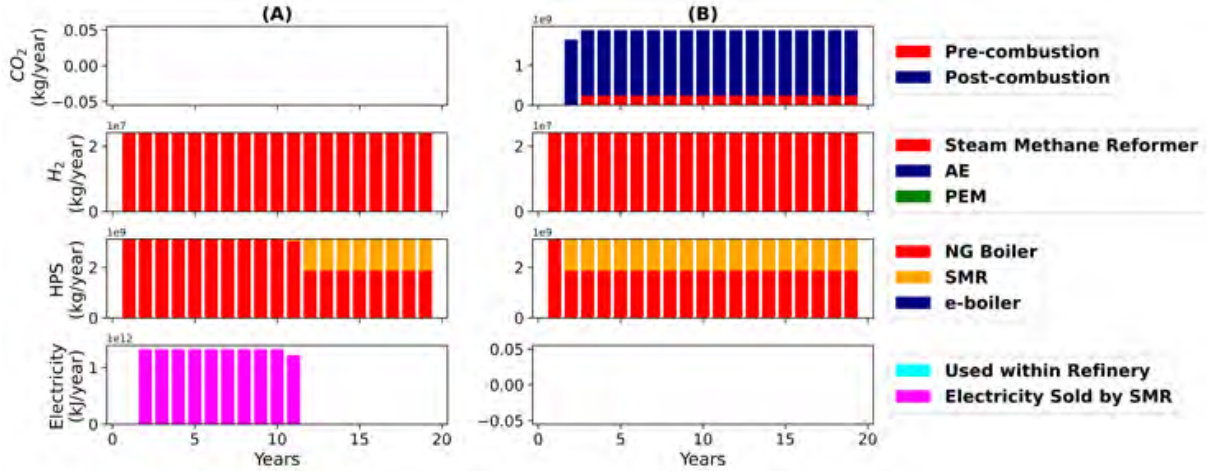

Fig. 4: Solutions obtained for case (1) without CO2 taxes

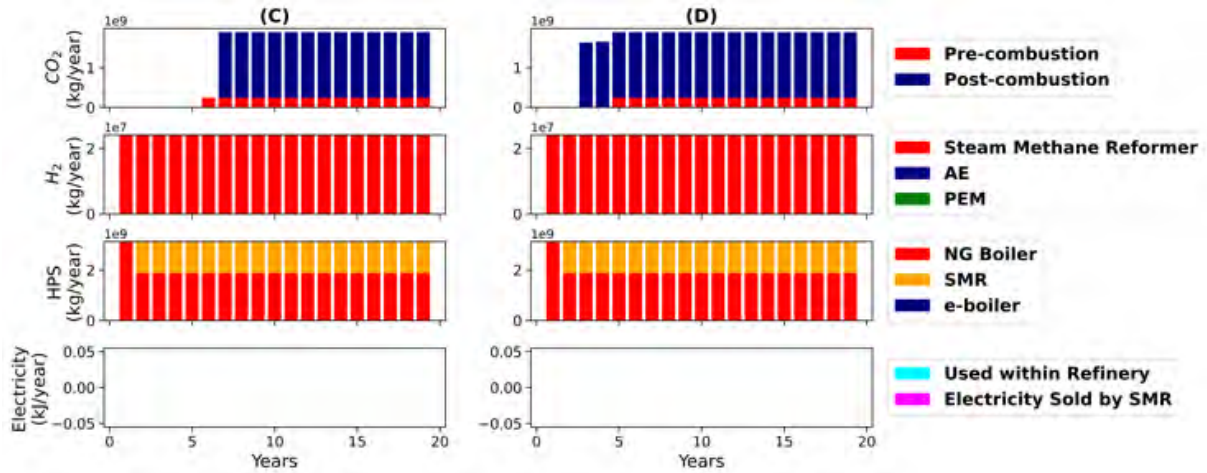

Fig. 5: Solutions obtained for case (1) with CO2 taxes

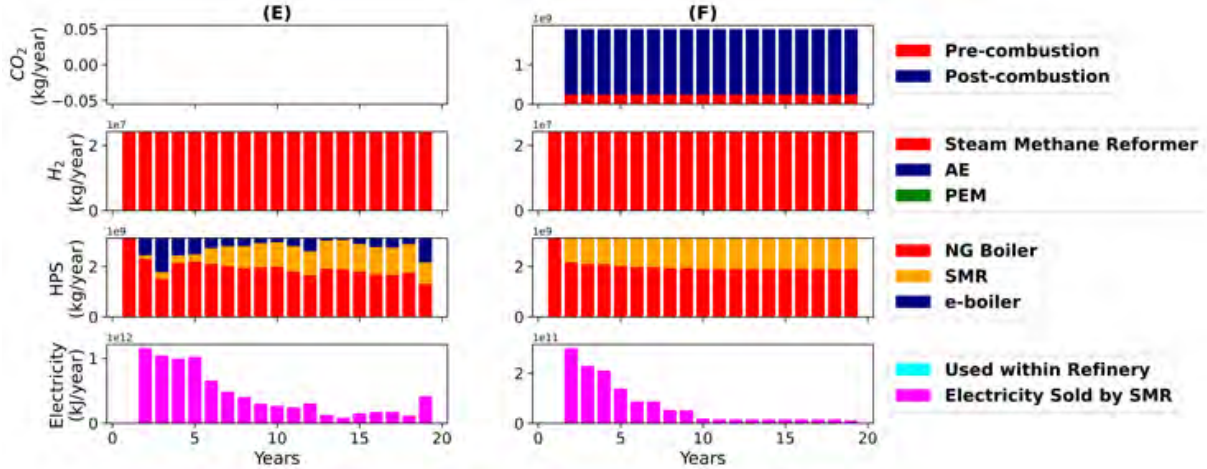

Fig 6: Non-dominated solutions under flexible operations; no carbon taxes

even for the rightmost point in the Pareto front (labeled as C in the fig. 3), as shown in Fig. 5. As we move leftward along the Pareto front, solutions involve earlier incorporation of pre and post combustion capture, as shown by the solution at point D in fig. 5.

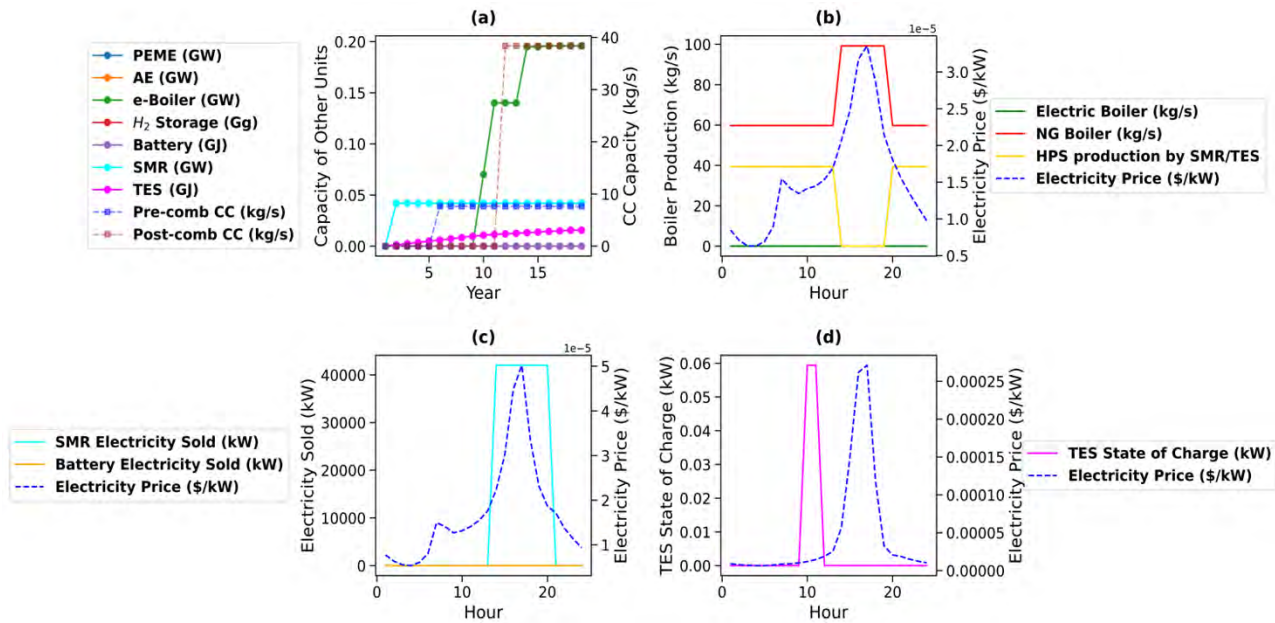

Fig 7: (a) Capacity expansion at a representative Pareto optimal solution (b) flexible HPS production over a representative day (c) Electricity sales from Small Modular Reactor across a representative day (d) Evolution of TES State of Charge across a representative day; for point G.

Case 2: Hourly Prices, Storage, and Flexible Operation

Case 2 considers hourly variations in electricity pricing and the option for enhanced flexibility with the possibility of installation of TES, battery, and hydrogen storage. Figure 6 shows that as the model exploits temporal price variation: hydrogen is still produced mainly by the steam methane reformer, while HPS is supplied by a mix of natural gas boiler, SMR, and e-boiler. E-boiler is chosen in this case, since flexible operation allows electricity to be bought from the grid when cheap for HPS generation. As emissions constraints tighten along the front, e-boiler usage is reduced to limit grid-related emissions, and SMR-generated heat and natural gas boiler in conjunction with capture units, are the main sources for HPS generation.

Figure 7 (a) shows investment decisions for a representative low-cost solution labeled as G in Fig. 3 when considering the possibility of flexible operations in the absence of carbon taxes. The SMR units are built by year 2, iPWR reactors are chosen; e-boiler units are installed from year 10, with expansions until year 13, pre-combustion capture capacity is installed by year 6, and post-combustion capture capacity is installed by year 12. TES capacity is installed starting from year 2, with capacity expansions until year 18. In contrast, no battery or hydrogen storage capacity is installed: under the assumed price profiles, price variations leveraged by these do not add enough value to justify their investment costs. Fig. 7b (hourly profiles, year 3, representative day 3) illustrates how the SMR, NG boiler, and e-boiler contribute towards steam generation. When electricity prices are low, SMR output is allocated more toward steam production; as prices rise, the NG boiler increasingly takes care of the steam load so that the SMR can be used for electricity generation and sales. The electricity-sales profile shown in fig. 7c (hourly profiles, year 3, representative day 1) also follows the price curve while respecting maximum ramp-up/ down rates: more electricity is sold during high-price hours. Finally, Fig. 7d (hourly profiles, year 7, representative day 2) shows the evolving state of charge for the TES unit. The profile roughly follows the electricity price signals, but at a time lead, which enables the TES unit to charge during the low-price durations while allowing the system to use heat generated from the discharge of TES during high price durations.

CONCLUSIONS

This work developed a multi-period superstructure optimization framework to evaluate the economic and environmental feasibility of integrating SMRs into an existing refinery for cogeneration of heat and electricity. Across all scenarios, the inclusion of SMRs as a decarbonization option lowers costs at comparable emissions or deeper emissions reductions at similar cost. They are predominantly deployed for high-pressure steam generation, reflecting the efficiency of direct nuclear heat utilization. Carbon taxes accelerate the adoption of SMRs and carbon capture units.

Flexible operation consistently reduces system

costs at comparable emissions levels, with TES emerging as a key enabler for leveraging the price variability, whereas battery and hydrogen storage are not selected under the assumed conditions. Despite the significant increase in model size, large-scale MIQCPs were solved to satisfactory convergence, demonstrating the tractability of the proposed approach. Overall, the results suggest that SMRs, particularly when paired with flexible operation and thermal storage, offer a cost-effective pathway for deep refinery decarbonization under a range of policy conditions.

NOMENCLATURE

Table 2: Abbreviations used in the text

SMR	Small Modular Reactor
TES	Thermal Energy Storage
iPWR	Integrated Pressurized Water Reactor
HTGR	High Temperature Gas-Cooled Reactor
MSR	Molten Salt Reactor
NG	Natural Gas
MIQCP	Multi-Integer Quadratically Constrained Problem

REFERENCES

1. Chattopadhyay S, Karthikeyan K, Gandhi R, Grossmann IE, Torres AI. Optimal retrofit of carbon capture and electrification technologies in oil refineries for reducing direct co₂ emissions. Ind. Eng. Chem. Res. 64:10131-10147 (2025). https://doi.org/10.1021/acs.iecr.4c04995
2. Cost estimating guidelines for Generation IV nuclear energy systems – Revision 4.2. Technical report, Generation IV International Forum, Economic Modeling Working Group. September 2007. URL: https://www.gen-4.org/gif/upload/docs/application/pdf/2013-09/emwg_guidelines.pdf
3. Mignacca B, Locatelli G. Economics and finance of small modular reactors: a systematic review and research agenda. Renewable and Sustainable Energy Reviews 118:109519 (2020). https://doi.org/10.1016/j.rser.2019.109519
4. Nuclear Energy Agency (NEA). Small Modular Reactors: Challenges and Opportunities. OECD Nuclear Energy Agency (2016).https://doi.org/10.1787/18fbb76c-en
5. Ingersoll DT. Deliberately small reactors and the second nuclear era. Progress in Nuclear Energy 51:589-603 (2009). https://doi.org/10.1016/j.pnucene.2009.01.003
6. Asuega A, Limb BJ, Quinn JC. Techno-economic analysis of advanced small modular nuclear reactors. Applied Energy 334:120669 (2023). https://doi.org/10.1016/j.apenergy.2023.120669
7. Wilding PR, Murray NR, Memmott MJ. The use of multi-objective optimization to improve the design process of nuclear power plant systems. Annals of Nuclear Energy 137:107079 (2020). https://doi.org/10.1016/j.anucene.2019.107079
8. Kumar A, Tsvetkov PV. A new approach to nuclear reactor design optimization using genetic algorithms and regression analysis. Annals of Nuclear Energy 85:27-35 (2015). https://doi.org/10.1016/j.anucene.2015.04.028
9. Locatelli G, Boarin S, Fiordaliso A, Ricotti ME. Load following of small modular reactors (SMR) by cogeneration of hydrogen: a techno-economic analysis. Energy 148:494-505 (2018). https://doi.org/10.1016/j.energy.2018.01.041
10. Hill D, Martin A, Martin-Nelson N, Granger C, Memmott M, Powell K, Hedengren J. Techno-economic sensitivity analysis for combined design and operation of a small modular reactor hybrid energy system. International Journal of Thermofluids 16:100191 (2022). https://doi.org/10.1016/j.ijft.2022.100191
11. Vanatta M, Patel D, Allen T, Cooper D, Craig MT. Technoeconomic analysis of small modular reactors decarbonizing industrial process heat. Joule 7:713-737 (2023). https://doi.org/10.1016/j.joule.2023.03.009
12. Karthikeyan K, Chattopadhyay S, Gandhi R, Grossmann IE, Torres AI. Pareto optimal solutions for decarbonization of oil refineries under different electricity grid decarbonization scenarios. Systems and Control Transactions 4:844-849 (2025). https://doi.org/10.69997/sct.102781
13. Tso WW, Demirhan CD, Heuberger CF, Powell JB, Pistikopoulos EN. A hierarchical clustering decomposition algorithm for optimizing renewable power systems with storage. Applied Energy 270:115190 (2020). https://doi.org/10.1016/j.apenergy.2020.115190
14. ENGIE Resources. Historical data reports. https://www.engieresources.com/historical-data#reports_anchor, 2025.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Exploring the Thermal Coupling of Solid Oxide Electrolysis and Ammonia Synthesis: A Plantwide Energy Integration Assessment

Alessandra Macchi[a], Federica Longobardi[b], Régis Anghilante[b], Andrea Isella[a], and Davide Manca[a]*

[a] PSE-Lab, Process Systems Engineering Laboratory, Dipartimento di Chimica, Materiali e Ingegneria Chimica "Giulio Natta", Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy
[b] Casale SA, Via Giulio Pocobelli 6, 6900 Lugano, Switzerland
* Corresponding Author: davide.manca@polimi.it.

ABSTRACT

The transition toward low-carbon ammonia production increasingly relies on highly efficient routes for renewable hydrogen generation, with Solid Oxide Electrolysis (SOE) representing a particularly promising solution. SOEs, operating with steam at elevated temperatures, offer intrinsic thermodynamic advantages and are attractive when integrated with processes that can effectively utilize or supply high-grade heat. This context opens the possibility for advanced energy coupling between hydrogen production and the exothermic ammonia synthesis process. This work investigates such an energy integration strategy by simulating a small-scale ammonia plant where hydrogen is produced through an SOE system thermally coupled to the ammonia synthesis loop. Specifically, high-grade heat available in the Haber-Bosch (HB) reactor outlet is recovered via a "Heat Recovery Steam Generator" (HRSG) to provide a substantial fraction (55 wt%) of the steam required by the electrolyzer. The assessment demonstrates that the integrated configuration significantly reduces the external electrical demand for steam generation by 84%, achieving a substantial enhancement in system-level electrical efficiency. Total plant energy efficiency increases from 57% to 62%, and the electricity-to-hydrogen ratio decreases from 45.9 to 42.1 kWh/kg. These improvements translate directly into a decrease in the Levelized Cost of Ammonia (LCOA) by 4–5%, underscoring the relevance of high-temperature electrolysis supported by effective plantwide heat recovery.

Keywords: Green ammonia, Haber-Bosch process, Hydrogen, Solid Oxide Electrolysis, Energy Integration

INTRODUCTION

Ammonia is a key industrial commodity, ranking first in global production volume, with an approximate annual yield of 185 Mt/y [1]. Its main synthetic route, the Haber-Bosch process, converts nitrogen and hydrogen into ammonia under severe high-temperature and high-pressure conditions. While nitrogen is readily obtained from air, the hydrogen feedstock is predominantly derived from fossil fuels, primarily via steam methane reforming or coal gasification. This dependence results in substantial greenhouse gas (GHG) emissions. Thus, conventional ammonia synthesis processes are highly CO_2-intensive, releasing an estimated 1.65 kg_{CO_2eq}/kg_{NH_3}, which accounts for approximately 440 Mt_{CO2eq} annually [2].

Decarbonization efforts for ammonia production inherently necessitate the substitution of fossil-based hydrogen with "green" hydrogen, to be generated via water electrolysis powered by renewable electricity. Electrolysis, an endergonic process, electrochemically dissociates water into hydrogen and oxygen, yielding a carbon-free hydrogen product when the electrical power source is renewable. State-of-the-art commercial technologies include low-temperature systems, such as alkaline (AEL) and proton-exchange-membrane (PEMEL) electrolyzers, which utilize liquid water. In contrast, solid oxide electrolyzers convert high-temperature steam into hydrogen, offering superior thermodynamic efficiency due to the reduced electrical energy demand when operating at elevated temperatures (*i.e.*, 500-900 °C).

The inherent thermal compatibility between the high-temperature heat available in the exothermic ammonia synthesis loop effluents and the high thermal demand of the SOE steam feed is central to this integration, providing a significant opportunity for system-level energy optimization via direct waste heat recovery.

This study investigates a 300 kg_{NH3}/h (small-scale) green ammonia plant configuration where the hydrogen feedstock is supplied by an SOE thermally coupled to the synthesis loop. Waste heat recovery from the HB reactor outlet is directly utilized to generate the required superheated steam feed for the electrolyzer. Therefore, the primary objective of this work is to provide a plantwide assessment of the energy integration potential, rigorously evaluating the energetic and techno-economic performance of this integrated configuration compared to a baseline, non-integrated scenario.

WATER ELECTROLYSIS

Water electrolysis is a key pathway for CO_2-free hydrogen production utilizing renewable electricity. Significant thermodynamic distinctions exist between low-temperature liquid-water electrolysis and high-temperature steam electrolysis, the latter being implemented exclusively in solid oxide electrolyzers.

Thermodynamics of Water Splitting

Water splitting is a non-spontaneous, endergonic reaction requiring an external electrical energy supply. The minimum electrical energy demand is governed by the Gibbs free energy change (ΔG), which decreases inversely with temperature. In the Solid Oxide Electrolyzer Cell (SOEC), supplying superheated steam at 800 °C reduces ΔG from approximately 237 kJ/mol for liquid water (ambient conditions) to about 189 kJ/mol. This corresponds to a drop in the minimum electrical energy requirement by roughly 20%. Crucially, this reduction is cell-specific and does not account for the additional thermal energy required to generate and superheat the steam feed. Consequently, the net energetic advantage of SOE technology is realized only when low-cost or waste heat sources are available for steam generation.

This thermodynamic reduction is also reflected in the cell-level reversible potential (a.k.a. Nernst potential), which defines the minimum theoretical voltage required under equilibrium conditions. The operating cell voltage must exceed the reversible potential to compensate for irreversible losses (*i.e.*, overpotentials), including ohmic resistance, activation barriers, and mass transport limitations. For low-temperature water electrolysis at ambient conditions, the reversible potential is approximately 1.23 V. Supplying superheated steam at 800 °C reduces the reversible potential to approximately 0.83 V, a decrease of nearly 30%. This substantial reduction underscores the potential for decreased electrical energy input per mole of hydrogen in high-temperature steam electrolysis.

Another important reference voltage at the cell level is the thermoneutral potential, which defines the thermoneutral operating condition where the supplied electrical energy exactly balances the total reaction enthalpy. At this potential, the cell operates with zero net heat exchange with the surroundings. Operating below this threshold leads to an endothermic regime, where the heat deficit for the electrolysis reaction is provided by the feed gases. As a result, the outlet gas temperature is lower than the inlet gas temperature. Conversely, operating above this value results in an exothermic regime, where the outlet gas temperature is higher than the inlet gas temperature. Operating near the thermoneutral condition is preferred as it minimizes internal temperature gradients, reducing thermal stress and enhancing long-term stability. Low-temperature electrolyzers (AEL and PEMEL) are generally constrained to the exothermic regime. In contrast, SOEs offer greater operational flexibility by utilizing all three regimes, which is particularly beneficial for process integration and the exploitation of waste heat.

Solid Oxide Electrolysis Technology: Single Cell Structure and Operation

The solid oxide electrolyzer is a high-temperature electrolysis unit operating at elevated temperatures ranging from 500 to 900 °C. The basic functional unit of this technology is the single Solid Oxide Electrolysis Cell (SOEC), schematically illustrated in **Figure 1**.

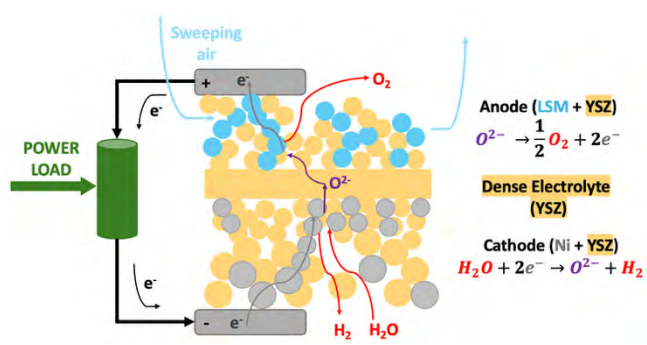

Figure 1. Structure of the solid oxide electrolysis cell.

In the SOEC, superheated steam is supplied to the cathode, where the hydrogen evolution reaction (HER) produces hydrogen (H_2) and oxide ions (O^{2-}). These ions migrate through the solid ceramic electrolyte to the anode, where the oxygen evolution reaction (OER) generates pure oxygen (O_2). Due to incomplete steam conversion, a portion of the feed steam exits the cathode along with the produced hydrogen, necessitating downstream separation.

The cell relies on an entirely ceramic composition, ensuring mechanical integrity and thermal stability. Key

materials include the Yttria-Stabilized Zirconia (YSZ) electrolyte, which provides high oxide-ion conductivity. The cathode consists of a Ni-YSZ cermet (ceramic-metal), where nickel acts as a catalyst for the HER and provides electronic conductivity. The anode is often composed of a strontium-doped lanthanum manganite (LSM) perovskite, which is thermally compatible with YSZ.

The defining operational parameter is the high-temperature operation, contrasting with low-temperature technologies (AEL and PEMEL). SOEs typically operate at pressures near atmospheric, as long-term stable performance at significantly higher pressures remains a technological challenge. The wide operating voltage range confirms the theoretical feasibility of operation in all three thermodynamic modes. This high-temperature operation provides a substantial energy advantage, reducing electrical power consumption by approximately 25-28% relative to low-temperature electrolyzers. Crucially, this efficiency gain is maximized when external low-cost or waste heat sources supply the thermal energy for steam generation. **Table 1** summarizes the SOE's key operational parameters.

Challenges associated with high-temperature operation include accelerated material degradation, which reduces cell lifetime and necessitates precise temperature control and smooth transients to maintain structural integrity. Furthermore, from an industrial perspective, current SOE unit capacities (up to 10 MW deployed) are significantly smaller than those of AEL (>1 GW) and PEMEL (>100 MW), indicating that scaling SOE technology for large-scale hydrogen production remains a significant commercial and engineering challenge [3].

Table 1: Operating parameters of solid oxide electrolysis technology [4].

OPERATING PARAMETER	Range of values
T [°C]	500-900
P [bara]	1-3
Cell type [-]	Cathode-supported cell (CSC)
Cell voltage [V]	1.1-1.5
DC electrical power consumption [kWh/kg$_{H_2}$]	48 (with no steam import) 38 (with steam import)
Voltage degradation rate [%/1000h]	0.55-1
Industrial maturity	< 10 MW deployed

SOLID OXIDE ELECTROLYZER MODEL

A steady-state, zero-dimensional (0D) model was developed to represent the electrochemical and process behavior of the solid oxide electrolyzer in a simplified yet physically consistent manner. A fully detailed electro-chemical-thermo-mechanical model is outside the scope of this work. Instead, the primary objective is to formulate a computationally efficient representation of the SOE suitable for implementation within the plantwide process simulation environment used for the integrated ammonia production system.

The model first describes the operation of the fundamental unit (*i.e.*, the single electrolysis cell) and subsequently extends this description to the stack level, where multiple cells are electrically connected in series to form an electrolyzer module of a defined size.

The core modeling assumption is that both the cell and the stack operate within the linear region of their polarization curves. Under this premise, the irreversible losses are not treated mechanistically but are collectively quantified by a single, comprehensive parameter: the Area Specific Resistance (ASR), in [Ω·cm^2]. The ASR term incorporates all irreversible losses arising during cell operation, including ohmic resistance, activation, and concentration overpotentials. The correlation for single cells' ASR is typically obtained by regression of experimental polarization curves, whose value is a function of operating conditions, primarily temperature and pressure. Consequently, the single cell operating voltage, E_{cell} [V], is defined by the following linearized relationship, where E_N [V] corresponds to the Nernst potential and $\hat{\imath}$ [A/cm^2] denotes the current density.

$$E_{cell} = E_N + ASR \bullet \hat{\imath} \tag{1}$$

At the stack level, additional irreversible losses arise due to the series electrical connection of multiple cells. To account for these effects, an additional constant resistance contribution of 0.1 Ω·cm^2 is incorporated into the cell ASR correlation when evaluating cell performance once integrated in a stack [5]. As a result of this approach, all cells in a stack are assumed to exhibit the same behavior. This methodology preserves computational simplicity while providing a realistic representation of the increased irreversibility observed in stack or short-stack configurations.

Cell-Level and Stack-Level Formulation

The operating model of the solid oxide electrolyzer is first formulated at the single-cell level and subsequently extended to the stack configuration. The model requires both cell-specific (and, analogously, stack-specific) parameters and the set of operating conditions associated with the selected operational mode.

At the cell level, the model utilizes several inputs:

the active cell area (where electrochemical reactions occur); the inlet hydrogen mole fraction at the cathode; the Faradaic efficiency; the steam-to-hydrogen ratio; and the ASR correlation that characterizes the cell's resistive behavior. The operating state is formally defined by temperature, pressure, and the specified operating regime.

Based on these inputs, the model determines the hydrogen production rate according to Faraday's law for electrolysis, the steam conversion at the cathode outlet, the electrical power required to sustain the imposed operating mode, and the resulting energetic efficiencies. A summary of the input data and the corresponding output variables is reported in **Table 2**.

When the modeling framework is extended to the stack level, the same governing equations are preserved, while specific input parameters are adjusted. Crucially, the ASR correlation is modified to account for the incremental resistive contribution, and the total number of cells connected in electrical series is specified. Under these conditions, the model computes the corresponding performance metrics on a stack basis, accurately capturing the linear voltage scaling imposed by series connectivity and the uniform current density that characterizes all cells within the assembly.

Table 2: Input parameters and resulting performance outputs from the cell-level operating model.

SOEC-Level Operating Model	
INPUT DATA	**OUTPUT RESULTS**
Cell-specific parameters	
Active cell area [cm^2]	H_2 production rate [mol/s]
Cell Faradaic efficiency [%]	
Steam-to-H_2 ratio [kg_{steam}/kg_{H_2}]	Steam conversion [%]
ASR law at the cell-level	Energetic (LHV and HHV) efficiency [%]
Operating conditions	
Temperature [°C]	Electrical power [W]
Pressure [bar]	
Operating regime	

SOE Simulation and Balance of Plant (BOP)

SOEs require several auxiliary process units for proper thermal, fluidic, and operational management. Accordingly, the entire hydrogen production section was modeled in UniSim® Design R500 to evaluate the associated mass and energy balances and overall performance. The resulting process flowsheet, which models all key components of the SOE unit and its Balance of Plant, is shown in **Figure 2**.

Since UniSim does not feature a dedicated solid oxide electrolyzer simulation item or tool, the electrochemical behavior of the SOE stack was integrated into the simulation environment via a spreadsheet interface. The SOE was represented through a Conversion Reactor coupled with the external SOE model described so far. Specifically, the 0D thermoneutral stack model directly supplies the Conversion Reactor with the key operating inputs: namely, the steam conversion and the electrical power required to ensure thermoneutral operation. Downstream of the reactor, a component separator replicates the intrinsic outlet stream partitioning of an SOE, separating the cathode effluent (hydrogen and unconverted steam) from the anode effluent (oxygen and sweeping air).

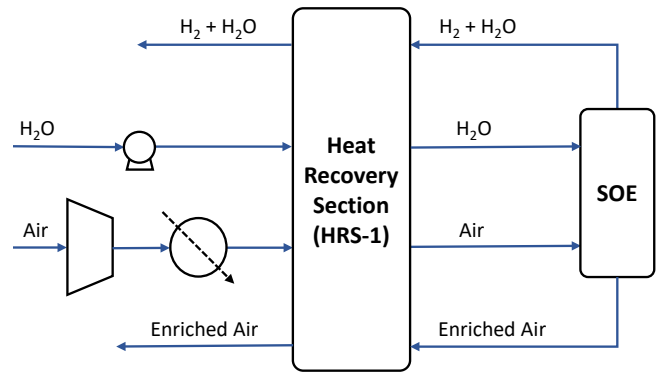

Figure 2. Schematic diagram of the hydrogen production section and the 1st heat recovery section (HRS-1).

At the process level, liquid water is fed to the hydrogen production section and must be conditioned to the SOE operating conditions. Preheating and vaporization are primarily achieved in the heat recovery section HRS-1, which recovers thermal energy from the high-temperature SOE exhaust streams. Superheated steam is fed to the electrolyzer, where conversion is deliberately limited (*i.e.*, never exceeding 90%). Unreacted steam leaves the cathode together with the produced hydrogen. Separation is achieved through cooling and condensation, which allows liquid water to be recovered and recycled.

Moreover, a minimum hydrogen fraction (≥10 mol%) was assumed at the cathode inlet to prevent oxidation of the metallic nickel phase in the Ni-YSZ cathode. Furthermore, an air sweep system was assumed at the SOEC anode to mitigate safety risks associated with the presence of pure oxygen at high temperatures. The literature does not always explicitly address these requirements [6]: in the present work, the necessary hydrogen fraction is provided by recycling a defined portion of the hydrogen-rich stream; simultaneously, hot air is supplied to the anode to dilute the produced oxygen, ensuring that the oxygen concentration in the anode exhaust remains at an acceptable level.

In such a process configuration, the final hydrogen product stream leaving the SOE section is delivered at approximately 40 °C and ambient pressure, containing a residual moisture content of ca. 7 mol%.

PROCESS SIMULATION OF THE GREEN AMMONIA PRODUCTION PLANT

In green ammonia configurations, the Haber-Bosch synthesis loop remains the central component of the plant; yet, the conventional grey-hydrogen production train is entirely replaced by the SOE section for renewable hydrogen sourcing. The hydrogen produced from the SOE is inherently carbon-free; however, it contains residual water, shifting the process focus from carbon removal to effective dehydration, as water is a strong poison for the ammonia synthesis catalyst. Therefore, water is first removed immediately downstream of the SOE through cooling to ambient temperature and condensation. Additional moisture is eliminated during interstage cooling of the hydrogen compressor, and final ppm-level purification is achieved with a molecular sieve dryer placed at the compressor outlet. A further distinction from conventional plants is that only hydrogen, nitrogen, and ammonia circulate within the synthesis loop. This feature eliminates the need for a purge stream, which is typically mandatory in grey ammonia plants due to the presence of inert species such as methane or argon, avoiding syngas losses.

The ammonia synthesis loop consists of three main sections: compression, reaction, and separation. In the compression section, hydrogen from the SOE and nitrogen from the Air Separation Unit (ASU) are independently compressed to high pressure (*i.e.*, the operating pressure of the entire loop) and subsequently mixed to form the ammonia syngas. The reaction section is centered on an ammonia converter operating at high temperature (350-550 °C) and high pressure (150-250 bar). Here, ammonia is formed over an iron-based catalyst, following Temkin-Pyzhev kinetics, and exhibits a relatively low per-pass conversion [7]. Finally, in the separation section, the ammonia product is removed from the unconverted gases through a two-step condensation sequence: first using cooling water and then via a mechanical refrigeration cycle that uses ammonia as refrigerant. The unreacted hydrogen and nitrogen are recycled to the reactor inlet.

Finally, the liquid ammonia product is obtained at ca. 8 bar and 15-20 °C, with a purity higher than 96.5 mol%.

Energy Integration Strategy

An energy integration strategy is proposed to enhance the energy efficiency of the SOE-based hydrogen production section by using steam produced in the ammonia synthesis loop.

The effluent from the ammonia reactor exits at high temperature, providing a substantial thermal source. Although a portion of this heat is inherently used to preheat the incoming ammonia syngas, the present work aims to recover as much of the remaining thermal potential as possible. To this end, a Heat Recovery section (HRS-2) is installed downstream of the reactor, as shown in **Figure 3**.

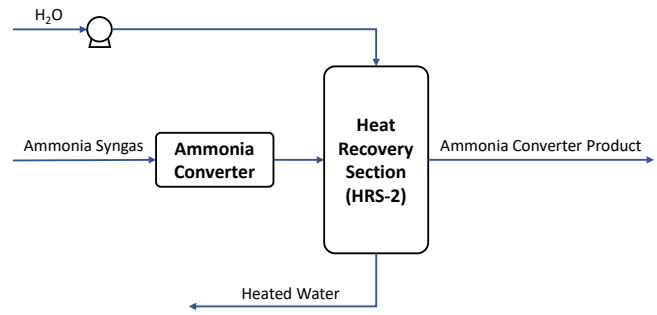

Figure 3. Schematic diagram of the ammonia synthesis section and the 2nd heat recovery section (HRS-2).

The introduction of the HRS-2 results in a significantly lower temperature of the product stream leaving the reaction section (*i.e.*, a drop of around 260 °C), which in turn decreases the thermal duty satisfied by the cooling water in the first ammonia condensation stage.

Integration Mechanism in the SOE BOP

Heat recovered in the first recovery section of the SOE unit (HRS-1) and in the second recovery section (HRS-2) downstream of the NH_3 converter are combined to generate all superheated steam requirements of the SOE section. Although the total steam demand of the SOE remains unchanged, a fraction is now provided by the ammonia synthesis loop. On the other hand, the remaining portion is still produced internally within the SOE Balance of Plant from liquid water under ambient conditions.

This modification substantially decreases the electric heating duty, otherwise required to raise the steam to the SOE operating temperature, thereby reducing the total energy consumption of the section and improving the overall energy efficiency of the hydrogen production system.

PERFORMANCE ANALYSIS OF SOE TECHNOLOGY AND ENERGY INTEGRATION

The performance of the SOE system was evaluated using the developed 0D electrochemical model and process simulations, examining operational behavior at the cell, stack, and system levels, as well as energy efficiency enhancements resulting from thermal integration with the ammonia synthesis loop.

Operational Modelling of SOE Technologies

Three literature-reported SOEC architectures (*i.e.*, Anghilante [5], Wendel-Braun [8], and Becker [9]) were simulated using the developed electrochemical model. Their geometric characteristics, operating specifications,

and permissible operating ranges are those reported in the aforementioned literature sources.

The behavior of a single cell is evaluated with the first model, with comparisons across the three models carried out under identical operating conditions: atmospheric pressure, 800 °C, and thermoneutral operation. This ensures that differences in operational outputs reflect intrinsic cell properties. At the cell level, results indicate that hydrogen productivity, steam conversion, and power demand strongly depend on intrinsic cell features, including the required inlet hydrogen molar fraction and the magnitude of internal irreversibility (described by the ASR). Conversely, energy efficiency is identical at 96% for all three cathode-supported cell types, as it is determined by imposing the thermoneutral operating regime.

The analysis is then extended to the stack level, following two approaches. In the first approach, each stack consists of 75 cells electrically connected in series. This approach quantifies the reduction in per-cell hydrogen productivity when cells operate as part of a stack. The additional irreversibility arising at the stack scale reduces individual cell productivity from 21% (*Anghilante SOEC*) to 29% (*Becker SOEC*). In the second approach, the electrolyzer's hydrogen production rate is fixed at 55 kg/h. This target value was selected to maintain the required electrical power within the limits reported in scientific literature for each SOE technology and to comply with industrial-scale applications below 10 MW. Under equal hydrogen output, the total electrical demand of the three electrolyzers is nearly identical (within 2% discrepancy) and slightly above 2.1 MW, resulting in comparable performance of the three SOE technologies under thermoneutral operation conditions.

Lastly, the analysis is extended to the system level, combining each SOE with its corresponding Balance of Plant through steady-state UniSim simulations, referring to hydrogen production at 55 kg/h. Since the electrolyzer power demand is almost the same across the three technologies, the focus shifts to the BOP power requirements. Overall, the BOP accounts for between 17% (*Anghilante SOE BOP*) and 20% (*Wendel-Braun SOE BOP*) of the total electrical consumption of the hydrogen production section. The dominant BOP contributor is the electric heater for steam superheating, representing 57% (*Becker SOE*) to 70% (*Wendel-Braun SOE*) of the BOP demand, making it the primary target for process energy integration. Considering both SOE and BOP contributions, the hydrogen production section exhibits a total electrical demand ranging from 2.5 MW (*Anghilante SOE system*) to 2.7 MW (*Wendel-Braun SOE system*). This corresponds to energy efficiency ranging from 68% (*Wendel-Braun SOE system*) to 73% (*Anghilante SOE system*). The resulting specific energy consumption of the electrolysis unit (defined as the total electrical energy required to produce one kilogram of hydrogen) ranges from 45.9 kWh/kg (*Anghilante SOE system*) to 49 kWh/kg (*Wendel-Braun SOE system*).

In this work, the SOEC model reported in Anghilante [5] was selected for subsequent integration with the ammonia production section.

Impact of Energy Integration on Plant-Level Performance

As mentioned earlier, the proposed energy integration strategy is based on two Heat Recovery Sections (HRS), the first in the SOE unit (HRS-1) and the second downstream of the ammonia reactor (HRS-2). This process integration leads to an 84% reduction in the electrical power required for steam superheating. This results in a substantial energy benefit for the renewable hydrogen production section, reducing its total energy demand by 8%. Correspondingly, the section's energy efficiency increases from 73% to 80%, and the specific energy consumption of the electrolysis unit ratio decreases by 8%, from 45.9 to 42.1 kWh/kg.

Table 3: Comparative Plant Performance Metrics (Integrated vs. Non-Integrated Scenarios) at a production capacity of 55 kg/h of hydrogen and 300 kg/h of ammonia.

Performance Metric	Case 1: Non Integrated	Case 2: Energy Integrated	$\Delta_{(2)-(1)}$%
SOE section: Electrical demand of ELH [kW]	259	42	-84%
SOE section: Energy efficiency [%]	73	80	+7%
Overall plant: Energy efficiency [%]	57	62	+5%
Specific energy consumption of electrolysis unit [kWh/kg$_{H_2}$]	45.9	42.1	-8%
Specific energy consumption of the plant [kWh/kg$_{NH_3}$]	9.13	8.32	-9%

The reduced energy demand in the SOE section also positively impacts the overall plant performance. The energy efficiency of the entire green ammonia plant rises from 57% to 62%; thus, the specific energy consumption of the power-to-NH_3 plant decreases by 9%, from 9.13 to 8.32 kWh/kg, confirming the effectiveness of the proposed integration strategy. A summary of the key performance enhancements achieved through the thermal integration is provided in **Table 3**.

TECHNO-ECONOMIC IMPACT OF THE ENERGY INTEGRATION

The techno-economic assessment provides a preliminary cost evaluation (AACE Class 4: Feasibility Estimate) aimed at quantifying the economic impact of the proposed energy integration strategy on the green ammonia production plant. This analysis, based on the UniSim-simulated process configurations, focuses on the estimation of capital and operating expenditures, acknowledging the inherent uncertainty (-30% to +50%) associated with Class 4 evaluations [10].

Capital Expenditure (CapEx) Analysis

Capital expenditure was estimated using the Module Costing Technique, with equipment costs updated to 2023 using the Chemical Engineering Plant Cost Index (CEPCI). The solid oxide electrolyzer is the dominant cost component, representing approximately 58% of the total plant CapEx with a specific investment of 2, 268 USD/kW [6]. Since the integration strategy does not modify the SOE operation or size, its cost remains unchanged across the two plant configurations.

The primary CapEx effect of the integration arises from the electric heater used for steam superheating. Because its capital cost scales directly with the electrical duty, the 84% reduction in power requirement results in a proportional decrease in heater cost, lowering the hydrogen section BOP capital cost by 52%.

Conversely, in the ammonia synthesis section, the required addition of HRS-2 increases the CapEx of the reaction section by approximately 13%. Considering both sections, the opposing contributions nearly compensate: the total plant CapEx varies only slightly (+0.4%), indicating that the proposed energy integration strategy has no significant impact on the overall investment cost.

Operating Expenditure (OpEx) Analysis

The operational expenditure assessment includes maintenance, raw materials, utilities, and electricity costs. Annual maintenance is assumed to be 2% of the corresponding CapEx, including a stack replacement cost set at 18.5% of the SOE CapEx with a five-year interval [6]. With the proposed energy integration strategy, raw material (demineralized water) consumption remains essentially unchanged. Utility costs are impacted: the lower reaction section effluent temperature in the integrated configuration reduces the cooling water demand of the first ammonia condenser by 81%, substantially lowering this cost component.

Electricity constitutes the dominant contribution to overall OpEx, accounting for approximately 77% of total operating costs. A continuous supply of renewable power is assumed through a Power Purchase Agreement (PPA), and electricity prices are represented by the Value Adjusted Levelized Cost of Electricity (VALCOE), which we evaluated for four countries (the United States, the European Union, China, and India) to capture geographical variability [11]. The energy integration strategy enables a decrease in electricity-related operating costs of approximately 7% across all countries considered, corresponding to an overall reduction of about 8% in total plant OpEx.

Levelized Cost Assessment

As a final step, we evaluated the Levelized Cost of green Hydrogen (LCOH) and the Levelized Cost of green Ammonia (LCOA). In the base configuration (without energy integration), the average LCOH is 4.55 USD/kg, decreasing to 4.31 USD/kg in the integrated configuration (*i.e.*, a reduction of approximately 5-6%). A similar trend is observed for LCOA, which decreases from an average value of 1, 024 USD/t to 979 USD/t (*i.e.*, a reduction of about 4-5%).

The strong dominance of renewable electricity cost, which alone accounts for more than half of the total ammonia production cost, ensures that the improvements achieved through energy integration directly translate into lower levelized costs.

The cost variability across regions is significant, mainly driven by the VALCOE. For the LCOH, China (reporting the highest VALCOE at 70 USD/MWh) yields a cost 9% higher than the calculated average value, while India (having the lowest VALCOE, at 50 USD/MWh) yields an LCOH 11% lower. An analogue geographical impact is observed for the LCOA: costs in China are approximately 8% higher than the average value, whereas costs in India are nearly 10% lower.

Overall, these results confirm that the proposed energy integration strategy not only enhances process efficiency but also yields non-negligible economic benefits across diverse geographic markets by mitigating the cost volatility associated with the primary operational expenditure driver (*i.e.*, the electric energy input).

CONCLUSIONS

We proposed and techno-economically assessed an energy integration strategy between a solid oxide electrolyzer for green hydrogen production and an ammonia synthesis loop. Our approach, which consisted of improving the heat integration of the plant with two heat recovery sections, one in the SOE unit and one downstream of the ammonia converter, proved highly effective. It resulted in a significant reduction (*i.e.*, 84%) in electric-heater duty. This modification successfully shifts the SOE system thermal load, resulting in a 7% increase in the SOE section's energy efficiency (from 73% to 80%) and a corresponding enhancement in the overall plant efficiency (from 57% to 62%).

At the economic level, the integration yields non-negligible economic benefits. The reduction in electrical demand translates directly into a decrease in the LCOH and LCOA, with average savings of 5-6% for LCOH and 4-5% for LCOA.

Lastly, this study is currently subject to several limitations, notably the assumption of steady-state operation (*i.e.*, continuous renewable electrical power input was considered) and the inherent uncertainty associated with a Class 4 cost assessment. However, despite these limitations, the results clearly indicate that solid oxide electrolysis, when appropriately integrated with high-temperature chemical processes, such as the Haber-Bosch loop, offers a highly promising route for efficient green ammonia production.

AUTHOR IDENTIFIERS

Author ORCIDs:
Isella A: 0000-0002-3699-4698
Manca D: 0000-0003-2055-9752

REFERENCES

1. IRENA and AEA. Innovation Outlook: Renewable Ammonia. International Renewable Energy Agency and Ammonia Energy Association (2022)
2. Isella A, Manca D. GHG emissions by (petro)chemical processes and decarbonization priorities—a review. Energies 15:7560 (2022). https://doi.org/10.3390/en15207560
3. Gambou F, Guilbert D, Zasadzinski M, Rafaralahy H. A comprehensive survey of alkaline electrolyzer modeling: electrical domain and specific electrolyte conductivity. Energies 15:3452 (2022). https://doi.org/10.3390/en15093452
4. Flis G, Wakim G. Solid oxide electrolysis: A technology status assessment. Clean Air Task Force (2023)
5. Anghilante R. Flexibilization and integration of solid oxide electrolysis units in power to synthetic natural gas plants. PhD Thesis, Institut National Polytechnique de Toulouse - INPT (2020)
6. Nami H, Rizvandi OB, Chatzichristodoulou C, Hendriksen PV, Frandsen HL. Techno-economic analysis of current and emerging electrolysis technologies for green hydrogen production. Energy Conversion and Management 269:116162 (2022). https://doi.org/10.1016/j.enconman.2022.116162
7. Nami H, Hendriksen PV, Frandsen HL. Green ammonia production using current and emerging electrolysis technologies. Renewable and Sustainable Energy Reviews 199:114517 (2024). https://doi.org/10.1016/j.rser.2024.114517
8. Wendel C, Braun R. Modeling and design of a novel solid oxide flow battery system for grid-energy storage. 10th European SOFC Forum. Lucerne, Switzerland (2012)
9. Becker WL, Braun RJ, Penev M, Melaina M. Production of fischer–tropsch liquid fuels from high temperature solid oxide co-electrolysis units. Energy 47:99-115 (2012). https://doi.org/10.1016/j.energy.2012.08.047
10. Turton R, Bailie RC, Whiting WB, Shaeiwitz JA. Analysis, Synthesis, and Design of Chemical Processes (3rd ed.). Pearson (2008)
11. IEA. World Energy Outlook 2024. International Energy Agency (2024)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Techno-Economic Assessment of Decarbonization Pathways for Methanol and Formaldehyde Production: A Superstructure Optimization Approach

Rafailia Mitraki[a]*, Muhammad Salman[a], Grégoire Léonard[a]

[a] University of Liège, Department of Chemical Engineering, Sart Tilman, Liège, Belgium.
* Corresponding Author: rafailia.mitraki@uliege.be.

ABSTRACT

This study aims to compare different pathways for achieving CO_2 emission reductions during the production of methanol and, subsequently, formaldehyde, i.e., its major derivative. An equation-oriented model of the formaldehyde sector is developed, incorporating a superstructure of various transition pathways including feedstock switching (biomass, biogas, waste), process electrification (Power-to-X), and CO_2 capture. The OSMOSE tool is used to evaluate the superstructure and compare the alternative production pathways on the basis of thermodynamic, environmental, and economic key performance indicators for future scenarios (2025 and 2050). Furthermore, to cope with the limitations of predefined pricing scenarios, a parameter sweep is performed, exploring a broader set of economic conditions and seeking to identify the zones of economic optimality associated with each configuration through the solving of a Mixed-Integer Linear Programming cost minimization problem, while generalizing the analysis results beyond a specific year or geographic context. Finally, the potential of industrial symbiosis implementation, between clinker and formaldehyde production facilities, as a cost-reduction measure for Power-to-X processes is assessed.

Keywords: Formaldehyde, Methanol, Energy Transition, Technoeconomic Analysis, Superstructure Optimization, Industrial Symbiosis

INTRODUCTION

The chemical industry is an energy-intensive and highly-emitting sector, representing 6.5% of the global greenhouse gases emissions [1]. Methanol (MeOH) is a key building block for solvents, plastics, and fuels. Its major application is its conversion to formaldehyde which accounted for 23% of global MeOH demand in 2023 [2]. Formaldehyde, of chemical composition CH_2O, is an organic chemical used as a raw material for the production of resins, polyurethanes, paints and coatings, pharmaceuticals, food intermediates, and chelating agents. It is produced from MeOH via two main routes: the silver process and the metal oxide process [3]. In the silver process, operated at high temperature (580–680°C) and under air deficiency, MeOH is converted into formaldehyde through catalytic dehydration ($CH_3OH \rightleftharpoons CH_2O + H_2$) using a silver-based catalyst. On the other hand, the metal oxide process consists of a direct catalytic MeOH oxidation ($CH_3OH + 0.5\,O_2 \rightarrow CH_2O + H_2O$), performed over a metal oxide catalyst (e.g. iron, molybdenum, or vanadium oxide). The reaction is implemented at medium temperatures (270–430°C) and in the presence of excess air, enabling very high MeOH conversion around 99%. Although European formaldehyde production in 2017 was almost evenly split between the silver and metal oxide processes, this study focuses on the latter given that recent investments tend to favor the metal oxide route [3].

Direct process emissions during formaldehyde production, stemming from MeOH or formaldehyde over-oxidation, are negligeable [3]. The environmental impact of formaldehyde is thus closely linked to the carbon footprint of MeOH feedstock. However, current methanol production is almost entirely fossil-based. Indeed, MeOH is predominately produced from natural gas with steam methane reforming (SMR, ~65%) or coal gasification (~35%, mainly in China), while biomass and other renewable feedstocks account for less than 1% of global production [4]. This results in a high carbon footprint ranging from 0.8 up to more than 4 t CO_2/t MeOH produced [5].

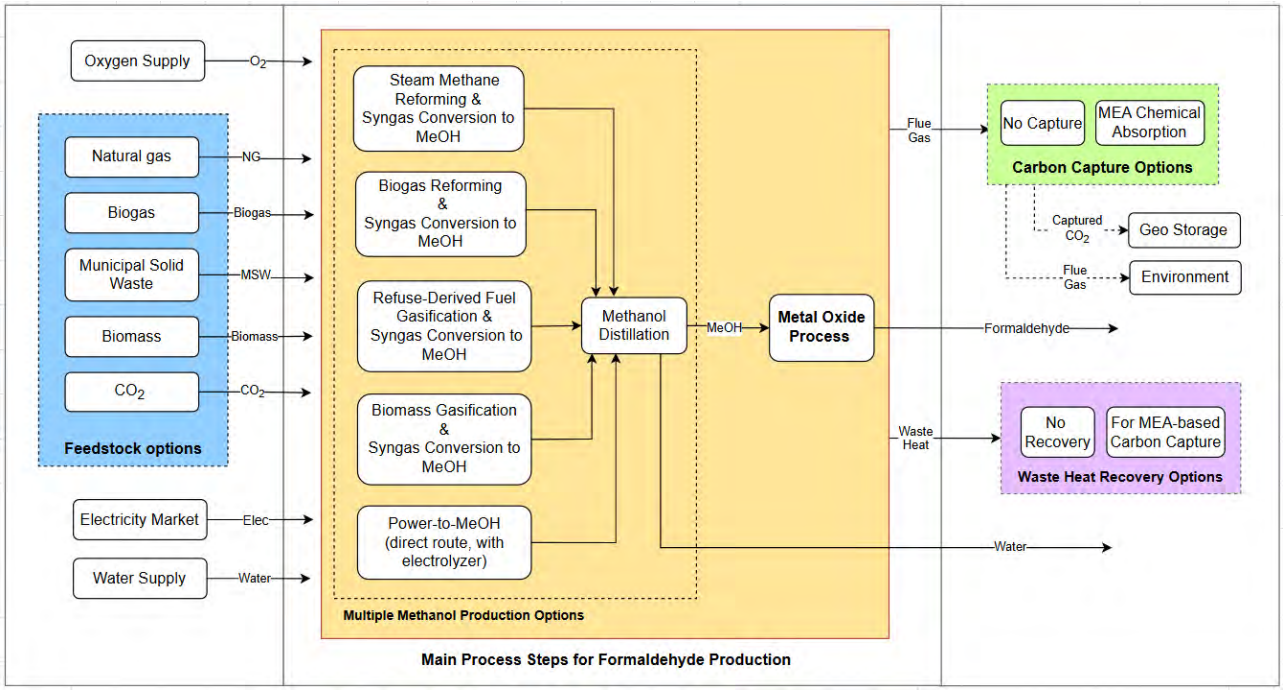

Figure 1. Superstructure for formaldehyde industry.

Hence, decarbonizing MeOH production through biomass use, CO_2 capture (CC), and process electrification [6] is essential for reducing the environmental impact of the European chemical sector.

METHODOLOGY

This study aims to compare various pathways for achieving CO_2 emission reduction during the production of methanol and, subsequently, formaldehyde, i.e. its major derivative. The methodological framework is adopted from our previous study [7]. Thorough literature review and data collection, from academic papers, institutional reports, and Best Available Techniques reference document [3], are first carried out, aiming at gathering information on the production process under study and applicable emission mitigation pathways. These pathways include energy efficiency improvement, feedstock switching, process electrification, and CO_2 capture and storage.

The collected data are used to develop an equation-oriented model of the formaldehyde sector, written in Lua language, consisting of mass and energy balances, as well as economic considerations such as annualized investment (CAPEX) and operating (OPEX) costs. This model incorporates a superstructure (depicted in **Figure 1**) of various MeOH production routes, including synthesis from syngas generated via (1) conventional Steam Methane Reforming (base case, *NG*), (2) biomass gasification (*BM*), (3) biogas tri-reforming (*BG*), (4) refuse-derived fuel (RDF, from municipal solid waste (*MSW*)) gasification, as well as (5) power-to-MeOH (*P2X*) process converting CO_2 and hydrogen (H_2, produced by alkaline electrolysis) into e-methanol. CO_2 in flue gases, from MeOH or formaldehyde production, may be emitted to the environment or captured through monoethanolamine (MEA)-based chemical absorption, using waste heat from the metal oxide process. Configurations abbreviations, process descriptions, and modeling details can be found in the Supplementary Sections **S1** and **S2**.

The production routes previously mentioned are compared based on different key performance indicators (KPIs) including specific energy demand (GJ/t CH_2O), CO_2 emissions (t CO_2/t CH_2O), and total cost (€/t CH_2O). Specific energy demand encompasses fuel and electricity requirements for an easier comparison across configurations. Specific total emissions include direct and indirect emissions. Direct, or scope 1, CO_2 emissions stem from product/feedstock over-oxidation or fuel combustion during formaldehyde or MeOH production. Biogenic CO_2 does not contribute to scope 1 emissions, while captured biogenic CO_2 is accounted as negative emission. Indirect emissions arise from feedstock supply and electricity production. CO_2 emission factors are listed in the Supplementary Section **S3**. Finally, the specific total cost (STC) encompasses both annualized CAPEX (for a greenfield scenario) and OPEX considerations including feedstock, fuel, electricity, CO_2 emission (no free allowances or revenues for negative emissions), as well as transport and storage (T&S) of captured CO_2.

The OSMOSE tool, developed by the IPESE group of EPFL [8], is used to evaluate the superstructure and a Mixed-Integer Linear Programming (MILP) problem is

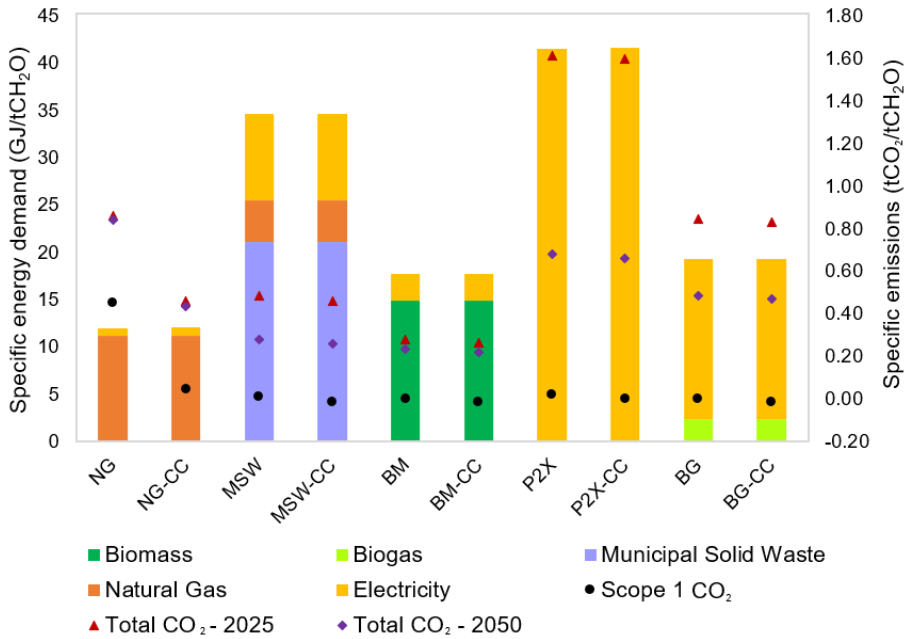

Figure 2. Specific energy demand (left axis) and scope 1 and total CO_2 emissions (right axis) in 2025 (electricity is sourced from the current Belgian grid) and 2050 (electricity is supplied by the 2050 European electricity mix) scenarios for different production routes.

solved to identify the optimal decarbonization configurations for two different energy and emissions scenarios: (1) the 2025 scenario, considering current energy and CO_2 prices in Belgium, and (2) the 2050 scenario, based on estimated future marginal production costs, considering higher ETS and lower energy prices [9-13]. Economic data for both scenarios are listed in **Table 1**.

Table 1: Economic parameters considered for 2025 and 2050 scenarios.

Scenario	2025	2050
Electricity (€/MWh)	78	58
Natural gas (€/MWh)	36	40
Biomass (€/MWh)	17	18
Biogas (€/MWh)	40	30
MSW (€/MWh)	9	9
ETS (€/tCO_2)	80	200
CO_2 T&S (€/tCO_2)	50	50
CO_2 supply (€/tCO_2)	80	80

The high volatility of energy prices—stemming from changing global political landscape, availability limitations, and regulatory shifts—prevents the accurate identification of future optimal production routes. To cope with the limitations of predefined pricing scenarios, a parameter sweep is performed, exploring a broader set of economic conditions and seeking to identify the zones of economic optimality associated with each configuration, while generalizing the analysis results beyond a specific year or geographic context. The price of natural gas (p_{NG}), electricity (p_E), municipal solid waste (p_{MSW}), biomass (p_{BM}), biogas (p_{BG}), and CO_2 supply (p_{CO2}), as well as the ETS price (p_{ETS}) are varied within the pricing sets specified in **Table 2**. For each of the 30, 000 resulting combinations, MILP optimization is performed, seeking total cost minimization. The results are analyzed to identify robust decarbonization pathways.

Table 2: Pricing sets considered for the parameter sweep.

Parameter	Set of values
p_{NG}	[10, 25, 50, 75, 100] €/MWh
p_E	[10, 25, 50, 75, 100] €/MWh
p_{MSW}	[-5, 0, 5] €/MWh
p_{BM}	[10, 25, 50, 75, 100] €/MWh
p_{BG}	[10, 25, 50, 75, 100] €/MWh
p_{CO2}	[0, 25, 75, 100] €/t
p_{ETS}	[75, 100, 200, 300] €/t

RESULTS

Energy demand and CO_2 emissions reduction

From **Figure 2**, total energy requirement (including fuel and electricity, and excluding feedstock) of conventional formaldehyde production (NG) amounts to 11.96 GJ/t CH_2O. Feedstock switching options result in higher total energy requirements: +61% for BG and +248% for

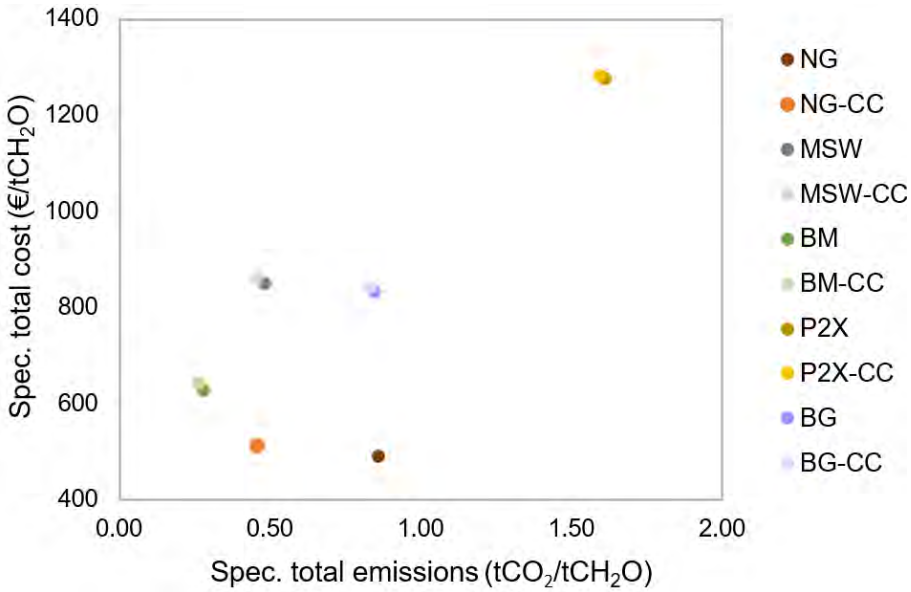

Figure 3. Specific total cost (€/t CH_2O) vs. total CO_2 emissions (t CO_2/t CH_2O) – 2025 scenario.

P2X due to H_2 production via alkaline electrolysis; +48% for BM and +189% for MSW due to lower process efficiency (e.g., drying requirements, lower carbon conversion) and additional energy requirements during MSW pre-treatment. Direct CO_2 emissions from formaldehyde production (i.e., 0.018 t CO_2/t CH_2O generated by over-oxidation) or fossil CO_2 emitted by MeOH production can be captured through chemical absorption with MEA. Heat integration between formaldehyde production and CC unit completely covers the additional heat demand and reduces the energy penalty of CO_2 capture to 0.01–1.1% compared to counterparts without CC.

While feedstock switching decreases drastically scope 1 CO_2 emissions (96–100%) compared to NG (0.452 t CO_2/t CH_2O), the contribution of indirect emissions, stemming from upstream material extraction and electricity production, results in a shift in trends as can be seen in **Figure 2**. Total CO_2 emissions breakdown can be found in Supplementary Material (Figure **S3**).

Based on current emission factors in Belgium (2025 scenario), BG and P2X routes are associated with total emissions, respectively, only 1% lower and 88% higher than NG due to large electric demand. The least emitting options are BM-CC (-69%) and BM (-68%), followed by NG-CC and MSW-CC (-47%). On the other hand, clear environmental advantage appears for feedstock switching options when low-carbon electricity (e.g., from the projected 2050 European electricity grid) is considered. Carbon footprint reduction between 45% and 77% can be achieved with alternative feedstocks, making NG-CC less environmentally competitive (-49% compared to base case), while P2X corresponds to a modest emission mitigation of 23%. While a 47–49% reduction in CO_2 emissions can be achieved by implementing CC in natural gas-based facilities, capturing CO_2 from electrified, bio-, and waste-based processes represents little interest in terms of carbon footprint (1–3% additional reduction) since only a limited amount of CO_2, stemming from formaldehyde production, is treated.

Scenario-based economic assessment

Economics being an important decision variable for industrial energy transition, scenario-based economic assessments are depicted in **Figure 3** and **Figure 5**. Despite the environmental benefit of alternative feedstocks, the economically optimal configuration in the 2025 scenario is NG (€489/t CH_2O), closely followed by NG-CC, reducing total CO_2 emissions by 47% and increasing STC only by 5%. High specific CAPEX (204–399% increase compared to NG for P2X and BM-CC respectively, as can be seen in **Figure 4**), low CO_2 ETS price, and expensive electricity (€78/MWh_e vs. €36/MWh_{NG}) result in alternative production pathways being economically uncompetitive compared to fossil-based chemicals. Indeed, configurations relying on MSW achieve similar environmental performance as NG-CC with a 67% higher cost. The current Belgian electricity mix results in P2X configurations being the least attractive, both from an economic and environmental perspective. On the other hand, BM and BM-CC achieve drastic emission mitigation (68–69%) at a reasonable cost (+28–32% compared to NG).

In the 2050 scenario, relying on lower energy costs and higher ETS price, NG-CC reduces the STC by 9% compared to the base case (€553/t CH_2O). Access to decarbonized electricity improves the environmental ranking of electricity-intensive configurations such as MSW,

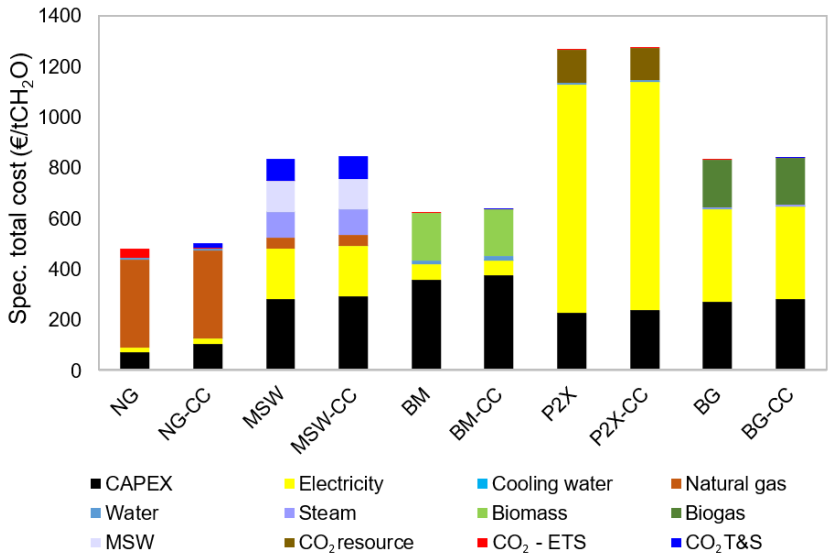

Figure 4. : Specific total cost breakdown – 2025 scenario.

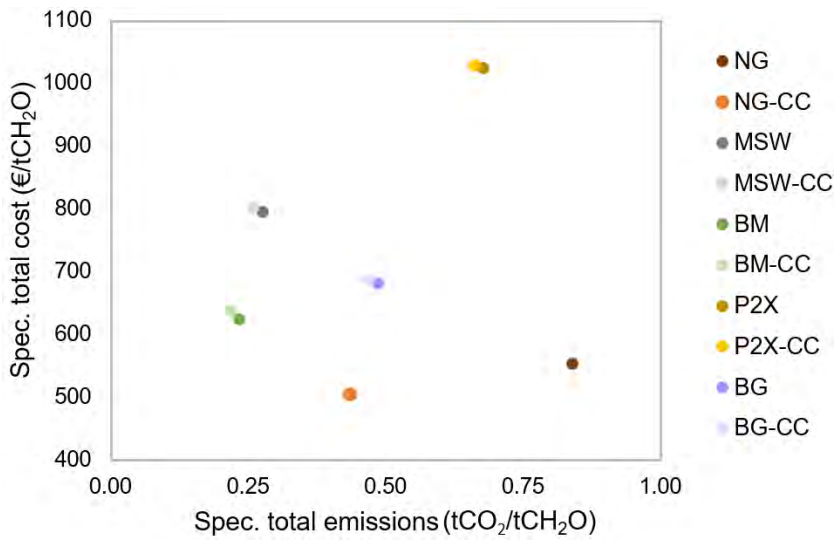

Figure 5. Specific total cost (€/t CH_2O) vs. total CO_2 emissions (t CO_2/t CH_2O) – 2050 scenario.

BG, and P2X. Biogas-based options benefit from cheaper electricity, reducing the cost gap with the base case down to only 23%, while production pathways relying on MSW are still 44% more expensive than NG. Despite lower emissions, formaldehyde produced from e-methanol feedstock remains the most expensive option with STC exceeding €1, 000/tCH_2O. Formaldehyde production from biomass represents the most promising non-fossil-based route, with a 13–15% increase in STC and 73–75% of emission reduction compared to NG.

Parameter sweep

The ranking of production pathways previously presented is valid solely for the commodities prices considered for the analysis. To cope with the high uncertainty of future energy and feedstock prices, a parameter sweep is performed, exploring a broad set of economic conditions. For each price combination, the configuration achieving the lowest production cost is identified. The results of the parameter sweep are represented in parallel coordinates plots (see **Figure 6** and Supplementary Section **S5**), enabling the identification of robust and flexible solutions for decarbonized formaldehyde production. As a general comment, CC implementation is favored over counterparts without CC for high ETS prices (≥€200/t CO_2). Moreover, as can be seen in **Figure 6**, natural gas price is a key parameter driving technology selection.

According to the performed analysis, SMR is the technology with the highest number of occurrences (14, 744), followed by biogas reforming (4, 884), MSW gasification (4, 837), and power-to-X (4, 071), while biomass gasification was found to be optimal only in 1, 464 times

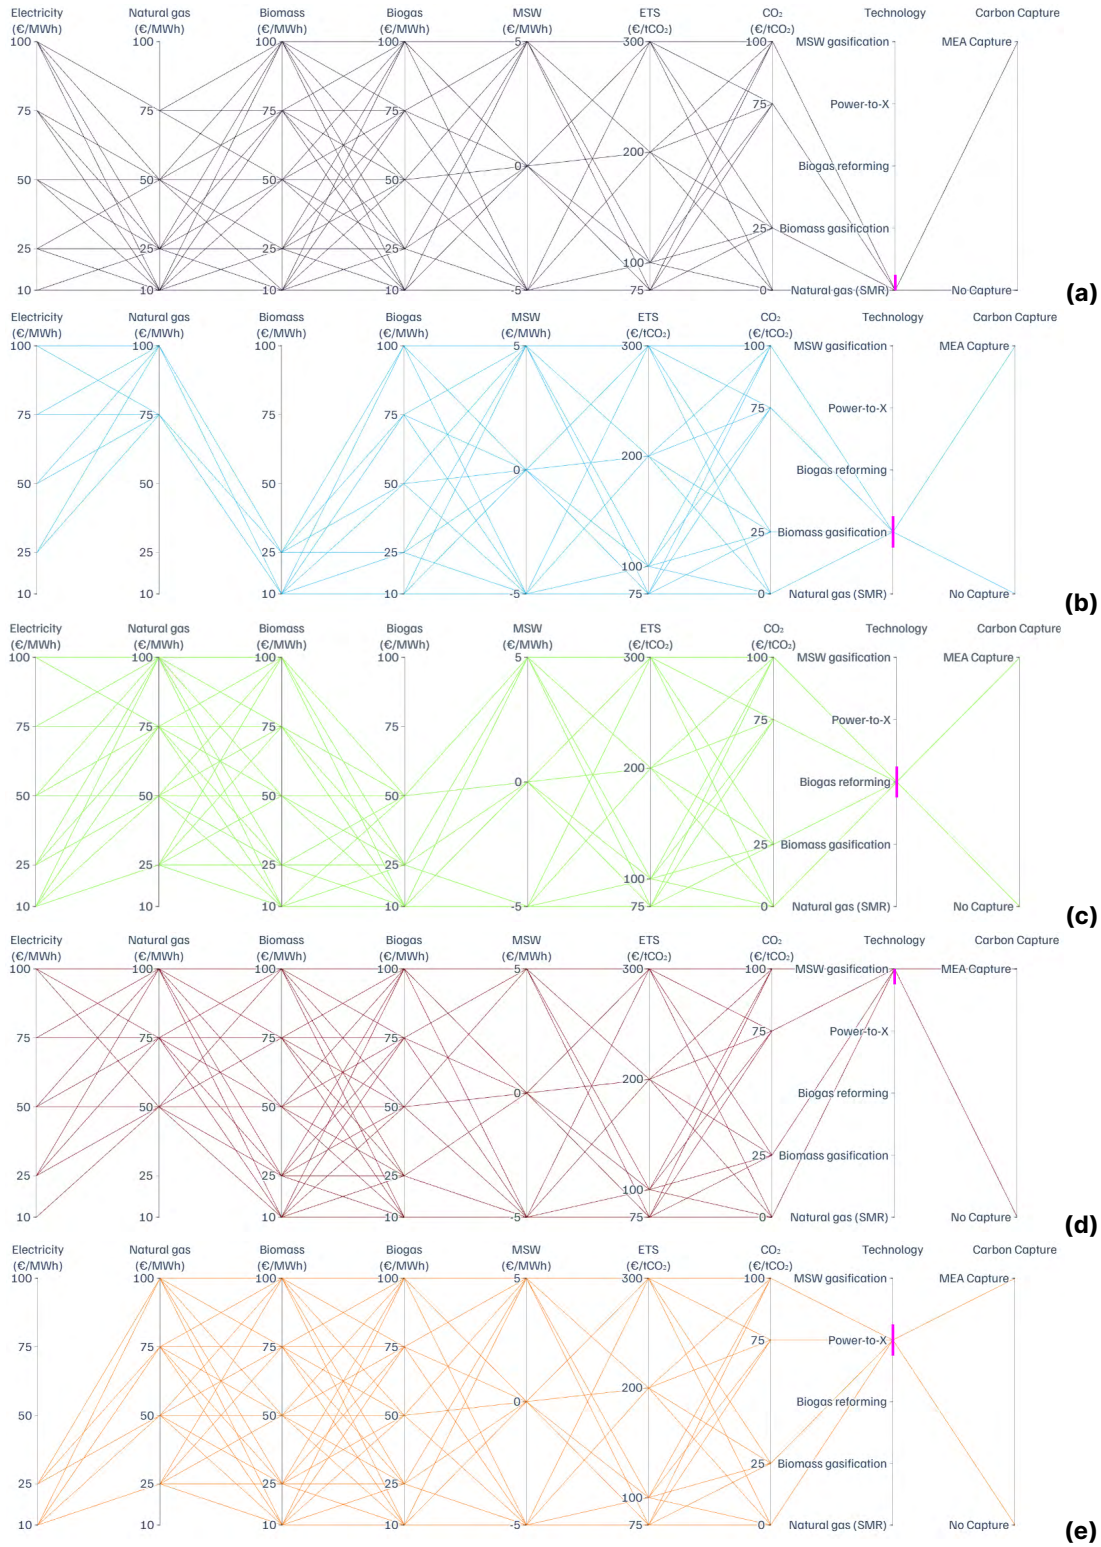

Figure 6: Solution space of economic optimality for (a) natural gas-, (b) biomass-, (c) biogas-, (d) MSW-based, and (e) electrified formaldehyde production routes.

out of 30, 000 scenarios tested.

Natural gas-based (**Figure 6**a) routes are identified as optimal for a wide range of economic conditions, making it a flexible and robust solution from an economic point of view. On the other hand, for BM and BM-CC (**Figure 6**b) to appear as economically optimal solutions,

market conditions must combine cheap biomass (≤€25/MWh) and expensive natural gas (≥€75/MWh). Moreover, electricity price should not be in the lower bracket (≥€25/MWh for a biomass price of €10/MWh; ≥€75/MWh when biomass is available at €25/MWh) to avoid favoring electricity-intensive configurations.

Biogas-based options (**Figure 6**c), which offer greater flexibility than biomass-dependent alternatives, become economically viable when natural gas prices exceed €25/MWh and €50/MWh for, respectively, BG-CC and BG. These options may remain optimal even at higher biogas prices (€50/MWh), provided that electricity prices are low to moderate (€25–50/MWh).

Economic optimality for P2X and P2XCC (**Figure 6**e) requires extremely cheap electricity (≤€25/MWh) combined with natural gas and biogas available at a higher price. Moreover, the lower the MSW price, the lower the CO_2 supply price should be to remain competitive with MSW-based options. On the contrary of other technologies, P2X-CC is always favored over P2X except when low ETS price (€75/tCO_2) is combined with free CO_2 feedstock.

Finally, a negative MSW price alone is not sufficient to ensure the economic optimality of MSW and MSW-CC configurations (**Figure 6**d). The availability of natural gas at a sufficiently high price (≥€50/MWh) is required.

As a conclusion, the parameter sweep revealed that NG and NG-CC are the most flexible and robust production pathways. This highlights the challenge of transitioning the chemical sector away from fossil feedstocks. However, biogas and waste-based technologies arise as interesting candidates for non-fossil MeOH and formaldehyde production.

Potential of industrial symbiosis for the production of e-chemicals

As previously mentioned, bio-based options arise as interesting candidates from an environmental and economic perspective for the energy transition in the chemical sector. However, availability issues are to be expected due to heavy reliance on biomass for industrial decarbonization [14]. Hence, to completely phase out chemicals production from fossil fuels, power-to-X processes should be included to some extent in the production mix. Nonetheless, the economic assessment reveals that e-MeOH feedstock represents the least cost-effective decarbonization pathway for formaldehyde production in a wide range of scenarios.

To address elevated costs, the potential of industrial symbiosis (IS) as a cost-reduction measure is investigated. IS [15] is a systemic approach that promotes collaboration among industries by resources exchanges, such as energy, water, and by-products, enhancing collective sustainability of industrial processes. In the present study, synergies between a formaldehyde production plant (0.1 Mt CH_2O/y) and nearby cement industry (0.9 Mt clinker/y) are considered. Commodities prices assumed for this analysis are listed in **Table 3**.

Table 3 : Economic parameters for IS implementation.

Commodities	Cost	
Coal	11.5	€/MWh
Municipal Solid Waste	3	€/MWh
Electricity	58	€/MWh
CO_2 ETS	200	€/tCO_2
CO_2 T&S	50	€/tCO_2

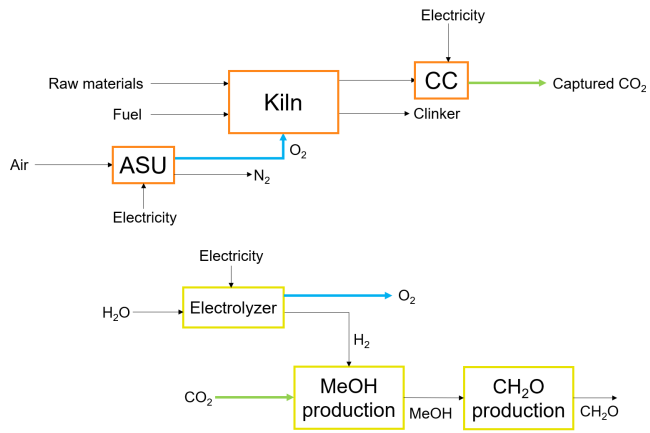

Figure 7. Case study for IS implementation between cement and formaldehyde industries.

The base case configuration of the system is depicted in **Figure 7**. In our previous study [16], oxy-fired kiln with cryogenic CO_2 capture arises as one of the most promising decarbonization options for the cement industry and hence is chosen as the production route for the IS case. In this configuration, heat is provided by burning a mixture of MSW and coal in oxygen produced by an Air Separation Unit (ASU). The generated CO_2 is subsequently captured via a cryogenic capture unit, before being transported and stored. On the other hand, the formaldehyde facility generates oxygen as a by-product of water electrolysis, which is necessary for hydrogen production. Meanwhile, CO_2 is required as a feedstock for MeOH production.

Competitiveness of clinker production and power-to-X process can be improved by mutualizing CO_2 capture unit and electrolyzer costs through CO_2 and oxygen exchanges, as highlighted in **Figure 7**. Such synergies result in an ASU that is no longer necessary, as the oxygen generated by the electrolyzer (0.17 Mt O_2/y) entirely covers the oxygen demand of the cement kiln (0.12 Mt O_2/y). Furthermore, the cement plant provides the entire CO_2 feedstock for methanol production (0.16 Mt CO_2/y), reducing the amount of CO_2 that needs to be transported and stored (0.59 Mt CO_2/y vs. 0.75 Mt CO_2/y without IS).

Removing the ASU from the system reduces total

electricity demand by 2% (-28 GWh_e/y), which represents 4.2 kt CO_2/y considering the current carbon intensity of Belgian electricity. Regarding the economics of the system, instead of buying CO_2 from the market at the ETS price (€200/t CO_2, representing the maximum marginal abatement cost industries are willing to bear, as well as the long-term target cost of CO_2 supplied by Direct Air Capture (DAC)), the formaldehyde facility is supplied by the cement plant at a much lower cost (i.e., capture cost of €27/t CO_2). Moreover, oxygen sales to the cement plant result in annual economies of about €11.5M. Despite the use of expensive oxygen (€91/t O_2 when produced via electrolysis vs. €30/t O_2 when produced with ASU), the cement plant benefits from CO_2 sales (€4.3M/y). Moreover, the reduced amount of CO_2 to be transported and stored represents additional economies of €8M/y. The implemented synergies and corresponding economies achieved are schematically represented in **Figure 8**.

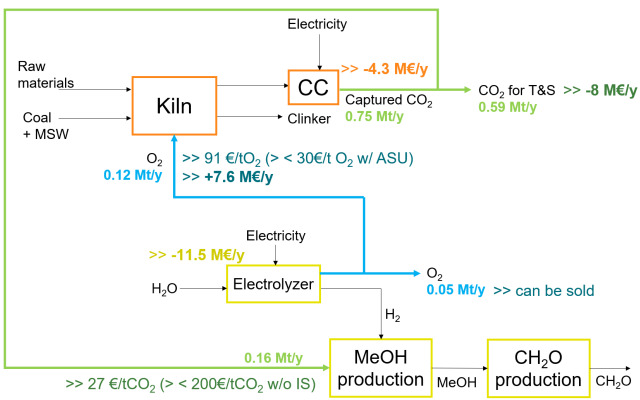

Figure 8. Synergies and corresponding economies achieved by IS implementation.

Hence, the implementation of synergies between these industries reduces the total production cost by 3% and by 32% for clinker and formaldehyde, respectively, compared to estimations without IS. Despite the decreased production cost in the IS case, P2X formaldehyde remains more expensive than other decarbonization options such as NG-CC and bio-based routes. Therefore, this option only makes economic sense when there is a scarcity of biomass, high natural gas prices, or low electricity prices. These conditions make P2X options difficult to appear as competitive decarbonization pathways for formaldehyde production.

Table 4: Specific total cost of clinker and formaldehyde with and without implementation of IS.

	No symbiosis	With symbiosis	
Clinker (€/$t_{clinker}$)	121	118	**-3%**
Formaldehyde (€/t_{CH2O})	1207	815	**-32%**
Total (M€/y)	229.8	187.5	**-18%**

CONCLUSION

This study provides a foundation for decision-making regarding CO_2 emission mitigation for formaldehyde production based on thermodynamic, environmental, and economic KPIs for future scenarios. For this purpose, an equation-oriented superstructure of formaldehyde industry was developed, encompassing various energy transition pathways including feedstock switching (biomass, biogas, waste), process electrification (Power-to-X), and CO_2 capture.

Despite requiring more energy (+ 48–247%) than the base case (11.96 GJ/tCH_2O), alternative feedstocks result in reductions of total CO_2 emissions up to 77% when low-carbon electricity is used, while emissions are cut by almost half when CC is implemented at NG-based facilities.

The economic assessment in the 2025 scenario identified NG as the most interesting configuration from an economic perspective, closely followed by NG-CC (5% cost increase, 47% emissions reduction compared to base case). Conversely, expensive and fossil-based electricity results in P2X being both the most expensive and most emitting configuration. The most promising bio-based options are BM and BM-CC: 28–32% cost increase, 65–67% emissions reduction compared to NG. If lower energy prices and higher ETS price are assumed (2050 scenario), NG-CC becomes the most cost-effective option, followed by biomass-based configurations, while electrification (P2X) remains the most expensive production route.

A parameter sweep, exploring a broad set of energy pricing scenarios, was also performed. An MILP problem was solved with the OSMOSE tool for each price combination to identify the solution space in which each configuration is associated with the lowest specific total cost. From the 30, 000 pricing scenarios tested, NG-based options (NG and NG-CC) were found to be the most flexible and robust production pathways, representing 49% of the optimal solutions identified. NG-CC exhibits reasonably good environmental performance, with emissions cuts reaching nearly half of current levels, despite its continued reliance on fossil fuels. The most interesting bio-based route is biogas reforming, appearing as optimal as long as moderate-to-high natural gas prices are combined with low-to-moderate electricity prices. On the contrary, stringent conditions for economic optimality arise for BM and P2X configurations requiring extremely low biomass and electricity prices respectively (≤25€/MWh).

Finally, the potential of industrial symbiosis as a cost-reduction measure for power-to-X process was assessed. Synergies between oxy-fired clinker production (requiring oxygen and capturing CO_2) and formaldehyde

production from e-MeOH (generating oxygen by electrolysis and requiring CO_2 as feedstock) result in reduced electricity demand (-2%) and reduced production costs (-3% for clinker and -32% for formaldehyde). However, P2X formaldehyde remains more expensive than other emission mitigation options, only making economic sense when biomass scarcity, high natural gas prices, and low electricity prices are observed.

These observations highlight the difficulty for the chemical sector to decarbonize and the importance of process efficiency improvements as well as policy frameworks and governmental funding to ensure competitiveness of low-carbon formaldehyde production routes.

DIGITAL SUPPLEMENTARY MATERIAL

Process descriptions, modeling details, and additional results can be accessed here: LAPSE:2026.0027

ACKNOWLEDGEMENTS

This paper is supported by the Belgian Fund for Scientific Research (FRS-FNRS) through a FRIA grant. The authors would like to acknowledge the Energy Transition Fund project 'TRILATE' organized by Belgium FPS economy. The authors also acknowledge François Maréchal for providing access to the OSMOSE tool.

AUTHOR IDENTIFIERS

Mitraki R.: 0009-0007-0990-0514
Salman M.: 0009-0001-5348-0701
Léonard G.: 0000-0003-2237-8306

REFERENCES

1. Isella A, Manca D. GHG emissions by (petro)chemical processes and decarbonization priorities—a review. Energies 15:7560 (2022). https://doi.org/10.3390/en15207560
2. Methanol Institute. https://methanol.org/methanol-price-supply-demand/index.html
3. Falcke H et al. Best available techniques (BAT) reference document for the production of large volume organic chemicals. Publications Office of the European Union (2017)
4. IRENA, Methanol Institute. Innovation outlook: Renewable methanol (2021)
5. Methanol Institute. Carbon footprint of methanol (2022)
6. Danish Energy Agency. Renewable fuels: Technology descriptions and projections for long-term energy system planning (2025)
7. Salman M, Flórez-Orrego D, Coppitters D, Mitraki R, Maréchal F, Léonard G. Decarbonising the glass industry: a comprehensive techno-economic assessment of low-emission pathways. Computers & Chemical Engineering 203:109329 (2025). https://doi.org/10.1016/j.compchemeng.2025.109329
8. Flórez-Orrego D, Ribeiro Domingos ME, Maréchal F. A systematic framework for the multi-time integration of industrial complexes and urban systems. In: Proceedings of the 7th International Conference on Contemporary Problems of Thermal Engineering (2022)
9. IEA. Outlook for Biogas and Biomethane: Prospects for Organic Growth (2020)
10. DG Energy. Quarterly Report on European Electricity Markets (2024)
11. Laguna JC, Moglianesi A, Vingerhoets P, Nijs W, Lodewijks P. PATHS 2050 – Scenarios Towards a Carbon-Neutral Belgium by 2050 (2023)
12. DG Energy. Quarterly Report on European Gas Market (2024)
13. Statista. https://www.statista.com/statistics/1322214/carbon-prices-european-union-emission-trading-scheme/
14. Lühmann M. Sustaining the european bioeconomy: the material base and extractive relations of a bio-based eu-economy. Bioeconomy and Global Inequalities :287-307 (2021). https://doi.org/10.1007/978-3-030-68944-5_14
15. Chertow MR. "uncovering" industrial symbiosis. J of Industrial Ecology 11:11-30 (2008). https://doi.org/10.1162/jiec.2007.1110
16. Salman M, Flórez-Orrego D, Maréchal F, Léonard G. Comparative assessment of fuel-switching options for the energy transition in the cement industry. In: 35th European Symposium on Computer Aided Process Engineering (ESCAPE 35): Book of Short Papers (2025)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Integrated Operating Strategies and Parameter Optimization for PEM Electrolyzers in Power-to-X Energy Systems

Luka Bornemann[a]*, Yifan Wang[b], and Martin Kaltschmitt[a]

[a] Hamburg University of Technology, Institute of Environmental Technology and Energy Economics, Hamburg, Germany
[b] RWTH Aachen University, Institute of Technical Thermodynamics, Aachen, Germany
* Corresponding Author: luka.bornemann@tuhh.de.

ABSTRACT

"Green" hydrogen production via polymer electrolyte membrane (PEM) electrolyzers must overcome significant energy penalties and high costs to become competitive in renewables-based energy systems. Adaptive operating strategies for PEM electrolyzers—by dynamically adjusting current density, pressure, and temperature—have demonstrated efficiency improvements in simple energy systems. However, their effectiveness in the context of complex power-to-X energy systems featuring variable downstream synthesis processes remains unclear. This work shows that integrated optimization of PEM electrolyzer operating parameters in conjunction with downstream methanation processes (MP) delivers substantial system-wide efficiency and cost benefits under dynamic hydrogen demand and pressure conditions. To demonstrate this, an equation-oriented process model of a PEM electrolysis system is embedded within a higher-level energy system model to compare sequential optimization (where the electrolyzer adapts to predetermined MP operating decisions) against integrated optimization (where electrolyzer and MP operating decisions are determined simultaneously). Sequential optimization delivers 7.3% operating cost savings compared to conventional fixed-parameter operation. In comparison, integrated optimization achieves 9.9% cost reductions and 7.8% decreases in electrolysis system electricity consumption. Operating pressure emerges as the most critical parameter, with electrolyzer electricity consumption exhibiting markedly higher sensitivity than MP electricity consumption. The analysis reveals fundamental trade-offs between component-level and system-level efficiency, demonstrating that prioritizing electrolyzer efficiency optimization yields superior system-wide performance. These findings establish that system-wide coordination through integrated optimization approaches substantially enhances the economic viability of "green" hydrogen supply in complex power-to-X energy systems.

Keywords: Scheduling, Energy Efficiency, Energy Systems, Hydrogen, Optimization

INTRODUCTION

Polymer electrolyte membrane (PEM) electrolyzers enable dynamic integration of hydrogen production into renewables-based energy systems [1]. However, energy-intensive operation and high production costs continue to impede the competitiveness and large-scale deployment of "green" hydrogen technologies [2].

Intelligent operating strategies offer substantial potential to improve both energy and cost efficiency of PEM electrolyzers [3]. Two key mechanisms enable these improvements: (1) dynamic adjustment of operating parameters (i.e., current density, pressure, temperature) to select the most energy-efficient operating point for each load, and (2) information exchange between the electrolyzer and downstream hydrogen applications (e.g., end-use pressure requirements), allowing the electrolyzer to adapt its operation accordingly. While these strategies may temporarily reduce electrolyzer stack efficiency, they yield net system-level gains by reducing hydrogen compression energy demand and lowering operating cost.

These benefits have been demonstrated primarily in relatively simple energy systems comprising photovoltaic units, a PEM electrolyzer, and hydrogen storage to meet a constant hydrogen demand [3]. In contrast, in broader power-to-X applications, PEM electrolyzers are coupled with various downstream synthesis processes (e.g., methanation, Fischer-Tropsch synthesis, Haber-Bosch synthesis), which impose time-varying hydrogen demands and variable pressure requirements [4]. Whether adaptive operating strategies can deliver comparable system-wide efficiency and cost benefits under these complex, dynamic conditions remains unclear.

This work addresses this research gap by extending a previously developed electrolyzer process model to account for downstream methanation process (MP) requirements with variable hydrogen demand and adjustable operating pressures. Two optimization approaches are applied: sequential optimization (where the electrolyzer adapts to predetermined methanation process operating decisions) and integrated optimization (where electrolyzer and methanation process operating decisions are determined simultaneously). The analysis investigates how flexible management of electrolyzer current density, pressure, and storage utilization affects hydrogen production, electricity consumption, and overall system operating costs. The specific contributions are threefold: (1) characterization of optimal electrolyzer operating strategies in integrated power-to-X energy systems with dynamic synthesis requirements, (2) quantitative comparison of sequential versus integrated optimization approaches, and (3) identification of trade-offs between component-level efficiency and system-level performance, providing insights into the practical potential of adaptive operating strategies in complex renewable hydrogen-based energy systems.

METHODOLOGICAL APPROACH

The optimal operating strategy of a PEM electrolyzer is determined through integrated energy system and process optimization. An equation-oriented process model of a PEM electrolysis system forms the methodological core, quantifying the effects of operating parameter selection (i.e., current density, temperature, pressure) on electricity consumption. This process model is embedded within a higher-level energy system model to derive an operational strategy that satisfies hydrogen demand from methanation processes, which impose both pressure and quantity requirements. The electrolyzer operating strategy encompasses time-variable hydrogen production and dynamic operating pressure adjustment.

Two optimization approaches are analyzed: sequential optimization, in which the electrolyzer adapts to predetermined methanation requirements, and integrated optimization, in which electrolyzer and methanation pressures are determined simultaneously to achieve system-level coordination.

Electrolysis system

A process model of a PEM electrolysis system [3] forms the methodological foundation. The system comprises the PEM electrolyzer itself and a multi-stage compressor. Throughout this work, 'electrolyzer' denotes the standalone electrolyzer unit, whereas 'electrolysis system' encompasses both the electrolyzer and the associated compressor.

The operational sequence begins with preheating of process water to the target stack operating temperature T_{stack}, followed by the electrolysis reaction within the cell. Hydrogen and oxygen, the reaction products, are separated through the membrane and subsequently directed to the positioned flash units. The separated hydrogen stream can be delivered through two outlets. The low-pressure (LP) outlet supplies hydrogen directly at the operating pressure p_{cat} for immediate use, while the high-pressure (HP) outlet enables compression to elevated pressure levels for storage and later use. Both outlets can operate simultaneously, allowing flexible allocation between direct supply and storage.

Electrical energy input occurs via the applied cell voltage U_{cell} and current I to the electrolyzer, as well as through compression work P_{comp} to the compressor. The electrolysis reaction generates waste heat $\dot{Q}_{waste}$, which can be recovered for preheating process water and external heat-integration applications. When the available waste heat proves insufficient for preheating requirements, supplementary heating P_{heat} must be provided by an electric heating element.

Cell voltage U_{cell} consists of three components: the open-circuit voltage U_{ocv}, activation overvoltage U_{act}, and ohmic overvoltage U_{res}. Reduced current densities j, decreased operating pressures p_{cat}, and increased temperatures T_{stack} contribute to lower cell voltage (Eq. (1) [3]).

$$U_{cell}(j, T_{stack}, p_{cat}) = U_{ocv}(T_{stack}, p_{cat}) + U_{act}(j, T_{stack}) + U_{res}(j, T_{stack}) \quad (1)$$

Aggregating all electrical input components yields the total electrical energy input to the electrolysis system (Eq. (2) [3]).

$$\dot{E}_{el}^{in} = U_{cell} I + P_{heat}(j, T_{stack}, p_{cat}) + P_{comp}(p_{cat}) \quad (2)$$

The electrolysis system efficiency η_{EL} (Eq. (3) [3]) quantifies the ratio of usable hydrogen energy to total electrical energy input. Hydrogen energy content is evaluated based on the lower heating value LHV_{H_2}. Diffusion and recombination losses $\dot{n}_{H_2}^{loss}$ reduce the total hydrogen production $\dot{n}_{H_2}^{tot}$. Beyond the previously discussed influences, the electrolysis system efficiency exhibits

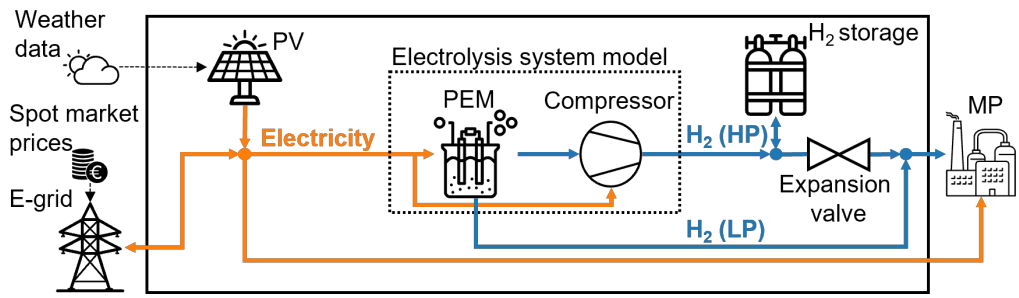

Figure 1. Energy system structure (E: Electricity, LP: Low-pressure, HP: High-pressure, MP: Methanation processes, PEM: Polymer electrolyte membrane electrolyzer, PV: Photovoltaic unit).

conflicting parameter dependencies: elevated pressures decrease compression work P_{comp}, whereas decreased pressures and elevated current densities minimize hydrogen losses $\dot{n}_{H_2}^{loss}$.

$$\eta_{EL} = \frac{LHV_{H_2}\left(\dot{n}_{H_2}^{tot}(j) - \dot{n}_{H_2}^{loss}(j, p_{cat})\right)}{\dot{E}_{el}^{in}} \tag{3}$$

Power-to-X energy system

The optimal electrolyzer operating strategy is analyzed within a higher-level energy system configuration (Figure 1). This energy system comprises a photovoltaic (PV) unit and a grid connection, enabling bidirectional electricity exchange (i.e., grid withdrawal during shortages and feed-in during PV surplus conditions), the electrolysis system and a hydrogen storage.

The primary objective of this energy system is to satisfy the electricity and hydrogen demands of two methanation processes. Hydrogen must be delivered at pressure levels between 1 and 20 bar (cf. Scenario definition). Direct supply via the low-pressure (LP) outlet without additional compression can be achieved by maintaining an electrolyzer operating pressure sufficiently elevated. Alternatively, hydrogen can be pressurized to 200 bar via the high-pressure (HP) outlet for temporary storage and subsequent withdrawal. This storage pathway provides operational flexibility to respond to time-variable electricity prices and fluctuating PV generation patterns. Depending on the optimization mode (cf. Scenario definition), direct hydrogen supply via the HP outlet is also possible.

Optimization problem

Determination of the electrolyzer operating strategy requires solving an energy system optimization problem (Eqs. (4) [3]). The objective function minimizes total operating costs C, representing the balance between electricity procurement expenditures and feed-in revenues (Eq. (4a)). Cost calculation incorporates electricity purchase prices $c_{el,t}^{buy}$, feed-in tariffs $c_{el,t}^{sell}$, and corresponding energy quantities $\dot{E}_{el,t}^{buy/sell}$, scaled by the time step duration Δt and the number of aggregated period occurrences n_t^p (i.e., scaling aggregated time series costs to annual values).

$$\min_{\mathbf{d},\dot{\mathbf{E}}} C = \Delta t \sum_{t\in\mathcal{T}} n_t^p \left(c_{el,t}^{buy}\dot{E}_{el,t}^{buy} - c_{el,t}^{sell}\dot{E}_{el,t}^{sell}\right) \tag{4a}$$

s.t.

$$\sum_{k\in\mathcal{K}} \left(\dot{E}_{k,e,t}^{out} - \dot{E}_{k,e,t}^{in}\right) + \dot{E}_{e,t}^{buy} - \dot{E}_{e,t}^{sell} = \dot{E}_{e,t}^{D}, \quad \forall e \in \mathcal{E}, \forall t \in \mathcal{T} \tag{4b}$$

$$\dot{E}_{PEM,\mathrm{H}_2,t}^{out} = \dot{n}_{\mathrm{H}_2,t} LHV_{\mathrm{H}_2}, \quad \forall t \in \mathcal{T} \tag{4c}$$

$$\dot{E}_{PEM,el,t}^{in} = U_{cell,t} I_t + P_{heat,t} + P_{comp,t}, \forall t \in \mathcal{T} \tag{4d}$$

$$\dot{E}_{MP,el,t}^{in} = \dot{E}_{MP,el,t}^{in,ref} + \frac{\Delta \dot{E}_{MP,el}^{in}\left(\dot{E}_{MP,\mathrm{H}_2,t}^{in}\right)}{\Delta p_{MP}}\left(p_{MP,t} - p_{MP,t}^{ref}\right), \quad \forall t \in \mathcal{T} \tag{4e}$$

$$p_{cat,t} \geq p_{MP,t}, \quad \forall t \in \mathcal{T} \tag{4f}$$

$$\mathbf{g}(\mathbf{d}, \dot{\mathbf{E}}) \leq \mathbf{0}, \tag{4g}$$

$$\mathbf{d} \in \mathcal{D} = \{T_{stack,t}, p_{cat,t}, j_t, p_{MP,t}\}, \ \forall t \in \mathcal{T} \tag{4h}$$

$$\dot{\mathbf{E}} \in \left\{\dot{E}_{k,e,t}^{out}, \dot{E}_{k,e,t}^{in}, \dot{E}_{el}^{buy}, \dot{E}_{el}^{sell}\right\}. \quad \forall k \in \mathcal{K}, \forall e \in \mathcal{E}, \forall t \in \mathcal{T} \tag{4g}$$

Energy balance constraints must be satisfied for all energy system components ($\forall k \in K$), energy carriers ($\forall e \in E$), and time steps ($\forall t \in T$) through input/output power flows $\dot{E}_{k,e,t}^{in/out}$, ensuring demand requirements $\dot{E}_{e,t}^{D}$ are met (Eq. (4b)).

Integration of the PEM electrolysis system into the energy system model occurs through its input/output power flows $\dot{E}_{PEM,H_2,t}^{in/out}$ (Eqs. (4c) and (4d)).

Methanation pressure variation (i.e., hydrogen demand pressure; hereafter termed methanation pressure for brevity) is governed by Eq. (4e). The methanation process operates at a characteristic curve, which is obtained through a pre-processing step (cf. System assumptions). This characteristic curve provides electricity demand changes $\Delta\dot{E}_{MP,el}^{in}$ in relation to methanation pressure changes Δp_{MP} for each hydrogen demand level. Within the optimization, time-dependent hydrogen demand $\dot{E}_{MP,H_2,t}^{in}$ serves as an exogenous input parameter, associated with the corresponding reference electricity

Table 2: Optimization modes (LP: Low-pressure, MP: Methanation process, EL: Electrolyzer).

Mode	EL LP outlet	MP pressure	EL temperature	EL pressure
Ref-seq		Fixed	80 °C	30 bar
Ref-int	✓	1 to 20 bar	80 °C	30 bar
pT-opt-seq	✓	Fixed	20 to 80 °C	1 to 30 bar
pT-opt-int	✓	1 to 20 bar	20 to 80 °C	1 to 30 bar

demand $\dot{E}^{in,ref}_{MP,el,t}$ and reference methanation pressure $p^{ref}_{MP,t}$. By applying linear interpolation, Eq. (4e) determines actual electricity demand through the adjustment of methanation pressure $p_{MP,t}$ relative to the reference pressure $p^{ref}_{MP,t}$.

Eq. (4f) ensures that electrolyzer pressure $p_{cat,t}$ is at least equal to the methanation pressure at the low-pressure outlet.

The constraint vector $\mathbf{g}$ encompasses all additional restrictions (e.g., energy storage capacities, load limitations, material balances). Decision variables include electrolyzer operating parameters (i.e., temperature $T_{stack,t}$, pressure $p_{cat,t}$, current density j_t), methanation pressure $p_{MP,t}$, and input/output power flows $\dot{\mathbf{E}}$ of all system components.

Solution of the resulting nonlinear programming optimization problem employs the open-source solver IPOPT, obtaining a local optimum. Identical initialization values are applied across all optimizations to ensure reliable, reproducible results and maintain consistency between different operating modes (cf. Scenario definition).

SYSTEM ASSUMPTIONS AND SCENARIO DEFINITION

System assumptions

Weather data from Hamburg, Germany, serves as time-dependent input for calculating PV system power availability [5]. Additionally, day-ahead electricity market prices are utilized, with both data sets based on the year 2023. Electricity prices comprise purchase prices and feed-in tariffs, with feed-in tariffs corresponding directly to day-ahead market prices [6]. Purchase prices include additional charges for non-household customers (i.e., grid fees and taxes) [7]. Hourly year-round data are aggregated into seven representative typical days (24 h each) using a hierarchical clustering method [8], yielding 168 time steps. To account for long-term storage, the storage formulation from [9] is applied.

Based on these temporal input data, pre-processing employs the energy system model and optimization approach from [4] to derive the methanation process operating strategy and generate an operating characteristic curve. The operating strategy provides time-dependent hydrogen $\dot{E}^{in}_{MP,\mathrm{H_2},t}$ and electricity demands $\dot{E}^{in,ref}_{MP,el,t}$, along with corresponding methanation pressures $p^{ref}_{MP,t}$, as exogenous input time series. The operating characteristic curve relates methanation pressure variations to electricity consumption across different load levels, enabling pressure adjustment according to Eq. (4e).

Table 1 presents the nominal capacities of all energy system components. The sizing of the PEM electrolysis system, PV unit and hydrogen storage reflect a cost-optimal hydrogen supply configuration [3], while methanation process capacities are derived from [4] with adjustments to match the electrolyzer capacity. The electrolyzer is assumed to operate at a maximum current density of 2 A/cm² at its nominal capacity of 1 MW [3].

Table 1: Nominal capacity of energy system components based on [3, 4].

Component	Nominal capacity
PEM electrolysis system	1 $\mathrm{MW_{H_2}}$
PV unit	2.5 MW
Hydrogen storage	49.3 MWh
Methanation processes	2x 0.45 $\mathrm{MW_{CH_4}}$

Scenario definition

Four optimization modes are defined to investigate the influence of integrated operating strategies (Table 2). These modes distinguish between the electrolyzer's intrinsic operational flexibility and its coordination with downstream methanation processes.

Regarding the electrolyzer, a distinction is made between fixed state-of-the-art operating parameters (i.e., 80 °C, 30 bar [3]; 'Ref' cases) and dynamically adjustable temperature and pressure within specified limits ('pT-opt' cases). The electrolyzer operating range spans 20 to 80 °C and 1 to 30 bar.

The optimization modes further differentiate the coordination approach. In the sequential approach ('seq'), the methanation process operating strategy (i.e., electricity demand, methanation pressure) is predetermined ($\dot{E}^{in}_{MP,el,t} = \dot{E}^{in,ref}_{MP,el,t}; p_{MP,t} = p^{ref}_{MP,t}$), and the electrolyzer subsequently determines its operating strategy accordingly. In the integrated approach ('int'), the electrolyzer operating strategy and methanation pressure (and consequently electricity demand) are determined

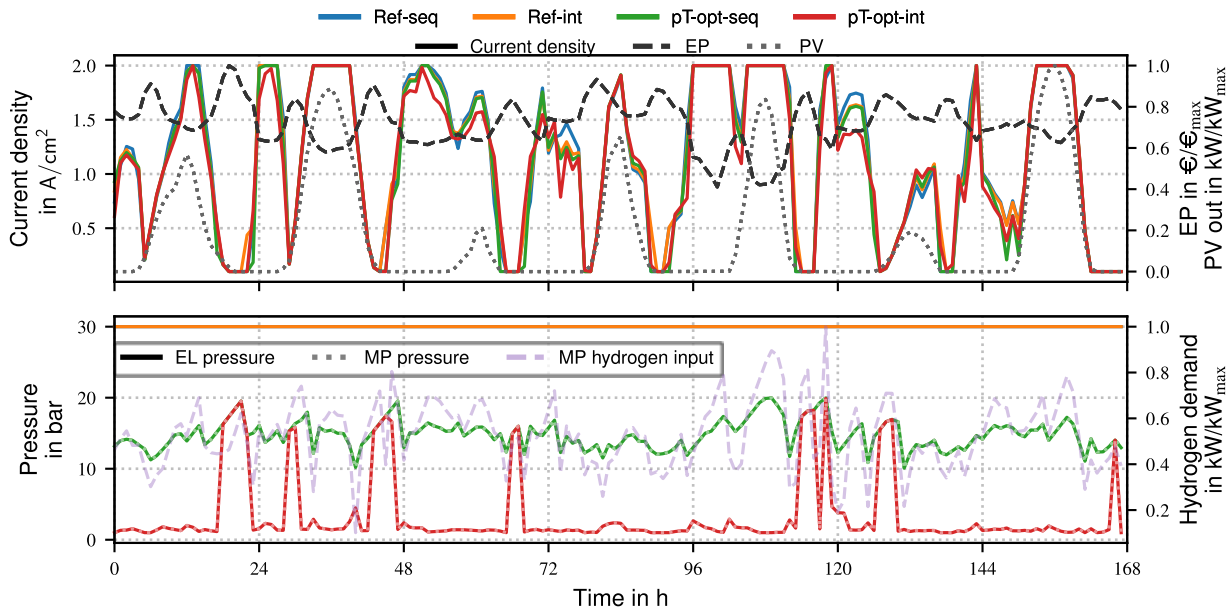

Figure 2: Temporal progression of current density, electrolyzer (EL) pressure, and methanation (MP) pressure across the four optimization modes. Electricity price (EP), photovoltaic (PV) output, and methanation hydrogen input are shown for reference, normalized to their respective maximum values. The blue, orange, and green MP pressure curves overlap.

simultaneously, enabling coordinated system-level optimization. Methanation pressure can be adjusted between 1 and 20 bar [4]. Lower pressure selection necessitates additional hydrogen compression within the methanation process to achieve the desired reaction pressure. These effects are captured in the operating characteristic curve derived from the methanation process model [4]. The time-dependent hydrogen demand of the methanation processes remains fixed across all optimization modes.

RESULTS AND DISCUSSION

Operational planning

Figure 2 illustrates the electrolyzer operating strategies across the four optimization modes, based on current density (i.e., load indicator) and operating pressure.

Current density. The current density of all modes is presented as solid lines in the upper part in Figure 2. Current density profiles—and thus hydrogen production timing and quantities—remain largely consistent across all optimization modes. High current densities (i.e., high operating loads) predominantly occur during periods of low electricity prices (e.g., at 32 h, 100 h) and/or high PV generation (e.g., at 84 h, 156 h).

The 'pT-opt-int' (cf. red line) mode exhibits minor deviations compared to other modes, with slightly lower current densities at certain times (e.g., at 50 h, 124 h) and higher values at others (e.g., at 134 h). This reduction stems from the electrolyzer's higher electrical efficiency at lower loads. Unlike other operating modes, operation at these lower loads is enabled by reducing diffusion losses through lower pressure, allowing the same usable hydrogen production to be maintained despite a lower current density.

Pressure. The lower part in Figure 2 illustrates electrolyzer and methanation pressure as solid and dotted lines, respectively.

The electrolyzer operates at a constant pressure of 30 bar in both 'Ref-seq' and 'Ref-int' modes by definition (cf. blue and orange lines, respectively). In the 'pT-opt-seq' mode (cf. green lines), pressure is reduced to match the methanation pressure within a variable range of ca. 10 to 20 bar. In all three modes, methanation pressure is prescribed exogenously and not optimized. In the 'Ref-seq' and 'pT-opt-seq' modes, this constraint results from the sequential optimization approach, whereas in the 'Ref-int' mode, the electrolyzer pressure necessarily exceeds the methanation pressure. Adjusting the electrolyzer pressure would not yield efficiency gains in this configuration (i.e., the methanation pressure already corresponds to the energetically optimal operating point from a methanation perspective).

Pressure optimization behavior differs in the 'pT-opt-int' mode (cf. red lines). Here, the electrolyzer operates predominantly at substantially lower pressures between 1 and 2 bar, with the methanation pressure adjusted accordingly. The reduced electrolyzer pressure

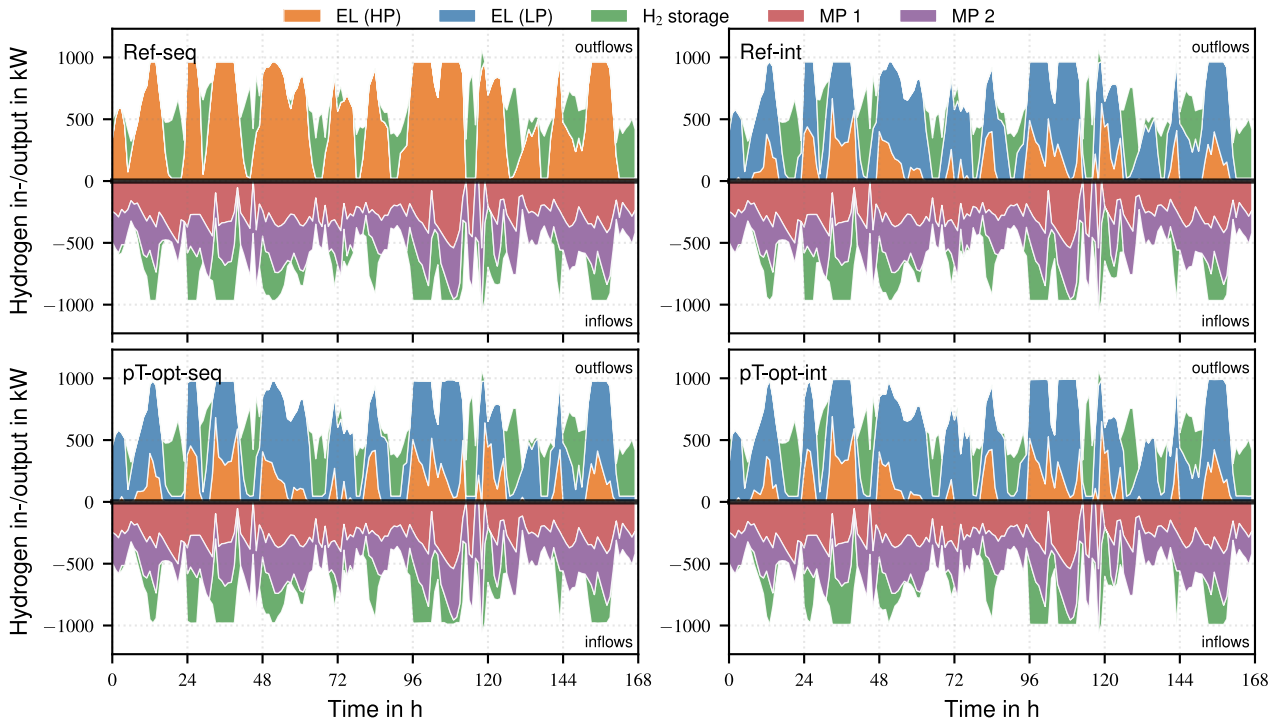

Figure 3: Temporal profiles of hydrogen inflows and outflows for production and utilization across all four optimization modes (EL: Electrolyzer, LP: Low-pressure outlet, HP: High-pressure outlet, MP: Methanation process).

increases electrolyzer efficiency, while the lower methanation pressure leads to higher electricity consumption in the methanation process. Overall, the absolute electricity savings of the electrolyzer outweigh the additional electricity demand of methanation during these periods. However, this trade-off is not universally beneficial. At isolated times when the electrolyzer operates at minimum load (e.g., at 20 h, 44 h), no pressure reduction is applied. Under these conditions, the absolute electricity savings of the electrolyzer are insufficient to compensate for the potential increase in methanation electricity consumption, particularly when methanation operates at elevated load.

Hydrogen production and storage. Figure 3 illustrates the temporal profiles of hydrogen inflows and outflows for production and utilization across all four optimization modes. Consistent with the current density profiles, electrolyzer hydrogen production remains nearly identical across all optimization modes (cf. orange and blue). As hydrogen demand from the methanation processes is additionally inflexible (cf. red and violet), hydrogen storage utilization exhibits very similar patterns across all modes (cf. green).

Consistent with current density patterns, the electrolyzer operates highly dynamically, switching between full load under favorable conditions (e.g., high PV output or low electricity prices) and minimum load under less favorable conditions. Full-load periods serve both direct methanation supply and hydrogen storage, whereas minimum-load periods rely primarily on storage discharge to meet methanation demand. This strategy ensures a stable and continuous hydrogen supply to the methanation processes over the entire time horizon.

All optimization modes except 'Ref-seq' allow the electrolyzer to supply hydrogen via both the low-pressure (LP) and high-pressure (HP) outlets. Hydrogen is either provided exclusively via the LP outlet (i.e., for direct methanation supply) or simultaneously via both outlets (i.e., for methanation supply and storage). At no point is hydrogen supplied solely via the HP outlet. This indicates that compression to 200 bar for storage is energetically unfavorable compared to direct supply at lower pressure levels, leading to a clear preference for maximizing direct hydrogen utilization.

Nevertheless, hydrogen storage remains economically attractive during periods of high PV availability and/or low electricity prices. The additional energy demand for compression is accepted during these periods, even though the electrolyzer capacity would, in principle, be sufficient to supply the methanation processes continuously without storage reliance.

The 'pT-opt-int' mode exhibits minor storage differences compared to other modes. At isolated time steps, less hydrogen is supplied via the HP outlet and thus

stored (e.g., at 60 h, 74 h). Over the entire period, only 21.3% of the produced hydrogen is provided via the HP outlet in the 'pT-opt-int' mode, compared to 23.3% in the 'Ref-int' mode and 24.2% in the 'pT-opt-seq' mode. As the 'pT-opt-int' mode operates the electrolyzer at significantly lower pressures, the subsequent compression effort is correspondingly higher. For specific time steps, the benefits of storage don't outweigh this increased compression energy demand, rendering high-pressure operation economically unattractive more frequently compared to the other modes.

Discussion. The optimization mode does not fundamentally influence the electrolyzer operating strategy with respect to current density. Instead, electrolyzer operation is primarily driven by cost-optimal production conditions, such as low electricity prices or high PV availability, rather than by temporal hydrogen demand. As a result, hydrogen production profiles and storage utilization remain largely consistent across all optimization modes, ensuring a stable hydrogen supply to the methanation processes despite highly dynamic electrolyzer operation.

Regarding pressure, the results indicate that the electrolyzer tends toward the lowest feasible operating pressure when sufficient flexibility is available. When integrated optimization is applied, information exchange between the electrolyzer and the methanation processes enables the determination of an energy- and cost-efficient hydrogen pressure at the overall system level. This leads not only to increased electrolyzer efficiency but also to changes in hydrogen routing, with a clear preference for direct low-pressure supply and a reduced reliance on energy-intensive high-pressure storage.

Consequently, integrated optimization approaches can yield substantial deviations from sequential approaches in optimal operating parameters, including pressure levels and storage usage. Although such approaches require increased planning effort and information exchange, they can enhance system-wide efficiency and contribute to a more cost-effective hydrogen supply and overall energy provision from an integrated energy system perspective.

Costs and electricity consumption

The economic and energetic implications of these operational strategies are quantified through operating cost and electricity consumption. Figure 4 presents the operating cost of the overall energy system as well as the electricity consumption of the electrolysis system and the methanation processes for all modes.

Compared to the reference case 'Ref-seq', the 'Ref-int' mode achieves operating cost savings of 2%. These savings are primarily attributable to the availability of the low-pressure outlet, which eliminates the need to compress the entire hydrogen output. As a result, the electricity consumption of the electrolysis system is reduced by 2.7%, while the electricity demand of the methanation processes remains unchanged.

The 'pT-opt-seq' mode enables 7.3% savings in total operating cost. These savings mainly result from reducing the electrolyzer operating pressure to the methanation pressure and optimizing the operating temperature. Both adaptions increase electrolyzer efficiency and reduce the electrolysis system's electricity consumption by 5.3%. As in the previous case, no impact on the electricity consumption of the methanation processes is observed.

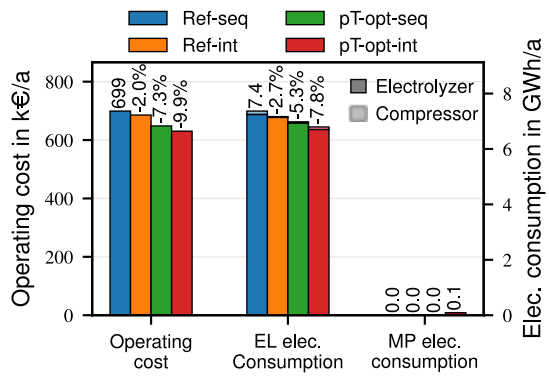

Figure 4. Operating cost and electricity consumption of the electrolysis system (EL) and methanation processes (MP) across all optimization modes.

The 'pT-opt-int' mode yields the highest operating cost savings, amounting to 9.9%. A further decrease in operating pressure drives the additional reduction compared to the 'pT-opt-seq' mode, as the methanation pressure is also adjusted under integrated optimization. While this results in a 7.8% reduction in the electricity consumption of the electrolysis system, the lower methanation pressure leads to an increase in the electricity demand of the methanation processes. In absolute terms, however, the electricity savings achieved by the electrolyzer outweigh the additional electricity consumption of the methanation processes.

Discussion. No single operating strategy is uniformly optimal for all energy system components within the analyzed energy system configuration. Instead, achieving a system-wide optimum requires balancing the differing sensitivities of the involved processes. In the considered optimization modes, electrolyzer electricity consumption is markedly more sensitive to operating pressure than methanation process electricity consumption. Consequently, operating strategies that prioritize electrolyzer performance—while accepting minor efficiency losses in methanation processes—prove advantageous from a system perspective. While sequential optimization approaches already enable noticeable cost and energy savings, the results demonstrate that additional

benefits can be realized through integrated planning by explicitly accounting for interactions between electrolyzer and methanation operation.

CONCLUSIONS

"Green" hydrogen production via PEM electrolyzers remains energy- and cost-intensive, limiting the competitiveness of renewable hydrogen-based energy systems. While adaptive operating strategies have demonstrated efficiency improvements in simple systems with constant hydrogen demand, their effectiveness in complex power-to-X configurations with dynamic downstream synthesis processes and variable pressure requirements has remained unclear. This study addresses this research gap by analyzing integrated optimization of PEM electrolyzer parameters in conjunction with downstream methanation processes, demonstrating that operating cost reductions of up to 9.9% are achievable compared to conventional fixed-parameter operation, with electrolysis system electricity consumption reduced by 7.8%.

Sequential optimization approaches, which adjust electrolyzer parameters independently based on predetermined downstream requirements, achieve 7.3% operating cost savings. In contrast, integrated optimization approaches that simultaneously coordinate electrolyzer and methanation operating conditions achieve an additional 2.6%-point reduction. This superior performance stems from the bidirectional information exchange between system components, enabling the determination of system-wide optimal operating pressures rather than component-specific set points. The operating pressure emerges as the most critical parameter, with electrolyzer electricity consumption exhibiting markedly greater sensitivity to pressure variation than methanation process electricity consumption. This results indicates that system-level optimization should prioritize electrolyzer efficiency, even at the cost of minor electricity consumption penalties in downstream synthesis processes.

These findings demonstrate the significant potential of system-wide optimization approaches for enhancing the cost-competitiveness of "green" hydrogen supply. Implementing control architectures that enable real-time coordination between electrolyzers and downstream synthesis units offers substantial economic benefits, particularly when electrolysis systems are designed for wide operating pressure ranges to fully exploit efficiency optimization potential. Future research should extend this analysis to additional synthesis pathways (e.g., Fischer-Tropsch synthesis and Haber-Bosch synthesis) to further enhance the system-level efficiency and economic viability of renewable hydrogen-based energy systems.

DIGITAL SUPPLEMENTARY MATERIAL

The model code and input data will be made available later in an online repository.

ACKNOWLEDGEMENTS

This work is funded by dtec.bw – Digitalization and Technology Research Center of the Bundeswehr. The support is gratefully acknowledged. dtec.bw is funded by the European Union – NextGenerationEU.

REFERENCES

1. Majumdar A, Haas M, Elliot I, Nazari S. Control and control-oriented modeling of PEM water electrolyzers: a review. International Journal of Hydrogen Energy 48:30621-30641 (2023). https://doi.org/10.1016/j.ijhydene.2023.04.204
2. Balcombe P, Chatenet M, Deseure J, Schäfer H, Staffell I. Markets and costs for hydrogen electrolysis. Green Energy and Technology :235-256 (2024). https://doi.org/10.1007/978-3-031-62411-7_9
3. Bornemann L, Lange J, Kaltschmitt M. Optimizing temperature and pressure in PEM electrolyzers: a model-based approach to enhanced efficiency in integrated energy systems. Energy Conversion and Management 325:119338 (2025). https://doi.org/10.1016/j.enconman.2024.119338
4. Wang Y, von der Assen N, Zhang Q. A decomposition method for optimizing the operation of power-to-X energy systems with detailed process models. (2025) https://doi.org/10.2139/ssrn.5421994
5. Pfenninger S, Staffell I. Long-term patterns of european PV output using 30 years of validated hourly reanalysis and satellite data. Energy 114:1251-1265 (2016). https://doi.org/10.1016/j.energy.2016.08.060
6. Bundesnetzagentur, SMARD.de. Market data. (2024) .
7. Eurostat. Electricity prices for non-household consumers - bi-annual data (from 2007 onwards). (2024) https://doi.org/10.2908/NRG_PC_205
8. Hoffmann M, Kotzur L, Stolten D. The pareto-optimal temporal aggregation of energy system models. Applied Energy 315:119029 (2022). https://doi.org/10.1016/j.apenergy.2022.119029
9. Kotzur L, Markewitz P, Robinius M, Stolten D. Time series aggregation for energy system design: modeling seasonal storage. Applied Energy 213:123-135 (2018). https://doi.org/10.1016/j.apenergy.2018.01.023

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator

and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Evaluating the Potential of Sustainable Aviation Fuel for Decarbonization of the Aviation Sector: An Agent-based Model

Geeta Joshi[a], Tejeswi Ramprasad[a], Harmandeep Singh[c], Narayanan Rajaraman[c], Vikrant Urade[c], Arnoud Higler[d], Rajagopalan Srinivasan[ab]*

[a] Department of Chemical Engineering, Indian Institute of Technology Madras, India
[b] American Express Lab for Data Analytics, Risk and Technology, Indian Institute of Technology Madras, India
[c] Shell India Markets Pvt. Ltd.
[d] Shell Global Solutions International BV.
* Corresponding Author: raj@iitm.ac.in

ABSTRACT

The aviation sector represents one of the most pressing challenges in the energy transition due to its strong reliance on energy-dense liquid fuels and established fuel infrastructure. Sustainable Aviation Fuel (SAF), particularly from agricultural residues, offers a near-term mitigation pathway; however, large-scale adoption is shaped by policy mandates, infrastructure expansion, market price formation, and passenger demand responses. These coupled dynamics are difficult to capture using aggregate or equilibrium-based models. This study develops an agent-based model to analyze SAF transition pathways and applies it to India's civil aviation system. Results show that SAF adoption emerges from the coordination between infrastructure entry, cost learning, and market responses rather than mandate ambition alone. Even moderate mandates fall short of intended adoption levels without timely infrastructure expansion, while aggressive mandates become infeasible under binding supply and price constraints. Passenger demand feedbacks further influence outcomes by linking fuel cost increases to airline operations and route-level allocation decisions. The findings highlight infrastructure coordination and price formation as critical leverage points for aviation decarbonization and demonstrate the value of agent-based models for evaluating realistic SAF transition strategies under policy and market uncertainty.

Keywords: Energy Systems, Agent-Based Modeling, Energy Transition, Sustainable Aviation Fuel, Aviation Decarbonization

INTRODUCTION

Decarbonization of the global energy system has become a central objective of climate mitigation efforts, with particular attention directed toward hard-to-abate sectors. The transport sector contributes nearly 28% of global CO_2 emissions, of which aviation accounts for approximately 2–3%, a share that is expected to increase due to sustained growth in air travel demand. Global aviation demand is projected to double relative to 2024 levels by the early 2040s, with annual growth rates in the range of 3.6–4.2% [1]. Unlike road transport, where electrification and efficiency improvements have already delivered substantial emissions reductions, aviation faces fundamental constraints related to energy density requirements, safety certification, long asset lifetimes, and fleet turnover [2, 3]. These constraints severely limit the near- to medium-term feasibility of alternative propulsion systems, positioning aviation as one of the most challenging sectors to decarbonize. A range of mitigation options is being pursued within aviation, including operational efficiency improvements, demand management, carbon offsetting mechanisms, and the development of hydrogen- and electric-powered aircraft. However, these pathways remain either incremental in their emissions impact or constrained by low technology readiness and infrastructure limitations at scale. Within this landscape, Sustainable Aviation Fuel (SAF) has emerged as the most

viable near-term pathway for reducing aviation emissions while maintaining compatibility with existing aircraft fleets and fuel infrastructure. As a drop-in fuel, SAF can be blended with conventional aviation turbine fuel at levels of up to 50% and used in current aircraft without modifications, enabling near-term emissions reductions without disruptive changes to aircraft or airport infrastructure. When produced from renewable and waste-based feedstocks, SAF can reduce life-cycle greenhouse gas emissions by up to 60–80% relative to fossil jet fuel [4].

LITERATURE REVIEW

SAF can be produced through multiple technological pathways, including Hydroprocessed Esters and Fatty Acids (HEFA) from waste oils and fats, Fischer–Tropsch synthesis via biomass or waste gasification, Alcohol-to-Jet (ATJ) conversion of bio- or waste-derived alcohols, and power-to-liquid routes based on renewable electricity and captured carbon. These pathways differ substantially in technology maturity, capital intensity, and cost trajectories, indicating that large-scale SAF deployment extends beyond a purely technological challenge. Techno-economic analyses consistently identify a substantial cost gap between SAF and conventional aviation turbine fuel. Across production pathways and regional contexts, SAF minimum selling prices are typically estimated to be 3 to 5 times higher than fossil jet fuel, with reported values in the range of USD 1.8–3.5 per kg at industrial scales of 400–600 million liters per year, compared to jet fuel prices of approximately USD 0.4–0.7 per kg [5, 6]. Capital investment dominates costs in gasification-based pathways, while feedstock costs account for more than 50% of operating expenditures in HEFA and ATJ routes. Although learning-by-doing, economies of scale, by-product valorization, and targeted policy support can reduce production costs over time, the magnitude and timing of such reductions remain uncertain. As a result, SAF deployment has been closely tied to policy intervention through blending mandates, tax credits, carbon pricing, and penalties for non-compliance.

Existing studies analyze the impacts of SAF mandates on airline behavior, airline adoption decisions shaped by interactions with airports and fuel suppliers, certification and chain-of-custody requirements for SAF producers, and the role of feed-in tariffs and capital subsidies in enabling biojet fuel supply chains [7, 8, 10]. Scenario-based modeling studies indicate that global SAF deployment faces persistent challenges related to high production costs, feedstock availability, and spatial concentration of supply. Co-benefit analyses further suggest that airlines adopt SAF primarily when SAF price premiums are comparable to carbon prices, while benefit–cost ratios decline when SAF availability is limited to a small number of airports, highlighting the importance of spatial scale and infrastructure coverage [9-11].

Despite these insights, much of the existing literature examines SAF adoption by focusing on individual actors or isolated system components, limiting its ability to capture emergent system behavior. Critical dynamics remain insufficiently explored, including passenger responses to SAF-induced ticket price changes, SAF market behavior under supply–demand imbalances, airline strategies for allocating SAF across airport networks, and the effectiveness of penalties under constrained supply conditions. Moreover, the interaction between short-term operational decisions and long-term strategic investments, which is central to understanding infrastructure lock-in and transition pathways, remains underdeveloped. These limitations point to the need to conceptualize SAF deployment as a socio-technical system in which outcomes emerge from interactions among heterogeneous actors operating under distinct objectives and constraints. The aviation fuel value chain comprises airlines, fuel producers, feedstock suppliers, airports, passengers, markets, and regulators whose decisions interact through non-linear economic, spatial, and temporal feedbacks. Fuel procurement choices influence market prices and investment incentives, infrastructure availability constrains feasible adoption pathways, and early policy signals shape long-term transition trajectories. Such adaptive and path-dependent dynamics are difficult to capture using aggregate, equilibrium-based, or centrally optimized modeling approaches. Agent-based modeling offers a suitable alternative by explicitly representing decentralized decision-making, heterogeneous behavior, and emergent system outcomes [13]. ABMs have been successfully applied in energy and transport systems to study technology diffusion, infrastructure evolution, and policy experimentation under uncertainty [14, 15].

Motivated by these gaps, this study develops Greening Aviation's Energy Requirements through Optimal SAF Supply (Green-AEROS), an agent-based model designed to analyze SAF adoption in aviation as an integrated, multi-actor transition process. Green-AEROS explicitly represents interacting agents across the aviation fuel value chain and links short-term operational behavior, such as route-level fuel blending and passenger demand responses, with long-term strategic decisions, including capacity expansion and infrastructure emergence. The model is applied to India's civil aviation sector, a rapidly growing aviation market with increasing policy interest in SAF deployment, to examine how regulatory ambition, market dynamics, and infrastructure evolution jointly shape SAF adoption pathways and long-term aviation decarbonization outcomes.

AGENT-BASED IMPLEMENTATION OF GREEN-AEROS

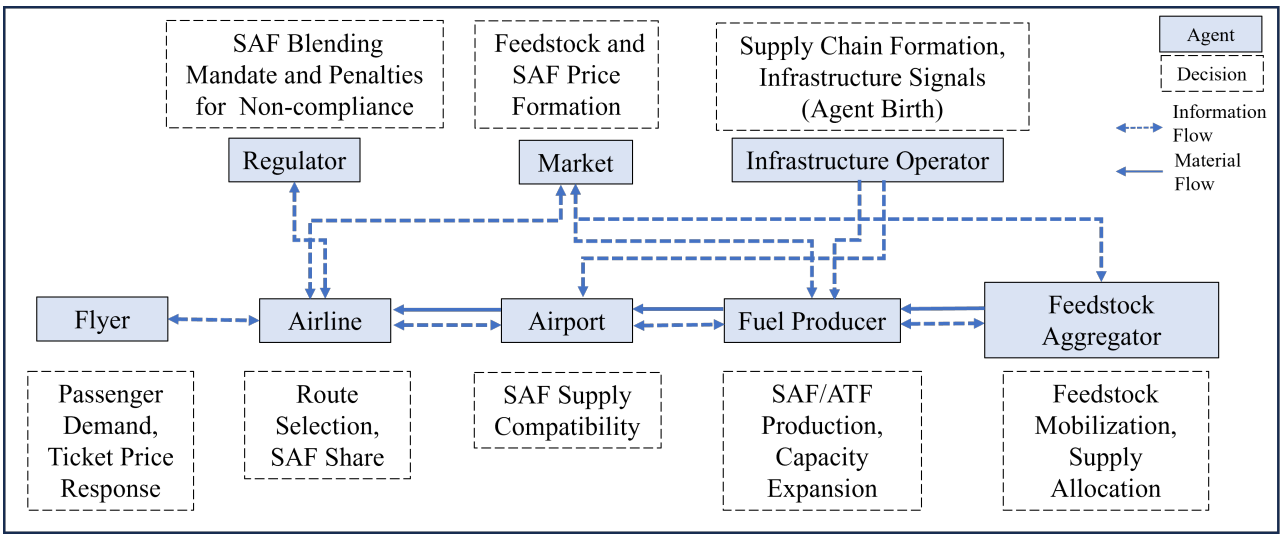

Figure 1. Conceptual architecture of Green-AEROS highlighting major agent decision interactions. Solid arrows denote physical flows such as, fuel use, and SAF supply across routes and airports. Dashed arrows indicate information, policy, price, and coordination signals governing market responses, regulatory enforcement, and infrastructure emergence.

Green-AEROS is an agent-based model developed to explore transition pathways for the aviation sector toward SAF. The model represents decentralized decision-making across the aviation fuel value chain by explicitly modeling 8 interacting agent types: flyers, airlines, airports, fuel producers, feedstock aggregators, infrastructure operator, market, and regulator **Figure 1**. Airline agents represent commercial carriers operating across multiple airports. To capture spatially localized behavior, each airline agent comprises airport-specific subagents that make route-level decisions, including passenger demand forecasting and ticket price formation. The allocation of SAF and ATF shares on routes originating from each airport is determined through a linear programming formulation that minimizes penalties, subject to a maximum SAF blending limit of 50% and airport-level SAF availability constraints. Baseline passenger demand is driven by exogenous socio-economic factors such as population growth and GDP per capita, which evolve over time. Flyer agents respond endogenously to ticket price changes, with price sensitivity expressed through the effective fuel cost resulting from SAF–ATF blending. Airport agents decide on SAF infrastructure activation based on airline demand for SAF and local availability from fuel producers. Fuel producer agent consists of two sub-agents: conventional refinery agents supplying ATF and SAF producer agents supplying SAF. SAF producer agents make decisions regarding production volumes, capacity expansion, technology deployment, and feedstock procurement. Production decisions are primarily driven by feedstock availability and learning effects associated with cumulative production. SAF is assumed to be transferred to existing refinery infrastructure for blending with ATF and subsequently transported to connected airports. This assumption reflects the utilization of existing fuel infrastructure and supports early-stage SAF deployment without requiring immediate system-wide retrofitting. Feedstock aggregator agents are spatially distributed according to regional agricultural residue availability and make decisions related to feedstock mobilization and supply network formation with SAF producers. The infrastructure operator agent governs the endogenous emergence of SAF supply chains by signaling the entry of new SAF production facilities and airport-level SAF infrastructure. This agent coordinates linkages between feedstock aggregators, SAF producers, and airports by optimizing the use of existing fuel infrastructure and minimizing system-level bottlenecks. The infrastructure operator signals the need for new SAF production capacity based on airline demand and the capacity expansion decisions of existing SAF producers. Regulator and market agents provide the policy and price-formation environment. The regulator sets national SAF blending mandates for airlines and enforces penalties for non-compliance, with penalties anchored to SAF prices discovered by the market agent. The market agent receives SAF demand from airlines and supply from SAF producers, as well as feedstock supply from feedstock aggregators and feedstock demand from SAF producers, and discovers SAF and feedstock prices in response to supply–demand imbalances. The resulting SAF prices evolve over time with changes in production scale, feedstock availability, and learning effects. To ensure numerical stability and avoid unrealistic price volatility, a min–max capping mechanism is applied to bound SAF price variations. Given the early stage of SAF deployment, the model assumes a uniform

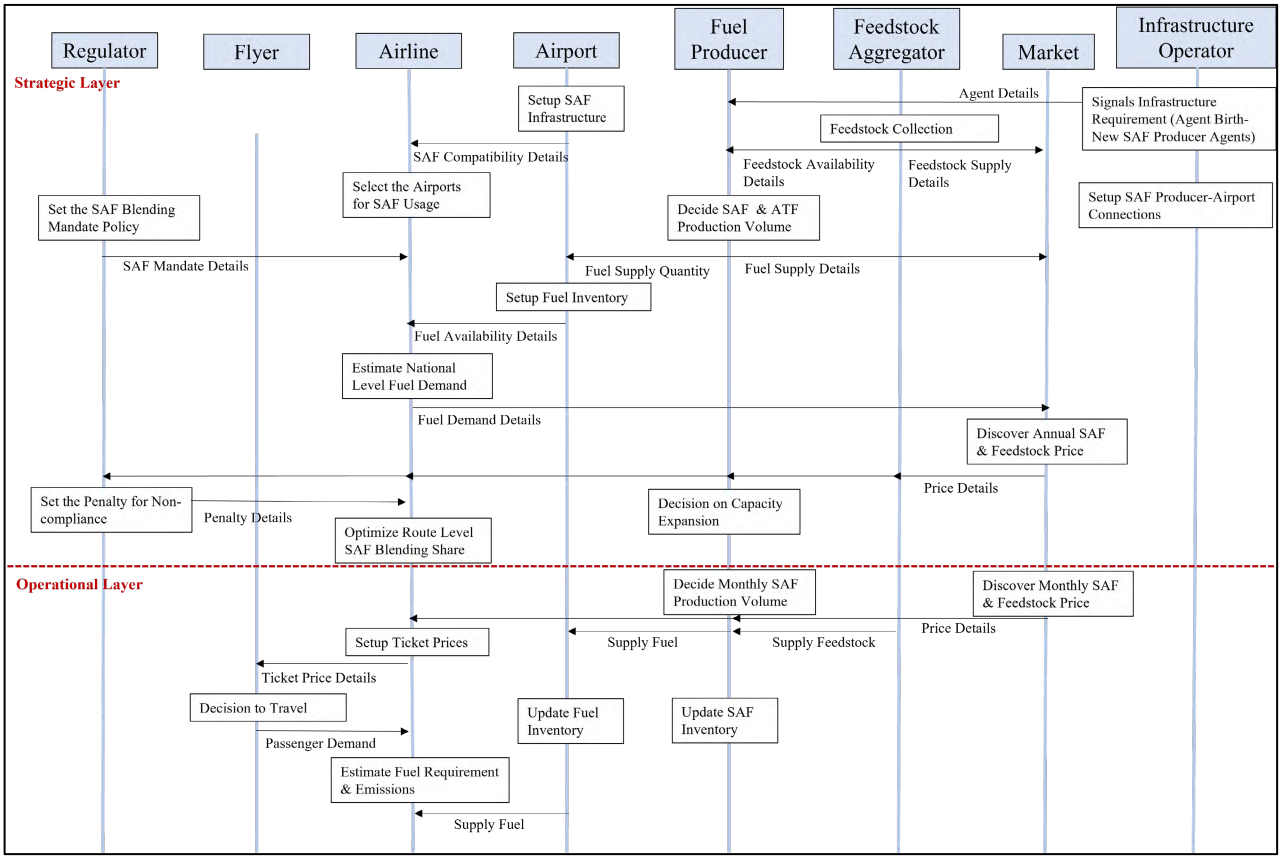

Figure 2. Agent-interaction and decision sequence in the Green-AEROS model.

national SAF price, while feedstock prices vary spatially to reflect regional availability. All agents are equipped with forecasting and optimization modules that enable forward-looking decision-making under uncertainty, incorporating costs, risks, incentives, and learning over time. Agent decisions in Green-AEROS occur at two levels, operational and strategic, and evolve over different time scales that jointly shape the transition dynamics. **Figure 2** presents the sequence of agent-level decisions and their interactions within each time step. It highlights the predefined order of decision updates and the exchange of information through shared system variables such as prices, demand, and capacity, thereby representing coordination across agents within a decentralized framework. For instance, at the strategic level, the infrastructure operator signals the birth of new SAF producer agents and determines SAF compatibility at connected airports, which is used by airline agents to optimize SAF blending shares at the route level. These decisions propagate to the operational level, where airlines set ticket prices that are communicated to flyer agents, whose responses capture passenger demand sensitivity to SAF adoption.

CASE STUDY: INDIA'S CIVIL AVIATION SECTOR

The performance of Green-AEROS is evaluated using a case study of India's Civil Aviation Sector (ICAS), the third-largest domestic aviation market globally, with more than 7, 100 daily aircraft movements and annual carbon dioxide emissions of approximately 30–35 million tonnes [16, 17]. Green-AEROS is implemented in Repast4Py and instantiated to represent 11 major airports (DEL, BOM, BLR, CCU, MAA, HYD, GOI, COK, LKO, AMD, and IDR) and 110 domestic routes, together capturing approximately 69% of domestic air traffic. A single representative airline agent is modeled to capture aggregate airline behavior, incorporating system-level risk preferences and strategic decision-making. Route-level passenger demand is forecast by the airline agent using exogenous socio-economic drivers, including gross domestic product, fuel prices, population, and per capita income. Conventional refineries are assumed to supply the required quantities of ATF, while SAF is produced from agricultural residues through the Alcohol-to-Jet pathway. Blending of SAF and ATF is assumed to occur at refineries located nearest to airports. Feedstock

aggregator agents are distributed across 14 states and 374 districts and supply agricultural residues derived from rice, maize, and sugarcane.

The number and capacity of SAF production sites evolve endogenously over time, driven by airlines' willingness to adopt SAF in order to reduce penalties for non-compliance and by feedstock availability within a procurement distance (P_d) from production sites which. Strategic decisions are modeled at an annual resolution, while operational decisions occur monthly over a 20-year simulation horizon spanning 2024–2044. In line with current policy timelines in India, SAF blending mandates are applied from 2027 onward, starting at 1% in 2027, increasing to 2% in 2028, and reaching 5% by 2030. The model is initialized with a single SAF production site with a capacity of 43.75 million liters, supplying SAF to one airport (DEL) at an initial SAF price ($Price_{SAF}$) of 275.23 INR per liter, with a lead time of two years before the facility becomes operational. SAF blending mandates are imposed exogenously as policy requirements; however, realized SAF adoption emerges endogenously from agent interactions and may fall short of mandated levels when supply-side, infrastructure, feedstock availability, or cost constraints become binding.

To examine the implications of increasing policy ambition, three policy configurations are evaluated. All policies follow an identical SAF blending trajectory until 2030, reflecting the current national mandate pathway. **Policy 1** represents a low-ambition pathway, in which the SAF blending requirement increases to a maximum of 5% by 2030 and remains constant thereafter. **Policy 2** reflects a moderate ambition pathway, in which the blending requirement continues to increase beyond 2030, reaching a maximum of 15% over the 20-year horizon. **Policy 3** represents a high-ambition pathway, in which the blending requirement rises further to a maximum of 30% over the same period. These policy configurations are used to assess how infrastructure emergence, market dynamics, and airline decision-making jointly shape feasible SAF adoption trajectories under progressively stronger regulatory targets.

RESULTS AND DISCUSSION

The evolution of SAF adoption under Policies 1, 2, and 3 by comparing imposed blending requirements with realized SAF shares emerging from agent interactions is illustrated in **Figure 3(a)**. For ICAS, passenger demand across 110 domestic routes increases from 100.4 million passengers in 2024 to 182 million by 2044, with total aviation fuel demand rising from 2.6 billion liters to 4.8 billion liters. Over the same period, ATF prices increase by 8.1%. Against this growth trajectory, early SAF uptake across all policies remains highly localized. In 2027, a national SAF blending level of 1% is achieved with a total SAF requirement of approximately 30 million liters, accomplished by deploying SAF on a single high-volume route (DEL–BLR), where SAF accounts for 24% of the fuel used on that route.

By 2030, the mandated blending level increases to 5%, yet realized SAF adoption remains limited to approximately 1.4% across all three policies. This divergence indicates that mandate escalation alone is insufficient in the absence of supporting infrastructure. From 2031 onward, realized adoption trajectories begin to diverge as infrastructure expansion is triggered endogenously by airlines' incentives to reduce penalties for non-compliance. Policies 2 and 3 achieve the 5% blending target by 2032, requiring approximately 206 million liters of SAF, while Policy 1 reaches the same level by 2033. This earlier achievement under Policies 2 and 3 reflects accelerated entry of SAF production sites and expansion of airport-level SAF compatibility driven by higher post-2030 policy ambition.

Table 1: Number of airports supplying SAF and SAF producers operating by 2044.

	Airports	SAF Producers
Policy 1	3	3
Policy 2	7	5
Policy 3	10	8

Under Policy 2, SAF adoption continues to increase steadily, reaching 15% blending by 2040 with a corresponding SAF requirement of 682 million liters. In contrast, under Policy 3, realized SAF adoption plateaus at a maximum of 25.15% by 2044 despite stronger policy signals. This saturation occurs even though additional SAF production capacity is available, indicating that airport-level blending constraints and limited route coverage restrict further demand growth. **Table 1** shows that achieving a 5% blending level by 2044 requires 3 SAF production sites supplying DEL, BLR, and BOM, whereas higher blending levels depend more strongly on expanding SAF compatibility across additional airports rather than increasing the number of producers alone. Under Policy 2, the addition of 2 SAF producers supplying HYD and MAA, along with SAF availability at LKO and GOI from existing SAF producers, enables higher system-wide adoption. Under Policy 3, even with 8 SAF production sites supplying 10 airports, a 16% SAF shortfall is observed by 2044.

These first-order insights indicate that mandate ambition alone does not determine realized SAF adoption. Instead, adoption is constrained by the timing and spatial alignment of SAF production sites, airport-level blending compatibility, and airline route-level allocation decisions. **Figure 3(b)** shows the regional distribution of India's SAF

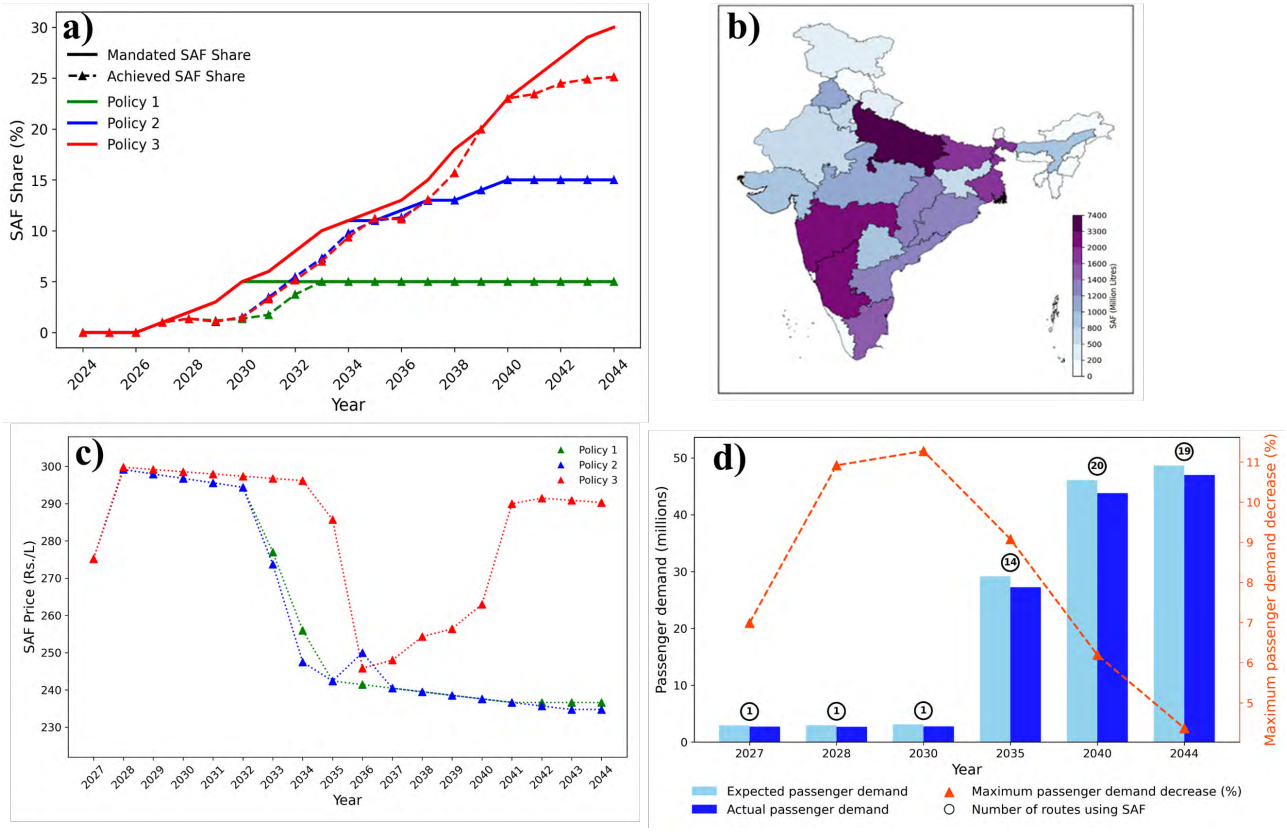

Figure 3. (a) SAF uptake under different policies. (b) Spatial distribution of potential SAF production from agricultural residues across India. (c) SAF market price evolution across policy pathways. (d) Passenger response to SAF adoption under Policy 2(only for routes using SAF).

production potential derived from agricultural residues of rice, maize, and sugarcane, which together amount to an estimated 7, 600 million liters annually. While this potential substantially exceeds the SAF requirement even under higher blending policies, it is unevenly distributed across regions and does not automatically translate into effective supply at demand centers. As a result, SAF adoption concentrates at airports with high route volumes and favorable blending conditions, even under a nationally uniform SAF price. This spatial mismatch between feedstock availability, production capacity, and airport demand explains why higher blending levels remain unattainable despite sufficient aggregate resource potential, highlighting the need for coordinated infrastructure expansion alongside policy mandates.

Figures 3(c) and 3(d) together reveal how SAF price dynamics translate into uneven passenger demand responses across spatial and temporal scales. In the early phase of Policy 2, limited SAF availability results in high blended fuel costs. In 2027, with a SAF price of 275.23 INR/L, airlines achieve the mandated 1% blending by concentrating SAF usage on a single high-volume route (DEL–BLR), where a 24% SAF share leads to a route-level passenger demand reduction of 6.8%. As SAF prices increase further by 8.2% in 2028 due to continued supply constraints, the maximum route-level demand reduction deepens to nearly 10%, despite SAF still being used on only one route. This illustrates how early SAF adoption manifests as localized demand shocks driven by price pressure rather than system-wide contraction.

As SAF supply expands and prices moderate, these dynamics evolve. By 2035, increased SAF availability lowers prices sufficiently to enable airlines to distribute SAF usage across 14 routes. While the maximum route-level demand reduction remains noticeable at 8.9%, the aggregate passenger demand reduction across all SAF-using routes is limited to 6.6%, indicating a diffusion of impacts rather than intensification. By 2042, SAF prices stabilize under Policy 2 as supply better aligns with demand, allowing SAF usage to expand to 19 routes. At this stage, both maximum and aggregate demand impacts decline further, with a maximum route-level reduction of 4.2% and an overall reduction of 3.4%. These results highlight a key emergent insight: passenger responses to SAF adoption are highly non-linear and scale-dependent. Localized impacts at the route level do not translate proportionally to system-wide outcomes, underscoring the importance of bottom-up modeling approaches in

capturing how micro-level airline allocation strategies shape macro-level demand trajectories during the SAF transition.

CONCLUSION

In the present study, Green-AEROS is developed as an agent-based model to examine Sustainable Aviation Fuel transition dynamics as an emergent outcome of decentralized decision-making across the aviation fuel value chain. Applied to India's civil aviation sector, this work constitutes one of the first agent-based, system-level representations of the SAF socio-technical system, explicitly integrating airlines, passengers, fuel producers, infrastructure evolution, markets, and regulation within a unified framework. The results show that realized SAF adoption is governed less by mandate ambition in isolation and more by the alignment of infrastructure emergence, supply-side learning, and market-mediated price signals. Moderate mandates fall short of intended adoption levels when SAF infrastructure does not emerge in a timely manner, while highly ambitious mandates become infeasible under binding supply, blending, and price constraints. By jointly representing passenger demand responses, airline operational choices, SAF production expansion, infrastructure birth, and endogenous price discovery, the study reveals tightly coupled policy–infrastructure–market feedbacks that shape realistic SAF transition pathways.

Future work will extend Green-AEROS along several dimensions to strengthen its applicability and policy relevance. Planned developments include representing multiple airline agents with heterogeneous route-selection strategies and incorporating evolving technology maturity to capture strategic diversity in SAF deployment decisions. The model will also be expanded to account for regional feedstock constraints arising from competing uses, which may significantly influence the spatial feasibility of SAF production. In addition, integrating detailed techno-economic assessments at the SAF production site level will enable explicit evaluation of long-term economic viability under alternative policy and market conditions. Given its modular structure, Green-AEROS is readily extendable to other national and regional aviation systems, providing a transferable modeling framework for analyzing SAF transitions across diverse geographies and policy contexts.

REFERENCES

1. IATA. Air Passengers Market Analysis (2025) https://www.iata.org/en/publications/economics/reports/air-passenger-market-analysis-november-2025/
2. Jensen LL, Bonnefoy PA, Hileman JI, Fitzgerald JT. The carbon dioxide challenge facing U.S. aviation and paths to achieve net zero emissions by 2050. Progress in Aerospace Sciences 141:100921 (2023). https://doi.org/10.1016/j.paerosci.2023.100921
3. Schäfer AW, Barrett SRH, Doyme K, Dray LM, Gnadt AR, Self R, O'Sullivan A, Synodinos AP, Torija AJ. Technological, economic and environmental prospects of all-electric aircraft. Nat Energy 4:160-166 (2018). https://doi.org/10.1038/s41560-018-0294-x
4. Watson MJ, Machado PG, da Silva AV, Saltar Y, Ribeiro CO, Nascimento CAO, Dowling AW. Sustainable aviation fuel technologies, costs, emissions, policies, and markets: a critical review. Journal of Cleaner Production 449:141472 (2024). https://doi.org/10.1016/j.jclepro.2024.141472
5. Teixeira AT, da Silva ACM, Cavalcante RM, Young AF. Process simulation and economic evaluation of the alcohol-to-jet production of sustainable aviation fuel in the brazilian context. Energy Conversion and Management 319:118947 (2024). https://doi.org/10.1016/j.enconman.2024.118947
6. Uddin MM, Lee U, Xu H, Li Y, Kwon H, Zhang Y, Smolinski S, Cai H, Tao L. Sustainable aviation fuel from ethanol: techno-economic analysis and life cycle analysis. Applied Energy 398:126373 (2025). https://doi.org/10.1016/j.apenergy.2025.126373
7. Prussi M. The sustainability dimension for sustainable aviation fuels (SAF): comparing regional and international approaches. Sustainability 17:8401 (2025). https://doi.org/10.3390/su17188401
8. Braun M, Grimme W, Oesingmann K. Pathway to net zero: reviewing sustainable aviation fuels, environmental impacts and pricing. Journal of Air Transport Management 117:102580 (2024). https://doi.org/10.1016/j.jairtraman.2024.102580
9. Martulli A, Brandt K, Allroggen F, Malina R. The potential scale-up of sustainable aviation fuels production capacity to meet global and EU policy targets. Nat Commun 16: (2025). https://doi.org/10.1038/s41467-025-66686-9
10. Cui Q, Lei Y, Li Y. Co-benefit scenario analysis on replacing aviation kerosene with sustainable aviation fuels in africa. Environmental Impact Assessment Review 104:107318 (2024). https://doi.org/10.1016/j.eiar.2023.107318
11. Woeldgen E, Teoh R, Stettler MEJ, Malina R. Sustainable aviation fuel deployment strategies in europe: supply chain implications and climate benefits. Environ. Sci. Technol. 59:12447-12457 (2025). https://doi.org/10.1021/acs.est.5c02364
12. Moncada JA, Verstegen JA, Posada JA, Junginger M, Lukszo Z, Faaij A, Weijnen M. Exploring the

emergence of a biojet fuel supply chain in brazil: an agent?based modeling approach. GCB Bioenergy 11:773-790 (2019). https://doi.org/10.1111/gcbb.12594
13. Hansen P, Liu X, Morrison GM. Agent-based modelling and socio-technical energy transitions: a systematic literature review. Energy Research & Social Science 49:41-52 (2019). https://doi.org/10.1016/j.erss.2018.10.021
14. Zhang Z, Jing R, Lin J, Wang X, van Dam KH, Wang M, Meng C, Xie S, Zhao Y. Combining agent-based residential demand modeling with design optimization for integrated energy systems planning and operation. Applied Energy 263:114623 (2020). https://doi.org/10.1016/j.apenergy.2020.114623
15. Venkataramanan VS, Shahrukh M, Papageorgiou DJ, Rajagopalan S, Srinivasan R. Identifying key drivers for a national transition to low carbon energy using agent-based supply chain models. Computers & Chemical Engineering 182:108541 (2024). https://doi.org/10.1016/j.compchemeng.2023.108541
16. Thummala V, Hiremath RB. Green aviation in india: airline's implementation for achieving sustainability. Cleaner and Responsible Consumption 7:100082 (2022). https://doi.org/10.1016/j.clrc.2022.100082
17. Ministry of Civil Aviation. Government of India. Annual Report 2024-25. https://www.civilaviation.gov.in/Publication/annual-reports

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Data-Driven Optimization Framework for the Design and Operation of Adaptive and Resilient Energy Supply Chain Networks under Uncertainty

Halil Iseri[ad], Funda Iseri[bd], Mahmoud El-Halwagi[bd], Eleftherios Iakovou[cd], Efstratios Pistikopoulos[bd*]

[a] Multidisciplinary Engineering, Texas A&M University, College Station, USA
[b] Chemical Engineering, Texas A&M University, College Station, USA
[c] Engineering Technology and Industrial Distribution, Texas A&M University, College Station, USA
[d] Texas A&M Energy Institute, College Station, USA
* Corresponding Author: stratos@tamu.edu.

ABSTRACT

Recent geopolitical disruptions and extreme weather events have underscored the importance of resilience in global energy supply chains, particularly for import-dependent economies pursuing ambitious energy transition targets. These events have exposed the limitations of supply chain designs focused solely on cost minimization that lack the flexibility and redundancy required for secure operation under stress. As energy systems evolve toward higher shares of variable renewable energy and increased demand uncertainty, episodic manual re-planning becomes inadequate, highlighting the need for modeling frameworks that integrate predictive modeling, optimization, and control to enable intelligent and adaptive supply-chain design and operations under uncertainty. This work presents a comprehensive data-driven modeling and optimization framework for adaptive energy supply-chain networks under evolving demand. The framework integrates three layers: (i) a machine-learning model for demand forecasting and scenario generation; (ii) a multi-period stochastic optimization model for strategic network design and operations; and a (iii) learning layer that monitors performance metrics and triggers strategic recourse when demand patterns shift significantly. The operational stage, which acts as a learning layer, is posed as a rolling horizon control problem determining necessary recourse decisions to adapt to changes in demand patterns. An illustrative case study, encompassing multiple energy generation hubs, energy carriers and transportation modes is shown to demonstrate the applicability of the framework.

Keywords: energy supply chain, intelligent systems, multi-scale modeling, hydrogen, ammonia, resilience

INTRODUCTION

Recent years have highlighted the critical importance of resilience in global energy supply chains, as geopolitical conflicts and natural disasters have exposed structural vulnerabilities, particularly in import-dependent economies pursuing ambitious energy-transition targets. These disruptions have revealed the risks of supply concentration and supply-chain network designs focused primarily on cost minimization, which often lack the flexibility and redundancy required to ensure secure and reliable operation under stress. As energy systems evolve toward higher shares of variable renewable energy and increased demand uncertainty, the intermittency of renewable generation and emerging electrification trends exacerbate spatio-temporal supply–demand mismatches, motivating energy supply-chain networks with seasonal storage as a multi-scale systems integration solution connecting resource-abundant regions to industrial demand centers [1, 9]. Moreover, the rising frequency of disruptions and changing demand patterns render episodic, manual re-planning inadequate, motivating the need for continuous, data-informed, and adaptive decision-making frameworks.

The process systems engineering community has developed a broad range of methodologies for planning

and operating energy and process systems under uncertainty. Foundational work established stochastic optimization frameworks for process synthesis and multi-period supply chain planning under uncertainty [2, 3], including applications to hydrogen and energy carrier supply chains [4]. At the operational level, model predictive control (MPC) has been widely adopted for real-time energy system scheduling using receding-horizon formulations [5]. More recently, data-driven approaches have emerged that integrate machine learning with mathematical programming, enabling forecasting-informed optimization and scenario-based planning under deep uncertainty [6–8]. Despite these efforts, the disruptions faced in the recent decade have revealed the risks of supply concentration and supply-chain network designs focused solely on cost minimization, lacking flexibility and redundancy. As the power grid has been expanding with intermittent renewable energy sources and the pace and frequency of disruptions rise in supply-chain networks and demand patterns change, episodic manual re-planning proves increasingly insufficient, motivating the need for systems that anticipate short-horizon changes and continuously update operating decisions. Addressing these challenges calls for a shift from static planning to continuous, data-informed, proactive operations capable of adapting as conditions evolve. Therefore, there is a clear need for modeling frameworks that integrate predictive modeling, optimization, and control to design and operate intelligent supply-chain networks.

In this research, we present a data-driven modeling and optimization framework for the design and operation of adaptive energy supply-chain networks under uncertainty. The framework integrates rolling demand forecasting and scenario generation, multi-period stochastic optimization, and a learning-based recourse mechanism to enable proactive adaptation of supply-chain structure and operations. An illustrative case study with multiple energy hubs, energy carriers, and transportation modes is used to demonstrate the applicability of the proposed approach.

METHODOLOGY

The framework features: (i) a machine-learning model that provides rolling demand forecasts and scenarios for stochastic optimization; (ii) a multi-period stochastic optimization framework for planning and operations; and (iii) a learning layer that tracks changes in demand patterns and informs strategic recourse decisions. The machine-learning layer supplies rolling-horizon predictions at each time step, and the stochasticity of demand is captured through representative scenarios. These scenarios are then utilized in the multi-period, multi-stage stochastic optimization step considering multiple production locations and transportation modes for planning the optimal supply chain network. The operational stage, which acts as a learning layer, is modeled as a rolling horizon control problem that determines the necessary recourse decisions to adapt to changes in short-term demand patterns and long-term forecasts. Figure 1 illustrates the overall structure of the framework and the superstructure of possible energy transmission solutions. The methodology follows a closed-loop architecture in which forecasting, optimization, control, and learning layers interact iteratively. At each reassessment epoch, updated demand information and operational decisions are revised over a receding horizon, and system performance indicators are evaluated to determine whether the existing strategic design remains adequate or whether strategic recourse actions are required.

In the proposed framework, strategic redesign decisions are not triggered by instantaneous forecast–realization discrepancies or ad hoc threshold rules, but rather by a structured interactions between rolling-horizon operational control and a learning-based policy layer enabling an “intelligent” network. At predefined reassessment epochs, the operational layer provides two complementary signals: (i) a backward-looking measure of persistent service stress; and (ii) a forward-looking adequacy indicator reflecting the predicted inability to meet future demand under the existing infrastructure. Only when both conditions indicate sustained capacity inadequacy is a strategic reassessment activated. While this adequacy gate determines *when* redesign is admissible, the selection of *how* to adapt the network structure is delegated to a learning-based recourse mechanism described in Section 2.5. Unlike rule-based triggering schemes, in which both the redesign trigger and the corrective action are fixed ex ante, the proposed framework decouples the admissibility of strategic redesign from the selection of recourse actions, enabling adaptive responses learned from operational experience, thus enabling data-driven and adaptive infrastructure evolution under uncertainty.

2.1 Mathematical Notation

The renewable energy supply chain is modeled through a superstructure, as shown in Figure 1(ii), that captures multiple production locations, energy carriers, conversion processes, transportation modes, and temporal scales. Renewable electricity generation, energy carrier production, storage, and transportation are explicitly represented to satisfy energy demand at designated sink locations. The renewable energy supply chain planning problem is modeled as a MILP problem and solved by the Gurobi optimizer. All processes and transportation modes are assumed to be powered by renewable electricity and the energy carriers produced in the

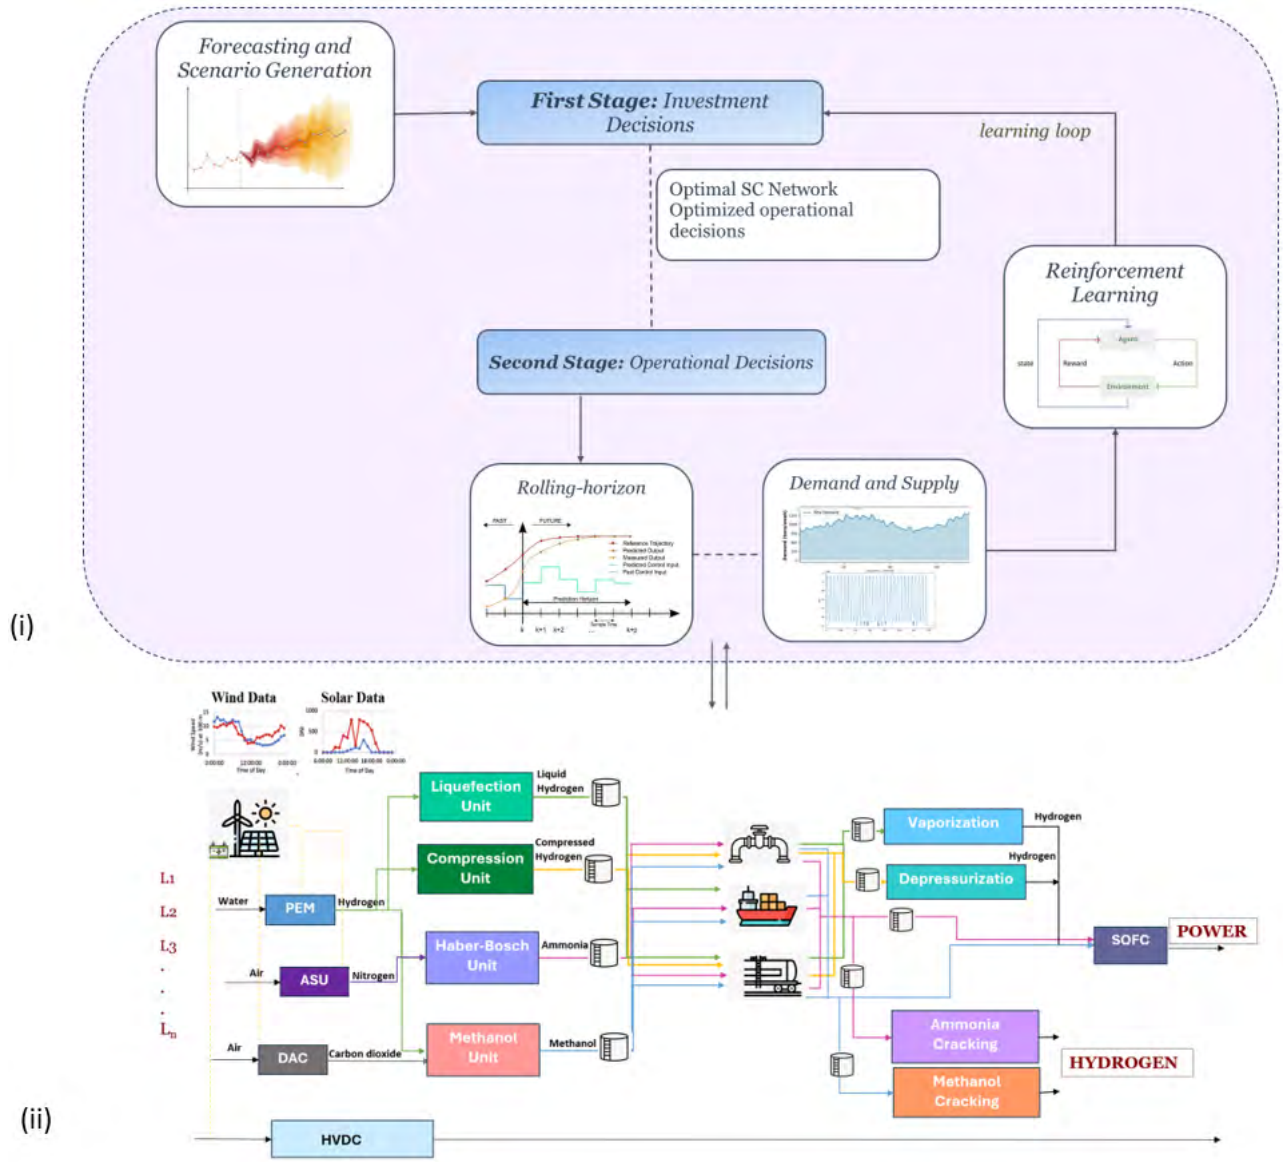

Figure 1. (i)Methodology and (ii) Superstructure of Energy Transmission Solutions

installed process plants.

The employed nomenclature is provided in Table 1. $\mathcal{L}$ represents the set production locations, $\mathcal{C}$ represents the set of energy carriers such as liquid hydrogen, methanol, ammonia and compressed hydrogen. $\mathcal{P}$ represents the set of processes to produce these energy carriers such as liquefaction, compression plants and air separation, direct air capture, methanol and Haber-Bosch units that are necessary in producing these energy carriers. The process superstructure is modelled through the parameters of production processes and utilization of mass and energy balance, capacity, transportation, demand and contingency constraints. $\mathcal{M}$ represents the set of transportation modes, either ship or pipe-line transport. $\mathcal{R}$ represents the set of renewable sources, solar, wind, etc.

Table 1 Notation Table for Sets

Symbol	Description
$l \in \mathcal{L}$	All locations
$c \in \mathcal{C}$	Energy carriers
$p \in \mathcal{P}$	Conversion processes
$r \in \mathcal{R}$	Renewable resources
$m \in \mathcal{M}$	Transport modes
$t \in \mathcal{T}$; $\tau \in \mathcal{T}^h$	Planning/Operational time steps
$k \in \mathcal{K} \subseteq \mathcal{T}$	Reassessment epochs
$\omega \in \Omega_k$	Scenarios at epoch k
$a \in \mathcal{A}$	Strategic action space

2.2 Demand Forecasting and Scenario Generation

Uncertainty in demand is represented through a machine learning based forecasting model. Using historical data, the forecasting layer generates demand trajectories by sampling quantile levels and mapping them through the cumulative distribution function (CDF) for each period. Then, scenarios corresponding to 25th, 50th and 75th percentiles aligned with low, medium and high demand levels are extracted. These scenarios are then embedded into the multi-period stochastic optimization framework to guide investment and operational decisions under demand uncertainty [6].

These probabilistic demand scenarios serve as the primary input to the strategic and operational optimization layers. By continuously updating scenario information as new observations become available, the framework maintains consistency between realized system behavior and future planning assumptions. In the operational layer, the forecasting model processes historical demand observations and generates regime-aware monthly energy demand forecasts on a rolling horizon. The model dynamically updates its feature set with each realized month's demand enabling adaptive learning.

2.3 Two-Stage Stochastic Optimization

Strategic investment decisions are determined using a two-stage stochastic optimization framework. First-stage decisions correspond to long-term infrastructure investments, including the portfolio of renewable power plants, production, transportation infrastructure, and storage capacities at each energy hub to generate power, produce energy carriers and to transport them to the demand center. Second-stage decisions represent operational responses to scenario-specific demand realizations, such as production and transportation dispatch.

The objective function minimizes the sum of investment costs and the expected operational cost across demand scenarios. Mass and energy balance, production and transportation capacity, contingency, and demand satisfaction constraints ensure physical and operational feasibility. Moreover, resilience constraints ensure decentralization of multiple energy hubs, oversizing of capacity and redundancy in transportation infrastructure. Investment lead-time constraints are explicitly enforced so that newly installed capacities become available only after realistic construction delays, preventing non-possible recourse actions in the operational layer. The resulting solution defines a robust baseline network design that balances cost efficiency and flexibility under demand uncertainty.

Objective function:

$$\min_{\Delta X_k,\{Y_\omega\}} \quad C^{inv}(\Delta X_k) + \sum_{\omega\in\Omega_k} \pi_{\omega k}\left[C^{op}(Y_\omega) + \rho \sum_{\tau\in\mathcal{T}_k^h} \sum_{l\in\mathcal{L}^{dem}} u_{l\tau\omega k}\right]$$

where u is the unmet demand characterized by the Value of Lost Load (VoLL).

Investment Cost (First-stage):

$$C^{inv}(\Delta X_k) = \sum_{l,r} c_r^{RE}\,\Delta K_{lrk}^{RE} + \sum_{l,p} c_p^{P}\,\Delta K_{lpk}^{P} + \sum_{l,c} c_c^{S}\,\Delta S_{lck} + \sum_{l} c_l^{hub}\left(z_{lk} - z_{l,k^-(k)}\right) + \sum_{i,j,c,m} c_{ijcm}^{Q}\,\Delta Q_{ijcmk}$$

Operating Cost (Second-stage):

$$C^{op}(Y_\omega) = \sum_{\tau\in\mathcal{T}_k^h}\left[\sum_{l,p} c_p^{op}\,y_{lp\tau\omega k} + \sum_{i,j,c,m} c_{ijcm}^{tr}\,q_{ijcm\tau\omega k} + \sum_{l,c} c_c^{h}\,I_{lc\tau\omega k}\right]$$

In addition to the aforementioned constraints, an improved resilience footprint is enforced at the strategic design stage through constraints promoting decentralization, oversizing, and redundancy, applied across all demand scenarios ω.

Decentralization is imposed by limiting supply concentration from each hub and geographic dispersion is further encouraged using a distance-weighted diversification constraint inspired by electric charge repulsion:

$$\sum_{(i,j)} w_{ij}\left(D_i^{\omega}+D_j^{\omega}\right) \le \gamma\, D_{\text{tot}}^{\omega} \forall\omega.$$

Oversizing is enforced by requiring aggregate production and transportation capacity to exceed delivered demand:

$$\sum_i C_i \ge (1+\beta^{\text{sys}})\, D_{\text{tot}}^{\omega} \forall\omega.$$

A *composite* resilience *constraint* combines structural diversity and capacity headroom:

$$\frac{\alpha_K}{|\mathcal{L}|-1}(K-1) + \frac{\alpha_H}{D^{\text{target}}}\left(\sum_i C_i - D_{\text{tot}}^{\omega}\right) \ge R_{\min} \forall\omega.$$

The composite resilience threshold is set to $R_{\min} = 0.5$, for the case study, which aims to provide a balanced contribution of hub diversification and capacity headroom. Higher values would enforce more conservative designs with additional hubs or excess capacity at increased system cost.

2.4 Rolling-Horizon Operational Control

The operational stage is posed as a rolling horizon control problem that determines the necessary recourse decisions to adapt to changes in short-term demand patterns and long-term forecasts. At each control step, the current system state, including installed capacities, inventories, and realized demand, is observed, and an optimization problem is solved over a finite prediction

horizon using the latest demand forecasts. Only decisions within the control horizon are implemented, after which the horizon is shifted forward and the procedure is repeated. This rolling-horizon structure enables the system to respond dynamically to short-term demand fluctuations and operational disruptions, while remaining consistent with long-term strategic investment decisions made in the first stage. Importantly, the rolling-horizon layer does not modify infrastructure decisions directly but operates within the feasible region defined by the strategic design layer, exploiting available production and transport capacity across energy hubs as imposed by resilience constraints. Operational disruptions are modeled as time-dependent availability multipliers that erode the effective capacity of affected assets during disruption windows, scaling renewable availability, process throughput capacity, and/or transport capacity at the corresponding hub. At time τ_0, given measured inventory and installed capacities, the optimization problem is solved over the prediction horizon, and the resulting decisions within the control horizon are implemented. For the case study, the prediction horizon is selected to be 24h while the control horizon is 6h.

2.5 Learning Layer and Strategic Recourse

The learning layer is the component that enables the system to learn from shifting demand patterns and disruptions to adaptively reconfigure its structure accordingly to ensure future feasibility. A bandit-based learning and monitoring component tracks system-level indicators derived from the rolling-horizon operational layer and evaluates whether observed deviations represent transient fluctuations or sustained structural inadequacy. Rather than directly controlling operations, the learning layer determines when strategic recourse action is warranted over predefined reassessment intervals (six months in the present study).

At each reassessment epoch k, the learning layer constructs a compact state vector,

$$\mathbf{s}_k = [\,\bar{u}_k, \sigma_{u,k}, \xi_k, \phi_k, \Psi_k\,],$$

which summarizes backward-looking service stress and forward-looking adequacy risk based on realized demand, forecast accuracy, and projected demand trends.

When sustained threshold violations are observed across reassessment periods, the learning layer activates a strategic recourse decision. In response, the two-stage stochastic design problem is re-solved using updated demand forecasts and newly generated demand scenarios, while respecting previously installed capacities and construction lead-time constraints. Recourse actions constrain the feasible investment set and include holding the current design, expanding capacity at existing hubs, or enabling new hub investments..

Table 2 State Vector Components

$\bar{u}_k = \frac{1}{W}\sum_{t=k-W+1}^{k}\sum_{l} u_{l,t}$	average unmet demand over window W
$\sigma_{u,k} = \mathrm{std}(u_{k-W+1},\ldots,u_k)$	Volatility of unmet demand
$\xi_k = \frac{\sum_l K_l^{\text{used}}}{\sum_l K_l^{\text{installed}}}$	capacity utilization
$\phi_k = \mathrm{MAPE}(\hat{D}_{k-W},\ldots,\hat{D}_k; D_{k-W},\ldots,D_k)$ *Mean Absolute Percentage Error	forecast accuracy
$\Psi_k = \frac{\hat{D}_k - D_{k-W}}{D_{k-W}}$	demand trend estimate

Table 3 Strategic Recourse Actions. The lite RL action $a_k \in \mathcal{A}$ recourse actions:

Action	Effect on Planning
HOLD	$\Delta X_k = 0, z_{lk} = z_{l,k^-(k)}$
EXPAND	$\Delta X_k \geq 0, z_{lk} = z_{l,k^-(k)}$
NEW Hub	$\Delta X_k \geq 0, z_{lk} \in \{0,1\}$ with $z_{l,k^-(k)} \leq z_{lk}$

The recourse decision is determined by using a contextual bandit formulation. Given the state vector $\mathbf{s}_k$, the selected action maximizes an upper confidence bound on the expected reward:

$$a_k = \arg\max_{a\in\mathcal{A}}\left(\hat{\theta}_a^{\top}\mathbf{s}_k + \alpha\sqrt{\mathbf{s}_k^{\top}A_a^{-1}\mathbf{s}_k}\right),$$

where the action-specific parameters $\hat{\theta}_a$ are updated online based on system performance following previous recourse decisions. The reward associated with action $a_k$is observed after investment effects become operational and is defined as follows at epoch $k+W$ (after investments from k are operational):

$$r_k = -\left(C^{inv}(\Delta X_k) + \sum_{t=k+L}^{k+L+W-1}\left[J_t^{op} + \rho\, U_t^{real}\right] + \lambda_{sw} \cdot 1\left[a_k \neq a_{k^-(k)}\right]\right)$$

The reward signal combines improvements in service adequacy with investment cost penalties, encouraging actions that reduce persistent unmet demand while avoiding unnecessary or overly aggressive investments. Through repeated interaction, the learning layer adapts the mapping between observed stress patterns and preferred strategic responses.

The overall information flow and interaction

between the forecasting, planning, operational control, and learning layers are summarized in the schematic shown below:

$$\begin{array}{rcl}
\text{Layer 0 (Forecasting)} & \xrightarrow{\Omega_k, \{\pi_{\omega,k}\}, \{D_{l,\tau,\omega,k}\}} & \text{Layer 1 (Planning)} \\
\text{Layer 1 (Planning)} & \xrightarrow{K^{RE}, K^{P}, Q, S, z} & \text{Layer 2 (MPC)} \\
\text{Layer 2 (MPC)} & \xrightarrow{U_t^{real}, J_t^{op}} & \text{Layer 3 (RL-lite)} \\
\text{Layer 3 (RL-lite)} & \xrightarrow{a_k \in \{\text{HOLD, EXPAND, NEW}\}} & \text{Layer 1 (Planning)}
\end{array}$$

ILLUSTRATIVE CASE STUDY

An illustrative renewable energy supply chain case study motivated by the increasing energy demand of highly industrialized European economies is considered. Germany is selected as the demand region due to its ambitious decarbonization targets, limited domestic variable renewable energy potential, and projected reliance on imported renewable energy. The supply chain planner aims to meet the monthly power demand discrepancies with respect to renewable energy generation and demand through a dedicated energy supply chain network considering candidate hub locations of Spain, Norway, Morocco, Egypt, and Saudi Arabia. The hourly renewable energy profiles of these locations are utilized in the multi-scale model to design and operate an energy supply chain network consisting of multiple energy generation hubs with optimal renewable portfolios, process facilities to produce dense energy carriers, and transportation infrastructure to meet the monthly energy demand.

The framework is evaluated through a series of instances that focus on the role of individual modeling components. First, a baseline case without resilience constraints is solved to find the lowest-cost transmission strategy for a known demand of 10 TWh per month. Second, resilience constraints that promote geographic diversification and capacity spread, as discussed in Section 2.3, are introduced to create a resilient design for the same demand. This instance aims to show the cost differential of the resilient network design vis-à-vis the minimum cost design. Third, demand uncertainty is incorporated using a two-stage stochastic optimization framework, in which demand forecasts and scenarios are generated from historical energy demand data (2019–2021) of Germany to inform strategic planning for system operation commencing in 2024. Finally, the complete intelligent and adaptive framework is evaluated by activating all modeling layers in the integrated closed loop setting. Building on the resilient network design under uncertainty obtained using demand scenarios generated from historical energy demand data, for system deployment starting in 2024, the framework jointly executes rolling-horizon forecasting, stochastic planning, operational control, and learning-based strategic recourse. Following commissioning, the system operation is simulated under a synthetic demand pattern that shows a steady increase trend, from 10 to 12 TWh/month over a six-month period with superimposed Gaussian noise to test the recourse decisions. This setup allows for demonstrating learning-triggered strategic actions in response to ongoing capacity shortfalls. In summary, four increasingly complex configurations are analyzed: (1) a deterministic cost-minimizing design; (2) a resilient deterministic design; (3) a stochastic resilient design that accounts for demand uncertainty; and (4) an intelligent and adaptive design and operation with learning-based reconfiguration. This progressive structure helps to capture the added value of resilience, stochastic planning, and adaptive responses.

RESULTS AND DISCUSSION

We provide the results of the four progressive cases to illustrate the value of resilience and adaptive operation.

Deterministic Demand Minimum Cost Design: The baseline energy transmission solution represents the minimum-cost network design obtained under the assumption of perfect information on renewable energy demand of 10 TWh/month. Under these conditions, the model invests in electrolyzer and hydrogen liquefaction facilities in Spain powered by solar photovoltaics and battery storage and 11 ships to transport liquid hydrogen to meet the monthly demand. However, the simulation of a random disruption lasting 72 hours in energy hub shows that the system experiences a severe shortfall with nearly 75% of demand left unmet and a substantial value of lost load.

Deterministic Resilient Network Design: In the second case of supply chain design under known demand, we utilize only the two-stage stochastic program component with the composite resilience requirement of R_{min} = 0.5 as explained in Section 2.3 with multiple and geographic diversified energy hubs. The optimal resilient supply chain for the known demand consists of two energy hubs located in Spain and Egypt, with Spain producing and exporting liquid hydrogen and Egypt supplying ammonia as depicted in Figure 2. The network operates 10 vessels from Spain and 2 from Egypt to meet the monthly demand. Despite the added diversification, the total system cost is observed to increase by only about 2% relative to the minimum-cost deterministic solution.

Moreover, when the same disruption is simulated at the Egypt hub as in the non-resilient design to observe if the operational recourse actions suffice to meet demand, the network ramps up production and transportation from Spain within capacity limits, compensating for the temporary loss of supply as demonstrated in Figure 3. As

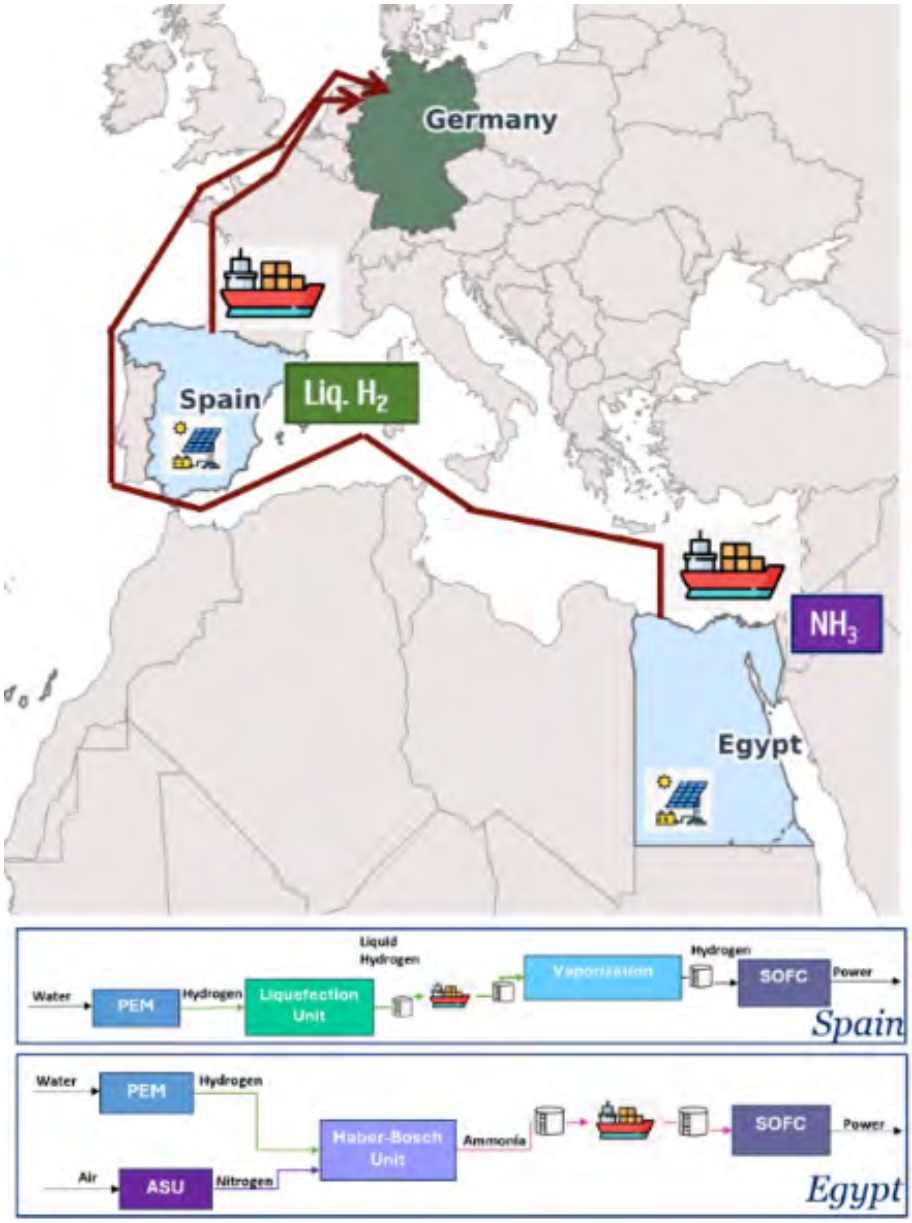

Figure 2. Resilient SC Network Configuration

operations at the Egypt hub recover, production gradually rebalances across both locations, and monthly demand is fully satisfied throughout the disruption period. This response confirms that a diversified and flexible network can absorb shocks to a certain extent and recover smoothly, reinforcing the idea that resilience through redundancy and geographic diversification is achievable at relatively low cost.

Stochastic Resilient Design: When demand uncertainty is introduced, the forecasting layer generates probabilistic demand scenarios that are embedded in the stochastic optimization model. The resulting network retains two production hubs in Spain and Egypt, consistent with the resilient design under known demand, but with moderately increased installed capacities to hedge against uncertainty. In this configuration, approximately 89% of demand is supplied from Spain and 11% from Egypt, reflecting the increased capacity at the Egyptian hub. Relative to the resilient deterministic design, accounting for demand uncertainty leads to an approximate 10% increase in total system cost. The cost comparisons of deterministic, deterministic resilient and stochastic resilient cases are depicted in Figure 4 with the selected energy carriers and locations.

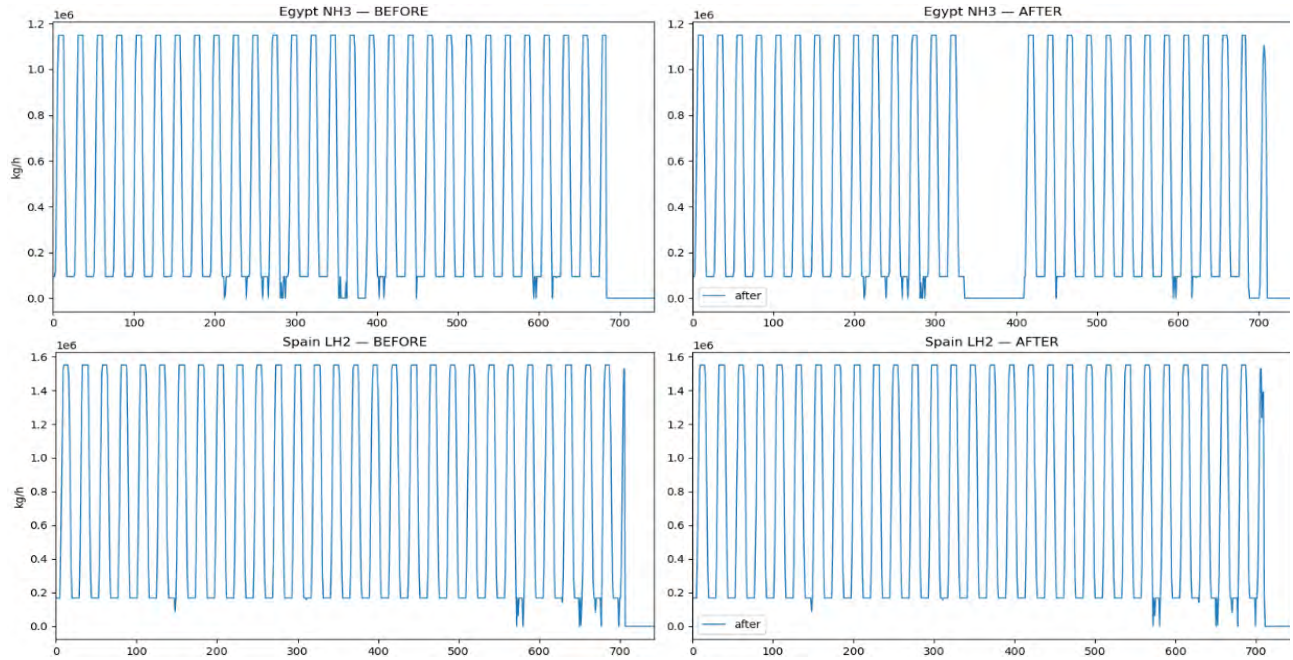

Figure 3. Recourse action to a hub disruption

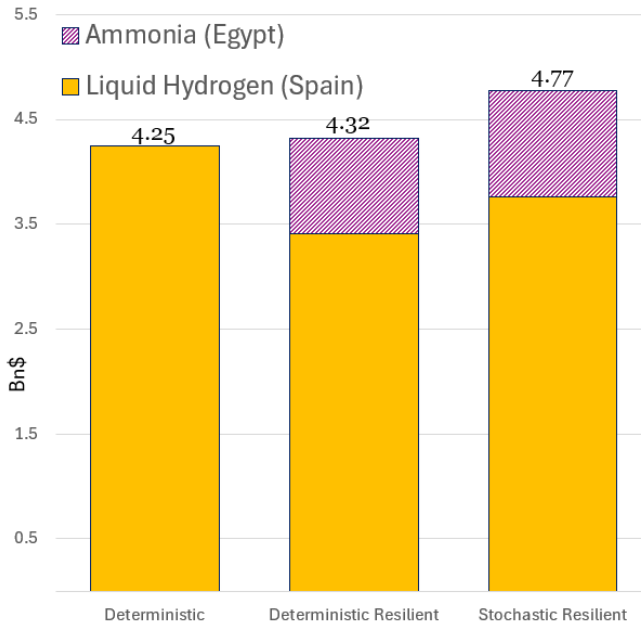

Figure 4. Design cost comparison

Intelligent and Adaptive Network Operation: Finally, the complete intelligent and adaptive framework is evaluated under evolving demand conditions to observe investment recourse decisions. First, probabilistic demand scenarios are generated using historical energy demand data (2019–2021) and provided as inputs to the two-stage stochastic optimization model described in Section 2.3. Solving this model yields an initial resilient network design comprising liquid hydrogen production in Spain supported by 10 ships and ammonia production in Egypt supported by 2 ships, corresponding to a mean projected monthly demand of approximately 10 TWh, which is consistent with the results of the resilient design under demand uncertainty. Once the strategic design is fixed, system operation is then simulated using the rolling-horizon control formulation described in Section 2.4. As the supply chain network becomes operational, a random demand data set is introduced. The simulated demand data set exhibits a persistent increasing trend from 10 to 12 TWh/month over a six-month period with superimposed Gaussian noise as depicted in Figure 5.

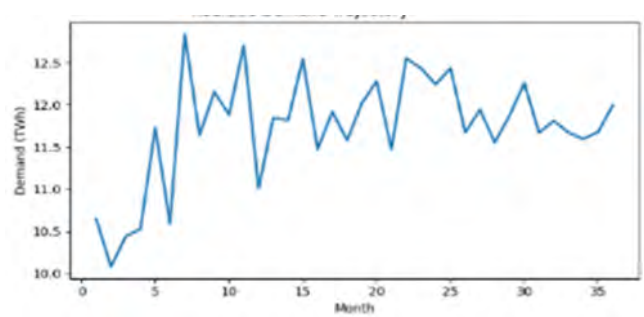

Figure 5. Simulated demand profile

At each operational step, the current system state including inventories, energy storage levels, and committed transportation flows, is observed, and short-term demand forecasts are incorporated to optimize production, storage, and transportation decisions over a finite prediction horizon. Only decisions within the control horizon are implemented before the horizon is advanced,

ensuring operational adaptability, while strictly respecting the production and transportation capacities dictated by the strategic design. At the sixth month of operation, corresponding to the predefined reassessment interval, system-level performance indicators obtained from the operational layer are evaluated. These include both backward-looking measures of persistent service stress, such as rolling unmet demand ratios, and forward-looking adequacy indicators reflecting the rolling-horizon controller's predicted inability to meet future demand under the installed infrastructure. After six months, the rolling unmet demand ratio approaches 25%, while both realized and forecasted future inadequacy indicators exceed approximately 6.7%, indicating that no feasible operational policy can satisfy projected demand within the existing network configuration. This persistent and forecast-consistent capacity inadequacy activates the learning-based recourse layer described in Section 2.5. The learning layer, implemented as a multi-armed bandit algorithm, evaluates admissible strategic actions, including maintaining the current design, expanding capacity at existing hubs, or enabling new hub investments, and triggers a re-solution of the two-stage stochastic design problem using updated demand scenarios. Re-solving the design problem under revised scenarios with a mean monthly demand of 11.5 TWh results in the activation of a new liquid hydrogen production hub in Morocco, along with an additional transportation asset. The redistribution of demand shares across energy hubs before and after the recourse action is shown in Figure 6. The introduction of the Moroccan hub improves both capacity adequacy and the geographic diversification of the supply chain, demonstrating the framework's ability to adapt strategic investment decisions over time by learning from realized demand patterns and proactively maintaining long-term energy security.

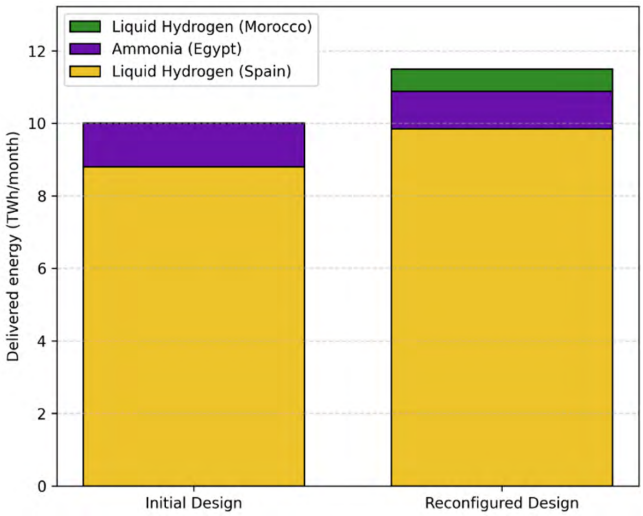

Figure 6. The final demand share of energy hubs

CONCLUSION

This work presents an initial attempt towards the design and operation of intelligent and adaptive energy networks capable of responding to the increasing volatility and uncertainty characteristic of modern energy systems. By integrating hybrid machine-learning forecasting, multi-period stochastic optimization, and rolling-horizon operational control with adaptive recourse mechanisms, the framework addresses critical gaps in existing approaches that rely on static planning paradigms. The proposed methodology enables energy supply chains to function as intelligent systems that continuously predict, optimize, learn from demand patterns, potentially anticipate disruptions, and adapt their configuration and operations as conditions evolve. Through its multi-layered architecture, the framework balances cost efficiency with resilience and security of supply, a balance increasingly vital for modern energy systems operating with increased variable renewable energy share and geopolitical uncertainties. The illustrative case study demonstrates the framework's applicability to a multi-carrier energy supply chain network, showing how the integration of predictive modeling, stochastic optimization, and adaptive control can transform energy supply chains from reactive networks requiring episodic manual intervention into proactive, data-informed systems capable of maintaining reliable operations in the face of sustained uncertainty and disruption. Future work will focus on enriching the learning layer with causal and context-aware signals to distinguish between transient operational shocks and persistent structural mismatches requiring new investments. While demonstrated here for renewable energy supply chains, the proposed closed-loop forecast, optimize, learn and adapt paradigm naturally extends to other large-scale systems with evolving uncertainty, including financial portfolios, critical material supply chains, and infrastructure planning.

ACKNOWLEDGEMENTS

The authors express their gratitude for the financial support from Texas A&M University and Texas A&M Energy Institute, United States. This work was partially supported by the National Science Foundation Convergence Accelerator Program under the Grant titled "Securing Critical Material Supply Chains by Enabling Photovoltaic Circularity (SOLAR)" with Grant No. 49100423C0005.

REFERENCES

1. Iseri H, Iseri F, El-Halwagi M, Iakovou E, Pistikopoulos EN. A multi-objective decision-making framework for renewable energy transportation. International Journal of Hydrogen

Energy 166:150789 (2025). https://doi.org/10.1016/j.ijhydene.2025.150789
2. Ierapetritou, M.G., Pistikopoulos, E.N., 1994. Novel optimization approach for stochastic planning models. Industrial & Engineering Chemistry Research 33(8), 1930–1942.
3. You F, Grossmann IE. Design of responsive supply chains under demand uncertainty. Computers & Chemical Engineering 32:3090-3111 (2008). https://doi.org/10.1016/j.compchemeng.2008.05.004
4. Hugo A, Pistikopoulos EN. Environmentally conscious long-range planning and design of supply chain networks. Journal of Cleaner Production 13:1471-1491 (2005). https://doi.org/10.1016/j.jclepro.2005.04.011
5. Kopanos GM, Pistikopoulos EN. Reactive scheduling by a multiparametric programming rolling horizon framework: a case of a network of combined heat and power units. Ind. Eng. Chem. Res. 53:4366-4386 (2014). https://doi.org/10.1021/ie402393s
6. Ning C, You F. Optimization under uncertainty in the era of big data and deep learning: when machine learning meets mathematical programming. Computers & Chemical Engineering 125:434-448 (2019). https://doi.org/10.1016/j.compchemeng.2019.03.034
7. Iseri F, Iseri H, Chrisandina NJ, Iakovou E, Pistikopoulos EN. Ai-based predictive analytics for enhancing data-driven supply chain optimization. J Glob Optim : (2025). https://doi.org/10.1007/s10898-025-01509-1
8. Iseri F, Iseri H, Shah H, Iakovou E, Pistikopoulos EN. Planning strategies in the energy sector: integrating bayesian neural networks and uncertainty quantification in scenario analysis & optimization. Computers & Chemical Engineering 198:109097 (2025). https://doi.org/10.1016/j.compchemeng.2025.109097
9. Iseri H, Iseri F, El-Halwagi M, Iakovou E, Pistikopoulos E. An enhanced decision-making framework for designing renewable energy supply chains. Lecture Notes in Networks and Systems :722-732 (2025). https://doi.org/10.1007/978-3-031-98304-7_78

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Powering AI Beyond the Grid: Optimal allocation and Behind the Meter Investment Portfolios for Data Centers

Mohamed Abdelhady[ad], Eleftherios Iakovou[bd], Efstratios N. Pistikopoulos[cd]

[a] Department of Multidisciplinary Engineering, Texas A&M University, College Station, TX, USA
[b] Department of Engineering Technology and Industrial Distribution, Texas A&M University, College Station, TX, USA
[c] Artie McFerrin Department of Chemical Engineering, Texas A&M University, College Station, TX, USA
[d] Texas A&M Energy Institute, Texas A&M University, College Station, TX, USA
* Corresponding Author: eiakovou@tamu.edu

ABSTRACT

The rapid expansion of AI data centers is straining electricity grids alarmingly, forcing data center planners to navigate two-pronged challenges: (1) lengthy interconnection queue delays undermining immediate grid access, and (2) volatile electricity prices that spike dramatically during high demand events. This convergence forces planners to reconsider traditional grid-only strategies. While behind-the-meter (BTM) generation offers a solution, existing research lacks comprehensive frameworks for identifying technology portfolios under combined uncertainties of grid access delays and market volatility. This study develops a two-stage stochastic optimization framework with binary capacity constraints co-optimizing data center location and BTM energy portfolios under these challenges. The model evaluates conventional (gas turbines), renewable (solar, wind, batteries), and emerging technologies (hydrogen fuel cells, small modular reactors) across four progressive scenarios spanning emission targets, demand flexibility, grid curtailment, land constraints, and queue delays, contrasting stochastic and deterministic solutions. Applied to a 5 GW data center expansion in ERCOT, three insights emerge: First, queue delays drive 2.7 GW bridging investments that transition to 92% grid reliance once interconnection is available, while stochastic optimization maintains higher BTM utilization against price volatility. Second, demand flexibility reduces required BTM capacity through load shifting during grid curtailment events, decreasing gas deployment from 2.1 to 1.7 GW. Third, land-constrained decarbonization shifts from wind-dominated portfolios to capital-intensive solar-hydrogen-SMR solutions, with stochastic optimization tripling SMR deployment to prioritize reliability under uncertainty.

Keywords: Data Centers, Energy Portfolio Optimization, Small Modular Reactors, Green Hydrogen, Power Grid

INTRODUCTION

Artificial intelligence is driving unprecedented electricity demand, with data centers projected to double energy consumption by 2026 [1]. In the Electric Reliability Council of Texas (ERCOT), interconnection requests surged from 2.6 GW (2022) to more than 200 GW (2030), extending grid connection congestion and challenges to connect on time [2]. Compounding this crisis, new regulations mandate load curtailment during grid stress events (Texas SB6 [3]), while electricity prices in energy-only markets fluctuate from $20/MWh to $5, 000/MWh [4]. These converging pressures—interconnection delays, mandatory curtailment risk, and price volatility—are forcing data center planners to fundamentally rethink their energy strategies [5]. In response, planners are exploring spatial diversification across regions combined with behind-the-meter generation portfolios to mitigate grid dependency and manage operational risks [6].

Behind-the-meter (BTM) generation offers a pathway to overcome grid constraints, but introduces complex multi-dimensional planning challenges. Planners must simultaneously decide: (1) where to locate facilities across regions with heterogeneous characteristics—zones differ in renewable resource quality, land availability, grid queue lengths, and regulatory requirements; (2) which technologies to deploy, portfolios may include solar, wind, gas turbines, small modular reactors (SMRs),

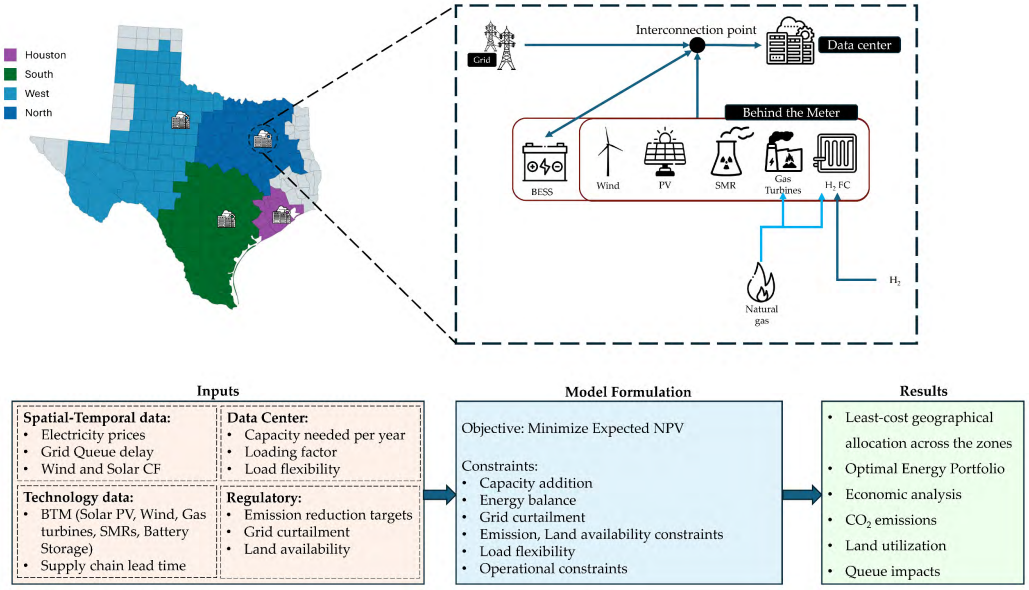

Figure 1. Methodological Framework.

fuel cells, and battery storage, each with distinct capital costs, lead times, and operational characteristics; and (3) how to operate under uncertainty, both grid interconnection delays and electricity prices are inherently stochastic, with queue timelines varying from months to years and market prices exhibiting extreme volatility. Additionally, facilities can leverage load flexibility to shift demand temporally, adding another operational dimension. Traditional deterministic planning frameworks fail to capture these interdependencies and uncertainties, potentially leading to suboptimal investments that either oversize capacity or underestimate exposure to regulatory and market risks.

This paper develops a two-stage stochastic optimization framework for multi-zone data center energy planning under uncertainty. Our model integrates spatial facility allocation, technology portfolio selection, and operational planning while explicitly accounting for: (1) stochastic grid interconnection queue delays across zones and time periods; (2) electricity price uncertainty stemming from historical market volatility; (3) zone-specific constraints including renewable capacity factors, land availability, and emission reduction targets; (4) technology-specific lead times affecting investment feasibility; and (5) load flexibility enabling temporal demand shifting. Contributions: We provide (1) a computationally tractable stochastic programming formulation for multi-zone energy infrastructure planning; (2) extensive case study analyses demonstrating how optimal portfolios shift under varying interconnection delays, emission policies, and grid curtailment policies; and (3) managerial insights on technology selection, spatial allocation, and operational flexibility for data facility planners.

METHODOLOGY

Problem Definition and System Architecture

We consider a multi-zone capacity expansion and operational planning problem for large-scale data centers subject to grid interconnection constraints, electricity price uncertainty, and emission reduction mandates. The system comprises $|Z|$ zones, each with distinct electricity price profiles, renewable resource availability (solar irradiance and wind patterns), land constraints, and time-varying grid interconnection queue delays. The planning objective is to determine: (1) optimal spatial allocation of facility capacity across zones; (2) behind-the-meter (BTM) generation portfolio composition; (3) energy storage deployment; and (4) grid interconnection capacity requests.

The BTM technology portfolio includes dispatchable sources (natural gas turbines, fuel cells with natural gas or hydrogen, small modular reactors) with high availability factors (90-95%), and variable renewable energy (solar PV, wind) with time-varying capacity factors. Battery energy storage systems provide operational flexibility. Critically, investment decisions must account for technology-specific lead times—the lag between capital commitment and operational availability—ranging from one to several years. Data center load comprises baseload (fixed continuous draw) and flexible load (temporally shiftable within daily windows), enabling demand-side response without compromising service levels.

Figure 1 illustrates the methodological framework. We develop a two-stage stochastic program that determines first-stage capacity investments under

uncertainty, with second-stage operational recourse as conditions realize. The deterministic case (single scenario) serves as a comparison benchmark.

Mathematical Formulation

Key Notation

Sets	
Z	Set of geographic zones
G	Set of BTM generation technologies
T	Hourly time periods
Y	Planning years
Parameters	
$C_g^{\{capex\}}, C_g^{\{FOM\}}, C_g^{\{VOM\}}, C_g^{\{fuel\}}$	Capital, fixed O&M, variable O&M, fuel costs
$P_{\{z,t\}}^{\{grid\}}(\omega), Q_{\{z,y\}}(\omega)$	Electricity price (scenario), queue delay (scenario)
$\lambda^{\{base\}}, \lambda^{\{flex\}}, \rho^{\{emis\}}, L_z^{\text{avail}}, \delta_{z,t}^{curtail}$	Load factor, flexibility fraction, Emission reduction target, Land availability, Grid curtailment
Δ^{DC}	Data center module size [MW]
$MinCap_g$	Minimum capacity addition per technology [MW]
Decision Variables:	
$\text{Cap}_{z,g,y}, CapTotal_{z,g,y}$	Capacity additions [MW], cumulative capacity [MW]
$StorCap_{z,y}$	Storage capacity [MWh]
$DCCap_{z,y}$	Facility capacity [MW]
$\text{Gen}_{z,g,t}(\omega)$	Electricity generation [MW]
$\text{Ch}_{z,t}(\omega), \text{Dis}_{z,t}(\omega)$	Battery charge and discharge in zone z at hour t [MW]
$\text{Flexload}_{z,t,y}(\omega)$	Flexible load [MW]
$Unserved_{z,t,y}(\omega)$	unserved energy [MW]
$y_{z,g,y}$	Binary variable for technology selection decisions
$n_{z,y} \in \mathbb{Z}_+$	Number of data center modules in zone z during year y

Objective Function. The objective minimizes total net present value (NPV) of system costs:

$$\min \sum_{y\in Y^{inv}}(1+r)^{-(y-y^{base})}\left[C^{inv}\cdot \Delta x_y + C^{FOM}\cdot x_y\right] + \sum_{\omega\in\Omega}\pi_\omega \sum_{y\in Y^{op}}(1+r)^{-(y-y^{base})}\left[C^{var}\cdot\left(g_y(\omega) + p_y(\omega)\right)\right] + UnservedPenalty(\omega) \quad (1)$$

Subject to:

$$CapTotal_{z,g,y} = InitCap_{z,g} + \sum_{y'\in Y^{inv}:y'+\tau_g\le y} Cap_{z,g,y'} \quad \forall z,g,y \quad (2)$$

$$\sum_{z\in Z} DCCap_{z,y} \ge D_y^{req} \ \forall z,y \quad (3)$$

$$\sum_{g\in G} Gen_{z,g,t,y}(\omega) + GridPurch_{z,t,y}(\omega) + Dis_{z,t,y}(\omega) - Ch_{z,t,y}(\omega) = Load_{z,t,y}^{base} + FlexLoad_{z,t,y}(\omega) - Unserved_{z,t,y}(\omega) \ \forall z,t,y,\omega \quad (4)$$

$$Gen_{z,g,t,y}(\omega) \le CapTotal_{z,g,y}\cdot CF_{z,g,t} \ \forall z,g,t,y,\omega \quad (5)$$

$$GridPurch_{z,t,y}(\omega) \le \left(1-\delta_{z,t}^{curtail}(\omega)\right)\cdot GridCapAvail_{z,t,y}(\omega) \ \forall z,t,y,\omega \quad (6)$$

$$SOC_{z,t,y}(\omega) = SOC_{z,t-1,y}(\omega) + Ch_{z,t,y}(\omega)\sqrt{\eta^{stor}} - \frac{Dis_{z,t,y}(\omega)}{\sqrt{\eta^{stor}}} \ \forall z,t,y,\omega \quad (7)$$

$$SOC_{(z,t,y)(\omega)} \le StorTotal_{z,y} \quad (8)$$

$$Ch_{(z,t,y)(\omega)} \le \frac{StorTotal_{z,y}}{d^{stor}} \ \forall z,t,y,\omega \quad (9a)$$

$$Dis_{(z,t,y)(\omega)} \le \frac{StorTotal_{z,y}}{d^{stor}} \ \forall z,t,y,\omega \quad (9b)$$

$$MinCap_g \cdot y_{z,g,y} \le Cap_{z,g,y} \ \forall z,t,y \quad (10)$$

$$Cap_{z,g,y} \le M\cdot y_{z,g,y} \ \forall z,t,y \quad (11)$$

$$DCCap_{z,y} = \Delta^{DC}\cdot n_{z,y} \ \forall z,y \quad (12)$$

$$y_{z,g,y} \in \{0,1\} \ \forall z,g,y \quad (13)$$

$$n_{z,y} \in \mathbb{Z}_+ \ \forall z, y\in Y^{inv} \quad (14)$$

$$All\ capacity\ and\ operational\ variables \ge 0 \quad (15)$$

The objective function (1) minimizes the expected net present value (NPV) of system costs across all scenarios, where r denotes the discount rate, Δx_y represents capacity additions in year y (aggregated across zones and technologies), and x_y is cumulative installed capacity. The terms $g_y(\omega)$ and $p_y(\omega)$ denote BTM generation and grid purchases under scenario ω, with π_ω representing scenario probabilities. The formulation captures annualized capital expenditures (C^{inv}, fixed O&M C^{FOM}, and variable operating costs including fuel and grid electricity procurement at scenario-specific prices (C^{var}).

Constraints (2) track capacity evolution with technology-specific lead times τ_g, ensuring capacity invested in year y' only becomes operational after the lead time elapses. Constraints (3) require total installed facility capacity to satisfy demand targets across all zones.

Equations (4) maintain hourly supply-demand equilibrium in each zone under each scenario, where $Load_{\{z,t,y\}}^{base}$ represents the fixed baseload component and

$FlexLoad_{(z,t,y)}(\omega)$ is the flexible load. Constraint (5) limits technology output by installed capacity and resource availability through the capacity factor $CF_{z,g,t}$.

Constraints (6) govern grid purchases, which are prohibited during curtailment events and constrained by time-varying interconnection capacity. The curtailment indicator $\delta_{z,t}^{curtail}(\omega)$ equals unity when electricity prices exceed zone-specific high percentile thresholds (e.g., 98th-99th percentile). Grid interconnection capacity $GridCapAvail_{(z,t,y)}(\omega)$ evolves based on queue delay schedules $Q_{(z,y)}(\omega)$ during the delay period in year y, only previously committed capacity is available; after the delay expires, newly requested capacity becomes accessible.

Equations (7)-(9) govern battery storage dynamics. Equation (7) tracks state-of-charge evolution with round-trip efficiency $\sqrt{\eta^{stor}}$, while constraints (8) and (9) enforce energy and power capacity limits respectively. The parameter d^{stor} denotes the power-to-energy duration (e.g., 4 hours for typical BESS).

Constraints (10) enforce minimum capacity requirements for selected technologies (gas turbines and SMRs), the capacity addition must be at least $MinCap_g$ (e.g., 100 MW). If no capacity is added ($y_{z,g,y} = 0$), then $Cap_{z,g,y} = 0$. While constraints (11) enforce the upper bound for capacity additions.

Constraints (12) enforce modular data center capacity: facility capacity must be integer multiples of a standard module Δ^{DC} (e.g., 200 MW), reflecting the discrete nature of data center construction. This forces data center capacity to take values of 0, 200, 400, 600, ... MW. The integer variable $n_{z,y}$ represents the number of modules deployed.

The model also includes constraints for load flexibility (daily energy budgets ensuring $\sum_{t \in T_d} FlexLoad_{z,t,y}(\omega)$ meets required daily flexible consumption with instantaneous capacity limits), emission reduction targets (expected value across scenarios constraining total annual emissions to $\rho^{emis} \cdot E_y^{baseline}$), land availability (zone-specific footprint limits), and storage cycling (end-of-year SOC matching initial conditions).

The formulation represents a two-stage stochastic program where each scenario $\omega \in \Omega$ specifies realizations of: (1) hourly electricity prices $P_{z,t}^{grid}(\omega)$ derived from historical years representing diverse market conditions (e.g., extreme weather events, high renewable penetration, volatile fuel prices), and (2) grid interconnection queue delays $Q_{z,y}(\omega)$ representing optimistic (short delays) versus pessimistic (extended delays) timelines. First-stage decisions (capacity investments $Cap_{z,g,y}, StorCap_{z,y}, DCCap_{z,y}, GridCap_{z,y}$, and discrete decisions $y_{z,g,y}, n_{z,y}$) are made before uncertainty realization; second-stage decisions (operational dispatch $Gen_{(z,g,t,y)}(\omega), GridPurch_{z,t,y}(\omega)$, storage operations, flexible load scheduling) adapt to realized conditions with scenario-specific recourse.

CASE STUDY

We apply the proposed framework to ERCOT, a deregulated energy-only market serving approximately 90% of Texas electricity demand. ERCOT presents an ideal testbed due to: (1) Electricity price volatility; (2) rapid data center load growth; (3) substantial renewable penetration exceeding 35% of generation capacity [2]; and (4) significant spatial heterogeneity in resource availability and grid constraints. The analysis considers four ERCOT load zones, HOUSTON, WEST, NORTH, and SOUTH, as shown in Figure 1.

The BTM technology portfolio encompasses solar PV, wind turbines, natural gas turbines (GT), fuel cells (natural gas and hydrogen-powered), small modular reactors (SMR), and battery energy storage (4-hour lithium-ion). Technology specifications follow NREL Annual Technology Baseline (ATB) 2023 [7], including capital costs, fixed and variable O&M, fuel costs, conversion efficiencies, and land footprint requirements. We model data center deployment planning with phased capacity expansion targeting 2.0 GW (2028), 4.0 GW (2029), and 5.0 GW (2030). Data center load comprises baseload (determined by load factor λ^{base}) and flexible load (flexibility ratio λ^{flex}), where $\lambda^{flex} = 0$ represents pure baseload demand.

Historical hourly day-ahead LMP data for each zone spanning are obtained from ERCOT market archives [8]. For stochastic optimization, scenarios combine historical price years with queue delay realizations. Grid curtailment periods are defined as hours exceeding zone-specific percentile thresholds ($\delta_{z,t}^{curtail}$), where grid purchases are prohibited.

We evaluate four progressive scenarios to isolate the impacts of key operational and regulatory constraints on technology selection and spatial deployment. All scenarios include grid interconnection queue delays (ranging from 0-9 months across zones), which represent the baseline infrastructure constraint. The scenarios progressively layer additional constraints:

Table 1: Scenario Definitions.

Parameter	S1	S2	S3	S4
Grid Queue	✓	✓	✓	✓
Curtailment	-	2%	2%	2%
Emission Reduction	-	-	70%	70%
Land Constraint	-	-	-	Limited
Demand Type	Base	Flex	Flex	Flex

All scenarios include grid interconnection queue

delays as a baseline infrastructure constraint. S1 (Baseline) isolates the impact of interconnection delays alone under pure baseload operation. S2 introduces grid curtailment (2% of hours at 98th percentile prices) and demand flexibility (80% load factor, 30% shiftable). S3 adds a 70% emission reduction target. S4 represents the most constrained case with limited land availability.

The stochastic optimization framework accounts for two primary sources of uncertainty: electricity price volatility and grid interconnection queue delays. These uncertainties are modeled independently from the sensitivity scenarios (S1-S4) described above, each sensitivity scenario is solved under both deterministic and stochastic frameworks.

Price Uncertainty: We construct price scenarios from three representative historical years of ERCOT day-ahead LMP data: 2021 (extreme prices including Winter Storm Uri), 2022 (volatile market conditions), and 2024 (normal operations). Each price scenario represents 8, 760 hourly prices across all four zones, capturing diverse market regimes including extreme weather events, high renewable curtailment periods, and fuel price volatility.

Queue Delay Uncertainty: Grid interconnection timelines are modeled with two realizations reflecting observed processing variability in ERCOT. The low queue (optimistic) scenario features delays decreasing from 9 months (2028) to 0 months (2030) as backlogs clear, while the extreme queue (pessimistic) scenario exhibits persistent delays from 15 months (2028) to 6 months (2030), representing prolonged interconnection congestion.

Scenario Set Construction: The complete stochastic scenario set Ω combines the three price years with two queue delay realizations, yielding six scenarios. Probability weights reflect the relative likelihood of market conditions: 2024 normal operations receive 60% weight (30% per queue scenario), 2022 volatile conditions receive 30% weight (15% per queue scenario), and 2021 extreme events receive 10% weight (5% per queue scenario). Within each price regime, queue delays are equally likely, capturing the inherent uncertainty in interconnection processing timelines.

Deterministic Baseline: For comparison, the deterministic model uses 2024 prices (normal operations) with the low queue scenario (9→0 months), representing a median planning assumption without explicit hedging against uncertainty.

To maintain computational tractability while preserving seasonal and operational variability, both deterministic and stochastic models employ temporal decomposition with representative days. We select 4 representative days per month using scenario-weighted K-means clustering on daily price profiles, wind capacity factors, and solar capacity factors. Representative days are weighted to preserve annual energy balances, with weights determined by cluster membership (number of days each representative day represents).

Both the deterministic and stochastic models are implemented in Python using Pyomo and solved with Gurobi.

RESULTS AND DISCUSSION

This section presents the optimization results in two parts: first, we establish deterministic baseline outcomes and sensitivity analysis across the four scenarios; second, we examine stochastic optimization results that account for price and queue delay uncertainty.

Deterministic Baseline Results

Figure 2 presents the deterministic optimization results across the four sensitivity scenarios, showing BTM technology capacity (2a), energy generation mix (2b), spatial distribution (2c), and battery storage deployment (2d). All deterministic cases use 2024 electricity prices and the low queue delay scenario (9→0 months) as the baseline planning assumption.

HOUSTON is excluded across all scenarios due to prohibitively high electricity prices. In the unconstrained baseline (S1), Figure 2c shows capacity concentrating in SOUTH (64%) and NORTH (20%), with WEST capturing 16%. The system deploys 2.1 GW of gas turbines and 0.5 GW of wind (Figure 2a) as interim capacity during queue delays. Once grid interconnection becomes available by 2030, grid purchases dominate at 92.2% of energy (Figure 2b), with gas declining to 3.0%. This "bridging" strategy invests in BTM capacity primarily for 2028-2029 queue periods, then transitions to grid-dominated operations. Storage remains minimal at 0.3 GWh (Figure 2d).

S2's demand flexibility enables complete spatial consolidation—100% of capacity in SOUTH (Figure 2c). The portfolio simplifies to 1.8 GW of gas turbines with zero renewables or storage (Figures 2a, 2d). Despite facing grid curtailment, load shifting enables 97.7% grid reliance (Figure 2b), demonstrating that operational flexibility substitutes for infrastructure diversification.

Emission constraints (S3) trigger spatial redistribution across WEST (52%), SOUTH (40%), and NORTH (8%) to access diverse wind resources. Figure 2a shows deployment of 6.8 GW wind, 1.5 GW solar, and 0.4 GW hydrogen fuel cells. Wind dominates energy supply at 47.5%, with solar contributing 9.3% and hydrogen fuel cells 8.5% (Figure 2b). Grid purchases persist at 35.1%, filling renewable generation gaps while maintaining the 70% emission reduction target. Storage increases to 0.9 GWh for renewable integration.

Land constrained (S4) force concentration in WEST (64%) and a technology shift to high-density solutions. Solar dominates with 4.7 GW due to superior power density (35 MW/km²) versus wind (0.4 GW, 5 MW/km²).

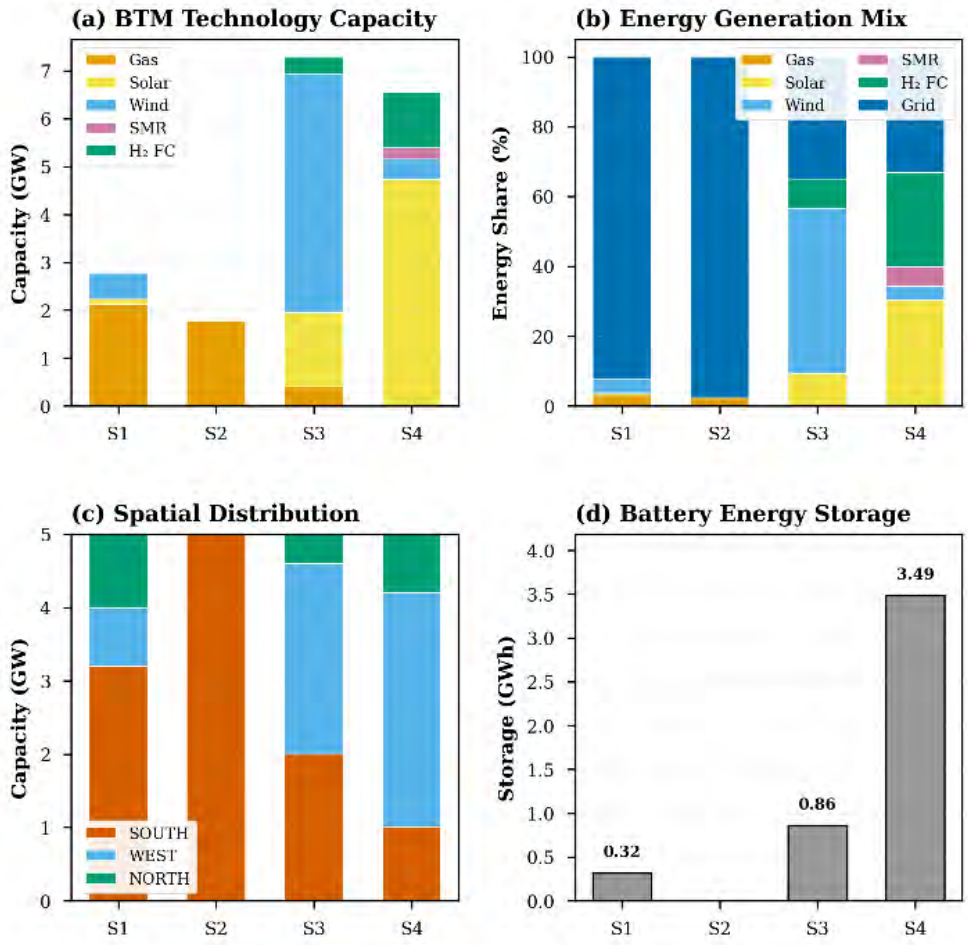

Figure 2. Deterministic optimization results.

Capital-intensive technologies emerge: hydrogen fuel cells (1.2 GW, 27.3% of energy) and SMRs (0.2 GW, 5.7%). Storage quadruples to 3.5 GWh—essential for managing diurnal solar patterns. This demonstrates deep decarbonization feasibility under severe land constraints, albeit requiring substantial storage and capital-intensive generation.

Stochastic Results

Figure 3 presents stochastic optimization results, accounting for price uncertainty (2021, 2022, 2024) and queue delay uncertainty (low vs. extreme scenarios). The stochastic framework fundamentally reshapes spatial allocation and technology selection compared to deterministic baselines.

Uncertainty triggers notable spatial redistribution. In S1, stochastic optimization shifts from deterministic SOUTH-NORTH concentration (64%-20%) to SOUTH-WEST distribution (76%-24%, Figure 3c), eliminating NORTH entirely. This reflects hedging against combined price and queue uncertainty—WEST's wind resources and moderate prices provide superior risk-adjusted returns compared to NORTH's reliance on delayed grid access.

S2 exhibits the most dramatic spatial reversal. Deterministic optimization placed 100% capacity in SOUTH; stochastic optimization distributes across SOUTH (68%) and WEST (32%). This geographic diversification hedges against curtailment events and queue delays by accessing WEST's renewable resources, even in a flexibility-enabled scenario. The stochastic portfolio adds 0.5 GW of wind (Figure 3a) absent in the deterministic case, demonstrating that uncertainty justifies renewable investment despite available flexibility.

Emission-constrained scenarios maintain similar spatial patterns but with subtle shifts. S3 stochastic allocation (WEST 56%, SOUTH 32%, NORTH 8%, HOUSTON 4%) closely mirrors deterministic distribution, confirming that emission constraints dominate spatial decisions regardless of uncertainty. However, S4 under land constraints exhibits a complete geographic flip: deterministic optimization favored WEST (64%), while stochastic optimization concentrates in SOUTH (56%), with WEST reduced to 24%. This reversal reflects uncertainty's

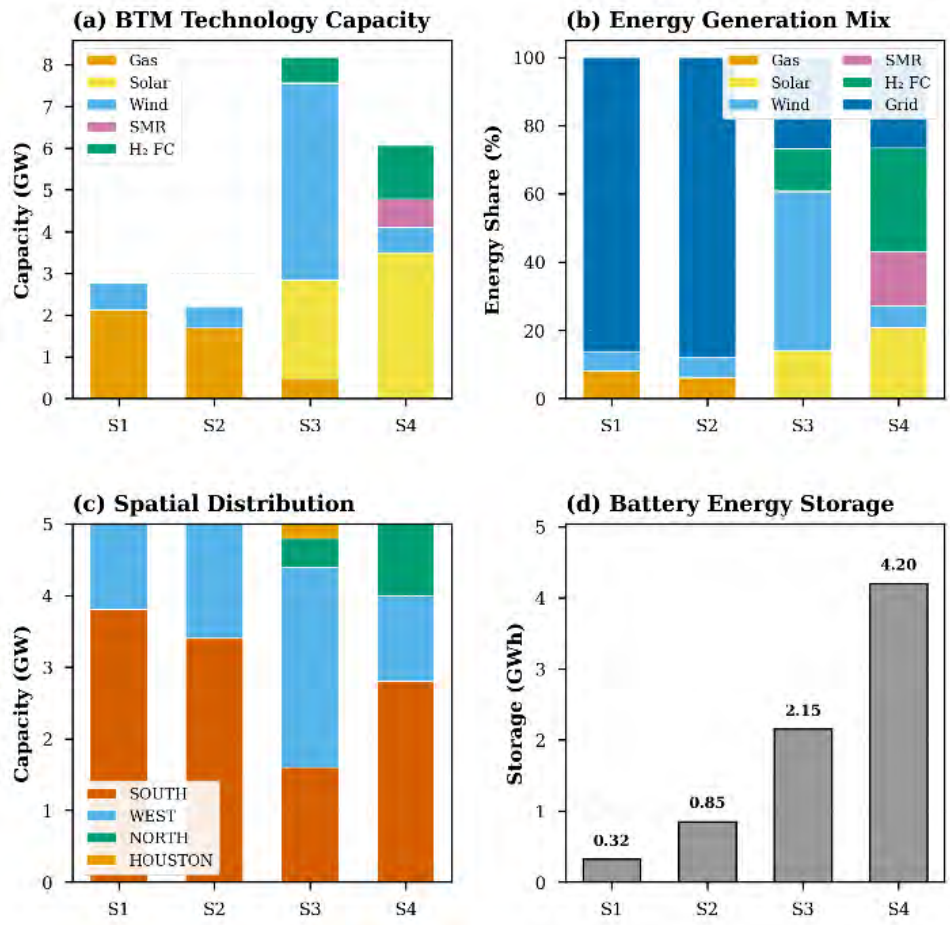

Figure 3. Stochastic optimization results.

interaction with land scarcity, SOUTH's operational characteristics provide superior resilience when both land and grid access are uncertain.

Stochastic optimization reveals distinct technology preferences. In S1, both deterministic and stochastic deploy 2.1 GW of gas and 0.6 GW of wind, but the stochastic energy mix shows gas maintaining 8.0% of generation through 2030 versus the potential for higher grid reliance under perfect foresight. This hedging premium protects against extreme price scenarios.

S2's technology portfolio remains stable: 1.7 GW gas and 0.5 GW wind in both frameworks. Generation shares (gas 5.9%, wind 6.0%, grid 88.1%) demonstrate that flexibility combined with moderate BTM investment provides robust performance across uncertainty scenarios.

S3's technology mix shifts notably under uncertainty. While total wind capacity (5.0 GW) remains comparable, solar increases to 2.4 GW (from 1.5 GW deterministic), and hydrogen fuel cells expand to 0.6 GW (from 0.4 GW). The stochastic energy mix (Figure 3b) shows wind at 46.7%, solar at 13.9%, and hydrogen fuel cells at 12.4%, a more diversified portfolio than wind-dominated deterministic alternatives, reflecting technology balancing to hedge resource variability.

The most significant shift occurs in S4. Stochastic optimization triples SMR deployment to 0.7 GW, with SMR energy share reaching 15.8% (Figure 3b). Solar decreases to 3.5 GW while hydrogen fuel cells increase to 1.3 GW (30.4% of energy). This substitution toward high-capacity-factor technologies (SMRs: 95% availability, hydrogen fuel cells: dispatchable) reflects uncertainty's premium on reliability under tight land constraints. SMRs provide consistent generation regardless of weather or price scenarios, justifying higher capital costs.

CONCLUSIONS AND MANAGERIAL INSIGHTS

This paper developed a two-stage stochastic optimization framework integrating binary capacity constraints, spatial facility allocation, and behind-the-meter energy portfolio selection under electricity price volatility and grid interconnection queue uncertainty. Applied to ERCOT data center planning, the framework reveals how progressive operational and regulatory constraints

reshape technology and spatial strategies.

The results demonstrate that optimal portfolios evolve systematically with constraint tightening. Under queue delays alone (S1), gas turbines serve as bridging capacity during 2028-2029 interconnection periods, with operations transitioning to 92% grid reliance by 2030. Demand flexibility (S2) simplifies portfolios but does not eliminate BTM investment needs, the system maintains 1.7 GW of gas and 0.5 GW of wind to hedge curtailment risk. Emission constraints (S3) trigger fundamental shifts toward spatially distributed wind portfolios (5.0 GW) supplemented by solar and hydrogen fuel cells, achieving 70% reduction while maintaining 27% grid purchases. Land constraints (S4) force capital-intensive, high-density solutions: solar-hydrogen-SMR portfolios with storage quadrupling to 3.5 GWh.

Stochastic optimization reveals that uncertainty drives technology diversification and dispatchable capacity preferences. While spatial patterns remain relatively robust across frameworks, technology selection exhibits significant sensitivity. Under land-constrained decarbonization (S4), stochastic optimization triples SMR deployment (0.2 to 0.7 GW) and increases hydrogen fuel cell energy share to 30.4%, reflecting uncertainty's premium on high-capacity-factor technologies. This substitution toward firm generation quantifies the hedging premium—investments justified by resilience across scenario distributions rather than expected-case optimization.

Among operational strategies, demand flexibility delivers consistent value across scenarios, enabling higher grid reliance and simplified portfolios. However, flexibility alone cannot eliminate infrastructure investment under curtailment and emission constraints. Storage emerges as scenario-specific: negligible in grid-dominated cases, moderate for wind integration, and essential (3.5 GWh) for high-solar land-constrained scenarios, scaling with renewable penetration rather than serving primarily as uncertainty hedge.

Spatial analysis confirms tight coupling between zone selection and technology choice. Emission constraints distribute capacity across zones (WEST 56%, SOUTH 32%, NORTH 8%) to access diverse wind resources, while land constraints concentrate deployment (SOUTH 56%) while forcing technology shifts toward high-density solutions. Binary capacity constraints (100 MW minimums for gas and SMRs, 200 MW data center modules) enforce realistic deployment patterns, ensuring solutions represent implementable strategies.

Future research should model the interaction between behind-the-meter investments and grid expansion planning, capturing how large-scale data center deployment decisions influence grid infrastructure development. Additionally, incorporating renewable energy credit markets and exploring multi-stakeholder coordination between data center operators and utilities would provide comprehensive policy insights for sustainable data center growth.

ACKNOWLEDGEMENTS

This work was partially supported by the National Science Foundation Convergence Accelerator Program under the Grant titled “Securing Critical Material Supply Chains by Enabling Photovoltaic Circularity (SOLAR)” with Grant No. 49100423C0005

AUTHOR IDENTIFIERS

Author ORCIDs:
Abdelhady M: 0000-0001-6524-1735
Iakovou E: 0000-0001-6010-3712
Pistikopoulos S: 0000-0001-6220-818X

REFERENCES

1. International Energy Agency. Electricity 2024: Analysis and Forecast to 2026. IEA (2024).
2. Public Utility Commission of Texas. Report on Electricity Supply and Demand in the ERCOT Region. PUCT (2024).
3. Texas Legislature. Senate Bill 6: relating to electricity planning and infrastructure costs for large loads. Texas Legislature (2025).
4. Grid Status. Exploring extreme prices in ERCOT with Grid Status. (2023). https://blog.gridstatus.io/exploring-extreme-prices-in-ercot-with-grid-status/
5. Abdelhady M, Iakovou E, Pistikopoulos EN, Robertson R. Meeting rising U.S. electricity demand from AI and data centers: an integrated technical and policy perspective. Mosbacher Institute White Paper Series 5:2 (2025) https://bush.tamu.edu/news/mosbacher/rising-us-electricity-demand-ai-data/
6. JLL. https://www.jll.com/en-us/insights/market-outlook/data-center-outlook
7. NREL. Annual Technology Baseline: The 2023 Electricity Update (2023). https://atb.nrel.gov/
8. Electric Reliability Council of Texas. Historical DAM clearing prices for capacity. ERCOT (2024). https://www.ercot.com/mp/data-products/data-product-details?id=NP4-181-ER

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Evaluating the potential of e-fuels for decarbonizing European truck transport: A techno-economic and life cycle approach

Marion Andritz[a*], Severin Sendlhofer[a], Rafailia Mitraki[b], Grégoire Léonard[b], Christoph Markowitsch[a]

[a] Chair of Process Technology and Industrial Environmental Protection, Technical University of Leoben, 8700 Leoben, Austria
[b] Department of Chemical Engineering, University of Liège, 4000 Liège, Belgium
* Corresponding Author: marion.andritz@unileoben.ac.at

ABSTRACT

Heavy-duty road transport remains a challenging sector to decarbonize, as full electrification of long-distance trucking is currently constrained by limitations in energy density and charging infrastructure. Alternative fuels such as hydrogen, biodiesel, and e-fuels are thus gaining increasing attention. In parallel, the cement industry is a major source of unavoidable, process-related CO_2 emissions, offering an opportunity to use captured industrial CO_2 as a feedstock for e-fuel production. This study evaluates the production of e-methanol and Fischer-Tropsch (FT) diesel from captured CO_2 at an Austrian cement plant as a base case. Several system configurations are analyzed, including different electricity supply options across Europe and the use of biogenic versus fossil CO_2. An integrated framework combining process simulation, techno-economic analysis, and life-cycle assessment is applied to compare both fuel pathways. Results show that the climate impact of e-fuels is highly dependent on the electricity mix. When non-renewable electricity is used, climate impacts are 100-440% higher than those of fossil diesel. In contrast, a wind-based Austrian scenario achieves the lowest impact, corresponding to a 44-50% reduction compared to fossil diesel. Overall, cement-plant-based power-to-liquid concepts offer limited near- to mid-term mitigation potential under current European energy conditions and deliver climate benefits primarily when based on biogenic CO_2, high process efficiencies, and low-carbon electricity. The study further highlights key differences between the fuels, with FT-diesel being a certified drop-in fuel, while e-methanol still requires technical and regulatory validation, underscoring the need for technology-specific and system-level assessments.

Keywords: Technoeconomic Analysis, Life Cycle Analysis, Process Design, Aspen Plus, Synthetic Fuels, Methanol, Fischer-Tropsch Synthesis, Power-to-Liquid

INTRODUCTION

The transport sector contributes approximately 30% to the EU's CO_2 emissions, with 72% attributable to road transport [1, 2]. While the decarbonization of domestic transport is progressing with an increasing deployment of electric vehicles, the electrification of long-distance freight transport poses several challenges due to limitations in energy density, vehicle payload, and charging times [3]. Therefore, alternative low-carbon fuels such as hydrogen (H_2), biodiesel, and alcohols are increasingly discussed as complementary solutions for heavy-duty transport applications [4, 5]. Among these options, e-fuels (also known as synthetic fuels produced from captured carbon dioxide (CO_2) and renewable H_2) are widely discussed as promising alternatives to fossil fuels. In this context, Fischer-Tropsch (FT) diesel represents a well-established e-fuel pathway, producing paraffinic diesel-range fuel that is fully compatible with existing engines and infrastructure. In contrast, alcohol-based e-fuels, particularly e-methanol, offer an alternative pathway characterized by comparatively simpler process steps and high system flexibility, as methanol is a well-established platform chemical that can be used both as an energy carrier and as a versatile chemical feedstock. While FT-diesel is functionally seen as a drop-in fuel for heavy-

duty diesel engines and is certified under EN 15940 [6], the large-scale deployment of e-methanol in road transport requires adapted engine concepts and further regulatory validation [7, 8].

Despite their potential, e-fuel production is highly energy-intensive, mainly attributable to the substantial H_2 demand [9]. As H_2 production via electrolysis dominates the energy consumption of the Power-to-Liquid (PtL) pathways, the environmental performance of e-fuels is strongly dependent on the carbon intensity of the electricity supply [10]. In addition to H_2 as key reactant, large quantities of CO_2 are required as a carbon source, shifting attention toward industrial point sources. In this context, the cement industry represents one of the largest sources of unavoidable, process-related CO_2 emissions, accounting for almost 4% of the European CO_2 emissions in 2022 [11]. These emissions originate mainly from the energy intensive calcination of limestone ($CaCO_3 \rightarrow CaO + CO_2$). Captured CO_2 from cement plants can theoretically serve as a potential carbon feedstock for e-fuel production, conceptually linking the decarbonization of both industry and transport sectors.

However, the utilization of fossil-derived industrial CO_2 for e-fuel production is subject to regulatory restrictions. Under the Renewable Energy Directive (RED) restrict, the eligibility of such pathways for renewable fuels of non-biological origin (RFNBOs) is limited to the period up to 2040, after which permanent CO_2 storage through carbon capture and storage (CCS) is required [12]. These regulatory conditions add a system-level constraint that must be considered when assessing cement-based power-to-liquid concepts. Thus, e-fuels represent a promising pathway for the decarbonization of heavy-duty road transport and cement industry, but imply distinct economic and ecological trade-offs, making a comparative assessment essential.

Against this background, the present study investigates whether e-methanol and FT-diesel produced from captured CO_2 at a cement plant can deliver economic and environmental benefits compared to fossil diesel. The assessment is conducted under realistic European electricity supply scenarios and the prevailing regulatory framework, enabling a comparative evaluation of both e-fuel pathways with respect to their potential role in a future low-carbon transport system.

METHODS

System scope and reference case

This study investigates industrial-scale power-to-liquid fuel production integrated into a cement plant. All assessments are based on a hypothetical, steady-state plant configuration representing a commercial-scale installation. The reference system is anchored to a cement plant located in Austria with annual CO_2 emissions of 750 $ktCO_2$/year [13]. Unless stated otherwise, plant scale, capture technology, electrolysis concept, and process design assumptions are kept constant across all cases to enable a consistent comparison.

Process routes considered

The investigated system includes CO_2 capture from cement plant flue gas using a monoethanolamine (MEA)-based solvent scrubbing unit and H_2 production via alkaline electrolysis (AEL), which are applied uniformly across all cases. MEA-based CO_2 capture is used as the reference technology. Alternative capture concepts may affect absolute results but are not expected to alter the relative comparison between the assessed fuel pathways.

Downstream of CO_2 capture and H_2 production, two alternative CO_2 conversion routes are examined: (i) direct hydrogenation of CO_2 to methanol, and (ii) conversion of CO_2 via the reverse water-gas shift reaction followed by FT-synthesis to diesel-range hydrocarbons. These two routes define the technological scope of the study and are assessed across all scenarios. Detailed process descriptions are provided in the following section.

Methodological approach

The methodological framework employed for the assessment of the economic and environmental performance of the evaluated fuel pathways is presented in the *Supplementary Material* (Figure S1), whereas the scope of the study is depicted in **Figure 1.** The integrated approach comprises three principal stages: (1) process design and simulation, (2) subsequent economic process evaluation within a techno-economic assessment (TEA), and (3) life cycle assessment (LCA). According to this framework, the investigated process routes are modelled in Aspen Plus® (V14) to generate mass and energy balances, which serve as the basis for plant design, equipment sizing, and the identification of potential energy savings through heat integration by applying a thorough pinch analysis. Equipment costs were estimated using the Aspen Process Economic Analyzer® (APEA). The resulting simulation data are subsequently used as input for the TEA, in which the process is evaluated based on the efficiency, stream data and economic indicators, while a custom-built Excel-based TEA tool integrates unit cost data with mass and energy balances to calculate capital (CAPEX), operating expenditures (OPEX), and net production costs (NPC). In parallel, the same data are transferred to the LCA to quantify environmental impacts (focusing on climate change) using the SimaPro software and ecoinvent database. This integrated approach is applied to both e-methanol (i) and FT-diesel (ii) pathways, enabling a transparent comparison with conventional diesel under identical system boundaries.

Scenario definition

Scenarios are defined as variations of external

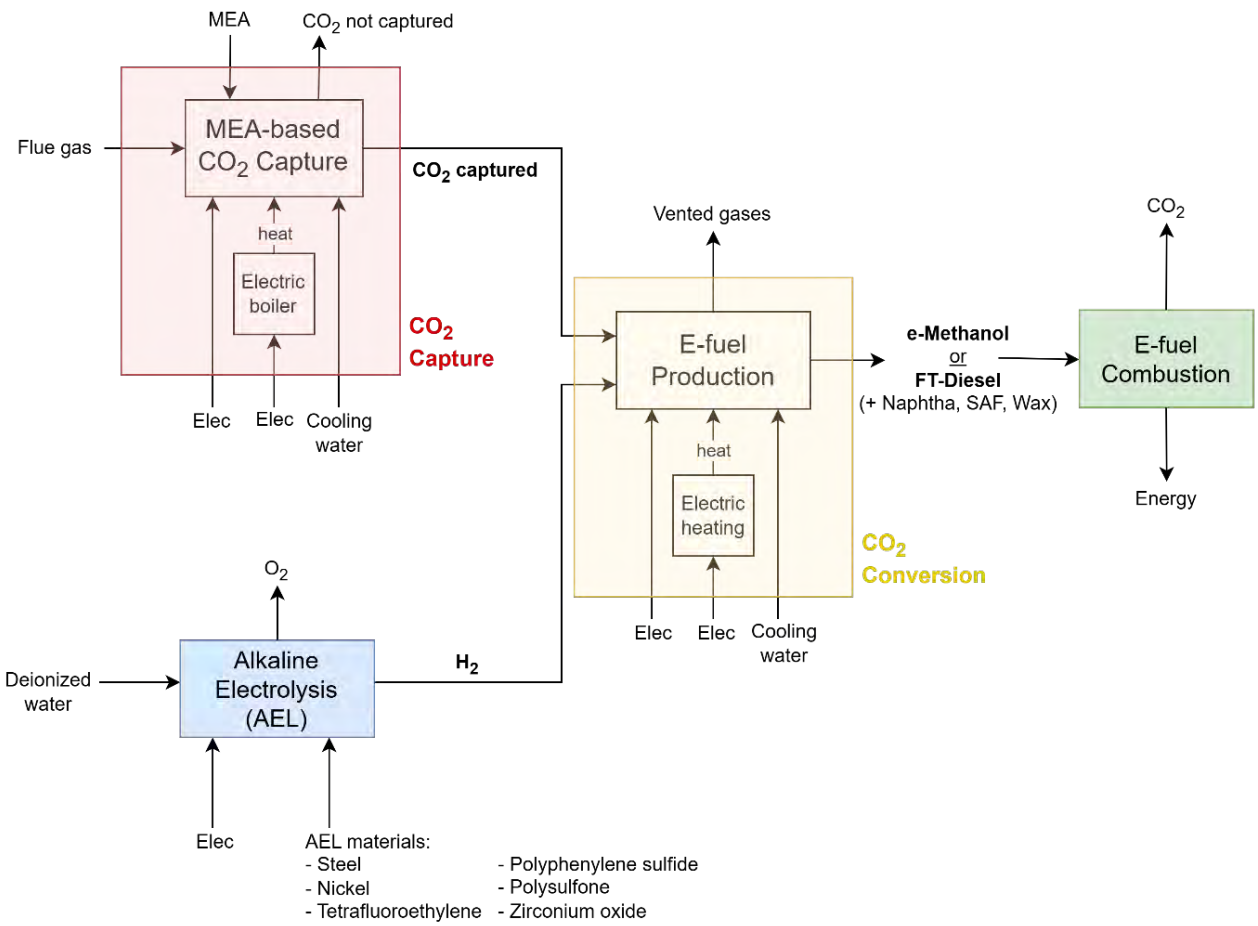

Figure 1. Scope of the study, consisting of four main blocks: (1) CO_2 capture from a cement plant; (2) H_2 production with alkaline electrolysis; (3) e-fuel production from CO_2 and H_2; and (4) e-fuel combustion.

boundary conditions applied to the reference system, while the plant configuration and process design remain unchanged. The comparison is carried out for Austria (base case) and other European countries, listed in the *Digital Supplementary Material* section S3, to assess the influence of the electricity grid mix on the climate change (CC) damage category (computed according to the GWP100 of the IPCC AR6 method). Grid-based electricity scenarios are evaluated to reflect current system conditions and constraints, rather than to imply compliance with additionality requirements under RFNBO regulation.

The scenarios differ with respect to (i) the carbon origin (fossil vs. biogenic CO_2) and (ii) the electricity supply for electrolysis and auxiliary processes. Hence, the following scenarios to mitigate CC of e-fuels in Austria are explored: (1) AT, fossil CO_2 from the cement plant is combined with H_2 produced using the electricity grid; (2) AT_BIO, biogenic CO_2 is combined with H_2 produced using the electricity grid; (3) AT_PV, fossil CO_2 is used as feedstock while the entire process (electrolysis, CO_2 capture, and CO_2 conversion) is powered by electricity from solar panels (PV); (4) AT_PV+BIO, biogenic CO_2 is used as feedstock for a process entirely powered by electricity from PVs; (5) AT_BIO+RENH2, biogenic CO_2 is used as feedstock and PV panels are powering the electrolyzer only, while remaining electricity demand is supplied by the grid; (6) AT_WIND, fossil CO_2 is used as feedstock while the entire process is powered by electricity from wind turbines.

Process simulation

The simulations are performed as steady-state flowsheet models, covering the whole process starting from CO_2 and H_2 to the final products. The methanol route was modeled as a three-stage fixed-bed reactor system with direct CO_2-to-methanol conversion (ratio of CO_2:H_2=1:3.08), following the configuration and operating conditions developed by [11] and [12]. Reaction kinetics for CO_2 hydrogenation were described using a Langmuir-Hinshelwood-Hougen-Watson (LHHW) model taken from [13]. Intermediate condensation between reactor stages was implemented to deduct H_2O and methanol, which enhances overall CO_2 conversion. Downstream of

the synthesis step, the water-methanol mixture is upgraded via a two-stage distillation, including one high- and one low-pressure distillation column.

The FT-diesel route consists of a high-temperature rWGS reactor operating at T = 950 °C and p = 11 bar(g) followed by low-temperature FT synthesis, adopting the process concept and product distribution models by [14]. Catalyst specific data for the rWGS reaction were derived from [15]. An elevated system pressure in the rWGS step was selected to avoid intermediate recompression prior to FTS. The FT reactor itself operates at a $CO:H_2$ ratio of 1:2.08, which is maintained by adjusting the H_2 feed. Unreacted gaseous components (CO_2, CO, H_2, CH_4, and light hydrocarbons) are recycled to the rWGS inlet, increasing carbon efficiency. Downstream of the FT reactor, a separation unit based on three distillation columns was implemented to obtain naphtha, kerosene, diesel, and wax fractions. Most of the wax is recycled in a hydrocracking unit to increase the diesel yield, while 10 wt% is withdrawn and sold as a by-product.

Techno-economic assessment

The techno-economic evaluation was based on the calculation of the NPC, expressed on both an energy (NPC_e in EUR/MWh) and mass basis (NPC_m in EUR/kg), which represents the product price required to achieve the economic break-even point. This approach is widely applied in the assessment of Power-to-X and CO_2 conversion processes and enables a consistent comparison of alternative process routes and plant configurations [14, 16]. Additionally, different energy pricing scenarios were considered to reflect different energy markets and price variability across Europe, which can be found in the Digital S*upplementary Material* (Table S1).

The NPC were determined by dividing the plant's total annualized cost by its annual product output. CAPEX values were derived from the major equipment costs obtained from process simulations, including reactors, heat exchangers, compressors, separation units, and utilities. These basic equipment costs were converted into total installed costs using standard installation and indirect cost factors according to the methodology originally proposed by Peters M. et al. [16].

Life cycle assessment

The impact assessment of e-methanol (i) and FT-diesel (ii) compared to fossil diesel is carried out according to a cradle-to-wheel approach using the Environmental Footprint (EF) 3.1 method [17] and the Ecoinvent 3.11 cut-off database [18] in Simapro v10.3. The combustion of 1 MJ of fuel, including emissions, is chosen as the functional unit (FU). The modeled production processes consist of four main blocks depicted in **Figure 1**. While the environmental impact assessment of H_2 production considers both energy consumption for operation and materials required for the electrolyzer, construction stages (i.e., facility and machinery) for MEA-based capture and e-fuel production units are excluded from the scope of the analysis due to lack of reliable data. Moreover, allocation methods are applied where co-products are produced: economic allocation for electrolysis (H_2/O_2) and energy-based allocation for the FTS to distribute the impact between FT-diesel and co-products (naphtha, kerosene, and wax). The modeling details can be found in the *Supplementary Materials* section S4.

Environmental impacts are evaluated using the climate change category (CC, kg CO_2-eq/FU), expressed relative to fossil diesel (=100%), and a single-score indicator (μPt/FU), representing the weighted aggregation of all impact categories (e.g., acidification, eutrophication, ecotoxicity, or water use).

RESULTS

The results of the process simulation, TEA, and LCA for the investigated e-methanol and FT-diesel pathways are presented in the following sections. Detailed results of the mass and energy balances, the calculation of the CAPEX and NPC, as well as the calculation of the climate change damage category from the LCA are provided in the *Supplementary Materials*.

Economic study

The following section summarizes the TEA results for the two investigated scenarios: e-methanol (i) and FT-diesel (ii). **Table 1** presents a breakdown of the main cost structure split into CAPEX and OPEX, and the resulting total and relative NPC values. Further details are available in the *supplementary material* (section S4).

From a techno-economic perspective, results indicate a clear differentiation between the two pathways. The FT-diesel route is characterized by substantially higher capital requirements. Total equipment costs amount to 1, 154 MEUR, compared to 343 MEUR for the e-methanol route. This reflects the additional synthesis, upgrading and separation units required along the FT pathway (pre-rWGS, hydrocracker, distillation). Investment in electrolysis is likewise higher for the FT route (717 MEUR vs. 631 MEUR), driven by increased demand for hydrogen. Although both pathways are based on an overall stoichiometric requirement of approximately 3 moles of H_2 per mole of CO_2, the FT diesel route exhibits electrolysis costs that are around 13% higher compared to scenario (i). This is mainly due to light hydrocarbons that are formed in the FT reactor and are recycled back to the rWGS unit, where they are partially converted into CO. This additional loop requires the addition of hydrogen for maintaining the targeted CO/H_2 ratio upstream of the FT reactor. This is further amplified by H_2 consumption in the downstream hydrocracker and slight losses in

the FT recycle loop. Consequently, the annualized capital cost (ACC) is almost twice as high for FT-diesel as for e-methanol (142 MEUR/a vs. 74 MEUR/a).

OPEX in both pathways are dominated by electrolysis-related costs. The annual replacement costs (ARC) of the electrolysis stack are 29 MEUR/a for the e-methanol route and 33 MEUR/a for the FT-diesel route, accounting for around 6.4% and 9.8% of total operating expenses (OPEX), respectively. Despite having a higher hydrogen demand and a higher share of electrolysis-related OPEX, the total OPEX of the FT-diesel route are lower than that of the e-methanol route (338 MEUR/a vs. 453 MEUR/a). This is primarily due to revenues from FT co-products (e.g. gasoline, naphtha, kerosene, wax and light gases), which are credited as sales and reduce the net OPEX of the FT pathway by around 7.7% compared to the single-product e-methanol route.

The interaction between higher CAPEX and lower OPEX results in slightly lower annual total net production costs (NPC_t) for the FT-diesel route than for e-methanol (513 vs. 556 MEUR/a). However, this apparent advantage is reversed when costs are normalized to the main product output. The NPC_m for FT-diesel is 7.16 EUR/kg, compared to 1.40 EUR/kg for e-methanol. The energy-specific NPC is 590 EUR/MWh for FT-diesel and 250 EUR/MWh for e-methanol.

Table 1: CAPEX and OPEX for e-methanol (i) and FT-diesel (ii) based on boundary conditions of the cement plant in Austria.

	Units	Process Route (i)	(ii)
CAPEX			
Total electrolysis costs	MEUR	631	717
Total equipment costs	MEUR	343	1, 154
Annualized capital cost [ACC]	MEUR/a	74	142
OPEX			
Annual replacement costs of electrolysis stack [ARC]	MEUR/a	29	33
Total OPEX	MEUR/a	453	338
NPC			
Total net-production costs [NPC_{total}]	MEUR/a	556	513
Mass specific [NPC_m]	EUR/kg	1.40	7.16
Energy specific [NPC_e]	EUR/MWh	250	590

To account for regional differences within Europe, country-specific electricity prices for 2025 were used to calculate the NPC_e, while all other techno-economic parameters and the process configuration remained constant. The resulting NPC spans a range of approximately 300-700 €/MWh for the e-methanol route and 700-1, 000 €/MWh for the FT-diesel route (**Figure 2** and **Figure 3**). Lower energy-specific production costs were observed in the Nordic countries (SE, FI, DK, EE, LV) as well as in south-western Europe (PT, ES).

Environmental study

The impacts in the CC damage category (kg CO_2 eq/FU), relative to fossil diesel (=100%), of e-fuels produced using electricity from different grid mixes (residual mixes with 2023 as reference year, except for Austria, market mix (2021) since residual mix is not available for this country) is depicted in **Figure 4**. Moreover, the overall environmental impact of e-fuels represented by the single score (µPt/FU), is depicted in **Figure 5**.

As expected, using non-renewable-sourced electricity to produce e-fuels results in a CC impact which is higher than the one of fossil diesel for the same energetic work. While CC of e-fuels are usually 100-300% higher than the fossil option, carbon-intensive mixes lead to even higher CC increases as can be observed for Germany (DE), Greece (GR), and Poland (PL). Indeed, the medium voltage electricity grid mix in Germany and Poland is dominated by hard coal (55% and 73%, respectively), while Greek electricity mix relies heavily on natural gas (52%). On the other hand, the lowest CC is achieved in France (FR) and Sweden (SE), respectively +34% and +35% compared to fossil diesel, attributable to two low-carbon grid mix profiles: Swedish electricity is heavily decarbonized through renewables (hydropower, wind turbines, and bioenergy, representing >50%) and nuclear energy (42%), whereas electricity is almost entirely produced from nuclear energy (88%) in France.

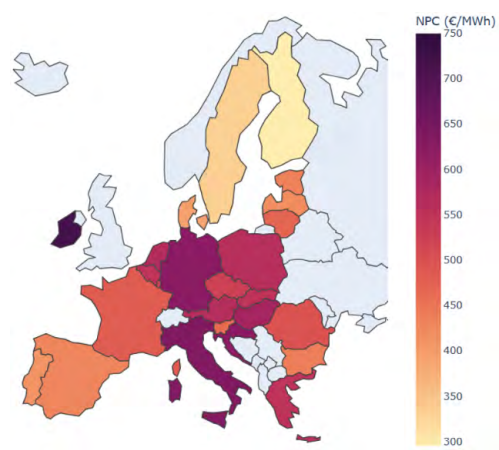

Figure 2. Influence on the NPC_e (€/MWh) for the e-methanol route (i) of regional electricity prices. The reference year for the economic assessment is 2025.

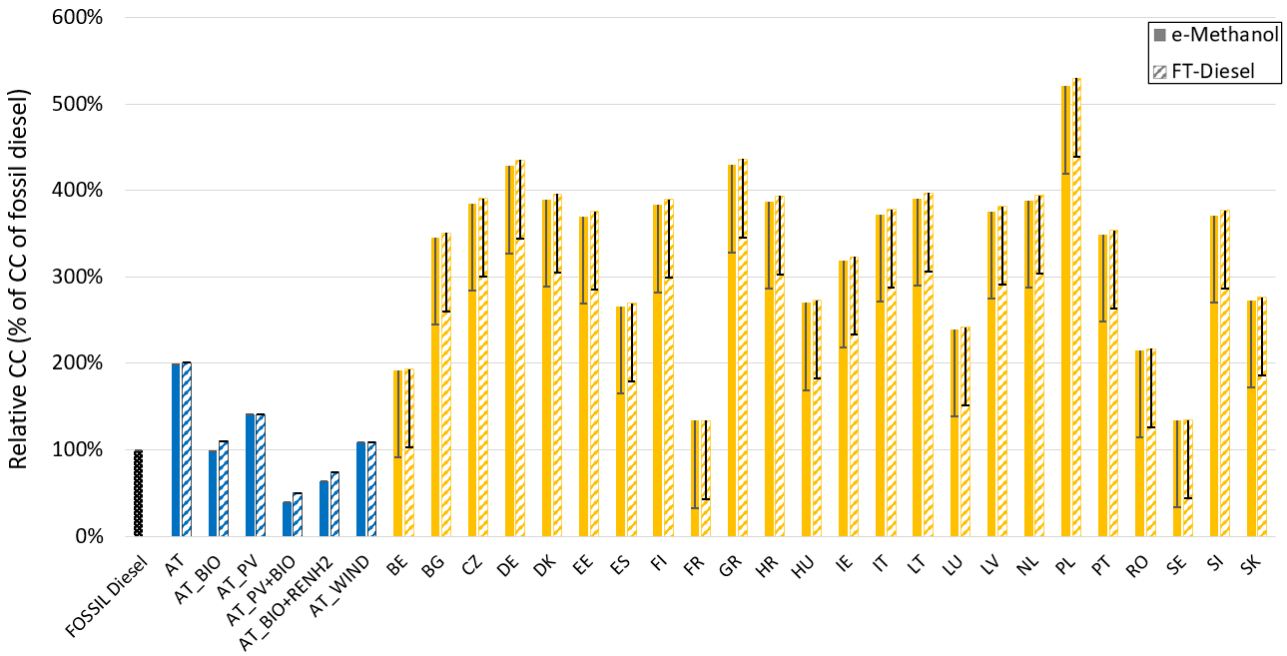

Figure 4. Damage assessment of CC for e-methanol (solid bars) and synthetic diesel (hatched bars) across Europe (fossil diesel = 100%), assessed using Environmental Footprint 3.1 method (v1.04).

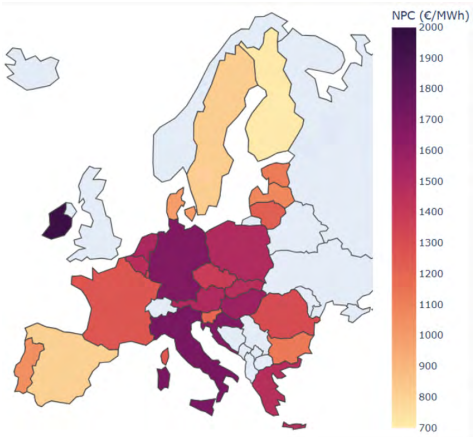

Figure 3. Influence on the NPC_e (€/MWh) for the FTS-diesel route (ii) of regional electricity prices. The reference year for the economic assessment is 2025.

When using the Austrian electricity grid mix, e-methanol combustion is associated with a CC of 0.178 kg CO_2 eq/FU, i.e., almost twice the one of fossil diesel. Nearly 40% of the CC is attributable to the combustion of the e-methanol, as CO_2 captured from the cement plant and used as feedstock is considered fossil (Figure S4).

If biogenic CO_2 is used instead, e.g., from a biogas production plant, the emissions during e-fuel combustion would not contribute to CC anymore since the characterization factor for biogenic carbon in EF 3.1. method is 0. Moreover, approximately 30% of the CC is linked to the AEL while electricity demand for CO_2 capture and e-methanol production further contribute 20%. These contributions could be reduced by using low-carbon electricity, e.g., PVs or wind turbines.

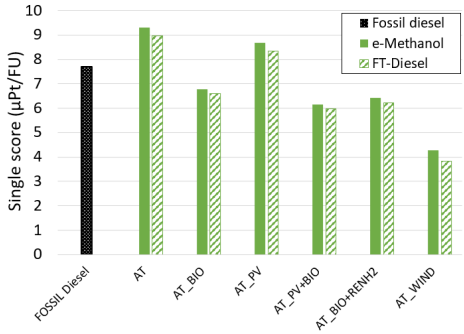

Figure 5. Single score of e-methanol (solid bars) and FT-diesel (hatched bars) for different Austrian scenarios, assessed using EF 3.1 method (v1.04).

The benefit of using biogenic CO_2 can be observed in **Figure 4**, represented by error bars for different European countries, excluding Austria where the AT_BIO scenario is used instead. Despite reducing the CC by 19-75% and 17-68% for, respectively, e-methanol and FT-diesel compared to initial estimations, switching towards biogenic carbon-rich feedstock alone is not sufficient in most cases to achieve a lower CC than fossil diesel. Only the French and Swedish electricity mixes enable net reductions of the CC for both e-fuels (56-67%) compared to the fossil alternative.

Utilizing PV-based electricity (AT_PV) decreases CC

by 29-30% compared to AT, but remains 41% more impacting than fossil diesel combustion. Moreover, as can be seen on **Figure 5**, AT and AT_PV scenarios are also associated with the highest single scores, indicating the worst environmental impact among all Austrian options. On the contrary, combining biogenic CO_2 and solar energy reduces CC by 25-36% when solar energy powers H_2 production only (AT_BIO+RENH2) and by 50-60%, compared to fossil diesel, when PVs cover process electricity demand entirely (AT_PV+BIO). Despite clear reduction in the CC damage category, the overall environmental burden mitigation compared to the fossil alternative is limited (17-23% lower single score) due to materials requirements for PVs.

Finally, the CC of e-fuels in AT_WIND scenario is only 9% higher than the one of fossil diesel. Such a small gap could be reduced by improving the energy efficiency of electrolysis or CO_2 capture and conversion processes, reducing the flow rate of vented flue gas, or increasing the capture efficiency to above 90%. However, AT_WIND scenario presents the lowest environmental impact (44-50% reduction compared to fossil diesel), suggesting this is the most sustainable Austrian configuration.

CONCLUSIONS

The limited electrification potential of long-distance heavy-duty transport underlines the need for low-carbon fuels with high energy density. While the utilization of captured CO_2 from cement plants is often discussed as a promising mitigation pathway, this study demonstrates that cement-based PtL concepts are fundamentally constrained by high energy demand, limited renewable electricity availability, large electrolyzer scale requirements, and regulatory uncertainty, particularly under RFNBO requirements restricting fossil CO_2 utilization beyond 2040. While both e-fuel pathways are technically feasible at process level, their system-level feasibility is constrained by low-carbon electricity availability, grid integration, and electrolyzer scale requirements, reaching hundreds of megawatts in the investigated configurations. Such capacities exceed current installation levels and raise practical challenges related to grid integration, investment risk, and operational flexibility [7, 12].

The techno-economic assessment of e-methanol and FT-diesel production from captured CO_2 at an Austrian cement plant highlights a clear differentiation between the two e-fuels, while both remain significantly more expensive than their fossil counterparts. For the Austrian base case, the NPC_e (dominated by electrolysis-related operating costs) amount to 565€/MWh for e-methanol and 1, 498€/MWh for FT-diesel, compared to approximately 81€/MWh for fossil methanol [19] and 155€/MWh for fossil diesel [20], with similar cost relationships observed across other European countries. Under current European market conditions, the production of e-fuels remains economically uncompetitive, even before considering additional limitations such as renewable electricity availability and infrastructure requirements, highlighting the continued need for funding mechanisms and process improvements to close the cost gap.[19, 20]

From an environmental perspective, the LCA highlights the dominant influence of the electricity supply on environmental impacts. When e-fuels are produced using residual or market-based grid electricity mixes, the CC impacts exceed that of fossil diesel in most investigated regions, with CC values 100% to more than 400% higher than fossil diesel for carbon-intensive grid mixes, consistent with previous findings on electricity mix sensitivity in PtL pathways [8, 21, 22]. Substitution from fossil to biogenic CO_2 feedstock alone is insufficient to achieve lower CC than fossil diesel in most cases. This is reflected in Austrian scenarios where using biogenic CO_2 with grid electricity improves environmental outcomes but remains above fossil diesel. However, the lower energy penalty associated with CO_2 capture from biogas streams, due to higher CO_2 concentration than in cement flue gases, suggests overestimated CC impacts of biogenic-CO_2-based e-fuels when uniform capture assumptions are applied [23]. Meaningful greenhouse gas emission reductions are achievable only when biogenic CO_2 is coupled with low-carbon electricity systems (e.g., AT_PV+BIO reducing CC impacts by 60% and 50% for e-methanol and FT-diesel, respectively, compared to fossil diesel). Consequently, the environmental viability of these process routes is strongly dependent on the availability and large-scale deployment of renewable electricity. However, other environmental impact categories should be considered to avoid burden displacement and make a sound decision about renewables deployment. Moreover, the present assessment evaluates absolute environmental impacts relative to fossil diesel and does not explicitly account for substitution effects. Although e-fuels are intended to replace fossil fuels and enable a secondary use of industrial CO_2, such system-level effects are not fully reflected in the current LCA. Thus, a comprehensive evaluation of substitution effects and their implications for fossil fuel replacement is therefore identified as an important subject for future studies.

Overall, the findings highlight that PtL assessments are highly sensitive to boundary-condition assumptions. Optimistic conclusions regarding climate benefits of e-fuels often rely on implicit assumptions of abundant low-cost, low-carbon electricity and large-scale electrolysis deployment, which are not yet realized within the European energy system. Consequently, meaningful evaluation of industrial CO_2-based e-fuel pathways requires realistic, system-consistent boundary conditions to ensure robust and policy-relevant conclusions [23, 24]. Results indicate that cement-plant-based PtL concepts offer

limited near- to mid-term mitigation potential under current European energy system conditions. Future PtL concepts should therefore prioritize biogenic CO_2 sources, efficiency improvements, and access to additional low-carbon electricity, while ensuring that synthetic fuels are deployed where they complement direct electrification and deliver the highest climate benefit. The fundamentally different characteristics of FT-diesel, a certified drop-in fuel for existing engines and infrastructure, and e-methanol, which is still undergoing technical and regulatory validation, highlight the necessity of technology-specific and system-level evaluation when defining their roles in a future low-carbon transport system. In this context, the study emphasizes that e-fuel pathways based on industrial CO_2 utilization can only be meaningfully assessed when grounded in realistic assumptions regarding renewable electricity availability and electrolyzer scale-up.

DIGITAL SUPPLEMENTARY MATERIAL

Abbreviations and modeling details can be accessed here: LAPSE:2026.0026

ACKNOWLEDGEMENTS

This work was supported by the Belgian Fund for Scientific Research (FRS-FNRS) through a FRIA grant. The authors also acknowledge Angélique Léonard and Sylvie Groslambert (University of Liège) for their valuable feedback on the life cycle assessment, as well as Markus Lehner (Technical University of Leoben) for critically reviewing the manuscript.

AUTHOR IDENTIFIERS

Andritz M.: 0009-0007-2150-6691
Sendlhofer S.: 0009-0009-3411-9070
Mitraki R.: 0009-0007-0990-0514
Léonard G.: 0000-0003-2237-8306
Markowitsch C.: 0000-0002-2016-9765

REFERENCES

1. Statista, https://www.statista.com/statistics/1171183/ghg-emissions-sector-european-union-eu/
2. European Environment Agency, https://www.eea.europa.eu/en/analysis/indicators/greenhouse-gas-emissions-from-transport
3. Jahangir Samet M, Liimatainen H, van Vliet OPR, Pöllänen M. Road freight transport electrification potential by using battery electric trucks in finland and switzerland. Energies 14:823 (2021). https://doi.org/10.3390/en14040823
4. Li C, Zhang J, Gu X, Yang J, Qin Q, Zhang W, Chandio AA. Assessing alternative fuels for heavy-duty trucks in the U.S. based on improved life cycle assessment. Sustainable Production and Consumption 58:1-13 (2025). https://doi.org/10.1016/j.spc.2025.06.003
5. Teixeira ACR, Machado PG, Collaço FMA, Mouette D. Alternative fuel technologies emissions for road heavy-duty trucks: a review. Environ Sci Pollut Res 28:20954-20969 (2021). https://doi.org/10.1007/s11356-021-13219-8
6. European Committee for Standardization. EN15940: Automotive fuels – Paraffinic diesel fuel from synthesis or hydrotreatment: Requirements and test methods (2016)
7. Marchese M, Giglio E, Santarelli M, Lanzini A. Energy performance of power-to-liquid applications integrating biogas upgrading, reverse water gas shift, solid oxide electrolysis and fischer-tropsch technologies. Energy Conversion and Management: X 6:100041 (2020). https://doi.org/10.1016/j.ecmx.2020.100041
8. Gao R, Zhang C, Jun KW, Kim SK, Park HG, Zhao T, Wang L, Wan H, Guan G. Transformation of CO2 into liquid fuels and synthetic natural gas using green hydrogen: a comparative analysis. Fuel 291:120111 (2021). https://doi.org/10.1016/j.fuel.2020.120111
9. Eyberg V, Dieterich V, Bastek S, Dossow M, Spliethoff H, Fendt S. Techno-economic assessment and comparison of fischer–tropsch and methanol-to-jet processes to produce sustainable aviation fuel via power-to-liquid. Energy Conversion and Management 315:118728 (2024). https://doi.org/10.1016/j.enconman.2024.118728
10. International Energy Agency, Global Hydrogen Review (2025)
11. d'Hubert X, Riley J. Main electrical requirements and their implementations to achieve the decarbonation of the cement plants. 2023 IEEE-IAS/PCA Cement Industry Conference (IAS/PCA) :1-11 (2023). https://doi.org/10.1109/ias/pca56501.2023.11303628
12. European Commission. https://energy.ec.europa.eu/topics/renewable-energy/renewable-energy-directive-targets-and-rules/renewable-energy-directive_en
13. C. Markowitsch. C2PAT - Carbon to Product Austria (2022)
14. Markowitsch C, Lehner M, Maly M. Comparison and techno-economic evaluation of process routes for lower olefin production via fischer–tropsch and

methanol synthesis. International Journal of Greenhouse Gas Control 129:103985 (2023). https://doi.org/10.1016/j.ijggc.2023.103985
15. Markowitsch C, Andritz M, Lindenthal L, Cotter T, Drexler H, Rameshan C, Lehner M. Process intensification of the rwgs reaction by a perovskite-based catalyst. Chemical Engineering Journal 500:156577 (2024). https://doi.org/10.1016/j.cej.2024.156577
16. Max S. Peters, et al. Plant Design and Economics for Chemical Engineers. McGraw-Hill (2003)
17. S. Andreasi Bassi, et al. Updated characterisation and normalisation factors for the Environmental Footprint 3.1 method. Publications Office of the European Union (2023) https://doi.org/10.2760/798894
18. Wernet G, Bauer C, Steubing B, Reinhard J, Moreno-Ruiz E, Weidema B. The ecoinvent database version 3 (part I): overview and methodology. Int J Life Cycle Assess 21:1218-1230 (2016). https://doi.org/10.1007/s11367-016-1087-8
19. Trading economics. https://tradingeconomics.com/commodity/methanol
20. European Commission. https://energy.ec.europa.eu/data-and-analysis/weekly-oil-bulletin_en
21. F. Bozzolo Lueckel, et al. Status and perspectives of non-biogenic renewable gases: synthesis report of work package 2 of the IEA Bioenergy intertask project (2022)
22. Gao R, Zhang C, Jun KW, Kim SK, Park HG, Zhao T, Wang L, Wan H, Guan G. Green liquid fuel and synthetic natural gas production via CO2 hydrogenation combined with reverse water-gas-shift and co-based fischer-tropsch synthesis. Journal of CO2 Utilization 51:101619 (2021). https://doi.org/10.1016/j.jcou.2021.101619
23. Pratschner S, Hammerschmid M, Müller S, Winter F. Evaluation of CO2 sources for power-to-liquid plants producing fischer-tropsch products. Journal of CO2 Utilization 72:102508 (2023). https://doi.org/10.1016/j.jcou.2023.102508
24. Colelli L, Dell'Aversano S, Bassano C, Vanga G, Gallucci K, Vilardi G. Liquid e-fuels for a sustainable future: a comprehensive review of production, regulation, and technological innovation. Energy Conversion and Management 347:120529 (2026). https://doi.org/10.1016/j.enconman.2025.120529

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Integration of computer aided design and emerging technology development based on a series of scale-up demonstration tests; Case study of thermal energy storage

Shoma Fujii[a], Yasunori Kikuchi[a]*

[a] The University of Tokyo, Institute of Future Initiatives, 7-3-1 Hongo, Bunkyo-ku, Tokyo, Japan
[b] The University of Tokyo Department of Chemical System Engineering, 7-3-1 Hongo, Bunkyo-ku, Tokyo, Japan
* Corresponding Author: shoma.fujii@ifi.u-tokyo.ac.jp

ABSTRACT

Early-stage system-level assessment of emerging technologies is essential for achieving climate neutrality and a circular economy; however, such assessments are often constrained by the lack of representative life cycle inventory data. In thermal energy systems, performance strongly depends on scale, making direct application of laboratory- or bench-scale experimental data potentially misleading in life cycle assessment (LCA). This study investigates the influence of experimental scale on system-level evaluation using a zeolite-based thermal energy storage (TES) system as a case study.LCAs were conducted using performance data from laboratory-, bench-, and pilot-scale experiments and compared with predicted commercial-scale performance derived from numerical simulations. The TES system stores waste heat via water vapor desorption from zeolite and generates pressurized steam using a moving-bed with indirect heat exchanging system. Heat recovery ratios of 36%, 50%, and 61% were obtained at laboratory, bench, and pilot scales, respectively, while commercial-scale performance was predicted to reach approximately 80% due to reduced heat losses. Accordingly, LCAs based on small-scale data indicated higher greenhouse gas (GHG) emissions than conventional technologies, whereas assessments using commercial-scale data demonstrated clear GHG reduction potential.The results highlight the risk of erroneous conclusions when small-scale data are directly used for system evaluation. By integrating experimentally validated numerical modeling with computer-aided process engineering, scale-dependent data can be transformed into system-representative inputs, enabling fair and consistent assessment of emerging technologies.

Keywords: Adsorption, Energy Systems, Modelling and Simulations, Life Cycle Analysis, Technoeconomic Analysis

INTRODUCTION

Achieving global visions such as climate neutrality and a circular economy requires the early adoption of emerging technologies and their system-level assessment using tools such as life cycle assessment (LCA) and techno-economic analysis (TEA) to quantify environmental and economic performance. In particular, LCA has been increasingly applied as a decision-support tool in the emerging technology development to guide research priorities [1, 2]. However, conducting LCA for emerging technologies is inherently challenging, as detailed life cycle inventory (LCI) data are often unavailable in existing databases, which predominantly represent mature, commercial-scale processes. For chemical and thermal processes, missing inventory data are frequently generated using process engineering approaches, including mass and energy conservation calculations, numerical modeling, and computer-aided process engineering (CAPE) tools [3, 4]. While such approaches enable the estimation of prospective system performance, a critical methodological challenge remains: determining which scale of

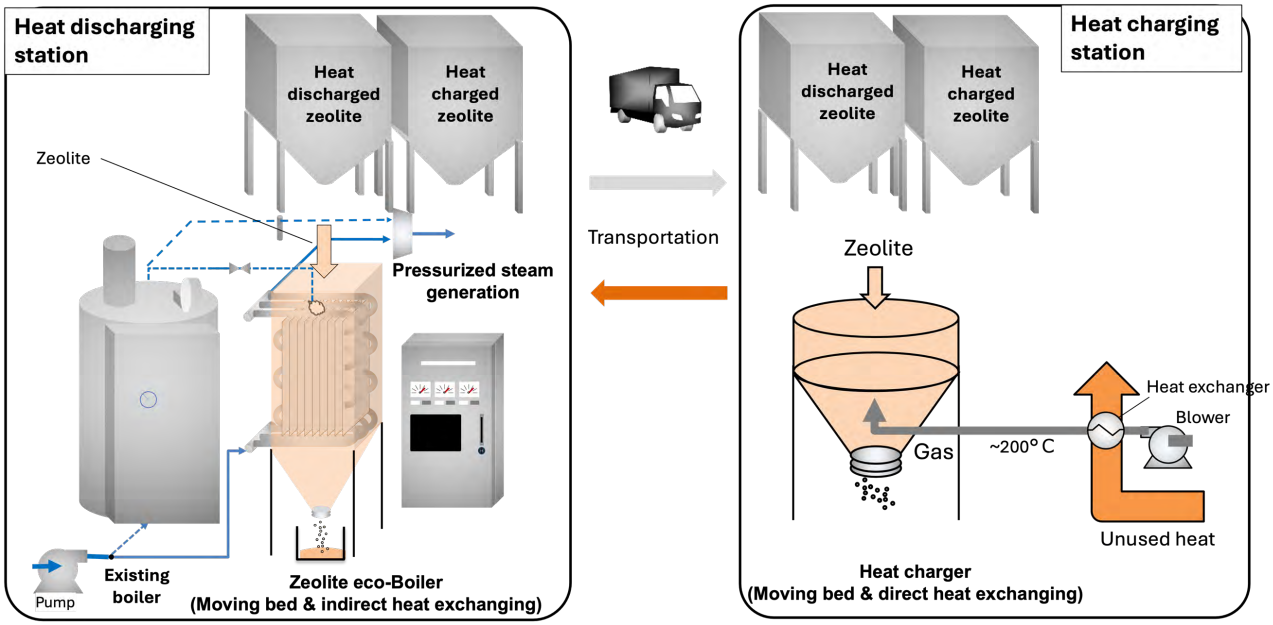

Figure 1. Schematic of thermal energy storage and transport system

experimental or modeled data—ranging from laboratory and bench scale to pilot and commercial scale—should be incorporated into system-level assessments. This choice is nontrivial, as key performance indicators such as efficiency, heat loss, and conversion yield often exhibit strong scale dependency [5].

For thermal energy systems in particular, efficiency generally improves with increasing scale due to reduced relative heat leakage and more favorable surface-to-volume ratios. As a result, system evaluations based only on laboratory- or bench-scale experimental data tend to underestimate achievable performance at commercial scale, potentially leading to overly conservative or even misleading conclusions regarding environmental benefits. In contrast, assessments based on full-scale demonstration or operational data provide higher accuracy but are typically time-consuming, capital-intensive, and impractical during the early research and development phases when key design decisions must be made.

To address this methodological gap, the present study employs a thermal energy storage system as a case study to systematically investigate how the scale of experimental data influences system-level evaluation outcomes. By explicitly examining the role of scale in LCA-relevant performance parameters and integrating experimental results with process-based modeling, this work aims to clarify the conditions under which small-scale data can be reliably transformed into system-representative inputs for fair and consistent decision-making in the development of emerging thermal energy technologies.

MATERIAL AND METHODOLOGY

LCAs using data from laboratory-, bench-, and pilot-scale demonstration tests were conducted and compared with predicted commercial-scale performance derived from numerical simulations in a case study of thermal energy storage technology. The objective was to quantify the uncertainty arising from scale differences and to explore methods to mitigate such uncertainty in the evaluation of emerging technologies.

Thermal energy storage technology enables the recovery and regeneration of waste heat across temporal and spatial gaps through the use of thermal energy storage materials[6]. Fig. 1 shows the schematic of thermal energy storage system using zeolite water vapor adsorption and desorption process[7]. Thermal energy is stored by desorbing water previously adsorbed in zeolite using waste heat. The water-desorbed zeolite is packaged and stored in sealed containers to prevent adsorption of ambient humid and is transported to the heat demand site when required. On the demand side, a heat discharged equipment named "zeolite eco-boiler" capable of continuous generation of pressurized steam is being developed by adopting a moving-bed reactor combined with an

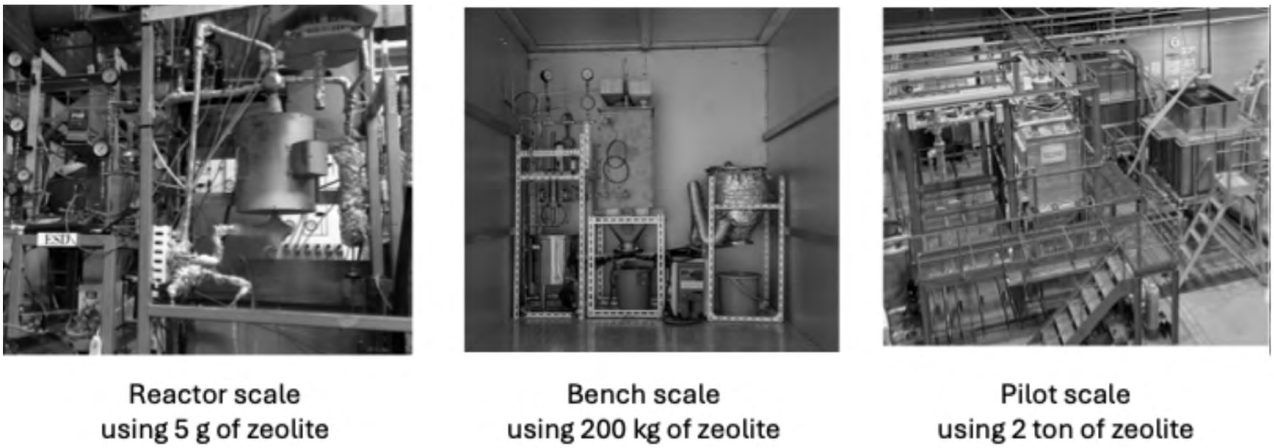

Figure 2. Demonstration tests at various scales for zeolite-based thermal energy storage technology

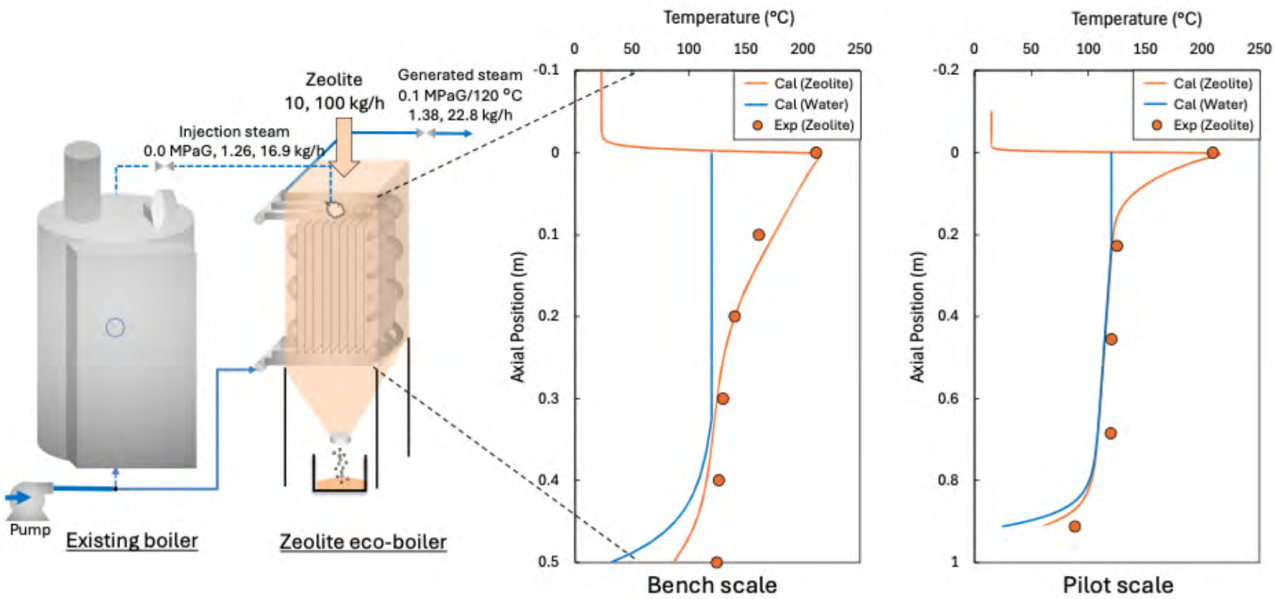

Figure 3. Experimental and numerical analysis results for bench- and pilot- scale zeolite eco-boiler demonstration test

indirect heat exchange system[8].

In the proposed system, dried zeolite particles are continuously supplied from the top of the chamber and discharged at a constant flow rate from the bottom, thereby forming a moving bed. Cooperative operation with an existing boiler is assumed. Steam generated from the existing boiler is injected into the chamber, where it is adsorbed by the zeolite, resulting in the release of adsorption heat. A plate-finned-tube heat exchanger is integrated within the chamber, and the released adsorption heat is transferred to feedwater flowing inside the heat exchanger to generate pressurized steam. As a result, the amount of steam generated exceeds the injected steam by approximately 20–30%, enabling a reduction in the fuel consumption of the existing boiler.

A series of experiments for zeolite eco-boiler across multiple scales: laboratory-scale tests using around 5 grams of storage material, bench-scale tests using around 200 kg[8, 9], and pilot-scale tests using around 2 tons[10] were conducted as shown in Fig .2.) Inventory data for LCA were generated from these experiments, as well as from performance data predicted for commercial-scale systems using numerical modeling that combines energy conservation (for zeolite, gas, and water), mass conservation (for zeolite, gas, and steam), and momentum conservation[11].

The life cycle boundary was kept consistent across all scales to ensure comparability of results. Additionally, the same numerical modeling framework was applied to predict large-scale system performance without any

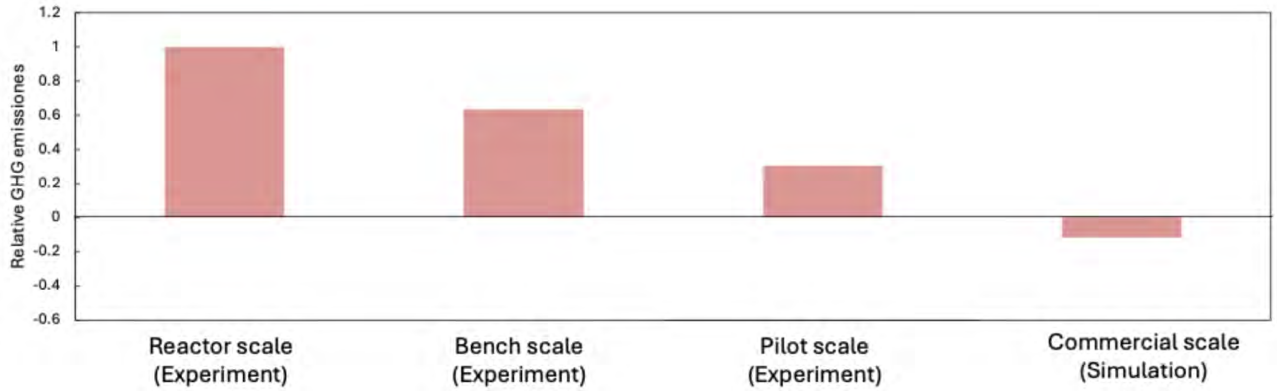

Figure 4. Results of lifecycle GHG emissions

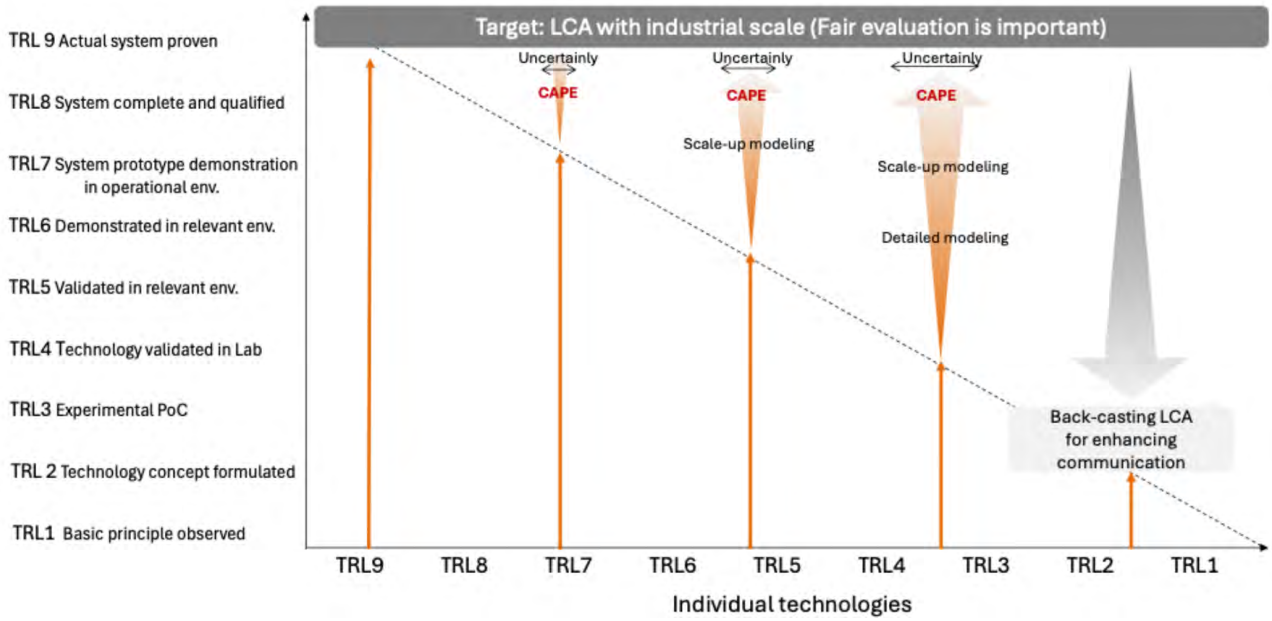

Figure 5. Relationship between experimental works, CAPE, and system-level assessment

fitting, by separately modeling heat and mass transfer phenomena. The life cycle boundary was defined to cover the entire system from equipment manufacturing to end-of-life disposal, and the operational lifetime of the equipment was assumed to be 15 years. The size of each component was set to that of a commercially scaled system designed through numerical simulations. Among the performance parameters, only the heat recovery rate—defined as the recovery rate of generated steam relative to the heat storage capacity and relevant to fuel reduction rate–was derived from experimental data obtained at different scales. The functional unit of LCA was defined as the annual steam supply delivered through boiler operation.

RESULT AND DISCUSSION

Fig. 3 shows the axial temperature distribution during steady-state operation for the bench scale and pilot scale of zeolite-eco Boiler, alongside the results predicted by numerical analysis. In both cases, continuous superheated steam generation and fuel reduction effects were demonstrated. Furthermore, it can be seen that the numerical analysis model incorporating heat and mass transfer simulates the experimental results for both scales.The results clearly demonstrated the scale dependency of TES system performance. The thermal recovery ratio was found to be 36% at the laboratory scale, 50% at the bench scale and 61% at the pilot scale. In contrast, for commercial-scale systems, the predicted thermal recovery ratio reached approximately 80%, indicating a substantial improvement with increasing scale due to reduced heat leakgae. These findings underscore the substantial discrepancy between small-scale exper–

imental data and real-world large-scale performance. Importantly, the numerical model developed in this study, which separately models heat and mass transfer phenomena without introducing scale-up correction factors, successfully reproduced the performance trends observed in scaled-up systems, thereby confirming the validity of the modeling approach.

The implications of these results were further examined through life cycle assessment. Fig. 4 shows the results of relative lifecycle GHG emissions compared to existing boilers. LCAs conducted using laboratory-, bench- and pilot- scale data indicated higher greenhouse gas (GHG) emissions compared to conventional technologies, primarily due to the lower heat recovery ratio at smaller scales. However, when performance data from commercial-scale systems were used, the results showed a clear reduction in GHG emissions, demonstrating the potential environmental benefits of thermal energy storage technology. This contrast reveals a critical methodological issue: using laboratory-, bench-, and pilot- scale data directly for system evaluation of developing technologies can lead to incorrect conclusions and misguided policy or investment decisions.

Furthermore, this study highlights the potential of combining numerical modeling and process engineering methods to bridge scale gaps in system evaluation. When only small-scale experimental data are available, validated numerical models that explicitly account for heat and mass transport phenomena can be used to extrapolate reliable large-scale performance estimates as shown in Fig. 3, thereby reducing uncertainty in environmental and techno-economic assessments. Such hybrid approaches enable more accurate and timely evaluations of emerging technologies without the need for time-intensive full-scale demonstrations.

Fig. 5 illustrates the relationship between experimental works, computer-aided process engineering (CAPE), and system-level assessment as a function of technology maturity. The vertical axis indicates the level of data used for system assessment, while the horizontal axis represents the technology readiness level (TRL) of the evaluated technology. For technologies with high TRLs (TRL 8–9), operational data from real-world systems are already available and major uncertainties have been eliminated; therefore, additional support from CAPE is generally unnecessary. In contrast, for technologies with intermediate TRLs (approximately TRL 3–7), direct use of experimental data to system assessment, as demonstrated in the thermal energy storage case study, may lead to erroneous decision-making, including reversals in the sign of greenhouse gas (GHG) impacts relative to existing technologies. Accordingly, CAPE that integrates modeling and process simulation is required to transform experimental data into system-representative data at a level suitable for fair and consistent system assessment. For technologies at very low TRLs (TRL 1–2), however, the uncertainty associated with scale-up predictions using CAPE becomes excessively large. In such cases, CAPE should instead be used to support goal setting by applying backcasting approaches, such as break-even analysis, to delineate the feasible performance range of the technology.

CONCLUSION

This study provides a framework that integrates experimentation, numerical modelling, process engineering, and life cycle assessment, thereby enhancing the reliability of system evaluations for emerging technologies. By demonstrating the impact of scale on both performance and environmental effects, it emphasises the importance of scale-aware methodologies in achieving climate neutrality and circular economy objectives. The approach presented here is extendable beyond thermal energy storage to diverse emerging technologies, providing a practical pathway to accelerate sustainable technology development while ensuring robust, information-based decision-making in the early stages of emerging technology development.

ACKNOWLEDGEMENTS

This work was supported by JSPS KAKENHI Grant-in-Aid for Early-Career Scientists (Grant Number 25K21417), JST COI-NEXT (Grant Number JPMJPF2003), JST PRESTO (Grant Number JPMJPR2278) and the Environment Research and Technology Development Fund (Grant Number JPMEERF20243G01) of the Environmental Restoration and Conservation Agency of Japan. The activities of the Presidential Endowed Chair for “Platinum Society” at the University of Tokyo are supported by Mitsui Fudosan Corporation, Sekisui House, Ltd., East Japan Railway Company, and Toyota Tsusho Corporation

AUTHOR IDENTIFIERS

Author ORCIDs:
Fujii S: 0000-0001-6440-3819
Kikuchi Y: 0000-0003-0002-8233

REFERENCES

1. Cucurachi S, van der Giesen C, Guinée J. Ex-ante LCA of emerging technologies. Procedia CIRP 69:463-468 (2018). https://doi.org/10.1016/j.procir.2017.11.005
2. Villares M, I??ldar A, van der Giesen C, Guinée J. Does ex ante application enhance the usefulness of LCA? a case study on an emerging technology for metal recovery from e-waste. Int J Life Cycle

Assess 22:1618-1633 (2017). https://doi.org/10.1007/s11367-017-1270-6
3. Piccinno F, Hischier R, Seeger S, Som C. From laboratory to industrial scale: a scale-up framework for chemical processes in life cycle assessment studies. Journal of Cleaner Production 135:1085-1097 (2016). https://doi.org/10.1016/j.jclepro.2016.06.164
4. Caduff M, Huijbregts MAJ, Koehler A, Althaus H, Hellweg S. Scaling relationships in life cycle assessment. J of Industrial Ecology 18:393-406 (2014). https://doi.org/10.1111/jiec.12122
5. Gavankar S, Suh S, Keller AA. The role of scale and technology maturity in life cycle assessment of emerging technologies: a case study on carbon nanotubes. J of Industrial Ecology 19:51-60 (2014). https://doi.org/10.1111/jiec.12175
6. Krönauer A, Lävemann E, Brückner S, Hauer A. Mobile sorption heat storage in industrial waste heat recovery. Energy Procedia 73:272-280 (2015). https://doi.org/10.1016/j.egypro.2015.07.688
7. Fujii S, Kanematsu Y, Kikuchi Y, Nakagaki T. Effect of bagasse drying on thermal energy storage utilizing zeolite water vapor ad/desorption at a sugar mill. Journal of Energy Storage 51:104495 (2022). https://doi.org/10.1016/j.est.2022.104495
8. Fujii S, Horie N, Nakaibayashi K, Kanematsu Y, Kikuchi Y, Nakagaki T. Design of zeolite boiler in thermochemical energy storage and transport system utilizing unused heat from sugar mill. Applied Energy 238:561-571 (2019). https://doi.org/10.1016/j.apenergy.2019.01.104
9. Fujii, S., Saito, S., Matsui, K., Nakagaki, T., Kikuchi, Y., Small-scale field demonstration of zeolite based mobile thermochemical energy storage, *Proc. the 16th IEA ES TCP Int'l Conf. Energy Storage*, 463-466 (2024)
10. Shibusawa, A., Baba, D., Kamata, H., Tanino, M., Aoki, H., Matsui, K., Nakagaki, T., Fujii, S., Kawakami, Y., Exploratory Design Study on Zeolite eco-Boiler through Scale-up Test Planning of Adsorption Thermal Storage system, Proc. the 29th Nat'l Symp. Power and Energy Systems, D133 (2025)
11. Fujii S, Nakagaki T, Kanematsu Y, Kikuchi Y. Prospective life cycle assessment for designing mobile thermal energy storage system utilizing zeolite. Journal of Cleaner Production 365:132592 (2022). https://doi.org/10.1016/j.jclepro.2022.132592

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimizing Heat Storage Integration for Solar Thermal Systems in Industrial Process Heat Networks

Håvard Falch[a], Henri Tande Espen[a] and Rahul Anantharaman[a*]

[a] SINTEF Energy Research, Trondheim, Norway
* Corresponding Author: Rahul.Anantharaman@sintef.no

ABSTRACT

European industry accounts for approximately 20% of total European CO_2 emissions, with heat demand representing one of the largest energy consumers. Solar thermal collectors offer an efficient renewable alternative to fossil fuels to cover the heat demand. However, due to the temporal mismatch between the solar thermal generation and process heat generation a thermal storage is needed to maximize the renewable utilization. This article presents a novel optimization framework for integrating an ideal heat storage with solar thermal systems in multiperiod heat exchanger network synthesis. We derive an analytical approach to optimize the heat storage by using physical insights from Pinch Analysis: heat can only charge the storage below the lowest pinch point in a given period and discharge above the highest pinch point. We show both how to do it for a storage of infinite size and of finite size, and that the infinite size storage is much more efficient to solve. The approach is validated using real industrial data from a lubricant plant and optimizing the implementation of solar thermal at the plant. By increasing the size of the storage more of the solar thermal heat is utilized, however, the increase in utilization becomes smaller as the storage becomes larger. While 100% solar thermal heat utilisation requires a storage size of 1040 kWh, 90% utilisation can be reached with a storage of 250 kWh, compared to only 70% without a storage.

Keywords: Heat Exchanger Network, Energy Storage, Renewable and Sustainable Energy, Solar Thermal

INTRODUCTION

European industry accounts for approximately 20% of total European CO_2 emissions, with heat demand representing one of the largest energy consumers. Currently, this demand is predominantly met through fossil fuel combustion [1]. Decarbonising industrial heat production is therefore critical to achieving European climate targets, and increasing the penetration of renewable heat sources offers a promising pathway forward.

Solar thermal (ST) collectors can efficiently convert solar radiation into low- to medium-temperature heat suitable for many industrial applications [2]. However, a fundamental challenge arises from the temporal mismatch between the availability of solar thermal energy and industrial heat demand. Solar radiation follows predictable daily and seasonal patterns, whereas industrial processes often exhibit different temporal profiles or require heat during periods of low solar availability. This mismatch can significantly limit the fraction of renewable heat that can be directly utilised, potentially undermining the economic viability of solar thermal installations.

Beyond solar thermal integration, industrial processes themselves exhibit fluctuations in heat demand and in the availability of hot process streams over time. These variations create opportunities to capture and store surplus process heat during periods of excess and to recover it during periods of heat deficit. Thermal energy storage systems can address both challenges simultaneously: storing excess solar thermal energy when production exceeds demand and capturing surplus process heat for later use when additional heat is needed.

Efficiently integrating heat sources, storage systems, and process demands in an industrial site requires sophisticated process integration methods. Heat exchanger network synthesis (HENS) provides a systematic framework for designing optimal heat exchanger layouts that minimise utility consumption while meeting all

process heating and cooling requirements. When dealing with time-varying conditions—whether from fluctuating solar production, variable process schedules, or both—multiperiod HENS becomes necessary to capture the full complexity of the system and optimise across all operational scenarios.

Previous work has explored the integration of thermal energy storage into HENS problems. Walmsley et al. [3] investigated the integration of industrial solar and gaseous waste heat into heat recovery loops using both constant and variable temperature storage. Oliveira et al. [4] addressed process integration in a multiperiod sugarcane biorefinery, demonstrating the importance of capturing temporal variations in process conditions. Olsen et al. [5] developed a methodology for thermal energy storage integration based on pinch analysis principles. However, efficient optimisation methods that can simultaneously determine optimal storage parameters—temperature, size, and initial charge—within multiperiod HENS frameworks remain an active area of research.

This work presents a novel analytical optimisation framework for integrating ideal heat storage systems within multiperiod HENS. By using physical insights from pinch analysis principles, we derive closed-form expressions and efficient solution procedures for determining optimal storage temperature and initial charge. We distinguish between infinite-size and finite-size storage cases, showing that the infinite-size formulation offers significant computational advantages and can serve as a screening tool before considering size constraints. The methodology is implemented within COMPHENS [5], an open-source HENS optimisation framework, enabling its application to industrial-scale problems.

We validate the approach using real operational data from a lubricant manufacturing plant in the Mediterranean region, considering solar thermal collector integration. The results quantify the relationship between storage size and solar thermal utilisation, revealing diminishing returns as storage capacity increases and identifying the minimum storage size required for complete renewable heat utilisation.

MODELLING

Heat exchanger network synthesis

The HENS modelling framework utilised in this work is the COMPHENS framework [6], which builds upon the seqHENS framework [7], [8]. This splits the problem into four subproblems: first, the minimum utilities required are determined; second, the minimum number of heat exchangers is determined; third, the optimal heat-load distributions are determined; and finally, the network generation and optimisation are performed. In this work, we focus primarily on the first subproblem, implementing heat storage as a preprocessing step to minimise utility costs. To be able to store and use the heat this also means that we will consider multiperiod HENS problems, where the heat storage links these periods together.

Heat storage

To derive an analytical approach to the heat storage implementation we consider an ideal heat storage, i.e. no losses and constant temperature. In addition, we assume that we have a cyclic process so the heat storage needs to start and end with the same charge level. In practice this means that there is no accumulation or depletion of heat over the cycle and that the heat storage is sized perfectly. It also means that the plant operates at consistent repeating cycles, e.g. over a day or a week.

Storage temperature

To determine the optimal temperature of the heat storage we use Pinch Analysis. A priori we know that it is only useful to charge the storage below the lowest pinch in each period and it is only useful to discharge it above the largest pinch in each period, otherwise we would need to increase the usage of hot or cold utilities in the period, either to directly charge the storage or to cover the extra heat or cold not available at the correct temperature somewhere else in the process. This already limits the storage temperature to be below the maximum pinch and above the minimum pinch of the overall problem. Furthermore, the optimal use of the heat storage is to charge and discharge it as much as possible over the periods to minimize the utility consumption. To determine the maximum chargeable or dischargeable heat in a single period at a given temperature, we use the Grand Composite Curve (GCC). The maximum chargeable heat is the heat that can be consumed to heat an extra cold stream at that temperature and create an extra pinch point at or below the given temperature (shifted up by $\Delta T_{min}/2$). Similarly, the maximum heat that can be discharged would be the maximum heat that could be delivered from an extra hot stream and create a single pinch point at or above the given temperature (shifted down by $\Delta T_{min}/2$).

By evaluating the maximum chargeable and dischargeable heat as functions of temperature for each period, we can identify the temperature at which these quantities are balanced across all periods. This balance point minimises utility consumption by maximising storage utilisation. For an infinite-size storage, this optimal temperature, $T_{storage}$, is found by solving:

$$T_{storage} = \underset{T}{\text{root}} \sum_p Q_p(T), \tag{1}$$

where T is the storage temperature, p denotes the period and Q_p the heat transferred to or from the storage in period p subject to the constraints outlined above, with charging of the storage having a positive sign and discharging having a negative sign. This equation finds the temperature where the net heat transferred to/from

storage over all periods equals zero, ensuring cyclic operation while maximizing the absolute heat transfer.

To derive Equation (1) we have assumed that the heat storage has infinite size and heat capacity, so it is always able to store the heat provided and deliver the heat requested in each period. The infinite storage assumption simplifies the problem significantly, as we only need to find where net heat transfer equals zero. However, for a finite storage, the optimization becomes more complex. In each period, charging is limited by the available capacity (maximum charge minus current charge), and discharging is limited by the current stored energy. This transforms the problem from finding a root to maximizing the total heat transferred subject to the cyclic constraint:

$$T_{storage} = \underset{T}{\operatorname{argmax}} \sum_p |Q_p(T)| \;\; s.t. \sum_p Q_p(T) = 0. \tag{2}$$

Initial storage charge and storage size

Similarly to the storage temperature the optimal initial charge is the one which maximizes the amount of heat that is transferred between the different periods, and therefore that is charged and discharged from the storage, while making sure that the start and end charge of the storage are the same and no extra utilities are used in any period. If the storage starts at a charge that is large enough such that it never reaches 0,

$$\min_p Q_p > 0, \tag{3}$$

then increasing the initial storage charge by a would just lift the charge in all periods by a. Doing this for a second initial charge b which also respects Equation (3) one can easily to find the initial charge, Q_i, where the storage just reaches 0 charge once at the end of a single period by linear extrapolation. Ensuring Equation (3) is respected can also easily be done by making sure a and b are larger than the maximum of the total cold and hot utilities when no heat storage is present. Equivalently, one could find the minimum initial charge where the storage is never fully emptied, but has an infinitesimal charge left at the minimum:

$$Q_i = \min_{q_i} q_i \;\; s.t. \min_p Q_p(q_i) > 0. \tag{4}$$

This represents the critical initial charge: any lower and the storage becomes fully depleted in at least one period (limiting heat transfer), any higher and the storage simply operates at an elevated baseline without transferring additional heat, essentially moving away from optimum. From the optimal initial charge, we can also find the optimal size of the storage, $Q_{storage}$, which just equals the maximum stored after any period:

$$Q_{storage} = \max_p Q_p. \tag{5}$$

If we again limit the storage to a finite size we end up with an optimization problem, this time finding the initial charge, Q_i, that maximizes the heat transferred:

$$Q_i = \underset{q_i}{\operatorname{argmax}} \sum_p |Q_p(\mathrm{q}_i)| \;\; s.t. \sum_p Q_p(q_i) = 0. \tag{6}$$

However, if the size of the storage is greater than the optimal size from Equation (5) we can instead use Equation (4) to find the initial charge. As these two equations can be solved simultaneously with only two evaluations, they are much faster than the optimization problem in Equation (6) and it will typically be worth checking them first even for a finite storage. In addition, if the storage temperature is not specified Equation (1) must be solved together with Equation (4) for an infinite storage, which is much faster than solving Equation (2) together with Equation (6) for a finite storage, increasing the speed advantage of checking the infinite storage first.

HENS integration

To add the heat storage to the HENS framework, we need to find both the storage temperature and the initial charge, given either an infinite or a finite storage. These depend on one another and can be solved simultaneously or sequentially. As finding the initial storage charge is very easy for the infinite-size storage, we do not want to redo this, so the storage temperature is solved each time for a given initial storage charge, both for the infinite and finite storage size.

With the heat storage temperature and initial charge we can find the heat transferred in each period. Based on the heat transferred we add a stream to each period of the HENS problem, with a hot stream being added when the storage is discharging heat and a cold stream being added when the storage is charging, with their heat capacity flow rate being dependent on the heat and duration of the period.

CASE

The case considers a chemical plant producing lubricants, located in the mediterranean region. The plant operates for 16 hours, five days a week. The heat demand consists of heating base oils and additives before they are pumped from storage tanks to mixing tanks, heating the mixing tanks, heating for an intermediate drum storage, and space heating. These heat demands are all relatively low temperature, and with little cooling demand, no internal heat exchange is currently in place. Currently, the heat demand of the plant is covered by a hot oil loop which is heated centrally by a gas boiler and then delivers the heat through several separate loops. The plant is looking into building solar thermal collectors to cover some of the heat demand and reduce the usage of the natural gas boiler.

To get an accurate picture of the time dependency of the heat demand a month-long measurement

campaign was performed and combined with operational data from the plant. The heat produced from the solar thermal collectors were modelled on an hourly basis and averaged into a morning, midday and afternoon production giving three periods of different heat production which was assumed to be the same each day. As the plant is not operational on weekends the solar thermal collectors were assumed to be defocused at this time and the heat storage could therefore not be charged on the weekends. This lead to a total heat production of 18.9 MWh from the solar thermal collectors over the period investigated.

The heat demand of the site and the heat production of the solar thermal collectors were used to optimize the solar thermal heat usage, and due to the large numbers of consumers with intermittent heat demand it was decided that the solar thermal heat would be connected to the existing hot oil loop alongside the gas boiler. To best utilise the solar thermal energy a heat storage is also required with the required heat storage capacity for this plant being investigated.

RESULTS

To investigate the impact of the heat storage on the amount of solar thermal energy used, the heat storage was optimised for varying sizes using the method described in the Methodology section. The resulting fraction of solar thermal heat utilised is shown in Figure 1. As expected, we see that by introducing the storage, we increase the fraction of the solar thermal heat that can be utilised. This increase is also nonlinear, with the gain in heat usage largest for small storages and flattening out for larger storages. For instance, the increase is greater from no storage to 50 kWh than from 50 kWh to 100 kWh, and this increase is again larger than the increase from 500 kWh to 750 kWh. This is because, for a small storage, increasing the size slightly allows it to store and discharge additional heat over multiple periods, whereas increasing the size of an already large storage only allows some additional heat to be transferred between a few periods. At the extreme is when all the solar thermal is utilised, as we see for a storage size of 1040 kWh, where increasing the size further would not increase the heat utilisation since there is no more heat to transfer between periods.

While a storage size of 1040 kWh achieves 100% solar thermal utilisation, the economic optimum may lie at a smaller capacity. The relationship between storage size and utilisation suggests strong diminishing returns: a 250 kWh storage (24% of the maximum size) captures over 90% of available solar thermal energy, whereas a 500 kWh storage (48% of the maximum size) captures approximately 95%. From a techno-economic perspective, the marginal benefit of the last 10% of utilisation must be weighed against the capital and operating costs of doubling storage capacity from 500 to 1000+ kWh. Additionally, the ratio of storage capacity to total heat production (1040 kWh / 18, 900 kWh ≈ 5.5%) indicates relatively efficient storage sizing, though this ratio would vary with different operational patterns and solar resource availability.

However, determining the optimal plant size would require a techno-economic analysis to evaluate the trade-off between increased heat utilisation and increased storage costs as size increases, which is outside the scope of this study.

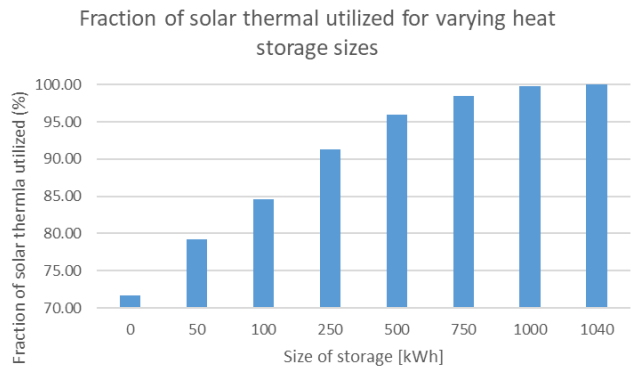

Figure 1. Fraction of solar thermal heat utilized as a function of the size of the heat storage.

CONCLUSION

A model for an ideal heat storage suitable for implementation in multiperiod HENS frameworks has been developed in this work. Based on pinch analysis principles, we have shown efficient methods for optimising both the temperature and initial charge of the storage, with particular emphasis on the infinite-size case, which requires only solving a root-finding problem and performing linear extrapolation, compared to two optimisation problems for finite-size storage. This computational advantage makes the infinite-size approach valuable as a screening tool to establish upper bounds on storage benefits.

Applying this methodology to optimise solar thermal integration at a lubricant manufacturing plant demonstrates that heat storage significantly enhances renewable energy utilisation, with diminishing returns as storage size increases. A relatively modest storage of 250 kWh achieves over 90% solar thermal utilisation, while complete utilisation requires 1040 kWh—representing only 5.5% of the total heat production over the analysed period. This efficient storage-to-production ratio suggests that well-designed thermal storage can enable high solar thermal penetration in industrial applications with manageable storage requirements.

Future work should extend this methodology to non-ideal storage systems that account for heat losses and

temperature variations, and conduct comprehensive techno-economic analyses to identify the economically optimal storage sizing, considering both capital costs and renewable energy benefits.

ACKNOWLEDGEMENTS

This work has been supported by the INDHEAP project, funded by the European Union under the Horizon Europe research and innovation programme, under the grant agreement number 101136140.

AUTHOR IDENTIFIERS

Author ORCIDs:
Falch H: 0009-0003-0797-5476
Anantharaman R: 0000-0001-9228-6197

REFERENCES

1. Thiel GP, Stark AK. To decarbonize industry, we must decarbonize heat. Joule 5:531-550 (2021). https://doi.org/10.1016/j.joule.2020.12.007
2. M. Spörk-Dür, "Solar Heat WorldWide Edition 2025, " AEE INTEC, Austria, Jun. 2025. Accessed: Jan. 30, 2026. [Online]. Available: https://www.iea-shc.org/Data/Sites/1/publications/Solar-Heat-Worldwide-2025.pdf
3. Walmsley TG, Walmsley MRW, Atkins MJ, Neale JR. Integration of industrial solar and gaseous waste heat into heat recovery loops using constant and variable temperature storage. Energy 75:53-67 (2014). https://doi.org/10.1016/j.energy.2014.01.103
4. Oliveira CM, Pavão LV, Ravagnani MASS, Cruz AJG, Costa CBB. Process integration of a multiperiod sugarcane biorefinery. Applied Energy 213:520-539 (2018). https://doi.org/10.1016/j.apenergy.2017.11.020
5. Olsen D, Liem P, Abdelouadoud Y, Wellig B. Thermal energy storage integration based on pinch analysis – methodology and application. Chemie Ingenieur Technik 89:598-606 (2017). https://doi.org/10.1002/cite.201600103
6. A. Subramanian, F. Holtorf, R. Anantharaman, and T. Gundersen, *CompHENS: Computational Tools for Heat Exchanger Network Synthesis*. (May 09, 2023). Zenodo. doi: 10.5281/zenodo.7914754.
7. R. Anantharaman and T. Gundersen, "Developments in the sequential framework for heat exchanger network synthesis of industrial size problems, " in *Computer Aided Chemical Engineering*, vol. 21, Elsevier, 2006, pp. 725–730. doi: 10.1016/S1570-7946(06)80131-8.
8. R. Anantharaman, "Energy Efficiency in Process Plants with emphasis on Heat Exchanger Networks: Optimization, Thermodynamics and Insight, " Doctoral thesis, Norges teknisk-naturvitenskapelige universitet, Fakultet for ingeniørvitenskap og teknologi, Institutt for energi- og prosessteknikk, 2011. Accessed: Mar. 27, 2025. [Online]. Available: https://ntnuopen.ntnu.no/ntnu-xmlui/handle/11250/234562

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Energy planning towards absolute environmental sustainability: identifying key demand-side sufficiency levers to stay within planetary boundaries using sensitivity analysis tool

Nicolas Ghuys[a], Diederik Coppitters[a], Anne van den oever[b], Mahdi Kchaou[a], Hervé Jeanmart[a] and Francesco Contino[a]

[a] UCLouvain, Louvain-la-Neuve, Belgium
[b] Vrij Universiteit Brussels, Brussels, Belgium
* Corresponding Author: nicolas.ghuys@uclouvain.be.

ABSTRACT

Human activities have already transgressed several planetary boundaries, yet energy system models remain largely focused on greenhouse gas mitigation, reflecting their original purpose of addressing climate change. Recent integrations of Planetary boundary-based Life Cycle Assessment into Energy System Optimisation Models show that even cost-optimal low-carbon pathways systematically violate multiple planetary boundaries, indicating that supply-side decarbonisation alone is insufficient for absolute environmental sustainability. At a 2050 horizon, where energy supply is largely decarbonised and technologies are assumed mature, further impact reductions through techno-economic optimisation become limited, positioning final energy demand as a key remaining lever for restoring feasibility under planetary constraints. To address this gap, we extend an Energy System Optimisation framework coupled with a Planetary Boundary framework by explicitly treating final energy demand as a decision variable and exploring energy sufficiency configurations within a multi-objective formulation minimising system cost and environmental pressures. The national-scale EnergyScopeTD model is coupled with the RHEIA uncertainty analysis framework to propagate uncertainty in final energy demands and to identify demand-side drivers through global sensitivity analysis. Applied to Belgium, the results show that individual mobility demand is the dominant driver of environmental pressures across multiple planetary boundaries, with consistently higher influence than other end-use demands. While large-scale vehicle fleet electrification substantially reduces climate impacts, it simultaneously shifts pressures toward other environmental dimensions, indicating that technological substitution alone, even combined with moderate sufficiency, is insufficient to achieve absolute sustainability. Overall, the results confirm that energy sufficiency, particularly in mobility, is necessary for operating energy systems within planetary limits, and that anchoring system design in absolute sustainability thresholds provides a robust basis for prioritising demand-side mitigation strategies.

Keywords: Energy system model, Optimisation, Planetary boundary, Sufficiency, Sensitivity analysis

INTRODUCTION

Background

With seven out of nine planetary boundaries already transgressed, ensuring that human activities remain within the Earth's safe operating space has become a central challenge for sustainability science [1, 2]. Although climate change has historically dominated environmental policy and energy planning, it is now well established that a narrow focus on greenhouse gas mitigation risks shifting environmental pressures toward other critical Earth system processes, including biodiversity

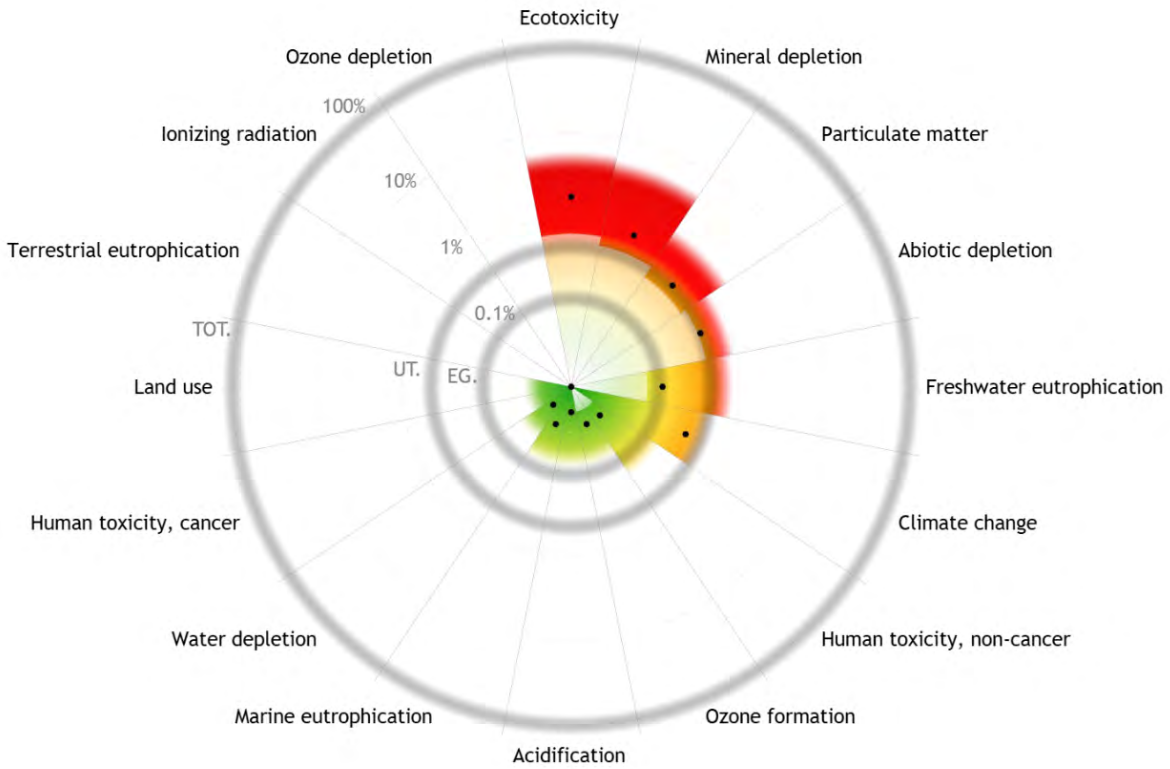

Figure 1. Each LCA indicator obtained from the belgian energy system designs satisfying the constraint of 20 Mt CO2,eq emissions in 2045, aligned with the work of Van den Oever [19], and a total system cost below 130% of the optimal-cost solution can vary within a solution space represented by the strongly colored areas. The indicator values corresponding to the optimal-cost solution are shown as black dots. Three boundaries are displayed to facilitate interpretation: moving from the center outward, the first corresponds to the egalitarian sharing principle (EG.), the second to the utilitarian sharing principle (UT.), and the outermost boundary represents the total PB (at 100%), as reported in Table 1. Because the PBs and the associated risk of exceeding them are uncertain, the boundaries are deliberately represented as blurred rather than as sharp, well-defined lines.

loss, freshwater degradation, and material depletion [3, 4]. Achieving climate neutrality alone is therefore insufficient to guarantee a long-term viable Earth system capable of sustaining ecosystems, essential resources, and human well-being [5].

Energy systems are central to this challenge. As energy transitions progress toward low carbon configurations dominated by renewable resources, electrification, and synthetic fuels, the importance of construction phase relative to the full life cycle impacts increases [6, 7]. In such systems, a growing share of environmental pressures originates from material extraction, infrastructure deployment, and upstream supply chains rather than from direct operational emissions [8]. This shift implies that future energy system designs must be assessed not only against climate targets but also against a broader set of environmental constraints reflecting absolute ecological limits.

In response, recent research has proposed integrating life cycle assessment and the planetary boundaries framework into energy system optimisation models, enabling the evaluation of energy transitions against absolute environmental thresholds rather than relative impact reductions [9, 10]. Building on this approach, a methodological framework was recently developed to embed prospective life cycle assessment and planetary boundary-based normalisation directly into national scale energy system optimisation. Applied to a Belgian case study, this framework showed that even under stringent carbon constraints, energy system configurations systematically violate several planetary boundaries [9] (shown in Figure 1).

These results point to a fundamental limitation of supply focused energy system optimisation and raise a critical question for energy system planning. If advanced, cost efficient, and climate compatible energy systems remain environmentally unsustainable when assessed against planetary boundaries, where does the remaining leverage for action lie?

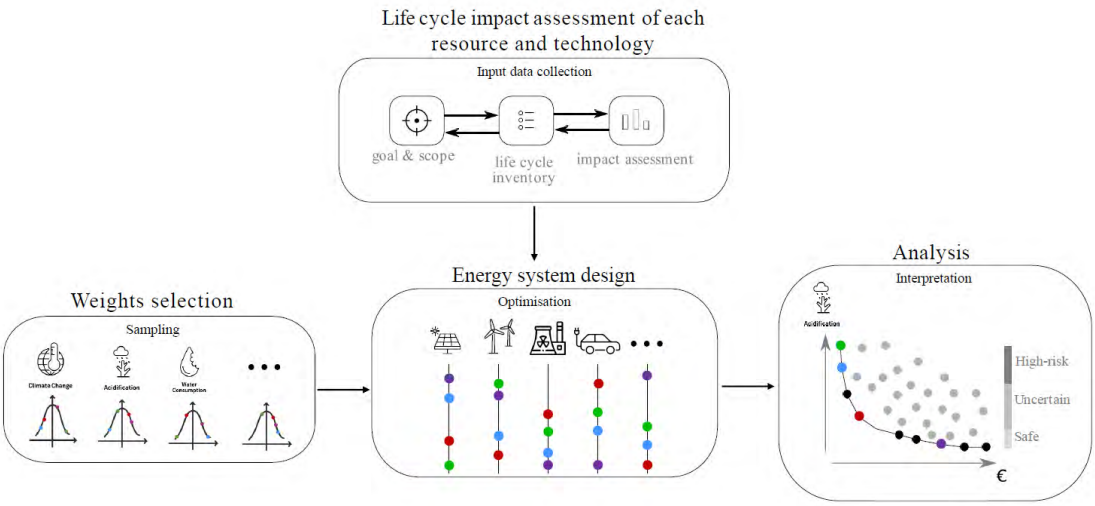

Figure 2. Schematic overview of the integrated modeling framework combining life cycle assessment and planetary boundary analysis with energy system optimisation. The process begins with comprehensive input data collection for the energy system, including technology performance, resource availability, and a prospective life cycle inventory (LCI) for all technologies and fuels. These data feed into the energy system design optimisation model (EnergyScopeTD), which has been augmented with life cycle impact assessment capabilities. Within the model, the LCI data are converted into environmental impacts across multiple planetary boundary (PB) indicators (e.g. climate change, land use, water use) and embedded into the optimisation structure. The model then performs a multi-criteria optimisation that jointly minimises the total system cost and a portfolio of environmental impact indicators, with each impact normalised by its PB threshold to ensure comparability. To explore trade-offs between cost and environmental objectives, a sampling of weight combinations is used, generating a broad Pareto front of optimised solutions rather than a single optimum.

Literature review

Energy system optimisation models have long been used to support strategic energy planning by identifying cost optimal pathways to reduce greenhouse gas emissions and meet climate targets [11, 12]. These models have traditionally emphasised supply side transformations, such as fuel switching, electrification, and the deployment of low carbon technologies, while treating final energy demand as an exogenous and fixed input. Within this paradigm, environmental performance is most often assessed through operational carbon emissions, reinforcing a narrow focus on climate mitigation that has been described as carbon tunnel vision [3, 4].

In parallel, life cycle assessment has been widely applied to quantify environmental impacts across multiple categories and life cycle stages, highlighting the risk of burden shifting when focusing on a single indicator. The integration of life cycle assessment into energy system models has substantially improved the environmental realism of scenario analysis by accounting for upstream supply chains and infrastructure related impacts [13]. However, conventional life cycle assessment remains a relative assessment tool and does not indicate whether a system is environmentally sustainable in absolute terms. The planetary boundaries framework addresses this limitation by defining scientifically grounded thresholds for key Earth system processes, providing an absolute reference against which environmental pressures can be evaluated [14, 15]. Recent studies coupling planetary boundaries with life cycle assessment have shown that many energy pathways compatible with climate targets still exceed planetary boundaries for indicators such as freshwater eutrophication, particulate matter formation, or material use [9]. Yet, the explicit integration of planetary boundaries into large scale energy system optimisation remains methodologically challenging due to issues related to indicator heterogeneity, normalisation, and burden sharing.

To address these challenges, recent methodological advances have introduced multi objective optimisation and goal programming formulations to jointly consider system cost and multiple environmental objectives, combined with systematic exploration of near optimal solution spaces [9, 10] (shown in Figure 2). Existing analyses of optimised energy systems provide valuable insights into technology trade-offs and system designs. However, these analyses almost systematically assume fixed final energy demand, thereby restricting the exploration of mitigation pathways to supply-side resources and technologies.

Research gap and contribution

Planetary boundary informed energy system optimisation consistently reveals a structural limitation. When multiple planetary boundaries are considered alongside climate targets, cost optimal designs systematically violate several boundaries, and reformulating optimisation objectives does not resolve this issue due to persistent antagonisms between impact categories [9]. Under such conditions, no feasible solution satisfies cost constraints, climate targets, and planetary boundaries within existing optimisation formulations.

Despite these findings, most energy system optimisation studies continue to treat final energy demand as fixed, focusing mitigation efforts almost exclusively on supply side technologies and resources. As a result, the potential of demand-side changes to alleviate environmental pressures under absolute sustainability constraints remains largely unexplored within quantitative energy system optimisation.

This paper addresses this gap by explicitly integrating demand side sufficiency into planetary boundary constrained energy system optimisation. By coupling energy system optimisation with uncertainty propagation and sensitivity analysis, the approach quantifies the relative influence of different demand reduction strategies on environmental pressures, enabling the identification of priority demand side levers without making normative assumptions on social acceptability.

METHODOLOGY

Overview of the ESOM–LCA–PB framework

This work builds upon the ESTD-PB modelling framework, which integrates a national-scale bottom-up energy system optimisation model (ESOM) with prospective life cycle assessment (LCA) and the planetary boundaries (PB) framework. The full mathematical formulation and data assumptions are documented in previous publications [9, 12]; the present section recalls only the elements required to understand the analysis introduced in this paper.

The model determines the optimal investment and operation of energy conversion, storage, and supply technologies required to satisfy exogenously defined final energy service demands at minimum total annualised system cost. The optimisation problem is formulated as a linear program with the following objective function:

$$\min \quad C_{tot} = \sum_{j \in TECH} \left(\tau_j C_j^{inv} + C_j^{maint}\right) + \sum_{i \in RES} C_i^{op}$$

where C_j^{inv}, C_j^{maint}, and C_i^{op} denote investment, maintenance, and operational costs, respectively, and τ_j is the annualisation factor associated with technology lifetime and discount rate. The optimisation is subject to energy balance, capacity, operational, resource availability, and policy constraints, including a climate constraint.

Environmental impacts are assessed using a cradle-to-grave life cycle assessment consistent with prospective background assumptions. For a given environmental impact category k, total system impacts are computed as the sum of impacts associated with resource extraction/prodution, technology construction, and system operation:

$$LCA_{tot}^k = \sum_{i \in RES} LCA_{res}^k(i) + \sum_{j \in TECH} \frac{LCA_{constr}^k(j)}{L_j} + \sum_{j \in TECH} LCA_{op}^k(j)$$

To enable an absolute interpretation of environmental performance, life-cycle impacts are normalised using the planetary boundaries framework:

$$f_{PB}^k = \frac{LCA_{tot}^k}{PB^k}$$

where PB^k denotes the global safe operating limit associated with Earth-system process k.

Uncertainty propagation and sensitivity analysis

To propagate uncertainty in final energy service demands through the integrated ESOM–LCA–PB framework, this study applies the open-source Python framework RHEIA for uncertainty propagation and global sensitivity analysis [18]. The methodology is used here without modification and is coupled to the energy system model to analyse the effects of demand variability on environmental indicators.

Selected final energy service demands are treated as uncertain input parameters, define as uniform distribution ranging between 70% and 100% of the nominal demand value, while all other model inputs are kept deterministic. These demands are characterised by prescribed probability distributions around a reference level and define a random input vector $\boldsymbol{\xi} = (\xi_1, \dots, \xi_d)$. For each realisation of $\boldsymbol{\xi}$, sampled via quasi-random Sobol sampling, the energy system optimisation problem is solved, and the resulting life-cycle impacts are computed and normalised using the planetary boundaries framework.
Given the computational cost of repeatedly solving a national-scale energy system optimisation model, uncertainty propagation is performed using Polynomial Chaos Expansion (PCE) surrogate modelling, as implemented in RHEIA. For each environmental indicator of interest, the model response is approximated by a truncated polynomial expansion:

$$\hat{Y}(\boldsymbol{\xi}) = \sum_{\alpha \in \mathcal{A}_{d,p}} u_\alpha \, \Psi_\alpha(\boldsymbol{\xi})$$

where Ψ_α are multivariate orthogonal polynomials consistent with the input distributions, $u_\alpha$are regression coefficients, and $\mathcal{A}_{d,p}$ denotes the set of multi-indices with total degree $|\alpha| \leq p$.

The coefficients u_α are estimated by regression using a limited number of model evaluations generated through space-filling sampling. The polynomial degree is increased until a predefined accuracy criterion is met, as assessed through leave-one-out cross-validation.

Once constructed, the PCE surrogate enables efficient evaluation of output statistics and supports variance-based global sensitivity analysis. For a given output Y, the total variance $\mathrm{Var}(Y)$ is decomposed into contributions associated with individual uncertain demands and their interactions. The first-order Sobol sensitivity index associated with demand $i$is defined as:

$$S_i = \frac{\mathrm{Var}_{\xi_i}(\mathbb{E}[Y \mid \xi_i])}{\mathrm{Var}(Y)}$$

and quantifies the contribution of variations in ξ_i to the total output variance. Total Sobol indices are also computed to capture higher-order interaction effects. All sensitivity indices are derived analytically from the PCE coefficients.

RESULTS AND DISCUSSION

Case study description

Based on the results presented in Figure 1, the analysis focuses on three environmental indicators identified as the most critical in the reference system: ecotoxicity, particulate matter formation, and freshwater eutrophication. Indicators related to mineral and abiotic depletion are not considered in this case study, as they reflect constraints on resource availability and long-term human activity sustainability rather than direct pressures on environmental systems, which constitute the core focus of this work. The analysis is conducted in two successive stages. First, a global sensitivity analysis is performed to rank final energy service demands according to their influence on each selected environmental indicator. To this end, a second-order Polynomial Chaos Expansion is constructed over the full set of final energy service demands, enabling the computation of Sobol sensitivity indices.

For presentation and synthesis purposes, Sobol indices are aggregated across the selected environmental indicators and across individual energy service categories. The results are reported as box plots, where each box represents the distribution of Sobol indices obtained for a given demand type, pooling the indices computed for all considered environmental indicators and corresponding demand components. In a second step, targeted demand reduction scenarios are defined based on the outcomes of the sensitivity analysis.

Demand Sobol's indices

The sensitivity analysis shows that variations in mobility demand are the dominant driver of environmental impacts in the energy system.

Across all considered environmental indicators, mobility demand consistently exhibits the highest influence, independent of the specific critical planetary boundary considered. This dominance is reflected in the limited dispersion of Sobol index values associated with mobility demand compared to other final energy service categories, indicating a systematic effect.

These results are consistent with previous findings by Gauthier Limpens [12], who identified individual mobility, particularly through the renewal of the Belgian vehicle fleet, as the largest cost component of the energy transition. Together, these findings suggest a potential alignment between economic and environmental objectives when defining transition pathways toward a more sustainable mobility system. This conclusion is further supported by previous work, who reported a strong covariance between electric vehicle deployment and critical environmental indicators [9]. The present analysis extends this evidence by demonstrating that such relationships are embedded in the global sensitivity structure of the model, with mobility demand emerging as the primary demand-side lever affecting both environmental pressures and total system cost.

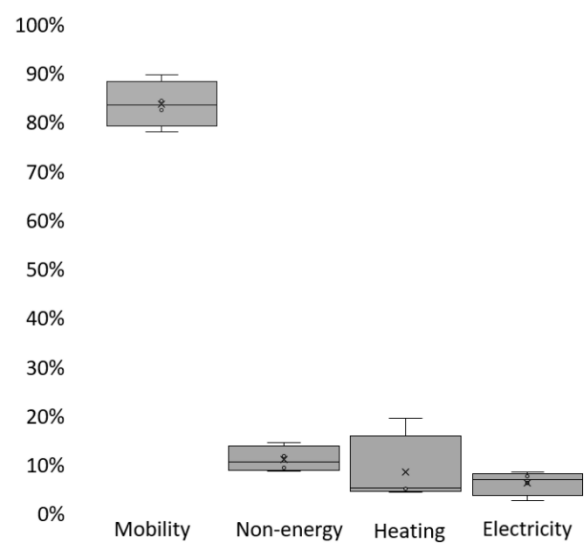

Figure 3. Distribution of total Sobol sensitivity indices associated with final energy demand categories across all environmental indicators.

Impact of end-use demand reduction

Building on the Sobol-based sensitivity analysis, mobility demand reduction emerges as the most impactful lever among the considered demand-side mitigation options. To explore this result further, the environmental

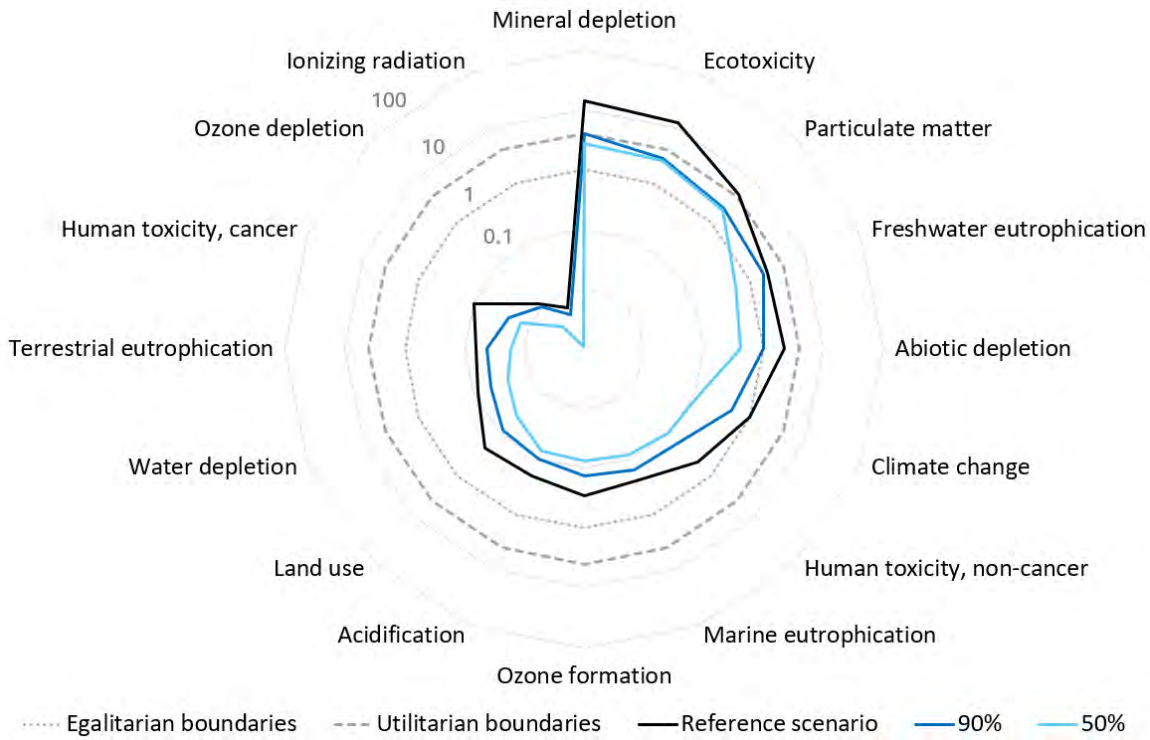

Figure 4. Each Life Cycle Assessment (LCA) indicator is shown for Belgian energy system configurations satisfying a constraint of 20 Mt CO_2-eq emissions in 2045, consistent with previous work [19], and a total system cost below 130% of the cost-optimal solution. The cost-optimal solution is indicated by the black line, while the 10% and 50% mobility demand reduction scenarios are shown in dark blue and light blue, respectively. Three planetary boundary thresholds are displayed to facilitate interpretation: from the center outward, the innermost boundary corresponds to the egalitarian sharing principle (EG., dotted), the intermediate boundary to the utilitarian sharing principle (UT., dashed), and the outermost boundary represents the full planetary boundary (100%).

implications of stepwise mobility demand reductions are assessed, considering two reduction levels (−10% and −50%) within the studied energy system. At this stage, neither the social acceptability nor the economic feasibility of such reductions is assessed. The analysis exclusively focuses on their effect on absolute environmental sustainability, with the objective of providing a quantified and robust benchmark against planetary boundaries that can later inform socially grounded sufficiency scenarios.

At the first reduction level (−10%), a system configuration complying with all planetary boundaries is achieved under the most permissive sharing principle, the utilitarian allocation, which is based on gross domestic product, favors the Belgian energy system in a global context by allowing differentiated environmental space across countries [9]. When sustainability is instead evaluated using a more constraining population-based sharing principle, the egalitarian sharing principle, three planetary boundaries remain exceeded: ecotoxicity, particulate matter formation and freshwater eutrophication. As noted previously, mineral and abiotic depletion indicators are excluded from the planetary boundary assessment, as they reflect constraints on the continuity of human activities rather than environmental system thresholds. At the second reduction level (-50%), two planetary boundaries remain exceeded under the population-based sharing principle, with only marginal improvements compared to the first reduction level.

DISCUSSION

The sensitivity analysis shows that mobility demand is the main driver of environmental impacts in the energy system across all planetary boundaries at risk. Its Sobol indices are consistently higher than those of other end-use demands and exhibit limited dispersion, indicating a dominant and systematic influence that is observed across the different critical environmental pressures considered. These findings are consistent with previous studies highlighting individual mobility as a key challenge of the energy transition, both economically and environmentally. While earlier work emphasized cost dominance

and correlations with environmental indicators, the present analysis shows that mobility demand structurally drives environmental pressures at the system level, suggesting an alignment between environmental sustainability and economic efficiency in mobility-oriented strategies. The demand reduction analysis shows that moderate mobility reductions substantially lower environmental pressures and enable planetary-boundary compliance under permissive allocation rules. However, under stricter population-based principles, several boundaries remain exceeded, indicating that technological substitution alone, even with modest sufficiency, is insufficient. Larger reductions further reduce impacts but show diminishing returns for the most constraining boundaries. Overall, mobility demand reduction emerges as a necessary but not sufficient condition for achieving absolute sustainability, requiring complementary demand-side measures in other sectors alongside targeted technological and structural interventions.

ACKNOWLEDGEMENTS

N.G acknowledges the support of the European Union's Horizon 2020 research and innovation program under grant agreement 101075660. D.C. acknowledges the support of the Fonds de la Recherche Scientifique - FNRS [CR 40016260].

REFERENCES

1. Steffen W. et al. Planetary boundaries: guiding human development on a changing planet. *Science* 347, 1259855 (2015) https://doi.org/10.1126/science.1259855
2. IPCC. Climate change 2021: The physical science basis. https://www.ipcc.ch/report/ar6/wg1/ (2021). Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change.
3. Algunaibet IM, Guillén-Gosálbez G. Life cycle burden-shifting in energy systems designed to minimize greenhouse gas emissions: novel analytical method and application to the united states. Journal of Cleaner Production 229:886-901 (2019). https://doi.org/10.1016/j.jclepro.2019.04.276
4. Laurent A, Olsen SI, Hauschild MZ. Limitations of carbon footprint as indicator of environmental sustainability. Environ. Sci. Technol. 46:4100-4108 (2012). https://doi.org/10.1021/es204163f
5. Algunaibet IM, Pozo C, Galán-Martín Á, Huijbregts MAJ, Mac Dowell N, Guillén-Gosálbez G. Powering sustainable development within planetary boundaries. Energy Environ. Sci. 12:1890-1900 (2019). https://doi.org/10.1039/c8ee03423k
6. Theilig K, Takser I, Reitberger R, Vollmer M, Lang W. Toward zero-emission buildings: a case study on a non-residential building in germany using life cycle assessment and carbon sequestration of green infrastructure. IOP Conf. Ser.: Earth Environ. Sci. 1196:012046 (2023). https://doi.org/10.1088/1755-1315/1196/1/012046
7. Creutzig F., et al. Planetary impact shifts from fossil fuels to material extraction. *arXiv* (2024) https://doi.org/10.22541/au.171901018.84864468/v1
8. Blanco H. et al. Life cycle assessment integration into energy system models: An application for power-to-methane in the eu. *Appl. Energy* 259 (2020) https://doi.org/10.1016/j.apenergy.2019.114121
9. Ghuys N., Coppitters D., Van den Oever A., Kchaou M., Contino F., Jeanmart H. Integrating planetary boundaries into energy system optimisation models for absolute environmental sustainability assessment: a methodological framework. Preprint (2025) https://doi.org/10.21203/rs.3.rs-8251250/v1
10. Schnidrig J, Souttre M, Chuat A, Maréchal F, Margni M. Between green hills and green bills: unveiling the green shades of sustainability and burden shifting through multi-objective optimization in swiss energy system planning. Journal of Environmental Management 370:122537 (2024). https://doi.org/10.1016/j.jenvman.2024.122537
11. Pfenninger S, Hawkes A, Keirstead J. Energy systems modeling for twenty-first century energy challenges. Renewable and Sustainable Energy Reviews 33:74-86 (2014). https://doi.org/10.1016/j.rser.2014.02.003
12. Limpens G, Moret S, Jeanmart H, Maréchal F. Energyscope TD: a novel open-source model for regional energy systems. Applied Energy 255:113729 (2019). https://doi.org/10.1016/j.apenergy.2019.113729
13. Volkart K, Mutel CL, Panos E. Integrating life cycle assessment and energy system modelling: methodology and application to the world energy scenarios. Sustainable Production and Consumption 16:121-133 (2018). https://doi.org/10.1016/j.spc.2018.07.001
14. Ryberg M. W. et al. Development of a life-cycle impact assessment methodology linked to the planetary boundaries framework. *Ecol. Indic.* 88, 250–262 (2018) https://doi.org/10.1016/j.ecolind.2017.12.051
15. Sala S, Crenna E, Secchi M, Sanyé-Mengual E. Environmental sustainability of european production and consumption assessed against planetary boundaries. Journal of Environmental

Management 269:110686 (2020). https://doi.org/10.1016/j.jenvman.2020.110686
16. Coppitters D, Tsirikoglou P, Paepe W, Kyprianidis K, Kalfas A, Contino F. RHEIA: robust design optimization of renewable hydrogen and derived energy carrier systems. JOSS 7:4370 (2022). https://doi.org/10.21105/joss.04370
17. van den Oever Anne (2025). Environmental life cycle assessment for decision-support in the development of advanced biofuels.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Assessing the potential of vehicle-to-grid (V2G) systems using dynamic simulation and life cycle assessment

Ayumi Yamaki[a], and Yasunori Kikuchi[abc]*

[a] The University of Tokyo, Institute for Future Initiatives, Tokyo, Japan
[b] The University of Tokyo, Department of Chemical System Engineering, Tokyo, Japan
[c] The University of Tokyo, Presidential Endowed Chair for "Platinum Society" Organization for Interdisciplinary Research Project, Tokyo, Japan
* Corresponding Author: ykikuchi@ifi.u-tokyo.ac.jp

ABSTRACT

The increasing deployment of variable renewable energy (VRE) is essential for achieving a sustainable society; however, its inherent variability poses challenges for maintaining a stable electricity supply. Vehicle-to-grid (V2G) technology enables bidirectional electricity exchange between electric vehicles (EVs) and the power grid and can enhance the utilization of renewable electricity by charging EVs during periods of VRE output curtailment. This study developed a regional V2G system model and evaluated its potential through energy flow simulations and life cycle assessment (LCA). The model explicitly considered hourly operation schedules of individual EVs, the spatial distribution of V2G infrastructure, and minimum output constraints of thermal power generation. The number of EVs is assumed to increase to up to 10, 000 units. In the energy flow simulations, EV charging and discharging were calculated on an hourly basis over one year. LCA was conducted to assess greenhouse gas (GHG) emissions of regional V2G system. Fuel consumption, VRE utilization, and GHG emissions were quantitatively evaluated. The results demonstrated that regional V2G systems could reduce VRE output curtailment and GHG emissions and, under certain conditions, could also reduce thermal power generation. The developed model enabled a quantitative evaluation of the potential of regional V2G systems. For practical deployment, integrated regional platforms would be required to coordinate electricity flows and to manage financial transactions. By assessing the potential of regional V2G systems, this study would provide valuable insights to support the design and implementation of environmentally sustainable V2G systems.

Keywords: Electric vehicle, Variable renewable energy, Output curtailment, Energy flow simulation

INTRODUCTION

To achieve a sustainable society, variable renewable energy (VRE) is being increasingly installed. However, the inherent variability of VRE makes it challenging to maintain a stable electricity supply. In some regions, VRE generation is curtailed when electricity generation exceeds demand. Electric vehicles (EVs), which are equipped with batteries, offer an opportunity to address this issue: charging EVs during periods of VRE output curtailment can enhance the effective utilization of renewable electricity. Furthermore, vehicle-to-grid (V2G) technology enables bidirectional electricity exchange between EVs and the power grid. Previous studies have shown that the deployment of V2G can reduce greenhouse gas (GHG) emissions, energy consumption and cost at national level [1].

However, the environmental impacts of introducing new energy supply technologies depend on the regional conditions, even for the same technology. Thus, energy systems should be designed considering the regional characteristics [2]. Tanegashima, an island in southern Japan, was the first region in Japan to experience VRE output curtailment. A previous study suggests that V2G deployment in Tanegashima has the potential to reduce GHG emissions [3]. However, that study does not consider the locations of V2G installation or the operational schedules of individual EVs.

In this study, we develop a regional V2G system model and evaluate its potential using energy flow simulations and life cycle assessment (LCA). A regional V2G system is defined as an energy system that integrates EVs and V2G technology within a target region. To evaluate the potential of the system, fuel consumption, VRE utilization, and GHG emissions are quantitatively evaluated under cases with and without EVs and V2G. The proposed model is applied to Tanegashima in Japan.

The model explicitly considers (i) hourly operation schedules of individual EVs, (ii) the spatial distribution of V2H/V2B infrastructure, and (iii) minimum output constraints of thermal power generation. By integrating dynamic energy flow simulation with LCA, the proposed approach enables a quantitative evaluation of the environmental performance of regional V2G systems.

MATERIALS AND METHODS

Model of Regional V2G system

This study constructed a model of a regional V2G system in which EVs and V2G technologies are deployed within a target region. To evaluate the potential of V2G, a large number of vehicle-to-home (V2H) and vehicle-to-building (V2B) systems were assumed to be installed. **Figure 1** (a) shows the concept of the regional V2G system constructed in this study. EVs were assumed to replace internal combustion engine vehicles (ICEVs) and to utilize electricity derived from VRE as much as possible.

Electricity supply in the model consisted of photovoltaic (PV) and wind turbine (WT) as VRE sources and thermal power generation as shown in **Figure 1** (a). Thermal power generation was used to balance electricity demand and supply. Total electricity consumption was estimated based on electricity demand in the target region. The minimum output of thermal power generation is assumed to be unchanged before and after EV deployment. EVs were assumed to charge proactively during periods of VRE output curtailment and to discharge electricity to the power grid via V2H/V2B systems when VRE generation was insufficient.

The model is based on several simplifying assumptions. All EVs are assumed to be use for commuter and to return home every night and to be connected to V2H/V2B systems whenever available. Electricity prices and user-driven charging behavior are not explicitly considered.

Individual vehicle operation schedules and the commuting distances between homes and workplaces were defined. The operation schedules were assumed to be identical for EVs and ICEVs. **Figure 1** (b) shows the vehicle operation schedule and the list of installed equipment. On workdays, individuals were assumed to work from 9:00 to 17:00, with typical driving time at 8:00 and at 18:00. Some individuals randomly worked earlier or later than these hours and may stop at other locations, such as shops, on their way home. Each individual was assumed to have 20 days of paid leave per year, with the dates assigned randomly. On weekends and paid leaves, individuals were assumed either to stay at home (i.e., not drive) or to visit one or two shops. Destinations, departure times, and return times on these days were determined randomly.

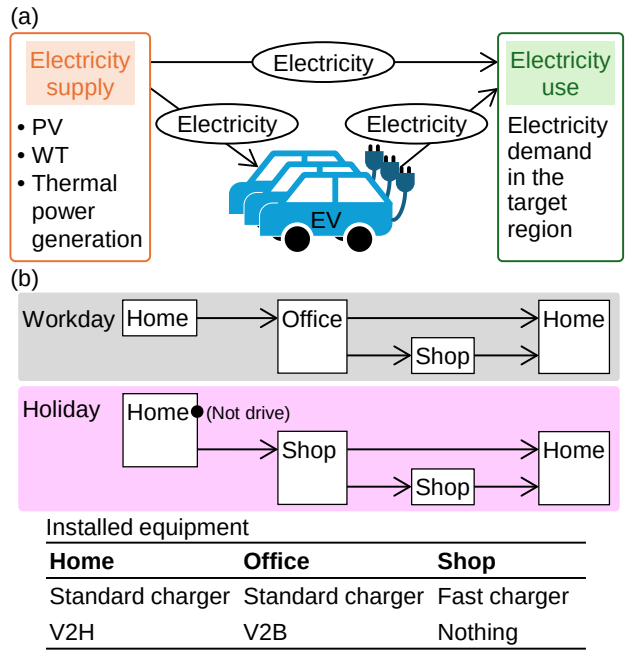

Figure 1. (a) Concept of regional V2G system. (b) EV drive schedule and installed equipment list.

If an individual owned an EV, their home was assumed to be equipped with either standard EV charger or V2H system. The offices were similarly assumed to have either standard EV charger or V2B system, as shown in **Figure 1** (b). At shops visited after work or on holidays, the presence of EV fast chargers was assigned randomly. When EVs were parked at locations equipped with V2H or V2B systems, they were always assumed to be connected to those systems.

Energy Flow Simulation

In the energy flow simulation, the amounts of EV charging and discharging were calculated on an hourly basis for one year. Based on these results, annual VRE output curtailment, EV electricity discharge to the power grid, electricity demand and thermal power generation were analyzed. Simulation cases were examined both with and without EVs and the V2G systems.

Electricity demand and supply must be balanced at all times to ensure stable grid operation; otherwise, disturbances in system frequency could compromise electricity supply. Before the introduction of EVs, the regional electricity demand at time t, $P^{\mathrm{use}}(t)$, was supplied by VRE-

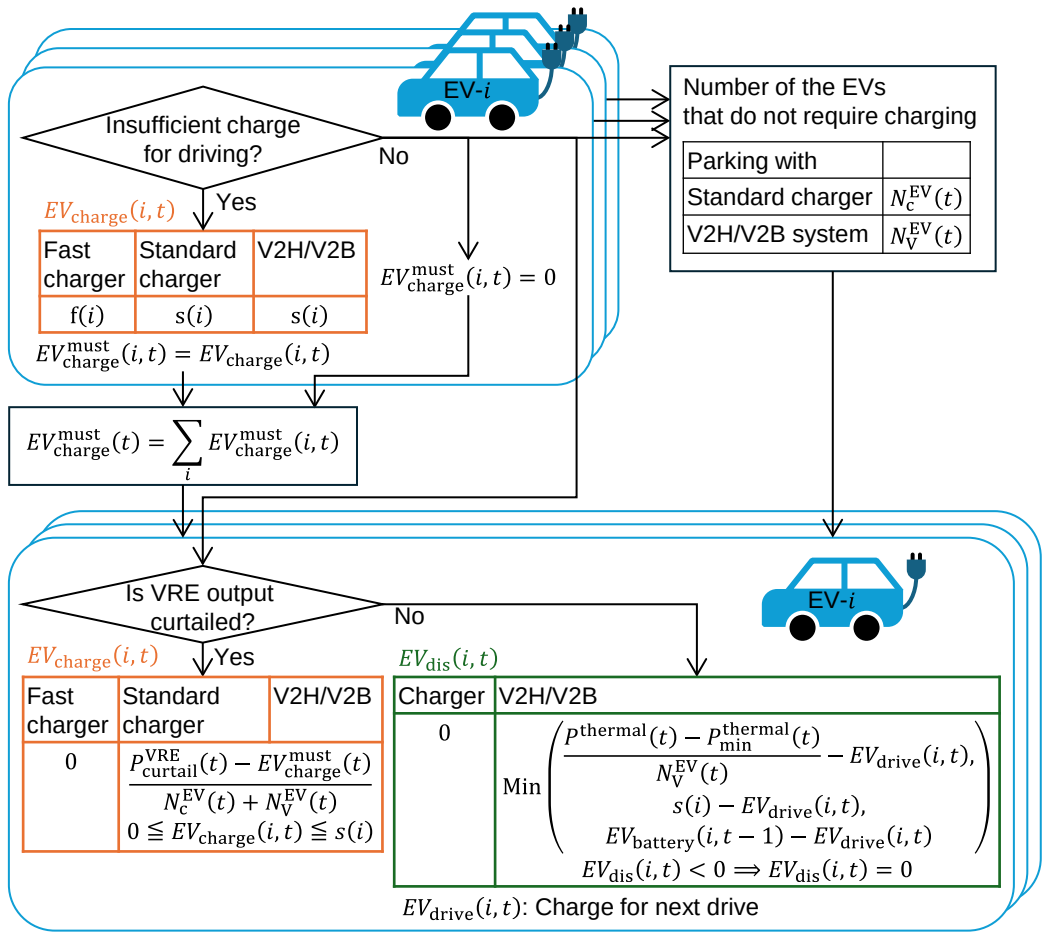

Figure 2. Procedure for determining the charge/discharge amount of EVs parked at the locations with chargers or V2H/V2B at time t. $f(i)$ and $s(i)$ represents charging speed of fast charger and standard charger, respectively.

based electricity, $P^{\mathrm{VRE}}(t)$, and thermal power generation, $P^{\mathrm{thermal}}(t)$ as expressed in Equation (1).

$$P^{\mathrm{use}}(t) \leqq P^{\mathrm{VRE}}(t) + P^{\mathrm{thermal}}(t) \tag{1}$$

In practice, left-hand and right-hand sides of Equation (1) are balanced. When VRE-based electricity generation was high, a portion of VRE electricity was curtailed because the output of thermal power plants could be reduced below a certain minimum level. The amount of VRE output curtailment, $P^{VRE}_{\mathrm{curtail}}(t)$, was calculated using Equation (2).

$$P^{VRE}_{\mathrm{curtail}}(t) = P^{\mathrm{VRE}}(t) + P^{\mathrm{thermal}}_{\mathrm{min}}(t) - P^{\mathrm{use}}(t) \tag{2}$$

Here, $P^{\mathrm{thermal}}_{\mathrm{min}}(t)$ represented the minimum output of thermal power generation. The minimum output was determined by various factors such as VRE electricity generation and electricity demand. In this study, $P^{\mathrm{thermal}}_{\mathrm{min}}(t)$ was assumed to remain before and after the installation of EVs. Similarly, baseline electricity demand excluding EV charging, $P^{\mathrm{use}}(t)$ was assumed to be the same regardless of EVs deployment.

Thermal power generation, $P^{\mathrm{thermal}}(t)$, was determined by the difference between the electricity demand and VRE electricity generation as shown in Equation (3). However, if this difference was smaller than minimum thermal output, thermal power generation was set equal to the minimum output:

$$P^{\mathrm{thermal}}(t) = \mathrm{Max}\left(P^{\mathrm{use}}(t) - P^{\mathrm{VRE}}(t), P^{\mathrm{thermal}}_{\mathrm{min}}(t)\right) \tag{3}$$

Simulation for EV charging and discharging

EV charging and discharging behavior was mathematically expressed as follows. While driving, an EV discharges electricity in proportion to travel distance. When parked, an EV may charge or discharge depending on whether the parking location was equipped with a charger (including V2H/V2B systems) whether VRE electricity generation was being curtailed. The battery state of EV-i at time t was calculated using Equation (4).

$$EV_{\mathrm{battery}}(i,t) = EV_{\mathrm{battery}}(i,t-1) + EV_{\mathrm{charge}}(i,t) - EV_{\mathrm{dis}}(i,t) \tag{4}$$

Here, $EV_{\text{battery}}(i,t)$ represented the amount of electricity stored in the battery of EV-i at time t, while $EV_{\text{charge}}(i,t)$ and $EV_{\text{dis}}(i,t)$ denoted the amount of charging and discharging at time t, respectively.

During driving, EV discharges according to the travel distance. When parked, the charging or discharging behavior was determined by whether the EV was parked at a location equipped with a charger or V2H/V2B system, and whether VRE output curtailment occurred at that time. If an EV was parked at a location without a charger or V2H/V2B, no charging and discharging occurs. When an EV was parked at a location equipped with a charger or V2H/V2B, the calculation procedure for $EV_{\text{charge}}(i,t)$ and $EV_{\text{dis}}(i,t)$ followed the process shown in **Figure 2**.

Before driving, EVs were charged based on the distance between the current parking location at time t and the next destination that had a charger or V2H/V2B. The charging start time and required charging amount were determined by the travel distance. The total amount of electricity that all EVs must charge before driving is defined as $EV_{\text{charge}}^{\text{must}}(t)$ (see **Figure 2**). EVs that did not need to charge at time t were available to function as distributed batteries for the power grid. The numbers of such EVs were counted as $N_{\text{c}}^{\text{EV}}(t)$ and $N_{\text{V}}^{\text{EV}}(t)$, depending on the parking location.

When VRE output was curtailed, EVs parked at the location equipped with standard chargers or V2H/V2B systems were charged. EVs were always connected to V2H or V2B when parked at homes or offices that had such equipment. When EVs were parked at homes or offices with standard chargers, they were charged at a high rate if the battery state of charge was below 50% or if VRE output curtailment occurred. When EVs were parked at locations with fast charger, which meant the EVs were parked at the shops, charging occurred only when the battery state of charge was low.

The amount of EV charging, $EV_{\text{charge}}(i,t)$ was adjusted according to VRE output curtailment, $P_{\text{curtail}}^{VRE}(t)$, and the numbers of EVs available as grid batteries, $N_{\text{c}}^{\text{EV}}(t)$ and $N_{\text{V}}^{\text{EV}}(t)$. The upper limit of $EV_{\text{charge}}(i,t)$ was set to $s(i)$, which represented the charging speed of the standard charger determined by charger performance.

EVs discharged electricity when VRE output was not curtailed and when they were connected to V2H or V2B systems. The discharge amount was calculated based on the controllable output of thermal power generation, $P^{\text{thermal}}(t) - P_{\text{min}}^{\text{thermal}}(t)$, the charging/discharging capacity of the charger, $s(i)$, and the amount of electricity stored in the EV battery, as shown in **Figure 2**. In addition, $EV_{\text{dis}}(i,t)$ was constrained to ensure that sufficient battery charge, $EV_{\text{drive}}(i,t)$, remained for the next driving event, thereby preventing situations in which the EV would need to be recharged immediately before driving due to insufficient remaining battery capacity.

The calculations were repeated on an hourly basis from 0:00 on January 1 to 24:00 on December 31. In cases with EV deployment, hourly values of VRE output curtailment, electricity discharged from EVs to the power grid, and total electricity demand (including EV charging) were calculated. Subsequently, thermal power generation output was determined to balance electricity demand and supply. Annual results were then compared between cases with and without the V2G system.

LCA for the V2G system evaluation

LCA was conducted to evaluate GHG emissions before and after the installation of EVs. This study targeted electricity supply for the target region as well as vehicle production and operation. The functional unit was set as the annual electricity demand of the region together with the production and operation of vehicles. Vehicle operation schedules were assumed to be identical for EVs and ICEVs. The VRE capacity was assumed to be the same before and after EV deployment.

Foreground data, such as the amount of thermal power generation, were obtained from the energy flow simulations. Data related to vehicle production were derived from a previous study [4]. Background data, including thermal power generation and fuel consumption of ICEVs, were sourced from the Japanese LCA Inventory Database for Environmental Analysis (IDEA) Version 3.5.1 [5] and ecoinvent v.3.11 [6].

Analysis Settings

This study focused on Tanegashima island, where VRE output was curtailed for the first time in Japan and where VRE output curtailment frequently occurs at present. Tanegashima has a population of approximately 27, 000. Several cases were defined to compare the potential of the regional V2G system. **Table 1** presents five cases representing different configurations of EV deployment and charging infrastructure.

Table 1: Case settings.

Case	EV	Home	Office
0	No	-	-
1	Yes	Standard charger	Standard charger
2	Yes	Standard charger	V2B
3	Yes	V2H	Standard charger
4	Yes	V2H	V2B

The installed capacities of VRE were set to be PV 13, 000 kW and WT 700 kW, reflecting the local conditions of Tanegashima [7]. Hourly electricity outputs from PV and WT were estimated using solar radiation and wind speed data. Electricity demand in Tanegashima was estimated by allocating the total electricity demand of Kyushu region [8] in proportion to population. The minimum output of thermal power generation, $P_{\text{min}}^{\text{thermal}}(t)$, was estimated

using a regression equation based on VRE output, $P^{\mathrm{VRE}}(t)$, electricity demand, $P^{\mathrm{use}}(t)$, and reported data on the minimum thermal power output released by electric power company [9].

The target vehicle type was the Japanese Kei car, which is the most commonly used vehicle in Tanegashima [3]. The number of EVs was assumed to increase to up to 10, 000 units. EV battery capacities and EV efficiencies were set based on an industrial data source [10].

RESULTS AND DISCUSSION

Overview of energy flow simulation with V2G

Energy demand and supply were calculated on an hourly basis using the energy flow simulation model. Under current conditions without EV deployment, VRE output curtailment occurred on 128 days per year. In comparison, 118 days of VRE output curtailment were reported in 2023 [9]. This close agreement suggests that the model developed in this study reasonably represents the actual electricity system of the region.

Figure 3 shows the temporal dynamics of electricity generation and demand on April 27 with 1, 000 EVs in case 4. The difference in electricity demand between cases with and without EVs is attributable to EV charging. On this day, VRE output was curtailed during the middle of the day. The installation of EVs reduced VRE output curtailment by EV charging. At that time, electricity demand increased in the morning due to EV charging before commuting to work, which led to increase in thermal power generation to meet this additional demand. In the evening, after returning home, EVs discharged electricity to the power grid.

Annual changes in electricity generation

The simulation shown in **Figure 3** was conducted for one year. Figure 4 shows the annual amounts of electricity demand, thermal power generation, VRE output curtailment, and EV discharge to the power grid. The number of installed EVs was varied from 0 to 10, 000 units.

EV deployment reduced VRE output curtailment (**Figure 4 (a)**). In case 4 with 10, 000-EVs, VRE output curtailment was reduced to 0.26%. The installation of V2H and V2B further decreased VRE output curtailment. EVs were able to charge electricity derived from VRE that would otherwise have been curtailed in case 0, because EV batteries had available capacity created by discharging electricity to the power grid on the previous day.

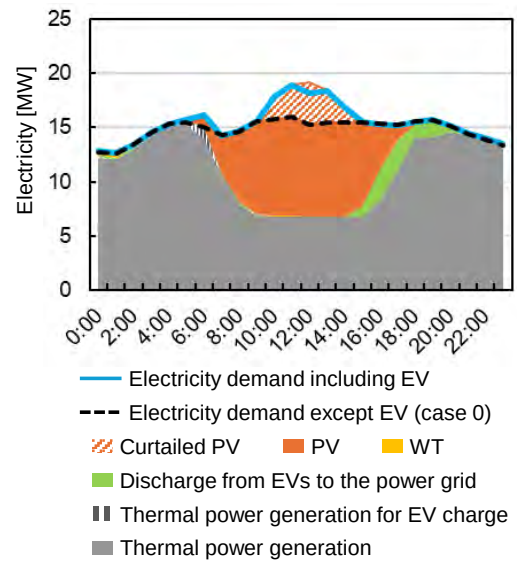

Figure 3. Temporal dynamics of electricity generation and demand on April 27th with 1, 000 EVs in case 4.

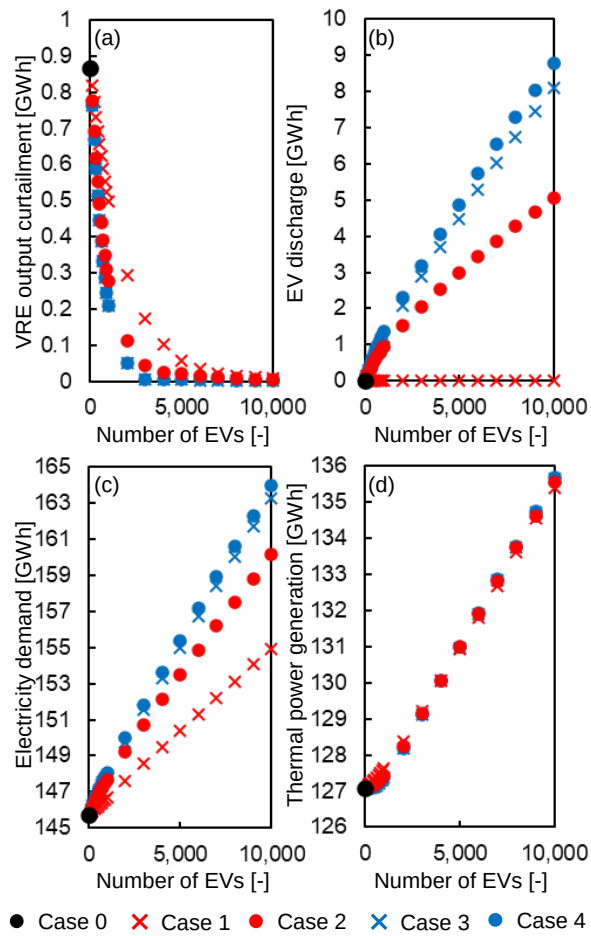

Figure 4. Effects of the regional V2G system. Note that the vertical axis scales differ. (a) VRE output curtailment. (b) EV discharge to the power grid. (c) Electricity demand. (d) Thermal power generation.

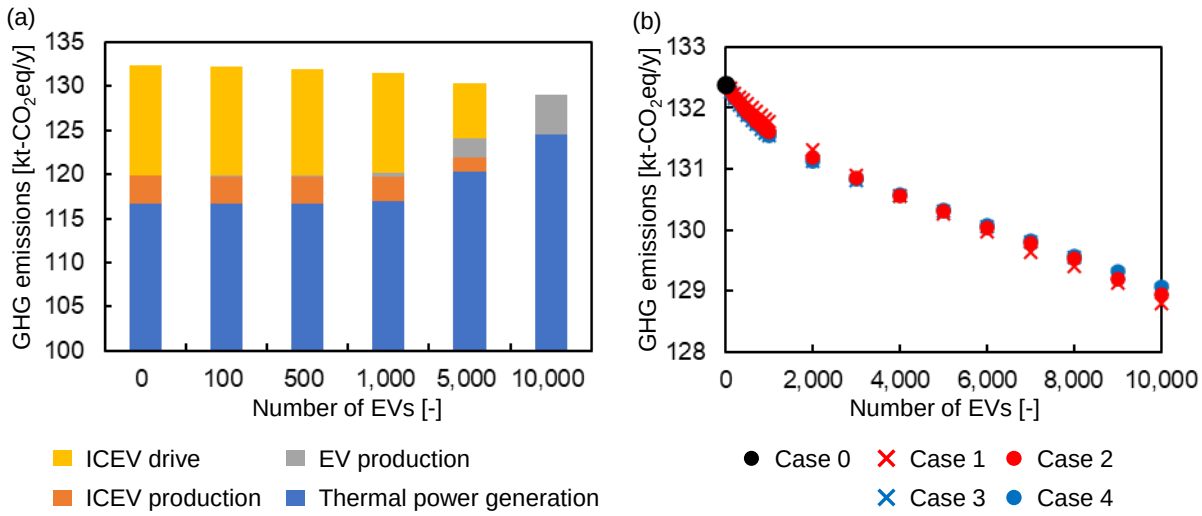

Figure 5. GHG emissions from the regional V2G system. Note that the vertical axis scales differ between (a) and (b). (a) GHG emissions from the regional V2G system in case 4. (b) GHG emissions from the regional V2G system among all cases.

By installing V2H and V2B, electricity stored in EV battery could be utilized at homes and offices (**Figure 4 (b)**). VRE output curtailment typically occurred during midday, so EVs mainly charged during this period and discharged in the evening. A comparison between cases 2 and 3 showed that EV discharge to the power grid was smaller in case 2. If EVs were parked at offices overnight, such as company owned vehicles, EV discharge to the power grid would be expected to increase.

The annual electricity demand in Tanegashima is shown in **Figure 4 (c)**. When EVs were deployed, total electricity demand increased due to electricity consumption for EV driving. In cases 2-4, electricity demand increased more than in case 1. This additional increase was caused by EV discharging to the power grid, which in turn created additional opportunities for EV charging.

Figure 4 (d) presents the annual amount of thermal power generation. In general, thermal power generation increased with EV deployment because EVs required electricity for driving and charging. However, when fewer than 200 EVs were installed in Cases 3 and 4, thermal power generation decreased compared with Case 0. This reduction occurred because EVs charged electricity that would otherwise have been curtailed and discharged electricity to the power grid during evening hours, thereby reducing the need for thermal power generation.

When fewer than 200 EVs were installed, thermal power generation in case 4 was the lowest among all cases (**Figure 4 (d)**). However, when a large number of EVs were deployed, thermal power generation in case 4 became the highest. This inversion phenomenon was observed at approximately 4, 000-EVs. As shown in **Figure 4 (a)**, VRE output curtailment in cases 2-4 was almost zero at this level of EV penetration. Therefore, when the number of EVs exceeds 4, 000, thermal power generation must increase to meet the electricity demand, because the surplus electricity that was curtailed in Case 0 had already been fully absorbed by EV charging. These results indicate that the environmental effectiveness of regional V2G systems strongly depends on the balance between available surplus VRE and additional electricity demand induced by EV deployment. Beyond a certain penetration level, further EV deployment no longer contributes to GHG reduction unless additional renewable capacity or system-level control is introduced.

These results suggested that the regional V2G system could reduce VRE output curtailment and had the potential to decrease thermal power generation by enabling electricity discharge from EVs to the power grid. However, it should be noted that EV operation also tended to increase thermal power generation due to EV driving.

GHG emissions of regional V2G system

GHG emissions from the regional V2G system were evaluated using LCA. **Figure 5 (a)** shows the GHG emissions for case 4. As the number of EVs increased, GHG emissions from thermal power generation and EV production also increased. However, total GHG emissions decreased because GHG emissions from ICEV operation were substantially reduced.

Figure 5 (b) compares GHG emissions all cases. When a small number of EVs were installed, case 4 exhibited the lowest GHG emissions among all cases. However, as the number of installed EVs increased, case 1 resulted in the lowest GHG emissions. This inversion phenomenon was observed at approximately 4, 000-EVs,

which is consistent with the trend observed for thermal power generation. At low EV penetration levels, GHG emissions could be reduced effectively because EVs were able to utilize surplus VRE that was curtailed in case 0. In contrast, at high EV penetration levels, electricity demand from EVs exceeded the amount of surplus VRE available, leading to increased reliance on thermal power generation.

GHG emissions associated with EV production were higher than those from ICEV production. Although this study did not include vehicles disposal, recycling could further reduce overall GHG emissions. In particular, recycling batteries used in EVs could reduce the environmental impact through appropriate recycling [11].

Possible role of regional V2G system

In regions distant from major urban centers, particularly remote islands, fuel for ICEV operation must be transformed from remote production facilities, resulting in relatively high fuel costs. If EVs and V2G systems were deployed and locally available VRE was utilized for EV operation, both regional GHG emissions and fuel cost could be reduced. In addition, increased utilization of locally generated energy had the potential to stimulate the local economy.

Furthermore, the installation of regional V2G systems was expected to promote local economic activity and contribute to regional revitalization. The local socio-economic impacts of adopting new technologies can be evaluated using Input-Output (IO) models [12]. The results of this study provided quantitative parameters that can be used to assess these local socio-economic effects.

Beyond reducing life cycle environmental impacts, V2G systems could also enhance energy resilience by mitigating the risk of electricity supply disruptions during natural disasters. In regions equipped with V2G systems, residents could continue to access electricity even in the event of blackouts caused by natural disasters such as typhoons. Therefore, regional V2G systems would be well suited for the development of resilient local energy systems.

Regional V2G systems could reduce VRE output curtailment and GHG emissions and, depending on the number of deployed EVs, may also decrease thermal power generation. In addition, such systems had the potential to stimulate the local economic activity and contributed to the development of resilient local energy systems.

Limitations of this study

The deployment of V2G system at the regional level was examined in this study. It assumed that a regional V2G system centrally controlled the charging and discharging of individual EVs by comprehensively considering the amount of regional VRE output curtailment, the number of EVs connected chargers and V2H/V2B systems in real time. Without such coordinated control, EVs may demand electricity beyond the available surplus VRE, potentially leading to unnecessary increases in thermal power generation. Under current conditions, individual V2H/V2B systems cannot access regional electricity system information in real time. Therefore, an advanced control system would be required to utilize regional V2G systems effectively.

Even though the model incorporated real-time information on regional electricity system, some EVs discharge electricity to the power grid that had been charged into their batteries during periods without VRE output curtailment. This behavior increased the load on thermal power generation.

The regional V2G system model developed in this study focused exclusively on Japanese Kei cars used for commuting purposes and did not account for differences in vehicle types or usage patterns. All EV owners and their workplaces were assumed to install identical charging equipment, such as standard chargers and V2H systems. Furthermore, the model could be enhanced by incorporating material flow analysis (MFA) to account for the timing of EV deployment, vehicle lifetimes including battery degradation, as well as variations in vehicle types, charging infrastructure, and usage patterns.

To enable the practical deployment of regional V2G systems, an economic framework must also be established. If EVs were allowed to charge and discharge freely at any location, individual EV owners could buy and sell electricity anywhere, leading to complex financial transactions. The development of a dedicated financial system, such as applications or platforms for V2G users, would help facilitate and accelerate the deployment of V2G systems.

For these reasons, regional V2G systems should manage regional electricity demand and supply by accounting for real-time system conditions, while coordinating the individual operation of V2H and V2B systems. Furthermore, the load of thermal power generation could be reduced by distinguishing whether the electricity stored in EV batteries was derived from VRE or from thermal power generation. Such an approach would promote more effective utilization of VRE. Therefore, integrated regional systems capable of controlling electricity flows and managing financial transactions were essential for the successful implementation of regional V2G systems.

CONCLUSIONS

This study developed a regional V2G system model and evaluated the potential of EVs and V2G systems using energy flow simulations and LCA. Fuel consumption, VRE utilization, and GHG emissions were quantitatively assessed to examine the potential of regional V2G

system.

The results revealed that the deployment of EVs and V2G systems could effectively reduce VRE output curtailment and GHG emissions. In the case of 10, 000-EVs equipped with V2H and V2B systems, VRE output curtailment was reduced to 0.26%. Moreover, when fewer than 200 EVs were installed with V2H, thermal power generation decreased compared with the case without EVs. These findings indicated that regional V2G systems had the potential to reduce VRE output curtailment, GHG emissions, and, under certain conditions, thermal power generation.

The model developed in this study enabled a quantitative evaluation of the potential of regional V2G systems. For the practical deployment of such systems, integrated regional platforms were required to coordinate electricity flows and manage financial transactions. By evaluating the environmental performance of regional V2G systems, this study provided valuable insights to support the design and the implementation of environmentally sustainable V2G systems.

ACKNOWLEDGEMENTS

This research contained the achievements of JST COI-NEXT (JPMJPF2003). Interlinkages and Innovation for Future Societies Research Unit at the University of Tokyo is supported by the Toyota Foundation. Activities of the Presidential Endowed Chair for “Platinum Society” at the University of Tokyo are supported by Mitsui Fudosan Corporation, Sekisui House, Ltd., the East Japan Railway Company, and Toyota Tsusho Corporation.

AUTHOR IDENTIFIERS

Author ORCIDs:
Yamaki A: 0000-0003-4162-7806
Kikuchi Y: 0000-0003-0002-8233

REFERENCES

1. Kabatepe B, Türkay M. A bi-criteria optimization model to analyze the impacts of electric vehicles on costs and emissions. Computers & Chemical Engineering 102:156-168 (2017). https://doi.org/10.1016/j.compchemeng.2016.11.026
2. Yamaki A, Fujii S, Kanematsu Y, Kikuchi Y. Life cycle greenhouse gas emissions of cogeneration energy hubs at japanese paper mills with thermal energy storage. Energy 270:126886 (2023). https://doi.org/10.1016/j.energy.2023.126886
3. Igarashi K, Kurishima H, Kikuchi Y. Evaluation of Vehicle-to-Grid System Using Renewable Energy in Tanegashima. *Papers on Environmental Information Science*. 36:87–92 (2022). https://doi.org/10.11492/ceispapers.ceis36.0_87
4. Hawkins TR, Singh B, Majeau?Bettez G, Strømman AH. Comparative environmental life cycle assessment of conventional and electric vehicles. J of Industrial Ecology 17:53-64 (2012). https://doi.org/10.1111/j.1530-9290.2012.00532.x
5. National Institute of Advanced Industrial Science and Technology. AIST-IDEA (National Institute of Advanced Industrial Science and Technology-Inventory Database for Environmental Analysis) v3.5.1. (2025).
6. ecoinvent. ecoinvent v3.11. (2025).
7. Organization for Cross-regional Coordination of Transmission Operators, JAPAN. https://www.occto.or.jp/assets/oshirase/shutsuryokuyokusei/2023/files/230427_kenshokekka.pdf
8. Kyushu Electric Power Co., Inc. https://www.kyuden.co.jp/td_power_usages/history202301.html
9. Kyushu Electric Power Co., Inc. https://www.kyuden.co.jp/td_power_usages/out_ctrl_history.html
10. Nissan Motor Co., Ltd, . https://www.nissan.co.jp/INFO/E_NOTE/SAKURA/index.html?utm_source=chatgpt.com
11. Kikuchi Y, Suwa I, Heiho A, Dou Y, Lim S, Namihira T, Mochidzuki K, Koita T, Tokoro C. Separation of cathode particles and aluminum current foil in lithium-ion battery by high-voltage pulsed discharge part II: prospective life cycle assessment based on experimental data. Waste Management 132:86-95 (2021). https://doi.org/10.1016/j.wasman.2021.07.016
12. Yuko O, Aya H, Kotaro O, Yuichiro K, Yasuhiro F, Yasunori K. Analyzing socio-economic effect induced by technology implementation on available renewables: a case study of tanegashima, japan. Journal of Life Cycle Assessment, Japan 15:360-376 (2019). https://doi.org/10.3370/lca.15.360

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Techno-economic assessment of green ammonia plants with multi-scale capacity

Ruitao Sun[a], Jie Li[a]*

[a] Centre for Process Integration, Department of Chemical Engineering, The University of Manchester, Manchester, UK
* Corresponding Author: jie.li-2@manchester.ac.uk.

ABSTRACT

Cost reduction of green ammonia production is critical to advancing the hydrogen-ammonia economy, as ammonia capable of cost-effective storage and transportation is a promising hydrogen carrier and energy carrier to alleviate the intermittency and geographical limitations of renewables. Optimisation and techno-economic assessment based on rigorous model are essential to accurately investigate techno-economic feasibility and fully explore optimisation potential. This work estimates the levelized cost of ammonia (LCOA) of an integrated system including a hydrogen generation process employing Proton Exchange Membrane (PEM) water electrolysis, a nitrogen generation process from flue gas recovery, and an ammonia synthesis process based on Haber-Bosch Process. To enhance the reliability of LCOA, detailed equipment sizing and costing is conducted according to stream data from rigorous modelling in Aspen Plus. A novel optimisation strategy is proposed to enhance the computational robustness by sequentially largening the group of optimised variables and improving the quality of initial points. Homotopy-continuation (HC) method is employed to ease the convergence difficulties for solving the flowsheet with multiple recycle streams and scaling up the plant capacities between 100 $tNH_3\ d^{-1}$ and 600 $tNH_3\ d^{-1}$ from an initially feasible point. Heat integration is also considered to realise the combined supply of heat and power by introducing a steam turbine power cycle and utilising the waste heat from exothermic ammonia synthesis.

Keywords: Green ammonia, Rigorous modelling, Process optimisation, Plant scale up

INTRODUCTION

Cost reduction of green ammonia production is critical to advancing the hydrogen-ammonia economy, as ammonia capable of cost-effective storage and transportation is a promising hydrogen carrier and energy carrier to alleviate the intermittency and geographical limitations of renewables. Currently, most current studies estimate the levelized cost of ammonia (LCOA) using oversimplified assumptions [1 – 3] or shortcut models [4, 5], which is inaccurate and unrealistic. In details, linear correlation is inappropriately used to estimate the LCOA in green ammonia systems with all plant scales, as it is derived in a basis of 300 $tNH_3\ d^{-1}$ using 70 USD MWh^{-1} of electricity [5]. Despite shortcut models could provide feasible initialisation for process simulation and optimisation, a technically feasible solution cannot be guaranteed when it is solely based on shortcut modelling, not to mention a reliable cost estimation. Moreover, optimisation scope is restricted to general terms such as equipment capacity and wind-solar mix ratio [1]. Therefore, optimisation and techno-economic assessment based on rigorous model are essential to accurately investigate techno-economic feasibility and fully explore optimisation potential.

Rigorous model of ammonia synthesis employing Neilsen's expression [6] is widely regarded to have the most accurate prediction of ammonia synthesis performance [7], whilst the application of it always focuses on single equipment simulation and cost assessment [8]. Exothermic nature in ammonia synthesis can be well captured by rigorous modelling, and utilisation of this waste heat can further reduce the LCOA to enhance the economic competitiveness.

In this work, detailed equipment sizing and costing are conducted according to stream data from rigorous modelling and in Aspen Plus. A novel optimisation

strategy sequentially largening the group of optimised variables and improving the quality of initial points will be employed to optimise the LCOA for plants between 100 $tNH_3\ d^{-1}$ and 600 $tNH_3\ d^{-1}$ from an initially feasible point.

METHODOLOGY

System description

An integrated system utilising off gas as the nitrogen source for green ammonia production was proposed, as depicted in **Figure 1** [10]. It encompasses an ammonia production system based on the Haber-Bosch process applying an industrial ammonia synthesis convertor (i.e., Topsoe S-200), an electrolysis-based hydrogen production process, and a stepwise nitrogen production technology where a pressure swing adsorption unit is facilitated by an amine-based carbon capture process for nitrogen enrichment. Apart from electric power driven by renewables such as solar and wind, this system also explores heat integration opportunities in ammonia purification process for further energy efficiency improvement and carbon mitigation.

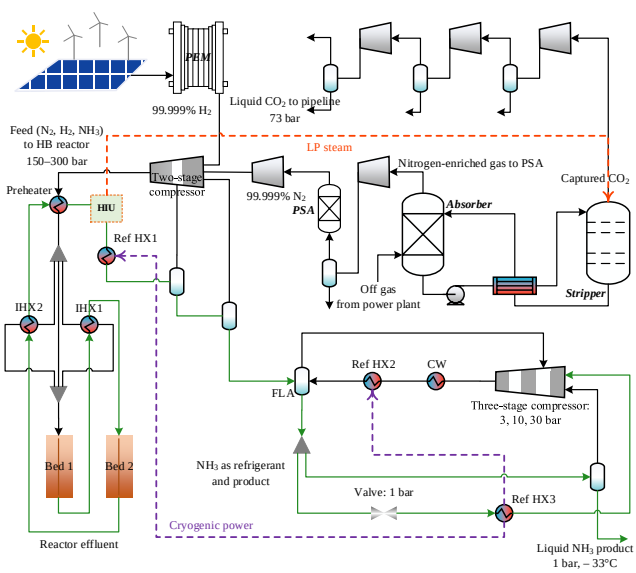

Figure 1. Flowsheet of the green ammonia production system.

Optimisation models

Modelling of ammonia synthesis

The modelling is conducted in Aspen Plus V12.1. The ammonia synthesis convertor (e.g., Topsoe S-200) is modelled by the combination of two internal heat exchangers (modelled by HEATX) and two plug flow reactors (PFR) as the catalytic beds.

The kinetic model of ammonia synthesis employs widely validated Neilsen's expression [6] throughout all sections (e.g. process simulation, optimisation and techno-economic analysis) as follows:

$$r_{NH_3} = \frac{\eta \cdot \kappa_2 \cdot \left(\alpha_{N_2} \cdot K_{eq}^2 - \alpha_{NH_3}^2 \cdot \alpha_{H_2}^{-3}\right)}{0.654 \cdot \left(1 + K_{a,eq} \cdot \alpha_{NH_3} \cdot \alpha_{H_2}^{-1.523}\right)^{1.308}} \cdot V \quad (1)$$

$$\kappa_2 = 7.89 \cdot 10^{10} \cdot e^{-\frac{11171}{RT}} \quad (2)$$

$$K_{a,eq} = -2.94 \cdot 10^{-4} \cdot e^{-\frac{24052}{RT}} \quad (3)$$

$$\log_{10} K_{eq} = 2.69 \cdot (1 - \log_{10} T) - 10^{-7} \cdot (552 \cdot T - 1.85 \cdot T^2) + \frac{2001.6}{T} \quad (4)$$

Objective function

The primary objective of this study is to evaluate and minimise the LCOA, defined in Eq. 5, which represents the average production cost per unit of ammonia (USD tNH_3^{-1}). The LCOA framework integrates both capital and operating expenditures over the lifetime of the plant, thereby providing a consistent economic metric for comparing alternative design (e.g., various plant scales and heat-integration strategies) and operating strategies (electricity prices and optimised operating conditions).

$$LCOA = \frac{ACCR \cdot CAPEX + OPEX}{m_{NH3}} \quad (5)$$

$$C_{TPI} = 1.18 \cdot \left(C_{TBM} + C_{site} + C_{buildings} + C_{offsite}\right) \quad (6)$$

$$C_{BM} = \frac{I}{I_0} \cdot C_{P0} \cdot \left[F_{BM} + \left(F_d \cdot F_p \cdot F_m - 1\right)\right] \quad (7)$$

$$C_{P0} = F \cdot exp\{A_0 + A_1 \cdot (ln\, S) + A_2 \cdot (ln\, S)^2\} \quad (8)$$

The capital component of the LCOA consists of the annualised investment cost associated with major process equipment. Equipment capital costs (Eq. 6–8) are derived from detailed sizing calculations based on rigorous process simulation. Specifically, heat-exchanger costs are correlated to the required heat-transfer area, vessel costs are evaluated as a function of material of construction and shell weight, and compression or expansion equipment costs are determined from shaft power requirements. All equipment sizes are obtained case-by-case from steady-state stream data generated under each candidate design or operating condition. More detailed sizing and costing equations and associated coefficients can be found in the Supplementary Materials.

Operating expenditure includes the cost of raw materials, utilities, and ancillary services necessary for plant operation as shown in Eq. 9. Key contributors are electricity consumption for compression and refrigeration, hydrogen and nitrogen feedstock costs, cooling-water demand, raw materials (e.g., MEA and water) and carbon tax. Variable operating costs are calculated based on process simulation outputs and unit prices, while fixed operating costs such as labour, insurance, and maintenance are incorporated as a fraction of the installed capital cost following standard chemical-engineering

economic practice. Necessary equations and cost assumptions can be found in the Supplementary Materials.

$$OPEX = C_{RM} + C_{utility} + C_{operation} + C_{maintenance} + C_{overhead} + C_{depreciation} + C_{carbon} \quad (9)$$

The optimisation problem is expressed as:

$Min\ LCOA$

$s.t.$

Mass and energy balance

Thermodynamic constraints: RKS-BM

Chemical equilibrium: Correlation of Gillespie and Beattie (Eq. 4)

Reaction kinetics: Neilson's expression (Eq. 1)

Unit models: PFR, Compressor, Heat exchanger, etc.

Boundary conditions: See **Table 1**

Table 1: Boundary conditions of decision variables for the optimisation problem.

Variables	Unit	Lower bound	Upper bound
PCOMP1	Bar	31	149
PCOMP2	Bar	150	300
PFLA	Bar	3	10
TPREHEAT	°C	140	220
TREFRIG-ERATION	°C	−25	0
TBED1OUT	°C	400	600
TBED2IN	°C	350	500

Solution strategy

Once a feasible scaled-up flowsheet is established, optimisation is performed using the Large-Scale Sparse Successive Quadratic Programming (LSSQP) algorithm. LSSQP is a gradient-based method widely adopted in the chemical process industry, noted for its fast convergence and robustness in solving large-scale nonlinear programming problems [11]. In this study, the derivative information automatically generated in EO mode is incorporated into the optimisation routine, enhancing computational efficiency and solution reliability.

To outlines the transition from an initial feasible point obtained in SM to the EO-based optimisation workflow, the overall procedure is illustrated in **Figure 2** and explained in the following:

Step 1: Generate an initial feasible solution in SM and synchronise the model in EO mode. In this step, all equations from unit models and thermodynamic property models are assembled, and the associated free variables are initialised.

Step 2: Scale up the plant by fixing the required product rate, freeing feed flowrates, scaling up the reactor volume, and initially fixing reactor outlet/inlet temperatures while freeing the internal heat exchanger areas. This ensures that the overall degrees of freedom are balanced (DOF = 0).

Step 3: Perform EO simulation using the LSSQP solver, with the Homotopy-Continuation (HC) method enabled to enhance convergence. If convergence fails, the HC step size is reduced and the simulation is repeated.

Step 4: Switch the selected constant variables to optimisation variables.

Step 5: Apply LSSQP optimisation. If convergence is successful, the optimised variable selections and flowsheet solutions are restored before expanding the set of optimised variables. If convergence fails, the optimisation set is reformulated based on the last feasible selection. After all possible combinations of a given number of optimisation variables have been tested unsuccessfully, the number of active optimisation variables is reduced.

Step 6: Repeat the optimisation–reset cycle until convergence is achieved across the full set of optimisation variables.

Step 7: Once convergence is achieved for all optimisation variables, the global optimal solution is obtained.

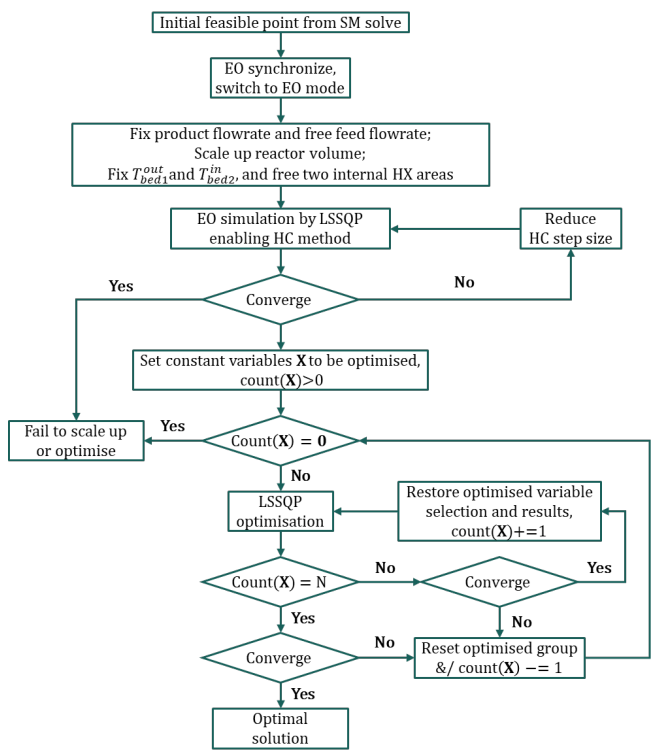

Figure 2. Solution strategy of the optimisation in Aspen Plus Equation-Oriented mode.

RESULTS

Distribution of optimized variables

As the cost of green ammonia production is highly dependent on hydrogen, which in turn originates from green electricity, the assumption regarding electricity price is critical. The International Renewable Energy Agency (IRENA) estimates that the cost of electrolysis-based hydrogen corresponds to electricity prices ranging from 20 USD MWh^{-1} to 65 USD MWh^{-1} over the period 2020–2050. For the optimisation of the base case, an electricity price of 20 USD MWh^{-1} is assumed.

Figure 3 presents the distribution of optimised decision variables for plant capacities of 100–600 tNH_3 d^{-1}. The box plots indicate the 30–70% percentile range, median, and full spread of solutions, revealing which variables consistently converge to well-defined optima and which remain flexible.

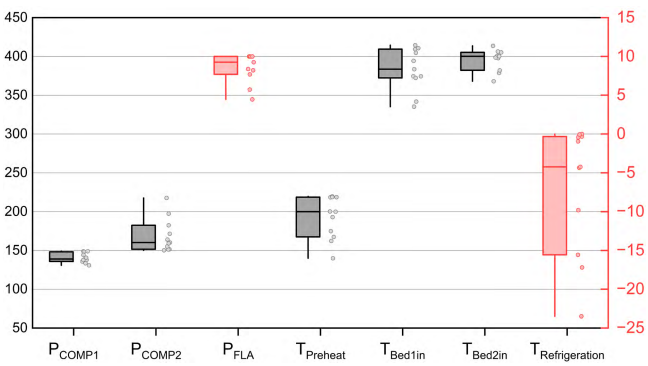

Figure 3. Distribution results of optimised variables

The discharge pressure of the high-pressure compressor (PCOMP2), corresponding to the ammonia synthesis loop pressure, consistently converges to 160–225 bar. This narrow range reflects the core economic trade-off in the Haber–Bosch process: although higher pressure enhances equilibrium conversion, the benefit is offset by higher compression energy demand and increased equipment cost. The medium-pressure compressor (PCOMP1) discharge pressure similarly clusters near its upper bound (ca. 150 bar), indicating a clearly defined optimal compression level to meet synthesis loop requirements.

The inlet temperatures of the catalyst beds (Tbed1 and Tbed2) also show tight distributions around 400 °C. This demonstrates that reactor thermodynamics dominate the optimal solution. The selected temperatures represent the classical compromise between improved reaction kinetics at high temperature and reduced equilibrium conversion.

In contrast, variables associated with downstream separation exhibit wide distributions. The flash pressure (PFLA) spans almost its full allowable range, and the refrigeration temperature varies broadly, including sub-zero conditions. These variables have only a minor impact on LCOA, leading to a relatively flat objective landscape and multiple economically equivalent operating points. The preheat temperature also shows moderate flexibility, as its effect can be compensated by internal heat exchange before the reactor.

Optimized LCOA in various plant scales

Figure 4 shows the dependence of LCOA on plant scale (100–600 tNH_3 d^{-1}) and electricity price (0–120 USD MWh^{-1}). The levelized cost of hydrogen (LCOH) is used as the hydrogen feedstock cost in the LCOA calculation. Plant scale affects both capital-related costs and raw material consumption, while electricity price directly influences electrolysis, compression energy demand, and indirectly the hydrogen production cost.

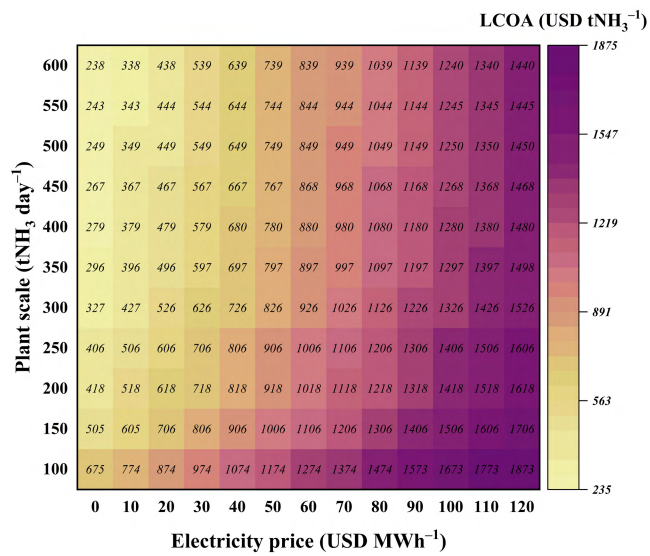

Figure 4. Optimised results, scale-up curve of the ammonia synthesis loop.

Electricity price is the dominant factor controlling LCOA across all plant sizes. Under zero-cost electricity, LCOA ranges from approximately 238 to 675 USD tNH_3^{-1} depending on scale. When electricity price increases to 120 USD MWh^{-1}, LCOA rises sharply to about 1440–1873 USD tNH_3^{-1}.

Plant scale has a secondary but still noticeable effect. Larger plants consistently achieve lower LCOA because capital-intensive equipment, such as reactors and compressors, is amortised over higher annual ammonia output. For instance, at 20 USD MWh^{-1} electricity, increasing capacity from 100 to 600 tNH_3 d^{-1} reduces LCOA from about 874 to 438 USD tNH_3^{-1}. At this electricity price, the corresponding LCOH is comparable to conventional hydrogen costs, and 600 tNH_3 d^{-1} represents a typical industrial ammonia plant scale. The resulting LCOA (ca. USD tNH_3^{-1}) is close to the current market price of ammonia (ca. 410 USD tNH_3^{-1}), supporting the credibility

of the economic model.

However, even at the largest scale, electricity prices above 40 USD MWh^{-1} drive LCOA well beyond commonly cited competitiveness ranges for green ammonia (400–600 USD tNH_3^{-1}). This indicates that economies of scale alone cannot compensate for high electricity costs, and that access to inexpensive renewable electricity remains the key requirement for economically viable green ammonia production.

To clarify how plant scale influences both the magnitude and structure of ammonia production cost, **Figure 5** and **Figure 6** decompose the LCOA into annualised capital cost (ACC), cost of manufacturing (COM), raw material, and utility costs. An electricity price of 20 USD MWh^{-1} is assumed for all plant capacities.

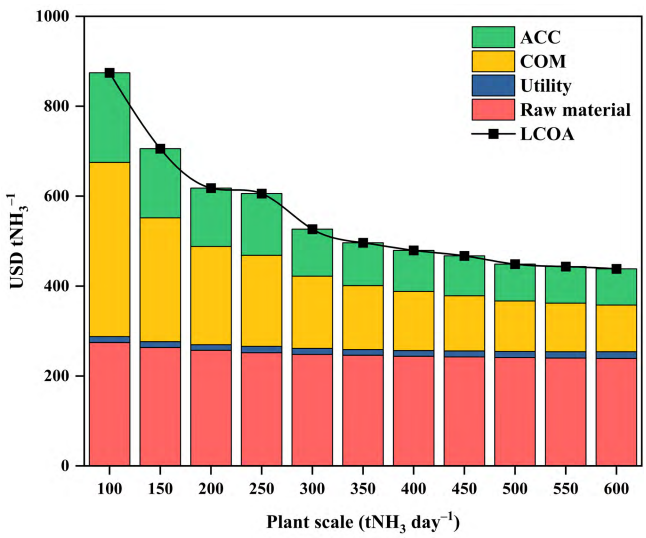

Figure 5. Optimised results, scale-up curve of the ammonia synthesis loop.

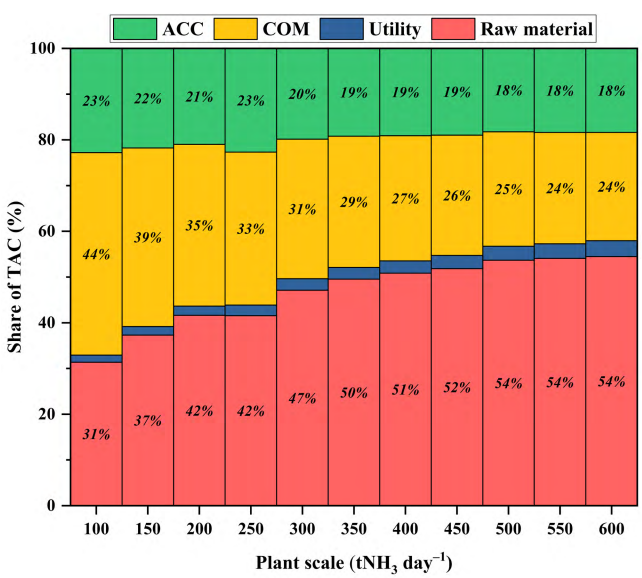

Figure 6. Breakdown of TAC.

Figure 5 shows a pronounced economy-of-scale effect, with total LCOA decreasing significantly as plant capacity increases from 100 to 600 $tNH_3\ d^{-1}$. At the smallest scale (100 $tNH_3\ d^{-1}$), LCOA is close to 900 USD tNH_3^{-1}, while at 600 $tNH_3\ d^{-1}$ it drops to about 440 USD tNH_3^{-1}. The reduction is steep between 100 and 300 $tNH_3\ d^{-1}$ and becomes more gradual beyond 400 $tNH_3\ d^{-1}$, indicating diminishing marginal benefits from further scale-up.

The stacked cost breakdown indicates that the decline in LCOA is mainly driven by reductions in ACC and COM. Both components decrease with scale because capital-intensive equipment is utilised more effectively and fixed costs are distributed over higher production rates. In particular, depreciation, maintenance, and equipment-related operating expenses are key contributors to COM. In contrast, raw material cost changes only slightly in absolute terms, and utility cost remains relatively small and nearly constant across plant sizes. Consequently, small-scale plants are dominated by capital and manufacturing burdens, whereas at large scale the cost structure becomes increasingly governed by raw material inputs.

Figure 6 presents the cost structure in percentage terms, showing the contribution of each component to the total annualised cost (TAC). A clear structural transition occurs as plant scale increases. At 100 $tNH_3\ d^{-1}$, COM is the largest contributor (44%), followed by raw materials (31%) and ACC (23%). With increasing capacity, the COM share steadily declines to about 24% at 600 $tNH_3\ d^{-1}$, while the ACC fraction decreases more moderately from about 23% to 18%.

In contrast, the raw material share increases significantly, rising from roughly 31% at small scale to more than 50% at large scale, eventually becoming the dominant cost component. The utility share remains small across all scales and increases only slightly. This behaviour is consistent with the optimisation results, where operating conditions tend to minimise compression work while maintaining relatively high ammonia synthesis conversion, thereby limiting additional utility penalties at larger scales.

CONCLUSION

This work estimates the levelized cost of ammonia (LCOA) of an integrated system including a hydrogen generation process employing Proton Exchange Membrane (PEM) water electrolysis, a nitrogen generation process from flue gas recovery, and an ammonia synthesis process based on Haber-Bosch Process. In particular, CO2 in flue gas is removed by Monoethanolamine (MEA)-based carbon capture process and nitrogen is enriched to 90%, contributing to a higher efficiency and cost reduction of nitrogen purification, compared with considering air as the nitrogen source with 79% of nitrogen. An

industrial ammonia synthesis convertor (e.g., Topsoe S-200), consisting of two reactor beds and two internal heat exchangers, is employed to make our work closer to a real-life project. To enhance the reliability of LCOA, detailed equipment sizing and costing is conducted according to stream data from rigorous modelling in Aspen Plus. A novel optimisation strategy is proposed to enhance the computational robustness by sequentially largening the group of optimised variables and improving the quality of initial points. Homotopy-continuation (HC) method is employed to ease the convergence difficulties for solving the flowsheet with multiple recycle streams and scaling up the plant capacities between 100 $tNH_3\ d^{-1}$ and 600 $tNH_3\ d^{-1}$ from an initially feasible point.

Overall, the optimisation results reveal a clear hierarchy in the factors governing the techno-economic performance of green ammonia production. Reactor operating conditions, particularly synthesis pressure and catalyst bed inlet temperature, consistently converge to narrow ranges, demonstrating that reactor thermodynamics and compression trade-offs dominate the optimal process configuration. In contrast, separation and auxiliary thermal variables show wide feasible ranges, indicating a relatively weak influence on the LCOA. From an economic perspective, electricity price is the most critical external parameter, exerting a nearly linear and dominant impact on LCOA across all plant scales due to the electricity-intensive nature of hydrogen production and compression. Plant scale introduces a strong but secondary economy-of-scale effect, substantially reducing LCOA by diluting capital and manufacturing costs, especially between 100 and 300 $tNH_3\ d^{-1}$. However, as capacity increases, the capital- and manufacturing-related expenses become less dominant, while raw material costs (primarily linked to electricity for hydrogen production) emerge as the controlling factor. These results indicate that, beyond industrial-scale capacities, further cost reductions depend more on access to low-cost renewable electricity and improvements in process efficiency than on additional scale-up, while fixing strongly converged reactor operating conditions can enhance optimisation robustness with minimal economic penalty.

ACKNOWLEDGEMENTS

Jie Li would like to appreciate the financial support from UKRI Impact Acceleration Account (IAA 424).

REFERENCES

1. Nayak-Luke RM, Bañares-Alcántara R. Techno-economic viability of islanded green ammonia as a carbon-free energy vector and as a substitute for conventional production. Energy Environ. Sci. 13:2957-2966 (2020). https://doi.org/10.1039/d0ee01707h
2. Grinberg?Dana A, Elishav O, Bardow A, Shter GE, Grader GS. Nitrogen?based fuels: a power?to?fuel?to?power analysis. Angew Chem Int Ed 55:8798-8805 (2016). https://doi.org/10.1002/anie.201510618
3. Ikäheimo J, Kiviluoma J, Weiss R, Holttinen H. Power-to-ammonia in future north european 100 % renewable power and heat system. International Journal of Hydrogen Energy 43:17295-17308 (2018). https://doi.org/10.1016/j.ijhydene.2018.06.121
4. Morgan E, Manwell J, McGowan J. Wind-powered ammonia fuel production for remote islands: a case study. Renewable Energy 72:51-61 (2014). https://doi.org/10.1016/j.renene.2014.06.034
5. Eric R. Morgan. (2013). Techno-economic feasibility study of ammonia plants powered by offshore wind Open Access Dissertations. 697.
6. NIELSEN A. Rate equation and mechanism of ammonia synthesis at industrial conditions. Journal of Catalysis 3:68-79 (1964). https://doi.org/10.1016/0021-9517(64)90094-6
7. Buzzi Ferraris G, Donati G, Rejna F, Carrà S. An investigation on kinetic models for ammonia synthesis. Chemical Engineering Science 29:1621-1627 (1974). https://doi.org/10.1016/0009-2509(74)87013-2
8. Burhan Kabir Suhan M, Naimur Rahman Hemal M, Shoukat Choudhury MAA, Ali Akkas Mazumder M, Islam M. Optimal design of ammonia synthesis reactor for a process industry. Journal of King Saud University - Engineering Sciences 34:23-30 (2022). https://doi.org/10.1016/j.jksues.2020.08.004
9. International Renewable Energy Agency. Green hydrogen cost reduction 2020
10. Sun R, Sun L, Li J. Techno-economic evaluation of green ammonia synthesis for renewable energy storage using rigorous models. Energy 336:138068 (2025). https://doi.org/10.1016/j.energy.2025.138068
11. Ma Y, Gao X, Liu C, Li J. Improved SQP and SLSQP algorithms for feasible path-based process optimisation. Computers & Chemical Engineering 188:108751 (2024). https://doi.org/10.1016/j.compchemeng.2024.108751

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Terawatts for Petabytes: Exploring the impact of AI data centres on Europe's net zero goals

Mohammad Hemmati, Vassilis M. Charitopoulos*

Department of Chemical Engineering, Sargent Centre for Process Systems Engineering, University College London (UCL), Torrington Place, WC1E 7JE, London, United Kingdom
* Corresponding Author: v.charitopoulos@ucl.ac.uk.

ABSTRACT

The unprecedented expansion of Artificial Intelligence is adding increasing electricity demand to Europe's power system. While incumbent plans pursue a net-zero future by 2050, they fail to consider the implications of large-scale AI-based data centres. In this study, a spatially explicit optimisation model is developed to assess how hyperscale data centres may reshape energy infrastructure investment, and emissions trajectories, across different AI demand growth scenarios. The results indicate that, after 2030, AI capacity deployment increasingly shifts toward regions with the ability to expand nuclear and gas-based generation, as firm and flexible power sources are essential for supporting the deployment of high-capacity AI data centres. By 2050, AI-driven electricity demand under high growth scenarios may reach up to 450 TWh, corresponding to 7% of total Europe's demand, with installed AI capacity reaching approximately 85 GW. This additional load leads to an increase of nearly 25 $MtCO_2$ in cumulative emissions between 2030 and 2050. Our analysis indicates that, depending on the AI growth scenario, meeting AI-related electricity demand by 2050 requires between 37 and 323 GW of additional capacity across Europe, ranging from the pessimistic to the lift-off scenario, including investments in nuclear (2-12 GW), gas (2-7 GW), wind (13-100 GW), solar (20-134 GW), and battery storage (0-70 GW).

Keywords: Data Centres, Artificial Intelligence, Net-Zero, Energy Systems, Capacity Expansion Planning, Sustainability

INTRODUCTION

Growing interest in generative artificial intelligence (GenAI) has triggered intense competition to advance the technology, driving the need for numerous high-intensity data centres, as well as greater electricity consumption [1]. To better understand the energy consumption of AI-based data centres, it is essential to consider the wide variation in electricity usage across different AI workloads. For example, a 100-word ChatGPT-4 response consumes approximately 0.14 kWh of electricity [2]. Based on OpenAI's reported figure of roughly 285 million daily ChatGPT users, the platform's energy consumption can be estimated at about 0.39 TWh per day, or roughly 14.4 TWh per year [3]. In addition to inference-related demand, model training represents a non-negligible source of electricity consumption. For example, the training of ChatGPT's 4 model is estimated to have required 50 GWh of electricity, equivalent to the output of a French nuclear reactor operating at full capacity for about two days, with further retraining cycles implying recurring energy requirements [4].

By end of 2024, data centres account for approximately 1.5% of global annual electricity consumption [5]. Looking at the Europe, the International Energy Agency (IEA) estimates that the current energy use of data centres, at around 70 TWh in 2024, will increase to approximately 115 TWh by 2030 [6]. McKinsey indicates that data centre electricity demand in Europe is expected to reach 150 TWh by 2030 [7], while the Independent Commodity Intelligence Services (ICIS) forecasts it will rise to 168 TWh by 2030 and further increase to 236 TWh by 2035 [8]. Although more precise forecasts for data centre demand are rare and prediction rates vary widely across the literature, nevertheless even under conservative scenario, the resulting demand could absorb as much as 20% of newly added renewable energy resource (RES) capacity in Europe and lead to emissions of

approximately 39 million tons of CO_2 annually, equivalent to nearly one-third of Italy's emissions in 2024 [9]. Consequently, the increasing demand of these infrastructures raises significant concerns regarding power system reliability and the Europe's Net Zero 2050 objectives.

Although only a limited number of studies have forecasted future data centre energy consumption and capacity in Europe up to 2030 [10, 11], no comprehensive approach has yet examined the impact of data centres on the future European energy system and national or continental development plans. Spatially explicit planning models such as PyPSA [12], and the ENTSO-E Ten-Year Network Development Plan (TYNDP) [13] do not incorporate emerging forms of geographically concentrated, high-density demand, such as AI data centres. To date, no planning-oriented framework has quantitatively evaluated how AI-driven electricity demand may influence Europe's grid expansion, infrastructure needs, or energy transition. Given high uncertainty in future AI trends, beyond even IEA forecasts, assessing Europe's current position, identifying regions suitable for hyperscale data centres, and understanding their potential impact on net-zero 2050 goals is essential [14, 15].

Motivated by the gaps identified above, this study develops a long-term expansion planning model to systematically evaluate how growing demand from AI may reshape Europe's power system development. A spatially explicit optimisation model is proposed, covering 33 countries and formulated as a Mixed-Integer Linear Program (MILP). The model incorporates current infrastructure plans, including 177 cross-border transmission projects, operational constraints, land-use restrictions for renewables, and climate targets, with an hourly temporal resolution. To reflect incumbent "myopic" planning under high uncertainty from data centres, a rolling-horizon solution approach is employed.

The main contributions of this work are as follows:

- Developing a European-scale modelling framework to evaluate whether existing power-sector development plans can accommodate AI-driven electricity demand growth.
- Investigating the effects of AI-driven demand on the spatial evolution of generation, transmission, and cross-border electricity flows across Europe.
- Identifying the optimal capacity of hyperscale data centres and the country's most suitable for their deployment under AI growth scenarios.
- Identifying the main trends in data centre-driven electricity demand and energy system development under various scenarios.

The remainder of this paper is organized as follows. First, we present the research methodology and problem description and discusses the main assumptions. Then we report and analyse the main results for the benchmark case and the alternative AI-driven demand scenarios on various elements of the future European energy system. Finally, concluding remarks are drawn.

METHODOLOGY

Problem Description

Europe's long-term energy strategy focuses on establishing a net-zero system through a deep transformation of the energy sector, driven by large-scale deployment of renewable, low-carbon, and storage technologies. Analyzing such a complex integrated system to obtain optimal decisions across the planning horizon requires a comprehensive spatial and temporal optimisation framework [16-18]. We employ a MILP model to formulate the capacity and transmission expansion planning problem for integrated European power system from 2025 to 2050. Subject to operational, technical, and climate policy restrictions [19], the model minimises the total system cost, our model includes investment in generation and transmission infrastructure, fixed and variable operating costs [20], emissions costs [21], value of loss of load (VOLL) costs, as ell as renewable curtailment costs. The overall form of the objective function is provided in Eq. (1).

$$TSC = \sum_t FCC_t + FOC_t + FFC_t + CC_t + VOLL_t + Curt_t + TIC \tag{1}$$

Where TSC is the total system cost over the planning horizon, FCC_t denotes facility investment cost, FOC_t denotes variable operating cost, FFC_t is fixed operating cost, CC_t is emissions cost, $VOLL_t$ is load shedding cost, $Curt_t$ is cost related to renewable curtailment, and finally TIC is the transmission investment cost.

An hourly resolution is adopted to accurately capture operational constraints, system flexibility requirements, and the temporal interaction between variable renewable generation and electricity demand, while maintaining computational tractability for the 33-country European power system. To this end, the planning horizon from 2025 to 2050 is discretized into five-year investment periods. Within each representative year, a time-reduction technique based on k-medoids clustering is applied to select 12 representative days, with three typical days assigned to each season, ensuring adequate coverage of seasonal variability [22, 23]. Operational variables, including generation dispatch, storage operation, and cross-border exchanges, are then resolved on an hourly basis for these representative days, with appropriate weighting factors used to approximate full-year system behavior. To provide a planning-consistent baseline, we model the evolution of the European electrical system

using the ENTSO-E TYNDP Global Ambition scenario [13]. To ensure a more realistic system representation, the model includes reliability (LOLE, loss of load expectation) constraints, must-run for thermal units, renewable curtailment costs, and land availability limits for renewable expansion [24]. It is worth noting that, in addition to conventional generation technologies, carbon capture and storage (CCS) options, including Combined-cycle gas turbine with CCS (CCGT-CCS) and bioenergy with CCS (BECSS), are considered as candidate technologies to facilitate compliance with decarbonization targets [25].

Regional Data Centre Capacity

Although obtaining precise current and future capacity data at the country level is highly challenging and no standardized forecasting methodology exists for installed data centre capacity, the EUDCA provides current capacity estimates and annual projections for hyperscale and colocation data centres across Europe according to Colocation and hyperscale data centre database, Pb7 Research [26]. The database is compiled and maintained by Pb7 Research using a combination of publicly available sources, targeted online searches, and validation checks. Where direct data are unavailable, values are estimated based on complementary sources (e.g., the number of racks or building layout). The estimates were obtained through data gathering and interviews with owners and operators of data centres throughout Europe. Overall, the reported capacities show reasonable consistency with the figures presented in the IEA report.

It should be noted that all reported values refer to IT capacity. These values are converted into total allocated capacity using the projected PUE values provided by the IEA (1.45 in 2024 and 1.29 in 2030) [27]. The 2024 capacity for each country is therefore used as the initial condition and baseline input for optimisation model.

Data Centre Load Profile

In addition to installed data centre capacity, both the associated energy consumption and the contribution of data centres to each country's total electricity demand must be considered. To this end, we incorporate realistic data centre load profiles into the model. Based on the dataset reported in, half-hourly percentage utilisation data for 87 data centres in the UK are available for the period 2023– July 2025 [28]. By clustering these load profiles, all data centres are classified into three voltage levels: low, medium, and high voltage. The reported utilisation percentage represents the ratio of the actual kVA import of each data centre to its contracted (maximum) kVA capacity from the perspective of the regional electricity network operator. Since hyperscale data centres with capacities exceeding 10 MW (at minimum) require high-voltage connections (110, 230, or 400 kV) and are typically connected to the transmission or sub-transmission network [29, 30], we derive a representative load profile from 57 data centres connected at high-voltage levels using K-medoids clustering. Fig. 1 illustrates the representative load profile of a high-voltage-connected data centre employed in the model.

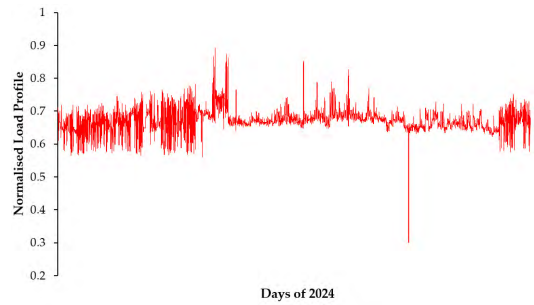

Figure 1. Historical hourly load profile of hyperscale data centre [28].

Hence, the energy consumption by AI is calculated as follows:

$$DemandAI_{t,c,h} = DCProfile_h \times CapAI_{t,c} \quad (2)$$

where $DemandAI_{t,c,h}$ is a variable representing AI electricity demand in year t, country c, and hour h, $CapAI_{t,c}$ denotes the optimal AI capacity in year t and country c, and $DCProfile_h$ is the normalized hourly demand profile used as model input.

RESULTS & DISCUSSION

In this section, we present the results of the proposed optimisation framework across all scenarios. First, a benchmark analysis for the base case validates the model by comparing its outputs with TYNDP projections. Next, AI-based data centres are incorporated to assess the impacts of AI-driven demand on capacity investment, generation mix, and emissions trajectories.

Base Case Analysis (Without AI)

The results of this scenario closely match the TYNDP Global Ambition trajectory [13]. Fig. 2-(a) illustrates the evolution of total installed capacity from 2025 to 2050, closely following the stated decarbonisation pathway. The model projects total installed capacity to reach approximately 4, 700 GW by 2050, which is about 3.3% higher than the TYNDP projection of 4, 548 GW. The resulting energy mix reproduces the main structural transitions of the TYNDP scenario, including the complete phase-out of coal and oil-fired generation by 2035 and the subsequent dominance of renewable energy sources. By 2050, the combined installed capacity of solar and wind reaches around 3, 300 GW, representing a 3.0% increase relative to the TYNDP estimate of 3, 205 GW. Regarding nuclear capacity, the model projects 141 GW by 2050, approximately 15% higher than the TYNDP forecast. This deviation is primarily driven by adequacy and

spinning reserve constraints, which increase the value of firm, low-carbon generation in maintaining system reliability. In term of spatial distribution of installed capacity, Fig. 2-(b), Germany will have the largest total capacity in Europe, at around 11, 000 GW, more than 50% of which is solar. Following Germany, Italy (490 GW), France (377 GW), the UK (333 GW), and Spain (286 GW) are next.

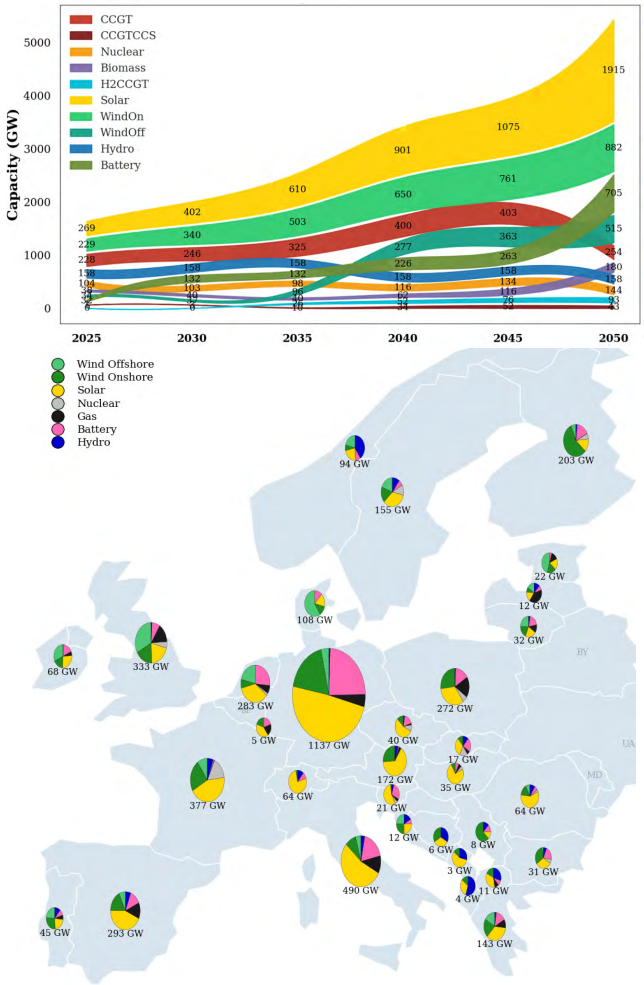

Figure 2. Base case analysis: **(a)** Installed capacity trajectories, **(b)** Country-level installed capacity.

The need for flexibility and firm capacity is clearly visible in countries such as the UK (50 GW CCGT, 10 GW CCGT-CCS), Italy (42 GW CCGT, 9 GW CCGT-CCs), Poland (40 GW CCGT), Spain (25 GW CCGT, 16 GW CCGT-CCS) where the model requires significant gas-based generation to meet reliability and reserve requirements. However, the gas installed in Germany and the Netherlands is mainly based on H2-CCGT technology (close to 60 GW in both countries) by 2050. Overall, Europe will have about 140 GW of nuclear capacity, with France leading at 62 GW, followed by the UK (18 GW), Sweden (17 GW), and Finland and Poland (10 GW each). In terms of wind capacity, the UK leads with 170 GW, Finland (130GW), followed by Denmark and the Netherlands (each 78 GW). High solar capacity in Southern Europe is also evident due to higher solar resource potential.

Impact of different Data Centres growth on the evolution of the power system

In this section, the impact of AI-driven electricity demand on Europe energy system are analysed under one lift-off scenario and two collapse scenarios (pessimistic–med and pessimistic–deflation). The lift-off scenario represents an optimistic pathway, characterized by sustained and rapid growth in AI deployment through 2050. In contrast, the deflation collapse scenario assumes that a contraction of data centre activity begins in 2035 and persists until 2045, while in the medium collapse scenario, the contraction starts in 2035, but AI demand partially recovers after 2040, albeit at a lower growth rate than in the lift-off scenario (cf. Fig. 4 red, dashed line).

Fig. 3 illustrates changes in the installed capacity under the three AI scenarios relative to the base case for the milestone years 2030, 2040, and 2050. Under the lift-off scenario, the system experiences a significant demand shock as early as 2030, requiring approximately 60 GW of additional onshore wind capacity, primarily in Germany, Spain, and Portugal. At the same time, around 10 GW of offshore wind and 35 GW of additional solar capacity are required compared to the base case. To accommodate the increased variability associated with higher renewable penetration, the system also requires approximately +37 GW battery by 2030.

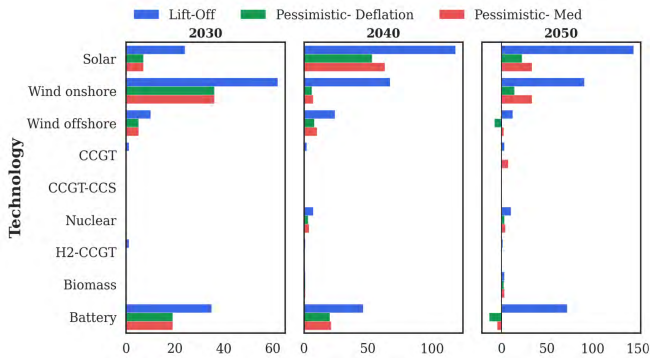

Figure 3. Change in Installed capacity under AI scenarios.

By 2040, the lift-off scenario requires approximately +130 GW of solar, +60 GW of onshore wind, and +10 GW of offshore wind compared to the base case. High AI-driven load also necessitates +3 GW of nuclear, primarily in France, to maintain system adequacy. By 2050, additional nuclear requirements range from 1 GW in the deflation scenario to 10 GW in the lift-off scenario, while lift-off and medium scenarios need 5 GW and 3 GW of CCGT, respectively. The growing reliance on gas generation under tighter emissions constraints may challenge long-term decarbonisation goals.

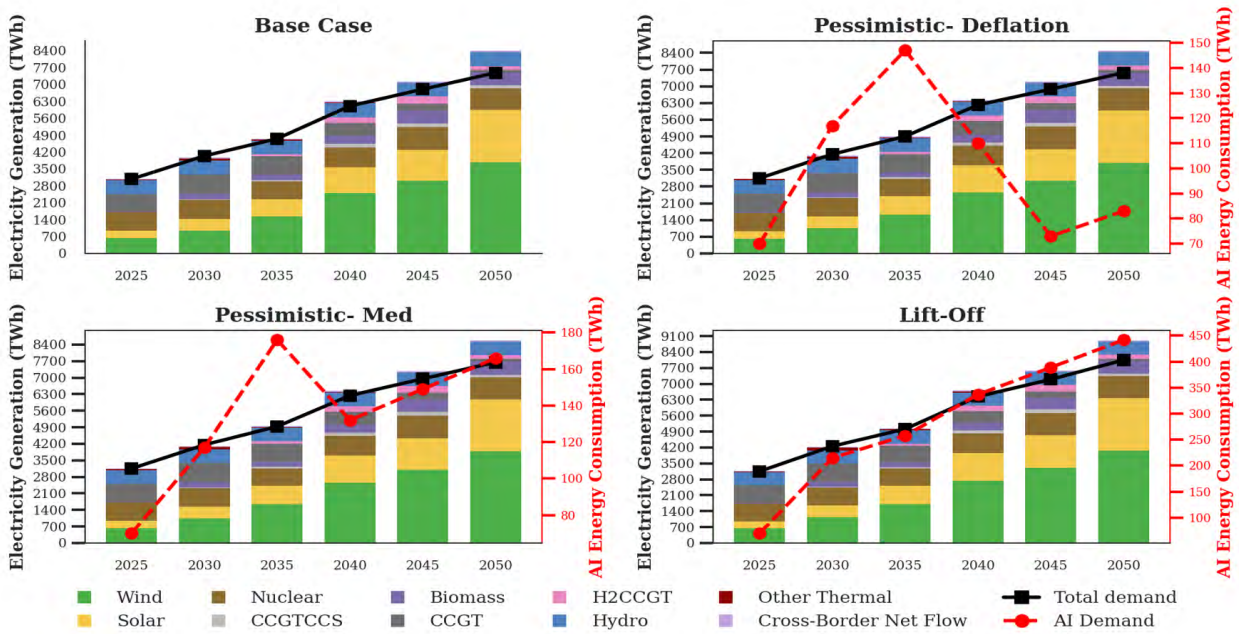

Figure 4: Energy mix trajectory under different scenarios, as well as AI energy consumption (red, dashed line)

Under the collapse scenarios, the system exhibits similar behavior until 2035. By 2040, in the deflation scenario, declining AI demand is primarily accommodated through additional renewable and storage capacity (+50 GW solar, +9 GW wind, and +15 GW battery). In contrast, the pessimistic-medium scenario requires not only renewables but also firm capacity additions (+2 GW nuclear and +3 GW CCGT). By 2050, the deflation scenario leads to partial retirement of offshore wind and battery capacity, with offshore wind reduced by about 7 GW in Finland and Denmark due to sustained AI contraction, although nuclear capacity remains 2 GW above the base case and gas capacity is unchanged. Under the pessimistic-medium scenario, renewed AI growth in the final decade drives additional investments of approximately +25 GW solar, +25 GW wind, +3 GW nuclear, and +4 GW CCGT relative to the base case.

Impact on power generation and interconnection

Fig. 4 compares the evolution of the European energy mix under different AI demand scenarios relative to the base case. The AI-driven electricity demand profile is highlighted in red for all three scenarios. Comparing this profile with total system load shows that, by 2050, AI demand accounts for approximately 450 TWh of total electricity consumption in the lift-off scenario, around nine times higher than in the deflation scenario and three times higher than in the medium scenario. As AI penetration increases, the capacity factors of nuclear and gas technologies rise by up to 15%, and 10 % respectively. As shown, unlike the base case, where coal and nuclear generation are almost entirely phased out after 2040, approximately 8 TWh of electricity between 2040 and 2045 is supplied by these technologies under the lift-off scenario. Cumulatively over the period 2030–2050, the system requires around +150 TWh of nuclear and +97 TWh of additional gas generation compared to the base case in lift-off scenario. In the final decade, tighter emissions constraints increase reliance on H_2-CCGT for flexibility, adding 26 TWh more while Wind captures the largest share of energy mix growth and emerges as the most dominant option for supplying AI-related demand, contributing more than +650 TWh over the same period. Overall, these results indicate that the highly concentrated electricity demand associated with hyperscale data centres can significantly alter the energy mix at critical planning milestones.

Optimal allocation of data centres

From a spatial optimisation perspective, AI-related electricity demand is primarily deployed in countries with sufficient potential to expand nuclear and thermal generation capacity. This pattern is consistent across all scenarios. In the lift-off scenario, Poland (14 GW), France (13 GW), and the UK (11 GW), experience high growth data centre deployment during the 2040–2050 period, driven by their ability to expand firm nuclear and gas capacity (Fig. 5-a). Under lift-off scenario, Portugal emerges as a favorable location for data centre development (3.4 GW), alongside Spain (6 GW), where a combination of low-cost solar electricity and available gas capacity exists.

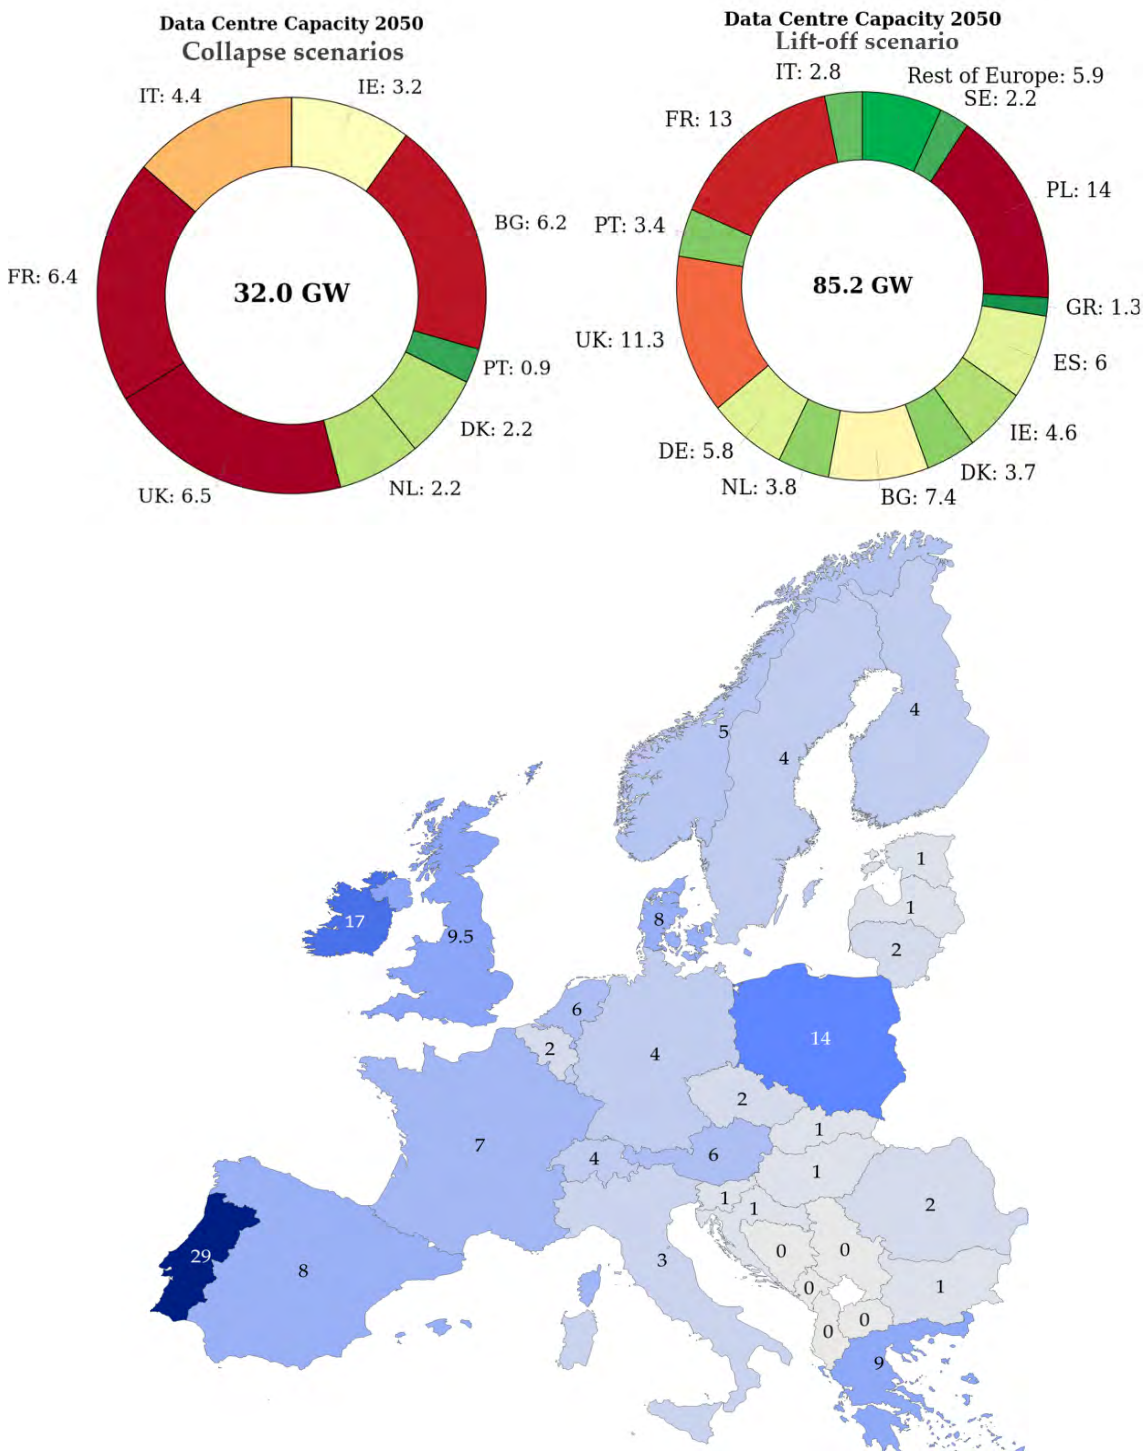

Figure 5: (a) Optimal data centre capacity, (b) Share of AI as a percentage of total demand per country (lift-off).

Between 2030 and 2040, Bulgaria and Greece also exhibit notable growth potential, reaching 7.5, and 1.5 GW in lift off scenarios. In contrast, Ireland, Germany, and the Netherlands face binding constraints from 2030 onward, or experience capacity reductions under collapse scenarios, due to stringent carbon limits, rapid growth in base electricity demand, and the absence of nuclear power options.

Under the pessimistic-medium collapse scenario, total AI-based data centre capacity reaches approximately 32 GW by 2050. Owing to the decline in AI demand prior to 2045, baseload nuclear generation is

largely sufficient to accommodate the additional load. As a result, France, the UK, and Bulgaria, benefiting from adequate nuclear expansion potential, together host nearly 20 GW of AI capacity. Italy, enabled by the deployment of CCGT-CCS in the final five years, accommodates around 4.4 GW. Furthermore, driven by the development of H_2-CCGT alongside wind capacity between 2035 and 2045, Ireland and the Netherlands recover from the post-2030 collapse and expand their AI capacity to approximately 3.2 GW and 2.2 GW, respectively. These results indicate that even under a market collapse scenario, Europe's power system continues to rely on nuclear and gas-based technologies, particularly CCGT-CCS alongside renewables to ensure system adequacy.

Spatial AI Demand Shares and Associated Emissions Trajectories

By 2050, AI-based data centres may account for approximately 3–7% of Europe's total electricity consumption, ranging from the pessimistic to the lift-off scenario. The share of AI-driven electricity demand in total national consumption varies widely: 23% in Portugal, 14% in Poland, 9.5% in the UK and Greece, and 7% in France (Fig. 5-b, lift-off scenario). Even at these levels, AI-driven demand has a significant impact on emissions. Across Europe, it contributes approximately 25 $MtCO_2$ cumulatively between 2030 and 2050, driven by an additional 10 TWh of oil- and coal-fired generation and 97 TWh of gas-fired generation. Country-specific impacts vary: Portugal sees only 0.02 $MtCO_2$, the UK ~4 $MtCO_2$, Germany and the Netherlands ~8 $MtCO_2$ combined due to higher gas and coal use, France ~0.2 $MtCO_2$, and Poland nearly 5 $MtCO_2$. These differences reflect each country's mix of AI load, existing capacity, and reliance on flexible fossil generation.

CONCLUSION

This paper investigates the effects of AI-driven demand growth on Europe's power system development toward 2050 using a comprehensive optimisation framework. Rapid AI expansion substantially increases generation capacity needs, placing high pressure on long-term complemented by flexible gas generation. Across all scenarios, only countries able to simultaneously expand renewable energy while retaining nuclear and gas options infrastructure planning can accommodate sustained AI growth. Between 2030 and 2050, AI demand contributes an additional 323 GW additional capacity from different generation technologies, as well as ~25 $MtCO_2$, posing a challenge to Europe's climate targets under the lift-off scenarios. While Poland, UK and France see significant capacity growth, Southern European countries (Spain, Portugal, Greece) also emerge as key locations for AI deployment under both lift-off and collapse scenarios.

These findings underscore that successfully integrating AI at scale will require the coordinated expansion of flexible, low-carbon generation and infrastructure across Europe. Ongoing work in our group focuses on systematic representations of growth uncertainty [31, 32] via means of multi-stage stochastic programming to elucidate further aspects of optionality.

Figure 5: (a) Optimal data centre capacity, (b) Share of AI as a percentage of total demand per country (lift-off).

ACKNOWLEDGEMENTS

Financial support under the UK EPSRC grants EP/T022930/1, EP/W003317/1 and EP/V051008/1 is gratefully acknowledged.

REFERENCES

1. Kamiya, G. and P. Bertoldi, *Energy consumption in data centres and broadband communication networks in the EU*. 2024: Publications Office of the European Union Luxembourg.
2. Demirtas C, Tiwari AK, Soyu Y?ld?r?m E, Shahbaz M. Does financial development support renewable energy consumption: evidence from the UK. Renewable Energy 243:122480 (2025). https://doi.org/10.1016/j.renene.2025.122480
3. Post, W. *How Much Energy Can AI Use?* 2025; WebPage.
4. de Roucy-Rochegonde, L. and A. Buffard, *AI, Data Centers and Energy Demand: Reassessing and Exploring the Trends*. 2025, Institut français des relations internationales (Ifri): Paris, France.
5. IEA, *World Energy Outlook 2025*. 2025, IEA: Paris.
6. Agency, I.E., *Energy and AI*. 2025, IEA: Paris.
7. Tommarello ND. Ai's appetite for energy: how data centers are affecting the environment. Computer 58:102-105 (2025). https://doi.org/10.1109/mc.2025.3603370
8. Analytics, I.P. Data centres: Hungry for power, Forecasting European power demand from data centres to 2035. 2025; WebPage.
9. Vanderbauwhede, W., Estimating the Increase in Emissions caused by AI-augmented Search. arXiv preprint arXiv:2407.16894, 2024.
10. Sachs, G. AI is Poised to Drive 160% Increase in Data Center Power Demand. 2024; WebPage.
11. Luers A. Net zero needs AI — five actions to realize its promise. Nature 644:871-873 (2025). https://doi.org/10.1038/d41586-025-02641-4
12. Hörsch J, Hofmann F, Schlachtberger D, Brown T. Pypsa-eur: an open optimisation model of the european transmission system. Energy Strategy Reviews 22:207-215 (2018). https://doi.org/10.1016/j.esr.2018.08.012

13. ENTSO-E, TYNDP Scenarios Methodology Report - Final Version. 2025.
14. Takci MT, Qadrdan M, Summers J, Gustafsson J. Data centres as a source of flexibility for power systems. Energy Reports 13:3661-3671 (2025). https://doi.org/10.1016/j.egyr.2025.03.020
15. Xiao T, Nerini FF, Matthews HD, Tavoni M, You F. Environmental impact and net-zero pathways for sustainable artificial intelligence servers in the USA. Nat Sustain 8:1541-1553 (2025). https://doi.org/10.1038/s41893-025-01681-y
16. Weidner T, Guillén-Gosálbez G. Planetary boundaries assessment of deep decarbonisation options for building heating in the european union. Energy Conversion and Management 278:116602 (2023). https://doi.org/10.1016/j.enconman.2022.116602
17. Karthikeyan K, Chattopadhyay S, Gandhi R, Grossmann IE, Torres AI. Pareto optimal solutions for decarbonization of oil refineries under different electricity grid decarbonization scenarios. Systems and Control Transactions 4:844-849 (2025). https://doi.org/10.69997/sct.102781
18. van Greevenbroek K, Grochowicz A, Zeyringer M, Benth FE. Trading off regional and overall energy system design flexibility in the net-zero transition. Nat Sustain 8:629-641 (2025). https://doi.org/10.1038/s41893-025-01556-2
19. Alanazi K, Mittal S, Hawkes A, Shah N. Renewable hydrogen trade in a global decarbonised energy system. International Journal of Hydrogen Energy 101:712-730 (2025). https://doi.org/10.1016/j.ijhydene.2024.12.452
20. Charitopoulos VM, Fajardy M, Chyong CK, Reiner DM. The impact of 100% electrification of domestic heat in great britain. iScience 26:108239 (2023). https://doi.org/10.1016/j.isci.2023.108239
21. Koltsaklis NE, Dagoumas AS, Kopanos GM, Pistikopoulos EN, Georgiadis MC. A spatial multi-period long-term energy planning model: a case study of the greek power system. Applied Energy 115:456-482 (2014). https://doi.org/10.1016/j.apenergy.2013.10.042
22. Efthymiadou ME, Charitopoulos VM, Papageorgiou LG. Optimal hydrogen infrastructure planning for heat decarbonisation. Chemical Engineering Research and Design 204:121-136 (2024). https://doi.org/10.1016/j.cherd.2024.02.028
23. Zhou X, Efthymiadou ME, Papageorgiou LG, Charitopoulos VM. Data-driven robust optimisation of hydrogen infrastructure planning under demand uncertainty using a hybrid decomposition method. Applied Energy 376:124222 (2024). https://doi.org/10.1016/j.apenergy.2024.124222
24. Koltsaklis NE, Georgiadis MC. A multi-period, multi-regional generation expansion planning model incorporating unit commitment constraints. Applied Energy 158:310-331 (2015). https://doi.org/10.1016/j.apenergy.2015.08.054
25. Bui M, Adjiman CS, Bardow A, Anthony EJ, Boston A, Brown S, Fennell PS, Fuss S, Galindo A, Hackett LA, Hallett JP, Herzog HJ, Jackson G, Kemper J, Krevor S, Maitland GC, Matuszewski M, Metcalfe IS, Petit C, Puxty G, Reimer J, Reiner DM, Rubin ES, Scott SA, Shah N, Smit B, Trusler JPM, Webley P, Wilcox J, Mac Dowell N. Carbon capture and storage (CCS): the way forward. Energy Environ. Sci. 11:1062-1176 (2018). https://doi.org/10.1039/c7ee02342a
26. EUDCA, State of European Data Centres 2025. 2025.
27. IEA. *Energy and AI Observatory*. 2025; WebPage.
28. Fajardo-Barco ES, Achicanoy-Martinez WO. Electricity demand forecasting in distribution circuits with incomplete data and differentiated demand profiles using spatiotemporal graph neural networks. 2025 IEEE 7th Colombian Conference on Automatic Control (CCAC) :1-6 (2025). https://doi.org/10.1109/ccac64704.2025.11259266
29. IEA, Powering Ireland's Energy Future. 2025.
30. Lopez F, Diguet C. Digital conquest. Policy Press (2026). https://doi.org/10.1332/policypress/9781529252156.001.0001
31. Zhou X, Charitopoulos VM. Scenario generation via moments-informed normalizing flows for stochastic optimization of local energy markets. Energy and AI 22:100649 (2025). https://doi.org/10.1016/j.egyai.2025.100649
32. Bounitsis GL, Papageorgiou LG, Charitopoulos VM. Stable optimisation-based scenario generation via game theoretic approach. Computers & Chemical Engineering 185:108646 (2024). https://doi.org/10.1016/j.compchemeng.2024.108646

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Simulation of Methanol Production from Biogas: Impact of Feedstock Composition and Stoichiometric Number Adjustment

Muhammad Zulkefal[a], Magne Hillestad[b], Truls Gundersen[a], and Bjørn Austbø[a]*

[a] Department of Energy and Process Engineering, Norwegian University of Science and Technology (NTNU), NO-7491 Trondheim, Norway
[b] Department of Chemical Engineering, Norwegian University of Science and Technology (NTNU), NO-7491 Trondheim, Norway
* Corresponding Author: bjorn.austbo@ntnu.no

ABSTRACT

Biogas offers a promising biogenic carbon source for renewable methanol, but differences in CH_4/CO_2 ratio across feedstocks and possible upstream CO_2 handling can shift syngas stoichiometry away from the methanol synthesis target range. This work quantifies how biogas composition and reformer operation influence the stoichiometric number (SN) and the associated conditioning requirement needed to meet methanol synthesis targets. A steady-state Aspen Plus® model of an integrated biogas-to-methanol process is used as the analysis framework. A base-case operating point is defined, followed by parametric evaluation of biogas CH_4/CO_2 ratio, reformer temperature, reformer pressure and steam-to-methane (S/C) ratio. The studied CH_4/CO_2 ratio range covers CO_2-rich to CH_4-rich cases that may occur across sites and upgrading levels. The resulting SN shifts are tracked and converted into a quantitative correction requirement to maintain the methanol design target (SN = 2.01). Temperature determines the upper limits of conversion and SN, while pressure and S/C ratio have secondary effects once high-temperature operation is reached. By comparison, the CH_4/CO_2 ratio has the strongest influence on syngas composition, defining H_2-deficiency and H_2-excess regimes. Methanol production increases as SN approaches the required target range but shows diminishing gains beyond it, indicating limited benefit from excess hydrogen under fixed synthesis conditions.

Keywords: biomethanol, biogas, stoichiometric number, eSMR

INTRODUCTION

International shipping is entering a transition period driven by tightening greenhouse-gas regulations and the need to reduce transport emissions. The International Maritime Organization (IMO) reports that international shipping emitted close to one billion tonnes of CO_2 in 2018 and that emissions are expected to grow with the continued rise of maritime trade [1]. To counter this, the IMO and the European Union (EU), among other stakeholders, have implemented measures to support the decarbonisation of marine transport, thus increasing the push for sustainable alternative fuels [2, 3].

Among the available alternatives, methanol has attracted strong interest as a practical short-term option for maritime decarbonisation. Methanol is in liquid form at ambient conditions, it can be stored and handled using established infrastructure, and it is already being used in dual-fuel marine engines [4]. The key point, however, is that methanol's decarbonisation benefit depends on how it is produced. When synthesised from renewable or biogenic carbon sources (biomethanol), methanol can offer substantially lower life-cycle emissions compared to conventional marine fuels [5]. Despite this potential, large-scale renewable methanol production remains limited and requires process configurations that combine robust syngas generation with efficient synthesis and purification.

Biogas-based methanol production is particularly attractive because it relies on biogenic carbon, but it also introduces an important process challenge: biogas composition varies widely with feedstock and digester

operation. Biogas typically contains a significant amount of both CH_4 and CO_2. In steam methane reforming (SMR) of biogas, the inlet carbon balance directly shapes the outlet H_2, CO, and CO_2 levels. Hence, changes in feed CH_4/CO_2 ratio strongly influences the resulting syngas composition and its suitability for downstream methanol synthesis [6]. In practice, this means that plants for biogas-based methanol production should be designed to handle varying feed conditions.

A lower direct-emissions route to generate syngas is electrified steam methane reforming (eSMR). In eSMR, high-temperature reforming heat is supplied electrically rather than by fuel combustion, which can reduce direct CO_2 emissions from syngas generation when the electricity is renewable or from other low-carbon sources [6]. However, electrifying the heat supply does not by itself guarantee a syngas composition suitable for methanol synthesis. The syngas composition at the reformer outlet remains sensitive to reformer temperature, pressure, and steam-to-carbon ratio and, critically, the CH_4/CO_2 ratio of the biogas feed, so syngas stoichiometry must be adjusted to meet methanol synthesis requirements.

Syngas quality for methanol synthesis is commonly assessed using the stoichiometric number (SN), $(H_2-CO_2)/(CO+CO_2)$, which reflects the balance between H_2 and carbon oxides (CO_x) in the synthesis gas. Stoichiometric methanol synthesis is equivalent to SN = 2, since this corresponds to the hydrogen-to-CO_x balance required by the overall methanol synthesis stoichiometry. Hence, values close to 2 are typically preferred for methanol production. When SN > 2, the syngas is hydrogen-rich, whereas SN < 2 indicates a carbon-rich mixture with insufficient hydrogen for the synthesis stoichiometry. For biogas reforming, external adjustment through the addition of H_2 or CO_2 (or, alternatively removal of CO_2 or H_2) may be necessary to meet the SN required for methanol synthesis.

The objective of this work is to quantify how biogas composition and reforming conditions influence synthesis gas stoichiometry and the downstream adjustment required to meet methanol design targets. The analysis is focused on steady-state behavior to give clear insight into the relationship between CH_4/CO_2 variability, reformer operation, and stoichiometric control. Specifically, an integrated Aspen Plus® model of a biogas-to-methanol process (eSMR → syngas conditioning → methanol synthesis loop → purification) is used to map how changes in CH_4/CO_2 ratio shift SN and to determine the magnitude and direction of the required conditioning action needed to reach the target SN used for methanol synthesis.

A CH_4/CO_2 ratio range of 1-5 is used as a bounding design range to cover differences across biogas sources and the effect of upstream CO_2 handling, rather than day-to-day variability for a single digester. Hence, the main focus is the syngas production. Process design, optimization, energy use, and cost are not covered in this study.

METHODS

Process Modeling and Simulation

Initially, a steady-state Aspen Plus® base-case model was developed to represent the integrated biogas-to-methanol pathway. The flowsheet comprises biogas compression and preheating, mixing with steam, syngas generation in the eSMR reactor, SN adjustment through addition of H_2 or CO_2, a methanol synthesis loop with recycle/purge, and downstream purification. To calculate thermodynamic properties, Peng–Robinson (PR) was used for the reforming section, Redlich-Kwong-Soave with modified Huron-Vidal mixing rules (RKSMHV2) for the high-pressure methanol synthesis loop, and the non-random two-liquid model (NRTL) for the low-pressure separation section. The main assumptions and key model parameters for the reformer and methanol reactor are listed in Table 1, while Figure 1 illustrates the overall biogas-to-methanol model in more detail.

Table 1: Reformer and methanol reactor key parameters and assumptions.

Block name	Block type	Parameter	Value
ESMR	RGibbs	Calculation mode: chemical equilibrium	
		Operating temperature	950 °C
		Operating pressure	25 bar
		Pressure drop	1 bar
BWR	RPlug	Reactor temperature	230 °C
		Reactor Pressure	75 bar
		Number of tubes	1100
		Length	7 m
		Diameter	0.04m
		Pressure drop	Ergun eq.
		Particle density	1775 kg/m^3
		Particle diameter	5.5 mm

A simplified biogas composition of 60 % CH_4 and 40 % CO_2 (60/40 biogas) was chosen for the base case, with a feed rate of 145 tonne/day. The biogas is compressed to the reformer pressure in a three-stage compressor with intercooling. The compressed biogas is then mixed with 172.7 tonnes/day steam to reach an H_2O/CH_4 (S/C) ratio of 3, selected to suppress carbon formation and ensure stable reforming. In the base case, the eSMR was operated at 950 °C and 25 bar. A target of SN = 2.01 was chosen for the methanol synthesis, potentially adding H_2 or CO_2 downstream to meet this requirement. A value

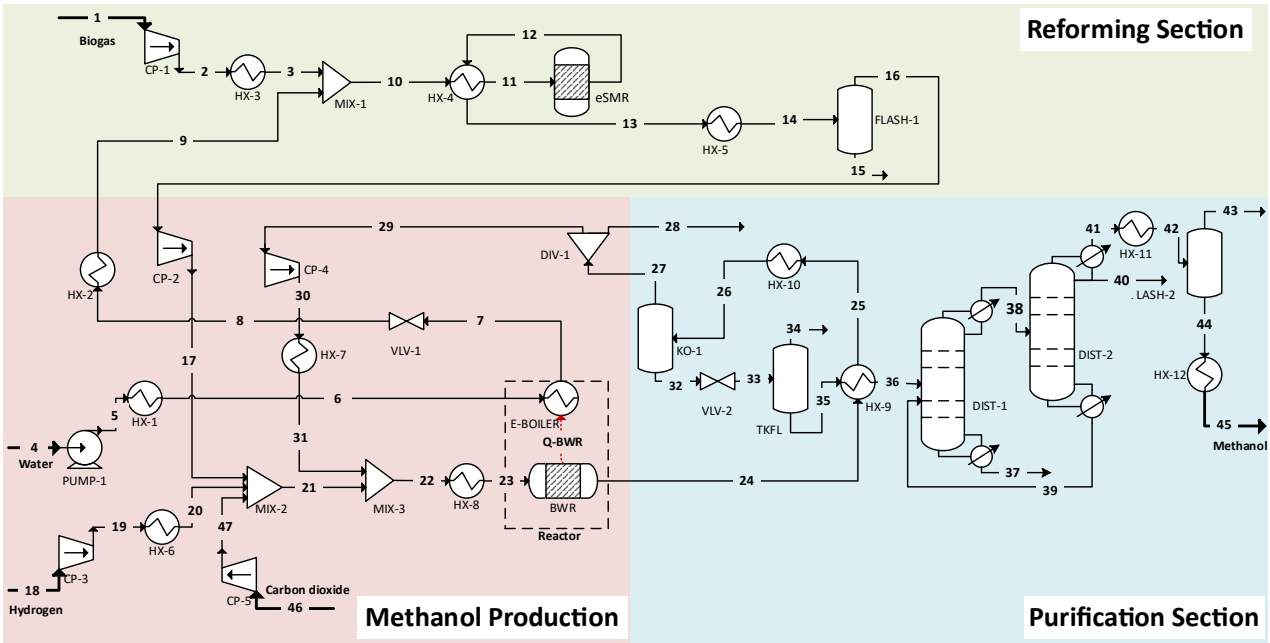

Figure 1: Detailed flowsheet for the biogas-to-methanol production base case

slightly above 2 was chosen as the reaction kinetics benefit from high H_2 partial pressure, and the catalyst benefits from a reducing atmosphere.

The eSMR operating temperature and pressure were selected based on the parametric sensitivity results, while the methanol reactor temperature and pressure were chosen to reflect industrial operating ranges reported by Nyári et al. [7]. The methanol synthesis catalyst and kinetic data were taken from Van-Dal et al. [8].

Reforming Section

The eSMR reactor was modelled in Aspen Plus® using an RGibbs equilibrium block. This choice is consistent with electrically heated reforming concepts, where heat is supplied rapidly and uniformly, minimizing thermal gradients and enabling operation close to thermodynamic equilibrium, as reported by Wismann et al. [9]. The following reforming reactions were modeled:

$$CH_4 + H_2O \leftrightarrow CO + 3H_2 \quad (1)$$

$$CO + H_2O \leftrightarrow CO_2 + H_2 \quad (2)$$

In conventional reforming processes aimed at H_2 production, a water–gas shift (WGS) unit is often added to increase H_2 yield. However, when producing methanol, SN plays the key role. Since the WGS reaction generates 1 mol CO_2 for every 1 mol H_2 produced, it has no net effect on SN and mainly adds extra equipment and utility demand. Therefore, a WGS unit was not included in the model.

Methanol Synthesis and Purification

The methanol reactor was modelled in Aspen Plus® using an RPlug reactor to represent a fixed-bed boiling water reactor (BWR). This reactor concept was selected because the exothermic reaction heat can be recovered as steam, which can be integrated to support the upstream steam demand in the reforming section. In addition, the BWR configuration provides effective temperature control, enabling operation close to the optimal reaction temperature.

The reactor kinetics were implemented using the Vanden Bussche-Froment model due to its validity at high-pressure [7], making it suitable for commercial operating conditions. The following reactions were modeled:

$$CO_2 + 3H_2 \leftrightarrow CH_3OH + H_2O \quad (3)$$

$$CO + 2H_2 \leftrightarrow CH_3OH \quad (4)$$

$$CO_2 + H_2 \leftrightarrow CO + H_2O \quad (5)$$

A recycle-to-fresh ratio of 4 was chosen, balancing conversion gains against increased compression duty. The reactor geometry was tuned to achieve a superficial gas velocity near 0.4 m/s, supporting hotspot mitigation while maintaining reasonable residence time and heat transfer performance in the catalyst bed.

Downstream purification was represented by a flash

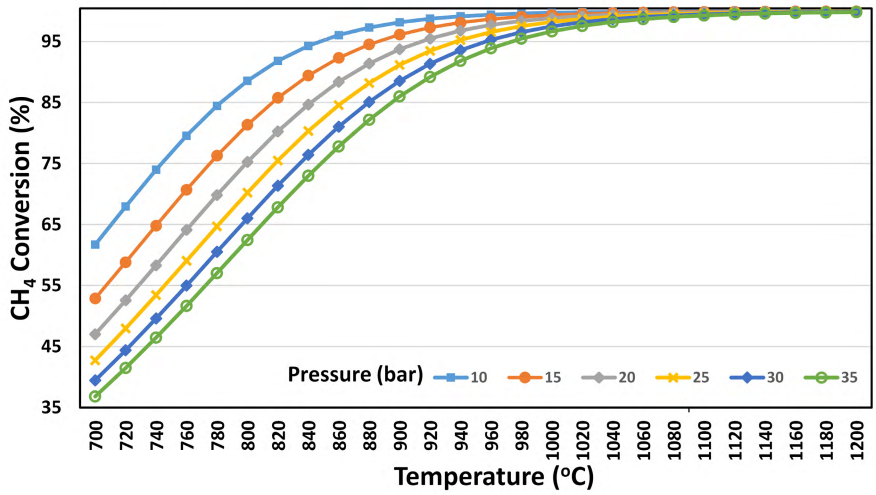

Figure 2. CH_4 conversion as function of reformer temperature for different reformer pressures (60/40 biogas, S/C = 3).

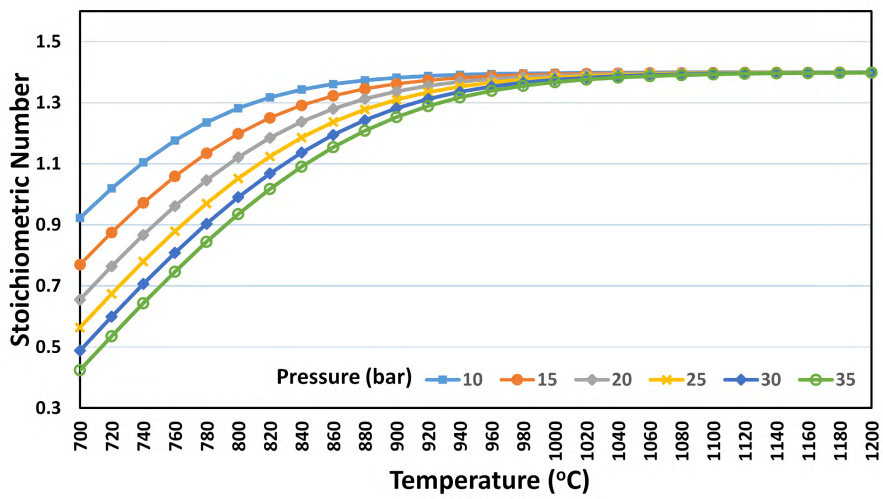

Figure 3. Reformer-outlet SN as a function of reformer temperature for different reformer pressures (60/40 biogas, S/C = 3).

separation (for non-condensable removal) followed by two distillation columns to obtain fuel-grade methanol (99.85 wt%) [10]. The purification section is included to close the overall process and verify the reported fuel-grade methanol purity, but it is not the main focus of this work.

Parametric Analyses

Parametric analyses were performed on a molar basis using the base-case model to examine how changes in feed composition, S/C ratio, reformer temperature and reformer pressure influence reformer-outlet syngas composition and downstream methanol production. Differences between the reformer outlet SN and the methanol design target were assessed to indicate the extent of composition adjustment required to reach the target stoichiometry, i.e., the amount of H_2 or CO_2 added (or removed, depending on the availability of the corresponding resources) to reach the target SN.

RESULTS AND DISCUSSION

Base Case Model

In the base case, the CH_4 conversion in the eSMR reactor was 96.2%, resulting in syngas with SN = 1.36. In order to reach SN = 2.01, 6.8 tonne/day green H_2 was added downstream. The final methanol production in the base case was 151 tonne/day. Table 2 summarises key overall conversion metrics for the base case model.

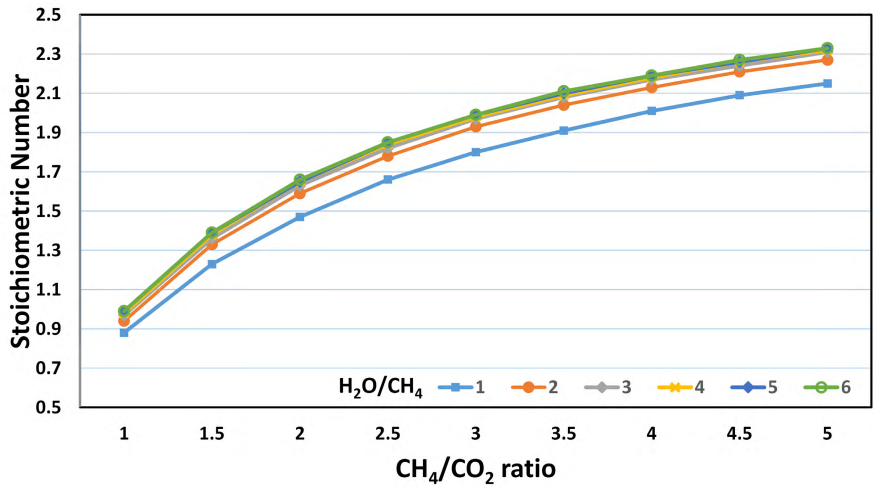

Figure 4. SN as function of biogas CH_4/CO_2 ratio for different S/C ratios (950 °C, 25 bar).

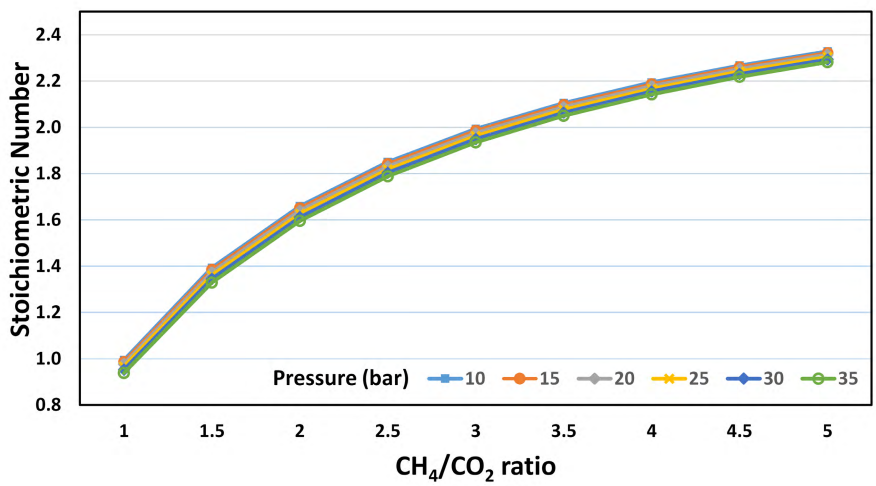

Figure 5. SN as function of biogas CH_4/CO_2 ratio for different reformer pressures (950 °C and S/C = 3).

Table 2: Base case conversion for the methanol section.

Parameter	Value
CO conversion (%)	98.8
CO_2 conversion (%)	87.9
H_2 conversion (%)	93.6
Biogas to methanol conversion (%)	92.9

Parametric Analysis

Figures 2 and 3 first establish the combined influence of temperature and pressure on reformer performance for the 60/40 biogas feed at S/C = 3. As expected for an equilibrium-limited endothermic reformer, increasing temperature increases CH_4 conversion and shifts the outlet composition towards higher H_2 production, which is directly reflected in a higher SN.

Pressure acts in the opposite direction, suppressing conversion and lowering SN, with the effect most evident at the lower end of the temperature range. At sufficiently high temperatures, the conversion and SN curves converge across pressures, indicating that moderate pressure changes have a diminishing influence. This behaviour motivates the selection of 950 °C as a reference operating point for the remaining parametric analysis, where conversion and SN are close to their upper-limit values, and pressure sensitivity is limited. This does, however, come at the cost of high energy input to the reformer.

With the reformer temperature fixed at 950 °C, Figures 4 and 5 highlight the role of feed composition through the biogas CH_4/CO_2 ratio. Figure 4 shows SN as a function of the CH_4/CO_2 ratio at constant reformer temperature (950 °C) and pressure (25 bar) while varying the

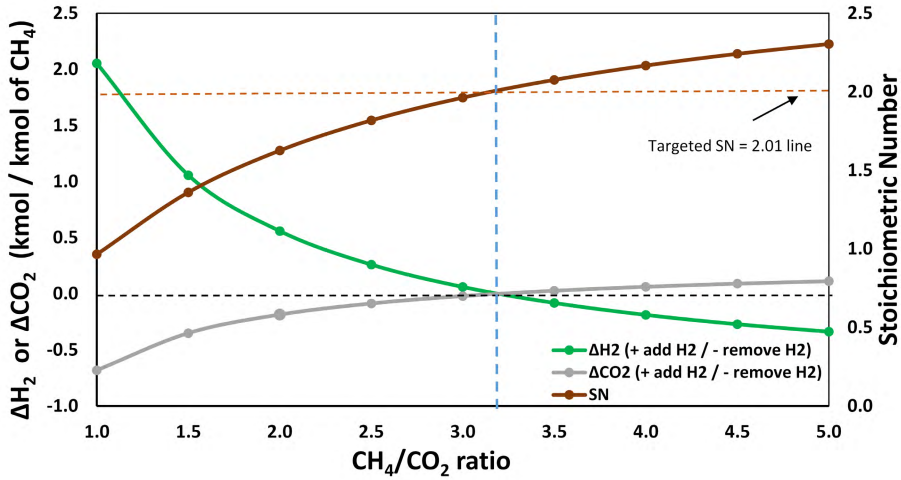

Figure 6. H_2 and CO_2 adjustment required to reach SN = 2.01 as a function of biogas CH_4/CO_2 ratio (950 °C, 25 bar, S/C = 3).

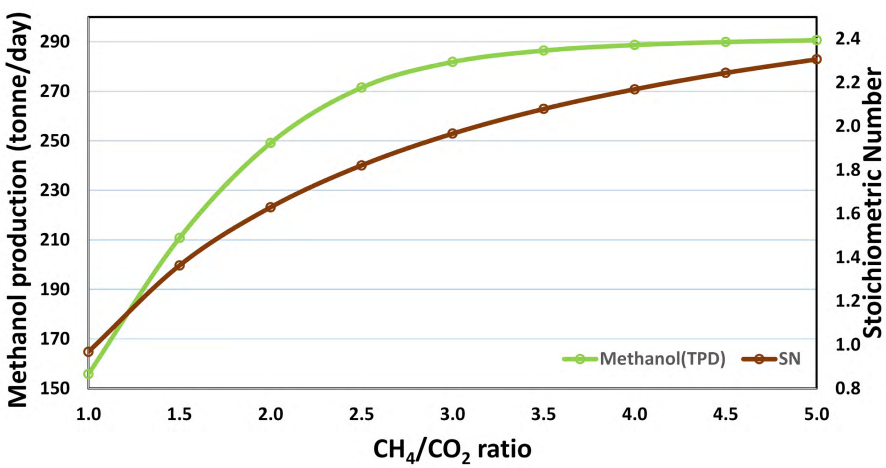

Figure 7. Methanol production rate and SN as functions of biogas CH_4/CO_2 ratio.

S/C ratio. SN increases steadily with increasing CH_4/CO_2 ratio for all S/C values, reflecting improved hydrogen availability as the methane fraction increases. Increasing S/C shifts SN upward across the studied composition range, but the effect reduces at higher S/C, indicating limited additional stoichiometric leverage once sufficient steam is supplied. Importantly, CO_2-rich feeds remain H_2-deficient (SN < 2.01) even at elevated S/C, demonstrating that steam adjustment alone cannot ensure methanol-grade stoichiometry over the full range of biogas compositions studied.

Figure 5 further illustrates this behaviour by showing the pressure sensitivity at constant temperature (950 °C) for S/C = 3. Across the 10-35 bar range studied, pressure has very little effect on SN compared to CH_4/CO_2 ratio. Together, the results in Figures 4 and 5 indicate that while S/C and pressure shift the absolute SN level, the dominant factor governing the transition between H_2-deficiency and H_2-excess regimes is the feed CH_4/CO_2 ratio. This observation motivates the introduction of an explicit stoichiometric correction step when feed composition varies, rather than relying solely on reformer operating adjustments.

In Figure 6, the required stoichiometric correction to reach SN = 2.01 is expressed as a function of the biogas CH_4/CO_2 ratio for a fixed reformer operating point (950 °C, 25 bar, S/C = 3). The maroon curve represents the reformer-outlet SN before any adjustment. The green

curve gives the hydrogen adjustment, ΔH_2 (kmol H_2 per kmol CH_4 in the feed) required to reach SN = 2.01, and the grey curve gives the corresponding required CO_2 adjustment, ΔCO_2(kmol CO_2 per kmol CH_4 in the feed). Positive values indicate addition of the respective component, while negative values indicate removal. The blue dashed vertical line marks the feed composition at which the reformer outlet already meets SN = 2.01, so no correction is required. For CO_2-rich feeds (left of the blue dashed line), the syngas is hydrogen-deficient, and the target can be reached by adding H_2 (or alternatively removing CO_2). For CH_4-rich feeds (right of the blue dashed line), the syngas becomes hydrogen-rich, and the correction reverses, requiring CO_2 addition (or alternatively H_2 removal). Overall, the figure links feed composition directly to the magnitude and direction of the adjustment required, while leaving the practical choice of correction route open, without prescribing a specific implementation route.

Finally, Figure 7 shows how changes in the CH_4/CO_2 ratio affect methanol production when the eSMR is operated at 950 °C, 25 bar, and S/C = 3, and no SN adjustment is applied upstream of the methanol reactor. Biogas flow was fixed at 500 kmol/h (to isolate the composition effect) while varying only the CH_4/CO_2 ratio. As the SN value increases with increasing CH_4/CO_2 ratio, the methanol production also increases quickly until the SN reaches the target for methanol production.

Once SN exceeds the target, the increase in methanol production becomes progressively smaller and eventually levels off. This indicates that further increase in SN provide limited additional benefit and can be interpreted as excess hydrogen relative to methanol demand. Together with Figure 6, this result highlights why maintaining SN close to the design value is important: it maximises methanol production while avoiding regimes where additional high-cost H_2 delivers diminishing returns.

CONCLUSION

This work quantified how biogas composition and reformer operating parameters can influence synthesis-gas stoichiometry and methanol production in an integrated biogas-to-methanol process under steady-state conditions. The results show that biogas composition, expressed through the CH_4/CO_2 ratio, is the dominant factor controlling the reformer-outlet stoichiometric number (SN).

Reformer temperature and pressure influence methane conversion and shift SN, but their impact on stoichiometry remains moderate compared with CH_4/CO_2 and becomes less significant as operation approaches high-temperature, equilibrium-limited conditions. When assessed over a bounding design range intended to bracket CO_2-rich to CH_4-rich cases across different sources and possible CO_2 upgrading levels, the process transitions between H_2-deficiency and H_2-excess regimes. Expressing the resulting SN deviation as a correction requirement provides a practical basis for design screening and reveals a crossover near $CH_4/CO_2 \approx 3.21$, where the reformer outlet meets the required SN without adjustment.

Propagating the same composition into the methanol section shows that methanol output improves as SN increases, but the gains diminish once SN exceeds the target value required for methanol production. This indicates that very high SN values mainly create surplus H_2 that is not converted to methanol, rather than delivering proportional increases in methanol production. Overall, the findings emphasise that robust biogas-to-methanol design should be centred on maintaining SN near the synthesis target and sizing conditioning capacity to accommodate plausible feed-composition differences, rather than relying on reformer operating adjustments alone.

Control, energy efficiency, heat integration, hydrogen supply cost, or economic trade-offs were not part of this work. These aspects depend strongly on feed composition and reformer conditions and should be addressed in future work.

ACKNOWLEDGEMENTS

Funded by the European Union

This project received funding from the European Union's Horizon Europe Research and Innovation Programme under grant agreement No 101136225. Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or CINEA. Neither the European Union nor the granting authority can be held responsible for them.

REFERENCES

1. International Maritime Organization. Fourth IMO Greenhouse Gas Study 2020 (2020).
2. Malmborg F. Advocacy coalitions and policy change for decarbonisation of international maritime transport: the case of fueleu maritime. Maritime Transport Research 4:100091 (2023). https://doi.org/10.1016/j.martra.2023.100091
3. Dolan G, Marquez C. Methanol as marine fuel - current status and future development. Marine Engineering 58:510-519 (2023). https://doi.org/10.5988/jime.58.510
4. Brynolf S, Fridell E, Andersson K. Environmental assessment of marine fuels: liquefied natural gas, liquefied biogas, methanol and bio-methanol. Journal of Cleaner Production 74:86-95 (2014).

https://doi.org/10.1016/j.jclepro.2014.03.052
5. Garcia G, Arriola E, Chen WH, De Luna MD. A comprehensive review of hydrogen production from methanol thermochemical conversion for sustainability. Energy 217:119384 (2021). https://doi.org/10.1016/j.energy.2020.119384
6. From TN, Rautenbach M, Partoon B, Østberg M, Wismann ST, Bentien A, Mortensen PM. Understanding the dynamic interplay of CH4/CO2 ratio in electrified reforming of biogas to renewable methanol production: a pilot plant story. Chemical Engineering Journal 518:164662 (2025). https://doi.org/10.1016/j.cej.2025.164662
7. Nyári J, Izbassarov D, Toldy ÁI, Vuorinen V, Santasalo-Aarnio A. Choice of the kinetic model significantly affects the outcome of techno-economic assessments of co2-based methanol synthesis. Energy Conversion and Management 271:116200 (2022). https://doi.org/10.1016/j.enconman.2022.116200
8. Van-Dal ÉS, Bouallou C. Design and simulation of a methanol production plant from CO2 hydrogenation. Journal of Cleaner Production 57:38-45 (2013). https://doi.org/10.1016/j.jclepro.2013.06.008
9. Wismann ST, Engbæk JS, Vendelbo SB, Bendixen FB, Eriksen WL, Aasberg-Petersen K, Frandsen C, Chorkendorff I, Mortensen PM. Electrified methane reforming: a compact approach to greener industrial hydrogen production. Science 364:756-759 (2019). https://doi.org/10.1126/science.aaw8775
10. Nyári J, Magdeldin M, Larmi M, Järvinen M, Santasalo-Aarnio A. Techno-economic barriers of an industrial-scale methanol ccu-plant. Journal of CO2 Utilization 39:101166 (2020). https://doi.org/10.1016/j.jcou.2020.101166

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Strategic Design of CO_2-Reuse Pathways for Sustainable Aviation Fuel: A Game-Theoretic Techno-Economic Analysis

Andrés I. Cárdenas[a], Víctor A. Soria[a], Ana I. Torres[ab]*

[a] Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, United States
[b] Wilton E. Scott Institute for Energy Innovation, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, United States
* Corresponding Author: aitorres@cmu.edu.

ABSTRACT

The aviation sector is difficult to decarbonize due to limits on aircraft electrification, making sustainable aviation fuel (SAF) a critical near-term solution. This study integrates Aspen-based process modeling with game-theoretic optimization to design a multi-agent SAF production network comprising coal gasification and CO_2-assisted natural gas reforming for syngas production, and Fischer–Tropsch (FT) synthesis for SAF production. Techno-economic parameters from Aspen simulations inform an agent-based model in which agents maximize their net present value subject to capacity and demand constraints. Three decision-making frameworks are compared: (i) social welfare optimization, (ii) cooperative bargaining - symmetric (equal bargaining power) and asymmetric (bargaining power weighted by agents' competitiveness outside cooperation), and (iii) competitive equilibria modeled as generalized Nash equilibrium. The results show that social welfare maximization excludes coal and yields the highest total profit, strongly favoring the FT agent's pay-off. Cooperative bargaining includes all agents in the system and promotes CO_2 and water recycling with only a 3.5% profit reduction; symmetric bargaining shifts profit toward coal, while asymmetric bargaining partially restores FT agents' share. Under bargaining, complete CO_2 recycling is achieved, integrating coal without direct emissions. Under strict competition, upstream syngas producers do not participate in the network, and reliance on external syngas reduces overall profit by 12.4%. These results show that game-theoretic modeling for process synthesis reveals strategic incentives for process design that are typically obscured in more traditional superstructure-based optimization frameworks.

Keywords: Sustainable aviation fuel, Optimization, Game Theory, Process Design

INTRODUCTION

Aviation is a major contributor to global greenhouse gas (GHG) emissions due to its intensive fuel use. Future demand for air travel is expected to rise [1], and alternatives such as electrification and sustainable aviation fuel (SAF) have been proposed to mitigate this impact. However, aircraft electrification faces significant challenges due to battery weight and volume constraints. Consequently, developing lower-carbon fuel pathways and efficient SAF production systems is critical.

There are multiple pathways to produce SAFs, all of them involving several processing steps. Optimization has been widely applied to find the best processing sequences under different objective functions (i.e. economic and/or environmental) [2, 3]. While these methods effectively identify optimal configurations, they assume centralized decision-making, overlooking the possibility of splitting the process across multiple industrial actors and the strategic interactions among them. Real-world processing networks involve multiple industries that buy and sell products to one another; these industries often have conflicting objectives, making cooperative and competitive behaviors essential to consider in optimal process and network design.

In this study, we analyze the design of sustainable aviation fuel (SAF) production networks under different decision-making paradigms. We integrate Aspen-based

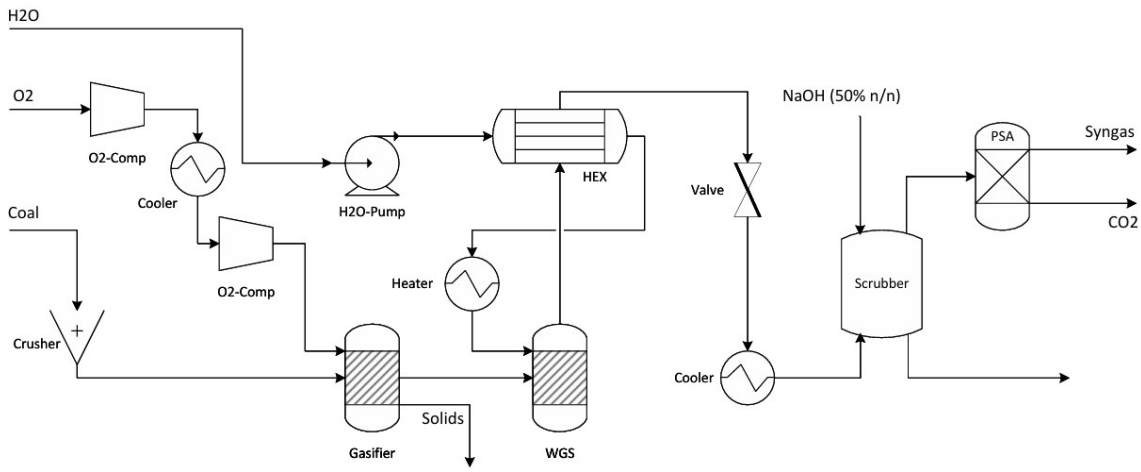

Figure 1: Coal gasification and syngas cleanup process flowsheet.

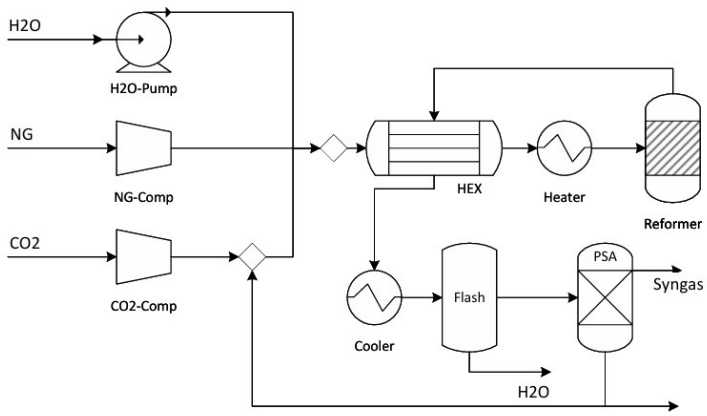

Figure 2: CO_2-assisted natural gas reforming process flowsheet

process simulations with game-theoretic frameworks to capture multi-agent interactions in a system comprising coal gasification, CO_2-assisted natural gas reforming, and Fischer–Tropsch aviation fuel synthesis. By comparing centralized social-welfare optimization, cooperative bargaining, and competitive equilibria, we reveal how these approaches influence system structure, profit allocation, and resource recycling.

This work shows how considering strategic behaviors of agents through game-theoretical models has implications on a SAF network optimal design and resource allocation, managing to find profitable and environmentally preferable solutions through cooperation rather than competition or centralized social-welfare optimization.

ASPEN-BASED PROCESS MODELS FOR SAF PRODUCTION

Three process flowsheets were designed and implemented in Aspen Plus to represent the agents in the SAF network: coal gasification, CO_2-assisted natural gas reforming, and Fischer–Tropsch synthesis of fuels. The process diagrams for each agent are provided in Fig. 1.

Coal gasification

The coal-based syngas route includes crushing, oxygen-blown gasification, H_2/CO ratio adjustment via water-gas-shift (WGS), and contaminant removal. Gasification was modeled using a Gibbs reactor, while WGS was implemented in an equilibrium reactor (REQUIL block) to maintain a H_2/CO ratio of approximately 2 for downstream synthesis. Heat integration between reactor effluents and water feeds improves energy efficiency. Syngas cleanup was modeled using a flash unit with the ELECNRTL property package to capture ionic interactions during NaOH absorption.

Natural gas and CO2 reforming

Reforming was modeled with CO_2 addition to control

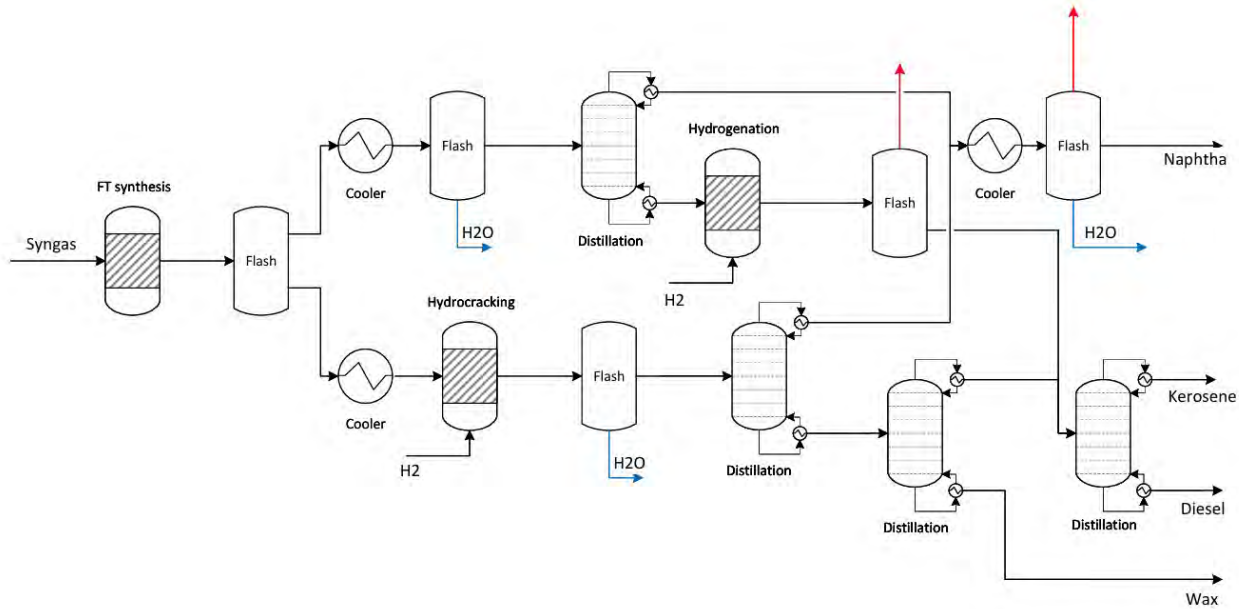

Figure 3: FT route of aviation fuel production process flowsheet. Red flows are light compounds to be burnt for energy recovery.

the H_2/CO ratio [4]. The flowsheet incorporates steam-methane reforming, autothermal reforming, and WGS reactions, with CO_2 recycling from the process to enhance carbon utilization. Natural gas composition includes methane (CH_4), ethane (C_2H_6), and propane (C_3H_8).

Aviation fuel (Fischer-Tropsch synthesis)

Aviation fuel production was modeled using the Anderson–Schulz–Flory (ASF) distribution to predict hydrocarbon chain growth according to Eq. (1) [5].

$$W_n = n(1-\alpha)^2\alpha^{n-1} \qquad (1)$$

Here, n is the carbon chain length, α is the chain growth factor and W_n is the mass fraction of the carbon chain of length n.

The model accounts for paraffin and olefin formation and water generation, ensuring realistic product distribution for jet fuel and co-products. Chain growth and olefin/paraffin ratios are extracted from literature [6].

Process Costing

Equipment and utility costs for each process were estimated using the Aspen Process Economic Analyzer (APEA). Major units such as reactors, heat exchangers, compressors, and separation systems were sized and costed directly from Aspen outputs. Additional data for reformers and gasifiers were taken from relevant literature [7].

MODELS FORMULATION

Agents' objectives and constraints

The payoff of each agent is determined as the Net Present Value (NPV) over a 15-year time horizon at a 15% discount rate. NPV accounts for equipment investment, utilities, feedstock purchases, and product revenues as shown in Eq. (2). Agents invest in capacity at the start of the horizon, with no expansion allowed thereafter.

$$NPV_i = \sum_{t=1}^{15}\left(DF_t \cdot (R_i - U_i - F_i) - C_i^{Cap} \cdot \left(\frac{Cap_i}{Cap_i^0}\right)^{0.6}\right) \qquad (2)$$

Here, $DF = \left(\frac{1}{(1+r)^t}\right)$ is the discount factor for the year t, R_i is the revenue from product sales, U_i is the utility cost, F_i is the feedstock cost, C_i^{Cap} is the cost for a given processing capacity of an agent; we use the 6/10 rule for economies of scale.

Material balances link feed and product flows using conversion factors ($\xi_{a,b}$) where b is the product and a a reference component in the feed; $\xi_{a,b}$ is obtained from ASPEN simulations as shown in Eq. (3). The set E contains all the chemicals that are consumed, produced and traded (coal, natural gas, water, syngas and CO_2).

$$F_{i,b}^{out} = \xi_{a,b}F_{i,a}^{in} \;\; a,b \in E, \forall i \in Agents \qquad (3)$$

Each agent's feedstock usage is limited by their installed capacity (Eq.(4)).

$$F_{i,feedstock}^{in} \leq Cap_i \;\; \forall i \in Agents \qquad (4)$$

Further, yearly aviation fuel demand must be met (Eq.(5)).

$$F^{out}_{AF,AF} = \text{DEMAND} \tag{5}$$

Finally, agents will not participate in a system that yields them negative profit (Eq. (6)).

$$NPV_i \geq 0 \; \forall i \in Agents \tag{6}$$

DECISION-MAKING FRAMEWORKS

To capture strategic interactions among agents, three decision-making paradigms are implemented, each reflecting a different degree of autonomy and conflict among decision-makers.

Centralized - Social welfare

This framework assumes that each agent takes actions that maximize the conglomerate's overall well-being (profit), even if those actions are detrimental to the individual's objectives. The system-level objective is defined as:

$$f^{SW} = \sum_{i \in agents} NPV_i \tag{7}$$

Cooperation - Nash bargain problem

Under this framework, agents remain autonomous but agree to collaborate through negotiation. This cooperation introduces strategic interdependence, as each agent's pay-off depends on the collective agreement. According to Nash's axiomatic framework, the bargaining solution is attained by maximizing the Nash product [8].

$$f^{NB} = \prod_{i \in agents} (NPV_i - d_i) \tag{8}$$

Here, the pay-off of each agent is their NPV, and d_i corresponds to the defection point for that agent. The defection point represents the pay-off an agent would receive by not being part of the bargain and can be interpreted as the bargaining power an agent possesses during the negotiation process.

For this study, two cases are considered: first, defection points are set to zero for all agents, assuming equal footing in the bargain, and that failure to reach an agreement results in no investment in the conglomerate. Secondly, the defection point for each agent is the payoff they receive if the competitive solution is implemented. We will refer to these models as the symmetric and asymmetric bargain problems, respectively.

Both scenarios provide insight into how forced agreements (no cooperation means no investment) and organic cooperation (cooperation occurs only if it is beneficial) shape the structure of the network.

Competition - Generalized Nash Equilibria

Under strict competition, agents act non-cooperatively, optimizing their own objectives while anticipating others' decisions. This creates strategic conflict, as each agent seeks their self-interest within shared constraints.

Since each agent has its own optimization goal, an integrated formulation of the multiple problems can be obtained by reformulating the optimization problems into their respective KKT conditions, yielding a Mathematical Program with Equilibrium Constraints (MPEC). Solving this model yields a Generalized Nash Equilibria (GNE), which is the point where no agent can improve their outcome by unilateral decisions [9].

The multiple optimization problems to be solved share the form presented in Eq. (9). Here $-i$ signifies all agents except i.

$$\left.\begin{array}{ll} minimize_{x_i} & f_i(x_i, x_{-i}) \\ s.t. & g_i(x_i, x_{-i}) \leq 0 \\ & h_i(x_i, x_{-i}) = 0 \\ & G(x_i, x_{-i}) \leq 0 \\ & H(x_i, x_{-i}) = 0 \\ & (x_i, x_{-i}) \in \mathbb{R}^n \end{array}\right\} \forall i \in Agents \tag{9}$$

Here, constraints that belong exclusively to agent i (g_i, h_i) and those shared across agents (G, H) must be clearly distinguished. When reformulating KKT conditions, it is essential to track which constraints are included in each agent's Lagrangian and to correctly associate the corresponding multipliers.

The KKT reformulation for each agent consists of the following components:

First-Order Conditions (FOCs):

$$\nabla_{x_i} f_i(x_i, x_{-i}) + \langle \begin{bmatrix} \rho_i^T \\ \nu_i^T \\ \lambda^T \\ \eta^T \end{bmatrix} \cdot \begin{bmatrix} \nabla_{x_i} g_i(x_i, x_{-i}) \\ \nabla_{x_i} h_i(x_i, x_{-i}) \\ \nabla_{x_i} G(x_i, x_{-i}) \\ \nabla_{x_i} H(x_i, x_{-i}) \end{bmatrix} \rangle = 0$$

Complementarity Conditions (ensuring only active inequalities contribute to the Lagrangian):

$$\rho_i^T \cdot g_i(x_i, x_{-i}) = 0, \lambda^T \cdot G(x_i, x_{-i}) = 0$$

Primal Feasibility:

$$\begin{gathered} g_i(x_i, x_{-i}) \leq 0, h_i(x_i, x_{-i}) = 0 \\ G(x_i, x_{-i}) \leq 0, H(x_i, x_{-i}) = 0 \\ (x_i, x_{-i}) \in \mathbb{R}^n \end{gathered}$$

Dual Feasibility:

$$\rho_i, \lambda \geq 0, \nu_i \in \mathbb{R}^{|\nu_i|}, \eta \in \mathbb{R}^{|\eta|}$$

MPECs are particularly hard to solve since complementarities may not respect constraint qualifications [10]. To overcome this, the problem is relaxed into a feasibility program by defining slacks on the First Order Conditions (ε_{FOC}) and in the complementarity constraints (ε_{Comp}), while smoothing the complementarity with a Fischer-Burmeister approximation (ϕ_{FB}) [10]. The objective function of this program is to minimize said slacks, approximating

the fulfillment of the equilibrium conditions. The complete relaxed problem is shown in equation (9).

$$\left.\begin{array}{c} \min \varepsilon_F + \varepsilon_C \\ s.t: \\ FOC_i(x) \in [-\varepsilon_F, \varepsilon_F] \\ g_i(x) \leq 0, h_i(x) = 0 \\ G(x) \leq 0, H(x) = 0 \\ \phi_{FB}(\rho_i^T \cdot g_i) \in [-\varepsilon_C, \varepsilon_C] \\ \phi_{FB}(\lambda^T \cdot G) \in [-\varepsilon_C, \varepsilon_C] \\ x \in \mathbb{R}^n; \varepsilon_F, \varepsilon_C \geq 0 \\ \rho_i, \lambda \geq 0;\ \nu_i \in \mathbb{R}^{|\nu_i|};\ \eta \in \mathbb{R}^{|\eta|} \end{array}\right\} \forall i \in Agents \tag{10}$$

The Fischer-Burmeister approximation is shown in Eq. (11)

$$\phi_{FB}(a \cdot b) \stackrel{\text{def}}{=} \sqrt{a^2 + b^2 + \delta} - (a + b) \tag{11}$$

Here, δ is a small positive constant to ensure that the derivative of the square root does not get undefined at the origin. The variables ρ_i^T, λ, ν_i and η are the dual multipliers of g_i, G, h_i and H respectively.

APPLICATION TO SAF PRODUCTION PROBLEM

To illustrate the proposed frameworks, an existing decarbonization initiative at Pittsburgh International Airport (PIT) is used as a case study. PIT aims to incorporate locally synthesized sustainable aviation fuel (SAF) into its operations. In alignment with this goal, the case study evaluates the production of syngas from coal and natural gas, readily available resources in Pennsylvania, as feedstock for SAF synthesis.

The system design includes cross-agent resource exchanges, such as the reuse of CO_2 generated during fuel synthesis in the natural gas reforming process, and the integration of water usage between agents. These interactions are modeled under the three frameworks discussed in the previous section.

RESULTS

The previous optimization problems correspond to Non-Linear Programs (NLPs). They are coded in Python using PYOMO [11] as the modelling tool and IPOPT [12] as the solver.

System structure

Figures 4, 5, and 6 illustrate the optimal system configurations under the three decision-making frameworks. The social welfare model favors complete syngas production from natural gas, entirely excluding its production from coal.

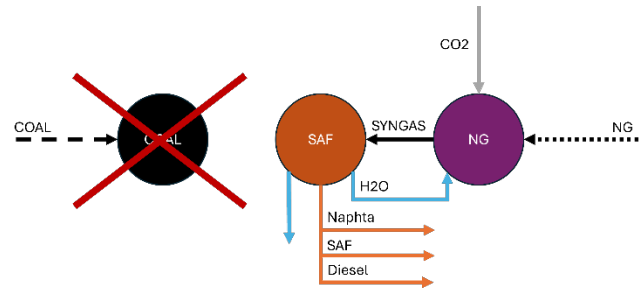

Figure 4. Social welfare system design.

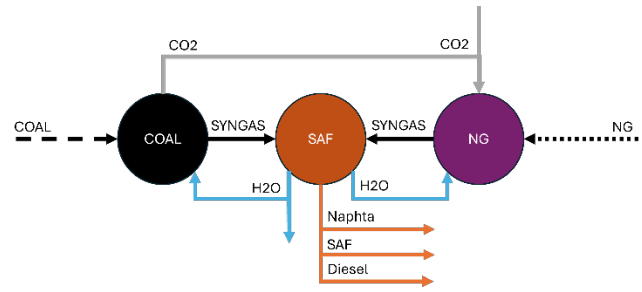

Figure 5. Cooperative system design.

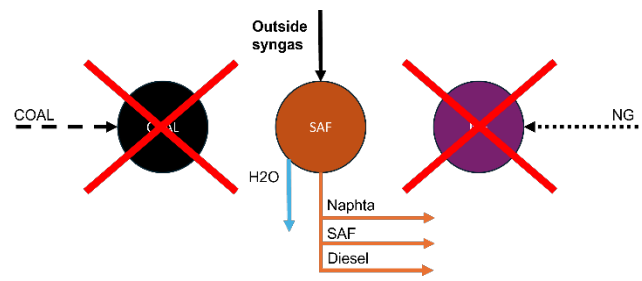

Figure 6. Competitive system design.

Under cooperative bargaining, either symmetric or asymmetric, coal is incorporated into the system. This inclusion reflects the bargaining dynamics: agents seek to maximize their individual payoffs while ensuring mutual benefit, which allows coal to enter the system, as it improves the payoff distribution among agents. Resource integration is also increased in this case, with the agents maximizing the reuse of CO_2 and water. This behavior arises because resource sharing enhances collective gains without compromising individual objectives under the Nash product formulation.

Under the strict-competition assumption, the system consisting only of coal, natural gas, and the SAF agent fails to admit a generalized Nash equilibrium. This indicates that SAF production cannot be sustained using solely these local resources when all agents behave competitively. An equilibrium emerges only when an external syngas supplier is introduced. In this case, the competitive formulation satisfies first-order optimality conditions within acceptable tolerances, and the complementarity conditions exhibit negligible residuals. The additional degree of freedom provided by the external supplier

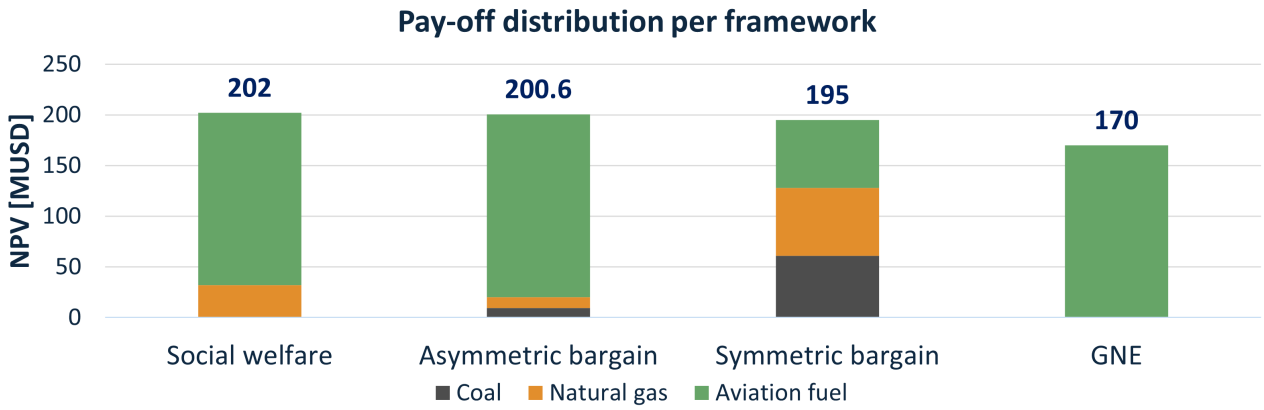

Figure 7. Profit distribution among agents under different frameworks. Social welfare maximizes total system profit and strongly favors the SAF agent. Cooperation (bargaining) redistributes profits, while competition (GNE) reduces overall profit and excludes upstream agents.

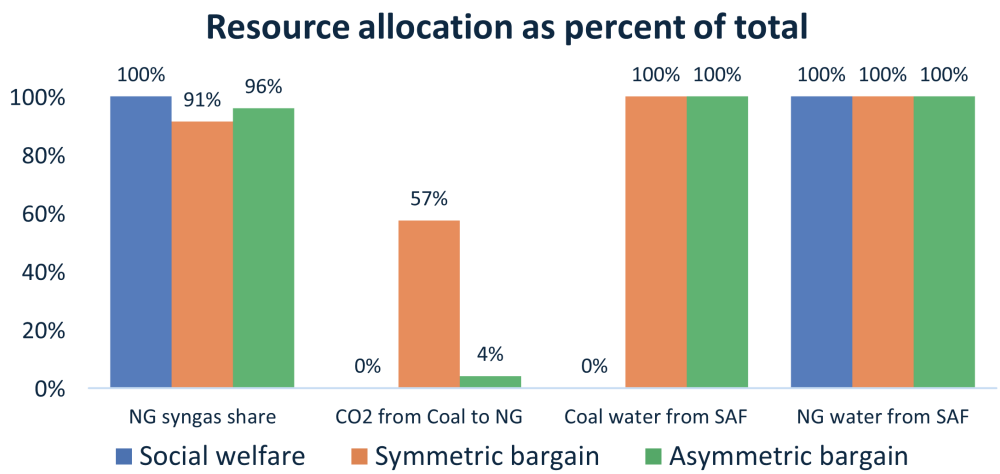

Figure 8. Resource distribution under different frameworks. Bairganing corresponds to the maximum reuse of resources (CO_2 and water) between the agents. Competition eliminates internal exchanges, favoring external sourcing and hence is excluded from the chart.

enables the SAF agent to meet its demand without full dependence on upstream agents. Consequently, the coal and natural gas agents are excluded from the equilibrium solution.

Profit and resource allocation

Figures 7 and 8 present the distribution of profits and resource use among agents under the three modelling frameworks.

The social welfare model achieves the highest overall profit, as expected, since its objective is to maximize the sum of all agents' NPVs, allowing for unfair distributions in which one agent receives most of the profit.

In contrast, the Nash bargaining solution redistributes profits across agents. Symmetrical bargaining reflects equal bargaining power for the agents, achieving an equal split of the profit share. An asymmetric bargain, on the other hand, distributes resources based on how competitive each agent is if no agreement is reached. It is important to note that, even though bargaining reduces the total profit of the system, it does so slightly: 99.3% and 96.5% of the maximum attainable system profit is achieved for asymmetric and symmetric bargains, respectively.

In the competitive solution, at the Nash equilibrium, the SAF agent attains a positive NPV, while coal and natural gas agents are excluded, with syngas sourced entirely from the external supplier. This outcome indicates that, in a fully competitive setting, local coal and natural gas agents lack sufficient competitiveness, suggesting that cooperative or centralized schemes may be required if these local feedstock sources are to be used. It is worth noting that the lowest overall NPV for the system and for the SAF agent is achieved in this competitive setting,

resulting in a 12.4% loss in overall system NPV relative to the social welfare maximum. A SAF agent aware of this outcome would then prefer a social welfare approach (i.e. to integrate the syngas production process within the same firm) or an asymmetric bargain where their competitive status is maintained. The SAF agent will discourage symmetric bargaining, since the cooperative solution provides the lowest payoff.

Resource exchanges also differ significantly across frameworks. CO_2 trading between agents occurs considerably only under symmetric bargaining, while water integration is maximized under social welfare and both symmetric and asymmetric bargaining compared to competition. This behavior aligns with theoretical expectations: competitive agents tend to minimize reliance on internal resource sharing and instead favor external procurement. These findings suggest that bargaining not only redistributes profit but also promotes resource reuse within the system, enhancing overall integration.

An important result is that bargaining enables a full recycle of CO_2 from the coal agent to the natural-gas agent, producing a configuration with no direct CO_2 emissions. In this setting, bargaining not only improves the distribution of profits with negligible loss in total system profit, but also leads to a more environmentally favorable configuration of the whole SAF network.

CONCLUSIONS

This study evaluates three decision-making paradigms, social welfare maximization, cooperative bargaining (with symmetric and asymmetric bargaining power), and strict competition to understand how strategic interactions influence network configuration, profit distribution, and resource sharing in a SAF production system.

Optimization of the social welfare achieves the highest total profit but excludes coal from the system, prioritizing the use of natural gas. Cooperative bargaining enables coal's inclusion and promotes internal resource sharing, generating positive NPVs for all agents with a minor decrease in the system's total profit of the system. However, the SAF agent benefits most from social welfare maximization and asymmetric bargaining. Hence, the SAF agent perceives no incentive to enter an equal-footing bargaining arrangement, whereas upstream agents favor some specific bargains to secure participation and positive payoffs. Coal benefits from any sort of arrangement, while natural gas will only be benefited by a symmetric bargain compared to the centralized social welfare solution.

The competitive solution requires the presence of an external syngas supplier, allowing the convergence to a Generalized Nash Equilibria. However, it excludes the coal and natural gas agents from the system, which is not desirable if the promotion of local resource usage is sought after.

Future work will extend this analysis to consider information asymmetry and coalition formation, offering a deeper understanding of strategic behavior in multi-agent systems.

ACKNOWLEDGEMENTS

This work was supported in part by a seed award from the Scott Institute for Energy Innovation.

AUTHOR IDENTIFIERS

Author ORCIDs:
Cárdenas A: 0009-0008-4423-4104
Soria V: 0000-0001-5742-892X
Torres AI: 0000-0001-8747-3406

REFERENCES

1. Cui Q, Sun XJ, Tang XY, Zhou Y, Zhang YX, Li Y. Altitude-dependent climate impacts and economic feasibility of alternative fuels in aviation from 2025 to 2050. iScience 28:113323 (2025). https://doi.org/10.1016/j.isci.2025.113323
2. Grossmann IE, Harjunkoski I. Process systems engineering: academic and industrial perspectives. Computers & Chemical Engineering 126:474-484 (2019). https://doi.org/10.1016/j.compchemeng.2019.04.028
3. Stephanopoulos G, Reklaitis GV. Process systems engineering: from solvay to modern bio- and nanotechnology.. Chemical Engineering Science 66:4272-4306 (2011). https://doi.org/10.1016/j.ces.2011.05.049
4. Saad JM, Williams PT. Manipulating the h 2 /CO ratio from dry reforming of simulated mixed waste plastics by the addition of steam. Fuel Processing Technology 156:331-338 (2017). https://doi.org/10.1016/j.fuproc.2016.09.016
5. Wang G, Wang YN, Yang J, Xu YY, Bai L, Xiang HW, Li YW. Modeling analysis of the fischer?tropsch synthesis in a stirred-tank slurry reactor. Ind. Eng. Chem. Res. 43:2330-2336 (2004). https://doi.org/10.1021/ie0308083
6. Bube S, Bullerdiek N, Voß S, Kaltschmitt M. Kerosene production from power-based syngas – a technical comparison of the fischer-tropsch and methanol pathway. Fuel 366:131269 (2024). https://doi.org/10.1016/j.fuel.2024.131269
7. S. McNaul, C. White, R. Wallace, T. Warner, H. S. Matthews, J. Ma, M. Ramezan, E. Lewis, "Hydrogen Shot Technology Assessment: Thermal Conversion Approaches, " National Energy Technology

Laboratory, Pittsburgh, December 5, 2023 (Revised).
8. Nash JF join(' '. The bargaining problem. Econometrica 18:155 (1950). https://doi.org/10.2307/1907266
9. Dreves A, Facchinei F, Kanzow C, Sagratella S. On the solution of the KKT conditions of generalized nash equilibrium problems. SIAM J. Optim. 21:1082-1108 (2011). https://doi.org/10.1137/100817000
10. Biegler, Lorenz T. Nonlinear programming: concepts, algorithms and applications to chemical processes. SIAM and Mathematical Optimization Society (2010).
11. Bynum ML, Hackebeil GA, Hart WE, Laird CD, Nicholson BL, Siirola JD, Watson JP, Woodruff DL. Pyomo overview. Springer Optimization and Its Applications :25-36 (2021). https://doi.org/10.1007/978-3-030-68928-5_3
12. Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math. Program. 106:25-57 (2005). https://doi.org/10.1007/s10107-004-0559-y

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Feasibility of Integrating Sugarcane-Derived Biogas into the Allam–Fetvedt Cycle for BECCS Power Generation

João Pedro B. Vasconcelos[a]*, Miguel D. Carvalho[a], and Jean F. Leal Silva[a]

[a] Universidade Estadual de Campinas - UNICAMP, Faculdade de Engenharia Química - FEQ, Campinas, São Paulo, Brazil
* Corresponding Author: joaopedro.borsoni@gmail.com; j218556@dac.unicamp.br

ABSTRACT

The development of energy technologies with low CO_2 emissions is increasingly important for achieving the United Nations Sustainable Development Goals. In this context, power plants based on the Allam-Fetvedt Cycle appear promising because this cycle (introduced in 2012) features inherent CO_2 capture. In this context, its association with biogas as fuel enables its application as a Bioenergy with Carbon Capture and Storage (BECCS) system. This study evaluates the technical, energetic, and economic feasibility of an Allam-Fetvedt Cycle power plant fueled by biogas. The methodology is based on detailed process simulations performed using Aspen Plus® v14. Two biogas production scenarios were assessed: Case 1 and Case 2, corresponding to the processing of 8 and 24 million tons of sugarcane per year, respectively. The economic analysis indicated a high capital investment, primarily in the Air Separation Unit (ASU) and the Balance of Plant (BOP). Nevertheless, a significant reduction in the Levelized Cost of Energy (LCOE) was observed, decreasing from 1443 R$/MWh (€0.227/kWh) in Case 1 to 848 R$/MWh (€0.134/kWh) in Case 2. Overall, the results suggest that integrating biogas into the Allam-Fetvedt Cycle represents a promising technological pathway for power generation with negative carbon emissions. Further research and optimization are required to improve economic viability and strengthen the role of this technology as a strategic option for decarbonizing the energy sector.

Keywords: BECCS, Biogas, Energy Efficiency, Technoeconomic Analysis, Biomethane

INTRODUCTION

Electricity generation is a fundamental basis for the economic and social development of modern society. Conventional Thermal Power Plants (TPP), primarily fueled by coal, account for a significant share of the global energy mix and operate predominantly according to the Rankine cycle. In contrast, natural gas-fired thermal power plants, particularly combined-cycle gas turbines (CCGT), have demonstrated substantial improvements in energy efficiency and environmental performance when compared to coal-fired power plants. Natural gas power plants are relatively more efficient and emit approximately 450 $gCO_2e\ kWh^{-1}$, about half the greenhouse gas (GHG) emissions of coal-fired power plants [1].

Despite this considerable reduction, the challenge of further decreasing power plant GHG emissions remains. In this sense, the next frontier for reducing CO_2 emissions in TPP is its integration with Carbon Capture and Storage (CCS) [2]. Capturing CO_2 from flue gases is challenging because of the abundance of N_2 that comes from the combustion air. Alternatively, oxy-combustion technology in power plants relies on combustion in an oxygen-rich environment, producing a flue gas composed of CO_2 and H_2O, thereby avoiding dilution with N_2. Due to the relatively high CO_2 purity after water is condensed, this CCS technology produces a stream ready for long-term CO_2 sequestration in underground formations, such as saline aquifers and depleted oil and gas reservoirs.

One of the technologies considered for oxy-combustion is the Allam-Fetvedt cycle. This cycle is based on a recuperated Brayton cycle operating at high pressure and uses supercritical CO_2 (sCO_2) as the working fluid under an oxy-combustion regime [2]. Its key distinguishing feature is the recirculation of sCO_2 to the combustion chamber. This recycled sCO_2 plays a dual role: it reduces the combustion flame temperature to a manageable level and dilutes the combustion products [2]. Furthermore,

sCO_2 is non-flammable and non-toxic, low cost, has a density approximately twice that of water vapor, and high volumetric heat capacity [2]. These properties allow a greater amount of energy to be transported per unit volume than water vapor. This allows the use of smaller equipment for the same function, reducing the overall installation footprint and cost while improving cycle efficiency [2]. Modeling studies using natural gas as fuel have indicated that the cycle can achieve efficiencies of up to 60% (LHV – lower heating value – basis), even when accounting for the energy consumption of the Air Separation Unit (ASU) [3].

Biogas, a mixture of about 60% CH_4 and 40% CO_2 produced via anaerobic fermentation of organic residue, emerges as a more sustainable alternative to natural gas. The Brazilian sugarcane industry generates large amounts of vinasse (stillage) as a byproduct of ethanol production from sugarcane. Vinasse has a high organic load and nutrient content, making it a favorable feedstock for biogas production [4]. Given the amount of vinasse currently produced by existing sugarcane mills in Brazil, the sugarcane sector alone has the potential to produce 15 billion Nm^3 of biomethane per year, enough to cover about two-thirds of the domestic natural gas demand [5]. This biogas could be used to increase the presence of the sugarcane sector in the electricity market, which is already significant since sugarcane mills sell surplus electricity produced from bagasse to the national grid. Moreover, biogas utilization in oxy-combustion systems with integral CO_2 capture is fully aligned with the Bioenergy with Carbon Capture and Storage (BECCS) concept. This integration offers a promising pathway to achieve negative emissions in electricity generation by combining sustainable bioenergy sources with CCS technology.

This work investigates the potential of the Allam-Fetvedt cycle using biogas in sugarcane biorefineries to produce renewable electricity. This combination allows biogas to be used without separating CH_4 and CO_2, thereby reducing the cost of biogas upgrading to biomethane. Moreover, it creates the opportunity to store both the CO_2 available in biogas and the CO_2 produced during the combustion of biomethane. To evaluate the potential of this technology, an integrated process was simulated to use the biogas produced as a co-product of ethanol production in sugarcane mills. Based on this simulation, a techno-economic analysis was conducted that considered both the levelized cost of electricity (LCOE) and the Carbon Abatement Cost (CAC).

METHODOLOGY

Allam–Fetvedt Cycle Model for Biogas

The evaluation of energy efficiency and the CO_2 capture in the Allam-Fetvedt cycle fueled by biogas followed a structured methodology, ranging from feedstock characterization to the integrated cycle performance analysis. The biogas yield and composition derived from vinasse were investigated. This analysis enabled defining the fuel flow rate required to supply a power plant operating at a thermal power load of around 200 MW_{th}. This value was established based on metrics from sugarcane processing plants in the State of São Paulo and served as a key assumption for subsequent steps of the study.

Two scenarios were considered for simulation. Scenario 1 (S1) represents a mill with a sugarcane processing capacity of 8 million tons per year, producing biogas for the TPP that uses the Allam-Fetvedt cycle. Scenario 2 (S2) is a hypothetical scenario that integrates three sugarcane mills supplying a single TPP, totaling 24 million tons of processed sugarcane per year. Table 1 presents the production parameters for both scenarios, and Table 2 summarizes biogas properties considered in this work

Table 1: Biogas Production Parameters (1G).

Parameters	Unit	S1	S2
Harvest and operation			
Harvest	$t_{cane}\ y^{-1}$	8, 0 x 10^6	2, 7 x 10^7
Operation	$d\ y^{-1}$	225	225
Mass flow rate [6]	$t_{cane}\ h^{-1}$	1, 481	4, 444
Conventional ethanol process yields			
Bagasse	kg t_{cane}^{-1}	252	252
Ethanol	L t_{cane}^{-1}	85	85
Vinasse	L t_{cane}^{-1}	1, 040	1, 040
Biogas	$Nm^3\ t_{cane}^{-1}$	22.4	22.4
Biogas flow rate			
Volume	$Nm^3\ h^{-1}$	33, 185	99.556
Mass	$t\ h^{-1}$	35.8	107.5
Thermal Power			
Available	MW_{th}	204.7	614.2

Table 2: Biogas properties.

Properties	Unit	Value
Biogas Composition [6]		
CH_4	% mol	65, 2
CO_2	% mol	32, 6
H_2S	ppm	60, 2
Density	$Kg\ Nm^{-3}$	1.08
LHV	$MJ\ kg^{-1}$	20.57
Desulfurization efficiency	%	99

Process modeling followed an iterative approach.

The simulations were carried out using Aspen Plus V14, and model validation was performed by comparing them with simulations using natural gas as fuel. The simulation was developed based on the patents reported in [9] and [10], as well as on previous studies [3, 4, 10]. After the validation step, the model parameters were adjusted to represent biogas as fuel. Given the supercritical operating condition and the mixture composition, the Peng-Robinson Equation of State (EoS) was employed. Previous studies indicate that this EoS provides slightly better performance compared to others, particularly for the calculation of volumetric properties of oxygen-rich mixtures [7]. The CO_2 Purification Unit (CPU) was also modeled using the Peng-Robinson EoS [8]. The main assumptions used in the simulation of the Allam-Fetvedt cycle and its auxiliary units are summarized in Table 3, while the performance indicators are presented in Table 4.

Table 3: Summary of the process assumptions

Assumptions	Unit	Value
Maximum Allowable Turbine Outlet Temperature	ºC	740
Turbine Inlet Pressure	bar	300
Turbine Outlet Pressure	bar	34
O_2 Molar Fraction in the Oxidant	%	13.3
O_2 Excess in Combustion	%	3
Combustion Chamber Inlet Pressure	bar	305
Pressure Drop in the Combustion Chamber	%	1.6

Table 4: Performance Assumptions

Assumptions	Unit	Value
Gas Turbine Isentropic Efficiency	%	90
Compressors Isentropic Efficiency	%	90
Pumps Hydraulic Efficiency	%	85
Mechanical Efficiency of All Turbomachinery	%	98

It is important to note that, in this work, specific operating conditions were defined for two key auxiliary facilities. The CO_2 Purification Unit (CPU) adopted in the model corresponds to systems previously studied and patented in references [11] and [12]. This configuration was selected due to its high CO_2 recovery rate, which can reach approximately 99%. Furthermore, to simplify the analysis, the Air Separation Unit (ASU) was not explicitly modeled. Instead, operating data were obtained from the literature [13] and are presented in Table 5.

For the simulation of the Allam-Fetvedt cycle in Aspen Plus, the use of Design Specifications (Design Specs) and customized nesting settings was required, along with a robust convergence method to handle the high complexity of the system. The Design Specs were applied to control the critical operating parameters of the cycle simulation, namely: 13.3% O_2 concentration, 3% excess O_2 in combustion, turbine outlet temperature of 740ºC, and a CO_2 recovery rate of 99.8%. Design Specs were nested outside, giving priority to stabilization of the mass and energy balance of all recycling loops before solving them [15]. The Broyden method was selected as the convergence algorithm to solve the nonlinear systems of equations associated with multiple Design Specs and recycling loops.

Table 5: Parameters and assumptions of the ASU.

Purity at 120 bar	**99, 5%***	**97%***	**95%***
Electricity demand (kJ_{el}/kgO_2) [13]	1.390	1.320	1.310
Heat recovery from the aftercooler air compressor (kJ_{th}/kgO_2) [14]	1.015	1.015	1.015

Economic Analysis

The economic analysis of the project estimated the Capital Expenditures (CAPEX) and Operating Expenditures (OPEX) using an adapted methodology based on Amec Foster Wheeler guidelines [2], combined with specific studies on biogas production [16]. In both scenarios, the CAPEX for biogas production followed a methodology based on capital cost studies of biogas plants. The cost adjustment was performed using the six-tenths rule, as expressed by Equation (1):

$$CAPEX_{Proposed} = CAPEX_{Base} \times \left(\frac{Q_{Proposed}}{Q_{Base}}\right)^{R} \quad (1)$$

where, $CAPEX_{Base}$ corresponds to the CAPEX at the reference capacity, and Q_{Base} e $Q_{Proposed}$ represent the reference and proposed plant sizes, respectively. The scaling exponent "R" was assumed to be 0, 6, according to [17].

The calculations were structured around three main indicators: Total Equipment Cost (TEC), Total Plant Cost (TPC) and Total Capital Requirement (TCR) [2]. The calculations already include cost updates based on plant cost indices and the necessary currency conversions [2]. The reference cost is presented in Table 6.

Table 6: Cost references for the main equipment of the Allam-Fetvedt Cycle.

Components	Reference Cost [8]
ASU	120.37
CPU	35.20
Turbine	32.55
Heat Exchangers and chillers	0.4
Compressors	21.50
Balance of plant (BOP)	125.82

The Total Plant Cost (TPC) comprises TEC plus additional installation cost and contingency, as defined by [2]. The F factor adopted in this study was 2, as in [2]. The TPC was calculated by Equation (2).

$$TPC = F \times TEC \quad (2)$$

The Lang Fang factor adopted was 4.73, and the location factor was 1.3. The TCR was calculated using Equation (3).

$$TCR = F_{Lang} \times F_{Locality} \times TEC \quad (3)$$

OPEX was classified into fixed and variable costs, based to the established assumptions. The fixed operational and maintenance costs, which do not vary with the production level, include: labor costs, estimated using the with the method presented by Turton et al. 2018 [17], in which the number of operators per shift was calculated by Equation 4; maintenance costs, assumed as 2.5% per year of the TPC for innovative technologies and 1.5% per year of TPC for conventional technologies (BOP) [2]; overhead costs, corresponding to 30% of labor costs, in addition to maintenance costs [2]; and insurance and local taxes, representing 0.5% of TPC [2].

$$N_{OL} = \sqrt{6.29 + 31.7P^2 + 0.23N_{np}} \quad (4)$$

where P is the number of process steps that involve solids and N_{np} is the number of non-particulate steps.

The variable costs depend directly on the production level, as detailed in Table 7. Indices used to update costs and exchange rates are presented in Table 8.

Table 7: Variable costs [2].

Item	Value
Operation at full power	7.884 h y^{-1}
Chemical product costs	0.06 € MW_{th}^{-1}
CO_2 transport and storage	20 € t^{-1}
Makeup water	0.38 US$ GJ^{-1}

Table 8: Indices and exchange rates.

Indicators	Value	Reference
Dollar (USD)	R$ 5.60	[18]
Euro (EUR)	R$ 6.35	
CEPCI 2015	556.8	
CEPCI 2022	816	[19]
CEPCI 2024	799	

To evaluate the performance of the Allam-Fetvedt cycle using biogas as fuel, environmental and economic indicators were applied. The net electrical efficiency (%) was determined considering that the net power output accounts for all auxiliary power consumptions. The calculation was performed using Equation (5).

$$\eta_{net} = \frac{Net\ Electrical\ Power\ Output}{Fuel\ Energy\ (LHV)} \times 100\% \quad (5)$$

The CO_2 capture rate (%) quantifies how much of the CO_2 generated during combustion is effectively captured by the system. It is calculated using Equation (6).

$$TC_{CO_2} = \frac{Massa\ de\ CO_2\ Capturada}{Massa\ Total\ de\ CO_2\ Produzida} \times 100\% \quad (6)$$

The CO_2 emissions indicator quantifies the mass of CO_2 released to the atmosphere per unit of electrical energy generated. The calculation was performed using Equation (7) and expressed as tCO_2 MWh^{-1}.

$$E_{CO_2} = \frac{Uncaptured\ CO_2\ Mass}{Net\ Electrical\ Power\ Generated} \quad (7)$$

The assessment of the amount of CO_2 captured relative to the electrical energy generated was performed using the calculation defined in Equation (8). The result is expressed in tCO_2 kWh^{-1}.

$$CO_{2\ captured} = \frac{Mass\ of\ CO_2\ Captured}{Net\ Electrical\ Power\ Generated} \quad (8)$$

The Carbon Abatement Cost (CAC) represents the additional cost required to reduce one tonne of CO_2 emissions when comparing a capture scenario with a non-capture reference scenario [2]. The calculation was performed using Equation (9), with the final unit expressed as R$ tCO_2^{-1}.

$$CAC = \frac{LCOE_{with\ CCS} - LCOE_{without\ CCS}}{Emissions_{without\ CCS} - Emissions_{with\ CCS}} \quad (9)$$

The LCOE represents the average cost of electricity production over the project's lifespan. It integrates CAPEX and OPEX, both discounted to present value, and divides them by the total electricity generated over the project's lifespan [2]. The calculation is carried out using Equation (10), with the result expressed in R$ MWh^{-1}.

$$LCOE = \frac{\sum_{t=1}^{n} \frac{CAPEX_t + OPEXt}{(1+r)^t}}{\sum_{t=1}^{n} \frac{Energia_Gerada_t}{(1+r)^t}} \quad (10)$$

Other financial parameters are detailed in Table 9, and they are based on [2] and other parameters relevant to the Brazilian sugarcane industry.

Table 9: Other financial parameters.

Economic Parameters	Value
Depreciation	12 y
Discount rate	8%
Inflation rate	Constant
Plant lifetime	25 y
Construction time	3 y
Capacity factor	90%
Years 1 and 2	65% e 85%
Start-up cost	2% TPC + 3 OPEX mo
Decommissioning cost	Zero

RESULTS AND DISCUSSION

To evaluate the influence of different operating conditions and flow rate on cycle performance, scenarios S1 and S2 were simulated. The process simulation flowsheet of the Allam-Fetvedt cycle in Aspen Plus is illustrated in Figure 1, and the CPU process flowsheet is shown in Figure 2. The operating parameters of the various key streams are shown in Table 10. The plant's net energy efficiency is presented in Table 11.

Table 10: Key process results for scenarios S1 and S2.

Parameters	S1	S2
Recycle flow rate (kg/s)	158.6	479.7
Combustion outlet temp. (ºC)	1, 152	1, 149
Turbine flow rate (kg/s)	337.1	1, 019
Turbine cooling flow (kg/s)	26.89	81.36
Turbine outlet temperature (ºC)	734.43	731.99
Turbine outlet pressure (bar)	34.00	34.00
Captured CO_2 flow rate (kg/s)	17.23	51.70
Oxygen flow rate (kg/s)	16.87	50.60

Table 11: Power Block Summary for Cases 1 and 2.

Parameters	S1	S2
Available thermal power (MW_{th})	184.3	555.8
Turbine gross electric power (MW_{el})	157.6	475.2
Total compression power (MW_{el})	27.83	84.15
ASU power (MW_{el})	23.44	70.33
CPU power (MW_{el})	2.29	6.86
Net electric power (MW_{el})	104.0	313.8
Net plant electrical efficiency (%)	56.43	56.47

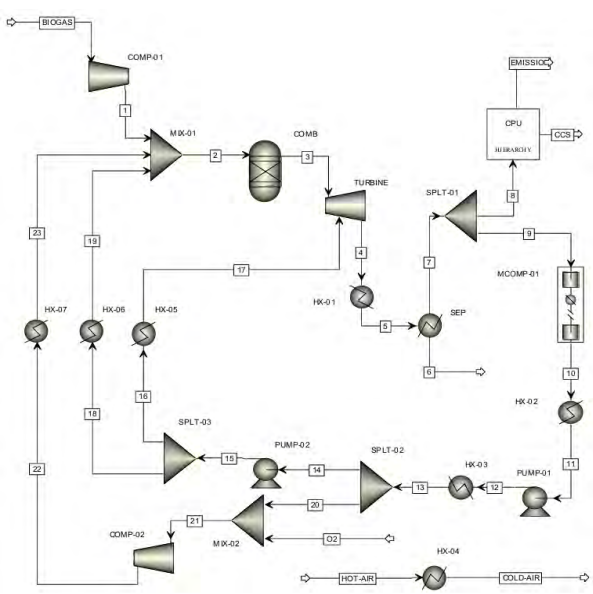

Figure 1. Process flowsheet diagram of the Allam-Fetvedt Cycle in Aspen Plus.

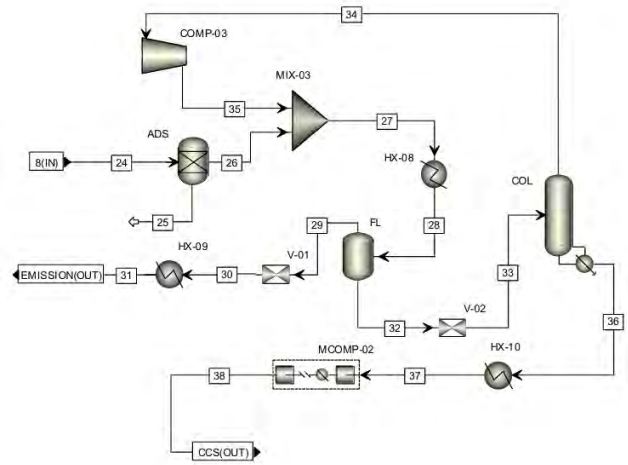

Figure 2. Process flowsheet diagram of the CPU in Aspen Plus.

Economic Analysis

According to [16], the base CAPEX considered was R$ 96.32 million for a biogas production capacity of 3.2 million $Nm^3\ y^{-1}$. Therefore, the calculated capital costs for the different study cases are presented in Table 12. The total CAPEX indicators for the Allam-Fetvedt cycle, including TEC, TPC, and TCR, are summarized in Table 13.

Table 12: Estimated CAPEX and OPEX for biogas production (million R$).

Item	Case 1	Case 2
CAPEX (M R$)	338, 97	655, 30
OPEX (M R$/year)	51, 66	99, 86

Table 13: Summary of CAPEX Indicators for the Allam-Fetvedt Cycle (Million R$).

Components	S1	S2
Estimated TEC for the Allam Cycle (Million R$)		
ASU	447	865
CPU	106	206
Turbine	178	345
Heat Exchangers	38, 4	72, 0
Recycle and O_2 Compressor	175	366
Balance of Plant (BOP)	535	745
Summary		
TEC	1.480	2.599
TPC	2.959	5.197
TCR	9.117	16.013

Based on the TEC breakdown, the ASU and BOP represent the largest portions of equipment cost in both cases. This reflects the complexity and scale of these units. Furthermore, in scenario S2, all values are higher for all components, which is expected due to the increased production capacity.

In Table 14, the OPEX analysis indicates that costs associated with innovative technologies (such as maintenance of the cooled turbine and recuperator heat exchangers) and the BOP are the highest contributors to operational costs in both cases. Also, for both scenarios, the labor costs remain the same. This occurs because Equation (4) considers only the number of equipment units. Finally, the total OPEX for S2 (R$ 442.82 million per year) is almost twice that of S1 (R$ 238.70 million per year), which is attributed to a production capacity three times larger.

Table 14: Results of OPEX (Million R$ y^{-1}).

Item	S1	S2
Labor	1.14	1.14
Technologies	74.0	129
Balance of plant (BOP)	44.4	77.8
Insurance and local taxes	14.8	25.9
CCS	8.83	26.5
Makeup water	3.64	0.91

Table 15 presents the cost metrics and CO_2 emissions in comparison with a reference Natural Gas Combined Cycle Gas Turbine (CCGT) power plant. It is also shown that, for both cases, the LCOE values are higher than those of the reference CCGT. This difference is mainly attributed to the high CAPEX inherent to the implementation of the Allam-Fetvedt cycle and carbon capture systems. On the other hand, a notable reduction in LCOE is observed when moving from S1 to S2, which is a direct result of economies of scale.

Table 15: Economic and environmental metrics.

Parameters	S1	S2
LCOE Biogas (R$ MWh^{-1})	1, 443	848
Specific Cost, TCP/kW (R$ kW^{-1})	28, 453	16, 560
Specific Cost, TCR/kW (R$ kW^{-1})	87, 662	51, 020
CO_2 Emission ($kgCO_2$ MWh^{-1})	1.2	1.2
CAC (R$ tCO_2^{-1})	2, 980	1, 271
Reference (CCGT) [2]		
LCOE Reference (R$ MWh^{-1})	406	
CO_2 Emission ($kgCO_2$ MWh^{-1})	348	

For comparison purposes, Table 16 presents the LCOE and CAC data for other advanced natural gas-based technologies with carbon capture. Regarding LCOE, although the Allam-Fetvedt cycle benefits from scalability, using biogas as fuel introduces additional complexities and costs that outweigh those of other natural gas-based carbon capture routes. Consequently, this explains the significantly higher LCOE of the Allam cycle when compared to the other technologies. A similar pattern is observed for the CAC values, with abatement costs in both scenarios S1 and S2 higher than those of the other carbon capture routes. Nevertheless, an improvement is again observed with increasing scale, highlighting the impact of economies of scale.

Table 16: Comparison of levelized cost of energy (LCOE) and carbon abatement cost (CAC)

Cycle/Case	LCOE (R$ MWh^{-1})	CAC (R$ t^{-1})
CCGT [2]	397	-
NET Power [2]	531	429
SCOC-CC [2]	589	622
Modified S-Graz [2]	595	644
CES [2]	604	673
Biogas/S1 – this work	1.443	2.980
Biogas/S1 – this work	848	1.271

CONCLUSION

The evaluation of energy efficiency and CO_2 sequestration in a thermal power plant operating with the Allam-Fetvedt cycle using biogas as fuel was conducted for two biogas production scenarios of different capacities: S1 (~104 MW_{el}) and S2 (~304 MW_{el}). The economic analysis revealed that the CAPEX of the Allam cycle in both cases is strongly influenced by the complexity and scale of the plants, particularly due to the ASU and BOP. OPEX was also significantly affected by the maintenance requirements of innovative technologies, while labor cost remained constant between the two cases. The comparison of the LCOE of the Allam-Fetvedt cycle using biogas with other advanced technologies also fueled by biogas indicates higher values. Similarly, the CAC analysis shows a high additional cost per tonne of CO_2 avoided, and although a considerable reduction is observed in S2, the CAC remains higher than that of other carbon capture routes. The main distinctive feature of this work lies in the use of a renewable energy source, which positions this technology within the BECCS concept and enables the achievement of net negative emissions, an outcome that fossil fuel-based technologies cannot attain. In summary, although economic viability still requires further optimization and demonstration at a large scale, this work consolidates the technology as a promising, technically validated pathway.

ACKNOWLEDGEMENTS

The authors acknowledge the support from the Human Resources Training Program of the Brazilian National Agency of Petroleum, Natural Gas and Biofuels (ANP) through the PRH-ANP Program (PRH-UNICAMP 29) and the São Paulo Research Foundation (FAPESP, grant numbers 2025/03698-6, 2025/21035-4, and 2025/09414-0).

We also thank EQS – Sustainable Energy and Chemistry Laboratory at the University of Campinas (UNICAMP), Brazil. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

AUTHOR IDENTIFIERS

Vasconcelos, J.P.B. (ORCID: 0009-0003-5533-6230)
Carvalho, M.D.
Leal Silva, J.F. (ORCID: 0000-0002-9915-5859)

REFERENCES

1. Tolmasquim MT. Energia Termelétrica: Gás Natural, Biomassa, Carvão, Nuclear. Rio de Janeiro: Empresa de Pesquisa Energética (EPE), 2016. Coordinated by Mauricio Tiomno Tolmasquim. ISBN 978-85-60025-05-3.
2. IEAGHG - IEA Greenhouse Gas TVL. Oxy-Combustion Turbine Power Plants. Cheltenham, UK, 2015.
3. Scaccabarozzi R, Gatti M, Martelli E. Thermodynamic analysis and numerical optimization of the NET power oxy-combustion cycle. Applied Energy 178:505-526 (2016). https://doi.org/10.1016/j.apenergy.2016.06.060
4. Brandão CM, Stradiotto NR. Comparative economic analysis of the utilization of biogas from sugarcane vinasse in electricity generation, transport fuel and natural gas substitution: a brazilian case study. Waste Biomass Valor 16:2497-2520 (2024). https://doi.org/10.1007/s12649-024-02829-1
5. Fernández DIT, Rodriguez CJC, Machín EB, Pedroso DT, de Carvalho Júnior JA. Brazil's biogas–biomethane production potential: a techno-economic inventory and strategic decarbonization outlook. Biomass 6:4 (2026). https://doi.org/10.3390/biomass6010004
6. Joppert CL, dos Santos MM, Costa HKM, dos Santos EM, Simões Moreira JR. Energetic shift of sugarcane bagasse using biogas produced from sugarcane vinasse in brazilian ethanol plants. Biomass and Bioenergy 107:63-73 (2017). https://doi.org/10.1016/j.biombioe.2017.09.011
7. Li H, Yan J. Impacts of equations of state (EOS) and impurities on the volume calculation of CO2 mixtures in the applications of CO2 capture and storage (CCS) processes. *Applied Energy*, 86, 2760–2770, .
8. Martinelli M, Chiesa P, Martelli E. Techno-economic assessment of the allam cycle for different plant sizes, oxygen purities and heat integration with external sources. Fuel 381:133383 (2025). https://doi.org/10.1016/j.fuel.2024.133383
9. Allam JR, Palmer M, Brown GW. System and Method for High Efficiency Power Generation Using a Carbon Dioxide Circulating Working Fluid. 2011. US Patent US20110179799A1.
10. Allam RJ, Palmer MR, Brown GWJ, Fetvedt JE, Forrest BA. System and Method for High Efficiency Power Generation Using a Carbon Dioxide Circulating Working Fluid. 2013. US Patent US20130213049A1.
11. White V, Allam JR. Purification of carbon dioxide. 2012. US Patent US8257476B2.
12. Strube R, Manfrida G. CO2 capture in coal-fired power plants—impact on plant performance. International Journal of Greenhouse Gas Control 5:710-726 (2011). https://doi.org/10.1016/j.ijggc.2011.01.008
13. Xin T, Xu C, Liu X, Li S, Yang Y. Thermodynamics study of a solar hybrid allam cycle integrated with methane reforming. AIP Conference Proceedings 2303:080008 (2020). https://doi.org/10.1063/5.0035144
14. Beauclair B. New horizons in gasification: the 12th european gasification conference (rotterdam, the netherlands, march 10–13, 2014). Green Processing and Synthesis 3:91-92 (2014). https://doi.org/10.1515/gps-2013-0105
15. Aspen Tech, Inc. Aspen Plus Steady State Simulation User Guide, Volume 2. Version 10.1-0., 1999.
16. Instituto 17. Biogas in Brazil: Economic Feasibility Analysis and Investment Potential. São Paulo, Brazil, 2022. p. 108. ISBN 978 65-997883-9-0.
17. Turton R, Bailie RC, Whiting WB, Shaeiwitz JA, Bhattacharyya D. Analysis, Synthesis, and Design of Chemical Processes. 4th ed. Upper Saddle River, NJ: Pearson, 2018. ISBN 978-0134177403.
18. Brazilian Central Bank (BCB). Currency conversion / Taxas de câmbio. Accessed in Jun 2025. Available at: https://www.bcb.gov.br/en/currencyconversion.
19. Chemical Engineering Magazine. 2024 CEPCI Updates: June (Prelim.) and May (Final). 2024. Available at: https://www.chemengonline.com/2024-cepci-updates-june-prelim-and-may-final/. Accessed Jun 2025.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimising Waste-to-Energy Power Generation in Trinidad and Tobago

Sherard Sadeek[a*], Thomas Hart[a], Arun Mangra[b], Daniel Chernick[a], Andrew Ross[a], Dhurjati Prasad Chakrabarti[b], Maria M. Papathanasiou[cd], Keeran Ward[a]

[a] University of Leeds, School of Chemical and Process Engineering, Leeds, United Kingdom
[b] The University of the West Indies, Department of Chemical Engineering, St.Augustine, Trinidad and Tobago
[c] Imperial College London, The Sargent Centre for Process Systems Engineering, London, Greater London, United Kingdom
[d] Imperial College London, Department of Chemical Engineering, London, Greater London, United Kingdom
* Corresponding Author: S.s.sadeek@leeds.ac.uk.

ABSTRACT

Trinidad and Tobago emits 11.3 metric tonnes of $CO_{2\text{-eq}}$ per capita per year, making it one of the highest per capita per year greenhouse gas (GHG) emitters globally. An estimated 87% of these emissions are linked to the industrial sector, including power generation. This study aims to reduce the national environmental impact of the power generation and waste disposal sectors, with the goal of reducing the country's reliance on natural gas and promoting sustainable power generation via waste-to-energy (WTE) pathways. This work seeks to implement a mixed-integer linear programming (MILP) framework, along with techno-economic analysis (TEA) and life-cycle assessment (LCA) to determine potential sustainable objectives with respect to T&T's power system. The study implements supply chain optimisation considering single objective optimisation (SOO) with life-cycle (LC) endpoint externalities included. The key constraint of the system was the production of sufficient electricity to sustain the national power grid. A WTE model was proposed in which incineration, gasification and anaerobic digestion (AD) were implemented to supplement the country's natural gas based power generation system. When considering LC endpoint externalities, inclusion of WTE along with improved natural gas power generation efficiency resulted in a maximum reduction of 31% in total annualised life-cycle cost (TALC) and 38% in total annualised GHG emissions (TAGE). Thus, indicating the potential of WTE as a key factor in achieving sustainable development goals in SIDS. This work provides a roadmap for the implementation of sustainable power generation practices for policymakers.

Keywords: mixed-integer linear programming, supply chain optimisation, life-cycle assessment, techno-economic analysis, sustainable power generation, waste-to-energy

INTRODUCTION

The Caribbean comprises of 13 island nations- most of which are classed as developing states. Across the Caribbean archipelago, energy is an expensive necessity as most of the region is heavily reliant on fossil fuels in meeting its energy needs. The deployment of clean affordable energy (SDG-7) is integral in combatting the effects of climate change and promoting climate action (SDG-13). A holistic approach is required to aid in this transition, considering all aspects of a robust and sustainable energy system supporting the development of sustainable communities (SDG-11). This transition, however, remains a challenging task as both economic and environmental constraints need to be optimised to guarantee SDGs are met across the wider global population.

Waste-to-energy (WTE) technology provides a potential solution to problems related to fossil fuel power generation in Trinidad and Tobago (T&T). At present, like most developing nations and small island developing states (SIDS), waste disposal in T&T is not considered to be sustainable due to an over-reliance on landfilling and open dumping. This over-reliance on landfilling has resulted in overcrowded landfill sites and is also a source of significant greenhouse gas (GHG) emissions along

with other environmental impacts such as toxic gases, toxins in soils and waterways.

Thus, the introduction of WTE technology can provide a potential solution to two problems evident in T&T and other developing nations. WTE technology can lower both environmental impacts of fossil fuel power generation and disposal of municipal solid waste (MSW).

By utilising a supply chain optimisation framework, guidance on achieving clean and affordable energy through whole systems thinking can be fostered and implemented locally and within the Caribbean region, and among SIDS globally. Hence our work aims to utilise mixed-integer linear programming (MILP) to minimise total annualised costs (TALC) via the use of single objective optimisation (SOO). Environmental indicators are considered in this study via the use of life-cycle externality costs related to endpoint indicator categories. One of the key novelties of the presented work is our MILP framework is tailored to resource availability, local infrastructure, energy system capabilities and waste disposal systems within the Caribbean. This is of critical importance as previous works in this field generally focus on larger and more developed nations.

METHODOLOGY

The model represented in this study focuses on SOO, with monetised impacts for endpoint indicators aligned to human health (HH), eco-system quality (EQ) and resource depletion (RD) being considered in addition to total annualised costs for both power generation and waste disposal. The general superstructure used in this model, which highlights the inputs and outputs of the system along with the technologies utilised, is shown in Figure 1. The scenarios represented in this study generally focus on altering the objective function between minimisation of TALC, minimisation of total annualised GHG emissions (TAGE) and maximisation of power generated from WTE sources.

The model involves the use of binary variables to represent each power plant and WTE technology, thus allowing the model the opportunity to select which power plants and WTE technology can be used. Each plant was constrained to a minimum capacity factor of 0.3. Key variables included the energy required by each plant for both peak and normal times (assumed to be 12 hours each), along with natural gas utilised. The system was constrained to meet the national electricity demand of 8.87 TWh, with a daily peak load of 1350 MW. There are four major classifications of waste considered in the model: organic, paper, plastic and textiles.

Scenario Description

Scenario 1

The Business-as-usual (BAU) case combines the present cases for both power generation and waste disposal in T&T. The model minimises TALC for the present power generation case (two single-cycle (SC) natural gas power plants and two combined cycle (CC) natural gas power plants). There is a total power capacity of 1950MW, 990MW and 960MW for SC and CC, respectively.

The waste collection system involves the movement of MSW from fourteen municipalities to any of three landfill sites. This case only considers landfilling as that is the only major waste disposal platform in the country at present.

Scenario 2

Scenario 2 is split into three subsections (Scenario 2A, 2B and 2C). The general layout and constraints of the model remain constant across all subsections, but the objective function differs. Scenario 2A minimises TALC, Scenario 2B minimises TAGE and Scenario 2C maximises the power generation attributed to WTE technology.

This Scenario considers an upgraded SC plant, resulting in an increased total of 1725MW CC grid power available. The power generated per $tonne_{MSW}$ for each WTE technology is defined in Table 1, and these values are utilised in Scenarios 2 and 3. Additionally, incineration, gasification and anaerobic digestion (AD) are implemented in this case. Scenario 2 also considers the recycling of all rigid plastics.

Scenario 3

Scenario 3 follows a similar formulation to Scenario 2; however, the system is constrained to ensure all organic waste is directed to AD. Scenario 3 considers a total of 1725 MW of CC grid (same as Scenario 2). TALC is minimised in this scenario as was done in Scenario 2A.

Table 1: Power generation for each WTE technology (MWh/$tonne_{MSW}$) (values used for both Scenario 2 and 3).

Waste Type	Incineration	Gasification	AD
Plastic	1.87	0.96	-
Paper	1.69	1.07	-
Textile	1.87	1.27	-
Organic	2.00	1.04	0.08

Model Formulation

Environmental Assessment

A life cycle assessment (LCA) was carried out to primarily examine environmental impacts related to endpoint indicators of human health (HH), ecosystem quality (EQ) and resource depletion (RD). However, GHG is also calculated via the midpoint impact category for global warming potential (GWP). This was done in alignment with T&T's policy focus on decarbonisation. A cradle-to-power generation gate approach was followed using the

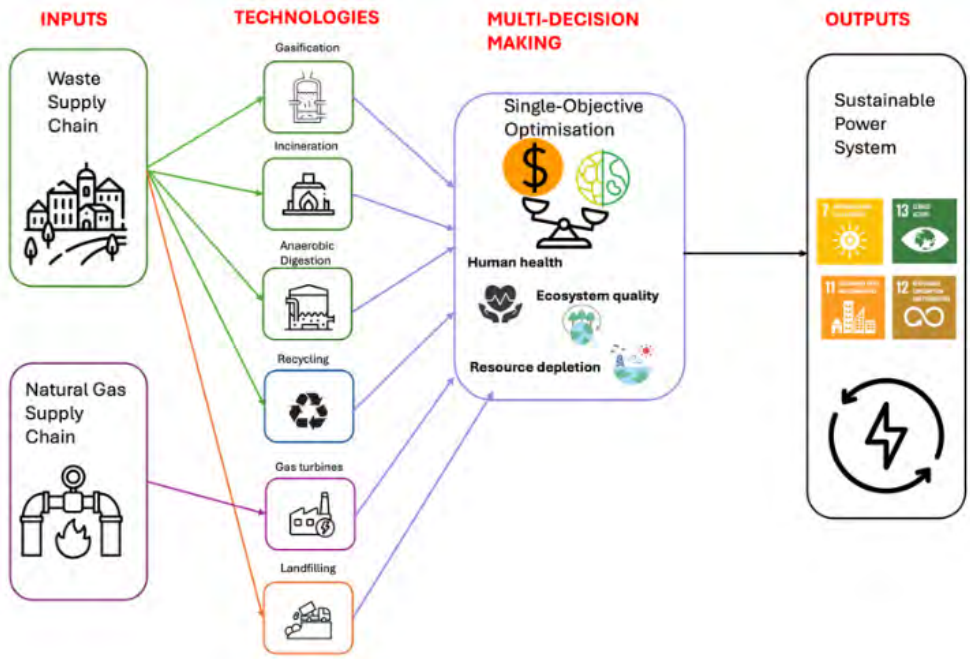

Figure 1. Superstructure highlighting waste disposal and power generation pathways and relevant SDGs impacted.

ISO 14040:2006 methodology [1]. When examining the life cycle impact assessment (LCIA), the ReCiPe 2016 Hierarchist (H) (endpoint) method within SimaPRo software was utilised.

The LCI steps compiled the inputs and outputs for each waste fraction (organic, plastic, paper, textiles) for incineration, gasification, recycling, AD and landfilling. Background inventories were derived from the EcoInvent v3.7.1 database using the Attributional Partly Optimised System (APOS) model. Foreground inventories for gasification were derived from Aspen Plus process models, which provide mass and energy balances for gasification. Inputs to the gasification models include MSW, process water and electricity for operations. This model includes emissions related to transport of MSW to waste technology sites, derived from EcoInvent v3.7.1. Total Annualised GHG Emissions (TAGE) is calculated by:

$$TAGE = \sum(F_{i,k}.e_{F,i,k}) + \sum(W_{i,k}.e_{W,i,k}) + \sum(\omega_{i,k}.e_{\omega,i,k}) + \sum(T_{i,k}.e_{T,i,k}) \quad (1)$$

$F_{i,k}$, $W_{i,k}$ and $\omega_{i,k}$ represent the feedstock, waste and residual flows in the system, with i representing inputs into the system, and k, the waste disposal technologies. The inputs, waste and residual system for transportation is denoted by $T_{i,k}$.The environmental impact of each input, output and residual flow is represented by $e_{F,i,k}$, $e_{W,i,k}$, $e_{\omega,i,k}$ and $e_{T,i,k}$, respectively.

Economic Assessment

The economic assessment is based on the Total Annualised Life-Cycle Cost (TALC), computed by estimating CAPEX and OPEX expenditure based on literature sources. This is inclusive of landfilling, sorting[2], incineration (with energy recovery) [3], gasification[4], plastic recycling [5], anaerobic digestion[6, 7], alongside natural gas power generation[8, 9]. Incineration, AD and recycling costs are estimated based on MSW throughput of the plant (in tonnes), but gasification is estimated based on the cost to run a 10 MW plant. The total costs for power generation via natural gas, waste disposal and transport, and WTE are combined with the externalities attributed to the environmental endpoint impact categories (HH, EQ, RD).

TAC_{Elec} is given by the following equation, accounting for each plant, *j*, and its relevant CAPEX and OPEX, along with the cost, C, of inputs, *i*, into the process:

$$TAC_{Elec} = CAPEX_j + OPEX_j + C_i \quad (2)$$

TAC_{WD} is given by the following equation, accounting for each waste technology, *k*, and its relevant CAPEX and OPEX, along with the cost, C, of inputs, *i*, into the process and the cost of transport in the system ($C_{Transport}$):

$$TAC_{WD} = CAPEX_k + OPEX_k + C_i + C_{Transport} \quad (3)$$

Monetised values for endpoint impact categories, *IC*, are calculated based on each waste fraction, *s*, at each WTE tech, *k*, and for the inputs and production attributed to natural gas power production. The variable $v_{s,k}$ represented the total waste of each type (organic, paper, plastic and textiles) treated by each given technology (incineration, recycling, gasification and AD). $Y_{i,j}$ defines the amount of power produced by each natural gas power

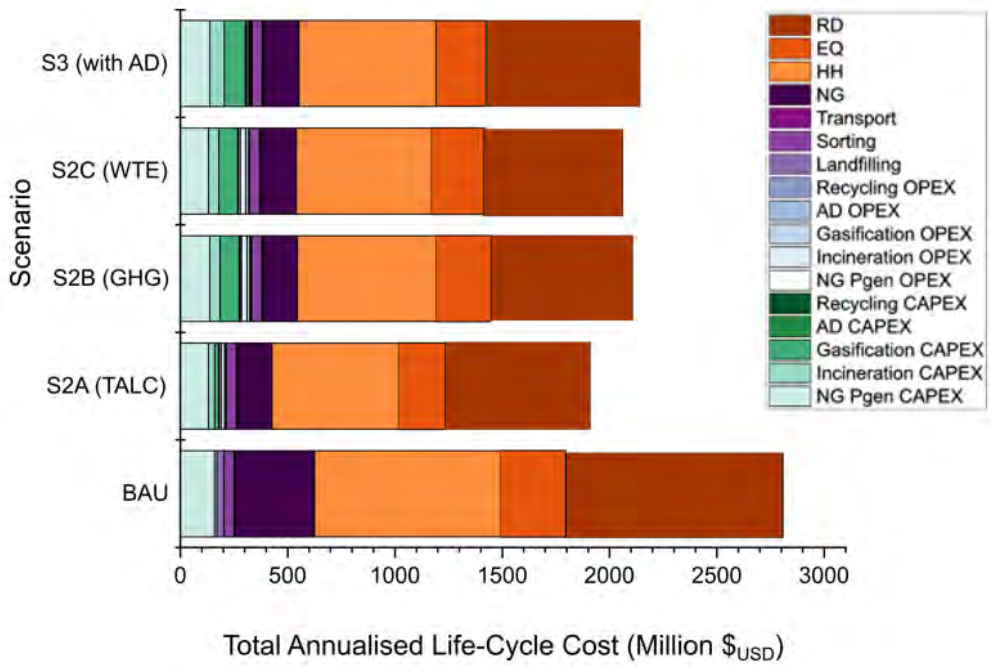

Figure 2. Total Annualised Life-Cycle Costs (TALC) for each Scenario, highlighting effects of varied objective funcitons on overall costs, inclusive of life-cycle externality costs.

plant. Endpoint indicators are related to the three categories, HH, EQ and RD, with $EP_{s,k}$ representing the impact for each waste disposal technology (including WTE) for every characterised waste fraction and $EP_{i,j}$ denoting the impact for natural gas combustion for both SC and CC plants. This is then multiplied by the monetisation factors for each category, Mon_{IC}, and the following equation is given:

$$LC_{IC} = \left[\sum_{s,k}(EP_{s,k} \,.\, v_{s,k}) + \sum_{j}(EP_{i,j} \,.\, Y_{i,j})\right].Mon_{IC} \quad (4)$$

The summation of TAC for both waste disposal and WTE technology (TAC_{WD}), natural gas power generation (TAC_{Elec}) and monetised life-cycle costs, LC_{IC}, is given in Equation 4:

$$TALC = TAC_{WD} + TAC_{Elec} + LC_{IC} \quad (5)$$

Constraints

The following equations, 6, 7 and 8, ensure the HH, EQ and RD values cannot surpass the BAU levels. This ensures no burden shifting occurs to artificially decrease TALC, by trading off performance in one endpoint indicator for another.

$$HH_{Scenario} \leq HH_{BAU} \quad (6)$$

$$EQ_{Scenario} \leq EQ_{BAU} \quad (7)$$

$$RD_{Scenario} \leq RD_{BAU} \quad (8)$$

Objective Function

The SOO model minimises TALC (Equation 2) for Scenario 2A, minimises TAGE for Scenario 2B (Equation 1) and maximises WTE for Scenario 2C.

Model Environment

The MILP model was simulated using the General Algebraic Modeling System (GAMS) version 34.3 on an Intel Core i7, 8th Gen. Processor (16 GB RAM).

RESULTS

The model primarily examined Total Annualised Life-cycle Costs (TALC), which combined overall costs related to the system and life-cycle externality costs related to HH, EQ and RD. Figure 2 highlights the TALC values across 3 scenarios. Scenario 1 represented the BAU case, whilst Scenario 2 highlighted an upgraded efficiency in natural gas power generation combined with WTE in the form of gasification, incineration and AD. Scenario 3 constrained all organic waste to be dispatched to AD. Scenarios 2 and 3 also included recycling of rigid plastics as part of the waste disposal system.

The TALC for the BAU case was found to be $2.74 billion, which included 31.3% for HH, 11% for EQ and 36.7% for RD indicators, respectively. Production of power from natural gas plants accounted for approximately 19% of the TALC and waste disposal via landfilling (including transport) accounted for the remainder.

In comparison to the BAU case, the optimised

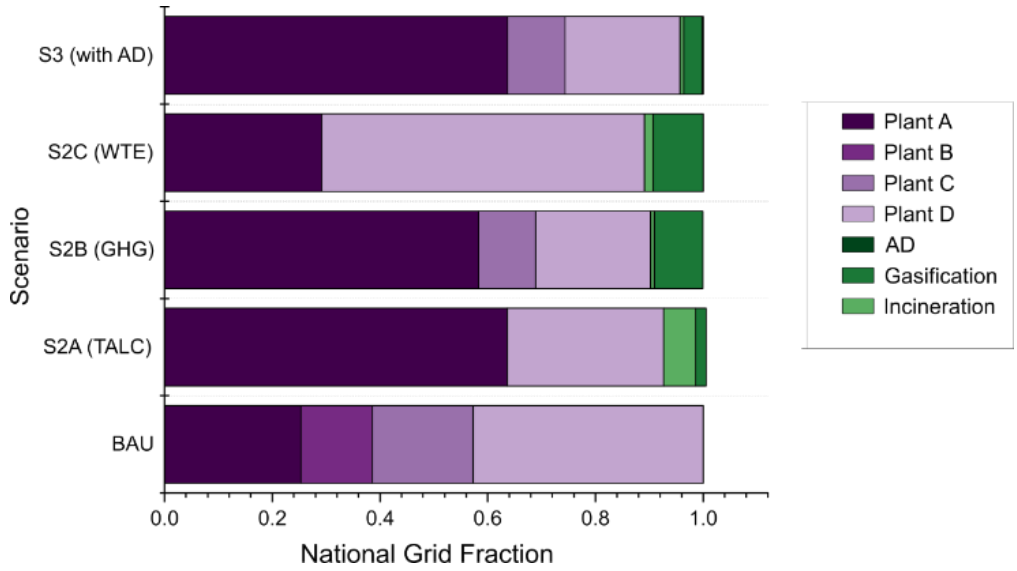

Figure 3. Power grid breakdown for each scenario exhibiting fossil-fuel and WTE sources.

(minimised TALC) values for Scenario 2A highlighted a reduction of 31.1% ($1.93 billion TALC). Externality costs attributed to HH, EQ and RD reduced by 31.6%, 29.9% and 33% compared the BAU values, respectively. Costs related to power generation accounted for approximately 16% of the TALC, with recycling, sorting, incineration and gasification representing the remaining 6% of TALC. The WTE from incineration and gasification represents 7.4% of the national grid presently. Due to the increase in power to supply the grid from WTE, only Plant A (765MW – CC) and Plant D (720MW – CC) were required and accounted for 8.2 TWh a year.

When the TAGE of the system was minimised for Scenario 2B, it was noted that TALC increased by 10% to $2.10 billion. This trade-off can be attributed to a 9.4% increase for the HH indicator and an 18.7% increase for EQ whilst there was subsequent decrease in RD of 2.5%, from Scenario 2A to 2B. The model reduced TAGE by increasing the proportion of waste allocated to gasification, thus avoiding potential TAGE emissions by increasing the WTE value to 9.7% of the national grid. Since GHG is prioritised, gasification is preferred to incineration. Production of electricity from gasification results in a reduced GHG emission value for each characterised waste fraction type (organic, paper, plastic and textiles). This increased use of gasification enhanced WTE output to the national grid as it yields approximately double the electricity per $tonne_{MSW}$ compared to incineration.

Scenario 2C explores the maximisation of WTE as the objective function of the model. This gives a TALC of $2.07 billion and a WTE accounting for 10.9% of the national grid. When compared to Scenario 2A, there is a 6.5% increase for the HH indicator and a 16% increase for EQ. There was a subsequent 3.14% decrease in RD. Across Scenarios 2A, 2B and 2C, it can be noted that as WTE increases RD decreases, which can be attributed to the reduced fossil fuel requirement to supply the national grid.

Scenario 3, which constrains all organic waste to be sent to AD, presents a TALC of $2.15 billion which represents a 10% increase when compared to Scenario 2A. The HH indicator only improves upon Scenario 2B and is the worst performing optimised scenario in terms of RD, recording a 5.5% increase compared to Scenario 2B. The increased HH and RD indicators can be attributed to the reduced WTE value of 4.3%. The outcome reflects the reduced power generation potential of AD in comparison to incineration and gasification. The use of AD and thus reduced gasification capacity in this scenario results in a reduction in EQ of (5.5 – 7.5) %. Despite being the worst performing of the optimised scenario with respect to TALC, the introduction of AD still exhibits a 22.4% reduction in comparison to the present BAU case and proves an improvement in terms of EQ when compared to scenarios 2B and 2C. This indicates that AD can be included as a potential WTE option in T&T, alongside gasification and incineration.

Minimisation of TAGE results in 38% GHG_{LC} reduction compared to the BAU case, with a reduction from 6.26 million tonne-$CO_{2\text{-eq}}$ to 3.86 million tonne-$CO_{2\text{-eq}}$. The model achieves this by significantly increasing the waste sent to gasification, from 94, 000 tonnes (S2A – min. TALC) to 426, 000 tonnes and reducing incineration, from 445, 000 tonnes to 47, 000 tonnes. This trade-off is due to gasification's higher electricity production potential compared to incineration. Natural gas power

generation is reduced by 2.5% with 3 combined-cycle power plants online, Plant A, Plant B and Plant C, and sharing the national grid load 58%, 11% and 21% respectively (Figure 3). Power produced from WTE sources accounts for 9.7% of the national grid. Overall, natural gas power generation accounts for 92% of the GHG emissions in the scenario, with incineration and gasification making up 7.5% and transport resulting in a negligible value.

The TAGE for this Scenario 2A (min. TALC) was found to be 4.05 million tonne-$CO2_{eq}$, signifying a reduction of approximately 35% compared to the BAU case and a 5% increase compared to Scenario 2B (min. TAGE). WTE sources account for approximately 9.8% of TAGE, with gasification making up 8.3% and incineration 1.5% of TAGE. Plant A, Plant D, incineration and gasification make up over 99% of TAGE for optimised cases.

CONCLUSIONS

We present the use of MILP framework tailored to Trinidad and Tobago to promote the sustainable energy transition within SIDS. This work demonstrates the potential of WTE as a cleaner power generation source within the Caribbean. Additionally, upgraded efficiency within the country's natural gas power generation sector can play a significant role in reducing environmental emissions.

The integration of waste-to-energy (WTE) in the form of gasification, incineration and anaerobic digestion (AD) proves to have a beneficial effect in terms of total annualised life-cycle costs (TALC) with reductions of 22 – 31 % for optimised scenarios, when compared to BAU. Similar effects are observed for total annualised GHG emissions (TAGE), with decreases in the range of 34 – 38 % in comparison to the BAU scenario. The scenarios are optimised using SOO, accounting for environmental emissions related to endpoint indicators which are monetised. Drastic reductions were recorded for all optimal scenarios due to the diminished use of natural gas to provide power. This reduction is attributed to increased resource efficiency from the SC to CC upgrade, coupled with 4 – 11 % of the grid being supplied by WTE.

These results validate the use of supply chain optimisation to guide decision making with respect to SDGs for the Caribbean and SIDS globally.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the following companies for their assistance with regards to data acquisition for this work: Power Generation Company of Trinidad and Tobago (Powergen), Trinidad Generation Unlimited (TGU) and Solid Waste Management Company Limited of Trinidad and Tobago (SWMCOL).

REFERENCES

1. International Organization for Standardization, ISO. 2016. Environmental Management - Life Cycle Assessment- Requirements and Guidelines.
2. Olafasakin O, Ma J, Bradshaw SL, Aguirre-Villegas HA, Benson C, Huber GW, Zavala VM, Mba-Wright M. Techno-economic and life cycle assessment of standalone single-stream material recovery facilities in the united states. Waste Management 166:368-376 (2023). https://doi.org/10.1016/j.wasman.2023.05.011
3. Kumar A, Samadder SR. A review on technological options of waste to energy for effective management of municipal solid waste. Waste Management 69:407-422 (2017). https://doi.org/10.1016/j.wasman.2017.08.046
4. IRENA. Renewable Energy Technologies: Cost Analysis Series. Biomass for Power Generation. (2012)
5. Bohm RA, Folz DH, Kinnaman TC, Podolsky MJ. The costs of municipal waste and recycling programs. Resources, Conservation and Recycling 54:864-871 (2010). https://doi.org/10.1016/j.resconrec.2010.01.005
6. Fei X, Jia W, Chen T, Ling Y. Life cycle assessment of food waste anaerobic digestion with hydrothermal and ionizing radiation pretreatment. Journal of Cleaner Production 338:130611 (2022). https://doi.org/10.1016/j.jclepro.2022.130611
7. Slorach PC, Jeswani HK, Cuéllar-Franca R, Azapagic A. Environmental sustainability of anaerobic digestion of household food waste. Journal of Environmental Management 236:798-814 (2019). https://doi.org/10.1016/j.jenvman.2019.02.001
8. Bhagaloo K, Ali R, Baboolal A, Ward K. Powering the sustainable transition with geothermal energy: a case study on dominica. Sustainable Energy Technologies and Assessments 51:101910 (2022). https://doi.org/10.1016/j.seta.2021.101910
9. Bhagaloo K, Baboolal A, Ali R, Razac Z, Lutchmansingh A, Mangra A, Muhammad T, Ward K. Resource efficiency as a guide to clean and affordable energy: a case study on trinidad and tobago. Chemical Engineering Research and Design 178:405-420 (2022). https://doi.org/10.1016/j.cherd.2021.12.026

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Harnessing waste heat in the optimal operation of power-to-X energy systems using detailed process models

Yifan Wang[a], Luka Bornemann[b], and Niklas von der Assen[a]*

[a] Institute of Technical Thermodynamics, RWTH Aachen University, Schinkelstr. 8, 52062 Aachen, Germany
[b] Institute of Environmental Technology and Energy Economics, Hamburg University of Technology, Eissendorfer Strasse 40, 21073 Hamburg, Germany
* Corresponding Author: niklas.vonderassen@ltt.rwth-aachen.de.

ABSTRACT

Power-to-X (PtX) technologies play a central role in renewables-based energy systems by enabling the conversion of renewable electricity into multiple energy carriers. However, due to the multiple energy conversion stages inherent to such energy systems, they often suffer low system efficiencies and high operational costs. In this context, waste product utilization offers significant potential for improving system performance. Directly integrating waste product utilization into energy system operational problems, however, is computationally challenging, as it requires high model granularity to capture waste product characteristics and introduces additional complex constraints.This work proposes a method to integrate waste heat utilization into operational optimization problems, aiming to improve the overall performance of PtX energy systems. Detailed process models, together with pinch analysis, are used to generate surrogate models for the thermal (by-)products and their associated temperature levels. The resulting optimization problem is decomposed into multiple subproblems. To balance problem complexity and solution optimality, two alternative approaches, “ex-post” and “co-optimized” are developed, differing in whether waste heat utilization is addressed ex post or jointly optimized with the energy system operation.The proposed method is applied to optimize the daily operation of a PtX energy system. Compared with a reference case without waste heat utilization, the ex-post approach achieves a cost reduction of 7% within identical computational time (0.6h). The co-optimized approach yields cost savings of up to 22% with a computational time of 3.4h. These results demonstrate the effectiveness of the proposed method for integrating waste heat utilization into energy system operational optimization.

Keywords: PEM electrolyzer, direct air capture plant, methanation process, MIQCQP, operational optimization, NLP, decomposition, waste heat utilization

1. INTRODUCTION

Global efforts to mitigate the impacts of the climate crisis and to reduce dependence on fossil resources are driving the development of innovative and sustainable energy systems. Among various alternatives, power-to-X (PtX) technologies have emerged as a promising solution, enabling the conversion of renewable electricity into alternative energy carriers and chemical products. When considered in a broader context, energy systems with integrated PtX technologies (PtX energy systems) offer significant advantages in terms of operational flexibility, decarbonization, and the provision of seasonal energy storage and supply.

However, as various process technologies and multiple energy conversion steps are integrated, the operation of PtX energy systems face significant challenges, namely low overall energy conversion efficiency and the resulting high operational costs. To improve the overall system efficiency, waste product utilization poses a promising solution [1, 2].

Nevertheless, integrating waste product utilization into the operational optimization problems for highly integrated energy system is a complex task. Generally, these optimization problems involve temporal coupling constraints across the entire operational time horizon

and are therefore usually computational intense. Moreover, the inclusion of complex process technologies often necessitates a high level of model granularity to ensure their operational feasibility, thereby introducing nonlinearities, nonconvexities, and, in some cases, dynamic algebraic equations into the optimization problem (challenge a). In addition, incorporating waste product utilization further increases problem complexity by introducing additional constraints that couple different product balances, rendering the resulting optimization problem highly nested and computationally intractable (challenge b).

Wang et al. [3, 4] developed a decomposition-based optimization method and successfully addressed the first complexity (challenge a). By incorporating detailed process models into energy system operational problems, the decomposition method successfully harmonizes the dynamic operation of a diverse set of process technologies within a PtX energy system. While high-fidelity process models ensure the feasibility of the resulting operational strategies, computational tractability is maintained.

Nevertheless, despite the optimized feasibility and flexibility, previous works have not fully exploited the potential of integrating detailed process models into energy system optimization. Specifically, auxiliary thermal (by-)products of different process technologies remain largely decoupled, often necessitating simultaneous and unnecessary external hot and cold streams purchases.

1.1 Contribution of this work

This work further advances the decomposition method to enable waste heat utilization in energy system operational optimization (challenge b). Built on prior work, we incorporate pinch analysis into the decomposition method. Correspondingly, the waste heat can be decomposed into thermal products with different temperature levels. As additional complicating constraints are introduced into the optimization problem, which could increase computational burden, we also develop two alternative solution approaches to balance the computational efficiency and optimality. In these approaches, waste heat utilization is either addressed ex post or jointly optimized with the energy system operation.

This work is structured as follows: The operational problem for a PtX energy system is stated in Section 2. Section 3 details the modeling of the optimization problem, while Section 4 describes the proposed solution approaches. The computational results are discussed in Section 5. Section 6 concludes the work.

2. PROBLEM STATEMENT

We adopt the PtX energy system from literature [4] to state our problem, as illustrated in **Figure 1.**

The energy system aims to fulfill electricity (el-) and heating (th-) demands of an office building via energy conversion or direct purchase from the grids (cf. the blue background in **Figure 1**). The th-demand for the office building typically requires a supply temperature at a medium temperature level (MT). A total of twelve energy system components are included: seven conversion components—photovoltaic panels (PVs), a proton exchange membrane (PEM) electrolyzer, a direct air capture (DAC) plant, two identical methanation processes (MPes), a boiler and a combined heat and power (CHP) engine—and five storage components for the products electricity, H_2, CO_2, CH_4 and heat, respectively.

The PtX technologies introduce diverse process characteristics into energy system operation (cf. the green background in **Figure 1**). The DAC plant operates as an adsorption-desorption batch process and requires an external hot stream at high temperature level (HT) after adsorption to raise the column temperature for the subsequent desorption step. The PEM electrolyzer is governed by electrochemical reaction, with conversion efficiency strongly influenced by operating temperature and pressure. The MPes involve strongly exothermic reactions, for which the catalyst performance is constrained by reaction temperature. Consequently, an external cold stream at low temperature (LT) level is required for the PEM electrolyzer and MPes to maintain their respective operating temperatures.

In summary, three thermal products at different temperature levels, HT, MT and LT, are involved in the studied PtX energy system. To date, these thermal products are treated independently during operation, which prevents the exploitation of thermal integration potential and leads to inefficient operational decisions. Therefore, the goal of this work is to develop modeling and solution approaches that enable the integration of waste heat utilization into the operational optimization for the PtX energy system (cf. the red background in **Figure 1**).

3. POWER-TO-X ENERGY SYSTEM MODELING

The modeling of the PtX energy and process systems (cf. the blue and green background colors in **Figure 1**) are adopted from our previous work [3–5]. For completeness, an overview of these models is provided in Section 3.1. This work further extends the modeling to address the heating circuit, as detailed in Section 3.2.

3.1 Energy and process systems modeling

We collect all twelve energy system components in the set $\mathcal{K}$ and model each component k ($k \in \mathcal{K}$) with its input-output relationship(s):

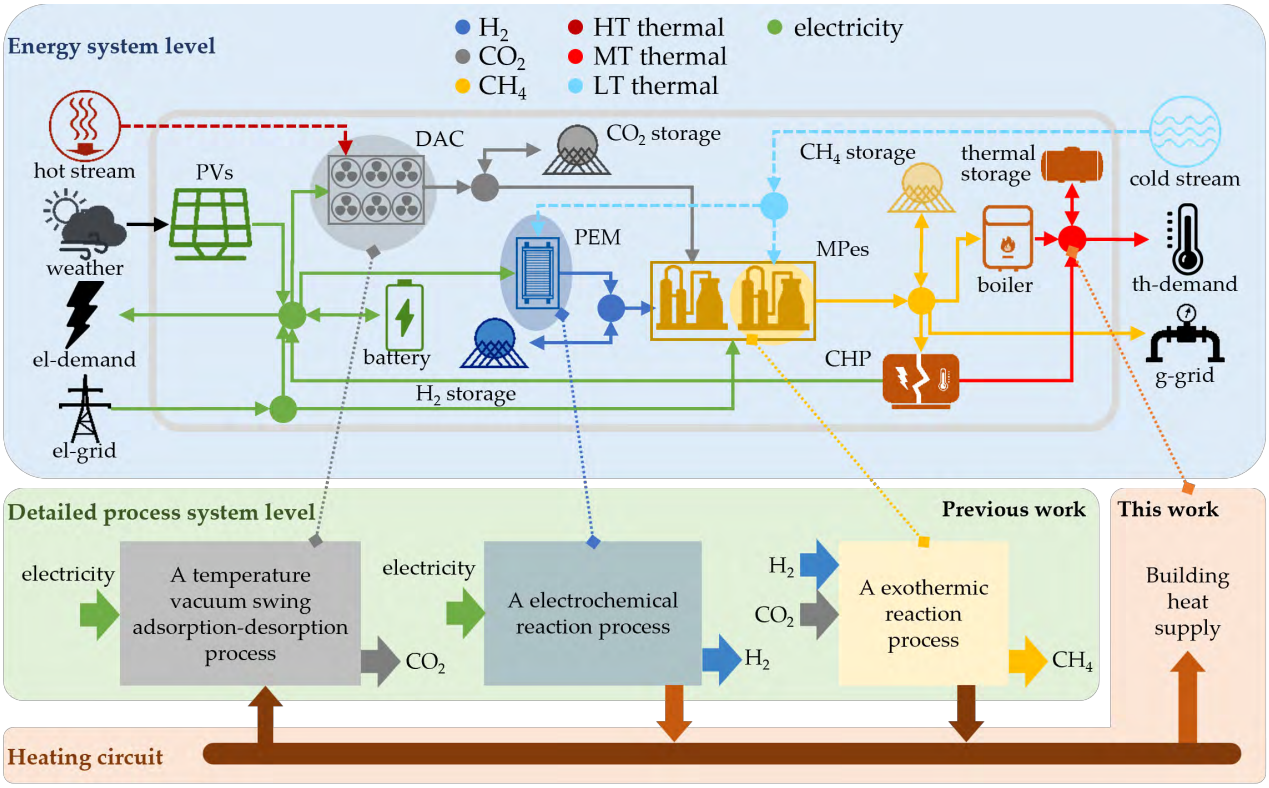

Figure 1. Structure of a power-to-X energy system, adapted from [4]. The complete energy system is indicated by a blue bacground, while the components that are modeled using detailed process models are highlighted with a green background. The red background represents the heating circuit for potential waste heat utilization proposed in this work. PVs: photovoltaic panels; DAC: direct air capture plant; PEM: proton exchange membrane electrolyzer; MPes: two identical methanation processes; CHP: a combined heat and power engine; HT, MT and LT: high, medium and low temperature levels, respectively; th: heating; el: electricity.

$$\dot{E}^{\text{out}}_{k,\tilde{e},t} = f(\dot{E}^{\text{in}}_{k,e,t}, \boldsymbol{x}^{\text{ES}}_k, \boldsymbol{y}^{\text{PS}}_k), \tag{1}$$

where $\dot{E}^{\text{out}}_{k,\tilde{e},t}$ and $\dot{E}^{\text{in}}_{k,e,t}$ denote the input and output product flows of component k at time step t ($t \in \mathcal{T}$), respectively. The subscripts e and $\tilde{e}$ indicate the product types ($e, \tilde{e} \in \mathcal{E} = \{\text{el}, \text{H}_2, \text{CO}_2, \text{CH}_4, \text{th}_{\text{LT}}, \text{th}_{\text{MT}}, \text{th}_{\text{HT}}\}$). $\tilde{e}$ and e are different if k is a conversion component and are identical if k is a storage component.

The surrogate vector $\boldsymbol{x}^{\text{ES}}_k$ collectively represents other decision variables of component k at the energy system level, such as on/off status of conversion components, charge/discharge decisions and the state of charge of storage components. These decision variables generally have a direct impact on the operational decisions of other energy system components. The surrogate vector $\boldsymbol{y}^{\text{PS}}_k$ denotes other decision variables of component k at the detailed process system level, such as the operating temperature and pressure of a reactor or the outlet temperature of a heat exchanger within the process of k. These decision variables ensure the feasibility of k and usually do not have a direct impact on other energy system components ($\mathcal{K}\backslash\{k\}$).

Note that, for notational convenience, this work adopts the following convention: Given a decision variable indexed over multiple sets, say $x_{k,e,t}$ with $k \in \mathcal{K}, e \in \mathcal{E}, t \in \mathcal{T}$, the bold notation $\boldsymbol{x}_k$ or $\boldsymbol{x}$ represents the vector containing all $x_{k,e,t}$ over subsets or the entire index set. For example, $\dot{\boldsymbol{E}}^{\text{in}}_{k,e}$ denotes $\dot{E}^{\text{in}}_{k,e,t}, \forall t \in \mathcal{T}$; $\dot{\boldsymbol{E}}^{\text{in}}$ indicates $\dot{E}^{\text{in}}_{k,e,t}, \forall k \in \mathcal{K}, e \in \mathcal{E}, t \in \mathcal{T}$.

Depending on the inherent operational complexity of component k, a detailed process model might be unnecessary. We collect components that require a detailed process model in the set $\mathcal{K}^{\text{PS}} \subseteq \mathcal{K}$. This work considers four key PtX technologies with detailed process models, i.e., $\mathcal{K}^{\text{PS}} = \{\text{PEM}, \text{DAC}, \text{MPes}\}$, where MPes collectively represents the two identical methanation processes. For all other components ($k \in \mathcal{K}\backslash\mathcal{K}^{\text{PS}}$), Eq. (1) generally takes a linear or quadratic form, with or without binary variables, and the vector $\boldsymbol{y}^{\text{PS}}_k$ is empty.

PEM electrolyzer. The detailed process model of the PEM electrolyzer is adopted from [5]. It comprises heat exchangers, electrolysis cells, flash units, a multi-stage compressor, all described by detailed electrochemical, heat and mass transfer models. Accordingly,

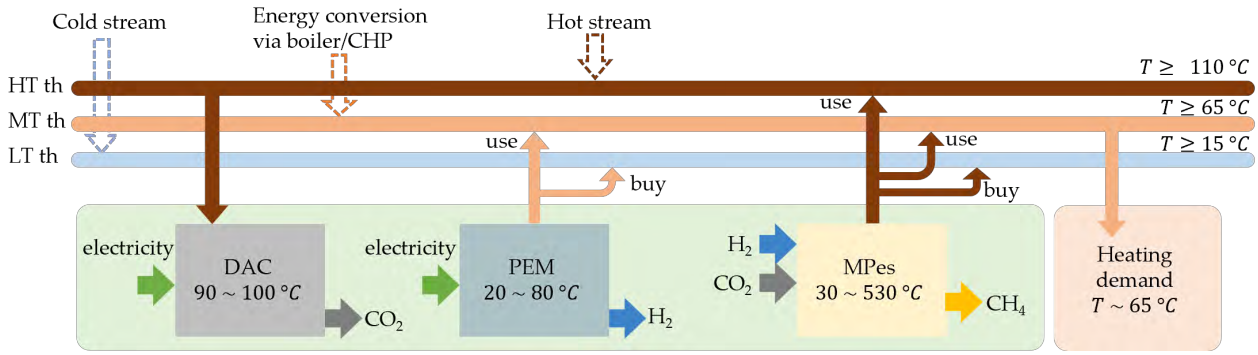

Figure 2. Heating circuit for waste heat utilization in power-to-X energy systems. A minimum allowable temperature difference ($\Delta T^{\min}$) of 10 K is assumed for all heat exchangers in the PtX energy system.

the following decision variables are included in Eq. (1):

$$\boldsymbol{y}_{\mathrm{PEM}}^{\mathrm{PS}} = \{\boldsymbol{U}^{\mathrm{cell}}, \boldsymbol{U}^{\mathrm{tn}}, \boldsymbol{I}, \boldsymbol{T}^{\mathrm{stack}}, \boldsymbol{p}^{\mathrm{cat}}, \cdots\}. \tag{2}$$

The waste heat $\dot{Q}_{\mathrm{PEM}}^{\mathrm{waste}}$ is decided by the cell voltage U^{cell}, the thermoneutral voltage U^{tn} and current I, which are related to other operating variables, such as the stack temperature T^{stack} and cathode pressure p^{cat}. As electrochemical reaction and physical properties are involved, Eq. (1) for the PEM electrolyzer includes multiple nonlinear formulations.

DAC plant. The detailed process model of the DAC plant is originally from [6]. It comprises a fan, an adsorption/desorption column, a vacuum pump and heat exchangers. The DAC plant operates through a steam-assisted temperature vacuum swing adsorption (S-TVSA) cycle, and progresses through five phases (adsorption, blowdown, heating, desorption, and cooling). The following decision variables are included in its detailed process model:

$$\boldsymbol{y}_{\mathrm{DAC}}^{\mathrm{PS}} = \{\boldsymbol{T}^{\mathrm{des}}, \dot{\boldsymbol{m}}^{\mathrm{ads}}, \boldsymbol{p}^{\mathrm{des}}, \cdots\}. \tag{3}$$

The external heat supply $\dot{Q}_{\mathrm{DAC}}^{\mathrm{supply}}$ for the DAC plant is required during the heating and desorption phases to ensure that the column reaches and maintains the desorption temperature T^{des}. Other relevant operating variables include, for example, the inlet air mass flow $\dot{m}^{\mathrm{ads}}$ and column pressure p^{des}. The original process model of the DAC plant incorporates detailed mass and heat transfer models, adsorption equilibrium models and is formulated as a system of differential-algebraic equations. To reduce computational burden, a shortcut search path is used, as applied in the literature [4].

Methanation processes. The PtX energy system includes two identical MPes. For each MP, the TREMP™ methanation process model from [3] is applied. The process model describes a four-stage Sabatier reaction, consisting of four adiabatic reactors (R1-R4) connected in series, seven adiabatic heat exchangers (HEX1-HEX7), two adiabatic compressors, and three isobaric, isothermal flash units. Accordingly, the following decision variables are included in Eq. (1):

$$\boldsymbol{y}_{\mathrm{MP}}^{\mathrm{PS}} = \{\boldsymbol{T}^{\mathrm{R1\text{-}R4}}, \boldsymbol{T}^{\mathrm{HEX1\text{-}HEX7}}, \cdots\}. \tag{4}$$

The waste heat of each MP $\dot{Q}_{\mathrm{MP}}^{\mathrm{waste}}$ is determined by the operating temperature of the reactors $T^{\mathrm{R1\text{-}R4}}$, which are primarily controlled by the inlet and outlet temperatures of the heat exchangers $T^{\mathrm{HEX1\text{-}HEX7}}$. Note that five h eat exchangers operate as cooling units and require external cold streams, while the remaining two operate as heating units and require external hot streams. Together, these seven heat exchangers enable internal heat integration within the MP and release only the residual waste heat (i.e., require only a cold stream supply). Because the detailed process model accounts for physical properties and catalytic chemical reactions, it introduces strong nonlinearities and nonconvexities into Eq. (1).

3.2 Waste heat utilization

To enable waste heat utilization, all thermal product flows must be considered along with their corresponding temperature levels. Therefore, we first extend the modeling. The proposed heating circuit for waste heat utilization in the PtX energy system is illustrated in **Figure 2**. We assume a minimum allowable temperature difference ($\Delta T^{\min}$) of 10 K for all heat exchangers in the studied PtX energy system.

According to the process characteristics, the operating temperature for the PEM electrolyzer ranges from 20 to 80 °C [5]. Under optimized operational conditions, with current densities ranging from 0.1 to 2 A/cm^2 (corresponding to 50–1000 kW H_2 at lower heating value), no waste heat is generated when the current density is below 0.98 A/cm^2 (480 kW H_2), and the stack temperature remains at 80 °C when the current density is above 0.88 A/cm^2 (433 kW H_2). The desorption temperature for the DAC plant is between 90 to 100 °C [6]. In MPes, the adiabatic reactors operate between 250 and 530 °C. Accordingly, the outlet temperature of the downstream cooling heat exchangers varies from 30 to 270 °C [3]. Additionally, the supply temperature for domestic heating

demand is typically around 65 °C [7].

Therefore, the HT waste heat of the MPes can be utilized to operate the DAC plant. Any surplus waste heat from the MPes can be further integrated into the thermal product balance at the MT level. At the same temperature level, the waste heat from the PEM electrolyzer can also be used for supplying the th-demand. As a result, the models for the thermal products are updated from:

$$\dot{\boldsymbol{Q}}^{\text{waste}}_{\text{PEM}} = \dot{\boldsymbol{E}}^{\text{out}}_{\text{PEM, th}_{\text{LT}}}, \tag{5a}$$

$$\dot{\boldsymbol{Q}}^{\text{waste}}_{\text{MP}} = \dot{\boldsymbol{E}}^{\text{out}}_{\text{MP, th}_{\text{LT}}}, \tag{5b}$$

to:

$$\dot{\boldsymbol{Q}}^{\text{waste}}_{\text{PEM}} = \dot{\boldsymbol{E}}^{\text{out}}_{\text{PEM, th}_{\text{LT}}} + \dot{\boldsymbol{E}}^{\text{out}}_{\text{PEM, th}_{\text{MT}}} + \dot{\boldsymbol{E}}^{\text{out}}_{\text{PEM, th}_{\text{HT}}}, \tag{6a}$$

$$\dot{\boldsymbol{Q}}^{\text{waste}}_{\text{MP}} = \dot{\boldsymbol{E}}^{\text{out}}_{\text{MP, th}_{\text{LT}}} + \dot{\boldsymbol{E}}^{\text{out}}_{\text{MP, th}_{\text{MT}}} + \dot{\boldsymbol{E}}^{\text{out}}_{\text{MP, th}_{\text{HT}}}. \tag{6b}$$

The bounds of the newly introduced decision variables are decided by the detailed process models:

$$\dot{\boldsymbol{E}}^{\text{out}}_{\text{PEM, th}_{\text{HT}}} = 0, \tag{7a}$$

$$0 \leq \dot{\boldsymbol{E}}^{\text{out}}_{\text{PEM, th}_{\text{MT}}} \leq \dot{\boldsymbol{E}}^{\text{out, max}}_{\text{PEM, th}_{\text{MT}}}, \tag{7b}$$

$$0 \leq \dot{\boldsymbol{E}}^{\text{out}}_{\text{MP, th}_{\text{HT}}} \leq \dot{\boldsymbol{E}}^{\text{out, max}}_{\text{MP, th}_{\text{HT}}}, \tag{7c}$$

$$0 \leq \dot{\boldsymbol{E}}^{\text{out}}_{\text{MP, th}_{\text{MT}}} \leq \dot{\boldsymbol{E}}^{\text{out, max}}_{\text{MP, th}_{\text{MT}}} + \dot{\boldsymbol{E}}^{\text{out, max}}_{\text{MP, th}_{\text{HT}}}, \tag{7d}$$

where

$$\dot{\boldsymbol{Q}}^{\text{waste}}_{\text{PEM}}, \dot{\boldsymbol{E}}^{\text{out, max}}_{\text{PEM, th}_{\text{MT}}} = g_{\text{PEM}}(\dot{\boldsymbol{E}}^{\text{in}}_{\text{PEM}}, \dot{\boldsymbol{E}}^{\text{out}}_{\text{PEM}}, \boldsymbol{y}^{\text{PS}}_{\text{PEM}}), \tag{8a}$$

$$\dot{\boldsymbol{Q}}^{\text{waste}}_{\text{MP}}, \dot{\boldsymbol{E}}^{\text{out, max}}_{\text{MP, th}_{\text{MT}}}, \dot{\boldsymbol{E}}^{\text{out, max}}_{\text{MP, th}_{\text{HT}}} = g_{\text{MP}}(\dot{\boldsymbol{E}}^{\text{in}}_{\text{MP}}, \dot{\boldsymbol{E}}^{\text{out}}_{\text{MP}}, \boldsymbol{y}^{\text{PS}}_{\text{MP}}). \tag{8b}$$

Here, the decision variable $\dot{\boldsymbol{E}}^{\text{out, max}}_{k,e}$ represents the maximum available amount of waste heat for utilization at different temperature levels.

4. OPERATIONAL OPTIMIZATION

The operational optimization problem is formulated by minimizing the operational cost, subject to Eqs. (1)–(4) and Eqs. (6)–(8), as well as product balance constraints. The resulting operational optimization is computational very challenging. Therefore, we apply the decomposition method proposed in [3, 4] and further develop it solve the operational problem with integrated waste heat utilization.

The decomposition method reformulates the original optimization problem into a set of subproblems by constructing surrogate models for the detailed process models, which are then solved iteratively. The surrogate models have either a linear or quadratic formulation. For brevity, the following description focuses on the newly developed part of the method. A comprehensive description of the complete method is provided in the previous work [3, 4].

Optimization problem at the energy system level:

$$\min_{\dot{\boldsymbol{E}}^{\text{buy,ES}}, \dot{\boldsymbol{E}}^{\text{sell,ES}}, \dot{\boldsymbol{E}}^{\text{in,ES}}, \dot{\boldsymbol{E}}^{\text{out,ES}}, \boldsymbol{x}^{\text{ES}}} C^{\text{ES}} = \sum_{t \in \mathcal{T}} \Delta t_t \sum_{\tilde{e} \in \mathcal{E}^{\text{ext}}} \left(c^{\text{buy}}_{\tilde{e},t} \dot{E}^{\text{buy,ES}}_{\tilde{e},t} - c^{\text{sell}}_{\tilde{e},t} \dot{E}^{\text{sell,ES}}_{\tilde{e},t} \right) \tag{9a}$$

$$\text{s.t.} \quad \sum_{k \in \mathcal{K}} \left(\dot{E}^{\text{out,ES}}_{k,e,t} - \dot{E}^{\text{in,ES}}_{k,e,t} \right) + \dot{E}^{\text{buy,ES}}_{e,t} - \dot{E}^{\text{sell,ES}}_{e,t} = \dot{E}^{\text{D}}_{e,t} \quad \forall e \in \mathcal{E}, t \in \mathcal{T}, \tag{9b}$$

$$\boldsymbol{g}^{\text{ES}}\left(\dot{E}^{\text{in,ES}}_{k,e,t}, \dot{E}^{\text{out,ES}}_{k,e,t}, \boldsymbol{x}^{\text{ES}} \right) \leq 0 \quad \forall k \in \mathcal{K}, e \in \mathcal{E}, t \in \mathcal{T}, \tag{9c}$$

Eqs. (6)–(7), (9d)

$$\dot{E}^{\text{buy,ES}}_{e,t}, \dot{E}^{\text{sell,ES}}_{e,t}, \dot{E}^{\text{in,ES}}_{k,e,t}, \dot{E}^{\text{out,ES}}_{k,e,t} \in \mathbb{R}^{+} \quad \forall k \in \mathcal{K}, e \in \mathcal{E}, t \in \mathcal{T}, \tag{9e}$$

$$\boldsymbol{x}^{\text{ES}} \in \mathbb{R}^{N_x}. \tag{9f}$$

Problem (9)

Optimization problem at the process system level:

$$\min_{\dot{\boldsymbol{E}}^{\text{buy,PS}}_{k}, \dot{\boldsymbol{E}}^{\text{in,PS}}_{k}, \dot{\boldsymbol{E}}^{\text{out,PS}}_{k}, \boldsymbol{y}^{\text{PS}}} C^{\text{PS}}_{k} = \sum_{t \in \mathcal{T}} \Delta t_t \sum_{\tilde{e} \in \mathcal{E}^{\text{ext}}} \left(c^{\text{buy}}_{\tilde{e},t} \dot{E}^{\text{buy,PS}}_{k,\tilde{e},t} \right) \tag{10a}$$

$$\text{s.t.} \quad \boldsymbol{g}^{\text{PS}}_{k,t}\left(\dot{E}^{\text{buy,PS}}_{\text{k},e,t}, \dot{E}^{\text{in,PS}}_{k,e,t}, \dot{E}^{\text{out,PS}}_{k,e,t}, \boldsymbol{y}^{\text{PS}}_{k,t} \right) \leq 0 \quad \forall e \in \mathcal{E}, t \in \mathcal{T}, \tag{10b}$$

$$\dot{E}^{\text{buy,PS}}_{k,e,t}, \dot{E}^{\text{in,PS}}_{k,e,t}, \dot{E}^{\text{out,PS}}_{k,e,t} \in \mathbb{R}^{+} \quad \forall e \in \mathcal{E}, t \in \mathcal{T}, \tag{10c}$$

$$\boldsymbol{y}^{\text{PS}}_{k} \in \mathbb{R}^{N_y}. \tag{10d}$$

Problem (10)

Pinch analysis at the process system level:

$$\dot{\boldsymbol{E}}^{\text{out,PS}}_{k,\text{th}_{\text{LT}}}, \dot{\boldsymbol{E}}^{\text{out,PS}}_{k,\text{th}_{\text{MT}}}, \dot{\boldsymbol{E}}^{\text{out,PS}}_{k,\text{th}_{\text{HT}}} = \text{pinchAnalysis}(\dot{\boldsymbol{E}}^{\text{in,PS}}_{e,t}, \dot{\boldsymbol{E}}^{\text{out,PS}}_{e,t}, \boldsymbol{y}^{\text{PS}}) \quad \forall k \in \mathcal{K}. \tag{11}$$

Table 1: Bounds on decision variables related to waste heat utilization in the operational optimization problem.

Ex-post approach		Co-optimized approach	
$s = 1$			
$\dot{E}^{\text{out,max}}_{\text{PEM,th}_{\text{MT}}} = 0$	(12a)	$\dot{E}^{\text{out,max}}_{\text{PEM,th}_{\text{MT}}} = \dot{E}^{\text{out,sur}}_{\text{PEM, th}_{\text{MT}}}$	(13a)
$\dot{E}^{\text{out,max}}_{\text{MP,th}_{\text{HT}}} = 0$	(12b)	$\dot{E}^{\text{out,max}}_{\text{MP,th}_{\text{HT}}} = \dot{E}^{\text{out,sur}}_{\text{MP, th}} - \dot{E}^{\text{out,sur}}_{\text{MP, th}_{\text{LT}}}$	(13b)
$\dot{E}^{\text{out,max}}_{\text{MP,th}_{\text{MT}}} = 0$	(12c)	$\dot{E}^{\text{out,max}}_{\text{MP,th}_{\text{MT}}} = \dot{E}^{\text{out,sur}}_{\text{MP, th}} - \dot{E}^{\text{out,sur}}_{\text{MP, th}_{\text{LT}}}$	(13c)
$s = 2$			
$\dot{E}^{\text{out,max}}_{\text{PEM,th}_{\text{MT}}} = 0$	(12d)	$\dot{E}^{\text{out,max}}_{\text{PEM,th}_{\text{MT}}} = \dot{E}^{\text{out,PS}}_{\text{PEM, th}_{\text{MT}}}$	(13d)
$\dot{E}^{\text{out,max}}_{\text{MP,th}_{\text{HT}}} = \dot{E}^{\text{out,PS}}_{\text{MP, th}_{\text{HT}}}$	(12e)	$\dot{E}^{\text{out,max}}_{\text{MP,th}_{\text{HT}}} = \dot{E}^{\text{out,PS}}_{\text{MP, th}_{\text{HT}}}$	(13e)
$\dot{E}^{\text{out,max}}_{\text{MP,th}_{\text{MT}}} = \dot{E}^{\text{out,PS}}_{\text{MP, th}_{\text{MT}}}$	(12f)	$\dot{E}^{\text{out,max}}_{\text{MP,th}_{\text{MT}}} = \dot{E}^{\text{out,PS}}_{\text{MP, th}_{\text{MT}}}$	(13f)
$s = 3$			
$\dot{E}^{\text{out,max}}_{\text{PEM,th}_{\text{MT}}} = \dot{E}^{\text{out,PS}}_{\text{PEM, th}_{\text{MT}}}$	(12g)	$\dot{E}^{\text{out,max}}_{\text{PEM,th}_{\text{MT}}} = \dot{E}^{\text{out,PS}}_{\text{PEM, th}_{\text{MT}}}$	(13g)
$\dot{E}^{\text{out,ES}}_{\text{MP,th}_{\text{HT}}}$ is fixed.	(12h)	$\dot{E}^{\text{out,ES}}_{\text{MP,th}_{\text{HT}}}$ is fixed.	(13h)
$\dot{E}^{\text{out,max}}_{\text{MP,th}_{\text{MT}}} = \dot{E}^{\text{out,PS}}_{\text{MP, th}_{\text{MT}}} + (\dot{E}^{\text{out,PS}}_{\text{MP, th}_{\text{HT}}} - \dot{E}^{\text{out,ES}}_{\text{MP,th}_{\text{HT}}})$	(12i)	$\dot{E}^{\text{out,max}}_{\text{MP,th}_{\text{MT}}} = \dot{E}^{\text{out,PS}}_{\text{MP, th}_{\text{MT}}} + (\dot{E}^{\text{out,PS}}_{\text{MP, th}_{\text{HT}}} - \dot{E}^{\text{out,ES}}_{\text{MP,th}_{\text{HT}}})$	(13i)

After decomposition, the operational problem consists of one optimization problem at the energy system level (Problem (9)), and multiple optimization problems focusing on the process system level (Problem (10)), one for each energy system component k ($k \in \mathcal{K}^{\text{PS}}$).

In Problem (9), the objective function minimizes the purchase of external energy consumption ($\dot{E}^{\text{buy, ES}}_{\tilde{e},t}$) and maximizes the feed-in revenues ($\dot{E}^{\text{sell, ES}}_{\tilde{e},t}$), weighted by their respective cost parameter, $c^{\text{buy}}_{\tilde{e},t}$, and $c^{\text{sell}}_{\tilde{e},t}$, and multiplied by the time-step length Δt_t (Eq. (9a)). Equation (9b) represents product balance constraints, while Eq. (9c) summaries the input-output relationships of each k. They correspond to Eq. (1) for $k \in \mathcal{K} \backslash \mathcal{K}^{\text{PS}}$ and to surrogate models for $k \in \mathcal{K}^{\text{PS}}$. For energy system component $k \in \mathcal{K}^{\text{PS}}$, the detailed process models are solved separately at the process system level in Problem (10) (cf. Eq. (10b)). The superscripts ES and PS indicate whether a decision variable or a constraint belongs to Problem (9) or (10), respectively.

Due to the inclusion of waste heat utilization, Eqs. (6)–(8) must be incorporated into the operational problem. Equations (6)–(7) can be directly integrated into the optimization problem at the energy system level (Problem (9)) as complicating constraints (cf. Eq. (9d)). Since Eqs. (8) involve detailed process models and cannot be directly integrated into Problem (9), we replace them by incorporating pinch analysis into the decomposition method, as shown in Eq. (11).

The semi-sequential approach introduced in the literature [4] is applied to solve the decomposed optimization problems. The main steps of the approach are presented in **Algorithm 1**.

The energy system components using detailed process models are addressed with a pre-defined sequence (lines 1 to 2). We first determine the operational decisions for MPes by iteratively solving Problems (9) and (10) until their solutions converge to identical operational decisions (lines 3 to 10). Subsequently, pinch analysis is performed for each MP to identify the bounds for waste heat utilization (line 11). We then fix these bounds, along with all decision variables related to MPes, in Problem (9) (lines 12 to 14) and proceed to the DAC plant in the next sequence (cf. line 2). These sequential steps are repeated until all $k \in \mathcal{K}^{\text{PS}}$ are addressed (lines 2 to 14). Consequently, we find an optimal operational strategy for the PtX energy system (line 15).

Algorithm 1 Decomposition method

```
1: s = 1
2: For k = {MPes, DAC, PEM}
3: i = 1
4: For i < i^max:
5: Solve Problem (9)
6: Solve Problem (10)
7: If converge:
8:  Goto Step 11
9: Else:
10:  i = i+1 and Goto Step 4
11: Do pinch analysis via Eq. (11)
12: Fix bounds via Eqs. (12) or (13)
13: Fix decision variables of all con-
converged components in Problem (9)
14: s = s+1 and Goto Next k
15: Find an optimal operational strategy
```

As mentioned earlier, incorporating waste heat utilization introduces additional complicating constraints into the optimization problem at the energy system level (cf. Eqs. (6)). These complicating constraints can significantly increase the computational effort required to solve Problem (9). To address this issue, we propose two solution approaches for handling these complicating constraints.

As shown in **Table 1**, the bounds, i.e., the maximum available waste heat at different temperature levels, are integrated into Problem (9) either "ex-post" or "co-optimized", depending on whether waste heat utilization is considered ex post or jointly with operational optimization.

In the ex-post approach, waste heat utilization is initially deactivated (cf. Eqs. (12a)–(12c)). After the operational decisions for MPes are determined, the bounds on the available thermal products that can be reused at HT and MT levels are identified via pinch analysis. These bounds are then passed to Problem (9) for the subsequent sequence (cf. Eqs. (12d)–(12f)). After the operational decisions for the DAC plant are determined, the bounds in Problem (9) are updated for the final component, PEM (cf. Eqs. (12g)–(12i)). In this approach, the complicating constraints, i.e., Eqs. (6), are inactive during $s = 1$, which reduces the complexity of Problem (9).

In the co-optimized approach, the complicating constraints are active throughout all sequences. Instead of fixing the maximum bounds to zero, feasible bounds are specified for the decision variables associated with waste heat utilization. This difference is illustrated by comparing Eqs. (12a)–(12d) with Eqs. (13a)–(13d). Moreover, because pinch analysis is performed at the end of each sequence, the feasible bounds for the first sequence are obtained from the surrogate model of component k ($k \in \{\mathrm{PEM}, \mathrm{MPes}\}$) (cf. Eqs. (13a)–(13c)).

Note that **Algorithm 1** can also handle non-convergence situations. As the focus of this work is the integration of waste heat utilization, a detailed description is omitted. Further information on handling non-convergence can be found in the literature [3, 4].

5. NUMERIC STUDY

5.1 Implementation and specification

We apply the proposed method to optimize the daily operation of the aforementioned PtX energy system. The resulting operational problem is solved via ex-post and co-optimized approaches, respectively. In addition, a reference case using the original decomposition method is included as a benchmark to evaluate the impact of waste heat utilization.

5.2 Results and discussion

We discuss the results of the numerical study from two aspects: computational performance and the resulting optimal operational strategy.

5.2.1 Computational performance

The computational effort of all three cases is given in **Table 2**.

In the reference and ex-post cases, the computational performance is very similar, with the operational optimization problem being solved within 40 min in both cases. This is reasonable, as the complexity of the optimization problems addressed in the two cases is nearly identical, with the ex-post case featuring slightly relaxed bounds on decision variables since the second sequence (cf. Eqs. (12e)–(12i) in **Table 1**).

Table 2: Comparison of computational performance. HU: heat utilization.

	w/o HU	**w/ HU**	
	Ref.	**Ex-post**	**Co-optimized**
Time [s]	2133 (~36 min)	2281 (~38 min)	12063 (~3.4 h)
Nr. of iteration ($\mathrm{MPes} + \mathrm{DAC} + \mathrm{PEM}$):			
	6+2+1=9	6+4+1=11	11+2+1=14
Average computing time for solving Problem (9) [s]:			
$s = 1$	256	259	1011

Table 3: Overview of the system efficiency and operational costs. HU: heat utilization.

	w/o HU	**w/ HU**	
	Ref.	**Ex-post**	**Co-optimized**
Efficiency	41 %	48 %	61%
Operational costs [€/d]:			
Total	5672	5266	4431
el-grid	4855	4545	3983
hot stream	793	747	448
cold stream	25	1	1
g-grid	0	-27	0

The co-optimized case requires significantly more computational time, approximately 3.4 hours. In particular, the operational decisions of MPes converge more slowly, necessitating twice as many iterations compared with the other two cases. Moreover, the inclusion of additional complicating constraints at the energy system level further increases the average solution time of Problem (9). Especially, in the first sequence ($s = 1$), the average solution time is about four times as long as in other cases. The average computing times in the subsequent sequences are comparable across the three cases and are therefore not presented here.

5.2.2 Optimal operational strategy

The overall operational performance of the PtX energy system is summarized in **Table 3**, where the system efficiency is defined as the ratio of useful output energy to total input energy.

The results indicate that waste heat utilization significantly improves the operational performance of the PtX energy system: The overall system efficiency

increases by up to 48.8 % and the total operational cost decreases by up to 21.9 % (cf. the co-optimized case).

In all three cases, electricity consumption accounts for over 85% of the total operational cost, while the purchase of the hot stream contributes more than 10%. When waste heat utilization is applied, additional savings are achieved not only through reduced reliance on external thermal supplies but also through lower electricity consumption.

To better understand the operational decisions on energy consumption, we examine the resulting operating strategy of selected energy system components. The total amounts of CH_4, CO_2, and H_2 produced for three cases are shown in **Figure 3**.

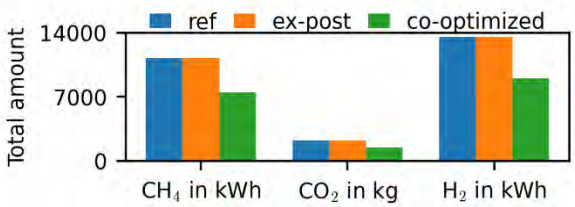

Figure 3. Total amounts of CH_4, CO_2, and H_2 produced.

Compared with the reference case, the amounts of CH_4, CO_2 and H_2 produced in the ex-post case remain unchanged. This is because the operational decisions of the MPes are determined before waste heat utilization is applied. As a result, the CH_4 output, H_2 and CO_2 inputs are identical in both cases. However, because less CH_4 is required for supplying th-demand, the surplus CH_4 is fed into the g-grid. In addition, the CHP engine operates and converts 6614 kWh redundant CH_4 into 2554 kWh electricity, leading to cost savings in electricity consumption. In contrast, in the reference case, the CHP engine remains inactive over the entire input time series.

In the co-optimized case, a significant decrease in the production of all three products is observed. To better understand this decrease, we further examine operational decisions related to waste heat utilization. **Figure 4** illustrates the distribution of waste heat from the MPes in both the ex-post and co-optimized cases. The lines indicate the available waste heat, while the bars show the integrated amounts at HT, MT and LT temperature levels.

In the ex-post case, 4257 kWh waste heat is generated from the MPes: 89.5 % at HT level, 5.4 % at MT level, and 5.1 % at the LT level, while the LT waste heat cannot be reused within the PtX energy system. About 52.2 % of the HT waste heat is integrated into the thermal product balance at the HT level to support the operation of the DAC plant. The remaining 47.8 % HT waste heat, together with the 5.4 % MT waste heat is supplied into thermal product balance at the MT level.

In the co-optimized case, although 89.4 % of the total 2805 kWh waste heat is available at the HT level, all of it is integrated into the MT level to support heating demand. As a result, less CH_4 is required, which in turn reduces the demand for H_2 and CO_2 within the energy system (cf. **Figure 3**), thereby significantly reducing the overall operational costs (cf. **Table 3**).

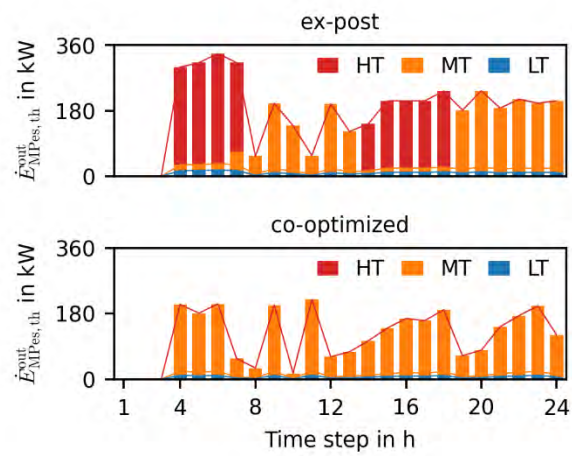

Figure 4. Waste heat generated by the MPes and its distribution among reuse options at each time step. Lines: available waste heat; bars: integrated waste heat.

6. CONCLUSION

This work presents an optimization method for integrating waste heat utilization into operational problems of PtX energy systems. Using detailed process models and pinch analysis, waste heat at multiple temperature levels is incorporated into the operational optimization problems. The proposed method includes two different solution approaches, namely ex-post and co-optimized, to handle the problem complexity and offer a trade-off between computational effort and solution optimality.

The proposed method is applied to optimize the daily operation of a PtX energy system adopted from literature. Compared with a reference case computed with literature setting (without waste heat utilization), the two proposed solution approaches improve operational performance to different extents while requiring different levels of computational effort. The ex-post case achieves an efficiency improvement of 7 p.p. and cost savings of 7.2 %, while requiring the same computational time as the reference case. When a higher computation effort is acceptable, the co-optimized approach yields significantly larger improvements, with an increase in system efficiency of 20 p.p. and cost savings of 21.9 %.

Overall, the proposed method successfully addresses waste heat utilization within the operational optimization problem and is generally applicable to renewables-based energy systems, especially those with PtX technologies integrated. Future research could focus on reducing the computational effort of the co-optimized

approach. Moreover, control strategies could be integrated into the proposed method to assess their applicability under real-world operating conditions.

DIGITAL SUPPLEMENTARY MATERIAL

The digital supplementary material of this work can be found in a Git repository: https://git-ce.rwth-aachen.de/ltt/ptg-es4.

ACKNOWLEDGEMENTS

This study is funded by the "Europäischer Fonds für regionale Entwicklung" (EFRE-20800350). The support is gratefully acknowledged.

AUTHOR IDENTIFIERS

Author ORCIDs:
Wang, Y.: 0009-0001-7914-5952
Bornemann, L.: 0009-0008-8064-1364
Von der Assen, N: 0000-0001-8855-9420

REFERENCES

1. Galvan-Cara AL, Bongartz D. Waste-heat upgrading from alkaline and PEM electrolyzers using heat pumps. Systems and Control Transactions 4:2063-2068 (2025). https://doi.org/10.69997/sct.192791
2. Bornemann L, Tuschewitzki W, Kaltschmitt M. Exploring the impact of synergistic oxygen and waste heat utilization from electrolysis in supporting a climate-neutral energy infrastructure in germany. Energy Conversion and Management: X 29:101504 (2026). https://doi.org/10.1016/j.ecmx.2025.101504
3. Wang Y, Bornemann L, Reinert C, von der Assen N. A method to bridge energy and process system optimization: identifying the feasible operating space for a methanation process in power-to-gas energy systems. Computers & Chemical Engineering 182:108582 (2024). https://doi.org/10.1016/j.compchemeng.2023.108582
4. Wang Y, von der Assen N, Zhang Q. A decomposition method for optimizing the operation of power-to-X energy systems with detailed process models. *Available at SSRN: http://dx.doi.org/10.2139/ssrn.5461101.* (2025).
5. Bornemann L, Lange J, Kaltschmitt M. Optimizing temperature and pressure in PEM electrolyzers: a model-based approach to enhanced efficiency in integrated energy systems. Energy Conversion and Management 325:119338 (2025). https://doi.org/10.1016/j.enconman.2024.119338
6. Postweiler P, Engelpracht M, Rezo D, Gibelhaus A, von der Assen N. Environmental process optimisation of an adsorption-based direct air carbon capture and storage system. Energy Environ. Sci. 17:3004-3020 (2024). https://doi.org/10.1039/d3ee02970k
7. Wahi P, Konstantinou T, Tenpierik MJ, Visscher H. Lower temperature heating integration in the residential building stock: a review of decision-making parameters for lower-temperature-ready energy renovations. Journal of Building Engineering 65:105811 (2023). https://doi.org/10.1016/j.jobe.2022.105811

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Comparative Life Cycle Assessment of Electrochemical and Conventional Regeneration Pathways in KOH-Based Direct Air Capture Systems

Georgia Ioanna Prokopou[a], Zoi Drakopoulou[b], Dominik Bongartz[cd], and Alexander Mitsos[eaf*]

a Process Systems Engineering (AVT.SVT), RWTH Aachen University, 52074 Aachen, Germany
b Department of Environmental and Resource Engineering, Technical University of Denmark, 2800, Copenhagen, Denmark
c Department of Chemical Engineering, KU Leuven, 3001 Leuven, Belgium
d EnergyVille, 3600 Genk, Belgium
e JARA-ENERGY, 52056 Aachen, Germany
f Energy Systems Engineering (IEK-10), Forschungszentrum Jülich, 52425 Jülich, Germany
* Corresponding Author: amitsos@alum.mit.edu.

ABSTRACT

Achieving net climate neutrality will likely require negative-emission technologies such as Direct Air Capture (DAC). Potassium hydroxide (KOH) absorption is one of the most mature DAC approaches, but it can cause significant emissions due to natural-gas-based thermal regeneration. Electrochemical regeneration methods, such as electrolysis and electrodialysis, have recently been proposed as alternatives, yet their relative performance and environmental impacts remain unclear. We present a comparative cradle-to-gate life cycle assessment (LCA) of three KOH-based DAC configurations: (i) the established Ca-looping thermal regeneration, (ii) the electrolysis regeneration (DAC-ELY), which co-produces hydrogen, and (iii) the electrodialysis regeneration (DAC-ED). The results show that, expectedly, electricity demand dominates life cycle impacts across all configurations. With the current German electricity mix, the established DAC has the lowest overall impacts, while DAC-ELY and DAC-ED exhibit higher global warming impacts due to the high electricity requirements of electrochemical regeneration. Infrastructure impacts are more noticeable for DAC-ED due to its low current density and large membrane areas. This study highlights that electrification alone does not guarantee lower life cycle impacts and that the sustainability of electrochemical DAC depends on process efficiency, material use, and the expected use of renewable electricity.

Keywords: Lifecycle Assessment, Direct Air Capture, Electrolysis, Electrodialysis, Electrochemical Regeneration

INTRODUCTION

Limiting global warming to 1.5 °C above pre-industrial levels requires reductions in greenhouse gas emissions [1]. Nevertheless, carbon-based resources are expected to remain in use in hard-to-abate sectors such as cement, steel, and aviation, where complete electrification is technically or economically challenging. As a result, residual CO_2 emissions are likely to persist, creating a challenge in achieving net-zero emission targets [2]. In this context, direct air capture (DAC) technologies, which remove CO_2 directly from the ambient air, have gained increased attention [2].

Several DAC technologies are currently under development in the market. The most mature approach is the absorption of CO_2 in a KOH solution with Ca-looping regeneration, which relies on high-temperature thermal regeneration. At the same time, novel electrochemical DAC concepts have recently emerged that replace thermal regeneration with electrically-driven processes, such as electrolysis-based regeneration and electrodialysis-based regeneration. These technologies offer potential advantages in terms of electrification and integration with renewable energy but are currently at lower

technology readiness levels [4].

Techno-economic assessments (TEAs) are widely used to evaluate emerging DAC technologies [5]. However, TEA alone does not capture the environmental implications of these systems. DAC technologies are both energy- and material-intensive, and their environmental performance is strongly influenced by factors such as the electricity supply and infrastructure requirements. A life cycle assessment (LCA) is therefore needed to identify potential environmental trade-offs.

Several studies have investigated the techno-economic and environmental performance of DAC technologies [5]-[8]. However, a comprehensive comparative LCA of the established KOH-based DAC with Ca-looping regeneration and the emerging KOH-based DAC technologies with electrochemical regeneration is still lacking in the literature.

To fill in this gap, we perform a comparative LCA of three KOH-based DAC technologies: DAC with Ca-looping regeneration, DAC with electrolysis-based regeneration, and DAC with electrodialysis-based regeneration. We thus aim to provide insights into their current environmental performance, identify key drivers of impacts, and highlight ways for improving the sustainability of electrochemical DAC systems.

SYSTEM DESCRIPTION

As mentioned earlier, three KOH-based DAC technologies are studied in the present work, which differ in the sorbent regeneration method: (i) the established Ca-looping regeneration (DAC-Established), (ii) the electrolysis regeneration (DAC-ELY), and (iii) the electrodialysis regeneration (DAC-ED). A schematic representation of the studied technologiesis given in Figure 1.

Direct Air Capture with Ca-looping Regeneration

This technology is commercialized by Carbon Engineering [3], exhibits a high TRL of 7-8 [9], and will be considered as the benchmark for the comparison of the other technologies. It consists of two connected chemical steps. In the first step, CO_2 from the atmosphere is brought in contact with a KOH solution in an air contactor unit, producing an aqueous solution of K_2CO_3. In the second step, the solution enters a pellet reactor in which it reacts with $Ca(OH)_2$, producing $CaCO_3$ and KOH. Then the $CaCO_3$ enters a calciner and is decomposed to CaO at high temperatures of 900 °C, while releasing CO_2. The heat needed for the release of CO_2 is provided by burning natural gas [10]. Finally, the CaO is hydrated in a slaker unit to regenerate the $Ca(OH)_2$ that is fed back to the pellet reactor. If the additional CO_2 released from the natural gas burning is not captured, then the process could potentially lead to positive direct CO_2 emissions.

Direct Air Capture with Electrolysis Regeneration

DAC with electrolysis regeneration is a technology at a low TRL but has been proposed as a promising alternative to the established thermal process [11], as it is powered by electricity and can integrate renewable energy [4].

The process consists of two coupled chemical steps, with the first being the same as in the DAC-Established. In the second step, KOH regeneration and CO_2 release are achieved in an electrochemical reactor. This reactor comprises an anode, an anolyte compartment, a cation exchange membrane (CEM), a catholyte compartment, and a cathode [11]. The $CaCO_3$ solution produced from the first step is fed to the anodic side of the electrochemical reactor. By applying current, the water oxidation reaction (WOR) occurs at the anode, where gaseous oxygen and protons are produced according to Reaction (1), while the water reduction reaction occurs at the cathode, producing gaseous hydrogen and hydroxide anions according to Reaction (2):

$$H_2O \rightarrow 2H^+ + 2e^- + 0.5O_2 \quad (1)$$

$$2H_2O + 2e^- \rightarrow H_2 + 2OH^- \quad (2)$$

Simultaneously, mainly K^+ migrate through the CEM from the anolyte to the catholyte. This results in an acidic environment in the anolyte and in an alkaline environment at the catholyte. The acidic environment in the anolyte drives the conversion of carbonate species into gaseous CO_2, which is released in the form of bubbles, while the alkaline environment in the catholyte enables the regeneration of the KOH [11]. Due to the production of O_2 and the release of CO_2 in the anolyte, an additional pressure swing adsorption unit (PSA) is required after the electrochemical reactor to separate the two gases [4].

Direct Air Capture with Electrodialysis Regeneration

Electrodialysis regeneration is another electrochemical approach that has been investigated for DAC. It exhibits a TRL of approximately 4 [5], and similarly to the electrolysis regeneration, the operating principle is based on the pH gradient across ion-exchange membranes.

As in the electrolysis regeneration, the process consists of two main steps, with the initial CO_2 capture step being identical to that of the DAC-Established process. In the second step, the $CaCO_3$ solution is fed into an electrodialysis regeneration unit, which exhibits some differences compared to the electrolysis regeneration. The electrodialysis reactor consists of a single anode and cathode, while in between multiple Bipolar Membranes (BPM) and CEM pairs are arranged in series [4].

By applying current, the reactions (1) and (2) occur at the anode and cathode, respectively. In addition, the

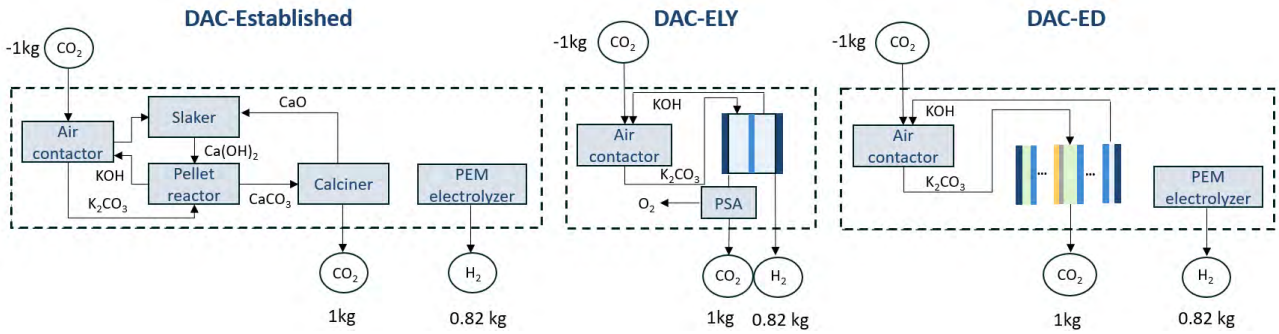

Figure 1. Schematic representation of the cradle-to-gate system boundaries for the DAC with Ca-looping regeneration (DAC-Established), the DAC with electrolysis regeneration (DAC-ELY), and the DAC with electrodialysis regeneration (DAC-ED).

BPMs, which consist of an anion exchange layer and a cation exchange layer, result in water dissociation into OH^- and H^+. This generates acidic and alkaline environments on opposite sides of the membrane, establishing a pH gradient across the CEM [4]. Similarly to the electrolysis regeneration, in the acidic compartments, CO_2 is released in the form of bubbles. In contrast to the electrolysis regeneration, CO_2 is released in chambers that are physically separated from the anode, where oxygen is generated. As a result, no additional PSA unit is required for gas separation. Moreover, the large number of BPM-CEM pairs between the two electrodes allows operation at a lower current density compared to the electrolysis unit. Consequently, the production rate of hydrogen is considered negligible compared to the release of CO_2 [4].

MATERIALS AND METHODS

The LCA was conducted in accordance with the ISO 14040/14044 framework, incorporating the main phases of goal and scope definition, life cycle inventory analysis, life cycle impact assessment, and interpretation [12].

Goal and Scope Definition

The goal of the present study is the comparative LCA of the three DAC technologies. We assess the environmental impacts by performing a cradle-to-gate analysis, including the infrastructure and air capture processes, while excluding end-of-life treatment due to the unknown fate of the final products, as they are similar among the technologies.

The DAC-ELY co-produces hydrogen as a valuable product during electrochemical regeneration. Under the operating current density of 0.14 A cm^{-2}, the hydrogen co-production is 0.82 kg H_2 per kg CO_2 captured [4]. To ensure functional equivalence among the compared systems, multifunctionality is addressed by expanding the system boundaries [13]. In this context, a PEM water electrolyzer is included within the system boundaries for the DAC-Established and the DAC-ED cases to generate 0.82 kg of H_2. Thus, the functional unit is 1 kg of CO_2 and 0.82 kg of H_2.

Life Cycle Inventory Analysis

The life cycle inventory analysis was performed using literature process modeling and experimental data for the different DAC technologies. GaBi [14] was employed as LCA software to calculate the environmental impacts. Background data for material and energy flows associated with process infrastructure were obtained from the Ecoinvent database [15]. For all the assessed technologies, the electricity supply for the operation was assumed to be provided by the German electricity grid (unless otherwise stated), and the lifetime of the equipment and the plant is assumed to be 20 years. In the following, the assumptions of the inventory modeling and the data used for the different DAC technologies are discussed in detail.

PEM Electrolyzer

A PEM electrolyzer operating at a current density of 1.5 A cm^{-2} is assumed to produce H_2 in the DAC-Established and the DAC-ED, according to Zachert et al. [16]. Infrastructure data for the PEM electrolyzer stack and balance of plant (BoP) are taken from Bareiß et al. [17]. Due to the lack of inventory data for iridium at the anode, platinum was assumed as a proxy based on available data from the Ecoinvent database [15]. Moreover, the Nafion membrane material was approximated as tetrafluoroethylene (TFE) [18].

Direct Air Capture with Ca-looping Regeneration

Operational data for the DAC-Established process are taken from Keith et al. [10]. The thermal energy required for sorbent regeneration is assumed to be supplied by natural gas combustion. The oxygen required for combustion is assumed to be provided by the PEM electrolyzer; therefore, no additional air separation unit was

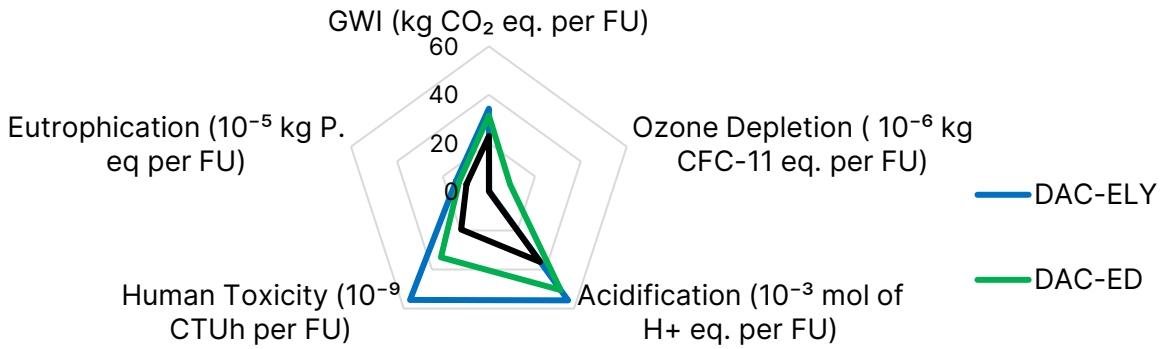

Figure 2: Comparative life cycle environmental impacts of the DAC with Ca-looping regeneration (DAC-Established), DAC with electrolysis regeneration (DAC-ELY), and DAC with electrodialysis regeneration (DAC-ED), assuming the German electricity grid. The DAC-Established exhibits the lowest life cycle impacts in all assessed categories, while the DAC-ELY shows the highest scores in all categories except ozone depletion. The functional unit (FU) is 1 kg CO_2 captured and 0.82 kg H_2 produced. GWI refers to the global warming impact.

included in the system boundaries [19]. Emissions from natural gas combustion were excluded from the inventory, as the CO_2 produced during sorbent regeneration was assumed to be fully captured. Only emissions associated with natural gas production and transportation were considered [19]. Infrastructure data for the air contactor and the regeneration units are taken from Madhu et al. [20]. The pellet reactor, slaker, and calciner were approximated as steel tanks for the purpose of infrastructure modeling [20].

Direct Air Capture with Electrolysis Regeneration

An electrolysis regeneration unit operating at a current density of 0.14 A cm^{-2} is assumed according to the operational data from Liu et al. [4]. A PSA unit for the separation of CO_2 and O_2 is included [4], with infrastructure data for the PSA taken from Burgers et al. [21]. Infrastructure data for the electrolysis regeneration unit stack and the BoP are obtained from Bareiß et al. [17], with modifications to reflect the specific operating conditions of the DAC system. In contrast to the PEM electrolyzer, a nickel catalyst is used at the cathode instead of platinum due to the alkaline operating environment [11]. Similarly to the PEM electrolyzer case, platinum is used as a proxy for iridium at the anode and TFE for the Nafion membrane.

Direct Air Capture with Electrodialysis Regeneration

An electrodialysis regeneration unit operating at a current density of 0.012 A cm^{-2} is assumed according to the operational data from Liu et al. [4]. Each regeneration unit is assumed to consist of 20 CO_2 recovery-alkali regeneration pairs [4]. Infrastructure data for the electrodialysis regeneration unit stack and the BoP are obtained from Bareiß et al. [17]. Similarly to the electrolysis regeneration case, a nickel catalyst is assumed at the cathode, while platinum was used as a proxy material at the anode. TFE is used as a proxy for the CEM, and the BPM is assumed to consist of two Nafion membranes [22].

Life Cycle Impact Assessment

We investigated five typical impact categories: global warming impact (GWI), ozone depletion, acidification, human toxicity, and eutrophication. The environmental impacts were quantified using the Environmental Footprint (EF) 2.0 impact assessment method [23]. The results are normalized to the functional unit of 1 kg CO_2 captured and 0.82 kg H_2 produced.

RESULTS AND DISCUSSION

Environmental Performance across Impact Categories

Figure 2 presents the comparative results for the different impact categories of the three DAC technologies assuming the German electricity grid. Overall, the DAC-Established technology exhibits the lowest life cycle impacts in all assessed categories. This can be attributed to its higher technological maturity and optimization level compared to the more novel electrochemical DAC processes. In contrast, the DAC-ELY shows the highest life cycle impacts in all categories except for ozone depletion. For ozone depletion, the DAC-ED has the highest impact, which can be attributed to the larger number of membranes in the electrodialysis configuration, leading to increased material-related ozone depletion contributions.

Since the global warming impact is one of the most important categories, we focus in the following subsections on its detailed analysis, including process- and infrastructure-level contributions, to explain the observed differences among the assessed technologies.

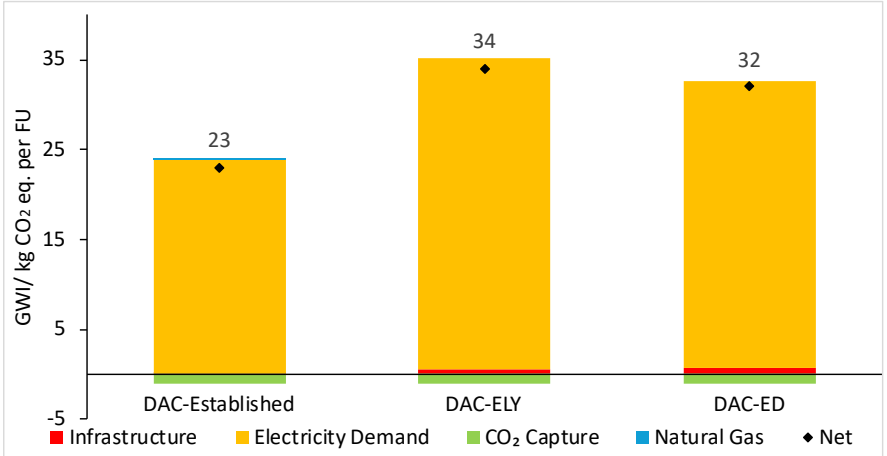

Figure 3: Cradle-to-gate global warming impact (GWI) distribution for the DAC with Ca-looping (DAC-Established), the DAC with electrolysis regeneration (DAC-ELY), and the DAC with electrodialysis regeneration (DAC-ED), assuming the German electricity grid. The main GWI contributor is the electricity demand in all three technologies. The DAC-ELY exhibits higher electricity demand compared to the DAC-Established and DAC-ED. The negative GWI is attributed to the direct removal of CO_2 from the atmosphere through the DAC process. The infrastructure contribution is approximately 2% of the total GWI in the DAC-ED. The functional unit (FU) is 1 kg CO_2 captured and 0.82 kg H_2 produced.

Contribution Analysis of the Global Warming Impact

Figure 3 presents the global warming impact distribution for the three DAC technologies. Overall, in all three technologies, the electricity demand dominates the GWI. There is a 47% and 38% increase in the net GWI of the DAC-ELY and the DAC-ED, respectively, compared to the DAC-Established. It should be noted that the high electricity demand in the DAC-Established and the DAC-ED is mainly due to the additional PEM-electrolyzer included to ensure a fair comparison of the three technologies via the system expansion approach.

It can also be seen that there are negative GWI corresponding to 1 kg of CO_2 eq. per kg of FU, reflecting the direct removal of CO_2 from the atmosphere through the DAC process. Moreover, the highest infrastructure GWI is observed for DAC-ED of about 2% of the total impact. Overall, the contribution of the infrastructure is negligible compared to the electricity demand. However, in case there is a larger penetration of renewable energy resources (as analyzed below), the relative contribution of the infrastructure may become more significant. Therefore, it is still of interest to examine the differences in the infrastructure contributions, as follows.

Contribution of Infrastructure to Global Warming Impact

Figure 4a) presents the distribution of infrastructure-related contributions to the global warming impact of the three DAC technologies. The regeneration unit is the main contributor to the infrastructure impact in the case of DAC-ELY and DAC-ED, while DAC-Established has a negligible infrastructure impact. The contribution of the air contactor is identical across all three technologies, reflecting the use of the same design assumptions. At the same time, the PEM-electrolyzer exhibits considerably lower infrastructure GWI compared to the electrochemical regeneration units in the DAC-ELY and DAC-ED systems. This is primarily due to the higher operating current density (1.5 A cm^{-2} compared to 0.14 A cm^{-2} vs 0.012 A cm^{-2}, respectively), which translates into a lower amount of material requirement per functional unit.

Figure 4b) presents the distribution of infrastructure-related contributions from the regeneration units to the global warming impact of the three DAC technologies. For both DAC-ELY and DAC-ED, the main contributors within the regeneration units are the iridium, titanium, and Nafion-based membrane materials. In the case of DAC-ED, the contribution associated with the Nafion membranes is higher than in the case of DAC-ELY. This can be attributed to two factors. First, the DAC-ED regeneration unit operates at a current density approximately eleven times lower than that of the DAC-ELY, resulting in higher material requirements per functional unit. Second, the DAC-ED requires 3 additional membranes per CO_2 recovery-alkali regeneration pair. Therefore, the influence of Nafion membranes on the overall infrastructure-related global warming impact is more pronounced in the DAC-ED system.

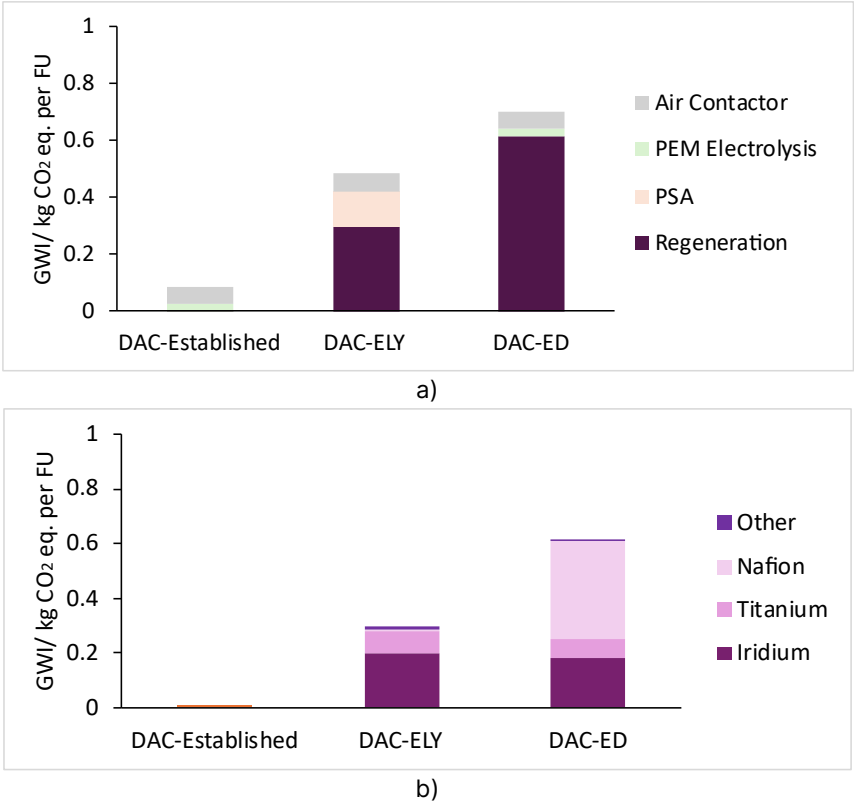

Figure 4: (a) Distribution of infrastructure-related global warming impact (GWI) and (b) distribution of infrastructure-related contributions from the regeneration units to the GWI for the DAC with Ca-looping (DAC-Established), the DAC with electrolysis regeneration (DAC-ELY), and the DAC with electrodialysis regeneration (DAC-ED). The regeneration unit has the highest contribution to the infrastructure of DAC-ED, with Nafion-based membranes being the dominant contributing material. The functional unit (FU) is 1 kg CO_2 captured and 0.82 kg H_2 produced.

Influence of the Carbon Footprint of Electricity on the Global Warming Impact

To assess the potential of the novel electrochemical DAC technologies, in Figure 5, we present the influence of the carbon footprint of electricity across different countries on the GW [24]. A lower electricity carbon footprint corresponds to a higher penetration of low-carbon energy sources in the respective electricity mix. The DAC-Established exhibits the lowest GWI across all electricity scenarios. Moreover, the DAC-ELY approaches the GWI of the DAC-ED at lower carbon footprints of electricity, since the electricity demand contribution becomes lower.

The results suggest that, in addition to the larger penetration of renewable electricity, further performance improvements are needed for DAC-ELY and DAC-ED to compete with the DAC-Established technology. Reductions in materials use and infrastructure-related impacts or operation at higher current densities will be necessary to realize the potential of electrochemical DAC technologies.

CONCLUSION AND OUTLOOK

We performed a comparative LCA of three DAC technologies: the established DAC technology with Ca-looping regeneration; the novel technology of DAC with electrolysis regeneration; and the DAC with electrodialysis regeneration. The results have shown that the DAC-Established outperforms the other technologies in all impact categories, when assuming the German electricity grid. DAC-ELY shows the highest GWI due to the higher electricity demand associated with the electrochemical sorbent regeneration. The infrastructure impact is higher in the case of DAC-ED due to both the use of additional membranes and the low operating current densities. Finally, it is concluded that DAC-ELY and DAC-ED need further performance improvement and a larger share of renewable energy to compete with the DAC-Established.

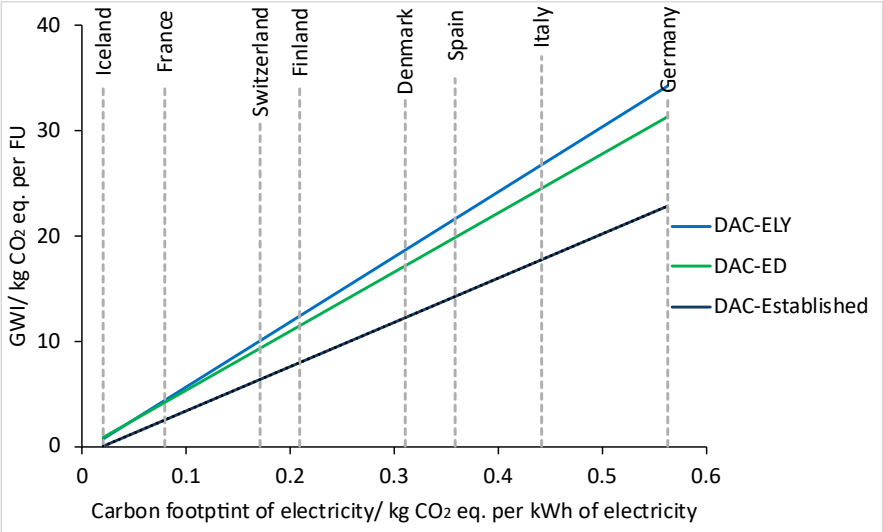

Figure 5: Global warming impact (GWI) with respect to the carbon footprint of electricity for the DAC with Ca-looping (DAC-Established), the DAC with electrolysis regeneration (DAC-ELY), and the DAC with electrodialysis regeneration (DAC-ED). DAC-Established exhibits the lowest GWI across all electricity mix scenarios. DAC-ELY approaches the GWI of DAC-ED in cases of larger penetration of alternative energy resources. The functional unit (FU) is 1 kg CO_2 captured and 0.82 kg H_2 produced.

However, several limitations of this study should be acknowledged. In cases where no dataset was available, proxies were used for the infrastructure, which may not fully capture the specific impacts of the actual materials. Moreover, upstream processes (e.g., metal extraction, refining) were modeled with the average European electricity mix, which may lead to changes if region-specific grids were applied.

Future work should extend the comparison of these DAC technologies beyond environmental performance to include TEA. This will allow for a more holistic evaluation of the three alternatives. Moreover, system optimization with respect to both environmental impacts and costs may help identify trade-offs and pathways for improving the overall performance of electrochemical DAC technologies.

AKNOWLEGMENTS

This work was funded by the German Federal Ministry for Education and Research (BMBF) within the project HyInnoHECC of SupplHyInno Rhineland (grant number 03ZU1115I). We are grateful to Lara Meuleneers for valuable discussions.

AUTHOR IDENTIFIERS

Author ORCIDs:
Prokopou G.I.: 0009-0007-2643-6899
Drakopoulou Z.: 0009-0009-8064-1898
Bongartz D.: 0000-0003-1790-0235
Mitsos A.: 0000-0003-0335-6566

REFERENCES

1. Kibria Nabil S, McCoy S, Kibria MG. Comparative life cycle assessment of electrochemical upgrading of CO ₂ to fuels and feedstocks. Green Chem. 23:867-880 (2021). https://doi.org/10.1039/d0gc02831b
2. de Jonge MMJ, Daemen J, Loriaux JM, Steinmann ZJN, Huijbregts MAJ. Life cycle carbon efficiency of direct air capture systems with strong hydroxide sorbents. International Journal of Greenhouse Gas Control 80:25-31 (2019). https://doi.org/10.1016/j.ijggc.2018.11.011
3. Keith DW, Ha-Duong M, Stolaroff JK. Climate strategy with co2 capture from the air. Climatic Change 74:17-45 (2005). https://doi.org/10.1007/s10584-005-9026-x
4. Liu G, Yang A, Darton RC. Numerical modeling and comparative analysis of electrolysis and electrodialysis systems for direct air capture. ACS Sustainable Chem. Eng. 12:3951-3965 (2024). https://doi.org/10.1021/acssuschemeng.3c06259
5. Young J, McQueen N, Charalambous C, Foteinis S, Hawrot O, Ojeda M, Pilorgé H, Andresen J, Psarras P, Renforth P, Garcia S, van der Spek M. The cost of direct air capture and storage can be reduced via strategic deployment but is unlikely to fall below stated cost targets. One Earth 6:899-917 (2023).

https://doi.org/10.1016/j.oneear.2023.06.004
6. Rosen N, Welter A, Schwankl M, Plumeré N, Staudt J, Burger J. Assessment of the potential of electrochemical steps in direct air capture through techno-economic analysis. Energy Fuels 38:15469-15481 (2024). https://doi.org/10.1021/acs.energyfuels.4c02202
7. Cuéllar-Franca RM, Azapagic A. Carbon capture, storage and utilisation technologies: a critical analysis and comparison of their life cycle environmental impacts. Journal of CO2 Utilization 9:82-102 (2015). https://doi.org/10.1016/j.jcou.2014.12.001
8. Hu Y, Gani R, Sundmacher K, Zhou T. Assessing the future impact of 12 direct air capture technologies. Chemical Engineering Science 298:120423 (2024). https://doi.org/10.1016/j.ces.2024.120423
9. Chauvy R, Dubois L. Life cycle and techno?economic assessments of direct air capture processes: an integrated review. Intl J of Energy Research 46:10320-10344 (2022). https://doi.org/10.1002/er.7884
10. Keith DW, Holmes G, St Angelo D, Heidel K. A process for capturing CO2 from the atmosphere. Joule 2:1573-1594 (2018). https://doi.org/10.1016/j.joule.2018.05.006
11. Bouaboula H, Belmabkhout Y, Zaabout A. Integrated CO2 capture and green hydrogen production: a promising approach for energy and cost reductions. Chemical Engineering Journal 509:161361 (2025). https://doi.org/10.1016/j.cej.2025.161361
12. von der Assen N, Voll P, Peters M, Bardow A. Life cycle assessment of co₂capture and utilization: a tutorial review. Chem. Soc. Rev. 43:7982-7994 (2014). https://doi.org/10.1039/c3cs60373c
13. von der Assen N, Jung J, Bardow A. Life-cycle assessment of carbon dioxide capture and utilization: avoiding the pitfalls. Energy Environ. Sci. 6:2721 (2013). https://doi.org/10.1039/c3ee41151f
14. Sphera. GaBi Education Software and Database. Edition 2024. https://www.gabi-software.com/
15. Frischknecht R, Jungbluth N, Althaus HJ, Doka G, Dones R, Heck T, Hellweg S, Hischier R, Nemecek T, Rebitzer G, Spielmann M. The ecoinvent database: overview and methodological framework (7 pp). Int J Life Cycle Assessment 10:3-9 (2004). https://doi.org/10.1065/lca2004.10.181.1
16. Zachert L, Suermann M, Bensmann B, Hanke-Rauschenbach R. Energetic evaluation and optimization of hydrogen generation and compression pathways considering PEM water electrolyzers and electrochemical hydrogen compressors. J. Electrochem. Soc. 168:014504 (2021). https://doi.org/10.1149/1945-7111/abcf1a
17. Bareiß K, de la Rua C, Möckl M, Hamacher T. Life cycle assessment of hydrogen from proton exchange membrane water electrolysis in future energy systems. Applied Energy 237:862-872 (2019). https://doi.org/10.1016/j.apenergy.2019.01.001
18. Curtin DE, Lousenberg RD, Henry TJ, Tangeman PC, Tisack ME. Advanced materials for improved PEMFC performance and life. Journal of Power Sources 131:41-48 (2004). https://doi.org/10.1016/j.jpowsour.2004.01.023
19. Liu CM, Sandhu NK, McCoy ST, Bergerson JA. A life cycle assessment of greenhouse gas emissions from direct air capture and fischer–tropsch fuel production. Sustainable Energy Fuels 4:3129-3142 (2020). https://doi.org/10.1039/c9se00479c
20. Madhu K, Pauliuk S, Dhathri S, Creutzig F. Understanding environmental trade-offs and resource demand of direct air capture technologies through comparative life-cycle assessment. Nat Energy 6:1035-1044 (2021). https://doi.org/10.1038/s41560-021-00922-6
21. Burgers I, Dehdari L, Xiao P, Li KG, Goetheer E, Webley P. Techno-economic analysis of PSA separation for hydrogen/natural gas mixtures at hydrogen refuelling stations. International Journal of Hydrogen Energy 47:36163-36174 (2022). https://doi.org/10.1016/j.ijhydene.2022.08.175
22. Pärnamäe R, Mareev S, Nikonenko V, Melnikov S, Sheldeshov N, Zabolotskii V, Hamelers HVM, Tedesco M. Bipolar membranes: a review on principles, latest developments, and applications. Journal of Membrane Science 617:118538 (2021). https://doi.org/10.1016/j.memsci.2020.118538
23. Schropp E, Campos-Carriedo F, Iribarren D, Naumann G, Bernäcker C, Gaderer M, Dufour J. Environmental and material criticality assessment of hydrogen production via anion exchange membrane electrolysis. Applied Energy 356:122247 (2024). https://doi.org/10.1016/j.apenergy.2023.122247
24. Deutz S, Bardow A. Life-cycle assessment of an industrial direct air capture process based on temperature–vacuum swing adsorption. Nat Energy 6:203-213 (2021). https://doi.org/10.1038/s41560-020-00771-9

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Comparative Techno-economic and Environmental Evaluation of Single-Step vs. Dual-Step CO_2-to-Methanol Processes using Multiobjective Optimization

Biswarup Mondal[a]*, Johannes Leipold[b], and Achim Kienle[ab]

[a] Max Planck Institute for Dynamics of Complex Technical Systems, Sandtorstraße 1, Magdeburg, D-39106, Germany
[b] Otto von Guericke University, Universitätsplatz 2, Magdeburg, D-39106, Germany
* Corresponding Author: mondal@mpi-magdeburg.mpg.de.

ABSTRACT

CO_2-to-methanol process is an attractive option to simultaneously reducing the anthropogenic CO_2 while producing value-added chemicals. In this work, two distinct CO_2-to-methanol process routes specifically, single step and dual step are evaluated based on their economic and environmental performance. First, a multiobjective optimization (MOO) framework is formulated to develop the optimal process configurations. Three conflicting objectives including methanol production rate, total annual cost (TAC) and carbon intensity of methanol are considered. For this MOO, the elitist non-dominated sorting genetic algorithm (NSGA-II) is employed to get the Pareto front. From the Pareto front, a balanced compromise solution is identified by the technique for order of preference by similarity to ideal solution (TOPSIS) with entropy information as weighting criteria. Then, the comparative performance analysis is conducted across the Pareto front. At the TOPSIS-selected configuration, the single step process offers economic advantages (1.97% lower levelised cost of methanol (LCoM)), but performs poorly on environmental perspective (3.03% higher carbon intensity of methanol). In contrast, at the cost-optimal point the dual step process is marginally favourable in both aspects (0.55% less LCoM and 1.26% reduction in carbon intensity of methanol). Overall, both processes exhibit comparable economic and environmental outcomes, with the dual step showing better environmental performance and reduced variability across the Pareto space.

Keywords: CO2-to-methanol, Process design, Multiobjetive optimization, Techno-economic assessment, Environmental performance

INTRODUCTION

Methanol has an extensive industrial demand, reaching around 112 million metric tons per year in 2024, owing to its versatile applications [1]. It serves as a building block for the production of numerous chemicals, such as formaldehyde, acetic acid, methyl tert-butyl ether (MTBE), and dimethyl ether (DME), among others. In addition, it exhibits favorable combustion properties, which make it a potential alternative fuel and energy carrier. Traditionally, methanol is produced from synthesis gas ($CO/H_2/CO_2$) using heterogeneous $Cu/ZnO/Al_2O_3$ catalyst. The synthesis gas is typically derived from fossil resources, with around 80% obtained through natural gas reforming and the remainder from coal gasification [2]. However, in recent time significant interests are being concentrated on 'green' methanol processes by replacing fossil resources with renewables. In this regard, CO_2 hydrogenation through the Power-to-Methanol concept emerges as an attractive option due to its dual benefit. It mitigates global warming by lowering or even reversing the CO_2 emission while simultaneously producing value added products. The CO_2 is obtained from industrial flue gases or, direct air capture and the green H_2 from electrolysis of water using renewable energy (e.g. solar, wind etc.

Main reactions, which are involved in the heterogeneously catalyzed synthesis of methanol are:

Hydrogenation of CO_2:

$$CO_2 + 3H_2 \xleftrightarrow{\Delta H = -49.4\, kJ.mol^{-1}} CH_3OH + H_2O \qquad (1)$$

Hydrogenation of CO:

$$CO + 2H_2 \xleftrightarrow{\Delta H = -90.6\ kJ.mol^{-1}} CH_3OH \quad (2)$$

Reverse water gas shift (RWGS):

$$CO_2 + H_2 \xleftrightarrow{\Delta H = 41.2\ kJ.mol^{-1}} CO + H_2O \quad (3)$$

Typically, hydrogenation of CO_2 is much faster than hydrogenation of CO [3]. Nevertheless, a CO rich feed is favorable because it consumes water during water gas shift, which limits the conversion of the resulting CO_2 hydrogenation to methanol. In contrast, from a chemical perspective, a CO_2 rich feed is not so favorable because it produces water during the reverse water gas shift.

In the standard process, a CO rich feed (i.e., synthesis gas) is used. Accordingly, two options are available for the new CO_2-to-methanol process, i.e., a single-step process converting the CO_2 rich feed directly to methanol or, a dual-step process, where CO_2 is first converted to CO in a separate reverse water gas shift reactor. Afterwards, the CO rich feed is converted in a second reactor to methanol following the standard process. The dual step process offers additional potential for improvement by using different catalysts and operating conditions in the different reactors [4]. In this context, a comparative techno-economic analysis is required to find out which process option is the best for CO_2-to-methnol production. To the best of our knowledge, limited investigations are available in this direction. Previous studies mainly compared both processes using relatively simple reaction kinetics and basic process analyses, often based on sensitivity studies [5-7]. Optimization is not applied for the rigorous evaluation of single and dual step processes. Therefore, in this contribution, a techno-economic analysis is performed using optimization and a rigorous kinetic model for both processes.

First, fixed bed reactor based single step and dual step processes are developed. Commercial $Cu/ZnO/Al_2O_3$ and $Ni/Al_{12}O_{19}$ catalysts are employed for the methanol synthesis and RWGS reactions, respectively. This study considers the mechanistic model of the methanol synthesis reaction based on the Langmuir-Hinshelwood kinetic model [8], with the corresponding model parameters adopted from Seidel et al., 2021 [9]. The kinetics of RWGS reaction are taken from literature [10]. Subsequently, the economic and environmental performance of the processes are assessed by total annual cost and carbon intensity of methanol. The economic analysis accounts for capital cost, operating cost, labour cost and raw material cost. Carbon intensity of methanol is defined as the total CO_2 emissions associated with the process per unit of methanol produced. Finally, the optimal process configuration, including design and operating parameters, is determined through a multiobjective optimization (MOO) technique using the elitist non-dominated sorting genetic algorithm-II (NSGA-II). The MOO problem is formulated with three conflicting objectives: maximizing methanol production rate, minimizing total annual cost and minimizing carbon intensity of methanol. To identify the final optimal point, the technique for order of preference by similarity to ideal solution (TOPSIS) with entropy information as a weighting criterion is implemented. Subsequently, both processes are evaluated based on their optimal configurations.

PROCESS DESCRIPTION

In this study, two process routes are analyzed/considered for the production of methanol from CO_2 namely, single step and dual step process. In both processes, pure CO_2 and H_2 are mixed and introduced into the reactor following a multistage compressor coupled with intermediate heat exchangers. This multistage compressor and intermediate heat exchanger arrangement are designed such that the desired operating pressure and temperature can be achieved. For the large scale production the feed flow rate is assumed to be fixed at 1000 mol/sec for both the processes.

Figure 1 demonstrates the single step process configuration. The compressed feed is combined with the recycle stream and preheated in a heat exchanger to attain the desired reaction temperature. Such heat integration provides a clear advantage, as reported in the literature [11].The mixed stream is introduced to an isothermal fixed bed type reactor. The reactor is packed with $Cu/ZnO/Al_2O_3$ catalyst and the specified temperature of the reactor is maintained by a water-cooled arrangement. The density of the catalyst is 770 kg/m^3 with the bed void fraction of 0.4. The outlet from the reactor is cooled down in a heat exchanger and introduced in a flash separator. The flash separator is operated at high pressure and low temperature at isothermal and isobaric condition. In the separator the gases are separated from liquid products. The gas stream mainly consists unreacted CO_2, CO and H_2; whereas liquid stream consists of methanol and water. A part of the gas stream is purged (1%) and the remaining is recycled (99%) back to the reactor. On the other hand, the liquid stream is then directed to a distillation column to get the pure methanol.

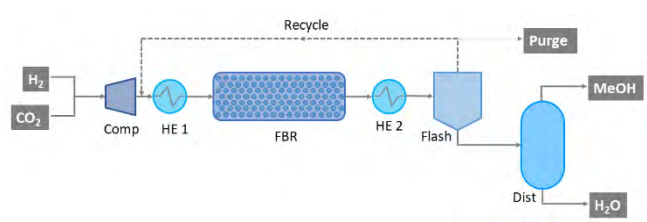

Figure 1. Schematic of the single step CO_2-to-methanol process route.

The dual step process configuration is shown in Figure 2, where the compressed feed is partially converted to CO in a RWGS reactor. The outlet stream of the RWGS reactor is then cooled down and introduced into a flash separator to separate the water. The remaining vapor stream, composed of CO_2, CO and H_2 is then compressed and subsequently mixed with the recycle stream. This steam is fed into the methanol synthesis reactor through a heat exchanger, where it is attained the desired temperature. The output of the second reactor is then cooled using a heat exchanger and directed into a flash separator. The vapor stream from the flash separator, mainly contains CO_2, CO and H_2 is recycled back after a small purge, similar to the single step process. The liquid stream of the flash separator contains methanol and water, which are separated from each other in a distillation column.

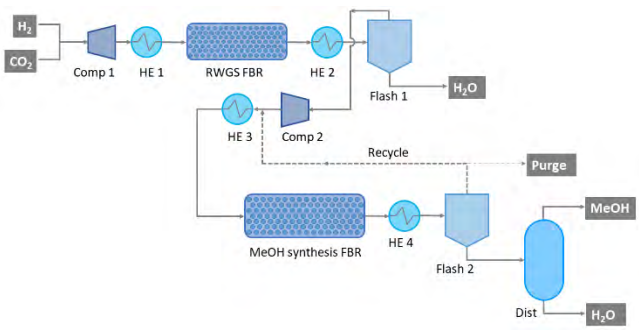

Figure 2. Schematic of the dual step CO_2-to-methanol process route.

MATHEMATICAL MODELING

This section highlights the mathematical model formulation of single and dual step process configurations. It consists modeling of reaction kinetics, fixed bed reactor and flash separator.

Reaction Kinetics

The kinetic model proposed by Seidel et al. [8], with reported parameters [9], is adopted here for the methanol synthesis reactor over $Cu/ZnO/Al_2O_3$ catalyst. Three reactions occurred simultaneously in the second reactor, as mentioned in Section 1. The corresponding rate expressions are described as follows,

$$r_{CO} = (1-\varnothing)k_1\left(p_{CO}p_{H_2}^2 - \frac{p_{CH_3OH}}{K_{P_1}}\right)\Theta_{CO}\Theta_{H_2}^4 \quad (4)$$

$$r_{CO_2} = \varnothing^2 k_2\left(p_{CO_2}p_{H_2}^2 - \frac{p_{CH_3OH}p_{H_2O}}{K_{P_2}p_{H_2}}\right)\Theta_{CO_2}^2\Theta_{H_2}^4 \quad (5)$$

$$r_{rwgs} = \frac{\varnothing}{1-\varnothing}k_3\left(p_{CO_2} - \frac{p_{CO}p_{H_2O}}{K_{P_3}p_{H_2}}\right)\Theta_{CO_2}\Theta_{CO} \quad (6)$$

The reaction constants are calculated from the extended Arrhenius equation, where *j* represents the reactions (1, 2 or 3).

$$k_j = exp\left(\beta_{2j-1} - \beta_{2j}\left(\frac{T_{ref}}{T} - 1\right)\right) \quad (7)$$

On the other hand, synthesis gas is produced in the RWGS reactor over the commercial Ni-catalyst (Ni/Al_2O_3). The reaction kinetics is given by:

$$r_{syn} = k_0 e^{\left(-\frac{E}{RT}\right)}\left(C_{CO_2} - \frac{C_{CO}C_{H_2O}}{K_c C_{H_2}}\right) \quad (8)$$

The model parameters are adopted from the work of [10].

Process Modeling

In this study, both RWGS and methanol synthesis reaction takes place in an isothermal fixed bed reactor (FBR). The reactor is operated at a constant pressure and considered no axial mixing and ideal gas phase. The FBR model is developed considering a cascade of CSTRs, where the number of stages correlates with the degree of back-mixing [12]. The corresponding material balance for each component in the gas phase is expressed as,

$$n^G\frac{dy_i}{dt} = \dot{n}_{feed}y_{i,0} - \dot{n}_{out}y_i + m_{cat}\sum_{j=1}^{NR}\nu_{i,j}r_j \quad (9)$$

Where, n^G denotes total mole present in the gas phase of the reactor. $\dot{n}_{feed}$ and $\dot{n}_{out}$ are molar flow rates of the feed and outlet of the reactor (mol/sec). y_i represents the gas phase mole fraction of component *i*. $\nu_{i,j}$ is the stoichiometric coefficient of component *i* in reaction *j*, r_j is the rate of reaction *j* (mol/sec/kg_{cat}) among the total number of reaction *NR*.

The summation of the component material balances over all species (*NC*) yields the total material balance, which is used to compute the outlet molar flow rate ($\dot{n}_{out}$).

$$0 = \dot{n}_{feed} - \dot{n}_{out} + m_{cat}\sum_{i=1}^{NC}\sum_{j=1}^{NR}\nu_{i,j}r_j \quad (10)$$

The outlet stream from the reactor is fed into an ideal flash separator, which is operated at a fixed temperature and pressure. The flash process is computed by the Rachford-Rice method, which is mathematically described by component material balance and total material balance.

$$Fz_i = Vy_i + Lx_i \quad (11)$$

$$F = V + L \quad (12)$$

$$y_i = K_i x_i = \frac{p_i^{sat}}{P}x_i \quad (13)$$

Here, the input flow rate to the flash separator is *F* (mol/sec), with composition *z*. *V* and *L* are the outlet flow rate of the vapor and liquid stream, respectively (mol/sec) with corresponding composition *y* and *x*. *K* is the vapor-liquid equilibrium coefficient of the ideal solution, which is the ratio of saturation pressure and total pressure. The saturation pressure is calculated using Antoine equation.

By substituting equation 13, into the material balance equations, we get the Rachford-Rice equation for the vapor fraction *δ=V/F*,

Table 1: Singificant decision variables with their ranges

Single step process		Dual step process	
Decision variables	**Range**	**Decision variables**	**Range**
Pressure of methanol synthesis reactor (bar)	25 – 85	Pressure of RWGS reactor (bar)	5 – 25
Temperature of methanol synthesis reactor (K)	450 – 650	Temperature of RWGS reactor (K)	900 – 1500
CO_2 mole fraction in feed	0.15 – 0.55	CO_2 mole fraction in feed	0.15 – 0.55
Catalyst amount methanol synthesis reactor (kg)	30000 – 150000	Catalyst amount of RWGS (kg)	100 – 5000
		Pressure of methanol synthesis reactor (bar)	35 – 85
		Temperature of methanol synthesis reactor (K)	450 – 650
		Catalyst amount of methanol synthesis Reactor (kg)	10000 – 150000

Table 2: TOPSIS-selected optimal process configurations

Single step process		Dual step process	
Decision variables	**Values**	**Decision variables**	**Values**
Pressure of methanol synthesis reactor (bar)	76.1	Pressure of RWGS reactor (bar)	18.07
Temperature of methanol synthesis reactor (K)	519.82	Temperature of RWGS reactor (K)	1418.82
CO_2 mole fraction in feed	0.2486	CO_2 mole fraction in feed	0.2496
Catalyst amount methanol synthesis reactor (kg)	150000	Catalyst amount of RWGS (kg)	1507.26
		Pressure of methanol synthesis reactor (bar)	69.2
		Temperature of methanol synthesis reactor (K)	483.35
		Catalyst amount of methanol synthesis reactor (kg)	108201

$$F(\delta) = \sum_{i=1}^{NC} \frac{z_i}{[1-\delta(1-K_i)]} - 1 \quad (14)$$

The Newton Raphson method is employed to find out the root of the above equation and correspondingly solve the flash calculations. The ODEs involved in the reactor analysis are solved by the ODE15s solver, and all simulation are conducted in the MATLAB R2022b environment.

PERFORMANCE ASSESSMENT

The processes are evaluated based on an economic and environmental perspective. For this, three primary performance indicators are used namely, methanol production rate (kg/hr), total annual cost (TAC) and carbon intensity of methanol (kg-eq CO_2/kg MeOH).

Total annual cost is expressed as:

$$TAC\ (€/yr) = (ACCR \times CAPEX) + OPEX \quad (15)$$

Where, CAPEX and OPEX represents capital and operating expenditures respectively. The annual capital charge ratio (ACCR) for this study is assumed 10%. The CAPEX includes direct process equipment cost (e.g., reactor, flash separator, distillation column, compressor and heat exchanger), indirect or, auxiliary unit cost, engineering, contingency and working capital. On the other hand, OPEX corresponding to fixed operating cost and variable operating cost. Fixed operating includes salary, supervision, salary overhead, maintenance, land, taxes and insurance and general overhead cost, whereas variable operating considers feed, catalyst and utility costs. The feed CO_2 and H_2 cost is 0.045 and 3.15 €/kg, respectively. Catalyst cost is 8.7 €/kg with 4 years of lifetime. The utilities includes electricity 78.51 €/MWh and cooling

water 0.03 €/ton. Steam methane reforming is considered as the H_2 source due to its lower cost and stable supply. The details of the cost functions can be found elsewhere [13] and the calculations are done based on the year 2024.

The carbon intensity is amount equivalent CO_2 emission to produce unit amount of methanol. This accounts the CO_2 emission due to the production of H_2, electricity and CO_2 present in the purge gas. Additionally, a secondary economic metric is computed from the first two i.e., methanol production rate and TAC, which is levelised cost of methanol (LCoM), expressed as,

$$LCoM\ (€/kg) = \frac{TAC}{Annual\ MeOH\ production} \tag{16}$$

MULTIOBJECTIVE OPTIMIZATION

The efficient design of the process and its performance rely on identifying the optimal design and operating parameters. In order to achieve this a multiobjective optimization strategy is incorporated based on non-dominated sorting genetic algorithm (NSGA-II). The primary performance indicators mentioned in the previous section are chosen as objectives, to take into account the economic and environmental aspects. Thus, the optimization problem is framed with the three conflicting objectives, including maximizing methanol production rate, minimizing TAC and minimizing carbon intensity of methanol.

RESULTS AND DISCUSSION

In this section, the optimization results and based on this the performance of both the processes are analyzed. The optimization is performed based on the stated objectives and a nonlinear constraint of LCoM < 2 €/kg. The significant decision variables and their feasible regions are primarily determined by sensitivity analysis to avoid the occurrence of infeasible solutions during the optimization. Accordingly, the lower and upper bounds are identified and listed in Table 1. Notably, the total feed flow rate is same for both the processes, while CO_2:H_2 molar ratio is varying as a decision variables. It ensures identical plant capacity for the fair comparison.

The optimization is carried out over the mentioned decision variable bounds with a population size and number of generations of 30. The NSGA-II parameters are adopted from the literature [14]. The stochastic variability of the algorithm is reduced by performing two independent runs. Then the non-dominated solutions from both runs are combined, and non-dominated sorting (NDS) is performed to find the 'true' Pareto optimal front. Here, the 'true' Pareto front consists of total of 53 and 51 non-dominated solutions for the single step and dual step processes, respectively. These solutions represent trade-offs among the competing objectives, such that improvement in one objective can only be achieved at the expense of deterioration in at least one other objective. However, the TOPSIS method with entropy weighting is employed to identify the decision-preferred compromise solution that reflects the balanced trade-off among all objectives. Accordingly, the design and operating parameters of the optimal process configurations are summarized in Table 2.

The Pareto front for the single step process is provided in Figure 3 (A), which demonstrates the trade-off among the three objectives. It can be noticed that all optimal solutions are in the range of 18171.3 to 27821.1 (kg/hr) for methanol production rate, 177.25 to 171.7 (million €) for TAC and 2.579 to 4.713 (kg-eq CO_2/kg MeOH) for carbon intensity of methanol. Correspondingly, the secondary economic metric, i.e., LCoM, varies from 0.7964 to 1.1811 €/kg within the Pareto space. These values represent percent variations of around 37% for methanol production, 3% for TAC and 45% for carbon intensity. Notably, methanol production and carbon intensity show a wide range of variation, because no greenness constraint is imposed apart from the economic constraint (LCoM < 2 €/kg), which preserves the natural trade-off among the objectives. The TOPSIS selected optimum point corresponds to methanol production rate of 27821.1 kg/hr, TAC of 177.25 €/yr and carbon intensity of methanol is 2.579 kg-eq CO_2/kg MeOH, that yields LCoM of 0.7964 €/kg.

On the other hand, Figure 3 (B) illustrates the Pareto front of the dual step process in a three dimensional space. The methanol production rate, TAC and carbon intensity of methanol are varied from 18916.4 to 28322.2 (kg/hr), 173.36 to 184.01 (million €) and 2.5 to 4.23 (kg-eq CO_2/kg MeOH), correspondingly, LCoM changes from 0.792 to 1.1456 €/kg across the Pareto space. For this case, the recommended optimum point yields objective values of 28322.2 kg/hr, 184.01 million €/yr and 2.501 kg-eq CO_2/kg MeOH. The corresponding LCoM is 0.8121 €/kg. An interesting trend can be observed in the Pareto front (Figure 3), where higher methanol production leads to lower carbon intensity. This trend is due to the trade-off between the three objectives and differences in operating conditions. Pareto points with lower TAC produce less methanol, resulting in higher CO_2 emission intensity, as it is the ratio of CO_2 emission to methanol production rate.

Now, a comparative performance analysis between the two processes based on the TOPSIS-selected optimal configuration are conducted. It is evident from the Pareto-curve trend that higher methanol production demands higher TAC but lower carbon intensity. The dual process achieves 1.80% higher methanol production rate than single step at the expense of 3.81% higher TAC. Consequently, LCoM from single step is 1.97% lower than

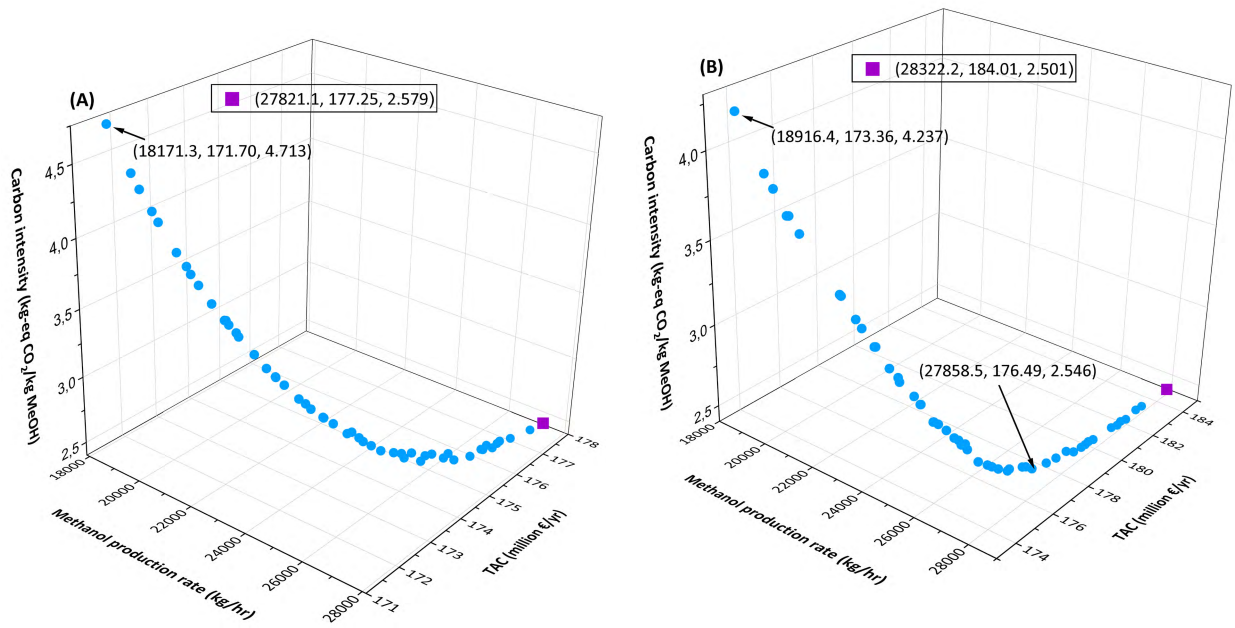

Figure 3: Pareto optimal front of the three objectives for the single step (A) and dual step (B) processes; violet squares in these plots corresond to the chosen optimal solution by TOPSIS

dual step process. However, from the environmental perspective dual step outperforms single step by 3.03% lower carbon intensity of methanol.

It is interesting to note that the TOPSIS-selected optimal point and cost-optimal point are the same for the single step process but not for the dual step process. Therefore, at the cost-optimal point (lowest methanol price), the dual-step process attains 0.55% lower LCoM and 1.26% lower carbon intensity than the single-step process. This indicates that, under strict cost-minimization objectives, the dual-step process offers a more favorable economic and environmental performance.

Further, we evaluated the impact of green H_2 on the TOPSIS selected optimal configurations. For the single-step process, TAC increases to 322.08 million €/yr, methanol production remains the same, carbon intensity reduces to 0.2861 kg-eq CO_2/kg MeOH and LCoM rises to 1.4471 €/kg. For the dual-step process, TAC is 328.65 million €/yr, methanol production rate remains unchanged, carbon intensity 0.2511 kg-eq CO_2/kg MeOH and LCoM 1.4505 €/kg. Results indicate that using green H_2 produced by electrolysis significantly reduces carbon intensity of methanol while increasing cost. However, the relative performance of both the processes will remain unchanged.

CONCLUSION

In this study, a comparative assessment of single step and dual step process is performed for the production of methanol from CO_2. The optimal process configurations are determined using the NSGA-II coupled with TOPSIS entropy method. This multiobjective optimization problem is formulated considering economic and environmental objectives namely, methanol production rate, TAC and carbon intensity of methanol. Based on TOPSIS-selected optimal configuration, single step process offers better economic performance of 1.97% lower LCoM than dual step process. However, the environmental performance is degraded significantly, accounting 3.03% more carbon intensity. Notably, at the cost-optimal configuration dual step process shows better performance in all the aspects, yielding 0.55% and 1.26% less LCoM and carbon intensity, respectively.

Overall, both processes exhibit comparable economic and environmental performance; however the dual step process demonstrates a narrower Pareto space, indicating slightly less variability in the trade-offs. In fact, for any economic trade-off across the Pareto front, including TOPSIS-selected, cost-optimal and highest price solutions the dual step consistently attains lower carbon intensity than the single step. This clearly suggests that the dual step is environmentally superior under equivalent economic criteria.

Here, the same catalyst is used for methanol synthesis reaction in both processes; however, recent advancements in catalyst research indicate the possibility of further improvement by employing more advanced catalysts for the single-step process. In this direction, it is crucial to investigate the impact of different catalysts,

additional process configurations and variable cost parameters for more comprehensive system-level insight.

ACKNOWLEDGEMENTS

Financial support by the German Research Foundation (DFG) is gratefully acknowledged through the priority program SPP 2331, Germany under grant KI 417/9-2.

AUTHOR IDENTIFIERS

Author ORCIDs:
Mondal B: 0000-0002-4868-3117
Leipold J: 0000-0002-1511-4794
Kienle A: 0000-0002-9011-0451

REFERENCES

1. The Methanol Institute. Methanol price and supply/demand. https://www.methanol.org/ (accessed October 2025)
2. Leonzio G, Zondervan E, Foscolo PU. Methanol production by CO2 hydrogenation: analysis and simulation of reactor performance. International Journal of Hydrogen Energy 44:7915-7933 (2019). https://doi.org/10.1016/j.ijhydene.2019.02.056
3. Slotboom Y, Bos MJ, Pieper J, Vrieswijk V, Likozar B, Kersten SRA, Brilman DWF. Critical assessment of steady-state kinetic models for the synthesis of methanol over an industrial cu/zno/al2o3 catalyst. Chemical Engineering Journal 389:124181 (2020). https://doi.org/10.1016/j.cej.2020.124181
4. Zang G, Sun P, Elgowainy A, Wang M. Technoeconomic and life cycle analysis of synthetic methanol production from hydrogen and industrial byproduct co₂. Environ. Sci. Technol. 55:5248-5257 (2021). https://doi.org/10.1021/acs.est.0c08237
5. Anicic B, Trop P, Goricanec D. Comparison between two methods of methanol production from carbon dioxide. Energy 77:279-289 (2014). https://doi.org/10.1016/j.energy.2014.09.069
6. Hren DT, Bogataj M, Nemet A. Methanol production via power-to-liquids: a comparative simulation of two pathways using green hydrogen and captured CO2. Processes 12:2843 (2024). https://doi.org/10.3390/pr12122843
7. Haghighatjoo F, Rahimpour MR, Farsi M. Techno-economic and environmental assessment of CO2 conversion to methanol: direct versus indirect conversion routes. Chemical Engineering and Processing - Process Intensification 184:109264 (2023). https://doi.org/10.1016/j.cep.2023.109264
8. Seidel C, Jörke A, Vollbrecht B, Seidel-Morgenstern A, Kienle A. Kinetic modeling of methanol synthesis from renewable resources. Chemical Engineering Science 175:130-138 (2018). https://doi.org/10.1016/j.ces.2017.09.043
9. Seidel C, Nikoli? D, Felischak M, Petkovska M, Seidel-Morgenstern A, Kienle A. Optimization of methanol synthesis under forced periodic operation. Processes 9:872 (2021). https://doi.org/10.3390/pr9050872
10. Samimi F, Karimipourfard D, Rahimpour MR. Green methanol synthesis process from carbon dioxide via reverse water gas shift reaction in a membrane reactor. Chemical Engineering Research and Design 140:44-67 (2018). https://doi.org/10.1016/j.cherd.2018.10.001
11. Leipold J, Kienle A. Nonlinear behavior of methanol synthesis through co2-hydrogenation. Chemical Engineering Journal 522:167176 (2025). https://doi.org/10.1016/j.cej.2025.167176
12. Leipold J, Seidel C, Nikolic D, Seidel-Morgenstern A, Kienle A. Optimization of methanol synthesis under forced periodic operation in isothermal fixed-bed reactors. Computers & Chemical Engineering 175:108285 (2023). https://doi.org/10.1016/j.compchemeng.2023.108285
13. Vázquez D, Guillén-Gosálbez G. Process design within planetary boundaries: application to CO2 based methanol production. Chemical Engineering Science 246:116891 (2021). https://doi.org/10.1016/j.ces.2021.116891
14. Mondal B, Jana AK. Optimal reflux splitting reactive distillation for algal biodiesel production: waste heat recovery through vapor recompression and organic rankine cycle. Separation and Purification Technology 292:121007 (2022). https://doi.org/10.1016/j.seppur.2022.121007

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Multi-Objective Optimization and Superstructure-Based Decision-Support Tool for Regional Low-Carbon Hydrogen Roadmaps: Methodology and Application to a region of Spain

Silvia Moreno[ab]*, Alejandro Aragón-García[a], Ángel L. Villanueva-Perales[ab], Bernabé Alonso-Fariñas[ab], and Pedro Haro[abc]

[a] Universidad de Sevilla, Chemical and Environmental Engineering Department, Seville, Spain
[b] Biorefinery Research Institute (I3B), Jaen, Spain
[c] Universidad de Sevilla, Laboratory of Engineering for Energy and Environmental Sustainability, Seville, Spain
* Corresponding Author: spuerta@us.es

ABSTRACT

Decarbonization of hydrogen-intensive industrial clusters is essential to meet the European Union's net-zero targets. Although hydrogen can replace fossil-based feedstocks and fuels in refineries and chemical industries, its production remains largely dependent on natural gas. Therefore, cost-effective and low-emission supply routes require a system-level approach that integrates regional resources, technologies, and industrial demand. This study applies a multi-objective optimization framework to design a low-carbon hydrogen supply system for Galicia (northwestern Spain), addressing two gaps in regional energy system modeling: model transferability across regions and integration of social criteria beyond techno-economic assessment. The model quantifies trade-offs between total system cost and greenhouse gas emissions, and an employment indicator is integrated via post-processing using TOPSIS. The results show that meeting 100% of the projected 2030 demand (105 kt H_2/a) yields a single feasibility-limited solution dominated by biomass gasification (69% of hydrogen production). When the demand was reduced to 60% coverage (63 kt H_2/a), 28 non-dominated solutions were obtained. The top-ranked compromise combines 55% PEM electrolysis and 45% biomass gasification, supported by a wind-dominated electricity mix, under equal weighting of economic, environmental, and social criteria. When either emissions or employment is prioritized, this compromise remains top-ranked, while the cost prioritization shifts preference to a solution with hydrogen production from PEM (51%) and alkaline (3%) electrolysis and biomass gasification (46%).

Keywords: Hydrogen, Optimization, Supply Chain, Biomass, Renewable and Sustainable Energy

INTRODUCTION

The deep decarbonization of hydrogen-intensive industrial clusters is required to align European industry with the legally binding climate objectives of the European Union for 2030 and climate neutrality by 2050 [1]. Hydrogen plays a structural role in refineries and chemical value chains, and near-term decarbonization options frequently rely on replacing fossil-derived hydrogen with low-carbon supply routes [2, 3]. Nonetheless, on a global scale, hydrogen production remains predominantly fossil-based, with natural gas still providing the largest share of dedicated hydrogen output [4]. This creates a dual challenge, i.e., low-carbon hydrogen must be implemented on a large scale, keeping costs under control, and avoiding the shift of carbon emissions throughout the life cycle [5].

A regional perspective is critical for this transition. The feasibility and performance of hydrogen pathways depend on location-specific factors such as renewable resource profiles, biomass availability, industrial demand, and infrastructure constraints [6, 7]. These drivers are

particularly relevant for regions with strong renewable and bio-based resources, where integrated local production of both hydrogen and electricity can reduce reliance on grid imports and improve system resilience. In Galicia (a region of $3.0 \cdot 10^4$ km2 and $2.7 \cdot 10^6$ inhabitants located in northwest Spain), hydrogen decarbonization has been actively discussed in relation to industrial needs and local resource potential, but economic, environmental performance and implementation barriers have constrained the pathway portfolio [8, 9].

Decision support for such systems requires methods that can capture discrete technological choices and competing system objectives. Optimization-based hydrogen supply chain (HSC) models, including multi-objective formulations that quantify trade-offs between cost and environmental indicators, have been widely used to support strategic planning [10–12]. Nevertheless, two practical gaps remain for regional deployment studies. Firstly, given the central role of spatial and temporal representation in energy system models, the practical applicability of optimization frameworks across regions with different resource portfolios and seasonal profiles remains an open question. Secondly, social aspects are frequently neglected or discussed qualitatively but are less often integrated in a transparent way that supports the selection of a balanced solution from a Pareto set [13–15].

This study addresses these gaps by applying a previously developed superstructure-based optimization model [16] to a new regional case study in Spain. The analysis assesses the model's applicability under region-specific resource and demand conditions, and its ability to generate feasible and interpretable cost-emission trade-offs for regional planning. As a novel contribution of this work, a quantitative social dimension is incorporated to enable the identification and ranking of preferred Pareto-optimal solutions, complementing the economic and environmental perspectives.

METHODOLOGY

This work applies a previously developed multi-objective optimization model to the region of Galicia to identify pathways for low-carbon hydrogen production based on renewable and biogenic resources. The methodological framework consists of four main steps: i) collection of regional resource availability data, ii) definition of techno-economic and environmental parameters for the technologies considered, iii) application of a cost-emissions multi-objective optimization model using the augmented e-constraint method, and iv) ex post evaluation of a social indicator to support the selection among Pareto-optimal solutions.

Regional data

The region of the case study (Galicia) is characterized by a high availability of renewable electricity resources and significant biogenic feedstock potential. Regional data on renewable energy potential, biogenic resource availability, and industrial water use were compiled from regional and national public reports [17, 18]. These data define upper bounds on resource exploitation and constrain the feasible technology deployment within the optimization model.

Renewable electricity and biogenic resource availability are modelled with monthly resolution to capture seasonal variability in generation and feedstock supply. Water availability for hydrogen production is constrained at the regional level based on official annual industrial water demand estimates [19]. The annual water availability is distributed evenly throughout the year to ensure consistency with the model's temporal structure.

A regional spatial representation is adopted. Transport is explicitly accounted for biomass-based resources. A representative transport distance is defined as the distance from the sub-region with the highest resource availability to the hydrogen production plant for each resource.

The hydrogen demand has been defined based on regional estimates for 2030 [8]. The monthly demand of hydrogen is assumed to be constant throughout the year to meet the time resolution of the model.

Techno-economic and environmental characterization of technologies

A superstructure of alternative technologies, including renewable-based and biogenic pathways, represents hydrogen production pathways. Each technology is characterized by a consistent set of technical, economic, and environmental parameters, including capital expenditure, operational costs, conversion efficiencies, and specific greenhouse gas emissions, which are expressed as CO_2-equivalent. All parameters are treated as deterministic and are kept consistent with those used in previous applications of the model in other European regions, allowing direct comparison.

Multi-objective optimization model

The optimization model adopted in this work was developed and documented in a previous study [16]. It determines the optimal configuration of hydrogen production pathways that satisfy hydrogen demand at minimum cost and minimum greenhouse gas emissions. The model is formulated as a mixed-integer linear programming problem and solved using Python-based optimization tools. Two objective functions are considered: i) minimization of total system cost, and ii) minimization of total greenhouse gas emissions. The Pareto frontier is generated using the augmented ε-constraint method, which ensures a well-distributed set of efficient solutions and

avoids weakly Pareto-optimal points [20]. The optimization model does not include social criteria, which are only introduced in a post-processing step.

Social indicator definition

A social indicator based on employment creation is introduced to differentiate among Pareto-optimal solutions. Employment creation is used as a proxy for social sustainability, as it is one of the most adopted metrics in the literature to represent the social dimension in hydrogen production assessments [21] and is relatively straightforward to quantify while remaining highly relevant for regional development.

For each Pareto-optimal solution i, the total employment is calculated a posteriori as the sum of the total number of jobs generated in the technologies active in that solution. The specific employment factors for each technology, expressed in jobs per unit of hydrogen produced or installed capacity, are assigned based on the values reported in the literature.

The total employment for solution i is calculated as follows:

$$S_i = \sum_k f_k \cdot H_{i,k} + \sum_r e_r \cdot P_{i,r} \quad (1)$$

where f_k is the employment factor for hydrogen production technologies (jobs/t H_2/a); $H_{i,k}$ is the annual hydrogen production from technology k (t H_2/a) in solution i; e_r is the employment factor for electricity generation technologies (jobs/MW_e); and $P_{i,r}$ is the installed capacity (MW_e) of electricity generation technology k in solution i. Technology-specific employment factors f_k and e_k and their literature sources are provided in the Supplementary Material (Table S3).

The social indicator is evaluated ex post and does not constrain the optimization problem. This choice avoids the introduction of highly uncertain social parameters into the feasible region and allows the economic-environmental trade-offs to be explored independently of social preferences.

Ranking of Pareto-optimal solutions (TOPSIS)

Subsequently, the set of Pareto-optimal solutions is ranked using the Technique for Order Preference by Similarity to the Ideal Solution (TOPSIS), considering the total cost, greenhouse gas emissions, and employment creation [22–24]. Cost and emissions are treated as criteria to be minimized, while employment is treated as a criterion to be maximized.

Let m be the number of Pareto-optimal solutions and n the number of indicators. The decision matrix is $X_{mxn} = [x_{ij}]$, where each row i corresponds to a solution and each column j to a criterion (e.g., cost, emission, employment). The criteria are first normalized using vector normalization as follows:

$$r_{ij} = \frac{x_{ij}}{\sqrt{\sum_{i=1}^{m} x_{ij}^2}} \quad (2)$$

Normalized values are then weighted using a weight vector w_j with $\sum_{j=1}^{n} w_j = 1$:

$$v_{ij} = w_j r_{ij} \quad (3)$$

The positive ideal solution A^+ and negative ideal solution A^- are defined as follows:

$$v_j^+ = \begin{cases} max_i v_{ij}, j \in B \\ min_i v_{ij}, j \in C \end{cases} \quad (4)$$

$$v_j^- = \begin{cases} min_i v_{ij}, j \in B \\ max_i v_{ij}, j \in C \end{cases} \quad (5)$$

where B denotes the benefit criteria (employment) and C cost criteria (economic cost and emissions). The distances to the ideal and anti-ideal are computed using the following Euclidean metrics:

$$D_i^+ = \sqrt{\sum_{j=1}^{n} \left(v_{ij} - v_j^+\right)^2} \quad (6)$$

$$D_i^- = \sqrt{\sum_{j=1}^{n} \left(v_{ij} - v_j^-\right)^2} \quad (7)$$

Finally, the TOPSIS closeness coefficient is given as follows:

$$CC_i = \frac{D_i^-}{D_i^+ + D_i^-} \quad (8)$$

Solutions are ranked in descending order of CC_i, with higher values indicating closer proximity to the ideal compromise across economic, environmental, and social dimensions.

Sensitivity analysis

A sensitivity analysis was conducted to test the robustness of the TOPSIS ranking to changes in the preference structure. The weight vector was modified to prioritize, the economic, environmental, or social criterion while the weights were kept normalized. Table 1 shows the different weights for each indicator in the different scenarios analyzed.

Table 1: Weights used for each indicator under different scenarios.

Scenario	Cost weight	Emission weight	Employment weight
Equal	0.33	0.33	0.33
Cost priority	0.80	0.10	0.10
Emission priority	0.10	0.80	0.10
Employment priority	0.10	0.10	0.80

RESULTS

Multi-objective optimization model

The projected 2030 hydrogen demand for Galicia (105 kt H_2/a) was initially imposed as an exogenous requirement. The model returns a single non-dominated solution at this demand level under the considered regional constraints, indicating a feasibility-limited regime. This solution has been reported for completeness in the Supplementary Material. To enable a comparative multi-objective assessment and apply the social dimension, additional runs were performed by scaling the demand to feasible regional supply levels. Accordingly, results are reported for a demand coverage level of 60% of the projected 2030 demand (i.e., 63 kt H_2/a), as higher coverage levels collapse the Pareto set to a single solution. The remaining demand could be met outside the modeled region by increasing the number of resources considered or using other conventional carbon capture technologies to reduce emissions. At this coverage level, the model yields a Pareto set with multiple non-dominated solutions (cf. figure 1), allowing the analysis of trade-offs including social performance.

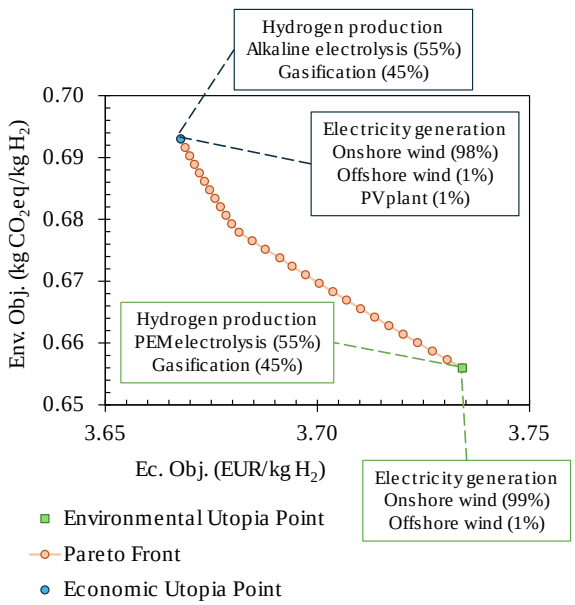

Figure 1. Pareto-optimal solutions obtained for a demand coverage level of 60% of the projected 2030 hydrogen demand (63 kt H_2/a) in Galicia. The percentages of technologies selected by the model for these solutions have been included in the economic and environmental utopia points.

Ranking of Pareto-optimal solutions (TOPSIS)

TOPSIS was applied to rank the 28 Pareto-optimal solutions (with a demand coverage level of 60%) using functional-unit indicators expressed per kg H_2 under equal weights for cost, emissions, and employment. The highest closeness coefficient was obtained from solution 0, which corresponds to the environmental optimum (3.73 €/kg H_2, 0.66 kg CO_2eq/kg H_2, $6.48 \cdot 10^{-6}$ jobs-year/kg H_2), followed by solutions 1 (3.66 €/kg H_2, 0.69 kg CO_2eq/kg H_2, $6.47 \cdot 10^{-6}$ jobs-year/kg H_2) and 2 (3.66 €/kg H_2, 0.69 kg CO_2eq/kg H_2, $6.46 \cdot 10^{-6}$ jobs-year/kg H_2). Solution 27, which coincides with the economic optimum, is the lowest-ranked solution. Overall, the TOPSIS order selects an interior compromise within the Pareto set. Solutions with lower costs tend to exhibit higher emissions and lower employment, whereas solutions with lower emissions tend to incur higher costs and higher employment. Under equal weights, the environmental optimum provides the closest balance to the ideal point across the three dimensions. Table 2 shows the top three solutions for the equal scenario's hydrogen production and electricity generation technologies.

Table 2: Share of hydrogen production and electricity generation technologies for the top three solutions in the ranking, assuming equal weighting for the three criteria considered.

Rank-Solution	H_2 production technologies	Electricity generation technologies
1º-0	PEM electrolysis 55% Gasification 45%	Onshore wind 99% Offshore wind 1%
2º-1	PEM electrolysis 51% Gasification 46% Alk. electrolysis 3%	Onshore wind 99% Offshore wind 1%
3º-2	PEM electrolysis 48% Gasification 45% Alk. electrolysis 7%	Onshore wind 99% Offshore wind 1%

Alk.: Alkaline
PEM: Proton exchange membrane

Sensitivity analysis

When the economic criterion is prioritized, solutions 1-3 become the top-ranked options, with solution 1 ranked first (3.66 EUR/kg H_2, 0.69 kg CO_2eq/kg H_2, $6.47 \cdot 10^{-6}$ jobs-year/kg H_2), followed closely by solutions 2 (3.66 EUR/kg H_2, 0.69 kg CO_2eq/kg H_2, $6.46 \cdot 10^{-6}$ jobs-year/kg H_2) and 3 (3.67 EUR/kg H_2, 0.69 kg CO_2eq/kg H_2, $6.46 \cdot 10^{-6}$ jobs-year/kg H_2). In this scenario, the lowest-ranked option shifts to the environmental optimum (solution 0) due to the strong emphasis on cost. When the environmental criterion is prioritized, the rankings shift toward low-emission solutions and solution 0 becomes the best-ranked option (3.73 EUR/kg H_2, 0.66 kg CO_2eq/kg H_2, $6.48 \cdot 10^{-6}$ jobs-year/kg H_2), while the economic optimum (solution 27) remains the worst-ranked option. When the social criterion is prioritized, the ordering becomes closer to the equal scenario, again selecting solution 0 as the top-ranked option, followed by solutions 1

and 3. Notably, solution 0 is the best-ranked option under both environmental and social priority scenarios, indicating that these two prioritizations converge on the same preferred configuration. Across scenarios, the worst-ranked solution is the economic optimum in all cases except under economic prioritization, where the environmental optimum becomes the lowest-ranked outcome as a direct consequence of assigning a dominant weight to cost.

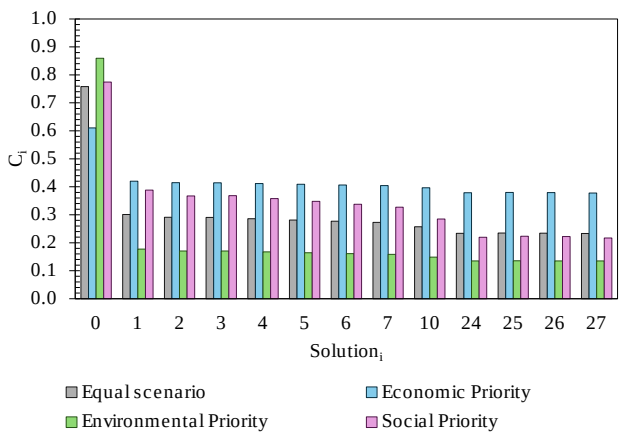

Figure 2. TOPSIS proximity coefficients for the best and worst ranked solutions in the equal scenario and sensitivity analysis cases.

Tables 3 and 4 show the contributions of each hydrogen production and electricity generation technology for the economic and the environmental priority scenario, respectively. The environmental priority case is closely aligned with the equal scenario; the corresponding solutions are shown in Table 2.

Table 3: Share of hydrogen production and electricity generation technologies for the top three economic priority scenario.

Rank-Solution	H_2 production technologies	Electricity generation technologies
1º-1	PEM electrolysis 51% Gasification 46% Alk. electrolysis 3%	Onshore wind 99% Offshore wind 1%
2º-2	PEM electrolysis 48% Gasification 45% Alk. electrolysis 7%	Onshore wind 99% Offshore wind 1%
3º-3	Gasification 46% PEM electrolysis 44% Alk. electrolysis 10%	Onshore wind 99% Offshore wind 1%

Table 4: Share of hydrogen production and electricity generation technologies for the top three social priority scenario.

Rank-Solution	H_2 production technologies	Electricity generation technologies
1º-0	PEM electrolysis 55% Gasification 45%	Onshore wind 99% Offshore wind 1%
2º-1	PEM electrolysis 51% Gasification 46% Alk. electrolysis 3%	Onshore wind 99% Offshore wind 1%
3º-3	Gasification 46% PEM electrolysis 44% Alk. electrolysis 10%	Onshore wind 99% Offshore wind 1%

CONCLUSIONS

This work applied a multi-objective optimization framework to design a low-carbon hydrogen supply system for Galicia. Quantified trade-offs between cost, emissions, and employment across optimal solutions. The model showed that meeting 100% of the projected 2030 demand (105 kt H_2/a) with the regional resources considered is feasible but yields a single solution. Under this full-coverage constraint, the system configuration is dominated by biomass gasification (69% of hydrogen production) and onshore wind power generation.

To enable a multi-criteria comparison and the integration of a social dimension through TOPSIS, the demand coverage was reduced to 60%, for which multiple efficient solutions exist. At 60% demand coverage (63 kt H_2/a), the model generated 28 non-dominated solutions. This Pareto set enabled a structured evaluation of competing objectives. The TOPSIS analysis with equal weights identified a stable compromise (solution 0), corresponding to the environmental optimum within the Pareto set. This solution combines 55% alkaline electrolysis with 45% biomass gasification. Its electricity supply is entirely wind-based (99% onshore wind and 1% offshore wind). This solution remained top-ranked when prioritizing either emissions or employment. This indicated robustness to stakeholder preferences. When prioritizing economic performance, the preferred option shifted to solution 1. This solution introduces electrolysis diversity, with 51% PEM electrolysis, 46% gasification, and 3% alkaline electrolysis. In all scenarios, the economic optimum was consistently the lowest-ranked option, except under economic prioritization where the environmental optimum became the worst-ranked outcome due to the dominant weight assigned to cost.

Overall, the framework proved transferable by producing an interpretable Pareto front and revealing binding regional constraints. The post-processing integration of employment via TOPSIS provided a transparent and repeatable basis for compromise selection. These results

provide a concrete evidence base for designing regional hydrogen strategies under resource constraints.

DIGITAL SUPPLEMENTARY MATERIAL

Supplementary material, including additional methodological details, input data, model assumptions, and extended results can be found at: https://psecommunity.org/LAPSE:2026.0033

ACKNOWLEDGEMENTS

This work was supported by project GH2T: grant PID2020-114725RA-I00 funded by MCIN/AEI/10.13039/50110 0 011033, project DOHVEN: SOL2024-31445 of the Programme FEDER Andalucía, and project SOL2024-31702 of the Programme FEDER Andalucía 2021–2027. The Ph.D. grant of S. Moreno from the FPU program of the Spanish Government (FPU22/00366) is acknowledged.

AUTHOR IDENTIFIERS

Author ORCIDs:
Moreno, S.: 0000-0001-5336-851X
Aragón-García, A.: 0009-0005-1527-5749
Villanueva-Perales, A.L.: 0000-0002-5403-7031
Alonso-Fariñas, B.: 0000-0001-9740-3364
Haro, P.: 0000-0003-1861-4062

REFERENCES

1. European Commission: Hydrogen. https://energy.ec.europa.eu/topics/eus-energy-system/hydrogen_en
2. International Energy Agency: The Future of Hydrogen. (2019). https://www.iea.org/reports/the-future-of-hydrogen
3. European Hydrogen Observatory: The European hydrogen market landscape. (2024). https://observatory.clean-hydrogen.europa.eu/
4. Iea: Global Hydrogen Review 2025. (2025). www.iea.org
5. Maniscalco MP, Longo S, Cellura M, Micciché G, Ferraro M. Critical review of life cycle assessment of hydrogen production pathways. Environments 11:108 (2024). https://doi.org/10.3390/environments11060108
6. Bolz S, Thiele J, Wendler T. Regional capabilities and hydrogen adoption barriers. Energy Policy 185:113934 (2024). https://doi.org/10.1016/j.enpol.2023.113934
7. Andersen, C.: Changing perspectives: Assessing regional potentials for green hydrogen infrastructure planning. AESOP Annual Congress proceedings. 36, (2024). https://proceedings.aesop-planning.eu/index.php/aesopro/article/view/2362
8. Pinotti G, Vence X. Hidrógeno verde y descarbonización en galicia: análisis multinivel de proyectos, motores y barreras para su desarrollo. RGE : (2025). https://doi.org/10.15304/rge.34.3.10430
9. Balsa Barreiro, J., Lois González, R.C., Miramontes Carballada, Á., Santana Rodrígues, A.P.: Galicia, https://atlasnacional.ign.es/wane/Galicia
10. Moreno-Benito M, Agnolucci P, Papageorgiou LG. Towards a sustainable hydrogen economy: optimisation-based framework for hydrogen infrastructure development. Computers & Chemical Engineering 102:110-127 (2017). https://doi.org/10.1016/j.compchemeng.2016.08.005
11. Quarton CJ, Samsatli S. The value of hydrogen and carbon capture, storage and utilisation in decarbonising energy: insights from integrated value chain optimisation. Applied Energy 257:113936 (2020). https://doi.org/10.1016/j.apenergy.2019.113936
12. Riera JA, Lima RM, Knio OM. A review of hydrogen production and supply chain modeling and optimization. International Journal of Hydrogen Energy 48:13731-13755 (2023). https://doi.org/10.1016/j.ijhydene.2022.12.242
13. Fodstad M, Crespo del Granado P, Hellemo L, Knudsen BR, Pisciella P, Silvast A, Bordin C, Schmidt S, Straus J. Next frontiers in energy system modelling: a review on challenges and the state of the art. Renewable and Sustainable Energy Reviews 160:112246 (2022). https://doi.org/10.1016/j.rser.2022.112246
14. Oliveira VESF, do Carmo BBT, da Costa Fontes FF. Hydrogen supply chains: modeling challenges and sustainability gaps. Applied Energy 401:126827 (2025). https://doi.org/10.1016/j.apenergy.2025.126827
15. Kachirayil, F., Weinand, J.M., Scheller, F., McKenna, R.: Reviewing local and integrated energy system models: insights

into flexibility and robustness challenges. (2022). https://arxiv.org/pdf/2202.13942
16. Moreno, S., Aragón-García, A., Villanueva-Perales, A.L., Alonso-Fariñas, B., Monaghan, R.F.D., Haro, P.: Multi-Objective Decision-Support Framework for Low-Carbon Hydrogen Supply Chains under Economic, Environmental, and Regulatory Constraints. SSRN Electronic Journal. WPS 6163626 (2026). https://ssrn.com/abstract=6163626
17. Pfenninger S, Staffell I. Long-term patterns of european PV output using 30 years of validated hourly reanalysis and satellite data. Energy 114:1251-1265 (2016). https://doi.org/10.1016/j.energy.2016.08.060
18. Staffell I, Pfenninger S. Using bias-corrected reanalysis to simulate current and future wind power output. Energy 114:1224-1239 (2016). https://doi.org/10.1016/j.energy.2016.08.068
19. Servizo de Planificación e Programación Hidrolóxica-Augas de Galicia: Plan hidrolóxico da demarcación hidrográfica de Galicia-Costa (Ciclo 2021-2027). (2023). https://augasdegalicia.xunta.gal/c/document_library/get_file?file_path=/portal-augas-de-galicia/plans/planHidroloxicoGC/2021_2027/Anexo_03_Usos_DD_gal.pdf
20. Mavrotas G. Effective implementation of the ?-constraint method in multi-objective mathematical programming problems. Applied Mathematics and Computation 213:455-465 (2009). https://doi.org/10.1016/j.amc.2009.03.037
21. Moreno S, Van den Auwelant E, Haro P, Alonso-Fariñas B. A conceptual framework for integrating social, economic, and environmental dimensions in the selection of hydrogen production pathways. Renewable and Sustainable Energy Reviews 230:116716 (2026). https://doi.org/10.1016/j.rser.2026.116716
22. Yoon K. A reconciliation among discrete compromise solutions. Journal of the Operational Research Society 38:277-286 (2017). https://doi.org/10.1057/jors.1987.44
23. Hwang, C.-L., Yoon, K.: Multiple Attribute Decision Making. 186, (1981). https://doi.org/10.1007/978-3-642-48318-9
24. Papathanasiou J, Ploskas N. TOPSIS. Springer Optimization and Its Applications :1-30 (2018). https://doi.org/10.1007/978-3-319-91648-4_1

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Electrified refineries in the Power Flow Network

Sampriti Chattopadhyay[a], Ana I. Torres[a], Ignacio E. Grossmann[a], Saif R Kazi[b]*

[a] Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh 15213, PA, USA
[b] Applied Mathematics & Plasma Physics Group, Theoretical Division, Los Alamos National Laboratory
* Corresponding author: skazi@lanl.gov

ABSTRACT

Industrial decarbonization has heightened interest in electrifying major chemical processes, but existing planning methods typically assume fixed electricity prices and overlook how industrial power use affects the grid. This work introduces a grid-aware optimization framework that captures two-way interactions between industrial electricity usage and the power flows within the grid. We use the DC Optimal Power Flow (DC-OPF) model to generate Locational Marginal Prices across refinery demand levels and embed a surrogate reflecting the relationship between the power demand and the prices into an operational optimization problem for a partially electrified refinery. The surrogate model is embedded within the optimization problem using disjunctive reformulations and off-the-shelf packages such as OMLT (Optimization and Machine Learning Toolkit). In a case study considering an oil refinery with installed electric boilers, electrolyzers, H_2 storage, and post-combustion carbon capture infrastructure, the grid-aware approach lowers operating costs by 7% relative to a price-taker model by anticipating how the refinery's own demand shifts electricity prices. The method is also shown to incorporate the effect of demand uncertainty at other grid nodes by embedding a surrogate model trained using data generated by a chance-constrained DC Optimal Power Flow.

Keywords: Energy Systems, Electricity & Electrical Devices, Surrogate Model, Refining, Process Operations

INTRODUCTION

The chemical industry is undergoing a fundamental transformation toward decarbonization and electrification, driven by the integration of renewable energy and novel low-carbon technologies. Investment and operational planning for this transition are interconnected, as the flexible operation of electrified systems enhances their economic viability. Previous literature typically utilizes multi-period superstructure optimization to determine optimal technology portfolios and operating schedules while satisfying emission targets and physical constraints [1–4, 24]. Stochastic and robust self-scheduling models outperform deterministic approaches by up to 10%, as they effectively hedge against price and renewable production uncertainty [5, 6].

Despite these advances, most optimization studies adopt a price-taker framework, assuming that individual facilities, whether acting on the supply or load side do not influence market prices. By relying on exogenous price profiles, these models overlook the impact of utility-scale loads (100+ MW) on the grid. In U.S. markets, Independent System Operators (ISOs) establish Locational Marginal Prices (LMPs) through a clearing process that dispatches resources optimizing for production cost subject to demand and transmission constraints. When large industrial facilities shift production or deploy storage, they directly alter the demand profile, shifting LMPs and potentially amplifying price volatility. Recent evaluations have shown that these price-taker assumptions in self-scheduling can lead to significant discrepancies when compared to system-scheduling solutions [7–9]. A few recent studies have more explicitly modeled the complex interactions between integrated energy systems (IESs) and the electric grid by embedding market-clearing simulations into facility models [10] or using machine-learning surrogates to capture price formation [11]. Price maker frameworks, on the other hand employ game-theoretic [13–15] or residual demand [12] methods. While previous literature has mostly focused on the supply side, recent work has begun to examine strategies such as demand bidding that have emerged as promising mechanisms for integrating large industrial consumers into power system operations. Tang et al [13] formulate

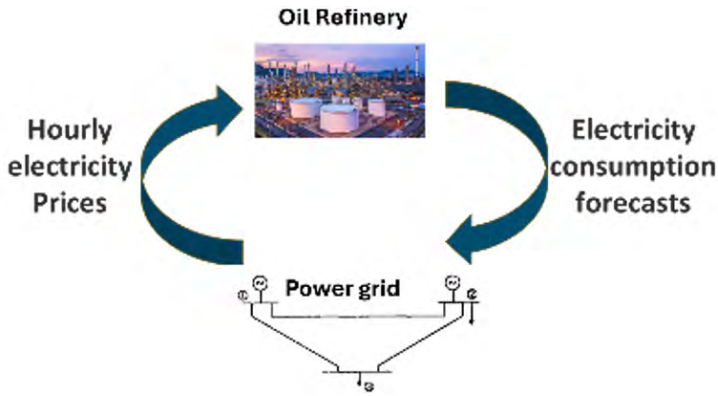

Figure 1. Interaction between industries and power grid

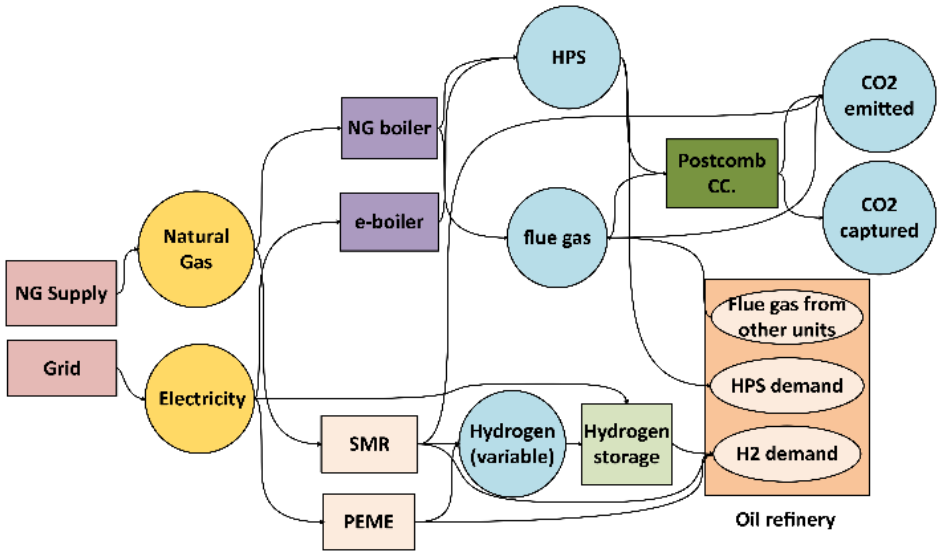

Figure 2. Flowsheet for partially electrified refinery

demand bidding within a modified AC optimal power flow framework, interpreting it as negotiation between operators and users. Their case study, implemented using a synthetic model of the Texas ISO, indicates that although the demand bidding framework decreases generation costs, its overall system benefits are restricted by transmission grid limitations and necessitate a level of data sharing not currently standard in the industry. This highlights a critical gap: the need for grid-aware demand-side strategies that allow industrial facilities to internalize grid limitations and price impacts within current data-sharing constraints.

The remainder of this paper is organized as follows. First, we introduce the problem statement. Then we present a multi-period linear program (LP) for operational decision-making in a partially electrified refinery model alongside standard and chance-constrained DC Optimal Power Flow (DCOPF) models to account for uncertainty. We then describe the surrogate modelling strategy capturing LMPs as a function of refinery load, which are embedded into the LP. Finally, we compare a price-taker baseline against these proposed scenarios and state concluding takeaways.

PROBLEM STATEMENT

Given a partially electrified refinery, where the refinery operator's decision depends on grid electricity prices, which in turn depend on power demand and supply (in the grid network). The public availability of historical load and price data from RTOs [23] allows the refinery to build surrogate models that characterize this relationship, effectively bypassing the need for full access to restricted grid-side data. Our goal is to develop a novel integrated approach that can be used by such refineries to make grid-aware decisions, capturing these price dynamics by embedding the trained surrogate relationship directly into the refinery's operational optimization.

METHODOLOGY

Optimization of refinery operations

The flowsheet for the considered partially electrified refinery is shown in Fig 1. The flowsheet is denoted by a resource task network [18-19] with circular nodes depicting the resource nodes and the rectangular nodes showing the task nodes. The steam requirements can be satisfied by a mix of a natural gas (NG) boiler and e-boiler. The H_2 requirements are satisfied by the steam methane reformer and proton exchange membrane electrolyzers (PEME). The optimization problem minimizing operational costs for such a refinery is modeled as a multiperiod linear program defined over discrete time periods $t \in T$. Material flows are represented by the $F^j_{k,t}$, where j denotes the technology or subsystem, and k shows the specific inlet or outlet stream.

The emissions from other parts of the refinery, such as the fluid catalytic cracking unit, furnaces, and gas-based units such as the NG boiler and steam methane reformer (SMR), can also be treated by the post-combustion capture unit. All equipment is assumed to operate flexibly within specified operational constraints.

Power outputs are denoted by P^o_t for an electrically driven device *o*. Storage levels are represented by L^s_t for each storage technology $s \in S$. These are the decision variables. Installed capacities for units and storage are given by Q_j, with minimum-load fractions κ_j. Conversion efficiencies are denoted by γ^j_k, and ramp-up/ramp-down limits by r^U_j and r^D_j. System demands, such as hydrogen demand $d_{H_2,t}$ and heat/power steam demand $d_{HPS,t}$, must always be met. External energy prices include the natural gas price $\$_{NG}$ and time-varying electricity price $\$_{elec.,t}$. A system-wide CO_2 emissions cap is enforced via a $CO2_cap$.

The optimization problem, which minimizes the discounted daily operating cost, is shown in Eq. (1) for steam generation, H_2 production, storage, and carbon capture. The hourly operating cost includes the natural gas cost, electricity cost, CO_2-capture solvent-related operating costs, and other operating expenses related to the storage systems shown in Eq. (2). The sets used in this formulation are defined in the nomenclature section at the end of this paper.

$$Min: \sum_{t \in T} \frac{OPEX_t}{(1+int)^{yr}} \tag{1}$$

$$OPEX_t = \left(\sum_{j \in Out(NG)} F^j_{NG,t}\ \$_{NG} + \sum_{j \in Out(grid)} P^j_t \times \$_{elec.,t} \right) + \sum_{j \in CC_{units}} \left(\frac{\$_{sol.,j}}{unit_{CO_2 capt.}} \right) F^j_{CO_2\ capt.,t} + \sum_{s \in S, j \in In(s)} \left(F^s_{H_2,t} \times \$OPEX_s \right) \tag{2}$$

Each unit must operate within its maximum installed capacity and above its minimum turn-down, described as a fraction ($\kappa \leq 1$) of the installed capacity, as denoted in Eq. (3) and (4). For the H_2 storage unit, the state of charge at any given time is lower than the total installed storage capacity, as shown in Eq. (5).

$$\kappa_o Q_o \leq P^o_t \leq Q_o\ \forall t \in T, \forall o \in O \tag{3}$$

$$\kappa_m Q_m \leq F^m_{k'_m,t} \leq Q_m\ \forall t \in T, \forall m \in M \tag{4}$$

$$L^s_t \leq Q_s\ \forall t \in T, \forall s \in S \tag{5}$$

Mass balances relating input and output component flows with respect to a key component is denoted by Eq. (6). The energy balances connecting power flows in electrically driven units to the corresponding mass flows of the key component are detailed in Eq. (7). Mass-balance on the H_2 storage unit is described in Equation (8) where the state of charge in consecutive time periods is obtained by the net H_2 inflow during a particular time. Equation (9) is a mass balance on the resource nodes.

$$F^m_{k,t} = \gamma^m_k F^m_{k'_m,t}\ \forall t \in T, \forall m \in M \tag{6}$$

$$\gamma^o F^o_{k'_o,t} = P^o_t \forall t \in T, \forall o \in O \tag{7}$$

$$F^s_{l,t} - F^s_{k,t} = L^s_{t+1} - L^s_t\ \forall t \in T, \forall s \in S \tag{8}$$

$$\sum_{m \in In(n)} F^m_{n,t} = \sum_{p \in Out(n)} F^p_{n,t}\ \forall t \in T, \forall n \in N \tag{9}$$

Equation (10) specifies hourly ramp-up/down limits of the equipment, such as electrolyzers and e-boilers, restricting the feasible space for flexible operations. Operational constraints on the H_2 storage nodes limit the maximum charge and discharge rates, as shown in Eqs. (11) and (12).

$$r^D_j \leq |F^j_{k'_j,t+1} - F^j_{k'_j,t}| \leq r^U_j\ \forall t \in T, \forall j \in J \tag{10}$$

$$F^s_{l_s,t} \leq \frac{Q_s}{\tau_{s,c}}\ \forall t \in T, l_s \in In(s), \forall s \in S \tag{11}$$

$$F^s_{k_s,t} \leq \frac{Q_s}{\tau_{s,d}}\ \forall t \in T, k_s \in Out(s), \forall s \in S \tag{12}$$

Hydrogen and steam demand from the refinery must always be met as shown in Eqs. (13) and (14). CO_2 emissions from low concentration flue gases released by other parts of the refinery, such as the FCC unit and furnaces that do not have any e-based alternatives, are fixed using Eq. (15).

$$\sum_{q \in Q} F^q_{H_2,t} \geq d_{H_2,t}\ \forall t \in T \tag{13}$$

$$\sum_{h \in H} F^h_{HPS,t} \geq d_{HPS,t}\ \forall t \in T \tag{14}$$

$$F^{plant}_{flue_gas,t} = F_o \forall t \in T \tag{15}$$

The emissions over the time horizon considered are limited by the $CO2_cap$ (Eq. (16)).

$$\sum_{q \in Q} \sum_{t \in T} F^q_{CO_2 emit.} \leq CO_2_cap \tag{16}$$

DCOPF

This is a linear approximation of the AC Optimal Power Flow problem. Here, the decision variable P_{G_i} represents the power generated by the generator i while the voltage angle at the bus i is given by θ_i. Parameters P_{D_i} denotes the real power demand at the bus i. The coefficients C_i' and C_i are the generator cost parameters in the quadratic cost function. The term B_{ij} refers to the line susceptance between buses i and j which is a line property measuring the ease of power flow, and N_i indicates the set of neighboring buses connected to the bus i. Parameters $P_{G_i}^{\min}$ and $P_{G_i}^{\max}$ specify the minimum and maximum real power generation limits for the generator i, respectively and $P_{ij}^{\max}$ shows the maximum magnitude of power flow through the line connecting buses i and j. Eqs. (17)-(20) show the associated quadratic program for the DCOPF problem. Eq. (17) shows the objective function, minimizing the total operating cost of all generators. Each generator's fuel cost is modeled as a quadratic function of its output. Eq. (18) states that the power balance must be satisfied at every bus i. According to the DC power flow model, real power flow through a transmission line depends linearly on the line susceptance and on the difference in voltage angles between the two buses it connects. Eq. (19) constraints that each generator must operate within its physical limits, whereas Eq. (20) restricts that the power flows through the lines are within allowed limits.

$$\min_{\{P_G,\theta\}} \varphi = \sum_{i\in G}\left(C_i' P_{G_i}^2 + C_i P_{G_i}\right) \tag{17}$$

$$P_{G_i} - P_{D_i} = \sum_{j\in N_i} B_{ij}\left(\theta_i - \theta_j\right) \forall i \tag{18}$$

$$P_{G_i}^{min} \le P_{G_i} \le P_{G_i}^{\max} \forall i \in G \tag{19}$$

$$-P_{ij}^{\max} \le B_{ij}(\theta_i - \theta_j) \le P_{ij}^{\max} \forall (i,j) \in L \tag{20}$$

The LMP is the electricity price at a specific grid node, representing the cost of delivering an additional unit of power, considering generation costs and transmission constraints as shown in Eq. (21). In practice, this is obtained from the Lagrange multipliers of the corresponding nodal power balance (Eq. (18)).

$$LMP_i = \lambda_i = \frac{\partial \varphi}{\partial P_{D_i}} \forall i \in G \tag{21}$$

Chance Constrained DCOPF (CC-DCOPF)

In contrast to the standard deterministic DC-OPF model, this model treats the load as the sum of nominal and uncertain components. Real-time reserve deployment is characterized by generators adjusting their output proportionally to the total demand error, governed by their participation factors, to maintain system balance. This uncertainty is managed using chance constraints, a method widely adopted in the power systems community. Unlike robust optimization, which conservatively targets the absolute worst-case scenario, chance constraints offer a less risk-averse framework by strictly enforcing reliability only up to a specified probability level [20-22]. The chance-constrained DCOPF formulation (shown through Eqs. (22)-(27) accounts for uncertainty in demand for buses. The objective function is augmented to include cost terms not only for scheduled generation but also for reserve deployment, represented by the additional decision variables, $\hat{P}_{G_i}$. While the power balance equations and transmission limits remain consistent with the traditional DCOPF, as shown in Eq. (23) and (25), the generation limits now bound both the base and reserve generation capacities, as shown in Eq. (24). The enhanced formulation introduces a risk-based reserve constraint that requires each generator's reserve contribution to satisfy Eq. (26), where the decision variable α_{G_i} denotes the generator i's participation factor, while the parameters σ and $z_{1-\varepsilon}$ capture the magnitude of uncertainty in demand and the risk-taking capacity, respectively [22].

$$\min \sum_{i\in G}\left(C_i' P_{G_i}^2 + C_i P_{G_i} + C_i' \hat{P}_{G_i}^2 + C_i \hat{P}_{G_i}\right) \tag{22}$$

$$P_{G_i} - P_{D_i} = \sum_{j\in N_i} B_{ij}\left(\theta_i - \theta_j\right) \forall i \tag{23}$$

$$P_{G_i}^{min} \le P_{G_i} + \hat{P}_{G_i} \le P_{G_i}^{\max} \forall i \in G \tag{24}$$

$$-P_{ij}^{\max} \le B_{ij}(\theta_i - \theta_j) \le P_{ij}^{\max} \forall (i,j) \in L \tag{25}$$

$$\hat{P}_{G_i} \ge z_{1-\varepsilon}\alpha_{G_i}\sigma \forall i \in G \tag{26}$$

$$\sum_{i\in G} \alpha_i = 1 \tag{27}$$

Eq. (26) is derived analytically by integration based on the relation shown in Eq. (28), stating that the probability of the reserve generation capacity being higher than the excess demand is greater than 1- ε. Assuming that the excess demand $\hat{P}_{D_i}$ follows a normal distribution with mean 0 and standard deviation σ_i, as shown in Eq. (29), the chance constraint is converted into an equivalent deterministic form expressed in terms of the standard normal quantile $\Phi^{-1}(1-\varepsilon)$ which is the same as $z_{1-\varepsilon}$. σ_i's are related to σ as described in Eq. (30), a normalization condition is introduced in Eq. (27) to ensure complete allocation of reserve responsibilities (participation factor) across the generating fleet. These additions enable the model to incorporate uncertainty-aware scheduling while preserving the structural simplicity of the conventional DCOPF framework.

$$\mathbb{P}\left(\left(\alpha_{G_i} \sum_i \hat{P}_{D_i}\right) - \hat{P}_{G_i} \le 0 \right) \ge 1 - \varepsilon \forall i \in G \tag{28}$$

$$=> \Phi^{-1}(1-\varepsilon)\alpha_{G_i}\sigma - \hat{P}_{G_i} \le 0 \ \forall i \in G \tag{29}$$

$$\hat{P}_{D_i} \sim N(0,\sigma_i) \qquad \sigma = \left(\sqrt{\sum_i \sigma_i^2}\right) \tag{30}$$

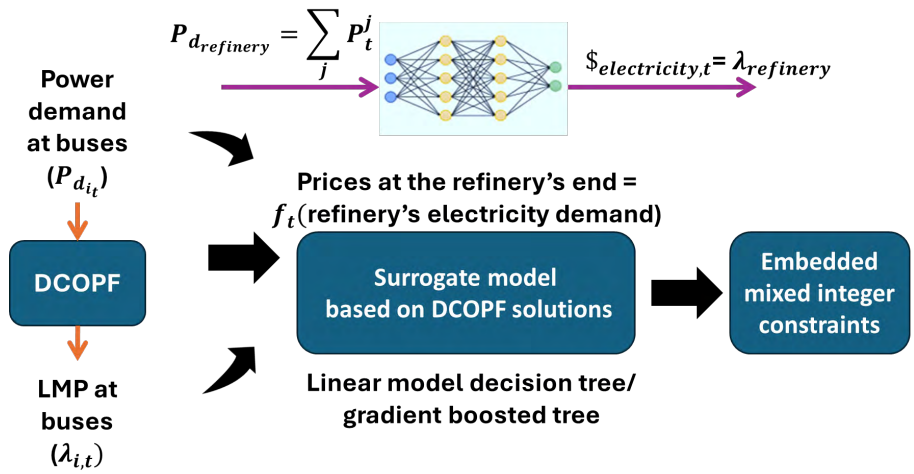

Figure 3. Surrogate model training and embedding

Training the surrogate model

In our framework, we vary the refinery's active-power demand P_{d_i} over a specified range while assuming that the demands at all other nodes are fixed and independent of the refinery's demand. For each time interval in which the non-refinery loads (e.g., residential demand) can be treated as fixed, we solve a series of DCOPF problems to generate training data and fit a separate surrogate model. Since the DCOPF is a quadratic program with linear constraints, its KKT conditions for a given active set constitute a system of linear equations in both the primal and dual variables. This implies that the nodal price is an affine function of the demand within a specific regime. However, as the active constraint set changes, the coefficients of this relationship shift, resulting in a global piecewise linear function. This structure motivates the use of linear-model decision trees and related surrogate models, such as gradient-boosted trees, to capture the varying linear regimes. In this work, we train such models on DCOPF-generated data and embed the resulting relationship into the refinery's operational optimization problem. This could be expressed as the disjunction shown in Eq. (31).

$$\vee \begin{bmatrix} Z_w \\ \underline{P_{w,t}} \le \sum_{j\in Out(Grid)} P_t^j \le \overline{P_{w,t}} \\ \$_{elec,t} = a_{w,t} \sum_{j\in Out(Grid)} P_t^j + b_{w,t} \end{bmatrix} \quad (31)$$

Each disjunct is indexed by w, and in every disjunct the refinery's total power demand, $\sum_{j\in \mathrm{Out(Grid)}} P_t^j$ is restricted to lie between the lower and upper bounds $\underline{P_{w,t}}$ and $\overline{P_{w,t}}$. Within each disjunct, the electricity price $\$_{\mathrm{elec},t}$ is modeled as a linear function of the total power consumption $\sum_{j\in \mathrm{Out(grid)}} P_t^j$, characterized by a slope $a_{w,t}$ and an intercept $b_{w,t}$.

Both the slope and intercept vary over time, as the applicable surrogate model is time-dependent. If we explicitly include equality constraints expressing the total electricity cost, $\$_{elec,t}$, as a function of demand, the objective function (2) still contains non-convex bilinear terms. These terms arise from expressing the total cost as a product of price and power consumption, which significantly increases the computational burden. While off-the-shelf tools like OMLT can automate the embedding of such relationships, a custom reformulation that substitutes the cost function directly into the objective allows for better convergence by avoiding bilinearities. We perform a convex hull reformulation, before substituting the cost term in the objective, Eq. (2). At a given time t, exactly one binary variable y_t^w is activated depending on the range in which the total power consumption lies. Depending on this, exactly one auxiliary variable P_t^w and its corresponding quadratic term, P_t^{w2} are nonzero, as shown in Eqs. (34) and (35). The sum of P_t^w over all disjuncts equates to the total consumption across all refinery units. The resulting model obtained by introducing Eqs. (33)-(35) and substituting the term in the RHS of Eq. (32) in the objective function, expressed IN Eq. (2) appropriately, is a Mixed Integer Quadratic Program (MIQP) with a convex relaxation.

$$\$_{elec.,t}\left(\sum_{j\in Out(grid)} P_t^j\right) = \sum_w (b_{w,t}P_t^w + a_{w,t}P_t^{w2})\ \forall t \in \mathrm{T} \quad (32)$$

$$\left(\sum_{j\in Out(grid)} P_t^j\right) = \sum_w P_t^w\ \forall t \in \mathrm{T} \quad (33)$$

$$\underline{P_{w,t}} y_t^w \le P_t^w \le \overline{P_{w,t}} y_t^w\ \forall t \in \mathrm{T} \quad (34)$$

$$\sum_w y_t^w = 1 \quad (35)$$

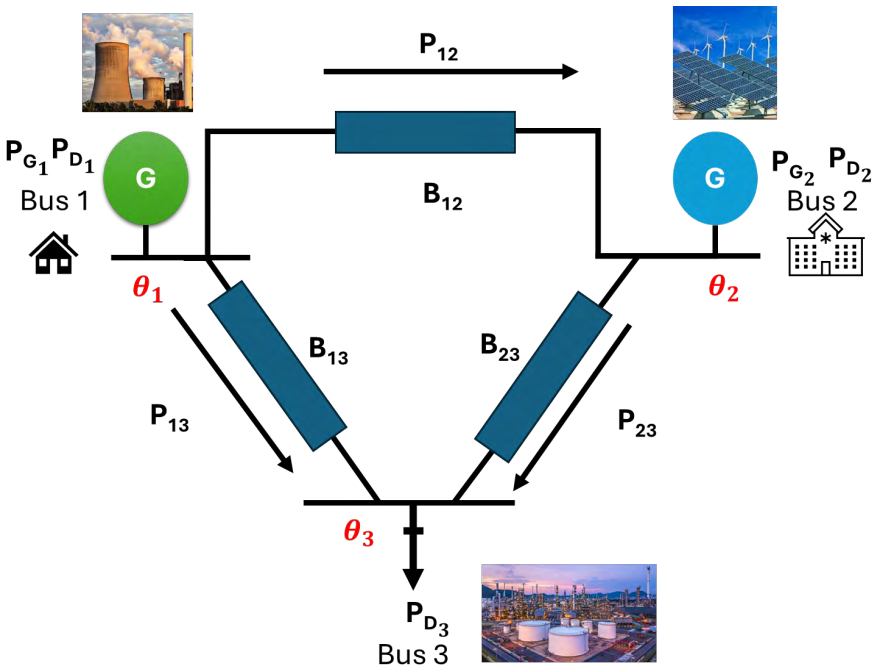

Figure 4. Considered Power grid network

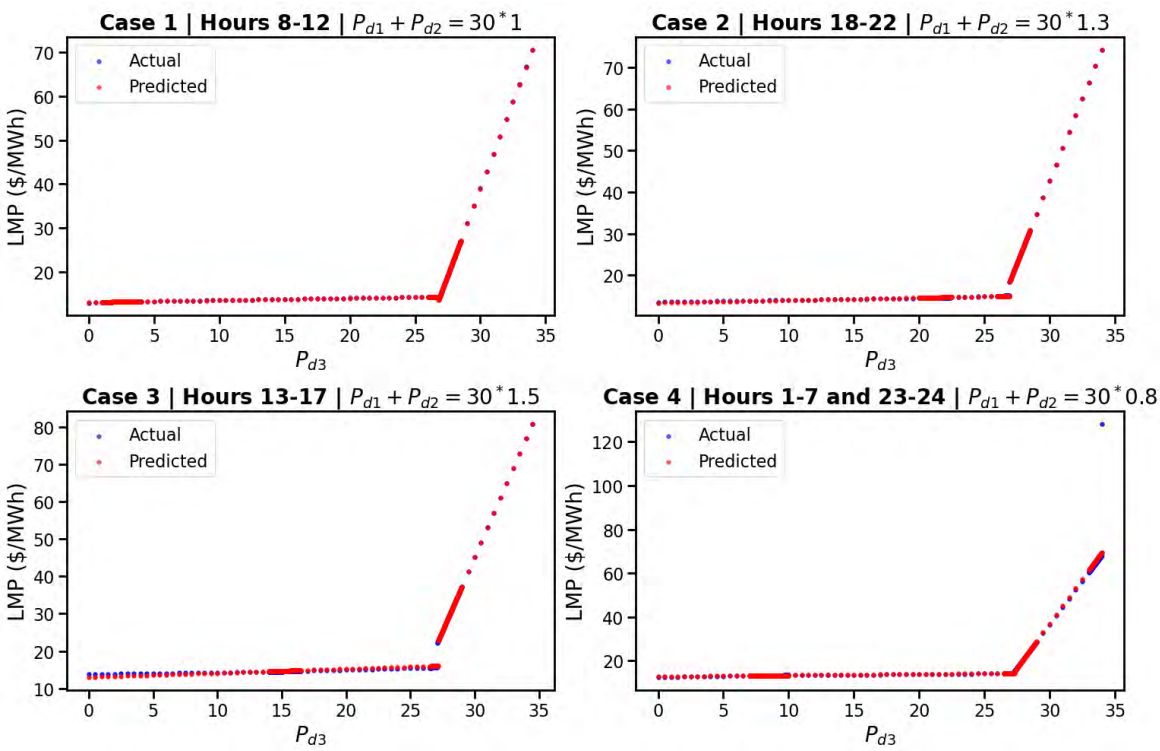

Figure 5. Fit for each time segment when demand at the other nodes are fixed using linear model decision trees Chance-constrained DCOPF embedding

RESULTS

We illustrate the proposed approach on a 3-bus system in which the refinery is at Bus 3. Fig. 4 presents the fits obtained using linear model decision tree surrogates for each time segment, with the

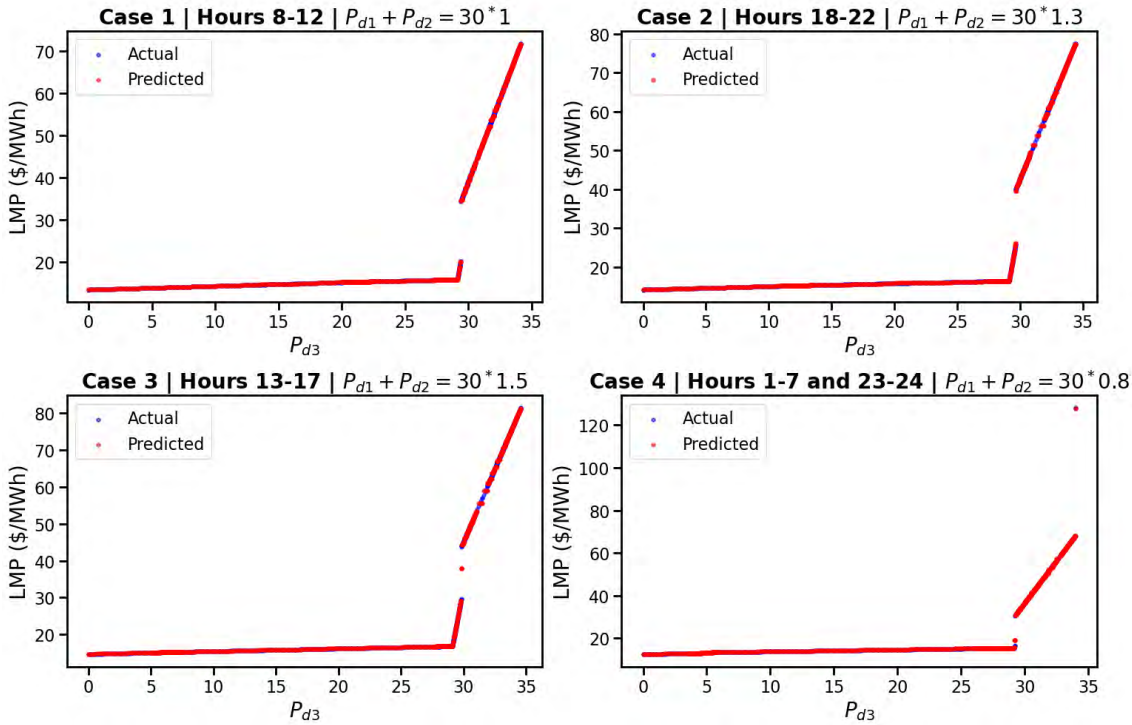

Figure 6. Fit obtained for each time segment using GBTs

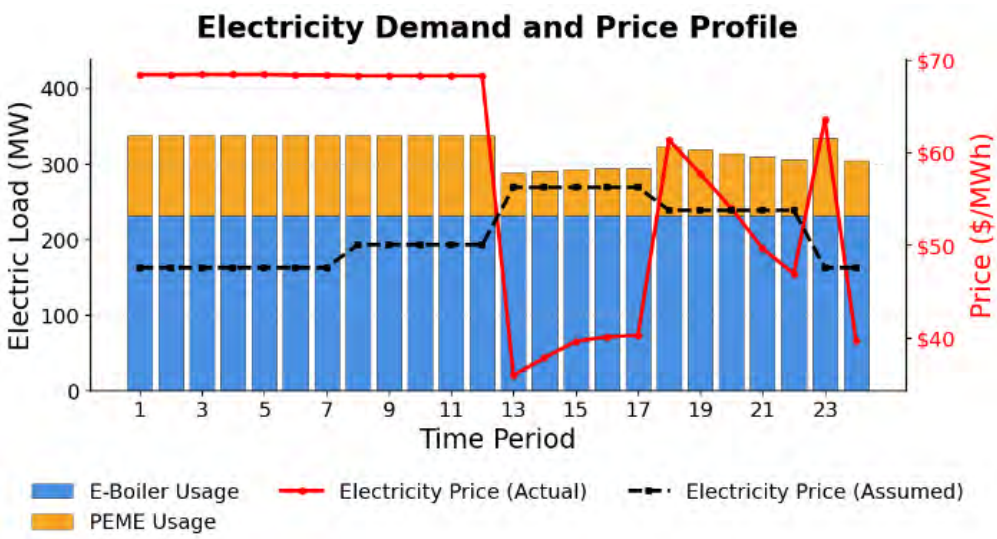

Figure 7. Base case electricity usage and prices

corresponding loads at the other buses annotated in the figure (shown in scaled units of 10 MW). Linear model decision trees work by recursively partitioning the data and then fitting a local linear model to each part, so the overall model captures piecewise-linear relationships. To ensure the accuracy of these captured transitions, the surrogate model was trained by sampling more densely around the sharp corners; detailed data regarding both the underlying refinery model and the surrogate fit accuracy shall be published in a forthcoming full-length paper. Since the slopes of all linear functions are positive, the relaxation of the previous formulation is convex. In this setting, we also observe in Fig 5, that the overall piecewise functions can also be represented as a maximum over the constituent linear functions, as shown in Eq. (36). Now, this term in the LHS of Eq. (37) can be substituted by an additional variable c_t in the objective of the refinery optimization problem, as shown in Eq. (2), and the max function can be represented by a series of W inequality functions as shown in Eq. (38). This removes

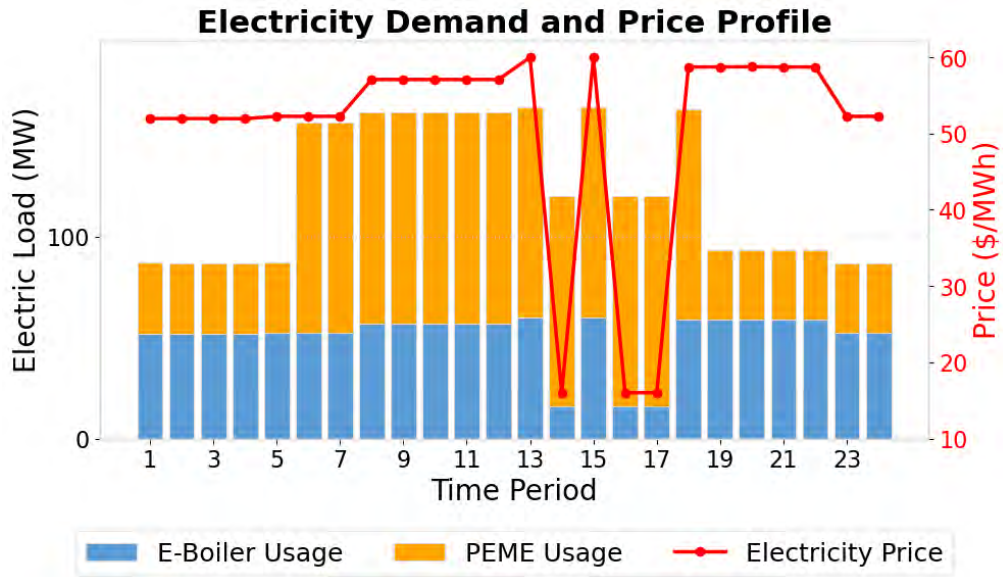

Figure 8. Load profile and prices using MIP formulation ranging from 5.9% to 6.7%. (using an embedding built from data generated from DCOPF solutions)

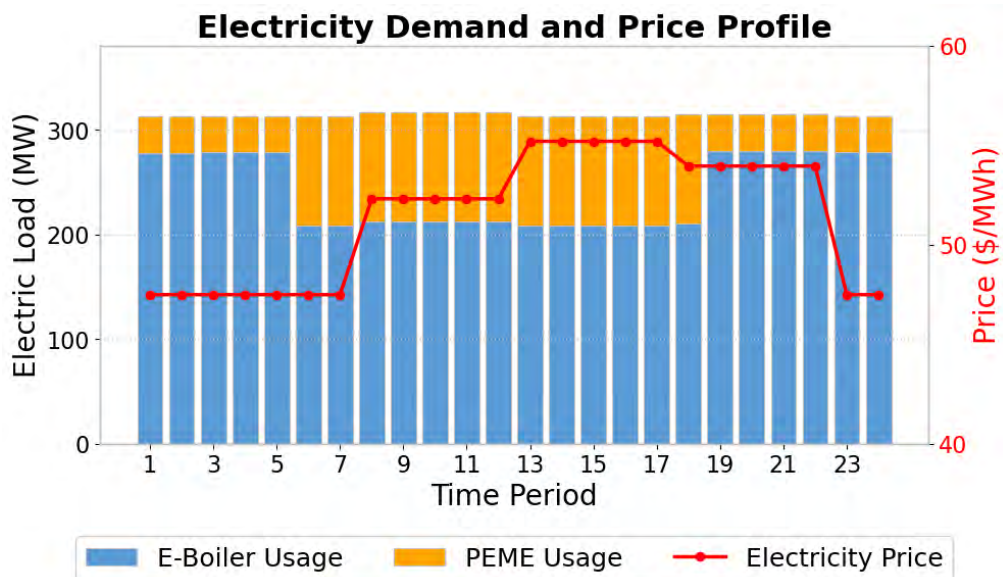

Figure 9. Load profile and prices using NLP formulation (using an embedding built from data generated from DCOPF solutions)

the need for additional binary variables and allows the NLP problem to be handled directly by any standard nonlinear solver.

$$\$_{elec.,t}\left(\sum_{j\in Out(Grid)} P_t^j\right) = \max_w a_{w,t}P_t^{w2} + (b_{w,t}P_t^w))\ \forall \mathrm{t} \in \mathrm{T} \quad (36)$$

$$\$_{elec.,t}\left(\sum_{j\in Out(Grid)} P_t^j\right) = c_\mathrm{t}\ \forall t \in T \quad (37)$$

$$c_\mathrm{t} \geq a_{w,t}P_t^{w2} + (b_{w,t}P_t^w)\ \forall w\ \forall \mathrm{t} \in \mathrm{T} \quad (38)$$

In contrast, when using CC DCOPF, expressed using solution Eq (22)-(27), to generate data for surrogate building, the linear model decision trees could not provide satisfactory fits, especially around the sharp corners, irrespective of the denser sampling. On the other hand, gradient-boosted trees provided a better fit. The obtained fit is shown in Fig 6. By varying the hyperparameters, Gradient-boosted trees (GBTs) are used to obtain satisfactory fits on the data. These generate predictions by combining many small decision trees sequentially, where each new tree focuses on correcting the errors of the previous ones. We used OMLT 1.2.2 [17] to generate a mixed Integer formulation.

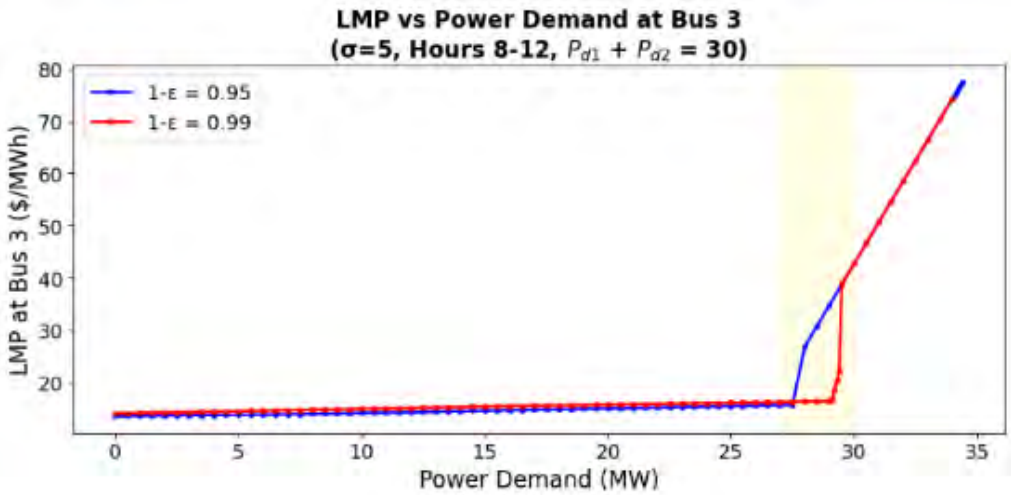

Figure 10. Chance constraint generated DCOPF data

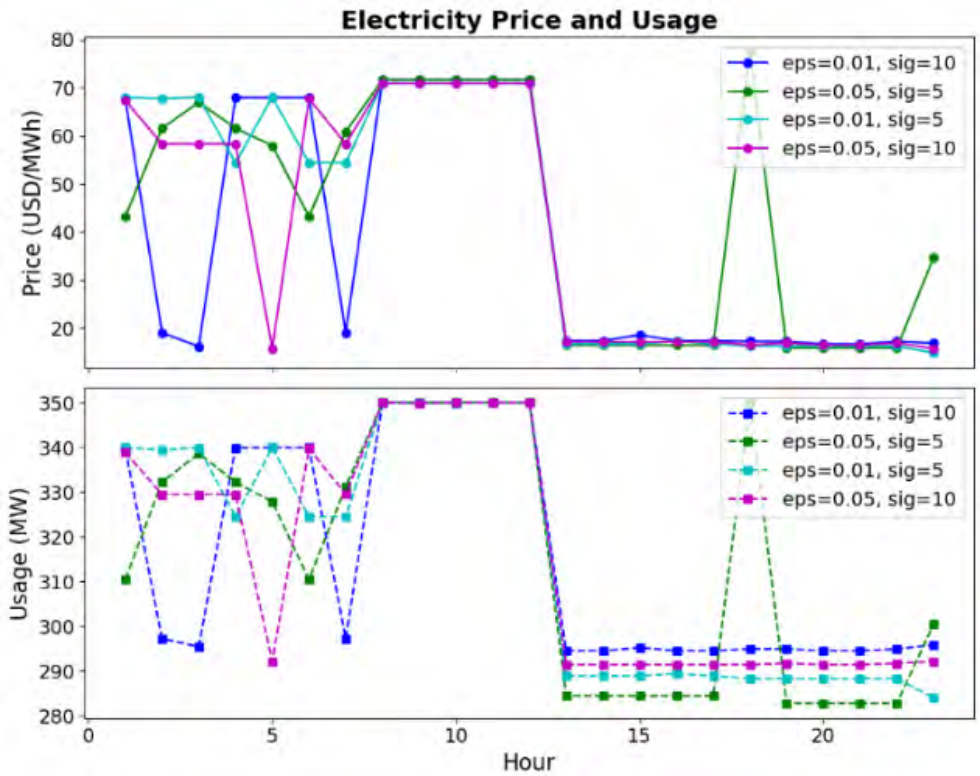

Figure 11. Optimal electricity price and Usage with varying hyperparameters within the chance-constrained model

DCOPF embedding

In the base case, when the price-taker assumption is in place as shown in Fig 7, the model attempts to leverage forecasted periods of low price (assumed price profile shown using the black dotted line) by ramping up production of hydrogen using electrolyzers. However, such actions raise the actual electricity price, which is shown using the solid red line. In this case, with the assumed costs, the optimizer predicted an objective of 0.767 MUSD. But when the demands from the refinery were added to the OPF model, the actual price profiles resulted in a cost of 0.824 MUSD being incurred, a 7.5% increase over the predicted cost. When using the 2 different formulations (MIQP and NLP) consisting of the embedded surrogate model, the objective function value in both these cases was 0.767 MUSD, about 7% lower than the 0.824 MUSD incurred cost in the base case.

The MIQP formulation consists of 1, 082 variables, of which 48 are binaries, and 989 constraints, out of

which 705 are equalities, and was solved using DICOPT 2.0 in 2.2 seconds. The NLP formulation consists of 1, 082 continuous variables and 989 constraints, of which 627 are equalities, and was solved using IPOPT 3.14.17 in 1.5 seconds. However, these two cases produce different load and price profiles that are degenerate solutions of the overall optimization problem (Fig 8 and 9). In particular, the NLP formulation yields a more stable load profile compared to the base case. The degeneracy arises because the model can exploit price variability through the use of hydrogen storage equipment and incorporate a hybrid approach using Proton Exchange Membrane Electrolyzer and Steam Methane Reformer.

A limitation of the previously developed framework is the assumption of perfectly known loads at other nodes and exact grid parameters. Accurately predicting loads at other buses is generally challenging. The resulting formulation obtained by embedding gradient boosted trees trained using optimal solutions to the CC DCOPF is a nonconvex mixed-integer quadratically constrained program (MIQCP) with 28, 360 constraints and 14, 046 variables, of which 3, 508 are binary (and integer). The problem is solved using the Gurobi Optimizer (version 12.0.1) with a time limit of 1, 800 seconds. We illustrate the impact of embedding a chance-constrained DCOPF within the refinery planning model, considering demand uncertainty, under varying risk tolerance (ϵ) and uncertainty levels (σ). Varying the σ and ϵ values as shown in Table 1, the solution converged with gaps

Lower values of ϵ and higher σ reflect reduced risk-taking capability and higher uncertainty in predictions of load at other nodes, respectively. As shown in Fig 10, while locational marginal prices within each linear segment are slightly higher with reduced ϵ, increased uncertainty causes the upper limit in constraint (24) to become binding at lower demand levels, moving the point of non-smoothness in the LMP–demand curve toward higher power values, as illustrated by the shaded region in Fig. 10. Since the optimal operating region includes the region consisting of this sharp corner, as shown in Fig. 10, higher σ values and lower ϵ values lead to lower overall objective function values due to the lower prices encountered in this region at similar loads (shown in Table 1).

Table 1. Objective value with changing hyperparameters

σ	1-ϵ	Objective Function (MUSD)
5	0.95	0.7237
10	0.95	0.6892
5	0.99	0.7046
10	0.99	0.6740

CONCLUSIONS

This work presents a framework for incorporating electricity market effects into the operational planning of an electrified oil refinery. These effects are captured by embedding constraints derived from relationships modeled with surrogate models, trained on DCOPF-based data. The surrogate models reflect the impact of refinery load on nodal electricity prices. The piecewise-linear relationship between refinery power demand and prices is approximated with linear model decision trees or gradient-boosted trees. In the illustrative example, the surrogates are embedded within the refinery optimization problem using reformulations of the resulting GDP, namely a MIQCP and an NLP reformulation, or off-the-shelf tools such as OMLT. These demonstrated improved objective values compared to a base case that ignores price feedback from the market, while providing optimal load profiles through coordinated hydrogen storage and hybrid electrolyzer and e-boiler operation. We extend this framework to incorporate uncertainty in load forecasts at other buses by embedding constraints that represent surrogate models trained using optimal solutions obtained from chance-constrained DCOPF. The results show that lower risk-taking capacity and higher uncertainty at the grid operator's end can counterintuitively reduce overall costs at the electrified refinery's end. Despite promising cost-reduction potential, two primary limitations remain for the illustrated framework. The use of a 3-bus system necessitates future testing on larger-scale networks to ensure scalability. Additionally, advancing from DCOPF to ACOPF-based surrogates is essential to capture the reactive power and voltage dynamics necessary for full physical fidelity.

ACKNOWLEDGEMENTS

This project was supported by the LDRD program project "Optimization and Physics Inspired Machine Learning Approaches" at Center for Nonlinear Studies in Los Alamos National Laboratory. Research conducted at Los Alamos National Laboratory is done under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy under Contract No. 89233218CNA000001. Report No. LA-UR-25-32297.

NOMENCLATURE

J = {PEME, SMR, NG boiler, e-boiler, Postcomb. CC}
M = {NG boiler, SMR, Precomb CC, Postcomb CC}
o = {e-boiler, PEME}
N = {NG, Hydrogen (var.), Rest of plant, Low conc. flue gas, High conc. flue gas, CO_2 emit., CO_2 capt.}
Q = {PEME, SMR, Hydrogen storage}
H = {NG boiler, e-boiler, Postcomb CC}
S = {Hydrogen storage}
T = {1, 2, ... 24}
G = Set of Buses consisting of generators

REFERENCES

1. Chattopadhyay S, Karthikeyan K, Gandhi R, Grossmann IE, Torres AI. Optimal retrofit of carbon capture and electrification technologies in oil refineries for reducing direct co₂ emissions. Ind. Eng. Chem. Res. 64:10131-10147 (2025). https://doi.org/10.1021/acs.iecr.4c04995
2. Corengia M, Torres AI. Coupling time varying power sources to production of green-hydrogen: a superstructure based approach for technology selection and optimal design. Chemical Engineering Research and Design 183:235-249 (2022). https://doi.org/10.1016/j.cherd.2022.05.007
3. Giannikopoulos I, Skouteris A, Allen DT, Baldea M, Stadtherr MA. Thermal electrification of chemical processes using renewable energy: economic and decarbonization impacts. Ind. Eng. Chem. Res. 63:12064-12082 (2024). https://doi.org/10.1021/acs.iecr.4c00737
4. Demirhan CD, Tso WW, Powell JB, Pistikopoulos EN. A multi-scale energy systems engineering approach towards integrated multi-product network optimization. Applied Energy 281:116020 (2021). https://doi.org/10.1016/j.apenergy.2020.116020
5. Jabari F, Mohammadi-ivatloo B, Mohammadpourfard M. Robust optimal self-scheduling of potable water and power producers under uncertain electricity prices. Applied Thermal Engineering 162:114258 (2019). https://doi.org/10.1016/j.applthermaleng.2019.114258
6. Attarha A, Amjady N, Dehghan S, Vatani B. Adaptive robust self-scheduling for a wind producer with compressed air energy storage. IEEE Trans. Sustain. Energy 9:1659-1671 (2018). https://doi.org/10.1109/tste.2018.2806444
7. Kazempour J, Hobbs BF. Value of flexible resources, virtual bidding, and self-scheduling in two-settlement electricity markets with wind generation—part I: principles and competitive model. IEEE Trans. Power Syst. 33:749-759 (2018). https://doi.org/10.1109/tpwrs.2017.2699687
8. Martinek J, Jorgenson J, Mehos M, Denholm P. A comparison of price-taker and production cost models for determining system value, revenue, and scheduling of concentrating solar power plants. Applied Energy 231:854-865 (2018). https://doi.org/10.1016/j.apenergy.2018.09.136
9. Orvis R, Aggarwal S. Refining competitive electricity market rules to unlock flexibility. The Electricity Journal 31:31-37 (2018). https://doi.org/10.1016/j.tej.2018.05.012
10. Gao X, Knueven B, Siirola JD, Miller DC, Dowling AW. Multiscale simulation of integrated energy system and electricity market interactions. Applied Energy 316:119017 (2022). https://doi.org/10.1016/j.apenergy.2022.119017
11. Jalving J, Ghouse J, Cortes N, Gao X, Knueven B, Agi D, Martin S, Chen X, Guittet D, Tumbalam-Gooty R, Bianchi L, Beattie K, Gunter D, Siirola JD, Miller DC, Dowling AW. Beyond price taker: conceptual design and optimization of integrated energy systems using machine learning market surrogates. Applied Energy 351:121767 (2023). https://doi.org/10.1016/j.apenergy.2023.121767
12. Shafiee S, Zamani-Dehkordi P, Zareipour H, Knight AM. Economic assessment of a price-maker energy storage facility in the alberta electricity market. Energy 111:537-547 (2016). https://doi.org/10.1016/j.energy.2016.05.086
13. Arteaga J, Zareipour H. A price-maker/price-taker model for the operation of battery storage systems in electricity markets. IEEE Trans. Smart Grid 10:6912-6920 (2019). https://doi.org/10.1109/tsg.2019.2913818
14. Nasrolahpour E, Kazempour SJ, Zareipour H, Rosehart WD. Strategic sizing of energy storage facilities in electricity markets. IEEE Trans. Sustain. Energy 7:1462-1472 (2016). https://doi.org/10.1109/tste.2016.2555289
15. Hobbs BF, Metzler CB, Pang JS. Strategic gaming analysis for electric power systems: an MPEC approach. IEEE Trans. Power Syst. 15:638-645 (2000). https://doi.org/10.1109/59.867153
16. Tang X, Petra CG, Baldea M, Baldick R. A grid-scale study of demand bidding by large industrial users. Computers & Chemical Engineering 205:109442 (2026). https://doi.org/10.1016/j.compchemeng.2025.109442
17. Ammari BL, Johnson ES, Stinchfield G, Kim T, Bynum M, Hart WE, Pulsipher J, Laird CD. Linear model decision trees as surrogates in optimization of engineering applications. Computers & Chemical Engineering 178:108347 (2023). https://doi.org/10.1016/j.compchemeng.2023.108347
18. Kondili E, Pantelides CC, Sargent RWH. A general algorithm for short-term scheduling of batch operations—i. MILP formulation. Computers & Chemical Engineering 17:211-227 (1993). https://doi.org/10.1016/0098-1354(93)80015-f
19. Pantelides CC. Unified frameworks for the optimal process planning and scheduling. In: Proceedings of the Second Conference on Foundations of Computer Aided Operations. 253–274 (1994).
20. Summers T, Warrington J, Morari M, Lygeros J. Stochastic optimal power flow based on convex

approximations of chance constraints. 2014 Power Systems Computation Conference :1-7 (2014). https://doi.org/10.1109/pscc.2014.7038376
21. Pena-Ordieres A, Molzahn DK, Roald LA, Wachter A. DC optimal power flow with joint chance constraints. IEEE Trans. Power Syst. 36:147-158 (2021). https://doi.org/10.1109/tpwrs.2020.3004023
22. Roald L, Oldewurtel F, Krause T, Andersson G. Analytical reformulation of security constrained optimal power flow with probabilistic constraints. 2013 IEEE Grenoble Conference :1-6 (2013). https://doi.org/10.1109/ptc.2013.6652224
23. PJM Interconnection. (n.d.). Data Viewer: Real-Time & Day-Ahead LMP., from https://dataviewer.pjm.com/dataviewer/pages/public/lmp.jsf
24. Mallapragada DS, Dvorkin Y, Modestino MA, Esposito DV, Smith WA, Hodge BM, Harold MP, Donnelly VM, Nuz A, Bloomquist C, Baker K, Grabow LC, Yan Y, Rajput NN, Hartman RL, Biddinger EJ, Aydil ES, Taylor AD. Decarbonization of the chemical industry through electrification: barriers and opportunities. Joule 7:23-41 (2023). https://doi.org/10.1016/j.joule.2022.12.008

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Techno-economic Analysis of Alternatives for Carbon Capture and Utilization and Green Ammonia Production from a Cement Plant Flue Gas

Miguel A. Pedro[a], Ana S. Amorim[b*], and Henrique A. Matos[c]

[a] Departamento de Engenharia Química, Instituto Superior Técnico, Universidade de Lisboa, 1049-001, Lisboa, Portugal
[b] C5Lab – Sustainable Construction Materials Association, 2795-242 Linda-a-Velha, Portugal
[c] CERENA, Departamento de Engenharia Química, Instituto Superior Técnico, Av. Rovisco Pais 1, 1049-001 Lisboa, Portugal
* Corresponding Author: aamorim@c5lab.pt

ABSTRACT

The manufacturing industry is the second largest emitter of CO_2, with the cement industry being one of the main contributors (7-8 % of the global emissions). Carbon capture and utilization (CCU) technologies are promising decarbonization solutions for the cement industry, addressing both fossil fuel-related (40 %) and process-derived emissions (60 %). Within a cement plant, producing synthetic natural gas (SNG) from captured CO_2 is particularly suitable, as it is sufficient to fully replace solid fuels in the rotary kiln. On the other hand, the use of zero-carbon fuels, such as green ammonia, is also recognized as a promising approach for decarbonization. In this work, a superstructure was developed to explore alternative routes for producing SNG and green ammonia from CO_2 and N_2 in cement plant flue gas, respectively. The routes were modelled in Aspen Plus® V14, and their economic viability was assessed. Currently, the most promising route, at a cost of 109 €/tonne of flue gas, involves CO_2 capture (90% efficiency) and storage, with the N2-rich stream emitted. However, all the routes remain uncompetitive compared with the full emission of the flue gas (22 €/tonne). Routes for green ammonia production from N_2 in flue gas were compared with those from atmospheric air but were found to be uncompetitive. Nevertheless, with future developments, the economic feasibility of integrating pathways for synthesizing high-value chemical products into the cement industry is expected to increase.

Keywords: Cement industry, SNG production, Green ammonia production, Aspen Plus, Techno-economic analysis

INTRODUCTION

Climate change continues to pose a pressing global challenge, driving extensive research and policy initiatives worldwide. The planet's temperature has been exponentially increasing over the past decades, reaching the value of 1.28 °C above the pre-industrial average in 2024 [1]. This translates to a new record high of 37.8 Gt of CO_2 emitted [2], which takes responsibility for global warming, as it is the main greenhouse gas (GHG). Relocating electricity and heat consumption to their respective sectors, the industry sector becomes the most impactful one for GHG emissions, accounting for 39% of global CO_2 emissions. The cement industry alone is responsible for around 7-8% of GHG emissions, as CO2 is a direct by-product of its production process, which explains the high CO_2 content in its flue gas. Carbon capture, utilisation and storage (CCUS) is considered a promising strategy for reducing these emissions. The literature on the treatment of cement plant flue gas focuses primarily on carbon capture and storage technologies [3]. However, studies have reported the possibility of producing high-value chemicals, such as formic acid, dimethyl ether and synthetic natural gas (SNG), despite having obstacles in terms of high energy and raw materials consumption [4]. Among these carbon utilization strategies, the conversion of the captured CO_2 into SNG through the Sabatier process, with the addition of green hydrogen and in the presence of metal catalysts, is particularly suitable for the cement industry, as the

produced SNG is enough to fully replace the solid fuels fed into the rotary kiln and precalciner of a cement plant [5].

Since the cement plant flue gas is composed of not only CO_2, but also N_2 and some minor compounds, the obtention of a high-purity CH_4 stream usually requires pre-treatment of the flue gas to remove NO_x, SO_x, HCl and dust particles and an additional capture unit for N_2 separation from CO_2 [6]. The best available technologies for pre-treatment include wet scrubbing with hydrated lime and selective catalytic reduction (SCR) for SOx and NOx removal, respectively, while the most mature carbon capture technology is absorption with monoethanolamine (MEA) [7-9].

Traditionally, ammonia (NH_3) is synthesized through the Haber-Bosch (H-B) process, which relies on fuel gasification, leading to large CO_2 emissions. Power-to-ammonia technologies offer a sustainable alternative, requiring nitrogen (N_2), usually obtained from an air separation unit (ASU), and green hydrogen from water electrolysis, which are both highly energy-intensive [10]. Because common cement plant flue gas has a high N_2 content, utilizing this component in an ammonia production process offers a promising alternative to bypass the ASU and reduce energy consumption. This way, transforming N_2 into NH_3, a versatile compound with fertilizer and energy applications, highlights the potential of generating valuable products from the cement plant flue gas.

Despite their potential, CCU and green ammonia production face economic and technical challenges, mainly due to the high cost of green hydrogen. Existing pilot plants study these technologies separately, missing potential synergies between the processes. This work introduces a novel approach by using industrial flue gas as an N_2 source, thereby eliminating the need for an ASU.

METHODS

The flue gas of a cement plant was considered to go through a set of processes defined in a superstructure to capture and either store or convert CO_2 into SNG and N_2 into ammonia. These processes were modelled using Aspen Plus® v14 and the techno-economic viability of each route of the superstructure was evaluated. The heat exchangers of the processes considered were designed using Aspen Exchanger Design & Rating (EDR)® [11].

Cement Plant Flue Gas

The typical composition of a cement plant flue gas, which was used as a basis for this work, was retrieved from the work of Schakel et al. [6]. However, only 10 % of the total mass flowrate of the cement plant flue gas was considered, to avoid ammonia market saturation.

Superstructure

The most mature technologies were integrated to assess the techno-economic feasibility of converting a typical cement plant's flue gas into SNG and green ammonia, reaching the superstructure shown in Fig. 1.

Before feeding the production processes, the flue gas must undergo pre-treatment to remove NO_x, SO_x, HCl, and dust particles. After this process, four options were considered: the separation of CO_2 from N_2 with 90 % efficiency ('MEA absorber (90 %)') or with approximately 100 % efficiency ('MEA absorber (~100 %)'), no separation of these components or direct emission of the flue gas. The 'MEA absorber (~100 %)' process was investigated to obtain a CO_2-free outlet, which can be sent directly to the H-B process. In the MEA absorption units, the stream is separated into a CO_2-rich and a N_2-rich stream. The CO_2-rich stream can either be sent to the 'Sabatier' unit to generate SNG or to storage, while the N_2-rich stream can be valorized into green ammonia in the H-B process or emitted. If the high-purity N_2 stream comes from a capturing unit with 90 % efficiency, a methanation step is required beforehand, since the catalysts used for ammonia production are poisoned by oxygen-containing compounds [12]. The 'MEA absorber (~100 %)' option was therefore considered to avoid the need for this additional methanation step. Regarding the option of no CO_2/N_2 separation, the pre-treated flue gas stream is subjected to methanation for converting the CO_2 to methane and sent to the Haber-Bosch process to generate green ammonia. The ammonia is then separated from the methane through distillation. All these routes were compared with the full flue gas emission and with the production of green ammonia through the hydrogenation of atmospheric air. Tab. 1 presents the nomenclature of the different superstructure paths.

Table 1: Nomenclature of the superstructure routes. MEA stands for monoethanolamine, and ASU stands for an air separation unit.

Separation	CO_2 stream	N_2 stream	Route
MEA absorber (90 %)	Sabatier	Haber-Bosch 1	$A90^{S}_{HB}$
		Emit	$A90^{S}_{E}$
	Storage	Haber-Bosch 1	$A90^{CCS}_{HB}$
		Emit	$A90^{CCS}_{E}$
MEA absorber (~100 %)	Sabatier	Haber-Bosch 2	$A100^{S}_{HB}$
		Emit	$A100^{S}_{E}$
	Storage	Haber-Bosch 2	$A100^{CCS}_{HB}$
		Emit	$A100^{CCS}_{E}$
No separation	Sabatier + Haber-Bosch 3		$A0^{S}_{HB}$
	Emit		E
ASU	-	Haber-Bosch 2	ASU_{HB}

The route nomenclature follows the rule: X^{i}_{j}. X represents the separation options of the pre-treated flue gas/air stream, i represents the options of the CO_2-rich stream generated after absorption and j represents the

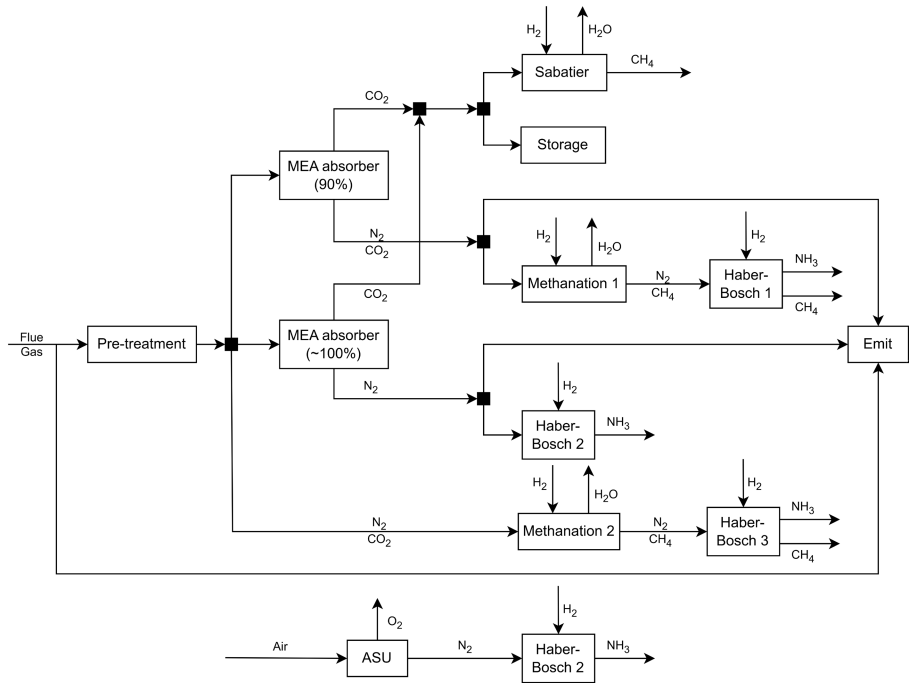

Figure 1. Superstructure for synthetic natual gas and green ammonia production from a cement plant flue gas.

options of the N_2-rich stream generated after absorption.

Pre-treatment

The pre-treatment process aims to remove O_2, NO_x, SO_2, HCl, and dust particles from the initial flue gas. For the removal of SO_2 and HCl, a 10 %wt. $CaCO_3$ aqueous solution was added to a scrubber, which was simulated in Aspen Plus® V14 by using a *RadFrac* model and implementing the equilibrium reactions described by Duque et al. [13]. The dust particles were also assumed to be removed by the scrubber. On the other hand, rate-controlled reactions were implemented to model the NO_x elimination with NH_3 in the SCR using a *RPlug* block, as suggested by Sweeney and Liu [14]. The O_2 elimination was modelled using a *RStoic* model, with its reaction with CH_4 following the work of Lu et al. [15]. The diameter of the scrubber was calculated by Eq. (1)

$$D_{scrubber} = \sqrt{\frac{4S}{\pi}} = \sqrt{\frac{4Q_{v,gas}}{\pi v_{gas}}} \quad (1)$$

, where $Q_{v,\,gas}$ is the gas volumetric flowrate and v_{gas} is the gas velocity inside the scrubber, which was considered 1 m/s [16]. The scrubber's height (L) was attained using Eq. (2) [16].

$$\frac{L}{D} = 3 \quad (2)$$

Regarding the SCR, a tubular flow reactor must be used, which has a usual residence time (τ) of 0.8 s [17]. This, combined with the gas flowrate going through the reactor, enables the calculation of the volume occupied by the gas (V_{gas}). The catalyst volume (V_{cat}), which can often be neglected for being considerably smaller than V_{gas}, can be calculated using the mass of catalyst inside the reactor (m_{cat}) and its density (ρ_{cat}), leading to Eq. (3).

$$V_{reactor} = V_{gas} + V_{cat} = Q_{v,gas} \bullet \tau + m_{cat} \bullet \rho_{cat} \quad (3)$$

Applying Eq. (2) to the formula of the volume of a tubular reactor, it is possible to calculate directly the diameter of the reactor (D), reaching Eq. (4) [16].

$$D = \sqrt{\frac{4\,V_{reactor}}{\pi L}} \quad (4)$$

Regarding the removal of O_2, Lu et al. [15] also suggests a tubular flow reactor, which can be designed using the previous expressions.

MEA absorption

The goal of the MEA absorption process is to separate N_2 from CO_2 in the treated flue gas. For this, two MEA absorption systems were simulated: one with 90 % carbon capture efficiency and the other with an almost complete CO_2 capture. The two simulations differ only on the reboiler duty needed for stripping the CO_2 from the solvent, to regenerate it. An AspenTech example [19] was adapted with the chemistry and reaction models retrieved from the work of Madeddu et al. [9] to build a rate-based model. The main equipment consists of two packed columns, which were simulated as RadFrac models with Sulzer Mellapak 250Y packing: an absorber, where the flue gas enters and contacts with the MEA solvent, which absorbs the CO_2, and a stripper, where the CO_2-rich solvent is separated from the CO_2 and pumped

back into the absorber. The design of both the columns followed the method of Madeddu et al. [9].

Methanation

The Sabatier process can have two objectives: to solely produce CH_4 or to prevent the poisoning of the downstream Haber-Bosch system by converting CO_2 into CH_4. For the simulation, the reaction of Eq. (5) was implemented with the kinetics retrieved from the study of Koschany et al. [20] with a $NiAl(O)_x$ catalyst.

$$CO_2 + 4H_2 \rightleftarrows CH_4 + 2H_2O \quad (5)$$

Before the reactors, the stream is compressed to 7 bar, not only because the reaction is favoured by higher pressure, but also because the natural gas distribution system operates at this pressure, and the product will be at the required conditions for integration. To achieve this, a three-stage compression system with interstage cooling is implemented. After the stream is compressed and mixed with H_2 in a 1:4 mole ratio, it is heated to 290-320 °C and goes through 3-4 Sabatier reactors, simulated by RPlug models, with condensation between each stage. As in the work of Marnate et al. [21], multitubular fixed-bed reactors were chosen for methanation. The dimensions of each tube were fixed at 2 m in height and 1 in. in diameter. Using Eq. (3), the required reactor volume can be calculated, assuming a catalyst density of 2300 kg/m^3 [22]. The mass of catalyst was adjusted to maintain the gas flowrate at 120 NL/h/g_{cat} [20]. Dividing the required reactor volume by the volume of each tube yields the required number of tubes.

Haber-Bosch

The H-B process aims to produce NH_3 from N_2. An AspenTech example model [19] was adapted, and reaction of Eq. (6) was implemented with the kinetics studied by Tripodi et al. [23] in iron-based catalysts.

$$N_2 + 3H_2 \rightleftarrows 2NH_3 \quad (6)$$

The inlet stream is mixed with H_2 at a 1:3 molar ratio and then compressed to 150 bar to favour the reaction [21], using four compression stages with intermediate cooling. The stream is heated to 400-540 °C before entering each reactor, depending on the optimal temperature found. After passing through 3 reactors with intermediate cooling, the product is condensed at -35 °C, and liquid NH_3 is separated from the unreacted gases that must be recycled back into the first reactor. The design of the Haber-Bosch reactors involved a sensitivity analysis varying the mass of catalyst and checking its influence on the conversion of each reactor: when the maximum conversion was reached, the value for the mass of catalyst was found. This way, employing Eqs. (3) and (4), it was possible to compute the dimensions of each reactor, using 2900 kg/m^3 [24] as ρ_{cat}.

Economic analysis

The cost of each piece of equipment for each other process was primarily estimated using the Aspen Process Economic Analyzer® (APEA) [25]. The equipment costs were then updated to 2025 using the corresponding CEPCI value [26]. The capital expenditure (CAPEX) and the operational expenditure (OPEX) were calculated using the factorial method, as described by Timmerhaus et al. [27]. The CAPEX was annualized using a 0.098 factor, assuming a plant lifespan of 20 years and a discount rate of 7.5%. The OPEX was calculated based on the values in Tab. 2. Adding the OPEX to the annualized capital cost enables the economic analysis of every process and, consequently, every superstructure route.

Table 2: Main values for the calculation of the OPEX and of the revenues [28-34].

Parameter	Value	Unit
H_2 price	5.1	€/kg
Electricity cost	0.06	€/kWh
Cooling water	0.01	€/m^3
Refrigerant	2.8	€/GJ
Steam	0.04	€/kg
NH_3 value	521	€/t
CH_4 value	1806	€/t
O_2 value	290	€/t
Storage cost	30	€/t
Emission cost	70	€/t

Although the Air Separation Unit was not modelled using Aspen Plus® v14, the associated costs were also evaluated. In this case, Eq. (7) [35] was used for the estimation of the equipment cost (EC), being $\dot{m}_{O_2}$ the mass flowrate of O_2 entering the system, in tonnes per day.

$$EC\ (M€_{2014}) = 22 \bullet \left(\frac{\dot{m}_{O_2}[t/d]}{432}\right)^{0.6} \quad (7)$$

Electric and steam consumptions of 230 kWh/t_{O2} and 16.2 kWh/t_{O2} were considered, respectively [35]. The value of the mass flowrate of O_2 used was calculated considering the mass flowrate of N_2 that enters the 'Haber-Bosch 2' process and the mole fractions of N_2 (~79 %) and O_2 (~21 %) in the air.

Since not all superstructure routes produce the same products, these were assumed to be sold to enable an economic comparison of the routes. In this sense, the revenue was calculated by multiplying the generated flowrate by the respective unitary values. In routes that do not generate any product, such as $A90_E^{CCS}, A100_E^S$ or $A0$, no revenue was calculated, but the emissions/storage fees were applied. Adding the costs, and subtracting the revenue, leads to the overall annual cost of each route. This value is then divided by the initial mass flowrate of the cement plant flue gas (42.8 t/h) to assess the cost of treating it per metric tonne.

RESULTS AND DISCUSSION

The economic results are detailed and discussed in this section for each individual process considered in the superstructure and for each route of the superstructure.

Pre-treatment

The dimensions of the scrubber, SCR and O_2 removal reactor are presented in Tab. 3.

Table 3: Design of the main equipment in the pre-treatment section.

Equipment	Height/Length (m)	Diameter (m)
Scrubber	13.0	4.3
SCR reactor	5.8	1.9
O_2 removal reactor	2.4	0.8

After determining the equipment dimensions, the total estimated cost was calculated at approximately 1.56 M€. The simulation enabled the calculation of utility and raw material consumption costs, which totaled 7.9 M€/y, indicating a profit, since steam is generated in this process and can be used in downstream processes.

MEA absorption

Subsequently to the pre-treatment process, the resulting stream can go through MEA absorption that can either capture 90 % or approximately 100 % of the CO_2 present. The dimensions of the absorber and stripper columns are similar for both capturing efficiencies and are presented in Tab. 4.

Table 4: Design of the main equipment in the capturing section.

Equipment	Height (m)	Diameter (m)
Absorber	26.9	3.4
Stripper	12.0	2.6

The equipment costs for both processes were 3.4 M€, while the total expenses for raw materials and utilities were 9.7 M€/y for the 'MEA absorption (90 %)' process and 76.6 M€/y for the 'MEA absorption (~100 %)' process. The large difference in operating costs reflects different steam consumption needs: the last process requires a higher heat duty to regenerate the MEA solvent because it absorbs more CO_2.

Methanation

The Sabatier processes have the objective of solely producing CH_4 ('Sabatier') or of eliminating the CO_2 content of the stream ('Methanation 1' and 'Methanation 2'), readying the stream for a subsequent H-B process. The design of these processes is presented in Tab. 5.

Table 5: Number of tubes (N_{tubes}) and mass of catalyst ($m_{catalyst}$) required for each reactor (R) of the Sabatier (S), Methanation 1 (M1) and Methanation 2 (M2) processes.

	N_{tubes}			$m_{catalyst}$ (kg)		
	S	M1	M2	S	M1	M2
R1	1855	1346	3151	240	185	451
R2	665	1300	1860	86	162	240
R3	460	1370	1768	56	161	215
R4	-	-	1867	-	-	214

The estimated equipment costs are 2.0, 3.7, and 5.2 M€ for the 'Sabatier', 'Methanation 1', and 'Methanation 2' processes, respectively, as they progressively handle higher flow rates and increase the dimensions of the respective equipment. Regarding the consumption of utilities and raw materials, the respective values are 86.2, 10.6, and 96.5 M€/y for the processes. The values differ due to the different needs for green H2 in each process: the inlet stream of the 'Methanation 1' process contains approximately 10 % CO2 relative to the other methanation processes, leading to the lowest green hydrogen consumption and thus the lowest operating costs.

Haber-Bosch

The design of the reactors of the Haber-Bosch process is presented in Tab. 6.

Table 6: Reactor (R) design and mass of catalyst ($m_{catalyst}$) required for each reactor (R) of the Haber-Bosch (HB) processes.

	Diameter (m) x Length (m)			$m_{catalyst}$ (kg)		
HB	1	2	3	1	2	3
R1	1.0x3.1	1.0x3.0	1.4x4.2	2	2	13
R2	1.3x3.9	1.2x3.7	1.8x5.4	10	9	34
R3	1.6x4.8	1.5x4.6	2.6x7.9	24	21	120

The estimated equipment costs for producing NH_3 through the H-B process are 11.7, 11.6, and 18.5 M€ for the 'Haber-Bosch 1', 'Haber-Bosch 2', and 'Haber-Bosch 3' processes, respectively. The 'Haber-Bosch 3' has the highest equipment costs because it handles the largest amount of inert gas (CH_4). This translates into the largest circulating flowrates and, consequently, the largest equipment dimensions.

As for the consumption of utilities and raw materials, all three processes have a value of approximately 211 M€/y, as their green H_2 consumption is similar because their respective inlet streams contain equal amounts of N_2.

Air Separation Unit

The estimated equipment cost of the ASU is 17.3 M€, while the estimated annual cost of electricity and steam is 0.9 M€.

Superstructure

Tab. 7 summarizes the results of the economic assessment for each possible superstructure route, each designated by its respective nomenclature (see Tab. 1).

It is evident that OPEX has a stronger influence on production costs than CAPEX, which is expected, since annualized CAPEX is 1-10% of OPEX. The high OPEX of the processes is mainly related to their high energy and green H_2 consumption.

Regarding route revenues, those that include the H-B process to produce green ammonia yielded higher revenues, due to the high N_2 content in the flue gas stream, which is converted into green ammonia. However, the routes that consider the emission of the CO_2-lean stream after the separation unit ($A90_E^S$,$A90_E^{CCS}$,$A100_E^S$ and $A100_E^{CCS}$), present fewer total costs than the ones that include a H-B system. This means that, despite having the highest revenues, the Haber-Bosch process is currently not economically viable, due to its large equipment costs and high green H_2 consumption.

When comparing the routes producing green ammonia from the cement plant flue gas with the benchmark technology to produce green ammonia (ASU_{HB}), the latter presents lower total costs, which means that it is currently still uncompetitive to utilize the N_2 from the cement plant flue gas instead of using an ASU.

Tab. 7 also reflects the lower costs of the routes including a CO_2 capturing unit with 90 % efficiency when compared with the ones comprising ~100 % MEA absorption systems. For example, route $A90_E^{CCS}$ is approximately 70 % less expensive than $A100_E^{CCS}$. This is mainly related to the high medium-pressure (MP) steam consumption required in the MEA absorption unit to capture approximately 100 % of the CO_2 emissions.

Among the routes that produce both green ammonia and SNG, route $A100_{HB}^S$ has the highest revenue value due to producing the largest amount of methane and ammonia. However, it also presents the highest total cost, since it not only requires the high green H_2 consumption of both the Haber-Bosch and the Sabatier processes, but it also contains the 'MEA absorption (~100 %)' system, which has a high utility demand. Route $A90_{HB}^S$ appears to be a more promising option for producing these two products, since it has lower energy requirements in the CO2-capture unit, leading to significantly lower OPEX. On the other hand, the route $A0$ exhibits the lowest OPEX value of the three, as it does not contain a CO_2 capturing unit. For this same reason, it needs to handle the total flue gas flowrate, which requires larger equipment dimensions, therefore it presents the highest CAPEX. The $A0$ route is the one which presents the lowest revenues, due to the need of higher purge flowrates in the process to accommodate the high flowrate of CH_4 in the H-B process. The total net cost of treating the flue gas in this route is similar to the $A90_{HB}^S$ route, indicating that the $A90_{HB}^S$ and $A0$ are the most promising routes for the simultaneous production of SNG and green ammonia from a cement plant's flue gas. However, these two routes present total costs approximately 40 times higher than the emission of the cement plant flue gas. This low economic attractiveness of the simultaneous production of e-methane and green ammonia from the cement plant flue gas is mainly related to the high costs of green hydrogen.

Although not producing methane or ammonia, route $A90_E^{CCS}$ is currently the less expensive path to treat the cement plant flue gas. Nonetheless, it is worth noticing the relatively low total cost of path $A90_E^S$, which includes methanation of the CO_2-rich outlet of the 'MEA absorption (90 %)' process and emission of the CO_2-lean outlet stream, as it presents the second lowest total costs while producing SNG through the Sabatier process. However, under current conditions, no route is economically competitive with the cement plant's total emissions from flue gas. The greatest contributor to the high cost of this route is the cost of green H_2, which is expected to decrease in the coming years [36]. By analysing the economic feasibility with the green H_2 price evolution, it is possible to conclude that the $A90_E^S$ route becomes the most profitable route when the green H_2 price reaches 1.81 €/kg and that its break even H_2 price is 1.51 €/kg of H_2.

Conclusions

Given the cement industry's elevated GHG emissions, this work proposes applying SNG and green ammonia production processes to cement plant flue gas. The most mature technologies were investigated and assembled into a superstructure comprising different routes, which were modelled using Aspen Plus®. Firstly, a pre-treatment was employed to remove unwanted components from the flue gas. Then, the possibility of separating N_2 from CO_2 was studied by applying a MEA absorption system with either 90 % or ~100 % capturing efficiency.

The obtained N_2-rich stream could be valorized to NH_3 via the Haber-Bosch process or emitted, while the CO2-rich stream could be sent to storage or converted to CH_4. Comparing the MEA units studied, as the complete CO_2 removal is approached, the system requires higher reboiler duties to strip the CO_2 from the solvent due to a substantial increase in MP steam consumption. Thus, the 90 % capture configuration appears to be more energetically and economically advantageous. The ~100% separation unit has the advantage of not requiring a subsequent methanation step to treat the N^2-rich stream before it enters the Haber-Bosch process. However, the cost increase of applying a ~100 % separation unit is greater than the cost of adding an extra

Table 7: Economic analysis of the superstructure routes.

Route	$A90^{S}_{HB}$	$A90^{S}_{E}$	$A90^{CCS}_{E}$	$A90^{CCS}_{HB}$	$A100^{S}_{HB}$	$A100^{S}_{E}$	$A100^{CCS}_{E}$	$A100^{CCS}_{HB}$	$A0$	ASU	E
EC (M€)	22.4	6.9	4.9	20.4	18.6	7.0	4.9	16.5	25.2	28.9	-
CAPEX (M€/y)	13.1	4.1	3.0	12.0	11.0	4.1	3.0	9.7	14.8	17.3	-
OPEX (M€/y)	477.7	152.0	31.0	356.7	587.5	280.2	122.9	430.1	468.9	332.7	-
CH_4 sales (M€/y)	66.5	66.5	-	-	73.8	73.8	-	-	62.4	-	-
NH_3 sales (M€/y)	111.1	-	-	111.1	113.6	-	-	113.6	101.9	113.6	-
O_2 sales (M€/y)	-	-	-	-	-	-	-	-	-	16.3	-
Total sales (M€/y)	177.6	66.5	-	111.1	187.4	73.8	-	113.6	164.3	129.3	-
Emissions (M€/y)	-	0.7	0.7	-	-	0.0	0.0	-	-	-	7.5
Storage (M€/y)	-	-	3.0	3.0	-	-	3.3	3.3	-	-	-
Total costs (M€/y)	313.2	90.3	37.6	260.5	411.1	210.6	129.1	329.5	319.5	220.7	7.5
€/t of treated FG	910	263	109	757	1195	612	375	958	928	-	22

methanation system.

The routes considering carbon capture with 90 % efficiency, with the emission of the N_2-rich stream and either the methanation or storage of the CO_2-rich stream are the most economically promising alternatives to treat the flue gas stream, with respective total costs of 263 and 109 €/t of flue gas. However, they are currently uncompetitive due to flue gas emissions (22 €/t).

Nevertheless, it is possible that H2-requiring routes become economically feasible as green hydrogen costs are expected to decline in the coming years. This trend indicates great potential to integrate pathways for synthesizing high-value chemical products into the cement industry, contributing to GHG mitigation and the transition to a more sustainable industrial sector.

ACKNOWLEDGEMENTS

Project co-financed by the PRR - Recovery and Resilience Plan by the European Union, in particular, the Mobilizing Agenda "Moving2Neutrality" (Project no. 32, with reference no. C644927397-00000038), and C5Lab – Sustainable Construction Materials Association [CENTRO-04-3559-FSE-000096; LISBOA-05-3559-FSE-000008; C05-i02/2022] and CERENA [FCT-UID/04028/2025]. This work was also supported by the Young Researcher and Innovator Conference grant awarded to Ana S. Amorim by COST Action on Techno-economic analysis of carbon mitigation technologies (TrANsMIT, CA21127).

AUTHOR IDENTIFIERS

Pedro MA: 0009-0008-8992-3077
Amorim AS: 0000-0002-9392-1081
Matos HA: 0000-0001-8128-7346

REFERENCES

1. NASA. https://science.nasa.gov/climate-change/
2. International Energy Agency. https://www.iea.org/reports/global-energy-review-2025
3. Plaza MG, Martínez S, Rubiera F. CO2 capture, use, and storage in the cement industry: state of the art and expectations. Energies 13:5692 (2020). https://doi.org/10.3390/en13215692
4. Ryu KH, Lee D, Hwang B. Sustainability assessment for the cement flue gas methanation process using dual functional material (DFM). Journal of Environmental Chemical Engineering 12:112451 (2024). https://doi.org/10.1016/j.jece.2024.112451
5. Chauvy R, Dubois L, Lybaert P, Thomas D, De Weireld G. Production of synthetic natural gas from industrial carbon dioxide. Applied Energy 260:114249 (2020). https://doi.org/10.1016/j.apenergy.2019.114249
6. Schakel W, Hung CR, Tokheim LA, Strømman AH, Worrell E, Ramírez A. Impact of fuel selection on the environmental performance of post-combustion calcium looping applied to a cement plant. Applied Energy 210:75-87 (2018). https://doi.org/10.1016/j.apenergy.2017.10.123
7. Schorcht F, Kourti I, Scalet BM, Roudier S, Sancho LD. Best Available Techniques (BAT) Reference Document for the Production of Cement, Lime and Magnesium Oxide (2010).
8. Badreldin A, Li Y. A critical appraisal of advances in integrated co₂ capture and electrochemical conversion. Chem. Sci. 16:2483-2513 (2025). https://doi.org/10.1039/d4sc06642a
9. Madeddu C, Errico M, Baratti R. CO2 capture by reactive absorption-stripping. Springer International Publishing (2019). https://doi.org/10.1007/978-3-030-04579-1
10. Narciso DAC, Pires JM, Fortunato J, Teixeira P, Castro PM, Pinheiro CIC, Matos HA. Design and operation of power-to-ammonia systems: a review. Energy Conversion and Management 327:119494 (2025). https://doi.org/10.1016/j.enconman.2025.119494

11. Aspen Technology, Inc. Aspen Exchanger Design and Rating, Version 14.2 (2023).
12. Krylova AV, Nefedova NV, Peev TM. Poisoning and Passivation of Ammonia Synthesis Catalyst at Various Temperatures. Studies in Surface Science and Catalysis, 34(C) (1987). https://doi.org/10.1016/S0167-2991(09)60397-9
13. Duque C, Montes C, Bustamante F, Ortiz A. Simulation of two alternatives for SO2 removal from wet cement kiln exhaust gases. Rev.Fac.Ing.Univ.Antioquia :49-57 (2013). https://doi.org/10.17533/udea.redin.14652
14. Sweeney AJ, Liu YA. Use of simulation to optimize no_x abatement by absorption and selective catalytic reduction. Ind. Eng. Chem. Res. 40:2618-2627 (2001). https://doi.org/10.1021/ie0005295
15. Lu H, Jiang Y, Abiodun O, Schideman L, Kuhn A, Yang H, Lu Y. Catalytic removal of oxygen impurities from pressurized oxy-combustion flue gas for the production of high-purity carbon dioxide. Energy Fuels 36:2701-2711 (2022). https://doi.org/10.1021/acs.energyfuels.1c04341
16. Coulson JM, Richardson JF. Coulson & Richardson's Chemical Engineering: Chemical Engineering Design. Ed: Elsevier (2005).
17. Woods DR. Rules of thumb in engineering practice. Wiley (2007). https://doi.org/10.1002/9783527611119
18. Leitão J. Haber-Bosch process alternatives for the production of green ammonia. Master's thesis, Instituto Superior Técnico (2023).
19. Aspen Technology, Inc. Aspen plus, Version 14 (2022).
20. Koschany F, Schlereth D, Hinrichsen O. On the kinetics of the methanation of carbon dioxide on coprecipitated nial(o). Applied Catalysis B: Environmental 181:504-516 (2016). https://doi.org/10.1016/j.apcatb.2015.07.026
21. Marnate K, Amorim AS, Sirigina D, Skorek-Osikowska A, Nazir SM. Techno-Economic Analysis of Carbon Capture & Utilization: Biogas to Methane. *EU Cost Action Virtual Mobility Grant* (2024).
22. Twigg M. Catalyst Handbook. Ed: Wolfe Publishing, Ltd (1989).
23. Tripodi A, Compagnoni M, Bahadori E, Rossetti I. Process simulation of ammonia synthesis over optimized ru/c catalyst and multibed fe + ru configurations. Journal of Industrial and Engineering Chemistry 66:176-186 (2018). https://doi.org/10.1016/j.jiec.2018.05.027
24. Nielsen A. Ammonia- Catalysis and Manufacture. Ed: Springer (1995).
25. Aspen Technology, Inc. Aspen Process Economic Analyzer (APEA), Version 14 (2022).
26. University of Manchester. CEPCI
27. Peters MS, Timmerhaus KD, West RE. Plant Design and Economics for Chemical Engineers. Ed: McGraw-Hill (2003).
28. Departamento de Engenharia Química. Projeto em Engenharia Química (PEQ). Course, 2024/2025.
29. Dionísio J. Com licença?. ptp 3: (2025). https://doi.org/10.51427/cet.ptp.2025.3.6
30. Gorman M. Fertilizers | S&P Global
31. International Energy Agency. Global hydrogen review (2024).
32. Smith E, Morris J, Kheshgi H, Teletzke G, Herzog H, Paltsev S. The cost of CO2 transport and storage in global integrated assessment modeling. International Journal of Greenhouse Gas Control 109:103367 (2021). https://doi.org/10.1016/j.ijggc.2021.103367
33. Trading Economics. EU Carbon Permits - Price - Chart - Historical Data - News
34. imarc. Bulk Oxygen Price Trend, Index, Chart and Forecast
35. De Lena E, Spinelli M, Gatti M, Scaccabarozzi R, Campanari S, Consonni S, Cinti G, Romano MC. Techno-economic analysis of calcium looping processes for low CO2 emission cement plants. International Journal of Greenhouse Gas Control 82:244-260 (2019). https://doi.org/10.1016/j.ijggc.2019.01.005
36. Curcio E. Techno-economic analysis of hydrogen production: costs, policies, and scalability in the transition to net-zero. International Journal of Hydrogen Energy 128:473-487 (2025). https://doi.org/10.1016/j.ijhydene.2025.04.013

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

CFD-based optimal design of a portable and stackable alkaline water electrolyser for hydrogen production

Akepogu Venkateshwarlu[a], Gianluca Li-Puma[ab], and Brahim Benyahia[a]*

[a] Department of Chemical Engineering, Loughborough University, LE11 3TU Leicestershire, UK
[b] Chemical and Environmental Engineering, University of Palermo, Piazza Marina, 61 90133 Palermo, Italy
* Corresponding Author: b.benyahia@lboro.ac.uk.

ABSTRACT

Hydrogen is increasingly recognized as a vital energy carrier for a sustainable future. Among the various methods for hydrogen production, alkaline water electrolysis (AWE) stands out as a well-established and commercially viable option. However, their more effective deployment requires more advanced, portable, and scalable designs. This study explores systematic model-based shape optimization of the next generation AWE based on computational fluid dynamic (CFD) aimed to enhance the hydrodynamics and electrochemical performance. Several design geometries and arrangements were proposed including flow baffles to enhance hydrodynamic and facilitate detachment of oxygen and hydrogen bubbles. The findings indicate that the optimal design and location of the baffles improve fluid mixing and enhance bubble detachment, resulting in a more uniform electrolyte distribution and decreased concentration polarization. Several key performance indicators were considered to analyse the performance of proposed designs including gas production rates, polarization curves, and fluid flow velocity profiles. The insights gained from this research offer valuable recommendations for optimizing flow field designs in alkaline water electrolyzers, aiming to enhance efficiency and operational robustness.

Keywords: Alkaline water electrolysis, zero-gap cell, CFD, Pyramidal pins, Mesh electrode, Multiphysics model

INTRODUCTION

Hydrogen is becoming essential in the transition to decarbonized energy systems, influenced by climate targets and energy security [1, 2]. The global demand for hydrogen is expected to increase significantly, shifting from fossil-based “grey hydrogen” to low-carbon “green hydrogen, ” produced through water electrolysis using renewable energy [3]. Hydrogen production via water electrolysis is a key enabling technology for future low-carbon energy systems[4]. There are three primary methods for water electrolysis: alkaline water electrolysis (AWE), proton exchange membrane electrolysis (PEM), and solid oxide electrolysis (SOEC) and among them, alkaline water electrolysis (AWE) is the most established for hydrogen production [1, 3–6]. It offers advantages such as lower material costs, reliance on non-noble catalysts, longer stack lifetimes, and scalability for large deployments, especially in regions abundant in renewable resources [5, 7, 8]. AWE's reliability and adaptability to variable operating conditions are crucial as electrolysis facilities align with fluctuating renewable energy supplies. This adaptability makes AWE an economically viable option for large-scale hydrogen production [9, 10]. As the world moves towards sustainable energy solutions, AWE is positioned to play a significant role in meeting the rising demand for green hydrogen[4].

Alkaline water electrolysis (AWE) faces significant challenges due to complex hydrodynamics, particularly the behavior of gas bubbles at the electrode-electrolyte interface [6, 11, 12]. These bubbles can create localized concentration gradients, increase ohmic losses, and generate thermal hotspots, ultimately diminishing energy efficiency [12]. Factors such as electrode shape, channel geometry, operating conditions, and cell integration strategies influence the distribution of gas and liquid, making the optimization of bubble behavior and internal flow essential for enhancing system performance [7, 9, 11]. Traditional AWE designs often utilize concave-convex spherical (SCC) flow channels to improve mixing; however, these can also lead to stagnant regions and bubble trapping, resulting in non-uniform mass transfer

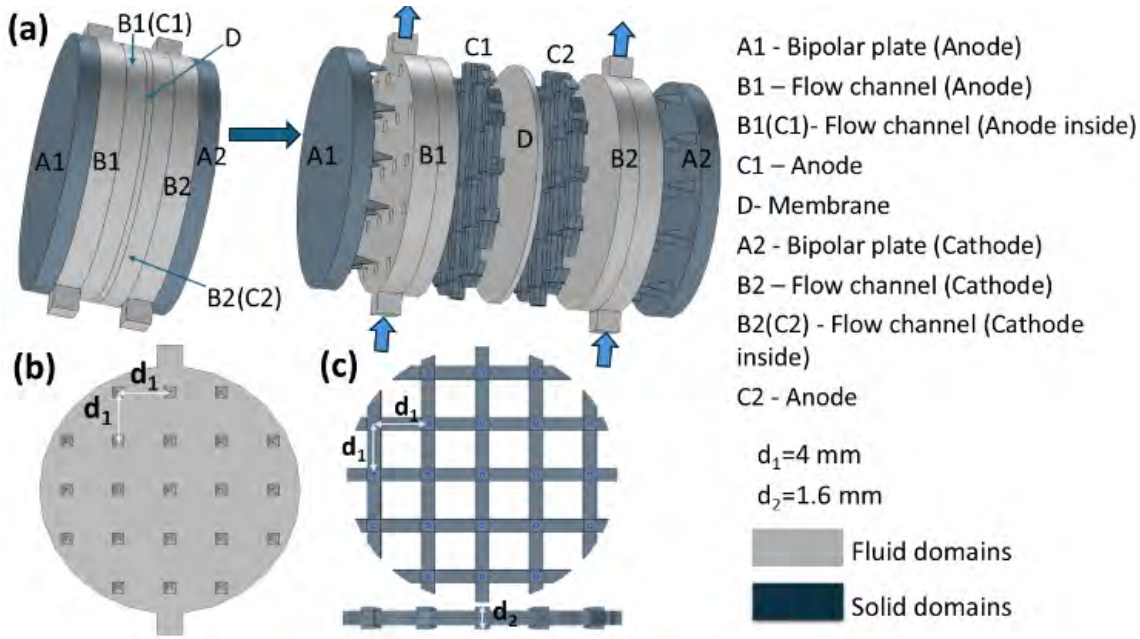

Figure 1. Schematics of the alkaline wate electrolyser with individual components, (b) arrangment of the pyramidal pins in the flow channel, and (c) mesh electrode.

[12]. A potential solution involves replacing porous electrodes with mesh-type structures, which may decrease bubble residence time and improve gas release [13]. Nonetheless, the interplay between mesh electrodes, SCC flow fields, and various operational variables have yet to be thoroughly investigated in a cohesive framework.

With increasing computational power, computational fluid dynamics (CFD) methods have gained widespread adoption across multiple disciplines. These methods have been widely implemented to support the optimal design and operation of multiphase systems across multiple scales, from microfluidic devices to mesoscale processes and full industrial-scale systems [14, 15]. Recent CFD applications highlight systematic geometry and shape optimization of small-scale packed-bed reactors with complex internal structures [16]. These methods have opened new opportunities for more effective in silico prototyping, enabling faster design, evaluation, and optimization of next-generation processes and technologies. The proposed CFD-based design framework of zero-gap alkaline electrolyzers allows more robust geometries and configuration through a detailed three-dimensional multiphysics simulation which addresses the limitations of the existing designs. The method incorporates crucial aspects such as heat transfer, gas-liquid transport, and electrochemical kinetics, facilitating a comprehensive assessment of how localized gradients affect system performance. The model evaluates critical parameters including cell voltage, temperature, and inlet flowrate collectively, revealing their interactions and providing insights for optimizing device performance beyond single-parameter analysis. Most importantly, the CFD-based approach reveals more effective alternatives to the traditional porous electrodes such as mesh-type electrodes, enhancing productivity and hydrodynamic performance particularly more effective bubble detachment and transport. The identified optimal designs featuring pyramidal pin structures on the bipolar plates. The method also informed the best trade-offs between mixing and bubble detachment. Ultimately the framework offers essential design insights laying the foundation for enhanced efficiency and durability of industrial green hydrogen production systems.

METHODOLOGY

A numerical model representing a single cell of an alkaline water electrolyzer is illustrated in **Figure 1**. The computational domain encompasses flow channels, cathode and anode mesh electrodes, bipolar plates, and a porous membrane region that separates the two half-cells. The electrodes are of non-porous made of Nickel metal. Pyramidal pin structures are integrated into the flow channels and attached to the bipolar plates, with sharp edges oriented towards the electrodes to enhance flow distribution and induce local turbulence. The geometric dimensions and physical properties are detailed in **Table 1** and **Table 2**.

Table 1: Geometric parameters of the alkaline water electrolyser model.

Description	Value
Cathode and anode radius	10 mm
Cathode thickness	1.6 mm
Anode thickness	1.6 mm
Flow channel radius	10.1 mm
Flow channel thickness	3.6 mm
Bipolar plate radius	10 mm
Bipolar plate thickness	0.5 mm
Membrane radius	10.1 mm
Membrane thickness	0.5 mm

Table 2: Physical parameters of the alkaline water electrolyser model.

Parameter	Value
Membrane porosity	0.5
Reference exchange current density (anode)	100 A/m^2
Reference exchange current density (cathode)	0.1
Transfer coefficient (anode)	0.5
Transfer coefficient (anode)	0.5
Inlet electrolyte velocity	0.01 m/s - 0.2 m/s
Outlet pressure	1 atm
Electrolyte conductivity	40 S/m
Bipolar plate conductivity	4.032e6 S/m
Nickel mesh electrode conductivity	13.8e6 S/m
Electrode active surface area	6.06 cm^2
Operating temperature	80 C
Cell voltage	1.3 V – 2.3 V
Electrolyte concentration (KOH)	3 M

Governing equations

The following equations are employed to solve the physics described above:

Electrochemical model

Alkaline water electrolysis splits water into hydrogen and oxygen using electrochemical reactions at the cathode and anode. An aqueous alkaline solution like KOH serves as the electrolyte, supplying hydroxide ions (OH^-) for charge transport, driven by electrical potential.

At cathode (hydrogen evolution reaction – HER):

$$2H_2O + 2e^- \rightarrow H_2 + 2OH^- \tag{1}$$

At anode (oxygen evolution reaction – OER):

$$4OH^- \rightarrow O_2 + 2H_2O + 4e^- \tag{2}$$

The total cell voltage is calculated by

$$E_{total} = E_{equi} + \eta_{acti} + \eta_{ohmic} \tag{3}$$

The equilibrium cell potential is calculated from the Nernst equation as follows:

$$E_{equi} = E_{0,ref}(T) + \frac{RT}{2F}\ln\left(\frac{p(O_2)^{0.5}p(H_2)}{p(H_2O)}\right) \tag{4}$$

The Butler-Volmer equation describes the relationship between the local current density and activation overpotential as follows:

$$i_{cathode} = a_{HER}i_{0,HER}(1-s_g)\left[exp\left(\frac{\alpha_c F\eta_{HER}}{RT}\right) - exp\left(\frac{(1-\alpha_c)F\eta_{HER}}{RT}\right)\right] \tag{5}$$

$$i_{anoode} = a_{OER}i_{0,OER}(1-s_g)\left[exp\left(\frac{\alpha_a F\eta_{OER}}{RT}\right) - exp\left(\frac{(1-\alpha_a)F\eta_{OER}}{RT}\right)\right] \tag{6}$$

The equations for the transport of electrons and ions are given below:

$$\nabla \bullet \boldsymbol{i}_s = \nabla \bullet \left(-\sigma_s^{eff}\nabla\phi_s\right) = S_{\phi,s} \tag{7}$$

$$\nabla \bullet \boldsymbol{i}_l = \nabla \bullet \left(-\sigma_l^{eff}\nabla\phi_l\right) = S_{\phi,l} \tag{8}$$

$$\sigma_l^{eff} = \sigma_l(1-s_g)^{1.5} \tag{9}$$

The ohmic overpotential resulting from the transport resistances of both electrons and ions is described as follows:

$$\eta_{ohmic} = i\left(\frac{\phi_l}{\sigma_l^{eff}} + \frac{\phi_s}{\sigma_s^{eff}}\right) \tag{10}$$

Mass transfer model

The Navier-Stokes equations are employed to determine the fluid velocity and pressure within the flow channels of the electrolyzer in the following manner.:

$$\frac{\partial\rho}{\partial t} + \nabla \bullet (\rho\boldsymbol{u}) = 0 \tag{11}$$

$$\rho\frac{\partial\boldsymbol{u}}{\partial t} + \rho(\boldsymbol{u} \bullet \nabla)\boldsymbol{u} = \nabla \bullet \left[-p\boldsymbol{I} + \mu\left(\nabla\boldsymbol{u} + (\nabla\boldsymbol{u})^T\right)\right] + \boldsymbol{F} \tag{12}$$

The phase transport equations have been adopted to describe the distribution of gas-liquid flow as follows:

$$\frac{\partial\rho_i s_i}{\partial t} + \nabla \bullet (\rho_i\boldsymbol{u}_i) = 0 \tag{13}$$

$$\frac{\partial\rho_i s_i}{\partial t} + \nabla \bullet \left[-\rho_i K\frac{K_{ri}}{\mu_i}(\nabla p_i - \rho_i\boldsymbol{g})\right] = Q_i \tag{14}$$

$$Q_{H_2} = \frac{M_{H_2}}{2F}i_v \tag{15}$$

$$Q_{O_2} = \frac{M_{O_2}}{4F}i_v \tag{16}$$

Heat transfer model

The energy equation has been employed to describe the thermal management inside the electrolyser:

$$(\rho C_p)_{eff}\frac{\partial T}{\partial t} + \rho C_p\boldsymbol{u} \bullet \nabla T + \left(-k_{eff}\nabla T\right) = S_{heat} \tag{17}$$

$$S_{heat} = |i_v||\eta_{acti}| + \frac{i_v^2}{\sigma_s} + \frac{i_v^2}{\sigma_l} \quad (18)$$

The governing equations delineated herein were resolved utilizing CFD through the COMSOL Multiphysics 6.4, employing the finite element method (FEM) [12]. A convergence criterion of relative error less than 0.0001 is applied, using the backward difference formula (BDF) method for time-stepping. This analysis was conducted under specific boundary conditions and underlying assumptions.

Boundary conditions

The KOH electrolyte solution enters the flow channels of anode and cathode at their inlets with a defined velocity, respectively. The cell voltage is applied on the bipolar plate and other side of the bipolar is made electric ground. The pressure at the outlets of the flow channels is 0 atm. The gravitational forces during the operation are neglected. The OH^- ions are allowed to pass through the porous membrane.

Assumptions

The flow is assumed to be laminar. There is no crossover of hydrogen and oxygen gases across the membrane. Overpotential losses due to contact resistance inside the electrolyser are assumed to be negligible.

RESULTS AND DISCUSSION

Validation

To assess the accuracy of the multi-physics alkaline electrolyzer model, operational parameters from existing literature were employed [12, 17], specifically at 353.15 K and 1 atm, as shown in **Figure 2**. A comparative evaluation of simulation outcomes alongside experimental data [17] and other literature simulations [12] shows a strong correlation within a current density range of 0.1 $\mathrm{A/cm^2}$ to 1.4 $\mathrm{A/cm^2}$. However, at higher current densities, simulated voltages exhibit slight underestimations, likely attributable to the omission of ohmic losses and diffusion voltage effects.

Effect of pyramidal pins on the velocity distribution

The study analyzed the effect of inlet electrolyte velocity on flow patterns in channels with 10 and 24 pyramidal pins on each side of the electrode, as shown in **Figures 3** and **4**. It has been observed that the flow patterns at velocities from 0.01 m/s to 0.2 m/s demonstrated that the pyramidal pin structures induced changes in the flow field, including local redirection, shear-layer formation, and wake recirculation. These alterations were found to depend significantly on both the inlet velocity and the number of the pins.

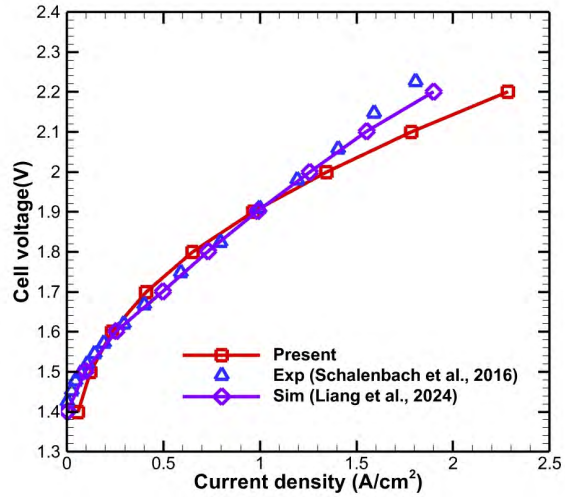

Figure 2. Comparison of the present simulation results against the experimental data [17] and other literature simulations [12].

At low velocity of 0.01 m/s, electrolyte motion was dominated by viscous forces. The 10 pins configuration caused minimal axial flow disruption (refer to **Figure 4(a)**), while the 24 pins setup resulted in increased tortuosity and narrower passages (refer to **Figure 4(a)**), leading to longer residence times and persistent recirculation zones, highlighting significant micro-circulation effects typically seen in creeping flow. Increasing inlet velocities of 0.05 m/s and 0.1 m/s intensified inertial effects, enhancing flow separation and wake formation. Higher pyramidal pin number improved hydrodynamic interactions, promoting vortex-dominated flow, which is crucial for enhancing bubble detachment and gas-liquid mass transfer in electrolyzer channels (refer to **Figures 3(b)-(c)** and **4(b)-4(c)**). At 0.2 m/s, strong convection dominated the flow field. In a 10 pins setup, limited recirculation occurred, while a 24 pins configuration enhanced flow complexity and inter-sphere mixing, resulting in maximum velocity uniformity and minimized small eddies with larger wake structures (refer to **Figures 3(d)** and **4(d)**).

Effect of inlet velocity on the gas volume fraction

The two-phase behavior in the channel arises from the hydrogen evolution reaction at the nickel mesh cathode in alkaline conditions. Gas bubbles form, grow, and detach due to nickel's wettability, leading to localized gas accumulation and high hydrogen volume near the cathode. Inlet electrolyte velocity impacts gas volume fraction distribution, as depicted in **Figures 5** and **6**. At 0.01 m/s, insufficient shear fails to remove bubbles from nickel mesh, increasing residence time and coalescence, leading to higher gas fractions near the cathode, reducing active surface area and hindering reactant replenishment

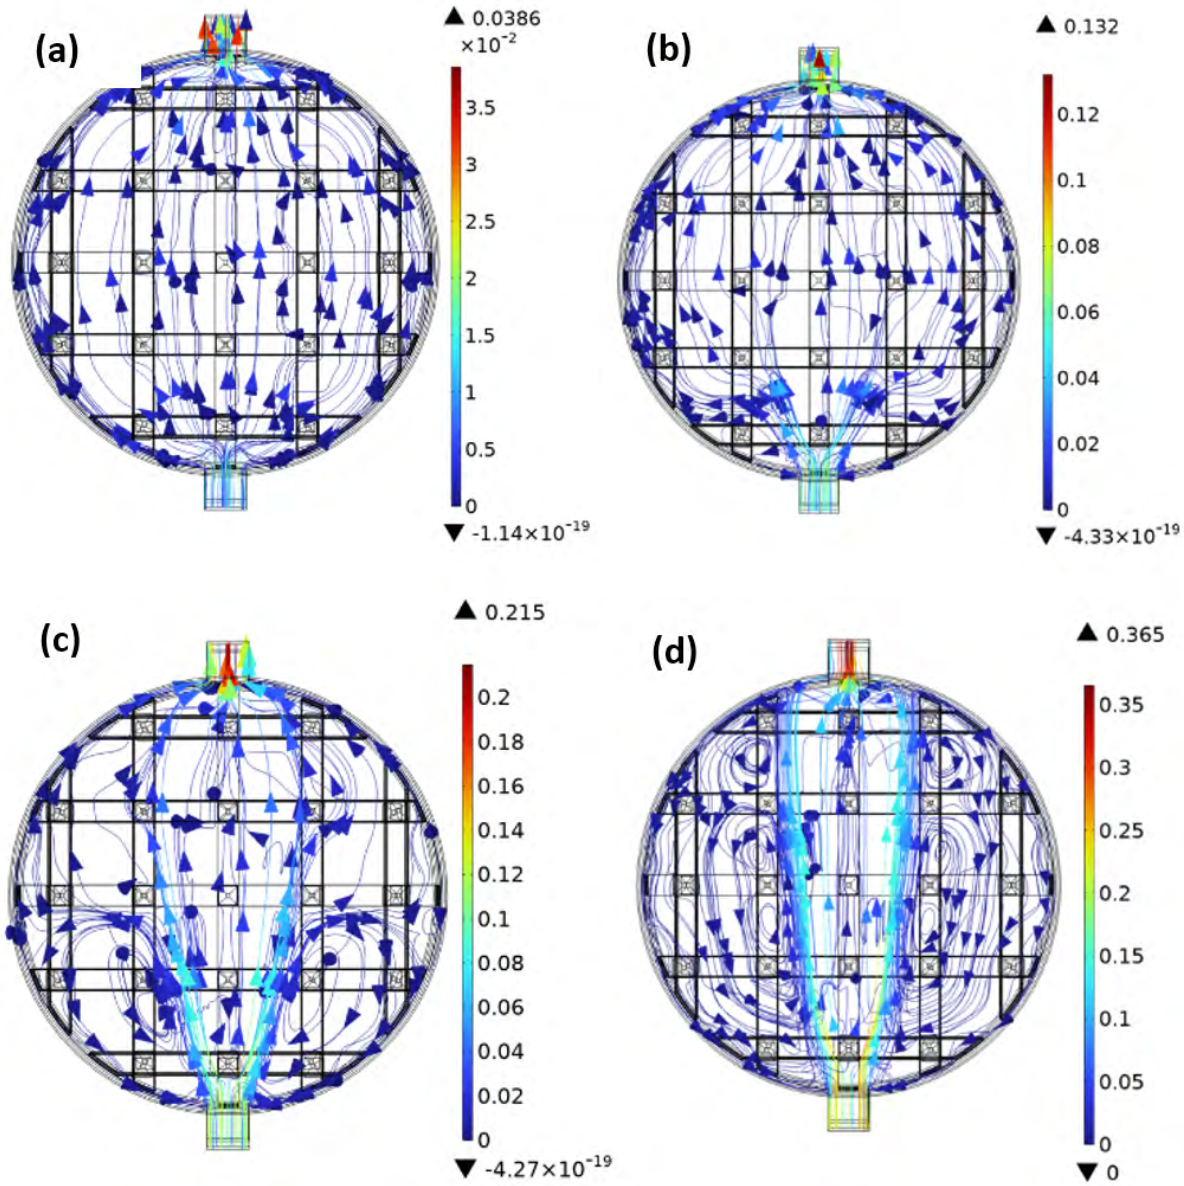

Figure 3. Velocity distribution in channels for 24 pyramidal pins, (a) $v_{in} = 0.01\ m/s$, (b) $v_{in} = 0.05\ m/s$, (c) (a) $v_{in} = 0.1\ m/s$, and (a) $v_{in} = 0.2\ m/s$.

(refer to **Figures 5(a)** and **6(a)**). Increasing inlet velocity to 0.05–0.1 m/s improved hydrogen bubble removal by enhancing convective forces, reducing gas volume fraction near the cathode. This facilitated electrolyte access to nickel sites, minimized gas holdup, and maintained efficient gas-evolving cathode performance without raising channel pressure losses (refer to **Figu res 5(b)-(c)** and **6(b)-4(c)**). At 0.2 m/s, the gas volume fraction was minimized, with rapid bubble detachment reducing interaction time with the nickel surface (refer to **Figures 5(d)** and **6(d)**). This inhibited coalescence, promoting effective mass transfer and bubble removal, though it may require higher pumping power for the system.

The study demonstrated minimal differences in gas volume fraction distributions within channels featuring 10 versus 24 pyramidal pins at varying velocities. While heightened pins number improved fluid tortuosity and localized micro-mixing, it did not significantly influence hydrogen bubble retention or removal on nickel mesh, underscoring the inlet velocity's critical impact on hydrogen gas holdup dynamics.

CONCLUSIONS

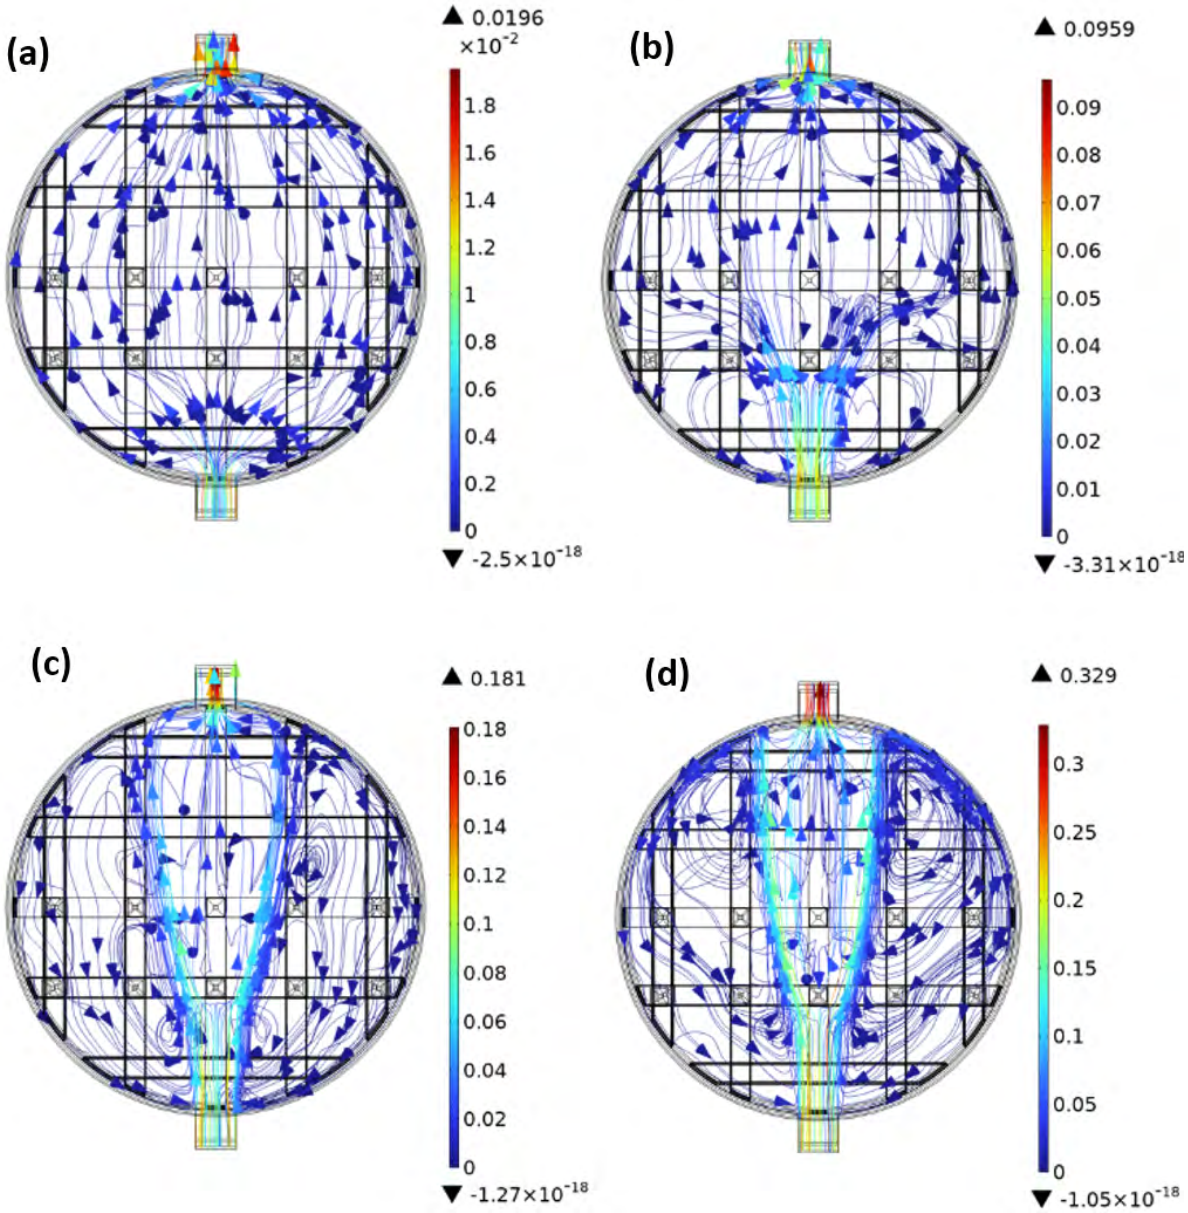

Figure 4. Velocity distribution in channels for 10 pyramidal pins, (a) $v_{in} = 0.01\ m/s$, (b) $v_{in} = 0.05\ m/s$, (c) (a) $v_{in} = 0.1\ m/s$, and (a) $v_{in} = 0.2\ m/s$.

A three-dimensional multiphysics model was developed to optimal design of a portable and stackable AWE for hydrogen production exhibiting two-phase flow dynamics and featuring nickel mesh electrodes, augmented by pyramidal pins to enhance flow disturbance and mixing. The simulations, conducted in COMSOL at a voltage of 1.8 V, 3 M KOH and a temperature of 80°C, were rigorously validated for reliability and accuracy. The findings indicated that the inlet elect rolyte velocity plays a pivotal role in hydrogen bubble detachment and gas holdup; specifically, lower velocities resulted in increased gas fractions due to bubble retention, whereas higher velocities facilitated bubble removal. Although increasing the number of pins improved flow tortuosity and micromixing, it did not significantly influence bubble detachment, suggesting that bubble dynamics are predominantly governed by bubble growth kinetics and convective shear forces. These findings imply that enhancing convective flow in proximity to gas-evolving electrodes can mitigate mass transfer limitations and enhance electrolyzer efficiency. Future research should prioritize the optimization of pyramidal pin design to maximize mixing while minimizing pressure drop challenges, advocating for multi-objective optimization to inform the

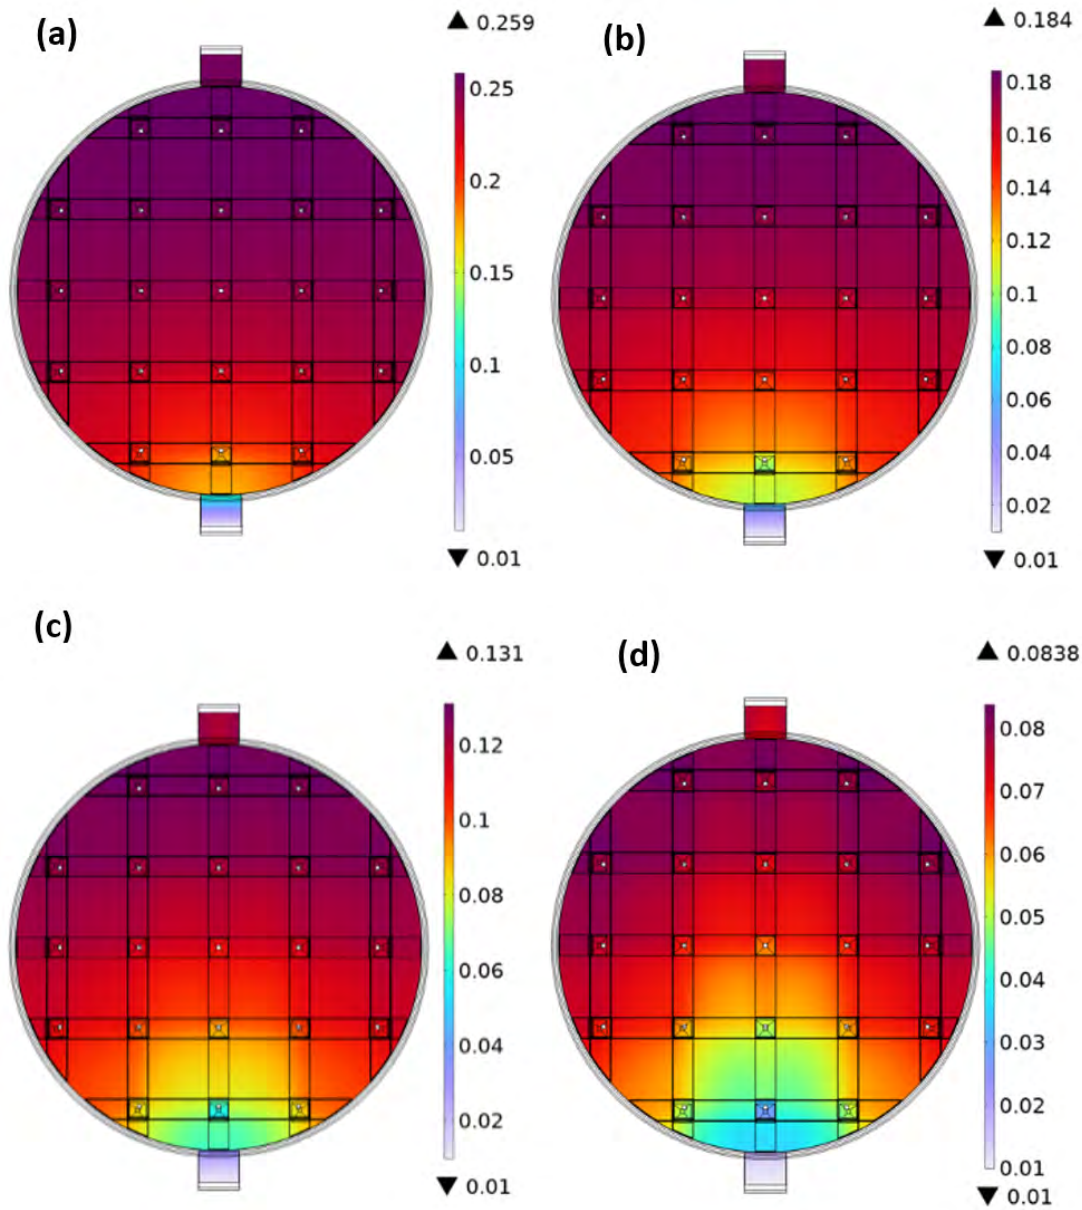

Figure 5. Gas volume fraction distribution in flow channels for 24 pyramidal pins, (a) $v_{in} = 0.01\ m/s$, (b) $v_{in} = 0.05\ m/s$, (c) (a) $v_{in} = 0.1\ m/s$, and (a) $v_{in} = 0.2\ m/s$.

development of efficient designs of next-generation alkaline electrolyzers.

ACKNOWLEDGEMENTS

This project has received funding from the European Union HORIZON Research and Innovation Actions under grant agreement ID 101122323 (UKRI reference No 10106958, and ERC/ HORIZON-HLTH-2021-IND-07, grant No 101057430).

DISCLAIMER

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

AUTHOR IDENTIFIERS

Author ORCIDs:
Venkateshwarlu, Akepogu: 0000-0002-6874-9531
Li Puma, Gianluca: 0000-0002-9168-6284

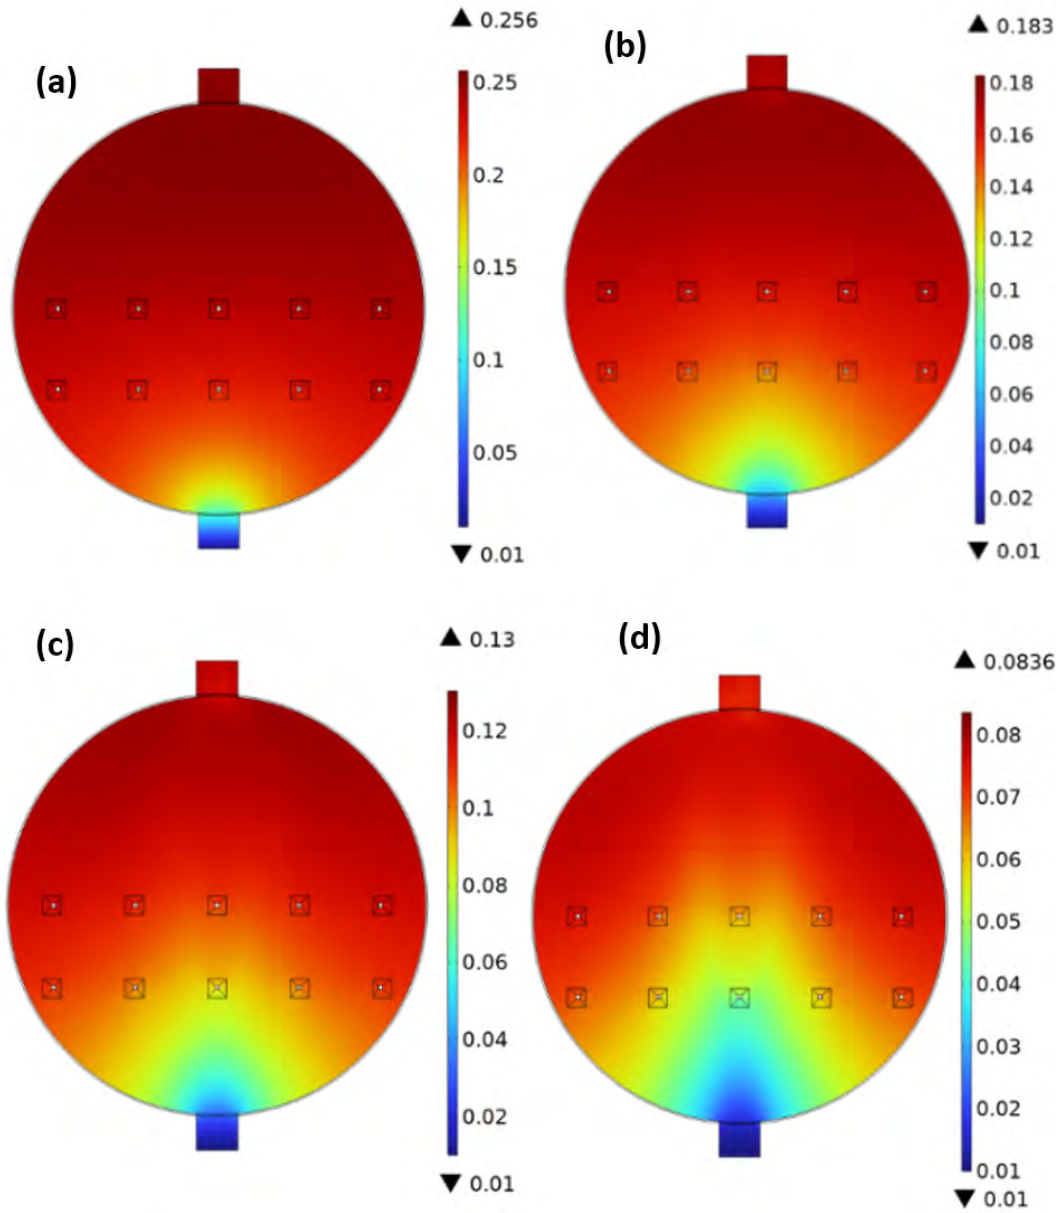

Figure 6. Gas volume fraction distribution in flow channels for 10 pyramidal pins, (a) $v_{in} = 0.01\ m/s$, (b) $v_{in} = 0.05\ m/s$, (c) (a) $v_{in} = 0.1\ m/s$, and (a) $v_{in} = 0.2\ m/s$.

Benyahia, Brahim: 0000-0002-5397-1498

REFERENCES

1. Shin H, Jang D, Lee S, Cho HS, Kim KH, Kang S. Techno-economic evaluation of green hydrogen production with low-temperature water electrolysis technologies directly coupled with renewable power sources. Energy Conversion and Management 286:117083 (2023). https://doi.org/10.1016/j.enconman.2023.117083
2. Riaz MA, Trogadas P, Aymé-Perrot D, Sachs C, Dubouis N, Girault H, Coppens MO. Water electrolysis technologies: the importance of new cell designs and fundamental modelling to guide industrial-scale development. Energy Environ. Sci. 18:5190-5214 (2025). https://doi.org/10.1039/d4ee05559d
3. Risco-Bravo A, Varela C, Bartels J, Zondervan E. From green hydrogen to electricity: a review on recent advances, challenges, and opportunities on power-to-hydrogen-to-power systems. Renewable and Sustainable Energy Reviews 189:113930 (2024). https://doi.org/10.1016/j.rser.2023.113930

4. Zhang W, Liu M, Gu X, Shi Y, Deng Z, Cai N. Water electrolysis toward elevated temperature: advances, challenges and frontiers. Chem. Rev. 123:7119-7192 (2023). https://doi.org/10.1021/acs.chemrev.2c00573
5. Luo S, Zhang T, Xu H, Zhang J, Zhao H, Yun J, Zhao H. Optimizing alkaline water electrolysis: a dual-model approach for enhanced hydrogen production efficiency. Energies 17:5512 (2024). https://doi.org/10.3390/en17215512
6. Zarghami A, Deen NG, Vreman AW. CFD modeling of multiphase flow in an alkaline water electrolyzer. Chemical Engineering Science 227:115926 (2020). https://doi.org/10.1016/j.ces.2020.115926
7. Jang D, Cho HS, Kang S. Numerical modeling and analysis of the effect of pressure on the performance of an alkaline water electrolysis system. Applied Energy 287:116554 (2021). https://doi.org/10.1016/j.apenergy.2021.116554
8. Hu S, Guo B, Ding S, Yang F, Dang J, Liu B, Gu J, Ma J, Ouyang M. A comprehensive review of alkaline water electrolysis mathematical modeling. Applied Energy 327:120099 (2022). https://doi.org/10.1016/j.apenergy.2022.120099
9. Gu H, Wu Q, Zhang M, Xu X, Yu J, Bai Y, Wang Y. Enhancing gas-liquid distribution uniformity in alkaline water electrolyzers by an innovative inlet array structure, a numerical study. Renewable Energy 256:124407 (2026). https://doi.org/10.1016/j.renene.2025.124407
10. Gong Y, Shen N, Zhang B, Xing P, Liu Q, Yu Z, Tan J, Ma X. Gas-liquid flow and electrochemical reactions involved in AEM water electrolytic cells: numerical simulation and parameter analysis. Fuel 409:137925 (2026). https://doi.org/10.1016/j.fuel.2025.137925
11. Hu R, Wen C, Ye Z, Qi Y, Zhang B, Kang K, Gao Y, Wang D, Tu Z. A comprehensive review of flow channel designs and optimizations for water electrolysis technology. Applied Energy 400:126643 (2025). https://doi.org/10.1016/j.apenergy.2025.126643
12. Liang H, Lin S, Zhao K, Liu P, Hu N, Song H, Yang Y, Zheng C, Zhang X, Gao X. Three-dimensional simulation of alkaline electrolyzer with an advanced partitioned point flow channel. International Journal of Hydrogen Energy 86:1252-1261 (2024). https://doi.org/10.1016/j.ijhydene.2024.09.015
13. Wei X, Umehara Y, Nakajima H, Ito K, Etoh A, Mori S. Effect of mesh size of ni wire mesh electrodes on alkaline water electrolysis performance: a study based on the observation of bubble departure behavior. International Journal of Hydrogen Energy 120:189-200 (2025). https://doi.org/10.1016/j.ijhydene.2025.03.267
14. Kimuli EN, Onyemelukwe II, Benyahia B, Rielly CD. Characterisation of axial dispersion in a meso-scale oscillatory baffled crystalliser using a numerical approach. In: Espuña A, Graells M, Puigjaner L, editors. *Computer Aided Chemical Engineering*. Elsevier. 40:223–228 (2017). https://doi.org/10.1016/B978-0-444-63965-3.50039-8
15. Benyahia B, Bandulasena MV, Bandulasena HCH, Vladisavljevi? GT. Experimental and computational analysis of mixing inside droplets for microfluidic fabrication of gold nanoparticles. Ind. Eng. Chem. Res. 60:13967-13978 (2021). https://doi.org/10.1021/acs.iecr.1c01960
16. Courtais A, Lesage F, Privat Y, Pelaingre C, Latifi AM. Cfd-based geometrical shape optimization of a packed-bed reactor combining multi-objective and adjoint system methods. Chemical Engineering Science 275:118728 (2023). https://doi.org/10.1016/j.ces.2023.118728
17. Schalenbach M, Tjarks G, Carmo M, Lueke W, Mueller M, Stolten D. Acidic or alkaline? towards a new perspective on the efficiency of water electrolysis. J. Electrochem. Soc. 163:F3197-F3208 (2016). https://doi.org/10.1149/2.0271611jes

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Modeling and experimental validation of a flat-conduit dense-phase receiver for concentrated solar power

Mustapha Hamdan[ab], Malak Hamdan[ab], Bogdan Dorneanu[a], and Harvey Arellano-Garcia[a]*

[a] Department of Process and Plant Technology, Brandenburg University of Technology Cottbus-Senftenberg, Cottbus, Germany
[b] Heat2X Ltd., London, United Kingdom
* Corresponding Author: arellano@b-tu.de

ABSTRACT

Thermal management and heat transfer optimization remain central challenges in next-generation concentrated solar power (CSP) systems employing solid particles for thermal energy storage and heat transfer. Conventional particle receiver concepts, such as fluidized beds and falling particle curtains are constrained by limited particle-wall contact, flow instabilities, and restricted operating temperature. This work presents a combined computational and experimental investigation of a gravity-driven dense-phase moving packed bed receiver featuring a flat conduit geometry and sub-millimeter particles. A multiphase modeling framework is developed and validated against pressure-drop measurements and particle velocity data obtained from dedicated experimental setups. The validated model is subsequently used to quantify dense-flow stability and thermal performance under indirect heating conditions. Results demonstrate stable dense-phase operation with particle volume fractions of approximately 0.6, leading to enhanced particle-wall transfer. Alumina particles achieve wall-to-particle heat transfer coefficients of up to 2.7 $kW \cdot m^{-2} \cdot K^{-1}$, more than twice those obtained with silica sand under comparable conditions. The findings indicate that flat-conduit dense-phase receivers provide a mechanically simple and thermally efficient alternative to conventional particle receiver concepts, supporting high-temperature and more efficient CSP systems.

Keywords: Energy Storage, Concentrated Thermal Solar, Computational Fluid Dynamics, Modelling and Simulations

INTRODUCTION

Concentrated solar power (CSP) technologies play a key role in the transition toward dispatchable renewable electricity and high-temperature industrial heat. In contrast to photovoltaic systems, CSP offers the intrinsic advantage of thermal energy storage, enabling decoupling of energy collection and power generation. While molten salt systems dominate current deployments, their operating temperature is limited by freezing, corrosion, and material compatibility [1].

Solid particles have therefore emerged as a promising alternative heat transfer and storage medium due to their thermal stability and ability to operate at temperatures exceeding 1000 K [2]. In addition, particles can simultaneously serve as heat transfer fluid and storage medium, offering the potential for simplified system layouts and higher energy storage density. These advantages have motivated extensive research into particle-based receiver concepts for next-generation CSP plants.

Several particle-based receiver concepts have been proposed, including fluidized beds, falling particle curtains, and mechanically driven systems [3]. Fluidized beds achieve high heat transfer coefficients but require large gas flow rates, increasing parasitic losses [4]. Falling particle receivers suffer from low residence times and sensitivity to flow instabilities [5]. Mechanically driven receivers can enhance residence time but introduce moving parts that reduce reliability and scalability at high temperatures [6].

Dense-phase, gravity-driven moving packed beds offer an attractive alternative by combining high particle packing density with mechanically simple operation. Heat transfer in such systems is dominated by particle-wall contact rather than gas-solid convection, making them particularly suitable for indirect heating configurations.

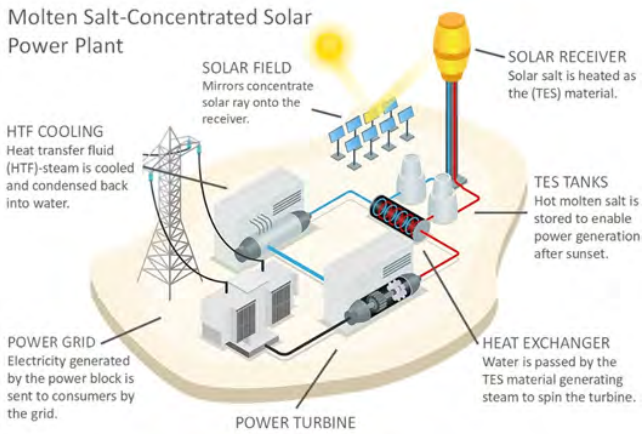

Figure 1. Concentrated solar power (CSP) plant.

Figure 1 illustrates the integration of a particle-based receiver within a CSP system with thermal energy storage (TES). Despite their potential advantages, systematic model-based and experimentally validated studies of dense-phase flat-conduit receiver geometries remain limited.

This work addresses this gap by developing and validating a multiphase modeling framework for a flat-conduit dense-phase particle receiver. The study combines computational fluid dynamics (CFD) modeling with targeted experimental validation to quantify dense-phase flow stability and thermal performance, with particular emphasis on the influence of particle material properties. The results provide insight into the design and operation of mechanically simple, high-performance particle receivers for future CSP applications.

RECEIVER CONCEPT AND OPERATING PRINCIPLE

The investigated receiver consists of a vertical, flat rectangular conduit through which solid particles flow downward under gravity in dense-phase conditions. Heat is supplied indirectly through the receiver wall, representing absorbed solar irradiation transferred to the particles primarily through conduction and contact-based mechanisms. The receiver is intended to operate as part of a particle-based CSP system, where it can be directly coupled to thermal energy storage and high-temperature power conversion units.

As shown schematically in Figure 2, the flat-conduit geometry maximizes particle-wall contact while maintaining a compact and modular design suitable for CSP integration. Unlike fluidized receivers, no external gas flow is required to sustain particle motion. Instead, particles are introduced at the top of the conduit and accelerate under gravity until a balance is reached between gravitational forces, interparticle friction, and wall shear.

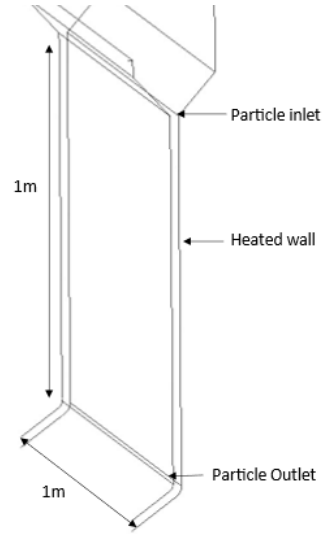

Figure 2. Proposed gravity-driven dense-phase solar receiver with indirect heating configuration.

This balance establishes a stable dense-phase flow regime with high and nearly uniform particle volume fractions across the conduit cross-section.

The flat geometry plays a central role in enhancing heat transfer performance. In dense-phase operation, heat transfer is governed primarily by conductive pathways through the particle network and by repeated particle-wall contact, rather than by convective gas-solid interactions. The planar wall surface increases the effective area available for heat exchange compared to cylindrical or free-falling configurations, while the narrow conduit width limits transverse particle motion and suppresses segregation. As a result, particles experience prolonged and repeated contact with the heated wall over the full receiver height, leading to efficient thermal absorption.

The gravity-driven operating principle further contributes to the robustness and simplicity of the receiver concept. Particle motion is established passively, without mechanical conveyance or fluidization, reducing parasitic energy consumption and mechanical complexity. This passive transport mechanism also decouples hydrodynamic stability from thermal operating conditions, enabling independent adjustment of particle residence time and wall temperature. Such flexibility is particularly advantageous for CSP systems operating under variable solar input, where stable flow behavior is required across a wide range of thermal loads.

MATHEMATICAL FORMULATION

Governing equations

The multiphase flow is modeled using a Eulerian-Eulerian approach [7], in which the gas phase (subscript *g*) and solid particle phase (subscript *s*) are treated as

interpenetrating continua. This approach is well suited for dense granular flows, where particle-particle interactions dominate and tracking individual particles becomes computationally prohibitive.

The local volume fraction of phase *k* is denoted by α_k, its density by ρ_k, and its velocity vector by u_k.

The continuity equation for each phase k is written as:

$$\frac{\partial}{\partial t}(\alpha_k \rho_k) + \nabla \bullet (\alpha_k \rho_k u_k)$$

Where *t* denotes time and $\nabla \bullet$ is the divergence operator.

The momentum balance for the solid phase is given by:

$$\frac{\partial}{\partial t}(\alpha_s \rho_s u_s) + \nabla \bullet (\alpha_s \rho_s u_s) = -\alpha_s \nabla p + \nabla \bullet \tau_s + \alpha_s \rho_s g + F_{sg}$$

Here, p is the pressure shared by both phases, τ_s represents the solid-phase stress tensor derived from the kinetic theory of granular flow, g is the gravitational acceleration vector, and F_{sg} denotes the interphase momentum exchange force, modeled using the Syamlal-O'Brien correlation [7].

The energy balance for the solid phase is expressed as:

$$\frac{\partial}{\partial t}(\alpha_s \rho_s h_s) + \nabla \bullet (\alpha_s \rho_s u_s h_s) = \nabla \bullet \left(k_{eff} \nabla T_s\right) + h_{ws} A_{ws}(T_w - T_s)$$

Where h_s is the specific enthalpy of the solid phase, T_s and T_w are the solid and wall temperatures, k_{eff} is the effective thermal conductivity of the packed bed, h_{ws} is the wall-to-particle heat transfer coefficient, and A_{ws} is the specific wall-particle contact area.

Wall-to-particle heat transfer is modeled using the Gunn correlation [8], and the effective thermal conductivity of the packed bed is computed using the Zehner-Bauer-Schlünder formulation [9].

Radiative heat transfer is not included explicitly in the present formulation. This assumption is justified by the moderate temperature range considered and by the objective of isolating conductive and contact-based heat transfer mechanisms in dense-phase operation. While radiation is expected to contribute at higher temperatures, its exclusion allows a clearer assessment of material and hydrodynamic effects and does not affect the comparative analysis between particle materials presented in this study.

The mathematical formulation captures the dominant physical mechanisms governing dense-phase flow and heat transfer in the flat-conduit receiver. By combining granular flow theory with established packed-bed heat transfer correlations, the model provides a physically consistent basis for evaluating receiver performance while remaining computationally tractable for parametric studies.

Table 1: Material properties.

Material	Density ($kg \cdot m^{-3}$)	Specific heat capacity ($J \cdot kg^{-1} \cdot K^{-1}$)
Air	1.225	1006
Alumina	3960	1250
Silica sand	2650	755

All simulations are performed using ANSYS Fluent. Transient simulations are performed to capture dense-phase dynamics. Grid-independence and time-step sensitivity analyses are conducted to ensure numerical stability. Indirect heating is imposed via prescribed wall temperature boundary conditions representative of CSP operation. Table 1 summarizes the thermophysical properties of silica sand and alumina particles used in the simulations. Particles had a size of 500 µm.

EXPERIMENTAL VALIDATION STRATEGY

The experimental validation strategy is designed to assess the predictive capability of the numerical model under operating conditions representative of gravity-driven dense-phase receiver operation. Rather than attempting full thermal validation at this stage, the focus is placed on validating the hydrodynamic behavior of the particle phase, which is a prerequisite for reliable heat transfer predictions. A transparent flat-conduit receiver prototype is used to obtain particle velocity measurements under controlled operating conditions.

High-speed imaging combined with particle tracking techniques is used to obtain local particle velocities at multiple axial locations along the receiver height. These measurements provide spatially resolved information on particle kinematics, allowing direct comparison with simulated velocity profiles. The experimental setup and measurement planes are shown in Figure 3a.

The validation focuses on particle velocity rather than temperature, as velocity directly governs residence time, particle-wall contact duration, and ultimately heat transfer performance. By ensuring that the model accurately reproduces dense-phase kinematics, confidence is established in its ability to predict thermal behavior when coupled with established heat transfer correlations.

Comparison between experimental and numerical results (Figure 4) shows good agreement, with average relative deviations on the order of 6%. Slight discrepancies are observed near the inlet region, where experimental uncertainties related to particle injection and optical tracking are highest. These deviations do not propagate downstream and do not affect the overall flow regime or residence time distribution.

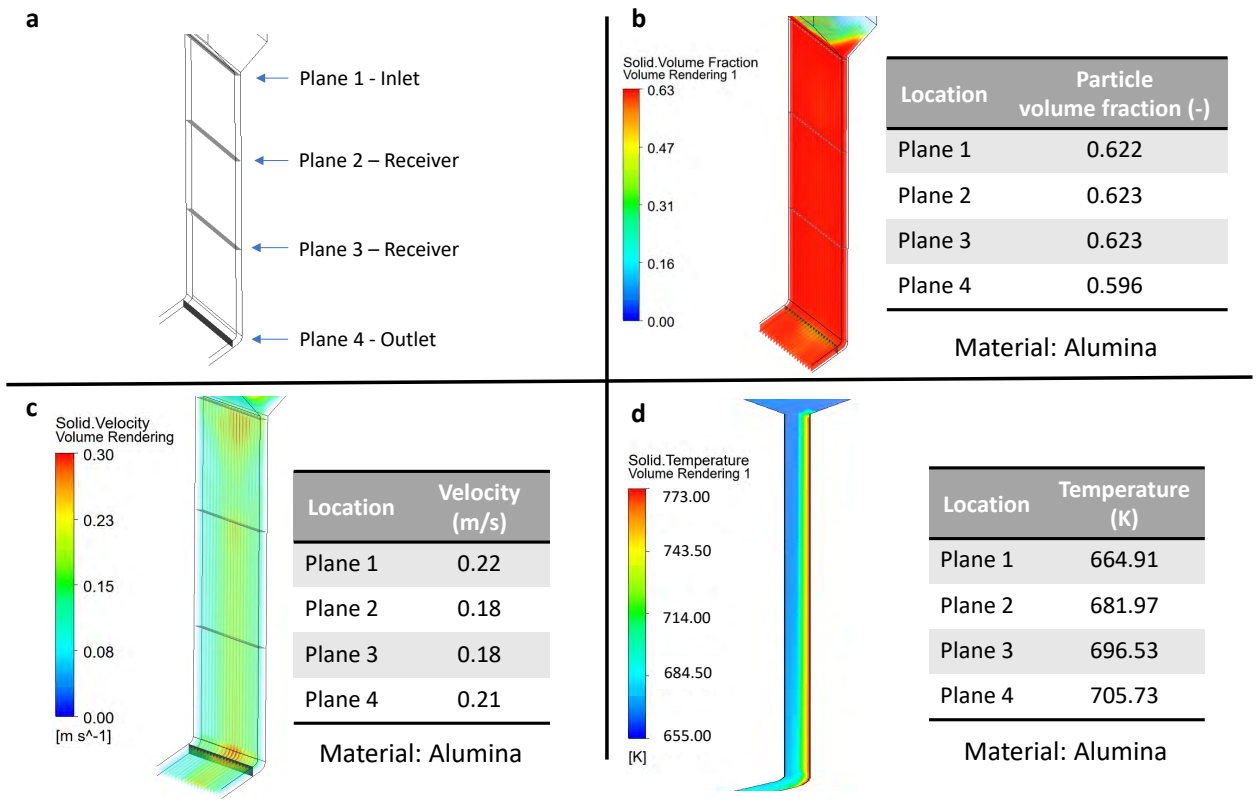

Figure 3. Simulated results at *t* = 15s showing: a) Measurement points along the receiver; b) Particle volume fractions; c) Particle velocity in the direction of gravity; and d) Particle temperature profile.

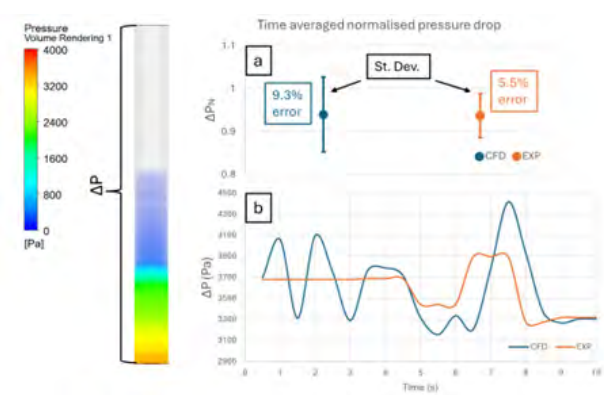

Figure 4. a) Comparison between CFD model (blue) and experimental (orange) results for the normalised bed pressure drop; b) Measurements of the pressure drop comparing experimental and CFD results.

This validation strategy reflects a deliberate modeling hierarchy. Hydrodynamic validation is performed first to confirm correct representation of dense-phase flow behavior. Thermal predictions are then evaluated numerically using well-established correlations for wall-to-particle heat transfer and effective packed-bed conductivity. This approach avoids conflating hydrodynamic modeling errors with uncertainties associated with high-temperature experimentation and radiative heat transfer.

The experimental validation confirms that the numerical framework captures the key dense-phase flow characteristics relevant to flat-conduit receiver operation. The validated model therefore provides a robust basis for assessing thermal performance and for supporting future experimental campaigns at elevated temperatures.

RESULTS AND DISCUSSION

Dense-phase flow characteristics

Simulations of the flat-conduit receiver show stable dense-phase flow with solid volume fractions in the range of 0.62-0.63. As illustrated in Figure 3c, particle velocities remain uniform across the conduit width, with no evidence of channeling or dead zones.

This homogeneous flow structure is a direct consequence of the flat geometry and gravity-driven operation, which promote continuous particle-wall contact and suppress lateral segregation. In contrast to falling particle receivers, where particle dilution and curtain instabilities frequently lead to non-uniform residence times, the dense-phase configuration ensures predictable particle trajectories and well-defined residence times. From a receiver design perspective, this stability is essential for

maintaining consistent thermal performance under fluctuating solar input.

Moreover, the absence of localized regions with reduced solid fraction indicates that heat transfer is not limited by intermittent particle contact, which is a common drawback in dilute-phase systems. These results confirm that the proposed flat-conduit configuration provides a robust hydrodynamic basis for indirect particle heating.

Thermal performance and material effects

Temperature fields and absorbed thermal power are summarized in Figure 5. For both particle materials, a clear temperature gradient develops between the heated wall and the particle core, indicating wall-dominated heat transfer under dense-phase conditions. However, alumina particles consistently reach higher temperatures throughout the conduit height compared to silica sand.

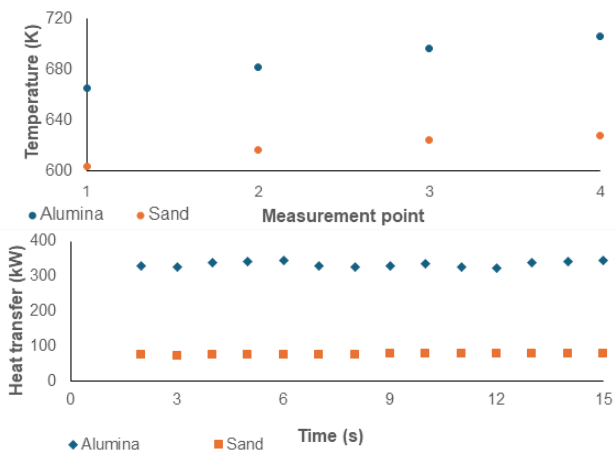

Figure 5. a) Particle temperature measurements at t = 15 s; b) Transient heat transfer rate from the receiver wall, comparing sand an alumina.

This behavior can be attributed to the higher density and thermal conductivity of alumina, which enhance both the effective thermal conductivity of the packed bed and the heat capacity flow rate of the solid phase. As a result, alumina particles are able to absorb and transport more thermal energy for a given residence time.

Quantitatively, alumina particles absorb approximately 334 kW of thermal power at a mass flow rate of 11.7 kg·s^{-1}, whereas silica sand absorbs about 168 kW at 9.5 kg·s^{-1} under identical conditions. These results demonstrate that material selection alone can nearly double thermal performance in dense-phase receivers.

From a system-level perspective, these results suggest that dense-phase receivers offer significant flexibility in performance tuning through material choice, enabling adaptation to different operating temperature ranges or storage requirements without major design modifications.

Wall-to-particle heat transfer coefficients

The resulting wall-to-particle transfer coefficients are shown in Figure 6. Alumina achieves peak values of approximately 2.7 kW·m^{-2}·K^{-1}, while silica sand reaches only 1.2 – 1.4 kW·m^{-2}·K^{-1}.

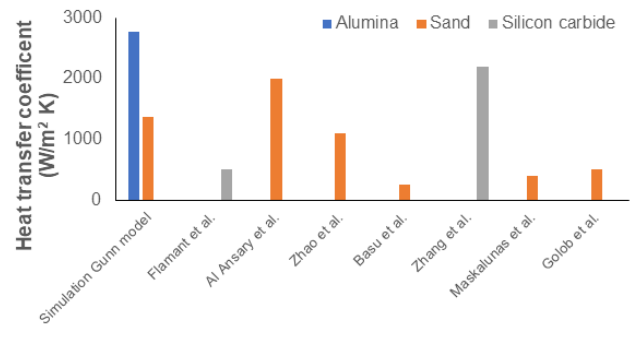

Figure 6. Wall-to-solid heat transfer coefficient from the receiver wall to the gas solid phase, comparing experimental results from literature.

These values compare favorably with experimental data reported in the literature for fluidized and falling particle receivers, while avoiding the need for auxiliary fluidization and its associated parasitic losses. The enhanced heat transfer observed in the dense-phase receiver is primarily driven by sustained particle-wall contact and high packing density, rather than by increased gas-solid convection.

The results also indicate diminishing marginal gains in heat transfer at very high packing densities, where increased particle-particle contact resistance can partially offset improvements in wall contact. This highlights the importance of balancing particle size, material properties, and conduit geometry when designing dense-phase receivers to achieve optimal thermal performance.

The heat transfer coefficients achieved in the present configuration demonstrate that dense-phase flat-conduit receivers can compete with, and in some cases outperform, more complex particle receiver concepts while maintaining mechanical simplicity.

CONCLUSIONS AND OUTLOOK

A flat-conduit dense-phase particle receiver for concentrated solar power applications has been investigated using a combined computational and experimental approach. A multiphase Eulerian-Eulerian model was developed to describe dense-phase hydrodynamics and particle-wall heat transfer and was validated against particle velocity measurements obtained from a transparent receiver prototype. The validation confirms that the model reliably captures gravity-driven dense-phase flow behavior relevant to the proposed receiver concept.

The results show that alumina particles significantly

outperform silica sand, achieving wall-to-particle heat transfer coefficients of up to 2.7 $kW \cdot m^{-2} \cdot K^{-1}$ and nearly doubling the absorbed thermal power under comparable operating conditions. These findings underline the critical role of particle material properties in dense-phase receiver performance and demonstrate that substantial efficiency gains can be achieved through material selection alone.

Beyond quantitative performance metrics, the study highlights the advantages of flat-conduit dense-phase receivers in terms of flow stability, predictability of residence time, and mechanical simplicity. These characteristics make the concept particularly attractive for integration with particle-based thermal energy storage systems and high-temperature CSP architectures.

The present work is subject to certain limitations that define its scope. Thermal predictions are based on conductive and contact-dominated heat transfer mechanisms under indirect heating conditions, and radiative heat transfer effects are not considered explicitly. While this assumption is appropriate for the objectives of the current study and enables a clear assessment of material and hydrodynamic effects, it also motivates further investigation at higher temperatures and under direct solar irradiation

Future work will therefore focus on extending the experimental validation to elevated temperature operation, incorporating radiative heat transfer into the numerical framework, and assessing long-term material behavior under cyclic thermal loading. In addition, scale-up studies and modular receiver integration will be investigated to evaluate system-level performance and economic feasibility. Coupling the validated model with optimization and control methodologies will further support the design of robust dense-phase receivers capable of operating under transient solar conditions.

AUTHOR IDENTIFIERS

Hamdan Malak: 0000-0001-8104-4539
Dorneanu B: 0000-0003-3553-6625
Arellano-Garcia H: 0000-0002-8297-0232

REFERENCES

1. Zhang HL, Baeyens J, Degrève J, Cacères G. Concentrated solar power plants: review and design methodology. Renewable and Sustainable Energy Reviews 22:466-481 (2013). https://doi.org/10.1016/j.rser.2013.01.032
2. Fernández AG, Ushak S, Galleguillos H, Pérez FJ. Thermal characterisation of an innovative quaternary molten nitrate mixture for energy storage in CSP plants. Solar Energy Materials and Solar Cells 132:172-177 (2015). https://doi.org/10.1016/j.solmat.2014.08.020
3. Al-Ansary H, El-Leathy A, Al-Suhaibani Z, Jeter S, Sadowski D, Alrished A, Golob M. Experimental study of a sand–air heat exchanger for use with a high-temperature solar gas turbine system. Journal of Solar Energy Engineering 134: (2012). https://doi.org/10.1115/1.4007585
4. Salatino P, Ammendola P, Bareschino P, Chirone R, Solimene R. Improving the thermal performance of fluidized beds for concentrated solar power and thermal energy storage. Powder Technology 290:97-101 (2016). https://doi.org/10.1016/j.powtec.2015.07.036
5. Imran Khan M, Asfand F, Al-Ghamdi SG. Progress in technology advancements for next generation concentrated solar power using solid particle receivers. Sustainable Energy Technologies and Assessments 54:102813 (2022). https://doi.org/10.1016/j.seta.2022.102813
6. Ho CK, Iverson BD. Review of high-temperature central receiver designs for concentrating solar power. Renewable and Sustainable Energy Reviews 29:835-846 (2014). https://doi.org/10.1016/j.rser.2013.08.099
7. Syamlal M, O'Brien TJ. Computer Simulation of Bubbles in a Fluidized Bed. *AIChE Symp Series* 82:22-31 (1989).
8. Gunn DJ. Transfer of heat or mass to particles in fixed and fluidised beds. International Journal of Heat and Mass Transfer 21:467-476 (1978). https://doi.org/10.1016/0017-9310(78)90080-7
9. Tsotsas E, Martin H. Thermal conductivity of packed beds: a review. Chemical Engineering and Processing: Process Intensification 22:19-37 (1987). https://doi.org/10.1016/0255-2701(87)80025-9

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Section 3

CAPEing with Uncertain Futures

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

The Value of Multi-Stage Stochastic Programming in Power Grid Capacity Expansion Planning

Sergio Bugosen[a], Tomas Valencia[b], Jean-Paul Watson[b], Chrysanthos E. Gounaris[a], and Carl D. Laird[a]*

[a] Carnegie Mellon University, Department of Chemical Engineering, Pittsburgh, PA, USA
[b] Lawrence Livermore National Laboratory, Livermore, CA, USA
* Corresponding Author: claird@andrew.cmu.edu.

ABSTRACT

This work develops a high-spatial resolution multi-stage stochastic programming (MS) model for power grid capacity expansion that co-optimizes generation, transmission, and energy storage system investments under uncertainty. Traditional two-stage stochastic programming (TS) models determine all investments in a single stage, limiting their ability to adapt to changing conditions such as evolving capital costs, policies, or supply chain disruptions. In contrast, the proposed MS formulation introduces sequential decision stages where partial information is revealed over time, allowing for adaptive, scenario-contingent investments. We compare TS and MS formulations using a modified IEEE 24-bus case study to quantify the Value of the Multistage Solution, which measures the economic benefit of allowing investment decisions to adapt over time as uncertainty is progressively resolved. Results show that while MS models are computationally more challenging, they achieve lower expected costs and yield flexible investment strategies that better respond to geopolitical and market uncertainties.

Keywords: Stochastic Programming, Capacity Expansion Planning, Power Systems Modeling, Mixed-Integer Linear Programming, Generation, Transmission, Energy Storage Systems

INTRODUCTION

Electric power systems require long lead times for the development of new infrastructure. Securing regulatory approvals, acquiring land, negotiating fuel supply contracts, and building new generation or transmission assets can take many years to a decade [1]. Planning entities must make critical investment decisions before the full trajectory of technological, policy, and market conditions is known. The challenge lies in acting under uncertainty: delaying commitments until perfect information arrives is untenable, since long construction lead times require proactive expansion to maintain adequate capacity and reliably satisfy electricity demand. This uncertainty is characterized by volatile economic conditions, the rapid growth of energy-intensive data centers, and the increasing penetration of renewable energy sources. Neglecting uncertainty in planning models can lead to stranded assets or a power grid that is vulnerable to blackouts [2].

To inform these long-term investment decisions, planners can rely on capacity expansion planning (CEP) models. CEP models provide optimal investment decisions such as the size, timing, geographic location, and type of technology of new generation, transmission, and storage needed to meet future electricity demand at a minimum capital and operating cost [3]. Traditional formulations often adopt a deterministic approach, in which the expansion of the power grid is planned for a single assumed future where parameters such as fuel prices, capital costs, and electricity demand are treated as known. A wide variety of research papers are available that solve CEP deterministically [4]. However, these plans are likely to fail if implemented, as the realization of these parameters may be different from what the planner initially expected [5].

Stochastic Programming (SP) can hedge against certain outcomes by explicitly embedding multiple future realizations in the decision model. The model will produce plans that remain robust even when reality diverges from any single forecast [6]. Among SP formulations, two-stage stochastic programming (TS) has become the standard approach for power grid expansion planning. Here, first-stage investment decisions are made before

uncertainty is realized, while second-stage recourse decisions adjust operations after uncertainty resolution. This yields a block-angular structure that can be solved using decomposition algorithms such as Progressive Hedging and Benders Decomposition [7, 8].

However, all investment decisions are made upfront in TS formulations, preventing the model, by definition, from adapting to information revealed in later stages. This limitation is critical in contexts where policies evolve over time, potentially leading to premature or suboptimal investments. To overcome this limitation, multi-stage stochastic programming (MS) introduces sequential decision stages that unfold over time. Partial information is revealed at each stage, and nonanticipativity constraints ensure consistency across scenarios that share the same history. This structure captures the dynamic nature of long-term planning, allowing decision-makers to defer and revise investments as uncertainties materialize [9].

In this work, we develop a high-spatial resolution MS stochastic programming model for power grid expansion that co-optimizes generation, transmission, and energy storage system (ESS) investments under geopolitical uncertainty. We then compare the TS and MS formulations to quantify the Value of the Multistage Solution (VMS) [10] on a modified IEEE 24-bus case study. Unlike the Value of the Stochastic Solution (VSS), the VMS does not measure the benefit of modeling uncertainty relative to a deterministic model. Instead, it quantifies the value of adaptability by allowing decisions to adjust over time as uncertainty is revealed. To our knowledge, this is the first study to report this metric in the context of power grid CEP, thus providing the basis for future research on decision-making under uncertainty.

BACKGROUND

A wide range of multistage stochastic CEP models has been proposed in the literature, each designed to address specific problem statements, system characteristics, and computational trade-offs. These models vary widely in the level of network granularity (regional vs. nodal), the treatment of investment decisions (continuous vs. discrete, with or without lead times), how operations are modeled (with or without unit commitment constraints), and the types of uncertainty (renewable power availability, loads, fuel prices, capital costs).

Liu et al. (2018) developed one of the first multistage, multiscale stochastic CEP models that co-optimizes generation, transmission, and ESS investments [11]. Their formulation is linear and highly aggregated, relying on a pipeline representation of transmission and continuous investment variables. These modeling choices allow the approach to scale well but limit its ability to represent discrete investment decisions, physical power-flow constraints, and delayed asset availability.

Thomé et al. (2019) propose an integrated generation–transmission planning framework that couples a MILP investment module with an SDDP-based multistage operational model [12]. Their formulation includes binary investment variables, power-flow physics, and detailed hydrothermal dispatch. However, the model does not represent construction lead times or ESS, factors that are becoming increasingly important in modern systems.

Hou et al. (2021) develop a multistage stochastic MILP linking long-term planning with hourly unit-commitment operations [13]. Their model co-optimizes generation and transmission expansion, but uses a six-region aggregated network without power-flow physics, continuous generator capacity additions, and no ESS or construction delays. In addition, uncertainty is limited to renewable availability, despite the important role of fuel prices, capital costs, and demand growth in CEP.

Apablaza et al. (2024) introduce a multistage stochastic model that co-optimizes transmission and energy storage investments using a nodal power-flow representation and explicit construction lead times [14]. Uncertainties are considered in fuel prices, load growth, and renewable availability. Because generator capacities are fixed, the model does not consider new generation investments and therefore cannot capture interactions between uncertain long-term conditions and generator siting, sizing, or technology choice.

Hole et al. (2025) present a stochastic planning model for New Zealand's hydro-dominated system that evaluates investments in wind and HVDC capacity using an SDDP-based operational framework [15]. The system is represented through an aggregated 11-node network with major corridors and key reservoirs, which supports large-scale hydrothermal modeling but omits finer spatial detail and storage technologies such as batteries.

These studies illustrate the breadth of modeling assumptions adopted in multistage stochastic CEP. Nonetheless, recent evidence highlights where model fidelity matters: Glista et al. (2025) show that low-spatial-resolution CEP models produce large errors and substantial generation under-investment relative to full-resolution formulations [16], while Watson et al. (2015) demonstrate that unit-commitment constraints have only limited influence on long-term system costs [2]. These insights help guide the modeling priorities of this work.

Our formulation is a nodal multistage stochastic MILP with the DC power flow approximation, discrete investment decisions for thermal generators and transmission lines, continuous sizing for renewable and ESS investments, and explicit construction lead times across stages. This design preserves the spatial accuracy needed for reliable expansion outcomes while avoiding unnecessary operational complexity, enabling the model to capture the spatial and temporal features that shape realistic long-term planning under uncertainty.

MODEL DEVELOPMENT

In this section, we develop a multi-stage stochastic programming model for power grid capacity expansion planning. We begin by describing the deterministic optimization model and subsequently extend it to MS and TS programs.

Deterministic program

The model described in this section is a multi-stage extension to the model described in [17]. For brevity, the detailed formulation, as well as notation for sets, parameters, and variables, is shown in the Supplementary Information. We refer to equation numbers in the Supplementary Information for completeness.

The objective function shown in Eq. (1s) minimizes the expected net present value (NPV) of total system costs over a planning horizon, defined as the sum of (i) capital investments in each investment stage for new generators, transmission lines, and ESS and (ii) the probability-weighted operational cost for all representative operating days. The operating component is comprised of fixed operation & maintenance (O&M) costs, variable O&M costs plus fuel expenses for dispatched generation, and a penalty for shortfalls in meeting the Renewable Portfolio Standard (RPS). In addition, each MWh of unserved energy is penalized with the Value of Lost Load (VOLL) – an economic proxy for the outage cost borne by customers. This penalty represents the loss of social welfare from electrical blackouts, and some system operators value it at between 10 and 30 thousand $/MWh in planning models.

Inter-stage linking is enforced through capacity accumulation constraints shown in Eqs. (2s) – (5s), which define installed capacity in each stage as the sum of previously available capacity and new investments, ensuring irreversibility of investment decisions. Consistent with this assumption, power plant retirements are not considered in this model.

Generation dispatch limits, shown in Eqs. (6s) – (8s), constrain power produced by newly built thermal and renewable generators by capacity installed in the previous stage to reflect construction lead times, with renewable output further limited by representative-day availability profiles. Existing generators are bounded by fixed nameplate capacities.

Power flow limits, given in Eqs. (9s) – (11s), bound flows on existing transmission lines by their thermal limits and enforce construction lead times by restricting flows on candidate lines according to capacity installed in the previous stage.

Power flow constraints, given in Eqs. (12s) – (14s), represent power-flow physics using a DC approximation of the AC power-flow equations, also known as the Bθ approximation, in which line flows are modeled as linear functions of phase angle differences. For candidate transmission lines, a disjunction implemented via big-M constraints enforces the DC power-flow relationship only when the line is available; in particular, the formulation links line availability to the prior-stage build decision, consistent with the assumed construction lead time.

ESS constraints, given in Eqs. (15s) – (19s), limit the state of charge by installed energy capacity and bound charging and discharging rates by power conversion capacity. The state of charge evolves according to net charging and discharging decisions, accounting for efficiency losses, with a cyclic condition imposed for each representative day. Unlike generation and transmission investments, ESS capacity is assumed to be immediately available upon installation, and therefore not subject to construction lead-time. This assumption is supported by the short project development timelines of modular battery technologies [18], which are negligible at the temporal resolution of the investment stages considered in this study.

The nodal balance constraint, given in Eq. (20s), equates the transmission inflows and local injections (generation and storage discharge) to transmission outflows and withdrawals (demand and storage charging), with load shed included to maintain feasibility under all operating conditions.

The RPS constraint, given in Eq. (21s), establishes a lower bound on the total power dispatch that must originate from renewable sources. Finally, via Eq. (22s) – (23s), the model establishes a limit on the number of new builds allowed on each bus due to physical space and interconnection capacity.

The formulation shown in the Supplementary Information can be compactly represented as

$$\min \sum_{k=1}^{K} (\alpha_k{}^{\top} u_k + \beta_k{}^{\top} y_k) \tag{1}$$

$$z_k = z_{k-1} + u_k \; \forall\, k = 1, \dots, K$$

$$B_k y_k \le b_k + G_k z_{k-1} \; \forall\, k = 2, \dots, K$$

$$E_k y_k \le r_k + T_k z_k \; \forall\, k = 1, \dots, K$$

$$Q_k y_k \le q_k \; \forall\, k = 1, \dots, K$$

$$D_k y_k = d_k \; \forall\, k = 1, \dots, K$$

$$0 \le u_k \le \bar{u}_k \; \forall\, k = 1, \dots, K$$

$$z_0 = \bar{z}$$

where u_k denotes the vector of investment decisions made in stage k, with elements representing the number of new thermal generators and candidate lines, as well as the capacity of new renewable generators and ESS. The vector z_k denotes the corresponding cumulative installed capacity available at the beginning of stage k, with z_0 denoting the initial system capacity prior to any new investments. The vector y_k collects all operational

variables in stage k, including generation dispatch, power flows, phase angles, ESS charging and discharging, and slack variables such as unserved energy and RPS non-compliance. The one-stage construction lead time is enforced by applying capacity-dependent operational limits using the lagged state z_{k-1}; specifically, the inequality block with matrices B_k and G_k ensures that capacity built in stage k cannot be used for dispatch or candidate line operation until stage $k+1$. This lead time is not enforced for ESS constraints, as shown in the inequality block with matrices E_k and T_k. The subsequent equality and inequality constraints represent constraints such as nodal balances, state-of-charge of ESS, and other bounds on the operational variables y_k and stage-wise build limits u_k. Finally, the vectors α_k and β_k contain investment and operational cost coefficients applied to u_k and y_k, respectively.

Multi-stage stochastic program

To capture the value of sequential decision-making under uncertainty, we now present the MS formulation in which investment decisions are allowed to adapt over time as uncertainty is gradually revealed.

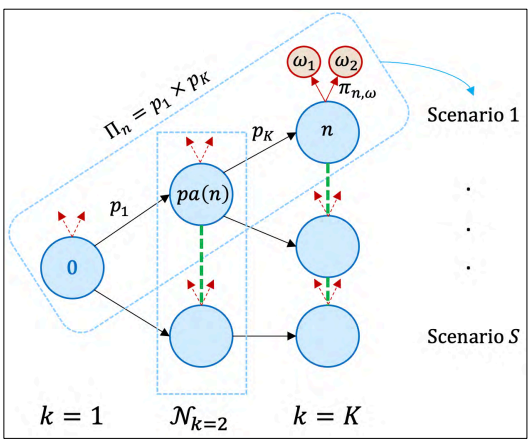

Figure 1. The scenario tree used in both MS and TS formulations. The TS model enforces stage-wise non-anticipativity of investment decisions (green dashed line), whereas the MS model allows node-dependent investments.

Uncertainty is represented by a scenario tree, as shown in Figure 1, whose nodes correspond to realizations of higher-level structural uncertainty, such as geopolitical events that affect long-term investment decisions. Let $\mathcal{N}$ denote the set of investment nodes of the scenario tree, with root node $n = 0$, and let $\mathcal{N}_k$ denote the set of investment nodes at stage k. Each node n has an associated probability Π_n, and $pa(n)$ denotes its parent node. Investment decisions u_n are made at each node and are therefore adaptive to the information revealed up to that node; nonanticipativity is enforced implicitly by the tree structure.

Short-term operational uncertainty such as fuel prices and load and renewable power availability profiles is represented through a finite set of operational realizations Ω_n with associated probabilities $\{\pi_{n,\omega}\}_{\omega\in\Omega_n}$. Operational decisions $y_{n,\omega}$ are variables that adapt to the variability within each node. The scenario tree captures strategic uncertainty that drives investment adaptability, while Ω_n captures operational variability conditional on each investment node. The formulation is as follows:

$$\min \sum_{n\in\mathcal{N}} \Pi_n \left(\alpha_n{}^{\top} u_n + \sum_{\omega\in\Omega_n} \pi_{n,\omega} \beta_{n,\omega}{}^{\top} y_{n,\omega} \right) \quad (2)$$

$$z_n = z_{pa(n)} + u_n \quad \forall\, n \in \mathcal{N}$$

$$B_n y_{n,\omega} \le b_{n,\omega} + G_n z_{pa(n)} \;\forall\, n \in \mathcal{N};\ \omega \in \Omega_n$$

$$E_n y_{n,\omega} \le r_{n,\omega} + T_n z_n \;\forall\, n \in \mathcal{N};\ \omega \in \Omega_n$$

$$Q_n y_{n,\omega} \le q_{n,\omega} \quad \forall\, n \in \mathcal{N};\ \omega \in \Omega_n$$

$$D_n y_{n,\omega} = d_{n,\omega} \quad \forall\, n \in \mathcal{N};\ \omega \in \Omega_n$$

$$0 \le u_n \le \bar{u}_n \qquad \forall\, n \in \mathcal{N}$$

$$z_0 = \bar{z}$$

Two-stage stochastic program

The TS model is obtained as a special case of the MS model by restricting investment decisions to be stage-wise non-anticipative. While the strategic scenario tree and the operational recourse structure is identical to that of the MS model, the TS model enforces that all investment decisions taken at the same investment stage are identical across strategic nodes. As a result, investment decisions do not adapt to the realized strategic uncertainty and are selected in a here-and-now fashion. In other words, the TS model does not allow any flexibility in the investment plan with respect to the realized state of the world.

Formally, let $\mathcal{N}_k \subset \mathcal{N}$ denote the set of investment nodes associated with investment stage k. The TS model imposes the following stage-wise non-anticipativity constraint:

$$u_n = u_m \qquad \forall\, n, m \in \mathcal{N}_k;\ k \in K \quad (3)$$

Equivalently, a single stage-level investment decision u_k may be introduced for each stage k.

$$u_n = u_k \qquad \forall\, n \in \mathcal{N}_k \quad (4)$$

Under this restriction, the complete investment plan $\{u_k\}_{k=1}^{K}$ is identical across all strategic scenarios, while operational decisions remain adaptive and are treated as recourse variables conditional on the realized operational uncertainty, such as fuel prices and load and renewable power availability profiles.

Table 1: Summary of scenarios considered in this case study.

Stage (k)	Narrative	Parameter materialization	Node		
2	Policies that favor rapid deployment of low-carbon technologies such as SMR, solar panels, and wind turbines. Reduced demand for fossil-fuel-based generation.	Fuel cost ↓	2		
	A disruption in renewable energy supply chain that shifts policy support toward fossil-fuel-based generation. Increased demand for fossil-fuel-based generation.	Fuel cost ↑↑	3		
	Broad advancements across all generation technologies. Moderately high fossil fuel costs, given no strong shift in generation mix.	Fuel cost ↑ Investment costs ↓5%	4		
3	Moderate demand growth; tamed datacenter growth	Demand ↑ (1.6x)	5	7	9
	High demand growth driven by AI datacenter proliferation	Demand ↑↑ (2x)	6	8	10

PROBLEM STATEMENT

We study a long-term power grid CEP under uncertainty to determine optimal generation, transmission, and ESS investments, and associated operational decisions, that minimize the expected NPV of total system costs.

The study is conducted on a modified IEEE 24-bus system over an eight-year planning horizon divided into three investment stages. Stage k=1 corresponds to the base year 2025, stage k=2 to 2029, and stage k=3 to 2033. Investment decisions include new open-cycle gas turbines (OCGT), combined-cycle gas turbines (CCGT), coal-fired power plants, small modular reactors (SMR), wind turbines, solar panels, 4h ESS, 2h ESS, and thirty-four candidate transmission lines.

Generators and transmission assets have a construction lead time of one stage, whereas ESS investments are assumed to be immediately available. System operations within each investment stage are represented using two representative days, one winter and one summer, to capture seasonal variability. This simplified choice is intentional, as the study is primarily methodological and designed as a proof of concept to demonstrate the proposed framework and analyze the VMS. In practical planning applications, a larger set of representative days would generally be needed to better capture operational variability. The value of lost load is set to $10, 000/MWh, the RPS is set to 40%, and the discount factor used in the NPV objective is 8%.

Uncertainty in the problem is revealed sequentially and operates at different time scales. The uncertainty structure is chosen to reflect commonly observed drivers of long-term uncertainty in power system expansion studies. At stage k=2, strategic uncertainty is realized in the form of geopolitical events that affect capital and fuel costs. We consider three geopolitical events with equal probabilities. At stage k=3, demand-side uncertainty is realized through load growth trajectories reflecting the expansion of data centers and AI-driven electricity demand. We consider two equiprobable load growth realizations at stage k=3, corresponding to moderate and high demand growth scenarios. A summary is provided in Table 1 and the full scenario tree in Figure 2.

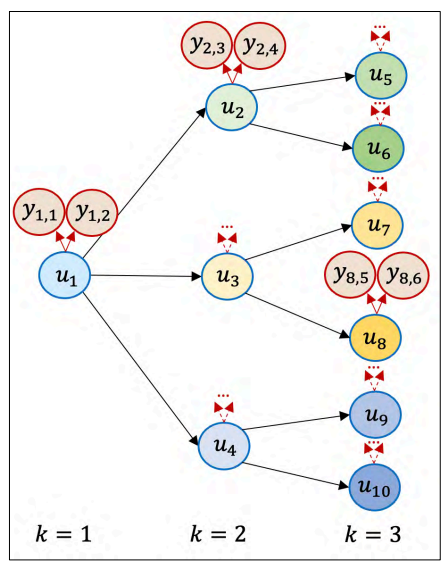

Figure 2. Problem-specific scenario tree.

IMPLEMENTATION

We compare two modeling paradigms. In the TS formulation, the entire investment plan across all stages is

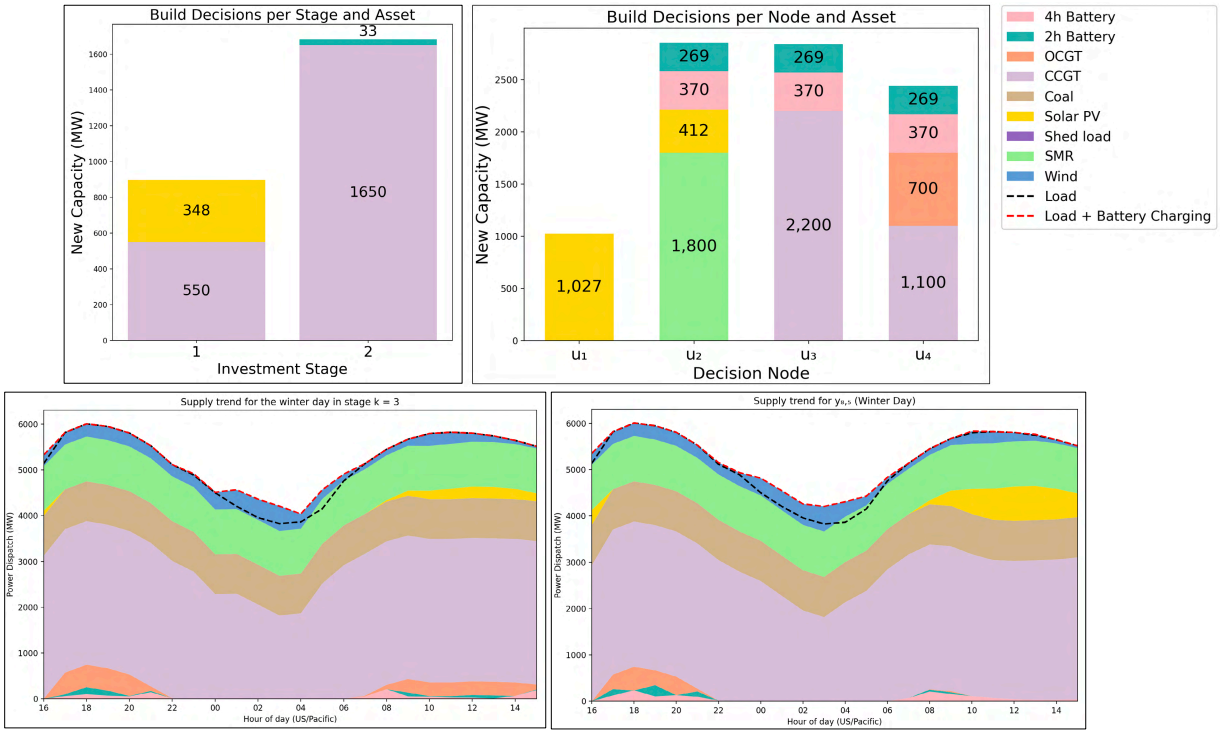

Figure 3. a) Top left corner: Investment decisions for the TS formulation; **b) Top right corner**: Investment decisions for the MS formulation; **c) Bottom left corner**: Supply trend for the TS formulation; **d) Bottom right corner**: Supply trend for the MS formulation;

selected prior to the realization of geopolitical events and load growth. In contrast, the MS formulation allows investment decisions to adapt sequentially: investments made at stage k=2 respond to the realized geopolitical event, while investments at stage k=3 further adapt to the realized load growth trajectory.

The formulations were implemented in Pyomo 6.9.5 [19]. The stochastic formulations were constructed using the `mpi-sppy` package [20], which was used to manage nonanticipativity constraints and assemble extensive-form models. The resulting MILP for TS and MS formulations were solved to optimality gaps of ~2% using academic license of Gurobi v12 [21] on an Apple M4 Pro workstation with 48 GB of RAM. In long-term infrastructure planning, optimality gaps on the order of 1 – 3% are commonly considered acceptable, as they are broadly consistent with the uncertainty in input cost data and are sufficient to support investment and policy analysis.

RESULTS

Investment decisions

Investment decisions for the TS and MS formulations are shown in Figures 3a and 3b, respectively. For brevity, we omit the expansion decisions for investment stage k=3. We also note that none of the optimal solutions recommend investing in new transmission lines. This outcome is consistent with the findings reported in [17], which show that when generation, transmission, and ESS are co-optimized, ESS often defers the need for new transmission expansions.

In the first stage, we note that the MS model invests in 1027 MW of solar panel capacity and defers 550 MW of early CCGT that the TS model builds, preserving the option to reallocate capital once uncertainty resolves. This difference highlights one advantage of the MS formulation: it acknowledges that future cost shocks will influence the relative attractiveness of thermal vs renewable investments. By waiting, the MS model preserves the option to redirect or delay capital depending on how conditions evolve at later stages.

In the second stage, the MS model selects one of three investment plans that directly reflect the realized geopolitical shock, allowing it to tailor its mix of thermal, renewable, and ESS additions to the actual cost in each branch. Importantly, these decisions correspond to forward-looking investments whose capacity becomes available in the next stage (except for ESS). This flexibility lets the MS model capitalize on low-cost opportunities in some branches while avoiding overly expensive technologies in others. In contrast, the TS model must commit to a single CCGT-heavy plan that applies to all futures, regardless of which shock occurs. Because it cannot adjust to branch-specific cost changes, the TS formulation

might overinvest in thermal capacity and overlook cheaper renewable or ESS options, incurring opportunity costs in branches where a different portfolio would have been more economical.

Operational decisions

The MS formulation also provides a cheaper operating schedule. For instance, when the disruption in renewable energy supply chain materializes, the dispatch merit order shifts in stage k=3, so CCGT and solar panels serve most hours while the existing OCGT is mostly idle. In contrast, the TS formulation relies more on OCGT, which has elevated fuel costs, given that it did not invest previously on solar panels. This can be seen in Figures 3c and 3d.

Value of multistage solution

As described in [10], we also denote by v_{TS} and v_{MS} the optimal objective values of the TS and MS formulations, respectively. Thus, the VMS is defined as:

$$\text{VMS} = \frac{v_{TS} - v_{MS}}{v_{MS}} \tag{5}$$

In addition, we can infer that $v_{TS} \geq v_{MS}$, since the MS formulation offers more flexibility in the investment decisions with respect to the set of strategic and operational sets of uncertainty [10]. The objective values obtained for the TS and MS formulations are 393, 857 USD/h and 381, 474 USD/h, respectively. To report costs on standard units, the optimal objective values are normalized by a representative system load, obtained by averaging loads along each root-to-leaf scenario path and then taking expectations over the scenario tree. For the TS and MS formulations we obtain objective values of 0.1179 and 0.1141 USD/kWh.

This reduction in expected electricity prices reflects the value of deferring and adjusting investments in response to realized outcomes. The resulting VMS, defined as the relative improvement of the MS solution over the TS solution, is 3.3%. This result quantifies the economic benefit of modeling investment decisions dynamically and highlights the importance of capturing intertemporal flexibility when planning under long-term uncertainty. For context, retail electricity prices in the United States have been increasing on the order of 2.5 % per year according to EIA data [22], indicating that the modeling difference is economically meaningful.

CONCLUSIONS

This paper demonstrates the value of multi-stage stochastic programming for long-term power grid capacity expansion planning. By allowing investment decisions to adapt to realized geopolitical events and operational uncertainty, the multi-stage formulation produces more flexible investment paths and improved operational schedules relative to a two-stage stochastic programming model. In the case study, this added flexibility yields a 3.3% reduction in expected electricity costs, highlighting the economic benefit of dynamically adjusting investments rather than committing to a single static plan.

From a practical standpoint, this multi-stage stochastic programming framework is best interpreted as a rolling-horizon planning tool. At each investment stage, planners can redefine the scenario tree based on expert judgment, policy insights, and emerging information about technology costs, fuel markets, and demand growth. In this setting, scenario construction plays a central role and reflects structured expert assessments of plausible future conditions. Although the multi-stage stochastic programming formulation is computationally more demanding than its two-stage counterpart, advances in decomposition methods are expected to support its use in realistic planning contexts.

DIGITAL SUPPLEMENTARY MATERIAL

The model formulation for the deterministic program can be found in the Living Archive for Process Systems Engineering with LAPSE ID 2026.0017.

ACKNOWLEDGEMENTS

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and was supported by the LLNL-LDRD Program under Project No. 25-SI-007.

REFERENCES

1. Jin S, Ryan SM, Watson JP, Woodruff DL. Modeling and solving a large-scale generation expansion planning problem under uncertainty. Energy Syst 2:209-242 (2011). https://doi.org/10.1007/s12667-011-0042-9
2. Munoz FD, Watson JP. A scalable solution framework for stochastic transmission and generation planning problems. Comput Manag Sci 12:491-518 (2015). https://doi.org/10.1007/s10287-015-0229-y
3. Cho S, Li C, Grossmann IE. Recent advances and challenges in optimization models for expansion planning of power systems and reliability optimization. Computers & Chemical Engineering 165:107924 (2022). https://doi.org/10.1016/j.compchemeng.2022.107924
4. Koltsaklis NE, Dagoumas AS. State-of-the-art generation expansion planning: a review. Applied

Energy 230:563-589 (2018). https://doi.org/10.1016/j.apenergy.2018.08.087
5. Munoz FD, Hobbs BF, Ho JL, Kasina S. An engineering-economic approach to transmission planning under market and regulatory uncertainties: WECC case study. IEEE Trans. Power Syst. 29:307-317 (2014). https://doi.org/10.1109/tpwrs.2013.2279654
6. Birge JR, Louveaux F. Introduction to stochastic programming. Springer New York (2011). https://doi.org/10.1007/978-1-4614-0237-4
7. Watson JP, Woodruff DL. Progressive hedging innovations for a class of stochastic mixed-integer resource allocation problems. Comput Manag Sci 8:355-370 (2010). https://doi.org/10.1007/s10287-010-0125-4
8. Rahmaniani R, Crainic TG, Gendreau M, Rei W. The benders decomposition algorithm: a literature review. European Journal of Operational Research 259:801-817 (2017). https://doi.org/10.1016/j.ejor.2016.12.005
9. Defourny B, Ernst D, Wehenkel L. Multistage stochastic programming. Decision Theory Models for Applications in Artificial Intelligence :97-143 (2012). https://doi.org/10.4018/978-1-60960-165-2.ch006
10. Huang K, Ahmed S. The value of multistage stochastic programming in capacity planning under uncertainty. Operations Research 57:893-904 (2009). https://doi.org/10.1287/opre.1080.0623
11. Liu Y, Sioshansi R, Conejo AJ. Multistage stochastic investment planning with multiscale representation of uncertainties and decisions. IEEE Trans. Power Syst. 33:781-791 (2018). https://doi.org/10.1109/tpwrs.2017.2694612
12. Thome F, Perez RC, Okamura L, Soares A, Binato S. Stochastic multistage co-optimization of generation and transmission expansion planning. arXiv (2019) https://doi.org/10.48550/arXiv.1910.01640
13. Hou S, Fan Y, Yi BW. Long-term renewable electricity planning using a multistage stochastic optimization with nested decomposition. Computers & Industrial Engineering 161:107636 (2021). https://doi.org/10.1016/j.cie.2021.107636
14. Apablaza P, Püschel-Løvengreen S, Moreno R, Mhanna S, Mancarella P. Assessing the impact of DER on the expansion of low-carbon power systems under deep uncertainty. Electric Power Systems Research 235:110824 (2024). https://doi.org/10.1016/j.epsr.2024.110824
15. Hole J, Philpott AB, Dowson O. Capacity planning of renewable energy systems using stochastic dual dynamic programming. European Journal of Operational Research 322:573-588 (2025). https://doi.org/10.1016/j.ejor.2024.12.031
16. Glista E, Knueven B, Watson J-P. From zonal to nodal capacity expansion planning: Spatial aggregation impacts on a realistic test-case. arXiv (2025) https://doi.org/10.48550/arXiv.2510.23586
17. Go RS, Munoz FD, Watson JP. Assessing the economic value of co-optimized grid-scale energy storage investments in supporting high renewable portfolio standards. Applied Energy 183:902-913 (2016). https://doi.org/10.1016/j.apenergy.2016.08.134
18. Numata Y, Gorin A, Speelman L, Shwisberg L, Gulli C. Fast, Flexible Solutions for Data Centers. RMI. https://rmi.org/fast-flexible-solutions-for-data-centers/
19. Bynum ML, Hackebeil GA, Hart WE, Laird CD, Nicholson BL, Siirola JD, Watson JP, Woodruff DL. Pyomo overview. Springer Optimization and Its Applications :25-36 (2021). https://doi.org/10.1007/978-3-030-68928-5_3
20. Knueven B, Mildebrath D, Muir C, Siirola JD, Watson JP, Woodruff DL. A parallel hub-and-spoke system for large-scale scenario-based optimization under uncertainty. Math. Prog. Comp. 15:591-619 (2023). https://doi.org/10.1007/s12532-023-00247-3
21. Gurobi Optimization, LLC. https://www.gurobi.com
22. U.S. Energy Information Administration. https://www.eia.gov

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Designing in an Unpredictable World: Novel Methods for Uncertainty Characterization, Quantification, and Optimization in Process Engineering

Diederik Coppitters[a]*, Antoine Laterre[a], Mahdi Kchaou[a], Kevin Verleysen[b], Panagiotis Tsirikoglou[c], Jerome Stock[d], Matthias Weigold[d], Konstantinos Kyprianidis[e], Ward De Paepe[f], Francesco Contino[a]

[a] Institute of Mechanics, Materials and Civil Engineering (iMMC), Université catholique de Louvain (UCLouvain), Place du Levant, 2, 1348 Louvain-la-Neuve, Belgium
[b] Thermo and Fluid dynamics (FLOW), Vrije Universiteit Brussel (VUB), Pleinlaan 2, 1050 Brussels, Belgium
[c] Limmat Scientific AG, Weinbergstrasse 31, 8006 Zürich, Switzerland
[d] Institute for Production Management, Technology and Machine Tools (PTW), Technical University of Darmstadt, Otto-Berndt-Str. 2, 64287 Darmstadt, Germany
[e] Mälardalen University, Universitetsplan 1, Västerås, Sweden
[f] Thermal Engineering and Combustion Unit, University of Mons, Rue de l'Épargne 56, Mons, 7000, Belgium
* Corresponding Author: diederik.coppitters@uclouvain.be.

ABSTRACT

Computer-Aided Process Engineering (CAPE) has transformed how we analyze, design, and optimize energy processes. Yet, even advanced models rest on uncertain ground: their reliability depends on how well future operating environments are described—environments that are dynamic, complex, and deeply uncertain. In practice, uncertainty is often reduced to local parameter variations, driven by limited data, computational burden, and overconservative robust formulations. This narrow treatment creates a false sense of confidence: Designs that perform well in theory often fail in real-world operation. In a century marked by economic, climatic, and technological volatility, designing under uncertainty is no longer optional; it is essential.We have developed approaches that place uncertainty at the core of energy process modeling and design. This paper provides an overview of these methods and how uncertainty can be explicitly represented, quantified, and embedded into the design process.We present approaches to characterize uncertainties under data scarcity, including imprecise probabilities to distinguish between aleatory and epistemic uncertainty. To propagate uncertainty through computationally intensive models, we introduce surrogate-assisted techniques that exploit structural sparsity, enabling the analysis of systems with large numbers of uncertain inputs (100+) while mitigating the curse of dimensionality. These methods are integrated into optimization frameworks that target expected performance, robustness, and, for the first time, antifragility—systems that can benefit from variability rather than merely withstand it. We illustrate these approaches across applications ranging from detailed process models to system-level energy analyses, advancing a shift in CAPE toward designs suited for an unpredictable world.

Keywords: Uncertainty assessment, Robust design optimization, Process simulation, Antifragility

INTRODUCTION

Computer-Aided Process Engineering (CAPE) tools, such as Aspen Plus, gPROMS, Thermoflex, and in-house high-fidelity thermodynamic simulators, have become central to process design and operation [1]. Engineers now evaluate complex heat integration networks, power-to-X routes, carbon capture systems, and flexible plants using nonlinear, tightly coupled models that embed thermodynamics, transport phenomena, and equipment constraints. These models can simulate performance with high accuracy, but they remain conditional on

assumptions about their operating environment, such as variations in ambient temperature and humidity, cooling-water temperature, heat-transfer coefficients, fuel composition (carbon, hydrogen, lower-heating value), electricity prices, CAPEX and OPEX of emerging units, policy constraints, market regimes, and evolving sustainability targets. Despite this, uncertainty in many CAPE studies is still mostly handled through deterministic base-case optimization followed by local sensitivity checks or a small scenario sweep [2]. This approach persists for understandable reasons: (i) data about the operating environment is fundamentally incomplete which makes rigorous uncertainty characterization difficult [3], (ii) brute-force Monte Carlo becomes infeasible for expensive nonlinear flowsheets [4], and (iii) worst-case robust formulations often drive designs toward conservatism that is difficult to defend in engineering and economic terms [5]. This narrow treatment of uncertainty can create confidence in designs that are fragile when exposed to operational variability and imperfect foresight.

Over the past decade, we have developed methods that place uncertainty at the center of CAPE modeling and design. This paper documents a probabilistic uncertainty workflow that we use for nonlinear, technology-detailed CAPE models under limited information. We focus on uncertainty characterization under limited data, surrogate-assisted uncertainty propagation that remains tractable for expensive simulations, and robust design optimization that explicitly trades off expected performance and sensitivity to variability. We also present antifragility as a paradigm shift for optimization under uncertainty and ways to address structural uncertainty, where uncertainty is not only about parameter values but also about what the problem formulation should be, for example when stakeholder preferences or acceptability thresholds are uncertain. Finally, we conclude this paper with a discussion and perspectives.

A FRAMEWORK FOR UNCERTAINTY-AWARE PROCESS DESIGN

In this section, we present a framework for uncertainty-aware process design (Figure 1). It covers methods for characterizing uncertainty, propagating uncertainty and robust design optimization.

Representing Uncertainty Under Limited Data

Uncertainty characterization translates knowledge about the parameter variation into a mathematical representation that will drive propagation and optimization. In CAPE, researchers repeatedly face the same obstacle [2]: Data about the operating environment is often sparse, inconsistent, or context-specific. For instance, electricity-price projections may exist only as a small set of scenarios, emerging unit costs may be based on few studies with incompatible assumptions, and ambient conditions may only be available at coarse data resolution. In this setting, analysts often default to common distributions such as uniform or normal [3]. In this way, they risk encoding strong assumptions, either by assigning zero probability to tail events or by extrapolating tails from limited central information. In nonlinear flowsheets, these assumptions can change optimized designs because feasibility, active constraints, and operating regimes can switch across the uncertain domain [6].

One way to represent uncertainty uses quantile-anchored parametric distributions by combining data with expert judgment. While frequently available, translating expert judgment into distributional parameters (e.g., mean, variance) is difficult: Even when experts can provide plausible ranges, they struggle to express how probability mass should be distributed within those ranges [7]. Expert sources are usually more consistent when asked for percentile statements [8]. Therefore, a pragmatic way is to elicit expert judgment on quantiles rather than attempt to obtain full distribution parameters directly, such as “a pessimistic value that is likely exceeded in only 5% of the scenarios.” Moreover, if the modeler commits to a distribution type (e.g., a lognormal distribution to represent the future electricity price), two distinct percentiles determine a unique distribution (in the common case of a two-parameter type), so the resulting distribution is fully anchored to explicit expert statements [9]. To illustrate this approach, we elicited two percentiles and distribution types from expert reports and characterized the corresponding distributions for grid electricity price, wind speeds and electricity demand in the assessment of a wind-powered community [8].

A second representation uses imprecise probabilities to treat distributional ambiguity explicitly. The same elicited quantiles may be consistent with a bounded distribution that has hard limits, or with a heavy-tailed distribution that assigns small but decision-relevant probability to extreme outcomes. When this ambiguity affects design choices, a single fitted curve can convey unwarranted certainty. Probability boxes (p-boxes) address this by representing uncertainty as a family of admissible distributions bounded by lower and upper cumulative distribution functions [10]. This representation is particularly useful when epistemic uncertainty dominates, because it treats uncertainty about the distribution itself as part of the model, not as an afterthought. It also supports a separation between aleatory variability, which reflects intrinsic randomness, and epistemic uncertainty, which reflects limited knowledge that could shrink with better information or clearer assumptions.

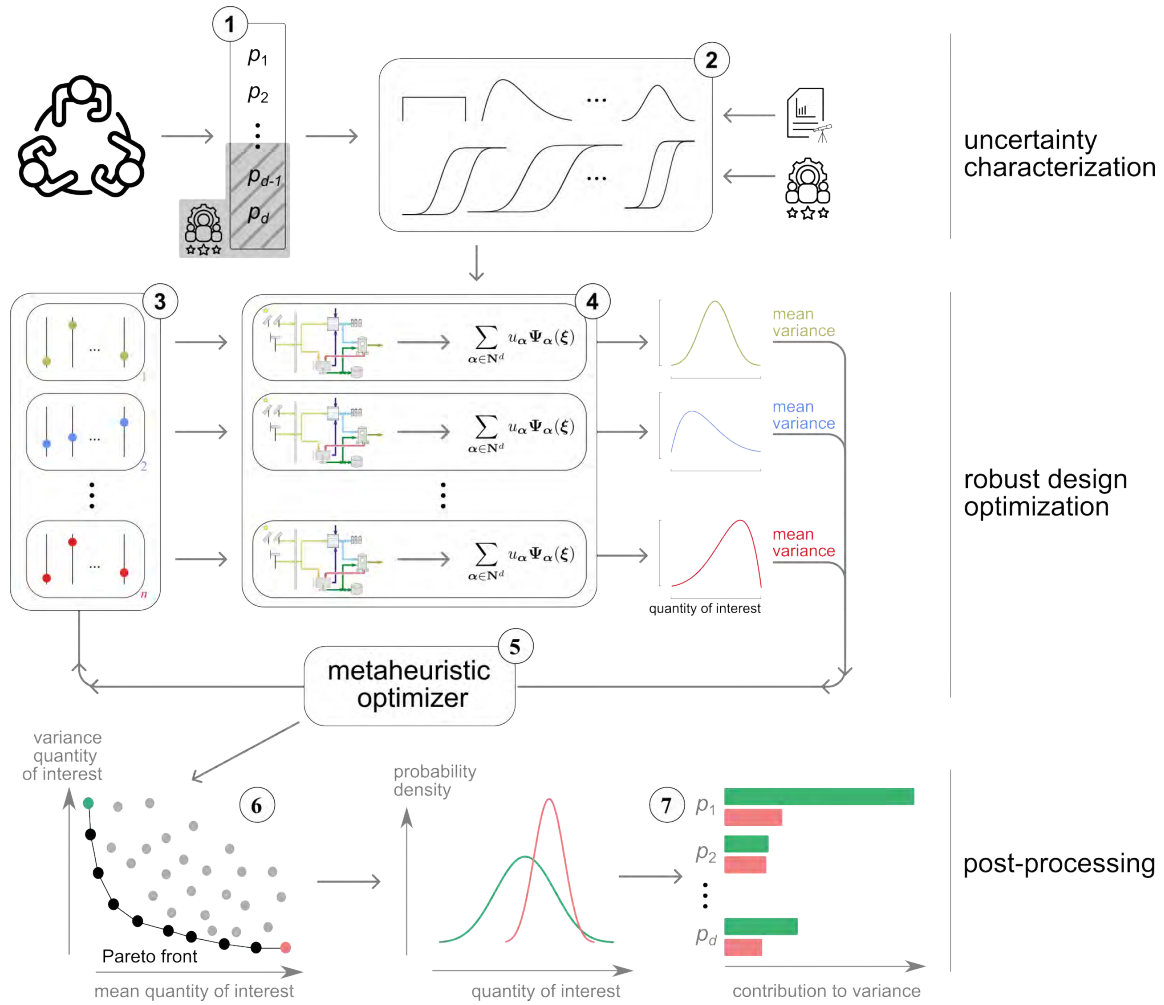

Figure 1. The schematic illustrates the proposed framework for uncertainty-aware process design. Experts first identify the parameters to be treated as uncertain (1). An optional expert-based screening may be applied to discard parameters expected to have negligible influence. Uncertainty is then characterized by combining data, projections, and expert judgment, using either precise probabilistic distributions or imprecise formulations when epistemic uncertainty is substantial (2). The optimization workflow then begins. An initial set of candidate designs is generated by exploring the design space, for example using Latin Hypercube Sampling (3). For each candidate design, the process model is evaluated under uncertainty by constructing a surrogate, typically a full or sparse Polynomial Chaos Expansion depending on model complexity and the number of retained uncertain inputs (4). Performance statistics such as the mean and variance, or bounds on these quantities in the imprecise case, are computed analytically from the surrogate and provided to a population-based metaheuristic optimizer as fitness values (5). The optimizer iteratively generates new design populations and updates the design set. After convergence, the resulting Pareto front quantifies the trade-off between expected performance and robustness (6). For selected Pareto-optimal designs, global sensitivity analysis is extracted analytically from the surrogate to identify the uncertain inputs that dominate performance variability and those with negligible influence (7).

We used probability boxes to represent the operating environment of a hybrid renewable energy system comprising a photovoltaic array, batteries, a heat pump, and thermal storage. For electricity and heat demand, aleatory uncertainty represented irreducible variability in future demand, while epistemic uncertainty reflected limited knowledge of occupant behavior and reliance on generic demand profiles. Energy prices were treated analogously, with aleatory uncertainty capturing plausible market evolution and epistemic uncertainty capturing ambiguity in tariffs and supplier choice. We propagated these probability boxes through the system model and used the aleatory component to drive robust design optimization of the levelized cost of exergy, because this

variability cannot be reduced through additional information and can only be mitigated through design choices. In contrast, we used global sensitivity analysis on the epistemic component to identify which uncertain assumptions most affected predicted performance and robustness, thereby indicating where additional information would most effectively reduce decision-relevant uncertainty [11].

In practice, a precise probabilistic formulation is pragmatic, as it is simpler to implement, propagate, and communicate, but it necessarily collapses epistemic uncertainty into a single predictive model. This is often acceptable for parameters such as ambient conditions, standardized fuel properties, or short-term market fluctuations, where uncertainty is largely aleatory or where epistemic uncertainty is small relative to its effect on design decisions. We propose an imprecise representation when epistemic uncertainty is significant, meaning that uncertainty about the parameters or assumptions of the distribution itself is large enough to affect design conclusions. Typical examples include electricity or carbon prices over multi-decadal horizons, CAPEX of emerging technologies, or demand projections driven by poorly specified user behavior. In such cases, explicitly separating aleatory variability, which must be mitigated through design, from epistemic uncertainty, which can be reduced through additional information or specification, avoids masking decision-relevant uncertainty, despite the added conceptual and computational complexity.

Uncertainty Propagation and Quantification

Once input uncertainties have been characterized, whether through precise or imprecise probabilities, the next step is to propagate that uncertainty through the model to quantify its impact on the outputs of interest. The most straightforward approach is Monte Carlo Simulation (MCS), which samples uncertain inputs from their assigned distributions, evaluates the model for each sample, and estimates output statistics. Although this approach is conceptually simple and asymptotically convergent, it becomes impractical for many CAPE applications. MCS typically requires a large number (~$10^{3\text{-}5}$) of model evaluations to resolve output variability and tail behavior, while high-fidelity CAPE models are computationally expensive (minutes to hours) and often involve many uncertain inputs (>10). As a result, naive Monte Carlo propagation quickly becomes computationally intractable.

An effective alternative is to move to surrogate-assisted Uncertainty Quantification (UQ). The core idea is to replace numerous brute-force evaluations of an expensive simulator with a cheaper approximation that captures the dominant input–output relationships. Among available surrogate approaches, Polynomial Chaos Expansion (PCE) is particularly well suited to UQ because it represents model outputs as an expansion in orthogonal polynomials of the uncertain inputs [12]. Once a PCE surrogate is trained, it can be evaluated orders of magnitude faster (typically milliseconds) than the original model. An additional benefit of PCE is that uncertainty propagation naturally yields statistical moments on the output (mean, variance) and a global sensitivity analysis can be performed at negligible additional cost by exploiting the PCE coefficients. In practice, high Sobol' indices guide where further knowledge is necessary. For example, we applied PCE to the integration of Direct Air Capture (DAC) in a power-to-gas flowsheet to quantify uncertainty in energy efficiency, exergy efficiency and levelized cost under uncertain ambient conditions, sorbent thermodynamics, electricity price, and CAPEX/OPEX. The resulting Sobol' indices identified that electricity price and DAC CAPEX dominated the variance of levelized cost, while thermodynamic model parameters and ambient conditions had a secondary influence on the overall performance [13].

In CAPE applications, the uncertainty space typically includes many input parameters, which makes a full PCE impractical because the number of required expansion coefficients, and thus training simulations, grows combinatorially with both polynomial degree (linked to the nonlinearity of the input-output relation) and the number of uncertain parameters. Dimensionality reduction is therefore unavoidable.

Screening based on expert judgment or coarse sensitivity checks provides a pragmatic first reduction, but it is inherently subjective and can overlook higher-order interactions that only emerge in nonlinear regimes. Multi-fidelity PCE strategies alleviate this by using low-order surrogates to rank influential inputs before constructing higher-order models, but their effectiveness depends on the reliability of these preliminary approximations, which can be poor when responses are highly nonlinear or discontinuous. In a long-term energy-system planning model with many uncertain techno-economic inputs, we used a low-order PCE as a preliminary step to reduce an initial set of uncertain parameters to a smaller set of critical parameters before building the higher-fidelity surrogate and revealed that a small subset of fuel-price and demand-growth assumptions dominated the uncertainty on the system costs [14].

Sparse PCE addresses these limitations by integrating screening directly into surrogate construction, retaining only influential polynomial terms and interactions while discarding the rest, thereby controlling dimensionality without requiring explicit pre-screening or multiple surrogate stages and significantly reducing the number of required training simulations [15]. We used sparse PCE for a Thermoflex®-based retrofit study of a large thermal power plant where 243 uncertain inputs made conventional PCE infeasible, and the resulting Sobol' indices showed that fuel composition, lower heating value, and boiler efficiency parameters dominated variability in

furnace temperature and efficiency, while most auxiliary parameters had negligible influence [16].

PCE can also be applied when input uncertainties are represented through probability boxes. In a parametric p-box setting, uncertainty in the statistical moments, such as the mean and variance, is represented by uniformly distributed variables to reflect bounded epistemic uncertainty. For any specific realization of these parameters, a corresponding formulation of the aleatory uncertainty is defined through the associated conditional distribution. Because this conditional formulation introduces dependence between epistemic and aleatory components, the aleatory variable is rewritten in standardized form, and the physical input is reconstructed by combining this standardized variable with the sampled mean and variance. This reformulation represents the uncertain input through a set of independent variables: one standardized variable capturing aleatory variability and uniform variables capturing epistemic uncertainty in the distribution parameters. A single PCE surrogate constructed on this augmented input space can then be sampled efficiently to compute bounds on the output distribution, yielding the probability box of the model response without resorting to a nested MCS [17]. In a PV–battery–heat-pump system with thermal storage, we represented mixed aleatory variability (demand and weather evolution) and epistemic ambiguity (occupant behavior and cost assumptions) with parametric p-boxes, then used the augmented PCE to propagate output bounds while using the aleatory part for design objectives and the epistemic part for sensitivity-guided information prioritization, showing that aleatory demand variability primarily affected the robustness, while epistemic uncertainty in electricity tariffs and user behavior dominated the epistemic uncertainty on the optimized costs [11].

Designing for Robustness

Design optimization under uncertainty is where all the pieces come together—and also where computational burden tends to peak. The goal is to identify design decisions (equipment sizes, operating conditions, etc.) that perform well across the range of uncertainties characterized earlier. A specific type of optimization under uncertainty is Robust Design Optimization (RDO), where the mean and variance of the quantities of interest obtained from UQ, such as costs, efficiencies and emissions, define the optimization objectives and constraints.

From a computational perspective, RDO inherits the computational cost of UQ. We therefore build on the surrogate-based propagation strategies introduced in the previous section and embed them within an optimization loop [19]. For a given candidate design, the design variables are held fixed, and the PCE is constructed by sampling the uncertain inputs around that design according to their prescribed probabilistic descriptions. Because the surrogate represents the full input–output map for that design, the statistical moments required for the objectives can be computed analytically from the PCE coefficients. Reliability-type constraints can be acquired in the same way by evaluating, for each design, the probability of violating constraints based on the propagated uncertainty [18]. In a high-fidelity RDO setting considered here, uncertainty propagation is thus performed separately for each candidate design evaluated by the optimizer. While this approach is computationally more demanding than global surrogates spanning both design and uncertainty spaces, it preserves accuracy in the presence of strong nonlinearities in the design space.

The optimization itself is performed using gradient-free, population-based algorithms, which are well suited to the black-box and often nonconvex nature of CAPE problems [20]. In our applications, multi-objective formulations are solved using NSGA-II, with objectives being the mean and variance of the quantity of interest, such as cost, efficiency and emissions. The resulting Pareto front quantifies the trade-off between the expected performance and robustness and provide a structured basis for optimized design selection.

The validity of this approach depends on the reliability of the surrogate models used during optimization. The fidelity is therefore enforced through an automated two-step procedure. First, the design space is explored using a space-filling sampling of the design variables, such as Latin Hypercube Sampling. For this initial set of designs, full PCEs are built for increasing polynomial orders, and the worst-case leave-one-out error over all sampled designs is evaluated. The smallest polynomial order that meets a predefined accuracy threshold is then selected and fixed for the remainder of the optimization. During the RDO iterations, sparse PCEs of this order are constructed for each candidate design, and their accuracy is monitored using leave-one-out error and stabilization of the estimated mean and standard deviation as training samples are added. If a candidate design leads to non-convergent or nonphysical model evaluations for one or more training samples, that design is classified as infeasible and assigned a penalized objective value.

We implemented the full workflow in an open-source framework called RHEIA [21] and applied it to several CAPE applications with distinct process structures and decision variables. In a grid-connected photovoltaic system with battery and hydrogen storage, RDO shifted the preferred designs from photovoltaic-only systems toward battery and hydrogen storage when robustness was valued. For a community-scale system, the robust PV–battery–hydrogen design increased the mean levelized cost of electricity from 269 €/MWh to 363 €/MWh (≈+35%), while reducing its standard deviation from 55.6 €/MWh to 41.9 €/MWh (≈ −25%), explicitly trading expected cost for reduced sensitivity to electricity-price

uncertainty [19]. In a wind-powered power-to-ammonia plant, RDO revealed a clear trade-off between productivity and robustness: the design maximizing mean annual ammonia production achieved 3.2 times higher average output, but its standard deviation of annual production was 2.6 times larger than that of the robust design, indicating substantially higher sensitivity to wind-speed and reactor-temperature uncertainties [22]. Finally, for a renewable-powered hydrogen refueling station and bus fleet, converting 54% of the diesel buses to hydrogen-powered buses reduced the standard deviation of the levelized cost by 36%, reduced the mean carbon intensity by 46%, and reduced its standard deviation by 51%, at the expense of an 11% increase in mean cost [23].

Limitations

The workflow has three main limitations related to computational cost, input dependency, and accuracy.

First, the computational cost can still be high because the RDO loop requires building, validating, and sampling a surrogate for each candidate design. This nested structure can have computational overhead because of expensive models, large populations, many generations, and higher polynomial orders.

Second, standard PCE construction typically assumes independent uncertain inputs. Many CAPE problems feature correlated drivers (e.g., demand–price coupling, weather covariation, or jointly evolving techno-economic assumptions). Ignoring dependence can misrepresent joint extremes, constraint-violation risk, and sensitivity indices. Accounting for dependence is possible (e.g., via transformations or copula-based modeling), but it introduces additional modeling choices.

Third, PCE accuracy degrades when outputs are non-smooth or discontinuous in the uncertain inputs, which can arise from regime switching, active-constraint changes, phase changes, and feasibility boundaries. In these cases, moment estimates, tail metrics, and Sobol' indices may become unreliable unless more complex surrogate strategies are used (e.g., piecewise surrogates or classification-plus-regression), which reduces the simplicity of the approach.

EMERGING DIRECTIONS IN UNCERTAINTY-AWARE DESIGN

While RDO provides a principled way to control sensitivity to uncertainty, it implicitly treats uncertainty as a source of risk. As a result, optimization strategies that focus exclusively on robustness tend to penalize all variability symmetrically, which can suppress configurations that would perform better under favorable operating conditions. In addition, RDO primarily addresses uncertainty in model parameters and implicitly assumes that the chosen model structure and objective formulation capture all decision-relevant aspects. In practice, however, factors such as social acceptance and stakeholder preferences may be difficult, if not impossible, to encode, yet they can strongly influence whether an optimized design can be implemented. The following sections therefore examine two complementary extensions: antifragility, which explicitly accounts for asymmetric responses to variability, and approaches that address structural and preference uncertainty by exploring sets of near-optimal alternatives rather than a single solution.

Antifragility: A Paradigm Shift in Designing Under Uncertainty

Optimization under uncertainty in CAPE has traditionally focused on robustness, that is, identifying designs whose performance is least sensitive to uncertain operating conditions. While effective when uncertainty is well characterized, this paradigm becomes limiting when the magnitude and structure of future techno-economic uncertainty are themselves poorly known. Moreover, robustness treats variability symmetrically: it controls downside risk, but it also suppresses upside opportunity.

An emerging alternative is antifragile design, which shifts the objective from minimizing harm due to uncertainty to exploiting favorable variability—an antifragile system does not merely withstand uncertainty; it benefits from it [24]. We have explored antifragility as an extension of RDO by introducing objectives that reward beneficial variability rather than penalizing variability indiscriminately [8]. One such metric is the Upside Potential Ratio (UPR), which favors designs that achieve substantially better-than-average outcomes while still limiting unacceptable losses. For common performance indicators, encouraging upside variability corresponds to the possibility of higher efficiency, lower cost, or lower emissions than expected under nominal conditions. We also propose using the skewness of the performance distribution as an objective to quantify antifragility. Positive skewness indicates limited downside exposure combined with rare but significant gains, and it captures how a design responds when the actual operating environment is more uncertain than anticipated.

Embedding these objectives within a multi-objective optimization framework yields designs with qualitatively different behavior from conventional robust solutions. In the design of a wind turbine for a community, antifragile optimization led to larger installed capacity than the robust design, deliberately accepting worse median performance and higher variability, but exhibited strongly asymmetric outcomes: when the actual uncertainty exceeded initial assumptions, it delivered substantially lower costs than both the deterministic and robust designs [8]. In a residential heat-pump dimensioning case for a single-family dwelling, antifragile objectives similarly shifted optimal designs away from those minimizing

median cost. Designs with higher antifragility featured larger thermal capacity and storage, which increased cost but limited downside exposure and enabled gains under volatile electricity prices and demand [25].

Addressing Structural and Preference Uncertainty

Decision-makers are rarely concerned only with the parameters and objectives explicitly modeled. Factors such as social acceptance, siting constraints, permitting risk, or political feasibility are difficult, if not impossible, to encode quantitatively, yet they can strongly influence whether an optimized design can be implemented. As a result, presenting only a single optimized solution, or even a narrow Pareto front, can be misleading. It embeds one particular problem formulation and worldview, while masking structural uncertainty, that is, uncertainty about the appropriate system layout or decision framing.

In the energy modeling literature, Modeling to Generate Alternatives (MGA) is an established response to this problem. MGA has been used primarily in Linear Programming settings to produce sets of near-optimal solutions that are structurally different, allowing decision-makers to choose among alternatives when unmodeled preferences become important [26].

Our work extended MGA concepts to nonlinear multi-objective optimization, generating portfolios of near-optimal designs under preference uncertainty and using these portfolios to reveal which design features are “must-haves” versus genuine degrees of freedom [27]. In practice, this was done by first computing the Pareto-optimal solutions and then solving additional optimization problems that were constrained to remain within a small, predefined degradation of the objective values while forcing variation in configuration choices and key design variables. The resulting set of designs was then analyzed by comparing configurations and parameter ranges across the ensemble. Applied to a Carnot-battery system, this procedure showed that some design choices, such as recuperated subcritical Heat Pump–Organic Rankine Cycle layouts for high-efficiency operation, appeared systematically across all near-optimal solutions and can be considered must-haves. Other choices, including temperature lift, superheating level, and specific operating points, varied widely with little impact on performance, indicating real choices. However, further convergence analysis and benchmarking are needed to assess the effectiveness of the method, and additional development of the post-processing is required to extract clearer and more compelling messages from the large volume of results.

CONCLUSION

Designing energy and process systems in an increasingly volatile operating environment requires uncertainty to be treated as a core modeling and optimization dimension rather than as a secondary sensitivity check. As operating environments become more volatile and assumptions harder to defend, design methods must explicitly acknowledge limited information, structural ambiguity, and asymmetric responses to variability. In this paper, we presented a coherent workflow for uncertainty-aware CAPE that spans uncertainty characterization under limited information, surrogate-assisted uncertainty propagation for computationally expensive models, and optimization frameworks that explicitly account for variability. Together, these methods advance a shift in CAPE from deterministic optimization toward design strategies that remain credible, informative, and actionable under uncertainty.

ACKNOWLEDGEMENTS

D.C. acknowledges the support of the Fonds de la Recherche Scientifique - FNRS [CR 40016260]. M.K. acknowledges the support of Wallonie-Bruxelles International – WBI.

AUTHOR IDENTIFIERS

Author ORCIDs:
Diederik Coppitters: 0000-0001-9480-2781
Antoine Laterre: 0000-0003-4226-1594
Mahdi Kchaou: 0009-0000-5840-409X
Kevin Verleysen: 0000-0002-7763-7208
Jerome Stock: 0009-0006-7908-069X
Matthias Weigold: 0000-0002-7820-8544
Konstantinos Kyprianidis: 0000-0002-8466-356X
Ward De Paepe: 0000-0001-5008-2946
Francesco Contino: 0000-0002-8341-4350

REFERENCES

1. Gani R, Cameron I, Lucia A, Sin G, Georgiadis M. Process systems engineering, 2. modeling and simulation. Ullmann's Encyclopedia of Industrial Chemistry : (2012). https://doi.org/10.1002/14356007.o22_o06
2. van der Spek M, Fout T, Garcia M, Kuncheekanna VN, Matuszewski M, McCoy S, Morgan J, Nazir SM, Ramirez A, Roussanaly S, Rubin ES. Uncertainty analysis in the techno-economic assessment of CO2 capture and storage technologies. critical review and guidelines for use. International Journal of Greenhouse Gas Control 100:103113 (2020). https://doi.org/10.1016/j.ijggc.2020.103113

3. Mavromatidis G, Orehounig K, Carmeliet J. A review of uncertainty characterisation approaches for the optimal design of distributed energy systems. Renewable and Sustainable Energy Reviews 88:258-277 (2018). https://doi.org/10.1016/j.rser.2018.02.021
4. Zhang J. Modern monte carlo methods for efficient uncertainty quantification and propagation: a survey. WIREs Computational Stats 13: (2020). https://doi.org/10.1002/wics.1539
5. Bertsimas D, Sim M. The price of robustness. Operations Research 52:35-53 (2004). https://doi.org/10.1287/opre.1030.0065
6. Sahinidis NV. Optimization under uncertainty: state-of-the-art and opportunities. Computers & Chemical Engineering 28:971-983 (2004). https://doi.org/10.1016/j.compchemeng.2003.09.017
7. Grigore B, Peters J, Hyde C, Stein K. Methods to elicit probability distributions from experts: a systematic review of reported practice in health technology assessment. PharmacoEconomics 31:991-1003 (2013). https://doi.org/10.1007/s40273-013-0092-z
8. Coppitters D, Contino F. Optimizing upside variability and antifragility in renewable energy system design. Sci Rep 13: (2023). https://doi.org/10.1038/s41598-023-36379-8
9. Cook JD. Determining distribution parameters from quantiles. UT MD Anderson Cancer Center Dept Biostatistics (2010)
10. Ferson S, Ginzburg LR. Different methods are needed to propagate ignorance and variability. Reliability Engineering & System Safety 54:133-144 (1996). https://doi.org/10.1016/s0951-8320(96)00071-3
11. Coppitters D, De Paepe W, Contino F. Robust design optimization of a photovoltaic-battery-heat pump system with thermal storage under aleatory and epistemic uncertainty. Energy 229:120692 (2021). https://doi.org/10.1016/j.energy.2021.120692
12. Sudret B. Global sensitivity analysis using polynomial chaos expansions. Reliability Engineering & System Safety 93:964-979 (2008). https://doi.org/10.1016/j.ress.2007.04.002
13. Coppitters D, Costa A, Chauvy R, Dubois L, De Paepe W, Thomas D, De Weireld G, Contino F. Energy, exergy, economic and environmental (4E) analysis of integrated direct air capture and CO 2 methanation under uncertainty. Fuel 344:127969 (2023). https://doi.org/10.1016/j.fuel.2023.127969
14. Rixhon X, Limpens G, Coppitters D, Jeanmart H, Contino F. The role of electrofuels under uncertainties for the belgian energy transition. Energies 14:4027 (2021). https://doi.org/10.3390/en14134027
15. Abraham S et al. A robust and efficient stepwise regression method for building sparse polynomial chaos expansions. J Comput Phys 332:461–474 (2017) https://doi.org/10.1016/j.jcp.2016.12.019
16. De Meulenaere R, Coppitters D, Sikkema A, Maertens T, Blondeau J. Uncertainty quantification for thermodynamic simulations with high-dimensional input spaces using sparse polynomial chaos expansion: retrofit of a large thermal power plant. Applied Sciences 13:10751 (2023). https://doi.org/10.3390/app131910751
17. Schöbi R, Sudret B. Global sensitivity analysis in the context of imprecise probabilities (p-boxes) using sparse polynomial chaos expansions. Reliab Eng Syst Saf 187:129–141 (2019) https://doi.org/10.1016/j.ress.2019.02.017
18. Paepe W, Coppitters D, Abraham S, Tsirikoglou P, Ghorbaniasl G, Contino F. Robust operational optimization of a typical micro gas turbine. Energy Procedia 158:5795-5803 (2019). https://doi.org/10.1016/j.egypro.2019.01.549
19. Coppitters D, De Paepe W, Contino F. Robust design optimization and stochastic performance analysis of a grid-connected photovoltaic system with battery storage and hydrogen storage. Energy 213:118798 (2020). https://doi.org/10.1016/j.energy.2020.118798
20. Tsirikoglou P, Abraham S, Contino F, Ba?ci Ö, Vierendeels J, Ghorbaniasl G. Comparison of metaheuristics algorithms on robust design optimization of a plain-fin-tube heat exchanger. 18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference : (2017). https://doi.org/10.2514/6.2017-3827
21. Coppitters D, Tsirikoglou P, Paepe W, Kyprianidis K, Kalfas A, Contino F. RHEIA: robust design optimization of renewable hydrogen and derived energy carrier systems. JOSS 7:4370 (2022). https://doi.org/10.21105/joss.04370
22. Verleysen K, Coppitters D, Parente A, De Paepe W, Contino F. How can power-to-ammonia be robust? optimization of an ammonia synthesis plant powered by a wind turbine considering operational uncertainties. Fuel 266:117049 (2020). https://doi.org/10.1016/j.fuel.2020.117049
23. Coppitters D, Verleysen K, De Paepe W, Contino F. How can renewable hydrogen compete with diesel in public transport? robust design optimization of a hydrogen refueling station under techno-economic and environmental uncertainty. Applied Energy 312:118694 (2022). https://doi.org/10.1016/j.apenergy.2022.118694
24. Taleb NN. 'antifragility' as a mathematical idea.

Nature 494:430-430 (2013). https://doi.org/10.1038/494430e
25. Frank M et al. A software framework for optimization under stochastic uncertainty to achieve antifragility. Procedia CIRP
26. Lombardi F, van Greevenbroek K, Grochowicz A, Lau M, Neumann F, Patankar N, Vågerö O. Near-optimal energy planning strategies with modeling to generate alternatives to flexibly explore practically desirable options. Joule 9:102144 (2025). https://doi.org/10.1016/j.joule.2025.102144
27. Laterre A, Coppitters D, Lemort V, Contino F. Designing small-scale rankine carnot batteries that suit your preferences: a near-optimal approach. Journal of Energy Storage 141:118650 (2026). https://doi.org/10.1016/j.est.2025.118650

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Enhancing Interpretability of Stochastic Programming Solutions: A Multiparametric Approach

Parth Brahmbhatt[a] and Styliani Avraamidou[a]*

[a] Department of Chemical and Biological Engineering, University of Wisconsin-Madison, Madison, WI, USA
* Corresponding Author: avraamidou@wisc.edu

ABSTRACT

Stochastic programming (SP) is a powerful framework for decision-making under uncertainty, but its practical adoption in industry is often hindered by the difficulty in understanding the causal relationships that drive optimal solutions. In the two-stage SP, strategic first-stage decisions are coupled with operational second-stage recourse decisions. When the number of scenarios under consideration is large, understanding the direct link between the uncertainty realization and optimal recourse strategy becomes computationally and cognitively demanding. Common approaches to improve interpretability include trained classification trees or scenario reduction, replacing the large scenario set with a representative subset. This is often achieved through post-hoc clustering (e.g., k-means) based on uncertainty realizations or optimal recourse decisions. While useful, these methods only provide a statistical approximation of the solution space and may fail to reveal the underlying structural properties of the recourse problem that drive optimal first-stage decisions. This work introduces a novel, deterministic approach to explainability using multiparametric programming (mp) within a Benders decomposition framework. We reformulate the recourse subproblem as a multiparametric linear program, generating an explicit map of Critical Regions (CRs), which are polyhedral partitions of the uncertainty space. This allows us to cluster scenarios analytically rather than statistically. We demonstrate this methodology on a supply chain planning under demand uncertainty. Our results show that 100 stochastic scenarios map to exactly six critical region clusters. This mapping allows us to explain optimal capacity planning decisions as a precise trade-off between specific operational modes, providing a fully transparent interpretation of the stochastic solution.

Keywords: Stochastic Optimization, Design Under Uncertainty, Supply Chain, Multiparametric Programming

INTRODUCTION

Optimization under uncertainty is central to Process Systems Engineering (PSE), particularly in applications such as supply chain management, energy systems, and capacity planning. Two-stage stochastic programming (SP) [1] is the standard approach for these problems, distinguishing between "here-and-now" decisions (made before uncertainty reveals itself) and "wait-and-see" recourse decisions (made after the uncertainty is realized).

Despite the mathematical maturity of SP, the resulting solutions often function as "black boxes" for decision-makers [2]. A typical large-scale SP solution provides a single optimal first-stage vector x^* that minimizes expected cost across thousands of scenarios. However, it rarely explains why this specific x^* is robust. For instance, in a capacity expansion problem, a manager might ask: "Why is the optimal capacity for Plant B set to 26 units rather than 30?" The answer lies buried in the complex interactions of thousands of second-stage variables.

Recent developments in Explainable AI [3] have attempted to bridge this gap. In the context of optimization, explainability usually relies on sensitivity analysis, which is inherently local, or post-hoc analysis using machine learning surrogates like classification trees [4].

A common approach to reduce complexity is scenario reduction [5], where a large set of scenarios is clustered into a smaller representative set. This is often achieved through post-hoc clustering [6] techniques (e.g., k-means). Rathi et al. [2] recently proposed

"recourse-based clustering, " where scenarios are grouped not by their input uncertain parameters (e.g., demand), but by the similarity of their optimal recourse decisions (y^*). While the method is improved compared to input-based clustering, it remains a statistical approximation, requiring the user to arbitrarily select the number of clusters (k), potentially obscuring the exact boundaries where the optimal decision strategy changes.

In this work, we propose a rigorous, analytical alternative to statistical clustering. We leverage the synergy between Benders Decomposition [7] and Multiparametric Programming (mp) [8]. By treating the first-stage variables and uncertainty as parameters in the second-stage subproblem, we can solve the subproblem explicitly [9]. The result is a set of Critical Regions (CRs). Each CR represents a partition of the parameter space where the optimal operating strategy (the set of active constraints) is identical. This allows us to:

- **Derive Explicit Policy Laws:** Within each cluster, the recourse decision is an explicit affine function of the uncertainty.
- **Cluster Scenarios Analytically:** Scenarios are grouped if they fall into the same CR.
- **Quantify Trade-offs:** We can explain first-stage decisions by summing the probability-weighted contributions of specific CRs.

PRELIMINARIES

Two-Stage Stochastic Programming

We consider the standard two-stage SP formulation, which can be a Linear Program (LP) or a Mixed Integer Linear Program (MILP):

$$\begin{aligned} \min_{\mathbf{x}} c^\top \mathbf{x} + \mathbb{E}_\xi[Q(\mathbf{x}, \xi)] \\ \text{s.t.} A\mathbf{x} \leq b;\ \mathbf{x} \geq 0 \end{aligned} \tag{1}$$

where $\mathbf{x}$ represents the first-stage decisions (e.g., capacity expansion). The vector c represents the cost coefficients, while A and b define the first-stage constraints. The recourse function $Q(\mathbf{x}, \xi)$ is defined as the optimal value of the second-stage problem for a specific uncertainty realization ξ:

$$\begin{aligned} Q(\mathbf{x}, \xi) = \min_{\mathbf{y}} q^\top \mathbf{y} \\ \text{s.t.} W\mathbf{y} \leq h(\xi) - T\mathbf{x};\ \mathbf{y} \geq 0 \end{aligned} \tag{2}$$

Here, $\mathbf{y}$ represents operational decisions (e.g., production rates, flows), and ξ captures uncertainties such as demand or prices. The q is the objective coefficient vector, and W (the recourse matrix), T (the technology matrix), and h (right-hand side uncertainty) define the recourse constraints.

Benders Decomposition

For large-scale problems, the two-stage SP is typically solved using a decomposition approach, such as Benders Decomposition [7], which partitions the problem into a Master problem and a set of sub-problems. The Master problem approximates the recourse cost using cuts [7]:

$$\begin{aligned} \min_{\mathbf{x},\eta} c^\top \mathbf{x} + \sum_{\omega \in \Omega} p_\omega \eta_\omega \\ \text{s.t.} \eta_\omega \geq \lambda_\omega^\top (h_\omega - T\mathbf{x})\ \forall \text{ cuts} \end{aligned} \tag{3}$$

Where p_ω is probability of scenario ω and λ are dual multipliers from the subproblem. Note that we assume relatively complete recourse holds for the problem. While efficient for computation, Benders decomposition does not inherently generate interpretable insights.

Multiparametric Programming Formulation

In our recent work [9], we reformulate the subproblem using multiparametric programming. In this work, we are exploring this method to gain interpretable insights.

In the recourse problem, the optimal second-stage decision $\mathbf{y}^*$ depends on two vectors: the fixed first-stage decision $\mathbf{x}$ and the random parameter realization ξ. We aggregate these into a single parameter vector $\theta = [\mathbf{x}, \xi]^\top$ [9]. The recourse problem becomes a multiparametric Linear Program (mp-LP) :

$$\begin{aligned} z(\theta) = \min_{\mathbf{y}} q^\top \mathbf{y} \\ \text{s.t.} W\mathbf{y} = F\theta + d;\ \mathbf{y} \geq 0 \end{aligned} \tag{4}$$

The solution to this mp-LP is not a single scalar value, but a set of critical regions. Each region CR_i is a convex polyhedron defined by a set of linear inequalities:

$$CR_i = \{\theta \in \Theta \mid E_i\theta \leq f_i\}$$

where E_i and f_i are a matrix and vector (boundary constraints of CR_i), respectively, derived from the optimality conditions (active constraints) of the subproblem. The parameter space Θ is partitioned into non-overlapping polyhedral regions $CR = \cup_i CR_i$. Within each region CR_i, the optimal decision variables y^* and the objective function z are explicit piecewise affine functions of θ:

$$\begin{aligned} \mathbf{y}^*(\theta) = H_i\theta + k_i, \forall\theta \in CR_i \\ z^*(\theta) = A_i\theta + B_i, \forall\theta \in CR_i \end{aligned} \tag{5}$$

Further details on this multiparametric-based Benders decomposition framework can be found in Brahmbhatt et al. [9]. This structure provides the "white-box" explainability: every scenario ξ maps to a specific CR_i, and that region dictates exactly which constraints are binding (e.g., "Demand constraint is active" vs "Capacity constraint is active").

METHODOLOGY

Our proposed framework integrates mp-programming into the post-optimality analysis of stochastic programming solutions.

Step 1: Parametric Solution- We treat the recourse problem as an mp-LP, considering uncertain parameters (e.g., the first stage decision $\mathbf{x}$ and demand ξ) as the varying parameters θ. We solve this using an mp algorithm [11]. This generates the set of all possible operational modes (critical regions) the system can experience under the optimal design.

Step 2: Solve the Full Stochastic Problem- We solve the original two-stage SP using either the mp Benders method [9] or a monolithic solver to obtain the optimal first-stage solution $\mathbf{x}^*$.

Step 3: Map Scenarios to Regions: Clustering- For a given optimal fixed first-stage decision $\mathbf{x}^*$, we project the full set of sampled scenarios Ω onto the generated critical regions.

$$\text{Cluster}_i = \{\, \omega \in \Omega \mid \xi_\omega \in CR_i \,\}$$

As shown in Figure 1, this mapping is treated as a point location problem. This process simply iterates through the generated regions to find the specific critical region (CR_i) where the boundary constraints are satisfied for a given scenario. Here in Figure 1, the demand of 3 different customers is uncertain in the 2nd stage. So, using parametric space for a fixed first-stage decision, we can cluster all demand scenarios in critical region clusters. Unlike k-means, which clusters based on Euclidean distance, this step clusters scenarios based on structural identity. Two scenarios are in the same cluster if and only if they share the exact same active constraints.

Step 4: Analytical Interpretation- We explain the optimal $\mathbf{x}^*$ by analyzing the probability mass accumulated in "stress" regions (regions with high penalties or recourse costs) versus "nominal" regions.

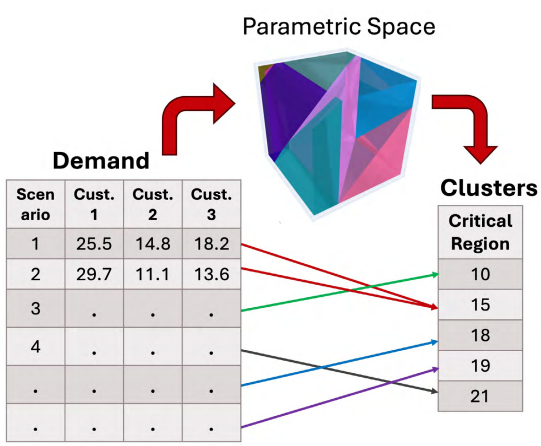

Figure 1. Recourse-based clustering of uncertain scenarios using mp critical regions.

CASE STUDY: SUPPLY CHAIN PLANNING

To demonstrate the efficacy of the proposed multiparametric framework, we utilize a canonical supply chain network design problem adapted from Rathi et al [2]. This problem represents a standard capacity planning challenge where strategic decisions must be made against uncertain customer demands.

Mathematical Formulation

We consider a network consisting of a set of manufacturing plants $i \in I$ and a set of customers $j \in J$. The problem is formulated as a two-stage stochastic linear program.

First-Stage Problem: The decision-maker determines the production quantities x_i for each plant i prior to the realization of uncertain demand. The objective is to minimize the total production cost plus the expected recourse cost. The C_i^{prod} are production costs and C_i^{max} the maximum production capacities of the plants, respectively.

$$\begin{aligned} &\min_x \sum_{i \in I} C_i^{prod} x_i + \mathbb{E}_\xi[Q(x,\xi)] \\ &\text{s.t.}\ 0 \le x_i \le C_i^{max}\ \forall\, i \in I \end{aligned} \tag{6}$$

Second-Stage (Recourse) Problem: Given the first-stage production levels x and a demand realization ξ (where $d_j(\xi)$ is the demand for customer j), the recourse problem optimizes transportation flows y_{ij} and external procurement $\bar{y}_j$. External procurement represents a penalty mechanism (e.g., outsourcing or lost sales) used when local production cannot satisfy demand.

$$\begin{aligned} Q(x,\xi) = \min_{y,\bar{y}} &\sum_{i \in I} \sum_{j \in J} C_{ij}^{trans} y_{ij} + \sum_{j \in J} C_j^{ext}\, \bar{y}_j \\ \text{s.t.} &\sum_{j \in J} y_{ij} \le x_i\ \forall i \in I\ (\text{Supply Capacity}) \\ &\sum_{i \in I} y_{ij} + \bar{y}_j \ge d_j(\xi)\ \forall j \in J\ (\text{Demand Satisfaction}) \\ &y_{ij}, \bar{y}_j \ge 0 \end{aligned} \tag{7}$$

Here, C_{ij}^{trans} is the transportation cost from plant i to customer j, and C_j^{ext} is the cost of external procurement, incentivizing satisfying the demand via internal production whenever possible.

Network Topology and Parameters

We apply this formulation to an instance consisting of two plants ($I = 2$, labeled A and B) and 3 Customers ($J = 3$, labeled 1, 2, and 3) in Figure 2.

First-Stage Parameters: The two plants differ in production efficiency. The production costs (C_A^{prod}, C_B^{prod}) are \$100/unit and \$150/unit, respectively. The maximum capacity (C_i^{max}) is 38.0 units for both plants.

Second-Stage Parameters: The unit transportation costs (C_{ij}^{trans}) from Plant A to Customers 1, 2, and 3 are \$90, \$67.5, and \$90, respectively. Similarly, the costs from Plant B to Customers 1, 2, and 3 are \$67.5, \$30, and \$67.5, respectively. The external procurement cost (C^{ext}) is set at \$500 per unit for each customer. External procurement is set significantly higher than production plus transport to penalize shortages.

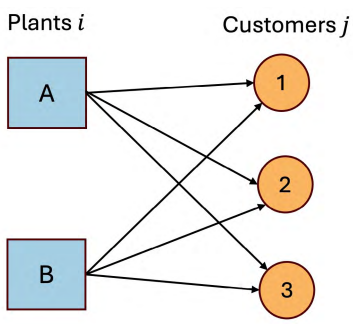

Figure 2. The network topology for 2 plants and 3 customers.

Uncertainty Modeling: The demand $d_j(\xi)$ for each of the three customers is uncertain. We approximate the uncertainty using a sample average approximation (SAA) approach with $N = 100$ scenarios. The demand for each customer j follows a uniform distribution $\mathrm{U}(10, 30)$.

The resulting problem has 2 first-stage variables, 900 second-stage variables (9 per scenario $\times$ 100 scenarios), and 700 constraints. While solvable by standard monolithic solvers, the focus of this study is on analyzing the solution structure using the multiparametric decomposition proposed in [9] and summarized in the previous section.

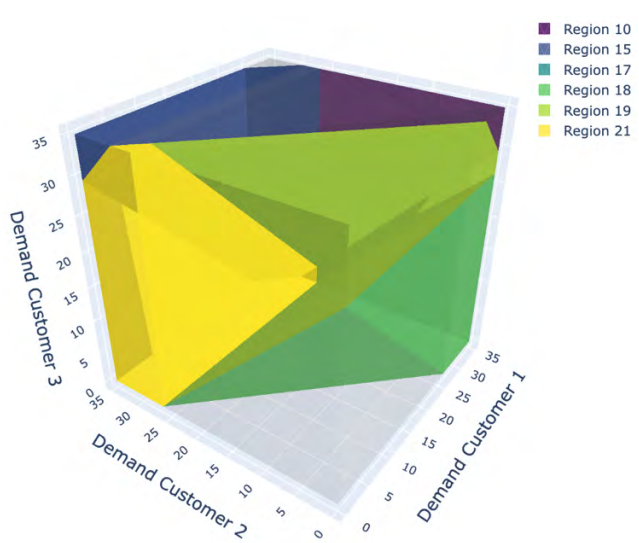

Figure 3. Critical regions in demand space given fixed first-stage decisions (each color represents a different critical region).

RESULTS AND DISCUSSION

The optimal first-stage solution found is Plant A, having a capacity of 38.00 units, and Plant B, having a capacity of 26.41 units, for a total production capacity of 64.41 units.

Critical Region Clusters Analysis

The multiparametric analysis was implemented in Python using the PPOPT solver [10], and it resulted in 23 critical regions. For the fixed first stage solution (x^*), Figure 3 shows the partitioned 3D demand space (Customer 1, 2, 3) with critical regions. Each color in Figure 3 represents a different critical region. While the statistical approach, like k-means, required selecting an arbitrary number of clusters (say k=5 or 6), our analytical approach identified exactly 6 populated critical regions that cover all 100 scenarios. As shown in Figure 1, the proposed clustering method uses the point location problem to find the critical region ID. This is what we call recourse-based clustering, which is shown in Figure 4. Crucially, it demonstrates that while the uncertainty space is continuous, the corresponding optimal recourse strategy is strictly piecewise affine and governed by exactly six distinct critical regions, each defined by a unique set of active constraints in the recourse (sub-problem) decision space.

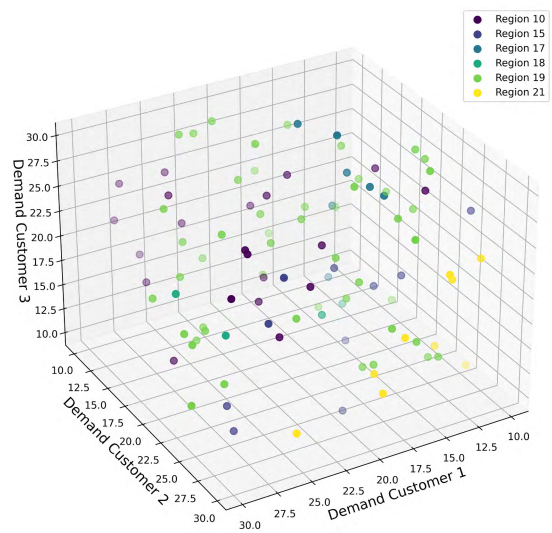

Figure 4. 3D Visualization of the 100 demand scenarios mapped to 6 Critical Regions. The recourse clustering is defined by polyhedral boundaries (active constraints) rather than simple distance.

A comparison of our recourse-based clustering with scenario realization-based clustering is presented in Figure 5, which shows the 100 scenarios clustered using the k-means clustering method, indicating that scenarios of the same type were grouped together within the same

cluster ID. But based on recourse-based clustering using mp critical region, we can say that scenarios that are similar in their uncertainty realizations may not lead to the same optimal recourse decisions. Figure 6 shows the clusters of recourse decision for all 100 scenarios based on the mp critical regions. Within the same critical region clusters, the scenarios have the same recourse decisions.

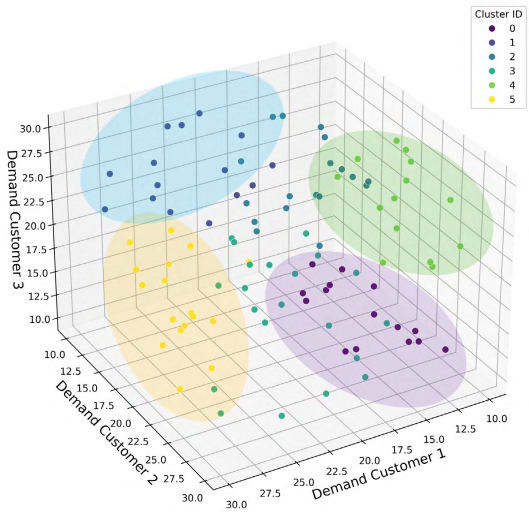

Figure 5. 3D Visualization of the 100 demand scenarios mapped to 6 clusters using k-means clustering.

Operational Modes (Clusters)

Figure 7 shows the average recourse decisions over different critical region clusters. The size of the square shows the production level, which is fixed for all scenarios. The size of the circle shows the average demand in each critical region cluster. As we can see Region 19 (Probability: 0.49) is the "Nominal Operation" mode. It contains nearly half of all scenarios. In this region, total demand is moderate, and the production from A and B is sufficient. No external procurement is required. Region 21 (Probability: 0.09): A transition region where demand is slightly higher, shifting transportation routes, but still manageable without external purchase. Regions 10 & 15 (Probability: 0.21 & 0.10): These are "Failure/Recourse" modes. In these regions, high simultaneous demand on the customer's side exceeds the available inventory (64.41 units). The system hits the active constraint $\sum y \leq x$, forcing the variable $\bar{y}$ (external procurement) to become non-zero. The average amount of this outside procurement is explicitly indicated by the red numbers next to the customer nodes in Figure 7. The probability here is calculated by the number of scenarios that got clustered in each specific region over the total scenarios.

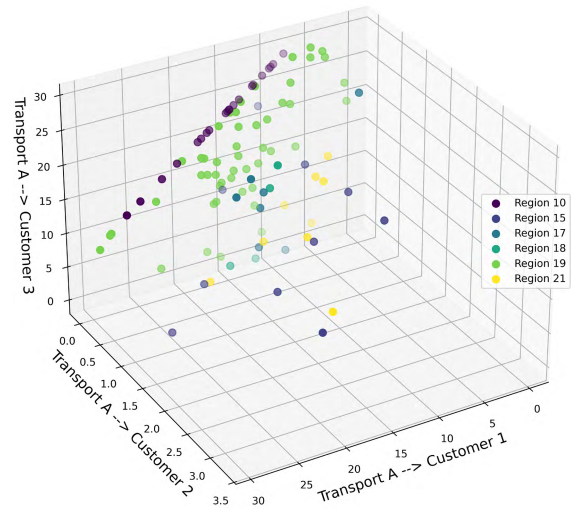

Figure 6. The recourse decisions of all 100 scenarios got clustered in 6 mp clusters. The 3 axes represent the transportation of the product from Plant A to customers 1, 2, and 3, respectively.

Explaining the Capacity Decision

The central explainability question for this SP solution is: Why is the optimal production at Plant B set to 26.41 units instead of its maximum 38.0 units? A human planner might argue: "We are hitting penalty costs in Regions 10 and 15 (31% of the time). Why not increase production at Plant B to avoid this?" Our mp based framework allows us to answer this with a precise cost-benefit calculation. As shown in Figure 8, for the higher demand for all customer cases (Regions 10 and 15), we have very high second-stage costs. But these 2 clusters (Region 10 and 15) have an overall low probability; they make a very low contribution to the overall second-stage cost. Current external procurement cost contribution from Regions 10 and 15 is approximately $2220 (weighted by probability). While increasing Plant B capacity and transportation of product to customers to cover these scenarios would cost roughly $3319 in additional first-stage investment.

It is mathematically cheaper to accept the "failure" (external procurement) in 31% of scenarios falling into Regions 10/15 than to pay for the capacity to cover them. This transforms the result from a simple number into a clear economic decision. The model intentionally incurs penalties in these specific high-demand regions rather than over-investing in capacity.

While the analysis above explains why we do not increase capacity to meet peak demand in Regions 10 and 15, a symmetrical question arises: Why is the optimal capacity for Plant B set to 26.41 units and not lower? The multiparametric framework allows us to answer this by analyzing the "nominal" operating modes, specifically Region 19 and Region 21. These regions have a combined probability of 0.58 and average demands (over all customers) of 56.51 and 60.10 units, respectively. Currently,

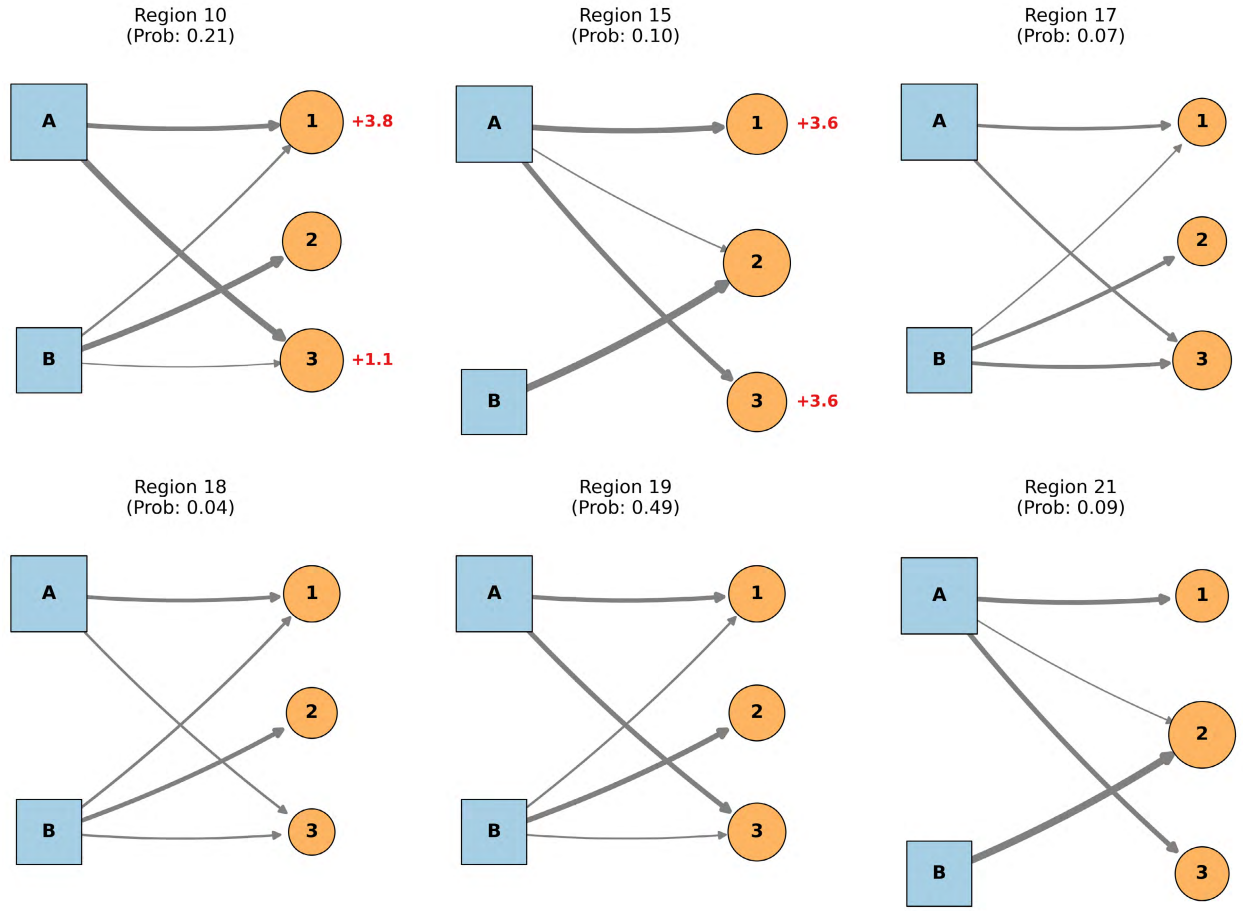

Figure 7. Visualization of recourse decisions averaged across the six Critical Region clusters. The square nodes represent manufacturing plants, where the size corresponds to the fixed production capacity. The circular nodes represent customers, with their size proportional to the average demand within that specific cluster. The red numbers adjacent to customer nodes indicate the specific quantity of external procurement required to satisfy

the total system capacity of 64.41 units is sufficient to satisfy these demands without external procurement. However, if we were to further reduce Plant B's capacity, we would force external procurement in these highly probable regions. Reducing capacity by one unit saves the guaranteed investment cost but incurs an expected penalty of $445.0 in additional external procurement costs across these clusters. This far exceeds the marginal cost of $182.32 to produce and transport that unit internally from Plant B. Thus, the optimal capacity of 26.41 units represents the precise equilibrium where the marginal savings of capacity reduction are outweighed by the expected penalties in the dominant operating regions.

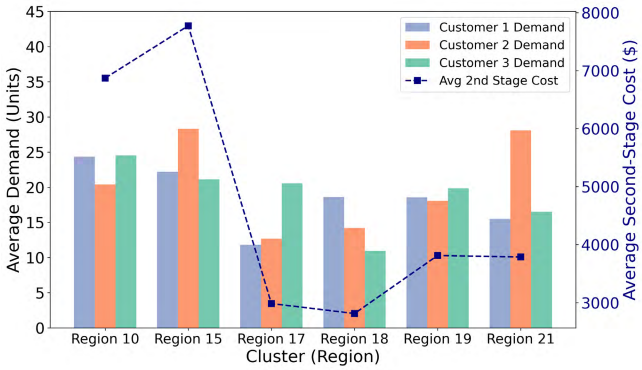

Figure 8. Average demand and average second-stage cost in different clusters.

CONCLUSION

We presented a methodology for enhancing the interpretability of stochastic programming solutions using multiparametric programming inspired by Rathi et al. [2].

By using the Benders subproblem to generate explicit critical regions, we achieve an objective, model-driven clustering of scenarios. Notably, while the presented case study demonstrated an LP formulation, this framework is fully applicable to problems involving Mixed-Integer Linear Programming (MILP) first-stage decisions.

Applied to a supply chain case study, this method successfully reduced 100 scenarios into 6 logical operating modes. It allowed us to explain the optimal first-stage production levels by rigorously quantifying the trade-off between guaranteed production costs and the expected recourse costs associated with specific critical regions. While this approach is effective, a key limitation is its scalability with respect to problem size. Specifically, as the number of uncertain parameters increases, the multiparametric solution becomes computationally intensive because the total number of critical regions grows combinatorially. Future work will focus on alternative mp algorithms, which generate only the relevant Critical Regions encountered by the Benders cuts, rather than exploring the entire parameter space, enabling scalability to larger networks.

ACKNOWLEDGEMENTS

This work is based upon work supported by the National Science Foundation under grant no. CMMI-2328160.

Furthermore, we would like to acknowledge the assistance provided by large language models (LLMs), including OpenAI's ChatGPT and Google's Gemini, in the preparation of this manuscript. These tools were used to improve the grammar and fluency of the text, as well as to assist with coding and commenting.

AUTHOR IDENTIFIERS

Author ORCIDs:
Brahmbhatt P: 0009-0006-1502-1359
Avraamidou S: 0000-0002-9334-9951

REFERENCES

1. Shapiro A, Dentcheva D, Ruszczynski A. Lectures on stochastic programming: modeling and theory, third edition. Society for Industrial and Applied Mathematics (2021). https://doi.org/10.1137/1.9781611976595
2. Rathi T, Gupta R, Pinto JM, Zhang Q. Enhancing explainability of stochastic programming solutions via scenario and recourse reduction. Optim Eng 25:795-820 (2023). https://doi.org/10.1007/s11081-023-09825-9
3. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell 1:206-215 (2019). https://doi.org/10.1038/s42256-019-0048-x
4. Bertsimas D, Stellato B. The voice of optimization. Mach Learn 110:249-277 (2020). https://doi.org/10.1007/s10994-020-05893-5
5. Medina-González S, Gkioulekas I, Dua V, Papageorgiou LG. A graph theory approach for scenario aggregation for stochastic optimisation. Computers & Chemical Engineering 137:106810 (2020). https://doi.org/10.1016/j.compchemeng.2020.106810
6. Beraldi P, Bruni ME. A clustering approach for scenario tree reduction: an application to a stochastic programming portfolio optimization problem. TOP 22:934-949 (2013). https://doi.org/10.1007/s11750-013-0305-9
7. Rahmaniani R, Crainic TG, Gendreau M, Rei W. The benders decomposition algorithm: a literature review. European Journal of Operational Research 259:801-817 (2017). https://doi.org/10.1016/j.ejor.2016.12.005
8. Pappas I, Kenefake D, Burnak B, Avraamidou S, Ganesh HS, Katz J, Diangelakis NA, Pistikopoulos EN. Multiparametric programming in process systems engineering: recent developments and path forward. Front. Chem. Eng. 2: (2021). https://doi.org/10.3389/fceng.2020.620168
9. Brahmbhatt P, Cole DL, Zavala VM, Avraamidou S. Benders decomposition using graph modeling and multi-parametric programming. Ind. Eng. Chem. Res. 64:21684-21700 (2025). https://doi.org/10.1021/acs.iecr.5c03189
10. Kenefake D, Pistikopoulos EN. Ppopt-multiparametric solver for explicit MPC. Comput Aided Chem Eng 51:1273-1278 (2022) https://doi.org/10.1016/B978-0-323-95879-0.50213-7
11. Bemporad A, Morari M, Dua V, Pistikopoulos EN. The explicit linear quadratic regulator for constrained systems. Automatica 38:3-20 (2002). https://doi.org/10.1016/s0005-1098(01)00174-1

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Exploring Robust Early-Stage Decisions in Energy Transitions Using Near-Optimal Pathways and Multi-Armed Bandits

Mahdi Kchaou[a]*, Diederik Coppitters[a], Francesco Contino[a]

[a] Institute of Mechanics, Materials and Civil Engineering (iMMC), Université catholique de Louvain (UCLouvain), Louvain-la-Neuve, Belgium
* Corresponding Author: mahdi.kchaou@uclouvain.be.

ABSTRACT

Although rare, unexpected events such as financial crises, geopolitical conflicts, and pandemics have reshaped reality in recent years. Despite their strong potential to affect the energy transition, such events are still largely overlooked in energy planning studies. Ignoring them can lead to poorly informed decisions that may jeopardize the transition. Identifying early-stage decisions that remain robust under unexpected events is therefore essential. To address this challenge, EnergyScope Pathway, a whole-energy system model with limited foresight, is applied to Belgium. To increase the likelihood of a successful transition, the Modeling to Generate Alternatives approach is used to diversify early-stage decisions in 2035. These alternatives are allowed to be up to 10% more expensive than the cost-optimal solution. However, the large number of alternative designs is difficult to navigate for decision makers. To address this, a decision-support framework based on the Multi-Armed Bandit framework is used to identify early-stage decisions that are most robust to future unexpected events. In this step, the remaining transition phases are optimized under unexpected events sampled within predefined impact ranges. The results show that, under normal conditions, there is a high degree of flexibility in the decision space for the 2030–2035 phase, with many technologies or resources that can be entirely omitted. However, robust early-stage decisions rely on a diverse energy generation portfolio, with a stronger emphasis on wind deployment, early mobility shifts toward battery electric vehicles, and the import of e-fuels. These insights can help decision makers steer the energy transition toward a robust path from the beginning. While Belgium is used as a case study, this framework is transferable to other contexts.

Keywords: Energy transition, Modeling to generate alternatives, Multi-armed bandit, Decision-making, Unexpected events

1. INTRODUCTION

The need for an energy transition has become increasingly urgent due to the impacts of climate change on the economy, food security, and human health [1]. To guide this transition, energy system modeling has emerged as a key tool for shaping policies and investment decisions. Despite the large number of studies on this topic, progress in the energy transition has been hindered by hesitation and inaction, largely driven by the many uncertainties surrounding energy projects [2]. Many existing energy planning studies ignore uncertainty altogether [3], making their solutions highly sensitive to inaccuracies in future conditions and ultimately prone to failure. A striking example is the collapse of General Electric's thermal power business, which lost approximately $193 billion between 2016 and 2018 due to incorrect assumptions favoring gas over renewables [4]. In the rare cases where uncertainty is considered, most models rely on local sensitivity analysis or probability-based approaches [5]. While useful, these methods fail to capture deeper and more complex forms of uncertainty, particularly those that are rare in nature but have significant consequences: unexpected events. These are largely mispredicted or unpredicted occurrences that deviate strongly from expectations and persist long enough to

influence strategic decisions [6]. Such events have already affected the energy sector in the past and are likely to do so again [7, 8]. They may arise from technological breakthroughs, supply-chain disruptions, global pandemics, geopolitical conflicts, or societal and political shifts. The COVID-19 pandemic, for instance, reshaped the global energy sector by altering fuel demand, accelerating the deployment of renewables, and highlighting vulnerabilities in global supply chains [9].

Despite their importance, unexpected events are poorly represented in energy planning tools [6]. Their low probability and the scarcity of data needed to characterize them make them difficult to incorporate into standard uncertainty analysis methods.

This paper addresses these challenges by proposing a novel framework to identify early-stage energy transition decisions that perform well under unexpected events.

The whole-energy, regional energy model EnergyScope Pathway [10] is applied to Belgium using a multistage formulation with limited foresight, ensuring that events occurring in later stages are not anticipated from the beginning but instead arise as surprises. The model optimizes the energy transition by minimizing costs, subject to constraints that ensure a carbon-neutral Belgium by 2050. The decision space is then expanded using the Modeling to Generate Alternatives (MGA) approach [11], which is applied to the early transition phase from 2030 to 2035 to generate a diverse set of near-optimal early-stage decisions. These decisions, such as not importing e-fuels or not relying on nuclear energy, increase the likelihood of a successful transition by providing more diverse starting conditions. Afterwards, these decisions are evaluated under unexpected events occurring in the remaining stages of the transition, from 2035 to 2050. Finally, a decision-making framework based on the Multi-Armed Bandit (MAB) algorithm [12] is used to identify the most robust early-stage decisions in 2035. Although the Belgian energy transition is used as a case study, the proposed framework is highly transferable and can be applied to other contexts, both within the energy sector and in other planning applications.

In this paper, we first present the EnergyScope Pathway model, the MGA approach, the generation of unexpected events, and the MAB approach (Section 2). The results section (Section 3) begins by illustrating the diversity of the near-optimal space, followed by the selection of representative early-stage decisions. The robust early-stage decisions are then identified using the MAB approach. Section 4 concludes the paper by summarizing the main takeaways.

2. METHODS

2.1 EnergyScope Pathway

The energy model used in this work is EnergyScope Pathway [10], a regional, whole-energy system model applied to the Belgian case. It covers most energy sectors, including electricity, heating, transport, and non-energy demand. It includes six demand categories and 112 technologies covering energy supply, storage and conversion, as well as import and export interactions with neighbouring countries. The model is formulated as a linear multistage optimization problem, with the objective of minimizing the total transition cost. In the pathway (multistage) formulation, the energy transition is divided into multiple phases, each with its own parameters and variables. The results of one phase are passed to the next as initial conditions, allowing the model to capture the evolution of investment decisions over time. The model can be used either with an hourly resolution, using a typical days approach [13], or with a monthly resolution, where each month is represented by a single time point. While the monthly formulation offers faster solution times, it is less detailed, as renewable intermittency and hourly demand variations are not explicitly captured. As a result, some technologies, such as storage technologies, are not included.

EnergyScope Pathway supports both perfect foresight and myopic foresight formulations. In the perfect foresight case, the entire transition is optimized in a single run, assuming full knowledge of future conditions, and decisions for all phases are optimized simultaneously. A global CO_2 emissions budget is imposed, allowing any decarbonization pathway as long as the final emissions target is met by the end of the transition. In the myopic foresight formulation, each phase of the transition is optimized on its own, sequentially, without knowledge of future conditions. Since each phase is optimized separately, a global CO_2 emissions constraint cannot be applied. Instead, a specific CO_2 emissions constraint is imposed for each one. These constraints follow an exponential decay formulation [14]:

$$e(t) = e_0(1 + (r + m)t)e^{-mt} \tag{1}$$

where e_0 124 MtCO2/yr) is the initial emission level, r is the initial slope parameter (set to zero), and m is the decay parameter. The parameter m is calibrated such that the trapezoidal approximation of emissions, evaluated in five-year steps between 2020 and 2050, matches the cumulative emissions budget 1.5 GtCO2 [15]. The resulting emissions levels at each five-year node are then used as constraints in the myopic optimization. In this work, the monthly and myopic options are chosen.

Additional details on the model can be found in [6, 10].

2.2 Early stage decisions

EnergyScope Pathway optimizes the energy transition by minimizing costs, delivering a single cost-optimal solution. While informative, this approach is not enough,

as it lacks the diversity of solutions required for decision makers with their different preferences. Moreover, relying on a single solution is risky, as it may not perform well under structural uncertainties, particularly high-impact ones such as unexpected events. For instance, under normal conditions, a maximum deployment level may be assumed for a given technology. The cost-optimal pathway will then size its PV deployment to this maximum in the early stages of the transition. Under unexpected events, however, this maximum deployment could be reduced by half, leaving the optimal solution unable to cope with such situation. To address this limitation, we rely on the Modeling to Generate Alternatives (MGA) [11] approach to explore the near-optimal space of early-stage decisions for the transition phase 2030–2035. MGA allows a small increase (slack) in cost in exchange for greater solution diversity. In this work, a 10% slack (cost increase) is allowed during the 2030–2035 phase relative to the cost of that phase, and all solutions satisfying this condition are identified as near-optimal. As a result, by 2035, rather than a single optimal energy system, a set of diverse energy systems with near-equivalent economic performance is identified. Considering such alternative solutions can lead to better performance when unexpected events occur. These alternatives may include, for example, energy systems with lower reliance on PV, reduced fossil fuels imports to accelerate early decarbonization, or higher e-fuels imports during the initial transition stages. These choices, among many others, provide policymakers with a broader decision space, rather than locking decisions into a single pathway. The focus on the 2030–2035 period is motivated by the need to provide actionable guidance for decisions that need to be taken today or in the near future, helping to place the energy transition on the right track from the beginning.

Another key element in defining the MGA problem is the choice of the MGA variables. These variables are allowed to vary within the feasible space to explore the near-optimal solutions. Ideally, the MGA variables would include all investment decisions. However, due to computational limitations, only a limited number of variables can be considered. In this work, the selected MGA variables are: battery electric vehicles (BEV) [Mpkm/h], nuclear capacity [GW], PV capacity [GW], wind capacity (the sum of onshore and offshore wind) [GW], e-fuels imports (the sum of imported renewable methanol, ammonia, and hydrogen) [GWh], and fossil resource imports (the sum of imported coal, diesel, and gas) [GWh]. These variables correspond to the strategic decisions that governments and policymakers can directly influence.

Using the MGA approach, the near-optimal space is defined as:

$$W = \{ x \in X \mid f(x) \leq f(x^{\wedge} *)(1 + \varepsilon) \} \quad (2)$$

where X is the feasible solution space, ε is the slack parameter, $f(x)$ denotes the objective function, and x^* is the optimal solution. In this study, $f(x)$ corresponds to the cumulative cost of the transition up to 2035, and ε is set to 10%. As a result, all solutions that are at most 10% more expensive than the minimum-cost solution are accepted. The exploration of the near-optimal space follows the Modeling All Alternatives (MAA) approach developed by Pedersen et al. [16]. This method maps the entire near-optimal space by defining its convex hull, which is the smallest polyhedron containing all near-optimal solutions. Since the original problem formulation is linear, the near-optimal space is guaranteed to form a polyhedron. The method starts with an initial set of near-optimal solutions, from which a first convex hull is constructed using the Quickhull algorithm. In the next iteration, new near-optimal points are generated to increase the volume of the convex hull. Once the convex hull is fully defined, it is then sampled since all points within it are valid near-optimal solutions. After generating a large set of near-optimal solutions, using all of them as early-stage decisions becomes computationally expensive. Therefore, a representative subset is selected using the k-means clustering approach, which groups solutions with similar characteristics. The centroid of each cluster is chosen as a representative solution, and the number of clusters is determined using the elbow method [17].

2.3 Unexpected events

Following the work of Coppitters et al. [6], six categories of unexpected events are identified. The first category concerns the availability of imported resources, including electrofuels, biofuels, and electricity. The second category is resistance to change, which includes PV and wind deployment, low-temperature heating renovations, and private mobility changes. The third category relates to exchange-rate deterioration, and the fourth to increases in energy demand. The fifth category concerns unicorn technologies, including small modular reactors (SMR), carbon capture and storage (CCS), geothermal power, ammonia-fired combined cycle gas turbines (CCGT), and ammonia cracking. The final category is the nuclear phase-out. For the first four categories, each subcategory is characterized by four parameters, each corresponding to a specific year of impact: 2035, 2040, 2045, and 2050. Each parameter can take one of six values: 0, 20, 40, 60, 80, or 100, where 0 indicates no impact and 100 represents full impact in the corresponding year. For the fifth category, each subcategory is characterized by a single parameter indicating when the technology becomes available: 2040, 2045, 2050, or after that. Similarly, the nuclear phase-out is represented by a single parameter indicating when the phase-out occurs: 2035, 2040, 2045, 2050, or after that.

From this parameter space, 1000 unexpected-event scenarios are sampled using Sobol sampling to ensure

uniform coverage of the space. For further details on the selection, parametrization, and generation of unexpected events, the reader is referred to [6].

2.4 Multi-armed bandit

After generating a set of early-stage decisions for the year 2035 using MGA (Subsection 2.2), the objective is to identify the most robust ones of these decisions under the unexpected events defined in Subsection 2.3. Policymakers can take one of these decisions in the near future. However, evaluating all possible combinations of early-stage decisions and unexpected event scenarios is computationally expensive, even for simplified models such as the monthly model used in this work. To address this computational challenge, a reinforcement learning approach based on the Multi-Armed Bandit (MAB) algorithm is used [12].

The MAB problem can be illustrated by a gambler playing slot machines with multiple arms, each returning random rewards. At each time step, the gambler pulls an arm and gets a reward. Based on the history of rewards, the MAB algorithm suggests which arm to pull next in order to maximize the expected reward. In the context of this work, each early-stage decision for the year 2035 represents an arm, and the reward corresponds to the final outcome of the energy transition. Each time a specific early-stage decision is selected, the energy system in 2035 is fixed according to the decision, and the remainder of the transition is optimized under a randomly selected unexpected event scenario drawn from the 1000 generated scenarios. This is where the uncertainty in the reward arises. By balancing exploration and exploitation, the MAB approach identifies the most robust early-stage decisions without the need to evaluate all possible scenario–decision combinations.

In this MAB problem, the Upper Confidence Bound (UCB) policy [12] is used. The policy defines the logic behind the selection of the next arm at each time step. Let K denote the number of arms. At each time step t, an arm $a_t \in \{1, \dots, K\}$ is pulled, resulting in a random reward r_t. The UCB policy defines an objective that combines exploitation, through the use of empirical mean of the previous rewards, and exploration, through an explicit uncertainty term that accounts for how often each arm has been pulled. This objective is maximized to select the next arm a_{t+1} to pull.
The reward is typically defined within the interval [0, 1]. For that, the reward in this work is defined as follows:

$$r_t = \begin{cases} \frac{C_{\text{optimal}}}{C_{\text{scenario}}}, & \text{if the CO_2 target is met by 2050} \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where C_{optimal} is the minimum (optimal) cost, under normal conditions and C_{scenario} is the cost of the scenario defined by the combination of an unexpected event scenario and the selected early-stage decision.

A summary of the methodology is illustrated in Figure 1.

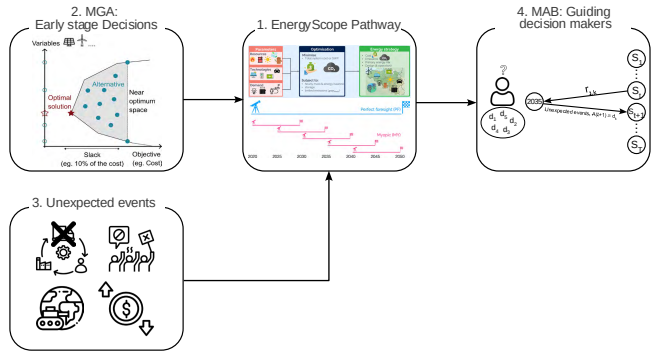

Figure 1: The methodology starts by setting up the energy transition model (1). Then, a set of alternative early-stage decisions for the year 2035 is generated (2), along with a set of unexpected-event scenarios (3). A MAB approach (4) is used to test the early-stage decisions against these unexpected events and identify the most robust ones.

3. RESULTS

3.1 Near-optimal space

The optimal early-stage decision, as well as the near-optimal decisions for the year 2035, are presented in Figure 2. The figure shows the probability mass functions (PMFs) of the near-optimal space together with the optimal values, indicated by red stars, for the MGA variables: BEV, nuclear capacity, PV capacity, wind capacity, e-fuels imports, and fossil resource imports.

The optimal solution is characterized by high nuclear and wind deployment, and fossil resource imports, medium level of integration for BEV and low PV deployment and e-fuels imports, as seen in the range of the near-optimal space. The near-optimal alternatives themselves span a wide range of values, indicating a high degree of flexibility and the existence of solutions that are structurally very different from the optimal one. For instance, most variables can reach zero or near-zero values, meaning that their corresponding technologies or resources can be entirely omitted (for example, some near-optimal systems have no wind deployment or almost no fossil fuels imports) while remaining at most 10% more expensive than the optimal system by 2035. However, the near-optimal solutions are not evenly distributed across this wide range. Looking at the PMFs of wind, PV, and nuclear capacity, distributions are skewed toward higher values, suggesting that higher deployment levels are more frequently favored among near-optimal solutions. This tendency does not always coincide with the optimal value, as seen for PV capacity. In contrast, the BEV PMF is normally distributed around the optimal value, indicating that the optimal BEV level is the most

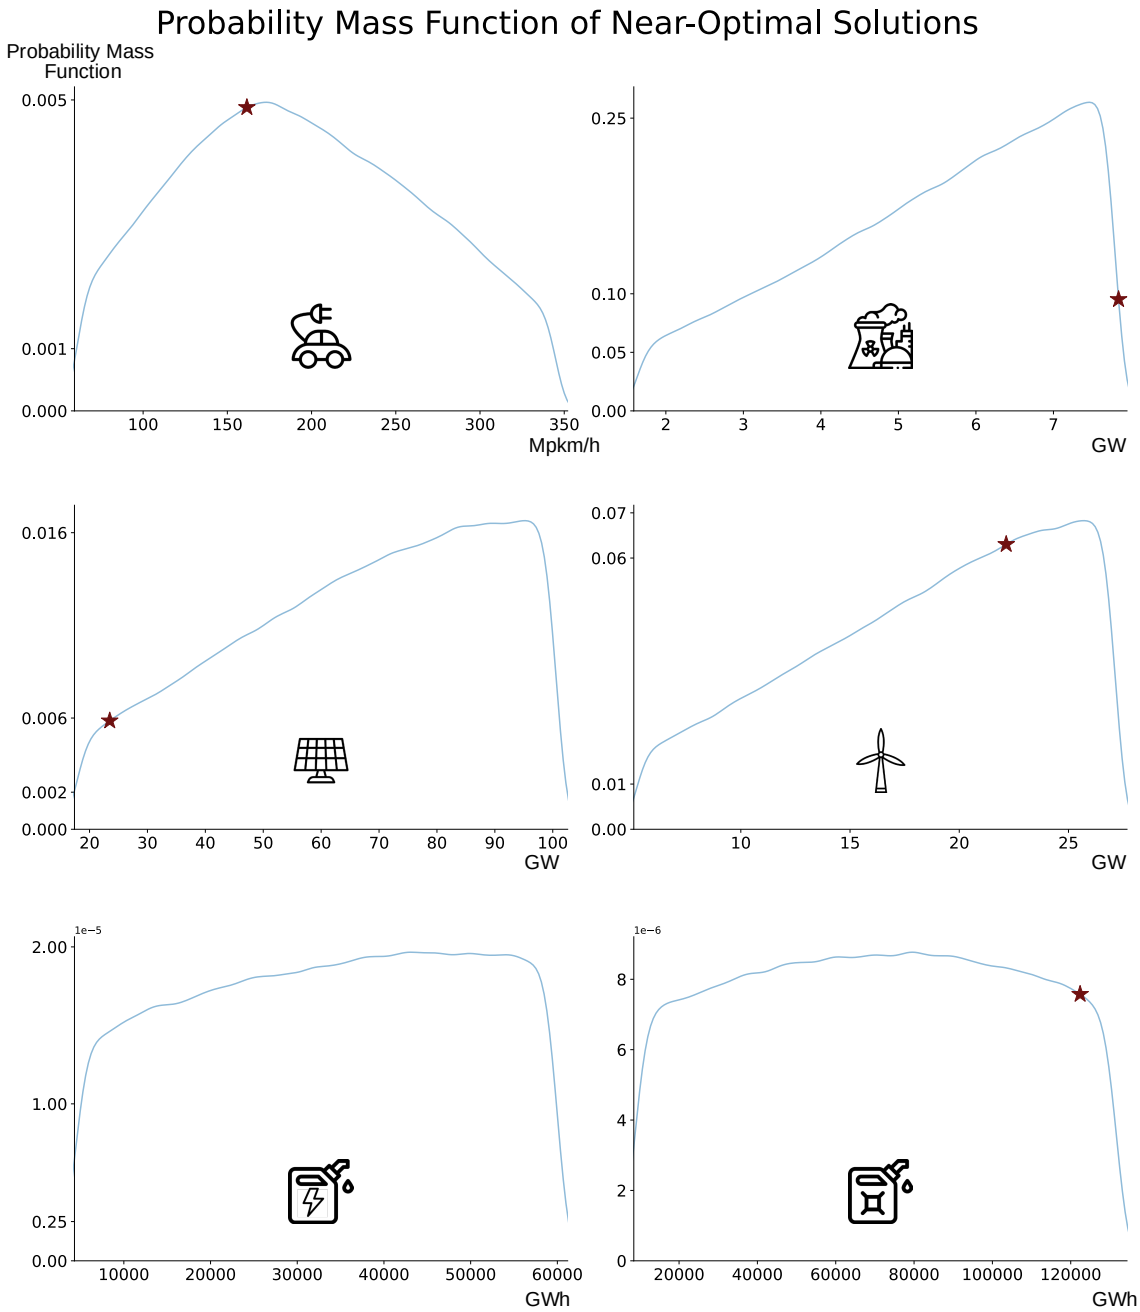

Figure 2. The probability mass function plots of the near-optimal space show a wide variety of feasible solutions. Despite this large space, solutions are not evenly distributed across the full range of values. PV, wind, and nuclear deployment exhibit strong skewness toward higher values. BEV is normally distributed around its optimal value (indicated by the red star), whereas fuels imports are more uniformly distributed across the near-optimal space.

favored option. Finally, for fuels imports, both distributions are close to uniform, with a slight skew toward higher values for e-fuels imports, indicating that near-optimal solutions are evenly spread across these dimensions of the decision space. At this stage, it should be kept in mind that this flexibility concerns only early-stage decisions specific to the year 2035 of the transition.

3.2 Early stage decisions selection

The MAA method generates thousands of near-optimal early-stage decisions for the year 2035. However, retaining the full set of solutions to study their robustness would be computationally expensive, as discussed in Section 2.2. To address this, the near-optimal space is reduced using representative solutions, selected as the

centroids of clusters obtained through k-means clustering. The elbow method suggested dividing the space into 5 to 10 clusters: for better interpretability, the space was finally clustered into 5 clusters.

Figure 3 illustrates the resulting heatmap of early-stage decisions. All MGA variables are normalized to the [0, 1] interval, where a value of 0 represents the absence of the technology, while 1 corresponds to its maximum level. The optimal solution is shown in the first row, followed by the five cluster centroids that characterize the near-optimal solution space. Across all early-stage decisions, high deployment of wind, solar, and nuclear technologies is a consistent feature, with at least two of these technologies reaching high levels in every case. BEV integration is generally moderate, with two exceptions: a low level in decision 2 and a slightly higher level in decision 5. Fossil fuels imports are similar across decisions, with mostly mid-level imports and higher imports in the optimal solution. In contrast, e-fuels imports are generally moderate, except in the optimal solution, where they are comparatively low.

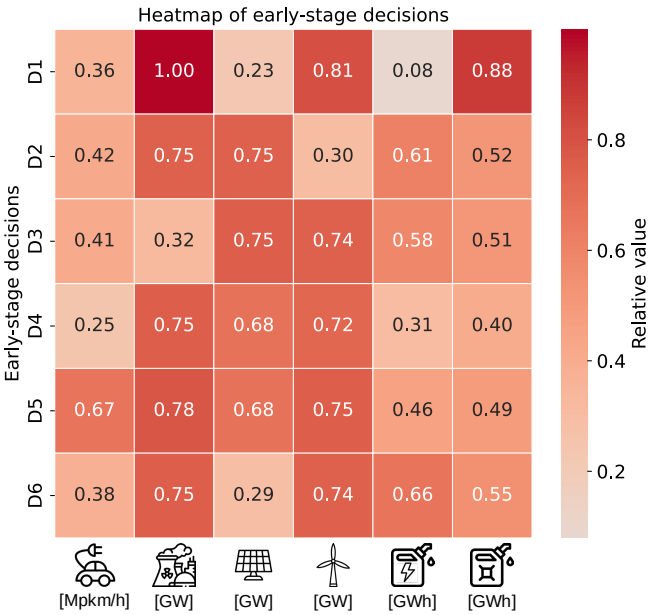

Figure 3. Heatmap of relative MGA variable values ([0, 1]) for the optimal solution and the five cluster centroids representing the near-optimal space. The figure shows consistent high deployment of wind, solar, and nuclear technologies across early-stage decisions, while BEV integration remains mostly moderate. Fossil fuels imports are present in all cases, with higher levels in the optimal solution, whereas e-fuels imports are generally moderate and lower in the optimal case.

The near-optimal early-stage decisions (decisions 2 to 6 in Figure 3) are then evaluated under normal conditions, without unexpected events. In 2035, the energy system is fixed according to each of these early-stage decisions, and the remainder of the transition is optimized under a cost-minimization objective. All of the decisions lead to a successful transition, in the sense that the 2050 CO_2 target is met in all cases. However, the associated total transition costs (all of the transition, up until 2050) differ. As shown in Figure 4, all decisions result in higher total costs than the total optimal solution (red dashed line). Nevertheless, total costs remain within the 10% slack of the total optimal solution (grey dashed line): below 110% of the minimum total transition cost. In terms of specific values, decisions 2, 3, and 4 have similar total costs, around 960 b€. Decision 5 leads to the highest total cost (≈ 968 b€), while decision 6 leads to the lowest total cost (≈ 943 b€).

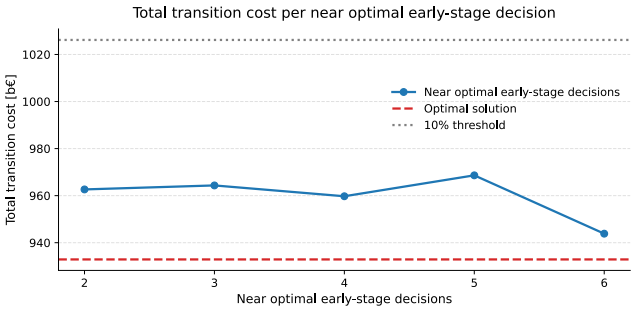

Figure 4. Total transition costs associated with the near-optimal early-stage decisions, compared to the optimal solution (red dashed line), under normal conditions. All near-optimal decisions result in higher total costs than the optimal case, while remaining below 110% of the minimum total transition cost (grey dashed line). Decision 2, 3 and 4 have similar costs, while decision 5 has the highest cost and decision 6 has the lowest cost.

3.3 Robust early stage decisions

The multi-armed bandit (MAB) method is used to identify the most robust early-stage decisions. The algorithm is run for 1000 time steps. At each step, a decision is selected (pulled) based on the UCB policy, and a random unexpected events scenario is applied starting from 2035. Figure 5 shows the cumulative number of pulls for each decision over time. Decision 6 clearly stands out, as it is selected most often, followed by decisions 3 and 5. In contrast, decisions 1 (the optimal decision for the 2030-2035 phase), 4, and 2 are selected least frequently. Similar conclusions can be drawn from the evolution of the average reward per decision, shown in Figure 6. Decision 6 reaches the highest average reward (≈ 0.49) (Eq. 3), followed by decisions 3 and 5 with average rewards of about 0.41. The weakest performing decisions are decisions 1, 4, and 2, with average rewards of approximately 0.35, 0.28, and 0.20, respectively.

One important observation is that all early-stage decisions share a common feature: they all have a mixed energy generation portfolio. However, what separates

well-performing decisions from poorly performing ones is the specific combination of technologies and resources that is chosen. When comparing decisions 2 and 3, both share similar characteristics, but decision 2 relies more on nuclear power, while decision 3 relies more on wind power. The better performance of decision 3 suggests that wind deployment provides greater robustness under uncertainty than relying mainly on nuclear power. Nuclear capacity can be entirely phased out during the transition, meaning that all previously installed nuclear capacity may eventually be dismantled. In contrast, even under resistance to further wind deployment, wind capacity installed earlier in the transition continues to be used.

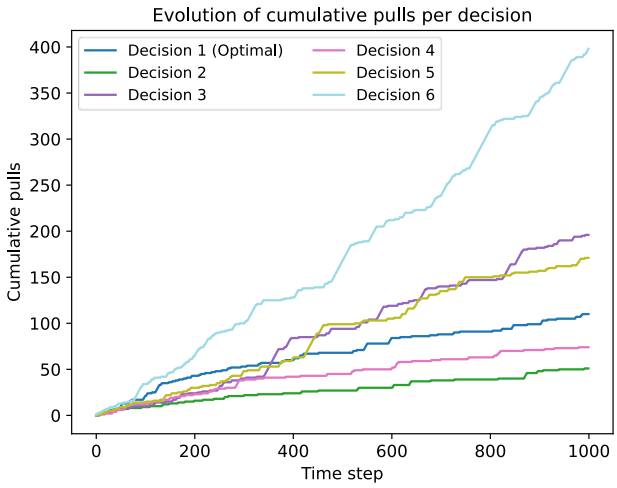

Figure 5. Cumulative number of pulls per decision over time. Decision 6 is selected most often, indicating the highest robustness, followed by decisions 5 and 3. Decisions 1, 4, and 2 are selected least frequently.

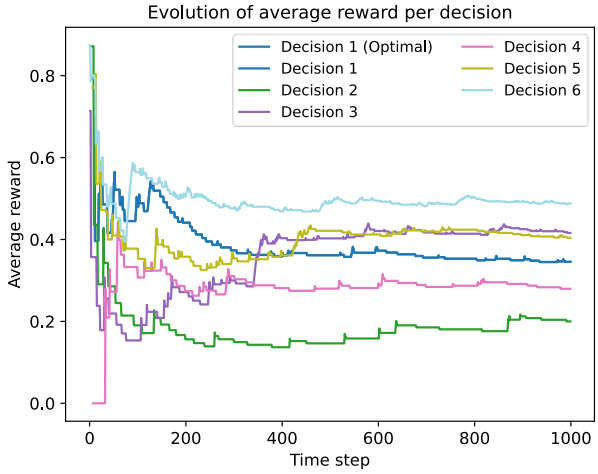

Figure 6. Evolution of the average reward per decision over time. Similar to the previous figure, decision 6 performs best, achieving the highest average reward over most time steps, followed by decisions 5 and 3. Decisions 1, 4, and 2 have the lowest final average rewards.

A similar pattern appears when comparing decisions 4 and 5. Although both have a diverse generation mix, decision 5 performs better because it includes higher BEV integration and higher e-fuels imports. This shows that having a mixed generation portfolio alone is not sufficient: early changes in mobility and use of e-fuels also play an important role.

Finally, decisions 1 and 6 also share some similarities, as both invest in the same technologies and resources, with the main difference being that decision 1 does not invest in e-fuels. Despite this similarity, their performance differs significantly. Decision 1 invests very heavily in a few technologies, with almost full nuclear and wind deployment and high fossil fuels imports. Decision 6 also invests strongly in nuclear, wind, and fossil fuels, but at more moderate levels, and additionally relies on e-fuels imports. This more balanced approach leads to better performance under uncertainty.

Overall, the results show that a diverse energy mix and strong wind deployment are essential, but they are not enough by themselves. Robust early-stage decisions also require early investment in e-fuels and mobility changes toward BEVs.

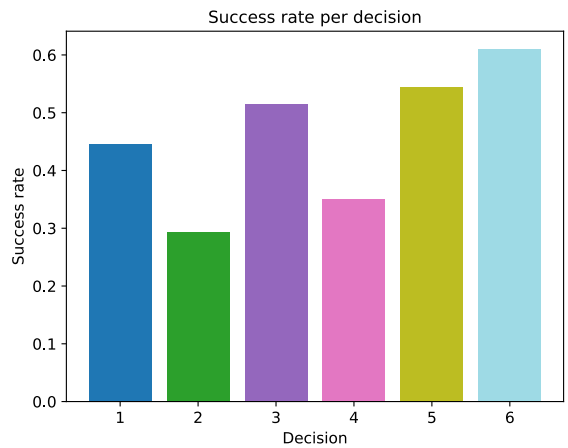

Figure 7. Success rate of each early-stage decision, defined as the ratio of successful trials meeting the 2050 CO_2 target to the total number of times each decision is selected. Decision 6 shows the highest success rate, followed by decisions 5, 3, 1, 4, and 2.

Digging deeper into the analysis of the results, the success rate is computed for each decision. The success rate is defined as the ratio of successful trials that meet the 2050 CO_2 target to the total number of times each decision is selected (Figure 7). In addition, the average total transition cost of the successful trials is calculated (Figure 8).

Consistent with previous results, decision 6 achieves the highest success rate at 61%, followed by decision 5 at 54%, decision 3 at 51%, decision 1 at 44%,

decision 4 at 35%, and decision 2 at 29%. When looking at the average total transition costs, decision 3 has the lowest cost at 1156 b€, followed by decision 6 at 1166 b€, decision 4 at 1172 b€, decision 1 at 1202 b€, decision 5 at 1255 b€, and decision 2 at 1372 b€. An important observation is that although decision 4 has one of the lowest average costs, it lacks robustness compared to decision 5, which has a higher average cost but performs better in terms of robustness. The same pattern appears when comparing decision 5 with decisions 1 and 3. Even though decisions 1 and 3 have lower average costs, they are less robust. This confirms that selecting the least costly option is not always the best strategy. Choosing alternatives that may be more expensive but ensure a robust transition leads to better outcomes.

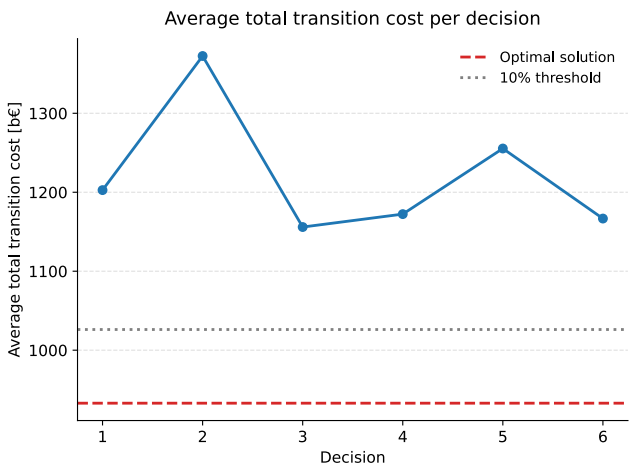

Figure 8. Average total transition cost of the successful trials for each early-stage decision. Although decision 3 exhibits the lowest average cost, it does not rank among the most robust options. Decision 6 combines relatively low costs with high robustness, while decision 5 shows higher average costs but stronger robustness compared to decisions 1, 3, and 4. This highlights that lower costs do not necessarily imply better robustness.

4 CONCLUSION

In this work, we explored the near-optimal space of early-stage decisions for the energy transition, focusing on the year 2035, and identified the most robust decisions under a broad range of unexpected events. To do so, we proposed a new decision-making framework that combines Modeling to Generate Alternatives (MGA) with a Multi-Armed Bandit (MAB) approach. This framework addresses the limitations of relying on a single cost-optimal pathway, which may perform poorly when unexpected events occur. In addition, the proposed approach reduces the computational effort required to identify robust early-stage transition strategies. Using the MGA approach on BEV adoption, nuclear, PV, and wind deployment, as well as e-fuels imports and fossil fuels imports, the near-optimal space reveals a wide range of solutions that differ structurally while remaining close in cost. For the period 2030–2035, almost any technology or resource among the MGA variables can be entirely omitted without exceeding 110% of the minimum cost for that phase. However, near-optimal solutions are not evenly distributed across the solution space: solar, wind, and nuclear deployment are skewed toward high levels, BEV adoption is centered around the optimal value, and e-fuels and fossil fuels imports are more uniformly distributed. After selecting five representative early-stage decisions from the large near-optimal space using k-means clustering, the MAB analysis shows that robust early-stage decisions are characterized by diverse energy generation portfolios, with a strong emphasis on wind deployment. In addition, early mobility changes toward BEVs and the use of e-fuels imports play an important role in improving robustness. These insights can support decision makers facing highly uncertain and unpredictable futures by highlighting actions that can be taken today to place the energy transition on a more robust path from the start.

Future work will define early-stage decisions as sets of alternatives rather than single solutions (entire clusters instead of centroids). This would better reflect real-world decision-making, where policies often guide the system toward general directions rather than fixed technology levels.

ACKNOWLEDGEMENTS

This work was supported by the Wallonie-Bruxelles International - WBI (M.K) and the Fonds de la Recherche Scientifique - FNRS [CR 40016260] (D.C).

REFERENCES

1. IPCC, 2023: Climate change 2023: Synthesis report. contribution of working groups i, II and III to the sixth assessment report of the intergovernmental panel on climate change [core writing team, h. lee and j. romero (eds.)]. IPCC, geneva, switzerland.
2. Stern N. Why are we waiting?. The MIT Press (2015). https://doi.org/10.7551/mitpress/10408.001.0001
3. Yue X, Pye S, DeCarolis J, Li FGN, Rogan F, Gallachóir BÓ. A review of approaches to uncertainty assessment in energy system optimization models. Energy Strategy Reviews 21:204-217 (2018). https://doi.org/10.1016/j.esr.2018.06.003
4. General Electric misread the energy transition: A cautionary tale. URL

https://ieefa.org/resources/general-electricmisread-energy-transition-cautionary-tale (accessed May 28, 2025).
5. Mavromatidis G, Orehounig K, Carmeliet J. A review of uncertainty characterisation approaches for the optimal design of distributed energy systems. Renewable and Sustainable Energy Reviews 88:258-277 (2018). https://doi.org/10.1016/j.rser.2018.02.021
6. Coppitters D, Wiest G, Göke L, et al. Identifying Dealbreakers and Robust Policies for the Energy Transition Amid Unexpected Events.
7. McCollum DL, Gambhir A, Rogelj J, Wilson C. Energy modellers should explore extremes more systematically in scenarios. Nat Energy 5:104-107 (2020). https://doi.org/10.1038/s41560-020-0555-3
8. Heinonen S, Karjalainen J, Ruotsalainen J, Steinmüller K. Surprise as the new normal – implications for energy security. Eur J Futures Res 5: (2017). https://doi.org/10.1007/s40309-017-0117-5
9. World Energy Council, Netherlands. About Black Swans and Green Futures (2021). https://publications.tno.nl/publication/34638137/0FQCYm/WEC-2021-about.pdf
10. Limpens G, Rixhon X, Contino F, Jeanmart H. Energyscope pathway: an open-source model to optimise the energy transition pathways of a regional whole-energy system. Applied Energy 358:122501 (2024). https://doi.org/10.1016/j.apenergy.2023.122501
11. Brill ED Jr. The use of optimization models in public-sector planning. Management Science 25:413-422 (1979). https://doi.org/10.1287/mnsc.25.5.413
12. Lai TL, Robbins H. Asymptotically efficient adaptive allocation rules. Advances in Applied Mathematics 6:4-22 (1985). https://doi.org/10.1016/0196-8858(85)90002-8
13. Thiran P, Jeanmart H, Contino F. Validation of a method to select a priori the number of typical days for energy system optimisation models. Energies 16:2772 (2023). https://doi.org/10.3390/en16062772
14. Victoria M, Zeyen E, Brown T. Speed of technological transformations required in europe to achieve different climate goals. Joule 6:1066-1086 (2022). https://doi.org/10.1016/j.joule.2022.04.016
15. Rixhon X. Exploration of Uncertainty-Aware Energy Transition Pathways, Ph.D. thesis, Université Catholique de Louvain (2024).
16. Pedersen TT, Victoria M, Rasmussen MG, Andresen GB. Modeling all alternative solutions for highly renewable energy systems. Energy 234:121294 (2021). https://doi.org/10.1016/j.energy.2021.121294
17. Ketchen DJ Jr, Shook CL. The application of cluster analysis in strategic management research: An analysis and critique. Strat Mgmt J 17:441-458 (1996). https://doi.org/10.1002/(sici)1097-0266(199606)17:6<441::aid-smj819>3.0.co;2-g

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimization-based Design, Simulation and Data-Driven Learning for Resilient Manufacturing Systems

Miriam Sarkis[a*] and Efstratios Pistikopoulos[ab]

[a] Texas A&M University, Energy Institute, College Station, 77843, Texas, United States
[b] Texas A&M University, Artie McFerrin Department of Chemical Engineering, College Station, 77843, Texas, United States
* Corresponding Author: msarkis@tamu.edu

ABSTRACT

Resilience is becoming a top priority across industrial sectors, with increasing pressures to assess it systematically. In this work, we present an optimization-based framework for proactive design and planning under uncertainty of multi-product manufacturing networks, and testing of the reactive strategies available to withstand unforeseen disruptions. Specifically, the design problem is formulated as a two-stage stochastic optimization, integrating multi-period planning and scheduling, aimed towards mitigation against uncertainty. Designs are then fixed and tested through simulated outcomes from *out-of-sample* uncertainty distributions, with feasibility of operation monitored through the time-to-recover post disruption. Infeasibility triggers a scenario-update procedure via K-means clustering, whereby critical uncertainty information based on simulated outcomes is integrated in the proactive planning step, including low-probability high-impact scenarios. Modular and non-modular designs are compared to quantify the value of recourse actions with respect to responsiveness to disruptions. Results highlight that adaptable modular designs enable shorter recovery times combined with lower cost commitments. Overall, the framework underpins the development of tools for the anticipation and mitigation of disruptive events and systematic re-design based on continuously updated uncertainty information.

Keywords: Multiscale Modelling, Planning & Scheduling, Design Under Uncertainty, Stochastic Optimization, Supply Chain, Resilience

INTRODUCTION

As disruptions become the new normal in manufacturing networks operation and management, there emerges a need to develop tools to test and quantify system resilience. In this context, process systems engineering opens opportunities to systematically design process systems under known uncertainty, operate them and test the capacity to withstand emerging unknown disruptive events. Following a disruptive event, data-driven methods offer the opportunity to learn and adapt process system design, integrating feedback from acquired knowledge regarding uncertainty[1].

Proactive strategies to improve resilience involve posing the design problem with respect to expected outcomes of uncertainty. For instance, a stochastic optimization problem is solved based on an approximation of fitted historical data of known demands which correspond to current available uncertainty information. This approach is well established across applications, including chemical supply chains, energy systems and pharmaceuticals [2, 3, 4].

During operation, design decisions are fixed and degrees of freedoms to withstand unforeseen scenarios and disruptive events depend on recourse actions available and the timescale of deployment of such measures. In this reactive planning step, available actions may range from adaptation of schedules to set up of additional capacity. To this end, processes with embedded modularity may allow flexibility to use repurpose and reconfigure equipment trains. Furthermore, modules can also correspond to parallel equipment lines which can be set up to rapidly scale-out/up capacity.

From a modelling perspective, simulation approaches such as rolling horizon [5] in this space are aimed at quantifying the reactive capabilities of a network/process design. In other words, reactive planning can be framed as a simulation of design performance as an *out-*

of-sample test of the first-stage decisions obtained based on the sampled uncertainty in the proactive planning step. Through simulation, significant deviations in target metrics can be monitored, such as service levels and time to recover. Acquired knowledge on system response can be then systematically incorporated in the overall design decisions through a learning step, iteratively re-solving the proactive problem and adjusting design-level decisions. To this end, there emerge opportunities to develop holistic methodological frameworks for the integration of feedback emerging from the reactive-planning part in the proactive layer to obtain better resilience-aware plans/design. This helps quantify resilience as the capacity to both *anticipate* and *respond* to uncertainty.

This study aims to provide a methodological iterative framework for doing so, illustrated on a pharmaceutical manufacturing network case, although transferable to any other batch manufacturing application, undergoing disruptions causing parameter fluctuations. In this work, modular versus non-modular manufacturing equipment cases are systematically compared to quantify the value of recourse capacity installation with respect to cost and adaptability to unforeseen uncertainty.

METHODS

Problem Statement

A capacity planning problem is considered, aimed at minimizing total network-level and scheduling costs. The objective is to evaluate candidate design alternatives based on the ability to anticipate uncertainty and respond under minimal cost and time delay. In this work, a pharmaceutical case study is considered, based on the multi-product manufacturing network shown in Figure 1a and manufacturing-level equipment network shown in Figure 1b. The problem considers:

- A set T of time periods t (daily);
- A subset TT of time periods ts (bi-monthly);
- A set S of states s,
- A subset S^p of product states s,
- A set I of tasks i,
- A set J of equipment j,
- A set A of scales a,
- A set U of parallel modules u,
- A set Ω of demand scenarios d.
- A set M of manufacturing facilities m
- A set K of storage nodes k
- A set L of demand zones l

In this work, multipurpose manufacturing is modelled. Additionally, parallel units can be installed for each equipment. Two alternatives are modelled, namely *modular* manufacturing, under the assumption that equipment can be installed with a time delay, and a *non-modular* option where equipment is installed once and cannot be scaled out in response to variability and disruptions. This effectively corresponds to comparing design with capacity-level recourse actions to fixed designs. The aim of the optimization is to determine the following decisions:

1. Investment in equipment in facilities (F_{jm}),
2. Scale and parallel units ($K_{jamu,ts,d}, Y_{jamu,ts,d}$)
3. Task allocation to equipment ($W_{ijamutd}$)
4. State inventories in manufacturing (S_{std})
5. Transport flows ($Q_{sjk,ts,d}, Q_{skl,ts,d}$)
6. Sales to meet demands ($Sales_{s,ts,d}$).

Decision 1 is scenario-independent and made at a *first stage*, prior to uncertainty realization. In the modular case, Decisions 2-6 are scenario-dependent and related to operations of the manufacturing network. Hence, these can be defined as *second-stage* decisions, adaptive to scenario realizations. In the non-modular case, Decision 2 is instead a first stage scenario-independent decision.

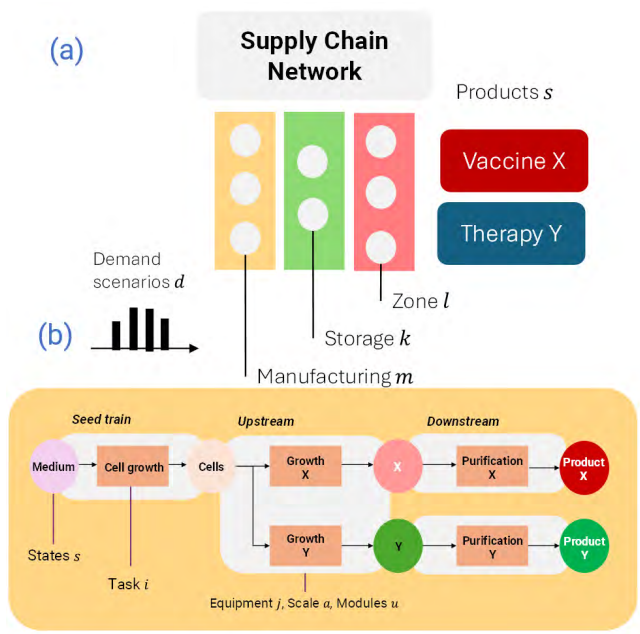

Figure 1. Optimization superstructure at (a) network level, (b) manufacturing process level, integrated supply chain network design and states-task network scheduling.

Algorithm 1: Framework Pseudo Code

Inputs:

- Initial scenario set Ω with probabilities ϕ
- Stress test sample size N_{stress}
- Time-to-recover threshold τ_{TTR}
- Initialize violation set $\Omega^{obs} \leftarrow \varnothing$

Initialize: $\epsilon \leftarrow 0$

while $\epsilon < 10$ do

Solve stochastic optimization problem $(\Omega^{\epsilon}, \phi^{\epsilon})$

Obtain first-stage variables α^{ϵ}, second-stage variables β^{ϵ}

Fix first-stage variables α^{ϵ}

Generate out-of-sample scenarios $\Omega^{stress}, |\Omega^{stress}| = N_{stress}$

for each sample $d = 1$ to N_{stress} **do**

Solve deterministic optimization problem

Evaluate time to recover TTR_{pd} for all products p

Compute $\widehat{TTR_d} = \max_p TTR_{pd}$

if $\widehat{TTR_d} \geq \tau_{TTR}$ **then**

Let $\Omega^{obs} \leftarrow \{D_1, \ldots, D_n\}$ for both products and $N = d$

Partition samples by regime into subsets $\Omega^{obs,r}$

Count regime occurrences $N_r = card(\Omega^{obs,r})$

Compute empirical regime probabilities $\pi_r = \frac{N_r}{N}$

Solve K-means $K_r := \min(K, N_r)$ for each regime r

Assign scenario realizations: $D_d^X := \mu_k^{X,r}, D_d^Y := \mu_k^{Y,r}$ for each (k, r)

Compute scenario probabilities $\phi_d := \pi_r p_k^r$

Update scenario tree $(\Omega^{\epsilon+1}, \phi^{\epsilon+1}) \leftarrow \{(D_d^X, D_d^Y, \phi_d)\}_{d=1}^{\Delta}$ with $\Delta = \sum_r K_r$

break

end if

end for

$\epsilon \leftarrow \epsilon + 1$

end while

Outputs: First-stage design vector α^*

Framework implementation

The workflow for the proactive, reactive planning and scenario updating problem is presented in Algorithm 1. First, a scenario-tree is postulated based on available uncertainty information and a stochastic optimization is solved (Step 3-6).

Iteratively, the obtained design is simulated for samples that are obtained from an unseen distribution, with the aim of testing design feasibility (Step 8-10). If $\widehat{TTR_d}$ is larger than a set threshold τ_{TTR}, the violation is deemed critical, hence the scenario tree is updated (Step 12-19), and the stochastic proactive problem resolved to obtain a new design. Ultimately, the iterative procedure terminates when a first-stage decision vector α^* is optimal with respect to the postulated scenario tree Ω^{ϵ} and feasible with respect to all disruptive scenarios Ω^{Stress}.

Process and demand data

Techno-economic models of manufacturing technologies of interest are developed using the flow sheeting software SuperPro Designer (Intelligen)[5]. Nominal capital, labor and variable cost and process tasks lengths (p_i) are presented in the SI document, for process volumetric scales $a = \{1000, 2000\}$. With respect to demand data, the optimization considers:

$$D_d^Y \sim \mathcal{U}(1000, 2000)\ \forall d = d1, \ldots, d5 \quad (1)$$

$$D_d^X \sim \mathcal{U}(10^6, 2 \times 10^6)\ \forall d = d1, \ldots, d5 \quad (2)$$

Whereby the demand samples are generated by sampling a uniform distribution which represents the knowledge based on historical data of demand patterns. Hence, D_d^X and D_d^Y correspond to samples for demand of product X and Y respectively. These correspond to the initialized scenario tree, with probabilities $\phi_d = \frac{1}{card(D)}$.

Manufacturing level formulation

The scheduling level is formulated as states-task network[6], modifying some constraints to account for scale-dependent (a) and parallel units (u).

Resource balance constraints model the relation between consumption and production ($B_{ijmautd}$) and storage (S_{smtd}) of each resource as follows, for every state s, time period t, facility m, equipment j, scale a and scenario d:

$$S_{smtd} = S_{s,t}^0 + S_{sm,t-1,d} + \sum_i \bar{\rho}_{is} \sum_{j,u,a} B_{ijmau,t-p_{is}d} - \sum_i \rho_{is} \sum_{j,u,a} B_{ijmautd}$$

$$\forall s, m, t, d \quad (3)$$

Here, it is worth noting that all $\bar{\rho}_{is}$ and ρ_{is} are equal to 1, due to each task modelled as a single state input/output entity (Figure 1b) (SI document).

The amount of material states processed in

equipment j, scale a and unit u through task i and time t in scenario d ($B_{ijmautd}$) is bounded by processing constraints, which also are tied to task allocation, via the binary of task-equipment allocation ($W_{ijmautd}$).

$$V_{ija}^{min} W_{ijmautd} \leq B_{ijmautd} \leq W_{ijmautd} V_{ija}^{max} \; \forall i, j, m, a, u, t, d \quad (2)$$

It is worth noting that each of these operational variables is scenario dependent d and made network-node dependent m, scale dependent a and unit u dependent through the respective indices.

Furthermore, task allocation is constrained by the following equation, which models the logic of equipment occupancy for equipment j at scale a and in unit u once a task is allocated, for the length of its duration p_i.

$$\sum_i \sum_{t'=t}^{t-p_i+1} W_{ijmautd} \leq 1 \; \forall i, j, m, a, u, t, d \quad (4)$$

Resource storage constraints model the bound on inventory for each state, with a capacity limit (C_s):

$$0 \leq S_{smtd} \leq C_s \; \forall s, m, t, d \quad (5)$$

Task allocation occurs only within installed equipment, which is modelled through the following constraint:

$$W_{ijmautd} \leq Y_{jmau,t_s(t),d} \; \forall i, j, m, a, u, t, d \quad (6)$$

Where $t_s(t) = \max\{t' \in T^m : t' \leq t\}$, hence corresponds to the month period ts that encompasses day t.

Equipment installation occurs only within installed facilities suited to produce equipment j at scale a (F_{jma}):

$$Y_{jmau,ts,d} \leq F_{jm} \; \forall i, j, m, a, u, ts \in TT, d \quad (7)$$

In the modular case modules can be additionally installed through investment decisions and available after a lead time of 60 days :

$$Y_{jmau,ts,d} = Y_{jmau,ts,d} + K_{jmau,ts-60,sd} \; \forall i, j, m, a, u, ts \in TT, d \quad (8)$$

In the non-modular case, no new investment is possible and availability is carried forward:

$$Y_{jmau,ts,d} = Y_{jmau,ts,d} \; \forall i, j, m, a, u, ts \in TT, d \quad (9)$$

Additionally, line installation is not scenario dependent, but rather an investment made as a first-stage decision, hence fixed across scenarios via the following set of additional constraints:

$$Y_{jmau,ts,d1} = Y_{jmau,ts,d'|d' \neq d1} \; \forall i, j, m, a, u, ts, d' \quad (10)$$

Network level formulation

At supply chain level, inventory balance constraints model the shipment from manufacturing facilities to storage facilities. For the states that are only produced (s^P):

$$S_{smtd} = S_{s,t}^0 + S_{sm,t-1,d} + \sum_i \bar{\rho}_{is} \sum_{j,u,a} B_{ijmau,t-p_{is}d} - \sum_k Q_{smk,ts,d} \; \forall s^P, m, t, d \quad (11)$$

Where, the outbound flow of manufacturing is considered as a consumption term ($Q_{smk,ts,d}$). Shipments to storage facilities k feed into the following equation which models the storage at each facility:

$$\sum_m Q_{smk,ts,d} = \sum_l Q_{skl,ts,d} \; \forall s \in S^P, k, ts \in TT, d \quad (12)$$

In this work, product inventory is not considered and these flow variables are matched only at time periods within the set TT, thus forcing shipments to only occur at time periods ts.

Transport flows from each storage facility feed into demand zones, where demand is fulfilled. Hence:

$$\sum_k Q_{skl,ts,d} = S_{sl,ts,d} \; \forall s \in S^P, l, ts \in TT, d \quad (13)$$

The following sales-demand constraint ensures the amount of doses sold for both product states, at each demand zone l and time period t within the mid term horizon is greater than the demand in each scenario:

$$Sales_{sl,ts,d} \delta_s \geq D_{sl,ts,td} \; \forall s \in S^p, l, ts \in TT, d \quad (14)$$

The demand in each demand zone is calculated based on the total demand sampled, previously computed:

$$D_{sl,ts,d} = \frac{D_d^s}{card(l)} \forall s \in S^p, l, ts \in TT, d \quad (15)$$

Objective function

The total cost of the supply chain is calculated as a function of the capital, labor, variable manufacturing costs and logistics costs:

$$\min Expected\ COST = \sum_{jm} c_j^{FAC} F_{jm} + [\sum_{ijtaumd} \phi_d c_{jai}^{MOD}(Y_{jmautd}) + \sum_{ijtaumd} \phi_d c_{jai}^{LAB}(W_{ijtaumd}) \sum_{ijtaumd} \phi_d c_{jai}^{VAR}(B_{ijtaumd}) + \sum_{std} \phi_d c_s^{STOR}(S_{std})] + \sum_{sltd} \phi_d c^P(\eta_{sltd}) + \sum_{smktd} c^T \lambda_{mk} Q_{smktd} + \sum_{skltd} c^T \lambda_{kl} Q_{skltd} \quad (16)$$

Where the expected cost is sum of scenario-independent investment and scenario-dependent weighted average sum over the operational, transport, storage decisions and, in the modular case, capacity expansion. Here, $c_{ja}^{FAC}, c_{jai}^{MOD}, c_{jai}^{LAB}, c_{jai}^{VAR}, c_s^{STOR}, c^P, c^T$ correspond to the capital, facilities dependent, labor, variable, storage and penalty cost coefficients respectively.

Optimization-based Simulation

The first-stage decisions are fixed and the violation of the sales demand constraint is monitored and penalized. Hence Equation 14 is modified as follows:

$$Sales_{sl,ts,d} \delta_s + \eta_{sl,ts,d} \geq D_{sl,ts,d} \; \forall s, l, ts \in TT, d = d1 \in \Omega^{Stress} \quad (17)$$

The time to recover TTR is the minimum time required to restore service after a disruption, corresponding to the

last time period and the first time period undergoing disruption:

$$TTR_{sd} = \max\{ts \in TT, |\sum_l \eta_{sl,ts,d} > 0\} - \min\{ts \in TT, |\sum_l \eta_{sl,ts,d} > 0\} \ \forall s, d = d1 \in \Omega^{Stress} \quad (18)$$

In the simulation step, the objective function (Equation 16) is modified with the addition of a penalty term calculated as function of the shortage of therapies.

$$\min Expected\ COST = \sum_{jm} c_j^{FAC} F_{jm} + [\sum_{ij,ts,aumd} \phi_d\, c_{jai}^{MOD}(Y_{jmau,ts,d}) + \sum_{ijtaumd} \phi_d\, c_{jai}^{LAB}(W_{ijtaumd}) + \sum_{ijtaumd} \phi_d\, c_{jai}^{VAR}(B_{ijtaumd}) + \sum_{std} \phi_d\, c_s^{STOR}(S_{std})] + \sum_{sl,ts,d} \phi_d\, c^P(\eta_{sl,ts,d}) + \sum_{smk,ts,d} c^T \lambda_{mk}(Q_{smk,ts,d}) + \sum_{skl,ts,d} c^T \lambda_{kl}(Q_{skl,ts,d}) \quad (19)$$

Upon simulation of reactive planning, the scenario set is reduced to $\{d1\}$, with demand levels assigned iteratively to out-of-samples generated in Ω^{stress}. Hence, demands are generated from out-of-sample distributions:

$$D_d^Y \sim \mathcal{U}(1000, 3000)\ \forall d = d1 \quad (20)$$

Whereby the bounds of the monthly demands for Product Y are modified (Figure 2a).

Additionally, the disruptive regime for Product X is generated based on conditional probability theory, considering a bimodal distribution of demands, where the modes are defined by the sample regime R_d (Figure 2b):

$$R_d \sim Bernoulli(p_{Disruption}), p_{Disruption} = 0.04\ \forall d = d1 \quad (21)$$

$$D_d^X = \begin{cases} Z_d^{disruption}, if\ R_d = 1 \\ Z_d^{Normal}, if\ R_d = 0 \end{cases} \quad (22)$$

Where:

$$Z_d^{disruption} \sim \mathcal{N}(\mu_{disruption}, \sigma^2_{disruption}),$$

$$\mu_{disruption} = 8 \times 10^7, \sigma^2_{disruption} = 2 \times 10^7 \quad (23)$$

$$Z_d^{Normal} \sim \mathcal{N}(\mu_{Normal}, \sigma^2_{Normal})$$

$$\mu_{Normal} = 7 \times 10^6, \sigma^2_{disruption} = 1 \times 10^6 \quad (24)$$

Scenario tree update

The scenario tree update mimics the concept of learning from the demand observed throughout the time horizon. K-means clustering allows systematic grouping collected data in set of K clusters, selected equal to the number of scenarios to build the tree for stochastic optimization and cluster heights used to scale scenario probabilities, whilst also accounting for regime information that enable partition of the samples and construction of regime-specific clusters. Here, $\Omega^{Obs} = \{(D_d^X, D_d^Y, r_d)\}_{d=1}^N$ is the set of demand samples d observed up to iteration ϵ with a recorded corresponding regime. The samples belonging to each regime can be defined as the subsets:

$$\Omega^{Obs,r} = \{(D_d^{X,r}, D_d^{Y,r}) : r_d = r\} \quad (25)$$

The number of occurrences of regime r is defined as:

$$N_r = \mathrm{card}(\Omega^{Obs,r}) \quad (26)$$

Hence, the empirical probability of regime r is the ratio of samples when the regime N_r realizes and the total N:

$$\pi_r = \frac{N_r}{N} \quad (23)$$

For each regime, K-means clustering is performed by solving:

$$\min_{\mu_\kappa^r} \sum_{\kappa=1}^{K_r} \sum_{d \in C_\kappa^r} \left\| \begin{pmatrix} D_d^X \\ D_d^Y \end{pmatrix} - \mu_\kappa^r \right\|^2 \quad (27)$$

Where the C_κ^r is the set of samples d assigned to cluster κ under regime r and μ_κ^r is the centroid of cluster κ under regime r, corresponding to the vector $\begin{pmatrix} \mu_\kappa^{X,r} \\ \mu_\kappa^{Y,r} \end{pmatrix}$ with respect to each product X and Y. The number of clusters per regime is therefore defined as:

$$K_r := \min(K, N_r) \quad (28)$$

Through Equation 28, if a new regime is discovered, N_r new leaves of the scenario tree are added if N_r is below that of the target number K of clusters, which in this work equals 3. Otherwise, K-means yields K scenarios for the each detected regime. The resulting scenario tree probabilities would be adjusted by regime-specific probabilities which are also computed based on collected data during the reactive stage. Therefore, once solved, K-means yields the new scenario realizations, which are reassigned as equal to the identified cluster centroids μ_r^k. Each cluster (k, r) defines a scenario d.

$$D_d^X := \mu_\kappa^{X,r} = \frac{1}{|C_\kappa^r|} \sum_{d \in C_\kappa^r} D_d^{X,r}$$

$$D_d^Y := \mu_\kappa^{Y,r} = \frac{1}{|C_\kappa^r|} \sum_{d \in C_\kappa^r} D_d^{Y,r} \quad (29)$$

New scenario probabilities can be derived from the empirical regime probabilities and the probabilities p_k^r of cluster k for each regime r are calculated as the number of samples that fall under cluster C_k^r and the occurrence of regime r:

$$p_\kappa^r = \frac{card(C_\kappa^r)}{N_r} \quad (30)$$

Hence, overall scenario probabilities are obtained as:

$$\phi_d := \phi_\kappa^r = \pi_r p_\kappa^r \quad (31)$$

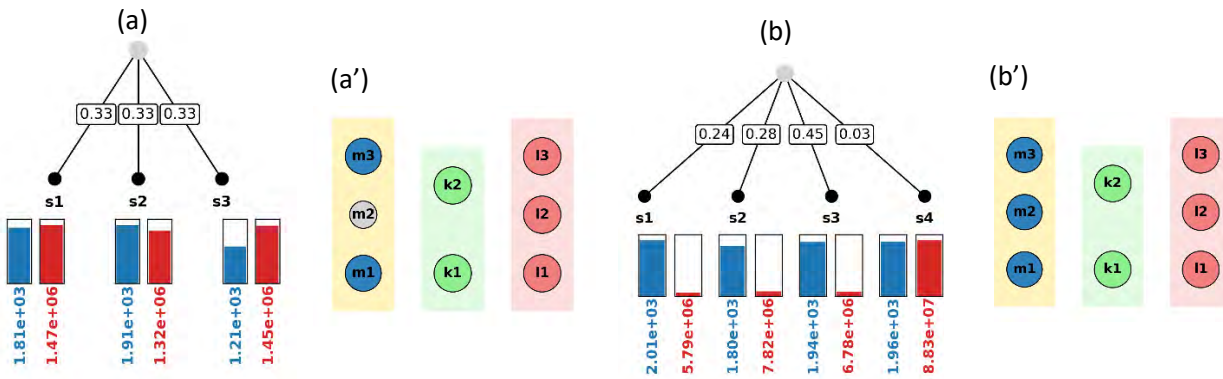

Figure 2. Inputs and first-stage outputs of iterative workflow (a) scenario tree for Iteration 1 (initialization, $\Omega = \Omega^1, \phi = \phi^1$), (a') network-level optimized design for iteration 1, (b) updated scenario tree for iteration 2 (K-means generated, $\Omega = \Omega^2, \phi = \phi^2$), (b') network-level optimized design for iteration 2.

COMPUTATIONAL STATISTICS

The framework is built in Python 3.11 implementing the PYOMO package. The MILP is solved with GUROBI v11 on a Intel(R) Core(TM) i7-10700 CPU @2.90 GHz machine with 64.0 GB RAM. Table 1 presents the computational time for the modular and the non-modular case. It is worth noting that both result in the same number of variables and different computational times across iterations, as the non-modular case is solely obtained by fixing binaries of equipment activation across scenarios (Equation 10), rather than reformulating the problem.

Table 1: Computational Statistics

Iteration	Metric	Modular	Non-modular
1	*CPUs*	236.1	170.11
	Continuous	522	522
	Integer	62,061	62,061
	Constraints	62,595	62,595
2	*CPUs*	690.78	443.14
	Continuous	696	696
	Integer	82,701	82,701
	Constraints	83,409	83,409

Statistics are reported for the proactive solution stage, The reactive stage was found to lead to computational times of $< 1s$, with CPUs reduced due to the fixing of the design degrees of freedom, speeding up convergence.

RESULTS & DISCUSSION

The proposed framework provides a platform to test design feasibility and adaptation capacity to unforeseen uncertainty. Figure 2a presents the initializing scenario tree, which corresponds to equiprobable scenarios sampled from the distributions presented in Equation 1-2. This results in first-stage network-level decisions which are equivalent in the modular and non-modular case (Figure 2a'), corresponding to two facilities installed, namely $m1$ and $m3$, minimizing transport cost and distance travelled, and fulfilling the expected demands. It is worth noting that first-stage equipment-level decisions for the modular and non-modular decisions differ and this is displayed in Figure 3a and Figure 3b, at time period $t = 0$, with different equipment lines installed in terms of scale and parallel units. This is due to the fact the modular equipment allows for ad-hoc installation of capacity as second-stage decisions, whilst the non-modular case considers that the installed equipment is fixed and should fulfill the postulated demand scenarios.

Fixing these case-specific first-stage decisions, stress testing allows the quantification of the adaptive capability of each alternative. In Figure 4a-a' the outcomes of stress testing are highlighted with respect to each product. For each demand realization highlighted in black, modular and non-modular options yield a feasible solution, whereby feasible corresponds to a network that ensures supply given the installed equipment capacity (non-modular case) or through installation of additional units. This is done minimizing cost and time to recover, and, if feasible, the time to recover is below the threshold of $\tau_{TTR} = 120\ days$. Instead, samples highlighted in red are those that trigger this infeasibility, corresponding to a 1900 doses of Y and critically, a 10-fold increase in demand for X, which in this case study corresponds to low-probability, high impact occurrence of a pandemic.

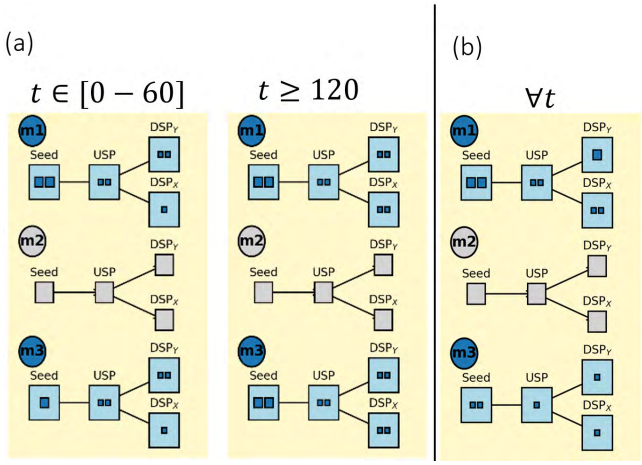

Figure 3. Equipment-level decisions: (a) modular case, (b) nonmodular case fixed equipment structures. Line activated (■), facility inactive (■). Number of lines u and scale a is illustrated for each process equipment type j.

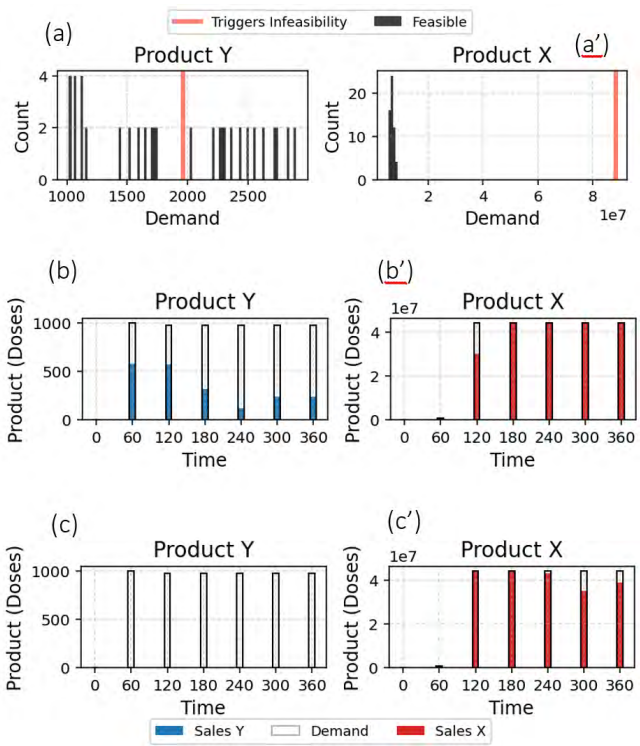

Figure 4. Infeasibility detection: (a) demand for Y, (a') demand for X, (b-b') time to recover in modular case, (c-c') time to recover in non-modular case.

As can be seen in Figure 4b-b', in the modular case, the disruption occurrence causes a $\widehat{TTR_d}$ with respect to product X of approximately $300\,days$ and a minimized $\widehat{TTR_d}$ of $60\,days$ with respect to X. Shortages are even more significant in the non-modular case, with delays of shortages of 120 days on the X side (Figure 4c') and unfulfilled demands on the Y side (Figure 4c).

The response capabilities of the network are seen in Figure 3a and b, where the adaptation of the second-stage decisions is displayed for $t > 60\,days$. In the modular case (Figure 3a), a scale-out of seed capacity and additional DSPx lines to cater for the increased X demand. This leads to the service levels seen in Figure 3b' for X, which are attainable without compromising on capacity for the manufacturing of Y. This then leads to a $\widehat{TTR_d}$ which still does not satisfy the feasibility check; however, it is less significant than in the fixed non-modular case. Figure 3b illustrates instead a fixed structure that is not scenario dependent on the entirety of the time horizon, hence the only degrees of freedoms to respond to disruptions become the adjustments of the operational planning decisions. This leads to more significant shortages and longer $\widehat{TTR_d}$ times in the non-modular case.

Additionally, Figure 2b highlights how post-infeasibility, K-means clustering generates a new scenario tree, systematically incorporating 1 new leaf, which corresponds to the detected 1 low-probability (0.03) high impact $(8.37e8\ doses)$ scenario representing the *disruption* regime, and adjusted probabilities based on the realized demands for the *normal condition* regime. Solving the stochastic optimization again leads to adjusted first-stage decisions, which also embed the activation of facility $m2$ (Figure 2b')

Based on the updated design, the simulated stress test is repeated, resulting in feasibility for all samples within Ω^{Stress}, for both the modular and non-modular case. Then, it becomes worth comparing the costs of optimality and feasibility. In Figure 5a, the costs for the modular case is displayed, with iteration 1 resulting in costs of $8.3M\ USD/y$ for facility capital-investment related decisions. The CapEx of the module investment $(90M\ USD/y)$, the variable $(11M\ USD/y)$ and labor costs $(6M\ USD/y)$ are second-stage because dependent on operation and scenario realization and therefore not-committed a priori. Instead, in the non-modular case, the equipment is installed and fixed, hence the total displayed CapEx of $12.8M\ USD/y + 131.9M\ USD/y$ is a cost-commitment taken prior to scenario realization and larger, as expected to cater for all demand occurrences, including increased demands with additional capacity.

Stress testing in Iteration 1, highlights that in both the modular and non-modular case the expected cost can increase, due to realization of larger demands not captured in the scenario tree. Specifically, the outlier on the modular case corresponds to the sample displayed in Figure 4b-b', where increased production levels to minimize shortages lead to higher costs. At Iteration 2 the expected optimal costs in both cases result higher based on an updated scenarios that capture the observed samples in the previous iteration, until infeasibility was detected. Added cost results from idle equipment in $m2$, catering for disruptions and increased demand resulting from it; expected modular costs remain lower.

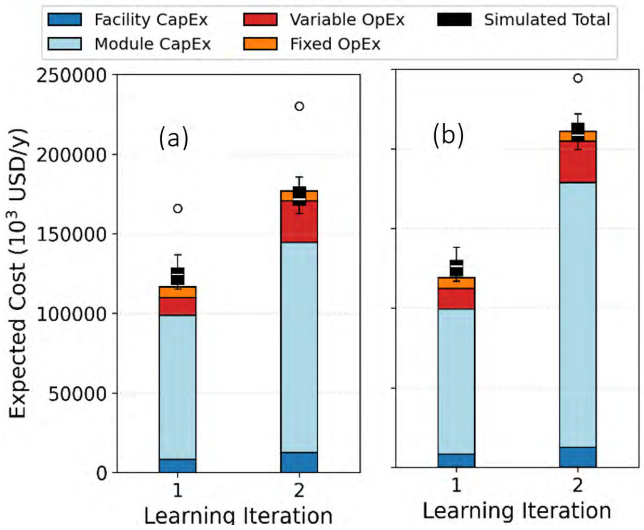

Figure 5. Expected and simulated costs for (a) modular, (b) non modular case over iterations

It emerges that flexibility of rapid scale-out allows minimization of cost commitments and better adaptability to disruption scenarios. Hence, the presented comparison between modular and non-modular equipment is fundamentally a theoretical comparison on benefit of recourse actions in ensuring cost-effective anticipation and response to disruptive events that impact the system.

Based on the results presented, there emerge avenues of development of the optimization-based assessment platform. First, the reactive planning part solves a deterministic optimization problem which assumes perfect foresight of the disruption realizing at time $t = 60$. Hence the TTR is computed optimistically, based on deterministic disruption anticipation and mitigation. Nevertheless, solving this part through a rolling horizon optimization would correspond to a more-realistic response characterization as part of the feasibility test. Second, expected cost was primarily used as performance indicator when comparing the two manufacturing alternatives. Sustainability metrics for capacity establishment and operation could be integrated in the optimization constraints: the former would account for footprint of facility establishment, multiplied by the selection of binary installment variables (e.g. $Y_{jmau,ts,d}$) and the latter accounting for footprint scaling production levels (e.g. $B_{ijtaumd}$) and resource use (e.g. heat and electricity).

Overall, the proposed tool enables process design resilience assessment prior to deployment at scale and during operation. On the one hand, the modelling platform could support comparative assessments, as shown in this study, to identify good design candidates that perform cost-effectively in response disruptions. On the other hand, the tool could be deployed as a back-end simulator, with testing for a continuously updated set of disruptive events and anticipation of re-design priorities prior to disruptions striking in real life scenarios.

CONCLUSIONS

Resilience is becoming a growing priority across industrial sectors, with pressures to design systems that are adaptive to evolving uncertainty from diverse sources. In this work, we develop a framework for proactive design under a priori postulated uncertainty and testing for adaptation to unforeseen uncertainty through simulation. This is achieved through an *out-of-sample* optimization-based simulation from distributions that describe occurrences of disruptive events impacting demand scale for the products. To update uncertainty information when significant service loss due to disruptions emerge, scenario clustering techniques are leveraged to allow updating of the scenario tree, updating uncertainty information worth integrating in the design steps.

The capabilities of the framework are shown comparing examples of modular and non-modular manufacturing equipment designs, highlighting the value of recourse capacity expansion which the former option offers, alongside to initial lower cost commitments. Overall, the tool enables comparative analyses towards the establishment of more resilient designs centered around continuous improvement via systematic uncertainty-related data, making it a valuable tool in real-time simulation of future disruptions through anticipation, mitigation and adaptation of process/network design.

DIGITAL SUPPLEMENTARY MATERIAL

Supplementary information including nomenclature and input data used is available at **LAPSE:2026.0044**.

ACKNOWLEDGEMENTS

Financial support from the Texas A&M Energy Institute is gratefully acknowledged.

REFERENCES

1. Chrisandina NJ, Vedant S, Iakovou E, Pistikopoulos EN, El-Halwagi MM. Metrics and methods for resilience-aware design of process systems: advances and challenges. Current Opinion in Chemical Engineering 43:100984 (2024). https://doi.org/10.1016/j.coche.2023.100984
2. Marques CM, Moniz S, de Sousa JP, Barbosa-Póvoa AP. A simula-tion-optimization approach to integrate process design and planning decisions under technical and market uncertainties. Comput Chem Eng. 106:796-813.(2017) https://doi:10.1016/j.compchemeng.2017.04.008
3. Efthymiadou ME, Charitopoulos VM, Papageorgiou

LG. Optimization approach for hydrogen infrastructure planning under uncertainty. Ind. Eng. Chem. Res. 64:7431-7451 (2025). https://doi.org/10.1021/acs.iecr.4c04211

4. Guillén?Gosálbez G, Grossmann IE. Optimal design and planning of sustainable chemical supply chains under uncertainty. AIChE Journal 55:99-121 (2008). https://doi.org/10.1002/aic.11662
5. Sarkis M, Shah N, Papathanasiou MM. Characterization of key manufacturing uncertainties in next generation therapeutics and vaccines across scales. J Adv Manuf Process. 5: e10158 (2023) https://doi:10.1002/amp2.10158
6. Kondili E, Pantelides CC, Sargent RWH. A general algorithm for short-term scheduling of batch operations—i. MILP formulation. Computers & Chemical Engineering 17:211-227 (1993). https://doi.org/10.1016/0098-1354(93)80015-f

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Machine Learning Implementation for Fermentation Quality Prediction in Wine Manufacturing

Matthew A.J. Hill and Dimitrios I. Gerogiorgis*

Institute for Materials and Processes (IMP), School of Engineering, University of Edinburgh, Edinburgh
* Corresponding Author: D.Gerogiorgis@ed.ac.uk.

ABSTRACT

Wine consumers are increasingly health- and environmentally conscious. At the same time, white wine and rosé drinkers favour freshness and varietal aromas, which requires low-temperature regimes that extend fermentation time and increase energy demand. Additionally, global warming accelerates grape ripening which increases alcohol level in wine. To reduce cost and alcohol levels while maintaining quality, predictive tools that forecast how fermentation conditions impact fermentation time, and primary and secondary metabolite concentrations, can provide practical benefits to wineries by expediting oenological decisions-making and in turn reducing energy demand. Additionally, literature highlights static models in smart manufacturing suffer from performance degradation with data drift. In light of this, we successfully developed and evaluated pipelines for the automated design and training of three ML methods - support vector regression, random forest and artificial neural networks – to predict fermentation outcomes from ten initial features. The dataset employed captures the fermentation of nine commercially available *Saccharomyces cerevisiae* strains, each performed in triplicate, across two synthetic media and four temperature regimes, sampled at nine time points. The targets we focused on included time to complete fermentation, Ethanol, Acetate, given its undesired spoilage effects and ten secondary metabolite concentrations, given their synergistic impact on aroma. We were able to recommend effective pipelines for each task and identified our feature set contained adequate predictive signal for ethanol, acetate and fermentation duration predictions, but inadequate for the multi-output regression of secondary metabolite concentrations, which exhibited target heterogeneity. Ultimately, embedding these pipelines into optimisation frameworks, offers an actionable route to tailoring wine characteristics to evolving consumer preference, while reducing energy demand.

Keywords: alcoholic fermentation, machine learning, artificial neural network, support vector regression, random forest regression, fermentation time, efficacy, secondary metabolite concentration

INTRODUCTION

In 2024 the International Organization of Vine and Wine attributed the six-decade low wine consumption, 214.2 million hectolitres, to high average prices and emerging consumer trends [1]. These trends include prioritizing health, sustainability, and preferring low-alcohol beverages. The wine industry therefore requires solutions to produce environmentally friendly, low-alcoholic and high-quality wines, whilst reducing cost.

This is a particularly significant challenge in white-wine and rosé production as consumers prefer freshness and varietal aromas [2]. This requires low-temperature fermentation as it fosters the development and preservation of volatile compounds [3]. However, low-temperature isothermal programs (12-19 °C) can extend the duration of fermentation and increase the risk of sluggish and stuck fermentations. Given that alcoholic fermentation is exothermic, the process requires wineries to actively cool fermentation vats. This cooling accounts for between 45% and 90% of the total energy consumption in manufacturing. In the case of white wine and rosé, where fermentation is prolonged, the process becomes environmentally unfriendly and costly [4]. To address this challenge Minebois et al. investigated a non-isothermal temperature regime on 9 commercially available *Saccharomyces cerevisiae* strains across two synthetic grape must (SGM) compositions. The non-isothermal

temperature regime, henceforth referred to as the time-varying temperature regime (TVAR), began fermentation at 23°C, to leverage *S. cerevisiae's* higher optimal growth temperature, and upon entry into the stationary phase of fermentation switched to 15°C [2]. The two SGM compositions contained 210 g L^{-1} and 250 g L^{-1} total sugars (glucose and fructose in racemic quantities) and hence-forth will be referred to as MS210 and MS250, respectively. Additionally, the TVAR was compared to three isothermal regimes (12°C, 18°C, 25°C). The published dataset captures the complex interaction between temperature, strain, and fermentation media.

Initial fermentation conditions strongly influence the final composition and therefore quality of wine. These conditions affect the concentrations of primary metabolites such as ethanol, acetate, and glucose, as well as the concentrations of secondary metabolites that define the wines aroma profile. Glucose contributes to sweetness and elevated acetate levels can produce vinegary notes [2]. Therefore, predictive tools that forecast the impact of fermentation conditions on fermentation time, and, primary and secondary metabolite concentrations can accelerate oenological decision-making, whilst ensuring quality and reducing both costs and energy demand. Considering this, we looked to employ three machine learning (ML) methods – support vector regression (SVR), random forest (RF), artificial neural networks (ANN) – to forecast these outcomes, given their proven ability to map complex interactions [5]. However, in recognition that literature adopts static ML models, often with arbitrarily selected architectures and hyperparameters, where performance degrades with data shifts [6]. We therefore developed and evaluated pipelines an actionable workflow for the automated design and training of these methods.

Consumers have become increasingly conscious of adverse health effects and roadside safety, and, as climate change increases alcohol level in wines by accelerating grape ripening, research on reducing alcohol content in wine has intensified [7]. Subsequently, the primary metabolites we chose to predict were ethanol and acetate, considering that bacteria from grapes have been found to oxidize ethanol to acetic acid, a volatile acid that causes wine spoilage [8]. We chose to predict secondary metabolite concentrations for their synergistic effects on wine's aroma profile, a key component of wine quality.

Machine Learning Methods

In our investigation we explored supervised learning methods which aim to find parameters, θ, that minimize the average loss over a finite training set of input-output pairs $\{(x_i, \mathrm{y_i})\}_{\mathrm{i=1}}^{\mathrm{N}}$. Where, $\mathrm{x_i}$ encodes strain identity, synthetic must strength (MS210 or MS20), temperature regime and initial metabolites, and $\mathrm{y_i}$ denotes our predicted targets.

Support Vector Machines (SVM)

Support Vector Machines (SVM) are large-margin classifiers that separate labeled data by maximizing the margin between classes [9]. For linear SVM, the decision function is $f(x) = \omega^{\mathrm{T}} x + b$, where ω is the weight vector and b is the bias. The parameters (ω, b), are obtained by solving the soft-margin optimization problem. SVR extends this principle to continuous targets by introducing Vapnik's ε-insensitive loss, where an ε-tube is fitted around the regression function [9]. Deviations larger than ε incur a penalty and the optimization problem becomes,

$$min_{\omega,\xi,\xi^*} \frac{1}{2}\|\omega\|^2 + C\sum_{i=1}^{N}(\xi_i + \xi_i^*) \quad (1)$$

s.t. the constraints (Eq. 2, 3, 4),

$$y_i - f(x_i) \le \varepsilon + \xi_i^* \quad (2)$$

$$f(x_i) - y_i \le \varepsilon + \xi_i^* \quad (3)$$

$$\xi_i, \xi_i^* \ge 0 \quad (4)$$

where the slack variables, ξ_i and ξ_i^*, measure how far sample i lies outside the ε-tube, C adds a linear penalty to those violations and y_i is the true target value. To capture nonlinearity, SVMs employ the kernel trick, replacing the inner product $\omega^{\mathrm{T}} x$ with a kernel function $K(x, x')$ that behaves like an inner product in some (potentially higher) feature space. A common choice is the radial basis function (RBF) kernel, $K(x, x') = e^{\left(-\gamma\|x - x'\|^2\right)}$, where γ is a scaling parameter that controls the width of the Gaussian kernel function.

Random Forest (RF)

Random Forest (RF) is an ensemble learning method that combines many decision trees using bootstrap aggregation (bagging). Each tree is fit to a bootstrap sample of the training data, and at every split the algorithm considers only a random subset of the input features. This randomization reduces correlation and tends to improve generalization [10]. For regression, the forest prediction for an input vector x is the average of individual tree predictions,

$$\hat{y}(x) = \frac{1}{T}\sum_{t=1}^{T} f_t(x) \quad (5)$$

where T is the number of trees in the forest and $f_t(x)$ is the prediction of t-th tree trained on its own bootstrap sample. Within each tree, splits are chosen to minimize the mean squared error (MSE) of the target values in the child nodes and a cost complexity pruning parameter, α, can be added to penalize larger trees.

Artificial and Recurrent Neural Networks (ANN)

An Artificial Neural Network (ANN) is a parametric

function approximator that learns the mapping $f_\theta: \mathbb{R}^d \rightarrow \mathbb{R}^m$, where d and m are the number of input and output dimensions, respectively. Given an input vector $x \in \mathbb{R}^d$ and parameters $\theta = \{\omega, b\}$, each neuron computes a weighted sum $f(x) = \omega^\top x + b$, to which a non-linear activation function is applied, this mathematically replicates a biological neuron [10]. Each neuron then passes the processed signal forward as its output and multiple neurons constructs a layer. If we denote k as a discrete training step, the output of a single neuron can be written as,

$$y(k) = F\left(\sum_{i=0}^{m} w_i\,(k) \bullet x_i\,(k) + b\right) \tag{6}$$

where $F\,(\bullet)$ is the activation function, $x_i\,(k)$ is the i-th input component and $y(k)$ is the output [11]. A common activation function that applies a nonlinear transform is the rectified linear unit function (ReLU), which is applied component-wise,

$$ReLU(z) = \begin{cases} z, z > 0 \\ 0, z \leq 0 \end{cases} \tag{7}$$

where z is the pre-activation function within Eq.6 [8]. A Feed-forward ANN (FFANN) is one in which information flows strictly in one direction from the input layer to the output layer. A common FFANN is a multilayer perceptron (MLP) where all layers are fully connected (dense) [11]. An MLP with two or more hidden layers is referred to as a deep neural network and can learn hierarchical feature representations.

When training, a backpropagation algorithm was employed to minimize the mean squared error between predicted and observed targets, Eq. 8 [12]. Backpropagation applies the chain rule to efficiently compute the gradient of the loss with respect to each weight and bias, Eq. 9. Weights are then updated with the Adam optimizer, an adaptive variant of stochastic gradient descent that uses exponentially decaying, bias-corrected first and second order moment loss gradient estimates, eq. 11.

$$\mathcal{L}(\theta) = \frac{1}{m} \sum_{j=1}^{m} \left(y_j - \hat{y}_j\right)^2 \tag{8}$$

$$\frac{\partial \mathcal{L}}{\partial \omega_{kj}} = \frac{\partial \mathcal{L}}{\partial \sum_i \omega_{kj} a_j + b_k} \bullet \frac{\partial \sum_i \omega_{kj} a_j + b_k}{\partial \omega_{kj}} \tag{9}$$

$$\omega_{kj}{}^{(t+1)} = \omega_{kj}{}^{(t)} - \eta \frac{\hat{m}_{t,kj}}{\sqrt{\hat{v}_{t,kj} + \epsilon}} \tag{10}$$

where $\mathcal{L}(\theta)$ is the MSE loss-function, y_j is the true target value, $\hat{y}_j$ is the models prediction, a_j is the activation output of neuron j, $\frac{\partial \mathcal{L}}{\partial \omega_{kj}}$ is the loss gradient with respect to weight $\omega_{kj}(k)$, t is the current training step, $\hat{m}_{t,kj}$ and $\hat{v}_{t,kj}$ are the bias corrected first and second order moments and η is the learning rate.

METHODOLOGY

Minebois et al.'s dataset contains samples of biomass, primary metabolites, amino-nitrogen species, and volatile compounds at nine time points (T_0- T_8), where T_8 marks the completion of fermentation. Data was collected for nine commercially available *S. cerevisiae* strains across two synthetic must strengths (MS210, MS250), and four temperature regimes. As mentioned above, this included a time-varying temperature regime (TVAR) that switches from 23°C to 15°C at one-third completion of fermentation and 3 isothermal regimes ($12°C, 18°C, 23°C$). We developed supervised ML pipelines to predict fermentation outcomes at T_8 from the initial measurements at T_0.

A compact set of ten variables were identified using literature and spearman's correlations including: strain, temperature regime, MS, cell counts, ethanol, glycerol, organic acids, and NH_4Cl. Categorical features were one-hot encoded, and numerical features were median-imputed and standardized [13]. For the SVR and ANN pipelines, numerical features were also converted to orthogonal components via PCA to reduce collinearity and dimensionality (95% of variance retained). In contrast, PCA was *not* applied to the RF pipeline because tree-based ensembles are inherently robust to collinearity, and this could discard predictive low-variance signals. The models were evaluated on several train-test splits (80:20%, 70:30%, 60:40%, 50:50%) to mimic data limitations and distributional shifts using a grouped splitting method, ensuring that each *Strain-MS-Temperature-regime* combination was contained entirely within either the training or testing set to prevent information leakage.

For the SVR and RF, hyperparameters were tuned on the training set using an exhaustive grid search method with a 5-fold grouped cross-validation (CV) technique. The hyperparameters were selected using a mean MAE fitness across folds. The selected model was refit on the full training set and evaluated on the unseen test set. For the ANN, we replaced the grid search with a genetic algorithm (GA) as a more efficient search of both the architecture and hyperparameter space. The 5-fold grouped CV prevented leakage by keeping all samples from the same combination within a single fold and allowed us to gauge how well the model generalizes to unseen conditions.

Pipeline recommendations were primarily based on mean CV MAE (predictive accuracy) and CV MAE range (robustness) across splits, with the test set used to check for overfitting, and quantified through CV-Test MAE difference (Generalization behaviour). Additionally, deeper analysis was performed on the models out-of-fold (OOF) CV predictions at the 80:20 split, where models were least-data limited.

Results were performed in triplicate, so predictions were averaged within each Strain-MS-Temperature-regime combination. Minebois et al. implicitly define time T_8 as the point at which CO_2 evolution has effectively

ceased and the cumulative weight-loss curve reaches a plateau.

As mentioned, the primary metabolite regressed included concentration ($g\,L^{-1}$) of ethanol and acetic acid. The targets were predicted using separate models to optimize hyperparameters. The secondary metabolite regression included concentrations (m$g\,L^{-1}$) of ethyl acetate, isobutyl alcohol, isoamyl acetate, isoamyl alcohol, ethyl hexanoate, ethyl octanoate, phenylethyl acetate, phenyl ethanol, octanoic acid, and decanoic acid. Given the time constraints we choose predict all secondary metabolite concentrations in one pipeline. However, *scikit learns* SVR function does not natively support multi-output regression, therefore, we employed a wrapper that fits one model per target. Additionally, because some targets were right skewed (e.g. Fig. 1) the pipelines assessed if applying a logarithm transform to the targets improves predictive accuracy.

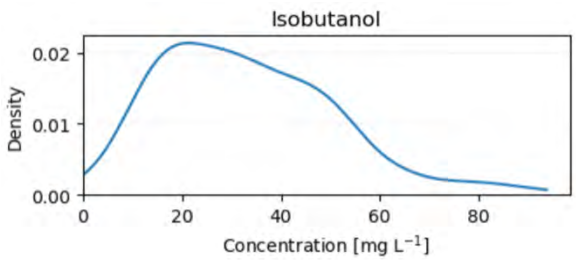

Figure 1: Kernel density estimate of Isobutanol concentrations.

RESULTS AND DISCUSSION

Across train-test splits, the Random Forest pipeline achieved the lowest mean cross-validated (CV) MAE (38.57h) and was the most robust to data partitioning (CV MAE range = 4.09h), outperforming the ANN (CV MAE = 40.59h, CV MAE range = 9.94h) and SVR (CV MAE = 42.5h, CV MAE range = 8.11h) pipelines. Generalization behaviour (CV-Test MAE gap) was also the most stable (-0.30h) indicating that the RF model performances transferred most consistently to the held-out test sets. Although larger training sets typically improve generalization, the RF pipeline achieved its lowest CV MAE at the 50:50 split (36.04h). This likely reflects split-composition effects in a small, heterogeneous, group-split dataset, where the presence/absence of harder regimes can dominate the CV error. Accordingly, we place stronger emphasis on 80:20 pipeline performance, as the split is the least data-limited and best reflects achievable performance given the feature set. The model at the 80:20 split achieved competitive CV metrics with strong test set performance (CV MAE = 39.65 h, test MAE = 32.68 h, both R^2 = 0.91). The model's configuration used 800 trees of unrestricted depth, that sampled 100% of feature at each split, with mild regularisation through minimum leaf size (4) and cost-complexity pruning ($\alpha = 1 \times 10^{-4}$), and chose to log transform targets. The 80:20 RF model, built in mean decrease in impurity (MDI) identified temperature (0.659) as the most important feature, consistent with the temperature-dependent yeast metabolism reported in literature [14].

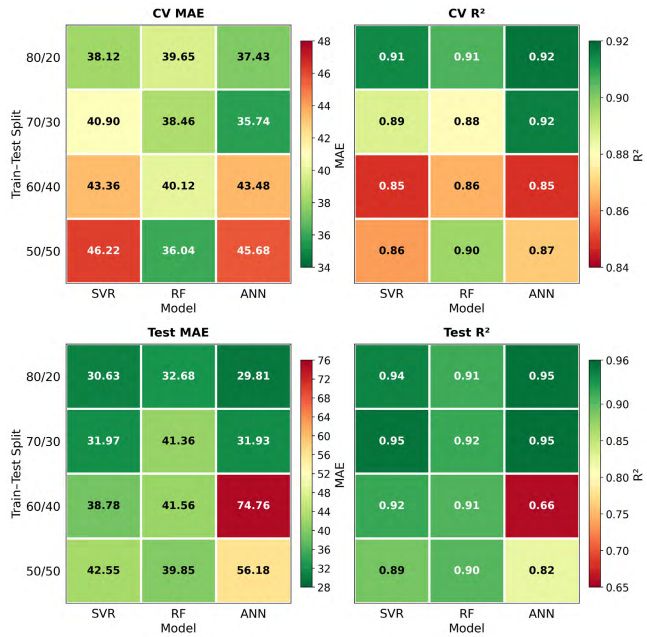

Figure 2: Cross-validated (CV) and test, MAE and R^2 scores for time to complete fermentation (hours) predictions across train-test splits and model pipelines.

Given the ANN's strong performance at higher splits (≥70:30), analysis of out-of-fold CV predictions relative errors (RE) and absolute relative errors (ARE) of both the RF and ANN is appropriate. Models at the 80:20 split were slightly positively biased, with RF achieving a closer to zero mean bias (RE = 0.63% vs +1.76%). RF was also more accurate than the ANN (mean ARE = 7.63% vs 8.60%), with an equivalent standard deviation (SD) (ANN: 0.104, RF: 0.101). Additionally, RF exhibited a heavier tail than the ANN (max. RE = +27.70% vs +22.04%, min. RE = -21.55% vs -18.99%), with a larger strain level bias. The RF model was found to be less biased to isothermal temperate regimes (12℃, 18℃, 25℃) and more biased under TVAR than the ANN (RE: +1.30% vs +1.10%), with the ANN being markedly more biased at 25℃ (RE = +3.98% vs -0.12%). RF exhibited a smaller MS level bias, with the ANN being notably more sensitive to MS210 (RE: +3.18% vs +0.78%), but both displayed poor predictive accuracy at MS210 compared to MS250 (MS210 vs MS250 mean ARE, **ANN: 9.38%** vs 7.95%, RF 9.50% vs **6.07%)**.

This supports recommending RF as the default pipeline, it was most robust to reduced training data and achieved the lowest mean error across splits. The 80:20 RF model was more accurate than the ANN overall, but

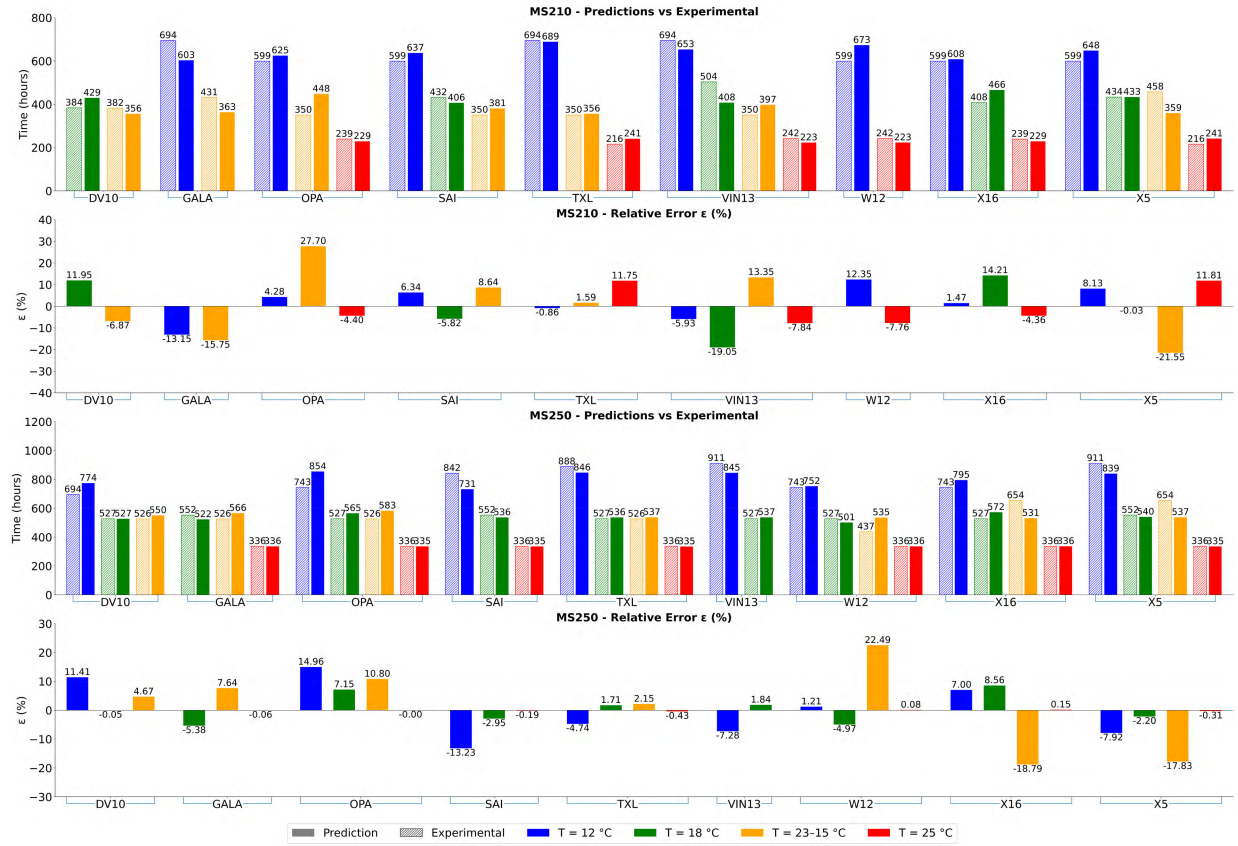

Figure 3. Random Forests 80:20 split out-of-fold cross-validation fermentation duration predictions relative

exhibited less stable tail behaviour. A key limitation is that Minebois et al. does not define an explicit quantitative stopping rule for complete fermentation, so time-to-completion may contain irreducible noise, regardless, the pipeline successfully predicted fermentation duration.

For the ethanol content (%) regression, across train-test splits ethanol was shown to be a learnable relationship on all three pipelines, but performance and stability varied. Over the splits, RF achieved the lowest mean CV MAE (0.563 %), outperforming both the ANN (0.566%) and SVR (0.605 %). However, SVR was the most robust across the train-test splits (smallest CV MAE range = 0.0319%) and exhibited the best generalisation behaviour (smallest CV-Test MAE gap: - 0.0799%).

Figure 3 shows RF achieved lower CV errors compared to SVR, across all splits, with RF being most competitive at the lower splits (≤60:40). The ANN marginally outperformed RF at the high splits (≥70:30) but was less robust overall. The best performing RF model occurred at the 50:50 split, again reflecting impact of the absence of difficult regimes in small, heterogeneous, group-split datasets. At 80:20 split the RF configuration employed 100 trees of restricted depth (max_depth = 10), considered 30% of feature set at each split, and applied regularisation through minimal cost-complexity pruning ($\alpha = 1 \times 10^{-2}$), and chose to log transform targets. The RF model at the 80:20 split's MDI identified must strength (0.695) as the most important feature, consistent with the fermentation's primary reaction. The 80:20 ANN exhibited competitive performance in Figure 3, with a four hidden-layer funnel configuration ($[48, 24, 12, 12]$), moderate drop-out rate (0.163), weak L2 regularisation ($\lambda = 3.36 \times 10^{-5}$) and a relatively high learning rate ($\eta = 3.64 \times 10^{-3}$), with the targets log transformed.

Analysis of out-of-fold CV predictions RE and ARE of both the RF and ANN at the 80:20 split shows that the ANN exhibited a slightly more negative mean bias (RE = -0.37% vs -0.06%), with the RF achieving a greater average predictive accuracy (mean ARE: RF = 3.07%, ANN = 3.56%). RF achieved a slightly tighter error distribution (SD = 0.0429 vs 0.0459), but both models exhibited stronger positive tails (max. RE: RF = +16.24%, ANN = +17.01%, min. RE: RF = -8.67%, ANN = -6.96%). At the temperature level, both models were unbiased under 12°C and TVAR, unpredicted at 18°C and overpredicted at 25°C, with the ANN exhibiting markedly larger under-predictions than RF at 18°C (RE = -1.86% vs -0.94%). However, RF was more accurate in absolute terms at every temperature regime. Strain level biases were small on both models (RE $\sim \pm 2\%$), but on harder to predict strains (OPA, X16) RF was consistently more accurate (lower mean ARE). Finally, both models were less

accurate at MS210 (mean ARE: ANN = 4.45%, RF = 4.00%) than at MS250 (mean ARE: ANN = 2.82%, RF = 2.28%), with RF being more accurate at both must strengths, with negligible MS level bias.

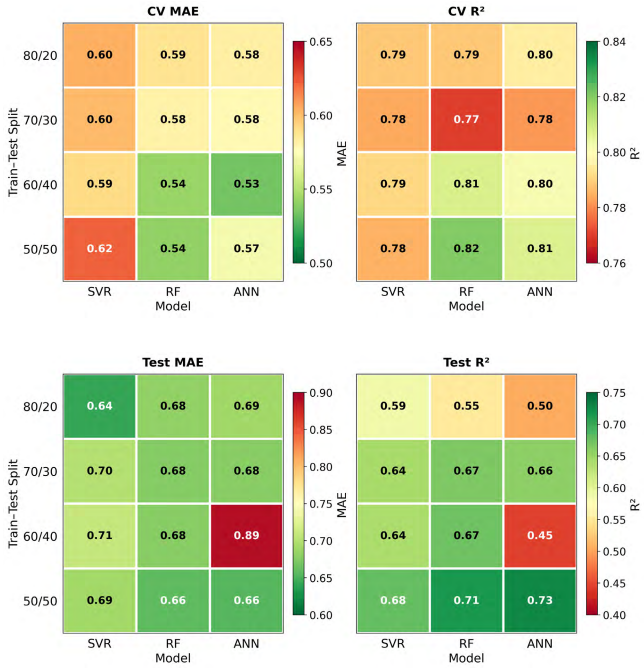

Figure 4: CV and test, MAE and R^2 scores for ethanol (%) predictions across train-test splits and model pipelines

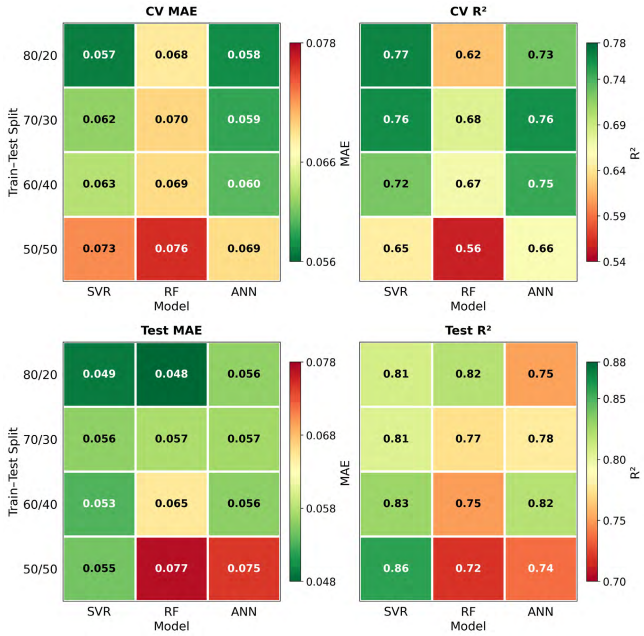

Figure 5: CV and test MAE and R^2 scores for acetate conc. (gL^{-1}) predictions across train-test splits and model pipelines.

Overall, the evidence supports recommending RF as the default pipeline because, although, SVR was the most robust to reduced training data, RF achieved the lowest mean error across splits and under the least data limited regime (80:20), RF achieved greater predictive accuracy, a tighter error distribution, and a smaller temperature and must strength dependency. However, the ANN produced fewer outlier (beyond $IQR \times 1.5$) predictions. A methodological refinement is to model ethanol yield (final – initial) instead of predicting final ethanol with initial ethanol included as a feature, which conflates baseline ethanol with yield.

For the acetate concentration task, across splits the ANN pipeline achieved the lowest mean CV MAE (0.0615 gL^{-1}) and the smallest CV-Test MAE gap (0.0007 gL^{-1}), therefore indicating stable generalization. The RF pipeline was the most robust to splits (smallest CV MAE range: 0.0075 gL^{-1}), however, achieves the worst predictive accuracy (mean CV MAE: 0.0708 gL^{-1}) and explained variance. The SVR pipeline achieved a close second in predictive accuracy across splits (mean CV MAE: 0.0638 gL^{-1}), whilst being the least robust (CV MAE range: 0.0155 gL^{-1}) and exhibiting the largest CV-Test MAE gap (0.0105 gL^{-1}), such that generalization was much stronger, indicating training **sampling noise** (test set was easier).

The SVR configuration at the 80:20 split used an **RBF kernel** with **C=100, ϵ=0.05$, γ=0.003**, and **no target transformation. This corresponds to a smooth nonlinear mapping with relatively strong penalization of training errors. The ANN at the 80:20 split used** a **shallow network** with one hidden layer ([5])**, a** low dropout rate (0.09), moderate L2 regularisation ($\lambda = 2.3 \times 10^{-3}$) and a relatively high learning rate ($\eta = 5.02 \times 10^{-3}$), with a log-transformed target.

Analysis on the 80:20 CV out-of-fold predictions indicate SVR was more positively biased than the ANN (RE = +1.25% vs +0.22%), but both achieved negative median RE (SVR −0.63%, ANN −1.52%), therefore on typical cases the models underpredicted, especially the ANN. The mean ARE highlights that the SVR was more accurate (mean ARE: 7.80% vs 9.15%) than the ANN, and given a typical condition is 52% more accurate (median ARE: 5.42% vs 8.22%). Additionally, the SVR's error distribution was tighter (SD = 0.0999 vs 0.1189). The ANN shows greater strain level bias than SVR, with SVR achieving a greater predictive accuracy for five of the nine strains (OPA, TXL, VIN13, X16, X5). The ANN exhibited particularly weak performance on X16. Both models were less accurate at **MS250** than **MS210**, with higher mean absolute relative error (ANN: 9.73% vs 8.46%, SVR: 8.39% vs 7.09%). Across MS, SVR was more accurate with a small positive bias and the ANN's bias changed sign, **overpredicting** at MS210 (mean RE = +2.56%) and **underpredicting** at MS250 (mean RE = −1.75%).

Finally, the evidence supports recommending the ANN pipeline when mean predictive accuracy, with moderate robustness across splits is preferred. However, the

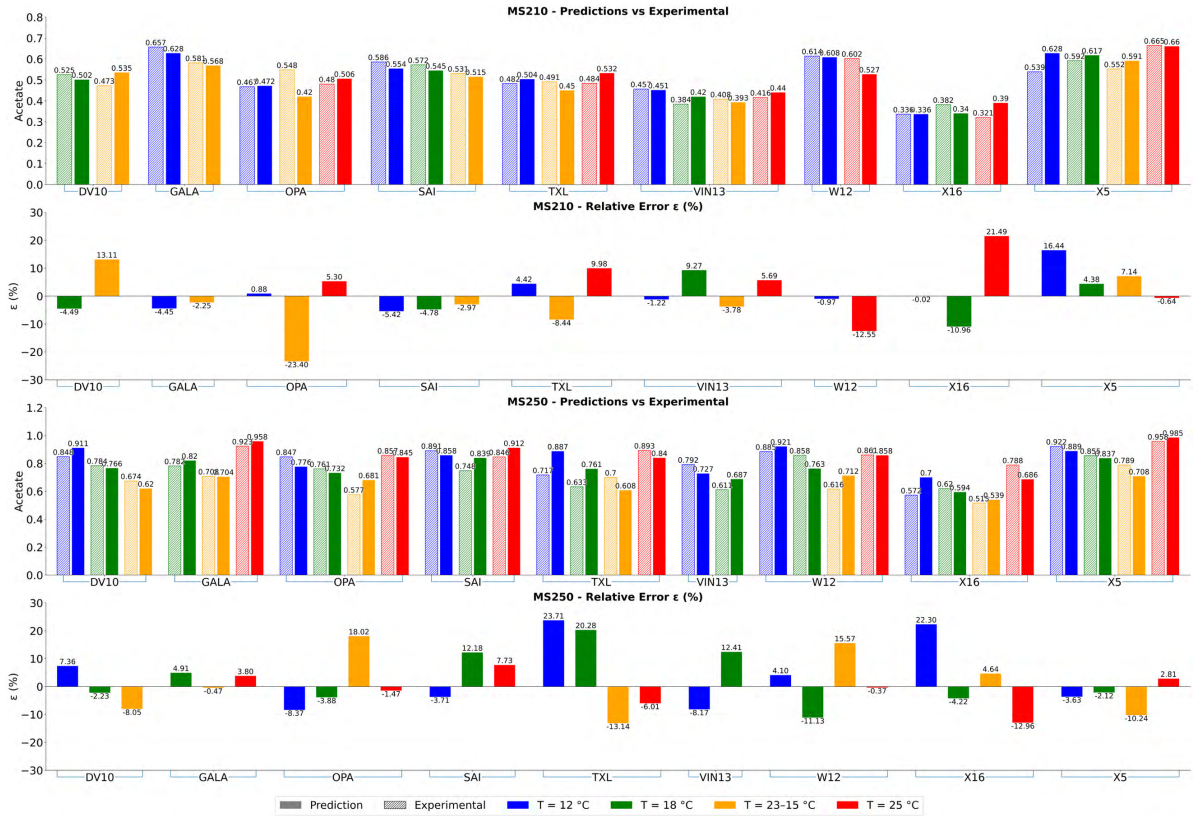

Figure 6. Support vector regression 80:20 split out-of-fold cross-validation acetate concentration (mgL^{-1}) predic–tions relative errors.

ANN's least data limited model (80:20) exhibited a wider error distribution and lower absolute predictive accuracy, with instability on the X16 strain, stronger MS dependence, and higher errors at **12°C and 25°C. Therefore, if predictive accuracy under non-data limited regimes is preferred, the SVR pipeline is recommended for its lower mean ARE, tighter error distribution and smaller MS and temperature dependence.**

For the secondary metabolite task, across the train test splits, SVR was both the most accurate and robust pipeline. On every split it achieved the lowest mean CV MAE across the 10 targets (4.97 mgL^{-1}) and highest mean CV R^2 (0.22 mgL^{-1}). However, the ANN pipeline achieved the smallest mean CV-test MAE gap (-0.05 mgL^{-1}), indicating the most stable generalization. Regardless SVR was the most competitive on the held-out test set.

Across the splits, the pipeline selected the same SVR configuration, an RBF kernel, $C = 100, \epsilon = 0.1$ and $\gamma = 0.001$, with a log-transform of the target variable, aside from at the 70:30 split, where the pipeline reduced ϵ (0.05). When coupled with the almost identical model configurations, the monotonic improvement of mean CV MAE and R^2 as training fraction increased is consistent with gradual loss of predictive signal under reduced sample size, rather than model instability.

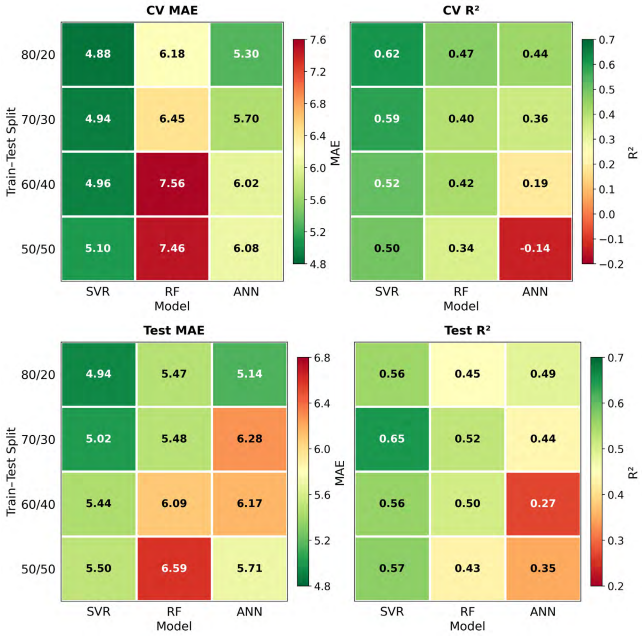

Figure 7: Mean CV and test, MAE and R^2 scores across secondary metabolite (mgL^{-1}) predictions from different train-test splits and model pipelines.

Although the mean results across the target show that the RF pipeline was most accurate and robust to train-test splits, the per target analysis at the 80:20 split reveals large target heterogeneity (Table 1). Specifically higher alcohols and their related esters (**PhenEtoh, Isobutanol, IsoamylAlcohol, PhenEthAcetate, IsoamylAcetate**) exhibit high learnability, whereas several medium-chain fatty acids and ethyl esters (OctanoicAcid, EthylHexanoate, DecanoicAcid, EthAcetate, EthOctanoate) remained difficult to learn.

Table 1: Per-target out-fold CV performance of SVR at the 80:20 train–test split, reported as normalized MAE (nMAE) and R^2.

Target	nMAE	R^2
PhenEtoh	0.144	0.886
Isobutanol	0.151	0.835
PhenEthAcetate	0.224	0.834
IsoamylAlcohol	0.115	0.814
IsoamylAcetate	0.286	0.759
OctanoicAcid	0.138	0.598
EthylHexanoate	0.237	0.444
DecanoicAcid	0.278	0.436
EthAcetate	0.133	0.343
EthOctanoate	0.271	0.206

The target learnability trend could be partially explained through feature set selection as medium-chain fatty acids and related ethyl esters depend on absent features like lipid/fatty-acid availability, thus limiting predictability.

Of the targets that exhibit high explained variance (R^2) PhenEthAcetate has an elevated nMAE which indicates predictions capture variance with poor practical accuracy. This can be caused through systematic bias and is common in targets with a small range, like PhenEthAcetate, where even modest absolute errors are large fraction of the data's range.

Analysing PhenEtoh, Isobutanol and IsoamylAlcohol out-offold CV predictions for their high learnability and low nMAE, we find that the three targets exhibit near-unbiased central behavior (median RE = -0.08-0.71%), but regime dependent heavy tails. For example, all targets performed relatively poorly on 12**°C** (larger positive bias, higher MARE, and outlier clustering). In terms of absolute accuracy and stability, the model performed best on IsoamylAlcohol (mean ARE = 11.05%, SD = 0.162, max. RE = +66.42%), while Isobutanol and PhenEtoh are less accurate on average (mean ARE = 16.07% and 16.32%) and with longer tails, especially Isobutanol which exhibited the largest upside tail risk (max RE = +157.58%). The mean ARE was higher at MS210 for IsoamylAlcohol and PhenEtoh regression, whereas Isobutanol regression showed a higher mean ARE at MS250. Additionally, the **SAI strain is repeatedly the most difficult to predict across the three targets,** indicating a strain specific failure.

The heterogeneous predictability across the ten metabolites indicates that a single, fixed model configuration is unlikely to be optimal for all targets. This reinforces concerns in the literature that applying static model (non-case specific) configurations limits industrial practicality.

Given the time constraints, we used a task-agnostic feature set and fitted one pipeline across all 10 targets. As a result, the chosen model configurations were likely a compromise that limited target instability on the least learnable targets, rather than refining configurations for highly learnable targets. We would therefore recommend future work expands the feature set to fatty acids/lipids, and include additional assimilable nitrogen sources (Lysine, Proline, Ornithine) and either, fit a pipeline per target or pipelines to groups of targets (higher alcohols, acetate esters, ethyl esters, medium-chain fatty acids).

CONCLUSION

This study successfully demonstrated that machine learning pipelines can map controllable fermentation conditions to fermentation kinetics and endpoint metabolite concentrations in *S. cerevisiae. Moving away from static ML models and motivated by industry constraints, we* developed pipelines that designed and trained three different ML models (SVR, RF, ANN), across four tasks (time to complete fermentation, and ethanol, acetate, and ten secondary metabolite concentrations). We evaluated the pipelines ability to cope with limited data and distributional shifts through evaluation on multiple train–test splits. We employed a *GroupKFold* splitting strategy to prevent information leakage and used five-fold cross-validation to improve generalization performance. Primary assessment criteria included mean MAE, CV MAE range, competitive R^2, and held-out tests were used to assess generalisation stability.

A compact ten variable feature set of operationally accessible predictors was identified: strain, temperature regime, must strength, biomass counts, early metabolites, and a key nitrogen source. This feature set was highly predictive for time, ethanol, and acetate tasks, supporting practical deployment, however, secondary metabolite predictions would likely benefit from additional predictors and per-target pipeline fits.

Performance was task-dependent, for **time-to-completion,** the RF pipeline is recommended as it

achieved the lowest mean CV error and highest robustness to split variation, although the ANN pipeline was competitive when sufficient training data was available (≥70:30). For **ethanol,** again the RF pipeline is the primary recommendation, as it achieved the lowest mean CV MAE with stable generalisation and the RF at the 80:20 split compared to the ANN achievd nearer-zero bias, with weaker temperature and must-strength dependence. For **acetate,** if robustness to split variations with competitive mean accuracy is the priority, the ANN is recommended. However, at the least data limited regime (80:20) the ANN exhibited lower accuracy and a wider error distribution than SVR, with stronger temperature and MS dependence. Therefore, if accuracy under high data availability is the priority, SVR is preferred due to lower mean ARE at the 80:20 split, a tighter error distribution and weaker MS/temperature dependence. For **secondary metabolites,** the per-target SVR models was both most accurate and most robust pipeline, highlighting the value of fitting separate models for targets with heterogeneous signal strength. Higher alcohols and associated acetate esters were highly learnable (R^2 ≥0.75), whereas several ethyl esters and medium-chain fatty acids remained weakly predicted, consistent with missing predictors such as lipids/fatty-acids.

Key methodological refinements for future work include, defining a quantitative fermentation endpoint, calculating and predicting ethanol yield (final–initial) to isolate production, adopting task-specific feature selection, expanding the secondary-metabolite feature space (Lysine, Proline, Ornithine, lipids/fatty acids), and increasing the GA search budget. Finally, we believe the recommended pipelines could be embedded in an optimisation framework to support industry and, to our knowledge, the present work was the first to train an machine learning models on a time-varying temperature regime in wine manufacturing.

ACKNOWLEDGEMENTS

D.I.G acknowledges support from UKRI/EPSRC (RAPID: ReAl-time Process Modelling & Diagnostics: Powering Digital Factories, EP/V028618/1) and a Royal Society International Exchanges grant (2023-25).

REFERENCES

1. International Organisation of Vine and Wine (OIV). State of the world vine and wine sector in 2024. *OIV*, April 2025. Accessed 22 November 2025.
2. Minebois R, Ramírez?Aroca L, Balsa?Canto E, Querol A. Toward less energy?consuming alcoholic fermentations in oenology: a laboratory?scale case study. Food Frontiers 6:2931-2941 (2025). https://doi.org/10.1002/fft2.70103
3. Du Q, Ye D, Zang X, Nan H, Liu Y. Effect of low temperature on the shaping of yeast-derived metabolite compositions during wine fermentation. Food Research International 162:112016 (2022). https://doi.org/10.1016/j.foodres.2022.112016
4. de Castro M, Baptista J, Matos C, Valente A, Briga-Sá A. Energy efficiency in winemaking industry: challenges and opportunities. Science of The Total Environment 930:172383 (2024). https://doi.org/10.1016/j.scitotenv.2024.172383
5. Gupta Y. Selection of important features and predicting wine quality using machine learning techniques. Procedia Computer Science 125:305-312 (2018). https://doi.org/10.1016/j.procs.2017.12.041
6. Sandeep Bharadwaj Mannapur . Understanding data drift and concept drift in machine learning systems. Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol 11:318-330 (2025). https://doi.org/10.32628/cseit25111239
7. Rodrigues AJ, Raimbourg T, Gonzalez R, Morales P. Environmental factors influencing the efficacy of different yeast strains for alcohol level reduction in wine by respiration. LWT 65:1038-1043 (2016). https://doi.org/10.1016/j.lwt.2015.09.046
8. Du Toit W. The enumeration and identification of acetic acid bacteria from south african red wine fermentations. International Journal of Food Microbiology 74:57-64 (2002). https://doi.org/10.1016/s0168-1605(01)00715-2
9. Montesinos López OA, Montesinos López A, Crossa J. Multivariate statistical machine learning methods for genomic prediction. Springer International Publishing (2022). https://doi.org/10.1007/978-3-030-89010-0
10. Biau G, Scornet E. A random forest guided tour. TEST 25:197-227 (2016). https://doi.org/10.1007/s11749-016-0481-7
11. Suzuki K. Artificial neural networks - methodological advances and biomedical applications. InTech (2011). https://doi.org/10.5772/644
12. Hameed AA, Karlik B, Salman MS. Back-propagation algorithm with variable adaptive momentum. Knowledge-Based Systems 114:79-87 (2016). https://doi.org/10.1016/j.knosys.2016.10.001
13. Mu, A.C., n.d. Introduction to Machine Learning with Python.
14. Molina AM, Swiegers JH, Varela C, Pretorius IS, Agosin E. Influence of wine fermentation temperature on the synthesis of yeast-derived volatile aroma compounds. Appl Microbiol Biotechnol 77:675-687 (2007). https://doi.org/10.1007/s00253-007-1194-3

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Global Optimization of Robust AC OPF

Yuhui Yin and Vassilis M. Charitopoulos*

Department of Chemical Engineering, Sargent Centre for Process Systems Engineering, UCL (University College London), Torrington Place, London WC1E 7JE, UK
* Corresponding Author: v.charitopoulos@ucl.ac.uk

ABSTRACT

Ensuring reliable operations of modern power systems under uncertainty remains a key challenge, particularly due to the non-convex nature of Alternating Current (AC) power flow equations and the presence of high-impact disturbances from load and renewable generation fluctuations. In this work, we address the robust AC Optimal Power Flow (AC OPF) problem by developing a robust spatial branch-and-bound (RsBB) algorithm. Robustness is achieved by identifying worst-case uncertainty realizations and iteratively incorporating robust cuts to eliminate constraint violations. To accelerate convergence and tighten bounds, Optimization-Based Bound Tightening (OBBT) and Feasibility-Based Bound Tightening (FBBT) techniques are integrated into the framework. The proposed method yields global robust solutions with certified optimality gaps below 0.01% across standard PGLib test cases.

Keywords: Nonconvex Robust Optimization, AC OPF, Cutting planes, Bound Tightening, Uncertainty, Global Optimization

INTRODUCTION

The AC optimal power flow (ACOPF) problem is a fundamental tool in power-system operations, computing minimum cost dispatch subject to power balance equations and network limits (e.g., voltage and thermal constraints) [1]. In practice, ACOPF is often solved deterministically using forecasts of demand and renewable injections. With the growing penetration of wind and solar, forecast errors can be substantial, so a deterministic operating point may violate constraints in real time (e.g., line overloads or voltage excursions) [2, 3]. Such violations jeopardize security and reliability and may push the system outside its secure operating limits. This motivates robust OPF formulations that remain feasible for all realizations of uncertainty within a prescribed set [4].

Existing OPF methods address uncertainty mainly via: (i) stochastic programming, which optimizes over sampled scenarios but can become computationally prohibitive as the number of scenarios grows, and is sensitive to scenario-generation and reduction errors [5, 6]; (ii) chance-constrained and distributionally robust chance-constrained models, which provide probabilistic security but typically rely on DC approximations or convex relaxations to gain tractability, thereby sacrificing exact AC feasibility and certifiable global optimality for the original nonconvex AC OPF [7, 8]; and (iii) robust optimization, which enforces feasibility for all uncertainty realizations, but in AC settings is still predominantly relaxation-based (SOCP/SDP/DC or convex restrictions), leading to conservatism and again losing guarantees of global optimality and exact AC feasibility [9-11].

Despite this progress, to the best of our knowledge there is still no method that simultaneously guarantees (a) robust feasibility for all values in a continuous uncertainty set and (b) certifiable global optimality for the full nonconvex ACOPF. Prior approaches either relax the AC physics or robustness to obtain tractability, or provide only scenario-based or probabilistic guarantees that do not cover the entire uncertainty set. This gap motivates the development of global-optimization framework tailored to robust ACOPF.

Such framework draws on several complementary lines of work. In semi-infinite global optimization, restriction and outer-approximation ideas provide a principled way to enforce constraints over continuous parameter domains [12]. In nonlinear robust optimization, generalized cutting-set procedures similarly generate many worst-case scenarios to certify feasibility over uncertainty sets [13]. On the algorithmic side, global solvers rely on systematic strengthening methods, including convex underestimators for trigonometric nonlinearities

[14] and optimality-based domain reduction to contract variable bounds during the search [15]. Complementing these methodological advances, recent power-system studies demonstrate that tailored global/MINLP strategies can be made effective at scale, e.g., for security-constrained unit commitment [16]. Together with broader developments in global optimization for MINLP [17], these ingredients provide the foundation for our robust spatial branch-and-cut approach for certifiable global robust ACOPF.

In this paper, we build on our previous work [18] and present a framework to globally solve the robust ACOPF via bound tightening and robust spatial branch-and-cut. The main contributions are as follows:

- BT–RsBB for global robust ACOPF: a robust spatial branch-and-cut scheme that iteratively refines convex relaxations while branching over the nonconvex region, converging to a globally robust optimal solution.
- Integrated bound tightening: a combination of optimization-based bound tightening (OBBT) and feasibility-based bound tightening (FBBT), where FBBT is applied within OBBT iterations to strengthen relaxations and shrink the search space at low overhead.
- Efficient trigonometric relaxations: an improved relaxation strategy that leverages precomputed sine-envelope parameters to rapidly generate tight linear relaxations of trigonometric constraints during the search.
- Parallel cut generation: concurrent generation of multiple cutting planes at nodes with uncertainty violations, improving node relaxations and accelerating pruning of robust-infeasible regions.

The remainder of the paper is organized as follows. Section **II** presents the ACOPF formulation and the relaxations used in the global algorithm. Section **III** details the robust spatial branch-and-cut procedure, including infeasibility testing, branching, cut generation, and bound tightening. Section **IV** reports numerical results and comparisons. Section **V** concludes.

PROBLEM FORMULATION

Deterministic AC-OPF (Polar Formulation)

We consider the standard AC optimal power flow (ACOPF) in polar coordinates, where each bus $k \in \mathcal{B}$ has voltage magnitude V_k and phase angle θ_k, and each generator $i \in \mathcal{G}$ controls real and reactive outputs (P_i^g, Q_i^g). Let (P_{km}, Q_{km}) denote the sending-end real/reactive line flows on branch $(k,m) \in \mathcal{L}$, with line parameters $(G_{km}, B_{km}, B_{km}^c)$[19]. The deterministic ACOPF minimizes a quadratic generation cost subject to (i) nodal real/reactive power balance (eq. 1b – 1c), (ii) voltage and generator operating limits (eq. 1d – 1f), (iii) nonlinear AC power-flow relations in polar form (eq. 1g), (iv) angle-difference bounds (eq. 1h), and (v) apparent-power (thermal) limits (eq. 1i). For compactness, we refer the interested read for the full formulation in [19-22], and succinctly present the problem below:

$$\min \quad \sum_{i\in\mathcal{G}} \left(c_{2i}(P_i^g)^2 + c_{1i}P_i^g + c_{0i}\right) \tag{1a}$$

$$\text{s.t.} \quad \sum_{i\in\mathcal{G}(k)} P_i^g - P_k^d - G_s V_k^2 = \sum_{km\in\mathcal{L}^{+(k)}} P_{km} + \sum_{km\in\mathcal{L}^{-(k)}} P_{mk}, \forall k \in \mathcal{B} \tag{1b}$$

$$\sum_{i\in\mathcal{G}(k)} Q_i^g - Q_k^d + B_s V_k^2 = \sum_{km\in\mathcal{L}^{+(k)}} Q_{km} + \sum_{km\in\mathcal{L}^{-(k)}} Q_{mk}, \forall k \in \mathcal{B} \tag{1c}$$

$$\underline{V}_k \le |V_k| \le \bar{V}_k, \forall k \in \mathcal{B} \tag{1d}$$

$$\underline{P}_i \le P_i^g \le \bar{P}_i, \forall i \in \mathcal{G} \tag{1e}$$

$$\underline{Q}_i \le Q_i^g \le \bar{Q}_i, \forall i \in \mathcal{G} \tag{1f}$$

$$c_{km} = V_k V_m cos(\theta_k - \theta_m), \forall (k,m) \in \mathcal{L} \tag{1g.1}$$

$$s_{km} = V_k V_m sin(\theta_k - \theta_m), \forall (k,m) \in \mathcal{L}$$

$$P_{km} = G_{km}V_k^2 - G_{km}c_{km} - B_{km}s_{km}, \forall (k,m) \in \mathcal{L} \tag{1g.2}$$

$$Q_{km} = -(B_{km} + B_{km}^c)V_k^2 - G_{km}s_{km} + B_{km}c_{km}, \forall (k,m) \in \mathcal{L} \tag{1g.3}$$

$$P_{mk} = G_{km}V_m^2 - G_{km}c_{km} + B_{km}s_{km}, \forall (k,m) \in \mathcal{L} \tag{1g.4}$$

$$Q_{mk} = -(B_{km} + B_{km}^c)V_m^2 + G_{km}s_{km} + B_{km}c_{km}, \forall (k,m) \in \mathcal{L} \tag{1g.5}$$

$$\theta_k - \theta_m \in [\underline{\theta}_{km}, \bar{\theta}_{km}], \forall (k,m) \in \mathcal{L} \tag{1h}$$

$$P_{km}^2 + Q_{km}^2 \le (\bar{S}_{km})^2, \forall (k,m) \in \mathcal{L} \tag{1i.1}$$

$$P_{mk}^2 + Q_{mk}^2 \le (\bar{S}_{km})^2, \forall (k,m) \in \mathcal{L} \tag{1i.2}$$

Note that, a reference angle constraint $\theta_{\text{ref}} = 0$ is imposed to ensure uniqueness of the polar representation [23].

Convex Relaxations for Global Optimization

We employ convex relaxations to provide valid lower

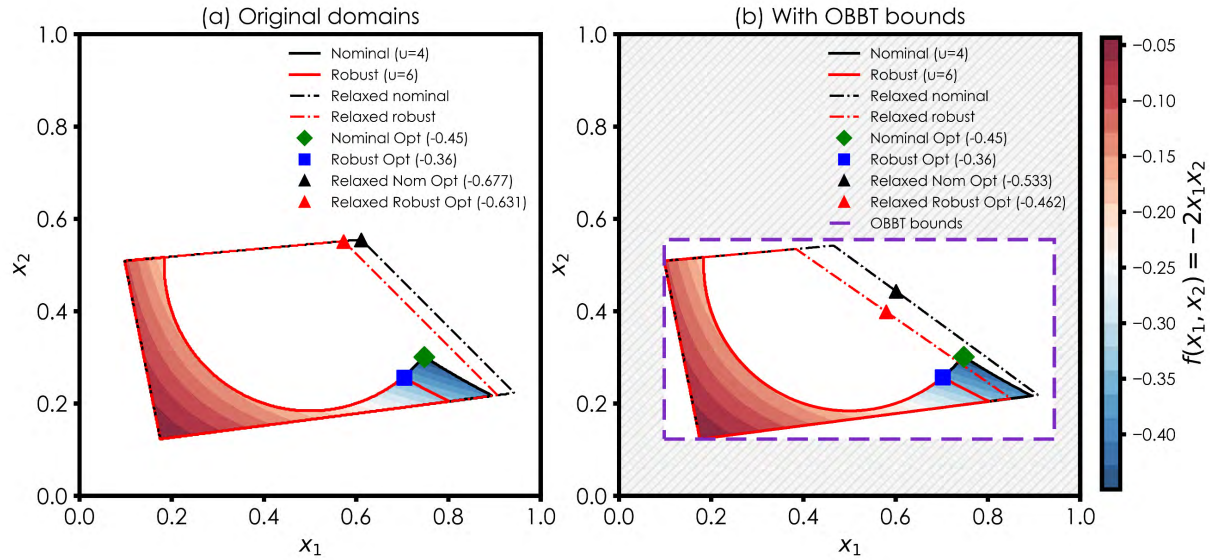

Figure 1. Illustration of the relationship between robust optimization and bound tightening: (a) original domains showing the nominal and robust feasible regions; (b) reduced domains after bound tightening, where the shaded region indicates the eliminated search space.

bounds within branch-and-bound. We adopt a quadratic-convex (QC) relaxation of the polar ACOPF [20, 22, 24]. The key idea is to lift the nonconvex bilinear and trigonometric terms and replace them with convex envelopes. In particular, we introduce lifted variables:

$$V_{km} \approx V_k V_m, V_{kk} \approx V_k^2, \cos_{km} \approx \cos(\theta_{km}) \tag{2}$$

$$\sin_{km} \approx \sin(\theta_{km}), \theta_{km} := \theta_k - \theta_m \tag{3}$$

We further define β_{km} and γ_{km} to represent the lifted products $V_{km}\cos_{km}$ and $V_{km}\sin_{km}$, respectively. Each lifted relation is relaxed using: (i) McCormick envelopes for bilinear products, (ii) convex quadratic envelopes for V_k^2, and (iii) convex/concave envelopes for $\sin(\cdot)$ and $\cos(\cdot)$ over the phase angle difference interval $\theta_{km} \in [\theta_{km}^l, \theta_{km}^u]$. Collectively, these constraints define a relaxation set $\mathcal{E}_{\mathrm{MC/QC}}(V^l, V^u, \theta^l, \theta^u)$. To further tighten the QC relaxation, we include a standard second-order cone (SOC) constraint

$$\beta_{km}^2 + \gamma_{km}^2 \leq V_{kk} V_{mm}, (k, m) \in \mathcal{L} \tag{4}$$

which couples the lifted terms and strengthens the QC bound.

Improved Trigonometric Relaxation

The tightness of the QC relaxation is often dominated by the quality of the envelopes related to the trigonometric functions present. Instead of a single generic sine envelope, we use a tighter piecewise outer approximation of $\sin(\theta)$ over $[\theta_{km}^l, \theta_{km}^u]$ that adapts to the interval sign pattern. Let the secant and tangent (supporting) lines be:

$$f^{\sec}(\theta^l, \theta^u) := \left(\frac{\sin(\theta^u) - \sin(\theta^l)}{\theta^u - \theta^l}\right)(\theta - \theta^l) + \sin(\theta^l) \tag{5}$$

$$f^{\tan}(t) := \sin(t) + \cos(t)(\theta - t) \tag{6}$$

If the interval does not cross zero, the global convexity/concavity of $\sin(\cdot)$ yields a strong two-sided envelope. If $\theta_{km}^l < 0 < \theta_{km}^u$, we combine a secant with a small set of tangent supports (e.g., at an endpoint and a midpoint on each active side) to obtain a strictly tighter polyhedral outer approximation:

$$\begin{aligned} \sin_{km} &\leq f^{\sec}(\theta_{km}^l, \theta_{km}^u), \\ \sin_{km} &\geq \max\{ f^{\tan}(t) : t \in T_{km} \} \end{aligned} \tag{7}$$

(or the symmetric variant with reversed inequalities, depending on the interval regime). For efficiency within RsBB, the tangent locations T_{km} and any regime-dependent transition parameters are precomputed and stored; during branching and bounds tightening, only the active constraints for the detected region are enforced, while interval bounds are updated via mutable parameters. This construction yields tighter node relaxations and accelerates convergence by reducing feasible domains and strengthening lower bounds.

ROBUST OPTIMIZATION METHODOLOGY

We consider the original nonconvex robust optimization problem

Table 1: Quality and runtime results of robust ACOPF (optimality gaps are expressed in %). Results are obtained under a 1% perturbation. Times are CPU seconds. For BT–RsBB, the reported time is shown as the summation of bound-tightening time and subsequent RsBB time. IR denotes the improved relaxation and RR denotes the reference relaxation.

Box uncertainty								
	Optimality gap (%)				CPU (s)			
Case	RsBB		BT-RsBB		RsBB		BT-RsBB	
	IR	RR	IR	RR	IR	RR	IR	RR
case3_lmbd	0.010	0.010	0.009	0.009	54.0	55.1	19.0	13.4
Case5_pjm	10.633	10.765	0.030	0.034	7200.0	7200.1	7200.0	7200.0
case14 ieee	0.034	0.034	0.009	0.009	7200.0	7200.0	27.2	26.2
case30 as	0.049	0.049	0.008	0.008	7200.0	7201.0	3160.5	3305.2
case73 ieee rts	0.026	0.026	0.010	0.010	7200.4	7200.0	258.2	246.7
case197 snem	0.030	0.030	0.010	0.010	7200.0	7202.0	1952.6	7200.0

Polyhedral uncertainty								
	Optimality gap (%)				CPU (s)			
Case	RsBB		BT-RsBB		RsBB		BT-RsBB	
	IR	RR	IR	RR	IR	RR	IR	RR
case3_lmbd	0.010	0.010	0.007	0.007	55.3	53.5	16.6	17.8
Case5_pjm	10.778	10.846	0.030	0.035	7200.0	7200.1	7200.0	7200.0
case14 ieee	0.034	0.034	0.009	0.009	7200.0	7200.0	33.0	29.8
case30 as	0.049	0.049	0.008	0.008	7200.0	7201.0	3127.1	3330.7
case73 ieee rts	0.027	0.027	0.010	0.010	7200.0	7209.9	258.1	245.6
case197 snem	0.030	0.030	0.010	0.010	7200.0	7201.1	1887.2	7200.4

$$\min_{x\in\mathcal{X}} f(x) \ \text{s.t.}\, g(x,u) \le 0, \forall u \in \mathcal{U}, \tag{8}$$

where x denotes the decision variables, $\mathcal{X}$ is a bounded feasible region, u is an uncertain parameter lying in a prescribed uncertainty set $\mathcal{U}$, and $g(\cdot,\cdot)$ collects the uncertain constraints. Since global solution of (8) is challenging for nonconvex models, we adopt the robust spatial branch-and-bound (RsBB) framework in [25], which provides finite convergence to a globally robust optimal solution under standard regularity assumptions (boundedness, continuity, and tractability of the inner maximization).

RsBB Framework: Master Problems and Infeasibility Test

At each spatial branch-and-bound node, RsBB maintains: (i) a nonconvex master problem that enforces robustness over a current finite subset of uncertainty realizations, and (ii) a convex relaxation that yields valid lower bounds. Specifically, for a cutting-set $\hat{\mathcal{U}} \subseteq \mathcal{U}$, the node master is

$$(P_{\text{nonc}}) \min_{x\in\mathcal{X}} f(x) \ \text{s.t.}\, g(x,u) \le 0, \forall u \in \hat{\mathcal{U}}, \tag{9}$$

whose robust-feasible solutions provide incumbents (upper bounds). Next, we solve a convex relaxation

$$(P_{\text{conv}}) \min_{x\in\mathcal{X}} f_{\text{conv}}(x) \ \text{s.t.}\, g_{\text{conv}}(x,u) \le 0, \forall u \in \hat{\mathcal{U}}, \tag{10}$$

where f_{conv} and g_{conv} are convex relaxations of f and g. The optimal value of P_{conv} is a valid node lower bound, and the global lower bound is the minimum across open nodes, as in standard spatial B&B [25].

Robust feasibility of a candidate x^* from P_{nonc} is checked by an infeasibility test (inner maximization)

$$(P_{\text{it}}) \ \max_{u\in\mathcal{U}} g(x^*,u), \tag{11}$$

which returns a worst-case realization u^*. If $g(x^*,u^*) > \varepsilon_{\text{it}}$, then x^* is not robustly feasible and the violating scenario u^* is added to $\hat{\mathcal{U}}$ (equivalently, we append the cut $g(x,u^*) \le 0$ to both P_{nonc} and P_{conv} at the node). This cutting-set loop is repeated until no violation is found (within ε_{it}), after which the resulting x^* is robust-feasible over $\mathcal{U}$ up to tolerance. The outer spatial branching loop guarantees global optimality by maintaining valid bounds and partitioning variable domains; full algorithmic details and convergence arguments follow [25].

Uncertainty in AC OPF

In this work, uncertainty enters through bus

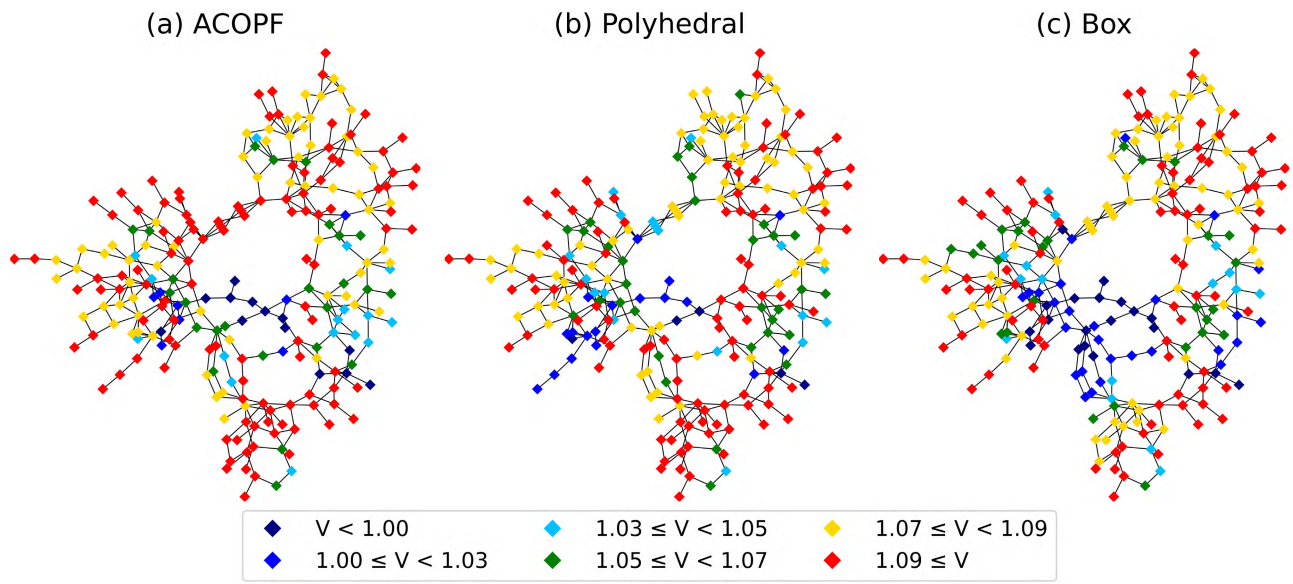

Figure 2. Voltage magnitude profiles for case240_pserc: (a) nominal ACOPF solution without uncertainty, (b) robust ACOPF under polyhedral uncertainty, and (c) robust ACOPF under box uncertainty, obtained by BT–RsBB with the improved relaxation. Bus colors indicate voltage bins (p.u).

demands. Let $(P^d_{k,\text{nom}}, Q^d_{k,\text{nom}})$ be nominal loads and define

$$\tilde{P}^d_k = (1+\xi_k)P^d_{k,\text{nom}},\ \tilde{Q}^d_k = (1+\xi_k)Q^d_{k,\text{nom}}, \xi \in \mathcal{U}. \qquad (9)$$

We consider standard convex uncertainty sets, either a box $\mathcal{U}_{\text{box}} = \{\xi : |\xi_k| \le \Psi_k, \forall k\}$, or a budgeted polyhedral set $\mathcal{U}_{\text{poly}} = \{\xi : |\xi_k| \le \Psi_k, \forall k, \sum_{k\in\mathcal{B}} |\xi_k| \le \Gamma\}$. Robustness is enforced by replacing the (deterministic) active/reactive power-balance equalities in (1) with robust constraints $\forall \xi \in \mathcal{U}$, and update violated constraints via P_{it} as described above. In our implementation, the inner maximization is solved using the current candidate operating point $(P^{g*}, Q^{g*}, V^*, \theta^*, \dots)$, and any violating realization ξ^* generates a cut corresponding to the power-balance constraint at ξ^*.

Bounds Tightening

To accelerate RsBB on the nonconvex AC OPF, we apply bound tightening to strengthen relaxations without materially increasing model size. As shown in Figure. 1, tightening contracts the variable domains and removes portions of the search space that cannot contain feasible solutions, thereby improve relaxation strength and reduce the need for branching. We use:

- **Feasibility-based bound tightening (FBBT):** inexpensive constraint propagation applied to the relaxation to remove infeasible portions of variable domains.
- **Optimization-based bound tightening (OBBT):** stronger tightening obtained by solving auxiliary minimization/maximization subproblems over the current convex relaxation to contract selected variable bounds (here, primarily V_k and θ_{km}).

FBBT is applied as preprocessing and after OBBT to propagate improvements through the constraint set. The resulting tighter bounds improve the quality of QC relaxations, reduce branching, and speed up convergence of the robust global algorithm.

NUMERICAL RESULTS

We evaluate the proposed robust ACOPF framework on benchmark networks and study three factors: (i) the effect of uncertainty set choice (nominal vs. box vs. polyhedral), (ii) the impact of bound tightening (RsBB vs. BT–RsBB), and (iii) the benefit of the improved trigonometric relaxation relative to a reference sine relaxation.

Test Cases and Computational Setting

Experiments are conducted on 17 test cases from PGLib-OPF [26] with fewer than 500 buses. Active and reactive demands are perturbed by $\xi_k \in [-0.01, 0.01]$ (1% level). We consider box $\mathcal{U}_{box} = \{\xi : |\xi_k| \le \Psi_k, \forall k\}$ and polyhedral set $\mathcal{U}_{\text{poly}} = \{\xi : |\xi_k| \le \Psi_k, \forall k, \sum_{k\in\mathcal{B}} |\xi_k| \le \Gamma\}$, and also report the deterministic nominal case. All models are implemented in Pyomo 6.9.2 (Python 3.10.18). Nonconvex NLP sub-problems are solved with IPOPT (tolerance 10^{-8}). The infeasibility test uses $\varepsilon_{\text{it}} = 10^{-4}$, and RsBB termination/fathoming uses a global relative tolerance $\varepsilon = 10^{-4}$. Bound tightening uses a hybrid FBBT + OBBT procedure with $\varepsilon_{\text{bt}} = 10^{-2}$ and at most 10 iterations; OBBT subproblems are run in parallel on 16 cores. Each RsBB

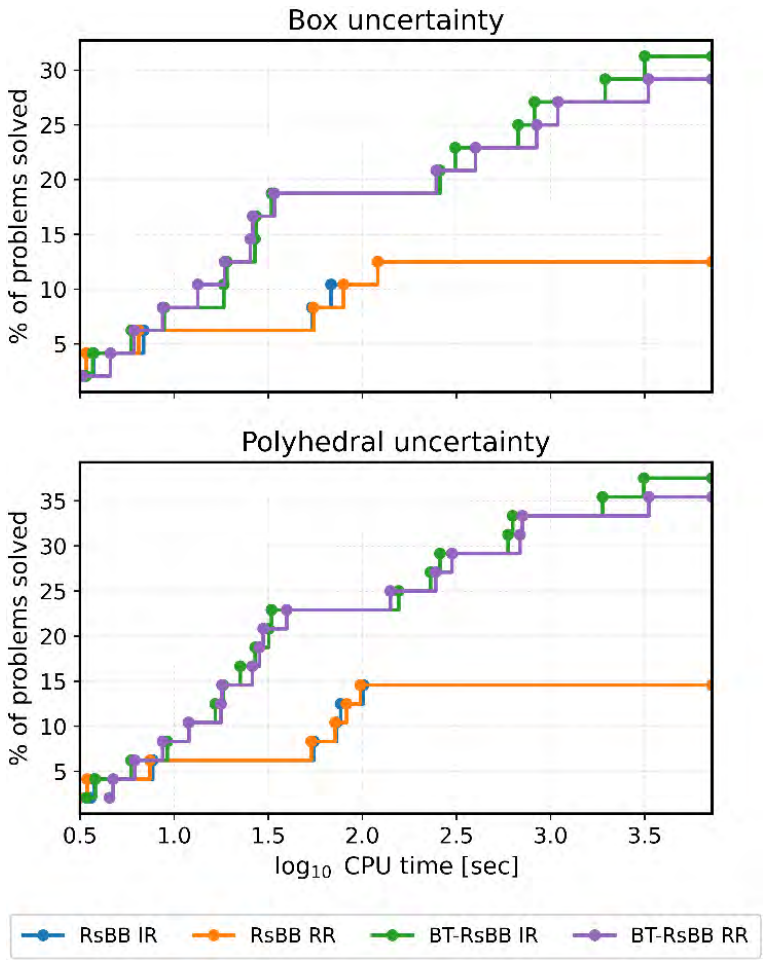

Figure 3. Percentage of instances certified within a given CPU time (log-scale) under a 2h time limit (7200 s). Curves compare RsBB vs. BT–RsBB with the improved relaxation (IR) and the reference relaxation (RR).

run has a 2 h time limit. Experiments are executed on a Windows workstation (Intel Core Ultra 7, 32GB).

Robust Operating Points and Computational Tractability

Table 1 reports robust AC OPF benchmark results under a 1% demand perturbation ($\xi = 0.01$) for both box and polyhedral uncertainty sets. For each test case, we compare RsBB against BT–RsBB (with FBBT+OBBT), and we also compare the improved relaxation (IR) with the reference relaxation (RR). Reported gaps are relative robust optimality gaps (in %), and times are CPU seconds (stopped at the 2h limit).

Overall, BT–RsBB consistently yields smaller gaps and/or more solved instances within the time limit than pure RsBB, confirming that bound tightening strengthens the relaxations and reduces the required branching effort. Across both uncertainty sets, IR tends to achieve smaller final gaps than RR for comparable runtimes, or reaches similar gaps faster, indicating that tighter trigonometric envelopes improve lower bounds and accelerate convergence. Finally, polyhedral uncertainty is generally easier than box uncertainty, since the budgeted structure is less conservative and typically produces weaker worst-case violations, leading to a smaller and faster global search.

Figure 2 illustrates the voltage magnitude profile for case240_pserc under: (a) nominal ACOPF, (b) polyhedral robustness, and (c) box robustness (BT–RsBB with the improved relaxation). Compared with the nominal dispatch, robust solutions reshape the voltage profile to preserve feasibility under demand deviations. The box robust solution is visibly more conservative than the polyhedral one, consistent with the fact that $\mathcal{U}_{\text{box}}$ allows simultaneous worst-case deviations at all buses, while $\mathcal{U}_{\text{poly}}$ limits the deviation via a budget Γ.

Figure 3 reports empirical solve-rate curves (fraction of instances certified globally optimal within a given CPU time, log-scale) for RsBB and BT–RsBB under both uncertainty sets, using either the improved relaxation (IR)

or the reference relaxation (RR).

Three trends are consistent across the benchmark suite:

- **Bound tightening is the main accelerator.** BT–RsBB solves substantially more instances than baseline RsBB within the 2 h limit. Tightened bounds shrink the search space and strengthen node relaxations, leading to fewer branch-and-bound nodes and faster gap closure.
- I**mproved trigonometric relaxation provides additional speedups.** IR yields systematically faster convergence than RR, reflecting tighter lower bounds and reduced relaxation error in the lifted sine/cosine constraints. The effect is moderate on the log-time plot but corresponds to noticeable reductions in wall-clock time on solved cases.
- **Polyhedral uncertainty is less conservative and easier computationally than box uncertainty.** Under Upoly, more instances are certified within the time limit. The budget constraint limits extremal scenarios, typically requiring fewer robust cuts and producing smaller search trees than the harsher Ubox setting. Overall, the numerical results indicate that combining RsBB with hybrid bound tightening is critical for scalability on robust ACOPF, while tighter trigonometric envelopes further improve performance. The choice of uncertainty set strongly affects both conservatism and computational difficulty, with box uncertainty being the most demanding.

CONCLUSION

In this paper, we developed an enhanced BT–RsBB framework to compute globally robust optimal solutions for nonconvex ACOPF under uncertainty. Building on the RsBB methodology, we improve practical efficiency via (i) parallel robust-cut generation in the infeasibility test, (ii) a hybrid bound-tightening module combining FBBT and OBBT, and (iii) a tighter trigonometric relaxation based on a precomputed mapping that strengthens node relaxations without substantially increasing model size.

We evaluate the method on 48 PGLib-OPF benchmark networks under three operating conditions (TYP, API, SAD) with both box and budgeted polyhedral uncertainty sets. The results show that BT–RsBB consistently outperforms baseline RsBB: bound tightening significantly contracts variable domains, strengthens relaxations, and accelerates robust gap closure, yielding more certified solutions within a two-hour limit. The improved sine relaxation further reduces runtimes by tightening lower bounds and decreasing the number of explored nodes.

The study also reveals strong sensitivity to operating conditions and uncertainty set geometry. Box uncertainty is systematically more conservative—and typically harder—than polyhedral uncertainty because it permits simultaneous worst-case deviations at all buses. Overall, our results provide a rigorous global benchmark for robust ACOPF with certified bounds and transparent convergence behaviour. Future work will target scalability through adaptive bound tightening, stronger branching and range-reduction strategies, and learning-based acceleration for cut selection, node selection, and relaxation refinement.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

ACKNOWLEDGEMENTS

Financial support from EPSRC grants EP/T022930/1, EP/W003317/1 and EP/V051008/1 and UCL's Research Excellence Scholarship is gratefully acknowledged.

AUTHOR IDENTIFIERS

Author ORCIDs:
Charitopoulos VM: 0000-0001-9051-917X
Yuhui Y: 0009-0000-5360-720X

REFERENCES

1. Carpentier J. Contribution à l'étude du dispatching économique. Bulletin de la Société Française des Électriciens 3(1):431–447 (1962).
2. Phan D, Ghosh S. Two-stage stochastic optimization for optimal power flow under renewable generation uncertainty. ACM Trans. Model. Comput. Simul. 24:1-22 (2014). https://doi.org/10.1145/2553084
3. Louca R, Bitar E. Robust AC optimal power flow. IEEE Trans. Power Syst. 34:1669-1681 (2019). https://doi.org/10.1109/tpwrs.2018.2849581
4. Lee D, Turitsyn K, Molzahn DK, Roald LA. Robust AC optimal power flow with robust convex restriction. IEEE Trans. Power Syst. 36:4953-4966 (2021). https://doi.org/10.1109/tpwrs.2021.3075925
5. Papavasiliou A, Oren SS, O'Neill RP. Reserve requirements for wind power integration: a scenario-based stochastic programming framework. IEEE Trans. Power Syst. 26:2197-2206 (2011). https://doi.org/10.1109/tpwrs.2011.2121095
6. Marley JF, Vrakopoulou M, Hiskens IA. An AC-QP optimal power flow algorithm considering wind

forecast uncertainty. 2016 IEEE Innovative Smart Grid Technologies - Asia (ISGT-Asia) :317-323 (2016). https://doi.org/10.1109/isgt-asia.2016.7796405
7. Venzke A, Halilbasic L, Markovic U, Hug G, Chatzivasileiadis S. Convex relaxations of chance constrained AC optimal power flow. IEEE Trans. Power Syst. 33:2829-2841 (2018). https://doi.org/10.1109/tpwrs.2017.2760699
8. Arab A, Tate JE. Distributionally robust optimal power flow via ellipsoidal approximation. IEEE Trans. Power Syst. 38:4826-4839 (2023). https://doi.org/10.1109/tpwrs.2022.3217941
9. Bai X, Qu L, Qiao W. Robust AC optimal power flow for power networks with wind power generation. IEEE Trans. Power Syst. 31:4163-4164 (2016). https://doi.org/10.1109/tpwrs.2015.2493778
10. Louca R, Bitar E. Robust AC optimal power flow. IEEE Trans. Power Syst. 34:1669-1681 (2019). https://doi.org/10.1109/tpwrs.2018.2849581
11. Yang H, Morton DP, Bandi C, Dvijotham K. Robust optimization for electricity generation. INFORMS Journal on Computing 33:336-351 (2021). https://doi.org/10.1287/ijoc.2020.0956
12. Mitsos A. Global optimization of semi-infinite programs via restriction of the right-hand side. Optimization 60:1291-1308 (2011). https://doi.org/10.1080/02331934.2010.527970
13. Isenberg NM, Akula P, Eslick JC, Bhattacharyya D, Miller DC, Gounaris CE. A generalized cutting?set approach for nonlinear robust optimization in process systems engineering. AIChE Journal 67: (2021). https://doi.org/10.1002/aic.17175
14. Caratzoulas S, Floudas CA. Trigonometric convex underestimator for the base functions in fourier space. J Optim Theory Appl 124:339-362 (2005). https://doi.org/10.1007/s10957-004-0940-2
15. Zhang Y, Sahinidis NV, Nohra C, Rong G. Optimality-based domain reduction for inequality-constrained NLP and MINLP problems. J Glob Optim 77:425-454 (2020). https://doi.org/10.1007/s10898-020-00886-z
16. Castelli, A. F., Harjunkoski, I., Poland, J., Giuntoli, M., Martelli, E., & Grossmann, I. E. (2024). Solving the security constrained unit commitment problem: Three novel approaches. *Int. J. Electr. Pow. Energy Sys.*, *162*, 110213.
17. Boukouvala F, Misener R, Floudas CA. Global optimization advances in mixed-integer nonlinear programming, MINLP, and constrained derivative-free optimization, CDFO. European Journal of Operational Research 252:701-727 (2016). https://doi.org/10.1016/j.ejor.2015.12.018
18. Marousi A, Charitopoulos VM. Global robust optimisation for non-convex quadratic programs: application to pooling problems. Systems and Control Transactions 4:1592-1597 (2025). https://doi.org/10.69997/sct.168949
19. Zimmerman RD, Murillo-Sánchez CE. MATPOWER user's manual. *Power Systems Engineering Research Center* (2016). [Online].
20. Available: https://matpower.org/docs/MATPOWER-manual.pdf https://doi.org/10.5281/zenodo.3236519
21. Coffrin C, Hijazi HL, Van Hentenryck P. Strengthening the SDP relaxation of AC power flows with convex envelopes, bound tightening, and valid inequalities. IEEE Trans. Power Syst. 32:3549-3558 (2017). https://doi.org/10.1109/tpwrs.2016.2634586
22. Taylor JA. Convex optimization of power systems. *Cambridge University Press* (2015).
23. Coffrin C, Hijazi HL, Van Hentenryck P. The QC relaxation: a theoretical and computational study on optimal power flow. IEEE Trans. Power Syst. 31:3008-3018 (2016). https://doi.org/10.1109/tpwrs.2015.2463111
24. Bynum M, Castillo A, Watson JP, Laird CD. Tightening mccormick relaxations toward global solution of the ACOPF problem. IEEE Trans. Power Syst. 34:814-817 (2019). https://doi.org/10.1109/tpwrs.2018.2877099
25. Chen C, Atamturk A, Oren SS. Bound tightening for the alternating current optimal power flow problem. IEEE Trans. Power Syst. 31:3729-3736 (2016). https://doi.org/10.1109/tpwrs.2015.2497160
26. Marousi A, Charitopoulos VM. Global and robust optimisation for non-convex quadratic programs. arXiv preprint (2025). https://arxiv.org/abs/2503.07310
27. Babaeinejadsarookolaee S, Birchfield A, Christie RD, Coffrin C, DeMarco C, Diao R, Ferris M, Fliscounakis S, Greene S, Huang R, et al. The power grid library for benchmarking AC optimal power flow algorithms. arXiv preprint(2019). https://arxiv.org/abs/1908.02788

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Section 4

Pharmaceutical and Biotechnological Systems

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Uncertainty-Aware Model Validation Framework for Pharmaceutical Process Development

Kensaku Matsunami[a], Yash Barhate[a], and Zoltan K. Nagy[a]*

[a] Purdue University, Davidson School of Chemical Engineering, West Lafayette, Indiana, United States
* Corresponding Author: znagy@purdue.edu.

ABSTRACT

Mathematical models are increasingly used in pharmaceutical process development within quality-by-design (QbD) frameworks to reduce experimental effort and enable rational process design. However, model validation is still often based on deterministic performance indicators, which do not explicitly account for experimental variability, measurement noise, and model uncertainty. This work proposes an uncertainty-aware framework for model validation in pharmaceutical processes that quantifies predictive reliability in probabilistic terms, consistent with regulatory concepts. The framework explicitly integrates uncertainty in operating conditions, measurements, and model parameters, and evaluates model performance based on the probability that prediction satisfy predefined acceptance criteria rather than on single-point accuracy indicators. An in-silico case study of crystallization was performed to demonstrate the approach, where synthetic experimental data with controlled uncertainty were generated. This enables systematic assessment of the effect of data quality, data quantity, and methods of uncertainty analysis on model validity results. Overall, the proposed framework provides a structured and robust approach for uncertainty-aware, regulatory-aligned model validation in pharmaceutical process development.

Keywords: Process Design, Design Under Uncertainty, Industry 4.0, Jacobian, Modelling and Simulations

INTRODCUTION

The increasing use of mechanistic, data-driven, and hybrid models in pharmaceutical process development reflects a broader shift toward model-based decision-making within quality-by-design (QbD) frameworks. Process models are increasingly used for process scale-up and design space exploration, contributing to significant reductions in experimental effort while enabling deeper process understanding [1]. From a regulatory perspective, however, model credibility and suitability must be systematically assessed in order to support risk-informed decision-making [2].

Since uncertainty plays a central role in process development, many process systems engineering studies have focused on uncertainty analysis and stochastic simulation. For example, Geremia et al. presented robust design spaces that capture the effects of process uncertainties on drug product quality [3], and Pessina et al. proposed a Bayesian framework for parametric uncertainty quantification in nonlinear protein crystallization models [4]. These studies primarily emphasize prediction of expected variability in product quality, rather than model credibility assessment.

Despite these advances, model validation in practice is still often based on deterministic goodness-of-fit metrics such as the coefficient of determination or the root mean square error. While these metrics quantify agreement between model predictions and experimental data, they do not explicitly account for experimental variability, measurement uncertainty, or uncertainty in model structure and parameters. Although Barhate et al. assessed the robustness of estimated parameters in crystallization by comparing parametric uncertainty with measurement uncertainty [5], model robustness and reliability remain difficult to quantify in a generalized manner. This gap limits the ability to perform risk-informed, context-dependent assessment of model credibility.

In this work, we present an uncertainty-aware model validation framework for pharmaceutical process development. An in-silico case study of cooling crystallization was used to demonstrate the approach, where

synthetic experimental datasets with controlled operational and measurement uncertainty were generated using a population balance model (PBM). By considering scenarios with different uncertainty levels and different amounts of experimental data, multiple validation criteria and model uncertainty quantification approaches were compared to assess the influence of data quality, data quantity, and uncertainty treatment on validation outcomes. The objective is to provide a structured and transparent methodology for uncertainty-aware, regulatory-aligned model validation in pharmaceutical process development.

PROPOSED METHODOLOGY

In-silico experimental data generation

Two sources of experimental variability were considered in the in-silico experiments. The first source is "operational variability" ($\boldsymbol{\varepsilon}_{\mathrm{ope}}$), which represents deviations between nominal setpoints and actual process conditions. The second source is "measurement variability" ($\boldsymbol{\varepsilon}_{\mathrm{meas}}$), which combines sampling effects and measurement errors. In this in-silico study, the function f_{GT} is treated as a ground-truth process model. The true response $\boldsymbol{y}_{\mathrm{true}}$ is defined as the noise-free outputs of f_{GT}, whereas the observed response $\boldsymbol{Y}_{\mathrm{meas}}$ denotes a synthetic observable generated from this model with imposed variability.

The true values and the observed values are defined as :

$$\boldsymbol{y}_{\mathrm{true}} = f_{\mathrm{GT}}(\boldsymbol{x}, \boldsymbol{\theta}) \tag{1}$$

$$\boldsymbol{Y}_{\mathrm{obs}} = f_{\mathrm{GT}}\left(\boldsymbol{x} + \boldsymbol{\varepsilon}_{\mathrm{ope}}\boldsymbol{x}, \boldsymbol{\theta}\right) + \boldsymbol{\varepsilon}_{\mathrm{meas}} \tag{2}$$

where $\boldsymbol{x}$ and $\boldsymbol{\theta}$ denote nominal operating conditions and model parameters, respectively.

Uncertainty-conscious predictions

After parameter estimation yielding $\hat{\boldsymbol{\theta}}$ of the predictive model f_{pred}, quality attributes can be predicted with uncertainty. Here, uncertainty is categorized into "parametric uncertainty" ($\boldsymbol{\varepsilon}_{\theta}$) and "experimental variability" ($\boldsymbol{\varepsilon}_{\mathrm{exp}}$). Model predictions are denoted by the parametric predicted values $\boldsymbol{Y}_{\mathrm{pred,par}}$ and the full predicted values $\boldsymbol{Y}_{\mathrm{pred,full}}$, corresponding to parametric uncertainty only and combined parametric and experimental uncertainty, respectively.

$$\boldsymbol{Y}_{\mathrm{pred,par}} = f_{\mathrm{pred}}\left(\boldsymbol{x}, \hat{\boldsymbol{\theta}} + \boldsymbol{\varepsilon}_{\theta}\right) \tag{3}$$

$$\boldsymbol{Y}_{\mathrm{pred,full}} = f_{\mathrm{pred}}\left(\boldsymbol{x}, \hat{\boldsymbol{\theta}} + \boldsymbol{\varepsilon}_{\theta}\right) + \boldsymbol{\varepsilon}_{\mathrm{exp}} \tag{4}$$

Parametric uncertainty quantification

As summarized in the literature [6, 7], several approaches exist to quantify parametric uncertainty, which differ in terms of their assumptions, computational cost, and the extent to which they explore the parameter space. These approaches can be categorized into derivative-based methods and sampling-based methods.

Derivate-based methods approximate uncertainty locally around the estimated parameter vector $\hat{\boldsymbol{\theta}}$. The Jacobian-based approach is the most common and relies on a first-order linearization of the residual function, which may lead to underestimation of uncertainty for strongly nonlinear models. The Hessian-based approach incorporates second-order curvature information of the objective function. For derivate-based methods, accurate parameter estimation is critical, as the covariance is computed locally around $\hat{\boldsymbol{\theta}}$.

In contrast, sampling-based approaches aim to characterize uncertainty more globally by repeatedly refitting the model to resampled datasets (e.g., bootstrap) or by sampling from a parameter distribution consistent with the data (e.g., Markov Chain Monte Carlo). These methods are generally more robust to nonlinearity and multimodality [4] but are computationally expensive. In this study, Jacobian- and Hessian-based covariance approximations were employed due to their computational efficiency within the proposed validation framework.

Regulatory-aligned model validation metrics

Threshold-based validation approach

US Food and Drug Administration (FDA) has proposed a "threshold-based" acceptance criterion for medical-device submissions to assess whether a model produces results that are sufficiently reliable from a safety standpoint [8]. This approach validates the closeness between simulation and experimental results relative to a safety threshold y_{th}. The comparison error E and the proximity Ω are computed for attribute i as:

$$E^{(i)} = \left|Y^{(i)}_{\mathrm{pred,par}} - Y^{(i)}_{\mathrm{obs}}\right| \tag{5}$$

$$\Omega^{(i)} = \left|y^{(i)}_{\mathrm{th}} - Y^{(i)}_{\mathrm{pred,par}}\right| \tag{6}$$

where a probability distribution of $Y^{(i)}_{\mathrm{obs}}$ is generated based on experimental data $y^{(i)}_{\mathrm{obs}}$ and $\varepsilon^{(i)}_{\mathrm{exp}}$. The model is considered validated if i) $\mathbb{E}[E] < \mathbb{E}[\Omega]$ and ii) the confidence intervals (CIs) of E and Ω do not overlap.

For general model validation, the safety threshold can be extended to acceptable prediction errors relative to experimental data. Using a predefined tolerance $\varepsilon^{(i)}_{\mathrm{th}}$, the threshold $y^{(i)}_{\mathrm{th}}$ is defined as $\left(1 \pm \varepsilon^{(i)}_{\mathrm{th}}\right) \cdot y^{(i)}_{\mathrm{obs}}$. Since this yields two-sided thresholds, the proximity is extended to Ω_{lower} and Ω_{upper}.

Model reliability assessment

Alternatively, model reliability under uncertainty is assessed as illustrated in Figure 1. A reliable model

Table 1: In-silico DoE used for parameter estimation and validation.

Purpose	Seed fraction [%]	Cooling rate [K s^{-1}]	Seed size [μm]	# of experiments
Parameter estimation	{0, 1.0}	{0.002, 0.008}	60 (if seeded)	4
Experimental variability	0.5	0.004	60	3
Interpolating validation	{0.25, 0.75}	{0.003, 0.006}	60	4
Extrapolating validation	{0.5, 1.5}	{0.001, 0.0015, 0.004, 0.010}	{50, 60, 70}	6

should exhibit a small deviation between $\boldsymbol{y}_{\mathrm{true}}$ and $\mathbb{E}[\boldsymbol{Y}_{\mathrm{pred,par}}]$ and a small statistical distance between the distributions of $\boldsymbol{Y}_{\mathrm{obs}}$ and $\boldsymbol{Y}_{\mathrm{pred,full}}$. Since neither $\boldsymbol{y}_{\mathrm{true}}$ nor distributions of $\boldsymbol{Y}_{\mathrm{obs}}$ is known in practice, reliability metrics are introduced to estimate model confidence.

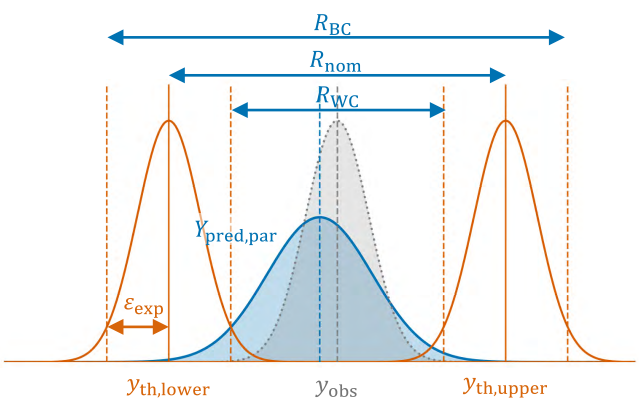

(a)

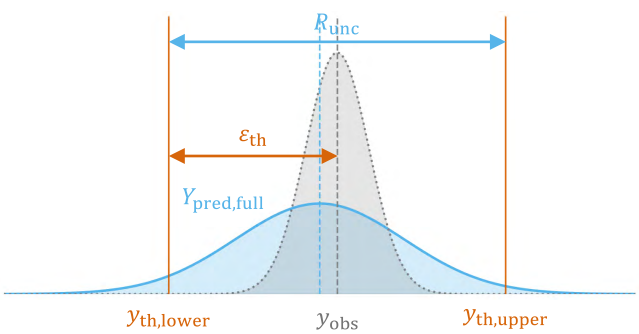

(b)

Figure 1. Concept of reliability scores based on (a) $\boldsymbol{Y}_{\mathrm{pred,par}}$ and (b) $\boldsymbol{Y}_{\mathrm{pred,full}}$.

First, predictions with parametric uncertainty $Y^{(i)}_{\mathrm{pred,par}}$ are compared against thresholds (Figure 1(a)). The nominal reliability score R_{nom} is defined as:

$$R^{(i)}_{\mathrm{nom}} = \mathbb{P}\left(y^{(i)}_{\mathrm{th,lower}} \le Y^{(i)}_{\mathrm{pred,par}} \le y^{(i)}_{\mathrm{th,upper}}\right) \tag{7}$$

To incorporate experimental variability, worst-case R_{WC} and best-case reliability scores R_{BC} are defined as:

$$R^{(i)}_{\mathrm{WC}} = \mathbb{P}\left(y^{(i)}_{\mathrm{th,lower}} + \varepsilon^{(i)}_{\mathrm{exp}} \le Y^{(i)}_{\mathrm{pred,par}} \le y^{(i)}_{\mathrm{th,upper}} - \varepsilon^{(i)}_{\mathrm{exp}}\right) \tag{8}$$

$$R^{(i)}_{\mathrm{BC}} = \mathbb{P}\left(y^{(i)}_{\mathrm{th,lower}} - \varepsilon^{(i)}_{\mathrm{exp}} \le Y^{(i)}_{\mathrm{pred,par}} \le y^{(i)}_{\mathrm{th,upper}} + \varepsilon^{(i)}_{\mathrm{exp}}\right) \tag{9}$$

While R_{WC} cannot be high when experimental variability is large, R_{BC} should be sufficiently high; otherwise $\boldsymbol{Y}_{\mathrm{pred,par}}$ deviates substantially from $\boldsymbol{y}_{\mathrm{true}}$ or parametric uncertainty is excessive.

An alternative reliability score R_{unc} is described based on $\boldsymbol{Y}_{\mathrm{pred,full}}$ (Figure 1(b)) as:

$$R^{(i)}_{\mathrm{unc}} = \mathbb{P}\left(y^{(i)}_{\mathrm{th,lower}} \le Y^{(i)}_{\mathrm{pred,full}} \le y^{(i)}_{\mathrm{th,upper}}\right) \tag{10}$$

CASE STUDY DESCRIPTION

To demonstrate the proposed validation framework, a case study was conducted using cooling crystallization of paracetamol. Both the experimental conditions and the true values of model parameters were defined based on literature data [9].

Model description

A one-dimensional population balance model (PBM) was used for synthetic experimental data generation, parameter estimation, and model validation. In this case study, the same PBM formulation was employed as both f_{GT} and f_{pred}. The number density of crystals of size L, denoted by $n(L)$, was described by considering crystal growth, primary nucleation, and secondary nucleation:

$$\frac{\partial n(L)}{\partial t} + \frac{\partial [Gn(L)]}{\partial t} = [B_{\mathrm{p}} + B_{\mathrm{s}}] \bullet \delta(L - L_{\mathrm{n}}) \tag{12}$$

where δ is a Dirac delta function and L_{n} is the nuclei size.

The growth rate G, the primary nucleation rate B_{p}, and the secondary nucleation rate B_{s} are defined as:

$$G = k_{\mathrm{g}}\sigma^{g}\exp\left(-\frac{E_{\mathrm{A}}}{RT_{\mathrm{K}}}\right) \tag{13}$$

$$B_{\mathrm{p}} = k_{\mathrm{p}}\sigma^{p} \tag{14}$$

$$B_{\mathrm{s}} = k_{\mathrm{s}}\sigma^{s}A_{\mathrm{c}} \tag{15}$$

where σ is the relative supersaturation, R is the gas constant, T_{K} is the temperature, A_{c} is the specific surface area of crystals. The seven model parameters

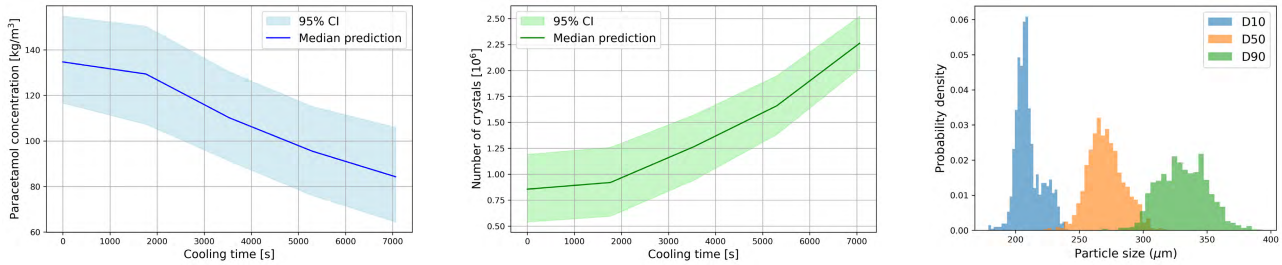

(a) (b) (c)

Figure 2. In-silico experimental data wih uncertainty at the centerpoint. (a) paracetamol concentration profile with 95% CIs, (b) crystal number profile with CIs, (c) probability densities of D_{10}, D_{50}, and D_{90} of the final products.

k_g, g, E_A, k_p, p, k_s, and s were estimated. All computational work was performed using PharmaPy, an open-source tool for the digital pharmaceutical process design [10].

In-silico design of experiments (DoE)

Table 1 summarizes the in-silico DoE used to generate synthetic experimental data, using the ground truth model. Two-level full factorial experiments in seed fraction and cooling rate were performed for parameter estimation, with three additional repeated runs at the center point to assess experimental variability. Two sets of validation runs were designed: four interpolating conditions and six extrapolating conditions. In the extrapolating conditions, seed size was also varied as an additional factor. Variability terms $\varepsilon_{\text{ope}}^{(\text{fra})}, \varepsilon_{\text{ope}}^{(\text{cl})}$, and $\varepsilon_{\text{ope}}^{(\text{size})}$ were assigned to seed fraction, cooling rate, and seed size, respectively.

During the in-silico experiments, both the paracetamol concentration and the total number of crystals were recorded at five normalized time points corresponding to 0, 0.25, 0.5, 0.75, and 1 times the total cooling duration, with measurement variabilities $\varepsilon_{\text{meas}}^{(\text{conc})}$ and $\varepsilon_{\text{meas}}^{(\text{num})}$. After completion of each experiment, product size quantiles D_{10}, D_{50}, and D_{90} were also generated, with associated measurement variabilities $\varepsilon_{\text{meas}}^{(\text{D10})}, \varepsilon_{\text{meas}}^{(\text{D50})}$, and $\varepsilon_{\text{meas}}^{(\text{D90})}$.

CASE STUDY RESULTS

The generated measured values $\boldsymbol{Y}_{\text{obs}}$ at the center point of the operating conditions are shown in Figure 2. Normal distributions were assumed for all variability terms. The standard deviations of operating variabilities were set to 5% for $\varepsilon_{\text{ope}}^{(\text{fra})}, \varepsilon_{\text{ope}}^{(\text{cl})}$, and $\varepsilon_{\text{ope}}^{(\text{size})}$, while the standard deviations of measurements variabilities were defined as $\sigma\left(\varepsilon_{\text{meas}}^{(\text{conc})}, \varepsilon_{\text{meas}}^{(\text{num})}, \varepsilon_{\text{meas}}^{(\text{D10})}, \varepsilon_{\text{meas}}^{(\text{D50})}, \varepsilon_{\text{meas}}^{(\text{D90})}\right)$ = (10 kg m^{-3}, 10^5, 3 μm, 5 μm, 8 μm).

Parametric uncertainty results

Parameter estimation was performed by minimizing a weighted least-squares objective function constructed from all available in-silico observations. To account for differences in magnitude and measurement reliability among these quantities, the weighting factors were defined based on the inverse of the experimental variability observed from the three repeated center-point runs.

Table 2: Estimated model parameters and 95% CIs using Jacobian and Hessian matrices.

Parameter	Estimated value	CI (Jacobian)	CI (Hessian)
k_g [10^4m s^{-1}]	2.92	[2.91, 2.92]	[2.71, 3.13]
g [-]	1.24	[1.24, 1.24]	[0.98, 1.50]
E_A [kJ mol^{-1}]	63.2	[63.1, 63.4]	[62.6, 63.8]
k_p [10^5m^{-3} s^{-1}]	4.84	[4.83, 4.85]	[3.56, 6.12]
p [-]	2.61	[2.60, 2.61]	[2.25, 2.97]
k_s [10^4m^{-2} s^{-1}]	10.2	[10.2, 10.2]	[10.0, 10.4]
s [-]	1.86	[1.86, 1.86]	[1.61, 2.11]

After parameter estimation, covariance matrices were computed using both the Jacobian and Hessian approximations. The resulting 95% CIs are summarized in Table 2. The Jacobian-based approach produced very narrow CIs, all within 0.5% of the estimated values. In contrast, the Hessian-based approach yielded substantially wider CIs, ranging from 0.9 and 26.5% of the estimated values, reflecting the nonlinear sensitivity of the model parameters.

In-silico validation experiments

At the validation conditions summarized in Table 1, both $\boldsymbol{y}_{\text{true}}$ and $\boldsymbol{Y}_{\text{obs}}$ were generated. Among the quality attributes, D_{90} was selected as the most critical attribute, and a relative tolerance of $\varepsilon_{\text{th}}^{(\text{D90})}$=15% was applied.

Figure 3(a) shows the probability density functions of $Y_{\text{pred,par}}^{(\text{D90})}$ for one representative validation condition. The expected value $\mathbb{E}\left(Y_{\text{pred,par}}^{(\text{D90})}\right)$ was very close to $y_{\text{true}}^{(\text{D90})}$, with a relative error of 0.8%. However, the Jacobian-based $Y_{\text{pred,par}}^{(\text{D90})}$ did not cover $y_{\text{true}}^{(\text{D90})}$, whereas $y_{\text{true}}^{(\text{D90})}$ fell within the support of the Hessian-based $Y_{\text{pred,par}}^{(\text{D90})}$. Similar

trends were observed in eight out of ten validation runs, indicating that the Jacobian approximation systematically underestimates parametric uncertainty.

The probability density functions of $Y_{\text{pred,full}}^{(\text{D90})}$, shown in Figure 3(b), were very similar for the Jacobian- and the Hessian-based approaches. This occurs because experimental variability dominates over parametric uncertainty for this attribute. Compared to $Y_{\text{pred,full}}^{(\text{D90})}$, the distribution of $Y_{\text{obs}}^{(\text{D90})}$ exhibited a broader distribution with weak bimodality. This discrepancy is attributed to underestimation of experimental variability and to the assumption of normally distributed noise.

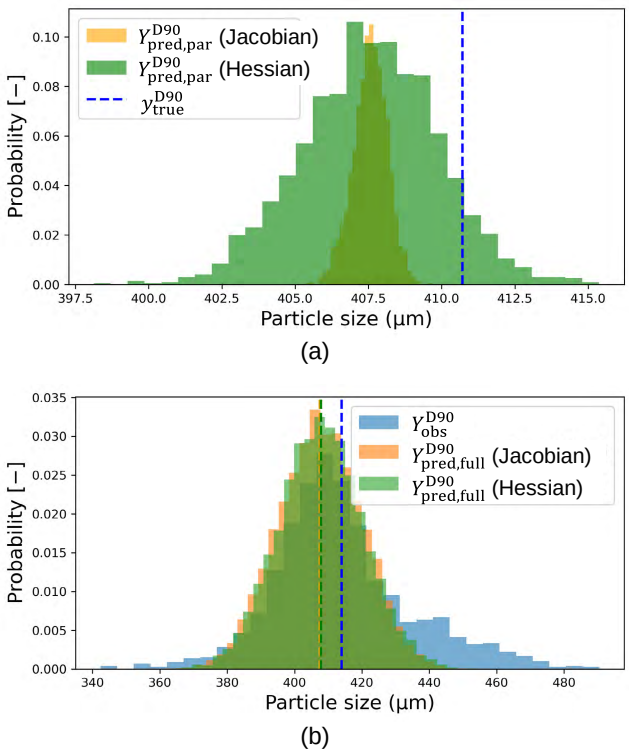

Figure 3. Probability density distributions of D_{90} at a validation run: (a) y_{true} and $\boldsymbol{Y}_{\text{pred,par}}$; (b) $\boldsymbol{Y}_{\text{obs}}$ and $\boldsymbol{Y}_{\text{pred,full}}$.

Validation metrics

For validation, it was assumed that only a single experimental observation $y_{\text{obs}}^{(\text{D90})}$, sampled from $Y_{\text{obs}}^{(\text{D90})}$, is available per validation condition, reflecting typical experimental practice. This observation was compared against the Hessian-based $Y_{\text{pred,par}}^{(\text{D90})}$ using the proposed validation metrics.

To assess the effect of experimental replication, three independent in-silico observations were generated per validation condition and their mean was used for validation, with the effective experimental variability ε_{exp} estimated separately for each condition rather than from center-point repeats.

Using the threshold-based approach, six out of ten validation conditions satisfied the validation criteria. Most failures were associated with large uncertainty in the comparison error $E^{(\text{D90})}$, which was primarily driven by experimental variability. When the mean of three repeated measurements was used for each condition, the number of validated conditions decreased to five, indicating that the validation outcome changed only marginally with replication in this dataset.

Table 3 summarizes the results of different validation metrics, including conventional deterministic metrics: the coefficient of determination (R^2), root mean square error ($RMSE$), and mean absolute percentage error ($MAPE$). Both deterministic and uncertainty-aware metrics changed with the number of samples; however, the interpretation may differ. The deterministic metrics are directly affected by experimental variability, meaning that validation conclusions based on unrepeated experiments can be strongly influenced by random experimental noise, making rational decision-making challenging.

Table 3: Summary of different validation metrics using a single sample and the mean of three repeated runs per condition.

Metrics	Single sample	Three samples
R^2 [-]	0.9066	0.9616
$RMSE$ [μm]	19.75	13.35
$MAPE$ [%]	5.65	3.98
$\mathbb{E}\left(R_{\text{nom}}^{(\text{D90})}\right)$ [%]	91.54	100
$\mathbb{E}\left(R_{\text{WC}}^{(\text{D90})}\right)$ [%]	61.04	81.14
$\mathbb{E}\left(R_{\text{BC}}^{(\text{D90})}\right)$ [%]	100	100
$\mathbb{E}\left(R_{\text{unc}}^{(\text{D90})}\right)$ [%]	91.29	98.38

In contrast, the proposed validation metrics could provide additional diagnostic insight into both model validity and data quality. The best-case reliability $R_{\text{BC}}^{(\text{D90})}$ remained at 100% across all cases, which is important because it avoids rejecting model validity solely due to large experimental variability. The worst-case reliability $R_{\text{WC}}^{(\text{D90})}$ was relatively low when only a single observation per condition was available. This indicates that there remains a non-negligible risk that the model could be incorrect, because a single noisy observation cannot reliably distinguish model mismatch from experimental scatter. When the mean of three repeated measurements was used, $R_{\text{WC}}^{(\text{D90})}$ increased, reducing the risk of an incorrect model conclusion. In this case, $R_{\text{WC}}^{(\text{D90})}$ acts as a conservative indicator of the remaining model risk under uncertainty.

Overall, these results highlight how uncertainty-aware validation metrics can improve interpretability of

model credibility and provide a structured way to assess the impact of experimental variability and replication on validation outcomes.

CONCLUSIONS

This work presented an uncertainty-aware framework for model validation in pharmaceutical process development from a regulatory perspective. By distinguishing between parametric uncertainty and experimental variability, the framework enables probabilistic rather than purely deterministic validation. The proposed reliability-based metrics extend traditional goodness-of-fit measures by quantifying the likelihood that predictions satisfy predefined acceptance criteria under uncertainty. An in-silico crystallization case study demonstrated how uncertainty quantification methods, experimental design, and replication influence validation outcomes. The results highlight limitations of local linear uncertainty approximations and the need to account for experimental variability when interpreting validation results.

Future work will focus on broader application of the framework across pharmaceutical unit operations, such as reactions and wet milling, as well as its application to real experimental datasets. Further development will also address systematic treatment of structural model uncertainty and improved characterization of experimental variability, including non-Gaussian noise and correlated disturbances. Ultimately, extending the framework toward integrated consideration of model, parameter, and experimental uncertainties can support transparent, risk-informed, and regulatorily aligned use of process models throughout pharmaceutical development and scale-up.

ACKNOWLEDGEMENTS

Funding from Takeda Pharmaceuticals International Co. is gratefully acknowledged.

AUTHOR IDENTIFIERS

Author ORCIDs:
Matsunami K: 0000-0002-8831-018X
Barhate Y: 0009-0007-6848-2244
Nagy ZK: 0000-0003-4787-6678

REFERENCES

1. Matsunami K, Martin Salvador P, Naranjo Gómez LN, Comoli GS, Charmchi I, Kumar A. Mechanistic modelling in pharmaceutical product and process development: a review of distributed and discrete approaches. Chemical Engineering Research and Design 218:8-24 (2025). https://doi.org/10.1016/j.cherd.2025.04.005
2. O'Connor TF, Chatterjee S, Lam J, de la Ossa DHP, Martinez-Peyrat L, Hoefnagel MHN, Fisher AC. An examination of process models and model risk frameworks for pharmaceutical manufacturing. International Journal of Pharmaceutics: X 8:100274 (2024). https://doi.org/10.1016/j.ijpx.2024.100274
3. Geremia M, Bezzo F, Ierapetritou MG. Design space determination of pharmaceutical processes: effects of control strategies and uncertainty. European Journal of Pharmaceutics and Biopharmaceutics 194:159-169 (2024). https://doi.org/10.1016/j.ejpb.2023.12.008
4. Pessina D, Calderon De Anda J, Heffernan C, Heng JYY, Papathanasiou MM. Integrated in vitro/in silico uncertainty quantification method for protein crystallization models. Ind. Eng. Chem. Res. 64:12025-12035 (2025). https://doi.org/10.1021/acs.iecr.4c04517
5. Barhate Y, Kilari H, Wu WL, Nagy ZK. Population balance model enabled digital design and uncertainty analysis framework for continuous crystallization of pharmaceuticals using an automated platform with full recycle and minimal material use. Chemical Engineering Science 287:119688 (2024). https://doi.org/10.1016/j.ces.2023.119688
6. Vanlier J, Tiemann CA, Hilbers PAJ, van Riel NAW. Parameter uncertainty in biochemical models described by ordinary differential equations. Mathematical Biosciences 246:305-314 (2013). https://doi.org/10.1016/j.mbs.2013.03.006
7. van Loosdrecht MCM, Nielsen PH, Lopez-Vazquez CM, Brdjanovic D. Experimental methods in wastewater treatment. Water Intelligence Online 15:9781780404752-9781780404752 (2016). https://doi.org/10.2166/9781780404752
8. Hariharan P, D'Souza GA, Horner M, Morrison TM, Malinauskas RA, Myers MR. Use of the FDA nozzle model to illustrate validation techniques in computational fluid dynamics (CFD) simulations. PLoS ONE 12:e0178749 (2017). https://doi.org/10.1371/journal.pone.0178749
9. Szilagyi B, Eren A, Quon JL, Papageorgiou CD, Nagy ZK. Application of model-free and model-based quality-by-control (qbc) for the efficient design of pharmaceutical crystallization processes. Crystal Growth & Design 20:3979-3996 (2020). https://doi.org/10.1021/acs.cgd.0c00295
10. Casas-Orozco D, Laky D, Wang V, Abdi M, Feng X, Wood E, Laird C, Reklaitis GV, Nagy ZK. Pharmapy: an object-oriented tool for the development of hybrid pharmaceutical flowsheets. Computers & Chemical Engineering 153:107408 (2021). https://doi.org/10.1016/j.compchemeng.2021.107408

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Comparison of Centralised and Decentralised Pharmaceutical Manufacturing Paradigms: An Agent-Based Simulation Study

Farshid Babaei[a], Mohammad Salehian[b], David Robins[a], Cameron J. Brown[b], Daniel Markl[b], Alastair J. Florence[b], Solomon Brown[a]*

[a] School of Chemical, Materials and Biological Engineering, University of Sheffield, Sheffield, UK
[b] CMAC, Strathclyde Institute of Pharmacy and Biomedical Sciences, University of Strathclyde, Glasgow, UK
* Corresponding Author: s.f.brown@sheffield.ac.uk.

ABSTRACT

Traditional centralised manufacturing offers efficient economies and broad market reach but faces increasing limitations with the rise of complex products requiring rapid localised delivery and greater supply chain resilience. The logistics demands of hospital-compounded therapies expose vulnerabilities in existing infrastructure, accentuating the need for rigorous evaluation of alternative paradigms. This study investigates the comparative performance of centralised and decentralised pharmaceutical manufacturing models, applying an agent-based simulation framework designed for specialised or time-sensitive drug product orders. The work implements an agent-based simulation to model both centralised and decentralised scenarios using key structural, resource, and demand parameters identified within the supply chain ecosystem. Comparison criteria include labour requirements, sustainability (as measured by environmental emissions and operational efficiency), and end-to-end supply chain lead times, informed by the geospatial distribution of manufacturers, hospitals, and clinics. The centralised case considers a single facility supplying major hospitals and several clinics from a designated hub, while the decentralised case models multiple smaller production sites supplying care centres more directly. Demand frequency, emergency inventory buffers, personnel allocations, and practical constraints are explicitly built into the simulation inputs. Preliminary simulation results reveal trade-offs between manufacturing paradigms across multiple performance dimensions. The decentralised model shows potential advantages in reducing supply chain lead times and improving responsiveness to localised demand surges, while the centralised model demonstrates efficiency gains in resource utilisation under steady-state conditions. The framework enables quantitative comparison of throughput, cost implications, delivery timeliness, and system resilience under varied operational scenarios. These findings will inform strategic design decisions for patient-centric and resilient pharmaceutical supply chains, facilitating adoption of flexible models capable of meeting modern healthcare delivery needs within the medicines manufacturing and supply ecosystem.

Keywords: Supply Chain, Modelling and Simulation, Intelligent Systems, Pharmaceutical Manufacturing

INTRODUCTION

Pharmaceutical supply chains are increasingly exposed to operational fragility, as demonstrated by recent global disruptions that have propagated rapidly through tightly coupled and inventory-lean networks. These failures highlight a structural dependence on centralised manufacturing architectures, where long replenishment paths and limited buffering capacity amplify the risk of shortages. [1-6].

Although decentralised and point-of-care manufacturing models are frequently proposed as a means of improving responsiveness and reducing logistical delays, existing literature provides limited quantitative

comparison of these paradigms under realistic demand and operational conditions. In particular, there is no systematic, simulation-based evaluation of how centralised and decentralised architectures differ in inventory behaviour, coverage performance, responsiveness, and environmental impact across multiple supply-chain echelons [7-9].

An increasingly prominent part of this discussion concerns the shift from batch based centralised production to decentralised or point of care (POC) manufacturing models, enabled by advanced processing technologies. Decentralised pharmaceutical manufacturing has attracted growing attention due to its potential to reduce transport delays, shorten replenishment lead times, and enable agile response to localised demand surges. Analyses across cell and gene therapy, monoclonal antibody production, and specialised medicines have highlighted those decentralised architectures, when underpinned by modern technologies, may achieve superior responsiveness and greater operational flexibility relative to traditional centralised paradigms [6, 10, 11].

This study addresses this gap by developing a multi-tier agent-based simulation of a representative pharmaceutical supply chain in Scotland. Two contrasting manufacturing configurations, a single large centralised plant and a set of smaller decentralised plants, are evaluated under identical downstream network structures. The model enables direct comparison of echelon-specific performance across inventory, service coverage, responsiveness, and emissions. The results provide an evidence-based assessment of trade-offs between efficiency, resilience, and sustainability across alternative manufacturing paradigms [6, 7, 9, 11].

METHODOLOGY

Pharmaceutical manufacturing paradigm Assessment

The pharmaceutical supply chain examined in this study (Figure 1) consists of four primary tiers: manufacturers, hospitals, clinics, and end consumers. The flow of products and information follows a demand-driven structure, where downstream demand signals propagate upstream to guide production and replenishment decisions. This pharmaceutical supply chain problem characterized by its structural configuration, connectivity among actors, and the policies and rules of engagement governing material and information flows, can be formulated as below:

Given

- Structure and connectivity between nodes, determining flow paths and network robustness.
- Policies and rules of engagement governing ordering, replenishment, and allocation.
- Emission rate associated with manufacturing activities.
- Demand variation arising from end consumers.

Evaluate

- Holding inventory as a cost indicator.
- Inventory coverage ratio showcasing the responsiveness of the system to meet demand reliably
- Emission levels as a measure of environmental impact generated by manufacturing segment

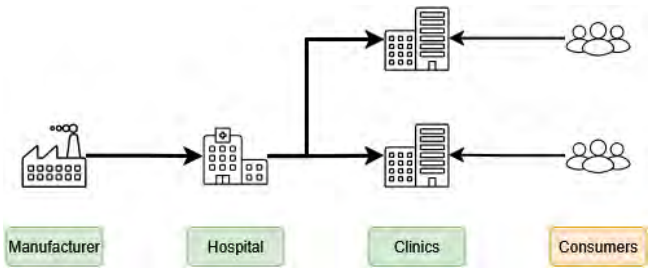

Figure 1. Conceptual model of the pharmaceutical supply chain.

Multi-Scale Agent-Based Model

The model, illustrated in Figure 2, is built on a foundational agent structure that is shared across all supply-chain entities (Figure 2-Top left). The application-level system built in the AnyLogic software (ver. 8.7.9), involves four primary agent types: manufacturer, distributor, health centre, and consumer, each instantiated from the same base agent template. As shown in Figure 2 (Top right), every agent in the logistic chain contains three core state variables: backlog, inventory, and a demand forecast, following system dynamic formulations reported in prior studies [12, 13]. Downstream demand is estimated using a first-order exponential smoothing process, and this forecast directly determines the agent's next order quantity. A set of behavioural parameters governs how each agent manages its inventory and backlog dynamics, shaping responsiveness and replenishment behaviour.

Interaction among agents occurs through order placement and shipment receipt, implemented via message-passing mechanisms that accommodate diverse ordering frequencies across tiers. Each agent type follows its own inventory policy, reflecting its position in the supply chain and influencing ordering decisions. To capture the logistics layer, the model incorporates explicit shipping and fleet components that represent transportation processes, queueing, and shipment delays, enabling a realistic depiction of material flows throughout the network.

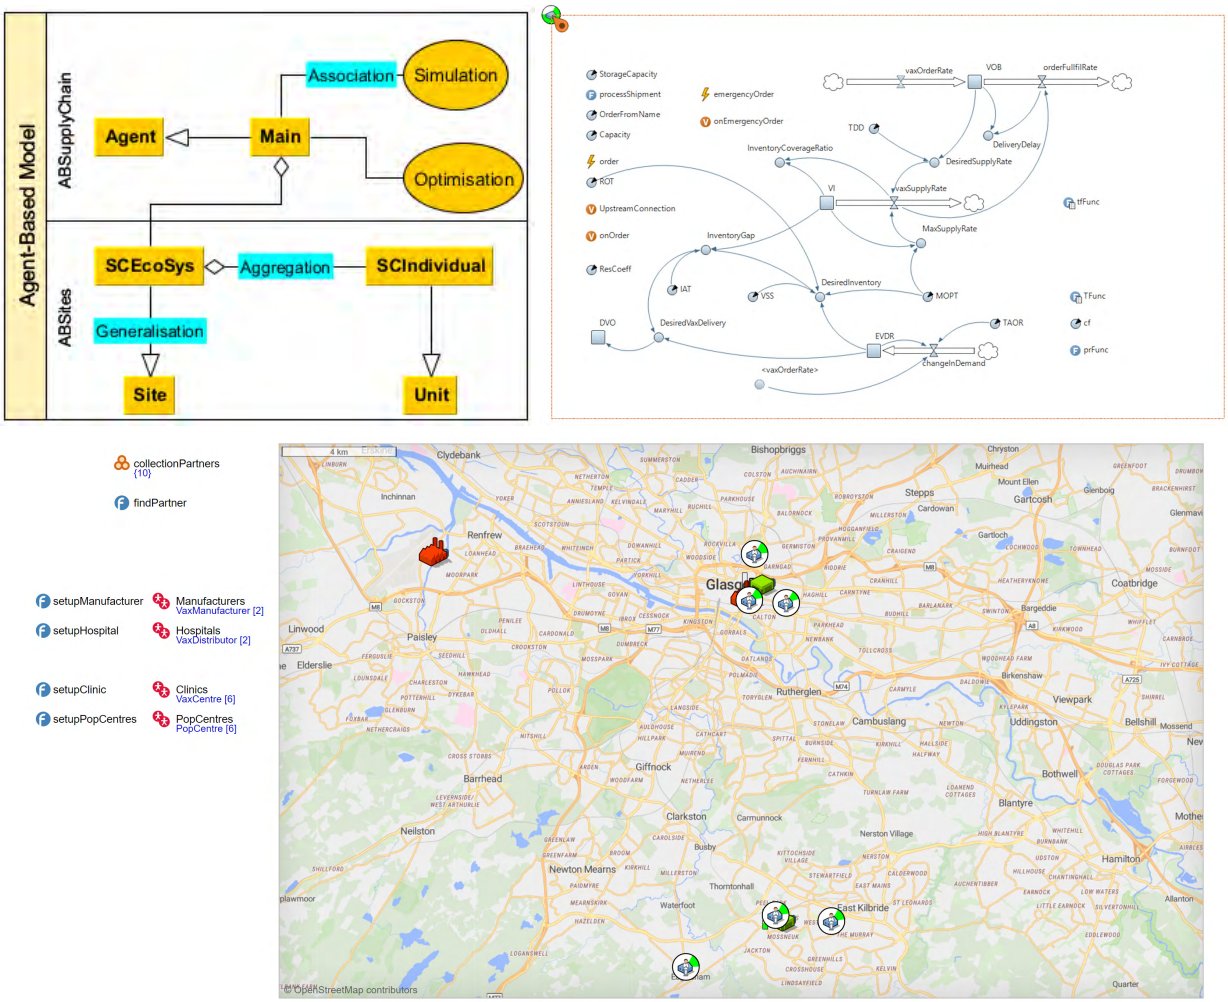

Figure 2. (Top left) The foundational representation of agent-based model (Top right) Dynamic agent behaviour representation (Bottom) Simulation model for Greater Glasgow pharmaceutical supply chain.

To represent the geographical distribution of supply chain entities, the model incorporates Geographic Information System (GIS) capabilities available within AnyLogic. These tools allow the spatial layout of manufacturers, distributors, and health centres to be embedded directly into the simulation environment, ensuring that distances, travel paths, and regional clustering are reflected in the system's behaviour. By linking agent locations to real or stylised geographic coordinates, the model captures spatial dependencies that influence shipment times, fleet utilisation, and overall logistics performance. This spatial layer complements the behavioural and structural components of the base model, enabling a more realistic depiction of how geography shapes SC dynamics.

The model uses a compact set of input data describing demand behaviour, operational policies, manufacturing capabilities, logistics characteristics, and sustainability factors. Demand inputs define the average and variable requirements of hospitals and clinics; operational inputs specify ordering rules, replenishment timings, and transport delays; manufacturing inputs capture production rates and process characteristics for both centralised and decentralised plants; and sustainability inputs include the emission factors associated with each manufacturing technology. Together, these inputs provide a consistent basis for comparing the two manufacturing paradigms under equivalent conditions.

Case study

The case study focuses on a pharmaceutical supply chain located in Scotland, modelled using two alternative manufacturing paradigms: a centralised configuration with a single large production plant, and a decentralised configuration with two smaller plants. In both scenarios, the downstream structure remains identical, consisting of two hospitals, six health centres, and six consumer centres. The overall supply-chain layout follows the structure

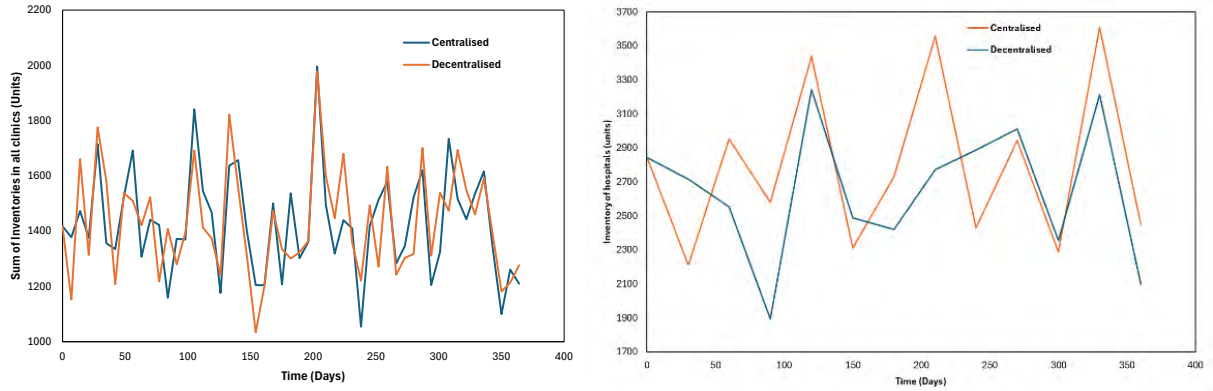

Figure 3: (left) Sum of inventory units in all clinics, and (right) Inventory of hospitals (Glasgow Royal Infirmary and Haimyres) in centralised and decentralised cases.

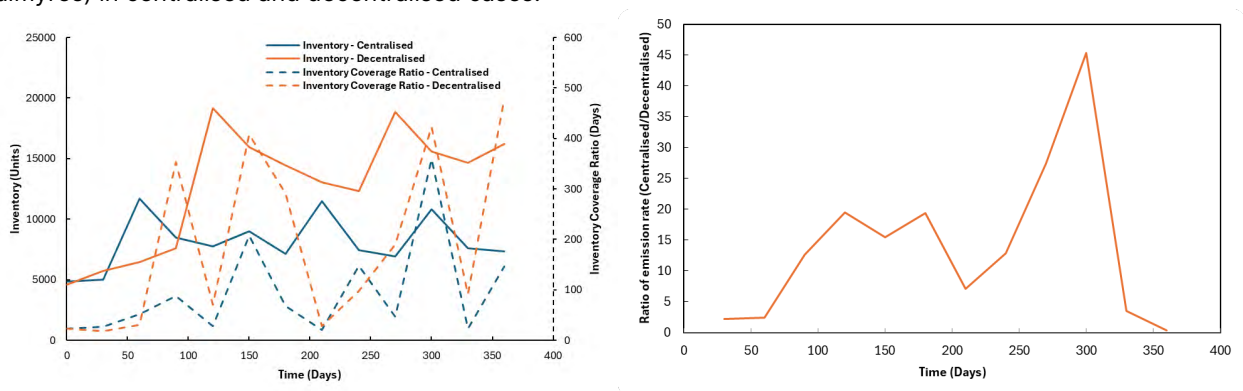

Figure 4: (left) Inventory units and inventory coverage ratio in centralised and decentralised cases, and (right) Ratio of emission rate of centralised over decentralised manufacturing cases.

reported in du Plessis, van Vuuren [14]. Under the centralised paradigm, hospitals place orders directly with the large plant, whereas in the decentralised paradigm each hospital orders from its corresponding small plant. Operational and logistical parameters for all tiers, as well as demand-variation characteristics, are taken from du Plessis, van Vuuren [14], ensuring behavioural realism. Emission parameters for manufacturing activities are also drawn from Hindiyeh, Altalafha [15] to capture sustainability impacts accurately. All agent locations in the model reflect their actual geographic coordinates, enabling spatially realistic routing, travel distances, and logistics performance within the Scottish context.

RESULTS AND DISCUSSION

The simulation outcomes reveal a structured set of performance differences between the centralised and decentralised pharmaceutical manufacturing paradigms. Performance differences emerge because clinics, hospitals, and manufacturers operate under distinct replenishment cycles, buffering rules, and demand-propagation dynamics. These structural differences cause each echelon to respond uniquely to changes in manufacturing configuration, leading to divergent inventory and coverage behaviours across the network. The analysis therefore demands a tiered interpretation rather than a monolithic comparison.

CLINIC LEVEL PERFORMANCE

Figure 3-Left shows the sum of all inventories at the clinic level in the centralised and decentralised cases. At the outermost tier (the six clinics) the two manufacturing architectures display no substantive divergence in overall inventory behaviour or service attainment. Although the decentralised system occasionally offers marginally tighter control of coverage during transient spikes, the broader simulation horizon demonstrates that clinics ultimately experience comparable stock levels and service stability under both paradigms. This is consistent with the observation that clinic level performance is heavily buffered by upstream nodes in the centralised case, while in the decentralised case it is moderated by shorter replenishment paths rather than by increased local inventory. Although decentralised manufacturing shortens physical

delivery distances, clinic-level performance remains comparable across both scenarios because clinics are heavily buffered by upstream inventory. In the centralised case, hospitals carry larger safety stocks to compensate for longer replenishment paths, while in the decentralised case shorter delivery routes offset lower upstream inventory. These opposing mechanisms result in similar clinic-level stability despite architectural differences.

Hospital Level Inventory

A more differentiated behaviour emerges at the hospital tier, where the simulation tracks the aggregated inventory of Glasgow Royal Infirmary and Hairmyres Hospital. The results show that the centralised scenario consistently yields slightly higher combined hospital inventories (Figure 3 - Right). While the magnitude of this increase is modest, its operational implications are not. Hospitals, as high cost stockholding environments, are disproportionately affected by incremental increases in buffer requirements. The elevated stock levels observed under the centralised architecture arise from the longer replenishment lead times from a single remote hub, compelling hospitals to hold additional safety stock. In contrast, under the decentralised configuration, more responsive resupply from CMAC and NMIS enables hospitals to operate with leaner inventory profiles, thereby reducing their holding cost burden.

Manufacturing Level Inventory

Figure 4-Left shows the inventory and coverage ratio over time for the manufacturers in the centralised (MMIC) and decentralised (CMAC + NMIS) cases. It can be seen that for the majority of the period considered, the decentralised case sees a higher total inventory (and therefore coverage ratio) maintained. This reflects the potential for greater fluctuations in order size in the decentralised case as a result of interactions between the larger number of participants. In the centralised case, the manufacturer is required to meet the demand of each hospital and so maintains an inventory sufficient to fulfil this requirement. In the decentralised case, each manufacturer could receive orders from either hospital, and does not know the orders placed to the order manufacturer. As a result, larger inventories are maintained as a means of ensuring resilience against downstream demand shocks. This effect is representative of a fundamental supply chain principle: the loss of risk pooling when production is divided across multiple smaller facilities. Each decentralised site must individually buffer against local variability, leading to a structurally heavier stock profile. The centralised facility, by contrast, benefits from statistically aggregated demand and can therefore operate more efficiently with lower safety stock requirements. Accordingly, while decentralisation may offer logistical and resilience benefits (as discussed below), it does so at the cost of increased working capital and reduced inventory efficiency at the manufacturing tier.

Coverage Ratios and Responsiveness

Despite its higher manufacturing inventory, the decentralised system exhibits stronger inventory coverage ratios, particularly during demand intensification periods. This reflects the inherent advantage of shorter physical and operational distances between production and consumption nodes. In practice, the decentralised units can respond more promptly to emergent demand fluctuations, thereby mitigating the risk of stockouts despite their smaller individual production capacity. This does not, however, translate into dramatically superior outcomes at the clinic tier (as previously noted), because the centralised system compensates through largescale throughput capability. Nevertheless, the decentralised architecture clearly delivers greater temporal responsiveness, and this confers a form of operational resilience that is particularly relevant for time sensitive or specialist pharmaceutical products.

Sustainability Metrics

Figure 4-Right shows the ratio of emissions in the centralised and decentralised cases. It can be seen that for the majority of the period considered, the centralised emissions are significantly greater than for the decentralised case. This is a result of the improved technologies used, but the variation in this ratio over time also shows a story of changing production rates. Spikes in the emissions ratios correspond to instances where centralised manufacturing increases in response to a low coverage ratio (see Figure 4-Left) whilst dips correspond to the inverse case with increased decentralised manufacturing.

This behaviour mirrors the technological assumptions embedded in the model: MMIC represents a traditional batch-based facility with correspondingly high emissions intensive cycles, whereas CMAC and NMIS operate small manufacturing processes with materially lower per unit emissions. The data confirms that even when emissions from both decentralised sites are combined, total environmental impact remains substantially below the centralised equivalent. This establishes that decentralisation provides a robust and widening sustainability advantage, conditional on the adoption of advanced continuous technologies rather than scaled-down legacy batch operations.

Beyond emissions, the manufacturing tier labour burden could also be compared across centralised and decentralised architectures, since site level carbon footprints can be translated into indicative fulltime equivalent (FTE) requirements using established energy and process intensity correlations. This linkage provides an

additional operational lens, suggesting that lower emission continuous manufacturing sites may also exhibit reduced labour intensity relative to traditional batch-based facilities of equivalent throughput.

Synthesis and Implications

The findings reveal that the comparative merits of centralised and decentralised manufacturing are multidimensional and echelon specific. The centralised architecture optimises inventory efficiency, while the decentralised architecture optimises responsiveness and sustainability. For strategic planning, this implies that pharmaceutical supply chains must prioritise between cost efficiency, resilience, and environmental performance, recognising that these objectives are not simultaneously maximised by any single configuration.

The results show that centralised manufacturing benefits from statistical risk pooling, which reduces total manufacturer-level inventory and stabilises upstream stock levels. This efficiency, however, comes at the cost of longer replenishment paths that require hospitals to maintain higher safety stocks and accept slower response during demand surges.

Decentralised manufacturing exhibits the opposite behaviour. Shorter transport distances and geographically distributed production enable faster recovery from demand spikes, as reflected by higher coverage ratios during high-demand periods. This responsiveness is complemented by significantly lower emissions due to continuous manufacturing technologies used at decentralised sites.

These patterns confirm that no single architecture is universally optimal. Centralisation minimises inventory, decentralisation maximises responsiveness and sustainability, and each performance dimension emerges from structural features of the network rather than parameter tuning alone.

CONCLUSIONS

This study demonstrates that the comparative performance of centralised and decentralised pharmaceutical manufacturing is echelon dependent rather than monolithic. At the periphery of the network (the clinics) both architectures converge to broadly similar service outcomes, indicating that geographical proximity alone does not yield decisive gains in routine delivery performance. At the hospital tier, the centralised configuration imposes a modest but persistent inventory uplift relative to the decentralised case, implying a higher holding cost burden in an inherently expensive stockkeeping environment. The effect is operationally material despite its limited magnitude and is consistent with longer replenishment paths from a single remote hub. At the manufacturing tier, the trade-off reverses. Decentralisation raises total upstream inventory, particularly pronounced at individual sites, reflecting the expected loss of risk pooling when capacity is fragmented across multiple smaller facilities. This escalates working capital requirements even as it improves responsiveness.

Critically, decentralisation delivers stronger inventory coverage ratios and greater temporal responsiveness during demand intensification, strengthening resilience against localised shortages. However, this advantage does not translate into a dramatic superiority at the clinic level, because the centralised system offsets through high throughput.

The environmental signal is unambiguous. The ratio of emissions intensity (centralised/decentralised) remains above unity throughout and increases as the simulation progresses, establishing a clear and widening sustainability advantage for decentralised production, conditional on the adoption of modern continuous processes rather than scaled down legacy batch operations.

There is no single dominant architecture. Centralisation minimises aggregate inventory via pooling, while decentralisation optimises responsiveness and sustainability. Strategic design must therefore be explicit about priorities: cost efficiency (inventory), resilience (coverage and surge handling), or environmental performance. A balanced network strategy may combine a limited number of advanced, decentralised nodes to secure responsiveness and carbon performance, while retaining selected centralised capabilities to preserve pooling efficiency, were economically critical.

The results should be interpreted within the bounds of the modelled demand cycles and behavioural rules. Extending the analysis to alternative demand regimes, richer disruption scenarios, and technology pathways would further illuminate design envelopes where hybrid architectures outperform either extreme. Incorporating explicit costings for transport, quality systems, and regulatory overheads across sites would allow a fuller economic assessment to complement the inventory, coverage, and emissions metrics reported here.

ACKNOWLEDGEMENTS

The authors from the University of Sheffield acknowledge funding provided through the UK–Southeast Asia Vaccine Manufacturing Research Hub. The Hub is funded by the Department of Health and Social Care using UK International Development funding and is managed by EPSRC. The views expressed in this publication are those of the authors and not necessarily those of the Department of Health and Social Care. The authors from the University of Strathclyde thank CMAC National Facility, Digital Medicines Manufacturing (DM^2) Research Centre co-funded by the Made Smarter Innovation challenge at UK Research and Innovation (grant ref:

EP/V062077/1), and UKRPIF (UK Research Partnership Institute Fund) capital award, SFC ref. H13054, from the Higher Education Funding Council for England (HEFCE).

REFERENCES

1. Aivazi MA, Rasekh HR, Peikanpour M, Peiravian F, Esmaeili S, Zarei L. The COVID-19 pandemic and the resilience of the pharmaceutical supply chain: lessons from past experiences and strategies for the future. Iran J Pharm Res 23: (2024). https://doi.org/10.5812/ijpr-152723
2. Jamil H, Leena Anjum , Nohri AR, Gul Sama , Ali Hamza , Saima Asif . RESILIENCE IN PHARMACEUTICAL SUPPLY CHAIN DURING PANDEMICS: a NARRATIVE REVIEW. IJLSS 3:96-104 (2025). https://doi.org/10.71000/bh4tn830
3. Howlett M, Niarchou E, Naughton BD. The response of the irish pharmaceutical manufacturing sector to the COVID-19 pandemic with a focus on lean and JIT. The Journal of Medicine Access 8: (2024). https://doi.org/10.1177/27550834241250279
4. Gupta G, Gupta M. Pharmaceutical supply chain role in managing customer service and medicine delivery amid shortages to USA customers. IJMER 6:22-35 (2025). https://doi.org/10.54660/ijmer.2025.6.1.22-35
5. Faggioni F, Valerio Rossi M, Sestino A. Supply chain resilience in the pharmaceutical industry: a qualitative analysis from scholarly and managerial perspectives. IJBM 18:129 (2023). https://doi.org/10.5539/ijbm.v18n1p129
6. Antonio Lozano-Diez J, Marmolejo-Saucedo JA, Rodriguez-Aguilar R. Designing a resilient supply chain: an approach to reduce drug shortages in epidemic outbreaks. EAI Endorsed Trans Perv Health Tech 6:e2 (2020). https://doi.org/10.4108/eai.13-7-2018.164260
7. Golan MS, Trump BD, Cegan JC, Linkov I. Supply chain resilience for vaccines: review of modeling approaches in the context of the COVID-19 pandemic. IMDS 121:1723-1748 (2021). https://doi.org/10.1108/imds-01-2021-0022
8. Das D, Datta A, Kumar P, Kazancoglu Y, Ram M. Building supply chain resilience in the era of COVID-19: an AHP-DEMATEL approach. Oper Manag Res 15:249-267 (2021). https://doi.org/10.1007/s12063-021-00200-4
9. Kiers J, Seinhorst J, Zwanenburg M, Stek K. Which strategies and corresponding competences are needed to improve supply chain resilience: a COVID-19 based review. Logistics 6:12 (2022). https://doi.org/10.3390/logistics6010012
10. Busola Fadojutimi . Exploring the critical factors in pharmaceutical supply chains revealed during COVID-19 and addressing missing links. World J. Adv. Res. Rev. 24:2407-2425 (2024). https://doi.org/10.30574/wjarr.2024.24.1.3283
11. Rahman T, Paul SK, Agarwal R, Shukla N, Taghikhah F. A viable supply chain model for managing panic-buying related challenges: lessons learned from the COVID-19 pandemic. International Journal of Production Research 62:3415-3434 (2023). https://doi.org/10.1080/00207543.2023.2237609
12. Andiç-Mortan E, Gonul Kochan C. Modeling a closed-loop vaccine supply chain with transshipments to minimize wastage and threats to the public: a system dynamics approach. JHLSCM 13:216-234 (2023). https://doi.org/10.1108/jhlscm-10-2021-0102
13. Gonul Kochan C, Nowicki DR, Sauser B, Randall WS. Impact of cloud-based information sharing on hospital supply chain performance: a system dynamics framework. International Journal of Production Economics 195:168-185 (2018). https://doi.org/10.1016/j.ijpe.2017.10.008
14. du Plessis, M., J. van Vuuren, and J. van Eeden, Reinforcement learning for information-sharing pharmaceutical supply chains.
15. Hindiyeh M, Altalafha T, Al-Naerat M, Saidan H, Al-Salaymeh A, Sbeinati L, Saidan M. Process modification of pharmaceutical tablet manufacturing operations: an eco-efficiency approach. Processes 6:15 (2018). https://doi.org/10.3390/pr6020015

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Developing predictive models for batch cooling crystallization of APIs with limited data availability

Mauro Davanzo[a], Emanuele Tomba[b], Enrico Carlassare[b], Riccardo Motterle[b], Massimiliano Barolo[a], Zoltan K. Nagy[c], and Fabrizio Bezzo[a]*

[a] University of Padova, Department of Industrial Engineering, Padova PD, Italy
[b] F.I.S. Fabbrica Italiana Sintetici S.p, a., Montecchio Maggiore VI, Italy
[c] Purdue University, Davidson School of Chemical Engineering, West Lafayette, Indiana, USA
* Corresponding Author: fabrizio.bezzo@unipd.it.

ABSTRACT

The objective of this work is to investigate strategies for the calibration of crystallization models aimed at predicting particle size distributions (PSDs) of active pharmaceutical ingredients (APIs) when using industrial datasets, which are limited in terms of number or information for the modeling exercise. In this work, the calibration task relies on two kinds of measurements, commonly performed in industrial crystallization practice: offline measurements of PSDs and API solute concentration carried out only at the beginning and at the end of experiments, and online measurements of chord length distributions (CLDs). Particularly, a strategy is proposed to use CLDs data from focused beam reflectance measurement (FBRM) probes as proxies of the PSD, which is the main key performance indicator for the model exercise. Industrial data concerning a seeded batch cooling recrystallization of an API in an organic solvent are used as a case study. The PharmaPy process simulator is used for parameter estimation and process simulation. Results demonstrate that, with proper data processing and feature extraction, all parameters can be estimated with sufficient precision. The model performance is satisfactory for most of the batch duration, even though some shortcomings highlight possible limitations in the data and/or in the model itself. From the industrial perspective, results pave the way for a quantitative usage of FBRM probes to enhance process understanding and to guide process development and scale-up.

Keywords: Pharmaceuticals, Crystallization, Population Balances, Modelling, Parameter estimation

INTRODUCTION

Crystallization is a key step in the downstream manufacturing of active pharmaceutical ingredients (APIs), since it is responsible for the isolation of the target compound as a stable solid phase. It has a huge impact on the final product's critical quality attributes (CQAs), which are mainly related to purity and size distribution of the obtained particles [1]. Models can be employed to predict such CQAs based on system properties and operating conditions, allowing for simulation, optimization and scale-up of crystallization unit operations. In this way, the application of model-based approaches can represent an answer to the increasing demand in research efficiency and quality assurance in pharmaceutical process design and setup [2].

A variety of phenomena can influence the evolution of crystal size distribution (CSD) during a crystallization process: nucleation and growth are typically the main ones, but also breakage and agglomeration of particles can occur. Whether each of these phenomena is triggered, and the mechanism according to which it may develop, depend on many different factors, e.g. the way and the rate at which the driving force is generated, the operation history, the equipment geometry and mixing, etc.

Population balance models (PBMs), combined with momentum, mass and energy balances, may allow for a comprehensive description of the crystallization process [3]. They are founded on the knowledge of the physics of the system, preserving interpretability and demanding less information to be calibrated with respect to data-driven alternatives. Moreover, they are more suitable to use such information, typically retrievable from

laboratory scale, to represent crystallizers at the industrial scale. For this reason, they are a very common and powerful tool for crystallization modelling [4].

Still, the calibration of PBMs is a non-trivial task, and often relies on multiple experiments characterized by different operating conditions, with measurements of particle size distribution (PSD) and concentration that allow to track the evolution of CQAs during the process [5, 6]. Performing such measurements can be difficult and time consuming, especially when it comes to PSD, due to the complexity either in treating samples for offline measurements or in getting the desired data for PBM calibration from online probes, such as FBRM or in process microscopy. However, industrial contexts are usually characterized by tight constraints on time and budget, and therefore the minimization of resources required for experimental campaigns is paramount. Therefore, developing calibration methods for PBMs that require fewer data points is important when available instrumentation, experimental setups, and sample sizes are limited.

For instance, exploiting measurements of crystals' chord length distributions (CLDs) is deemed as an interesting solution to maximize the extraction of information for PBMs calibration, even if in industrial practice they are typically used as qualitative indicators only [7]; in fact, they are often available and can be continuously recorded by process analytical technology (PAT) tools installed inline. However, CLDs are inherently different from CSDs, and then system-specific transformations should be applied in order to be able to use those quantities for parameter estimation [8]. In this work a more general strategy is proposed, based on the assumption that the CLDs collected by focused beam reflectance measurement (FBRM) probes can be compared to CSDs obtained as model outputs through their derivatives. Thus, trends in CLD data are combined with concentration and PSD data gathered only for the initial and final conditions of each batch, in order to calibrate PBMs with good predictive capability for industrial purposes.

MATERIALS AND METHODS

The proposed approach is presented in this section, starting from the definition of the batch cooling crystallization model. After that, the strategy for model calibration with limited datasets is illustrated, together with criteria for the assessment of fitting quality. Finally, the application to the industrial case study concerning API crystallization are commented on.

Model structure

Mechanistic models are chosen over data-driven models, since their calibration needs less data, they are interpretable and they are typically more suitable for scale-up problems. In particular, population balance models are a staple in the field, since they allow to describe the evolution of distributed properties, like particle sizes. Hence, the following formulation of the population balance equation (Equation 1) is derived from the general form for the case of a simple batch cooling crystallization in which mixing can be assumed ideal, volume of the slurry can be considered constant, 3D particles are characterized by a single equivalent dimension, and there are no significant dissolution, breakage or agglomeration phenomena:

$$\frac{\partial n(t,L)}{\partial t} + G \bullet \frac{\partial n(t,L)}{\partial L} = B \bullet \delta \bullet (L - L_0) \quad (1)$$

Here, n represents the CSD, depending on time t and the internal coordinate L, B and G represent the nucleation and growth kinetic terms, respectively, δ is the Kronecker's delta and L_0 is the minimum value of the internal coordinate L, namely the minimum characteristic size of a particle. As typically done [4], the terms for nucleation (B) and growth (G) kinetics are defined as Arrhenius-like laws based on the relative supersaturation.

$$B = B_s = e^{ln(k_{bs})} \bullet \mu_3^j \bullet S(T)^{bs} \quad (2)$$

$$G = e^{ln(k_g)} \bullet S(T)^g \quad (3)$$

$$S = \frac{c_{tgt} - c_{sat}(T)}{c_{sat}(T)} \quad (4)$$

$$c_{sat} = A + \frac{B}{T} + C \bullet ln(T) \quad (5)$$

This latter is denoted with S (Equation 4) and, in turn, depends on the difference between the concentration of the target compound in the liquid (c_{tgt}) and its solubility (c_{sat}, Equation 5), which varies with the temperature (T). If the crystallization is seeded, it is reasonable to take into account the effect of the total volume of particles already present in the system (μ_3) on the generation of new crystals, and thus to formulate B as a secondary nucleation term (B_s, Equation 2).

K_{bs}, j, bs, k_g and g are kinetic parameters specific to each crystallization process, while A, B and C are coefficients of the Apelblat equation, used to describe the solubility curve. In addition, the model includes the equation for the mass balance, based on concentrations, coupled with an equation to track the evolution of liquid volume:

$$\frac{dc_i}{dt} = -\frac{\rho_{crys} \bullet k_v \bullet (3G \bullet \mu_2 + B \bullet V_{liq} \bullet L_0^3)}{V_{liq}} \bullet (\delta - \frac{c_i}{\rho_{liq}}) \quad (6)$$

$$\frac{dV_{liq}}{dt} = -\frac{\rho_{crys} \bullet k_v \bullet (3G \bullet \mu_2 + B \bullet V_{liq} \bullet L_0^3)}{\rho_{liq}} \quad (7)$$

where, c_i is the concentration of the compound i, V_{liq} is the liquid volume, k_v and μ_2 are the shape factor and the total surface area of the crystals, while ρ_{cryst} and ρ_{liq} are the densities of the crystals and of the liquid, respectively. The energy balance is not considered, instead, since temperature is assumed to be perfectly controlled

and to follow a known temperature profile. The standard method of moments (SMOM) is used to discretize the partial differential equations, obtaining a closed form of the set of equations to be solved [3].

Model calibration

In order to estimate the model parameters, pre-processing of available datasets is required. In this way, the model calibration strategy is refined by excluding possible outliers in the data and by selecting which model responses or features should be considered. Total counts and particles' population weighted mean diameter are chosen as core quantities for the calibration method, together with initial and final API concentration. In particular, derivatives over time of such core quantities are employed, based on the assumption that they are comparable for CLDs and CSDs. In this way, trends of FBRM measurements can be used in a quantitative manner, combined with the few samples that are available for the other types of measurements. Average size allows to solve potential consistency issues between the information provided by CLD trends and that provided by initial and final PSD measurement, so that the maximum amount of information can be exploited. The objective function (OBJ) to be minimized is defined as a weighted mean square error (WMSE) that compares model responses and experimental data:

$$\begin{aligned} OBJ = & \frac{w_0}{N_{CLD}} \sum_i^{N_e} \sum_j^{N_{CLD}} \left(\mu_{0,m,ij}{}' - \mu_{0,CLD,ij}{}'\right)^2 \\ & + \frac{w_1}{N_{CLD}} \sum_i^{N_e} \sum_j^{N_{CLD}} \left(\left(\frac{\mu_{1,m,ij}}{\mu_{0,m,ij}}\right)' - \left(\frac{\mu_{1,CLD,ij}}{\mu_{0,CLD,ij}}\right)' \right)^2 \\ & + \frac{w_2}{N_{PSD}} \sum_i^{N_e} \sum_j^{N_{PSD}} \left(\left(\frac{\mu_{1,m,ij}}{\mu_{0,m,ij}}\right) - \left(\frac{\mu_{1,PSD,ij}}{\mu_{0,PSD,ij}}\right) \right)^2 \\ & + \frac{w_3}{N_c} \sum_i^{N_e} \sum_j^{N_c} \left(c_{m,ij} - c_{ij}\right)^2 \end{aligned} \quad (8)$$

Here, w coefficients indicate the weights attributed to each of the terms in the summation, N_e indicates the number of experiments, N_{CLD}, N_{PSD} and N_c indicate the number of samples available for each type of measurement, c indicates concentrations and μ_0 and μ_1 indicate moments of size distributions. Subscripts of moments refer to the kind of size distribution (PSD or CLD) and the applied pre-treatments that are used to perform the calculations. All the quantities reported with the subscript m, instead, correspond to model outputs. It is important to highlight that the first two terms involve derivatives over time of total counts and particles' mean diameter, respectively, instead of their absolute values.

The SciPy implementation of adaptive Nelder-Mead algorithm is used for parameter estimation; multiple starts from different initial conditions reduce the possibility to incur local minima. Note that, for the same reason, terms in Equations 2 and 3 are formulated by employing dimensionless transformations of all parameters, and logarithmic versions of the pre-exponentials. The precision of the parameters' estimates is evaluated by means of a t-test, according to Equation 9.

$$t_{\theta,i} = \frac{\hat{\theta}_i}{\sqrt{\nu_{\theta,ii}}} > t_{ref}(1 - \alpha/2, N_{sp} - N_\theta) \quad (9)$$

Each parameter estimate $\hat{\theta}_i$ is considered statistically meaningful only if the related t-value t_{θ_i} exceeds a specific reference value t_{ref}. This latter is identified as the percentile with α confidence of a standardized t-distribution, characterized by a number of degrees of freedom which is obtained as the difference between the number of samples N_{sp} and N_θ. Parameter uncertainties are computed based on the diagonal terms of the variance-covariance matrix $\nu_{\theta,ii}$, which in turn is retrieved from the local sensitivities of responses with respect to parameters, thanks to common assumptions of local linearity [10].

Case study

Model and methods implementations are built around tools and functionalities provided by the PharmaPy library, which serves as core process simulator [11]. Procedures are tested on real data collected in industrial laboratories for a seeded batch cooling recrystallization of an API in an organic solvent. The experimental campaign includes five experiments: three runs are distinguished only by different temperature profiles, while the other two are characterized by either different batch size or seed quantity. Indirect estimates of initial and final concentration are available for all the experiments, together with measurements of the seed and final product PSDs obtained by laser diffraction on the dried product. In addition, the evolution of the CLDs is tracked online by FBRM probes during the entire batches. Furthermore, for all the experiments, some data points are removed from the calibration set mainly due to probe fouling and cleaning, or to the unwanted recording of the addition of seed material.

RESULTS

In this section, results that were achieved by applying the presented methodology to the industrial case study are discussed. The model was calibrated on the data gathered during the experimental campaign described in the previous paragraph.

Table 1 shows the estimates of the kinetic parameters. It can be noticed that their values are outside the typical ranges [4]. In particular, the exponent in the growth term is higher than the exponent in the nucleation term, implying that the growth phenomenon is faster and requires higher levels of supersaturation to be triggered, with respect to nucleation. Even though this is not commonly verified in crystallization systems, it is actually coherent with the behavior of the available experimental

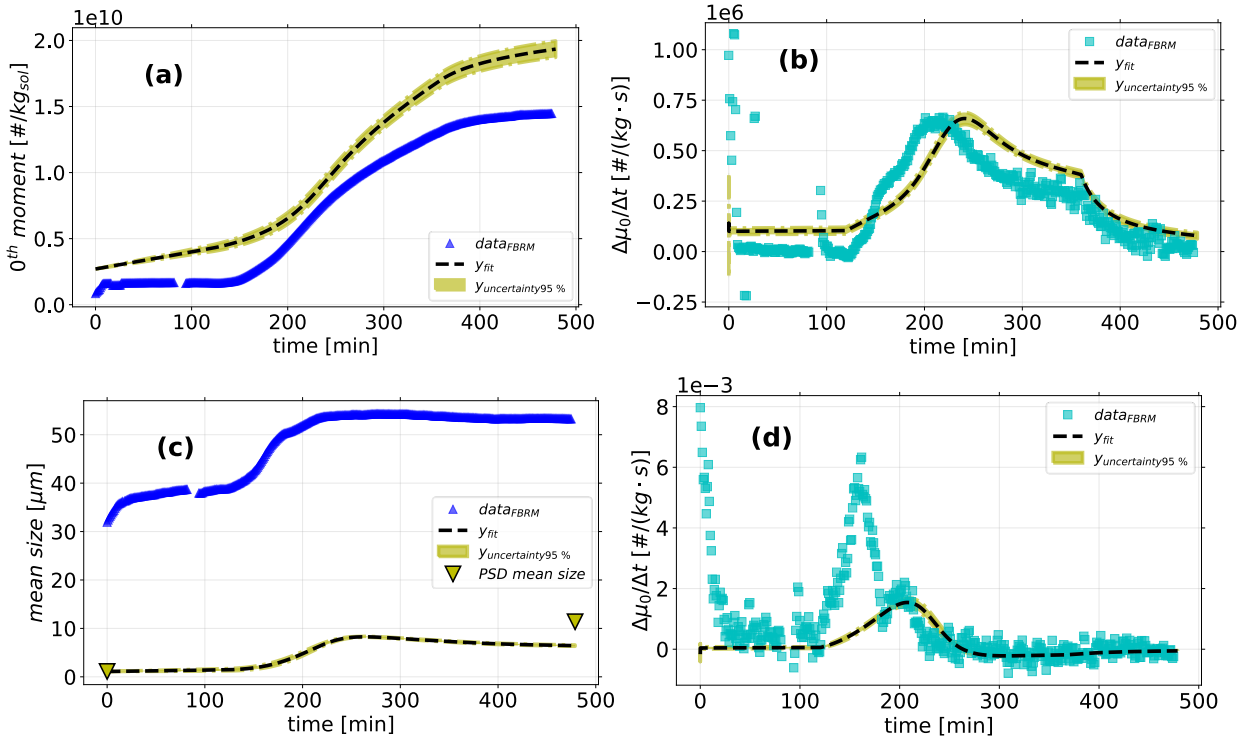

Figure 1: Fitting results for the first experiment of the campaign. Each graph shows: the quantity derived from chord length distribution data (*data*$_{FBRM}$, dots), the curve of model output, or feature, obtained from simulation with estimated values of kinetic parameters (y_{fit}, dashed line), 95 % confidence bounds ($y_{uncertainty,95\ \%}$, yellow area). **a)** total counts. **b)** Derivative of total counts. **c)** Average size, with initial and final point values computed from laser diffraction measurements (*PSD mean size*, yellow triangles). **d)** Derivative of average size.

data gathered by the FBRM probes.

Table 1: Estimated values of the kinetic parameters with the respective units, 95 % confidence intervals (95 % CI) and t-values. t_{ref} indicates the reference threshold value for the t-test.

Parameter	Units	$\hat{\theta}$	95% CI	t-value
$log\ (k_{bs})$	$\frac{\#}{kg \bullet s}$	21.79	0.13	333.72
b_s	-	0.54	0.09	11.37
j_s	-	0.57	0.06	17.96
$log(k_g)$	$\frac{\mu m}{s}$	-5.92	0.07	165.57
g	-	1.61	0.18	17.18
				$\lvert t_{ref} \rvert = 1.96$

Moreover, computed uncertainties are relatively low, as proven by the very high t-values that easily overcome the threshold reference value for all the parameters; for this reason, all of them can be considered as statistically significant. This is a non-obvious finding because population balance models, and especially Arrhenius-like kinetic equations, often suffer from strong correlation among parameters, which undermines parametric precision. In this case, the large number of available samples allows decreasing the uncertainty on the estimates, and it represents one of the main advantages of exploiting data from online probes.

Fitting results for the first experiment of the campaign are represented in Figure 1. Each blue dot in Figure 1.a represents the total number of particles (or 0th moment) calculated from the CLD measured by the FBRM probe at a specific time point. The same holds for blue dots in Figure 1.c, except that they refer to population weighted mean size. In both cases, the dashed lines represent the corresponding curves obtained from model outputs after simulating with the estimated values of kinetic parameters. Furthermore, 95 % confidence bounds of the outputs are calculated by following a Monte Carlo approach, and are reported as yellow areas. In the case of average size (c), the initial and final values derived from the laser diffraction measurements are also included (yellow triangles). Figures 1.b and 1.d show the derivatives of what is displayed in 1.a and 1.c, respectively.

From the profiles it can be seen that the main trends in the experimental data are fitted by the model, in particular for the total counts. However, the increase in average size in the first half of the batch is not captured,

and this leads also to a slight underestimation in the prediction of the end-point value. In general, even by testing slightly different calibration setups, one part of the process is always less accurately described, highlighting the limitations in the available measurement techniques and, possibly, in the structure of the model itself. Offsets in the plots of the non-differentiated quantities are expected, since the model was trained on the derivatives of CLD data, but used information from the laser diffraction measurements as initial conditions.

The outcomes are deemed satisfactory, and demonstrate that FBRM measurements can be used as complementary information to more reliable data about the starting point of the process and product properties. The proposed strategy can be very valuable from the industrial perspective, especially for early-stage process development, where resources are limited and maximizing information from scarce data can be a key advantage.

CONCLUSIONS

First-principles models based on population balances are powerful tools that can enhance the understanding and aid the development of crystallization processes for API production. Identifying a proper structure and achieving an accurate calibration for PBMs can be a challenging task, particularly in an industrial environment, where thorough experimental campaigns may not be affordable, and advanced instrumentation and setup may not be available. Thus, being able to exploit in a quantitative manner trends in data that are usually interpreted as qualitative indications of process evolution, like CLDs measured by FBRM probes, can be crucial. The results of this work demonstrate the potential of strategies built around the usage of such CLD data, in combination with a minimum amount of offline measurements and with proper data processing, feature extraction and numerical routines. In fact, statistically precise estimations for kinetic parameters of basic PBMs formulations that include nucleation and growth mechanisms are obtained, together with good prediction capabilities for early-stage process development purposes. Future work will be focused on assessing the impact that measurement availability can have on model performances, which will be evaluated both in calibration and in validation. Similar approaches will be applied on different datasets and by varying the amount and type of data used for parameter estimation, with the aim of finding an optimum trade-off between model quality and experimental effort.

ACKNOWLEDGEMENTS

This research was funded by Ministero dell'Università e della Ricerca (MUR) through the scholarship 8646 (course: DOT1319737, cycle: XXXIX), as part of PNRR program, Missione 4, componente 2 "Dalla Ricerca all'Impresa" – Investimento 3.3. The author also acknowledges F.I.S. Fabbrica Italiana Sintetici S.p.A for their support and contribution to the work.

AUTHOR IDENTIFIERS

Author ORCIDs:
Davanzo M: 0009-0002-0739-7541
Barolo M: 0000-0002-8125-5704
Nagy ZK: 0000-0003-4787-6678
Bezzo F: 0000-0003-1561-0584

REFERENCES

1. Myerson AS. Preface to the second edition. Handbook of Industrial Crystallization :xiii (2002). https://doi.org/10.1016/b978-075067012-8/50002-1
2. Destro F, Barolo M. A review on the modernization of pharmaceutical development and manufacturing – trends, perspectives, and the role of mathematical modeling. International Journal of Pharmaceutics 620:121715 (2022). https://doi.org/10.1016/j.ijpharm.2022.121715
3. Omar HM, Rohani S. Crystal population balance formulation and solution methods: a review. Crystal Growth & Design 17:4028-4041 (2017). https://doi.org/10.1021/acs.cgd.7b00645
4. Maldonado DA, Vassileiou A, Johnston B, Florence AJ, Brown CJ. Data mining crystallization kinetics. Digital Discovery 1:621-635 (2022). https://doi.org/10.1039/d2dd00033d
5. Szilágyi B, Eren A, Quon JL, Papageorgiou CD, Nagy ZK. Digital design of the crystallization of an active pharmaceutical ingredient using a population balance model with a novel size dependent growth rate expression. from development of a digital twin to in silico optimization and experimental validation. Crystal Growth & Design 22:497-512 (2021). https://doi.org/10.1021/acs.cgd.1c01108
6. Barhate Y, Kilari H, Wu WL, Nagy ZK. Population balance model enabled digital design and uncertainty analysis framework for continuous crystallization of pharmaceuticals using an automated platform with full recycle and minimal material use. Chemical Engineering Science 287:119688 (2024). https://doi.org/10.1016/j.ces.2023.119688
7. Orosz Á, Szilágyi B. Maximizing similarity: using correlation coefficients to calibrate kinetic parameters in population balance models. Heliyon 10:e39851 (2024). https://doi.org/10.1016/j.heliyon.2024.e39851
8. Szilágyi B, Agachi P?, Nagy ZK. Chord length

distribution based modeling and adaptive model predictive control of batch crystallization processes using high fidelity full population balance models. Ind. Eng. Chem. Res. 57:3320-3332 (2018). https://doi.org/10.1021/acs.iecr.7b03964

9. Franceschini G, Macchietto S. Model-based design of experiments for parameter precision: state of the art. Chemical Engineering Science 63:4846-4872 (2008). https://doi.org/10.1016/j.ces.2007.11.034
10. Casas-Orozco D, Laky D, Wang V, Abdi M, Feng X, Wood E, Laird C, Reklaitis GV, Nagy ZK. Pharmapy: an object-oriented tool for the development of hybrid pharmaceutical flowsheets. Computers & Chemical Engineering 153:107408 (2021). https://doi.org/10.1016/j.compchemeng.2021.107408

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

An in silico/in vitro approach for uncertainty-aware hybrid models for template-induced protein crystallisation systems

Daniele Pessina[ab], Jerry Y. Y. Heng[bc], and Maria M. Papathanasiou[ab]*

[a] Sargent Centre for Process Systems Engineering, Imperial College London, SW7 2AZ, United Kingdom
[b] Department of Chemical Engineering, Imperial College London, SW7 2AZ, United Kingdom
[c] Institute of Molecular Science, Department of Chemical Engineering, Imperial College London, SW7 2AZ, United Kingdom
* Corresponding Author: maria.papathanasiou11@imperial.ac.uk

ABSTRACT

Crystallisation is a promising and scalable alternative to chromatography for biologics purification. However biologics such as proteins and peptides often crystallise only in narrow operating windows, limiting process flexibility. Template-induced crystallisation can lower supersaturation requirements and expand feasible operating ranges, yet the template dependence of nucleation and growth kinetics remains difficult to parametrise mechanistically. To address this, we develop and experimentally validate uncertainty-aware hybrid models for lysozyme crystallisation on hydroxyl- and carboxyl-functionalised silica templates. A mechanistic population-balance model is coupled to a data-driven regressor that maps operating conditions and template variables to effective nucleation and growth rates. We compare a neural network baseline against a structured neural power-law surrogate, which embeds a supersaturation-dependent power-law form. Both hybrid models are trained in-the-loop via differentiable simulation, and variational inference is used to obtain posterior parameter distributions and calibrated predictive uncertainty. Across cross-validation and off-grid tests at previously unseen combinations of temperature and template loading, the hybrid models accurately reproduce solute concentration dynamics and capture key particle-size trends, while the neural power-law surrogate provides improved robustness and faster uncertainty quantification. These results support hybrid, uncertainty-aware PBMs as practical tools for prediction, design-space exploration, and comparison of template-enabled protein crystallisation processes.

Keywords: Crystallisation, Hybrid Models, Uncertainty-aware

INTRODUCTION

Crystallisation is emerging as a viable downstream alternative to chromatography for the purification and formulation of biologics-based therapeutics. As demand for protein- and peptide-based drugs continues to increase, improvements in upstream productivity have shifted the bottleneck to the purification train where conventional chromatographic steps are capital- and material-intensive and can be challenging to scale. In contrast, crystallisation produces a highly pure, easily filterable solid form of the active pharmaceutical ingredient with favourable stability, storage and handling characteristics, and is inherently well suited to scale-up and continuous processing [1]. Yet despite its extensive use for small-molecule APIs, translating crystallisation into a robust, industrially relevant unit operation for macromolecules remains difficult, and largely empirical due to the complexity of protein crystallisation, resulting in poor control of the process and its' quality attributes [2].

To overcome such limitations and widen the conditions at which crystallisation occurs, template-assisted crystallisation has attracted growing interest. In templated systems, a finely dispersed solid phase, such as a porous or nanostructured material with tailored surface chemistry, is introduced into the solution to promote nucleation. Experimental studies have shown that varying template loading, pore structure and surface

functionalisation can lower the supersaturation required for nucleation, shift metastable zone boundaries and improve control over crystal size and habit for both small molecules and biologics [3–6]. In the context of protein crystallisation, such templates offer a route to mitigate the harsh supersaturation conditions and long induction times often required to obtain crystals, and introduce the possibility of designing more flexible, template-enabled crystallisation processes for biologics. However, these benefits come at the cost of additional complexity: the nucleation behaviour now depends on a richer set of process and material descriptors, including template loading, surface charge, chemistry, pore architecture, and their coupling with solution conditions such as pH, buffer and precipitant concentration.

Mechanistic population balance models (PBMs) provide a powerful tool to describe crystallisation, combining mass and energy balances with kinetic expressions for nucleation, growth, aggregation and breakage. They can be used to rationalise experimental trends, perform *in-silico* process development, and support scale-up and control of both small-molecule and protein crystallisers [7]. In conventional applications, nucleation and growth rates are described with a relatively small number of parameters using, for example, classical nucleation theory and empirical power-law kinetics. These expressions, while still imperfect, are often sufficient to capture the behaviour of unseeded or seeded cooling and anti-solvent crystallisations under well-characterised conditions. For template-induced systems, the effective nucleation mechanism involves a combination of adsorption, confinement, local changes in supersaturation and potentially intermediate dense-liquid or amorphous phases, all strongly modulated by template morphology and surface chemistry [4]. The unique, condition-dependent nature of protein crystallization means that static kinetic expressions can lack the flexibility to describe different templated systems. Since a unified framework for these interactions has yet to be established, there is no assurance that a model structure identified for one process will hold for another. This lack of transferability necessitates a more robust modeling approach that can internalize these dependencies without requiring a complete mechanistic rebuild for every new template loading or chemistry.

The core challenge is therefore the explicit parametrisation of the kinetics in terms of experimentally tunable template variables. For each new combination of template material, surface functionalisation and loading, one would, in principle, need to postulate a new kinetic expression, re-estimate kinetic parameters and re-assess identifiability and uncertainty. In practice, the scarcity and noise of protein crystallisation data, the stiffness and non-linearity of PBMs, and the lack of a clear mechanistic consensus on template-induced nucleation make this programme difficult to sustain. Mismatch between the assumed kinetic law and the true, template-dependent behaviour can yield models that appear to fit calibration data but fail to generalise to new template loadings or process conditions, undermining their usefulness for process design and risk assessment.

HYBRID MODELLING

Hybrid modelling offers an attractive route to overcome the obstacles of repeated parametrisation and uncertainty quantification for each unique template chemistry and loading. In a hybrid formulation, a data-driven regressor is embedded directly within the mechanistic PBM to predict effective kinetic rates from the current process state and operating variables, while the population balance and mass balance equations continue to enforce conservation and encode the underlying physics. Rather than writing down an explicit expression for the nucleation rate as a function of supersaturation and template loading, we instead learn this relationship from data using a flexible model class. This architecture is particularly appealing for template-induced protein crystallisation, where relevant inputs naturally include both solution variables (e.g., temperature, current protein concentration) and template descriptors (e.g., loading). The model can possibly be augmented with further process variables (buffer and precipitant conditions, average pore diameter, zeta potential, surface functional group) which might be difficult to integrate in an explicit kinetic expression. By learning a mapping from this input space to effective nucleation and growth rates, hybrid models can internalise complex, template-specific effects without sacrificing the interpretability and extrapolation structure provided by the mechanistic balances. Nielsen et al. successfully applied this methodology to model a crystallisation process, and noted that while hybrid models have higher data requirements than purely mechanistic models, they are effective in cases where there is high-frequency measurements from PAT and incomplete knowledge, such that a mechanistic model could not have been properly formulated [8]. McKay et al. applied a parallel hybrid methodology for the control of a fed batch crystalliser, notably with high data availability and a large number of generated batches from a reference in-silico model. In this case, the data-driven regressor is successfully used to correct the mispredictions of a mechanistic model, formulated with kinetic parameters with 5% deviation from the reference parameters used to generate the original data [9]. In both cases, training data for the data driven regressor had to be carefully collected with simulations of the reference model to train the data-driven counterpart. It would therefore be advantageous to be able to train the whole hybrid model ‘in-the-loop’, without lengthy data generation tasks.

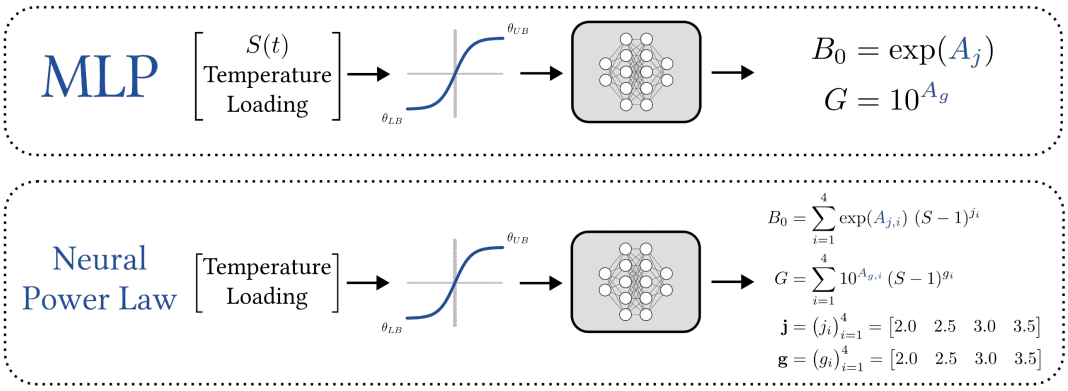

Figure 1. Structure of the two proposed regressors to be embedded in the mechanistic model.

VARIATIONAL INFERENCE

For hybrid models to be actionable in process development, point predictions are insufficient. Protein crystallisation experiments are inherently noisy, and the kinetic identification problem must account for both irreducible experimental noise (aleatoric) and limited confidence in regressed neural network parameters (epistemic). Uncertainty-aware formulations that quantify the posterior variability in kinetic surrogates, and propagate it through the PBM, are essential for assessing confidence in predicted trajectories and prioritizing future experiments. Traditional Bayesian techniques, such as Markov Chain Monte Carlo, are often computationally prohibitive for the high parameter counts typical of neural networks. Variational Inference (VI) offers a scalable alternative by reframing the inference problem as an optimization task [10]. The true but intractable posterior distribution of network weights $p(\omega|D)$ is approximated with a simpler, variational distribution $q_\theta(\omega)$ such as a multivariate Gaussian. The parameters θ are then optimized to minimize the Kullback–Leibler (KL) divergence from the true posterior, which is equivalent to maximizing the Evidence Lower Bound (ELBO).

$$\mathrm{ELBO} = \boldsymbol{E}_{q_\theta(\omega)}[\log p(\omega|D)] - \mathrm{KL}(q_\theta || p(\omega)) \tag{1}$$

The first term encourages the model to fit the experimental concentration and crystal size data, while the second acts as a Bayesian regularizer. By sampling from the optimized $q_\theta(\omega)$ during the ODE integration, we can propagate weight uncertainty through the population balance equations to generate predictive intervals for the entire process trajectory.

In this work, we address these challenges by developing and experimentally validating uncertainty-aware hybrid PBMs for lysozyme crystallisation in the presence of porous silica nanoparticle templates with different surface functionalisations. We consider two hybrid architectures in which a multi-layer perceptron neural network learns effective nucleation and growth rates as functions of process and template variables, while a mechanistic population balance model governs the evolution of the crystal size distribution and solute concentration. The data-driven components are trained in-the-loop via auto-differentiable programming, and variational inference is used to recover posterior distributions over the hybrid model parameters and corresponding predictive uncertainty. By validating these models on unseen template loadings and comparing their accuracy and calibration across distinct template chemistries, we aim to demonstrate that hybrid, uncertainty-aware formulations can bridge the gap between rich, template-enabled experimental design spaces and mechanistically grounded, in-silico process models for protein crystallisation.

METHODOLOGY

Experimental Methodology

This architecture is applied to experimental datasets of template-induced lysozyme crystallisation. Lysozyme was crystallised in a 0.1 M sodium acetate buffer (pH = 4.2) containing 50mg/ml NaCl. Experiments are conducted in both template-free solutions and in the presence of irregularly shaped, porous silica nanoparticles functionalised with carboxyl and or hydroxyl surface groups. For each tested condition, three independent replicates are performed. Crystallisations are carried out at three temperatures (17, and 21°C) and three nanoparticle loadings (1, 5 and 10 mg/ml) and 2 template chemistries (Hydroxyl and Carboxyl functionalised). Validation experiments were carried out at intermediate loading and temperatures (19°C, 7.5 mg/ml), to assess the interpolative ability of the hybrid models. Two initial protein concentrations are tested per temperature–loading combination. Reactions were conducted in 20 ml volumes

under constant stirring at 75 rpm. Protein concentration was monitored every 30 minutes using a UV Nanodrop spectrophotometer. Upon batch termination, the slurry is filtered, dried, and the crystal size distribution is measured using a laser-diffraction particle size analyser.

In-Silico Methodology

Eq 1 shows the population balance with negligible aggregation and breakage, while Eq2 describes the mass balance. To solve this hyperbolic PDE, a method of moments transformation is used to recover the moments of the crystal size distribution. Rather than defining explicit equation for nucleation and growth, neural networks are embedded in the mechanistic model framework. By implementing the model in an auto-differentiable language such as JAX, Pytorch or Julia, the neural networks can be trained by minimising a parameter estimation loss function and a gradient-based optimiser. In both regressors sigmoid input and output transformations are implemented to the predicted rates remain in feasible ranges that won't make the ODE simulation unstable when the network is first initialised. They also serve to improve the gradient conditioning and smoothness.

$$\frac{\partial n}{\partial t} + \frac{\partial (Gn)}{\partial L} = B_0 \delta(L - L_{min}) \quad (2)$$

$$\frac{dc}{dt} = -3\rho_c k_v \int_0^\infty GnL^2 \, dL \quad (3)$$

$$d_{43} = \frac{\mu_4}{\mu_3} \quad (4)$$

Figure 1 shows the two proposed data driven kinetic regressors. The MLP regressor takes the supersaturation at each ODE call, and the temperature and loading as inputs. Those are passed to a sigmoid function, scaled on global lower and upper bounds (Table 1). The sigmoid-transformed inputs pass through a neural network, which outputs a nucleation and growth coefficient which are finally passed to exponential functions for a physically-feasible scaling. The Neural Polynomial, is also proposed as an alternative architecture. In this architecture the neural network only takes temperature and loading as input, and outputs coefficients of a sum of power law expressions for nucleation and growth. Training of both regressors is done in a 2-stage approach to improve the conditioning of the regressors, in which the hybrid models are first trained on 50% of the initial timepoints before including the whole dataset. Both neural networks have 1 layer, 150 neurons, and were trained with the following schedule: Phase 1 LR = 8×10^{-4}, Steps = 40; Phase 2 LR = 2×10^{-3}, Steps = 200. Variational inference is performed with 5000 steps and LR = 1×10^{-3}. The prior distribution is centred around the weights of the trained hybrid model regressor, with 10% standard deviation.

Table 1. Bounds for sigmoid transformations

Name	Lower Bound	Upper Bound
Temperature [°C]	13	27
Template Loading [mg/ml]	0	30
Solute conc. [mg/ml]	0	25
A_j [-]	-15	30
A_g [-]	-15	-5

RESULTS AND DISCUSSION

Hold-one-out Validation

Hold-one-out (HOO) validation was performed using the training experiments, where each experiment is left out in turn and predicted using a model trained on the remaining experiments. As a result, some of the validation cases are extrapolating with respect to the initial conditions (and occasionally the transient state trajectory), while others are more interpolative. However, each validation experiment still shares the same discrete temperature and template loading values with at least some of the training experiments. In other words, across the HOO benchmark the model is never forced to extrapolate temperature or loading; the “hard” extrapolation is primarily in the initial condition and the resulting early-time dynamics.

It is important to note that HOO is an inherently biased validation procedure in data-scarce settings [11] as leaving out any single experiment can change the empirical mean and variance of the training set. This means HOO effectively introduces a fair amount of extrapolation into the evaluation, which is notoriously challenging for any regressor embedded in a stiff mechanistic simulator, but it also makes the benchmark quite strict. In practice, this is useful here because generalisation is often presented as a key advantage of mechanistic PBMs, and HOO provides a direct stress test of whether the hybrid formulation can retain that behaviour under limited data.

Table 2. Hold-one-out validation mean absolute error for hydroxyl and carboxyl systems, MLP and Neural P-L models.

Model	Conc. MAE [mg/ml]	PS MAE [micron]
	Hydroxyl System	
MLP	1.402	4.089
Neural P-L	1.439	4.078
	Carboxyl System	
MLP	1.876	5.506
Neural P-L	1.468	3.558

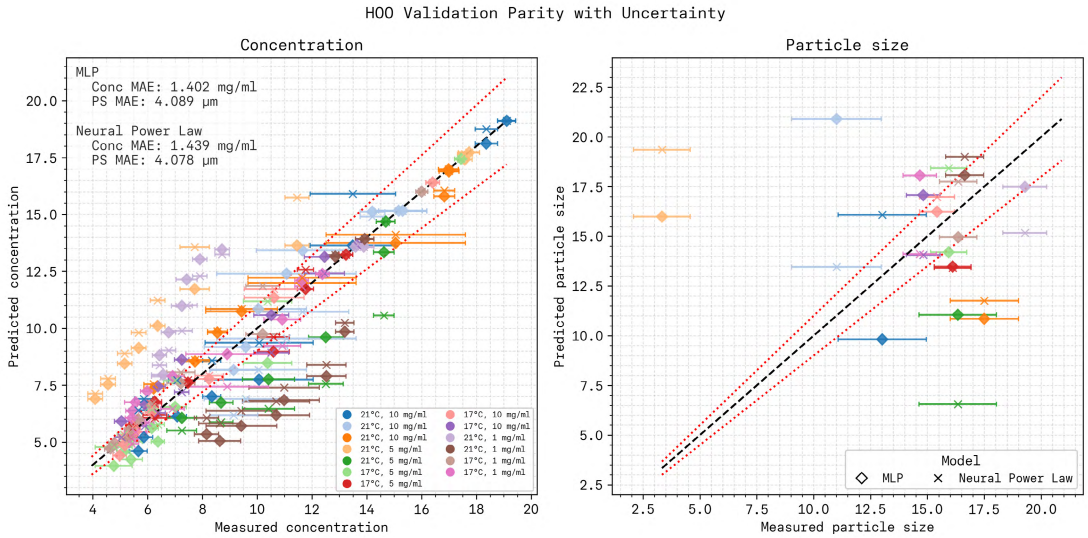

Figure 2: Hold-one-out Validation parity plot for the Hydroxyl template system.

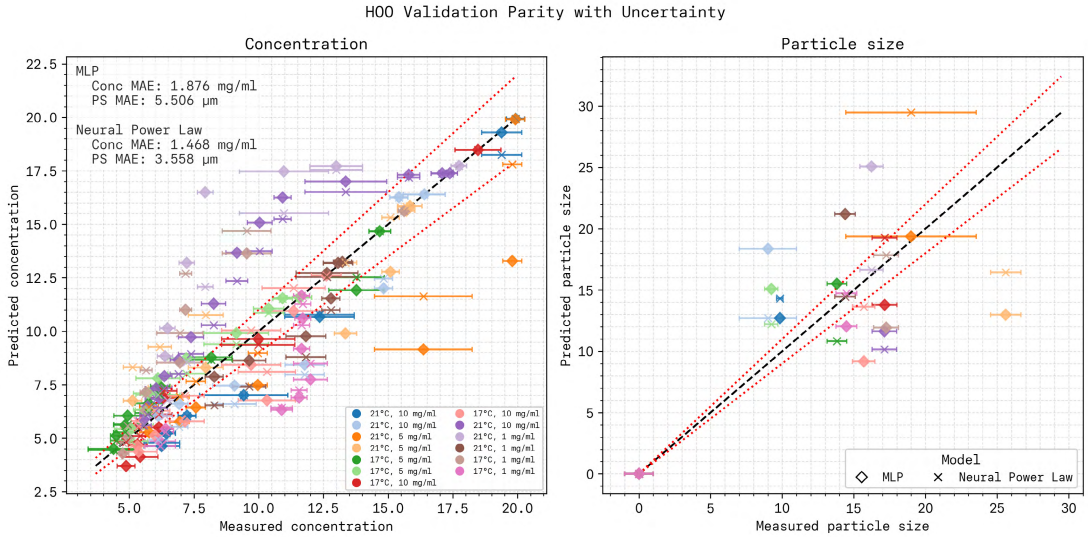

Figure 3: Hold-one-out Validation parity plot for the Carboxyl template system.

The HOO parity plots are shown in Figure 2 (hydroxyl- template) and Figure 3 (carboxyl- template), and error metrics are reported in Table 2. Across both datasets (hydroxyl- and carboxyl- functionalised templates) and both hybrid architectures, we observe consistent trends. First, the overall performance in the HOO benchmark is similar between the two models. Both regressors typically recover the main features of the concentration trajectories and the end-point particle size metric with comparable mean absolute errors. In both hydroxyl and carboxyl datasets, most experiments are predicted reasonably well, but there are 2–3 validation experiments where the prediction is noticeably poorer. This pattern suggests that a small subset of experiments sit in a more difficult region of the design space (e.g., sharper nucleation, stronger nonlinearity, or less informative measurements), and these become “failure modes” for both hybrid parametrisations.

Second, even though the two regressor setups are structurally different, the predicted validation trajectories are often very similar, especially for concentration. This is quite striking: despite one model predicting rates directly as a function of state (via the MLP) and the other enforcing a more structured supersaturation dependence (via the neural power-law surrogate), both end up producing closely aligned concentration dynamics across most left-out cases. One plausible interpretation is that, within the range of operating conditions covered by the training grid, both models are converging to essentially the same effective kinetic relationship. The neural power-law explicitly constrains the supersaturation dependence to a sum of power laws, while the MLP may be

learning a function that is effectively equivalent over the observed supersaturation range. In that sense, the two architectures may be different parametrisations of the same learned dynamics, rather than fundamentally different learned behaviours.

In contrast, particle size predictions are generally less accurate than concentration predictions in the HOO benchmark, as reflected by higher relative errors. This is not unexpected given the experimental data structure: each experiment contributes a full concentration time series (multiple points) but only a single average particle size value is available at batch termination. The loss therefore contains substantially more information about concentration than about size, and the hybrid model is asked to infer size dynamics largely through the mechanistic coupling rather than through direct dense measurements. This imbalance renders the particle size objective harder to learn and easier to over/under-shoot, particularly for the more challenging left-out experiments. For carboxyl templates specifically, the neural power-law surrogate tends to yield slightly better particle size errors, which suggests that the additional structural bias in the kinetics help stabilise the size predictions when size data are sparse.

Table 3: training and variational inference times for the 2 models.

Model	Avg. HOO training time [s]	Variational Inference time [s]
MLP	62 ± 8.5	1498
Neural P-L	51 ± 12	910

By calling the network only once per simulation, rather than once per ODE call, the computational graph is simplified leading to improved memory usage and faster training. Secondly, this simplifies the kinetic identification problem by removing the time dependence of the inference and training problem. The hybrid model training and uncertainty quantification runtimes are reported in Table 3. The neural power-law regressor also delivers a ~40% speed-up under the same variational inference settings, which is consistent with its reduced computational graph (the regressor being called once per simulation rather than at every ODE evaluation), and is a relevant advantage when scaling these methods to larger experimental grids, or more complex (like multidimensional) models.

Overall, the HOO results indicate that both hybrid PBMs learn the core crystallisation dynamics reasonably well, with concentration being the easier and more consistently captured output. The remaining gaps, most clearly in particle size and in a small number of difficult experiments, are likely dominated by the combination of sparse size supervision and the intrinsic ambiguity of mapping limited measurements to effective nucleation/growth kinetics in template-induced systems.

Off-grid Validation

Table 4: Off-grid validation mean absolute error for hydroxyl and carboxyl systems, MLP and Neural P-L models.

Model	Conc. MAE [mg/ml]	PS MAE [micron]
	Hydroxyl System	
MLP	1.56	2.48
Neural P-L	0.75	1.5
	Carboxyl System	
MLP	0.73	0.26
Neural P-L	0.64	3.24

Next, we examine the interpolative validation experiments, with predictions shown in Figure 4 (hydroxyl-template) and Figure 5 (carboxyl- template), and error metrics are reported in Table 4. These validation runs were carried out at 19 °C and 7.5 mg/ml template loading, i.e. conditions chosen halfway between the training temperatures and loadings. Unlike the HOO benchmark, where each validation experiment still shared the same discrete temperature and loading values with some of the training data, these validation conditions are unique and are not present anywhere in the training set. This means the models are genuinely required to interpolate jointly in temperature and template loading, rather than only interpolating in initial condition or trajectory shape.

Across both template chemistries, the two hybrid architectures recover the overall dynamics, but with consistent biases. In both systems, the MLP tends to underpredict the concentration during the batch. The resulting trajectories often follow the correct qualitative shape, but are shifted downward relative to the measurements, suggesting that the learned effective kinetics lead to overly fast depletion (or equivalently an overestimation of net crystallisation rate) at the interpolated condition. By comparison, the neural power-law surrogate is generally closer to the measured concentration trajectories, and for the hydroxyl templates (Figure 4), it gives an almost perfect fit over the entire horizon.

For the carboxyl template system (Figure 5), the two models produce broadly similar concentration trajectories, and the concentration MAE is comparable between them despite the MLP underprediction. The difference is more visible in the particle size metric: the neural power-law surrogate slightly underpredicts the particle size, while the MLP predicts particle size more accurately, in some cases essentially matching the measurement within experimental error bounds. This is consistent with the trade-off observed in other parts of the evaluation: the more flexible direct-rate MLP can sharpen the size prediction when it lands in the right region of rate space,

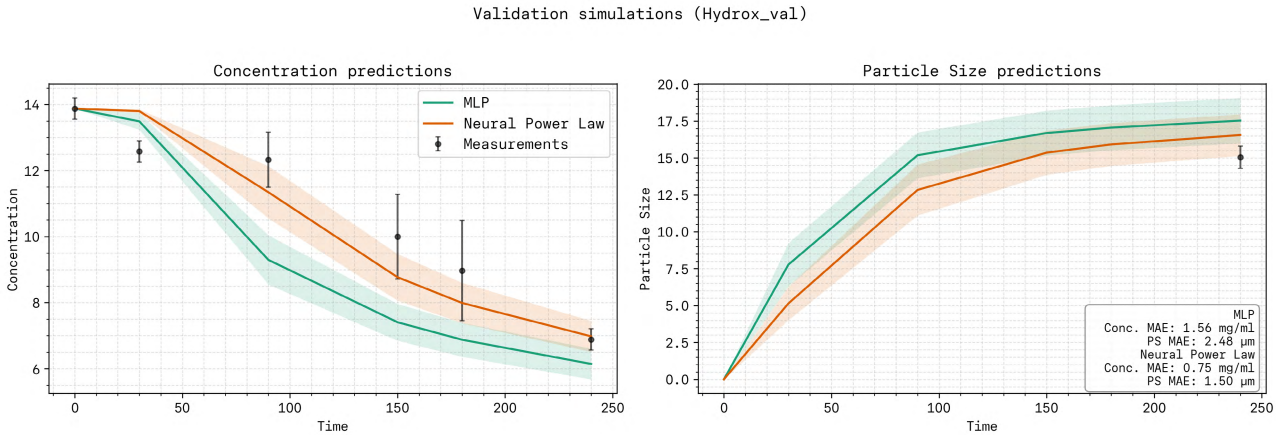

Figure 4. Off-grid validation prediction plots for the Hydroxyl template system at T = 19°C, Loading = 7.5 mg/ml

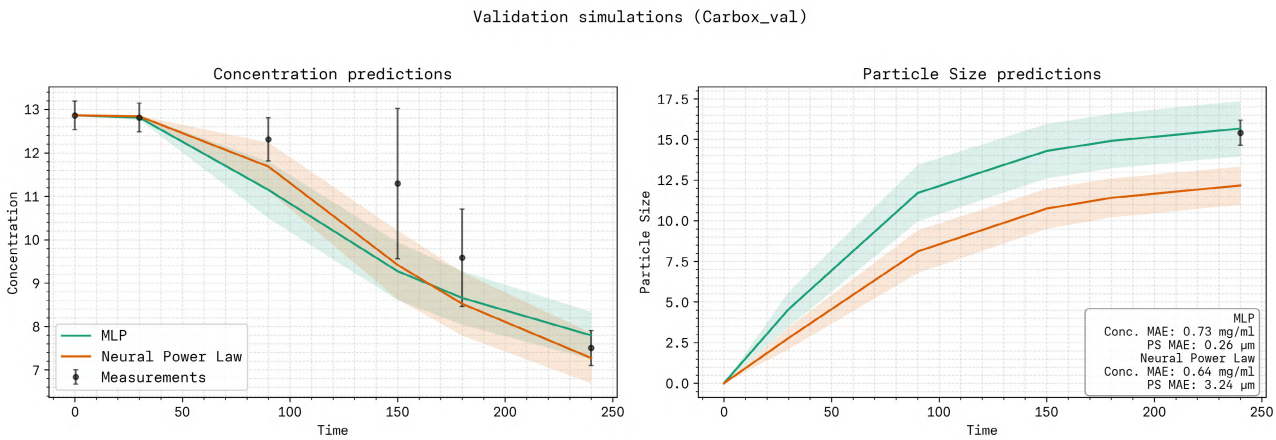

Figure 5. Off-grid validation prediction plots for the Carboxyl template system at T = 19°C, Loading = 7.5 mg/ml

whereas the more structured surrogate can prioritise capturing the concentration evolution, sometimes at the cost of a small bias in the final size metric.

For hydroxyl templates, the separation between the models is clearer. The neural power-law surrogate matches the validation concentration trajectory extremely well, while the MLP again underpredicts concentration across the batch. For particle size, both models slightly overpredict, but the neural power-law surrogate has the better overall outcome when considering both the mean prediction and the uncertainty intervals. Importantly, even where the point predictions differ, both models' confidence intervals appear reasonably calibrated: the predictive bands are typically consistent with the observed experimental scatter and remain within the reported measurement error. This is particularly relevant here because the validation condition is unseen, so interval calibration is a stronger test than simply matching a single mean trajectory.

Taken together, the interpolative validation suggests that both hybrid PBMs can generalise to new temperature–loading conditions without re-fitting, but the two kinetic surrogates exhibit distinct failure modes. The MLP appears prone to a systematic concentration underprediction at the interpolated condition, while the neural power-law surrogate is more robust in concentration interpolation and tends to produce more conservative uncertainty. Meanwhile, the MLP can provide sharper particle size predictions when the learned kinetics align well with the true size dynamics at that condition. This highlights that model choice should not rely purely about average error: for template-enabled process design, the preferable surrogate will depend on whether the primary objective is accurate concentration trajectory prediction (e.g. yield/endpoint concentration targets) or more accurate particle size outcomes (e.g. downstream filtration and formulation constraints), and on whether conservative uncertainty quantification is preferred for risk-aware decision making.

Beyond the specific comparison of the two architectures, the results presented here raise several broader open questions for hybrid PBM development. First, the data requirements of these models remain to be systematically assessed. The experimental grid used in this work is deliberately compact, and while both regressors demonstrate reasonable generalisation, it is not yet clear

how much additional experimental coverage would be needed to meaningfully improve the learned kinetics, nor at what point diminishing returns set in. Closely related is the question of data quality: in protein crystallisation, measurements are subject to heteroscedastic noise, batch-to-batch variability, and limited temporal resolution, all of which propagate through the training loss. The extent to which measurement quality, as opposed to sheer quantity, limits the fidelity of the trained regressor is an important practical consideration that is not addressed here. Finally, the hybrid models as presented remain black-box at the kinetic level: the learned regressors capture effective nucleation and growth rates, but these rates are not readily interpretable or transferable in the way that a compact, explicit kinetic expression would be. Distilling the trained surrogate into a human-readable form, for example via symbolic regression on the learned regressor outputs, would allow practitioners to extract mechanistic insight from the flexible model, compare it against established kinetic forms, and stress-test the resulting expression outside the data-supported region. These open questions, data sufficiency, data quality, training robustness, and interpretability, represent natural directions for further development of hybrid PBMs for template-induced crystallisation.

CONCLUSIONS

This work proposes a workflow for the development and experimental validation of uncertainty-aware hybrid population balance models for template-induced protein crystallisation.

Across the hold-one-out benchmark, the two hybrid architectures present very similar effective dynamics, despite making different modelling choices about how kinetics are captured. Improving the HOO performance likely requires either more experimental coverage, since both regressors still have capacity and will generally benefit from additional temperature–loading combinations and denser particle-size information, or strategies that effectively densify the dataset without simply overfitting. One plausible direction is data augmentation through mechanistic fits on subsets of experiments to generate richer trajectories; however, this must be approached carefully, because any model–data mismatch in the mechanistic fits would be inherited by the augmented dataset and may ultimately degrade calibration.

The off-grid validation is arguably the more demanding test: here the models interpolate in temperature and loading, rather than relying on having seen the same discrete operating values during training. Under this stress test, the Neural Power Law structure with its explicit decoupling of supersaturation dependence, tends to deliver more robust concentration predictions and more conservative uncertainty, while the MLP can produce sharper particle-size predictions. Overall, these results support the central conclusion of this study: uncertainty-aware hybrid PBMs can learn useful, template-dependent kinetic relationships even under severe data scarcity and measurement sparsity and can generalise to unseen operating conditions without re-training.

Several directions for future work emerge from these findings. The data requirements of hybrid PBMs, both in terms of quantity and quality, remain to be characterised more rigorously: understanding the minimum experimental coverage needed to train a reliable regressor, and quantifying how measurement noise and temporal resolution affect the learned kinetics, would provide actionable guidance for experimental planning. There remains also an opportunity to explore interpretability further. A key advantage of purely mechanistic PBMs is that the kinetic relationships are compact and human-readable; hybrid models trade some of that transparency for flexibility. One natural next step is therefore to treat the trained kinetic surrogate as a discovery tool and distil it into an explicit functional form, for example by applying symbolic regression to the learned regressor heads, so that the final outcome is both accurate in the data-supported region and easier to communicate, compare, and stress-test for process design.

AKNOWLEDGEMENTS

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) for the Imperial College London Doctoral Training Partnership (DTP) and by AstraZeneca UK Ltd through a CASE studentship award.

AUTHOR IDENTIFIERS

Pessina D: 0009-0004-3107-6857
Heng Jerry Y. Y. : 0000-0003-2659-5500
Papathanasiou Maria M.: 0000-0002-8886-0624

REFERENCES

1. Hekmat D. Large-scale crystallization of proteins for purification and formulation. Bioprocess Biosyst Eng 38:1209-1231 (2015). https://doi.org/10.1007/s00449-015-1374-y
2. Ferreira J, Araújo S, Ferreira A, Teixeira J, de Campos JM, Rocha F, Castro F. Insulin nucleation kinetics in an oscillatory flow-based platform: protein crystallization as a highly reproducible separation process. Chemical Engineering Research and Design 203:453-466 (2024). https://doi.org/10.1016/j.cherd.2024.01.057
3. Saridakis, E. *et al.* Protein crystallization facilitated by molecularly imprinted polymers. *Proc. Natl.*

Acad. Sci. **108**, 11081–11086 (2011).
4. Parambil JV, Poornachary SK, Heng JYY, Tan RBH. Template-induced nucleation for controlling crystal polymorphism: from molecular mechanisms to applications in pharmaceutical processing. CrystEngComm 21:4122-4135 (2019). https://doi.org/10.1039/c9ce00404a
5. Pessina D, Anda JC, Heffernan C, Tian T, Watson O, Heng JY, Papathanasiou MM. Model-based approach to template-induced macromolecule crystallisation. Systems and Control Transactions 4:2504-2509 (2025). https://doi.org/10.69997/sct.131246
6. McCue C, Girard HL, Varanasi KK. Enhancing protein crystal nucleation using in situ templating on bioconjugate-functionalized nanoparticles and machine learning. ACS Appl. Mater. Interfaces 15:12622-12630 (2023). https://doi.org/10.1021/acsami.2c17208
7. Barhate Y, Kang YS, Nazemifard N, Renner B, Yang Y, Papageorgiou C, Nagy ZK. Data-driven digital design of pharmaceutical crystallization processes. Systems and Control Transactions 4:2604-2609 (2025). https://doi.org/10.69997/sct.128994
8. Nielsen RF, Nazemzadeh N, Sillesen LW, Andersson MP, Gernaey KV, Mansouri SS. Hybrid machine learning assisted modelling framework for particle processes. Computers & Chemical Engineering 140:106916 (2020). https://doi.org/10.1016/j.compchemeng.2020.106916
9. McKay A, Ghosh D, Zhu L, Xi L, Mhaskar P. A novel linear hybrid model predictive control design: application to a fed batch crystallization process. Digital Chemical Engineering 3:100033 (2022). https://doi.org/10.1016/j.dche.2022.100033
10. Fox CW, Roberts SJ. A tutorial on variational bayesian inference. Artif Intell Rev 38:85-95 (2011). https://doi.org/10.1007/s10462-011-9236-8
11. Austin GI, Pe'er I, Korem T. Distributional bias compromises leave-one-out cross-validation. Sci. Adv. 11: (2025). https://doi.org/10.1126/sciadv.adx6976

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Capturing mixing effects on aggregation kinetics of monoclonal antibodies during viral inactivation

T. Marella[a], F. Cenci[b], P. Thompson[c], M. Muhieddine[b], F. Bezzo[a]*

[a] CAPE-Lab – Computer-Aided Process Engineering Laboratory. Department of Industrial Engineering. University of Padova, Padova, Italy.
[b] GSK Ware Research and Development, Ware, UK.
[c] GSK R&D, Biopharmaceutical Drug Substance Development (BDSD), King of Prussia, USA.
* Corresponding Author: fabrizio.bezzo@unipd.it

ABSTRACT

Mathematical models play a central role in biopharmaceutical manufacturing, especially within the Quality by Design framework. For these models to be effectively used in optimization tasks, they must be both reliable and capable of delivering results in an affordable computational time. This work proposes a strategy to model aggregate formation during viral inactivation in the context of monoclonal antibody downstream processing. These units often display mixing-sensitive behavior because aggregation kinetics is controlled by local pH, whose spatial heterogeneities arise from titrant addition at a defined feed point. To address this challenge, compartment models (CMs) are employed. This modeling approach captures spatial inhomogeneities within the unit by leveraging flow-exchange information derived from a single steady-state Computational Fluid Dynamics (CFD) simulation involving only the solution of mass, momentum and turbulence equations. Results obtained by comparing compartment models with both perfectly mixed models and full CFD simulations including aggregation kinetics demonstrate that CMs can reproduce the CFD results with good approximation, while reducing computational time by orders of magnitude.

Keywords: Compartment Models, Computational Fluid Dynamics, Monoclonal Antibodies, Downstream Bioprocessing

INTRODUCTION

Mathematical modeling has become a cornerstone in the modernization of pharmaceutical manufacturing, enabling process optimization, quality-by-design strategies, and the development of digital twins [1]. To effectively serve this role, models must ensure robustness and reproducibility under real manufacturing conditions, especially because this level of reliability is increasingly required by regulatory frameworks.

Beyond supporting unit-operation design, mathematical models are equally essential at the system level, where they allow the integrated optimization of process configurations, resource allocation, and operational strategies.

Achieving this goal calls for an accurate description of the relevant phenomena across multiple scales, from process-level behavior down to local fluid dynamics [2]. This aspect is particularly critical in unit operations where fluid flow and mixing have a decisive impact on process outcomes.

In the context of biopharmaceutical manufacturing for monoclonal antibodies (mAb) production, one of the most critical operations where mixing plays a central role is the viral inactivation step. This step takes place in the downstream section of the process, during which purification and product quality refinement are carried out. This operation is designed to ensure product safety by lowering the solution pH through the addition of an acidic titrant and holding this condition for a minimum amount of time, thereby inactivating acid-labile viruses. However, low pH conditions simultaneously trigger aggregation kinetics for certain monoclonal antibodies [3]. Aggregation is a key risk, as it may impact product quality and reduce yield, and its kinetics depend strongly on the local pH distribution within the vessel. Therefore, spatial heterogeneities in mixing directly translate into variations in aggregation rates, linking hydrodynamics with critical

quality attributes, especially for proteins that are prone to low-pH aggregation.

A model of the unit in which the viral inactivation step is performed must therefore be able to capture the influence of imperfect mixing on both pH distribution and aggregation kinetics in order to provide reliable and accurate predictions. Additionally, the application of a unit model inside process-level optimization tasks requires the capability of this model to provide predictions within a significantly reduced computational time. This combination of physical fidelity and computational efficiency is essential to enable model-based decision-making in biopharmaceutical development and manufacturing.

Due to their simplicity and fast computational times, models that assume perfectly mixed conditions inside the vessel are commonly used in process-level optimization tasks. However, these models fail to represent the effect of fluid dynamics on the aggregation kinetics, leading to an underestimation of aggregation risks and potentially misleading performance predictions. On the other hand, when a detailed hydrodynamic description is required, computational fluid dynamics (CFD) models that solve both the flow behavior and the transport of species are generally employed. While this kind of approach can be highly accurate, it requires substantially higher computational time and resources, which limits its applicability for system-level optimization and for extensive exploration of complex physical phenomena [4].

To overcome the challenge of developing hydrodynamics-oriented yet time-efficient models, a multi-scale approach should be adopted for the modeling of the vessel. One promising way to implement such a multi-scale strategy is through compartment models (CMs). In this modeling approach, the vessel is discretized into a network of interconnected compartments, each assumed to be ideally mixed [5]. The exchange of mass between compartments is not modeled directly by solving the full fluid-flow equations, but is instead derived from a single CFD simulation that solves only the flow and turbulence equations [6]. This makes it possible to capture the large-scale hydrodynamics at a reduced computational cost, while retaining spatial heterogeneities and flow non-idealities within the equipment. Compartment models therefore provide a powerful intermediate between overly simplified ideal-mixing assumptions and the full resolution of CFD simulations [7][8][9]. This makes them particularly useful for describing the pH distribution inside the unit and, in turn, its effect on aggregation kinetics at a significantly reduced computational cost.

This work aims to examine the applicability of compartment modeling to the viral inactivation step in mAb production processes, with the goal of developing a model that incorporates local pH gradients to improve predictive accuracy compared to perfectly mixed models, while providing results in a substantially shorter time than a full CFD simulation. Such a model can be readily used for process-level optimization or for evaluating alternative operating conditions that do not affect the underlying fluid dynamics.

MODELING WORKFLOW

The modeling workflow adopted in this study is described in this section.

The development of the compartment model follows a sequential procedure. First, a CFD simulation resolving only the fluid flow (momentum, continuity, and turbulence equations) is performed. The process geometry is then discretized into compartments, after which the mass-flow exchanges between compartments are computed from the CFD results. Finally, the corresponding material balances are solved. The following sections detail each step of this workflow.

Fluid flow only CFD simulation

A steady-state CFD simulation resolving the fluid flow (momentum, continuity, and turbulence equations) serves as the starting point for the development of the compartment model.

Two different vessel geometries representative of equipment used during industrial viral inactivation at different scales are here considered. Comparing these two configurations allows to investigate how vessel geometry and characteristic scales influence mixing patterns, pH distribution, and consequently aggregation behavior.

The first geometry, referred to as the *R&D scale* vessel, consists of a cylindrical unbaffled tank with a conical bottom and a maximum working volume of 200 L, operated here with 90L of liquid. The nominal impeller speed is set to 230 RPM. The second geometry, referred to as the *Clinical scale* vessel, features a square cross-sectional area and a maximum working volume of 1000 L, and it is operated with 550L of liquid. The nominal impeller speed is set to 180 RPM.

An unstructured polyhedral mesh was selected for the simulations of both geometries. The meshes were generated using Ansys® Fluent, release 25.1 watertight geometry meshing workflow. The resulting meshes comprise 377348 computational cells for the *R&D scale* vessel and 1261684 cells for the *Clinical scale* one. CFD simulations were run using Ansys® Fluent, release 25.1.

The fluid inside the vessel comes the prior purification step. It contains the product, process impurities, and buffer components. The density is 1007 kg/m^3, while the viscosity is 10^{-3} Pa·s.

To model turbulence, the SST k-ω model is selected, retaining the default parameters. To model the impeller rotation, the Multiple Reference Frame approach is used; the rotating reference frame is shown in Figure 1 for the *R&D scale* vessel and in Figure 2 for the *Clinical*

scale one.

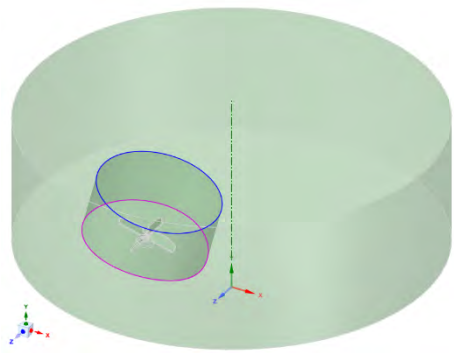

Figure 1. 3D geometry of the *R&D scale* vessel.

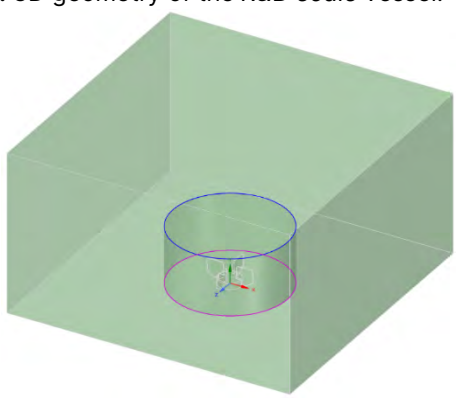

Figure 2. 3D geometry of the *Clinical scale* vessel.

The simulations are performed under steady-state conditions and the resulting fluid-flow behavior was validated against experimental operating data from the vessel, using blend time as the metric for comparison.

Key information required to build the compartment model (e.g. cell centroid coordinates, volume, velocity components, data concerning neighbouring cells) are exported in text files using User Defined Functions.

Compartment model

Starting from the results of the steady-state CFD simulation describing the detailed fluid flow behavior, the compartment model is built through a sequence of steps: first, the geometry is divided into compartments; next, the fluxes between compartments are computed; and finally, the material balances within each compartment are solved. All implementations were carried out in a Python 3.13 environment.

Compartmentalization

The compartmentalization procedure is the preliminary step in constructing the compartment model. Its purpose is to divide the system geometry into compartments, each representing a volume in which the perfectly mixed assumption is applied. Adjacent compartments are then allowed to exchange mass with one another.

In this study, a basic compartmentalization procedure is employed. More complex approaches, e.g. accounting for the flow patterns inside the tank, can be found in the literature [10][11][12].

The selected compartmentalization procedure consists of defining the number of compartments to be generated along the different coordinate axes. In this way, the system is compartmentalized so that all compartments have approximately the same volume. Due to the distinct geometries of the two vessels, a cylindrical coordinate system (r, θ, y) was used for the *R&D scale* tank whereas a cartesian coordinate system (x, y, z) was adopted for the *Clinical scale* tank. An example of compartment division with 4 compartments along each coordinate is visualized in Figure 3 using cylindrical coordinates for the *R&D scale* vessel, in Figure 4 using cartesian coordinates for the *Clinical scale* one.

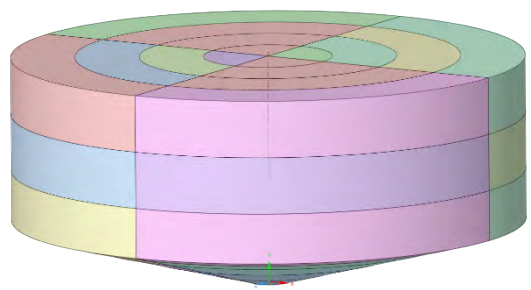

Figure 3. Example of compartmentalization with four compartments along the different coordinates for the *R&D scale* vessel using cylindrical coordinates.

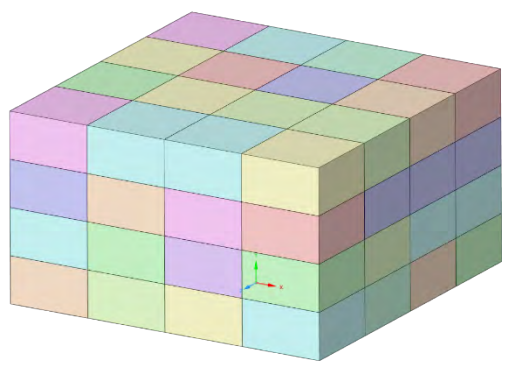

Figure 4. Example of compartmentalization with four compartments along the different coordinates for the *Clinical scale* tank using cartesian coordinates.

After this subdivision, the CFD mesh cells are grouped based on whether their centroids fall within the coordinate bounds of each compartment [13]. Consequently, at the end of this step, a mapping is obtained that associates each CFD mesh cell with its corresponding compartment.

Fluxes computation

To compute the material exchanges between compartments, the face flux data exported from the CFD

simulation are used. For each mesh face *k*, the CFD solver provides a flux value f_k.

Faces that connect cells belonging to two different compartments are identified by examining the compartments associated with their neighboring CFD cells. Let compartment *i* contain the set of CFD mesh cells $\mathcal{C}_i$, and compartment j the set $\mathcal{C}_j$. A face *k* contributes to the exchange between compartments *i* and *j* if it connects a cell $c_a \in \mathcal{C}_i$ and a cell $c_b \in \mathcal{C}_j$. Based on this identification, the set of faces contributing to the flux from compartment *i* to compartment *j* can be defined as:

$$\mathcal{F}_{i\to j} = \{k \mid k \text{ connects a cell in } \mathcal{C}_i \text{ to a cell in } \mathcal{C}_j\} \quad (1)$$

A schematic illustrating how cells and faces are assigned to each set is shown in Figure 5.

The total flux from compartment *i* to compartment *j* is then computed as the sum of the corresponding face fluxes:

$$F_{i\to j} = \sum_{k\in\mathcal{F}_{i\to j}} f_k \quad (2)$$

Consequently, the overall outgoing and incoming fluxes for compartment *i* are defined as:

$$F_i^{out} = \sum_{j=1}^{N} F_{i\to j} \quad (3)$$

$$F_i^{in} = \sum_{j=1}^{N} F_{j\to i} \quad (4)$$

where *N* is the total number of compartments.

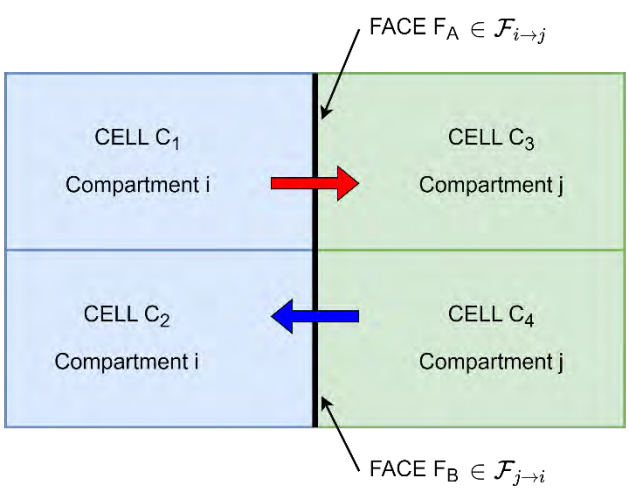

Figure 5. Schematics of assignments of cells and faces. $C_1, C_2 \in \mathcal{C}_i, C_3, C_4 \in \mathcal{C}_j$. Face F_A divides cells $C_1 \in \mathcal{C}_i$ and $C_3 \in \mathcal{C}_j$, and is assigned to set $\mathcal{F}_{i\to j}$ since mass flux is from C_1 to C_3; face F_B divides cells $C_2 \in \mathcal{C}_i$ and $C_4 \in \mathcal{C}_j$, and is assigned to set $\mathcal{F}_{j\to i}$ since mass flux is from C_4 to C_2.

Material balances

The compartment model then solves a system of ordinary differential equations, composed of one equation per species in each compartment. The mass balances are expressed as follows:

$$\frac{dM_{i,c}}{dt} = \sum_{k=1}^{N}(F_{k\to c} * w_{k,i}) - w_{c,i} * F_{c,outgoing} + R \quad (5)$$

where $M_{i,c}$ is the mass of species *i* in compartment *c*, *N* is the total number of compartments, $w_{c,i}$ is the mass fraction of species *i* in compartment *c*, *R* represents reaction or source terms for species *i* in compartment *c*.

The resulting system of material balances is solved using the *solve_ivp* function from the *scipy* library, with the BDF method selected as the numerical solver.

The solution returns the mass profiles of each species within all compartments as a function of time. To enable comparison across different models, the total mass of each species in the tank is then computed at every timestep by summing its mass over all compartments.

All simulation results will be compared to that of a perfectly mixed tank.

CASE STUDY

Some monoclonal antibodies are prone to aggregation under low-pH conditions. A simplified kinetic expression describing this behavior was derived from the work of Arosio et al. (2011) [3]. The simplified aggregation mechanism can be represented as:

$$M + M \xrightarrow{k_{aggr}} A \quad (6)$$

The kinetic constant k_{aggr} is assumed to depend on pH according to the following relationship:

$$k_{aggr} = (k_{aggr}^0) \cdot \exp(-\beta(pH - 3.334)) \quad (7)$$

where k_{aggr}^0 and β are model parameters estimated from experimental data. Their numerical values are reported in Table 1.

Table 1: Aggregation kinetics model parameters.

Parameter	Value
k^0_{aggr}	4.518 [m³/(kg h)]
β	8.468 [-]

The reaction rate can therefore be expressed on a mass basis as:

$$R = k_{aggr} \cdot [M]^2 \quad (8)$$

where *[M]* is the mass concentration of the mAb.

All kinetic expressions were formulated on a mass basis to simplify their implementation in the CFD simulations, since no information on the molecular weight of the mAb was available.

To incorporate the aggregation kinetics into both the CFD simulations and the compartment models, three species must be considered: the monoclonal antibody, referred to as the monomer; the product formed through aggregation, referred to as the aggregate; and the

oxonium ion, which is included to represent the pH distribution inside the tank through its definition:

$$pH = -\log_{10}[H_3O^+] \quad (9)$$

where $[H_3O^+]$ is the molar concentration of oxonium ion expressed in mol/L. Since the total fluid density is assumed to remain constant, an additional species, referred to as the solvent, is considered implicitly. Its mass is not obtained by solving a dedicated balance equation; instead, it is computed at each time step as the difference between the total fluid mass and the sum of the masses of all explicitly modeled species.

The viral inactivation process operates in fed-batch mode, as a titrant is continuously added to the vessel to reduce the pH to the desired target value. In principle, this would result in a time-varying liquid volume inside the tank. However, to simplify the set up of the simulations, the process was modeled under the assumption of constant volume.

In the CFD simulations, the reaction rate is implemented as a source term for monomer or aggregate in all fluid zones using Equation (5). The continuity, momentum, and turbulence equations were not solved during this stage, and the flow variables were fixed to the values obtained from the initial steady-state simulation.

All simulations were carried out in transient mode. The simulations were initialized using typical values reported in experimental studies of aggregation kinetics for pH, monomer concentration, and aggregate concentration. These initial conditions were assumed to be homogeneous throughout the vessel volume. The process was simulated until the pH reached the final target value, corresponding to a decrease of 1.55 relative to the initial condition.

RESULTS

This section presents the main results obtained in assessing the performance of the compartment models, comparing them against both the perfectly mixed models and the CFD simulations. The CFD results are taken as the most representative approximation of real process behavior, as no experimental data were available for direct validation.

For the compartment model, different numbers of compartments were investigated in order to evaluate how the degree of geometric discretization influences the model predictions. The number of compartments selected along each coordinate is reported in Table 2. At the finest discretization level, the total number of compartments differs between the *R&D scale* and *Clinical scale* vessels due to the conical bottom of the *R&D scale* tank, which results in some compartments containing no CFD mesh cells and therefore being excluded from the model.

Table 2: Different degrees of compartmentalization considered.

Number of compartments			Tot. comp.
x/r	y	z/θ	
2	2	2	8
4	4	4	64
8	4	8	256
12	8	12	1152 (Clin.) 1068 (R&D)

Two scenarios were considered in the performance assessment. In the first scenario, a slow titrant addition was modeled: the amount of titrant required to reach the target pH was fed over a 30 minutes interval. The simulation was then run for a total of 35 minutes to allow both the addition to be completed and the pH to homogenize throughout the tank. In the second scenario, the titrant was added over a 0.5 seconds interval, representing an almost instantaneous addition. In this case, the simulations were performed over a 35 seconds time window to ensure that mixing and pH homogenization could be observed across the entire vessel. The second scenario, characterized by a fast titrant addition, was included because it is expected to generate stronger local inhomogeneities within the tank. This, in turn, is anticipated to amplify the differences in the predictions obtained from the different modeling approaches.

Scenario 1: Slow titrant addition

The time profiles of the total monomer mass inside the vessel, normalized with respect to the initial value, as predicted by the different models for the first scenario, are shown in Figure 6 for the *R&D scale* vessel and in Figure 7 for the *Clinical scale* one.

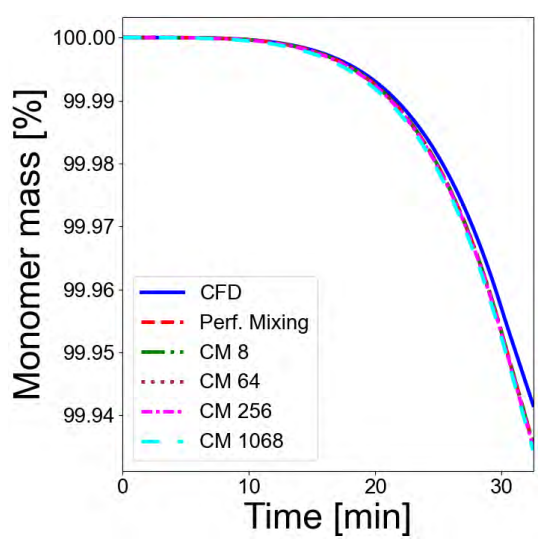

Figure 6. Scenario 1: total monomer mass profiles for the R&D scale vessel.

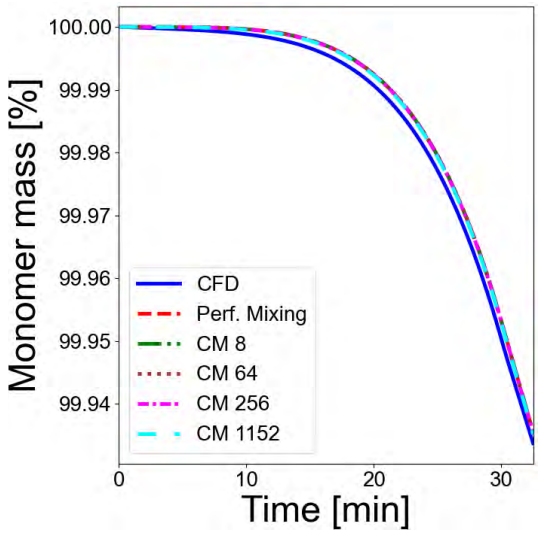

Figure 7. Scenario 1: total monomer mass profiles for the *Clinical scale* vessel.

The profiles obtained with the perfectly mixed model and the compartment models are nearly identical and closely match the CFD results, with discrepancies below 0.01%. This means that for scenario 1 both vessels behave as perfectly mixed systems. Consequently, in this scenario, both CFD and compartment models do not provide a strategic advantage with respect to the perfectly mixed model in predicting the unit behavior.

Scenario 2: Fast titrant addition

The time profiles of the total monomer mass inside the vessel, normalized with respect to the initial value, as predicted by the different models for the second scenario, are shown in Figure 8 for the *R&D scale* vessel and in Figure 9 for the *Clinical scale* one.

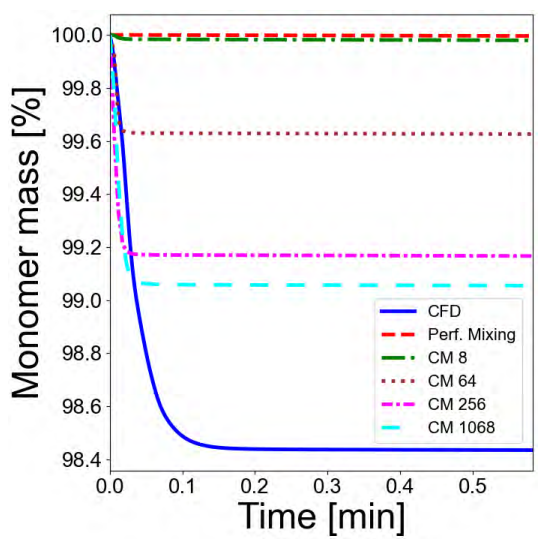

Figure 8. Scenario 2: total monomer mass profiles for the *R&D scale* vessel.

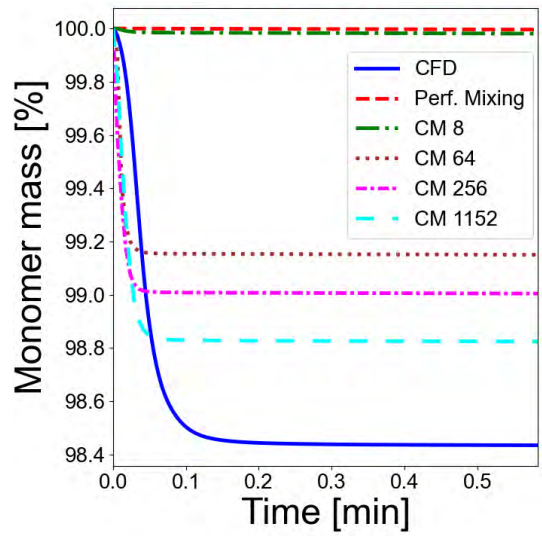

Figure 9. Scenario 2: total monomer mass profiles for the *Clinical scale* vessel.

As expected, this second scenario in which the titrant is added almost instantaneously inside the tank leads to the formation of pronounced localized pH inhomogeneities within the tank. These local variations are inherently captured by the CFD simulations but are not represented by the perfectly mixed model, which assumes instantaneous homogenization throughout the vessel. As a consequence, in this second scenario the CFD simulation predicts a pronounced initial aggregation that is not captured by the perfectly mixed approach. The compartment models, on the other hand, are able to reproduce this early aggregation behavior with increasing accuracy, as their predictions progressively approach the CFD results when a finer level of compartment discretization is employed.

Table 3 reports the computational times for both the R&D scale and the Clinical scale vessels.

Table 3: Computational times for the model solution under the fast titrant addition scenario.

Model	Computational time	
	R&D scale	Clinical scale
CFD	16.75 [h]	75 [h]
Perf. Mixed	0.05 [s]	0.05 [s]
CM 8	0.05 [s]	0.05 [s]
CM 64	0.13 [s]	0.16 [s]
CM 256	1.17 [s]	1.51 [s]
CM 1068 or 1152	39.40 [s]	56.30 [s]

A substantial reduction in computational cost can be observed: while the CFD simulations required days, the compartment models provided the predictions in less than a minute at the highest level of discretization, and in less than two seconds for the coarser discretizations. The reported computational times do not include the

duration of the preliminary steps required for model preparation, as these must be performed only once before the model can be repeatedly solved. The preliminary steady-state CFD simulations, required for both the CFD runs and the compartment models with kinetics included, took approximately 4 hours for the R&D-scale vessel and 8 hours for the clinical-scale vessel. The compartmentalization step took approximately 5 seconds, independent of the discretization level and the vessel scale.

CONCLUSIONS

This study explored the integration of fluid-flow information obtained from CFD simulations into modeling frameworks that require significantly lower computational effort. The proposed procedure was applied to a viral inactivation process within the context of downstream biopharmaceutical manufacturing.

The results demonstrated that compartment models are able to provide predictions that are closer to the CFD results than those obtained with perfectly mixed models, while requiring substantially less computational time. This makes these models suitable for system-level optimization studies and, more generally, for scenarios in which multiple simulations are needed.

These models also offer the capability to rapidly evaluate different process operating conditions, provided that such variations do not affect the underlying fluid-flow behavior.

Opportunities for further improvement lie primarily in the compartmentalization procedure, where more advanced techniques could enhance both model performance and reliability.

ACKNOWLEDGEMENTS

T.M. and F.B. acknowledge GSK R&D for financial support.

DECLARATION OF COMPETING INTEREST

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: F. Cenci, M. Muhieddine and P. Thompson are employees of GSK and may hold share options and/or shares in GSK. The authors declare no other competing conflicts of interest.

AUTHOR IDENTIFIERS

Author ORCIDs:
Tiziana Marella: 0009-0008-3119-7818
Francesca Cenci: 0000-0002-3360-0470
Fabrizio Bezzo: 0000-0003-1561-0584

REFERENCES

1. Destro F, Barolo M. A review on the modernization of pharmaceutical development and manufacturing – trends, perspectives, and the role of mathematical modeling. International Journal of Pharmaceutics 620:121715 (2022). https://doi.org/10.1016/j.ijpharm.2022.121715
2. Jaworski Z, Zakrzewska B. Towards multiscale modelling in product engineering. Computers & Chemical Engineering 35:434-445 (2011). https://doi.org/10.1016/j.compchemeng.2010.05.009
3. Arosio P, Barolo G, Müller-Späth T, Wu H, Morbidelli M. Aggregation stability of a monoclonal antibody during downstream processing. Pharm Res 28:1884-1894 (2011). https://doi.org/10.1007/s11095-011-0416-7
4. Le Nepvou De Carfort J, Pinto T, Krühne U. An automatic method for generation of cfd-based 3D compartment models: towards real-time mixing simulations. Bioengineering 11:169 (2024). https://doi.org/10.3390/bioengineering11020169
5. Jourdan N, Neveux T, Potier O, Kanniche M, Wicks J, Nopens I, Rehman U, Le Moullec Y. Compartmental modelling in chemical engineering: a critical review. Chemical Engineering Science 210:115196 (2019). https://doi.org/10.1016/j.ces.2019.115196
6. Bezzo F, Macchietto S, Pantelides CC. A general methodology for hybrid multizonal/cfd models. Computers & Chemical Engineering 28:501-511 (2004). https://doi.org/10.1016/j.compchemeng.2003.08.004
7. Delafosse A, Collignon ML, Calvo S, Delvigne F, Crine M, Thonart P, Toye D. Cfd-based compartment model for description of mixing in bioreactors. Chemical Engineering Science 106:76-85 (2014). https://doi.org/10.1016/j.ces.2013.11.033
8. Farzan P, Ierapetritou MG. A framework for the development of integrated and computationally feasible models of large-scale mammalian cell bioreactors. Processes 6:82 (2018). https://doi.org/10.3390/pr6070082
9. Maldonado de León H, Straathof A, Haringa C. Dynamic compartment models: towards a rapid modeling approach for fed-batch fermentations. Chemical Engineering Science 308:121396 (2025). https://doi.org/10.1016/j.ces.2025.121396
10. Tajsoleiman T, Spann R, Bach C, Gernaey KV, Huusom JK, Krühne U. A CFD based automatic method for compartment model development. Computers & Chemical Engineering 123:236-245 (2019).

https://doi.org/10.1016/j.compchemeng.2018.12.015
11. Laborda VPI, Puiman L, Groves T, Haringa C, Nielsen LK. Unsupervised learning bioreactor regimes. Computers & Chemical Engineering 194:108891 (2025). https://doi.org/10.1016/j.compchemeng.2024.108891
12. Le Nepvou De Carfort J, Puig I Laborda V, Nielsen LK, Gernaey KV, Krühne U. Flow-informed clustering of bioreactor volumes to build cfd-based compartment models. Chemical Engineering Science 320:122539 (2026). https://doi.org/10.1016/j.ces.2025.122539
13. Haringa C, Tang W, Noorman HJ. Stochastic parcel tracking in an euler–lagrange compartment model for fast simulation of fermentation processes. Biotech & Bioengineering 119:1849-1860 (2022). https://doi.org/10.1002/bit.28094

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Generative AI Approach to Inverse Design for Continuous Pharmaceutical Manufacturing

Consuelo Del Pilar Vega-Zambrano, Vassilis M. Charitopoulos*

Department of Chemical Engineering, The Sargent Centre for Process Systems Engineering, University College London, Torrington Place, London WC1E 7JE, UK
* Corresponding Author: v.charitopoulos@ucl.ac.uk

ABSTRACT

Continuous pharmaceutical manufacturing (CM) offers improved quality assurance, operational agility, and supply resilience, yet process development remains dominated by expensive trial-and-error experimentation and high-dimensional space exploration. Motivated by ICH Q13, we develop a generative inverse-design framework that maps target product quality to feasible process recipes for an integrated twin-screw wet granulation and segmented fluidized-bed drying line. The framework integrates three components: (i) a Conditional Variational Autoencoder (CVAE) generator that proposes process parameter sets conditioned on desired Critical Quality Attributes (CQAs), (ii) a Gaussian Process (GP) surrogate validator that screens candidates for manufacturing feasibility, and (iii) SHapley Additive exPlanations (SHAP) to interpret the generated designs. Training data were produced from a validated gPROMS digital twin of the Diamond Pilot Plant (DiPP) ConsiGma-25 line, covering liquid -to-solid ratio, drying temperature, drying time and air flowrate, with CQAs including granule moisture content, average particle size and porosity. The trained CVAE generated ~50, 000 candidate recipes and learned constrained feasible regions of the design space. Across seven operating scenarios, generated recipes achieved low deviations from targets. For a target-center case, deviations were 2.0% (moisture), 0.4% (particle size) and 0.5% (porosity). Edge cases remained acceptable, with the largest deviation observed for moisture in a high-moisture scenario (8.8%). SHAP analysis highlighted drying time and liquid-to-solid ratio as the dominant drivers of moisture, while liquid-to-solid ratio governed particle size and porosity. Overall, the approach enables rapid, explainable exploration of CM design spaces, reducing experimental burden and supporting QbD-aligned development for integrated continuous processes.

Keywords: Design space, Generative Artificial Intelligence, Pharmaceutical manufacturing, Inverse Design, Quality by Digital Design, Conditional Variational Autoencoder

INTRODUCTION

Continuous manufacturing (CM) is increasingly adopted in the pharmaceutical sector to improve product quality, manufacturing agility and supply resilience. Unlike traditional batch processing, CM operates closer to a steady state and can enable tighter coupling between process analytical technology, advanced process control and real-time release testing. This evolution is supported by the broader shift toward digital applications (online/real-time, model-based decision support and control) and by the growing role of AI across the pharmaceutical lifecycle [1-4]

Regulatory agencies have actively encouraged CM adoption; in particular, ICH Q13 consolidates scientific and regulatory considerations for CM development, implementation, operation and lifecycle management, and emphasizes systematic approaches to define and maintain a state of control [5] Within the Quality-by-Design (QbD) paradigm (e.g., ICH Q8), the central engineering task is to connect Critical Process Parameters (CPPs) to Critical Quality Attributes (CQAs) and to identify a design space (DS) that guarantee product quality [6]

Despite these advantages, CM process development and scale-up remain resource intensive. Integrated drug-product flowsheets frequently involve several unit operations. As a result, CPP-CQA relationships are highly nonlinear and can exhibit strong cross-unit interactions.

A representative example is the integrated twin-screw wet granulation and segmented fluidized-bed drying line used in continuous tableting: granulation conditions govern nucleation, growth and consolidation through liquid bridge formation and shear, while drying conditions determine moisture removal and influence granule structure and porosity [7]. Because these mechanisms interact across units, traditional design-of-experiments (DoE) campaigns can become prohibitively expensive, requiring many experiments (or high-fidelity simulations) to explore multidimensional spaces and capture cross-unit effects.

Inverse design reframes this workflow by treating product specifications as the starting point. Instead of asking "what CQAs result from these CPPs?", the inverse problem asks "which CPPs can achieve a target set of CQAs?". This aligns naturally with QbD because target CQAs and acceptable ranges define the desired product profile, and the process recipe should be selected from within a feasible and validated design space. However, inverse design in CM is challenging: several distinct CPP combinations can yield similar CQAs (a one-to-many mapping), feasible regions can be non-convex, and manufacturing constraints (equipment limits, residence-time windows, safe operating envelopes) are difficult to encode explicitly.

Deep generative models offer an alternative route to solve inverse problems. Variational autoencoders (VAEs) learn probabilistic latent representations enabling stable training and sampling of new candidates through latent-space exploration [8].Conditional variants (CVAEs) extend VAEs by learning conditional distributions given labels or targets, providing direct controllability for inverse design [9]. In materials discovery, deep generative models have been widely explored to accelerate inverse design by encoding structure-property relationships into latent spaces that can be navigated to propose candidates with targeted characteristics [10]. In pharmaceutical sciences, generative models have achieved notable success in drug discovery and are now being explored for digital formulation optimization[11-12]. Related work in particle engineering has shown that CVAEs can generate particle geometries to match target sphericity and packing fraction, illustrating the suitability of conditional generation for constrained inverse tasks [13]

Despite this progress, the application of generative inverse design to integrated, industrial-style CM unit-operation networks remain limited. Existing CM studies often focus on data augmentation. For example, Ma et al. combined GAN-based data augmentation with surrogate models and multi-objective optimization for an integrated synthesis-crystallization case study [14]. In GMP-relevant CM, however, robustness, feasibility screening and interpretability are essential: candidate recipes must remain within validated operating envelopes and the rationale for recipe selection should be explainable to support process understanding and lifecycle management.

In this work, we propose a three-component framework tailored to CM inverse design. First, a CVAE acts as a generator that learns the conditional distribution of CPP recipes given target CQAs, enabling controllable, one-to-many generation. Second, an independent Gaussian Process (GP) surrogate acts as a validator to screen generated candidates for feasibility and to quantify predictive uncertainty; this reduces the risk of selecting recipes that exploit model artefacts or extrapolate beyond the validated region. Third, we apply SHapley Additive exPlanations (SHAP) [15] to the surrogate validator (GP model) that maps CPPs to CQAs. The resulting feature attributions explain how changes in CPPs drive the validator's predicted CQAs and can be used to interpret candidate CPP sets generated by the CVAE.

We demonstrate the framework using a validated gPROMS flowsheet model representative of the Diamond Pilot Plant (DiPP) ConsiGma-25 continuous tableting line, integrating twin-screw wet granulation with segmented fluidized-bed drying [7-16] .The case study covers four CPPs (liquid-to-solid ratio, drying temperature, drying time and air flowrate) and three CQAs (granule moisture content, average particle size and porosity). Results show that the CVAE can generate large sets of candidate recipes that cluster around desired CQAs while maintaining operational diversity, and that feasibility and interpretability layers provide robust screening and actionable process insight. Together, these capabilities support data-efficient CM development aligned with ICH Q13 and QbD principles.

METHODOLOGY

Figure 1 summarizes the Conditional Variational Autoencoder (CVAE) architecture adopted for inverse design. The encoder maps process parameters together with target CQAs into a low-dimensional latent representation, and the decoder reconstructs process parameters conditioned on the target CQAs. Sampling in latent space enables generation of multiple diverse recipes for the same target product profile.

Data generation and preprocessing

Training data were generated using a validated gPROMS FormulatedProducts flowsheet model of the integrated twin-screw wet granulation and segmented fluidized-bed drying line [7-16]. A space-filling Sobol sequence was used to sample the CPP design space efficiently and to avoid clustering typical of purely random sampling [17]. A dataset of 300 formulations was used with 210 for training and 90 for validation. The same dataset was used to train the GP and CVAE models. Training on the same data as the CVAE ensures consistency in the

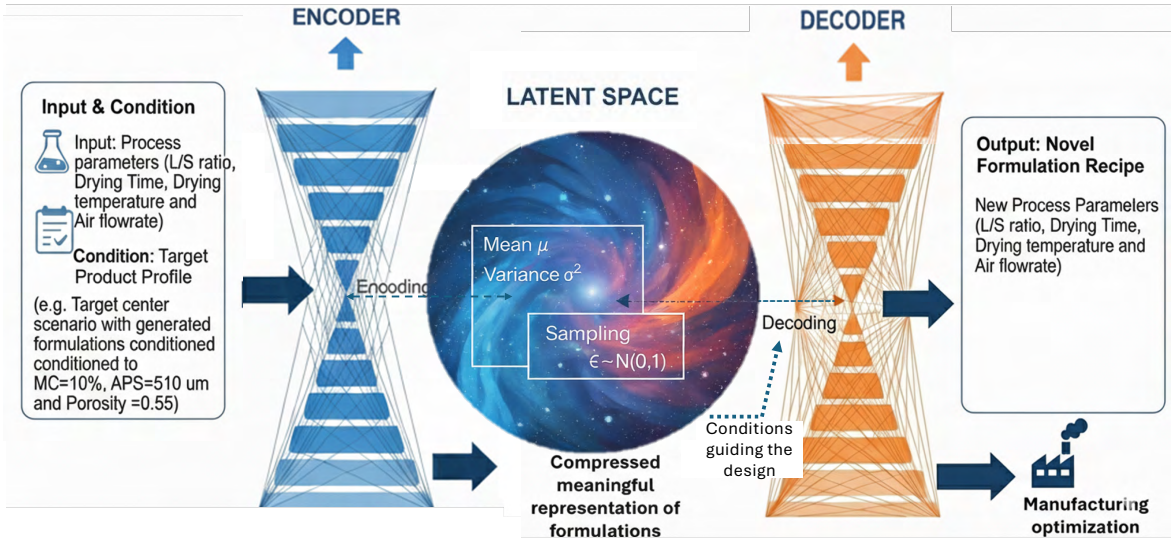

Figure 1. Conditional Variational Autoencoder (CVAE) architecture adapted for inverse design.

modeling framework and leverages the full available training information for robust surrogate predictions. For the GP model, the validation dataset was used to evaluate the GP's generalization performance on unseen data.

Inputs and outputs were standardized prior to training; CQAs were scaled to facilitate conditioning and multi-output learning.

CVAE generator for inverse design

We formulate inverse design as learning the conditional distribution $p(\boldsymbol{u}|\boldsymbol{y})$, where $\boldsymbol{u}$ denotes CPPs and $\boldsymbol{y}$ denotes target CQAs. The CVAE is trained by minimizing a composite loss function (Eq. 1) consisting of an MSE reconstruction term and a Kullback-Leibler divergence regularizer [8], [9]

$$L_{CVAE}(\theta,\ \varphi;\ x,\ y) = L_{recon}(x,\ \hat{x}) + \beta(D_{KL}(q_{\varphi}(z \mid x,\ y) \mid\mid p(z))) \quad (1)$$

Where:

- $\boldsymbol{x}$: vector of critical process parameters (CPPs), scaled to the range [0, 1].
- $\hat{\boldsymbol{x}}$: reconstructed vector of critical process parameters (CPPs).
- $L_{recon}(\boldsymbol{x}, \hat{\boldsymbol{x}})$: reconstruction loss that enforces similarity between the original CPPs and the reconstructed CPPs. In this work, the mean squared error (MSE) is used.
- $\boldsymbol{y}$: vector of conditioning critical quality attributes (CQAs), representing the desired product quality targets.
- $q_{\varphi}(\mathbf{z} \mid \mathbf{x}, \mathbf{y})$: approximate posterior distribution learned by the encoder network. It is assumed to follow a multivariate Gaussian distribution with diagonal covariance
- p(**z**): prior distribution over the latent variables, defined as a standard normal distribution N(0, 1).
- β: weighting coefficient controlling the trade-off between reconstruction accuracy and latent space smoothness/continuity of the latent space.

Latent space and regularisation

The CVAE introduces a low-dimensional latent variable z that captures the unobserved factors and one-to-many variability in feasible CPP recipes that can produce the same target CQAs. This is important for inverse design because multiple CPP combinations can satisfy a single product specification; sampling different z values allows the model to generate diverse candidate recipes for the same y, rather than collapsing to one deterministic solution.

To make latent-space sampling meaningful, the latent distribution must be well-structured. The Kullback–Leibler (KL) divergence term regularises the encoder's approximate posterior $q\varphi(z|x,\ y)$ toward a simple prior $p(z)$=N(0, 1). This discourages the encoder from memorising training samples (overfitting) and helps avoid irregular latent regions ("holes") where sampling would decode to unrealistic CPPs. As a result, the decoder can reliably generate feasible recipes by sampling $z \sim p(z)$ and conditioning on the desired CQAs y.

A key practical requirement in CM is recipe diversity under constraints. By sampling the latent space, the CVAE generates a distribution of recipes rather than a single optimum. This flexibility can support operational preferences, robustness to disturbances and equipment

limitations while still meeting target CQAs.

GP validator for feasibility assessment

To screen AI-generated recipes, we train a Gaussian Process (GP) surrogate model that predicts CQAs from CPPs and provides predictive uncertainty. Candidate recipes are accepted when predicted CQAs fall within specification windows and uncertainty remains below a threshold. This validator step reduces the risk of selecting recipes outside the validated region and provides an additional safety layer for GMP-relevant deployment.

Explainability with SHAP

To support interpretability, we apply SHapley Additive exPlanations (SHAP) to the surrogate validator (GP) used to predict CQAs from CPPs. SHAP decomposes each predicted CQA into additive contributions from individual CPPs, providing global trends (feature importance across the design space) and local explanations for specific CPP candidates. In the continuous manufacturing context, this supports "Quality by explainability" [18] by identifying which CPPs most strongly drive each predicted CQA and therefore should be prioritized for monitoring and control.

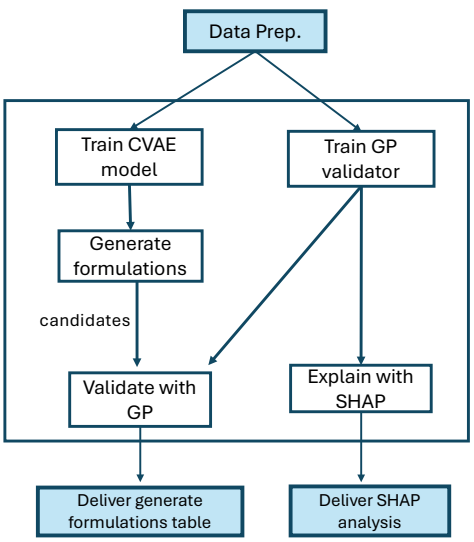

Figure 2. Generative inverse design workflow: CVAE generation.

CASE STUDY: INTEGRATED WET GRANULATION AND DRYING

The case study considers an integrated twin-screw granulator coupled with a segmented fluidized-bed dryer from the ConsiGma-25 continuous tableting line at the Diamond Pilot Plant (DiPP) [7]. The process model, implemented in gPROMS, captures nonlinearities and cross-unit couplings relevant to granule formation and moisture removal.

Four CPPs are considered: liquid-to-solid ratio (L/S), drying temperature (Tem), drying time (DT) and air flowrate (AF). Three CQAs are targeted: granule moisture content (MC), average particle size (APS) and porosity (Poro).

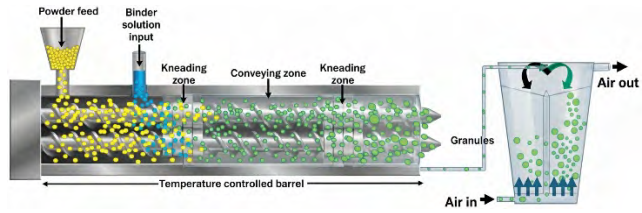

Figure 3. Integrated twin-screw wet granulator and segmented fluidized-bed dryer (ConsiGma-25 line, DiPP).

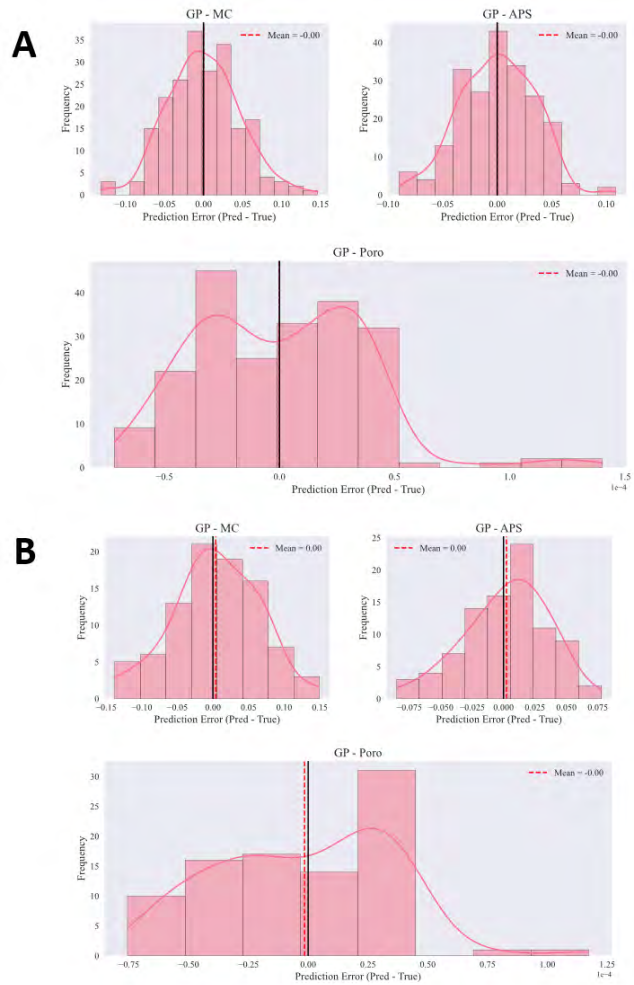

Figure 4. Prediction error distribution for the Gaussian Process validator. A. Training B. Validation

The inverse design task is to generate CPP recipes that meet specified targets (e.g., MC≈10%, APS≈510 µm, Poro≈0.55) while remaining operable for the integrated process.

RESULTS AND DISCUSSION

Validator accuracy

The GP validator achieved high accuracy on training and validation data, with prediction errors tightly centered around zero for all CQAs (Fig. 4).

Explainability and process insight

SHAP analysis (Fig. 5) indicates that drying time (DT) and liquid-to-solid ratio (L/S) exert the strongest influence on moisture content, consistent with heat and mass transfer fundamentals in drying.

In contrast, APS and porosity are primarily governed by L/S, aligning with granulation theory where liquid availability controls nucleation, coalescence and consolidation. These explanations help translate data-driven recipes into actionable engineering understanding and support monitoring and control strategies centered on the most influential CPPs.

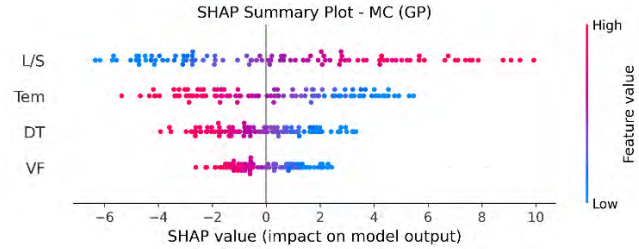

Figure 5. SHAP summary plot highlighting dominant CPP drivers for the CQAs.

Design-space coverage and conditional generation

The trained CVAE generated approximately 50, 000 candidate recipes. Figure 6A shows that generated recipes explore the broader CPP design space while clustering into feasible regions that correspond to target formulations. The generator naturally produced multiple distinct CPP combinations that meet the same target CQAs, illustrating the one-to-many nature of the inverse design

Figure 6. A. CVAE design-space exploration with clustered target formulation regions (colored) within the broader feasible set (grey). B. Close-up of Target center scenario with generated formulations conditioned to MC=10%, APS=510 um and Porosity =0.55.

mapping.

This inverse-design workflow complements established DS determination and uncertainty-aware DS methodologies by providing a generative mechanism to propose diverse CPP candidates that can then be validated within a DS framework [19-23]

Figure 6B illustrates example generated recipes conditioned on a representative target profile (MC=10%, APS=510 µm, Poro=0.55). Recipes remain diverse in CPP space while producing similar CQAs after feasibility screening, enabling operational flexibility.

Robustness across operating scenarios

Robustness was tested across seven operating scenarios that probe challenging specification demands (including low/high moisture, low/high APS, and low/high porosity).

Figure 7 summarizes percentage deviations from targets across scenarios. The target-center scenario achieved deviations of 2.0% (MC), 0.4% (APS) and 0.5% (Poro). The low-moisture scenario also performed well (3.1%, 0.3%, 0.7%).

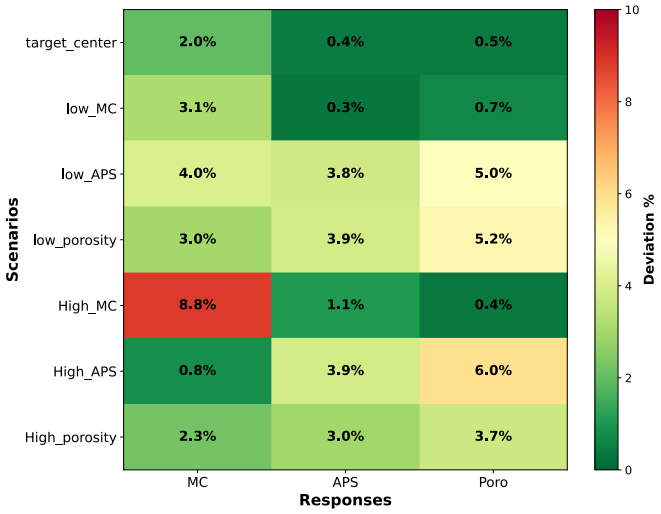

Figure 6. Deviation study for generated recipes across seven operating scenarios.

The most challenging case was high moisture content, where MC deviation increased to 8.8% while APS and porosity remained close to targets (1.1% and 0.4%). Overall, results indicate that the CVAE can learn constrained feasible regions and produce recipes that satisfy tight specifications across a range of industrially relevant conditions.

CONCLUSIONS AND FUTURE WORK

We presented a generative inverse-design framework for integrated continuous pharmaceutical manufacturing that combines a CVAE generator, a GP feasibility validator and SHAP-based explainability. Using a realistic wet granulation-drying case study, the framework generated large sets of candidate recipes (~50k) and achieved low deviations from targets across seven scenarios (best-case: target-center: 2.0% MC, 0.4% APS, 0.5% porosity; worst-case: 8.8% MC in the high-moisture scenario). Explainability analysis provided physically consistent insights, highlighting drying time and liquid-to-solid ratio as key levers for moisture control and L/S as the dominant driver for particle size and porosity.

Future work will focus on hybrid generative architectures, as well as experimental validation of selected recipes on pilot-scale equipment.

ACKNOWLEDGEMENTS

Financial support from EPSRC grants EP/V051008/1, EP/V034723/1, EP/W003317/1 and Peruvian Government through PRONABEC is gratefully acknowledged

AUTHOR IDENTIFIERS

Author ORCIDs:
Vega-Zambrano CDP: 0000-0003-0155-2088
Charitopoulos VM: 0000-0001-9051-917X

REFERENCES

1. Pantelides CC, Pereira FE. The future of digital applications in pharmaceutical operations. Current Opinion in Chemical Engineering 45:101038 (2024). https://doi.org/10.1016/j.coche.2024.101038
2. Jayakrishnan A, Phang HC, Ying Tan VX, Kee PE, Loke YH, Mod Razif MRF, Yee KM, Md Noh SM, Ming LC, Gan SH, Liew KB. Towards pharmaceutical industry 5.0: impact of artificial intelligence in drug discovery and development. CPB 27: (2025). https://doi.org/10.2174/0113892010396035250710120501
3. Nagy B, Galata DL, Farkas A, Nagy ZK. Application of artificial neural networks in the process analytical technology of pharmaceutical manufacturing—a review. AAPS J 24: (2022). https://doi.org/10.1208/s12248-022-00706-0
4. Destro F, Barolo M. A review on the modernization of pharmaceutical development and manufacturing – trends, perspectives, and the role of mathematical modeling. International Journal of Pharmaceutics 620:121715 (2022). https://doi.org/10.1016/j.ijpharm.2022.121715
5. International Conference on Harmonisation (ICH). Continuous Manufacturing of Drug Substances and Drug Products, Q13 (2022).
6. International Council for Harmonisation (ICH). ICH

Q8(R2): Pharmaceutical Development (2009).
7. Wang LG, Omar C, Litster J, Slade D, Li J, Salman A, Bellinghausen S, Barrasso D, Mitchell N. Model driven design for integrated twin screw granulator and fluid bed dryer via flowsheet modelling. International Journal of Pharmaceutics 628:122186 (2022). https://doi.org/10.1016/j.ijpharm.2022.122186
8. Kingma DP, Welling M. Auto-encoding variational bayes. arXiv:1312.6114 (2013) https://arxiv.org/abs/1312.6114
9. Sohn K, Lee H, Yan X. Learning structured output representation using deep conditional generative models. Adv Neural Inf Process Syst 28 (2015).
10. Fuhr AS, Sumpter BG. Deep generative models for materials discovery and machine learning-accelerated innovation. Front. Mater. 9: (2022). https://doi.org/10.3389/fmats.2022.865270
11. Tong X, Liu X, Tan X, Li X, Jiang J, Xiong Z, Xu T, Jiang H, Qiao N, Zheng M. Generative models for de novo drug design. J. Med. Chem. 64:14011-14027 (2021). https://doi.org/10.1021/acs.jmedchem.1c00927
12. Hornick T, Mao C, Koynov A, Yawman P, Thool P, Salish K, Giles M, Nagapudi K, Zhang S. In silico formulation optimization and particle engineering of pharmaceutical products using a generative artificial intelligence structure synthesis method. Nat Commun 15: (2024). https://doi.org/10.1038/s41467-024-54011-9
13. Qian Y, Li S. Inverse design of particle shapes with target sphericity and packing fraction using variational autoencoders. Engineering Applications of Artificial Intelligence 162:112509 (2025). https://doi.org/10.1016/j.engappai.2025.112509
14. Ma Y, Li W, Liu J, Shang G, Yang H, Gong J, Nagy ZK, Benyahia B. Digital design and optimization of the integrated synthesis and crystallization process using data?driven approaches. AIChE Journal 71: (2025). https://doi.org/10.1002/aic.18931
15. Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. Adv Neural Inf Process Syst 30 (2017). https://doi.org/10.48550/arXiv.1705.07874
16. Siemens Process Systems Engineering, 2025. gPROMS.https://www.psenterprise.com/products/gproms.
17. Bratley P, Fox BL. Algorithm 659. ACM Trans. Math. Softw. 14:88-100 (1988). https://doi.org/10.1145/42288.214372
18. Vega-Zambrano C, Diangelakis NA, Charitopoulos VM. Data-driven model predictive control for continuous pharmaceutical manufacturing. International Journal of Pharmaceutics 672:125322 (2025). https://doi.org/10.1016/j.ijpharm.2025.125322
19. Geremia M, Bezzo F, Ierapetritou MG. Design space determination of pharmaceutical processes: effects of control strategies and uncertainty. European Journal of Pharmaceutics and Biopharmaceutics 194:159-169 (2024). https://doi.org/10.1016/j.ejpb.2023.12.008
20. Bano G, Facco P, Ierapetritou M, Bezzo F, Barolo M. Design space maintenance by online model adaptation in pharmaceutical manufacturing. Computers & Chemical Engineering 127:254-271 (2019). https://doi.org/10.1016/j.compchemeng.2019.05.019
21. Meng Q, Bogle D, Charitopoulos VM. Probabilistic design space exploration and optimization via bayesian approach for a fluid bed drying process. European Journal of Pharmaceutical Sciences 210:107116 (2025). https://doi.org/10.1016/j.ejps.2025.107116
22. Chiplunkar R, Pauzi SM, Sachio S, Papathanasiou MM, Kontoravdi C. Probabilistic design space identification for upstream bioprocesses under limited data availability?. Systems and Control Transactions 4:2561-2567 (2025). https://doi.org/10.69997/sct.166359
23. Boukouvala F, Muzzio FJ, Ierapetritou MG. Design space of pharmaceutical processes using data-driven-based methods. J Pharm Innov 5:119-137 (2010). https://doi.org/10.1007/s12247-010-9086-y

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Development of Symbolic Regression-Based ATR-FTIR Calibration Models

Fernando A. R. D. Lima[abc]*, Inga S. Nordhus[a], Marcellus G. F. de Moraes[c], M. Enis Leblebici[d], Argimiro R. Secchi[bc], Mauricio B. de Souza Jr.[bc] and Idelfonso Nogueira[a]

[a] Chemical Engineering Department, Norwegian University of Science and Technology, Trondheim, 793101, Norway
[b] EPQB, School of Chemistry, Universidade Federal do Rio de Janeiro, Av. Horácio Macedo, 2030, CT, Bloco E, 21941-914, Rio de Janeiro, RJ – Brazil
[c] Chemical Engineering Program, PEQ/COPPE – Universidade Federal do Rio de Janeiro, Av. Horácio Macedo, 2030, CT, Bloco G, G115, 21941-914, Rio de Janeiro, RJ – Brazil
[d] KU Leuven, Center for Industrial Process Technology, Diepenbeek, Belgium
* Corresponding Author: farrais@eq.ufrj.br.

ABSTRACT

Accurate calibration of spectroscopic measurements is essential for reliable real-time monitoring and control of crystallization processes. In this work, calibration strategies for Attenuated Total Reflectance Fourier Transform Infrared (ATR-FTIR) spectroscopy were systematically evaluated for concentration monitoring in batch cooling crystallization of paracetamol in ethanol. Linear regression (LR), Partial Least Squares Regression (PLSR), Principal Component Regression (PCR), and symbolic regression (SR) were compared using both peak-based features and full spectral representations. Peak-based models provided a transparent baseline, with peak-area-based models consistently outperforming peak-height-based models. For LR, incorporating multiple absorption bands reduced the mean squared error (MSE) by nearly one order of magnitude compared to single-peak models. Using the same peak-based inputs, SR further improved performance, reducing prediction bias at high concentrations and yielding higher coefficients of determination ($R^2 > 0.99$) compared to LR. A substantial improvement was achieved when full spectral information was used. Among all evaluated approaches, SR with unprocessed spectra yielded the best overall performance, achieving an R^2 of 0.996 and an MSE of 1.4×10^{-6} on the validation dataset. This model also demonstrated strong generalization on an independent solubility test dataset, closely reproducing the reference solubility curve over the full temperature range with minimal deviation. In contrast, PCR and PLSR models showed increased sensitivity to preprocessing choices and exhibited larger errors on the test dataset. SR provided an accurate, robust, and interpretable calibration framework for ATR-FTIR, with reduced reliance on spectral preprocessing and potential for real-time process analytical technology and control applications.

Keywords: PAT, Preprocessing, Process monitoring, PLSR, Crystallization

INTRODUCTION

Crystallization is a key separation and purification unit operation in the pharmaceutical and fine chemicals industries, where stringent quality requirements demand tight control over product attributes such as purity, crystal size distribution (CSD), polymorphic form, and morphology [1]. In batch cooling crystallization, these attributes are influenced by the temporal evolution of supersaturation, which governs nucleation and crystal growth kinetics. Poor supersaturation control can lead to excessive primary or secondary nucleation, agglomeration, fouling, and batch-to-batch variability, ultimately compromising product quality and process robustness [2]. As a result, real-time monitoring and control of crystallization processes have been a long-standing focus of research in process systems engineering and control [3].

The introduction of process analytical technology (PAT) frameworks has significantly advanced the monitoring and understanding of crystallization systems [4].

PAT tools enable real-time or near-real-time access to critical process information, supporting data-driven decision-making and advanced control strategies [5]. Commonly applied PAT techniques in crystallization include attenuated total reflectance Fourier transform infrared (ATR-FTIR) spectroscopy, Raman spectroscopy, focused beam reflectance measurement (FBRM), particle vision and measurement (PVM), and ultrasonic spectroscopy [6]. Spectroscopic techniques such as ATR-FTIR and Raman are primarily used for monitoring solute concentration and supersaturation, whereas chord length distribution (CLD) measurements from FBRM and image-based measurements from PVM provide indirect information related to particle size and shape. The complementary use of multiple PAT tools has enabled more comprehensive characterization of crystallization dynamics and facilitated the development of model-based monitoring and control frameworks [7].

Among these techniques, ATR-FTIR spectroscopy has become one of the most widely adopted PAT tools for concentration measurement in crystallization processes [8]. Its popularity stems from its non-invasive operation, fast sampling rates, robustness in slurry environments, and compatibility with laboratory and pilot-scale crystallizers. However, ATR-FTIR does not provide direct concentration measurements; instead, it generates spectra that must be translated into quantitative concentration estimates through calibration models [7]. The reliability of any ATR-FTIR-based monitoring or control strategy therefore depends fundamentally on the accuracy, robustness, and generalization capability of the calibration model [4, 6].

Early ATR-FTIR calibration approaches were predominantly based on linear regression (LR) models using selected peak heights or peak areas, often combined with baseline correction and explicit temperature compensation [9]. While these methods are straightforward to implement and interpret, they rely heavily on expert-driven feature selection and may fail to capture nonlinearities arising from temperature effects, solute–solvent interactions, and spectral overlap. To address these limitations, multivariate statistical techniques such as partial least squares regression (PLSR) have become the standard approach to develop ATR-FTIR calibration models [6]. PLSR enables dimensionality reduction while maximizing covariance between spectral data and concentration and has been successfully applied in a wide range of crystallization studies [10, 11].

Despite their widespread use, conventional calibration approaches exhibit several limitations. Linear models can be overly restrictive, while PLSR models are often sensitive to data preprocessing choices, including spectral filtering, scatter correction, and wavelength range selection [12]. Furthermore, latent-variable models such as PLSR may suffer from reduced interpretability, as the physical meaning of latent components is not always clear. These challenges become particularly pronounced when calibration models are expected to operate across varying temperatures, solvent compositions, or process conditions, motivating the exploration of more flexible and robust modeling strategies [4, 6].

In recent years, machine learning techniques have been increasingly investigated as alternatives or complements to classical calibration methods in spectroscopic PAT applications [13]. Approaches such as support vector regression, Gaussian process regression, and artificial neural networks (ANNs) have demonstrated improved predictive performance in some cases, particularly when nonlinear effects are significant [6]. However, many of these methods function as black-box models, limiting their interpretability and complicating their deployment in regulated industrial environments [4]. Moreover, their computational complexity and reliance on large, representative datasets can pose challenges for real-time implementation and long-term maintenance.

Within this context, symbolic regression (SR) has emerged as a promising interpretable machine learning approach for process modeling and calibration [7]. SR aims to identify explicit mathematical expressions that relate input variables to outputs, without assuming a predefined model structure [14]. By balancing model accuracy and complexity, SR can uncover nonlinear relationships while producing closed-form equations that are transparent and amenable to physical interpretation [15]. These characteristics make SR particularly attractive for PAT-enabled crystallization systems, where interpretability, robustness, and ease of deployment are critical.

Although SR has been successfully applied to a variety of chemical engineering problems, including model correction, system identification, and control-oriented modeling [14-16], its application to ATR-FTIR calibration for crystallization monitoring remains limited in the open literature [7]. Most existing studies continue to rely on LR or latent-variable methods, despite growing interest in interpretable scientific machine learning frameworks [4, 6]. This gap highlights the need for a systematic evaluation of SR as a calibration methodology and a direct comparison with established approaches under consistent experimental conditions.

In this work, ATR-FTIR calibration models were developed for paracetamol–ethanol systems using SR and benchmarked against conventional LR, PLSR, and and principal component regression (PCR) models. As presented in Figure 1, different input representations were investigated, including peak-based features, temperature, and principal components derived from limited spectral ranges. The influence of common preprocessing strategies was assessed, and the resulting models were evaluated in terms of predictive accuracy, robustness, and generalization capability. The main objective was to

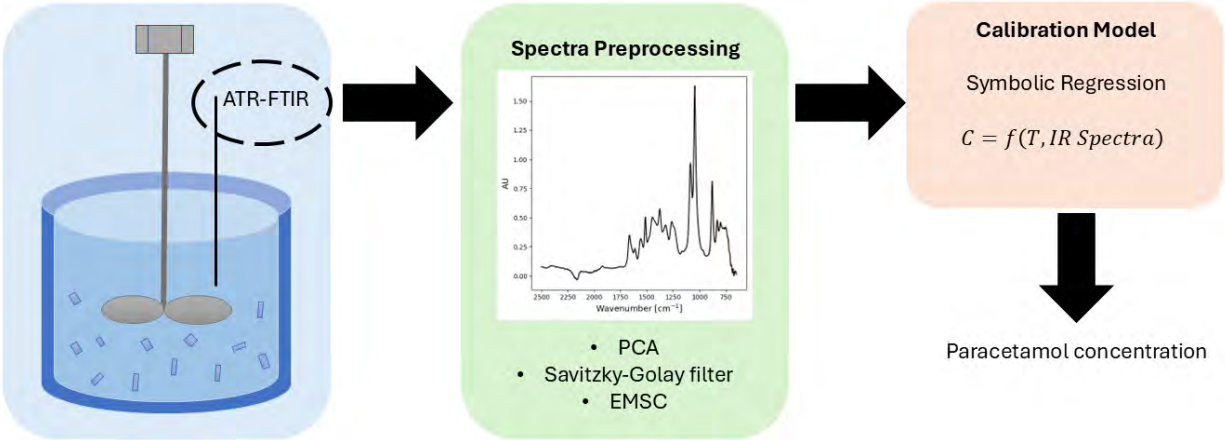

Figure 1. Symbolic regression approach proposed to develop ATR-FTIR calibration models.

assess the suitability of SR as a practical and interpretable alternative for ATR-FTIR calibration in crystallization monitoring and control applications.

DATA ACQUISITION

The experimental dataset used in this work was obtained from the batch crystallization study reported in Lima et al. [7]. In that work, paracetamol–ethanol mixtures were prepared gravimetrically and monitored in situ using ATR-FTIR spectroscopy, with complementary PAT tools (FBRM and temperature sensing) available in the same setup. Infrared spectra were acquired using a ReactIR system over 648–2500 cm^{-1} at a sampling rate of one spectrum per minute.

For calibration-model development, the dataset comprises 107 ATR-FTIR spectra paired with reference concentrations covering 0 to 0.5 g paracetamol/g ethanol and temperatures spanning from 10 to 70 °C. The data originate from four calibration experiments designed to systematically vary composition and temperature, including dilution sequences where ethanol was added to a fully dissolved paracetamol–ethanol solution and stepwise addition of paracetamol into ethanol. Across these experiments, spectra for a given concentration level were acquired at multiple temperatures, ensuring that the calibration models capture the coupled influence of concentration and temperature on spectral response. This dataset was randomly divided into calibration and validation subsets, with 70 % of the samples used for model training and the remaining 30 % reserved for validation.

In addition to the calibration and validation datasets, an independent test dataset was used to evaluate model generalization. This test set corresponds to an equilibrium solubility experiment reported in Lima et al. [7], where a paracetamol–ethanol mixture was held at 10 °C for 3 h to ensure equilibrium and then heated at 0.1 K/min up to 55 °C. The resulting equilibrium concentrations were compared against the solubility model reported by Griffin et al. [17], providing an external check that model-predicted concentrations reproduce physically meaningful temperature–solubility trends beyond the calibration split.

SPECTRAL PREPROCESSING AND FEATURE CONSTRUCTION

Preprocessing approaches were systematically tested to quantify how common chemometric treatments and feature representations affect calibration accuracy and robustness. The starting point was a peak-based calibration strategy derived from chemically meaningful absorption bands of paracetamol in the ATR-FTIR spectra. Both peak height and peak area were evaluated as inputs. Prior to peak extraction, baseline subtraction was applied to reduce offsets and slow drifts in the spectra. Peak areas and peak heights were then computed over the following wavenumber intervals: 1584-1532 cm^{-1}, 1532-1492 cm^{-1}, and 1292-1184 cm^{-1}.

Beyond the peak-based representation, full-spectrum representations were also investigated to evaluate whether multivariate inputs improve performance relative to hand-crafted features. In addition to the full spectral range, a reduced window between 600 and 1800 cm^{-1} was considered, as it concentrates the most informative absorbance features for paracetamol while excluding regions more affected by noise or solvent interference. For noise attenuation, Savitzky–Golay filtering was tested using SciPy [18]. To compensate for baseline shifts and multiplicative scattering effects, extended multiplicative signal correction (EMSC) was evaluated using the chemometrics Python package [19], both alone and combined with smoothing. When spectral vectors (full or reduced) were used as model inputs, dimensionality reduction via principal component analysis (PCA) was performed using scikit-learn [20], and the resulting

components were used as regressors together with temperature.

All preprocessing variants were applied consistently across training, validation, and test datasets, and preprocessing parameters were kept fixed across modeling approaches to avoid information leakage and ensure that observed performance differences reflect the calibration models rather than data handling choices.

CALIBRATION MODEL DEVELOPMENT

Four calibration strategies were investigated to calculate the paracetamol mole fraction and compared in this work: LR, PLSR, PCR, and SR. LR models were constructed using peak-based spectral features obtained after baseline subtraction. Both peak heights and peak areas were considered as explanatory variables, together with temperature to account for its influence on spectral response and solubility. Model parameters were estimated using ordinary least squares as implemented in scikit-learn [20]. These linear models were directly compared to SR models using the same peak-area and peak-height inputs, enabling an explicit assessment of whether SR provides benefit over the conventional peak-based calibration strategy.

PLSR and PCR models were developed using spectral information in combination with temperature, using implementations available in scikit-learn [20]. Depending on the preprocessing strategy, either full-spectrum data or reduced spectral windows were used as inputs. For PCR, PCA was first applied to the spectral matrix, and a fixed number of principal components were retained as regressors in a linear model together with temperature. For PLSR, latent variables were extracted by maximizing covariance between spectral inputs and concentration. In both PCR and PLSR, the number of retained components/latent variables was selected a priori by inspecting the cumulative explained variance of the spectral decomposition and the stability of prediction performance on the held-out validation split, and then fixing a single value used consistently across all subsequent comparisons.

SR was employed as an interpretable machine learning approach to develop nonlinear calibration models without assuming a predefined functional form. SR models were developed using the PySR framework [21], which searches for explicit mathematical expressions relating spectral inputs and temperature to solute concentration. Different input representations were explored, including peak-based features (peak height/area) and principal components derived from full or reduced spectra. Model complexity was controlled by limiting expression depth and length to promote models suitable for real-time monitoring and control applications.

Across all calibration strategies, model performance was evaluated using identical data splits and performance metrics, ensuring a consistent and fair comparison between linear, latent-variable, and symbolic approaches.

RESULTS AND DISCUSSION

Peak-based Calibration Models

Initially, calibration models based on manually selected spectral features were investigated. Peak-based models were constructed using both peak height and peak area extracted from characteristic paracetamol absorption bands after baseline subtraction. The models were initially developed using a single absorption band and subsequently extended by including additional chemically relevant peaks, specifically the 1292–1184 cm^{-1}, 1532–1492 cm^{-1}, and 1584–1532 cm^{-1} regions. These models represent the conventional approach adopted in ATR-FTIR calibration and provide a benchmark against which multivariate and symbolic regression models can be assessed. The SR models were trained with 200 iterations and a population size of 20.

Table 1: Performance of the LR and SR models to calculate the paracetamol mole fraction for the validation set, considering temperature and peak areas or heights as inputs

Peak Area			
Model	**Number of Inputs**	**R^2**	**MSE**
LR	2	0.982	1.40×10^{-5}
LR	3	0.984	1.31×10^{-5}
LR	4	0.990	7.74×10^{-6}
SR	2	0.985	1.21×10^{-5}
SR	3	0.991	7.16×10^{-6}
SR	4	0.994	4.41×10^{-6}
Peak Height			
Model	**Number of Inputs**	**R^2**	**MSE**
LR	2	0.986	1.08×10^{-5}
LR	3	0.986	1.14×10^{-5}
LR	4	0.993	5.75×10^{-6}
SR	2	0.976	3.49×10^{-3}
SR	3	0.992	5.97×10^{-6}
SR	4	0.997	2.37×10^{-6}

Table 1 summarizes the performance of the LR and SR models developed using peak-based spectral features and temperature as inputs. For both peak height and peak area representations, increasing the number of peaks used as inputs leads to a systematic improvement in model accuracy, as reflected by increasing R^2 values and decreasing mean squared error (MSE). This trend is observed consistently for both LR and SR models and indicates that incorporating multiple chemically relevant

absorption bands provides complementary information, reducing sensitivity to noise and local spectral distortions associated with individual peaks.

For LR models, the improvement with additional peaks is particularly evident when moving from two to four inputs, where the MSE decreases significantly for both peak height and peak area representations. This reflects the limited expressive capability of linear models, which rely on sufficiently rich input information to approximate the underlying concentration–temperature relationship. Eq. 1 shows the resulting LR formulation, where the predicted concentration is expressed as a weighted linear combination of peak intensities and temperature. While this structure is transparent and easy to interpret, it cannot explicitly represent nonlinear interactions between spectral features and temperature.

$$X_{para} = -0.000191T + 0.271\, h_{1256} + 0.356 h_{1512} - 0.559 h_{1560} + 0.0447 \quad (1)$$

On the other hand, the SR model, given in Eq. 2, introduces nonlinear combinations of peak intensities and temperature through higher-order and multiplicative terms. These nonlinear structures allow the model to capture temperature-dependent changes in spectral response and interactions between absorption bands that are not accessible to linear regression. As a result, SR achieves lower MSE and higher R^2 values than LR for the same set of peak-based inputs, particularly when three or four peaks are included. The model described in Eq. 2 was obtained using the three peak heights as inputs, but only the most relevant inputs stay in the equation in symbolic regression.

$$X_{para} = h_{1512}[\,(0.714 + 0.430 h_{1512})^4 - 0.00553 h_{1560} T] \quad (2)$$

Figure 2 presents the parity plots for LR and SR models using four peak heights and temperature as inputs. Both models show good agreement between predicted and reference concentrations. However, the SR model exhibited reduced bias at higher concentrations and improved clustering around the parity line. This behavior was consistent with the nonlinear structure of Eq. 2, which enables better extrapolation toward the solubility limit compared to the linear formulation in Eq. 1.

Despite these improvements, peak-based models remain fundamentally limited by the reduced information content of manually selected spectral features. Although SR partially mitigates this limitation through nonlinear functional forms, the performance gains remain modest when compared to multivariate models based on full or reduced spectral representations. Nevertheless, the peak-based analysis establishes a clear and interpretable baseline, demonstrating that SR can extract additional predictive value from low-dimensional inputs while preserving model transparency and suitability for real-time implementation.

Spectra-based Calibration Models

To overcome the limitations associated with manually selected peak-based features, calibration models based on spectral representations were subsequently investigated. In this approach, the ATR-FTIR spectra were used directly as model inputs together with temperature, allowing the calibration models to exploit distributed spectral information related to both concentration and temperature effects. The comparative performance of the spectra-based calibration models is summarized in Table 2 for the validation set, while representative parity plots for the validation dataset are shown in Fig. 3.

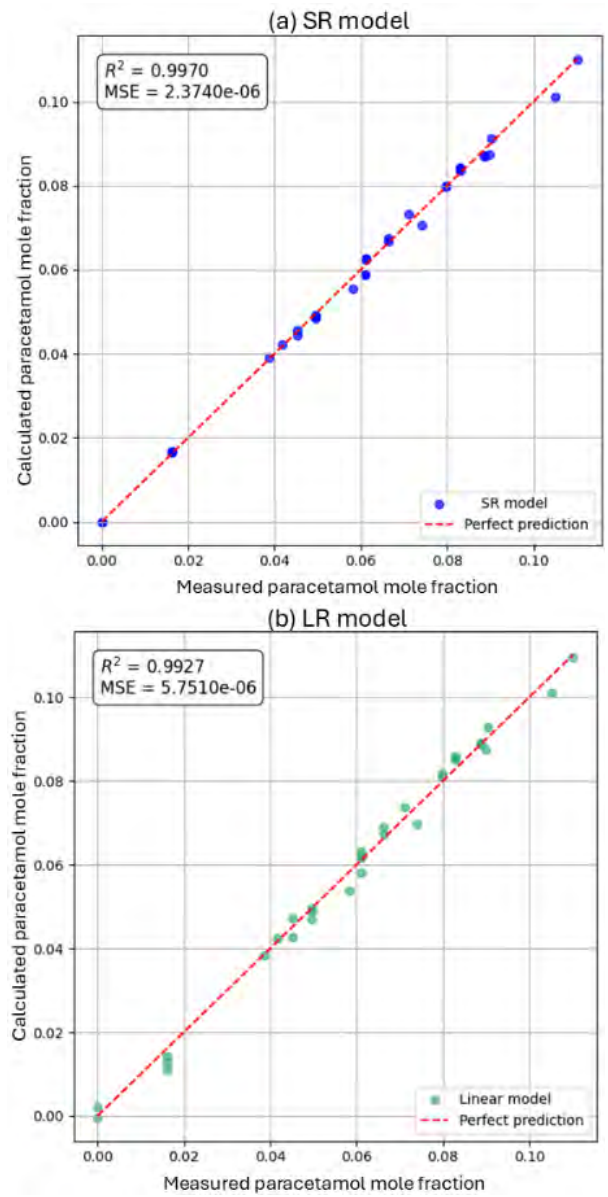

Figure 2. Comparison between the predicted and measured values of the paracetamol mole fraction for the validation set, considering SR and LR models using four peak heights and temperature as input.

Table 2: Performance of spectra-based calibration models considering different preprocessing strategies for the validation set

	Metric	Full spectra			Trimmed spectra		
		SR	**PLSR**	**PCR**	**SR**	**PLSR**	**PCR**
No preprocessing	R^2	0.996	0.974	0.971	0.992	0.974	0.9708
	MSE	1.40×10^{-6}	9.83×10^{-6}	1.12×10^{-5}	3.18×10^{-6}	9.83×10^{-6}	1.12×10^{-5}
SG filter	R^2	0.997	0.998	0.997	0.998	0.998	0.997
	MSE	1.15×10^{-6}	8.48×10^{-7}	1.03×10^{-6}	8.06×10^{-7}	9.60×10^{-7}	1.20×10^{-6}
EMSC	R^2	0.994	0.994	0.992	0.992	0.994	0.992
	MSE	2.31×10^{-6}	2.32×10^{-6}	2.91×10^{-6}	3.10×10^{-6}	2.32×10^{-6}	2.91×10^{-6}
EMSC + SG filter	R^2	0.995	0.997	0.996	0.992	0.992	0.996
	MSE	1.83×10^{-6}	1.03×10^{-6}	1.62×10^{-6}	3.11×10^{-6}	2.91×10^{-6}	1.48×10^{-6}

The influence of spectral preprocessing was systematically evaluated using common chemometric approaches, including Savitzky–Golay (SG) filtering, EMSC, and spectral range reduction. As presented in Table 2, these preprocessing steps noticeably affect linear and latent-variable models, but have a limited impact on SR. Notably, the best overall performance was obtained with SR trained on the full, unprocessed spectra, indicating that SR can internally construct nonlinear combinations of spectral features that effectively account for baseline variations, noise, and temperature-dependent distortions without requiring explicit preprocessing. In several cases, applying preprocessing led to negligible improvements or slight degradations, reinforcing that extensive preprocessing is not necessary for the SR workflow in this dataset.

As shown in Table 2, spectra-based models consistently outperformed peak-based strategies across all modeling approaches. This improvement is particularly pronounced for SR, which achieves the lowest prediction errors and highest coefficients of determination. The parity plots in Fig. 3 presents the improved agreement between predicted and measured concentrations when full spectral information is used, with reduced scatter and minimal systematic bias across the calibration and validation datasets.

The calibration model obtained by SR using the raw spectra and temperature as inputs is described in Eq. 3. $\hat{h}_i$ is the absorbance intensity at wavenumber i cm^{-1}. All SR models were trained with 200 iterations and a population size of 20.

$$X_{para} = \hat{h}_{1656}(1.148 - \hat{h}_{1056}) + 2.425(\hat{h}_{2128}^2 - 0.00419) \quad (3)$$

Eq. 3 highlights the ability of SR to identify compact and interpretable nonlinear relationships directly from raw spectral inputs. The model combines contributions from multiple spectral regions through both linear and nonlinear terms, capturing interactions that are not explicitly represented in linear or latent-variable models.

The generalization capability of the spectra-based models was further assessed using an independent test dataset, corresponding to equilibrium solubility measurements. The performance on this test dataset is shown in Fig. 4, in which model predictions are compared against reference solubility model as a function of temperature from Griffin et al. [17]. The SR model trained on the full, unprocessed spectra shows close agreement with the reference solubility behavior over the investigated temperature range, with only minor deviations near the solubility limit. The performance of the SR model confirms the extrapolation capability beyond the calibration split and demonstrates that the model preserves physically meaningful concentration–temperature trends under equilibrium conditions.

PCR and PLSR models exhibited more dependence on preprocessing choices. Although dimensionality reduction and scatter correction can improve calibration and validation metrics, increased deviation from the reference solubility trend was observed on the test dataset, as shown in Fig. 4, indicating reduced robustness when

extrapolating toward equilibrium conditions. Overall, the spectra-based analysis demonstrates that combining full ATR-FTIR spectra with SR yielded the most accurate and robust calibration strategy investigated in this study, and that its accurate performance on the independent solubility test dataset supports its suitability for reliable spectroscopic monitoring and control of crystallization processes.

CONCLUSION

This study evaluated different calibration strategies for ATR-FTIR-based concentration monitoring, comparing conventional peak-based models with spectra-based approaches using LR, latent-variable methods, and SR. Peak-based models provided a transparent baseline, with improved performance when multiple absorption bands and peak areas were used. SR consistently outperformed LR for the same peak-based inputs by capturing nonlinear interactions between spectral features and temperature, although performance remained limited by the reduced information content of manually selected peaks.

The use of full spectral information led to a substantial improvement in calibration accuracy and robustness. Among all approaches, SR applied directly to the unprocessed spectra yielded the best performance and demonstrated generalization on an independent solubility test dataset. The ability of SR to deliver compact, interpretable models without extensive preprocessing highlights its potential as a practical and robust calibration methodology for real-time spectroscopic monitoring and control of crystallization processes.

ACKNOWLEDGEMENTS

This study was financed in part by CAPES - Finance Code 001 and by Petrobras S.A. (Cooperation term no. 0050.0125244.23.9). Professors MB de Souza Jr. and AR Secchi are grateful to financial support from CNPq (Grants No. 311153/2021-6 and 300744/2025-0).

AUTHOR IDENTIFIERS

Author ORCIDs:
Lima FARD: 0009-0004-5025-3804
Nordhus IS: -
Moraes MGF: 0000-0003-1580-3669
Leblebici ME: 0000-0003-4599-9412
Secchi AR: 0000-0001-7297-3571
Souza MB Jr: 0000-0002-1090-8958
Nogueira IBR: 0000-0002-0963-6449

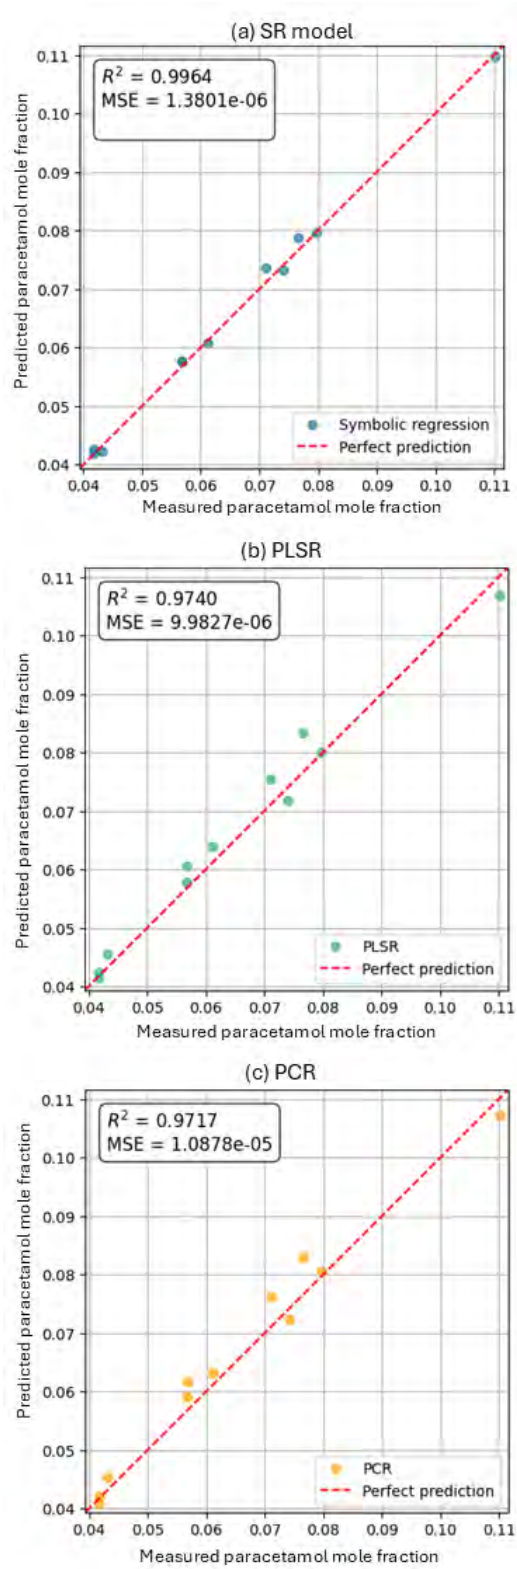

Figure 3. Comparison between the predicted and measured values of the paracetamol mole fraction for the validation set, considering SR, PLSR and PCR models using all spectra and temperature as input.

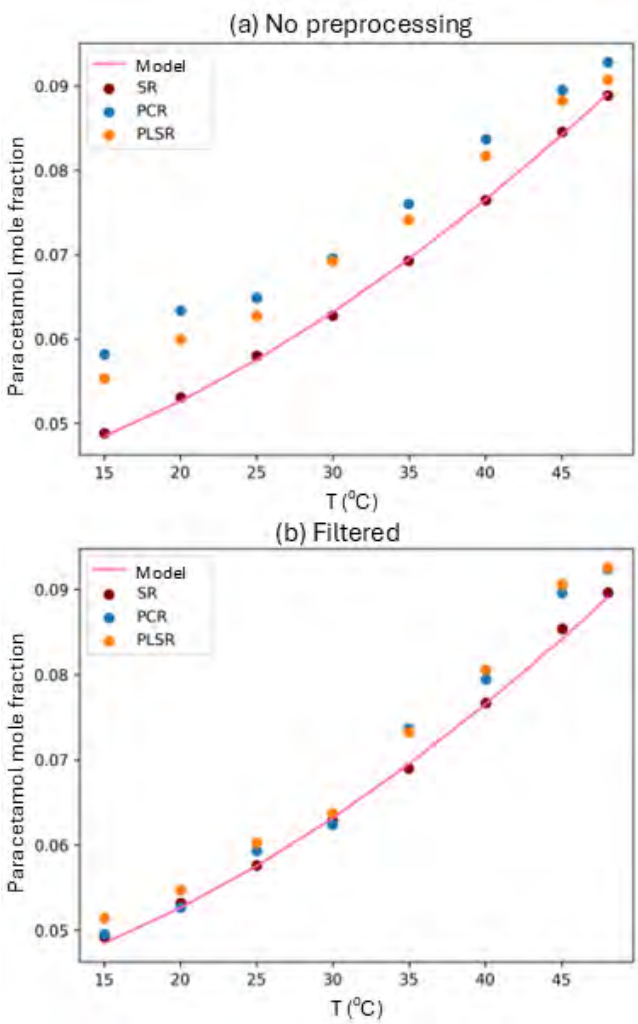

Figure 4. Comparison between the predictions of SR, PLSR and PCR to the solubility model of Griffin et al. [17], considering no preprocessing and filtering.

REFERENCES

1. Gao Z, Rohani S, Gong J, Wang J. Recent developments in the crystallization process: toward the pharmaceutical industry. Engineering 3:343-353 (2017). https://doi.org/10.1016/j.eng.2017.03.022
2. Braatz RD. Advanced control of crystallization processes. Annual Reviews in Control 26:87-99 (2002). https://doi.org/10.1016/s1367-5788(02)80016-5
3. Nagy ZK, Braatz RD. Advances and new directions in crystallization control. Annu. Rev. Chem. Biomol. Eng. 3:55-75 (2012). https://doi.org/10.1146/annurev-chembioeng-062011-081043
4. Lima FARD, de Moraes MGF, Barreto AG Jr, Secchi AR, Grover MA, de Souza MB Jr. Applications of machine learning for modeling and advanced control of crystallization processes: developments and perspectives. Digital Chemical Engineering 14:100208 (2025). https://doi.org/10.1016/j.dche.2024.100208
5. Simon LL, et al. Assessment of recent process analytical technology (PAT) trends: a multiauthor review. *Org Process Res Dev* 19:3–62 (2015) https://doi.org/10.1021/op500261y
6. Xiouras C, Cameli F, Quilló GL, Kavousanakis ME, Vlachos DG, Stefanidis GD. Applications of artificial intelligence and machine learning algorithms to crystallization. Chem. Rev. 122:13006-13042 (2022). https://doi.org/10.1021/acs.chemrev.2c00141
7. Dias Lima FAR, Fernandes de Moraes MG, Resende Secchi A, de Souza MB Jr, Grover MA. Experimental nonlinear model predictive control of crystal size and yield in batch cooling crystallization enabled by soft sensor and symbolic-based calibration model. Ind. Eng. Chem. Res. 64:23582-23600 (2025). https://doi.org/10.1021/acs.iecr.5c03894
8. Zhang F, Du K, Guo L, Huo Y, He K, Shan B. Progress, problems, and potential of technology for measuring solution concentration in crystallization processes. Measurement 187:110328 (2022). https://doi.org/10.1016/j.measurement.2021.110328
9. Swinehart DF. The beer-lambert law. J. Chem. Educ. 39:333 (1962). https://doi.org/10.1021/ed039p333
10. Lindenberg C, Krättli M, Cornel J, Mazzotti M, Brozio J. Design and optimization of a combined cooling/antisolvent crystallization process. Crystal Growth & Design 9:1124-1136 (2008). https://doi.org/10.1021/cg800934h
11. Trampuž M, Tesli? D, Likozar B. Process analytical technology-based (PAT) model simulations of a combined cooling, seeded and antisolvent crystallization of an active pharmaceutical ingredient (API). Powder Technology 366:873-890 (2020). https://doi.org/10.1016/j.powtec.2020.03.027
12. Zhang F, Liu T, Wang XZ, Liu J, Jiang X. Comparative study on ATR-FTIR calibration models for monitoring solution concentration in cooling crystallization. Journal of Crystal Growth 459:50-55 (2017). https://doi.org/10.1016/j.jcrysgro.2016.11.064
13. Lu M, Rao S, Yue H, Han J, Wang J. Recent advances in the application of machine learning to crystal behavior and crystallization process control. Crystal Growth & Design 24:5374-5396 (2024). https://doi.org/10.1021/acs.cgd.3c01251
14. Lima FARD, de Moraes MGF, Rebello CM, Barreto AG Jr, Secchi AR, de Souza MB Jr, Nogueira IBR.

Interpretable and uncertainty-aware machine learning for trustworthy prediction in batch crystallization. Chemical Engineering and Processing - Process Intensification 215:110350 (2025). https://doi.org/10.1016/j.cep.2025.110350
15. Rebello CM, Costa EA, Fontana M, Schnitman L, Nogueira IBR. Interpretable scientific machine learning approach for correcting phenomenological models: methodology validation on an ESP prototype. Ind. Eng. Chem. Res. 63:19030-19050 (2024). https://doi.org/10.1021/acs.iecr.4c02104
16. Santana VV, Costa E, Rebello CM, Ribeiro AM, Rackauckas C, Nogueira IBR. Efficient hybrid modeling and sorption model discovery for non-linear advection-diffusion-sorption systems: a systematic scientific machine learning approach. Chemical Engineering Science 282:119223 (2023). https://doi.org/10.1016/j.ces.2023.119223
17. Griffin DJ, Kawajiri Y, Rousseau RW, Grover MA. Using MC plots for control of paracetamol crystallization. Chemical Engineering Science 164:344-360 (2017). https://doi.org/10.1016/j.ces.2017.01.065
18. Virtanen P, Gommers R, Oliphant TE, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods* 17:261–272 (2020) https://doi.org/10.1038/s41592-019-0686-2
19. Mehmood T, Liland KH, Snipen L, Sæbø S. A review of variable selection methods in partial least squares regression. Chemometrics and Intelligent Laboratory Systems 118:62-69 (2012). https://doi.org/10.1016/j.chemolab.2012.07.010
20. Pedregosa F, Varoquaux G, Gramfort A, et al. Scikit-learn: machine learning in Python. *J Mach Learn Res* 12:2825–2830 (2011)
21. Cranmer K. Interpretable machine learning for science with PySR and symbolic regression. *Mach Learn Sci Technol* 4:015018 (2023) https://doi.org/10.1088/2632-2153/ac9f7c

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Beyond Solid-Phase: Comparative Assessment of Liquid Phase Oligonucleotide Synthesis with Single- and Dual-Stage Diafiltration

Alberto Saccardo, Rachel Ha, Zoe Fang, and Benoît Chachuat*

The Sargent Centre for Process Systems Engineering, Department of Chemical Engineering, Imperial College London, UK
* Corresponding Author: b.chachuat@imperial.ac.uk.

ABSTRACT

Oligonucleotides are short, sequence-defined nucleic acid chains with major therapeutic and diagnostic potential. Their industrial production is currently dominated by solid-phase oligonucleotide synthesis (SPOS), which suffers from mass-transfer limitations, limited scalability, lack of real-time process monitoring, and high process mass intensity. Membrane-enhanced liquid-phase oligonucleotide synthesis (LPOS) has emerged as a scalable alternative, in which oligonucleotide chains are grown on soluble anchors and organic solvent nanofiltration is used (OSN) to remove excess reagents and by-products between each reaction steps. However, diafiltration also introduces a risk of large cumulative product loss over multiple addition cycles, which requires fine-tuning of design and operational strategies in practice. This paper presents the results of a comparative assessment of two LPOS variants with either a single- or dual-stage diafiltration against a state-of-the-art SPOS, within a unified dynamic modelling framework. A recent mechanistic kinetic model of SPOS is transferred to LPOS, including a description of OSN diafiltration dynamics and internal recycles. A feasibility analysis is conducted to identify operating regimes under which LPOS outperforms SPOS in terms of selected yield and impurity performance indicators and to expose the underlying trade-offs. Overall, the results establish LPOS with dual-stage diafiltration as a promising configuration towards more flexible and sustainable oligonucleotide manufacturing.

Keywords: Dynamic Modelling, Feasibility Analysis, Oligonucleotide Synthesis, Liquid-Phase Synthesis, Organic Solvent Nanofiltration, Membrane Cascade

INTRODUCTION

Oligonucleotides are short sequence strands of DNA or RNA that are central in modern biotechnology and medicine [1]. Their high selectivity and specificity allow targeted modulation of gene expression, underpinning applications ranging from antisense and RNA interference therapies to gene editing, molecular diagnostics, and polymerase chain reaction [2]. The growing clinical success of oligonucleotide-based drugs, together with increasing demand for personalised and precision medicine, has intensified the need for efficient, scalable and sustainable manufacturing routes.

Industrial oligonucleotide production is currently dominated by solid-phase oligonucleotide synthesis (SPOS), typically based on phosphoramidite chemistry [3]. In SPOS, nucleotide monomers are sequentially added to a growing chain immobilised on an insoluble resin, with each cycle comprising detritylation, coupling, oxidation, and capping steps [1]. Excess reagents and by-products are removed through repeated washing steps. However, SPOS suffers from mass-transfer limitations due to heterogeneous reaction conditions, which requires large excess of costly reagents and extensive solvent use and leads to high process mass intensity (PMI) [3]. In addition, SPOS scalability is constrained by resin handling and the lack of real-time monitoring hinders advanced process control strategies.

Liquid-phase oligonucleotide synthesis (LPOS) has emerged as a promising alternative to overcome these limitations [4-6]. In LPOS, the oligonucleotide chain is grown on a soluble multibranched support (homostar), allowing reactions to proceed under homogeneous

conditions. Unreacted reagents and by-products are removed between the reaction steps using organic solvent nanofiltration (OSN), while retaining the supported oligonucleotide due to the large size of the homostar complex. This configuration improves mass transfer, facilitates inline monitoring, and provides greater operational flexibility compared with SPOS. However, even small loss of supported oligonucleotides through the membrane during each cycle can significantly reduce the product yield after multiple cycles of addition. Multi-stage membrane configurations with intermediate recycle offer a strategy to enhance product recovery, but recycling of intermediate streams between diafiltration stages can also increase impurity retention. Careful process design and fine-tuned operational strategies thus become essential for LPOS to outperform SPOS.

Through this paper, we apply a consistent dynamic modelling framework to quantify and compare the performance of LPOS against SPOS. We adapt a recent mechanistic kinetic model of SPOS [7] to predict LPOS performance in both single- and dual-stage diafiltration configurations. We conduct a feasibility analysis to characterize operating regions in which LPOS outperforms SPOS in terms of product yield and impurities. These results provide clear insight into the conditions under which membrane-enhanced LPOS can provide a scalable and more sustainable route for oligonucleotide manufacture.

METHODOLOGY

Description of SPOS and LPOS Processes

Since SPOS is the industrial benchmark for oligonucleotide manufacture, it is used as the reference for comparison in this work. In SPOS, oligonucleotides are assembled via sequential phosphoramidite chemistry on an insoluble polymeric support, typically a functionalised polystyrene resin. The growing chain remains covalently bound to the solid phase throughout the synthesis, while reagents and by-products are convectively transported through the resin bed with the solvent [7]. Each synthesis cycle comprises four main steps: detritylation, coupling, oxidation (or thiolation), and capping [7]. Detritylation removes the 5′-dimethoxytrityl (DMT) protecting group, generating a reactive hydroxyl site. During coupling, an activated phosphoramidite monomer reacts with the deprotected chain, followed by oxidation of the internucleotide linkage. Unreacted hydroxyl groups are finally capped to prevent further chain extension, thereby limiting impurity formation. Excess reagents and by-products are removed through repeated washing steps between the reactions.

In LPOS by contrast, the oligonucleotide chains are grown onto a soluble multibranched support (homostar), enabling all reactions to occur under homogeneous conditions [4-6]. The reactor is coupled to an OSN unit for removing small molecules—including acids, activators, unreacted monomers, and by-products—while retaining the large homostar-supported oligonucleotides. Chain extension proceeds similarly to SPOS through repeated cycles of detritylation, coupling, and quenching, except that each reaction step is now followed by a constant volume diafiltration stage, during which fresh solvent is added at the same rate as permeate removal. In addition to eliminating diffusion limitations and facilitating inline monitoring of the liquid phase, LPOS therefore decouples the reaction times for both deprotection and coupling steps from the purification steps [4-6].

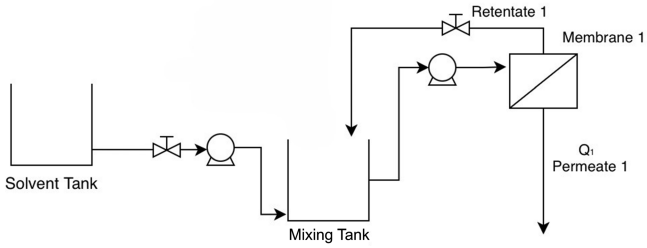

Figure 1. LPOS-1S process layout with single-stage OSN diafiltration (adapted from Kim et al. [5]).

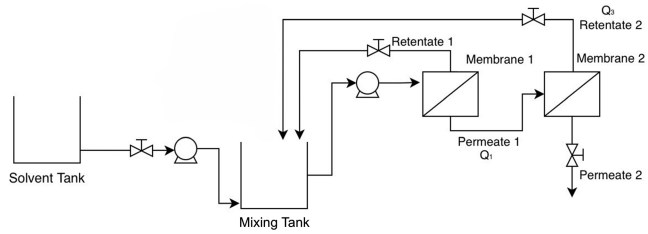

Figure 2. LPOS-2S process layout with dual-stage OSN diafiltration (adapted from Kim et al. [5]).

Although it is effective at removing small molecules, insufficient rejection of the supported oligonucleotide by the OSN membrane can incur large product loss over multiple cycles of addition [8]. To mitigate this loss, a LPOS configuration with dual-stage diafiltration (LPOS-2S, Figure 2) is investigated alongside single-stage diafiltration (LPOS-1S, Figure 1). In this configuration, the retentate from the first OSN unit is fed to a second OSN unit, and the retentate streams from both stages are returned to the reactor. The addition of a second stage improves the retention of supported oligonucleotides and increases product yield across multiple cycles [5, 6], but it also increases the residence time—and therefore the residual concentration—of other solutes that may interfere with subsequent coupling steps. The two-stage configuration therefore introduces a key trade-off between enhanced product recovery and increased impurity retention.

In this work, a unified dynamic modelling framework is used to describe the LPOS-1S and LPOS-2S process variants and enable their comparison with SPOS in terms of yield and impurity in a systematic feasibility analysis.

LPS transferred model

The mechanistic kinetic model of SPOS proposed by Pfister et al. [7]. is here transferred to LPOS. The key structural modification is the removal of resin-related transport phenomena. The SPOS model distinguishes concentrations between the bulk liquid and resin pores, relying on solid-phase balances and corresponding mass-transfer terms. In the LPOS model by contrast, all the solutes are in a well-mixed liquid phase. Intraparticle diffusion, pore–bulk gradients, and resin-phase balances are therefore neglected, and reaction rates are written directly in terms of bulk concentrations.

For a target sequence comprising N building blocks, the LPOS process is represented as a sequence of N repeated synthesis cycles. Within cycle $n = 1, \dots, N$, anchored oligonucleotides are described by their protected and deprotected states, denoted as OP_k and OD_k, respectively, and indexed by the chain length k. Small molecules include the acid catalyst used for deprotection and the building block added during coupling. The model is formulated as a multistage dynamic system, with state discontinuities at stage transitions describing pulse additions of reagents.

At the start of the synthesis, each homostar branch carries an unextended anchored chain, such that

$$OP_0(0) = HS \bullet NB \tag{1}$$

where HS (mol m^{-3}) denotes the homostar concentration, and NB (-) the number of branches per support. The total concentration of anchored oligonucleotides is given by

$$OT(t) = \sum_{k=0}^{N} \left(OP_k(t) + OD_k(t) \right) \tag{2}$$

Deprotection

In [7], detritylation is described as a reversible acid-catalysed reaction, reflecting possible re-attachment of the trityl group under resin-bound conditions. In practice, detritylation is operated far from equilibrium and the deprotection step can thus be modeled as an apparent irreversible reaction [4]. The balance equations describing the conversion of protected chains of lengths $k = 0, \dots, n-1$ into deprotected chains during cycle n are given by

$$\dot{OD}_k(t) = -\dot{OP}_k(t) = k_D OP_k(t) DCA(t) \tag{3}$$

where k_D (m^3 mol^{-1} s^{-1}) is the deprotection kinetic constant, and DCA (mol m^{-3}) denotes the acid catalyst concentration. The latter is introduced at the start $t_{D,n}$ of each deprotection stage as a pulse addition,

$$DCA(t_{D,n}^{-}) + DCA^{\text{in}} \tag{4}$$

Coupling

The coupling reaction follows standard phosphoramidite chemistry, with two LPOS-specific simplifications. First, the chain-length-dependent reduction in coupling reactivity included in the SPOS model to account for steric hindrance within the resin is omitted, as such effects are not observed under homogeneous liquid-phase conditions. Second, no capping step is included since homogeneous mixing and tighter control over reaction times in LPOS allow incomplete coupling to be minimised without deliberate chain termination. During the coupling stage of cycle n, deprotected anchored chains react with the n-th building block BB_n to form protected chains of length $k+1$,

$$OD_k + BB_n \rightarrow OP_{k+1}$$

The corresponding balance equations are:

$$\dot{OP}_{k+1}(t) = -\dot{OD}_k(t) = k_C BB_n(t) OD_k(t) \tag{5}$$

where k_C (m^3 mol^{-1} s^{-1}) denotes the coupling kinetic constant. The building block is introduced at the start $t_{C,n}$ of each coupling stage via a pulse addition,

$$BB_n(t_{C,n}^{+}) = BB_n(t_{C,n}^{-}) + BB^{\text{in}} \tag{6}$$

The auxiliary reactions (e.g. oxidation or quenching) are assumed to occur on much faster timescales than the coupling and diafiltration and are therefore not included in the dynamic model.

Single-Stage Diafiltration

After each reaction step, excess reagents and by-products are removed using OSN membrane technology [6]. Constant-volume diafiltration is used, where solvent is added at the same rate as permeate removal. Under perfect mixing conditions, the concentrations of solute $X \in \{OP_k, OD_k, BB_n, DCA\}$ in the retentate is described by the following balance equation

$$\dot{X}(t) = -\frac{Q_1}{V_1}(1 - r_X) X(t) \tag{7}$$

where V_1 (m^3 s^{-1}) denotes the retentate volume of the membrane unit, r_X (-) the observed membrane rejection of species X, and the permeate flow rate Q_1 (m^3) is defined as

$$Q_1 = AK\Delta p_1 \tag{8}$$

with A (m^2) the membrane area, K the membrane permeability (m s^{-1} bar^{-1}), and Δp_1 (bar) the transmembrane pressure difference.

Dual-Stage Diafiltration

The dual-stage diafiltration is described by the following coupled balance equations

$$\dot{X}_1(t) = \frac{1}{V_1}\left[-Q_1 X_1 (1 - r_X) + (Q_1 - Q_2) X_2\right] \tag{9}$$

$$\dot{X}_2(t) = \frac{1}{V_2}\left[Q_1 X_1 (1 - r_X) - Q_2 X_2 (1 - r_X) - (Q_1 - Q_2) X_2\right] \tag{10}$$

where V_1 and V_2 (m^3) denote the retentate volumes of the first and second stages, respectively, $X_1(t)$ and $X_2(t)$ (mol m^{-3}) the corresponding concentrations of each solute $X \in \{OP_k, OD_k, BB_n, DCA\}$, and the same membrane in terms of the rejection parameter r_X is assumed in both stages. In addition to the permeate flowrate from the first stage given by Eqn. (8), the permeate flowrate from the second stage is such that

$$Q_2 = (1-\rho)Q_1 \tag{11}$$

where the recycle ratio ρ is treated as an independent operating variable. The solute concentrations in the second-stage retentate are initially set to $X_2(0) = 0$. Then, at the end of each diafiltration stage, the retentates from both stages are recombined to set the initial condition of the subsequent reaction step [6].

Process Performance Indicators

The process performance is quantified in terms of three metrics:

- The overall yield, Y, is defined as the ratio of correctly assembled, fully protected oligonucleotides after N synthesis cycles to the theoretical maximum amount supported on the homostar:

$$Y = \frac{OP_N(t_{F2,N})}{HS \cdot NB} \tag{12}$$

  where $t_{F2,n}$ corresponds to the end of the final diafiltration step.

- The product purity, P, measures the fraction of correctly sequenced oligonucleotides relative to the total amount of anchored material present at the end if the synthesis:

$$P = \frac{OP_N(t_{F2,N})}{OT(t_{F2,N})} \tag{13}$$

  Achieving a high product purity is essential, as sequence impurities directly compromise therapeutic effectiveness. Moreover, sequence errors and desired products exhibit very similar physicochemical properties, which makes downstream separation particularly challenging. As a result, impurities need to be minimised within the synthesis process itself.

- The maximal residual building block ratio, R, is defined as the largest ratio between the total concentration of remaining building blocks and the total anchored material, following each coupling step:

$$R = \max_{n=1,\dots,N} \left[\frac{\sum_{k=1}^{n} BB_k(t_{C,n})}{OT(t_{C,N})} \right] \tag{14}$$

A low R value prevents the accumulation of unreacted building blocks during each cycle and thus limits the risk of unwanted side reactions in subsequent cycles, as per standard operational strategies in iterative oligonucleotide synthesis.

Overall, good process performance should seek to maximise the overall yield Y and product purity P, while simultaneously minimising the maximal residual building block ratio R.

Computational Aspects

We conducted feasibility analysis using the Python module NSFeas, interfaced from the C++ library MAGNUS (v4.0, https://github.com/omega-icl/magnus), which implements a nested sampling algorithm [9]. The method maintains a population of live points of given size that it steers towards the feasible operating domain through adaptive sampling. Each sample is labelled based on the extent to which their performance indicators Y, P and R violate the respective target thresholds, and those having the highest violation are replaced with newly drawn samples with lower violation and stored as dead points for subsequent reconstruction of the feasible domain. The algorithm is interrupted as soon as the live points all belong to the feasible domain.

Simulation of the dynamic LPOS model with either single- or dual-stage diafiltration was implemented using the Python module ODESLV, interfaced from the C++ library CRONOS (v4.0, https://github.com/omega-icl/cronos). Expression trees for the model equations and all the necessary derivatives are constructed using the library MC++ (v4.0, https://github.com/omega-icl/mcpp). ODESLV calls the solver CVODES of the library SUNDIALS (v7.1.1) [10] for numerical ODE integration, with relative and absolute tolerances set to 10^{-7}.

Table 1: Design and model parameter values used in this work. *Equivalent amounts are relative to the initial anchored oligonucleotides $OT(0) = HS \cdot NB$.

Parameter	Nominal value
N	20 cycles
HS	5.0 mol m^{-3}
NB	4
k_D	0.65 m^3 mol^{-1} s^{-1}
k_C	1.5 × 10^{-4} m^3 mol^{-1} s^{-1}
BB^{in}	1.5 eq*
DCA^{in}	9.0 eq*
A	8.5 × 10^{-2} m^2
V_1	9 × 10^{-4} m^3
Δp_1	8.9 bar
K	5.6 × 10^{-7} m^3 s^{-1} bar^{-1}
r_{BB}	0.40
r_{DCA}	0.33

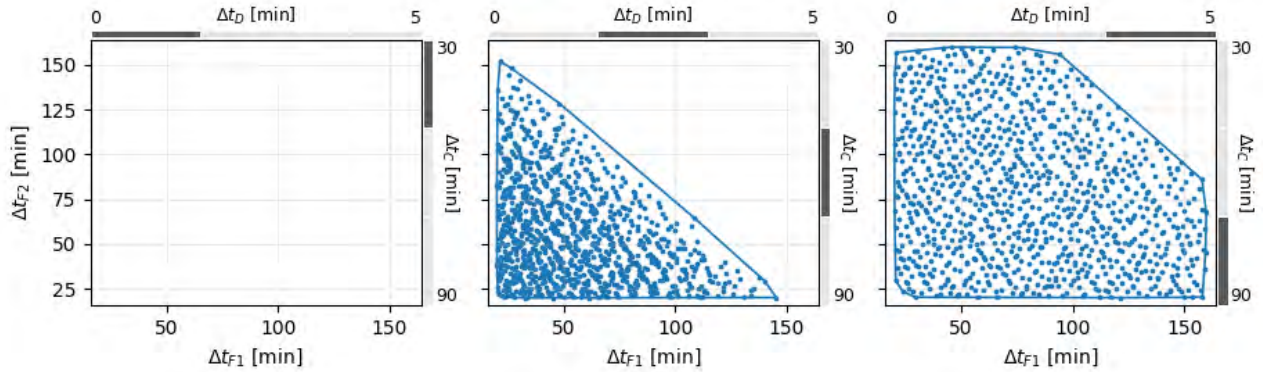

Figure 3: Feasible operating region of LPOS with single-stage diafiltration in terms of the yield constraint $Y \geq 0.86$.

RESULTS

We present the results of a case study for $N = 20$ synthesis cycles, using the parameter values in Table 1. The kinetic parameter values and reagent concentrations are adapted from Pfister et al. [7], and the LPOS-specific values are taken from [6, 8]. Unless otherwise noted, a homostar rejection of 99.90% is assumed for the OSN membrane. Due to the absence of diafiltration steps and neglecting any detachment of oligonucleotides from the resin support, the yield and purity of SPOS were both estimated at 86%, consistent with the specifications of the experimental system reported in Pfister *et al.* [7, 11], while the residual building block content trapped within the resin pores was estimated at 2%. These threshold values are used in the ensuing feasibility analysis of LPOS with OSN diafiltration.

A key advantage of LPOS over SPOS lies in the ability of exerting tighter control over the reaction steps, enabling higher product purity. Figure 3 compares the feasible operating regions of LPOS with single-stage diafiltration for the benchmark purity constraint $P \geq 0.86$ (blue points & line) as well as a tighter purity $P \geq 0.95$ (orange points & line). Projection onto the subspace $(\Delta t_C, \Delta t_D)$ of reaction times reveals that both regions are lower bounded by a convex combination of minimal deprotection and coupling times, as incomplete reactions directly reduce the formation of the desired product. Overall, this first comparison also shows that LPOS can achieve purities well above the SPOS purity benchmark of 86%.

Figure 4 shows the feasible operating region of LPOS with single-stage diafiltration for the benchmark yield constraint $Y \geq 0.86$. The four-dimensional feasible region is represented on a trellis chart, with the diafiltration times Δt_{F1} and Δt_{F2} on inner axes, and the deprotection and coupling times Δt_D and Δt_C indicated by grey-shaded ranges on outer axes. Since longer diafiltration steps (inner axes) result in greater loss of anchored oligonucleotides through the membrane, the feasible region is upper bounded by a maximal combined diafiltration duration, $\Delta t_{F1} + \Delta t_{F2} \leq 250\,\text{min}$. This maximal duration is seen to increase for longer deprotection and coupling times (outer axes), as bringing the reactions closer to completion during each addition cycle also increases the overall product yield.

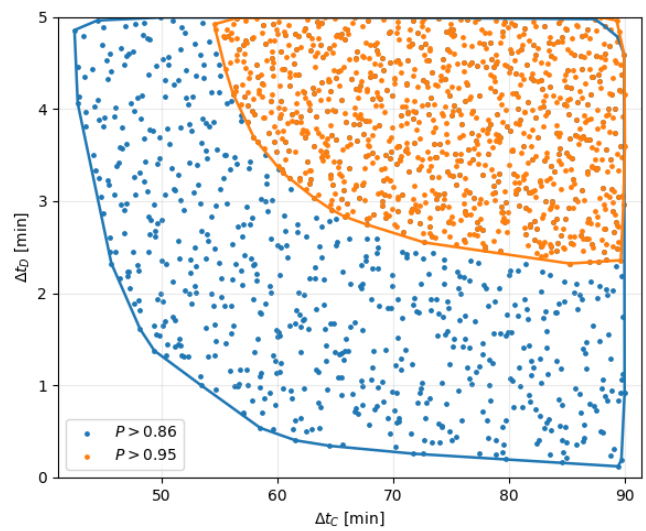

Figure 4. Feasible operating region of LPOS with single-stage diafiltration for two purity constraints $P \geq 0.86$ and $P \geq 0.95$.

Figure 5 shows a projection of the feasible region for the residue constraint $R \leq 0.02$ onto the subspace $(\Delta t_{F1}, \Delta t_{F2})$ of diafiltration times. In contrast to the yield constraint in Figure 4, this region is now lower bounded by a minimal diafiltration duration, $\Delta t_{F1} + \Delta t_{F2} \geq 215\,\text{min}$, as short diafiltration times can leave large amount of residual building blocks in the solvent. The feasible region of $R \leq 0.02$ is also less sensitive to the durations of the reaction steps: with 1.5 equivalents of building block added at each coupling step, even complete conversion of anchored oligonucleotides leaves a residual excess of 0.5 equivalents, so that the removal of residual building

blocks is primarily governed by the diafiltration times rather than the reaction times. In practice, the improved reaction control afforded by LPOS compared to SPOS greatly reduces the need for operating with large excess of reagents.

Taken together, Figures 4 & 5 highlight the opposing effect that the combined diafiltration duration $\Delta t_{F1} + \Delta t_{F2}$ has on the feasibility constraints $Y \geq 0.86$ and $R \leq 0.02$. Fine-tuning is thus critical if LPOS is to outperform SPOS simultaneously in terms of yield and purity. This trade-off is captured in Figure 6 by the projection of the combined feasible region for $Y \geq 0.86, P \geq 0.86$ and $R \leq 0.02$ onto the subspace $(\Delta t_{F1}, \Delta t_{F2})$ of diafiltration times.

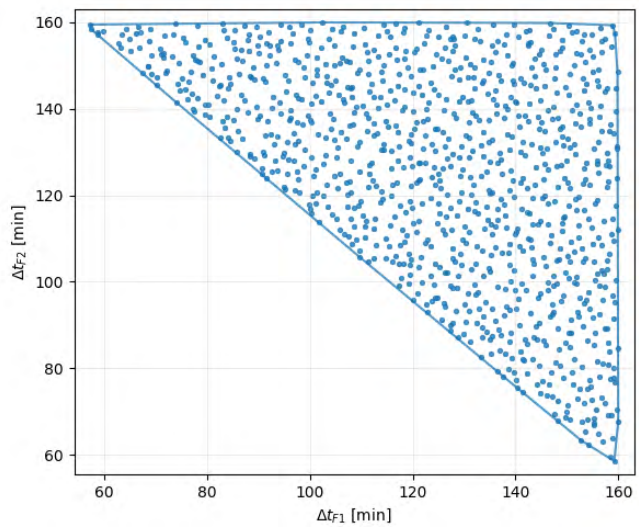

Figure 5. Feasible operating region of LPOS with single-stage diafiltration in terms of the maximal residual building block constraint.

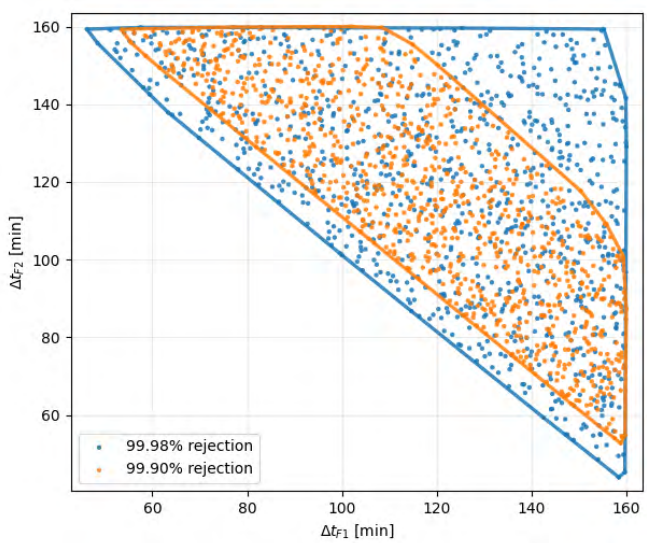

Figure 6. Feasible operating region of LPOS with single-stage diafiltration for the combined constraints $Y \geq 0.86, P \geq 0.86, R \leq 0.02$, comparison for two homostar rejection values of 99.90% and 99.98%.

For high homostar rejection values of 99.98% (blue points and lines, Figure 6), extending the combined diafiltration time, e.g. beyond 300 min in each cycle, does not result in appreciable product loss. But for lower rejection values of 99.90% (orange points and lines), the feasible region is clearly bracketed between a minimal and maximal combined diafiltration duration, e.g. between 210–270 min. Although homostar rejection in OSN membranes nowadays exceed 99.95% [4], this comparison demonstrates the strong interplay between the membrane selection and operational strategy on LPOS performance.

In principle, the implementation of dual-stage diafiltration can provide additional operational flexibility in LPOS, as it enhances product retention, while still allowing low residual building block levels. To establish this potential, the feasible operating regions on Figure 7 are computed for the tighter combined constraints $Y \geq 0.95, P \geq 0.95, R \leq 0.01$, considering homostar rejections of 99.90% and 99.25% by the OSN membrane (assumed identical in both stages). Notably, no feasible operating region exists for single-stage diafiltration under these combined constraints. Recall that dual-stage diafiltration introduces two extra operational parameters: the retentate volume V_2 of the second membrane stage (set to V_2= 5x10^{-5} m^3 on Figure 7), and the recycle ratio ρ. The extent of these feasible regions, here projected onto the subspace $(\Delta t_{F1} + \Delta t_{F2}, \rho)$, establishes that the performance of LPOS with dual-stage diafiltration can significantly exceed that of SPOS, even with relaxed membrane specifications.

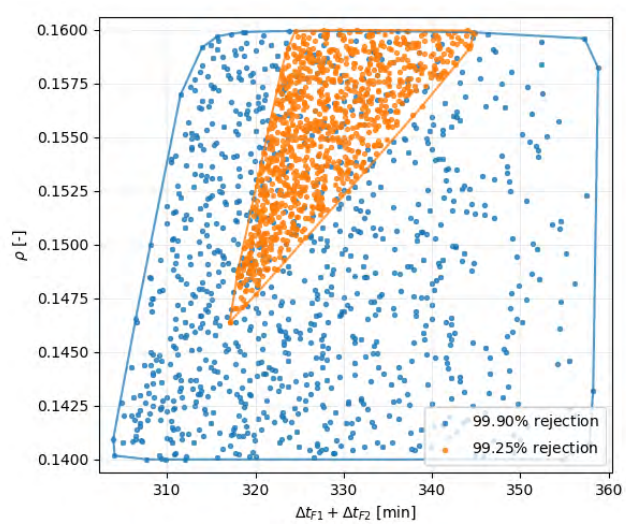

Figure 7. Feasible operating region of LPOS with dual-stage diafiltration for the tighter constraints $Y \geq 0.95, P \geq 0.95, R \leq 0.01$, comparison of two homostar rejection values 99.90% and 99.25% with a second-stage diafiltration volume V_2= 5x10^{-5} m^3.

The boundary of the feasible region on Figure 7 shows a clear interplay between the recycle ratio and the combined diafiltration time when the OSN membrane's rejection of homostar is low (orange points and line). In this case, the maximal diafiltration times (right boundary of projected feasible region) coincide with higher recycle ratios in order to avoid excessive loss of homostar-supported oligonucleotides through the membrane; whereas the minimal diafiltration times (left boundary of project feasible region) is driven by the stringent constraint $R \leq 0.01$, which is also contingent on the extents of the chemical reactions. The span of the feasible operating region (blue points and line) expands significantly as the membrane's rejection of homostar-supported oligonucleotides increases. Clearly, a higher rejection reduces the loss of oligonucleotides for a given recycle ratio, leading to higher building block consumption and in turn allowing for shorter diafiltration durations to achieve sufficient building block removal.

DISCUSSION

The comparison on Figure 6 establishes the potential of LPOS to outperform SPOS even with a single-stage diafiltration. However, the analysis also confirms that this configuration is highly sensitive to the diafiltration duration settings in an interplay with the OSN membrane's rejection properties. Diafiltration introduces a fundamental trade-off: longer diafiltration times improve the removal of unreacted building blocks, but they also increase the overall loss of product through the membrane; whereas shorter diafiltration times preserve the product yield, but they limit the washout of excess reagents. As a result, the feasible operating region of LPOS with single-stage diafiltration can be quite narrow and its high sensitivity to membrane selection requires careful tuning to balance excessive product loss with sufficient removal of unreacted reagents.

The deployment of dual-stage diafiltration with LPOS shows great promises for mitigating these challenges. The second stage increases effective product retention and provides additional flexibility through the recycle ratio and second-stage volume, allowing high yield and purity to be achieved while maintaining low residual building block levels. Consequently, LPOS with dual-stage diafiltration is predicted to significantly outperform nominal SPOS performance over a wide range of operating conditions, even for OSN membranes with low homostar rejection by today's standards. The dynamic model developed through this work predicts that the product yield and purity with LPOS could even reach values as high as 99.5%.

CONCLUSIONS

We conducted a model-based feasibility analysis to compare the performance of LPOS with both single- and dual-stage diafiltration against state-of-the-art SPOS. We transferred a mechanistic model of SPOS to LPOS by including OSN diafiltration dynamics and defined consistent performance indicators and threshold. While LPOS with single-stage diafiltration can outperform SPOS, its performance is highly sensitive to diafiltration times and membrane rejection properties. With dual-stage diafiltration, LPOS provides greater flexibility and can achieve higher product yield and purity, making it a promising alternative. Future work will expand the analysis to account for uncertainty in the kinetic and membrane parameters and enable probabilistic characterisation of the feasible operating space, and will seek to validate the model-based predictions experimentally.

ACKNOWLEDGEMENTS

This work was supported by Innovate UK under grant number 10103756 (Sustainable bioprocess for oligonucleotide manufacture using Nanostar Sieving - BioNanostar). The authors gratefully acknowledge Exactmer Ltd for providing them with information about the LPOS process.

AUTHOR IDENTIFIERS

Author ORCIDs:
Saccardo A: 0000-0002-9262-2207
Chachuat B: 0000-0003-4780-9686

REFERENCES

1. Pichon M, Hollenstein M. Controlled enzymatic synthesis of oligonucleotides. Commun Chem 7: (2024). https://doi.org/10.1038/s42004-024-01216-0
2. Vinjamuri BP, Pan J, Peng P. A review on commercial oligonucleotide drug products. Journal of Pharmaceutical Sciences 113:1749-1768 (2024). https://doi.org/10.1016/j.xphs.2024.04.021
3. Tolomelli A, Ferrazzano L, Cabri W. Sustainability in tides chemistry. Royal Society of Chemistry (2024). https://doi.org/10.1039/9781837674541
4. Gaffney PRJ, Jarrett-Wilkins C, Yeo J, Nylander C, Mansha K, Mukadam A, Gupta L, Akbar S, Jahan B, Li Volsi A, Pember-Jones A, So J, Murphy N, Rupsys E, Oxley A, Pypec M, Bertulli C, Felines N, O'Brien DM, Di Silvestro A, Ray B, Cuccato D, Khalily MA, Madugundu GS, Pereira I, Moore DN, Kleczewska N, Cook M, Diorazio LJ, Telford A, Benstead D, Radinov R, Brooks J, Nechev LV,

Parga C, Venturoni F, Haber J, MacLeod C, Cassey B, Sutherland E, Walton HP, Livingston AG. Fully liquid phase oligonucleotide synthesis. Org. Process Res. Dev. 30:98-120 (2025). https://doi.org/10.1021/acs.oprd.5c00346

5. Kim JF, Freitas da Silva AM, Valtcheva IB, Livingston AG. When the membrane is not enough: a simplified membrane cascade using organic solvent nanofiltration (OSN). Separation and Purification Technology 116:277-286 (2013). https://doi.org/10.1016/j.seppur.2013.05.050
6. Chen W, Sharifzadeh M, Shah N, Livingston AG. Iterative peptide synthesis in membrane cascades: untangling operational decisions. Computers & Chemical Engineering 115:275-285 (2018). https://doi.org/10.1016/j.compchemeng.2018.04.024
7. Pfister DE, Tilland A, Larue L, Kobl K, Weber P, Olbrich M. Kinetic modeling of solid-phase oligonucleotide synthesis: mechanistic insights and reaction dynamics. Org. Process Res. Dev. 29:2298-2309 (2025). https://doi.org/10.1021/acs.oprd.5c00199
8. Saccardo A, Chachuat B. Model-based analysis of membrane-enhanced liquid-phase oligonucleotide synthesis. IFAC-PapersOnLine 59:277-282 (2025). https://doi.org/10.1016/j.ifacol.2025.07.158
9. Kusumo KP, Gomoescu L, Paulen R, García Muñoz S, Pantelides CC, Shah N, Chachuat B. Bayesian approach to probabilistic design space characterization: a nested sampling strategy. Ind. Eng. Chem. Res. 59:2396-2408 (2019). https://doi.org/10.1021/acs.iecr.9b05006
10. Hindmarsh AC, Brown PN, Grant KE, Lee SL, Serban R, Shumaker DE, Woodward CS. SUNDIALS. ACM Trans. Math. Softw. 31:363-396 (2005). https://doi.org/10.1145/1089014.1089020
11. NittoPhase® HL solid support from Kinovate. https://www.kinovate.com/downloads/04%20_NPHL_flyer_0415.pdf

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

In silico solvent selection for green and cost-effective pregabalin crystallisation

Matthew Blair and Dimitrios I. Gerogiorgis*

Institute for Materials and Processes (IMP), School of Engineering, University of Edinburgh, Edinburgh, EH9 3FB, UK
* Corresponding Author: D.Gerogiorgis@ed.ac.uk

ABSTRACT

Identifying cost-optimal yet environmentally friendly crystallisation processes in the production of small molecule pharmaceuticals is a highly complex task, since multiple solvent systems often exist which could be used to purify a given drug to a similar standard. It is, however, rarely possible to test each of these solvent systems within a laboratory setting, since this would be time-consuming and incur large material costs. It has, therefore, been suggested that process modelling tools should be used to screen different crystallisation processes available to produce a new drug prior to studying them experimentally – essentially allowing a shortlist of promising process candidates to be created. Indeed, it has been shown by several authors that this sort of work can be conducted without any need to establish the crystallisation kinetics associated with each drug-and-solvent combination considered: the reason being that crystallisation processes may be defined as simple solid-liquid equilibria (SLE) problems using thermodynamic models alone. Considering this, the present work has used SLE modelling to assess the cost and environmental impact of different solvent systems available for the crystallisation of pregabalin – an anticonvulsant and anxiolytic drug used to treat epilepsy, anxiety and neuropathic pain. Moreover, it has used entirely data-led approaches to achieve this: leveraging established costing methodologies and green metrics (i.e., Scope 1 and 2 CO_2e emissions) to assess their economic viability and carbon footprint. This modelling framework can be used to identify viable crystallisation strategies for other small molecules.

Keywords: Pregabalin, technoeconomic analysis, crystallisation, environmental impact analysis, solvent selection.

1. INTRODUCTION

In 2024, pregabalin – a generic anticonvulsant drug and anxiolytic agent commonly used to treat epilepsy, anxiety and neuropathic pain [1-3] – had a global market of 1.8 billion USD, and it is expected that this market will continue to grow at a compound annual growth rate (CAGR) of 5.1 % between 2025 and 2034 [4]. Consequently, this work aims to identify technoeconomically optimal – and environmentally benign – manufacturing routes that can be used for its production. To achieve this, it has used in-silico technoeconomic analyses and environmental impact assessments to establish which solvents (of those commonly used in the crystallisation of small molecule pharmaceuticals) and process conditions allow cost-effective – yet green – pregabalin manufacturing.

2. PROCESS MODELLING STRATEGY

Designing crystallisation processes for pharmaceuticals can be an arduous process, since often several solvent systems will exist which could be used to perform the purification required by a given process – each of which will hold unique crystallisation kinetics. Recent studies [5, 6] have shown that promising crystallisation setups can be identified for new processes entirely in silico, without the need to establish kinetics (thereby greatly reducing the number of experiments required during process design). Thermodynamic modelling alone [5-7] can be used to achieve this by simply ranking each system based on their expected trends in process costs as API solubility varies under different operating conditions.

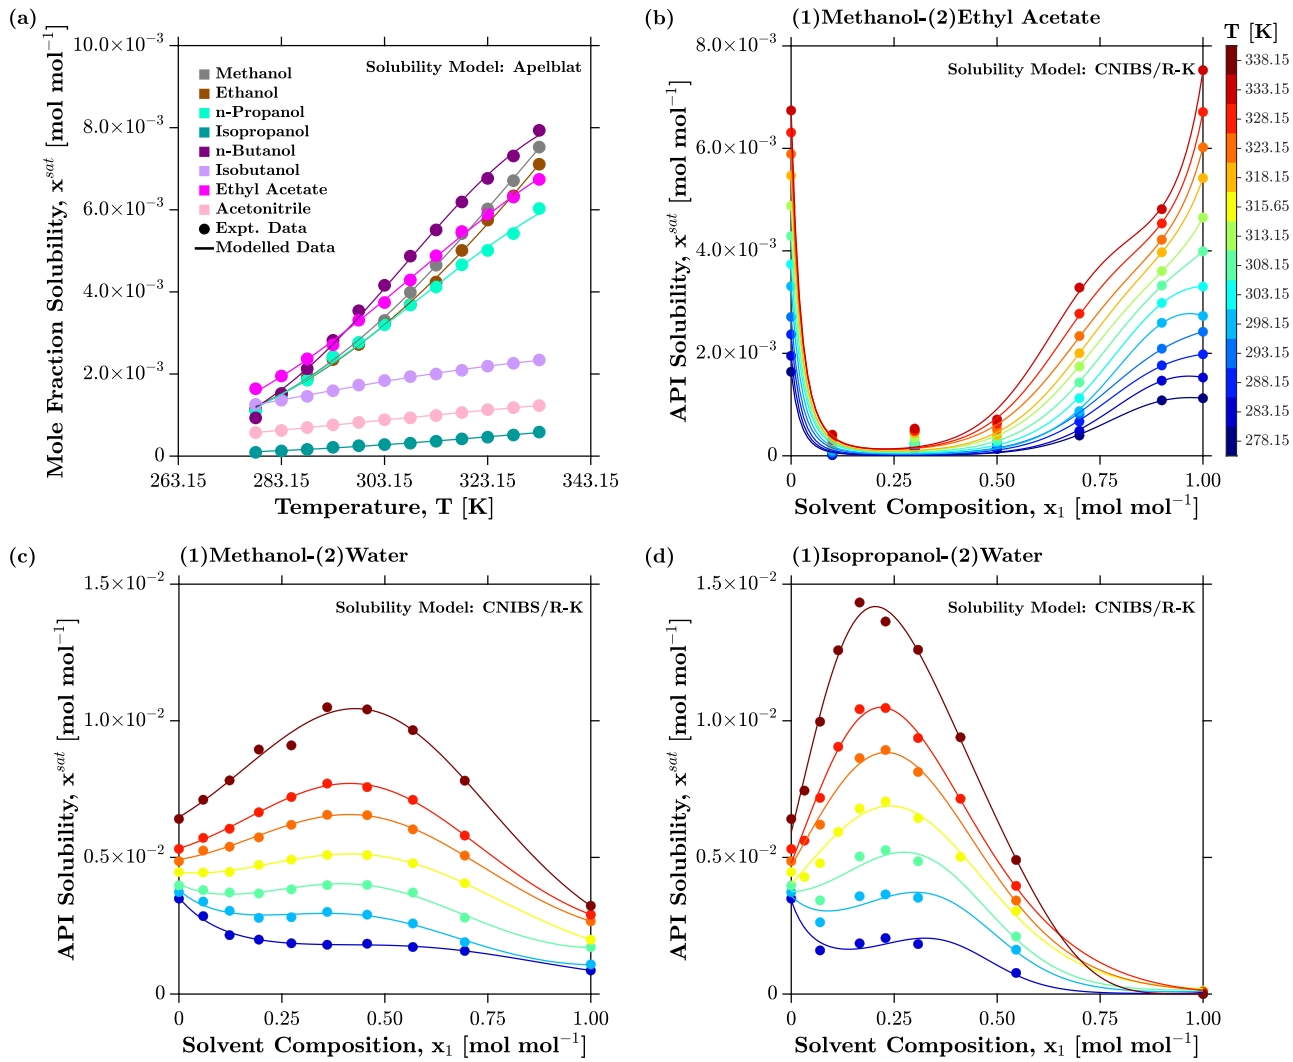

Figure 1. Pregabalin solubility in (a) pure and (b-d) mixed solvents [8-9].

This paper uses the modified Apelblat equation (Eq. 1) to model the temperature dependent solubility of pregabalin in pure solvents (methanol, ethanol, n-propanol, isopropanol, n-butanol, isobutanol, ethyl acetate, acetonitrile), whilst it has used the Combined Nearly Ideal Binary Solvent/Redlich-Kister (CNIBS/R-K) model (Eq. 2) to study its solubility in solvent mixtures (i.e., (1) methanol-(2) ethyl acetate, (1) methanol-(2) water, (1) isopropanol-(2) water) with varying molar compositions [8, 9] (Fig. 1).

$$\ln x^{sat} = A + \frac{B}{T} + C\ ln(T) \tag{1}$$

$$\ln x^{sat} = B_0 + B_1 x_1 + B_2 x_1^2 + B_3 x_1^3 + B_4 x_1^4 \tag{2}$$

Material balances can be computed for all processes, assuming that they will produce 100 metric tonnes of pregabalin over an annual operating time of 6000 h (m_{annual}) – thus allowing solid-liquid equilibrium (SLE) principles (Eqs. 3-6) to be used to determine the mass of pregabalin that would be (theoretically) produced per each batch (m_{BC}) in each crystallisation process proposed under various different operating conditions (i.e., solvent mixture composition, T_0 and T_f) [5, 7].

$$n_{prod,0} = x_{prod,0}^{sat} n_{S_T,0} \tag{3}$$

$$n_{prod} = x_{prod}^{sat} n_{S_T} \tag{4}$$

$$n_{prod}^{cryst} = n_{prod}^{0} - n_{prod} \tag{5}$$

$$m_{prod}^{cryst} = M_{prod} n_{prod}^{cryst} \tag{6}$$

It is important to note, however, that it has been assumed in this work that all process mixtures considered are well-mixed, fully suspended, and have uniform temperature and concentration profiles. Moreover, it has been assumed that each crystallisation process takes 12 hours (t_{cryst}) – since the thermodynamic models used

herein (Eqs. 1-2) have no time element – whilst the clean-in-place requirements of any equipment has been set to 30 minutes (t_{CIP}) [7, 10, 11].

This framework may be used to model both pure and mixed solvent systems; however, when dealing with mixed solvent systems, it is important to note that the total amount of solvent in each system must also obey the relationships expressed in Eqs. 7-14. In other words, the total amount of solvent present in the initial and final states of each system will be the sum of any solvents and antisolvents present (Eqs. 7-8) – noting that the composition of the initial and final states will be different if an antisolvent is added throughout the crystallisation process (Eqs. 9-12). Hence, the solute-free compositions of each solvent system must be held between 0 and 1 in each state (Eq. 13), whilst also summing to 1 (Eq. 14).

$$n_{S_T,0} = n_{S_1,0} + n_{S_2,0} \tag{7}$$

$$n_{S_T} = n_{S_1} + n_{S_2} \tag{8}$$

$$n_{S_1} = n_{S_1,0} + n_{S_1,+} \tag{9}$$

$$n_{S_2} = n_{S_2,0} + n_{S_2,+} \tag{10}$$

$$n_{S_1}^{+} \geq 0 \tag{11}$$

$$n_{S_2}^{+} \geq 0 \tag{12}$$

$$0 \leq \{x_1, x_2\} \leq 1 \tag{13}$$

$$x_1 + x_2 = 1 \tag{14}$$

Eqs. 15-16 should be used to determine the size of any equipment required by each process – noting that only standard sizes of crystalliser (i.e., 5000, 10000, 20000 L) which have a 70 v/v % fill capacity may be used. Moreover, parallel production lines (N_{PT}) should be permitted whenever a single production line proves incapable of meeting throughput requirements for a specific process. It is important to note, however, that – since other unit operations (e.g., reactions, filtrations and liquid-liquid extractions) are likely to precede crystallisation steps in production processes – this work has also stipulated that each production line designed may only complete one batch cycle per day to prevent unrealistic manufacturing processes (i.e., those which require unreasonably large numbers of batches per day) from being developed. Eqs. 17-18 are used to determine the cycle time (t_{BC}) and heating requirements of each process, whilst Eq. 19 are used to compute theoretical product recovery.

$$V_{sys} = V_{liq} + V_{cryst} \tag{15}$$

$$m_{annual} = m_{BC} N_{BC} N_{PT} \tag{16}$$

$$t_{BC} = t_{cryst} + t_{CIP} \tag{17}$$

$$Q \approx \bar{C}_{P_{S_T}} \bar{n}_{S_T} \Delta T \tag{18}$$

$$r_{API} = 100\left(\frac{n_{prod,0} - n_{prod}}{n_{prod,0}}\right) \tag{19}$$

Physical properties required to implement this model (e.g., the molecular weights, melting points and boiling points of pregabalin and any solvents) have been taken from the literature wherever possible [8, 9, 12]. Thus, established temperature dependent correlations have been used to assess the density and specific heat capacities of each solvent [13]. However, it should be noted that the initial temperature of each process should always be greater than or equal to its final temperature (Eq. 20), whilst respecting the limits specified in Eq. 21 ($T_{OS,1} = 5\,K; T_{OS,2} = 10\,K$) at all times, so as to make sure that: (i) crystallisation occurs, whilst (ii) the freezing and/or evaporation of any solvent does not [5, 6]. Likewise, it has been stipulated that (in order to avoid high shear forces during mixing) at least 4 mL of solvent should be used per gram of API being processed in each system [3].

$$T_0 \geq T_f \tag{20}$$

$$\max\begin{Bmatrix} T_{SM}^{\min} \\ T_{mp} + T_{os_1} \\ T_{phys}^{\min} \end{Bmatrix} \leq \{T_0, T_f\} \leq \min\begin{Bmatrix} T_{SM}^{\max} \\ T_{bp} - T_{os_2} \\ T_{phys}^{\max} \end{Bmatrix} \tag{21}$$

Each of the mathematical terms presented in Eqs. 1-21 are defined as described by the following statements. Term x^{sat} denotes the mole fraction (mol mol^{-1}) of API in a given solvent system when saturated, whilst terms $x_{prod,0}^{sat}$ and x_{prod}^{sat} denote the fraction of drug product dissolved in a system during the initial and final states of a crystallisation process (mol mol^{-1}). Terms x_1 and x_2 then denote the solute-free mole fraction of solvents 1 and 2 in each solvent mixture (mol mol^{-1}). A (–), B (K) and C (–) are correlation coefficients in the Apelblat equation, whilst B_0, B_1, B_2, B_3 and B_4 are CNIBS/R-K coefficients (–). Term T then denotes the temperature of each system, such that T_0 and T_f may be used to denote their initial and final temperatures. Meanwhile, terms T_{bp} and T_{mp} have been used to represent the boiling and melting point temperatures of each system. Terms $T_{SM}^{min}, T_{SM}^{max}, T_{phys}^{min}, T_{phys}^{max}$ then symbolise the minimum and maximum temperatures of the solubility (i.e., Apelblat, CNIBS/R-K) and physical property models used herein – leaving T_{os_1} and T_{os_2} to represent the temperature offsets applied to each system to prevent freezing and boiling, respectively. Term ΔT then represents the temperature difference between the initial and final states of a system. Meanwhile, terms n_{S_T} and $n_{S_T,0}$ denote the total moles of solvent used in the initial and final states of each process (mol); $n_{S_1,0}, n_{S_1}, n_{S_2,0}$ and n_{S_2}represent the initial and final amounts of any solvents and antisolvents in a system (mol); $n_{S_1,+}$ and $n_{S_2,+}$ signify the amount of solvent or antisolvent added to each

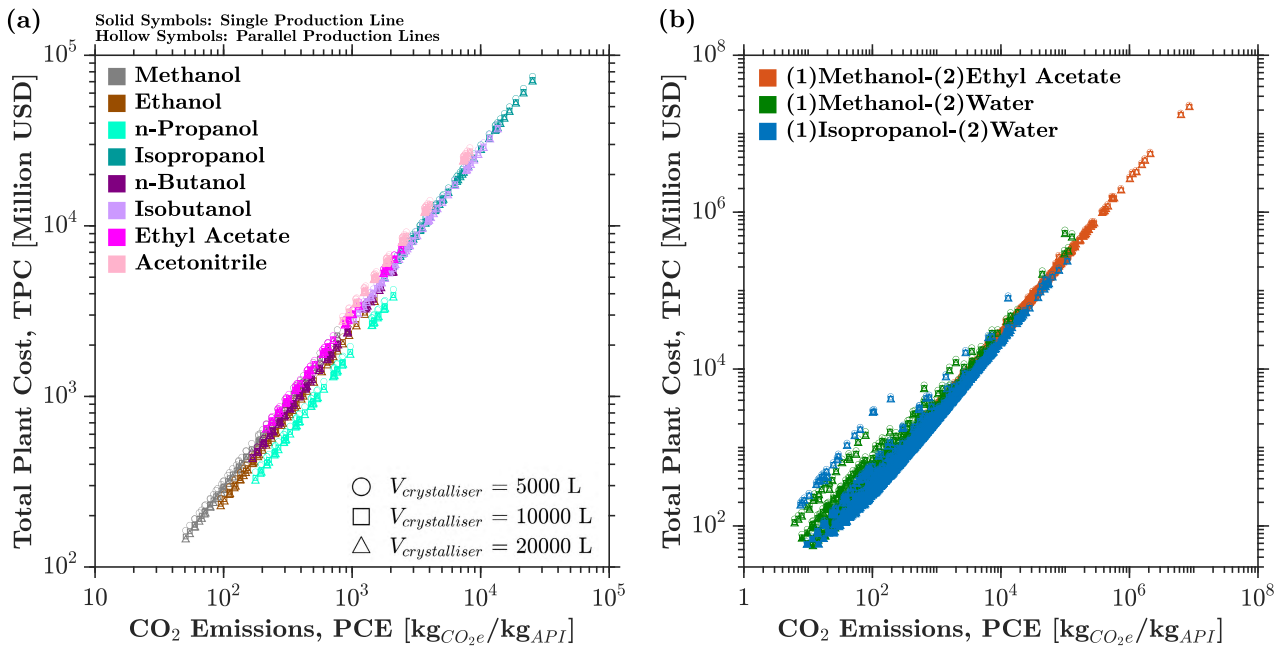

Figure 2. Relationship observed between the total cost and (scope 1 and 2) carbon emissions of each process.

system between the initial and final states (mol); $n_{prod,0}$ and n_{prod} give the amount of API product dissolved in this each system in each state (mol); and n_{prod}^{cryst} represents the amount of API crystallised during each process (mol). Meanwhile, $\bar{n}_{S_T}$ represents the average moles of solvent associated with a given system over its batch operating time. N_{BC} then denotes the annual number of batch cycles required by each production line. Symbols m_{prod}^{cryst} (kg) and M_{prod} (g mol^{-1}) then describe the mass and molar mass of the API product. Term V_{sys} then symbolises the total volume of a batch in a given process, whilst V_{liq} and V_{cryst} denote its liquid (mother-liquor) and crystal components (L). Q (kJ) then represents the heating requirements of each process, whilst r_{API} (%) gives the API recovery of each system. $\bar{C}_{P_{S_T}}$ (J mol^{-1} K^{-1}) then represents the average specific heat capacity of a system over its batch operating time.

3. CARBON EMISSIONS ANALYSIS

Carbon emission factors (PCE) have been used to assess the environmental impact of each process (Eq. 22) [7, 10] – assuming that any NO_x and SO_x generated by each process can be entirely abated by on-site scrubbing and catalytic conversion units. To allow this, it has been assumed that any waste will be incinerated and experience complete combustion, with every 1 kWh of electricity used producing 0.20705 kg of CO_2e [14].

$$PCE = \frac{m_{total\ CO_2e}}{m_{total\ prod}} \tag{22}$$

Terms $m_{total\ prod}$ and $m_{total\ CO_2e}$ in Eq. 22 denote total API production and CO_2e emissions (kg).

4. TECHNOECONOMIC ASSESSMENT

The cost of each process has been evaluated using an established costing methodology where: (i) capital costs ($CapEx$) are taken as the sum of any fixed (FCC) and working (WCC) capital costs (Eq. 23); whilst (ii) operating costs ($OpEx$) are set equal to the sum of any fixed ($FCOP$) and variable ($VCOP$) costs of production (Eq. 24).

$$CapEx = FCC + WCC \tag{23}$$

$$OpEx = FCOP + VCOP \tag{24}$$

Thus, the total cost of each process (TPC) has been calculated via Eq. 25 – assuming a plant lifetime (t_{LTP}) of 20 years and an annual discount rate (I_{dr}) of 5.5 % [10, 15].

$$TPC = CapEx + \sum_{n=1}^{t_{LTP}} \frac{OpEx}{(1+I_{dr})^n} \tag{25}$$

Within this, the FCC costs for each process have been estimated by taking the sum of any inside battery limits, outside battery limits, engineering and contingency costs – ultimately allowing any WCC costs to be calculated as 15 % of the FCC. Meanwhile, $VCOP$ costs have been calculated by assessing the raw material, waste disposal, and utility requirements of each process.

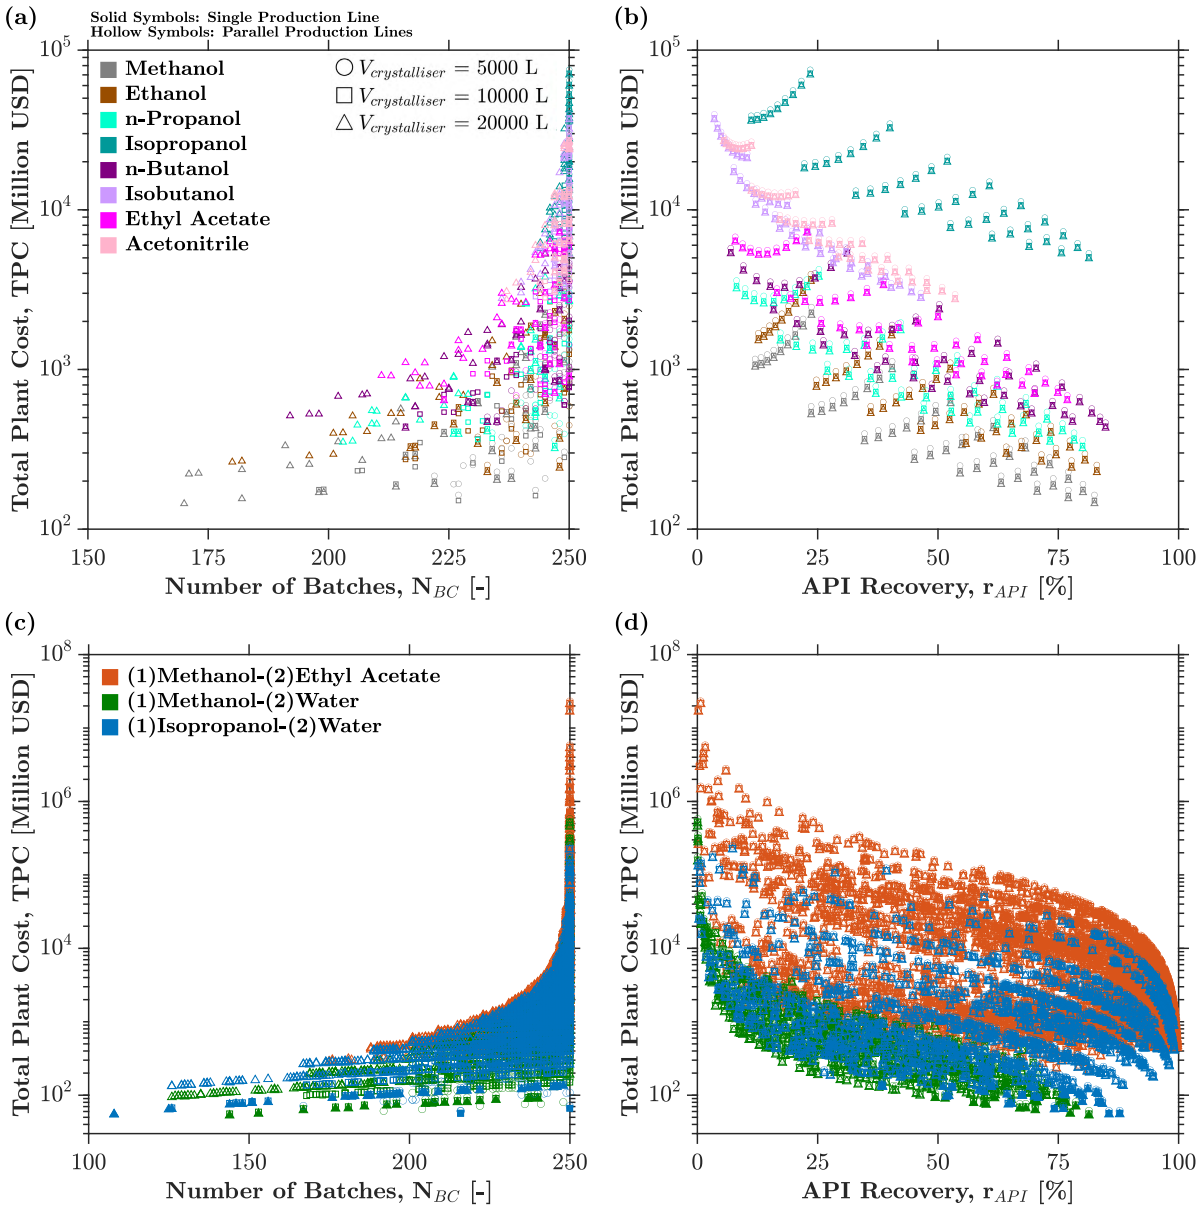

Figure 3. Impact of solvent system and crystalliser volume on the batch processing requirements (left), API recovery (right), and process costs associated with the production of pregabalin when using (a, b) pure and (c, d) mixed solvent systems to bring about its crystallisation under varying process conditions ($x_{1,0}$, $x_{1,f}$ T_0, T_f).

Any $FCOP$ costs have then been evaluated by assessing any supervision and management, labour, and maintenance costs incurred alongside any costs arising from property taxes, rental payments, insurance premiums and direct salary overheads.

Considering this, an energy efficiency of 40 % has been assumed for any electrical equipment designed in his work. Meanwhile, any raw materials required by each process have been priced using: (i) average vendor prices published by market research firms (e.g., the IMARC group [16]); and (ii) national import and export reports across major markets. Whereas waste disposal costs have been calculated in line with recommendations made by the World Health Organization (1999) [12] and Jolliffe and Gerogiorgis (2017) [11].

5. DISCUSSION OF RESULTS

The modelling framework developed in Sections 2-4 has been used to map the design space associated with each crystallisation process considered herein. From this, it has been shown that for each solvent system: (i) the

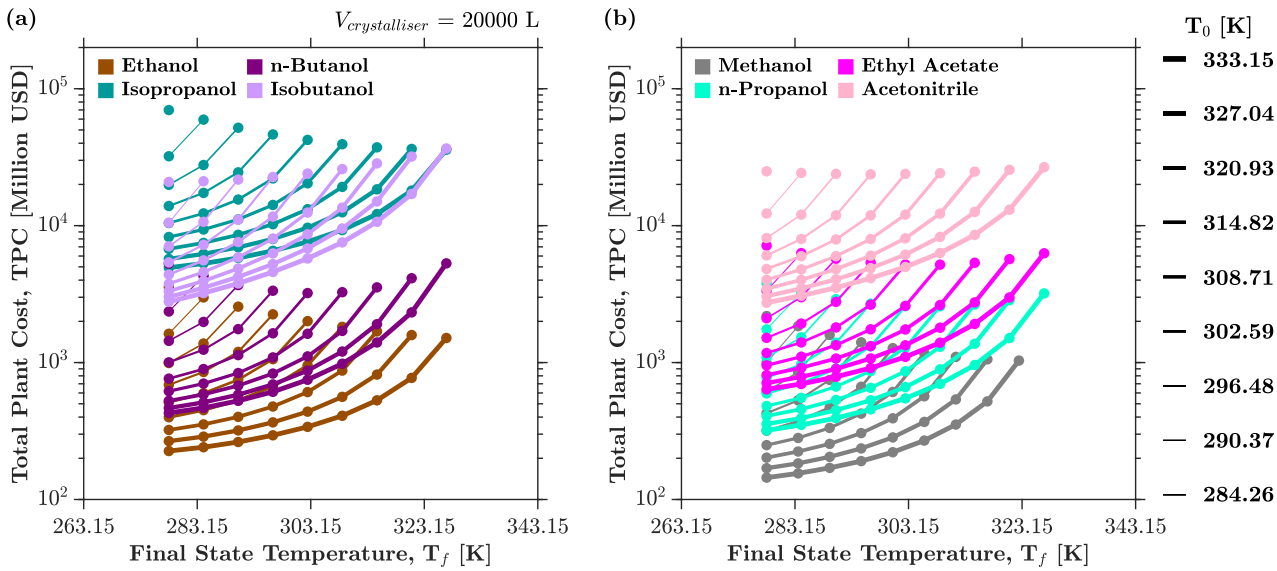

Figure 4. Impact of T_0 and T_f on the total cost of each pure solvent system.

cost of crystallisation is proportional to its expected carbon emissions, and (ii) its total costs drop when larger crystallisers are used (Fig. 2). Hence, crystalliser volume should be maximised (i.e., set to 20000 L) to simultaneously minimise process costs and environmental impact. Fig. 3 has been used to highlight this concept by demonstrating that increasing crystalliser volume: (i) reduces process costs, and (ii) simplifies manufacturing strategy (i.e., the number of parallel production lines required by a process, and the number of batches that these lines must produce to meet throughput targets). It is, therefore, clear that maximising crystalliser volume leads to cheaper, faster, and simpler pregabalin production.

Upon fixing the crystalliser volume for each solvent system at 20000 L, relationships between process temperature (T_0, T_f) and process cost can be determined for each of the pure solvent systems (Fig. 4). Meanwhile, the impact of any changes in process temperature (T_0, T_f) and the initial ($x_{1,0}$ and $x_{2,0}$) and final ($x_{1,f}, x_{2,f}$) compositions of the mixed solvent systems can be established (Figs. 5-6). In this regard, identical trends can be seen for each of the pure solvent systems whereby maximising their initial temperature whilst minimising their final temperature minimises costs (Table 1) – with methanol offering the most cost-effective process overall on account of its high API solubility (Fig. 1) and low purchase cost [16]. However, much more complicated trends exist when dealing with the mixed solvent systems – specifically because their cost depends strongly upon their initial and final solvent compositions (Fig. 5). For example, looking at the (1) methanol-(2) ethyl acetate system in Fig 5, it can be concluded that maximising its initial temperature whilst using an initial solvent composition ($x_{1,0}$) of either 0.0 or 1.0 mol mol^{-1} will yield reduced process costs regardless of the final temperature used. On the other hand, when looking at the (1) methanol-(2) water system, it can be seen that (again, regardless of the final temperature used) employing an initial solvent composition ($x_{1,0}$) of 0.30 mol mol^{-1} whilst maximising its initial temperature gives lower costs. Meanwhile, the (1) isopropanol-(2) water system benefits from maximising the initial temperature and using an initial solvent composition ($x_{1,0}$) of 0.10 mol mol^{-1}. Interestingly, however, it can be seen that, for all three systems, setting the initial and final solvent compositions equal to one another generally reduces costs. To illustrate this point, Fig. 6 has been used to examine the impact of changing the final temperature and solvent composition on each system more closely. From this, it can be seen that minimising the final temperature of each mixed solvent system minimises costs (just as seen for the pure solvent systems), whilst maintaining a constant solvent composition throughout the process also reduces costs. This is because introducing fresh antisolvent to each system raises operating costs whilst often only marginally reducing API solubility.

Considering this, optimal process conditions for each of the pure and mixed solvent systems could be identified as shown in Table 1. Hence, it could be deduced that the mixed solvent systems containing (1) methanol-(2) water and (1) isopropanol-(2) water will likely offer the lowest process costs, whilst methanol is expected to be the best performing pure solvent. However, as can be seen rom Table 1, the two mixed solvent systems greatly outperform methanol in terms of their overall cost and carbon emissions. Interestingly, however, the (1) methanol-(2) water system is expected to offer lower manufacturing costs but higher emissions than the (1) isopropanol-(2) water system. In fact, upon

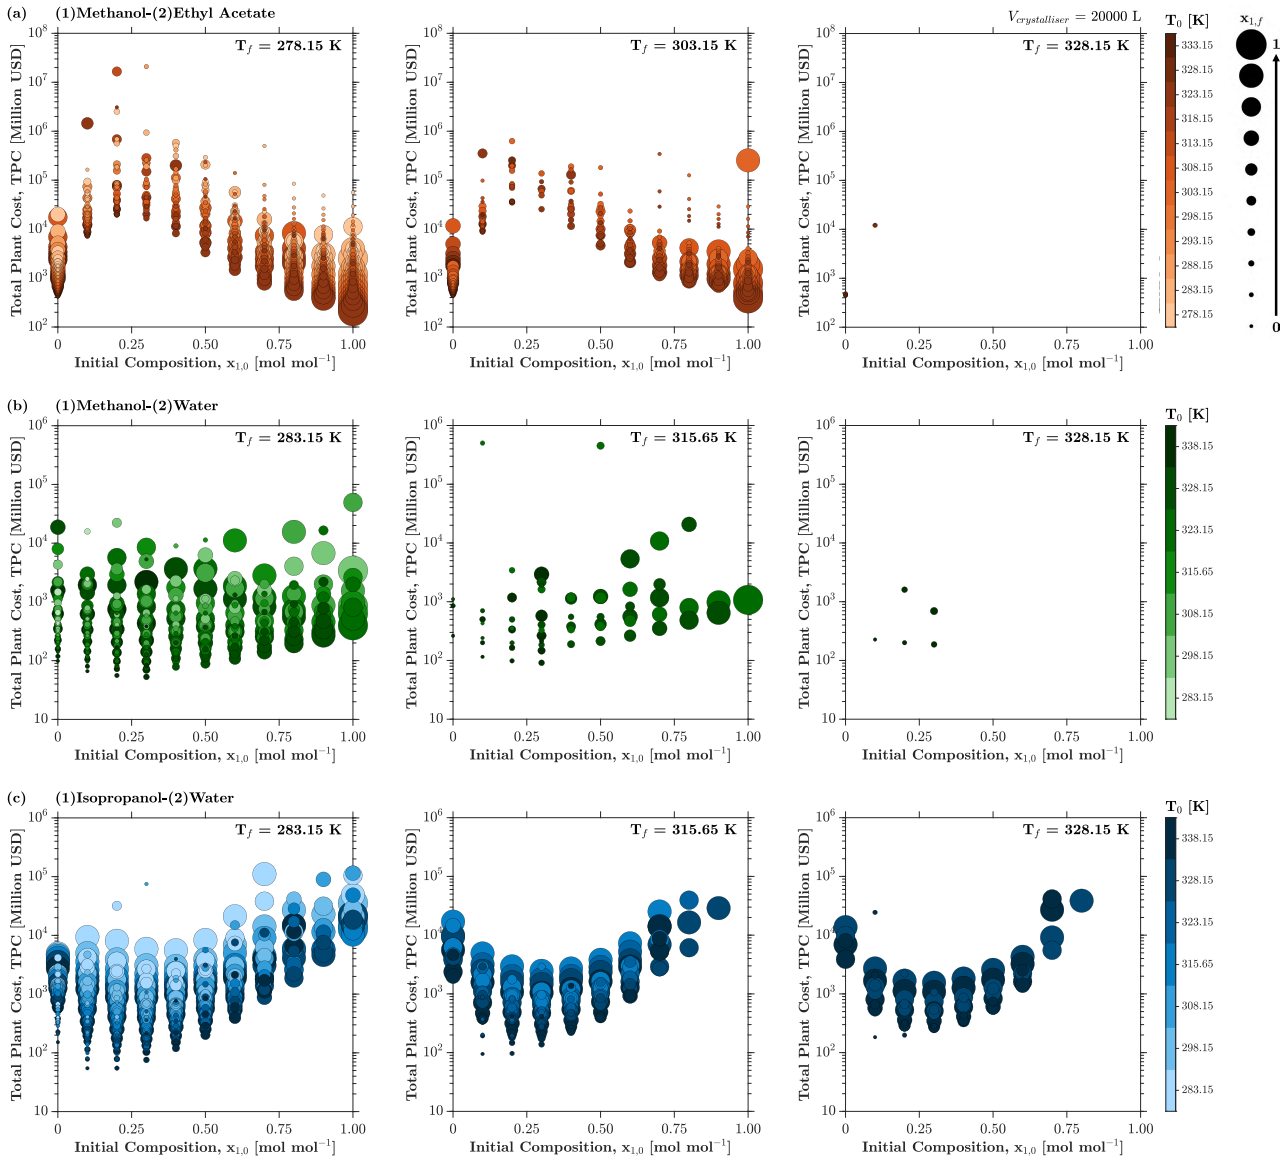

Figure 5. Impact of solvent composition ($x_{1,0}$ and $x_{1,f}$) and operating temperature (T_0 and T_f) on the total cost of each mixed solvent system.

closer inspection of the results in Table 1, it can be seen that the (1) methanol-(2) water system will only offer a saving of 3.21 % when compared with the (1) isopropanol-(2) water system, yet will emit 25.13 % more CO_2e. Thus, the (1) isopropanol-(2) water system should be selected as the best performing system.

It is important to note, however, that these results have been gathered using thermodynamic principles alone, and do not consider the crystallisation kinetics of each system as a consequence. Hence, they should only be used as a guide to rank promising systems before conducting experimental trials – thereby streamlining experimental campaigns.

6. CONCLUSIONS

This study has demonstrated that pharmaceutical crystallisation processes can be visualised in silico using thermodynamics principles alone for pure and mixed solvent systems. Further to this, it has shown how this approach can be used to identify promising process configurations (and operating conditions) when presented with a wide range of potential manufacturing routes for a new drug. Indeed, it has then used this framework to identify cost-optimal and environmentally friendly manufacturing routes for pregabalin production by carefully considering a wide range of process conditions ($x_{1,0}, x_{1,f}$ T_0, T_f) and solvent systems (i.e., 9 pure solvents, 3 mixed solvents): ranking each of these routes from best to worst based on their expected emissions and total manufacturing costs. This ultimately resulted in a mixed solvent system (i.e., (1)

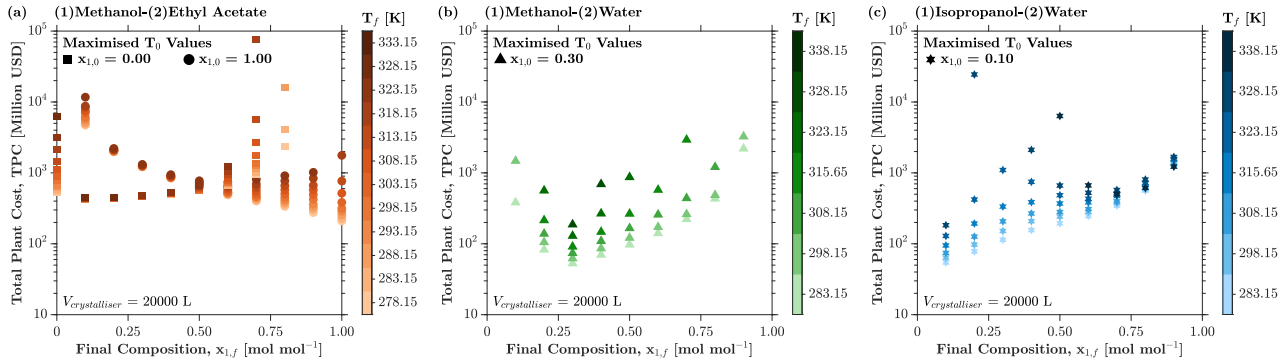

Figure 6. Impact of $x_{i,f}$ and T_f on process cost for each mixed solvent system, using maximised T_0 and optimised $x_{i,0}$.

Table 1: Cost-optimal process conditions for each crystallisation process ($V_{crystalliser}$ = 20000 L).

	T_0 [K]	T_f [K]	$x_{1,0}$ [mol/mol]	$x_{1,f}$ [mol/mol]	*PCE* [kg_{CO_2e}/kg_{API}]	*TPC* [Million USD]
(1) Methanol-(2) water	**328.15**	**283.15**	**0.30**	**0.30**	**11.60**	**52.78**
(1) Isopropanol-(2) water	**328.15**	**283.15**	**0.10**	**0.10**	**9.27**	**54.53**
Methanol	333.15	278.15	–	–	50.70	144.44
(1) Methanol-(2) ethyl acetate	318.15	278.15	1.00	1.00	58.01	210.09
Ethanol	333.15	278.15	–	–	94.89	226.02
n-Propanol	333.15	278.15	–	–	176.92	319.26
n-Butanol	333.15	278.15	–	–	167.98	429.01
Ethyl acetate	333.15	278.15	–	–	214.47	637.34
Acetonitrile	333.15	278.15	–	–	843.98	2734.47
Isobutanol	333.15	278.15	–	–	1029.57	2784.17
Isopropanol	333.15	278.15	–	–	1785.30	4935.44

isopropanol-(2) water) being selected as the most promising process available (c.f. Table 1).

More importantly, however, this work has served to demonstrate how the enhanced solubility trends often seen for solvent mixtures (Fig. 1) do not promise that antisolvent crystallisation will yield cost-optimal results. In fact, it shows that adding an antisolvent to a system could increase its process costs quicker than it decreases an API's solubility, ultimately causing the volume of solvent available for the API to dissolve into to increase faster than API recovery. As such, it is clear that new crystallisation processes should be studied in-silico (using models such as those proposed herein) prior to conducting experimental trials if one wishes to streamline their experimental campaigns by reducing their search space.

ACKNOWLEDGEMENTS

M. Blair acknowledges financial support from an EPSRC Industrial CASE PhD Fellowship. D. I. Gerogiorgis acknowledges financial support from UKRI/EPSRC (RAPID: ReAl-time Process Modelling and Diagnostics: Powering Digital Factories, EP/V028618/1) and Royal Society International Exchanges Programme (IES\R2\232014, 2023–25) grants.

REFERENCES

1. Kavoussi R. Pregabalin: from molecule to medicine. European Neuropsychopharmacology 16:S128-S133 (2006). https://doi.org/10.1016/j.euroneuro.2006.04.005
2. National Health Service (NHS). Pregabalin. Available at: https://www.nhs.uk/ (accessed: Nov. 28, 2025)
3. Ben-Menachem E. Pregabalin pharmacology and its relevance to clinical practice. Epilepsia 45:13-18 (2004). https://doi.org/10.1111/j.0013-9580.2004.455003.x
4. Global Market Insights. Pregabalin market - by formulation, by application, by distribution channel & global forecast, 2025 – 2034 (2025). Available at: https://www.gminsights.com/

5. Watson OL, Jonuzaj S, McGinty J, Sefcik J, Galindo A, Jackson G, Adjiman CS. Computer aided design of solvent blends for hybrid cooling and antisolvent crystallization of active pharmaceutical ingredients. Org. Process Res. Dev. 25:1123-1142 (2021). https://doi.org/10.1021/acs.oprd.0c00516
6. Sheikholeslamzadeh E, Chen CC, Rohani S. Optimal solvent screening for the crystallization of pharmaceutical compounds from multisolvent systems. Ind. Eng. Chem. Res. 51:13792-13802 (2012). https://doi.org/10.1021/ie3014742
7. Blair M, Gerogiorgis DI. Solvent selection, sustainability analysis, technoeconomic evaluation and optimisation of batch cooling crystallisation for flurbiprofen production. Computers & Chemical Engineering 199:109116 (2025). https://doi.org/10.1016/j.compchemeng.2025.109116
8. Cogoni G, de Souza B, Croker DM, Frawley PJ. Solubility of (s)-3-(aminomethyl)-5-methylhexanoic acid in pure and binary solvent mixtures. J. Chem. Eng. Data 61:587-593 (2016). https://doi.org/10.1021/acs.jced.5b00736
9. Tabassum K, Huang X, Yifu Z, Feng Y, Lamichhane S, Wang T, Wang N, Hao H. Solubility, hansen solubility parameter, molecular dynamic simulation and thermodynamic analysis of pregabalin in eight pure solvents and one binary solvent at various temperature. Journal of Molecular Liquids 433:127855 (2025). https://doi.org/10.1016/j.molliq.2025.127855
10. Blair M, Molaei Chalchooghi M, Cox RJ, Gerogiorgis DI. Design space visualization and technoeconomic evaluation of a batch manufacturing process for the green production of an anti-cancer drug (adavosertib) precursor. Org. Process Res. Dev. 29:2178-2199 (2025). https://doi.org/10.1021/acs.oprd.5c00056
11. Diab S, Gerogiorgis DI. Technoeconomic evaluation of multiple mixed suspension-mixed product removal (MSMPR) crystallizer configurations for continuous cyclosporine crystallization. Org. Process Res. Dev. 21:1571-1587 (2017). https://doi.org/10.1021/acs.oprd.7b00225
12. Kim S, Chen J, Cheng T, Gindulyte A, He J, He S, Li Q, Shoemaker BA, Thiessen PA, Yu B, Zaslavsky L, Zhang J, Bolton EE. Pubchem 2023 update. Nucleic Acids Research 51:D1373-D1380 (2022). https://doi.org/10.1093/nar/gkac956
13. Yaws, C. L., Chemical Properties Handbook: Physical, Thermodynamic, Environmental, Transport, Safety, and Health Related Properties for Organic and Inorganic Chemicals. McGraw-Hill (1999)
14. Department for Energy Security and Net Zero. Greenhouse gas reporting: conversion factors 2024 (2024). Available at: https://www.gov.uk/
15. Towler, G. P., Sinnott, R. K. Chemical Engineering Design: Principles, Practice and Economics of Plant and Process Design, 3rd Edition. Butterworth-Heinemann (2022)
16. IMARC Group. Pricing report 2024: price trend, chart, market analysis, news, demand, historical and forecast data (2024). Available at: https://www.imarcgroup.com/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Molecular Similarity Coefficient in Chemical Design and Analysis

Youquan Xu[a], Zhijiang Shao[a], Abdulelah S. Alshehri[b], Mansour S. Alhoshan[b], Anjan K. Tula[a*]

[a] Key Laboratory of Industrial Control Technology, College of Control Science and Engineering, Zhejiang University, Hangzhou 310027, China
[b] Chemical Engineering Department, College of Engineering, King Saud University, Riyadh 11421, KSA
* Corresponding author: anjantula@zju.edu.cn.

ABSTRACT

Computer-aided molecular design (CAMD) is an efficient product design method that is gradually attracting attention at present. It mainly uses data mining technology to extract information from the existing chemical molecular data and use this information to generate potential excellent molecules. However, the key that CAMD can truly provide accurate and reliable results lies in the efficient utilization of chemical data. In this paper, a series of chemical data analysis methods based on molecular similarity are proposed to enhance the data utilization efficiency of CAMD, which mainly includes 3 applications: adaptive modeling, reliability assessment and advanced data preprocessing including molecular recommendation, data consistency test and data augmentation. We propose specific methodology for each application, and use multiple cases to verify the effect. The results show that molecular similarity can help to improve the accuracy of property prediction at the data level, provide quantification for property prediction reliability, recommend potential excellent molecules with similar structures, locate data mistakes and perform reliable data augmentation, finally enhancing the data utilization efficiency of CAMD.

Keywords: Molecular design, Similarity, Reliability quantification, Data preprocessing, Property prediction

1. INTRODUCTION

The rapid development of industry has put forward higher requirements for efficient chemical product design. Due to the time-consuming, labor-intensive and easy to miss the potential excellent molecules in the traditional expert-based product molecular design, and the rapid development of modern computer technology, computer-aided molecular design (CAMD) has gradually attracted attention. CAMD mainly includes two parts: property prediction and molecular structure optimization. In these two parts, many researchers have achieved a lot of outcomes. In property prediction, a variety of models has been proposed to improve the prediction accuracy, such as group contribution models [1], machine learning-based models [2]and neural network-based models [3]. In molecular structure optimization, researchers have proposed numerous applications for different products such as fragrance [4], solvent [5] and refrigerant [6].

The core of the above advances in CAMD is to extract the characteristics of molecules and analyze the relationships between molecules. On the other hand, molecular similarity is an indicator used to measure the correlation between molecules, so it can play an important role in the data analysis of CAMD. The calculation of molecular similarity mainly includes two parts: molecular fingerprint and similarity measure method. Molecular fingerprint is used to convert molecular structure into mathematical description such as vector, and then the similarity measure method is used to quantify the molecular similarity. Different molecular fingerprints and similarity formulas adapt to different scenarios. Since the similarity of molecular structures often leads to the similarity of properties, there are many applications of chemical data analysis based on molecular similarity, such as database searching and virtual screening, including drug discovery for anti-cancer agents [7], COVID-19 therapeutics [8], and constrained optimization of drug like molecules [9].

Although some progress has been made in CAMD and molecular similarity, the current application of molecular similarity mainly focuses on the retrieval of known datasets. There are still many unexplored uses of

molecular similarity for CAMD that need to design potential novel molecules. Therefore, this paper proposes 3 main applications of molecular similarity to enhance CAMD: (i) adaptive modeling, (ii) reliability assessment, and (iii) advanced data preprocessing.

2. SIGNIFICANCE OF MOLECULAR SIMILARITY IN CAMD

In the molecular structure chemical space, closer molecules tend to present similar properties or follow similar rules, so these similar molecules can provide more valuable information to each other, that is, molecular similarity can be used as a measure of the value of information between two molecules. This provides a basis for molecular similarity to enhance data analysis in CAMD.

Since group contribution (GC) method provides the important advantages of simplicity, quick property predictions without requiring substantial computational resources, and being convenient to be integrated into optimization formulation, this paper uses the GC method [10] to describe the molecular structure, that is, the molecular structure M can be expressed as follows:

$$M = [n_1, n_2, n_3, \cdots \cdots, n_{220}] \tag{1}$$

The molecular structure M can be represented as a 220-dimensional vector, and each element n_i of this vector represents the number of the ith group contained in M. The GC method can also achieve compatibility with the existing CAMD method. In order to more accurately measure information value between molecules, this paper uses the molecular similarity coefficient (MSC) proposed by Xu et al. [11] in this paper. Here, the value of MSC ranges from [0, 1], where closer to 1 means two molecules are more similar, and closer to 0 means less similar. An MSC value of 1 indicates that the two molecules are exactly the same, and an MSC value of 0 indicates that two molecules do not have any common groups, namely, cannot provide any valuable information between each other. This MSC integrates molecular chemical characteristics and can effectively deal with issues such as molecular type and molecular size, so as to more accurately measure the value of information that molecules can provide to each other, which is the most critical property of molecular similarity for the data analysis applications proposed in this paper. Next, we discuss the application of molecular similarity: adaptive modeling, reliability assessment, and enhanced data preprocessing in detail.

3. THE APPLICATION OF MOLECULAR SIMILARITY IN CAMD

3.1. Adaptive modeling

The first application of molecular similarity to CAMD data analysis is adaptive modeling. The main modeling work in CAMD is the molecular property prediction model. The basic flowchart of this adaptive modeling method for property prediction is shown in Figure 1.

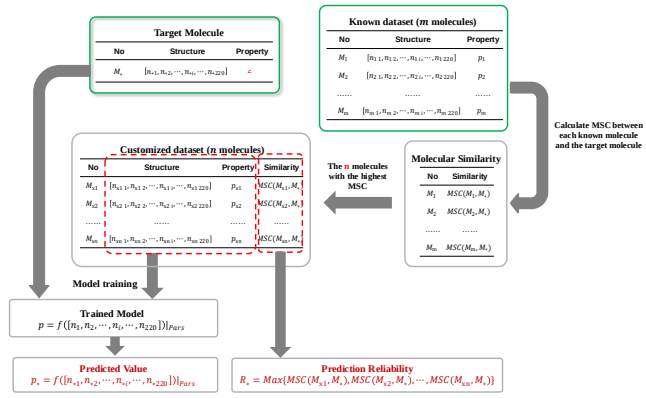

Figure 1. Adaptive property modeling method

As can be seen in Figure 1, for a target molecule M_*, the MSC between it and all the m molecules $(M_1, M_2, \ldots\ldots, M_m)$ in the known dataset are calculated, those are, $MSC(M_1, M_*), MSC(M_2, M_*)$, ..., $MSC(M_m, M_*)$. Then the n molecules with the highest MSC to M are screened out as the customized training set of M, those are $MSC(M_{x1}, M_*), MSC(M_{x2}, M_*)$, ..., $MSC(M_{xn}, M_*)$, which are used for the model training and the property prediction of M, that is p_*.

General machine learning methods will directly use all available known molecular dataset to train the model, so as to obtain a property prediction model with fixed parameters, which is then used to predict the properties of all molecules. However, this approach lacks targeted utilization of dataset. The ultimate goal of the property prediction model is to predict the properties of new molecules individually. In this process, we are most concerned about whether the property prediction of each molecule is accurate. Moreover, when all the known molecular data are used for training, the molecules dissimilar to the target molecule will also be used, and the influence of these molecules on the model will affect the property prediction accuracy for the target molecule. Therefore, it necessitates a more targeted method to predict the property of a single target molecule, that is, adaptive modeling.

The basic idea of adaptive modeling method is that when predicting the properties of a target molecule, we can first construct a customized training dataset for the target molecule according to the similarity between the target molecule and the known molecules, and use the customized training dataset to train the model and predict the properties of the target molecule. Through this method, we can use the information most relevant and valuable to the target molecule for modeling. Different target molecules will adaptively customize their own

training dataset, so as to improve the accuracy of property prediction from the data level. This prediction accuracy improvement has been proved by Xu et al. 's research through a variety of property datasets, and the detailed tests can refer to their work [11].

3.2. Reliability Assessment

Another application of molecular similarity lies in reliability assessment. In CAMD, the molecular evaluation is often performed based on the predicted properties. However, the property prediction error is inevitable and it will directly affect the judgment of the molecular performance and even lead to a large difference between the predicted property of the designed molecule and its experimental verification value, losing the guiding significance for experiment. In addition, the property prediction error can only be obtained after experimental verification, which means that high costs have been paid. Therefore, it is important to assess the prediction reliability before experimental verification. To this end, this paper proposes a novel application of molecular similarity in reliability assessment.

The basic idea of this method is to quantitatively evaluate the property prediction reliability of a target molecule according to its information correlation with its training information. The flowsheet is also shown in Figure 1. In section 3.1, we have introduced customizing the training data set according to molecular similarity. Here, we can still employ this concept. Since these customized training data can provide the most valuable information for the target molecule, we can use the comprehensive relevance of this information to design a reliability quantification index for the target molecule. If the comprehensive similarity between the customized training dataset and the target molecule is large, the more information value these training molecules can provide for the target molecule, and the more reliable the property prediction of the target molecule will be. Conversely, if the training molecules customized by similarity are still molecules very different from the target molecule, the information value of these molecules for the target molecule will be very low, and the unreliability of the prediction accuracy will be higher. Here, we choose the maximal MSC of the customized training molecules as the comprehensive similarity, which is used to characterize the comprehensive information value of the customized training data, that is, the reliability index of the target molecular prediction, which is defined as R. As shown in Figure 1, for a target molecule M_*, its R is calculated according to Eq. (2),

$$R_* = Max\{MSC(M_{x1}, M_*), MSC(M_{x2}, M_*), \cdots, MSC(M_{xn}, M_*)\} \quad (2)$$

The value range of R is also [0, 1], the closer R is to 1, the more reliable the prediction is, and the closer R is to 0, the less reliable the prediction is. If a molecule has an R of 1, it is a molecule already existing in the known dataset.

Then, we need to verify if R can indeed be used as a quantitative indicator of prediction reliability. In this regard, Xu et al. have verified the performance of the reliability index R in quantification of prediction reliability through tests in a variety of property datasets. Specific testing details can be referred to their work [11].

3.3. Advanced data preprocessing

Another application of molecular similarity lies in enhancing data preprocessing. This paper introduces 3 applications of molecular similarity in data preprocessing, which are molecular recommendation, data consistency test and data augmentation.

3.3.1. Molecular Recommendation

In molecular design, the recommendation of potential excellent molecules is very meaningful, which can help to narrow the search scope of molecular chemical space, and give priority to the performance evaluation of some molecules, so as to find excellent molecules more efficiently. In the process of molecular design, we often have known molecules with excellent properties at hand. Since molecules with similar structures often show similar properties, we can use these known excellent molecules as reference molecules, and then recommend potential excellent molecules based on molecular similarity. Specifically, molecular recommendation can be divided into single reference molecule-based and multiple reference molecules-based, as shown in Figure 2. Then they are introduced separately as follows.

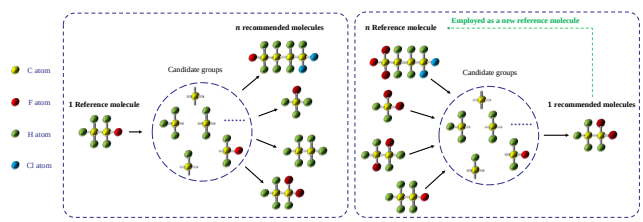

Figure 2. The flowsheet of molecular recommendation

From the figure, in the process of molecule generation through the candidate groups, we can generate molecules similar to the reference molecule by adding the similarity constraint between the generated molecule and the reference molecule. In addition, molecular recommendation in CAMD cannot be limited to screening molecules from known molecular dataset, and it is necessary to break through this limitation to achieve the recommendation of potential novel molecular structures. Therefore, the proposed molecular recommendation method is based on the concept of optimization, that is, generating the optimal potential recommended molecular structure that best satisfies the constraints. The specific molecular

Table 1: The recommended molecules using a single reference molecule

No	CH_3	CH_2	CH	C	COOH	MSC	Name	Application
	2	14	1	0	1	1.00	Isooctadecanoic acid [a]	Cosmetics, emollient
1	2	15	1	0	1	0.94		
2	2	13	1	0	1	0.93		
3	2	16	1	0	1	0.88	2-octyldodecanoic acid [b]	Emulsifier, Lubricant
4	2	12	1	0	1	0.87	2-Hexyldecanoic acid [c]	Emulsifier, Lubricant
5	2	17	1	0	1	0.83		
6	2	11	1	0	1	0.80		
7	2	18	1	0	1	0.79		
8	2	19	1	0	1	0.75		
9	2	10	1	0	1	0.73	2-hexyloctanoic acid [d]	Cosmetic agents
10	2	20	1	0	1	0.71		

a. https://www.specialchem.com/cosmetics/inci-ingredients/isostearic-acid?lr=beyable
b. https://parchem.com/chemical-supplier-distributor/2-octyl-dodecanoic-acid-040581?q=&_pos=2&_sid=64799ef8f&_ss=r
c. https://parchem.com/chemical-supplier-distributor/2-hexyldecanoic-acid-085101?q=&_pos=1&_sid=05130b18d&_ss=r
d. https://www.thegoodscentscompany.com/data/rw1446781.html#tosafiu

recommendation optimization model is as follows:

Objective:

$$\text{Max: } Obj_{dummy} = 1 \tag{3}$$

Subject to:

$$M = [n_1, n_2, n_3, \cdots \cdots, n_{220}] \tag{4}$$

$$\sum(\upsilon_j - 2)n_j + 2q = 0 \tag{5}$$

$$\sum_{i \neq j} n_i \geq (\upsilon_j - 2)n_j + 2 \quad j = 1, 2, 3, \cdots, 220 \tag{6}$$

$$ng_{lower} \leq \sum n_j \leq ng_{upper} \tag{7}$$

$$nf_{lower} \leq \sum n_j - n_{CH3} - n_{CH2} - n_{CH} - n_C \leq nf_{upper} \tag{8}$$

$$Mw_{lower} \leq \sum(Mw \times n_j) \leq Mw_{upper} \tag{9}$$

$$n_{lower} \leq n_j \leq n_{upper} \tag{10}$$

$$MSC(M, M_{reference}) \geq 0.7 \tag{11}$$

Here, υ_j and n_j refer to the valency and number of the jth group. Eqs. (5)-(10) are structural constraints about the octet rule and size related boundary. Eq. (5) indicates that the sum of the number of free electrons of all groups constituting a molecule is -2q. q is a constant, q=1 for acyclic molecules, q=0 for monocyclic molecules, and q=-1 for bicyclic molecules. Eq. (6) states that the total number of free electrons across all instances of the jth group must not exceed that of all other groups, ensuring a neutral molecule with a total charge of zero. Eq. (7) limits the total number of groups in the molecule to be between ng_{lower} and ng_{upper}. Eq. (8) limits the total number of functional groups in the molecule to be between nf_{lower} and nf_{upper}. Eq. (9) ensures that the molecular weight falls within the range of Mw_{lower} and Mw_{upper}. Eq. (10) limits the number of the jth group to be between n_{lower} and n_{upper}. The constraint shown in Eq. (11) limits that only molecules whose similarity to the reference molecule is larger than 0.7 can be recommended. Finally, we define a dummy objective function shown in Eq. (3) to generate all the feasible solutions. Consequently, the molecular recommendation model can be solved through NIP solver and output all feasible molecules.

Then we show a case. This case is about emollient molecule recommendation. In order to design new emollient molecules, we need to recommend some potential molecules which have a high probability of having excellent performance to be emollient. Besides, we already have some known emollient molecules. Therefor, in this case, we just use a known excellent emollient molecule as a reference to recommend some potential emollient molecules through the proposed method. The reference molecule employed here is Isooctadecanoic acid. The recommended molecules are shown in Table 1:

From the table, we can see that the MSC between these molecules and Isooctadecanoic acid remains above 0.7, and we find that among the 10 recommended molecules, molecules 3 (2-octyldodecanoic acid), molecule 4 (2-Hexyldecanoic acid), and molecule 9 (2-hexyloctanoic acid) are also excellent emollient molecules, which proves that the proposed molecular recommendation method can effectively generate molecular with similar excellent performance to the reference molecule.

However, the single reference molecule-based recommendation has a shortcoming, that is, the structural

diversity of the recommended molecules is relatively limited. For example, in the above cases, the generated molecules are all homologs of the target molecule. Therefore, we need to use multiple structurally diverse reference molecules, so as to achieve more diverse molecular recommendation. The corresponding molecular recommendation optimization model is shown below:

Objective:

$$\text{Max: } \sum_i MSC(M, M_{reference_i}) \tag{12}$$

Subject to:

$$M = [n_1, n_2, n_3, \cdots\cdots, n_{220}] \tag{13}$$

$$\sum(v_j - 2)n_j + 2q = 0 \tag{14}$$

$$\sum_{i \neq j} n_i \geq (v_j - 2)n_j + 2 \; j = 1, 2, 3, \cdots, 220 \tag{15}$$

$$ng_{lower} \leq \sum n_j \leq ng_{upper} \tag{16}$$

$$nf_{lower} \leq \sum n_j - n_{CH3} - n_{CH2} - n_{CH} - n_C \leq nf_{upper} \tag{17}$$

$$Mw_{lower} \leq \sum(Mw \times n_j) \leq Mw_{upper} \tag{18}$$

$$n_{lower} \leq n_j \leq n_{upper} \tag{19}$$

$$MSC(M, M_{reference_i}) < 1 \; for \; \forall i \tag{20}$$

The model uses the sum of MSC between the generated molecule and all reference molecules as the objective function to generate the recommended molecule with the highest comprehensive similarity with all reference molecules, as shown in Eq. (12). At the same time, the constraint shown in Eq. (20) requires that the generated molecule is different from all reference molecules. Eqs. (13-19) are the same with Eqs. (4-10).

Then we show another case. This case is about the design of environmentally friendly refrigerants. Here, we employ several known environmentally friendly refrigerants as reference molecules to recommend some potential molecules which has high possibility of having excellent environmental-related properties such as low GWP and ODP. The reference molecules are shown in Table 2.

Table 2: List of reference refrigerants

No	SMILES
1	FCF
2	FC(F)F
3	FCC(F)(F)F
4	FC(F)C
5	FC(F)C(F)(F)F
6	C=C(F)C(F)(F)F
7	FC=CC(F)(F)F
8	CCC
9	CC(C)C
10	FC(F)C(F)(F)C
11	FCC(F)F

Through the solving of the molecular recommendation model, we can generate a molecule with the highest comprehensive similarity to these reference molecules. After that, we can continue to insert the recommended molecule into the reference molecule list and use this updated reference molecule list to recommend a new molecule again. 10 rounds of molecular recommendation are repeated according to the above steps, thus generating 10 recommended molecules as shown in Table 3.

Table 3: The recommended molecules using multiple reference molecules

No	SMILES	ASHRAE Name	Environmental Impact
1	C(F)(F)(F)F	R14	GWP=7390, ODP=0
2	FC(F)C(F)(F)C(F)(F)F		
3	FC(F)(F)C(F)(F)F	R116	GWP=12200, ODP=0
4	FC(F)C(F)(F)F	R125	GWP=3500, ODP=0
5	C=C(F)C(F)F		
6	FC=CC(F)F		
7	FCC(F)(F)F	R134a	GWP=1430, ODP=0
8	FCC(F)F		
9	FCC(F)(F)C(F)F	R245ca	GWP=693, ODP=0
10	FCCC(F)(F)C(F)F		

As we can see from the table, these molecules maintain structural similarity with all the reference molecules but also have difference. In addition, we found that among the 10 recommended molecules, molecules 1 (refrigerant R14), 3 (refrigerant R116), and 4 (refrigerant R125) have an ODP of 0, but they have a higher GWP. Molecules 7 (refrigerant R134a) and molecule 9 (refrigerant R245ca) have 0 ODP and low GWP. Therefore, this case can show that the proposed molecular recommendation method can effectively generate excellent molecules with similar performance to multiple reference molecules.

3.3.2. Data Consistency Test

Another application of molecular similarity in data preprocessing is data consistency test. The correctness of data is the basis of data analysis, and effective data mistake localization is of great significance. Here, we propose a molecular similarity-based method for data mistake localization, namely data consistency test. As we have mentioned above, using the most similar molecules to as the training set can provide the most valuable information for the target molecule, thereby improving the accuracy of property prediction and quantifying prediction reliability. Then, if the reliability index R of the target

molecule is high, the prediction accuracy should have high possibility to be high. However, if we find that the prediction error for the target molecule is very large in this situation, we have reason to suspect that the data of the target molecule and its customized training molecules may have mistakes. For example, when using the critical temperature (Tc) dataset to conduct the adaptive modeling for molecules in test dataset, we found that the prediction reliability R of a molecule $CH_3(CH_2)_{16}COOH$ in the test set was 0.9437, but the prediction error was 10.55% after comparing its predicted value with the real data record in the original dataset. Using the above analysis, we exported the property data of $CH_3(CH_2)_{16}COOH$ and its customized training molecule in Table 4.

Table 4: The customized training dataset of $CH_3(CH_2)_{16}COOH$

No	MSC	Tc (real)/K	CH_3	CH_2	COOH
$CH_3(CH_2)_{16}COOH$	1.00	799	1	16	1
Molecule 1	0.94	474.2	1	17	1
Molecule 2	0.94	590	1	15	1
Molecule 3	0.89	809	1	18	1
Molecule 4	0.88	785	1	14	1
Molecule 5	0.82	693	1	13	1
Molecule 6	0.76	568	1	12	1
Molecule 7	0.70	591.15	1	11	1
Molecule 8	0.64	618	1	10	1
Molecule 9	0.58	654	1	9	1
Molecule 10	0.52	456.86	1	8	1

As can be seen, these molecules belong to the homologue, the only difference being the number of carbon atoms. According to the law of chemistry, there is a monotonic trend in Tc and the number of carbon atoms of these molecules. Therefore, we plotted Tc versus the number of carbon atoms for these molecules in Figure 3 (a). It can be seen that there is no monotonic law, indicating that the property data of these molecules are highly likely to be erroneous. Therefore, we checked these data through ICAS software [12], and eventually, we did find some data mistakes, and the corrected data are shown in Table 5.

Table 5: The customized training dataset of $CH_3(CH_2)_{16}COOH$ (corrected)

No	MSC	Tc (real)/K	CH_3	CH_2	COOH
$CH_3(CH_2)_{16}COOH$	1.00	799	1	16	1
Molecule 1	0.94	812[a]	1	17	1
Molecule 2	0.94	793[a]	1	15	1
Molecule 3	0.89	821[a]	1	18	1
Molecule 4	0.88	785	1	14	1
Molecule 5	0.82	775[a]	1	13	1
Molecule 6	0.76	765[a]	1	12	1
Molecule 7	0.70	754[a]	1	11	1
Molecule 8	0.64	734[a]	1	10	1
Molecule 9	0.58	732[a]	1	9	1
Molecule 10	0.52	722.1[a]	1	8	1

a: checked in ICAS (Integrated Computer Aided System).

Similarly, we plot the corrected Tc versus carbon atoms as shown in Figure 3 (b).

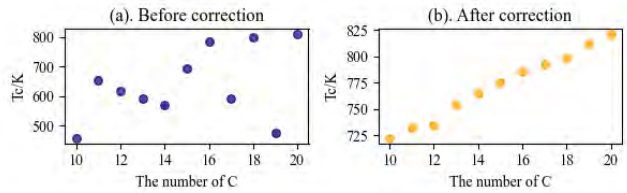

Figure 3. The data of Tc with the number of C

As we can see, there is a monotonic trend between Tc and the number of carbon atoms now. Furthermore, we use the corrected data to predict Tc of $CH_3(CH_2)_{16}COOH$, and find that the prediction error is 0.42% now, which indicates that the data have been corrected.

Based on the above case, we propose the data consistency test method. Its flowsheet is shown in Figure 4.

The basic idea of this method is to score the inconsistency of each molecule by analyzing the degree of inconsistency between each molecule and other molecules in the known dataset, and sort all molecular data according to the inconsistency score. The molecule with a higher inconsistency score has a higher probability of having data mistakes, and the more recommended to users for verification. The most critical step in the above flowsheet is how to analyze the inconsistency score for each molecular data. So, we propose a data inconsistency evaluation method based on molecular similarity, as shown in Figure 5.

First of all, we take out a molecule M_i from the known dataset. Using other known molecules as candidate training data, the training molecules of M_i are customized based on molecular similarity, and the property prediction and reliability assessment are carried out. The prediction value is compared with the property data of M_i to calculate the prediction error. Next, we set up an inconsistency scoring mechanism. Initially, all the molecules have an inconsistency score of 0. When molecule M_i satisfies the inconsistency condition, that is, when the prediction reliability of M_i is greater than 0.9 but the prediction error is greater than 5% or the prediction reliability is greater than 0.7 but the prediction error is greater than 10%, the inconsistency score of M_i and all molecules in its customized training dataset is increased by 1. According to the above method, each molecule in the known dataset is evaluated for inconsistency. Finally, the molecules with the highest inconsistency are recommended for verification of data mistake.

Next, we show a case. The case was to conduct data consistency test for a critical volume (Vc) dataset and screen out the potential data mistakes from this dataset. A part of the dataset is shown in Table 6.

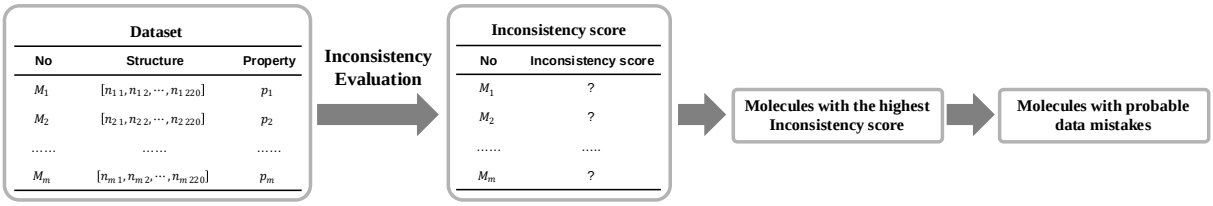

Figure 4. The workflow of Data Consistency Test

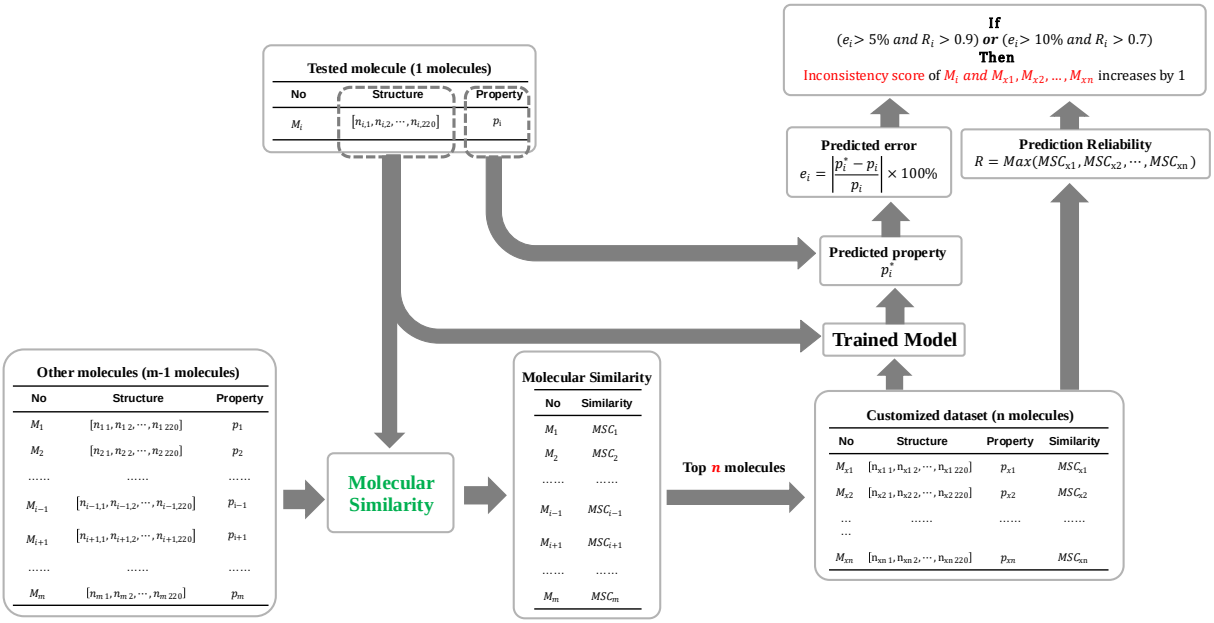

Figure 5. The algorithm of inconsistency evaluation

Table 6: A Part of the dataset of critical volumes (Vc)

No	SMILES	Vc/(mL/mol)
1	C=CCCCCCCCCCC	684
2	C=COCC	263
3	O=N(=O)c1ccccc1C	441
4	FC(F)C	179
5	CCCCOC(=O)C(=C)C	481
...	...	...

After conducting the proposed data consistency test in this dataset, the inconsistency evaluation results can be obtained, shown in Table 7.

Table 7: The result of Data Consistency Test

SMILES	Vc(mL/mol)	Inconsistency score
CCCCCCCCCCCCCCCC(=O)O	1059	4
CCCCCCCCCCCCCCCCC(=O)O	1130	4
CCCCCCCCCCCCCCCCCCCCC(=O)O	1346	3
CCCCCCCCCCCCCCCCCC(=O)O	1251	3
CCCCCCCCCCCCCCCCCCC(=O)O	1070	3
CCCCCCCCCCCCCCCCCCCCCC(=O)O	1485	2
CCCCCCCCCCCCCCC(=O)O	1002	2
CCCCCCCCCCCCCC(=O)O	811	1
CCCCCCCCCCCC(=O)O	705	1
CCCCCCCCCCC(=O)O	653	1
C=CCCCCCCCCCC	684	0
C=COCC	263	0
.....	.....	.....

From this result, we can see that there is a total of 10 molecules with an inconsistency score greater than 0, and these molecules are the molecules that need to be paid attention to. Next, we try to correct the property data of them and the result is shown in Table 8.

Table 8: The result of data correction

SMILES	Vc/(mL/mol)	Correct Vc(mL/mol)
CCCCCCCCCCCCCCCC(=O)O	1059	917
CCCCCCCCCCCCCCCCC(=O)O	1130	969
CCCCCCCCCCCCCCCCCCCCC(=O)O	1346	1130
CCCCCCCCCCCCCCCCCC(=O)O	1251	1020
CCCCCCCCCCCCCCCCCCC(=O)O	1070	1070
CCCCCCCCCCCCCCCCCCCCCC(=O)O	1485	Missing
CCCCCCCCCCCCCCC(=O)O	1002	864
CCCCCCCCCCCCCC(=O)O	811	811
CCCCCCCCCCCC(=O)O	705	705
CCCCCCCCCCC(=O)O	653	653

It can be seen that there are mistakes in the property data for a total of 5 molecules out of the 10 most inconsistent molecules. In addition, the corrected property data of CCCCCCCCCCCCCCCCCCCCCCC(=O)O is missing, which will be introduced in the next application again. This result demonstrates the effectiveness of the proposed data consistency test method. Additionally,

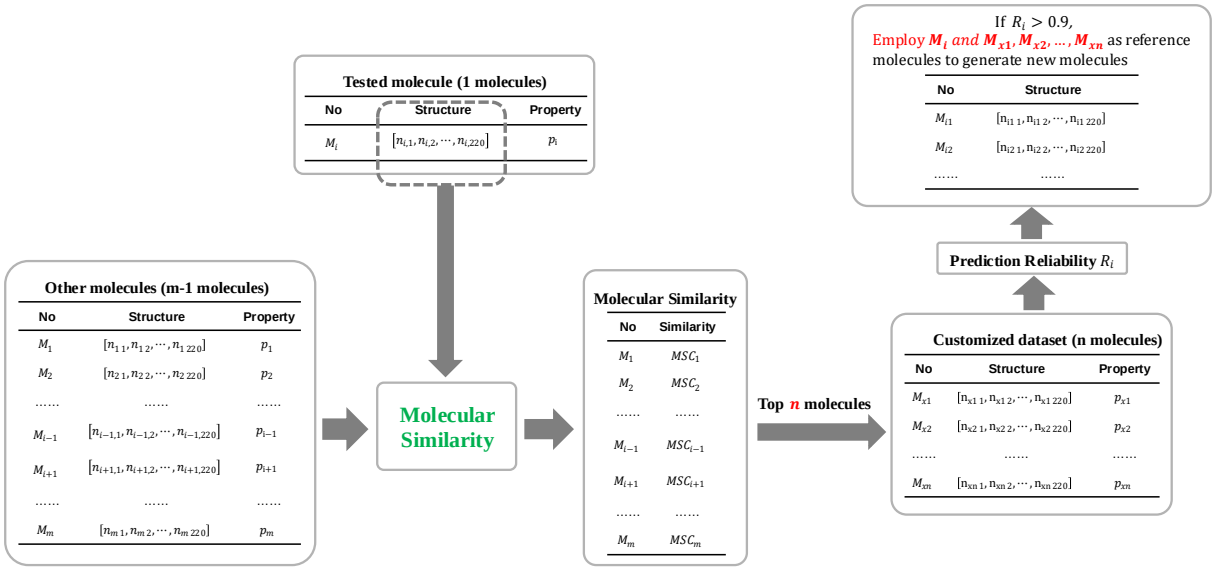

Figure 6. The flowsheet of data augmentation

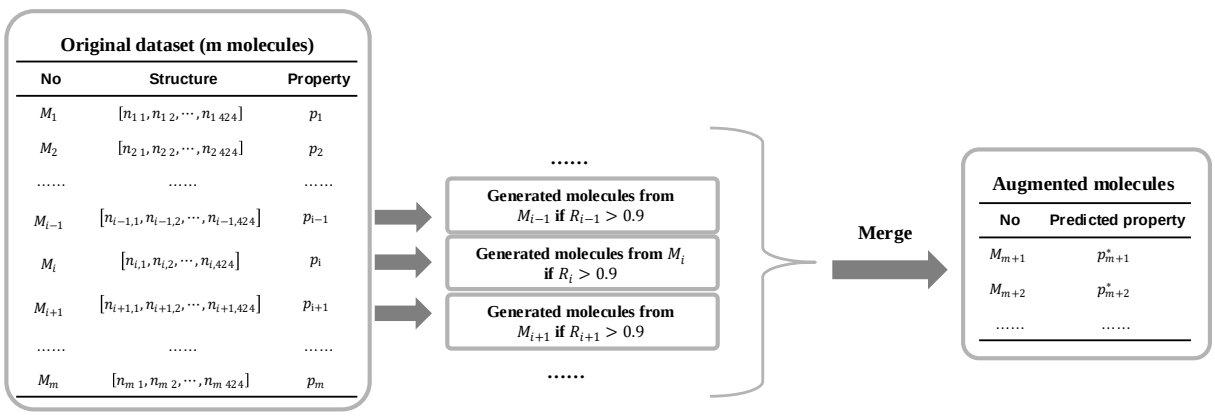

Figure 7. The merge of augmented molecules

this method essentially provides a quantitative assessment of the probability that the data has mistakes. It does not mean that the data with high inconsistency must have mistakes or the data with low inconsistency must not have mistakes. In addition, we only evaluated and counted the inconsistency of molecules with high prediction reliability. This is because when the actual prediction error of molecules with high prediction reliability is large, we have higher confidence to conclude that the data have mistakes. For molecules with low prediction reliability, their training molecules can provide low information value itself, even if the data is correct, the prediction error may still be high, so there is no great confidence to evaluate the data based on the large prediction error.

3.3.3. Data Augmentation

Data augmentation is a very important data enhancement method in chemical data analysis. For the augmentation of chemical molecular dataset, we need to generate some augmented molecules with high prediction reliability, so that these molecules can be insert in the original molecules reliably. According to the reliability assessment method introduced above, if a target molecule can collect customized training molecules with high similarity to it from the known dataset, then its property prediction reliability will be significantly high. Therefore, we can use this conclusion in turn to generate molecules with high comprehensive similarity to the known molecules, and then use the adaptive modeling method to predict the properties of them. Their prediction reliability will be very high, so, we can then use these generated molecules and their predicted properties as augmentations to the original dataset. This is the basic idea of the proposed data augmentation method.

First, we present a case. This case is about the augmentation of missing data. In the case of data consistency test, after the data correction, there is still a molecule CCCCCCCCCCCCCCCCCCCCCC(=O)O in the

dataset whose real property data is missing. However, when we use other known molecules as candidate training dataset for adaptive modeling and reliability assessment of CCCCCCCCCCCCCCCCCCCCCCC(=O)O, we can obtain a predicted Vc of 1232 mL/mol for it with a prediction reliability of 0.9 which is very high. Therefore, we can try to directly use this prediction value as an augmentation for its missing property data. To see if this method is feasible, we performed the data consistency test on the dataset. The results are shown in Table 9.

Table 9: The data consistency test result for the dataset after the data augmentation

SMILES	Vc(mL/mol)	Inconsistency score
C=CCCCCCCCCCC	684	0
C=COCC	263	0
O=N(=O)c1ccccc1C	441	0
FC(F)C	179	0
CCCCOC(=O)C(=C)C	481	0
OCCCCCCCCCC	600	0
Clc1ccc(Cl)cc1	351	0
CC=Cc1ccccc1	407	0
CCC(Cl)C	305	0
CCCCN(CCCC)CCCC	735	0
.....	.....	.....

It can be seen that the inconsistency scores of all the molecular data are 0, which shows the consistency of the dataset and the feasibility of the above augmentation method for missing data.

Based on the results of the above case and the basic idea, we propose the data augmentation flowsheet shown in Figure 6. For a known molecular dataset, we first take out one of the molecules M_i, and use other known molecules except M_i as candidate training dataset to conduct reliability assessment. If the reliability of M_i is high, for example greater than 0.9, then the similarity between $M_i$and its customized training molecules will be high. Then, we can use M_i and its customized training molecules as reference molecules to perform molecular recommendation introduced above. For example, we generate 10 recommended molecules, which will have high similarity with these reference molecules. Through the above method, we can try to generate augmented molecules from each known molecule. Next, we merge all the augmented molecules as shown in Figure 7. Firstly, the duplicated augmented molecules are removed, and then the adaptive modeling and reliability assessment are carried out for them. We then remove augmented molecules with reliability of 1, since they are already present in the original known dataset rather than new molecules. Finally, we selected the augmented molecules with reliability index greater than 0.9 as the final augmented molecules, and added them to the original molecular dataset to complete the data augmentation.

Next, we show another case. This case uses the Vc dataset after data correction and missing data augmentation described above. In this way, the original dataset can be clean and correct. Then we use the proposed data

Table 10: The augmented molecules

No	SMILES	Predicted Vc(mL/mol)	R	Real Vc(mL/m
1	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	2147.06	0.97	
2	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	2032.88	0.97	
3	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	1918.69	0.97	
4	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	1804.50	0.97	
5	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	1690.32	0.97	
6	CCCCCCCCCCCCCCCCCCCCCCCCCCCC	1576.13	0.96	
7	CCCCCCCCCCCCCCCCCCCCCCC(=O)O	1284.53	0.95	
8	OCCCCCCCCCCCCCCCCCCCCCC	1171.85	0.95	
9	CCCCCCCCCCCCCCCCCCCCC(=O)O	1179.25	0.95	
10	OCCCCCCCCCCCCCCCCCCCC	1067.81	0.95	1070
11	C=CCCCCCCCCCCCCCCCCCCC	1221.02	0.95	
12	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	2204.16	0.95	
13	C=CCCCCCCCCCCCCCCCCC	1104.87	0.94	1100
14	OCCCCCCCCCCCCCCCCCC	964.37	0.94	960
15	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	1975.78	0.94	
16	C=CCCCCCCCCCCCCCCCC	988.73	0.94	992
17	OCCCCCCCCCCCCCCCC	859.73	0.93	854
18	OCCCCCCCCCCCCCC	755.08	0.93	749
19	CCCCCCCCCCCCC(=O)O	760.97	0.92	758
20	CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC	2261.25	0.92	
21	CCCCCCCCCCCCCCCCCCCCCCCCC(=O)O	1337.17	0.91	
22	OCCCCCCCCCCCCCCCCCCCCCCC	1221.73	0.91	

augmentation method to expand this dataset. Finally, we can obtain the augmented molecules shown in Table 10. From this result, we can see that a total of 22 augmented molecules was generated, and their prediction reliability is very high. Then, we tried to verify their predicted property and finally obtained some of the real property values, as shown in the table. We can see that the prediction errors of the verified augmented molecules are all less than 1%, and the average error is 0.47%, which indicates that the proposed data augmentation method can indeed generate augmented molecules with reliable property information, so as to form reliable and accurate information enhancement to the original dataset.

4. CONCLUSION

This paper presents several applications of molecular similarity in data analysis for CAMD, including adaptive modeling, reliability assessment, and data preprocessing. Of course, none of these applications are isolated; they can be used in conjunction with each other. For example, the molecules generated by data augmentation cannot guarantee the correctness completely. In this case, the data consistency test can be used to further remove the incorrect data. Besides, adaptive modeling and reliability assessment have high requirements for the correctness of data. Therefore, the data consistency test method can be used to screen data mistakes for them. Finally, the molecular recommendation method can generate potential excellent molecules, and the adaptive modeling and reliability assessment can generate accurate property prediction and reliability index for them. Using these two indices can screen out molecules with both excellent performance and high reliability, thereby improving the reliability and accuracy of CAMD. In addition, the specific algorithm framework of each application introduced in this paper is not fixed. They can be adjusted and expanded according to the specific molecular fingerprint, similarity formulas and application scenario, which are also worthy of improvement in the future. In summary, this paper presents multiple applications of molecular similarity to enhance the efficiency of data utilization, laying the foundation for more efficient, accurate and reliable chemical data analysis and molecular design.

ACKNOWLEDGEMENTS

Financial support from the Natural Science Foundation of China (No. 22150410338) is gratefully acknowledged.

REFERENCES

1. Hukkerikar AS, Sin G;, Abildskov J;, Sarup B, Gani R. General rights Development of pure component property models for chemical product-process design and analysis). Development of pure component property models for chemical product-process design and analysis. APA; 2017.
2. Alshehri AS, Tula AK, You F, Gani R. Next generation pure component property estimation models: with and without machine learning techniques. AIChE Journal 68: (2021). https://doi.org/10.1002/aic.17469
3. Zhang J, Wang Q, Lei Y, Shen W. An interpretable 3D multi-hierarchical representation-based deep neural network for environmental, health and safety properties prediction of organic solvents. Green Chem. 26:4181-4191 (2024). https://doi.org/10.1039/d3gc04801b
4. Zhang L, Mao H, Liu L, Du J, Gani R. A machine learning based computer-aided molecular design/screening methodology for fragrance molecules. Computers & Chemical Engineering 115:295-308 (2018). https://doi.org/10.1016/j.compchemeng.2018.04.018
5. Jonuzaj S, Akula PT, Kleniati P, Adjiman CS. The formulation of optimal mixtures with generalized disjunctive programming: a solvent design case study. AIChE Journal 62:1616-1633 (2016). https://doi.org/10.1002/aic.15122
6. Sahinidis NV, Tawarmalani M, Yu M. Design of alternative refrigerants via global optimization. AIChE Journal 49:1761-1775 (2004). https://doi.org/10.1002/aic.690490714
7. Al-Shar'i NA. The design of TOPK inhibitors using similarity search, molecular docking, and MD simulations. Journal of Biomolecular Structure and Dynamics 43:8456-8467 (2024). https://doi.org/10.1080/07391102.2024.2319107
8. Gurung AB, Ali MA, Lee J, Farah MA, Al-Anazi KM. An updated review of computer?aided drug design and its application to COVID?19. BioMed Research International 2021: (2021). https://doi.org/10.1155/2021/8853056
9. Fu T, Xiao C, Li X, Glass LM, Sun J. MIMOSA: multi-constraint molecule sampling for molecule optimization. AAAI 35:125-133 (2021). https://doi.org/10.1609/aaai.v35i1.16085
10. JOBACK KG, REID RC. ESTIMATION OF PURE-COMPONENT PROPERTIES FROM GROUP-CONTRIBUTIONS. Chemical Engineering Communications 57:233-243 (2007). https://doi.org/10.1080/00986448708960487
11. Xu Y, Shao Z, Tula AK. Accurate property predictions and reliability quantification in molecular design based on molecular similarity. Computers & Chemical Engineering 201:109241 (2025).

https://doi.org/10.1016/j.compchemeng.2025.109241

12. Gani R, Hytoft G, Jaksland C, Jensen AK. An integrated computer aided system for integrated design of chemical processes. Computers & Chemical Engineering 21:1135-1146 (1997). https://doi.org/10.1016/s0098-1354(96)00324-9

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Pareto Front Guided Sampling for Efficient Bioprocess Experimentation

Stricker Samuel[ab*], Lucas Francisco dos Santos[c], Claus Wirnsperger[d], Alessandro Butté[d], Antonio del Rio Chanona[ab], Mehmet Mercangöz[ab], Gonzalo Guillén Gosálbez[c]

[a] Imperial College London, Department of Chemical Engineering, London, United Kingdom
[b] Sargent Center for Process Systems Engineering, London, United Kingdom
[c] ETH Zurich, Department of Chemical- and Bioprocess Engineering, Zurich, Switzerland
[d] DataHow AG, Zurich, Switzerland
* Corresponding Author: sam.stricker25@imperial.ac.uk

ABSTRACT

This work presents Pareto Front Guided Sampling (PFGS), a model-guided Design of Experiments (DoE) strategy for bioprocess development that makes the exploration–exploitation trade-off explicit and integrates human expertise into experiment selection. Starting from an initial experimental design, PFGS fits a probabilistic surrogate and then proposes new experiments by solving a multi-objective design problem that simultaneously rewards (i) high predicted performance (posterior mean) and (ii) high information gain (posterior uncertainty). Rather than collapsing this trade-off into a single acquisition value, PFGS generates a Pareto set of candidate experiments, that reflect different balances between improvement-seeking and learning. To prevent wasted runs, an automated screening step is performed to remove candidates in (i) low predicted-mean regions unlikely to yield near-optimal performance and (ii) low-uncertainty regions already well explained by the surrogate, concentrating effort on promising yet under-explored areas and encouraging coverage of multiple near-optimal basins. For industrial application, PFGS was extended to batch selection by choosing a diverse subset of Pareto candidates to support parallel experimentation without substantial loss in sample efficiency, and a termination rule to stop experimentation once user-defined accuracy thresholds are met. PFGS is benchmarked against Latin Hypercube Sampling (LHS) and Bayesian Optimisation (BO) on two representative problems: (i) a bioprocess case study with shallow performance gradients, where naive exploitation is unreliable and non-adaptive designs are inefficient, and (ii) a multi-modal landscape, where methods that focus on a single basin can miss alternative optima. Across these settings, PFGS allocates experimental resources more effectively than non-adaptive baselines and, in the multi-modal case, uniquely identifies and characterises all stationary regions while maintaining strong overall performance. By combining Pareto-transparent decision support, automated filtering, batch practicality, and principled stopping, PFGS provides an interpretable and resource-efficient DoE methodology for biopharmaceutical process optimisation.

Keywords: Design of Experiments (DoE), Bayesian Optimization (BO), Bioprocesses, Optimization, Pareto Front

INTRODUCTION

The biopharmaceutical industry faces mounting pressure to optimize processes due to elevated development costs, stringent regulatory requirements, and the inherent complexity of bioprocesses. Digital tools, including hybrid modeling and advanced statistical methodologies, have become essential for reducing development timelines and ensuring consistent product quality. Design of experiments has emerged as a critical methodology for systematically identifying and optimizing the critical factors that influence bioprocess outcomes, thereby enabling researchers to efficiently navigate complex experimental domains [1].

Classical approaches, such as Full Factorial Design (FFD), remain widely utilized in the pharmaceutical

industry owing to their simplicity and effectiveness in systematically mapping investigation spaces [2]. However, these methodologies prove insufficient for the demands of modern biopharmaceutical development, which increasingly necessitate rapid iterative refinement and stringent cost management [1]. Classical DoE methods typically necessitate a large number of experiments, which introduces substantial costs and extended timelines. Furthermore, conventional DoE frameworks, including FFD, are fundamentally limited in their capacity to model highly nonlinear processes, as they are restricted to approximating response functions using first- or second-order polynomial equations. Although LHS provides improved coverage for complex response surfaces, it similarly lacks adaptive and iterative refinement capabilities.

In contrast, model-based DoE approaches based on BO are becoming increasingly relevant because they (i) accommodate model uncertainty as actionable information for decision-making, and (ii) enable iterative experimental design optimization through systematic incorporation of knowledge from preceding experiments [1]. Methods such as Batch Bayesian Optimization (BBO) facilitate simultaneous optimization of multiple parallel experiments. Despite their considerable potential, these approaches remain underutilized in industrial settings, primarily due to challenges associated with integrating complex algorithms into established workflows [3]. Moreover, these methods are primarily designed for optimization rather than comprehensive mapping of the optimal region, which is essential for developing robust models capable of generalizing across diverse operational conditions.

Problem Statement

This work addresses the development of a comprehensive Design of Experiments framework designed to efficiently guide experimentation in bioprocesses. The overarching objectives are:

- to optimize an objective function (e.g., maximize product concentration),
- to systematically explore the investigation space such that the underlying, often nonlinear, process behavior can be accurately modeled with minimal experimental burden, and
- to accommodate practical constraints inherent to industrial operations (e.g., parallel experimentation or discretization).
- to provide transparent, human-interactable results that facilitate domain expert engagement and enable systematic incorporation of prior knowledge into the experimental design process.

The primary goal is to efficiently acquire process knowledge in regions of the investigation space near optima. This necessitates balancing between exploitation of known promising regions and exploration of uncertain or under-sampled areas, while simultaneously accounting for process nonlinearity and operational constraints of real-world systems. Furthermore, the practical requirements for adherence to experimental timelines and parallelization of resources motivate the development of algorithms capable of proposing multiple diverse, high-performing experimental candidates concurrently.

A fundamental challenge lies in the inherent trade-off between exploitation and exploration. Optimization-focused sampling strategies tend to concentrate proposed experiments in regions of high performance, potentially overlooking broader system dynamics and behavior. Additionally, designs proposed by such algorithms like BO are often not easily human-interpretable and often feel random and unfounded. Conversely, space-filling design approaches enhance model generalizability but may allocate resources inefficiently to regions of limited practical interest. This inefficiency is further compounded by the curse of dimensionality, whereby the investigation space expands exponentially with increasing numbers of input factors, rendering uniform sampling impractical. Effective DoE must therefore balance these competing objectives, emphasizing relevant design subspaces while maintaining sufficient global representativeness.

Industrial applications impose additional constraints that must be addressed: input variables may exhibit discretization due to measurement limitations, categorical factors introduce discontinuities, and operational feasibility considerations and Critical Quality Attributes (CQAs) substantially constrain the accessible design space. Consequently, any newly proposed DoE framework must integrate three fundamental aims: (i) identifying high-performing regions within a multi-dimensional investigation space, (ii) ensuring sufficient coverage to reliably capture process behaviour, and (iii) adhering to practical constraints, like batch operating modes, typical of industrial bioprocessing.

In summary, there is a critical need for integrated methodologies that efficiently balance exploration, exploitation, and practical operational considerations to accelerate biopharmaceutical process development under constraints of cost and regulatory complexity. Importantly, such advancement does not necessarily require the development of novel or more sophisticated acquisition functions for BO. Rather, existing algorithms should be adapted to enhance human-machine interaction and facilitate the systematic incorporation of domain expertise, thereby improving user accessibility and comprehensibility. This emphasis on human-centered design is essential, as human decision-making remains a predominant factor in contemporary process development

workflows.

METHODS

Classical Design of Experiment Techniques

Full Factorial Design is a statistical experimental methodology in which each factor is evaluated at n discrete levels, and all possible combinations of factor levels are systematically evaluated within the investigation space (an example is presented in Figure 1 to inform model development. This approach enables investigation of both the effects of each variable. However, the experimental complexity of FFD increases exponentially with the number of factors, rendering this method impractical in high-dimensional investigation spaces [4]. Despite this limitation, Fractional Factorial Design, a subset of FFD utilizing a reduced set of the experimental runs that only mildly mitigates the problem of the number of experiments, remains widely utilized in pharmaceutical research applications. Its continued adoption is attributable to its simplicity and experimental efficiency. Furthermore, the results obtained from factorial designs can be directly interpreted by researchers without requiring the development of a formal predictive model [2].

One of the most widely employed Monte Carlo-based sampling techniques for systematic space exploration is Latin Hypercube Sampling. Applied across diverse scientific and engineering disciplines, LHS is extensively utilized for uncertainty quantification and reliability analysis [5]. The foundational methodology, stratified sampling, partitions the investigation space into distinct strata, which are weighted according to their respective conditional probabilities. In LHS, each dimension is subdivided into n_s mutually non-overlapping subspaces. From these subspaces, random pairings are performed across dimensions such that each value from each dimension is selected exactly once without replacement, yielding n_s experimental designs of dimensionality h_d. A representative example of a LHS in two dimensions is illustrated in **Figure 1.** LHS and its variants are generally preferred for their computational efficiency and demonstrated capacity to capture nonlinear behavior across the investigation space with a minimal number of experiments.

Bayesian Optimization

Bayesian optimization is employed for the optimization of computationally expensive, black-box functions, such as in materials design applications [6]. BO utilizes a probabilistic surrogate model, typically a Gaussian process (GP) regressor, to approximate the objective function and to systematically guide the selection of promising candidate points [7]. The method balances exploration (sampling regions of maximum uncertainty) and exploitation (optimizing the objective function) through the implementation of acquisition functions [6], which are subsequently maximized using standard optimization algorithms. A widely employed acquisition function is the Upper Confidence Bound (UCB), which manages the exploration-exploitation trade-off by adopting an optimistic perspective regarding uncertainty. UCB is computed according to:

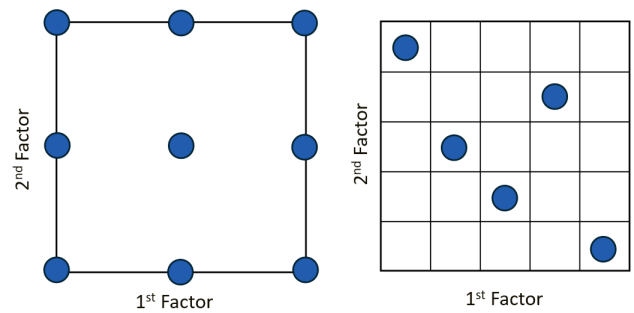

Figure 1 Example of a 3-level FFD (left) and a LHS (right) in two dimensions.

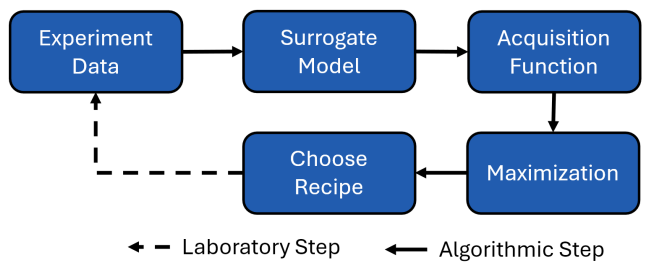

Figure 2 Overview of BO framework. BO uses the experimental data to train a surrogate model, from which an acquisition function is derived. This function is then maximized to find the most promising recipe.

$$\mathrm{UCB}(\boldsymbol{x}) = \mu(\mathbf{x}) + \beta * \sigma(\mathbf{x});\ \beta > 0 \qquad (1)$$

where μ and σ represent the mean and standard deviation of the surrogate model at point $\boldsymbol{x}$, respectively, while β is a hyperparameter that controls the exploration versus exploitation balance [8]. Alternatively, more sophisticated acquisition functions include Expected Improvement and Probability of Improvement. A representative workflow for BO is illustrated in **Figure 2**.

Batch Bayesian Optimization

The simultaneous sampling of multiple experimental designs using a BO approach is referred to as BBO [9]. BBO extends the principles of standard BO by enabling the concurrent evaluation of multiple candidate points in parallel, rather than sequentially. This parallelization offers substantial advantages in scenarios where multiple experiments or simulations can be conducted simultaneously, such as in distributed computing environments or laboratory settings equipped with multiple testing units. The primary challenge in BBO is that the acquisition function must be modified between successive point

selections to prevent redundant sampling of the same design point [9].

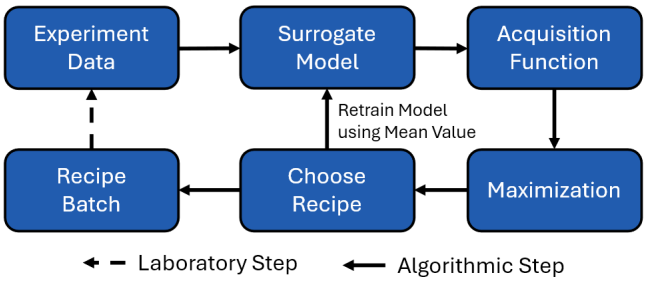

Figure 3 Overview of BBO. In q-point BO experimental data is used to generate a surrogate model from which an acquisition function is then maximised. The mean of the optimal point $\boldsymbol{x}$ is then added to the experimental batch and the surrogate model is retrained. This iterative procedure is performed until q points are sampled.

In classical BO the maximum of the acquisition function is identified and an experiment at the corresponding point $\mathbf{x}$ is conducted to obtain the true response y. This result is subsequently used to retrain the surrogate model. In q-point BO, a BBO method, no experiment is performed after the first BO iteration. The retraining procedure is performed by employing the mean prediction of the surrogate model $\mu(\boldsymbol{x})$ as a surrogate true value. After iterating this procedure n_q times, where n_q represents the batch size, all n_q experiments are executed in parallel, and the surrogate model is retrained with the empirically obtained values [9]. This approach is called BBO with Kriging Believer. The procedural flow of this algorithm is visualized in **Figure 3**.

Pareto Front Guided Sampling

In this work, a novel DoE methodology termed Pareto Front Guided Sampling was developed. The fundamental principle underlying PFGS is the generation of a Pareto front that explicitly addresses the exploration-exploitation trade-off. The methodology proceeds as follows: the surrogate model is initially trained on available experimental data. Subsequently, a Pareto front is generated to identify candidate points representing optimal compromises between competing objectives, like exploration (posterior mean) and exploitation (posterior uncertainty). An initial experimental design is selected from this Pareto front, and additional candidate designs are systematically identified based on the maximum Euclidean distance between candidate parameter vectors within the normalized, predefined investigation space. Once a sufficient number of candidate designs has been identified, the corresponding experiments are executed in a controlled laboratory setting, and the surrogate model is updated with the empirically obtained results. A comprehensive schematic representation of the PFGS algorithm is presented in **Figure 4**.

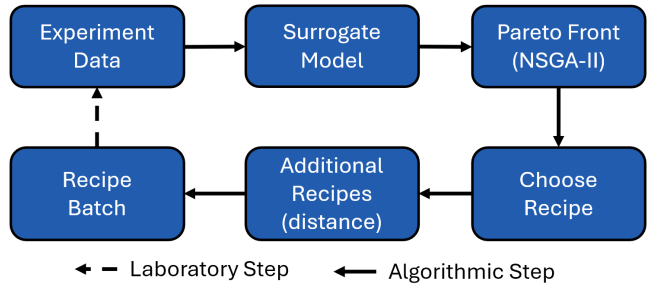

Figure 4 Overview of PFGS framework. The PFGS algorithm uses experiments to train the surrogate model from which a Pareto front is generated. The recipes are selected and experiments performed.

Pareto Front Generation

The Pareto front is generated by simultaneously optimizing the mean prediction of the surrogate model (representing exploitation) and the corresponding standard deviation (representing exploration) using a Multi Objective Optimization (MOO) Algorithm. The resulting Pareto front is subsequently refined through the elimination of candidate points from two distinct regions, as illustrated in **Figure 5**. First, candidate points characterized by low mean prediction are systematically removed, as these regions are of limited interest for identifying optimal process conditions. Second, candidate points located in regions of low standard deviation are removed. This elimination strategy serves two purposes. On the one hand it mitigates the influence of process measurement noise by focusing on dominant trends. On the other hand it ensures that the resulting sampling points are appropriately distributed throughout the investigation space, leveraging the property that standard deviation increases with distance from previously sampled experimental designs.

Recipe Selection

To determine the next experimental design in the batch, a candidate point must be selected from the remaining Pareto front (white region in **Figure 5**). In principle, all remaining points on the Pareto front are equally suitable for sampling. However, the specific selection depends on user preferences and objectives. The selection of the first recipe employs the utopia point as a reference. The utopia point is defined as the hypothetical point simultaneously maximizing both the mean prediction and standard deviation. Ideally, this utopia point (green) would be directly sampled; however, since it does not exist in the investigation space, the geometrically nearest point is selected using the Euclidean distance calculated in the normalized parameter space. For subsequent selections within the batch setting, the remaining candidate designs are sequentially identified from the Pareto front by maximizing the Euclidean distance (in the input space)

to previously selected experimental designs and executed experiments.

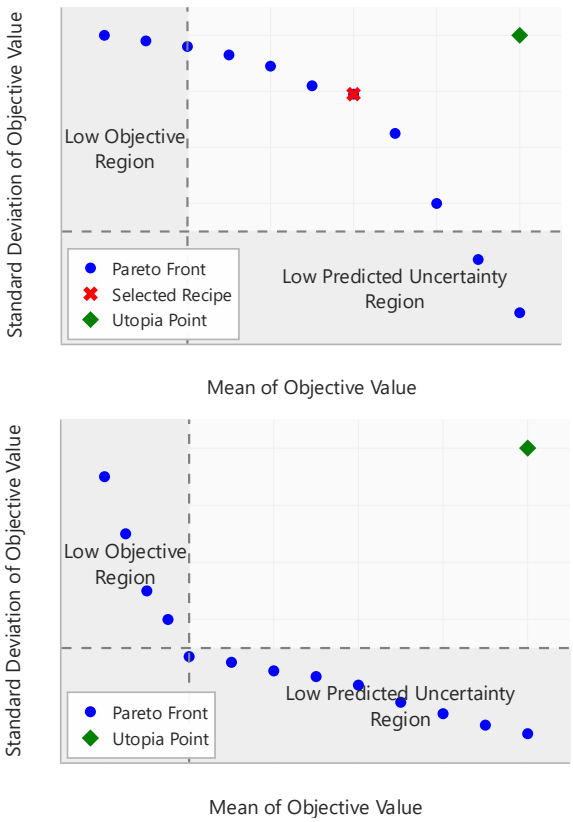

Figure 5 (top) The generated Pareto front using the model's mean and standard deviation. The gray areas are not sampled due to the low objective values or low uncertainty. (bottom) The Pareto front shown in blue is moving towards lower standard deviations as the model's knowledge is improving, resulting in the case that no point of the Pareto front is selected.

Stopping Criteria

The algorithmic structure enables termination of experimental design generation when the surrogate model achieves a specified level of accuracy within the target region of interest. **Figure 5** illustrates the Pareto front at an early sampling stage, where numerous candidate points occupy the white region from which the algorithm draws samples. However, as iterations progress and the model improves through successive refinement, the Pareto front gradually shifts toward regions of lower standard deviation. Eventually, no candidate points remain within the white sampling region.

This termination condition can occur at two distinct stages of the algorithm. First, termination may occur during batch selection, where insufficient candidate points remain in the sampling region to complete the full batch. For example, if five experimental designs were requested but only three viable candidates remain in the white region, the algorithm terminates early, resulting in an adapted (reduced) batch size. This indicates that continued sampling of the investigation space is not justified for the current model state. To obtain the remaining designs, either the elimination barrier that constrains sampling to high-performing regions must be relaxed, or a higher accuracy threshold must be set. Alternatively, the three available designs can be executed, and sampling can resume in the subsequent iteration. Second, termination may occur when no candidate points remain in the white sampling region at the initiation of an iteration. In this case, the algorithm ceases entirely, indicating that the surrogate model has achieved a sufficient level of training for the specified objectives.

Case Study

To demonstrate the efficacy of the proposed PFGS algorithm, two case studies of varying complexity were selected. To facilitate interpretability and explainability, the dimensionality was intentionally constrained to two input parameters (pH and temperature) and one output parameter (product concentration) across both case studies. The studies should represent simplified practical objectives for bioprocesses. In both cases no noise was added to the data.

Bioprocess

The first case study is based on a simplified Chinese Hamster Ovary (CHO) cell bioprocess model described in [10] that models a product concentration by varying pH and temperature. A visualization of the process response surface is presented in **Figure 6**. This case study presents the challenge of a relatively shallow gradient across most of the investigation space, which consequently renders optimization computationally challenging.

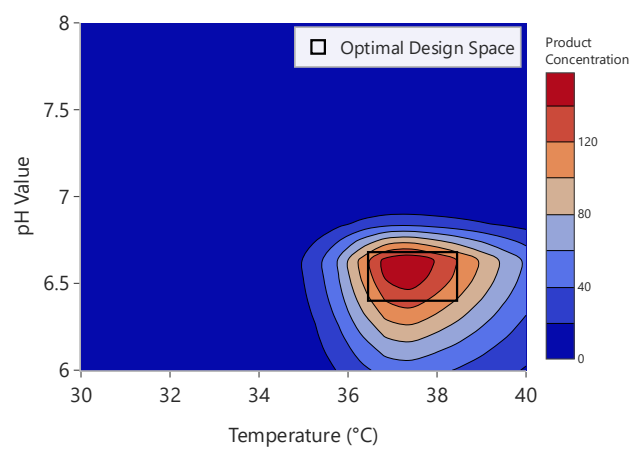

Figure 6 The optimization landscape of the bioprocess optimization problem.

Adapted Himmelblau Function

A second case study was developed using the well-established Himmelblau function [11]. The function was adapted to a maximization problem with an investigation space equivalent to that employed in the first case study. This case study presents the distinct challenge of systematically identifying and characterizing multiple local maxima. From a practical perspective, the global maximum (located in the upper left region) may be characterized by excessively constrained input variable ranges, rendering implementation infeasible in an industrial context. Conversely, the broader local maxima positioned in the right region of **Figure 7** offer more favorable practical operating conditions.

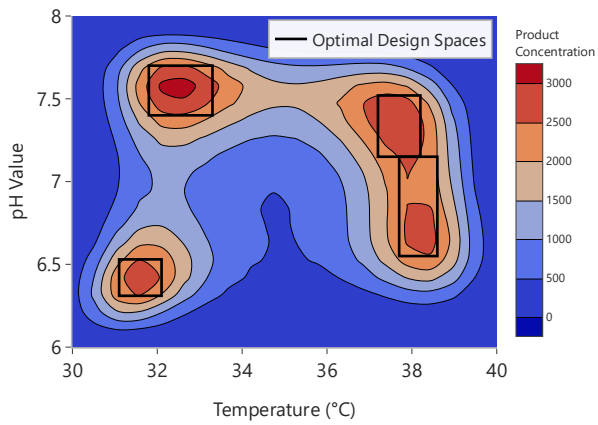

Figure 7 The adapted Himmelblau function to the optimization capabilities of the algoritms when multiple optima are present.

Experimental Setup

To evaluate the algorithm's performance, the optimal design regions are identified and delineated by black boxes. These regions are then used to generate test datasets by systematic sampling from the ground truth function. The resulting test datasets are presented in **Figure 8**.

Three distinct algorithms were compared in this study: LHS, BO, and the proposed PFGS. For both BO and PFGS, batch implementations were additionally evaluated. An initial sampling phase was conducted using LHS for all approaches, and a GP was trained on the resulting experimental data. Following initialization, the Relative Root Mean Squared Error (RRMSE) shown in **Equation 2** was calculated using the two previously defined test datasets.

$$RRMSE = \frac{\sqrt{mean((y_{predicted} - y_{test})^2)}}{std(y_{test})} \quad (2)$$

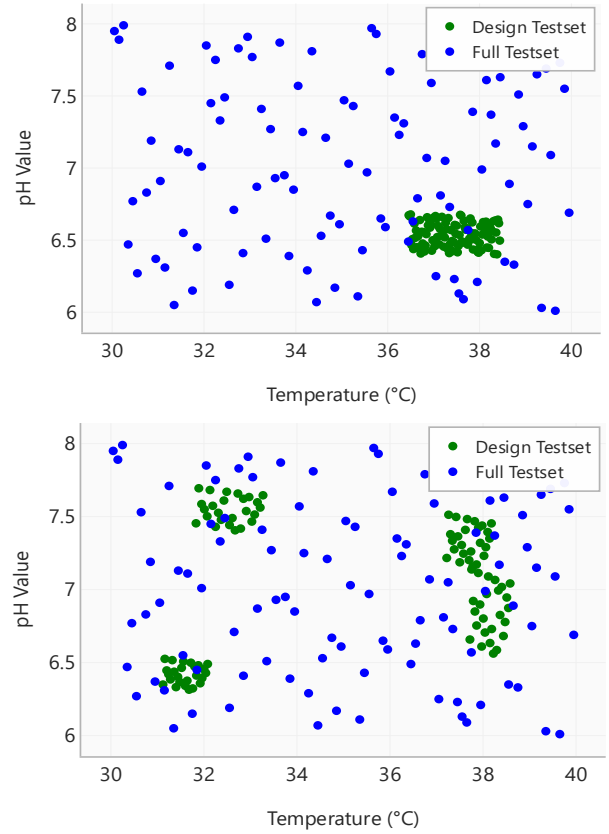

Figure 8 The test set for the classical Bioprocess (top) and the adapted Himmelblau function (bottom). The blue points correspond to the full dataset and the green ones to the design dataset.

$y_{predicted}$: test set values predicted with GP

y_{test}: true test set values

Subsequently, experimental designs were generated according to the respective algorithmic approaches, a GP was trained and the RRMSE was recalculated after each sample or batch of samples using the testset in order to track model improvement.

Implementation

In this study the data was sampled using LHS sampling implementation of the pyDOE2 package [12]. For all algorithms a GP was trained using botorch [13]. The MOO of the PFGS was performed using the NSGA-II algorithm [14] of the pymoo package [15], and BBO was used in the Kriging Believer setting with and UCB acquisition function with $\beta = 1.96$ and maximized using particle swarm optimization [16] with the pyswarm package [17]. The code is available online on the [Autonomous Industrial Systems Laboratory GitHub Page](#).

RESULT & DISCUSSION

Results were generated using the same five random

seeds initializations across the five independent runs for each algorithm. All plots display the mean RRMSE calculated from the five replicate runs and the 95% confidence interval. However, for the PFGS algorithm in the first case study, the stopping criterion was occasionally satisfied prior to completion of all five runs, resulting in mean RRMSE values calculated from fewer than five independent replicates. For automated boundary selection in PFGS, threshold values constraining the objective function and uncertainty metrics must be specified. In the first case study, a minimum objective value of 20 and a standard deviation threshold of 10 were established. For the adapted Himmelblau function, minimum objective and standard deviation thresholds of 1500 and 1000, respectively, were employed. These threshold values were derived from the initial Pareto front. However, in a human-in-the-loop setting, these parameters could be iteratively adjusted during each sampling iteration. For reproducibility and standardization purposes, such manual adjustments were not implemented in this study.

Bioprocess

The first bioprocess case study (**Figure 9**) clearly demonstrates that algorithmic guidance is essential for efficient identification of the optimal region within the investigation space. LHS exhibited the poorest performance, which is attributable to the relatively small region of practical interest compared to the total investigation space size. Consequently, LHS allocated experimental resources inefficiently to regions of negligible interest.

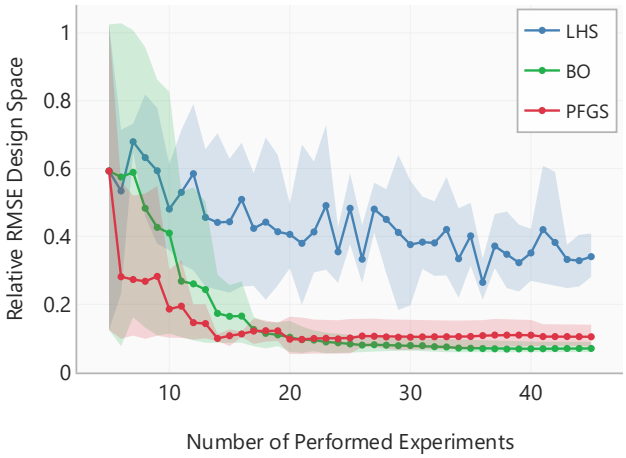

Figure 9 Bioprocess case study: Performance on the Design space test set of LHS, BO and PFGS with one experiment taken at each iteration.

Figure 10 presents the results for the batch optimization setting, where five parallel experiments were conducted simultaneously in each iteration. Both PFGS and BBO demonstrated comparable performance trajectories. When these batch results are compared to the sequential sampling approach (one experiment per iteration, shown in **Figure 9**), a modest decline in model performance is observed upon transition to the batch setting. This performance reduction is a well-established phenomenon in batch optimization, arising from the delayed model updating schedule; the surrogate model is updated only after completion of all five experiments within each batch, rather than after each experiment.

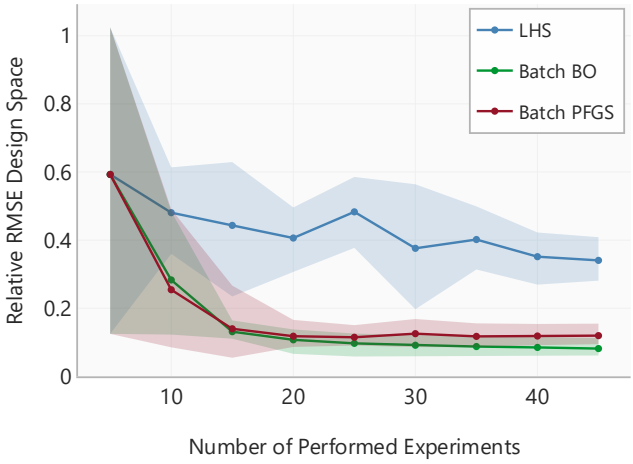

Figure 10 Performance on the Design space test set of LHS, Batch BO and Batch PFGS with five experiments taken at each iteration.

Adapted Himmelblau Function

This case study presents a more complex landscape characterized by three local optima and one global optimum. BBO encounters limitations in this setting, as it is fundamentally designed to identify a single maximum rather than to comprehensively map regions of interest. This limitation is evident in **Figure 11**, where the BO algorithm has a high RRMSE. As shown in **Figure 12** BO converges to a single optimum and focuses subsequent sampling efforts on local exploitation rather than exploring the different optima like PFGS. LHS demonstrated substantially improved performance in this case compared to the first example, attributable to the distributed spatial arrangement of the optima across the investigation space. However in needs to be mentioned that LHS is not suited in a practical setting due to its noniterative nature. Finally, PFGS allocates experimental resources more efficiently and does not wastefully concentrate sampling in subregions of negligible interest. It demonstrated comparable performance to LHS while possessing the distinct advantage of identifying and systematically sampling all four optima.

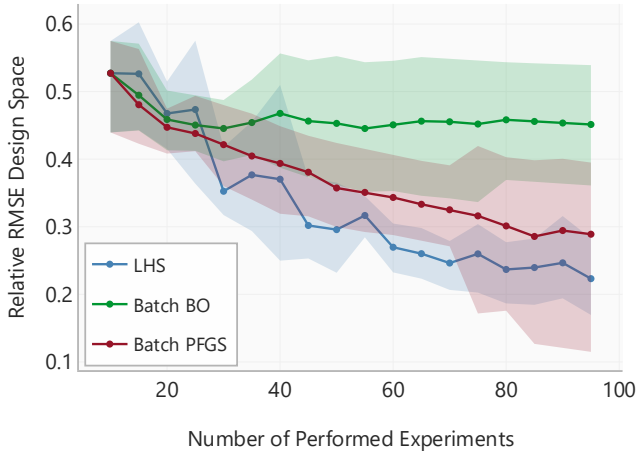

Figure 11 Himmelblau Case Study: Performance on the Design space test set of LHS, BBO and PFGS with one experiment taken at each iteration.

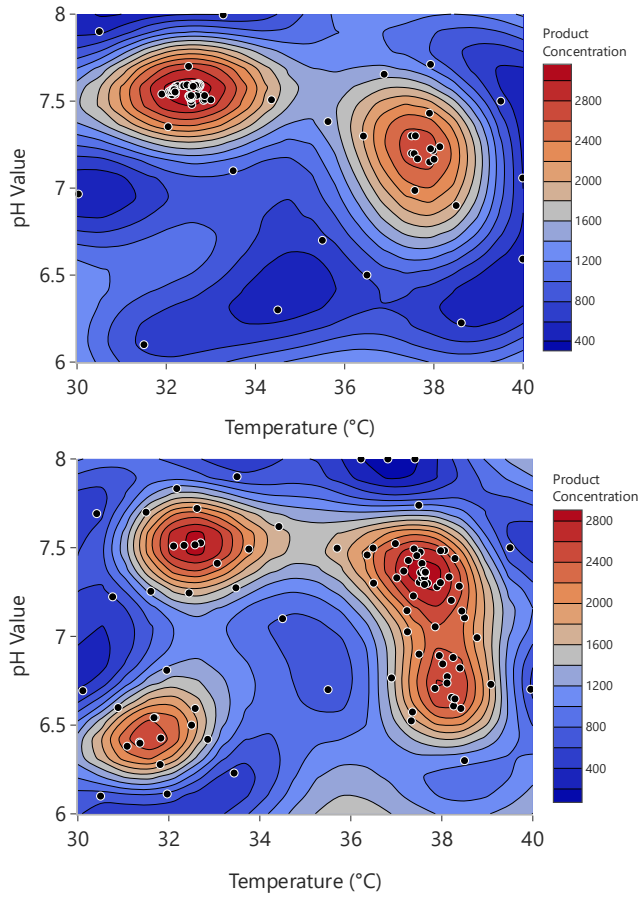

Figure 12 The location of sampled points (95 total) in the investigation space for one seed of Batch BO (top) and Batch PFGS (bottom). Additionaly a contour plot of the trained GP on those samples is visible.

The performance across the entire investigation space is illustrated in **Figure 13**. Both optimization-focused algorithms demonstrate comparatively poor performance in regions of low objective values. This outcome reflects the inherent trade-off previously described wherein improved accuracy in high-performing regions is achieved at the cost of reduced sampling coverage in regions of low objective values. However, this sampling bias is not problematic from a practical standpoint, since in industrial applications, experimental resources devoted to regions of limited practical interest represent an inefficient allocation.

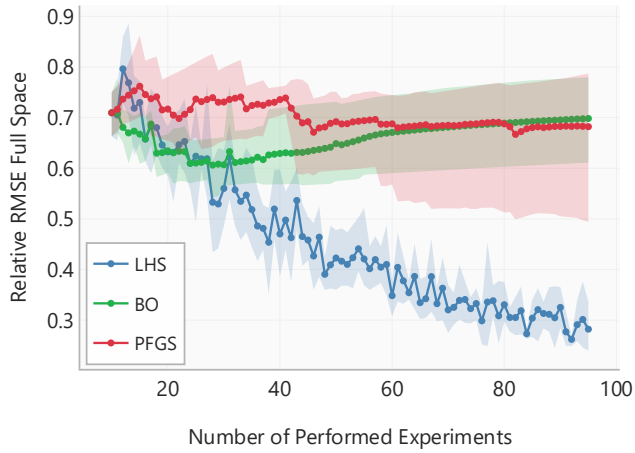

Figure 13 The performance of the algorithms on the full space (blue points in **Figure 8** (bottom) for the Himmelblau case study).

CONCLUSION

This work presents Pareto Front Guided Sampling, a novel DoE methodology that systematically addresses the fundamental trade-off between exploitation and exploration in nonlinear bioprocess optimization. The proposed algorithm generates a Pareto front representing simultaneous optimization of the surrogate model's predictive accuracy (mean prediction) and uncertainty quantification (standard deviation), enabling intelligent selection of experimental designs that balance the identification of high-performing with consideration of under-sampled regions.

In the first bioprocess case study, characterized by shallow gradients across the investigation space and one optima, PFGS showed similar performance to BO, efficiently allocating experimental resources to regions of practical interest. However, the ability of PFGS to better distribute experiments might still lead to increased performance relative to BO, particularly in high-dimensional settings or under substantial process noise; these aspects were not investigated in this study and remain directions for future work. In the second case study with multiple local and global optima, PFGS outperformed BO and was able to uniquely identify and systematically sampling all stationary points, a capability unavailable to classical BO approaches optimized for single-point convergence. Furthermore, the batch implementation of PFGS

maintains performance while accommodating the practical requirement for parallel experimentation, a critical consideration in industrial bioprocessing workflows. It is crucial to emphasize that PFGS does not outperform classical BO strategies in a strict sense. Rather, within bioprocessing contexts, broader coverage of the optima landscape is often desirable, as it facilitates identification of optimal design spaces. In this framework, the uncertainty threshold effectively biases the acquisition toward increased exploration.

The algorithmic framework incorporates automatic termination criteria, enabling the method to cease experimentation once the surrogate model achieves specified accuracy thresholds within regions of interest. This feature substantially reduces experimental burden and development costs while respecting practical constraints inherent to industrial bioprocesses, including parameter discretization, categorical factors, and critical quality attributes.

OUTLOOK

Future extensions of this work will incorporate constrained optimization and multi-objective formulations to address more complex bioprocess scenarios requiring simultaneous optimization of multiple competing performance criteria. Additionally, multi-fidelity optimization strategies leveraging experimental data from distinct reactor types and scales are envisioned, enabling more efficient knowledge transfer across different bioprocessing modalities. In this study, the analysis was constrained to two-dimensional investigation spaces to facilitate the interpretability and validation of the PFGS methodology. However, future investigations should systematically evaluate algorithm performance in high-dimensional investigation spaces representative of realistic bioprocess scenarios with numerous controllable parameters.

ACKNOWLEDGEMENTS

This research was conducted during a Master's thesis research at ETH Zurich in collaboration with DataHow AG and was subsequently further developed during doctoral studies at Imperial College London. This work was supported by financial contributions from DataHow AG and the Department of Chemical Engineering Scholarship from Imperial College London.

AUTHOR IDENTIFIERS

Author ORCIDs:
Stricker S: 0009-0008-2261-8430
Francisco dos Santos L: 0000-0001-5931-4272
Wirnsperger C: 0009-0000-8511-1428
Butté A: 0000-0003-3506-3792
Del Rio Chanona A: 0000-0003-0274-2852
Mercangöz M: 0000-0002-4449-0414
Gosàlbez G. G: 0000-0001-6074-8473

REFERENCES

1. Kasemiire A, Avohou HT, De Bleye C, Sacre PY, Dumont E, Hubert P, Ziemons E. Design of experiments and design space approaches in the pharmaceutical bioprocess optimization. European Journal of Pharmaceutics and Biopharmaceutics 166:144-154 (2021). https://doi.org/10.1016/j.ejpb.2021.06.004
2. Chowdary KPR, Ravi Shankar K. Optimization of pharmaceutical product formulation by factorial designs: case studies. Journal of Pharmaceutical Research 15:105 (2016). https://doi.org/10.18579/jpcrkc/2016/15/4/108815
3. Möller J, Kuchemüller KB, Pörtner R. Bioprocess intensification with model-assisted doe-strategies for the production of biopharmaceuticals. Physical Sciences Reviews 9:2925-2945 (2023). https://doi.org/10.1515/psr-2022-0105
4. Antony J. Full factorial designs. Design of Experiments for Engineers and Scientists :63-85 (2014). https://doi.org/10.1016/b978-0-08-099417-8.00006-7
5. Shields MD, Zhang J. The generalization of latin hypercube sampling. Reliability Engineering & System Safety 148:96-108 (2016). https://doi.org/10.1016/j.ress.2015.12.002
6. P. I. Frazier, “A Tutorial on Bayesian Optimization, ” Jul. 08, 2018, *arXiv*.
7. Frazier PI, Wang J. Bayesian optimization for materials design. Springer Series in Materials Science :45-75 (2015). https://doi.org/10.1007/978-3-319-23871-5_3
8. Archetti F, Candelieri A. The acquisition function. SpringerBriefs in Optimization :57-72 (2019). https://doi.org/10.1007/978-3-030-24494-1_4
9. N. Hunt, “Batch Bayesian Optimization, ” 2020.
10. Siska M, Pajak E, Rosenthal K, del Rio Chanona A, von Lieres E, Helleckes LM. A guide to bayesian optimization in bioprocess engineering. Biotech & Bioengineering 123:805-830 (2026). https://doi.org/10.1002/bit.70129
11. D. M. Himmelblau, *Applied Nonlinear Programming*. McGraw-Hill, 1972.
12. R. Sjögren and D. Svensson, *pyDOE2 Design of Experiments for Python*. (2018).
13. M. Balandat *et al.*, “BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization, ” Dec. 08, 2020, *arXiv*.
14. Deb K, Pratap A, Agarwal S, Meyarivan T. A fast and elitist multiobjective genetic algorithm: NSGA-

II. IEEE Trans. Evol. Computat. 6:182-197 (2002). https://doi.org/10.1109/4235.996017
15. Blank J, Deb K. Pymoo: multi-objective optimization in python. IEEE Access 8:89497-89509 (2020). https://doi.org/10.1109/access.2020.2990567
16. Kennedy J, Eberhart R. Particle swarm optimization. Proceedings of ICNN'95 - International Conference on Neural Networks 4:1942-1948 (None). https://doi.org/10.1109/icnn.1995.488968
17. James V Miranda L. Pyswarms: a research toolkit for particle swarm optimization in python. JOSS 3:433 (2018). https://doi.org/10.21105/joss.00433

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Sensitivity-Based Comparison of Resource Competition Models for Optogenetic Gene Circuit Design

Pratham Kapavarapu, Satyajeet S. Bhonsale, Simen Akkermans, Jan F.M. Van Impe*

BioTeC+ Chemical & Biochemical Process Technology & Control, KU Leuven, Ghent, Belgium
* Corresponding Author: jan.vanimpe@kuleuven.be.

ABSTRACT

Managing cellular resources, especially transcription and translation machinery, constitutes a significant constraint in the effective synthesis of useful bio-compounds from synthetic gene circuits. Although light-based optogenetic control approaches provide precise temporal and spatial control that can balance growth and production. The increased complexity adds on to the competition for cellular resources such as ribosomes and RNA polymerases. The optimization of regulatory elements in bioengineering is a vital task, as the selection of promoters and ribosome binding sites (RBS) directly affects transcriptional initiation rates and translation efficiency, hence influencing resource allocation. The vast number of potential parameter combinations requires systematic approaches to narrow the design space and determine which bioparts most significantly influence system performance. Furthermore, it is crucial to ascertain the uncertainty regarding which bioparts exert the most significant influence on the performance of engineered gene expression. In this study, we conduct an exploratory analysis by combining two established global gene expression models with two optogenetic TCS regulation models within a 2×2 framework. The first resource model (McBride & Del Vecchio, 2021) illustrates resource competition via coupled transcription-translation dynamics, but the second model (Santos-Navarro, 2021) identifies translation as the primary resource-limiting factor. We illustrate the utility of Sobol Global Sensitivity Analysis (GSA) to evaluate which model parameters (bioparts) influence target gene expression in the CcaS/CcaR optogenetic system regulating GFP expression. The findings indicate that the formulation of resource competition, rather than regulatory complexity, fundamentally dictates the rankings of parameter sensitivity. The sensitivity ranking explicitly indicates which design parameters require tight control and which may be relaxed, resulting in a significant reduction in design space.

Keywords: Optogenetic Control, Resource Competition, Global Sensitivity Analysis

INTRODUCTION

Optogenetic systems have gained traction in bioprocess engineering as compelling alternatives to chemical inducers, largely because they avoid contamination issues and simplify downstream processing [2, 12, 13]. The CcaS/CcaR two-component system derived from Synechocystis sp. PCC 6803 represents a prime example—it offers reversible, light-responsive regulation while imposing less metabolic load than conventional IPTG or arabinose-based approaches [2, 5]. Yet designing optogenetic bioprocesses presents an inherent difficulty: competition for cellular resources between engineered circuits and endogenous pathways intensifies considerably under industrial settings characterized by dense cultures and prolonged production timescales [3, 4]. Synthetic gene circuits applied for bioproduction must contend with resource depletion dynamics that can substantially shift system performance across extended production windows.

Two distinct modelling approaches have been established to capture the resource competition dynamics in gene expression. The McBride & Del Vecchio method tackles competition at the gene level, incorporating constraints of both RNA polymerase and ribosome availability via coupled demand parameters [6]. The Santos-Navarro RRS (Ribosome-Recruitment-Strength) model focuses exclusively on ribosomal competition, positing that

translation is the primary limitation in gene expression [11]. Analyzing these differing assumptions on cellular resource allocation provides a chance to investigate how model architecture and mechanistic intricacies affect design recommendations.

Sobol-based sensitivity approaches offer a flexible and computationally efficient means to assess parameter influence on elaborate bioprocess systems [1, 10]. Sensitivity analysis identifies parameters that consistently influence system thereby supporting robust design decisions, especially in synthetic gene circuits where experimental validation for all tunable parts is laborious. In this work, we establish a 2×2 sensitivity framework comparing both resource competition models coupled with optogenetic TCS (Two-Component System) regulation at varying levels of mechanistic detail, as outlined in Table 1. Using the CcaS/CcaR system regulating GFP production in E. coli as our case study, we examine how structural uncertainty in model assumptions propagates to parameter importance rankings and derive practical guidelines for gene circuit optimization.

Table 1. Model framework combinations used in the 2×2 sensitivity analysis. Each configuration pairs a resource competition model (columns) with a TCS (Two Component System) regulatory description (rows) at varying mechanistic details.

		Resource Competition Models	
		McBride & DelVecchio	Santos-Navarro RRS
TCS Models	Simple TCS	Model 1: McBride + Simple TCS	Model 3: RRS + Simple TCS
	Detailed TCS	Model 2: McBride+ Detailed TCS	Model 4: RRS + Detailed TCS

METHODS

Resource Competition Models

McBride & Del Vecchio Model

The McBride model represents resource competition through gene-level demand weights that account for both RNA polymerase and ribosome [6]. For gene i, the resource demand weight is calculated as:

$$w_i = D_i \times \frac{k_{tx,i}}{k_{tx,ref}} \times \left(1 + \frac{k_{tl,i}}{k_{tl,ref}} \times \frac{P_{tot}}{R_{tot} \cdot \delta}\right)$$

where D_i is the gene copy number, $k_{tx,i}$ and $k_{tl,i}$ are transcription and translation rates, P_{tot} and R_{tot} are total RNA polymerase and ribosome counts, and δ is the mRNA degradation rate. The resource availability factor is then computed as:

$$\rho_{McBride} = \frac{1}{1 + \sum_j w_j \cdot F_j}$$

where F_j represents the regulatory state (constitutive: $F_j = 1$; regulated: $F_j =$ Hill function).

Santos-Navarro RRS Model

The Santos-Navarro model focuses on mRNA-level ribosome competition through effective RBS strength parameters [11]. The effective ribosome sequestration strength for gene i is:

$$K_{C0,i} = K_{C0,ref} \times \frac{\mathrm{TIR}i}{\mathrm{TIR}ref}$$

where $K_{C0,ref} = 0.04\ molecule^{-1}$ for ribosomal genes and $K_{C0,ref} = 0.1\ molecule^{-1}$ for non-ribosomal genes, with $\mathrm{TIR}_{ref} = 6030$ representing the RBS strengths of non-ribosomal genes providing calibration scaling [6, 9, 11]. The resource availability factor becomes:

$$\rho_{RRS} = \frac{1}{1+\sum_j K_{C0,j} \times m_j}$$

where m_j is the mRNA concentration of gene j .

Tw-Component System Models

Simplified TCS Model

The simplified model assumes rapid equilibration of the CcaS photocycle, with light-dependent active fraction y:

$$\frac{d[\mathrm{CcaR}]}{dt} = k_{tl,CcaR} \cdot m_{CcaR} \cdot \rho - k_{phos,eff} \cdot [\mathrm{CcaR}] + k_{dephos} \cdot [\mathrm{CcaR\text{-}P}] - \gamma \cdot [\mathrm{CcaR}]$$

$$\frac{d[\mathrm{CcaR\text{-}P}]}{dt} = k_{phos,eff} \cdot [\mathrm{CcaR}] - k_{dephos} \cdot [\mathrm{CcaR\text{-}P}] - \gamma \cdot [\mathrm{CcaR\text{-}P}]$$

where $k_{phos,eff} = k_{phos} \times y \times [\mathrm{CcaS}_{tot}]$ and the regulatory function is $F_{GFP} = \frac{[\mathrm{CcaR\text{-}P}]^n}{K_d^n + [\mathrm{CcaR\text{-}P}]^n}$.

Detailed TCS Model

The detailed model explicitly tracks CcaS photocycle dynamics:

$$\frac{d[\mathrm{CcaS}_g]}{dt} = JCcaS - k_1[\mathrm{CcaS}_g] + (k_2 + k_{dr})[\mathrm{CcaS}_a] + k_{phos}[\mathrm{CcaS}_a][\mathrm{CcaR}] - \gamma[\mathrm{CcaS}_g]$$

$$\frac{d[\mathrm{CcaS}_a]}{dt} = k_1[\mathrm{CcaS}_g] - (k_2 + k_{dr})[\mathrm{CcaS}_a] - k_{phos}[\mathrm{CcaS}_a][\mathrm{CcaR}] - \gamma[\mathrm{CcaS}_a]$$

where $JCcaS = k_{tl,CcaS} \cdot m_{CcaS} \cdot \rho$ represents CcaS protein synthesis rate while $CcaS_g$ and $CcaS_a$ represent ground and active state respectively, with photoconversion rates k_1 and k_2 determined by wavelength and intensity [7].

Gene Expression Dynamics

Table 2: Parameters used in the Simulations.

Parameter Type	Parameter Symbol	Description
Common Parameters	AU_CcaS	CcaS promoter strength
	AU_CcaR	CcaR promoter strength
	AU_GFP_max	GFP promoter strength (activated)
	AU_GFP_basal	GFP promoter strength (basal)
	TIR_CcaS	CcaS RBS strength
	TIR_CcaR	CcaR RBS strength
	TIR_GFP	GFP RBS strength
	γ	Protein degradation rate
	δ	mRNA degradation rate
	D_{CcaS}	CcaS gene copy number
	D_{CcaR}	CcaR gene copy number
	D_{GFP}	GFP gene copy number
TCS Parameters	K_d	Activation threshold (half-maximal CcaR-P)
	n_{hill}	Hill coefficient (cooperativity)
	k_{phos}	Phosphorylation rate
	k_{dephos}	Dephosphorylation rate
	k_{dr}	Dark reversion rate (detailed TCS only)
	k_1	Light activation rate (detailed TCS only)
	k_2	Light deactivation rate (detailed TCS only)

For all models, mRNA dynamics follow:

$$\frac{dm_i}{dt} = k_{tx,i,eff} \cdot D_i - \delta \cdot m_i$$

where $k_{tx,i,eff} = k_{tx,i,basal} + (k_{tx,i,max} - k_{tx,i,basal}) \times F_i$ for regulated genes and where $k_{tx,i,max}$ directly correlates to promoter strength (AU_gene). Target protein dynamics are governed by:

$$\frac{d[\mathrm{GFP}]}{dt} = k_{tl,GFP} \cdot m_{GFP} \cdot \rho - \gamma \cdot [\mathrm{GFP}]$$

where $\mathrm{X_i}$ is protein concentration of gene i and $k_{tl,i}$ is translation rate for gene i and directly corresponds to RBS strength of the gene (TIR_gene).

Sobol Global Sensitivity Analysis

Global sensitivity analysis provides a systematic framework for quantifying parameter importance in complex systems with multiple uncertain inputs [10]. Let Y be the system output (steady-state GFP concentration) that depends on a set of uncertain parameters $\boldsymbol{X} = \{X_1, X_2, \dots, X_k\}$. The Sobol method decomposes the total output variance into contributions from individual parameters and their interactions: $\mathrm{Var}(Y) = \sum_i V_i + \sum_{i<j} V_{ij} + \sum_{i<j<l} V_{ijl} + \cdots + V_{1,2,\dots,k}$ where V_i represents the variance contribution from parameter $X_i$alone, $V_{ij}$from the interaction between parameters X_i and X_j, and so forth. First-order sensitivity indices quantify the main effect of each parameter:

$$S_i = \frac{V_i}{\mathrm{Var}(Y)} = \frac{\mathrm{Var}[E(Y|X_i)]}{\mathrm{Var}(Y)}$$

Total-order sensitivity indices capture all effects involving parameter X_i, including interactions:

$$S_{T_i} = 1 - \frac{\mathrm{Var}[E(Y|X_{\sim i})]}{\mathrm{Var}(Y)}$$

where $X_{\sim i}$ denotes all parameters except X_i. Sensitivity indices are computed using Saltelli's sampling scheme with N=256 base samples, resulting in $N \times (2k + 2)$ toal model evaluations where k is the number of parameters. All simulations were performed to $t = 500$ minutes to ensure steady state was reached. The analysis was implemented using the SALib Python library. The main parameters used in the simulations are listed in Table 2.

RESULTS AND DISCUSSION

McBride Resource Competition Models

In this analysis, we consider the McBride resource competition model integrated with both simplified and detailed TCS descriptions. The TCS parameter values were drawn from published characterizations of the CcaS/CcaR system [7], resource competition parameters from [6], [11], and promoter and RBS strengths were computed using the Salis Lab RBS Calculator [9].

Steady-state GFP expression was assessed at constant green light and served as the output metric for sensitivity analysis. Figure 1 illustrates the variations in the sensitivity profiles between the simplified and detailed TCS implementations. In Model 1 (McBride + Simple TCS),

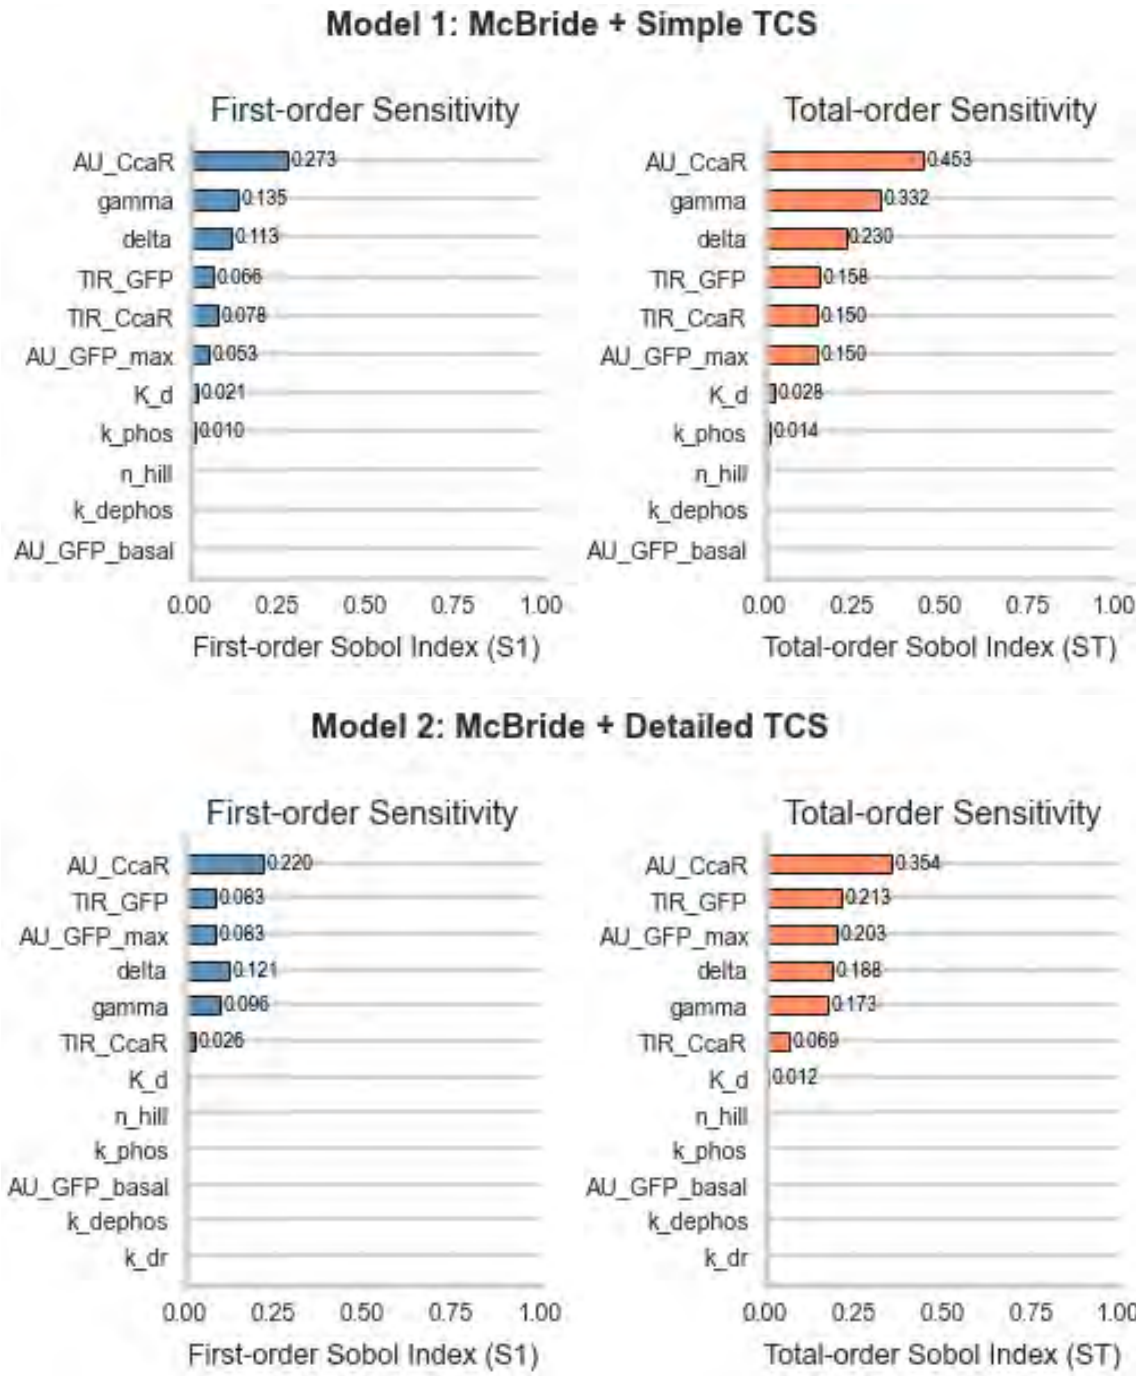

Figure 1. Global sensitivity analysis of McBride models showing first-order indices (S_1, blue bars) and total-order indices (ST, orange bars) for Model 1: McBride + Simple TCS (on top) and Model 2: McBride + Detailed TCS (bottom). CcaR promoter strength (AU_CcaR) dominates sensitivity in both models, with detailed TCS implementation showing increased sensitivity to target gene parameters (TIR_GFP, AU_GFP_max).
Santos-Navarro RRS Model

the CcaR promoter strength (AU_CcaR) exhibits the highest sensitivity index at ST = 0.45, followed by protein degradation (γ, ST = 0.33) and mRNA degradation (δ, ST = 0.23). Target gene parameters demonstrate a moderate effect, with TIR_GFP and TIR_CcaR presenting ST values in the range of 0.15–0.16. This pattern indicates that transcriptional regulation through CcaR availability primarily controls system output while degradation kinetics contribute secondarily.

In Model 2 (McBride + Detailed TCS), AU_CcaR

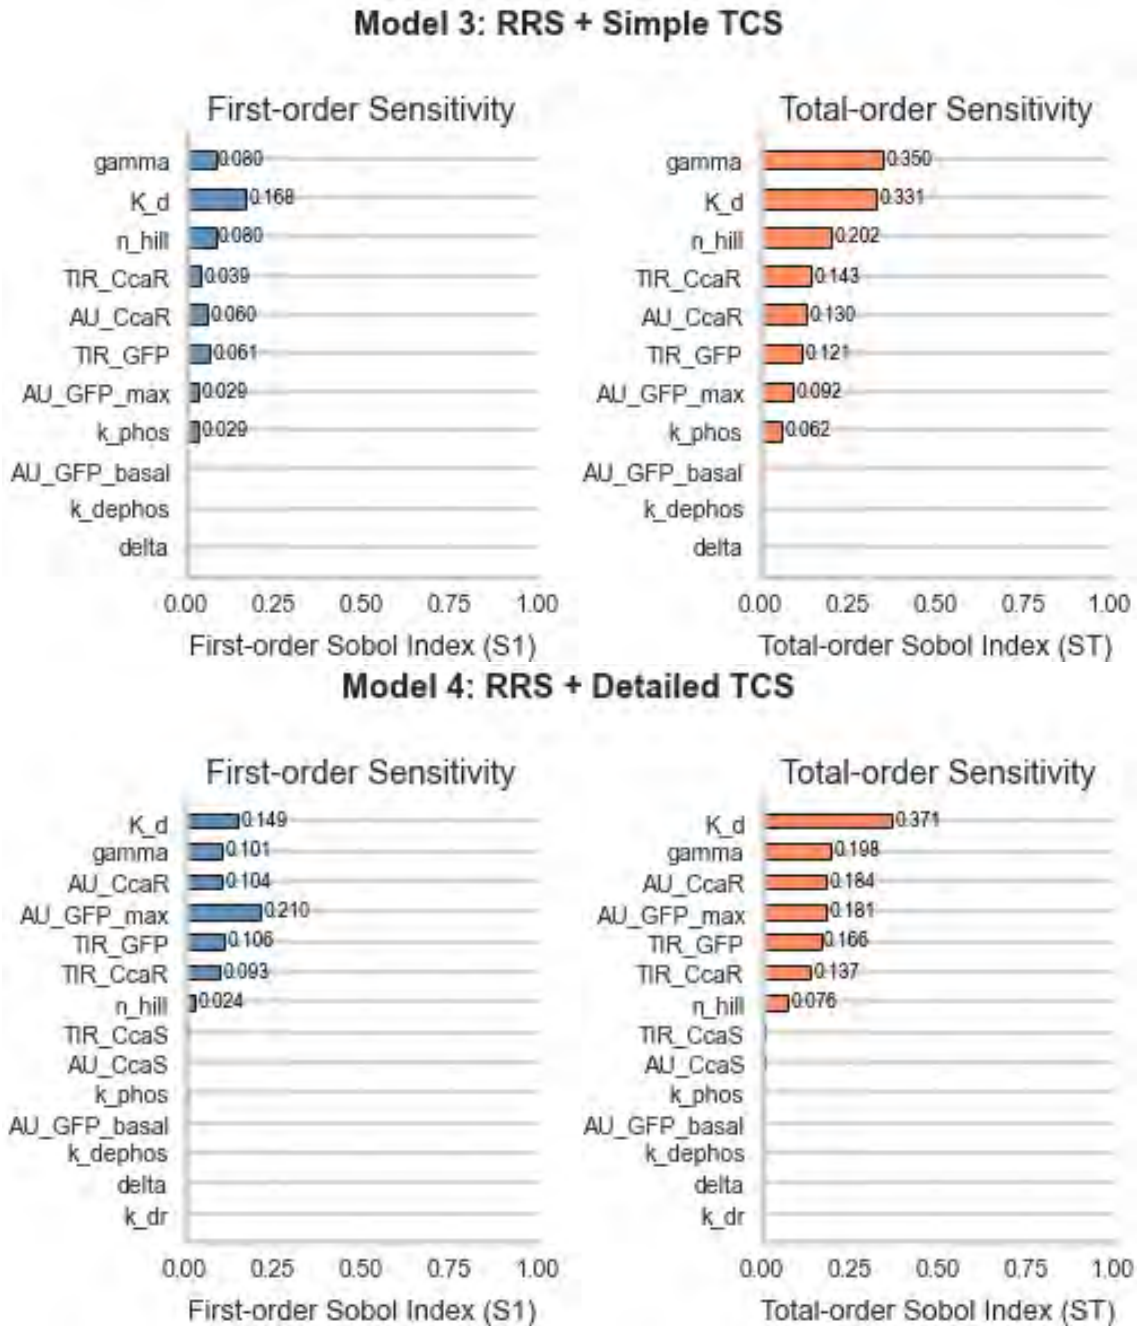

Figure 2. Sobol sensitivity analysis of Santos-Navarro RRS resource competition models. Global sensitivity analysis showing first-order indices (S_1, blue bars) and total-order indices (ST, orange bars) for Model 3: RRS + Simple TCS (on top) and Model 4: RRS + Detailed TCS (bottom). TCS activation threshold (K_d) emerges as the dominant parameter in both RRS formulations.

remains as the most sensitive parameter (ST = 0.35), albeit with a reduced extent. A notable shift in the sensitivity indices emerges as target gene parameters become more significant, with TIR_GFP as the second most influential parameter (ST = 0.21), followed by AU_GFP_max (ST = 0.19). Degradation parameters (δ, γ) exhibit a decline in relative importance. This shift suggests that explicit photocycle dynamics redirects sensitivity toward downstream expression elements while maintaining transcriptional dominance.

The Santos-Navarro resource competition models reveal a distinct sensitivity landscape compared to McBride models, despite same parameter set, as shown in Figure 2. For Model 3 (RRS + Simple TCS), protein

degradation (γ) emerges as the dominant parameter with ST = 0.39, while the TCS activation threshold (K_d) exhibits substantial influence (ST = 0.33). The Hill coefficient (n_{hill}) gains notable importance (ST = 0.20), reflecting heightened sensitivity to cooperative binding dynamics. TCS regulatory parameters (TIR_CcaR, AU_CcaR) show moderate sensitivity (ST = 0.13–0.14), while target gene parameters (TIR_GFP, AU_GFP_max) display reduced influence compared to McBride models.

In Model 4 (RRS + Detailed TCS), K_d becomes the most influential parameter (ST = 0.37), representing a shift toward signal transduction-dominated sensitivity. Protein degradation (γ) maintains high sensitivity (ST = 0.20), followed by balanced contributions from AU_CcaR (ST = 0.18), AU_GFP_max (ST = 0.18), and TIR_GFP (ST = 0.17). TIR_CcaR also retains moderate importance (ST = 0.14). Notably, mRNA degradation (δ) becomes negligible in both RRS models (ST < 0.05), contrasting sharply with its importance in McBride models. This reflects the RRS framework's emphasis on translational resource competition, where mRNA steady-state levels exert reduced influence on system dynamics.

CONCLUSION

We put forward a sensitivity-driven approach for studying light-controlled gene expression systems, through a comparative sensitivity-based evaluation of resource competition models. The approach described above systematically identifies essential design parameters and quantifies the impact of model structural assumptions for optimising synthetic gene circuits. Notably, certain parameters demonstrate model-invariant behaviour: CcaS expression (AU_CcaS, TIR_CcaS), basal GFP levels (AU_GFP_basal), and photocycle kinetics (k_{dephos}) consistently exhibit negligible sensitivity (ST < 0.05), indicating these are not rate-limiting under either framework.

From a practical standpoint, this approach enables substantial design space reduction as parameters with ST < 0.1, comprising approximately 40–50% of each model, can be fixed at nominal values, thereby directing experimental effort toward high-impact parameter targets while minimizing combinatorial burden.

Though experimental validation of model predictions to confirm both simulated system behavior and sensitivity rankings remains to be seen. The framework's modularity readily allows for adaptation to other bacterial species and various genetic circuit architectures. Integration with automated strain engineering platforms [8] represents a particularly promising direction for accelerating sensitivity-informed optimization of bioproduction systems.

ACKNOWLEDGEMENTS

This work was supported by the ALFAFUELS project, funded by the European Union (Grant Agreement No. 101122224), and by the Research Foundation Flanders FWO (μ4C Project, Grant No. G0A9H26N).

AUTHOR IDENTIFIERS

Author ORCIDs:
Pratham Kapavarapu: 0009-0004-4002-4178
Satyajeet S. Bhonsale: 0000-0001-9734-4122
Simen Akkermans: 0000-0003-2904-0439
Jan F.M. Van Impe: 0000-0002-5904-1638

REFERENCES

1. Bhonsale, S., Munoz Lopez, C. A. and Van Impe, J. (2019) "Global Sensitivity Analysis of a Spray Drying Process, " Processes, 7(11), p. 562. at: https://doi.org/10.3390/pr7110562
2. Camsund D, Lindblad P, Jaramillo A. Genetically engineered light sensors for control of bacterial gene expression. Biotechnology Journal 6:826-836 (2011). https://doi.org/10.1002/biot.201100091
3. Ceroni F, Algar R, Stan GB, Ellis T. Quantifying cellular capacity identifies gene expression designs with reduced burden. Nat Methods 12:415-418 (2015). https://doi.org/10.1038/nmeth.3339
4. Hartline CJ, Schmitz AC, Han Y, Zhang F. Dynamic control in metabolic engineering: theories, tools, and applications. Metabolic Engineering 63:126-140 (2021). https://doi.org/10.1016/j.ymben.2020.08.015
5. Hirose Y, Shimada T, Narikawa R, Katayama M, Ikeuchi M. Cyanobacteriochrome ccas is the green light receptor that induces the expression of phycobilisome linker protein. Proc. Natl. Acad. Sci. U.S.A. 105:9528-9533 (2008). https://doi.org/10.1073/pnas.0801826105
6. McBride CD, Del Vecchio D. Predicting composition of genetic circuits with resource competition: demand and sensitivity. ACS Synth. Biol. 10:3330-3342 (2021). https://doi.org/10.1021/acssynbio.1c00281
7. Olson EJ, Tzouanas CN, Tabor JJ. A photoconversion model for full spectral programming and multiplexing of optogenetic systems. Mol Syst Biol 13: (2017). https://doi.org/10.15252/msb.20167456
8. Pushkareva A, Beltrán J, Díaz-Iza H, Arboleda-García A, Boada Y, Vignoni A, Picó J. Towards automation of the design-build-test-learn (DBTL) bioengineering cycle: application to the testing and characterization of standard bioparts. XLIV Jornadas de Automática: libro de actas: Universidad de Zaragoza, Escuela de Ingeniería y

Arquitectura, 6, 7 y 8 de septiembre de 2023, Zaragoza :453-458 (2023). https://doi.org/10.17979/spudc.9788497498609.453

9. Salis HM, Mirsky EA, Voigt CA. Automated design of synthetic ribosome binding sites to control protein expression. Nat Biotechnol 27:946-950 (2009). https://doi.org/10.1038/nbt.1568
10. Saltelli A, Ratto M, Andres T, Campolongo F, Cariboni J, Gatelli D, Saisana M, Tarantola S. Global sensitivity analysis. the primer. Wiley (2008). https://doi.org/10.1002/9780470725184
11. Santos-Navarro FN, Vignoni A, Boada Y, Picó J. RBS and promoter strengths determine the cell-growth-dependent protein mass fractions and their optimal synthesis rates. ACS Synth. Biol. 10:3290-3303 (2021). https://doi.org/10.1021/acssynbio.1c00131
12. Tabor JJ, Levskaya A, Voigt CA. Multichromatic control of gene expression in escherichia coli. Journal of Molecular Biology 405:315-324 (2011). https://doi.org/10.1016/j.jmb.2010.10.038
13. Zhao EM, Suek N, Wilson MZ, Dine E, Pannucci NL, Gitai Z, Avalos JL, Toettcher JE. Light-based control of metabolic flux through assembly of synthetic organelles. Nat Chem Biol 15:589-597 (2019). https://doi.org/10.1038/s41589-019-0284-8

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Genome to Production: A Multiscale Model for Bioprocess Design

Rajiv Kailasanathan[ab], Mohammad Reza Boskabadi[a], Abhishek Sivaram[a], Seyed Soheil Mansouri[a]*

[a] Department of Chemical and Biochemical Engineering, Technical University of Denmark, Kongens Lyngby, Denmark
[b] Novo Nordisk Foundation Center for Biosustainability, Technical University of Denmark, Kongens Lyngby, Denmark
* Corresponding Author: seso@kt.dtu.dk.

ABSTRACT

Bioprocesses are inherently multiscale, spanning intracellular metabolism to production-scale reactors. Simulation models that integrate these scales offer potential strategies to study the effect of changing metabolic states and enable efficient integration of biological knowledge gathered from lab-scale experiments. In this study, we demonstrate the potential of such simulation model towards the production of mevalonate, an important pharmaceutical drug compound produced through fermentation of a fungal species *Aspergillus terreus.* We integrate a genome-scale metabolic model of the organism with a plant-wide simulation model for the bioprocess that encompasses several upstream and downstream unit operations. Through this integration, we identify potential targets for metabolic engineering towards increased product flux and simultaneously estimate the associated oxygen requirements. This framework serves as a foundation for developing digital twins of bioprocesses that bridges strain engineering with process design and operations.

Keywords: Simulation, Biosystems, Fermentation, Optimization, Multiscale Modelling, Metabolic models

INTRODUCTION

Bioprocesses are multiscale phenomena, where organic reactions catalysed by enzymes (10^{-9} m) occur within cells (10^{-6} m) that exist in complex chemical environments inside fermenters (10^{0}–10^{4} m^3). Digital models of bioprocesses typically approximate this complexity using parameterized kinetics (e.g.: Michaelis-Menten, Hill) without detailed descriptions of intracellular metabolism. The basis of modern industrial biotechnology is formed on the ability to engineer intracellular metabolism, resulting in development of strains with metabolic capabilities distinct from wild-type. Genome-scale metabolic networks (GEMs) provide this intracellular resolution, allowing us to study the characteristics of altered metabolic states. Strain design based on computational tools are at the forefront of modern systems biology research.

Modern bioprocess development follows the Design-Build-Test-Learn (DBTL) cycle at both strain and process levels. These cycles are interdependent: a strain optimised under idealised conditions may underperform in production-scale fermenters, while processes designed around wild-type kinetics may fail to exploit engineered hosts. Currently, these cycles occur independently. Metabolic engineers use genome-scale models under steady-state assumptions, while process engineers develop dynamic simulations from empirical data. Neither captures the feedback between intracellular metabolism and process operating constraints.

An integrated framework embedding GEMs within dynamic process models closes this gap. Candidate strains can be evaluated in silico under realistic conditions before costly experiments, while operating strategies can be tailored to host-specific metabolic capabilities. Such integration supports early techno-economic assessment by propagating flux predictions through to process metrics—titres, yields, productivities—and identifies which metabolic states satisfy process objectives under realistic operating conditions.

In this study, we demonstrate integration of a plant-wide simulation model with a genome-scale metabolic model of *Aspergillus terreus* for mevalonate production. Through this framework, we investigate intracellular flux distributions under process conditions, determine metabolic reorganisation required for high production, and predict oxygen requirements inaccessible from empirical

models alone. This framework demonstrates a step toward bioprocess digital twins, where mechanistic metabolic models inform process design and enable the DBTL cycle to operate across scales.

METHODS

The proposed approach is implemented on the KTB1 simulator, which is a MATLAB/Simulink benchmark for continuous biomanufacturing that couples an upstream CSTR bioreactor with realistic downstream operations (cell removal/clarification, filtration, buffer tanks, and chromatography)[1]. The KTB1 simulator considers a regulatory controller, considering noise and delay for sensors and actuators. This simulator provides physics-based unit models, instrumentation hooks, and disturbances to study end-to-end monitoring, control, and soft-sensor design in an integrated upstream–downstream setting. The bioreactor module is constructed from a kinetic model of *Aspergillus terreus* that characterizes the yield of mevalonate from lactose and adenine[2]. A genome-scale metabolic model of *Aspergillus terreus* (iJLi454)[3] is used to provide the intracellular resolution into the process. The model comprises 1451 reactions and 1160 metabolites, and the metabolic pathways required to produce the product precursor, mevalonate, are already present in the model.

Metabolic regimes for high product productivity

A key question is whether the current metabolic state fully exploits the cell's biosynthetic capacity, or whether alternate flux distributions could achieve higher yields under identical process constraints. To investigate the metabolic flux redistribution required to achieve high mevalonate production, flux sampling was performed on the GEM under two scenarios: (1) mevalonate production constrained to 13% of the theoretical maximum, representing the current CSTR operating point predicted by empirical kinetics, and (2) mevalonate production constrained to 90% of the theoretical maximum, representing a high-production target. Both scenarios were subject to identical constraints on biomass formation rate (μ = 0.00149 h^{-1}), lactose uptake (0.00906 mmol/gDW/h), and adenine uptake (0.00106 mmol/gDW/h), derived from the KTB-1 model at steady state. Constraints were applied with ±5% tolerance to ensure numerical feasibility.

Flux sampling was performed using the OptGP algorithm[4] implemented in COBRApy, which employs a hit-and-run approach to uniformly sample the feasible solution space. For each scenario, 5000 flux distributions were collected with a thinning parameter of 500 steps between samples to minimise autocorrelation. Sampling was parallelised across 8 processes.

Flux distributions for 47 reactions spanning glycolysis, pentose phosphate pathway (PPP), TCA cycle, glyoxylate shunt, and anaplerotic pathways were compared between scenarios. Statistical significance was assessed using the Mann-Whitney U test with Bonferroni correction for multiple comparisons. Reactions were considered significantly altered if the adjusted p-value was less than 0.05.

Oxygen requirements for continuous operation of the bioprocess

The oxygen requirement of the process is not considered by the KTB-1 system, and the kinetic model assumes that the reactor is supplied with enough oxygen to sustain its metabolism. In this study, we use the genome-scale reconstruction to estimate the oxygen requirement for the continuous operation of the process. The steady-state specific rates from KTB-1 simulator were used to constrain flux bounds in the genome-scale metabolic model (GEM). Three fluxes were constrained based on CSTR predictions: lactose uptake, adenine uptake (nitrogen source), and biomass formation rate. Specific rates from the CSTR model (g/gDW/h) were converted to GEM-compatible units (mmol/gDW/h) using molecular weights of 342.3 g/mol for lactose, 70 g N/mol for adenine (based on nitrogen content), and 148.16 g/mol for mevalonate. Constraints were applied with a ±1% tolerance to account for numerical precision.

Flux variability analysis (FVA) [5] was performed on oxygen uptake and mevalonate production reactions to determine the feasible flux ranges under CSTR constraints. Subsequently, a Pareto analysis [6] was conducted by systematically varying the oxygen uptake constraint across its feasible range and maximising mevalonate production at each level, generating the oxygen–mevalonate trade-off frontier. All simulations were performed using COBRA Toolbox in MATLAB, and COBRApy with the GLPK solver [7].

RESULTS

Metabolic flux redistributions at high product production

Specific reaction rates of lactose uptake, adenine uptake, mevalonate production, and the specific biomass production rate were extracted from KTB-1 simulator. After constraining the genome-scale metabolic model iJL1454 with the extracted rates, the model was optimized for maximum mevalonate production, and the resultant rate was compared with current specific productivity of the process. It was found that the current process operates at 13% of the maximum flux predicted by the model. This suggests that either the current operating conditions do not fully exploit the metabolic capacity for mevalonate synthesis, or that regulatory constraints not captured by the stoichiometric model limit production

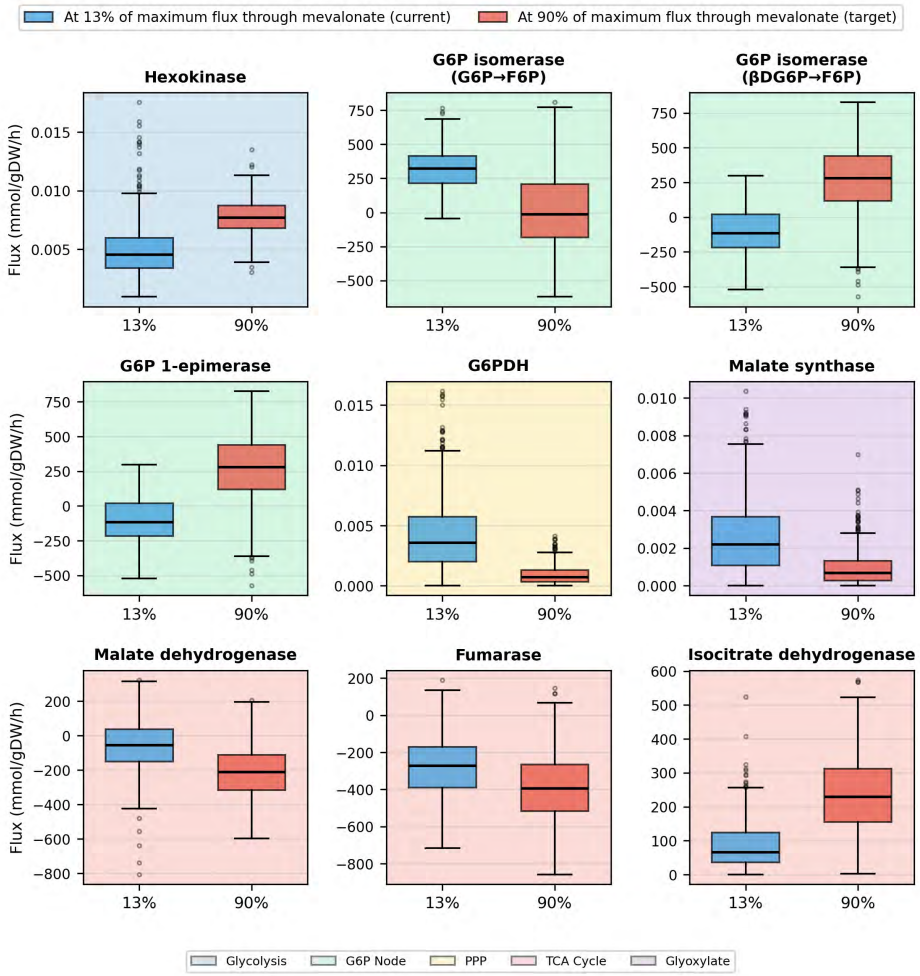

Figure 1. Box-plots of the distributions of reactions in the central carbon metabolism at the current efficiency (13%) and at 90% of the theoretical maximum of mevalonate production flux. The glucose-6-phosphate node shows metabolic reorganization, where the canonical enzyme is seen to be likely decreased in magnitude, while the other isoforms simultaneously show large changes in their distributions.

in vivo. This provides scope for metabolic engineering strategies that can increase the specific flux through the mevalonate production pathway.

To compare the systemic characteristics in the metabolism of a high-yielding strain, two models were constructed: In one, the flux through the mevalonate pathway was constrained to the current value (13% of theoretical maximum), and in another, the flux was constrained to 90% of the maximum.

Flux sampling confirmed that both the 13% and 90% mevalonate scenarios were feasible under the conditions specified by the KTB-1 simulator, with biomass flux tightly controlled (0.00142 ± 0.000003 h^{-1} in both cases). Comparison of flux distributions revealed significant changes in several reactions in the central carbon metabolism, indicating widespread metabolic reorganisation when shifting from low to high mevalonate production (Figure 1).

At the pathway level, the citric acid cycle showed the largest average flux change, followed by glycolysis. Anaplerotic reactions and glyoxylate shunt showed relatively minor variations, indicating that these pathways do not influence the high productivity strain. The flux through the upper glycolysis (hexokinase) is increased in the higher-yielding strain. However, at the glucose-6-phosphate (G6P) node, flux through the canonical isomerase (G6P to fructose-6-phosphate) decreased, while alternate isomerases (G6P 1-epimerase and β-G6P) increased their activity. At lower glycolysis (G6PDH) the flux is seen to have decreased.

In the mitochondria, the largest changes were observed in malate dehydrogenase, fumarase, and isocitrate dehydrogenase, indicating an altered TCA cycle that has reorganized to meet the energetic and co-factor demands of the product. These findings demonstrate the value of model-driven metabolic engineering.

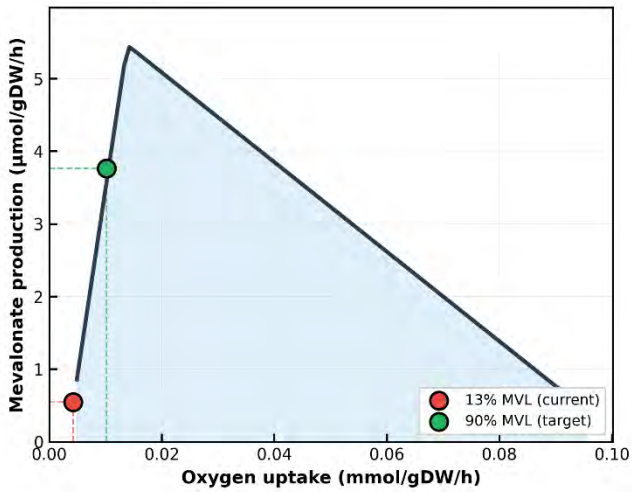

Figure 2. Pareto front showing the trade-off between oxygen uptake and mevalonate production under CSTR-derived constraints on lactose uptake, adenine uptake, and biomass formation. The solid line represents the maximum achievable mevalonate flux at each oxygen uptake level determined by FBA. The shaded region indicates the feasible operating space. Red circle marks the current operating point (13% of theoretical maximum mevalonate); green circle marks the target operating point (90% of maximum). Dashed lines indicate the minimum oxygen uptake required to sustain each production level.

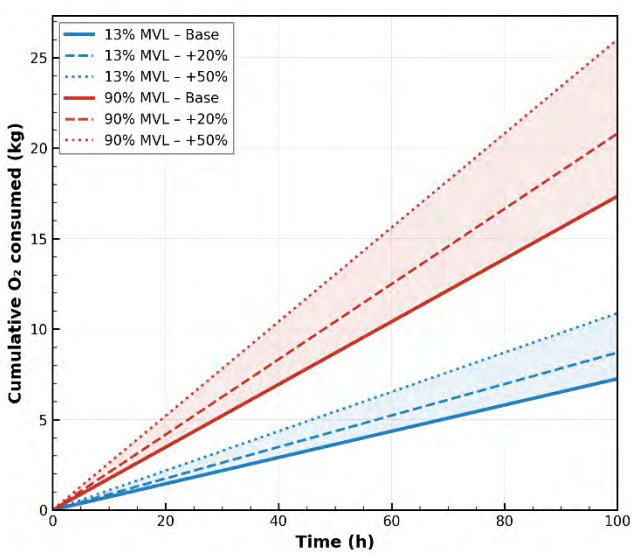

Figure 3. Sensitivity analysis on the cumulative oxygen requirements over 100 hours of steady state operations at the current productivity and at 90% of the theoretical maximum.

Oxygen demand – mevalonate production trade-offs

The GEM analysis confirmed that the steady-state operating point (growth rate = 0.00149 h^{-1}, lactose uptake = 0.00906 mmol/gDW/h, adenine = 0.00106 mmol/gDW/h) is metabolically feasible. Under these constraints, FVA revealed that the oxygen uptake rate can range from 0.0071 to 0.0964 mmol/gDW/h, while meeting the metabolic demands for product production.

The Pareto analysis identified a trade-off between oxygen efficiency and mevalonate productivity. Maximum mevalonate production (0.0054 mmol/g_DW/h) required an oxygen uptake of 0.0145 mmol/g_DW/h, approximately thrice the minimum oxygen requirement (Table 1). The minimum oxygen operating point yielded only 9% of the maximum achievable mevalonate flux, indicating that oxygen availability strongly influences product formation capacity. The Pareto front is visualized in Figure 2.

Table 1: Key operating points identified from the pareto front.

Operational mode	Oxygen uptake (mmol/gDW/hr)	Mevalonate flux (max)(µmol/gDW/hr)
Maximum mevalonate	0.0145	5.4
Minimum oxygen	0.0049	0.485

As seen from the Pareto analysis, any strategies undertaken to increase the product productivity under the current conditions would also be associated with increased oxygen demand. Through steady state simulations of the KTB-1 model, we visualize the cumulative oxygen demands over 100 hours of steady state operation while operating at 13% of the theoretical maximum product flux (current) as well as at 90% of the maximum. In both cases, the minimum oxygen consumption was computed by minimizing the flux through oxygen uptake while tightly constraining the fluxes through biomass, lactose uptake, and mevalonate production. In addition, a sensitivity analysis on the metabolic respiration efficiency was performed by assuming that the cells require 20% and 50% more oxygen than the minimum requirement. These results are visualized in Figure 3.

The results suggest that the minimum volumetric oxygen demand increase by 139% when the product flux increases from 13% of the theoretical maximum to 90% of the maximum. When the cells are metabolically inefficient, and they consume 50% more oxygen than the minimum, the difference in the cumulative oxygen demand over 100 hours of system operation correspond to an increased requirement of around 473.15 mols (15.1 kg) of oxygen.

DISCUSSION

The integrated framework provides quantitative predictions that are inaccessible from either model in isolation. The KTB-1 model tracks extracellular concentrations and predicts product titres, but lacks crucial information about oxygen demand, which is a critical parameter for reactor design and operating costs. By constraining the GEM with CSTR-derived specific rates, we determined that achieving 90% of the theoretical maximum mevalonate production requires 139% increase in the minimum oxygen demand compared to the current operation. Through the integrated framework, we can gather information essential for sizing aeration equipment and estimating utility costs during early-stage process development.

The flux sampling analysis further identified specific metabolic engineering targets. The substantial flux reorganization at the glucose-6-phosphate node and the shift in malate dehydrogenase activity between compartments suggest that cofactor balancing and carbon routing are key bottlenecks limiting mevalonate production. These predictions can guide the DBTL cycle by prioritizing genetic interventions. The current results suggest that modifying G6P isomerase expression or engineering redox shuttle capacity are potential metabolic engineering strategies that can improve the yield of the process.

From a techno-economic perspective, the framework enables the comparison of strain candidates not only on theoretical maximum yields, but on their associated oxygen costs and process requirements. A strain achieving high titers, but demanding excessive aeration, may prove uneconomical at scale. A multi-scale integration as demonstrated here could allow us to study such trade-offs *insilico,* before committing to expensive experimental campaigns.

CONCLUSIONS

This study demonstrates that integrating genome-scale metabolic models with dynamic process simulations provides mechanistic insight into bioprocess operation that neither approach offers alone. By bridging intracellular metabolism with reactor-scale mass balances, the framework enables prediction of oxygen requirements, identification of metabolic bottlenecks, and quantitative comparison of strain candidates under process-relevant conditions. These capabilities support rational decision-making in both strain and process engineering, allowing the DBTL cycle to operate across scales. As bioprocess development increasingly relies on engineered strains with non-native pathways, such integrations will become necessary as empirical correlations from wild-type strains may be insufficient. This work represents a step toward comprehensive digital twins for industrial biotechnology, where mechanistic models inform experiment design, accelerate scale-up, and reduce the cost and the risk of bringing new bioproducts to market.

ACKNOWLEDGEMENTS

This work was funded by the Novo Nordisk Foundation through grant number NNF20CC0035580.

AUTHOR IDENTIFIERS

R K: https://orcid.org/0009-0009-8915-7111
M R B: https://orcid.org/0009-0009-9619-3735
A S: https://orcid.org/0000-0003-0345-9776
S S M: https://orcid.org/0000-0002-0924-4947

REFERENCES

1. Boskabadi MR, Ramin P, Kager J, Sin G, Mansouri SS. Kt-biologics i (KTB1): a dynamic simulation model for continuous biologics manufacturing. Computers & Chemical Engineering 188:108770 (2024). https://doi.org/10.1016/j.compchemeng.2024.108770
2. Ebrahim A, Lerman JA, Palsson BO, Hyduke DR. Cobrapy: constraints-based reconstruction and analysis for python. BMC Syst Biol 7: (2013). https://doi.org/10.1186/1752-0509-7-74
3. Gudmundsson S, Thiele I. Computationally efficient flux variability analysis. BMC Bioinformatics 11: (2010). https://doi.org/10.1186/1471-2105-11-489
4. Liu J, Gao Q, Xu N, Liu L. Genome-scale reconstruction and in silico analysis of aspergillus terreus metabolism. Molecular BioSystems 9:1939-1948 (2013). https://doi.org/10.1039/c3mb70090a
5. Megchelenbrink W, Huynen M, Marchiori E. Optgpsampler: an improved tool for uniformly sampling the solution-space of genome-scale metabolic networks. PLoS ONE 9:e86587 (2014). https://doi.org/10.1371/journal.pone.0086587
6. Shoval O, Sheftel H, Shinar G, Hart Y, Ramote O, Mayo A, Dekel E, Kavanagh K, Alon U. Evolutionary trade-offs, pareto optimality, and the geometry of phenotype space. Science 336:1157-1160 (2012). https://doi.org/10.1126/science.1217405

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Modelling & optimization of recombinant protein production in a microbial cultivation with tunable induction

Philipp Pably[a]*, I Gede Eka Perdana Putra[a], Gerd Seibold[b], Jakob K. Huusom[a], and Julian Kager[a*]

[a] DTU, Chemical and Biochemical Engineering, Kgs. Lyngby, Denmark
[b] DTU, Bioengineering, Kgs. Lyngby, Denmark
* Corresponding Author: jukager@kt.dtu.dk

ABSTRACT

Recombinant protein production in *Escherichia coli* is a widely used system in industry for biopharmaceuticals, enzymes or other proteins. For protein expression, lactose poses as a more favorable and cost-effective induction agent over the common IPTG trigger. It imposes less stress on the cells and is fully metabolizable by the strain used. Therefore, lactose serves as an additional substrate source and adds a degree of freedom through tunable induction levels. To harness this opportunity, a physiological bioprocess model was created, describing the growth and production dynamics of this 2-feed system. Green fluorescent protein is expressed as a model protein in a fed-batch process using glucose as the main substrate and lactose as the digestible inducer. A suitable production kinetic is chosen by fitting a number of models to a collected dataset. The resulting model is used to highlight opportunities for improved process design and control of a 2-stage fed-batch process. It emphasizes the benefit of model-based methods to directly optimize for productivity over traditional design of experiment approaches.

Keywords: Fermentation, Modelling and Simulations, Biosystems, Process Design, Dynamic Modelling

INTRODUCTION

Escherichia coli (*E. coli*) is one of the most extensively used host systems to produce recombinant proteins. The product palette ranges from biopharmaceuticals over industrial enzymes to functional proteins for emerging food applications [1]. The organism offers a manifold of advantages such as fast growth, cheap substrate sources or broad availability of biotechnological tools [2, 3].

The common procedure in triggering the production of recombinant proteins utilizes the T7 promoter system. Naturally, it is activated by the presence of lactose in the media, while in industry the analogue chemical isopropyl β-d-1-thiogalactopyranosid (IPTG) is used. It is non-digestible by the cells and serves as a one-shot addition that induces a strong transcription of the encoded protein [4]. This triggers high levels of productivity but also elevated stress on the cells, giving rise to the formation of misfolded protein aggregations called inclusion bodies [5]. Using lactose as an digestible inducer allows for tunable induction, which provides an opportunity for improved process design. Previous studies have shown the positive impact of manipulating physiological parameters such as the specific uptake rates of glucose and lactose for improving productivity and inclusion body content [6]. Nevertheless, the used strains lacked the ability of metabolizing galactose, the first metabolic product of digesting lactose. This leads to metabolite accumulation, blocking further digestion of the inducer and even re-induction after the depletion of lactose, also called “Lazarus effect” [7]. To enable a fully controllable induction process, this study employed a strain that retains the full galactose metabolism.

Commonly, fed-batch processes are designed in two phases. After an initial batch-phase, a fed-batch regime is started with the aim to grow the biomass in the reactor. Microbial cells act as a catalyst for the production, therefore a higher concentration is favorable for increased output. In the last phase, the induction agent is added, and the cells are triggered to produce the encoded protein. The feeding rates and substrate levels in the reactor need to be efficiently controlled, as a surplus can lead to repression or byproduct accumulation [8].

The tunable induction system provides an opportunity to optimize the protein expression dynamically

throughout the changing process conditions of a cultivation. Capturing the intrinsic connections in a model lays the foundation for improved process design and advanced control strategies to ensure higher productivity over extended periods or move towards continuous operation [9].

For this purpose, the interaction of the two substrates onto cell growth and product expression needs to be studied. Previous works focused on describing the effects of two available sugars on the respective uptake rates and their interaction through classical design of experiment [6, 7]. This work expanded this approach to create a novel model, which describes the productivity as a function of the two independent substrate additions. Several experiments were carried out to retrieve data for the characterization and modelling of the presented strain's ability to produce green fluorescent protein (GFP) when using separate glucose and lactose feeds. The parameterized model was then utilized to dynamically optimize the two feed-rates for a 2-stage fed-batch process to improve production processes for different design objectives.

MATERIAL & METHODS

This contribution suggests a physiological growth and productivity model for a microbial process focusing on the impact of separate glucose and lactose feeds. The data was obtained through lab-scale cultivations in a glass stirred-tank bioreactor system (Biostream International, The Netherlands) with a maximum working volume of $2\,L$. In total 5 experiments were conducted and used for parameter estimation and validation. The modelling steps were implemented through custom Python scripts and followed the iterative workflow for the setup of mechanistic models by Daume et al. [10].

Cultivation conditions

The cultivation uses DeLisa minimal media [11] consisting of 13.3 g/L KH_2PO_4, 4 g/L $(NH_4)_2HPO_4$, 1.7 g/L citric acid, 1.2 g/L $MgSO_4{\cdot}7H_2O$, 0.1 g/L Fe(III) citrate, 8.4 mg/L EDTA, 1.3 mg/L $Zn(CH_3COO)_2{\cdot}2H_2O$, 4.5 mg/L thiamine HCl, 50 mg/L kanamycin, and 5 mL/L of a trace element solution. The trace element solution contained 2.5 mg/L $CoCl_2{\cdot}6H_2O$, 15 mg/L $MnCl_2{\cdot}4H_2O$, 1.5 mg/L $CuCl_2{\cdot}4H_2O$, 3 mg/L H_3BO_3, and 2.5 mg/L $Na_2MoO_4{\cdot}2H_2O$.

Strain

The *Escherichia coli* strain HMS174 (DE3) was used with a pET28a(+) plasmid incorporated to produce His-tagged green fluorescent protein (GFP) after induction. Production was triggered through the lac-operon by the presence of lactose. The cell line retained all enzymes of the Leloir pathway, enabling it to metabolize the inducer as well as all its primary metabolites, such as galactose.

Cultivation procedure

The cultivations were carried out in several steps. First, a cryogenic vial containing the cells was transferred from the cell bank to a shake flask containing 50 mL Luria-Bertain (LB) medium and incubated in a shaking cabinet (Infors HT Multitron, Switzerland) at 37 °C and 270 rpm. After 5-8 h 1 mL of the grown culture was transferred to 100 mL of DeLisa minimal media in a 500 mL shake flask and further incubated overnight.

On the next day the preculture media was transferred to the bioreactor system containing DeLisa media to form a total of 1 L batch volume to start the main cultivation. The bioreactor system used feedback control to maintain the temperature at 37 °C and the pH at 6.75 through the addition of 12.5 w-% ammonia. The dissolved oxygen level was kept above 30 % through a split-range PID controller acting on the stirring speed and the sparger airflow.

The reactor cultivations consisted of a batch phase, lasting until the initial substrate was depleted. Afterwards a fed-batch phase followed through the separate addition of a 300 g/L glucose feed and a 100 g/L lactose feed with peristaltic pumps. The feed-rates were implemented as feed-forward control to match certain specific uptake rates q for each experiment, resulting in exponential feeding. The added amounts of substrate were monitored through balances, and the recorded signal was used to correct the actual additions for the data analysis.

The rates investigated are listed in Tab. 1, where the different setpoints were maintained for several hours before switching to the next ones. The selected setpoint combinations resulted in 13 unique cultivation conditions. 9 conditions of 3 experiments were used to parameterize the proposed model structures and 2 experiments using 4 different setpoint combinations for validation and model selection.

Table 1: Specific uptake rate setpoints for parameterization (P) and validation (V) cultivation experiments.

Experiment		Type	Setpoint in g/g/h		
			1	2	3
1	$q_{S,Gluc}$	P		0.15	
	$q_{I,Lac}$		0	0.075	0.15
2	$q_{S,Gluc}$	P	0.325	0.125	
	$q_{I,Lac}$		0	0.0375	0.25
3	$q_{S,Gluc}$	P	0.325	0.375	
	$q_{I,Lac}$		0	0.0375	0.25
4	$q_{S,Gluc}$	V	0.375	0.15	0.075
	$q_{I,Lac}$			0.225	
5	$q_{S,Gluc}$	V		0.1125	
	$q_{I,Lac}$			0.225	

Sample analysis

An autosampler (Biolab A/S, Denmark) collected cultivation samples automatically every 1-2 h, which were immediately cooled to 4 °C and stored for analysis.

Biomass concentration was determined in duplicates through dry cell weight (DCW) analysis. 1.8 mL of sample were centrifuged for 5 min at 10, 000 rcf, decanted and subsequently washed with 1 mL of 0.9 % (w/v) NaCl solution. The initial supernatant was frozen and the cell pellet was dried at 105 °C for >72 h.

The supernatant sample was analyzed for its metabolites content. The glucose concentration was obtained through an enzymatic assay (Megazyme International Ireland Ltd., Ireland). The absorbance changes of a standard curve and the diluted cultivation samples in triplicates were measured with a 96-well microplate reader (Tecan, Switzerland). Lactose, galactose, acetate and ethanol concentration were obtained through high-performance liquid chromatography (HPLC) using the Aminex HPX-87H column (Bio-Rad, USA).

GFP concentrations were obtained in triplicates from diluted cultivation samples with a microplate reader (Tecan, Switzerland), using black flat-bottom 96-well plates. The samples were excited at 485±10 nm and the emissions at 530±10 nm were measured. The resulting fluorescence measurements were converted from RFU to ng GFP/well using a standard curve prepared with the GFP quantification kit (Abcam, Cambridge, UK).

Process model

The *E. coli* process was modeled as a state-space model. The chosen state vector to describe the dynamics was $x = [V_R, c_X, c_S, c_I, c_P]^T$, consisting of the reactor volume V_R in L, the concentrations for biomass c_X, the main substrate glucose c_S, the inducing agent lactose c_I in g/L, as well as the product concentration c_P in mg/L.

$$\frac{d}{dt}\begin{bmatrix} V_R \\ c_X \\ c_S \\ c_I \\ c_P \end{bmatrix} = \begin{bmatrix} F_{in,S} + F_{in,I} - F_{out} \\ \left(\mu - \frac{F_{in,S}+F_{in,I}}{V_R}\right) \bullet c_X \\ c_{S,in} \bullet \frac{F_{in,S}}{V_R} - q_S \bullet c_X - c_S \bullet \frac{F_{in,S}+F_{in,I}}{V_R} \\ c_{I,in} \bullet \frac{F_{in,I}}{V_R} - q_I \bullet c_X - c_I \bullet \frac{F_{in,S}+F_{in,I}}{V_R} \\ q_P \bullet c_X - c_P \bullet \frac{F_{in,S}+F_{in,I}}{V_R} \end{bmatrix} \quad (1)$$

The state update vector dx/dt in eq. (1) contains the mass balance of the system together with the biological kinetics of microbial growth and product expression. The glucose consumption q_S was modeled as a Monod kinetic on the glucose concentration in eq. (2). The specific lactose uptake rate q_I was setup as a Monod kinetic on the lactose concentration, but with an additional inhibition term related to the specific glucose uptake rate, see eq. (3). This followed the assumption that glucose is the more bioavailable substrate source and the cells prefer it over the disaccharide lactose, similar to Wurm et al. [12].

$$q_S = q_{S,max} \bullet \frac{c_S}{c_S + K_S} \quad (2)$$

$$q_I = q_{I,max} \bullet \frac{c_I}{c_I + K_I} \bullet \frac{K_{IS}}{q_S + K_{IS}} \quad (3)$$

The specific growth rate μ was set up as a combination of the conversion rates of the two available substrate sources minus the maintenance rate q_m. The collective substrate conversion rate was then multiplied with a lumped biomass yield in eq. (4).

$$\mu = (q_S + q_I - q_m) \bullet Y_{X/S} \quad (4)$$

The formulation of the specific product formation rate q_P was found by testing different common kinetic expressions for bioprocesses, listed in Tab. 2. The general structure considered the influence of the two substrates, glucose and lactose influencing a maximum productivity $q_{P,max}$, following eq. (5).

$$q_P = q_{P,max} \bullet \mathrm{f}(c_S \; or \; q_S) \bullet \mathrm{f}(c_I \; or \; q_I) \quad (5)$$

Model identification and parameterization

The model was parameterized in two main steps. First the growth model was fitted to the process data without the product formation dynamics. The measured values for DCW, glucose and lactose concentration were used to minimize the normalized root mean squared error (NRMSE) to the simulation values.

In a second step the tuned growth model was extended with different models for the specific product formation q_P, based on Schügerl and Bellgardt et al. [13]. The model structures listed in Table 2 were tuned to fit the offline GFP measurements and the resulting NRMSE values for the product concentration were used to identify the best model formulation.

Process optimization

The final parameterized model was used to design an optimized cultivation process. The cultivation layout consisted of an initial batch followed by an exponential fed-batch phase by adding glucose with the setpoint $q_{S,grow}$. The objective of this phase was to increase the biomass content, therefore the setpoint was fixed with the maximum substrate uptake rate. After a certain time t_{ind} the process focus was switched to production where glucose and lactose were added with the new setpoints of $q_{S,ind}$ and $q_{I,ind}$ respectively. These 3 parameters defining the induction phase were tuned to optimize the process for 3 different objectives, by maximizing the performance indicators listed in Tab. 3. The process was designed to start with $V_{R,0} = 1\,L, c_{X,0} = 0.1\,g/L$ and $c_{S,0} = 8\,g/L$. The feeding was stopped after $100\,g$ of combined substrate was added and fully consumed. With the parameterized biomass/substrate yield $Y_{X/S}$ the expected final biomass was ~$31.5\,g/L$.

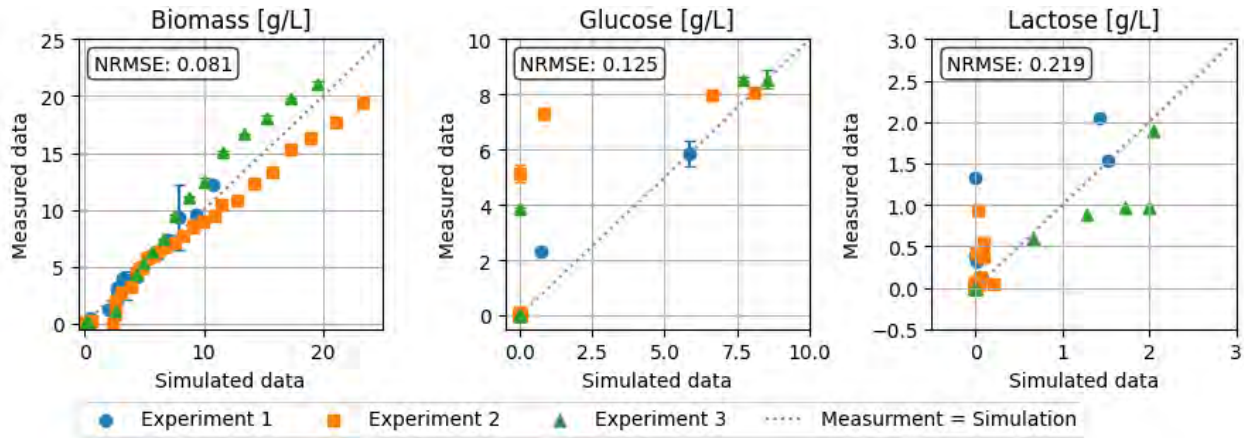

Figure 1. Comparison of growth model fit with NRMSE values for each variable of the parameterization datasets.

Table 2: Potential productivity models

Model	Equation
P1	$q_P = q_{P,max} \bullet \frac{q_I}{q_I+K_{PI}}$
P2	$q_P = q_{P,max} \bullet \frac{c_I}{c_I+K_{PI}}$
P3	$q_P = q_{P,max} \bullet \frac{{q_I}^{k_{PN}}}{{q_I}^{k_{PN}}+{K_{PI}}^{k_{PN}}}$
P4	$q_P = q_{P,max} \bullet \frac{{c_I}^{k_{PN}}}{{c_I}^{k_{PN}}+{c_{PI}}^{k_{PN}}}$
P5	$q_P = q_{P,max} \bullet \frac{q_I}{q_I+K_{PI}} \bullet \frac{K_{PS}}{q_S+K_{PS}}$
P6	$q_P = q_{P,max} \bullet \frac{q_I}{q_I+K_{PI}} \bullet \frac{{K_{PS}}^{k_{PN}}}{{q_S}^{k_{PN}}+{K_{PS}}^{k_{PN}}}$
P7	$q_P = q_{P,max} \bullet \frac{q_I}{q_I+K_{PI}} \bullet \frac{q_S}{q_S+K_{PS}}$
P8	$q_P = q_{P,max} \bullet (\frac{q_I}{q_I+K_{PI}} + \alpha \bullet \frac{q_S}{q_S+K_{PS}})$
P9	$q_P = q_{P,max} \bullet \frac{q_I}{q_I+K_{PI}} \bullet (1 + \beta \bullet \frac{q_S}{q_S+K_{PS}})$
P10	$q_P = q_{P,max} \bullet \frac{q_I}{q_I+K_{PI}} \bullet \left(1 + \beta \bullet \frac{q_S}{q_S+K_{PS}}\right) + q_{P,\mathrm{base}}$
P11	$q_P = q_{P,max} \bullet \frac{q_I}{q_I+K_{PI}} \bullet \frac{q_S}{q_S+K_{PS}} + q_{P,\mathrm{base}}$
P12	$q_P = q_{P,max} \bullet \frac{c_I}{c_I+K_{PI}} \bullet \frac{K_{PS}}{c_S+K_{PS}}$

Table 3: Investigated process objectives

Objective	Symbol	Equation
Total product	P	$c_{P,end} \bullet V_{R,end}$
Product concentration	c_P	$c_{P,end}$
Space-time yield	STY	$c_{P,end}/t_{end}$

RESULTS & DISCUSSION

The model building process was separated into a procedure for the growth dynamics on the two separate substrate feeds and the selection of a fitting productivity kinetic. The final model structure was then used for a process optimization simulation study.

Growth model parameterization

The described growth model was parameterized by minimizing the NRMSE between model and offline samples for DCW, glucose and lactose concentration, using all available datasets. The resulting parameter values are listed in Tab. 4.

Table 4: Fitted growth model parameters.

Parameter	Symbol	Value	Unit
Max specific glucose uptake rate	$q_{S,max}$	2.350	g/g/h
Monod constant glucose	K_S	0.014	g/L
Max specific lactose uptake rate	$q_{I,max}$	0.456	g/L
Monod constant lactose	K_I	0.005	g/L
Monod constant lactose-glucose inhibition	K_{IS}	0.270	g/L
Specific maintenance coefficient	q_m	0	g/g/h
Biomass/substrate yield	$Y_{X/S}$	0.296	g/g

The performance of the model is visualized in Fig. 1 through a simulated vs measured plot for the 3 variables. The data showed a satisfying fit for all variables, where the biomass was described the best with a remaining NRMSE of 0.081. The glucose concentration showed a similarly low level of 0.125, the deviations stemmed mostly from a slower growth dynamic in the batch phase of the measured data in Experiment 2. The lactose model showed the highest remaining error of 0.219. Overall, the model fit was deemed as satisfactory in face of the low amount of experimental data.

Production model parameterization

The production model was obtained by testing

different kinetic formulations, listed in Tab. 2. Their performance to describe the productivity was evaluated through the NRMSE and the conditional number, using experiment 1-3 for parameterization and 4-5 as validation sets. Fig. 2 depicts the NRMSE values for the parameterization and validation datasets for each model formulation.

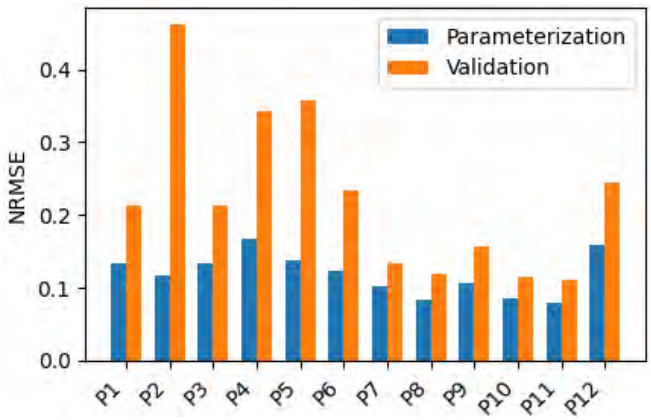

Figure 2. Productivity model NRMSE values for parameterization and validation datsets.

The chosen productivity kinetic for the final model was P11, with the parameter values given in Tab. 5. It showed the overall lowest remaining error of the validation dataset compared to the other options with 0.111. The selected model structure in eq. 6 consisted of a positive Monod type influence for both substrate uptake rates q_S and q_I., with an additional constant base expression rate $q_{P,base}$.

$$q_P = q_{P,max} \cdot \frac{q_I}{q_I + K_{PI}} \cdot \frac{q_S}{q_S + K_{PS}} + q_{P,base} \tag{6}$$

Table 5: Fitted growth model parameters.

Parameter	Symbol	Value	Unit
Max specific GFP productivity	$q_{P,max}$	2.017	mg/g/h
Monod constant glucose-product	K_{PS}	0.361	g/g/h
Monod constant lactose-product	K_{PI}	0.147	g/g/h
Base specific GFP productivity	$q_{P,base}$	0.06	mg/g/h

The inclusion of a base product expression rate hints at a certain leakiness of the expression system. This means that the microbes showed GFP expression without an induction trigger present. Further, the model fit concluded that both substrate uptake rates have a positive influence on productivity. Nevertheless, the growth model included an inhibition term in the lactose uptake related to high glucose uptakes, limiting the double enhancing effect of high substrate feed-rates. This influence of glucose and lactose uptake rates onto the modeled specific product expression rate values are shown in Fig. 3.

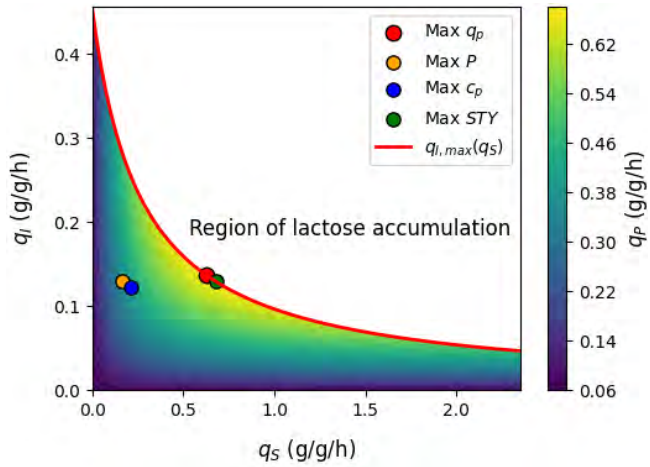

Figure 3. Contour plot of the parameterized productivity q_P as a function of the substrate uptake rates q_S and q_I, including the optimal setpoints for different objectives.

Process optimization – Simulation study

The model structure suggested that the combination of moderate glucose and lactose uptake rates results in the highest specific productivity. Nevertheless, the biomass concentration is the catalyst in translating the specific rates to absolute product amounts in the reactor. Since high substrate feed-rates favor biomass formation, a non-trivial optimization problem is presented.

A global optimization algorithm was applied to optimize the induction phase and different optimal feeding strategies for each of the 3 objectives listed in Tab. 3 were found. All simulations were limited to add a maximum of 100 g of combined substrate to the process. The optimal specific substrate uptake rates are marked in the contour plot in Fig. 3 and the resulting objective values for each process are compared in Tab. 6.

Table 6: Performance values of optimized processes.

Objective function	P_{end} [mg]	$c_{P,end}$ [mg/L]	STY [mg/L/h]
max(P_{end})	121.18	74.41	2.17
max($c_{P,end}$)	119.43	75.77	2.49
max(STY)	82.51	57.30	3.70

The resulting profiles in Fig. 4 show that the optimal feeding strategies are dependent on the objective. The highest absolute product content was achieved with lower setpoints for the two specific substrate uptake rates. With a focus on higher space-time-yield a faster feeding regime was proposed. All the processes included

only a short initial growth phase, where solely the glucose feeding was active at the maximum substrate uptake rate to build more initial biomass. Therefore, higher objective values in all categories were achieved when both feeds were active simultaneously for the fed-batch phase.

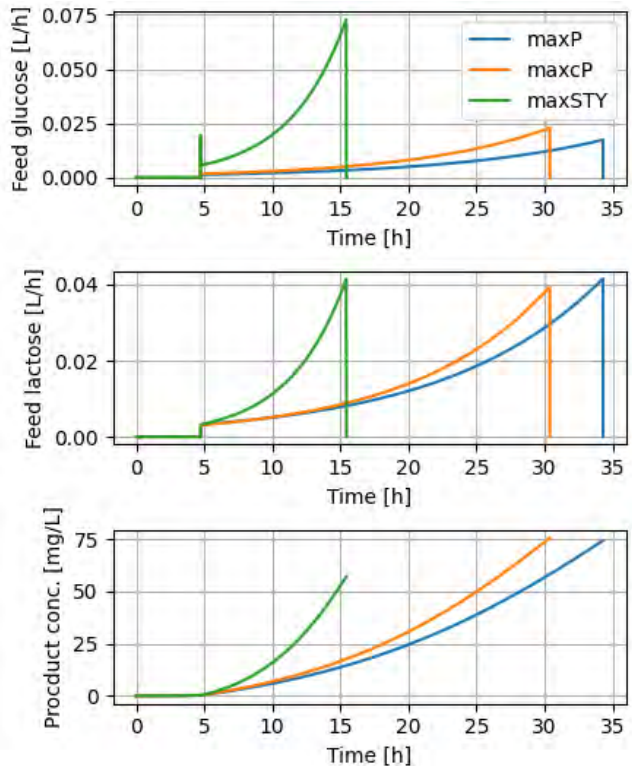

Figure 4. Optimized feeding strategies and product concentration for three different process objectives for the 2-feed process.

CONCLUSION

A novel physiological model is presented for growth and productivity of an *E. coli* process which uses two independent feed sources, glucose as main substrate and lactose as digestible inducer and opens the possibility of tunable induction. The experimental data of 5 lab-scale processes was used to identify and parameterize a suitable model structure. The results showed that both substrate uptake rates were enhancing the product formation, which is in line with the observations of Wurm et al. [6, 12] where low glucose uptake resulted in lower productivity. The lactose uptake on the other hand was repressed by increased glucose uptake rates, which created a non-trivial optimization problem.

By optimizing the individual feed-rates for a simulated 2-stage fed-batch process, the dynamic model suggested moderate substrate and inducer uptake rates result in the highest productivity. Different objective functions were investigated and showed that for a high space-time yield setpoints close to the maximum specific product formation rate were favorable, while for high titers a slower feeding regime is preferred.

This contribution established a physiological bioprocess model for recombinant protein production in *E. coli* with tunable induction through separate glucose and lactose feeds. The resulting model presents an opportunity for improved process design and advanced process control such as described in Ulonska et al. [9] and shows potential for enabling continuous operation modes.

ACKNOWLEDGEMENTS

This work obtained funding from the Novo Nordisk Foundation, Start Package grants for faculty recruitment Q4 2022; assistant professor on bioprocess digitalization and control (NNF220C0081250).

REFERENCES

1. Rani AK, Katiyar R, Rathore AS. Bioprocessing of inclusion bodies from E. coli. to produce bioactive recombinant proteins. Biochemical Engineering Journal 203:109188 (2024). https://doi.org/10.1016/j.bej.2023.109188
2. Baneyx F. Recombinant protein expression in escherichia coli. Current Opinion in Biotechnology 10:411-421 (1999). https://doi.org/10.1016/s0958-1669(99)00003-8
3. Francis DM, Page R. Strategies to optimize protein expression ine. coli. CP Protein Science 61: (2010). https://doi.org/10.1002/0471140864.ps0524s61
4. Marbach A, Bettenbrock K. Lac operon induction in escherichia coli: systematic comparison of IPTG and TMG induction and influence of the transacetylase laca. Journal of Biotechnology 157:82-88 (2012). https://doi.org/10.1016/j.jbiotec.2011.10.009
5. Lin W, Huang S, Chou CP. Degp?coexpression minimizes inclusion?body formation upon overproduction of recombinant penicillin acylase in escherichia coli. Biotech & Bioengineering 73:484-492 (2001). https://doi.org/10.1002/bit.1083
6. Wurm DJ, Veiter L, Ulonska S, Eggenreich B, Herwig C, Spadiut O. The E. coli pet expression system revisited—mechanistic correlation between glucose and lactose uptake. Appl Microbiol Biotechnol 100:8721-8729 (2016). https://doi.org/10.1007/s00253-016-7620-7
7. Kittler S, Kopp J, Veelenturf PG, Spadiut O, Delvigne F, Herwig C, Slouka C. The lazarus escherichia coli effect: recovery of productivity on glycerol/lactose mixed feed in continuous biomanufacturing. Front. Bioeng. Biotechnol. 8:

(2020). https://doi.org/10.3389/fbioe.2020.00993
8. Stülke J, Hillen W. Carbon catabolite repression in bacteria. Current Opinion in Microbiology 2:195-201 (1999). https://doi.org/10.1016/s1369-5274(99)80034-4
9. Ulonska S, Waldschitz D, Kager J, Herwig C. Model predictive control in comparison to elemental balance control in an E. coli fed-batch. Chemical Engineering Science 191:459-467 (2018). https://doi.org/10.1016/j.ces.2018.06.074
10. Daume, S., Kofler, S., Kager, J., Kroll, P., and Herwig, C. Generic Workflow for the Setup of Mechanistic Process Models. *Animal Cell Biotechnology: Methods and Protocols* 2095: 189-211 (2020) https://doi.org/10.1007/978-1-0716-0191-4
11. DeLisa, M. P., Li, J., Rao, G., Weigand, W. A., and Bentley, W. E. Monitoring GFP-operon fusion protein expression during high cell density cultivation of Escherichia coli using an on-line optical sensor. *Biotechnology and Bioengineering* 65.1: 54-64 (1999) https://doi.org/10.1002/(SICI)1097-0290(19991005)65:1%3C54::AID-BIT7%3E3.0.CO;2-R
12. Wurm DJ, Hausjell J, Ulonska S, Herwig C, Spadiut O. Mechanistic platform knowledge of concomitant sugar uptake in escherichia coli BL21(DE3) strains. Sci Rep 7: (2017). https://doi.org/10.1038/srep45072
13. Schügerl, K., Bellgardt, K.-H. (Eds.) Bioreaction Engineering. Springer Berlin Heidelberg.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Automatic k_La determination in stirred tank reactors by model-based design of experiments

Ana Helena V. Caetano[a]*, Krist V. Gernaey[a], and Julian Kager[a]

[a] DTU, Department of Chemical and Biochemical Engineering, Kgs. Lyngby, Denmark
* Corresponding Author: ahvca@kt.dtu.dk.

ABSTRACT

The volumetric gas-liquid mass transfer coefficient (k_La) is a key performance parameter in stirred tank reactors and is commonly determined through extensive experiments across the operational space. This work presents an automatic, closed-loop framework for k_La determination based on model-based design of experiments (MBDoE), in which agitation and aeration inputs are adapted in real time.During each experiment, dissolved oxygen data is collected and used to estimate the parameters of a Van't Riet k_La relation. The parameter uncertainty is quantified using the covariance matrix, and the experiments are iteratively selected based on D-optimality or E-optimality MBDoE, until a threshold of $RSE_i < 0.15$ is reached for all parameters. The MBDoE approach is evaluated through repeated runs and compared against random designs, full factorial (FF) design, and a full grid design.The results demonstrate that the closed-loop MBDoE framework can significantly reduce the number of experiments required to characterize k_La while maintaining accurate predictions of gas-liquid mass transfer over the design space. D-optimal MBDoE converges in fewer than four experiments on average, while E-optimal MBDoE converges in less than six experiments on average, translating into approximately a 60% reduction of the number of experiments as compared to a full factorial DoE design.

Keywords: Optimization, Modelling, Numerical Methods, Model-based design of experiments, Gas-liquid mass transfer, Stirred tank reactors

INTRODUCTION

k_La identification as an experiment design problem

The volumetric gas-liquid mass transfer coefficient (k_La) is a key performance parameter of stirred tank reactors and its determination can be vital to avoid oxygen limitations and to optimize bioprocesses. In practice, the k_La can be affected by several different factors such as tank geometry, operational conditions and properties of the mixture/fluid/medium [1]. Consequently, k_La is typically determined experimentally by running multiple experiments under different combinations of operational inputs, followed by fitting an empirical correlation to the experimental data [2].

Fundamentally, this is an experimental design problem. Each experiment results in different excitation profiles of the operational inputs and consequently in different information content. Therefore, each experiment provides different information about the underlying k_La model. Because experiments in stirred tanks can be time-consuming and resource-intensive, reducing the number of experiments required to obtain a reliable k_La estimate is important for rapid bioreactor characterization and process development.

Limitations of classical design strategies

Classical k_La characterization strategies rely on predefined experimental designs, such as full factorial designs and stepwise variation of the input variables [3]. Although these approaches can give a good coverage of the design space, they do not exploit the information gained during the experiments, *i.e.*, later experiments are not adapted based on what has already been learned, which can lead to redundancy.

Model-based design of experiments (MBDoE) defines a structured way to iteratively select the input trajectories of the experiments that maximize information about the system. Along the experiments there is a

constant and iterative retrofitting of the prototype model with new data, which ensures minimal experimental redundancy and inclusion of information gathered throughout the process [4]. Thus, a closed-loop approach is established, in which experimental conditions are adapted in real time based on continuously updated parameter estimates and uncertainty information.

Contributions of this work

In this work, we present an automatic, closed-loop framework for k_La determination in stirred tank reactors using MBDoE. The system iteratively performs experiments in which agitation and aeration are varied, the experiment is run, the data of the dissolved oxygen trajectory is collected in real time, the parameters of the Van't Riet k_La correlation [5] are estimated, and the next most informative experiment to perform is computed.

The methodology is evaluated on a stirred-tank k_La case study, using D-optimal and E-optimal MBDoE strategies and it is compared against random designs, full factorial DoE designs, and a grid of 100 experiments with uniform stepwise variation of the two inputs covering the full design space (referred to as full-grid). The performance is assessed in terms of the number of experiments required to reach a target level of parameter uncertainty, as well as the accuracy of the resulting k_La predictions across the design space.

MATERIALS AND METHODS

Experimental system

The experiments were conducted in a lab-scale Biostream glass stirred tank reactor, equipped with a top mounted Rushton impeller, a dissolved oxygen (DO) probe, mass flow controller for air and nitrogen supply, and a Biostream tower. The tower is connected to the Lucullus software (Securecell, Urdorf, Switzerland), thus facilitating data acquisition and remote control. A Rest API connection is used to communicate between the software and python scripts, thus allowing for a full closed-loop automated system.

The reactor was operated in a gas-liquid system using 2L of water without biological consumption, so that oxygen transfer is governed exclusively by physical mass transfer.

The agitation speed and the gas flow rate were used as the manipulated variables. For each experiment, the DO concentration was recorded continuously, from which the experimental k_La corresponding to the operational inputs was calculated. This was computed by performing a dynamic oxygen balance model (Eq. 1) and using nonlinear regression techniques.

$$\frac{dC}{dt} = k_L a\,(C^* - C) \tag{1}$$

Where C represents the DO concentration at time t and C^* is the DO concentration at saturation.

The experiments are performed using the gassing-out method [5], where nitrogen gas is sparged through the liquid to strip it of DO before a k_La experiment.

The design space is restricted to air flow [0.4 : 4.0] L/min and stirrer speed [100 : 1000] rpm.

k_La model, parameter estimation and uncertainty

The volumetric mass transfer coefficient is described using a Van't Riet correlation [5], which relates k_La to the power input P (directly related to agitation speed) and superficial gas velocity v_g (directly related to aeration rate), and reactor geometry, where V_R is the reactor volume (Eq. 2).

$$k_L a = K\left(\frac{P}{V_R}\right)^{\alpha} v_g^{\beta} \tag{2}$$

For each experiment, the model parameters (K, α, β) were estimated by minimizing the normalized root-mean-square error (NRMSE) between the measured and simulated k_La.

Parameter uncertainty was quantified by relative standard errors (RSE), computed based on the covariance matrix (Cov) of local sensitivities. The local sensitivity matrix (S) is computed based on central differences with a relative step of 1%. Standard errors (SE) are computed to quantify the spread of the estimate around its value and scaled by the parameter value (Eq. 3 – 4).

$$SE_i = \sqrt{Cov(\hat{\theta})_{ii}} \tag{3}$$

$$RSE_i = \left|\frac{SE_i}{\hat{\theta}_i}\right| \tag{4}$$

Closed-loop MBDoE workflow

The k_La identification was performed using an iterative, closed-loop MBDoE framework (Figure 1). Each run started from a predefined initial experiment, defined at the centre of the design space. After each experiment, the collected DO data were used to re-estimate the model parameters. Based on the calibrated model, sensitivities and uncertainties are computed with respect to each of the parameters. The updated model is then used to solve the MBDoE optimization problem and suggest the next most informative experiment. For this, a genetic algorithm is used, by sampling, *in silico* and across the design space, new experiments to perform. The objective is to maximize the expected information gain of the next experiment. The resulting inputs are then applied to the reactor, generating a new DO time series, and the procedure is repeated. The iteration can be terminated when all the model parameters are below a threshold of parameter uncertainties or a threshold of the maximum allowed number of experiments performed has been reached.

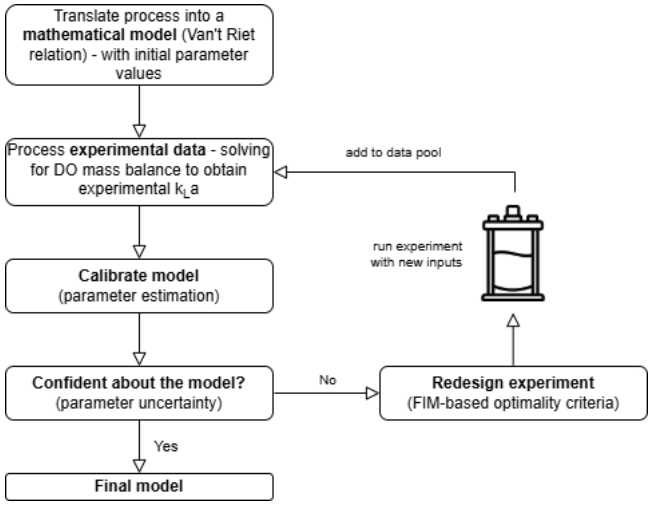

Figure 1. Closed-loop MBDoE worflow for the k_La model.

Design strategies

The information content given by the closed-loop MBDoE framework was quantified using two different optimality criteria: D-optimality and E-optimality. D-optimality or determinant optimality, aims to maximize the determinant of the Fisher Information Matrix (FIM), thereby minimizing the overall parameter uncertainty. E-optimality or eigenvalue optimality, aims to maximize the smallest eigenvalue of FIM, thus targeting a balanced identifiability across all parameters.

$$\Phi_{D-opt} = \det(\mathrm{FIM}) \tag{5}$$

$$\Phi_{E-opt} = \lambda_{min}(\mathrm{FIM}) \tag{6}$$

To benchmark the performance of the MBDoE framework, other alternative designs were evaluated, such as random uniform design, full factorial (FF) design and full grid design.

All design methods were evaluated using the same benchmark $RSE_i < 0.15$. For each method that does not rely on pre-setting the operational inputs, 10 independent runs were performed, so to account for stochastic effects. The FF DoE were performed in triplicates.

All final parameter estimates were evaluated against a full grid reference by computing the NRMSE of the predicted k_La surface over the design space.

RESULTS AND DISCUSSION

Number of experiments to converge

The primary performance metric evaluated was the number of experiments required to satisfy the stopping criterion $RSE_i < 0.15$ for all model parameters. Figure 2 summarizes the mean and standard deviation of the required number of experiments over 10 independent runs for each design strategy.

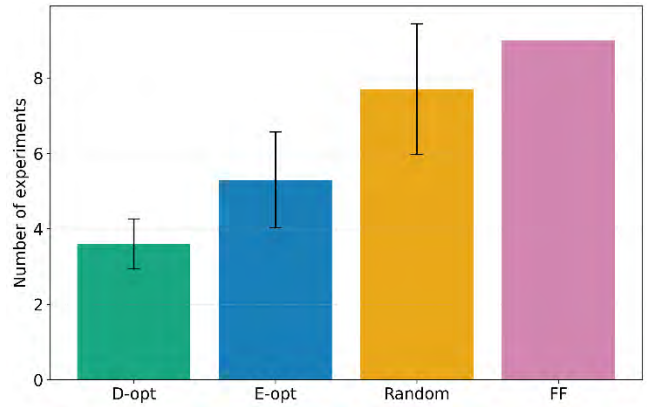

Figure 2. Number of experiments required to reach $RSE_i < 0.15$ for all parameters.

Both closed-loop MBDoE strategies substantially reduced the number of experiments necessary to reach the target precision as compared to the benchmark designs. In particular, the D-optimal MBDoE method converged in significantly fewer experiments with a mean of 3.60 ± 0.66 number of experiments. The E-optimal MBDoE strategy required 5.30 ± 1.27 experiments on average, while the random strategy required 7.70 ± 1.73 experiments to reach the same parameter uncertainty threshold. The full factorial design required 9 experiments by construction, while the full grid consisted of 100 experiments.

These results show a substantial reduction in experimental effort when using a closed-loop MBDoE design. D-optimal MBDoE reduced the number of experiments by 60% and 47% respectively, compared to a full factorial design and a random strategy. The relatively small standard deviation of the D-optimal MBDoE further indicates robustness to stochastic variation between runs.

Parameter uncertainty collapse

Figure 3 shows the evolution of the RSE of the parameters K, α, β of the Van't Riet relation as a function of the experiment number, averaged over 10 independent runs.

For the D-optimal and E-optimal MBDoE strategies, the parameter uncertainties decrease rapidly during the first few experiments. Random designs reduce uncertainty more slowly and with larger variability, reflecting the lack of adaptation to the collected data thus far.

These results show how real-time recalibration of parameter estimates pushes the MBDoE framework to select highly informative agitation and aeration inputs.

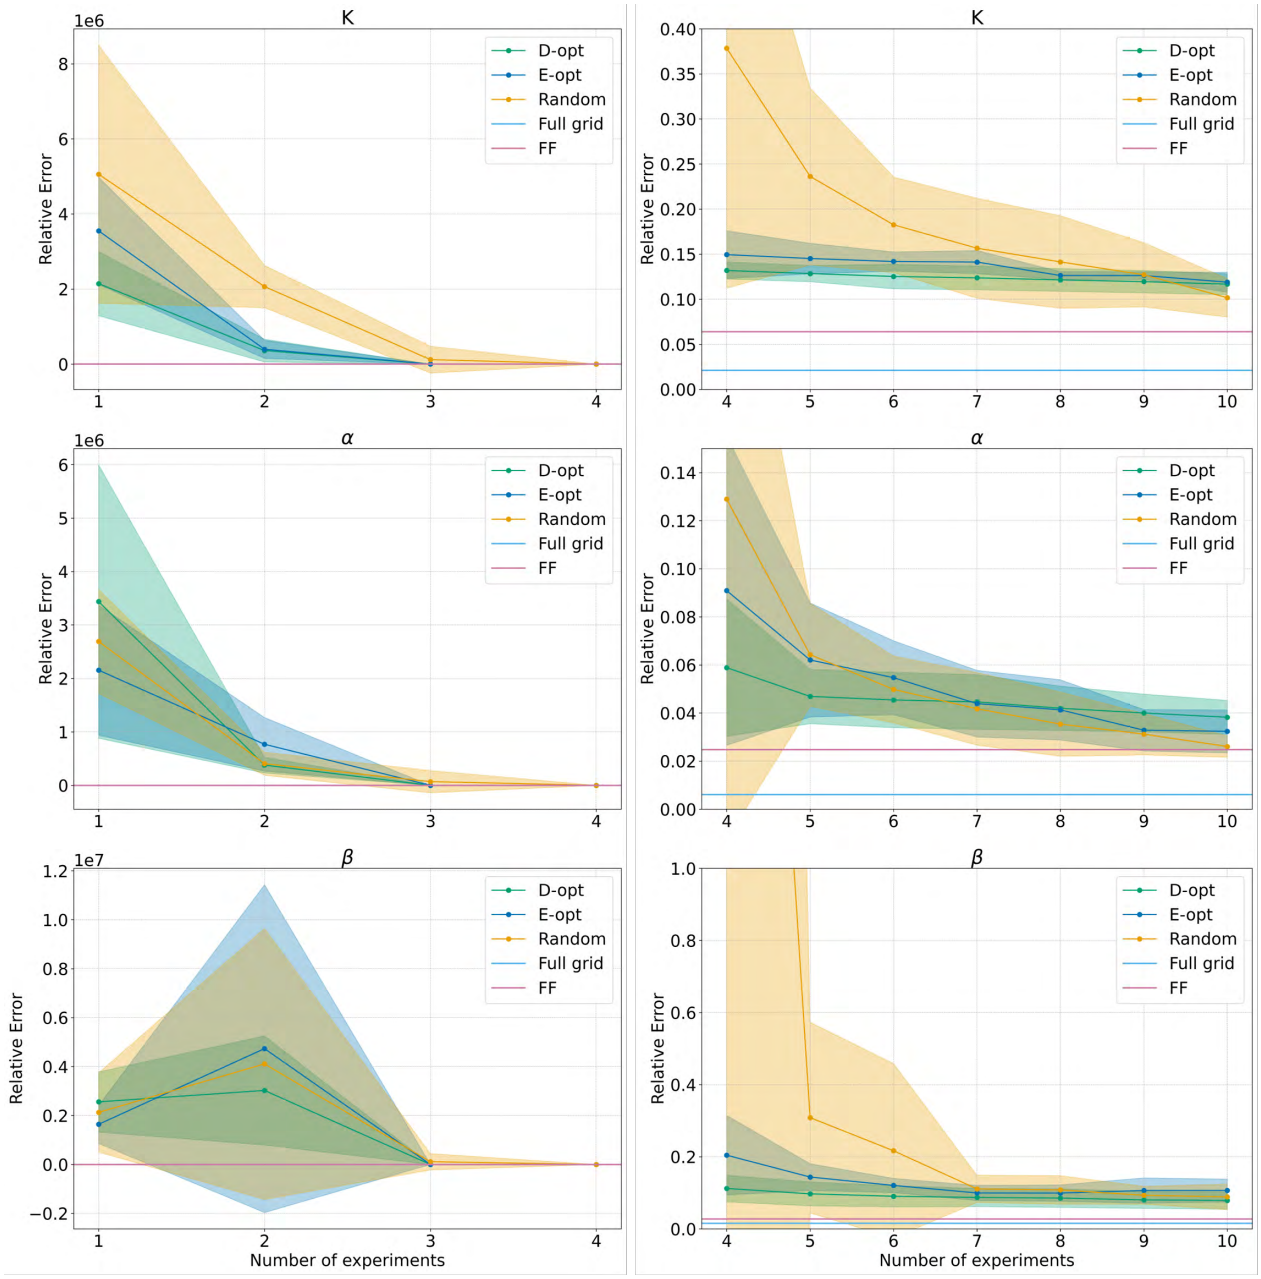

Figure 3. Evolution of the parameters RSE along the experiments for 10 independent runs.

k_La surface accuracy

Although the closed-loop algorithm terminates based on parameter uncertainty, the practical objective of the k_La characterization is an accurate prediction of this coefficient over the operational space. In order to assess this, we compare the k_La surfaces that were obtained by the different designs when they reached convergence.

For each design strategy, the final parameter estimates were averaged across the ten independent runs. This parameter set is the one used in the model (Eq. 2) to compute the k_La surface over the operational space of agitation and aeration. The resulting k_La surfaces are compared pointwise to the reference k_La surfaces by NRMSE (Figure 4).

The k_La surface references used are (a) the model fit to the full grid design of 100 experiments and (b) the model fit to the full factorial (FF) dataset. When compared against the full factorial reference, the D-optimal MBDoE strategy achieved the lowest deviation, with an NMRSE of 0.020. The E-optimal MBDoE and random strategies showed larger discrepancies, with NRMSE values of 0.092 and 0.122, respectively. This shows that the

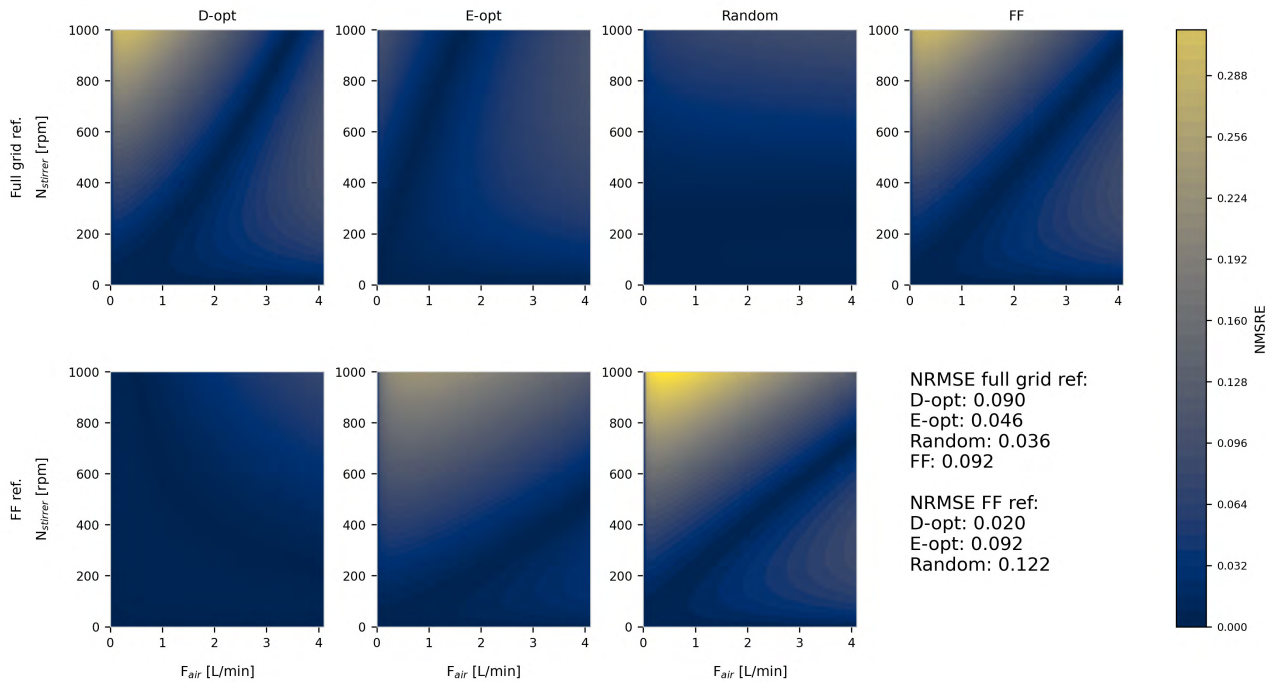

Figure 4. NRMSE of the predicted k_La surfaces in reference to the full grid design (upper row) and full factorial design (lower row). The aggregated NRMSE per design and per reference are listed on the bottom right.

closed-loop MBDoE approach achieves a similar k_La surface as the one learned by the FF dataset using substantially fewer experiments – 3.60 ± 0.66 as opposed to 9.

However, when comparing to the full grid reference, all design strategies show similar NRMSE ranges, although neither FF nor a MBDoE approach resulted in a better accuracy. Because different methods converge to different parameter estimates, these results further suggest that multiple parameter sets can produce identical k_La predictions across the design space.

Identifiability and model sloppiness

The differences observed between the two reference comparisons highlight a fundamental limitation of the Van't Riet model parametrization used in this study. Although this model provides an empirical description of k_La which is widely accepted, its parameters are structurally non-identifiable (over the explored design space)[2, 6].

This consequently results in the situation that distinct parameter vectors obtained by fitting the Van't Riet model to different experimental datasets, can generate very similar k_La surfaces. In such cases, low parameter uncertainty does not necessarily imply a better model. The accuracy of the predicted k_La surface can provide a more meaningful measure of the model quality.

Nevertheless, it is important to note that, in practical applications, the true parameter values are not known a priori and thus the objective of MBDoE is not to identify a unique "true" parameter set, but to obtain a model that can predict the k_La behaviour reliably with minimal experimental burden. Hence, the here proposed closed-loop MBDoE framework remains effective.

These results further indicate that the MBDoE stopping-criterion is better interpreted as an indication of prediction confidence rather than an indication of convergence to unique parameter vectors for a Van't Riet relation.

CONCLUSIONS

This work presented an automatic, closed-loop MBDoE framework for k_La determination in stirred tank reactors. By iteratively recalibrating the parameter estimates and uncertainty information from time-series data of dissolved oxygen, this approach autonomously selects agitation and aeration inputs that maximize information gain.

The results demonstrated that closed-loop MBDoE can substantially reduce the number of experiments necessary to reach a confidence threshold as compared to other design strategies, namely full factorial DoE, full grid over the operational space, and random inputs selection. In particular, a D-optimal MBDoE approach achieves the set uncertainty threshold of $RSE_i < 0.15$ with a 60% reduction in the required number of experiments as compared to a full factorial design.

Furthermore, comparison of the predicted k_La surfaces against full factorial and full grid designs showed

that, although the parameter values obtained differ between design strategies, the resulting k_La predictions over the design space are often comparable. This reflects the lack of identifiability of the Van't Riet relation and further suggests that the uncertainty-based stopping criterion indicates mostly a convergence of predictive behaviour rather than convergence to unique model parameters values.

Future and ongoing work shall extend the proposed framework to include additional operational variables, such as temperature, pressure, ionic strength, and viscosity, all important factors to gas-liquid mass transfer [1].

The methodology here suggested provides a foundation for autonomous reactor characterization, in particular for situations where experimental efficiency is of the most importance.

ACKNOWLEDGEMENTS

Financial support was provided by the Novo Nordisk Foundation Start Package grant for Bioprocess Control and Digitalization (NNF220C0081250).

AUTHOR IDENTIFIERS

Author ORCIDs:
Caetano AHV: 0009-0008-2951-4159
Gernaey KV: 0000-0002-0364-1773
Kager J: 0000-0001-7274-6678

REFERENCES

1. Garcia-Ochoa F, Gomez E. Bioreactor scale-up and oxygen transfer rate in microbial processes: an overview. Biotechnology Advances 27:153-176 (2009). https://doi.org/10.1016/j.biotechadv.2008.10.006
2. Schaepe S, Kuprijanov A, Sieblist C, Jenzsch M, Simutis R, Lübbert A. Kla of stirred tank bioreactors revisited. Journal of Biotechnology 168:576-583 (2013). https://doi.org/10.1016/j.jbiotec.2013.08.032
3. Aroniada M, Maina S, Koutinas A, Kookos IK. Estimation of volumetric mass transfer coefficient (kla)—review of classical approaches and contribution of a novel methodology. Biochemical Engineering Journal 155:107458 (2020). https://doi.org/10.1016/j.bej.2019.107458
4. Franceschini G, Macchietto S. Model-based design of experiments for parameter precision: state of the art. Chemical Engineering Science 63:4846-4872 (2008). https://doi.org/10.1016/j.ces.2007.11.034
5. Van't Riet, K. Review of Measuring Methods and Results in Nonviscous Gas-Liquid Mass Transfer in Stirred Vessels Introduction Stirred vessels are frequently employed to achieve a. *Ind. Eng. Chem. Process Des. Dev* 18, 357 (1979).
6. Heinrich M, Arutjunjan R, Timmer J. On the different flavours of practical identifiability. Current Opinion in Systems Biology 42:100556 (2025). https://doi.org/10.1016/j.coisb.2025.100556

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Decarbonizing API Manufacturing: Conceptual Design and Scale-up Analysis of Continuous-Flow Electrosynthesis for Ibuprofen Production

Tuse Asrav[a]*, Merlin Alvarado-Morales[a], Gürkan Sin[a]

[a] Process and Systems Engineering Center (PROSYS), Department of Chemical and Biochemical Engineering, Technical University of Denmark (DTU), 2800, Kgs. Lyngby/Denmark
* Corresponding Author: tusas@kt.dtu.dk.

ABSTRACT

The decarbonization of pharmaceutical manufacturing is critical for achieving the industry's net-zero targets, and electrochemistry is emerging as a promising green technology that could play a key role in this transition. This work evaluates a continuous-flow electrochemical route for ibuprofen synthesis through electrochemical carboxylation of 1-chloro-(4-isobutylphenyl) ethane as a low-carbon alternative that can be directly coupled with renewable electricity. Experimental studies have demonstrated the selective formation of ibuprofen using a silver cathode in the ionic liquid N-methyl-N-propylpiperidinium bis(trifluoromethanesulfonyl)imide (PP13 TFSI). While the reaction mechanism is based on laboratory-scale, batch experiments, this study develops a conceptual design and scale-up methodology for the continuous route to provide an evaluation of the industrial feasibility of this electrochemical pathway through a rigorous plant-wide simulation in AVEVA® Process Simulation. Global sensitivity analysis is employed to identify key operating variables and evaluate their impact on reactor and process performance, energy consumption, and ionic liquid recovery. These insights provide a robust foundation for informed decision-making in process intensification, demonstrating the technical viability and scalability of continuous flow electrosynthesis as a sustainable alternative to conventional API manufacturing.

Keywords: Process Design, Simulation, Renewable and Sustainable Energy, Pharmaceutical Manufacturing

INTRODUCTION

As the pharmaceutical industry is a significant contributor to global emissions, increasing the focus on decarbonizing active pharmaceutical ingredient (API) manufacturing is crucial to achieving global net-zero goals. For this purpose, renewable energy sources should be utilized, and electrosynthesis is one of the green alternatives that offers an electrified synthetic pathway to replace carbon-intensive routes.

However, the feasibility and economic viability of achieving high throughput from electrosynthesis for the pharmaceutical industry remain underexplored. While batch reactors are frequently used for exploration in research due to their simple setup, flexibility, and predictability, they have significant limitations for industrial production. Flow reactors provide improved mass transfer, enhanced safety, and greater control over reaction parameters, ensuring the reproducibility required for pharmaceutical standards. Additionally, they are easier to scale up by using modular stacked-cell configurations. Considering the advantages, the transition from batch to continuous flow reactors is needed to achieve a successful scale-up, and the reliable implementation of this transition depends on the conscious design of the electrochemical flow reactor [1].

This work proposes a conceptual design for a greener, continuous electrochemical alternative to the established Hoechst process [2] for ibuprofen synthesis, which was evaluated in our previous work [3]. Mena et al. [4] demonstrated that using a silver cathode and ionic liquid (IL) allow for selective ibuprofen synthesis from electrochemical carboxylation of 1-chloro-(4-isobutylphenyl) ethane in a divided cell by suppressing side reactions without sacrificial metals. It is demonstrated that using dried ionic liquid, N-Methyl-N-propylpiperidinium

bis(tri-fluoromethanesulfonyl) imide (PP13 TFSI), overcomes significant drawbacks found using other electrolytes, such as imidazolium-based ionic liquids or organic solvents like dimethylformamide (DMF). The main advantage of using PP13 TFSI is the decreased electrochemical reduction potential which improves the energy efficiency of the process [4]. Even though the ionic liquids are more expensive than the conventional electrolytes, their recovery and recyclability capability, along with other advantages, show the potential of industrial scale-up using flow chemistry.

This study presents a rigorous plant-wide simulation of this greener, electrochemical route using AVEVA® Process Simulation. By translating laboratory cyclic voltammetry (CV) data into a modular, scalable filter-press reactor model and expanding the experimental separation logic into a full-scale downstream separation and purification stages, this work provides a methodology for early-stage conceptual design and scale-up analysis.

A global sensitivity analysis is performed to assess the impact of key operating parameters on the performance of the process. This analysis identifies the feasible operating ranges and quantifies the trade-offs between space-time yield, overall process yield, specific energy consumption, and electrolyte recovery efficiency. By identifying the operational limits and quantifying key performance indicators (KPIs), this study provides a comprehensive evaluation of the scalability of the process, offering a robust foundation for transitioning from laboratory-scale electrosynthesis to industrial production.

METHODOLOGY

The process flowsheet, illustrated in Figure 1, is developed in AVEVA® Process Simulation using mass and energy balances to model and simulate the steady-state performance of the individual unit operations.

Reactor Modelling & Scaling Up

To achieve industrial-scale throughput, we follow the technical specifications of the commercial filter-press reactor, Electro Prod Cell, marketed by Electrocell AB, which supports up to 16 m^2 of projected electrode area per module by using multiple electrodes, with a single electrode area of 0.4 m^2 [5]. To achieve higher yields, we assume that we have three reactors in series reaching a total projected electrode area of 48 m^2, each including modular 40 stacked cells having the highest possible area of Electro Prod Cell.

Each reactor is modeled as a steady-state Plug Flow Reactor unit in AVEVA® with uniform mass transfer characteristics. The reactor is modeled using a divided cell configuration, where the cathodic and anodic compartments are separated by a membrane to prevent product crossover and undesired secondary reactions. While the physical system consists of both compartments, we specifically model the cathodic compartment and its associated electrochemical kinetics in this study, as the electrochemical carboxylation of the organic chloride occurs at the silver cathode. The total active electrode area (A_e) and the total catholyte channel volume (V) are calculated as:

$$A_e = LW\sigma_{active}N_{cells} \quad (1)$$

$$V = LWd_sN_{cells} \quad (2)$$

where L and W represent the length and width of each cathode electrode, both set to 0.633 m to support for the maximum allowable area based on Electro Prod Cell specifications, and d_s represents the cathode-to-membrane spacing, which is set to minimum allowable distance 0.001 m, σ_{active} accounts for the effective surface area in contact with the electrolyte, set to 1.

In alignment with the experiments conducted by Mena et al. [4], it is assumed that carbon dioxide (CO_2) is pre-dissolved into the electrolyte until reached saturation before entering the reactor. This assumption allows for the treatment of electrolysis as a single-phase liquid process, simplifying the process by eliminating the complexities associated with gas bubble formation within the narrow electrolyte channels.

To optimize the industrial operation, the cells in one reactor module are configured in parallel for fluid flow and in series for electrical connectivity. The parallel fluid arrangement ensures that the total pressure drop across the stack is minimized while keeping uniform residence time and reactant distribution across all 40 cells per module. On the other hand, electrical connectivity is configured in series to ensure a uniform current across all cells, allowing the stack to operate at an optimized voltage under galvanostatic control.

Electrochemical Kinetics & Parameter Estimation

The electrochemical carboxylation follows an electrochemical-chemical-electrochemical (ECE) mechanism as summarized in Figure 2.

$$Step\ 1: \quad R-Cl + e^- \rightarrow [R-Cl]^{\bullet -}$$

$$Step\ 2: \quad [R-Cl]^{\bullet -} \rightarrow R^{\bullet} + Cl^-$$

$$Side\ Reaction: \quad R^{\bullet} + H \rightarrow R-H$$

$$Step\ 4: \quad R^{\bullet} + CO_2 + e^- \rightarrow R-COO^-$$

$$Step\ 5: \quad R-COO^- + HCl \rightarrow R-COOH + Cl^-$$

Figure 2: Reaction Mechanism for Electrochemical Carboxylation of Organic Chlorides

The initial electron reduction produces a radical

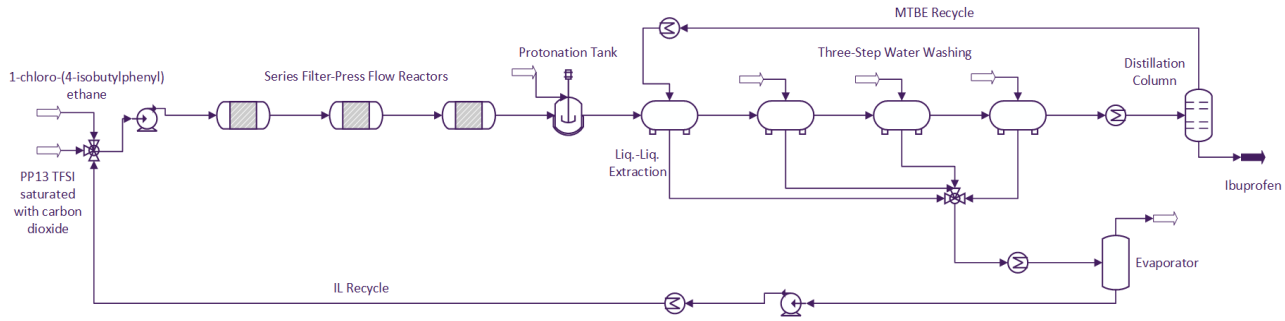

Figure 1. Ibuprofen Synthesis via Electrochemical Carboxylation Process Flowsheet.

anion, which undergoes rapid, irreversible C-Cl bond cleavage. The resulting intermediate can either form ibuprofen anion or a side-product, isobutyl ethyl benzene (IBEB). Following the second electron transfer, the resulting carboxylate anion is protonated to yield the neutral ibuprofen molecule.

The half-cell kinetic parameters are estimated using the established engineering correlations, based on data provided from Mena et al. [4]. The rate equation is derived from Faraday's Law and adapted for plug flow behavior [6].

$$\frac{dC_j}{d\tau} = \upsilon_j \eta \frac{C_j k_{eff}}{1+\frac{k_{eff}}{k_m}} \frac{A_e}{V} \tag{3}$$

Where C_j is the concentration of the compound j, τ is the residence time in the reactor, υ_j is the stoichiometric coefficient for the compound j, η is the Faradaic efficiency, k_{eff} is the effective electrochemical rate constant and k_m is the mass transfer coefficient.

The formation rates for ibuprofen and IBEB are determined by multiplying the absolute consumption rate by a selectivity ratio (S_R) of 0.9 and $(1 - S_R)$, respectively. This ensures a closed mass balance where 90% of the converted reactant is directed toward the target API, ibuprofen.

Because the chemical step is extremely fast and lacks a return peak in cyclic voltammetry, the reaction can be modeled using a single effective rate constant. This simplifies the multi-step sequence into a single rate-determining approximation suitable for plant-wide simulation in AVEVA®. The effective kinetic rate constant (k_0) is estimated using the Tafel approximation of the Butler-Volmer equation.

$$k_{eff} = k_0 \exp\left(\frac{-\alpha_c nF(E_{op,c} - E_{eq,c})}{RT}\right) \tag{4}$$

Where α_c is the cathodic charge transfer coefficient, n is the number of electrons required, F is the Faraday's constant, E_{op} is the cathodic operating potential, E_{eq} is the standard equilibrium potential, R is the universal gas constant, and T is the reaction temperature, which is 25 ºC.

To estimate the standard cathodic rate constant for irreversible reactions, the Klingler-Kochi equation [7] is used, and the charge transfer coefficient (α_c) and cathodic peak potential $(E_{p,c})$ values are taken from the cyclic voltammetry (CV) results provided by Mena et al. [4] for the silver cathode and PP13 TFSI ionic liquid as 0.37, and -1.66 V, respectively.

$$k_{0,c} = \sqrt{\frac{D\alpha_c nF\upsilon_{scan}}{RT}} exp\left(\frac{\alpha_c nF(E_{p,c} - E_{0,c})}{RT} + 0.78\right) \tag{5}$$

Where D is the diffusion coefficient and υ_{scan} is the scan rate used for the CV experiment, which is 0.5 V/s.

The mass transfer coefficient is determined by the fluid dynamics inside the filter-press cell, and we assume that we have turbulent promoters to enhance the mass transfer.

$$k_m = \frac{ShD}{d_H} \tag{6}$$

Where the Sherwood number (Sh) is a function of the Reynolds number (Re) and Schmidt number (Sc).

We assume that the process operates under stage-wise galvanostatic conditions, where a constant current is maintained across the 40-cell stack. To estimate the required operating potential for the nth reactor unit for galvanostatic operation, the following equation is used.

$$E_{op,c,n} = E_{eq,c} - \frac{RT}{\alpha_c nF} \ln\left(\frac{i_n k_{m,c}}{k_{0,c}(i_{lim,n} - i_n)}\right) \tag{7}$$

Where i_n is the applied current density, calculated as a fixed ratio of the limiting current density and the limiting current density $(i_{lim,n})$ is calculated by $nFC_b k_{m,c}$ for each reactor unit due to mass-transport limitations.

The total specific cathodic energy consumption per unit mass of product is determined by summing the power requirements of the three individual reactor modules. Each module's voltage $(V_{stack.n})$ is calculated by adding the activation overpotentials and ohmic losses $(V_{ohmic,n})$ of its 40 stacked cells.

$$Specific\ Energy\ Consumption = \frac{\sum_{n=1}^{3}(V_{stack,n} I_n)}{\dot{m}_{ibuprofen}} \tag{8}$$

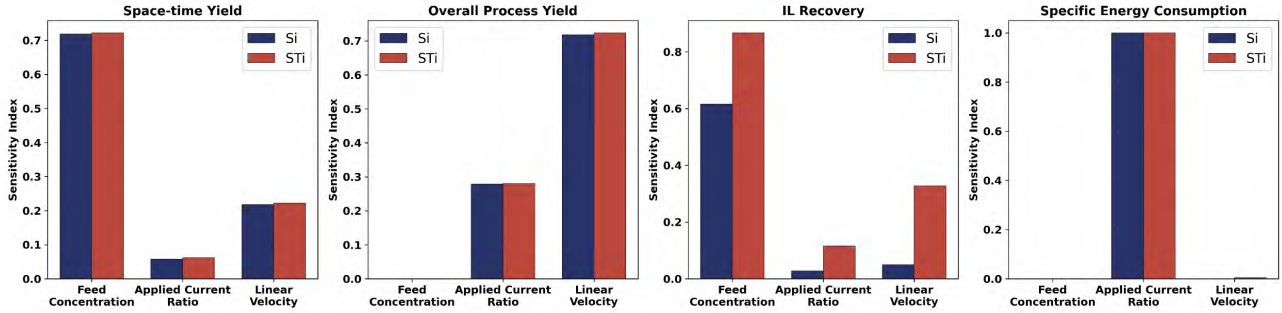

Figure 3. Global Sensitivity Analysis Results.

$$V_{stack,n} = 40\left(|E_{op,c,n}| + V_{ohmic,n}\right) \quad (9)$$

$$V_{ohmic,n} = i_n\left(\frac{d_s}{\kappa}\right) \quad (10)$$

Where κ is the electrolyte conductivity.

Downstream Processes

To model phase equilibria, the NRTL (Non-Random Two-Liquid) property method is utilized for components with established experimental binary parameters. To accurately capture the non-ideal behavior of the ionic liquid and its interactions with organic solutes, the Modified UNIFAC (Dortmund) Group Contribution Method is used to estimate the required thermodynamic parameters. For vapor-liquid equilibrium calculations, the vapor phase is modeled using the Ideal Gas Law, as the system operates at pressures where gas-phase non-ideality is negligible. These enable a rigorous calculation of activity coefficients and phase equilibria during both the reaction and downstream stages.

The reactor effluent enters a Liquid-Liquid Extraction (LLE) unit, where methyl tert-butyl ether (MTBE) is used to separate the ibuprofen from the ionic liquid. In this model, the recovery efficiency for ibuprofen is set at a fixed design specification of 90% for the LLE unit and 92% for the subsequent washing stages to simulate a standardized quality for downstream processes. Under these constraints, initial separation yields a recovery of approximately 30% of the ionic liquid. To enhance this, the ibuprofen-rich stream undergoes a 3-step deionized water wash. Finally, the multi-stage separation achieves a total recovery of the ionic liquid of approximately 70%.

This downstream configuration scales up the laboratory logic established by Mena et al. [4], where ether extraction and a single water wash achieved approximately 80% ionic liquid recovery in a laboratory scale, batch setting.

To ensure the ionic liquid is suitable for recycling and reuse as the electrolyte, an evaporator is used to dry the recovered ionic liquid by removing residual water and ether. Simultaneously, the ibuprofen-rich stream is sent to a distillation column where MTBE is recovered for recycling, and the final ibuprofen product is purified to pharmaceutical standards.

SENSITIVITY ANALYSIS RESULTS AND DISCUSSION

The global sensitivity analysis is conducted using the Sobol variance-based method to estimate first-order sensitivity indices (S_i), which quantify the individual contribution of each parameter to the output variance, excluding interaction effects, and total-order sensitivity indices (S_{Ti}), which quantify the total contribution of each parameter, including all interactions with other parameters [8]. The analysis is performed using a quasi-Monte Carlo (QMC) Sobol sequence sampling [9] and the Saltelli design [10] with a base sample size of 2048, resulting in 10240 simulation runs to construct the required matrices for the three input variables. The indices are then calculated using the Jansen estimator [11].

For this study, the reactor dimensions are kept constant based on the industrial specifications of the Electro Prod Cell. The sensitivity analysis focuses on operational variables, specifically the feed concentration, applied current ratio, and electrolyte linear velocity, to evaluate the process performance.

The feed concentration consisting of 1-chloro-(4-isobutylphenyl) ethane dissolved in CO_2-saturated ionic liquid (PP13 TFSI) is assumed to be varied between 0.2 and 0.3 kmol/m^3. This range is chosen to balance the need for high industrial throughput with the solubility limits and viscosity constraints of the ionic liquid. To evaluate the reactor's performance near its physical limits, the applied current density is defined as a fraction of the limiting current density. The ratio is assumed to be varied between 0.7 and 0.9. This range is selected to investigate the mass-transfer controlled regime, which is of particular interest for industrial scale-up. The linear velocity is assumed to be varied between 0.02 and 0.03 m/s, comparably lower since we use ionic liquids that are highly viscous.

To provide a comprehensive evaluation, four key

performance indicators (KPIs) are selected: space-time yield (STY), which provides a measurement for reactor productivity; overall process yield, which assesses both the upstream and downstream efficiency; specific cathodic energy consumption, which quantifies the electrical energy requirements; and ionic liquid recovery to evaluate the economic viability and sustainability. The results from the global sensitivity analysis are presented in Figure 3.

Results reveal that the space-time yield is dominated by the feed concentration, while the feed concentration has a negligible impact on the overall process yield. This behavior is caused by the first-order kinetics of electrochemical synthesis. The rate of reaction scales linearly with bulk concentration; therefore, increasing feed concentration directly increases the amount of product produced per volume per unit time. However, because conversion and yield are fractional metrics, they remain unchanged. Additionally, because the downstream recovery ratios for ibuprofen are fixed, the overall process yield also remains unchanged. This suggests that the process can achieve higher throughputs by increasing concentration without sacrificing the chemical efficiency of the reactor, considering the solubility limits of the ionic liquid.

While both electrolyte linear velocity and applied current ratio theoretically influence the reaction rate, the Sobol analysis unveils a higher impact of the linear velocity on the overall process yield compared to the applied current ratio. Lowering the velocity increases the residence time, which proves to be the dominant factor in ensuring sufficient duration for the electrochemical conversion.

Even though there is a trade-off since lowering the velocity also decreases the space-time yield, the results guides for a strategic pathway for optimization. Because feed concentration has a more dominant effect on space-time yield than velocity does, by increasing feed concentration, the plant can offset the productivity losses occurred from lowering the velocity. This strategy can allow us to achieve high chemical efficiency and high industrial throughput, reaching the target production without the loss of performance associated with increased residence time.

In contrast to the other input variables, the ionic liquid recovery showed high parameter interactions, with higher total-order indices compared to first-order indices. As the feed concentration increases, the mass of ibuprofen in the effluent also increases. To maintain the fixed recovery specifications, the process scales the consumption of MTBE and water. As more solvent flows through the system, it physically carries away more of the ionic liquid, either as tiny droplets or by slightly dissolving it. This creates a direct trade-off: as production capacity increases, the efficiency of the ionic liquid recycling decreases.

Finally, the Sobol analysis reveals that the specific energy consumption is exclusively governed by the applied current ratio, with a total-order index of approximately 1.0. While the feed concentration and electrolyte linear velocity both fundamentally influence the limiting current density, their impact on the energy intensity is neutralized. Consequently, the only variable that can change the cost per kilogram is the applied current ratio.

As the applied current ratio increases, the cell operates closer to its mass-transfer limit, in this region the specific energy consumption behaves asymptotically, as the cell voltage increases significantly without a matching proportional increase in mass yield. This result proves that the energy cost is decoupled from production throughput; we can increase the plant's capacity by raising the feed concentration without observing an increase in the electricity cost per unit of product.

CONCLUSION AND FUTURE WORK

This work presents a rigorous modeling and simulation framework for the design and scale-up of continuous electrochemical flow reactors, applied to an electrosynthesis route for ibuprofen production. A modular filter-press reactor is selected to reflect industrial electrochemical practice and is modeled using a plug flow reactor, as the electrolyte flows through narrow channels with plug-flow-like behavior and uniform mass transfer characteristics along the electrode surface. For the smooth transition from laboratory-scale data to industrially relevant design parameters, we use governing equations to define electrochemical kinetics, mass-transfer limitations, hydrodynamics, and reactor geometry. The simulation of the process flowsheet demonstrates the technical feasibility and scalability of continuous electrochemical carboxylation for ibuprofen synthesis.

Global sensitivity analysis is used as a guidance tool to explore the simulated process and identify the operating variables that most strongly govern productivity, efficiency, energy intensity, and electrolyte recovery. The results reveal clear trade-offs between throughput, residence time, energy efficiency, and ionic liquid losses that are inherent to the scale-up of electrochemical flow reactors. Importantly, the analysis shows that production capacity can be increased through feed concentration without penalizing specific energy consumption, while electrolyte recovery emerges as a critical bottleneck at higher throughputs.

Beyond demonstrating process feasibility, this study highlights that rigorous reactor and downstream processes modeling combined with the global sensitivity analysis is a necessary early-stage step for scale-up and decision making toward economically viable electrochemical API manufacturing. By identifying the key

operating variables with dominant effects and feasible operating windows, the proposed framework provides a robust foundation for reliable techno-economic and sustainability analysis.

Future work will extend the current framework to include full-cell kinetics, incorporating anodic reactions, additional voltage losses through membrane resistance, and anodic potential. Based on the results from the global sensitivity analysis, the optimization will be performed, and techno-economic and life cycle analysis will be conducted comparing the green pathway with the current manufacturing pathway for the ibuprofen synthesis.

ACKNOWLEDGEMENTS

This work is part of the European Marie Skłodowska-Curie network MiEl. The MiEl project received funding by the European Union under the Grant Agreement no. 101073003. Views and opinions expressed are however, those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

AUTHOR IDENTIFIERS

Author ORCIDs:
Asrav T: 0009-0001-8807-6450
Alvarado-Morales M: 0000-0003-3689-4563
Sin G: 0000-0003-0513-4502

REFERENCES

1. Regnier M, Vega C, Ioannou DI, Noël T. Enhancing electrochemical reactions in organic synthesis: the impact of flow chemistry. Chem. Soc. Rev. 53:10741-10760 (2024). https://doi.org/10.1039/d4cs00539b
2. Elango V, Murphy MA, Smith BL, Davenport GK, Mott GN, Zey EG, Zey GL. Method for producing Ibuprofen. US Patent US 4981995A. (1990)
3. Asrav T, Alvarado-Morales M, Sin G. Simulation-based Techno-economic Optimization of Continuous Ibuprofen Manufacturing under Uncertainty. *ChemRxiv*. [Preprint]. (2025) https://doi.org/10.26434/chemrxiv-2025-j23jg-v2
4. Mena S, Sanchez J, Guirado G. Electrocarboxylation of 1-chloro-(4-isobutylphenyl)ethane with a silver cathode in ionic liquids: an environmentally benign and efficient way to synthesize ibuprofen. RSC Adv. 9:15115-15123 (2019). https://doi.org/10.1039/c9ra01781j
5. ElectroCell A/S. Electro Prod Cell. ElectroCell. https://www.electrocell.com/products/electrochemical-flow-cells/electro-prod-cell
6. Scott K. Reactor modelling for electrochemical processes. J of Chemical Tech & Biotech 54:257-266 (2007). https://doi.org/10.1002/jctb.280540308
7. Klingler RJ, Kochi JK. Electron-transfer kinetics from cyclic voltammetry. quantitative description of electrochemical reversibility. J. Phys. Chem. 85:1731-1741 (2002). https://doi.org/10.1021/j150612a028
8. Sobol IM, Kucherenko SS. Global sensitivity indices for nonlinear mathematical models. Review. *Wilmott*, 56–61 (2005) https://doi.org/10.1002/wilm42820050114.
9. Sobol' IM. On the distribution of points in a cube and the approximate evaluation of integrals. USSR Computational Mathematics and Mathematical Physics 7:86-112 (1967). https://doi.org/10.1016/0041-5553(67)90144-9
10. Saltelli A, Annoni P, Azzini I, Campolongo F, Ratto M, Tarantola S. Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index. Computer Physics Communications 181:259-270 (2010). https://doi.org/10.1016/j.cpc.2009.09.018
11. Jansen MJW. Analysis of variance designs for model output. Computer Physics Communications 117:35-43 (1999). https://doi.org/10.1016/s0010-4655(98)00154-4

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Multi-Level Hybrid EKF–Machine Learning Soft Sensor for Robust Bioprocess Monitoring

Mohammad Reza Boskabadi[a], Rajiv Kailasanathan[ab], Luis Ricardez-Sandoval[c] and Seyed Soheil Mansouri[a*]

[a] Department of Chemical and Biochemical Engineering, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark
[b] Novo Nordisk Foundation Centre for Biosustainability, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark
[c] Department of Chemical Engineering, University of Waterloo, Waterloo, N2L 3G1, ON, Canada
* Corresponding Author: seso@kt.dtu.dk

ABSTRACT

Real-time monitoring of bioprocesses is hindered by sparse, heterogeneous measurements of key biological states, such as biomass, substrate, and product concentrations. Extended Kalman Filter (EKF)–based soft sensors offer a physics-grounded solution but are sensitive to limited observability, sensor bias, and process-model mismatch—conditions common in industrial fermentations. This work proposes a Levelized Hybrid Estimation Architecture (LHEA) that systematically enhances physics-based state estimation through increasing robustness and adaptivity while preserving model transparency and regulatory interpretability. The approach is evaluated using the KTB1 benchmark simulation model for continuous lovastatin production under an industrially realistic, cost-constrained sensor configuration combining dissolved oxygen, biomass proxy, volume measurements, and sparse HPLC product assays. Three estimator levels are investigated: (L1) a baseline EKF, (L2) a bias-augmented EKF for sensor-drift robustness, and (L3) a hybrid EKF with physics-constrained residual learning driven by sparse assays. Results show that L1 provides stable state reconstruction under nominal conditions but is sensitive to bias and model mismatch. L2 effectively isolates measurement bias, while L3 adapts to evolving process kinetics and achieves the lowest estimation errors under mismatch. These findings demonstrate that a structured, levelized integration of machine learning can significantly enhance soft-sensor reliability without sacrificing interpretability, providing a practical pathway toward robust digital twins.

Keywords: Biosystems, Process Monitoring, Hybrid Modelling, Soft Sensor

INTRODUCTION

Real-time monitoring and optimization of biomanufacturing processes remains an open challenge due to the limited availability of reliable online sensors for critical variables such as biomass, substrates, and product concentrations [1]. While measurements such as dissolved oxygen, pH, and flow rates are typically available at high frequency, key quality attributes and performance indicators are often accessible only through infrequent and costly offline assays [2]. In continuous biomanufacturing, long operating horizons, tight quality constraints, and economic pressures further exacerbate these challenges. As a result, soft sensors must operate under sparse, asynchronous, and noisy measurements, where sensor drift, calibration errors, and process disturbances can significantly degrade estimation performance.

Model-based state estimation methods, particularly Extended Kalman Filters (EKFs), have been widely applied to infer unmeasured states in bioprocesses using first-principles dynamic models [3], [4]. When accurate models and informative measurements are available, EKFs can provide effective real-time monitoring of biological systems. However, in realistic biomanufacturing environments, model mismatch arising from uncertain kinetics, unmodelled physiological shifts, and slowly varying parameters is unavoidable. Combined with sensor drift and measurement bias, these effects can lead to filter inconsistency, biased estimates, or even divergence in the estimation [5], [6]. Addressing robustness to both model mismatch and measurement imperfections, therefore, remains a central challenge in the development of

reliable soft sensors for bioprocess monitoring.

Machine-learning (ML)–based soft sensors have attracted increasing attention due to their ability to capture complex nonlinear relationships directly from process data [7]. Despite their predictive power, purely data-driven approaches raise concerns related to interpretability, extrapolation beyond training data, sensitivity to operating conditions, and regulatory acceptance [8]. In regulated biomanufacturing contexts, particularly biopharmaceuticals, the lack of explicit physical structure makes it difficult to guarantee mass-balance consistency, diagnose estimation failures, or ensure reliable performance during process regime changes. These limitations motivate hybrid approaches that combine the transparency and structure of first-principles models with the adaptability of data-driven learning [9].

This study presents a *levelized hybrid Extended Kalman Filter–Machine Learning (EKF–ML) approach* for robust soft sensing in bioprocesses, where robustness refers to the preservation of estimator stability and acceptable estimation accuracy in the presence of sensor drift, measurement sparsity, and structural mismatch between assumed and true process kinetics. The methodology is evaluated using the KTB1 MATLAB/Simulink benchmark, a realistic simulator for continuous biomanufacturing that couples an upstream CSTR bioreactor with downstream operations and incorporates sensor noise, delays, and regulatory control [10]. The main contributions of this work are fourfold:

(i) the development of a *robust soft-sensing model* that enhances state estimation accuracy and reliability under realistic measurement limitations;

(ii) the introduction of a *levelized estimation structure*, enabling systematic investigation of increasing model–data hybridization and facilitating future methodological extensions;

(iii) a *comprehensive case study* demonstrating the practical value of the approach under nominal operation, sensor bias, and process-regime changes; and

(iv) the release of an *open-source implementation* to support transparency, reproducibility, and further research in hybrid soft-sensor design (https://github.com/Boskabadi/LHEA).

METHODOLOGY: LEVELIZED HYBRID ESTIMATION ARCHITECTURE (LHEA)

We consider a nonlinear bioprocess described by a first-principles mass-balance model:

$$\dot{x}(t) = f(x(t), u(t)), \tag{1}$$

where $x(t) \in \mathbb{R}_{\geq 0}^{n_x}$ denotes the process states (e.g., volume, biomass, substrates, product, dissolved oxygen), $u(t) \in \mathbb{R}_{\geq 0}^{n_x}$ represents known non-negative manipulated inputs (e.g., feed flow rates and compositions), and $t \in \mathbb{R}_{\geq 0}$is time. The nonlinear function $f(\cdot)$represents the reactor mass-balance equations derived from conservation laws and kinetic expressions, and is assumed to be continuously differentiable with respect to $x$within the physically admissible operating region. Only a subset of states is measured online at heterogeneous sampling rates, while key quality attributes (e.g., product concentration) are available through sparse laboratory assays, e.g. at the end of the batch. State estimation is performed using an Extended Kalman Filter (EKF) as a baseline estimator. The novelty of this work lies not in the EKF formulation itself, but in a systematic, levelized extension of the estimator to address practical challenges in bioprocess monitoring, i.e. weak observability, sensor drift, and model mismatch.

Level 1 (L1): Physics-based soft sensor

Level 1 implements a standard physics-based EKF using the mechanistic model and available multi-rate measurements. The process dynamics are given by

$$\dot{x}(t) = f(x(t), u(t)), \tag{2}$$

$$y_k = h_k(x) + v_k, \tag{3}$$

where $h_k(\cdot)$represents the subset of sensors available at the time t_k, and v_k is measurement noise. In the multi-rate implementation, each measurement channel is incorporated only when its corresponding sampling time is reached, so the EKF evolves continuously by model prediction between measurement arrivals and applies discrete corrections at the sensor and assay instants. For numerical robustness, the covariance update is implemented using the Joseph form [11], with covariance symmetrization after propagation and update. L1 serves as a baseline physics-consistent soft sensor, providing real-time state estimates under nominal conditions but remaining sensitive to sensor bias and model mismatch.

Level 2 (L2): Bias-robust estimation via state augmentation

Level 2 extends the EKF by explicitly estimating slowly varying sensor biases. This is achieved by augmenting the state vector as:

$$x_a = \begin{bmatrix} x \\ b \end{bmatrix}, \tag{4}$$

where b represents unknown measurement bias terms associated with selected sensors (e.g., volume or dissolved oxygen probes). The bias dynamics are modeled as a random walk,

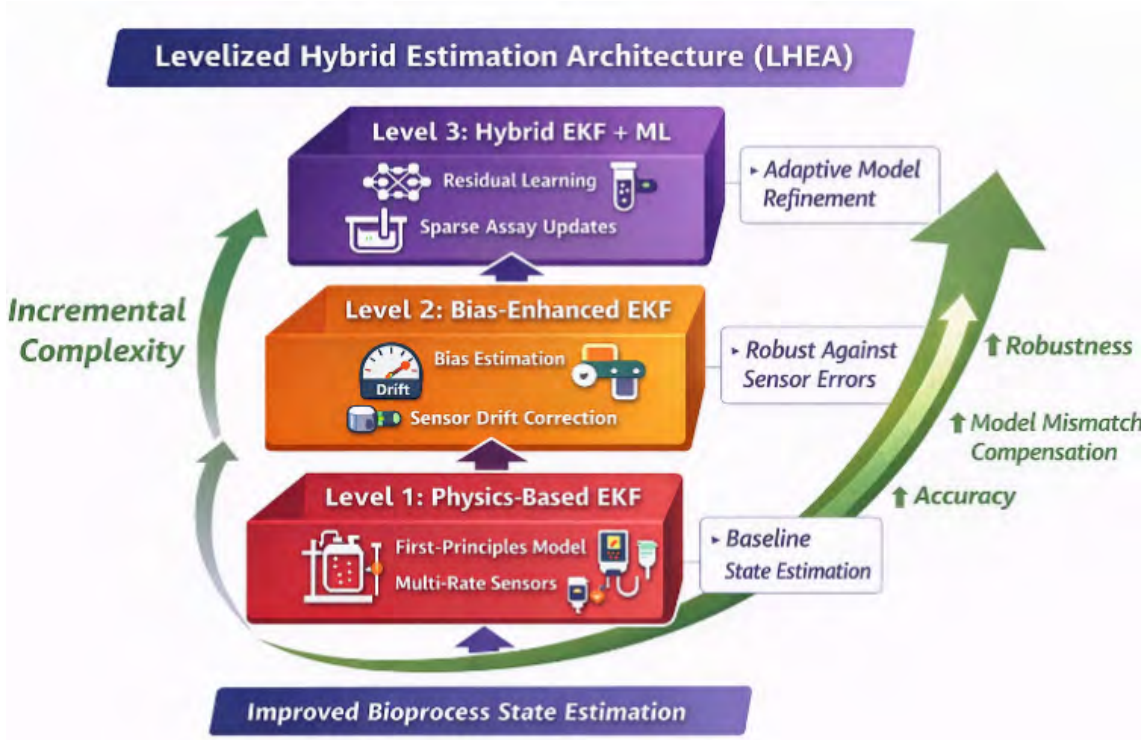

Figure 1. Conceptual illustration of the proposed Levelized Hybrid Estimation Architecture (LHEA). The framework progresses from a physics-based EKF (L1, red), to a bias-robust augmented EKF (L2, orange), and finally to a hybrid EKF–ML estimator (L3, purple). Sparse and multi-rate measurements are incorporated at all levels.

$$\dot{b}(t)=wb(t), \tag{5}$$

reflecting slow drift due to fouling, calibration loss, or aging. Measurements affected by bias are written as

$$y_k = h(x_k) + b + v_k. \tag{6}$$

Importantly, neither the presence, magnitude, nor onset of bias is assumed to be known a priori. Only the possibility of bias in specific measurement channels is assumed based on engineering knowledge. The EKF jointly estimates process states and bias terms using innovation statistics, allowing sensor drift to be absorbed without corrupting the underlying state estimates. At this level, the process model remains unchanged, and no data-driven components are introduced.

Level 3 (L3): Physics-constrained hybrid EKF–ML

Even with bias handling, first-principles models inevitably exhibit a mismatch due to uncertain kinetics, physiological adaptation, or operating-regime changes. Level 3 introduces a physics-constrained residual learning mechanism within the EKF to compensate for such a mismatch while preserving interpretability. The process model is augmented as

$$\dot{x}(t) = f\big(x(t), u(t)\big) + r\theta\Big(\phi\big(x(t)\big)\Big), \tag{7}$$

where $r_\theta(\cdot)$is a learned residual correction and $\phi(x)$is a low-dimensional, interpretable feature vector (e.g., biomass and dissolved oxygen). The residual is restricted to selected dynamics (e.g., product formation) and is bounded to ensure physical plausibility, meaning that it (i) does not violate mass-balance structure, (ii) preserves state positivity, and (iii) remains bounded within a small fraction of the nominal kinetic rates. These constraints prevent violations of positivity, mass conservation, or known stoichiometric relationships.

The residual parameters θ are updated online only when sparse laboratory assays become available (e.g., HPLC product measurements every 24 h), using a lightweight adaptive law(e.g., normalized gradient or RLS-type update), triggered only when sparse, high-fidelity laboratory assays become available. Learning therefore, occurs within the same run, without offline retraining. The feasibility of this adaptation is ensured by (i) fast online measurements that constrain the state trajectory, (ii) sparse but informative assays that anchor slow dynamics, and (iii) the low dimensionality of the residual parameterization. The preceding observability analysis confirms that inclusion of sparse product measurements renders the system observable over the estimation horizon,

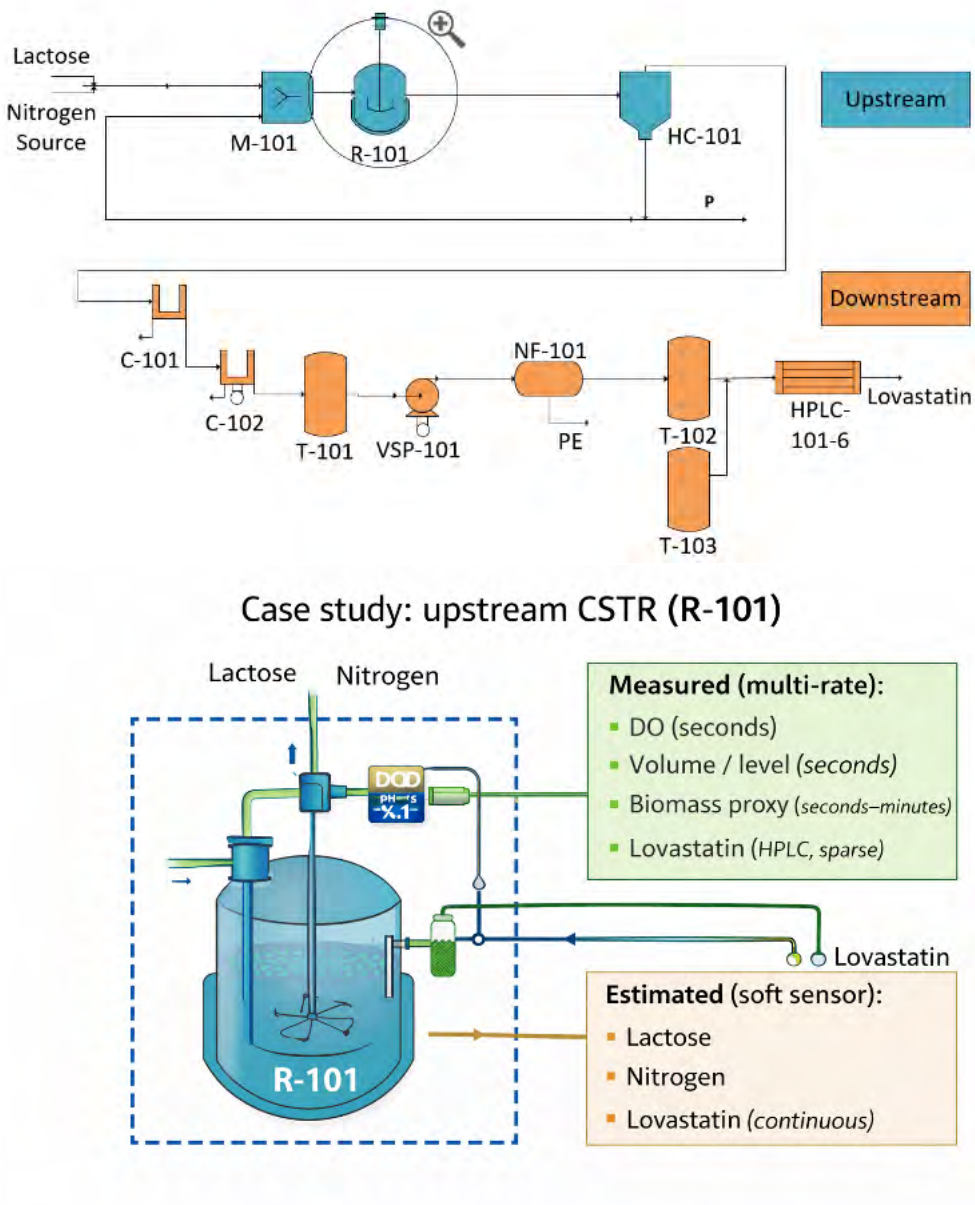

Figure 2. (a) Overall process flow diagram (PFD) of the KTB1 continuous biomanufacturing benchmark, including upstream fermentation and downstream separation units for lovastatin production. (b) Zoomed-in schematic of the CSTR reactor (R-101) highlighting available online sensors, sparse laboratory assays, and the hybrid EKF-based soft-sensor role for real-time state estimation.

enabling meaningful online learning.

In addition, slowly varying kinetic effects can be represented through a scalar scaling parameter:

$$q_{\text{eff}}(t) = q_{\text{nom}}\, q_{\text{scale}}(t), \tag{8}$$

which is updated using assay information and constrained within physically meaningful bounds.

Figure 1 summarizes the proposed *Levelized Hybrid Estimation Architecture (LHEA)* and illustrates how estimation capability is progressively enhanced through structured hybridization. Starting from a purely physics-based estimator (L1), each subsequent level introduces additional mechanisms to address practical limitations encountered in bioprocess monitoring, including sensor bias, model mismatch, and sparse measurements. The color coding highlights the incremental nature of the

framework: L1 (red) provides baseline real-time state reconstruction using first-principles models; L2 (orange) augments the estimator with bias states to improve robustness against sensor drift without altering the underlying process model; and L3 (purple) incorporates physics-constrained, sparsely trained machine-learning residuals to compensate for unmodelled kinetics or slowly varying parameters. Multi-rate and sparse measurements (green) are consistently integrated across all levels, reflecting realistic PAT availability. As illustrated by the upward trend in the figure, each level systematically increases estimation robustness, tolerance to model mismatch, and predictive accuracy while preserving physical interpretability.

CASE STUDY: CONTINUOUS LOVASTATIN PRODUCTION IN THE KTB1 SIMULATOR

Process Description

This study uses the KTB1 simulation platform, implemented in MATLAB/Simulink, as a representative case study for continuous biomanufacturing (Figure 2 a). KTB1 models the continuous production of the Activate Pharmaceutical Ingredient (API) lovastatin and includes realistic upstream and downstream unit operations. The upstream section consists of a mixer, a continuous stirred-tank bioreactor (CSTR), pumps, and a hydrocyclone, while the downstream section includes centrifuges, buffer tanks, a variable-speed pump, nanofiltration, and HPLC columns for product isolation. The simulator is able to account for sensor noise, actuator delays, and regulatory control, making it a suitable benchmark for evaluating advanced monitoring and soft-sensing strategies under industrially relevant conditions [10].

The focus of this work is the upstream CSTR reactor (R-101), where lovastatin is produced by *Aspergillus terreus*. Continuous monitoring of key states inside the reactor—particularly substrate and product concentrations—is essential for maintaining process stability, ensuring consistent product quality, optimizing yield, and enabling early fault detection. However, continuous monitoring in bioprocesses remains challenging, and this difficulty is especially pronounced for secondary metabolites such as lovastatin [12].

Continuous monitoring of lovastatin inside the reactor is challenging due to several practical limitations. Online sensors are affected by noise, drift, and delays, which can significantly degrade estimation performance during long continuous runs. Moreover, lovastatin is produced at low concentrations in dilute fermentation broths, resulting in weak analytical signals. Direct in-line quantification remains limited, and product measurements typically rely on sparse at-line or offline HPLC assays. In addition, slow biological kinetics and strong upstream–downstream coupling in continuous operation amplify the impact of measurement delays, making adequate feedback control difficult.

The dynamic CSTR model in KTB1 explicitly represents six key reactor states: biomass concentration, nitrogen (adenine), lactose (substrate), lovastatin (product), dissolved oxygen, and liquid volume/level. Table 1 summarizes these states, their symbols and units, typical measurement availability, preferred sensing technologies, and representative sampling cadences reported in the literature. While dissolved oxygen and volume/level are routinely available at fast sampling rates, substrate and product concentrations are typically measured only sparsely or inferred via soft sensors.

Table 1. Reactor states considered in the CSTR model, including symbols, units, measurability, typical sensing technologies, and realistic sampling cadences.

State	Symbol & unit	Preferred sensors & cadence	Key refs
Biomass	C_X (g/L)	Capacitance (5–15 s); OD/Turbidity (1–5 s)	[13]
Nitrogen	C_N (g/L)	HPLC-UV 260 nm (10–30 min); UV/Vis or Raman + PLS (30–120 s)	[14]
Lactose	C_LAC (g/L)	HPLC/enzymatic (5–30 min); Raman/NIR + PLS (30–120 s)	[15]
Lovastatin	C_MEV (g/L)	HPLC (30–120 min); LC-MS for intermediates; Raman soft-sensor (1–3 min)	[16]
Dissolved oxygen	C_DO (g/L)	Optical DO probe (1–5 s); polarographic (legacy)	[17]
Liquid volume/level	V (L)	Level sensor 1–5 s + geometry	[17]

Observability Analysis and Sensor Scenarios

To systematically assess how realistic sensor configurations affect state reconstruction, a multi-rate observability analysis was performed on the CSTR model. Nine sensor scenarios were defined, reflecting progressively richer combinations of online, at-line, and sparse measurements commonly encountered in industrial practice. These scenarios are summarized in Table 2.

Table 2. Sensor configuration scenarios used for the observability analysis, showing included measurements, and sampling cadences

Scenario	Sensors included	Sampling cadence
1	DO	5 s
2	DO, Volume/Level	DO: 5 s, Level: 5 s
3	DO, Biomass proxy (capacitance)	DO: 5 s, Biomass: 10 s
4	DO, Biomass proxy, Volume/Level	DO: 5 s, Biomass: 10 s, Level: 5 s
5	DO, Lactose (Raman)	DO: 5 s, Lactose: 60 s
6	DO, Biomass proxy, Lactose	DO: 5 s, Biomass: 10 s, Lactose: 60 s
7	DO, Biomass proxy, Lactose, Volume/Level	DO: 5 s, Biomass: 10 s, Lactose: 60 s, Level: 5 s
8	Scenario 7 + Product (MEV, HPLC)	DO: 5 s, Biomass: 10 s, Lactose: 60 s, Level: 5 s, MEV: 30 min
9	Scenario 4 + Product (MEV, HPLC)	DO: 5 s, Biomass: 10 s, Level: 5 s, MEV: 24 h

The observability rank for each scenario was evaluated over a 24-hour horizon using a base integration step of 1 s. The results are summarized graphically in Figure 3. As shown, DO-only sensing (Scenario 1) yields limited observability (rank 4), as oxygen dynamics alone are insufficient to uniquely reconstruct all material balances. Adding either volume/level measurements (Scenario 2) or a biomass proxy (Scenario 3) increases the observability rank to 5, demonstrating the importance of dilution effects and growth-related information. However, neither configuration alone achieves full state observability.

Scenario 4, which combines DO, biomass proxy, and volume/level measurements at realistic online cadences, consistently achieves an observability rank of 5 out of 6. This indicates that all states except the product concentration are practically observable, a result that aligns well with industrial experience: product formation dynamics are only weakly coupled to fast online sensors. The inclusion of substrate measurements (Scenarios 5–7) does not fundamentally improve observability unless combined with complementary sensors, highlighting that additional measurements must contribute independent dynamic information to be effective.

Full observability (rank 6) is achieved only in Scenarios 8 and 9, where sparse product (lovastatin) measurements are included, because the product state cannot be uniquely inferred from fast online measurements alone. Occasional product assays provide independent information that anchors the slow product dynamics and closes the observability gap when fused with faster measurements. Notably, Figure 3 shows that even very slow product assays—on the order of tens of minutes to several hours—are sufficient to close the observability gap when fused with faster online measurements over an extended horizon. Based on these results, Scenario 9 was selected as the primary sensor configuration for the subsequent estimation studies. This scenario reflects a cost-constrained yet industrially realistic sensing setup and serves as a stringent test case for the proposed levelized hybrid estimation architecture.

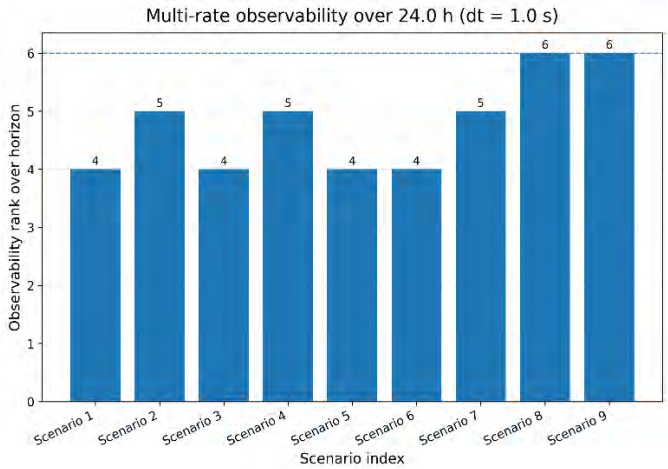

Figure 3. Multi-rate observability rank of the CSTR reactor states over a 24 h horizon for different sensor configurations (Scenarios 1–9). The dashed line indicates full observability (rank 6). Scenarios 8 and 9 have full observability.

RESULTS AND DISCUSSION

Considering the observability analysis in Figure 3, prior experience with similar industrial fermentations, and economic constraints associated with advanced sensing, *Scenario 9* was selected for subsequent estimation studies. This scenario combines volume/level measurement, a capacitance-based biomass proxy, dissolved oxygen, and sparse HPLC product assays, representing an industrially realistic and cost-constrained sensor configuration.

EXP1. Nominal estimation performance

Figure 4 compares the lovastatin concentration estimates obtained under nominal operating conditions. The normal reactor trajectory, shown as the blue line, remains nearly constant at approximately 1.2 g/L over the operating horizon. The Level 1 estimate, based solely on the first-principles model together with the available online measurements and sparse 24 h product assays, remains numerically stable and exhibits small discrete corrections at the assay instants. These update events confirm that the baseline EKF does assimilate the sparse

laboratory measurements; however, between assay times, a mild downward drift is still observed, indicating sensitivity to residual process-model mismatch. The Level 2 bias-augmented EKF estimate follows the Level 1 trajectory closely and shows the same assay-driven corrections, with only marginal deviation from Level 1 under nominal conditions. This behavior indicates that, in the absence of systematic measurement bias, the additional bias states remain weakly excited and therefore provide limited improvement in estimation accuracy. In contrast, the Level 3 hybrid estimator (dashed red), in which the physics-based EKF is augmented with residual learning driven by sparse HPLC product assays, remains tightly bound around the normal reactor trajectory, with small corrective adjustments occurring at assay times. These results demonstrate that, under nominal conditions, even when all estimators receive the sparse assays, the primary performance improvement arises from adaptive compensation of process-model mismatch rather than sensor bias handling.

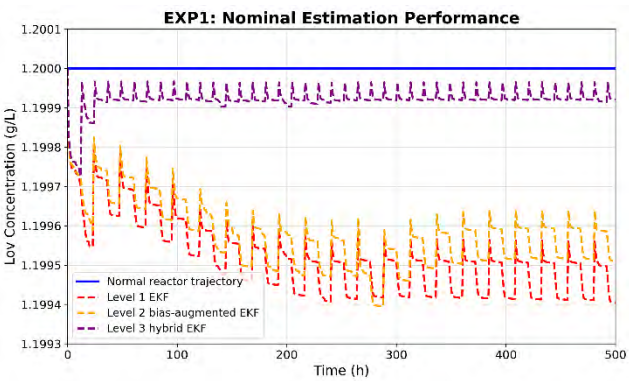

Figure 4. Lovastatin concentration estimates under nominal operation. The blue curve represents the true lovastatin concentration. Level 1 (EKF), Level 2 (bias-augmented EKF), and Level 3 hybrid are overlaid for direct comparison

EXP2. Robustness to measurement bias

Figure 5 evaluates estimator robustness to sensor bias by introducing a step change only in the measured reactor volume (5000 → 5050 L at t = 50 h), while the true process remains unchanged. The normal reactor trajectory, shown as the blue line, remains consistent with the nominal operating case. The Level 1 estimate (dashed red) directly follows the biased volume signal; once the bias is introduced, its lovastatin estimate progressively deviates upward from the normal reactor trajectory, indicating that the standard EKF propagates the biased measurement through the process model and mass balances. The Level 2 bias-augmented EKF estimate (dashed orange) effectively absorbs the artificial measurement shift through its bias states, maintaining a lovastatin estimate consistent with the nominal case. The Level 3 estimator (dashed purple) exhibits similar stability under this scenario, as no kinetic mismatch is present. These results demonstrate that the primary contribution of Level 2 is robustness to sensor bias and drift rather than improved nominal accuracy.

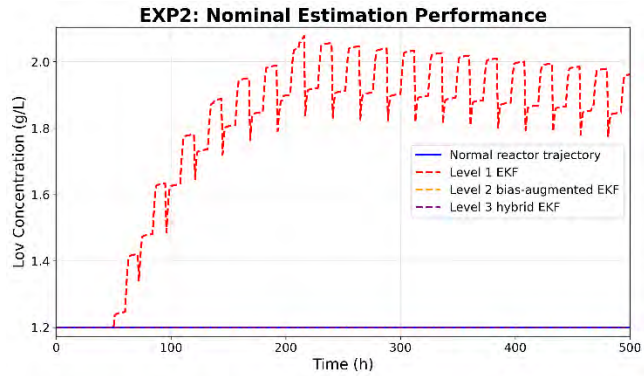

Figure 5. Lovastatin concentration estimates under a measurement bias scenario. The normal reactor trajectory (blue) is shown together with Level 1 EKF (dashed red), Level 2 bias-augmented EKF (dashed orange), and Level 3 hybrid EKF (dashed purple).

EXP3. Adaptation to process-model mismatch

Figure 6 examines estimator performance under a true process-model mismatch, introduced by increasing the intrinsic product formation rate $q_{\max,\mathrm{MEV}}$by 20% at $t = 50$h. The resulting reactor lovastatin concentration, shown as the blue line, exhibits a gradual but persistent deviation from the nominal trajectory due to the altered kinetics. Under this mismatch, the Level 1 estimate (dashed red) and the Level 2 bias-augmented EKF estimate (dashed orange) are repeatedly corrected by the sparse assay updates; however, both remain limited by the fixed process model and therefore do not fully recover the transient lovastatin evolution. In contrast, the Level 3 estimator (dashed purple) tracks the changes in the trajectory much more closely by combining physics-based prediction with a bounded residual correction learned from the assay information. Notably, the Level 3 estimate is not simply a sample-and-hold reconstruction: it evolves between assay times and provides a smoother and more accurate transient estimate than a zero-order-hold (ZOH) baseline. The ZOH baseline yielded an RMSE of 3.24×10^{-4}g/L and an MAE of 2.10×10^{-4}g/L, whereas Level 3 achieved lower errors, corresponding to a 41.3% reduction in RMSE and a 38.0% reduction in MAE relative to ZOH. These results demonstrate that residual learning provides a capability fundamentally distinct from bias handling, enabling accurate estimation under evolving bioprocess kinetics.

Table 3 summarizes the quantitative estimation performance of the three estimator levels across the three

evaluation scenarios: nominal operation, measurement bias, and kinetic mismatch. Under nominal conditions, all estimators achieve low prediction errors, with Level 1 and Level 2 showing comparable RMSE and MAE values, reflecting the adequacy of the baseline physics-based EKF when model assumptions hold. When a measurement bias is introduced, the error metrics of Level 1 increase substantially, whereas Level 2 maintains nearly unchanged performance due to its explicit bias-state formulation.

Under kinetic parameter uncertainty or change, Level 3 provided the lowest estimation error, while Levels 1 and 2 remain more strongly affected by the structural model discrepancy. Overall, the results confirm that each estimator level contributes a distinct and complementary capability: baseline state reconstruction (L1), robustness to sensor bias (L2), and adaptation to process-model mismatch (L3), supporting the proposed levelized estimation strategy.

Table 3. Quantitative comparison of lovastatin concentration estimation accuracy (RMSE, MAE) for Level 1–3 estimators under nominal operation, measurement bias, and kinetic mismatch scenarios.

Level	Index	EXP1	EXP2	EXP3
Level 1	RSME	0.00049	0.64920	0.00121
	MAE	0.00048	0.59454	0.00099
Level 2	RSME	0.00043	0.00043	0.00102
	MAE	0.00042	0.00042	0.00082
Level 3	RSME	8.5e-5	8.3e-5	0.00019
	MAE	7.9e-5	7.9e-5	0.00013

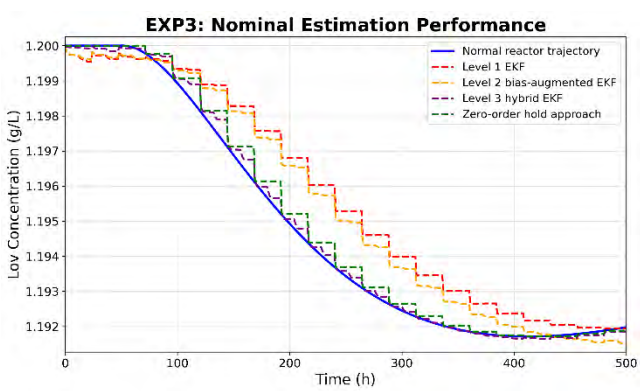

Figure 6. Lovastatin concentration under a true process–model mismatch. The reactor trajectory (blue) is shown together with Level 1 EKF (dashed red), Level 2 bias-augmented EKF (dashed orange), Level 3 hybrid EKF (dashed purple), and zero-order hold baseline (dashed green).

Perspective and Future Work

This work demonstrates the first three layers of the proposed Levelized Hybrid Estimation Architecture (LHEA), showing how physics-based EKF estimation can be systematically enhanced to improve robustness and accuracy under realistic sensing constraints. While this study focuses on state reconstruction, bias robustness, and physics-constrained residual learning, several extensions remain for future work. These include learned measurement models that convert auxiliary PAT signals (e.g., Raman or capacitance) into soft measurements, innovation-adaptive noise tuning for automated robustness, and physics-regularized surrogate models for operation under severe regime changes. Integrating the levelized estimators with closed-loop control and validating the framework on experimental or pilot-scale data are also important next steps. Together, these directions complete the envisioned framework and position LHEA as a scalable, interpretable pathway for incorporating machine learning into high-reliability soft sensing for continuous biomanufacturing

DIGITAL SUPPLEMENTARY MATERIAL

Digital supplementary material and codes associated with this paper are available at the following link: https://github.com/Boskabadi/LHEA

ACKNOWLEDGEMENTS

Luis Ricardez-Sandoval acknowledges the support from the Canada Research Chairs (CRC) program.

AUTHOR IDENTIFIERS

Author ORCIDs:
Mohammad Reza Boskabadi: 0009-0009-9619-3735
Rajiv Kailasanathan: 00009-0009-8915-7111
Luis Ricardez-Sandoval: 0000-0001-9867-6778
Seyed Soheil Mansouri: 0000-0002-0924-4947

REFERENCES

1. Kailasanathan R, Clements WR, Boskabadi MR, Gibford S, Papadakis E, Savoie CJ, Mansouri SS. Quantum enhanced ensemble gans for anomaly detection in continuous biomanufacturing. Ind. Eng. Chem. Res. 65:5046-5057 (2026). https://doi.org/10.1021/acs.iecr.5c04183
2. M. Drobnjakovic, R. Hart, B. Kulvatunyou, N. Ivezic, and V. Srinivasan, "Current challenges and recent advances on the path towards continuous biomanufacturing, " *Biotechnol. Prog.*, vol. 39, no. 6, p. e3378, 2023.
3. Liñán DA, McCormick-Mantilla JA, Ricardez-Sandoval LA. Economic nonlinear model predictive control and scheduling of multiple fed-batch fermenters in a lignocellulosic biorefinery.

Computers & Chemical Engineering 201:109223 (2025). https://doi.org/10.1016/j.compchemeng.2025.109223
4. Ghodba A, Richelle A, McCready C, Ricardez-Sandoval L, Budman H. A robust batch-to-batch optimization framework for pharmaceutical applications. Computers & Chemical Engineering 193:108935 (2025). https://doi.org/10.1016/j.compchemeng.2024.108935
5. Valipour M, Ricardez-Sandoval LA. Abridged gaussian sum extended kalman filter for nonlinear state estimation under non-gaussian process uncertainties. Computers & Chemical Engineering 155:107534 (2021). https://doi.org/10.1016/j.compchemeng.2021.107534
6. Valipour M, Ricardez-Sandoval LA. Assessing the impact of EKF as the arrival cost in the moving horizon estimation under nonlinear model predictive control. Ind. Eng. Chem. Res. 60:2994-3012 (2021). https://doi.org/10.1021/acs.iecr.0c06095
7. Boskabadi MR, Murugaiah M, Nielsen TR, Sivaram A, Sin G, Mansouri SS. Virtual sensor for sustainable large-scale process monitoring. Ind. Eng. Chem. Res. 64:3902-3917 (2025). https://doi.org/10.1021/acs.iecr.4c03342
8. Strielkowski W, Vlasov A, Selivanov K, Muraviev K, Shakhnov V. Prospects and challenges of the machine learning and data-driven methods for the predictive analysis of power systems: a review. Energies 16:4025 (2023). https://doi.org/10.3390/en16104025
9. R. Gani *et al.*, "Challenges and opportunities for process systems engineering in a changed world, " *Comput. Aided Chem. Eng.*, vol. 49, pp. 7–20, 2022.
10. Boskabadi MR, Ramin P, Kager J, Sin G, Mansouri SS. Kt-biologics i (KTB1): a dynamic simulation model for continuous biologics manufacturing. Computers & Chemical Engineering 188:108770 (2024). https://doi.org/10.1016/j.compchemeng.2024.108770
11. Bar-Shalom Y, Li XR, Kirubarajan T. Estimation with applications to tracking and navigation. Wiley (2002). https://doi.org/10.1002/0471221279
12. Kroll P, Sagmeister P, Reichelt W, Neutsch L, Klein T, Herwig C. ex situonline monitoring: application, challenges and opportunities for biopharmaceuticals processes. Pharmaceutical Bioprocessing 2:285-300 (2014). https://doi.org/10.4155/pbp.14.22
13. Bergin A, Carvell J, Butler M. Applications of bio-capacitance to cell culture manufacturing. Biotechnology Advances 61:108048 (2022). https://doi.org/10.1016/j.biotechadv.2022.108048
14. Popović I, Dončević L, Biba R, Košpić K, Barbalić M, Marinković M, Cindrić M. Advancements in adenine nucleotides extraction and quantification from a single drop of human blood. Molecules 29:5630 (2024). https://doi.org/10.3390/molecules29235630
15. Gerzon G, Sheng Y, Kirkitadze M. Process analytical technologies – advances in bioprocess integration and future perspectives. Journal of Pharmaceutical and Biomedical Analysis 207:114379 (2022). https://doi.org/10.1016/j.jpba.2021.114379
16. Garavaglia M, McGregor C, Bommareddy RR, Irorere V, Arenas C, Robazza A, Minton NP, Kovacs K. Stable platform for mevalonate bioproduction from co₂. ACS Sustainable Chem. Eng. 12:13486-13499 (2024). https://doi.org/10.1021/acssuschemeng.4c03561
17. Gibford SM, Boskabadi MR, Savoie CJ, and Mansouri SS. Quantum Synthetic Data Generation for Industrial Bioprocess Monitoring. arXiv preprint (2025) https://arxiv.org/pdf/2510.17688

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Understanding the Impact of Ribbon Splitting on Tablet Properties Using a Hybrid Mechanistic–Machine Learning Framework

Shumaiya Ferdoush[a], Mohammad Shahab[b], Xinle Zhang[b], Jayden A. Pierce[c], Emma Jeffries[b], Adaugo Ufomba[b], Zoltan K. Nagy[b], Gintaras V. Reklaitis[b], Marcial Gonzalez[ad*]

[a] School of Mechanical Engineering, Purdue University, West Lafayette, IN 47907, USA
[b] Davidson School of Chemical Engineering, Purdue University, 47907, IN, USA
[c] School of Materials Engineering, Purdue University, 47907, IN, USA
[d] Ray W. Herrick Laboratories, Purdue University, West Lafayette, IN 47907, USA
* Corresponding Author: marcial-gonzalez@purdue.edu .

ABSTRACT

Roller compaction is widely used in pharmaceutical manufacturing to improve powder flowability and enable robust tablet production. Although often treated as producing a homogeneous granule population, ribbons may undergo splitting during compaction, generating structurally distinct granules that affect downstream tableting. This study investigates the impact of ribbon splitting on tablet critical quality attributes (CQAs) for 10% and 20% acetaminophen (APAP) formulations. Reduced-order models (ROMs) proposed by Bachawala et al. [1] were applied to predict tablet density, elastic recovery, tensile strength, and tablet weight under split and non-split conditions. Although ribbon splitting alters the granule size distribution (GSD) and ribbon density, tablet CQAs such as tensile strength, elastic recovery, and tablet density are accurately predicted by the existing ROM framework, provided that GSD and ribbon density are known. In contrast, tablet weight predictions deteriorate when split and non-split data are combined, reflecting the sensitivity of die filling to granule packing changes induced by splitting. Separating the datasets improves weight prediction for non-split granules, while split granules remain challenging to model. A multitask Gaussian Process regression model is further used to analyze splitting as a process disturbance, highlighting the need for extended or data-driven approaches to accurately predict tablet weight under split conditions.

Keywords: Ribbon splitting, Roller compaction, Tableting, Granule Size Distribution, Reduced Order Models

INTRODUCTION

Roller compaction followed by milling is widely used in pharmaceutical manufacturing to enhance powder flowability and ensure robust tablet production, where powders are first consolidated into ribbons and then milled into granules. Although often treated as producing a single homogeneous ribbon/granule population, several studies have shown that ribbons exhibit spatial heterogeneity in solid fraction and stress state and may undergo distinct failure modes, such as longitudinal or transverse splitting upon exiting the rolls, leading to materially different structures entering downstream processes such as tableting [2], [3], [4], [5], [6].

In particular, Shahab et al. [6] demonstrated that ribbon splitting originates during the roller compaction stage as the ribbon exits the rolls. They attributed this behavior to non-uniform stress and density distributions across the ribbon thickness, coupled with differential elastic recovery, which generate tensile stresses exceeding the fracture strength of the compacted material. As a result, two physically distinct ribbon states may be produced under otherwise similar operating conditions: intact (non-split) ribbons and ribbons that undergo longitudinal splitting. These structural differences are subsequently inherited by the granules formed during milling, leading to variations in granule morphology and packing characteristics [7], [8], [9]. While Shahab et al. identified

and characterized split and non-split ribbons based on their fracture behavior at the roll exit during roller compaction, prior studies have not assessed whether ribbon splitting itself selectively influences tablet properties such as tablet weight, tablet density, or tensile strength.

Bachawala et al. [1] developed a reduced-order mechanistic framework linking granule size distribution, packing fraction, ribbon density, and compaction behavior to predict tablet critical quality attributes (CQAs), including density, tensile strength, and weight, demonstrating strong predictive capability across process conditions. However, the model assumes that ribbons are fully characterized by density and size distribution, without explicitly accounting for ribbon formation pathways or splitting behavior.

From a manufacturing standpoint, this assumption may be limiting, as ribbon splitting modifies the internal structure and effective packing of granules, potentially influencing die filling and mass deposition during tableting. As a result, some tablet properties may be sensitive to splitting, while others may remain unaffected.

The objective of this work is to examine the extent to which ribbon splitting influences tablet properties, specifically tablet weight, tablet density, and tensile strength, in order to identify which critical quality attributes are sensitive to splitting phenomena. Furthermore, building on the reduced-order modeling framework proposed by Bachawala et al. [1], this study evaluates the applicability and robustness of the existing models under split and non-split ribbon conditions, and assesses whether additional data-driven approaches, such as machine learning, are required to accurately predict tablet properties in cases where the reduced-order models are insufficient.

MATERIALS AND METHODS

Preparation and Characterization of Granules

Granules were produced using an Alexanderwerk WP120 roller compactor from two formulations: (i) 90% w/w microcrystalline cellulose (MCC; Avicel PH-102, International N&H USA, Inc.) with 10% w/w acetaminophen (APAP; Grade 0048, Mallinckrodt Pharmaceuticals), and (ii) 80% w/w Avicel PH-102 with 20% w/w APAP. Powder blends (3 kg) were prepared in a 5-L tote blender and pre-mixed for 30 min, then fed by a Coperion Ktron feeder into a Gericke GCM-250 continuous blender operating at 150 rpm before conveyance to the roller compactor. No lubricants or glidants were added.

The experimental campaign followed Shahab et al. [6], where roll gap, roll pressure, and roll speed were varied while screw feed speed was adjusted to achieve target set points. A three-level Box–Behnken design of experiments was employed to reduce experimental runs while enabling robust analysis [2], [3]. Thirty-two granule batches (sixteen per formulation) were produced by varying hydraulic pressure (30–90 bar), roll gap (1.4–2.4 mm), and roll speed (3.5–5.5 rpm). Ribbons were milled using the integrated two-stage hammer mill with 2.5 and 1.25 mm screens. Ribbon density was measured using a GeoPyc 1360 pycnometer, while granule samples collected at the compactor outlet were analyzed for GSD using a Canty SolidSizer and used for tablet production on a rotary tablet press. The measured GSD was modeled using a bimodal Weibull distribution.

$$P(x) = \alpha \left[\frac{k_1}{\lambda_1} \left(\frac{x}{\lambda_1} \right)^{k_1 - 1} e^{-\left(\frac{x}{\lambda_1}\right)^{k_1}} \right] + (1-\alpha) \left[\frac{k_2}{\lambda_2} \left(\frac{x}{\lambda_2} \right)^{k_2 - 1} e^{-\left(\frac{x}{\lambda_2}\right)^{k_2}} \right] \quad (1)$$

where α denotes the fraction of the small mode and (λ_1, k_1) and (λ_2, k_2) are the scale and shape parameters of the small and large modes, respectively. The mean granule size is then calculated as a weighted average of the mean sizes of each mode ($\mu_{GSD} = \alpha\mu_1 + (1-\alpha)\mu_2$) with $\mu_1 = \lambda_1 \Gamma(1 + 1/k_1)$, $\mu_2 = \lambda_2 \Gamma(1 + 1/k_2)$, and $\Gamma(\cdot)$ denotes the gamma function.

Preparation and Characterization of Tablets

Granules were tableted on a Natoli NP-400 rotary press using doubly convex D-type tooling, producing thirty two tablet batches. The dosing position and pre-compression thickness were fixed at 9 mm and 6 mm, respectively, with a feeder speed of 35 rpm. The main compression thickness was varied between 2.5 and 3.5 mm, while the turret speed was adjusted from 15 to 20 rpm. After manufacturing, tablets were stored in Ziploc bags for 24 h to allow elastic relaxation and were then tested using a SOTAX AT4 tablet tester to measure tablet weight, tensile strength, tablet density, and elastic recovery following the work of [1]. The tablet volumes at different compaction stages [10] are determined from tablet geometry following [1], where t^{fill}, $t^{in\text{-}die}$, and t^{tablet} represent the fill depth, in-die thickness, and final tablet thickness, respectively, and h is the punch cup depth. Using the tablet weight and true density, the in-die and final tablet relative densities are then calculated. Tablet tensile strength is obtained from the measured crushing force F using Pitt's model as recommended by Razavi et al. [11], with the full expression available in the cited literature [1]. It is worth mentioning that many of these quantities are not available online during manufacturing, motivating the need for predictive reduced-order models to estimate Critical Material Attributes (CMAs) and CQAs for process monitoring and control.

Tableting Reduced-Order Models

Mechanistic and semi-mechanistic reduced-order models (ROMs) provide a practical balance between physical interpretability and computational simplicity, enabling their use in process design, optimization, and control of pharmaceutical tableting [12], [13]. The ROM

framework describes die filling, compaction, unloading, and ejection using physically based equations that link material and process variables to tablet CQAs [1]. The ROM framework was employed by Bachawala et al. [1], [14] to establish a quantitative link between dry granulation and the downstream tableting process. The mathematical expressions for die-filling efficacy coefficient (η), the elastic recovery (ϵ_ρ), the tablet density (ρ^{tablet}), and the Leunberger's tablet tensile strength (σ_t) are listed below:

$$\eta = \frac{W}{\rho_b V^{fill}} \tag{2}$$

$$\epsilon_\rho = \epsilon_0 \left(\frac{\rho^{in-die} - \rho_{c,\epsilon}}{1-\rho_{c,\epsilon}} \right)^n \tag{3}$$

$$\rho^{tablet} = \rho^{in-die}(1-\epsilon_\rho) \tag{4}$$

$$\sigma_t = \sigma_0 \left[1 - \left(\frac{1-\rho^{tablet}}{1-\rho_{c,\sigma_t}} \right) e^{(\rho^{tablet} - \rho_{c,\sigma_t})} \right] \tag{5}$$

Here, ρ_b is the bulk density of the powder blend. The parameter ϵ_0 represents the elastic recovery at full compaction (zero porosity), $\rho_{c,\epsilon}$ is the minimum in-die relative density at which elastic recovery begins, and n accounts for the nonlinearity of this behavior. Tablet tensile strength σ_t depends on the tablet relative density after elastic recovery ρ_{tablet}, the maximum tensile strength at full densification σ_0, and the critical relative density ρ_{c,σ_t} required for forming a mechanically stable tablet.

Model Selection and Parameter Estimation

Bachawala et al. [1] developed a reduced-order modeling (ROM) framework that propagates dry granulation attributes—ribbon density and granule size distribution—into tableting ROMs using normalized rational functions, with model selection guided by the Akaike Information Criterion (AIC). In this work, the same framework is adopted to evaluate the applicability and robustness of these ROMs under split and non-split ribbon conditions, with emphasis on predicting tablet weight, density, and tensile strength. Upstream CMAs and critical process parameters (CPPs) are coupled to the tableting ROM parameters through nonlinear mappings of normalized variables, $X = \left((x-lb)/(ub-x)\right)^{r_X}$, where lb and ub are variable bounds and r_X accounts for nonlinear effects. Each ROM parameter f is expressed as a bivariate rational function,

$$f(X,Y) = \frac{p_1XY + p_2X + p_3Y + p_4}{q_1XY + q_2X + q_3Y + 1} \tag{6}$$

with parameter set $\theta = p_1, p_2, p_3, p_4, q_1, q_2, q_3, r_X, r_Y$ estimated from experimental data. Model complexity is controlled by progressively simplifying the rational function to form a library of candidate models, with selection based on:

$$AIC = n \ \ln \frac{SSE}{n} + 2Np \tag{7}$$

where n is the number of data points, SSE the sum of squared errors, and N_p the number of parameters. The lower the AIC values, the better the model. The optimal model structure m and its parameters θ are given by:

$$m, \theta = \arg\min_{m \in F} \left[n \ln \left(\frac{1}{n} \min_{\theta^m \in D_m} \left(\sum_{i=1}^{n} \left(M^i(\theta^m) - E^i \right)^2 \right) \right) \right] + 2N_p^m \tag{8}$$

Here, $M^i(\theta^m)$ denotes the model prediction, E^i the experimental value, n the number of data points, N_p^m the number of parameters, and D_m the feasible parameter space. Parameter estimation is carried out using MATLAB's genetic algorithm, followed by `fmincon` and `GlobalSearch`, to reduce sensitivity to initial guesses and improve convergence to the global optimum.

Although the 10% APAP formulation matches that of Bachawala et al. [1], parameters are re-estimated here because a different MCC 102 manufacturer was used, and all ribbon density and GSD data were measured experimentally, rather than predicted by machine learning models as in Huang et al. [15]. Moreover, the tablets produced in this study exhibit higher densities compared to those reported in the authors' work.

Machine Learning framework

Gaussian Process (GP) regression is a Bayesian nonparametric method for modeling nonlinear input–output relationships with uncertainty quantification. A GP defines a prior over latent functions $f(x) \sim GP\big(m(x), k(x,x')\big)$, where $m(x)$ is the mean function and $k(x,x')$ the covariance kernel, and observations are modeled as $y = f(x) + \varepsilon$, with $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. Given training data, posterior predictions yield analytical expressions for both predictive means and variances, with hyperparameters learned by maximizing the marginal likelihood.

To jointly model correlated tablet properties, multitask Gaussian Process regression (MT-GPR) was implemented using the Intrinsic Coregionalization Model (ICM), where the joint covariance between outputs t and t'at inputs x and x' is defined as $K\big((x,t),(x',t')\big) = k_{data}(x,x') \cdot B_{t,t'}$, where k_{data} captures input-space similarity, while B is a low-rank coregionalization matrix that encodes inter-task correlations, enabling information sharing across outputs while preserving task-specific variability. The input kernel was specified as an additive combination of a linear kernel and a Matérn kernel with automatic relevance determination to capture both global trends and smooth nonlinear effects. Task-specific observation noise was modeled using a multitask Gaussian likelihood with weakly informative priors, with continuous inputs standardized prior to training and binary indicators left unscaled. This multitask formulation enables coherent joint inference and uncertainty propagation across tablet properties, which is essential for analyzing process disturbances affecting multiple CQAs simultaneously.

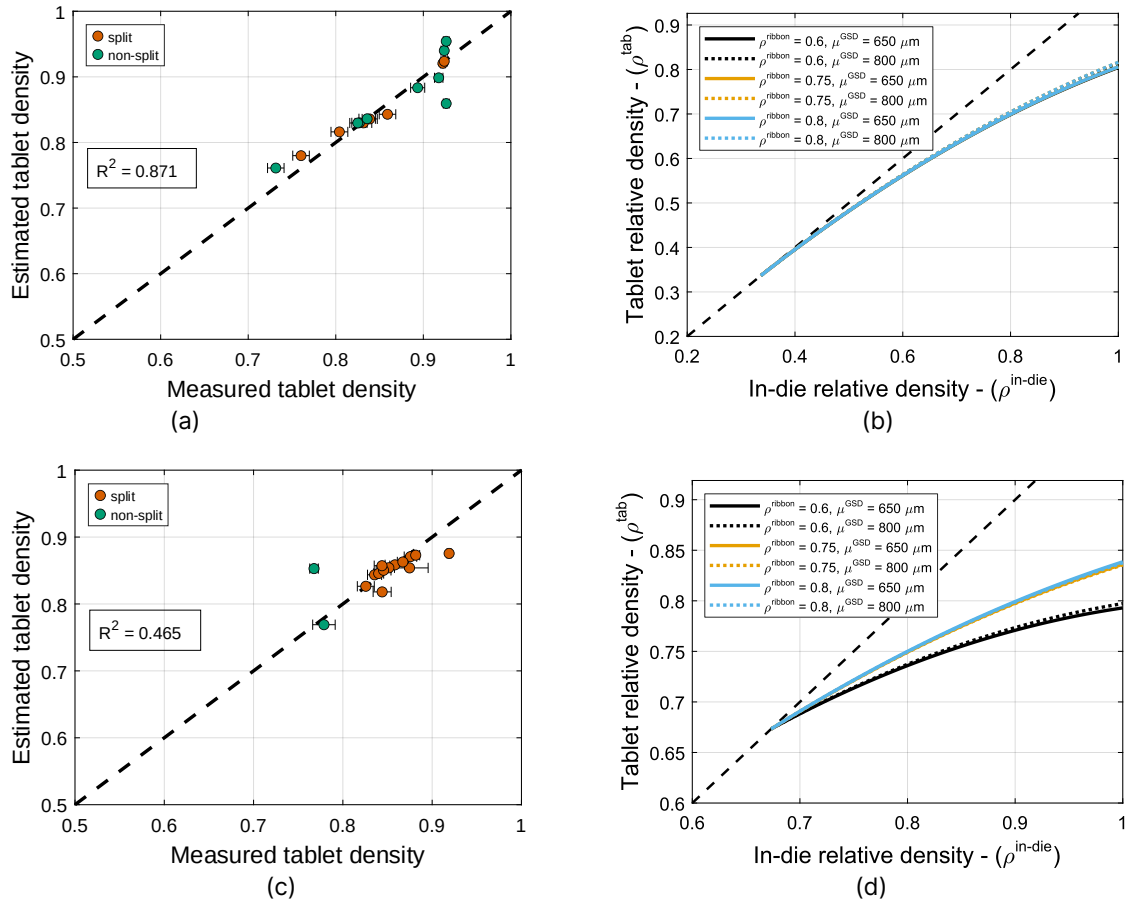

Figure 1: The top row corresponds to the 10% APAP formulation and the bottom row to the 20% APAP formulation, (a), (c) goodness of the tablet density prediction, and (b), (d) prediction of the tablet density model.

RESULTS AND DISCUSSION

Impact of Splitting on Granule Size Distribution and Ribbon Density

Ribbon splitting during roller compaction produces structurally weakened ribbons with internal cracks, resulting in reduced effective ribbon density. These defects promote enhanced fragmentation during milling, shifting the granule size distribution toward finer particles and increasing fines content. In contrast, non-split ribbons retain higher structural integrity and generate coarser, more stable granule populations. Similar relationships between ribbon failure modes, granule size outcomes, and ribbon mechanical integrity have been reported in prior studies on roller compaction and ribbon fracture behavior [6].

Rational Function-Based Approach for Integrating Dry Granulation with Tableting

In this section, the framework of the previous section is used to integrate dry granulation with tableting as demonstrated by Bachawala et al [1]. The ROMs for tablet weight, tablet density, elastic recovery, and tablet tensile strength are presented sequentially. The best-performing models proposed by Bachawala et al. [1] have been adopted for parameter estimation and analysis of the results.

Elastic recovery and tablet density

Ribbon relative density ρ^{ribbon} and mean granule size μ^{GSD} are selected as the two upstream CMAs governing the coupling between unit operations. Their normalized forms are defined as:

$$X = \left(\frac{\rho^{ribbon}-\rho_{c,rib}}{1-\rho^{ribbon}}\right)^{r_X}, \qquad Y = \left(\frac{\mu^{GSD}}{D/6-\mu^{GSD}}\right)^{r_Y} \tag{9}$$

The ROM parameters for the elastic recovery model (equation (3)) are expressed as

$$\epsilon_0 = \frac{p_1XY+p_2X+p_4}{q_1XY+q_2X+1}, \qquad \rho_{c,\epsilon} = p_4 \tag{10}$$

For the best model, $r_X = r_Y = 1$ in Equation (9). Parameters were estimated simultaneously for split and non-split data as described previously.

Figure 1 summarizes the results for both formulations. The tablet density model yields $R^2 = 0.87$ for 10% APAP (Figure 1(a)) and a lower $R^2 = 0.46$ for 20% APAP (Figure 1(c)), attributed to the limited range of measured densities, which could be improved by expanding the experimental space. The relationship between tablet density and in-die density is shown in Figures 1(b) and 1(d).

Ribbon splitting reduces effective ribbon density and shifts the GSD toward finer particles. Since ribbon density and granule size are included as model inputs, the ROMs of Bachawala et al. capture trends accurately for both split and non-split conditions for 10% APAP formulations.

In Equation (3), the fitted exponent n is 1.20 for the 10% APAP formulation and 1.003 for the 20% APAP formulation. For Equation (10), the estimated parameters for 10% APAP are $p_1 = 4.39, p_2 = 4.72, p_4 = 0.201, q_1 = 28.64$, and $q_2 = 18.64$, while for 20% APAP they are $p_1 = 2.44, p_2 = 5.67, p_4 = 0.423, q_1 = 14.45$, and $q_2 = 37.19$. In Equation (10), $\rho_{c,\epsilon} = p_4$ is estimated as 0.336 for 10% APAP and 0.673 for 20% APAP.

Tablet tensile strength

Nordström et al [16] showed that tablet tensile strength depends on the coupled effects of particle size ratio and the fraction of large particles rather than mean granule size, whereas most earlier studies focused on monodisperse or narrowly sieved granules, reporting either stronger tablets for smaller granules [17] or no significant effect of granule size on tensile strength [18]. Hence, Bachawala et al. expanded the model library to include dependence on ribbon density ρ^{ribbon} and packing fraction ϕ through the normalized variables

$$X = \left(\frac{\rho^{ribbon} - \rho_{c,rib}}{1 - \rho^{ribbon}}\right)^{r_X}, Z = \left(\frac{\phi - \phi_{min}}{\phi_{max} - \phi}\right)^{r_Z} \quad (12)$$

The ROM parameters for the tensile strength model (equation (5)) are expressed as

$$\sigma_0 = \frac{p_1 XZ + p_2 X + p_3 Z + p_4}{q_1 XZ + q_2 X + q_3 Z + 1}, \qquad \rho_{c,\sigma_t} = \frac{p_1 XZ + p_2 X + p_3 Z + p_4}{q_1 XZ + q_2 X + q_3 Z + 1} \quad (13)$$

Parameter estimation results are reported in Table 1, with predictions shown in Figure 2 and 3 for both 10% and 20% formulations. The ROMs exhibit excellent predictive performance for both formulations ($R^2 \approx 0.99$). Although ribbon splitting modifies the GSD and consequently influences the packing fraction, this effect is inherently accounted for in the tensile strength model through the inclusion of packing fraction as an input, enabling the ROMs to accurately capture the experimental trends.

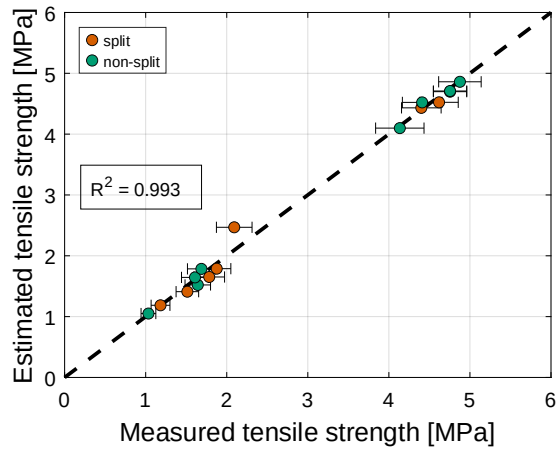

Figure 2: Goodness of the tensile strength prediction for 10% APAP.

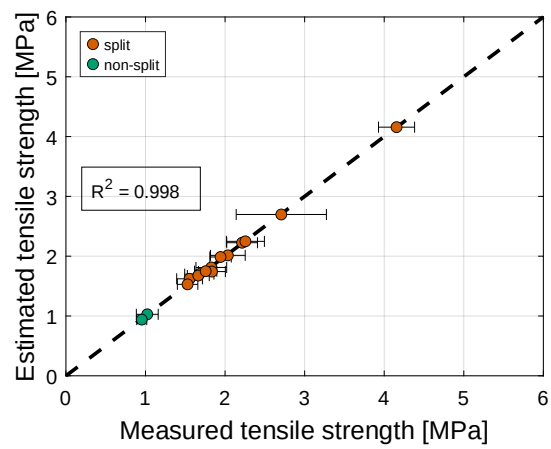

Figure 3: Goodness of the tensile strength prediction for 20% APAP.

Table 1: Estimated parameters for the tensile strength models are listed below.

Equations	Parameters 10% APAP	Parameters 20% APAP
Equation 12	$r_X = 3.21$ $r_Z = 0.05$	$r_X = 9.98$ $r_Z = 5.72$
Equation 13 σ_0	$p_1 = 215.07$ MPa $p_2 = 71.86$ MPa $p_3 = 59.51$ MPa $p_4 = 0.08$ MPa $q_1 = 0.04$ $q_2 = 15.66$ $q_3 = 5.34$	$p_1 = 1.28$ MPa $p_2 = 345.79$ MPa $p_3 = 0.34$ MPa $p_4 = 422.84$ MPa $q_1 = 6.54e-04$ $q_2 = 145.89$ $q_3 = 0.09$
Equation 13 ρ_{c,σ_t}	$p_1 = 326.06$ MPa $p_2 = 0.004$ MPa $p_3 = 6.79$ MPa $p_4 = 8.91e-05$ MPa $q_1 = 336.12$ $q_2 = 19.19$ $q_3 = 6.99$	$p_1 = 0.07$ MPa $p_2 = 166.87$ MPa $p_3 = 0.003$ MPa $p_4 = 0.997$ MPa $q_1 = 0.08$ $q_2 = 235.86$ $q_3 = 0.005$

Tablet weight

For tablet weight, coupling is introduced through the packing fraction ϕ, determined from:

$$W = \eta\rho_b V^{fill} = \phi(X,Y)\rho_{true}\rho_{ribbon}V_{fill} \quad (14)$$

The packing fraction is modeled using a simplified rational function:

$$\phi(X,Y) = p_4 \frac{q_2X+1}{q_1XY+q_2X+1} \quad (15)$$

with $X = (1-\alpha) * (1-\mu_1{}^3/\mu_2{}^3)$ and $Y = \alpha$.

Here, μ_1/μ_2 is the size ratio of the bimodal distribution. This formulation is motivated by the analytical expression proposed by Brouwers and co-workers for the packing fraction of bimodal hard spheres, which relates packing efficiency to the size ratio and composition of the mixture [19], [20].

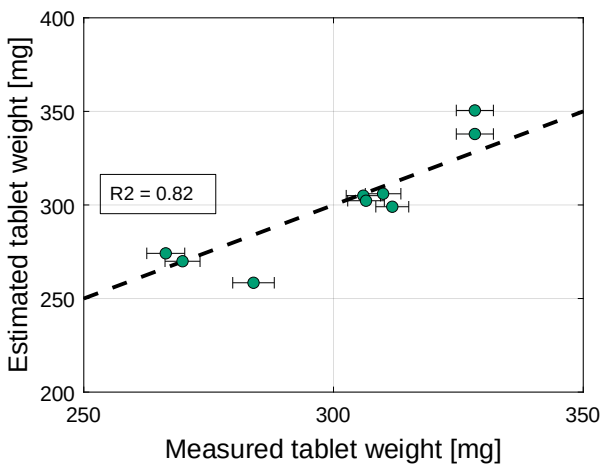

Figure 3: For 10% APAP, goodness of the tablet weight prediction for non-split data.

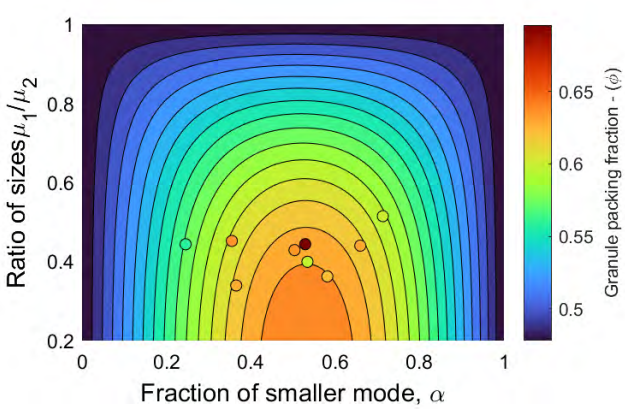

Figure 4: Predictions of the packing fraction model for 10% APAP (non-split data).

Parameter estimation was performed using combined split and non-split data, and the tablet weight model showed limited accuracy with $R^2 \approx 0.6$. Although the tablet weight model depends on packing fraction ϕ and ribbon density ρ, ribbon splitting alters granule shape, surface roughness, porosity, and cohesion, which modify die filling dynamics and effective fill efficiency in ways not fully captured by these two inputs, leading to lower accuracy in weight prediction [14], [21], [22].

The datasets were therefore separated, and parameters were estimated independently for split and non-split granules. For the non-split 10% APAP formulation, tablet weight prediction improves markedly (R^2 = 0.82, Figure 3), and the contour plot (Figure 4) reproduces the trend reported by Bachawala et al, with a maximum packing fraction at α ≈ 0.7 and increasing packing fraction with decreasing μ_1/μ_2 for 0.6 < α < 0.9. By contrast, predictions for split granules show lower accuracy (R^2 = 0.552) and are not shown here, indicating that the current tablet weight model does not adequately capture split behavior and may require modification or alternative approaches such as machine learning. For the 20% APAP formulation, limited non-split data prevented separation, resulting in lower prediction accuracy (R^2 = 0.616, Figure 5) as the model is dominated by split conditions.

The estimated parameters (Equation (15)) for the 10% APAP formulation are $p_4 = 0.48, q_1 = -1.23$, and $q_2 = 0.32$. For the 20% APAP formulation, the corresponding parameters are $p_4 = 0.63, q_1 = -0.10$, and $q_2 = 999.99$. The value of $\rho_{c,rib}$ was kept constant at 0.566, and the values of ϕ_{min} and ϕ_{max} were set to 0.367 and 0.75, respectively.

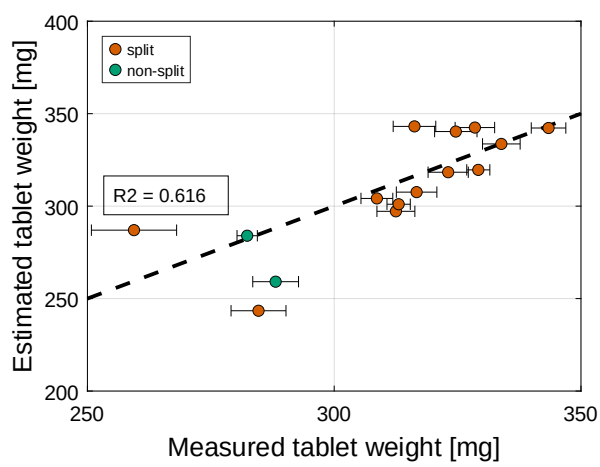

Figure 5. Goodness of prediction for tablet weight (both split and non-split data included) for 20% APAP.

Impact of Splitting on Tablet Properties

To quantify the impact of splitting on tablet properties, a multitask Gaussian Process Regression (MT-GPR) model was implemented to simultaneously predict four CQAs: tablet weight (W), elastic recovery (ER), relative density (RD), and tensile strength (TS). The multitask

formulation exploits correlations among outputs, improving robustness relative to independent single-output models. The covariance kernel was defined as an additive combination of linear and Matérn ($\nu = 2.5$) kernels, $K(x, x') = K_{\text{Linear}} + K_{\text{Matérn}}$, enabling the capture of both linear trends and smooth nonlinear dependencies in process variables (e.g., roll gap, pressure, and speed). Automatic relevance determination (ARD) was employed to weight input contributions. The model was implemented in GPyTorch using a multitask Gaussian likelihood, and hyperparameters were optimized by maximizing the exact marginal log-likelihood with the Adam optimizer (learning rate 0.03, 2000 iterations). All continuous inputs and outputs were standardized prior to training to ensure numerical stability and interpretability.

Table 2: The predictive performance of the MT-GPR model and independent GP models.

Model	RMSE	MAE	NLPD	Coverage (90%)
Multitask GP	**0.41**	**0.37**	4.4	**0.91**
Ind. GPs	0.58	0.39	**1.21**	0.77

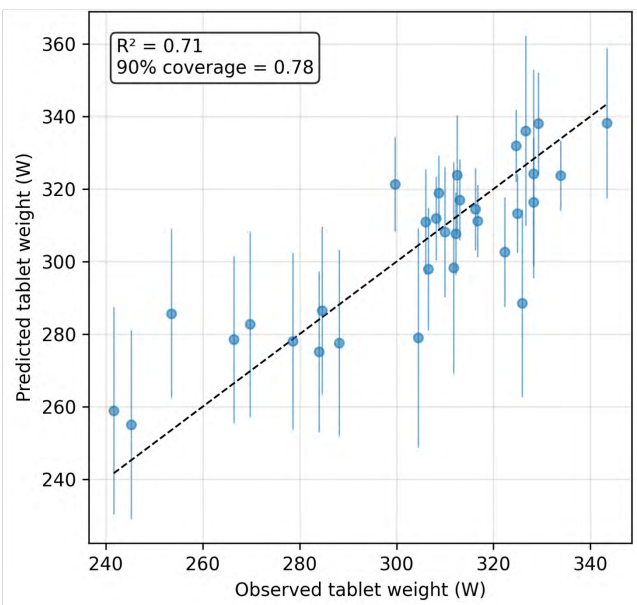

Figure 6. Goodness of prediction for tablet weight.

The predictive performance of the MT-GPR model was evaluated against independent GP models using leave-one-out cross-validation (Table 2), with RMSE, MAE, negative log predictive density (NLPD), and 90% predictive interval coverage as metrics. For tablet weight, the goodness of MT-GPR prediction is shown in Figure 6. The MT-GPR achieved lower overall RMSE and MAE, indicating improved average accuracy, and demonstrated superior uncertainty calibration through improved predictive interval coverage. While independent GPs achieved lower RMSE for tablet weight (0.591), the multitask model provided substantially improved accuracy for elastic recovery (0.027) and relative density (0.027), reflecting the inherent trade-off in multitask learning whereby information sharing benefits noisier outputs at the cost of modest degradation for dominant responses. MT-GPR yields coherent joint predictions, and consistent uncertainty estimates across all tablet properties, which are not guaranteed with independent models, making it better suited for multivariate, uncertainty-aware analysis of coupled process disturbances such as ribbon splitting.

To quantify splitting effects, a binary split indicator was included, and a counterfactual impact analysis was performed by predicting tablet properties across the experimental design under forced control (Split = 0) and failure (Split = 1) scenarios. The resulting differences (Figure 7) represent the predicted change in tablet properties due to splitting. Splitting emerges as a significant process disturbance, primarily reducing tablet weight (−3.16 mg ± 6.43) and tensile strength (−0.49 MPa ± 0.8). The large variability indicates strong interactions between splitting and roll conditions, localizing quality risk within specific regions of the design space. Splitting is associated with over-compaction or high elastic recovery in the nip region, producing structurally compromised, work-hardened ribbons that generate finer, bimodal GSDs during milling. Elevated fines degrade hopper flow and die filling consistency, reducing tablet weight and increasing variability, while reduced deformability and bonding area lead to lower tensile strength.

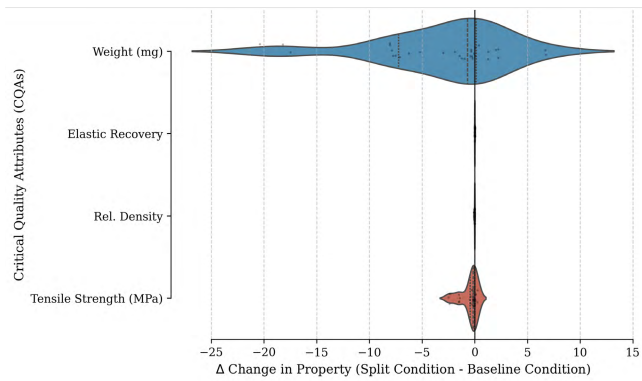

Figure 7. Impact of splitting on tablet properties using counterfactual impact analysis.

CONCLUSION

Ribbon splitting was investigated as a process disturbance in dry granulation and its impact on downstream tablet CQAs was quantified using reduced-order and machine-learning models. Although ribbon splitting alters the granule size distribution (GSD) and ribbon density, tablet CQAs such as tensile strength, elastic recovery, and tablet density are accurately predicted by the

existing ROM framework, provided that GSD and ribbon density are known. In contrast, tablet weight is strongly affected by splitting due to changes in granule packing and die filling. Separating split and non-split data improves weight prediction for non-split ribbons, but split granules remain challenging to model mechanistically. These findings highlight the need for extended or data-driven approaches to ensure robust tablet weight prediction under splitting conditions.

ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation grant No. 2140452, titled "Right First Time Manufacturing of Pharmaceuticals (RIFTMaP)". The authors gratefully acknowledge Natoli Engineering Company for providing the Natoli NP-400 tablet press and Alexanderwerk (Germany) for supplying the WP120 roller compactor used in this study.

REFERENCES

1. Bachawala S, Lagare RB, Delaney AB, Nagy ZK, Reklaitis GV, Gonzalez M. Rational function-based approach for integrating tableting reduced-order models with upstream unit operations: dry granulation case study. Pharmaceuticals 17:1158 (2024). https://doi.org/10.3390/ph17091158
2. Mahmah O, Adams MJ, Omar CS, Gururajan B, Salman AD. Roller compaction: ribbon splitting and sticking. International Journal of Pharmaceutics 559:156-172 (2019). https://doi.org/10.1016/j.ijpharm.2019.01.031
3. Akseli I, Iyer S, Lee HP, Cuitiño AM. A quantitative correlation of the effect of density distributions in roller-compacted ribbons on the mechanical properties of tablets using ultrasonics and x-ray tomography. AAPS PharmSciTech 12:834-853 (2011). https://doi.org/10.1208/s12249-011-9640-z
4. Lück M, Klinken-Uth S, Kleinebudde P. Kinetics of elastic recovery in roll compaction. International Journal of Pharmaceutics: X 8:100303 (2024). https://doi.org/10.1016/j.ijpx.2024.100303
5. Miguélez-Morán AM, Wu CY, Seville JPK. The effect of lubrication on density distributions of roller compacted ribbons. International Journal of Pharmaceutics 362:52-59 (2008). https://doi.org/10.1016/j.ijpharm.2008.06.009
6. Shahab M, Sixon D, Pierce JA, Yang MH, Gonzalez M, Nagy ZK, Reklaitis GV. Enhanced ribbon quality in roller compaction process by mitigating splitting through a machine-learning framework. International Journal of Pharmaceutics 685:126174 (2025). https://doi.org/10.1016/j.ijpharm.2025.126174
7. Perez-Gandarillas L, Perez-Gago A, Mazor A, Kleinebudde P, Lecoq O, Michrafy A. Effect of roll-compaction and milling conditions on granules and tablet properties. European Journal of Pharmaceutics and Biopharmaceutics 106:38-49 (2016). https://doi.org/10.1016/j.ejpb.2016.05.020
8. Kazemi P, Khalid MH, Pérez Gago A, Kleinebudde P, Jachowicz R, Szl?k J, Mendyk A. Effect of roll compaction on granule size distribution of microcrystalline cellulose–mannitol mixtures: computational intelligence modeling and parametric analysis. DDDT Volume11:241-251 (2017). https://doi.org/10.2147/dddt.s124670
9. Miguélez-Morán AM, Wu CY, Dong H, Seville JPK. Characterisation of density distributions in roller-compacted ribbons using micro-indentation and x-ray micro-computed tomography. European Journal of Pharmaceutics and Biopharmaceutics 72:173-182 (2009). https://doi.org/10.1016/j.ejpb.2008.12.005
10. Gonzalez M. Generalized loading-unloading contact laws for elasto-plastic spheres with bonding strength. Journal of the Mechanics and Physics of Solids 122:633-656 (2019). https://doi.org/10.1016/j.jmps.2018.09.023
11. Razavi SM, Gonzalez M, Cuitiño AM. General and mechanistic optimal relationships for tensile strength of doubly convex tablets under diametrical compression. International Journal of Pharmaceutics 484:29-37 (2015). https://doi.org/10.1016/j.ijpharm.2015.02.030
12. Ferdoush S, Gonzalez M. Semi-mechanistic reduced order model of pharmaceutical tablet dissolution for enabling industry 4.0 manufacturing systems. International Journal of Pharmaceutics 631:122502 (2023). https://doi.org/10.1016/j.ijpharm.2022.122502
13. Ferdoush S, Gonzalez M. A two-stage mechanistic reduced-order model of pharmaceutical tablet dissolution: population balance modeling and tablet wetting functions. International Journal of Pharmaceutics 664:124635 (2024). https://doi.org/10.1016/j.ijpharm.2024.124635
14. Bachawala S, Nasilowski DT, Gonzalez M. Rational function-based approach for integrating tableting reduced-order models with upstream unit operations: lubricants and glidants case study. Pharmaceuticals 18:1514 (2025). https://doi.org/10.3390/ph18101514
15. Y. S. Huang *et al.*, "A Machine Learning-assisted Hybrid Model to Predict Ribbon Solid Fraction, Granule Size Distribution and Throughput in a Dry Granulation Process, " in *Computer Aided Chemical Engineering*, vol. 52, Elsevier B.V., 2023, pp. 813–830. doi: 10.1016/B978-0-443-15274-0.50130-X.

16. Nordström J, Alderborn G, Frenning G. Compressibility and tablet forming ability of bimodal granule mixtures: experiments and DEM simulations. International Journal of Pharmaceutics 540:120-131 (2018). https://doi.org/10.1016/j.ijpharm.2018.02.006
17. HERTING M, KLEINEBUDDE P. Studies on the reduction of tensile strength of tablets after roll compaction/dry granulation. European Journal of Pharmaceutics and Biopharmaceutics 70:372-379 (2008). https://doi.org/10.1016/j.ejpb.2008.04.003
18. Mitra B, Hilden J, Litster J. Assessment of intragranular and extragranular fracture in the development of tablet tensile strength. Journal of Pharmaceutical Sciences 107:2581-2591 (2018). https://doi.org/10.1016/j.xphs.2018.05.011
19. Brouwers HJH. Random packing fraction of bimodal spheres: an analytical expression. Phys. Rev. E 87: (2013). https://doi.org/10.1103/physreve.87.032202
20. Brouwers HJH. Packing fraction of crystalline structures of binary hard spheres: a general equation and application to amorphization. Phys. Rev. E 78: (2008). https://doi.org/10.1103/physreve.78.011303
21. Sun C(, Himmelspach MW. Reduced tabletability of roller compacted granules as a result of granule size enlargement. Journal of Pharmaceutical Sciences 95:200-206 (2006). https://doi.org/10.1002/jps.20531
22. Dong R, Goodwin DJ, Nassar J, Dhenge R, Matrali SSH, Hodnett N, Zeitler JA. Roller compaction: measuring ribbon porosity by terahertz spectroscopy and machine learning. International Journal of Pharmaceutics 667:124852 (2024). https://doi.org/10.1016/j.ijpharm.2024.124852

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Physics-Informed Neural Networks for NIR Spectroscopy Analysis of Pharmaceutical Tablet Properties

Xinle Zhang[a], Shumaiya Furdoush[b], Marcial Gonzalez[bc], Gintaras V. Reklaitis[a]*

[a] Purdue University, Davidson School of Chemical Engineering, West Lafayette, IN, USA
[b] Purdue University, School of Mechanical Engineering, West Lafayette, IN, USA
[c] Purdue University, Ray W. Herrick Laboratories, West Lafayette, IN, USA
* Corresponding Author: greklait@purdue.edu.

ABSTRACT

In pharmaceutical process engineering, accurate prediction of tablet properties is crucial for ensuring product quality, optimizing manufacturing efficiency, and advancing sustainable production practices. This study presents a physics-informed neural network (PINN) framework for predicting the physical properties of pharmaceutical tablets from near-infrared (NIR) spectra. The PINN framework integrates revised Kubelka-Munk theory and physical constraints to ensure physically consistent predictions while requiring less training data than conventional artificial neural networks. Tablets were manufactured using acetaminophen and microcrystalline cellulose formulations with varying compositions and compression settings. The PINN framework successfully predicts critical quality attributes, including tensile strength, porosity, and density. It offers a data-efficient, interpretable solution for pharmaceutical tablet quality control.

Keywords: Physics-Informed Neural Networks, Machine Learning, Industry 4.0, Near Infrared Spectroscopy, Pharmaceutical Tablets.

INTRODUCTION

Tablet quality attributes including size, hardness, tensile strength, porosity and density are critical for ensuring mechanical integrity, dissolution performance and bioavailability in pharmaceutical manufacturing [1,2]. The US Food and Drug Administration (FDA) requires rigorous monitoring of these properties as critical quality. However, different tablet compression machines, operating conditions and formulation compositions can introduce substantial quality variations. Traditional tablet testing methods are time consuming and destructive and only provides limited inspection into the sampled tablets.

Near-infrared (NIR) spectroscopy enables rapid, non-destructive and continuous analysis of tablets, which is suitable for process monitoring [3, 4]. Current NIR calibration methods employ chemometric techniques such as partial least squares regression (PLSR) or machine learning models like artificial neural networks (ANNs). However, in predicting physical properties, PLS models often struggles with accuracy, and ANNs requires substantial training datasets and offer limited model interpretability [5, 6].

The NIR reflectance of tablets is fundamentally governed by light scattering and absorption in porous media. The revised Kubelka-Munk theory (RKM), developed by Yang and Kruse, describes scattering-induced photon path variation (SIPV) to the original Kubelka-Munk theory. By introducing the SIPV factor μ, the RKM theory accounts for changes in the effective light path length in highly scattering samples [7]. This provides a physics-based framework for relating spectral features to tablet microstructure. Physics-informed neural networks (PINNs) integrate physical laws into the training process of neural networks, which can improve data efficiency and prediction consistency [8].

In this study, we present a PINN model that incorporates RKM equations to predict tablet properties based on NIR spectra. The model learns the intrinsic absorption and scattering coefficients, as well as the SIPV factor, while predicting tablet critical quality attributes. We compare the performance of the PINN-RKM model with conventional ANN model and evaluate the data requirements for both approaches. The results show that integrating physical constraints reduces the need for training data while maintaining prediction accuracy.

METHODS

Experiments

A powder formulation was prepared for the tablet production: 80% microcrystalline cellulose (Avicel PH-102, IFF, USA) and 20% acetaminophen (Mallinckrodt Pharmaceuticals). The powders were blended for 30 minutes in a five-liter tote blender and then processed through an Alexanderwerk WP120 roller compactor to produce granules. Ten granulation sets were produced for each formulation with varying roller compaction settings (roller speed, hydraulic pressure, and roll gap), while three granulation sets were further used to evaluate the proposed models.

Tablets were manufactured using a Natoli NP-400 tablet press with D-type tooling (D = 7.94 mm, h = 0.3302 mm) at a constant turret speed of 20 rpm, feed-frame speed of 35 rpm and precompression thickness of 6mm. Three main compression thickness settings (3.0, 3.5, and 4.5 mm) were used for each granule set.

Data Acquisition

For each tableting condition, 20 tablets were sampled and measured sequentially. Each tablet was measured three times. Near-infrared (NIR) reflectance spectra (1100–2100 nm) were acquired at-line with the probe in direct contact with the tablet surfaces. External referencing was performed before each dataset using a certified 99% reflectance Spectralon standard. An example of raw tablet spectra is shown in Figure 1. After the NIR measurement, the tablets were characterized using an AT4 tablet tester to determine their weight, hardness, diameter, and thickness. Tensile strength was calculated from the hardness measurements using the method described in Bachawala et al [].

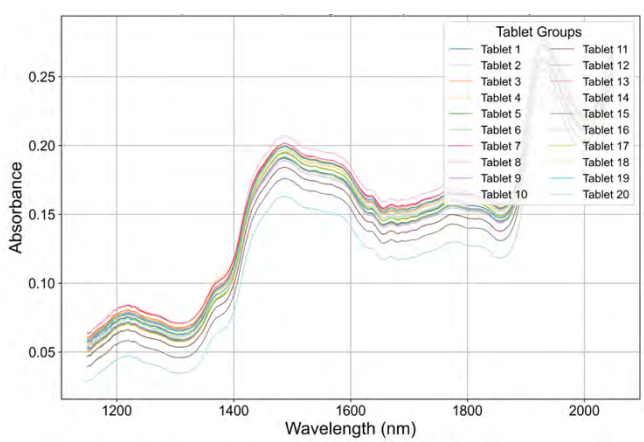

Figure 1. Raw NIR spectra from one tableting conditon.

Model Architecture

The revised Kubelka-Munk (RKM) theory describes effective absorption and scattering coefficients by $K = \mu * \alpha * a$ and $S = \frac{1}{2} * \mu * \alpha * s$, where a and s are the intrinsic absorption and scattering coefficients. α is the diffuse light factor, and $\mu = \sqrt{s/a}$ is the SIPV factor that characterizes the effective photon path length. The reflectance for an infinitely thick sample is $R_\infty = 1 + \left(\frac{K}{S}\right) - \sqrt{(K/S)^2 + 2(K/S)}$. This theory links NIR reflectance to tablet physical properties through the intrinsic optical coefficients.

The RKM theory describes light propagation in pharmaceutical tablets more accurately than the original Kubelka-Munk theory because it explicitly accounts for path variation. When photons scatter through an opaque and heterogeneous media, they do not follow straight paths. Instead, photons go through multiple scattering events, leading to path length distributions that deviate from the classical two-flux assumptions in original KM theory. The SIPV factor μ captures this phenomenon. When scattering dominates ($s > a$), photons travel longer effective distances ($\mu > 1$). This makes the RKM theory ideal for PINN implementations because it provides a mechanistic link between measurable NIR spectra and the underlying tablet microstructure. Therefore, the PINN can learn physically consistent representations of tablet properties rather than purely empirical correlations.

The PINN-RKM architecture we developed consists of multiple sub-networks that simultaneously predict tablet physical properties and optical coefficients from NIR spectra. Physics parameter networks predict porosity, density, and tablet thickness. A composition network estimates component fractions at absorption bands identified through SHapley Additive Explanations (SHAP) analysis of wavelength importance, which revealed key bands around 1150, 1450, 1730, 1940, and 2050 nm corresponding to characteristic chemical bonds in the formulation. Optical coefficient networks predict wavelength-dependent intrinsic coefficients a and s, which are then used to calculate the SIPV factor. A property prediction network maps the predicted physics parameters to target tablet properties.

The total loss function integrates data-driven and physics-informed terms: $\mathcal{L}$ $\mathcal{L}_{total} = \mathcal{L}_{porperty} + w_1 * \mathcal{L}_{spectrum} + w_2 * \mathcal{L}_{physics} + w_1 * \mathcal{L}_{bounds}$, where $\mathcal{L}_{porperty}$ is the mean squared error between predicted and measured tablet properties, $\mathcal{L}_{spectrum}$ quantifies reconstruction error between measured and RKM predicted reflectance spectra, $\mathcal{L}_{physics}$ penalizes violations of physical constraints, and $\mathcal{L}_{bounds}$ enforces physically reasonable parameter ranges. Physical constraints include the porosity-density relationship, SIPV factor, and parameter bounds for porosity, density, intrinsic absorption coefficient, intrinsic scattering coefficient, thickness and light penetration depth. The constraint weights were empirically tuned to balance prediction accuracy and physical consistency.

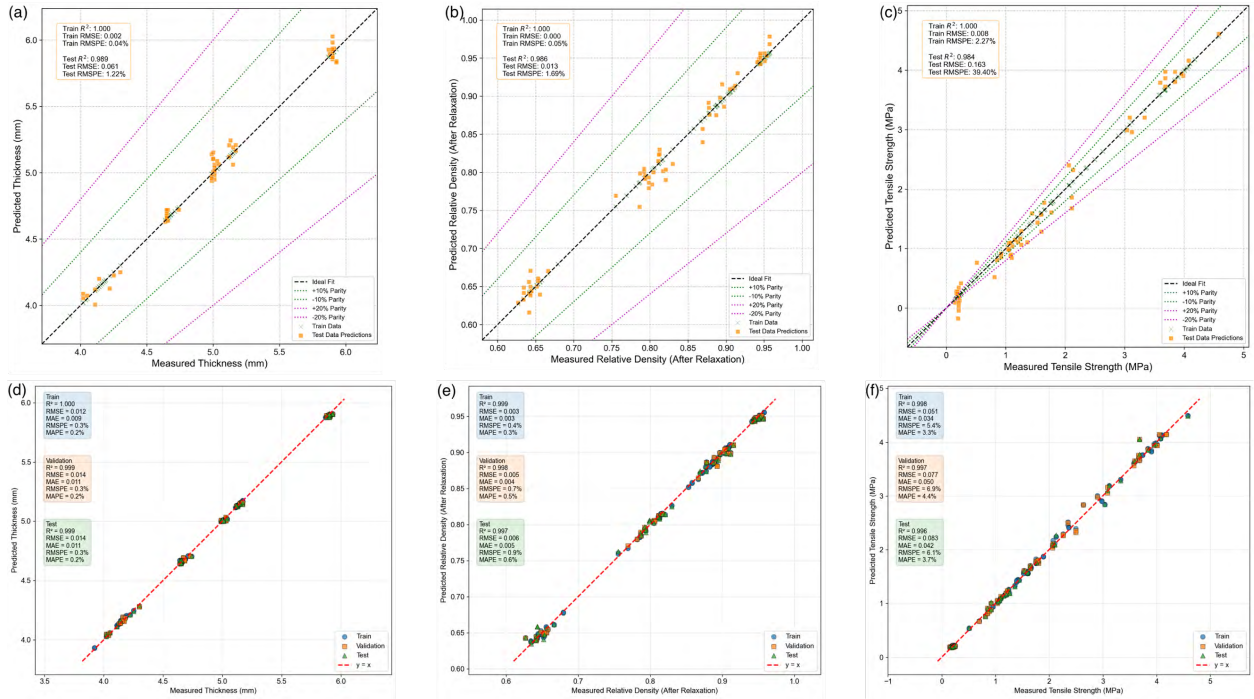

Figure 2. Parity plots of the ANN models (a) tensile strength (b) relative density (c) thickness, and PINN-RKM models (d) tensile strength (e)) relative density (f) thickness.

The network was implemented in PyTorch and trained using the Adam optimizer with learning rate 5×10-4 for 1500 epochs with early stopping (patience = 300 epochs). Data were split 70/30 for training and testing, with standard scaling applied to both NIR spectra and tablet properties. The PINN performance was evaluated against a conventional feedforward ANN with identical network capacity.

RESULTS AND DISCUSSION

Table 1 compares the performance of the proposed a PINN-RKM model with that of a standard ANN in predicting three critical tablet properties from near-infrared (NIR) spectra. The PINN-RKM model demonstrated higher predictive accuracy for all three tablet properties compared to the standard ANN approach. As shown in Figure 2, both models achieved high R^2 values, indicating strong correlations between predictions and reference values. However, the PINN-RKM model has substantially lower prediction errors. PINN-RKM model was able to achieve an improvement in tensile strength predictions, where it reduced the root mean square error (RMSE) by 49% (from 0.163 to 0.083) compared to the standard ANN model. For relative density and thickness predictions, the PINN-RKM model achieved reductions in RMSE of 54% and 77%, respectively. This demonstrates its consistent superiority across properties with different physical characteristics and measurement scales.

Table 1: Modeling results.

Model	TestR^2	RMSE	RMSPE(%)
Tensile Strength			
PINN-RKM	0.996	0.083	6.1
ANN	0.984	0.163	39.4
Relative Density			
PINN-RKM	0.997	0.006	0.9
ANN	0.986	0.013	1.69
Thickness			
PINN-RKM	0.999	0.014	0.3
ANN	0.989	0.061	1.22

Beyond improved accuracy, the PINN-RKM framework provides physically interpretable parameters that offer mechanistic insights into tablet microstructure and light-matter interactions. Figure 3 illustrates the wide range of physical properties and the correlations between the properties and the tablet tensile strength.

As shown in Figure 3, the tablet porosity ranged from 0.06 to 0.21 with a mean of 0.12, while the density distribution peaked at a mean of 1.2. These variations directly influenced the SIPV factor, which showed a mean value of 2.18. A clear positive correlation was observed between porosity and the SIPV factor. As porosity increases, the increased number of air-solid interfaces

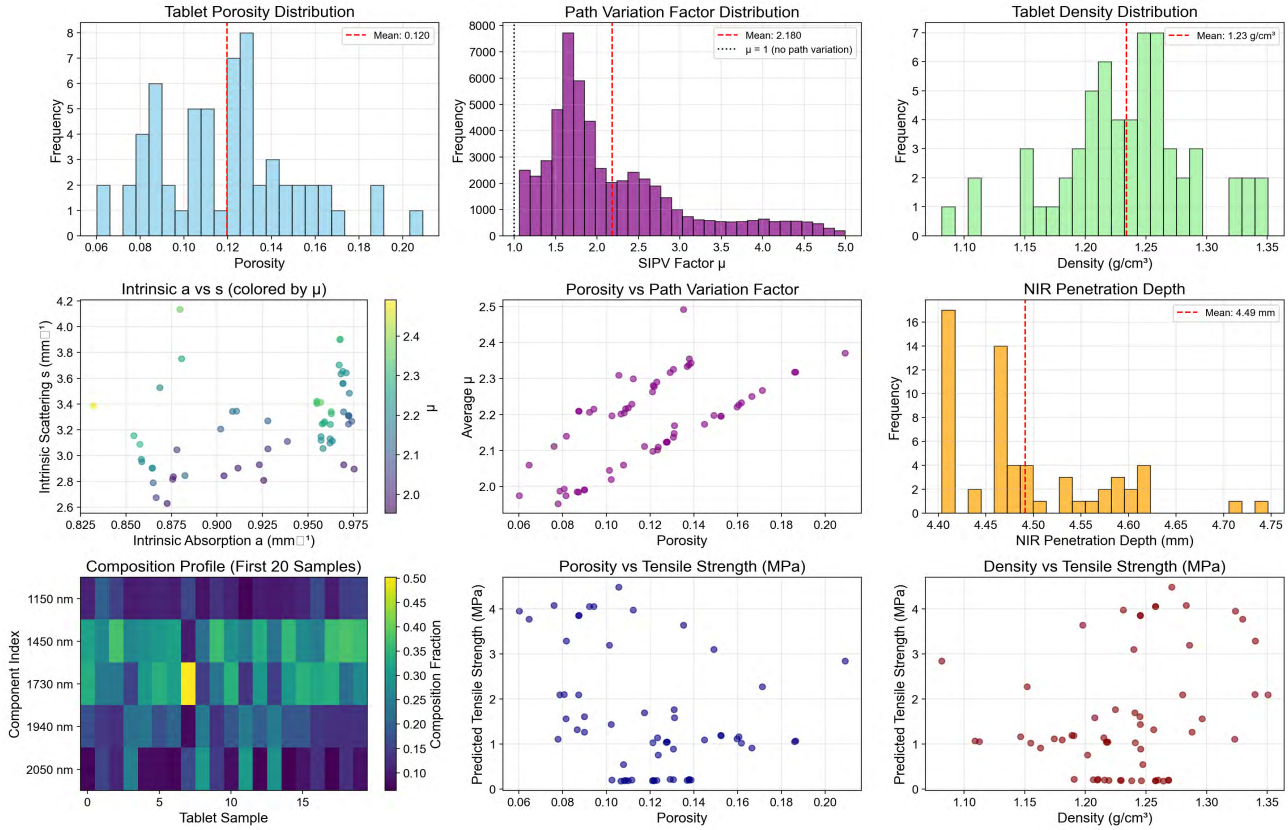

Figure 3. PINN-RKM model mechanistic interpretations.

within the tablet enhances light scattering, thereby extending the total photon pathlength. Furthermore, the NIR penetration depth aligns with the tablet thickness settings, suggesting that the NIR spectra effectively sampled the internal bulk properties of the tablets. However, previous research indicates that the NIR penetration depth in pharmaceutical compacts typically does not exceed 3 mm. Although our model results indicate effective sampling up to the maximum experimental thickness, this discrepancy highlights the need for further study to confirm the accuracy of depth predictions.

CONCLUSION

This research introduces a novel Physics-Informed Neural Network (PINN) architecture integrated with Revised Kubelka-Munk (RKM) theory for the non-destructive characterization of pharmaceutical tablets. By including fundamental light-propagation physics into the deep neural networks, the model successfully transitioned from a traditional empirical correlation to a robust, physics-aware predictive framework. The PINN-RKM model demonstrated superior predictive performance across all evaluated Critical Quality Attributes (CQAs). Specifically, the framework achieved an R2 of 0.999 for thickness and 0.998 for tensile strength, significantly outperforming conventional artificial neural. In summary, the integration of RKM physics provides a significant advancement. This method not only enhances predictive accuracy but also offers a high degree of model interpretability, providing a solid solution for process monitoring and quality assurance in pharmaceutical manufacturing.

REFERENCES

1. Yu LX, Amidon G, Khan MA, Hoag SW, Polli J, Raju GK, Woodcock J. Understanding pharmaceutical quality by design. AAPS J 16:771-783 (2014). https://doi.org/10.1208/s12248-014-9598-3
2. Affairs O of R. Oral Solid Dosage Forms Pre/Post Approval Issues (1/94). FDA 2022.
3. De Beer T, Burggraeve A, Fonteyne M, Saerens L, Remon JP, Vervaet C. Near infrared and raman spectroscopy for the in-process monitoring of pharmaceutical production processes. International Journal of Pharmaceutics 417:32-47 (2011). https://doi.org/10.1016/j.ijpharm.2010.12.012
4. Burns DA, Ciurczak EW. Handbook of near-infrared analysis. 3rd ed. Boca Raton: CRC Press; 2008..
5. Dou Y, Zou T, Liu T, Qu N, Ren Y. Calibration in non-linear NIR spectroscopy using principal

component artificial neural networks. Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy 68:1201-1206 (2007). https://doi.org/10.1016/j.saa.2007.01.021

6. Nagy B, Szabados-Nacsa Á, Fülöp G, Turák Nagyné A, Galata DL, Farkas A, Mészáros LA, Nagy ZK, Marosi G. Interpretable artificial neural networks for retrospective qbd of pharmaceutical tablet manufacturing based on a pilot-scale developmental dataset. International Journal of Pharmaceutics 633:122620 (2023). https://doi.org/10.1016/j.ijpharm.2023.122620
7. Yang L, Kruse B. Revised kubelka–munk theory i theory and application. J. Opt. Soc. Am. A 21:1933 (2004). https://doi.org/10.1364/josaa.21.001933
8. Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics 378:686-707 (2019). https://doi.org/10.1016/j.jcp.2018.10.045

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Probabilistic design spaces from small DoEs – *A boundary-focused workflow using quantile surrogates*

Tobias Overgaard[ab]*, Emmanouil Papadakis[b], Maria-Ona Bertran[b], and Maria M. Papathanasiou[cd]

[a] Department of Applied Mathematics and Computer Science, Technical University of Denmark, 2800 Kongens Lyngby, Denmark
[b] CMC & Product Supply, PS API, Novo Nordisk A/S, 4400 Kalundborg, Denmark
[c] The Sargent Centre for Process Systems Engineering, Imperial College London, London SW7 2AZ, United Kingdom
[d] Department of Chemical Engineering, Imperial College London, London SW7 2AZ, United Kingdom
* Corresponding Author: tobov@dtu.dk.

ABSTRACT

Probabilistic design spaces enable pharmaceutical manufacturers to balance regulatory compliance and operational efficiency. In this context, the "edge-of-failure" separating compliant from non-compliant operation is not a fixed line, but an uncertain region driven by model parameter uncertainty. Traditional methods typically map this probability of failure across the whole design space, which is a computationally expensive task. We propose a more direct approach, reformulating the problem using quantile functions to search for a deterministic boundary at a desired confidence level. These functions are approximated via Gaussian process surrogates using a novel adaptive sampling strategy. An industrial peptide acylation case study identifies the probabilistic design space using 13 experiments versus 18 from a traditional DoE; a 28% reduction. The framework quantifies trade-offs between operational robustness and yield optima.

Keywords: Probabilistic design space, Adaptive sampling, Pharmaceutical manufacturing, Surrogate modeling

INTRODUCTION

In pharmaceutical manufacturing, a design space (DS) is *"the multidimensional combination and interaction of material attributes and process parameters that have been demonstrated to provide assurance of quality"* [1]. Once approved, a DS allows manufacturers to adjust process settings within its bounds without additional regulatory submissions, enabling operational flexibility [2]. A DS can be expressed mathematically as

$$\{\mathbf{x} \in \mathcal{K} \mid \forall \ell \in \{1, \dots, L\}, g_\ell(\mathbf{x}, f(\mathbf{x}, \theta)) \leq 0\}, \quad (1)$$

where $\mathcal{K} \subseteq \mathbb{R}^P$ is the knowledge space, and the process model $\mathbf{y} = f(\mathbf{x}, \theta)$ predicts CQAs $\mathbf{y} \in \mathbb{R}^Q$ with fixed model parameters $\theta \in \mathbb{R}^R$. Each constraint g_ℓ represents a CQA specification limit or a process constraint.

Ignoring model parameter uncertainty yields a deterministic design space (DDS), as in Eq. (1). In practice, parameter uncertainty can reduce or reshape the DS, motivating methods to define a probabilistic design space (PDS). In this context, the model parameters follow a conditional density $p(\theta \mid \mathbf{x})$, which can be inferred via first-order Taylor series approximation [3] or Bayesian parameter estimation [4]. Under model parameter uncertainty, the constraints in Eq. (1) become probabilistic

$$\{\mathbf{x} \in \mathcal{K} \mid \forall \ell \in \{1, \dots, L\}, \mathbb{P}(g_\ell(\mathbf{x}, f(\mathbf{x}, \mathbf{\Theta})) \leq 0) \geq C\}, \quad (2)$$

where $\mathbf{\Theta}$ is a random variable with density $p(\theta \mid \mathbf{x}), C$ is a reliability level, and

$$\mathbb{P}(g_\ell(\mathbf{x}, f(\mathbf{x}, \mathbf{\Theta})) \leq 0) = \int_{\{\theta \mid g_\ell(\mathbf{x}, f(\mathbf{x}, \theta)) \leq 0\}} p(\theta \mid \mathbf{x}) \, \mathrm{d}\theta. \quad (3)$$

The central challenge in PDS identification is the computational cost of evaluating Eq. (3). The objective of this paper is to propose a method that bypasses this cost by reformulating Eq. (2) as a deterministic problem.

Sampling-based approaches to PDS identification

Wet-lab experimentation is a costly and time-consuming task in pharmaceutical development. Therefore, the literature favors model-based methods for mapping the DS [5]. Knowledge-driven models are often preferred over purely empirical ones as they embed physical principles, require less calibration data, and extrapolate more reliably beyond the training domain [6].

Standard PDS identification typically employs a Monte Carlo framework [7]. The knowledge space $\mathcal{K}$ is discretized, and the model $f(\mathbf{x}, \theta)$ is evaluated by drawing samples of θ at each discrete point $\mathbf{x}$. The feasibility probability in Eq. (3) is estimated as the fraction of simulations satisfying all CQA constraints. The PDS is then defined as the set of inputs $\mathbf{x}$ where this probability meets the reliability level C.

While robust and widely applicable, this approach carries a computational cost proportional to the product of grid points and parameter samples, which often proves prohibitive for complex models or large knowledge spaces. To combat this, dimension reduction techniques like partial least squares (PLS) are used to project high-dimensional inputs into a tractable latent space prior to PDS mapping [19]. Computational overhead can also be managed through targeted exploration; nested sampling [8], for instance, accelerates convergence by iteratively discarding low-probability points in favor of candidates near the target threshold C. Furthermore, adaptive surrogate-driven sampling [9] can reduce simulation requirements by learning a surrogate of the feasibility probability and direct new evaluations to the most informative regions (e.g., near the feasibility boundary [10]). Common surrogates include Gaussian process (GP) regression [11] to drive acquisition via uncertainty estimates, or artificial neural networks (ANNs) [12], which have recently been deployed to approximate expensive mechanistic models under material uncertainty [20]. Alternative approaches replace the inner sampling over θ by a fast optimizer inspired by flexibility analysis [13], though it tends to yield more conservative PDS estimates.

After identifying the PDS, a common next step is to inscribe the largest axis-aligned (hyper)rectangle within it - the *acceptable operating region* (AOR) - as this facilitates simple communication about process flexibility to regulatory authorities. However, a fundamental limitation persists across the aforementioned methods: they rely on grid discretization to generate probability maps. If explicit analytical expressions for the boundary constraints existed, AOR identification could be formulated directly as a volume-maximizing problem, akin to flexibility index calculations [13]. Because current sampling-based approaches lack these explicit expressions, AOR identification typically demands complex post-processing via computational geometry [5]. This highlights a critical gap: *the need for a methodology that generates explicit boundary constraints to streamline AOR identification and reveal flexibility-limiting boundary points.*

Quantile-based PDS formulation

To provide practitioners with explicit boundaries, we propose a new formulation of the PDS, inspired by methods from reliability-based design optimization [14, 17]. We utilize the quantile function, which is defined as

$$Q_\ell^{(C)}(\mathbf{x}) = \inf\{q \in \mathbb{R} \mid \mathbb{P}(g_\ell(\mathbf{x}, f(\mathbf{x}, \mathbf{\Theta})) \leq q) \geq C\}, \quad (4)$$

for each $\ell = 1, \ldots, L$. Effectively, $Q_\ell^{(C)}$ acts as the inverse of Eq. (3). Leveraging the properties of $Q_\ell^{(C)}$, we can recast the probabilistic constraints in Eq. (2) into

$$\mathbb{P}(g_\ell(\mathbf{x}, f(\mathbf{x}, \mathbf{\Theta})) \leq 0) \geq C \;\Leftrightarrow\; Q_\ell^{(C)}(\mathbf{x}) \leq 0, \quad (5)$$

for all $0 < C \leq 1$. Using Eq. (5), PDS identification turns into a deterministic problem, where the zero-level set $\{\mathbf{x} \mid Q_\ell^{(C)}(\mathbf{x}) = 0\}$ explicitly defines the "edge-of-failure" boundary at a fixed reliability level C. Unlike raw probabilities, the quantile constraint provides a clearer, more useful signal for targeted experimentation: its sign indicates feasibility, and its magnitude measures the distance to the boundary.

The challenge lies in learning a surrogate for these quantiles based on data from a small DoE. This work addresses three practical questions:

(i) **Surrogate construction:** *How do we build quantile surrogates that account for input and model uncertainty?* We employ GP surrogates trained with a novel three-phase adaptive sampling scheme: (1) sample along the PDS boundary, (2) target intersection points between the boundary and the AOR to identify binding constraints, and (3) sample near KPI optima to balance performance with robustness.

(ii) **Experiment selection**: *Which experiments most effectively reduce uncertainty?* We target the intersection of the AOR and the feasibility boundary to pinpoint edge-of-failure zones.

(iii) **Stopping criteria:** *How many experiments are required?* We monitor the convergence of the size of the AOR to determine a minimum viable dataset.

The remainder of this paper details the methodology, applies it to an industrial API case study, and validates the results against traditional approaches.

METHODOLOGY

This paper introduces a novel methodology to identify PDSs using quantile surrogates, as summarized in Fig. 1. The core innovation is to replace expensive Monte Carlo simulations with a surrogate model that predicts the quantiles of the CQA constraints to account for both process and model uncertainty. To use the methodology, we assume the following two requirements are met:

(a) **Process model:** A mechanistic model f exists, mapping process parameters $\mathbf{x}$ and model parameters θ to the relevant CQAs.

(b) **Uncertainty characterization**: We assume a small initial DoE has already been performed. This data

is used to fit the model f and will result in a conditional probability density $p(\theta \mid \mathbf{x})$ that quantifies the uncertainty in the model parameters.

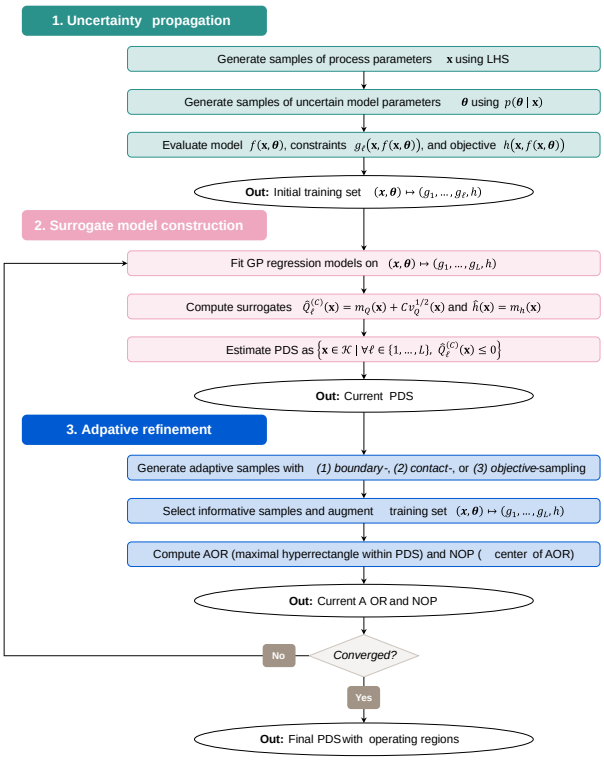

Figure 1. PDS identification methodology.

Step 1: Uncertainty propagation

The first step generates an initial training set to capture how parameter uncertainty propagates to the constraints g_ℓ, as well as to the objective function h, which describes process KPIs like yield or capacity. To obtain analytical expressions of uncertainty, we require Gaussian model parameters. Since this may not be true in general, we assume a transformation T (e.g., Box-Cox) exists such that the transformed model parameters $\mathbf{Z} := T(\boldsymbol{\Theta})$ are distributed as $\mathbf{Z} \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\Sigma})$. This leads to analytical results for the model uncertainty and avoids expensive Monte Carlo sampling in subsequent steps [16].

We generate an initial training set $\mathcal{D}_0$ by sampling I_0 process parameter points using Latin hypercube sampling (LHS) for uniform coverage. At each process point, we draw J_0 model parameter realizations from the conditional density. In practice, we typically find that $I_0 \in [5, 15]$ and $J_0 \in [1, 10]$ are sufficient. The model f is evaluated at these points to compute the constraints and objective. The surrogate fits to the output w_{ij}, which is defined according to the function being modeled:

$$w_{ij} = \begin{cases} g_\ell(\mathbf{x}_i, T^{-1}(\mathbf{z}_{ij})), & \text{constraint model l,} \\ h(\mathbf{x}_i, T^{-1}(\mathbf{z}_{ij})), & \text{objective model.} \end{cases} \tag{6}$$

This yields an initial training set of input-output pairs $\mathcal{D}_0 = \{(\mathbf{v}_{ij}, w_{ij})\}_{1 \le i \le I_0, 1 \le j \le J_0}$, where the input vector $\mathbf{v}_{ij} := [\mathbf{x}_i;\ \mathbf{z}_{ij}]$ concatenates both the deterministic process settings and the transformed model parameters.

Step 2: Surrogate model construction

Given the initial training set $\mathcal{D}_0$, we employ GP regression to model the constraint functions g_ℓ and objective h. GP regression is preferred for its strength in uncertainty quantification [11]. The GP assumes a zero-mean prior and a squared exponential (SE) kernel function with automatic relevance determination (ARD) [15]. This allows the model to learn unique length-scales for each input dimension, effectively identifying the importance of each process and model parameter.

Analytical moment matching: Standard GP prediction assumes deterministic inputs. However, our inputs include uncertain model parameters $\mathbf{Z}$. While mapping a Gaussian input through a nonlinear GP generally yields a non-Gaussian predictive distribution, we approximate the resulting distribution with a Gaussian using moment matching [16]. As a result, we can derive closed-form expressions for the mean $\tilde{m}_w(\mathbf{x}^*)$ and variance $\tilde{s}_w(\mathbf{x}^*)$

$$\begin{aligned} \tilde{m}_w(\mathbf{x}^*) &:= \boldsymbol{\beta}^\top \mathbf{q}(\mathbf{x}^*), \\ \tilde{s}_w(\mathbf{x}^*) &:= \alpha^2 - \operatorname{tr}\Big((\mathbf{K} + \sigma_\varepsilon^2 \mathbf{I})^{-1} \mathbf{Q}(\mathbf{x}^*) \Big) \\ &\quad + \boldsymbol{\beta}^\top \mathbf{Q}(\mathbf{x}^*) \boldsymbol{\beta} - \tilde{m}_w^2(\mathbf{x}^*). \end{aligned} \tag{7}$$

We refer the reader to [16] for the definitions of $\boldsymbol{\beta}, \mathbf{q}$ and $\mathbf{Q}$ and the full derivation of Eq. (7). The other parameters include the signal variance α^2 of the SE kernel and the observation noise σ_ε^2. The covariance matrix $\mathbf{K}$ is computed by evaluating the SE kernel across all input pairs.

Quantile surrogate construction: Using these analytical moments, we construct the C-level quantile surrogate. For a required reliability level C (e.g., 95%), the surrogate is defined as the upper confidence bound of the predictive distribution

$$\hat{Q}_\ell^{(C)}(\mathbf{x}) = \tilde{m}_w(\mathbf{x}) + \Phi^{-1}(C) \sqrt{\tilde{s}_w(\mathbf{x})}, \tag{8}$$

where Φ^{-1} is the quantile function of the Gaussian distribution. The PDS can then be directly approximated as

$$\left\{ \mathbf{x} \in \mathcal{K} \mid \forall \ell \in \{1, \dots, L\}, \hat{Q}_\ell^{(C)}(\mathbf{x}) \le 0 \right\}, \tag{9}$$

thereby accounting for both input and model uncertainty. Unlike Eq. (8), the objective surrogate is expressed as $\hat{h}(\mathbf{x}) = \tilde{m}_w(\mathbf{x})$, as it approximates only the expected value rather than a quantile. At this stage, all surrogates are built on $\mathcal{D}_0$, and the next step is to refine them near AOR-critical regions and the process optimum.

Step 3: Adaptive refinement

Surrogates based on space-filling designs often lack accuracy at the specific boundaries that define

feasibility. We define a three-phase adaptive sampling strategy to combat this. At each iteration, the algorithm selects a new design point $\mathbf{x}^*$ and parameter realization $\mathbf{z}^*$ to evaluate, refits the surrogates, and updates the PDS.

- **Phase 1: Boundary-guided sampling:** Early iterations focus on sharpening the PDS boundaries. We search for process settings where the constraints are closest to their critical values (the zero-level set). This amounts to minimizing the worst-case constraint violation, i.e.., $\max_{\ell} |\hat{Q}_{\ell}^{(C)}(\mathbf{x})|$. To enable gradient-guided optimization, we build a smooth proxy of the worst-case constraint violation, given by

$$r(\mathbf{x}) = \frac{1}{\eta}\log\Big(\sum_{\ell=1}^{L} \exp\big(\eta\, \hat{Q}_{\ell}^{(C)}(\mathbf{x})\big)\Big), \tag{10}$$

  where $\eta > 0$ is a sharpness parameter. As $\eta \to \infty, r(\mathbf{x})$ approaches $\max_{\ell} \hat{Q}_{\ell}^{(C)}(\mathbf{x})$. We optimize

$$\mathbf{x}^* = \underset{\mathbf{x}\in\mathcal{K}}{\operatorname{argmin}}\, r^2(\mathbf{x}), \tag{11}$$

  to target regions where $r(\mathbf{x}) \approx 0$, which correspond to the zero-level set of the worst-case constraint.
- **Phase 2: Contact-guided sampling**: Once the quantile boundary is defined, we focus on the AOR - the largest hyperrectangle inscribed within the PDS. The accuracy of the AOR depends on the so-called *contact points*, which we define as the points where the rectangle intersects the PDS boundary. We first solve an optimization problem to maximize the AOR volume. We then sample the edge points of this hyperrectangle that are predicted to lie on the quantile boundary. This refines the "edge-of-failure" regions that ultimately limit the size of the operating window.
- **Phase 3: Objective-guided sampling**: Finally, the strategy aims at optimizing process performance. We maximize the objective surrogate subject to the quantile constraints. This ensures the PDS is accurate in the region of optimal operation, verifying that the identified optimum is truly robust to uncertainty.

The Next Iteration: Once the best process point $\mathbf{x}^*$ is identified (from any of the three phases), we must select which parameter realization $\mathbf{z}^*$ to simulate. Simply sampling the mean is insufficient. Instead, we draw candidates from the parameter distribution $\mathcal{N}(\boldsymbol{\mu}(\mathbf{x}^*), \boldsymbol{\Sigma})$ and select the one that minimizes the *boundary uncertainty*. Mathematically, we identify active constraint $\ell^* \in \{1, \dots, L\}$ (defined as the one nearest to $\mathbf{x}^*$) and compute

$$\mathbf{z}^* = \underset{\mathbf{z}}{\operatorname{argmin}} \frac{|m_{\ell^*}([\mathbf{x}^*;\mathbf{z}])|}{\sqrt{s_{\ell^*}([\mathbf{x}^*;\mathbf{z}])}}, \tag{12}$$

where m_{ℓ^*} and s_{ℓ^*} are the GP posterior mean and variance for constraint ℓ^* [15]. This criterion selects parameters where the constraint prediction is near the boundary (small $|m_{\ell^*}|$) yet highly uncertain and poorly informed (large s_{ℓ^*}) [17]. If multiple parameter samples are desired (as in Step 1), we select those close to the minimum.

The new simulation is added to the training set

$$\mathcal{D}_i \leftarrow \mathcal{D}_{i-1} \cup \begin{cases} \{\mathbf{v}^*, g_{\ell}(\mathbf{x}^*, T^{-1}(\mathbf{z}^*))\}, & \textit{constraint l,} \\ \{\mathbf{v}^*, h(\mathbf{x}^*, T^{-1}(\mathbf{z}^*))\}, & \textit{objective.} \end{cases} \tag{13}$$

and all surrogates are refitted. Convergence is assessed differently for each phase: An error-based stopping criterion (ESC) is used for boundary sampling (see [18] for details), AOR size stability is used for contact sampling, and optimizer convergence is used for the objective phase. When one phase converges, we either advance to the next or analyze the final PDS. If not, we can add more digital samples or update with new experimental data.

CASE STUDY

In this paper, we consider an acylation reaction step, which is a critical synthetic step in the production of a long-acting peptide-based API. The process attaches a hydrophobic side chain to the peptide intermediate to form the final API. Several competing reactions must be balanced to achieve high product purity and yield while minimizing the formation of impurities.

Reaction scheme and knowledge-driven model

The mechanistic model comprises seven coupled ordinary differential equations tracking the concentrations of the species: SM (starting material), P (desired product), A (impurity from undesired acylation), HI (hydrophobic impurities from di-acylation), HSC (hydrolyzed side chain reagent), SC (side chain reagent), and P^* (degraded product). The reaction scheme is:

$\mathrm{SM} + \mathrm{SC} \xrightarrow{k_1} \mathrm{P}$ (desired acylation at primary site)

$\mathrm{SM} + \mathrm{SC} \xrightarrow{k_2} \mathrm{A}$ (undesired acylation at secondary site)

$\mathrm{P} + \mathrm{SC} \xrightarrow{k_2} \mathrm{HI}$ (di-acylation at secondary site)

$\mathrm{A} + \mathrm{SC} \xrightarrow{k_1} \mathrm{HI}$ (di-acylation at primary site)

$\mathrm{SC} \xrightarrow{k_3} \mathrm{HSC}$ (hydrolysis of side chain reagent)

$\mathrm{P} \xrightarrow{k_4} \mathrm{P}^*$ (product degradation)

We assume that reactions at the same site happen at identical rates regardless of the other site's state, so only two acylation rates are required. We also assume that both rates depend on pH and temperature, $k_i = k_i(\mathrm{pH}, T)$ for $i = 1, 2$. The hydrolysis and degradation rates follow $k_i = k_i(T)10^{\mathrm{pH}-14}$ for $i = 3, 4$, reflecting their hydroxide-catalyzed mechanisms. All rate parameters are log-transformed and estimated using least-squares regression on the 18 available LHS experiments

$$
\begin{aligned}
\log(k_1(\mathrm{pH}, T)) &= \beta_{1,1} + \beta_{1,2}/T + \beta_{1,3}\mathrm{pH} + \beta_{1,4}\mathrm{pH}^2 + \varepsilon_1, \\
\log(k_2(\mathrm{pH}, T)) &= \beta_{2,1} + \beta_{2,2}/T + \beta_{2,3}\mathrm{pH} + \beta_{2,4}\mathrm{pH}^2 + \varepsilon_2, \\
\log(k_3(T)) &= \beta_{3,1} + \beta_{3,2}/T + \varepsilon_3, \\
\log(k_4(T)) &= \beta_{4,1} + \beta_{4,2}/T + \varepsilon_4,
\end{aligned} \tag{14}
$$

where the βs denote regression coefficients, and the εs are zero-mean Gaussian errors. The $1/T$ term captures Arrhenius temperature dependence, and log-transformation ensures approximate Gaussianity (as required by Step 1 of the methodology), with means parameterized by pH and T. As an example, take $\mathrm{pH} = 11.5, T = 2.50°\mathrm{C}$. Then $\log([k_1, k_2, k_3, k_4]^\top) \sim \mathcal{N}(\boldsymbol{\mu}(11.5, 2.50), \boldsymbol{\Sigma})$ with

$$
\begin{aligned}
\boldsymbol{\mu}(11.5, 2.50) &= [-4.26, -8.11, 1.11, -0.950]^\top, \\
\boldsymbol{\Sigma} &= 10^{-3}\,\mathrm{diag}(5.96, 0.856, 915, 269).
\end{aligned} \tag{15}
$$

For brevity (and model confidentiality), we do not report the full distributions across all (pH, T). By construction, the rates are modeled as independent random variables (no covariance entries in $\boldsymbol{\Sigma}$).

RESULTS AND DISCUSSION

We demonstrate the methodology on the peptide acylation case study. Based on the reaction scheme, a standard kinetic model was formulated as a system of ordinary differential equations. With this, Requirement (a) of the methodology (**process model**) is fulfilled. Furthermore, the fitted rate parameters in Eq. (14) follow a Gaussian distribution, thereby meeting Requirement (b) (**uncertainty characterization**).

Design space characterization

We characterize the PDS by two CQAs - hydrophobic impurities $\leq 2\%$ (constraint 1) and purity $\geq 70\%$ (constraint 2) over $\mathrm{pH} \in [10.5, 12.5]$ and $T \in [-3, 10]°\mathrm{C}$. We set $\mathbf{x} := [\mathrm{pH}, T]^\top$ and $\mathbf{z} := \log([k_1, k_2, k_3, k_4]^\top)$ and calibrate with 13 experiments (this choice is justified later). The reliability level is set at $C = 0.975$. For each sample of $\mathbf{x}$, we draw five $\mathbf{Z} \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\Sigma})$. The surrogate is initialized with 10 LHS points in $\mathbf{x}$, then refined adaptively.

Fig. 2 shows the PDS evolution at 15/30/40 iterations using the three-phase sampling strategy. At iteration 15 (Fig. 2**(a)**), boundary-guided samples (black ×) cluster on $\hat{Q}_1^{(C)} = 0$ and $\hat{Q}_2^{(C)} = 0$ (solid lines), quickly capturing the quantile boundary. Separation from the mean contours (dashed lines) reflect model uncertainty.

By iteration 30 (Fig. 2**(b)**), the AOR (black □) is identified, and contact-focused samples (orange ×) intersect the contours. The PDS geometry changes only slightly, with the mean purity contour shifting toward the quantile, indicating reduced uncertainty; the AOR's upper-left corner remains unsampled because it is unconstrained. The nominal operating point (NOP) lies around $\mathrm{pH} \approx 11.8$ and $T \approx 1.6°\mathrm{C}$ (black ★).

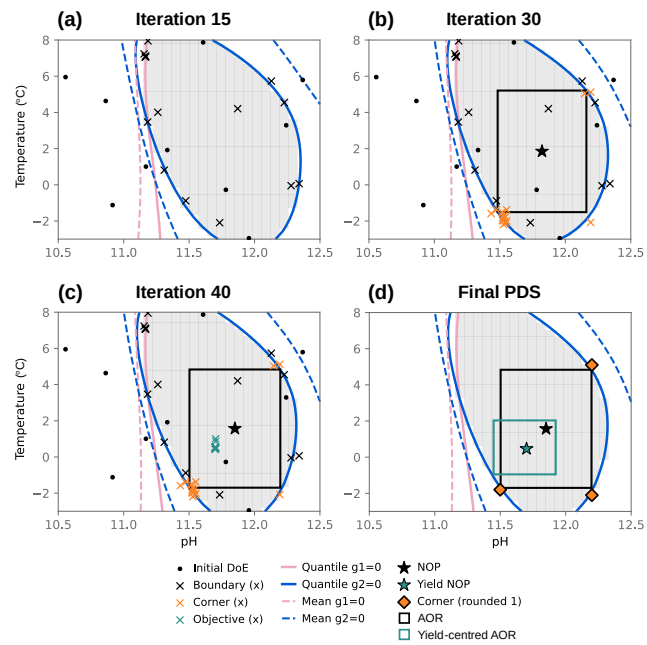

Figure 2. Snapshots of the PDS identification scheme.

At iteration 40 (Fig. 2**(c)**), we introduce a yield objective. Objective-guided samples (green ×) cluster near $\mathrm{pH} \approx 11.7$ and $T \approx 0.5°\mathrm{C}$, producing a yield-centered NOP (green ★). The AOR has shifted slightly to the right, possibly reflecting selection among equally sized rectangles.

The resulting PDS (Fig. 2**(d)**) shows two sets of NOPs and AORs: the yield-centered AOR (green □) has area ≈ 1.8 versus ≈ 4.6 for the maximal-flexibility AOR (black □), quantifying the robustness–performance trade-off. Maximizing yield narrows the window and raises violation risk, yet the flexibility-oriented AOR still contains the yield optimum, enabling optimal operation within robust ranges. Three edge-of-failure points (orange ♦) are found as local averages of the corner-point clusters. Using these as suggestions for follow-up experiments will most efficiently sharpen the boundary between feasible and infeasible operation.

Convergence analysis and experimental strategy

Fig. 3 shows convergence of the three-phase adaptive sampling. The error-based stopping criterion (ESC) in Fig. 3**(a)** converges within ~10 boundary-guided iterations: the purity error (e_2) drops from ~18% to ~7.5% and remains below 10%, while the impurity constraint (e_1) is reliably identified within five iterations. We extend to 15 iterations to confirm stability of the estimates.

The NOP stabilizes after 15 iterations of contact-guided sampling (30 iterations in total). Another 10 iterations of objective-guided sampling shows that the NOP remains locked at $\mathrm{pH} \approx 11.8$ and $T \approx 1.6°\mathrm{C}$ (Fig. 3**(b)**). Overall, we ran 40 iterations (five parameter draws each)

plus 10 LHS burn in runs, totaling 210 evaluations of f.

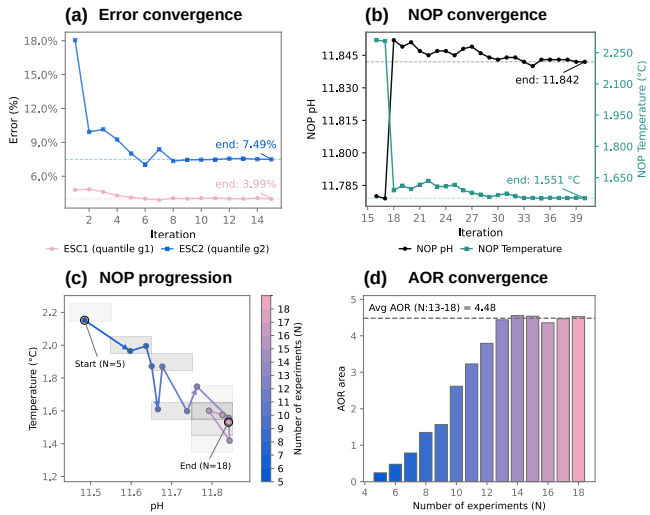

Figure 3. Convergence analysis.

Fig. 3**(c)** uses the edge-of-failure points (orange ♦) to guide laboratory experimentation. We start with $N = 5$ experiments (the minimum to estimate Eq. (14)) and iteratively add new experiments by choosing (pH, T) settings near these points. As N increases, the NOP converges; a color gradient shows experiment order, and the gray background partitions points into bins by rounding them to one decimal place; darker shading indicates a higher concentration of points. By $N = 10$, the NOP is near its final location, with most runs at $pH \approx 11.8$ and $T \approx 1.6°\mathrm{C}$.

Fig. 3**(d)** confirms that 13 experiments suffice to estimate the PDS: tracking the AOR area for $N \in [5, 18]$ shows stabilization at $N = 13$, with an average area of ~ 4.48; additional runs will not expand the feasible region.

Validation and practical considerations

We validated the PDS against a DDS derived from standard Monte Carlo [7]. Rather than sampling across an uncertain parameter space, we fix the log-transformed rates at their expected values $\boldsymbol{\mu}(\mathbf{x})$ and compute $\max(g_1, g_2)$ on a 250×250 grid (totaling 62500 evaluations of f). We report that the PDS mean boundary closely aligns with the DDS boundary, while the PDS quantile boundary is strictly enclosed by the DDS, ensuring a conservative buffer against parameter uncertainty.

Implementation alongside traditional workflows: Traditional (large factorial) DoEs lack adaptivity and expend runs in infeasible or low-information regions; our approach concentrates experiments at feasibility boundaries to most effectively reduce uncertainty and define the operating window.

In process development, a small confirmation study is often included between process characterization and validation to address residual gaps; the proposed methodology in this paper can guide this stage. Furthermore, it naturally accommodates incremental updates with new data from scale-up and commercial manufacturing, making it well-suited for lifecycle management.

Batch-wise experimentation: As shown in Fig. 3, targeting edge-of-failure points reduced experiments from 18 to 13; an important saving given costly intermediates/reagents and reaction times of up to 17 hours. In practice, single-run experimentation is often infeasible, as samples are sent for analysis at central facilities. This can delay results for hours or days. Batch-sequential design can address this by bundling the experiments, with adaptive refinement between batches.

Model extensions: CQA data may exhibit significant analytical variation. To account for this uncertainty, implementing stochastic differential equations (rather than ordinary differential equations) could provide a more robust analysis by explicitly modeling measurement noise alongside process dynamics.

CONCLUSION

We present a quantile-based, adaptive methodology for probabilistic design space identification that efficiently characterizes feasible operating regions under parameter uncertainty; it yields explicit boundary constraints, identifies edge-of-failure points for targeted experimentation, and provides convergence criteria for reliable characterization.

Applied to industrial API manufacturing, the methodology identified the probabilistic design space in 13 rather than 18 experiments (28% reduction), saving costly intermediates and time. It also computed the largest inscribed operating window at the most flexible point (AOR area ≈ 4.5) and at the yield optimum (AOR area ≈ 1.8), capturing the robustness–performance trade-off.

Future extensions include: (i) strategies for cases where the process optimum lies on the feasibility boundary (vanishing flexibility); and (ii) multi objective formulations using convex combinations of yield, capacity, and unit cost to map trade-offs and define operating windows spanning multiple modes. A mode-switching policy could alternate between these optima as conditions change, enabling real-time optimization while maintaining compliance. In summary, quantile-based design space identification provides a fast route to process characterization that balances product quality and process efficiency.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge funding from Novo Nordisk A/S, PS API manufacturing PhD program.

REFERENCES

1. FDA. Guidance for Industry: Q8(R2) Pharmaceutical Development. 2009.
2. Facco P, Dal Pastro F, Meneghetti N, Bezzo F, Barolo M. Bracketing the design space within the knowledge space in pharmaceutical product development. Ind. Eng. Chem. Res. 54:5128-5138 (2015). https://doi.org/10.1021/acs.iecr.5b00863
3. Franceschini G, Macchietto S. Model-based design of experiments for parameter precision: state of the art. Chemical Engineering Science 63:4846-4872 (2008). https://doi.org/10.1016/j.ces.2007.11.034
4. Coleman MC, Block DE. Bayesian parameter estimation with informative priors for nonlinear systems. AIChE Journal 52:651-667 (2005). https://doi.org/10.1002/aic.10667
5. Sachio S, Kontoravdi C, Papathanasiou MM. A model-based approach towards accelerated process development: a case study on chromatography. Chemical Engineering Research and Design 197:800-820 (2023). https://doi.org/10.1016/j.cherd.2023.08.016
6. Moshiritabrizi I, McMullen JP, Wyvratt BM, McAuley KB. A comparative study of strategies for incorporating uncertainty in design space determination for pharmaceutical manufacturing. Ind. Eng. Chem. Res. 64:21658-21668 (2025). https://doi.org/10.1021/acs.iecr.5c03037
7. García-Muñoz S, Luciani CV, Vaidyaraman S, Seibert KD. Definition of design spaces using mechanistic models and geometric projections of probability maps. Org. Process Res. Dev. 19:1012-1023 (2015). https://doi.org/10.1021/acs.oprd.5b00158
8. Kusumo KP, Gomoescu L, Paulen R, García Muñoz S, Pantelides CC, Shah N, Chachuat B. Bayesian approach to probabilistic design space characterization: a nested sampling strategy. Ind. Eng. Chem. Res. 59:2396-2408 (2019). https://doi.org/10.1021/acs.iecr.9b05006
9. Kucherenko S, Giamalakis D, Shah N, García-Muñoz S. Computationally efficient identification of probabilistic design spaces through application of metamodeling and adaptive sampling. Comput. Chem. Eng. 2020;132:106608. doi:10.1016/j.compchemeng.2019.106608Geremia
10. Geremia M, Bezzo F, Ierapetritou MG. A novel framework for the identification of complex feasible space. Computers & Chemical Engineering 179:108427 (2023). https://doi.org/10.1016/j.compchemeng.2023.108427
11. Ding C, Ierapetritou M. A novel framework of surrogate-based feasibility analysis for establishing design space of twin-column continuous chromatography. International Journal of Pharmaceutics 609:121161 (2021). https://doi.org/10.1016/j.ijpharm.2021.121161
12. Metta N, Ramachandran R, Ierapetritou M. A novel adaptive sampling based methodology for feasible region identification of compute intensive models using artificial neural network. AIChE Journal 67: (2020). https://doi.org/10.1002/aic.17095
13. Laky D, Xu S, Rodriguez JS, Vaidyaraman S, García Muñoz S, Laird C. An optimization-based framework to define the probabilistic design space of pharmaceutical processes with model uncertainty. Processes 7:96 (2019). https://doi.org/10.3390/pr7020096
14. Moustapha M, Sudret B, Bourinet JM, Guillaume B. Quantile-based optimization under uncertainties using adaptive kriging surrogate models. Struct Multidisc Optim 54:1403-1421 (2016). https://doi.org/10.1007/s00158-016-1504-4
15. Rasmussen CE, Williams CKI. Gaussian processes for machine learning. The MIT Press (2005). https://doi.org/10.7551/mitpress/3206.001.0001
16. Deisenroth MP. Efficient Reinforcement Learning using Gaussian Processes. PhD thesis, KIT; 2010.
17. Kim J, Song J. Quantile surrogates and sensitivity by adaptive gaussian process for efficient reliability-based design optimization. Mechanical Systems and Signal Processing 161:107962 (2021). https://doi.org/10.1016/j.ymssp.2021.107962
18. Wang Z, Shafieezadeh A. ESC: an efficient error-based stopping criterion for kriging-based reliability analysis methods. Struct Multidisc Optim 59:1621-1637 (2018). https://doi.org/10.1007/s00158-018-2150-9
19. Bano G, Facco P, Bezzo F, Barolo M. Probabilistic design space determination in pharmaceutical product development: a bayesian/latent variable approach. AIChE Journal 64:2438-2449 (2018). https://doi.org/10.1002/aic.16133
20. Meng Q, Bogle D, Charitopoulos VM. Probabilistic design space exploration and optimization via bayesian approach for a fluid bed drying process. European Journal of Pharmaceutical Sciences 210:107116 (2025). https://doi.org/10.1016/j.ejps.2025.107116

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Combined PBM-PBPK Modeling for Optimized Integrated Oral Solid Dosage Form and Dosing Strategy Design

Meng-Hua Yang[a], Francesco Rossi[a], and Gintaras V. Reklaitis[a], Zoltan K. Nagy[s]*

[a] Davidson School of Chemical Engineering, Purdue University, West Lafayette, IN, USA
* Corresponding Author: znagy@purdue.edu.

ABSTRACT

The formulation of oral solid dosage forms can have a significant impact on drug bioavailability, particularly for poorly soluble drugs. However, traditional formulation development relies heavily on extensive experimental testing, which limits its efficiency and effectiveness in oral drug product design. In this study, we present an integrated framework to support rational formulation design and exploration of optimal dosage regimens. This framework combines population balance-based tablet disintegration and dissolution modeling with physiologically based pharmacokinetic (PBPK) modeling to link critical material attributes (CMAs) with the pharmacokinetic response. The anti-coagulant drug rivaroxaban is selected as a model compound for calibration and deployment of the framework, enabling systematic investigation of the effects of crystal size distribution (CSD) and tablet porosity on in vivo performance. The results demonstrate that CSD has a pronounced impact on in pharmacokinetics, whereas tablet porosity exhibits a smaller but non-negligible effect. Furthermore, optimization is implemented to identify the optimal dose amount under a given formulation for producing the desired pharmacokinetic profile for average patient group, demonstrating the potential of this framework for the digital design of both drug efficacy and treatment strategies.

Keywords: PBM, PBPK, Optimal Dose regimen

1. INTRODUCTION

The performance of oral solid dosage forms can be critically influenced by formulation attributes such as particle size distribution (PSD) of active pharmaceutical ingredient (API), solubility, excipient type, and porosity of dosage form. These properties collectively governed dissolution behavior and, in turn, oral drug absorption. Although formulation effects on dissolution have been widely studied, formulation development still relies largely on traditional trial-and-error experimentation, resulting in substantial resource demand [1]. The challenge arises from the complex, combined effects of CMAs, and the gap between *in vitro* dissolution and in *vivo* performance [2].

Model-informed formulation development has therefore emerged as an important strategy in order to bridge formulation properties and *in vitro* dissolution with *in vivo* performance. By integrating experimental data with computational approaches, this method provides us an effective tool for understanding properties impact. *In vitro-in vivo* correlation (IVIVC) is one of the most widely used approaches. It involves establishing empirical or semi-empirical relationships between dissolution profiles and absorption kinetics, which is inferred from plasma drug concentration data [3]. Although IVIVC can be useful in certain formulation development contexts, it depends on great amounts of paired in vitro and in vivo data to build robust correlation. It is also typically specific to a given drug, and offers limited mechanistic interpretability.

To address the limitations of IVIVC, hybrid approaches that combine dissolution modeling or IVIVC concepts with physiologically based pharmacokinetic (PBPK) modeling have been introduced. PBPK modeling is a mechanistic modeling approach that predicts drug concentration-time profiles. It is based on mass balance equations defined over structured compartments, which incorporate relevant physiological process and drug-specific physicochemical parameters. Several studies

have integrated PSD–dependent dissolution model with PBPK model. For example, Stefan et al. coupled a Noyes–Whitney–based dissolution model with a PBPK framework to study PSD effects on cilostazol absorption [4]. Pepin et al. incorporated the API PSD into the commercial PBPK software, and demonstrated its application for regulatory support [5]. Although these studies successfully link PSD to in *vivo* performance, there is still a lack of systematic frameworks that can evaluate the combined effect of multiple CMAs on pharmacokinetic response.

This work aims to present a framework that integrates predictive modeling with optimization methods to support both formulation design and the development of dose regimen. A PBM-based dissolution model coupled with a PBPK model is developed. Using this combined model, a preliminary investigation is conducted to evaluate the influence of CSD and tablet porosity on plasma drug concentration for rivaroxaban. Additionally, the dose is optimized to achieve a target pharmacokinetic profile.

2. PROPOSED METHODOLOGY

2.1 The digital design framework

A framework is proposed for the integrated digital design of pharmaceutical formulations and dosing regimen, as illustrated in Figure 1. Within this digital design framework, a hybrid model is developed to predict systemic pharmacokinetics from formulation properties, while optimization approaches can be applied to explore optimal formulation design variables and dosing strategies, either independently or simultaneously.

The hybrid model is developed using a two-stage, sequential integrated strategy. In the first stage, a PBM-based tablet disintegration and dissolution model is independently employed to simulate the dissolution behavior of an immediate-release tablet. This model accounts for the impact of CMAs, including tablet porosity and API PSD. Parameters set a priori represent stomach-specific physiological conditions such as solubility, while adaptive parameters are calibrated using in vitro dissolution data under gastric-relevant conditions. Upon completion of tablet disintegration, the outputs from the first stage, including the dissolved API concentration and PSD, are used as initial conditions for the second stage.

In the second stage, the PBM-based dissolution model is embedded within a PBPK framework to describe the dissolution of remaining API particles in the gastrointestinal (GI) lumens. The PBM models predicts time-dependent dissolved API concentrations which act as state variables governing the absorption rate in the PBPK model. The PBPK model accounts for patient physiology and whole-body absorption, distribution, metabolism, and excretion (ADME). Physiological parameters are obtained from literature or databases, while a subset of adaptive parameters is calibrated using in vivo clinical plasma concentration data.

This two-stage strategy is based on the assumptions that (i) tablet disintegration predominantly occurs in the stomach, and (ii) GI transit after the disintegration is complete. These assumptions are based on the consideration that disintegration occurs rapidly relative to gastrointestinal transit and absorption timescale, thereby, treated as a fast initial process.

2.2 Stage I: The PBM-based tablet dissolution and disintegration model

The PBM-based tablet disintegration and dissolution model, adopted from a previous study [6], describes the temporal evolution of crystal populations in different status. It is formulated using two population balance equations (PBEs): one for the number density of bound crystals within the tablet, $n_{CB}(t,L)$, and one for released crystals $n_{CR}(t,L)$, both as functions of time t and particle size L:

$$\frac{\partial n_{CB}(t,L)}{\partial t} = -S_1(t) n_{CB}(t,L) \tag{1}$$

$$\frac{\partial n_{CR}(t,L)}{\partial t} = -\frac{\partial (n_{CR}(t,L) D_c(t))}{\partial L} + S_1(t) n_{CB}(t,L) \tag{2}$$

Here, $S_1(t)$ denotes the tablet disintegration rate, which is described using the Washburn equation. In Eq. (3), θ represents the contact angle, γ is surface tension, and δ denotes to viscosity of liquid. The effective pore length, L_{eff}, is estimated from cubit root of its total pore volume. Tablet porosity (ε) is correlated with effective capillary radius (r_e) through empirical relationship.

$$S_1(t) = \frac{1}{L_{eff}} \sqrt{\frac{r_e(\varepsilon)\gamma \cos\theta}{2\delta}} \cdot \frac{1}{2\sqrt{t}} \tag{3}$$

$D_c(t)$ represents the Noyes-Whitney based dissolution ra, in which k_c is the dissolution rate constant, C_{API} is the concentration of API, and C_s is solubility:

$$D_c(t) = -k_c\left(1 - \frac{C_{API}}{C_s}\right) \tag{4}$$

The concentration of API dissolved in liquid can be calculated as follows:

$$\frac{dC_{API}}{dt} = -3\rho_c k_v \int_0^\infty L^2 n_{CR}(t,L) D_c(t)\, dL \tag{5}$$

where k_v corresponds to crystal shape factor and ρ_c represents crystal density.

2.3 Stage II: The coupled PBM-PBPK model

The PBPK model represents the human body using multiple anatomically relevant compartments, including: brain, lung, adrenal, thyroid, heart, liver, pancreas, spleen, kidney, skin, fat, skeleton, and the gastrointestinal (GI) tract, which are interconnected through blood flow or gastrointestinal transit. The bronchial is grouped

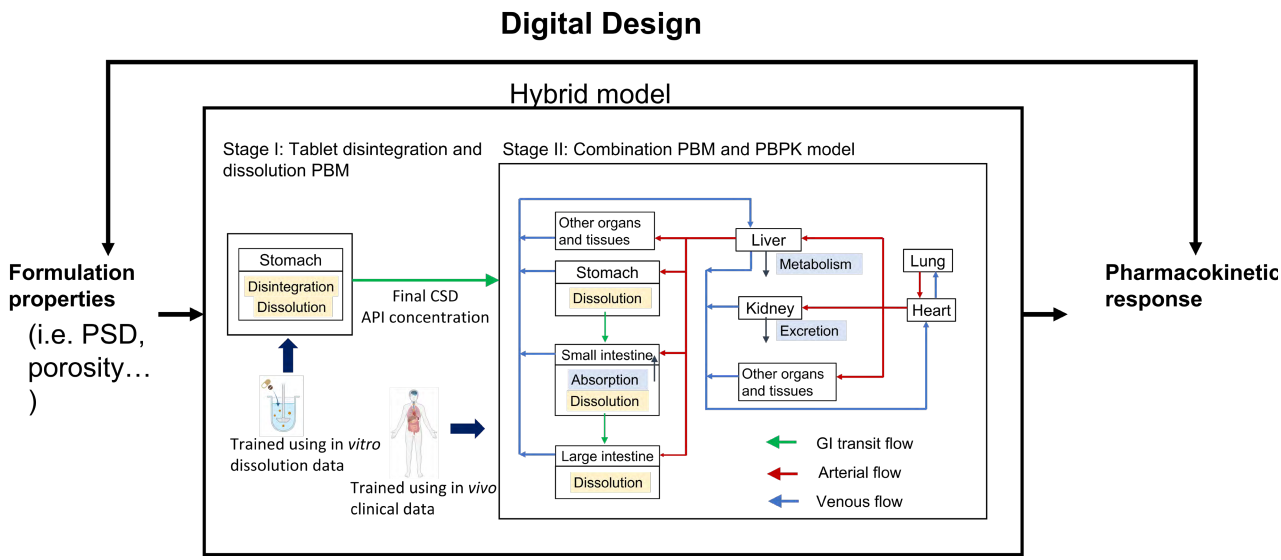

Figure 1. Digital design framework enabling the integrated optimization of drug formulation and dosing regimen design.

within the lung compartment; coronary and heart are unified due to their similar composition. The remaining organs and tissues such as bladder, gonads are lumped into a single compartment because of the low perfusion index (i.e. their minimal blood flow).

Each compartment is divided into three components: plasma, blood cell, and tissue sub-compartments. Drug exchange within a compartment occurs between plasma and blood cells and between plasma and tissue. Under the assumption that GI dissolution and absorption is the rate determining step, intra-compartment drug distribution is assumed to reach rapid equilibrium, thus characterized using partition coefficient K_{eq}^{B} (blood cell-to-plasma) and K_{eq}^{T} (tissue-to-plasma). Only the unbound fraction of drug can across cell membranes freely, thereby participating in the process of intra-compartment transport, metabolism and excretion. Consequently, the unbound plasma drug concentration in compartment i, $C_{u,i}(t)$, is selected as the primary state variable in the PBPK model. The corresponding total plasma drug concentration can be calculated from $C_{u,i}(t)$ using Eq. (6), where f_u is the constant unbound fraction.

$$C_{u,i}(t) = f_u C_i(t) \tag{6}$$

Based on mass-balance principles, the temporal evolution of the unbound plasma drugs in compartment i is governed by convective blood flow transport and can be expressed as:

$$\frac{dC_{u,i}(t)}{dt} = F_{in,i}C_{u,i+1}\varphi_{i+1} - F_{out,i}C_{u,i}\varphi_i \tag{7}$$

Here, $F_{in,i}$ and $F_{out,i}$ denote the volume arterial inflow and venous outflow rates of blood. The weighting factor φ_i is a function of plasma, blood, and tissue volumes as well as the partition coefficients, and is used to calibrate the unbound plasma drug concentration to effective concentration associated blood flow, ensuring consistency in mass-balance calculations.

Three compartments are considered to represent the physiological structure of the GI tract: the stomach, small intestine, and large intestine. Within each compartment, a luminal sub-compartment is introduced to describe intraluminal particle transport, dissolution, and drug absorption. The GI compartments are interconnected through GI transit and systemic blood flow (Figure 2). A single luminal component is used for stomach and large intestine. However, small intestine lumen is subdivided into three segments — duodenum, jejunum, and ileum — to capture regional differences in luminal conditions, including surface area and absorption capacity. Within each luminal segment, a PBE is employed to describe the evolution of API particle size distribution influence by dissolution ($D_c(t)$) and transit flow.

$$\frac{\partial(n_j V_{L,j})}{\partial t} = f_{j-1}n_{j-1} - f_j n_{j-} \frac{\partial(n_j D_c V_{L,j})}{\partial L} \tag{8}$$

where n_j denotes the API number density in j^{th} lumen, f_j and $V_{L,j}$ represents the luminal transit volume flow rate and the liquid volume, respectively.

The luminal liquid phase is assumed to consist primarily of solvent. The luminal liquid volumes of the small and large intestines are assumed to be constant due to bal-anced inflow and outflow. In contrast, the luminal volume of stomach is time-dependent variable because it lacks the upstream gastrointestinal supply. The PBE is coupled with a solute mass balance to describe the mass

of dissolved API ($m_{API,j}$) in each luminal compartment. This accounts for contributions from dissolution and luminal transit. It is noted that for the small intestine lumens, additional drug absorption terms are incorporated to represent systemic uptake:

$$\frac{dm_{API,j}}{dt} = f_{j-1}\frac{m_{API,j-1}}{V_{L,j-1}} - f_j\frac{m_{API,j}}{V_{L,j}} - \int_0^\infty k_c\rho_c L^2 n_j D_c V_{L,j}\, dL - P_1 - P_2 \quad (9)$$

Drug absorption is governed by two primary mechanisms: passive and active transport. Passive transport (P_1) is modeled as a permeability-limited flux driven by the concentration gradient across the intestinal membrane. The Passive transport rate is modeled by Fick's law equation, in which permeability (K) is considered. Active transport (P_2) process is mediated by membrane transporters and is modeled using Michaelis–Menten kinetics, where the kinetics is controlled by maximum transport velocity (V_{max}) and protein-drug affinity (K_m).

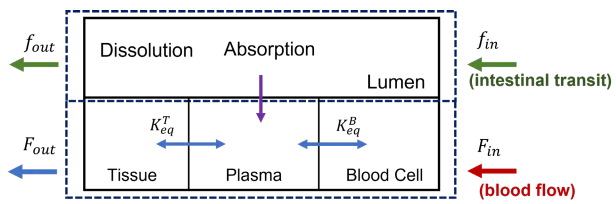

Figure 2. Small intestine compartment scheme.

Drug metabolism and excretion are modeled within the liver and kidney compartments, respectively. Since the hepatic metabolism is also an enzyme-mediated bioprocess, it is described using Michaelis–Menten equation. The excretion is considered in the kidney compartment by using the total renal clerance.

2.4 Dose regimen optimization formulation

For this study, the optimization problem is expressed as a non-linear constrained design problem. The objective function is to minimize the integral of the squared deviation between the predicted plasma concentration profile and a desired target concentration. The proper definition can be seen in following equations:

$$\min_x J(x,z,t) \quad (10)$$

$$s.t. \frac{dz}{dt} = f(x,z,t)\ t \in [0,T] \quad (11)$$

$$x_{lb} \le x \le x_{ub}$$

$$g_i(x,z,t) \le 0$$

Eq. (11) represents the combined PBM-PBPK model as a differential-algebraic equation (DAE) system. The decision variable x denotes the dose amount, where x_{lb} and x_{ub} represent the lower and upper bounds, respectively. The treatment horizon T is defined and considered fixed. In this case, the dosing interval is set to a constant value. The constraint equation g_i represent the toxicity limit. This optimization problem is solved using PyNomad.

3. RESULTS AND DISCUSSION

Within this proposed framework, rivaroxaban, an anticoagulant, is selected as the model drugs. Rivaroxaban is classified as Biopharmaceutics Classification System (BCS) class II drug. BCS class II drugs have low solubility and high permeability, for which dissolution is considered to be the rate-determining step for absorption. The PBM-based tablet disintegration and dissolution model is first calibrated using published in vitro dissolution profiles (Figure 3) [7]. These two dissolution profiles were obtained from dissolution tests of commercial rivaroxaban tablet (Xarelto) in fasted-state simulated gastric fluid (FaSSGF) and intestinal fluid (FaSSIF). The parameters that set a *prior* are summarized in Table 1. The parameters of Washburn equation (Eq. (3)) were obtained from literatures [8].

For the input CMAs parameters, the CSD is assumed to follow a Gaussian distribution with a mean size reported in the literature [7], and the tablet porosity is assumed to be 0.85. The two fitted parameters, including dissolution rate constant and effective capillary radius constant, are shown in Table 2. The fitted results (Figure 3) show good agreement between the PBM model predictions and dissolution profiles.

Table 1: Parameters of the disintegration and dissolution PBM set to literature values

Parameter	Description	Value
C_{sg}	Solubility in FaSSGF	0.011mg/mL
C_{si}	Solubility in FaSSIF	0.0121mg/mL
ρ_c	Crystal density	1491 kg/m^3
γ	Surface tension	0.00723 N/m
δ	Liquid viscosity	0.001Pa s
θ	Contact angle	88.2°

Table 2: Estimated parameter for the disintegration and dissolution PBM

Parameter	Description	Value
k_d	Dissolution rate constant	0.039 μm/s
r_e	Effective capillary radius	1509 μm

The corresponding average plasma drug concentration data were obtained from 40 patients under fasted conditions following a single 15-mg dose of rivaroxaban [7]. For PBPK model, the body characteristic parameters were calculated based on an average standard adult BMI value. Cardiac output and tissue mass were taken from published data [9] and rescaled to the reference BMI using allometric relationships.

Blood flow in each compartment is assigned

proportionally to tissue volume while ensuring the total flow matched cardiac output. For drug-specific parameters, tissue-to-plasma partition coefficients (K_eq^T) are predicted using the Rodgers and Rowland model [10], with the compound type set to neutral. The blood-to-plasma partition coefficients (K_eq^B) are set to be 0.8, following previous report [11]. The unbound fraction (f_u) is set as 0.1 [12].

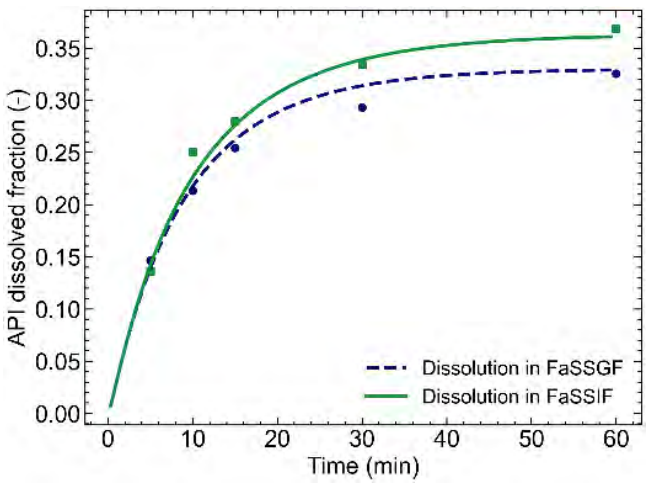

Figure 3. Fitted dissolution profiles under different physiological conditions.

For physiological processes, absorption is assumed to occur solely via passive transport. Given the high permeability of rivaroxaban, passive transport is expected dominate the absorption process. The permeability (K) is set as an adaptive parameter. Hepatic metabolism is modeled using Michaelis-Menten equation, with K_m fixed based on reported value for CYP3A4, which is one of the main enzymes responsible for rivaroxaban metabolism in liver [13]. To prevent overfitting, renal clearance for kidney excretion is also assigned based on literature [12]. The overall estimated model parameters are summarized in Table 3.

Table 3: Estimated parameter of PBPK

Parameter	Description	Value
V_{max}	Maximum velocity	39.5 μg min^{-1} mL^{-1}
K_1	Permeability 1	0.032 min^{-1}
K_2	Permeability 2	0.084 min^{-1}
f_i	Flow rate constant	0.002 min^{-1}

The fitted results (Figure 4) demonstrate that the mean predicted plasma drug concentration profile of rivaroxaban from PBM-PBPK model showed good agreement with the observed data. The predicted area under the curve from the time of administration up to 48 hours (AUC_{0-48h}) is 1667 h μg/L and the (AUC_{0-48h}) for observed data is 1710 h μg/L. The plasma drug concentration increases initially due to dissolution and absorption, reaching a maximum around 2 hr after administration. It then gradually decreases as the metabolism and excretion dominate drug elimination.

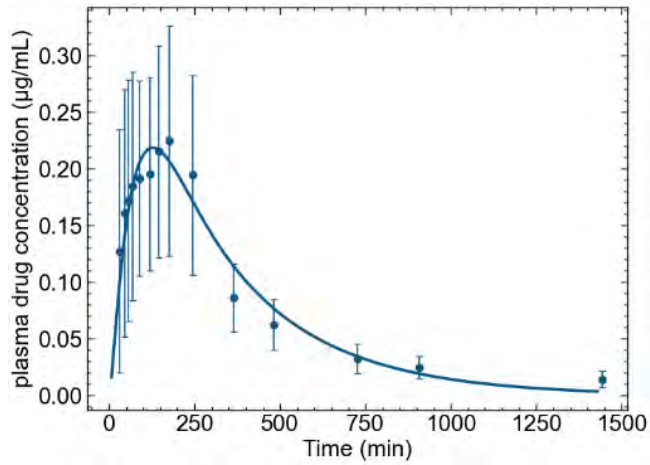

Figure 4. Fitted plasma drug concentration profile.

CSD is considered one of the most impactful tablet properties on dissolution [6], an increase in mean size reduces the dissolution due to the smaller available surface area. This slower dissolution decreases the drug concentration in the intestinal lumen, which in turn lowers the absorption rate. As a result, larger crystal sizes lead to lower maximum plasma concentration and a longer time to reach that maximum (Figure 5(a)). In contrast, tablet porosity primarily influences drug release through disintegration rate. The higher tablet porosity (smaller relative density) results in increased maximum plasma concentration, suggest the impact of initial concentration.

The calibrated model is further employed in the optimization problem. Figure 6 demonstrates the optimal dose regimen, in which the optimal single dose is 20.6 mg adminitered at a dose interval of 4 h over a 24 hr treatment period. In this study, the optimization problem is considered to focus only on the dose amount; however, within the proposed framework, it is possible to consider both the dose regimen and formulation properties. On one hand, the optimal dose regimen can be investigated under a fixed formulation design. On the other hand, a multi-objective optimization problem can be formulated to investigate the minimization of both cost (i.e., the amount of drug) and dosing regimen by considering both dose amount, dosing interval, and formulation properties as decision variables. The therapeutic window can be considered as a constraint not only for the maximum limit but also for the lower limit after the first dose.

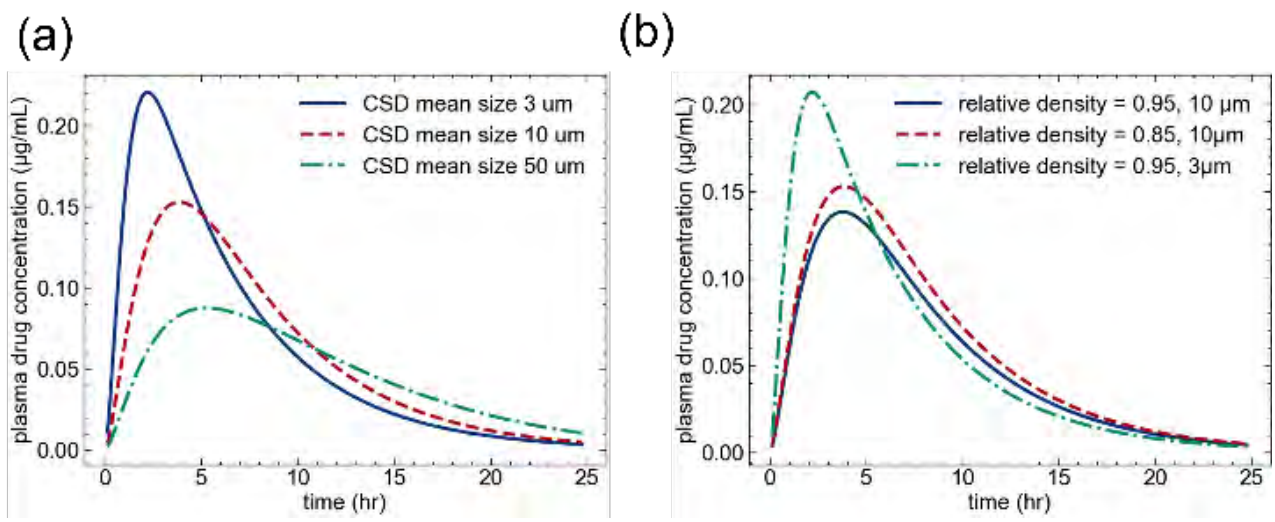

Figure 5. Impacts of (a) CSD and (b) tablet porosity on plasma drug concentrations.

4. CONCLUSION

This work presented an integrated PBM–PBPK digital design framework for rational formulation development and dosing strategy optimization of oral solid dosage forms. By mechanistically linking critical material attributes, including crystal size distribution and tablet porosity, to systemic pharmacokinetics, the framework enables quantitative assessment of formulation–performance relationships beyond traditional approaches. Application to rivaroxaban demonstrated that crystal size distribution has a dominant impact on plasma exposure. Furthermore, optimization of dose amount illustrated the framework's capability to support model-informed dosing decisions, highlighting its potential for data-efficient formulation and regimen design.

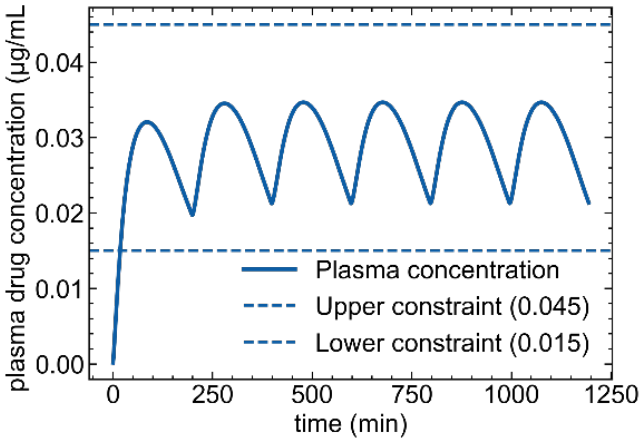

Figure 6. Dose regimen under the optimal dose amount

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support received from the National Science Foundation (NSF) under Grant No. 2132142.

REFERENCES

1. Aungst BJ. Optimizing oral bioavailability in drug discovery: an overview of design and testing strategies and formulation options. Journal of Pharmaceutical Sciences 106:921-929 (2017). https://doi.org/10.1016/j.xphs.2016.12.002
2. Zaborenko N, Shi Z, Corredor CC, Smith-Goettler BM, Zhang L, Hermans A, Neu CM, Alam MA, Cohen MJ, Lu X, Xiong L, Zacour BM. First-principles and empirical approaches to predicting in vitro dissolution for pharmaceutical formulation and process development and for product release testing. AAPS J 21: (2019). https://doi.org/10.1208/s12248-019-0297-y
3. Lu Y, Kim S, Park K. In vitro–in vivo correlation: perspectives on model development. International Journal of Pharmaceutics 418:142-148 (2011). https://doi.org/10.1016/j.ijpharm.2011.01.010
4. Willmann S, Thelen K, Becker C, Dressman JB, Lippert J. Mechanism-based prediction of particle size-dependent dissolution and absorption: cilostazol pharmacokinetics in dogs. European Journal of Pharmaceutics and Biopharmaceutics 76:83-94 (2010). https://doi.org/10.1016/j.ejpb.2010.06.003
5. Pepin XJH, Flanagan TR, Holt DJ, Eidelman A, Treacy D, Rowlings CE. Justification of drug product dissolution rate and drug substance

particle size specifications based on absorption PBPK modeling for lesinurad immediate release tablets. Mol. Pharmaceutics 13:3256-3269 (2016). https://doi.org/10.1021/acs.molpharmaceut.6b00497

6. Ma Y, Li W, Liu J, Shang G, Yang H, Gong J, Nagy ZK, Benyahia B. Digital design and optimization of the integrated synthesis and crystallization process using data?driven approaches. AIChE Journal 71: (2025). https://doi.org/10.1002/aic.18931
7. Ishihara M, Kambayashi A. Physiologically based biopharmaceutics modeling to predict bioequivalence of rivaroxaban formulated as conventional tablets, OD tablets, and fine granules. Journal of Drug Delivery Science and Technology 114:107631 (2025). https://doi.org/10.1016/j.jddst.2025.107631
8. Ferdoush S, Gonzalez M. A two-stage mechanistic reduced-order model of pharmaceutical tablet dissolution: population balance modeling and tablet wetting functions. International Journal of Pharmaceutics 664:124635 (2024). https://doi.org/10.1016/j.ijpharm.2024.124635
9. Valentin J. Basic anatomical and physiological data for use in radiological protection: reference values. Ann ICRP 32:1-277 (2002). https://doi.org/10.1016/s0146-6453(03)00002-2
10. Rodgers T, Rowland M. Mechanistic approaches to volume of distribution predictions: understanding the processes. Pharm Res 24:918-933 (2007). https://doi.org/10.1007/s11095-006-9210-3
11. Deepika D, Kumar V. The role of “physiologically based pharmacokinetic model (PBPK)” new approach methodology (NAM) in pharmaceuticals and environmental chemical risk assessment. IJERPH 20:3473 (2023). https://doi.org/10.3390/ijerph20043473
12. Terrier J, Gaspar F, Gosselin P, Raboud O, Lenoir C, Rollason V, Csajka C, Samer C, Fontana P, Daali Y, Reny J, . Apixaban and rivaroxaban's physiologically?based pharmacokinetic model validation in hospitalized patients: a first step for larger use of a priori modeling approach at bed side. CPT Pharmacom & Syst Pharma 12:1872-1883 (2023). https://doi.org/10.1002/psp4.13036
13. Mueck W, Stampfuss J, Kubitza D, Becka M. Clinical pharmacokinetic and pharmacodynamic profile of rivaroxaban. Clin Pharmacokinet 53:1-16 (2013). https://doi.org/10.1007/s40262-013-0100-7

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Experiments & Modelling of Batch Fermentation of *Fusarium venenatum* on Glucose-Fructose Mixtures

Tom Vinestock[a], Miao Guo[a]*

[a] King's College London, Department of Engineering, London, United Kingdom
* Corresponding Author: miao.guo@kcl.ac.uk

ABSTRACT

Single-cell protein (SCP) fermentation efficiently converts carbohydrates into high-protein food products but typically relies on purified glucose. Better understanding of SCP growth on mixed sugar substrates could allow for use of lower-cost, less processed sugars and waste-derived feedstocks. Glucose–fructose mixtures are particularly relevant, as these are the main sugars in sucrose hydrolysates and are also common in many food and beverage waste streams. In this study, the growth of *Fusarium venenatum* on glucose–fructose mixtures was investigated experimentally and modelled using a lagged dual-substrate Monod framework incorporating inhibition of fructose growth by glucose. Batch fermentations were conducted at a fixed total sugar concentration (15 g/L) with four different initial substrate compositions. Model parameters were estimated using both single-experiment and multi-experiment fitting strategies using differential evolution and wild bootstrap uncertainty analysis. A shared parameter set reproduced biomass growth and substrate depletion across all conditions with only a modest increase in prediction error relative to individual fits. Bootstrap analysis revealed substantial correlation among uptake and inhibition parameters, while yield coefficients were well constrained. These results demonstrate that mixed-sugar growth of *F. venenatum* can be described using a unified kinetic model, supporting the feasibility of transitioning from glucose-only to mixed-substrate SCP fermentation.

Keywords: Batch Process, Biomass, Biosystems, Fermentation, Modelling and Simulations, Single Cell Protein, Dual Substrate Growth

BACKGROUND

Single-cell protein (SCP) fermentation provides an efficient route for converting carbohydrate into high-protein food products and has the potential to reduce land use, greenhouse gas emissions, and resource intensity relative to conventional animal-based protein production [1]. Among SCP microorganisms, filamentous fungi such as *Fusarium venenatum* are of particular industrial interest because their mycelial morphology imparts a fibrous, meat-like texture desirable in food applications [2]. Industrial SCP processes commonly rely on refined carbohydrate feedstocks, most commonly glucose [3]. However, the use of single, purified substrates increases upstream processing requirements and cost and limits the flexibility of fermentation systems to utilise heterogeneous, low-cost and waste-derived carbon sources [4].

There are many potential feedstocks that are composed of mixtures of sugars rather than a single monosaccharide. Glucose–fructose mixtures are especially relevant, because sucrose hydrolysates and many sugar-rich food and beverage waste streams, including fruit-processing and soft-drink manufacturing waste streams are rich in these monosaccharides [5]. When glucose and fructose are both present, there will be an interaction between their metabolisms. Quantifying this interaction is essential for constructing accurate growth models and therefore, for the rational design of fermentation strategies based on mixed-sugar feedstocks.

In this work, we estimate the parameters of a dual-substrate growth model describing the growth of *F. venenatum* on glucose–fructose mixtures. By explicitly characterising mixed-substrate kinetics, this study provides a basis for assessing the feasibility of transitioning from glucose-based operation to mixed-sugar fermentation and so supports the development of more flexible, sustainable and economically robust SCP production

processes. This will help to inform industrial process design and operation, helping to derisk the transition to dual substrate fermentation, from existing glucose-based processes.

METHODOLOGY

Experimental Methodology

50 ml of fermentation media, as defined in Table 2, was transferred into a 250 ml baffled flask. Total sugar concentration was fixed at 15 g/L. Four different initial substrate compositions were investigated: 15 g/L glucose, 10 g/L glucose + 5 g/L fructose, 5 g/L glucose + 10 g/L fructose, and 15 g/L fructose. This was then inoculated with *F. venenatum* A3/5 spores to give an initial spore concentration of 10, 000/ml in each flask. Spores were generated by inoculating a 30g/l glucose/low-nitrogen RHM media with a plate culture provided by Quorn Foods, incubating this broth in a shaker-incubator at 28°C for 4 days, and then straining to remove fungal filaments.

To ensure sufficient oxygen transfer to the fermentation broth, oxygen and compressed air were blended to give a headspace oxygen concentration of 200% of the oxygen concentration of air (~42% oxygen by volume). This enriched air was humidified before it entered the fermentation flask to prevent evaporation losses. The oxygen flow rate was ~0.2 LPM and the air flow rate was ~0.5 LPM, for a total flow rate of ~0.7 LPM, or ~0.2 LPM /flask. This is substantially in excess of metabolic requirements, meaning that the headspace gas composition should remain unchanged during the fermentation. The flasks were vented to maintain atmospheric pressure throughout the fermentation.

PreSens SFR vario readers (PreSens, Germany) were used to measure biomass concentration, pH, and broth dissolved oxygen concentration during the fermentation runs. The biomass concentration measure is based on the reflectance of the fermentation broth, while pH and dissolved oxygen are determined based on pre-calibrated optical sensor spots. Biomass and pH concentration were recalibrated after each experiment using the start and end points. At the endpoint, to determine the biomass concentration, the fermentation broth from each flask was vacuum filtered through a pre-dried, pre-weighed fine fibre-glass filter. It was then rinsed with deionised water and then dried to constant mass in a 65 °C oven, before the filter was reweighed using a precision balance.

Each experiment was conducted in triplicate. During each run, 300 µL samples were collected from the flasks at regular intervals, with sampling concentrated in the exponential growth phase, for substrate analysis by high-performance liquid chromatography (HPLC). The HPLC parameters are given in Table 3. Before the fermentation experiments, the linear relationship between peak area and substrate concentration was confirmed, and calibration curves for glucose and fructose were established over the full concentration ranges observed. The fermentation was allowed to continue until HPLC analysis showed the sugar substrates had been fully depleted.

Oxygen uptake rate (OUR) was calculated according to Equation 1, where kLa is the volumetric mass transfer coefficient, C_{O_2} is the measured dissolved oxygen concentration in the fermentation broth, and $C^*_{O_2}$ is the equilibrium dissolved oxygen concentration based on Henry's Law. kLa was determined using the gassing-out method, in which dissolved oxygen was first reduced to near zero by nitrogen sparging and then monitored during re-aeration using a PreSens DO sensor.

$$OUR = kLa * (C^*_{O_2} - C_{O_2}) \quad (1)$$

Biomass growth rate was estimated by applying a Savitzky–Golay differentiating filter [6] to the biomass concentration time series, in order to reduce the effect of noise in the biomass concentration measurement. The specific growth rate was estimated by linear regression of $ln(X)$ against time within a centred sliding window, where X is the biomass concentration, with the local slope taken as the instantaneous growth rate.

Table 1: Experimental Parameters

Parameter	Value
Broth Volume	50 ml
Shaker Orbit	20 mm
Shaker Rotational Speed	220 RPM
Incubation Temperature	28 °C
Headspace Oxygen Concentration	200 % air saturation
Volumetric Mass Transfer Coefficient (kLa)	96 /h

Table 2: Fermentation Media Components

Media Component	Amount per Flask (ml)
Rank Hovis McDougal Media (RHM) [7] with 30 g/l KH_2PO_4	46
Polypropylene Glycol (5% w/w) Solution	0.5
Tween-80 (1% w/w) Solution	0.5
Glucose (25% w/v)	3-x
Fructose (25% w/v)	x
x∈ {0, 1, 2, 3}	

Figure 1. Photograph of experimental set-up. The blue base units are the PreSens Shake Flask Readers. The media bottles visible on the right humidify the enriched air before it enters the baffled flasks.

Table 3: HPLC Parameters

Parameter	Value
Column	Bio-Rad Aminex HPX-87H
Detector	Shimadzu RID-20A
Mobile Phase	0.001 M H_2SO_4
Run Time	35 min
Column Temperature	30°C
Flow Rate	0.4 mL/min
Injection Volume	5 mL

Computational Methodology

Biomass growth was modelled using a lagged dual-substrate Monod [8] framework, with glucose assumed to be the preferred substrate. Accordingly, the model includes inhibition of fructose utilisation by glucose, but no inhibition of glucose utilisation by fructose. During the lag phase of duration λ, biomass growth and substrate consumption were assumed to be negligible, such that $\dot{X} = \dot{S}_1 = \dot{S}_2 = 0$. After the end of the lag phase, the system dynamics are given by:

$$\dot{X} = (\mu_1 + \mu_2) X \tag{2}$$

$$\dot{S}_1 = \frac{-\mu_1}{Y_1} X \tag{3}$$

$$\dot{S}_2 = \frac{-\mu_2}{Y_2} X \tag{4}$$

with specific growth rates:

$$\mu_1 = \mu_{max,1} \frac{S_1}{S_1 + K_{S1}} \tag{5}$$

$$\mu_2 = \mu_{max,2} \frac{S_2}{S_2 + K_{S2}} \frac{1}{1 + \frac{S_1}{K_I}} \tag{6}$$

Here, X denotes biomass concentration, $S_1$and S_2 are the concentrations of glucose and fructose, respectively, $\mu_{\max,1}$ and $\mu_{\max,2}$ are the maximum specific growth rates, $K_{S1}$and $K_{S2}$are the Monod half-saturation constants, K_I is the glucose inhibition constant for fructose growth and $Y_1$and Y_2 are yield coefficients. A step-function lag phase, as described in [9], was included to account for delayed onset of growth. The $\frac{1}{1+\frac{S_1}{K_I}}$ term models catabolite repression of fructose by glucose, as in [10]. Note that K_I is the minimum glucose concentration sufficient to halve the fructose-related specific growth rate, μ_2. Therefore, smaller values of K_I indicate stronger inhibition.

Model parameters were estimated using two fitting strategies: (i) individual fits to each experiment, and (ii) a combined multi-experiment fit. In the multi-experiment fit, seven kinetic parameters ($\mu_{\max,1}, \mu_{\max,2}, K_{S1}, K_{S2}, K_I, Y_1$, and Y_2) were constrained to be identical across all four experiments, while the lag time and initial biomass concentration were estimated separately for each experiment, giving a total of 15 fitted parameters. The values obtained from this shared-parameter fit were subsequently compared with those obtained from the single-experiment fits to assess the consistency and robustness of the kinetic estimates.

Parameter estimation was performed by minimising a weighted sum of squared errors across biomass and substrate measurements from triplicate flasks. To account for the higher sampling frequency of biomass measurements relative to substrate measurements, substrate residuals were weighted by $w_S = \frac{n_X}{n_S}$, where n_X is the number of biomass measurements and n_S is the total number of substrate measurements. This weighting equalises the aggregate contribution of the substrate data to that of the biomass data.

Equation 7 defines the single-flask error, with $i \in \{1, \dots, 4\}$ indexing the initial substrate condition, $f \in \{1, 2, 3\}$ the flask replicate, k the biomass observation time index, and j the substrate observation time index. The error for each initial substrate condition was then $E_{\mathrm{i}} = \sum_{f=1}^{3} E_i$. This is used as the objective function for the single experiment fit. For the multi-experiment fit, the objective function was $E_{\mathrm{tot}} = \sum_{i=1}^{4} E_i$.

$$E_{i,f} = \sum_{k=0}^{n_x-1} \left(X_{i,f}^{pred}[k] - X_{i,f}^{obs}[k] \right)^2 + \sum_{j=0}^{n_s-1} w_s \left(S_{i,f,1}^{pred}[j] - S_{i,f,1}^{obs}[j] \right)^2 + \sum_{j=0}^{n_s-1} \mathrm{w_s} \left(S_{i,f,2}^{pred}[j] - S_{i,f,2}^{obs}[j] \right)^2 \tag{7}$$

Table 4: Kinetic Parameter Bounds

Parameter	Lower Bound	Upper Bound	Type	Unit
$\mu_{max,1}$	0.01	0.6	Global for shared fit	/h
$\mu_{max,2}$	0.01	0.6		/h
K_{S1}	0.001	5.0		g/L
K_{S2}	0.001	5.0		g/L
K_I	0.001	5.0		g/L
Y_1	0.01	1.0		g/g
Y_2	0.01	1.0		g/g
λ	0	60.0	Experiment specific	h
X_0	0	0.5		g/L

Optimisation was performed using differential evolution as implemented in SciPy [11], using the best1bin strategy, a population size of 50, mutation dithering over [0.5, 1.0], and a crossover probability of 0.7, followed by local polishing with L-BFGS-B. The bounds used for each parameter are set out in Table 4.

Confidence intervals for the parameter estimates were constructed using the wild bootstrap method [12], which does not require the assumption of homoscedastic residuals. Per-observation residuals were first computed from the fitted model on the original data. For each of 50 bootstrap replicates, a synthetic response dataset $y_n^*$was generated by adding randomly signed residuals to the model fitted values. Specifically, if $\hat{y}_n$ denotes the predicted value at observation n and $\hat{e}_n = y_n - \hat{y}_n$the corresponding residual, then the bootstrap data were constructed as $y_n^* = \hat{y}_n + w_n \hat{e}_n$, where $w_n \in \{-1, +1\}$ is a Rademacher random variable taking each value with equal probability. This procedure preserves the magnitude and structure of the original residuals while randomly reversing their sign, thereby maintaining the mean trend and accommodating heteroscedastic error behaviour. Each synthetic dataset was then re-fitted using the same differential evolution procedure, and the distribution of the resulting parameter estimates across all replicates was used to construct 95% bootstrap confidence intervals.

RESULTS & DISCUSSION

Experimental Results & Discussion

The experimental results in Figure 2 show good repeatability within an experiment, particularly for the substrate and biomass concentrations, as well as pH. There are greater differences between replicates for oxygen concentration, oxygen uptake rate (OUR), biomass growth rate and specific growth rate. With respect to biomass growth rate and specific growth rate (SGR), part of this is probably due to noise in the biomass concentration readings. However, despite this, many of the results are as would be expected: the peak OUR and SGR are highest for the pure glucose substrate, and monotonically reduce as the initial glucose concentration is reduced in favour of fructose. The SGR results suggest, in contrast to the lag model, which assumes that SGR is maximum at the very start of the exponential growth phase, that there is a gradual ramp up in SGR, until the falling substrate concentrations begin to reduce SGR again.

The length of the lag phase doesn't follow the expected pattern, with a shorter lag phase for the fructose-glucose mixture than for the pure glucose. This is believed to be driven by differences in the age and viability of the inoculation conidia.

The HPLC results show some anomalies, with the substrate concentrations not always monotonically decreasing with time. This could be due to heterogeneity in the fermentation broth, or HPLC-related factors such as peak overlap with other broth components, baseline drift, or small variations in sample preparation and integration that introduce apparent fluctuations in the measured concentrations.

Computational Results & Discussion

The MAE for the multi-experiment fit was larger than that of the single-experiment fits, as expected due to the reduced number of free parameters. However, despite the single-experiments fit providing many more degrees of freedom, since the single-experiment fits allow independent estimation of kinetic parameters for each dataset, whereas the multi-experiment fit enforces a shared parameter set, the difference in the magnitude of the MAE between the two is slight, at c. 30% higher total MAE for the multi-experiment fit compared to the single experiment fit. MAE for the multi-experiment is still below 0.5 g/l.This suggests that the model is flexible enough to describe growth rates with various initial glucose and fructose concentration, without requiring re-estimation of kinetic parameters for each initial condition. From a process modelling perspective, this supports the use of a globally parameterised model for predictive simulations and control design.

Figure 4 shows that there is considerable variation in the single experiment parameter estimates, suggesting weak identifiability from individual datasets, particularly as some of the parameters estimated based on single experiments seem to be outside the plausible range. For example, the Y_2 estimate from the Glu 10 g/l + Fru 5 g/l is nearly 1. The parameters estimated on data from all four experiments are more reliable. However, the $K_{S,1}$estimate would be expected to be smaller than the $K_{S,2}$estimate, as glucose is the favoured substrate.

Comparison of the median $\mu_{\max,1}$and $\mu_{\max,2}$ estimates (Figures 4 and 6) with the observed specific growth rates under pure glucose and pure fructose

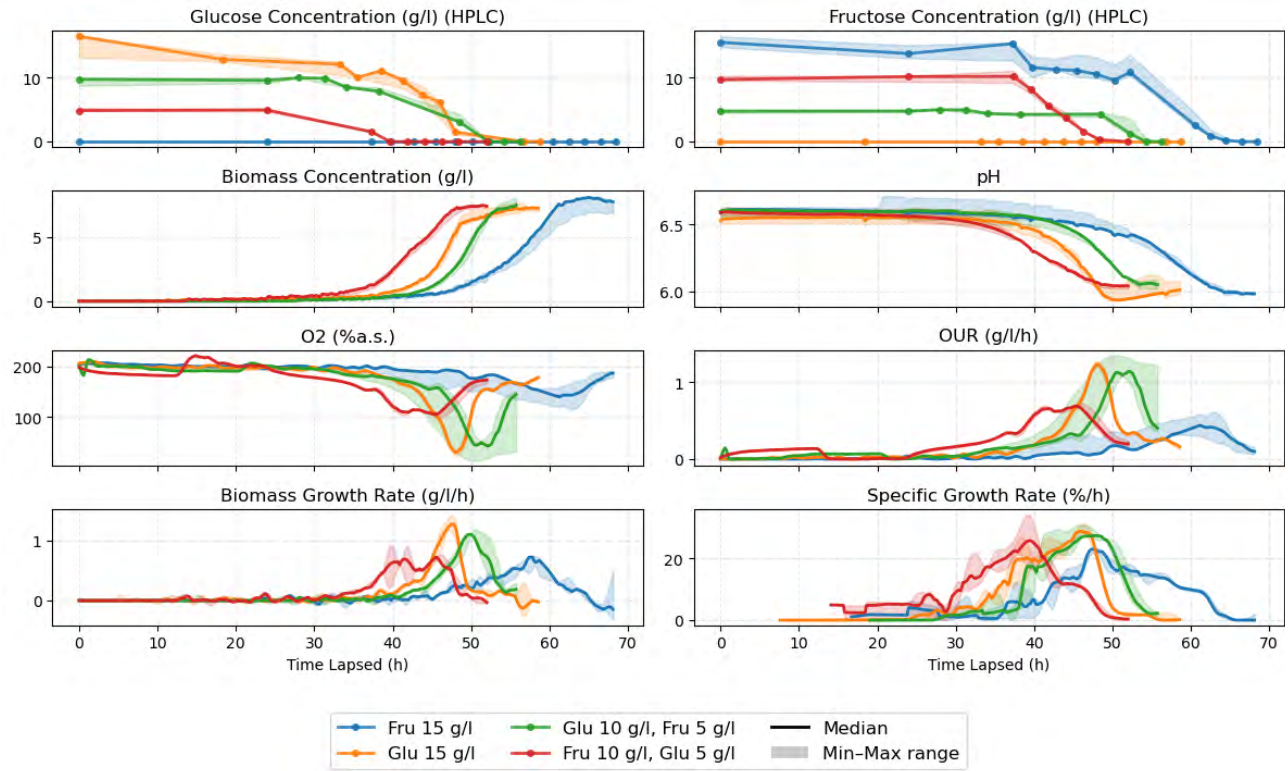

Figure 2. Experimental results for glucose-fructose mixed-substrate batch fermentations. Substrate and biomass concentrations, as well as pH and O_2, were directly measured. OUR, biomass growth rate and specific growth rate were calculated based on the measurements as explained in the Methodology.

conditions (Figure 2) shows that the peak specific growth rates exceed the corresponding estimates. This reflects the modelling assumption of a step change in growth rate at the end of the lag phase, which is not supported by the data. Instead, the specific growth rate increases gradually to a maximum, leading to underestimation of the true peak growth rate by $\mu_{\max,1}$ and $\mu_{\max,2}$.

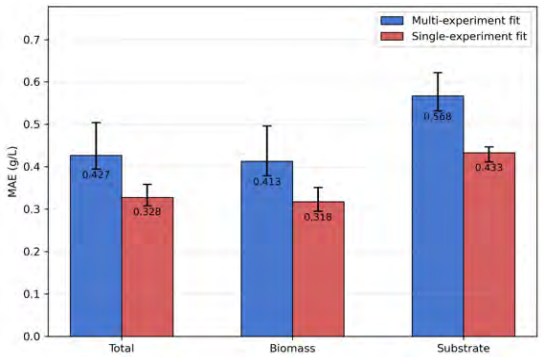

Figure 3. Mean absolute error (MAE) of multi-experiment and single-experiment model fits evaluated against observed biomass and substrate concentrations. Bar heights represent the median MAE across 20 wild bootstrap replicates and error bars denote the 95% confidence interval.

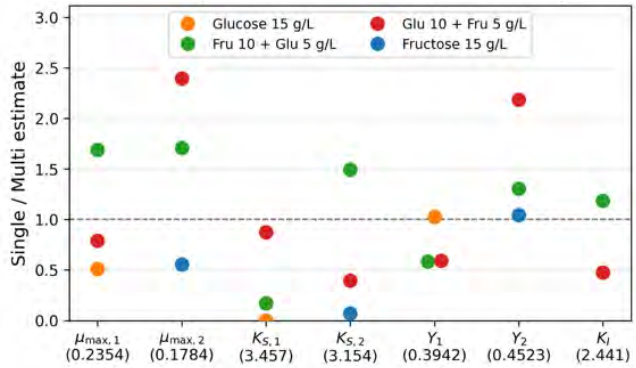

Figure 4. Single-experiment parameter estimates (median of 20 bootstrap replicates) for each experimental condition, normalised by the corresponding multi-experiment estimate (median of bootstrap replicates). A value of unity indicates agreement between the two fitting approaches.

Importantly, the multi-experiment model reproduces the qualitative and quantitative features of the growth and substrate depletion dynamics across all 4 experiments (Figure 5), capturing the timing and shape of the exponential growth phase and the coupled consumption of glucose and fructose. Deviations between the

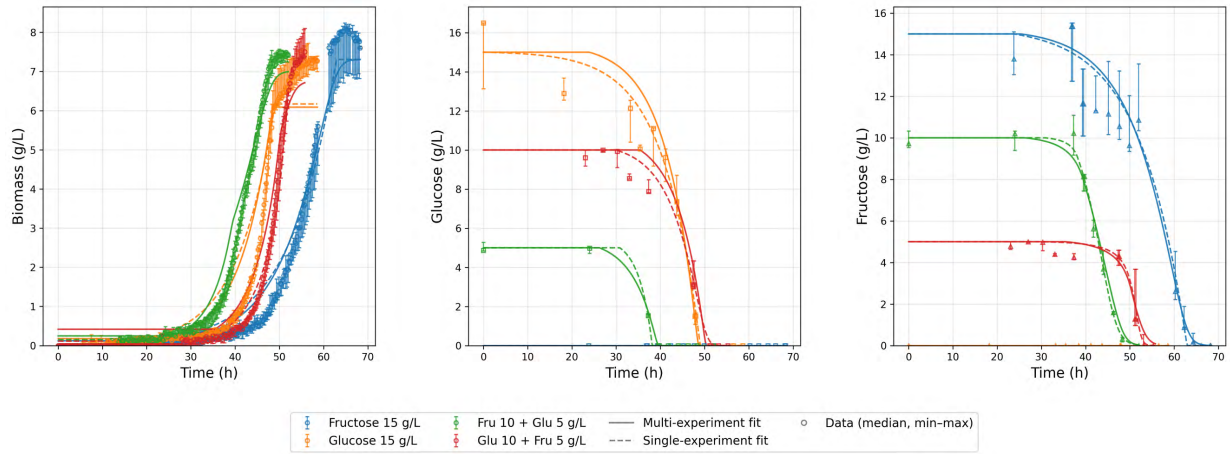

Figure 5. Comparison of single-experiment and multi-experiment model fits against original experimental data for a lagged dual-substrate Monod kinetic model, for batch fermentation across 4 different initial substrate conditions.

model predictions from the shared parameters and the experiment-specific parameters are generally small.

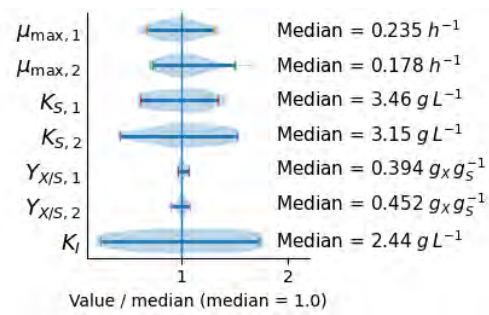

Figure 6. Median value and 95% confidence intervals for glucose-fructose dual substrate growth parameters based on wild bootstap method with 20 replicates.

The bootstrap results show considerable residual uncertainty in many of the parameter estimates. With the exception of the two yield coefficients, the redundancy in the model and the limited amount of substrate concentration data make it difficult to disentangle the effects of $\mu_{\max,1}$ and $K_{S,1}$ in the case of glucose growth, and $\mu_{\max,2}$, $K_{S,2}$ and K_I in the case of fructose growth. This can be seen in the Spearman Correlation Heatmap, which shows strong positive correlation between $\mu_{\max,1}$ and $K_{S,1}$, as well as $\mu_{\max,2}$ and $K_{S,2}$, and strong negative correlation between K_I and both $\mu_{\max,2}$ and $K_{S,2}$. As a result, the observed fructose uptake can be equivalently explained by a reduced $\mu_{\max,2}$, a high effective $K_{S,2}$, or strong glucose inhibition (low K_I), limiting mechanistic interpretability of the fitted parameters.

CONCLUSION

This study investigated the growth of *F. venenatum* on glucose–fructose mixtures and quantified the associated kinetics using a lagged dual-substrate Monod model with glucose inhibition of fructose utilisation. Although noise in derived rates and minor non-monotonicity in HPLC measurements introduced uncertainty, the overall trends were reproducible across replicates and aligned with expected physiological behaviour. Notably, the specific growth rate exhibited a gradual increase during early exponential growth rather than an instantaneous maximum, indicating that a strict step-change assumption at the end of the lag phase may be an oversimplification of the underlying dynamics.

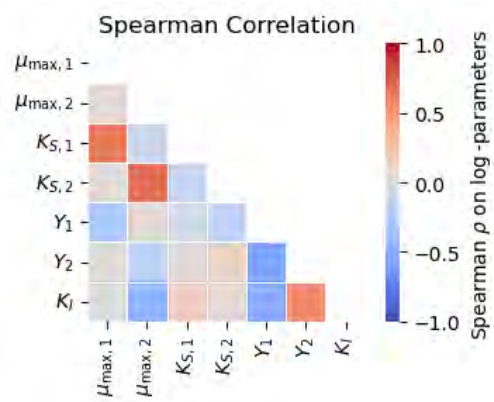

Figure 7. Heatmap of pairwise Spearman rank correlation coefficients between kinetic parameter estimates across wild bootstrap replicates, indicating the degree of parameter interdependence in the fitted model.

Parameter estimation using both single-experiment and multi-experiment fitting demonstrated that a common set of kinetic parameters can describe growth across a range of initial substrate compositions with only a modest loss of accuracy relative to individually fitted models. This supports the suitability of a unified dual-

substrate framework for representing mixed-sugar growth in *F. venenatum*. Wild bootstrap analysis revealed substantial uncertainty and correlation among several kinetic parameters, particularly those governing fructose uptake and inhibition, reflecting structural redundancy in the model and limited resolution in the available substrate data. By contrast, yield coefficients were comparatively well identified, suggesting that overall carbon conversion efficiency is more robustly constrained than individual uptake kinetics.

Together, these results provide quantitative evidence that glucose–fructose mixtures can be modelled within a mechanistic kinetic framework and that mixed-substrate operation can be described without re-parameterising the model for each feed composition.

This contributes to the development of predictive tools for SCP fermentations using heterogeneous sugar streams and supports the feasibility of transitioning from glucose-only processes to mixed-sugar feedstocks.

FUTURE WORK

Future work will investigate the applicability of this model to continuous fermentation of *F. venenatum* on glucose-fructose mixtures. This is a more industrially relevant mode of operation, and will help to address the key question of whether significant simultaneous co-utilisation of glucose and fructose can be realised in industrially relevant conditions.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge funding from Quorn and the BBSRC, under project reference 2725958. Additionally, the authors would like to thank Mark Taylor (Quorn Foods) for technical advice, and material support, such as access to the *F. venenatum* A3/5 strain.

AUTHOR IDENTIFIERS

Author ORCIDs:
Vinestock T: 0009-0002-7689-2837
Guo M: 0000-0001-7733-5077

REFERENCES

1. Good Food Institute, Fermentation: State of the industry report. (2024)
2. Finnigan TJA, Theobald HE, Bajka B. Mycoprotein: a healthy and sustainable source of alternative protein-based foods. Annual Review of Food Science and Technology 16:105-125 (2025). https://doi.org/10.1146/annurev-food-111523-121802
3. Yang Z, Li L, Chen Y, Zhu L, Zhu Z, Jiang L. Trends in customizable single-cell protein production enabled by synthetic biology: carbon-negative biomanufacturing and engineered functionalities. J. Agric. Food Chem. 73:32957-32969 (2025). https://doi.org/10.1021/acs.jafc.5c08537
4. Vlaeminck E, Uitterhaegen E, Quataert K, Delmulle T, De Winter K, Soetaert WK. Industrial side streams as sustainable substrates for microbial production of poly(3-hydroxybutyrate) (PHB). World J Microbiol Biotechnol 38: (2022). https://doi.org/10.1007/s11274-022-03416-z
5. Thiviya P, Gamage A, Kapilan R, Merah O, Madhujith T. Single cell protein production using different fruit waste: a review. Separations 9:178 (2022). https://doi.org/10.3390/separations9070178
6. Savitzky A, Golay MJE. Smoothing and differentiation of data by simplified least squares procedures.. Anal. Chem. 36:1627-1639 (2002). https://doi.org/10.1021/ac60214a047
7. Whittaker JA. Sugar, soil and sequencing: Studies in *Fusarium venenatum*. PhD Thesis, Univ Nottingham (2022) https://repository.nottingham.ac.uk/handle/123456789/59078
8. Monod J. The growth of bacterial cultures. Annu Rev Microbiol 3:371-394 (1949). https://doi.org/10.1146/annurev.mi.03.100149.002103
9. Buchanan RL, Whiting RC, Damert WC. A comparison of the Gompertz, Baranyi, and three-phase linear models for fitting bacterial growth curves. Int J Food Microbiol 36(2–3):151–159 (1997) https://doi.org/10.1016/S0168-1605(97)01286-6
10. Schmitt E, Bura R, Gustafson R, Ehsanipour M. Kinetic modeling of moorella thermoacetica growth on single and dual-substrate systems. Bioprocess Biosyst Eng 39:1567-1575 (2016). https://doi.org/10.1007/s00449-016-1631-8
11. Virtanen P, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat Methods 17:261–272 (2020) https://doi.org/10.1038/s41592-019-0686-2
12. Wu CFJ. Jackknife, bootstrap and other resampling methods in regression analysis. Ann Stat 14(4):1261–1295 (1986) http://www.jstor.org/stable/2241454.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Towards Digital Threads for FAIR, Trustworthy, and Human-Centric Bioprocess Development

Jonas M. Karsten[a]*, Ernesto C. Martínez[a], Mariano N. Cruz Bournazou[a]

[a]Technische Universität Berlin, Bioprocess Engineering, Ackerstraße 76, 13355 Berlin, Germany
* Corresponding Author: j.karsten@tu-berlin.de

ABSTRACT

Decisions taken throughout a bioprocess lifecycle are often guided by heuristic knowledge that is difficult to summarize and sort, scattered across heterogeneous tools and documents, and partly retained as tacit expert mental models alongside fragmented computational models. This fragmentation remains a central barrier to reproducibility, transparent provenance, and systematic reuse of prior learning across comparable development projects. In this paper, it is argued that a key missing link toward Bioprocessing 5.0 is the digitalization of FAIR knowledge through a Cognitive Digital Thread that couples semantic knowledge graphs with AI methods to connect experimental data, protocols, workflows, and decision rationale with mathematical models and digital twins in a machine-actionable and auditable manner. A digitalization roadmap is outlined as a sequence of capability stages—from local device and data integration, to reproducible workflow execution and metadata capture, to semantic knowledge digitalization and explainable reasoning, to secure cross-institutional exchange of FAIR data and knowledge (optionally anchored via permissioned blockchain and smart contracts for authorship, versioning, and IP provenance), and finally to human-in-the-loop collaborative environments for model-informed bioprocess design and control. The result is a layered conceptual architecture that clarifies required functions and dependencies and provides a practical framework for implementing cognitive decision support across the bioprocess lifecycle.

Keywords: Bioprocess development, Industry 5.0, Cognitive Digital Thread, Information Management, Supply Chain, Blockchain, Knowledge Graphs

INTRODUCTION

Industry 4.0 introduced powerful technologies with connectivity, automation, digital twins (DTw) [1, 2], and data-driven control, that raised expectations for more robust, scalable, and faster bioprocess development by reducing uncertainty across scales and development stages through continuously linking evidence, models, and decisions. However, while physical and digital integration of systems has significantly advanced, expectations have largely fallen short in practice. While device integration, automation and data acquisition have improved, the inherent complexity of biological systems, limited and heterogeneous datasets, and missing semantic structure are currently the bottleneck in the development of bioprocesses. Experimental workflows in bioprocess engineering still rely heavily on manual steps, which are often not captured, resulting in experimental knowledge that is incompletely documented, metadata fragmented across institutional data silos, and critical decisions that are rarely traceable across development stages. The absence of semantic connectivity between devices, data streams, experimental workflows, and computational models across laboratories and IT systems means that data are exchanged without shared meaning [4], which undermines the reproducibility, robustness, scalability, and cross-institutional collaboration [3]. This results in experiments that cannot be reproduced, which imposes substantial hidden costs in research and development, as experiments must be repeated, process knowledge must be re-established during scale-up and technology transfer, and previously generated insights are lost due to insufficient data contextualization. Beyond waste of resources, these inefficiencies increase the risk of failure during development, delay time-to-market, and limit systematic learning

across projects and institutions. These limitations expose a fundamental mismatch between technology-centric Industry 4.0 concepts and the knowledge-intensive nature of bioprocess engineering, motivating a shift towards the human-centric Industry 5.0 paradigm [5, 6]. In this context, the concept of the Cognitive Digital Thread (CDTh) can be understood as a qualitative extension beyond Industry 4.0: a human-centric digital infrastructure that links data, models, decisions, and expertise across the bioprocess lifecycle, addresses semantic fragmentation, and establishes the foundations for reproducibility and provenance in experimentation and throughout the complete developmental pipeline. This contextualization turns datasets gathered at different experimentation scales into a reproducible development path by preserving their context from experimental planning through execution context, and up to analysis. Taken from system engineering, the digital thread in bioprocess engineering materializes as a cognitive cyber-physical system, in which physical experimentation, computational models, and human decision-making are tightly coupled through continuous digital feedback. Finally, to foster collaborative research and knowledge transfer across institutions, the system must be extended with a blockchain-enabled infrastructures [7, 8] in which smart contracts ensure transparent and verifiable provenance of intellectual property, authorship, and contributions.

From Digital Thread to Cognitive Digital Thread

A Digital Thread (DTh) is a structured and persistent information framework that links data, models, and decisions across the full lifecycle of a system, product, or process so that information remains traceable, interpretable, and reusable over time. The concept emerged in the early 2010s, most prominently in aerospace and defense [9], where long lifecycles, safety critical requirements, and fragmented engineering workflows created a need for end-to-end accountability and configuration control. In these settings, the core use case was to reduce information loss between lifecycle phases, enable traceability from requirements to implementation and maintenance, support compliance and auditability, and eliminate redundant data entry and manual translation across tools. Closely aligned with Model Based Systems Engineering and Product Lifecycle Management, Digital Thread approaches have since been applied beyond aerospace, where complex product structures and strict quality requirements similarly benefit from lifecycle wide traceability [10].

A Cognitive Digital Thread (CDTh) [11] extends this continuity beyond connectivity by making domain knowledge and human decision making first class, machine actionable elements of the thread to reinforce contextual understanding and provenance. This cognitive dimension arises from combining semantic representations such as KGs with semantic web technologies [4], interoperable lifecycle linking mechanisms, and AI supported reasoning and learning, so that the system can store artifacts and contextualize them, relate them, and reactivate them for decision support. In bioprocess engineering this is particularly valuable because biological variability, small and heterogeneous datasets, and pervasive manual steps often lead to poorly captured context, fragmented metadata, and persistent data silos, which hamper reproducibility and limit technology transfer. A CDTh addresses these bottlenecks by explicitly linking experimental conditions, workflow intent, model assumptions, digital objects, and experimental or computational outcomes into an evolving knowledge base. By this, the CDTh supports reproducible workflows, interoperable models/DTs, and traceable decision provenance across scales and organizations. DTs, for example, fit naturally within the structure of the CDTh [1, 2] as lifecycle specific virtual representations of physical processes, while the CDTh provides the overarching “information flow” that connects multiple twins with experimental evidence and design rationale. In collaborative settings, the CDTh can be complemented by permissioned blockchain infrastructures to anchor provenance and governance for shared digital objects (models, designs, datasets, etc.) enabling knowledge exchange across institutional boundaries while preserving authorship and accountability is mandatory for trust [7, 8].

Enabling technologies for semantic digitalization in a DTh include FAIR-by-design data practices with persistent identifiers, rich metadata, and controlled vocabularies captured at execution time, as well as metadata standards and ontologies that provide structured terminology for process descriptions, experimental context, and model lineage. Workflow and device integration are required to automate the capture of run context and intent, while provenance mechanisms such as audit trails and versioned artifacts preserve who generated which data, when, with which methods, models, and assumptions. Together, these components operationalize the reproducibility loop of bioprocess development, where experiments, models, and decisions continuously inform each other through explicit, machine-readable context.

It is worth noting the difference between a DTh and a digital twin (DTw). The latter is an adaptive representation of the system based on information artifacts or models that continuously resembles the structure, state, and behavior of a natural, engineered, or social system. For accuracy, this representation must be dynamically updated with data from its physical twin counterpart to guarantee its predictive capability for real-time decision-making in monitoring, optimization and control tasks. The bidirectional interaction between the virtual and the physical twins is central to the concept of DTw [1].

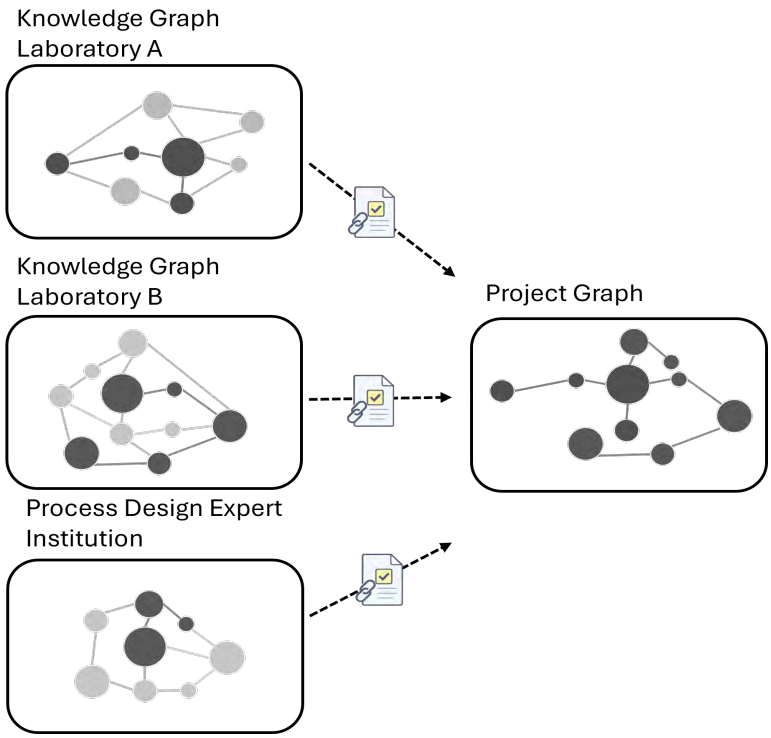

Figure 1. Federated knowledge integration across multiple local knowledge graphs. Selected nodes and relations from each local graph are curated and merged into a shared project graph that serves as the project's single source of truth. Access control, versioning, and provenance of the contributed graph fragments are governed by smart contracts.

Blockchain enabled IP provenance

Blockchain technology is a distributed digital ledger that is jointly maintained by multiple participants in a network. Instead of relying on a single central authority, the decentralized ledger is updated through a consensus process, and each new record is cryptographically linked to previous records. This structure makes the ledger tamper-proof, since altering past entries would require changing the entire chain consistently across the network. In collaborative research settings where trust is paramount, blockchain technology is therefore primarily useful as a shared, auditable record of contributions, IP ownership and transfer transactions, and full traceability of decision making [7, 8].

A blockchain-enabled CDTh can enable trusted data and knowledge exchange between research institutions and companies of different sizes by providing a provenance layer while keeping sensitive assets protected under local control. High volume experimental datasets, models, process designs and protocol packages remain off chain in controlled repositories such as institutional storage, encrypted object stores, or secure data spaces. The blockchain stores on chain only compact, verifiable evidence, descriptive metadata, including cryptographic hashes, timestamps, identifiers, and signed metadata that describe authorship, ownership, version, and intended usage context. Smart contracts provide the governance logic for automated collaboration by recording transfers as auditable events and by encoding usage conditions, such as attribution requirements and rules for version supersession, while retrieval of off-chain artifacts is mediated through secure authorization mechanisms [7, 8].

To exchange knowledge related digital objects in collaborative bioprocess engineering, such as a DTw of a bioreactor or a scale up ready process model package, IP NFTs can build on this infrastructure by representing each object as a uniquely identifiable digital asset on a consortium blockchain. The token acts as a persistent reference to the dataset and its contextual metadata, while transfer events record who created the asset, who shared it, and when a specific version was exchanged, enabling traceability of dataset lineage and authorship across institutional boundaries [7, 8]. Smart contracts can register new versions over time and trigger transparent update events that reference evolving off chain states, which is particularly relevant for DTw that are refined as new experimental results and model updates become available [1, 2].

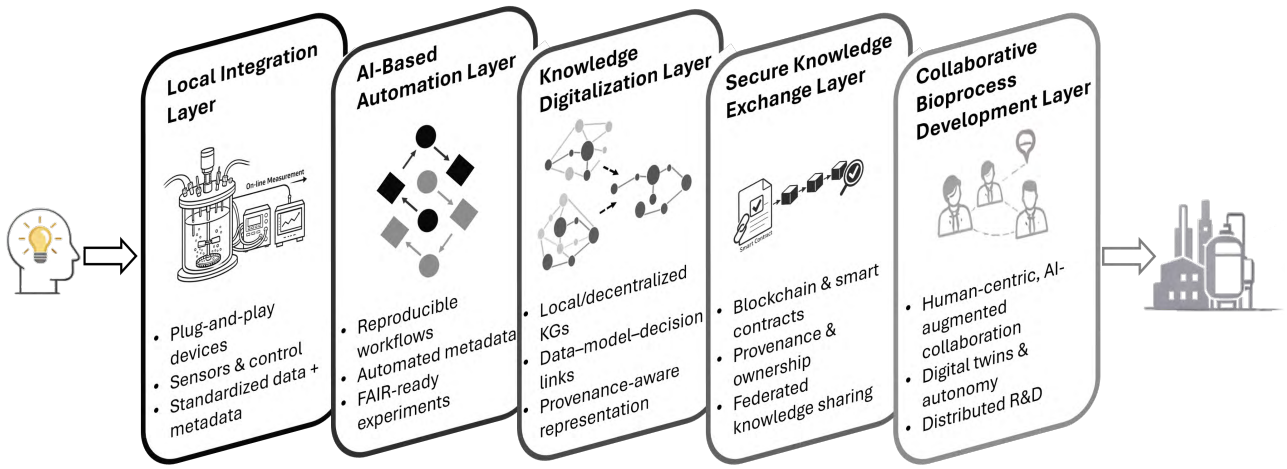

Figure 2. Five-layer architecture of a Cognitive Digital Thread for bioprocess engineering. Including (1) Local Integration (device connectivity, standardized data capture), (2) AI-Based Automation (reproducible workflows, automated metadata generation), (3) Knowledge Digitalization (local/decentralized knowledge graphs linking data–models–decisions with provenance awareness), (4) Secure Knowledge Exchange (governance of provenance, ownership, and access via blockchain and smart contracts), and (5) Collaborative Bioprocess Development (human-centric, AI-augmented collaboration). Together, the architecture enables a coherent digital pathway from scientific hypotheses to efficient biotechnological production.

From Silos and Lakes to Semantic Graphs

Digitalization in bioprocess engineering still operates in a fragmented landscape: experimental evidence is distributed across instruments, vendor-specific software, ELN and LIMS entries, spreadsheets, locally maintained scripts, and experimental inputs computed by several control and optimization algorithms [3]. These data silos constrain reuse and comparison of experiments, hinder systematic learning across projects, and complicate technology transfer because critical context and decision rationale remain scattered or implicit. Data lakes can centralize heterogeneous files and time series streams, yet they often lack sufficiently rich, standardized, and machine-actionable metadata, which limits interpretability for both domain experts and AI pipelines. In response, the trajectory toward semantic graphs treats contextualization as the primary design goal: Knowledge Graphs (KG) represent entities such as runs, strains, media, equipment, sampling events, analytics, models, and decisions as linked objects with explicit relationships and time-dependent context [4, 12]. This graph-based structure supports reproducibility as reconstruction, because the experimental and computational conditions that produced a result become queryable and transferable across scales and institutions [12].

At the project or facility level, CDThs can connect run-specific artifacts such as ELN records, instrument logs, bioreactor control data, Manufacturing Execution System events, and analytical Quality Control results into a continuous trace. When captured through standardized pipelines and FAIR-aligned metadata, each cultivation run, for example, is automatically linked to its design parameters, execution context, and outcomes. Aggregating multiple project-level threads into a KG then enables cross-project insight through context-rich queries, for example by retrieving: i) all experiments that used a given cell line, media formulation, ii) reactor configuration, iii) identifying recurring quality deviations across studies. In this way, the KG becomes a semantic layer that connects current and future projects while preserving traceability back to the original sources.

A natural extension of local KGs is the concept of a shared KG (Figure 1) that supports collaboration across laboratories, companies, and research institutions by linking knowledge products rather than merely exchanging raw data [3]. In this setting, partners contribute selected, contextualized graph fragments such as model versions, curated metadata, protocol templates, or decision rationales into a shared semantic layer that acts as an authoritative, lifecycle-spanning reference for a joint development effort. In practice, the shared graph becomes the project graph and serves as the single source of truth for all jointly relevant artifacts, while local graphs remain the source for private or institution-specific knowledge. To support IP-sensitive collaboration, the shared graph should include explicit nodes that capture IP-related metadata (e.g., ownership, licensing scope, usage constraints, attribution requirements) linked to the contributed knowledge fragments. A permissioned blockchain infrastructure can complement this approach by enabling targeted sharing through smart contracts,

which record immutable provenance of contributions and govern which specific knowledge artifacts are released, under which conditions, and to whom [7, 8]. This allows collaborators to retrieve precisely scoped information without granting access to an entire local knowledge base, preserving sovereignty while still enabling traceable reuse and attribution. Importantly, a shared KG does not require a single global uniform ontology across all participants. Interoperability can be achieved when shared artifacts are sufficiently contextualized and machine-readable, using persistent identifiers, explicit relationships, and semantic mappings that preserve meaning across heterogeneous ontologies, so that the intent and interpretation of the exchanged knowledge remain clear even when internal conceptual models differ [4].

This shift from siloed records and weakly annotated repositories toward semantically grounded local and shared graphs provides the informational substrate for a Cognitive Digital Thread, because cognition in bioprocess engineering depends on linking data to intent, assumptions, and decision provenance, and on enabling humans and AI systems to retrieve, compare, and reuse knowledge rather than isolated datasets.

CDTH IN A 5 LAYER ARCHITECTURE

To translate the previously outlined CDTh vision into a technically and organizationally implementable concept for bioprocess engineering, an architectural structure is required [10]. We propose a Five Layer Architecture (Figure 2) that describes how a CDTh can be constructed as a stepwise transformation from heterogeneous laboratory and production connectivity into a distributed, trustworthy, and human-centered decision environment. Its central premise is that cognition does not emerge from data availability or automation alone, but from the combination of reliable data acquisition, explicit contextualization, semantically grounded knowledge representation, governed collaboration, and human-in-the-loop decision processes. To operationalize these requirements, the CDTh is decomposed into a coherent dependency chain of modular capabilities.

Layer 1: Local Integration Layer as the foundation for interoperability and data integrity

The Local Integration Layer ensures plug-and-play compatibility of laboratory instruments, sensors, and control systems. It enables seamless data acquisition and device interoperability while capturing high-frequency process data and operational metadata in standardized formats. In bioprocess research environments, this is particularly challenging due to extreme heterogeneity across device vendors, device generations, communication protocols, and update capabilities. Legacy equipment is often difficult to integrate because older hardware cannot support modern interface standards. Even when connectivity is technically possible, heterogeneity in exported data products remains a critical issue. As of today, there is no general standard for communication between laboratory devices in biotechnology.

In addition, semantic inconsistencies in units, abbreviations, and references can cause misinterpretation even when data transport is syntactically correct. These issues tend to be more pronounced in academic laboratories than in industrial facilities, as research environments are often organically grown and frequently reconfigured to accommodate changing project needs.

For Layer 1 to support FAIR compliance, operational metadata must be captured systematically. This includes standardized instrument and method metadata, measurement context and uncertainty, and consistent time references. Robust identifiers for runs, workflows, and responsible agents are required to ensure machine readability and unambiguous linkage of records across the Digital Thread.

Layer 2: AI Based Automation Layer as reproducible experimental execution

Layer 2 provides reproducible experimental workflows and automated metadata capture to support FAIR compliance, but its function goes beyond documentation [12]. It turns experimental execution into a knowledge-generating process by explicitly linking experimental intent, model-based decision support, and reproducible workflows. As a result, experimental data are consistently accompanied by structured contextual information required for interpretation, comparability, and downstream reasoning.

A prerequisite for this layer is machine-readable data and context. Without a consistent semantic foundation, AI-enabled automation can execute tasks, but it cannot support trustworthy cognitive decisions or maintain verifiable control.

Layer 2 therefore requires machine-controllable experimentation through programmable devices, liquid handling systems, robotic assistants, or integrated laboratory infrastructures that connect cultivation with analytics via automated sampling and transport. A workflow management software (WMS) or workflow execution system can operationalize this requirement by orchestrating end-to-end experimental runs across devices via APIs, enforcing workflow templates and parameter sets, and automatically capturing execution logs and deviations as structured metadata. Reproducibility depends on minimizing manual interventions or capturing them explicitly in standardized form. Formal intent capture enables decision provenance, improves comparability across experiments, and supports explainable AI.

An essential element of this layer is the

reproducibility of AI workflows. Reproducibility in this context requires the reconstruction of objectives, decision rules, model versions, and decision outcomes under the same operational context. Semantic mechanisms must therefore provide explicit descriptions of objectives, context, and assumptions, while organizational mechanisms include workflow templates, parameter sets, and audit trails that act as metadata scaffolding to reduce manual burden while ensuring comparability and reproducibility.

Explainable AI is a core requirement, because AI components must provide transparent, human-understandable rationales for predictions and recommendations so that scientists and engineers can validate or challenge them. This supports human-in-the-loop oversight, improves failure analysis, strengthens governance of AI-assisted automation, and supports more robust transfer learning across development projects [13].

This layer is also relevant for rule-based systems, model predictive control, multi-agent systems, workflow engines, and orchestration tools that interface with Layer 1 via device APIs. Human-in-the-loop operation remains essential for safety, exception handling, and oversight of unintended system behavior [14].

Layer 3: Knowledge Digitalization Layer as the semantic backbone and the first truly cognitive layer

Layer 3 establishes local and decentralized KGs that represent entities such as process runs, organisms, materials, measured attributes, models, and decisions as nodes, with edges encoding relationships, context, and temporal dependencies [12]. By linking data, models, and decisions to underlying metadata, and assumptions, the graph makes experimental outcomes interpretable beyond isolated datasets and enables decision provenance across development steps. Within a CDTh, Layer 3 forms the semantic backbone that transforms bioprocess data into reproducible, explainable, and transferable knowledge.

A critical requirement is that KGs must be dynamic [1]. Bioprocess knowledge evolves continuously as new experiments are executed, model structures are updated, and assumptions are revised. Dynamic graphs support event-driven updates, time-dependent states, and evolving validity of models and decisions, so that the knowledge state relevant to a specific result can be reconstructed and compared against later revisions. Typical use cases include representing experimental trajectories, decision paths with their constraints, and explicit model version histories linked to the evidence that motivated updates [11].

Layer 3 also provides the interface between human reasoning and AI systems. It supports explainable AI by providing structured context, enables domain-specific querying by making relationships between evidence, models, and decisions explicit. This reduces redundancy by allowing teams to locate and reuse previously generated knowledge that would otherwise remain trapped across repositories, and it supports hypothesis generation by exposing patterns across runs, organisms, and operating regimes.

Semantic querying across projects enables comparability and transfer learning, for example by retrieving analogous processes or comparable organisms via structured relationships. Importantly, federation does not imply centralized control or raw data sharing. Instead, it relies on selective exchange of contextualized graph fragments and explicit semantic alignment where knowledge is shared, while local sovereignty and access control remain intact.

Layer 4: Secure Knowledge Exchange Layer as an optional trust and governance layer

In collaborations with limited institutional trust, intellectual property sensitivity, or strict audit requirements, Layer 4 adds a permissioned blockchain and smart contracts to provide verifiable provenance, attribution, and enforceable governance for shared knowledge artifacts across organizational boundaries. The purpose of this layer is not to store knowledge itself, but to make the terms and history of knowledge exchange tamper-evident and auditable.

When required, the core architectural principle of this layer is the separation of trust from knowledge. KGs store content, semantics, and context, while Layer 4 provides registries, signatures, governance and verification mechanisms. This preserves local sovereignty over sensitive assets while enabling traceable reuse in joint projects [7, 8].

A particularly important function of this layer is KG versioning. Because KGs are dynamic, context-dependent, and decision-relevant, it is critical to identify which knowledge state was valid at the time of a decision, which model version supported an experiment, and which assumptions were later revised [11].

Smart contracts operationalize these requirements by acting as version registries and trust anchors. They identify registered graph states or exchanged artifacts via cryptographic hashes and immutable metadata (timestamps, authorship, project scope), while keeping large datasets and models off-chain in controlled repositories. Smart contracts can further encode and enforce access rights and usage policies such as automatic attribution recording, purpose limitation, and rules for version supersession, thereby supporting reproducibility, accountability, and explainable reasoning despite ongoing knowledge evolution.

By making usage conditions explicit and enforceable, Layer 4 complements social trust with machine-

verifiable traceability and supports collective intelligence without forcing loss of control.

Layer 5: Collaborative Bioprocess Development Layer as the socio technical locus of system cognition

Finally, layer 5 integrates perception, contextualization, goal orientation, learning, and human interaction into actionable decision processes, making cognition operational. Lower layers provide the prerequisites as structured data and metadata (Layer 1), reproducible and standardized workflows (Layer 2), explicit knowledge representation (Layer 3), and governed collaboration where required (Layer 4). Layer 5 combines these elements into a shared decision space that supports human-centered, AI-augmented collaboration across distributed laboratories while preserving accountability and reproducibility.

At this level, the CDTh connects automated experimentation systems, DTw environments, simulation workflows, and expert reasoning interfaces across laboratories and sites. Cross-laboratory orchestration is enabled by standardized workflows from Layer 2, federated KGs from Layer 3, and trust mechanisms from Layer 4, allowing joint objectives and information flows to be aligned without enforcing central control over local assets.

A defining feature of Layer 5 is the human-in-the-loop operation. AI systems contribute recommendations, alternatives, and uncertainty estimates, while humans retain decision authority and responsibility. Cognition emerges as co-cognition rather than full autonomy. Humans validate model assumptions, interpret results, and override system decisions. Layer 5 operationalizes intentionality by enabling objectives to be formulated, revised, and prioritized, and by making trade-offs explicit. It supports context-sensitive decision-making across organizational boundaries by combining distributed knowledge with local expertise and by accounting for uncertainty and knowledge gaps. Learning across time is enabled by making prior project experience retrievable and by embedding decisions in versioned knowledge contexts.

Because Layer 5 is the primary user-facing locus of the CDTh, it must also reduce cognitive burden rather than amplify it. Role-specific views, prioritization, and consistent templates are required to build digital dexterity and avoid transformation fatigue, ensuring that transparency and traceability translate into faster, more robust decisions instead of information overload.

CONCLUDING REMARKS

In this review, we positioned Bioprocessing 5.0 as a human-centered, knowledge-driven, and collaborative development strategy that responds to the unmet promises of Industry 4.0 in bioprocess engineering [5, 6]. We conclude that this vision will not be fully realized without a CDTh that unifies data, knowledge, and decision provenance across the bioprocess lifecycle [11]. The CDTh enables end-to-end traceability of decisions by explicitly linking data, experimental context, models, and human intent, and it supports reproducible experimental and computational workflows through machine-actionable, audit-ready execution traces. Its semantic backbone, implemented through KGs, enables interoperable models that can be transferred, compared, and reused across projects and scales, while explainable AI is essential to make algorithmic recommendations transparent and thereby trustworthy for scientific validation and regulatory adoption [4, 12, 13]. To extend these capabilities beyond organizational boundaries, blockchain-based mechanisms can serve as a trust anchor for collaboration, where smart contracts provide verifiable provenance of intellectual property, authorship, and contributions without requiring the sharing of raw data [7, 8]. With the combined use of FAIR-aligned data platforms, semantic KGs, explainable AI, and collaborative software agents embedded in human-in-the-loop decision environments, the field can transcend current limitations in reproducibility, scalability, and knowledge transfer [14]. Taken together, this review provides a framework for future trends in bioprocess digitalization and charts a path toward an Industry 5.0 paradigm that blends human creativity with machine precision in collaborative bioprocess development [5, 6].

ACKNOWLEDGEMENTS

We gratefully acknowledge the financial support of the German Federal Ministry of Research, Technology and Space (03DPC0101A – BioBlock StEDI).

AUTHOR IDENTIFIERS

Author ORCIDs:
Karsten JM: 0009-0005-2004-7345
Martínez E: 0000-0002-2622-1579
Cruz Bournazou MN: 0000-0001-9461-4414

REFERENCES

1. Akroyd J, Mosbach S, Bhave A, Kraft M. Universal digital twin - a dynamic knowledge graph. DCE 2: (2021). https://doi.org/10.1017/dce.2021.10
2. Su C, Tang X, Han Y, Wang T, Jiang D. Cognitive digital twin in manufacturing process: integrating the knowledge graph for enhanced human-centric industry 5.0. International Journal of Production Research :1-22 (2024). https://doi.org/10.1080/00207543.2024.2435583

3. Erickson J, Baker J, Barrett S, Brady C, Brower M, Carbonell R, Charlebois T, Coffman J, Connell?Crowley L, Coolbaugh M, Fallon E, Garr E, Gillespie C, Hart R, Haug A, Nyberg G, Phillips M, Pollard D, Qadan M, Ramos I, Rogers K, Schaefer G, Walther J, Lee K. End?to?end collaboration to transform biopharmaceutical development and manufacturing. Biotech & Bioengineering 118:3302-3312 (2021). https://doi.org/10.1002/bit.27688
4. Uschold M, Gruninger M. Ontologies and semantics for seamless connectivity. SIGMOD Rec. 33:58-64 (2004). https://doi.org/10.1145/1041410.1041420
5. Leng J, Sha W, Wang B, Zheng P, Zhuang C, Liu Q, Wuest T, Mourtzis D, Wang L. Industry 5.0: prospect and retrospect. Journal of Manufacturing Systems 65:279-295 (2022). https://doi.org/10.1016/j.jmsy.2022.09.017
6. Barata J, Kayser I. Industry 5.0 – past, present, and near future. Procedia Computer Science 219:778-788 (2023). https://doi.org/10.1016/j.procs.2023.01.351
7. Van Woensel W, Seneviratne O. Semantic interoperability on blockchain by generating smart contracts based on knowledge graphs. Blockchain: Research and Applications 7:100320 (2026). https://doi.org/10.1016/j.bcra.2025.100320
8. Wang S, Huang C, Li J, Yuan Y, Wang FY. Decentralized construction of knowledge graphs for deep recommender systems based on blockchain-powered smart contracts. IEEE Access 7:136951-136961 (2019). https://doi.org/10.1109/access.2019.2942338
9. West TD, Pyster A. Untangling the digital thread: the challenge and promise of model?based engineering in defense acquisition. INSIGHT 18:45-55 (2015). https://doi.org/10.1002/inst.12022
10. Madni AM, Sievers M. Model?based systems engineering: motivation, current status, and research opportunities. Systems Engineering 21:172-190 (2018). https://doi.org/10.1002/sys.21438
11. Wu S, Wang G, Lu J, Huang J, Qiao J, Yan Y, Kiritsis D. Cognitive digital thread tool-chain for model versioning in model-based systems engineering. Advanced Engineering Informatics 67:103490 (2025). https://doi.org/10.1016/j.aei.2025.103490
12. Bai J, Mosbach S, Akroyd J, Lapkin AA, Kraft M. From Platform to Knowledge Graph: Evolution of Laboratory Automation, J. Amer. Chem. Soc., 2:2, 292–309 (2022). https://doi/10.1021/jacsau.1c00438
13. del Águila Escobar RA, del Carmen Suárez-Figueroa M, Fernández López M, Villazón Terrazas B. Bridging text and knowledge: explainable AI for knowledge graph classification and concept map-based semantic domain discovery with OBOE framework. Applied Sciences 15:12231 (2025). https://doi.org/10.3390/app152212231
14. David J, Lobov A, Jarvenpaa E, Lanz M. Enabling the digital thread for product aware human and robot collaboration - an agent-oriented system architecture. 2021 20th International Conference on Advanced Robotics (ICAR) :1011-1016 (2021). https://doi.org/10.1109/icar53236.2021.9659352

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Safe and Sustainable by Design Pharmaceuticals through Combined Computer-Aided Retrosynthesis, Techno-Economic Analysis, and Life Cycle Assessment

Shang Gao[a], Brahim Benyahia[a*]

[a] Loughborough University, Department of Chemical Engineering, Epinal Way, Loughborough, Leicestershire, LE11 3TU, UK
* Corresponding Author: B.Benyahia@lboro.ac.uk

ABSTRACT

Recent advances in computer-aided retrosynthesis (CAR), flow chemistry, and continuous manufacturing collectively offer new opportunities to enable environmentally sustainable development and manufacturing practices across the pharmaceutical development and manufacturing value chain. However, the implementation of these methods and technologies remains scattered and fragmented, preventing full realization of their potential to address one of the most urgent needs in the pharmaceutical and related sectors. This work introduces a holistic digital framework for the design and optimization of an end-to-end manufacturing process for paracetamol (acetaminophen). The framework integrates Green-by-Design synthetic and purification routes of the active pharmaceutical ingredient (API) aims to deliver cost efficiency and robust quality, safety, and environmental sustainability assurance. The approach integrates AI-driven CAR with plant wide modelling, Techno-Economic Analysis (TEA), and prospective cradle-to-gate prospective Life Cycle Assessment (LCA) to evaluate designs options and greener, chemically feasible, and more selective synthetic pathways. An end-to-end mathematical model is implemented in gPROMS Formulated Products® to optimize operating and design parameters and generate inventory data for TEA/LCA. The framework also embeds Quality by Digital Design (QbDD) to identify CPPs and CMAs that influence CQAs, enabling definition of a robust design space. Finally, a multicriteria decision-aiding approach is implemented to determine the best trade-offs among yield, productivity, resource efficiency, cost, and environmental performance, providing a transferable Safe-and-Sustainable-by-Design (SSbD) methodology for developing greener and more resilient pharmaceutical manufacturing systems.

Keywords: Computer-Aided Retrosynthesis, Life Cycle Analysis, Technoeconomic Analysis, Modelling and Simulations, Safe-and-Sustainable-by-Design, Quality by Digital Design

INTRODUCTION

Pharmaceutical manufacturing is among the most solvent- and utility-intensive industries and is frequently characterized by very high waste-to-product ratios. The development and manufacturing of pharmaceutical products is a long complex and fragmented journey, characterized by early-stage discovery and clinical trials, followed by process development, scale-up and optimization, regulatory validation, and commercial manufacturing. Reducing the environmental footprint of pharmaceutical manufacturing require the development of changes that cut across synthesis, separations, and solvent recovery. Early LCA work on API manufacture highlighted that solvents could dominate material usage and that solvent production and waste management can drive a large share of life-cycle emissions, making solvent choice/recycle central to greener manufacturing [1-4]. More recent investigations confirm that, across pharmaceutical LCAs, energy consumption and synthetic and purification processes are frequently the key contributors to the overall environmental impacts, while noting persistent limitations in inventory transparency and system boundary consistency [5, 6].

Advances in computer-aided retrosynthesis (CAR) have

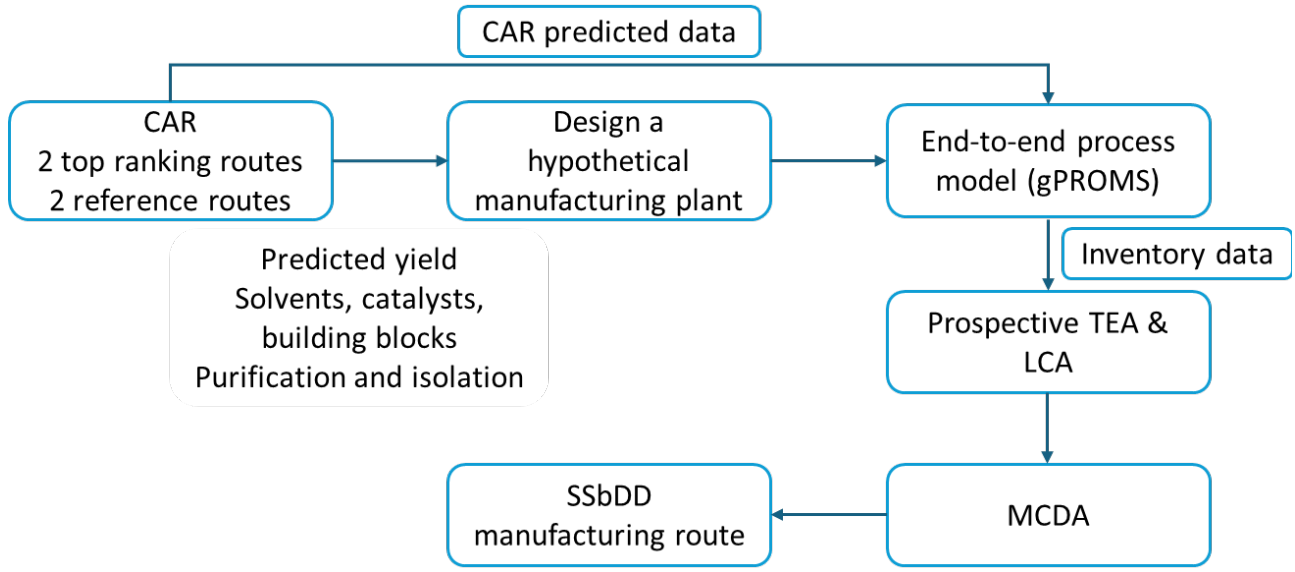

Figure 1. Overview of the proposed SSbD workflow integrating CAR, process modelling, QbDD, TEA and LCA. End-to-End Process Modelling

enabled rapid generation of feasible and Green-by-Design synthetic routes and systematic route comparison, moving beyond expert-based and trial-and-error based synthetic route design [5-6]. On the manufacturing side, continuous flow and telescoped processing have demonstrated clear potential to improve safety and footprint by enabling tighter control, improved heat/mass transfer, and avoidance of isolating unstable intermediates; for paracetamol specifically, multistep continuous-flow synthesis and automated optimization studies illustrate both feasibility and the value of algorithmic design-space exploration [7]. Renewables-oriented thinking is also gaining traction in pharmaceutical manufacturing, with recent work explicitly proposing and evaluating renewable routes using green chemistry analysis [8]. However, current practice still largely treats synthetic route design, process development, and sustainability/techno-economics (LCA/TEA) as separate activities, limiting the ability to steer route selection using end-to-end sustainability and robustness metrics early in development.

The emergence of digital frameworks such as Quality by Digital Design (QbDD) is transforming the pharmaceutical manufacturing landscape, placing model-based and data-driven methods at unprecedented levels of importance and integration. [16]

This work aims to address the critical gaps highlighted above by developing and validating an integrated SSbD digital framework that couples: (i) CAR-assisted Green-by-Design synthetic route generation and ranking, (ii) end-to-end synthesis and purification process modelling, and (iii) integrated prospective TEA and cradle-to-gate LCA, while embedding QbDD concepts to ensure built-in quality and safety assurance. The proposed approach aims to connect design space definition with sustainability performance. By linking CAR decisions to process simulation outputs and inventory generation, the framework enables multi-objective optimization over yield, resource efficiency, cost, and environmental impacts, supporting early, data- driven selection of greener and more robust paracetamol manufacturing solutions. Paracetamol (acetaminophen), a high-volume essential medicine, is selected to validate the proposed framework. Its substantial environmental footprint linked to large-scale production makes it a suitable case study for demonstrating the value of the holistic methodology.

METHODOLOGY

The proposed end-to-end SSbD workflow is captured by the flow diagram shown in figure 1. Paracetamol, the top-selling over-the-counter pharmaceutical with a global market volume of nearly 275 tonnes and a value of USD 876.48 million in 2024, is used to demonstrate the proposed framework. The proposed end-to-end manufacturing interstates process of Paracetamol involves aromatic functionalization and acetamidation chemistry, with well-known trade-offs between selectivity, solvent usage, and purification effort. Multiple synthesis concepts are therefore feasible, making paracetamol particularly suitable for route screening and comparative sustainability analysis. Candidate synthesis routes are structured into four route families as shown in Table 1:

(i) Route A: CAR (SYNTHIA) baseline / conventional-like route (used as an internal benchmark)
(ii) Route B: CAR (SYNTHIA) SSbD-preferred route

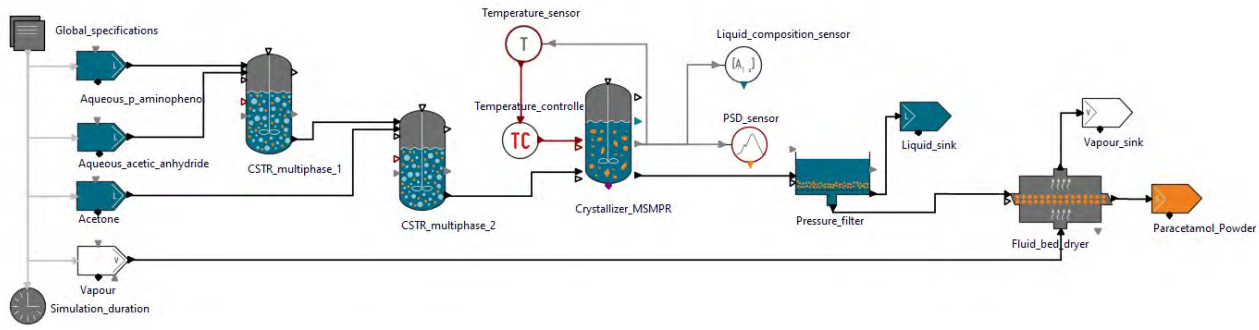

Figure 2 End-to-end flow diagram of paracetamol manufacture

(screened for lower expected separation burden and improved operability)
(iii) Route C: Renewable phenol-based route [10]
(iv) Route D: Bio-waste-derived concept (β-pinene-based continuous route) [13]

Table 1: Selected synthesis routes for paracetamol

Routes	# Steps	Yield	Solvents
A	2	88%	Acetic anhydride
B	2	92%	Acetic acid
C	3	85%	Ammonium acetate
D	4	90%	Methanol

These routes form the basis for CAR-based screening and subsequent process modelling within a common superstructure. CAR is used as the entry point of the workflow to generate feasible retrosynthetic pathways for paracetamol using SYNTHIA®, with pathway exports supporting transparent route documentation [8, 9]. In contrast to conventional retrosynthesis, route evaluation is extended beyond chemical feasibility to include sustainability-relevant attributes. Each candidate route is assigned a composite screening score based on normalized indicators representing step count, feedstock origin, reagent/solvent hazard proxies, telescoping potential, a selectivity-related impurity risk proxy, and compatibility with continuous processing. The highest-ranked routes are retained for detailed process modelling, ensuring direct propagation of early-stage design decisions into downstream simulation, TEA and LCA.

The selected total paracetamol synthetic routes are integrated to an end-to-end manufacturing superstructure developed in gPROMS Formulated Products® as shown in figure 2. The superstructure comprises raw material, solvent, catalyst feed streams, reaction, quenching and conditioning, API isolation, crystallization, solid–liquid separation, drying, formulation and solvent recovery where feasible. Route-specific differences are represented through variations in reaction sequence, solvent systems, operating conditions, and separation difficulty. Reaction selectivity is captured using impurity load proxies, which directly affect yield losses and solvent demand in purification. This abstraction enables quality-sustainability coupling while remaining suitable for early-stage design studies. The model generates consistent mass and energy balances for all unit operations, producing inventory-quality data for the prospective TEA and LCA.

Quality by Digital Design

To deliver built-in quality assurance for safety, a QbDD approach is embedded within the modelling framework to systematically identify relationships between CPPs, CMAs, and CQAs for paracetamol manufacture. CQAs considered at the API stage include product purity, levels of impurities, and residual solvent content. Candidate CPPs and CMAs include reaction temperatures, residence times, solvent choice, crystallization conditions, and solvent recycling ratios. Sensitivity analyses are used to quantify their influence on CQAs, yield, and resource consumption. Variations that degrade quality typically increase solvent usage and energy demand, thereby linking quality robustness directly to sustainability performance.

Integrated Techno-Economic Analysis and Life Cycle Assessment

Prospective TEA and LCA are addressed based on the inventory data extracted from the gPROMS simulations. The TEA captures the capital and operating expenditures, with particular emphasis on solvent consumption, recovery efficiency, and energy demand.

A cradle-to-gate life cycle assessment (LCA) is performed to quantify the environmental impacts of each manufacturing alternative associated with shortlisted total synthetic routes. The proposed LCA is consistent with the ISO 14040/44 standards, covering four key steps: goal and scope definition, inventory analysis, impact

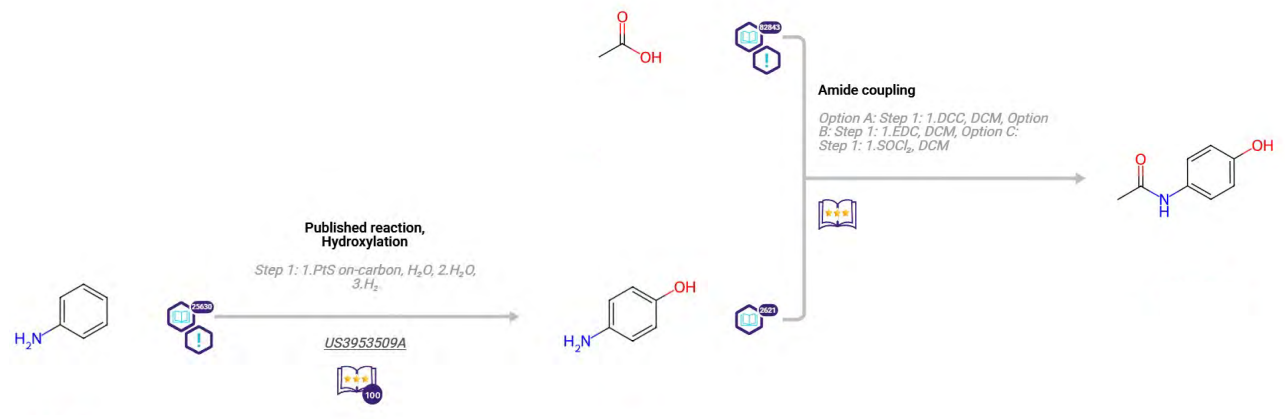

Figure 3. Route B from SYNTHIA® retrosynthesis

assessment, and interpretation. Environmental impacts were assessed across multiple categories, including global warming potential, energy use, and waste generation.

Goal and Scope

Perform the LCA for all case model simulations with the functional unit of 1 kg paracetamol products.

Life cycle inventory

Based on the process simulation outputs, detailed inventory data were gathered for each route. This data encompasses material usage, energy consumption, waste generation, and emissions at each processing stage.

Impact assessment

The impact assessment is based on the method ReCiPe 2016 available in SimaPro with the Ecoinvent 3 environmental footprint database. The ReCiPe 2016 captures three endpoint categories (human health, ecosystems, and resources) and seventeen midpoint categories. These different categories provide a comprehensive basis to compare the different design options.

Interpretation of results

Finally, the results obtained from the previous steps are interpreted and the environmental hotspots are identified.

The overall LCA was structured with a cradle-to-gate system boundary for API manufacture. The boundary includes raw material production (reagents, solvents and auxiliaries), on-site utilities (electricity/steam and cooling), and treatment of process waste streams associated with synthesis, purification, and solvent recoveryAPI production: The environmental impacts associated with the synthesis of the API, including raw material extraction, energy use, solvent recovery and waste generation during API production.

Multicriteria Decision-Making

A multicriteria decision aiding approach is implemented to identify the best compromises between environmental impacts, economic performance, yield, and robustness of CQAs. Integrating this approach to SSbD delivers a recommended paracetamol synthesis route together with a robust operating and design space that satisfies quality constraints while minimizing economic and environmental burdens.

RESULTS AND DISCUSSION

CAR route screening and selection

CAR was applied using SYNTHIA® Retrosynthesis to generate feasible synthetic paracetamol pathways along with transparent route documentation. Route A represents a SYNTHIA-generated baseline pathway selected primarily based on the feasibility criteria (short step count and readily available starting materials), while Route B represents a SYNTHIA-generated SSbD-preferred pathway selected using additional screening for lower expected separation burden (telescoping potential and reduced impurity risk proxy) and improved operability. Additionally, these routes are benchmarked against two published synthetic routes namely: Route C (renewables-based routes) [10] and Route D (bio-waste β-pinene concept) [13]. These four routes were each translated into an integrated synthesis–purification superstructure for consistent simulation, inventory generation, and TEA/LCA based comparison.

- Route A (CAR baseline): SYNTHIA® retrosynthesis

Table 2: KPIs for different routes

Route	Fresh solvent use (kg/kg API)	Energy use (MJ/kg API)	GWP (kg CO2-eq/kg API)	Production cost (Dollars/kg API)
Route A	25	120	18	3
Route B	15	85	11	2.55
Route C	18	95	13	2.85
Route D	14	80	10	2.7

route (2 steps) selected as a feasibility/benchmark-like pathway; expected higher purification burden and solvent demand due to conservative isolations.

- Route B (CAR SSbD): SYNTHIA® retrosynthesis route (3 steps) down-selected through sustainability-based screening (telescoping potential and reduced impurity-risk proxy); expected reduced separation burden and lower solvent/energy intensity as shown in Figure 3.

End-to-end process-based performance

The integrated process model provides consistent comparison of synthesis + purification under a common functional unit (kg API). In the current dataset, overall yields range from 85–92%, indicating that route choice affects not only productivity but also purification recovery efficiency:

- Route A (CAR baseline): 88%
- Route B (CAR SSbD): 92%
- Route C (renewables literature): 85%
- Route D (β-pinene literature): 90%

The higher yield exhibited by Route B indicates that either the route exhibits enhanced selectivity or /and reduced losses during isolation. Conversely, Route C (renewables-based) shows relatively lower yield which may be due to the nature of renewable feed stocks. While yield strongly correlated to the process economics, it should not be treated as an isolated objective in SSbD, as a route with marginally lower yield can still be preferred if it substantially reduces solvent intensity and utilities, because solvent/utility hotspots often dominate cradle-to-gate impacts and cost. Hence, the SSbD decision requires multi-criteria decision aiding to address the conflicting or competing nature of these objectives. Other KPIs including solvent intensity, energy intensity, environmental intensity and economic intensity are as shown in Table 2.

Solvent intensity

Solvent intensity is an emergent property of selectivity/impurity formation, crystallization efficiency, and recycle strategy. Routes that reduce impurity load reduce the number/volume of wash and recrystallization steps and therefore reduce both solvent purchase and solvent recovery duty. As a direct indicator of separation burden and a major driver for both OpEx and environmental impacts (solvent production + waste handling + recovery duties). The four routes show strong differentiation as outline below:

- Route A: 25 kg solvent/kg API
- Route B: 15 kg solvent/kg API (40% reduction vs A)
- Route C: 18 kg solvent/kg API (28% reduction vs A)
- Route D: 14 kg solvent/kg API (44% reduction vs A)

These reductions provide critical insight: the best-performing routes (B and D) are those that are expected to require less intensive recrystallization/washing and/or have more effective recycle structures. For CAR, this is where Route B demonstrates tangible value, CAR selection can target reaction sequences and solvent sets that reduce downstream separation intensity.

Energy intensity

Energy intensity (MJ/kg API) integrates heating/cooling duties for reaction temperature control, crystallization cooling, solvent recovery, and drying. The four routes in the dataset show as below:

- Route A: 120 MJ/kg API
- Route B: 85 MJ/kg API (29% reduction vs A)
- Route C: 95 MJ/kg API (21% reduction vs A)
- Route D: 80 MJ/kg API (33% reduction vs A)

The close alignment between solvent reductions and energy reductions is expected because solvent recovery and drying duties often scale with solvent throughput. Hence, reducing solvent makeup yields a compounding benefit: lower fresh solvent demand reduces recovery loads, which reduces utility demands and associated emissions. In SSbD optimization, solvent intensity is not only a material KPI; it is also a strong proxy for energy intensity and therefore for climate impacts. This justifies (i) using solvent-related screens already at

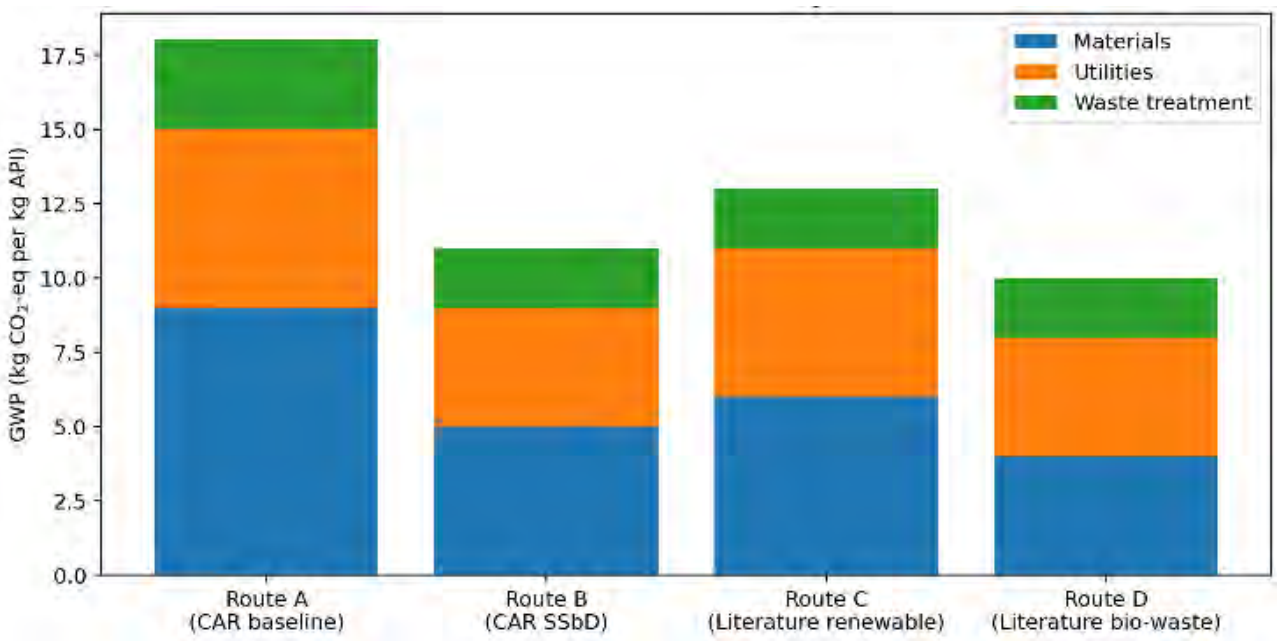

Figure 4. GWP hotspot breakdown

the CAR stage and (ii) prioritizing route selections compatible with solvent recycle/recovery strategies.

Environmental intensity

With the predicted inventories as inputs, the hotspot breakdown can be analysed by converting gPROMS outputs (fresh reactants/solvents, steam/electricity/cooling, waste flows) into LCI flows and multiplying by characterization factors (ReCiPe 2016) to obtain the contribution scores. This decomposition is particularly important for CAR-derived Route B: if Route B reduces GWP primarily via utilities, it suggests the improvement arises from reduced recovery duty; if it reduces GWP primarily via materials, it suggests the route reduces fresh solvent/reagent demand. For climate change impacts, the draft cradle-to-gate GWP comparison shows a clear route dependence as below:

- Route A: 18 kg CO_2-eq/kg API
- Route B: 11 kg CO_2-eq/kg API (~39% reduction vs A)
- Route C: 13 kg CO_2-eq/kg API (~28% reduction vs A)
- Route D: 10 kg CO_2-eq/kg API (~44% reduction vs A

To move beyond ranking and support actionable SSbD decisions, hotspot interpretation is essential. Another effective reporting structure decomposes total GWP into materials, utilities, and waste treatment contributions as shown in Figure 4: The baseline CAR Route A exhibits the highest total GWP, driven primarily by the materials contribution, consistent with higher fresh input requirements (notably solvent and reagent supply chains) and the propagation of separation intensity into upstream burdens. In contrast, the CAR Route B (SSbD-preferred) achieves a lower total GWP, with reductions observed in both materials and utilities components, indicating that improvements are not solely attributable to energy efficiency but also to reduced fresh material demand (e.g., solvent makeup) and associated recovery duties.

Among the literature comparators, Route D (bio-waste derived concept) shows the lowest total GWP in this breakdown, reflecting reduced contributions from both materials and utilities relative to the baseline. Route C (renewables-based concept) displays intermediate performance, suggesting that feedstock substitution alone may not guarantee minimum impacts unless coupled with process designs that reduce solvent throughput and utility demand. Overall, the environmental intensity indicates that end-to-end impacts are governed by coupled effects of solvent/material intensity and utility demand, supporting the SSbD rationale for combining CAR route selection with process modelling and integrated TEA/LCA.

Techno-economic comparison

Techno-economic analysis (TEA) was applied to quantify the manufacturing cost implications of route selection under a consistent cradle-to-gate system boundary (API synthesis + purification) and a common

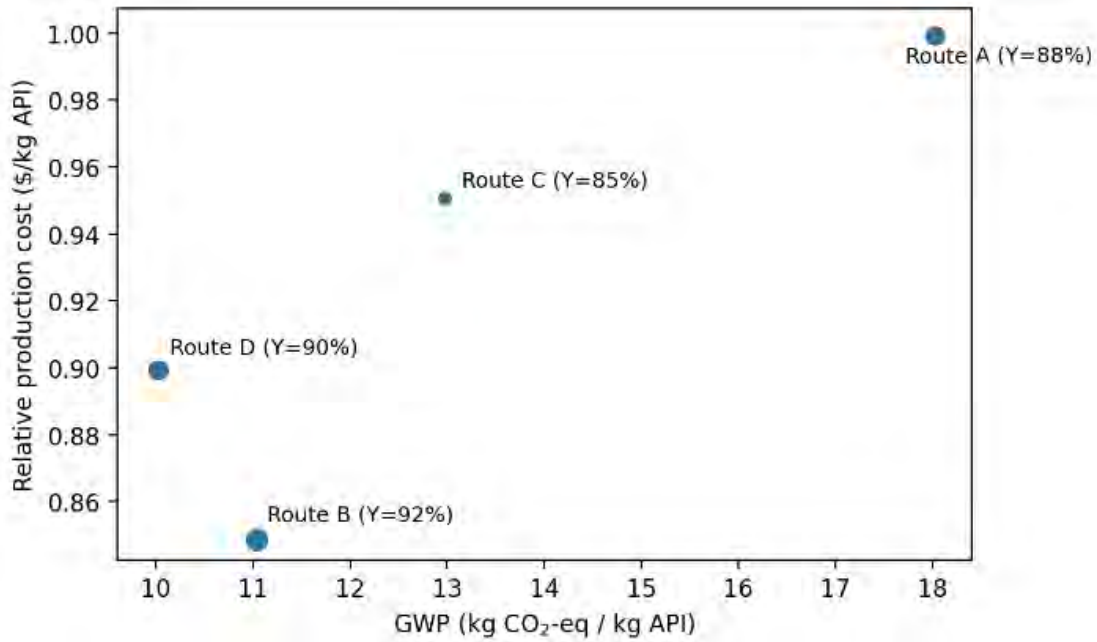

Figure 5. Comparison of the different scenarios in the KPI space.

functional unit (cost per kg paracetamol API). To avoid over-interpretation at conceptual stage, results are reported on a relative cost basis normalised to Route A (CAR baseline) (= 1.00). The comparative TEA indicates that route selection yields meaningful cost differences, with the ranking:

- Route A (CAR baseline): 1.00
- Route B (CAR SSbD): 0.85
- Route C (literature renewable): 0.95
- Route D (literature bio-waste): 0.90

The dominant cost drivers are consistent with the process-level KPIs reported in previous sections. Specifically, solvent management (fresh solvent purchase, solvent recovery, and solvent losses to waste streams) and utilities (steam/electricity for heating, cooling, and drying) are the primary contributors to operating expenditure (OpEx), while recovery and drying capacity influence capital expenditure (CaPEx) proxies. In this context, Route A exhibits the highest relative cost due to its elevated fresh solvent makeup and energy intensity, which increase both (i) recurring material and utility costs and (ii) equipment sizing requirements for solvent recovery and drying.

In contrast, the CAR Route B shows the lowest cost in the current assessment, consistent with (i) reduced solvent intensity (40% lower than Route A) and (ii) reduced energy demand (≈29% lower than Route A), which together decrease solvent purchasing, recovery duty, and drying load. Importantly, Route B also shows the highest yield), improving effective throughput per unit feed and lowering cost per kg API. These results support the SSbD premise that CAR-guided route selection can deliver economically favourable alternatives when the screening explicitly targets reduced separation burden and improved end-to-end efficiency.

The two literature comparators provide additional context for route selection. Route D (bio-waste derived) is economically competitive (0.90) while exhibiting the lowest GWP, indicating that sustainability-oriented route concepts can remain cost-viable when solvent and utility demands are constrained. Route C (renewables-based) shows intermediate cost performance (0.95), suggesting that feedstock substitution alone may not ensure minimum manufacturing cost unless accompanied by downstream separation designs that minimise solvent throughput and recovery duty.

Multi-criteria decision aiding

Figure 5 shows the conflicting nature of the multicriteria decision problem associated with the proposed SSbD using 2 critical KPIs: the relative production cost and cradle-to-gate GWP, with each point annotated by overall yield.

In this assessment, Route A occupies the unfavourable region with the highest cost (normalized to 1.00) and highest climate impact (18 kg CO_2-eq/kg API) and is therefore dominated by multiple alternatives. In contrast, Route B is positioned in a strongly favourable region, combining the lowest relative cost (0.85) with a substantially reduced GWP (11 kg CO_2-eq/kg API) and the highest

yield (92%). The location of Route B indicates that CAR-informed route selection can produce pathways that improve both economic and environmental performance simultaneously, consistent with the SSbD hypothesis that reductions in separation intensity (solvent and recovery duty) deliver compounding benefits across TEA and LCA.

Among the published reference synthetic routes, Route D exhibits the lowest GWP (10 kg CO_2-eq/kg API) while maintaining competitive cost (0.90) and high yield (90%), placing it close to the lower-left region of the MCDA space. Route C remains intermediate in both cost and environmental performance. Importantly, Figure 6 indicates a non-trivial trade-off between Route B and Route D: Route B offers the lowest cost and highest yield, whereas Route D offers the lowest climate impact. This provides a clear SSbD decision interpretation: selection depends on the weighting of economic vs environmental objectives and on robustness considerations. In practice, such a MCDA representation is preferable to single-metric ranking because it transparently communicates which improvements are achievable without sacrifice and where genuine trade-offs remain.

While Figure 5 provides a two-objective SSbD decision surface (cost vs climate impact), route selection for pharmaceutical manufacture typically requires consideration of multiple performance dimensions simultaneously, including yield, solvent intensity, and utility demand. Figure 6 therefore provides a normalized, multi-metric summary across five key indicators (Yield, Solvent, Energy, GWP, and Cost), with values normalized such that 1 indicates best performance among the compared routes. This representation supports rapid identification of "balanced" candidates and clarifies whether a route's apparent advantage is narrow (e.g., only cost) or systemic (across multiple indicators).

The radar plot indicates that Route A performs worst (near-zero normalized scores) in solvent intensity, energy demand, GWP, and cost, confirming that it serves primarily as a baseline reference rather than an optimal SSbD choice. Route B exhibits consistently high normalized scores across all indicators, reflecting a comparatively balanced pathway that is simultaneously high-yield and low in both operating burden (solvent/energy) and system-level metrics (cost/GWP). Route D shows particularly strong normalized performance in solvent and energy intensity and in GWP, with somewhat weaker (but still competitive) cost performance relative to Route B. Route C shows intermediate performance and highlights that renewable feedstock concepts may require additional process intensification or separation optimization to reach best-in-class end-to-end metrics.

From a QbDD perspective, Figure 6 can be interpreted as evidence that routes with reduced separation intensity (lower solvent and energy) are inherently favourable not only for LCA and TEA but also for quality robustness. In practice, reduced purification demand generally increases robustness margins because fewer or less severe isolation steps reduce sensitivity to upstream variability (e.g., selectivity fluctuations propagating into impurity burden). Within the proposed SSbD framework, this motivates a coupled optimization strategy in which CPPs/CMAs (e.g., reagent stoichiometry, temperature, crystallization endpoint, recycle ratio) are tuned to satisfy CQA constraints while simultaneously minimizing solvent and energy penalties. Therefore, multi-metric results such as Figure 6 provide a compact screening tool to identify routes that are likely to admit a robust design space under QbDD constraints, rather than merely optimizing nominal performance. A preliminary QbDD screening for Route B indicates a constrained operating window in which near-stoichiometric acetic acid usage and high 4-aminophenol purity jointly maximize yield, highlighting the coupling between input material quality (CMA) and stoichiometric control (CPP) in determining CQA performance as shown in Figure 7.

CONCLUSIONS

This work presented an integrated SSbD framework for paracetamol manufacturing that couples CAR with model-based end-to-end process design, TEA, and prospective LCA, while embedding Quality by Digital Design (QbDD) principles to support robust decision-making. The case study compared four alternative route candidates: two CAR/SYNTHIA®-derived routes (Route A as a baseline, conventional-like pathway and Route B as an SSbD-preferred alternative) and two literature comparators (renewables-based and bio-waste derived concepts).

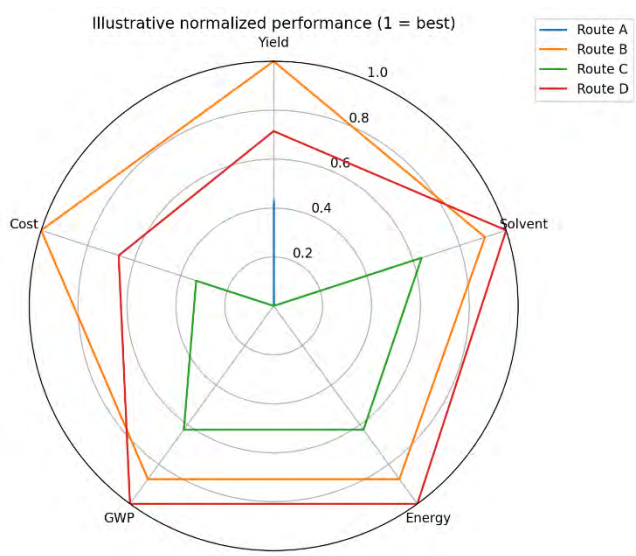

Figure 6. Normalized multi-criteria performance comparison (1 = best) for yield, solvent intensity, energy intensity, GWP, and relative production cost.

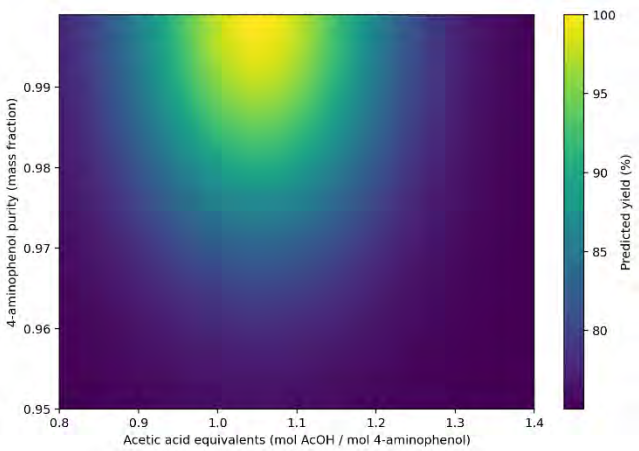

Figure 7. Design space analysis for route B.

Across the evaluated KPIs, route choice was shown to strongly influence end-to-end performance through coupled effects on separation intensity, solvent makeup, and utility demand. The results indicate that Route A is dominated by the alternative candidates, while Route B (CAR SSbD route) and Route D (bio-waste literature route) occupy the preferred region of the cost–impact trade-off space. Hotspot analysis further indicates that reductions in environmental burden are primarily achieved by lowering materials-related contributions (notably solvent/reagent supply chains) and utilities-related contributions (solvent recovery and drying duties), reinforcing the importance of solvent management as a central SSbD lever. Multi-criteria comparison provided a transparent basis for decision-making by revealing where improvements can be achieved simultaneously and where genuine trade-offs remain between cost, climate impact, and yield.

It is important to note that the mapping between reaction-level impurity proxies and downstream separation intensity in the present framework represents a comparative screening abstraction rather than a strictly linear mechanistic relationship. In pharmaceutical crystallization systems, minor variations in specific impurity concentrations can significantly alter nucleation kinetics, crystal habit, polymorphic outcome, and filterability. Consequently, the current modelling approach should be interpreted as an early-stage digital screening tool intended to prioritize candidate routes. Experimental validation, including crystallization studies and solid-state characterization, remains essential to confirm separation behaviour and product quality attributes prior to scale-up.

Overall, the study demonstrates that integrating CAR route generation with process simulation and prospective TEA/LCA enables earlier, more quantitative screening of “Greener-by-Design” options than conventional sequential development workflows. Future work will explore more placeholders with SYNTHIA-exported route descriptors and gPROMS-derived inventories and will extend the framework by explicitly quantifying robustness under uncertainty (CPP/CMA variability) to link QbDD design space definition with sustainability and economic outcomes.

AUTHOR IDENTIFIERS

Author ORCIDs:
Gao S: 0009-0001-0541-1015

REFERENCES

1. Gazis TA, Wuyts J, Moutsiou A, Volpin G, Ford MJ, Teixeira RI, Wheelhouse KMP, Natho P, Žnidarši?-Plazl P, Jost S, Luisi R, Benyahia B, Maes BUW, Vilé G. Towards greener-by-design fine chemicals. part 1: synthetic frontiers. Chem. Soc. Rev. 55:619-674 (2026). https://doi.org/10.1039/d5cs00929d
2. Ma Y, Gao S, Mehta N, Fu Q, Li W, Benyahia B. solecos : a data-driven platform for sustainable and comprehensive solvent selection in pharmaceutical manufacturing. Green Chem. 27:12621-12641 (2025). https://doi.org/10.1039/d5gc04176g
3. Zhao Y, Li W, Zhao C, Fang L, Yang H, Ma Y. Data-driven point adjusted jouyban-acree-artificial neural network hybrid model for predicting solubility of active pharmaceutical ingredients in binary solvent mixtures. Ind. Eng. Chem. Res. 63:16529-16544 (2024). https://doi.org/10.1021/acs.iecr.4c02110
4. Gao S, Benyahia B. Robust techno-economic analysis, life cycle assessment, and quality and sustainability by digital design of three alternative continuous pharmaceutical tablet manufacturing processes. Systems and Control Transactions 4:2568-2573 (2025). https://doi.org/10.69997/sct.104102
5. Teixeira RI, Benyahia B. Design and optimization of a shared synthetic route for multiple active pharmaceutical ingredients through combined computer aided retrosynthesis and flow chemistry. Chemical Engineering Research and Design 216:367-375 (2025). https://doi.org/10.1016/j.cherd.2025.03.004
6. Teixeira, RI, Andresini, M, Luisi, R, Benyahia, B (2024) Computer-Aided Retrosynthesis for Greener and Optimal Total Synthesis of a helicase-primase inhibitor Active Pharmaceutical Ingredient. *Journal of the American Chemical Society Gold*, 4, 4263–4272.
7. Chen Z, Lian JZ, Zhu H, Zhang J, Zhang Y, Xiang X, Huang D, Tjokro K, Barbarossa V, Cucurachi S,

Dong B. Application of life cycle assessment in the pharmaceutical industry: a critical review. Journal of Cleaner Production 459:142550 (2024). https://doi.org/10.1016/j.jclepro.2024.142550
8. Satta M, Passarini F, Cespi D, Ciacci L. Advantages and drawbacks of life cycle assessment application to the pharmaceuticals: a short critical literature review. Environ Sci Pollut Res : (2024). https://doi.org/10.1007/s11356-024-33964-w
9. Boyall SL, Clarke H, Dixon T, Davidson RWM, Leslie K, Clemens G, Muller FL, Clayton AD, Bourne RA, Chamberlain TW. Automated optimization of a multistep, multiphase continuous flow process for pharmaceutical synthesis. ACS Sustainable Chem. Eng. 12:15125-15133 (2024). https://doi.org/10.1021/acssuschemeng.4c05015
10. Park J, Evans C, Maier J, Hatzell M, France S, Sievers C, Bommarius AS. Renewables-based routes to paracetamol: a green chemistry analysis. ACS Sustainable Chem. Eng. 12:16271-16282 (2024). https://doi.org/10.1021/acssuschemeng.4c05353
11. SYNTHIA®. Retrosynthesis software workflow and export functions (PDF/RDF/SVG). https://www.synthiaonline.com/product/workflow (accessed 2026-01-21)
12. MilliporeSigma. SYNTHIA® Retrosynthesis Software product description (expert-coded rules; large starting material catalog). https://www.sigmaaldrich.com/GB/en/services/software-and-digital-platforms/synthia-retrosynthesis-software (accessed 2026-01-21)
13. Hallamasek S, Ubbenjans V, Lapkin AA. Life cycle assessment of a process for paracetamol flow synthesis from bio-waste derived ?-pinene. Sustainable Chemistry and Pharmacy 40:101629 (2024). https://doi.org/10.1016/j.scp.2024.101629
14. ICH Harmonised Guideline. Pharmaceutical Quality System Q10. https://database.ich.org/sites/default/files/Q10%20Guideline.pdf (accessed 2026-01-21)
15. World Health Organization (WHO). WHO Model Lists of Essential Medicines (includes paracetamol/acetaminophen). https://www.who.int/groups/expert-committee-on-selection-and-use-of-essential-medicines/essential-medicines-lists (accessed 2026-01-21)
16. Mustoe CL, Turner AJ, Urwin SJ, Houson I, Feilden H, Markl D, Al Qaraghuli MM, Chong MWS, Robertson M, Nordon A, Johnston BF, Brown CJ, Robertson J, Adjiman C, Batchelor H, Benyahia B, Bresciani M, Burcham CL, Cardona J, Cottini C, Dunn AS, Fradet D, Halbert GW, Henson M, Hidber P, Langston M, Lee YS, Li W, Mantanus J, McGinty J, Mehta B, Naz T, Ottoboni S, Prasad E, Quist PO, Reynolds GK, Rielly C, Rowland M, Schlindwein W, Schroeder SLM, Sefcik J, Settanni E, Siddique H, Smith K, Smith R, Srai JS, Thorat AA, Vassileiou A, Florence AJ. Quality by digital design to accelerate sustainable medicines development. International Journal of Pharmaceutics 681:125625 (2025). https://doi.org/10.1016/j.ijpharm.2025.125625

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Section 5

Modelling and Simulation

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

An Open-Source IDAES Framework for Simulating Inductively Heated Adsorption Processes

Sudip Sharma[a], Thomas A. Adams II[a]*

[a] Norwegian University of Science and Technology (NTNU), Department of Energy and Process Engineering, Trondheim, Norway
* Corresponding Author: thomas.a.adams@ntnu.no

ABSTRACT

Magnetic Inductive Swing Adsorption (MISA) is a carbon dioxide capture process similar to Temperature Swing Adsorption that uses direct electromagnetic heating instead of classic heating systems for the regeneration step of the process. However, the lack of validated dynamic models hinders process optimization. This work introduces an open-source MISA model in the IDAES framework, incorporating Specific Absorption Rate (SAR) physics (SAR $\propto B^2$) to capture electromagnetic heating. Binary Sips isotherm parameters for Fe_3O_4@HKUST-1 were fitted to experimental data, achieving high statistical agreement (R2 > 0.996, RMSE < 0.022 mol/kg). Comprehensive validation was performed against adsorption isotherms, dynamic breakthrough curves, and desorption profiles. The model predicts breakthrough time with only 9% error and saturation time with 6% error. Crucially, the coupled thermal transport and SAR heating model capture temperature evolution during desorption within 5% error across all field strengths. Although the use of Linear Driving Force kinetics introduces minor systematic overprediction, the model successfully bridges the gap between laboratory feasibility and industrial design. This validated tool enables the first systematic investigation of cycle configurations, providing a platform for techno-economic comparisons and scale-up of energy-efficient magnetic carbon capture.

Keywords: Carbon Capture, Adsorption, Modelling and Simulations, Metal Organic Framework (MOF)

1. INTRODUCTION

Mitigating anthropogenic carbon dioxide emissions remains critical for achieving global climate targets, with post-combustion carbon dioxide (CO_2) capture representing a key near-term strategy for industrial decarbonization [1]. Traditional Temperature Swing Adsorption (TSA) processes regenerate solid adsorbents through external heating, but suffer from fundamental limitations; poor thermal conductivity creates temperature gradients, slow heat penetration extends cycle times, and thermal resistances cause energy losses often exceeding 3-4 GJ/tonne CO_2 [2].

Magnetic Inductive Swing Adsorption (MISA) addresses these limitations through direct volumetric heating of magnetic nanoparticles embedded in adsorbent materials [2]. When exposed to alternating magnetic fields, Fe_3O_4 nanoparticles generate heat through hysteresis and relaxation mechanisms, bypassing thermal resistances inherent in conventional TSA. Bellusci *et al.* [2] demonstrated that Fe_3O_4@HKUST-1 composites achieve 95% CO_2 recovery in about 2.5 minutes, almost four times faster than Temperature Swing Adsorption (TSA) [2].

Despite experimental demonstrations of MISA feasibility, validated dynamic models for process design remain lacking. Existing fixed-bed models in process simulation frameworks do not include electromagnetic heating capabilities, and proprietary models are not publicly available. This absence of validated open-source tools represents a barrier to systematic exploration of MISA cycle configurations and scale-up necessary for industrial implementation.

This work presents a new, validated open-source MISA model using the Institute for the Design of Advanced Energy Systems (IDAES) framework [3]. The model extends the standard `FixedBed1D` class with Specific Absorption Rate (SAR) physics for electromagnetic heating. Binary Sips isotherm parameters were fitted to experimental data [2, 4] using Python Adsorption Parameter Estimation Package (pyAPEP) [5], achieving R^2 >

0.996 for both CO_2 and Nitrogen (N_2). Comprehensive validation was performed against three independent datasets: adsorption isotherms (273-403 K), dynamic breakthrough curves (15-50% CO_2), and magnetic desorption profiles (11.1-12.6 mT). The validated implementation enables systematic investigation of cycle design and optimization, providing the computational foundation for techno-economic analysis and scale-up studies.

2. MATERIALS AND METHODS

2.1 Adsorbent Material and Isotherms

The adsorbent material used in this study is a Fe_3O_4@HKUST-1 magnetic composite developed by Bellusci *et al.* [2]. This material consists of the metal organic framework (MOF) HKUST-1 ($Cu_3(BTC)_2$) impregnated with 17 wt% magnetite (Fe_3O_4) nanoparticles. The magnetic nanoparticles enable direct heating of the solid phase through electromagnetic induction, while MOF provides high CO_2 adsorption capacity.

The adsorption equilibrium for both CO_2 and N_2 was modeled using the Sips isotherm, which combines Langmuir and Freundlich characteristics:

$$q_{(eq,i)(T,P_i)} = \frac{n_{\infty,i}(b_i(T)\,P_i)^{c_i}}{1+(b_i(T)\,P_i)^{c_i}} \quad (1)$$

where $q_{(eq,i)}$is the adsorbed amount of component i (mol/kg), $n_{\infty,i}$ is saturation capacity, $b_i(T)$ is the affinity parameter as a function of T, c_i is the heterogeneity parameter, and P_i is the partial pressure. The affinity parameter exhibits temperature dependence according to:

$$b_i(T) = b_{ref,i}\exp\left[\frac{\Delta H_{b,i}}{R}\left(\frac{1}{T_{ref}}-\frac{1}{T}\right)\right] \quad (2)$$

where $\Delta H_{b,i}$ is the isosteric heat of adsorption and T_{ref} = 298 K.

Experimental adsorption isotherms for both CO_2 and N_2 were obtained from Bellusci *et al.* [2, 4] by manual extraction of gas uptake (mol/kg) versus pressure (kPa) data across multiple temperatures (273-403 K for CO_2, 273-308 K for N_2) using a plot digitation tool. The extracted dataset served as the training set for parameter estimation.

The binary Sips isotherm parameters were determined through nonlinear regression using pyAPEP [5], which performs global nonlinear regression to simultaneously optimize all parameters.

Figure 1 demonstrates excellent agreement between the fitted model and experimental data. The fitted parameters (Table 1) reveal the expected selectivity: CO_2 exhibits significantly higher saturation capacity. The heats of adsorption (ΔH_b = 23.5 kJ/mol for CO_2, 15.3 kJ/mol for N_2) are physically reasonable for physisorption and correctly predict decreased loading with increasing temperature and is close to the experimental value of 23 kJ/mol for CO_2 [2]. The validated isotherm provides the thermodynamic foundation for all subsequent dynamic simulations.

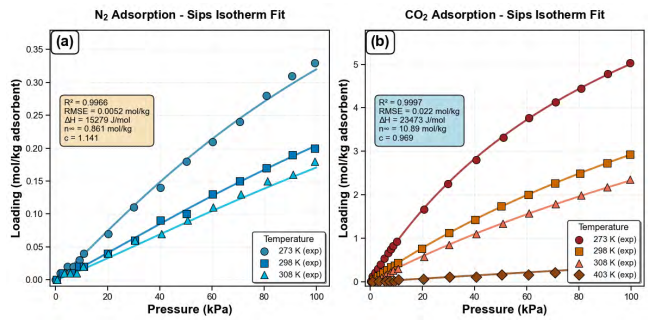

Figure 1. Validation of Sips isotherm model against experimental adsorption data from Bellusci et al. [2, 4]. (a) N_2 adsorption (273-308 K, R^2=0.9970). (b) CO_2 adsorption (273-403 K, R^2=0.9997). Lines: model predictions; symbols: experimental measurements.

Table 1 presents the fitted parameters along with statistical measures of goodness-of-fit. The high R^2 values (>0.996) and low RMSE values indicate excellent agreement with data across the measured temperature and pressure ranges (298 – 403K, 0.01-100 kPa). Details are available in the supplementary material.

Table 1. Sips Isotherm Parameters found

Parameter	CO_2	N_2	Unit
n_∞	10.89	0.926	mol/kg
b_{ref}(298K)	4.24 × 10^{-6}	1.57 × 10^{-6}	Pa^{-1}
c	0.969	1.128	–
ΔH_b	23, 473	15, 184	J/mol
R^2	0.9997	0.997	–
$RMSE$	0.022	0.0049	mol/kg

2.2 Mathematical Model

2.2.1 System Description and Assumptions

The Magnetic Inductive Swing Induction process is modeled as a one-dimensional fixed bed adsorber operated in transient mode. The model captures the coupled dynamics of gas flow, heat transfer, and adsorption/desorption in a cylindrical packed bed containing Fe_3O_4@HKUST-1 composite adsorbent. The key modeling assumptions are:

- **One-dimensional plug flow:** Radial concentration and temperature gradients are negligible compared to axial variations
- **Dynamic operation:** Transient behavior captured through time-dependent terms
- **Two-phase system** with local thermal non-equilibrium between gas and solid phases

- **Ideal gas behavior:** Gas mixture obeys the ideal gas law
- **Uniform bed properties:** Bed porosity and particle size throughout the bed
- **Linear driving force (LDF) kinetics:** Adsorption rate proportional to difference between equilibrium and actual loading
- **Negligible axial dispersion:** Convective transport dominates over diffusive/dispersive transport in the gas phase
- **Local pressure equilibrium:** Pressure within particle pores instantly equilibrates with the bulk gas pressure

2.2.2 Conservation Equations -Standard Framework

We implemented the model in IDAES using the `FixedBed1D` class, which generates standard conservation equations [6]. The governing equations include:

Gas phase material balance for each component i (CO_2, N_2):

$$A_b \cdot \varepsilon_b \cdot \frac{\partial(\rho_{mol} \cdot y_i)}{\partial t} = -\frac{\partial(F_{mol,i})}{\partial x} - A_s \cdot \rho_p \cdot r_{ads,i} \quad (3)$$

Gas phase energy balance:

$$A_b \cdot \varepsilon_b \cdot \frac{\partial(\rho_{mol} \cdot h_{mol,g})}{\partial t} = -\frac{\partial H_g}{\partial x} - \frac{6h_{gs}}{d_p} \cdot A_s \cdot (T_g - T_s) \quad (4)$$

Momentum Balance (Ergun Equation):

$$-\frac{dP_g}{dx} = \frac{150\mu_i(1-\varepsilon_b)^2 u}{\varepsilon_b{}^3 d_p^2} + \frac{1.75\rho_g(1-\varepsilon_b)u^2}{\varepsilon_b{}^3 d_p} \quad (5)$$

Solid Phase material balance for adsorbed component i:

$$A_s \cdot \rho_p \cdot \frac{\partial(q_i)}{\partial t} = A_s \cdot \rho_p \cdot r_{ads,i} \quad (6)$$

where geometric parameters include bed porosity ε_b[-], bed cross-sectional area A_b $[m^2]$, solid cross-sectional area $A_s = (1-\varepsilon_b)\,A_b$ $[m^2]$, and particle diameter $d_p[m]$. Gas phase variables include molar density ρ_{mol} $[mol/m^3]$, mass density ρ_g $[kg/m^3]$, species mole fraction y_i [-], superficial velocity u $[m/s]$, temperature T_g [K], absolute pressure P_g $[Pa]$, viscosity μ_g $[Pa \cdot s]$, and specific enthalpy $h_{mol,g}$ $[J/mol]$. Enthalpy flow rate $H_g = u \cdot A_b \cdot \rho_{mol} \cdot h_{mol,g}$ $[W]$ and species molar flow rate $F_{mol,i} = u \cdot A_b \cdot \rho_{mol} \cdot y_i$ $[mol/s]$. Solid phase variables include particle density ρ_p $[kg/m^3]$ adsorbed loading q_i $[mol/kg]$, temperature T_s $[K]$, and adsorption rate r_{ads}, i $[mol/(kg \cdot s)]$. The gas-solid heat transfer coefficient is h_{gs} $[W/(m^2 \cdot K)]$.

Bed porosity (void fraction) is calculated using the following relation [7]:

$$\varepsilon_b = 0.38 + 0.073\left[1 - \frac{\left(\frac{d_b}{d_p} - 2\right)^2}{\left(\frac{d_b}{d_p}\right)^2}\right] \quad (7)$$

where d_b is bed diameter and d_p is the particle diameter.

2.2.3 Inductive Heating Model

The key innovation extends the standard IDAES `FixedBed1D` model with inductive heating capability. When an alternating magnetic field is applied, the Fe_3O_4 nanoparticles within the composite generate heat through magnetic losses (hysteresis and relaxation mechanisms). This enables direct volumetric heating of the solid phase, bypassing the thermal resistances inherent in conventional TSA.

Specific Absorption Rate (SAR)

The heating power is characterised by the Specific Power Loss (SPL) or Specific Absorption Rate (SAR), which represents the power dissipated per unit mass of material. Based on magnetic relaxation theory[8, 9], SAR exhibits a quadratic dependence on the applied magnetic field amplitude:

$$L(B) = L_{ref}\left(\frac{B}{B_{ref}}\right)^2 \quad (8)$$

B is the applied magnetic field amplitude (mT), B_{ref} = 12.6 mT is the reference field amplitude, L is the SAR as a function of B, and L_{ref} = 1.6 W/g is the experimentally measured SAR of the Fe_3O_4@HKUST-1 composite at the reference field amplitude [2]. This quadratic scaling captures the fundamental physics of magnetic power dissipation in the superparamagnetic regime.

Modified Solid Energy balance

$$A_s \cdot \rho_p \cdot c_{p,s} \cdot \frac{\partial T_s}{\partial t} = \frac{6h_{gs}}{d_p} \cdot A_s \cdot (T_g - T_s) - A_s \cdot \rho_p \sum_i [\Delta H_{ads,i} \cdot r_{ads,i}] + Q_{ind} \cdot A_s - Q_{wall} \quad (9)$$

where Q_{ind} $[W/m^3]$ is the volumetric inductive heating rate defined in Eq. (10), $\Delta H_{ads,i}$ $[J/mol]$ is the heat of adsorption for species i (negative for exothermic adsorption), and Q_{wall} $[W/m]$ is the heat loss to ambient defined in Eq. (11)

$$Q_{ind} = L(B) \cdot \rho_p \quad (10)$$

$$Q_{wall} = h_{wall} \cdot \pi \cdot d_b \cdot (T_s - T_{ambient}) \quad (11)$$

where h_{wall} is the wall heat transfer coefficient and $T_{ambient}$ = 298 K is ambient temperature.

This formulation has three key advantages for MISA modeling:

- **Direct solid heating:** Energy is deposited directly in the solid phase, eliminating gas-to-solid heat transfer resistance
- **Controllable heating:** SAR is proportional to B^2,

allowing precise control via magnetic field amplitude

- **Rapid thermal response:** Inductive heating timescale (seconds) is much faster than conductive heating (minutes) in conventional TSA

When B = 0 (adsorption phase), Q_{ind} = 0 and the model reduces to the standard fixed-bed model but with wall heat loss. This enables modeling of complete MISA cycles where the magnetic field is switched on only during the desorption step.

2.2.4 Constitutive Relations

The model is completed with the following constitutive relations:

Adsorption equilibrium: Binary Sips isotherm (Section 2.1)

Adsorption kinetics: Linear Driving Force (LDF) model:

$$r_{ads,i} = k_i(q_{eq,i} - q_i) \quad (12)$$

where k_i is the mass transfer coefficient, $q_{eq,i}$ is the equilibrium loading from the Sips isotherm, and $q_i$is the average solid-phase loading concentration of component i within the adsorbent particle at a given time. The mass transfer coefficient (k_{CO_2} = $0.15\ s^{-1}$) was determined by fitting experimental breakthrough curve data. The N_2 coefficient (k_{N_2} = $0.20\ s^{-1}$) was set higher than CO_2, consistent with N_2's weaker adsorption (lower heat of adsorption) and minimal contribution to overall bed dynamics due to its low saturation capacity (0.93 mol/kg vs. 10.89 mol/kg for CO_2).

Gas phase thermodynamics: Ideal gas equation of state. Pure component properties are calculated from temperature-dependent correlations: heat capacity from NIST Shomate equations [10], and while viscosity and thermal conductivity use temperature-dependent power law correlations [11]. Mixture heat capacity is calculated as mole-fraction weighted average of component values. Mixture viscosity and thermal conductivity are calculated using Wilke mixing rules [12].

Heat transfer: The gas-solid heat transfer coefficient (h_{gs}) is calculated from particle-scale Nusselt number correlations as described in the IDAES documentation [6].

2.3 Numerical Implementation

The mathematical model was implemented using IDAES framework built on the Pyomo optimization modeling environment [13]. IDAES provides standardized property package interfaces, robust initialization strategies for DAE systems, and integration with state-of-the-art solvers. The IDAES `ControlVolume1D` and `FixedBed1D` classes automatically generate standard conservation equations (material, energy, momentum), while custom constraints were implemented for inductive heating and wall heat loss.

Table 2. Model Parameters

Parameter	Value	Units
Column Geometry		
Length (L)	1.5	cm
Diameter (D)	1.0	cm
Bed porosity (ε_b)	0.39	–
Adsorbent Properties		
Particle diameter (d_p)	0.5	mm
Particle density (ρ_p)	1170	kg/m^3
Solid heat capacity ($C_{p,s}$)	1070	$J/(kg \cdot K)$
Operating Conditions - Adsorption		
Feed temperature (T_{in})	303	K
Feed pressure (P_{in})	1.3	bar
CO_2 concentration	15	$mol\%$
Molar flow rate	3.33×10^{-5}	mol/s
Magnetic field (B)	0	mT
Operating Conditions - Desorption		
Feed temperature (T_{in})	303	K
CO_2 concentration	15	$mol\%$
Molar flow rate	3.33×10^{-5}	mol/s
Magnetic field (B)	11.1 - 12.6	mT
Inductive Heating Parameters		
SAR at 12.6 mT (L_{ref})	1.6	W/g
Frequency (f)	190	kHz
Fe_3O_4 content	17	$wt\%$
Heat Transfer		
Wall heat transfer Coeff. (h_{wall})[a]	20	$W/(m^2 \cdot K)$
Ambient temperature (T_{amb})	298	K

[a] Values chosen to be in line with the literature [2, 16]

The reactor length domain (0 ≤ x ≤ L) was discretized using the backward finite difference method:

$$\frac{\partial \psi}{\partial x|_i} \approx \frac{\psi_i - \psi_{i-1}}{\Delta x} \quad (13)$$

This discretization converts the PDEs into a DAE system in time. The backward difference scheme ensures numerical stability for convection-dominated flows and preserves causality in the flow direction.

A two-phase solution strategy was employed.

Phase 1 (Initialization): The model was initialized using hierarchical block triangularization with IPOPT [14]. State variables were temporarily fixed and constraints progressively activated in this order: thermophysical

properties, hydrodynamics, mass balances, energy balances.

Phase 2 (Dynamic Integration): The spatially discretized DAE system was integrated forward in time using the PETSc library [15] with backward Euler method:

$$\frac{\partial \psi}{\partial t}\Big|_{n+1} \approx \frac{\psi_{n+1} - \psi_n}{\Delta t} \tag{14}$$

where Δt is the time step (adaptive). Backward Euler provides stability for stiff systems, essential for capturing fast thermal transients during inductive heating.

Table 2 summarizes the operating conditions, bed geometry, and material properties used in all simulations.

2.4 Cycle Configuration

A simple MISA process operates as a three-step cyclic process: adsorption, inductive desorption, cooling.

Step 1 – Adsorption: The regenerated bed, initially equilibrated with pure N_2 at 303 K, is exposed to the adsorption operation condition given in Table 2. CO_2 selectively adsorbs onto the Fe_3O_4@HKUST-1 composite until the bed approaches saturation, typically occurring at 135-145 seconds under the conditions in Table 2.

Step 2 - Magnetic Desorption: Upon completion of adsorption, the magnetic field is activated (B = 11.1-12.6 mT, f = 190 kHz) while maintaining the same feed flow. Electromagnetic energy is directly absorbed by the Fe_3O_4 nanoparticles, causing rapid localized heating of the solid adsorbent. This direct solid-phase heating mechanism bypasses the thermal resistances inherent in conventional TSA, where heat must be conducted through the reactor walls and gas phase. The bed temperature rises rapidly from 303 K to approximately 440 K (167°C), driving CO_2 desorption. At B = 12.6 mT (L = 1.6 W/g), approximately 95% of adsorbed CO_2 is recovered within 2.11 minutes—a 4-fold reduction compared to conventional TSA [2].

Step 3 - Cooling: The magnetic field is deactivated, and the bed is cooled by continuing gas flow (pure N_2) until thermal equilibrium is restored at 303 K. Cold N_2 (convection) and wall heat loss (h_{wall}) to ambient contribute to cooling. Once cooled, the system begins the next adsorption cycle.

This cycle configuration serves as the basis for the validation studies in Section 3.

3 MODEL VALIDATION AND DISCUSSION

The model demonstrated two independent tests: dynamic breakthrough behavior, magnetic desorption performance. This demonstration approach ensures the model accurately and captures both thermodynamic equilibrium and transport-limited dynamic behavior under both adsorption and desorption conditions.

3.1 Breakthrough Validation

Dynamic breakthrough simulations were performed to validate the coupled mass transfer and equilibrium behavior of the model under transient conditions. The breakthrough experiments reported were replicated numerically using the operating conditions in Table 2. A step change in feed composition from pure N_2 to a binary mixture (15% CO_2, 85% N_2) was introduced at t = 0 to a freshly regenerated bed initially at 303 K. The outlet CO_2 concentration was monitored until bed saturation.

3.2.1 Detailed Validation at 15% CO_2

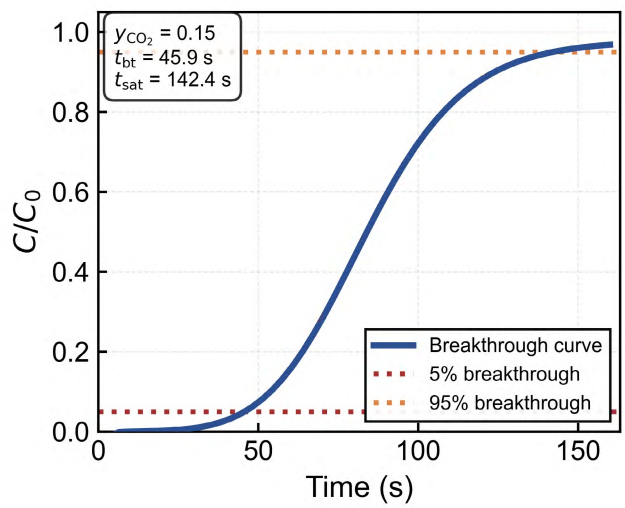

Figure 2. Solid line: model prediction. Orange dashed lines indicate the model's calculated 5% breakthrough (t = 45.9 s) and 95% saturation (t = 142.4 s) times. These are compared against the experimental benchmark values of 42 s and 138 s, respectively, as reported in the text of [2].

Figure 2 shows the high-resolution breakthrough curve validation for the 15% CO_2 feed case over a 0–160 second timeframe. Breakthrough curves express outlet concentration normalized by inlet (C/C_0). Time domain was discretized into 50 points.

Comparison with experimental data demonstrates strong agreement. The model achieves 9% error in breakthrough time and 6% error in saturation time relative to experimental values from [2]. This close agreement validates both the Sips equilibrium isotherm parameters (Table 1) and the LDF mass transfer coefficients, confirming that the model accurately captures the coupled thermodynamic and adsorptive behavior of the system.

The slight overprediction of breakthrough times may be attributed to factors like: i) the LDF approximation assumes a uniform concentration driving force within particles, while actual intraparticle profiles may be more complex [17] ii) bed channeling or radial flow maldistribution in the experimental setup, which are not captured by the 1D model [18].

3.2.2 Effect of Feed CO_2 Concentration

To further demonstrate the model's predictive capability across a wider operating range, breakthrough simulations were performed at three different CO_2 feed concentrations: 15%, 30%, and 50% (balance N_2). Figure 3 shows the simulated breakthrough curves over a 0–15 minute timeframe for all three cases. As expected, increasing CO_2 concentration results in earlier breakthrough and faster approach to saturation due to the higher driving force for adsorption and the more rapid consumption of available adsorption sites.

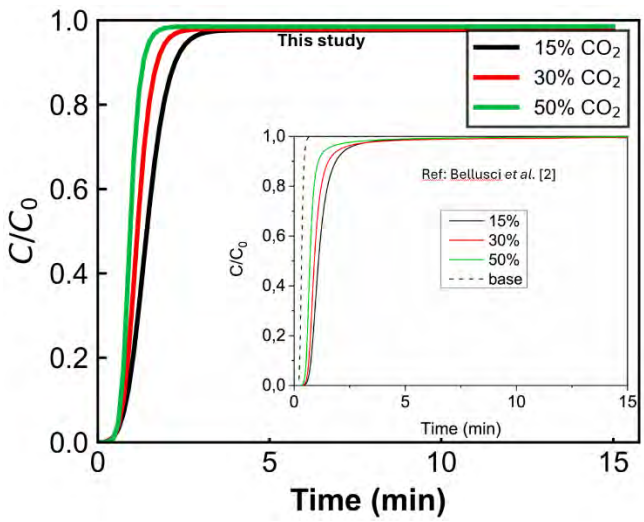

Figure 3. Breakthrough behavior at varying CO_2 feed concentrations (15%, 30%, 50%).

High-resolution experimental data for quantitative validation at 30% and 50% CO_2 were not available in [2]. However, the model captures the expected theoretical behavior: increasing the CO_2 feed concentration from 15% to 50% reduces the saturation time by a factor of 1.6. This response is consistent with the competitive adsorption physics described by the binary Sips isotherm, where higher partial pressures drive faster saturation.

3.2 Desorption Validation

Desorption simulations were performed to demonstrate the inductive heating capability and thermal coupling in the model. The desorption experiments were also replicated numerically by first saturating the bed with 15% CO_2 at 303 K, then activating the magnetic field while maintaining the same feed flow rate and composition. Three different magnetic field intensities were tested: B = 11.1, 11.8, and 12.6 mT at f = 190 kHz.

3.2.1 CO_2 Desorption Profiles

Figure 4 presents the simulated outlet CO_2 concentration profiles. Following bed saturation, the magnetic field is activated at t=0, and the outlet concentration is monitored. The peak outlet concentration reaches 28.3 vol% at 12.6 mT, demonstrating near-complete regeneration.

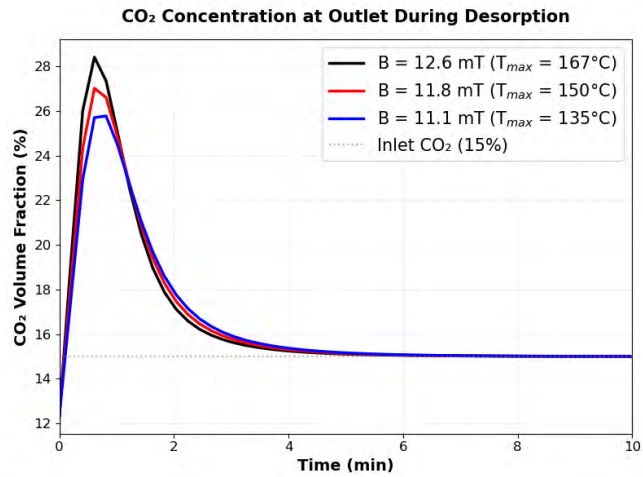

Figure 4. CO_2 desorption curve at different applied magnetic fields with maximum temperature attained at each magnetic field.

These results demonstrate the quadratic B^2 scaling relationship in the SAR model (Equation 8), consistent with Rosensweig's theory [8]. To substantiate the model's accuracy, Table 3 compares the results with experimental data from Bellusci et al. [2]. The model achieves high quantitative accuracy, predicting peak desorption magnitudes with <6% error across all field strengths. While the simulation predicts a slightly broader desorption profile due to continuum mass-transfer assumptions, the peak magnitude and total regeneration times show strong agreement with experimental values. We note that the experimental values shown in Table 3 are not precise because they were digitally determined from a low-resolution figure, so peak error is only valid to one significant figure.

Table 3. Quantitative comparison of simulated desorption against experimental data [2].

Field	Peak CO_2 (vol%)	Peak Error	Time to Peak (min)
mT	Exp./Sim.	(%)	Exp./Sim.
12.6	29.0/28.3	2	0.50/0.65
11.8	26.5/27.0	2	0.48/0.65
11.1	24.5/25.8	5	0.49/0.65

3.2.2 Temperature Evolution

Figure 5 shows the simulated bed temperature profiles during magnetic desorption for the same three field strengths. The temperature evolution directly reflects the electromagnetic heating rate determined by the SAR value. At B = 12.6 mT, the bed temperature rises rapidly from 303 K to a peak of approximately 440 K (167°C). At the lower field strengths, the peak temperatures are

proportionally lower: ~423 K (150°C) at B = 11.8 mT and ~408 K (135°C) at B = 11.1 mT.

Table 4 provides a quantitative comparison of the predicted peak temperatures against experimental measurements. The model produces close match to experimental data, with errors under 5% for all three field strengths. At the highest field strength (B = 12.6 mT), the model predicts 440 K (167°C) versus the experimental value of 448 K (175°C), representing a 4.5% error. This close agreement validates both the SAR heating model (Equation 8) and the thermal energy balance (Equation 9), confirming that the model accurately captures the coupled heat generation, accumulation, and loss mechanisms.

Table 4. Temperature and Regeneration Validation

B (mT)	T_{max} Exp.	T_{max} Model	Error (%)
12.6	448 K (175°C)	440 K (167°C)	4.5
11.8	428 K (155°C)	423 K (150°C)	3.2
11.1	403 K (130°C)	408 K (135°C)	3.8

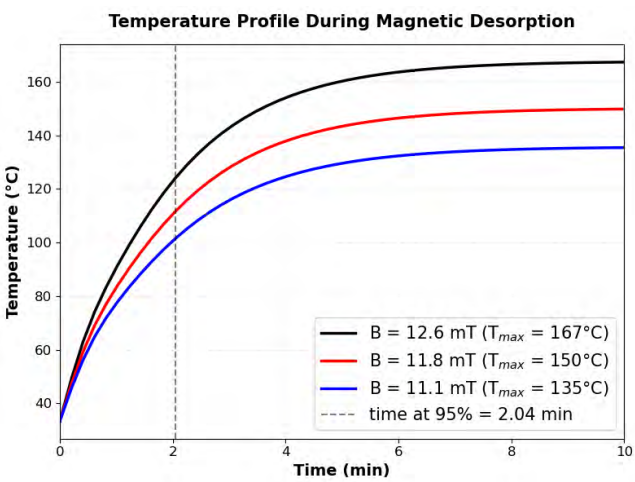

Figure 5. Temperature evolution at different B values and 95% regeneration time.

The slight underprediction of peak temperatures may be attributed to: (*i*) radial temperature gradients not captured by the 1D axial model - the thermocouple may measure local radial hot spots near the wall or center that differ from the radially-averaged temperature predicted by the 1D formulation [18], (*ii*) uncertainty in the solid heat capacity ($Cp,s = 1070\, J/(kg \cdot K)$), which was assumed constant but varies with temperature, and (*iii*) heat losses through the reactor wall (h_{wall} = 20 $W/(m^2 \cdot K)$), which may be overestimated, leading to lower predicted peak temperatures.

Despite these minor discrepancies, the strong correlation between predicted and experimental temperatures across three different field strengths confirms that the model correctly captures the physics of electromagnetic heating and thermal transport in the MISA system. The validated thermal model enables confident prediction of energy requirements and regeneration performance under various operating conditions.

3.3 Model Limitations and Applicability

While the model accurately captures thermodynamic limits, two approximations introduce kinetic deviations. First, the LDF approximation (k_i= 0.15 s^{-1}) simplifies intra-particle transport, resulting in a broader desorption profile and slower post-peak decay compared to the sharp experimental release [18]. Second, the continuum heating assumption causes a minor onset lag (~0.15 min), as volume-averaged temperatures respond slower than instantaneous local "hot spots." Despite these broadening effects, the model's thermodynamic accuracy (<6% peak error) and robust regeneration time predictions validate its utility for cycle optimization and scale-up.

4 CONCLUSION AND FUTURE WORK

This work presents the first validated open-source IDAES model for Magnetic Inductive Swing Adsorption, extending the standard `FixedBed1D` framework with electromagnetic heating capability via Specific Absorption Rate physics. Multi-level validation against experimental data confirms the model accurately captures coupled thermodynamic-kinetic-thermal behavior across different conditions for CO_2 capture (273-440 K, 15-50% CO_2, 11.1-12.6 mT), with breakthrough predictions within 9%, temperature predictions within 5%, and isotherm fits exceeding R^2 = 0.996. The validated quadratic $L \propto B^2$ scaling confirms MISA achieves four-fold faster regeneration than conventional TSA.

While the LDF kinetic approximation introduces ~10% systematic error in breakthrough times and validation remains limited to laboratory scale, the model provides a robust foundation for cycle optimization and techno-economic analysis. Future work will leverage this platform to identify optimal cycle configurations through multi-objective optimization, compare MISA economics against conventional TSA/PSA technologies, and guide scale-up to pilot demonstration.

The open-source implementation enables model extensions and facilitates MISA advancement from laboratory demonstrations toward industrial deployment for energy-efficient CO_2 capture.

ACKNOWLEDGEMENTS

This project was funded by the NTNU-MIT Energy Research Programme (project NMERP-C2).

SUPPLEMENTARY MATERIAL

Supplementary material can be downloaded at: https://PSEcommunity.org/LAPSE:2025.0725.

AUTHOR IDENTIFIERS

Author ORCIDs:
Sharma Sudip: 0000-0003-4049-6031
Adams Thomas A., II: 0000-0002-9871-9851

REFERENCES

1. Raganati F, Ammendola P. Co₂ post-combustion capture: a critical review of current technologies and future directions. Energy Fuels 38:13858-13905 (2024). https://doi.org/10.1021/acs.energyfuels.4c02513
2. Bellusci M, Albino M, Masi A, Peddis D, Innocenti C, Varsano F. High porosity-magnetic composite materials for magnetic induction swing adsorption (MISA): improvement of performance properties. Materials Chemistry and Physics 311:128525 (2024). https://doi.org/10.1016/j.matchemphys.2023.128525
3. Lee A, Ghouse JH, Eslick JC, Laird CD, Siirola JD, Zamarripa MA, Gunter D, Shinn JH, Dowling AW, Bhattacharyya D, Biegler LT, Burgard AP, Miller DC. The idaes process modeling framework and model library—flexibility for process simulation and optimization. J Adv Manuf & Process 3: (2021). https://doi.org/10.1002/amp2.10095
4. Bellusci M, Masi A, Albino M, Peddis D, Petrecca M, Sangregorio C, La Barbera A, Varsano F. Fe3o4@hkust-1 magnetic composites by mechanochemical route for induction triggered release of carbon dioxide. Microporous and Mesoporous Materials 328:111458 (2021). https://doi.org/10.1016/j.micromeso.2021.111458
5. Ga S, An N, Joo C, Kim J. Pyapep: an all-in-one software package for the automated preparation of adsorption process simulations. Computer Physics Communications 291:108830 (2023). https://doi.org/10.1016/j.cpc.2023.108830
6. Fixed bed 1d reactor — IDAES v2.8.0. [Accessed: Nov. 7 2025] https://idaes-pse.readthedocs.io/en/2.8.0/reference_guides/model_libraries/gas_solid_contactors/unit_models/fixed_bed_1D.html
7. Adams II TA, Barton PI. A dynamic two-dimensional heterogeneous model for water gas shift reactors. International Journal of Hydrogen Energy 34:8877-8891 (2009). https://doi.org/10.1016/j.ijhydene.2009.08.045
8. Rosensweig RE. Heating magnetic fluid with alternating magnetic field. Journal of Magnetism and Magnetic Materials 252:370-374 (2002). https://doi.org/10.1016/s0304-8853(02)00706-0
9. Déjardin JL, Vernay F, Kachkachi H. Specific absorption rate of magnetic nanoparticles: nonlinear AC susceptibility. Journal of Applied Physics 128: (2020). https://doi.org/10.1063/5.0018685
10. Informatics NO of D and. NIST chemistry webbook. [Accessed: Nov. 8, 2025] https://webbook.nist.gov/chemistry/
11. Green DW, Southard MZ Eds. Perry's Chemical Engineers' Handbook. 9th Edition Edition. McGraw-Hill Education (2019). ISBN 978-0-07-183408-7
12. Wilke CR. A viscosity equation for gas mixtures. The Journal of Chemical Physics 18:517-519 (1950). https://doi.org/10.1063/1.1747673
13. Bynum ML, Hackebeil GA, Hart WE, Laird CD, Nicholson BL, Siirola JD, Watson J-P, Woodruff DL. Pyomo — Optimization Modeling in Python. Springer International Publishing (2021). ISBN 978-3-030-68927-8. https://doi.org/10.1007/978-3-030-68928-5
14. Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math. Program. 106:25-57 (2005). https://doi.org/10.1007/s10107-004-0559-y
15. Abhyankar S, Brown J, Constantinescu EM, Ghosh D, Smith BF, Zhang H. PETSc/ts: a modern scalable ode/dae solver library. arXiv (2018). https://doi.org/10.48550/arXiv.1806.01437
16. Joss L, Gazzani M, Mazzotti M. Rational design of temperature swing adsorption cycles for post-combustion CO2 capture. Chemical Engineering Science 158:381-394 (2017). https://doi.org/10.1016/j.ces.2016.10.013
17. Glueckauf E, Coates JI. 241. theory of chromatography. part IV. the influence of incomplete equilibrium on the front boundary of chromatograms and on the effectiveness of separation. J. Chem. Soc. :1315 (1947). https://doi.org/10.1039/jr9470001315
18. Shafeeyan MS, Wan Daud WMA, Shamiri A. A review of mathematical modeling of fixed-bed columns for carbon dioxide adsorption. Chemical Engineering Research and Design 92:961-988 (2014). https://doi.org/10.1016/j.cherd.2013.08.018

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Development of a Predictive Model for Microbial Growth under Variable Conditions Using a Multilayer Perceptron Neural Network: Application *to Candida guilliermondii*

Jazmín Cortez-González[a], Juan Gabriel Segovia-Hernández[b], Salvador Hernández[b], Varinia López-Ramírez[a], Arturo Hernández-Aguirre[c], Rodolfo Murrieta-Dueñas[a*]

[a] Tecnológico Nacional de México / ITS Irapuato (ITESI), División de Ingeniería Química, ; Carretera Irapuato-Silao Km. 12.5, C.P:36821 Irapuato, Guanajuato, MEXICO
[b] Departamento de Ingeniería Química, División de Ciencias Naturales y Exactas, Universidad de Guanajuato, Noria Alta S/N, C.P. 36050, Guanajuato, Guanajuato, México.
[c] Centro de Investigación en Matemáticas (CIMAT) A.C., Departamento de Ciencias Computacionales, Callejón de Jalisco s/n, 36240, Mineral de Valenciana, Guanajuato, Gto., México.
* Corresponding Author: jazmin.cg @irapuato.tecnm.mx

ABSTRACT

In the field of biochemical process design, the accurate modeling of microbial growth is essential for the development and optimization of biological reactors used in the production of high-value compounds. Achieving this objective requires a detailed understanding of how environmental factors—such as pH and nutrient availability—influence microbial dynamics across the four distinct growth phases: lag, exponential, stationary, and death. Traditionally, reactor design relies heavily on the Monod model, which provides a simplified representation of microbial growth, focusing primarily on the exponential phase under constant operating conditions (1). However, this model presents substantial limitations when applied to dynamic environments where key parameters vary over time. To overcome these constraints, the present study proposes a data-driven modeling approach using a multilayer perceptron (MLP) artificial neural network for the prediction of microbial growth trajectories under varying pH conditions and substrate compositions. The yeast strain *Candida guilliermondii* was selected as the model microorganism due to its industrial relevance. Experimental growth data were collected through optical density measurements using a Multiskan™ FC Microplate Photometer (Thermo Scientific), covering a pH range from 6.0 to 8.5 and two substrate scenarios: pure xylose, and a 1:2 glucose–xylose mixture. The experimental data were used to train the MLP neural network, which generated predictive models capable of estimating growth behavior under the specified input conditions. This modeling approach enables the simulation of microbial growth curves at any given time point within the defined parameter space, providing a more flexible and comprehensive tool compared to classical models. The results of this study demonstrate that the proposed MLP-based model is a powerful computational tool for both the design and real-time control of bioprocesses. The integration of these tools into predictive control strategies is a key aspect of the bioprocess research. The model's versatility and its ability to predict microbial behavior make it a very important tool for bioreactor control.

Keywords: Artificial Intelligence, Modelling and Simulations, Machine Learning, Optimization, Biomass, microbial growth

INTRODUCTION

The design and operation of biological reactors critically depend on an accurate representation of microbial growth kinetics. These kinetics directly determine reactor productivity, stability, and control perfor-

mance. Among the most widely adopted kinetic formulations, the Monod model remains a foundational tool due to its simplicity and interpretability. Nevertheless, its validity is constrained to dilute systems, constant environmental conditions, and predominantly the exponential growth phase. Consequently, significant uncertainty arises when Monod-based models are employed for predicting microbial behavior under dynamic conditions, particularly during the lag and stationary phases.

In recent years, surrogate modeling techniques based on artificial intelligence and machine learning have emerged as promising alternatives for representing complex biological phenomena. Neural network–based models are capable of learning nonlinear relationships directly from experimental data without requiring explicit mechanistic assumptions. This makes them particularly suitable for systems involving multiple substrates, inhibitory effects, high substrate concentrations, and environmental variability.

Previous studies have explored surrogate models for microbial and microalgal growth under limited and mostly static conditions. For instance, Lee et al. (2015) reviewed kinetic models for microalgae cultivation, while Akkermans and Van Impe (2018) investigated the effect of pH inhibition using mechanistic approaches. More recent works have applied data-driven techniques to microbial systems; however, these efforts typically focus on one or two growth phases or assume fixed operating conditions. Notably, Murrieta-Dueñas et al. (2022) proposed a neural network approach to describe exponential growth under dynamic conditions, yet a comprehensive representation of the full growth trajectory remains largely unexplored.

To the best of the authors' knowledge, no previous study has demonstrated the prediction of multiple microbial growth phases under simultaneously varying pH and substrate conditions using a single surrogate model. Therefore, this work aims to fill this gap by developing an MLP-based surrogate model capable of predicting the complete microbial growth trajectory of *Candida guilliermondii* under dynamic environmental conditions.

METHODOLOGY

Figure 1 illustrates the overall methodology adopted in this study, which consists of two main stages: (i) experimental data acquisition and (ii) surrogate model development.

Experimental Design

Microbial growth experiments were conducted under controlled laboratory conditions to generate high-resolution kinetic data. The experimental matrix included pH values ranging from 6.5 to 8.5 and two substrate formulations: pure xylose and a glucose–xylose mixture (1:2). Cultivations were performed over a 24-hour period to capture the lag, exponential, and stationary growth phases. Optical density measurements were collected every 15 minutes in triplicate to ensure reproducibility and statistical robustness.

Surrogate Model Development

The second stage involved the construction of a surrogate model using a multilayer perceptron neural network. The MLP was trained to approximate the nonlinear mapping between the input variables—time and pH—and the output variable, optical density (absorbance). Different network architectures were evaluated by varying the number of neurons and activation functions. Model performance was assessed using standard statistical metrics, including AIC, MSE, RMSE, and the coefficient of determination (R^2).

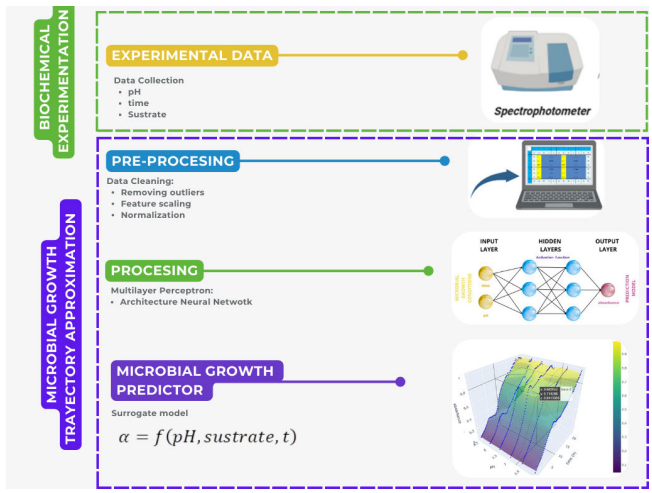

Figure 1. Methodology of Microbial growth trajectory approximation.

Biochemical Experimentation

Microbial growth is conventionally described through four distinct phases: lag, exponential, stationary, and death. The lag phase corresponds to the period immediately following inoculation, during which microorganisms adapt to the new environmental conditions and activate the metabolic pathways required for growth. Once adaptation is complete, the system enters the exponential phase, characterized by rapid cell division and substrate consumption at an approximately constant specific growth rate. As nutrients become depleted and inhibitory metabolites accumulate, the system transitions into the stationary phase, where cell growth and death rates are approximately balanced. Finally, the death phase is marked by a progressive decline in viable biomass, resulting in a significant reduction or cessation of microbial growth. These phases constitute a generalized kinetic framework that applies to a wide range of microorganisms, including bacteria,

fungi, yeasts, and microalgae.

In this study, the kinetic behavior of the yeast *Candida guilliermondii* was selected as the model system. This microorganism is widely distributed in natural environments and can be readily isolated from fruits and berries with high sugar content. In addition, it is commonly found as part of the native microflora of human skin and mucous membranes. Beyond its ecological prevalence, C. guilliermondii is of considerable biotechnological interest due to its ability to act as a biological control agent and its well-documented capacity to bioconvert xylose into xylitol, a high-value polyol with applications in the food and pharmaceutical industries. This metabolic capability relies on the efficient production of enzymes and cellular energy required for pentose assimilation.

Under optimal conditions, C. guilliermondii exhibits maximal growth at approximately pH 7 and a temperature of 28 °C. However, industrial bioprocesses often operate under suboptimal or fluctuating environmental conditions. For this reason, the present work explores a broader pH range, from 6.5 to 8.5, in order to capture the adaptive kinetic response of the microorganism. These controlled pH variations enable the observation of different growth regimes over a 24-hour cultivation period, thereby generating a comprehensive dataset that spans multiple growth phases.

Given the relevance of C. guilliermondii in xylose bioconversion, two substrate formulations were evaluated: a culture medium containing xylose as the sole carbon source and a mixed-substrate medium consisting of a glucose–xylose blend at a 1:2 ratio. The inclusion of glucose aims to assess its influence on reducing the lag phase and accelerating the onset of exponential growth, a strategy commonly employed to enhance process productivity in mixed-sugar fermentations.

All experiments were conducted in triplicate to ensure reproducibility and statistical reliability. Experimental analyses were performed at the Microbial Interaction Laboratory of the National Technological Institute of Mexico, Irapuato campus. Microbial growth was monitored via optical density measurements at 600 nm, which served as a proxy for biomass concentration. Measurements were recorded every 15 minutes over a 24-hour period at a constant temperature of 28 °C using a Thermo Scientific Multiskan™ FC Microplate Photometer. Cultivations were carried out in multiwell plates, allowing for high-throughput data acquisition and consistent experimental conditions across all runs.

Microbial Growth Trajectory Approximation

The second stage of the proposed methodology focuses on learning the underlying data patterns from the experimentally obtained microbial growth kinetics of C. guilliermondii. To this end, a data-driven surrogate model based on machine learning was developed using a multilayer perceptron (MLP) neural network. The well-established capabilities of modern artificial intelligence techniques, particularly neural networks, provide an efficient and flexible framework for predicting microbial growth under previously unobserved combinations of operating conditions.

The MLP architecture employed in this study consists of three main components: (i) an input layer, (ii) one or more hidden layers, and (iii) an output layer. Input signals are propagated forward through the network, where each neuron applies a nonlinear transformation to a weighted linear combination of its inputs. The final output of the network is a scalar value representing the predicted absorbance. In this work, the MLP is used as a nonlinear function approximator and predictive model to estimate the absorbance α as a function of time t and the pH of the culture medium.

From a theoretical standpoint, a multilayer perceptron is a universal approximator of continuous nonlinear functions, which makes it particularly suitable for modeling complex biological and biochemical processes. Formally, the MLP implements the functional mapping described by Equation (1), where the output is expressed as a composition of nonlinear activation functions parameterized by a set of weights and biases.

This modeling approach allows the MLP to implicitly capture phase transitions and nonlinear growth dynamics without explicitly segmenting the growth curve, thereby offering a significant advantage over classical kinetic models when dealing with dynamic and non-ideal operating conditions.

From a formal perspective, the MLP implements the function described by Equation 1.

$$y = f(x, \omega) \qquad \text{Eq. (1)}$$

Where f is a composition of nonlinear functions parameterized by weights and biases, ω.

The experimental setup is as follows: The number of neurons in the hidden layer was varied to determine the optimal topological architecture for our artificial neural network. Three activation functions—Sigmoid, ReLu, and Tanh—were analyzed. We implemented the Adam optimizer to automatically adjust the learning rate of each parameter, allowing for faster and more stable convergence. The number of neurons is as follows: 2, 4, 6, 8, and 10. We used a dataset of 540 feature vectors. We use the widely tested least squares function as the fitness function during the training stage. We validate our model using standard growth factors. We trained the machine learning model using three pH values (6.5, 7.5, 8, and 8.5) and substrate formulations in the culture medium (xylose at 100% and glucose-xylose at a ratio of 1:2). Our model predicts absorbance values for any combination of pH between 6.5 and 8.5, substrate (xy-

Table 1. Summary of the metrics that were obtained from the different architecture of the multilayer perceptron (MLP).

Configuration	Architecture		AIC			MSE			RMSE			R^2		
	Hidden layers	Neurons	Training	Validation	Testing	Training	Validation	Testing	Training	Validation	Testing	Training	Validation	Testing
1	2	2	-1822.761	-244.757	-518.618	0.008	0.007	0.007	0.088	0.082	0.083	0.910	0.937	0.926
2	2	6	-1825.331	-151.625	-451.464	0.006	0.006	0.005	0.076	0.075	0.070	0.932	0.948	0.947
3	4	2	-1583.158	-164.052	-408.564	0.014	0.019	0.015	0.117	0.138	0.122	0.842	0.822	0.839
4	4	6	-1775.455	4.179	-307.745	0.004	0.004	0.004	0.065	0.067	0.063	0.950	0.958	0.958
5	6	2	-855.842	-42.732	-183.566	0.086	0.111	0.093	0.294	0.334	0.305	0.000	-0.042	0.000
6	6	6	-463.842	349.268	208.434	0.086	0.111	0.093	0.294	0.334	0.305	0.000	-0.042	0.000
7	8	2	-831.842	-18.732	-159.566	0.086	0.111	0.093	0.294	0.334	0.305	0.000	-0.042	0.000
8	8	6	-1484.352	363.763	39.159	0.004	0.007	0.004	0.062	0.083	0.066	0.956	0.936	0.953
9	10	2	-807.842	5.269	-135.566	0.086	0.111	0.093	0.294	0.334	0.305	0.000	-0.042	0.000
10	10	6	-1146.487	522.436	243.249	0.006	0.006	0.006	0.077	0.076	0.077	0.931	0.946	0.935

lose or glucose-xylose), and timestamp (hours) between 0 and 24.

RESULTS

In this work, a machine-learning–based framework was developed for the estimation and prediction of microbial growth kinetics, as schematically illustrated in Fig. 6. The proposed framework comprises two main stages: model training and model testing. During the training stage, experimentally obtained kinetic data are used to fit the parameters of the neural network and to quantify the training error, expressed in terms of the mean squared error (MSE train). This stage results in the construction of the HC-MLP predictor model, which encapsulates the learned nonlinear relationship between the input variables and the microbial growth response.

To rigorously assess the generalization capability of the trained model, the testing stage employs an independent replicate of the experimental kinetic data that was not used during training. The prediction error obtained at this stage, quantified as MSE test, provides an unbiased evaluation of the model's predictive performance. Based on the results of the testing stage, the final surrogate model is obtained. This surrogate model generates a continuous response surface that represents the MLP approximation of the output variable—optical density (absorbance)—as a function of the input variables, time (t) and pH, for culture media containing xylose as the carbon source.

A case study was conducted to evaluate the predictive capability of the proposed MLP framework. The objective of this analysis was to assess the model's ability to accurately predict microbial growth for arbitrary combinations of pH and time within the studied domain. Specifically, the model was used to estimate absorbance values over a time horizon ranging from 0 to 24 h and pH values between 6.5 and 8.5, assuming xylose as the sole substrate. To identify the optimal network topology, the number of neurons in the hidden layer was systematically varied, allowing for an evaluation of the trade-off between model complexity and predictive accuracy.

To select the most appropriate neural network architecture, several complementary performance metrics were employed, including the Akaike information criterion (AIC), mean squared error (MSE), root mean squared error (RMSE), and the coefficient of determination (R^2). The AIC was used as the primary criterion for model selection, as it simultaneously accounts for goodness of fit and model parsimony by penalizing excessive complexity. Lower AIC values indicate a more favorable balance between predictive accuracy and information loss. The MSE quantifies the average squared deviation between predicted and observed values and was primarily used to compare model performance during the training phase. The RMSE, expressed in the same units as the experimental measurements, provides an intuitive measure of prediction error and was used for both validation and testing stages. Finally, the coefficient of determination (R^2) evaluates the proportion of variance in the experimental data explained by the model, with values approaching unity indicating excellent agreement between predictions and observations.

Table 1 summarizes the values of all performance metrics for the several MLP architectures evaluated in this study. Each metric is reported for the training, validation, and testing stages, providing a comprehensive basis for selecting the optimal neural network configuration.

As evidenced by the results summarized in Table 1, configuration 2 exhibits the most robust and consistent performance across all selected evaluation metrics. This configuration yields the lowest Akaike information criterion (AIC) value during the training phase (–1825.33), indicating an optimal trade-off between model fidelity and structural complexity. Comparable AIC behavior is observed in the validation and testing phases, which

further suggests that the selected architecture generalizes well beyond the training dataset.

In terms of predictive accuracy, configuration 2 achieves one of the lowest mean squared error (MSE) values, equal to 0.005842, demonstrating a high-quality fit to the experimental data. A similar trend is observed for the root mean squared error (RMSE), with values of 0.074527 and 0.069864 obtained for the validation and testing stages, respectively. These low RMSE values, expressed in the same units as the measured absorbance, confirm the model's ability to accurately reproduce microbial growth dynamics across the studied domain. Furthermore, the coefficient of determination (R^2) for this configuration approaches unity in all stages, indicating that the surrogate model explains nearly all of the variance observed in the experimental measurements.

Taken together, these results indicate that configuration 2 provides the most favorable balance between predictive performance, model parsimony, and computational efficiency. Consequently, the optimal MLP architecture corresponds to a network with two hidden layers comprising six neurons each, which minimizes the risk of overfitting while maintaining high predictive capability.

Figure 2 illustrates the three-dimensional response surface generated by the surrogate model corresponding to configuration 2 of the multilayer perceptron neural network. In this representation, the experimental data points are shown as blue markers, while the continuous surface depicts the MLP-based approximation of the microbial growth trajectory within the defined data domain. The smoothness and continuity of the predicted surface, together with its close agreement with the experimental observations, demonstrate the model's ability to capture the underlying nonlinear kinetics of *Candida guilliermondii* growth as a function of time and pH.

As illustrated in Figure 2, red neuronal perception can effectively predict the four phases of microbial growth: lag, exponential, stationary, and death. This observation is consistent across the pH range analyzed in the experimental setting (6.5–8.5). The resulting model is able to predict the trajectory of microbial kinetics from a point within the experimental data space. This trajectory can be verified by the gray line shown as an example in Figure 2. This trajectory can be verified at any point in the analyzed space and yields the corresponding coordinate values. Therefore, if we are aware of this condition, we can predict the microorganism's growth phase and implement the appropriate control measures to ensure the highest possible yield of the biological reaction.

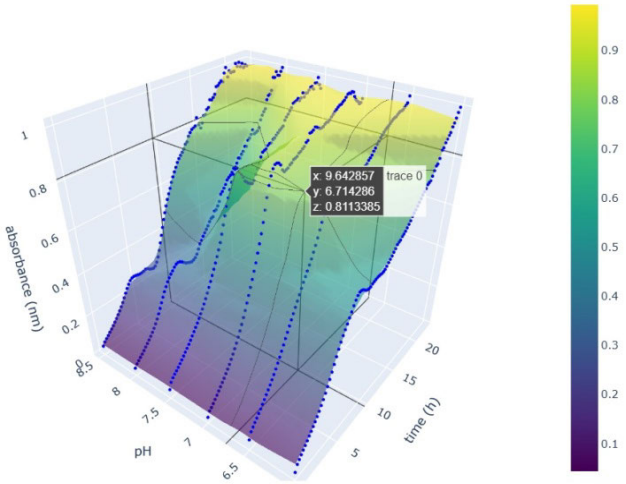

Figure 2. Best surrogate model of configuration 2, for predicted microbial growth kinetics trajectory of *Candida guillermondii.*

CONCLUSIONS

This study presents a robust data-driven framework for modeling microbial growth kinetics under dynamic environmental conditions using a multilayer perceptron neural network. Unlike classical kinetic models, the proposed surrogate model successfully captures nonlinear microbial behavior across multiple growth phases without relying on restrictive mechanistic assumptions.

The results demonstrate that the MLP-based model provides accurate predictions of microbial growth trajectories for *Candida guilliermondii* under varying pH and substrate conditions. This approach significantly enhances the flexibility and predictive capability of kinetic modeling tools, making it particularly suitable for applications in bioreactor design, optimization, and real-time control.

From an engineering perspective, the developed surrogate model enables rapid scenario analysis and supports the automation of experimental workflows, thereby reducing experimental costs and development time. Future work will focus on extending the model to include additional environmental variables, such as temperature and substrate inhibition, as well as integrating the surrogate model into advanced process control and digital twin frameworks for industrial bioprocesses.

AUTHOR IDENTIFIERS

Cortez-González J: 0000-0003-1613-0371
Segovia-Hernández J.G: 000-0003-4175-9809
Hernández S: 0000-0002-4598-3392
López-Ramírez V: N/A
Hernández-Aguirre A: 0000-0002-3744-9827
Murrieta-Dueñas R: 0000 0001 5759 2500

REFERENCES

1. Akkermans S, Van Impe JF. Mechanistic modelling of the inhibitory effect of ph on microbial growth. Food Microbiology 72:214-219 (2018). https://doi.org/10.1016/j.fm.2017.12.007
2. Lee E, Jalalizadeh M, Zhang Q. Growth kinetic models for microalgae cultivation: a review. Algal Research 12:497-512 (2015). https://doi.org/10.1016/j.algal.2015.10.004
3. Moreno?Paz S, Schmitz J, Martins dos Santos VAP, Suarez?Diez M. Enzyme?constrained models predict the dynamics of saccharomyces cerevisiae growth in continuous, batch and fed?batch bioreactors. Microbial Biotechnology 15:1434-1445 (2022). https://doi.org/10.1111/1751-7915.13995
4. Murrieta-Dueñas R, Serrano-Rubio JP, López-Ramírez V, Segovia-Dominguez I, Cortez-González J. Prediction of microbial growth via the hyperconic neural network approach. Chemical Engineering Research and Design 186:525-540 (2022). https://doi.org/10.1016/j.cherd.2022.08.021
5. Öksüz HB, Buzrul S. MONTE CARLO ANALYSIS FOR MICROBIAL GROWTH CURVES. J microb biotech food sci 10:418-423 (2020). https://doi.org/10.15414/jmbfs.2020.10.3.418-423

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Dynamic optimization of glucose feed in cell cultivation for monoclonal antibody production process design balancing productivity and impurity generation

Kosuke Nemoto[a], Yuki Yoshiyama[a], Mizuki Morisasa[b], Junshin Iwabuchi[b], Yusuke Hayashi[a], Sara Badr[a], Hirokazu Sugiyama[a]*

[a] The University of Tokyo, Department of Chemical System Engineering, 7-3-1, Hongo, Bunkyo-ku, Tokyo, Japan
[b] Chitose Laboratory Co., Ltd., Tech & Biz Development Division, KSP R&D C432, 3-2-1, Sakado, Takatsu-ku, Kawasaki, Kanagawa, Japan
* Corresponding Author: sugiyama@chemsys.t.u-tokyo.ac.jp

ABSTRACT

This work presents the dynamic optimization of glucose feed in cell cultivation considering the balance between productivity and impurity generation. We first developed a mechanistic model considering cell growth promotion by glucose and cell growth inhibition by osmolarity for a newly developed, high-productivity CHO-MK cell line. For model development, fed-batch cultivation experiments were conducted at a 250 mL scale under three different glucose feeding profiles. Results from a single-objective dynamic optimization, using the glucose feed profile as a design variable, were compared to those from multi-objective problem settings with varying weights assigned to productivity and final impurity concentrations. Simulation results suggested different glucose feed profiles depending on the priority given to the mAb and impurities, where the main difference was in the generated viable cell density profiles. Productivity-focused profiles employed a low–high–intermediate feeding strategy, in which the total glucose feed over the cultivation period was set to a balanced level, with an initially low rate, followed by a gradual increase, and a later decrease. Profiles focusing on impurity reduction adopted a feeding strategy that maintained a high glucose feed throughout the cultivation period. This work demonstrated the power of mechanistic-model-based process design and is expected to encourage further model-based process design in biopharmaceutical manufacturing.

Keywords: Process Design, Dynamic Modelling, Simulation, Biosystems

INTRODUCTION

Monoclonal antibody (mAb) drugs are effective for the treatment of serious diseases, leading to rapid expansion of the mAb market and the development of new technologies, including host cell lines such as the newly developed high-productivity CHO-MK cell line [1]. In the cultivation step, cells grow and produce mAb while consuming nutrients such as glucose. Impurities such as host cell protein (HCP) are released from viable and dead cells, whereas impurities such as charge variants are acidic or basic modifications of the mAb. These impurities need to be removed during the succeeding purification step, which is costly, making it crucial to maximize mAb yield while minimizing impurity concentrations during cultivation. To achieve this objective, feed rate of glucose, which serves as a primary nutrient for cells, has been reported to play an important role. Insufficient glucose supply can lead to its depletion as cell density increases, whereas excessive supply can result in high osmolarity, both of which negatively affects cell growth [2].

Design of the cultivation step should carefully consider dynamic changes and cell phases (such as stationary and exponential growth phases) by accordingly adjusting process parameters. Several studies have proposed experimental approaches for controlling glucose feed rates, in which feeding strategies were adjusted in response to changes in measurable physical parameters (such as glucose concentrations [3] and oxygen uptake rates [4]). Although these experimental approaches can

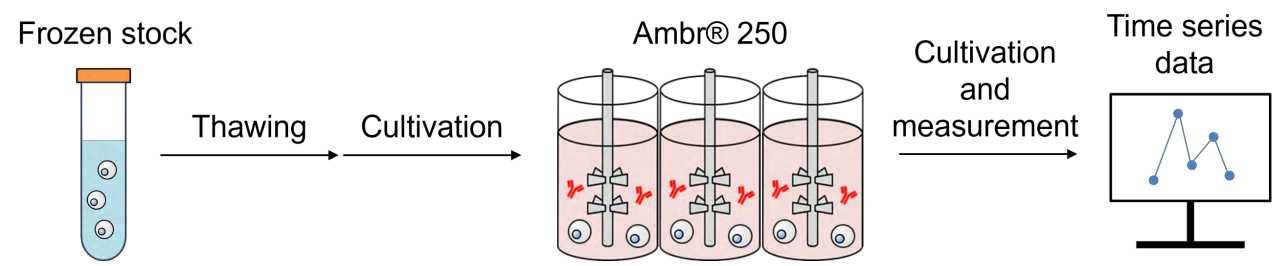

Figure 1. Overview of the cultivation experiments.

be useful, they are often limited in exploration range due to the reliance on empirically determined setpoints. As an alternative, mathematical models offer a powerful means to efficiently explore combinations of process parameters. In fed-batch cultivation, variables such as glucose feed rates [5], and in perfusion cultivation, bleed rates [6], have been employed as design variables to maximize the final mAb product concentration. While some studies considered the balance between productivity and impurity concentrations such as aggregates through multi-objective optimization [7], generally the consideration of quality attributes in these works remain limited.

This work presents dynamic optimization of glucose feed in cell cultivation considering the balance between mAb productivity and the generation of multiple impurities. We first developed a mechanistic model that considers cell growth promotion by glucose and cell growth inhibition by osmolarity for a newly developed, high-productivity CHO-MK cell line. For model development, fed-batch cultivation experiments were conducted at a 250 mL scale under three different glucose feeding profiles. The developed model was then applied to dynamic optimization with the glucose feed profile as a design variable. Single-objective optimization aimed at maximizing mAb productivity and multi-objective optimization considering impurities were compared. The results suggested different glucose feed profiles depending on the relative priority of mAb and impurities, with the main differences observed in the generated viable cell density profiles.

METHODS

Experimental methods

Figure 1 represents an overview of the experimental methods. To obtain experimental data for modeling, cultivation experiments were conducted using a newly developed CHO-MK cell line [1] was used. Cells were extracted from a frozen stock, thawed, and cultured in flasks for one week to increase cell density. Subsequently, cultivation was carried out with Ambr® 250 at a 250 mL scale. Initial solution volume (V_0 [L]) was set to 0.17 in all experiments. Three fed-batch cultivation experiments, each lasting one week, were conducted at a 250 mL scale using Ambr® 250 under different glucose feed rate conditions ($F_{\mathrm{in,glc}}$ [V_0 g day^{-1}]: 6, 11, and 15), hereafter referred to as Experiments (A), (B), and (C), respectively. The values of $F_{\mathrm{in,glc}}$ were set within the range of 6–15 V_0 g day^{-1} commonly used by cultivation experts. In the experiments, medium feed rate ($F_{\mathrm{in}} = 8$, 10 % V_0 day^{-1}), feed start time ($t_{\mathrm{in}} = 48$ h), glucose feed start time ($t_{\mathrm{in,glc}} = 48$ h), agitation rate ($\omega = 1$, 200 rpm), and dissolved oxygen ($DO = 20$ %) were set as specified. In addition to viable cell density, the concentrations of multiple components (e.g., mAb, impurities, and metabolites) were measured offline once or twice per day Table 1 shows the values of $F_{\mathrm{in,glc}}$ in the three experiments.

Table 1: Values of $F_{\mathrm{in,glc}}$ in Experiments (A)–(C).

Experiments	(A)	(B)	(C)
$F_{\mathrm{in,glc}}$ [V_0 g day^{-1}]	6	11	15

Model development

Mechanistic model equations

A mechanistic model was developed as shown by Eqs. (1)–(6) in Table 2, where V [L], X_{v} [cells L^{-1}], X_{d} [cells L^{-1}], $[\mathrm{GLC}]$ [mmol L^{-1}], $[\mathrm{mAb}]$ [g L^{-1}], $[\mathrm{Main}]$ [g L^{-1}], $[\mathrm{Acid}]$ [g L^{-1}], $[\mathrm{HCP}]$ [g L^{-1}], Osm [mOsm kg^{-1}] are the solution volume, viable, and dead cell densities, glucose, mAb, main/acidic/basic charge variant, and HCP concentrations, as well as osmolarity. μ [h^{-1}] and μ_{d} [h^{-1}] represent specific cell growth and death rates, respectively. μ_{max} [h^{-1}] is the maximum value of μ described by the Monod type equations. $q_{\mathrm{p,a}}$ [g^{-1} cell] and $q_{\mathrm{p,b}}$ [g cell^{-1} h^{-1}] are the growth-dependent and growth-independent constants for mAb production, respectively. $k_{i,\mathrm{lysis}}$ [h^{-1}] is the rate constant of the dead cell lyses and the HCP dissolution, where i includes dead cell and HCP. k_{acid} [h^{-1}] is the rate constant for acidic charge variant formation. $Y_{\mathrm{hcp}/X_{\mathrm{v}},X_{\mathrm{d}}}$ [g cell^{-1}] is the mass of HCP released from a viable cell and a lysed dead cell. f_{osm} [–] is the growth inhibition term due to osmolarity, taking a value in the range of 0–1. K_{osm} [mOsm kg^{-1}] represents the minimum osmolarity of the

Table 2: Mechanistic model equations.

$$\frac{d(VX_\mathrm{v})}{dt} = (\mu - \mu_\mathrm{d})VX_\mathrm{v} \quad (1)$$

$$f_\mathrm{osm} = \begin{cases} \frac{Osm_\mathrm{lim} - Osm}{Osm_\mathrm{lim} - K_\mathrm{osm}}\Big|_{Osm < Osm_\mathrm{lim}} \\ 0|_{Osm \geq Osm_\mathrm{lim}} \end{cases} \quad (2)$$

$$\mu = \mu_\mathrm{max} f_\mathrm{osm}\left(\frac{[\mathrm{GLC}]}{K_\mathrm{glc} + [\mathrm{GLC}]}\right) \quad (3)$$

$$\frac{d(V[\mathrm{mAb}])}{dt} = \left(\frac{\mu - \mu_\mathrm{d}}{q_\mathrm{p,a}} + q_\mathrm{p,b}\right)VX_\mathrm{v} \quad (4)$$

$$\frac{d(V[\mathrm{Acid}])}{dt} = k_\mathrm{acid}V[\mathrm{Main}] \quad (5)$$

$$\frac{d(V[\mathrm{HCP}])}{dt} = Y_{\mathrm{hcp}/X_\mathrm{v}}VX_\mathrm{v} + Y_{\mathrm{hcp}/X_\mathrm{d}}k_{X_\mathrm{d},\mathrm{lysis}}VX_\mathrm{d} - k_\mathrm{hcp,lysis}(V[\mathrm{HCP}])^n \quad (6)$$

Table 3: Design problem for single-objective and multiobjective dynamic optimization.

$$\max \alpha\frac{[\mathrm{mAb}](\boldsymbol{x}) - [\mathrm{mAb}]_\mathrm{min}}{[\mathrm{mAb}]_\mathrm{max} - [\mathrm{mAb}]_\mathrm{min}} + \beta\frac{[\mathrm{Acid}]_\mathrm{max} - [\mathrm{Acid}](\boldsymbol{x})}{[\mathrm{Acid}]_\mathrm{max} - [\mathrm{Acid}]_\mathrm{min}} + \gamma\frac{[\mathrm{HCP}]_\mathrm{max} - [\mathrm{HCP}](\boldsymbol{x})}{[\mathrm{HCP}]_\mathrm{max} - [\mathrm{HCP}]_\mathrm{min}} \quad (7)$$

subject to

$$\boldsymbol{x} = \left(F_\mathrm{in,glc}\big|_{48 \leq t < 72}, F_\mathrm{in,glc}\big|_{72 \leq t < 96}, F_\mathrm{in,glc}\big|_{96 \leq t < 120}, F_\mathrm{in,glc}\big|_{120 \leq t < 144}, F_\mathrm{in,glc}\big|_{144 \leq t}\right)$$

$$F_\mathrm{in,glc} \in \{6, 7, \ldots, 14, 15\}$$

$$\alpha + \beta + \gamma = 1 \qquad t\,[\mathrm{h}] \leq 168 \qquad F_\mathrm{in}\,[\%\,V_0\,\mathrm{day}^{-1}] = 8$$

solution. Osm_lim [mOsm kg^{-1}] is the osmolarity value which represents osmolarity levels with zero growth levels. In this work, osmolarity was considered to depend on the number of dissolved particles per unit weight of solution, based on a previous study [8]. It was expressed as a linear function of the viable and dead cell densities and the glucose concentration which are described in the model.

Model fitting and validation

Model fitting and cross-validation were performed using cultivation data from Experiments (A)–(C). Model parameters were estimated for two cases: one case in which Experiments (A) was used as the validation batch in Case 1, and one case in which all three experiments were used for training (Case 2). The system of ordinary differential equations was solved using the Python package solve_ivp with a time step of 0.001 h. Parameter estimation was performed using the Nelder-Mead optimization method to maximize the sum of coefficients of determination (R^2) between the calculated and measured values of all system components. Parameter uncertainty was analyzed by computing standard errors via inversion of the Hessian matrix, using the Python package lmfit.

Multi-objective dynamic optimization

By using the developed mechanistic model, a design problem was defined as shown in Eq. (7) in Table 3 for performing optimization by adjusting the glucose feed rate $F_\mathrm{in,glc}$ every 24 h starting from the initiation of feeding (48 h after the start of cultivation). In the optimization, both single-objective optimization to maximize mAb productivity and multi-objective optimization considering impurities were conducted. Glucose feed rates on days 2–6 were determined. The values of $F_\mathrm{in,glc}$ were set in the range of 6 to 15 V_0 g day^{-1} with an increment of 1 V_0 g day^{-1} (a total of 10 values) to ensure that the exploration remained within the applicable range of the model. As a result, a total of 10^5 (100, 000) profiles were explored, with the total glucose feed over the entire cultivation period ranging from 30 to 75 V_0 g. The simulation was performed from t = 2 h, which corresponds to the starting time of data measurements in the cultivation experiments, until t = 168 h. The measured values t = 2 h were used as initial conditions. The update frequency of $F_\mathrm{in,glc}$ (i.e., every 24 h) was determined based on expert recommendations. The feed start time (t = 48 h) was set to the same value

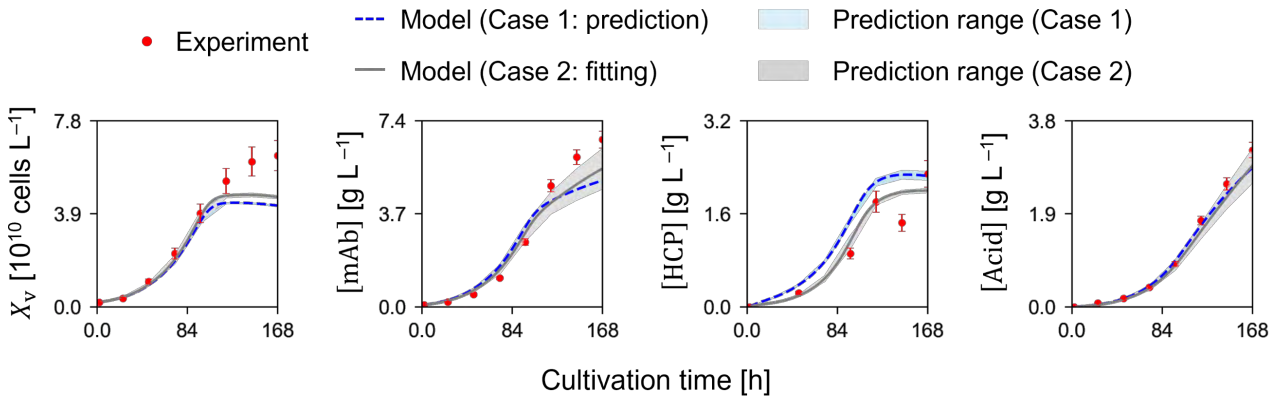

Figure 2. A part of the fitting and validation results for Experiment (A).

Table 4: Average values of R^2 and NRMSE across all modeled components for the validation batch in Case 1 and for all training batches in Case 2.

Case (Validation batch)	1 (A)	2 (All data used for training)
R^2	0.84	0.95
NRMSE	0.13	0.13

as that used in the model training experiments (Experiments (A)–(C)). $[\mathrm{mAb}]_{\max}, [\mathrm{mAb}]_{\min}, [\mathrm{Acid}]_{\max}, [\mathrm{Acid}]_{\min}, [\mathrm{HCP}]_{\max}$, and $[\mathrm{HCP}]_{\min}$ represent the maximum and minimum calculated values of $[\mathrm{mAb}], [\mathrm{Acid}]$, and $[\mathrm{HCP}]$, respectively, at the end of cultivation (t = 168 h) among the 100, 000 profiles considered. α, β, and γ are the weighting values for the normalized values of each component in the objective function, and the sum is always equal to 1. α, β, and γ represent the priorities of mAb productivity and the burden of acidic variants and HCP in the purification process, respectively. To assess the impact of the chosen three weighting factors, two scenarios where (α, β, γ) = (1, 0, 0) and (0.4, 0.3, 0.3) were investigated as Scenarios (i) and (ii), respectively. In Scenario (i), the weighting of α was set to 1, aiming for single-objective optimization considering only mAb productivity. In contrast, in Scenario (ii), multi-objective optimization was performed.

RESULTS AND DISCUSSION

Modeling results

Model fitting and validation were performed for the two cases (Cases 1 and 2), each using different experimental data as training sets. Figure 2 shows the fitting and validation results of the mechanistic model for Experiment (A) as an example of the obtained results. In each figure, the results for Cases 1 and 2 are represented in blue and gray, respectively. Calculated results for the training and validation data are shown as solid and dashed lines, respectively. The shaded prediction range represents the interval between the minimum and maximum predicted profiles, calculated using the estimated parameter values and their standard errors under the assumption that the parameters follow a normal distribution. Table 4 summarizes the R^2 and normalized root mean square error (NRMSE) values averaged over all components for the validation batch in Case 1 and for the three experimental datasets in Case 2. High predictive accuracy was confirmed in the cross-validation of Case 1, with an R^2 value of 0.84. The model parameters estimated in Cases 1 and 2 took similar values, demonstrating the robustness of the model. In the subsequent simulations, the parameter values estimated in Case 2 were used, where all three experimental datasets were used for training.

Multi-objective dynamic optimization

The mechanistic model was applied to find the solution to the single-objective (Scenario (i)) and multi-objective (Scenarios (ii)) design problems (Eq. (7)) shown in Table 3. The glucose feed profiles that maximized the objective function in Scenario (i) and Scenario (ii) were $(F_{\mathrm{in,glc}}|_{48\le t<72}, F_{\mathrm{in,glc}}|_{72\le t<96}, F_{\mathrm{in,glc}}|_{96\le t<120}, F_{\mathrm{in,glc}}|_{120\le t<144}$, and $F_{\mathrm{in,glc}}|_{144\le t})$ = (6, 13, 15, 15, 11) and (15, 15, 15, 15, 15), respectively. Figure 3 shows a part of the simulation results based on the optimized profile in each scenario, where

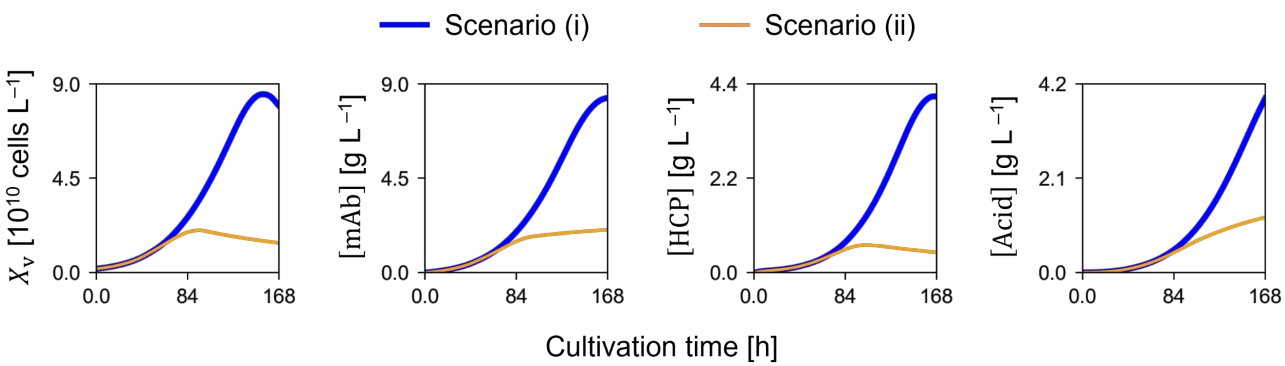

Figure 3. Simulation results based on the optimized profile in the two scenarios.

different glucose profiles and simulation outcomes were observed. In the two scenarios, based on the combination of weighting values, glucose feed profiles were appropriately set to adjust viable cell density, which would directly affect mAb and impurity concentrations, for achieving a suitable balance between nutrient availability and osmolarity. The following points were considered as the factors that led to the selection of the optimized profile in the two scenarios.

<Scenario (i)>

In the optimized profile, the total glucose feed over the entire cultivation period, which was explored within the range of 30–75 V_0 g, was found to take a value of 60 V_0 g. This value is considered to reflect a balance aimed at avoiding both glucose depletion due to excessively low feeding and increased osmolarity due to excessively high feeding. The optimized profile exhibited a low-high-intermediate feeding strategy: the feed rate was kept low (6 V_0 g day^{-1}) during the early stage of cultivation (t = 48–72 h), then gradually increased (t = 72–96 h: 13 V_0 g day^{-1}; t = 96–144 h: 15 V_0 g day^{-1}), and subsequently reduced toward the late stage of cultivation (t = 144–168 h) to 11 V_0 g day^{-1}. In this low–high–intermediate strategy, glucose feed rates were kept low during the early stage of cultivation to avoid an increase in osmolarity. Subsequently, during the mid-to-late stage when cells grow exponentially, the feed rates were increased, implying that a sufficient amount of glucose was supplied.

<Scenario (ii)>

In the optimized profile, the total glucose feed was 75 V_0 g with glucose feed rates maintained at high levels throughout the cultivation period. This outcome can be explained as follows: high values of $F_{\mathrm{in,glc}}$ lead to growth inhibition due to elevated osmolarity, resulting in a decrease in viable cell density as well as a concomitant decrease in dead cell density. Since mAb and HCP originate from viable and dead cells, and acidic variants originate from mAb, the reduction in viable cell density is considered to contribute to the decrease in impurity concentrations observed in the simulation results.

CONCLUSIONS AND OUTLOOK

This work presented dynamic optimization of glucose feed in cell cultivation considering the balance between productivity and impurity concentrations. A mechanistic model was developed to describe concentration profiles of mAb as well as multiple impurities, considering cell growth promotion by glucose and cell growth inhibition due to osmolarity for a newly developed, high-productivity CHO-MK cell line. The developed model enabled the optimization of glucose feed profiles to maximize mAb yield, as well as multi-objective optimization considering multiple impurities. The simulation results indicated that distinct glucose feed profiles were obtained depending on whether priority was placed on mAb production or impurity reduction, with the primary difference reflected in the resulting viable cell density profiles. In the productivity-focused profile, the total glucose feed over the entire cultivation period was set to a balanced level (60 V_0 g) within the explored range of 30–75 V_0 g, and a low–high–intermediate feeding strategy was adopted. In contrast, the impurity-focused profile employed a strategy that maintained the maximum glucose feed within the explored range (15 V_0 g day^{-1}) throughout the cultivation period. This work demonstrated the power of mechanistic-model-based process design in biopharmaceutical manufacturing. The future work includes experimental validation of the proposed low-high-intermediate feeding strategy.

DIGITAL SUPPLEMENTARY MATERIAL

The raw experimental data used for Figure 2 are available as Table S1 in LAPSE:2026.0040.

ACKNOWLEDGEMENTS

This work was supported by the Japan Agency for Medical Research and Development (AMED) under Grant Number JP21ae0121015 and by JST SPRING under Grant Number JPMJSP2108, and Leadership Development Program for Ph.D., the University of Tokyo. The authors would like to thank Dr. Takayuki Horiuchi for providing general support for this work.

REFERENCES

1. Masuda K, Kubota M, Nakazawa Y, Iwama C, Watanabe K, Ishikawa N, Tanabe Y, Kono S, Tanemura H, Takahashi S, Makino T, Okumura T, Horiuchi T, Nonaka K, Murakami S, Kamihira M, Omasa T. Establishment of a novel cell line, CHO-MK, derived from chinese hamster ovary tissues for biologics manufacturing. Journal of Bioscience and Bioengineering 137:471-479 (2024). https://doi.org/10.1016/j.jbiosc.2024.02.005
2. Fan Y, Jimenez Del Val I, Müller C, Lund AM, Sen JW, Rasmussen SK, Kontoravdi C, Baycin?Hizal D, Betenbaugh MJ, Weilguny D, Andersen MR. A multi?pronged investigation into the effect of glucose starvation and culture duration on fed?batch CHO cell culture. Biotech & Bioengineering 112:2172-2184 (2015). https://doi.org/10.1002/bit.25620
3. Sukwattananipaat P, Kuroda H, Yamano?Adachi N, Omasa T. Metabolomic characterization of monoclonal antibody?producing chinese hamster lung (CHL)?YN cells in glucose?controlled serum?free fed?batch operation. Biotech & Bioengineering 121:2848-2867 (2024). https://doi.org/10.1002/bit.28777
4. Reyes S, Pham PL, Durocher Y, Henry O. cho stable pool fed?batch process development of sars ? cov ?2 spike protein production: impact of aeration conditions and feeding strategies. Biotechnology Progress 41: (2024). https://doi.org/10.1002/btpr.3507
5. Jones W, Gerogiorgis D. Dynamic optimisation of fed-batch bioreactors for mabs: sensitivity analysis of feed nutrient manipulation profiles. Processes 11:3065 (2023). https://doi.org/10.3390/pr11113065
6. Jones W, Gerogiorgis DI. Dynamic simulation, optimisation and economic analysis of fed-batch vs. perfusion bioreactors for advanced mab manufacturing. Computers & Chemical Engineering 165:107855 (2022). https://doi.org/10.1016/j.compchemeng.2022.107855
7. Jones W, Gerogiorgis DI. Dynamic optimization of an integrated cultivation?aggregation model for mab production. Biotech & Bioengineering 121:2716-2727 (2024). https://doi.org/10.1002/bit.28761
8. Martín-Calderón JL, Bustos F, Tuesta-Reina LR, Varona JM, Caballero L, Solano F. Choice of the best equation for plasma osmolality calculation: comparison of fourteen formulae. Clinical Biochemistry 48:529-533 (2015). https://doi.org/10.1016/j.clinbiochem.2015.03.005

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Development of ANN-based models for dye removal through electrochemical advanced oxidation techniques

Zaira J. Mosqueda-Huerta[a], Oscar D. Lara-Montaño[b], Juan Manuel Peralta-Hernández[c], and Fernando I. Gómez-Castro[a]*

[a] Universidad de Guanajuato, Campus Guanajuato, División de Ciencias Naturales y Exactas, Departamento de Ingeniería Química, Guanajuato, Guanajuato, México
[b] Universidad Autónoma de Querétaro, Campus Amealco, Facultad de Ingeniería, Amealco de Bonfil, Querétaro, México
[c] Universidad de Guanajuato, Campus Guanajuato, División de Ciencias Naturales y Exactas, Departamento de Química, Guanajuato, Guanajuato, México
* Corresponding Author: fgomez@ugto.mx.

ABSTRACT

In this work, artificial neural networks are used to represent advanced oxidation processes, such as electrochemical oxidation, electro-Fenton, and photoelectro-Fenton, for the degradation of two dyes. The effect of treatment time, initial dye concentration, and current density on the degradation percentage is studied. Additionally, a network is developed to include discrete variables, such as treatment type and dye type, as input features, enabling it to predict the system's performance across different technologies and pollutants. Operating conditions are optimized using the universal ANN as a surrogate model and the adaptive differential evolution algorithm to maximize dye removal efficiency. According to the results, after optimizing the architecture of the artificial neural networks using Bayesian optimization, deviations of 3.9% or less are obtained for purple RL removal predictions, while for Green A, deviations of 8% or less are obtained. These models may serve as a basis for more robust representations in the future, including the effects of electrochemical device size, among other relevant variables.

Keywords: artificial neural networks, dye removal, electrochemical processes

INTRODUCTION

Water is a highly important resource for humankind, fundamental for the development of life on the planet. Unfortunately, water scarcity has increased over the last few years in several countries [1, 2], underscoring the urgent need for strategies to reduce the water footprint of human activities. Among the actions to reduce such a footprint, the treatment and reuse of industrial wastewater have gained importance, as they reduce the need for fresh water in industrial processes [3]. These actions are highly relevant, as approximately 25-30% of the total wastewater discharged into water channels comes from the industrial sector [4, 5].

In particular, the textile industry discharges large volumes of polluted water, including dyes and other chemicals that must be removed to prevent damage to aquatic plants and animals [6]. Another industry with a high contribution of dyes to wastewater is the tannery industry, due to the materials used to color leather in the post-tanning stage [7]. Electrochemical advanced oxidation processes, such as electrochemical oxidation (EO), electro-Fenton (EF), and photoelectro-Fenton (PEF), have proven to be highly efficient technologies for degrading dyes from wastewater [8]. Electrochemical oxidation involves the direct oxidation of the pollutant at the anode surface, mediated by hydroxyl radicals [9]. On the other hand, the EF and PEF processes use the Fenton reaction to generate hydroxyl radicals from hydrogen peroxide. In the electro-Fenton approach, H2O2 is generated internally by electrical power [10], whereas the PEF approach involves exposing the solution to ultraviolet light [11].

Rigorous mathematical models have been reported to represent the performance of electrochemical advanced oxidation technologies. Cañizares et al. [12] visualized the system as three interconnected stirred-tank reactors, one where a chemical reaction occurred, the

other two corresponding to the cathodic and anodic reaction zones. Similarly, Chen et al. [13] presented a model that accounted for mass transfer due to electromigration, surface adsorption and desorption, and oxidation at the anode. Unfortunately, the available models have diverse limitations stemming from the assumptions they make. Since kinetic models require information on reaction stoichiometry and kinetic parameters, it is necessary to understand how each type of pollutant reacts. Moreover, it is necessary to estimate the mass transfer coefficients for the pollutant under analysis. Additionally, using actual wastewater introduces additional complexity in obtaining the parameters required to solve the mathematical model and obtain reliable information about the system. As an alternative to overcome such limitations, the use of artificial intelligence-based models has been proposed in recent years to represent the interactions among variables in systems whose mechanistic models are complex or inaccurate. Among the artificial intelligence tools used to represent the relationships between variables in a given system, artificial neural networks (ANNs) are powerful and widely used in recent years for a variety of applications. ANNs mimic the learning capacities of the human brain to analyze and process data. By this approach, patterns in complex systems can be identified by providing sufficient data, and it is theoretically capable of approximating any nonlinear functionality [14]. Thus, these techniques can be adequate for representing the relationships among the relevant variables associated with electrochemical advanced oxidation processes.

CASE STUDY

The removal of two dyes commonly used in the tannery industry (Green A and Purple RL) was investigated using three electrochemical advanced oxidation processes: EO, EF, and PEF. For each experiment, 100 mL of dye solution at initial concentrations of 30, 50, and 80 ppm was treated. The electroactive area was 2 cm², using a Dimensionally Stable Anode–Boron Doped Diamond (DSA–BDD) anode (electroactive area 2 cm²) and a Pt cathode, with Na_2SO_4 as the supporting electrolyte. The pH was adjusted to 3.0 and maintained, and three current densities were applied: 25, 35, and 50 mA cm^{-2}. All experiments were conducted over 60 minutes, with samples collected every 10 minutes (Figure 1). The percentage of pollutant removal was determined as the response variable by UV–Vis spectrometry. The full set of variables used to develop the ANN models and their values are summarized in Table 1.

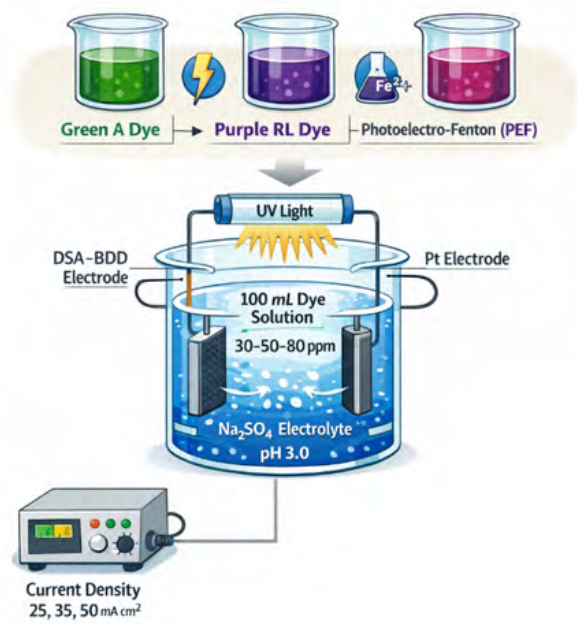

Figure 1. Schematic of the EO, EF, and PEF system for removing Green A and Purple RL dyes, showing the DSA–BDD anode and Pt cathode, Na_2SO_4 electrolyte at pH 3.0, current densities (25–50 mA cm^{-2}), UV light for PEF, and UV–Vis monitoring of dye removal.

Table 1: Variables used in the ANN models.

Variable	Description	Values
D	Dye type (1 = Green A, 2 = Purple RL)	1, 2
P	Electrochemical process	1, 2, 3
C	Initial dye concentration (ppm)	30, 50, 80
I	Current density (mA cm^-2)	25, 35, 50
t	Treatment time (min)	0, 10, 20, 30, 40, 50, 60
Removal (%)	Dye removal efficiency	0-100

METHODOLOGY

A total of 378 experimental data points were obtained (63 points for each Dye-Process combination). The input variables selected for the ANN models were dye concentration, current density, and treatment time. For the ANN models developed at the colorant level, the type of electrochemical process was additionally included as an input variable, encoded numerically as EF = 1, EO = 2, and PEF = 3.

The experimental dataset was randomly split into training and test sets at 80–20%. Prior to training, all continuous input variables were standardized using a StandardScaler to ensure comparable magnitudes and to

improve training stability. The output variable (% removal) was kept on its original physical scale.

Artificial neural network modeling

Two ANN modeling strategies were adopted in this work. First, individual ANN models were developed for each dye–process combination, yielding 6 neural networks in total. These models were used exclusively for predictive purposes, enabling the assessment of removal behavior under fixed electrochemical conditions without optimization. The optimal architecture for each ANN was determined using Bayesian optimization (Table 2).

All networks were trained by backpropagation, and the training optimizer (Adam or stochastic gradient descent, SGD) was selected as part of the Bayesian optimization search space.

Second, a single universal ANN model was developed by integrating both dyes and all electrochemical processes into one predictive framework. In this model, the dye type and the electrochemical process were included as discrete input variables, while dye concentration, current density, and treatment time were experimentally defined at discrete levels but treated as continuous variables within the ANN and optimization framework. A repair mechanism was applied during optimization to ensure that all evaluated solutions corresponded to experimentally feasible operating conditions. The universal ANN was optimized and used as a surrogate model for process optimization.

Table 2: Search space for Bayesian optimization of ANN architectures.

Variables	Values
Number of layers	[1 - 4]
Number of neurons	[1 - 16]
Activation function	[ReLU, ELU, Leaky ReLU]
Optimizer	[Adam, SGD]

Physical consistency and output constraints

Although the ANN output layer employed a linear activation function to avoid saturation and preserve sensitivity near optimal conditions, no explicit output constraint was enforced during training. Consequently, the model may sometimes predict removal efficiencies that exceed the physically feasible range, especially when extrapolating outside of the experimental domain. To ensure physical consistency, predictions were clipped post-processing to the range 0–100% during the optimization stage.

It is important to note that this clipping operation does not improve the predictive performance of the ANN nor correct any model deficiencies. Instead, it introduces a limitation in the interpretation of predictions: while clipped values prevent unphysical results, they mask potential overestimation in regions where the model extrapolates beyond the available data. Therefore, clipped predictions should be viewed as upper-bound estimates rather than definitive predictions, and experimental validation remains essential to confirm the actual removal efficiencies under the predicted conditions. In the future, further model refinements could address this limitation and improve accuracy, potentially eliminating the need for post-processing clipping.

Optimization of operating conditions

The universal ANN model was used as a surrogate objective function to determine the operating conditions that maximize dye removal efficiency. The optimization was performed using the adaptive differential evolution algorithm (JADE). JADE was used for operating-condition optimization with the following parameters: population size of 20, 50 generations, mutation factor (default value), and crossover rate (default value). The convergence criterion was based on a threshold for improvement in the fitness value.

The optimization was conducted independently for each dye, with the dye identifier fixed prior to execution. The decision variables included the electrochemical process, initial dye concentration, current density, and treatment time. Discrete variables were handled through a repair mechanism, which involved rounding the values to the nearest experimentally feasible levels. This repair mechanism ensured that all candidate solutions corresponded to experimentally feasible conditions.

Model validation and interpretation

The predictive performance of the ANN models was evaluated using the test dataset through quantitative error metrics, including the Mean Absolute Error (MAE) and standard deviation of the predicted removal efficiencies. Given the limited dataset size, no cross-validation or random splits were performed during this analysis. Although, the performance metrics consistently demonstrated a reliable model prediction across the dataset. The standard deviation was We chose ordinal encoding as a practical approach to represent the process and dye types while avoiding a large increase in feature dimensionality. One-hot encoding could potentially be more appropriate for certain models, but we believe the ordinal encoding sufficiently captured the relationships between the variables without overcomplicating the model. We will explore alternative encoding strategies, such as one-hot encoding, in future work to better understand its impact on the model performance and accuracy calculated from the residuals, defined as the difference between experimental and predicted removal efficiencies for all test samples.

Parity plots were also generated to compare experimental and predicted removal efficiencies. For the optimization results, the ANN-derived optimal conditions

Table 3: Optimal ANN architectures and predictive performance for individual dye–process models.

Dye	Process	Hidden layers	Neurons per layer	Activation function	Optimizer	Standard deviation of prediction error (%) (test set)
Green A	EF	1	11	Leaky ReLU	SGD	2.669
Green A	EO	2	14-9	Leaky ReLU	Adam	7.963
Green A	PEF	4	14-12-1-6	ELU	Adam	2.040
Purple RL	EF	1	6	ELU	SGD	3.832
Purple RL	EO	3	13-12-8	ELU	Adam	3.733
Purple RL	PEF	2	16-12	ELU	Adam	6.635

Table 4: Optimal architecture and performance metrics of universal ANN model.

Dye	Hidden layers	Neurons per layer	Activation function	Optimizer	Standard deviation of prediction error (%) (test set)
Scope Both Dyes	4	10-16-13-9	Leaky ReLU	SGD	3.977

Table 5: Optimal operating conditions predicted by the universal ANN models and corresponding experimental validation.

Dye	Process	Concentration (ppm)	Current density (mA cm^-2)	Time (min)	ANN-predicted removal (%)	Experimental removal (%)
Green A	PEF	30	50	60	100.0	95.0
Purple RL	PEF	30	50	50	100.0	100.0

were experimentally validated, and any discrepancy between predicted and observed removal efficiencies was attributed to model interpolation and extrapolation effects inherent to data-driven approaches.

RESULTS

Performance of individual ANN models.

Six individual ANN models were developed for each dye–process combination. These models were not used for operating-condition optimization and were used solely to evaluate the predictive capability of the ANN approach under fixed electrochemical conditions.

The optimal architectures of the individual ANN models, obtained via Bayesian optimization, are summarized in Table 3. The low standard deviations across all models indicate satisfactory predictive performance.

Parity plot comparing experimental and predicted removal efficiencies for Green A (PEF process) is presented in Figure 2, showing good agreement across the evaluated operating range.

Universal ANN model

A single universal ANN model integrating both dyes and all electrochemical processes was developed and optimized. The optimal architecture and performance metrics of the universal ANN model are reported in Table 4.

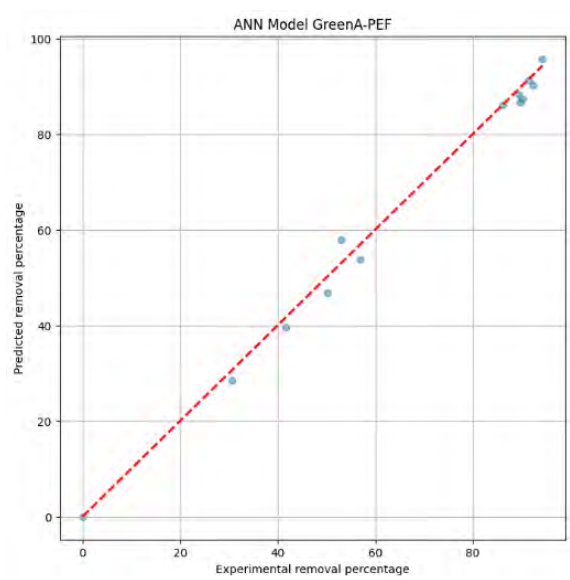

Figure 2. Parity plot for individual ANN models.

Despite the increased complexity of modeling multiple dyes and processes, the universal ANN demonstrated strong predictive capability, as evidenced by the parity plot shown in Figure 3.

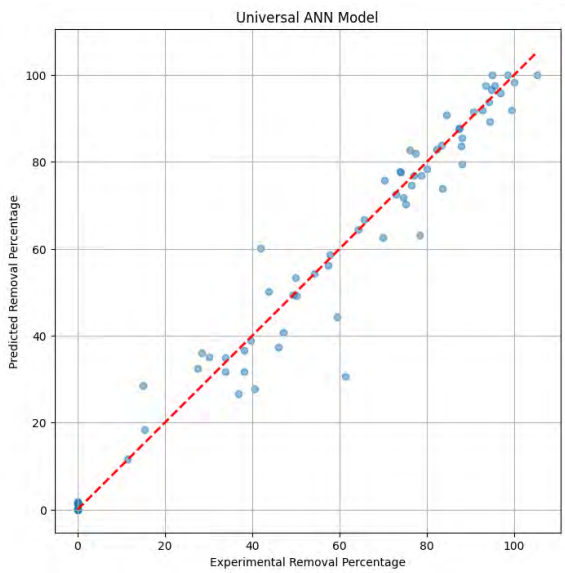

Figure 3. Parity plot for Universal ANN model.

Optimization of operating conditions

The optimized universal ANN model was employed to determine the operating conditions that maximize dye removal efficiency. The optimization was performed independently for each dye by fixing the dye identifier and maximizing the ANN-predicted removal efficiency with respect to the operating variables.

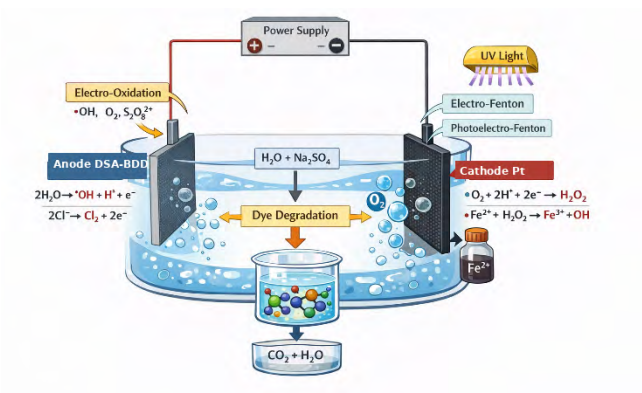

Figure 4: Electrochemical system for dye degradation showing EO, EF, and PEF pathways, with •OH generation at the DSA–BDD anode, H_2O_2 production and Fenton reactions at the Pt cathode, and UV-assisted Fe^{2+} regeneration leading to dye mineralization.

Table 5 summarizes the optimal operating conditions predicted by the ANN-based optimization framework, along with the corresponding predicted removal efficiencies. Experimental validation confirmed that the optimized conditions yielded high removal efficiencies, though slightly lower than the ANN-predicted theoretical maxima.

Overall, Figure 4 shows how the combination of EO, EF, and PEF processes sustains the continuous generation of highly reactive species in the system.

The formation of •OH radicals at the anode and through the Fenton reaction at the cathode, together with the photochemical regeneration of Fe^{2+} under UV irradiation, maintains a strongly oxidizing environment that accounts for the progressive degradation of the dyes. This behavior highlights the importance of the synergy between electrochemical and photochemical pathways in driving the transformation of contaminants into increasingly simpler products.

CONCLUSIONS

Artificial neural networks-based models have been developed to represent the performance of three electrochemical advanced oxidation techniques in the degradation of dyes. Both the individual and the universal models predict accurately the performance of the processes, with deviations lower than 10%. This indicates that these models can properly predict the complex performance of the electrochemical advanced oxidation techniques, opening further research areas for more detailed analysis.

ACKNOWLEDGEMENTS

Authors acknowledge the support of Universidad de Guanajuato and Universidad Autónoma de Querétaro for the development of this research.

AUTHOR IDENTIFIERS

Author ORCIDs:
Mosqueda-Huerta ZJ: Not Available
Lara-Montaño OD: 0000-0003-2055-1216
Peralta-Hernández JM: 0000-0003-2733-4126
Gómez-Castro FI: 0000-0003-4906-063X

REFERENCES

1. Huang Z, Yuan X, Liu X. The key drivers for the changes in global water scarcity: water withdrawal versus water availability. Journal of Hydrology 601:126658 (2021). https://doi.org/10.1016/j.jhydrol.2021.126658
2. Birnbaum A, Lamontagne J, Wild T, Dolan F, Yarlagadda B. Drivers of future physical water scarcity and its economic impacts in latin america

and the caribbean. Earth's Future 10: (2022). https://doi.org/10.1029/2022ef002764
3. Esmaeeli A, Sarrafzadeh MH, Zeighami S, Kalantar M, Bariki SG, Fallahi A, Asgharnejad H, Ghaffari SB. A comprehensive review on pulp and paper industries wastewater treatment advances. Ind. Eng. Chem. Res. 62:8119-8145 (2023). https://doi.org/10.1021/acs.iecr.2c04393
4. Dutta D, Arya S, Kumar S. Industrial wastewater treatment: current trends, bottlenecks, and best practices. Chemosphere 285:131245 (2021). https://doi.org/10.1016/j.chemosphere.2021.131245
5. Mao G, Han Y, Liu X, Crittenden J, Huang N, Ahmad UM. Technology status and trends of industrial wastewater treatment: a patent analysis. *Chemosphere* 288:132483 (2022) https://doi.org/10.1016/j.chemosphere.2021.132483
6. Adane T, Adugna AT, Alemayehu E. Textile industry effluent treatment techniques. Journal of Chemistry 2021:1-14 (2021). https://doi.org/10.1155/2021/5314404
7. Hossain MI, Soliman MM, El-Naggar ME, Sultan MZ, Kechi A, Abdelsalam NR, Abu-Saied MA, Chowdhury M. Synthesis and characterization of graphene oxide-ammonium ferric sulfate composite for the removal of dyes from tannery wastewater. Journal of Materials Research and Technology 12:1715-1727 (2021). https://doi.org/10.1016/j.jmrt.2021.03.097
8. Herrera-Chávez S, Pacheco-Álvarez M, Kadier A, Brillas E, Peralta-Hernández JM. Efficient electrochemical advanced degradation of red CL and red WB dyes from the tanning industry using a boron-doped diamond anode. Chemosphere 363:142825 (2024). https://doi.org/10.1016/j.chemosphere.2024.142825
9. Isaev AB, Shabanov NS, Magomedova AG, Nidheesh PV, Oturan MA. Electrochemical oxidation of azo dyes in water: a review. Environ Chem Lett 21:2863-2911 (2023). https://doi.org/10.1007/s10311-023-01610-5
10. Sobczak M, Bujnowicz S, Bili?ska L. Fenton and electro-fenton treatment for industrial textile wastewater recycling. comparison of by-products removal, biodegradability, toxicity, and re-dyeing. Water Resources and Industry 31:100256 (2024). https://doi.org/10.1016/j.wri.2024.100256
11. Behrouzeh M, Mehdi Parivazh M, Danesh E, Javad Dianat M, Abbasi M, Osfouri S, Rostami A, Sillanpää M, Dibaj M, Akrami M. Application of photo-fenton, electro-fenton, and photo-electro-fenton processes for the treatment of DMSO and DMAC wastewaters. Arabian Journal of Chemistry 15:104229 (2022). https://doi.org/10.1016/j.arabjc.2022.104229
12. Cañizares P, García-Gómez J, Lobato J, Rodrigo MA. Modeling of wastewater electro-oxidation processes part I. general description and application to inactive electrodes. Ind. Eng. Chem. Res. 43:1915-1922 (2004). https://doi.org/10.1021/ie0341294
13. Chen Z, Wang X, Feng H, Chen S, Niu J, Di G, Kujawski D, Crittenden JC. Electrochemical advanced oxidation of perfluorooctanoic acid: mechanisms and process optimization with kinetic modeling. Environ. Sci. Technol. 56:14409-14417 (2022). https://doi.org/10.1021/acs.est.2c02906
14. Schweidtmann AM, Esche E, Fischer A, Kloft M, Repke J, Sager S, Mitsos A. Machine learning in chemical engineering: a perspective. Chemie Ingenieur Technik 93:2029-2039 (2021). https://doi.org/10.1002/cite.202100083

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Model Screening and Identifiability Analysis of Methanol Synthesis Kinetics: Information-Guided Evaluation of Operating Conditions

Eblagh Ahmad[a], Biasin Alberto[b], Nardi Luca[b], Federico Galvanin[a]*

[a] University College London, Department of Chemical Engineering, London, United kingdom
[b] Casale SA, R&D division, Basic Research Department, Lugano, Switzerland
* Corresponding Author: f.galvanin@ucl.ac.uk

ABSTRACT

Reliable kinetic models are essential for the design, optimisation and operation of methanol synthesis reactors in power-to-X applications. However, parameter estimation is frequently performed without prior assessment of parametric identifiability or the information content of experimental conditions, often resulting in poorly constrained parameters and inefficient experimental campaigns. This study introduces a systematic, pre-calibration, information-driven framework for identifiability analysis and experimental design. The framework integrates local sensitivity analysis, structural and practical identifiability metrics and a Sequential Information-Driven Experimental Selection (SIDeS) strategy to guide experiment selection prior to parameter estimation. The methodology is applied to four literature kinetic models for methanol synthesis, spanning varying levels of mechanistic detail. A Sobol-sampled design space is first screened for feasibility and local information content, followed by cumulative Fisher information analysis. Results demonstrate that conventional screening campaigns provide limited gains in parameter precision, whereas SIDeS identify compact experimental sequences that significantly improve practical identifiability within a constrained experimental budget. Cross-model comparison reveals that increased mechanistic description does not necessarily translate to improved identifiability under fixed experimental resources. The proposed framework supports informed model selection, reduces unnecessary calibration effort, and provides actionable design-of-experiments recommendations for methanol synthesis systems.

Keywords: Methanol Synthesis, Kinetic Modelling, Identifiability Analysis, Fisher Information Matrix, Design of Experiments

1. INTRODUCTION

Methanol is a versatile chemical used as a fuel, hydrogen carrier, and precursor to a wide range of industrial products. Its synthesis from captured CO_2 and green hydrogen represents a key component of emerging power-to-X technologies aimed at reducing carbon emissions. Accurate kinetic models are central to the design, optimisation, and simulation of methanol synthesis reactors, yet predictive modelling remains challenging due to the complex reaction network and variability across catalysts and operating conditions [1].

The literature reports several kinetic expressions and mechanistic formulations developed under different assumptions, datasets, and catalyst systems [1, 2, 3, 4]. As a result, reported parameter values exhibit significant variability, reflecting differences in operating windows, data quality, and estimation methodologies. Historically, parameter estimation has relied on nonlinear regression of reactor-specific datasets, with limited assessment of model transferability. Moreover, calibration is typically performed without prior evaluation of structural or practical parameter identifiability, i.e. whether model parameters can be uniquely inferred from the available measurements under the given experimental condition, nor of the information content of the underlying experimental conditions [5]. This limits confidence in estimated parameters and hampers the design of informative

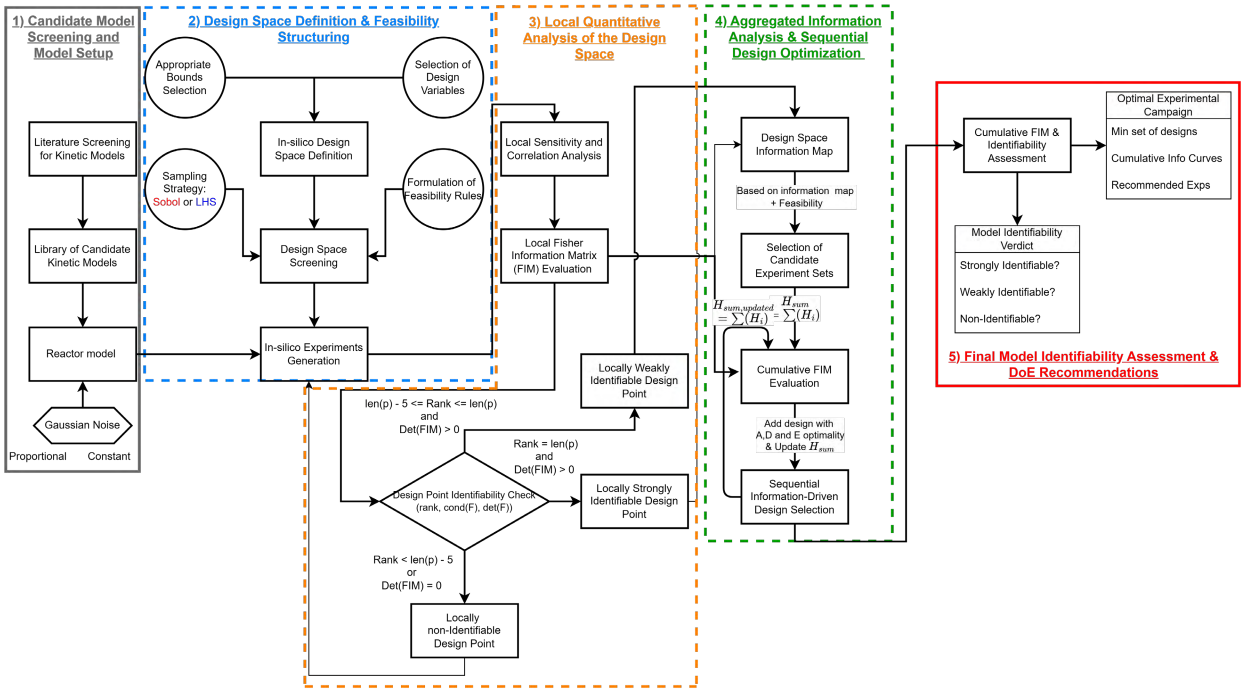

Figure 1. Overview of the proposed pre-calibration identifiability and sequential design framework. The workflow comprises: (1) model setup and candidate screening, (2) design space definition and feasibility constraints, (3) local sensitivity-based information analysis, (4) cumulative FIM-based SIDeS design selection, and (5) final identifiability assessment and DoE recommendations.

experiments.

Model-based design of experiment (MBDoE) techniques has been widely applied to improve parameter precision and guide experiment selection for kinetic models using Fisher-information-based criteria. Early contributions focused on optimal experiment design for parameter estimation in reaction systems, while more recent studies have explored identifiability analysis and information-driven experiment selection strategies for reactor models and catalytic systems. For example, De Luca et al. (2016) [8] investigated experiment redesign strategies for improving parameter estimation, while Deussen et al. (2022) [9] analysed identifiability issues in kinetic model calibration workflows. More recent work by Friso and Galvanin (2024) [10] further explored advanced MBDoE methodologies for reaction engineering applications. However, these approaches are typically applied during or after calibration, whereas the present work focuses on pre-calibration screening of kinetic models and operating conditions, aiming to identify informative experimental regions before parameter estimation is performed.

This study introduces a systematic, pre-calibration, information-driven framework integrating local sensitivity analysis, structural identifiability, practical identifiability under assumed noise, and a sequential information-driven experimental selection (SIDeS) strategy [6, 7]. The framework is applied to four literature kinetic models of methanol synthesis, ranging from the classical Graaf et al. model [1] and Vanden Bussche and Froment (VBF) ([2] to more recent models proposed by Slotboom [3] and Seidel [4]. The aim is to characterise the identifiability properties of each model, identify informative regions of the design space, and provide experiment design recommendations prior to calibration, thereby supporting model selection and reducing unnecessary calibration effort.

Section 2 describes the proposed methodology. Section 3 presents the identifiability and information analysis results and discusses the SIDeS-based design recommendations. Section 5 summarises the main conclusions.

2. METHODOLOGY

This section outlines the proposed pre-calibration, information-driven framework, beginning with the definition of candidate kinetic models and the reactor representation (Section 2.1), followed by information and identifiability analyses across the design space (Section 2.2 – 2.5), and finally the application of SIDeS to reach a final identifiability decision and experiment-selection recommendations. An overview of the workflow is provided in Figure 1, which organises the methodology into five consecutive stages, that are detailed in the following

sections.

2.1. Candidate Models and Reactor Representations

2.1.1 Reactor model

A one-dimensional pseudo-homogenous packed-bed reactor model was used to simulate species and temperature profiles under non-isothermal conditions. The model incorporates a non-isothermal energy balance with axial thermal dispersion. Radial dispersion, catalyst deactivation, and external or internal mass-transfer limitations were neglected as standard simplifying assumptions [11].

Non-ideal gas behaviour was modelled using the Soave-Redlich-Kwong (SRK) equation of state [11]. Component fugacity's were defined as the product of partial pressure and SRK-derived fugacity coefficients, while equilibrium constants were evaluated using the correlations of Graaf and Winkelman [12]. These thermodynamic relationships were applied consistently across all kinetic models to ensure a uniform physical basis for comparison and identifiability analysis.

2.1.2 Candidate kinetic models

Four literature kinetic models were examined: the classical dual-site LHHW formulation of Graaf et al. [1], the single-site Vanden Bussche–Froment (VBF) model which omits direct CO hydrogenation [2], the three-site 10-parameter refit of Slotboom et al. [3], and the CO_2-oriented mechanistic model of Seidel et al [4]. These models span different mechanistic assumptions and levels of parameterisation, ranging from detailed adsorption structures to reduced-order kinetic expressions, providing a representative basis for identifiability screening.

Kinetic and adsorption parameters appearing explicitly in the rate expressions were treated as uncertain parameters. Outlet molar flowrates of CH_3OH, CO_2, CO, H_2 and H_2O, together with the outlet temperature and pressure, were used as model observables for computing local sensitivities and Fisher information metrics. Identifiability analysis was performed solely on model-predicted responses without the need of experimental data (a priori identifiability). A concise comparison of the four kinetic models is provided in Table 1.

Table 1: Summary of the kinetic models considered in this study, including parameter counts and key mechanistic characteristics.

Model	No of Params	Type
Graaf et al.	12	Type 1, Dual-site LHHW
VBF	9	Type 2, Single-site
Slotboom	10	Type 2, 3-site mechanism
Seidal et al.	14	Type 1, 3-site mechanism

All four kinetic models were evaluated under identical reactor and thermodynamic assumptions to ensure that differences in identifiability arise solely from the kinetic formulations.

2.2. Design Space Definition and Feasibility Filter

A multidimensional design space was constructed by incorporating key operating variables relevant to methanol synthesis, including reactor inlet temperature, inlet pressure, gas hourly space velocity (GHSV), and inlet composition expressed via the stoichiometric number (SN). Additionally, bounds were selected to reflect industrially relevant operating windows [3, 10], and a low-discrepancy Sobol method [13] was used to generate uniformly distributed inlet candidate design points. Each candidate point corresponds to a complete inlet-condition vector, which is used to generate in-silico reactor outputs for subsequent sensitivity and information analysis.

Static feasibility constraints were imposed prior to simulation, including SN > 2 to ensure hydrogen-rich operation [3], adherence to sum-to-one constraints for inlet compositions and temperature and pressure limits consistent with industrial practice [3, 14]. Each feasible candidate was then evaluated using the full reactor model, and points resulting in solver failure, negative concentrations or thermodynamic inconsistencies under the SRK EOS were discarded. This two-stage feasibility filter ensured that the resulting design set was both physically meaningful and operationally realistic.

The final feasible design space forms the basis for local sensitivity analysis and the subsequent sequential design optimisation in SIDeS.

2.3. Local Sensitivity Analysis and Structural Identifiability

Local parameter sensitivities were computed to quantify response of model outputs to perturbations in each kinetic parameter. For a given design point and nominal parameter vector θ, the sensitivity matrix is defined by the following [15]:

$$S_{ij} = \frac{\partial y_i}{\partial \theta_j} \tag{1}$$

Where y_i denotes a simulated model response. In this study, the response vector $\boldsymbol{y}$ comprises steady-state outlet measurements, namely the outlet molar fractions of the species together with outlet temperature and pressure. Sensitivities were computed using a forward finite difference method, with the resulting matrix forming the basis of the information-based identifiability analysis.

Structural identifiability is assessed using the noise-free standard Fisher Information Matrix (FIM):

$$H_\theta = S^T S \tag{2}$$

The rank and conditioning of the FIM provides information on parameter identifiability. A full-rank FIM and non-singular (non-zero determinant) indicates all parameters are linearly independent and locally identifiable at the given point, whereas rank deficiency reveals structurally unidentifiable or redundant parameter combinations. Additionally, the trace of the FIM is used as a scalar metric of the overall information content, which in this case reflects the aggregate magnitude of the kinetic parameter sensitivities. Furthermore, eigen-analysis of the FIM was deployed to examine the null-space directions associated with weakly identifiable parameters. On this basis, operating points were classified as i) strongly identifiable when the FIM was full-rank and well-conditioned; ii) weakly identifiable when near-singular behaviour or pronounced null-space directions were present; and iii) non-identifiable when the determinant approached zero or the rank was deficient [6, 7].

2.4. Practical Identifiability

Practical identifiability quantifies parameter uncertainty in the presence of measurement noise. All practical uncertainty calculations were performed around nominal parameter values obtained from the literature for each kinetic model. Two noise structures were considered: a constant Gaussian noise model and a proportional Gaussian noise model, each defining the corresponding diagonal weighting matrix $\boldsymbol{\Sigma_y}$ used in the weighted FIM. An additive noise model was assumed for temperature ($\sigma_T = 1\,\mathrm{K}$), whereas composition measurements y_i were described by Gaussian noise with standard relative deviation of $\sigma_{y_i} = 2\%$. For a given design point, the weighted FIM is defined by [15, 9]:

$$H_\theta = S^T\, \Sigma_y\, S \tag{3}$$

Parameter uncertainties were obtained from the variance-covariance matrix $\boldsymbol{V}_\theta$, which is computed from FIM as follows:

$$V_\theta = H_\theta{}^{-1} \tag{4}$$

Where the diagonal entries $V_{\theta,ij}$ represent the variances of the kinetic parameters.

Practical identifiability was assessed using the relative standard error (RSE%):

$$\mathrm{RSE_j} = \frac{\sqrt{V_{\theta,jj}}}{\theta_j} \times 100\% \tag{5}$$

Each RSE has a corresponding *t*-statistic in the form:

$$\mathrm{t_j} = \frac{\theta_j}{\sqrt{V_{\theta,jj}}} \tag{6}$$

Low RSE values and large $|\, t_j \,|$ indicate practically identifiable parameters, whereas large uncertainties imply weak or non-identifiability. For each parameter, practical identifiability was further assessed by comparing the computed t-statistic to the corresponding critical value at a 95% and 99% confidence level. The effective degrees of freedom were taken as the number of in-silico experiments accumulated at that stage of the cumulative FIM evaluation [9, 10, 15].

2.5. Cumulative Information Metrics

When multiple experiments are considered, the total information available for parameter estimation is represented by the cumulative FIM, which is defined as:

$$H_{\theta,\mathrm{cum}} = \sum_{k=1}^{\mathrm{N_{exp}}} H_\theta^{(k)} \tag{7}$$

Where $\boldsymbol{H}_\theta^{(k)}$ denotes the weighted FIM for the kth design point. The cumulative variance-covariance is obtained from:

$$V_{\theta,\mathrm{cum}} = H_{\theta,\mathrm{cum}}^{-1} \tag{8}$$

This allows the evaluation of cumulative RSE% and t-statistics as additional experiments are incorporated. The trace of $\boldsymbol{H}_{\theta,\mathrm{cum}}$ is used as a scalar measure of overall information content to visualize information gain across experiment sequences. This cumulative information is the underpinning principle behind the SIDeS strategy described in section 2.6 below [13].

2.6. SIDeS Design Strategy

The Sequential Information-Driven Experimental Selection (SIDeS) procedure was used to construct an experiment sequence that maximises parameter identifiability while minimising the number of required experiments. SIDeS operate on the feasible design space defined in Section 2.2, in which each candidate point has an associated local FIM and information metric.

The algorithm begins with no prior experiments. The first experiment is selected as the design point with the highest local information content, based on the chosen scalar objective. In this work, the scalar objective was defined as the reduction in the worst-case parameter RSE, favouring experiments that most improve the least identifiable parameter.

The cumulative FIM is then initialised using this first experiment, and subsequent experiments are added in a "greedy" manner. At each iteration, every remaining feasible design point is evaluated in terms of its marginal information contribution, defined as the improvement in the selected identifiability objective when its FIM contribution is added to the current cumulative FIM, given by Eq. (9). In this study, identifiability objectives are evaluated through practical identifiability metrics derived from the inverse cumulative FIM, namely the parameter %RSE (Eq. 5) and their associated t-statistics (Eq. 6). This

criterion accounts for both the magnitude and direction of the parameter sensitivities, ensuring that additional experiments provide complementary, non-redundant information rather than reinforcing already saturated directions.

After selecting the kth design, the cumulative FIM is updated as follows:

$$H_{\theta,\mathrm{cum}}{}^{(n+1)} = H_{\theta,\mathrm{cum}}{}^{(n)} + H_{\theta}{}^{(k)} \tag{9}$$

Here, n denotes the iteration index of the sequential design procedure, such that $\boldsymbol{H}_{\theta,\mathrm{cum}}^{(n)}$ represents the cumulative FIM after n experiments have been selected.

The procedure terminates when either the maximum number of design steps (the experiment budget) is reached or when the user-defined identifiability criterion is satisfied, such as achieving a specified structural rank or meeting practical identifiability thresholds for the parameters of interest.

All simulations were implemented in Python using the SciPy scientific computing ecosystem and executed on a workstation equipped with an Intel i7 processor and 32 GB RAM.

3. RESULTS AND DISCUSSION

3.1 Case Study and Design Space

The case study focuses on the Graaf et al. kinetic model, which remains one of the most widely adopted reference formulations for methanol synthesis in the literature, integrated in a non-isothermal packed-bed reactor model with axial dispersion.

For the purposes of identifiability analysis, the uncertain kinetic and adsorption parameters of the Graaf model are collected into a parameter vector expressed as

$$\theta_k = [k_{0,CO}, Ea_{CO}, k_{0,CO_2}, Ea_{CO_2}, k_{0,RWGS}, Ea_{RWGS},$$

$$\Delta S^o_{ads,CO}, \Delta H^o_{ads,CO}, \Delta S^o_{ads,CO_2}, \Delta H^o_{ads,CO_2}, A_{\frac{K_{H_2O}}{K_{H_2}^{0.5}}}, B_{\frac{K_{H_2O}}{K_{H_2}^{0.5}}}]$$

Where the elements correspond to kinetic and adsorption-related parameters appearing in the Graaf rate expressions. These parameters constitute the target of the subsequent identifiability and experimental design analyses.

A Sobol low-discrepancy sampling scheme was employed to generate 50 candidate design vectors, each defined by inlet composition, inlet temperature, inlet pressure, and gas hourly space velocity (GHSV). These vectors constitute a discrete set of in-silico candidate experiments used to probe the design space prior to any experimental execution.

The local information content of each candidate design is quantified using the trace of the FIM, which is widely employed as a scalar measure of parameter sensitivity magnitude [14]. While the trace provides a convenient aggregate metric, it does not account for directional information or parameter correlations. Directional diagnostics, such as eigenvalue-based criteria, are therefore considered in later stages of the framework.

Figure 2 illustrates the candidate design space (three of eight design variables shown for visual clarity), with each point coloured by its local FIM trace value. The Sobol sampling strategy yields a well-distributed coverage of the admissible design space. Despite this uniform geometric coverage, the associated information content exhibits substantial variability, with only a limited subset of designs providing comparatively high local information.

Executing all candidate experiments without regard to their information content would be highly cost-inefficient and representative of unstructured screening experimental practice. This observation motivates the need for information-driven experiment selection strategies that explicitly account for both model structure and parameter sensitivities. The application of the Sequential Information-Driven Experimental Selection (SIDeS) framework is therefore investigated in the following sections.

3.2 Identifiability Analysis from screening experiments

This section presents the identifiability analysis of the Graaf et al. kinetic model from screening experiments, conducted prior to the application of the SIDeS framework. Identifiability is assessed assuming all candidate experiments are performed without optimisation or reordering, following the original Sobol-generated order. Herein, a preliminary identifiability analysis on Sobol generated screening experiments is presented to provide a diagnostic reference against which the benefits of the information-driven experiment selection can be assessed later.

As shown in Figure 2, local information content and single-experiment identifiability vary substantially across the candidate design space. While certain operating conditions exhibit comparatively high Fisher information trace values, most individual experiments remain locally weak or non-identifiable when assessed in isolation. These single-experiment identifiability results are presented for diagnostic purposes only, but, in practice, reliable parameter identification requires the aggregation of information across multiple experiments.

Having established the limitations of single-experiment identifiability, the cumulative effect of aggregating experimental information is now examined. Cumulative identifiability is assessed using practical metrics that quantify parameter precision under the assumed noise model. In the following, relative standard errors (%RSE) are used to evaluate how parameter uncertainty evolves

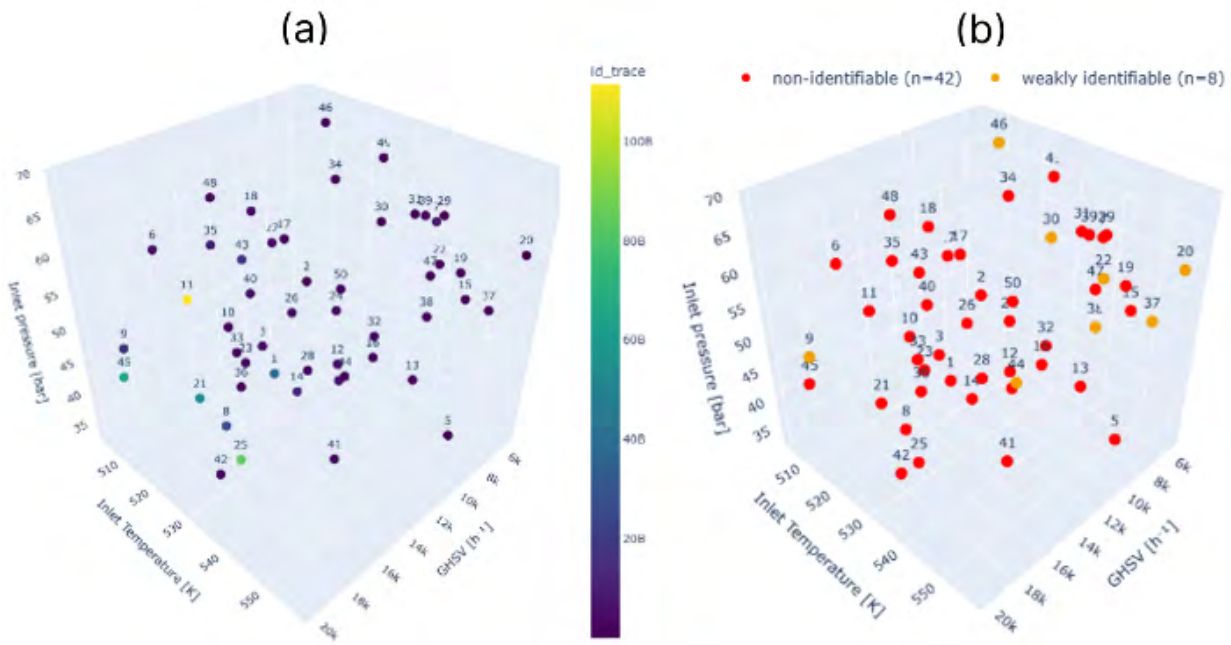

Figure 2. Candidate design space generated using Sobol sampling. **a):** Points coloured by the local Fisher information trace, indicating the magnitude of parameter sensitivities ats each individual design. **b):** The same design space coloured by local identifiability verdicts based on single-experiment Fisher information analysis. Each point corresponds to a feasible in-silico experiment and is labelled by its design identifier.

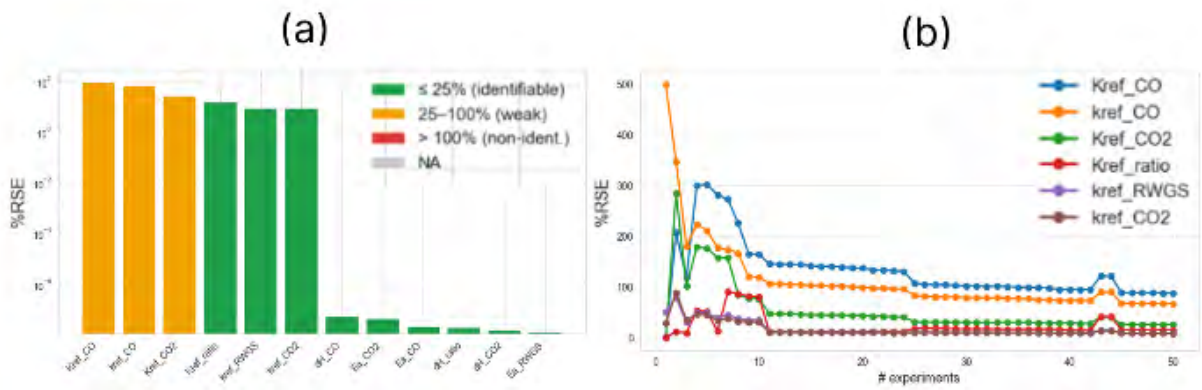

Figure 3. a) Evolution of cumulative parameter relative standard error (%RSE) as a function of number of experiments b) Final cumulative % RSE for each kinetic parameter after all candidate experiments are added.

as experiments are progressively added without optimal design, providing insight into the effectiveness of screening experiments. The cumulative evolution of parameter uncertainty across screening experiments is illustrated in Figure 3.

Figure 3.a illustrates the evolution of cumulative parameter %RSE as a function of the number of experiments. While uncertainty decreases as additional experiments are incorporated, convergence is slow, and highly non-uniform across the parameters. Importantly, most of the uncertainty reduction is achieved within the first approximately 10 screening experiments, beyond which further additions yield diminishing returns.

The final cumulative identifiability status after the full set of screening experiments is summarised in figure 3.b. Although some parameters achieve acceptable precision, others remain weakly identifiable according to the chosen %RSE thresholds. This result highlights that increasing the number of experiments alone is insufficient in guaranteeing practical identifiability, and that the distribution of experimental information across parameters is highly uneven.

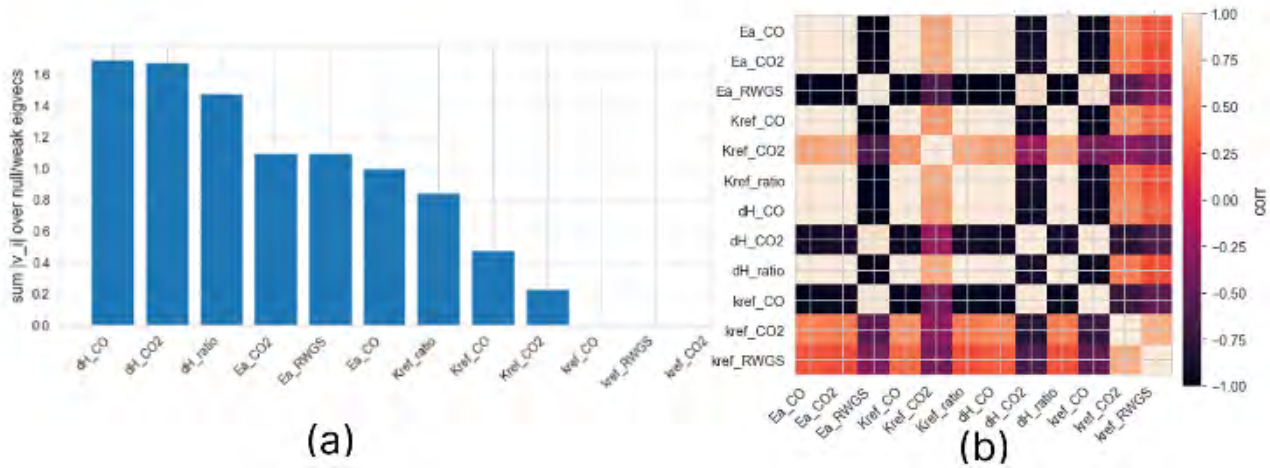

Figure 4. a) cumulative null-space loadings, indicating contribution of each parameter to weakly excited eigen-directions of FIM. b) Cumulative parameter correlation matrix after 50 experiments

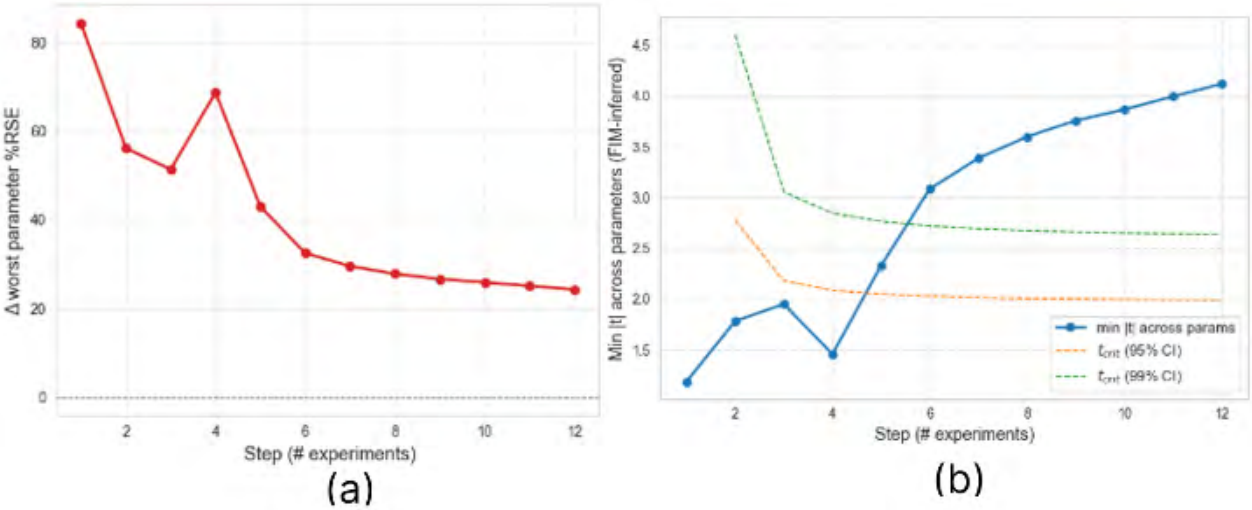

Figure 5. a) Evolution of the worst parameter %SE as a function of SIDeS iteration b) Convergence of minimum absolute t-statistic across parameters, with critical values of 95% and 99% shown for reference.

To further interpret the observed limited increment in parameter precision, null-space loadings and cumulative parameter correlations of the FIM are examined in Figure 4.a and 4.b, respectively. Strong and persistent correlations are observed between several kinetic parameters, indicating that specific parameter combinations remain poorly distinguishable even as additional experiments are added. Similarly, null-space analysis reveals dominant weakly excited directions associated with the corresponding parameters. This confirms that unstructured experimental screening fails to systematically resolve structural collinearity and weakly informed parameter subspaces.

Together, these results demonstrate that uniform coverage of feasible design space does not translate into uniform information gain across parameters. Screening Experiments without SIDeS leads to inefficient use of experimental effort, motivating the need for an information-driven experiment selection strategy capable of explicitly targeting weakly excited parameter directions. This motivation is at the cornerstone of the application of the SIDeS framework, which is examined next.

3.3 SIDeS Results

The SIDeS framework was applied to the same candidate design used in Section 3.2 to ensure a fair comparison against screening experiments. A maximum experimental budget of 12 experiments was imposed. The experimental selection objective was defined as the minimisation of the worst-case parameter %RSE at each iteration. Experiments were then selected sequentially using a greedy strategy without repetition, with FIM cumulatively updated after each addition. The SIDeS algorithm executed successfully, reaching full identifiability for all

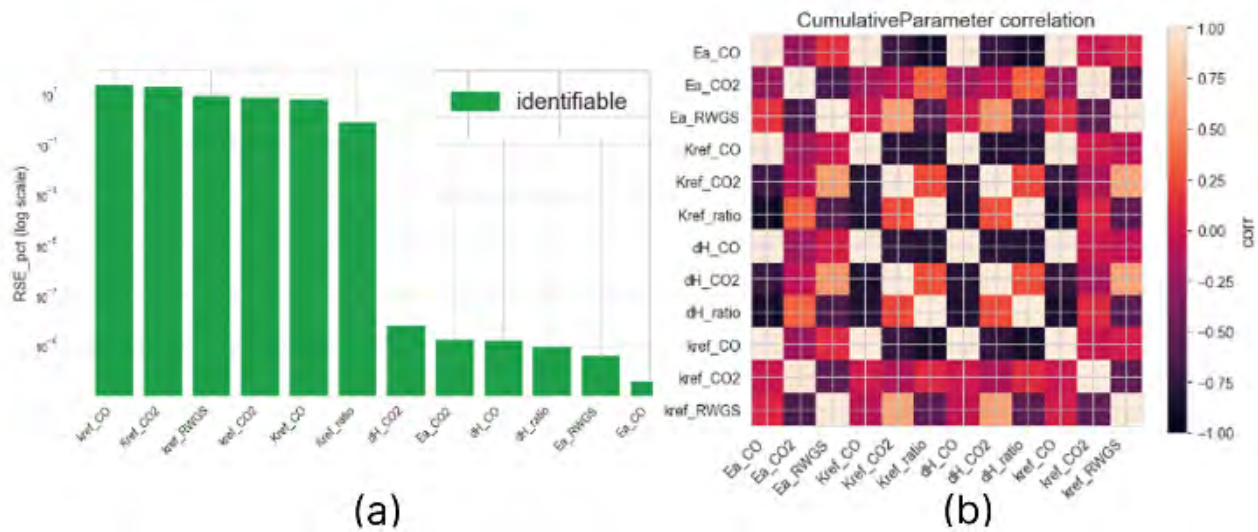

Figure 6. a) Final parameter %RSE after SIDeS campaign. b) Cumulative parameter correlation matrix after SIDeS campaign.

Table 2: Cross-model identifiability and SIDeS outcomes.

Kinetic Model	# Params	# Active Sites	Baseline Identifiability	SIDeS steps to full Iden	Worst %RSE (Final)	Dominant Design Bias	Residual Correlation
Graaf	12	2	Weakly Iden	12	24%	High Temp,	Moderate
VBF	9	1	Weekly Iden	8	18%	high GSV	Low
Slotboom	10	3	Weekly Iden	>12	35%	High P, mid T	Moder-
Seidel	14	3	Non Iden	>12	>100%	Broad pattern	ate/high
						No clear bias	Very high

models after reaching the pre-selected budget of 12. The chosen sequence of experiments based on candidate design ID was [41, 16, 36, 11, 45, 39, 25, 37, 50, 8, 7, 1].

Figure 5.a shows the evolution of the worst-case parameter %RSE as experiments are sequentially selected by the SIDeS. A rapid reduction appears in uncertainty during the initial iterations (up to step 6). This is followed by diminishing improvements as the experimental budget increases. This behaviour reflects an initial exploration phase, in which dominant information directions are rapidly resolved, followed by refinement phase where additional experiments yield progressively smaller marginal gains. This is complimented by the evolution of minimum absolute t-statistic across parameters shown in Figure 5.b. The t-statistics are computed from FIM-inferred standard errors under the assumed Gaussian noise model and provide an additional screening metric for practical identifiability. As experiments are added, the minimum $|t|$ increases and exceeds the critical values associated with common confidence levels of 95% and 99%. This indicates that all parameters become statistically distinguishable from zero within the local linear approximation.

Figure 6.b presents the cumulative parameter correlation matrix after SIDeS. Compared to correlation structure obtained from screening experiments, the strong parameter couplings are significantly reduced, particularly among adsorption and kinetic pre-exponential terms. This structural improvement confirms that SIDeS actively select experiments that resolve poorly identifiable parameter directions and does not just reduce overall uncertainty. This leads to a more informative and balanced parameter estimation problem.

The final parameter uncertainties obtained from the SIDeS-selected experimental campaign are summarised in Figure 6.a. In contrast to the experimental screening of all 50 candidate experiments presented in Section 3.2, SIDeS achieve practical identifiability for all kinetic parameters using only 12 experiments.

These result highlights the inefficiency of uninformed experimental accumulation and demonstrates the effectiveness of information-driven design in extracting parameter-relevant information, leading to substantially reduced experimental effort.

Furthermore, Examination of the SIDeS-selected

experimental sequence reveals clear qualitative trends in the regions of the design space identified as most informative for the Graaf et al. model. Within the admissible operating window, the algorithm exhibits a pronounced tendency to prioritise experiments at elevated inlet temperatures, often in combination with moderate-to-high pressures and higher space velocities. These conditions enhance sensitivity contrasts in Arrhenius and adsorption terms, thereby improving distinguishability among kinetic parameters. Importantly, the selected experiments are not confined to a single local optimum; instead, SIDeS distribute experiments across complementary regions of the design space, progressively exciting parameter directions that remain weakly informed at earlier stages. This behaviour reflects the algorithm's emphasis on directional information gain rather than purely scalar sensitivity magnitude. Importantly, these trends are specific to the Graaf model and the explored design space and are reported as qualitative insights rather than prescriptive operating guidelines.

3.4 Cross-model Comparison

This section compares the identifiability characteristics and information requirements of the four candidate kinetic models considered, under identical reactor, thermodynamics and design-space assumptions.

Table 2 summarises the cross-model identifiability outcomes obtained. Although all kinetic models describe the same methanol synthesis systems, substantial differences are observed in baseline identifiability, achievable parameter precision, and experimental information requirements. These results emphasize the strong dependence of practical identifiability on not only the experimental design bounds, but also the kinetic model structure. As such, the comparison thus provides a systematic basis for assessing the trade-off between mechanistic detail and experimental feasibility prior to model calibration.

Looking at Table 2, a clear inverse relationship is observed between model complexity and identifiability. The Graaf and VBF models, which contain fewer parameters and simpler active-site assumptions, exhibit weak baseline identifiability, but can be rendered fully practically identifiable using the SIDeS framework within modest experimental budgets. In contrast, the Slotboom and Seidel models, which introduce additional parameters and multi-site reaction mechanisms, display significantly poorer identifiability characteristics.

The Seidel model, comprising the largest parameter set, remains non-identifiable even after 12 optimally selected experiments, with worst-case %RSE exceeding 100%. This observation is accompanied by persistently high residual parameter correlations, indicating fundamental redundancy within the kinetic formulation. The variation in the budget required to achieve identifiability for the four kinetic models, including cases where identifiability was not attained, indicate that even information-driven experiment selection cannot fully compensate for intrinsic structural complexity. Within the investigated design space and experimental budget, the results suggest that increasing mechanistic detail does not necessarily yield improved practical identifiability and may be associated with higher information requirements.

Furthermore, analysis of the selected experiments reveals distinct model-dependent biases in the informative regions of the design space. The Graaf model consistently favours high temperature and high GHSV conditions, whereas VBF model priorities higher pressure combined with moderate inlet temperatures. In contrast, the Slotboom and Seidel models do not exhibit a favouritism to any region of the design space. Particularly for Seidel, where the absence of design bias suggests no single operating regime sufficiently resolves its parameter interdependencies.

From a practical perspective, these results demonstrate that model choice should be guided not only by mechanistic completeness but also by identifiability and experimental tractability. Models such as Graaf et al. and VBF offer a favourable balance between descriptive capability and calibrability, whereas more complex formulations may require significantly large or specialised experimental campaigns to achieve reliable parameter estimates. The proposed framework therefore enables informed kinetic model selection prior to experimental execution. Therefore, the risk of under-determined calibration and wasted experimental effort is reduced.

CONCLUSION

This work presented an information-driven framework for pre-calibration identifiability assessment and experimental screening in methanol synthesis kinetics. By combining local sensitivity analysis, Fisher information–based metrics and the Sequential Information-Driven Experimental Selection (SIDeS) strategy, the framework enables systematic screening of kinetic models and operating regions prior to parameter estimation. Application to four literature kinetic models showed that unstructured screening experiments yield diminishing gains in parameter precision, whereas SIDeS identify compact experimental campaigns that substantially improve practical identifiability within a fixed budget. Cross-model comparison further indicated that increased mechanistic description does not necessarily lead to improved identifiability under constrained experimental resources. The proposed framework is intended as a precursor to full model-based design of experiments (MBDoE), providing an informed starting point for optimal experimental campaign design in power-to-methanol applications.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the funding support from UCL and Casale SA.

REFERENCES

1. Graaf GH, Stamhuis EJ, Beenackers AACM. Kinetics of low-pressure methanol synthesis. Chemical Engineering Science 43:3185-3195 (1988). https://doi.org/10.1016/0009-2509(88)85127-3
2. Bussche KMV, Froment GF. A steady-state kinetic model for methanol synthesis and the water gas shift reaction on a commercial cu/zno/al2o3catalyst. Journal of Catalysis 161:1-10 (1996). https://doi.org/10.1006/jcat.1996.0156
3. Slotboom Y, Bos MJ, Pieper J, Vrieswijk V, Likozar B, Kersten SRA, Brilman DWF. Critical assessment of steady-state kinetic models for the synthesis of methanol over an industrial cu/zno/al2o3 catalyst. Chemical Engineering Journal 389:124181 (2020). https://doi.org/10.1016/j.cej.2020.124181
4. Seidel C, Jörke A, Vollbrecht B, Seidel-Morgenstern A, Kienle A. Kinetic modeling of methanol synthesis from renewable resources. Chemical Engineering Science 175:130-138 (2018). https://doi.org/10.1016/j.ces.2017.09.043
5. Walter, E.; Luc Pronzato. IDENTIFICATION OF PARAMETRIC MODELS; Springer, 1997.
6. Asprey SP, Macchietto S. Designing robust optimal dynamic experiments. Journal of Process Control 12:545-556 (2002). https://doi.org/10.1016/s0959-1524(01)00020-8
7. Asprey SP, Macchietto S. Statistical tools for optimal dynamic model building. Computers & Chemical Engineering 24:1261-1267 (2000). https://doi.org/10.1016/s0098-1354(00)00328-8
8. De-Luca R, Galvanin F, Bezzo F. A methodology for direct exploitation of available information in the online model-based redesign of experiments. Computers & Chemical Engineering 91:195-205 (2016). https://doi.org/10.1016/j.compchemeng.2016.03.016
9. Deussen P, Galvanin F. A model-based experimental design approach to assess the identifiability of kinetic models of hydroxymethylfurfural hydrogenation in batch reaction systems. Chemical Engineering Research and Design 178:609-622 (2022). https://doi.org/10.1016/j.cherd.2021.12.028
10. Friso A, Galvanin F. An optimization-free fisher information driven approach for online design of experiments. Computers & Chemical Engineering 187:108724 (2024). https://doi.org/10.1016/j.compchemeng.2024.108724
11. H Scott Fogler. *Elements of Chemical Reaction Engineering.*; 5th ed.; Prentice Hall, 2016.
12. Graaf GH, Winkelman JGM. Chemical equilibria in methanol synthesis including the water–gas shift reaction: a critical reassessment. Ind. Eng. Chem. Res. 55:5854-5864 (2016). https://doi.org/10.1021/acs.iecr.6b00815
13. Sobol' IM. On the distribution of points in a cube and the approximate evaluation of integrals. USSR Computational Mathematics and Mathematical Physics 7:86-112 (1967). https://doi.org/10.1016/0041-5553(67)90144-9
14. Bisotti F, Fedeli M, Prifti K, Galeazzi A, Dell'Angelo A, Manenti F. Impact of kinetic models on methanol synthesis reactor predictions: in silico assessment and comparison with industrial data. Ind. Eng. Chem. Res. 61:2206-2226 (2022). https://doi.org/10.1021/acs.iecr.1c04476
15. Franceschini, G., Macchietto, S., Model-based design of experiments for parameter precision: State of the art, Chemical Engineering Science, 63 (2008) 4846–4872. https://doi.org/10.1016/j.ces.2008.07.034

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Process Modeling Approach for Water and Energy Optimization in Geologic Hydrogen Extraction

Caroline Kaitano[a], and Thokozani Majozi[a]*

[a] University of the Witwatersrand, School of Chemical and Metallurgical Engineering, Braamfontein, Gauteng, South Africa
* Corresponding Author: thokozani.majozi@wits.ac.za.

ABSTRACT

Geologic hydrogen has emerged as a promising low-carbon energy vector, but its sustainable recovery requires effective stimulation and production strategies. This study presents an integrated process-modeling framework for evaluating hydrogen extraction from hydraulically stimulated reservoirs. The framework combines fracture propagation, damage evolution, Darcy-scale multiphase flow, permeability–aperture dynamics, and a dual-porosity dual-permeability (DPDP) representation to simulate hydrogen production in fractured source rock systems. In addition to production dynamics, the framework tracks operational water and energy inputs and incorporates a simplified reaction-extent formulation to represent hydrogen generation under data-limited conditions. The model was benchmarked against published shale gas production results to evaluate its ability to reproduce fracture-controlled production behavior and was subsequently applied to a representative multi-well hydrogen development scenario. Simulation results indicate that hydrogen production is strongly fracture-dominated, with fracture pathways accounting for more than 98% of total modeled output. Under the simulated conditions, the system produced approximately 1.13 million kg of hydrogen, with an estimated operational energy intensity of 3.4 MJ/kg H_2 and water consumption of 0.02 m^3/kg H_2. These values represent model-based estimates within the defined operational system boundary and illustrate the potential performance indicators that the framework can generate. Per-well analysis reveals variability in water and energy efficiency across the multi-well system, highlighting the importance of fracture design and reservoir characterization. The simulations also indicate that matrix contributions to total recovery are limited under the assumed conditions, suggesting potential benefits from strategies that enhance matrix reactivity or transport processes. Overall, the framework provides an integrated approach for exploring hydrogen production dynamics together with associated resource requirements. As field and experimental data for geologic hydrogen systems become available, the framework can be further refined through improved parameterization, sensitivity analysis, and calibration to support future technical, environmental, and economic assessments.

Keywords: Geologic hydrogen, Extraction, Modelling and Simulations, Optimization, Energy, Water

INTRODUCTION

The intensifying global energy crisis, propelled by geopolitical instability and the objectives of the 2015 Paris Agreement to restrict warming to below 2 °C and achieve net-zero emissions by 2050, highlights the pressing necessity for a swift transition to reliable, sustainable energy sources [1]. A significant issue in this shift is the identification of carbon-free, economically feasible, and sustainable energy sources [2].

Hydrogen has emerged as a cornerstone of the global energy transition, offering a versatile, low-carbon energy vector capable of decarbonizing hard-to-abate sectors [3, 4]. While manufactured hydrogen from electrolysis or reforming dominates current discourse, these pathways remain constrained by high costs, scalability challenges, and significant energy or carbon footprints [5]. In contrast, geologic hydrogen naturally generated in the subsurface through processes such as serpentinization, radiolysis, and iron oxidation has recently attracted

attention as a potentially abundant and low-emission alternative [6]. The accidental discovery and subsequent exploitation of native hydrogen in Bourakébougou, Mali, demonstrated the feasibility of harnessing subsurface hydrogen resources, catalyzing global interest in both natural accumulations and engineered stimulation approaches [6–8].

Recent research has introduced the concept of stimulated geologic hydrogen production, sometimes termed "orange hydrogen, " which seeks to actively engineer subsurface conditions to accelerate hydrogen-generating reactions in iron-rich ultramafic rocks [9, 10]. Techniques such as hydraulic fracturing, thermal treatment, chemical stimulation, and electrical reservoir stimulation have been proposed to enhance permeability, reactive surface area, and geochemical kinetics [11]. Numerical modeling studies indicate that stimulation can increase hydrogen yields by several orders of magnitude compared to unstimulated systems, with fracture networks and damage zones playing a decisive role in enabling efficient hydrogen migration and recovery [11, 12]. However, despite promising laboratory and pilot-scale results, critical uncertainties remain regarding the efficiency, sustainability, and resource requirements of such systems [13].

A key challenge lies in balancing water and energy inputs with hydrogen outputs. Stimulated systems require a continuous water supply to sustain serpentinization, while compression and fracturing impose significant energy costs [14]. Recent circularity studies suggest that integrating waste heat recovery and water recycling can reduce operational burdens, yet quantitative assessments of water–energy–hydrogen balances at the reservoir scale remain scarce. Moreover, while fracture-dominated pathways have been shown to account for the majority of production [15], the contribution of the rock matrix is often minimal, and this raises questions about how stimulation design and reservoir characterization can be optimized.

Against this backdrop, this study introduces an integrated process-modeling framework for evaluating geologic hydrogen recovery in hydraulically stimulated reservoirs. The framework combines fracture propagation, damage evolution, Darcy-scale multiphase flow, permeability–aperture dynamics, and a dual-porosity dual-permeability (DPDP) representation to simulate hydrogen production from fractured source rock systems.

Existing fractured-reservoir models primarily focus on fluid flow and production forecasting, while reactive transport models emphasize geochemical reactions and species transport in porous media. These approaches are typically applied independently and are not specifically designed to evaluate the coupled production dynamics and resource requirements associated with geologic hydrogen extraction.

The framework developed in this study integrates fracture-controlled flow behavior with a simplified representation of hydrogen-generating reactions and explicit accounting of water and energy inputs. This system-level integration enables quantitative evaluation of water–energy–hydrogen trade-offs in hydraulically stimulated reservoirs. It provides a basis for exploring operational constraints and efficiency considerations in geologic hydrogen development.

By linking subsurface production dynamics with resource-use metrics, the framework provides a tool for estimating production performance, evaluating resource consumption requirements, and informing future optimization of stimulation strategies. In the absence of field-scale geologic hydrogen datasets, the framework is intended as a first-order modeling approach that can be refined as new experimental and field observations become available.

The objectives of this study are therefore to (i) develop the integrated modeling framework, (ii) demonstrate its application to hydraulically fractured hydrogen systems, and (iii) evaluate production performance together with associated water and energy requirements.

FRAMEWORK DEVELOPMENT

This study develops an integrated process modelling framework to quantify hydrogen generation, transport, and recovery in hydraulically stimulated geologic hydrogen systems. The approach couples (i) fracture propagation and damage evolution, (ii) dual-porosity dual-permeability (DPDP) fluid flow, (iii) multiphase reactive transport, (iv) permeability-aperture evolution, and (v) thermal feedback from serpentinization reactions.

Governing Equations

In order to simulate the generation and transport of geologic hydrogen in the source rock, mass and energy conservation were solved.

$$\rho C_p \frac{\partial T}{\partial t} + \nabla \bullet \left(\rho_f C_{p,f} \vec{v} T - k \nabla T\right) = q_{rad} + \Delta H_r R_H \qquad (1)$$

Where ρ is the rock density, C_p is the rock's specific heat capacity, ρ_f is the fluid density, $C_{p,f}$ is the fluid specific heat capacity, $\vec{v}$ is the fluid velocity (m/s), T is the temperature, k is the thermal conductivity, q_{rad} is the radiogenic heat generation, ΔH_r is the heat from serpentinization (J/mol) and R_H is the generation rate from serpentinization (mol/m^3.s).

$$\frac{\partial C_H}{\partial t} + \nabla \bullet \left(-D_H \nabla C_H + \vec{v} C_H\right) = R_H - R_{microbial} - R_{leak} \qquad (2)$$

Where D_H is hydrogen diffusivity, C_H is hydrogen concentration (mol/m^3), $R_{microbial}$ is the microbial

hydrogen consumption, and R_{leak} is the loss due to leakage through fractures or faults.

$$R_H = \frac{\varepsilon_{max}}{b} \bullet \frac{exp\,(\frac{-t-t_{1/2}}{b})}{(1+exp\,(\frac{-t-t_{1/2}}{b}))^2} \quad (3)$$

$$\varepsilon(t) = \frac{\varepsilon_{max}}{1+exp(\frac{-t-t_{1/2}}{b})} \quad (4)$$

Where ε_{max} is the maximum reaction extent(%), t is the time in days, $t_{1/2}$ is the time to reach half-reaction extent (days), b is the rate constant in days, and $\varepsilon(t)$ is the reaction extent.

The fracture propagation with damage evolution and the aperture evolution are accounted for using equations (5)-(7).

$$\frac{dD}{dt} = A_d \bullet (\sigma - \sigma_{crit})_+^n \quad (5)$$

Where D is the damage variable, A_d is the rate constant for damage growth, σ is the applied stress, and σ_{crit} is the threshold for damage initiation.

$$k(D) = k_o \bullet (1 + \alpha_k D) \quad (6)$$

$$a(D) = a_o \bullet (\beta D) \quad (7)$$

Where a(D) is the evolving fracture aperture and α_k and β are sensitivity coefficients.

Mass transfer between the fracture and the fracture and matrix is accounted for using the dual permeability model.

$$Q_{f\to m} = \alpha \bullet (P_f - P_m) \quad (8)$$

Where $Q_{f\to m}$ is the flowrate between fracture and matrix (m^3/s), α is the transfer coefficient (s^{-1}), and P_f, P_m is the pressures in fractures and matrix.

The temperature feedback and fluid temperature change in the fractures is accounted for using equations (9) and (10), respectively.

$$Q_{rxn} = \Delta H_r R_H \to T^{n+1} = T^n + \frac{\Delta t}{\rho_r C p_r} Q_{rxn} \quad (9)$$

$$T_{out} = T_g - (T_g - T_{in}) \bullet exp(\frac{-h_c A}{\dot{m} C_{p,f}}) \quad (10)$$

Where T_{out} is the outlet fluid temperature, T_g is the rock matrix temperature, $\dot{m}$ is the fluid mass flowrate (kg/s), A is the contact area (m^2), and h_c is the convective heat transfer coefficient.

Numerical Implementation

The governing equations describing fluid flow, hydrogen transport, and reaction extent were solved numerically using a finite-difference discretization on a structured grid. The spatial domain representing the reservoir and fracture system was discretized into grid blocks that capture both matrix and fracture regions within the dual-porosity dual-permeability framework. Grid resolution was selected to resolve fracture–matrix interactions while maintaining computational efficiency adequately.

Time integration was performed using an implicit time-stepping scheme to ensure numerical stability for the coupled flow, transport, and reaction processes over the simulation period. The time step size was adjusted to capture the early-time production dynamics and longer-term stabilization behavior observed in the system.

All simulations were conducted over a 365-day production period using the parameter set described in Table 1. The computational implementation was developed in a numerical simulation environment capable of solving coupled nonlinear equations representing subsurface flow and reaction processes. Simulation runtimes for the scenarios presented in this study were on the order of minutes on a standard workstation, indicating that the framework is computationally efficient for scenario analysis and preliminary system evaluation.

Table 1: Parameters used in this study

Parameter	Value
Number of wells	5
Initial reservoir pressure	50 MPa
Reservoir thickness	50 m
Rock density	3300 kg/m³
Rock heat capacity	870 J/kg·K
Rock thermal conductivity	3.45 W/m·K
Matrix porosity	0.02
Matrix permeability	5×10^{-12} mD
Hydrogen viscosity	8.90×10^{-6} Pa·s
Enthalpy of reaction	–520 J/mol
Matrix diffusivity	5×10^{-8} m²/s
Fracture diffusivity	1×10^{-5} m²/s
Reaction extent	0.30
Half-life	60 days
Hydrogen heat capacity	14, 300 J/kg·K
Hydrogen density	80 kg/m³

FRAMEWORK VALIDATION

To evaluate the proposed framework, published shale gas production results were used as a reference dataset for benchmarking. Since direct field-scale data for geologic hydrogen are not yet available, shale gas was selected due to its analogous subsurface extraction processes and reliance on hydraulic fracturing. Parameters from Zhao et al. [16], who developed a variable-fidelity radial basis function (VFRBF) surrogate-assisted model for multi-objective fracture stage placement optimisation, were incorporated into the framework to generate simulated production curves. This allowed the framework outputs to be compared against a well- characterised set of published results. The resulting

cumulative production curves for a single well are shown in Figure 1, illustrating the temporal evolution of gas production predicted by the framework. Quantitative performance metrics further confirm the close alignment, with R^2 =0.9928, RMSE=0.2689, and MAE=0.2108. These values indicate that the framework reproduces single-well production dynamics with high fidelity, explaining more than 99% of the variance in the surrogate data while maintaining average errors below 0.3 million m³.

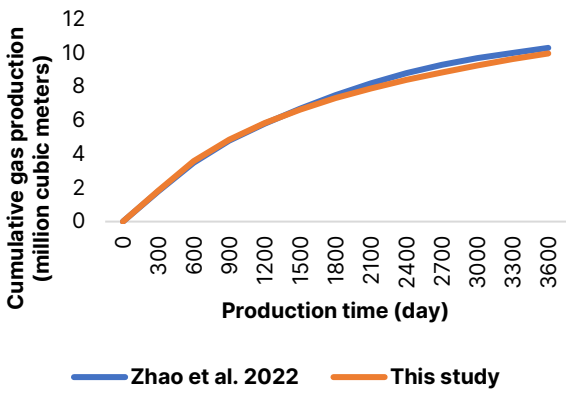

Figure 1. Cumulative production curves for a single well comparing Zhao et al. [16] and this study. The framework demonstrates strong agreement with the VFRBF model, supported by performance metrics: R^2 = 0.9928, RMSE = 0.2689, and MAE = 0.2108.

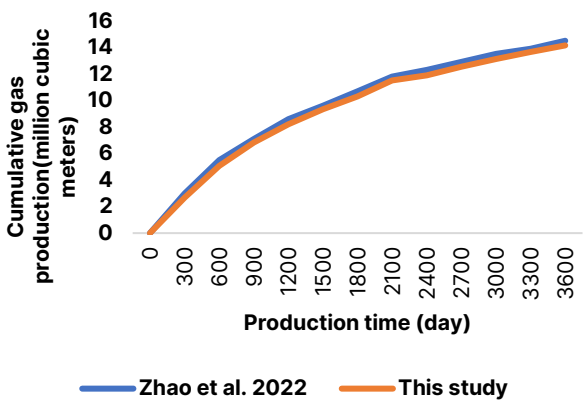

Figure 2. Cumulative production curves for multiple wells comparing Zhao et al. [16] and this study. The outputs from this framework show close alignment with the VFRBF model, validated by performance metrics: R^2 = 0.9940, RMSE = 0.3333, MAE = 0.3154.

The cumulative production curves produced by the framework show strong agreement with the results reported by Zhao et al. [16], suggesting that the framework can reliably capture the dynamics of resource extraction and production decline in hydraulically fractured reservoirs. A multi-well comparison is presented in Figure 2, highlighting consistent trends and close alignment between the aggregated outputs and the results reported by Zhao et al. over the production period. Statistical validation supports this observation, with R^2 =0.9940, RMSE=0.3333, and MAE=0.3154. These metrics show that the framework explains over 99% of the observed variance at the multi-well scale, with deviations representing less than 2% of total cumulative production.

It should be noted that this comparison primarily evaluates the framework's ability to reproduce production decline behavior in hydraulically fractured reservoirs. The benchmarking exercise, therefore, tests the production forecasting component of the framework but does not constitute direct validation of the coupled reactive transport processes, thermal feedback mechanisms, hydrogen generation kinetics, or the water–energy accounting components included in the model. Because field-scale datasets for geologic hydrogen systems are not yet available, validation of these processes remains limited. The present comparison should therefore be interpreted as a model-to-model benchmarking exercise focused on production dynamics, providing an initial consistency check rather than a comprehensive validation of all coupled processes.

While shale gas differs from hydrogen-bearing formations in geochemical properties, the observed alignment supports the framework's validity as a first-order estimation tool. This validation establishes confidence in the framework's applicability and provides a foundation for future refinement once empirical data on geologic hydrogen become available.

FRAMEWORK APPLICATION

Following benchmarking with shale gas data, the framework was applied to geologic hydrogen extraction scenarios to demonstrate its utility in estimating resource usage under realistic subsurface conditions. The application considered hydrogen generation from source rock formations, parameterized with representative geologic and operational inputs. In this case, a multi-well development consisting of five hydraulically fractured wells was assumed, reflecting typical field-scale deployment strategies in unconventional reservoirs. Table 1 shows the parameters used in this study.

The framework integrated source rock properties, fracture design parameters, and well spacing to simulate cumulative hydrogen production and associated resource requirements. The multi-well configuration allowed assessment of inter-well interactions and aggregate production behavior, providing insight into scalability and efficiency. Results indicated that the framework successfully captures the production dynamics across the five-well system, generating a cumulative production curve that highlights both individual well contributions and overall field performance. This application

demonstrates the framework's adaptability to geologic hydrogen contexts and establishes a basis for future refinement once empirical hydrogen production data become available.

The reaction extent formulation used in this study is intended as a phenomenological representation of hydrogen generation from subsurface water–rock interactions. Because field-scale kinetic data for geologic hydrogen systems remain limited, the model employs simplified parameters that capture the dominant temporal behavior of reaction-driven resource generation.

In this formulation, the half-life parameter represents an effective kinetic timescale governing the rate at which reactive minerals contribute to hydrogen generation under reservoir conditions. The maximum reaction extent parameter represents the fraction of the reactive mineral inventory that can realistically participate in hydrogen-producing reactions, accounting for constraints such as mineral accessibility, transport limitations, and thermodynamic equilibrium.

These parameters should therefore be interpreted as effective model parameters rather than direct measurements of intrinsic reaction kinetics. Their values were selected to represent plausible reaction timescales reported in studies of water–rock reactions and hydrogen-generating geochemical systems. Because uncertainties remain in these processes, the present formulation is intended as a first-order approximation suitable for framework development. Future work will include systematic sensitivity analysis and calibration as experimental and field data for geologic hydrogen systems become available.

The simulation results reveal a stark contrast between hydrogen production from the matrix and fracture domains. As shown in Figure 3, the fracture system dominates hydrogen output, with production rates rapidly rising to approximately 4000 kg/day within the first 100 days and remaining stable through 365 days. In contrast, matrix production remains negligible throughout the simulation period, indicating a limited contribution to the overall yield. This disparity underscores the critical role of fracture-enhanced transport and reactivity in enabling scalable hydrogen recovery from source rocks. The results demonstrate the framework's sensitivity to diffusivity and permeability parameters and highlight the importance of fracture network design in optimizing hydrogen extraction efficiency.

Figure 4 illustrates the temporal evolution of hydrogen output and water utilization efficiency over a 365-day simulation period. Hydrogen production rate increases steadily during the initial 180 days, reaching a plateau at approximately 4000 kg/day, consistent with fracture-dominated flow behavior observed in earlier results. Simultaneously, water efficiency, defined as kilograms of hydrogen produced per cubic meter of water used, rises gradually and reaches about 57 kg H_2/m^3 by day 365. This trend reflects improving resource utilization over time, as the system becomes more productive without proportional increases in water input.

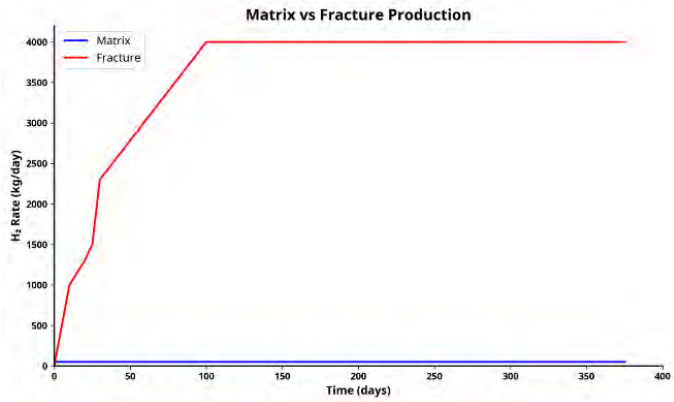

Figure 3. Hydrogen production rate over time from matrix and fracture domains

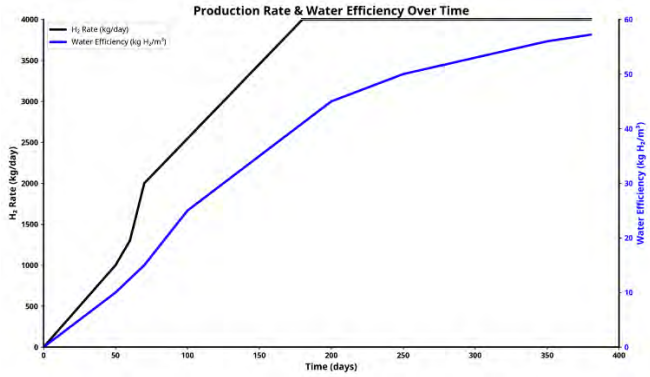

Figure 4. Hydrogen production rate and water efficiency over time

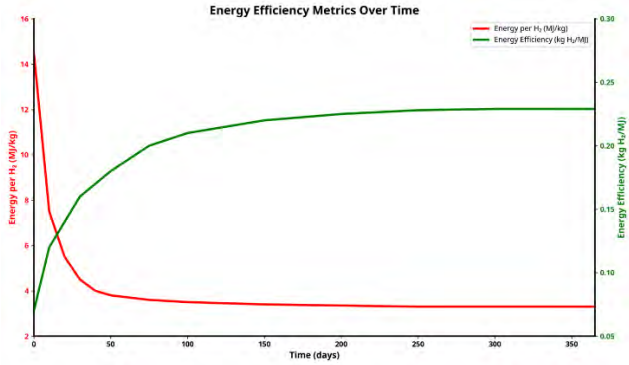

Figure 5. Energy efficiency metrics over time during hydrogen extraction

The energy and water estimates reported in this analysis are defined within a well-field operational system boundary. The accounting includes the primary operational energy inputs associated with subsurface hydrogen extraction, specifically hydraulic fracturing operations, fluid pumping during production, and gas compression required for surface handling of the produced

gas stream. Upstream processes such as drilling equipment manufacturing, infrastructure construction, and downstream processes, including hydrogen purification, storage, or long-distance transport, are not included. Consequently, the reported energy and water metrics represent direct operational requirements rather than a full life-cycle assessment. When comparisons are made with conventional hydrogen production pathways, they should therefore be interpreted as contextual benchmarks rather than strict life-cycle comparisons.

The co-evolution of these metrics highlights the framework's ability to capture both production dynamics and sustainability indicators. The plateau in hydrogen rate suggests stable fracture performance, while the continued rise in water efficiency implies that early water investment yields increasingly favorable returns. These insights are critical for evaluating the viability of geologic hydrogen extraction, especially in water-constrained environments. The results reinforce the importance of fracture design and operational timing in optimizing both output and efficiency.

Figure 5 provides insight into the evolving energy performance of the hydrogen extraction system. Over the 365-day simulation, the energy required per kilogram of hydrogen decreases steadily from around 16 MJ/kg to approximately 4 MJ/kg, indicating improved thermodynamic efficiency as the system stabilizes. Concurrently, the energy efficiency, expressed as kilograms of hydrogen produced per megajoule of input energy, increases and plateaus near 0.25 kg H_2/MJ. This inverse relationship reflects the system's optimization over time, where less energy is needed to sustain high hydrogen output. These trends underscore the framework's ability to capture not only production dynamics but also energy sustainability, offering a valuable metric for evaluating the long-term viability of geologic hydrogen operations.

The key performance indicators over the 365-day simulation period provide a comprehensive snapshot of the system's overall efficiency and sustainability. Total hydrogen production reached 1, 132, 287 kg, demonstrating the scalability of the multi-well fracture-based extraction approach. Water consumption averaged 0.02 m^3 per kilogram of hydrogen, translating to an overall water efficiency of 57 kg H_2/m^3. This aligns with the trends observed in the water efficiency graph, confirming that the system maintains high productivity with minimal water input.

Energy consumption averaged 3.40 MJ per kilogram of hydrogen, yielding an overall energy efficiency of 0.29 kg H_2/MJ. These values reflect the system's ability to optimize energy use over time, as previously shown in the declining energy-per-kilogram and rising energy efficiency metrics. Together, these indicators validate the framework's capacity to simulate not only production dynamics but also resource efficiency, offering a robust foundation for evaluating geologic hydrogen extraction under realistic operational constraints.

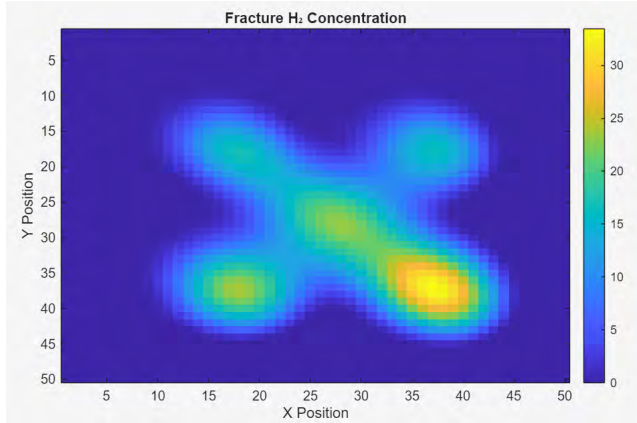

Figure 6. Fracture H_2 concentration distribution

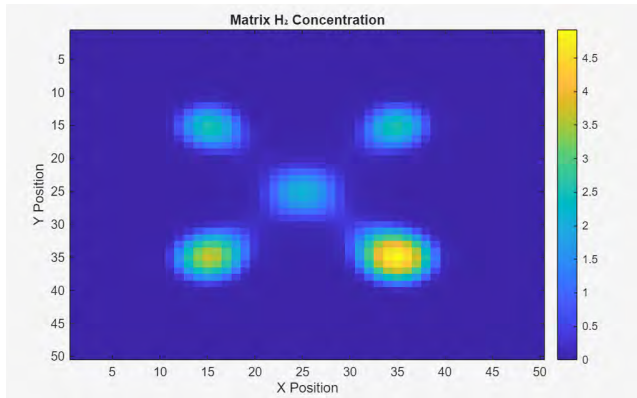

Figure 7. Matrix H_2 concentration distribution

The fracture hydrogen concentration distribution shown in Figure 6 demonstrates the dominant role of fracture networks in controlling recovery efficiency. Concentrations along the fracture plane are relatively uniform and continuous, reflecting the high permeability and aperture evolution captured in the dual-porosity dual-permeability (DPDP) framework. This behavior confirms that fracture-dominated pathways serve as the primary conduits for hydrogen transport, accounting for the vast majority of cumulative production. The spatial spread of hydrogen within the fracture system highlights how fracture propagation and connectivity enable rapid mobilization of hydrogen, minimizing diffusion limitations and ensuring efficient recovery at low water and energy costs.

By contrast, the matrix hydrogen concentration distribution in Figure 7 reveals a different behavior. Concentrations are localized and attenuated, with values significantly lower than those observed in fractures. The patchy distribution reflects the limited contribution of the matrix to overall recovery, a consequence of lower permeability and slower diffusion processes. This observation aligns with the model's prediction that the matrix

contributes minimally to total recovery, underscoring the need for enhancement strategies such as chemical stimulation to increase reactivity or thermal stimulation to accelerate diffusion and desorption. The figure also highlights the heterogeneity of reservoir performance, where matrix zones act more as secondary storage rather than primary flow pathways.

Taken together, the two figures validate the DPDP representation by illustrating the duality of reservoir behavior: fractures act as high-efficiency flow channels. At the same time, the matrix functions as a low-efficiency reservoir. This duality explains the efficiency trade-offs identified in the system-wide benchmarks, where fracture-driven recovery achieves specific energy consumption of 3.4 MJ/kg H_2 and specific water consumption of 0.02 m^3/kg H_2, an order of magnitude lower than conventional hydrogen production methods. At the same time, the limited matrix contribution points to opportunities for optimization through stimulation design and reservoir characterization. These results emphasize that while fracture systems provide the cornerstone of geologic hydrogen recovery, unlocking the potential of the matrix remains a critical frontier for future research and development.

CONCLUSION

This study presents an integrated modeling framework for estimating resource usage associated with geologic hydrogen extraction via hydraulic fracturing. By adapting published shale gas production results for benchmarking and applying the framework to a multi-well hydrogen development scenario, the model demonstrates its capability to simulate fracture-controlled production dynamics and associated resource utilization under representative subsurface conditions. The results highlight the dominant role of fracture networks in controlling hydrogen recovery and illustrate how production behavior, water use, and energy inputs evolve during reservoir operation.

A key component of the framework is the phenomenological reaction-extent formulation used to represent hydrogen-generating reactions. Although simplified relative to detailed geochemical models, this approach provides a practical means of exploring hydrogen generation and production efficiency in the absence of well-constrained field-scale kinetic data. Under the simulated multi-well scenario, the framework produced indicative performance metrics including cumulative hydrogen production of approximately 1.13 million kg, a water utilization efficiency of about 57 kg H_2/m^3, and an operational energy intensity of roughly 3.40 MJ/kg H_2, corresponding to an energy efficiency of 0.29 kg H_2/MJ. These values represent model-based estimates under the assumed conditions and illustrate the type of system-level performance indicators that the framework can generate.

The framework is intended as an exploratory tool for assessing production dynamics and resource requirements under data-limited conditions rather than as a fully validated predictive model. The benchmarking exercise against published shale gas results primarily evaluates the framework's ability to reproduce production decline behavior, while other components, including reactive transport processes and hydrogen generation kinetics, remain subject to uncertainty due to the limited availability of field-scale geologic hydrogen datasets.

Despite these limitations, the framework demonstrates how fracture-controlled production behavior, simplified reaction modeling, and resource accounting can be integrated within a single modeling structure. This integration provides a basis for evaluating water–energy–hydrogen trade-offs and exploring operational scenarios relevant to geologic hydrogen development. As experimental and field data become available, the framework can be refined through improved parameterization, sensitivity analysis, and calibration to support more detailed technical, environmental, and economic assessments of geologic hydrogen extraction.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the financial support provided by the National Research Foundation (NRF) of South Africa and the Faculty of Engineering and the Built Environment at the University of the Witwatersrand (Wits). Their funding and support made this research possible.

AUTHOR IDENTIFIERS

Author ORCIDs:
Kaitano C: 0000-0003-0303-9608
Majozi T: 0000-0003-0864-2393

REFERENCES

1. Boafo J, Dotsey S, Spencer R. Energy transition diplomacy: the eu’s pursuit of africa's critical minerals for renewable energy at whose expense?. Energy Research & Social Science 127:104213 (2025). https://doi.org/10.1016/j.erss.2025.104213
2. Dawood F, Anda M, Shafiullah GM. Hydrogen production for energy: an overview. International Journal of Hydrogen Energy 45:3847-3869 (2020). https://doi.org/10.1016/j.ijhydene.2019.12.059
3. Osman AI, Mehta N, Elgarahy AM, Hefny M, Al-Hinai A, Al-Muhtaseb AH, Rooney DW. Correction to: hydrogen production, storage, utilisation and environmental impacts: a review. Environ Chem Lett 20:2213-2213 (2022).

https://doi.org/10.1007/s10311-022-01432-x
4. Hasiuk F, Conley DM. EXPLORING GEOLOGIC HYDROGEN : A NEW FRONTIER FOR AFFORDABLE, RELIABLE ENERGY SECURITY. 2025.
5. Abdelghany MB, Shafiqurrahman A, Al-Durra A, Mohamed M, Hu J, Vaccaro A, Gao F, El Moursi MS. Hydrogen energy systems for decarbonizing smart cities and industrial applications: a review. Renewable and Sustainable Energy Reviews 226:116370 (2026). https://doi.org/10.1016/j.rser.2025.116370
6. Zgonnik V. The occurrence and geoscience of natural hydrogen: a comprehensive review. Earth-Science Reviews 203:103140 (2020). https://doi.org/10.1016/j.earscirev.2020.103140
7. Maiga O, Deville E, Laval J, Prinzhofer A, Diallo AB. Trapping processes of large volumes of natural hydrogen in the subsurface: the emblematic case of the bourakebougou H2 field in mali. International Journal of Hydrogen Energy 50:640-647 (2024). https://doi.org/10.1016/j.ijhydene.2023.10.131
8. Prinzhofer A, Moretti I, Françolin J, Pacheco C, D'Agostino A, Werly J, Rupin F. Natural hydrogen continuous emission from sedimentary basins: the example of a brazilian h2-emitting structure. International Journal of Hydrogen Energy 44:5676-5685 (2019). https://doi.org/10.1016/j.ijhydene.2019.01.119
9. Osselin F, Soulaine C, Fauguerolles C, Gaucher EC, Scaillet B, Pichavant M. Orange hydrogen is the new green. Nat. Geosci. 15:765-769 (2022). https://doi.org/10.1038/s41561-022-01043-9
10. Templeton AS, Ellison ET, Kelemen PB, Leong J, Boyd ES, Colman DR, Matter JM. Low-temperature hydrogen production and consumption in partially-hydrated peridotites in oman: implications for stimulated geological hydrogen production. Front. Geochem. 2: (2024). https://doi.org/10.3389/fgeoc.2024.1366268
11. Egert R, Neupane G, Jin W. On the viability of stimulated hydrogen generation from iron-rich formations. Geoenergy Science and Engineering 256:214130 (2026). https://doi.org/10.1016/j.geoen.2025.214130
12. Awosiji V, Graham S, Mukerji T. The concept of a blind hydrogen system. International Journal of Hydrogen Energy 158:150492 (2025). https://doi.org/10.1016/j.ijhydene.2025.150492
13. Angelina NM, Isah A, Samouei H, Okoroafor ER. Can stimulated geological hydrogen generation systems be circular?. SPE Annual Technical Conference and Exhibition : (2025). https://doi.org/10.2118/227940-ms
14. Mao S, Yu S, Xu J, Chen H, Zhao W, Blunt MJ, Kang Q, Gross M, Chen B, Van Wijk J, Yuan Q, Gao K, Kazi SR, Mehana M. Geologic hydrogen: a review of resource potential, subsurface dynamics, exploration, production, transportation, and research opportunities. Energy Environ. Sci. 18:9991-10035 (2025). https://doi.org/10.1039/d5ee02910d
15. Jiang L, Ma Y, Wu J, Liu J, Liu T. Modeling production decline for fractured wells with non-uniform fracture properties: a semi-analytical approach based on double-segment fracture model. Processes 13:3627 (2025). https://doi.org/10.3390/pr13113627
16. Zhao G, Yao Y, Wang L, Adenutsi CD, Feng D, Wu W. Optimization design of horizontal well fracture stage placement in shale gas reservoirs based on an efficient variable-fidelity surrogate model and intelligent algorithm. Energy Reports 8:3589-3599 (2022). https://doi.org/10.1016/j.egyr.2022.02.228

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimal Simulation of an Electrodialysis Reactor for the Desalination and Regeneration of Multi-Ionic Wastewater

Vicent Ayala-Andreu[a], Miguel A. Montiel[a], Vicente Montiel[a], Juan A. Labarta[bc]*

[a] Applied Electrochemistry and Electrocatalysis (LEQA) Research Group, University Institute of Electrochemistry, University of Alicante. PO 99, E-03080, Alicante, Spain
[b] Computer Optimization of Chemical Engineering Processes and Technologies (COnCEPT) Research Group, University Institute of Chemical Process Engineering, University of Alicante, Po99, Alicante, E-03080, Spain
[c] Department of Chemical Engineering, University of Alicante, Po99, Alicante, E-03080, Spain
* Corresponding Author: ja.labarta@ua.es.

ABSTRACT

The objective of the present work is to optimize the simulation of an electrodialysis reactor for the desalination and regeneration of multi-ionic wastewater with high salt contents and conductivities, within the framework in the Sustainable Development Goal 6 (clean water and sanitation) and remarking the Electrodialysis (ED) as a highly energy-efficient and sustainable technology. The mathematical modelling has been carried out by using a semiempirical model that involves an algebraic system of differential equations, including mass and charge balances (taking into account the ions present in the wastewater: Na^+, Ca^{2+}, Mg^{2+}, Cl^-, SO_4^{2-}, and HCO_3^-), and the total electrodialysis stack voltage considering ohmic drops (in the dilute and concentrate compartments), the potential of membrane in each cell pair, and the electrode potentials. In the simulation process, different theoretical and experimental parameters are necessary such as number of cells, membrane working areas, efficiency, diffusion coefficients, molar conductivity at infinite dilution, etc.The experimental parameters have been obtained initially with a specially designed batch reactor printed with 3D technology, to imitate the hydrodynamics and the behavior of the industrial final reactor used. Additionally, to increase the accuracy of the simulation results, same parameters such as the resistance of the membranes (r_m), the transport numbers and the mass transfer coefficient (k_m) can be optimized by using a metaheuristic derivative-free optimizer, specifically a proprietary version of Particle Swarm Optimization (PSO) implemented in MatLab®.The results obtained yield a satisfactory reproduction of the observed experimental behavior, analyzing the evolution over time of the intensity, conductivity, and concentration of the concentrated and diluted currents, respectively. In addition, the evolution of the different contributions of the stack voltage and the membrane resistances are shown, in order to confirm the coherence of the results obtained.

Keywords: Wastewater, Modelling, Membranes, Process Design, Particle Swarm Optimization

INTRODUCTION

Wastewater Treatment Plants (WWTP) constantly receive water from industrial process that can present salinity levels from 1 wt % up to values similar to seawater (3.5 wt %) [1], including also domestic uses such as detergents or water softening processes [2], or waste derived from constructions and transportation activities [3]. This problem also occurs in coastal areas due to marine intrusion. On the other hand, waste water with high salinity degree released into the environment may cause damage by contaminating soils, surfaces, groundwater, or any water body with which they come into contact [4] There is evidence that it affects ecosystem processes by altering the decomposition of organic matter, biomass production, and reducing biodiversity [5].

Thus, in order to reuse any treated water as regenerated water for instance in urban cleaning tasks, garden watering, agricultural uses, etc., regulations must be followed to mitigate their impact. According to the FAO

(Food and Agriculture Organization), for conductivities higher than 3 mS/cm water must be under several restrictions and controls, between 0.7 and 3 mS/cm those restrictions are lower, and with electrical conductivities under 0.7 mS/cm there are any restriction [6].

In this context, electrodialysis (ED) appears as a very energy efficient and sustainable wastewater desalination treatment for those type of salinized water compared to other processes [7, 8]. Additionally, ED process is driven by an electric field and neutral species present in wastewater that may be associated with beneficial elements for plants, such as urea or organic nitrogen, are almost not altered by the process [9, 10, 11].

The ED process is conducted in a membrane stack (electrodialysis reactor) where two kind of exchange membranes: N+1 cation exchange membranes (CEM) and N anion exchange membranes (AEM) are alternatively located generating a series of N+1 compartment between two different consecutives membranes. A potential difference is applied across the membrane stack using two electrodes placed at the sides of it. The feed solution flows through the membranes, generating a dilute channel and a concentrate channel as a result of ion migration. The cathode is negatively charged, so cations migrate toward it by passing through the CEMs, but they cannot pass through the AEMs. Likewise, the positively charged anode attracts anions, which migrate toward it under the same selective transport mechanism [12].

To scale up an ED process without the need to perform many experimental tests at each stage of the scaling procedure, it is absolutely necessary to have a mathematical model to simulate the main characteristics of the process.

However, electrochemical processes are particularly complex to model due to their multiscale and multiphysics nature. In such processes, phenomena of different origins coexist, requiring the simultaneous consideration of thermodynamic and kinetic aspects together with the distribution of electric potential and ionic transport in the solution [13, 14]. In this sense, there are ED mathematical models capable of predicting the desalination rate and the current density [15] depending on the water composition, as well as the electrical energy consumption of the process [16].

In the case of the model presented by Ortiz et al. [16], parameters such as membrane resistance (Rm) and the potential difference between electrodes are introduced and assumed to be constant. The potential difference depends on the overpotential (μ), which also depends on the current density (j), the exchange current density (j0), and the Tafel slope (b). All of these parameters, including membrane resistance, vary depending on the type of electrolyte, its concentration and the hydrodynamics of the reactor. Treating these parameters as constants leads to a loss of accuracy in the results, for that, it is needed to determinate these parameters under specific operating conditions of each reactor or experiment.

Thus, this work presents a new modified and optimized version of the model developed by Ortiz [16, 17] further extended to a multi-ionic wastewater system. As case study, a Wastewater Treatment Plant in Santa Pola (Spanish east cost) has been selected where, even with effective tertiary treatment, the final water has conductivities between 4 and 7 mS/cm that hinder its reuse for purposes such as irrigation. Although the regulations considered these waters have high conductivity, for a desalination process they are relative low salinities.

MATHEMATICAL MODEL

The proposed model is structured into two blocks: one that defines mass transfer through mass balances equations, and another that is responsible for providing all the electrical data via the calculation of the total potential.

Mass balance

It is necessary to know the ionic composition of the solution, both in the reactor (concentrated and diluted compartments) and in the concentrated and diluted product tanks. To this end, the corresponding mass balances over the time are established. Assuming that the electrodialysis reactor and the reservoirs or tanks follow a perfectly mixed reactor model and that the compartments are equivalent, the following mass balances can be established for each cation (*i*) present in the concentrate and dilute compartments of the electrodialyzer, respectively:

$$NV_k \frac{dC_{conc,i}}{dt} = Q_{conc}C^0_{conc,i} - Q_{conc}C_{conc,i} + NA(J_i^{cat} - J_i^{an}) \quad (1)$$

$$NV_k \frac{dC_{dil,i}}{dt} = Q_{dil}C^0_{dil,i} - Q_{dil}C_{dil,i} - NA(J_i^{cat} - J_i^{an}) \quad (2)$$

where N is the number of cells, V_k is the volume of the compartments (m^3), t is time (s), $C^0_{dil,\,i}$ and $C^0_{conc,\,i}$ are the initial concentrations or inlet concentrations of the ion i in the dilute and concentrate compartments, respectively, and $C_{dil,\,i}$ and $C_{conc,\,i}$ are the corresponding outlet concentrations (mol/m^3). Q_{dil} and Q_{conc} are the flow rates of the dilute and concentrate compartments (m^3/s), A is the electrode area (m^2), and $J_{cat,\,i}$ and $J_{an,\,i}$ are the fluxes of cations and anions of species i through the cation-exchange and anion-exchange membranes, respectively (mol/m^2).

For anions, the expression would be equivalent, but with the opposite sign for the ionic fluxes:

$$NV_k \frac{dC_{conc,i}}{dt} = Q_{conc}C^0_{conc,i} - Q_{conc}C_{conc,i} +$$

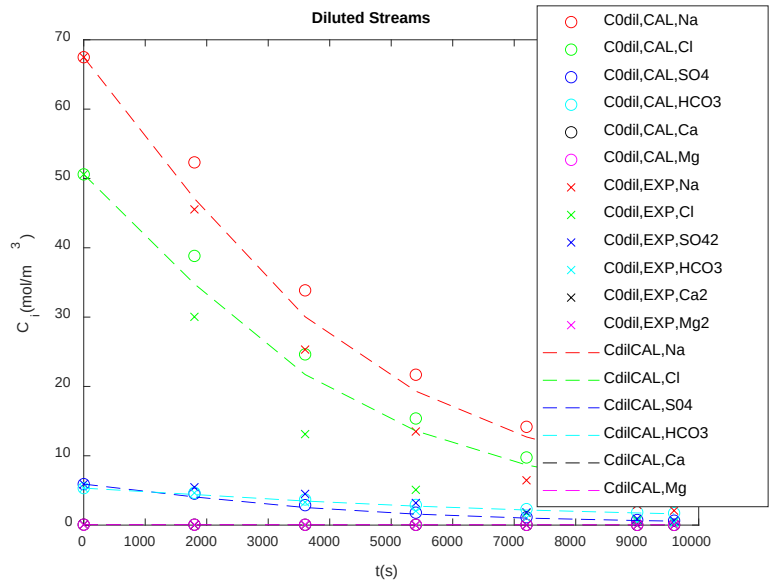

a)

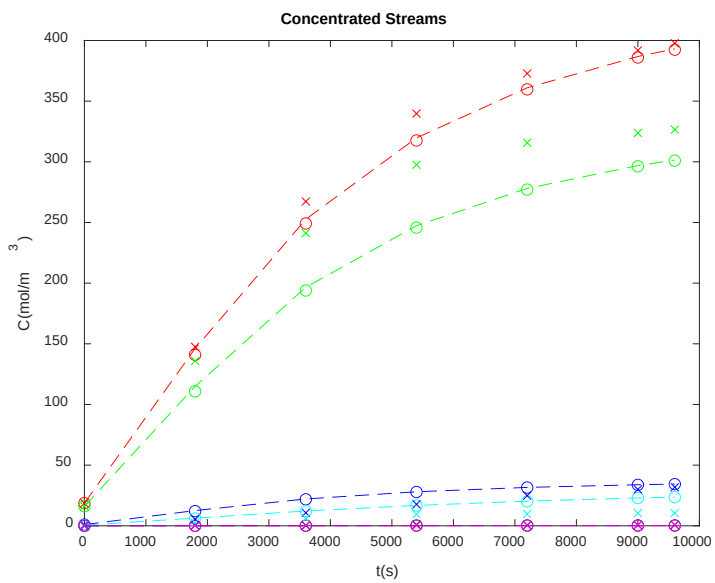

b)

Figure 1. Composition profiles obtained for a sample with a initial conductivity of 6 mS/cm.

$$+NA(J_i^{an} - J_i^{cat}) \quad (3)$$

$$NV_k \frac{dC_{dil,i}}{dt} = Q_{dil}C^0_{dil,i} - Q_{dil}C_{dil,i} - NA(J_i^{an} - J_i^{cat}) \quad (4)$$

The ionic fluxes for ions i are given by the sum of the migration and diffusion components:

$$J_i^{cat} = \sum J_{i,mig}^{cat} + \sum J_{i,dif}^{cat} = \frac{\phi t_i^{cat}}{z_i F} j - D_i^{cat} \frac{C^{w,cat}_{conc,i} - C^{w,cat}_{dil,i}}{l_m^{cat}} \quad (5)$$

$$J_i^{an} = \sum J_{i,mig}^{an} + \sum J_{i,dif}^{an} = \frac{\phi t_i^{an}}{z_i F} j - D_i^{an} \frac{C^{w,an}_{conc,i} - C^{w,an}_{dil,i}}{l_m^{an}} \quad (6)$$

where ϕ is the experimental current efficiency, t^{cat}_i and t^{an}_i are the transport numbers of ion i in the cation-exchange and anion-exchange membranes, respectively, z_i is the ionic charge, F is Faraday's constant (C/mol), j is the current density (A/m^2), D^{cat}_i and D^{an}_i are the diffusion coefficients of ion i in the cation-exchange and anion-exchange membranes, respectively (m^2/s), $C^{w,\,cat}_{conc,\,i}$ and $C^{w,\,cat}_{dil,\,i}$ are the concentrations of ion i at the surface of the cation-exchange membrane for the concentrate and dilute compartments, respectively (mol/m^3), $C^{w,\,an}_{conc,\,i}$ and $C^{w,\,an}_{dil,\,i}$ are the analogous concentrations for the anion-exchange membrane, and l^{cat}_m and l^{an}_m are the thicknesses of the cation-exchange and anion-exchange membranes (m). This expression is equivalent for cations and anions, the only difference lies in the values of the corresponding parameters.

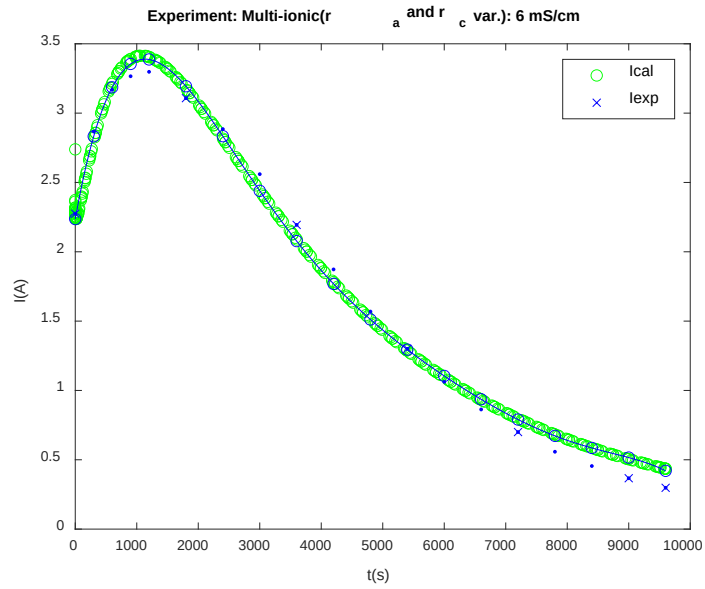

c)

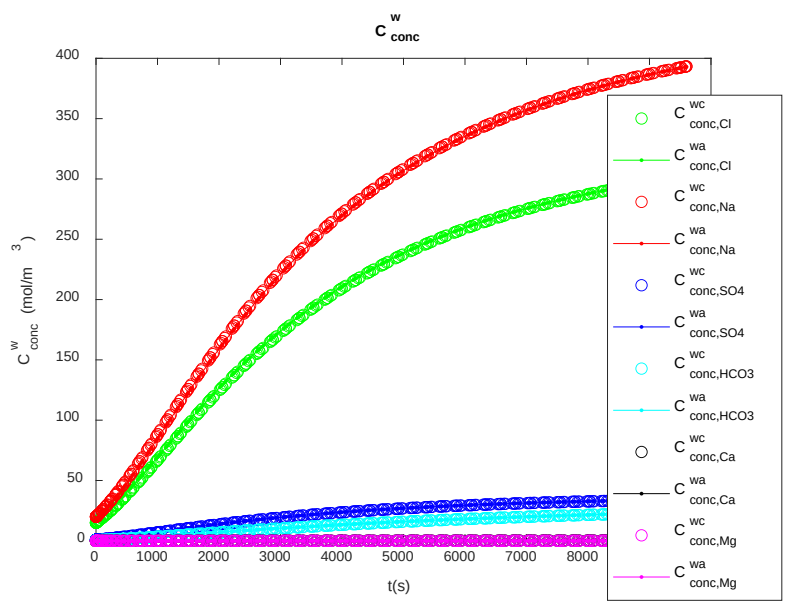

d)

Figure 1(cont.). Intensity, membrane compositionc and potential profiles obtained for a sample with an initial conductivity of 6 mS/cm.

The corresponding concentrations in the membranes can be calculated as follows:

$$C_{dil,i}^{wa} = C_{dil,i} - \frac{\phi j}{|z_i|Fk_m}\left(t_{i,contra\text{ or }co}^{ma} t_{contra\text{ or }co}^{a} - t_i^{dil}\right) \quad (7)$$

$$C_{conc,i}^{wa} = C_{conc,i} + \frac{\phi j}{|z_i|Fk_m}\left(t_{i,contra\text{ or }co}^{ma} t_{contra\text{ or }co}^{a} - t_i^{conc}\right) \quad (8)$$

$$C_{dil,i}^{wc} = C_{dil,i} - \frac{\phi j}{|z_i|Fk_m}\left(t_{i,contra\text{ or }co}^{mc} t_{contra\text{ or }co}^{c} - t_i^{dil}\right) \quad (9)$$

$$C_{conc,i}^{wc} = C_{conc,i} + \frac{\phi j}{|z_i|Fk_m}\left(t_{i,contra\text{ or }co}^{mc} t_{contra\text{ or }co}^{c} - t_i^{conc}\right) \quad (10)$$

where the value of the transport numbers $t^{c\text{ or }a}_{contra}$ is 0.999, with $t^{c\text{ or }a}_{co}$ = 1-$t^{c\text{ or }a}_{contra}$.

In the case of transport numbers, the general equations are:

$$t_i^{conc\text{ or }dil} = \frac{|z_i|u_i\ C_i^{conc\text{ or }dil}}{\sum_i |z_i|u_i\ C_i^{conc\text{ or }dil}} \quad (11)$$

$$t_{app,i,contra\text{ or }co}^{ma\text{ or }mc} = \frac{|z_i|y_i^{contra\text{ or }co} C_{conc\text{ or }dil,i}^{wa\text{ or }wc}}{\sum_i |z_i|y_i^{contra\text{ or }co} C_{conc\text{ or }dil,i}^{wa\text{ or }wc}} \quad (12)$$

where y_i is a term related to ionic mobility in the membranes (if it were in solution, the notation would be u_i), and C_i is the ion concentration in the concentrate or dilute solution. This concentration may vary depending on whether it is in the membrane or in the solution [18].

The mass balances for the concentrate and dilute

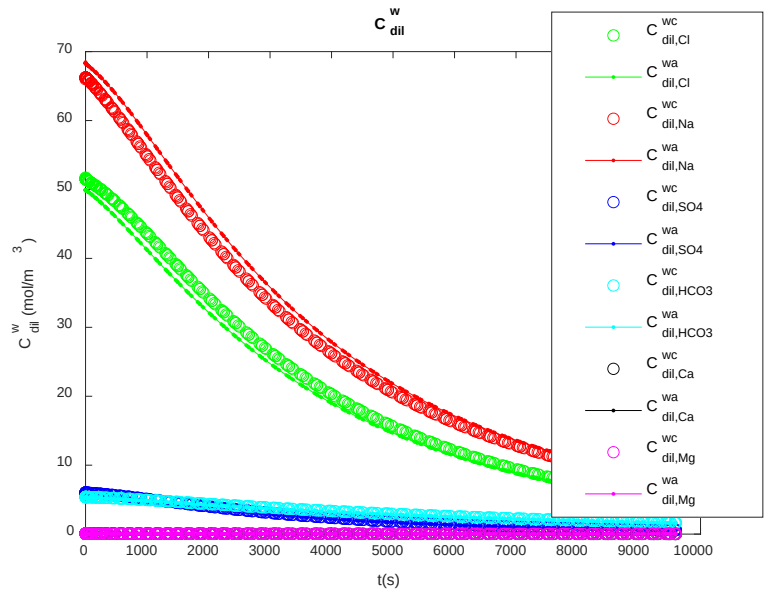

e)

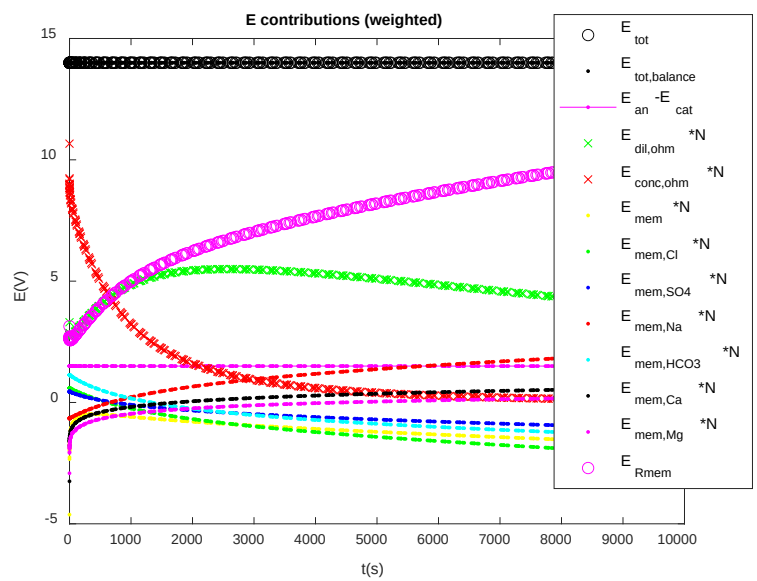

f).

Figure 1(cont.). Intensity, membrane compositionc and potential profiles obtained for a sample with an initial conductivity of 6 mS/cm.

tanks can be expressed as follows:

$$\frac{C^0_{conc,i}V^T_{conc}}{dt} = Q_{conc}\left(C_{conc,i} - C^0_{conc,i}\right) \quad (13)$$

$$\frac{C^0_{dil,i}V^T_{dil}}{dt} = Q_{dil}\left(C_{dil,i} - C^0_{dil,i}\right) \quad (14)$$

where V^T_{dil} and V^T_{conc} are the volumes of the dilute and concentrate tanks, respectively (m^3).

Electrical potential difference in the electrodialysis stack

The current intensity (I) that flows through the electrodialyzer is related to the applied electrical potential. The total electrical potential drop for an electrodialyzer can be obtained as follows:

$$\varphi_{tot} = (\varphi_{an} - \varphi_{cat}) + N\left(\varphi^{dil}_{ohm} + \varphi^{conc}_{ohm} + \varphi_{mem}\right) + R_{mem} \cdot \frac{I}{A} \quad (15)$$

$$\text{with: } \varphi^{dil}_{ohm} = \frac{L\,I}{\Lambda_{dil}\,A};\ \varphi^{\text{conc}}_{\text{ohm}} = \frac{\text{L I}}{\Lambda_{\text{conc}}\,\text{A}} \quad (16)$$

$$\varphi_{mem} = \sum_i^{ions} \frac{RT}{z_i F} ln\left(\frac{C^{wc}_{conc,i}}{C^{wc}_{dil,i}}\frac{C^{wa}_{conc,i}}{C^{wa}_{dil,i}}\right) \quad (17)$$

where φ_{an} and φ_{cat} are the anode and cathode potentials, respectively (V), φ^{dil}_{ohm} is the ohmic drop in the dilute compartment (V), φ^{conc}_{ohm} is the ohmic drop in the

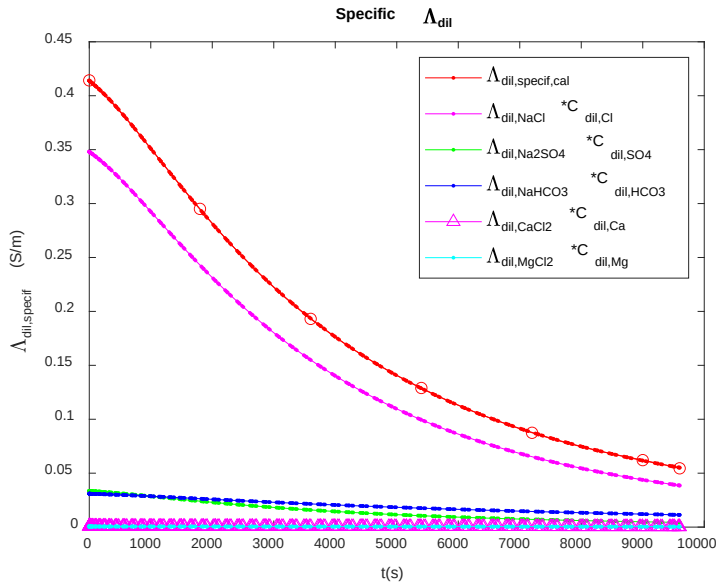

g)

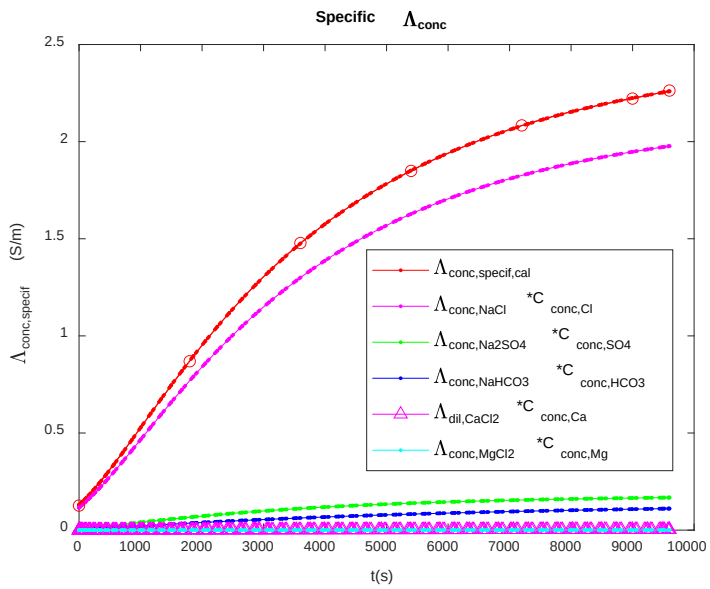

h)

Figure 1.(cont.). Intensity, membrane compositionc and potential profiles obtained for a sample with an initial conductivity of 6 mS/cm.

concentrate compartment (V), φ_{mem} is the total membrane potential at each unit cell (V) due to the presence of all the ions present in the water, considering ideal mixture, Λ_{dil} and Λ_{conc} are the total specific conductivities (S/m) of the concentrated or diluted solution taking into account all the individual salts present in the solution (Table 1), respectively, L is the membrane gap, and $R_{mem} \cdot l/A$ is the term associated with the total resistance contributed by all the membranes (N+1 CEM, and N AEM), which has been obtained by correlation the experimental intensity values (obtained in A 3D printed reactor [18]) with the conductivity of the diluted solution [eq. 18] by using the proposed model and a proprietary version of Particle Swarm Optimization (PSO) implemented in MatLab®, that has been successfully used for the optimal design of different unit operations [19].

$$R_{mem} = Nra + (N+1)rc \tag{18}$$

with: $ra = 58.04\Lambda_{dil}^{-1.66}$; $rc = 120.46\Lambda_{dil}^{-1.42}$ and $\Lambda_{dil} = \sum_{s}^{Salts} \Lambda_s \, C_{dil,s}$.

RESULTS

To solve the obtained algebraic system of differential equations, additional equations and parameters are

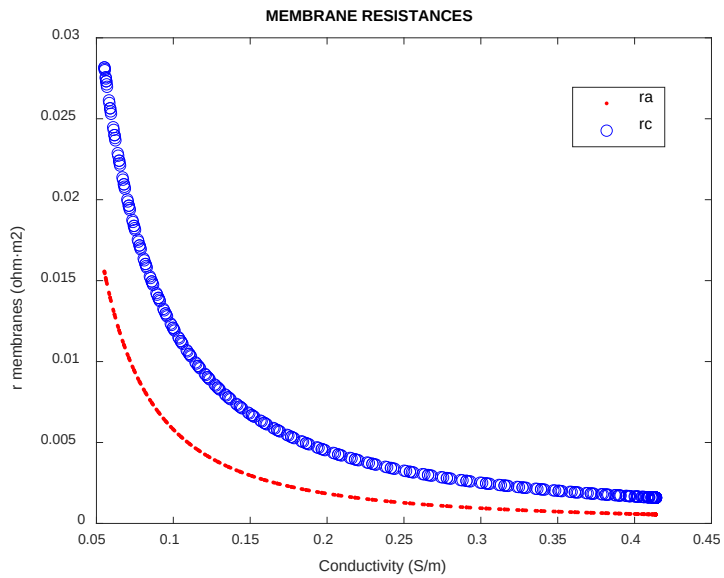

i)

Figure 1.(cont.). Specific conductivities and membrane resistences profiles obtained for a sample with an initial conductivity of 6 mS/cm.

required (Tables 1-3), such as the characteristic parameter y, related to ion mobility in the membrane, the ion mobility in the solution u; the Falkenhagen equation for the molar conductivity of each salt of the solution (and its parameters B_0, B_1 and B_2), among others. Figures 1 show the main profiles obtained for the simulation of a case study with an initial conductivity of 6 mS/cm (concentration of the concentrated and diluted streams, intensity and, the different contributions of the stack voltage). As can be seen, the proposed model satisfactorily represents the experimental trends, analyzing the different evolutions over time, especially in the initial section up to approximately 3000 s, when a conductivity of less than 2.5 mS/cm is reached, which is the limit allowed for safe reuse of regenerated water in irrigation. However, future work should incorporate e.g. the non-ideality of the liquid mixture by using the corresponding activity coefficients in the calculation of the membrane potential.

Table 1. Parameters for Falkenhagen equation, with U= 4 Å. $\Lambda_{salt} = \Lambda_0 - (B_1\Lambda_0 + B_2)\frac{C_{salt}^{0.5}}{1+B_0 U C_{salt}^{0.5}}$.

Salt	$\Lambda_0 \cdot 10^4$	B_0	B_1	$B_2 \cdot 10^4$
NaCl	126.45	0.358	0.240	60.32
Na_2SO_4	129.80	1.010	2.269	120.65
$NaHCO_3$	117.94	0.510	0.544	109.98
$CaCl_2$	135.77	1.336	1.986	120.65
$MgCl_2$	129.34	1.356	1.973	120.64

Table 2. Values of the parameters used in the model.

Parameter	Value	Parameter	Value
A (m^2)	0.02	Q (m^3/s)	5.56 10^{-5}
L_{an} (m)	$1.5 \cdot 10^{-4}$	φ_{tot} (V)	14
L_{cat} (m)	$1.6 \cdot 10^{-4}$	φ_{an} - φ_{cat} (V)	1.5
L (m)	0.001	Φ(exp.)	0.87
N	10	k_m	0.0005
V_k ($m^{3)}$	$1.96 \cdot 10^{-5}$	F (C/mol)	96485
V^T_{conc} (m^3)	0.004	R (J(mol K)	8.31
V^T_{dil} (m^3)	0.025	T (K)	298.15

Table 3. Data for streams without pressure change. Molar ionic conductivities and limiting diffusion coefficients in water at 298 K [20]. Characteristic parameter (y) related to ion mobility in ion-exchange membranes [17], and ionic mobilities (u) in water at 298 K [21, 22].

Ion	$\lambda \cdot 10^{-4}$ (m^2S/mol)	$D \cdot 10^{-9}$ (m^2/mol)	y *contra / co*	$u \cdot 10^{-8}$ (m^2/sV)
Ca^{2+}	118.94	1.584	0.3/0.1	6.17
Mg^{2+}	106	1.412	0.1/0.03	5.50
Na^+	50.08	1.334	0.7/0.1	5.19
Cl^-	76.31	2.032	1/0.1	7.90
HCO^{-3}	44.5	1.185	0.5/0.3	4.61
$SO2^{-4}$	160	2.13	1/0.1	8.29

ACKNOWLEDGEMENTS

The authors gratefully acknowledge financial support by the University of Alicante's Vice-rectorate of Research own program for the promotion of R+D+I in 2021, EURODIA Industries (EURODIA1-21Y) and to the Ministry of Science and Innovation from Spain under the project PID2024-155232NB-I00 (Net-0-ChemI).

Cofinanciado por la Unión Europea

author identifiers

Author ORCIDs:
Ayala-Andreu V: 0009-0002-7724-2799
Montiel MA: 0000-0003-3097-0975
Montiel V: 0000-0003-2353-2993
Labarta JA: 0000-0002-4870-2031

REFERENCES

1. Intrasungkha N, Keller J, Blackall LL. Biological nutrient removal efficiency in treatment of saline wastewater. *Water Sci Technol* 39 (6): 183-190 (1999) https://doi.org/10.1016/S0273-1223(99)00138-9
2. Corral AM, Val del Río Á, Campos Gómez JL. Treatment and valorisation of saline wastewater: principles and practice. IWA Publishing (2021). https://doi.org/10.2166/9781789060645
3. Singh N, Poonia T, Siwal SS, Srivastav AL; Sharma HK, Mittal SK. Challenges of water contamination in urban areas. *Current Directions in Water Scarcity Research* 9: 173- 202 (2022). https://doi.org/10.1016/b978-0-323-91838-1.00008-7
4. Lefebvre O, Moletta R. Treatment of organic pollution in industrial saline wastewater: a literature review. Water Research 40:3671-3682 (2006). https://doi.org/10.1016/j.watres.2006.08.027
5. Berger E, Frör O, Schäfer RB. Salinity impacts on river ecosystem processes: a critical mini-review. Phil. Trans. R. Soc. B 374:20180010 (2018). https://doi.org/10.1098/rstb.2018.0010
6. Ayers RS, Westcot DW. Water quality for agriculture. *FAO Irrigation and Drainage Paper* Vol. 29 (1) (1985) https://www.fao.org/4/T0234e/T0234e00.htm
7. Patel SK, Biesheuvel PM, Elimelech M. Energy consumption of brackish water desalination: identifying the sweet spots for electrodialysis and reverse osmosis. ACS EST Eng. 1:851-864 (2021). https://doi.org/10.1021/acsestengg.0c00192
8. Patel SK, Lee B, Westerhoff P, Elimelech M. The potential of electrodialysis as a cost-effective alternative to reverse osmosis for brackish water desalination. Water Research 250:121009 (2024). https://doi.org/10.1016/j.watres.2023.121009
9. Gong Y, Wang XL, Li-xin Y. Process simulation of desalination by electrodialysis of an aqueous solution containing a neutral solute. Desalination 172:157-172 (2005). https://doi.org/10.1016/j.desal.2004.06.200
10. Campione A, Gurreri L, Ciofalo M, Micale G, Tamburini A, Cipollina A. Electrodialysis for water desalination: a critical assessment of recent developments on process fundamentals, models and applications. Desalination 434:121-160 (2018). https://doi.org/10.1016/j.desal.2017.12.044
11. Li H, Li M, Wang H, Tan M, Zhang G, Huang Z, Yuan X. A review on migration and transformation of nitrogen during sewage sludge thermochemical treatment: focusing on pyrolysis, gasification and combustion. Fuel Processing Technology 240:107562 (2023). https://doi.org/10.1016/j.fuproc.2022.107562
12. Van der Bruggen B. Advances in electrodialysis for water treatment. Advances in Membrane Technologies for Water Treatment :185-203 (2015). https://doi.org/10.1016/b978-1-78242-121-4.00006-x
13. Yang X, Zhuang Y, Zhu J, Le J, Cheng J. Recent progress on multiscale modeling of electrochemistry. WIREs Comput Mol Sci 12: (2021). https://doi.org/10.1002/wcms.1559
14. Ding L, Azimi G. Mathematical modeling of rare earth element separation in electrodialysis with adjacent anion exchange membranes and ethylenediaminetetraacetic acid as chelating agent. Sci Rep 14: (2024). https://doi.org/10.1038/s41598-024-62885-4
15. Wright NC, Shah SR, Amrose SE, Winter AG V. A robust model of brackish water electrodialysis desalination with experimental comparison at different size scales. Desalination 443:27-43 (2018). https://doi.org/10.1016/j.desal.2018.04.018
16. Ortiz JM, Sotoca JA, Expósito E, Gallud F, García-García V, Montiel V, Aldaz A. Brackish water desalination by electrodialysis: batch recirculation operation modeling. Journal of Membrane Science 252:65-75 (2005). https://doi.org/10.1016/j.memsci.2004.11.021
17. ORTIZ J, EXPOSITO E, GALLUD F, GARCIAGARCIA V, MONTIEL V, ALDAZ A. Desalination of underground brackish waters using an electrodialysis system powered directly by photovoltaic energy. Solar Energy Materials and Solar Cells 92:1677-1688 (2008). https://doi.org/10.1016/j.solmat.2008.07.020
18. Ayala-Andreu V, García-García V, Labarta JA,

Montiel V, Montiel MA. Determination of critical parameters for the mathematical modelling of an electrodialysis process using a 3d-printed reactor. Chemical Engineering Journal 530:173277 (2026). https://doi.org/10.1016/j.cej.2026.173277

19. Mekidiche Z, Labarta JA, Javaloyes-Anton J, Caballero JA. From power to heat: strategies for electrifying distillation for sustainable chemical processes. Applied Thermal Engineering 257:124316 (2024). https://doi.org/10.1016/j.applthermaleng.2024.124316
20. Haynes WM, Lide DR, Bruno TJ. CRC handbook of chemistry and physics. CRC Press (2016). https://doi.org/10.1201/9781315380476
21. Varcoe JR, Atanassov P, Dekel DR, Herring AM, Hickner MA, Kohl PA, Kucernak AR, Mustain WE, Nijmeijer K, Scott K, Xuk T, Zhuangl L. Anion-exchange membranes in electrochemical energy systems. *Energy Environ. Sci.* 7 (10): 3135-3191 (2014) http://dx.doi.org/10.1039/C4EE01303D
22. Atkins P, Paula J. Physical Chemistry. 8th ed. Vol. 1. *WH Freeman and Company* (2006), pp. 1-1053.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Data-Efficient Symbolic Regression Framework for Automated Interpretable Bioprocess Modelling

Luca Riezzo[a], Alexander Rogers[a], Harry Kay[a], Dongda Zhang*

[a] Department of Chemical Engineering, The University of Manchester, Manchester, M13 9PL, UK
* Corresponding Author: dongda.zhang@manchester.ac.uk

ABSTRACT

Bioprocess modelling, optimisation and scale-up are central components for improving sustainable manufacturing within pharmaceutical and chemical industries. However, developing accurate bioprocess digital twins remains a challenging process. Conventional mechanistic models are difficult to construct because of limited mechanistic understanding and large complexity of cellular metabolisms. While data-driven models have gained popularity, they often require large amounts of experimental data that is often time consuming to obtain and lack any quantitative description of the process. Hybrid modelling methods have emerged as promising alternatives however fail to provide physical insight to the root cause of model error. This work therefore presents a promising solution by developing a data-efficient symbolic regression (SR) based framework to enable the automated discovery of interpretable bioprocess models. A universal kinetic model backbone was used to capture overall process behaviour, while SR was applied to strategically uncover the structures of critical kinetic terms within the backbone. Two frameworks, embedding SR directly into the kinetic model backbone or identifying time-varying parameter profiles prior to SR, were benchmarked using an in-silico yeast fermentation case study. The results demonstrated that independently identifying individual kinetic terms was crucial for recovering the ground-truth model, while refining SR-generated candidates through a novel local iterative structural correction strategy significantly improved convergence to the true kinetic expressions, surpassing model-based design of experiments in data efficiency. This study therefore enables automated yet interpretable model construction for small-data bioprocess applications, paving the way towards augmented intelligence driven bioprocess modelling and accelerating digital twin development for process optimisation and control.

Keywords: Machine learning, Data intelligence, Biochemical reaction kinetics, Symbolic Regression, Augmented intelligence

INTRODUCTION

Accurate bioprocess models are critical for control, optimisation and accelerating the transition from laboratory experiments to production scale operations. Mechanistic models are widely adapted for modelling bioprocess kinetics due to their interpretability and ability for providing insight into the governing reaction kinetics. Represented using ordinary differential equations, they are derived using prior knowledge of underlying chemical reactions that occur within the biological cell. Typical strategies to construct mechanistic models for biochemical reaction networks include using the Monod and adapted Monod kinetic structures. This approach finds prominent application within bioprocess modelling and is widely used to simulate the evolution of key biological states over time. However, identifying the correct underlying mechanistic structure is often a difficult task due to the large network of complex metabolic pathways that contribute to apparent stochastic behaviour.

In response to the challenges faced with mechanistic approaches, data-driven modelling has gained popularity as a viable solution in literature. Black-box models can use large datasets gained from running experiments to make accurate kinetic predictions without requiring prior knowledge of underlying reaction mechanisms. Machine learning-based models have seen success in modelling chemical and biochemical reaction kinetics [1]

however, their practical applications remain limited. This is because data-driven models often require large amounts of data which is a scarce resource in bioprocess industry due to large expenses and risks involved when running experiments on novel systems.

Recently, hybrid modelling has emerged as a promising alternative to address the challenges with both mechanistic and data-driven modelling by combining machine learning techniques with mechanistic knowledge. Prior knowledge is embedded using a mechanistic backbone while the data-driven component leverages data to account for model error that may arise from any incomplete physical understanding. In this way, hybrid modelling has been able to demonstrate advantages by making accurate predictions for complex biochemical reaction kinetics [2] and offer a suitable platform to model the dynamics of similar biological strains by applying transfer learning [3]. However, despite the success in literature, there exist various practical challenges when applying hybrid modelling methods for modelling novel biological strains. The prominent disadvantage lies with providing no mechanistic insight into the root cause of model error, that is essential for in applications that require a high degree of explainability and limits human understanding of underlying reaction kinetics.

Motivated by the need for deeper mechanistic understanding, automated model structure identification methods have been proposed in literature for overcoming the limitations of hybrid modelling. Among these methods the sparse identification of nonlinear dynamics (SINDy) has gained most attention and has seen successful application for identifying accurate model structures for chemical reaction systems [4]. By leveraging sparse regression techniques, SINDy algorithm can construct kinetic rate equations using a library of predefined candidate functions found most likely within literature. However, application for biochemical processes remains limited by the requirement for a suitably large library of candidate functions that is particularly challenging to obtain if there exists no prior physical knowledge of biological strains. This has motivated the development of more flexible model structure identification techniques such as symbolic regression (SR) that uses genetic programming to explore a much broader solution space. the advantage of SR is that the models identified are a function of only predefined operators and variables and not limited by imposed structural forms. Presently SR has been successful in discovering chemical rate equations without a-priori [5] and simple unstructured bioprocess models [6] however the application for complex biochemical systems remains relatively unexplored due to significant challenges to guarantee interpretability of resultant models for real-world biological systems.

Therefore, to address the described challenges, this work utilises symbolic regression (SR) to develop an automated model structure identification framework that can uncover underlying kinetic information with high data-efficiency. Through this approach, is it possible to simultaneously accelerate the development of accurate bioprocess models while also deepening our mechanistic understanding of emerging processes.

METHODOLOGY

Symbolic regression and information criteria

Symbolic regression (SR) is a data-driven technique that leverages genetic programming to simultaneously identify both functional form and numerical constants that best describe a dataset. SR explores a solution space of equations represented as expression trees built from predefined mathematical operators and variables without imposing a set structural form. At each iteration, a set of expressions are drawn at random from an initialised population and undergo tournament selection process. The expressions are by default ranked according to their mean squared error or a fitness function of the user's definition where the highest scoring models proceed to the next iteration with a chance of undergoing mutation or crossover. In this work, SR was implemented using the PySR package, which combines the evolutionary search with gradient-based optimisation in an iterative evolve–simplify–optimise process. During this cycle, candidate equations are periodically simplified to remove redundancy, and their numerical constants are refined using local gradient-based methods. Rather than producing a single model, this approach generates a set of candidate equations across varying levels of complexity, forming a Pareto front that captures the trade-off between model simplicity and fitting accuracy.

Because symbolic regression produces many candidate models, this work adopts the corrected Akaike Information Criterion (AICc) as a systematic method for model selection, balancing goodness of fit with model simplicity. AICc evaluates models based on their residual error, number of parameters, and sample size, applying an additional correction term that increasingly penalises overly complex models when data is limited. This makes it particularly suitable for bioprocess applications, where experimental data is often scarce. By selecting the model with the lowest AICc, the framework identifies the best trade-off between accuracy and parsimony, reducing overfitting and providing a consistent, objective approach to choosing the most plausible model without relying on prior mechanistic assumptions.

Approach 1 - Concentration space

Specifically, SR was implemented in two approaches and both methods were compared within this work. The first framework was designed to incorporate concentration profiles within the SR model discovery process by

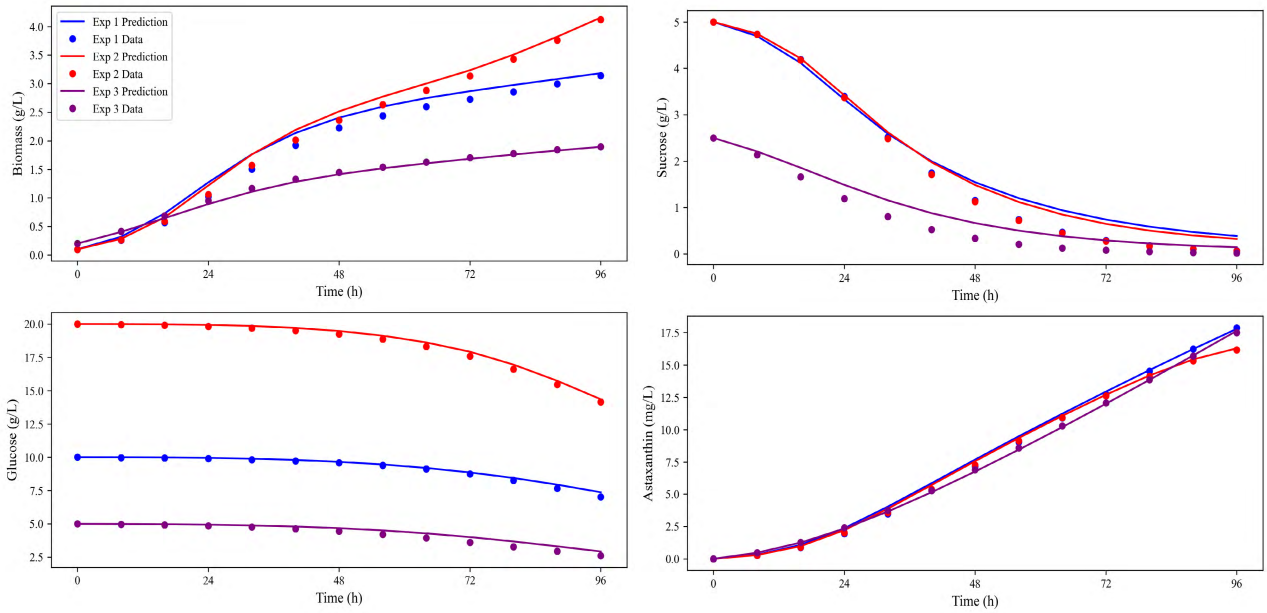

Figure 1. Fitting performance of the complete kinetic model using expressions discovered by SR for μ_{m1} and μ_{m2}. Colours blue, red and purple represent experiment 1, 2, and 3 respectively. (a), (b), (c) and (d) shows the fitting result for concentrations of biomass, sucrose, glucose and astaxanthin, respectively.

embedding SR expressions directly into the kinetic model, ensuring that candidate models are evaluated based on how well they reproduce observed concentration dynamics. In this method, symbolic expressions representing key kinetic terms are iteratively generated and inserted into the model, and the resulting system is solved using numerical integration to simulate state variable evolution under experimental conditions. Model performance is assessed by comparing predicted and experimental concentration profiles, with better-fitting models guiding the evolutionary search process. Through repeated refinement, this approach produces a set of candidate models that balance accuracy and complexity, with the final selection made using the AICc criterion to identify the most plausible representation of the underlying system dynamics.

Approach 2 - Gradient space

The second framework identifies kinetic model structure by initially treating key lumped parameters as time-varying quantities, estimated directly from experimental data alongside constant parameters within a nonlinear ODE-based model. This is achieved through an optimisation problem that minimises the error between model predictions and experimental observations while incorporating a regularisation term to ensure smooth, physically realistic parameter evolution over time.

Once these time-resolved parameter profiles are obtained across all experiments, they are used as target data for SR, which searches for analytical expressions that relate the parameters to the system's state variables. Candidate expressions are evaluated based on how well they reproduce the estimated parameter trajectories, and an iterative evolutionary process generates a set of models representing different accuracy–complexity trade-offs. Finally, the AICc criterion is applied to select the most plausible expressions, ensuring an optimal balance between model fit and simplicity.

CASE STUDY

To evaluate whether SR could recover accurate kinetic information (i.e. ground-truth kinetic model structure), a case study was constructed using a yeast-based fermentation process for astaxanthin production. Using a kinetic model developed in our previous work [7], three in-silico experiments were generated at different initial conditions of biomass and substrate concentration to act as experimental data. The ground truth kinetic model is proposed in Equation 1.

$$\frac{dX}{dt} = \frac{\mu_{m1} S_1 X}{(S_1 + K_{c1} X)\cdot(1 + k_1 S_2)} + \frac{\mu_{m2} S_2 X}{(S_2 + K_{c2} X)\cdot(1 + k_2 S_1)} - \mu_d X \quad (1a)$$

$$\frac{dS_1}{dt} = \frac{-Y_{S1} \mu_{m1} S_1 X}{(S_1 + K_{c1} X)\cdot(1 + k_1 S_2)} \quad (1b)$$

$$\frac{dS_2}{dt} = \frac{-Y_{S2} \mu_{m2} S_2 X}{(S_2 + K_{c2} X)\cdot(1 + k_2 S_1)} \quad (1c)$$

$$\frac{dP}{dt} = \frac{\alpha_1 \mu_{m1} S_1 X}{(S_1 + K_{c1} X)\cdot(1 + k_1 S_2)} + \frac{\alpha_2 \mu_{m2} S_2 X}{(S_2 + K_{c2} X)\cdot(1 + k_2 S_1)} + \beta X - k_d X^2 \quad (1d)$$

where X, S_1, S_2, and P denote concentrations of biomass, two substrates, and astaxanthin, respectively.

To effectively discover the kinetic information of

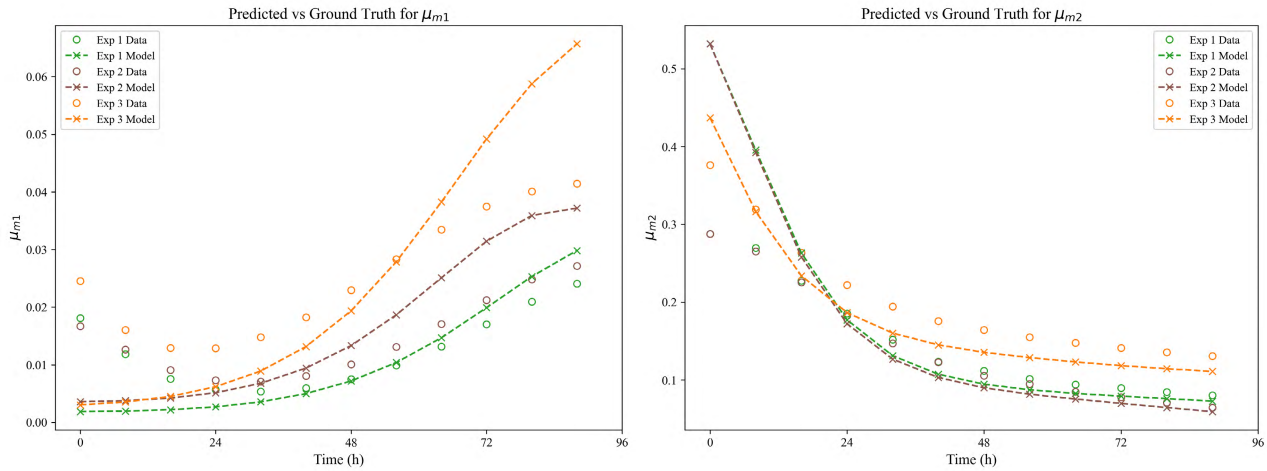

Figure 2. Comparison of the trends of μ_{m1} (a) and μ_{m2} (b) over time. Colours green, brown and orange represent experiments 1, 2, and 3 respectively.

this in-silico case study, a typical Monod model shown in Equation 2 was initially used to simulate a scenario where no prior strain-specific physical knowledge was available. SR was therefore applied to recover ground truth knowledge by automatically modifying and improving the Monod model structure and retrieve the expressions of two key kinetic parameters μ_{m1} and μ_{m2} in Equation 2.

$$\frac{dX}{dt} = \frac{\mu_{m1} S_1 X}{S_1 + K_{s1}} + \frac{\mu_{m2} S_2 X}{S_2 + K_{s2}} - \mu_d X \tag{2a}$$

$$\frac{dS_1}{dt} = \frac{-Y_{S1} \mu_{m1} S_1 X}{S_1 + K_{s1}} \tag{2b}$$

$$\frac{dS_2}{dt} = \frac{-Y_{S2} \mu_{m2} S_2 X}{S_2 + K_{s2}} \tag{2c}$$

$$\frac{dP}{dt} = \frac{Y_{P1} \mu_{m1} S_1 X}{S_1 + K_{s1}} + \frac{Y_{P2} \mu_{m2} S_2 X}{S_2 + K_{s2}} + \beta X - k_d X^2 \tag{2d}$$

Using this case study, if the expressions identified by SR can transform the kinetic backbone into the ground truth kinetic model by identifying the true expressions for μ_{m1} and μ_{m2} as shown in equation 3, it can be said that the proposed framework has the potential to systematically uncover novel relationships for emerging bioprocesses where prior mechanistic understanding is incomplete.

$$\mu_{m1,goal} = \frac{\mu_{m1} \cdot (S_1 + k_{S1})}{(S_1 + K_{c1} X) \cdot (1 + k_1 S_2)} \tag{3a}$$

$$\mu_{m2,goal} = \frac{\mu_{m2} \cdot (S_2 + k_{S2})}{(S_2 + K_{c2} X) \cdot (1 + k_2 S_1)} \tag{3b}$$

RESULTS AND DISCUSSION

Results of Approach 1

The results, presented in figure 1, showed that the generated models were able to accurately simulate state variable concentrations for each experiment. This was reflected in a low mean absolute percentage error (MAPE) value of 8.9% averaged across each experiment and each of the four state profiles (biomass, sucrose, glucose and astaxanthin concentration).

Despite success in simulating the concentration profiles, the framework of approach 1 was unable to uncover the exact ground truth model structures. It was discovered that, even though SR was tasked with discovering only two lumped terms, the potential solution space remained extensive and many models that had a strong fitting toward the experimental data exhibited significant structural divergence from the ground truth nevertheless. Further analysis deemed that the discrepancy stemmed from the failure to reproduce the temporal evolution of the specific growth rate terms μ_{m1} and μ_{m2} themselves because subtle differences in the gradient space can be diminished after integration, leading to different gradient profiles producing similar integral profiles. Figure 2 displays the profiles of μ_{m1} and μ_{m2} over time and highlights the discrepancy between the ground truth and predicted profiles from SR models.

Results of Approach 2

The results, presented in figure 3, showed that the models generated using Approach 2 were also able to accurately simulate the state variable concentrations dynamics and reduce the model error (MAPE) by 60% with an average value of 3.5%.

Importantly, the fitting in the gradient space depicting the profiles of μ_{m1} and μ_{m2} through time, shown in figure 4, experienced a significant improvement. the models generated by SR were able to almost exactly reproduce the ground truth profiles with MAPE values of 0.08% and 0.39% for μ_{m1} and μ_{m2} respectively, averaged across all three experiments. Ultimately this led to the recovery of the ground truth expressions for both parameters. It should be noted that, for simple model structures such as μ_{m2}, SR was able to identify the correct physical structure

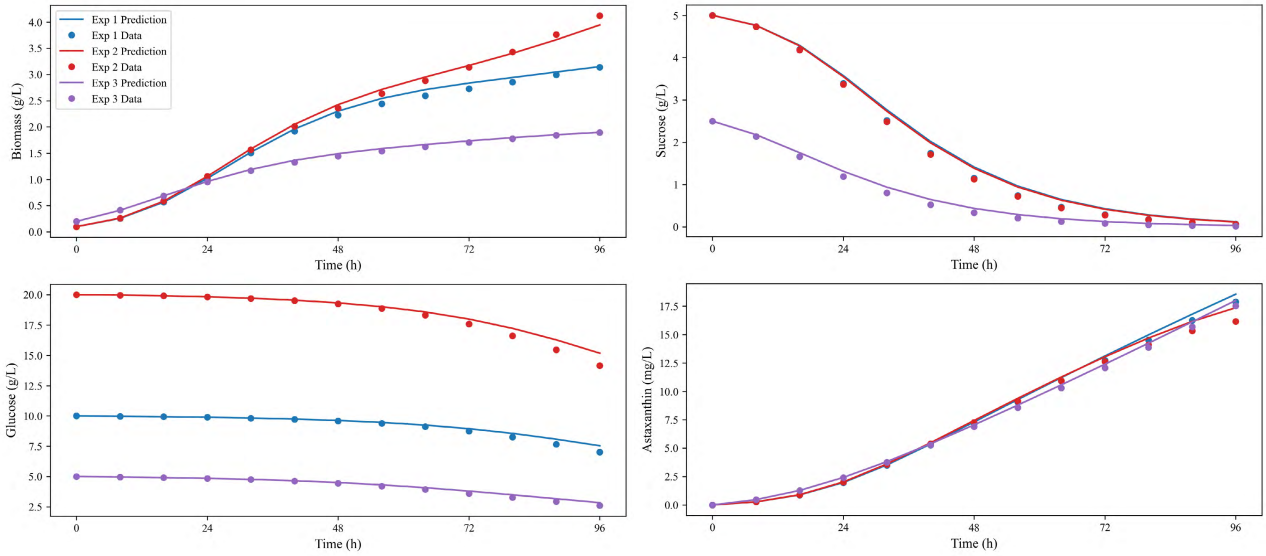

Figure 3. Fitting performance of the complete kinetic model using expressions discovered by SR for μ_{m1} and μ_{m2}. Colours blue, red and purple represent experiment 1, 2, and 3 respectively. (a), (b), (c) and (d) shows the fitting result for concentrations of biomass, sucrose, glucose and astaxanthin, respectively.

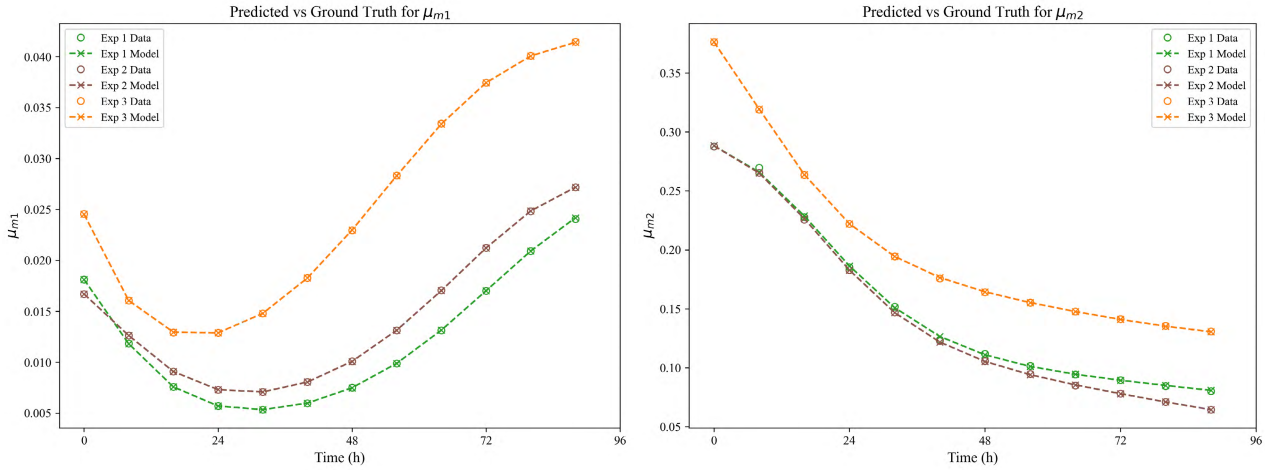

Figure 4. Comparison of the trends of μ_{m1} (a) and μ_{m2} (b) over time. Colours green, brown and orange represent experiments 1, 2, and 3 respectively.

without the need of running additional experiments. However, for complex model structures such as μ_{m1}, as evidenced by their evolution over time, SR was unable to find the ground truth expressions on the first iteration due to the large solution space of possible expressions. Nonetheless, by iteratively freezing and mutating specific terms within the best expression selected using AICc, the SR generated expressions were able to converge to the ground truth structure.

CONCLUSION

This study presents a symbolic regression–based framework for automated kinetic model structure identification that aims to improve mechanistic understanding of bioprocess systems while maintaining interpretability. By integrating symbolic regression with a general kinetic model backbone, two complementary strategies were explored, revealing that while the concentration-based approach could accurately reproduce system dynamics, it struggled to distinguish between structurally different models. In contrast, the gradient-based approach proved more effective and computationally efficient, successfully recovering underlying kinetic expressions when combined with AICc-based model selection and targeted refinement strategies.

Although the framework shows strong potential for data-efficient and interpretable model discovery, limitations remain, including the lack of validation under noisy data, untested performance on more complex systems, and challenges associated with limited experimental data in real-world scenarios. Despite this, the framework remains valuable as a knowledge discovery tool, capable of revealing meaningful relationships that can guide hypothesis generation even when exact model identification is not possible.

A key strength of the approach is its flexibility for expert involvement, aligning with augmented intelligence principles. Domain knowledge can be incorporated to guide model development, constrain the search space, and improve model plausibility, while quantitative tools such as information criteria can assess the impact of these interventions. Future work will focus on integrating such human-in-the-loop strategies, alongside methods like sparse regression, to further enhance the efficiency and interpretability of bioprocess model discovery.

AUTHOR IDENTIFIERS

Author ORCIDs:
Riezzo L: 0009-0007-3309-628X
Rogers A: 0000-0003-2298-6520
Kay H: 0009-0005-9565-5687
Zhang D: 0000-0001-5956-4618

REFERENCES

1. Park S, Kim S, Park C, Kim J, Lee D. Data?driven prediction models for forecasting multistep ahead profiles of mammalian cell culture toward bioprocess digital twins. Biotech & Bioengineering 120:2494-2508 (2023). https://doi.org/10.1002/bit.28405
2. Pennington O, Ríos SE, Sebastian MT, Dickson A, Zhang D. Dynamic multiscale hybrid modelling of a CHO cell system for recombinant protein production. IFAC-PapersOnLine 58:133-138 (2024). https://doi.org/10.1016/j.ifacol.2024.08.326
3. Riezzo L, Kay H, Feng Y, Jing K, Zhang D. Accelerating bioprocess digital twin development by integrating hybrid modelling with transfer learning. Chemical Engineering Journal 511:162018 (2025). https://doi.org/10.1016/j.cej.2025.162018
4. Brunton, S.L., Proctor, J.L. and Kutz, J.N., “Discovering governing equations from data by sparse identification of nonlinear dynamical systems”, *Proceedings of the National Academy of Sciences of the United States of America*, 113(15), pp. 3932–3937, (2016) https://doi.org/10.48550/arXiv.1509.03580
5. Jul-Rasmussen P, Chakraborty A, Venkatasubramanian V, Liang X, Huusom JK. Hybrid AI modeling techniques for pilot scale bubble column aeration: a comparative study. Computers & Chemical Engineering 185:108655 (2024). https://doi.org/10.1016/j.compchemeng.2024.108655
6. Forster T, Vázquez D, Müller C, Guillén-Gosálbez G. Machine learning uncovers analytical kinetic models of bioprocesses. Chemical Engineering Science 300:120606 (2024). https://doi.org/10.1016/j.ces.2024.120606
7. Vega?Ramon F, Zhu X, Savage TR, Petsagkourakis P, Jing K, Zhang D. Kinetic and hybrid modeling for yeast astaxanthin production under uncertainty. Biotech & Bioengineering 118:4854-4866 (2021). https://doi.org/10.1002/bit.27950

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Modelling Pressure Effects in Boiling Brazed Aluminum Heat Exchangers: A Software Comparison

Hamza Karim[a], Rim Khodr[a], Rodolphe Sardeing[b] and Gaëtan Becker[a]*

[a] Alfa Laval Golbey, 25 bis rue du Fort, 88190 Golbey, France
[b] Fives ProSim, 51 rue Ampère, Immeuble Stratège A, 31670 Labège, France
* Corresponding Author: gaetan.becker@alfalaval.com.

ABSTRACT

Brazed aluminum heat exchangers (BAHX) are key cryogenic equipment, but their simulation is sensitive to phase change modelling. This work benchmarks ProSec, CO-ProSec Reaction, and Aspen EDR on an ethylene-chiller BAHX using ethane in thermosiphon mode. ProSec's interpolation scheme of tabulated data is validated against CO-ProSec Reaction's approach that employs a thermodynamic model for the evaluation of thermodynamic and physical properties; an initial 25% pressure drop gap is traced to different void-fraction models (slip vs homogeneous). Against Aspen EDR, default duty/quality are ~15% higher; differences are mainly due to nucleate-boiling treatment and distributor pressure drop modelling. Harmonizing options reduces duty mismatch to ~5%.

Keywords: Energy, Modelling and Simulations, Cryogenics, Brazed Aluminum Heat Exchanger

INTRODUCTION

Brazed Aluminum heat Exchangers (BAHX) are one of the most widely used equipment in various industries. They are a core technology in air separation units, industrial gas processing, natural gas treatment and liquefaction (LNG), petrochemical and refinery services.

BAHX are valued for their high thermal efficiency, low weight, compact footprint and limited routine maintenance requirements. Their outstanding performance derives from an exceptionally large heat-transfer area per unit volume: typically area densities are on the order of 1000-1500 $m^2.m^{-3}$, which is much higher than conventional technologies, such as shell and tubes heat exchangers (40 to 70 $m^2.m^{-3}$). In addition, a single core can accommodate up to twenty streams enabling complex

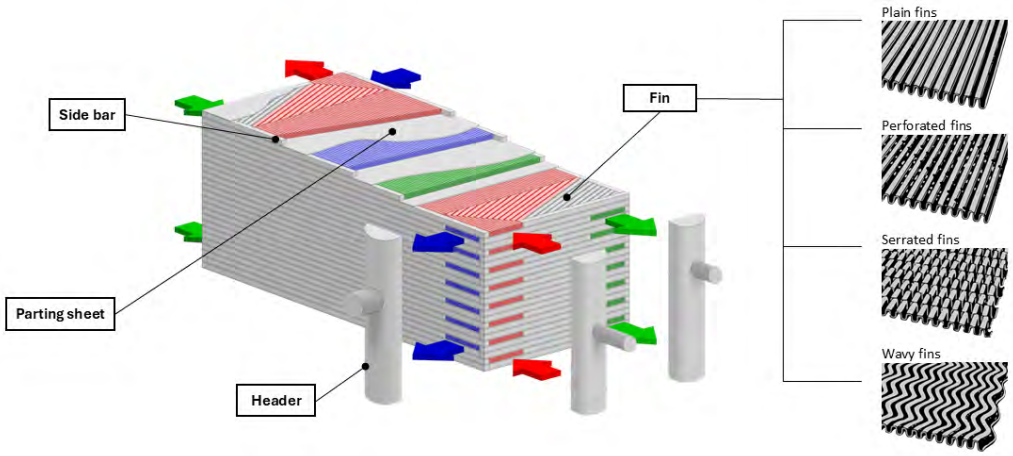

Figure 1. (left) Structure of a Brazed Aluminum Heat Exchanger, (right) Types of fins encountered in BAHX

thermal integration within one exchange and very small temperature approaches(2–5 °C) [1–4].

The geometry of a BAHX core is defined in detail by ALPEMA standards [2]. As schematically shown in **Figure 1**, a BAHX consists of a plate-fin core to which headers are attached. The headers distribute the process streams into the core, where the heat exchange takes place.

The core is built as a stack of passages, ranging from a few tens to more than one hundred layers. Each passage is dedicated to a given stream and comprises: (i) inlet distribution fins (distributors) to spread the incoming flow across the core width, (ii) heat-transfer fin sections where the fluid flows through the exchanger and exchanges heat with adjacent passages above and below, and (iii) side bars (closure bars) that seal the passage and contain the flow as shown in **Figure 2**. Adjacent passages are separated by thin parting sheets, across which passage-to-passage heat transfer occurs

To tailor the heat-transfer performances and pressure-drop behavior, several fin typologies are used beyond the standard straight fins, including perforated fins, serrated fins, and wavy fins.

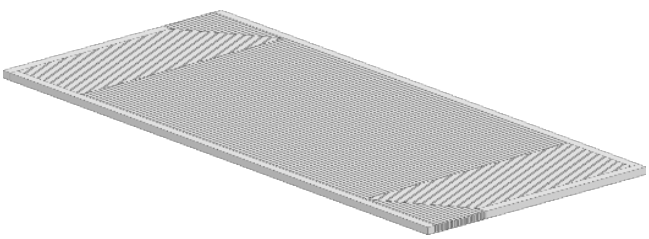

Figure 2. Schematic representation of a single passage

This construction makes BAHX intrinsically multiscale, structurally complex, and strongly heterogeneous, features that underpin their superior performance but also complicate modelling.

Indeed, the same features that make BAHX attractive also make their thermal prediction non-trivial. Strong stream-to-stream coupling, the influence of fin efficiency, axial conduction through parting sheets and side bars, sensitivity to maldistribution and phase change can all affect predicted temperature profiles, pinch locations, and duty. Consequently, selecting an appropriate simulation tool and understanding the modelling assumptions it embeds is essential for credible performance prediction and for robust design and troubleshooting.

To this end, we present a benchmark study comparing three different simulation tools on a representative case with a boiling stream. The first, ProSec, is an in-house software developed at Alfa Laval Golbey in collaboration with Fives ProSim, drawing on more than 30 years of research and industrial know-how. Based on robust numerical methods, ProSec simulate BAHX and analyze local profiles of numerous physical parameters (temperature, pressure, heat transfer coefficient, etc.) and allows the user to fine tune the design of the equipment to reach the best possible performances. It also accounts for pressure effects on thermodynamic properties through a bi-pressure interpolation scheme. The second, CO-ProSec Reaction (version 3.7.8), is a commercial version distributed by Fives ProSim [5]; it evaluates properties directly from thermodynamic models, offering a fully continuous treatment. The third tool, Plate-Fin in Aspen Exchanger Design & Rating (EDR) (version 14) distributed by AspenTech [6], adopts alternative modelling strategies together with its own heat-transfer and pressure-drop correlations . Due to the large diffusion of AspenTech's software in the industrial community, EDR is often the commercial software chosen to simulate heat exchangers. Therefore, comparisons with EDR is a necessary step for validating any development.

The remainder of this paper is organized as follows. Section 1 describes the pressure interpolation method implemented in ProSec. Section 2 introduces the case study used for the software comparison. Section 3 compares ProSec and CO-ProSec Reaction: both tools share the same correlations and numerical resolution framework, but CO-ProSec Reaction evaluates thermodynamic and transport properties from a thermodynamic model via Simulis Thermodynamics through a CAPE-OPEN interface, whereas ProSec relies on tabulated data. This comparison therefore isolates and validates the interpolation strategy. Section 4 presents a critical comparison between ProSec and Aspen EDR, with the objective of identifying the origin of the observed differences and critically discuss them.

PRESSURE INTERPOLATION METHOD

In phase change operation, pressure variations directly govern the saturation temperature and affect vapor quality and enthalpy, thereby influencing the local thermal driving forces. This is important when a thermosiphon stream is present. In a thermosiphon loop, the circulation rate is not imposed by a pump; it results from the balance between the hydrostatic head in the downcomer and the pressure losses across the loop components (downcomer line, BAHX core, and return line).

Figure 3 shows a typical temperature profile for a thermosiphon stream. The fluid enters the BAHX at approximately the saturation temperature at the surface of the liquid in the column. However, due to the hydrostatic pressure in the downcomer, the inlet pressure is higher and the liquid is typically subcooled with respect to the local saturation temperature. The temperature then increases (while static pressure decreases) until the local boiling point is reached, which depends on the pressure profile along the passage. Beyond this point, an evaporation zone develops and the thermodynamic state follows the saturation line as vaporization proceeds [7].

Accurate prediction of pressure drop and pressure-dependent thermodynamic properties is therefore required to reproduce the local thermal profile. This is particularly relevant for BAHX: while aluminum alloys generally retain high tensile strength at cryogenic temperatures compared with carbon steels, they can be relatively susceptible to fatigue damage under cyclic thermal loading [8]. Reliable local temperature profiles are thus important to assess whether significant thermal stresses may occur.

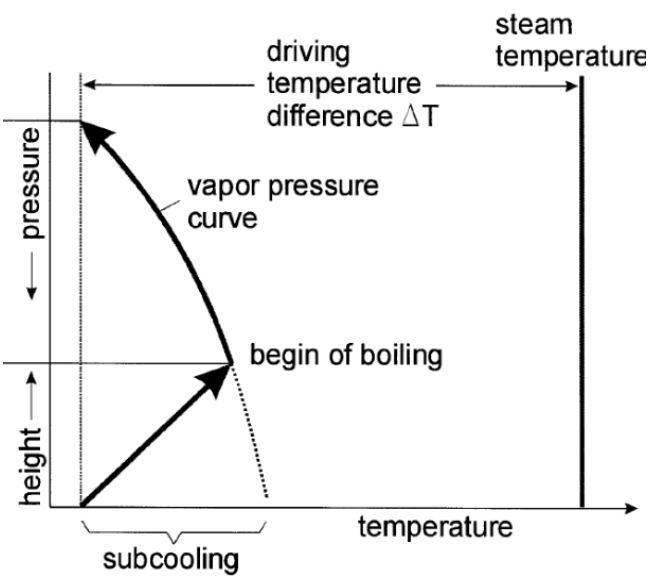

Figure 3. Characterisitic temperature profile for a thermosiphon stream [7]

In ProSec, the user provides for each stream tabulated values of temperature, T, vapor quality, x and enthalpy, h at two pressure levels P_1 and P_2. The program then evaluates the temperature and vapor quality at a specified enthalpy h and pressure P (with $P_1 \leq P \leq P_2$).

This is illustrated in **Figure 4**. In this example, the user provides two enthalpy tables, (T, h_1) and (T, h_2), corresponding to pressures P_1 and P_2. The objective is to determine the temperature $T(h, P)$. First, the interpolation factor r is computed on a logarithmic pressure basis:

$$r = \frac{\ln(P/P_1)}{\ln(P_2/P_1)} \tag{1}$$

The bubble-point and dew-point enthalpies (h_b, h_d) at pressure P are then obtained by interpolation from the corresponding values at P_1 and P_2:

$$h_b = (1-r) \times h_{b_1} + r \times h_{b_2} \tag{2}$$

$$h_d = (1-r) \times h_{d_1} + r \times h_{d_2} \tag{3}$$

Next, the enthalpy h is compared to h_b and h_d to identify the fluid region (liquid, vapor – liquid, or vapor). In the example shown in the figure, $h < h_b$ indicates that the fluid is in the liquid region at pressure P.

Two auxiliary enthalpies h_m and h_n are then defined at P_1 and P_2, respectively, such that the enthalpy shift between them is consistent with the corresponding saturation boundary (i.e, $h_n - h_m = h_{b_2} - h_{b_1}$ in the liquid region). Depending on the region, h_m and h_n are defined as follows:

liquid region:

$$h_m = h - (h_b - h_{b_1}) \tag{4}$$

$$h_n = h + (h_{b_2} - h_b) \tag{5}$$

vapor region:

$$h_m = h - (h_d - h_{d_1}) \tag{6}$$

$$h_n = h + (h_{d_2} - h_d) \tag{7}$$

vapor-liquid region:

$$h_m = h - (1-y)(h_b - h_{b_1}) - y(h_d - h_{d_1}) \tag{8}$$

$$h_n = h + (1-y)(h_{b_2} - h_b) + y(h_{d_2} - h_d) \tag{9}$$

$$y = \frac{h - h_b}{h_d - h_b} \tag{10}$$

Once h_m and h_n are known, the corresponding temperatures $T(h_m, P_1)$ and $T(h_n, P_2)$ are obtained by cubic spline interpolation of the tabulated data at P_1 and P_2. Finally, the temperature $T(h, P)$ is computed as:

$$T(h, P) = (1-r)T(h_m, P_1) + rT(h_n, P_2) \tag{11}$$

The same interpolation approach is used for the vapor quality x. This is only required when the fluid lies in the vapor–liquid region. In the liquid and vapor regions, the vapor quality is simply 0 and 1, respectively.

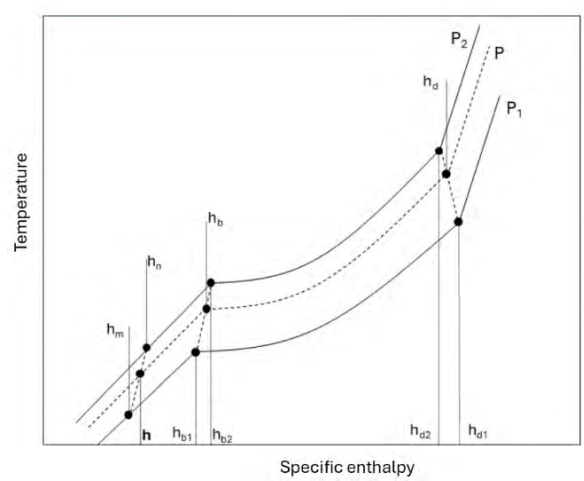

Figure 4: Schematic representation of the interpolation of the temperature between two sets of enthalpic profiles at P_1 and P_2 for a given enthalpy value, h and pressure P in the case of a diphasic mixture

CASE STUDY: ETHYLENE PRODUCTION

Ethylene is a major building block of the chemical industry. It is widely used as a feedstock for plastics, fibers, and other organic chemicals that support key sectors such as packaging, transportation, and construction. Indeed, global ethylene consumption exceeded 150 million tonnes per year in 2017 [9].

Today, ethylene is produced mainly by steam cracking of ethane or naphtha. The overall process is commonly divided into three sections: (i) pyrolysis, where the feed is cracked in furnaces to form light olefins and the effluent is rapidly quenched to limit secondary reactions; (ii) compression and primary treatment, where the cracked gas is cooled, compressed in multiple stages, and purified (notably removal of water and acid gases) prior to cryogenic processing; and (iii) cryogenic product recovery, where chilling trains and distillation towers separate hydrogen/methane and C_2 components, and ultimately split ethylene from ethane, typically at very low temperatures [10].

The case study considered in this paper belongs to the cryogenic recovery section and focuses on ethylene chilling, a duty where BAHX technology is particularly well suited. In the investigated exchanger, an ethylene stream is cooled using ethane as a refrigerant operating in thermosiphon mode. The inlet conditions and operating parameters of both streams are summarized in **Table 1** Thermodynamic properties are evaluated using a Peng–Robinson equation of state coupled with the Soave Boston–Mathias alpha function.

Details of the BAHX internal geometry (core dimensions, stacking arrangement, and fin selection) are proprietary to Alfa Laval Golbey and are therefore not reported here. The exchanger uses in-house fins developed from extensive industrial experience. The corresponding hydraulic and thermal performances - expressed through the Fanning friction factor (C_f) and Colburn factor (C_j) - were characterized experimentally using an internal C_f–C_j test bench, and these measured correlations are used as inputs to the simulations.

Table 1: Details of the streams input data used in the simulation of the BAHX

Parameter	HOT (ethylene)	COLD (ethane)
$\dot{M}$ (kg.h^{-1})	157000	209000
T_{in} (°C)	-24.6	-54.1
P_{in} (bar)	10.44	4.012
Flow Direction	Downward	Upward

COMPARAISON OF PROSEC WITH CO-PROSEC REACTION

Table 2 compares the predicted exchanger performances obtained with ProSec and CO-ProSec Reaction for the thermosiphon stream. Overall, the agreement is very good: the outlet temperatures are virtually identical, and the outlet vapor fraction and exchanged duty differ by about 3%. In contrast, the predicted pressure drop shows a noticeable deviation, up to 25%. This difference is unexpected, since both tools are intended to use the same pressure drop correlations

Table 2: Global performances comparison between ProSec and CO-ProSec Reaction for the thermosiphon stream

Parameter	ProSec	CO-ProSec Reaction
T_{out} (°C)	-53.29	-53.25
x_{out} (kg.kg^{-1})	0.565	0.547
Q (kW)	14490.4	14036.5
DP (mbar)	133.4	106.8

A plausible explanation is the different treatment of properties in the two tools. In ProSec, the influence of pressure on thermodynamic properties (temperature, quality, enthalpy) is accounted for, whereas transport properties (density, viscosity, thermal conductivity, and heat capacity) are interpolated only with respect to temperature. On the other hand, in CO-ProSec Reaction, properties are evaluated directly from the thermodynamic model at each mesh point using the local temperature and pressure. If pressure significantly affects transport properties - particularly density - then the converged pressure drop profile may differ, even when the same correlations are used.

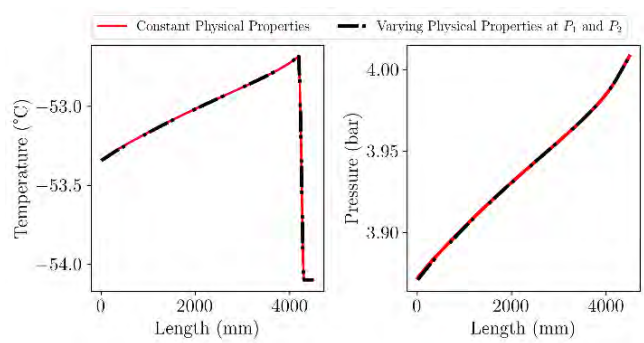

Figure 5: Influence of physical properties variations with respect to pressure on the BAHX performance

To test this hypothesis, an additional simulation was carried out in Aspen EDR, which allows interpolation of both thermodynamic and transport properties with respect to temperature and pressure (the user provides property tables at two pressure levels, and the program performs the interpolation)

Figure 5 compares two Aspen EDR cases. The black

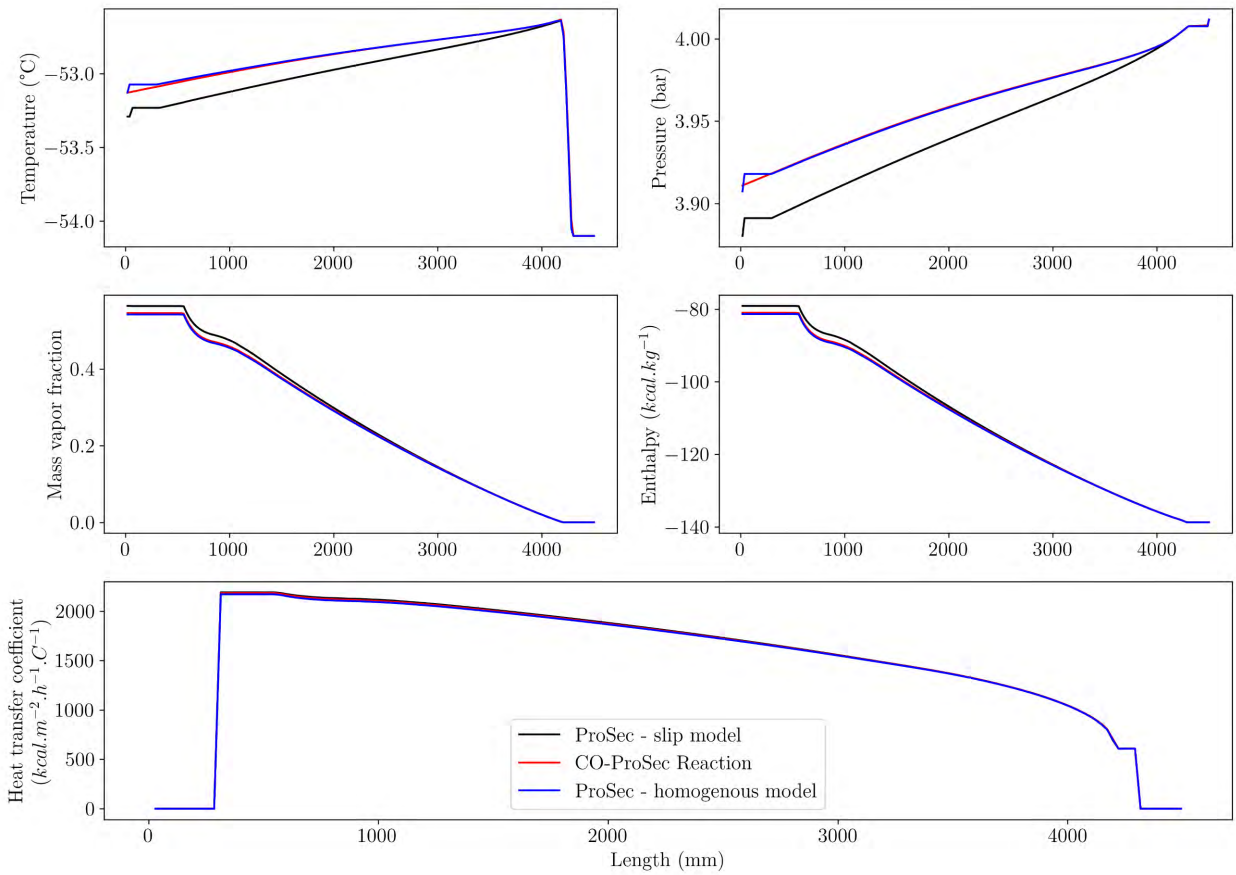

Figure 6: Comparison of the different profiles along the length of the BAHX between ProSec with slip model, CO-ProSec Reaction and ProSec with homogeneous model for the thermosiphon stream

curves correspond to identical transport properties at both pressure levels (mimicking a temperature-only dependence, as in ProSec). The red curves correspond to pressure-dependent transport properties (mimicking the CO-ProSec Reaction approach). In this case, the temperature and pressure profiles overlap, indicating that accounting for pressure dependence of transport properties has no measurable impact on the results. This is consistent with the small overall pressure drop (≈0.1 bar), for which pressure-induced variations of transport properties are expected to be negligible.

These findings indicate that the deviation must originate from another modelling difference between the two tools. To investigate further, the pressure profile was recomputed manually by separating the frictional and static contributions and applying the same frictional and gravitational formulations for both tools. The reconstructed profile matched the ProSec simulation closely, but not the CO-ProSec Reaction results. This pointed to a difference in the treatment of the two-phase static pressure drop.

The key parameter is the void fraction, defined as the volumetric fraction of vapor in the two-phase mixture. Void fraction is required to evaluate two-phase mixture density, and therefore directly affects the gravitational pressure drop [11, 12]. Numerous void-fraction models exist in the literature; they are commonly grouped into homogeneous (no-slip), slip-ratio based, and drift-flux approaches, each relying on different assumptions [12].

In ProSec, void fraction is computed using a slip-based correlation from the HTFS Handbook [13]. In contrast, CO-ProSec Reaction (version used in this study) uses a homogeneous model for void fraction in the computation of the static pressure drop. **Figure 6** compares the profiles along the exchanger length obtained with CO-ProSec Reaction (red curve) and with ProSec using the HTFS correlation (black curve). The pressure profiles show significant differences, which in turn affect the temperature, vapor quality, and enthalpy profiles through the coupling between pressure and saturation conditions in the two-phase region.

When ProSec was configured to use the homogeneous model (blue curve) for void fraction, its results became identical to those of CO-ProSec Reaction as shown in **Figure 6**. This provides strong confidence in the pressure interpolation approach implemented in ProSec: once

the two tools share the same void-fraction treatment, the simulations are consistent. It is also worth noting that since version 3.7.11, CO-ProSec Reaction implements the same void-fraction correlation as ProSec, which is a necessary step to ensure consistent modelling assumptions between the two tools.

This comparison highlights the importance of void-fraction modelling for two-phase flow in BAHX passages. Channel hydraulic diameters typically range from ~1 to 5 mm, whereas many widely used correlations were developed for larger, conventional channels. While two-phase flow in mini-channels has attracted increasing research attention [14, 15], further work is still needed -particularly studies that connect fundamental understanding to industrial geometries and operating conditions - to improve predictive capability in BAHX applications.

COMPARAISON OF PROSEC WITH ASPEN EDR

As noted earlier, comparison with Aspen EDR is essential because it is widely used in industry; benchmarking against it therefore provides a relevant reference point for assessing any simulation approach. Aspen EDR employs its own numerical solution strategy as well as its own set of heat-transfer and pressure-drop correlations.

The first comparison was performed using the default settings in both ProSec and Aspen EDR. This choice is motivated by two practical reasons. First, in typical workflows, Aspen EDR users often rely on the default options for thermal and hydraulic calculations rather than tuning correlation sets or numerical controls. Second, quantifying agreement under default configurations provides a clear baseline from which the dominant sources of divergence can be identified and then reduced in a controlled manner.

As shown in **Table 3**, the default simulations exhibit non-negligible deviations. Aspen EDR predicts a heat duty and outlet vapor quality that are approximately 15% higher than those predicted by ProSec. Aspen EDR also predicts a higher total pressure drop. These differences indicate that the default modelling choices (heat transfer and pressure drop correlations) have a measurable impact on the predicted performance.

To interpret these discrepancies, axial profiles of the cold (boiling) stream - temperature, pressure, vapor quality, enthalpy, and heat-transfer coefficient - were extracted from both simulations and are compared in **Figure 7**. The differences are most apparent in the heat-transfer coefficient: Aspen EDR predicts higher values than ProSec, which directly drives higher local heat flux, a larger enthalpy rise, and therefore a higher outlet vapor quality and duty. This behavior is due to the default correlation choices for the heat transfer coefficient: Aspen EDR uses an HTFS boiling correlation, whereas ProSec uses an in-house boiling correlation (confidential).

Table 3: Comparison of the results of global performances between the different configurations of ProSec and Aspen EDR. See text for the meaning of different case studies

	T_{out} (°C)	x_{out} (kg.kg^{-1})	Q (kW)	DP (mbar)
ProSec - default	-53.29	0.565	14490	133.4
EDR - default	-53.37	0.650	16612	145.8
ProSec – HTFS85	-53.29	0.544	13946	133.3
EDR – NNB	-53.36	0.580	14968	144.9
ProSec - DP	-53.36	0.554	14212	142.3

Differences are also observed in the pressure-drop prediction. Aspen EDR yields a larger pressure drop, which is mainly attributed to the treatment of the distributor pressure drop (see **Figure 7** pressure profile plot for the different zones inside the BAHX). Aspen EDR relies on HTFS handbook methods for distributor pressure drop, while ProSec uses correlations and practices derived from Alfa Laval Golbey's internal experience. A closer look at the pressure-drop breakdown in Aspen EDR indicates that it includes a static contribution in the distributors, whereas this contribution is not accounted for in ProSec. In the Alfa Laval workflow, a more detailed and comprehensive distributor pressure-drop evaluation (including static effects) is performed using a separate proprietary tool, which explains the difference in scope between the two ProSec and Aspen EDR calculations.

ProSec also offers the option to use the HTFS correlation for evaluating the heat-transfer coefficient. In practice, this option is mainly applied to condensing streams, as internal experience indicates that the in-house correlation is better suited to boiling flows. For consistency with Aspen EDR, the HTFS correlation was therefore applied to all streams in ProSec.

Interestingly, this change did not reduce the discrepancy between the two tools but increased it. Under this configuration, Aspen EDR predicts a duty about 20% higher than ProSec, as shown in **Table 3** (ProSec – HTFS85). **Figure 7** explains this trend: in ProSec, switching from the default correlation to HTFS85 reduces the predicted boiling heat-transfer coefficient, which lowers the duty and enlarges the gap with Aspen EDR.

Further examination of Aspen EDR documentation gave more information about the default correlation. In Aspen EDR, the default boiling model includes an additional nucleate boiling contribution to the overall heat transfer coefficient, whereas the HTFS correlation used in ProSec accounts essentially for the convective component only. In boiling flow inside channels, heat transfer is commonly described as the combination of convective

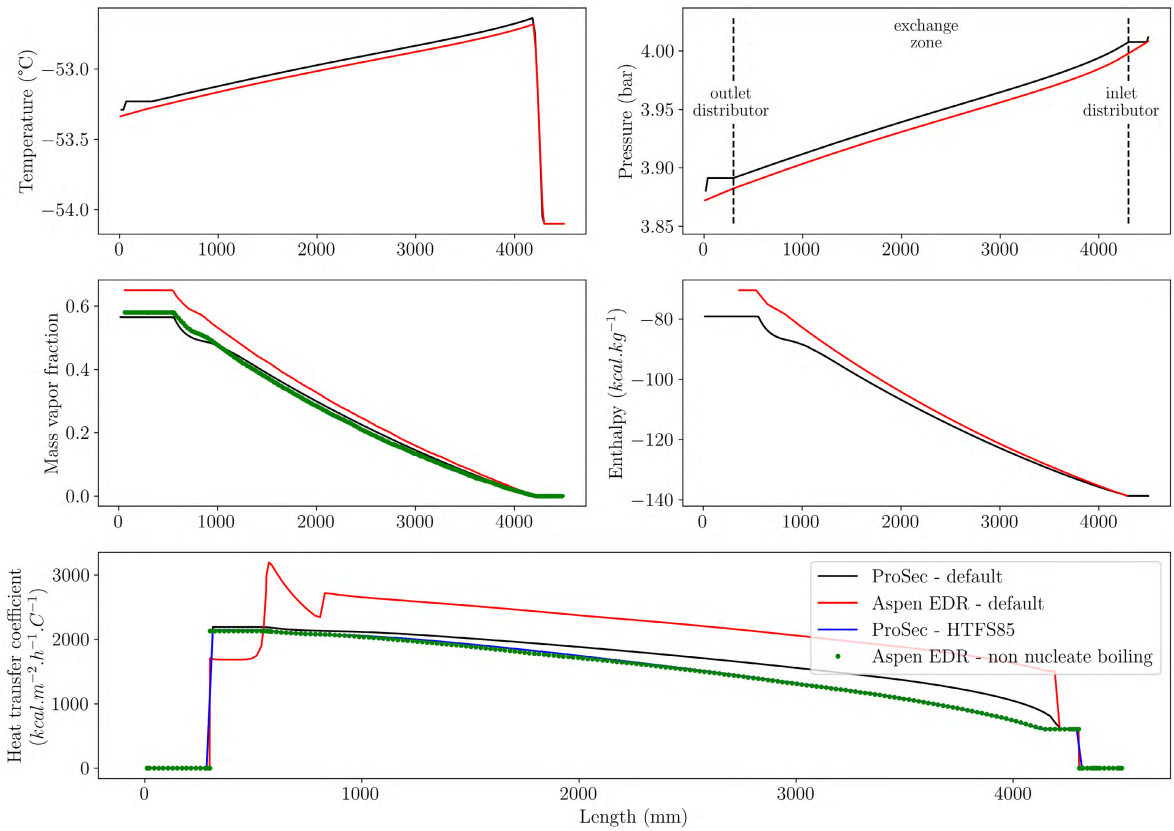

Figure 7. Comparison of the different profiles along the length of the BAHX between ProSec and Aspen EDR under their default settings for the thermosiphon stream. "PoSec – HTFS85" represents Prosec when using the HTFS correlation, "Aspen EDR – non nucleate boiling" is when the option of nucleate boiling is deactivated in Aspen EDR.

boiling and nucleate boiling, with their relative importance depending on heat flux, wall superheat, flow regime, and vapor quality. Nucleate boiling tends to be most influential at low to moderate qualities and sufficiently high wall superheat, while convective boiling becomes increasingly dominant as quality increases. For further information, interested readers are encouraged to read the following references [16–19].

Aspen EDR provides an option to disable the nucleate boiling contribution and retain only the convective component. The resulting case (EDR–NNB) is also reported in Table 3. When EDR–NNB is compared to ProSec – HTFS85, the duty difference decreases substantially (from ~20% to ~7%). This is also visible in Figure 7, where the heat transfer coefficient profiles of EDR – NNB and ProSec – HTFS85 nearly overlap.

To further reduce differences, the ProSec inlet pressure was adjusted so that the pressure at the beginning of the exchange zone matches that in the Aspen EDR simulation (case ProSec – DP). This alignment decreases the remaining gap, yielding a duty difference of about 5% (Table 3). Nevertheless, even when comparable pressure levels and cold-side boiling heat transfer coefficients are obtained, Aspen EDR still predicts a slightly higher duty.

This remaining difference is attributed to the treatment of the condensing hot stream, as shown in Figure 8. The two simulations predict similar heat transfer coefficients over most of the condensation range, except near the onset of condensation. In ProSec, the heat transfer coefficient is linearized between its value at a quality of 0.9 and the single-phase value to avoid numerical instabilities. In Aspen EDR, the two-phase correlation is retained up to the transition to single-phase flow. This reflects different modelling philosophies: ProSec prioritizes numerical robustness in the vicinity of phase-boundary transitions, where certain correlations can become ill-conditioned, whereas Aspen EDR maintains the two-phase correlation closer to the single-phase limit.

Overall, this alignment of modelling options shows that most of the initial discrepancies between ProSec and Aspen EDR can be traced to identifiable assumptions. Once these elements are harmonized, the two tools provide comparable trends and global performances.

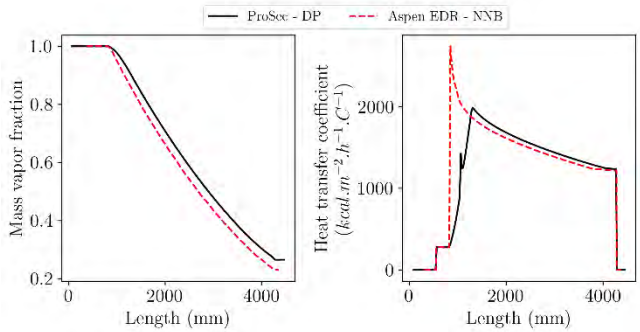

Figure 8: Mass vapor fraction and heat transfer coefficient profiles of the hot stream for both ProSec – DP and Aspen EDR – NNB simulations.

CONCLUSION

This paper presented a benchmark of three BAHX simulation tools (ProSec, CO-ProSec Reaction and Aspen EDR) on a thermosiphon boiling case. ProSec's pressure-interpolation approach was validated against CO-ProSec Reaction; the initial 25% pressure-drop discrepancy was traced to different void-fraction models for the static two-phase contribution (slip model in ProSec versus homogeneous in CO-ProSec Reaction) and disappeared once the same model was used.

Comparison with Aspen EDR showed larger default deviations (≈15% higher duty and outlet quality, and higher pressure drop), mainly due to differences in boiling heat transfer modelling (nucleate contribution included by default in EDR) and in the treatment of distributor pressure drop. When nucleate boiling was disabled in EDR and pressures were aligned at the exchange-zone inlet, agreement improved to within ~5%, with remaining differences linked to condensation modelling near phase-boundary transitions.

Overall, the results emphasize that comparisons between BAHX simulation tools are only meaningful when the main physical assumptions are made consistent. More broadly, the work highlights the need for continued refinement and validation of two-phase correlations in BAHX passages with millimetric hydraulic diameters, where correlation choice can materially affect predicted performance and engineering conclusions.

REFERENCES

1. Picard F (2008) Modélisation et simulation thermo-hydrauliques en trois dimensions des échangeurs multifluides à plaques brasées. Institut National Polytechnique (Toulouse)
2. Thome JR. The new 3rd edition of the ALPEMA plate-fin heat exchanger standards. Heat Transfer Engineering 31:1-2 (2010). https://doi.org/10.1080/01457630903263176
3. PRASAD BSV. The performance prediction of multistream plate-fin heat exchangers based on stacking pattern. Heat Transfer Engineering 12:58-70 (1991). https://doi.org/10.1080/01457639108939764
4. Averous D (2000) Modélisation et simulation des échangeurs de chaleur multifluides à plaques brasées en configurations complexes. Toulouse, INPT
5. ProSec software | Fives ProSim. https://www.prosim.net/en/product/prosec-simulation-of-brazed-plate-fin-heat-exchangers/. Accessed 14 Jan 2026
6. Aspen Exchanger Design and Rating (EDR) | AspenTech. https://www.aspentech.com/en/products/engineering/aspen-exchanger-design-and-rating. Accessed 6 Jan 2026
7. Arneth S, Stichlmair J. Characteristics of thermosiphon reboilers. International Journal of Thermal Sciences 40:385-391 (2001). https://doi.org/10.1016/s1290-0729(01)01231-5
8. R.D S, A.C M (2011) Aluminum Plate-Fin Heat Exchanger Thermal Fatigue Study. AIChE
9. Gao Y, Neal L, Ding D, Wu W, Baroi C, Gaffney AM, Li F. Recent advances in intensified ethylene production—a review. ACS Catal. 9:8592-8621 (2019). https://doi.org/10.1021/acscatal.9b02922
10. REN T, PATEL M, BLOK K. Olefins from conventional and heavy feedstocks: energy use in steam cracking and alternative processes. Energy 31:425-451 (2006). https://doi.org/10.1016/j.energy.2005.04.001
11. Cai C, Mudawar I, Liu H, Xi X. Assessment of void fraction models and correlations for subcooled boiling in vertical upflow in a circular tube. International Journal of Heat and Mass Transfer 171:121060 (2021). https://doi.org/10.1016/j.ijheatmasstransfer.2021.121060
12. Márquez-Torres L, Ochoa-Pineda J, Pico P, Valdés JP, Becerra D, Pinilla A, Pereyra E, Ratkovich N. Comparison of 63 different void fraction correlations for different flow patterns, pipe inclinations, and liquid viscosities. SN Appl. Sci. 2: (2020). https://doi.org/10.1007/s42452-020-03464-w
13. R.A. S (1989) Two phase pressure drop in the straight channels of plate-fin heat exchangers. HTFS Handb.
14. Pietrzak M, P?aczek M. Void fraction predictive methods in two-phase flow across a small diameter channel. International Journal of Multiphase Flow 121:103115 (2019). https://doi.org/10.1016/j.ijmultiphaseflow.2019.103115

15. Han J, Liu Y, Chu W, Zhao C, Bo H. Experimental study on visualized flow boiling in a narrow rectangular channel. International Communications in Heat and Mass Transfer 138:106383 (2022). https://doi.org/10.1016/j.icheatmasstransfer.2022.106383
16. Kreith F. CRC handbook of thermal engineering. CRC Press (1999). https://doi.org/10.1201/9781420050424
17. Agostini B (2002) Etude expérimentale de l'ébullition de fluide réfrigérant en convection forcée dans des mini-canaux. Université Joseph Fourier (Grenoble; 1971-2015)
18. Desrats C (2006) Etude expérimentale de l'ébullition d'hydrocarbures dans les échangeurs à plaques et ailettes: analyse des écoulements et des transferts de chaleur par une approche locale. Aix-Marseille 1
19. Feldman A (1996) De l'ébullition en convection forcée dans des canaux d'échangeurs à plaques et ailettes. Université Henri Poincaré-Nancy 1

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

From Drift to Adaptation to the failed ML model: Transfer Learning in Industrial MLOps

Waqar Muhammad Ashraf[ab], Talha Ansar[c], Fahad Ahmed[d], Jawad Hussain[d], Muhammad Mujtaba Abbas[c], and Vivek Dua[a]

[a] The Sargent Centre for Process Systems Engineering, Department of Chemical Engineering, University College London, Torrington Place, London WC1E 7JE, UK
[b] The Alan Turing Institute, British Library, 96 Euston Road, London NW1 2DB, UK
[c] Department of Mechanical, Mechatronics and Manufacturing Engineering (New Campus), University of Engineering and Technology (UET), Lahore 54000, Pakistan
[d] Shandong Huatai Electric Operation & Maintenance (Private) Limited, Sahiwal Coal Fired Power Plant, Sahiwal, 57000, Pakistan
* Corresponding Author: v.dua@ucl.ac.uk

ABSTRACT

Model adaptation to production environment is critical for reliable Machine Learning Operations (MLOps), less attention is paid to developing systematic framework for updating the ML models when they fail under drift. This paper compares the transfer learning enabled model update strategies including ensemble transfer learning (ETL), all-layers transfer learning (ALTL), and last-layer transfer learning (LLTL) for updating the failed feedforward artificial neural network (ANN) model. The flue gas differential pressure across the air pre-heater unit installed in a 660 MW thermal power plant is analyzed as a case study since it mimics the batch processes due to load cycling in the power plant. Updating the failed ANN model by three transfer learning techniques reveals that ETL provides relatively higher predictive accuracy for the batch size of 5 days than those of LLTL and ALTL. However, ALTL is found to be suitable for effective update of the model trained on large batch size (8 days). A mixed trend is observed for computational requirement (hyperparameter tuning and model training) of model update techniques for different batch sizes. These empiric insights obtained from the batch process-based industrial case study can assist the MLOps practitioners in adapting the failed models to drifts for the improved monitoring of industrial processes.

Keywords: model failure, data drift, concept drift, industrial processes, MLOps

INTRODUCTION

The industrial operations have nonlinear characteristics, meaning that change in the process conditions over time and equipment changeover may result in changing patterns of correlations between the process variables. The “static” data-driven models trained on “fixed data-distributions” may not handle non-stationary data streams of industrial processes and are not suitable for their deployment for long-term monitoring tasks. This is attributed to data drift (shift in covariance due to change in data distributions between the input variables of process) and/or concept drift (the underlying relationships between the input and output variables of process are changed) that drive the model failure.

The data-driven models are retrained on the new data streams in order to maintain their predictive accuracy for new data-distributions. The potential caveat with this approach is the “catastrophic forgetting”, meaning that the model parameters have been tuned to the new dataset and may not predict well the historical data the model was trained on [1]. Initially, Geoffrey E. Hinton hypothesized that small changes in model parameters for their tuning to new data may collectively cancel out their effect on historical data [2]. However, McCloskey and Cohen [3] and Ratcliff [4] showed that sequential training of simple neural networks on disjoint sets of data results in significant forgetting even when the models are updated on small amount of new data.

To address the potential issue of catastrophic

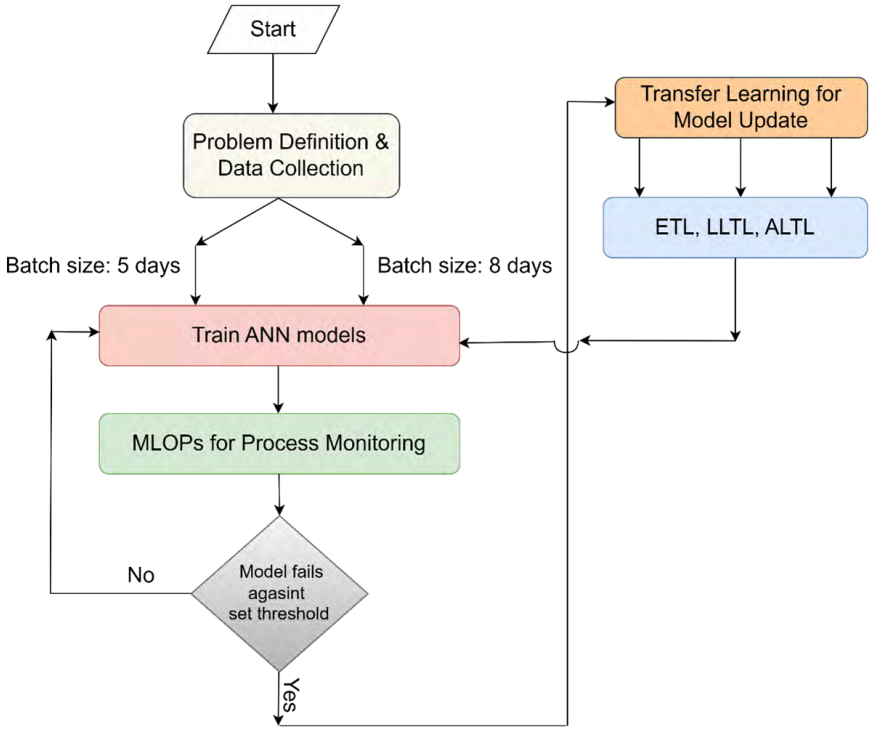

Figure 1. Schematic on transfer learning-based model update framework for process monitoring.

forgetting, continual learning approach is used in updating the model parameters and it balances the trade-off between the plasticity (learn new dynamics) and stability (retaining historical information) through experience replay and regularization techniques during the model update mechanism [5]. On the other hand, online learning retrains the model on a single or batch of data continuously once the model is productionised [6]. Online learning offer high adaptability to the model to the incoming data stream or concept drift; however, it suffers from catastrophic forgetting. Just-In-Time-Learning is another model update paradigm where a local model is trained on the incoming data stream and similar data queried from the historical data [7]. Just-In-Time-Learning offers high accuracy to approximate nonlinear dynamics of system. However, it has higher computational and memory requirement and may not extrapolate when system enters in the new operating model where historical data is not available. To enable the high predictive accuracy and reducing the computational expense for the models in production, transfer learning approach is also implemented where parameters of trained model (source domain maintained on historical data) are transferred for adapting them with the "target" domain (drift in data) through minimal requirement of data [8].

The literature studies identify the importance of updating the model for their adaptability to the changing industrial operations. In [9], the neural network parameters in all layers are tuned upon the detection of drift and resulting updated model was deployed for digital twin applications. In another study [10], the researchers have allocated the epoch budget and bounds on learning rate for updating the model parameters completely. Similarly, transfer learning is also implemented for model adaptation to source domain [11]. However, it is noted that much of research attention is paid to update the model parameters, the effect of batch-size on updating the model parameters by transfer learning-based model update techniques is not studied systematically. More importantly, it remains largely unknown how transfer learning driven model update techniques including ensemble transfer learning (ETL), all-layers transfer learning (ALTL), and last-layer transfer learning (LLTL) for neural networks in combination with batch size and how it would affect the predictive accuracy of the model. These fundamental issues about the selection of model update techniques considering plasticity vs stability and the computational expense are not investigated systematically to obtain mechanistic insights about addressing these issues.

Motivated by the current research gaps, this paper investigates the effect of batch size and transfer learning driven model update techniques (ETL, ALTL, LLTL) to address the concept drift for updating the feedforward ANN model. The flue gas differential pressure (DP) across the air preheater installed in a 660 MW thermal power plant is considered as a case study. Flue gas DP is correlated

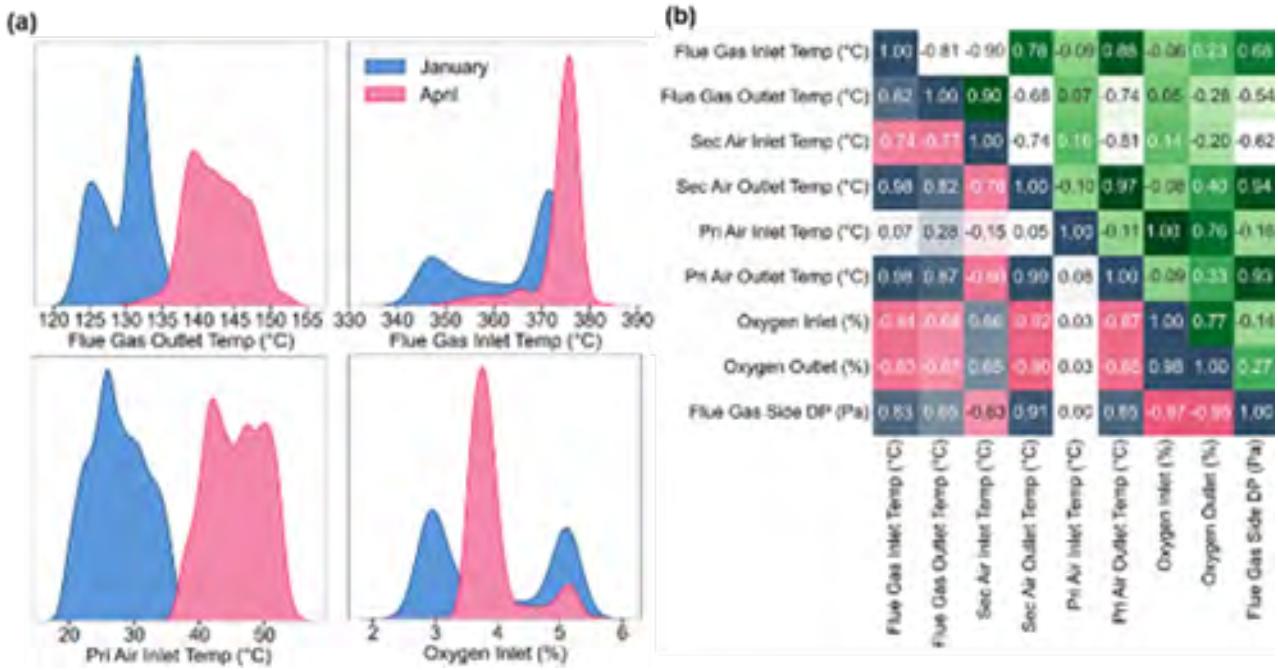

Figure 2. Characteristics of data collected for heat exchange across air preheater. (a) Data drift is observable in data collected for month of January and April. (b) Correlation structure between the operating variables is also changed for the two months (lower diagonal: January, upper diagonal: April) and data & concept drift is present in the collected data.

with load cycling (ramp-up and ramp down) and is treated as a batch process which is observable from the data visualizations. This allows to transfer the findings from this case study to the dynamics of batch processes which are carried out in various chemical engineering applications. A systematic comparison is made to reveal the effect of batch size on the predictive accuracy of updated model by three techniques, changes in the weight space of the updated models as well as their computational requirement. Moreover, the evolution in the feature importance is also investigated by SHapley Additive Explanations (SHAP) when the models are updated to reveal their stability for industrial monitoring tasks. These fundamental and empiric insights obtained through the industrial case study can assist the MLOps practitioners in adapting the failed models to drifts for the improved monitoring of industrial processes.

METHOD

Figure 1 shows the transfer learning enabled model update framework developed in this work to monitor the flue gas DP across the air preheater. Two months of data is collected separately for the month of January and April for the model development on the two batch sizes of 5 days and 8 days. Initially, a set of two batches (for 5 days and 8 days each) is used to train ANN models under hyperparameters optimization (learning rate, number of neurons in the hidden layer). The models are evaluated on coefficient of determination (R^2) and root-mean-squared-error (RMSE) on the test data (20% of set). Later, the trained models are deployed as MLOps to predict the Flue Gas DP for the rest of batch data. When daily average error (RMSE and MAE) remains above the twice of test data (observed during model training) for three times consecutively, model update mechanism is triggered. A set of the batch data where model fails and the consecutive previous batch data (buffer data for model retraining) is used to update the model by transfer learning driven approaches (ETL, LLTL and ALTL). Later, the models trained by the three approaches are put into MLOps again to predict the rest of the collected data. This allows to compare the predictive accuracy of the model against the ground truth already available in the collected data. The stability of the models in terms of evolution of feature importance for the trained and updated models under different batches and model update approaches are investigated by SHAP framework. The feature evolution is affected by the process dynamics; however, they can be corrupted by the ineffective model update and mapping created in the model. The computational time consumed to train and update the models are also recorded to investigate the suitability of transfer learning enabled model update approaches for MLOps.

RESULTS: MONITORING DIFFERENTIAL PRESSURE ACROSS AIR PREHEATER

Data Characteristics

The operating variables across the air preheater for modelling flue gas DP (Pa) are as follows: flue gas inlet temperature (°C), flue gas outlet temperature (°C), secondary air inlet temperature (°C), secondary air outlet temperature (°C), primary air inlet temperature (°C), primary air outlet temperature (°C), and oxygen inlet (%), oxygen outlet (%). The historical operation data of the variables is collected at the averaging interval of 10 minutes and 8 minutes, respectively, for the month of January and April in order to explore the capacity and efficacy of models to handle different sampling frequency of data collection for the monitoring task. This allows us to mimic the real-life situation when sampling frequency of data collection is affected by data storage and maintenance costs and in turn investigate how the model responds to this transition yet ensuring acceptable predictive accuracy.

Figure 2(a) presents the data-distribution profiles of a few operating variables for the months of January and April. The asymmetric data-distribution as well as data drift are visually observable for the variables, and are further quantified by population stability index (PSI) and Cramér-von Mises (CvM). PSI ranges from 1.89 to 23.15 and CvM statistics from 72.4 to 658.9 with p-values less than 0.001 confirm drift in the data observed in two datasets sampled at multi time-scale sampling. Primary air inlet temperature exhibits the strongest data drift across both metrics (PSI=23.15, CvM=658.9) and flue gas DP (PSI=1.89) and oxygen outlet (CvM=72.4) show the relatively smaller data drift, yet noticeable for potential model failure. The correlation between the pair of variables (input-input & input-output) is computed for the two months separately (shown on Figure 2(b)) and reveals that correlation between the pair of variables is changed in the two datasets (data and concept drift are present in the data). This indicates different operating modes of operation are observed for the air preheater and this is useful for the monitoring evaluation capacity of the trained model once deployed in production.

Training of ANN model and update under two different batch sizes

Feedforward ANN models are trained on a set of two batches of 5 days and 8 days separately, taken from the data of January. These two batch sizes are considered based on the operating characteristics of air pre-heater as influenced by power generation mode and soot deposition in the heat exchanger plates. During the training of model, hyperparameters of the ANN model (hidden layer neurons, learning size, number of hidden layers) are optimized and models are evaluated on the test data (20% of initial data). The model trained on set of 5-day batches yielded a test R^2 of 0.9824 (RMSE 0.0381), while the model on set of 8-day batches achieved a test R^2 of 0.9807 (RMSE 0.0385). After the training phase, the models trained on batch size of 5 days and 8 days are deployed in production to monitor flue gas DP across the air preheater. The monitoring accuracy of the two models is depicted on Figure 3(a, b) respectively. It is noted that the ANN model trained on batch size of 5 days exhibited relatively superior performance in predicting the Flue Gas DP in the month of January than those of ANN model trained on batch size of 8 days.

Between April 7 and 10, the two models crossed the human-defined error threshold, i.e., crossing twice the test error limit successively for three times. The mapping of actual and model-based responses (reference) is shown in the bottom of Figure 3(a, b). The failure of models on different sampling frequency of data collection reveals that the trained models do not scale well when data sampling frequency changes with respect to those the model was trained on. The failure of model can also be attributed to process dynamics and drift (data and concept) as depicted on Figure 2. The model failure triggers the model update procedure built on three techniques including ETL, ALTL, and LLTL and we carried it out for batch size of 5 days and 8 days separately (batch data where model fails and the consecutive previous batch). It is important to note here that during model update, only learning rate is optimized iteratively in the narrow range and without changing the model architecture. Later, the models trained on three model update techniques are put into MLOps to monitor Flue Gas DP across the air preheater for the rest of data of April.

The mapping of actual and model-based responses obtained from models trained by ETL, ALTL and LLTL is depicted on Figure 3(c, d) for the batch size of 5 days and 8 days respectively. Similarly, parity plots are also constructed to directly compare the actual and model-based responses as shown on Figure 3(e, f). The RMSE and MAE for ETL are as follows: 0.0564, 0.0414 (5-day), 0.0652, 0.0541 (8-day); ALTL: 0.0674, 0.0563 (5-day), 0.0384, 0.0315(8-day); and LLTL 0.0400, 0.0321 (5-day), 0.0537, 0.0438 (8-day). The visualization of actual and model-based responses in conjunction with errors computed for the models trained on three update techniques allows us to inspect local predictive performance of the trained model once productionized in industrial settings.

ETL approach seems to monitor flue gas DP with relatively more accuracy in comparison with ALTL and LLTL for the model trained on batch size of 5 days (Figure 3(a)). On the other hand, ALTL is found to have comparatively better predictive performance for the model trained on batch size of 8 days (Figure 3(b)). For small

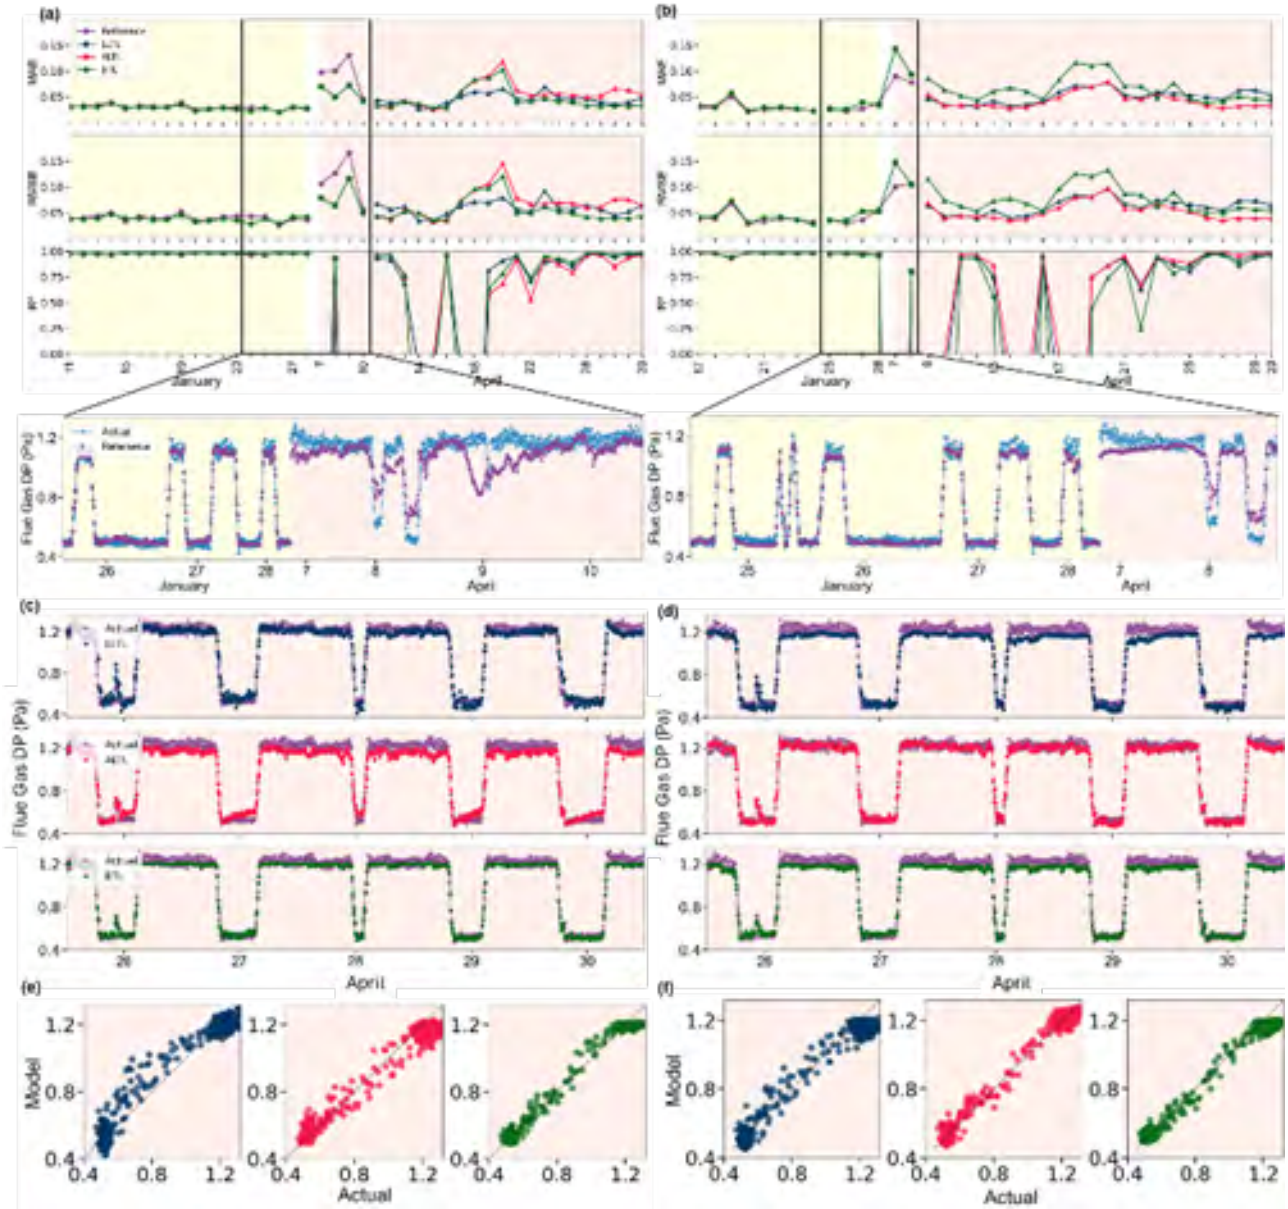

Figure 3. MLOps of ANN models trained on batch sizes of (a) 5 days and (b) 8 days to monitor flue gas DP across air preheater. The trained ANN model exhibited good monitoring performance in the month of January. However, they failed during 7-10 April and model update techniques are implemented. Actual and model-based responses are also mapped corresponding to the failure zone. Peformance of models updated by LLTL, ALTL, ETL techniques on remaining April data is shown in (c, d). (e, f) parity plot to compare actual and model-based response for three model update techniques in the remaining april data.

batch sizes, ensemble may average out the predictive responses and provide stable performance. However, when batch size increases, ALTL updates the model parameters across all layers and model is adapted with the recent data stream. This increases the risk of "forgetting" the operation history and model is poised to adapt with the recent operation dynamics.

Insights about layer-wise weight space distribution after model update

The layer-wise weight space of the models updated with three techniques is retrieved and is visualised to obtain insights on how model parameters are tuned or varied after the model update phase corresponding to batch size of 5 days and 8 days. The weight-space distribution

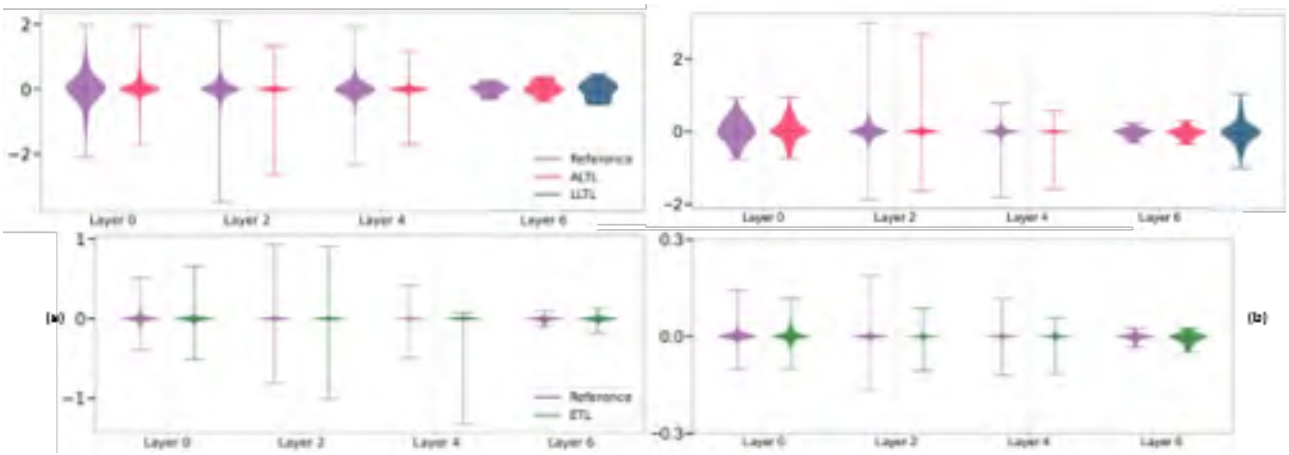

Figure 4. Weight space distribution changes with model update for the batch size of (a) 5 days, and (b) 8 days. In common, last-layer weight space is widened to adapt with data-drift.

is visualized as violin plot for each layer of ANN and is depicted on Figure 4(a, b) corresponding to batch size of 5 days and 8 days respectively. In general, it is found that weight space distribution for last-layer of ANN gets widened for the models trained on batch sizes of 5 days and 8 days. This can be attributed to back-propagation of error signal starting from last layer of ANN that tunes the model parameters. For ALTL, the weight space of hidden layers seems to be compressed than those of reference distribution (pre-trained model space) that can explain the improved generalization capability of the model while adapting with drift observed in the data distribution collected for the month of April. However, ensembles approach seems to depict mixed results in widening or compressing the weight space that can be explained by a number of ensembles digesting the data based on the batch sizes.

Evolution of feature importance

The feature importance to predict flue gas DP is established by SHAP and is used to investigate the evolution in feature importance after the model update. The reference column represents the feature importance before the model update that serves as a baseline to compare the feature importance order after the model update. Secondary air outlet temperature is found to be the most significant feature affecting flue gas DP for models updated on batch size of 5 days and 8 days, respectively, as shown on Figure 5(a, b). After the model update, secondary air outlet temperature remains the most significant feature as established for models updated by LLTL, ALTL and ETL approaches. Multiple ETL columns explicitly track the sequential updates across successive batch intervals, where each column corresponds to a distinct version produced after an update and the resulting feature importance order. The consistent significant order of features, especially for most significant operating variable, by three model update techniques depicts the accuracy of feature importance analysis as well as stability of the trained models to capture the process dynamics. It also reveals that feature importance may remain the same even when concept and data drift are observed in the data (refer to Figure 2).

Although the magnitude of feature importance is varied for each model update approaches, the feature importance order remains nearly the same, especially for significant variables. This may depict the stable tuning of model parameters, irrespective of batch-sizes, and is also evident as stable error profiles depicted on Figure 3(a, b).

Computational Performance of model update techniques

The computational time required to train a reference model as well as its update by three techniques for the batch size of 5 days and 8 days is recorded and is presented on Figure 6(a, b). For a reference model, comparatively larger computational time is observed (hyperparameter tuning time: 1113 s, training time: 28.4 s) for model trained on batch size of 8 days than those of model trained on 5 days (hyperparameter tuning time: 868 s, training time: 14.4 s). It can be attributed to non-smooth convergence to error minimum on the fewer batches. However, models trained on batch size of 5 days, in general, consumed large computational time for model update by LLTL, ALTL and ETL approaches than those of models trained on batch size of 8 days. This observation may reflect that stochastic error convergence in the baseline model training produces the parametric values that adapt quickly on new data streams as compared to smoothly converged models. A mixed trend is observed in comparing the computational time used for updating

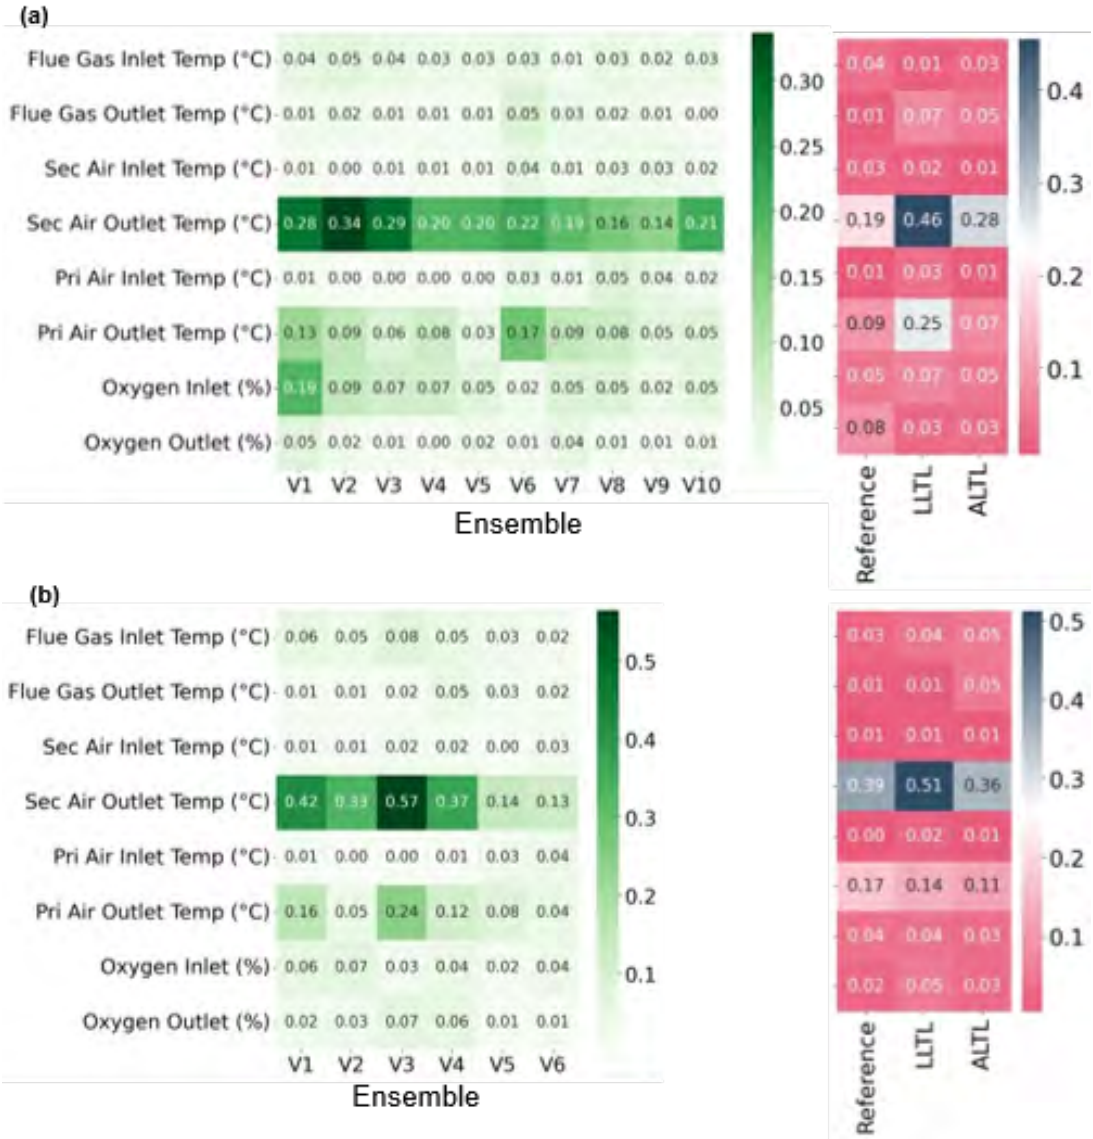

Figure 5. Evolution of SHAP-based feature importance with model update for ETL, LLTL and ALTL with respect to batch size of (a) 5 days and (b) 8 days. While Secondary Air Outlet Temperature (°C) is consistently the most dominant feature, but during periods of information scarcity, Primary Air Outlet Temperature (°C) exhibits a transient increase in importance to compensate for the lack of information of the new regime (V6 in (a) and V3 in (b)). Once sufficient data is accumulated, the model stabilizes, and Sec Air Outlet Temperature (°C) regains its primary significance.

the models for the same batch size. However, hyperparameter tuning time remains larger than those of training time irrespective of batch sizes.

CONCLUSION

We provide insights into the selection of model update techniques when feedforward ANN model fails in MLOps due to drift (data and concept). The effect of batch size (5 days or 8 days) on the predictive accuracy of feed forward ANN model is investigated to monitor flue gas DP across air preheater (treated as batch process) installed in a 660 MW thermal power plant. We find that model failure is driven by drift and ensemble approach implemented for batch size of 5 days (small batch size) provides superior monitoring performance than those of LLTL and ALTL techniques. This can be attributed to widening of weight space across the layers of ANN models in the ensemble which is otherwise compressed for LLTL and ALTL techniques. The comparison of computational time requirements (hyperparameter tuning and model training) for model update techniques reveals that ETL

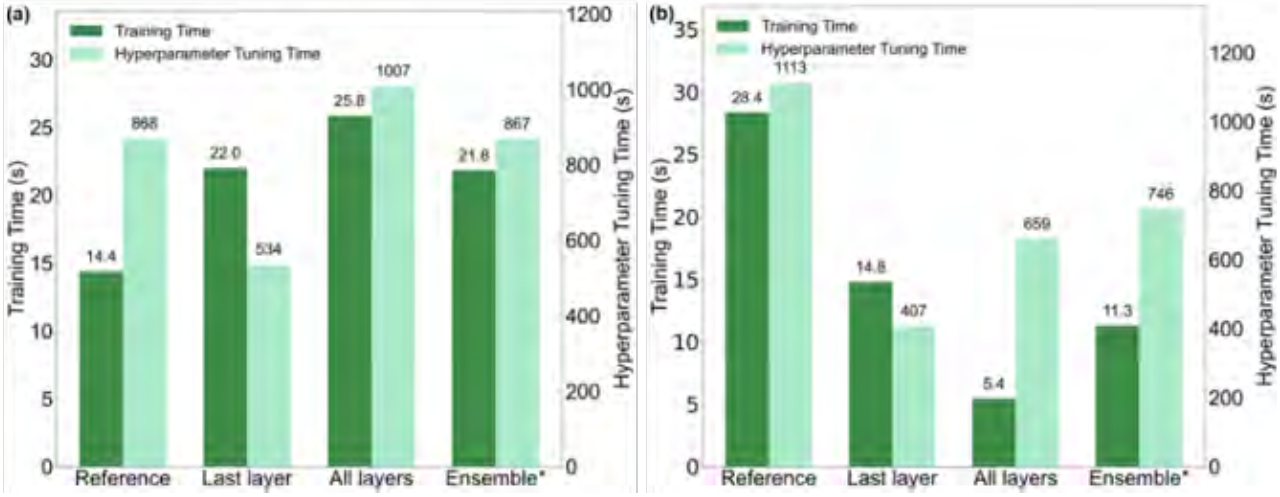

Figure 6. Comparing the computational time consumed for hyperparameter tuning and model training for batch size of (a) 5 days and (b) 8 days.

takes longer to update the model for the batch size of 5 days. However, ALTL, which depicts superior monitoring performance for the large batch size, takes the lowest time for updating the model once hyperparameters have been optimised. This aspect may have wider implications for MLOps practices since the models having stochastic error convergence profiles may quickly adapt with drift (data and concept) than the models which initially have smooth error convergence profiles. These fundamental but empiric insights obtained from model update for industrial batch processes highlight the role of batch size of data and the selection of model update techniques for industrial monitoring applications. These findings can also help troubleshoot the MLOps pipeline and dealing with failed ML model for adapting it to volatile industrial production environments.

REFERENCES

1. Kadlec P, Grbić R, Gabrys B. Review of adaptation mechanisms for data-driven soft sensors. Computers & Chemical Engineering 35:1-24 (2011). https://doi.org/10.1016/j.compchemeng.2010.07.034
2. Hinton, G.E., Distributed representations. 1984.
3. McCloskey M, Cohen NJ. Catastrophic interference in connectionist networks: the sequential learning problem. Psychology of Learning and Motivation :109-165 (1989). https://doi.org/10.1016/s0079-7421(08)60536-8
4. Ratcliff R. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions.. Psychological Review 97:285-308 (1990). https://doi.org/10.1037/0033-295x.97.2.285
5. Shah SRB, Schwung A. Continual learning by gradient monitoring for remaining useful lifetime estimation. Knowledge-Based Systems 323:113721 (2025). https://doi.org/10.1016/j.knosys.2025.113721
6. Prapas I, Derakhshan B, Mahdiraji AR, Markl V. Continuous training and deployment of deep learning models. Datenbank Spektrum 21:203-212 (2021). https://doi.org/10.1007/s13222-021-00386-8
7. Ge Z, Song Z. A comparative study of just-in-time-learning based methods for online soft sensor modeling. Chemometrics and Intelligent Laboratory Systems 104:306-317 (2010). https://doi.org/10.1016/j.chemolab.2010.09.008
8. Jia R, Zhang S, You F. Transfer learning for end-product quality prediction of batch processes using domain-adaption joint-y PLS. Computers & Chemical Engineering 140:106943 (2020). https://doi.org/10.1016/j.compchemeng.2020.106943
9. Rebello CM, Nogueira IBR. Digital twins in chemical engineering: an integrated framework for identification, implementation, online learning, and uncertainty assessment. Computers & Chemical Engineering 200:109178 (2025). https://doi.org/10.1016/j.compchemeng.2025.109178
10. Metcalfe B, Acosta-Pavas JC, Robles-Rodriguez CE, Georgakilas GK, Dalamagas T, Aceves-Lara CA, Daboussi F, Koehorst JJ, Corrales DC. Towards a machine learning operations (mlops) soft sensor for real-time predictions in industrial-scale fed-batch fermentation. Computers & Chemical

Engineering 194:108991 (2025). https://doi.org/10.1016/j.compchemeng.2024.108991

11. Theisen MF, Meesters GMH, Schweidtmann AM. Transferring graph neural networks for soft sensor modeling using process topologies. Systems and Control Transactions 4:40-45 (2025). https://doi.org/10.69997/sct.185977

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Design of a Chemical Heat Pump based on Methylcyclohexane, Toluene and Hydrogen

Rajalakshmi Krishnadoss[a], Félix Le Bot[b], and Thomas A. Adams II[a]*

[a] Norwegian University of Science and Technology, Department of Energy and Process Engineering, Trondheim, Norway
[b] ENSTA Paris, Department of Energy, Transports and Space, Paris, France
* Corresponding Author: thomas.a.adams@ntnu.no.

ABSTRACT

The conceptual design and performance of a novel Methylcyclohexane-Toluene-Hydrogen based chemical heat pump was studied using steady state simulations. The distillation operating parameters of the chemical heat pump were optimized to maximize the Coefficient of Performance based on heat quantity (COP) and its corresponding Coefficient of Performance based on electric work input (COP_W) was calculated. The best operating temperature ranges of the endothermic and exothermic reactor are 200°C-225°C and 250°C-275°C respectively. An endothermic temperature of 200°C and an exothermic temperature of 250°C results in a COP of 0.1357 and a COP_W of 13.3. By integrating this chemical heat pump with a vapor compression heat pump COP increased to 0.1445 while COP_W reduced to 4.9.

Keywords: Chemical heat pump, Energy Efficiency, Toluene, Methylcyclohexane, Hydrogen

INTRODUCTION

Chemical heat pumps convert low temperature heat to high temperature heat with only a small amount of electricity or work requirements. Organic chemical heat pumps are based on reversible organic reactions. The key unit operations in a chemical heat pump include an endothermic reactor, an exothermic reactor, and a distillation column [1]. Low temperature heat input (primarily waste heat) is given to the endothermic reactor and high temperature heat is released from the exothermic reactor. Because the reactions approach equilibrium, the distillation column separates the reactants from the products to establish the necessary driving forces. Ideally, the distillation column reboiler can be powered by waste heat at the same or lower temperature as the endothermic reactor.

The most researched organic chemical heat pump is the isopropanol + acetone based chemical heat pump. It is generally best used for lifting heat from about 80-90°C to about 150-210°C [1]. In this work, we examine a chemical set (methylcyclohexane and toluene) that could lift heat from about 200°C to about 250°C that has not been studied previously for chemical heat pump applications. Drying in textile and distillation in chemical industries are some of the processes that require heat in the range 200°C-250°C [2]. The methylcyclohexane and toluene based reversible system has been widely explored in other applications, such as a liquid organic hydrogen carrier. In that application, the hydrogen is stored in the form of methycyclohexane by using hydrogenation of toluene. Hydrogen is released during the dehydrogenation of methyl cyclohexane as shown:

$$C_7H_{14} \rightleftharpoons C_7H_8 + 3H_2 \tag{1}$$

Usman *et al.* [3] studied the dehydrogenation kinetics of methylcyclohexane in the temperature range of 340°C-380°C. Castano *et al.* [4] studied the toluene hydrogenation kinetics in the temperature range of 100°C-250°C. These studies provide evidence that the reversible reaction is feasible within the desired temperature ranges suitable for our proposed chemical heat pump application. Toluene and methylcyclohexane have a difference in boiling point by around 10°C, which enables their separation in a distillation column. Additionally, toluene and methylcyclohexane also meet the ozone depletion potential and global warming potential screening criteria as per [5].

In this work, we designed a novel methylcyclohexane–toluene–hydrogen (MTH) based chemical heat pump (CHP) and determined its key performance indicators using Aspen Plus process simulations.

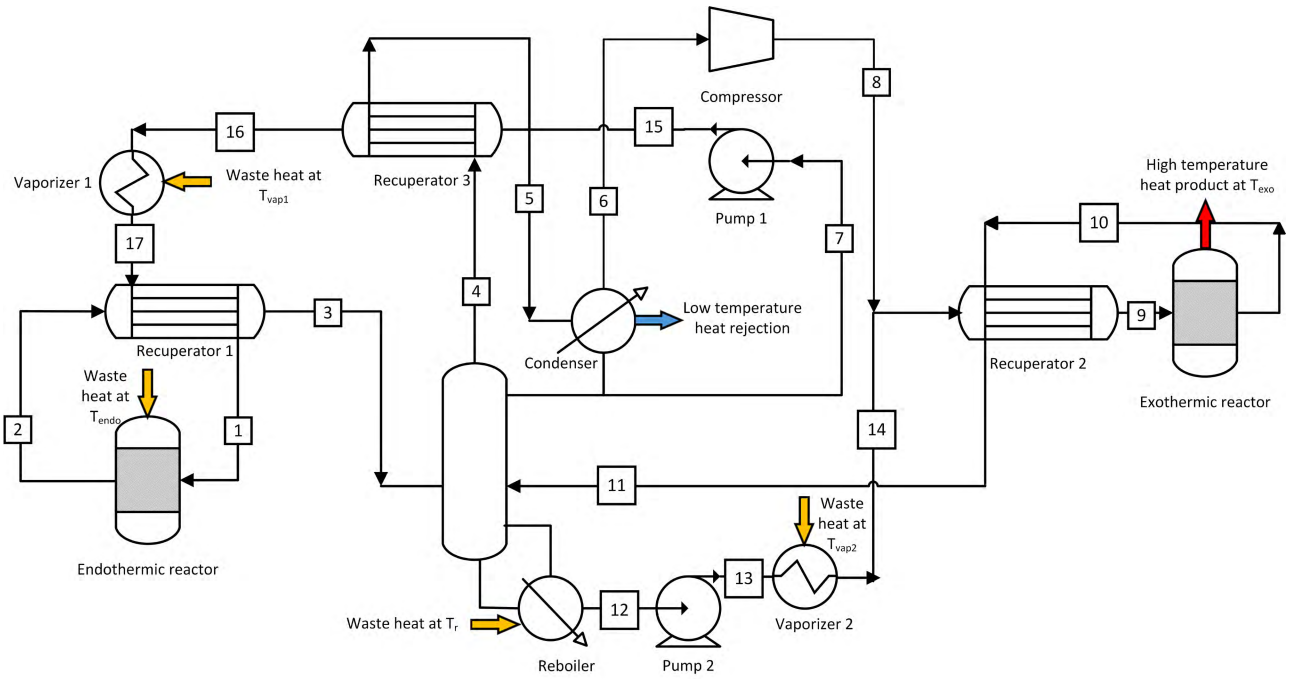

Figure 1. Process scheme of MTH-CHP

Table 1: Stream details of MTH-CHP at for T_{endo} = 200°C, T_{exo} = 250°C when COP is maximized

Stream number	Temperature, °C	Pressure, bar	Mole Flow ratio (Stream X / Stream 1)	Methylcyclohexane mole fraction	Hydrogen mole fraction	Toluene mole fraction
1	190	1.23				
2	200	1.13	1.00	0.6174	0.0028	0.3798
3	113	1.10	1.10	0.5321	0.0923	0.3756
4	92	1.06	1.10	0.5321	0.0923	0.3756
5	86	1.06	1.70	0.4419	0.2903	0.2678
6	35	1.01	1.70	0.4419	0.2903	0.2678
7	35	1.01	0.54	0.0638	0.9098	0.0264
8	59	1.30	1.00	0.6174	0.0028	0.3798
9	240	1.28	0.54	0.0638	0.9098	0.0264
10	250	1.18	0.61	0.1146	0.7969	0.0885
11	83	1.15	0.52	0.2002	0.7580	0.0418
12	107	1.10	0.52	0.2002	0.7580	0.0418
13	107	1.35	0.08	0.4729	-	0.5271
14	113	1.30	0.08	0.4729	-	0.5271
15	35	1.30	0.08	0.4729	-	0.5271
16	82	1.30	1.00	0.6174	0.0028	0.3798
17	110	1.25	1.00	0.6174	0.0028	0.3798
			1.00	0.6174	0.0028	0.3798

SYSTEM DESCRIPTION

The process scheme of the MTH-CHP is shown in Figure 1 with corresponding stream conditions in Table 1. It consists of an endothermic reactor, an exothermic reactor, a distillation column, recuperators, pumps and a compressor. The methylcyclohexane dehydrogenation takes place in the endothermic reactor and the toluene hydrogenation occurs in the exothermic reactor. Both the reactor effluents, containing all the chemicals, enter the distillation column for separation. The distillation column bottom consists of mostly toluene since it is the heaviest component. The distillate consists of hydrogen and methylcyclohexane. The hydrogen and methylcyclohexane need to be separated since hydrogen is required in

the exothermic reactor and methylcyclohexane is required in the endothermic reactor. A partial condenser facilitates this separation. Because the reactions happen in the gas phase, methylcyclohexane in the partial condenser liquid product and toluene at the distillation bottom and need to be fully vaporized (vaporizer 1 and 2 respectively). To reduce the heat duty of vaporizer 1, recuperator 3 is present to recover some of the heat from the distillate. Recuperator 1 and recuperator 2 recover heat from the endothermic and exothermic reactor outlet respectively.

The MTH-CHP requires waste heat at four different temperatures as shown in Figure 1. Out of all the four, the heat input at the endothermic reactor will have the highest temperature since the methylcyclohexane dehydrogenation is only feasible above 150°C at atmospheric pressure. The temperature of the reboiler, vaporizer 1 and vaporizer 2 is determined by the bubble point of the mixture at those pressure conditions.

METHODOLOGY

Modelling and simulation

A steady state simulation model was created using the software Aspen Plus V14. The PRMHV2 (Peng–Robinson equations of state with modified Huron–Vidal mixing rules) property method is chosen since the model phase equilibrium predictions are in close agreement with the NIST experimental data [10], as shown in Figure 2. The flow sheet was solved using Equation Oriented (EO) mode. Convergence was achieved by first solving the flowsheet in sequential modular mode with one or more loops broken (a.k.a. manual tearing), such that the solver converges but there is a small error in the global mass and energy balances. Then the loops are closed programmatically using the connections feature in EO, and the previous inaccurate solution is used as an excellent initial guess. The final converged EO solution satisfies all mass and energy balances.

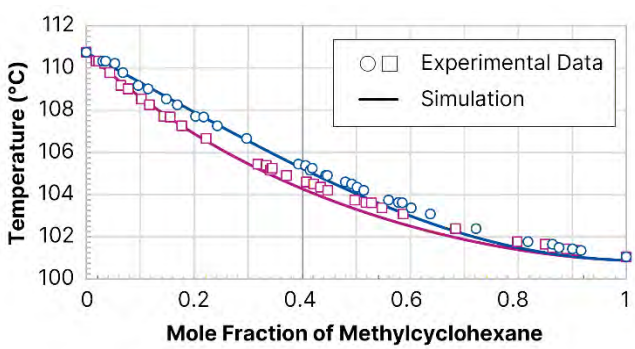

Figure 2. A comparison of the PRMHV2 vapor-liquid equilibria predictions for methylcylohexane and toluene vs. experimental data.

RADFRAC was chosen for modelling the distillation column in equilibrium made. The Design Spec feature in RADFRAC was used to meet the column top and bottom purity specs. The partial condenser temperature in RADFRAC was chosen as 35°C which maximized the separation between hydrogen and organics. The operating pressure of condenser is set as atmospheric. For both the endothermic and exothermic reactor, RGIBBS was used, which gives the maximum conversion possible since chemical equilibrium is assumed. The reactor assumes isothermal condition. 100% selectivity was assumed for both the reactors (in other words, we ignore potential side reactions not mentioned in this paper so as to achieve an optimistic performance bound). All the recuperators and vaporizers were modelled using HEATX and HEATER models respectively. The vapor fraction of vaporizer 1 is considered as 0.95 to reduce its heat duty, while the recuperator present in the downstream handles the remaining heat load. The approach temperature, defined as the temperature difference between hot inlet and cold outlet stream, is set to be 10°C for all the recuperators. The COMPR model was used for the compressor and the default isentropic efficiency of 0.72 was assumed. Pumps were modelled using PUMP with the default pump efficiency. Pumps and compressor are present to account for the pressure drops in the system. The pressure drop of the reactor, vaporizer, partial condenser and recuperators were assumed to be 0.1, 0.05, 0.05 and 0.025 bar respectively. A tray pressure drop of 0.005 bar per tray was considered [6].

For brevity, only the most relevant details are contained in this paper. The model has been released to the public as supplementary material (see link at end of article). Full details of the model and steam converged conditions of a representative case can be found there.

Performance analysis

The COP of MTH-CHP is calculated by the formula:

$$COP = \frac{Q_{exo}}{Q_{endo}+Q_{vap1}+Q_{vap2}+Q_r} \quad (2)$$

where Q_{exo} is the rate of high temperature heat produced from the exothermic reactor, Q_{endo} is the rate of the heat consumed in the endothermic reactor needed for the reaction and to raise the feed stream to reaction temperature, Q_r is the reboiler heat duty, Q_{vap1} is the heat duty of vaporizer 1 and Q_{vap2} is the heat duty of vaporizer 2 (all in kW).

To enable the comparison with mechanical heat pumps, the COP_w for a chemical heat pump is defined as follows:

$$COP_W = \frac{Q_{exo}}{W_{in}} \quad (3)$$

where $W_{in} = W_c + W_{p1} + W_{p2}$ and W_c, W_{p1}, W_{p2} is the power

of compressor, pump 1 and pump 2 in kW.

Optimization framework

The performance of MTH-CHP depends on many design parameters such as endothermic and exothermic reactor temperatures and distillation column parameters. Hence, multiparameter optimization is carried out to maximize the performance. The optimization problem is defined as follows:

$$\max_{d} COP$$

$$f(x) = 0 \quad \text{[all equations in Aspen Plus, \& eq 2]}$$

$$g(x) \leq 0 \quad \text{[all inequalities in Aspen Plus]}$$

$$L \leq d \leq U \quad \text{[all decision variable bounds]}$$

The particle swarm optimization algorithm (PSO) was used for optimization using the method described in [7]. There are 6 total degrees of freedom. 4 degrees of freedom were taken as decision variables d described in Table 2. The remaining 2 degrees of freedom are the total number of stages in the distillation column (N_{stages}), and the ratio of the molar flow rate of hydrogen to toluene at the exothermic reactor inlet ($r_{H2/tol}$). These two were fixed or taken as parameters. The objective function is theoretically maximized at infinite N_{stages}, and so reasonable upper bounds much be chosen as described later. Although out of the scope of this work, they can be decision variables if the objective function were based on economic or other metrics that balance capital (N_{stages}) versus operating ($r_{H2/tol}$) tradeoffs.

Table 2 Decision variables considered for optimization.

d No	Description	Lower Bound	Upper Bound
1	Mole fraction of toluene in the partial condenser condensate stream	0.01	0.99
2	Mole fraction of methylcyclohexane in distillation bottom stream	0.01	0.99
3	Feed Tray Location of endothermic feed (tray 1 is partial condenser)	2	$N_{stages} - 1$
4	Feed Tray Location of exothermic feed (tray 1 is partial condenser)	2	$N_{stages} - 1$

The number of particles considered for all runs is 40. The maximum number of iterations is set as 100. Each case was run for at least twice to ensure that the result was consistent. The arbitrary upper and lower bounds for d_1 and d_2 were wide enough such that d_1 and d_2 at the optimum was never at a bound.

Table 3 COP variation with number of stages at T_{endo} = 200°C, T_{exo} = 250°C, $r_{H2/tol}$ =3

N_{stages}	d_3	d_4	COP
5	4	2	0.1176
8	5	3	0.1181
10	7	3	0.1183
15	8	4	0.1183
20	9	2	0.1183

Prior to solving the full optimization problem, N_{stages} was set to 10 based on a preliminary sensitivity study. In that preliminary study, N_{stages} was varied around a base case of a "very good" design candidate with T_{endo} = 200°C, T_{exo} = 250°C, d_1 = 0.43 and d_2 = 0.45. The values for d_1 and d_2 were chosen in accordance with our previous work for a different chemical heat pump system where "sloppy distillation" is optimal from efficiency point of view [7]. $r_{H2/tol}$ was set as 3, which is stoichiometrically balanced (eq. 1). Then, for each value of N_{stages} considered in the preliminary sensitivity study, the other two decision variables d_3 and d_4 were optimized using simple enumeration of the possible feed tray location combinations, to find the combination that yielded the maximum COP. It was found that the COP increased with N_{stages} up to 10 stages. $N_{stages} > 10$ did not improve the COP meaningfully. The results are populated in Table 3. N_{stages} is relatively small mainly because the vaporizer duty represents the largest quantity of heat load and reboiler duty is not the dominant term.

RESULTS AND DISCUSSION

The optimization problem described in the previous section was solved for different $r_{H2/tol}$ of 3, 6, 9 and 12 and for T_{endo} = 200°C, T_{exo} = 250°C. Performance was recorded in terms of COP and COP_W.

Table 4 Solutions to the maximize COP problem for each $r_{H2/tol}$ value taken as a parameter, with the corresponding COP_W shown.

$r_{H2/tol}$	COP	COP_W
3	0.1202	19.6
6	0.1341	16.7
9	0.1357	13.3
12	0.1357	10.4

Toluene hydrogenation which is the exothermic reaction is favored at higher $r_{H2/tol}$ and there is a COP increase initially. However, the COP does not increase

further when $r_{H2/tol}$ is above 9 because a major portion of the heat released from the exothermic reactor is used for heating the excess H_2 to exothermic reaction temperature. COP_W reduces with higher $r_{H2/tol}$ due to higher work requirement in compressor. Based on these results, $r_{H2/tol}$ is fixed at 9 for further runs.

Finally, T_{endo} and T_{exo} were varied and the maximized COP obtained by the optimization procedure described earlier and its corresponding COP_w is populated in Table 5. Increasing the T_{exo} for the same T_{endo} reduces COP because the heat liberation from the exothermic reactor (Q_{exo}) reduces at higher temperatures. The rate of methylcyclohexane dehydrogenation (Q_{endo}) increases with increase in endothermic reactor temperature and the rate of toluene hydrogenation decreases with increase in exothermic reactor temperature. This leads to more toluene and hydrogen in the system at higher endothermic and exothermic reactor temperatures. This causes increase in the compressor load and hence there is a drop in COP_W at higher reactor temperatures. For the same reason, the toluene concentration in the partial condenser condensate increases, while the methylcyclohexane concentration in the distillation bottom decreases as the reactor temperature increases (Table 6). Energy flows per unit of product heat produced are shown in Table 7 for a representative case.

T_{exo} of 300°C and above is not feasible for the MTH-CHP. This is because the sensible heat required to heat the feed inlet stream from 290°C to 300°C is more than the exothermic heat liberated at 300°C.

Table 5 Variation of maximized COP with T_{endo} and T_{exo}

T_{endo}, °C	T_{exo}, °C	COP	COP_W
175	225	0.0632	18.3
200	250	0.1357	13.3
200	275	0.0766	4.5
225	275	0.1614	4.8

Table 6: Decision variable values for the maximized COP case

T_{endo}, °C	T_{exo}, °C	d_1	d_2	d_3	d_4
175	225	0.2307	0.7371	6	7
200	250	0.3798	0.4729	5	2
200	275	0.7490	0.1285	8	2
225	275	0.7607	0.1277	9	4

Integration with vapor compression heat pump

The vaporizer 1 heat duty represents the majority of heat consumption as shown in Table 7. A design improvement of MTH-CHP could consist of including a conventional vapor compression heat pump. The function of a vapor compression heat pump is to take part of the partial condenser heat duty as its heat source and deliver heat at a higher temperature to vaporizer 1. Thus, this hybrid system not only reduces the hot utility demand of vaporizer 1, but also the cold utility demand of the partial condenser, although at the penalty of additional work/electricity requirements.

Table 7: Heat duties and electric power per unit of product heat delivered for T_{endo} = 200°C, T_{exo} = 250°C for the maximized COP case

Ratio	Value
Q_{vap1}/Q_{exo}	5.38
Q_{vap2}/Q_{exo}	0.38
Q_{endo}/Q_{exo}	1.36
Q_r/Q_{exo}	0.25
W_{in}/Q_{exo}	0.08

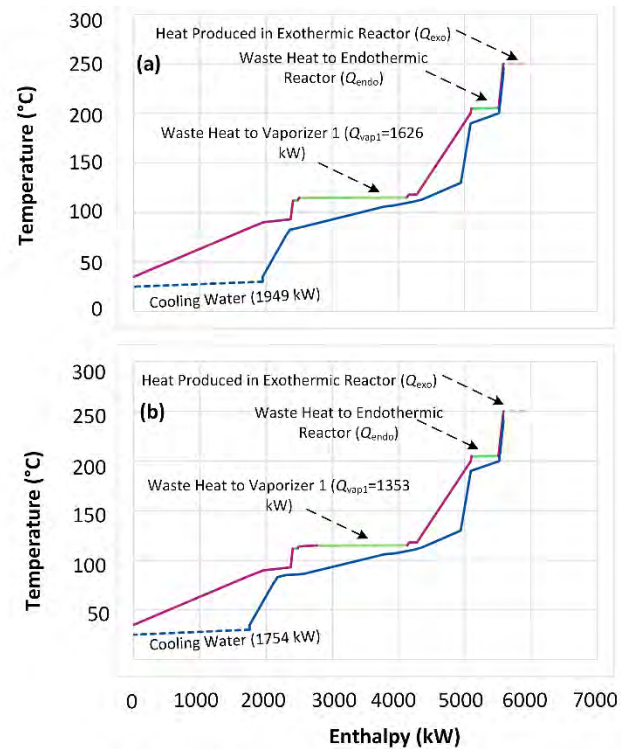

Figure 3. (a) Composite curve of MTH-CHP for T_{endo} = 200°C, T_{exo} = 250°C . (b) Composite curve of the hybrid heat pump design when 10% of the partial condenser heat duty is used as a heat source for the vapor compression heat pump. Both systems produce the same amount of exothermic heat, but the second system requires less waste heat consumption.

To look at the possibilities of mechanical heat pump integration, a composite curve was constructed for the MTH-CHP, using Aspen Energy analyzer. Figure 3a shows the composite curve of MTH-CHP without any heat integration for T_{endo} of 200°C and T_{exo} of 250°C. For this composite curve, a pinch temperature of 5°C was

chosen. Figure 3b shows what the composite curve would be if 10% of the partial condenser heat were upgraded and used for vaporizer 1 with a vapor compression heat pump, with a pinch temperature of 5°C. We assume a COP_w of 3.5 for the vapor compression heat pump where the heat source temperature is around 85°C and heat sink temperature is around 115°C [8]. This increases the high-quality heat product generated per unit waste heat consumed by about 14%. This was repeated for 5% and 25% of partial condenser duty upgrading, with results summarized in Table 8.

Table 8: Performance result of hybrid heat pump

Heat source of vapor compression heat pump	COP	% Increase	COP_W
25% of condenser duty	0.1955	44%	1.4
10% of condenser duty	0.1546	14%	3.0
5% of condenser duty	0.1445	6.5%	4.9

The COP presented in Table 8 reflects the result of heat integration already present as per Figure 1 as well as the contribution from the mechanical heat pump integration. The results show that this hybrid system allows the increase of COP at the cost of COP_w. Increasing the scale of the vapor compression heat pump to 5% -10% of the partial condenser duty would result in COP_w being in the same range as that of many mechanical heat pumps. Currently there is at least one mechanical heat pump at TRL 9 that can supply heat up to 280°C [9].

CONCLUSIONS

The performance of novel chemical heat pump based on methylcyclohexane, hydrogen, toluene was studied. The most promising results are obtained at T_{endo} of 200°C and T_{exo} of 250°C, where we achieve a COP of 0.136 and a COP_w of 13. Increasing the temperature lift beyond 50°C significantly reduces the COP for this chemical heat pump. Furthermore, a new configuration of hybrid system combining both chemical heat pump and mechanical vapor compression heat pump was studied. This configuration resulted in higher COP at reduced COP_w. Very few commercial heat pumps currently operate within this range, and future economic assessments will determine whether this MTH-CHP is commercially viable.

DIGITAL SUPPLEMENTARY MATERIAL

The aspen plus simulation file of the MTH-CHP is available in LAPSE at https://psecommunity.org/LAPSE:2026.0022.

ACKNOWLEDGEMENTS

The authors would like to thank the Energy and Process Engineering Department at the Norwegian University of Science and Technology for financial support. We also thank Prof. Lars O. Nord and Lisa Marie Njå for their valuable input.

AUTHOR IDENTIFIERS

Author ORCIDs:
Krishnadoss R: 0009-0007-0138-1414
Le Bot F: 0009-0007-5659-5823
Adams TA II: 0000-0002-9871-9851

REFERENCES

1. KlinSoda I, Piumsomboon P. Isopropanol–acetone–hydrogen chemical heat pump: a demonstration unit. Energy Conversion and Management 48:1200-1207 (2007). https://doi.org/10.1016/j.enconman.2006.10.006
2. Brückner S, Liu S, Miró L, Radspieler M, Cabeza LF, Lävemann E. Industrial waste heat recovery technologies: an economic analysis of heat transformation technologies. Applied Energy 151:157-167 (2015). https://doi.org/10.1016/j.apenergy.2015.01.147
3. Usman MR, Aslam R. Dehydrogenation of methylcyclohexane for on-board hydrogen use: initial rate kinetics over 1.0 wt% pt/?-al2o3. Arab J Sci Eng 39:615-620 (2013). https://doi.org/10.1007/s13369-013-0694-y
4. Castaño P, Arandes JM, Pawelec B, Fierro JLG, Gutiérrez A, Bilbao J. Kinetic model discrimination for toluene hydrogenation over noble-metal-supported catalysts. Ind. Eng. Chem. Res. 46:7417-7425 (2007). https://doi.org/10.1021/ie070094m
5. Agromayor R, Nord LO. Fluid selection and thermodynamic optimization of organic rankine cycles for waste heat recovery applications. Energy Procedia 129:527-534 (2017). https://doi.org/10.1016/j.egypro.2017.09.180
6. Woods D. Homogeneous Separation (pp. 86–136). In: Rules of Thumb in Engineering Practice. Wiley-VCH (2007). ISBN 978-35-27312-20-7
7. Krishnadoss R, Dyrland IA, Adams TA II. Design and optimization of alcohol-ketone-hydrogen chemical heat pumps. Computers & Chemical

Engineering 199:109158 (2025). https://doi.org/10.1016/j.compchemeng.2025.109158
8. SPH Sustainable Heat GmbH. "High temperature heat pump ThermBooster." https://spheat.de/thermbooster/?lang=en (accessed 28 Jan, 2026).
9. Zühlsdorf B et al. Task 1 report - High Temperature Heat Pumps. Heat Pump Centre (2023). ISBN 978-91-89821-34-7
10. Tyminski B, Klepanska A. Vapor-liquid equilibriums in the methylcyclohexane-toluene system. *Inz. Chem* 7: 193-205 (1977).

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

An Integrated Process of Multi Effect Distillation Based Desalination with Renewable Energies: Evaluation of Power Generation Efficiency and Freshwater Production Cost

Mohammed Adam[a], Mudhar A. Al-Obaidi[b], and I. M. Mujtaba[a*]

[a] Chemical Engineering Department, Faculty of Management, Sciences and Engineering, University of Bradford, West Yorkshire BD7 1DP, UK
[b] Middle Technical University, Technical Instructor Training Institute, Baghdad 10074, Iraq
* Corresponding Author: I.M.Mujtaba@bradford.ac.uk

ABSTRACT

As the demand of freshwater continuous to rise, the water desalination can act as prominent solution to address the global water scarcity. However, significant environmental concerns arise as the process is mainly powered by fossil fuel, which contributes to greenhouse gas emissions and thus leading to global warming. This research intends to explore the potential of renewable energy sources to effectively reduce freshwater production cost. Specifically, it intends to estimate the power generated and freshwater production cost of a multi-effect distillation (MED) desalination process, powered by solar and wind energy sources in addition to a comparison against the performance of an MED process powered by fossil fuel based conventional steam boiler. The comparison of energy efficiency and freshwater production cost is conducted in two different UK locations: Wick in the north of Scotland and Watchet in the south of England. MED process model with both wind power and solar energy models are developed to estimate the freshwater production cost while assuring 12kg/s of freshwater for small coastal communities. The results ascertain that both energy sources are able to produce thermal energy required for the MED process. Statistically, for a full day, 4029 kWh and 16383 kWh amounts of electricity can be generated in Watchet and Wick, respectively. Referring to wind powered MED, the freshwater production cost in Watchet is estimated at \$20.49 per m³ of freshwater whereas it is reduced to \$5.56/m³ at Wick as due to stronger and more consistent wind speeds, which is less than a third of the Watchet freshwater production cost. However, solar powered MED shows a much smaller difference between the locations: \$1.87/m³ in Watchet and \$3.11/m³ in Wick. Based on the results of freshwater production cost, the solar energy for MED process is more efficient in Watchet, whereas wind energy for MED process is more efficient in Wick.

Keywords: Desalination, Multi effect distillation, Hybrid, Wind Energy, Solar Energy

1. INTRODUCTION

As water demands continues to rise due to population growth, and agricultural demand, desalination process has become the main solution, especially in areas experiencing declining freshwater supplies. In the United Kingdom, high amount of rainfall has supported freshwater reservoirs and rivers. However, due to climate change, with warmer summers and less rainfall, the level of water in reservoir have been declined. Desalination process is considered as a long-term strategy to help meet water demand and provide a reliable source of clean water. Three desalination technology that are used for desalination processes are MED, multistage flash (MSF) and reverse osmosis (RO) [1]. Desalination is the removal of salts and impurities from seawater to produce safe water to drink. MED is a thermal process that uses multiple effects, which allows seawater to evaporate and

condense at progressively lower pressures and temperatures [2]. This process efficiently recycles heat between effects, which makes it more energy efficient. The Middle East and Africa regions dominates the global desalination market, by owning majority of the installed desalination technologies due to its dry climate, low amount of rainfall, high evaporation rates, and increasing water scarcity. The largest MED plant around the globe was constructed in Saudi Arabia called the Jubail Plant, which serves dual purpose facility by its combined cycled based power generation and MED plant that produces 800, 000 m^3/day of freshwater [4].

There are few examples of articles regarding MED powered by Solar Energy, using Solar Collectors. Dilli and Kumar [7], presented an experimental performance evaluation of a solar powered MED system with a parallel feed configuration. The main target was to assess how real environmental and operational factors such as solar irradiance, temperature, pressure, and steam flow rate can affect the plant's desalination capacity and thermal performance. The researchers used a 10 m^3/day of freshwater output solar MED plant. The MED system included four effects and was powered using solar heated steam. Experiments were conducted under variable solar irradiance conditions throughout the day. Parameters such as brine and distillate flow, evaporator pressure and temperature, and steam mass flow were measured simultaneously. The system achieved a maximum GOR of 3.6, indicating efficient heat reuse across effects. As solar irradiance increased by 24%, evaporation rates significantly increased. Freshwater production was higher in the afternoon despite higher morning solar irradiation, with higher steam production rates directly proportional to improved GOR and overall plant efficiency. This research offered valuable understanding of real world thermal performance, however, its limitations are the system's small scale and a lack of energy storage. The absence of backup storage means that productivity remains tightly linked to real time solar conditions.

Muftah et al. [11] explored how solar energy can be used to power MED plants to make freshwater in a more sustainable way. It focused on a plant in Zuara, Libya, and compared four different ways to run the system. The 1st method used solar energy only during daylight hours. The 2nd method was to combine solar power during the day with boiler operating at night. The 3rd method used solar power along with a thermal energy storage system to allow 24-hour process. The researchers developed a validated mathematical model based on real empirical data. The findings showed that the first method has a freshwater production cost of \$1.40/m^3 and no emissions. Also, the second and third methods introduced a freshwater production cost of \$0.79/m^3 with moderate emissions, and \$0.97/m^3 with no emissions, respectively.

Blanco and Alarcón [5] presented the Aquasol Project, which was built in Spain as a pilot plant for research project funded by the EU to test and prove the technology of solar powered MED with combined hybrid heat source. The Aquasol developed a high capacity solar thermal desalination system that combines renewable energy with MED to reduce fossil fuel dependence and achieve no brine discharge. The main goal was to improve efficiency, reduce energy consumption, and enhance sustainability of seawater desalination at scale. The findings stated that at feed water flow rate of 8 m^3/h, the plant can produce 3 m^3/h of distilled water resulting a thermal power consumption of 190 kW_{th}. It achieved a PR over 9. These results emphasised that the pilot plant operated successfully showing both high freshwater output and very efficient thermal usage. The Aquasol plant highlighted the advantages of hybridisation in achieving high freshwater output, improved efficiency, and zero brine discharge due to a solar powered brine dryer. However, its reliance on fossil fuel backup weakens its long term sustainability.

There are few examples of articles below regarding MED powered by Wind Energy, using Wind Turbine. Rostami et al. [12] presented a hybrid desalination system that uses both waste heat and electricity from a wind turbine. It combined three desalination technologies: A humidification dehumidification unit, MED and MVC systems. The key aim was to minimise liquid brine discharge while improving freshwater productivity by exploiting waste heat, reducing reliance on external energy inputs. The system has a waste heat recovery subunit, which captures heat from the wind turbine cooling system. This recovered waste heat drives a closed air, open water HDH desalination unit. The 90% rejected brine from the HDH unit was fed into the 7 effect MED-MVC for further water recovery. This allowed to a maximise the freshwater productivity while minimising the brine discharge. The findings were that freshwater output increased by 18% across all wind speeds compared to prior layouts. In this regard, Specific Work Consumption (SWC) dropped from 33.41 to 21.27 kWh/ as wind speed increased. When compared to a RO system powered by electricity from wind turbine, this hybrid system of MED-MVC exhibited lower SWC and better recovery ratio, especially under higher wind conditions.

Khalilzadeh and Hossein Nezhad [10] proposed a desalination system that captures waste thermal energy from a 7.58 MW wind turbine to produce steam for a six-effect MED unit. The target was to enhance overall plant efficiency by turning waste turbine heat into freshwater, particularly in remote, wind-rich areas lacking potable water supplies. To capture waste heat from turbine's cooling system was fulfilled by the closed loop heat exchanger. The extracted heat generated steam at 100 ℃ at standard pressure, was fed in a 6 effect MED unit. The

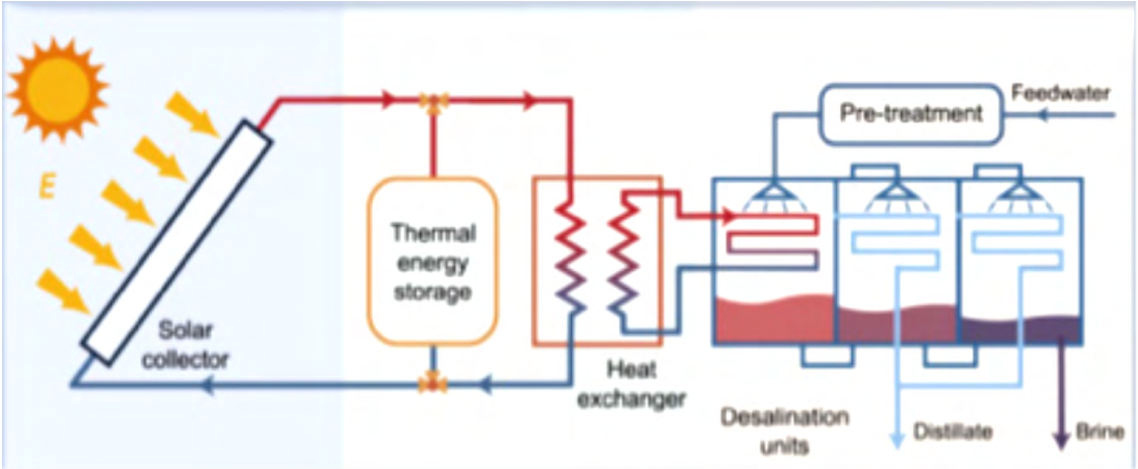

Figure 1. The concept of full model of Solar Powered MED [15]

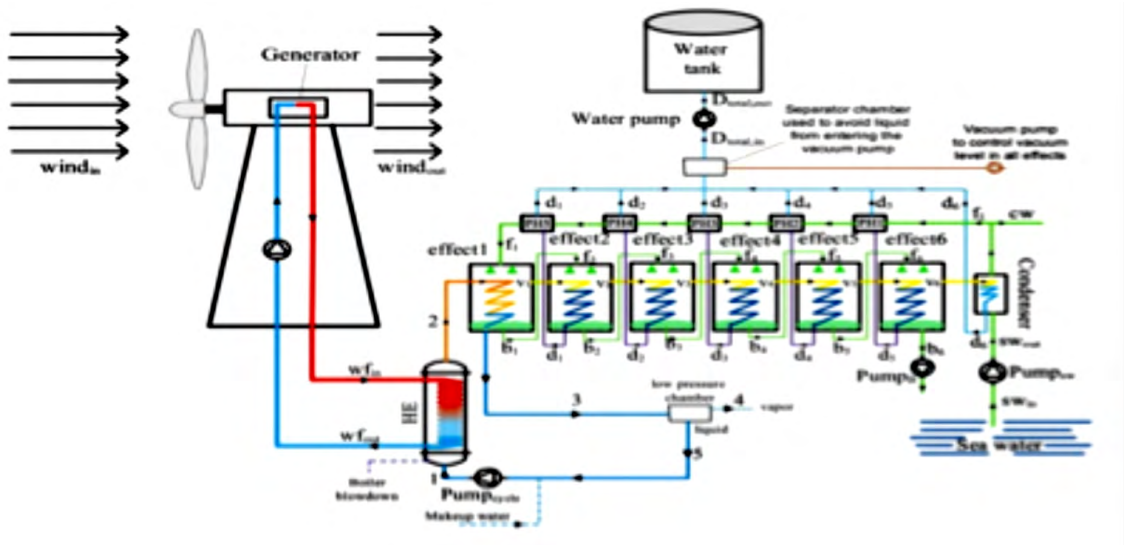

Figure 2. The concept of full model of Wind Powered MED [10]

power required for the pumps and pre-heaters were fulfilled by the electricity generated by the same turbine. The system produced 45 m^3/day of freshwater, with an estimated freshwater production cost of \$16.676/m^3, and a payback period of 6.33 years, based on its capital and operational costs.

2. MODEL OF SOLAR AND WIND POWERED MED

Solar irradiance is the total amount of light energy that reaches the Earth since it is the output of light energy that is received from the Sun. In UK, solar irradiance levels are moderate but still sufficient for effective solar energy production. In the North of Scotland, the solar irradiance can reach up to 71.8 W/m^2, whereas in the South of England, the solar irradiance can reach up to 128.4 W/m^2. Therefore, in this research, coastal towns from South of England called Watchet and coastal town in North of Scotland called Wick were selected for their data of Solar Irradiance and Wind Speed. Thermal energy from a solar collector is used to heat water, which acts as the Heat Transfer Fluid (HTF). The HTF flows through a temperature-controlled pipeline. If its temperature exceeds the required temperature, the fluid is transferred to a thermal energy storage tank. From the tank, the HTF at the appropriate temperature is then supplied to a heat exchanger, which is used in the MED process. Figure 1 shows the concept of this process.

Wind speed can be converted into either electrical energy by using Wind Turbines or thermal energy from the waste heat, which is lost from the wind turbine Generator and Gear components. Figure 2 shows the concept of this process.

3. METHODOLOGY

3.1 Model Equations

3.1.1 MED Process Equations

The MED mode equations produced by Filippini et al [9] is presented below. The MED model composes a set of algebraic and differential equations representing the thermodynamic and mass transport characteristics of a MED system. The model is based on fundamental mass and energy balance principles and includes equations for heat transfer, boiling point elevation, and salinity effects. Heat exchangers within pre-heaters and condensers is modelled using logarithmic mean temperature differences, with the model designed to simulate thermal performance under different input conditions.

Feed Flowrate

$$Mf = \frac{Ms.\lambda Ts}{Q_{sensible}+Q_{Latent}} \quad (1)$$

Sensible Heat in First Effect

$$Q_{sensible} = Mf \int_{t1}^{T1} cp\,(T1, x1)dt \quad (2)$$

Latent Heat in First Effect

$$Q_{Latent} = D1\,.\,\lambda(Tv1) \quad (3)$$

Temperature drop among effects

$$\Delta T = \frac{Ts - Tb}{n} \quad (4)$$

Temperature drop among pre-heaters

$$\Delta T = \Delta t \quad (5)$$

Feed temperature in first effect

$$t1 = tn + (n-1)\Delta t \quad (6)$$

Temperature of vapor phase

$$Tv = T - BPE\,(T, x) \quad (7)$$

Flowrate of flashed distillate

$$D_{flash,i} = \propto B_{i-1} \quad (8)$$

Fraction of distillate by flashing

$$\propto = \frac{cp\,(T_{mean} - x_{mean})\,\Delta T}{\lambda(T_{mean})} \quad (9)$$

Mean Temperature

$$T_{mean} = \frac{T1 + Tb}{2} \quad (10)$$

Mean Salinity

$$x_{mean} = \frac{xf + xb}{2} \quad (11)$$

Fraction of distillate by evaporation

$$\beta = \frac{\propto [xb(1-\propto)^n - xf]}{(xb - xf)[1-(1-\propto)^n]} \quad (12)$$

Flowrate of evaporated distil

$$D_{i,boiled} = \beta M_D \quad (13)$$

Total distillate

$$D_i = D_{i,boiled} + D_{i,flashed} \quad (14)$$

Rejected brine flowrate

$$B_i = B_{i-1} - D_i \quad (15)$$

Salinity profile in the effects

$$x_i = \frac{x_{i-1}B_{i-1}}{B_i} \quad (16)$$

Area of i-th effects

$$A_{ev,i} = \frac{Q_i}{U_{ev,i}\,.\Delta T_{ev,i}} \quad (17)$$

Heat load in i-th effect

$$Q_i = D_{boiled,i-1}\,.\,\lambda(T_{v,i-1}) \quad (18)$$

Temperature drop in Heat Exchanger

$$\Delta T_{ev,i} = \Delta T - BPE_{i-1} \quad (19)$$

Area of i-th pre – heater

$$Mf\,.\int_{t_{i+1}}^{t_i} cp(t, xf)dt = U_{ph,i}\,.A_{ph,i}\,.\Delta t_{log,i} \quad (20)$$

Logarithmic temperature difference in pre-heaters

$$\Delta t_{log,i} = \frac{\Delta t}{log\,(\frac{Tv_i - t_{i+1}}{Tv_i - t_i})} \quad (21)$$

Area of Final Condenser

$$Q_{COND} = U_{COND}\,.A_{COND}\,.\Delta T_{log,COND} \quad (22)$$

Heat load in final condenser

$$Q_{COND} = D_n\,.\lambda(Tv_{n)} \quad (23)$$

Logarithmic Temperature difference in final condenser

$$\Delta T_{log,COND} = \frac{tn - Tw}{log\,(\frac{Tv_n - Tw}{Tv_n - tn})} \quad (24)$$

3.1.2 Solar Collector Process Equations

The Solar model equations presented below. The solar thermal model is formulated to calculate the thermal energy produced from solar irradiance and from the storage tank which will be transferred to be used in the MED model. It is composed of a set of algebraic equations describing collector performance, and the total energy gained by solar irradiance and thermal storage.

Solar Thermal power gained from Solar collectors

$$Q_{solar} = N_{collector}\,.A_{Collector}.\,\eta_{Collector}\,.\frac{I}{1000} \quad (25)$$

Total Heat supplied to MED

$$Q_{MED} = Q_{solar} + Q_{storage} \quad (26)$$

Heat Required by MED for desalination to work

$$Q_{Required} = \dot{X}_{Freshwater} \cdot \lambda_{Evap} \cdot C_p \cdot \Delta T \quad (27)$$

Energy Balance of Thermal Tank

$$\frac{dE_{storage}}{dt} = Q_{storage} - Q_{Charge} - Q_{Loss} \quad (28)$$

Thermal Tank Temperature with Stored HTF

$$T_{Tank} = \frac{T_{Ambient} + E_{storage}}{\rho_{HTF} \cdot C_p \cdot V_{Tank}} \quad (29)$$

Temperature of HTF from Thermal Tank

$$T_{Storage} = T_{Tank} \quad (30)$$

Temperature of HTF from Solar Collectors

$$T_{Collector} = T_{Exit_HX} \cdot \frac{Q_{solar}}{\dot{m}_{HTF} \cdot C_p} \quad (31)$$

Heat Charged into Thermal tank from excess HTF from Solar collector

$$Q_{Charge} = A \cdot \eta_{Collector} \cdot \frac{I}{1000} \quad (32)$$

Number of Collectors based on Desired Freshwater

$$N_{collector} = \frac{Q_{Required}}{Q_{solar}} \quad (33)$$

Mass Balance of HTF

$$\dot{m}_{HTF,in} = \dot{m}_{HTF,Out} \quad (34)$$

3.1.3 Wind Turbines Process Equations

The Wind model equations presented below. The wind energy model developed for this work is based on the electrical and thermal energy outputs from wind turbines. It uses a series of algebraic formulations. The model captures mechanical power generation, gear and generator losses, and waste heat recovery, with them being useful for thermal applications.

Rotor Swept Area of Wind Turbine

$$A_{Rotor} = \pi \cdot r^2 \quad (35)$$

Max power available in the wind

$$P_{Wind_Max} = \frac{1}{2} \cdot \rho_{Air} \cdot A \cdot \frac{V^3}{1000} \quad (36)$$

Mechanical power extracted by wind turbines

$$P_{Turbine} = N_{Turbine} \cdot P_c \cdot P_{Wind_Max} \quad (37)$$

Electrical power output after losses

$$P_{Electric} = P_{Turbine} \cdot \eta_{Gear} \cdot \eta_{Generator} \quad (38)$$

Waste heat generated

$$Q_{Waste} = P_{Turbine} \cdot (1 - \eta_{Gear}) + \eta_{Gear} \cdot (1 - \eta_{Generator}) \quad (39)$$

Overall efficiency of wind to electric system

$$\eta_{Systems} = \frac{P_{Electric}}{P_{Wind_Max}} \quad (40)$$

Temperature Rise in HTF due to Waste heat

$$\Delta T_{HTF} = \frac{Q_{Waste}}{\dot{m}_{HTF} \cdot C_p} \quad (41)$$

Required thermal energy to produce desired freshwater output

$$Q_{Required} = \dot{X}_{Freshwater} \cdot \lambda_{Vaporisation} \cdot C_p \cdot \Delta T \quad (42)$$

Number of Wind Turbine based on Desired Freshwater

$$N_{collector} = \frac{Q_{Required}}{Q_{Waste}} \quad (43)$$

Mass Balance of HTF

$$\dot{m}_{HTF,in} = \dot{m}_{HTF,Out} \quad (44)$$

3.1.4 Economic Model Equations

The economic equation model below is developed by Al-Obaidi et al. [3] will be used to calculate the cost of freshwater produced from Solar Powered MED process, Wind Powered MED process and Steam Generated MED process. The economic equation model provides the equation whereas the Table 1 provides economic parameters to be used for the economic equations.

Freshwater Cost

$$FWC_{MED} = \frac{TAC}{\dot{X}_{Freshwater} \cdot THY \cdot 3600} \quad (45)$$

Total Annual Cost

$$TAC = AOC + CRF \cdot TCC \quad (46)$$

Total Capital Cost

$$TCC = Capex_{Direct} + Capex_{Indirect} \quad (47)$$

Direct CAPEX

$$Capex_{Direct} = Capex_{Equipment} + Capex_{Civil_Work} \quad (48)$$

Indirect CAPEX

$$Capex_{Indirect} = 0.25 \cdot Capex_{Direct} \quad (49)$$

Equipment Cost

$$Capex_{Equipment} = C_{Intake} + C_{MED} + C_{Cond} \quad (50)$$

Civil work Cost

$$Capex_{Civil_Work} = 0.15 \cdot Capex_{Equipment} \quad (51)$$

Seawater Intake and Pre – treatment Cost

$$C_{Intake} = \frac{K_{Intake} \cdot 24.3600 \cdot \dot{m}_{Seawater}}{\rho} \quad (52)$$

MED Plant Cost

$$C_{MED} = K_{MED} \cdot C_{mat_MED} \cdot (A_{MED})^{0.64} \quad (53)$$

Final Condenser Cost

$$C_{Cond} = K_{Cond} \cdot C_{mat_Cond} \cdot (A_{Cond})^{0.8} \quad (54)$$

Annual Operating Cost

$$AOC = AOC_{Chem} + AOC_{Lab} + AOC_{Pow} + AOC_{Man} + AOC_{Steam} \quad (55)$$

Cost of Chemical Treatment

$$AOC_{Chem} = \frac{C_{Chem} \cdot THY \cdot 3600 \cdot \dot{m}_{Seawater}}{\rho} \quad (56)$$

Cost of Human Labour

$$AOC_{Lab} = \frac{C_{Lab} \cdot THY \cdot 3600 \cdot X_{Freshwater}}{\rho} \quad (57)$$

Cost of Power

$$AOC_{Power} = \frac{C_{Power} \cdot THY \cdot 100 \cdot X_{Freshwater}}{\rho \cdot \eta} \cdot f(\Delta P) \quad (58)$$

Cost of Manutention

$$AOC_{Man} = 0.002 \cdot TCC \quad (59)$$

Cost of External Steam

$$AOC_{Steam} = \frac{C_{Steam} \cdot THY \cdot (T_s - 40) \cdot \dot{m}_{Steam}}{80} + 0.005 \cdot TCC \quad (60)$$

Capital Recovery Factor

$$CRF = \frac{Ir(1+Ir)^{life}}{(1+Ir)^{life} - 1} \quad (61)$$

Table 1: Economic Parameters Developed by Al-Obaidi et al. [3]

Variables	Unit	Value	Description
THY	Hr/Yr	8760	Total Hour Per Year
K_{Intake}	$\$day/m^3$	50	Seawater Intake
K_{MED}	-	1.4	Coeff of MED
K_{Cond}	-	2.8	Coeff of Condenser
C_{Chem}	$\$/m^3$	0.024	Chemical Treatment
C_{Lab}	$\$/m^3$	0.05	Labour
C_{Power}	$/kWh	0.09	Power
C_{Steam}	$/kg	0.004	External Steam
C_{mat_MED}	$\$/m^2$	3644	Material of MED
C_{mat_Cond}	$\$/m^2$	500	Material of Condenser
Ir	-	0.07	Interest Rate
$life$	Yr	25	Life of the Plant
$f(\Delta P)$	-	3571	Pressure Losses
η	-	0.75	Efficiency of Power Generation

The parameters in Table 2 are taken from Filippini et al [9]. These value are fixed inputs that are used for the gPROM simulation. They MED model uses 10 effects, which improves thermal efficiency by reusing heat. The low Top brine Temperature (Tn from last effect) indicates MED is optimised for low temperature heat sources. Higher water quality (Tw) suggests effective salt separation by MED. The Mev shows that vapour is not being mechanically compressed to raise the pressure and temperature.

The Solar Collector Specifications are provided in Table 3. These are fixed inputs that are used for the gPROM simulation. The area is large, which allows more solar energy to be absorbed. The efficiency is quite high suggesting most of sunlight being converted into thermal energy rather than being lost. The high absorptivity helps to maintain high thermal performance especially when the solar irradiance is low.

Table 2: MED unit Parameters Used in gPROMS developed by Filippini et al. [9]

Variables	Units	Values
Number Of Stages	-	10
Tn	-	40
Xn	-	60000
Tw	ppm	25
Tphn	-	36
Pm	-	1300
Pcrit	kpa	22089
Tcrit	kpa	647.286
Mev	Kg/s	0

Table 3: Data Specification of a Solar Collector [13]

Variables	Units	Values
Area	m^2	10.1
Efficiency	%	83.8
Absorption	%	96
Max Pressure	bar	10
Cost of 1 Solar Collector	$	635

The Wind Turbine Specification are provided in Table 4. These are fixed inputs that are used for the gPROM simulation. The diameter of rotor suggests that bigger the diameter, the more power generated. The gear and generator efficiency indicates both components are very efficient with minimal energy loss. The power coefficient is the efficiency of converting wind energy to mechanical energy by the rotor.

Table 4: Data Specification of a Wind Turbine [14]

Variables	Units	Values
Diameter of Rotor	m	52
Efficiency Of Gear	%	98
Efficiency Of Generator	%	90
Power Coefficient	%	59.3
Cost of 1 Wind Turbine	$	38100

The Data specification of the steam boiler is provided in Table 5. These are fixed inputs that are used for the gPROM simulation for the Steam Boiler Powered MED process. The steam Min-Max is the value of how much steam in kg/s can be produced by the boiler.

Table 5: Data Specification of Steam Boiler [6]

Variables	Units	Values
Steam Min - Max	Kg/s	0.347 - 7.78
Max Temperature	°C	235
Max Pressure	bar	30
Cost of Steam Boiler	$	100, 000

Finally, the Data specification of the thermal tank is provided in Table 6. These are fixed inputs that are used for the gPROMS simulation which is combined with the Solar code. The insulation is made up of foam and its thickness is 25mm to reduce heat loss.

Table 6: Data Specification of Thermal Tank [8]

Variables	Units	Values
Volume Capacity	L	7000
Cost of Storage Tank	$	4360
Insulation	mm	25

4. RESULTS AND DISCUSSION

4.1 Economical Results

Based on the analysis, the initial investment almost fully determines the Levelized Cost of Water (LCOW). The Annual Operating Cost (AOC) remains below 3% of the total annualised cost for all renewable installations. This indicates that even though the energy source is free, the primary economic barrier in the UK is still the high initial cost of producing it. Wind power is particularly location-sensitive. Watchet's lower average wind speed (5.1 m/s) requires a large 1, 588 turbines to meet the system's heat demand, increasing the capital cost to $86.97 million. The turbine count in Wick decreases by 74% when speeds approach 8.0 m/s, resulting in a far more affordable water cost of $5.56/m³ instead of $20.49. Although the 25-year payback time helps keep the price gap acceptable, Wick requires almost twice as much land area as Watchet to produce the same 12 kg/s.

The research shows that the initial investment of wind power industry is where the two locations economic differences are most noticeable. The main factor driving Watchet's costs, is the high CAPEX recovery $7.58M/yr, which results directly from the need for 1, 588 turbines to compensate for lower wind speeds. Whereas, Wick, shows better resource efficiency; the higher wind speeds require less infrastructure, which enables the annualised CAPEX recovery to drop dramatically to $2.06M/yr. This shows how natural resources directly reduces the project's financial expense. Additionally, when considering all renewable energy sources, the Annual Operating Cost represents a small fraction of the total annualized cost, typically staying under 5%. This suggests that the financial feasibility of these desalination plants is mainly influenced by the upfront installation costs rather than ongoing maintenance expenses. While solar costs tend to be more stable, Wick notes a rise in the annualised repayment from $0.69 million to $1.15 million—because a larger collector field is needed in northern areas to meet the same thermal demand, due to reduced solar irradiance.

Table 7 provides the freshwater production cost ($/m³) for Wick and Watchet, while Figure 3 depicts a breakdown of annualised cost, CAPEX against OPEX.

Table 7: Freshwater production cost ($/m3) for Wick and Watchet

Parameters	Watchet Solar	Watchet Wind	Wick Solar	Wick Wind
Units	7, 255	1, 588	13, 055	412
Area, m^2	73, 275	-	131, 855	-
TCC, $M	6.62	86.97	11.92	22.56
TCC MED, $M	1.4	1.35	6.27	1.39
AOC, $k/yr	17.77	173.95	28.36	45.13
AOC MED, $k/yr	7.35	2.74	7.42	2.82
TAC, $M/yr	0.71	7.75	1.18	2.10
Cost of Freshwater, $\$/m^3$	1.89	20.49	3.12	5.56

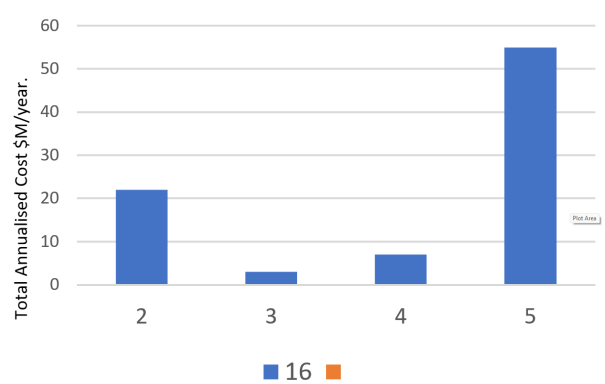

Figure 3. Breakdown of Annualised Cost, CAPEX vs OPEX

4.1.1 Comparison of Solar Powered freshwater production cost between Simulated Results and Case Study

The freshwater production costs for solar powered MED plants vary significantly between the case study of Libya, Muftah et al [11], and the current results for Watchet, and Wick. In Libya, the cost of freshwater is

around \$0.97 m^3, which is lower than the costs calculated for the UK locations This relatively low cost can be due to lower capital cost, such as, the costs of materials, equipment, and installation. Plus, labour and operational costs in Libya are considerably less expensive compared to the UK, reducing the overall expenses of the plant. In contrast, Watchet freshwater cost, \$1.87 $m^3$and Wick freshwater costs \$3.11 m^3, which reflects the higher economic and operational costs of the UK. Capital costs are higher due to more expensive equipment and installation procedures. Labor costs, including skilled technicians and maintenance, are costly in UK, increasing the operational expenses. The UK also imposes stricter environmental, safety, and quality regulations, which lead to additional compliance costs that must be factored into the project budget. The case study at Libya, they used large solar collectors with an area of 72m^2, which costs \$6600 each, whereas, I used solar collectors with an areas of 10.1m^2, which costs \$635 each. Despite, the lower cost of solar collectors, the cost of freshwater production calculated was higher than the results retrieved from Muftah et al [9]. In Libya, the solar irradiance is much higher compared to UK, so each solar collector in Libya can generate more thermal energy efficiently, reducing the numbers of collectors needed per unit of freshwater produced. Whereas, in UK, lower solar irradiance, requires more solar collectors to generate the same amount of thermal energy which increases the operational and capital costs leading to higher freshwater cost. To demonstrate the validity of the current simulation, the model was compared against the solar-MED case study in Libya by Muftah et al. [11] and the wind-MED study by Khalilzadeh and Hossein Nezhad [10]. While the calculated costs for Watchet of \$1.89/m^3, and Wick of \$3.12/m^3, are higher than the Libyan result of \$0.97/m^3, this difference validates the model's sensitivity to geographic and economic factors rather than indicating an error. Reassessing the resource requirements reveals that in Libya, the higher solar irradiance means fewer solar collectors are required to produce the same amount of freshwater, in contrast, the UK it has lower solar irradiance which requires 7, 255 collectors in Watchet and 13, 055 in Wick to meet the 12 kg/s demand, resulting in a TCC ranging from \$6.6M to \$11.9M. Furthermore, the model's validity is confirmed in the wind industry, where Wick's \$5.56/m^3, cost is far less than the \$16.68/m^3, reported in research by Khalilzadeh [10]. The switch from a 6-effect MED system used by Khalilzadeh [10] to the more effective 10-effect MED design used in this work, which achieves a higher Gained Output Ratio (GOR) and increased waste heat utilization, provides technical backing for this improvement. Using a 25-year lifecycle and a CRF of 0.0858, the model correctly shows that although the UK has higher initial costs, long-term payback of high-efficiency systems makes renewable water production potentially cost-effective.

4.1.2 Comparison of Wind Powered freshwater production cost between Simulated Results and Case Study

The freshwater cost powered by Wind Energy in Watchet is \$20.49 m^3 and in Wick it costs \$5.56 m^3 whereas, the Khalilzadeh and Hossein Nezhad [10] case study produces freshwater powered by Wind Energy which costs \$16.68m^3. The cost of freshwater in Watchet was 23% higher than the case study which can be due to less consistent wind in Watchet compared to Wick and case study location. This shows that more turbines were required to produce the same amount of thermal energy, increasing the TCC and AOC of Watchet. The cost of freshwater in Wick was cheaper than the case study as my MED model uses 10 effects, which may have caused to achieve a higher thermal efficiency compared to the 6 effect MED process used in the case study. Adding more effects in MED desalination improves the recovery ratio and reduces specific energy consumption, which allows more freshwater to be produced per unit of energy. This increased efficiency can contribute to lowering the overall cost of freshwater in my project compared to the pilot plant. Plus, the higher number of effects allows better utilisation of available waste heat, reducing the operational costs. Plus, the wind speed is major factor, as higher wind speed leads higher power output, which leads to more waste heat available for the MED process. Wick has lower cost of freshwater compared to the case study is because, the case study Muftah et al [11], has a short payback period of 6.33 years, which requires the initial investment of the capital cost to be recovered over a small period. This results in a higher annualised cost, which forces the freshwater cost to be higher because each year the cost of freshwater has to cover large fraction of the initial investment.

5. CONCLUSIONS

This research aligned with the inclination of transition from fossil fuels for seawater desalination into Renewable energy such as Solar and Wind energy. To conduct this research, specific models of MED, Solar and Wind energy were developed. Specifically, the freshwater production costs of MED powered by Solar and Wind energy were estimated for Watchet and Wick cities in the UK with a comparison against the freshwater production cost to MED powered by Conventional Steam boiler. The results demonstrated that using Watchet and Wick dataset of solar irradiance and wind speed, both energy sources were able to produce thermal energy required to heat the HTF of MED process. However, Solar energy for MED process is more efficient in Watchet, whereas Wind energy for MED process is more effective in Wick. Watchet is located South of England where there is high solar irradiance. Whereas, Wick is located in North of

Scotland which has a climate of strong and consistent wind speeds throughout the year and has low solar irradiance and is mainly cloudy especially in winter season. By using the same 750 kW turbine for both Watchet and Wick, we can ensure that the cost difference is caused strictly by the wind speeds, not the infrastructure. This demonstrates that Wick is more cost-effective due to its natural resources rather than only having a better infrastructure. It simplifies the study and keeps the focus on how location affects the cost of freshwater.

REFERENCES

1. Al-hotmani OMA, Al-Obaidi MA, Li JP, John YM, Patel R, Mujtaba IM. A multi-objective optimisation framework for MED-TVC seawater desalination process based on particle swarm optimisation. Desalination 525:115504 (2022). https://doi.org/10.1016/j.desal.2021.115504
2. Al-hotmani OMA, Al-Obaidi MA, John YM, Patel R, Manenti F, Mujtaba IM. Minimisation of energy consumption via optimisation of a simple hybrid system of multi effect distillation and permeate reprocessing reverse osmosis processes for seawater desalination. Computers & Chemical Engineering 148:107261 (2021). https://doi.org/10.1016/j.compchemeng.2021.107261
3. Al-Obaidi MA, Filippini G, Manenti F, Mujtaba IM. Cost evaluation and optimisation of hybrid multi effect distillation and reverse osmosis system for seawater desalination. Desalination 456:136-149 (2019). https://doi.org/10.1016/j.desal.2019.01.019
4. Allouhi A, Almohammadi KM. Towards green desalination: a multi-site analysis of hybrid renewable energy integration in saudi arabian RO plants. Desalination 592:118087 (2024). https://doi.org/10.1016/j.desal.2024.118087
5. Blanco, J. and Alarcón, D. (2005) Improving the efficiency of high capacity solar thermal seawater desalination system.
6. Bosch (2025) Steam boiler UL-S - world's best selling Bosch boiler.https://www.boschindustrial.com/gb/en/ocs/commercial-industrial/universal-steam-boiler-ul-s-ul-sx-669474-p/
7. Dilli B, Kumar SVSP. Solar parallel feed multi effect distillation plant: an experimental study. Desalination and Water Treatment 317:100147 (2024). https://doi.org/10.1016/j.dwt.2024.100147
8. Enduramaxx (2025) *7, 000 Litre Insulated Water Tank - Low Profile*. https://enduramaxx.co.uk/enduramaxx/7000-litre-vertical-insulated-water-tank-with-25mm-insulation/?attribute_insulation-thickness=25mm
9. Filippini G, Al-Obaidi MA, Manenti F, Mujtaba IM. Design and economic evaluation of solar-powered hybrid multi effect and reverse osmosis system for seawater desalination. Desalination 465:114-125 (2019). https://doi.org/10.1016/j.desal.2019.04.016
10. Khalilzadeh S, Hossein Nezhad A. Utilization of waste heat of a high-capacity wind turbine in multi effect distillation desalination: energy, exergy and thermoeconomic analysis. Desalination 439:119-137 (2018). https://doi.org/10.1016/j.desal.2018.04.010
11. Muftah AK, Zili-Ghedira L, Abugderah MM, Hassen W, Becheikh N, Alshammari BM, Kolsi L. Sustainable water production: solar energy integration in multi-effect desalination plants. Water 17:647 (2025). https://doi.org/10.3390/w17050647
12. Rostami S, Ghiasirad H, Rostamzadeh H, Kalan AS, Maleki A. A wind turbine driven hybrid HDH-MED-MVC desalination system towards minimal liquid discharge. South African Journal of Chemical Engineering 44:356-369 (2023). https://doi.org/10.1016/j.sajce.2023.03.007
13. Wagner & Co (n.d.) *LBM Solar Collector*. https://downloadcenter.wagner-solar.com/download.php?c_download=3010&c_event=download
14. Wind--Turbine-Models (2025) *Page Restricted*. https://en.wind-turbine-models.com/turbines/2358-wind-world-w5200-750
15. Zheng Y, Caceres Gonzalez RA, Hatzell KB, Hatzell MC. Large-scale solar-thermal desalination. Joule 5:1971-1986 (2021). https://doi.org/10.1016/j.joule.2021.07.005

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Dynamic Modelling of Renewable-driven CO_2 Methanation using Recurrent Neural Networks

M. Andrea Pappagallo[ab], Diego A. Romero Lombo[c], Mattia Vallerio[c], Emanuele Moioli[ac]*

[a] Paul Scherrer Institut, Center for Energy and Environmental Sciences, Villigen, Switzerland
[b] École Polytechnique Fédérale de Lausanne, Institute of Chemical Science and Engineering, Lausanne, Switzerland
[c] Politecnico di Milano, Dipartimento di Chimica, Materiali e Ingegneria Chimica 'Giulio Natta', Milano Italy
* Corresponding Author: emanuele.moioli@polimi.it.

ABSTRACT

A recurrent neural network (RNN) model for a CO_2 methanation reactor was developed based on synthetic data generated from a validated mechanistic model of the same unit. The model was used to predict the main properties of the reactor – methane productivity and hotspot temperature – during a dynamic operation of the unit. The dynamic profile of feedstock availability was simulated taking into account the H_2 flow that can be produced from PV-powered water electrolysis using solar irradiation profiles over a year in Milan, Italy. The dataset therefore consists of 366 data instances (one for each day), each composed of one datapoint per minute of sunlight. The best agreement between the predictions from the RNN and the target output values from the mechanistic model was found using a shallow RNN of 20 hidden-layer neurons, trained with a batch size of 10 and an 80/20 training-testing split. This showed that RNNs can constitute a reliable tool for dynamic surrogate modelling of energy conversion reactors. The surrogate model was embedded in an energy system model to provide dynamic predictions of the energy conversion efficiency in the reactor, providing a more realistic performance assessment compared to the average efficiency models commonly used in the literature.

Keywords: Dynamic surrogate modelling, Recurrent neural networks, Machine learning, Energy systems analysis, Chemical reaction engineering

INTRODUCTION

The global transition towards decarbonised energy systems is reshaping the design and operation of chemical reactors within integrated energy networks. As renewable technologies become central to energy production, their inherent intermittency introduces temporal fluctuations that destabilise the conventional steady-state operation of chemical processes [1]. This shift demands new modelling approaches capable of capturing transient dynamics to ensure process efficiency, stability, and safety under variable operating conditions.

Among the technologies driving decarbonisation, the catalytic methanation of carbon dioxide (CO_2) has gained relevance within Power-to-Gas (PtG) systems. This process converts renewable hydrogen and captured CO_2 into synthetic natural gas (SNG), which can be injected into existing gas grids or used as a carbon-neutral fuel [2]. By enabling renewable energy storage and carbon utilisation, CO_2 methanation provides a key link between the renewable power and gas sectors. However, the highly exothermic nature of the reaction ($\Delta H_{298K} = -165.0 \frac{kJ}{mol}$) poses challenges for heat management and reactor stability, particularly under the intermittent and variable conditions imposed by renewable power fluctuations [3]. These dynamic changes in feed composition, flowrate, and temperature cannot be reliably represented through steady-state modelling approaches.

In this context, dynamic reactor modelling has emerged as a critical tool for understanding and optimising CO_2 methanation under varying conditions. By explicitly accounting for temporal variations in system inputs, dynamic modelling enables the prediction of phenomena such as transient temperature evolution, heat accumulation, and reaction rate fluctuations, providing the founda-

tion for safe and efficient reactor operation in non-stationary regimes [4]. Nevertheless, traditional mechanistic models, despite their physical accuracy, are computationally demanding due to the need to solve complex coupled differential equations across multiple spatial and temporal scales.

To overcome these limitations, recent advances in machine learning have opened new pathways for developing data-driven models that replicate mechanistic behaviour with significantly lower computational cost [5]. Recurrent neural networks (RNNs), in particular, have demonstrated great potential for modelling sequential and time-dependent chemical processes by capturing the effect of dynamic inputs on dynamically operated reactors [6].

This study develops and applies an RNN model trained on data from a validated mechanistic simulation to reproduce the dynamic behaviour of a CO_2 methanation reactor fed with hydrogen produced via photovoltaic-powered electrolysis. The main objective is to quantify the impact of seasonal renewable fluctuations on reactor performance and stability, providing a computationally efficient modelling framework to support future Power-to-Gas technologies.

While the tool of RNNs has been used in the literature to construct hybrid, discrepancy-based models [6, 7], i.e., dynamic mechanistic models where the machine learning tool is used to bridge the gap between the first-principles equations and the real-life phenomenon, this work aims at using the machine learning tool to construct a fully black-box model for the dynamic calculation of key performance variables in a chemical reactor as a function of present and past values of one or more selected input variables. This with the aim to realise a model as efficient as possible for the design of an energy conversion reactor capable of dealing with non-constant input quantities.

METHODS

The methodology adopted in this study consists of three main steps. First, a mechanistic reactor model was developed and validated to generate a synthetic dataset to support the machine learning model development. Subsequently, a surrogate dynamic model based on recurrent neural networks was designed and trained using this dataset. Finally, the model outputs were integrated into a black-box CO_2 methanation model to assess the dynamic behaviour of a PtG system and evaluate the performance and applicability of the predicted results.

Mechanistic reactor modelling

The kinetic expression for the CO_2 methanation reaction is derived from the work of Pétremand *et al.* (2025) [8]. Additionally, the behaviour of the system around the thermodynamic equilibrium is characterised by adding two kinetic expressions for the methanation of CO and the water-gas shift reaction from the work of Xu and Froment (1989) [9]. The rates of these reactions are only significant when the composition of the system is slightly different from that at the thermodynamic equilibrium, but not enough for the rate of CO_2 methanation to bridge such gap on its own. For the rest of the integration, their influence is negligible.

Table 1: Reactions in the kinetic model for CO methanation.

Reaction	Equation
CO_2 methanation	$CO_2 + 4H_2 \rightleftharpoons CH_4 + 3H_2O$
CO methanation	$CO + 3H_2 \rightleftharpoons CH_4 + 2H_2O$
Water-gas shift	$CO + H_2O \rightleftharpoons CO_2 + H_2$

The CO_2 methanation reactor considered in this study is a plate-type unit operating at 8 bar and thermally regulated with boiling water at 230 °C. The feed consists of biogas with a CO_2:CH4 ratio of 1:1, and hydrogen is supplied to achieve an inlet molar H_2/CO_2 ratio of 4. The dynamic behaviour of the reactor is represented using a one-dimensional dynamic model in which mass balances are pseudohomogeneous, while heat balances are heterogeneous.

The gas-phase component balance is expressed in Equation (1):

$$\frac{\partial F_i}{\partial t} + v_g \frac{\partial F_i}{\partial z} = v_g A \mathcal{D}_i \frac{\partial^2 c_i}{\partial z^2} + \rho_c (1-\varepsilon_b) v_g A \sum_{j=1}^{N_R} \nu_{ij} r_j \quad (1)$$

The energy balance equations for the gas and solid phases, respectively, are expressed in Equations (2) and (3):

$$\rho_g \hat{c}_{P_{mix}} \left(\frac{\partial T}{\partial t} + v_g \frac{\partial T}{\partial z}\right) = v_g A \frac{\partial}{\partial z}\left(\Lambda \frac{\partial T}{\partial z}\right) - v_g a_s h_s (T - T_s) - v_g a_e U (T - T_e) \quad (2)$$

$$(1-\varepsilon_b)\rho_s c_s \frac{\partial T_s}{\partial t} = -(1-\varepsilon_b)\rho_s \sum_{j=1}^{N_R} \Delta h_j^R r_j + v_g a_s h_s (T - T_s) \quad (3)$$

At the reactor inlet, the Danckwerts boundary conditions are applied, while zero-gradient Neumann conditions are imposed at the outlet by setting the second derivative of concentration and temperature to zero.

Hydrogen production from photovoltaic-powered electrolysis is determined using the efficiency η_{el} of the proton exchange membrane (PEM) electrolyser, defined as:

$$F_{H_2} = \frac{\eta_{el} W_{el}}{HHV_{H_2}} \quad (4)$$

In this formulation, the electrical power W_{el} corresponds to the installed photovoltaic capacity derived

from satellite observations for Milan, Italy, and an efficiency of 75% was assumed for the electrolyser.

Artificial neural network development

An artificial neural network (ANN) is a machine learning model that processes information through interconnected computational units known as neurons. Each neuron processes input signals (x), turning them into outputs (y) through two steps:

1. A linear combination of weights W and biases b;
2. A nonlinear activation function θ.

In the simplest ANN architecture, called a feed-forward neural network (FFNN), the neurons are organised in subsequent layers. The inputs to the network are processed by an input layer, one or more hidden layers, and one output layer [10].

The network parameters are initialised randomly, then updated iteratively to minimise the error between the actual and predicted outputs, in a procedure called training. Each iteration of the procedure is called epoch, and the training ends when a sufficient agreement between the training outputs and the network prediction is reached.

This study employs a recurrent neural network (RNN), which extends the conventional feedforward architecture by incorporating feedback connections to process sequential data and capture temporal dependencies relevant to reactor dynamics. A gated recurrent unit (GRU) network is used, as its gating mechanisms preserve information over longer time sequences through the hyperbolic tangent ($\tanh$) and the sigmoid (σ) activation functions [11]. The GRU is defined by its gating mechanism, which at each time step k computes the reset gate $\boldsymbol{r}^{(k)}$, update gate $\boldsymbol{z}^{(k)}$, candidate gate $\boldsymbol{g}^{(k)}$, and the hidden state $\boldsymbol{h}^{(k)}$ (in the following equations, $\circ$ represents element-wise multiplication):

$$\boldsymbol{r}^{(k)} = \sigma\left(\boldsymbol{W}_r\boldsymbol{x}^{(k)} + \boldsymbol{R}_r\boldsymbol{h}^{(k-1)} + \boldsymbol{b}_r\right) \tag{5}$$

$$\boldsymbol{z}^{(k)} = \sigma\left(\boldsymbol{W}_z\boldsymbol{x}^{(k)} + \boldsymbol{R}_z\boldsymbol{h}^{(k-1)} + \boldsymbol{b}_z\right) \tag{6}$$

$$\boldsymbol{g}^{(k)} = \tanh\left(\boldsymbol{W}_g\boldsymbol{x}^{(k)} + \boldsymbol{r}^{(k)} \circ \left(\boldsymbol{R}_g\boldsymbol{h}^{(k-1)} + \boldsymbol{b}_g\right)\right) \tag{7}$$

$$\boldsymbol{h}^{(k)} = \left(\boldsymbol{1} - \boldsymbol{z}^{(k)}\right) \circ \boldsymbol{g}^{(k)} + \boldsymbol{z}^{(k)} \circ \boldsymbol{h}^{(k-1)} \tag{8}$$

Afterwards, the GRU hidden layer is linked to a fully connected output layer defined by its own weight vector and bias, yielding the network output $\boldsymbol{y}^{k}$:

$$\boldsymbol{y}^{(k)} = \boldsymbol{W}_y\boldsymbol{h}^{(k)} + \boldsymbol{b}_y \tag{9}$$

The RNN tool was used to produce a surrogate, data-driven model of the reactor where the only input feature is the hydrogen inlet flowrate, while the CO_2 is assumed to always be fed in stoichiometric amounts according to the Sabatier reaction (CO_2 methanation, Table 1). As output features, three performance variables for the methanation reactor were chosen in the hotspot temperature, the methane output flowrate, and the methane purity at the outlet.

As this study assesses a regression task, the predicted outputs are assessed against the target values using the mean squared error (MSE), which quantifies the average squared difference between the predictions y_i and target values $\hat{y}_i$:

$$\mathcal{L}_{MSE} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \tag{10}$$

Integration into an energy system model

The trained RNN was integrated into a model for the design of a system to provide a constant supply of H_2 to satisfy the yearly energy demand from fuel-cell-powered vehicles. In this model, it is assumed that H_2 is primarily used in decarbonised mobility. When H_2 is available in excess, it is stored in the form of methane. Then, methane is reformed to produce hydrogen when H_2 supply is not sufficient [12] The energy demand is assumed as a problem datum and marked as W_{req}. The model equates W_{req} with the energy converted from solar to calculate the minimum PV panel surface required to meet the demand.

The energy from the PV panels is stored and converted according to three steps:

1. **electrolysis** to produce H_2,
2. **methanation** to convert the H_2 into CH_4,
3. **reforming** to reconvert the methane to hydrogen for use on vehicles.

Each of these steps is characterised by an efficiency value, resulting in the following equation for the conversion of the available energy W_{av}:

$$W_f = W_{av} \bullet \eta_{el} \bullet \eta_{meth} \bullet \eta_{ref} \tag{11}$$

While the efficiencies connected to the electrolysis and reforming steps are assumed constant ($\eta_{el} = 0.75, \eta_{ref} = 0.7$), the value of η_{meth} can either be a constant empirical value of 0.7, or calculated for each day individually with the help of the RNN model (in which case it is denoted by η_{RNN}):

$$\eta_{RNN} = \frac{F_{CH_4}^{out} \bullet HHV_{CH_4}}{F_{H_2}^{in} \bullet HHV_{H_2}} \tag{12}$$

The available energy for conversion is determined from the solar irradiation intensity I in the dataset:

$$W_{av} = I \bullet S_{PV} \tag{13}$$

The equivalence between W_{req} and W_f is then used to determine the required PV panel surface S_{PV}. Results are compared between the constant assumption of $\eta_{meth} = 0.7$ and the average of the daily values derived from the RNN model (η_{RNN}).

RESULTS AND DISCUSSION

Mechanistic reactor model validation

The mechanistic model was validated by reproducing the transient and steady-state behaviour of the reactor reported by Moioli *et al.* (2024) [13]. The transient thermal behaviour of the CO_2 methanation reactor was evaluated by analysing the temperature profiles every second minute during start-up. The modelling case used for the validation shown in Figure 1 has been obtained at a coolant temperature of 230°C, reactor pressure of 8 bar, and with an inlet composed of hydrogen and biogas in stoichiometric ratio, with the molar ratio of CO and CH_4 in the biogas assumed to be 1/1. The case in the figure refers to a step change in the inlet flowrate from 0% to 50% of the nominal design conditions of the reactor (3000 h^{-1}).

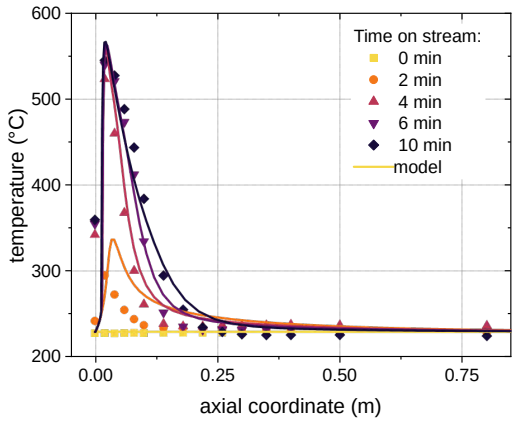

Figure 1: Temperature profile of the reactor during start-up at 50% load. Continuous lines represent the model results, scattered points the experimental data from [13].

At time-zero, the catalyst bed is preheated to 230°C. Upon feeding of CO_2, the reaction is initiated immediately, generating a hotspot whose axial width and peak temperature increase over time, reaching steady-state in under ten minutes. The mechanistic model successfully reproduces the start-up dynamics, closely matching the experimental observations. The discrepancies between the experimental and modelled profiles can be attributed to the difference between the catalysts used in the pilot-scale plant and in the work of Pétremand *et al.* [8].

The mechanistic model was subsequently validated through four simulations of the CO_2 methanation reactor, assessing its response to load variations, temperature evolution, concentration profiles, and axial temperature behaviour under different feed gas temperatures. In all cases, the reactor exhibited rapid thermal adaptation to changing operating conditions, and the model successfully reproduced key experimental reaction dynamics, including hotspot development and the attainment of grid-compliant methane. Overall, the simulated responses demonstrate the reliability and flexibility of the mechanistic model, confirming its suitability for generating representative training data.

RNN development and performance evaluation

The RNN was trained using three target variables: hotspot temperature, methane dry molar fraction, and methane molar flow. These descriptors were selected as key indicators because they represent the thermal behaviour of the reactor, product quality, and productivity under fluctuating hydrogen supply. To ensure balanced learning across variables of different magnitudes, a standardisation procedure that transforms all target variables into a comparable numerical range was applied:

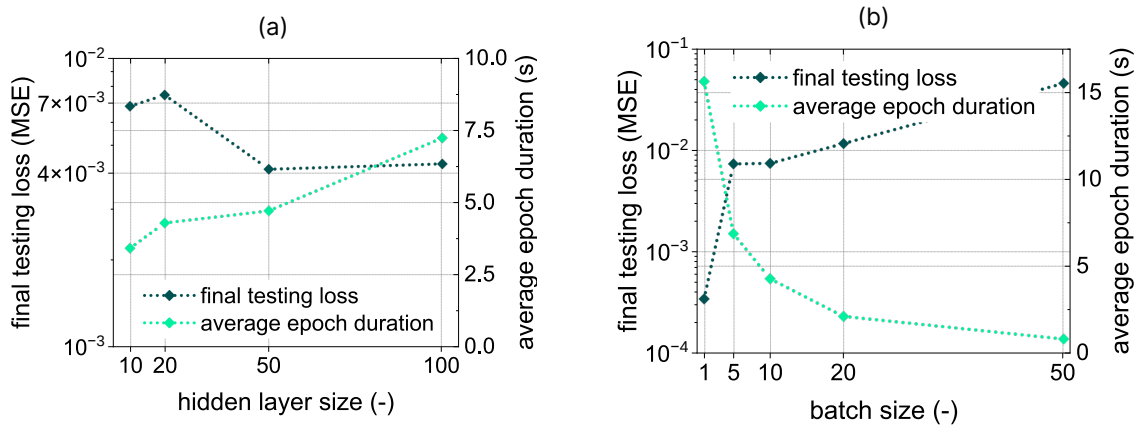

Figure 2: Effect of the number of neurons in the hidden layer (a) and of the training batch size (b) on the final performance of the RNN on the testing dataset – measured by mean squared error – and on the time required for the training.

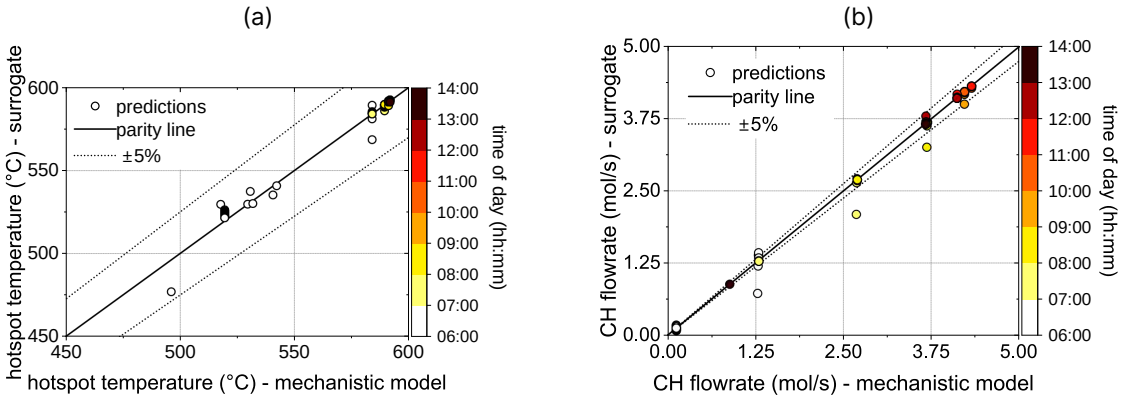

Figure 3: Parity plot for the hotspot temperature [°C] (a) and for the methane flowrate [mol/s] (b), comparing the values predicted by the RNN-based model (y axis) with those from the mechanistic model (x axis).

$$x_{i,sc} = \frac{(x_i - \mu_i)}{\sigma_i} \tag{14}$$

Where μ_i is the mean and σ_i the standard deviation of feature i across all datapoints. The scaling of the dataset allowed for stable convergence and close agreement between the predictions and the mechanistic model.

A sensitivity analysis was then performed to evaluate the influence of the hidden layer size, and batch size on model performance, whose results are reported in Figure 2. An increase in the hidden layer size did not lead to a clear improvement in the performance, but rather to a slight increase in the time needed for training. In machine learning theory, it is well established that increasing the number of neurons in hidden layers generally improves the accuracy of a neural network, but also that an excessive parametrisation may increase the risk of overfitting [6, 11]. The layer size for the following evaluations was therefore limited to 20 as a compromise to achieve relative accuracy while maintaining a low risk of overfitting and a manageable training time.

The sensitivity analysis on the batch size showed a clear trade-off between computational time and prediction accuracy: larger batch sizes reduced epoch duration but increased the testing loss. A batch size of 10 was adopted for the following evaluations, due to lower values also increasing the risk of overfitting.

The parity plots presented in Figure 3 offer an effective visual assessment of the model performance by comparing predicted values with their corresponding reference data for both output variables: the hotspot temperature, in °C, and the methane output of the reactor, in mol/s. Again, a strong agreement is shown between the predictions from the RNN model and those from the mechanistic model.

The parity plots show most data points lying close to the parity line, indicating that the RNN accurately reproduces the nonlinear behaviour of the mechanistic model under dynamic conditions. The highest deviations appear for the methane output at the extremes of the daytime range (lighter yellow datapoints), reflecting a reduced accuracy when the reactor operates at very low space velocities. These discrepancies remain limited and do not significantly affect overall predictive performance.

The parity plots confirm that the RNN captures both the steady-state and transient responses of the CO_2 methanation reactor with high fidelity.

Figure 4 compares the entire daily profile of the outlet flowrate of methane predicted by the mechanistic ground truth model and the RNN-based model, respectively. The strong agreement shown in the parity plot in Figure 2b between the predictions by the RNN and the mechanistic model is again made visible, confirming the ability of the network to capture the variation of key performance variables following a change in the hydrogen supply to the CO_2 methanation reactor.

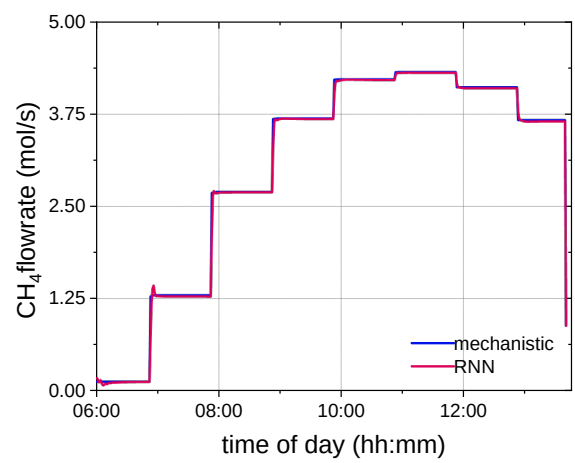

Figure 4: Comparison between the methane flowrate values [mol/s] predicted by the mechanistic model and the RNN-based model.

Table 2: Comparison of technical indicators obtained from the power-to-gas energy system using three modelling approaches.

	Constant efficiency	Average RNN efficiency
PV surface required $[\mathrm{m}^2]$	1001	989
Active days per year of the storage system $[\mathrm{d}]$	257	254
Required amount of energy carrier (CH_4) per year $[\mathrm{ton}]$	9.20	8.98
Required amount of biogas per year $\left[\frac{\mathrm{m}^3}{\mathrm{h}}\right]$	0.732	0.718

Integration into the energy system model

The recurrent neural network was integrated into the energy system model by applying the daily power-to-gas efficiency profile calculated in Equation (12). These daily efficiency values were incorporated into the macroscopic energy balance described in Equation (11), linking the dynamic reactor behaviour with the overall energy system evaluation. For comparison, the RNN results were assessed against those obtained using a constant efficiency value for the methanation block, as commonly done in energy system analysis [14]. The results are summarised in Table 2.

When comparing the constant efficiency evaluations with the ones conducted using a dynamic efficiency profile, some notable differences are found. First, the higher efficiencies calculated by the RNN model result in a 1.2% lower surface of PV panels needed to meet the energy demands, and a lower need for the energy storage system, which is used 3 days fewer than what calculated using the standard efficiency value.

This reduces the amount required over a year for both the energy carrier (-2.4%) and the biogas (-1.9%), leading to a more efficient use of resources.

These results show that integrating a dynamic efficiency profile, calculated with the help of a model, directly improves the energy balance and storage behaviour of the system model, whereas assuming constant efficiencies fails to capture the daily variations in performance. Including these variations provides a more realistic representation of the response of the PtG system to the intermittency typical of renewable energy sources.

CONCLUSIONS

This study developed and evaluated a recurrent neural network (RNN) framework for the dynamic modelling of a CO_2 methanation reactor powered by renewable hydrogen. The research aimed to bridge the gap between detailed mechanistic modelling, which provides high physical fidelity but is computationally intensive, and data-driven modelling, which enables fast, flexible, and adaptive simulations. The methodology consisted of three main stages: mechanistic model validation, RNN development, and a system evaluation within a Power-to-Gas (PtG) configuration.

In the first stage, a one-dimensional mechanistic reactor model using pseudohomogeneous component balances and heterogeneous energy balances was developed and validated with established CO_2 methanation kinetics. The model successfully reproduced the steady-state and transient responses reported in the literature [13], including the rapid thermal adaptation to load changes in the reactor, the evolution of the hotspot, and the corresponding concentration profiles. This validation confirmed the model as a reliable tool generating dynamic data representative of operating conditions influenced by fluctuating hydrogen inputs.

The second stage focused on developing and training the RNN using synthetic data produced by the mechanistic simulations. The hotspot temperature, methane dry molar fraction, and methane molar flow were selected to represent the thermal behaviour of the reactor, product quality, and productivity. Standardisation was essential to achieve balanced learning and stable convergence. A sensitivity analysis of network architecture showed that an increase in parametrisation (hidden layer size, network depth) had a limited influence on predictive accuracy, whereas the evaluation of the batch size revealed a clear trade-off between computational time and model performance. The final architecture, consisting of one hidden layer with 20 neurons and trained using a batch size of 10, provided the best balance between accuracy and efficiency. The trained RNN reproduced the nonlinear reactor dynamics with high fidelity, as demonstrated by strong agreement between predicted and mechanistic model outputs.

In the third stage, the RNN was integrated into a black-box PtG model to evaluate the system performance under daily and seasonal variability in hydrogen production from photovoltaic electrolysis. Compared to a constant efficiency assumption, the inclusion of dynamic daily efficiencies significantly altered the predicted behaviour of the chemical storage system. The dynamic case showed substantial decreases in the operating days needed to meet the energy demand, with a consequent decrease in the usage of methane and biogas, and in the PV panel surface required, demonstrating that simplified

efficiency assumptions notably overestimate the operational demands of the system.

Overall, this work demonstrates that RNNs offer an effective and computationally efficient surrogate for dynamic reactor modelling. The framework developed here constitutes a practical methodological approach that combines the accuracy of mechanistic models with the speed of data-driven ones, supporting future efforts in developing efficient optimisation and control strategies for flexible, renewable-driven chemical processes.

ACKNOWLEDGEMENTS

The authors acknowledge the support from the Swiss National Science Foundation (SNSF) Ambizione project: 'Moving catalyst vs. multi-catalyst: determination of the best reactor for the processing of unconventional feedstock' with grant number No. 20912.

TABLE OF VARIABLES

Mechanistic model

A	reactor cross-section	m^2
a_s	specific surface for interphase heat exchange	m^{-1}
c_i	concentration of species i	mol/m^3
$\hat{c}_{P_{mix}}$	mass-specific heat of gas mixture	J/kg•K
c_s	specific heat of solids	J/kg•K
$\mathcal{D}_i$	axial dispersion coefficient of species i	m^2/s
F_i	molar flowrate of component i	kmol/s
HHV_i	higher heating value of species i	J/mol
h_s	interphase heat exchange coefficient	$W/m^2 \cdot K$
$\dot{m}$	total mass flowrate	kg/s
r_j	rate of reaction j	$kmol/m^3 \cdot s$
T	temperature	K
T_e	coolant temperature	K
T_s	temperature of solid phase	K
t	time	s
U	global heat transfer coefficient	$W/m^2 \cdot K$
v_G	gas velocity	m/s
W_{el}	available electrical power	W
z	axial reactor coordinate	m
Δh_j^R	enthalpy of reaction j	J/mol
ε_b	bed void fraction	-
η_{el}	electrolyser efficiency	-
ν_{ij}	stoichiometric coefficient of species i in reaction j	-
ρ_C	catalyst density	kg/m^3

Surrogate model

$\boldsymbol{b}_g$	bias of gate g/output layer y
$\boldsymbol{g}^{(k)}$	candidate gate output at timestep k
$\boldsymbol{h}^{(k)}$	network hidden state at timestep k
$\boldsymbol{R}_g$	recurrent weights of gate g
$\boldsymbol{r}^{(k)}$	reset gate output at timestep k
$\tanh(\bullet)$	hyperbolic tangent activation function
$\boldsymbol{x}^{(k)}$	network input at timestep k
$\boldsymbol{y}^{(k)}$	network output at timestep k
$\boldsymbol{W}_g$	input weights of gate g/output layer y
$\boldsymbol{z}^{(k)}$	update gate output at timestep k
$\sigma(\bullet)$	sigmoid activation function

Energy system model

F_i^{in}	inlet flowrate of component i	kmol/s
F_i^{out}	outlet flowrate of component i	kmol/s
I	solar intensity	W/m^2
S_{PV}	total surface of PV panels	m^2
W_{av}	available power from PV plant	W
W_f	enthalpy content of final product	W
η_{el}	electrolyser efficiency	-
η_{meth}	thermal efficiency of methanation (constant hypothesis)	-
η_{RNN}	thermal efficiency of methanation (from surrogate model)	-

REFERENCES

1. Cegla M, Semrau R, Tamagnini F, Engell S. Flexible process operation for electrified chemical plants. Current Opinion in Chemical Engineering 39:100898 (2023). https://doi.org/10.1016/j.coche.2023.100898
2. Koschany F, Schlereth D, Hinrichsen O. On the kinetics of the methanation of carbon dioxide on coprecipitated nial(o). Applied Catalysis B: Environmental 181:504-516 (2016). https://doi.org/10.1016/j.apcatb.2015.07.026
3. Kreitz B, Wehinger GD, Turek T. Dynamic simulation of the CO2 methanation in a micro-structured fixed-bed reactor. Chemical Engineering Science 195:541-552 (2019). https://doi.org/10.1016/j.ces.2018.09.053
4. Tauer G, Kern C, Jess A. Transient effects during dynamic operation of a wall?cooled fixed?bed reactor for co₂methanation. Chem Eng & Technol 42:2401-2409 (2019). https://doi.org/10.1002/ceat.201900367
5. Schweidtmann AM, Esche E, Fischer A, Kloft M, Repke J, Sager S, Mitsos A. Machine learning in chemical engineering: a perspective. Chemie Ingenieur Technik 93:2029-2039 (2021).

https://doi.org/10.1002/cite.202100083
6. Wu Z, Rincon D, Christofides PD. Process structure-based recurrent neural network modeling for model predictive control of nonlinear processes. Journal of Process Control 89:74-84 (2020). https://doi.org/10.1016/j.jprocont.2020.03.013
7. Zhang D, del Río Chanona EA. Machine learning and hybrid modelling for reaction engineering. Royal Society of Chemistry (2023). https://doi.org/10.1039/9781837670178
8. Pétremand H, Witte J, Kröcher O, Moioli E. Machine learning-based kinetic modeling of the co₂ methanation reaction over an industrial catalyst. Ind. Eng. Chem. Res. 64:19864-19875 (2025). https://doi.org/10.1021/acs.iecr.5c02947
9. Xu J, Froment GF. Methane steam reforming, methanation and water-gas shift: I. Intrinsic kinetics. AIChE J 1989;35:88–96.
10. Yang L, Padding JT, Kuipers JAM. Two-fluid modelling for cylindrical fluidized beds using kinetic theory for rough spheres. Particuology 50:25-34 (2020). https://doi.org/10.1016/j.partic.2019.05.006
11. Zarzycki K, ?awry?czuk M. LSTM and GRU neural networks as models of dynamical processes used in predictive control: a comparison of models developed for two chemical reactors. Sensors 21:5625 (2021). https://doi.org/10.3390/s21165625
12. Madi H, Schildhauer T, Moioli E. Comprehensive analysis of renewable energy integration in decarbonised mobility: leveraging power-to-x storage with biogenic carbon sources. Energy Conversion and Management 321:119081 (2024). https://doi.org/10.1016/j.enconman.2024.119081
13. Moioli E, Senn P, Østrup S, Hütter C. Results from the operation of an efficient and flexible large-scale biogas methanation system. Energy Adv. 3:131-142 (2024). https://doi.org/10.1039/d3ya00436h
14. Elhaus N, Volkmann M, Kolb S, Schindhelm L, Herkendell K, Karl J. Techno-economic evaluation of anaerobic digestion and biological methanation in power-to-methane-systems. Energy Conversion and Management 315:118787 (2024). https://doi.org/10.1016/j.enconman.2024.118787

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Chemical Language Transformers for the Inverse Design of Novel Surfactants

Alexander W. Rogers[a], Ruediger Zillmer[b], Amanda Lane[b], Adam Kowalski[b] and Dongda Zhang[ab]*

[a] The University of Manchester, Department of Chemical Engineering, Manchester, UK
[b] Unilever, R&D Port Sunlight, Liverpool, UK
* Corresponding Author: dongda.zhang@manchester.ac.uk.

ABSTRACT

Rapid, sustainable redesign of large functional molecules demands efficient exploration of vast chemical spaces. Chemical language models (CLMs), especially transformers, can learn long-range structure-property relationships and enable fast candidate generation after training. However, inverse molecular design is ill-posed – many structures can meet the same target – and conditioned generation often decodes to invalid or off-spec molecules. To address this challenge, we propose a CLM-based inverse design framework that optimises latent representations toward target properties and explicitly evaluates round-trip fidelity, i.e., whether decoded candidates remain on-target after decoding and forward re-evaluation. To improve reliability, we introduce post-decoding beam re-ranking using round-trip consistency and a predictor-guided minimal-edit repair step that corrects invalid near-misses while preserving closeness to the target property. We demonstrate the approach on surfactant critical micelle concentration (CMC) design, benchmarking existing large pretrained CLMs against our lightweight domain-trained CLM. The framework produces a high proportion of valid and diverse molecules (~90%) while maintaining target property error near 1%. Moreover, atom-level saliency analysis confirms that the generated structures follow established surfactant design rules, supporting interpretable structure–property control. Overall, the framework provides an efficient and broadly applicable solution to reliable inverse design of novel functional molecules.

Keywords: Inverse molecular design, chemical language models, transformers, interpretable AI, surfactants

1. INTRODUCTION

Across home and personal care, food, and speciality chemicals, many formulation industries face a portfolio-level redesign challenge driven by sustainability targets and tighter timelines [1]. Large fractions of core ingredient families – including surfactants – must be updated while preserving cost and performance. Yet developing new functional molecules remains slow and expensive: synthesis-test-validate cycles are labour-intensive, and even high-throughput workflows remain constrained by synthesis, characterisation, and safety testing [2]. Virtual screening can prioritise candidates, but for large, flexible amphiphiles, simple descriptor filters can be weakly informative while high-fidelity simulation becomes computationally intractable at scale [3]. This motivates tools that rapidly propose a small, diverse set of chemically valid, on-target candidates for portfolio-scale redesign.

Machine learning offers a route to accelerate both property prediction and molecular generation. In this work we use SMILES strings, which represent molecules as compact token sequences compatible with modern sequence models [4]. For the forward task, structure-property prediction is commonly addressed using descriptor/fingerprint-based regressors, kernel methods, graph neural networks, and other supervised deep learning models [5]. The reverse task – finding structures that meet target properties – is ill-posed: many distinct molecules can satisfy the same specification. Existing inverse-design approaches include brute-force screening, Bayesian optimisation over enumerated libraries and evolutionary-inspired mutations to strings or graphs [6, 7] but maintaining validity and reliable property control remains challenging over large combinatorial spaces.

To overcome these limitations, a more scalable approach is latent-space optimisation, where embeddings

represent molecules. A differentiable model learns to map the embeddings to the properties, and gradients are used to steer the embeddings toward the target property [8]. Autoregressive decoding generates a candidate token-by-token, each token conditioned on the previously generated tokens and the conditioning signal – a prompt or latent embedding – while beam search maintains the top-scoring partial sequences until an end-of-sequence token is produced. Sequence models, such as recurrent neural networks (RNNs) [9] and long short-term memory (LSTM) [10] networks, have been effective in de novo drug and organic molecule design. Transformers model long-range interactions via self-attention, which is valuable for large molecules whose properties depend on distributed motifs (e.g., head-tail interactions), and benefit from large-scale pretraining on unlabelled molecular structure corpora [11]. Chemical language transformers have been used in crystal property prediction [12], small-molecule inverse design and property prediction [13], and polymer inverse design and property prediction [14]. For surfactants, however, systematic transformer-based inverse design remains comparatively underexplored: recent studies use reinforcement learning [15] for inverse surfactant design and foundation models for phase-diagram prediction [16]. Overall, best practices for transformer-driven large-molecule inverse design remains an open area of research.

Most importantly, when a generator is conditioned – via a prompt or latent embedding – to meet a specified target, candidates must still satisfy the target after decoding and re-evaluation with the forward property predictor. Therefore, a core contribution of this research is to formalise and quantify this requirement as the round-trip fidelity. This behaviour is seldom reported in inverse-design studies [6] yet is critical for practical adoption, or else there is a risk of wasting experimental validation later. Accordingly, this work formalises, diagnoses, and mitigates round-trip fidelity failures for surfactant inverse design. We introduce two control strategies: (1) post-decoding re-ranking of beam candidates using round-trip consistency, and (2) predictor-guided minimal-edit repair that converts invalid near-miss SMILES into valid molecules while preserving property fidelity. We demonstrate the approach on a surfactant critical micelle concentration (CMC) design case study, benchmarking pretrained chemical language models and a lightweight domain-trained CLM, analysing novelty/diversity and atom-level saliency consistency with established surfactant structure–property design rules.

The remainder of this paper is organised as follows: Section 2 presents the model, latent optimisation, round-trip metrics, re-ranking, and repair; Section 3 describes the dataset and CMC case study; Section 4 reports the results; and Section 5 concludes with future directions.

2. METHODOLOGY

We first outline the backbone of our framework. A transformer-based encoder-decoder model was implemented with a property regressor to enable inverse molecular design. The encoder maps discrete molecular structures, represented as tokenised SMILES strings, to a continuous embedding space. A regressor maps this embedding to the molecular properties of interest, and the decoder reconstructs SMILES from the embeddings. Inverse design is performed by optimising over the latent embedding space to match the target property, decoding the optimised embedding to candidate SMILES, then applying round-trip validation, re-ranking, validity screening, and minimal-edit repair to improve reliability.

2.1. Constructing the Chemical Language Transformer Model

Building on this, our encoder-decoder transformer model was constructed for molecular representation and generation, trained from scratch on a large, unlabelled, augmented domain-specific dataset of molecule structures. The *HuggingFace* [17] *EncoderDecoderModel* class was used to build and manage the transformer architecture, enabling integration of our tokenisers, embedding layers, and downstream decoding pipelines.

SMILES strings are tokenised using byte-pair encoding trained on a large surfactant-like corpus (Section 3) to capture repeated chemical motifs and shorten sequences for long amphiphiles. Sequences are padded/truncated to a maximum length of 128. The model architecture consisted of two transformer stacks: a 2-layer encoder and a 2-layer decoder, each with 2 attention heads and an embedding size of 32. The number of transformer stacks (i.e., 8, 4, 2) and the embedding size (i.e., 128, 64, 32) were reduced until an elbow point was observed in SMILES reconstruction accuracy.

2.2. Property Regressor and Multitask Model Structure

Within this architecture, the encoder produces a contextual embedding that captures general chemical syntax, motifs and patterns. As illustrated in Figure 1, this shared representation then feeds the two heads of the multitask model structure: a decoder for SMILES reconstruction and a regressor that maps the embedding to the molecular properties of interest.

Let $\boldsymbol{Z} \in \mathbb{R}^{L \times d}$ denote the encoder output (contextual token embeddings). To predict the property of interest, $\boldsymbol{Z}$ is flattened to a vector $\boldsymbol{z} \in \mathbb{R}^{L \cdot d}$, concatenated with any relevant covariate conditions $\boldsymbol{c}$ (e.g., temperature) and then passed to a differentiable regressor $r([\boldsymbol{z}; \boldsymbol{c}]) \mapsto \hat{y}$. In this work, we use a linear regressor, while partial least squares (PLS) will later be used for compression and manifold-based optimisation (Section 2.6).

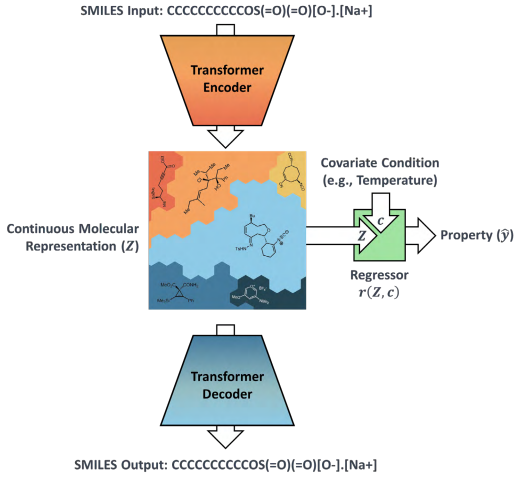

Figure 1. Multitask transformer architecture: a SMILES is encoded as an embedding that is then reconstructed or used to predict the molecular property under a covariate.

2.3. Transformer Pretraining and Multitask Finetuning

Training proceeded in two stages. The encoder–decoder was first pretrained on the large unlabelled surfactant-like corpus for SMILES reconstruction. It was then finetuned on a smaller labelled dataset using a multitask loss combining reconstruction and regression:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{recon}} + \lambda_{\text{reg}} \cdot \mathcal{L}_{\text{reg}} \tag{1}$$

Where $\mathcal{L}_{\text{recon}}$ is the cross-entropy loss between the predicted and ground-truth SMILES tokens, and $\mathcal{L}_{\text{reg}} = (\boldsymbol{y} - \hat{\boldsymbol{y}})^2$ is the mean squared error between the predicted, $\hat{\boldsymbol{y}}$, and true property, $\boldsymbol{y}$, values, and $\lambda_{\text{reg}} = 1$ is a weighting coefficient. Pretraining (~28 hours) and finetuning (~30 minutes) were performed on V100 (20 GB) and RTX 4060 (8 GB) GPUs, respectively.

2.4. Inverse Design via Latent Embedding Optimisation

Once the transformer-regressor was trained, we then performed inverse molecular design by optimising in the continuous embedding space to identify candidate molecules with a desired target property, as illustrated in Figure 2. Given a target value y_{target}, inverse design seeks an embedding, $\boldsymbol{Z} \in \mathbb{R}^{L \times d}$, whose predicted property matches the target, $y_{\text{target}} \approx r([\boldsymbol{z}; \boldsymbol{c}])$, as shown in Eqn. 2a. We use gradient-based optimisation and encourage diversity using a swarm of candidates $\{\boldsymbol{Z}^{(n)}\}$ optimised in parallel. Two penalties, $P(\boldsymbol{z})$, defined in Eqn. 2b were added to the total loss: (i) a quadratic repulsion term and (ii) a penalty on a point's squared Mahalanobis distance, D_M, [18], to the centre of the training dataset distribution.

$$\mathcal{L}^{(n)} = \left(r([\boldsymbol{c}; \boldsymbol{z}]) - y_{\text{target}}\right)^2 + P(\boldsymbol{z}) \tag{2a}$$

$$P(\boldsymbol{z}) = \lambda_1 \sum_{n \neq m} \left\| \boldsymbol{z}^{(n)} - \boldsymbol{z}^{(m)} \right\|_2^2 + \lambda_2 D_M(\boldsymbol{z}^{(n)}, \bar{\boldsymbol{z}}) \tag{2b}$$

$$D_M(\boldsymbol{z}, \bar{\boldsymbol{z}}) = (\boldsymbol{z} - \bar{\boldsymbol{z}})^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} (\boldsymbol{z} - \bar{\boldsymbol{z}}) \tag{2c}$$

The squared Mahalanobis distance, D_M, is defined in Eqn. 2c, where $\boldsymbol{\Sigma} \in \mathbb{R}^{L \cdot d \times L \cdot d}$ is the covariance matrix, and $\bar{\boldsymbol{z}} \in \mathbb{R}^{L \cdot d}$ is the mean flattened embedding over the training dataset. Pairwise quadratic repulsion spreads swarm members apart regardless of direction, discouraging collapse to a single solution. Mahalanobis distance is distribution-aware, pulling designs toward chemically plausible regions of the embedding space supported by the training dataset, accounting for anisotropy and correlated dimensions.

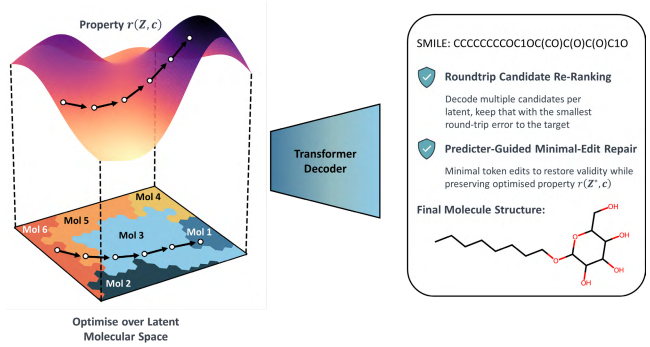

Figure 2. Inverse-design workflow: optimise over a learned property surface, $r(\boldsymbol{z}; \boldsymbol{c})$, toward the target, decode to SMILES, re-rank by round-trip error, and repair invalid near-misses.

2.5. Round-Trip Consistency Metric

After optimisation, each candidate embedding is decoded to SMILES. Because decoding may not preserve the intended latent state, we perform a round-trip check: for each optimised embedding, $\boldsymbol{Z}^*$, we (1) decode $\boldsymbol{Z}^*$ to generate a SMILES string $\boldsymbol{s}$, (2) re-encode $\boldsymbol{s}$ using the encoder to obtain the round-trip optimised embedding, $\boldsymbol{Z}^*_{\text{rt}}$, and then (3) flatten both embeddings and predict their properties using the regressor: $\hat{y}^* = r([\boldsymbol{z}^*; \boldsymbol{c}]), \hat{y}^*_{\text{rt}} = r([\boldsymbol{z}^*_{\text{rt}}; \boldsymbol{c}])$. Two consistency metrics were defined. The first was the round-trip MAPE between $\hat{y}^*_{\text{rt}}$ and y_{target}, quantifying how well the decoded molecule retained the desired property. To quantify how much the decoder perturbed the embedding representation, the second was the percentage Euclidean distance (PED) between $\boldsymbol{Z}^*$ and $\boldsymbol{Z}^*_{\text{rt}}$, defined in Eqn. 3. A high PED signals structural inconsistency introduced during decoding.

$$\text{PED} = 100 \times \frac{\|\boldsymbol{Z}^* - \boldsymbol{Z}^*_{\text{rt}}\|_2}{\|\boldsymbol{Z}^*\|_2} \tag{3}$$

To reduce decoding sensitivity, we use beam search to generate multiple candidates and re-rank candidates using round-trip consistency, selecting low-drift and/or low-error candidates, as reported in Section 4.

2.6. Latent-Manifold Optimisation

We perform inverse design in a latent space that captures the dominant directions of variation relevant to the target property. This was achieved by projecting the full embedding into a low-dimensional manifold and optimising within this space. To learn a suitable manifold, partial least squares (PLS) regression was fitted on the flattened embedding vectors, $\boldsymbol{z} \in \mathbb{R}^{L \cdot d}$, and the target property, y. This yields the linear projection in Eqn. 4a, with rotations $\mathbf{R} \in \mathbb{R}^{L \cdot d \times k}$, loadings $\mathbf{P} \in \mathbb{R}^{L \cdot d \times k}$, and scores $\boldsymbol{t} \in \mathbb{R}^{k}$, such that $\boldsymbol{t}$ lies in a k-dimensional subspace aligned with property-relevant variation. We reduced k until an elbow point in SMILES reconstruction and property prediction accuracy was observed. A linear head was then fitted on the scores concatenated with the covariate condition, as shown in Eqn. 4b.

$$\boldsymbol{t} = (\boldsymbol{z} - \bar{\boldsymbol{z}})\mathbf{R} \tag{4a}$$

$$\hat{y} = r([\boldsymbol{t}; \boldsymbol{c}]) \tag{4b}$$

At inverse design time, $\boldsymbol{c}$ was fixed and $\boldsymbol{t} \in \mathbb{R}^{k}$ was optimised subject to the property loss and penalty terms described before (e.g., swarm repulsion and distribution-aware tether), now in the latent score space. Afterwards, the optimised scores were then mapped back to a flattened embedding vector using Eqn. 5, reshaped back to a matrix and decoded to SMILES.

$$\boldsymbol{z}^{*} = \bar{\boldsymbol{z}} + \boldsymbol{t}^{*}\mathbf{P}^{\mathrm{T}} \tag{5}$$

To improve reconstruction fidelity, the multitask transformer can optionally be finetuned on the manifold-projected embeddings to better align the decoder with the subspace geometry.

2.7. Predictor-Guided Minimal-Edit Repair

Invalid SMILES are handled with a lightweight repair step. We first apply simple deterministic normalisations (e.g., minor punctuation/bracket fixes) and test validity with RDKit sanitisation. If still invalid, we perform a bounded search over minimal token edits (prioritising punctuation-like tokens commonly responsible for SMILES failures, expanding to atom tokens only if needed). Valid repaired candidates are scored by proximity to y_{target} and the closest candidate is returned.

3. SURFACTANT CASE STUDY AND DATASET PREPARATION

We evaluated the framework on surfactant critical micelle concentration (CMC) design – a nonlinear, emergent property influenced by both molecular structure and environmental conditions such as temperature.

The labelled dataset contains 965 surfactant SMILES with experimentally measured CMC values and temperatures [19]. CMC (mM) is log-transformed, and the dataset is split into train/validation/test = 70/15/15. To construct a domain-relevant pretraining corpus, each labelled surfactant was augmented with structurally similar molecules retrieved from PubChem [20] using a Tanimoto fingerprint similarity threshold of 0.60 yielding >15 million unique valid SMILES. All SMILES were sanitised and canonicalised using RDKit to remove invalid structures and duplicates. A BPE tokeniser was trained on this augmented corpus to compress frequent motifs (e.g., headgroups and alkyl segments), improving sequence efficiency for long amphiphiles.

4. RESULTS AND DISCUSSION

4.1. Property Prediction and Reconstruction Performance

Table 1. Comparison transformer models' performance.

Model	Parameters (million)	Train R^2	Test R^2	Reconstruction Similarity
ChemBERTa	77	0.85	0.83	-
SMI-TED	289	0.95	0.81	80%
MolT5	250	**0.94**	**0.80**	84%
Our CLM	**0.43**	0.89	**0.80**	**89%**

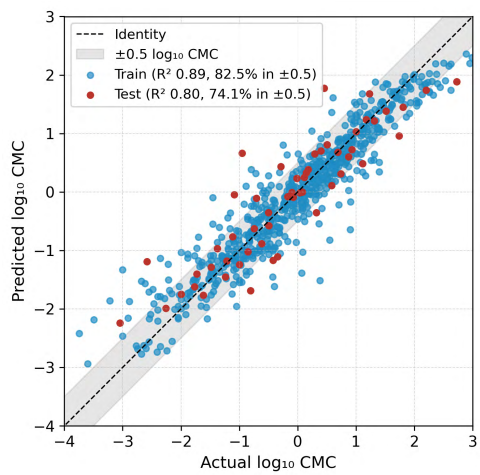

Figure 3. Predicted vs experimental $\log_{10}$ CMC. Shaded band marks ±0.5 $\log_{10}$ CMC with 82.5% and 74.1% of the training and testing points falling within this tolerance.

We first verified that the models used in the inverse-design loop provide (1) sufficient CMC prediction accuracy and (2) adequate encode–decode reconstruction, since round-trip evaluation depends on decoding candidates and re-evaluating them. Table 1 compares three pretrained chemical language models (ChemBERTa [11], SMI-TED [21], MolT5 [22]) to our CLM trained on an augmented surfactant dataset. Across models, property prediction test performance is similar ($R^2 \approx 0.80$–0.83), but

reconstruction is weaker for the pretrained models (80–84% exact matches) than for our CLM (89%), consistent with domain mismatch on long, motif-rich surfactant SMILES. Despite being much smaller (0.43M parameters), our CLM provides adequate accuracy. Figure 3 illustrates the train/test parity using our CLM where 74.1% of test points fall within ±0.5 $\log_{10}$ CMC, meeting typical industrial screening tolerances. Therefore, henceforth, we will use our CLM.

4.2. Diagnosing and Mitigating Round-Trip Inconsistencies

We evaluated inverse design on 50 target $\log_{10}$ CMC values evenly spaced from −2 to +2 at 25 °C. For each target, we optimise a latent embedding to match the target under the forward regressor, decode to SMILES, and re-evaluate the decoded molecule. A consistent failure mode emerged: embeddings could be optimised to match the target, yet decoded molecules were often off-target. To diagnose the cause, we calculated embedding drift as defined in Eqn. 3. Figure 4 shows drift occurs primarily on the first decode–re-encode (e.g., PED ≈ 33% after round-trip 1, falling to ~0% by round-trip 3), consistent with greedy decoding "snapping" to a nearby high-probability molecule that is stable under subsequent round-trips but not the intended latent optimum.

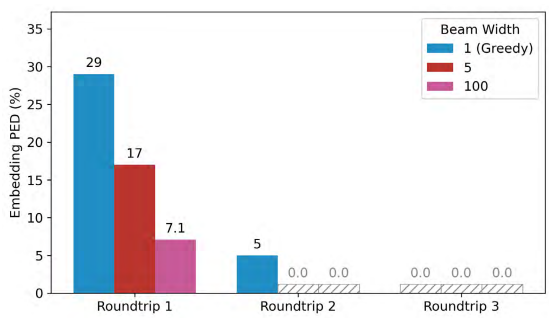

Figure 4. Impact of beam width on round-trip embedding stability, where PED is percentage Euclidian distance.

To mitigate snapping, we use beam search and a post-decoding re-ranking step: generate candidates (beam widths 1, 5, 100), re-encode each, and select the candidate with the smallest PED. Re-ranking improves both fidelity and stability: at beam width 100, average round-trip MAPE decreases from 20% to 5.1% and PED decreases from 29% to 7.1%. With re-ranking in place, decoder fine-tuning further improves round-trip fidelity (MAPE 5.1% → 0.7%), showing that round-trip accuracy – rarely reported in prior work – is crucial to ensure that the designed molecules reflect the optimisation objectives.

4.3. Molecular Structure Validity

SMILES validity remains a bottleneck: many near-target candidates are invalid due to small syntactic/chemical errors. Table 2 summarises the fidelity–validity trade-off under three post-processing strategies applied after decoding.

Table 2. Effect of validity filtering and SMILES repair on the round-trip target error.

	Round-Trip MAPE (%)	SMILES Validity (%)
No Post-Processing	0.7	6%
With Validity Filtering	56%	14%
Property-Guided Repair	**1.1%**	**90%**

Beam-decoded candidates were initially selected by target property error – this is reported as the no post-processing result in Table 2 – yielding low error (MAPE 0.7%) but very low validity (6%). Instead, from each beam of decoded sequences, we retained only chemically valid SMILES (i.e., via RDKit sanitisation). Within this valid subset, we selected the molecule with the lowest round-trip drift relative to the optimised embedding. This increased validity to 14% but introduced a trade-off: the reduced candidate pool inflated target property deviation, with round-trip MAPE rising from 0.7% to 56%. This highlights a tension between enforcing strict validity and maintaining property fidelity; ideally, validity should be enforced during decoding, or by using less aggressive sanitisation techniques. We addressed this with predictor-guided minimal-edit repair: invalid near-misses are minimally edited to restore validity, then re-ranked by round-trip drift, recovering ~84% of invalid outputs and achieving 90% validity with a round-trip MAPE of 1.1%.

4.4. Property Trends and Design Interpretability

We finally check that the designed molecules are chemically plausible and interpretable. Figure 5 shows representative designs across the target range, paired with nearest neighbours from the training set. No exact matches to the training set were observed; average nearest-neighbour Tanimoto similarity was 0.31, and adjacent-target designs remain diverse (average Tanimoto 0.20). To our knowledge, all of the generated surfactants were entirely novel. The generated molecules follow established surfactant trends: molecules generated under lower CMC targets favour longer hydrophobic tails, while molecules generated under higher CMC targets favoured shorter tails and more polar/ionic headgroups (e.g., amines, ammonium salts, carboxylates) [23].

It is essential to understand why a model makes a specific prediction and which substructures influence the property of interest. To this end, gradient-based atom-level saliency [24] was carried out in this work. The saliency scores are visualised in Figure 5, where atoms with score > 0 that locally increase $\hat{y}$ are shaded red, while atoms with score < 0 that locally decrease $\hat{y}$ are shaded

blue. Saliency agrees with existing surfactant design rules: hydrophobic alkyl tails decrease CMC (blue), while polar/ionic headgroups increase CMC (red) [23]. Branching/heteroatoms in the tail tend to raise CMC, while linker effects (ethers/esters) are modest and context dependent. Furthermore, these saliency motifs align between designed and nearest-neighbour structures (e.g., the same headgroup and the same tail segments receive similar signs), reinforcing that the model has learned chemically meaningful structure-property relationships.

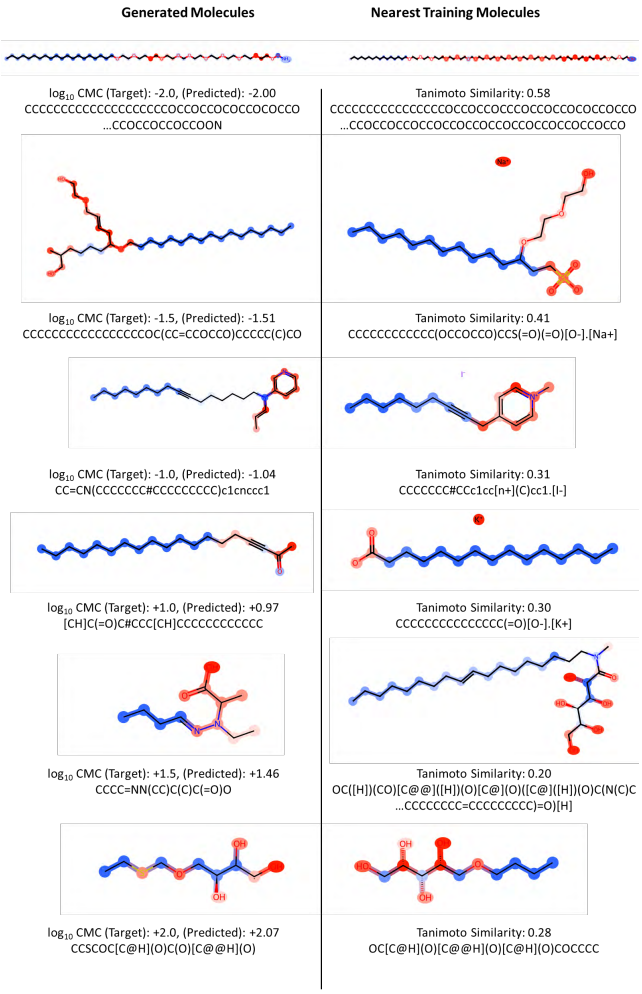
Generated Molecules
Nearest Training Molecules
log10 CMC (Target): -2.0, (Predicted): -2.00
CCCC

information system. 1. introduction to methodology and encoding rules. J. Chem. Inf. Comput. Sci. 28:31-36 (2002). https://doi.org/10.1021/ci00057a005
5. Wieder O, Kohlbacher S, Kuenemann M, Garon A, Ducrot P, Seidel T, Langer T. A compact review of molecular property prediction with graph neural networks. Drug Discovery Today: Technologies 37:1-12 (2020). https://doi.org/10.1016/j.ddtec.2020.11.009
6. Meyers J, Fabian B, Brown N. De novo molecular design and generative models. Drug Discovery Today 26:2707-2715 (2021). https://doi.org/10.1016/j.drudis.2021.05.019
7. Kwon Y, Kang S, Choi YS, Kim I. Evolutionary design of molecules based on deep learning and a genetic algorithm. Sci Rep 11: (2021). https://doi.org/10.1038/s41598-021-96812-8
8. Gómez-Bombarelli R, Wei JN, Duvenaud D, Hernández-Lobato JM, Sánchez-Lengeling B, Sheberla D, Aguilera-Iparraguirre J, Hirzel TD, Adams RP, Aspuru-Guzik A. Automatic chemical design using a data-driven continuous representation of molecules. ACS Cent. Sci. 4:268-276 (2018). https://doi.org/10.1021/acscentsci.7b00572
9. Segler MHS, Kogej T, Tyrchan C, Waller MP. Generating focused molecule libraries for drug discovery with recurrent neural networks. ACS Cent. Sci. 4:120-131 (2017). https://doi.org/10.1021/acscentsci.7b00512
10. Kim K, Kang S, Yoo J, Kwon Y, Nam Y, Lee D, Kim I, Choi YS, Jung Y, Kim S, Son WJ, Son J, Lee HS, Kim S, Shin J, Hwang S. Deep-learning-based inverse design model for intelligent discovery of organic molecules. npj Comput Mater 4: (2018). https://doi.org/10.1038/s41524-018-0128-1
11. S. Chithrananda, G. Grand, B. Ramsundar, ChemBERTa: Large-Scale Self-Supervised Pretraining for Molecular Property Prediction, (2020). https://doi.org/10.48550/arXiv.2010.09885
12. C. Xu, S. Zhu, V. Viswanathan, CLOUD: A Scalable and Physics-Informed Foundation Model for Crystal Representation Learning, (2025). https://doi.org/10.48550/arXiv.2506.17345
13. Chen LY, Li YP. Uncertainty quantification with graph neural networks for efficient molecular design. Nat Commun 16: (2025). https://doi.org/10.1038/s41467-025-58503-0
14. Xu C, Wang Y, Barati Farimani A. Transpolymer: a transformer-based language model for polymer property predictions. npj Comput Mater 9: (2023). https://doi.org/10.1038/s41524-023-01016-5
15. Nnadili M, Okafor AN, Olayiwola T, Akinpelu D, Kumar R, Romagnoli JA. Surfactant-specific ai-driven molecular design: integrating generative models, predictive modeling, and reinforcement learning for tailored surfactant synthesis. Ind. Eng. Chem. Res. 63:6313-6324 (2024). https://doi.org/10.1021/acs.iecr.4c00401
16. E. Soares, Z. Sumer, R.L. Anderson, REPRESENTING SURFACTANTS BY FOUNDATION MOD-, (2025).
17. T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T.L. Scao, S. Gugger, M. Drame, Q. Lhoest, A.M. Rush, HuggingFace's Transformers: State-of-the-art Natural Language Processing, (2020). https://doi.org/10.48550/arXiv.1910.03771
18. P.C. Mahalanobis, Reprint of: Mahalanobis, P.C. (1936) "On the Generalised Distance in Statistics., " Sankhya A 80 (1936) 1–7. https://doi.org/10.1007/s13171-019-00164-5
19. P. Mukerjee, K. Mysels, Critical micelle concentrations of aqueous surfactant systems:, NIST (1971). https://www.nist.gov/publications/critical-micelle-concentrations-aqueous-surfactant-systems (accessed September 19, 2025).
20. S. Kim, J. Chen, T. Cheng, A. Gindulyte, J. He, S. He, Q. Li, B.A. Shoemaker, P.A. Thiessen, B. Yu, L. Zaslavsky, J. Zhang, E.E. Bolton, PubChem 2025 update, Nucleic Acids Res. 53 (2025) D1516–D1525. https://doi.org/10.1093/nar/gkae1059
21. E. Soares, E.V. Brazil, V.Y. Shirasuna, D. Zubarev, R. Cerqueira, K. Schmidt, SMI-TED: A large-scale foundation model for materials and chemistry, (2024). https://openreview.net/forum?id=Yq8At31hLi&utm_source=chatgpt.com (accessed August 18, 2025).
22. C. Edwards, T. Lai, K. Ros, G. Honke, K. Cho, H. Ji, Translation between Molecules and Natural Language, (2022). https://doi.org/10.48550/arXiv.2204.11817
23. Rosen MJ, Kunjappu JT. Surfactants and interfacial phenomena. Wiley (2012). https://doi.org/10.1002/9781118228920
24. D. Smilkov, N. Thorat, B. Kim, F. Viégas, M. Wattenberg, SmoothGrad: removing noise by adding noise, (2017). https://doi.org/10.48550/arXiv.1706.03825

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Modeling Framework Integrating Data Trends and Reference Information for Predicting Temperature-Dependent Thermophysical Properties

Shuai Zhang[a], Abdulelah S. Alshehri[b*], Mansour S. Alhoshan[b], and Anjan Tula[a*]

[a] State Key Laboratory of Industrial Control Technology, College of Control Science and Engineering, Zhejiang University, Hangzhou 310027, China.
[b] Chemical Engineering Department, College of Engineering, King Saud University, Riyadh 11421, KSA
[*] Corresponding author: mhoshan@KSU.EDU.SA.
[*] Corresponding author: anjantula@zju.edu.cn.

ABSTRACT

The availability of temperature-dependent physicochemical property data forms the cornerstone of process simulation, optimization, and sustainable molecular and product design. However, a critical data gap persists, as experimental measurements are accessible for only a small subset of known chemicals. This renders experimental characterization resource-prohibitive, often compelling reliance on empirical estimation methods. Moreover, although many models offer single-point predictions at fixed temperatures, accurately modeling continuous temperature-dependent behavior remains challenging. Conventional methods frequently overlook intermediate variations, resulting in limited extrapolation capability. To overcome these limitations, we introduce a mechanism-guided hybrid modeling framework that integrates physical insights into data-driven models. This framework is built on two strategies. Strategy I targets trend correction by generating a continuous representation from discrete single-point predictions, incorporating descriptors and slopes. Strategy Ⅱ addresses bias removal by anchoring a baseline to a high-accuracy point estimate and fitting the remaining deviations. The framework's effectiveness is evidenced by evaluations across ten thermophysical properties: Strategy I achieves MSE reductions of 19.23% and 20.33% for the quantitative structure-property relationship and group contribution methods, respectively. Strategy II provides a more substantial improvement, attaining an 81.63% MSE reduction for the gradient boosting decision tree regression model. This work demonstrates that incorporating trend and slope constraints facilitates physically consistent, bias-corrected, and accurate predictions, offering a scalable approach to bridge the data gap and accelerate computer-aided engineering and design.

Keywords: Temperature-dependent property prediction, Hybrid modeling, Mechanistic constraints, Machine learning, Slope-based correction, Bias correction

1. INTRODUCTION

The characteristics of organic compounds, especially the physical properties related to energy and phase behavior, are of great significance in the field of process engineering and serve as the foundation for process design and optimization. However, due to the scarcity of experimental data and the high cost of conducting experimental measurements, property prediction models thus represent the most viable and practical means to estimate the required properties in the initial design phase.

Benefiting from the relative abundance of standardized data points, single-point prediction models (such as boiling point and critical temperature) have seen significant development. In stark contrast, the development of models for temperature-dependent properties remains a notably under-explored area. This research gap is not only due to the scarcity of comprehensive experimental

data but, more fundamentally, to the significant increase in modeling complexity. Predicting a single value, by comparison, constitutes a fundamentally simpler task than learning a continuous temperature-dependent function. The latter demands models that can simultaneously capture the interplay of molecular structure and thermal dynamics, adhere to underlying physical laws, and maintain generalizability across both chemical space and continuous temperature domains. This pronounced complexity, coupled with the scarcity of comprehensive data, has therefore intensified the need for accurate and efficient models capable of estimating thermophysical properties for compounds lacking experimental measurements.

The most classic predictive models include: group contribution (GC) models[1], quantitative structure-property relationship (QSPR) models[2], and machine learning (ML)models[3]. In the prediction of temperature-dependent properties, GC models are often used to regress the coefficients of thermodynamic approximation or correlation functions. This workflow and its essentials are well captured by early foundational work [4]. Early temperature-dependent predictive models include:

(1)Heat capacity polynomial

$$C_p = [\sum(a) - 37.93] + [\sum(b) + 0.210]T + [\sum(c) - 3.91 \times 10^{-4}]T^2 + [\sum(d) + 2.06 \times 10^{-7}]T^3 \quad (1)$$

(2)Liquid-viscosity (Andrade-type) relation

$$\eta_L = mw \times \exp\left\{\frac{\sum(\eta_a) - 597.82}{T} + \sum(\eta_b) - 11.202\right\} \quad (2)$$

For models without explicit functional forms, a standard approach to predicting temperature-dependent properties is to treat temperature as an additional input feature, concatenated with molecular descriptors or structural information. This approach is widely used in various modeling contexts. For instance, in QSPR modeling, Sosnowska et al. [5] integrated QSPR-predicted enthalpy of vaporization into the Clausius-Clapeyron equation to characterize the temperature-dependent vapor pressure of persistent organic pollutants. This data-driven strategy has enabled predictions even under extreme conditions, as demonstrated by Yin et al. [6], who predicted gas heat capacity across an exceptionally wide temperature range. Similarly, tree-based ensemble methods such as gradient boosting regression trees and light gradient boosting machine have been employed to predict liquid heat capacity [7] and solubility [8] across varying temperatures by treating temperature as an input variable.

These models share a fundamental limitation in handling temperature-dependent properties: the prevalent methodology essentially treats data points from the same compound at different temperatures merely as an expanded dataset, processing each (temperature, property) pair as an independent sample. This approach inherently disregards the intrinsic physical continuity of the temperature-property curve. Consequently, conventional models trained on such discrete (structure, temperature) pairs fail to learn the underlying continuous functional relationship, resulting in physically inconsistent predictions that often exhibit incorrect curve shapes and poor extrapolation beyond the training temperature domain.

The above theoretical limitations are clearly reflected in practical predictions. As shown in Figure 1, the GC model, constrained by its fixed functional form, fails to accurately capture experimental trends: Figure 1A demonstrates its constant-trend prediction for heat capacity that deviates from the experimental data, while Figure 1B and 1C show increasingly amplified errors as the temperature range extends. Figure 1D further illustrates its overall failure to follow the actual property trend.

In contrast, while QSPR models offer greater flexibility, they introduce a different challenge due to their purely data-driven nature. As shown in Figure 2, all four subfigures (A-D) display abrupt, physically unsupported changes in predictions. Notably, Figures 2B and 2C exhibit nearly non-monotonic segments that contradict fundamental thermodynamic principles.

2. PROPOSED HYBRID MODELING FRAMEWORK

2.1 Strategy I – Trend-Guided modeling

Table 1: Linearity assessment of ten temperature dependent thermophysical properties. Data were obtained from the pure component analysis module of Aspen Plus. For each property, the R^2 value is averaged across all substances, where each substance's R^2 is calculated from the linear regression of property values over temperature points.

Property (abbreviation)	Mean R^2	Count Of R^2<0.95
Heat capacity(C_p)	0.9768	27*
Heat of vaporization($\Delta Hvap$)	0.9843	8
Internal energy(U)	0.9932	0
Liquid density(ρ_L)	0.9904	0
Pure component enthalpy(H)	0.9948	0
Pure component exergy(Ex)	0.9909	0
Gibbs free energy(G)	0.9964	0
Isentropic exponent(γ)	0.9927	0
Vapor pressure(P_{vap})	0.9942	0
Volume(V)	0.9731	4

*17 substances with R^2 between 0.9 and 0.95, 6 substances with R^2 between 0.8 and 0.9, and 4

substances with R^2 below 0.8.

As demonstrated in the table above, fitting a linear trend to the experimental data for ten temperature-dependent properties reveals a strong linear relationship in all cases, with an average coefficient of determination (R^2) exceeding 0.97. (Note that vapor pressure data was log-transformed prior to fitting.) Consequently, high-accuracy predictions at two physically meaningful temperature points, such as the critical temperature (T_c) and the normal boiling point (T_b), are employed to estimate the property's temperature-dependent tendency.

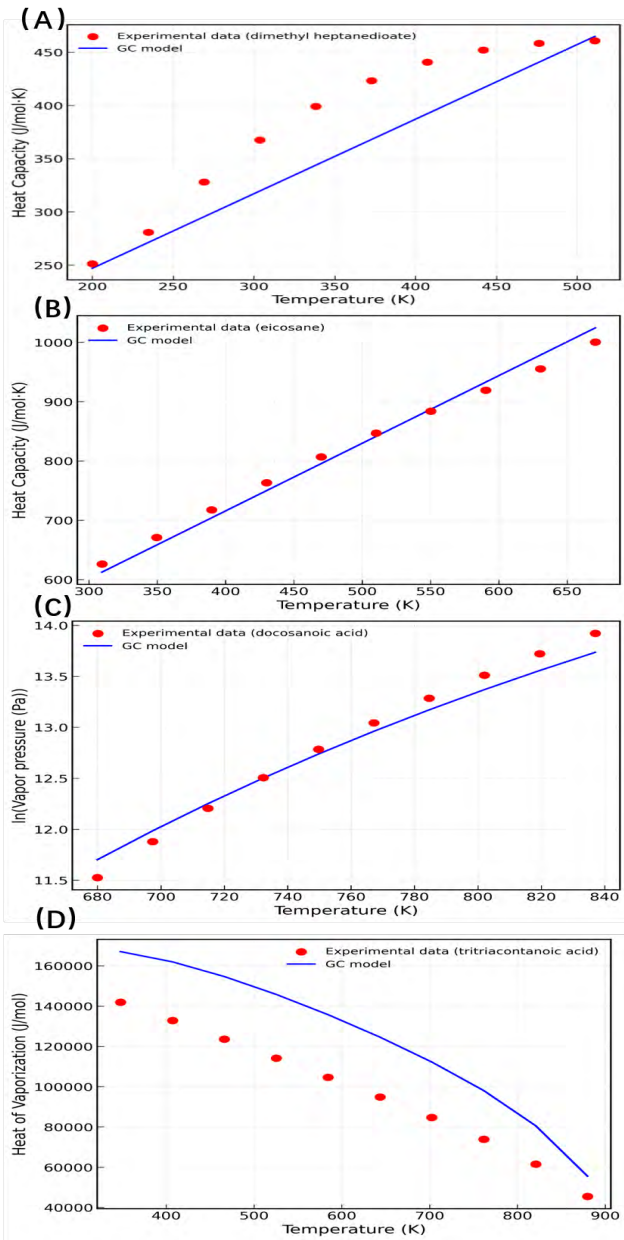

Figure 1. Prediction Performance of the GC Model for Thermodynamic Properties.(A, B) Heat capacity[9]. (C)Vapor pressure[10].(D)Heat of vaporization[9]. Model details are available in the Supplementary Information.

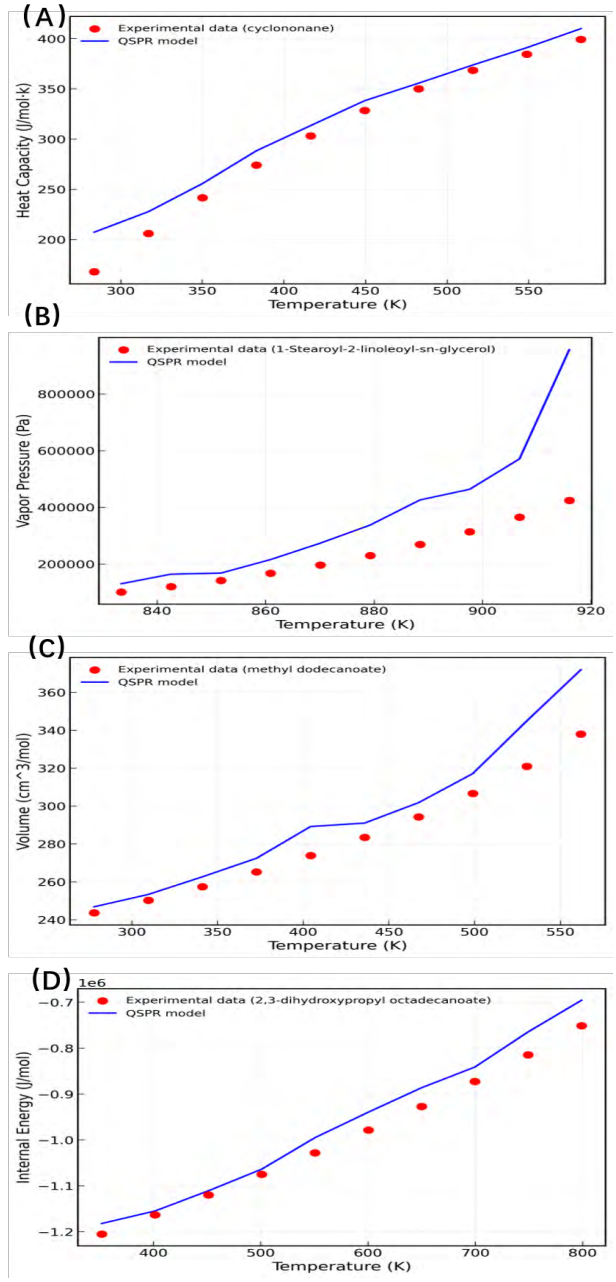

Figure 2. Prediction Performance of the QSPR Model for Thermodynamic Properties.(A) Heat capacity.(B)Vapor pressure.(C) Volume. (D) Internal Energy. Molecular descriptors derived from 2D structures downloaded from PubChem were used as input features. After data cleaning, appropriate descriptors were selected, and a random forest regression model was employed. Model details are available in the Supplementary Information.

2.1.1 Application of strategy I for explicit models

As presented in Figure 3, the modeling of temperature dependence for explicit models is accomplished through the following two-step framework:

Step 1: Construct four foundational models from molecular structure: two for predicting the reference

temperatures T_1 and T_2, and two for estimating the corresponding properties y_{r1} and y_{r2} at these points. The slope between these reference points is then derived as follows:

$$tendency = \frac{y_{r2}-y_{r1}}{T_2-T_1} \tag{3}$$

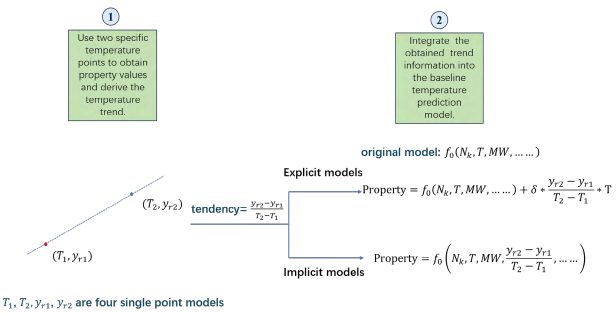

Figure 3. Flowchart of the Trend-Guided Hybrid Modeling Framework. Step 1: Construct models to predict two reference temperature points and their corresponding target properties, then use these four models to derive the predicted slope. Step 2: For existing models with an explicit expression, incorporate this slope as an additional temperature interaction term. For models without an explicit expression, use the slope as an additional input feature.

Step 2: Incorporate this slope into a baseline explicit model by introducing a temperature correction term. The augmented model takes the form:

$$\hat{y} = f_0(N_k, T, MW, \dots\dots) + \delta * \frac{y_{r2}-y_{r1}}{T_2-T_1} * \mathrm{T} \tag{4}$$

Here, $f_0(N_k, T, MW, \dots\dots)$denotes the original mechanistic or empirical model. Its temperature-inclined correction is introduced via a term weighted by a coefficient δ and scaled by the temperature T. This design allows the model to adaptively capture temperature-driven behavior.

2.1.2 Application of strategy I for ML models

The temperature-dependent modeling for ML models follows a two-step procedure:

Step 1: Construct four foundational models from molecular structure: two for predicting the reference temperatures T_1 and T_2, and two for estimating the corresponding properties y_{r1} and y_{r2} at these points. The slope between these reference points is then derived as follows:

$$tendency = \frac{y_{r2}-y_{r1}}{T_2-T_1} \tag{5}$$

Step 2: Incorporate the temperature tendency by adding the slope as a new input feature to the original data-driven model. This allows the model to learn temperature-dependent relationships directly, while its architecture remains unmodified. The resulting enhanced model formulation is:

$$\hat{y} = f_0\left(N_k, T, MW, \dots\dots, \frac{y_{r2}-y_{r1}}{T_2-T_1}\right) \tag{6}$$

Here, $f_0(N_k, T, MW, \dots\dots)$represents the original implicit model that learns structure-property relationships directly from data.

2.1.3 Physical significance

The physical significance of Strategy I stems from its fundamental departure from discrete-point prediction. By introducing a "temperature tendency" derived directly from experimental data points, the method incorporates prior physical intuition about continuous, smooth property evolution. This data-driven yet physics-aware approach not only improves robustness against experimental noise but also ensures more reliable extrapolation and interpolation. Crucially, it bridges the gap between discrete experimental observations and continuous theoretical physical models.

2.2 Strategy II –Bias-Corrected hybrid modeling

The conceptual foundation of Strategy II originates from a critical observation regarding purely data-driven machine learning methods. While these algorithms excel at identifying complex patterns, we propose that a more effective approach is to first establish a predictive baseline model, rather than feeding all data directly into a black-box model. This baseline, which uses linear regression with group features as inputs, captures the fundamental trend; subsequently, a GBDT model is applied specifically to learn the deviations between the baseline predictions and the actual measured values. This two-step process is illustrated in the flowchart of Figure 4.

2.2.1 Baseline model construction

Supported by the linearity assessment of ten temperature-dependent thermophysical properties (Table 1), a linear model is adopted as the baseline. Two preliminary models are first built to predict a reference temperature, T_1, and the corresponding property value at that temperature, y_{r1}, directly from molecular structure. This reference prediction, y_{r1}, serves as the intercept of the linear baseline model. The slope of the baseline is then determined using group contribution methods, giving the following expression for the baseline prediction $\hat{y}$:

$$\hat{y} = y_{\mathrm{r1}} + \sum_K A_K N_K (T - T_1) \tag{7}$$

Here, y_{r1} is the reference prediction for the property at temperature T_1, A_K denotes the group contribution coefficients, N_K represents the counts of molecular groups, and $T - T_1$ is the temperature deviation from the reference point.

2.2.2 Bias learning and compensation

Although the baseline model is physically reasonable, it cannot fully capture all complex, nonlinear relationships. To address this limitation, the deviation between the baseline prediction and the actual experimental value is calculated as $b = y - \hat{y}$. A machine learning model is subsequently employed to predict this systematic bias:

$$\hat{b} = f_1(N_k, T, MW \dots \dots) \tag{8}$$

This model, f_1, learns the residual errors of the baseline using a set of features including molecular descriptors, temperature, and molecular weight. The final, refined prediction is obtained by compensating the baseline output with the predicted bias:

$$\check{y} = \hat{y} + b \tag{9}$$

This step combines the robustness of the physical baseline with the flexibility of machine learning, enabling the hybrid model to capture nonlinear deviations that the baseline alone cannot represent.

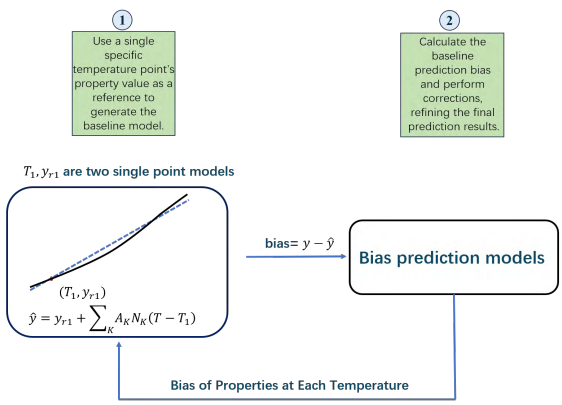

Figure 4. Flowchart of the Bias-Corrected Hybrid Modeling Framework. Step 1: Construct a prediction model for a reference temperature point and its corresponding target property to serve as an anchor point. Then, using the group contribution method, regress a straight line through this anchor point, where the method determines the slope. This line is taken as the predictive prior model for the property. Step 2: Calculate the deviation between the predictions of this prior model and the actual observed values. By predicting this deviation, apply a feedback compensation to the baseline model.

2.2.3 Physical significance

This hybrid methodology offers distinct advantages over purely black-box machine learning. Its primary strength lies in enforcing physically realistic predictions by anchoring them to a group-contribution baseline that adheres to thermodynamic principles. This constraint mitigates the generation of unrealistic outcomes—a common issue in purely data-driven models, particularly during extrapolation beyond the training domain. The framework also enhances interpretability: the baseline provides a physically meaningful starting point, while the machine learning component is dedicated to modeling only the residual deviations from this baseline. This separation allows researchers to distinguish between the fundamental physical relationships embedded in the baseline and the more complex, data-driven patterns captured by the machine learning model. Furthermore, by decomposing the problem into these two complementary tasks, the approach achieves greater data efficiency and robustness, as the machine learning model is tasked specifically with correcting the baseline's limitations rather than learning the entire underlying relationship from scratch.

3. RESULTS AND DISCUSSION

3.1 Data and evaluation framework

Table 2: Summary of experimental data points for the ten evaluated temperature-dependent properties, For each substance, data was collected at 10 different temperature points. Experimental data were obtained from Aspen Plus.

Property	Symbol	Datapoints
C_p	$J/mol \cdot k$	2070
$\Delta Hvap$	J/mol	2040
U	J/mol	2070
ρ_L	g/ml	2070
H	J/mol	2070
Ex	J/mol	2050
G	J/mol	2050
γ	1	2070
P_{vap}	Pa	2090
V	cm^3/mol	2080

To evaluate the proposed enhancement strategies, we conducted a systematic performance comparison across several modeling frameworks. In this notation, the suffix "_SI" indicates the application of Strategy I, and "_SII" indicates Strategy II. Using this labeling scheme, we compared the following model variants: GC vs. GC_SI, GBDT vs. GBDT_SII, and QSPR vs. QSPR_ SI. All model variants are clearly identified with these labels in the subsequent results and discussion sections.

Table 2 summarizes the experimental data points for the ten evaluated temperature-dependent properties. For each substance, data were collected at ten distinct temperature points to comprehensively capture the property's behavior across a temperature range. These experimental data, serving as the foundation for all evaluations, were sourced from Aspen Plus—a widely used process simulation tool known for providing reliable thermophysical property data for diverse compounds.

Table 3. Overall predictive performance of original and enhanced models for ten thermophysical properties. Original models are labeled as GC, GBDT, and QSPR; models enhanced by Strategy I and Strategy II are denoted by the suffixes "_SI" and "_SII", respectively. For specific models and raw data, please refer to the digital supplementary material.

Properties		GC	GC_SI	GBDT	GBDT_SII	QSPR	QSPR_SI
C_p	R^2	0.9918	0.9924	0.9971	0.9981	0.9994	0.9994
	MSE	16.08	14.89	5.60	3.66	1.19	1.14
P_{vap}	R^2	0.9033	0.9194	0.8476	0.9914	0.9886	0.9901
	MSE	63.04	52.58	99.36	5.59	7.41	6.43
$\Delta Hvap$	R^2	0.9260	0.9534	0.9713	0.9941	0.9929	0.9961
	MSE	64.55	40.66	24.98	5.14	6.15	3.40
V	R^2			0.9931	0.9985	0.9991	0.9992
	MSE			331.45	73.67	41.83	37.78
U	R^2			0.9965	0.9996	0.9991	0.9992
	MSE			50.58	6.36	12.59	11.51
H	R^2			0.9965	0.9996	0.9991	0.9992
	MSE			49.85	6.18	13.17	11.00
Ex	R^2			0.9728	0.9987	0.9992	0.9992
	MSE			119.77	5.90	3.59	3.53
γ	R^2			0.9600	0.9956	0.9931	0.9955
	MSE			354.26	39.37	61.24	40.08
ρ_L	R^2			0.9636	0.9957	0.9932	0.9950
	MSE			67.49	8.00	12.53	9.23
G	R^2			0.9843	0.9975	0.9957	0.9971
	MSE			50.27	8.05	13.80	9.27

3.2 Comparative performance analysis

Model performance was quantitatively assessed using the R^2 to evaluate the overall goodness of fit and the MSE to measure the magnitude of absolute error.

The overall performance of the models in predicting the ten target properties is summarized in Table 3. Note that results from the GC-based model are only reported for heat capacity, vapor pressure, and heat of vaporization; for the remaining properties, no established GC models were available, and those entries are left blank. Based on the results in Table 3, a consistent improvement is observed across all original models after incorporating the proposed strategies, evidenced by increased R^2 values and reduced MSE. This dual improvement reflects both a better overall fit and a systematic reduction in prediction errors, thereby statistically validating the effectiveness of Strategy I and Strategy II.

Table 4: MSE improvement results using Strategy I and Strategy II: Percentage improvement in MSE relative to the original model for each property. Blank entries indicate cases where the corresponding strategy-model combination was not applicable. Original models are denoted as GC, gradient boosting decision tree regression (GBDT), and QSPR; models enhanced with Strategy I and Strategy II are labeled with the suffixes "_SI" and "_SII", respectively.

Properties	GC_SI	GBDT_SII	QSPR_SI
C_p	7.40%	34.64%	4.20%
P_{vap}	16.59%	94.37%	13.23%
$\Delta Hvap$	37.01%	79.42%	44.72%
V		77.77%	9.68%
U		87.42%	8.58%
H		87.60%	16.48%
Ex		95.07%	1.67%
γ		88.89%	34.55%
ρ_L		88.15%	26.34%
G		83.99%	32.83%
Mean	20.33%	81.63%	19.23%

To further quantify the improvement, we compared the percentage reduction in MSE before and after applying the proposed strategies. As shown in Table 4, the strategies yield substantial and broadly consistent error reduction across most models. With Strategy I, the average MSE decreased by 19.23% for QSPR-based models and by 20.33% for GC models. For GBDT models, Strategy II delivered the most pronounced improvement, achieving an 81.63% reduction in MSE. These results underscore the efficacy of the Bias-Corrected Hybrid Modeling Framework for complex data-driven modeling tasks.

3.3 Model interpretation and physical insights

Figure 5 compares the original and enhanced models across three temperature-dependent properties. In Figure 5A, the GC model—originally constrained by a fixed linear form—exhibits markedly improved alignment with the experimental trend after incorporating Strategy I, yielding more flexible predictions that reduce deviations at both ends of the temperature range. Figure 5B shows that for the QSPR model, Strategy I preserves its inherent capacity to capture nonlinearity while imposing a physically guided constraint that suppresses unphysical, abrupt variations. The remaining panels consistently illustrate the advantages gained by integrating the proposed physics-informed features, further confirming the general effectiveness of the enhancement strategies.

The comparative results highlight a consistent advantage of incorporating physical knowledge into data-driven models. While purely statistical approaches learn patterns directly from the data, they lack awareness of the underlying thermodynamic structure governing temperature-dependent properties. The proposed strategies address this gap by embedding mechanistic priors into the learning process, enabling the hybrid framework to recover physically meaningful trends that conventional models often fail to capture. As illustrated by the predictions, many original models struggle to reproduce the experimental temperature variation because they are trained across multiple compounds using a global loss function that produces parameter sets representing statistical compromises. This leads to persistent prediction gaps for specific compounds—systematic biases that cannot be eliminated without incorporating physical mechanisms.

The physics-informed strategies introduced in this work effectively resolve these issues by embedding fundamental temperature-dependent behavior into the predictive process. In particular, Strategy I incorporates a physically interpretable slope term extracted from high-accuracy reference points such as T_b and T_c. This slope simultaneously acts as a physics-informed feature—encoding the direction and magnitude of temperature effects—and as a molecularly grounded descriptor, since it reflects influences such as intermolecular forces, molecular weight, functional group contributions, and polarizability. By doing so, Strategy I transforms the learning task from predicting an unconstrained surface in (structure, T) → property space to predicting deviations around a physically anchored linearized manifold, analogous to embedding a first-order thermodynamic Taylor expansion as a prior.

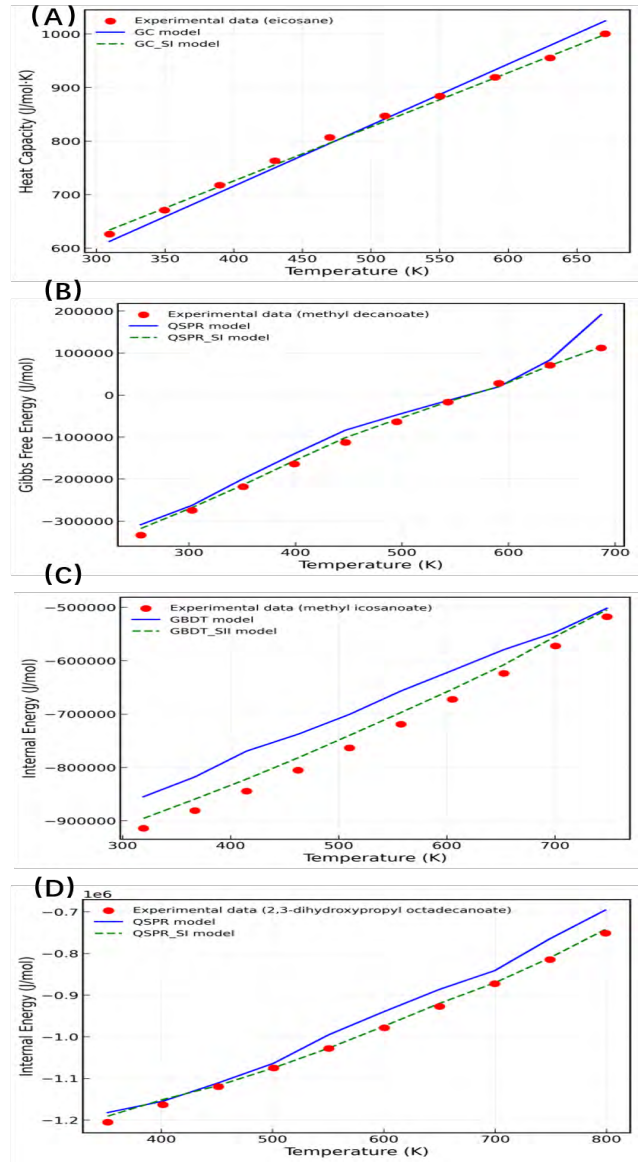

Figure 5. (A) GC vs. GC_SI for heat capacity. (B) QSPR vs. QSPR_SI for internal energy. (C) GBDT vs. GBDT_SII for internal energy. (D) QSPR vs. QSPR_SI for Gibbs free energy. (The suffix "_SI" denotes enhancement by Strategy I, "_SII" by Strategy II). For specific models and raw data, please refer to the digital supplementary material.

Similarly, Strategy II leverages a baseline GC model that encodes known physical relationships between

molecular groups and temperature tendencies. Although the baseline is approximate, it captures the dominant trend with clear physical interpretability, while the machine-learning component predicts only the deviation from this baseline. This decomposition substantially reduces the learning difficulty: the ML model focuses solely on complex residual patterns rather than the entire thermodynamic mapping. The baseline also prevents thermodynamic inconsistencies, such as incorrect concavity or unphysical changes in slope. As a result, the overall error decreases because ML is used to refine the nonlinear residual structure rather than reconstruct the entire temperature dependence from scratch. This approach closely mirrors the "Δ-learning" philosophy used in quantum chemistry machine learning, and our framework adopts the same logic for temperature-dependent thermophysical properties.

4 CONCLUSIONS

This paper presents a predictive framework for temperature-dependent properties that advances existing methodologies by addressing a fundamental limitation: the prevalent treatment of experimental data as discrete points rather than as manifestations of an underlying continuous function. By incorporating physically meaningful temperature trends and systematically correcting biases within data-driven workflows, the framework enables models to learn complete temperature-dependent functional relationships. Beyond predictive accuracy, it ensures thermodynamic consistency, enhances interpretability, and improves extrapolation reliability—attributes that are critical for process simulation, optimization, and computer-aided molecular design.

DIGITAL SUPPLEMENTARY MATERIAL

The complete source code, datasets, and implementation details for all models presented in this study have been made openly available. The repository contains the input features for all 10 properties across 200+ compounds, the code for all comparison models. The TRGPACK repository can be accessed at: https://github.com/SZ-ZJU/TRGPACK.

ACKNOWLEDGEMENTS

Financial support from the Natural Science Foundation of China (No. 22150410338) is gratefully acknowledged. The authors would like to extend their appreciation to the Deanship of Scientific Research at King Saud University for funding this work through ISPP program (ISPP25-11).

REFERENCES

1. Hukkerikar AS, Sarup B, Ten Kate A, Abildskov J, Sin G, Gani R. Group-contribution+ (GC+) based estimation of properties of pure components: improved property estimation and uncertainty analysis. Fluid Phase Equilibria 321:25-43 (2012). https://doi.org/10.1016/j.fluid.2012.02.010
2. Yalamanchi KK, van Oudenhoven VCO, Tutino F, Monge-Palacios M, Alshehri A, Gao X, Sarathy SM. Machine learning to predict standard enthalpy of formation of hydrocarbons. J. Phys. Chem. A 123:8305-8313 (2019). https://doi.org/10.1021/acs.jpca.9b04771
3. Alshehri AS, Tula AK, You F, Gani R. Next generation pure component property estimation models: with and without machine learning techniques. AIChE Journal 68: (2021). https://doi.org/10.1002/aic.17469
4. JOBACK KG, REID RC. ESTIMATION OF PURE-COMPONENT PROPERTIES FROM GROUP-CONTRIBUTIONS. Chemical Engineering Communications 57:233-243 (2007). https://doi.org/10.1080/00986448708960487
5. Sosnowska A, Barycki M, Jagiello K, Haranczyk M, Gajewicz A, Kawai T, Suzuki N, Puzyn T. Predicting enthalpy of vaporization for persistent organic pollutants with quantitative structure–property relationship (QSPR) incorporating the influence of temperature on volatility. Atmospheric Environment 87:10-18 (2014). https://doi.org/10.1016/j.atmosenv.2013.12.036
6. Yin J, Jia Q, Yan F, Wang Q. Predicting heat capacity of gas for diverse organic compounds at different temperatures. Fluid Phase Equilibria 446:1-8 (2017). https://doi.org/10.1016/j.fluid.2017.05.006
7. Shan Y, Wu Q, Yuan H, Liu W. Develop machine learning-based model and automated process for predicting liquid heat capacity of organics at different temperatures. Fluid Phase Equilibria 584:114132 (2024). https://doi.org/10.1016/j.fluid.2024.114132
8. Ye Z, Ouyang D. Prediction of small-molecule compound solubility in organic solvents by machine learning algorithms. J Cheminform 13: (2021). https://doi.org/10.1186/s13321-021-00575-3
9. Ceriani R, Gani R, Meirelles AJA. Prediction of heat capacities and heats of vaporization of organic liquids by group contribution methods. Fluid Phase Equilibria 283:49-55 (2009). https://doi.org/10.1016/j.fluid.2009.05.016
10. Ceriani R, Meirelles AJA. Predicting vapor–liquid equilibria of fatty systems. Fluid Phase Equilibria 215:227-236 (2004).

https://doi.org/10.1016/j.fluid.2003.08.011

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Exploring Molecular Pretraining and Mechanism-Aware Modeling for Reaction Yield Prediction

Yongkyu Lee[ab], Won Bo Lee[b], Lauren Ye Seol Lee[a]*

[a] Department of Chemical Engineering, University College London, Torrington Place, London, WC1E 7JE, United Kingdom Torrington Place, London, WC1E
[b] Seoul National University, Department of Chemical and Biological Engineering, Gwanak-ro 1, Gwanak-gu, Seoul 08826, Republic of Korea
* Corresponding Author: lauren.lee@ucl.ac.uk.

ABSTRACT

Accurately predicting chemical reaction yields can accelerate reaction optimization by prioritizing promising conditions; however, the complexity of multicomponent reactions and the limited availability of high-quality datasets remain significant challenges. While machine learning has achieved substantial progress in molecular property prediction, reaction-level modeling requires representations that capture three-dimensional structures, intercomponent interactions, and mechanistically critical features. In this study, we investigate the effective extension of molecular pretraining to reaction yield prediction from two complementary perspectives. First, we apply a reaction-aware architectural design based on molecular representation models pretrained via partial denoising to multicomponent reactions. Each reaction component is encoded as a three-dimensional stereoisomer embedding and concatenated into a reaction-level representation, with multi-head attention modeling intercomponent dependencies. We show that this approach achieves robust performance and generalizes well to reactions involving previously unseen components. Second, we design an architecture that extracts intermolecular relational embeddings through mechanistically meaningful pre-supervised learning using electron source-sink annotations derived from reaction mechanisms, followed by an active fine-tuning process. This mechanism-aware extension leverages large-scale reaction data to provide informative inductive bias and improves performance in data-constrained settings. We further confirm that this approach outperforms models that simply combine individual molecular embeddings, as it directly captures chemically meaningful information relevant to reaction outcomes. These results suggest promising directions for developing machine-learning models for reaction yield prediction that are more closely aligned with underlying chemical mechanisms.

Keywords: Deep learning, Reaction yield, Transfer learning

1. INTRODUCTION

Chemical reactions play a central role in diverse fields, including the synthesis of new compounds, process optimization, and materials development. Accurate prediction of chemical reaction outcomes can significantly reduce experimental costs and enable more efficient optimization strategies. However, reaction outcome prediction remains highly challenging due to the complex interplay of numerous factors, such as reaction conditions, catalysts, and the structural properties of reactants. Consequently, predicting reaction outcomes a priori, without performing physical experiments, remains an open and difficult problem.

Meanwhile, recent advances in deep learning have demonstrated strong performance in single-molecule property prediction tasks [1, 2], and the rapid growth of large-scale chemical reaction databases [3-5] has highlighted the potential of data-driven approaches to advance reaction modeling. Inspired by natural language processing, the BERT architecture has been applied to reaction yield prediction through large-scale pretraining on reaction corpora such as the Pistachio dataset[4]. The

resulting YieldBERT [6] model has shown superior performance compared to existing data-driven methods and has been widely adopted as a benchmark in subsequent studies [7]. However, as molecules are represented as linearized SMILES strings, key structural characteristics—such as explicit atomic connectivity and three-dimensional conformational information—are not directly encoded, which may limit the ability of such models to fully capture the chemical features governing reaction.

At the single-molecule level, numerous graph-based approaches have been developed to better capture intrinsic molecular characteristics [2, 8]. By representing molecules as graphs, these methods naturally encode atomic connectivity and, in some cases, incorporate three-dimensional coordinates into molecular embeddings. Pretraining strategies based on denoising techniques [2] have demonstrated the ability to learn physically meaningful representations beyond simple masking-based objectives. At the same time, reaction modeling inherently involves multiple interacting molecular species, and existing studies often address this challenge by aggregating or concatenating embeddings from individual molecules; however, it remains unclear whether such straightforward combinations are sufficient to fully capture complex inter-molecular interactions.

Accordingly, in this work we investigate reaction yield prediction using approaches specifically designed to address two essential aspects: intrinsic molecular representations that capture graph structure and three-dimensional conformational information, and data representations and architectures that explicitly model inter-molecular interactions. First, we adapt and extend a denoising-based molecular property prediction model to reaction-level prediction, focusing on capturing intrinsic structural characteristics of individual molecules. This approach demonstrates stable generalization in scenarios involving unseen products. Second, from a complementary perspective, we incorporate mechanistically relevant supervision by leveraging electron source-sink annotations derived from reaction mechanisms. Rather than treating interactions implicitly, the model is pretrained to identify atomic sites that play decisive roles in electron flow, and this mechanistic information is subsequently integrated into reaction-level embeddings during fine-tuning. Taken together, these results suggest that explicitly reflecting fundamental mechanistic aspects of chemical reactions, rather than relying solely on molecular aggregation or architectural complexity, provides a promising direction for developing more chemically informed machine-learning models for reaction yield prediction.

2. METHODS

2.1 Structure-Aware Molecular Representation for Reaction Yield Prediction

2.1.1 Overview

To incorporate three-dimensional molecular conformations and chemically meaningful representations for reaction, without performing expensive geometry optimization, we build upon the fractional denoising molecular pretraining framework, denoted as Frad, introduced by Ni *et al* [2]. Frad is a denoising-based pretraining approach based on TorchMD-Net [9] that learns structure-aware molecular representations by reconstructing molecular conformation perturbed from their equilibrium geometries. A key feature of Frad is a chemically aware noise (CAN) designed to capture physically relevant perturbations, including rotations and vibrational modes. This enables the model to encode conformational variability beyond small harmonic distortion and has been shown to produce molecular embeddings that improve molecular property prediction and remain robust to the quality of input geometry, including those generated using inexpensive classical methods. In this work, the pretrained Frad encoder is extended from single molecule property prediction to reaction-level modeling and reaction yield prediction on benchmark reaction datasets. Each reaction is represented as a set of molecular components, and their embeddings are aggregated to obtain a reaction-level representation, while treating the relationship between reaction components implicitly.

2.1.2 Input Construction and Model Architecture

Reactions are represented using three-dimensional structure-based encoding produced by the pretrained Frad molecular encoder, in which each molecular component is treated as an independent molecular graph. To accommodate reactions with a variable number of components, a dummy molecular graph with atomic number zero and fixed coordinates at (0, 0, 0) is introduced to represent missing species. This allows all reactions to be expressed in a consistent input format. Each component-specific graph is then processed independently by the pretrained encoder to generate a fixed-dimensional molecular embeddings. To adapt the Frad designed only for single-molecule for reaction-level modeling, the component embeddings are concatenated in a fixed order and pass through a multi-head attention module. This attention mechanism serves as a learnable aggregation function that adaptatively capture correlations among reaction constituents. The resulting interaction-aware representation is subsequently passed through two multi-layer perceptron (MLP) layers to produce a scalar prediction of the reaction yield (see Figure 1).

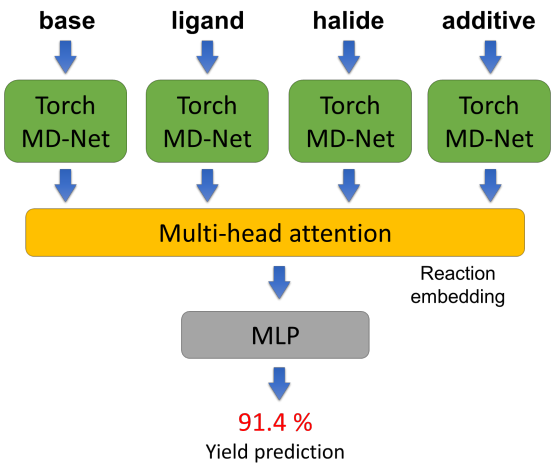

Figure 1. Reaction yield prediction using structure-aware molecular representations (SA-YIELD).

2.1.3 Datasets and Evaluation Protocol

Model performance for reaction yield prediction in this section is evaluated on the Buchwald-Hartwig coupling dataset [3], which was generated through high-throughput experimentation (HTE). Each reaction is defined by its constituent molecular components and represented using SMILES strings. In the HTE study, reaction conditions, such as concentration and temperature, are fixed across the dataset and are therefore not included as model inputs. The reactions involve four molecular components: aryl halide, ligand, base, and additive.

Reactions with zero yield frequently do not produce an isolable product, and including the product information as model input may introduce information leakage. To avoid this, each reaction instance is defined using only the four reactant-side components—aryl halide, ligand, base, and additive—while excluding the product from the input representation. This formulation ensures that yield prediction is based solely on information available prior to reaction execution and aligns with realistic experimental design constraints.

Model performance is evaluated under different generalization settings: random and component-wise unseen splits. For random split setting, reactions are randomly divided into training, validation, and test sets with a 6:2:2 ratio. To assess extrapolation to unseen molecular components, component-wise hold-out splits are applied. For bases and ligands, which consist of three and four distinct species, respectively, splits are generated by holding out one component at a time for testing, while using the remaining two bases and three ligands for training. For aryl halides and additives, which exhibit greater chemical diversity, unseen splits are generated by randomly selecting one quarter of the component combinations as the test set. Performance metrics reported for unseen component prediction are obtained by averaging results across all corresponding splits.

2.2 Interaction-Aware Pretraining Approach via Supergraph Representation

2.2.1 Overview

Many reaction prediction models form reaction representations by combining features derived from individual components, leaving cross-component relationship to be inferred implicitly from yield supervision. While the structure-aware formulation in Section 2.1 enriches molecular embeddings with conformation information, it similarly does not incorporate explicit mechanism derived relation across reaction components. In many cases, reaction outcomes depend on how specific atomic sites across different molecules participate in bond formation and electron transfer. Even for the same molecular species, the atoms involved in bond formation or braking can vary depending on the co-reactants, which limits what can be inferred from representations learned for each molecule in isolation.

To address this limitation, we introduce a mechanism-derived reaction modelling framework that incorporate explicit relational structure and supervision. First, reactions are represented using a pairwise supergraph formulation in which molecular components are decomposed into all combinatorial pairs. Second, electron source-sink annotations derived from reaction mechanisms are used as a pretraining objective to introduce mechanism-inspired relational bias into the learned representation. The pretrained model is then fine-tuned on a smaller reaction yield dataset, during which the predicted source-sink probabilities are incorporated to adaptively modulate how atomic information is aggregated into supergraph embedding, rather than relying on uniform pooling.

2.2.2 Supergraph Representation for interaction awareness

To encode intermolecular relationship explicitly, reactions are represented using a hierarchical, relation-centric supergraph formulation rather than as an aggregation of individual molecular embeddings. The model architecture is adapted from the HiMol [11] hierarchical graph framework, which augments atom-level graph with higher-order (e.g., fragment- or molecular-level) structural nodes to enable information exchange across multiple scales via deep message-passing neural network (D-MPNN)[12].

In the work, the hierarchical representation is extended from single molecules to reactions by constructing pairwise molecular supergraph. For a reaction involving multiple molecular components, all combinatorial molecular pairs are enumerated, and each pair is encoded as an independent supergraph. Within each supergraph, atom-level nodes capture local chemical structure, while higher-level nodes aggregate information at molecular-

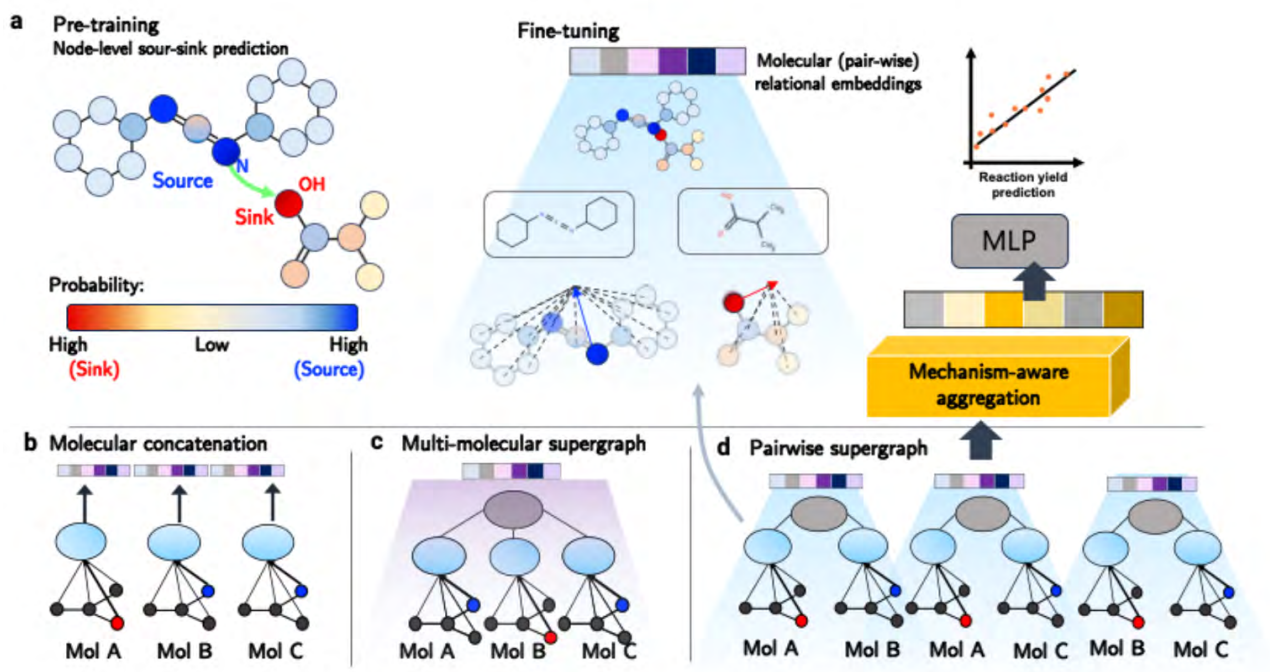

Figure 2. Mechanism-aware relational modeling for reaction yield prediction. (a) Source-sink pretraining and fine-tuning: predicted atom-level source-sink probabilities are used to weight atomic contributions during pooling when constructing pairwise relational embeddings for yield prediction. (b) Molecular concatenation baseline. (c) Multi-molecular supergraph. (d) Pairwise (combinatorial) supergraph formulation with mechanism-aware aggregation.

and reaction-levels. Message passing is then performed across all hierarchical levels of nodes, enabling information flow among atom-, molecular-, and reaction-level representations.

This pairwise construction is motivated by the observation that, even in multicomponent reactions, mechanistically relevant relationships are more naturally and effectively represented as pairwise interactions between specific components rather than as uniformly distributed multimolecular correlations. Key mechanistic process, such as electron flow and bond rearrangement, typically arises from interactions between two reacting species at a time, rather than from simultaneous many-body effects. Representing reactions using a pairwise supergraphs therefore provides a closer abstraction of the underlying reaction mechanism and aligns direclty with the pretraining objective, which is designed to capture electron source-sink relationships defined between interacting molecular pairs.

For reaction-level, embeddings extracted from all molecular pairs are integrated via a masked attention pooling. This aggregation yields a fixed-dimensional reaction representation and allows different molecular pairs to be weighted adaptively. A schematic overview of the pairwise supergraph representation and its integration into a reaction-level embedding is shown in Figure 2.

2.2.3 Source–Sink–Aware Pretraining and Active Finetuning for Reaction Yield Prediction

Reaction yield datasets are often limited in size, which constrains the ability of supervised models to learn complex intermolecular relationships reliably from yield labels alone. To mitigate this limitation, a transfer learning strategy is employed in which the pairwise supergraph encoder is first trained on a source-sink prediction task to introduce chemically motivated inductive bias, and is subsequently transferred to downstream reaction yield prediction.

Pretraining is performed using the Mech-USPTO-31k dataset [10], which provides reaction mechanisms annotated in an arrow-pushing format and include atomic- or bond-level labels identifying electron sources and sinks based on electron pair transfer. These annotations provide atom-level supervision that depends on the joint context of reacting molecular pairs and are therefore well aligned with the pairwise supergraph representation. Source-sink prediction is formulated as a node-level classification tasks on pairwise supergraphs, encouraging the model to assign higher likelihood to atomic sites associated with electron transfer in a reaction context.

Following pretraining, the model is fine-tuned on a liquid-phase amide coupling dataset reported by Zhang *et al* [13]. This dataset comprises reaction between

amines and carboxylic acids conducted in solvents in the presence of base and and coupling agent. During fine-tuning, pretrained parameters are used as initialization, and the source–sink prediction head, trained during pre-training as a classification objective with cross-entropy loss, is directly reused to generate softmax-normalized atom-level probabilities. These probabilities are employed as node-wise weighting factors during supergraph embedding aggregation, biasing the pooled representation toward atomic regions inferred to be more mechanistically relevant to electron transfer. The resulting weighted pairwise embeddings are then aggregated to form the final reaction-level representation for yield prediction. It is worth noting that the full distribution of source-sink probabilities is retained during pooling rather than selecting a single reactive atom, which improves robustness when reactive sites are distributed across multiple atoms or are uncertain at the atom level.

This pretraining-fine-tuning strategy allows mechanism-derived information learned from large-scale reaction data to influence downstream yield prediction indirectly, without requiring mechanistic annotations in the yield datasets. Figure 2 summarize our workflow.

2.3 Implementation Details

All experiments for the structure-aware reaction models (Section 2.1) were conducted on an NVIDIA A100 GPU, with batch size of 32 and the learning rate 5×10^{-5}. Each molecular component embedding had dimension 256, matching the dimensionality of the pretrained Frad encoder. Unless otherwise stated, all remaining hyperparameters, including the optimizer, scheduler, and activation functions, were kept identical to those used in the original Frad implementation [2]. All Frad encoders were initialized from checkpoints pretrained on the PCQM4Mv2 dataset. Reaction-level representations used a four-head multi-head self-attention module, with component embedding used as query, key, and value. Attention outputs were concatenated before being passed to the subsequent fully connected layers. Each task was trained independently and required approximately 12 hours to complete.

Experiments for the pairwise supergraph model with source-sink pretraining (Section 2.2) were conducted on a single NVIDIA RTX 4090 GPU. For source-sink pretraining on Mech-USPTO-31k, a random 0.85:0.15 train-validation split was used with AdamW optimizer (learning rate 1×10^{-4}, batch size 256). Source-sink prediction was formulated as a graph-wise node selection task, where each atom in a given (super)graph is treated as a candidate class and the number of classes thus varies with graph size. Two parallel heads (source and sink) are trained with cross-entropy loss (with non-reactive nodes masked), and predictions are obtained by taking the argmax over atoms for each head. For yield prediction, the pretrained model was fine-tuned on the amide coupling dataset[13] using a random 0.85:0.15 split, learning rate 1×10^{-6} and batch size 16. Reactions were represented as a hierarchical graph with atom-, molecule-, and molecule-pair-level nodes, using a five-layer D-MPNN backbone, with embeddings projected to dimension 256. Pairwise embeddings were aggregated using masked four-head attention pooling. During fine-tuning, predicted source-sink probabilities were used for probability-weighted pooling. Dropout rates of 0.1 (encoder) and 0.5 (yield head) were applied.

3. RESULTS

3.1 Reaction Yield Prediction with Structure-Aware Molecular Representations

This section evaluates whether incorporating three-dimensional, structure-aware molecular representation improve reaction yield prediction, with particular emphasis on generalization to reactions involving previously unseen components. The proposed structure-aware reaction model (SA-YIELD), described in Section 2.1, is evaluated on Buchwald-Hartwig coupling reaction dataset.

First, we note that SA-YIELD achieves a MAE of 4.18 under a random split setting. This performance is comparable to that reported for Yield-BERT [5], which achieves an average MAE of approximately 4.0 on the same dataset. Considering that Yield-BERT relies on representations pretrained on reaction SMILES from the USPTO patent corpus (approximately 3.7 million reactions), these results indicate that incorporating structure-aware molecular representations at the reaction level can achieve competitive yield prediction accuracy without requiring geometries derived from crystal structures or large-scale reaction pretraining.

Next, we compare the SA-YIELD to representative baseline models, sequence-based Yield-BERT [5] and structure-aware CRYSTAL-YIELD7. All performance of components splits are summarized in Table 1. In this setting, prediction accuracy decreases substantially for all models, as it assesses extrapolation to reactions containing held-out components. For example, the MAE is approximately 4 in random split, whereas it increases into a substantially higher range for reactions involving unseen components, reaching their largest values for unseen aryl halides up to 24.

This behavior reflects the fact that aryl halides span a broader chemical space and play a central role in determining reactivity in Buchwald–Hartwig couplings. As primary electrophilic partners, they directly determine the identity of the reacting electrophile and play a central role in bond formation. Withholding such components therefore induces a pronounced shift in the underlying yield distribution, which limits extrapolation based solely on patterns learned from the remaining reaction

components. This effect is substantially more pronounced than for other reaction partners, such as ligands, and exposes limitations in models that encode reactants in isolation, without explicitly accounting for how reactivity depends on interacting reaction partners.

Despite these challenges, SA-YIELD consistently achieve the lowest MAE across all unseen-component categories. The average MAE is reduced from 24.4 for Yield-BERT and 15.65 for CRYSTAL-YIELD to 13.5, with the largest improvement observed for aryl halides—the most challenging unseen component, where SA–YIELD achieves an MAE of 18.8. Comparisons between models with (SA-YIELD) and without pretraining (SA-YIELD-RAW) indicate that these gains primarily arise from reaction-level architectural design choices rather than from denoising-based molecular pretraining alone.

While the Frad framework effectively captures three-dimensional molecular structure, molecular-level pretraining alone provides limited benefit for extrapolative prediction. This suggests that accurately modeling reaction outcomes requires representations that go beyond intrinsic molecular features. This motivates a shift toward reaction-level architectures and supervision strategies that explicitly reflect reaction mechanisms, which we explore in the following section.

Table 1: Reaction yield prediction performance on the Buchwald-Hartwig coupling dataset under component-wise unseen splits. Reported values are mean absolute error (MAE) in yield percentage points (mean ± s. e. across splits where applicable)

Model	Base	Lig-and	Halide	Addi-tive	Aver-age
Yield-BERT	24.3 ± 1.6	24.3 ± 1.4	24.7 ± 2.1	24.1 ± 0.7	24.4
CRYSTAL-YIELD	13.4 ± 0.3	11.7 ± 2.2	21.3 ± 3.3	16.2 ± 0.4	15.65
SA-YIELD	11.8 ± 1.6	11.3 ± 2.6	18.8 ± 0.8	12.1 ± 1.3	13.5
SA-YIELD-RAW	11.5 ± 1.5	11.7 ± 3.0	19.5 ± 1.4	12.1 ± 1.3	13.7

3.2 Reaction Yield Prediction from explicit relation embedding

3.2.1 Source-Sink Pretraining Performance

To assess the effectiveness of mechanism-aware pretraining, the pretrained model was evaluated on the electron source-sink identification task defined by the Mech-USPTO-31k dataset. The evaluation follows the first-step source-sink labeling scheme described in Section2.2.3. Consistent with the dataset formulation, source–sink prediction is treated as a graph-wise node selection problem, in which each atomic node within a pairwise molecular supergraph is assigned a relative score, and the source and sink atoms are identified by selecting the highest-scoring nodes for each role. These scores are normalized via a softmax operation during training with cross-entropy loss, allowing the model to learn a single source and a single sink atom per supergraph.

As summarized in Table 2, our model demonstrates high accuracy in both source and sink predictions. On the test set comprising 4, 705 reactions, 4, 288 cases (91.2%) ere correctly predicted for both source and sink components. Partial predictions were observed in 264 cases with only the source correctly identified and in 137 cases with only the sink correctly identified, indicating greater difficulty in sink-site prediction. Notably, only 16 reactions (0.3%) were misclassified for both sites, highlighting the robustness of the model in distinguishing reactive and non-reactive atomic positions. This reliability allows the predicted source–sink signals to be used for active fine-tuning. Rather than treating source-sink prediction as a standalone objective, the present model incorporates source-sink supervision within a pairwise molecular supergraph architecture that explicitly encodes interactions between reacting species. In this framework, source-sink prediction is not an endpoint, but a means of biasing learned representations toward interaction-dependent atomic features.

Table 2: Prediction Outcomes on the Test Set

Outcome	Counts	%
Both correct	4288	91.1
Source only	264	5.6
Sink only	137	2.9
None	16	0.3

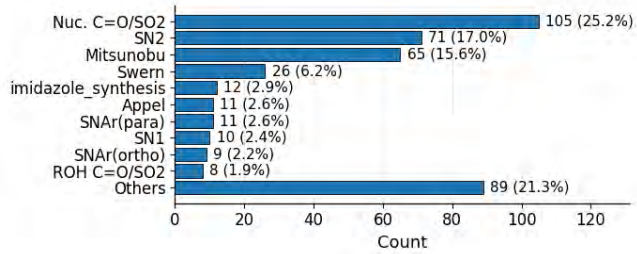

Figure 3. Reaction-Class Distribution of Failed Predictions

To examine the limitations of the pre-training approach, we analyzed the reaction classes associated with prediction errors, including both complete and partial mispredictions (Figure 3). Prediction errors were frequently observed in nucleophilic attacks on (thio)carbonyl or sulfonyl groups, S_N2 reactions, and Mitsunobu-type transformations, accounting for 105 (25.2%), 71 (17.0%), and 65 cases (15.6%), respectively. As discussed in Ref 10, S_N2 reactions are inherently challenging to annotate using discrete electron-flow representations, as bond formation and cleavage occur concertedly and multiple mechanistic interpretations can be chemically

plausible. In nucleophilic attack reactions involving (thio)carbonyl or sulfonyl groups, misassignments frequently involved intramolecular sites rather than external nucleophiles, reflecting a fundamental limitation of the molecular representation, as the 2D graph encoding does not capture steric accessibility or approach geometry. For Mitsunobu reactions, detailed inspection revealed that a substantial fraction of errors originated from the symmetric structure of diethyl azodicarboxylate, in which two nitrogen atoms are chemically equivalent, and the model frequently selected the alternative symmetric site that was treated as incorrect under the current labeling scheme despite being mechanistically equivalent. These observations indicate that many Mitsunobu-related mispredictions arise from annotation artifacts rather than genuine mechanistic failures, suggesting that the effective prediction accuracy of our model is likely higher than the nominal value reported here.

Despite these limitations, the pretrained model exhibits informative probabilistic behavior. Even when the highest predicted source or sink probability does not coincide with the ground-truth, the model frequently assigns relatively high probabilities to other chemically plausible atomic sites. Because the downstream framework does not rely on discrete source-sink assignments but instead incorporate the full probability distribution during the pooling, such partially correct predictions can still contribute useful interaction-relevant information.

3.2.2 Effect of Mechanism-Aware Pretraining and Active Finetuning

Reaction yield prediction performance on the amide coupling dataset is summarized in Table 3, comparing different reaction embedding strategies with and without mechanism-aware pretraining. Three embedding formulations are evaluated. Molecular concatenation aggregates independently encoded molecular embeddings without explicit intermolecular message passing and serves as a non-relational baseline. Multi-molecular supergraph encodes all reaction components within a single graph, enabling message passing across all atoms simultaneously. Pairwise supergraph, as described in Section 2.2.2, constructs relational embeddings by encoding molecular components as combinatorial molecular pairs. The first two settings are included as ablation baselines to assess the effect of explicitly modeling pairwise intermolecular relationships.

When trained without source-sink pretraining, the model based on molecular concatenation achieve slightly better performance than the pairwise supergraph. In this setting, reaction yield labels alone provide limited information about which molecular pairs and atomic regions are responsible for determining the reaction outcome. As a result, the model is unable to reliably identify and exploit pair-specific interaction patterns from the amide coupling dataset, which contains on the order of 500 reactions. which limits the effectiveness of relational representations and can lead to degraded generalization performance.

In contrast, when the pairwise supergraph encoder is initialized using source-sink pretraining and actively leveraging the mechanistic information, a substantial improvement in prediction accuracy is observed. This improvement reflects the ability of the large-scale pretraining to bias the model toward interaction-dependent atomic features that are relevant to reaction outcomes. Because source-sink annotations encode electron-transfer roles that arise from interactions between specific reacting species, the resulting inductive bias is naturally aligned with pairwise relational representations and cannot be readily captured by models that operate solely on independently encoded molecules.

Finally, comparison between relational formulations shows that the pairwise (combinatorial) supergraph outperforms the multi-molecular supergraph when combined with mechanism-aware pretraining. This behavior is consistent with the structure of the source-sink supervision, which is defined at the level of interacting molecular pairs rather than across all reaction components simultaneously. Encoding all molecules within a single supergraph can dilute pair-specific information, whereas the pairwise construction preserves interaction-specific contexts prior to aggregation. In summary, these results confirm that the proposed framework effectively pretrains mechanistically meaningful interaction information and selectively deploys it during fine-tuning, providing a principled and data-efficient strategy for yield prediction.

Table 3: Effect of embedding strategy and source-sink pretraining on amide coupling yield prediction. Reported values are MAE on the held-out test set

Embedding strategy	No pre-training	With pre-training
Molecular concatenation	0.17	-
Multi-molecular supergraph	-	0.16
Pairwise supergraph	0.19	0.14

4. CONCLUSION

In this study, two complementary approaches for reaction yield prediction were developed and evaluated, incorporating three-dimensional molecular structure and mechanism-derived relational information. First, a denoising-based molecular representation model pretrained on three-dimensional conformational data was extended from single-molecule property prediction to reaction-level modeling. This structure-aware formulation improved robustness under component-wise extrapolation, although molecular-level pretraining alone provided limited gains when transferred directly to reaction yield

prediction.

Second, a mechanism-aware framework was introduced that represents reactions using a pairwise supergraph formulation and leverages electron source-sink annotations for pretraining. By biasing representations toward interaction-dependent atomic regions associated with electron transfer, this approach improved yield prediction in data-constrained settings and outperformed direct training without such pretraining. These results indicate that reaction-level learning benefits from aligning representation granularity and pretraining tasks with mechanistically motivated constructs.

The findings suggest that yield prediction models benefit from moving beyond aggregation of isolated molecular features and instead incorporating interaction-specific, mechanism-derived information. While the present study used two-dimensional graph representations for mechanistic pretraining and was limited by available dataset scale, it establishes a basis for integrating three-dimensional structure and richer mechanistic annotations in future work.

ACKNOWLEDGEMENTS

This research was supported by Korea Institute for Advancement of Technology (KIAT) grant funded by the Ministry of Trade, Industry & Energy (MOTIE), Korea Government (RS-2024-00436106, Human Resource Development Program for Industrial Innovation).

AUTHOR IDENTIFIERS

Author ORCIDs:
Lee Y: 0009-0007-5662-0494
Lee WB: 0000-0001-7801-083X
Lee L: 0000-0003-1847-7838

REFERENCES

1. Jiang J, Ke L, Chen L, Dou B, Zhu Y, Liu J, Zhang B, Zhou T, Wei G. Transformer technology in molecular science. WIREs Comput Mol Sci 14: (2024). https://doi.org/10.1002/wcms.1725
2. Ni Y, Feng S, Hong X, Sun Y, Ma WY, Ma ZM, Ye Q, Lan Y. Pre-training with fractional denoising to enhance molecular property prediction. Nat Mach Intell 6:1169-1178 (2024). https://doi.org/10.1038/s42256-024-00900-z
3. Ahneman DT, Estrada JG, Lin S, Dreher SD, Doyle AG. Predicting reaction performance in C–N cross-coupling using machine learning. Science 360:186-190 (2018). https://doi.org/10.1126/science.aar5169
4. Pistachio (Nextmove Software). http://www.nextmovesoftware.com/pistachio.html.
5. Lowe DM. Extraction of chemical structures and reactions from the literature. In: University of Cambridge, editor.; 2012.
6. Schwaller P, Vaucher AC, Laino T, Reymond JL. Prediction of chemical reaction yields using deep learning. Mach. Learn.: Sci. Technol. 2:015016 (2021). https://doi.org/10.1088/2632-2153/abc81d
7. King-Smith E. Transfer learning for a foundational chemistry model. Chem. Sci. 15:5143-5151 (2024). https://doi.org/10.1039/d3sc04928k
8. Fang X, Liu L, Lei J, He D, Zhang S, Zhou J, Wang F, Wu H, Wang H. Geometry-enhanced molecular representation learning for property prediction. Nat Mach Intell 4:127-134 (2022). https://doi.org/10.1038/s42256-021-00438-4
9. Singh J, Singh C. Learning invariant representations for equivariant neural networks using orthogonal moments. 2022 International Joint Conference on Neural Networks (IJCNN) :1-8 (2022). https://doi.org/10.1109/ijcnn55064.2022.9892412
10. Chen S, Babazade R, Kim T, Han S, Jung Y. A large-scale reaction dataset of mechanistic pathways of organic reactions. Sci Data 11: (2024). https://doi.org/10.1038/s41597-024-03709-y
11. Zang X, Zhao X, Tang B. Hierarchical molecular graph self-supervised learning for property prediction. Commun Chem 6: (2023). https://doi.org/10.1038/s42004-023-00825-5
12. Yang K, Swanson K, Jin W, Coley C, Eiden P, Gao H, Guzman-Perez A, Hopper T, Kelley B, Mathea M, Palmer A, Settels V, Jaakkola T, Jensen K, Barzilay R. Analyzing learned molecular representations for property prediction. J. Chem. Inf. Model. 59:3370-3388 (2019). https://doi.org/10.1021/acs.jcim.9b00237
13. Zhang C, Lin Q, Yang C, Kong Y, Yu Z, Liao K. Intermediate knowledge enhanced the performance of the amide coupling yield prediction model. Chem. Sci. 16:11809-11822 (2025). https://doi.org/10.1039/d5sc03364k

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Synergistic integration of direct air capture in bioenergy systems

Nor Syuriaty Jaafar[a], Norhuda Abdul Manaf[ab*], Noor Fatina Emelin Nor Fadzil[a], and Nilay Shah[b*],

[a] Department of Chemical and Environmental Engineering, Malaysia-Japan International Institute of Technology (MJIIT), Universiti Teknologi Malaysia, Kuala Lumpur, Malaysia
[b] Department of Chemical Engineering, Imperial College London, SW7 2AZ, United Kingdom
* Corresponding Author: norhuda.kl@utm.my, n.shah@imperial.ac.uk

ABSTRACT

The present work aims to demonstrate the synergy achieved through the integration of biomass gasification with a direct air capture (DAC) system to maximize overall CO_2 removal capacity, while simultaneously converting waste into value-added products (hydrogen) and supplying the energy required for DAC operation (BG-H_2P-DAC). The proposed configuration is modeled using Aspen Plus to investigate the synergistic interactions and key performance indicators of the BG-H_2P-DAC system. Parametric analyses are conducted by varying gasification temperature, air inlet flow rate, and amine concentration and flow rate. The results indicate that increasing the monoethanolamine (MEA) concentration from 10% to 40% leads to a gradual decline in CO_2 capture efficiency, accompanied by a reduction in CO_2 slip. The system achieves a net specific electricity consumption of 0.0293 MWh/t CO_2, confirming that the electricity generated from the integrated steam power cycle is sufficient to fully offset the electrical requirements of the DAC process. The regeneration heat requirement at a steam temperature of 150 °C is 1.32 MWh/t CO_2, which represents the total net thermal demand of the DAC unit.

Keywords: Carbon Capture, Aspen Plus, Energy, Hydrogen, Environment

INTRODUCTION

Deep decarbonization pathways aimed at limiting the world's surface temperature increases to no more than 1.5°C have placed significant pressure on the global energy system. One promising avenue to achieve this target is by deploying the renewable sources for clean energy alongside negative emission technologies (NETs).

Sustainable biomass plays a vital role towards the transition to decarbonized economy with broad applications [1]. Gasification of biomass features an effective technology to enhance thermochemical efficiency and curtailing environmental impact. The benefits result from the syngas production composes of hydrogen that serves as the primary fuel in the combustion process. Which render vital potential for multiple products such as methanol, hydrogen, power, and freshwater production[2].

Direct Air Capture (DAC) has emerged as a game changing NET due to its ability to capture CO_2 directly from the atmosphere (demand-side) [3], in contrast to other carbon capture technologies that primarily focus on the supply-side. DAC operates by removing CO_2 directly from ambient air through chemical processes using liquid solvents or solid sorbents usually requiring high and low regeneration temperatures.

Few studies have explored the potential of integrating biomass gasification, DAC couple/hybrid with other technologies. A novel integrated sustainable aviation fuel (SAF) production system combining DAC, CO_2 electrolysis, and biomass gasification was proposed and evaluated by [3]. The system achieved an overall energy efficiency of 54.7 % and a negative global warming potential estimated at -339.7 g CO_2-e/MJ SAF. In another study, [4] designed a high-temperature direct air capture (HTDAC) system integrated with Fischer-Tropsch synthesis (FTS) for renewable jet fuel production. The proposed configuration resulted in a fuel specific energy consumption about 122 MJ/kg jet fuel with DAC accounting for about 16 % of the total energy demand. From a thermodynamic perspective, the theoretical minimum energy required to separate CO_2 from air using a DAC

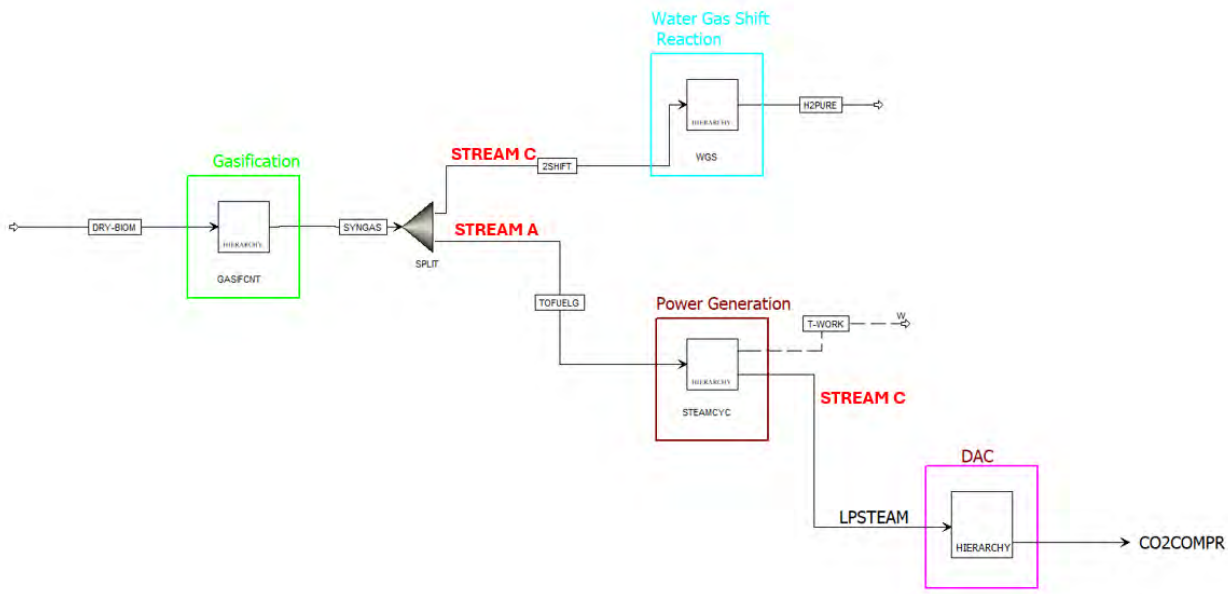

Figure 1: Hierarchies block in the Aspen Plus interface for BG-PH2-DAC system.

system and deliver it in pure form at 1 atm has been reported to be approximately 0.5 GJ per tonne of CO_2 [5]. The integration of nuclear small modular reactor (SMR) with a low-temperature solid sorbent DAC system was evaluated by [6]. Based on the DAC plant's output and energy requirement, a thermal input of 51.25 MW was required necessitating the SMR to supply 100°C steam at this capacity. Utilizing this amount corresponds to an estimated 15% reduction in electricity generation. An innovative low-CO_2 emission system that integrates biomass gasification, monoethanolamine (MEA)-based CO_2 capture, and methanol synthesis for waste-to-energy and power-to-fuel applications was proposed by [7]. This system generates 75 MW of electricity with the capture unit exhibits a capture efficiency of 59.24 % requiring 1.92 MW of power per kilogram of CO_2 captured. While extensive research has been conducted on various technology integration configurations, limited research has focused on integrated application of biomass gasification combined cycle systems for simultaneous power and hydrogen production coupled with DAC (BG-PH_2-DAC).

Research contribution

Biomass gasification has been evident to make a significant contribution to achieve net-zero/negative emissions [8, 9]. Synergizing biomass gasification with a combine cycle system has been proven to enhance power generation efficiency and environmental performance compared to a conventional coal fired power plant [10].

A major challenges associate with DAC lies in its high CO_2 capture cost [11] and substantial demand of low carbon thermal and electrical energy [6]. Where, electricity is required to operate fans and auxiliary equipment, while a considerable amount of heat to regenerate the capture medium and release the captured CO_2. Despite its considerable mitigation potential, the large-scale commercialisation and deploying DAC of systems remain challenging. Although recent improvements have brought DAC closer to commercial feasibility, high regeneration energy and substantial energy requirement continue to represent a critical barrier to widespread adoption [5, 12].

The ability of biomass gasification to generate power in its downstream processes offer a significant advantage for integrated industrial ecosystem that include DAC. Present work demonstrates this synergy at a computational level (technology readiness level 2) serving as a preliminary contribution toward the development of a smart industrial ecosystem. Furthermore, this work highlights how the low-carbon energy required for DAC can be accommodated by the power produce from gasification system. The high-purity hydrogen produced via the water-gas shift reaction in biomass gasification enhances the efficiency of industrial ecosystems by enabling multiple energy vectors from renewable biomass sources.

METHODOLOGY

Process description

A simplified diagram of BG-PH_2-DAC system proposed in this study is illustrated in Figure 1. The diagram shows the hierarchical blocks used in Aspen Plus simulation consisting of biomass gasification, water gas shift (WGS), steam cycle and DAC processes. In this work, enriched air is employed as the gasification agent to produce syngas. A study has confirmed that using air as a gasification agent can produce high-purity H_2 [13], demonstrating the technical feasibility of this approach. Based on Figure 1, the syngas is controlled and divided into multiple streams for the co-production of electricity (Stream A), low pressure steam (Stream B) and hydrogen (Stream C). For Stream A, the syngas is cooled before entering the combustion unit, where it is burned with compressed air. The hot combustion gases drive a gas turbine and waste heat is recovered through a series of heat exchangers to generate power in combine cycle. At the down-stream process, energy required by the DAC is accommodated by generated power (assume as a pseudo stream/connection) and bleed steam is used for solvent regeneration at stripper. For Stream C, the syngas undergoes a water-gas shift reaction to increase its hydrogen content.

Model development

Biomass integrated gasification combine cycle coupled with water gas shift (BG-PH_2) process

A simplified biomass integrated gasification combine cycle model is adapted from a study conducted by [14], while WGS process is based on the simplified model developed in [15, 16]. Biomass gasification is modelled in Aspen Plus using a non-conventional solid approach combined with equilibrium-based reactors. The process consists of biomass decomposition, gasification, particulate removal, and syngas cleaning prior to its use in a combined cycle gas turbine. In gasification model, woodchip is used as fuel with their proximate and ultimate composition is based on [13, 17]. Based on the experimental study conducted by [13], utilization of woodchips as fuel source can produce almost H_2 purity of ≥ 99.977 vol.%

The WGS reactor converts CO and steam into H_2 and CO_2 through the reversible chemical reaction: $CO + H_2O \rightleftarrows CO_2 + H_2$. During this process, raw syngas is mixed with steam to achieve high purity H_2. The cooled syngas mixer then enters the RStoic reactor containing iron-based catalyst where perfect mixing is assumed for modelling purposes. For simplification, a single-stage WGS reactor is employed in this study following the process suggested by [16]. This process involves high temperature stage between 300 – 500°C. Finally, the tail gas stream (CO_2 + unconverted gases) is removed by the scrubber and high purity H_2 is released on the top scrubber column. Details of the model development and Aspen Plus unit operations involved is presented in Appendixes S1 and S2.

Direct air capture with CO_2 compression

The DAC model is developed and adapted from [5, 11, 18] utilizing the mature MEA absorption process modified to suit DAC environmental conditions. This technology is deployed as it holds a promise as a low-cost and scalable solution for DAC [19]. Details of the model development and Aspen Plus unit operations involved is presented in Appendixes S1 and S2.

Hybrid BG-PH2-DAC

The individual models are interconnected as illustrated in Figure 1. A splitter unit operation is used to divide the syngas between H_2 production and power generation at 50% ratio. For steam supply to the reboiler, a MHeatX unit operation is employed. In the stripper column (Radfrac), reboiler option is disable as steam is supplied externally from a low-pressure steam. The individual models are validated independently before being integrated to ease the convergence process.

RESULT AND DISCUSSION

Synergistic analysis of BG-PH2-DAC

Synergistic analysis through sensitivity evaluation is conducted by manipulating several operating variables such as gasification temperature, air inlet (at DAC), and amine solvent concentration and flowrate. The analysis is conducted based on nominal operating condition of BG-PH_2-DAC plant as tabulated in Appendix S2.

Figures 2(a)-(b) show the strong influence of gasifier temperature on syngas composition and the subsequent enhancement of hydrogen yield through the WGS reaction during biomass gasification. In Figure (a), increasing the gasifier temperature from 300 °C to around 700 °C markedly increases the H_2 and CO mole fractions, indicating enhanced devolatilization, char gasification, and endothermic reactions such as steam reforming [20] and the Boudouard reaction [21]. The H_2 mole fraction reaches a maximum at approximately 700 °C, beyond which it slightly declines, suggesting the increasing dominance of competing high-temperature reactions and which shift the equilibrium[20]. In contrast, [20] obtained the highest H_2 fraction at 800 °C, suggesting the existence of optimal condition for efficient H_2 production. The CO_2 shows a decreasing trend with temperature, reflecting its consumption in gasification and affect of the endothermic character of Boudouard reaction [22].

After the WGS reactor (Figure 2(b)), H_2 started to produce at 300 °C and increased drastically until 700 °C, with H_2 mole fraction significantly increasing and stabilizing at around 34–35% for temperatures above 700 °C, while CO_2 correspondingly increases due to CO conversion. This plateau indicates that the WGS reaction approaches equilibrium and all methane had been converted into hydrogen. Similar trend was observed in the

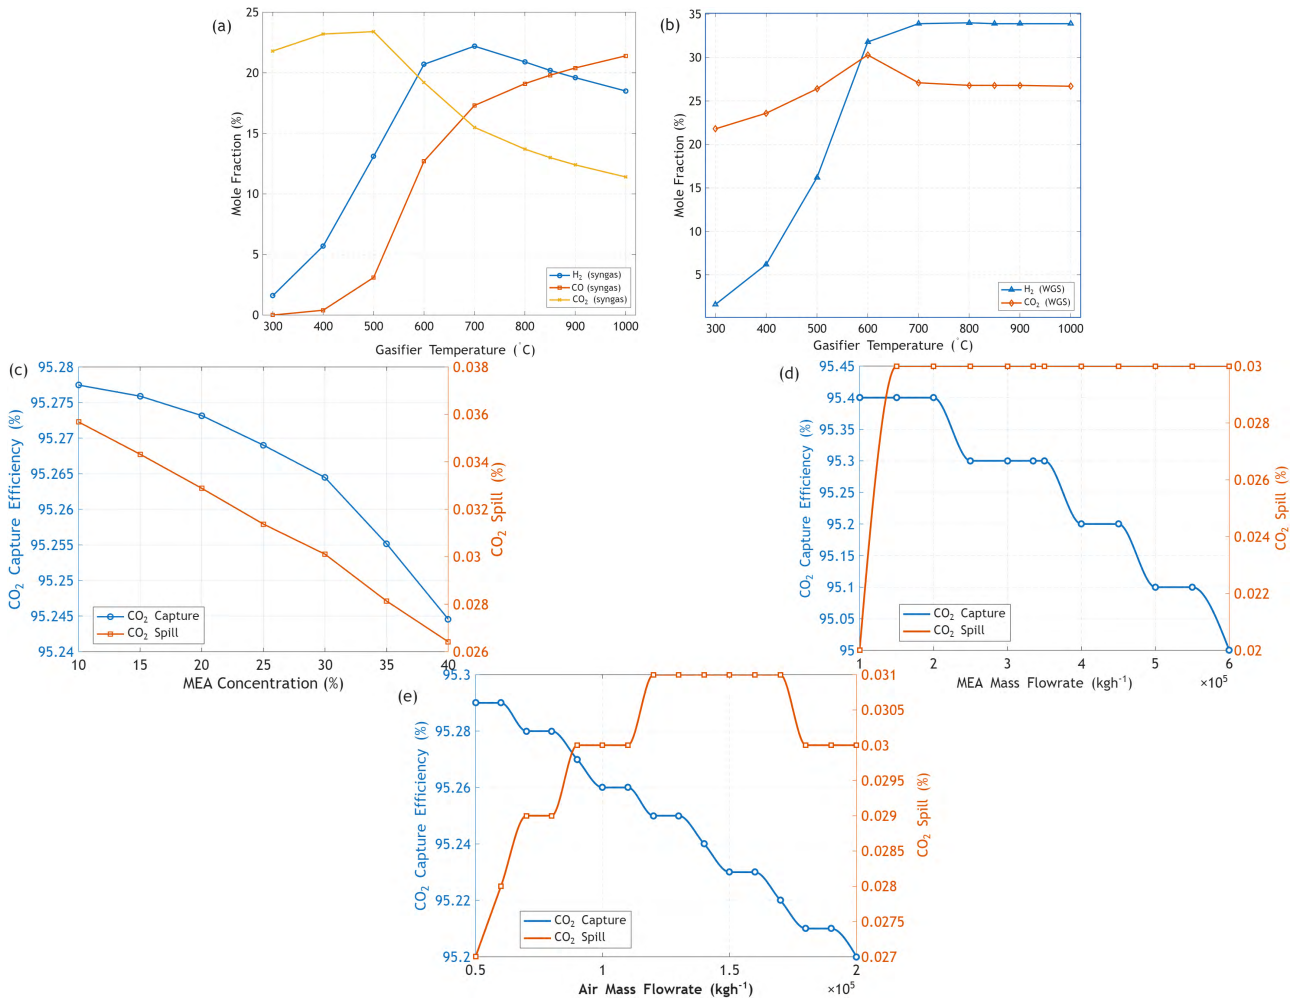

Figure 2. Synergisitc effect of BG-PH$_2$-DAC BG at nominal operating condition, in which a single parameter is varied (x-axis) while all other operating conditions are maintained constant. (a) Variation of gasification temperatue with syngas composition, (b) Variation of gasification temperature with gas composition after WGS, (c) Variation of MEA concentration with CO_2 capture rate and CO_2 spill (top absorber), (d) Variation of MEA mass flowrate with CO_2 capture rate and CO_2 spill (top absorber), (e) Variation of air mass flowrate with CO_2 capture rate and CO_2 spill.

study conducted by [15].

Figures 2 (c)–(e) present a synergistic analysis of the downstream process (DAC), highlighting the sensitivity of CO_2 capture efficiency and CO_2 slip (off-gas at absorber column) to key operating variables. In Figure 2(c), increasing the MEA concentration from 10% to 40% results in a gradual decline in CO_2 capture efficiency, accompanied by a reduction in CO_2 slip. This trend suggests that while higher MEA concentrations enhance the solvent's CO_2 absorption capacity, mass transfer limitations and high liquid-to-gas volumetric flow ratio may hinder overall capture efficiency under DAC conditions [23]. Additionally, higher lean loading (i.e. higher MEA concentration) can lead to low productivity (low capture efficiency) due to increased pressure drop within the column [11]. Figure 2(d) shows the effect of MEA mass flow rate, where an increase initially sustains high capture efficiency within the range of 95.0 – 95.4%. This behavior indicates that, within this specific range of MEA flowrate, the flowrate ensures complete wetting of the packing and effective gas–liquid contact for reaction and capture [23]. Meanwhile, CO_2 slip remains relatively stable, implying that solvent availability is no longer the limiting factor. As illustrated in Figure 2(e), increasing air mass flow rate exhibits minimal impact on capture efficiency, exhibiting a slight decreasing trend accompanied by a corresponding increase in CO_2 slip. This trend indicates that higher air flow rates may minimize the required contactor inlet area [23], though this may subsequently diminish the capture efficiency. This interpretation is based on the qualitative trend rather than a quantitative approach.

Key performance analysis of BG-PH2-DAC

Table 1 summarizes the key performance indicators of the integrated BG-PH_2-DAC system operating under nominal conditions (Appendix 2). The system demonstrates marginal hydrogen production, achieving a specific H_2 yield of 0.0292 kg H_2/biomass with an almost complete hydrogen purity (~100%). This indicates efficient syngas conversion and effective downstream purification. Simultaneously, the process generates 66.25 kgCO_2/biomass, of which approximately 98% is successfully captured, highlighting the strong integration between biomass gasification, hydrogen production, and DAC subsystem.

The specific CO_2 capture energy requirement is quantified at 4.74 MJ/kg CO_2, indicating a moderate thermal energy intensity for solvent regeneration. This result is consistent with the findings of [24] who reported a regeneration duty of approximately 4.0 MJ/kg CO_2 when increasing the reboiler temperature from 123 to 150 °C. This minor deviation can be attributed to differences in solvent operating conditions and process integration strategies. In terms of power generation, the system delivers a gross electrical output of approximately 6.5 MWe, corresponding to a specific gross power output of 0.302 MWh/t CO_2.

As the present work employed a simplified BG-PH_2-DAC plant design, the electricity consumption for the air contactor is adopted from [6] at 0.232 MWh/t CO_2. Electricity consumption for CO_2 compression is calculated at 0.04 MWh/t CO_2 with auxiliary pumping at 0.0007 MWh/t CO_2. This results in a total DAC electricity demand of 0.2727 MWh/t CO_2. It can be observed that the pumping energy requirement in this study is slightly lower than the 0.007 MWh/t CO_2 reported in [24], as only a single pump is considered in the present configuration. Larger discrepancies are noted in CO_2 compression work, where other studies reported higher values of 0.0944 MWh/t CO_2 (0.34 MJ/kg CO_2) [25] and 0.132 MWh/t CO_2 (0.475 MJ/kg CO_2) [26].These differences are primarily attributed to variations in compressor design assumptions and isentropic efficiencies. Overall, the electrical energy demand (for DAC) reported here is comparable with the value reported by [27] with a difference of 0.0327 MWh/t CO_2. As a benchmark, a standalone DAC system typically requires about 1.4 MWh/t CO_2 under the external supply of power and heat sources [24].

The resulting net specific electricity of 0.0293 MWh/t CO_2 confirms that the electricity generated from the integrated combined power cycle is sufficient to fully offset the electrical requirements of the DAC process. This internal energy integration substantially enhances overall system efficiency and reinforces the feasibility of the BG-PH_2-DAC configuration. In this work, net specific electricity is defined as the gross power generated from the steam cycle minus the total electrical demand of the DAC unit. On the thermal side, the desorption heat requirement at a steam temperature of 150 °C is 1.32 MWh/t CO_2, which constitutes the total net thermal demand of the DAC unit. This value is slightly higher than the 1.14 MWh/t CO_2 reported in the literature [27], reflecting differences in solvent regeneration conditions and heat integration assumptions. Overall, the combination of high CO_2 capture efficiency, high-purity H_2 production, and effective internal energy utilization underscores the technical viability and energetic coherence of the BG-PH_2-DAC system.

CONCLUSION

This study provides a preliminary insight for the BG-PH_2-DAC system that aim to anchor the decarbonization effort while producing multiple energy vectors. A simplified configuration proposed in this work evident the feasibility of integrated system to sufficiently accommodate electricity and thermal energy within the system. Existing work can be enhanced by considering comprehensive plant design with the integration of pressure swing adsorption and detailed combined cycle system to reflect actual operation design.

DIGITAL SUPPLEMENTARY MATERIAL

Digital supplementary is available in Living Archive for Process Systems Engineering at http://PSEcommunity.org/LAPSE:2026.0037

Table 1: Key performance analysis of BG-PH_2-DAC at nominal operating condition (extracted and derived from Aspen stream summary).

Key Performance	Value	Unit
Specific H_2 production	0.0292	kg H_2/biomass
H_2 purity	~100	%
Specific CO_2 production	66.25	kg CO_2/biomass
CO_2 capture rate	~98	%
Specific CO_2 capture	4.74	MJ/kg CO_2
Gross power output	~6.5	MW_e
Energy requirement		
Electrical		
Specific power output	0.302	MWh/tCO_2
Air contactor (DAC)	0.232 [6]	MWh/tCO_2
CO_2 compressor (DAC)	0.04	MWh/tCO_2
Pump (DAC)	0.0007	MWh/tCO_2
Thermal		
Desorption heat (at steam 150°C)	1.32	MWh/tCO_2
Net specific (thermal)	1.32	MWh/tCO_2
Net specific (electricity)	0.0293	MWh/tCO_2

ACKNOWLEDGEMENTS

This study was supported by the Ministry of Higher Education under Fundamental Research Grant Scheme (FRGS/1/2024/TK08/UTM/02/17).

AUTHOR IDENTIFIERS

Abdul Manaf N: 0000-0002-4548-6425
Shah N: 0000-0002-8906-6844

REFERENCES

1. Saravanan A, Ragini YP, Karishma S, Hemavathy RV, Jyotsna M. A review on advancing sustainable energy: the role of biomass and bioenergy in a circular economy. Sustainable Futures 10:100835 (2025). https://doi.org/10.1016/j.sftr.2025.100835
2. Lu P. Biomass gasification for multi-generation: methanol, hydrogen, power, and freshwater production with exergy and economic insights. Thermal Science and Engineering Progress 68:104321 (2025). https://doi.org/10.1016/j.tsep.2025.104321
3. Kim D, Won W, Park J. An integrated process for sustainable aviation fuel production via direct air capture, carbon dioxide electrolysis, and biomass gasification. Energy Conversion and Management 346:120465 (2025). https://doi.org/10.1016/j.enconman.2025.120465
4. Paulsen MM, Petersen SB, Lozano EM, Pedersen TH. Techno-economic study of integrated high-temperature direct air capture with hydrogen-based calcination and fischer–tropsch synthesis for jet fuel production. Applied Energy 369:123524 (2024). https://doi.org/10.1016/j.apenergy.2024.123524
5. Barzagli F, Giorgi C, Mani F, Peruzzini M. Screening study of different amine-based solutions as sorbents for direct co₂ capture from air. ACS Sustainable Chem. Eng. 8:14013-14021 (2020). https://doi.org/10.1021/acssuschemeng.0c03800
6. Slesinski D, Litzelman S. How low-carbon heat requirements for direct air capture of CO2 can enable the expansion of firm low-carbon electricity generation resources. Front. Clim. 3: (2021). https://doi.org/10.3389/fclim.2021.728719
7. Nouri A, Tohidi F, Chitsaz A. Integration of biomass gasification and mea-based CO2 capture for sustainable power and methanol production: energy, exergy, and economic analysis and optimization. Energy 319:135062 (2025). https://doi.org/10.1016/j.energy.2025.135062
8. Lu X, Cao L, Wang H, Peng W, Xing J, Wang S, Cai S, Shen B, Yang Q, Nielsen CP, McElroy MB. Gasification of coal and biomass as a net carbon-negative power source for environment-friendly electricity generation in china. Proc. Natl. Acad. Sci. U.S.A. 116:8206-8213 (2019). https://doi.org/10.1073/pnas.1812239116
9. Jaiganesh N, Kuo PC, Champatan V, Gopi G, Kumar RA, Aravind PV. Negative emission power plants: techno-economic analysis of a biomass-based integrated gasification solid oxide fuel cell/gas turbine system for power, heat, and biochar co-production—part 2. Front. Energy Res. 10: (2022). https://doi.org/10.3389/fenrg.2022.826227
10. Darmawan A, Tokimatsu K, Aziz M. Integrated gasification combined-cycle power plants utilizing biomass waste: energy-efficient recovery. Green Energy and Technology :109-136 (2025). https://doi.org/10.1007/978-981-97-9360-0_4
11. de Joannis P, Castel C, Kanniche M, Favre E, Authier O. Direct air capture by monoethanolamine absorption with heat pump enhancements. Ind. Eng. Chem. Res. 64:2208-2225 (2025). https://doi.org/10.1021/acs.iecr.4c02033
12. Kiani A, Jiang K, Feron P. Techno-economic assessment for CO2 capture from air using a conventional liquid-based absorption process. Front. Energy Res. 8: (2020). https://doi.org/10.3389/fenrg.2020.00092
13. Gubin V, Kadlez D, Bartik A, Steiner L, Zeitlhofer J, Thelen F, Hofbauer H, Müller S. Hydrogen production from woody biomass via fixed-bed gasification at pilot-scale. International Journal of Hydrogen Energy 114:462-474 (2025). https://doi.org/10.1016/j.ijhydene.2025.02.357
14. Jana K, De S. Biomass integrated combined power plant with post combustion co2 capture – performance study by ASPEN plus®. Energy Procedia 54:166-176 (2014). https://doi.org/10.1016/j.egypro.2014.07.260
15. Amrana, U.I., A. Ahmadb, and M.R. Othman, Kinetic based simulation of methane steam reforming and water gas shift for hydrogen production using aspen plus. Chemical Engineering, 2017. **56**.
16. Bosch, C. and W. Wild, *Hydrogen production.* Industrial & Engineering Chemistry, 1915. **7**(1): p. 83-83.
17. Chung H, Lim Y, Son SH. Industrial-scale biogas-to-hydrogen: integrated steam biogas reforming–water-gas shift modeling and multi-objective optimization for energy-efficient hydrogen production. International Journal of Hydrogen Energy 195:152520 (2025). https://doi.org/10.1016/j.ijhydene.2025.152520
18. Soares Chinen A, Morgan JC, Omell B, Bhattacharyya D, Tong C, Miller DC. Development

of a rigorous modeling framework for solvent-based co₂ capture. 1. hydraulic and mass transfer models and their uncertainty quantification. Ind. Eng. Chem. Res. 57:10448-10463 (2018). https://doi.org/10.1021/acs.iecr.8b01471

19. Kiani A. Session 24. oral presentation for: amine based liquid capture technology for direct air capture of CO2 – an overview on technology development. Australian Energy Producers Journal 65: (2025). https://doi.org/10.1071/ep24424
20. Liu S, Ben-Abdelwahed A, Gaborieau C, Lachaud J. Biomass char gasification for hydrogen production: a thermodynamic equilibrium analysis. Fuel Processing Technology 276:108279 (2025). https://doi.org/10.1016/j.fuproc.2025.108279
21. González JF, Román S, Bragado D, Calderón M. Investigation on the reactions influencing biomass air and air/steam gasification for hydrogen production. Fuel Processing Technology 89:764-772 (2008). https://doi.org/10.1016/j.fuproc.2008.01.011
22. Li Z, Yang T, Yuan S, Yin Y, Devid EJ, Huang Q, Auerbach D, Kleyn AW. Boudouard reaction driven by thermal plasma for efficient CO2 conversion and energy storage. Journal of Energy Chemistry 45:128-134 (2020). https://doi.org/10.1016/j.jechem.2019.10.007
23. Li G, Yao J. Direct air capture (DAC) for achieving net-zero CO2 emissions: advances, applications, and challenges. Eng 5:1298-1336 (2024). https://doi.org/10.3390/eng5030069
24. Kiani A, Jiang K, Feron P. Techno-economic assessment for CO2 capture from air using a conventional liquid-based absorption process. Front. Energy Res. 8: (2020). https://doi.org/10.3389/fenrg.2020.00092
25. Keith DW, Holmes G, St Angelo D, Heidel K. A process for capturing CO2 from the atmosphere. Joule 2:1573-1594 (2018). https://doi.org/10.1016/j.joule.2018.05.006
26. Sabatino F, Grimm A, Gallucci F, van Sint Annaland M, Kramer GJ, Gazzani M. A comparative energy and costs assessment and optimization for direct air capture technologies. Joule 5:2047-2076 (2021). https://doi.org/10.1016/j.joule.2021.05.023
27. Slesinski D, Litzelman S. How low-carbon heat requirements for direct air capture of CO2 can enable the expansion of firm low-carbon electricity generation resources. Front. Clim. 3: (2021). https://doi.org/10.3389/fclim.2021.728719

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

An Adaptive Framework for Robust Energy Forecasting under Concept Drift and Feature Uncertainty

Francesco Marcato[a], Alessio Santecchia[a], Manuel Ruivo de Oliveira[a], Francesco Silvestri[b], Rafael Castro-Amoedo[c]*

[a] Emissium Labs Unipessoal LDA, Alcácer do Sal, Portugal
[b] AIDA Lab (Algorithms for Intelligent Data Analytics), Department of Information Engineering, University of Padova
[c] Centro de Recursos Naturais e Ambiente (CERENA), Instituto Superior Técnico, Lisbon, Portugal;
* Corresponding Author: rafael.amoedo@tecnico.ulisboa.pt

ABSTRACT

The rapid integration of renewable energy sources is increasing the volatility and non-stationarity of modern power systems, posing significant challenges for data-driven forecasting models. In particular, concept drift and uncertainty in exogenous inputs such as weather forecasts can severely degrade predictive performance over time. This work proposes a lightweight two-layer forecasting framework that decouples prediction from adaptation. A traditional offline regression model is augmented by an online meta-learner that continuously generates adaptive meta-features, enabling the system to respond to structural changes and noisy inputs without repeated retraining. The framework is evaluated on two real-world case studies. First, concept drift is addressed in nuclear power production forecasting, where abrupt and gradual capacity changes are inferred through an online meta-learner. Second, feature uncertainty is mitigated in day-ahead solar production forecasting by correcting noisy weather forecast inputs. Across both scenarios, the proposed approach consistently outperforms single-layer baselines, reducing root mean squared error by up to 10% and maintaining robust performance over multi-year horizons without retraining. The results demonstrate that meta-learning provides a practical and computationally efficient mechanism for improving forecast robustness in non-stationary energy systems, with applicability to a wide range of power-system forecasting problems.

Keywords: Algorithms, Big Data, Modelling, Artificial Intelligence, Energy

INTRODUCTION

The increasing penetration of renewable energy sources is fundamentally changing the operating conditions of modern power systems. Weather-dependent generation, evolving market mechanisms, and policy-driven structural changes are introducing significant non-stationarity into electricity production and demand time series. Accurate forecasting remains essential for grid operation, market participation, and emerging carbon-aware applications [1, 2], yet conventional data-driven models often struggle when historical patterns no longer hold.

Supervised learning models trained offline remain widely used in energy forecasting applications [3]. While advanced deep learning architectures and sliding-window gradient boosting ensembles represent the current state-of-the-art for raw predictive accuracy, their high computational overhead and requirement for continuous retraining limit their scalability in real-time operational environments. However, these models implicitly assume stationarity in the underlying data-generating process. Violations of this assumption lead to systematic degradation of predictive performance over time, motivating the development of adaptive forecasting strategies.

Concept Drift and Adaptive Forecasting

Concept drift refers to changes in the statistical relationship between input variables and target outputs over time as described by Tsymbal [4], who formalised the concept and identified several drift patterns, including abrupt, gradual, and incremental drift. In energy systems, concept drift arises from structural events such as power plant commissioning or decommissioning, fuel

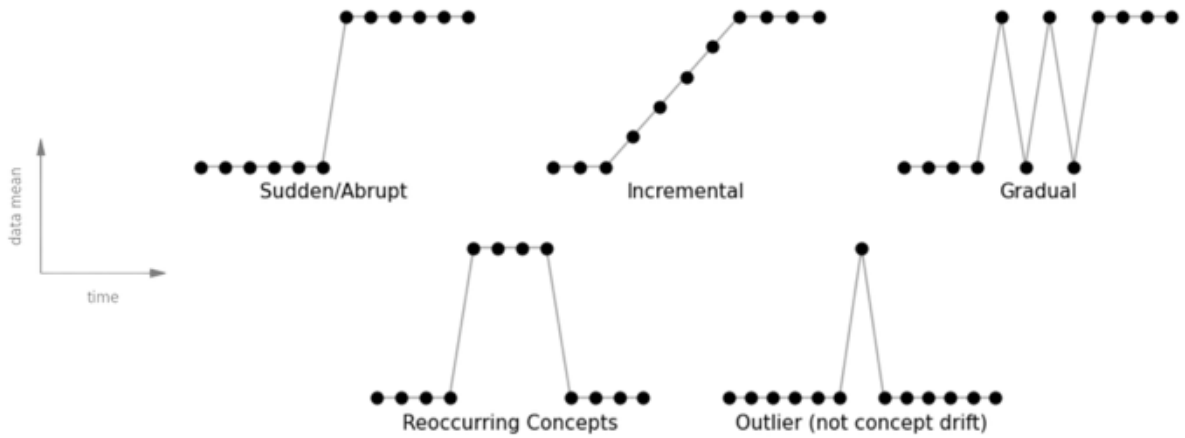

Figure 1. – Canonical drift patterns

switching, and regulatory interventions.

These drift types are illustrated schematically in **Figure 1**, which provides a conceptual reference for the experimental scenarios studied later in this work, and illustrates typical concept drift patterns, motivating the need for adaptive forecasting strategies capable of handling both abrupt and gradual changes.

Several approaches have been proposed to address concept drift in regression problems, including sliding-window retraining [5], weighted learning [6], and fully online adaptation [7]. While effective, retraining-based methods are computationally expensive, fully online learners may suffer from instability. On the other hand, Ensemble-based adaptive methods offer a compromise by combining robustness and adaptability.

Adaptive ensemble methods have gained significant attention for non-stationary data streams. Gomes et al. [8] introduced the Adaptive Random Forest (ARF), which integrates online bagging, per-tree drift detection, and background learners to enable localised adaptation. ARF has been successfully applied to a variety of streaming regression problems [9].

Efficient open-source implementations of these methods are provided by modern streaming machine-learning libraries, such as the River library in Python [10], facilitating their deployment in operational forecasting systems.

Uncertainty in Exogenous Inputs

In addition to structural drift, which refers to physical changes in the system (like equipment aging) rather than the statistical shifts of concept drift, forecasting accuracy is strongly affected by uncertainty in exogenous inputs. Renewable generation forecasting, particularly for solar and wind power, relies heavily on meteorological predictions that are inherently uncertain. Errors and biases in weather forecasts propagate directly into production forecasts, often resulting in large prediction errors [11].

Various strategies have been proposed to mitigate this effect, including probabilistic forecasting [12], noise injection during training [13], and ensemble-based uncertainty modelling [14]. More recently, meta-learning approaches have been explored to correct biased, noisy, or missing inputs online without retraining the main forecasting model [15].

Contribution

This paper proposes a two-layer meta-learning framework designed to address both concept drift and feature uncertainty without repeated retraining of the main forecasting model. The key contributions are:

1. A unified two-layer architecture that separates prediction and adaptation through meta-feature generation.
2. An online meta-learning strategy that adapts to structural changes and noisy inputs in streaming data.
3. Validation on real-world energy datasets covering both capacity-driven concept drift and weather-driven feature uncertainty.

METHODOLOGY

The proposed framework combines an online meta-learner with a base forecaster that undergoes offline training. In the present implementation, the base forecaster is a linear regression model trained using ordinary least squares (implemented using river [0.23.0], scikit-learn [1.7.2]). All experiments were conducted in Python [3.12.10]. The model is trained once on historical data, and its coefficients remain fixed during the evaluation

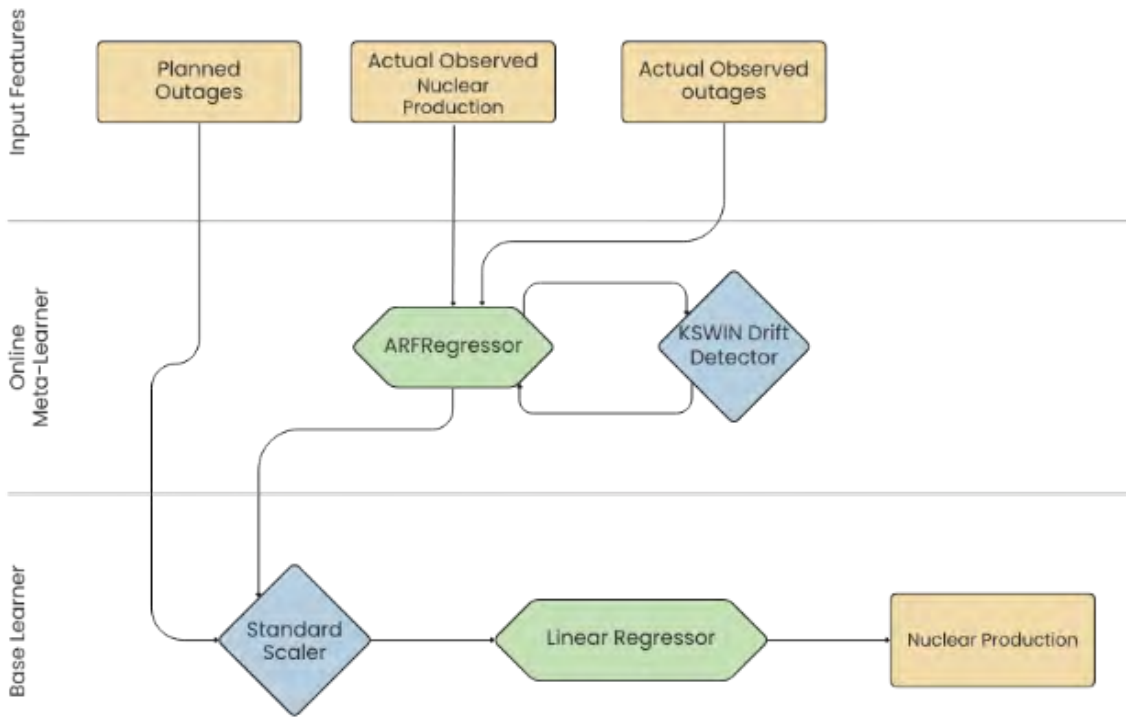

Figure 2. – Graphical representation of the two-layer framework, with an online meta-learner estimating operational capacity, which is appended to the offline forecaster as a meta-feature

period. The choice of a simple regression model for the base forecaster was deliberate. While modern forecasting often employs complex ensembles, utilizing a static offline baseline allows the isolated impact of the adaptive meta-feature to be measured cleanly, proving that our lightweight layer can rescue legacy models without the need for computationally expensive sliding-window retraining. This designation signifies that the model is pretrained on historical data and its internal weights remain fixed throughout the deployment phase. To enable adaptation without retraining the forecaster, the system augments the input with meta-features generated online.

Predictions $\hat{y}_t$ at time t are computed as:

$$\hat{y}_t = f(x_t, m_t),$$

where $x_t$denotes the original feature vector and m_t is a meta-feature generated by:

$$m_t = g(z_t),$$

with $g(\cdot)$ representing the meta-learner and z_t auxiliary information such as recent observations, forecast errors, or uncertain inputs. During deployment, the framework operates sequentially on the data stream. At each time step, the meta-learner receives the most recent observations and updates its internal state. The resulting meta-feature is appended to the original feature vector and passed to the fixed offline forecaster, which produces the final prediction. This streaming workflow enables continuous adaptation without modifying the parameters of the base forecasting model.

The overall architecture of the proposed two-layer framework is shown in **Figure 2**, which illustrates the interaction between the fixed offline forecaster and the online meta-learner. The separation of prediction and adaptation is central to the proposed approach and enables efficient handling of non-stationarity. In the nuclear forecasting case study (see below), the meta-learner estimates a latent variable representing the effective operational capacity of the nuclear fleet. This estimate is derived incrementally from the production time series and the outage signal available in the data stream. At each time step, the Adaptive Random Forest updates its estimate using the newly observed production value, producing a capacity-related feature that reflects the current operational state of the system. This estimate is then appended to the original feature vector as a single additional input to the offline forecaster.

The meta-learner is implemented using an Adaptive Random Forest operating in streaming mode as described by Gomes et al. [7]. The Adaptive Random Forest implementation used in this work is based on the River library [10]. Across both case studies the ARFRegressor always uses the default (n_models=10) unless otherwise stated. In the first case study, we override the default detector with KSWIN, whereas in the second case study, we

use the default ARF configuration, which includes ADWIN detectors. The detectors act to identify distributional changes and trigger local model updates. Unless otherwise specified, the remaining hyperparameters follow the default configuration provided in the River implementation. Internal drift detection mechanisms allow individual trees within the ensemble to adapt locally, while the forecaster remains fixed, reducing computational overhead compared to retraining-based approaches. Models are compared using a set of metrics, including the coefficient of determination (R^2), the Root Mean Squared Error (RMSE), the Mean Absolute Error (MAE), the Mean Squared Error (MSE) and the Symmetric Mean Absolute Percentage Error (sMAPE).

All experiments were implemented in Python using the River streaming machine learning library (version 0.23.0). To ensure exact reproducibility of the Adaptive Random Forest's stochastic processes, the global random seed was fixed at 42 across all experiments. The full experimental workflow, including data preprocessing and model configuration, follows the methodology described above and can be reproduced using the same datasets and parameter settings.

Case Study I: Concept Drift in Nuclear Power Forecasting

The first case study evaluates the framework under structural concept drift in nuclear power production. Nuclear generation is generally stable but subject to abrupt changes due to plant shutdowns, making it a suitable testbed for drift handling.

The meta-learner is tasked with estimating the operational capacity of the nuclear fleet, a latent variable not directly observable from production data alone. This estimated capacity is provided as a meta-feature to an offline regression model trained prior to major structural changes. Experiments were conducted on national-level nuclear production data for Germany, Switzerland, and the Netherlands, sourced from the ENTSO-E transparency platform [17]. The dataset consists of hourly aggregated production measurements spanning the period 2018–2024. Production values were aligned with outage information obtained from the European Energy Exchange (EEX). The resulting time series was cleaned for missing values and resampled to hourly resolution to ensure temporal consistency across all variables.

Case Study II: Feature Uncertainty in Solar Power Forecasting

The second case study addresses feature uncertainty in day-ahead solar power forecasting. Solar generation forecasts rely heavily on weather predictions, which are inherently uncertain and may exhibit systematic bias.

In this scenario, the meta-learner operates as a feature uncertainty corrector. It learns the relationship between forecasted and observed solar radiation in an online fashion and outputs corrected radiation estimates as meta-features. The forecaster uses these corrected inputs to predict solar production.

The framework was evaluated on aggregated Swiss solar production data using three configurations: a baseline model trained on forecasted weather (using data from 2023-01-01 to 2024-01-01), the proposed two-layer model with corrected features, and an upper-bound model trained on observed weather over the same training interval. Solar radiation forecasts were obtained from the Visual Crossing Weather API and correspond to day-ahead meteorological predictions. Observed solar radiation values were aligned with the forecasts at hourly resolution. The forecasting model uses the radiation forecast together with the meta-corrected radiation estimate produced by the meta-learner.

RESULTS AND DISCUSSION

Structural concept drift: Nuclear Power Production

Nuclear power production constitutes a particularly challenging forecasting problem under concept drift. Unlike renewable generation, which is primarily driven by stochastic exogenous variables, nuclear production is dominated by structural constraints, namely installed capacity and plant availability. When these constraints change due to policy decisions or outages, the mapping between historical inputs and future outputs is fundamentally altered.

This effect is particularly evident in Germany, where the progressive nuclear phase-out resulted in multiple abrupt capacity reductions over the study period. The timing and magnitude of these events are shown in **Figure 3**, which highlights the strongly non-stationary nature of the production time series. Switzerland and the Netherlands exhibit milder forms of drift, characterised by gradual capacity changes and shorter outage periods.

These three countries, therefore, provide a spectrum of drift regimes, allowing the robustness of the proposed framework to be assessed under increasing levels of structural disruption.

We compared our proposed Two-Layer Framework (TLF) against a Single-Layer (SL) baseline defined as an offline linear regression model using only the outage signal as input. In the TLF, an online Adaptive Random Forest meta-learner estimates a latent capacity-related feature from the data stream (hourly resolution), and a second, offline linear regression main model predicts production using both outage and the meta-learner's output. Models were trained on 2018-01-01 to 2020-01-01 and evaluated on 2020-01-01 to 2024-01-01. Forecasts were generated sequentially in a walk-forward evaluation setting. At each time step in the test period, the meta-

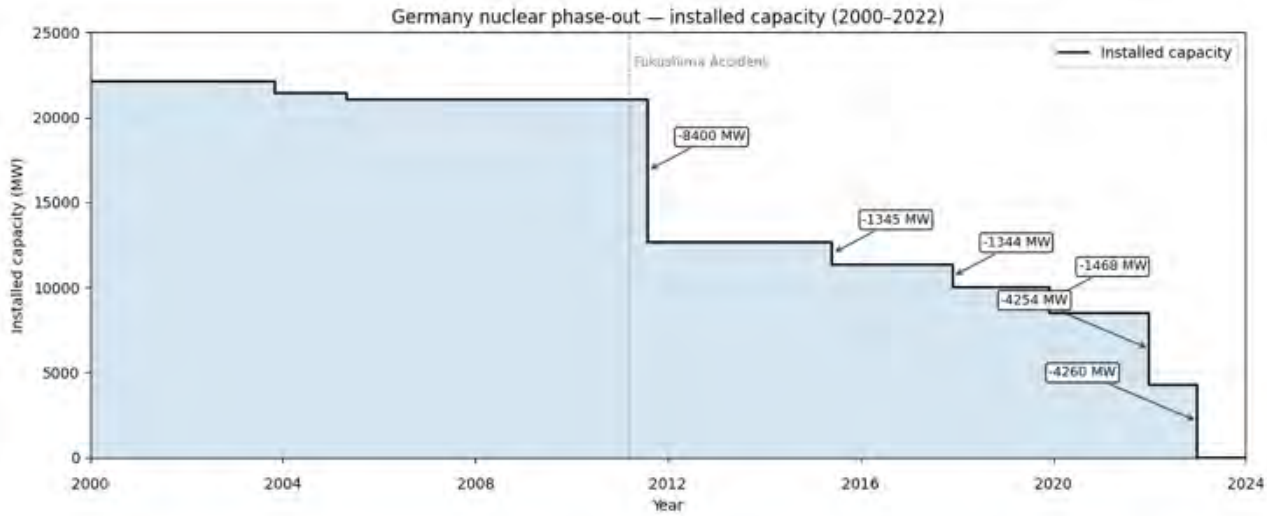

Figure 3. – Timeline for Germany's nuclear power plants phase-out culminating in April 2023. Source : Mez et al. [16].

learner was updated using the newly observed production value before the next prediction was produced.

Model performance was evaluated using several complementary error metrics, including RMSE, MAE, sMAPE, and the coefficient of determination (R^2). These metrics capture both absolute prediction accuracy and relative forecasting performance across different production scales. A quantitative comparison of forecasting performance is reported in **Table 1**, which aggregates RMSE, MAE, sMAPE and R^2 metrics for Germany, Switzerland, and the Netherlands.

Table 1 – Performance comparison of a single-layer (SL) traditional model and our two-layer framework with meta-learner (TLF) for Germany (DE), Switzerland (CH) and the Netherlands (NL).

Country	Model	R^2 [-]	RMSE [MW]	MAE [MW]	sMAPE [%]
DE	**SL**	-1.5547	4525	3714	71.8
	TLF	0.9824	375	293	40.0
CH	**SL**	0.6966	369	348	14.4
	TLF	0.9732	110	34	1.8
NL	**SL**	0.9644	26	10	12.3
	TLF	0.9698	24	6	9.5

Table 1 shows that the two-layer framework consistently outperforms the single-layer baseline in the presence of concept drift. In the Netherlands, where structural changes are limited, both models achieve similar performance, indicating that the meta-learning layer does not degrade accuracy under quasi-stationary conditions. This is an important result, as it demonstrates that the proposed approach introduces no penalty when adaptation is not strictly required.

In Switzerland, where gradual capacity changes and recurring outages occur, the two-layer framework yields a clear improvement over the baseline. RMSE reductions on the order of several percentage points are observed, alongside a noticeable increase in R^2. This indicates that the meta-learner is able to track slow structural changes effectively, providing the forecaster with a more accurate representation of the system's operational state.

Germany represents the most challenging case. Here, the baseline model rapidly loses predictive power following major shutdown events, as it continues to rely on outdated capacity assumptions. As shown in **Table 1**, the baseline model for Germany yields a severely negative R^2 (-1.5547). It is important to note that the coefficient of determination (R^2) measures the model goodness-of-fit and is not strictly equivalent to the squared Pearson correlation coefficient (r^2). Consequently, R^2 is not bounded at zero and can drop below -1 when a model's predictions perform substantially worse than simply predicting the historical mean of the observed data. In this case, because the static baseline model continuously relies on outdated capacity assumptions after abrupt phase-out events, its predictions exhibit extreme systematic bias, mathematically resulting in this highly negative value. The two-layer framework significantly mitigates this degradation, achieving substantially lower error metrics over extended periods. While performance still deteriorates during extreme shutdown phases, the meta-learning approach preserves a meaningful level of predictive accuracy where the baseline effectively fails.

While aggregate metrics provide a useful summary,

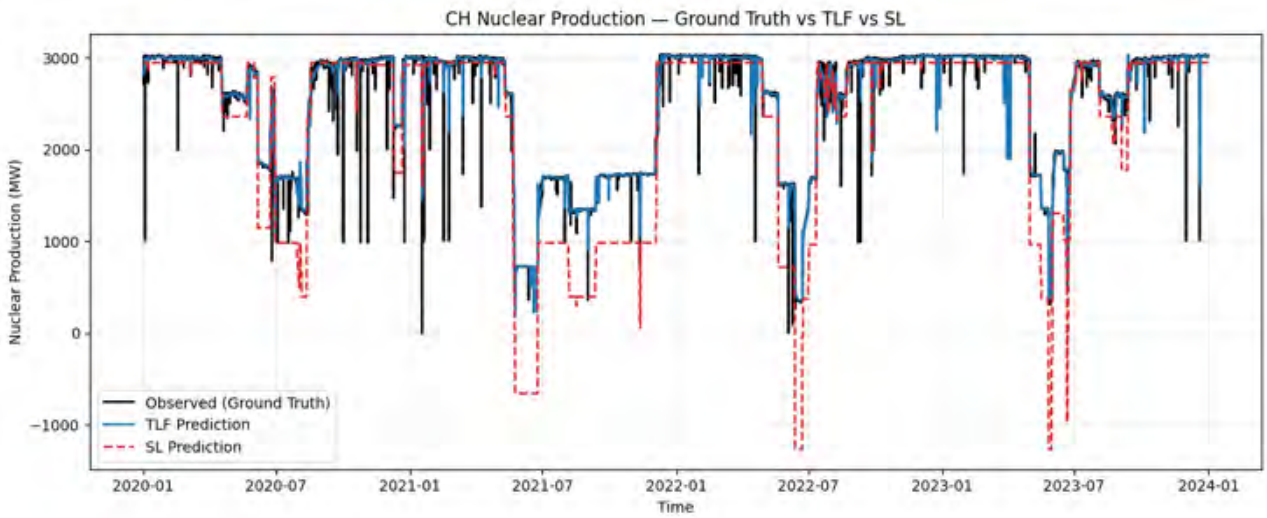

Figure 4. – Predictions of the Swiss nuclear production using the two-layer framework. The model tracks structural change better than the single-layer baseline.

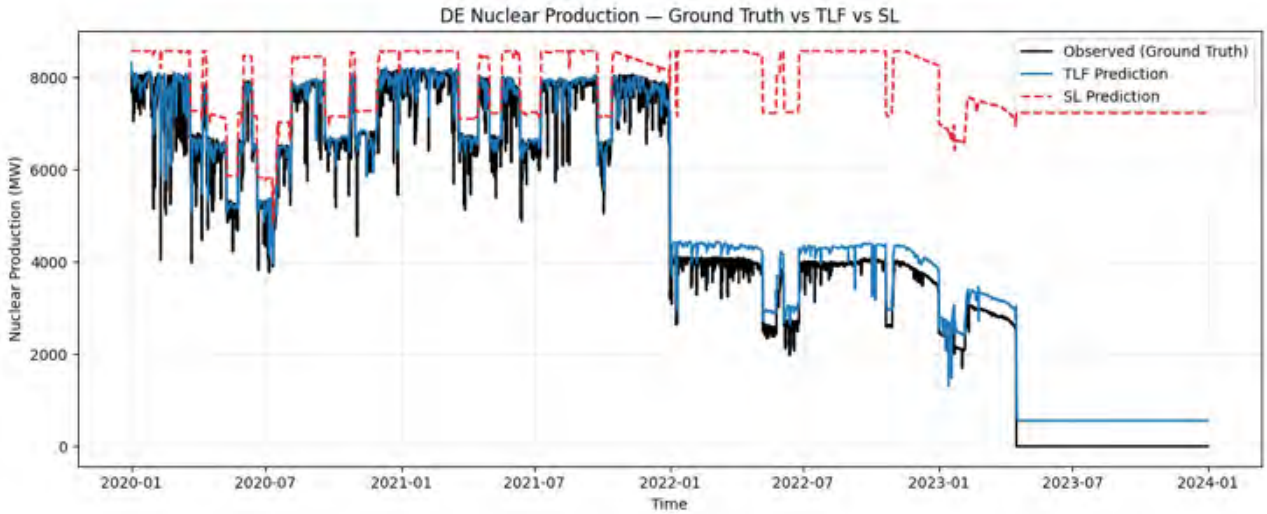

Figure 5. – Predictions of the German nuclear production using the two-layer framework. Despite improvements over the baseline, large abrupt capacity shutdowns remain challenging.

they obscure the temporal dynamics of adaptation. These dynamics are illustrated in **Figures 4 and 5**, which show predicted and observed nuclear production for Switzerland and Germany, respectively.

Figure 4 demonstrates how the two-layer framework maintains close alignment with observed production over time, even as capacity evolves. In contrast, the baseline model exhibits systematic bias during extended outage periods, reflecting its inability to adapt to changing structural constraints.

Figure 5 provides a more nuanced picture. Following abrupt shutdown events, both models experience an immediate increase in error. However, the two-layer framework adapts more rapidly, reducing bias and variance over subsequent time steps. This behaviour reflects the ability of the Adaptive Random Forest meta-learner to detect distributional change and update its internal representation accordingly.

Importantly, the results also reveal the **limits of adaptation**. When nuclear capacity approaches zero, the latent-variable assumption underpinning the meta-learner becomes invalid. In such cases, the model cannot fully recover, highlighting a fundamental limitation of any approach that relies on historical production patterns. This observation is critical and underscores the importance of explicitly acknowledging failure modes in adaptive forecasting systems.

From an engineering perspective, these results demonstrate that the proposed framework offers a practical alternative to frequent retraining. Structural capacity

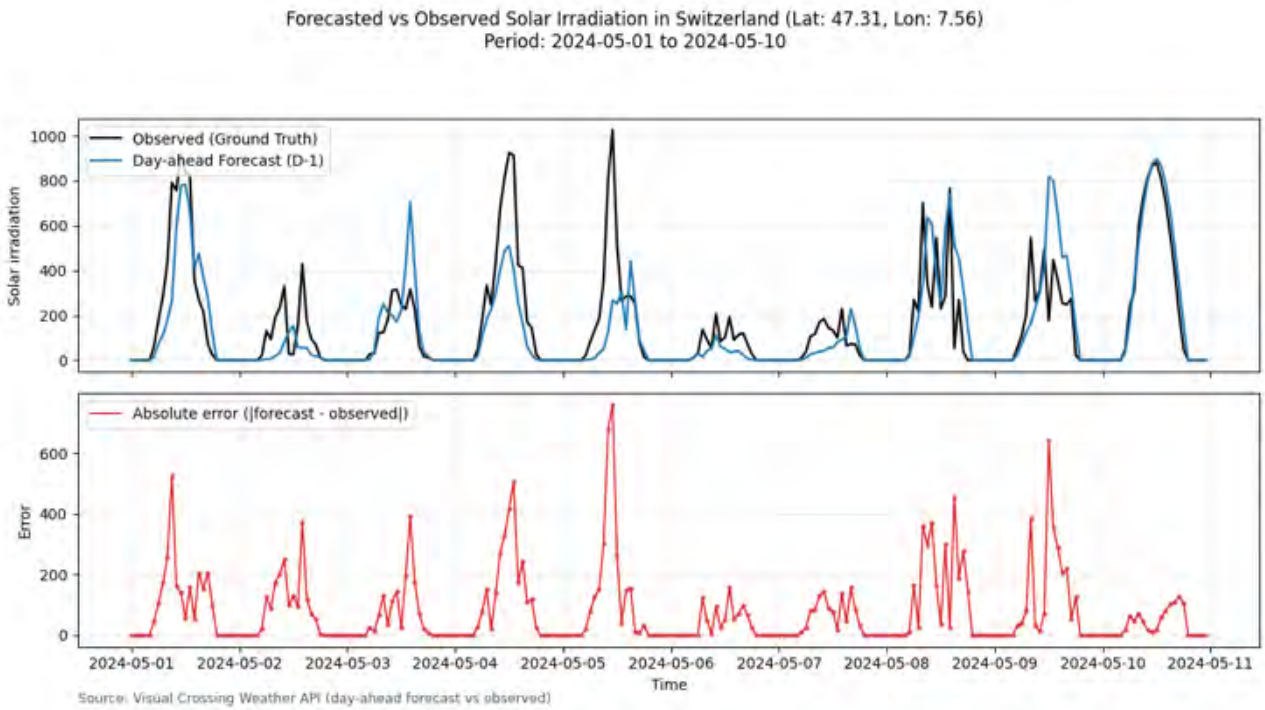

Figure 6. – Forecasted vs. observed solar irradiation in Switzerland (Lat: 47.31, Lon: 7.56) over the period May 1 to May 10, 2024. Data source: Visual Crossing Weather API. The top panel shows day-ahead forecasts (blue) deviating from observations (black), while the bottom panel tracks the absolute error, demonstrating recurrent inaccuracies.

changes are handled internally by the meta-learner, reducing the need for manual intervention or model redevelopment. This is particularly relevant for operational forecasting systems managing multiple assets or regions simultaneously.

Moreover, the results suggest that the two-layer approach is especially well suited to **moderate-to-strong drift regimes**, where structural changes occur but do not completely invalidate the historical relationship between inputs and outputs.

Feature uncertainty: Solar Power Production

In contrast to nuclear power, solar power production is dominated by uncertainty in exogenous inputs, particularly weather forecasts. Even when the underlying physical relationship between irradiance and production remains stable, errors in meteorological predictions introduce noise and bias that directly affect forecast accuracy.

This effect is illustrated in **Figure 6**, which compares forecasted and observed solar radiation. The figure reveals both random dispersion and systematic deviations, indicating that forecast errors are neither negligible nor purely stochastic.

Forecasting performance is summarized in **Table 2**. We compare the Baseline model against our Two-Layer Framework, as well as an "Upper Bound" (Oracle) model. The Upper Bound model was trained and tested using observed (perfect) weather data rather than forecasts; it represents the theoretical limit of performance achievable if meteorological predictions were error-free.

Table 2 – Performance comparison of the Baseline, Two-Layer, and Upper Bound models for Swiss solar production forecasting (Test set: 01.01.24 to 01.01.25). Metrics are computed on hourly aggregated data.

Model	R^2 [-]	RMSE [MW]	MAE [MW]
Baseline	0.7183	436.52	215.84
Two-Layer	0.7711	393.50	198.43
Upperbound	0.7740	391.34	185.34

Table 2 shows that the baseline model suffers from elevated RMSE and MAE due to error propagation from uncertain inputs. Introducing the meta-learning layer leads to a consistent reduction in both metrics, with RMSE improvements approaching 10%(reducing from

437 to 394 MW). While the two-layer framework does not fully match the upper-bound model, it recovers a substantial fraction of the lost performance.

This result is particularly significant, as it demonstrates that meaningful gains can be achieved **without access to perfect information** and without retraining the main forecasting model.

Beyond average error metrics, the two-layer framework also exhibits improved behaviour during extreme events. The analysis shows that large forecast errors are disproportionately reduced, indicating that the meta-learner is especially effective at correcting severe input deviations.

This observation is consistent with the ensemble nature of the Adaptive Random Forest, which can capture nonlinear relationships between forecasted and observed weather variables. As a result, the meta-learning layer acts as a dynamic error filter, attenuating the impact of extreme weather forecast inaccuracies.

Comparing the solar and nuclear case studies reveals an important distinction between **structural non-stationarity** and **stochastic uncertainty**. In the solar case, the underlying system remains stable, and the meta-learner operates primarily as a noise corrector. In the nuclear case, the system itself evolves, requiring the meta-learner to infer latent structural variables.

Despite these differences, the same two-layer architecture proves effective in both settings. This highlights the generality of the proposed framework and supports the claim that decoupling prediction and adaptation is a powerful design principle for energy forecasting under non-stationarity.

Overall Discussion and Lessons Learned

Across both case studies, several key insights emerge:

1. **Adaptation without retraining is feasible** for a wide range of non-stationary energy forecasting problems.
2. **Meta-features provide a flexible interface** between adaptive components and fixed forecasters.
3. The framework performs best under **partial but persistent non-stationarity**, while extreme regime shifts remain challenging.
4. The computational efficiency of the approach makes it attractive for large-scale operational deployment.

These findings position the proposed two-layer framework as a practical compromise between static models and fully online learning systems.

CONCLUSIONS

This paper proposes a two-layer meta-learning framework for energy forecasting under non-stationary conditions, addressing both structural concept drift and uncertainty in exogenous inputs. By decoupling prediction from adaptation, the framework enables lightweight online adjustment without retraining the main forecasting model.

Results from real-world case studies demonstrate that the proposed approach consistently improves forecasting robustness across heterogeneous settings. In nuclear power production forecasting, the meta-learning layer effectively tracked capacity-driven structural changes, substantially reducing error relative to static baselines and maintaining meaningful predictive performance under moderate and strong drift. In solar power forecasting, the framework mitigated error propagation from uncertain weather inputs, recovering a significant portion of the performance gap to an upper-bound model trained on observed meteorological data.

An important outcome of this work is the demonstration that adaptation can be localised within a meta-layer, avoiding the computational cost and operational complexity associated with frequent retraining or fully online learning. This makes the proposed framework particularly suitable for large-scale operational forecasting systems.

The framework's limitations become apparent under extreme regime shifts that fundamentally invalidate historical relationships, and the present work is restricted to point forecasting. Future research will therefore focus on extending the approach to probabilistic forecasting, improving robustness to extreme structural breaks, and evaluating its applicability to additional energy system use cases.

Overall, the results indicate that two-layer meta-learning provides a practical and effective strategy for improving energy forecasting performance in evolving power systems.

ACKNOWLEDGEMENTS

This work was supported by the European Union's Horizon 2020 research and innovation program under the Marie Sklodowska-Curie grant agreement No 101150337 – DMOptfES and under the project EGRID (Innovative Tracking, Monitoring and Reduction of Electricity Carbon Emissions) supported by COMPETE2030-FEDER-01073700.

AUTHOR IDENTIFIERS

Marcato, F: 0009-0003-5161-4635
Ruivo de Oliveira, M: 0009-0008-1142-3777
Santecchia, A.: 0000-0002-6089-0776
Francesco Silvestri: 0000-0002-9077-9921
Castro-Amoedo R.: 0000-0001-8287-8295

REFERENCES

1. Santecchia A, Kantor I, Castro-Amoedo R, Maréchal F. Industrial flexibility as demand side response for electrical grid stability. Front. Energy Res. 10: (2022). https://doi.org/10.3389/fenrg.2022.831462
2. Santecchia A, Castro-Amoedo R, Nguyen T, Kantor I, Stadler P, Maréchal F. The critical role of electricity storage for a clean and renewable european economy. Energy Environ. Sci. 16:5350-5370 (2023). https://doi.org/10.1039/d3ee02768f
3. Ali KH, Sigalo M, Das S, Anderlini E, Tahir AA, Abusara M. Reinforcement learning for energy-storage systems in grid-connected microgrids: an investigation of online vs. offline implementation. Energies 14:5688 (2021). https://doi.org/10.3390/en14185688
4. Tsymbal, A. (2004). The Problem of Concept Drift: Definitions and Related Work. Tech. Rep. Department of Computer Science, Trinity College, Dublin.
5. Widmer G, Kubat M. Learning in the presence of concept drift and hidden contexts. Machine Learning 23:69-101 (1996). https://doi.org/10.1023/a:1018046501280
6. Klinkenberg R. Learning drifting concepts: example selection vs. example weighting. IDA 8:281-300 (2004). https://doi.org/10.3233/ida-2004-8305
7. Ikonomovska E, Gama J, Džeroski S. Learning model trees from evolving data streams. Data Min Knowl Disc 23:128-168 (2010). https://doi.org/10.1007/s10618-010-0201-y
8. Gomes HM, Bifet A, Read J, Barddal JP, Enembreck F, Pfharinger B, Holmes G, Abdessalem T. Adaptive random forests for evolving data stream classification. Mach Learn 106:1469-1495 (2017). https://doi.org/10.1007/s10994-017-5642-8
9. Heitor M. Gomes et al. "Adaptive random forests for data stream regression". In: *ESANN 2018 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning* (2018), pp. 267–272.
10. Putatunda S. Supervised learning for streaming data. Practical Machine Learning for Streaming Data with Python :57-96 (2021). https://doi.org/10.1007/978-1-4842-6867-4_3
11. Pinson P. Wind energy: forecasting challenges for its operational management. Statist. Sci. 28: (2013). https://doi.org/10.1214/13-sts445
12. Hong T, Pinson P, Fan S, Zareipour H, Troccoli A, Hyndman RJ. Probabilistic energy forecasting: global energy forecasting competition 2014 and beyond. International Journal of Forecasting 32:896-913 (2016). https://doi.org/10.1016/j.ijforecast.2016.02.001
13. Zhang, Z., Zhang, H., Zhang, Y., & Wang, Y. "Noise-intensification data augmented machine learning for day-ahead wind power forecast". Energy Conversion and Management 266 (2022), 115803. https://doi.org/10.1016/j.enconman.2022.115803
14. Ren Y, Suganthan PN, Srikanth N. Ensemble methods for wind and solar power forecasting—a state-of-the-art review. Renewable and Sustainable Energy Reviews 50:82-91 (2015). https://doi.org/10.1016/j.rser.2015.04.081
15. Nishanthi M, Pandiyan P. Meta-learning-based imputation in solar energy systems: bridging the gap between missing data and forecasting reliability. J. Electr. Eng. Technol. 21:991-1010 (2025). https://doi.org/10.1007/s42835-025-02420-1
16. Mez L, Okamura L, Weidner H. The ecological modernization capacity of japan and germany. Springer Fachmedien Wiesbaden (2020). https://doi.org/10.1007/978-3-658-27405-4
17. ENTSO-E. "ENTSO-E Transparency Platform". URL: https://transparency.entsoe.eu/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Analysis and comparison of technologies for the regeneration of a capture solution in DAC absorption systems

Grazia Leonzio[a], Nilay Shah[b]

[a] University of Cagliari, Department of Mechanical, Chemical and Materials Engineering, Cagliari, Italy
[b] Sargent Centre for Process Systems Engineering, Department of Chemical Engineering, Imperial College London, London SW7 2AZ, UK
* Corresponding Author: grazia.leonzio@unica.it.

ABSTRACT

Direct air capture is gaining significant interest due to its potential to achieve carbon-negative emissions. With the aim to reduce the energy consumption and in line with the electrification of chemical processes, the absorption direct air capture system is integrated into a bipolar membrane electrodialysis cell stack for solvent regeneration and carbon dioxide release. The scheme solution is characterized by carbon dioxide bubbles inside the cell reducing efficiency so that other regeneration schemes have been proposed in the literature. A direct comparison of those is missing in the existing state of the art: the present work wants to fill this gap. In addition to the above process, the reaction of the rich solution with a weak organic acid and the use of both nanofiltration and reverse osmosis membranes are considered in other process schemes. The three case studies are modelled in Aspen Plus software with the aim to compare the energy consumption and total cost while the environmental impact is evaluated with the life cycle assessment methodology. Results show that the best capture process from economic and environmental point of views is that using only a bipolar membrane electrodialysis for solvent regeneration and carbon dioxide release despite a higher energy consumption for the cell stack. Here total capture cost and environmental impact are respectively 480 $/ton$CO_2$ and -0.9593 kgCO_2eq/kgCO_2 while the energy requirement in the BPMED stack is 408 kJ/molCO_2.

Keywords: Carbon Dioxide Capture, Modelling and Simulations, Life Cycle Analysis, Technoeconomic Analysis, Aspen Plus, OpenLCA

INTRODUCTION

Climate phenomena have caused profound social, economic and environmental damage, highlighting the urgency for systemic solutions. There is a strong scientific consensus that anthropogenic greenhouse gas emissions, particularly carbon dioxide (CO2), are the primary drivers of climate change. Hene, CO2 needs to be actively removed from the environment using carbon dioxide removal technologies (CDR). In this context, direct air capture (DAC) is a crucial and emerging carbon dioxide removal (CDR) technology for achieving net-zero emissions by balancing atmospheric CO2 release with removal. Different DAC technologies have been developed in the current state-of-the art such as absorption, adsorption, ion exchange resin, mineral carbonation, membrane, cryogenic separation [1]. These are mainly based on the use of thermal energy for regeneration of the capture agent [1, 2]. Recently, significant attention of the research community has been focused on electrochemical methods for regeneration: electrodialysis, electrolysis, proton coupled electron transfer, electro swing adsorption have been investigated [3, 4]. In the electrodialysis based process an absorption-based direct air capture column is connected to an electrodialysis stack for the release of CO2 and regeneration of the alkaline solvent recycled back to the absorption column [5]. Several research have been presented in the literature for this process scheme in order to evaluate the cost and the energy consumption as those reported by Sabatino et al.

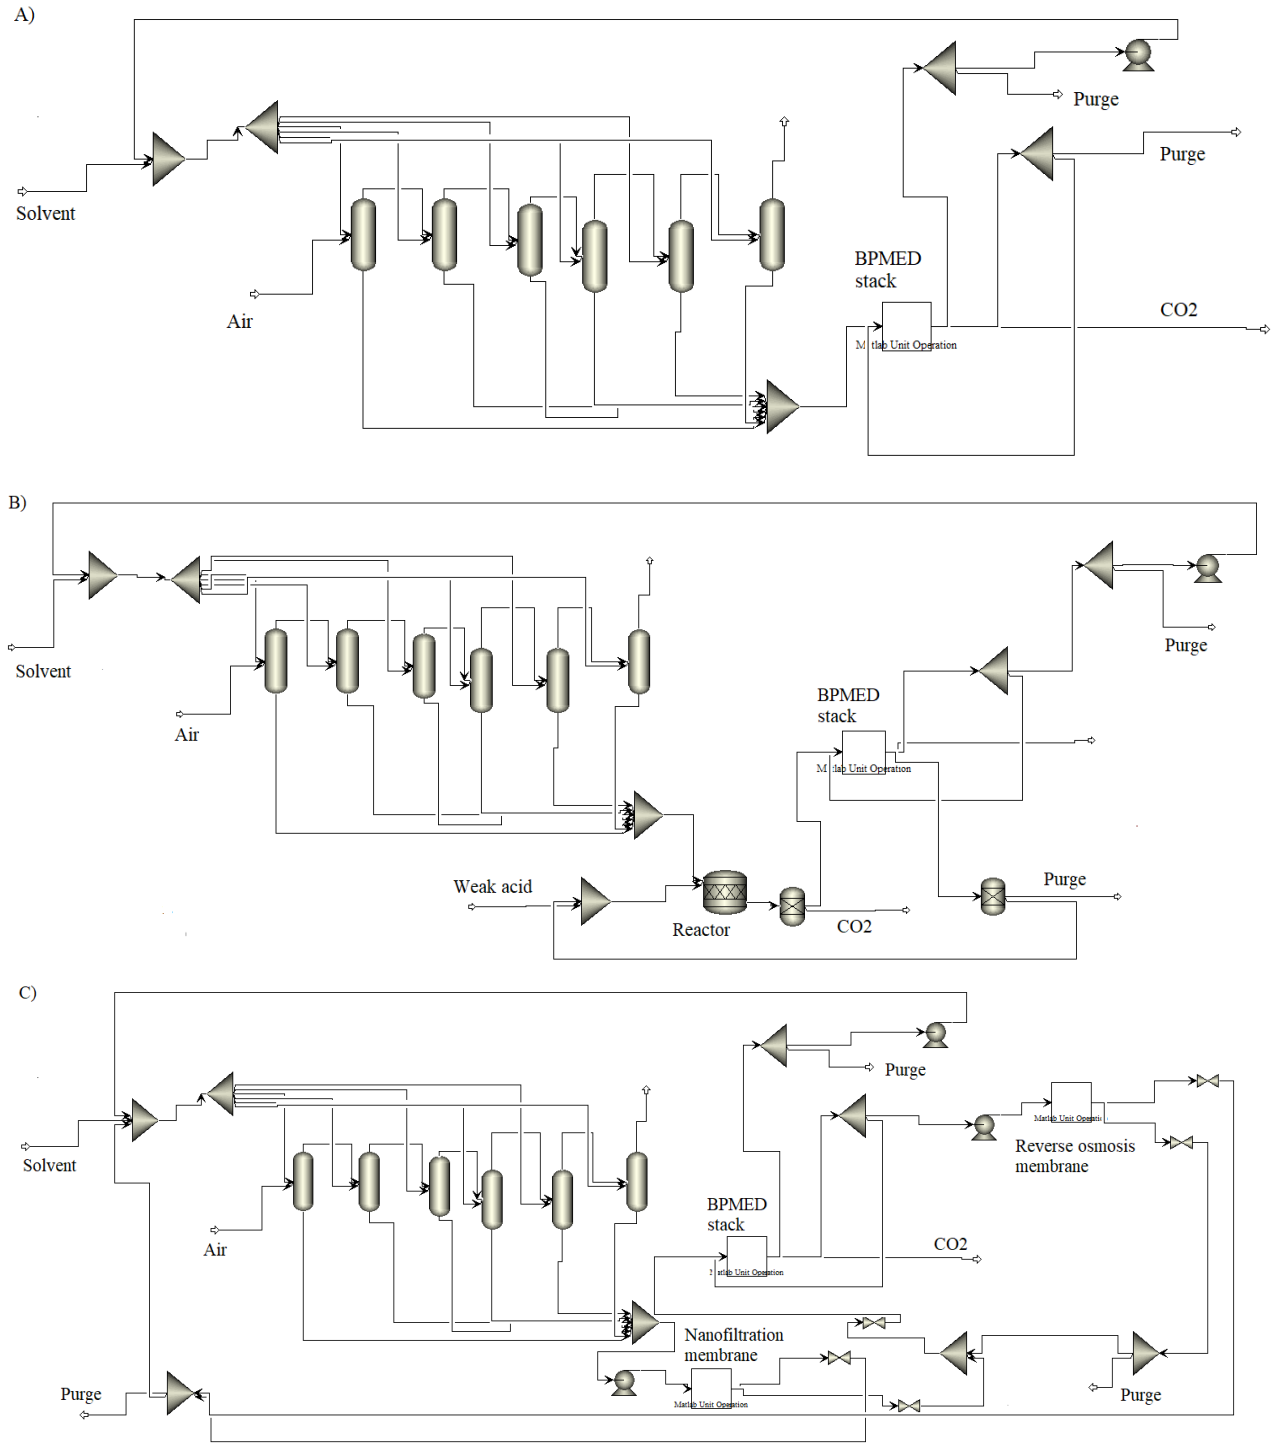

Figure 1: Process flowsheet of the investigated process: a) conventional; b) based on reaction between bicarbonates/carbonates and weak acid; c) based on BPMED with nanofiltration and reverse osmosis membrane

[6, 7], Young et al. [8], Leonzio and Shah [4], Leonzio et al. [9]. However, this process scheme can suffer from CO2 bubble formation inside the cell causing an increase of electrical energy demand. CO2 bubbles can, in fact, obstruct flow channels, block active membrane area, increase local resistance, exacerbate mass-transport limitations, current crowding and local heating [10]. To overcome this problem, in a previous work we proposed a system where carbonates and bicarbonates from the absorption column react with an organic weak acid to

release CO_2, producing its acid salt that is sent to the electrodialysis stack for both acid and alkaline solvent regeneration [11]. However, many other different alternatives have been proposed by Kemp et al. [12] to regenerate the capture solution and these are all based on the electrification of the process (direct and indirect protonation, nanofiltration and direct protonation, nanofiltration and indirect protonation, regeneration with CO2 electroreduction for product generation, CO2 electroreduction for product generation only, chiller crystallisation and direct protonation, nanofiltration and reverse osmosis, system with filtration unit and pH adjustment of electrodialysis feed stream, system with filtration unit and electrodialysis, schematic of electrodialysis unit with gas diffusion electrode). For the best our knowledge, a direct comparison of these solutions is not present in the literature yet. Therefore, the present research wants to overcome this gap. In particular, in this research, we aim to compare the following direct air capture technologies based on the use of an absorption column but with the regeneration of the rich solution by using a bipolar membrane electrodialysis stack, reacting carbonates and bicarbonates with a weak organic acid (formic acid or a mixture between formic and acetic acid) and by using nanofiltration and reverse osmosis technology. The proposed process schemes are simulated in detail and the material and energy balances are used to conduct an economic and environmental analysis finding the total capture cost and environmental burden.

MODELLING

Description of the capture plant

Three different process schemes are compared in this work: "conventional" scenario with the DAC absorption column connected to the BPMED stack, "reaction between bicarbonates/carbonates and weak acid" scenario where the BPMED stack is used for the regeneration of the acid salt obtained with the reaction between bicarbonates/carbonates and formic and/or acetic acid, "nanofiltration and reverse osmosis" scenario where these membranes are used to concentrate bicarbonates/carbonates sent to the BPMED connected to the DAC absorption column. These schemes are shown respectively in Figure 1a, 1b, 1c. In all process schemes, six absorption columns, in a cross-flow configuration, are considered according to the Carbon Engineering plant [6, 7]. An ambient air flow rate of 150 ton/h and a liquid flow rate according to a liquid to gas ratio equal to 0.6 are considered in the feed [6]. The liquid flow rate is based on a mixture of an aqueous solution of potassium hydroxide (KOH) (1 M) and potassium carbonate (K_2CO_3) (0.01 M), the latter added with the aim to reduce toxicity and corrosivity [13, 14]. From the absorption columns, the rich solution with carbonates and bicarbonates is sent to the BPMED stack for solvent regeneration and CO_2 release in the "conventional" scenario. In the "reaction between bicarbonates/carbonates and weak acid" scenario, the rich solution from the absorption column is sent to a reactor for the reaction with formic acid, in one case, and formic/acetic acid mixture, in the other case. Here, CO2 is released while the acid salt is sent to the BPMED stack for the regeneration of both acid and absorption solvent. In the "nanofiltration and reverse osmosis" scenario, the rich solution from the absorption column, is sent to the nanofiltration unit to concentrate carbonates sent to the BPMED stack, while the permeate solution rich in hydroxide ions is recycled back to the absorption column. In the same scheme, the brine stream from the BPMED is sent to the reverse osmosis unit to concentrate bicarbonates recycled back to the BPMED while the permeate rich in hydroxide ions is sent to the absorption column. All processes are modelled in Aspen Plus software by using RadFrac packed absorption columns, stoichiometric reactor and Matlab Cape Unit operation for the BPMED stack, nanofiltration and reverse osmosis membranes. For the BPMED cell stack, material balance equations, equilibrium equations and electroneutrality equations are used to solve the system and find the product stream composition. The following parameters are defined: membrane area equal to 1.785 m2 and current density equal to 800 A/cm2 [6], number of cells of 105. For membranes, by fixing the rejection factor for each compound and using material balances allow to find the product stream composition and needed membrane area. For the nanofiltration unit water permeability and transmembrane pressure are fixed to $9 \cdot 10^{-11}$ m/(Pas) and $5.4 \cdot 10^{6}$ Pa [15, 16]. For the reverse osmosis unit water permeability and transmembrane pressure are fixed to $1 \cdot 10^{-10}$ m/(Pas) and 1.1•10 Pa [15, 16]. The electrolyte Non-Random Two-Liquid (e-NRTL) thermodynamic model is used for the simulation in Aspen Plus of the overall capture plant [6, 7].

Energetic and economic analysis

Regarding the energetic analysis, the needed electric power for the BPMED is evaluated with the following correlations (See Eq. 1-4) [6, 7]:

$$E_{cell} = i \bullet R_{cell} + E_{BP} \quad (1)$$

$$R_{cell} = R_{base} + R_{CEM} + R_{acid} + R_{BP} \quad (2)$$

$$R_{base/acid} = \frac{\delta}{\sigma_{base/acid}} \quad (3)$$

$$P_{cell} = i \bullet A \bullet N_{cell} \bullet E_{cell} \quad (4)$$

with E_{cell} the voltage drop across the BPMED cell, E_{BP} the water-splitting potential of the BPM (equal to 0.829 V for ΔpH of 14 [6]) (V), R_{cell} the electrical resistance of a unit area of a single cell (Ωm^2), i the current density (A/m^2), R_{base}, R_{acid}, R_{CEM} and R_{BP} respectively the resistance

of basic and acid solution and cation and bipolar membrane (these two last are respectively equal to $1.5 \bullet 10^{-4}$ Ωm^2 and $5.2 \bullet 10^{-4}$ Ωm^2 [6]) (Ωm^2), $\sigma_{base/acid}$ the electrical conductivity of acid and base solutions, δ the distance between the BPM and cation membrane (1.5 mm [6]) (m), A the membrane area (m^2), N_{cell} the number of cells, P_{cell} the power of the BPMED unit (W). In the economic analysis the CO_2 capture cost is evaluated from the sum of capital (CAPEX) and operating (OPEX) costs. The CAPEX in the form of total overnight capital (TOC) ($) is evaluated on the basis of the bare erected cost (BEC), as the sum of the purchase cost of all plant components, taken from the literature and from Aspen Plus Economic Analyzer (APEA): absorption columns from Sabatino et al. [6, 7], reactor and pump from APEA, BPMED stack from Virruso et al. [17]. For the total cost of membranes in the stack, an ion exchange membrane cost and bipolar membrane cost respectively are assumed at 145 $/$m^2$ and 1300 $/$m^2$ [18]. For nanofiltration and reverse osmosis, a value of 134 $/m2 and 170 $/m2 is respectively assumed [19, 20]. In particular, the total overnight capital (TOC) is evaluated in accordance with NETL guidelines [21]. The OPEX is evaluated as the sum of fixed and variable operating and maintenance cost (O&M) ($/year). In the variable O&M cost, the cost of raw materials and utilities is taken into account assuming a price for aqueous hydroxide solvent of 580 $/ton [22], potassium carbonate of 600 $/ton, sustainable formic acid from CO_2 of 1738 $/ton [23], sustainable acetic acid from CO_2 of 1170 $/ton [24] and renewable electricity from wind onshore in the UK of 0.05 $/kWh [25]. The specific CO_2 capture cost is the sum TOC and OPEC divided the amount of captured CO_2.

Environmental analysis

The LCA is conducted with OpenLCA software (version 1.11) and Ecoinvent database 3.9 accordingly with ISO 14040-14044 [26, 27]. This methodology includes four phases (goal and scope definition, life cycle inventory, life cycle impact assessment (LCIA) and life cycle interpretation). The goal of the study is the comparison of the three investigated processes from an environmental point of view with the scope to understand which of them is the most favorable. To this aim the functional unit is 1 tonCO2 that is captured with cradle-to-gate system boundary (the end use of captured CO_2 is not considered in the analysis). In the inventory phase, inventory data (e.g. material and energy balances) for the processes are defined as taken from the simulation in Aspen Plus. For formic acid and acetic acid, a sustainable production process from CO_2 is considered as that reported in Cuevas-Castillo et al. [23] and Regis et al. [24]. After these stages, the LCI results are organized into the usual LCA impact categories and then into the impact indicators at the midpoint level. The life cycle impact assessment is performed with the ReCiPe Midpoint (H) 2016 to obtain a general overview of the impact of the studied process on the climate change impact category [28]. In the interpretation phase, a comparison of the obtained results is carried out.

RESULTS AND DISCUSSION

Results of process modelling

In all simulations, the kinetic model of the absorption column is validated with the experimental data of Keith et al. [13] while the BPMED is validated with Sabatino et al. [6, 7] data.

From the developed simulations, the energy consumption (especially that in the BPMED stack) and the material balances are obtained and compared in each investigated process scheme.

For the same amount of air and solvent flow rate fed into the process, the energy consumption in the BPMED and the amount of captured CO2 are respectively of 2314 kWh/tonCO2 and 1142 tonCO2/year for the conventional process, 1907 kWh/tonCO2 and 959 tonCO2/year for the process using formic acid reacting with the rich solution, 1943 kWh/tonCO2 and 949 tonCO2/year for the process using the mixture formic/acetic acid reacting with the rich solution, 2468 kWh/tonCO2 and 1253 tonCO2/year for the process using a nanofiltration with reverse osmosis membrane.

The higher efficiency of the system regenerating CO2 from the electrodialysis of carbonates and bicarbonates is due the fact that in this system the total dissolved organic carbon is conserved while in the acid reaction-based process CO2 is produced not after a redistribution of carbonates because it is a chemical transformation.

Moreover, it is clear that the process with the highest electricity consumption in the BPMED stack is that using membranes due to the higher concentrate rich solution from which a higher amount of CO2 is captured compared to other case studies. Considering the energy consumption of pumps in both membranes, the overall energy needs is higher and with a value of 3752 kWh/tonCO2. Then, despite the higher amount of CO2 that is captured, the process based on the use of membranes has a higher energy requirement.

It is interesting to see that both processes based on CO2 recovery from rich solution through the BPMED stack are characterized by a greater value of energy consumption in the cell stack due to the CO2 bubble formation. A higher value is for the process with membranes due to a higher concentrate solution (435 kJ/molCO2 vs 408 kJ/molCO_2 for the process using only the electrolytic stack).

The process with the lowest energy consumption is that based on the reaction between formic acid and

bicarbonates/carbonates: the BPMED stack requires 336 kJ/molCO2. The lower energy consumption for the system using formic acid is due to a very low specific energy required to regenerate this acid from its salt: the range is between 0.3 and 0.8 kWh/kg of formic acid [29]. On the other hand, the energy consumption for recovery of acetic acid from its salt can be up to 15.7 kWh/kg of acetic acid [30]. Then the presence of acetic acid in the mixture increases the energy demand for CO2 capture.

Overall, for the process based on the regeneration of CO2 via BPMED, values of these electrical consumptions are in agreement with that reported in the literature by Sabatino et al. [6, 7] and Eisaman et al. [30] where energy consumption up to 3789 kWh/tonCO2 are reported when varying current density and composition of the rich solution. Similar values are reported by Vallejo Castano et al. [32], Shu et al. [33] analyzing the same process under different sensitivity analyses. For the other two processes, a direct comparison with the literature cannot be conducted because a similar process has not been considered for CO2 capture from the air yet.

Results of economic analysis

Results of the economic analysis show that the process based on only BPMED stack costs 480 $/tonCO2, that based on membranes costs 1179 $/tonCO2 while the others based on formic acid reaction and formic/acetic acid reaction cost respectively 510 $/tonCO2 and 519 $/tonCO2. It is clear that the process using membranes is the most expensive (59% more than that based on the regeneration using only BPMED that is the cheapest one), despite the highest amount of captured CO2 (1253 tonCO2/year). This is due to the highest impact of nanofiltration cost on total capital cost (due to a higher treated flow rate, 292 modules with 37 m2 of area are needed for the specific separation while 4 modules of reverse osmosis membrane with an area of 40 m2 are required).

A deeper analysis shows that the system using membranes has the highest values of TOC, total fixed and variable O&M respectively of 0.55 M$/year (437 $/tonCO2), 0.65 M$/year (520 $/tonCO2) and 0.28 M$/year (222 $/tonCO2). For the process scheme using a BPMED stack for the regeneration TOC, total fixed and variable O&M costs are 0.19 M$/year (165 $/tonCO2), 0.18 M$/year (161 $/tonCO2) and 0.17 M$/year (153 $/tonCO2). TOC, total fixed and variable O&M costs are respectively 0.22 M$/year (229 $/tonCO2), 0.1 M$/year (111 $/tonCO2) and 0.16 M$/year (170 $/tonCO2) for the process using formic acid reacting with the rich solution. On the other hand, TOC, total fixed and variable O&M costs for the process using a mixture acid solution are respectively 0.22 M$/year (231 $/tonCO2), 0.1 M$/year (112 $/tonCO2) and 0.17 M$/year (176 $/tonCO2).

It is evident that despite a higher operating costs (due to the higher energy consumption) compared to the process based on acid reaction, the process with the use of only BPEM stack for CO2 regeneration has the lowest capital cost due to the lower amount of used equipment. On the other hand, the process with nanofiltration and reverse osmosis has the highest values of capital and operating costs compared to all solutions.

The cost for the process based on only BPMED is in line (e.g. the same order of magnitude) with studies reported in the literature from Sabatino et al. [6, 7], Iizuka et al. [34]. For the other processes a comparison with the literature can not be done as similar works have not been considered.

Results of environmental analysis

It is possible to see that values of climate change for the process with the use of only BPMED, using formic acid and formic/acetic acid reaction and based on membranes are respectively equal to -0.9593 kgCO2eq/kgCO2, -0.9378 kgCO2eq/kgCO2 and -0.9238 kgCO2eq/kgCO2, -0.9327 kgCO2eq/kgCO2. Then, recovering CO2 from carbonates via BPMED ensures the lowest environmental burden in terms of climate change, followed by the process whit a reaction between bicarbonates/carbonates and formic acid, by the system using membranes and then followed by the process with a reaction between bicarbonates/carbonates with the acetic/formic acid mixture.

A breakdown of results shows that the most significant input on total climate change is the electricity consumption for the conventional and membrane based processes (28.53 kgCO2eq/tonCO2 and 46.25 kgCO2eq/tonCo2 respectively), while for the other processes based on chemical reaction to release CO2, the potassium carbonate production has the highest impact (35.14 kgCO2eq/tonCO2 for the process using only formic acid and 35.52 kgCO2eq/tonCO2 for the process using the mixture acetic/formic acid).

Results for the system based on electrodialysis of carbonates are in agreement with the literature [4]. For the other systems, a direct comparison with the state-of-the-art is not possible.

CONCLUSIONS

With the aim of electrifying chemical processes for CO2 capture from the air, different schemes are proposed and investigated in this work. The DAC absorption column is related to a BPMED stack in one case, to a BPMED stack and nanofiltration and reverse osmosis membranes in another case study and to a reactor for the reaction between the rich solution and formic acid and formic/acetic acid mixture in the other two case studies.

All processes are simulated in Aspen Plus to have material and energy balances required to conduct an economic and environmental analysis, finding the total

capture cost and the climate change impact.

A similar comparison analysis has not been considered in literature yet, underlining the novelty of the process.

Results show that the best process for the regeneration of solvent and release of CO2 is that based on the use of BPMED stack characterized by 480 $/tonCO2 and -0.9593 kgCO2eq/kgCO2. On the other hand, due to the higher impact of membrane cost, the process with the nanofiltration and reverse osmosis has the highest costs (1179 $/tonCO2) and relatively higher environmental burden (-0.9327 kgCO2eq/kgCO2).

ACKNOWLEDGEMENTS

The authors would like to thank the University of Cagliari and Imperial College London.

REFERENCES

1. Leonzio G, Fennell PS, Shah N. Analysis of technologies for carbon dioxide capture from the air. Applied Sciences 12:8321 (2022). https://doi.org/10.3390/app12168321
2. Sanz-Pérez ES, Murdock CR, Didas SA, Jones CW. Direct capture of co₂ from ambient air. Chem. Rev. 116:11840-11876 (2016). https://doi.org/10.1021/acs.chemrev.6b00173
3. Bouaboula H, Chaouki J, Belmabkhout Y, Zaabout A. Comparative review of direct air capture technologies: from technical, commercial, economic, and environmental aspects. Chemical Engineering Journal 484:149411 (2024). https://doi.org/10.1016/j.cej.2024.149411
4. Leonzio G, Shah N. Sustainability analysis of electrochemical direct air capture technologies. RSC Sustainability 3:4632-4650 (2025). https://doi.org/10.1039/d5su00227c
5. Sabatino F, Mehta M, Grimm A, Gazzani M, Gallucci F, Kramer GJ, van Sint Annaland M. Evaluation of a direct air capture process combining wet scrubbing and bipolar membrane electrodialysis. Ind. Eng. Chem. Res. 59:7007-7020 (2020). https://doi.org/10.1021/acs.iecr.9b05641
6. Sabatino F, Mehta M, Grimm A, Gazzani M, Gallucci F, Kramer GJ, van Sint Annaland M. Evaluation of a direct air capture process combining wet scrubbing and bipolar membrane electrodialysis. Ind. Eng. Chem. Res. 59:7007-7020 (2020). https://doi.org/10.1021/acs.iecr.9b05641
7. Sabatino F, Gazzani M, Gallucci F, van Sint Annaland M. Modeling, optimization, and techno-economic analysis of bipolar membrane electrodialysis for direct air capture processes. Ind. Eng. Chem. Res. 61:12668-12679 (2022). https://doi.org/10.1021/acs.iecr.2c00889
8. Young J, McQueen N, Charalambous C, Foteinis S, Hawrot O, Ojeda M, Pilorgé H, Andresen J, Psarras P, Renforth P, Garcia S, van der Spek M. The cost of direct air capture and storage can be reduced via strategic deployment but is unlikely to fall below stated cost targets. One Earth 6:899-917 (2023). https://doi.org/10.1016/j.oneear.2023.06.004
9. Leonzio G, Thill A, Shah N. A global sensitivity analysis for a bipolar membrane electrodialysis capturing carbon dioxide from the air. Systems and Control Transactions 4:456-461 (2025). https://doi.org/10.69997/sct.160159
10. Nagasawa H, Yamasaki A, Iizuka A, Kumagai K, Yanagisawa Y. A new recovery process of carbon dioxide from alkaline carbonate solution via electrodialysis. AIChE Journal 55:3286-3293 (2009). https://doi.org/10.1002/aic.11907
11. Leonzio, G., Xing, L., Shah, N., Analysis of new strategies integrating bipolar membrane electrodialysis and absorption systems in DAC applications. Carbon Capture Science & Technology, under review.
12. Kemp, K.W., Ostericher, A.L., Olmstead, D.E., 2022. Systems and methods for capturing carbon dioxide and regenerating a capture solution, US 2022/0362707 A1
13. Keith DW, Holmes G, St Angelo D, Heidel K. A process for capturing CO2 from the atmosphere. Joule 2:1573-1594 (2018). https://doi.org/10.1016/j.joule.2018.05.006
14. Rouxhet A, Kim SM, Leonard G. Study of direct air capture (DAC) using a KOH/K2CO3 absorbing solution for CO2 capture. SSRN Journal : (2022). https://doi.org/10.2139/ssrn.4273837
15. Maeda Y. Fouling of reverse osmosis (RO) and nanofiltration (NF) membranes by low molecular weight organic compounds (lmwocs), part 1: fundamentals and mechanism. Membranes 14:221 (2024). https://doi.org/10.3390/membranes14100221
16. Ma Q, Lei Q, Liu F, Song Z, Khusid B, Zhang W. Evaluation of commercial nanofiltration and reverse osmosis membrane filtration to remove per?and polyfluoroalkyl substances (PFAS): effects of transmembrane pressures and water matrices. Water Environment Research 96: (2024). https://doi.org/10.1002/wer.10983
17. Virruso G, Cassaro C, Culcasi A, Cipollina A, Tamburini A, Bogle IDL, Micale GDM. Multi-scale modelling of an electrodialysis with bipolar membranes pilot plant and economic evaluation of its potential. Desalination 583:117724 (2024). https://doi.org/10.1016/j.desal.2024.117724

18. Ankoliya D, Mudgal A, Sinha MK, Davies P, Park K, Alegre RR, Patel V, Patel J. Techno-economic analysis of integrated bipolar membrane electrodialysis and batch reverse osmosis for water and chemical recovery from dairy wastewater. Journal of Cleaner Production 420:138264 (2023). https://doi.org/10.1016/j.jclepro.2023.138264
19. https://www.bigbrandwater.com/nf90-2540.html?utm_source=copilot.com
20. https://www.acquahome.it/en/03-vessel-e-membrane/membrane-filmtec/filmtec-bw30-2540-reverse-osmosis-membrane.2.10.35.gp.1587.uw?utm_source=copilot.com
21. Spallina V, Pandolfo D, Battistella A, Romano MC, Van Sint Annaland M, Gallucci F. Techno-economic assessment of membrane assisted fluidized bed reactors for pure h 2 production with CO 2 capture. Energy Conversion and Management 120:257-273 (2016). https://doi.org/10.1016/j.enconman.2016.04.073
22. Intratec, 2025 available at: https://www.intratec.us/solutions/primary-commodity-prices/commodity/potassium-hydroxide-prices?utm_source=chatgpt.com
23. Cuevas-Castillo GA, Michailos S, Hughes K, Ingham D, Navarro-Pineda F, Pourkashanian M. Techno-economic and life cycle assessment of power-to-formic acid production using direct air capture and green hydrogen. Cleaner Engineering and Technology 26:100929 (2025). https://doi.org/10.1016/j.clet.2025.100929
24. Regis F, Demichelis F, Monteverde A, Fino D. Techno-economic-environmental analysis of CO2 fermentation to acetic acid. Chemical Engineering Journal 520:166138 (2025). https://doi.org/10.1016/j.cej.2025.166138
25. OurWorldInData, 2025, available at: https://ourworldindata.org/grapher/levelized-cost-of-energy?country=~GBR
26. ISO 14040, 2009. Environmental Management – Life Cycle Assessment – Principles and Framework. European Committee for Standardisation, Brussels.
27. ISO 14044, 2006. Environmental Management – Life Cycle Assessment – Requirements and Guidelines. European Committee for Standardisation, Brussels
28. Eke V, Sahu T, Ghuman KK, Freire-Gormaly M, O'Brien PG. A comprehensive review of life cycle assessments of direct air capture and carbon dioxide storage. Sustainable Production and Consumption 55:217-241 (2025). https://doi.org/10.1016/j.spc.2025.02.017
29. Handojo L, Wardani AK, Regina D, Bella C, Kresnowati MTAP, Wenten IG. Electro-membrane processes for organic acid recovery. RSC Adv. 9:7854-7869 (2019). https://doi.org/10.1039/c8ra09227c
30. Wang Y, Zhang N, Huang C, Xu T. Production of monoprotic, diprotic, and triprotic organic acids by using electrodialysis with bipolar membranes: effect of cell configurations. Journal of Membrane Science 385-386:226-233 (2011). https://doi.org/10.1016/j.memsci.2011.09.044
31. Eisaman MD, Alvarado L, Larner D, Wang P, Littau KA. CO2 desorption using high-pressure bipolar membrane electrodialysis. Energy Environ. Sci. 4:4031 (2011). https://doi.org/10.1039/c1ee01336j
32. Vallejo-Castaño S, Bianchi G, Shu Q, Mathiasin E, Saakes M, Hamelers HVM, Kuntke P. Efficiency and energy consumption analysis of bipolar membrane electrodialysis for electrochemical CO ₂ capture. J. Mater. Chem. A 13:37544-37557 (2025). https://doi.org/10.1039/d5ta04909a
33. Shu Q, Sin CS, Tedesco M, Hamelers HVM, Kuntke P. Optimization of an electrochemical direct air capture process with decreased CO2 desorption pressure and addition of background electrolyte. Chemical Engineering Journal 470:144251 (2023). https://doi.org/10.1016/j.cej.2023.144251
34. Iizuka A, Hashimoto K, Nagasawa H, Kumagai K, Yanagisawa Y, Yamasaki A. Carbon dioxide recovery from carbonate solutions using bipolar membrane electrodialysis. Separation and Purification Technology 101:49-59 (2012). https://doi.org/10.1016/j.seppur.2012.09.016

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

From Plastic Waste to Platform Chemicals: Aspen Plus Modeling of Polystyrene Conversion Through Hydrothermal Processing into Value-added Chemicals

MohammadSina HajiHashemi[a], Corinna Schulze-Netzer[a], and Thomas A. Adams II*[a]

[a] Norwegian University of Science and Technology, Department of Energy and Process Engineering, Trondheim, Norway
* Corresponding Author: thomas.a.adams@ntnu.no.

ABSTRACT

Polystyrene (PS) recycling remains limited despite large waste volumes, largely because many routes struggle to recover high-value chemicals at scale. This work develops a full process concept that upgrades PS through low-pressure hydrothermal processing (HTP) and directs the resulting aromatic oil to separation and downstream conversion. The superstructure includes HTP at 30 bar and 350°C, three distillation columns for toluene/ethylbenzene/styrene splits, ethylbenzene dehydrogenation to boost styrene yield, steam reforming of the heavier C_{9+} fraction to syngas, and an ICI-type (Imperial Chemical Industries) methanol loop with inter-bed quench. The integrated flowsheet was simulated in Aspen Plus V14 using a consistent NRTL-RK (Non-Random Two-Liquid-Redlich–Kwong) property framework. Deep-vacuum operation (≈100–400 mbar) was applied where needed to limit styrene polymerization. Sensitivity and optimization focused on meeting a syngas stoichiometric number near 2; the best case occurred at ~998.8°C and 20.2 bar with ~37.8 t/h steam. Normalized to the polystyrene feed, the model predicts ~0.13 kg toluene, 0.16 kg styrene, and 1.27 kg methanol per kg PS, with only trace benzene/ethanol, at a water-to-PS ratio of ~3.77 kg/kg. When heat recovery is excluded, the overall heat demand corresponds to ~84 MJ per kg of PS processed. The flowsheet is technically consistent, but heat integration and a full TEA/LCA (Techno-economic Assessment/Life-cycle Assessment) will ultimately determine its viability.

Keywords: Plastic Waste, Hydrothermal Processing, Superstructure, Simulation, Recycling

INTRODUCTION

The amount of generated plastic waste has surpassed 400 million tonnes annually and is growing at an alarming pace. Only 9% of this generated waste is recycled, and 12% is burned, while the rest is landfilled [1]. In this vein, polystyrene (PS) constitutes 7 wt% of worldwide plastic trash, equating to about 28 million tonnes generated each year [1]. However, around 1% of the PS-generated waste is recycled due to constraints in conversion techniques [2]. PS is generally used in the food packaging industry. Among polystyrene waste streams, expanded polystyrene (EPS; 'Styrofoam') is especially problematic because its intrinsic material properties promote environmental release. Its very low density and brittle morphology facilitate wind-driven dispersal from collection points and landfills, while minor mechanical stress readily causes fracture and fragmentation. As a result, EPS can rapidly generate microplastics that persist and spread over wide areas, making improved capture and controlled recycling or upgrading routes particularly important for mitigating this pathway [3].

Several studies in the literature showed that hydrothermal processing (HTP) is a promising way to convert PS into a wide range of chemicals in the literature. Gentilcore *et al.* have developed a low-pressure hydrothermal processing method to convert PS into oils, achieving an impressive yield of 96–99% oil with only 1–2% char at pressures of 2–3 MPa based on their experiments. They suggest and implement a reversible kinetic model involving monomers and poly-aromatics, confirm predictions, co-process PS with polyolefins, and assess greater

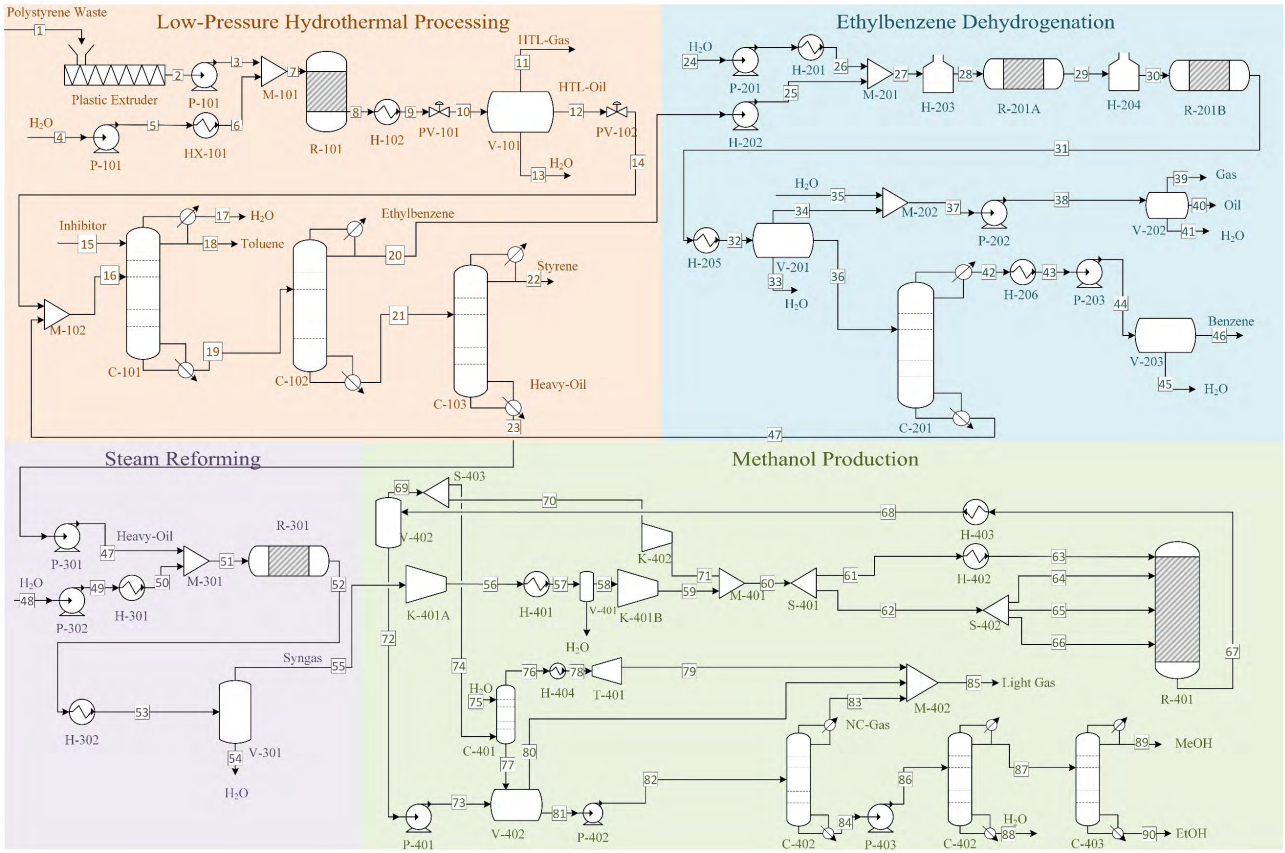

Figure 1. Superstructure design for polystyrene waste into value-added chemicals through low-pressure hydrothermal processing, ethylbenzene dehydrogenation, steam reforming, and the methanol production unit

energy recovery along with reduced environmental impacts compared to pyrolysis/incineration and landfilling [4]. The composition of the generated oil has become the basis for our simulation. It mostly contains monomers (C_7-C_9), dimers (C_{10}-C_{17}), and heavy aromatics (C_{16}-C_{20+}). In the present work, we propose a novel conceptual design of a process that converts the oil into styrene and convenient co-products like toluene and methanol, shown in Figure 1. Toluene and styrene (monomer) are valuable products of the reactor and can be separated using distillation columns.

Ethylbenzene, as the other fraction of the generated oil (15–20 wt%), can be transformed into styrene through the process of dehydrogenation, resulting in an increased yield of styrene, the most valuable product in our process. The remaining components of the composition can be converted into syngas via the steam reforming process. The produced syngas serves as the precursor for methanol production in our design concept, noting that methanol can further be converted to additional plastic precursors through well-known methanol-to-olefins processes.

Compared with mechanical recycling, solvent-based recovery, pyrolysis, and gasification, this study adopts low-pressure hydrothermal processing because it produces a high oil yield with low char formation and generates an aromatic-rich stream suitable for downstream upgrading [4]. In the present design, toluene, styrene, and methanol were selected as the main target products due to their high economic value and substantial production volumes. In addition, the heavy oil fraction (C_{9+}) accounted for about 65% of the produced oil, making it the dominant intermediate stream. Direct purification of this fraction is less attractive, so it was routed to steam reforming for syngas generation and then to methanol synthesis.

This is the first such study to present a complete design concept, addressing some particular challenges such as determining designs for difficult separation processes, and the relationships of important design parameters and operating conditions on key performance indicators like yield, efficiency, energy use, and environmental impact.

We use Aspen Plus to simulate the process and compute the mass and energy flows for the complete process. Optimization techniques are used to help identify design parameters that achieve maximum yield or other important objectives. This demonstrates a pathway

for the recycling of PS waste, which could contribute to a more sustainable industry environment.

METHODS

Process Description

Figure 1 presents the superstructure process flow diagram for the conversion of polystyrene (PS) waste into value-added chemicals across four processing units. As shown, PS waste is transformed into toluene, ethylbenzene, styrene monomer, benzene, syngas, methanol, and ethanol via low-pressure hydrothermal processing, ethylbenzene dehydrogenation, steam reforming, and a methanol synthesis section. The integrated process was simulated in Aspen Plus V14.

Low-pressure Hydrothermal Processing

In the introduced process, PS waste (stream 1) and water (stream 4) are mixed and enter the low-pressure (at 30 bar and 350°C) hydrothermal reactor to produce an aromatic-rich oil that serves as the primary feed to downstream separation and conversion units. The reactor products contain mostly aromatic hydrocarbons from C_7 to C_{20+}, which are going to be separated in three distillation columns from C-101 to C-103.

In modern styrene plants (ethylbenzene dehydrogenation process), the crude product is separated by vacuum distillation to avoid excessive temperatures that could polymerize styrene [5]. Then, all of the columns in the first and second units operate under deep vacuum (on the order of 100–400 mbar top pressure) to keep bottom temperatures low and minimize styrene polymerization. Even in a vacuum condition, styrene is prone to free-radical polymerization, especially at elevated temperatures. Polymerization is exothermic and self-accelerating; at ambient temperature, styrene reacts slowly, but rates increase rapidly with heat. In fact, the rate of styrene polymerization (uninhibited) roughly doubles for every 10°C rise [6]. In practice, polymerization inhibitors are always used when handling styrene (and other vinyl aromatic fractions). In this process, 4-hydroxy-TEMPO (2, 2, 6, 6-tetramethylpiperidine 1-oxyl) is used as an inhibitor to suppress undesired free-radical polymerization of styrene. This inhibitor showed a great thermal stability based on thermogravimetric analysis [7].

Toluene (stream 18) and styrene (stream 22) are considered the main products in the separation unit. Ethylbenzene (stream 20) is used as the feedstock for the downstream process of ethylbenzene dehydrogenation to produce more styrene. This configuration leverages hydrothermal depolymerization to create a flexible liquid platform, while assigning downstream pathways according to product value and integration opportunities.

Ethylbenzene Dehydrogenation

Ethylbenzene is purified (stream 20) as an aromatic fraction from upgraded HTP oil and can be directed to a traditional styrene production unit. The feed is vaporized and diluted with process steam, preheated using hot reactor effluent, and subsequently raised to the desired inlet temperature in a fired heater. Dehydrogenation to styrene and hydrogen is conducted in a series of low-pressure, adiabatic vapor-phase reactors. The endothermic and equilibrium-limited nature of the reaction favors conversion at elevated temperatures and lower ethylbenzene partial pressure; steam primarily serves as a diluent to shift equilibrium and mitigate the production of by-products such as benzene, toluene, light gases, and CO_x, which become more pronounced under extreme conditions. The temperature decrease between beds is mitigated by interstage reheating, followed by cooling of the effluent for phase separation and recovery of unreacted ethylbenzene for recycling (stream 47).

Steam Reforming

Heavier oil containing C_{9+} (stream 23) is used in the process of steam reforming for syngas production. The oil is vaporized and exposed to steam over a reforming catalyst at an elevated temperature (around 1000°C) and regulated pressure to transform hydrocarbons into a syngas mixture. The main products are H_2 and CO, while CO_2 and leftover CH_4 (stream 55) are produced via the water–gas shift reaction and incomplete reforming. The steam-to-carbon ratio, temperature, and pressure are established to reduce the formation of coke and methane while optimizing the stoichiometric number ratio (SN) for the following methanol synthesis.

Methanol Production

Syngas (stream 55) from the upstream reforming section is provided to a methanol synthesis loop that is based on the low-pressure process of the Imperial Chemical Industries (ICI) process. Before entering the catalytic reactor, the fresh gas is compressed in two stages to approximately 80 bar, combined with recycled converter gas to regulate the inlet temperature, and partially preheated. To regulate the exotherm and maintain equilibrium in each bed, divided fractions of cool syngas are injected between catalyst beds in a four-bed quench reactor. The crude methanol (stream 67) is recovered in a knockout cylinder after the reactor effluent is chilled and partially condensed at 260°C and 76 bar. The unreacted gas (stream 70) is recompressed for recycling, while a portion (stream 74) is water-scrubbed to recover residual methanol. A minor purge is implemented to prevent inert build-up.

Modelling and Simulation

The integrated process described in the previous section was simulated in Aspen Plus V14. NRTL-RK (Non-Random Two-Liquid-Redlich–Kwong) was used across

HTP, ethylbenzene dehydrogenation, steam reforming, and methanol synthesis to keep one consistent thermodynamic framework. NRTL captures non-ideal liquid behavior in water–organics and aromatic mixtures, while RK provides a reasonable vapor-phase model for high-temperature, high-pressure gas streams.

In the proposed design, different reactor models were selected according to the available mechanistic information and the objective of each unit. The hydrothermal processing step was modeled with an RYIELD reactor because it involves complex reaction pathways and lacks detailed kinetic data, so the model was based on experimental product yields. In contrast, RPLUG reactors were used for ethylbenzene dehydrogenation and methanol synthesis because reliable kinetic expressions are available. Steam reforming was modeled with an REQUIL reactor because the heavy C_{9+} fraction was treated as a lumped feed, and the main aim was to estimate equilibrium syngas composition for methanol production.

An RYIELD reactor model (R-101) was chosen for the hydrothermal processing based on experimental data from Gentilcore et al. [4]. Three RADFRAC models in equilibrium mode (from C-101 to C-103) were selected to purify toluene, ethylbenzene, and styrene, respectively. A sensitivity analysis was considered for each distillation column to vary the feed-stage number and calculate the reboiler heat duty. The feed-stage number of each distillation column was determined based on the minimum reboiler heat duty.

Ethylbenzene was mixed with steam in the ratio of 1:11 (EB: ST) to minimize coke formation and higher efficiency based on the available experimental data [8]. The temperature of both furnaces (H-203 and H-204) was set to 620°C to balance the conversion and selectivity, and serve as a catalyst-activity benchmark [9]. Two RPLUG models were selected to kinetically simulate ethylbenzene conversion. The kinetic equations and parameters proposed by Elnashaie *et al.* were validated based on the experimental data [10]. These reactions are presented as follows:

Main reaction:

$$(C_6H_5)CH_2CH_3 \leftrightarrow (C_6H_5CH)CH_2 + H_2 \quad (1)$$

Side reactions:

$$(C_6H_5)CH_2CH_3 \rightarrow C_6H_6 + C_2H_4 \quad (2)$$

$$(C_6H_5)CH_2CH_3 + H_2 \rightarrow (C_6H_5)CH_3 + CH_4 \quad (3)$$

$$2H_2O + C_2H_4 \rightarrow 2CO + 4H_2 \quad (4)$$

$$H_2O + CH_4 \rightarrow CO + 3H_2 \quad (5)$$

$$H_2O + CO \rightarrow CO_2 + H_2 \quad (6)$$

The selectivity of styrene is defined as follows:

$$Selectivity = \frac{F_{ST,out} - F_{ST,in}}{F_{EB,in} - F_{EB,out}} \quad (7)$$

Where $F_{ST,\,out}$ and $F_{ST,\,in}$ are molar flow rates of styrene inside (first bed) and outside (second bed) of the reactor, respectively. Moreover, the products of this reactor cooled down to separate water, light gases, and oily material in V-201. A RADFRAC model equilibrium-mode(C-201) was chosen to purify the generated benzene as a by-product in this process.

To optimize the ethylbenzene dehydrogenation section for styrene production, two targets were considered: styrene selectivity and the styrene molar flow rate at the end of the second bed. Because economics are not assessed here, neither metric alone is sufficient to define a practical optimum. In principle, the styrene molar flow rate would continue to increase with reactor length, whereas the highest selectivity would occur with a vanishing second-bed volume, since byproduct formation rises as conversion is pushed further. These issues happen when considering only maximum selectivity or styrene molar flow rate for the optimization. To balance these competing trends, the objective function was defined as the product of styrene selectivity and styrene molar flow rate. The optimization was then applied to the effluent of the second bed, R-201B (stream 31), to identify operating conditions that improve styrene output while limiting the loss of selectivity to side reactions.

An R-Equil reactor model was used for energy and mass balance of the process of steam reforming. The general formula for the steam reforming of the oil was defined as follows:

$$C_nH_m + nH_2O \leftrightarrow nCO + (n + \frac{m}{2})H_2 \quad (8)$$

The desired syngas must have a stoichiometric number (SN) close to 2 (on a molar basis). For methanol synthesis, an SN close to 2 corresponds to a near-stoichiometric hydrogen availability for converting COx to MeOH, while accounting for the CO2 “hydrogen penalty” associated with RWGS. In industrial loop design, values around 2.0–2.1 are commonly reported as optimal: if SN is higher than this range, excess H2 tends to accumulate in the recycle and must be removed by a larger purge, reducing loop efficiency; if SN is lower than 2, the make-up gas becomes hydrogen-limited, which typically lowers per-pass MeOH yield and increases water formation and recycle/purge requirements. Therefore, SN close to 2 was used here as a design target for the make-up gas composition, reflecting established catalyst-operability and loop-design practice rather than an economic optimum. Stoichiometric number (SN) is defined as follows:

$$SN = \frac{H_2 - CO_2}{CO + CO_2} \quad (9)$$

Where H_2, CO_2, and CO are the molar flowrates of each component in stream 55. Reactor (R-301) temperature, pressure, and injected steam flow rate (stream 48) were the variables in the process of steam reforming. These

variables were used in Aspen Plus to optimize the required syngas for methanol production to achieve the SN equal to two and the highest hydrogen to methane ratio.

Synthesis gas was pressurized in a 2-stage compressor (K-401 A and B) up to 80 bar. The introduced reactor in this process is based on the ICI method, which is described in the previous section. The following equations belong to the methanol reactor (R-401):

$$CO_2 + 3H_2 \leftrightarrow CH_3OH + H_2O \quad (10)$$

$$CO_2 + H_2 \leftrightarrow \mathrm{CO} + H_2O \quad (11)$$

$$2CO + 4H_2 \rightarrow C_2H_5OH + H_2O \quad (12)$$

$$2CH_3OH \rightarrow CH_3OCH_3 + H_2O \quad (13)$$

Four PLUG reactors were considered to simulate the process of methanol production using kinetic equations in the literature [11]. To evaluate the accuracy of the results, the model was validated based on experimental data [12]. In addition, a separate optimization block was introduced to maximize methanol production by varying the dimensions of the catalyst beds. The methanol loop was modelled as four insulated beds (i.e., zero external heat duty). Because the main synthesis reactions are exothermic, operating at high temperature can shift the equilibrium away from methanol and may even promote net methanol consumption at sufficiently high conversion and temperature. The bed sizes, therefore, cannot be increased arbitrarily; each bed must be long enough to achieve meaningful conversion, but not so long that the temperature rise drives the system into an unfavourable equilibrium region. Since the reactor consists of a single shell (one diameter) with multiple beds of different lengths, several geometrical combinations could deliver similar methanol molar flow rates when economics are not considered. To obtain a unique and comparable solution at this stage, the reactor diameter was fixed at 5 m, and the length of each bed was treated as the decision variable to identify the configuration that maximizes the methanol molar flow rate.

Table 1 summarizes the decision variables for each unit that were adjusted to determine the unit's optimal operating conditions for product maximization.

RESULTS AND DISCUSSION

Validation of Kinetic Mechanisms

Table 2 compares the outlet molar flow rates from the Aspen Plus kinetic model with the corresponding industrial plant data for ethylbenzene dehydrogenation. The simulation reproduces the key species reasonably well: ethylbenzene is predicted within 6.9% and styrene within 10.4%, indicating that the implemented kinetics capture the main conversion and selectivity trends. Benzene is also close (11.1%), whereas toluene shows the largest deviation (28.5%), likely reflecting its low concentration and higher sensitivity to side reactions and operating details not fully represented in the model.

Table 1: Decision variables for optimization

Variable	Lower Bound	Upper bound
Ethylbenzene Dehydrogenation		
R-201A Diameter	0.9 m	7 m
R-201B Diameter	0.9 m	7 m
R-201A Length	2 m	9 m
R-201B Length	2m	6 m
Stream 27 Pressure	0.4 bar	8 bar
Steam Reforming Conditions		
R-301 Pressure	20 bar	35 bar
R-301 Temperature	800°C	1000°C
Steam Flowrate	15, 000 kg/hr	80, 000 kg/hr
Methanol Reactor Bed Lengths		
R-401 First Bed	1 m	2 m
R-401 Second Bed	2 m	5 m
R-401 Third Bed	1 m	3 m
R-401 Fourth Bed	1 m	3 m

Table 3 compares the reactor outlet composition and temperature from the Aspen Plus kinetic model with industrial plant data. Overall agreement is good: methanol is predicted within 5.3% relative error, while the main inerts and light gases are captured with small-to-moderate deviations (N_2: 0.34%, CH_4: 10.5%, CO: 13.2%). The outlet temperature is reproduced within 2.35%, indicating that the heat release and reaction rate levels are broadly consistent. The largest mismatch is observed for CO_2 (22.7%), suggesting higher sensitivity to the CO/CO_2 shift balance under the selected kinetics and operating conditions.

Table 2: Validation of the methanol synthesis reactor model against industrial data

Outlet Component (kmol/h)	Industrial Data [10]	Aspen Plus (This Model)	Error (%)
Benzene	1.5	1.35	11.11
Toluene	2.03	1.58	28.48
Ethylbenzene	19.45	18.20	6.86
Styrene	15.57	17.37	10.36

Sensitivity and Optimization

Figure 2 illustrates a three-dimensional sensitivity map of the syngas stoichiometric number (SN) for the steam reforming of hydrothermal-processed heavy oil, characterized by a lumped C_9–C_{20+} proportion. The horizontal axes represent reactor pressure and reactor temperature, while the vertical axis indicates the resultant SN. Each point represents a steady-state scenario achieved by altering the controllable variables: reactor

temperature, reactor pressure, and the syngas to oil ratio introduced into the reactor. The point color denotes the steam flow rate to oil ratio, as specified by the color scale. It is worth mentioning that heavier fractions (e.g., C_{19} and C_{20}) cannot fully convert into syngas in the steam reformer at pressures lower than 20 bar. Therefore, we put 20 bars for the lower bound in the optimization block, as shown in Table 1.

Table 3: Comparison of industrial and simulated outlet flowrates for ethylbenzene dehydrogenation

Outlet Component	Industrial Data [12]	Aspen Plus (This Model)	Error (%)
MeOH (mol%)	4.74	4.99	5.3
CO_2 (mol%)	2.83	2.19	22.7
CO (mol%)	1.88	1.63	13.2
CH_4 (mol%)	12.86	11.50	10.6
N_2 (mol%)	2.35	2.34	0.3
Outlet T (K)	544.7	531.89	2.4

The results reveal significant sensitivity of SN to both pressure and temperature, highlighting a distinct interaction between the two variables. The highest stoichiometric number occurs at 520°C, 50 bar, and 10 steam to oil ratio. For any given temperature and pressure, the vertical dispersion of points suggests that syngas injection presents a viable method to adjust SN without changing the main operational parameters. Elevated steam flow is often associated with increased SN, which corresponds with the effects of steam on reforming and water-gas shift processes.

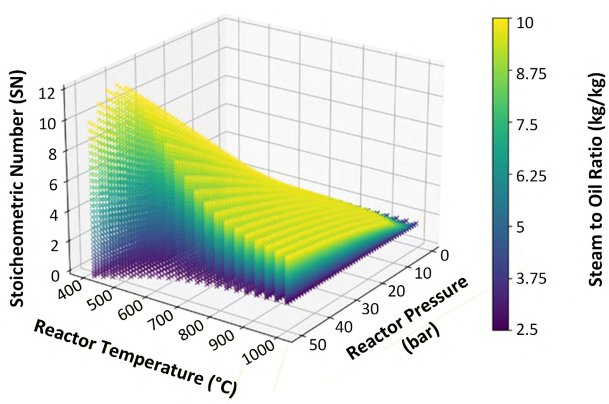

Figure 2. Sensitivity analysis of syngas stoichiometric number (SN) in steam reforming of hydrothermal-processed heavy oil (C_9–C_{20+}) as a function of reactor temperature, pressure, and syngas injection (steam flow rate as colour scale).

Figure 3 shows the sensitivity of the H_2/CH_4 molar ratio to the total required heat duty (Q_t) to oil flow rate ratio, across the evaluated operating cases. Points are colour-coded by reactor pressure (20–50 bar), highlighting how pressurisation shifts the H_2/CH_4–Q_t relationship. All of the points indicate a stoichiometric number between 2 and 2.01. The dataset follows a strongly non-linear trend: H_2/CH_4 remains low and changes gradually at lower heat duties, then increases sharply beyond roughly 7.5-8 of Q_t to oil flow rate ratio. To minimize the accumulation of methane in the methanol reactor, it is important to find the optimum point with the lowest methane in the syngas.

Table 4 summarizes the optimized operating and design variables that maximize the production of the two main value-added products, styrene and methanol, for the specified feedstock flow rate. The reported values represent the best-performing configuration under the present model assumptions and constraints, and they provide a consistent basis for comparing the two synthesis sections within the overall process.

Table 5 shows the summary of the input and output of the proposed process flow diagram. It also suggests that the proposed design has the possibility of further investigation through economic analysis, energy evaluation, and life cycle assessment.

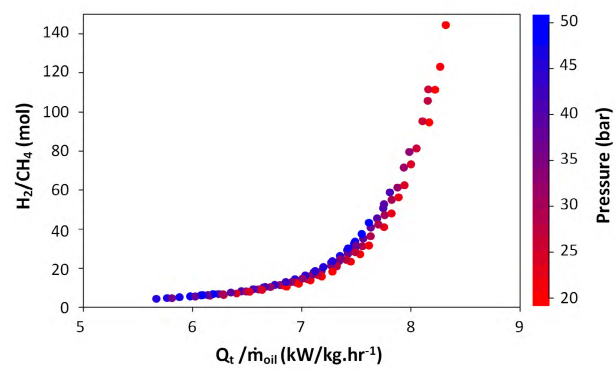

Figure 3. Sensitivity of H_2/CH_4 ratio to the total required heat duty, considering the reactor pressure at SN≈2

Following preliminary heat integration, the process required 85.01 MW of hot utility and 70.93 MW of cold utility. These values remained high mainly because the final distillation column (C-403), used to separate only 49 kg/h of ethanol, imposed substantial energy demands, with condenser and reboiler duties of 23.06 MW and 23 MW, respectively. After removing the final methanol distillation column and carrying out a new heat integration, the external utility demand decreased to 59.76 MW for heating and 47.4 MW for cooling. This reduction is significant because purchased heat represents the main variable operating cost of the process. Based on 12 tonne/h of polystyrene feed, 59.76 MW is equivalent to 4.98 MWh_{th} per tonne of PS. Assuming steam is generated from natural gas at an efficiency of about 80%, the

corresponding fuel demand is 6.23 MWh_{fuel} per tonne of PS [13]. Using an EU non-household natural gas price of 6.24 €/100 kWh, excluding VAT, the estimated fuel cost is about 389 €/tonne PS [14]. In addition, direct combustion emissions, estimated using the IPCC emission factor for natural gas, reach about 1.26 tonne CO_2/tonne PS. At a carbon price of 67.89 €/tonne CO_2, this adds roughly 85 €/tonne PS, giving a combined heat and carbon cost of about 474 €/tonne PS. Although this improves the economic position of the process, profitability still depends strongly on product value, particularly the netbacks of styrene and toluene, as well as capital control in the heat exchanger network and separation section.

Table 4: Decision variables for optimization

Variable	Optimal Value
Ethylbenzene Dehydrogenation	
R-201A Diameter	1.31 m
R-201B Diameter	0.98 m
R-201A Length	2.60 m
R-201B Length	3.55 m
Stream 27 Pressure	3.03 bar
Steam Reforming Conditions	
R-301 Pressure	20.2 bar
R-301 Temperature	989°C
Steam Flowrate	37, 780 kg/hr
Methanol Reactor Bed Lengths	
R-401 First Bed	1.23 m
R-401 Second Bed	2.45 m
R-401 Third Bed	1.86 m
R-401 Fourth Bed	1.71 m

Table 5: Main Process Inputs, Utility Requirements (before heat integration), and Product Outputs

Variable	Final Value
Inputs to Process	
Polystyrene Waste	12, 000 kg/hr
Total Water	45, 300 kg/hr
Required Heat	148 MW
Required Cooling	134. MW
Outputs from Process	
Toluene	1, 600 kg/hr
Benzene	10.2 kg/hr
Styrene	1910 kg/hr
Methanol	15, 200 kg/hr
Ethanol	49 kg/hr

CONCLUSIONS

An integrated route for upgrading polystyrene waste via low-pressure hydrothermal processing was modelled in Aspen Plus and connected to separation, ethylbenzene dehydrogenation, steam reforming, and an ICI-type methanol loop. In the reforming section, the syngas stoichiometric number (SN) was highly sensitive to the temperature–pressure, and steam flow rate combination. The optimized case (≈1000°C, 20.2 bar, 37, 800 kg/h steam) hit SN~2 while limiting methane carry-over; it provides an ideal syngas for methanol production. At steady state, 12, 000 kg/hr PS could yield about 1, 603 kg/hr toluene, 1, 909 kg/h styrene, and 15, 193 kg/hr methanol, with minor benzene and ethanol, but the gross energy demand is still 279.7 MW when heat recovery is excluded. The concept looks technically sound; the next step is heat integration and a full techno-economic and LCA check.

SUPPORTING INFORMATION

The Aspen Plus simulation file is available at: https://psecommunity.org/LAPSE:2026.0025

AUTHOR IDENTIFIERS

Author ORCIDs:
HajiHashemi M: 0000-0002-9421-4004
Schulze-Netzer C: 0000-0002-4445-7243
Adams TA II: 0000-0002-9871-9851

REFERENCES

1. Geyer R, Jambeck JR, Law KL. Production, use, and fate of all plastics ever made. Sci. Adv. 3: (2017). https://doi.org/10.1126/sciadv.1700782
2. Recycling Today. Polystyrene recycling programs expand despite bans. April 24 P, 2019. [accessed 2025 Oct 15]. Available from: https://www.recyclingtoday.com/news/polystyrene-recycling-programs-expand-despite-bans/
3. Chaukura N, Gwenzi W, Bunhu T, Ruziwa DT, Pumure I. Potential uses and value-added products derived from waste polystyrene in developing countries: a review. Resources, Conservation and Recycling 107:157-165 (2016). https://doi.org/10.1016/j.resconrec.2015.10.031
4. Gentilcore C, Jin K, Barzallo G, Vozka P, Wang NHL. Low-pressure hydrothermal processing for conversion of polystyrene into oils. Journal of Environmental Chemical Engineering 12:113836 (2024). https://doi.org/10.1016/j.jece.2024.113836
5. Tarafder A, Rangaiah GP, Ray AK. Multiobjective optimization of an industrial styrene monomer manufacturing process. Chemical Engineering Science 60:347-363 (2005). https://doi.org/10.1016/j.ces.2004.07.120
6. Chen CC, Duh YS, Shu CM. Thermal polymerization of uninhibited styrene investigated by using microcalorimetry. Journal of Hazardous Materials 163:1385-1390 (2009).

https://doi.org/10.1016/j.jhazmat.2008.07.151
7. Darvishi A, Rahimpour MR, Raeissi S. A theoretical and experimental study for screening inhibitors for styrene polymerization. Processes 7:677 (2019). https://doi.org/10.3390/pr7100677
8. Dimian AC, Bildea CS. Energy efficient styrene process: design and plantwide control. Ind. Eng. Chem. Res. 58:4890-4905 (2019). https://doi.org/10.1021/acs.iecr.8b05560
9. Liao S, Tang Q, Zhou G, Lu B, Liu W. Industrial application of macrocat-201s catalyst with low steam–oil ratio in large-scale styrene unit. Catalysts 15:308 (2025). https://doi.org/10.3390/catal15040308
10. Elnashaie SSEH, Abdalla BK, Hughes R. Simulation of the industrial fixed bed catalytic reactor for the dehydrogenation of ethylbenzene to styrene: heterogeneous dusty gas model. Ind. Eng. Chem. Res. 32:2537-2541 (2002). https://doi.org/10.1021/ie00023a016
11. Salkuyeh YK, Adams TA II. A new power, methanol, and DME polygeneration process using integrated chemical looping systems. Energy Conversion and Management 88:411-425 (2014). https://doi.org/10.1016/j.enconman.2014.08.039
12. Bisotti F, Fedeli M, Prifti K, Galeazzi A, Dell'Angelo A, Manenti F. Impact of kinetic models on methanol synthesis reactor predictions: in silico assessment and comparison with industrial data. Ind. Eng. Chem. Res. 61:2206-2226 (2022). https://doi.org/10.1021/acs.iecr.1c04476
13. U.S. Department of Energy (DOE), Industrial Technologies Program. Consider Installing High-Pressure Boilers with Backpressure Turbine-Generators. Steam Tip Sheet #22. DOE/GO-102006-2267 (January 2006). Available from:
14. https://www.nrel.gov/docs/fy06osti/39324.pdf
15. Eurostat. EU household gas prices rise in the second half of 2024. Eurostat News Article (6 May 2025). [accessed 2026 Mar 26]. Available from:
16. https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20250506-3

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Uncertainty Prioritisation for Water-Energy-Food-Land Nexus Optimisation

Md Shamsul Alam[a], I. David L. Bogle[a], Vivek Dua[a*]

[a] Department of Chemical Engineering, The Sargent Centre for Process Systems Engineering, University College London, Torrington Place, WC1E 7JE, London, UK
*Corresponding author: v.dua@ucl.ac.uk.

ABSTRACT

The interdependence among energy, water, food, and land sectors has been addressed through the concept of Energy-Water-Food-Land nexus (EWFLN), where interconnections between different sectors generate complex feedback loops. In the field of EWFLN, transitioning from deterministic to stochastic approach is considered as the natural choice for policy makers. Working with a large number of uncertain parameters can make a stochastic system complex. Therefore, Identification of the significant uncertain parameters in the system would create a more acceptable model before transforming from a deterministic to a stochastic approach. This study incorporates uncertainty prioritisation in the mathematical model for the optimisation of EWFLN. The specific objectives of this study include creating a mathematical model, determining and prioritising uncertain parameters, and prescribing appropriate policy recommendations. This study points out clearly which specific parameters should be taken into consideration for risk hedging. The sensitivity analysis results present a quantitative hierarchical ranking of uncertain parameters based on their influence on the objective function. The results demonstrate that energy price, solar radiation flux and crop yield are the most sensitive among them all. Moreover, it provides information about prioritised uncertain parameters to the policy makers to avoid dangerous risk of food, water, and energy scarcity.

Keywords: Resource allocation, uncertain parameter, stochastic approach

INTRODUCTION

The existence of human beings is under threat without proper utilisation of basic lifesaving resources like food, water, and energy. Food is produced by the agriculture, livestock, and aquaculture sectors. On the whole, production of food cannot be dispensed without water and energy. Therefore, for ensuring food security, energy and water security must also be ensured. EWFLN, a widely used scholarly concept, demonstrates complex interactions [1, 2, 3], which can be used in an innovative way to manage optimal resources in a sustainable way [4]. According to the UN World Water Development report 2024, the agricultural sector consumes 70% of global freshwater while the domestic sector accounts for approximately 10-12% [5]. However, the International Renewable Energy Agency and Food and Agriculture Organisation [6] report that the food system including 'production to table' accounts for about 30% of total global energy consumption.

In this study of optimisation of biomass energy production in the water-food-energy-land nexus system, it is inevitable to have uncertainties in the determination of parameters due to insufficient information, variations of natural resources, fluctuations of socio-economic conditions as well as subjective judgment [7, 8]. Different authors have used different techniques to manage uncertainties in the parameters in the field of optimisation of EWFLN. This optimisation framework, combined with scenario analysis of population growth, has been designed to assess EWFLN systems in the study of Peña-Torres et al. [9].

Xie et al. [10] studied an inexact stochastic-fuzzy programming model with a view to allocating irrigation water and land resources with multiple uncertainties such as surface water and groundwater availability, the amount of random effective rainfall and the imprecise water requirements of crops. The results revealed that

water allocation is explicitly affected by uncertainties, and the results also provided information for the adjustment or justification of the existing water and land resources management. Cansino-Loeza, Tovar-Facio and Ponce-Ortega [11] provided a mathematical model for disadvantaged rural areas for their availability of basic needs to produce water, food and energy sustainably with provision for determining the mix of technology consistent with demand and nutritional requirements. This work incorporates mixed integer stochastic linear program with population change, water and energy demand change as uncertain parameters. The results provided positive feedback that enables the community to be self-sufficient and provides a balanced diet for inhabitants using renewable technologies considering waste treatment.

After reviewing all the literature, the authors believe that prioritisation of uncertainty in the EWFLN has not been fully explored and addressed. Therefore, prioritisation of uncertain parameters through sensitivity analysis will be presented and its impact will be discussed in this study.

NOMENCLATURE

Index	Description
o	Sources of water
a	Agriculture sectors
L	Livestock sectors
mt	Livestock types
c	Crop types
d	Domestic sectors
t	Time period in months
A	Water and energy consuming sectors
p	Biomass gasification power plant
sc	Sectors such as agriculture, livestock, domestic, energy, water
al	Agriculture and livestock sector

Acronym	Description
$Actors$	Water and energy consuming sectors
AgL	Agriculture and livestock
ARW_A	Effective area for rainwater harvesting in the sectors
$ASolarSector_A$	Effective area for solar panel in the sectors
$ASolaroofSurAgMax_a$	Maximum limit of effective area of solar panel installed in surface water pumping location
αF	Water efficiency coefficient for energy generation
$CostNexus$	Total cost in all sectors
$AgL^{dom}_{al,d,t}$	Raw food sent to domestic from agriculture and livestock sector
$AgL^{ext}_{al,t}$	Raw food sent to market from agriculture and livestock sector
$ALndagb_a$	Area of land dedicated for crop production in agriculture sector
$ConfCrpToBmas_c$	Conversion factor: crop to biomass
$CofeedToMeat_{mt}$	Monthly amount feed required to produce unit amount of meat
$Cost_{sc}$	Cost in an individual sector
$E^{To-Actor}_{A,t}$	Monthly amount of energy supplied to actors
$E^{\text{Plant}-Actor}_{p,A,t}$	Energy distribution to actors from power plant
$E^{Solar-Actor}_{A,t}$	Energy distribution to actors from solar panel
$E^{Solar-Ext}_{A,t}$	Monthly energy sold from solar panel to external grid
$ProfitNexus$	Overall nexus profit
$\alpha AgWst_c$	Fuel efficiency coefficient for energy production
$F^{Rain-Actor}_{A,t}$	Water harvested in all sectors such as domestic, livestock
$WRnflux_{A,t}$	Rainwater precipitation flux in the sectors
$FoodDemand^{dom}_{d,t}$	Amount of food required in domestic sector
$F^{To-Actor}_{A,t}$	Total water Input to a sector
$F^{WSource-Actor}_{o,A,t}$	Monthly water distribution from a water source to actors
$F^{Wsource-\text{plant}}_{p,t}$	Monthly water distribution in power plant

$GRfood_{al,t}^{AgL}$	Production of raw food such as crop from agriculture and meat from livestock
$GE_{p,t}$	Monthly energy generation in a gasification power plant
$GE_{A,t}^{Solar-Actor}$	Solar energy generation in the sectors
MI_{sc}	Amount of input resource sold in a sector
MR_{sc}	Amount of resource sold by an individual sector
$MAgWst_c$	Amount of biomass type c fed for energy generation per month
$MCrp_{a,c,s}$	Amount of crop produced in agriculture sector in a season
$MAgWaste_{a,c,s}$	Amount of crop waste produced in agriculture sector in a season
$MFeed_{mt,t}$	Monthly amount of feed required for meat production in livestock sector
$MLiv_{mt,t}$	Monthly amount of meat production in livestock sector
$SRflux_t$	Solar radiation flux in month t
$Wsource$	Water source such as river, deep well
$RevNexus$	Total revenue from all sectors
Rev_{sc}	Revenue in an individual sector
$UCostI_{sc}$	Unit cost of input in a sector
$UCostWatAg$	Cost of water (per m^3) consumed in agriculture sector
UPR_{sc}	Price of resource output from individual sector
$UPWatAg$	Price of water (per m^3) distributed to agriculture sector
$UPEngyToExt$	Price of energy sold to external power grid
$URCrp_c$	Price of crop type c
$URAgWaste_c$	Price of agriculture waste (per ton)
$CofPrdAg_c$	Yield of crop
Z	Overall profit of the nexus

METHODS

This section includes superstructure, the mathematical model for maximising profit of the nexus, and the solution method. The superstructure includes addition of modular gasification biomass energy plant, solar power system in all sectors, water sources including rainwater and their overarching input/output network.

Mathematical model

Objective function

Maximisation of overall profit, expressed in terms of revenue minus cost, of the nexus is a single objective function for this study.

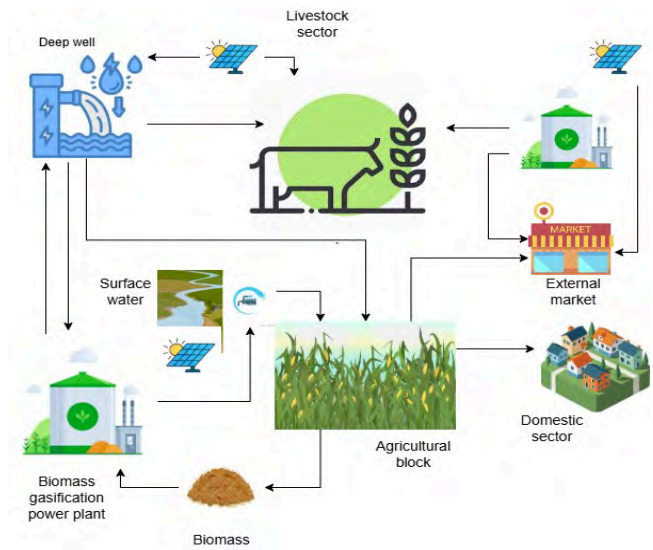

Figure 1. Superstructure of EWFLN

$$Max\ Z = RevNexus - CostNexus \tag{1}$$

$$RevNexus = \sum_{sc} Rev_{sc} \tag{2}$$

$$CostNexus = \sum_{sc} Cost_{sc} \tag{3}$$

$$Rev_{sc} = UPR_{sc}.MR_{sc} \tag{4}$$

$$Cost_{sc} = UCostI_{sc}.MI_{sc} \tag{5}$$

Water mass balance

In this mathematical model, the term 'actors' is defined as the sectors in the nexus which consume resources delivered by other sectors of the nexus. The amount of water distributed in actors is equal to the amount of water sent from water sources and the amount of rainwater harvested in that sector.

$$F_{A,t}^{To-Actor} = \sum_{o} F_{o,A,t}^{WSource-Actor} + F_{A,t}^{Rain-Actor} \quad (6)$$

$$F_{A,t}^{Rain-Actor} = ARW_A . WRnflux_{A,t} \quad (7)$$

Energy balance

Amount of energy generation in a biomass gasification power plant is expressed in terms of amount of water and biomass utilised for this purpose.

$$GE_{p,t} = \alpha F . F_{p,t}^{Wsource-\text{plant}} \quad (8)$$

$$GE_{p,t} = \alpha AgWst_c . MAgWst_c \quad (9)$$

Total energy in consuming bodies (actors) is equal to the amount of energy flow from generation units to actors and the amount of solar energy generated in that sector.

$$E_{A,t}^{To-Actor} = \sum_{p} E_{p,A,t}^{\text{Plant}-Actor} + E_{A,t}^{Solar-Actor} \quad (10)$$

$$GE_{A,t}^{Solar-Actor} = ASolarSector_A . SRflux_t \quad (11)$$

$$GE_{A,t}^{Solar-Actor} = E_{A,t}^{Solar-Actor} + E_{A,t}^{Solar-Ext} \quad (12)$$

Food mass balance in agriculture and livestock sectors

Crops from agriculture sectors and meat from livestock sectors are sources of food which are distributed to domestic sectors and external markets. Production of crops, biomass from crops, and meat are linearly related to the respective production coefficient factors.

$$MCrp_{a,c,s} = CofPrdAg_c . ALndagb_a \quad (13)$$

$$MAgWaste_{a,c,s} = MCrp_{a,c,s} . ConfCrpToBmas_c \quad (14)$$

$$MFeed_{mt,t} = CofeedToMeat_{mt} . MLiv_{mt,t} \quad (15)$$

$$AgL_{al,d,t}^{dom} = FoodDemand_{d,t}^{dom} \quad (16)$$

$$GRfood_{al,t}^{AgL} = AgL_{al,d,t}^{dom} + AgL_{al,t}^{ext} \quad (17)$$

Solution method

The model is implemented in four mouza (clusters of land) with four domestic sectors in Kazla union, Sariakandi upazila of Bogura districts of Bangladesh. Data was collected and implemented in GAMS software with CPLEX v.22.1.0.0 as solver for the LP. After running the model we have shortlisted the important parameters which can affect the objective function through looking at the shadow price of equations and the reduced cost of decision variables. In such a way we have screened out parameters which have no impact on the objective function and sorted 46 parameters which might have importance in shifting the value of the objective function. Then we have taken each parameter one at a time and run the model with ±20% deviation from the base parameter value and the percentage shift in the objective function is calculated. In this way we have sorted important parameters and prioritised them sequentially. Out of these, we have selected the three most sensitive parameters for further analysis. With these three most important parameters, we have made 125 joint independent scenarios; each parameter having five scenarios. The scenarios L20, L10, Bse, H10, and H20 are 20% lower, 10% lower, 10% higher, and 20% higher than the base scenario 'Bse' respectively. Thus, we have run the model with 125 scenarios to find the impact of uncertain parameters on the objective function.

RESULTS

In Table 1 and Figure 2, the term "swing" refers to the percentage change of overall nexus profit resulting from changes in parameter value relative to the base parameter value. The swing of a parameter is from 20% lower and 20% higher than the base parameter value. The most sensitive parameters are 1) UPEngyToExt 2) SRFlux(t) (Figure 2, table 1) which decrease the energy sector's parameters. The second range of parameter falls are in the agriculture sector's parameters which are more than three times lower than the energy sector parameters and include 1) CofPrdAg(c) 2) URCrp(c) and 2) Alndagb(a). The third range is cost parameters which affect negatively the overall nexus profit and include 1) UCostLabr 2) UnituseLabr 3)UPWatAg, and 4) UCostWatAg. The fourth category affect minutely the overall nexus profit and include 1) ASolaroofSurAgMax(a) 2) ConfCrptoBmas(c) 3) URAgWaste(c). The range of prioritised parameters start with the first prioritised parameter "UPEngytoExt" and end with the twelfth prioritised parameter "URAgWaste(c)". In this study, we have not considered the parameters as prioritised which have a percentage swing of less than 1. The percentage swing of the most prioritised parameters, UPEngyToExt, SRflux(t) and CofPrdAg(c), are 36, 35 and 10% respectively. In practice, UPEngyToExt, SRflux(t) and CofPrdAg(c) are independent, and variation of any of them does not have any relationship with variation of other parameters. Now, we shall analyse how variation of these three independent parameters combined affects the overall profit of the nexus.

Figure 3 shows a range of values of overall profit resulting from 125 joint uncertain parameter scenarios of the three parameters. The lowest profit is in the LLL (lowest value of all 3 parameters) joint scenario which is 17.43 million USD while the highest profit is observed in HHH (Highest value of all 3 parameters) joint parameter

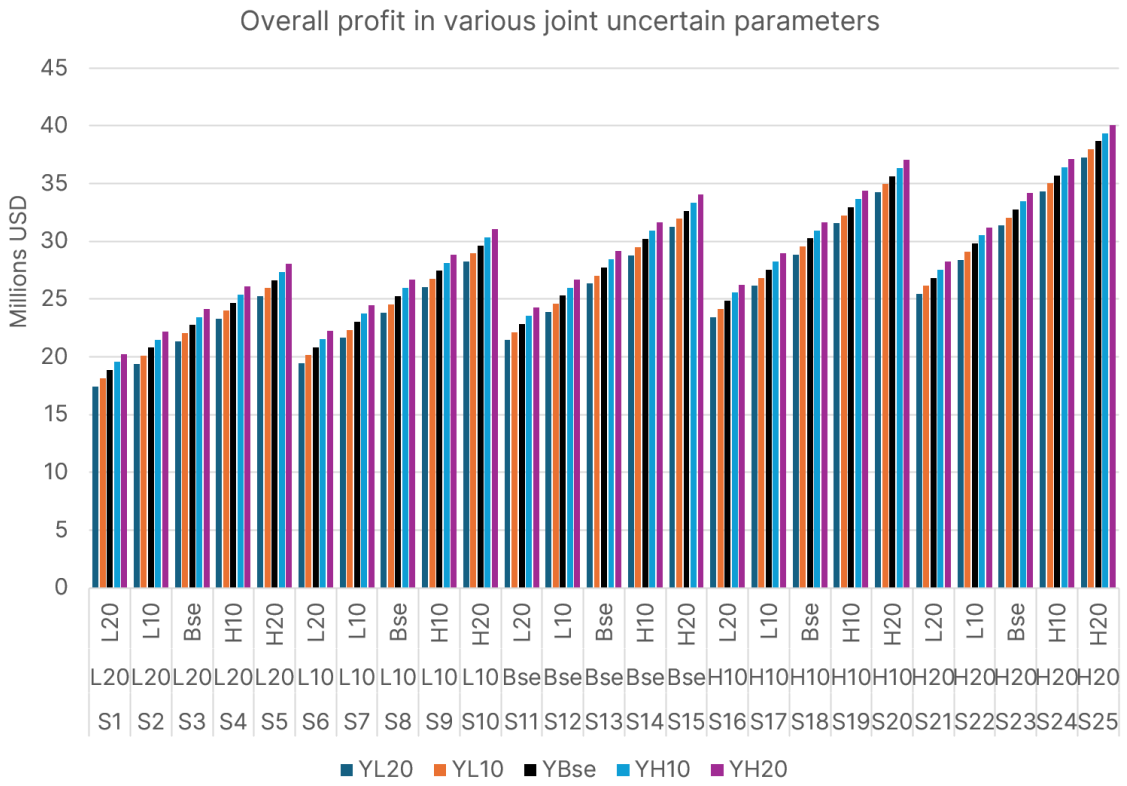

Figure 3. Overall profit of the nexus in different scenarios

scenario which is 40 million USD. The results indicate that the maximum profit(from HHH scenario) is approximately 2.3 times the minimum profit (from LLL scenario).

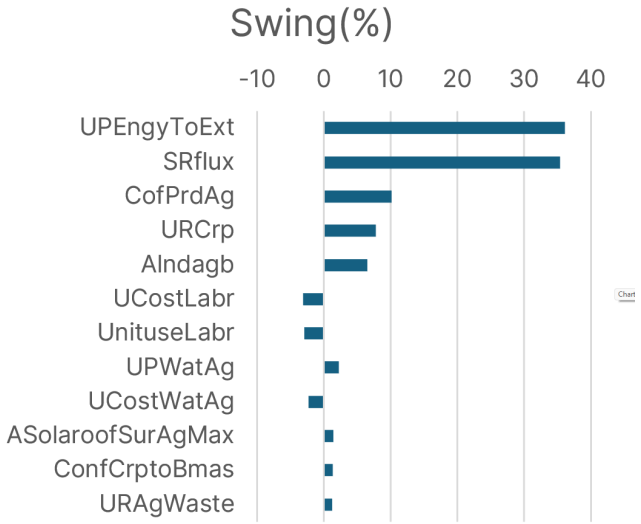

Figure 2. Tornado graph for sensitivity analysis

Table 1. Prioritised parameters

Serial	Parameters	Swing (%)
1	UPEngyToExt	36.08876914
2	Flux	35.35711568
3	CofPrdAg	10.12992412
4	URCrp	7.798400723
5	AIndagb	6.551808129
6	UCostLabr	-3.087921957
7	UnituseLabr	-2.921180162
8	UPWatAg	2.267019713
9	UCostWatAg	-2.267019713
10	ASolaroofSurAgMax	1.490959063
11	ConfCrptoBmas	1.335354998
12	URAgWaste	1.240125902

CONCLUSION

Uncertainty prioritisation is an important field of study in the realm of stochastic optimisation. This study provides a structured uncertainty prioritisation in the framework of the optimisation of EWFLN by maximising

the overall nexus profit as objective function. Within a range of candidate parameters in the mathematical model of EWFLN, we conducted a one-at-a-time (OAT) sensitivity analysis which revealed 12 important prioritised parameters. We have observed that the three parameters such as UPEngytoExt, SRflux(t) and CofCrpPrd(c) are the most sensitive parameters within these 12 parameters in this study. Therefore, with three parameters, each with 5 scenarios, 125 discrete joint parameter scenarios are constructed. Solving the optimisation model with each scenario has generated scenario wise overall nexus profit results. Visualised in the bar chart (Fig 3), the results show that the profit response is not the same across the joint scenario space. Three parameters are in the positive direction of profit maximisation; the lowest profit is 17.43 million USD while the highest profit is 40.00 million USD.

Practically, these findings focus on targeted uncertain parameters: the top ranked parameters are the most valuable for improvement of the objective function. Moreover, they are significantly important for risk hedging. The limitation includes the independence assumption among the three most prioritised parameters, discretized uncertainty ranges and reliance on OAT impact and observation of the list file (of GAMS code) for initial screening out. The future work should consider correlated uncertainties, probabilistic weighting of scenarios and stochastic optimisation incorporating these prioritised uncertain parameters.

ACKNOWLEDGEMENTS

This research is supported by “Strengthening Government through Capacity Development of the BCS Cadre Officials (Phase-II)” project implemented under the Ministry of Public Administration, Government of the People's Republic of Bangladesh

REFERENCES

1. Chamas Z, Abou Najm M, Al-Hindi M, Yassine A, Khattar R. Sustainable resource optimization under water-energy-food-carbon nexus. Journal of Cleaner Production 278:123894 (2021). https://doi.org/10.1016/j.jclepro.2020.123894
2. Hooda PS, Edwards AC, Anderson HA, Miller A. A review of water quality concerns in livestock farming areas. Science of The Total Environment 250:143-167 (2000). https://doi.org/10.1016/s0048-9697(00)00373-9
3. Wicaksono A, Jeong G, Kang D. Water, energy, and food nexus: review of global implementation and simulation model development. Water Policy 19:440-462 (2017). https://doi.org/10.2166/wp.2017.214
4. Alam MS, Bogle IDL, Dua V. Optimisation of biomass-energy-water-food nexus under uncertainty. Systems and Control Transactions 4:86-91 (2025) http://doi.org/10.69997/sct.117909
5. United Nations. The United Nations World Water Development Report 2024: Water for Prosperity and Peace. UNESCO (2024) ISBN 978-92-3-100657-9.
6. International Renewable Energy Agency and Food and Agriculture Organization of the United Nations. Renewable energy for agri-food systems: Towards the Sustainable Development Goals and the Paris Agreement (2021). https://doi.org/10.4060/cb7433en
7. Azadeh A, Vafa Arani H, Dashti H. A stochastic programming approach towards optimization of biofuel supply chain. Energy 76:513-525 (2014). https://doi.org/10.1016/j.energy.2014.08.048
8. Purkus A, Röder M, Gawel E, Thrän D, Thornley P. Handling uncertainty in bioenergy policy design – a case study analysis of UK and german bioelectricity policy instruments. Biomass and Bioenergy 79:64-79 (2015). https://doi.org/10.1016/j.biombioe.2015.03.029
9. Peña-Torres D, Boix M, Montastruc L. Multi-objective optimization and demand variation analysis on a water energy food nexus system. Computers & Chemical Engineering 180:108473 (2024). https://doi.org/10.1016/j.compchemeng.2023.108473
10. Xie YL, Xia DX, Ji L, Huang GH. An inexact stochastic-fuzzy optimization model for agricultural water allocation and land resources utilization management under considering effective rainfall. Ecological Indicators 92:301-311 (2018). https://doi.org/10.1016/j.ecolind.2017.09.026
11. Cansino-Loeza B, Tovar-Facio J, Ponce-Ortega JM. Stochastic optimization of the water-energy-food nexus in disadvantaged rural communities to achieve the sustainable development goals. Sustainable Production and Consumption 28:1249-1261 (2021). https://doi.org/10.1016/j.spc.2021.08.005

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Data Transformation Techniques and its Influence in Hybrid Model Performance

Juan Federico Herrera-Ruiz[a], Carlos Eduardo Robles-Rodriguez[b], Cesar Arturo Aceves-Lara[b], Javier Fontalvo[a] and Oscar Andrés Prado-Rubio[a*]

[a] Grupo de Investigación en Aplicación de Nuevas Tecnologías (GIANT) Departamento de Ingeniería Química, Universidad Nacional de Colombia sede Manizales, Campus La Nubia, Manizales, 170003, Colombia
[b] Toulouse Biotechnology Institute, Toulouse, France
*Corresponding author: oaprador@unal.edu.co

ABSTRACT

The global transition toward sustainable energy has intensified research into biofuels, with bioprocess optimization playing a central role in achieving decarbonization goals. Biobutanol, in particular, is a high-value molecule for sustainable fuel applications due to its superior energy density and compatibility with existing infrastructure. However, model-based optimization of its production is hindered by traditional semi-structured kinetic models that often suffer from limited predictive robustness. To address this challenge, within this study we developed a hybrid modeling framework for *Clostridium saccharoperbutylacetonicum* that integrates mechanistic mass-balance equations with Gaussian Processes (GPs) aiming to describe the biobutanol formation rate. Here, we investigate the effect of data normalization techniques on hybrid model's prediction capabilities comparing min-max normalization, z-score normalization, and no transformation. For each data treatment strategy, 8, 000 hybrid models were trained using experimental fermentation datasets, and ensemble predictors were constructed to mitigate variability and enhance generalization. Results demonstrate that hybrid models significantly outperform parametric baselines, with the ensemble approach leveling the performance differences across all pretreatment techniques. Ensembles achieved validation RMSE reductions of approximately 10.1% to 10.4%, which are two times higher than the best individual models regardless of the scaling method used. Despite similar RMSE reductions, prediction envelopes varied along data transformations revealing tradeoffs between uncertainty and coverage. Notably, targeted hybridization of the butanol reaction rate propagated performance improvements to other states, such as glucose and biomass. These findings suggest that for hybrid modeling, ensemble design provides greater benefits for robustness than any data normalization technique.

Keywords: Biofuels, Industry 4.0, Butanol, Ensemble Learning, Hybrid Modeling

1. INTRODUCTION

Driven by the sustainable development goals, research into optimizing bioprocesses for bio-based fuels has shifted toward digitalized frameworks. Recent advancements in sensing technologies and computational power have enabled the integration of sophisticated modeling and simulation tools into biomanufacturing [1]. These developments, together with modern modeling and simulation tools, have catalyzed digitalization of biomanufacturing processes. In this context, hybrid modelling stands as a powerful paradigm, combining parametric (i.e. mechanistic representations) with non-parametric (i.e. data-driven) methods to mitigate their respective limitations [2] and accelerate advanced optimal design and operation within process system engineering [3].

Despite hybrid model's enhanced prediction capabilities, their formulation requires biological and operational data as inputs to inform the data-driven subcomponents. Traditionally, machine learning and computer science literature has stated that normalization techniques such as z-score normalization and min-max normalization have great influence on model performance of

data-driven techniques [4, 5]. Several studies have confirmed this especially when dealing with metabolomic data [6–8]. However, the effect of data transformation techniques in the performance of hybrid models remains largely unexplored [9].

This work evaluates the influence of data transformation techniques on the predictive robustness of hybrid models applied to Acetone-Butanol-Ethanol (ABE) fermentation. Herein, different hybrid model's structures are developed to improve state predictions of ABE fermentation by *Clostridium saccharoperbutylacetonicum*, with a particular focus on biobutanol as a key molecule for sustainable biofuel applications [10]. Although semi-structured kinetic models for butanol fermentation are well established, they are frequently overparameterized, which limits both structural and practical identifiability and degrades predictive robustness. Therefore, a hybrid modeling framework is here proposed, coupling mechanistic mass-balance equations with Gaussian Processes (GPs) used to represent the biobutanol formation rate. Furthermore, this study evaluates how different data transformation techniques affect hybrid model's performance of individual and ensemble models. This work is structured into four sections. Section 1 presents the case study and the goals of this contribution. Section 2 outlines modeling details, including data transformations. Section 3 analyzes the results, focused on performance of the various hybrid models. Section 4 presents conclusions and a summary of the key findings.

2. MODELING DETAILS

2.1 Experimental Data and Data preprocessing

This study utilized four fermentation datasets of *C. saccharoperbutylacetonicum* using glucose as the primary carbon source [11]. Data was extracted from the figures from the original paper [11]. Experimental measurements included glucose, acetone, acetate, butyrate, biomass, and butanol. Data were structured into three training sets and one validation set to rigorously evaluate predictive performance across different experimental runs. To ensure high-quality inputs for the hybrid model, raw signals underwent outlier detection and correction using a moving median window and Savitzky-Golay smoothing, respectively [12].

2.2 Data Transformation Techniques

Data was transformed using min-max normalization and z-score normalization, which are shown in Eq (2) and Eq (3), respectively:

$$X_{min-max} = \frac{X - min(X)}{max(X) - min(X)} \quad (2)$$

$$X_{z-score} = \frac{X - \bar{X}}{\sigma(X)} \quad (3)$$

Where X represents a vector of data, $\bar{X}$ represents their average values, and $\sigma(X)$ represents the vector's standard deviation. The transformation was applied to both inputs and outputs of the data-driven component. As a base line, data without transformation was also used.

2.3 Hybrid Model Formulation

The proposed hybrid model architecture uses Shinto's ABE parametric model as a baseline [11]. The mass balances in batch conditions consist of 19 reaction rates, 16 states and 46 parameters. Only 6 metabolites were experimentally monitored: glucose, acetone, acetate, biomass, butyrate, and butanol all expressed in Mm. For the hybrid model, Butanol's reaction rate $r_{butanol}$ (mM/h) was replaced by a GP, as seen in Eq(1):

$$\frac{dButanol}{dt} = r_{butanol} = GP(W) \quad (1)$$

Where W represents the features used as inputs for the GPs. Initial Glucose concentration along instantaneous glucose and acetone concentrations were used as features for training the GPs and making predictions.

2.4 Training Strategy and Evaluation

For each data transformation technique, a predictor ensemble strategy was proposed. Experimental data are employed to train the Gaussian Process (GP) models, while model-based inputs (outputs from ODEs) are used for evaluating hybrid model's performance in both training and validation sets. Specifically, at each integration step, the ODE solver provided a set of predicted concentrations as inputs to the GP to facilitate dynamic simulation. After integrating the responses, they were algebraically averaged to make predictions. This ensemble approach served to enhance predictive robustness and mitigate potential overfitting [13, 14]. For each data transformation technique, 8000 hybrid models were trained to ensure robustness. This number of models was chosen by manual inspection and sensitivity analysis as a tradeoff between the possibility of finding models with good performance and balancing computation times. Root Mean Squared Error (RMSE) for all variables relative to experimental data quantified model performance. For each hybrid model, RMSE reduction was calculated taking the full parametric Shinto model as baseline. RMSE reduction is defined in Eq. (4).

$$RMSE_{red} = \frac{RMSE_{Parametric\ Model} - RMSE_{Hybrid\ Model}}{RMSE_{Parametric\ Model}} \quad (4)$$

The optimal predictor ensemble size was the number of models that maximized the objective function defined in Eq. (5) using the predictor average.

$$max\ (Prediction_{Enhancement} - Overfitting) \quad (5)$$

$$Prediction_{Enhancement} = (RMSE_{red,train} + RMSE_{red,valid})$$

$$Overfitting = |RMSE_{red,train} - RMSE_{red,valid}|$$

Table 1. RMSE reduction on training and validation sets when compared against parametric model and number of optimal models for the ensemble predictor.

Data Transformation technique	Ensemble Model			Best individual Model	
	%RMSE reduction training	%RMSE reduction validation	Number of optimal models Ensembled	%RMSE reduction training	%RMSE reduction validation
No transformation	16.3	10.2	28	22.3	2.2
Min-Max normalization	18.2	10.4	8	20.8	4.5
Z-score normalization	16.5	10.1	13	21.6	6.6

This formulation balances two competing objectives: the first term maximizes the total error reduction across both training and validation sets, while the second term imposes a penalty for overfitting during calibration and for poor generalization performance.

Moreover, prediction envelopes and relative deviations of average values were used for approximating prediction uncertainty.

2.5 Computational Hardware Details

All computations were performed on an Intel® Core™ i9-14900HX CPU 2.20GHz, with 24 cores, 32 GB of RAM, and a 16 GB NVIDIA GeForce RTX 4090 GPU. Calculations were performed in MATLAB 2024a®.

3. RESULTS

Hybrid models' performance outperformed the parametric model baseline. The best individual model reduced RMSE by up to 21.6% (training) and 6.6% (validation), using z-score normalization, while ensemble models achieved reductions up to 18.2% and 10.4%, for training and validation respectively, using min-max normalization. Table 1 presents RMSE reduction for training and validation, as well as the number of optimal models used by the ensemble when using each data transformation technique.

There are differences in performance of the best individual models in the validation set, however, the ensemble approach increased generalization capabilities across the three techniques, leveling their performance. For ensemble models, performance across the different data transformation techniques was very similar for both training and validation. In training, no-transformation and z-score normalization yielded RMSE reductions of around 16.5%, with min-max normalization achieving higher performance, presenting reductions up to 18.2%. However, in validation there was no significant difference between the pretreatment techniques: 10.2% (no-transformation), 10.1% (z-score), and 10.4% (min-max).

Optimal ensemble size requirements varied, 28 models for no-transformation, 8 for min-max normalization, and 13 for z-score normalization. Training time was also comparable across methods, with 8000 models requiring 80–84 hours for calibration and less than 30 seconds for prediction in all cases. On the other hand, it's notable that out of 8000 calibrations, optimal ensemble sizes represent less than 1% of the total models calibrated. Moreover, when calibrating 4000 models, ensemble sizes varied between 4 and 9 (less than 1%). This lower number of models also deteriorated prediction capabilities, with RMSE reductions of the validation set being around 6-7%.

Figure 1 presents the behavior of the different ensemble predictors obtained using the three data transformation techniques. As expected from the results from Table 1, the behavior is very similar between them, which suggests that the ensemble approach is able to capture the dynamics of different models and minimize their differences while enhancing the predictive power. It is important to note that although butanol's reaction rate was hybridized, performance improved in other key states, such as glucose and biomass, since butanol concentration affects them. For these metabolites, individual RMSE reductions in the validation set ranged between 4% through 14%. Figure 2 presents the predictor confidence intervals, shown as envelopes of butanol prediction of the individual models used for the ensemble approach at each time-step. The hybrid model using data without transformation presents the widest predictor confidence intervals, due to having more than twice individual models when compared to z-score or more than triple when using min-max normalization. It is observed that confidence intervals are not symmetric, showcasing that the individual models' error is not distributed evenly. This might be attributed to both the non-parametric techniques as well as structural deficiencies in the parametric components of the model.

Moreover, when using no transformation, the ensemble envelope contains 94% of the experimental data, in contrast to 86% and 78% when using min-max and z-score normalization. This highlights that while their

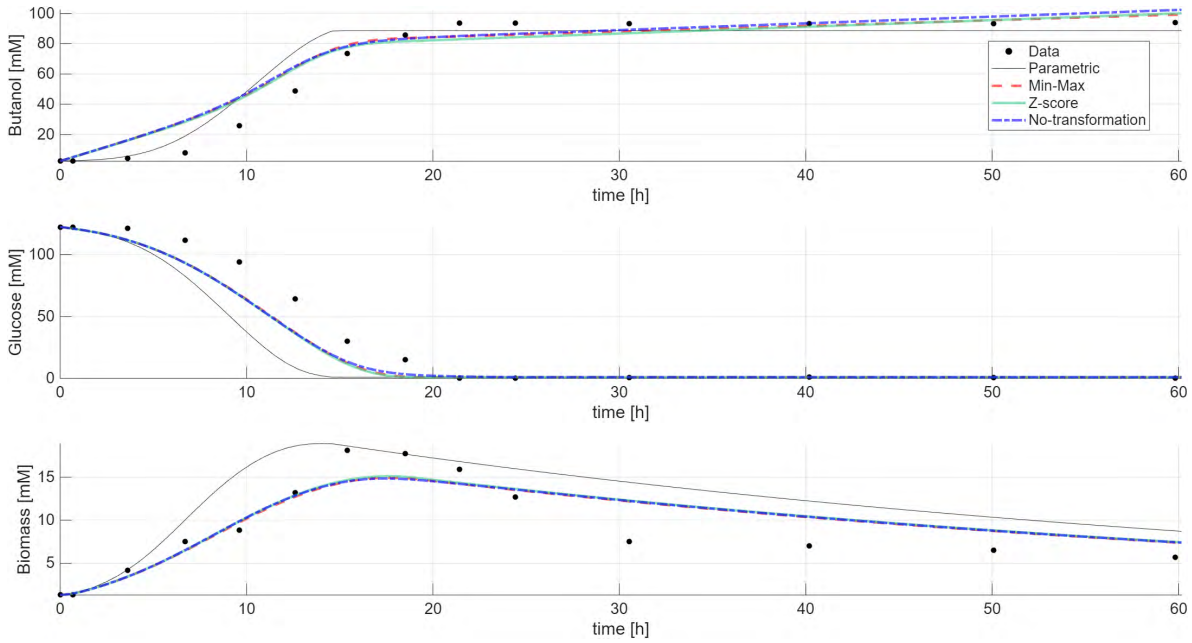

Figure 1. Ensemble hybrid models' behavior in the validation experiment according to each type of data transformation technique compared against real experimental data and the parametric model: a) butanol, b) glucose, c) biomass.

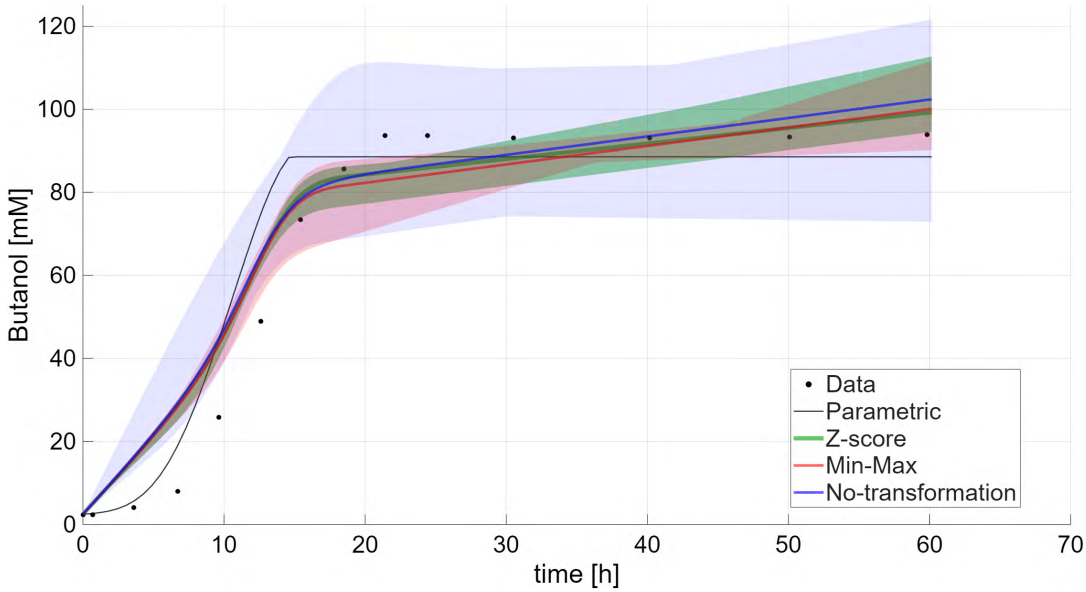

Figure 2. Hybrid models' prediction intervals for butanol in the validation set.

average predictions are similar, their uncertainty varies significantly. On the other hand, butanol's envelope of predictions also affects other variables, as seen in Figure 3 and Figure 4 for glucose and biomass in the validation set, respectively. As seen, prediction intervals follow the same trends as those for butanol.

Table 2 presents the average absolute and relative deviations using the different normalization techniques for the three metabolites. The absolute deviations remain relatively low for all metabolites, with min-max normalization and z-score normalization presenting deviations of less than 1 mM for glucose and biomass, and around 5 mM for butanol. These deviations are on average less than 5% of the predicted values, with the exception of biomass' deviation when using min-max normalization. On the other hand, using no data transformation results in higher absolute and relative deviations, being up to five times those obtained by the normalization techniques. This is in part due to ensemble size: when using no data transformation, the optimal ensemble is

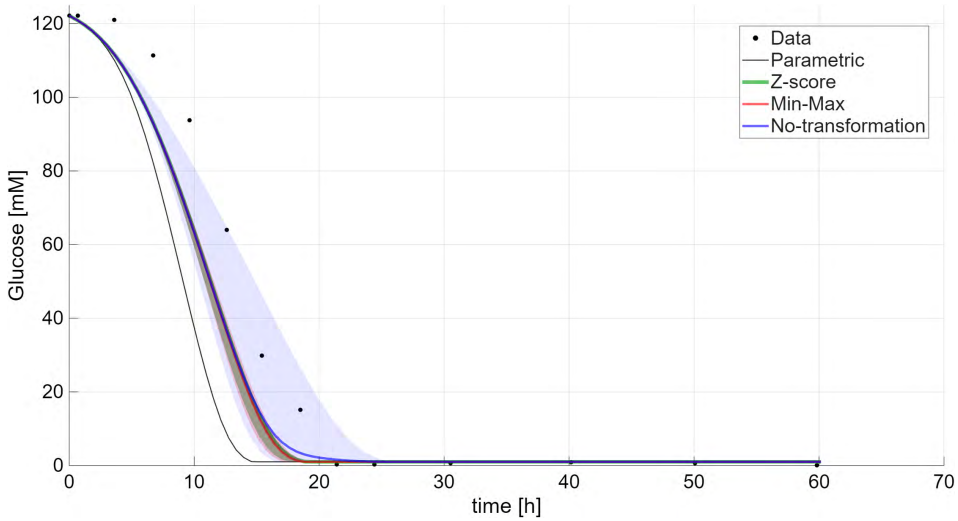

Figure 3. Hybrid models' prediction intervals for glucose in the validation set

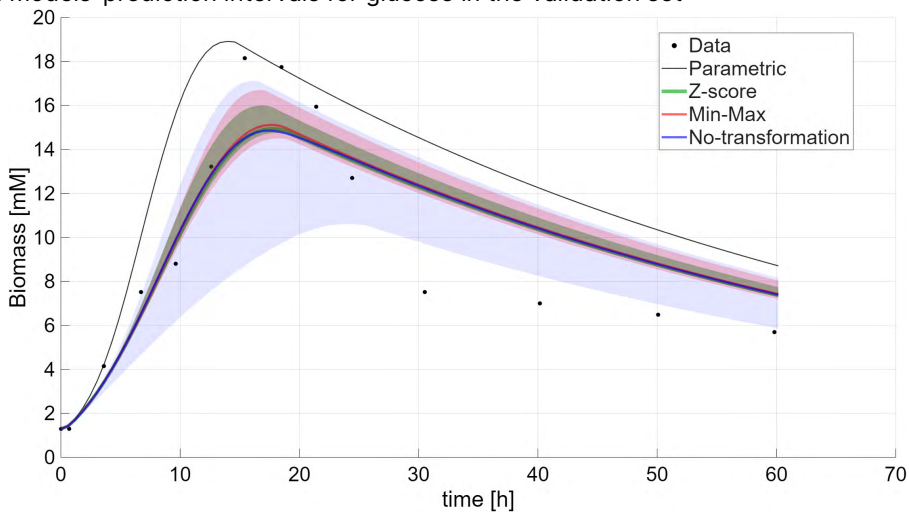

Figure 4. Hybrid models' prediction intervals for biomass in the validation set

Table 2. Average absolute and relative deviations (%) of the prediction envelopes from the average predicted value by the ensemble model

	Absolute deviation (mM)			Relative deviation (%)		
Metabolite	**Min-Max**	**Z-score**	**No-transformation**	**Min-Max**	**Z-score**	**No-transformation**
Glucose	0.674	0.535	3.736	5.85	4.03	24.80
Butanol	5.594	5.200	17.414	4.72	4.40	12.27
Biomass	0.620	0.367	1.901	5.67	4.03	24.84

twice as large as the ensemble obtained using z-score normalization and more than three times as the one obtained when using min-max.

However, as seen in Figure 2, Figure 3 and Figure 4, the prediction envelope is not symmetric, as such, the models present zones where they have different relative deviations. This asymmetry can be explained by the fact that GPs numeric errors don't propagate linearly towards the rest of the states. Figure 5 and Figure 6 present the relative deviation of the prediction envelope as a function of time in the validation experiment for biomass and butanol respectively, with the continuous lines representing the average values. Figure 5 shows that when using z-score and min-max normalization, for biomass model deviation is low, with maximum individual values around 10% deviation. In contrast, when using no-pretreatment, relative deviations are higher, up to +20% and -40%. Similarly, Figure 6 shows that when using no transformation, butanol's relative deviation can vary by as much as +60% and -20% in the early stages of the fermentation, with lower values towards the end (+20% through -20%). On the other hand, when using z-score

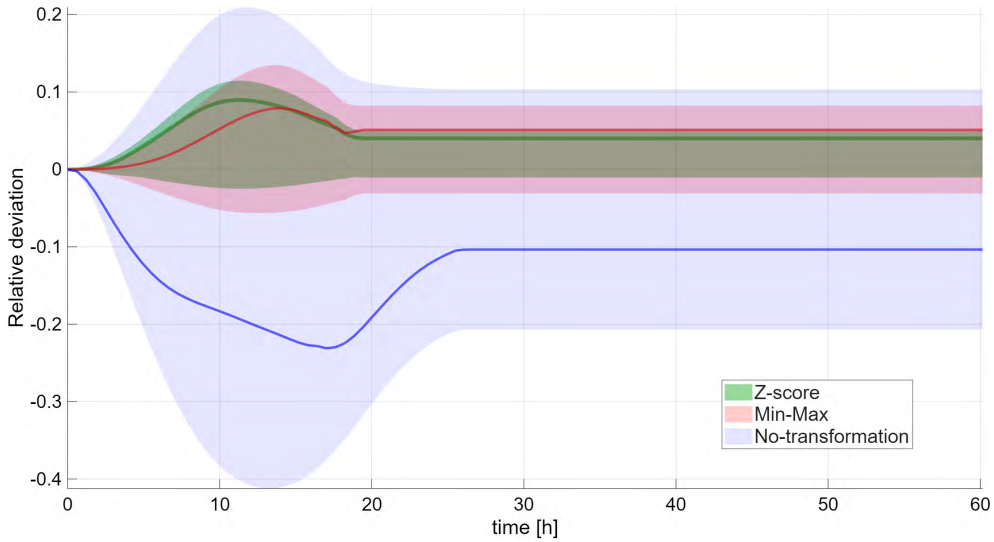

Figure 5. Hybrid models' prediction relative deviation for biomass in the validation set. Lines represent average prediction interval's width

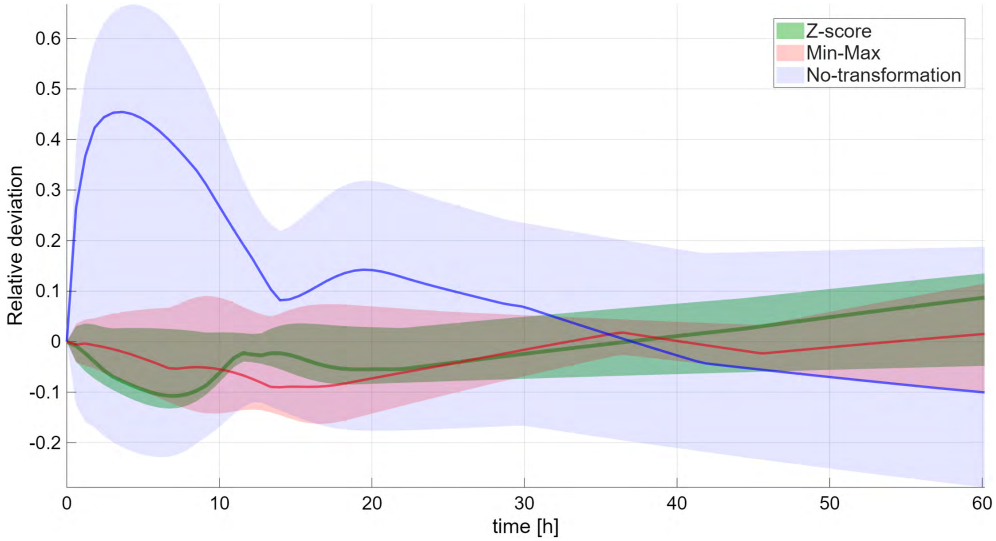

Figure 6. Hybrid models' prediction relative deviation for butanol in the validation set. Lines represent average prediction interval's width

and min-max normalization, predictions remain less dispersed, with relative deviations between -15% in the early stages of the fermentation and +15% towards the end. Higher deviations in early stages might be explained due to complexities for measuring metabolites and biomass at low concentrations, especially during lag phase and early exponential growth. Moreover, due to low values, high relative deviations might come from low absolute deviations, amplifying numerical error. As such, specific reaction rates change rapidly and error propagation might greatly affect prediction accuracy of the GPs, propagating uncertainty in a feedback loop.

Moreover, as observed in Figure 5 and Figure 6, hybridizing one reaction rate might enhance predictive capabilities, but also propagates uncertainty towards other metabolites and states in the system. As such, hybridization strategies should evaluate models as a whole in kinetic coupled systems, to enhance overall process understanding.

4. CONCLUSIONS

This study investigated the influence of data transformation techniques on the predictive robustness of hybrid models for ABE fermentation. The integration of Gaussian Processes into a mechanistic mass-balance framework successfully addressed the limitations of traditional kinetic models, specifically mitigating the issues of overparameterization and limited practical identifiability.

Main takeaways from this study are:

- **Ensemble-driven robustness:** while individual model performance was highly sensitive to data pretreatment—with z-score normalization yielding the most accurate single candidates, the implementation of an ensemble strategy effectively leveled these disparities. The ensemble approach achieved consistent validation RMSE reductions of approximately 10.1% to 10.4%, ggregardless of the normalization technique applied.
- **Prediction uncertainty vs. Coverage:** the analysis of prediction intervals (**Figure 5** and **Figure 6**) revealed trade-offs. Normalization techniques (z-score, min-max) provided superior narrower prediction intervals, particularly during the kinetically dynamic early fermentation stages, maintaining relative deviations within ±15%. In contrast, models trained on raw data exhibited significant volatility during the exponential phase (deviations up to +60%) but offered wider uncertainty envelopes that encapsulated 94% of experimental data, compared to ~80% for normalized models. Consequently, despite having similar performance indexes, uncertainty and coverage differ between data transformation techniques, making model selection more complex than just performance indexes.
- **Kinetic propagation:** although the hybridization specifically targeted the biobutanol formation rate, accuracy gains and conversely, uncertainty propagated to interconnected state variables such as glucose and biomass. This confirms that the hybrid framework functions as a cohesive dynamic system, where refining a single critical reaction rate improves the global fidelity of the bioprocess digital twin.

Finally, while normalization remains important for narrowing early-phase uncertainty intervals in individual predictors, the ensemble design serves as the primary driver of generalization and robustness. Future development of hybrid bioprocess models should therefore prioritize ensemble models over the exhaustive optimization of data scaling protocols.

ACKNOWLEDGEMENTS

Juan Federico Herrera-Ruiz received funding by the Colombian General Royalty System (Sistema General de Regalías de Colombia) and Ministry of Science, Technology and Innovation of Colombia (Ministerio de Ciencia, Tecnología e Innovación) with the projects BPIN: 2022000100035, and contract 112721-029-2025 for mobility in partnership with ECOS NORD.

This work was partially supported by the Ecos Nord Programme France-Amérique Latine (Colombie), reference ECOS Nord N° C25S01.

This work was partially supported by the French National Research Agency (ANR) under the "Investissements d'Avenir" program, reference ANR-18-EURE-0021.

AUTHOR IDENTIFIERS

Author ORCIDs:
Herrera-Ruiz Juan Federico: 0009-0007-3784-3860
Robles-Rodríguez Carlos Eduardo: 0000-0002-2436-3653
Aceves-Lara César Arturo: 0000-0001-6291-3655
Fontalvo Javier: 0000-0002-7868-8736
Prado-Rubio Oscar Andrés: 0000-0003-1419-542X

REFERENCES

1. Schweidtmann AM, Zhang D, von Stosch M. A review and perspective on hybrid modeling methodologies. Digital Chemical Engineering 10:100136 (2024). https://doi.org/10.1016/j.dche.2023.100136
2. Gargalo CL, Malanca AA, Aouichaoui ARN, Huusom JK, Gernaey KV. Navigating industry 4.0 and 5.0: the role of hybrid modelling in (bio)chemical engineering's digital transition. Front. Chem. Eng. 6: (2024). https://doi.org/10.3389/fceng.2024.1494244
3. de Azevedo CR, Díaz VG, Prado?Rubio OA, Willis MJ, Préat V, Oliveira R, von Stosch M. Hybrid semiparametric modeling: a modular process systems engineering approach for the integration of available knowledge sources. Systems Engineering in the Fourth Industrial Revolution :345-373 (2020). https://doi.org/10.1002/9781119513957.ch14
4. Maharana K, Mondal S, Nemade B. A review: data pre-processing and data augmentation techniques. Global Transitions Proceedings 3:91-99 (2022). https://doi.org/10.1016/j.gltp.2022.04.020
5. Singh D, Singh B. Investigating the impact of data normalization on classification performance. Applied Soft Computing 97:105524 (2020). https://doi.org/10.1016/j.asoc.2019.105524
6. Li J, Brenner M, Pierides I, Wessner B, Franzke B, Strasser EM, Waldherr S, Wagner KH, Weckwerth W. Machine learning and data-driven inverse modeling of metabolomics unveil key processes of active aging. npj Syst Biol Appl 11: (2025). https://doi.org/10.1038/s41540-025-00580-4
7. van den Berg RA, Hoefsloot HC, Westerhuis JA, Smilde AK, van der Werf MJ. Centering, scaling,

and transformations: improving the biological information content of metabolomics data. BMC Genomics 7: (2006). https://doi.org/10.1186/1471-2164-7-142
8. Craig A, Cloarec O, Holmes E, Nicholson JK, Lindon JC. Scaling and normalization effects in NMR spectroscopic metabonomic data sets. Anal. Chem. 78:2262-2267 (2006). https://doi.org/10.1021/ac0519312
9. Herrera?Ruiz JF, Fontalvo J, Prado?Rubio OA. Hybrid modeling for bioprocesses: architectures, applications, and perspectives. Engineering Reports 7: (2025). https://doi.org/10.1002/eng2.70502
10. Malik MAI, Usman M, Waqas Rafique M, Raza S, Saleem MW, Abbas N, Sajjad U, Hamid K, Rezaul Karim M, Abul Kalam M. Managing energy transition alongside environmental protection by making use of ai-led butanol powered SI engine optimization in compliance with sdgs. Heliyon 10:e29698 (2024). https://doi.org/10.1016/j.heliyon.2024.e29698
11. Shinto H, Tashiro Y, Yamashita M, Kobayashi G, Sekiguchi T, Hanai T, Kuriya Y, Okamoto M, Sonomoto K. Kinetic modeling and sensitivity analysis of acetone–butanol–ethanol production. Journal of Biotechnology 131:45-56 (2007). https://doi.org/10.1016/j.jbiotec.2007.05.005
12. Sánchez-Rendón JC, Matallana LG, Morales-Rodriguez R, Prado-Rubio OA. Enhanced kinetic model parameters for xylitol bioproduction from Candida mogii ATCC 18364, 2024, p. 2497–502. https://doi.org/10.1016/B978-0-443-28824-1.50417-8
13. Herrera-Ruiz JF, Fontalvo J, Prado-Rubio OA. Hybrid model development for succinic acid fermentation: relevance of ensemble learning for enhancing model prediction. Systems and Control Transactions 4:1896-1901 (2025). https://doi.org/10.69997/sct.153338
14. Pinheiro JMH, Oliveira SVB, Silva THS, Saraiva PAR, Souza EF, Godoy RV, Ambrosio LA, Becker M. The impact of feature scaling in machine learning: effects on regression and classification tasks. IEEE Access 13:199903-199931 (2025). https://doi.org/10.1109/access.2025.3635541

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Simulation and analysis of carbon capture process using piperazine for large scale biomass-fired power plant

Shengyuan Huang[a], Olajide Otitoju[a], Yao Zhang[a], Meihong Wang[a]*

[a] School of Chemical, Materials and Biological Engineering, The University of Sheffield, Sheffield S1 3JD, UK
* Corresponding Author: meihong.wang@sheffield.ac.uk.

ABSTRACT

Environmental concerns caused by CO_2 emissions has attracted much attention by researchers worldwide. CO_2 can be captured from large single sources such as power plants to reduce the CO_2 emission. Solvent-based post-combustion carbon capture (PCC) process for large scale biomass-fired power plant could achieve negative carbon emission. However, capture level is commonly set at 90% in many studies. The small fraction of residual CO_2 is still a large amount due to the high flue gas flowrate. In this study, a piperazine-based PCC process at 95% capture level for biomass-fired power plant was studied. The process was simulated in Aspen Plus® V11, validated and scaled up. The energy performance results showed that when the capture level is increased to 95%, the reboiler duty rises to 4.07 GJ/tCO_2, corresponding to an increase of approximately 13.7% compared to the 90% case. This additional regeneration energy demand is offset by the reduction in residual CO_2 emissions from flue gas or 0.23 million tons extra CO_2 captured each year (from 3.86 million tons CO_2/year to 4.09 million tons CO_2/year). It is feasible for improving negative emission performance in BECCS systems. More analysis (e.g. different configurations and economic analysis) will be performed to further discuss the high capture level process.

Keywords: carbon capture, chemical absorption, process simulation, technical assessment, negative emission technologies

1. INTRODUCTION

1.1 Background and motivation

Biomass energy with carbon capture and storage (BECCS) has been widely recognized as a negative emission technology in which renewable energy generation is combined with permanent removal of atmospheric carbon dioxide. BECCS is increasingly considered as an essential component in long-term carbon emission reducing strategies.

From process engineering, the effectiveness of BECCS is mainly dependent on the efficiency of the solvent-based PCC system. While a capture level of 90% is commonly adopted in PCC studies from the IEA report [1], the remaining CO_2 in the treated flue gas still represents a direct emission to the atmosphere. Reducing the residual CO_2 concentration through higher capture levels becomes increasingly relevant for BECCS applications that aim to achieve meaningful negative emissions.

1.2 Literature review

The solvent-based PCC is widely used from Wang et al. [2]. The process is based on chemical absorption to capture CO_2 from the flue gas. In most system studies, a capture level of 90% is typically adopted as the default design point. It avoids the steep energy penalty and equipment scale-up associated with very high capture level. In Canepa et al. [3], MEA-based PCC process for combined cycle gas turbine (CCGT) power plant was studied under 90% capture level. Exhaust gas recirculation (EGR) technology was proposed to reduce the flue gas flowrate to solvent-based PCC plant, therefore reducing sizes of key equipment (mainly absorber) energy consumption for solvent regeneration. Luo and Wang [4] analysed the solvent-based PCC for nature gas combined cycle (NGCC) power plant. They found that carbon capture becomes economically justified only when the carbon price exceeds about €100–120 per ton CO_2. Higher fuel and CO_2 transport costs will further increase the required carbon price to maintain high capture levels. Hu et al. [5] applied the PCC using MEA for treating flue

gas from thermal cracking furnace in ethylene manufacturing under the same capture level. Through integrating excess process heat can supply regeneration energy and reduce CO_2 capture cost from about \$80/tCO_2 to \$70/tCO_2

Study of Michailos and Gibbins [6] explored solvent-based PCC process to facilitate 95–99% capture level from CCGT flue gas. The results show that there are significant differences in performance and energy consumption under different capture levels. The increasing capture level above 90% is technically possible but often comes with higher costs and higher energy consumption used, especially for post-combustion systems. Many studies historically capture level around 85–90%, leaving residual emissions uncaptured.

BECCS studies commonly treat solvent-based PCC as a key enabling block. IEAGHG has discussed higher capture rates for biomass as a method for lower residual emissions, while also noting that solvent-based PCC can face a high cost and energy penalty as capture level increases beyond typical benchmarks. More recent work has started to examine biomass-fired power plant with PCC configurations and solvent choices in more detail. For example, Ren et al. [7] investigated techno-economic performance and optimisation of piperazine-based PCC and CO_2 compression for large-scale supercritical biomass-fired power plants, discussing solvent choice and configuration effects specifically for biomass contexts.

1.3 Novelty of this study

Several studies on the solvent-based PCC process under 90% capture level has been discussed in the literature review section. For large-scale biomass-fired power plants, even a relatively small fraction of residual CO_2 can translate into a substantial emission due to the high flue gas flow rate. For the flue gas from a 550 MW biomass-fired power plant, increasing the capture level from 90% to 95% raises the daily amount of captured CO_2 nearly 700 tonnes. The capture cost of DAC commonly exceeds USD 600 per tonne. This additional captured amount shows great potential. Within this context, a capture level of 95% sets as the critical engineering setting. It could reduce the residual CO_2 emissions compared with the widely adopted 90% level, while remaining within a range that may still be technically achievable using advanced solvents and well-designed solvent-based PCC systems. The energy performance of solvent-based PCC process operating at such high capture levels has not been systematically assessed for biomass-fired power plants.

2. PROCESS DESCRIPTION

In this study, a standard solvent-based PCC process is considered for application to a large-scale biomass-fired power plant. The process configuration follows a conventional closed-loop absorption–regeneration system using piperazine (PZ) as the reactive solvent and is representative of industrial PCC designs targeting high CO_2 capture level. The flowsheet diagram of solvent-based PCC is shown in the Figure 1.

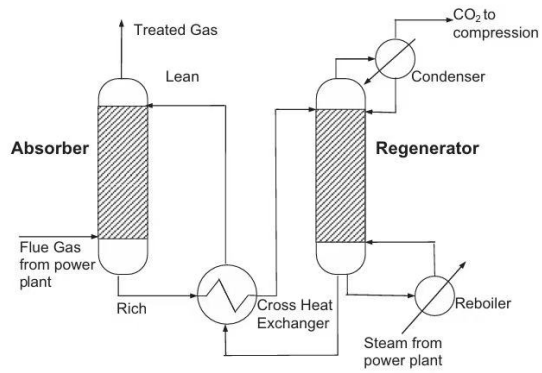

Figure 1. The process of standard PCC process from Lawal et al. [8]

The overall process consists of an absorber, a stripper equipped with a reboiler, solvent circulation pumps, and a set of heat exchangers for temperature control and heat recovery. Flue gas exiting the biomass-fired power plant is introduced at the bottom of the absorber. As Lawal et al. [8] and Otitoju et al. [9] mentioned, the gas flows upward through the packed bed, where it comes into counter-current contact with lean PZ solvent introduced at the top of the column. CO_2 is removed from the flue gas through combined gas–liquid mass transfer and chemical reaction with the solvent, resulting in a treated gas stream with reduced CO_2 concentration leaving the top of the absorber.

The CO_2-rich solvent collected at the bottom of the absorber is pumped to the regeneration section. Steam generated in the reboiler raises the temperature of the solvent, reversing the absorption reactions and releasing CO_2 from the liquid phase. The regenerated lean solvent exits the bottom of the stripper and is routed back to the absorber after cooling.

The solvent used in this study is 30 wt% PZ. Relative to MEA, PZ exhibits higher CO_2 absorption capacity and faster reaction kinetics, allowing operation at higher rich loadings. PZ solvent has better performance when targeting high capture levels, as they reduce the extent of solvent circulation and absorber height required to achieve a given removal target.

3. MODEL DEVELOPMENT, MODEL VALIDATION AND SCALE-UP

3.1 Model development

In this study, the DAC process using PZ was developed in Aspen Plus ® V11. The absorber and stripper are all rate-based models using RadFrac Block in the software. The reactions and physical properties used for rate-based model was the same as Otitoju et al. [9]. The details are not shown here due to page limit.

3.2 Model validation

The developed PCC model was validated by pilot-scale experimental data. The experimental cases were chosen from the pilot plant experimental results in the University of Texas at Austin from Otitoju et al. [9]. The flue gas conditions were stabled as 0.165 m^3/s and 12 vol% of CO_2 in flue gas. The pilot plant consists of a packed-bed absorber and stripper equipped with structured packing. The absorber was designed as packed columns with a diameter of 0.427 m and a total height of 6.1 m. The column contained two 3.05 m sections of Mellapak 2X packing. In the present study, Mellapak 2X packing was selected in the simulation [7].

Model validation used five experimental cases operated under different conditions. These cases were selected with different operating parameters, including variations in liquid-to-gas (L/G) ratio, lean solvent loading, and solvent circulation rate from Dugas [10]. The CO_2 capture level was selected as the primary output variable for validation. For each experimental case, the simulated capture level was compared with the corresponding experimental value, and the relative error was calculated. As shown in the Figure 2, across the five validation cases, the relative error between model prediction results and experimental data remained within ±10%. This level of agreement indicates that the rate-based model can simulate the CO_2 absorption process with reliability. The validation results therefore support the use of the developed model for subsequent scale-up and high capture level analysis.

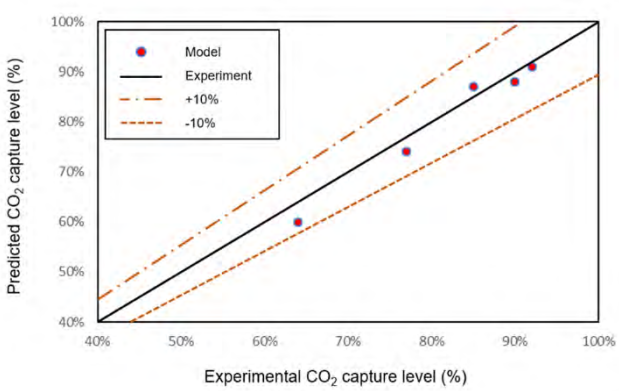

Figure 2. Experimental data versus model prediction for CO_2 capture level for the 5 experimental cases

3.3 Model scale-up

Following model validation at pilot scale, the PCC model was scaled up to deal with the flue gas from a large biomass-fired power plant with an electrical output of 550 MW. The solvent composition was maintained at 30 wt% PZ. The scale-up of the absorber and stripper was performed using established packed-column scale up design principles. The absorber diameter was determined based on flue gas flowrate as 647.5 kg/s and hydrodynamic constraints using the Generalized Pressure Drop Correlation (GPDC) [9]. This approach makes sure that column operation remains below flooding limits while maintaining acceptable pressure drop under large flue gas flowrates from biomass-fired power plants.

Table 1 presents the dimensions of the absorber and stripper calculated using the GPDC method and compares the cases operating at 90% and 95% capture levels. As shown in the Table 1, increasing the capture level requires the absorber to remove a larger fraction of CO_2 from the flue gas. To achieve this, a higher solvent circulation rate and a longer effective contact time between the gas and liquid phases are required, which leads to an increase in absorber size compared with the 90% capture case. The higher solvent flow rate also results in an increased regeneration demand, which in turn increases the required size of the stripper.

Table 1: Comparison of 90% capture level scale up simulation results with 95%

Scale up results	90% capture level	95% capture level
Absorbers		
Diameter (m)	10.4	11.9
Packing height (m)	25	28
Stripper		
Diameter (m)	8	10
Packing height (m)	20	20

4. SENSITIVITY ANALYSIS

The influence of lean solvent loading on the energy performance of the PZ-based PCC process was analysed using the validated standard PCC model. Two capture levels, 90% and 95%, were considered in order to examine how the operating strategy affects reboiler duty under different removal targets. For each capture level, the lean loading was varied with same specified overall CO_2 capture level.

As shown in Figure 3, the reboiler duty exhibits the relationship with lean loading for both capture levels. At relatively low lean loading, deeper solvent regeneration is required to maintain the target CO_2 removal, which leads to a higher reboiler duty. As lean loading increases, the regeneration requirement is gradually relaxed, resulting in a decrease in reboiler duty. However, beyond a certain lean loading range, further increases in lean loading cause the reboiler duty to rise again. This trend can be

explained by the corresponding change in solvent circulation rate. Figure 4 shows that the solvent flow rate increases continuously with increasing lean loading for both capture levels. At higher lean loading, a larger solvent circulation rate is required for same CO_2 removal efficiency, as the driving force for absorption in the absorber is reduced. The increased circulation rate leads to higher sensible and latent heat requirements in the regeneration section, which offsets the benefit of reduced regeneration depth and results in an overall increase in reboiler duty.

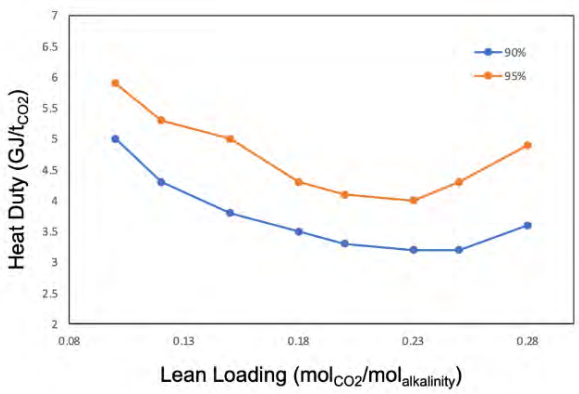

Figure 3. Energy consumption at various lean loadings for different capture level (%)

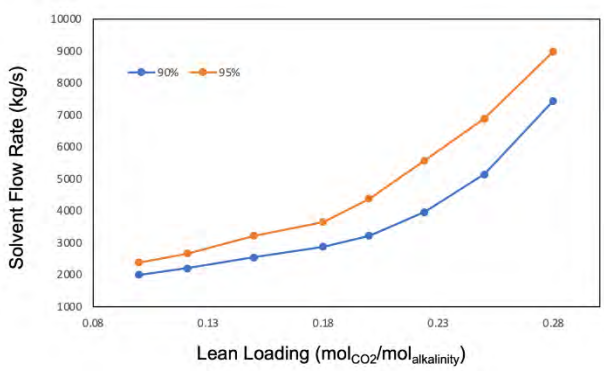

Figure 4. Solvent flow rate at various lean loading for different capture level (%)

For the 90% capture case, the minimum reboiler duty of approximately 3.5 GJ/tCO_2 is achieved at a lean loading of around 0.22–0.24 mol CO_2/mol alkalinity. A similar trend is observed for the 95% capture case, although the reboiler duty becomes higher across the entire lean loading range due to the higher regeneration requirement. The minimum reboiler duty for the 95% capture case is approximately 4.0 GJ/tCO_2, occurring at a similar lean loading range.

Comparing these two capture levels, the results showed that increasing the capture level from 90% to 95% systematically shifts the energy performance curve upward. This is attributed to the combined effects of deeper solvent regeneration and higher solvent circulation rates required at the higher capture level.

Overall, the sensitivity analysis shows that lean loading is important in determining the regeneration energy demand of the PCC process using PZ. An optimal lean loading range exists where the trade-off between regeneration depth and solvent circulation is balanced. Although operating at 95% capture inevitably increases the reboiler duty, the additional energy requirement remains moderate within the identified optimal lean loading range, supporting the feasibility of high capture operation in BECCS applications.

5. RESULTS FROM ENERGY ASSESSMENT

The energy performance of the PZ-based PCC process was evaluated by comparing operation at 90% and 95% CO_2 capture levels. The specific reboiler duty, expressed in GJ/tCO_2, was used as the primary performance indicator, as solvent regeneration represents the dominant energy consumption in the process.

Figure 5 presents the variation of reboiler duty with capture level. At a capture level of 90%, the reboiler duty is 3.58 GJ/tCO_2 [7]. When the capture level is increased to 95%, the reboiler duty rises to 4.07 GJ/tCO_2, corresponding to an increase of approximately 13.7%. This result confirms that achieving higher CO_2 removal requires additional thermal energy for solvent regeneration.

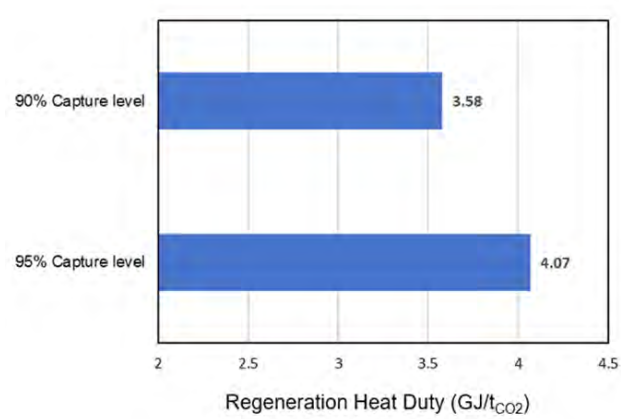

Figure 5. Regeneration heat duty of different configurations of PCC processes using PZ

From a process perspective, the increase in reboiler duty can be explained by the combined effects of deeper solvent regeneration and changes in column operation. To achieve a higher capture level, the lean solvent leaving the stripper must reach a lower CO_2 loading, which directly increases the thermal requirement in the reboiler. At the same time, the absorber must provide sufficient contact time and solvent capacity to remove a larger

fraction of CO_2 from the flue gas. This is reflected in the increased absorber dimensions and solvent circulation rates observed for the 95% capture case compared with the 90% case. The higher solvent circulation rate required at 95% capture also affects the regeneration section. An increased liquid flow to the stripper leads to a larger stripper size and higher sensible and latent heat demands during regeneration. These effects together contribute to the observed increase in reboiler duty. The results showed that higher capture efficiency can be achieved, but with the increasing cost of regeneration energy demand.

Despite this increase, the additional energy requirement is still acceptable. For the flue gas from the 550MW biomass-fired power plant as 647.5 kg/s, the increasing capture level as 5% could capture additional 7.55kg/s of CO_2. It results in an annual captured amount over 0.23 Mt. In other words, at 90% capture level, such a solvent-based PCC would capture 3.86 Mt CO_2 each year. At 95% capture level, such a solvent-based PCC would capture 4.09 Mt CO_2 each year. An increase of 13.7% in reboiler duty represents a relatively small penalty when considered against the substantial reduction in residual CO_2 emissions achieved by moving from 90% to 95% capture. In the context of BECCS, where the objective is to maximise net negative emissions, reducing residual emissions from the treated flue gas is particularly important. Existing biomass-fired power plants and PCC infrastructure do not require complete reconstruction. This indicates that the higher capture level does not introduce impractical equipment requirements under the operating assumptions considered. Consequently, the additional regeneration energy associated with 95% capture can be regarded as acceptable for BECCS applications, provided that the objective is to minimise residual emissions and enhance negative emission performance.

6. CONCLUSION AND RECOMMENDATIONS FOR FUTURE WORK

Operating a PZ-based PCC process at a 95% capture level is technically feasible and does not incur a prohibitive energy penalty relative to the commonly used 90% benchmark. The additional captured CO_2 by setting at 95% capture level is about 653 tonnes per day, corresponding to an annual increase of more than 0.23 Mt (from 3.86 Mt CO_2 to 4.09 Mt CO_2). The reboiler duty rises to 4.07 GJ/tCO_2, corresponding to an increase of approximately 13.7% than 90% case. These results provide the necessary research foundation for the subsequent integration of the BECCS system with direct air capture technology. This study will first conduct further configurations, such as utilizing advanced flash stripper (AFS), intercooling system and rich-split process. This study will also perform economic optimizations for the system. Optimal design and operation for high capture level biomass-fired power plant process using PZ can be achieved.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the financial support of the EU RISE project OPTIMAL (Grant Agreement No: 101007963) and CATALYSE (Grant Agreement No: 101183092).

REFERENCES

1. International Energy Agency (2021). IEA World energy outlook 2021 [Online]. https://www.iea.org/reports/world-energy-outlook-2021
2. Wang M, Lawal A, Stephenson P, Sidders J, Ramshaw C. Post-combustion CO2 capture with chemical absorption: a state-of-the-art review. Chemical Engineering Research and Design 89:1609-1624 (2011). https://doi.org/10.1016/j.cherd.2010.11.005
3. Canepa R, Wang M, Biliyok C, Satta A. Thermodynamic analysis of combined cycle gas turbine power plant with post-combustion CO ₂ capture and exhaust gas recirculation. Proceedings of the Institution of Mechanical Engineers, Part E: Journal of Process Mechanical Engineering 227:89-105 (2012). https://doi.org/10.1177/0954408912469165
4. Luo X, Wang M. Optimal operation of mea-based post-combustion carbon capture for natural gas combined cycle power plants under different market conditions. International Journal of Greenhouse Gas Control 48:312-320 (2016). https://doi.org/10.1016/j.ijggc.2015.11.014
5. Hu G, Li X, Liu X, Hu J, Otitoju O, Wang M, Du W, Ye Z, Long J, Qian F. Techno-economic evaluation of post-combustion carbon capture based on chemical absorption for the thermal cracking furnace in ethylene manufacturing. Fuel 331:125604 (2023). https://doi.org/10.1016/j.fuel.2022.125604
6. Michailos S, Gibbins J. A modelling study of post-combustion capture plant process conditions to facilitate 95–99% CO2 capture levels from gas turbine flue gases. Front. Energy Res. 10: (2022). https://doi.org/10.3389/fenrg.2022.866838
7. Ren J, Otitoju O, Gao H, Liang Z, Wang M. Techno-economic analysis and optimisation of piperazine-based post-combustion carbon capture and CO2 compression process for large-scale biomass-fired power plants through simulation. Fuel 381:133340

(2025). https://doi.org/10.1016/j.fuel.2024.133340
8. Lawal A, Wang M, Stephenson P, Koumpouras G, Yeung H. Dynamic modelling and analysis of post-combustion CO2 chemical absorption process for coal-fired power plants. Fuel 89:2791-2801 (2010). https://doi.org/10.1016/j.fuel.2010.05.030
9. Otitoju O, Oko E, Wang M. Technical and economic performance assessment of post-combustion carbon capture using piperazine for large scale natural gas combined cycle power plants through process simulation. Applied Energy 292:116893 (2021). https://doi.org/10.1016/j.apenergy.2021.116893
10. Dugas ER. Pilot plant study of carbon dioxide capture by aqueous MEA. M.S.E. Thesis, University of Texas, Austin, USA, 2006.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Differentiable Programming for Cyclic Adsorption Processes

Alex Glover[ab], Maria M. Papathanasiou[ab], and Ronny Pini[ab]*

[a] Department of Chemical Engineering, Imperial College London
[b] The Sargent Centre for Process Systems Engineering, Imperial College London
* Corresponding Author: r.pini@imperial.ac.uk.

ABSTRACT

The design of cyclic adsorption processes is computationally expensive as it involves screening many process designs, each of which involve a time-consuming simulation to reach cyclic steady state. In this work, we demonstrate how differentiable programming can be used to accelerate both the simulation and attainment of cyclic steady state for a four-step pressure vacuum swing adsorption (PVSA) process to concentrate carbon dioxide from flue gas. A mechanistic one-dimensional dynamic adsorption model was implemented in JAX, enabling automatic differentiation and just-in-time compilation for efficient solution and accurate sensitivity evaluation. The latter was exploited to implement a Newton-based direct determination method for accelerated convergence to cyclic steady state, avoiding repeated cycle simulations. Across 4096 designs sampled in a six-dimensional design space, the direct determination method converged in an average of 4.6 iterations, compared to 145 cycles required by successive substitution. When combined with the computational gains from the JAX framework, this resulted in an overall speed-up of over 20 times relative to the conventional MATLAB-based implementation.

Keywords: Differentiable Programming, JAX, cyclic adsorption processes, mechanistic modelling, gradient-enhanced acceleration

INTRODUCTION

Cyclic adsorption processes are used throughout a range of industrial settings for gas separations and have seen an increase in interest for their use in carbon capture applications. In these processes, a solid adsorbent is used to selectively adsorb components of a gas stream, due to differences in adsorption capacity. The adsorbent is regenerated through changing the column conditions such as with a decrease in pressure and/or an increase in temperature. Thus, adsorption processes consist of a sequence of distinct steps, which repeat in a cyclic manner.

Computational process models are a powerful and widely used tool for screening potential designs, where operating parameters such as the pressure set points, or duration of the steps are varied to optimise the process performance [1]. Such models describe the dynamic behaviour of these processes through systems of partial differential equations (PDEs) which represent mass, energy and momentum balances. Despite the inherently dynamic nature of adsorption processes, they are typically designed to operate at cyclic steady state (CSS): an operating point where one cycle behaves identically to the next. Therefore, determination of this CSS is a key part of evaluating every candidate process design. Typical methods involve simulating the process for multiple cycles, from startup conditions, until CSS is reached. Depending on the process and operating conditions, this procedure may require tens, hundreds or thousands of cycles to converge [2, 3]. In isolation, this procedure often requires just a few minutes of CPU time to evaluate the performance of the design at CSS. However, process design and optimisation studies typically evaluate thousands of designs due to having several operational parameters to evaluate simultaneously. Thus, for one process, the design may require hundreds or thousands of CPU hours [4, 5]. The computational demand becomes even larger when different cycle configurations or adsorbents are compared, or if uncertainty quantification is required, as the total number of high-fidelity simulations increases substantially, leading to a large cumulative computational cost even when parallel resources are

available, limiting the use of high-fidelity process models. As such, many authors have turned to short-cut models, where axial gradients are neglected[6, 7] or trained machine learning surrogates using data from high-fidelity process models [8, 9].

The conventional approach to reach CSS described above relies on successive substitution, in which the final state of one cycle is used as the initial condition for the next, which benefits from simulating the process dynamics. However, there have been several methods developed within the literature to accelerate the determination of CSS. Gradient-based methods, where the Jacobian matrix of the output state of the cycle with respect to the input state is used in a Newton iteration, have been used to converge to the CSS in fewer iterations [3, 10–12]. However, the evaluation of these gradients, typically through solving the forward sensitivity or adjoint ordinary differential equations [13, 14], require an added computational requirement at each iteration. Quasi-Newton methods, which do not require a full Jacobian evaluation at each iteration have been used to reduce this demand [10, 15]. Gradient-free approaches, such as algebraic methods introduced by Subraveti et al. [2], where the next iteration is based on the three preceding inputs have been used for the fixed point iteration. Machine learning models have also been trained to predict CSS [16]. Rather than the conventional method of solving the PDE system using the method of lines and integrating the resulting ordinary differential equations (ODEs), the full discretization of the PDE system (in both the spatial and time domain) reduces the system of equations to a set of algebraic equations, which can be combined with the CSS criterion to avoid repeated simulation [17]. However, authors have reported convergence difficulty with this method [12]. To address this issue, Tsay et al. developed a pseudo-transient model based on the fully discretized system to simulate and optimize cyclic processes[18].

Differentiable programming offers a promising alternative for addressing the above challenges by enabling the efficient computation of exact gradients of model outputs with respect to inputs. Recent advances in differentiable frameworks, such as Google's JAX and MIT's Julia, have made automatic differentiation techniques accessible to develop end-to-end differentiable mechanistic models. Such models have been applied to crystallisation processes, computational fluid dynamics, chemical reaction networks and molecular dynamics simulations[19–22].

In this work, we exploit these capabilities to develop a differentiable model of a pressure vacuum swing adsorption (PVSA) process for carbon capture. Due to the high computational effort required for the design and optimisation from the repeated simulations to obtain CSS, this study explores the use of JAX to accelerate process simulations. The first method for the acceleration of the solution of the ODE system is using just-in-time compilation and automatic differentiation within the ODE solver. By leveraging exact sensitivities, the second method implements a Newton-based direct determination method for rapid convergence to CSS [3], through a discrete sensitivity method.

METHODOLOGY

Mechanistic Adsorption Model

The case study for this work is a four-step pressure vacuum swing adsorption process for the separation of a 15% carbon dioxide in nitrogen feed, using Zeolite 13X as the adsorbent (Figure 1). This case study has been chosen for this work as it has been a system of interest throughout many optimisation studies, and due to its relevancy to post-combustion carbon capture applications [5, 23–25]. The process model describes the coupled heat and mass transfer in the adsorption bed through a system of partial differential and algebraic equations. The following assumptions are made in deriving the model equations:

- Only spatial gradients in the axial direction of the column are considered.
- The gas phase is ideal.
- Thermal equilibrium between the solid, adsorbed and gas phases.
- Constant heat capacities of all phases.
- Constant density of solid phase and the column wall.
- Constant column void fraction and adsorbent particle size.
- Constant mass dispersion coefficient and effective thermal conductivity.
- The pressure drop along the column is described by Darcy's law.

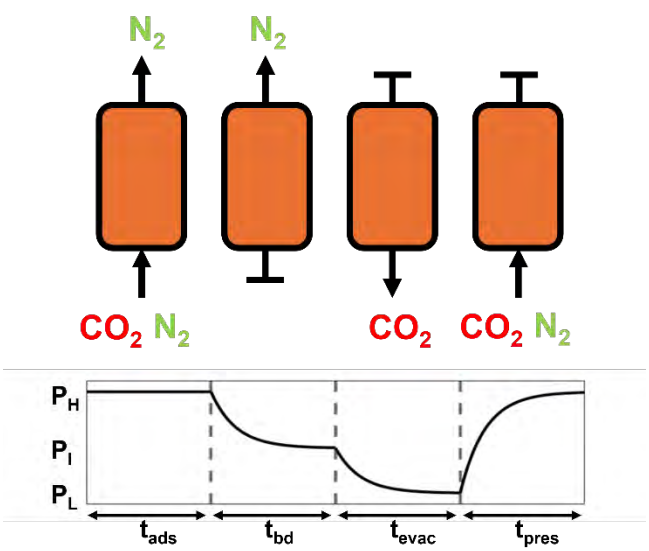

Figure 1. Schematic of the four step PVSA cycle. The four

Table 1: Adsorption model equations

Overall Material Balance

$$\frac{\partial p}{\partial t} - \frac{p}{T}\frac{\partial T}{\partial t} = -T\frac{\partial}{\partial z}\left(\frac{pv}{T}\right) - RT\frac{\rho_b}{\epsilon}\sum_{i=1}^{n_c}\frac{\partial q_i}{\partial t}$$

Component Material Balance

$$\frac{\partial y_i}{\partial t} = \frac{T}{p}D_L\frac{\partial}{\partial z}\left(\frac{p}{T}\frac{\partial y_i}{\partial z}\right) - \frac{T}{p}\left(\frac{\partial}{\partial z}\left(\frac{y_i pv}{T}\right) - y_i\frac{\partial}{\partial z}\left(\frac{pv}{T}\right)\right) - \frac{RT}{p}\frac{\rho_b}{\epsilon}\left(y_i\sum_{i=1}^{n_c}\frac{\partial q_i}{\partial t} - \frac{\partial q_i}{\partial t}\right)$$

Column Energy Balance

$$\left(C_{p,g}\frac{P}{RT} + \frac{\rho_b}{\epsilon}\left(C_{p,s} + C_{p,a}\sum_{i=1}^{n_c}q_i\right)\right)\frac{\partial T}{\partial t} = \frac{K_z}{\epsilon}\frac{\partial^2 T}{\partial z^2} - \frac{C_{p,g}}{R}\left(\frac{\partial}{\partial z}(pv) - T\frac{\partial}{\partial z}\left(\frac{pv}{T}\right)\right) - \frac{\rho_b}{\epsilon}\sum_{i=1}^{n_c}\left(T\left(C_{p,a} - C_{p,g}\right) - (-\Delta H_i)\right)\frac{\partial q_i}{\partial t} - \frac{2h_{in}}{\epsilon r_{in}}(T - T_w)$$

Column Wall Energy Balance

$$\rho_w C_{p,w}\frac{\partial T_w}{\partial t} = K_w\frac{\partial^2 T_w}{\partial z^2} + \frac{2r_{in}h_{in}}{r_{out}^2 - r_{in}^2}(T - T_w) - \frac{2r_{out}h_{out}}{r_{out}^2 - r_{in}^2}(T_w - T_a)$$

Steady State Momentum Balance (Darcy's Law)

$$-\frac{\partial P}{\partial z} = \frac{150}{4r_p^2}\left(\frac{1-\epsilon}{\epsilon}\right)^2\mu v$$

Linear Driving Force

$$\frac{\partial q_i}{\partial t} = k_i(q_i^* - q_i)$$

Adsorption Equilibrium (Dual Site Langmuir)

$$q_i^* = \frac{q_{b,i}b_i c_i}{1 + \sum_{i=1}^{n_c} b_i c_i} + \frac{q_{d,i}d_i c_i}{1 + \sum_{i=1}^{n_c} d_i c_i}$$

$$b_i = b_{i,0}e^{-\frac{U_{b,i}}{RT}}$$

$$d_i = d_{i,0}e^{-\frac{U_{d,i}}{RT}}$$

steps include adsorption, forward blowdown, reverse evacuation and forward pressurisation

The one-dimensional PDE model which describes the heat and mass balances within the column is illustrated in Table 1. Auxiliary equations include the adsorption isotherms which describe the thermodynamic equilibrium between the gas and the adsorbed phase. For the Zeolite 13X adsorbent used in this work, the extended dual-site Langmuir model is used to quantify this relationship for the carbon dioxide and nitrogen mixture. The parameters used for this model are the same as those reported in Ward and Pini [5]. For each step of the four-step cycle, the PDE system is solved with the boundary conditions relevant to the step, as described in [5]. The diagram in Figure 1 illustrates this 4-step cycle with the pressure profile given below. The cycle begins with adsorption at the high pressure set point. Blowdown to the intermediate pressure set point removes most of the adsorbed nitrogen. The carbon dioxide product is extracted in the evacuation step before the column is repressurized with the feed.

Numerical Solution and Sensitivities

A finite volume discretization (10 finite volumes) of the spatial domain of PDE system was used to reduce the system to a set of ODEs that were then solved sequentially for each step of the cycle. Due to stiffness of the

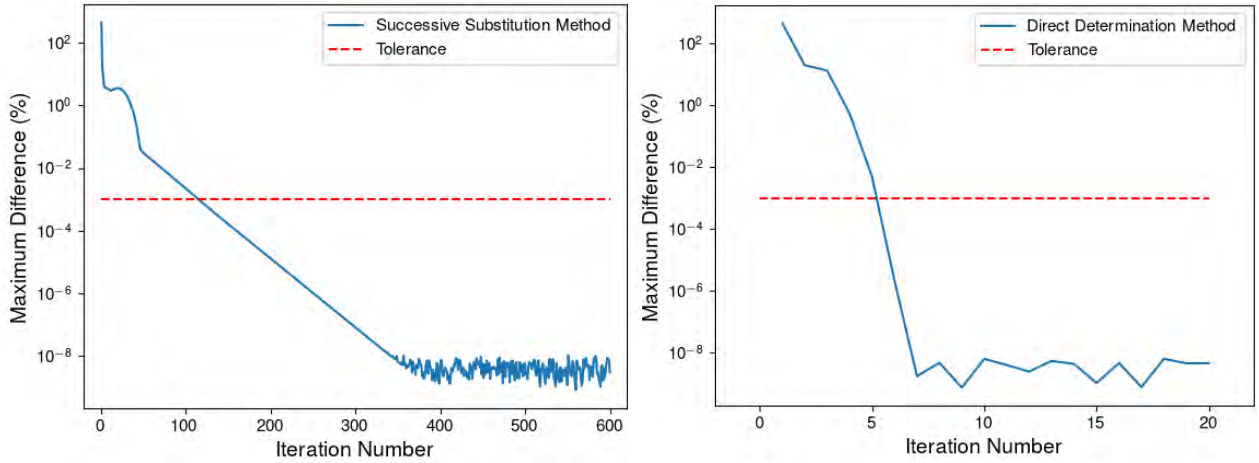

Figure 2. Convergence to cyclic steady state for successive substitution (left) and direct determination (right) methods.

ODEs, arising from diffusion, heat conduction, and adsorption source terms, we emplan implicit time integration scheme. Specifically, a numerical differentiation formula (NDF) method similar to MATLAB's ode15s solver was implemented in JAX [26]. This implementation uses automatic differentiation (AD) to efficiently evaluate the Jacobian of the state gradients with respect to the state variables: a key part of the implicit solver.

Sensitivities of the system states with respect to the initial conditions are required for the direct determination of CSS. Rather than using continuous sensitivity equations, a discrete forward sensitivity method was used, in which sensitivities are obtained by differentiating the numerical integration scheme itself. This ensures that the computed gradients are consistent with the discrete trajectory produced by the time integrator and enables the reuse of Jacobian information already computed by the implicit solve. The use of automatic differentiation within the JAX differentiable framework makes the evaluation of sensitivities computationally effective for the use of the direct determination method for reaching CSS in adsorption processes. All the simulations were performed on an Intel Core i5-10500 CPU with 8GB of RAM.

Direct Determination of Cyclic Steady State

The CSS is defined as a point where each cycle operates identically to the next. In mathematical models this can be monitored through the convergence of the overall mass balance, the key performance indicators, or the values of the state variables themselves. In this work, we define the cycle error as the difference between the vector of state variables at the start and end of a cycle (Equation 1). The state variable vector $\boldsymbol{x}$ includes the pressure, column temperature, wall temperature, gas phase mole fraction and adsorbed quantity at every finite volume element. The subscript k represents the iteration number, with $\boldsymbol{x}_k^0$ denoting the initial state of the cycle and $\boldsymbol{x}_k$ the final state.

$$\boldsymbol{e}_k = \boldsymbol{x}_k - \boldsymbol{x}_k^0 \tag{1}$$

The CSS is reached when the max norm (also known as the L^∞) of this error (normalised to the input state) is less than 10^{-5}. To attain the CSS, the successive substitution method uses the final state of the cycle as the initial state for the next cycle, following the dynamics of the process (Equation 2).

$$\boldsymbol{x}_{k+1}^0 = \boldsymbol{x}_k \tag{2}$$

As the design of adsorption processes is typically at the CSS and the dynamics are not of interest, the direct determination method described by Croft and LeVan is used, where a Newton iteration, based on the gradient of the error vector with respect to the initial state is used for the initial state of the next iterate (Equation 3).

$$\boldsymbol{x}_{k+1}^0 = \boldsymbol{x}_k^0 - \frac{\partial \boldsymbol{e}_k}{\partial \boldsymbol{x}_k^0}^{-1} \boldsymbol{e}_k \tag{3}$$

Within the literature, authors have mentioned the computational burden of solving for the sensitivities as a bottleneck for the implementation of this method [27]. The availability of gradients, accurate to the state solution integration, which can be computed efficiently in the JAX ecosystem, allows for the implementation of gradient-based methods.

RESULTS AND DISCUSSION

Convergence f the successive substitution and direct determination methods was compared for the PVSA process, with the initial condition corresponding to a bed saturated with the feed gas. The design parameters were taken from the optimal cost design reported by Ward and

Pini [5].

For the successive substitution method, convergence to the specified CSS tolerance required 116 cycles (Figure 2, left plot). Beyond this point, further cycling continued to reduce the error, but at a significantly slower rate. To reach an error level comparable to the numerical tolerance of the ODE solver, nearly 350 full process cycles were required. This slow asymptotic convergence behaviour is characteristic of successive substitution, which exhibits linear convergence near the fixed-point. In contrast, the direct determination method only required 6 Newton iterations, with solver level accuracy being achieved within 7 iterations (Figure 2, right plot). Compared to the fixed-point iteration, close to quadratic convergence is expected for this direct determination method as it includes the Newton iteration.

The convergence properties of cyclic adsorption processes can be characterised by the dominant Floquet multiplier, which is also known as the spectral radius, and is the largest eigenvalue of the Jacobian matrix for the cycle, $\frac{\partial \boldsymbol{x}_k}{\partial \boldsymbol{x}_k^0}$. When this value is less than unity, the CSS is stable, but as this value approaches one, the cycle dynamics become increasingly sluggish, such that more iterations are required to reach CSS when the successive substitution method is used. For the adsorption process considered in this work, the dominant Floquet multiplier was calculated as 0.952 at the CSS. For processes exhibiting even slower cycle dynamics, often seen in air purification PSA processes, systems with strongly adsorbing impurities, or rapid PSAs cycles, the direct determination method is therefore expected to yield even greater performance compared to the successive substitution method. Conversely, processes with temperature swings and longer cycle times, such as direct air capture processes, typically exhibit faster cycle dynamics such that fewer cycle simulations are required to reach CSS when using the successive substitution method.

While the single-design study illustrates the potential of the direct determination method, practical process design and optimisation studies require reliable performance of this method across a wide range of operating conditions. To assess robustness and scalability, both methods were therefore applied to a set of 4096 designs generated using Sobol sampling in a six-dimensional design space. The sampled parameters include the step durations for the adsorption, blowdown and evacuation steps, the high and intermediate pressure set points and the feed velocity. These parameters are commonly selected in adsorption process optimisation studies, where the low-pressure set-point is fixed at 0.05 bar, and 20s for the pressurisation time. The bounds of the design parameters are given in Table 2.

Table 2: Bounds for design parameters.

Parameter	Lower Bound	Upper Bound
$t_{ads}\ [s]$	20	100
$P_H\ [bar]$	1	10
$v_F\ [ms^{-1}]$	0.1	2
$t_{bd}\ [s]$	30	100
$P_I\ [bar]$	0.07	3
$t_{evac}\ [s]$	30	100

For each design, the CSS was computed using three different approaches:

1. Successive substitution in JAX
2. Direct determination in JAX
3. Successive substitution in MATLAB which is representative of standard practice.

For all implementations, simulations were run to the same CSS criteria. Between the two successive substitution methods, the solution method differs only in the evaluations of the Jacobian within the implicit ODE solver, which are computed through automatic differentiation in JAX and via numerical approximation in MATLAB. The comparison with an equivalent model in MATLAB was included to benchmark performance against the current standard approach for modelling cyclic adsorption processes. For the direct determination method, the ODE solver was modified such that the Jacobian was evaluated at every accepted time step, rather than being reused over multiple steps when the sensitivities are not required.

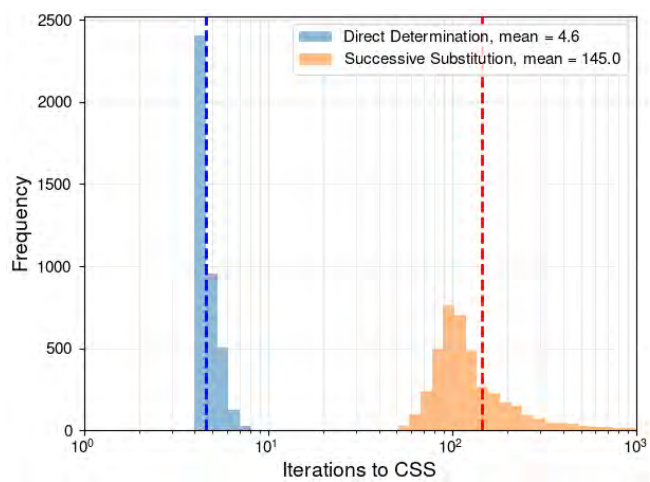

Figure 3. Distribution of iterations required to reach cyclic steady state for the 4096 Sobol-sampled designs, for the two methods. The number of iterations are plotted on a logarithmic axis.

Figure 3 shows the histogram of iterations required to reach CSS for the successive substitution and the direct determination methods across all Sobol-sampled designs. For every design, both methods converged to the same CSS, as verified by identical purity and recovery

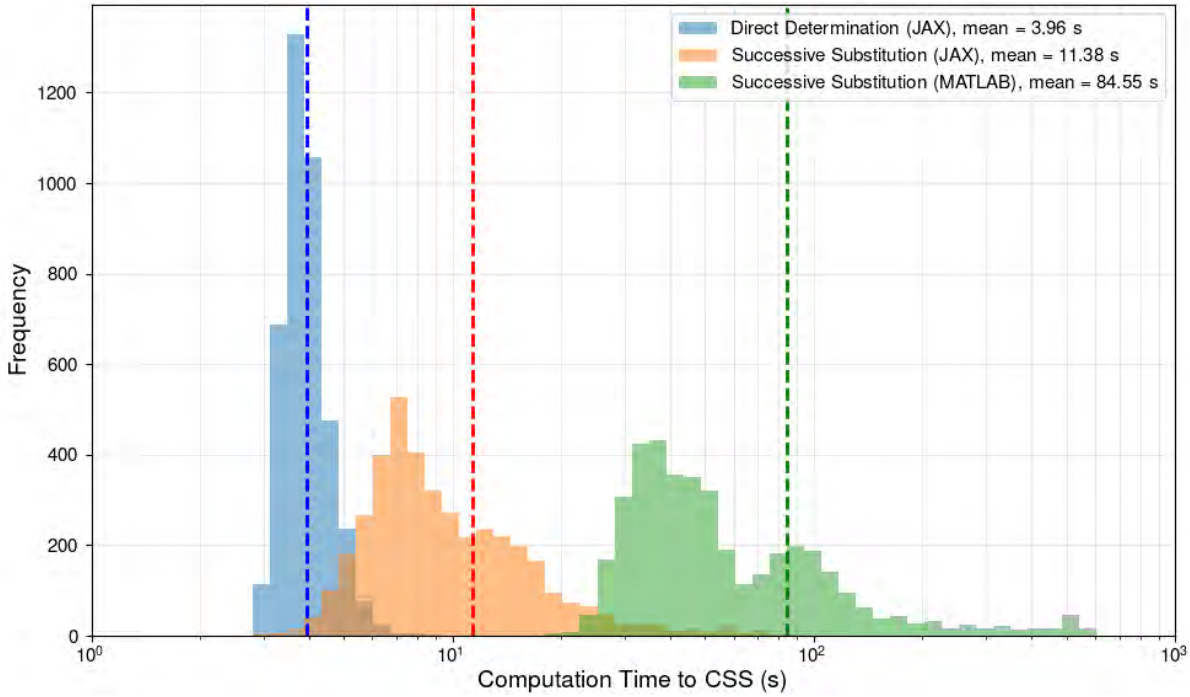

Figure 4. Distribution of CPU time required to reach cyclic steady state for 4096 Sobol-sampled designs, for each of the three approaches. The simulation time is plotted on a logarithmic axis

values at convergence. On average, the direct determination method converged to CSS in 4.6 iterations compared with 145 cycles for the successive substitution method. The spread of iterations was substantially larger for the successive substitution method. Most of the direct determination simulations converged in 4 iterations, with very few requiring more than 6. However, the successive substitution method required between 50 and 1000 iterations across the sampled design space. This behaviour highlights the robustness of direct determination which avoids the need to explicitly follow slow cycle dynamics to CSS.

In this work, the direct determination method was shown to attain CSS with significantly fewer iterations than the successive substitution method. This reduction in iteration count comes at the cost of additional computation per iteration due to the evaluation of state sensitivities. As such, computational time was also compared for the methods across all Sobol-sampled designs (Figure 4). For the process considered here, the direct determination method achieved convergence to CSS approximately three times faster than the successive substitution on average, requiring less than 4 seconds of CPU time per design. For processes with faster cycle dynamics, where fewer successive substitution iterations are required, the relative benefit of direct determination may be reduced. However, the automatic differentiation in JAX allows for this method to be accessible, accurate and efficient.

Also shown in Figure 4 is the comparison with the "standard" MATLAB implementation. The same solution method in JAX compared to MATLAB achieved nearly an eightfold reduction in solution time due to the just-in-time compilation and automatic differentiation available in JAX. These results the one-off compilation cost of 26 seconds for the JAX just-in-time compilation, which is incurred only during the first solution of the ODE system. For optimisation studies that require evaluating lots of model designs, this compilation overhead is negligible compared to the total computational cost.

CONCLUSIONS

In this work, a mechanistic model of a cyclic adsorption process for post combustion capture was developed within the JAX differentiable framework. A key limitation of adsorption process modelling is the computational time required to evaluate the dynamic process model as it typically involves solving a system of semi-discretized PDEs several hundred times to simulate the process from startup to cyclic steady state. This cost is exacerbated in design and optimisation studies where thousands of candidate designs must be evaluated. The differentiable model was used to accelerate the solution of the resulting ODE system, and to enable the efficient computation of accurate sensitivities. These sensitivities were exploited to implement a Newton-based direct determination method to accelerate the convergence to the cyclic

state, in which gradients of the final state of the cycle with respect to the initial state were used. This was made effective due to the automatic differentiation tools within the JAX environment.

Both the direct determination method and the traditional successive substitution method which simulates the actual dynamics of the process were compared for 4096 Sobol samples in a six-dimensional design space. Both methods converged to the same CSS, with the direct determination method requiring over an order of magnitude fewer iterations due to the quadratic convergence of the Newton method. Despite the additional computational effort required per iteration for sensitivity evaluation, the overall time to reach cyclic steady state was significantly reduced, with the direct determination method achieving convergence approximately three times faster on average compared to the successive substitution method.

Performance gains were also obtained during the solution of the ODE system for the JAX framework when compared to an equivalent MATLAB implementation which is representative of the current modelling practice. The combination of the direct determination method and the efficient ODE solution yielded a total speed-up of over 20 times compared to the MATLAB implementation. These results demonstrate that differentiable programming is a useful tool for accelerating the design of adsorption processes. As well as gradient-based methods for the determination of CSS, this framework is further extendable to gradient-based optimisation.

ACKNOWLEDGEMENTS

A.G. gratefully acknowledges funding from the EPSRC DTP scholarship granted to the Department of Chemical Engineering at Imperial College London.

AUTHOR IDENTIFIERS

Author ORCIDs:
Alex Glover: 0009-0007-8104-3916
Ronny Pini: 0000-0002-9443-3573
Maria M. Papathanasiou: 0000-0002-8886-0624

REFERENCES

1. Biegler LT, Jiang L, Fox VG. Recent advances in simulation and optimal design of pressure swing adsorption systems. Separation & Purification Reviews 33:1-39 (2005). https://doi.org/10.1081/spm-120039562
2. Subraveti SG, Karimi K, Gazzani M, Anantharaman R. A study on accelerated convergence of cyclic steady state in adsorption process simulations. Systems and Control Transactions 3:215-221 (2024). https://doi.org/10.69997/sct.137508
3. Croft DT, LeVan MD. Periodic states of adsorption cycles—i. direct determination and stability. Chemical Engineering Science 49:1821-1829 (1994). https://doi.org/10.1016/0009-2509(94)80067-7
4. Peh SB, Farooq S, Zhao D. A metal-organic framework (mof)-based temperature swing adsorption cycle for postcombustion CO2 capture from wet flue gas. Chemical Engineering Science 250:117399 (2022). https://doi.org/10.1016/j.ces.2021.117399
5. Ward A, Pini R. Efficient bayesian optimization of industrial-scale pressure-vacuum swing adsorption processes for co₂ capture. Ind. Eng. Chem. Res. 61:13650-13668 (2022). https://doi.org/10.1021/acs.iecr.2c02313
6. Ajenifuja A, Joss L, Jobson M. A new equilibrium shortcut temperature swing adsorption model for fast adsorbent screening. Ind. Eng. Chem. Res. 59:3485-3497 (2020). https://doi.org/10.1021/acs.iecr.9b05579
7. Ward A, Li K, Pini R. Assessment of dual-adsorbent beds for CO 2 capture by equilibrium-based process design. Separation and Purification Technology 319:123990 (2023). https://doi.org/10.1016/j.seppur.2023.123990
8. Pai KN, Prasad V, Rajendran A. Generalized, adsorbent-agnostic, artificial neural network framework for rapid simulation, optimization, and adsorbent screening of adsorption processes. Ind. Eng. Chem. Res. 59:16730-16740 (2020). https://doi.org/10.1021/acs.iecr.0c02339
9. Subraveti SG, Li Z, Prasad V, Rajendran A. Machine learning-based multiobjective optimization of pressure swing adsorption. Ind. Eng. Chem. Res. 58:20412-20422 (2019). https://doi.org/10.1021/acs.iecr.9b04173
10. Ding Y, LeVan MD. Periodic states of adsorption cycles III. convergence acceleration for direct determination. Chemical Engineering Science 56:5217-5230 (2001). https://doi.org/10.1016/s0009-2509(01)00127-0
11. Kvamsdal HM, Hertzberg T. Optimization of PSA systems—studies on cyclic steady state convergence. Computers & Chemical Engineering 21:819-832 (1997). https://doi.org/10.1016/s0098-1354(96)00315-8
12. Jiang L, Biegler LT, Fox VG. Simulation and optimization of pressure?swing adsorption systems for air separation. AIChE Journal 49:1140-1157 (2004). https://doi.org/10.1002/aic.690490508
13. Leis JR, Kramer MA. Algorithm 658. ACM Trans. Math. Softw. 14:61-67 (1988). https://doi.org/10.1145/42288.214371

14. Li S, Petzold L (1999) Design of new Daspk for Sensitivity Analysis
15. Smith OJ IV, Westerberg AW. Acceleration of cyclic steady state convergence for pressure swing adsorption models. Ind. Eng. Chem. Res. 31:1569-1573 (2002). https://doi.org/10.1021/ie00006a021
16. Pai KN, Prasad V, Rajendran A. Experimentally validated machine learning frameworks for accelerated prediction of cyclic steady state and optimization of pressure swing adsorption processes. Separation and Purification Technology 241:116651 (2020). https://doi.org/10.1016/j.seppur.2020.116651
17. Nilchan S, Pantelides CC. On the optimisation of periodic adsorption processes. Adsorption 4:113-147 (1998). https://doi.org/10.1023/a:1008823102106
18. Tsay C, Pattison RC, Baldea M. A pseudo?transient optimization framework for periodic processes: pressure swing adsorption and simulated moving bed chromatography. AIChE Journal 64:2982-2996 (2017). https://doi.org/10.1002/aic.15987
19. Lee I, Buse S, Winfree E (2025) Differentiable Programming of Indexed Chemical Reaction Networks and Reaction-Diffusion Systems. LIPIcs Vol 347 DNA 31 347:4:1-4:23. https://doi.org/10.4230/LIPICS.DNA.31.4
20. Alsubeihi M, Jessop A, Moseley B, Fonte CP, Rajagopalan AK. Modern, efficient, and differentiable transport equation models using JAX: applications to population balance equations. Ind. Eng. Chem. Res. 64:4541-4553 (2025). https://doi.org/10.1021/acs.iecr.4c04208
21. Bezgin DA, Buhendwa AB, Adams NA. Jax-fluids: a fully-differentiable high-order computational fluid dynamics solver for compressible two-phase flows. Computer Physics Communications 282:108527 (2023). https://doi.org/10.1016/j.cpc.2022.108527
22. Schoenholz SS, Cubuk ED. JAX, M.D. a framework for differentiable physics*. J. Stat. Mech. 2021:124016 (2021). https://doi.org/10.1088/1742-5468/ac3ae9
23. Haghpanah R, Majumder A, Nilam R, Rajendran A, Farooq S, Karimi IA, Amanullah M. Multiobjective optimization of a four-step adsorption process for postcombustion co₂capture via finite volume simulation. Ind. Eng. Chem. Res. 52:4249-4265 (2013). https://doi.org/10.1021/ie302658y
24. Hao Z, Caspari A, Schweidtmann AM, Vaupel Y, Lapkin AA, Mhamdi A. Efficient hybrid multiobjective optimization of pressure swing adsorption. Chemical Engineering Journal 423:130248 (2021). https://doi.org/10.1016/j.cej.2021.130248
25. Pai KN, Haghpanah R, Edsall W. Superstructure model for the simultaneous design and optimization of a pressure vacuum swing adsorption process cycle. Ind. Eng. Chem. Res. 64:627-638 (2024). https://doi.org/10.1021/acs.iecr.4c02842
26. Shampine LF, Reichelt MW. The MATLAB ODE suite. SIAM J. Sci. Comput. 18:1-22 (1997). https://doi.org/10.1137/s1064827594276424
27. Ko D, Siriwardane R, Biegler LT. Optimization of a pressure-swing adsorption process using zeolite 13X for co₂ sequestration. Ind. Eng. Chem. Res. 42:339-348 (2002). https://doi.org/10.1021/ie0204540

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Bayesian Optimization Framework for Agrochemical Formulation Design

Yipei Zhao[a] Tong Liu[b], Robin Wesley[c] and Joan Cordiner[a]

[a] School of Chemical, Materials and Biological Engineering, University of Sheffield, Sheffield S1 3JD, United Kingdom
[b] School of Computer Science, University of Sheffield, Sheffield S1 4DP, United Kingdom
[c] Syngenta, Jealott's Hill International Research Centre, Bracknell RG42 6EY, United Kingdom
* Corresponding Author: j.cordiner@sheffield.ac.uk.

ABSTRACT

Manufacturing kinetically stable products remains a challenge in the agrochemical industry. Current agrochemical formulation design relies on semi-empirical and trial-and-error methods. The inconsistency is caused by the lack of a mechanistic understanding of the formulation, making the design a black-box optimisation problem. In addition, validating the ground truth of the high-dimensional design space is expensive, driving chemists to explore possible solutions using data-driven methods. We proposed a Bayesian optimisation framework employing a Gaussian process as the surrogate model to intelligently guide the screening of the design space. The uniqueness of our framework is the application to the classification task to increase the number of hits of stable formulation recipes. The framework was tested on a provided industry dataset with a focus on emulsifiable concentrates. The performance reached a comparable accuracy with only ~25% of the data being sampled and hit more stable formulations than a Monte Carlo search. Our framework accelerates the discovery of stable recipes and guides formulation screening.

Keywords: Agrochemical Formulation, Bayesian Optimisation, Gaussian Processes, Space-Filling Designs, Machine Learning

INTRODUCTION

Developing stable and effective agrochemical formulations – specifically Emulsifiable Concentrates (EC) remains a challenging bottleneck in product development. The formulation design space is high-dimensional and is characterised by limited mechanistic understanding of the complex chemical interactions between active ingredients, surfactants, and solvents when diluted in water [1]. Conventionally, these complexities are addressed through iterative trial-and-error campaigns [2]. However, these methods are inefficient for exploring vast compositional design spaces and often fail to identify optimal formulations within tight industrial timelines. Moreover, the identification of subsequent candidate formulations is non-trivial. As empirical models often fail to generalise, researchers are forced into unguided experimental loops that waste materials and stall development.

Modelling this black-box function presents common constraints in chemical engineering: the gradient information (∇f) of the chemical interaction is usually unknown; thus, gradient descent methods cannot be applied [3]. Another challenge is that the cost of evaluations (experimental validation) is also expensive in labour and time to verify the function sufficiently.

To address these limitations, we present a data-driven, closed-loop Bayesian optimisation (BO) framework that systematically accelerates the formulation development process. Unlike the traditional design of experiments (DoE), our approach employs BO to actively learn the composition-stability relationship. By treating the formulation problem as a classification task (homogeneous vs. inhomogeneous), the framework iteratively proposes the next set of experimental candidates to maximise the probability of identifying a homogeneous recipe.

METHODOLOGY

Mathematically, the challenge can be formulated into $x^* \in \arg\max_{x \in X} f(x)$ where the design space is represented by $\mathcal{X} \subset R^d$, and the black-box function is

represented by matching the input x with output $f(x)$. Modelling the black-box function f is usually expensive in the chemical engineering field, due to the following constraints:

- **No analytical form.** The underlying physics/chemical reactions between materials are difficult to evaluate and express in closed-form equations.
- **Derivative-free.** The gradient information $\nabla f(x)$ is unknown. The gradient descent approach can't be applied.
- **Expensive evaluations.** Conducting experiments to verify the function can be costly.

Bayesian Optimisation

BO is a sample-efficient, iterative strategy designed for the global optimisation of expensive, black-box objective functions. Unlike Random Search or traditional DoE, which typically rely on static, pre-determined sampling plans, BO is inherently adaptive. It intelligently guides subsequent experiments by balancing the model's explored region (exploitation) with the unexplored region (exploration) [4]. This allows the algorithm to focus the search on the most promising areas of the design space, achieving the optimal solution with significantly fewer experimental runs than conventional methods.

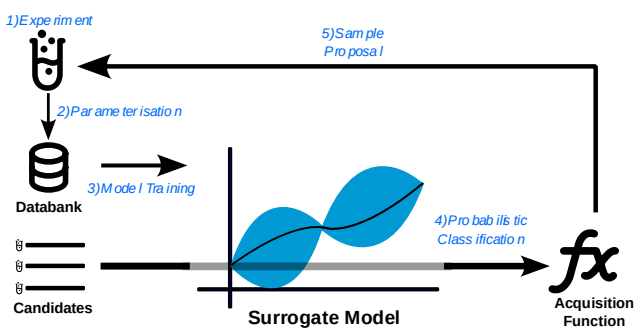

Figure 1. The iterative cycle of BO.

The BO procedure is an iterative strategy designed to optimise computationally expensive black-box functions. The workflow is characterised by the following stages:

(1) **Initial Sampling:** An initial set of experimental observations is collected.

(2) **Data Parameterisation:** The raw experimental inputs and corresponding outputs are transformed into a structured database for mathematical modelling.

(3) **Surrogate Model Construction:** A probabilistic surrogate model—most commonly a Gaussian Process (GP)—is trained on the available dataset to approximate the underlying objective function.

(4) **Acquisition Function Evaluation:** An acquisition function is employed to balance the exploration-exploitation trade-off. It utilises the surrogate model's mean predictions and associated predictive uncertainty to identify the most promising candidate for the next iteration.

(5) **Iterative Optimisation:** The suggested candidate is experimentally validated, the resulting data is appended to the training set, and the model is updated. This cycle repeats until a predefined termination criterion is reached.

The BO contains two core components: a surrogate model and an acquisition function. The surrogate model acts as the inference tool of the optimisation process, updating the probabilistic model given sampled data $\mathcal{D}_t = \{(x_i, y_i)\}_{i=1}^{t}$. The acquisition function is the decision maker, selecting the next point x_{t+1} to sample by balancing exploitation (sampling regions known to be more promising) and exploration (sampling regions with higher uncertainty) using a predefined hyperparameter.

A **Gaussian process** is usually considered as the gold standard surrogate model as they provide reliable uncertainty estimates, analytic tractability and sample efficiency [4]. The Gaussian process is defined by a mean function $m(x)$ and a covariance function (also known as the kernel) $k(x, x')$:

$$f(x) \sim \mathcal{GP}\big(m(x), k(x, x')\big), \tag{1}$$

both mean functions $m(x)$ and covariance function $k(x, x')$ are the prior beliefs of the objective function. They are critical in training as they assume the smoothness and periodicity, before any data is observed by the model. The mean function is usually considered to be $m(x) = 0$, while the selection of the covariance function is trickier. The Matérn kernel family offers a good balance between smoothness and noise-handling. The Matérn kernel $k_{a,h}: \mathcal{X}: \mathcal{X}$ is defined by:

$$k_{a,h}(x, x') = \frac{1}{2^{\alpha-1}\,\Gamma(\alpha)} \left(\frac{\sqrt{2\alpha}||x - x'||}{h}\right)^{\alpha} K_{\alpha} \left(\frac{\sqrt{2\alpha}||x - x'||}{h}\right)^{\alpha}, \tag{2}$$

Where Γ is the gamma function, and K_a is the modified Bessel function of the second kind of order α [5]. More details of the Gaussian process can be found in [6].

The other core component of BO is the acquisition function. In this work, we employed the Upper Confidence Bound (UCB) as the acquisition function. The UCB is defined as:

$$\alpha_{UCB}(x) = \mu(x) + \kappa \cdot \sigma(x), \tag{3}$$

where $\mu(x)$ and $\sigma(x)$ are the mean and standard deviation of the prediction from the surrogate model, respectively,

and κ is a predefined hyperparameter to balance exploitation and exploration. Compared to other acquisition functions, UCB is mathematically straightforward, intuitive, and based solely on the prediction and its uncertainty.

Space-Filling Designs

Space-filling designs aim to encourage the diversity included in the dataset, resulting in not only saving computational calculations but also leading to a more sophisticated model in interpolation/extrapolation by providing a better coverage or representation of the design space [7], [8]. In this work, we considered 3 space-filling design methods with random sampling.

Latin Hypercube Sampling (LHS) is a space-filling design strategy that improves upon simple random sampling by ensuring one-dimensional uniformity. For each of the k input variables, the range is divided into n equally spaced intervals, with exactly one observation per interval. Mathematically, a design point d_{ij} (representing the i-th run and j-th factor) is calculated using the formula $d_{ij} = \frac{l_{ij}+(n-1)/2+u_{ij}}{n}$, while l_{ij} is a value from a permutation of n levels (specifically taken as $-(n-1)/2$ to $(n-1)/2$) which determines the specific interval, and u_{ij} is a random number between 0 and 1 that provides a random position within that interval. This stratification ensures that the variance of the LHS estimator is lower than that of simple random sampling, specifically by subtracting the variance contributions of the main effects. As one of the most popular space-filling design strategies, LHS offers a reduction of the variance of the sample mean compared to random sampling and provides better coverage of the range of each input [9].

Sobol Sequence is a quasi-random number generator which offers lower discrepancy compared to a random number generator. A pseudo-random number generator can cause the data points to cluster together, leaving a large area unexplored. Unlike random sampling, where the discrepancy decreases probabilistically, Sobol sequences are constructed using radical inversion in base-2 to ensure the discrepancy decreases, resulting in a faster conversion rate [10].

Neither the LHS nor the Sobol sequence directly select samples from the existing database; instead, they generate a list of theoretical coordinates to sample. To sample from the existence database, samples are selected by comparing the distance between each data point to match the theoretical points proposed by both generators with the same dimension. Euclidean distance $d(p,q) = ||p-q||^2 = \sum_{i=1}^{m}(p_i - q_i)^2$ is used to measure the distance between data points pairwise.

Maximin distance design is more intuitive. By starting with a random sample, the algorithm iteratively selected samples that maximise the minimum Euclidean distance (can be substituted by other distance measures) from the sampled data points. This algorithm ensures that no data point will be close to other data points pairwise, thus ensuring a maximum spread of the samples.

CASE STUDY

To evaluate our framework, we utilise an experimental dataset provided by Syngenta. The data was generated using a customised automated robotic platform specifically designed for high-throughput formulation design and screening. Our goal is to validate the framework using an experimental dataset by sampling only a proportion of the data and testing the model with unseen data points. In the context of this case study, our optimisation objective is to maximise the probability of identifying a kinetically stable (homogeneous) formulation using BO. The underlying objective is to optimise for the most stable formulation composition. The optimisation variable defining our design space is the continuous weight percentages of the raw materials comprising the formulation product. Consequently, unguided formulation screening is time and cost-intensive; thus, we employed BO to accelerate the screening process.

Formulation Composition and Design Space

The dataset focuses on EC formulations. These formulations consist of a mixture of active ingredients (AI), solvents, and emulsifiers. To ensure a robust evaluation across a diverse chemical space, a total of 25 raw materials were tested:

- Active ingredients: Five unique AI and control formulations containing no active ingredients.
- Solvent Systems: Five unique solvent systems, comprising blends of two or three solvents (from 11 unique solvents) at various ratios.
- Emulsifier Systems: A primary anionic emulsifier along with six different co-emulsifier systems, each a blend of two emulsifiers selected from a pool of nine non-ionic emulsifiers.

Data Acquisition

Following preparation, the EC formulations were diluted in water at a 1:100 ratio to simulate practical manufacturing conditions. Kinetic (short-term) homogeneity was assessed 24 hours after dilution using a high-throughput image module.

To maximise information gain and capture subtle stability characteristics, the photography module utilised a multi-exposure technique. Each final data point is a concatenated image with a resolution of 9328 ×2550 pixels. These high-resolution images are an aggregation of 13 separate captures taken at varying exposures and

Category	Components	Count	Percentage/Details
Dataset Statistics	Total Samples (N)	811	100%
	Homogeneous	341	42.05%
	Inhomogeneous	470	57.95%
Formulation Components	Active Ingredients (AI)	5	Plus AI-free controls
	Solvent Systems	5	11 unique solvents
	Emulsifier Systems	6	9 unique emulsifiers
Total Unique Materials		**25**	

Table 1. A detailed description of the dataset and label distributions.

laser intensities.

Data Labelling & Distribution

The labels of all the formulations are evaluated via an automated image analysis tool designed to classify formulations based solely on visual characteristics. This tool achieved a binary classification accuracy of 94.83%, providing a highly reliable and scalable method for processing large image batches.

As summarised in Table 1, the final dataset consists of 811 samples. The distribution is relatively balanced, with 341 samples (42.05%) labelled as homogeneous and 470 samples (57.95%) labelled as inhomogeneous.

The model's input is compositional data, consisting of the weight percentages of each unique material. The input has been transformed using the centred log ratio (CLR) transformation and a standard scaler to make the model more sophisticated. The CLR transformation is used to handle the feature of compositional data (all data points have a constant sum) and preserves the mean distances [11]. The dataset was partitioned into two subsets using a 70:30 ratio. Seventy per cent of the data constituted the sampling pool, from which all initial and subsequent samples were drawn. The remaining 30% functioned as a held-out test set to evaluate model performance.

Optimisation Strategy

Since we are working on a classification rather than a regression model, we employed a different approach than Gaussian process regression. The construction and training of Gaussian process models are carried out using the Python library GPflow [12]. A Gaussian process classifier (GPC) can be slightly different to a regression task. A latent function must be applied to convert the output to the desired likelihood, which is a Bernoulli likelihood for a binary classification problem.

We have chosen Matérn kernel with $\alpha = 5/2$ (also known as the Matérn 5/2 kernel) as it provides a smoother result than other common hyperparameters. The GP models' other parameters, such as length scale and variance will be optimised by the GPflow library. The hyperparameter κ of UCB was chosen to be 1.5 to balance exploitation and exploration.

RESULT

We applied our framework to the provided dataset with an initial sample size of 100 and ran 50 iterations. As illustrated in Figure 2(a), the diverse sampling strategies fluctuate in performance but generally demonstrate the ability to achieve high predictive accuracy with significantly reduced data requirements. While the baseline model required 500 random samples to achieve an F1 score of 71.9%, the BO framework achieved comparable F1 scores (ranging between 0.6 and 0.8) with only 100 to 150 samples. Among the space-filling designs, the Sobol Sequence demonstrated strong potential, peaking above the baseline in the final iterations. This framework reduced the experimental burden by approximately 70% compared to the baseline random sampling required to achieve the same F1 score.

This underperformance may be attributed to the tendency of Maximin distance designs to prioritise the boundaries of the design space. Consequently, the central regions remain undersampled, leading to poorer generalisation in those areas compared to Sobol or Random initialisation.

Figure 2(b) details the efficiency of the framework in locating stable (homogeneous) formulations. The effectiveness of the BO approach is evident when compared to the Monte Carlo baseline (grey dashed line), which represents a pure random search without model guidance. The figure shows that the Gaussian model using random sampling achieves the highest hit rate (35 found), while demonstrating no superiority over LHS (33 Maximin distance (26 found). Whereas all Gaussian models found more homogeneous samples than a pure random search (19 found).

The result demonstrates the advantages of our framework, as all four different initial samplers captured more homogeneous formulations than the Monte Carlo method. This highlights the limitations of traditional trial-and-error approaches, which suffer from inconsistency and the inability to intelligently guide the search toward promising compositional regions.

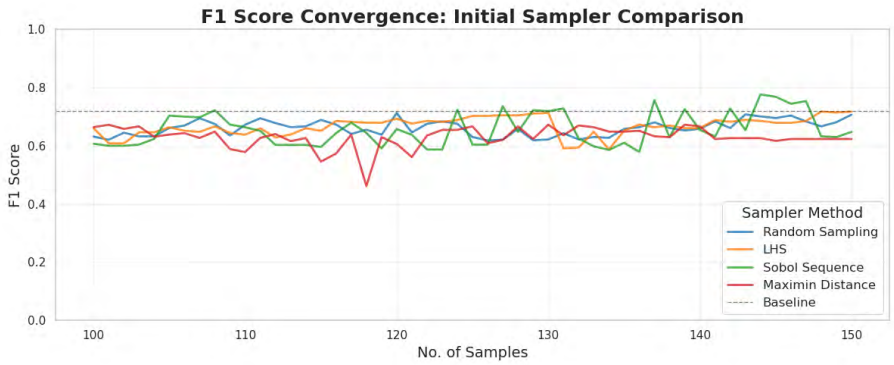

(a)

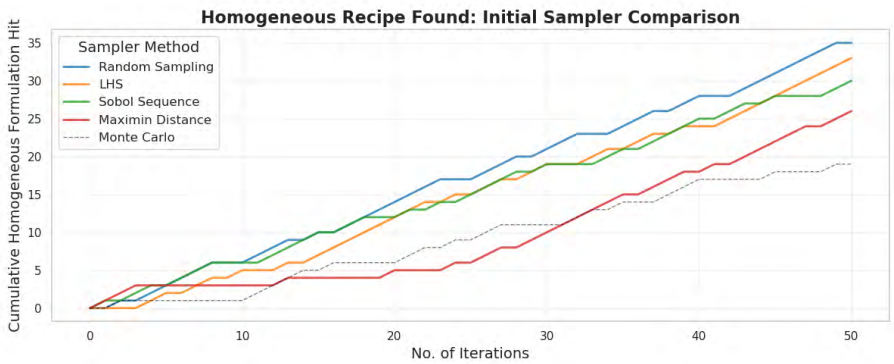

(b)

Figure 2. (a) The change of the F1 score over iterations. The grey dashed line is the baseline model trained using a Gaussian process with random sampling with 500 sampled data points, which leads to an F1 score of 71.9%. **(b)** The cumulative number of homogeneous formulations identified over 50 iterations using different initial sampling strategies compared to a Monte Carlo baseline.

CONCLUSIONS AND FUTURE WORK

This study presented a data-driven, closed-loop BO framework designed to address the inefficiencies in traditional trial-and-error practice for agrochemical formulation design. By treating the formulation design of EC as a binary classification task (homogeneous vs. found), the Sobol sequence (30 found), and outperforming inhomogeneous, we successfully demonstrated the framework's ability to navigate a high-dimensional design space with limited mechanistic understanding using Gaussian process models.

The next step will be applying the framework to a larger dataset to validate its performance against a dataset with higher dimensions and more samples. In addition, the idea can be applied to a multi-class classification task, but perhaps with slightly different approaches.

DIGITAL SUPPLEMENTARY MATERIAL

Due to commercial confidentiality, specific compositional data cannot be provided. However, we have included representative images captured by the camera module of the automated formulation system: https://github.com/yipeizhao/ESCAPE-material.

ACKNOWLEDGEMENTS

This is an iCASE awarded project, sponsored by Syngenta and the UK Engineering and Physical Sciences Research Council (EPSRC) under the grand code EP/Y528808/1.

REFERENCES

1. Tadros TF, Ed., *Emulsion formation and stability*. in Topics in colloid and interface science. Wiley-VCH (2013). https://doi.org/10.1002/9783527647941
2. Montgomery DC, Design and analysis of experiments, Ninth edition. John Wiley & Sons, Inc. (2017).
3. Jones DR, Schonlau M. Efficient Global Optimization of Expensive Black-Box Functions'.
4. Snoek J, Larochelle H, Adams RP. Practical Bayesian Optimization of Machine Learning

Algorithms. arXiv preprint (2012). arXiv:1206.2944.
5. Kanagawa M, Hennig P, Sejdinovic D, Sriperumbudur BK. Gaussian Processes and Kernel Methods: A Review on Connections and Equivalences. arXiv preprint (2018). arXiv:1807.02582.
6. Rasmussen CE, Williams CKI, *Gaussian processes for machine learning*, 3. print. In: Adaptive computation and machine learning. MIT Press (2008).
7. Joseph VR. Space-filling designs for computer experiments: a review. Quality Engineering 28:28-35 (2016). https://doi.org/10.1080/08982112.2015.1100447
8. Huang S. Surrogates: gaussian process modeling, design, and optimization for the applied sciences. Journal of Quality Technology 53:440-441 (2020). https://doi.org/10.1080/00224065.2020.1764416
9. Lin CD, Tang B. Latin Hypercubes and Space-filling Designs. arXiv preprint (2012) *arXiv*: arXiv:2203.06334.
10. Michel ES. Population Ecology in Practice. John Wiley & Sons, inc. (2020). ISBN: 9780470674147.
11. Galletti A, Maratea A. Numerical stability analysis of the centered log-ratio transformation. 2016 12th International Conference on Signal-Image Technology & Internet-Based Systems (SITIS) :713-716 (2016). https://doi.org/10.1109/sitis.2016.119
12. de G. Matthews AG, van der Wilk M, Nickson T, Fujii K, Boukouvalas A, Léon-Villagrá P, Ghahramani Z, Hensmen J. GPflow: A Gaussian process library using TensorFlow, *J Machine Learning Res* 18:1–6 (2017)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Variational Bayesian Neural Networks for Modelling and Uncertainty Quantification in Bioprocessing

George Spencer[a], Harini Narayanan[b], Claus Wirnsperger[b], Alessandro Butté[b], Cleo Kontoravdi[a], Maria M. Papathanasiou[a*]

[a] The Sargent Centre for Process Systems Engineering, Imperial College London, London, United Kingdom, SW7 2AZ
[b] DataHow AG, Hagenholzstrasse 111, 8050, Zürich, Switzerland
* Corresponding Author: maria.papathanasiou11@imperial.ac.uk.

ABSTRACT

Biopharmaceutical manufacturing requires systematic identification, understanding and control of critical process parameters (CPPs) and critical quality attributes (CQAs) While deterministic machine learning models have proven valuable for describing complex systems and providing insights into process behaviour, the extension to probabilistic frameworks allows for the capture of intrinsic biological and process variability, improving robustness, understanding and safety. This work introduces a framework for autoregressive variational Bayesian neural networks (BNNs) that integrate uncertainty quantification into data-driven modelling. The approach is demonstrated on a multicolumn countercurrent solvent gradient purification (MCSGP) process for monoclonal antibody purification. The BNN is trained autoregressively on high fidelity data generated using a detailed mechanistic process model data. A heteroscedastic variance head is included to model signal-dependent observation noise, while the epistemic uncertainty is captured through stochastic network parameters. Model performance is assessed using both deterministic accuracy metrics - MSE and probabilistic scores namely coverage and continuous ranked probability scores (CRPS). Results show that the proposed method accurately recovers noiseless concentration profiles while maintaining well calibrated predictive distributions. Additionally, the framework is tested under multiple noise regimes to test robustness of the proposed approach.

Keywords: Modelling and Simulations, Algorithms, Bayesian Neural Networks, Variational Inference, Uncertainty Quantification, Chromatography

INTRODUCTION

Biologics, such as monoclonal antibodies (mAbs), play a central role in the treatment of diseases, including cancers and autoimmune disorders. The commercial importance of mAbs is substantial, with six of the top ten best pharmaceutical products being mAb-based therapies [1]. As a result, there is an increasing demand for cost minimisation and high production efficiency, while regulators are pushing for transparency, robustness and explainability through initiatives such as Quality by Design (QbD) and the drive towards Biopharma 4.0. In this context, process modelling provides a valuable tool for supporting design, operation and decision making, by improving understanding and operability of complex bioprocesses.

Data-driven modelling allows engineers with access to process data and often limited mechanistic process knowledge to model unknown phenomena. This is pertinent in upstream bioprocessing where many underlying mechanisms remain partially characterised. In downstream processing, where process knowledge is richer, data-driven modelling can offer significant improvement in computational time for inference, making them attractive for online monitoring, optimisation, and control. However, in the pursuit of Biopharma 4.0, deterministic model inference alone is insufficient to support risk-aware decision making. Biological systems are often characterised by substantial inherent variability and therefore process decisions must account for uncertainty to satisfy stringent regulatory constraints and ensure consistent product quality.

In this setting, Bayesian Neural Networks (BNNs) offer a promising solution, by retaining the predictive

capabilities of artificial neural networks, while providing uncertainty estimates via probabilistic inference. In contrast to deterministic models, BNNs provide predictive distributions rather than point estimates, enabling quantification of uncertainty. This includes aleatoric uncertainty arising from measurement noise or process variability, and epistemic uncertainty associated with limited data and uncertainty in model parameters. Gaussian Processes (GP) are commonly utilised for this purpose, however lack scalability in dynamic systems such as bioprocessing [2].

Many algorithms exist for the training of BNNs, with Markov Chain Monte Carlo (MCMC) methods often being regarded as the gold standard [3]. Despite their wide application, these exact methods can be slow to converge and may become computationally prohibitive for large networks. Approximate inference methods offer a scalable alternative that is compatible with gradient based optimisation. In this work, the applicability of variational inference for autoregressive BNNs in bioprocess modelling is assessed using ion-exchange chromatography as a use case. The proposed framework is designed to enable probabilistic prediction and decompose aleatoric and epistemic uncertainties. This distinction is important, as only epistemic uncertainty can be reduced by more data and targeted experimentation, whereas the aleatoric uncertainty represents intrinsic process variability that must instead be managed through robust operation and control.

THE USE CASE OF MULTICOLUMN COUNTERCURRENT SOLVENT GRADIENT PURIFICATION (MCSGP) PROCESS

The case study of interest is the twin-column Multicolumn Countercurrent Solvent Gradient Purification (MCSGP) process presented by Aumann and Morbidelli [4]. The process is an ion exchange chromatography used as a purification step in mAb production. The system separates three components, namely the product, the weak impurities, and the strong impurities.

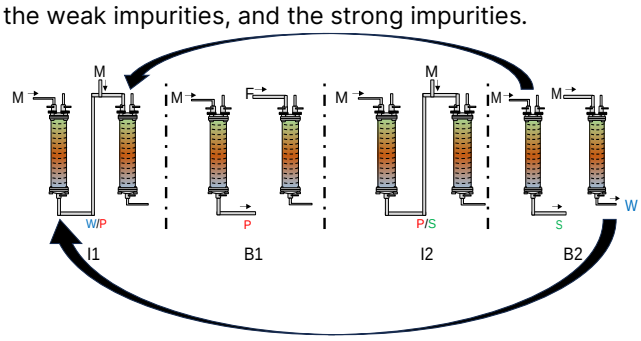

Figure 1: Diagram of MCSGP operating cycle

Operation proceeds through four phases: Two interconnected (I1 and I2) and two batch (B1 and B2) (Figure 1). During the interconnected phases, outlet streams containing off specification product fractions are recycled between columns to enhance recovery, while batch phases allow for the elution of high purity product. Strong and weak impurity fractions are discarded. After the cycle shown in Figure 1, the roles of the columns reverse and the process is repeated.

A high fidelity, experimentally validated, process model, originally developed by Müller-Späth et al. [5] was utilised to generate data of the MCSGP process. This lumped kinetic model, comprises over 4000 ordinary differential equations [6], after spatial discretization, representing a computational bottleneck for real time applications and uncertainty quantification (UQ).

METHODOLOGY

Data Generation and Process Model

Data is generated according to the strategy proposed by Michalopoulou et al. [7]. In which Sobol sampling is used to sample nine design variables through ranges presented in Table 1.

Table 1: Inputs and Bounds for Data Generation

Variable	Lower Bound	Upper Bound	Units
Weak Feed Concentration	0.05	0.08	mg/ml
Product Feed Concentration	0.3	0.5	mg/ml
Strong Feed Concentration	0.03	0.05	mg/ml
Batch Flowrates	0.1	1.0	ml/min
I1 Flowrate	0.1	1.0	ml/min
I2 Flowrate	0.1	1.0	ml/min
Modifier Feed Concentration	1.0	3.0	mg/ml
Initial Modifier Concentration Phase I1 column 1	2.0	3.0	mg/ml
Initial Modifier Concentration Phase I1 column 2	0.8	1.2	mg/ml

The model is simulated to Cyclic Steady State (CSS) and input-output data, consisting of these nine design variables and the output concentrations of each species for a full CSS cycle is taken to create the dataset. For quantifying aleatoric uncertainty in this methodology, heteroscedastic noise is added to the measured concentrations of species.

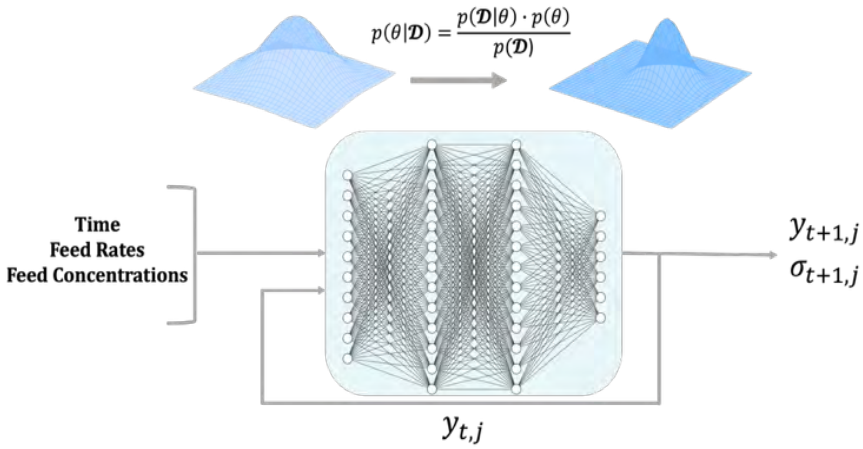

Figure 2. Autoregressive Bayesian Neural Network Architecture

Autoregressive Bayesian Neural Network

A dataset of multivariate time sequences $\left\{\left(x_{1:T}^{(n)}, y_{1:T}^{(n)}\right)\right\}_{n=1}^{N}$, where $x_t \in \mathbb{R}^{d_x}$ are exogenous inputs and $y_t \in \mathbb{R}^{d_y}$ are the targets, is constructed. The model is trained autoregressively to predict $\hat{y}_t^{(n)} = f_\theta\left(x_{t-C:t}^{(n)}, \hat{y}_{t-C:t-1}^{(n)}\right)$. Using a fixed context length C. Predictions are generated sequentially, with the previous C predicted outputs fed back as inputs for subsequent time steps.

The predictor, f_θ, is a Bayesian neural network with a factorised Gaussian variational posterior over weights and biases. For each layer ℓ, the posterior is parameterised as shown in Eq. (1) and (2).

$$q(W_\ell) = \mathcal{N}\left(\mu_{W_\ell}, \sigma_{W_\ell}^2\right) \tag{1}$$

$$q(b_\ell) = \mathcal{N}\left(\mu_{b_\ell}, \sigma_{b_\ell}^2\right) \tag{2}$$

Where W_ℓ and b_ℓ denote the weights and biases of layer ℓ of the network. The standard deviations are parameterised via Eq. (3).

$$\sigma = log\left(1 + exp(\rho)\right) \tag{3}$$

With, ρ, a learned parameter for each weight and bias, to ensure positivity and improved training stability. The BNN includes a heteroscedastic variance head that predicts the log-variance, $log\left(\sigma_{t,j}^2\right) \forall j = 1, \ldots, d_y$, for each output dimension at each timestep. A schematic of the network is displayed in Figure 2.

Variational Inference

The training of the BNN approximates the posterior distribution over network parameters (θ) given by Eq. (4)

$$p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta)p(\theta)}{p(\mathcal{D})} \tag{4}$$

Where, $p(\mathcal{D})$ is the evidence or the marginal likelihood shown in Eq. (5).

$$p(\mathcal{D}) = \int p(\mathcal{D}|\theta)p(\theta)\, d\theta \tag{5}$$

Which requires integration over all possible parameter realisations and is non-conjugate and therefore intractable to compute analytically requiring an approximate inference algorithm. In this case variational inference was chosen as it can be seamlessly integrated into gradient descent optimisation algorithms.

Variational inference introduces a parameterised variational distribution and seeks a member of this family that is the closest to the true posterior, defined by Eq. (6) the Kullback-Leibler (KL) divergence.

$$q_\phi^*(\theta) = \arg\min_{q_\phi} D_{KL}\left(q_\phi(\theta) \,\middle|\, p(\theta \mid \mathcal{D})\right) \tag{6}$$

Direct evaluation of Eq. (6) is infeasible, as it requires the unknown posterior, therefore optimisation is performed indirectly by maximising the evidence lower bound (ELBO) defined in Eq. (7).

$$\mathcal{L}_{\mathcal{ELBO}} = \mathbb{E}_{q_\phi(\theta)}[\log p(\mathcal{D} \mid \theta)] - \mathrm{KL}\left(q_\phi(\theta)|p(\theta)\right) \tag{7}$$

The ELBO defines a lower bound on the log-evidence and is related to the intractable KL divergence in Eq. (8).

$$log\ p(\mathcal{D}) = \mathcal{L}_{\mathcal{ELBO}} + \mathrm{KL}\left(q_\phi(\theta)||p(\theta|\mathcal{D})\right) \tag{8}$$

Since the KL divergence on the right-hand side of Eq. (8) is non-negative, maximising the ELBO is the equivalent to minimising the KL divergence between the variational approximation and the true posterior. The ELBO consists of two competing terms, the expected log-likelihood term that encourages the model to explain the observed data, while the KL term acts as a regulariser towards the prior distribution.

The variational loss function applied in this algorithm is defined by Eq. (9).

$$\mathcal{L} = -E_{q(\theta)}\left[\sum_{t=C+1}^{T} \log p\left(\hat{y}_t \mid \hat{\mu}_t, \hat{\sigma}_t^2\right)\right] + \beta \mathrm{KL}(q(\theta)||p(\theta)) \quad (9)$$

To enable gradient based optimisation of the ELBO, the expectation with respect to the variational posterior is approximated using Monte Carlo sampling combined with the reparameterization trick [8]. Specifically, for layer ℓ, weights and biases are sampled from Eq. (10) and Eq. (11).

$$W_\ell = \mu_{W_\ell} + \sigma_{W_\ell}\epsilon_\ell \quad (10)$$

$$b_\ell = \mu_{b_\ell} + \sigma_{b_\ell}\epsilon'_\ell \quad (11)$$

Where $\epsilon \sim \mathcal{N}(0, \mathrm{I})$. This formulation expresses the stochastic network parameters as deterministic functions of the variational parameters and auxiliary noise variables (ϵ), allowing the Monte Carlo estimated ELBO to be computed via automatic differentiation.

Model Training

The training data consists of 400 cyclic steady state sequences generated from the mechanistic model, with the nine varied design variables and three output concentration profiles. A further 200 sequences are used for validation, and 90 sequences are held out for testing and visualisation.

The BNN receives 13 input features at each time step. These consist of the nine design variables, the most recent prediction of the three outlet concentrations, and time. The network contains three fully connected hidden layers with sizes 64, 64 and 32, respectively, each followed by ReLU activations. The architecture and model hyperparameters were optimised using Tree Structured Parzen Estimator based Bayesian Optimisation [9]. The final output layer consists of six neurons, predicting the mean and heteroscedastic noise estimate for each concentration.

During training the KL divergence term is computed independently for each layer and scaled by the batch size to avoid batch-size dependence or the variational objective. To improve optimisation stability, the KL weight β, is linearly annealed from 0 to 1 during an initial warmup phase. After warmup, full variational regularisation is applied.

The ELBO is approximated using $\mathcal{K}$ Monte Carlo samples of the network parameters. For each particle independent random keys are used to sample weights and biases, producing $\mathcal{K}$ autoregressive rollouts.

During training teacher forcing is employed to stabilise the learning. At each time step, the history value fed back to the model is chosen according to Eq. (11).

$$y_t^{hist} = \begin{cases} y_t, & with \text{ probability } p_{TF} \\ \hat{\mu}_t, & with \text{ probability } (1 - p_{TF}) \end{cases} \quad (11)$$

Where p_{TF} is the probability of teacher forcing. This probability is annealed linearly over a fraction of the total optimisation steps in training, gradually transitioning from fully guided rollouts towards fully self-predicted trajectories.

When physical bounds are supplied the predictive mean is mapped through a sigmoid transformation (Eq. 12)

$$\hat{\mu}_t = \mathrm{a} + (b - a)\sigma(\hat{\mu}_t^{raw}) \quad (12)$$

Where a and b represent the lower and upper bounds and $\sigma(\cdot)$ is the sigmoid function. In this case bounds are chosen as [0, 1] mg/ml. This deterministic and differentiable mapping enforces hard constraints on the mean prediction while preserving gradient based optimisation of the variational loss. Since epistemic uncertainty arises from variability in the network parameters, constraining the sampled mean predictions also constrains the resulting epistemic uncertainty distribution. In contrast, the heteroscedastic variance head remains unbounded, allowing the model to freely represent the aleatoric noise.

Model Inference

At inference time, predictive uncertainty is obtained via Monte Carlo sampling of the learned variational posterior. In all experiments, 1000 parameter samples are drawn. For each sample, a full autoregressive rollout of the concentration trajectories is generated.

The predictive mean is computed as the mean trajectories across all samples. Epistemic uncertainty is quantified as the variability across these sampled mean predictions. This sampling-based formulation allows the predictive distribution to be non-Gaussian and enables direct estimation of prediction quantiles.

Aleatoric uncertainty is obtained from the heteroscedastic output head. Since the variance prediction itself depends on stochastic network parameters, the aleatoric estimate exhibits epistemic uncertainty as a second order effect. The total predictive uncertainty is therefore composed of both epistemic and aleatoric contribution, combined under a conditional independence assumption.

Computation

In this work the BNN architecture and training algorithm were constructed using JAX, to enable efficient automatic differentiation and hardware acceleration [10]. Just-in-time (JIT) compilation was used extensively to optimise the autoregressive rollout and ELBO computation, allowing entire sequence scans to be compiled into fused kernels. Vectorisation was performed across batches and Monte Carlo particles to enable efficient parallel evaluation of stochastic network samples during training and inference. Gradient descent was performed with AdamW using the Optax suite [11], [12]. All experiments are performed using an NVIDIA RTX 4090 GPU.

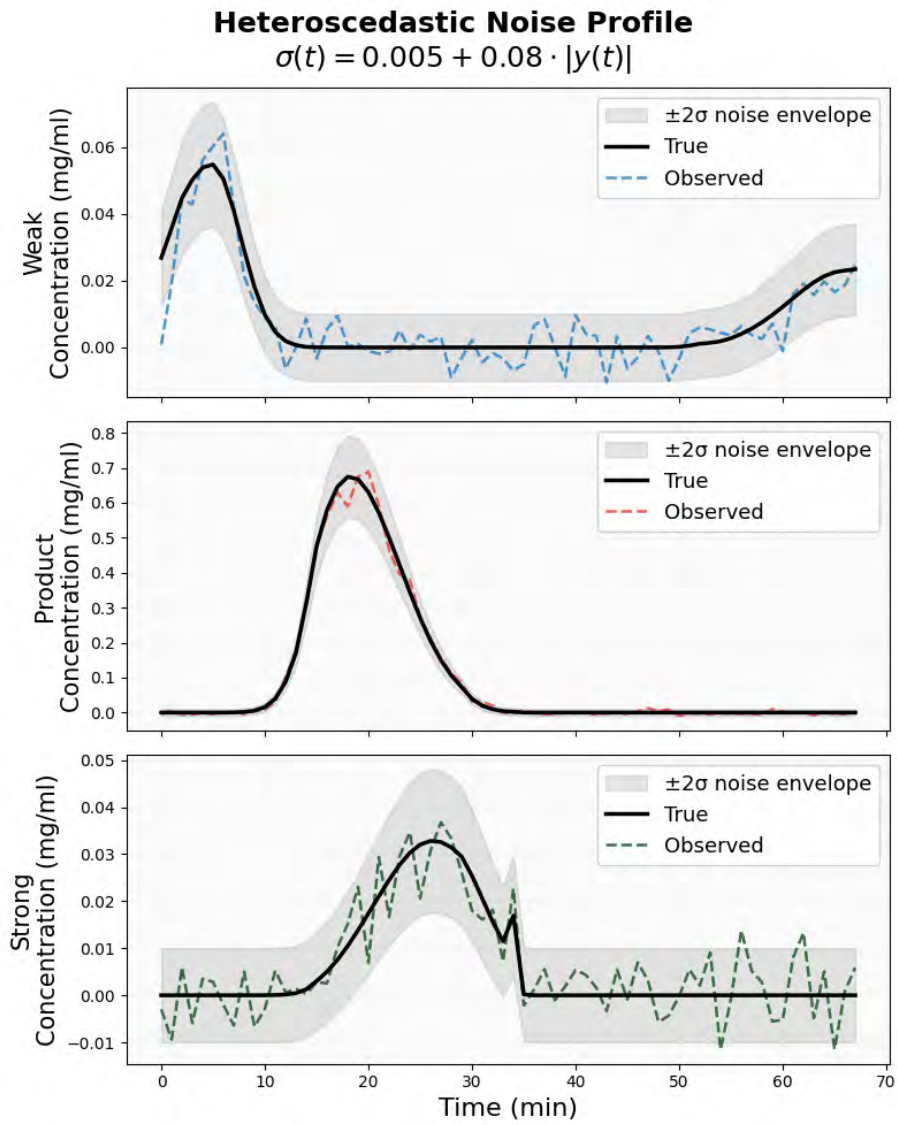

Figure 3. Nominal heteroscedastic noise profiles applied to outlet concentrations.

RESULTS AND DISCUSSION

Nominal Test Case

The nominal test case takes the train, test and validation sets and applies noise to the outputs before training and testing. The noise model in this scenario is Gaussian heteroscedastic noise with noise standard deviation defined as a sum of a constant baseline component and a signal magnitude dependent component given by Eq. (13)

$$y_{t,j}^{noisy} = y_{t,j} + \epsilon_{t,j},\ \epsilon_{t,j} \sim \mathcal{N}\left(0, \left(\sigma_0 + \alpha|y_{t,j}|\right)^2\right) \quad (13)$$

Where σ_0 represent the baseline noise and α controls the signal dependent noise, and where j represents each component. For this nominal case $\sigma_0 = 0.005, \alpha = 0.08$ was used, an example profile is shown in Figure 3.

Validation metrics of deterministic prediction accuracy and uncertainty quantification capability are displayed in Table 3. The mean squared error (MSE) quantifies the accuracy of the predictive mean, while the empirical coverage, the frequency at which the model's predicted intervals contain the observed values, evaluates the calibration of the predictive intervals. The continuous ranked probability score (CRPS) [13] is a proper scoring rule that assess the difference between the predicted and observed cumulative distributions.

Figure 4 displays a representative outlet concentration prediction profile from the test set. The predictive means closely follow the true, noiseless, profiles. While the predictive uncertainty bands capture the epistemic and aleatoric uncertainty throughout time. A subset of

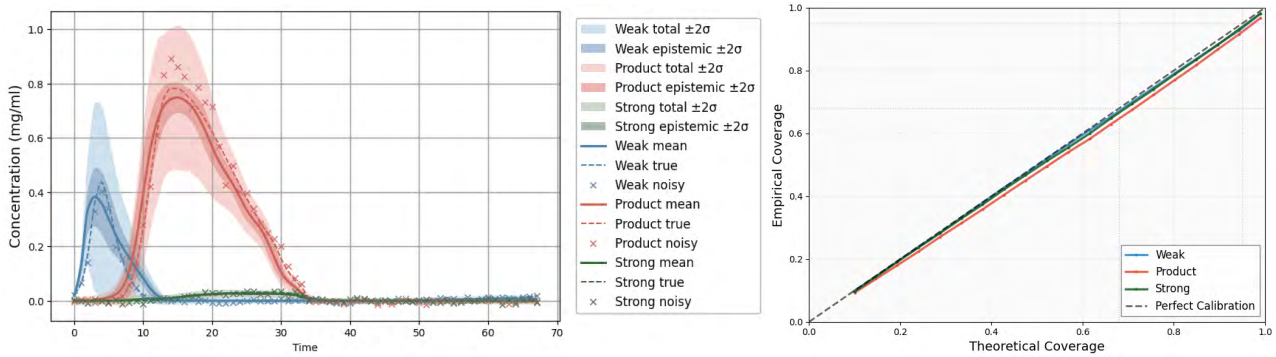

Figure 4. Left**:** Predicted outlet concentration profiles for representative test sequence. Right: Calibration plots for nominal case, over all test set sequences.

noisy realisations are marked with a cross to demonstrate the capability of the model to capture the heteroscedastic noise.

Table 2: Validation Metrics for Nominal Case

Metric	Weak Impurities	Product	Strong Impurities	Units
MSE	6.32e-4	1.79e-3	1.37e-4	[mg/ml]^2
95% Coverage	93.4%	92.2%	93.8%	-
CRPS	5.99e-3	1.13e-2	4.03e-3	mg/ml

The probabilistic calibration of the predictive distributions is assessed via calibration curves shown in Figure 4, across all sequences in the test set, for each component. The empirical coverage closely matches the nominal confidence levels across all components, with slight overconfidence in the product profiles. Overall, the nominal test case demonstrates that the predictive distribution from the BNN achieves accurate predictions and is able to filter noise to predict close to true noiseless values with low MSE, within its predicted epistemic uncertainty. While the noise levels are accurately captured by the heteroscedastic head, with a mean absolute error in noise levels of 8e-4.

Secondary Noise Model

To better reflect realistic UV sensor behaviour typically observed in chromatographic systems, and test the algorithm in different, more complex and realistic noise regimes, a secondary noise model is introduced. Unlike the nominal case, which assumes temporally independent noise observations, this model incorporates additional structure, including temporal correlation and baseline drift. In this case there is an additive noise given by Eq. (14)

$$y_{t,j}^{noisy} = y_{t,j} + \epsilon_{t,j} \tag{14}$$

Where $\epsilon_{t,j}$ is composed of three contributions:

(i) The heteroscedastic white noise displayed in Eq. (15)

$$\epsilon_{het,t,j} \sim \mathcal{N}\left(0, \left(\sigma_0 + \alpha|y_{t,j}|\right)^2\right) \tag{15}$$

(ii) Temporally correlated first order autoregressive noise ($AR(1)$) shown in Eq. (16)

$$\epsilon_{ar,t,j} = \rho\epsilon_{ar,t-1,j} + \eta_{t,j},$$
$$\eta_{t,j} \sim \mathcal{N}\left(0, (\sigma_{ar})^2\right) \tag{16}$$

Where $\rho \in (-1, 1)$ controls the strength of temporal correlation and σ_{ar} determines the magnitude of the correlated noise.

(iii) Baseline drift modelled as a random walk given by Eq. (17)

$$d_{t,j} = d_{t-1,j} + \xi_{t,j} \tag{17}$$

with $\xi_{\{t,j\}}$ denoting a zero mean Gaussian increment.

Figure 5 illustrates the new noise profiles on an example sequence. This test case uses: $\sigma_0 = 0.002, \alpha = 0.02, \rho = 0.95, \sigma_{ar} = 0.002$.

The model architecture, training epochs and data set size are identical to the nominal test case for comparability. Quantitative validation metrics obtained under the secondary regime are reported in Table 3. Compared to the nominal case, point prediction error increases modestly across all components, reflecting the increased complexity in filtering temporally correlated noise and baseline drift. Despite this, the empirical 95% coverage remains close to the nominal level across all components, indicating that the predictive uncertainty adapts to the new noise regime. The increase in CRPS reflects a reduced predictive sharpness and increased uncertainty, which is consistent with the more challenging noise model.

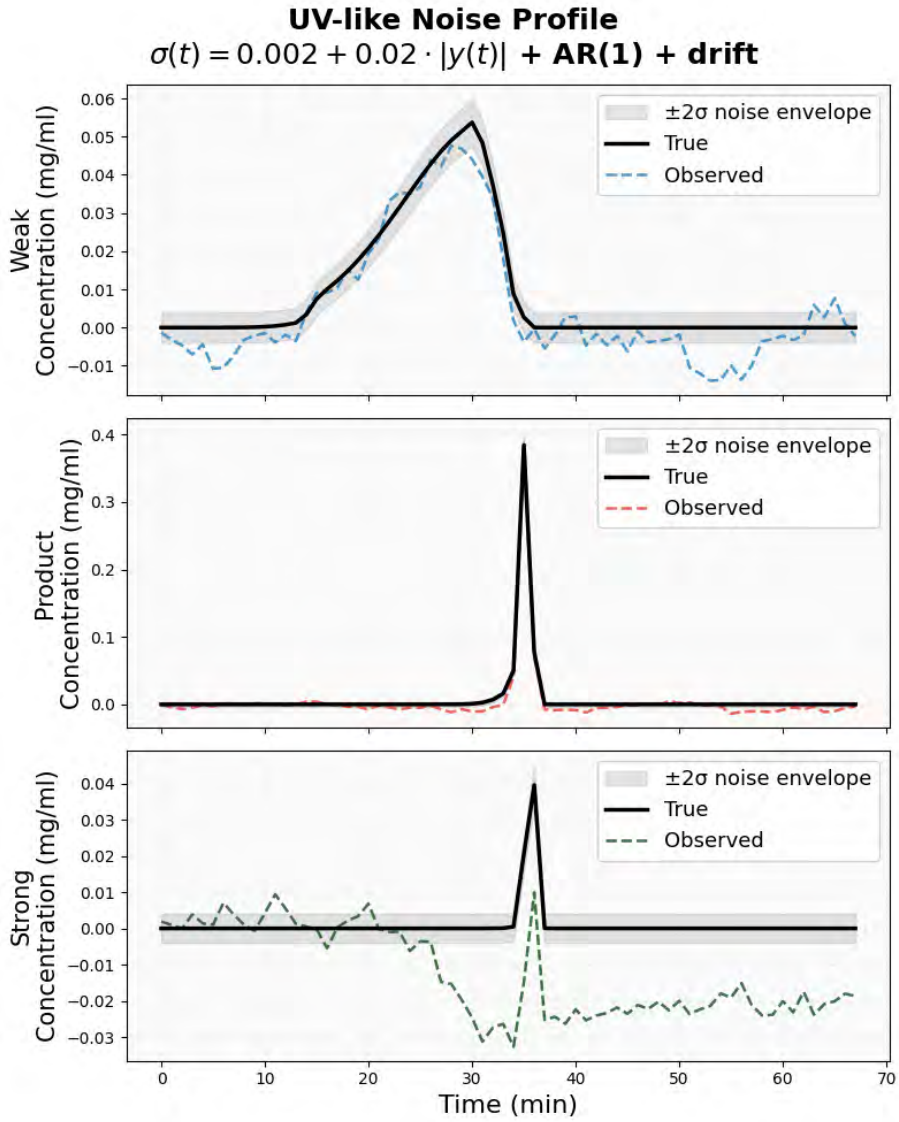

Figure 5. Secondary noise model profiles applied to outlet concentrations.

Figure 6 shows a representative prediction under the secondary noise model. Despite the added complexity of the noise regime, the predictive mean is still able to track the noiseless truth, while the uncertainty bands adapt to the new regime. The predictive uncertainty increases in regions effected by baseline drift, reflecting reduced confidence in the inferred signal.

The right-hand side of Figure 6 shows the calibration of the network remains tight across all components, with slight overconfidence in some regions.

Table 3: Validation Metrics for Secondary Case

Metric	Weak Impurities	Product	Strong Impurities	Units
MSE	7.54e-4	2.73e-3	2.40e-4	$[mg/ml]^2$
95% Coverage	91.8%	92.6%	90.6%	-
CRPS	1.14e-2	1.85e-2	9.23e-3	mg/ml

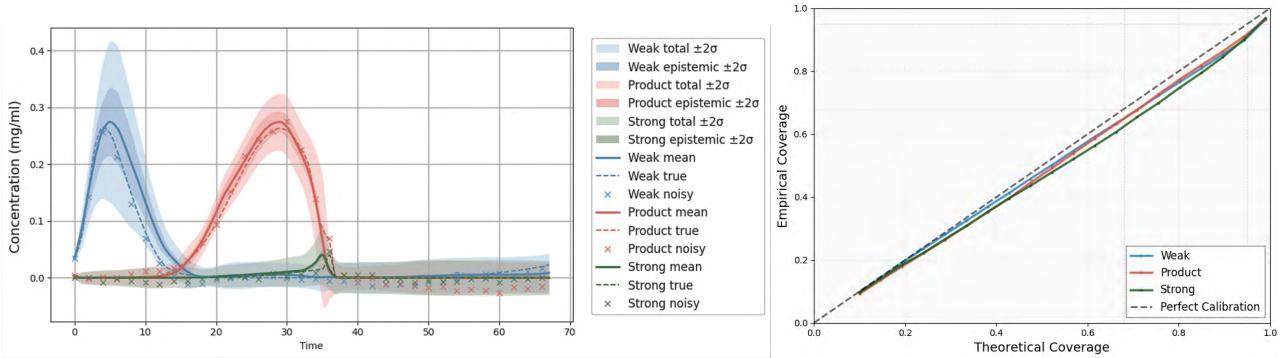

Figure 6. Left: Predicted outlet concentration profiles for representative test sequence, under secondary noise model. Right: Calibration plots for secondary case, over all test set sequences.

CONCLUSION AND OUTLOOK

This work has demonstrated the applicability of autoregressive Bayesian neural networks trained via variational inference for modelling and uncertainty quantification in downstream bioprocessing operations. Using the MCSGP case study, the proposed framework was shown to produce accurate predictive means while also providing well calibrated uncertainties. The decomposition of predictive uncertainty into epistemic and aleatoric components enables meaningful interpretation of model confidence.

The nominal test case demonstrated that the BNN is able to accurately recover noiseless concentration profiles from heteroscedastic observations, while maintaining reliable uncertainty calibration. The introduction of a secondary, more complex, noise model showed the framework remains robust under more challenging, but practically relevant conditions.

A key strength of the proposed approach is the autoregressive formulation, which allows the model to generate predictions sequentially without reliance on cyclic steady state conditions. Looking forward, the framework will be tested under fully dynamic operation or transient disturbances, allowing deployment in online monitoring, optimisation or control.

A particularly promising direction in the application of the framework is in upstream bioprocessing, where biological variability arising from cellular mechanisms introduces substantially higher uncertainty than downstream steps. Eventually, targeting an uncertainty-aware holistic modelling pipeline connecting the upstream and downstream processing units.

Training data requirements remain a key limitation for Bayesian Neural Networks over using Gaussian Processes. Future work will explore smaller models and lower training data availability, to demonstrate applicability in the data sparse environments often found in bioprocessing. Particularly looking at the effect of training data availability on epistemic uncertainty estimates.

ACKNOWLEDGEMENTS

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) for the Imperial College London Doctoral Training Partnership (DTP)

AUTHOR IDENTIFIERS

Author ORCIDs:
Spencer G: 0009-0002-5375-874X
Narayanan H: 0000-0003-3580-284X
Wirnsperger C: 0009-0000-8511-1428
Butté A: 0000-0003-3506-3792
Kontoravdi C: 0000-0003-0213-4830
Papathanasiou M. M: 0000-0002-8886-062

REFERENCES

1. Quinteros DA, Bermúdez JM, Ravetti S, Cid A, Allemandi DA, Palma SD. Therapeutic use of monoclonal antibodies: general aspects and challenges for drug delivery. Nanostructures for Drug Delivery :807-833 (2017). https://doi.org/10.1016/b978-0-323-46143-6.00025-7
2. M. Sendera, A. Sorkhei, and T. Kuśmierczyk, 'Revisiting the equivalence of Bayesian neural networks and Gaussian processes: on the importance of learning activations', in *Proceedings of the Forty-First Conference on Uncertainty in Artificial Intelligence*, in UAI '25, vol. 286. Rio de Janeiro, Brazil: JMLR.org, Jul. 2025, pp. 3675–3700.
3. Akhare D, Luo T, Wang JX. Hybridndiff-uq: uncertainty quantification for hybrid neural differentiable modeling. Theoretical and Applied Mechanics Letters 16:100609 (2026). https://doi.org/10.1016/j.taml.2025.100609
4. Krättli M, Steinebach F, Morbidelli M. Online control

of the twin-column countercurrent solvent gradient process for biochromatography. Journal of Chromatography A 1293:51-59 (2013). https://doi.org/10.1016/j.chroma.2013.03.069
5. Müller?Späth T, Krättli M, Aumann L, Ströhlein G, Morbidelli M. Increasing the activity of monoclonal antibody therapeutics by continuous chromatography (MCSGP). Biotech & Bioengineering 107:652-662 (2010). https://doi.org/10.1002/bit.22843
6. Papathanasiou MM, Avraamidou S, Oberdieck R, Mantalaris A, Steinebach F, Morbidelli M, Mueller?Spaeth T, Pistikopoulos EN. Advanced control strategies for the multicolumn countercurrent solvent gradient purification process. AIChE Journal 62:2341-2357 (2016). https://doi.org/10.1002/aic.15203
7. Foteini Michalopoulou and Maria M. Papathanasiou, 'Assessment of data-driven modeling approaches for chromatographic separation processes'. Accessed: Apr. 28, 2025. [Online]. Available: https://aiche.onlinelibrary.wiley.com/doi/epdf/10.1002/aic.18600
8. D. P. Kingma, T. Salimans, and M. Welling, 'Variational Dropout and the Local Reparameterization Trick', in *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 2015. Accessed: Feb. 02, 2026. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2015/hash/bc7316929fe1545bf0b98d114ee3ecb8-Abstract.html
9. T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, 'Optuna: A Next-generation Hyperparameter Optimization Framework', Jul. 25, 2019, *arXiv*: arXiv:1907.10902. doi: 10.48550/arXiv.1907.10902.
10. J. Bradbury *et al.*, *JAX: composable transformations of Python+NumPy programs*. (2018). [Online]. Available: http://github.com/jax-ml/jax
11. I. Loshchilov and F. Hutter, 'Decoupled Weight Decay Regularization', arXiv.org. Accessed: Jan. 08, 2026. [Online]. Available: https://arxiv.org/abs/1711.05101v3
12. DeepMind *et al.*, *The DeepMind JAX Ecosystem*. (2020). [Online]. Available: http://github.com/google-deepmind
13. Matheson JE, Winkler RL. Scoring rules for continuous probability distributions. Management Science 22:1087-1096 (1976). https://doi.org/10.1287/mnsc.22.10.1087

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Multi-objective simulation-based optimisation of pharmaceutical process systems

Artemis Tsochatzidi[a], Francesca Cenci[b], Magdalini Aroniada[b] and Lazaros G. Papageorgiou[a]*

[a] The Sargent Centre for Process Systems Engineering, Department of Chemical Engineering, UCL (University College London), Torrington Place, London, UK
[b] GlaxoSmithKline (GSK) Research and Development, Park Road, Ware, UK
* Corresponding Author: l.papageorgiou@ucl.ac.uk

ABSTRACT

The pharmaceutical industry is placing growing emphasis on sophisticated process modeling to enhance the efficiency of drug design and production pipelines. Optimal control over these models can significantly improve manufacturing performance by lowering costs, boosting productivity, and ensuring rigorous quality compliance. However, the intricate nature and heavy computational load of these models often require the adoption of more practical or simplified alternative strategies for optimisation such as simulation-based approaches. In this work, we introduce a simulation-based framework including a "top-level" gradient-based mathematical programming optimisation model coupled with a "low-level" simulation scheme, to optimise multi-scale drug substance manufacturing flowsheets. The proposed framework optimises critical quality attributes, such as yield and purity, including green metrics such as process mass intensity, aligning with digital platforms (e.g. gPROMS) used in the pharmaceutical industry. By constructing Pareto fronts, we capture the balance between conflicting objectives, providing valuable insight for informed decision-making. The results indicate that this simulation-based approach can effectively optimise complex process behaviors while improving computational efficiency as no preliminary data is needed. By applying the framework to a real-world case study, we demonstrate its potential in optimising complex pharmaceutical manufacturing flowsheets.

Keywords: Optimisation, Simulation, Algorithms, Industry 4.0, Process Operations

INTRODUCTION

The pharmaceutical industry has undergone substantial transformation over the past decades, driven in large part by advances in process modelling that have accelerated drug development and manufacturing. Mechanistic flowsheet models are often highly complex, comprising tightly coupled unit operations, nonlinear dynamics, and numerous decision variables [3, 4]. Optimising such flowsheets with traditional optimisation methods, such as gradient-based approaches, is therefore challenging. To address these challenges, different kinds of methods such as surrogate-based [2] or simulation-based techniques [1] have been developed, enabling the optimisation of highly complex flowsheets with increased dimensionality.

Surrogate-based optimisation methods require *a priori* data to train surrogate models that approximate the underlying process behaviour, often necessitating substantial offline computational time. In contrast, simulation-based optimisation approaches do not require an explicit surrogate model to be trained in advance, instead relying directly on repeated evaluations of the high-fidelity process model during the optimisation procedure. However, these heuristic approaches often suffer from stochasticity, or limited convergence guarantees [1, 7].

Motivated by the need to combine data-free simulation-based optimisation with the advanced optimisation capabilities of surrogate-based methods, we propose a simulation-optimisation framework for pharmaceutical manufacturing flowsheets that integrates with the GAMS optimisation language (via gamspy) and aligns with commonly used software in the pharmaceutical industry (e.g. gPROMS, via gO:Python).

The rest of the manuscript is structured as follows. In the methodology section, the general workflow and mathematical model of the "top-level" mathematical programming component are presented, followed by the

computational methodology section that presents the proposed algorithm. The drug substance manufacturing Case Study is then presented while the results of the application of the proposed methodology on the Case Study are latter discussed. Concluding remarks are highlighted at the end of the manuscript.

METHODOLOGY

The objective of the proposed framework is to optimise the operating conditions of complex pharmaceutical manufacturing processes using a simulation–optimisation strategy. A mechanistic flowsheet model representing the pharmaceutical manufacturing process was developed internally and its predictive capability was validated against experimental data. Although the experimental data cannot be disclosed for confidentiality reasons, the validated flowsheet model is considered representative of the process and can therefore be used to generate reliable in silico experiments. The optimisation objective is to determine the operating conditions of the manufacturing process, such as temperatures and solution volumes, that achieve multi-objective optimisation of key performance indicators including yield, purity, and sustainability metrics. However, the mechanistic flowsheet model is highly complex and computationally intensive, which makes direct optimisation intractable.

To address this challenge, we propose a simulation–optimisation framework that combines a high-fidelity simulator with a gradient-based mathematical programming model. The mechanistic flowsheet model is repeatedly queried to evaluate the values of the critical quality attributes and their sensitivities with respect to the decision variables. At each iteration, these quantities are extracted from the simulator and treated as fixed parameters in the top-level Linear Programming (LP) optimisation model presented below [5]. The framework is formulated in a general manner and can therefore be applied to any mechanistic flowsheet model.

Nomenclature

Indices	
m	Features ($m = m_1, m_2, \ldots, M$)
k	Iterations ($k = k_1, k_2, \ldots, K$)
r	Responses ($r = r_1, r_2, \ldots, R$)
Sets	
V	Set including principal response *r*
$\bar{O}$	Set of responses *r* to be maximised
$\underline{O}$	Set of responses *r* to be minimised
Parameters	
T_r	Threshold level for response *r*
$\bar{Y}_{kr}(X'_{mk})$	Value of response *r* obtained by simulation at iteration *k*
X'_{mk}	Value of X_m used in simulation iteration *k*
$\partial\bar{Y}_r / \partial X'_{mk}$	Partial derivative of Y_r with respect to X_m obtained at iteration *k*
μ	Parameter determing the explored region
α_r	Weight for response *r* in the objective function
Variables	
X_m	Optimal value of feature *m*
Y_r	Objective value of response *r*
δ_r	Slack variable for response *r*

Constraints

The objective value of response *r* is approximated using a first-order Taylor expansion of the solution obtained at the previous iteration *k-1* (Eqs. (1, 2)).

$$Y_r \geq \bar{Y}_{kr}(X'_{mk}) + \sum_m \frac{\partial \bar{Y}_r}{\partial X'_{mk}}(X_m - X'_{mk})\ \forall\, r \in \underline{O}, k = k-1 \quad (1)$$

$$Y_r \leq \bar{Y}_{kr}(X'_{mk}) + \sum_m \frac{\partial \bar{Y}_r}{\partial X'_{mk}}(X_m - X'_{mk})\ \forall\, r \in \bar{O}, k = k-1 \quad (2)$$

The objective value of a non-primary response *r* to be minimised needs to be below a certain threshold with the addition of a slack variable δ_r (Eq. (3)), that is minimised in the objective function.

$$Y_r \leq T_r + \delta_r\ \forall\, r \in \underline{O} \setminus V \quad (3)$$

The objective value of a non-primary response *r* to be maximised needs to be above a certain threshold with the extraction of a slack variable δ_r (Eq. (3)), that is minimised in the objective function.

$$Y_r \geq T_r - \delta_r\ \forall\, r \in \bar{O} \setminus V \quad (4)$$

For the Taylor expansion to be valid as well as to prevent the optimisation from extrapolating too far from the previous iteration, the concept of trust region is used [5]. The trust region is a local domain that the optimisation is allowed to explore. A small μ (10%) is chosen to determine the range of the region explored (Eq. (5)).

$$X'_{m,k-1}(1-\mu) \leq X_m \leq X'_{m,k-1}(1+\mu)\ \forall m \quad (5)$$

The slack variable δ_r needs to be positive for every response r.

$$\delta_r \geq 0\ \forall\, r \in R \quad (6)$$

Objective function

In the objective function, a principal response *r* is optimised, while the remaining responses are incorporated through the slack variables δ_r (Eq. (7)).

$$\max_{X_m} \sum_{r \in V} Y_r - \sum_{r \in R \setminus V} a_r \bullet \delta_r \quad (7)$$

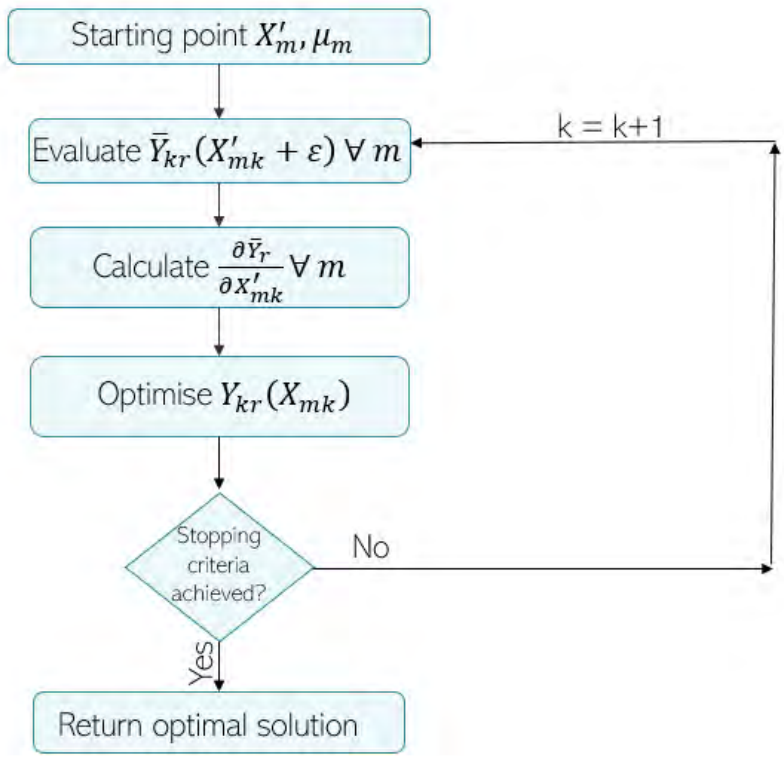

Figure 1. Proposed algorithm.

COMPUTATIONAL METHODOLOGY

Figure 1 illustrates the proposed computational methodology. Starting from a random initial point, a mechanistic flowsheet model is used to evaluate the values of the objectives. By introducing a small perturbation (ε) and re-performing the simulation for every perturbated feature m, the derivatives $\partial \bar{Y}_r / \partial X'_{mk}$ can be estimated using the finite differences method (Eq. 8).

$$\frac{\partial \bar{Y}_{kr}}{\partial X'_{mk}} = \frac{\bar{Y}_{kr}(X'_{mk}+\varepsilon) - \bar{Y}_{kr}(X'_{mk})}{\varepsilon}\ \forall r, m, k \tag{8}$$

The information is then used in the Linear Programming (LP) mathematical model described in the methodology section from which the optimal values of the decision are decided. These are fed back into the flowsheet for the next simulation, thereby forming an iterative simulation–optimisation loop that continues until convergence criteria are satisfied. The iterative procedure is terminated when at least one of the following stopping criteria is satisfied: (i) a prescribed maximum number of iterations is reached or (ii) over T consecutive iterations, the changes in all response values remain below a specified tolerance, implying practical convergence of the solution.

CASE STUDY

The proposed framework is designed to efficiently optimise complex industrial processes and is demonstrated through a pharmaceutical manufacturing Case Study. The system model represents an API production flowsheet (Figure 2), comprising a dynamic batch reactor, a liquid–liquid extraction unit, and a crystallisation stage. In the reaction step, an exothermic process partially dissolves suspended solids, forming two liquid phases: a product-rich organic layer and an impurity-rich aqueous layer. The phases are separated, followed by repeated washing of the aqueous phase. Product recovery is completed through solvent addition, seeding, and controlled cooling during crystallisation. The integrated flowsheet model consists of over 25, 000 equations, including both algebraic and differential components. A total of 22 decision variables spanning reaction, extraction, and crystallisation stages are optimised with respect to Yield, Impurity D and Process Mass Intensity (PMI) [6]. The decision variables are distributed across the main stages of the manufacturing process. Specifically, 11 variables correspond to the reaction stage (A1–A4, B1–B5, E1–E2), 3 variables are associated with the extraction stage (C1–C3), and 8 variables relate to the crystallisation stage (D1–D8). These variables represent key operating conditions of the flowsheet, including quantities of solvents or reactants, operating temperatures, addition or residence times, and mixing conditions. For more details on the Sensitivity Analysis for the Case Study the reader is

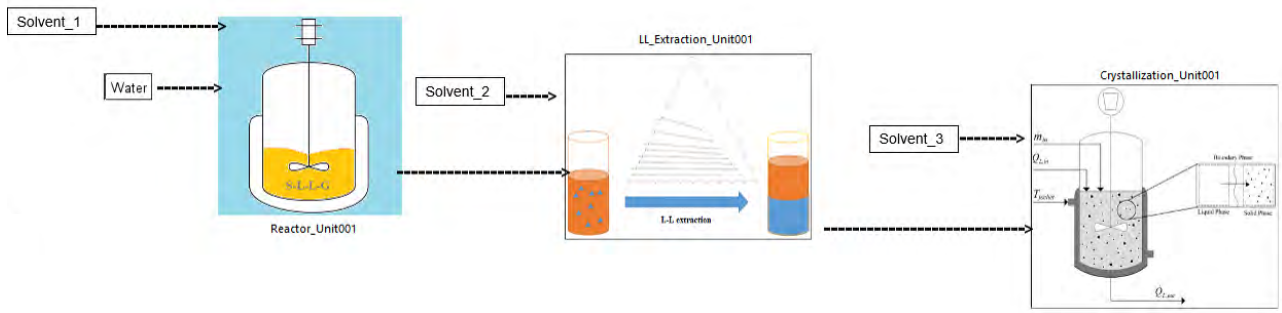

Figure 2. Active Pharmaceutical Ingredient Case Study (API)

referred to [8, 9].

RESULTS AND DISCUSSION

The proposed approach has been applied to the Case Study presented in the previous section. The goal of the multi-objective optimisation was to maximise Yield while keeping the Impurity D levels below a threshold ($T_{ImpurityD} \leq 0.0215\ \%w/w$). The maximum number of iterations is set to 50, with the convergence window $T = 3$ consecutive iterations.

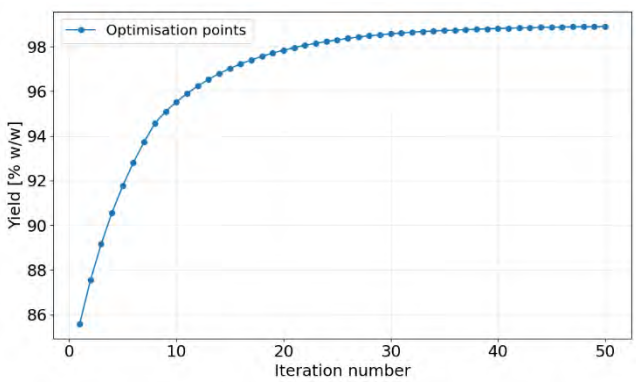

Figure 3.a: Yield [% w/w] vs Iteration number

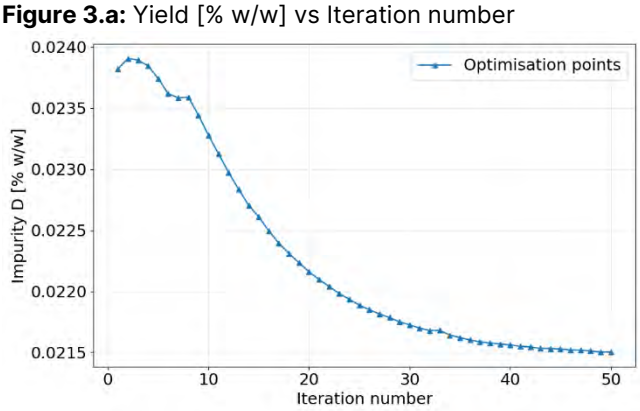

Figure 3.b: Impurity D [% w/w] vs Iteration number

Figures 3a and 3b demonstrate the profile of the validated $\bar{Y}_r$ from the mechanistic flowsheet model for responses Yield and Impurity D respectively. Starting from a random sample (Yield ≈ 86 % w/w, Impurity D ≈ 0.024 %w/w), the proposed algorithm managed to get to a value of Yield ≈ 98.9% w/w while respecting the Impurity D threshold. Yield reaches its maximum value early in the applied algorithm, whereas Impurity D requires the full iteration limit to satisfy the convergence threshold.

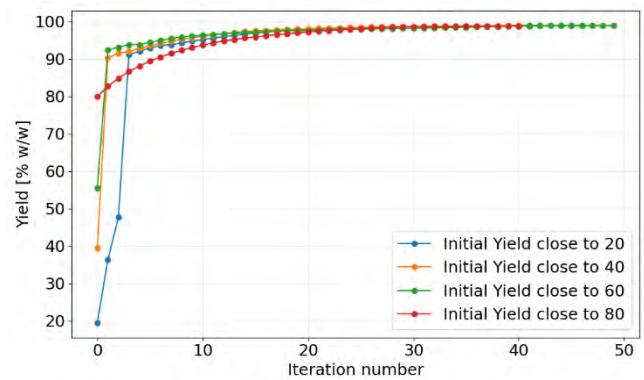

Figure 4.a: Yield [% w/w] vs Iteration number for different initial samples

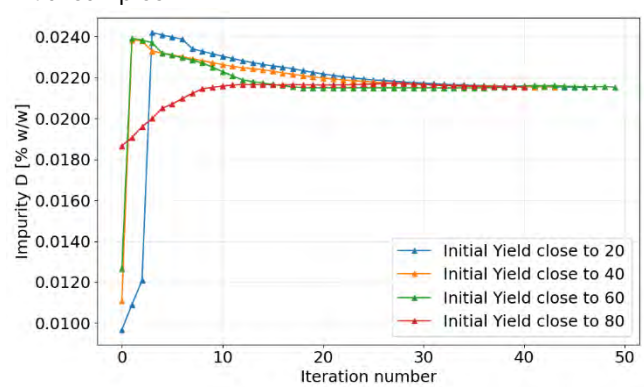

Figure 4.b: Impurity D [% w/w] vs Iteration number for different initial samples

To demonstrate the effectiveness of the algorithm in different initial points, we choose 4 representative scenarios to start the proposed approach. Those scenarios correspond to initial samples with Yield values close to

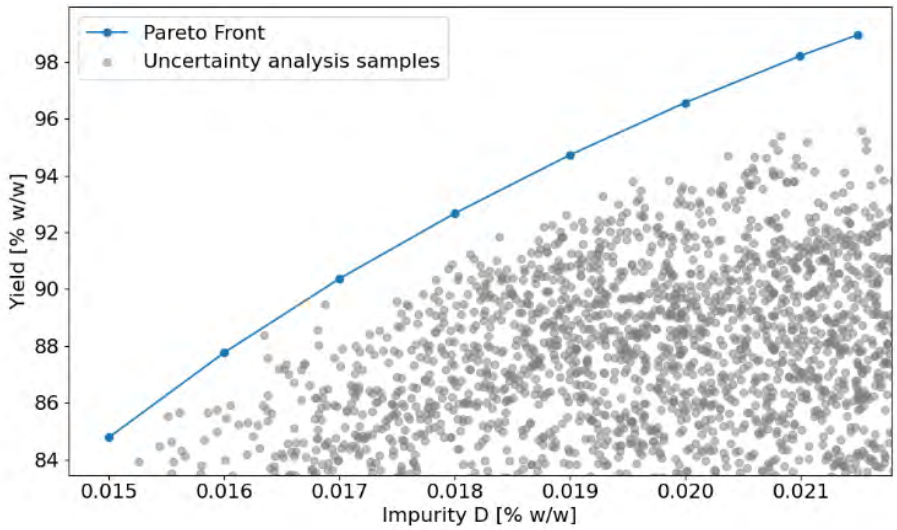

Figure 5.a: Pareto Front: Yield [% w/w] vs Impurity D [%w/w]

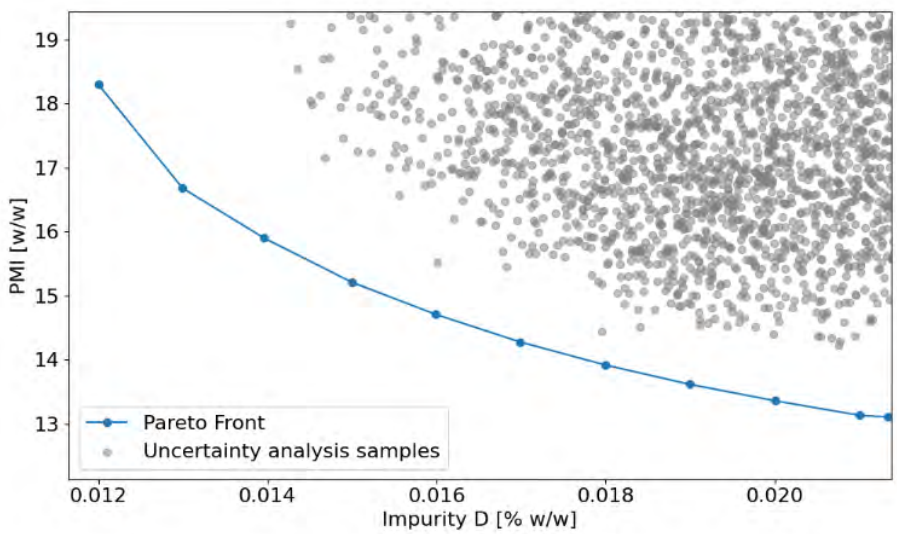

Figure 5.b: Pareto Front: PMI [w/w] vs Impurity D [%w/w]

20, 40, 60 and 80 %w/w. As depicted in Figures 4a and 4b, the algorithm can obtain almost identical solutions when terminated, with the number of iterations needed varying between 40 to 50.

Motivated by industrial requirements to evaluate trade-offs between competing objectives, we apply the proposed methodology to generate Pareto fronts that map the space of optimal compromises. These fronts provide actionable insight into objective interactions, enabling more informed decisions. By adjusting the threshold for Impurity D and applying the proposed approach for multiple values, we can derive a spectrum of optimal validated Yield values with respect to Impurity D thresholds. This is demonstrated in Figure 5a where the gray samples result from an uncertainty analysis carried out in gPROMS software using quasi-random Sobol sampling over 5, 000 uncertainty scenarios. The blue points are the samples derived from the proposed approach and represent the limits between the feasible and non-feasible space.

Figure 5b illustrates the results of the proposed approach applied to a different optimisation scheme: Minimisation of PMI while keeping Impurity D up to a threshold. These objectives are competing (as the case of Yield and Impurity D) meaning that it is difficult to have a low value of PMI along with high purity in the same sample. However, the proposed approach managed to achieve up to 33.8% improved samples compared to the ones obtained from the uncertainty analysis.

Each point on the Pareto front represents a validated operating condition obtained from the mechanistic flowsheet model. These points correspond to feasible

solutions where improvement in one objective cannot be achieved without deteriorating at least the other. These fronts serve as a valuable tool for decision-making, to understand the limits imposed by the process. The final selection of an operating point depends on the preferences of the decision maker, such as the relative importance assigned to product quality, process efficiency, or sustainability metrics.

CONCLUDING REMARKS

In this study, we proposed a simulation-optimisation framework for complex system models in the pharmaceutical industry. The framework integrates mathematical programming principles with a simulation-based approach to enable systematic exploration of process trade-offs. The applicability of the proposed approach is demonstrated by an industrial case study characterised by an existing mechanistic flowsheet model. The results demonstrate that the approach can efficiently identify Pareto-optimal operating conditions for multi-objective optimisation.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the industrial partner GlaxoSmithKline (GSK) for provided data and financial support of this project.

AUTHOR IDENTIFIERS

Author ORCIDs:
Tsochatzidi A.: 0000-0001-5861-8303
Cenci F.: 0000-0002-3360-0470
Aroniada M.: 0009-0008-9273-4314
Papageorgiou L.G: 0000-0003-4652-6086

REFERENCES

1. Amaran S, Sahinidis NV, Sharda B, Bury SJ. Simulation optimization: a review of algorithms and applications. Ann Oper Res 240:351-380 (2015). https://doi.org/10.1007/s10479-015-2019-x
2. Boukouvala F, Ierapetritou MG. Surrogate-based optimization of expensive flowsheet modeling for continuous pharmaceutical manufacturing. J Pharm Innov 8:131-145 (2013). https://doi.org/10.1007/s12247-013-9154-1
3. Christodoulou C, Sorensen E, García-Muñoz S, Mazzei L. Mathematical modelling of water absorption and evaporation in a pharmaceutical tablet during film coating. Chemical Engineering Science 175:40-55 (2018). https://doi.org/10.1016/j.ces.2017.09.021
4. Diab S, Christodoulou C, Taylor G, Rushworth P. Mathematical modeling and optimization to inform impurity control in an industrial active pharmaceutical ingredient manufacturing process. Org. Process Res. Dev. 26:2864-2881 (2022). https://doi.org/10.1021/acs.oprd.2c00208
5. Gjerdrum, J., Samsatli, N.J., Shah, N., Papageorgiou, L. G. Optimisation of policy parameters in supply chain applications. Int. J. Logist. Res. Appl. 8:15–36 (2005)
6. Jimenez-Gonzalez C, Ponder CS, Broxterman QB, Manley JB. Using the right green yardstick: why process mass intensity is used in the pharmaceutical industry to drive more sustainable processes. Org. Process Res. Dev. 15:912-917 (2011). https://doi.org/10.1021/op200097d
7. Laky DJ, Casas-Orozco D, Laird CD, Reklaitis GV, Nagy ZK. Simulation–optimization framework for the digital design of pharmaceutical processes using pyomo and pharmapy. Ind. Eng. Chem. Res. : (2022). https://doi.org/10.1021/acs.iecr.2c01636
8. Tsochatzidi A, Cenci F, Aroniada M, Papageorgiou LG. A novel surrogate-based optimisation framework for pharmaceutical process systems. International Journal of Pharmaceutics 682:125861 (2025). https://doi.org/10.1016/j.ijpharm.2025.125861
9. Tsochatzidi A, Liapis GI, Cenci F, Aroniada M, Papageorgiou LG. Surrogate-based multi-objective optimisation via tree regression. Computers & Chemical Engineering 207:109538 (2026). https://doi.org/10.1016/j.compchemeng.2025.109538

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Modeling, Simulation, and Optimization of an Anion Exchange Membrane Cell for Ammonia Electrolysis

Laureen E. Hernández Lefrán, Leonardo A. Cáceres Avilez, Antonio E. Bresciani, Claudio A. Oller do Nascimento, Rita M. Brito Alves*

Universidade de São Paulo, Escola Politécnica, Department of Chemical Engineering, São Paulo, São Paulo, Brazil
* Corresponding Author: rmbalves@usp.br

ABSTRACT

The need to reduce greenhouse gas emissions and diversify energy sources has driven the development of hydrogen (H_2) as a leading carbon-free energy vector with high gravimetric energy density (33.3 kWh/kg) for stationary and transport applications. However, its storage and transportation remain challenging. Therefore, ammonia emerges as a promising hydrogen carrier due to its high energy density and ease of liquefaction and transport. This work presents the development of a phenomenological based mathematical model for an ammonia electrolysis cell operating in a zero-gap configuration with an anion exchange membrane. The model considers the contributions of different overpotentials (activation, ohmic, concentration) and incorporates empirical and semi-phenomenological expressions adjusted to experimental data from Zhang et al [1]. The model shows both qualitative and quantitative agreement with experimental measurements, achieving a determination coefficient (R^2) of 0.9874. Beyond electrochemical modeling, the cell model is integrated into a nonlinear programming (NLP) optimization framework to minimize the total annualized cost (TAC) of the system. For hydrogen production of 1250 kg/h, the optimal solution corresponds to 353 K, 0.82 A/cm^2, and 1302 cells, yielding a TAC of about 45.0 MUSD/year.

Keywords: Hydrogen production, Ammonia electrolysis, Zero gap cell, Mathematical model, Optimization

INTRODUCTION

Hydrogen (H_2) is widely regarded as a promising energy vector due to its high gravimetric energy density and carbon-free end use [2]. However, its low volumetric energy density and wide flammability range pose challenges for storage and long-distance transportation, requiring energy-intensive solutions such as high-pressure compression or cryogenic liquefaction [2-3].

A promising alternative to overcome these limitations involves storing hydrogen in chemical compounds that can be easily transported and stored, with hydrogen released on demand through chemical or electrochemical processes. Among the available hydrogen carriers, ammonia (NH_3) stands out due to its high hydrogen content (17 wt.%), relatively mild liquefaction conditions (~10 bar at ambient temperature), and well-established global infrastructure [4-5].

In this context, ammonia electrolysis has emerged as an attractive pathway for hydrogen production, offering a significant thermodynamic advantage over conventional water electrolysis [5]. Recent studies highlight that anion exchange membrane (AEM) electrolyzers are particularly suitable for ammonia electrolysis due to their chemical compatibility with NH_3 and their ability to operate under alkaline conditions. Furthermore, zero-gap cell configurations have been identified as the most promising architecture, as they minimize internal ohmic losses while enabling compact and scalable designs for industrial deployment [6]. However, despite the growing number of experimental studies, the availability of mathematical models for system-level analysis and scale-up remains limited.

The present work addresses this gap by developing a mathematical model of an anion exchange membrane ammonia electrolyzer within a compact zero-gap configuration.

The Methods section describes the electrolyzer configuration, the modeling approach and assumptions, the parameter estimation procedure, and the formulation

of the optimization problem, minimizing the total annualized cost (TAC). The model was parameterized using experimental data from the literature and optimized using nonlinear programming (NLP) to minimize the total annualized cost (TAC), evaluating the trade-off between the number of cells and the applied current density for large-scale hydrogen production. The Results section presents the model fitting, followed by the optimal operating and design conditions obtained from the TAC minimization. Finally, the main findings and implications are summarized in the Conclusions.

METHODS

A typical anion exchange membrane (AEM) electrolyzer cell consists of an anion exchange membrane sandwiched between an anode and a cathode, both comprising catalyst and gas diffusion layers, as well as bipolar plates containing engraved flow channels (Figure 1).

In the present system, an aqueous KOH solution is supplied to the cathode, while a mixture of NH_3 and KOH is fed to the anode. At the cathode, water is reduced to hydrogen and hydroxide ions via the hydrogen evolution reaction. The hydroxide ions migrate through the AEM toward the anode, where ammonia electro-oxidation produces nitrogen and water. The generated gaseous products are removed through the respective flow channels.

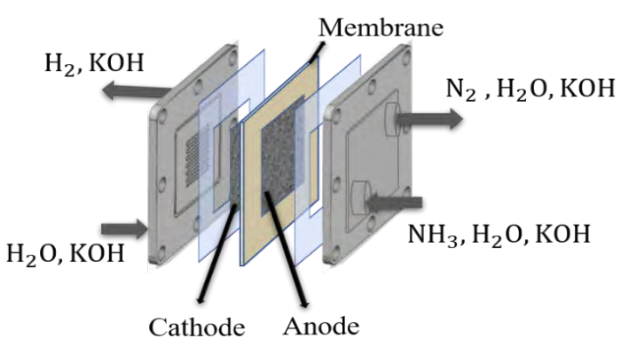

Figure 1. Conceptual schematic of the anion exchange membrane ammonia electrolysis cell.

Mathematical modeling

In contrast to previously reported ammonia electrolysis models that focus either on detailed reaction mechanisms or highly resolved transport phenomena [7-8], the present work adopts a physics-based zero-dimensional modeling approach that is suitable for system-level analysis and optimization. The electrolyzer is treated as a lumped system and spatial gradients of temperature, concentration, and current density are neglected. The model formulation is based on global electrochemical relations and simplified thermal and pressure descriptions, under the following assumptions:

1. Steady-state operation at constant temperature and pressure;
2. Catalyst layers are treated as interfacial surfaces;
3. Ohmic losses are described using Ohm's law;
4. The electrolyte is an aqueous KOH solution;
5. Identical active areas and uniform current density at both electrodes;
6. Gaseous species behave as ideal gases and are uniformly distributed;
7. Uniform pressure along gas flow channels;
8. Electrolyte ohmic resistance in the electrode gap is neglected due to the zero-gap configuration;
9. Gas bubble coverage effects are not explicitly modeled, as forced electrolyte flow and the zero-gap architecture promote rapid bubble detachment.

The model builds upon established formulations previously developed for anion exchange membrane water electrolysis (AEMWE) [9–12] and alkaline water electrolysis (AWE) [13-15].

The cell voltage is calculated as the sum the reversible cell potential (V_{rev}) and the different overpotential contributions, namely activation (η_{act}), concentration (η_{con}), and ohmic (η_{ohm}) overpotentials, as expressed in Equation 1.

$$V_{cell} = V_{rev} + \eta_{act} + \eta_{con} + \eta_{ohm} \quad (1)$$

The reversible cell potential is calculated using the Nernst equation as expressed in Equation 2:

$$V_{rev} = V_{rev}^0 + \frac{R \cdot T}{z \cdot F} \ln \frac{a_{H_2}^3 \cdot a_{N_2}}{a_{NH_3}^2} \quad (2)$$

Where R is the universal gas constant, F is Faraday's constant, a_k denotes the thermodynamic activity of species k, and z is the number of electrons transferred in the overall ammonia electrolysis reaction. The standard reversible potential (V_{rev}^0) is obtained from the semi-empirical correlation proposed by Estejab et al. [16]. The thermodynamic activities of hydrogen and nitrogen are assumed to be equal to their respective partial pressures, which are calculated following the procedure described in [13].

The activity of ammonia is assumed to be equal to that of the alkaline electrolyte composed solely of KOH, since the ammonia concentration employed is relatively low. The activity of water in aqueous KOH solutions is calculated according to the correlation reported by Balej [17].

The activation overpotential is expressed as the sum of the anodic and cathodic activation contributions, derived from the Butler-Volmer formulation, and written in logarithmic form as Equation 3:

$$\eta_{act} = \frac{2.3 \cdot R \cdot T}{z \cdot \alpha_{cat} \cdot F} \log \frac{i}{i_{cat}^0} + \frac{2.3 \cdot R \cdot T}{z \cdot \alpha_{ano} \cdot F} \log \frac{i}{i_{ano}^0} \quad (3)$$

where i is the operating current density, i_{cat}^0 e i_{ano}^0 are the exchange current densities at the cathode and anode, respectively, and α_{cat} and α_{ano} are the corresponding charge transfer coefficients. The exchange current densities are modeled as temperature-dependent quantities

following an Arrhenius-type relationship (Equations 4-5):

$$i^0_{cat} = i^0_{cat,ref} \bullet \exp\left[-\frac{\Delta G_{cat}}{R} \bullet \left(\frac{1}{T} - \frac{1}{T_{ref}}\right)\right] \quad (4)$$

$$i^0_{ano} = i^0_{ano,ref} \bullet \exp\left[-\frac{\Delta G_{ano}}{R} \bullet \left(\frac{1}{T} - \frac{1}{T_{ref}}\right)\right] \quad (5)$$

where $i^0_{cat,ref}$ and $i^0_{ano,ref}$ represent the exchange current densities at reference temperature T_{ref}, and ΔG_{cat} and ΔG_{ano} denote the activation free energies associated with the cathodic and anodic reactions, respectively.

In the present model, the charge transfer coefficients α_{cat} and α_{ano} are assumed to be temperature-dependent. Similar behavior has been reported in previous electrochemical modeling studies [9–11]. In this work, empirical correlations as functions of temperature were obtained by regression of experimental polarization data and are described by third-order polynomial expression. The nonlinear temperature dependence of α observed for ammonia electrolysis contrasts with the commonly assumed linear behavior in water electrolysis and reflects the higher mechanistic complexity of the ammonia oxidation reaction.

The concentration overpotential is derived from the Nernst equation combined with Fick's law to account for mass transport limitations associated with gas diffusion through the porous electrodes. Following an approach from Chang et al. [11], the concentration overpotential is expressed as Equation 6:

$$\eta_{conc} = \frac{2.3 \bullet R \bullet T}{n_{cat} \bullet F} \log \frac{C_{H_2,cat}}{C_{H_2,low}} + \frac{2.3 \bullet R \bullet T}{n_{ano} \bullet F} \log \frac{C_{N_2,ano}}{C_{N_2,low}} \quad (6)$$

where $C_{N_2,ano}$ and $C_{H_2,cat}$ are the concentrations of hydrogen and nitrogen, respectively, at the membrane–electrolyte interface, while $C_{N_2,low}$ and $C_{H_2,low}$ denote the reference concentrations at working conditions. The parameters $n_{cat} = 2$ and $n_{ano} = 6$, represent the number of electrons required to produce one mole of H_2 and N_2, respectively.

The ohmic overpotential accounts for electrodes and membrane resistances within the electrolyzer and is calculated by modeling the different cell components as electrical resistances connected in series. The effective cell resistance (R_{cell}) is obtained as the sum of the electrode resistances ($R_{electrodes}$), and the separator resistance, which in the present case corresponds to the anion exchange membrane ($R_{membrane}$). Accordingly, the ohmic overpotential is given by Ohm's law (Equation 7):

$$\eta_{ohm} = I \bullet R_{cell} \quad (7)$$

The total electrode resistance is calculated according to the Equation 8:

$$R_{electrodes} = \rho_{ef,cat} \bullet \frac{\delta_{cat}}{A_e} + \rho_{ef,ano} \bullet \frac{\delta_{ano}}{A_e} \quad (8)$$

where $\rho_{ef,cat}$ and $\rho_{ef,ano}$ are the effective resistivities of the cathode and anode, respectively, δ_{cat} and δ_{ano} are the corresponding electrode thicknesses, and A_e is the active electrode area. The expressions used to calculate the effective resistivities are taken from the correlations reported in [13].

The membrane resistance is given by Equation 9:

$$R_m = \frac{\delta_m}{\sigma_m \bullet A_m} \quad (9)$$

where δ_m is the membrane thickness, A_m is the active membrane area, and σ_m is the ionic conductivity of the membrane material.

Experimental data and parameter estimation

Experimental polarization data reported by Zhang et al. [1] were used to estimate the model parameters. The experiments were conducted in a zero-gap membrane electrode assembly (MEA) electrolyzer with an active area of 5 cm^2, operating between 323 and 363 K. The anode was supplied with an aqueous solution containing 3 M KOH and 3 M NH_3, while the cathode was fed with 3 M KOH, both at a flow rate of 3 mL/min and atmospheric pressure.

Model parameters, including anodic and cathodic exchange current densities, activation energies, and membrane thickness, were estimated through nonlinear least-squares regression. Parameter estimation was performed using a routine developed in Python (Python Software Foundation, USA), employing the curve_fit function from the SciPy library together with NumPy [18]. The optimization was carried out using the Trust-Region Reflective algorithm.

Model performance metrics

Model performance was evaluated by calculating the determination coefficient (R^2), mean absolute error (MAE), mean absolute percentage error (MAPE), and root mean square error (RMSE). Parameter uncertainty and identifiability was assessed using the bootstrap method, confirming the robustness of the estimated parameters.

Optimization framework

To evaluate the techno-economic performance of the zero-gap ammonia electrolyzer and its scale-up to plant level, the developed electrochemical model was embedded into a nonlinear optimization framework to minimize the total annual cost (TAC), implemented in Python using the Pyomo package (version 6.10.0) [19]. The formulation accounts for electrochemical, thermal, and economic trade-offs, enabling the identification of optimal operating and design conditions.

The TAC (Equation 10) accounts for both capital and operating expenditures associated with the electrolyzer system.

Capital expenditure ($CAPEX$, Equation 11) was assumed to be proportional to the number of cells (N), the cost of a cell per area (C_{cell}), and the cell area (A_{cell}).

Operating expenditure (OPEX, Equation 12) includes

electricity and steam consumption. The capital recovery factor (CRF, Equation 13) was calculated assuming a plant lifetime (n) of 20 years and interest rate (j) of 10%. The parameters used to calculate the expenditures are list in Table 1.

$$\mathrm{TAC} = \mathrm{CRF} \bullet \mathrm{CAPEX} + \mathrm{OPEX} \tag{10}$$

$$\mathrm{CAPEX} = \mathrm{N} \bullet \mathrm{C_{cell}} \bullet \mathrm{A_{cell}} \tag{11}$$

$$\mathrm{OPEX} = \mathrm{C_{elect}} \bullet \mathrm{P} \bullet \mathrm{h_{op}} + \mathrm{C_{steam}} \bullet \mathrm{Q_t} \bullet \mathrm{h_{op}} \tag{12}$$

$$\mathrm{CRF} = \frac{\mathrm{j(1+j)^n}}{\mathrm{(1+j)^n - 1}} \tag{13}$$

where C_{elect} is the electricity cost (USD/MWh), P is the total electrical power (MW), h_{op} is the annual operating time (h/year), C_{steam} is the steam cost (USD/GJ) and Q_t is the total thermal demand (GJ).

The thermal energy required to heat the electrolyte was estimated using a surrogate model derived from process simulations performed in Aspen Plus v.12, using a temperature-dependent polynomial correlation. These simulations were used to calculate the specific heat requirement per unit mass of electrolyte solution, Q_{spec}(GJ/kg), under representative operating conditions.

Within the optimization framework, the total thermal demand was then obtained by scaling the specific heat requirement with the electrolyte mass flow per cell and the total number of cells, Equation 14:

$$Q_t = Q_{spec} \bullet \dot{m}_{sol} \bullet N \tag{14}$$

where $\dot{m}_{sol}$ is the electrolyte mass flow rate per cell.

Table 1. Main parameters of the optimization problem.

Parameter	Value	Reference
$\mathrm{h_{op}}$ (h/year)	8000	Adopted
$\mathrm{C_{cell}}$ (USD/m^2)	16200	[20]
$\mathrm{C_{steam}}$ (USD/GJ)	12.85	[21]
$\mathrm{C_{elect}}$ (USD/MWh)	136.95	[21]
$\mathrm{A_{cell}}$ (m^2)	3.14	[1]

The main decision variables of the optimization problem are:

- Current density, i (A/cm^2);
- Number of cells, N;
- Operating temperature, T (K).

These variables simultaneously control the electrochemical performance of the cell and the overall plant sizing, directly affecting hydrogen production, energy consumption, and total system cost.

The optimization problem was formulated as the minimization of TAC, subject to electrochemical, physical, thermal, operational, and hydrogen production constraints, as expressed in Equations 15-20.

$$\min \mathrm{TAC}, (\text{Equation 10}) \tag{15}$$

$$\text{s. t. cell voltage equation}, Eq.\,(1) \tag{16}$$

$$\mathrm{H_2}\ \text{production}: \frac{\mathrm{I}}{2 \bullet \mathrm{F}} MM_{H_2} \geq 1250\ \mathrm{kg/h} \tag{17}$$

$$\mathrm{T^l} \leq \mathrm{T} \leq \mathrm{T^u} \tag{18}$$

$$\mathrm{N^l} \leq \mathrm{N} \leq \mathrm{N^u} \tag{19}$$

$$\mathrm{i^l} \leq \mathrm{i} \leq \mathrm{i^u} \tag{20}$$

where T is the operating temperature of the electrochemical cell, which was allowed to vary between $\mathrm{T^l}$=323 K and $\mathrm{T^u}$=353 K. The current density i was evaluated in the range $\mathrm{i^l}$=0 to $\mathrm{i^u}$=2 A/cm^2, and the number of electrochemical cells was varied between $\mathrm{N^l}$=0 and $\mathrm{N^u}$=400 000.

A minimum hydrogen production target of 1250 kg/h was imposed to represent large-scale centralized production, consistent with the hydrogen demand of multiple refueling stations, where a single station typically consumes about 1000 kg/day of hydrogen.

The resulting nonlinear programming (NLP) problem was solved using the IPOPT solver [22], which is well suited for constrained nonlinear optimization with continuous decision variables.

RESULTS

Modeling results

The model parameters were estimated through nonlinear regression using the experimental polarization data. The regressed model was then compared with the experimental measurements to assess its ability to reproduce the observed electrochemical behavior.

Figure 2 presents the experimental polarization curves together with the corresponding model predictions at different operating temperatures between 323 and 353 K. Over this temperature range, the model reproduces the experimental cell voltages with good agreement across the investigated current density domain. At higher temperatures beyond the analyzed range, deviations are expected due to the reduced ammonia solubility at elevated temperatures, which affects mass transport phenomena.

Statistical goodness-of-fit indicators as R^2= 0.9874, and MAE=0.0126, MAPE= 2.72%, and RMSE = 0.0155, indicate a strong agreement between the regressed model and the experimental data within the range of operating

conditions considered.

The estimated parameters are reported in Table 2. The obtained values for the exchange current densities at reference temperature, fall within ranges commonly reported in the literature for alkaline electrolysis [9], [13], [23], varying between 10^{-3} and 10^{-11} A/cm², supporting the physical consistency of the proposed modeling approach. Similarly, the estimated activation free energies are close to values previously reported in the literature [23], [24].

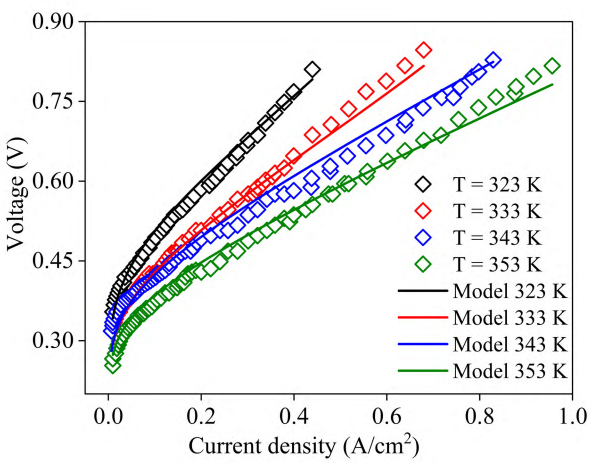

Figure 2. Comparison of experimental data with model predictions.

Table 2. Estimated model parameters from non-linear regression.

Parameters	Value
$i^0_{ano,ref}$ (A/cm²)	$7.64 \cdot 10^{-7}$
ΔG_{ano} (J/mol)	24935.61
$i^0_{cat,ref}$ (A/cm²)	$3.32 \cdot 10^{-5}$
ΔG_{cat} (J/mol)	16524.13
δ_m (cm)	0.0034

Optimization results

Solving the nonlinear optimization problem resulted in a unique set of operating and design conditions that minimize TAC while satisfying the minimum hydrogen production constraint. Figure 3 illustrates the variation of TAC with current density for the different operating temperatures considered. A well-defined minimum is observed at 353 K, indicating that this temperature provides the most favorable trade-off between improved electrochemical performance and increased thermal and operating costs.

At the optimal temperature of 353 K, the optimal current density was found to be 0.82 A/cm². This value represents a compromise between higher hydrogen production rates and the increase in activation, concentration, and ohmic overpotentials at elevated current densities. The resulting optimal cell voltage is approximately 0.81 V, leading to a total plant electrical power consumption of about 27.1 MW. The optimal current density and cell voltage obtained in this work are in good agreement with limits reported for ammonia electrolysis, which indicate that economically viable operation requires current densities close to 1 A/cm² while maintaining the applied cell potential below approximately 1 V [25].

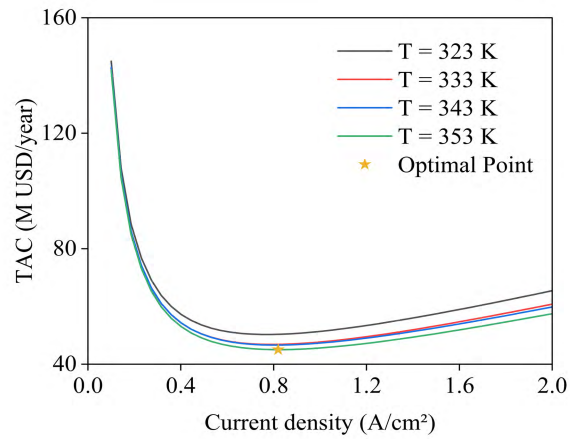

Figure 3. Variation of TAC at different temperatures.

The corresponding optimal number of cells was determined to be N = 1302, as illustrated in Figure 4 at 353 K.

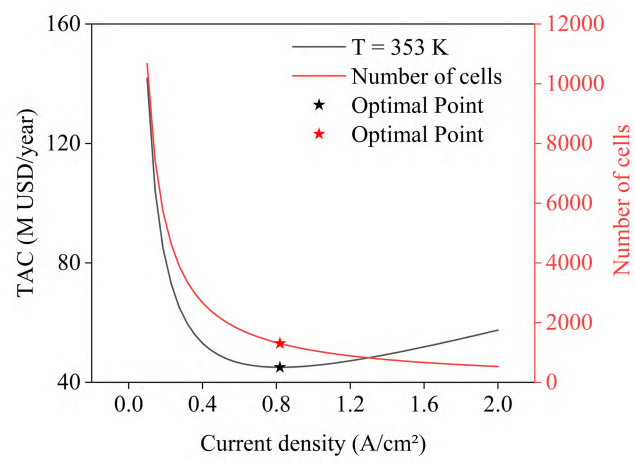

Figure 4. Optimization of TAC and number of cells at 353 K.

The estimated capital expenditure (CAPEX) under optimal conditions amounts to approximately 132.5 million USD, dominated by the installed electrochemical active area. The annual operating expenditure (OPEX) was calculated as approximately 29.5 million USD/year, primarily driven by electricity and steam consumption. The resulting minimum TAC is approximately 45.0 million

USD/year. Under these optimal conditions, the system achieves the target hydrogen production rate of 1250 kg/h. The estimated cost of hydrogen production is approximately 4.5 USD/kg H_2, excluding ammonia feedstock cost, and it falls within the range reported values for alkaline water electrolysis [26]. A complete techno-economic assessment including ammonia supply, storage, logistics, and sensitivity analysis is left for future work.

CONCLUSIONS

This work developed a mathematical model of a zero-gap AEM ammonia electrolyzer integrated into a techno-economic optimization framework. The model was fitted to experimental polarization data and describes the main factors affecting cell voltage, energy efficiency, and hydrogen production, showing good agreement (R^2 = 0.9874). For the base case, optimal configuration corresponds to a temperature of 353 K, a current density of 0.82 A/cm^2, and 1302 cells, resulting in a TAC of about 45.0 MUSD/year.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge Petronas and the São Paulo Research Foundation (FAPESP) (#2020/15230-5 and #2024/08858-9) for their financial support, ANP (Brazilian National Oil, Natural Gas, and Biofuels Agency) for the strategic importance of its support under the R&D levy regulation, and the Research Center for Greenhouse Gas Innovation (RCGI). The authors also thank the National Council for Scientific and Technological Development (CNPq) (#304346/2025-0). This study was financed in part by the Personnel Coordination for the Improvement of HigherLevel Personnel-Brazil (CAPES)-Finance Code 001.

AUTHOR IDENTIFIERS

Lefrán LEH: 0009-0006-9712-3454
Avilez LAC: 0000-0001-8320-863X
Bresciani AE: not available
Oller CAO: 0000-0001-6110-0249
Alves RMB: 0000-0002-1914-5133

REFERENCES

1. Zhang K, Han Y, Zhao Y, Wei T, Fu J, Ren Z, Xu X, Zhou L, Shao Z. Energy-efficient and cost-effective ammonia electrolysis for converting ammonia to green hydrogen. Cell Reports Physical Science 5:102171 (2024). https://doi.org/10.1016/j.xcrp.2024.102171
2. Mansilla C, Bourasseau C, Cany C, Guinot B, Le Duigou A, Lucchese P. Hydrogen applications: overview of the key economic issues and perspectives. Hydrogen Supply Chains :271-292 (2018). https://doi.org/10.1016/b978-0-12-811197-0.00007-5
3. Niaz S, Manzoor T, Pandith AH. Hydrogen storage: materials, methods and perspectives. Renewable and Sustainable Energy Reviews 50:457-469 (2015). https://doi.org/10.1016/j.rser.2015.05.011
4. Lucentini I, Garcia X, Vendrell X, Llorca J. Review of the decomposition of ammonia to generate hydrogen. Ind. Eng. Chem. Res. 60:18560-18611 (2021). https://doi.org/10.1021/acs.iecr.1c00843
5. Baeck J. S-injective modules. Rev. Real Acad. Cienc. Exactas Fis. Nat. Ser. A-Mat. 118: (2023). https://doi.org/10.1007/s13398-023-01514-7
6. Serrano-Jiménez J, Martín C, Pinzón M, Sánchez P, de la Osa AR. Exploring the potential of ammonia electrolysis for hydrogen production: from lab-performance to stack architectures. Current Opinion in Chemical Engineering 51:101204 (2026). https://doi.org/10.1016/j.coche.2025.101204
7. Estejab A, Daramola DA, Botte GG. Mathematical model of a parallel plate ammonia electrolyzer for combined wastewater remediation and hydrogen production. Water Research 77:133-145 (2015). https://doi.org/10.1016/j.watres.2015.03.013
8. Diaz LA, Botte GG. Mathematical modeling of ammonia electrooxidation kinetics in a polycrystalline pt rotating disk electrode. Electrochimica Acta 179:519-528 (2015). https://doi.org/10.1016/j.electacta.2014.12.162
9. Gunuru M, Paravada M. Mathematical modeling of an AEMWE: investigating optimum operating conditions of an alkaline-fed AEM electrolyzer. International Journal of Hydrogen Energy 157:150320 (2025). https://doi.org/10.1016/j.ijhydene.2025.150320
10. Gomez Vidales A, Millan NC, Bock C. Modeling of anion exchange membrane water electrolyzers: the influence of operating parameters. Chemical Engineering Research and Design 194:636-648 (2023). https://doi.org/10.1016/j.cherd.2023.05.004
11. Chang SC, Gu RE, Chan YH. Parameter analysis of anion exchange membrane water electrolysis system by numerical simulation. Energies 17:5682 (2024). https://doi.org/10.3390/en17225682
12. D. Lee, M. Kim, and I. Moon, *Mathematical Modeling of Anion Exchange Membrane Water Electrolyzer*, vol. 52. Elsevier Masson SAS, 2023.
13. Abdin Z, Webb CJ, Gray EM. Modelling and simulation of an alkaline electrolyser cell. Energy 138:316-331 (2017). https://doi.org/10.1016/j.energy.2017.07.053
14. Hu S, Guo B, Ding S, Yang F, Dang J, Liu B, Gu J,

Ma J, Ouyang M. A comprehensive review of alkaline water electrolysis mathematical modeling. Applied Energy 327:120099 (2022). https://doi.org/10.1016/j.apenergy.2022.120099
15. Garibaldi L, Blanco-Aguilera R, Berasategi J, Martinez-Agirre M, Giorgi G, Bracco G, Penalba M. Holistic and dynamic mathematical model for the assessment of offshore green hydrogen generation and electrolyser design optimisation. Energy Conversion and Management 294:117488 (2023). https://doi.org/10.1016/j.enconman.2023.117488
16. Estejab A, Daramola DA, Botte GG. A semi-empirical model of ammonia electrolysis in comparison to water electrolysis. Meet. Abstr. MA2012-02:108-108 (2012). https://doi.org/10.1149/ma2012-02/2/108
17. BALEJ J. Water vapour partial pressures and water activities in potassium and sodium hydroxide solutions over wide concentration and temperature ranges. International Journal of Hydrogen Energy 10:233-243 (1985). https://doi.org/10.1016/0360-3199(85)90093-x
18. P. Virtanen *et al.*, "SciPy 1.0: fundamental algorithms for scientific computing in Python, " *Nat. Methods*, vol. 17, no. 3, pp. 261–272, Mar. 2020.
19. Bynum ML, Hackebeil GA, Hart WE, Laird CD, Nicholson BL, Siirola JD, Watson JP, Woodruff DL. Pyomo overview. Springer Optimization and Its Applications :25-36 (2021). https://doi.org/10.1007/978-3-030-68928-5_3
20. Le Bideau D, Chocron O, Mandin P, Kiener P, Benbouzid M, Sellier M, Kim M, Ganci F, Inguanta R. Evolutionary design optimization of an alkaline water electrolysis cell for hydrogen production. Applied Sciences 10:8425 (2020). https://doi.org/10.3390/app10238425
21. Eurostat, "Electricity and natural gas price statistics., " *Stat. Explain. 2025*.
22. Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math. Program. 106:25-57 (2005). https://doi.org/10.1007/s10107-004-0559-y
23. Rodríguez J, Amores E. CFD modeling and experimental validation of an alkaline water electrolysis cell for hydrogen production. Processes 8:1634 (2020). https://doi.org/10.3390/pr8121634
24. Gwak J, Choun M, Lee J. Alkaline ammonia electrolysis on electrodeposited platinum for controllable hydrogen production. ChemSusChem 9:403-408 (2015). https://doi.org/10.1002/cssc.201501046
25. Long H, Sit CY, Paliwal A, Cheng T, Gupta JK, Reynolds MA, Son YJ, Zhang K, Gewirth AA, Kenis PJA. Alkaline ammonia electrolysis in a membrane electrode assembly cell: parameter optimization and dynamic operation. EES Catal. 4:465-482 (2026). https://doi.org/10.1039/d5ey00324e
26. M. R. Cerqueda, "Selección natural en el hidrógeno verde : Solo quedarán los mejores Autor, " 2024.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Coarse-Graining Algorithm for Complex Chemical Reaction Networks

Yi Tao[ab], Tong Qiu[ab]

[a] Department of Chemical Engineering, Tsinghua University, Beijing 100084, China
[b] Beijing Key Laboratory of Industrial Big Data System and Application, Tsinghua University, Beijing 100084, China
* Corresponding Author: ilmaple@163.com

ABSTRACT

Simulation and modeling technologies play an increasingly critical role in enhancing feedstock utilization, improving product quality, and enhancing enterprise competitiveness in the refining industry. Molecular-level kinetic models for refining processes typically involve highly detailed and complex data, significantly increasing computational costs. Additionally, the data granularity within these models fails to align with the practical requirements of industrial production. To address these, this study proposes the Reaction Network Coarse-Graining Integration (RNCGI) algorithm to achieve flexible adjustments to model complexity and data granularity. Species with similar molecular structures and kinetic behaviors within the reaction network are identified and integrated. Subsequently, a kinetic parameter mapping approach is developed to preserve the critical reaction behaviors. Finally, the reaction coarse-graining process significantly reduces the number of reactions within the network. The RNCGI algorithm was validated using the naphtha catalytic reforming process. The original complete network with 206 species nodes was integrated into two versions with different data granularities: a medium-scale network with 42 nodes and a small-scale network with 5 nodes. Both coarse-grained networks calculated the product composition accurately and provided clearer insights into core species transformation mechanisms compared to the original network. Additionally, the RNCGI algorithm ensures the completeness of reaction routes. By flexibly handling detailed data in molecular-level models, this work preserves the necessary complexity of models tailored to industrial needs, supporting the development of refinery-wide production planning and scheduling strategies.

Keywords: Coarse-graining, Reaction network, Modelling and Simulations, Model Reduction, Algorithms

INTRODUCTION

In response to the rising challenge of declining crude oil quality and stringent carbon reduction targets, the petroleum refining industry is undergoing transformative upgrades toward intelligent operations [1]. Mechanistic simulation and modeling technologies enable the construction of detailed process models that accurately predict critical process indicators, such as product composition. These comprehensive refining models improve product quality while enhancing enterprise competitiveness [2].

Reaction kinetic models form the core foundation of refinery process modeling, and the accurate simulation of key industrial indicators serves as a fundamental criterion for model validation. While the widely employed lumped models [3, 4] offer a simplified representation of reaction processes, their oversimplification leads to poor extrapolation performance when feedstock composition varies. To enhance simulation accuracy, the concept of molecular management [5, 6] has been introduced, facilitating the development of refined molecular-level kinetic models [7, 8, 9] and enabling a comprehensive analysis of reaction mechanisms. However, the number of species and reactions within the refined models significantly increase, which leads to a dramatic rise in computational costs. Moreover, the microscopic granularity of molecular-level models often fails to align with the macroscopic data requirements of practical operations, thereby constraining their applicability for industrial production.

Effective reaction kinetic models should balance industrial applicability with the appropriate level of

complexity, retaining only the components necessary to fulfill process demands. Reaction mechanism reduction is a promising method to lower computation costs and improve model interpretability. A common approach involves identifying less important species or reactions and subsequently removing them from the reaction networks. Chen et al. [10] focused on reaction network structures and applied molecular-level oil characterization data to classify and remove species, thereby significantly reducing computation time. Considering both network structures and simulation results, Fang et al. [11] developed the Network Flow Analysis Algorithm (NFAA) to rank the importance of species and reactions, enabling reaction reduction and parameter adjustments. However, few studies focus on the mismatch in data granularity between molecular-level models and practical production. To address this, we propose an information integration method as an alternative to conventional information reduction in this study. The Reaction Network Coarse-Graining Integration (RNCGI) algorithm was developed to identify and integrate species or reactions exhibiting similar kinetic behaviors. This method facilitates both size reduction and data granularity adjustments in reaction networks. A reaction network of the naphtha catalytic reforming process was employed to validate the RNCGI algorithm. The results demonstrate that the algorithm can flexibly adjust the size and data granularity of the reaction network while ensuring high simulation accuracy. The coarse-grained network, with significantly reduced complexity, preserves the key routes and provides a clear representation of the core species transformation mechanisms.

METHODS

The Reaction Network Coarse-Graining Integration (RNCGI) Algorithm

From the perspective of network science, the kinetic model of a refining process can be transformed into a directed weighted reaction network $G = (V, E, w)$. V represents the node set, corresponding to species in the model. E denotes the directed edge set, corresponding to the reactions, and each edge is directed from reactants to resultants. Each directed edge is assigned a weight w, which quantifies the connection strength between nodes and corresponds to the reaction kinetic information in the model. To achieve scale reduction and data granularity adjustment of complex networks, we have developed the Reaction Network Coarse-Graining Integration (RNCGI) algorithm. Taking the original reaction network as the input, the algorithm generates a coarse-grained reaction network through three sequential steps: node coarse-graining, kinetic parameter mapping, and edge coarse-graining. Assuming that all reactions in the original complex network are first-order reversible reactions, the flowchart of the RNCGI algorithm is illustrated in Fig. 1, including the following three steps. The corresponding pseudo-code is presented in Fig. 2.

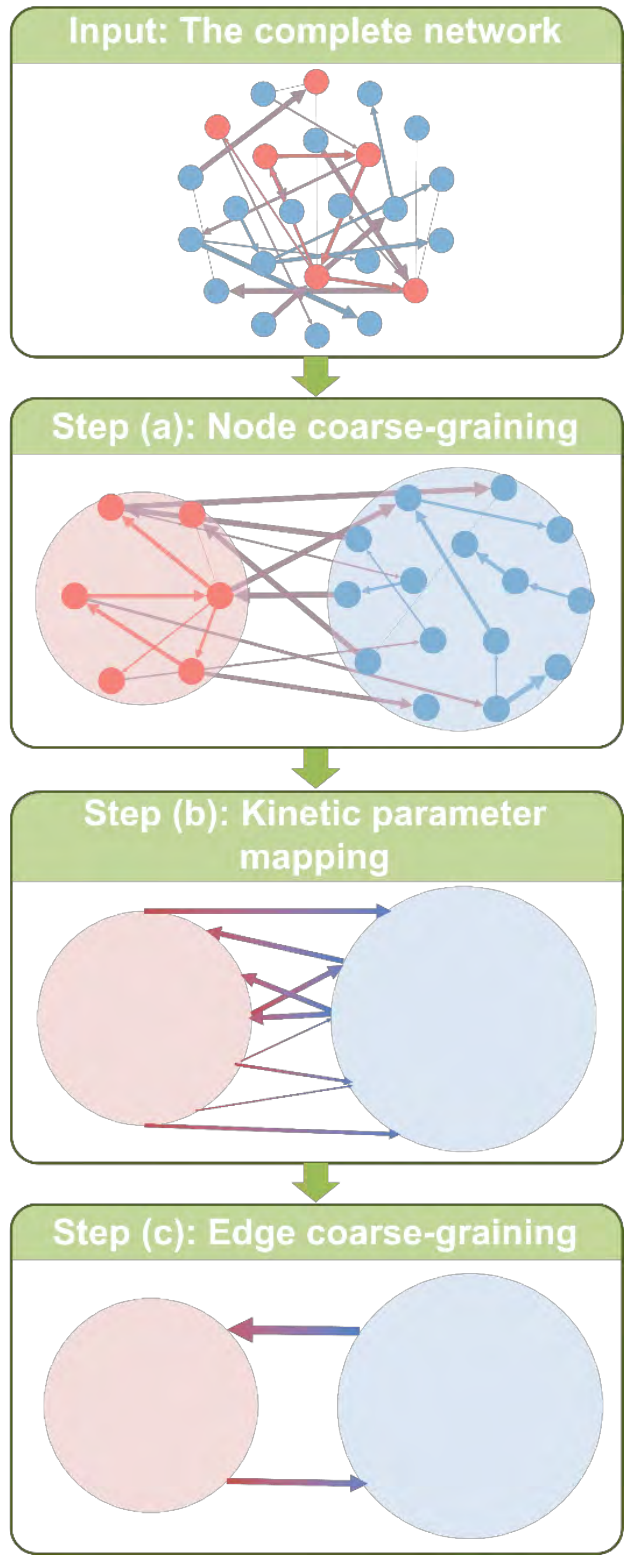

Figure 1. Flowchart of the RNCGI algorithm

Algorithm 1 The Reaction Network Coarse-Graining Integration (RNCGI) Algorithm

Input: The original complete reaction network and the kinetic model $Comp = \{V^{comp}, E^{comp}, k^{comp}, C_{V^{comp}}^{t_0}\}$

Output: The coarse-grained reaction network and the kinetic model $Coa = \{V^{coa}, E^{coa}, k^{coa}, C_{V^{coa}}^{t_0}\}$

V---species nodes in the network

E---reaction edges in the network

k---kinetic parameters in the kinetic model

$C_V^{t_0}$---molar concentrations of species nodes V in the feedstock

Superscript: $comp$---the original complete reaction network and the kinetic model $Comp$

coa---the coarse-graining reaction network and kinetic model Coa

Function Coa = RNCGI($Comp$)

1. Node coarse-graining

$A = \{a_1, a_2, \dots, a_n\}$ $\{a_1, a_2, \dots, a_n\} \in V^{comp}$ ▷ Define the node coarse-graining rules and generate $A \in V^{coa}$

$C_A^{t_0} = \sum_{i=1}^{n} C_{a_i}^{t_0}$ ▷ Adjust the feedstock composition $C_{a_i}^{t_0}$ at time t_0 and generate $C_A^{t_0} \in C_{V^{coa}}^{t_0}$

2. Kinetic parameter mapping

for a reaction $a_i \rightleftharpoons b_j$ with $r = -k_{+,a_i,b_j}^{comp} \times C_{a_i} + k_{-,a_i,b_j}^{comp} \times C_{b_j}$ ▷ The calculation equation for reaction rate r

$C_{a_i}^{mean} = \frac{\sum_{p=1}^{N-1}(t_{p+1}-t_p)\times C_{a_i}^{p+1}}{t_N - t_1}$ ▷ Calculate the time-weighted average molar concentrations $C_{a_i}^{mean}$

$\tilde{k}_{+,a_i,b_j}^{coa} = \frac{k_{+,a_i,b_j}^{comp} \times C_{a_i}^{mean}}{C_A^{mean}} = \frac{k_{+,a_i,b_j}^{comp} \times C_{a_i}^{mean}}{\sum_{i=1}^{n} C_{a_i}^{mean}}$ ▷ Mapping equation for rate constants $\tilde{k}_{+,a_i,b_j}^{coa}$ and $\tilde{k}_{-,a_i,b_j}^{coa}$ that account for concentration changes due to node coarse-graining

$\tilde{k}_{-,a_i,b_j}^{coa} = \frac{k_{-,a_i,b_j}^{comp} \times C_{b_j}^{mean}}{C_B^{mean}} = \frac{k_{-,a_i,b_j}^{comp} \times C_{b_j}^{mean}}{\sum_{j=1}^{m} C_{b_j}^{mean}}$

end

3. Edge coarse-graining

for all reaction edges between a_i and b_j, where $a_i \in A$ and $b_j \in B$

replace all reactions $a_i \rightleftharpoons b_j$ with a single new reaction $A \rightleftharpoons B$ ▷ Edge coarse-graining and generate E^{coa}

$k_{+,A,B}^{coa} = \sum_{i=1}^{n}\sum_{j=1}^{m} \tilde{k}_{+,a_i,b_j}^{coa} + \sum_{i=1}^{n}\sum_{j=1}^{m} \tilde{k}_{-,b_j,a_i}^{coa}$ ▷ Calculate the rate constants for Coa and generate $\{k_{+,A,B}^{coa}, k_{-,A,B}^{coa}\} \in k^{coa}$

$k_{-,A,B}^{coa} = \sum_{i=1}^{n}\sum_{j=1}^{m} \tilde{k}_{-,a_i,b_j}^{coa} + \sum_{i=1}^{n}\sum_{j=1}^{m} \tilde{k}_{+,b_j,a_i}^{coa}$

end

Return $Coa = \{V^{coa}, E^{coa}, k^{coa}, C_{V^{coa}}^{t_0}\}$

Figure 2: The corresponding pseudo-code for the RNCGI algorithm.

Step (a): Node coarse-graining. The first step of the RNCGI algorithm is to identify which nodes in the original complex network can be integrated into a single new node in the coarse-grained network, which is to define the node coarse-graining rules. All species nodes constitute the species library of this reaction system. Multi-dimensional structural features of species, such as the number and linkage pattern of aromatic rings, naphthenic rings and side-chain carbons, serve as classification criteria for these species nodes. The dominant reaction mechanisms and industrial production requirements enable the identification of critical versus secondary structural features. When the node set $\{a_1, a_2, \dots, a_n\}$ in the original network is integrated into a single new node A, the reactions among nodes $\{a_1, a_2, \dots, a_n\}$ are neglected.

Species nodes sharing identical critical structural features but differing only in minor secondary structural features are likely to exhibit similar reaction behaviors, and can therefore be grouped into the same category and integrated into a single node. This approach minimizes the network scale while preserving the dominant species transformation mechanisms, and aligns the network with the data granularity required for practical industrial production. For example, n-decane can undergo cyclization to form naphthene isomers with varying side-chain structures, which can subsequently be dehydrogenated to produce the corresponding aromatic isomers (Fig. 3(a)). Considering the similar reaction mechanisms of species with different side-chain structures, the naphthene and aromatic isomers can be integrated into new nodes, respectively (Fig. 3(b)). In contrast, the molecular structures of paraffins and naphthenes differ

significantly. Therefore, integrating these two types of nodes could lead to the interrupted reaction routes and the loss of important reaction categories, such as cyclization, in the coarse-grained network (Fig. 3(c)).

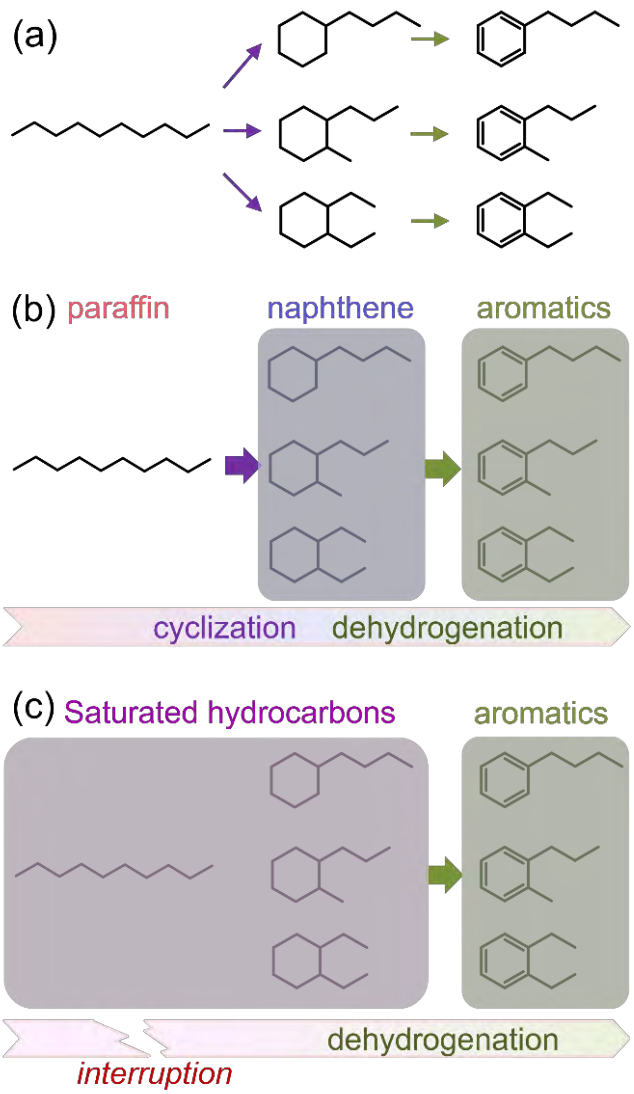

Figure 3. Examples of different node coarse-graining rules. (a) the original reaction routes; (b) the complete routes derived by appropriate node coarse-graining rules; (c) the interrupted routes derived by inappropriate rules

Once the node coarse-graining rules are defined, it is necessary to adjust the feedstock composition at time t_0. The molar concentration of node A in the feedstock of the coarse-grained network is equal to the sum of the molar concentrations of the integrated nodes $\{a_1, a_2, \ldots, a_n\}$ in the original network (Equation 1).

$$C_A^{t_0} = \sum_{i=1}^{n} C_{a_i}^{t_0} \tag{1}$$

Step (b): Kinetic parameter mapping. The molar concentrations of the species nodes are altered due to the node coarse-graining step, thereby introducing perturbations in the reaction behaviors. Consequently, the kinetic parameters of the coarse-grained network could be mapped from those of the original network to preserve the accuracy of overall kinetic behaviors. Suppose the original species node sets $\{a_1, a_2, \ldots, a_n\}$ and $\{b_1, b_2, \ldots, b_m\}$ are integrated respectively into new species nodes A and B. For a first-order reaction involving species a_i and b_j (Equation 2), the reaction rate r is calculated in the original network as Equation 3. C_{a_i} and C_{b_j} are the molar concentrations of a_i and b_j. Theoretically, to preserve the reaction rates after node coarse-graining, the forward and reverse rate constants k_{+,a_i,b_j}^{coa} and k_{-,a_i,b_j}^{coa} at time t_x in the coarse-graining network should satisfy the mapping relationships in Equation 4 with respect to the original forward and reverse rate constants k_{+,a_i,b_j}^{comp} and k_{-,a_i,b_j}^{comp}. n and m are the number of species integrated into the nodes A and B, respectively. As the molar concentrations $C_{a_i}^{t_x}$ and $C_{b_j}^{t_x}$ vary dynamically with time t_x, the theoretical values of k_{+,a_i,b_j}^{coa} and k_{-,a_i,b_j}^{coa} are inherently time-dependent. To effectively preserve the key characteristics of reaction behaviors, the time-weighted average molar concentrations $C_{a_i}^{mean}$ at N discrete time points (Equation 5) are used to calculate the approximate rate constants $\tilde{k}_{+,a_i,b_j}^{coa}$ and $\tilde{k}_{-,a_i,b_j}^{coa}$ in the coarse-grained network (Equation 6).

$$a_i \rightleftharpoons b_j \tag{2}$$

$$r = -k_{+,a_i,b_j}^{comp} \bullet C_{a_i} + k_{-,a_i,b_j}^{comp} \bullet C_{b_j} \tag{3}$$

$$k_{+,a_i,b_j}^{coa} = \frac{k_{+,a_i,b_j}^{comp} \bullet C_{a_i}^{t_x}}{C_A^{t_x}} = \frac{k_{+,a_i,b_j}^{comp} \bullet C_{a_i}^{t_x}}{\sum_{i=1}^{n} C_{a_i}^{t_x}} \tag{4-1}$$

$$k_{-,a_i,b_j}^{coa} = \frac{k_{-,a_i,b_j}^{comp} \bullet C_{b_j}^{t_x}}{C_B^{t_x}} = \frac{k_{-,a_i,b_j}^{comp} \bullet C_{b_j}^{t_x}}{\sum_{i=1}^{m} C_{b_j}^{t_x}} \tag{4-2}$$

$$C_{a_i}^{mean} = \frac{\sum_{p=1}^{N-1} (t_{p+1} - t_p) \bullet C_{a_i}^{p+1}}{t_N - t_1} \tag{5}$$

$$\tilde{k}_{+,a_i,b_j}^{coa} = \frac{k_{+,a_i,b_j}^{comp} \bullet C_{a_i}^{mean}}{C_A^{mean}} = \frac{k_{+,a_i,b_j}^{comp} \bullet C_{a_i}^{mean}}{\sum_{i=1}^{n} C_{a_i}^{mean}} \tag{6-1}$$

$$\tilde{k}_{-,a_i,b_j}^{coa} = \frac{k_{-,a_i,b_j}^{comp} \bullet C_{b_j}^{mean}}{C_B^{mean}} = \frac{k_{-,a_i,b_j}^{comp} \bullet C_{b_j}^{mean}}{\sum_{i=1}^{m} C_{b_j}^{mean}} \tag{6-2}$$

Step (c): Edge coarse-graining. After node coarse-graining, multiple edges may exist between the same pair of nodes in the network, corresponding to multiple identical reactions. To ensure the simplicity and consistency of the coarse-grained network, these edges are integrated into a single representative edge directed from node A to B, which represents the reaction in Equation 7. Considering the directions of all the integrated edges, the forward and reverse rate constants for the coarse-grained new reaction $k_{+,A,B}^{coa}$ and $k_{-,A,B}^{coa}$ are calculated by

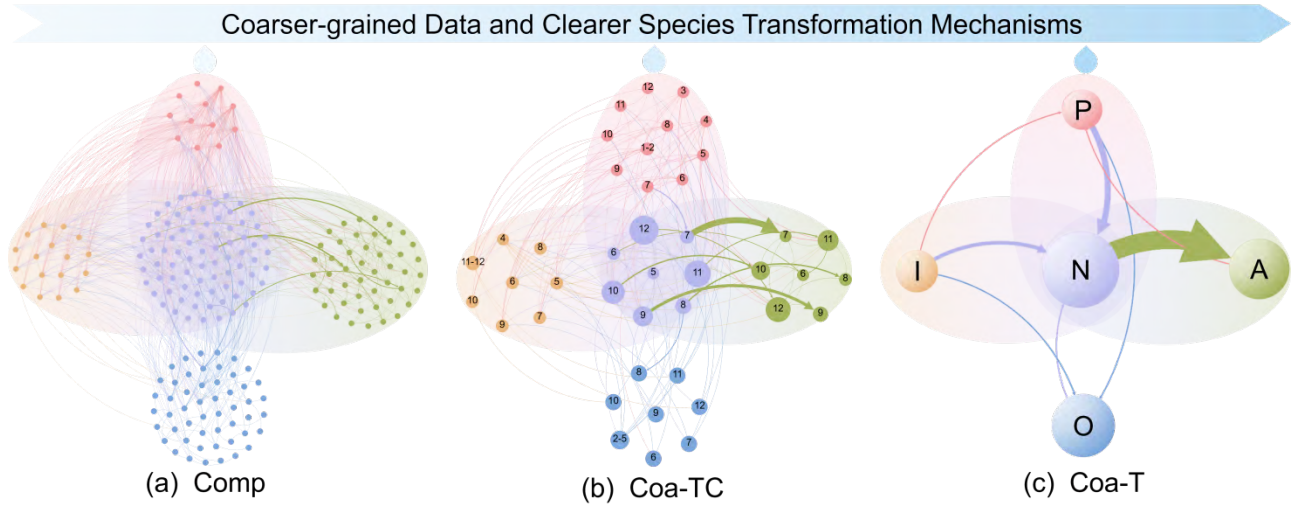

Figure 4. Comparison of three reaction networks. (a) the original complete network (Comp); (b) the medium-scale coarse-graining network (Coa-TC); (c) the small-scale coarse-graining network (Coa-T)

aggregating the respective rate contributions (Equation 8). $\tilde{k}^{coa}_{+,b_j,a_i}$ and $\tilde{k}^{coa}_{-,b_j,a_i}$ denote the approximate forward and reverse rate constants, respectively, after kinetic parameter adjustment for the reaction where b_j serves as a reactant and a_i as a resultant.

$$A \rightleftharpoons B \tag{7}$$

$$k^{coa}_{+,A,B} = \sum_{i=1}^{n}\sum_{j=1}^{m}\tilde{k}^{coa}_{+,a_i,b_j} + \sum_{i=1}^{n}\sum_{j=1}^{m}\tilde{k}^{coa}_{-,b_j,a_i} \tag{8-1}$$

$$k^{coa}_{-,A,B} = \sum_{i=1}^{n}\sum_{j=1}^{m}\tilde{k}^{coa}_{-,a_i,b_j} + \sum_{i=1}^{n}\sum_{j=1}^{m}\tilde{k}^{coa}_{+,b_j,a_i} \tag{8-2}$$

Example: The Catalytic Reforming Reaction Network

Catalytic reforming of naphtha is a crucial process for increasing the octane number of gasoline and producing high-value-added aromatics. Taking a self-constructed molecular-level catalytic reforming network as an example, the RNCGI algorithm was employed to derive coarse-grained networks at different levels of data granularity. The original complex network (Comp) consists of 206 species nodes and 674 reaction edges, in which molecular structures of species are described via the structure-oriented lumping (SOL) method. The edge weight in the network represents the cumulative molar flow rate CF_h^0 (Equation 9) of the corresponding reaction h throughout the complete reaction extent. ν_0 is the constant volumetric flow rate, r_h is the reaction rate (Equation 3), and t is the reaction time.

$$CF_h^0 = \nu_0 \int_0^t r_h \, \mathrm{d}t \tag{9}$$

The catalytic reforming process primarily involves species with carbon numbers ranging from C1 to C12, and the detailed composition of these light components could be determined using advanced oil characterization techniques. The two-dimensional PIONA (paraffin, isoparaffin, olefin, naphthene, and aromatic) compositions of feedstocks and products represent key industrial parameters in the reforming process. Based on practical purposes, species with the same PIONA type and carbons in the Comp network can be integrated into a single new node, while neglecting more detailed structural information from the SOL method. The RNCGI algorithm results in a coarse-grained network, termed Coa-TC (type and carbon-based coarse-graining).

On the other hand, coarser-grained data, such as the aromatic yield, are key considerations in the development of production planning and plant-wide scheduling strategies. To address this, the RNCGI algorithm allows for flexible adjustment of data granularity, using only PIONA types as the basis for node coarse-graining. This approach results in a coarser-grained network, Coa-T.

RESULTS AND DISCUSSIONS

Network Scale and Kinetic Accuracy

Fig. 4 shows the scales of the Comp, Coa-TC, and Coa-T networks. In the Coa-TC network (Fig. 4(b)), species nodes are distinguished by the PIONA types and carbon numbers. The node colors represent different PIONA types, serial numbers indicate the carbon numbers, and node sizes reflect the number of original nodes that are integrated into the corresponding coarse-grained node. The Coa-TC network consists of 42 species nodes and 131 reaction edges, accounting for only about 20% of the complexity of the original Comp network. Furthermore, the species concentrations within all nodes directly correspond to the two-dimensional PIONA composition, ensuring the clear representation of feedstocks and products. At different temperatures, the product components

calculated by the Coa-TC network are in good agreement with the original values in the Comp network, with the maximum absolute error less than 1.5% (Fig. 5). This demonstrates that the RNCGI algorithm ensures the accuracy of the coarse-grained network.

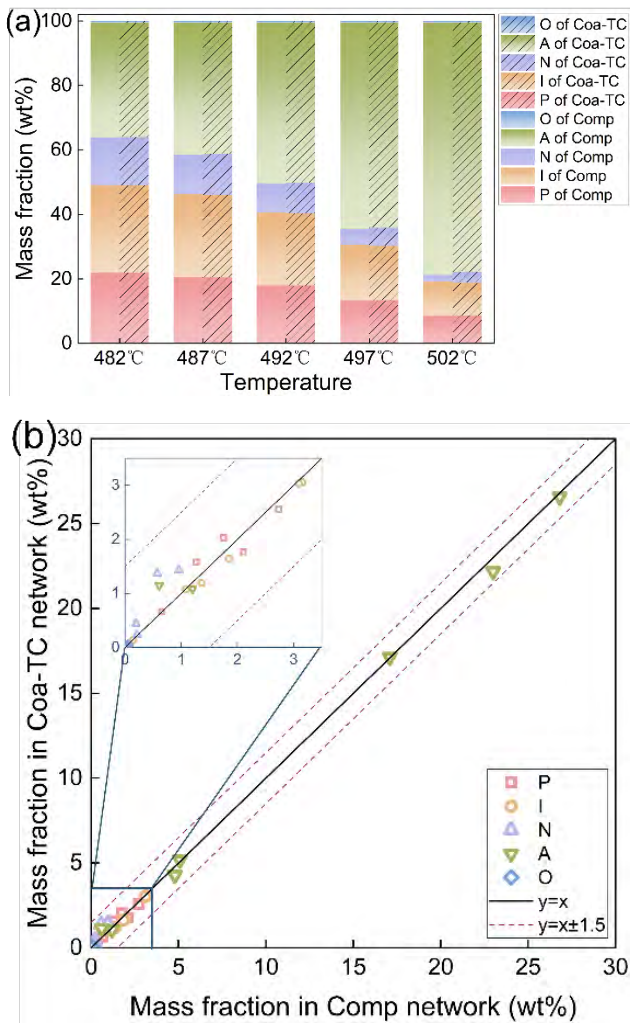

Figure 5. The Calculation results of the Coa-TC network (a) PIONA compositions at different temperatures; (b) two-dimensional PIONA compositions at 502 ℃

In contrast, the Coa-T network only consists of five species nodes (P, I, N, A, O) and eight reaction edges (Fig. 4(c)). In the Coa-T network, the calculated PIONA compositions across different temperatures show strong consistency with those obtained from the Comp network, with absolute errors maintained within 5% (Fig. 6). Compared to the Comp network, the Coa-T network exhibits a significantly reduced scale while better enhancing alignment with industrial needs for aromatics production, providing a simplified and accurate foundation for plant-wide scheduling and process optimization.

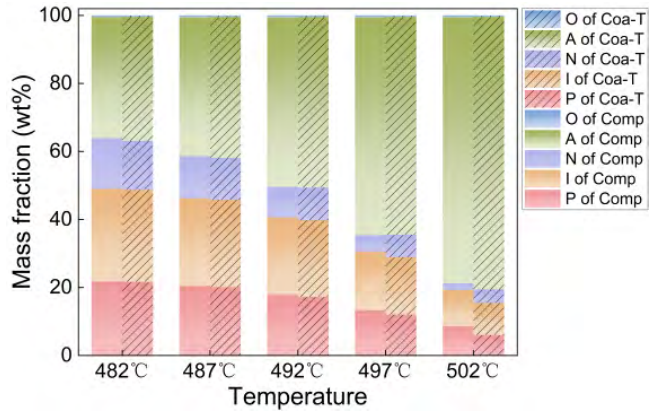

Figure 6. The Calculation results of the Coa-T network

Species Transformation Mechanisms

One of the primary objectives of catalytic reforming is to produce aromatics. The aromatic content in the product increases significantly with rising temperatures. Fig. 4 compare the data granularity and species transformation mechanisms among the Comp, Coa-TC, and Coa-T networks. Across different levels of data granularity, these networks reveal a consistent reaction route for the accumulation of aromatics: chain alkanes undergo cyclization to form naphthenes, which are subsequently aromatized. In this route, chain alkanes, naphthenes and aromatics serve as the main reactants, intermediates and products, respectively. Furthermore, the essential species transformation mechanisms, from chain alkanes to naphthenes and finally to aromatics, become more clearly evident in the coarser-graining network structures (shaded area in Fig. 4).

Additionally, the RNCGI algorithm demonstrates superior performance compared to traditional information reduction methods in reducing reaction networks while maintaining simulation accuracy. Table 1 evaluates reduced networks derived from different algorithms with respect to network scale and simulation accuracy. The reaction deletion (RD) network is obtained by using the NFAA algorithm [11], which ranks reaction importance and remove reactions with low importance. When the RD network is reduced to 198 species nodes and 278 reaction edges, the calculated mass fractions of iso-paraffins and aromatics with 9 carbons (denoted as I7 and A7, respectively) in the product exhibit significant deviations compared to the Comp network. In contrast, the Coa-TC network, generated using the RNCGI algorithm, contains fewer species nodes and reaction edges while achieving higher simulation accuracy compared to the RD network. This enhanced performance is attributed to the kinetic information integration in the RNCGI algorithm, which preserves critical reaction routes while reducing the overall network scale.

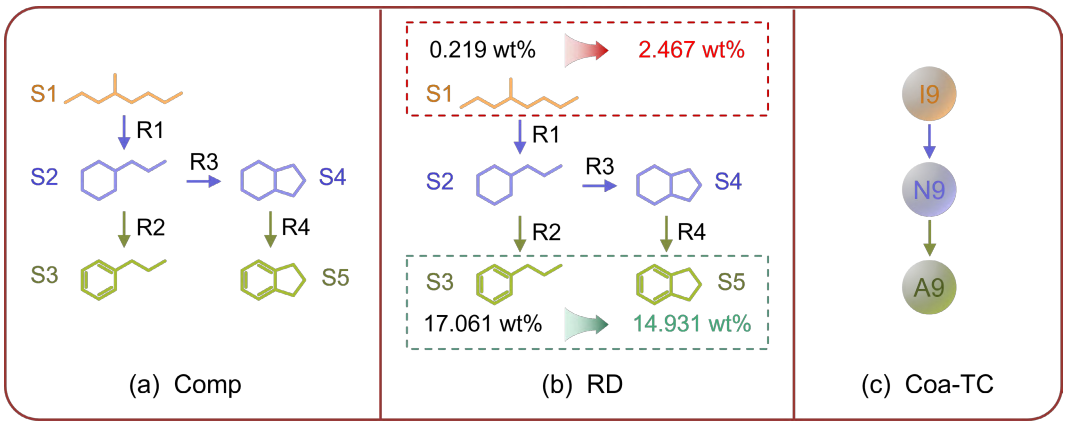

Figure 7. Comparison of speceis transformation mechanisms within different networks. (a) reaction routes in the Comp network; (b) reaction routes in the RD network; (c) reaction routes in the Coa-TC network

Table 1. The scale and simulation accuracy of reduced networks derived from different algorithms

Metrics		**Networks**	
		Coa-TC	RD
Absolute error	I9	0.087 wt%	2.417wt%
	A9	0.094 wt%	2.299wt%
Network scale	Node number	42	198
	Edge number	131	278

From the perspective of species transformation mechanisms, the generation of A9 in the Comp network primarily proceeds via cyclization and aromatization of species with the same carbon number (Fig. 7(a)). In the RD network (Fig. 7(b)), the cyclization reaction of I9 (reaction R1) is removed. The interruption of this route reduces the content of intermediate naphthene species (S2), thereby diminishing the reaction flow rates of its subsequent aromatization and side-chain cyclization reactions. These affected kinetic behaviors lead to a decreased production of aromatic species (S3 and S5), while the iso-paraffin species (S1) accumulate excessively due to the interruption of transformation routes. In contrast, the RNCGI algorithm replaces information deletion with an information integration strategy, preserving critical reaction routes in the Coa-TC network (Fig. 7(c)). This approach ensures accurate kinetic behaviors even with a significant reduction in network scale.

CONCLUSIONS

This study proposes a novel reaction network coarse-graining integration (RNCGI) algorithm, which comprehensively considers the accuracy, simplicity, and applicability of reaction models. The node coarse-graining rules are constructed specific to various industrial demands, and kinetic parameters mapping and reaction coarse-graining are followed to derive a coarse-grained network. In comparison to the original Comp network, the resulted Coa-TC and Coa-T networks exhibit a significantly reduced scale and a data granularity that is more closely aligned with real industrial production. Therefore, these coarse-grained networks enable the rapid and accurate calculation of critical industrial indicators. The RNCGI algorithm employs information integration instead of deletion, consistently preserving key reaction routes across various degrees of coarse-grained networks. By avoiding reaction routes interruption, this approach ensures the reliability of species transformation mechanisms within the network.

Future work will likely focus on integrating this algorithm into the refinery-wide scheduling and optimization software in a modular form, providing practical support for achieving comprehensive automation in modern refinery operations.

AUTHOR IDENTIFIERS

Author ORCIDs:
Yi Tao: 0009-0006-8012-7029
Tong Qiu: 0000-0003-4294-9671

REFERENCES

1. Hasheminasab H, Gholipour Y, Kharrazi M, Streimikiene D, Hashemkhani S. A dynamic sustainability framework for petroleum refinery projects with a life cycle attitude. Sustainable Development 28:1033-1048 (2020). https://doi.org/10.1002/sd.2054
2. Bi K, Zhang S, Zhang C, Li H, Huang X, Liu H, Qiu T. Knowledge expression, numerical modeling and optimization application of ethylene thermal cracking: from the perspective of intelligent manufacturing. Chinese Journal of Chemical

Engineering 38:1-17 (2021). https://doi.org/10.1016/j.cjche.2021.03.033
3. Padmavathi G, Chaudhuri KK. Modelling and simulation of commercial catalytic naphtha reformers. Canadian J Chem Eng 75:930-937 (1997). https://doi.org/10.1002/cjce.5450750513
4. Smith RJ. Kinetic analysis of naphtha reforming with platinum catalyst. Chemical Engineering Progress, 55: 76-80 (1959).
5. Ren Y, Liao Z, Sun J, Jiang B, Wang J, Yang Y, Wu Q. Molecular reconstruction: recent progress toward composition modeling of petroleum fractions. Chemical Engineering Journal 357:761-775 (2019). https://doi.org/10.1016/j.cej.2018.09.083
6. Van de Vijver R, Devocht BR, Van Geem KM, Thybaut JW, Marin GB. Challenges and opportunities for molecule-based management of chemical processes. Current Opinion in Chemical Engineering 13:142-149 (2016). https://doi.org/10.1016/j.coche.2016.09.006
7. Quann RJ, Jaffe SB. Building useful models of complex reaction systems in petroleum refining. Chemical Engineering Science 51:1615-1635 (1996). https://doi.org/10.1016/0009-2509(96)00023-1
8. Chen Z, Dang F, Zhao X, Gao J, Xu C, Wang G, Zhang L. A molecular-level kinetic model for one-step catalytic conversion of crude oil to light olefins. Chemical Engineering Science 316:121971 (2025). https://doi.org/10.1016/j.ces.2025.121971
9. Chen Z, Xie Y, Xu C, Zhang L. Scale-up of complex molecular reaction system by hybrid mechanistic modeling and deep transfer learning. Nat Commun 16: (2025). https://doi.org/10.1038/s41467-025-63982-2
10. Chen Z, Wang Y, Li Y, Wu J, Zhang L. Explicit molecule-based reaction network simplification: theory and application on catalytic reforming. Chemical Engineering Science 277:118833 (2023). https://doi.org/10.1016/j.ces.2023.118833
11. Fang Z, Qiu T, Chen B. Improvement of ethylene cracking reaction network with network flow analysis algorithm. Computers & Chemical Engineering 91:182-194 (2016). https://doi.org/10.1016/j.compchemeng.2016.04.020

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Automated Construction of Bayesian Networks of Chemical Process for Dynamic Risk Assessment

Kai Yin[ab], Hao Kang[c], Jinsong Zhao[a]*

[a] Tsinghua University, Department of Chemical Engineering, Beijing, 100084, China
[b] State Key Laboratory of Chemical Engineering and Low-carbon Technology, Engineering Department, Beijing, 100084, China
[c] Zhejiang Petroleum & Chemical Co., Ltd, Zhoushan, Zhejiang, China
* Corresponding Author: jinsongzhao@tsinghua.edu.cn.

ABSTRACT

Chemical processes are characterized by high complexity and inherent hazards, necessitating systematic and comprehensive risk assessment methodologies. However, traditional risk assessment approaches are often limited to static analyses and small-scale models, resulting in insufficient coverage and a lack of detailed, unit-level insights. This paper proposes an automated Bayesian network–based dynamic risk assessment (DRA) framework for full-process risk analysis in chemical plants. The proposed method automatically extracts knowledge from established risk analysis documents, such as HAZOP reports, and integrates it with dynamic process monitoring data, expert knowledge, and reliability databases. By effectively leveraging multi-source heterogeneous data, an equipment-level causal inference network is constructed automatically, addressing the high labor demand and limited scalability of conventional DRA approaches. The proposed framework is validated through an industrial FCC case study at a large petrochemical enterprise in Zhejiang, China. The results demonstrate that the method can effectively construct Bayesian causal networks for chemical processes and quantitatively evaluate both the likelihood and severity of accident consequences under various failure scenarios, thereby enabling systematic and dynamic risk assessment of chemical processes.

Keywords: dynamic risk assessment, Bayesian network, multi-source heterogeneous data, automated construction

1. INTRODUCTION

Chemical processes are inherently complex and fraught with various risks due to their dynamic nature and the involvement of hazardous materials. Traditionally, risk assessments in these environments have been static, relying on predefined assumptions and historical data. However, as industrial systems grow more complex and real-time monitoring becomes more prevalent, these conventional approaches are no longer sufficient to capture the dynamic behavior and evolving risks that can occur during operation.

Dynamic risk assessment (DRA) has emerged as a more suitable methodology to address these challenges. The main goal of DRA is to provide a real-time, adaptable framework for assessing the risks associated with chemical processes. This approach continuously updates risk levels based on variations in operating conditions, equipment degradation, human–machine interactions, and control system performance. By incorporating time-dependent factors, DRA offers a more comprehensive understanding of risk propagation and its evolution over time, enabling better decision-making in high-risk environments.

According to the definition by Khan et al.[1], DRA is a methodology that updates process risk estimates based on the performance of control systems, protection barriers, inspection and maintenance activities, human factors, and operational procedures. A key distinction between DRA and traditional risk assessment lies in the explicit incorporation of time-dependent factors. Importantly, risk analysis should quantify both the probability and the consequences of hazardous scenarios. Approaches that focus exclusively on occurrence probability without assessing consequence severity result in incomplete and potentially misleading risk evaluations.

In 2006, Meel and Seider[2] developed an integrated framework combining precursor information with real-time failure data to dynamically assess accident sequences, marking one of the earliest practical applications of DRA in chemical process safety.

Building upon tree-based models, subsequent research increasingly focused on dynamic risk assessment. Kalantarnia et al.[3] advanced DRA by integrating Bayesian failure modeling with consequence analysis. Their framework consists of two key modules: (i) an accident analysis module, which incorporates real-time process parameters into dynamic probability estimation using fault tree analysis; and (ii) a probability updating module, which iteratively updates probabilities through Bayesian inference as new data become available.

Between 2000 and 2009, Bayesian networks (BN) were widely adopted in DRA due to their natural compatibility with probabilistic updating and causal reasoning. Montani et al.[4] introduced dynamic Bayesian networks (DBN) and developed the DBNet tool to automatically convert dynamic fault trees into DBNs, representing a significant milestone in BN-based DRA. Kohda and Cui et al.[5] proposed a DBN-based dynamic reconfiguration method for safety monitoring systems, using time slices to model operational states and sensor reliability degradation. More recently, Shan et al.[6] employed fuzzy Bayesian networks (FBN) to quantify gas pipeline leakage risk, while Guo et al.[7] proposed a fuzzy dynamic Bayesian network (FDBN) framework for uncertainty propagation in DRA.

Machine learning methods have been increasingly applied to DRA. Adedigba et al.[8] used artificial neural networks (ANNs) for real-time accident probability prediction, while Wang et al.[9] applied long short-term memory (LSTM) networks for real-time risk forecasting. Other techniques, such as principal component analysis (PCA)[10] and support vector machines (SVM)[11], have also been explored.

Existing DRA methodologies mainly depend on Event Tree Analysis (ETA), Fault Tree Analysis (FTA), and Bow-Tie Analysis. Despite their effectiveness, most current approaches suffer from several limitations: incomplete risk characterization, limited network scale, and high reliance on expert knowledge. In particular, the construction of BN structures and the assignment of initial probabilities often require substantial manual effort, which hinders scalability and industrial applicability.

Although these studies collectively illustrate the evolution of DRA methodologies, most existing approaches still rely heavily on manual model construction and expert-driven probability assignment. This limitation is particularly pronounced for methods such as FBN, which are highly dependent on expert judgment. To address these challenges, this paper proposes an automated framework for extracting knowledge from multi-source heterogeneous data and constructing BN for DRA.

2. PRELIMINARIES

This section introduces the key techniques and algorithms employed in the proposed framework, including time-series Transformer-based causal discovery, sentence Transformer for heterogeneous textual data, and Bayesian parameter estimation. These techniques jointly enable the automated construction of a Bayesian causal network for DRA.

2.1 Time-Series Transformer-Baesd Causal Discovery

Causal discovery from process data plays a critical role in constructing the structural framework of a causal BN. In complex chemical process such as Fluid Catalytic Cracking (FCC), process variables are strongly coupled, dynamically evolving, and often influenced by delayed and indirect interactions. Traditional correlation-based methods are insufficient to capture such causal relationships.

In this work, causal discovery is performed using a time-series Transformer model[12] applied to DCS historical data. Transformer architectures, originally developed for sequence modeling, are particularly well suited for capturing long-range temporal dependencies and nonlinear interactions among multivariate time-series signals. By leveraging self-attention mechanisms, the model can learn how past values of one variable contribute to the future evolution of another.

Specifically, Granger causality is estimated within the Transformer framework: a variable X is considered to Granger-cause another variable Y if incorporating the historical information of X significantly improves the prediction of Y compared to using the history of Y alone. By evaluating such predictive contributions across all monitored variables, a directed causal graph is obtained, which serves as the initial structure of the BN. This data-driven approach enables scalable and automated identification of causal relationships in large-scale industrial systems without relying solely on expert-defined structures.

2.2 Sentence Transformer

To bridge the semantic gap between Bayesian network node descriptions expressed in unstructured natural language and reliability database entries provided in structured textual form, this study employs the Sentence-BERT (SBERT) framework for semantic vector representation. Conventional BERT models typically adopt a cross-encoder architecture for sentence similarity tasks, in which sentence pairs are jointly encoded. Although effective, this approach incurs high computational cost and is therefore unsuitable for large-scale database matching

and real-time retrieval. SBERT addresses this limitation by reformulating pretrained BERT models into a Siamese architecture with shared weights, allowing each sentence to be encoded independently into a fixed-dimensional embedding.

Given a Bayesian network node description A and a candidate database entry B, SBERT processes the two texts separately and generates contextualized token-level representations using BERT encoders. A pooling operation is then applied to aggregate token embeddings into a sentence-level vector. In this work, mean pooling is adopted, where the final sentence embedding u is computed as the average of all token embeddings:

$$u = \frac{1}{N}\sum_{i=1}^{N} h_i \quad (3)$$

where hi denotes the embedding of the ith token and N is the sentence length.

To ensure that the resulting embeddings exhibit meaningful geometric properties in the semantic space, SBERT is typically fine-tuned using a supervised classification objective. For a sentence pair (A, B), the corresponding embeddings u and v are concatenated with their absolute difference, and fed into a Softmax classifier trained on semantic entailment, contradiction, and neutrality labels. This training strategy encourages semantically similar sentences to be mapped closer together in the embedding space while pushing dissimilar sentences further apart.

During inference, Bayesian network node names and reliability database entries are mapped into high-dimensional vectors and their semantic correspondence is quantified using cosine similarity:

$$sim(A, B) = \frac{u \cdot v}{\|u\| \cdot \|v\|} \quad (4)$$

The resulting similarity score ranges from [−1, 1], with higher values indicating stronger semantic similarity. This approach not only captures direct lexical overlap but also effectively identifies implicit semantic associations, such as those between instrumentation tags and their corresponding process semantics, thereby enabling robust matching between Bayesian network nodes and heterogeneous reliability data sources.

2.3 Bayesian Parameter Estimation

Parameter learning for conditional probability tables (CPTs) in BN is performed using Bayesian inference with Dirichlet priors. A commonly used baseline approach is maximum likelihood estimation (MLE), which estimates probabilities by simple frequency counting. For example, if node X has parent Y, and the dataset contains *N* instances where Y=High, among which X=Alarm occurs *n* times, then:

$$P(X = Alarm|Y = High) = \frac{n}{N} \quad (4)$$

However, MLE suffers from the zero-frequency problem. If a critical event has not occurred in the dataset, MLE assigns it a probability of zero. In process safety applications, “not observed” does not imply “impossible, ” and MLE completely ignores prior knowledge from reliability databases such as OREDA.

To overcome these limitations and integrate prior knowledge with observational data, Bayesian estimation is adopted. In discrete BN, node states follow a multinomial distribution, whose conjugate prior is the Dirichlet distribution. This approach is commonly referred to as Dirichlet smoothing.

A key parameter in Bayesian estimation is the Equivalent Sample Size (ESS), which quantifies the confidence placed in prior probabilities. ESS represents the number of virtual observations implied by the prior. Given a prior probability $P_{prior}(X_i = k|Pa_i = j)$, the corresponding pseudo-count is:

$$a_{ijk} = P_{prior}(X_i = k|Pa_i = j) \times ESS \quad (1)$$

The updated CPT is then computed as:

$$\theta_{ijk} = \frac{N_{ijk}+\alpha_{ijk}}{N_{ij}+\alpha_{ij}} \quad (2)$$

where N_{ijk} is the observed count from process data and α_{ij} is the sum of pseudo-counts. When large amounts of data are available, the influence of the prior diminishes. When data are scarce, the prior prevents zero-probability artifacts.

3. PROPOSED FRAMEWORK

This chapter provides a detailed description of the complete implementation workflow of the proposed automated dynamic risk assessment framework. The framework is designed to address two fundamental limitations of conventional risk assessment approaches: first, the subjectivity and high labor cost associated with manual construction of BN; and second, the inherent latency of static risk data, which fails to reflect the real-time operating state of chemical processes. As illustrated in Fig. 3, the implementation workflow is tightly integrated, starting from preprocessing of DCS data, proceeding through causal discovery and multi-source heterogeneous data extraction, and ultimately achieving the automated construction of a risk-oriented causal network for chemical processes.

3.1 Data Preprocessing and Transformer-Based Causal Discovery

The first stage of the framework aims to automatically extract the topology of the BN from the vast amount of Distributed Control System (DCS) data in chemical processes.

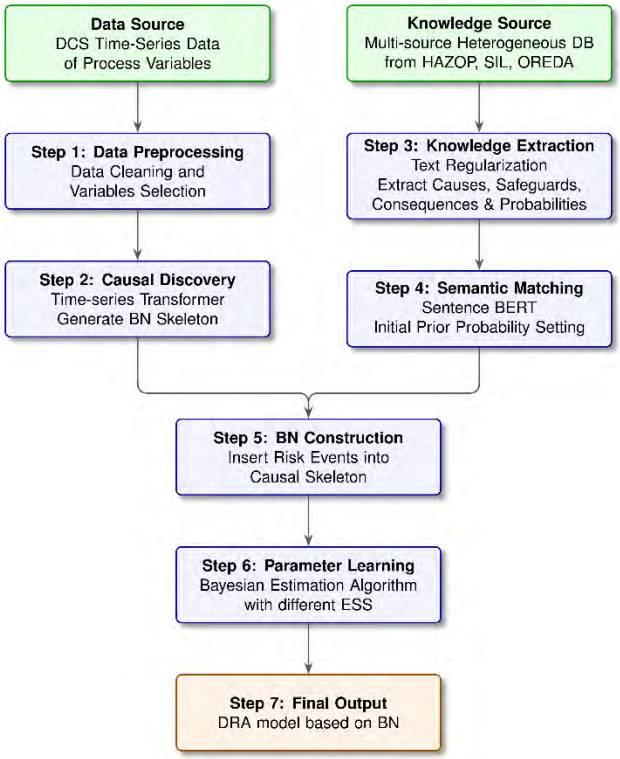

Figure 1. Framework of automated Bayesian causal network construction.

Initial preprocessing of the raw DCS time-series data involves selecting the key variables that have a critical impact on the safety of the process. In real-world chemical production plants, a single plant consists of thousands of instrument points that indicate the parameters of different systems and equipment. During the risk assessment process, some variables are closely related to the safe operation of the plant, while others are used to indicate production efficiency. In our risk analysis, we focus primarily on those variables that are tightly connected to production safety.

To extract these key variables, we collect a list of key variables from the HAZOP analysis table. The structure of the HAZOP analysis table is shown in Table 1. The HAZOP analysis report provides risk data on the occurrence probabilities of causes and consequences for a specific unit's risk scenarios, which are determined through brainstorming sessions by an expert panel. In the HAZOP analysis process, expert panels analyze all nodes in the P&ID diagrams and evaluate the risks in the process based on deviations observed in DCS data. We extract high-risk variables identified by the expert group during the HAZOP analysis, and then select these variables from the thousands of DCS data points for causal discovery.

After preprocessing, the study employs a Time-series Transformer model to capture generalized Granger causality relationships between process variables and construct the structural backbone of the BN. The Transformer architecture utilizes a self-attention mechanism to compute an attention weight matrix that quantifies the influence of historical states of variable j on the future state of variable i. This influence is expressed as the attention weight α_{ij}. By setting a significance threshold δ, if $\alpha_{ij}>\delta$, a causal link from node j to node i is determined. The algorithm iterates through all pairs of process variables, generating an adjacency matrix that describes the global causal relationships of the system. This matrix is directly transformed into a directed acyclic graph (DAG) of the Bayesian network, successfully identifying implicit control loop couplings. This approach not only significantly reduces the time cost of manual network construction but also ensures the objectivity and accuracy of the network structure.

3.2 Mapping Risk Events from HAZOP to BN

Relying solely on node identifiers is insufficient to fully capture all the risks in a chemical process. To comprehensively reflect all risks and their evolution paths within the system, we match key node identifiers from the HAZOP report to potential causes and consequences of deviations, as well as the protective measures within the risk evolution paths. Using text normalization techniques, specific matching rules are established to link the risk causes, consequences, and protective measures from the HAZOP analysis table to the deviations of key variables. This enables the exploration of the risk evolution path, where risk sources cause deviations in key variables, bypass protective measures, and ultimately lead to severe consequences. This process forms a Bayesian network that incorporates risk causes, consequences, protective measures, DCS variables, and risk outcomes.

3.3 Automated Extraction of Risk Data for BN Initialization

To accurately obtain the occurrence probabilities of causes and consequences, as well as the Probability of Failure on Demand (PFD) of protection layers, the method extracts failure data from the OREDA database, SIL verification reports, and HAZOP analysis reports, constructing a comprehensive and targeted database. SIL verification reports provide failure probability analyses of protection measures for a specific unit, confirmed through expert discussions based on national standards, industry guidelines, and expert experience, thus providing targeted reliability data for protection layers. HAZOP analysis reports contain risk data on the occurrence probabilities of causes and consequences for unit-specific risk scenarios, determined through expert brainstorming sessions and analyzed according to the HAZOP procedure and expert judgment, offering targeted risk assessment data. The OREDA database, a widely used Norwegian reliability database, supplements failure probability data and risk assessment information that may not be

Table 1: Structure Example of a HAZOP table

Designator	Deviation	Cause	Protection Layer	Consequence	Probability	Severity
FI10101	High	HV10101 closed on failure	The feed will be cut off by interlock TICA10101.	Resulting in regenerator overtemperature.	3	5
TI10103	High	R0102 broke down	The settler is equipped with multiple-point temperature display.	The wax oil solidified, causing a pipeline blockage.	2	2
PIC10104	Low	P0101 broke down	P0101 one as main, one as standby.	The knockout drum was damaged due to overpressure.	3	5

captured in the SIL verification or HAZOP reports.

In the process of filling the BN with information from the database, a pre-classified Sentence-BERT (SBERT) algorithm is used to match entries, eliminating the need for manual risk probability assessment. The specific matching algorithm, as described in section 2.2, first classifies risk events into causes, protection measures, and consequences. This classification is pre-labeled according to the HAZOP analysis table during the insertion of causal events in section 3.2. The protection measures are further categorized by different equipment types, and the corresponding PFDs are matched from the database. Based on the risk event names in the Bayesian network, failure entries from the risk database are matched, allowing for the assignment of initial probabilities to risk events and completing the initialization of the Bayesian network.

3.4 Parameter Learning

After completing the initialization of the Bayesian network, parameter learning is performed using the Bayesian inference algorithm described in section 2.3. The initial probability distribution for risk event nodes is derived from data in the risk database, while a normal distribution is used as the initial probability distribution for critical node data. Actual accident occurrence data and DCS data are used as observational samples for parameter learning, allowing for the refinement of probability distributions and conditional probability tables (CPT) to better reflect real conditions.

3.5 Bayesian Inference

After the topology construction and parameter learning of the Bayesian causal network are completed, the network can be used for probabilistic inference. This network integrates risk event evidence with real-time DCS data streams to perform DRA. The inference can be performed in two modes, forward prognostic inference and backward diagnostic inference. By inputting observational data for a subset of nodes and evidence of initiating risk sources, the probabilities of different consequences can be inferred, enabling dynamic risk assessment. Conversely, given evidence of observed consequences, backward inference can be performed to trace and identify the most likely underlying causes that led to those outcomes.

4. CASE STUDY

The proposed automated Bayesian causal network construction method was applied to the fluid catalytic cracking process of a petrochemical enterprise in Zhejiang, China, to validate the feasibility of the approach.

According to the HAZOP analysis report provided by the enterprise, a total of 769 risk scenarios were identified. Among the alarm tag nodes involved in these risk events, some nodes did not have corresponding DCS data available. Consequently, 94 tag variables that were both included in the HAZOP analysis and recorded in the DCS were selected for further analysis. A Time-Series Transformer was applied to these variables for causal discovery, with the significance threshold δ set to 1.05. In total, 141 causal relationships among the variables were identified, and the resulting causal adjacency matrix is shown in Fig. 2(a). To verify the accuracy of the causal discovery algorithm, the causal relationships among these variables were manually assessed. Together with expert engineers, the production structure of the unit was analyzed to identify causal relationships based on expert knowledge, and the corresponding expert-derived causal neighborhood is illustrated in Fig. 2(b). The two matrix is generated by different ways, so the designator names are not completely same. But labels are corresponded in sequence. Treating the expert-identified causal relationships as the ground truth, a comparison with the data-driven causal matrix shows that the proposed causal discovery algorithm achieves an accuracy of 94.61%. The accuracy was calculated by the similarity

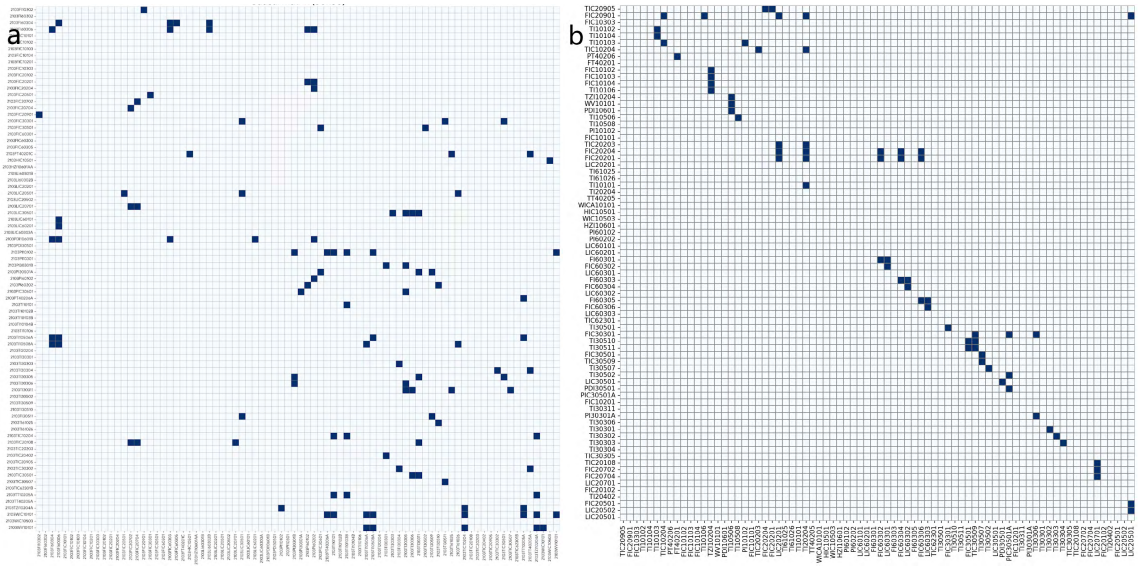

Figure 2. Causal adjacency matrix of 94 tag variables. (a) Discovered by Time-Series Transformer. (b) Discovered by expert knowledge

of two adjacent matrix.

Based on the results of causal discovery, the nodes and connections in the adjacency matrix were transformed into nodes and edges of a Bayesian network, forming its structural backbone. Using this backbone, risk events extracted from the HAZOP analysis were inserted into the Bayesian network as risk nodes, resulting in the incorporation of 260 risk scenarios. The complete Bayesian network generated is shown in Fig. 3 and consists of a total of 358 nodes, including 98 alarm tag nodes, 260 event nodes, 62 cause nodes, 61 consequence nodes, and 137 protection layer nodes, with 426 dependency relationships. The proportions of different node types in the Bayesian network are illustrated in Fig. 3.

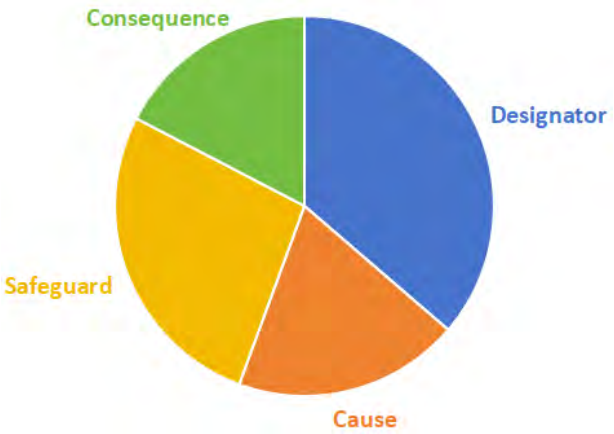

Figure 3. Nodes proportions in BN.

Based on the Bayesian network framework established using DCS time-series data and risk events, risk probabilities are matched from the risk database using a Transformer–BERT approach. The cause and consequence nodes extracted from the HAZOP analysis table can be fully and accurately associated with their corresponding probabilities. In contrast, the failure probabilities of protection layer nodes need to be matched from the risk database, with a matching accuracy of over 80%. For some protection measures, accurate and precise failure probabilities cannot be obtained from the risk database. However, this level of uncertainty is acceptable, as subsequent parameter learning is performed using actual operational data. These operational data include DCS time-series data of critical variables as well as fault records from real plant operations.

The final BN auto-constructed is shown in Fig.4, where the yellow nodes stand for hazardous causes, gray for protection layers, red for consequences, and blue for designators. This BN can be used for both forward risk inference and backward root cause tracing. Since the FCC unit considered in this case study is relatively new and no severe accident events have occurred, unplanned shutdowns are used as representative risk scenarios for analysis and validation. Taking the low–low reactor temperature interlock feed cut-off event as an example, the inference result was shown in Fig. 5. Backward risk tracing indicates that the event may have been caused by changes in feed flow rate, an increase in reactor nozzle pressure, or a failure-open condition of WICA10201. The most probable cause is lower feed flow rate, which is further caused by an increase in reactor nozzle pressure. These inferred causes are consistent with those identified in the corresponding accident analysis report.

5. CONCLUSION

Figure 4. Auto-constructerd BN of FCC unit.

Figure 5. Risk evolution inferred by part of the BN.

This paper presents an automated framework for constructing Bayesian causal networks to enable dynamic risk assessment in complex chemical processes. By integrating time-series Transformer–based causal discovery with multi-source risk knowledge from HAZOP reports and reliability databases, the proposed approach effectively eliminates the heavy reliance on manual modeling and expert-driven network construction in conventional Bayesian network–based risk assessment.

The proposed method automatically identifies causal structures from DCS time-series data, systematically embeds risk causes, consequences, and protection layers extracted from HAZOP analyses, and assigns node probabilities through semantic matching using Sentence-BERT. Combined with Bayesian parameter learning based on real operational data, the resulting network is able to continuously adapt to process dynamics and reflect real-time risk evolution.

A case study on an industrial fluid catalytic cracking unit demonstrates that the proposed framework can accurately construct a large-scale Bayesian causal network and support both forward risk prediction and backward root cause diagnosis. The results confirm that the method significantly reduces modeling effort while providing a comprehensive and reliable representation of dynamic risk propagation in chemical plants. Overall, this work offers a practical and scalable solution for automated dynamic risk assessment and has strong potential for industrial safety management applications.

ACKNOWLEDGEMENTS

This work was supported by the National Key R&D Program of China under Grant 2024YFC3013600.

REFERENCES

1. Khan F, Hashemi SJ, Paltrinieri N, Amyotte P, Cozzani V, Reniers G. Dynamic risk management: a contemporary approach to process safety management. Current Opinion in Chemical Engineering 14:9-17 (2016). https://doi.org/10.1016/j.coche.2016.07.006
2. Meel A, Seider WD. Plant-specific dynamic failure assessment using bayesian theory. Chemical Engineering Science 61:7036-7056 (2006). https://doi.org/10.1016/j.ces.2006.07.007
3. Kalantarnia M, Khan F, Hawboldt K. Modelling of BP texas city refinery accident using dynamic risk assessment approach. Process Safety and Environmental Protection 88:191-199 (2010). https://doi.org/10.1016/j.psep.2010.01.004
4. Montani S, Portinale L, Bobbio A, et al. DBNet, a tool to convert dynamic fault trees into dynamic

Bayesian networks[M]//Technical Report, Dipartimento di Informatica, Università del Piemonte Orientale. 2005.
5. Kohda T, Cui W. Risk-based reconfiguration of safety monitoring system using dynamic bayesian network. Reliability Engineering & System Safety 92:1716-1723 (2007). https://doi.org/10.1016/j.ress.2006.09.012
6. Shan X, Liu K, Sun PL. Risk analysis on leakage failure of natural gas pipelines by fuzzy bayesian network with a bow-tie model. Scientific Programming 2017:1-11 (2017). https://doi.org/10.1155/2017/3639524
7. Guo X, Ji J, Khan F, Ding L, Tong Q. A novel fuzzy dynamic bayesian network for dynamic risk assessment and uncertainty propagation quantification in uncertainty environment. Safety Science 141:105285 (2021). https://doi.org/10.1016/j.ssci.2021.105285
8. Adedigba SA, Khan F, Yang M. Dynamic failure analysis of process systems using neural networks. Process Safety and Environmental Protection 111:529-543 (2017). https://doi.org/10.1016/j.psep.2017.08.005
9. Wang Q, Han J, Chen F, Zhang X, Yun C, Dou Z, Yan T, Yang G. Real-time risk prediction of chemical processes based on attention-based bi-lstm. Chinese Journal of Chemical Engineering 75:131-141 (2024). https://doi.org/10.1016/j.cjche.2024.06.026
10. Zadakbar O, Imtiaz S, Khan F. Dynamic risk assessment and fault detection using principal component analysis. Ind. Eng. Chem. Res. 52:809-816 (2012). https://doi.org/10.1021/ie202880w
11. Li S, Chen G, Men J, Li X, Zhao Y, Xu Q, Zhao J. A novel dynamic risk assessment method for hazardous chemical warehouses based on improved SVM and mathematical methodologies. Journal of Loss Prevention in the Process Industries 89:105302 (2024). https://doi.org/10.1016/j.jlp.2024.105302
12. Bi X, Wu D, Xie D, Ye H, Zhao J. Large-scale chemical process causal discovery from big data with transformer-based deep learning. Process Safety and Environmental Protection 173:163-177 (2023). https://doi.org/10.1016/j.psep.2023.03.017

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

CO_2 Conversion: Three-Dimensional Modelling of Gas Diffusion Electrodes

Cristina González-Fernández[a]*, Camilo Peralta[b], Jose Antonio Abarca[a], Esther Santos[b], Guillermo Díaz-Sainz[a], Ángel Irabien[a]

[a] Departamento de Ingenierías Química y Biomolecular, Universidad de Cantabria, Avda. Los Castros s/n, 39005 Santander, Spain
[b] Apria Systems S.L. Parque Empresarial de Morero, 39611 Guarnizo, Spain
* Corresponding Author: cristina.gonzalezfdez@unican.es

ABSTRACT

Electrochemical reduction of CO_2 ($ERCO_2$) in gas diffusion electrode (GDE)-based electrolyzers represents a potential strategy for global decarbonization, achieving simultaneously the valorization of this abundant carbon resource. While significant progress has been achieved in enhancing CO_2 conversion in these systems, further advances are required to enable their practical implementation at the industrial scale. Physics-based simulations offer a powerful tool to guide the optimization of design and operating parameters as well as for the efficient scale-up of CO_2 electrolyzers. In this work, we have developed a three-dimensional multiphysics model of the cathodic compartment of a GDE electrolyzer for $ERCO_2$ to formate. For that purpose, the software COMSOL Multiphysics has been used. The model is experimentally validated, confirming its accuracy at reproducing current density and Faradaic efficiency at cathode potentials in the range -1.2 V and -1.7 V. Moreover, kinetic parameters are fitted to experimental data performing several parametric sweeps to minimize discrepancies between simulated and measured current densities. We obtain charge transfer coefficients ($\alpha_{c,k}$) of 0.07 and 0.34, together with exchange current densities ($i_{0,k}$) of 30 mA·cm^{-2} and 10^{-4} mA·cm^{-2} for $ERCO_2$ and the competing hydrogen evolution reaction, respectively. Finally, the model is used to predict formate concentration under varying applied potential conditions. Collectively, our three-dimensional multiphysics model reliably predicts CO_2 conversion to formate, thus representing a useful tool for guiding system optimization and scale-up. Moreover, the systematic methodology followed for developing the model can be readily extended to the design and analysis of other electrochemical cells beyond CO_2-to-formate conversion.

Keywords: Carbon Dioxide, Computational Fluid Dynamics, Catalysis, Modelling and Simulations

INTRODUCTION

Anthropogenic CO_2 emissions have triggered severe environmental concerns and represent a serious concern worldwide. Hence, there is an urge to implement actions that alleviate these emissions and develop alternative pathways for CO_2 utilization. [1-3] In this context, carbon capture and utilization (CCU) is recognized as a potential strategy for curbing energy and industry-related emissions. [4, 5]

Electrochemical reduction of CO_2 ($ERCO_2$) represents a promising CCU technology for CO_2 conversion into value-added products and fuels. [4, 6, 7] Thereby, $ERCO_2$ aligns with circular economy by enhancing the sustainability of industrial processes using waste CO_2 as a raw material. [5, 7]

The application prospects of $ERCO_2$ stem from its appealing advantages, including (i) mild operating conditions (room temperature and pressure), (ii) use of non-toxic green electrolytes, (iii) obtention of various carbon-based product (e.g., carbon monoxide, formic acid, ethanol, etc.) by controlling the reaction conditions, (iv) simplicity of electrochemical systems, (v) process repeatability, (vi) scalability, and (vii) possibility of integrating with low-carbon energy sources (e.g., solar, wind, geothermical, etc.) photovoltaics. [2, 4, 6]

From the different products that can be derived from $ERCO_2$, formate ($HCOO^-$) has received special attention due to its industrial relevance and possible use as fuel source. Hence, significant research efforts have been devoted in recent years to convert CO_2 into $HCOO^-$, thus advancing toward large-scale implementation. [7, 8] This progress has been mainly focused on the electrolyzer configuration, catalyst, electrode fabrication and optimizing reaction conditions. In this sense, different cathode configurations have been proposed to enhance electrolyzer performance and ensure closer alignment to industrial requirements (i.e., operating at current densities higher than 200 $mA \cdot cm^{-2}$). [2, 5, 7] In this regard, gas diffusion electrodes (GDEs) exhibit outstanding performance, since they alleviate mass transfer limitations. Particularly, GDEs provide a three-phase boundary (gas-liquid-solid), thus facilitating the accessibility of CO_2 to the catalyst surface. GDEs comprise two layers: (i) the gas diffusion layer (GDL), which acts as a porous support that promotes CO_2 mass transport, and (ii) the catalyst layer (CL), which is deposited onto the GDL and is the component of the GDE where the electrochemical reactions occur. [3, 5, 7]

Optimizing the performance of GDE electrolyzers requires rationalizing the effect of multiple process, operation and design variables. Hence, addressing this complex investigation through a trial-and-error approach could result impractical. [9, 10] Physics-based modeling provides fundamental understanding about the phenomena underlying $ERCO_2$; hence, it can guide electrolyzer optimization prior to their fabrication. Several one- and two-dimensional models of GDE electrolyzers have been developed, which have provided important insights so far. [10, 11] Nevertheless, advancing electrolyzers design and scale-up calls for gaining a more realistic understanding of the system performance so that currently relevant issues, such as the effect of flow distributors, CO_2 partial pressure in the feed or pH, could be reliably addressed.

In this study, a three-dimensional multiphysics model is developed for the cathodic compartment of a GDE electrolyzer for $ERCO_2$ to $HCOO^-$. The model has been first validated by comparing simulation predictions of current density and Faradaic efficiency with experimental data of our previous work. Moreover, the kinetic parameters of the involved electrochemical reactions ($ERCO_2$ and hydrogen evolution reaction) have been obtained. Finally, the model is used to predict formate concentration under varying potential conditions. Overall, the developed model provides a powerful framework for gaining insights into the different phenomena underlying electrochemical CO_2 valorization and represents a useful tool for guiding optimization and scale-up of GDE-electrolyzers.

MODEL DEVELOPMENT

System description

The cathodic compartment of the GDE-electrolyzer modeled in this study derives from the work of Díaz-Sainz et al. as illustrated in Figure 1. [12] In this system, gaseous CO_2 and 0.5M KCl + 0.45M aqueous $KHCO_3$ solution are fed through the cathode gas and catholyte channels, respectively. CO_2 flows through the GDL into the CL, where it meets the electrolyte and dissolves in it. The $ERCO_2$ and the competing hydrogen evolution reaction (HER) take place in the CL, yielding the desired $HCOO^-$ (eq. 1) and H_2 (eq. 2), respectively:

$$CO_2 + 2e^- \rightarrow HCOO^- + OH^- \; (E_0 = -0.02 \; vs. RHE) \quad (1)$$

$$2H^+ + 2e^- \rightarrow H_2 \; (E_0 = 0 \; V \; vs. RHE) \quad (2)$$

Subsequently, the electrolyte along with $HCOO^-$ and unreacted CO_2 exit the cell through the catholyte outlet.

In addition to the electrochemical reactions, homogeneous reactions take place in the electrolyte:

$$CO_{2(aq)} + H_2O \leftrightarrow H^+ + HCO_3^- \quad (3)$$

$$HCO_3^- \leftrightarrow H^+ + CO_3^{2-} \quad (4)$$

$$CO_{2(aq)} + OH^- \leftrightarrow HCO_3^- \quad (5)$$

$$HCO_3^- + OH^- \leftrightarrow H_2O + CO_3^{2-} \quad (6)$$

$$H_2O \leftrightarrow H^+ + OH^- \quad (7)$$

These reactions are fast, and thus, they are considered as equilibrium reactions. They balance the concentration of H^+, OH^- and carbon-containing species. [13]

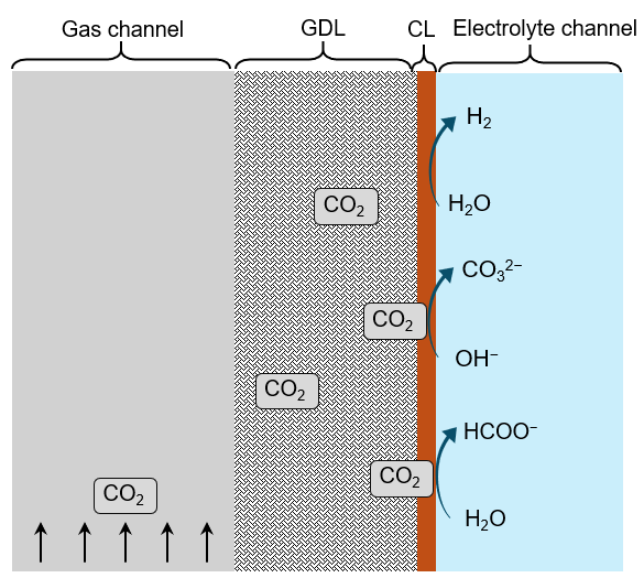

Figure 1. Schematic diagram of the $ERCO_2$ to $HCOO^-$ process modeled in this work.

As seen from Figure 1, this study explicitly models

four distinct domains, namely: gas and catholyte channels, GDL and CL. The cell dimensions and operating conditions have been included in Table 1:

Table 1: Model parameters for cell geometry and operating conditions.

Parameter	Value	Unit
Cell length and width	30.2	mm
GDL thickness	325	µm
CL thickness	15	µm
Gas and catholyte channels thickness	2.2	mm
Temperature	298.15	K
Pressure	1	atm
Feed gas flow rate	200	$ml \cdot min^{-1}$
Catholyte flow rate per geometric surface area	0.57	$ml \cdot min^{-1} \cdot cm^{-2}$
Cathode applied potential	-1.2 to -1.7	V (vs. Ag/AgCl)

Model assumptions

The following assumptions are considered in the developed model:

- The system operates at steady state and isothermal (298.15 K) conditions
- Fluid flow remains in the laminar regime, with the gas phase treated as weakly compressible and the liquid phase as incompressible flow
- The physical properties of GDL and CL are isotropic with constant porosity and permeability.
- Zero-slip and no-species flux are applied to the side walls of the cell
- Kinetics of electrochemical reactions are described by Butler-Volmer equations

Governing equations

The governing equations for modeling $ERCO_2$ to $HCOO^-$ in a GDE-electrolyzer include both transport equations and electrochemical reaction kinetics. A brief description of these equations is provided below.

Transport equations

The transport equations include mass, momentum and charge as follows. In these equations, EC (electrolyte channel), CG (gas channel), GDL (gas diffusion layer), and CL (catalyst layer) denote the domains where they are applied.

$$\nabla \cdot (\rho_g u_g) = 0 \, (\mathrm{GC}) \tag{8}$$

$$\nabla \cdot (\rho u) = 0 \, (\mathrm{EC}) \tag{9}$$

$$\rho_g(u_g \nabla) = \nabla p_g + \mu_g\left(\nabla u_g + (\nabla u_g)^T\right) - \frac{2}{3}\mu_g(\nabla u_g)(GC) \tag{10}$$

$$\rho(u\,\nabla)u = \nabla p + \mu(\nabla u + (\nabla u)^T)\,(EC) \tag{11}$$

$$\frac{\rho_g}{\varepsilon} u_g \cdot \nabla \frac{u_g}{\varepsilon} = -\nabla p_g + \nabla \cdot \left[\frac{\mu_g}{\varepsilon}\left(\nabla u_g + (\nabla u_g)^T\right) - \frac{2}{3}\mu_g(\nabla u_g) I\right] - \frac{\mu_g}{k} u_g (GDL/CL) \tag{12}$$

$$\nabla \cdot n_i = 0 \, (EC/GC) \tag{13}$$

$$\nabla \cdot n_i = R_i \, (GDL/CL) \tag{14}$$

$$n_i = -D_i \nabla C_i + Z_i F \frac{D_i}{RT} C_i \nabla \varphi_l + C_i u \, (EL/CL) \tag{15}$$

$$\nabla \cdot i = 0 \, (GDL/CL) \tag{16}$$

where ρ, μ, u and p are the density, viscosity, velocity and pressure of the electrolyte, whereas ρ_g, μ_g, u_g and p_g represent the density, viscosity, velocity and pressure of the gaseous stream; n_i, D_i and C_i are the total mass flux, diffusivity and local concentration of species i (CO_2, $HCOO^-$, H_2, H_2O, H^+, OH^-); φ_l is the electrolyte potential; *i* represents the current density of species; F denotes the Faraday's constant; Z_i is the electron transfer number; R is the gas constant and T the temperature.

Gas-phase transport is modeled as the sum of convection and diffusion. Diffusion is driven by gradients in species mass fraction and mixture molar mass, and it is governed by an effective diffusivity. This diffusivity combines Stefan–Maxwell diffusion (multicomponent gas interactions) and Knudsen diffusion (pore-scale effects), whose resistances are added in parallel.

$$n_i = j_i + \rho_i u_g \tag{17}$$

$$j_i = -\rho_g D_i^{eff}\left(\nabla \omega_i + \omega_i \frac{\nabla M_n}{M_n}\right) \tag{18}$$

$$\frac{1}{D_i} = \frac{1}{D_i^m} + \frac{1}{D_i^k} \tag{19}$$

$$D_i^m = \frac{1-\omega_i}{\sum_{n \neq i} \frac{y_n}{D_{in}}} \tag{20}$$

$$D_i^k = \frac{2 r_{p,m}}{3}\sqrt{\frac{8RT}{\pi M_i}} \tag{21}$$

where n_i is the total mass flux of species i, j_i is the diffusive mass flux of species i, ρ_i is the mass density of species i, u_g is the average mass velocity of the gas mixture, ρ_g is the density of the gas mixture, ω_i is the mass fraction of species i, D_i^{eff} is the effective diffusion coefficient of species i, M_i is the molar mass of species i, M_n is the average molar mass of the gas mixture, D_i is the overall diffusion coefficient of species i, D_i^m is the Stefan–Maxwell diffusivity, and D_i^k is the Knudsen diffusivity.

Electrochemical kinetics

The electrochemical kinetic reactions occurring at the cathode ($ERCO_2$ to $HCOO^-$ and HER) are modeled through the widely used Butler-Volmer equations. Particularly, concentration-dependent and concentration-

independent Butler-Volmer equations have been considered for $ERCO_2$ and HER, respectively:

$$i_{HCOO^-} = -i_{0,HCOO-}\left(\frac{c_{CO_2,aq}}{c^{ref}_{CO_2,aq}}\right)^{\gamma} exp\left(-\frac{\alpha_{c,ERCO2}F}{RT}\eta_{ERCO_2}\right) \quad (22)$$

$$i_{HER} = -i_{0,HER}\, exp\left(-\frac{\alpha_{c,HER}F}{RT}\eta_{HER}\right) \quad (23)$$

$$n_i = \varphi_s - \varphi_l - (U_k^0 - \frac{2.303}{F}\mathrm{pH}) \quad (24)$$

$$R_k = \frac{\nu_k i_k}{n_k F} \quad (25)$$

$$FE = \frac{i_k}{i_{total}} \times 100 \quad (26)$$

$$i_{total} = \sum_k i_k = i_{HCOO^-} + \mathrm{i_{HER}} \quad (27)$$

where $i_{0,k}$, $\alpha_{c,k}$ and η_k are the exchange current density, charge transfer coefficient and activation overpotential for reaction k; $c_{CO_2,aq}$ is the local CO_2 concentration; $c^{ref}_{CO_2,aq}$ represents the reference concentration of CO_2 (1M); γ is the rate order of CO_2 to obtain $HCOO^-$; U_k^0 is the reduction potential of reaction k; φ_s denotes the solid potential; φ_l is the electrolyte potential; ν_k represents the stoichiometric coefficient of the species; n_k is the number of participating electrons; i_{total} is the total current density.

CO_2 absorption

At the GDL-CL interface, gaseous CO_2 and liquid electrolyte meet. Hence, CO_2 is absorbed into the electrolyte, which is described as follows:

$$R_{PT,CO_2} = a_V k_{GL,CO_2} M_{CO_2}(H_{CO_2} p_G y_{CO_2} - C_{CO_2(l)}) \quad (28)$$

where a_V is the specific interfacial area, k_{GL,CO_2} is the gas-to-liquid mass transfer coefficient, H_{CO_2} represents the Henry's law constant, y_{CO_2} is the mole fraction of gaseous CO_2 and $C_{CO_2(l)}$ is the concentration of dissolved CO_2 in the electrolyte.

Numerical method

The previous equations are solved at steady-state in COMSOL Multiphysics 6.2. Particularly, the "Free and porous media flow, Brinkman", "Transport of concentrated species", "Transport of diluted species", and "Secondary current distribution" physics interfaces are selected to account for the different phenomena underlying $ERCO_2$. A physics-controlled mesh was used; the resulting mesh contains 42844 boundary and 219785 domain elements, with an average element quality close to unity. The model is solved using MUMPS solver with a relative tolerance of 5×10^{-4} through a segregated approach.

To enhance model's convergence, the stationary solver was run in three sequenced steps. First, a "Current distribution and initialization" step is performed. Second, fluid flow is solved since it is not affected by other interfaces of the model. Then, species transport coupled to flow and current distribution are computed simultaneously to derived the final results. Finally, parametric sweeps are conducted for exploring the influence of several operating conditions on cell performance.

Results and discussion

Model validation and parametric estimation

Our previous work on the design of filter-press reactors for $ERCO_2$ to $HCOO^-$ is considered for obtaining the kinetic parameters of the involved electrochemical reactions. [12] Starting from literature-based initial values [3, 4, 6], a trial-and-error method is applied by performing parametric sweeps in COMSOL Multiphysics. In particular, the kinetic parameters for $ERCO_2$ and HER are iteratively modified, and the computational results are compared with the experimental data. These parameters are refined to minimize discrepancies between simulated and experimentally measured current densities and Faradaic efficiencies. Table 2 gathers the kinetic parameters (namely charge transfer coefficient, α_c and exchange current density, i_0) that have been obtained for $ERCO_2$ and HER. These values align closely with previously reported kinetic parameters for the involved electrochemical reactions.

Table 2: Electrochemical kinetic parameters.

Reaction	Parameter	Value	Unit
$ERCO_2$	α_c	0.07	-
	i_0	30	$mA\cdot cm^{-2}$
HER	α_c	0.34	-
	i_0	10^{-4}	$mA\cdot cm^{-2}$

The kinetic parameters derived for $ERCO_2$ and HER lead to successful model validation, as evidenced from Figure 2 which illustrates the simulated and experimental current density and Faradaic efficiency for varying cathode potential. Thereby, simulation results exhibit good agreement to experimental data for potentials within the range of -1.2 V to -1.7 V, where CO_2 electrolyzers for $HCOO^-$ typically operate, as experimentally demonstrated. Particularly, the deviation between experimental and modeled current density and Faradaic efficiency is below 3%, which demonstrates the reliability of the developed model for capturing the phenomena underlying $ERCO_2$.

Electrochemical production of $HCOO^-$

The validated model is used to investigate species distribution in the electrolyzer and predict $HCOO^-$ concentration at the system outlet when applying different cathode potentials. Figure 3 depicts 2D contours of $HCOO^-$ concentration within the electrode and catholyte regions at cathode potentials of -1.2 and -1.7 V. At both potentials, the highest $HCOO^-$ concentration is localized near the edges of the electrode surface. Under an applied

potential of -1.2 V, $HCOO^-$ concentration remains uniform at low values in the catholyte. Conversely, increasing the cathode potential to -1.7V results in a pronounced concentration gradient across the cell width, and significantly higher $HCOO^-$ concentration in the catholyte. This behaviour reflects enhanced electroreduction kinetics at more negative applied potentials.

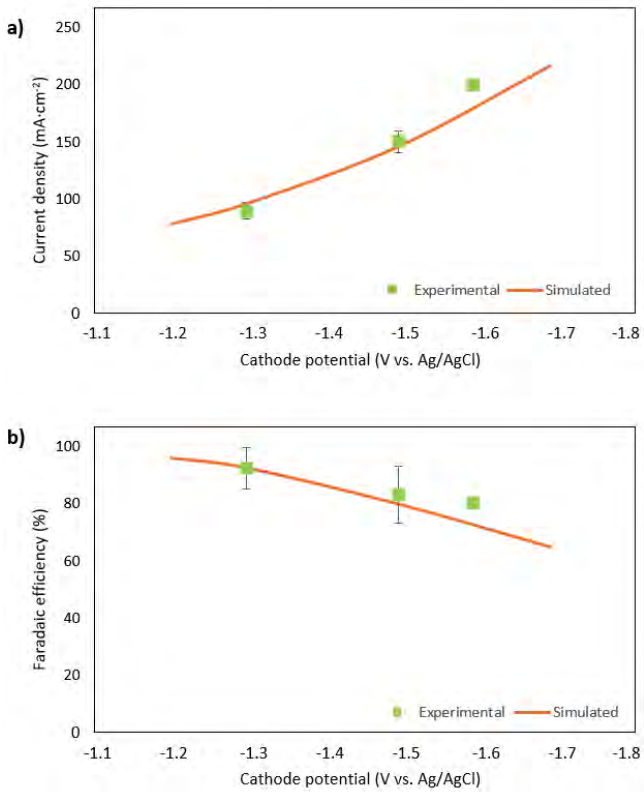

Figure 2. Experimental and simulated results for (a) current density and (b) Faradaic efficiency.

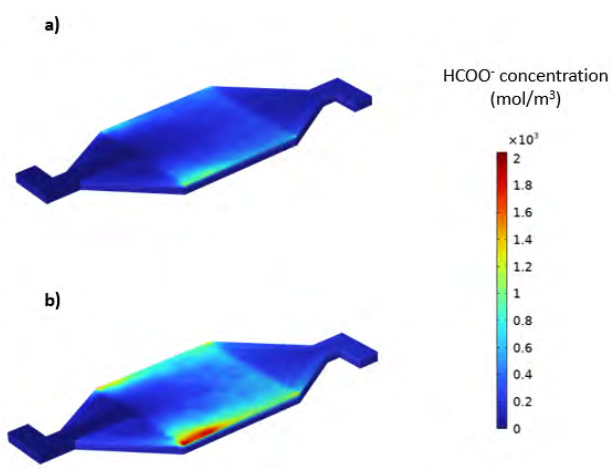

Figure 3. Distribution of $HCOO^-$ concentration in the CL and catholyte at two applied cathode potentials: (a) -1.2V, and (b) -1.7V.

Finally, the model is exploited to predict $HCOO^-$ concentration at the system outlet for different applied cathode potentials (Figure 4). It can be easily noticed that there is a strong dependence of $HCOO^-$ production on the applied potential. Particularly, $HCOO^-$ concentration increases from 2 to 3.7 $g \cdot L^{-1}$ when the cathode potential varies in the range -1.2V - -1.7V, representing approximately a two-fold increase. While negative potentials favor CO_2 conversion they entail increased energy consumption and enhancement of side reactions. Hence, predicting $HCOO^-$ concentration at various cathode potentials can guide the optimization of the operating conditions to achieve a trade-off between the desired concentration of the target product and overall system efficiency.

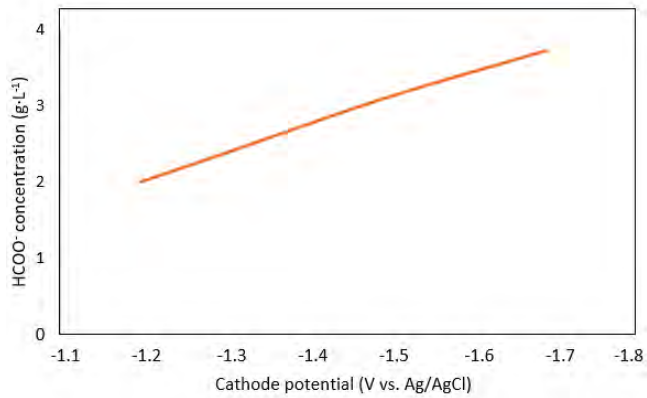

Figure 4. $HCOO^-$ concentration at the system outlet.

CONCLUSION

This study introduces a three-dimensional multiphysics model of the cathodic compartment of a GDE electrolyzer for $ERCO_2$ to $HCOO^-$. The model reliably reproduces experimental observations, since deviations lower than 3% are reported for current density and Faradaic efficiency toward $HCOO^-$ at different applied potentials. Additionally, the three-dimensional model provides insights into the distribution of the involved species in the electrolyzer and enables prediction of $HCOO^-$ concentration at the system outlet under varying cathode potentials. Hence, the developed model can be used to systematically investigate the effect of several design and operating variables on the system performance, including CO_2 partial pressure in the feed, inlet flow rate, or the geometry of flow distributors. Moreover, it represents a useful tool for guiding system optimization and scale-up. Additionally, the systematic methodology followed for developing the model can be readily extended to the design and analysis of other electrochemical cells beyond CO_2-to-formate conversion.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge financial support through project PCI2024-155027-2, funded by MICIU/AEI/10.13039/501100011033 and co-funded by the European Union. The present work is related to CAPTUS Project. This project has received funding from the European Union's Horizon 2020 - Research and Innovation Framework Programme under grant agreement No 101118265. Jose Antonio Abarca gratefully acknowledges the predoctoral research grant (FPI) PRE2021-097200. Camilo Peralta gratefully acknowledges funding from the EU's Horizon 2021 program under the Marie Skłodowska-Curie Doctoral Networks (MSCA-DN) grant agreement No 101072830 (ECOMATES).

AUTHOR IDENTIFIERS

Author ORCIDs:
Cristina González-Fernández: 0000-0002-1571-057X
Camilo Peralta: 0009-0002-6587-4342
José Antonio Abarca: 0000-0003-0120-8682
Esther Santos: 0000-0001-9370-141X
Guillermo Díaz-Sainz: 0000-0002-8687-300X
Ángel Irabien: 0000-0002-2411-4163

REFERENCES

1. Gao D, Arán-Ais RM, Jeon HS, Roldan Cuenya B. Rational catalyst and electrolyte design for CO2 electroreduction towards multicarbon products. Nat Catal 2:198-210 (2019). https://doi.org/10.1038/s41929-019-0235-5
2. He R, Xu N, Hasan IMU, Peng L, Li L, Huang H, Qiao J. Advances in electrolyzer design and development for electrochemical co₂ reduction. EcoMat 5: (2023). https://doi.org/10.1002/eom2.12346
3. O. A. El-Shafie, R. M. El-Maghraby, J. Albo, S. E. K. Fateen, A. Abdelghany. Modeling and numerical investigation of the performance of gas diffusion electrodes for the electrochemical reduction of carbon dioxide to methanol. *Ind. Eng. Chem. Res.*, 59(47), 20929-20942 (2020) https://dx.doi.org/10.1021/acs.iecr.0c02358
4. Offong AC, Hanak DP. Modeling and simulation of microfluidic electrolytic cells for CO2 electro-reduction to formic acid: the influence of a bi-sn catalyst and ionic liquid electrolyte on cell performance. Next Energy 7:100276 (2025). https://doi.org/10.1016/j.nxener.2025.100276
5. Abarca JA, Díaz-Sainz G, Irabien A. Inhibiting salt precipitation on the gas diffusion electrode surface in gas-phase CO2 electroreduction to formate by using an acidic anolyte. Journal of CO2 Utilization 86:102897 (2024). https://doi.org/10.1016/j.jcou.2024.102897
6. Agliuzza M, Pirri CF, Sacco A. A comprehensive modeling for the co₂ electroreduction to CO. J. Phys. Energy 6:015004 (2023). https://doi.org/10.1088/2515-7655/ad0a39
7. Abarca JA, Warmuth L, Rieder A, Dutta A, Vesztergom S, Broekmann P, Irabien A, Díaz-Sainz G. GDE stability in co₂ electroreduction to formate: the role of ionomer type and loading. ACS Catal. 15:8753-8767 (2025). https://doi.org/10.1021/acscatal.5c02052
8. Díaz-Sainz G, Abarca JA, Alvarez-Guerra M, Irabien A. Exploring the impact of partial pressure and typical compounds on the continuous electroconversion of CO2 into formate. Journal of CO2 Utilization 81:102735 (2024). https://doi.org/10.1016/j.jcou.2024.102735
9. Bahreini M, Désilets M, Pahija E, Legrand U, Guo J, Fink AG. Investigation of co₂ reduction to formate in an industrial-scale electrochemical cell through transient numerical modeling. Ind. Eng. Chem. Res. 63:18187-18198 (2024). https://doi.org/10.1021/acs.iecr.4c03239
10. Weng LC, Bell AT, Weber AZ. A systematic analysis of cu-based membrane-electrode assemblies for co₂ reduction through multiphysics simulation. Energy Environ. Sci. 13:3592-3606 (2020). https://doi.org/10.1039/d0ee01604g
11. Agarwal VG, Haussener S. Quantifying mass transport limitations in a microfluidic CO2 electrolyzer with a gas diffusion cathode. Commun Chem 7: (2024). https://doi.org/10.1038/s42004-024-01122-5
12. Díaz-Sainz G, Alvarez-Guerra M, Solla-Gullón J, García-Cruz L, Montiel V, Irabien A. CO2 electroreduction to formate: continuous single-pass operation in a filter-press reactor at high current densities using bi gas diffusion electrodes. Journal of CO2 Utilization 34:12-19 (2019). https://doi.org/10.1016/j.jcou.2019.05.035
13. Corpus KRM, Bui JC, Limaye AM, Pant LM, Manthiram K, Weber AZ, Bell AT. Coupling covariance matrix adaptation with continuum modeling for determination of kinetic parameters associated with electrochemical CO2 reduction. Joule 7:1289-1307 (2023). https://doi.org/10.1016/j.joule.2023.05.007

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Multiscale Modeling of PHBV Production: Explicit Polymerization Modeling and Improved Prediction of Chain Length Distributions

Stefan Hempfling[a]*, Rudolph Kok[b], Stefanie Duvigneau[c], Achim Kienle[bc], and Robert Dürr[a]

[a] Engineering Mathematics, Magdeburg-Stendal University of Applied Sciences, Magdeburg, Germany
[b] Institute of Automation Engineering, Otto von Guericke University Magdeburg, Magdeburg, Germany
[c] Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany
* Corresponding Author: stefan.hempfling@h2.de

ABSTRACT

Multiscale models provide a powerful framework to link bioprocess operation conditions with polymer microstructure, yet their predictive capability for polymer attributes such as chain length distributions (CLDs) remains limited. In this work, an advanced multiscale modeling framework for the microbial production of poly(3-hydroxybutyrate-co-3-hydroxyvalerate) (PHBV) in *Cupriavidus necator* is presented, targeting the quantitative prediction of polymer microstructure. The model consistently integrates a structured macroscopic kinetic description of substrate uptake, biomass growth, and copolymer accumulation with an explicitly formulated microscopic polymerization model resolving initiation, propagation, termination, and depolymerization reactions of living and dead chains. A central contribution of this study is the quantitative calibration of the polymerization kinetics based on experimental size-exclusion chromatography (SEC) data. Polymerization rate constants were identified by fitting simulated CLDs to measurements obtained at multiple diagnostic time points during fed-batch cultivations using fructose and propionic acid as carbon sources. Parameter estimation was performed using a hybrid multi-start optimization strategy combining simulated annealing with Gauss–Newton refinement. Compared to literature-based parameter sets, the optimized model accurately reproduces both the temporal evolution of the weight-average molecular weight and the full CLD shape, including the emergence of high-molecular-weight tails. The results demonstrate that explicit polymerization modeling combined with parameter identification is essential for quantitatively linking process dynamics to PHBV microstructure. The proposed framework provides a sophisticated basis for model-based optimization of feed strategies and supports quality-by-design approaches for sustainable biopolymer production.

Keywords: Modelling and Simulations, Polymers, Optimization, Multiscale Modelling, Process Design

INTRODUCTION

Microbial polyhydroxyalkanoates (PHAs) are biodegradable and bio-based polymers that represent a sustainable alternative to fossil-based plastics [1]. Among them, the copolymer poly(3-hydroxybutyrate-co-3-hydroxyvalerate) (PHBV) is of particular interest due to its enhanced flexibility, lower melting temperature, and reduced crystallinity compared to the homopolymer poly(3-hydroxybutyrate) (PHB) [2, 3]. These characteristics improve processability and broaden the application range from packaging to agricultural and biomedical materials. However, despite these advantages, the microbial production of PHAs remains expensive compared to petroleum-based polymers, with reported production costs for PHAs typically exceeding 2.0 US$/kg, compared to average market prices of approximately 1.5 US$/kg for conventional plastics such as polypropylene, creating a strong need for improved process design and optimization strategies [4, 5].

Research on PHBV production should not only focus on maximizing polymer yield at the process level but also

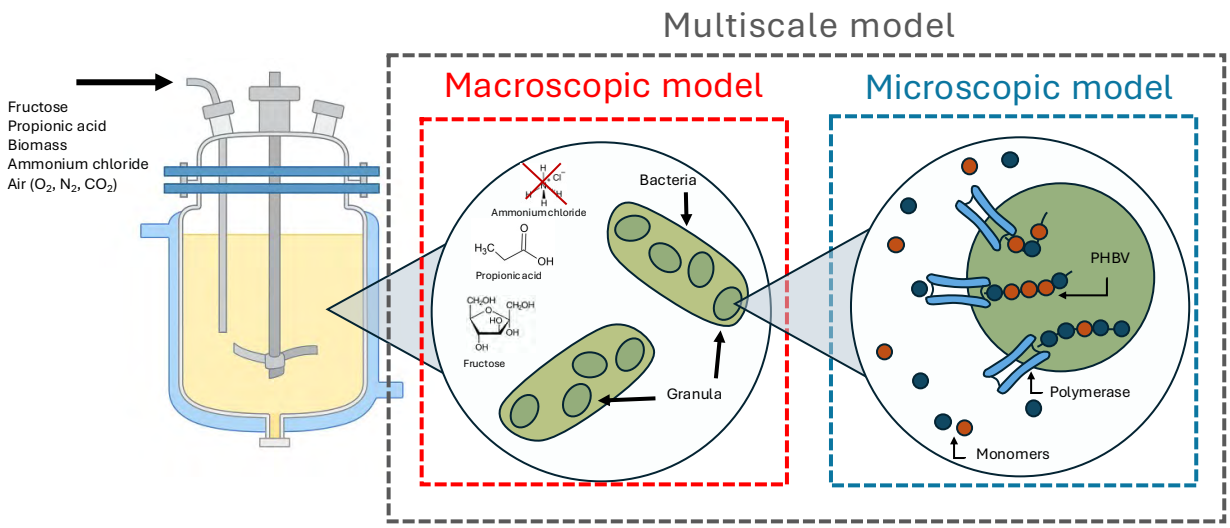

Figure 1: Schematic illustration of the multiscale modeling framework for PHBV production in a bioreactor. The macroscopic model (red box) describes reactor-scale dynamics, including substrates (fructose, propionic acid, ammonium chloride), biomass growth and product formation. The microscopic model (blue box) represents intracellular polymerization processes, where monomers are converted into PHBV by polymerases and stored within granules. Coupling both scales enables a consistent description of process dynamics and polymer structure evolution. Adapted from [11].

on tailoring end-use properties, which are determined by polymer quality attributes such as the chain length distribution (CLD), weight-average molecular weight (M_w), and polydispersity index (PDI). While molecular averages are commonly reported, they provide only limited insight into the underlying polymer microstructure. In contrast, the detailed information given in the full CLD strongly influences material properties such as mechanical strength and thermal stability. Accurate prediction of CLDs is therefore essential for enabling quality-by-design (QbD) approaches in biopolymer production [6]. Figure 1 provides an overview of the considered fed-batch process and the associated multiscale modeling framework, linking macroscopic dynamics with intracellular polymerization mechanisms.

Previous kinetic models of PHBV production successfully describe biomass growth, substrate consumption, and overall polymer accumulation; however, the polymer is described only by its total concentration, which limits the prediction of polymer quality properties [7]. To overcome this limitation, multiscale modeling approaches have been proposed that couple macroscopic process kinetics with microscopic polymerization models. A first multiscale framework for PHB demonstrated that cultivation conditions can be systematically linked to evolving CLDs by integrating process kinetics with polymerization modeling [8].

In contrast to PHB, PHBV production involves copolymerization of two monomer species, hydroxybutyrate (HB) and hydroxyvalerate (HV), whose relative incorporation depends on metabolic regulation and feed strategy. The use of propionic acid as a precursor for HV formation introduces additional complexity due to inhibitory effects on biomass growth and altered monomer availability. As a result, PHBV production exhibits a strong coupling between macroscopic process conditions and polymer microstructure, which poses additional challenges for predictive modeling.

The present work builds upon an existing multiscale PHBV modeling framework [9] and advances it towards quantitative prediction of polymer microstructure. The proposed approach provides an explicit formulation of the microscopic polymerization model, incorporating living and dead chain dynamics, monomer pools, and multiscale coupling. It also integrates these elements consistently with an updated macroscopic kinetic description within a single simulation environment. In this study, fructose and propionic acid are used as carbon sources, and the polymerization kinetics are quantitatively calibrated by identifying rate constants through fitting simulated chain length distributions to size-exclusion chromatography (SEC) data. This combined model formulation and parameter identification enables improved prediction of molecular weight evolution and CLD shape and establishes a consistent link between process dynamics and polymer quality.

MODEL FRAMEWORK

The multiscale modeling framework combines a macroscopic description of reactor-scale process dynamics with a microscopic representation of intracellular polymerization mechanisms. An overview of the considered process and the interaction between both modeling levels is shown in Figure 1.

The macroscopic kinetics of substrate uptake, biomass growth, and PHBV accumulation are based on the updated structured kinetic model proposed by Kok et al. [10], which refines the original formulation introduced by Duvigneau et al. [7]. The model describes the temporal evolution of key process variables, including fructose, propionic acid, ammonium chloride, biomass, and polymer concentrations. It explicitly accounts for CO_2-dependent metabolic activity, inhibitory effects of propionic acid on biomass growth, and steric limitations arising from intracellular polymer accumulation in granulas. These mechanisms allow the model to capture characteristic features of fed-batch PHBV production under varying substrate availability. For clarity, the macroscopic balance equations are not explicitly repeated and are adopted from the validated formulations reported in Kok et al. [10] and Duvigneau et al. [7].

The microscopic model resolves the intracellular polymerization of hydroxybutyrate (HB) and hydroxyvalerate (HV) monomers into PHBV chains. The polymerization kinetics and the associated multiscale coupling are implemented using the commercial software package Predici [12], which provides a mechanistic description of initiation, propagation, termination, and depolymerization reactions. The model explicitly tracks living and dead polymer chains as a function of their chain length, thereby enabling the direct simulation of chain length distributions.

Microscopic polymerization model

Living polymer chains terminated with either HB or HV are denoted as $LPHB_i$ and $LPHV_i$, respectively, where i represents the chain length and the last capital letter (B or V) corresponds to the terminal monomer (HB or HV) in the chain. Their temporal evolution is governed by initiation, propagation via monomer addition, and chain termination. The corresponding population balance equations read

$$\frac{d[LPHB_i]}{dt} = \delta_{\{i,1\}} k_{ini}[HB_{syn}] + k_{m2}[HB_{poly}]([LPHB_{i-1}] + [LPHV_{i-1}]) - k_{m2}([HB_{poly}] + [HV_{poly}])[LPHB_i] - k_c[LPHB_i] \quad (1)$$

$$\frac{d[LPHV_i]}{dt} = \delta_{\{i,1\}} k_{ini}[HV_{syn}] + k_{m2}[HV_{poly}]([LPHB_{i-1}] + [LPHV_{i-1}]) - k_{m2}([HB_{poly}] + [HV_{poly}])[LPHV_i] - k_c[LPHV_i] \quad (2)$$

Here, $\delta_{i,1}$ denotes the Kronecker delta accounting for chain initiation at i=1. Living chains grow by monomer addition with rate k_{m2} and are shifted from active to dead (inactive) population with rate k_c. Dead polymer chains, denoted as $DPHB_i$ and $DPHV_i$, are formed via change of living chains and subsequently shortened by depolymerization with rate k_{dep}. Their dynamics are described by

$$\frac{d[DPHB_i]}{dt} = k_c[LPHB_i] - k_{dep}[DPHB_i] + k_{dep} f_{HB}([DPHB_{i+1}] + [DPHV_{i+1}]) \quad (3)$$

$$\frac{d[DPHV_i]}{dt} = k_c[LPHV_i] - k_{dep}[DPHV_i] + k_{dep} f_{HV}([DPHB_{i+1}] + [DPHV_{i+1}]) \quad (4)$$

Since the microscopic model tracks only the terminal monomer unit of each chain, the identity of the newly exposed terminal unit after a depolymerization step is inferred statistically. The relative fractions of HB and HV as penultimate monomers are assumed to be proportional to the HB/HV composition predicted by the macroscopic model,

$$f_{HB} = \frac{c_{HB}}{c_{HB} + c_{HV}} \quad (5)$$

$$f_{HV} = \frac{c_{HV}}{c_{HB} + c_{HV}} \quad (6)$$

This approximation allows HB-HB, HB-HV, HV-HB and HV-HV terminal combinations to occur during depolymerization without explicit sequence tracking.

The intracellular pools of polymer-forming monomers HB_{poly} and HV_{poly} are consumed during initiation and propagation and replenished by the macroscopic monomer formation rates M_{HB}^{+} and M_{HV}^{+}, yielding

$$\frac{d[HB_{poly}]}{dt} = -k_{m1}[HB_{poly}] - k_{m2}[HB_{poly}] \sum_{i=1}^{\infty}([LPHB_i] + [LPHV_i]) + M_{HB}^{+} \quad (7)$$

$$\frac{d[HV_{poly}]}{dt} = -k_{m1}[HV_{poly}] - k_{m2}[HV_{poly}] \sum_{i=1}^{\infty}([LPHB_i] + [LPHV_i]) + M_{HV}^{+} \quad (8)$$

The monomer formation rates are defined in [8] and are adapted for both monomers in the present article. The monomer-synthase complexes are formed reversibly prior to chain initiation and are governed by

$$\frac{d[HB_{syn}]}{dt} = k_{m1}[HB_{poly}] - k_{ini}[HB_{syn}] \quad (9)$$

$$\frac{d[HV_{syn}]}{dt} = k_{m1}[HV_{poly}] - k_{ini}[HV_{syn}] \quad (10)$$

Multiscale coupling

To ensure consistency between the macroscopic and microscopic descriptions, both model layers are coupled via a mass-conservation constraint on the total HB and HV monomer mass [8]. At each time point, the polymerized monomer mass in the microscopic model is required to match the monomer mass predicted by the macroscopic kinetics. Instead of solving the resulting differential–algebraic system explicitly, a coupling strategy is employed. Here, the monomer formation rates M_{HB}^{+} and

M_{HV}^{+} supplied to the polymerization module are dynamically adjusted such that the microscopic polymer mass follows the macroscopic target.

A key advance of the present work is the consistent implementation of the entire multiscale workflow within a single simulation environment. Both the macroscopic process model, the polymerization kinetics, and the coupling strategy are implemented directly in Predici, enabling numerically robust dynamic simulations and providing a unified basis for subsequent parameter estimation and model validation.

PARAMETER ESTIMATION AND VALIDATION

A central novelty of this work is the explicit parameter identification for the polymerization kinetics. Literature-based rate constants adopted from PHB studies proved insufficient to reproduce experimentally observed PHBV chain length distributions (CLDs) [9]. Therefore, rate constants for initiation, propagation, termination, and degradation were identified directly from experimental size-exclusion chromatography (SEC) data measured at multiple sampling times during the fed-batch process.

The parameter estimation problem was formulated as a minimization of the deviation between simulated and experimental CLDs. Prior to error evaluation, both experimental and simulated CLDs were normalized to their respective maxima, and the simulated distributions were interpolated onto the experimental molar-mass grid. At each sampling time t_k, the discrepancy was quantified using the root mean square error (RMSE),

$$RMSE(t_k) = \sqrt{\frac{1}{N}\sum_{j=1}^{N}\left(\widetilde{w}_{sim}\left(M_j^{exp}, t_k\right) - \widetilde{w}_{exp}\left(M_j^{exp}, t_k\right)\right)^2}, \quad (11)$$

where $\widetilde{w}_{sim}$ and $\widetilde{w}_{exp}$ denote the normalized weight-based CLDs of simulation and experiment, respectively, evaluated at the experimental molar-mass points M_j^{exp}.

To obtain a single scalar performance measure over the entire process, the RMSE was averaged over time,

$$\langle RMSE \rangle = \frac{1}{t_f - t_0}\int_{t_0}^{t_f} RMSE(t)dt, \quad (12)$$

and evaluated numerically using the available sampling times.

A hybrid multi-start optimization strategy combining simulated annealing with local Gauss-Newton refinement was employed to ensure robust convergence and avoid trapping in local minima. Model performance was assessed by comparing simulated and experimental CLDs as well as derived quantities such as the weight-average molecular weight and the polydispersity index. The resulting parameter set reproduces both the detailed CLD shapes and the average CLD-derived values, considering experimental uncertainty caused by sample preparation, storage time, and possible impurities.

RESULTS

Figure 2 compares experimental and simulated chain length distributions (CLDs) at 18 h and 36 h of the fed-batch cultivation. Using literature-based polymerization parameters, the model underestimates both, the shift of the distribution towards higher molecular weights and the overall distribution width. In particular, the high-molecular-weight region is insufficiently captured, indicating an imbalance between propagation, termination, and degradation kinetics. By contrast, the optimized parameter set reproduces the experimentally observed CLD shapes with good agreement at both sampling times. This includes not only the position of the main peak but also the emergence of a pronounced high-molecular-weight tail at 36 h, demonstrating that the temporal evolution of polymerization dynamics is captured more accurately.

The improved agreement highlights the sensitivity of CLDs to polymerization rate constants and underlines that parameter values derived from PHB systems cannot be transferred directly to PHBV without adjustment. The copolymerization of hydroxybutyrate (HB) and hydroxyvalerate (HV), combined with the influence of propionic acid on monomer availability, leads to altered growth and termination behavior that must be reflected at the polymerization level.

Figure 3 shows the corresponding time courses of the weight-average molecular weight (M_w) and the polydispersity index (PDI). The optimized parameter set reproduces both the magnitude and temporal evolution of M_w significantly better than the literature-based parameters. In particular, the increase in M_w during the later stages of cultivation is captured consistently, reflecting the accumulation of longer polymer chains observed experimentally. Experimental PDI values exhibit comparatively large scatter, which can be attributed to uncertainties associated with SEC measurements, sample preparation, storage effects, and potential impurities. Nevertheless, the simulated PDI remains within a physically plausible range and follows the experimentally observed trend.

To quantitatively assess the improvement achieved by parameter optimization, the deviation between simulated and experimental CLDs was evaluated using the root mean square error (RMSE). Figure 4 compares the time-resolved RMSE for the literature-based (BASIC) and optimized (OPT) parameter sets. While both simulations capture the overall temporal trend, the optimized parameter set consistently exhibits lower RMSE values over the course of the cultivation. The time-averaged RMSE is

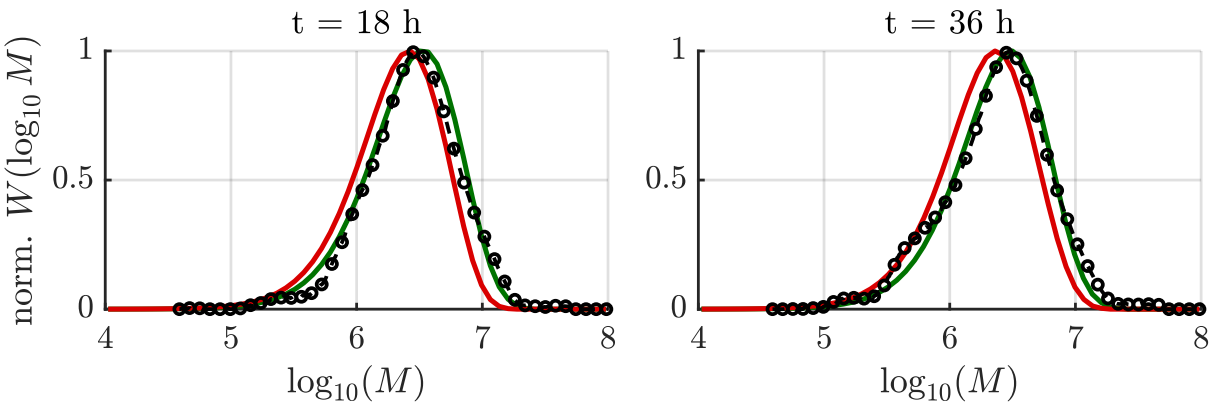

Figure 2: Comparison of experimental and simulated chain length distributions at 18 h and 36 h. Experiment: black, dashed with markers; basic model: red; optimized model: dark green.

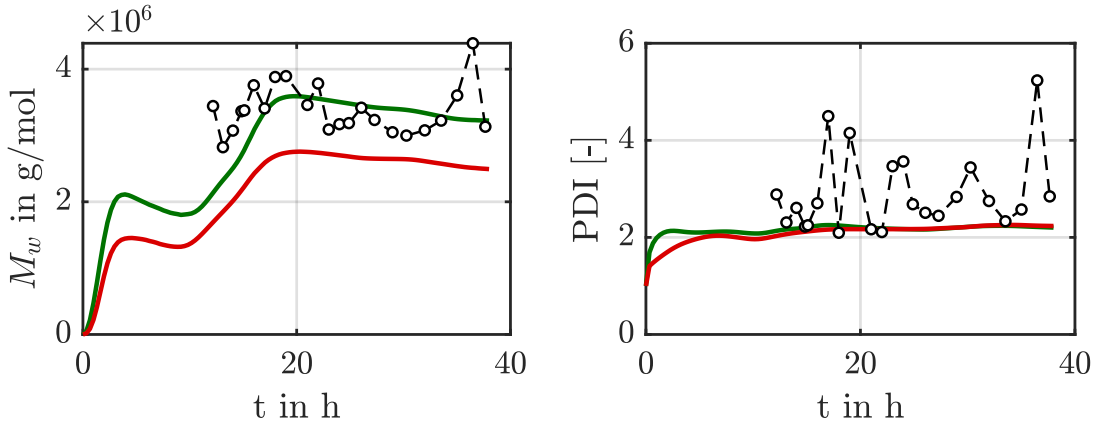

Figure 3: Time courses of M_W (left) and PDI (right). Experiment: black, dashed with markers; basic model: red; optimized model: dark green.

reduced from 0.088 for the BASIC parameter set to 0.058 for the optimized case, corresponding to a relative error reduction of approximately 34%. This quantitative improvement confirms that the enhanced agreement observed in the CLD comparisons is not only visual but persists over the entire process duration.

Table 1: Comparison between the original and updated parameter set

Parameter	Unit	Original value	Updated value
k_{ini}	[h^{-1}]	6.3 x 10^3	1.26 x 10^4
k_{m1}	[h^{-1}]	1.1 x 10^{-4}	5.16 x 10^{-5}
k_{m2}	[L h^{-1} mol^{-1}]	8.6 x 10^6	3.58 x 10^7
k_c	[h^{-1}]	1.40	8.26
k_{deg}	[h^{-1}]	83	6.85 x 10^2

The identified polymerization parameters are summarized in Table 1, comparing literature-based values [13] with the optimized parameter set. The changes in initiation, propagation, termination, and degradation rate constants reflect the altered kinetics required to reproduce PHBV copolymerization under propionic-acid-based feeding conditions.

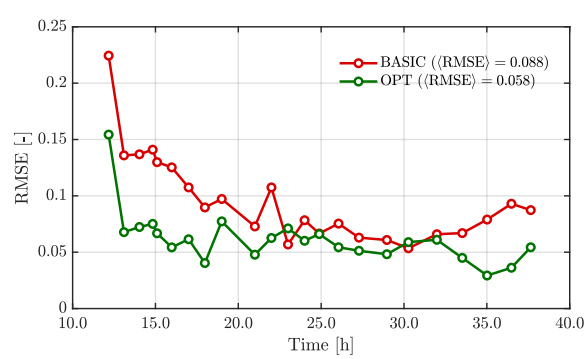

Figure 4: Time-resolved RMSE between simulated and experimental chain length distributions for the polymerization parameter sets. The legend indicates the time-averaged RMSE.

Overall, these results demonstrate that explicit parameter optimization at the polymerization level is essential for achieving quantitative agreement between simulated and experimental polymer microstructure. While macroscopic kinetics determine the overall monomer supply, the polymerization parameters govern how this supply is translated into chain length distributions and molecular averages. By reproducing both CLD shapes and derived quantities such as M_w, PDI, and RMSE, the proposed multiscale framework establishes a consistent and quantitative link between process conditions and polymer quality.

CONCLUSION AND OUTLOOK

This study builds upon a previously published multiscale PHBV modeling framework [9] and extends it towards quantitative prediction of both process dynamics and polymer microstructure by explicitly formulating the intracellular polymerization model and calibrating its kinetics. The proposed framework provides a mechanistic description of living and dead chain dynamics, monomer pools, and depolymerization. It also refines the polymerization rate constants by directly fitting simulated chain length distributions to experimental data. As a result, the predictive accuracy for the weight-average molecular weight and the overall CLD shape is substantially improved.

The established framework enables *in-silico* exploration of feeding strategies for tailoring PHBV quality and provides a consistent basis for investigating continuous and semi-continuous process scenarios. Such capabilities are particularly relevant for scale-up, where feed composition and timing can be systematically adjusted to steer polymer composition and molecular properties. Moreover, the modeling approach is not restricted to PHBV and can be extended to other copolymers as well as mixed-substrate fermentations, thereby broadening its applicability in industrial biotechnology.

Overall, this work advances the integration of bioprocess modeling and polymer science and supports quality-by-design approaches for sustainable biopolymer production. Future research should focus on open-loop model validation and the development of advanced control strategies for polymer quality. In this context, surrogate modeling approaches provide a promising pathway to reduce model complexity while preserving essential process dynamics, thereby enabling efficient model-based optimization and control [11]. These steps would further enhance the industrial relevance of the proposed framework and contribute to establishing biopolymers as competitive alternatives to conventional plastics.

ACKNOWLEDGEMENTS

This research received funding from the European Fund for Regional Development (EFRE) and the federal state of Saxony-Anhalt within the project PHA4VALUE (ZS/2023/12/182209), which is hereby gratefully acknowledged. The author would like to thank the University of Applied Sciences Magdeburg-Stendal, and all colleagues involved in the project.

AUTHOR IDENTIFIERS

Author ORCIDs:
Stefan Hempfling: 0000-0003-4759-0186
Rudolph Kok: 0009-0006-4061-106X
Stefanie Duvigneau: 0000-0002-7829-5190
Achim Kienle: 0000-0002-9011-0451
Robert Dürr: 0000-0001-5397-9302

REFERENCES

1. Sudesh K, Abe H, Doi Y. Synthesis, structure and properties of polyhydroxyalkanoates: biological polyesters. Progress in Polymer Science 25:1503-1555 (2000). https://doi.org/10.1016/s0079-6700(00)00035-6
2. Kim JH, Kim BG, Choi CY. Effect of propionic acid on poly (?-hydroxybutyric-co-?-hydroxyvaleric) acid production byalcaligenes eutrophus. Biotechnol Lett 14:903-906 (1992). https://doi.org/10.1007/bf01020626
3. Bugnicourt E, Cinelli P, Lazzeri A, Alvarez V. Polyhydroxyalkanoate (PHA): review of synthesis, characteristics, processing and potential applications in packaging. Express Polym. Lett. 8:791-808 (2014). https://doi.org/10.3144/expresspolymlett.2014.82
4. Koller M. A review on established and emerging fermentation schemes for microbial production of polyhydroxyalkanoate (PHA) biopolyesters. Fermentation 4:30 (2018). https://doi.org/10.3390/fermentation4020030
5. Guzman-Lagunes F, Wongsirichot P, Winterburn J, Guerrero Sanchez C, Montiel C. Polyhydroxyalkanoates production: A challenge for the plastic industry.
6. Guzmán-Lagunes F, Wongsirichot P, Winterburn J, Guerrero Sanchez C, Montiel C. Polyhydroxyalkanoates production: a challenge for the plastic industry. Ind. Eng. Chem. Res. 62:18133-18158 (2023). https://doi.org/10.1021/acs.iecr.2c04614
7. Troiano G, Nolan J, Parsons D, Van Geen Hoven C, Zale S. A quality by design approach to developing and manufacturing polymeric nanoparticle drug

products. AAPS J 18:1354-1365 (2016). https://doi.org/10.1208/s12248-016-9969-z
8. Duvigneau S, Dürr R, Behrens J, Kienle A. Advanced kinetic modeling of bio-co-polymer poly(3-hydroxybutyrate-co-3-hydroxyvalerate) production using fructose and propionate as carbon sources. Processes 9:1260 (2021). https://doi.org/10.3390/pr9081260
9. Duvigneau S, Dürr R, Wulkow M, Kienle A. Multiscale modeling of the microbial production of polyhydroxyalkanoates using two carbon sources. Computers & Chemical Engineering 160:107740 (2022). https://doi.org/10.1016/j.compchemeng.2022.107740
10. Duvigneau S, Wilisch-Neumann S, Dürr R, Kienle A. Modeling and experimental validation of poly(3-hydroxybutyrate-co-3-hydroxyvalerate) chain length distribution. *Comput Aided Chem Eng* 52:2221–2226 (2023) https://doi.org/10.1016/B978-0-443-15274-0.50354-1
11. Kok RL, Findeisen R, Kienle A, Duvigneau S. Improving the production of PHB in Cupriavidus necator by optimal control. *Comput Aided Chem Eng* 53:2503–2508 (2024) https://doi.org/10.1016/B978-0-443-28824-1.50418-X
12. Dürr R., Otto E., Kok R., Hempfling S., Duvigneau S., Kienle A., Bück A. Surrogate Modeling of Microbial PHA-Biopolymer Synthesis.
13. *Dürr R, Otto E, Kok R, Hempfling S, Duvigneau S, Kienle A, Bück A. Surrogate modeling of microbial pha-biopolymer synthesis. Ind. Eng. Chem. Res. 64:18640-18655 (2025). https://doi.org/10.1021/acs.iecr.5c01118*
14. Wulkow M. Computer aided modeling of polymer reaction engineering—the status of predici, i?simulation. Macro Reaction Engineering 2:461-494 (2008). https://doi.org/10.1002/mren.200800024
15. Penloglou G, Roussos A, Chatzidoukas C, Kiparissides C. A combined metabolic/polymerization kinetic model on the microbial production of poly(3-hydroxybutyrate). New Biotechnology 27:358-367 (2010). https://doi.org/10.1016/j.nbt.2010.02.001

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Evaluating Extrapolation of Modular Hybrid Process Models for Pilot-Scale Batch Separation Processes

Søren Villumsen[a], Jakob K. Huusom[a], Xiaodong Liang[a], Jens Abildskov[a*]

[a] Dept. of Chemical and Biochemical Engineering, Technical University of Denmark, Søltofts Plads 228a, 2800 Kgs. Lyngby, Denmark
* Corresponding Author: ja@kt.dtu.dk

ABSTRACT

Hybrid process models are increasingly used for real-time decision support in dynamic operations, where models must remain reliable under changing operating conditions. In such settings, models are often required to extrapolate beyond previously observed batch trajectories, yet conventional validation strategies may fail to reveal weaknesses in extrapolative behavior. This work investigates the effect of hybrid model structure on extrapolative performance using a pilot-scale batch crystallizer as a test case. A structured set of mechanistic and hybrid models is evaluated using nested batchwise leave-one-out cross-validation (NBLOOCV), in which entire batches are withheld to assess extrapolation across operating regimes. For this specific application, the results show that hybrid models employing simple linear correction terms consistently outperform more flexible neural network–based formulations under extrapolation, despite comparable training performance. For the studied process, the findings highlight the importance of validation strategies that reflect intended model use and demonstrate that modest, physically consistent hybridization provides a robust and interpretable basis for operational decision support.

Keywords: Hybrid modeling, Pilot-scale, Crystallization, Data-driven, Neural network, Separation Processes

1. INTRODUCTION

Process models are widely used in chemical and biochemical engineering to support monitoring, prediction, optimization, and operational decision making. Such models are particularly useful in equilibrium-based separation processes, as the evolution of interdependent states can be difficult to predict. In real-time batch separations, models are not used to reconstruct historical trajectories but to support time-sensitive decisions under uncertainty, such as when to intervene or adjust operational parameters. In this context, a model's practical value is determined not only by its average predictive accuracy, but by its stability and reliability when operating conditions deviate from those previously observed.

Mechanistic models provide a natural starting point for such applications, as they encode mass and energy balances and offer transparent extrapolation behavior. However, mechanistic descriptions of pilot- and industrial-scale equipment are often incomplete due to unmodeled effects, uncertain parameters, and simplifying assumptions [1]. Hybrid modeling approaches seek to address these limitations by augmenting mechanistic models with data-driven components to improve predictive performance.

Recent advances in machine learning have accelerated interest in neural network–based (NN) hybrid models for batch separation processes, because they can flexibly approximate complex discrepancies between models and processes [2]. However, some hybrid model structures can yield performance worse than baseline, necessitating screening procedures to compare them [3]. Additionally, such methods must be applied carefully, as NN-based models can overfit to training trajectories [4], and it is a known challenge to ensure good extrapolation [1, 2]. Poor screening and poor extrapolation performance can lead to model failure when models are used to predict future trajectories under operating conditions that differ from those previously observed.

Whether such failure modes are revealed depends on the validation strategy used during model development. Conventional cross-validation approaches that use validation data to guide hyperparameter selection are well-suited when training and deployment data are

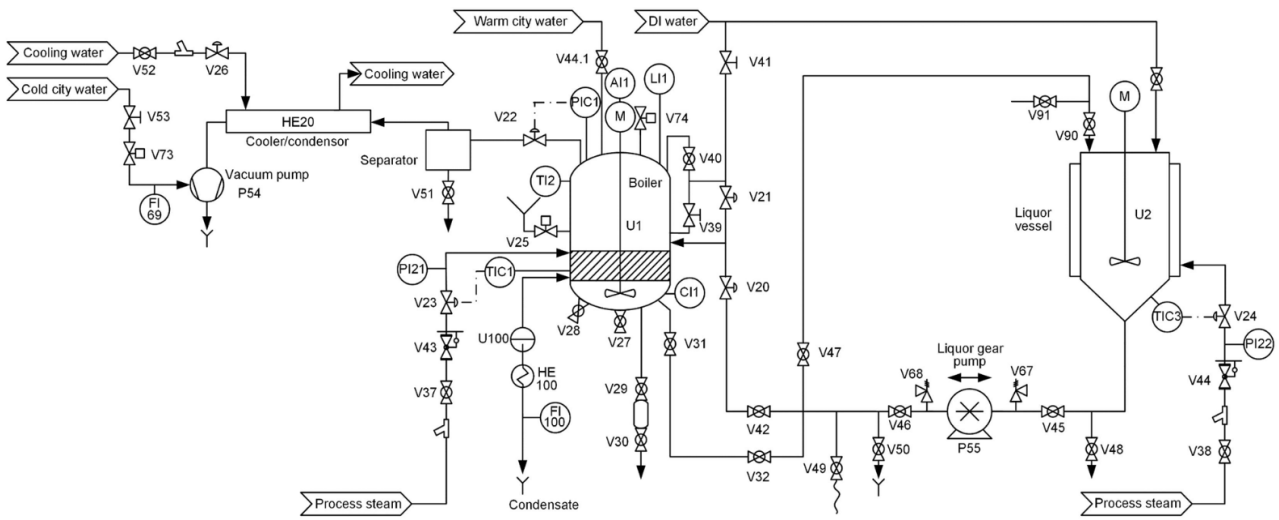

Figure 1. Process and Instrumentation diagram for the pilot-scale batch crystallizer.

drawn from similar distributions. However, when models are intended for extrapolation, which is common in low-data settings, these strategies may overestimate robustness [5]. Reporting validation performance on data that has influenced model selection leads to optimistic bias, an issue highlighted by Varma and Simon (2006), who showed that nested cross-validation is required to obtain unbiased generalization estimates [5].

In batch process modelling, segmenting a batch for both training and validation may yield a similar effect. To address this, this work employs nested batchwise leave-one-out cross-validation (NBLOOCV), in which entire batches are withheld as independent test cases. This extends the nested cross-validation framework of Varma and Simon to batch processes, ensuring that performance metrics reflect a model's generalization ability.

The objective of this study is to use a pilot-scale evaporative crystallizer as a case study to examine how different hybrid model structures behave under extrapolation. Pilot-scale batch crystallization provides a controlled yet industrially relevant setting in which operator interventions, unmeasured disturbances, and limited batch data jointly challenge model robustness. The work investigates how linear, nonlinear, and neural network–based correction terms interact with a mechanistic model when evaluated using NBLOOCV, and how these interactions affect trajectory prediction and downstream estimates of crystallization onset.

By focusing on model behavior under extrapolation rather than interpolation performance, this study aims to clarify when increased model flexibility improves predictive utility—and when it instead introduces instability that undermines operational trust.

2. METHODOLOGY

At the Technical University of Denmark, a batch crystallizer is routinely used to crystallize citric acid in water for educational purposes. A process and instrumentation diagram of this unit is shown in Figure 1. The crystallization tank is equipped with a temperature, weight, pressure, and conductivity sensor. The steam supply is equipped with a pressure sensor.

For operation, an initial mass of 150-200 kg of approximately 50% w/w citric acid solution is added to the crystallizer tank. A vacuum is pulled, reducing vessel pressure to ~0.15 bar; steam at ~1-2 bar gauge is supplied; and agitation is engaged to maximize evaporation rate. These must be continuously adjusted by the operator due to foaming and overflow risks. Every 20 minutes, samples are taken from the V28 valve, and titration is used to determine the citric acid concentration. Evaporation is continued until crystallization occurs, indicated by a sharp decrease in conductivity. Agitation is stopped, and crystallization is allowed to proceed, whereafter the separation is considered complete.

This batch crystallization process is difficult to describe using simple mechanistic models alone. Evaporation rates and heat transfer are affected by unmeasured disturbances, operator interventions, and equipment-specific effects, motivating the use of hybrid modelling. The process exhibits substantial batch-to-batch variability due to differences in initial conditions, steam and vacuum pressure profiles, and operator actions, resulting in non-identical trajectories that closely resemble variability observed in industrial batch operations. This setting is therefore deliberately chosen to reflect conditions under which hybrid models are most likely to be deployed in practice, but also most vulnerable to failure if validation is insufficient.

Data from seven batches of this operation were collected. The batch operation was segmented into the post-start-up operational phase. Using results from at-line samples, the initial conditions were established as

Table 1: Overview of mechanistic and hybrid model structures

Model Group	Evaporation Rate Term	Heat Source Term
Mechanistic	$r_{\text{evap}} = k_{\text{evap}}\left(p^{(\text{sat})}(T) \cdot x_w - p_{\text{vessel}}\right)$	$Q_{\text{in}} = \text{UA}\,(T_{\text{steam}} - T)$
Linear Correction	$r_{\text{evap}} = f_{\text{lin}}(\text{T}, x_w) \cdot k_{\text{evap}}\left(p^{(\text{sat})}(T) \cdot x_w - p_{\text{vessel}}\right)$	$Q_{\text{in}} = g_{\text{lin}}(\text{T}, x_w) \cdot \text{UA}\,(T_{\text{steam}} - T)$
NN Correction	$r_{\text{evap}} = f_{\text{NN}}(\text{T}, x_w) \cdot k_{\text{evap}}\left(p^{(\text{sat})}(T) \cdot x_w - p_{\text{vessel}}\right)$	$Q_{\text{in}} = g_{\text{NN}}(\text{T}, x_w) \cdot \text{UA}\,(T_{\text{steam}} - T)$
Structural	$r_{\text{evap}} = k_{\text{evap}}\left(f_{\text{poly}}(T) \cdot x_w - p_{\text{vessel}}\right)$	$Q_{\text{in}} = \text{UA}\left(g_{\text{poly}}(p_{\text{steam}}) - T\right)$
Full NN Term	$r_{\text{evap}} = f_{\text{NN}}(T, x_w, p_{\text{steam}}, p_{\text{vessel}})$	$Q_{\text{in}} = g_{\text{NN}}(T, x_w, p_{\text{steam}})$

the temperature, mass, and citric acid concentration at the time of the first sample. The time until crystallization is between 45 and 90 minutes at this point.

2.1 Mechanistic Modelling

A mechanistic model for the evaporation rate, r_{evap}, and temperature T, change can be expressed as,

$$\frac{dm}{dt} = r_{\text{evap}} = k_{\text{evap}}\left(p^{(\text{sat})}(T) \cdot x_w - p_{\text{vessel}}\right), \tag{1}$$

$$\frac{dT}{dt} = \frac{Q_{\text{in}} - r_{\text{evap}} \cdot \Delta H_{\text{vap}}}{m \cdot C_p}, \tag{2}$$

where the evaporation rate depends on the difference between the liquid's partial pressure, determined by saturated vapor pressure $p^{(\text{sat})}(T)$, and the molar water fraction x_w, and the vessel pressure p_{vessel}, scaled by an empirical coefficient k_{evap}. The temperature change is given by the heat input Q_{in} minus the energy required for evaporation: $r_{\text{evap}} \cdot \Delta H_{\text{vap}}$. An Antoine relationship is used to determine the saturated vapor pressure, $\log \text{p}^{(\text{sat})}(T) = A + B/(T - C)$ where A, B, and C are experimentally determined for a given solution. The heat balance is fitted using a linear driving force $Q_{\text{in}} = \text{UA}\,(T_{\text{steam}} - T)$, where UA is the overall heat transfer coefficient and T_{steam} is the steam temperature, which may be determined from the measured steam pressure assuming saturated steam. Additionally, Eqn. 1 may be augmented, replacing x_w by water activity a_w, to more accurately describe the observed relationship of $p^{(\text{sat})}(T)$ and x_w in low water fraction solutions. $a_w = 1 - k_{a_w} \cdot (1 - x_w)^2$, where k_{a_w} is empirically found.

Determining Crystallization Onset

To determine when crystallization will occur, the expected concentration of citric acid (CA) is continuously calculated. The concentration of saturated anhydrous citric acid in water as a function of temperature in °C has been determined empirically by,

$$c_{\text{CA, sat}} = 0.246 \cdot T + 59.17, \tag{3}$$

for temperatures greater than 40 °C. During operation, crystallization is not observed at this saturated concentration but instead requires a degree of supersaturation, which is not known in advance and will be determined from the data. In each fold, the training batches will be used to fit an offset to Eqn. 3, yielding the supersaturation curve. This will be used to predict the time of crystallization for the held-out batch.

2.2 Introducing Data-Driven Components

The mechanistic model will be augmented with data-driven components to improve predictive performance. An overview of the different structures is presented in Table 1, and includes:

- Correction models: the mechanistic model is corrected using a learned multiplicative factor. Both linear and NN-based correction factors are tested.
- Feature learning: Antoine parameters and steam pressure-temperature relationship are learned.
- Full NN: the entire term is approximated by a NN.

Besides models shown in Table 1, additional mechanistic and corrective hybrid models were evaluated, including models with:

- A heat loss term: $Q_{\text{net}} = Q_{\text{in}} - \text{UA}_{\text{loss}}(T - 25°\text{C})$
- Antoine values for saturated citric acid solution [1]
- Water activity substitute in Eqn. 1.

These variants reflect increasing levels of physical fidelity and parameter uncertainty.

2.3 Nested Batchwise Leave-one-out Cross-Validation

To evaluate model performance under extrapolative conditions, nested batchwise leave-one-out cross-validation (NBLOOCV) is employed. A schematic overview of the procedure is shown in Figure 2. In NBLOOCV, one complete batch is withheld as the outer-loop test set, while all remaining batches are used for model training and inner-loop validation. Hyperparameter tuning and model selection are performed exclusively within the inner loop, ensuring that the test batch remains fully independent.

This validation structure reflects how hybrid models are intended to be used in real-time batch operation. In practice, a model structure and its associated hyperparameters are selected using historical batches, trained on all available data, and then deployed to predict the trajectory of a future batch. NBLOOCV mirrors this setting by treating each batch as an indivisible unit and

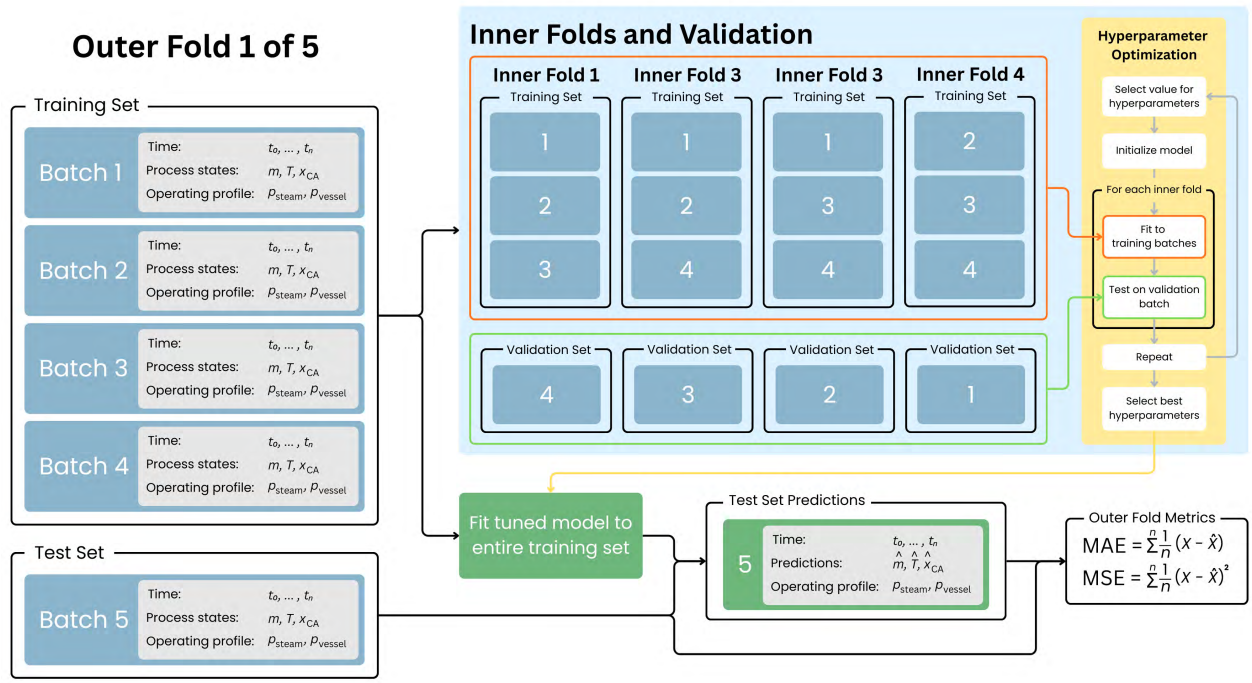

Figure 2. Schematic of the nested batchwise leave-one-out cross-validation (NBLOOCV) for a case of five total

evaluating performance on a completely unseen batch.

Within each outer fold, the remaining batches are further split using batchwise leave-one-out cross-validation in the inner loop. Here, one batch is used for validation, while the remaining batches are used for training; this procedure is repeated for all available training batches. The model structure and hyperparameters that yield the best average inner-loop performance are selected, retrained on all inner-loop batches, and evaluated on the withheld outer-loop test batch. This yields an unbiased estimate of performance on an independent batch.

Because batch operating conditions can differ substantially, evaluation on a single held-out batch may produce high-variance estimates of generalization error. NBLOOCV mitigates this by repeating the procedure such that each batch serves once as the independent test batch, after which performance metrics are aggregated across outer folds. As more batches become available, the framework can be extended to leave-*n*-out cross-validation to reduce computational cost while preserving test-set independence.

NBLOOCV is not required in all settings. When training and deployment data are drawn from similar operating regimes, conventional cross-validation provides meaningful estimates of generalization error. However, in low-data settings where future batches deviate from historical trajectories, validation performance can substantially overestimate robustness if the same data influence both model selection and evaluation. This mirrors pilot- and industrial-scale batch operations, where limited data and batch-to-batch variability make extrapolation unavoidable.

By withholding entire batches in the outer loop, NBLOOCV deliberately evaluates extrapolative predictions. As a result, performance metrics obtained using NBLOOCV reflect a model's ability to generalize across operating regimes, rather than its ability to interpolate within previously observed batch trajectories.

2.4 Model Training

Model training is performed using a staged optimization strategy that promotes numerical stability and efficient convergence. Initially, time-differenced residual minimization is performed on the mass and temperature gradients to obtain initial parameter estimates. Next, a Nelder–Mead optimization is performed using the full trajectory and initial conditions to refine the parameter values and the correction term parameters.

For NN-based models, the training is instead performed with gradient-based methods: Adam optimization is used for 150 epochs with a learning rate of $2.5 \cdot 10^{-3}$ and L-BFGS optimization for 25 epochs using a learning rate of 1.0. Unified NN models were trained that jointly predict both evaporation and heat input within a single network, such that the functions f and g in Table 1 are represented by one neural network. All model structures were evaluated both with a learned component applied to a single model term while retaining the baseline mechanistic term for the remaining term, and with learned components applied to both terms.

Hyperparameter search was performed using a

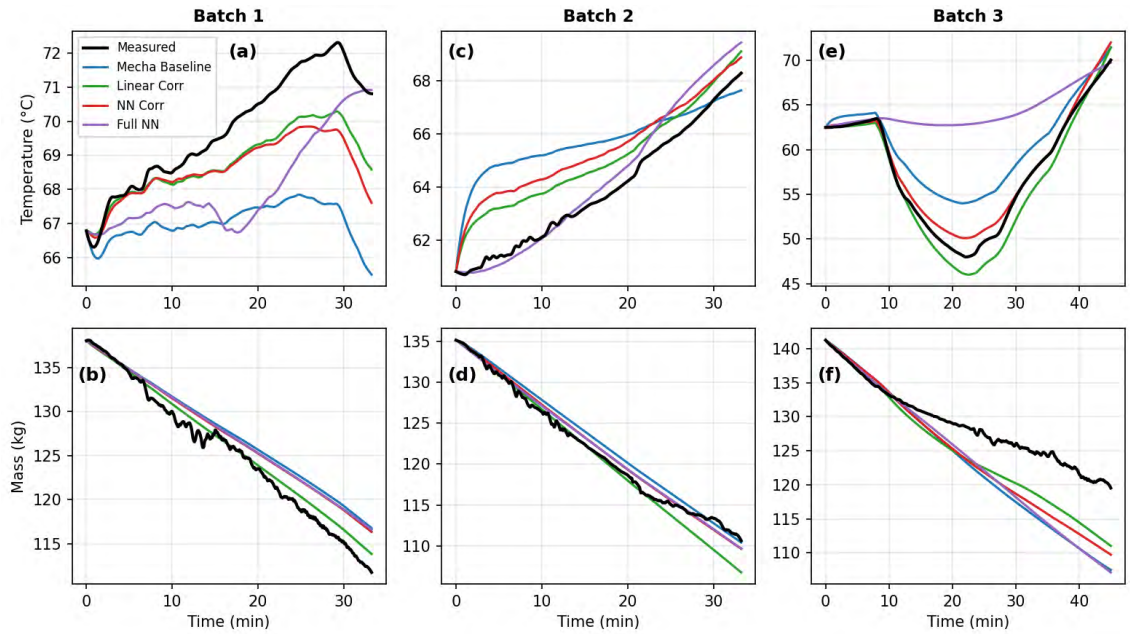

Figure 3. Example batches and the prediction of various hybrid models given the initial conditions.

Bayesian Tree-structured Parzen estimator optimizing the number of hidden layers: $\{1, 2\}$, the number of neurons per hidden layer: $\{2, 8\}$, and the weight decay: $\{10^{-5}, 10^{-2}\}$. The $\tanh$ activation function was used. The optimization metric for Nelder-Mead and gradient-based methods was $\sqrt{T_{\text{norm}}^2 + m_{\text{norm}}^2}$, where T_{norm} and m_{norm} are the normalized temperature and mass, respectively, based on the approximate means of batches, 70 °C and 120 kg.

3. RESULTS & DISCUSSION

Table 2 shows the predictive performance of representative models from each model class, evaluated using NBLOOCV. Training and outer-fold test losses are reported.

The computational cost of NBLOOCV varies substantially across model classes. For example, the linear correction models require on the order of 10 seconds per outer fold, while NN correction and full NN models require on the order of 10–30 minutes per outer fold. These costs are not directly comparable because the linear correction cost reflects a single training run, whereas NN costs include 10 iterations of inner-loop hyperparameter search. With 7 outer folds, the full NBLOOCV procedure for NN-based models takes on the order of 2–3 hours in total.

The mechanistic model using Antoine parameters for water serves as a simple baseline. The mechanistic model incorporating literature values for citric acid-water and water activity already performs worse, indicating that the Antoine parameters for pure water are more representative of the observed dynamics.

The linear correction models performed best in cross-validation across all explored model types. The NN correction models perform worse than the linear correction model but achieve lower average error than the baseline, albeit with higher variance. The full NN models perform worse again with even higher variance. The worst-performing cross-validation models are feature learning models, which perform notably worse than other model types.

In Figure 3, three of the seven batches and their trajectories predicted by select mechanistic and hybrid models are shown. The trajectories are not presented as evidence of strong absolute predictive accuracy, but to illustrate how model structure governs extrapolative behavior. Qualitatively similar trajectories are predicted by the mechanistic and corrected models for temperature, where the linearly corrected model is typically closest to the observed trajectories. The full NN-based models typically do not produce a qualitative match to the evolution of temperature. All models predict a constant decrease in mass, although some deviations are observed that none of the models successfully predict. The relevant distinction is therefore not which model performs best in an absolute sense, but that linear corrections fail in a physically consistent manner while NN-based models can deviate sharply and without physical justification.

3.1 Model Structure and Trajectory Fidelity

The strong performance of linearly corrected hybrid models suggests that the underlying mechanistic model already captures the dominant process dynamics, requiring only modest adjustments to account for unmodeled effects. This further implies that the discrepancy between the mechanistic model and the process exhibits a

Table 2: Averaged R^2 and losses for the select models of all structures. The outer fold test reports a standard deviation.

Model Type and Description	Tunable Parameters	Training		Outer Fold Test	
		R^2 $[T/m]$	loss $[\cdot 10^{-2}]$	R^2 $[T/m]$	loss $[\cdot 10^{-2}]$
Mechanistic baseline	2	$.17/.71$	2.77	$.13/.66$	3.32 ± 1.01
Mech. Lit. with a_w†	3	$< 0/.63$	4.82	$< 0/.61$	3.42 ± 1.09
Linear Correction	6+2*	$.64/.89$	1.85	$.78/.80$	1.88 ± 0.61
NN correction	8-52+2*, **	$.71/.89$	1.79	$.54/.77$	2.27 ± 1.96
Feature Learning	6+2*	$.43/.76$	2.33	$< 0/< 0$	13.4 ± 2.49
Full NN	12-52**	$.98/.97$	1.20	$.27/.73$	2.91 ± 2.77

†Uses the mechanistic model with water acitivty substituion and Antoine parameters for sautrated CA-water solution.
*Data-driven parameters + mechanistic parameters, **Span of parameters as a result of the per-fold hyperparameter running.

similar functional structure across operating regimes, resulting in comparable training and cross-validation performance.

In contrast, both the neural-network correction model and the fully data-driven NN achieve lower training error but substantially worse cross-validation performance. Because batches differ in operating conditions and trajectory evolution, held-out folds frequently require extrapolation beyond the training regime. For NN-based models, this leads to noticeable degradation in certain batches, lowering outer-fold metrics and resulting in high fold-to-fold variance. This implies that when extrapolation is required, fully data-driven models lack structural constraints to guide their behavior, increasing the likelihood of large prediction errors.

More generally, nonlinear correction terms are not constrained to preserve the functional relationships learned during training when applied outside the data range. As a result, corrections that improve fit within the training regime may scale differently or alter their functional behavior under extrapolation, leading to inconsistent interactions with the mechanistic model and deviations from physically plausible trends.

A similar effect is observed in feature-learning model structures. Although polynomial regressions represent only a modest increase in complexity relative to linear models, their behavior remains highly dependent on the training data used to estimate coefficients. When evaluated outside the training regime, small fitting deviations can be amplified by nonlinearity. While these models are more constrained than neural networks, their embedding within nonlinear relations—such as exponential Antoine-type expressions—can still lead to large extrapolation error. Notably, the feature-learning model that kept the mechanistic evaporation term and learned only the heat input Q_{in} performed similarly to NN correction models, indicating improved stability when learned components are confined to well-conditioned substructures.

Taken together, these results show that NBLOOCV effectively exposes extrapolation-induced failure modes that remain hidden under conventional validation strategies. These failures are consistent with a low-data, high-variability batch setting, in which flexible models are insufficiently constrained during extrapolation.

3.2 Predicting Crystallization Onset

Having established how model structure influences extrapolative robustness, crystallization onset prediction is used here as a downstream test of trajectory reliability, rather than as a standalone modeling objective. Using the linearly corrected hybrid model, a mean supersaturation offset of 3.8% w/w beyond the saturation curve was estimated from training batches and applied to predict crystallization onset in held-out batches. For all but one batch, the onset was predicted within 10 minutes (approximately 16% relative error). One batch (Figures 3e and 3f) exhibited delayed crystallization at lower temperatures, suggesting that the assumed supersaturation offset is invalid at low temperatures.

The purely mechanistic model achieved comparable accuracy (RMSE 8.6 min versus 6.7 min for the linear hybrid), indicating that the dominant process dynamics are already captured mechanistically. The hybrid model's advantage lies primarily in more consistent trajectory prediction across batches. The results show that when trajectory prediction is robust across batches—as enforced by NBLOOCV—simple hybrid corrections can translate into modest but reliable improvements in event timing.

3.3 Practical Considerations for Hybrid Model Design

The appropriate design of a hybrid model is inherently tied to its intended use. In real-time batch-process applications, models should support time-sensitive operational decisions under uncertainty. In this context, reliability, stability across operating regimes, and predictable failure modes are often more important than minimizing average prediction error.

The results of this study illustrate how these considerations directly influence hybrid model design. Hybrid models employing simple, structurally consistent correction terms modify mechanistic behavior in a predictable manner and therefore exhibit stable extrapolative behavior across varying batch conditions. This stability is essential when models are used during operation rather than for retrospective analysis. In contrast, hybrid models

incorporating nonlinear correction terms or fully data-driven components exhibit regime-dependent behavior, as reflected in poorer, more variable performance across NBLOOCV outer folds. While such models may perform well under specific conditions, their responses outside the training regime are less constrained, increasing the risk of misleading predictions when operators rely on the model most.

Models that occasionally produce large, unexpected errors can undermine trust and limit adoption, even when average performance appears strong. In contrast, models whose behavior aligns with physical intuition are easier to validate and act upon. The linearly corrected hybrid models evaluated here exemplify this trade-off, yielding reliable improvements over the mechanistic baseline without introducing unstable extrapolation.

Finally, the remaining limits in predictive accuracy indicate that not all errors should be addressed by increasing model complexity. Batch-to-batch variations in evaporation behavior and variability in crystallization onset reflect unmeasured disturbances and intrinsic process variability that cannot be resolved from available data. In such cases, increasing model complexity risks degrading robustness. For real-time applications, these uncertainties may be addressed more naturally through adaptive techniques or by communicating predictions as ranges rather than as point estimates.

Although demonstrated on a single crystallizer, the findings are expected to generalize to other batch separation processes characterized by limited data and batch-to-batch variability, such as chromatographic separation. In these settings, the same tension between flexibility and extrapolative robustness arises, and the case for modest, physically consistent hybridization is similarly desired.

4. CONCLUSION

This study examined how the validation strategy influences the selection of hybrid model structures in a case-study pilot-scale crystallizer. Nested batchwise leave-one-out cross-validation (NBLOOCV) was used to test extrapolation across batch trajectories. Here, hybrid models with simple linear correction terms consistently outperformed neural network–based models, which exhibited unstable behavior despite strong training performance. These results demonstrate that increased model flexibility does not guarantee improved predictive power when extrapolating.

The findings indicate that the hybrid model structure should be selected based on the intended operational use, particularly whether extrapolation across batches is required. For real-time applications, models must behave predictably when conditions deviate from historical data. NBLOOCV provides a practical mechanism for identifying extrapolation-induced failure modes during hybrid model selection. In the case studied, simple, linear hybrid corrections provide a robust foundation for operational decision support.

ACKNOWLEDGEMENTS

The authors thank the DTU-KT Pilot Plant for collecting the data, and Novo Nordisk Foundation for grant NNF19SA0035474.

AUTHOR IDENTIFIERS

Villumsen S: 0009-0006-0865-0122
Huusom JK: 0000-0002-1367-7235
Liang X: 0000-0002-2007-546X
Abildskov J: 0000-0003-1187-8778

REFERENCES

1. Gargalo CL, Malanca AA, Aouichaoui ARN, Huusom JK, Gernaey KV. Navigating industry 4.0 and 5.0: the role of hybrid modelling in (bio)chemical engineering's digital transition. Front. Chem. Eng. 6: (2024). https://doi.org/10.3389/fceng.2024.1494244
2. Schweidtmann AM, Zhang D, von Stosch M. A review and perspective on hybrid modeling methodologies. Digital Chemical Engineering 10:100136 (2024). https://doi.org/10.1016/j.dche.2023.100136
3. Frandsen J, Santana VV, Jul?Rasmussen P, Nogueira IBR, Huusom JK, Gernaey KV, Abildskov J. A systematic screening of neural network?based hybrid models of adsorption in chromatography processes. AIChE Journal 71: (2025). https://doi.org/10.1002/aic.70045
4. Jul-Rasmussen P, Chakraborty A, Venkatasubramanian V, Liang X, Huusom JK. Hybrid AI modeling techniques for pilot scale bubble column aeration: a comparative study. Computers & Chemical Engineering 185:108655 (2024). https://doi.org/10.1016/j.compchemeng.2024.108655
5. Varma S, Simon R. Bias in error estimation when using cross-validation for model selection. BMC Bioinformatics 7: (2006). https://doi.org/10.1186/1471-2105-7-91
6. de Kruif CG, van Miltenburg JC, Sprenkels AJJ, Stevens G, de Graaf W, de Wit HGM. Thermodynamic properties of citric acid and the system citric acid-water. Thermochimica Acta 58:341-354 (1982). https://doi.org/10.1016/0040-6031(82)87109-8

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Exploiting Input-Space Separation in Kolmogorov–Arnold Networks to Prevent Catastrophic Forgetting in Industrial NIR Systems

Imam M. Iqbal*, Isabell Viedt, and Leon Urbas

Technische Universität Dresden, Process-to-Order Group, Dresden, Germany
* Corresponding Author: imam_mujahidin.iqbal@tu-dresden.de.

ABSTRACT

Near-infrared (NIR) sorting systems in waste sorting plants operate under multiple settings, creating distinct input–output relationships that challenge predictive modeling. Conventional neural networks, such as multilayer perceptron (MLP), often suffer from catastrophic forgetting under continual training, limiting reliability across settings. This study evaluates Kolmogorov–Arnold Networks (KAN) for continual regression modeling of multi-setting NIR systems. KAN assign nonlinear transformations to network edges using localized spline grids, enabling structural isolation between input regions. We introduce controlled input-space manipulations (shifting successive settings to adjacent or non-overlapping grid regions) and compare KAN performance with MLPs of comparable parameter count. We also examine single-input versus multi-input configurations to assess dimensionality effects. Results show that KANs with sufficient input-space separation maintain previously learned knowledge with perfect resistance to forgetting, whereas overlapping inputs induce interference. MLP forgetting depends on learning rate and training duration and cannot be fully avoided without additional methods. Single-input KANs achieve comparable accuracy to multi-input models in this system, suggesting limited benefit from additional inputs. These findings demonstrate that KAN's structural locality, combined with controlled input-space alignment, provides a practical and robust approach for continual learning in industrial NIR systems.

Keywords: Artificial Intelligence, Machine Learning, Industry 4.0, Modelling, Process Monitoring

INTRODUCTION

Near-infrared (NIR) sorting systems in waste sorting plants operate under multiple parameter settings to accommodate variations in material composition, throughput, and product quality. These setting changes induce distinct input–output relationships, making reliable output prediction across operating conditions essential for disturbance detection, quality monitoring, and advanced control applications.

Classical models for waste sorting equipment typically rely on linear or split-factor formulations with a small number of dominant inputs [1, 2]. While interpretable, their limited expressiveness restricts their ability to capture nonlinear and setting-dependent behavior. Data-driven models, including multilayer perceptron (MLP), have demonstrated improved predictive accuracy by learning nonlinear relationships directly from data [3].

In industrial practice, however, data from different operating settings often become available sequentially, while historical data may be inaccessible or impractical to reuse. Under such continual learning conditions, neural networks such as MLP are prone to catastrophic forgetting [4], where training on new settings degrades performance on previously learned ones [5]. This behavior limits their reliability for long-term deployment in adaptive industrial systems.

Kolmogorov–Arnold Networks (KAN) have recently been proposed as an alternative neural architecture that assigns nonlinear transformations to network edges rather than nodes. Owing to their localized, grid-based parameterization, KAN may reduce interference between learned representations when trained sequentially. Initial evidence from synthetic benchmarks suggests that KAN

can better preserve previously learned functional components under certain continual learning scenarios [6]. Whether these properties extend to real-world industrial regression tasks remains an open question.

Motivated by these observations, this work investigates continual regression modeling for multi-setting NIR sorting systems using KAN. To explicitly examine the role of input-space structure in catastrophic forgetting. In addition, the study evaluates whether KAN can maintain predictive performance under reduced-input configurations, reflecting practical constraints in industrial modeling.

The contributions of this paper are summarized by the following research questions:

RQ1: How does input-space separation affect catastrophic forgetting in Kolmogorov–Arnold Networks under continual training?

RQ2: To what extent does KAN retain predictive performance on previously learned settings compared to MLP under identical sequential training protocols?

RQ3: Can KAN achieve comparable or superior prediction accuracy using a reduced set of input features compared to multi-input configurations?

METHOD

Problem Definition and Data Description

A NIR sorting unit is an optical-based separation system that exploits differences in infrared spectral absorption to distinguish between material types. During operation, the NIR unit periodically measures the spectral absorption, position, and surface area of materials transported on a conveyor belt. Based on the configured operating setting, air jets located at the end of the belt are actuated to divert materials into different output bins. As a result, the NIR sorting unit also functions as a material surface flow meter. Measurements from two subsequent units provides inflow and outflow of upstream unit.

The modeling task considered in this work is to predict the outflow of the NIR unit for a given inflow under different operating settings. Each flow stream consists of 11 material types. To limit scope and ensure clarity, this study focuses on predicting the outflow of a single target material, noting that the proposed methodology is directly extensible to other material fractions.

Inflow and outflow data were collected from the NIR unit and downstream processing equipment in an European waste sorting plant over approximately ten days equivalent of continuous operation. During this period, the system operated under two distinct settings. Setting A was active for the majority of the operating time, while Setting B was active for approximately three days. Measurements were recorded at a sampling interval of one minute, resulting in datasets of size 10, 187 and 3, 771 for Settings A (Dataset A) and B (Dataset B), respectively. The box plot of each dataset is presented in Figure 1.

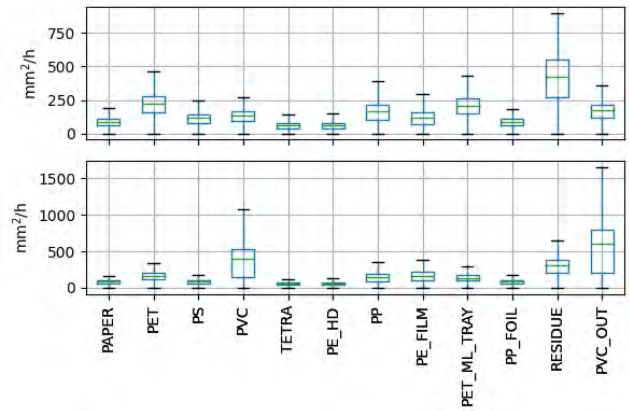

Figure 1. Flow Distribution in Dataset A (Top) and Dataset B (Bottom).

Continual Learning

To emulate a continual learning scenario, the model is trained sequentially on data from different operating settings. In the first phase (P_1), training is performed using Dataset A. In the second phase (P_2), the model is further trained on Dataset B without access to Dataset A. This setup reflects practical industrial constraints, where historical data may be unavailable when new operating conditions are introduced.

For each setting, the available data are randomly split into 70% for training, 15% for validation, and 15% for hold-out testing. The validation set is used exclusively for early stopping to mitigate overfitting with Patience of 20 steps and Max Steps of 500. Each training is repeated over five independent runs with different random seeds to account for stochastic variability in initialization and training. Model performance is evaluated on the hold-out test sets of all settings after each phase.

Prediction accuracy is quantified using coefficient of determination (R^2) for straightforward interpretation of predictive performance and forgetting behavior.

Kolmogorov Arnold Network

Kolmogorov Arnold Networks are a class of neural architectures inspired by the Kolmogorov-Arnold representation theorem, which states that any continuous multivariate function can be expressed as a finite sum of univariate functions composed with affine transformations. In contrast to MLP, where nonlinearity is applied at the nodes and edges are linear, KAN assign nonlinear transformations to the edges while nodes perform simple summation. See [6] for architectural illustrations.

Each edge in a KAN is parameterized by a learnable univariate function, typically represented using spline-based basis functions defined over a finite grid. These spline functions are locally supported, meaning that

parameter updates primarily affect specific regions of the input space determined by the grid structure and spline order k. As a result, changes induced by new training data may remain confined to localized subsets of parameters.

This parameter locality contrasts with the globally shared weights in MLP and provides a structural mechanism by which KAN may reduce interference between sequentially learned tasks. In continual learning settings, such locality suggests that learning new operating conditions could have a limited impact on previously learned input-output mappings, motivating the investigation of KAN for multi-setting industrial regression problems.

Input-Space Manipulation Strategies

To leverage KAN's grid-based parameter locality in resisting catastrophic forgetting, this study introduces deliberate input-space manipulation strategies. The objective is to control how data from different operating settings occupy the input domain and, consequently, how they align with the spline grids of the KAN model. These manipulations are introduced solely as an experimental mechanism inspired by synthetic continual learning benchmarks and are not assumed to arise naturally in industrial NIR data.

Specifically, multiple input-shifting configurations are examined. In all shifted cases, the input features of Dataset B are translated by a fixed offset relative to Dataset A. Shifts of 2000, 4000, and 6000 mm^2/h are evaluated for spline orders k=1 and k=2, corresponding to different degrees of separation between the input regions associated with successive training phases. A baseline configuration with no input shift is also included for reference.

For all configurations, the grid width is chosen such that each dataset predominantly occupies a single grid region (i.e., 2000 mm^2 per region). Grid knot positions are fixed, while the grid range and number of knots vary according to the applied input shift. This ensures that changes in forgetting behavior can be attributed to controlled input-space separation rather than adaptive grid refinement.

By comparing these configurations, the study systematically evaluates how increasing degrees of input-space separation affect catastrophic forgetting behavior in KAN under continual training.

Experimental Design

Effect of Input-Space Separation in KAN

The combinations of input shift magnitude and spline order k, together with their corresponding grid configurations, are summarized in Table 1. Across all cases, a consistent KAN architecture and training configuration is employed to isolate the effect of input-space separation. Specifically, all models use a single-layer KAN with width of [11, 1], a zero base-function, and fixed spline scaling and bias terms (i.e., spline scale and bias are not trainable).

Table 1: Input Space Manipulation Cases and The Corresponding KAN architecture.

Shift (mm^2/h)	0	2000	4000	6000
Case k=1	K1S0	K1S1	K1S2	K1S3
Case k=2	-	K2S1	K2S2	K2S3
Grid	1	2	3	4
Max Grid	2000	4000	6000	8000

KAN vs MLP under Continual Training

To address RQ2, MLPs are trained under the same continual learning protocol as KANs and evaluated for resistance to catastrophic forgetting. Because MLPs do not exhibit inherent parameter locality, multiple MLP models are trained using different learning rates during the second training phase (L_{P_2}) to examine the effect of learning rate on forgetting behavior and to identify the best achievable performance under this constraint.

To ensure a fair comparison of model capacity, MLP architectures with approximately matched numbers of trainable parameters (n_{PAR}) are used. In particular, an MLP with structure of [11, 4, 1] is evaluated and compared against the KAN configuration exhibiting the strongest forgetting resilience identified in RQ1 (denoted KAN-K1S2). All MLPs are trained using the original input space, as input shifting is not expected to induce parameter modularity in standard MLP architectures.

Reduced-Input vs Multi-Input Configurations

To address RQ3, the predictive performance of KANs under reduced input dimensionality is examined. Two reduced-input KAN models, using a single dominant input and spline orders k=1 and k=2 (denoted KAN-K1RI and KAN-K2RI), are trained under the same continual learning protocol. Their performance is compared against the best-performing multi-input KAN models for the corresponding spline orders to assess whether comparable accuracy and forgetting resilience can be achieved with fewer input features.

Forgetting Metrics

Forgetting is quantified by measuring the change in predictive performance on the initial operating setting after training on a new setting. Specifically, the coefficient of determination R^2 on the test set of Dataset A is evaluated before and after Phase 2 training. The difference between these values is used as the primary metric.

Models exhibiting a strongly negative R^2 on Dataset A after Phase 2 are classified as having undergone catastrophic forgetting and are labeled as "Forgetful". Models maintaining near-identical R^2 values before and after

Table 2: Input Space Manipulation Results.

Case	K1S0	K1S1	K1S2	K1S3	K2S1	K2S2	K2S3
S_{P_1}	29±4	28±3	27±4	27±4	87±18	77±23	110±50
S_{P_2}	56±15	45±19	60±22	64±26	68±8	66±20	50±18
R^2_{A,P_2}	0.76 ±0.00	-2.99 ±0.10	0.93 ±0.00	0.93 ±0.00	-456.05 ±14.49	0.72 ±0.06	0.93 ±0.00
R^2_{B,P_2}	0.96 ±0.00	0.96 ±0.00	0.96 ±0.00	0.96 ±0.00	0.97 ±0.00	0.97 ±0.00	0.97 ±0.00
ΔR^2_A	-0.17 ±0.00	Forgetful	Resilient	Resilient	Forgetful	-0.21 ±0.06	Resilient

Phase 2 are classified as "Resilient".

RESULTS AND DISCUSSIONS

Effect of Input-Space Separation

Table 2 summarizes the results of the input-space separation experiments. All configurations triggered early stopping in both Phase 1 and Phase 2, indicating stable convergence under the continual training protocol. Across all cases, predictive performance on the newly learned setting (Dataset B) remained high, with second-phase coefficients of determination (R^2_{B,P_2}) exceeding 96%. In contrast, performance retention on the previously learned setting (Dataset A) varied substantially across configurations.

Complete resilience to catastrophic forgetting is observed in cases K1S2, K1S3, and K2S3, where the coefficient of determination on Dataset A remains unchanged after Phase 2 training. In contrast, cases K1S1 and K2S1 exhibit complete forgetting, characterized by a severe degradation in R^2_{A,P_2}. The baseline cases K1S0 and K2S2 demonstrate partial resilience, retaining some predictive capability on Dataset A but with a noticeable reduction in performance.

These results indicate that catastrophic forgetting in KANs is strongly governed by the degree of overlap between input regions learned in successive training phases. When the input space of the new setting overlaps with, or is insufficiently separated from, the previously learned region, parameter updates propagate to shared spline basis functions, leading to interference and forgetting. Conversely, shifting the input space by at least k grid lengths from the boundary of the previously occupied region consistently prevents forgetting.

This behavior can be explained by the local support of spline basis functions in KAN. For a spline of order k, each spline basis function spans k+1 adjacent grid intervals. Consequently, basis functions located near grid boundaries influence neighboring regions outside the immediate input domain. If new training data falls within k grid intervals of a previously learned region, shared spline coefficients are updated, inducing forgetting. Ensuring a separation of at least k grid lengths guarantees that new data activates a disjoint set of spline basis functions, thereby preserving previously learned mappings.

These findings are consistent with observations reported by [7] and [8], who showed that architectures with localized or region-specific activations exhibit task-specific feature detectors. When the activation regions associated with different tasks do not overlap, earlier knowledge remains intact; when overlap occurs, forgetting increases proportionally to the extent of shared activation. In this context, KAN's spline-based edge functions act as localized feature detectors whose spatial separation directly governs forgetting behavior.

KAN vs MLP under Continual Training

Table 3 compares the forgetting behavior of MLPs under varying second-phase learning rates with that of a KAN model of comparable parameter count. For MLPs, reducing the second-phase learning rate resulted in smaller ΔR^2_A values, indicating slower forgetting; however, this was accompanied by a decline in second-phase test performance on data B (R^2_{B,P_2}). Early stopping was not triggered for any MLP configuration, implying continued parameter updates throughout the second training phase.

Despite learning-rate tuning, none of the MLP configurations fully avoided catastrophic forgetting. Instead, forgetting was mitigated only by slowing the optimization process, confirming that forgetting in MLPs is governed by optimization duration over shared parameters.

KAN achieved the highest R^2_{B,P_2} while maintaining a comparable R^2_{A,P_2} to the best-performing MLP, with substantially lower result variability. This indicates that KAN's improved forgetting resilience arises from its localized spline-based representations rather than from optimization hyperparameters.

It is acknowledged that several established techniques such as regularization-based methods, replay strategies, or parameter-isolation approaches can further improve forgetting resilience in MLPs. However, learning-rate control was deliberately selected here to

isolate the effect of optimization dynamics alone, without introducing additional mechanisms that explicitly modify the network structure or training data. This ensures a fair comparison focused on intrinsic architectural differences between MLP and KAN.

Table 3: KAN vs MLP in Continual Learning Results.

Case	K1S2	MLP1	MLP2	MLP3
L_{P_2}	1	1×10^{-4}	2×10^{-5}	1×10^{-6}
n_{PAR}	44	53	53	53
S_{P_1}	27±4	323±167	323±167	323±167
S_{P_2}	60±22	500±0	500±0	500±0
R^2_{A,P_2}	0.93 ±0.00	0.54 ±0.05	0.89 ±0.02	0.93 ±0.00
R^2_{B,P_2}	0.96 ±0.00	0.95 ±0.00	0.92 ±0.01	0.90 ±0.00
ΔR^2_A	Resilient	-0.38 ±0.05	-0.04 ±0.02	Resilient

Input Dimensionality Analysis

Table 4 summarizes the input dimensionality analysis. Across both k=1 and k=2, multi-input KAN models consistently achieve slightly higher R2 values than their single-input counterparts; however, the performance gap remains marginal.

This indicates that interactions between multiple material inputs contribute only weakly to predictive performance in the studied setting. Single-input models with k=2 exhibit a small but consistent improvement over k=1, suggesting the presence of mild nonlinear correlations in the input–output relationship.

Given the negligible accuracy gain from additional inputs and higher model complexity, a single-input KAN with k=1 emerges as the most effective configuration, delivering comparable predictive accuracy with substantially fewer trainable parameters and reduced model complexity.

Table 4: Input Dimensionality Analysis Result.

Case	K1S2	K2S3	K1RI	K2RI
n_{INPUT}	11	11	1	1
n_{PAR}	44	66	4	6
S_{P_1}	27±4	110±50	20±0	22±1
S_{P_2}	60±22	50±18	26±7	22±2
R^2_{A,P_2}	0.93 ±0.00	0.93 ±0.00	0.92 ±0.00	0.92 ±0.00
R^2_{B,P_2}	0.96 ±0.00	0.97 ±0.00	0.96 ±0.00	0.97 ±0.00
ΔR^2_A	Resilient	Resilient	Resilient	Resilient

Practical Implications for Industrial Modeling

From a practical perspective, the proposed input-space shifting strategy offers a simple yet effective mechanism for maintaining predictive accuracy across multiple operating settings without requiring access to historical data.

The results of this study suggest that, when using KANs, deliberate alignment of new operating data to previously unused regions of the input space can substantially mitigate catastrophic forgetting. This enables long-term model reuse across settings using a single model instance, reducing maintenance effort and simplifying deployment. Moreover, the strategy is purely on data-level and does not require architectural modification or specialized continual learning algorithms, making it attractive for industrial adoption.

Limitations

Several limitations of the proposed approach should be acknowledged i.e.:

- The input-space manipulation strategy operates exclusively at the input level and explicitly exploits the localized grid structure of shallow KAN. As network depth increases, interactions between layers may reduce the degree of effective parameter locality, potentially diminishing the strength of input-space separation.
- The strategy relies on deliberate and controlled input transformations, which may not always align with naturally occurring process variations. While suitable for controlled modeling scenarios, its applicability to systems with complex, high-dimensional, or poorly bounded input spaces requires further investigation.
- Finally, the present study focuses on a single real-world system and a limited number of operating settings. While sufficient for mechanism validation, broader conclusions regarding generalization across domains should be drawn cautiously.

Future Work

The following are several directions for future research emerge from this work:

- Extending the experimental setup to deep KAN would allow assessment of forgetting resilience in deep KAN.
- Evaluating the proposed mechanism on other architectures that employ localized representations, such as radial basis function models, would help determine whether the observed behavior is specific to KANs or reflects a broader architectural principle.
- Embedding shallow KAN as base layer of MLP and applying input-space separation might enhance the forgetting resilience of the resulting complex

model.

- Combining input-space manipulation with complementary continual learning techniques, such as replay buffers or regularization-based methods, may further enhance robustness while maintaining feasibility.

CONCLUSION

Shifting input regions by at least k grid lengths prevent catastrophic forgetting in KAN, while overlapping or insufficiently separated regions cause interference in spline coefficients. This demonstrates that KAN's local-support spline functions can be systematically exploited to retain prior knowledge in continual training.

KAN maintains high accuracy on both old and new settings, whereas MLP forgetting depends on learning rate and training duration and cannot be fully avoided without additional techniques. This shows that KAN's structural locality is key to its superior continual learning performance.

Multi-input KAN models offer only marginal gains over single-input models, indicating minimal effect from multi-material interactions. Single-input KAN with k=1 achieves comparable accuracy with fewer parameters, making it the most efficient choice for practical deployment.

Together, the results demonstrate that Kolmogorov–Arnold Networks provide an effective and robust framework for continual learning in multi-setting NIR systems, highlighting its practical suitability for industrial deployment where both adaptability and model reliability are critical.

ACKNOWLEDGEMENTS

As part of the Federal Government's seventh energy research program 'Research for an environmentally friendly, reliable and affordable energy supply', research funding is provided by the BMWK and Project Management Jülich (PtJ) under the title of project EnSort (FKZ 03EN2083B).

REFERENCES

1. Kleinhans K, Hallemans M, Huysveld S, Thomassen G, Ragaert K, Van Geem KM, Roosen M, Mys N, Dewulf J, De Meester S. Development and application of a predictive modelling approach for household packaging waste flows in sorting facilities. Waste Management 120:290-302 (2021). https://doi.org/10.1016/j.wasman.2020.11.056
2. Tanguay-Rioux F, Legros R, Spreutels L. On the limits of empirical partition coefficients for modeling material recovery facility unit operations in municipal solid waste management. Journal of Cleaner Production 293:126016 (2021). https://doi.org/10.1016/j.jclepro.2021.126016
3. Kroell N, Maghmoumi A, Dietl T, Chen X, Küppers B, Scherling T, Feil A, Greiff K. Towards digital twins of waste sorting plants: developing data-driven process models of industrial-scale sensor-based sorting units by combining machine learning with near-infrared-based process monitoring. Resources, Conservation and Recycling 200:107257 (2024). https://doi.org/10.1016/j.resconrec.2023.107257
4. McCloskey M, Cohen NJ. Catastrophic interference in connectionist networks: the sequential learning problem. Psychology of Learning and Motivation :109-165 (1989). https://doi.org/10.1016/s0079-7421(08)60536-8
5. van de Ven GM, Soures N, & Kudithipudi D. Continual Learning and Catastrophic Forgetting (Version 1). arXiv (2024). https://doi.org/10.48550/ARXIV.2403.05175
6. Liu Z, Wang Y, Vaidya S, Ruehle F, Halverson J, Soljačić M, Hou TY, & Tegmark M. KAN: Kolmogorov-Arnold Networks (Version 5). arXiv (2024). https://doi.org/10.48550/ARXIV.2404.19756
7. Lee A, Gomes HM, Zhang Y, Kleijn WB. Kolmogorov-arnold networks still catastrophically forget but differently from MLP. AAAI 39:18053-18061 (2025). https://doi.org/10.1609/aaai.v39i17.33986
8. Rahman MM, Wang G, Zhou K, Chen M, & Yang F. Catastrophic Forgetting in Kolmogorov-Arnold Networks (Version 1). arXiv (2025). https://doi.org/10.48550/ARXIV.2511.12828

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Renewables to X: Micro-Reactor Pathways towards Methanol and Dimethyl Ether Production

David T. Hren[a], Andreja Nemet[a]*

[a] University of Maribor, Faculty of Chemistry and Chemical Engineering, Maribor, Slovenia
* Corresponding Author: andreja.nemet@um.si

ABSTRACT

Renewable-to-X products such as methanol (MeOH) and dimethyl ether (DME) offer scalable, carbon-neutral options for decentralized chemical production. Microreactors, with superior heat and mass transfer, provide more controllable reaction environments. This improved control enhances selectivity and conversion, making microreactors particularly well suited for intensifying CO_2/CO hydrogenation within a Power to X framework for synthetic products. However, an assessment of MeOH and DME synthesis routes under microreactor operation is still lacking. To address this gap, a microreactor-scale model was developed where two reactor configurations were analyzed: i) parallel configuration, in which MeOH synthesis and subsequent dehydration to DME take place in the same reactor, and ii) series configuration, in which MeOH synthesis is carried out in the first reactor, followed by MeOH dehydration to DME in a second reactor. To capture realistic process behavior, the simulations incorporated non-isothermal, non-isobaric operation and fugacity-based reaction rates. The comparison reveals that reaction decoupling in the series configuration enables higher overall conversion of MeOH to DME, while the parallel route is constrained by strong kinetic and thermal coupling between hydrogenation, dehydration, and water formation. Despite requiring fewer unit operations, the parallel route yields mixed MeOH–DME product streams, whereas the series pathway favors higher DME productivity. The results highlight a fundamental trade-off between process intensification and conversion efficiency in microreactor-based Renewable-to-X systems.

Keywords: Methanol, Dimethyl Ether, Process Synthesis, MATLAB, Reaction, microreactor systems

INTRODUCTION

Renewable-to-X processes enable the conversion of CO_2- and CO-derived synthesis gas into liquid chemical energy carriers such as methanol (MeOH) and dimethyl ether (DME), which are relevant target products for decentralized Power-to-X systems [1]. DME synthesis can be produced through different process configurations that vary in reaction coupling, heat management, and handling of intermediate species, leading to distinct reactor and process performance characteristics. DME can be produced either in parallel, where MeOH synthesis and dehydration proceed concurrently from synthesis gas in a single reactor [2], or in series of reactors, in which MeOH is first formed and then dehydrated to DME in a downstream reactor stage [3].

Microreactor technology offers intensified heat and mass transfer, narrow residence time distributions, and improved temperature control [4], which are advantageous for kinetically coupled reaction networks such as CO_2/CO hydrogenation and MeOH dehydration. These characteristics enable systematic investigation of reaction pathway integration and separation under well-defined operating conditions. While microreactors have been widely applied for reaction intensification studies, comparative analyses of alternative DME synthesis routes under consistent microreactor-scale conditions remain limited.

This work addresses this gap by developing a unified microreactor modeling framework for two alternative DME production routes: (i) DME synthesis via parallel MeOH formation and dehydration reactions within a single microreactor, and (ii) DME synthesis via MeOH production followed by dehydration in a second

microreactor. Identical kinetic descriptions and operating conditions are employed to enable a consistent, side-by-side assessment of conversion, selectivity, and reactor performance, providing a quantitative basis for evaluating pathway selection in Renewable-to-X process design.

METHOD AND MODELS

Microreactor geometry and operating conditions

The investigated system was a tubular microreactor operated under steady-state conditions. Two reactor volumes were considered, namely 4 mL and 16 mL, corresponding to reactor lengths of 3.16 m and 12.63 m, respectively, at a constant inner diameter of 1.27 mm. The feed flow rate into the system varied between 0.5 and 10 mL/in, depending on the syringe pump configuration. The catalyst was modeled as a thin porous layer uniformly coated on the reactor wall with a thickness of 20 µm. Effective catalyst bulk densities were specified according to the reaction system under consideration: 440 kg/m^3 for MeOH synthesis [5], 693.96 kg/m^3 for MeOH dehydration to DME [2], and 1019 kg/m^3 for DME synthesis [6], based on reported experimental and apparent density data.

Flow model and species balances

Axial plug flow was assumed, neglecting radial gradients in concentration and temperature due to the small reactor diameter. Species balances were formulated following the approach of Nestler et al. [7], accounting for convective transport and heterogeneous catalytic reactions. For each species *i*, the molar flow rate variation along the axial coordinate *z* was given by Eq. 1.

$$\frac{d\dot{n}}{dz} = A_R \rho_{\text{cat}} \sum_j \nu_{i,j} \cdot r_j \tag{1}$$

In this Eq. 1 A_R is the reactor cross-sectional area, ρ_{cat} is the catalyst bulk density, $\nu_{i,j}$ is the stoichiometric coefficient, and r_j is the reaction rate of reaction *j*.

Reaction kinetics

Reaction kinetics were described using established literature models for MeOH synthesis, MeOH dehydration (DME production in series configuration), and DME synthesis in parallel configuration, enabling a consistent and mechanistically grounded comparison of alternative pathways. All kinetic expressions were implemented in terms of fugacities to account for non-ideal gas-phase behavior, as described by the Peng–Robinson equation of state [8].

Methanol synthesis in series configuration

MeOH synthesis was described by reaction network as presented in Eq. 2. It consisted of CO hydrogenation, reverse water–gas shift reaction and CO_2 hydrogenation.

$$\begin{gathered} \mathrm{CO} + 2H_2 \rightleftarrows \mathrm{CH_3OH} \\ \mathrm{CO_2} + 2H_2 \rightleftarrows \mathrm{CO} + H_2O \\ \mathrm{CO_2} + 3H_2 \rightleftarrows \mathrm{CH_3OH} + \mathrm{H_2}O \end{gathered} \tag{2}$$

MeOH synthesis from CO/CO_2 and H_2 was modeled using the kinetic formulation proposed by Park et al. [9], which builds upon Langmuir–Hinshelwood–Hougen–Watson expressions originally developed by Graaf et al. [10]. The model accounts for parallel hydrogenation of CO and CO_2 as well as the reverse water-gas shift reaction, incorporating competitive adsorption of reactants and products on $Cu/ZnO/Al_2O_3$ catalysts. Reaction rates are formulated as functions of species fugacities, ensuring thermodynamic consistency under elevated pressures. Temperature-dependent equilibrium constants and kinetic parameters were taken from Park et al. [9], while adsorption constants were adopted from Graaf et al. [10].

DME synthesis (methanol dehydration) in series configuration

In the series configuration, MeOH is dehydrated over an acidic catalyst, typically γ-Al_2O_3, according to Eq. 3.

$$2\mathrm{CH_3OH} \rightleftarrows \mathrm{CH_3OCH_3} + \mathrm{H_2O} \tag{3}$$

The kinetics of MeOH dehydration to DME were described using the model developed by Berčič and Levec [11]. Their original formulation expresses the reaction rate as a function of species concentrations, which is strictly valid for ideal reaction mixtures. Since non-ideal gas-phase behavior is considered in this work, the kinetic expression was reformulated in terms of partial pressures. At relatively low pressures, fugacities can be approximated by partial pressures, enabling a consistent transformation using the ideal gas relation $C_i = P_i/RT$ [12], [13]. The resulting pressure-based rate expression was implemented, while kinetic and adsorption constants were adopted from the original work [11]. The equilibrium constant for MeOH dehydration was evaluated using a temperature-dependent correlation from the literature [14].

DME synthesis in parallel configuration

In the parallel DME synthesis route, MeOH synthesis reactions (Eq.1-3) and MeOH dehydration (Eq.4) take place simultaneously within a single reactor over a bifunctional catalyst. This configuration provides a thermodynamic advantage compared to standalone MeOH synthesis, as continuous consumption of MeOH shifts the equilibrium toward higher overall conversion [15].

For the parallel route, the kinetic model proposed by Behloul et al. [6] was employed. The model combines MeOH formation and dehydration kinetics within a unified framework, accounting for reaction coupling and shared

intermediate species. No additional empirical coupling terms were introduced; instead, the interaction between hydrogenation and dehydration reactions emerges naturally through shared species, adsorption terms, and reaction enthalpies.

At each axial discretization step, local temperature, pressure, and composition were used to evaluate fugacities, kinetic rate constants, and adsorption terms. Reaction rates were then coupled to the species and energy balances, enabling fully coupled simulation of mass transfer, heat effects, and reaction kinetics along the reactor length. Identical kinetic formulations were applied across all microreactor configurations to ensure that observed differences in performance arise solely from pathway selection and reactor configuration.

Energy balance

An axial energy balance was formulated following the methodology of Nestler [7], extended to include external cooling. The energy balance accounts for reaction enthalpies, sensible heat effects, and heat exchange with the reactor surroundings. Reaction enthalpies were evaluated using Kirchhoff's law with temperature-dependent heat capacities [16]. At each axial discretization step, reaction rates, enthalpy changes, and temperature were updated consistently for all reaction networks.

Thermodynamic model

Non-ideal gas behavior was described using the Peng–Robinson equation of state [8]. For each component, temperature-dependent attractive and repulsive parameters were computed from critical properties and acentric factors. Mixture parameters were obtained using classical van der Waals mixing rules with binary interaction parameters. At each axial position, the cubic Peng–Robinson equation was solved to determine the mixture compressibility factor, selecting the largest real root corresponding to the gas phase. Fugacity coefficients were then calculated and used to evaluate component fugacities in the kinetic rate expressions.

Pressure drop model

Pressure drop along the microreactor was calculated using the model proposed by Hessel for microchannels [17]. Local losses were neglected, and the pressure drop was expressed as a function of friction factor, channel length, hydraulic diameter, and fluid velocity. The formulation was adapted to use dynamic viscosity instead of density to account for strong temperature dependence. Gas-phase viscosity was calculated using the Chapman–Enskog [18] method for low-pressure cases, while for the high-pressure cases the viscosity was calculated using the method developed by Chung et al. [19].

CASE STUDY

The simulation study consisted of a sensitivity analysis of two different pathways. For the MeOH and parallel DME configuration, the reactor temperature varied from 210 to 260 °C with 5 °C increments, and the pressure varied from 50 to 90 bar with 5 bar increments. The inlet volumetric flow rate varied between 0.5 and 10 mL/min with a step size of 0.5 mL/min. For the DME production in series reactor configuration, the temperature range was shifted to 290–360 °C in 5 °C increments, and the reactor pressure was fixed at 2.1 bar; the inlet flow rate corresponded to the MeOH outlet stream from the MeOH reactor simulation. The temperature and pressure ranges were chosen in accordance to the ranges at which the kinetic models were developed. Three different synthesis gas compositions were investigated, with H_2:CO_2:CO ratios of 0.8:0.2:0, 0.8:0.15:0.05, and 0.7:0.2:0.1. For the DME production in the series configuration, the feed composition was fixed at a MeOH:DME:H_2O ratio of 0.96:0.01:0.03. The feed compositions chosen were the most studied ones in literature. All simulations were conducted for reactor volumes of 4 mL and 16 mL to assess scale effects, reflecting the capacities of the available modules.

Series reactions configuration

The simulation results are presented for each process to examine the effects of operating conditions, reactor size, and process configuration on conversion and product formation. For the methanol reactor system the cases with the highest methanol output in terms of mmol/h and the highest methanol molar fraction in the outlet were discussed.

Methanol synthesis

For MeOH synthesis, the 16 mL reactor consistently outperformed the 4 mL configuration in both reactant conversion and MeOH production rate across the investigated operating conditions. The highest MeOH concentration at the outlet, 14.2 mol%, was achieved at 255 °C, 90 bar, 0.5 mL/min, with a H_2:CO_2:CO feed ratio of 0.7:0.2:0.1. The highest total MeOH output, 68.9 mmol/h, was achieved with a H_2:CO_2:CO feed ratio of 0.8:0.15:0.05 at 260 °C, 90bar, and 10 mL/min.

Although the feed ratio 0.7:0.2:0.1 delivered the highest outlet MeOH fraction (14.2 mol%) at low flow (0.5 mL/min), this primarily reflects long residence time and does not indicate production capacity. Therefore, the ratio 0.8:0.15:0.05, which yielded the maximum MeOH molar output (68.9 mmol/h) at high throughput (10 mL/min), was selected for further analysis as it represents the relevant design point from a productivity perspective. Figure 1 shows the effect of temperature (Figure 1a) and inlet flow rate (Figure 1b) on MeOH synthesis performance for the 16 mL reactor operated at 90 bar with a H_2:CO_2:CO feed ratio of 0.8:0.15:0.05.

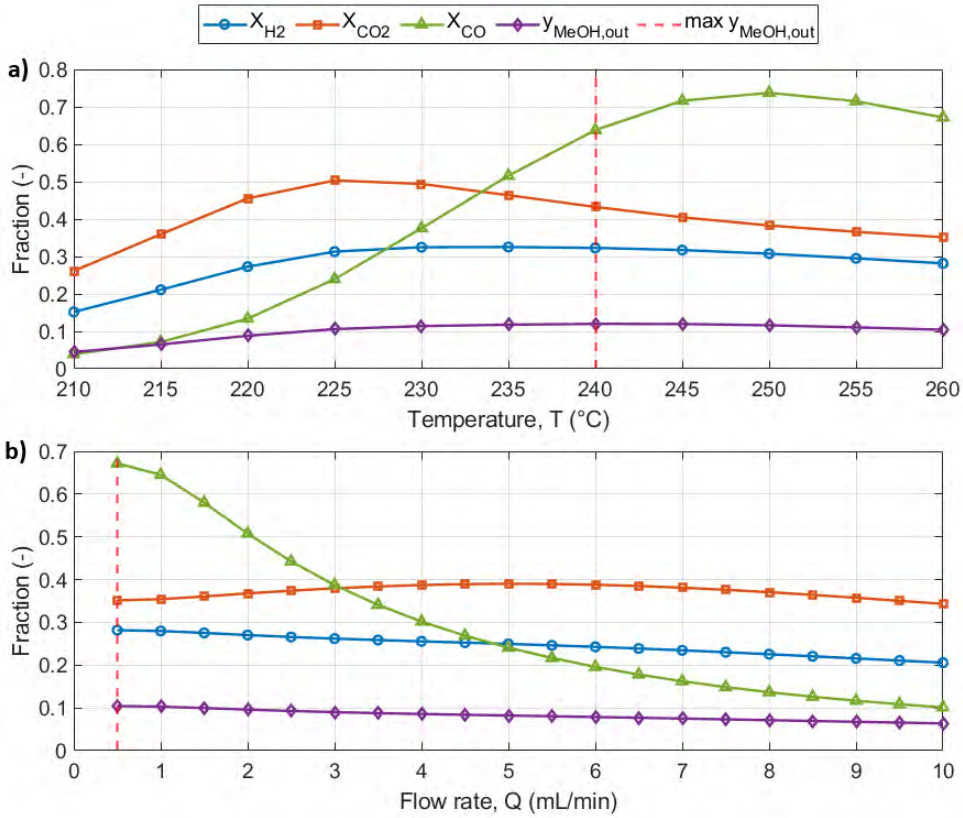

Figure 1. Effect of a) temperature, and b) inlet flow rate on reactant conversion and MeOH production for MeOH synthesis

For the temperature sensitivity analysis (Figure 1a), the inlet flow rate was fixed at 0.5 mL/min. Under these conditions, the maximum MeOH outlet fraction of 12.0% was obtained at 240 °C, resulting in a MeOH output of 6.1 mmol/h.

The flow rate sensitivity analysis (Figure 1b) was conducted at a fixed temperature of 260 °C, as the highest possible MeOH output was achieved at that temperature. The highest H_2 and CO conversions were observed at the lowest inlet flow rate of 0.5 mL/min. CO_2 conversion reached its maximum value of 39.0% at an intermediate flow rate of 5 mL/min. Although reactant conversions decreased with increasing inlet flow rate, the maximum MeOH production rate of 68.9 mmol/h was obtained at 10 mL/min, corresponding to CO_2 and CO conversions of 34.3% and 10.1%, respectively, and a MeOH outlet fraction of approximately 6.3 mol%. Although the MeOH outlet fraction under high-throughput operation (6.3 mol%) is substantially lower than the maxima observed at low flow rates (12.0–14.2 mol%), the corresponding MeOH production rate is more than an order of magnitude higher, increasing from approximately 6.1 mmol/h to 68.9 mmol/h.

Figure 2 shows the axial molar fraction values of species as well as the molar flow rate of the product for the MeOH synthesis reactor for the configurations that yield the highest molar MeOH flow rate (Figure 2a) and the configuration yielding a MeOH fraction of 12% (Figure 2b).

Figure 2a shows that in the case of maximum MeOH output, the reaction does not approach equilibrium, as indicated by the still relatively steep increase in MeOH flow and MeOH mol%. It can also be seen that CO mol% remains relatively unchanged from input to output, suggesting that most MeOH is produced from CO_2. Figure 2b shows a different trend: MeOH production begins to slow after 5 m. The mol% of CO starts to decrease after 4 m of reactor length, while the amount of CO_2 begins to increase. At 4 m of reactor length, the amount of water in the system remains approximately constant even though MeOH mol% continues to increase, indicating that the reverse water-gas shift reaction starts to consume water to produce hydrogen and CO_2, which also explains the increasing amount of CO_2 in the reactor.

DME synthesis

Based on the MeOH reactor performance at maximum MeOH output (68.9 mmol/h), the DME synthesis pathway is evaluated using the corresponding MeOH-rich outlet stream as feed. Figure 3 shows the effect of

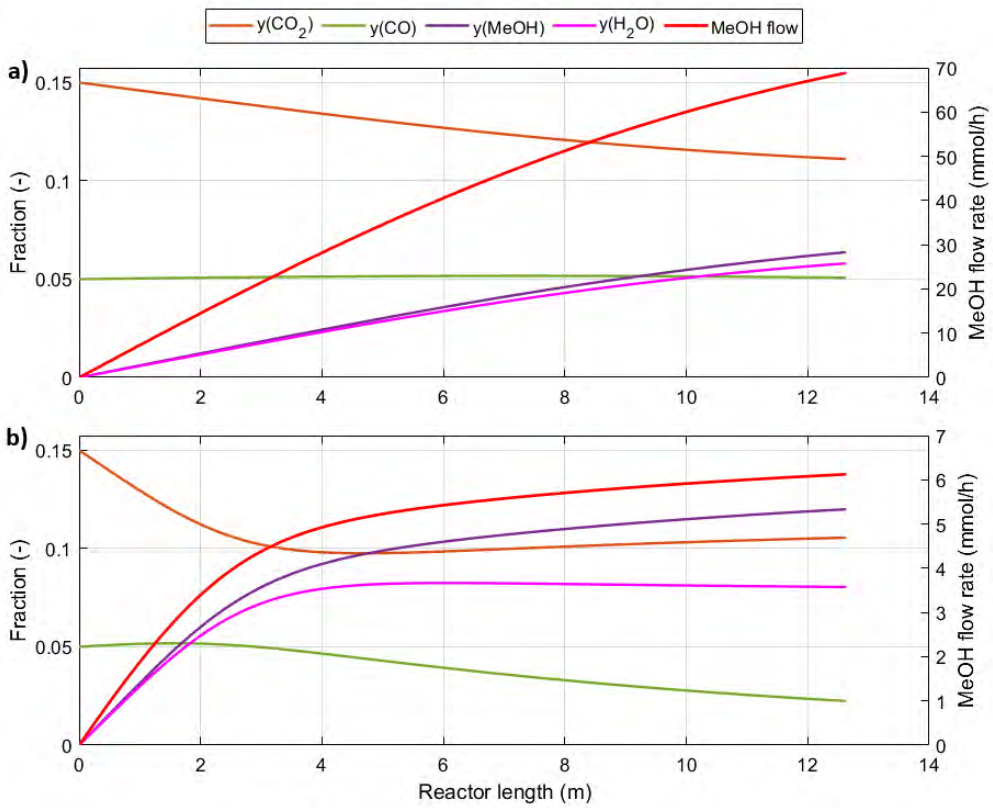

Figure 2. Species mol% and MeOH flow rate based on reactor length for a) the case for maximum MeOH output and b) the case for maximum mol% of MeOH in output

temperature on MeOH conversion and DME outlet fraction for DME synthesis pathway in series operated at 2.1 bar with reactor volumes of 4 mL and 16 mL.

For both reactor sizes, increasing the temperature from 290 to 360 °C resulted in a monotonic increase in MeOH conversion and DME formation, indicating that MeOH dehydration is strongly temperature-activated under the investigated conditions. In the 4 mL reactor, operation at 360 °C resulted in a MeOH conversion of 49.70%, corresponding to a DME production rate of 17.83 mmol/h. In the 16 mL reactor, operation at 360 °C led to a MeOH conversion of 87.37%, with a corresponding maximum DME production rate of 30.80 mmol/h. The substantially higher conversion and DME output achieved in the 16 mL reactor highlight the strong influence of reactor volume on DME production in series configuration performance.

Figure 4 presents axial species concentration profiles for this configuration to assess reactor utilization.

MeOH is rapidly consumed near the reactor inlet, with its molar fraction decreasing sharply in the initial portion of the reactor, indicating fast dehydration kinetics under these conditions. Correspondingly, the molar fractions of DME and water increase steeply at short reactor lengths and approach asymptotic values further downstream.

Beyond approximately the first 6 m of reactor length, the molar fractions of DME and water change only slightly, while MeOH approaches a low residual level. This behavior indicates that the dehydration reaction is mainly completed within the upstream section of the reactor, and that additional reactor length contributes minimally to further DME formation under the selected operating conditions.

Parallel reactions configuration

For DME synthesis in parallel configuration, as before, the 16 mL reactor consistently outperformed the 4 mL configuration in both reactant conversion and DME production rate across the investigated operating conditions. The best results were achieved at the highest investigated pressure of 90 bar using a H_2:CO_2:CO feed ratio of 0.7:0.2:0.1. Figure 5 shows the effect of temperature (Figure 5a) and inlet flow rate (Figure 5b) on DME synthesis performance for the 16 mL reactor operated at 90 bar with a H_2:CO_2:CO feed ratio of 0.7:0.2:0.1.

In the temperature sensitivity analysis (Figure 5a), the inlet flow rate was fixed at 0.5 mL/min. Under these conditions, the maximum H_2 conversion of 51.4% was obtained at 255 °C, along with the highest CO_2 conversion

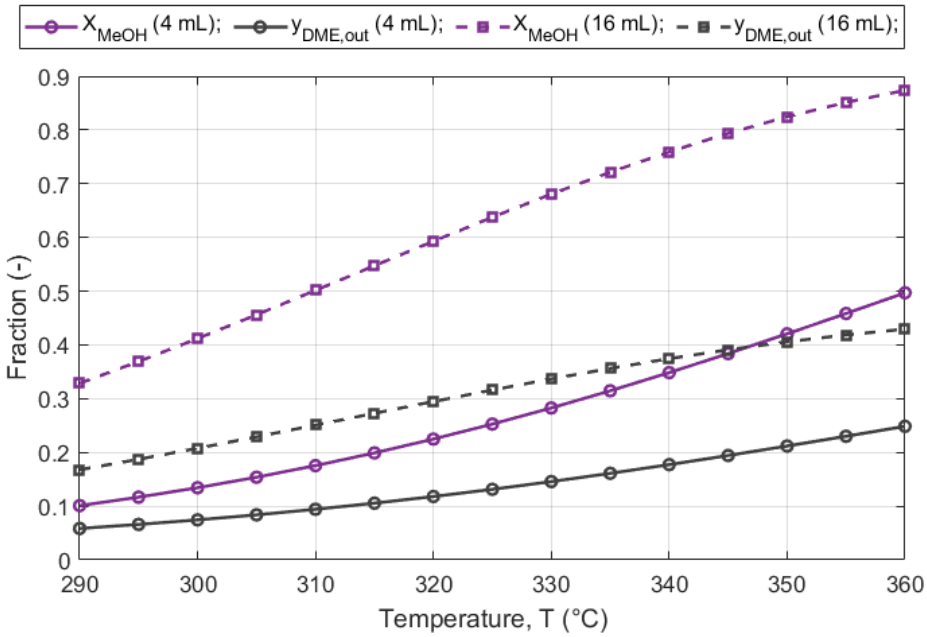

Figure 3. Effect of temperature on MeOH conversion and DME outlet fraction for DME synthesis in series configuration

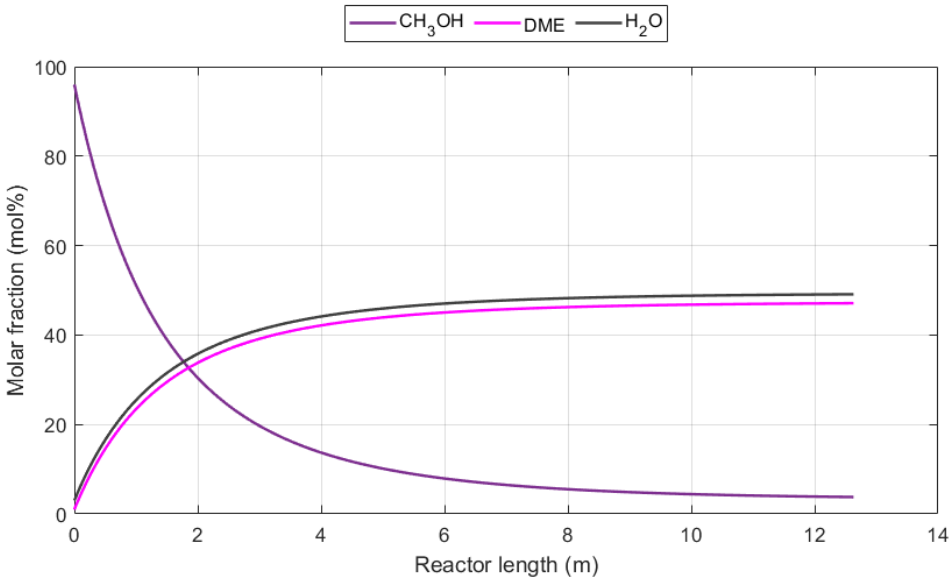

Figure 4. Axial molar fraction profiles of MeOH, DME, and water

of 29.2% and CO conversion of 92.3%. The DME outlet fraction peaked at 9.1 mol% (3.3 mol% MeOH) at the same temperature of 255 °C, reaching an output of 3.92 mmol/h.

The flow-rate sensitivity analysis shown in Figure 5b was conducted at a fixed temperature of 260 °C. The highest conversions and the highest DME fraction were achieved at the lowest inlet flow rate. The configuration with the maximum DME production rate (8.48 mmol/h) occurred at an inlet flow rate of 2.5 mL/min and temperature of 260 °C, but resulted in a substantially lower DME fraction of 3.4 mol% (4.3 mol% MeOH).

Figure 6 shows the axial molar fraction values of species and the molar flow rate of the product for the DME synthesis reactor for the configurations that yield the highest molar DME flow rate (Figure 6a) and the configuration yielding a DME fraction of 9.1 mol% (Figure 6b).

In Figure 6a, which corresponds to the highest DME

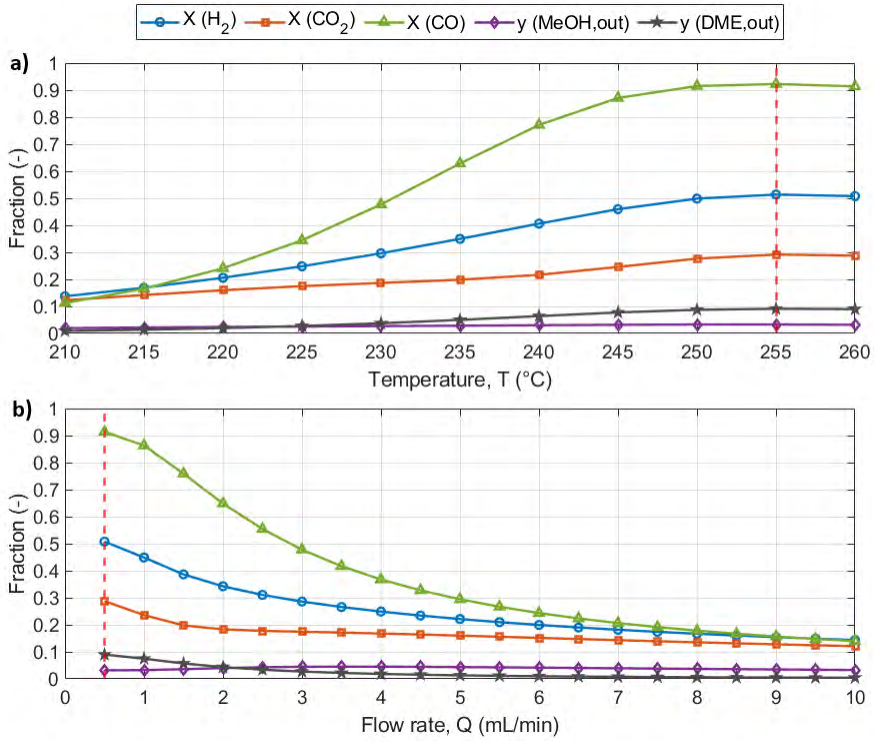

Figure 5. Effect of a) temperature, and b) inlet flow rate on reactant conversion and DME production

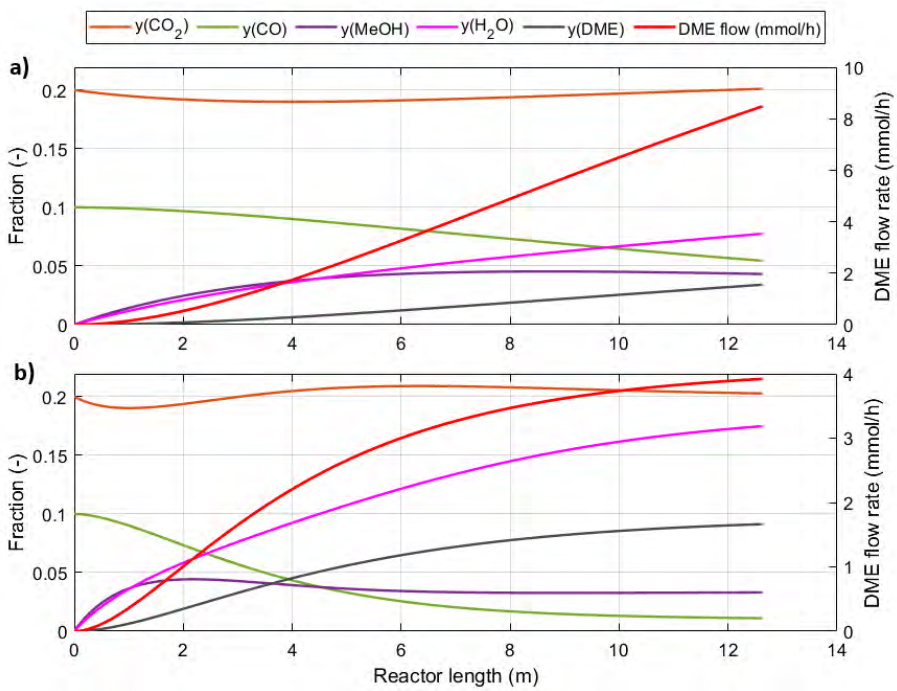

Figure 6. Species mol% and DME flow rate based on reactor length for a) the case for maximum DME output flow rate and b) the case for maximum mol% fraction of DME in output

outlet, it is shown that at the start (up to 2 m), CO_2 is being consumed to produce MeOH, with MeOH dehydration beginning after 2 m. After 2 m of reactor length, the CO mol% starts to decrease while CO_2 mol% starts to increase, indicating that CO is now being utilized to produce MeOH. Although MeOH is being consumed to produce DME, its value remains slightly above 4 mol%. At the reactor output, it can be seen that reactions are still

occurring, as indicated by the continued steep increase in water and DME mol%.

Figure 6b shows a different trend, as the DME mol% at the end of the reactor slowly stops increasing, which is also evident from the DME flow line. The MeOH mol% indicates that MeOH is predominantly produced up to 2 m of reactor length, after which it is consumed until reaching an equilibrium between production and consumption at around 8 m of reactor length.

Comparison of parallel and series DME production configuration

In the DME route with parallel configuration, where MeOH synthesis and dehydration occur simultaneously in a single reactor, the highest DME outlet fraction was achieved at low volumetric flow rate. At 255 °C, 90 bar, and 0.5 mL/min, the DME mole fraction in the outlet stream reached approximately 9.1 mol%, while the corresponding DME production rate (3.92 mmol/h) remained limited due to the low throughput. Increasing the inlet flow rate resulted in a higher amount of DME produced, with a maximum of approximately 8.5 mmol/h observed at 260 °C and 2.5 mL/min; however, this increase was accompanied by a substantial decrease in the DME fraction to 3.3 mol%.

In contrast, the series route, consisting of a MeOH synthesis reactor followed by an intermediate separation step and a downstream dehydration reactor operated at 2.1 bar, achieved significantly higher DME production. In the dehydration stage, MeOH conversion increased monotonically with temperature, reaching approximately 87% at 360 °C in the 16 mL reactor. Under these conditions, the maximum DME production rate reached approximately 31 mmol/h, about 3.5–4 times higher than the maximum value obtained in the parallel route for the same reactor volume.

The observed difference in performance can be attributed to reaction decoupling in the series configuration allowing the best operating conditions for both reactions. The dehydration reactor is more efficient when operating at elevated temperature and low pressure, conditions that strongly favor MeOH conversion to DME, while MeOH production reaction favor high pressure and peak temperature. In the parallel route, MeOH synthesis, dehydration, and water formation occur simultaneously within a single reactor, causing excess water produced from the MeOH synthesis reaction to further inhibit the MeOH dehydration reaction to DME.

CONCLUSIONS

A microreactor modeling framework was used to evaluate parallel and series production routes for MeOH and DME under identical reactor-scale assumptions. The series configuration consisting of a MeOH synthesis reactor followed by an intermediate separation step and a downstream dehydration reactor, enabled independent optimization of MeOH formation and MeOH dehydration. In the 16 mL reactor, this separation allowed high MeOH availability at the inlet of the dehydration stage, resulting in MeOH conversions of approximately 87% and a maximum DME production rate of about 31 mmol/h at elevated dehydration temperature and low pressure. In contrast, the parallel route integrates MeOH synthesis and dehydration within a single microreactor, resulting in the simultaneous presence of MeOH and DME in the outlet stream. While this configuration achieved a maximum DME production rate of approximately 8.5 mmol/h, the highest outlet DME mole fraction of about 9% was obtained only at low throughput, where MeOH conversion to DME was favored over absolute production rate. At higher flow rates, increased MeOH slip limited the extent of dehydration, leading to lower DME-to-MeOH ratios despite higher total throughput. Overall, while the parallel route offers a more compact process with a single reactor and no intermediate separation, the series reactors route demonstrates substantially higher DME productivity due to optimized and independent operation of each reaction step.

This study represents an initial step toward the optimal configuration and operation of dimethyl ether (DME) production in a microreactor environment. A kinetic model of the process was developed to a level suitable for validation against laboratory experiments. Once validated, the model can be embedded in a superstructure-based optimization framework (e.g., GAMS) to identify the optimal reactor and separation network topology, as well as the associated operating variables, for DME production. To enable mathematical programming, the governing dynamic model (ODE/DAE/PDE, depending on the chosen reactor representation) can be discretized into a finite-dimensional algebraic form using direct transcription methods such as orthogonal collocation (e.g., collocation on finite elements), resulting in a MINLP formulation.

ACKNOWLEDGEMENTS

The authors acknowledge financial support from the Slovenian Research Agency (project J2-60044, research core funding P2-0414 and PhD Research Fellowship No. 1000-24-0552).

AUTHOR IDENTIFIERS

Andreja Nemet ORCID: 0000-0002-2036-5053

REFERENCES

1. Dieterich V, Buttler A, Hanel A, Spliethoff H, Fendt

S. Power-to-liquidviasynthesis of methanol, DME or fischer–tropsch-fuels: a review. Energy Environ. Sci. 13:3207-3252 (2020). https://doi.org/10.1039/d0ee01187h
2. Fazlollahnejad M, Taghizadeh M, Eliassi A, Bakeri G. Experimental study and modeling of an adiabatic fixed-bed reactor for methanol dehydration to dimethyl ether. Chinese Journal of Chemical Engineering 17:630-634 (2009). https://doi.org/10.1016/s1004-9541(08)60255-4
3. Nasrollahi F, Bakeri G, Ismail AF, Rahimnejad M, Imanian M. Development of a model for dimethyl ether non-adiabatic reactors to improve methanol conversion. Korean J. Chem. Eng. 30:1867-1875 (2013). https://doi.org/10.1007/s11814-013-0138-0
4. Tanimu A, Jaenicke S, Alhooshani K. Heterogeneous catalysis in continuous flow microreactors: a review of methods and applications. Chemical Engineering Journal 327:792-821 (2017). https://doi.org/10.1016/j.cej.2017.06.161
5. Portha JF, Parkhomenko K, Kobl K, Roger AC, Arab S, Commenge JM, Falk L. Kinetics of methanol synthesis from carbon dioxide hydrogenation over copper–zinc oxide catalysts. Ind. Eng. Chem. Res. 56:13133-13145 (2017). https://doi.org/10.1021/acs.iecr.7b01323
6. Behloul CR, Commenge JM, Castel C. Simulation of reactors under different thermal regimes and study of the internal diffusional limitation in a fixed-bed reactor for the direct synthesis of dimethyl ether from a co₂-rich input mixture and h₂. Ind. Eng. Chem. Res. 60:1602-1623 (2021). https://doi.org/10.1021/acs.iecr.0c05535
7. Nestler F, Schütze AR, Ouda M, Hadrich MJ, Schaadt A, Bajohr S, Kolb T. Kinetic modelling of methanol synthesis over commercial catalysts: a critical assessment. Chemical Engineering Journal 394:124881 (2020). https://doi.org/10.1016/j.cej.2020.124881
8. Peng DY, Robinson DB. A new two-constant equation of state. Ind. Eng. Chem. Fund. 15:59-64 (2002). https://doi.org/10.1021/i160057a011
9. Park N, Park MJ, Lee YJ, Ha KS, Jun KW. Kinetic modeling of methanol synthesis over commercial catalysts based on three-site adsorption. Fuel Processing Technology 125:139-147 (2014). https://doi.org/10.1016/j.fuproc.2014.03.041
10. Graaf GH, Stamhuis EJ, Beenackers AACM. Kinetics of low-pressure methanol synthesis. Chemical Engineering Science 43:3185-3195 (1988). https://doi.org/10.1016/0009-2509(88)85127-3
11. Bercic G, Levec J. Catalytic dehydration of methanol to dimethyl ether. kinetic investigation and reactor simulation. Ind. Eng. Chem. Res. 32:2478-2484 (2002). https://doi.org/10.1021/ie00023a006
12. Prausnitz JM. Fugacities in high?pressure equilibria and in rate processes. AIChE Journal 5:3-9 (2004). https://doi.org/10.1002/aic.690050103
13. H. S. Fogler, Elements of chemical reaction engineering, Fifth edition. in Prentice-Hall international series in the physical and chemical engineering scien-ces. Boston Columbus Indianapolis: Prentice-Hall, Pear-son, 2016.
14. Lacher JR. The chemical thermodynamics of organic compounds (stull, daniel R.; westrum, edgar F.; sinke, gerard C.). J. Chem. Educ. 47:A300 (1970). https://doi.org/10.1021/ed047pa300.2
15. Aguayo AT, Ereña J, Mier D, Arandes JM, Olazar M, Bilbao J. Kinetic modeling of dimethyl ether synthesis in a single step on a cuo?zno?al₂o₃/?-al₂o₃ catalyst. Ind. Eng. Chem. Res. 46:5522-5530 (2007). https://doi.org/10.1021/ie070269s
16. P. W. Atkins and J. de Paula, Atkins' physical chemistry, 9th ed. Oxford: Oxford university press, 2010.
17. Hessel V, Renken A, Schouten JC, Yoshida J. Micro process engineering. Wiley (2013). https://doi.org/10.1002/9783527631445
18. S. Chapman and T. G. Cowling, The mathemati-cal theory of non-uniform gases: an account of the kine-tic theory of viscosity, thermal conduction, and diffusion in gases, 3rd ed. in Cambridge mathematical library. Cambridge ; New York: Cambridge University Press, 1990.
19. Chung TH, Ajlan M, Lee LL, Starling KE. Generalized multiparameter correlation for nonpolar and polar fluid transport properties. Ind. Eng. Chem. Res. 27:671-679 (2002). https://doi.org/10.1021/ie00076a024

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Dynamic modeling of fouling development during dead-end filtration of dusty superheated steam

Felipe de Oliveira[a*], Wijtze Nijhuis[a], Marcel Meinders[b], Edwin Zondervan[a]

[a] University of Twente, Department of Chemical Engineering, Enschede, Overijssel, the Netherlands
[b] Wageningen University & Research, Department of Food Technology, Wageningen, Gelderland, the Netherlands
* Corresponding Author: f.deoliveira@utwente.nl

ABSTRACT

This work proposes a parsimonious dynamic filtration model for superheated steam containing paper-derived dust, suitable for parameter identification, prediction, and future optimization under limited observability. The model is based on Darcy's law, with the total resistance expressed as the sum of the intrinsic filter resistance and a time-dependent fouling contribution. Experimental data obtained from a dedicated superheated-steam filtration setup were used for parameter estimation and model validation under a single operating condition. Assuming a linear dust dosing rate, the model yields limited agreement with experimental data (R^2 = 0.24). By estimating the time-varying solid loading through minimization of the sum of squared errors between measured and predicted pressure drop, the agreement improves significantly (R^2 = 0.94). This demonstrates that uncertainties in the dust dosing rate strongly affect pressure drop predictions. The proposed model provides a foundation for extending the analysis to broader process conditions, enabling statistical characterization of fitted parameters, integration with a filter cleaning model, and the development of optimization strategies for long-term, energy-efficient operation of closed-loop superheated steam drying systems.

Keywords: Gas filtration model, superheated steam filtration, cake filtration, fouling model, paper industry

INTRODUCTION

We live in an era where greenhouse gas emissions must be drastically reduced. Much of these emissions result from industrial energy use. In 2023, the European Union's top three energy-consuming industries were chemical and petrochemical, non-metallic, and paper, pulp, and printing [1]. In paper manufacturing, the drying step accounts for about 70% of total energy consumption, highlighting a clear opportunity for efficiency improvements [2]. Superheated steam drying (SHSD) offers a promising alternative to conventional air drying, potentially reducing energy demand by 70% and CO_2 emissions by 88% [2].

During SHSD, dust released from the paper contaminates the circulating steam, which can reduce process efficiency and cause operational issues in closed-loop systems. Effective control of these contaminants is therefore essential. Filtration showed potential as a separation method due to its high efficiency for fine particles, operational flexibility, and well-established technology. To support the industry in developing a long-term, energy-efficient filtration process, a filtration model should first be developed, which can later be coupled with a filter cleaning model to guide optimization strategies.

Filtration is a commonly used and extensively studied separation process. Numerous researchers have contributed to the modeling of filtration dynamics, with a focus on the filter performance with fouling development. One of the most recognized models in this field is the Hermia model, introduced in the 1980s [3]. This model describes fouling mechanisms associated with different pore-blocking behaviors and has gained popularity due to its simplicity and ease of application. The model was originally developed for constant-pressure filtration. Later, it was extended to constant-flux operation [4]. Under constant flux, the flow rate remains fixed, with the fouling manifesting as increases in resistance and, consequently, pressure drop. The pressure drop evolution can be expressed as:

$$\frac{d(\Delta P)}{dt} = K\Delta P^n \quad (1)$$

ΔP represents the transmembrane pressure drop over the filter, t is the filtration time, K is the blocking constant, which is typically determined by fitting the model to experimental data, and n is an exponent that defines the prevailing fouling mechanism (e.g., complete blocking, standard blocking, intermediate blocking, or cake filtration). Cake filtration is characterized by a linear increase in pressure drop. In contrast, intermediate, standard, and complete blocking exhibit nonlinear, concave-up pressure drop increases, with progressively steeper slopes from intermediate to standard blocking and the steepest increase observed for complete blocking [5].

Although widely used, the Hermia model has limitations. As it is fundamentally empirical, it does not capture the physical mechanisms.

In contrast to the Hermia model, other filtration models exhibit significantly greater complexity. An example is the model developed by Köry et al., which describes dead-end cake filtration using poroelasticity theory [6]. This approach accounts for the compressibility of both the filter medium and the growing cake layer. To capture these effects, dimensionless parameters are introduced that depend on specific material properties, such as elastic moduli and the permeability–strain relationship. As a result, the model requires detailed material characterization, which makes its application challenging in typical industrial settings due to the complexity and difficulty of obtaining such data.

Other models provide a balance between simplicity and physical insight, offering moderate complexity. One such model is that developed by Saleem et al., which extends Darcy's law to porous filtration systems [7]. The resulting expression includes two additive resistance terms: one representing the flow resistance of the filter medium, characterized by its permeability (B_m) and thickness (L_m), and the other representing the resistance of the deposited cake layer, characterized by its permeability (B_c) and thickness (L_c). Both terms scale with the fluid viscosity (μ) and filtration velocity (u).

$$\Delta P = \frac{1}{B_m}\mu L_m u + \frac{1}{B_c}\mu L_c u \tag{2}$$

While this model features a reasonable level of complexity and improves upon purely empirical formulations, it does not account for the wide range of operational conditions typically encountered in industrial processes. As a result, there remains a need for a moderate-complexity model that integrates physics-based principles with empirical flexibility, enabling practical implementation in industrial settings. Such a model should be based on parameters that are both physically meaningful and experimentally accessible, enabling accurate predictions under varying process conditions to support design and operational decision-making, particularly for superheated steam applications.

The novelty of this study lies in addressing this need for a superheated steam filtration model, where high temperatures and dynamic particle behavior introduce additional challenges.

The goal of the present work is to develop a parsimonious dynamic model for the dead-end filtration of superheated steam containing paper-like solid particles. The model integrates physically based formulations with empirical calibration to describe filtration behavior under the investigated operating condition. The framework can later be extended to a broader range of operating conditions, with the model refined using empirical correlations derived from statistical analysis of the fitted parameters, coupled with a filter-cleaning model, and ultimately used to support dynamic optimization and model-based control frameworks for industrial applications, aiming to reduce energy consumption and operational costs.

METHODOLOGY

Filtration model development

Darcy's equation describes the volumetric flow (Q, in m³/s) through a filter material over the effective filter area (A, in m). It states that the flow is proportional to the transmembrane pressure drop (ΔP, in Pa), which acts as the driving force, and is inversely proportional to the total filter resistance (R_t, in m⁻¹) and the fluid viscosity (μ, in Pa·s):

$$\frac{Q}{A} = \frac{\Delta P}{\mu R_t} \tag{3}$$

Due to the low penetration of the filter material, it is expected that the filtration is done at the surface of the filter, also known as cake filtration. For this reason, the total resistance is the sum of the filter resistance (R_f, in m⁻¹) and cake resistance (R_c, in m⁻¹):

$$R_t = R_c + R_f \tag{4}$$

The resistance of the cake can be written as the specific resistance (α, in m⁻²) multiplied by the cake thickness (l, in m):

$$R_c = \alpha\, l \tag{5}$$

The specific resistance is a material-dependent parameter that can, in principle, be estimated using the Carman–Kozeny equation. This equation considers the solid's specific surface area, the void fraction, and the Kozeny constant. However, given the objective of developing a model with parameters that are easily measurable for practical implementation, the specific resistance is determined from experiment.

The cake thickness is dependent on the deposited dust mass, the effective filter area, and the density of the dust cake. For a given mass, a larger filter area or a higher dust density results in a thinner cake. Assuming uniform

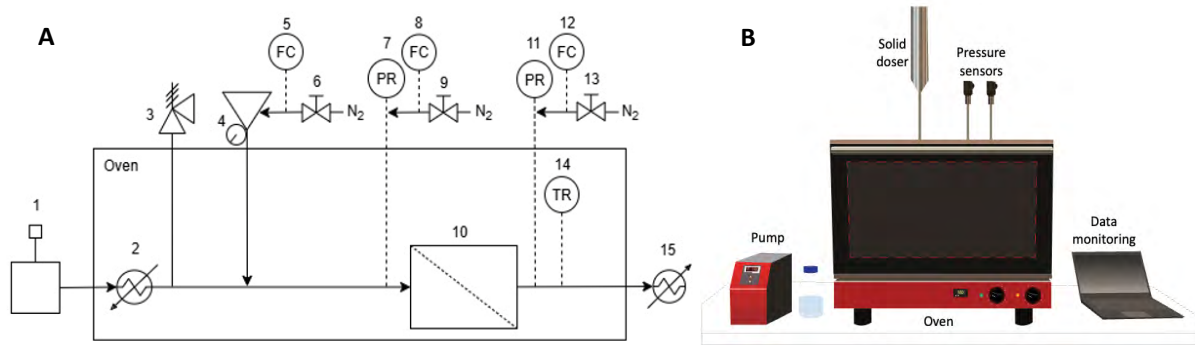

Figure 2. A- Piping and Instrumentation diagram of the set up where: 1 Pump; 2 Heating element; 3 Pressure relief; 4 Dust doser; 5, 8, 12 N2 flow controller (FC); 6,9,13 N2 gate valve; 7, 11 Pressure recorder (PR); 12 Dead-end filter in tri clamp; 14 Temperature recorder (TR); 15 Condenser. B- Overview of the main components of the system.

deposition, the cake thickness increases proportionally with the accumulated dust mass over time, reflecting the growth of the dust layer on the filter surface. The cake thickness as a function of time is given by:

$$l(t) = \frac{Q\,C_f\,t}{A\,\rho_c} = \frac{\dot{M}\,t}{A\,\rho_c} \tag{6}$$

with C_f [kg/m³] is the foulant concentration in the fluid, t [s] is the filtration time, ρ_c is the cake density [kg/m³], and $\dot{M} = \frac{dM}{dt}$ is the mass flow into the filter [kg/s].

By substituting equations 6, 5, and 4 in equation 3 and rewriting, the pressure drop is expressed as:

$$\Delta P\,(t) = \frac{Q}{A}\mu(\alpha\frac{M\,(t)}{A\,\rho_c} + R_f) \tag{7}$$

From Equation 7, the fitting parameters, filter resistance and specific cake resistance, were calibrated based on experimental data. The remaining variables were determined or measured, with further details provided in the next section.

A schematic is provided below to aid in visualizing the system. Dust-laden steam enters from the left-hand side and passes through the filter, where solid particles are retained, forming a cake layer. The filter acts as a resistive barrier, represented in the model by the filter resistance term. The total pressure drop across the system is the sum of the resistances contributed by the filter and the accumulated cake layer.

The main assumptions for the development of the filtration model are the following:

- Gas flow is well distributed;
- Uniform and incompressible cake and filter;
- No gradients within the filter;
- Cake filtration is the dominant mechanism;
- Solids do not influence the fluid viscosity;
- No condensation or leakages of steam;
- No chemical reactions;

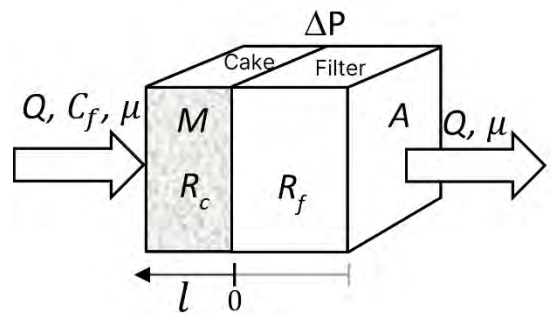

Figure 1. Diagram of the filtration mechanism.

Experimental investigation

Experimental studies were conducted to adjust the fitting parameters in the filtration model, namely the specific resistance and the filter resistance. The setup consisted of the production of superheated steam, dosing of solid particles, filtration of the dusty-laden steam, and the monitoring of the system, especially the temperature and the pressure drop across the filter. It was located inside an oven to avoid the condensation of the steam. Figure 2 shows the schematic of the filtration system.

The filter used was a Polytetrafluoroethylene (PTFE) filter, chosen for its ability to withstand high temperatures and its low dust adhesion, which ensures better long-term performance [8]. It had a surface area of 0.0028 m² and was mounted in the setup using a tri-clamp system, allowing for easy assembly and preventing leakage in the filtration section.

Nitrogen was used to prevent steam from entering the solid doser and to assist in transporting the solids at a flow rate of 3.5 l/min. The pressure sensors, connected to the system via capillaries and located outside the oven, were supplied with 0.05 l/min nitrogen per sensor

to prevent steam condensation, since they could not withstand the oven temperature.

The parameters in equation 7 were obtained through a combination of measurements, calculations, and calibration. Thermophysical properties of the gas phase were evaluated at local temperature and atmospheric pressure using the CoolProp library [9]. Temperature-dependent gas densities were used to determine the volumetric flow rates of nitrogen and steam. The dynamic viscosity of each pure gas was evaluated using Sutherland's law and combined into a mixture viscosity using a composition-dependent mixing rule. The effect of dust particles was not considered in the fluid properties.

The density of the filter cake was determined experimentally using a volumetric method, in which the dust was placed in a volumetric flask and gently compacted to reduce interparticle voids before weighing. This approach is based on the expected low void fraction of the cake resulting from the assumption that there is a uniform and distributed deposition of particles during filtration.

The pressure drop across the filter was measured directly using upstream and downstream sensors. The total mass was determined at the end of each run by weighing the clean filter and the filter after filtration. Cellulose fibers were used as the dust to mimic paper-like particulate matter. The cumulative deposited mass at any time, $M(t)$, and the total deposited mass, $M_{\max}$, were initially assumed based on a linear dosing profile. The actual time-dependent dosing was then reconstructed using an inverse estimation approach by minimizing the sum of squared errors between the pressure drop predicted by the model and the experimental measurements.

The fitting parameters were determined in two steps. First, the filter resistance, R_f, was obtained under steam filtration conditions. Second, the specific cake resistance, α, was obtained during the dust filtration phase using the filter resistance previously obtained.

Model validation was performed by comparing measured pressure drop data with the output of the calibrated filtration model. Four independent filtration experiments were conducted under identical operating conditions: a process temperature of 160 °C, a water pump flow rate of 5 ml/min, and an average mass of solids retained in the filter of 7.5 g. Two runs were used exclusively to construct the experimental reference, obtained by averaging their pressure drop profiles. The remaining two runs were used for model calibration.

RESULTS AND DISCUSSION

Model validation

The filtration system was used to assess the performance of the calibrated model by comparing its predictions with the experimental data. Table 1 presents the process conditions and the parameters used for the pressure drop predictions.

Some degree of variability is inherent to the experimental setup and should be considered when assessing reproducibility. The fabric filter material is not perfectly identical between experiments, and small variations in its structure may occur. In addition, the permeate flux and gas viscosity depend on the system temperature; although the temperature was maintained within a narrow range, minor fluctuations were observed. Small variations in the water pump operation may also lead to slight deviations in the steam flow rate. These factors can introduce limited experimental variability and are considered potential sources of deviation between repeated runs.

Using the average filter resistance and specific cake resistance, together with the remaining operational parameters calibrated from the experimental runs, the filtration model was used to predict the pressure drop, as shown in Figure 3. The solid addition was assumed to occur linearly over time, resulting in a linear pressure drop trend for the fitted model.

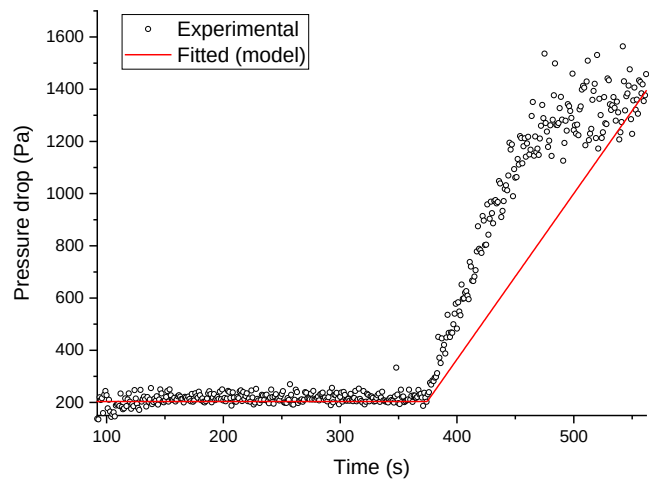

Figure 3. Comparison between experimental data and fitted model predictions assuming linear solid dosing.

From Figure 3, two main regions can be identified: a superheated steam filtration region (up to approximately 375 s) and a dust filtration region (from 375 s onward). During the steam filtration stage, the model predicts the pressure drop well in comparison with the experimental data. Deviations observed at the beginning and end of this region are attributed to experimental limitations, mainly the time required to reach equilibrium in the steam production. The coefficient of determination (R^2) was evaluated for the steam filtration region by comparing the squared residuals between experimental and modeled pressure drop values to the variance of the experimental data. In this regime, the experimental pressure drop remains nearly constant, as no fouling occurs. Because the variance of the experimental data is very low, even small deviations between the model and experimental result in

negative R^2 values, indicating that R^2 is not a suitable performance metric here. Instead, model performance was assessed using the root-mean-square error (RMSE), which yielded 24 Pa.

In the dust filtration region, the experimental data show a clear linear increase in pressure drop up to approximately 450 s. Beyond this point, the slope gradually decreases, and the pressure drop approaches a quasi-steady value. The fitted model predictions, which assume linear dust dosing, do not accurately capture the pressure drop increase during this transient phase. The R^2 in the dust filtration region is 0.25. However, the model successfully predicts the final pressure drop level. Experimental investigations indicate that the solid mass retained in the filter is measured only at the end of each run, while the dust dosing rate is difficult to control and not constant over time. This uncertainty represents a challenge for predicting the pressure drop evolution.

This behavior is also likely to be observed in industrial settings. Although operating conditions are controlled to ensure steady production, variations in foulant concentration are expected due to upstream fluctuations, changes in raw materials, or transient operating conditions. Similarly, the flow rate of the filtered steam is generally maintained around a target value but may exhibit small variations. Together, these variations can influence the pressure drop evolution across the filter.

Given the known process parameters (Table 1) and the experimental pressure drop profile, the dust dosing profile $M(t)$ was estimated by minimizing the sum of squared errors (SSE) between model predictions and experimental data:

$$SSE = \sum_{i=1}^{n} [\Delta P_{model}(t_i, M(t_i)) - \Delta P_{exp}(t_i)]^2 \qquad (8)$$

The dosing profile was obtained using the L-BFGS-B algorithm, subject to constraints: (i) non-decreasing dust mass, (ii) bounds $M(t)$ between zero and the total measured dust mass $M_{\max}$, and (iii) the final accumulated mass is the measured dust mass. The resulting profile represents the dust accumulation that best matches the filtration model with the measured pressure drop data. The resulting dust dosing profile is shown in Figure 4.

By using the adjusted dust dosing profile, a comparison between the experimental data and the fitted model predictions was performed, as shown in Figure 5A. A significantly improved accuracy is obtained, with a R^2 of 0.94, which is evident in the parity plot presented in Figure 5B.

The parity plot also highlights larger deviations at the extremities of the run, particularly around the transition points. The highest spread is observed in the region where the maximum dust mass ($M_{\max}$) is accumulated in the filter.

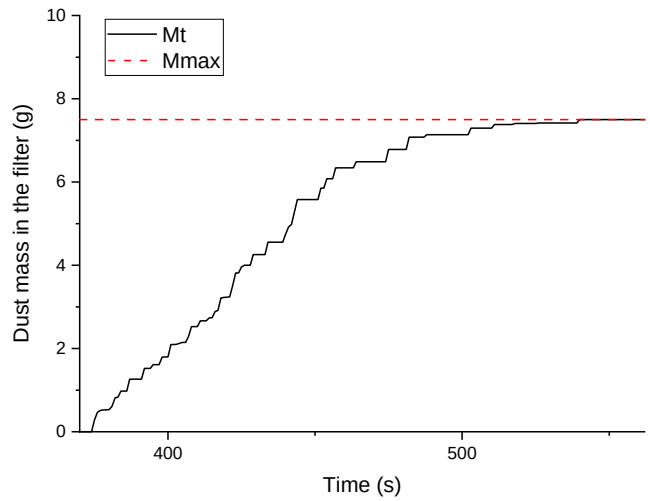

Figure 4. Dust dosing profile obtained by minimizing the sum of squared errors between experimental data and model predictions.

At higher dust loadings, the filter cake is subjected to increased compressive stresses, which can modify its internal structure and, consequently, its transport properties, including permeability and porosity [10]. Furthermore, high dust accumulation may promote local compaction, cracking, and channeling within the cake, resulting in non-uniform permeability and flux redistribution across the filter. These structural variations can induce flow instabilities and transient pressure fluctuations, contributing to the increased scatter observed in the experimental data at high dust loadings.

CONCLUSIONS & OUTLOOK

Overall, the validation suggests that the model provides a reasonable prediction of the system, while also indicating areas where improvements could be made. Accounting for potential effects such as cake compressibility and possible time-dependent variations in filter and specific cake resistance may further enhance predictive accuracy. Additionally, maintaining a fully consistent solid dosing profile, where only the final point, but not the full path, is controlled, introduces variability that affects model performance.

As no models for the filtration of superheated steam containing paper-like dust have been reported, this model provides a useful starting point for exploring this field and guiding future development. With validation under a single operating condition, the next step is to extend the model to a broader range of process scenarios, using statistical analysis of the fitted parameters to capture the influence of varying conditions. The filtration model will then be coupled with a filter-cleaning model to capture system performance over time. With this framework in place, a sensitivity analysis will be performed to identify the parameters and operational factors with the

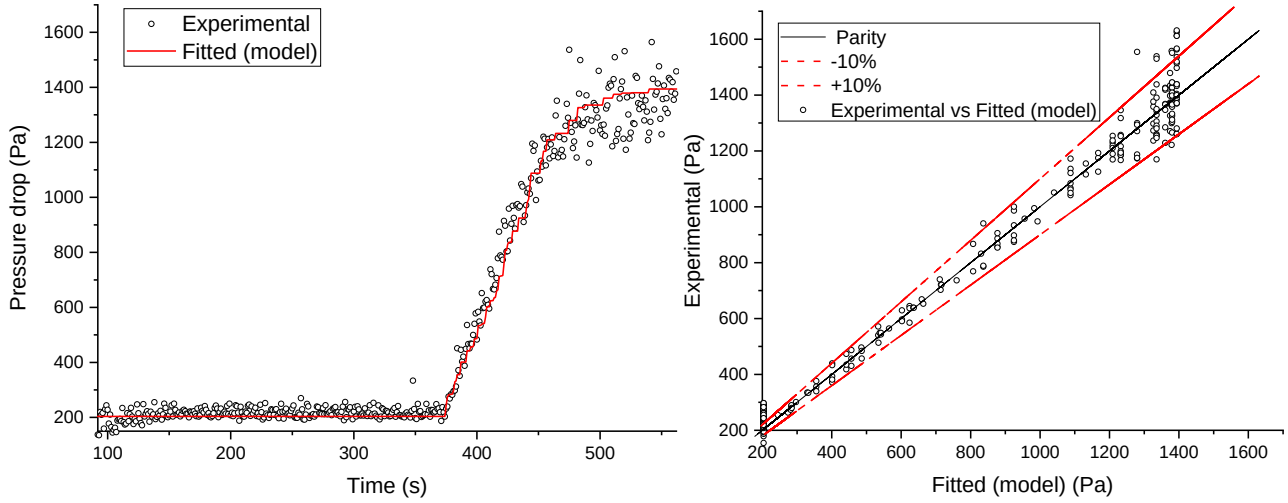

Figure 5. A- Comparison between experimental data and fitted model predictions assuming an adjusted solid dosing profile. B- Parity plot comparing experimental data and fitted model predictions, including ±10% deviation.

greatest impact on the process. These insights will guide the development of a dynamic optimization and model-based control strategies that account for variability and uncertainty in operation, such as fluctuations in solid loading or particle composition, supporting energy-efficient and cost-effective industrial applications. The proposed structure is compatible and commonly used in Computer-Aided Process Engineering.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of the SteamDry consortium (grant agreement 101137906) and the technical assistance provided by the staff of the University of Twente.

AUTHOR IDENTIFIERS

Author ORCIDs:
Oliveira F: 0000-0002-1160-4366
Nijhuis WJ: 0009-0005-6700-7180
Meinders MBJ: 0000-0001-9903-9826
Zondervan E: 0000-0003-2659-1622

REFERENCES

1. Eurostat. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Final_energy_consumption_in_industry_-_detailed_statistics (accessed 2025-01-21)
2. Wilk V, Knöttner S, Drexler-Schmid G, Barz T. Superheated steam drying for paper production: process efficiency assessment. 36th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems (ECOS 2023) :1411-1421 (2023). https://doi.org/10.52202/069564-0128
3. Hermia J. Constant pressure blocking filtration law: application to power law non-newtonian fluids. *Trans Inst Chem Eng* **60**:183 (1982)
4. Hlavacek M, Bouchet F. Constant flowrate blocking laws and an example of their application to dead-end microfiltration of protein solutions. Journal of Membrane Science 82:285-295 (1993). https://doi.org/10.1016/0376-7388(93)85193-z
5. Liu QF, Kim SH. Evaluation of membrane fouling models based on bench-scale experiments: a comparison between constant flowrate blocking laws and artificial neural network (anns) model. Journal of Membrane Science 310:393-401 (2008). https://doi.org/10.1016/j.memsci.2007.11.020
6. Köry J, Krupp A, Please C, Griffiths I. Optimising dead-end cake filtration using poroelasticity theory. Modelling 2:18-42 (2021). https://doi.org/10.3390/modelling2010002
7. Saleem M, Krammer G, Khan RU, Tahir MS. Influence of operating parameters on cake formation in pilot scale pulse-jet bag filter. Powder Technology 224:28-35 (2012). https://doi.org/10.1016/j.powtec.2012.02.016
8. Xie B, Li S, Chu W, Liu C, Hu S, Jin H, Zhou F. Improving filtration and pulse-jet cleaning performance of metal web filter media by coating with polytetrafluoroethylene microporous membrane. Process Safety and Environmental Protection 136:105-114 (2020). https://doi.org/10.1016/j.psep.2020.01.002
9. Bell IH, Wronski J, Quoilin S, Lemort V. Pure and pseudo-pure fluid thermophysical property evaluation and the open-source thermophysical

property library coolprop. Ind. Eng. Chem. Res. 53:2498-2508 (2014). https://doi.org/10.1021/ie4033999

10. Chen YS, Hsiau SS. Influence of filtration superficial velocity on cake compression and cake formation. Chemical Engineering and Processing: Process Intensification 48:988-996 (2009). https://doi.org/10.1016/j.cep.2009.01.006

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Robust Design of Transient Flow Experiments for the Identification of Kinetic Models in Flow Reactor Systems with Catalyst Deactivation

Jinwen Cui[ab], Federico Galvanin[a]*

[a] University College London, Department of Chemical Engineering, London, WC1E 7JE, United Kingdom
[b] Tsinghua University, Department of Chemical Engineering, Beijing, 100084, China
* Corresponding Author: f.galvanin@ucl.ac.uk

ABSTRACT

Catalyst deactivation significantly affects reactor performance, process efficiency, and economic viability in chemical processes. The precise estimation of kinetic and deactivation parameters in transient tubular reactors is essential but remains challenging due to strong parameter correlations, nonlinear dynamics, and limited prior knowledge of parameter values. Model-based Design of Experiments techniques for improving parameter precision (MBDoE-PP) has been shown to enhance parameter identifiability by optimally designing informative transient experiments even when catalyst deactivation occurs. However, MBDoE-PP is highly sensitive to parameter misspecification and can lead to suboptimal or infeasible solutions under model uncertainty, in which is the typical case in reaction systems exhibiting catalyst deactivation. In this work, a robust MBDoE framework for parameter precision (RMBDoE-PP) is proposed to explicitly account for process-model parameter mismatch (PMPM) during the experimental design stage. In the framework, stochastic sampling is used to evaluate metrics of the Fisher Information Matrix (FIM) and the experimental design is carried out considering a worst-case scenario, where the globally least informative parameter corresponding to the minimum information content across the uncertainty domain is identified. The performance of MBDoE-PP and RMBDoE-PP is assessed in the sequential closed-loop experimental design framework under varying levels of PMPM. Results demonstrate that RMBDoE-PP consistently yields tighter confidence intervals and more accurate parameter estimates with fewer experimental runs, confirming its superior robustness for deactivation kinetics.

Keywords: Design of Experiments, Catalyst Deactivation, Model-based Design of Experiments (MBDoE), Robustness, Dynamic modelling, Transient Experiments, Parameter Estimation.

1. INTRODUCTION

Catalysts are essential to chemical manufacturing and are extensively used across industrial processes, where they are typically operated over long time horizons. In practice, effective catalyst screening and selection require not only high initial activity but also a quantitative understanding of catalyst stability and deactivation behavior under realistic operating conditions. Catalyst deactivation, defined as the gradual loss of catalytic activity with time on stream, is an inevitable phenomenon that poses significant challenges to process operation and optimization, often leading to reduced productivity and adverse economic impacts [1][2]. Industrial decisions regarding catalyst selection, utilization, and replacement scheduling therefore strongly depend on the accurate quantification and prediction of catalyst activity evolution. Reliable identification of deactivation behavior enables optimized catalyst usage, avoids premature disposal or excessive replacement, reduces catalyst waste, and ultimately improves process profitability.

However, catalyst activity cannot be directly measured, and the identification of deactivation kinetics is inherently complex [3]. This complexity arises from the time-varying nature of catalyst deactivation, uncertainty in reaction kinetics, strong parameter correlations, and limited prior knowledge of model parameters and

operating conditions. Dynamic mechanistic models are commonly employed to describe and predict reactor behavior in the presence of catalyst deactivation, but the development of reliable deactivation models remains challenging. To improve parameter identifiability while reducing experimental effort, transient reactor operation has been increasingly explored which can generate rich time-dependent data without the need to wait for steady-state conditions, thereby reducing experimental time and preventing information loss.

Model-Based Design of Experiments for parameter precision (MBDoE-PP) provides a systematic framework for addressing these challenges by optimally designing experiments to maximize information and reduce the number of experiments needed for model identification [4]. In pevious studies the effectiveness of MBDoE-PP under transient conditions was clearly demonstrated, but in systems where catalyst deactivation was not significantin [5]. In our recent work, MBDoE-PP has been shown to be a powerful approach for accurately identifying deactivation kinetics, but, the effectiveness of MBDoE-PP critically depends on the quality of the starting parameter estimates used during the experimental design stage [6]. When a significant process-model parameter mismatch (PMPM) is present, or when little or no prior knowledge on parameters is available, experimental designs optimized at incorrect nominal values may be poorly informative, leading to slow convergence and unreliable parameter [7]. To address these limitations, a robust Model-Based Design of Experiments (RMBDoE-PP) framework is adopted in this work. Robust experimental design approaches aim to identify experimental conditions that maintain satisfactory performance over a range of plausible parameter values, rather than relying on the assumption that nominal parameter estimates are accurate [8]. However, designs based on expected information may be biased by a few highly informative parameter realisations, masking poor performance elsewhere. The proposed RMBDoE-PP addresses this limitation by targeting robustness, considering the worst-case scenarios, ensuring reliable information content regardless of initial parameter guesses. In this study, MBDoE-PP and RMBDoE-PP are investigated and compared for the design of transient experiments in a tubular reactor with catalyst deactivation.

2. METHODOLOGY

The reactor model consists of a system of dynamic partial differential and algebraic equations (DAE) characterized by a set of estimated parameters θ, state variables $\boldsymbol{x}(\boldsymbol{t})$ and the time varying controls $u(t)$ and time invariant controls $w(t)$, which can be used to predict the behaviour of the reactor system. The model enables the prediction of experimentally measurable values $\hat{\mathbf{y}}$:

$$f\left(\dot{\boldsymbol{x}}(\boldsymbol{t}), \boldsymbol{x}(\boldsymbol{t}), \boldsymbol{u}(\boldsymbol{t}), \boldsymbol{w}, \theta, t\right) = 0 \tag{1}$$

$$\hat{\mathbf{y}}(\boldsymbol{t}) = \boldsymbol{g}\left(\boldsymbol{x}(\boldsymbol{t})\right) \tag{2}$$

The Fisher Information Matrix (FIM) for a sample point in the design space can be expressed as follows [9][10]:

$$FIM_t = Q_t^T \cdot \Sigma_y{}^{-1} \cdot Q_t \ = [V_\theta]^{-1} \tag{3}$$

where Q_t is the matrix of sensitivities whose q_{ij} element can be defined as

$$q_{ij} = \frac{\partial \hat{y}_i}{\partial \theta_j} \tag{4}$$

and V_θ is the variance-covariance matrix of model parameters, $\Sigma_y{}^{-1}$ is the inverse matrix of measurement errors Σ_y which can be defined as

$$\Sigma_y = \begin{bmatrix} \sigma_1{}^2 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \sigma_y{}^2 \end{bmatrix} \tag{5}$$

where σ_i is the standard deviation of measurement error for the i_{th} response.

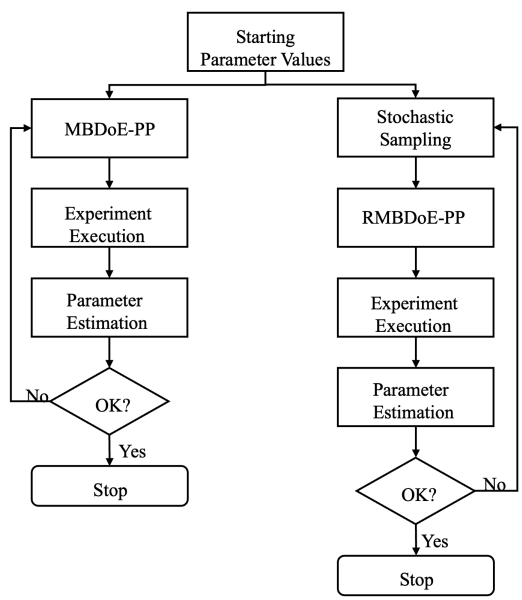

Figure 1. The Model Validation Procedure of MBDoE-PP and RMBDoE-PP.

Figure 1 illustrates the iterative model validation and experimental design workflow adopted for both MBDoE-PP and RMBDoE-PP frameworks. In the presence of limited or no prior knowledge of the model parameters, the same set of initial parameter guesses is used to initialise both design strategies. At each iteration, optimal experimental conditions are first determined based on the current parameter information and subsequently applied to a worst-case scenario, which is assumed to represent the true operating conditions of the real

system. In-silico experimental data are then generated by simulating the model responses under additive random measurement noise. Parameter estimation is performed using the generated data, and the resulting estimates are evaluated through a set of statistical model adequacy and parameter significance tests. The iterative procedure is repeated until predefined precision and identifiability criteria are satisfied, at which point the identification process is terminated.

This iterative framework follows the general philosophy employed in [11][12], where successive experiments are designed based on information gained from previous iterations. More recent studies, such as [13], further developed feedback-based experimental design strategies by explicitly incorporating updated parameter knowledge into subsequent design steps. In the present work, this iterative structure is combined with an A-optimality criterion, which aims to minimise the trace of the parameter variance–covariance matrix [9], equivalently maximising the integrated trace of the Fisher Information Matrix over the experimental horizon, thereby improving overall parameter estimation robustness.

2.1 Standard MBDoE-PP formulation

For MBDoE-PP, the experimental design problem is formulated as an optimization task where the objective is to determine the optimal time-invariant input variables, denoted as $w(t)$, that maximize the integrated trace of the Fisher Information Matrix (FIM). The trace of the FIM, which is the sum of its diagonal elements, serves as a crucial metric for evaluating the information gained from an experiment, particularly in scenarios where frequent sampling is implemented. This metric provides an effective means of assessing the overall sensitivity of the system to variations in its parameters, ensuring that the designed experiment captures maximal information for parameter estimation. The trace of the FIM can also be leveraged for the dynamic assessment of information gain under the assumption of a constant sampling time interval Δt. By integrating the trace of the FIM over the duration of the experiment, a cumulative measure of information can be obtained [14]

$$\varphi^{opt}_{MBDoE-PP} = \arg\max_{\varphi} \left\{ \int_0^{t_h} tr(FIM)dt \right\} \tag{6}$$

where φ^{opt} is the optimal experimental design vector, containing all the decision variables that can be manipulated during a dynamic experiment. The design accounts for the time-dependent nature of the system's response and allows for an optimal design strategy that ensures the greatest reduction in parameter uncertainty. By optimizing the criterion objective (6), the design ensures that the chosen experimental conditions provide the most informative dataset possible, enhancing the precision and reliability of parameter estimation.

2.2 Determination of the worst-case scenario

Worst-case (minimax) experimental design seeks to optimise performance under the most adverse parameter realisation within a predefined uncertainty set [6]. In this study, a worst-case scenario is explicitly constructed to assess the robustness of MBDoE-PP under parameter uncertainty. A parameter domain $\boldsymbol{\Theta}$ is defined to represent plausible model parameters, and multiple initial guesses are selected within this domain to compute corresponding MBDoE-PP optimal experimental conditions. For each design, a stochastic simulation is employed as an evaluation tool, where model parameters are uniformly sampled within $\boldsymbol{\Theta}$ and the cumulative information metric is assessed under fixed experimental conditions. The worst-case scenario is defined as the parameter yielding the globally minimum cumulative information across all evaluated MBDoE-PP designs and is adopted as the true value for subsequent analysis. This formulation provides a conservative and systematic benchmark for assessing the robustness of experimental design strategies under adverse operating conditions.

2.3 RMBDoE

RMBDoE extends MBDoE-PP framework by explicitly accounting for process model parameter mismatch (PMPM) in the optimization. Instead of solving the MBDoE optimization problem based on a local FIM evaluation computed at a nominal parameter set, RMBDoE considers a distribution of FIM metric values based on the expected variability in the parameter space $\boldsymbol{\Theta}$. To obtain the distribution, a stochastic simulation is employed and embedded in the MBDoE optimization, where multiple random samples of parameter sets are drawn from a distribution centered around nominal parameter values with a specified variance to reflect the expected uncertainty range. Each sampled parameter set is then used to compute the FIM metric for a given experimental design. Here the objective of RMBDoE is formulated to optimize experimental conditions such that the ratio between the mean of the integral of the FIM trace (μ_I) over the corresponding standard deviation (σ_I) across all sampled parameter sets is maximised:

$$\varphi^{opt}_{RMBDoE-PP} = \arg\max_{\varphi} \left(\frac{\mu_I}{\sigma_I} \right) \tag{7}$$

with

$$\mu_I = \frac{\sum_{i=1}^{N_s} \left[\int_0^{t_h} tr(FIM)dt \right]_i}{N_s} \tag{8}$$

$$\sigma_I = \sqrt{\frac{1}{N_s} \sum_{i=1}^{N_s} \left(\left[\int_0^{t_h} tr(FIM)dt \right]_i - \mu_I \right)^2}$$

and where N_s is the number of scenarios used in the embedded stochastic simulation. The geometrical interpretation of criterion (7) is that the model-based design optimization is seeking an optimal set of experimental

conditions generating a flat and high information surface in the parameter space $\boldsymbol{\Theta}$ so that the design can be insensitive to parametric uncertainty.

2.4 Parameter Estimation

The estimation of model parameters is performed using the maximum likelihood estimation (MLE) method, which ensures that the chosen parameters maximize the probability of observing the collected experimental data [15]. The parameter estimation performance is assessed by using a χ^2 lack-of-fit test which provides a quantitative measure of the agreement between the model predictions and the observed experimental data,

$$\chi^2 = \sum_{i=1}^{N_{exp}} \sum_{j=1}^{N_M} \frac{(y_{ij}-\hat{y}_{ij})^2}{\sigma_{ij}^2} \tag{9}$$

and a standard parameter *t*-value test to assess parameter precision, with *t*-values t_i calculated from

$$t_i = \frac{\hat{\theta}_i}{\sqrt{v_{ii}}} \tag{10}$$

where v_{ii} is the variance of the i_{th} parameter (i.e. the i_{th} diagonal element of the parameter variance-covariance matrix). The *t*-value for each parameter is tested against a tabulated reference value related to a Student's t-distribution. The quality of each experimental design is assessed based on the statistical reliability of the fitted model and estimated parameters. χ^2 test and t-value test are the two key tests used for evaluation of parameter estimates: (1) the χ^2 test is used to assess how well the model provides an adequate representation of the experimental data, by comparing the computed χ^2 with the reference value χ^2_{ref}. If the χ^2 result from designed experiments is lower than χ^2_{ref}, the kinetic model provides a statistically adequate fit to the simulated experimental data, otherwise, the model is considered inadequate for describing the observed behavior; (2) the t-value test indicates the confidence interval in each parameter estimate and all the *t*-values of model parameters must exceed a critical reference *t*-value to be considered statistically significant.

3. CASE STUDY

The optimal design of transient dynamic experiments in a tubular reactor with catalyst deactivation using MBDoE-PP is illustrated through a case study on the oxidation of o-xylene to phthalic anhydride[16]. The reaction takes place in a gas phase tubular reactor which is packed with decaying catalyst particles and where a cooling jacket is positioned around the reactor, as shown in Figure 2. The same experimental setup is also used in other kinetic modelling studies[17][18].

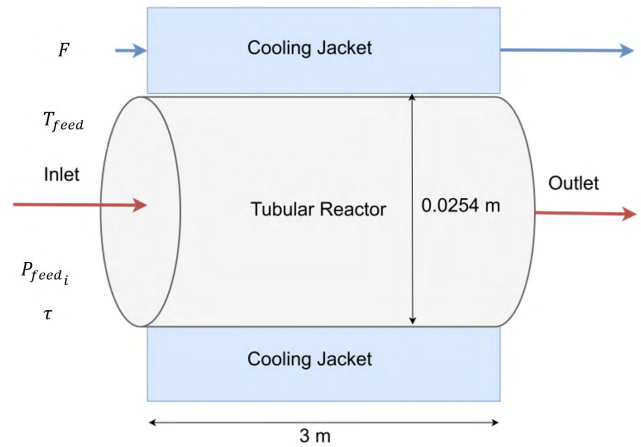

Figure 2. Sketch of the experimental system and typical dynamic response for the tubular reactor with cooling jacket operating in transient state. The set of inputs u is defined by the following designed variables: i) initial concentration of the reactants P_{feed_i}; ii) residence time τ; iii) feed temperature of cooling fluid T_{feed} and iv) coolant flowrate in the cooling jacket F.

The mass balance of the system is defined as

$$\frac{\partial C_i}{\partial t} = -v_g * \frac{\partial C_i}{\partial z} + \varepsilon * D_z * \frac{\partial^2 C_i}{\partial z^2} + \varepsilon * D_r * \left(\frac{\partial^2 C_i}{\partial r^2} + \frac{1}{r}\frac{\partial C_i}{\partial r}\right) + v_i * \rho_b * R_{rxn} \; z \in [0, L], r \in [0, R] \tag{11}$$

$$R_{rxn} = A * e^{-\frac{E}{RT}} * (RT)^2 * C_A * C_B$$
$$z \in (0, L), r \in (0, R) \tag{12}$$

The tubular reactor is non-isothermal, but it is assumed that heat losses can be neglected and that the heat of reaction goes into the bulk fluid phase rather than into heating up the catalyst particles. The energy balance in the gas phase accounts for the axial convection heat transport, the axial and radial dispersion energy, and the heat generated or consumed by the chemical reaction as reported below.

$$\rho_f * C_{pf} * \frac{\partial T}{\partial t} = -\rho_f * C_p * v_g * \frac{\partial T}{\partial z} + k_z * \frac{\partial^2 T}{\partial z^2} + k_r\left(\frac{\partial^2 T}{\partial r^2} + \frac{1}{r}\frac{\partial T}{\partial r}\right) + \rho_b R_{rxn}(-\Delta H) \tag{13}$$

$$z \in (0, L), r \in (0, R)$$

To address the high correlation, a reparameterization of the original Arrhenius equation including a reference temperature is needed[19] to minimize the correlation between the two kinetic parameters as follows:

$$k = A_{ref} * e^{-\frac{E}{R}\left(\frac{1}{T} - \frac{1}{T_{ref}}\right)} \tag{14}$$

$$A_{ref} = A * e^{-\frac{E}{RT_{ref}}} \tag{15}$$

As the reaction occurs in the gas phase and assumes ideal gas behavior, the ideal gas law is applied, and the

RT term is incorporated into the expressions for the molar concentration or partial pressure of each component.

$$R_{rxn} = k * (RT)^2 * C_A * C_B \quad (16)$$

The Levenspiel's Deactivation Kinetic Model (LDKM) is selected to represent the loss of catalyst activity (a) overtime[20]. It is assumed that the catalyst activity changes along with the radius and the length, but it is assumed to be the same for all of the components which can be expressed as

$$-\frac{da}{dt} = \psi_d * a \quad (17)$$

$$R_{rxn_{DM1}} = k_{DM1} * (RT)^2 * C_A * C_B * a \quad (18)$$

In Equation 17, ψ_d is assumed to be the deactivation parameter that needs to be estimated.

In this study, the gPROMS modelling environment is used for simulation, optimization as well as for parameter estimation and to optimally design dynamic experiments [21]. The mathematical model of the reactor system is established assuming that the measurable variables are the concentrations of each component at the outlet ($C_{out,i}$). Upper and lower bounds on experimental conditions defining the experimental design space are shown in Table 1.

Table 1. Definition of the experimental design space

Manipulated Variables	Lower Bound	Upper Bound
F	$0.08\,\frac{kg}{s}$	$0.12\,\frac{kg}{s}$
τ	$2.74\ s$	$4.10\ s$
T_{feed}	$570\ K$	$625\ K$
$P_{feed,A}$	$0.880\ kPa$	$1.32\ kPa$

The experimental design vector variables are implemented using a time-invariant input function within a fixed time interval. Four experimental conditions are manipulated simultaneously within this framework to evaluate their combined effects on the system's information which allow for the interaction between conditions to be captured, enhancing the sensitivity of the design. In MBDoE it is assumed that a very frequent sampling in time is possible, i.e. n_{sp} = 361 sampling points can be collected within a fixed time horizon of $\tau_{horizon} = 3hrs$. The experimental data is generated in-silico by assuming that are corrupted by a normally distributed noise with a zero mean and $\Sigma_y = \begin{bmatrix} 0.002 & 0 & 0 \\ 0 & 0.4 & 0 \\ 0 & 0 & 0.02 \end{bmatrix}$. This level of noise amounts to about 10% error on the average values of $C_{out,i}$.

The predefined model parameter domain Θ is defined as [0.5, 1.5]. The estimated parameters are the ratio between the reference value θ_{ref_i} of parameters and the estimated one as

$$\hat{\theta}_{ratio_i} = \frac{\hat{\theta}_i}{\theta_{ref_i}} \quad (19)$$

where $\hat{\theta}_i$ is the *i*-th parameter [*i*=1, 2, 3], $\hat{\theta}_1$ is the estimated activation energy, $\hat{\theta}_2$ is the estimated Arrhenius constant and $\hat{\theta}_3$ is the deactivation parameter through equations (14-17). To identify the potential worst-case scenario, MBDoE-PP evaluations are performed for multiple initial parameter guesses to identify the respectively optimal experimental conditions (Table 2). Stochastic simulation is then used to evaluate the information distribution over the predefined parameter domain (Figure 3), and the set of parameters yielding the minimum information level are identified as the worst-case scenario. MBDoE-PP seeks to maximize the integral of the trace FIM, optimizing input profiles under the assumption of fixed nominal parameters. In contrast, RMBDoE-PP maximizes the ratio of the mean to standard deviation of the FIM across a parameter distribution here assumed as $\pm 20\%$ of its initial guess, reducing information loss due to poor initial guesses. Once the designed experiment is executed under the optimized input conditions, parameter estimation is performed using the newly acquired data to update the model. The two optimal experimental design approaches, MBDoE-PP and RMBDoE-PP, are compared in terms of statistical tests (χ^2 test to assess the lack-of-fit and t-test to evaluate the precision in parameter estimates, see Section 2.4 for details).

4. RESULT

4.1 Identification of the worst-case scenario

MBDoE-PP optimisations were conducted using multiple sets of initial parameter guesses, resulting in a corresponding set of optimal experimental conditions, which are summarised in Table 2. The column denotes the initial guesses for the three kinetic parameters. A single value indicates identical initial guesses for all parameters, while the grouped column "0.5; 0.6; 0.7" corresponds to multiple starting points that yields the same optimal experimental design. For each optimal design, a fixed set of 50 parameters uniformly sampled over the predefined parameter domain were generated to ensure a consistent evaluation framework across all designs. These parameters were used to compute the cumulative information metric under the respective optimal experimental conditions.

The resulting information distributions are presented in Figure 3, which illustrate the variation in information content as a function of parameter realisation and experimental design. From these distributions, the globally worst-case scenario was defined where the

parameter has the minimum cumulative information across all evaluated designs and parameter samples was identified.

These determined true parameter values $\hat{\theta}_{ratio_i} = (1.43, 1.20, 1.19)$ are used to represent the true system for subsequent in-silico experiment execution and parameter estimation studies. Selecting the true value in this manner ensures that subsequent analyses are conducted under the most adverse yet admissible conditions for MBDoE-PP, thereby providing a conservative and meaningful basis for assessing estimation performance and robustness.

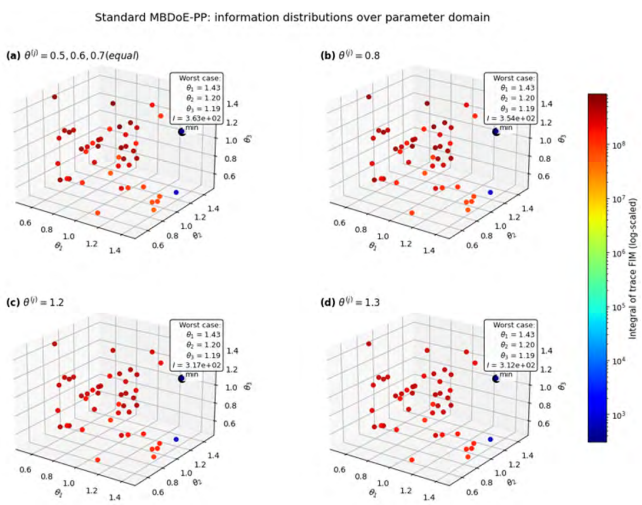

Figure 3. The MBDoE-PP information ditribution over parameter domian $\boldsymbol{\Theta}$

Table 2. The optimal experimental conditions computed by MBDoE-PP under different starting parameter values.

Design variable	MBDoE-PP Starting Parameter Value			
	$0.5; 0.6; 0.7$	0.8	1.2	1.3
$F\ (\frac{kg}{s})$	0.12	0.12	0.12	0.12
$\tau\ (s)$	4.10	4.10	4.10	4.10
$T_{feed}(10^2 K)$	5.70	5.73	5.84	5.86
$P_{feed,A}\ (kPa)$	1.32	1.32	1.32	1.32

4.2 MBDoE-PP and RMBDoE-PP

The guessed parameter vector $(1, 1, 1)$ is selected to initialise both MBDoE-PP and RMBDoE-PP at the center point of the predefined domain. Figure 4 illustrates the initial guessed values for both MBDoE techniques, the parameter uncertainty domain considered in the RMBDoE-PP formulation, and the location of the identified worst-case true value. In the robust design, parametric uncertainty is explicitly represented by a hyper-rectangular domain defined as $\pm 20\%$ around the nominal values, shown as the cubic region in the parameter space. Within this domain, 50 parameter realisations are uniformly sampled and used to evaluate the cumulative information metric under the corresponding optimal experimental conditions shown in Table 3.

The sampled parameter points within the cube represent admissible uncertainty scenarios accounted for during the robust design stage. In contrast, the identified worst-case true value lies outside the predefined uncertainty domain. This point corresponds to the globally minimum information obtained from the stochastic evaluation of all MBDoE-PP designs and therefore represents the most adverse possible parameter set for parameter estimation.

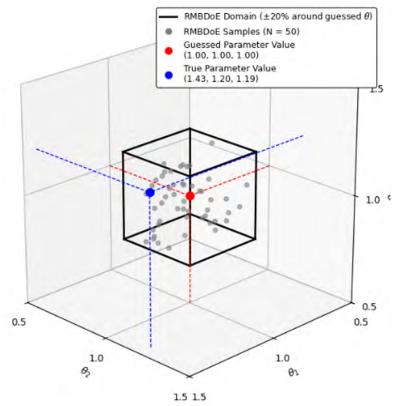

Figure 4. Guessed parameter values for both MBDoE-PP techniques and the true parameter value to represent the real system.

Table 3. Optimal experimental conditions computed by MBDoE-PP and RMBDoE-PP.

	MBDoE-PP	RMBDoE-PP
$F\ (\frac{kg}{s})$	0.12	0.08
$\tau\ (s)$	4.10	2.74
$T_{feed}(10^2 K)$	5.70	6.07
$P_{feed,A}\ (kPa)$	1.32	0.88

4.3 Parameter Estimation

The results of the χ^2 test and statistical test results for both MBDoE-PP techniques are shown in Table 4, confirm good model-data agreement, with all total χ^2 values falling below the reference value of $1.16 * 10^3$, indicating satisfactory model–data agreement. Despite this overall adequacy, significant differences are observed in the quality of parameter estimation between the two design strategies. Under the standard MBDoE-PP framework, parameter estimation performance is poor: none of the parameters are estimated with satisfactory precision. The estimation of activation energy ($\hat{\theta}_1$) converges to its lower bound, and the estimation of the deactivation parameter ($\hat{\theta}_3$) exhibits poor precision. The estimation of Arrhenius constant ($\hat{\theta}_2$) fails to reach the reference t-value,

Table 4. Kinetic and deactivation parameter estimation results from MBDoE-PP and RMBDoE-PP in the presence of PMPM in terms of estimated values and a-posteriori statistics. Superscript asterisk ($*$) indicates t-value failing the t-test and ($**$) indicates that the parameter lies on the bounds and is excluded from the statistical analysis.

	Param. estimate + conf. interval (95%)	***t*-values ($t_{ref} = 1.65$)**	$\chi^2(\chi^2_{ref} = 1.16 * 10^3)$
MBDoE-PP	$\hat{\theta} = \begin{bmatrix} 0^{**} \\ 3.87 \pm 3.02 \\ 5.62 * 10^{-7^*} \pm 1.24 * 10^2 \end{bmatrix}$	$\begin{bmatrix} * \\ 1.28^* \\ 4.52 * 10^{-9^*} \end{bmatrix}$	$1.00 * 10^3$
RMBDoE-PP	$\hat{\theta} = \begin{bmatrix} 1.34 \pm 8.71 * 10^{-2} \\ 1.14 \pm 6.17 * 10^{-2} \\ 2.10 \pm 1.05 \end{bmatrix}$	$\begin{bmatrix} 1.53 * 10^1 \\ 1.85 * 10^1 \\ 1.98 \end{bmatrix}$	$1.08 * 10^3$

indicating poor precision achieved under the optimal experimental conditions.

In contrast, the RMBDoE-PP framework yields substantially better estimation performance. The kinetic parameters are estimated with narrow confidence intervals and *t*-values are well above the reference value, demonstrating statistically significant and reliable identification. The two kinetic parameters remain highly correlated (the correlation coefficient is 0.99), but both parameters are estimated with high precision and statistically significant t-values. This behaviour is consistent with the use of an A-optimality criterion, which prioritises overall parameter precision rather than minimising parameter correlations. The deactivation parameter also exhibits good precision, with its *t*-value exceeding the significance threshold and low correlation realised with the kinetic parameters ($\hat{\theta}_1$) and ($\hat{\theta}_1$), -0.26 and -0.16, respectively. This suggests that RMBDoE-PP markedly enhances robustness against PMPM with the catalyst deactivation.

4.4 CPU Time

The CPU time required by both MBDoE techniques are here compared. For the standard MBDoE-PP approach, the optimization step requires 4.86 minutes, followed by 3.38 minutes for parameter estimation. In contrast, the RMBDoE-PP approach incurs a substantially higher computational cost, with 27.88 hours (1.67×10^3 minutes) required for the robust optimisation step and 7.40 minutes for parameter estimation.

These results highlight a clear trade-off between computational effort and parameter estimation quality. While RMBDoE-PP achieves substantially higher parameter precision within the first iteration, this improvement comes at the expense of a significantly increased computational burden. Conversely, standard MBDoE-PP requires considerably less CPU time but yields negligible parameter precision in the first iteration, implying that multiple additional iterations and experiments would be required to reach comparable levels of precision. Consequently, the overall experimental and computational burden associated with standard MBDoE-PP may increase substantially when high parameter precision is desired.

5. CONCLUSION AND FUTURE WORK

The accurate modelling of tubular reactor operated under transient conditions with catalyst deactivation requires reliable estimation of both kinetic and deactivation parameters. In this work, standard MBDoE-PP and RMBDoE-PP are assessed under the presence of process-model parameter mismatch (PMPM). A worst-case scenario is identified and considered in the study in which the “true” system is defined as the parameter point generating the lowest information among all standard MBDoE-PP designs within a predefined parameter domain.

Under the worst-case scenario, standard MBDoE-PP exhibits poor parameter estimation performance, with low statistical significance and limited precision, indicating practical non-identifiability of some model parameters. In contrast, RMBDoE, which accounts for the PMPM, delivers substantially improved parameter estimates, achieving acceptable precision despite the presence of PMPM and limited prior information. These results demonstrate the advantage of incorporating parameter uncertainty directly into the experimental design stage when robustness against worst-case scenarios is required. However, the study highlights a clear trade-off between computational cost and estimation robustness and confirms the importance of robust design strategies for dynamic reactor systems affected by catalyst deactivation. Future work will focus on the development of hybrid design frameworks that combine standard MBDoE-PP and RMBDoE-PP, aiming to balance computational efficiency with robustness by adaptively integrating information gained during the experiment design and parameter estimation process.

ACKNOWLEDGEMENTS

UCL Department of Chemical Engineering for the visiting position of Jinwen Cui.

AUTHOR IDENTIFIERS

Author ORCIDs:
Jinwen Cui: 0009-0008-3198-0909
Federico Galvanin: 0000-0002-3296-1581

REFERENCES

1. Forzatti P. Catalyst deactivation. Catalysis Today 52:165-181 (1999). https://doi.org/10.1016/s0920-5861(99)00074-7
2. Argyle M, Bartholomew C. Heterogeneous catalyst deactivation and regeneration: a review. Catalysts 5:145-269 (2015). https://doi.org/10.3390/catal5010145
3. Shakor ZM, Al-Shafei EN. The mathematical catalyst deactivation models: a mini review. RSC Adv. 13:22579-22592 (2023). https://doi.org/10.1039/d3ra02912c
4. Asprey SP, Macchietto S. Statistical tools for optimal dynamic model building. Computers & Chemical Engineering 24:1261-1267 (2000). https://doi.org/10.1016/s0098-1354(00)00328-8
5. Waldron C, Pankajakshan A, Quaglio M, Cao E, Galvanin F, Gavriilidis A. Model-based design of transient flow experiments for the identification of kinetic parameters. React. Chem. Eng. 5:112-123 (2020). https://doi.org/10.1039/c9re00342h
6. Asprey SP, Macchietto S. Designing robust optimal dynamic experiments. Journal of Process Control 12:545-556 (2002). https://doi.org/10.1016/s0959-1524(01)00020-8
7. Galvanin F, Barolo M, Bezzo F, Macchietto S. A backoff strategy for model?based experiment design under parametric uncertainty. AIChE Journal 56:2088-2102 (2009). https://doi.org/10.1002/aic.12138
8. Pronzato L, Walter E. Robust experiment design via stochastic approximation. Mathematical Biosciences 75:103-120 (1985). https://doi.org/10.1016/0025-5564(85)90068-9
9. Franceschini G, Macchietto S. Model-based design of experiments for parameter precision: state of the art. Chemical Engineering Science 63:4846-4872 (2008). https://doi.org/10.1016/j.ces.2007.11.034
10. Zullo L. Computer-aided design of experiments: an engineering approach. PhD thesis. University of London; 1991.
11. Buzzi-Ferraris G, Forzatti P. A new sequential experimental design procedure for discriminating among rival models. Chemical Engineering Science 38:225-232 (1983). https://doi.org/10.1016/0009-2509(83)85004-0
12. Galvanin F, Macchietto S, Bezzo F. Model-based design of parallel experiments. Ind. Eng. Chem. Res. 46:871-882 (2007). https://doi.org/10.1021/ie0611406
13. Barz T, Arellano-Garcia H, Wozny G. Handling uncertainty in model-based optimal experimental design. Ind. Eng. Chem. Res. 49:5702-5713 (2010). https://doi.org/10.1021/ie901611b
14. Galvanin F, Boschiero A, Barolo M, Bezzo F. Model-based design of experiments in the presence of continuous measurement systems. Ind. Eng. Chem. Res. 50:2167-2175 (2011). https://doi.org/10.1021/ie1019062
15. Bard Y. Comparison of gradient methods for the solution of nonlinear parameter estimation problems. SIAM J. Numer. Anal. 7:157-186 (1970). https://doi.org/10.1137/0707011
16. BOND G. Mechanism of the oxidation of o-xylene to phthalic anhydride. Journal of Catalysis 116:531-539 (1989). https://doi.org/10.1016/0021-9517(89)90119-x
17. Bendaho D, Kaid N, Alqahtani S, Alshammari BM, Menni Y, Chamkha AJ, Kolsi L. Cfd-aided enhancement of propylene oxide decomposition in a tubular reactor with a corrugated cooling jacket. Case Studies in Thermal Engineering 61:105004 (2024). https://doi.org/10.1016/j.csite.2024.105004
18. Memon AA, Memon MA, Bhatti K, khan I, Alshammari N, Al-Johani AS, Hamadneh NN, Andualem M. Thermal decomposition of propylene oxide with different activation energy and reynolds number in a multicomponent tubular reactor containing a cooling jacket. Sci Rep 12: (2022). https://doi.org/10.1038/s41598-022-06481-4
19. Schwaab M, Lemos LP, Pinto JC. Optimum reference temperature for reparameterization of the arrhenius equation. part 2: problems involving multiple reparameterizations. Chemical Engineering Science 63:2895-2906 (2008). https://doi.org/10.1016/j.ces.2008.03.010
20. Levenspiel O. Experimental search for a simple rate equation to describe deactivating porous catalyst particles. Journal of Catalysis 25:265-272 (1972). https://doi.org/10.1016/0021-9517(72)90227-8
21. Process Systems Enterprise. gPROMS® v.2022.1.0. London; 2022. https://www.psenterprise.com/products/gproms

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Process modelling and multi-objective optimisation of solid sorbent-based direct air capture

Toluleke E. Akinola[a], Meihong Wang[a]*

[a] The University of Sheffield, School of Chemical, Materials and Biological Engineering, Sheffield S1 3JD, South Yorkshire, United Kingdom
* Corresponding Author: meihong.wang@sheffield.ac.uk

ABSTRACT

Direct Air Capture (DAC) is recognised as a critical climate mitigation technology necessary for achieving global net-zero emissions by balancing difficult-to-avoid emissions. Despite its importance, the commercial deployment of DAC technology is currently challenged by the substantial energy demands and resultant extremely high operational costs associated with handling the low atmospheric CO_2 concentration. This study aims to address these two challenges through dynamic process modelling, simulation and rigorous multi-objective optimisation of a solid sorbent-based temperature vacuum swing adsorption (S-TVSA) cycle. The system utilises an advanced amine-functionalized sorbent, selected for its favourable low-temperature regeneration kinetics.A technical performance assessment was conducted using a first-principle mathematical model to accurately simulate mass and heat transfer in the adsorbent beds. To improve system viability, a multi-objective optimisation with NSGA-II in MATLAB was performed, maximising CO_2 productivity and minimising energy demand. The optimisation targeted key parameters like adsorption/desorption time, temperature, and pressure. The Pareto frontier shows trade-offs, identifying optimal points that minimise W_{eq} and energy use. These results provide a pathway to reduce energy costs and enhance the economic viability of solid-sorbent DAC technology.

Keywords: Direct air capture, amine-functionalised sorbent, process modelling, performance evaluation, Process optimisation, temperature vacuum swing adsorption

1. INTRODUCTION

A significant driver of climate change is the rapidly rising atmospheric concentration of greenhouse gases including CO_2 [1]. Atmospheric CO_2 levels have increased from 280 ppm in the pre-industrial era to 425 ppm in 2025 [2]. As a result, global temperatures have risen, posing a threat to the environment. To limit the temperature rise to 1.5–2°C, achieving net-zero CO_2 emissions by 2050 is essential [3].

Solid-based direct air capture (DAC) technology plays a key role in meeting net-zero targets [4]. However, the diluted atmospheric CO_2 concentration poses challenges for meeting energy demand and for the cost of CO_2 extraction, as large volumes of air must be processed to extract a significant amount of CO_2 [1, 4, 5].

Previous studies [1, 5] demonstrated that a steam-assisted temperature-vacuum swing adsorption (S-TVSA) process can achieve high-purity CO_2 (≥ 95%) from ambient air. In S-TVSA, CO_2 is initially adsorbed onto an adsorbent under ambient conditions. Subsequently, the pressure is reduced, and steam is used to increase the temperature, releasing CO_2. The process alternates between adsorption and regeneration to capture CO_2. Designing these processes is complex due to the nonlinear mass and heat transfer within the adsorbent column [19]. Operational decisions, such as cycle step durations, temperatures, pressures, and flow rates, directly influence performance metrics. This study aims to optimise these operational parameters to maximise productivity, reduce energy consumption, and ensure high CO_2 purity. This study utilises 3-aminopropylmethyldiethoxysilane-functionalized nanofibrilated cellulose (APDES-NFC), an advanced amine- functionalized sorbent due to its (i) favourable low temperature kinetics and (ii) reliable performance in the presence of moisture [1, 5].

2. METHODOLOGY

2.1 Model development

2.1.1 Adsorption Equilibrium

The co-adsorption of CO_2 and H_2O on the APDES-NFC surface was modelled using a combination of the modified Toth isotherm for CO_2 and the Guggenheim-Anderson-de Boer (GAB) isotherm for H_2O. The modified Toth isotherm model is defined as:

$$q_{co_2}(T, p_{co_2}, q_{H_2O}) = n_s(T, q_{H_2O}) \frac{b(T, q_{H_2O}).\, p_{co_2}}{\left[1 + (b(T, q_{H_2O}).\, p_{co_2})^{t(T)}\right]^{1/t(T)}} \quad (1)$$

$$n_s(T, q_{H_2O}) = \left[n_{s_0} \exp\left[\chi\left(1 - \frac{T}{T_0}\right)\right]\right] . \left[\frac{1}{1 - \gamma_{q_{H_2O}}}\right] \quad (2)$$

$$b(T, q_{H_2O}) = \left[b_0 \exp\left[\frac{\Delta H_0}{RT_0}\left(\frac{T}{T_0} - 1\right)\right]\right] . \left[1 + \beta_{q_{H_2O}}\right] \quad (3)$$

$$t(T) = t_0 + \alpha\left(1 - \frac{T_0}{T}\right) \quad (4)$$

Where q_{co_2} is specific adsorbed amount of CO_2 (mol/kg), n_s is the saturation capacity of CO_2 (mol/kg), p_{co_2} is CO_2 partial pressure (kPa), b_0 is the adsorption equilibrium constant (kPa). The GAB isotherm model is defined as:

$$q_{H_2O}(x) = c_m \frac{c_G K_{ads} x}{(1 - K_{ads} x)(1 + (c_G - 1) K_{ads} x)} \quad (5)$$

Table 1 GAB and Modified Toth isotherm fitted parameters [1, 5]

GAB isotherm	
c_m (mol/kg)	36.4800
c_G	0.1489
K_{ads}	0.5751
Modified Toth isotherm	
n_{s0} (mol/kg)	2.3800
b_0(kPa^{-1})	70.7000
t_0	0.4148
T_0 (K)	296.0000
ΔH_0 (kJ/mol)	-57.0470
α	-1.6060
χ	0.0000
γ (kg/mol)	0.0061
β (kg/mol)	28.907

2.1.2 Model development of fixed-bed adsorption

The key assumptions are as follows:

- 1-D model in axial direction
- Solid and gas phases are in thermal equilibrium
- Mass transfer rate described using the linear driving force (LDF) model
- Pressure drop was described using Ergun equation
- Heat and mass transfer coefficients were assumed to be constant
- Heat capacities and heat of adsorption were assumed to be constant
- Only CO_2 and H_2O are adsorbed by the solid adsorbent.

The mass and energy balance of the fixed-bed adsorption column is described as follows:

Mass Balance

$$\epsilon_t \frac{\partial c_i}{\partial t} + \frac{\partial}{\partial z}(v c_i) + (1 - \epsilon_b)\rho_p \frac{\partial q_i}{\partial t} = 0 \quad i = 1, \dots, n \quad (6)$$

LDF

$$\frac{\partial q_i}{\partial t} = k_i(q_i^* - q_i) \quad i = 1, \dots, n \quad (7)$$

Energy balance

$$(\epsilon_t C_g + \rho_b C_s + \rho_b C_{ads}) \frac{\partial T}{\partial t} - \epsilon_t \frac{\partial P}{\partial t} + v C_g \frac{\partial T}{\partial z} + \rho_b \sum_{j=1}^{n} \Delta H_j \frac{\partial q_j}{\partial t} = -\frac{4 h_L}{d_{in}} (T - T_w) \quad (8)$$

Wall Energy Balance

$$\frac{\partial T_w}{\partial t} = \frac{2}{C_w(r_{out}^2 - r_{in}^2)} \left[h_L r_{in}(T - T_w) - h_w r_{out}(T_w - T_{amb})\right] \quad (9)$$

2.2 Process simulation of the TVSA cycle

Figure 1 shows the steam-assisted temperature-vacuum swing adsorption (S-TVSA) cycle. The cycle consists of four steps: (1) adsorption, where a fan blows atmospheric air (mostly N_2 and O_2) at T_L and p_H through a packed column until CO_2 saturation; (2) blowdown, where the inlet is closed, and the column is evacuated to pressure p_L using a vacuum pump, releasing air back into the atmosphere; (3) heating, where heat at T_H is transferred via an external jacket, with the gaseous stream mainly

being air released back; and (4) desorption, with a steam purge providing additional temperature and vacuum swings[1].

Table 2 Adsorption Column specifications

Column	
Column length L (m)	0.0100
Inner diameter (m)	0.0800
Outer diameter (m)	0.0820
Bed density (kg/m^3)	55.4000
Particle size (m)	0.0075
Wall heat capacity (J/kg/K)	513.0000
Solid heat capacity (J/kg/K)	2070.0000
Fluid Heat transfer coefficient (W/m^2/K)	3.0000
Wall Heat transfer coefficient (W/m^2/K)	26.0000
CO_2 heat of adsorption (kJ/mol)	-57.0000
H_2O heat of adsorption (kJ/mol)	-49.0000
CO_2 Heat capacity (J/K/mol)	42.4600
H_2O Heat capacity (J/K/mol)	73.1000
CO_2 Mass transfer coefficient (1/s)	0.0002
H_2O Mass transfer coefficient (1/s)	0.0020

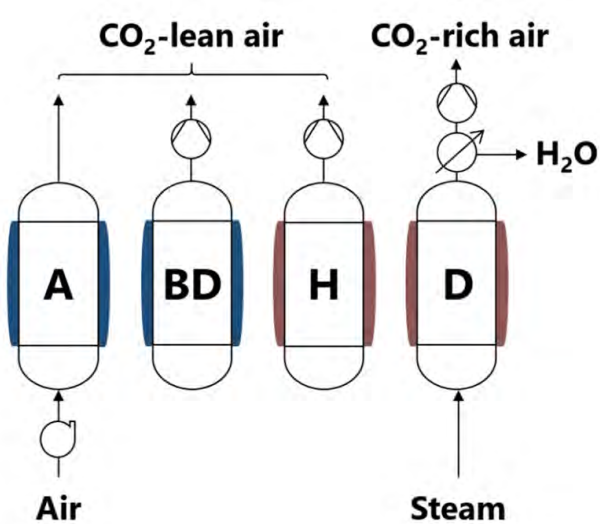

Figure 1. S-TVSA cycle [1]

The boundary conditions for each step is reported in Table 3.

Table 3 S-TVSA boundary conditions[1, 5]

Step	Time (s)	Boundary condition
Adsorption	13, 772	$p\mathrm{I}_{Z=L} = p_H$
		$u\mathrm{I}_{Z=0} = u_{air}$
		$T\mathrm{I}_{Z=0} = T_{air}$
		$y_i\mathrm{I}_{Z=0} = y_{i,air}$
Blowdown	30	$p\mathrm{I}_{Z=L} = p(t)$
		$u\mathrm{I}_{Z=0} = 0$
		$\frac{\partial T}{\partial z}\mathrm{I}_{Z=0} = 0$
		$\frac{\partial y_i}{\partial z}\mathrm{I}_{Z=0} = 0$
Heating	704	$p\mathrm{I}_{Z=L} = p_L$ (0.05 bar)
		$u\mathrm{I}_{Z=0} = 0$
		$\frac{\partial T}{\partial z}\mathrm{I}_{Z=0} = 0$
		$\frac{\partial y_i}{\partial z}\mathrm{I}_{Z=0} = 0$
Desorption	30, 000	$p\mathrm{I}_{Z=L} = p_L$
		$u\mathrm{I}_{Z=0} = u_{steam}$
		$T\mathrm{I}_{Z=0} = T_{steam}$
		$y_i\mathrm{I}_{Z=0} = y_{i,steam}$

2.3 Performance assessment of S-TVSA

2.3.1 Key Performance Indicators

The performance of the adsorption process was assessed using four key indicators: CO_2 Purity, CO_2 productivity, thermal energy use, and electrical energy use. The CO_2 Purity in the desorption step is defined as:

$$Pu_{CO_2}(\mathrm{mol\%}) = 100 \times \frac{n_{CO_2}^{prod}}{n_{CO_2}^{prod} + n_{H_2O}^{prod} + n_{N_2}^{prod}} \tag{10}$$

Where n_i^{prod} is the number of moles of component i produced from the desorption step. CO_2 productivity is defined as the capture rate per unit volume of the adsorption bed over the duration of a complete cycle. The CO_2 productivity is defined as:

$$Pr_{CO_2}(kg/m^3/day) = \frac{m_{CO_2}^{prod}}{V_{bed} \cdot t_{cycle}} \tag{11}$$

Where $m_{CO_2}^{prod}$ is the mass CO_2 produced during the desorption step (kg), V_{bed} is the adsorption bed volume (m^3), and t_{cycle} is the cycle time (s). Thermal energy consumption comprises the jacket heating required during the heating and desorption steps and the heat supplied

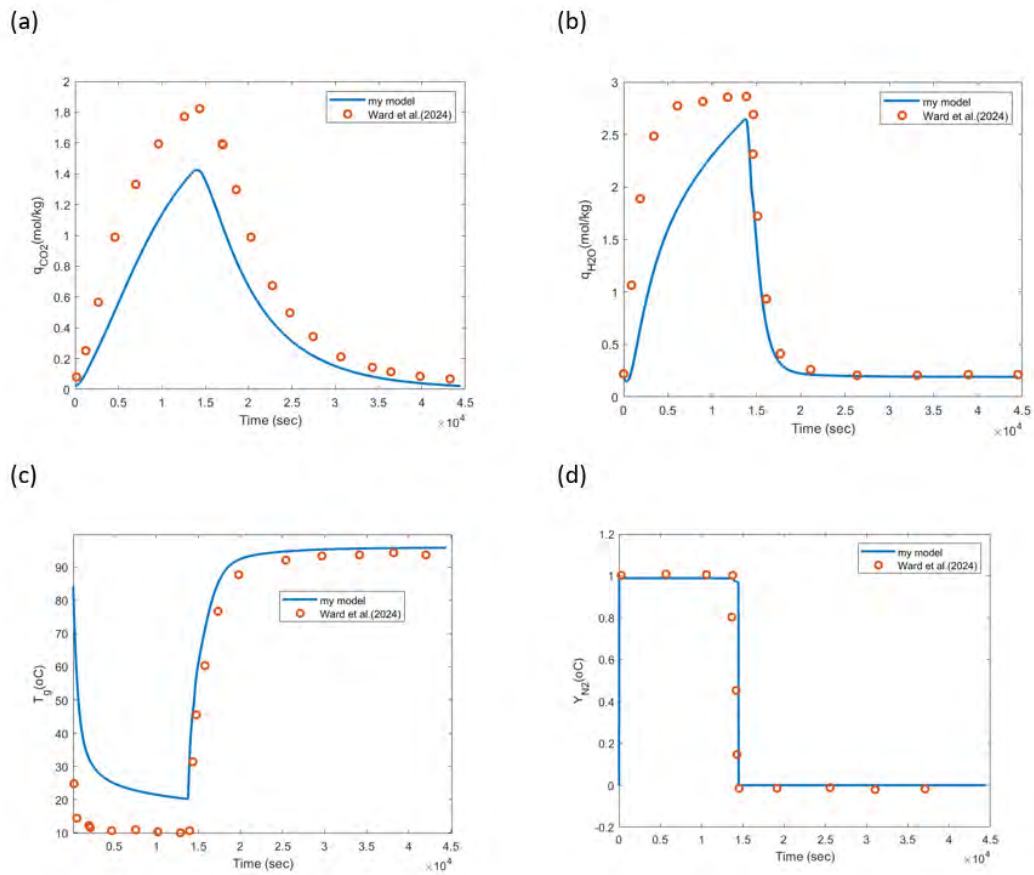

Figure 2: TVSA Cycle performance compared with Ward et al[5]. (a) CO_2 loading (b) H_2O loading (c) gas temperature (T_g)(d) N_2 composition

by the steam during the desorption step. The specific thermal consumption (Q_{therm}) is defined:

$$Q_{therm}(MJ/kg) = \frac{Q_{heat}^{Jacket} + Q_{des}^{jacket} + Q_{des}^{steam}}{m_{CO_2}^{prod}} \quad (12)$$

Where Q_i^{Jacket} is the heat (MJ) required in the jacket for step i, and Q_{des}^{steam} is the heat (MJ) supplied by the stream during the desorption step. The electrical energy consumption includes the energy required to operate the compressor during both adsorption and desorption steps, as well as the energy required for the vacuum pump to create and maintain a vacuum during the blowdown and heating stages. The electrical energy consumption is defined as:

$$E_{elec}(MJ/kg) = \frac{E_{ads} + E_{bd} + E_{heat} + E_{des}}{m_{CO_2}^{prod}} \quad (13)$$

Where E_i is the electrical energy (MJ) consumption of step i. To facilitate representing the total energy duty of the capture process as a single value for process optimisation, we use the concept of specific equivalent work to combine electrical and thermal energy consumption. The specific equivalent work is defined as:

$$W_{eq}(MJ/kg) = E_{elec} + \eta_{turb}\left(1 - \frac{T_L}{T_H}\right) \quad (14)$$

Where η_{turb} is the turbine isentropic efficiency of a hypothetical turbine undertaking the conversion between electrical and thermal work. η_{turb} was assumed to be 75%.

2.3.2 Process optimisation

The solid-based DAC performance was optimised by maximising system productivity and minimising energy use. CO_2 productivity quantifies the system volume required to achieve a specified CO_2 capture rate, thereby affecting the land footprint. Since DAC must be deployed globally for a significant climate impact, high productivity minimises land use. The aim is to minimise energy consumption in a steam-assisted DAC process. The multi-objective optimisation was formulated as follows:

$$\min_{\theta}[-Pr_{CO_2} + \phi, W_{eq} + \phi] \quad (15)$$

$$s.t\ \ \theta_L \leq \theta \leq \theta_U$$

$$\phi = 0.25 \times [\max(0, 95 - Pu_{CO_2})]^2$$

Where θ is the vector of the operating variables, θ_L and θ_U are the lower and upper operating variable limits (see Table 4). A constraint was imposed on the CO_2 purity, requiring a minimum of 95%. ϕ is the penalty function.

Table 4 Lower and upper limits

Parameter	Lower Limit	Upper Limit
t_{ads}(s)	10, 000.000	20, 000.000
t_{heat}(s)	500.000	1, 500.000
t_{des}(s)	20, 000.000	40, 000.000
p_L(bar)	0.050	0.500
T_{des}(°C)	80.000	100.000
v_s(m/s)	0.001	0.005
v_f(m/s)	0.005	0.010

We deployed the NSGA-II algorithm from MATLAB's Global Optimization Toolbox as the gamultiobj function, running for 100 generations with a population of 140. It was initialized with 140 operating points and their key performance indicators, determined by quasi-random Sobol sampling within bounds in Table 4.

3. RESULTS

3.1 TVSA Cycle performance

The behaviour of the S-TVSA cycle was analysed through CO_2 and H_2O loading profiles in comparision with Ward et al.[5]. Figure 2(a) shows CO_2 loading increasing during the adsorption phase (13, 772 s) at 20 °C and 1 bar.

Figure 2(b) presents H_2O loading, indicating sorbent stability under humid conditions, which is crucial for DAC applications. Thermal dynamics are shown in Figure 2(c), with the gas temperature (T_g) rising to approximately 95 °C during heating and desorption, driven by steam injection and external heating. Figure 2(d) tracks N_2, which evacuates rapidly during blowdown and desorption, maintaining high CO_2 purity.

3.2 Energy consumption

The assessment (Figure 3) shows thermal energy (Q_thermal) as the dominant component of total energy, accounting for jacket heating and steam regeneration (7.04 MJ/kg). Electrical energy (0.37 MJ/kg) primarily powers blowers and vacuum pumps across various phases. The Specific Equivalent Work (W_{eq}) determined is 7.41 MJ/kg.

3.3 Process optimisation

The optimisation results, generated using the NSGA-II algorithm with 100 generations and a population of 140, are summarised in the Pareto Curve (Figure 4). The results demonstrate a clear trade-off between CO_2 productivity and energy consumption. Initial points derived via Sobol sequence sampling within the bounds helped identify an optimal operating window that balances throughput with energy efficiency. The optimal solution yielded a productivity range of 32.6-33.4kg/m^3/day. The corresponding energy requirement for the optimal points ranges from 7.0 to 7.6 MJ/kg.

The multi-objective optimisation reveals a key link between the S-TVSA cycle's physical performance and the economic viability of solid-based DAC. The Pareto front marks the 'best possible' performance. Increasing CO_2 productivity raises specific work (W_{eq}) and reduces adsorbent bed volume (V_{bed}), lowering capital costs and land requirements. The 33.4 kg/m^3/day point boosts productivity and cuts CAPEX; the 7.0 MJ/kg point minimises OPEX by reducing thermal energy ($Q_{thermal}$).

Thermal energy far exceeds electrical energy, making the heat source critical. Waste heat or low-cost geothermal energy for the heating and desorption steps could greatly lower OPEX, favouring high-productivity points.

4. CONCLUSION

This study developed and validated a robust fixed-bed model and an S-TVSA cycle for solid-sorbent-based DAC. The process optimisation with NSGA-II revealed that thermal energy is the main factor driving the energy penalty. Key findings include:

- APDES-NFC demonstrates reliable co-adsorption of CO_2 and H_2O, making it suitable for environments with varying humidity.
- $Q_{thermal}$ is the dominant factor in W_{eq}, necessitating advanced regeneration strategies or materials with higher working capacities.
- Optimal solutions achieved CO_2 productivity between 32.6 and 33.4 kg/m^3/day. The corresponding specific equivalent work is between 7.0 to 7.6 MJ/kg. This provides a pathway to balancing the trade-off between plant size (CAPEX) and energy demand (OPEX).
- Findings indicate that solid-sorbent DAC's economic feasibility largely relies on low-carbon, low-cost thermal energy for desorption.

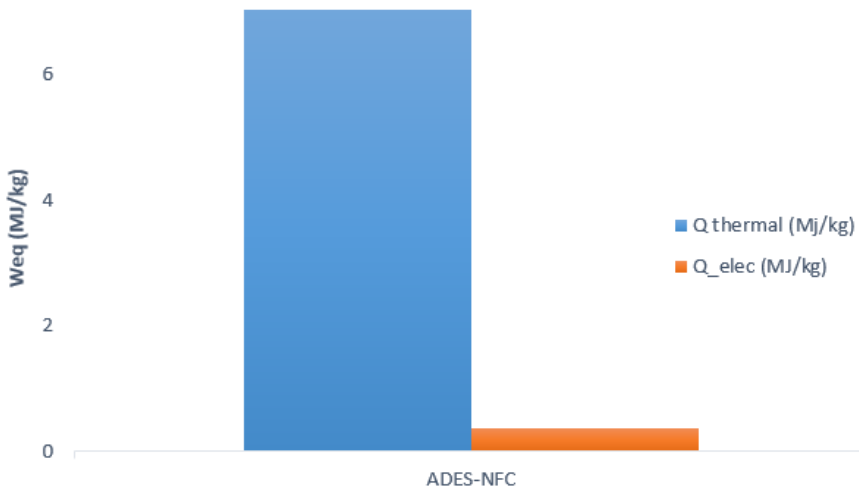

Figure 3. Specific equivalent work

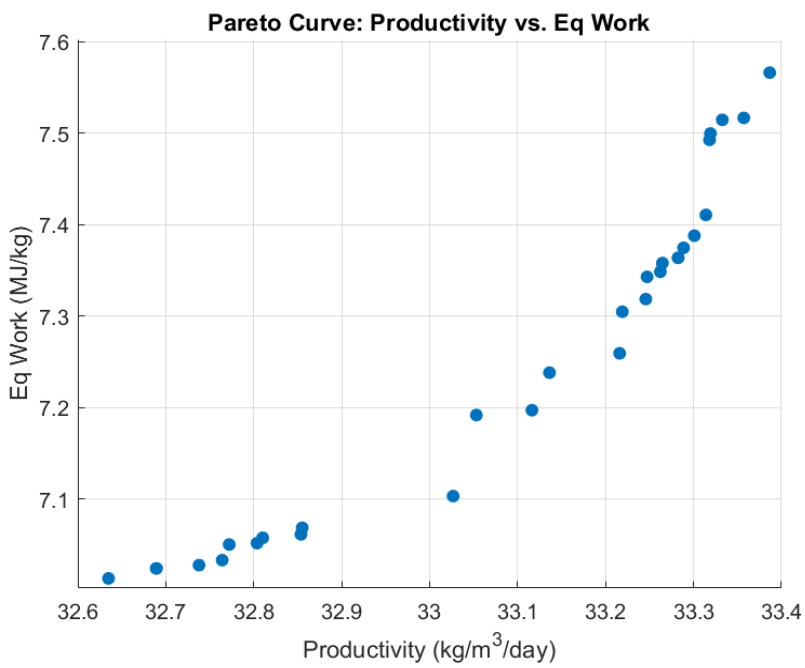

Figure 4. Pareto Curve

Future work will focus on a detailed parametric analysis and the impact of water-enhanced amine-functionalised adsorbents on the DAC performance.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the financial support of EU-RISE project CATALYSE (Ref: Grant Agreement No: 101183092).

AUTHOR IDENTIFIERS

Author ORCIDs:
Toluleke E Akinola: 0000-0003-4379-9425
Meihong Wang: 0000-0001-9752-270X

REFERENCES

1. Stampi-Bombelli V, van der Spek M, Mazzotti M. Analysis of direct capture of $${\hbox {CO}}_{2}$$ from ambient air via steam-assisted temperature–vacuum swing adsorption. Adsorption 26:1183-1197 (2020). https://doi.org/10.1007/s10450-020-00249-w
2. NOAA Research. Trends in Atmospheric Carbon Dioxide (CO_2) 2025. https://gml.noaa.gov/ccgg/trends/weekly.html (accessed October 24, 2025).
3. IEA. Net Zero by 2050. IEA, Paris 2021. https://www.iea.org/reports/net-zero-by-2050, Licence: CC BY 4.0 (accessed October 24, 2024).
4. IEA. Direct Air Capture 2022. IEA, Paris 2022. https://www.iea.org/reports/direct-air-capture-2022, Licence: CC BY 4.0.
5. Ward A, Papathanasiou MM, Pini R. The impact of design and operational parameters on the optimal performance of direct air capture units using solid sorbents. Adsorption 30:1829-1848 (2024). https://doi.org/10.1007/s10450-024-00526-y

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Uncertainty Quantification of Stochastic Gene Expression

Francisca Pizarro Galleguillos[abc], Satyajeet S. Bhonsale[a], and Jan F.M. Van Impe[a]*

[a] KU Leuven, Department of Chemical Engineering, BioTeC+, Gent, Belgium
[b] Universitat Politècnica de València, Automatics and Industrial Informatics Research Institute (AI2), Valencia, Spain
[c] Institute for Integrative Systems Biology (I2SysBio), Spanish National Research Council (CSIC), Computational Synthetic Biology Group, Valencia, Spain
* Corresponding Author: jan.vanimpe@kuleuven.be.

ABSTRACT

Stochastic gene regulatory networks exhibit complex dynamics that require efficient methods for parameter inference and uncertainty quantification. In this work, we propose a surrogate modelling framework that combines a partial integro-differential equation (PIDE) formulation with polynomial chaos expansions (PCE) to efficiently approximate the stochastic dynamics of gene expression models under parametric uncertainty. The approach represents the time evolution of low-order statistical moments as polynomial functions of uncertain kinetic parameters, enabling fast evaluations and tractable inference. The method is demonstrated on a self-regulating gene network, achieving accurate parameter estimation and a reduction of approximately two orders of magnitude in computational cost compared to direct PIDE-based optimisation.

Keywords: Surrogate Model, Modelling and Simulations, Optimization

INTRODUCTION

Stochasticity is an inherent feature of cellular processes, playing a crucial role in various biological phenomena. A prominent example is gene expression, where small numbers of DNA and regulatory molecules lead to pronounced stochastic effects [1]. These fluctuations contribute to the significant cell-to-cell variability observed even in genetically identical cell populations, influencing key processes such as differentiation, phenotypic switching, adaptation, and evolution [1-3]. Advanced single-cell microscopy techniques, combined with mathematical models for single-gene expression, have led to essential insights into elucidating the sources of intrinsic noise in cells [4]. At the same time, they have highlighted new challenges for quantitative modelling, particularly in statistical inference and uncertainty quantification.

Stochastic gene expression is commonly modelled as a reaction network governed by the chemical master equation (CME) [5]. Despite its simple formulation, the CME is generally intractable and exact analytical solutions exist only for a few simple systems. On the other hand, the stochastic simulation algorithm (SSA) provides a numerical approximation of the CME solution. However, stochastic simulations are computationally demanding, making systematic analysis challenging. As an alternative to computing full probability distributions using the CME or its approximations, such systems are often analysed through low-order statistical moments of mRNA or protein abundance [6].

Gene regulatory network (GRN) models are necessarily approximations of biological reality and are therefore affected by multiple sources of uncertainty. In many practical applications, model parameters are only approximately known, as direct experimental measurements of kinetic reaction rates are difficult to obtain. Moreover, even when reliable estimates are available, *in vivo* parameters may be influenced by a multitude of cellular and environmental factors, leading to significant parametric variability [7]. Furthermore, incomplete knowledge of network topology and reaction mechanisms can introduce structural uncertainty, while the intrinsic randomness of molecular interactions gives rise to aleatory variability. If not properly accounted for, these effects can substantially impact model predictions and lead to misleading interpretations in biological applications.

In this work, we combine an efficient partial integro-differential equation (PIDE) framework for modelling stochastic gene regulatory networks with polynomial chaos expansion (PCE) for uncertainty quantification under

parametric uncertainty. The PCE represents the statistical moments of the system as polynomial functions of the model parameters, enabling fast moment evaluation. We demonstrate the effectiveness of the PCE-based approach on a positive autoregulation model by estimating all kinetic parameters of the network. Compared to standard stochastic simulation–based methods, the proposed framework achieves up to a 100-fold reduction in computational cost, making it particularly suitable for inference and uncertainty quantification analyses.

METHODS

PIDE Model

We consider a gene regulatory network comprising n genes $\boldsymbol{G} = \{\mathrm{DNA}_1, \dots, \mathrm{DNA}_n\}$, each responsible for the synthesis of a distinct protein type $\boldsymbol{X} = \{X_1, \dots, X_n\}$ through their respective messenger RNAs $\boldsymbol{M} = \{\mathrm{mRNA}_1, \dots, \mathrm{mRNA}_n\}$, following the central dogma. The number of molecules of mRNA and proteins is represented by vectors, $\boldsymbol{m}, \boldsymbol{x} \in R^n$, where $\boldsymbol{m} = (m_1, \dots, m_n)$ and $\boldsymbol{x} = (x_1, \dots, x_n)$, respectively. A schematic representation of the gene network is shown in Figure 1.

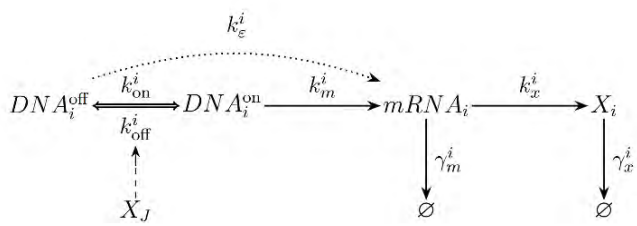

Figure 1. Schematic representation of the gene expression model.

Each gene switches between an active DNA_i^{on} and an inactive DNA_i^{off} state. In the active state, mRNA is transcribed at a rate k_m^i and translated into a protein at a rate k_x^i. A gene is activated at a rate k_{on}^i and inactivates at a rate k_{off}^i, with an additional leakage transcription rate k_ε^i. mRNA and protein degrade at rates γ_m^i and γ_x^i, respectively. Under the assumption of protein-burst condition, the time evolution of the protein probability density function can be described by the following PIDE [4]:

$$\frac{\partial p}{\partial t}(t, \boldsymbol{x}) = \sum_{i=1}^{n} \frac{\partial}{\partial x_i}\left[\gamma_x^i x_i p(t, \boldsymbol{x})\right] + \sum_{i=1}^{n}\left(k_m^i \int_0^{x_i} \beta_i(x_i - y_i) c_i(\boldsymbol{y}_i) p(t, \boldsymbol{y}_i) \mathrm{d}\, y_i\right) \tag{2}$$

$$p(0, \boldsymbol{x}) = p_0(\boldsymbol{x}).$$

The burst condition is represented by β_i, with

$$\beta_i(x_i - y_i) = \omega_i(x_i - y_i) - \delta(x_i - y_i), \tag{3}$$

where $\boldsymbol{y}_i$ is obtained from $\mathbf{x}$ by replacing its i-th component by y_i. Additionally, burst size is modelled by an exponential distribution, such that

$$\omega_i(x_i - y_i) = \frac{1}{b_i} \exp\left(-\frac{(x_i - y_i)}{b_i}\right), \tag{4}$$

where $b_i = \frac{k^i{}_x}{\gamma^i{}_m}$. The initial condition $p_0(x)$ is a probability density function assumed to be normally distributed. Here, the function $c_i(x)$ represents the regulatory control of a gene i by other proteins in the network, and it is typically modelled using Hill functions

$$\rho_{ij(x_i)} = \frac{x_j^{H_{ij}}}{x_j^{H_{ij}} + K_{ij}^{H_{ij}}}, \tag{5}$$

where H_{ij} is the Hill coefficient indicating the cooperativity in the regulation, and K_{ij} is the corresponding Hill constant. In the case of regulation by a single protein, the transcription rate is proportional to the fraction of active promoter, expressed as:

$$c_i(x_j) = \left[1 - \rho_i(x_j)\right] + \rho_i(x_j)\varepsilon_i. \tag{6}$$

Here ε_i is the leakage constant from the inactive promoter. The effect of multiple proteins regulating a gene can be modelled by combining individual Hill functions using logical operations [4].

The PIDE is solved numerically using SELANSI, which computes the solution via an efficient and scalable semi-Lagrangian method, providing the temporal evolution of the protein probability density function [5]. The PIDE approach provides an accurate and computationally efficient approximation of the chemical master equation (CME) solution compared to the standard Gillespie stochastic simulation algorithm. A computational mesh consisting of 5000 time points and 1200 spatial nodes is employed.

Polynomial Chaos Expansion

PCE is a spectral method for uncertainty quantification that represents random variables as expansions in orthogonal polynomial bases for which stochastic information is available [6]. The PCE expansion of a random variable u can be described as:

$$u(\theta) = \sum_{j=0}^{L} u_j \boldsymbol{\Phi}_j(\theta), \tag{7}$$

where L is the total number of polynomial terms up to a degree $d, \boldsymbol{\Phi}_j(\theta)$ are multivariate orthogonal polynomials in the random variable θ, u_j the PCE coefficients, and j is a multi-index. The choice of the polynomial family follows the Askey scheme and depends on the probability distributions of the input parameters. Although classical PCE assumes standard distributions, the approach can be extended to arbitrary input measures. The PCE coefficients

are computed using a non-intrusive regression-based approach. The underlying model is evaluated at a finite set of sampling points in the parameter space, and the coefficients are obtained by solving an ordinary least-squares problem using the corresponding model evaluations.

Moment-based PCE

Since tracking the full probability distribution is computationally expensive, a surrogate model tracking the moments of the gene circuit is constructed. To this end, the first k raw moments are computed by integrating the probability density obtained from the PIDE formulation at selected time instants

$$M_i = \int x^k p(x) dx. \tag{8}$$

Each moment $M_i(t, \theta)$ is then represented as a polynomial chaos expansion in the parametric space. Specifically, the i-th moment is approximated as

$$M_i(t, \theta) \approx \sum_{j=0}^{L} m_{i,j}(t)\, \mathbf{\Phi}_j(\theta), \tag{9}$$

Where $m_{i,j}$ are the corresponding deterministic PCE coefficients.

RESULTS

Moment-based PCE

In this case study, we consider a single-gene circuit with self-regulation, in which feedback arises from the binding of the protein encoded by the gene to its own regulatory regions. This is

$$\rho_{1(x_1)} = \frac{x_1^{H_{11}}}{x_1^{H_{11}} + K_1^{H_{11}}}, \tag{10}$$

and H_{11} has a negative sign because positive self-regulation is considered. The network dynamics are governed by the transcription rate k_m, the translation rate k_x, and the mRNA and protein degradation rates γ_m and γ_x. However, since the mean burst size is given by $b_1 = k_x/\gamma_m, k_x$ and γ_m are not independently identifiable. Consequently, only three independent network parameters can be estimated, namely k_m, k_x and γ_x. These parameters are assumed to be uncertain and are modelled as independent normally distributed random variables with $N\sim(\mu, (0.1\mu)^2)$. Their mean values used in the simulations are listed in Table 1, and the PIDE solutions corresponding to the nominal values are shown in Figure 2.

Simulation data are normalised to the interval $[0, 1]$ to facilitate the PCE construction by avoiding large differences in numerical magnitude between data points. The corresponding probability density functions are normalised consistently. The first five raw moments of protein expression are computed from the distributions obtained via the PIDE stochastic model and used as training data for the surrogate.

Table 1: Kinetic parameters used in the PIDE simulation.

Parameter	Description	Value
k_m	Translation rate	8
k_x	Transcription rate	400
γ_m	mRNA degradation rate	60
γ_x	Protein degradation rate	1
H_{11}	Hill coefficient	4
K_1	Hill constant	20
ε_1	Leakage constant	0.15

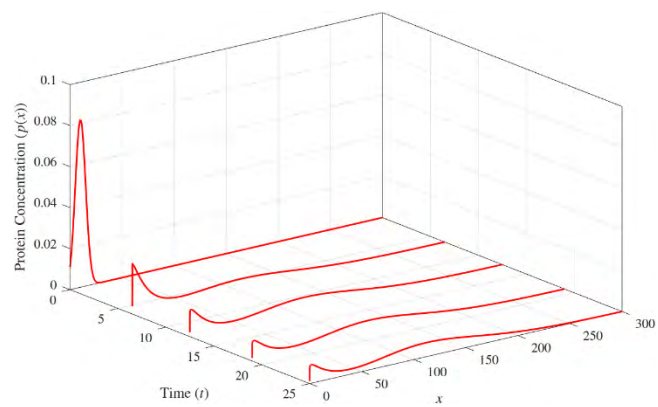

Figure 2: Time evolution of PIDE simulation with nominal values. Bimodal behaviour characteristic of positive autoregulation is observed.

A second-order PCE is employed, with an experimental design consisting of 300 model evaluations and their corresponding moments. Based on the assumed distributions of the uncertain parameters, the polynomial basis is chosen as a set of normalised Hermite polynomials. The surrogate is fitted to snapshot data collected at 10 discrete time steps, and the PCE coefficients are computed using the least angle regression (LAR) algorithm, as implemented in UQLab [8]. The choice of the polynomial degree, the size of the experimental design, and the regression algorithm may vary depending on the level of accuracy required. In practice, these choices involve a trade-off between computational cost and surrogate performance. In the present case, the generation of the training samples represents the most computationally expensive step and requires approximately 20 minutes of CPU time. Once trained, however, the PCE surrogate provides an efficient representation of the model response over the entire parameter space, enabling rapid evaluations. Moreover, key statistical quantities such as the mean, variance, and Sobol sensitivity indices can be obtained directly from the PCE coefficients at negligible additional cost.

Parity plots shown in Figure 3 demonstrate strong

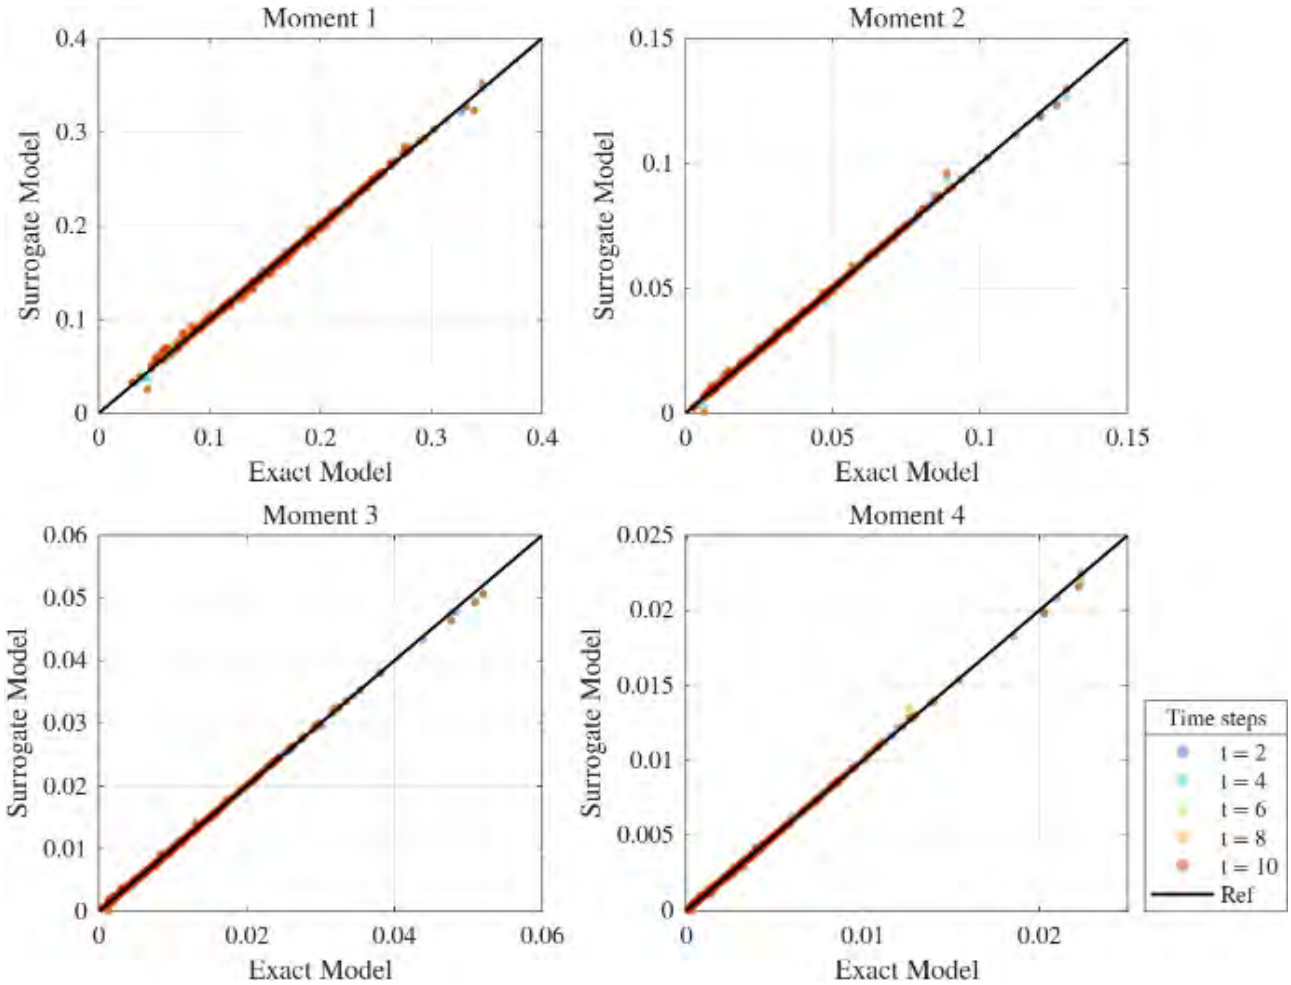

Figure 3. Parity plots comparing the first four moments predicted by the PCE surrogate (dots) against the validation dataset (continuous line). A second-order PCE constructed from 300 training samples is used. The zeroth moment is omitted for clarity, and selected data points at representative time instants are shown.

agreement between the moments predicted by the PCE surrogate and those obtained from the validation dataset across all considered time points. In particular, data points closely follow the identity line for the first four moments, indicating that the surrogate accurately captures both the magnitude and temporal evolution of the stochastic moment dynamics.

To quantitatively assess the accuracy of the PCE surrogate, leave-one-out (LOO) cross-validation errors, modified LOO errors, and root-mean-square errors (RMSE) were evaluated at all time points. Since the zeroth moment is identically equal to 1, it was excluded from the error analysis. Considering only moments of order greater than zero, the surrogate errors remained low throughout the simulation, with the global RMSE staying below 1.5×10^{-3}. The largest contribution to the total error arose from the first moment, whose RMSE increased at early times and then stabilised at approximately 1.5×10^{-3}, while the remaining moments were predicted more accurately, with RMSE values below 1.0×10^{-4}. Small deviations are observed only in a limited number of cases, with the first moment contributing most to the total error due to its comparatively larger variability. Nevertheless, the remaining moments are reproduced with very high accuracy, and the overall results confirm that the surrogate provides an accurate approximation of the moment dynamics over the entire time horizon. Further improvements may be achieved through alternative sampling strategies, different regression techniques for computing the PCE coefficients, or larger experimental designs.

Moment-based approaches provide an efficient alternative for characterising stochastic dynamics by representing the system through the evolution of its moments [9]. By focusing on their dynamics, these methods enable the quantification of fluctuations over time while substantially reducing computational cost compared to full distribution-based analyses, as will be demonstrated

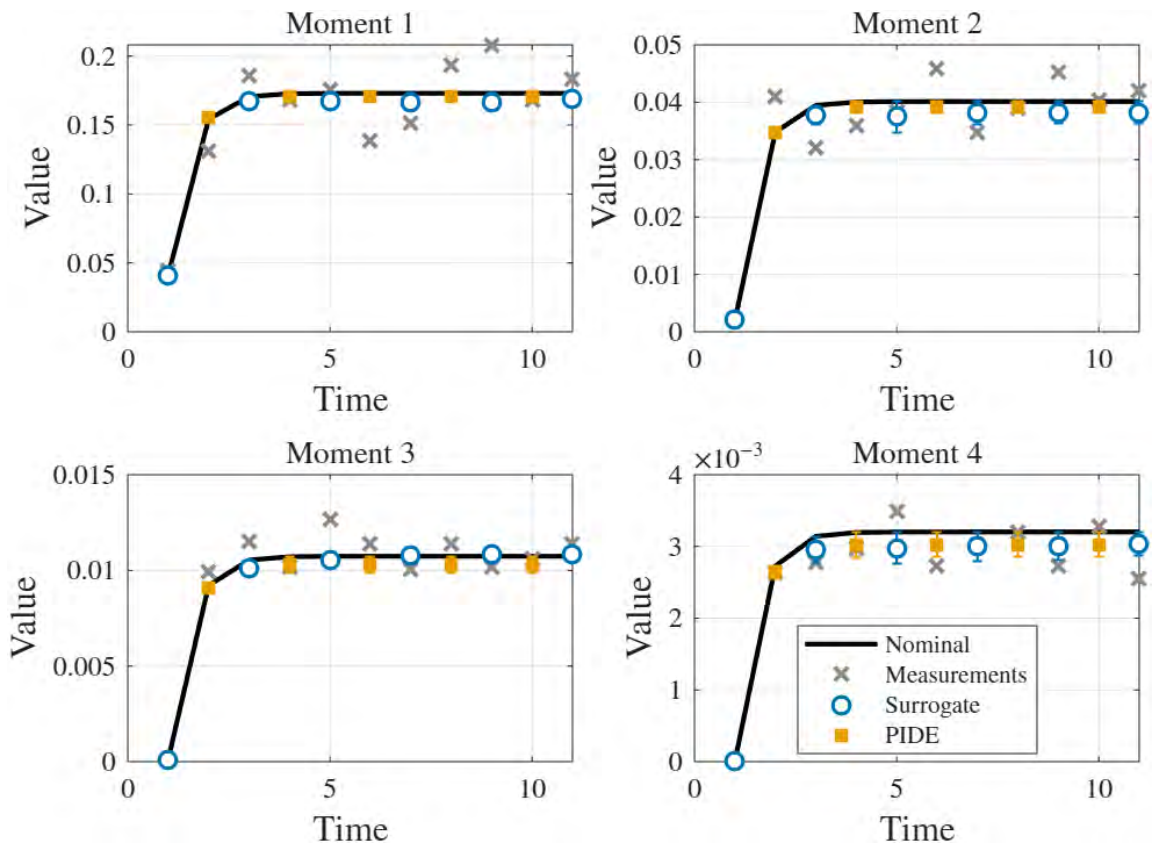

Figure 4: Nominal values, measurements, and estimated moments obtained using the PIDE and PCE models. The zeroth moment is omitted for clarity, and 10 data snapshots are considered.

in the following section.

Parameter Inference

Following the validation of the PCE surrogate, it is employed to estimate the model parameters using a multi-start optimisation strategy. Randomised initial guesses within predefined parameter bounds are used in a constrained optimisation algorithm to minimise a weighted least-squares objective. Noisy measurements are generated from the moment dynamics, assuming 10% multiplicative Gaussian noise. The PCE surrogate enables rapid evaluation of the forward model, making the multi-start procedure computationally feasible while jointly estimating all kinetic parameters of the regulatory network. The parameter set yielding the lowest objective value across all optimisation runs is selected as the final estimate.

Under 10% relative noise, the PCE surrogate accurately captures the underlying moment dynamics. The estimated parameters reproduce the temporal evolution of the measured moments and remain in close agreement with the reference PIDE solution, as shown in Figure 4.

Table 2 summarises the estimated parameters and their relative errors obtained with the PIDE and PCE models, while the resulting probability density functions simulated using the optimised parameters are shown in Figure 5.

Table 2: Estimated parameters and relative errors from the PIDE and PCE model.

Param.	PIDE	PCE	Rel. Error PIDE	Rel. Error PCE
k_m	8.92	8.34	11.5%	4.2%
k_x	385.4	367.3	3.6%	8.1%
γ_x	1.086	1.14	8.6%	14.0%

While the PIDE approach yields accurate parameter estimates with relative errors below 12%, the surrogate-based inference exhibits moderate deviations for some parameters, particularly those associated with stronger coupling and reduced practical identifiability. Nevertheless, the parameter estimates obtained with the PCE surrogate remain within acceptable accuracy and successfully reproduce the temporal evolution of the observed moments. Moreover, the probability density functions reconstructed using the parameters inferred from the PCE surrogate closely match those obtained with the reference PIDE model and the nominal parameter set despite parameter deviations.

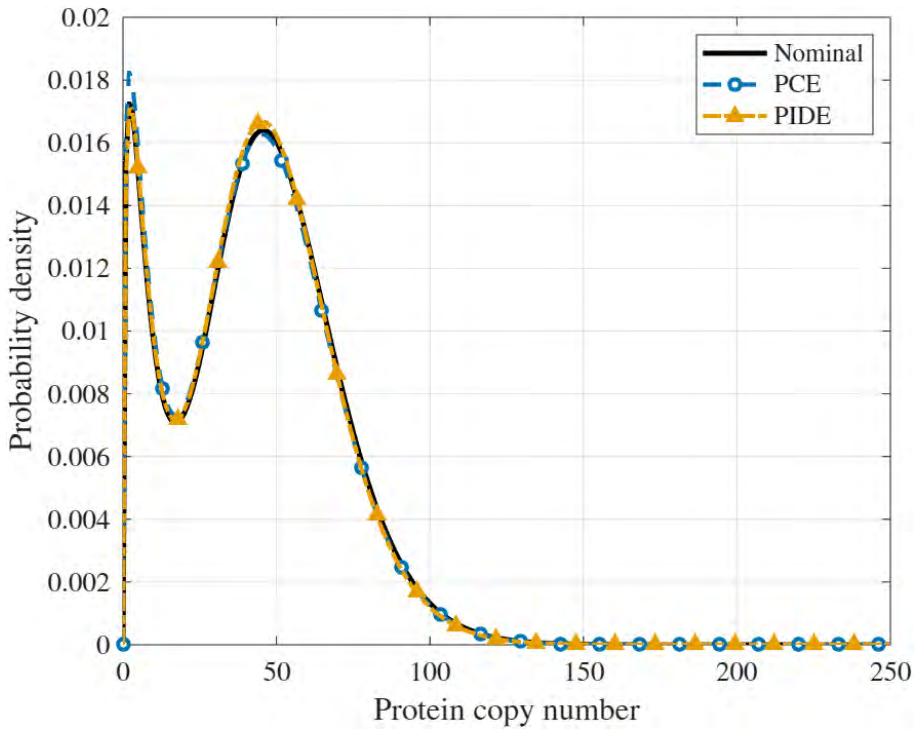

Figure 5: Probability density functions obtained using the optimised parameters for the PIDE and PCE models. The PCE-based estimates closely match the reference PIDE solution and preserve the multimodal structure of the distribution.

Beyond accuracy, the use of the PCE surrogate provides substantial computational advantages. For a multi-start optimisation with 20 initial guesses and parallel execution, the total optimisation time using the PCE surrogate was 1.8 minutes, compared to 180.9 minutes when relying on the PIDE model, corresponding to an approximately 100-fold reduction in computational cost. Once constructed, the surrogate can be readily employed for both forward and inverse uncertainty quantification tasks. Furthermore, the non-intrusive nature of the surrogate modelling approach facilitates the inclusion of multiple uncertain parameters, making the proposed framework well-suited for efficient inference and uncertainty quantification in stochastic gene regulatory networks.

CONCLUSIONS AND FUTURE RESEARCH

In this work, we develop a surrogate modelling framework for uncertainty quantification in stochastic gene regulatory networks. The proposed approach combines PIDE-based simulations for computing probability density functions with Polynomial Chaos Expansions (PCE) to construct an efficient surrogate model. The central idea is to represent the time evolution of statistical moments as functions of uncertain input parameters.

The effectiveness of the proposed method is demonstrated using a one-dimensional self-regulating gene network as a case study. This system is selected for its conceptual simplicity while still exhibiting bimodal behaviour depending on the uncertain input parameters. Despite its low dimensionality, the example highlights that accurate joint inference of all kinetic parameters can be achieved with a substantial reduction in computational cost. These computational gains are expected to become even more pronounced for more complex gene regulatory networks involving multiple interacting genes.

Future research directions include extending the proposed framework to more advanced uncertainty quantification tasks, such as Bayesian parameter inference, and to larger gene regulatory networks with multiple uncertain parameters. Additional reductions in training cost may be achieved through more efficient sampling strategies and improved experimental design.

ACKNOWLEDGEMENTS

This work was supported by the European Union's Horizon 2022 research and innovation program "FAIROmics" under the Marie Skłodowska-Curie grant agreement No. 101120449. SB was supported by the Research Foundation Flanders (FWO) under grant No.

G0A9H26N.

AUTHOR IDENTIFIERS

Author ORCIDs:
FP: 0009-0002-0291-0772
SB: 0000-0001-9734-4122
JVI: 0000-0002-5904-1638

REFERENCES

1. Elowitz MB, Levine AJ, Siggia ED, Swain PS. Stochastic gene expression in a single cell. Science 297:1183-1186 (2002). https://doi.org/10.1126/science.1070919
2. Brophy JAN, Voigt CA. Principles of genetic circuit design. Nat Methods 11:508-520 (2014). https://doi.org/10.1038/nmeth.2926
3. Raj A, van Oudenaarden A. Nature, nurture, or chance: stochastic gene expression and its consequences. Cell 135:216-226 (2008). https://doi.org/10.1016/j.cell.2008.09.050
4. Pájaro M, Alonso AA, Otero-Muras I, Vázquez C. Stochastic modeling and numerical simulation of gene regulatory networks with protein bursting. Journal of Theoretical Biology 421:51-70 (2017). https://doi.org/10.1016/j.jtbi.2017.03.017
5. Raharinirina NA, Peppert F, von Kleist M, Schütte C, Sunkara V. Inferring gene regulatory networks from single-cell rna-seq temporal snapshot data requires higher-order moments. Patterns 2:100332 (2021). https://doi.org/10.1016/j.patter.2021.100332
6. Ghanem R, Higdon D, Owhadi H. Introduction to uncertainty quantification. Handbook of Uncertainty Quantification :3-6 (2017). https://doi.org/10.1007/978-3-319-12385-1_1
7. Schnoerr D, Sanguinetti G, Grima R. Approximation and inference methods for stochastic biochemical kinetics—a tutorial review. J. Phys. A: Math. Theor. 50:093001 (2017). https://doi.org/10.1088/1751-8121/aa54d9
8. Marelli S, Sudret B. Uqlab: a framework for uncertainty quantification in matlab. Vulnerability, Uncertainty, and Risk :2554-2563 (2014). https://doi.org/10.1061/9780784413609.257
9. Eduardo Hernández-García M, Gómez-Schiavon M, Velázquez-Castro J. Stability analysis under intrinsic fluctuations: a second-moment perspective of gene regulatory networks. Phys. Biol. 22:066001 (2025). https://doi.org/10.1088/1478-3975/ae0b22

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Practical Identifiability and Optimal Experiment Design for Hybrid Cybernetic Models: An *E.Coli* Case Study

Stylianos Floros, Satyajeet S. Bhonsale, Simen Akkermans, Jan F.M. Van Impe*

BioTeC+ Chemical & Biochemical Process Technology & Control, KU Leuven, Ghent, Belgium
* Corresponding Author: jan.vanimpe@kuleuven.be.

ABSTRACT

Developing predictive, high-resolution microbial models that retain mechanistic insight remains a central challenge in biochemical engineering. This paper addresses the challenge of preserving metabolic information while ensuring model accuracy and proper statistical definition. It employs the Hybrid Cybernetic Modelling (HCM) approach to integrate metabolic regulation with genome-scale information, and dynamically predict *E.Coli* phenotypes. The main aim of this work is to explore HCM's parameter identifiability to advance its accuracy and robustness for a limited set of data. To bypass the computational burden of computing the elementary flux modes, the opt-yield Flux Balance Analysis (opt-yield FBA) is employed to identify a physiologically relevant set of yield-maximising metabolic pathways. Metabolic Yield Analysis (MYA) then reduces this to four key pathways which capture 99% of the original metabolic yield space. The cybernetic model is formulated where "artificial enzymes" are allocated among these key competing pathways, based on a return-on-investment criterion. Initial parameter estimates are derived through non-linear least squares regression for an *in silico* data set of biomass, glucose and acetate concentration. Practical identifiability analysis is employed to explore parameter estimates. Specifically, profile likelihood & sensitivity analysis results lead to model reduction by eliminating the combined acetate-biomass producing pathway. To further improve parameter uncertainty, Optimal Experiment Design (OED) is used to obtain an additional informative experiment. Finally, the calibrated model is validated against a new artificial dataset.

Keywords: Metabolic Engineering, Cybernetic Model, Practical Identifiability, Optimal Experiment Design

INTRODUCTION

Mathematical modelling of intracellular processes is gaining popularity towards developing metabolic and biochemical engineering applications [1]. Recent breakthroughs suggest the formulation of multi-scale modelling approaches that avail of mechanistic information for intracellular metabolism, towards accurately describing changes in microbial communities.

Typically, such models adopt a "grey-box" structure. Specifically, population-based models that come with a form of Ordinary Differential Equations (ODEs) are directly informed by mechanistic models (e.g. Genome-scale Models (GEMs)) for intracellular metabolism [2]. When parametrised, they are calibrated based on limited experimental data to offer informative predictions for bioprocesses.

However, the systematic employment of such models is often hindered by parameter identifiability issues. Many reasons lead to excess parameter uncertainty, which hinders the causal link between intracellular metabolism and extracellular phenotype [3]. As a result, this poses a major threat to the fundamental understanding of complex biological phenomena.

This paper aims to advance the understanding of parameter identifiability in multi-scale models. It employs the Hybrid Cybernetic Model (HCM) to capture temporal changes of biomass, acetate and glucose concentrations within an E.Coli case study [4]. It performs practical identifiability and sensitivity analysis to further explore the estimated parameters' relationship with model inputs/outputs. The study further proceeds with model reduction and Optimal Experiment Design (OED), towards expanding the knowledge on HCM's responses.

MATERIALS & METHODS

Case Study Design

This study chooses *Escherichia Coli str. K-12* as a model organism for *in silico* experimentation. First, an artificial data set is generated based on *E.Coli* metabolism. The *in silico* environment resembles a bio-reactor that operates under batch-ideal conditions. This assumption serves as the basis for all subsequent modelling and analysis.

Dynamic - Flux Balance Analysis

To simulate the temporal changes in cellular biomass production, metabolite consumption and secretion, the *e_coli_core* GEM for *E.Coli K-12* is selected [5].

Dynamic-Flux Balance Analysis (d-FBA) further aids in generating the artificial data set. The initial conditions of the *in silico* experiment are: 0.001 $\mathrm{gDW/L}$ biomass, 10 $\mathrm{mmol/L}$ glucose and 0.35 $\mathrm{mmol/L}$ acetate. To impose glucose availability as an *in silico* constraint for cellular uptake of glucose, a Michaelis-Menten expression is used:

$$V_{glc} = -V_{max} \bullet \frac{S}{K_s + S} \tag{1}$$

where V_{glc}, V_{max} are the actual and maximum specific uptake rate ($\mathrm{mmol/gDW*h}$), equal to 10, and K_s is the half-maximum saturation constant ($\mathrm{mmol/L}$), equal to 5. Oxygen uptake was further constrained (-2, +1000) $\mathrm{mmol/gDW*h}$ to enable secretion of both acetate and biomass. The sampling interval is 1 hour, and the simulation lasts for 36 hours, until complete depletion of substrate.

Hybrid Cybernetic Model

The HCM is a dynamic approach for metabolic modelling, initially presented by [4]. It comes as a set of ODEs, which stem from GEM networks. It incorporates a regulatory process on internal metabolism by using "artificial enzymes" and cybernetic variables [6]. This approach describes the cellular attempt to make optimal use of its resources, among competing pathways, in response to extracellular nutrient availability.

Extracellular Metabolite Balance

The extracellular metabolite balance is given by the following Equations:

$$\frac{dm_{ex}}{dt} = S \bullet Z^T \bullet \mathbb{I} \bullet v \bullet r \bullet c \tag{2}$$

$$r_i = k_i^{max} \bullet e_{rel,i} \bullet \left(\frac{s}{K_i + s}\right) \tag{3}$$

where m_{ex} $(\mathrm{mmol/L})$ denotes the extracellular metabolite concentration and r $(\mathrm{mmol/gDW*h})$ is a vector for all pathway specific uptake rates, S $(n_{metabolites}$ x $n_{reactions})$ is the GEM stoichiometric matrix, Z $(n_{reactions}$ x $n_{pathways})$ extracts key yielding pathways from S, c $(\mathrm{gDW/L})$ is the biomass concentration, v [-] is a vector for all pathways cybernetic enzyme activity, k_i^{max} $(\mathrm{mmol/gDW*h})$ is the maximum specific uptake rate, e_{rel} [-] is the relative enzyme level for every pathway, s $(\mathrm{mmol/L})$ is the substrate concentration and K $(\mathrm{mmol/L})$ is the half-maximum saturation constant.

Biomass Balance

The biomass concentration is given by:

$$\frac{dc}{dt} = \mu \bullet c \tag{4}$$

where μ $(1/h)$ is the total specific growth rate, given by the weighted sum of all fluxes.

Enzyme Balance

As stated, key "artificial enzymes" are allocated to every reaction of the network's metabolic pathways and catalyse substrate conversion. The relative enzyme balance is:

$$\frac{de_{rel}}{dt} = a + (\mathbb{I} \bullet u \bullet r_{e,rel}) - (\mathbb{I} \bullet \beta + \mu) \bullet e_{rel} \tag{5}$$

$$r_{e,rel,i} = k_{e,i} \bullet \left(\frac{s}{K_i + s}\right) \tag{6}$$

where a $(1/h)$ and β $(1/h)$ are the enzymatic constitutive synthesis and degradation rate, respectively. Moreover, u (-) is the cybernetic variable for enzyme synthesis, $\mathbb{I}$ is the identity matrix, $r_{e,rel}$ $(1/h)$ is the specific enzyme synthesis rate and $k_{e,i}$ $(1/h)$ is the maximum enzyme synthesis rate.

Cybernetic Laws

The allocated enzymes are themselves governed by two cybernetic variables, u and v. Both are bounded [0, 1] and are determined through optimal control theory [7]. Therefore, two cybernetic laws are formulated:

$$u = \frac{p_i}{\sum p_i} \tag{7}$$

$$v = \frac{p_i}{\max(p_i)} \tag{8}$$

Specifically, u regulates enzyme synthesis by following the "Matching Law ", and v regulates enzyme activity by following the "Proportional Law". For both cybernetic variables, the cellular metabolic goal p is defined as follows:

$$p_i = f \bullet r_i \tag{9}$$

Where f (-) is the substrate's number of carbon atoms, and r is each pathway's specific uptake rate.

Pathway Matrix - Z

The *e_coli_core* GEM contains 72 metabolites and 95 reactions; the stoichiometric matrix S is of size [72 x 95]. To reduce the total pathways, and therefore parameters, a key-pathway identification routine takes place.

Opt-yield Flux Balance Analysis

An opt-yield FBA algorithm was employed to identify a physiologically relevant set of yield-maximising metabolic pathways [8]. Yield is defined as the ratio of product formation rate (r_p) to substrate uptake rate (r_s). In short, opt-yield FBA transforms the nonlinear fractional objective of yield $\left(Y = \frac{r_p}{r_s}\right)$, into a linear one $(r_p - Y_{temp} \bullet r_s)$, by introducing a Y_{temp} temporary yield, and optimises the latter iteratively until the maximal yield is found. The algorithm was applied to the *e_coli_core* GEM under aerobic conditions, and bound glucose (-10, -0.1) $mmol/gDW*h$ to compute the yield space for biomass and acetate, having glucose as substrate.

Metabolic Yield Analysis

Subsequently, MYA was used to reduce the opt-yield FBA pathways to a minimal set that covers 99% of the original convex yield space. A bottom-up approach was used select the minimum number of biochemical pathways, by sequentially comparing the spanned convex hull volume between the reduced and original yield space [9]. The minimal number of pathways is used for formulating the HCM through the Z matrix.

Parameter Estimation

For the HCM, 5 unknown parameters (k^{max}, K, a, β, k_e) exist. Every unknown parameter vector has size equal to the number of columns (i.e., minimal pathways) of Z. Values for a, β, and k_e are selected from the literature [6-8].

Nonlinear Least Squares Regression

Model parameters k^{max} and K were estimated with least squares regression, minimising the sum of squares errors between model predictions and in silico data for biomass, glucose and acetate concentrations. The objective function was defined as:

$$LSE(\theta) = \sum_i^N \left(y_{sampled,i} - y_{model,t}(\theta)\right)^2 \quad (10)$$

where θ denotes the vector of unknown parameters, $y_{sampled}$ represents the measured data, and $y_{model}(\theta)$ the model predictions. To assess the goodness of fit, the Mean Squared Error (MSE) was used:

$$MSE = \frac{1}{n_{sampled} - n_\theta} \bullet LSE(\theta) \quad (11)$$

Here, n signifies either the number of sampled data or parameters.

Parameter Identifiability

To evaluate parameter practical identifiability, we employed the profile likelihood approach [10, 11]. For each parameter θ, the profile likelihood $\chi^2_{PL}(\theta_i)$ was computed by re-optimising all other parameters while keeping θ_i fixed at a series of values:

$$\chi^2_{PL} = \min\left[\chi^2(\theta)\right] \quad (12)$$

where, $\chi^2(\theta) = LSE(\theta)$ under the assumption of normally distributed residuals. The likelihood-based confidence interval (CI) at level α was defined by the threshold:

$$\Delta_a = \chi^2(\alpha, df = 1) \quad (13)$$

where α is the quantile of the χ^2 distribution with df degrees of freedom. A parameter was considered practically non-identifiable if its profile likelihood remained below Δ_α in either the increasing or decreasing direction, indicating an infinitely extended confidence bound in that direction. CIs were extracted from the parameter values were χ^2_{PL} crossed the threshold Δ_α.

Local Sensitivity Analysis

All parameters were examined for their influence on model output via local sensitivity analysis. The sensitivity matrix S was computed numerically as:

$$S_{i,j} = \frac{\partial y_i}{\partial \theta_j} \quad (14)$$

Parameters with consistently low sensitivity and flat profile likelihoods were considered for model reduction.

Optimal Experiment Design

To reduce parameter uncertainty, OED was performed using the Modified E-optimality criterion [12]. The Fischer Information Matrix (FIM) was approximated from local sensitivities, evaluated at the nominal parameter estimates:

$$FIM(\theta) = \sum_{k=1}^{N_t} S(\theta, t_k)^T \bullet Q^{-1} \bullet S(\theta, t_k) \quad (15)$$

Where $S(\theta, t_k) = \partial y/\partial\theta|t_k$ is the sensitivity matrix at time t_k, and Q is the covariance matrix of measurement errors. The modified E-optimality criterion aims to minimise the condition number of the FIM, which is equivalent to minimising the ratio of the largest to the smallest eigenvalue:

$$E_{criterion} = \min\left[\frac{\lambda_{max}(FIM(\theta))}{\lambda_{min}(FIM(\theta))}\right] \quad (16)$$

An additional *in silico* experiment was designed by solving this optimisation problem under feasible operational bounds, yielding an informative input profile (e.g., a time-varying substrate feed) that was simulated to generate new data for re-estimation.

Numerical Simulations

All computations were performed in Spyder v6 IDE virtual environment, using the Python 3.10.6 programming language. A numerical simulation for the whole pipeline lasts approximately 7 hours and 15 minutes.

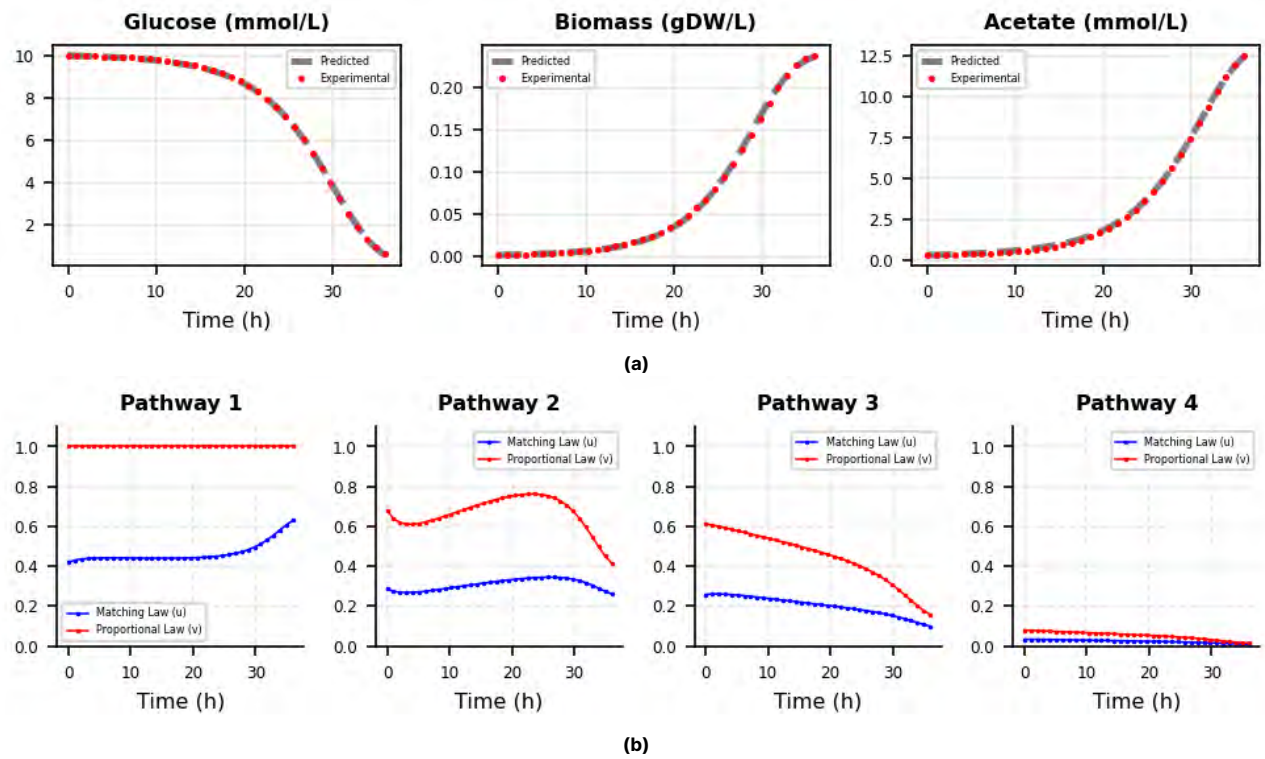

Figure 2: Model fit of HCM for 4 pathways. **(a)** HCM predictions for glucose, biomass and acetate states, using optimal parameters **(Table 1)**, **(b)** Simulated cybernetic variables (u, v) for model's fit.

RESULTS & DISCUSSION

Yield Space

The opt-yield FBA results in 22 physiologically relevant metabolic pathways, which together span the - acetate/biomass over glucose - yield space. Then, MYA identifies 4 pathways containing 99% of the original convex hull. From these, 2 identified pathways correspond to high-yield acetate and biomass production (pathway 1 & 2, respectively), where a third one represents a no-growth state of complete respiration process.

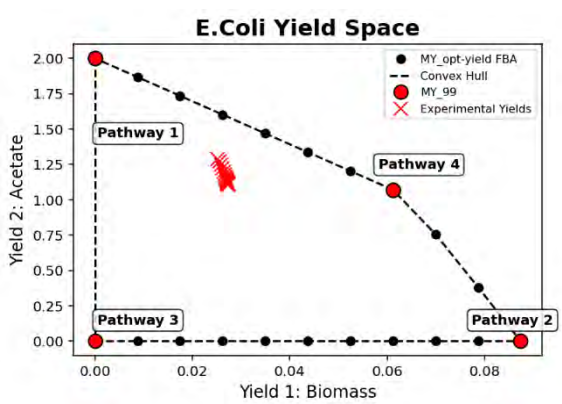

Figure 1. Opt-yield FBA and MYA convex yield space.

With pathway 4, which contains a set of intra-extracellular reaction fluxes for combined acetate & biomass secretion. All four are forming a Z-matrix and include all the necessary knowledge towards describing trade-offs between growth and by-product formation.

HCM – 4 Pathways

Model Fit

The HCM accurately reproduces the batch dynamics of biomass growth, glucose uptake and acetate secretion, while preserving metabolic mechanistic interpretability. Model predictions achieve a low MSE for all measured states (Table 1). Meanwhile, cybernetic variables allocate enzyme activity in a physiologically consistent way; growth-oriented pathways dominate during high glucose, and acetate-reassimilation pathways dominate for low carbon conditions.

Table 1: Optimal parameter values, HCM (4 paths).

Pathway	k_{max} (mmol/gDW*h)	K (mmol/L)	MSE
1	0.715	1.404	0.002
2	0.717	6.798	
3	0.587	5.257	
4	0.100	10.00	

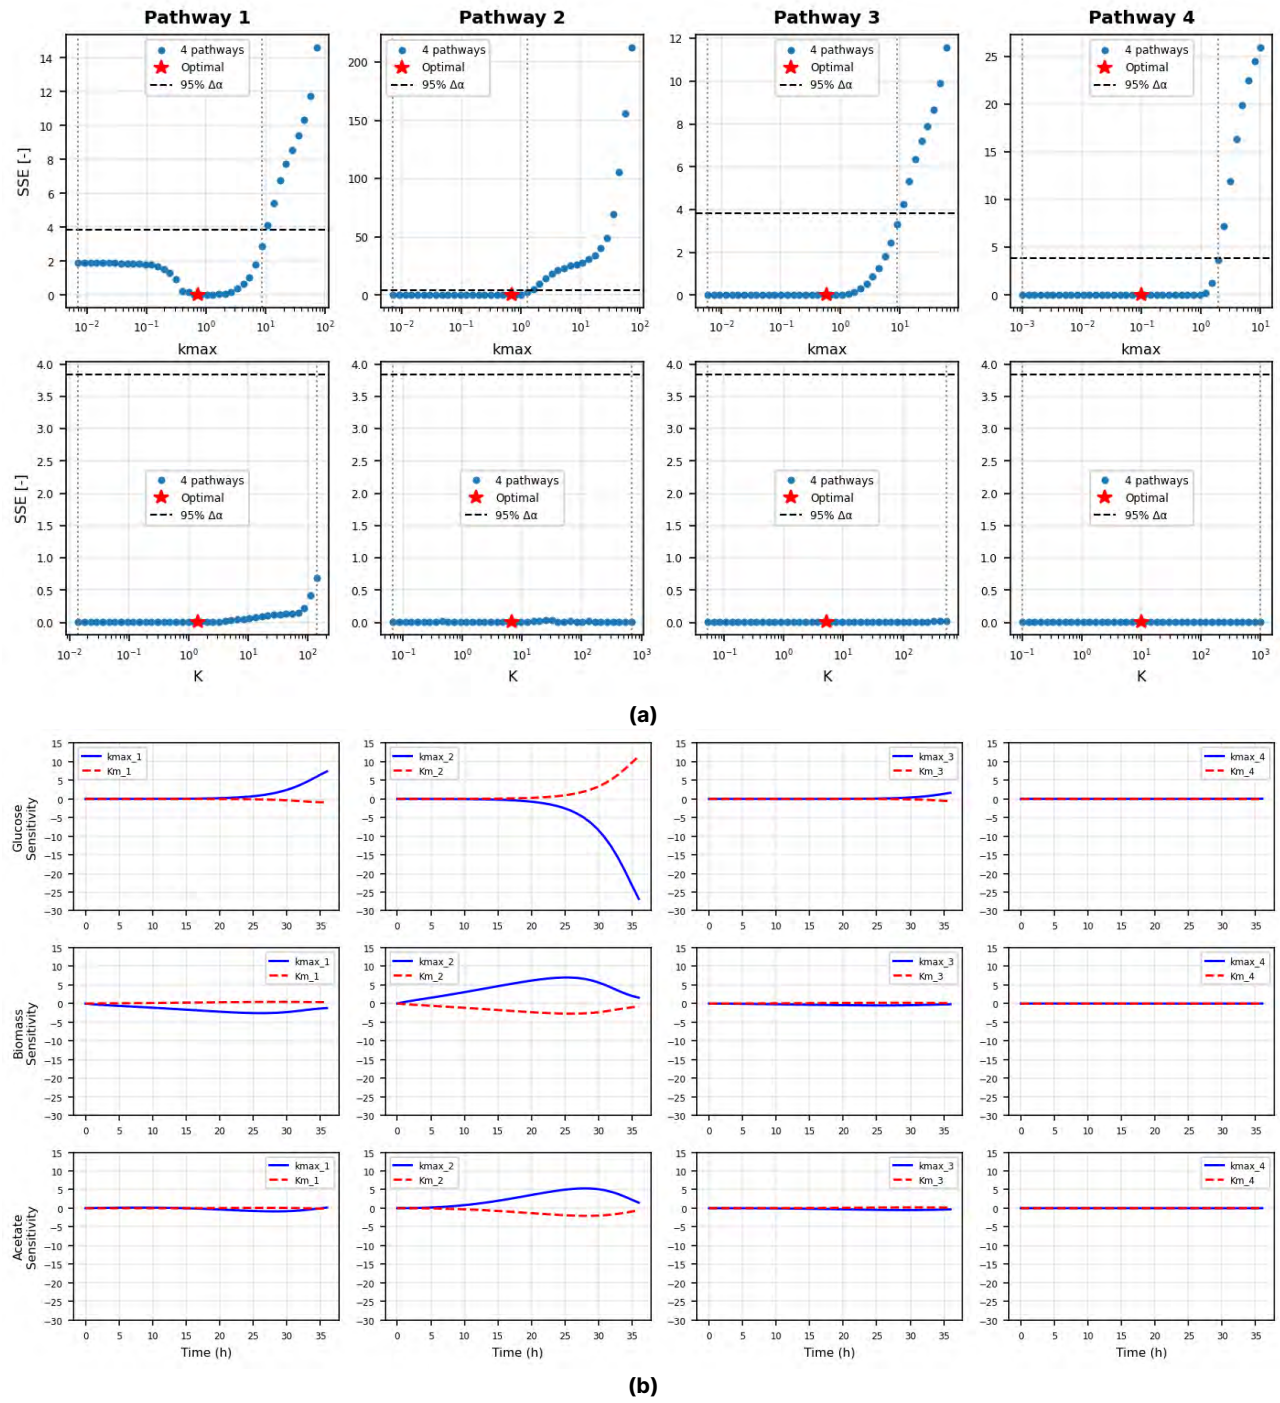

Figure 3: Parameter identifiability and sensitivity results, of HCM, for optimal parameters **(Table 1)**. **(a)** Profile Likelihoods. **(b)** Normalised sensitivity profiles; impact of perturbed parameter on model outputs.

Within all pathways, the 4th stands out; it is least sensitive to substrate concentration, has a low exchange rate constant (K_4 = 10, k_4^{max}=0.1), while pathway enzyme synthesis and activity are approaching 0 over time.

Identifiability & Sensitivity Analysis

Profile likelihoods and local sensitivities reveal that not all pathway-specific uptake parameters are practically identifiable with the initial experiment, motivating the need for model reduction and additional input data. Similar identifiability issues have been reported for large kinetic and population-balance models, where flat profile likelihoods indicate over-parametrisation and the need for targeted experiments through OED [10, 12].

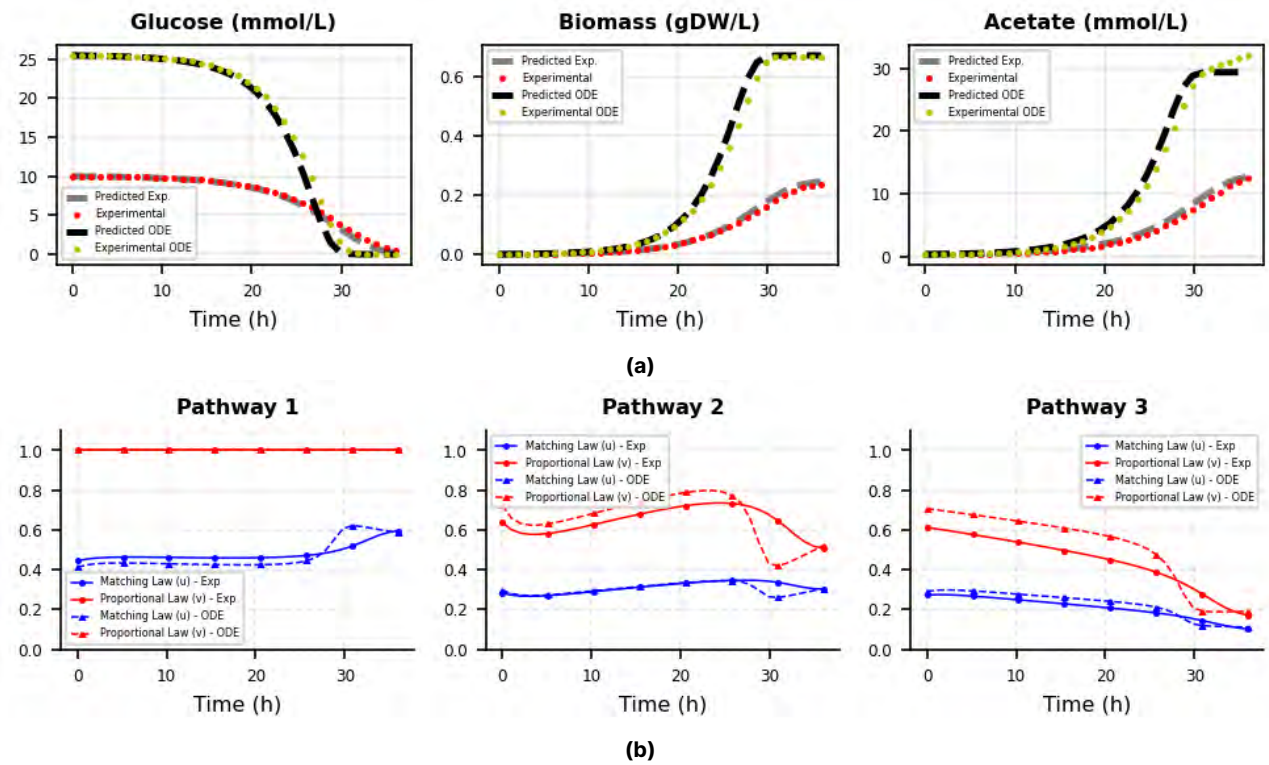

Figure 4: Model fit of ODE-enriched HCM for 3 pathways. **(a)** HCM predictions for glucose, biomass and acetate states, using optimal parameters **(Table 2)**. **(b)** Simulated cybernetic variables (u, v) for model's fit.

HCM – 3 Pathways

Model Reduction & OED

Model reduction from 4 to 3 pathways is motivated by the combined flat profile likelihoods (K_4, Figure 3a) and low local sensitivities (k^4_{max}, K_4 Figure 3b) for the acetate-biomass co-producing route, indicating practical non-identifiability and limited influence on measured outputs. Removing this pathway promises a more contained model structure, with chances of parameter identifiability [13].

Furthermore, on the structurally reduced 3-pathway HCM, OED is employed using the modified E-criterion to improve the conditioning of the FIM matrix. Specifically, it aims to target FIM directions that are poorly informed by the initial dataset. OED results in an optimal initial glucose concentration (i.e., 25.47 g/L), thus dFBA is performed again to design this new input, yielding a more differentiated activation of the 3 remaining pathways, over time.

Table 2: Optimal parameter values, ODE-HCM (4 paths).

Pathway	k_{max} (mmol/gDW*h)	K (mmol/L)	MSE
1	0.826	1.730	1e-4
2	0.666	4.823	
3	0.660	5.303	

Model Fit

In the joint estimation for the merged dataset, the 3-pathway HCM seem to play a complementary role. The high-yield, glucose-to-biomass route dominates early growth and remains highly activated throughout the experiment. Additionally, acetate overflow and respiration pathways are initially active, but their enzymatic promotion and activity are decreasing proportionally with time (i.e., substrate depletion) (Figure 3). This synergy leads to a very low MSE (Table 2) while maintaining physiologically plausible values for six estimated parameters.

Identifiability & Sensitivity Analysis

Model reduction, by itself, improves the practical identifiability of maximum exchange rate constants for pathways 1 and 2, since both profiles cross the $\Delta\alpha$ threshold. However, it only induces a slight change of slope for the rest of the parameters, with partial or unsuccessful crossing of the threshold, still indicating practical non-identifiability for k^3_{max}, K_1, K_2 and K_3 (Figure 5a). The OED experiment induces a clear curve across the 95% CI for three parameters, notably making K_2 practically identifiable. The 95% CI of previously identifiable parameters become narrower, but the kinetic parameters that belong to pathway 3 and K_1 remain practically non-identifiable.

Multiple reasons may contribute to this result. For example, the lack of sufficient/informative data could

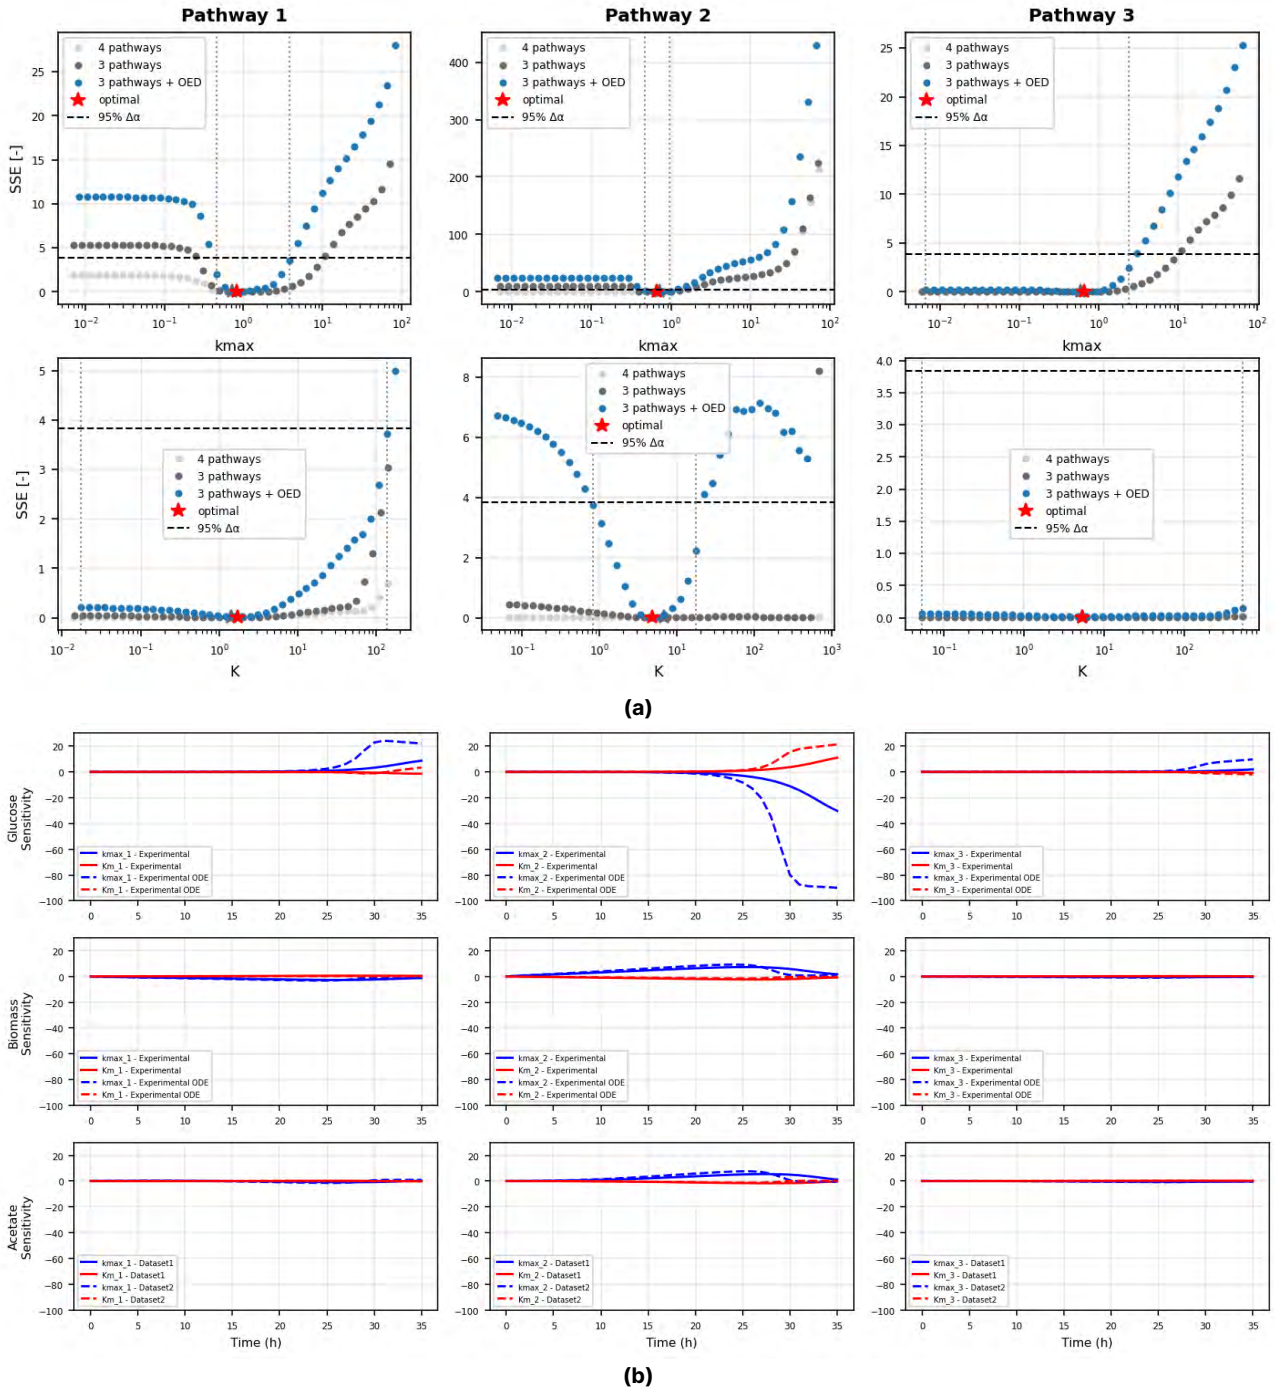

Figure 5: Parameter identifiability & sensitivity results; ODE-assisted HCM, for optimal parameters **(Table 2)**. **(a)** Profile Likelihoods, **(b)** Normalised sensitivity profiles; impact of perturbed parameters on model outputs.

lead to incomplete parameter identifiability. As local sensitivities show (Figure 5b), the OED dataset improves the sensitivity of model output over parameter perturbations mostly for pathway 2. However, it didn't succeed in separating the profiles for the pathway 1 & 2. This could require more measurements (e.g. CO_2 for pathway 3), or an acetate-driven OED design is needed to improve pathway contributions between 1 & 2. Additionally, if one revisits the HCM Michaelis-Menten formulation for pathway exchange rate, it is of intuitive understanding that the k_{max} and K parameters are countering each other during the profile likelihood analysis. Specifically, when the K value becomes too low, the parameter is negligible towards the value of k_{max}; thus, the parameter is not

informed by the data. On the other hand, when K achieves a large value, substrate concentration is negligible, thus only the $\frac{k_{max}}{K}$ ratio can be identifiable. It is therefore a matter of proper experiment design, as well as inclusion of multiple measurements, to achieve a practically identifiable profile for the K parameters, as the K_2 profile likelihood signifies.

CONCLUSIONS

The proposed workflow shows that HCMs can be systematically reduced and still capture microbial metabolic dynamics. Combining methods like opt-yield FBA, MYA and profile likelihood analysis yields a statistically well-defined 3-pathway HCM that retains almost all the original yield space. A modified E-criterion provides an additional informative experiment that improves confidence bounds, measurement sensitivity and parameters' practical identifiability. Together, these steps showcase how integrated analysis enables the reliable employment of mechanistic models for robust prediction and control of microbial bioprocesses.

ACKNOWLEDGEMENTS

Stylianos Floros was supported by the Research Foundation – Flanders (FWO) under grant No. 1SACZ26N. Satyajeet S. Bhonsale and Simen Akkermans were supported by the Research Foundation – Flanders (FWO) under grant No. G0A9H26N.

AUTHOR IDENTIFIERS

Author ORCIDs:
Stylianos Floros: 0009-0000-2217-7542
Satyajeet S. Bhonsale: 0000-0001-9734-4122
Simen Akkermans: 0000-0003-2904-0439
Jan F.M. Van Impe: 0000-0002-5904-1638

REFERENCES

1. Nielsen J, Petranovic D. Modeling for understanding and engineering metabolism. QRB Discovery 6: (2025). https://doi.org/10.1017/qrd.2025.1
2. Pennington O, Espinel Ríos S, Sebastian MT, Dickson A, Zhang D. A multiscale hybrid modelling methodology for cell cultures enabled by enzyme-constrained dynamic metabolic flux analysis under uncertainty. Metabolic Engineering 86:274-287 (2024). https://doi.org/10.1016/j.ymben.2024.10.013
3. Gábor A, Villaverde AF, Banga JR. Parameter identifiability analysis and visualization in large-scale kinetic models of biosystems. BMC Syst Biol 11: (2017). https://doi.org/10.1186/s12918-017-0428-y
4. Kim JI, Varner JD, Ramkrishna D. A hybrid model of anaerobic e. coli GJT001: combination of elementary flux modes and cybernetic variables. Biotechnology Progress 24:993-1006 (2008). https://doi.org/10.1002/btpr.73
5. King ZA, Lu J, Dräger A, Miller P, Federowicz S, Lerman JA, Ebrahim A, Palsson BO, Lewis NE. Bigg models: a platform for integrating, standardizing and sharing genome-scale models. Nucleic Acids Res 44:D515-D522 (2015). https://doi.org/10.1093/nar/gkv1049
6. Song H, Ramkrishna D. Prediction of metabolic function from limited data: lumped hybrid cybernetic modeling (L?HCM). Biotech & Bioengineering 106:271-284 (2010). https://doi.org/10.1002/bit.22692
7. Song HS, Ramkrishna D. Prediction of dynamic behavior of mutant strains from limited wild-type data. Metabolic Engineering 14:69-80 (2012). https://doi.org/10.1016/j.ymben.2012.02.003
8. Luo H, Li P, Ji B, Nielsen J. Modeling the metabolic dynamics at the genome-scale by optimized yield analysis. Metabolic Engineering 75:119-130 (2023). https://doi.org/10.1016/j.ymben.2022.12.001
9. Song H, Ramkrishna D. Reduction of a set of elementary modes using yield analysis. Biotech & Bioengineering 102:554-568 (2008). https://doi.org/10.1002/bit.22062
10. Bhonsale SS, Stokbroekx B, Van Impe J. Assessment of the parameter identifiability of population balance models for air jet mills. Computers & Chemical Engineering 143:107056 (2020). https://doi.org/10.1016/j.compchemeng.2020.107056
11. Nimmegeers P, Bhonsale S, Telen D, Van Impe J. Optimal experiment design under parametric uncertainty: a comparison of a sensitivities based approach versus a polynomial chaos based stochastic approach. Chemical Engineering Science 221:115651 (2020). https://doi.org/10.1016/j.ces.2020.115651
12. Bhonsale S, Nimmegeers P, Akkermans S, Telen D, Stamati I, Logist F, Van Impe JFM. Optimal experiment design for dynamic processes. Simulation and Optimization in Process Engineering :243-271 (2022). https://doi.org/10.1016/b978-0-323-85043-8.00010-6
13. Raue A, Karlsson J, Saccomani MP, Jirstrand M, Timmer J. Comparison of approaches for parameter identifiability analysis of biological systems. Bioinformatics 30:1440-1448 (2014). https://doi.org/10.1093/bioinformatics/btu006

14. Johnson KA, Goody RS. The original michaelis constant: translation of the 1913 michaelis–menten paper. Biochemistry 50:8264-8269 (2011). https://doi.org/10.1021/bi201284u

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Semi-Supervised Generative Augmentation Improves Surfactant Surface Tension Prediction from Limited Experimental Data

Gabriela C. Theis Marchan[a], Kyle Territo[a], Jose A. Romagnoli[a]*

[a] Louisiana State University, Department of Chemical Engineering, Baton Rouge, Louisiana 70803, United States
* Corresponding Author: jose@lsu.edu.

ABSTRACT

Predictive modeling of surfactant properties is constrained by limited experimental datasets, a common challenge in specialty chemical development where property measurements require specialized equipment and significant time investment. In this study, we address data scarcity through a semi-supervised generative augmentation framework that leverages both labeled and unlabeled molecular data for surface tension prediction. We implemented a two-stage variational autoencoder (VAE) training strategy using a curated database of 600 non-ionic surfactants. First, 461 unlabeled surfactant structures were used for VAE pre-training to learn latent representations capturing molecular connectivity patterns and amphiphilic relationships. Second, 125 molecules with surface tension measurements were used for fine-tuning to embed property-structure relationships. Our stratified generation framework produces surfactants matching target property distributions (Wasserstein distance = 0.030, KS statistic = 0.152) through stratified sampling and multi-criteria filtering, achieving 88.5% structural validity. We systematically evaluated augmentation effectiveness across seven levels (×1.25 to ×5) using graph convolutional networks with five random seeds per level. Results show optimal performance at ×3 augmentation (375 training molecules), where test R^2 improved from 0.60 (baseline, 125 molecules) to 0.78 (+28%), with RMSE decreasing to 0.056 mN/m. Multi-seed analysis demonstrates stable and reproducible training (test R^2 = 0.776 ± 0.030). This framework achieves performance equivalent to tripling the experimental dataset, providing a practical and scalable approach for specialty chemical property prediction where structural data is abundant but property labels are scarce.

Keywords: surfactants, surface tension, data augmentation, VAE, GCN, semi-supervised learning

1. INTRODUCTION

Surfactants are amphiphilic molecules central to formulation chemistry, with applications spanning detergents, cosmetics, pharmaceuticals, and enhanced oil recovery.[1] Their interfacial properties, particularly surface tension at the critical micelle concentration (CMC), directly govern product performance. However, experimental characterization requires specialized tensiometer equipment and careful measurement protocols, [2] creating a bottleneck in surfactant discovery, where thousands of candidate molecules might be considered, but only dozens can be practically tested.

Machine learning offers a pathway to accelerate surfactant discovery by predicting properties from molecular structure.[3-5] Graph convolutional networks (GCNs) have demonstrated strong performance in molecular property prediction tasks.[6] However, these models require substantial training data to learn robust structure-property relationships, conflicting with the reality of specialty chemical development where experimental datasets are typically small (<200 molecules) due to measurement costs and domain specificity.

Data augmentation addresses this limitation by synthetically expanding training sets. While traditional augmentation in computer vision leverages known invariances, [7, 8] molecular augmentation requires generating chemically valid structures with properties matching

target distributions. Generative models, particularly variational autoencoders (VAEs), offer a principled approach by learning continuous latent representations enabling controlled generation of novel structures.[7, 9] When coupled with property prediction, VAEs can generate molecules matching specific property targets.[10] However, training effective molecular VAEs requires substantial data, presenting a fundamental challenge for small experimental datasets.

We address this through semi-supervised learning, leveraging the asymmetry between structural data availability (abundant) and property measurements (scarce). Our approach uses unlabeled surfactant structures to pre-train a molecular VAE, learning valid connectivity patterns, then fine-tunes on property-labeled molecules to embed structure-property relationships. This enables generation of synthetic training data that is both chemically valid and property-relevant.

In this work, we present a systematic evaluation of semi-supervised generative augmentation for non-ionic surfactant surface tension prediction. We demonstrate that intelligent augmentation can improve predictive performance equivalent to tripling the experimental dataset, while identifying fundamental limitations in generating molecules at property extremes and providing comprehensive validation through chemical space analysis and multi-seed evaluation. Our analysis reveals why intermediate augmentation levels show instability and provides guidelines for optimal augmentation strategies. Figure 1 summarizes the proposed semi-supervised augmentation workflow, including two-stage VAE training, stratified generation, and multi-seed GCN retraining.

2. METHODS

Figure 1 provides an overview of the full pipeline. Briefly, a baseline GCN is trained on labeled data and used as a property oracle to pseudo-label VAE-generated molecules; stratified augmentation sets are then constructed and used to retrain GCN models under multiple random seeds.

2.1 Dataset and Preprocessing

We curated a database of 600 unique non-ionic surfactant molecules from literature sources.[1, 3, 5, 11, 12] Of these, 139 molecules had experimentally measured surface tension, while 461 had only structural information (SMILES strings). Surface tension measurements ranged from 18.0 to 47.5 mN/m. All values were normalized to [0, 1] using min-max scaling.[13]

The 139 labeled molecules were randomly split (seed =2024) into training (125 molecules, 90%) and test (14 molecules. 10%) sets. The split achieved good property distribution balance, with training set mean of 0.460 (std=0.179) and test set mean of 0.471 (std=0.123) in normalized units. All molecules were validated for chemical correctness using RDKit (2024.9.6), [14] and SMILES were canonicalized. The test set was held out for all experiments and used only for final evaluation.

2.2 Baseline GCN model

Our surface-tension baseline is a molecular GCN regressor built on graph representations $G = (V, E)$, with atoms as nodes and bonds as edges.[15] Atom states are initialized from RDKit-derived descriptors capturing local chemistry (e.g., elemental identity and bonding environment). The network updates these states through two GraphConv layers (256 channels) using neighborhood aggregation with learnable weights and ReLU nonlinearities, enabling information to propagate across the molecular structure. Node embeddings are then summarized via global mean pooling to yield a fixed-dimensional molecular fingerprint. This pooled representation is mapped to the final prediction by a compact fully connected regressor (256→256→1) that outputs the normalized surface tension. Training employed early stopping using validation RMSE for checkpoint selection, and performance was reported using R^2, RMSE, and MAE on the held-out experimental test set.

2.3 Semi-supervised VAE training

A VAE was used to learn a continuous generative representation of non-ionic surfactant structures from their SMILES-based encoding. Training followed a two-stage strategy to exploit the larger pool of unlabeled structures. In Stage 1, the VAE was pre-trained on the 461 unlabeled molecules to learn a chemically plausible latent manifold. In Stage 2, the model was fine-tuned on the 125 labeled training molecules to align the latent space with the labeled domain used for surface-tension learning. The fine-tuned VAE parameters were saved and reused for all subsequent generation experiments.

2.4 Stratified latent-space generation and filtering

Generation operates exclusively through the VAE decoder: latent vectors z ~ $\mathcal{N}$(0, I) are sampled directly without an encoder pass and decoded to candidate SELFIES strings. To construct a property-relevant synthetic pool, we performed **stratified latent-space generation** using the fine-tuned VAE coupled to the baseline GCN as a property oracle. The target surface-tension distribution was first estimated from the **training labels** by binning normalized surface tension into **10 intervals** and computing the empirical proportion in each bin. Synthetic generation then proceeded iteratively: latent

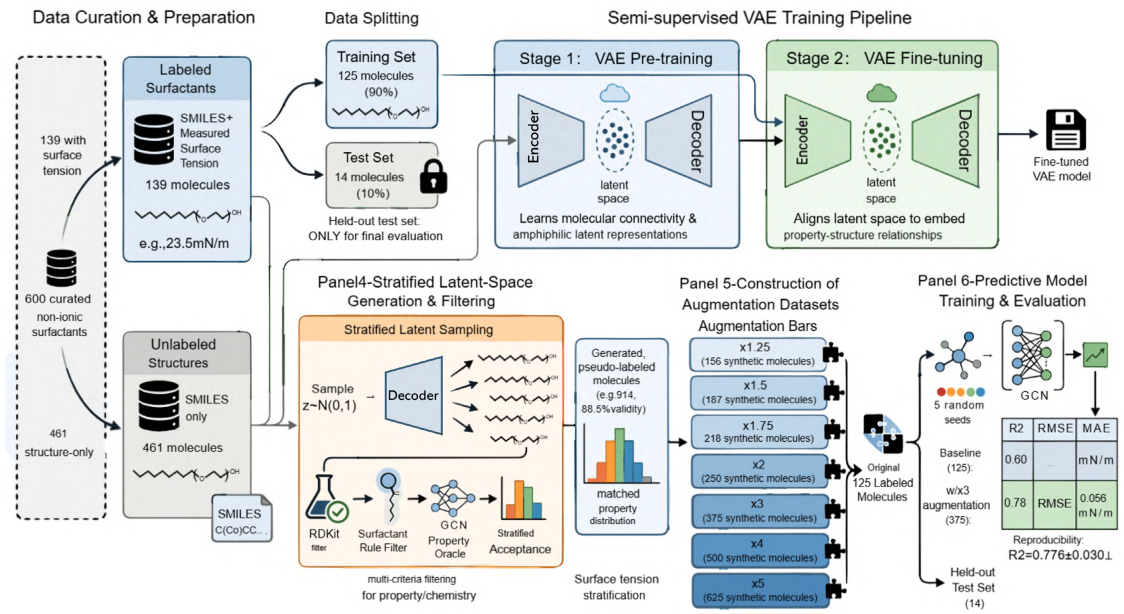

Figure 1. Semi-supervised generative augmentation pipeline for surface-tension prediction. Labeled surfactants (n=139) are split into training (n=125) and held-out test (n=14). A VAE is pre-trained on unlabeled structures (n=461) and fine-tuned on the labeled training set. Synthetic molecules are generated via stratified latent sampling, filtered for chemical validity and surfactant plausibility, and pseudo-labeled using a baseline GCN predictor. Stratified, independent augmentation sets (×1.25–×5) are constructed from the pseudo-labeled pool and combined with the labeled training data to retrain GCN models across five random seeds; performance is reported on the fixed experimental test set. Note that Panel 4 (Stratified Latent Sampling) uses only the VAE decoder: latent vectors z are sampled directly from $\mathcal{N}(0,I)$ and decoded into candidate molecules; the encoder is not used during generation.

vectors $z \sim \mathcal{N}(0, I)$ were sampled, decoded to candidate molecules, converted from SELFIES to SMILES, and filtered by RDKit sanitization and radical checks; candidates were then assigned a predicted surface tension using the trained GCN, and accepted only if their prediction fell into a bin that had not yet reached its target quota.

To improve coverage of underfilled bins (particularly near the distribution tails), generation used a **multi-phase adaptive strategy** with increasing sampling temperature and targeted sampling around previously successful latent regions for difficult bins. In addition, a surfactant-specific quality filter was applied during acceptance using a rule-based evaluator that screens for a plausible hydrophobic tail, presence of hydrophilic head groups, and basic structural stability heuristics, with a minimum surfactant score threshold applied during generation.

In the final generation run, the procedure targeted 1000 molecules but accepted **914** after reaching the maximum attempt budget; the RDKit-valid fraction among decoded candidates was **88.5%**, and the generated pseudo-label distribution closely matched the training distribution (Wasserstein distance **0.0298**). Generated molecules spanned predicted normalized surface tension from **0.002 to 0.799**.

2.5 Augmentation construction and retraining protocol

From the pseudo-labeled generated pool (SMILES, Prediction), we constructed augmentation datasets at seven levels (×1.25–×5) using independent stratified sampling. For each level, generated molecules were partitioned into bins defined by the training-label quartiles of normalized surface tension, and molecules were sampled from each bin in proportion to the empirical bin frequencies of the real training set. If a bin contained fewer generated candidates than required, all available candidates were retained and the remaining quota was filled by sampling from the residual pool. Each augmentation level was generated with a deterministic level-specific seed to ensure reproducibility while keeping levels non-nested.

For each augmentation level, the synthetic subset was concatenated with the 125 labeled training molecules (using pseudo-labels for synthetic samples), and a GCN was retrained using the baseline hyperparameters. Robustness was assessed across five random seeds (42, 104, 2021, 2024, 4592) with early stopping based on validation RMSE; results are reported as mean ± standard

deviation of test R^2, RMSE, and MAE on the unchanged experimental test set.

3. RESULTS AND DISCUSSION

3.1 Baseline GCN performance

Using only the labeled training set (125 molecules), the baseline GCN achieved a test R^2 of 0.599 with RMSE = 0.075 and MAE = 0.052 in normalized units (Table 1). The gap between training and test performance (train $R^2 = 0.868$) highlights the limited generalization achievable with small labeled datasets and motivates the use of generative augmentation.

Table 1: Baseline GCN performance on surface tension normalized.

data	R^2	RMSE	MAE
train	0.868	0.065	0.046
test	0.599	0.075	0.052

3.2 Generated dataset validity and distribution matching.

The semi-supervised VAE generation pipeline produced 914 accepted molecules after RDKit validation and rule-based surfactant filtering, corresponding to an overall structural validity of 88.5%.

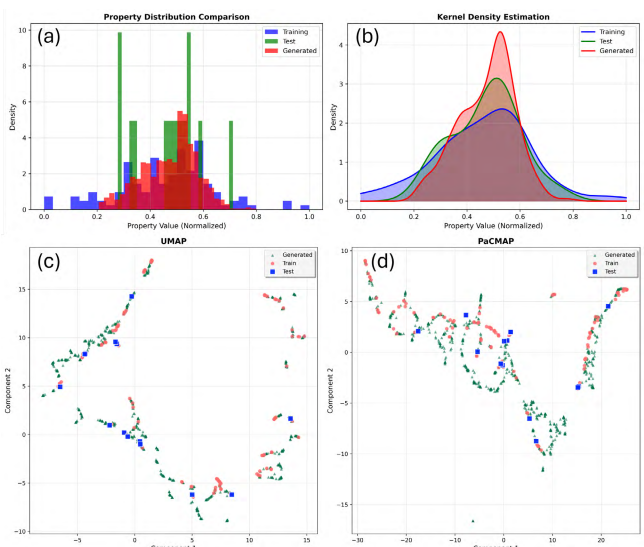

Figure 2. Generated data validity, distribution matching, and chemical-space coverage. (a) Normalized surface-tension distributions for training, test, and generated (pseudo-labeled) molecules. (b) Kernel density estimates showing close agreement between generated and empirical distributions. (c) UMAP and (d) PaCMAP projections illustrating that generated samples overlap the empirical domain while modestly expanding chemical-space coverage.

The pseudo-labeled surface-tension distribution of the generated pool closely matched the empirical training distribution (Wasserstein distance = 0.0298; KS statistic = 0.152), indicating that stratified sampling successfully controls property coverage (Figure 2a-b). The chemical space projections in (Figure 2c-d) reveal that VAE-generated molecules predominantly occupy regions already covered by the training data, with modest expansion into adjacent chemical space. This behavior reflects the fine-tuning strategy employed in Stage 2 of VAE training, where the model learns to generate structures similar to the 125 labeled surfactants used for property embedding. While this ensures generated molecules remain within the interpolation regime of the oracle GCN (reducing extrapolation error in pseudo-labeling), it also limits the framework's ability to discover truly novel surfactant chemotypes or explore underrepresented regions of structure-property space.

The narrower spread of the generated distribution visible in Figure 2a-b warrants discussion in the context of VAE mode collapse. Mode collapse in VAEs occurs when the decoder maps only a restricted region of the latent space, producing low-diversity outputs concentrated near the distribution mean and underrepresenting the tails. In our case, however, the reduced spread is primarily a consequence of two deliberate design choices rather than pathological collapse: (1) Stage 2 fine-tuning on only 125 labeled surfactants biases the latent manifold toward the well-represented central property region, and (2) the stratified generation strategy explicitly targets the empirical training distribution rather than attempting to extrapolate beyond it. The low Wasserstein distance (0.030) and KS statistic (0.152) confirm that the generated molecules faithfully reproduce the shape and central tendency of the training distribution. That said, we acknowledge that tail regions, corresponding to very low and very high surface tension, remain underrepresented in the generated pool, consistent with limited tail coverage in the fine-tuning data and mild posterior compression in those latent regions. Future work incorporating β-VAE annealing or tail-targeted latent sampling could improve extremal coverage.

UMAP projections (Figure 2c) show that generated molecules form clusters closely aligned with training data, suggesting the VAE successfully learned the manifold of non-ionic surfactant structures. However, the test molecules; which were held out entirely from VAE training; occupy positions slightly separated from the main training cluster, indicating that real chemical diversity extends beyond what the generative model captures. This observation highlights a key limitation: the VAE can densify coverage within known chemical neighborhoods but struggles to generate out-of-distribution molecules that would expand predictive

capabilities into genuinely unexplored territory.

The PaCMAP projection (Figure 2d), which emphasizes global structure preservation, further confirms that generated molecules largely interpolate between existing training examples rather than extrapolating to new regions. This interpolative behavior is generally desirable for augmentation in small-data regimes, as it reduces the risk of generating unrealistic molecules or introducing covariate shift that degrades model performance. However, it also means that this augmentation approach is most effective for improving prediction accuracy within the existing property range and chemical domain, rather than enabling discovery of surfactants with properties beyond those already measured experimentally.

3.3 Augmentation level controls predictive performance.

We evaluated seven augmentation levels (×1.25–×5) under five random seeds per level. Performance improved substantially with moderate augmentation and peaked at ×3, where the model achieved mean test $R^2 = 0.776 \pm 0.030$ and RMSE $= 0.0559 \pm 0.0039$ (Figure 3; Table 2).

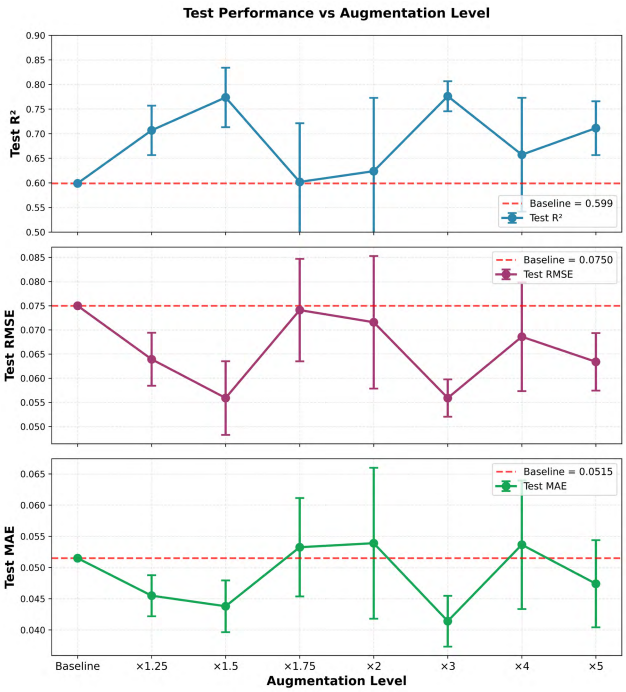

Figure 3. Test performance as a function of augmentation level. Mean ± standard deviation across five random seeds for test (top) R², (middle) RMSE, and (bottom) MAE for augmentation levels ×1.25–×5. Dashed lines indicate the baseline performance trained on labeled data only.

Relative to the baseline, ×3 augmentation increases R^2 by 0.177 (≈29%) and reduces RMSE by ≈25%, indicating that adding pseudo-labeled samples can substantially improve generalization. Performance was non-monotonic across augmentation levels, with intermediate levels (e.g., ×1.75–×2) exhibiting increased variance and lower mean test R^2, suggesting that beyond a certain point, augmentation may introduce pseudo-label noise and/or distributional drift that offsets gains from larger sample size.

Table 2: Test performance across augmentation levels (mean ± std over 5 seeds). RMSE is reported in MinMax-scaled units.

Level	Train size	Test R2	Test RMSE
Baseline	125	0.599	0.075
X1.25	156	0.707 ± 0.050	0.064 ± 0.005
X1.5	187	0.774 ± 0.060	0.056 ± 0.008
X1.75	218	0.602 ± 0.119	0.074 ± 0.011
X2	250	0.624 ± 0.149	0.072 ± 0.014
X3	375	0.776 ± 0.030	0.056 ± 0.004
X4	500	0.657 ± 0.116	0.069 ± 0.011
X5	625	0.711 ± 0.055	0.063 ± 0.006

3.4 Reproducibility and Understanding Augmentation Limits.

At the optimal ×3 level, test performance was consistent across seeds with relatively small dispersion in test RMSE and MAE (Figure 4). This indicates that the observed gains are not driven by a favorable random split or initialization, supporting the stability of the proposed augmentation framework under training stochasticity.

The non-monotonic relationship between augmentation level and predictive performance reveals fundamental trade-offs in generative augmentation strategies. While the optimal ×3 level demonstrates clear benefits, intermediate levels (×1.75–×2) exhibit performance instability with increased variance and occasional drops below baseline. This behavior suggests a transition regime where competing effects interact: modest augmentation provides insufficient regularization benefit to overcome pseudo-label noise, while the synthetic samples have not yet achieved the diversity and coverage necessary to meaningfully expand the training distribution.

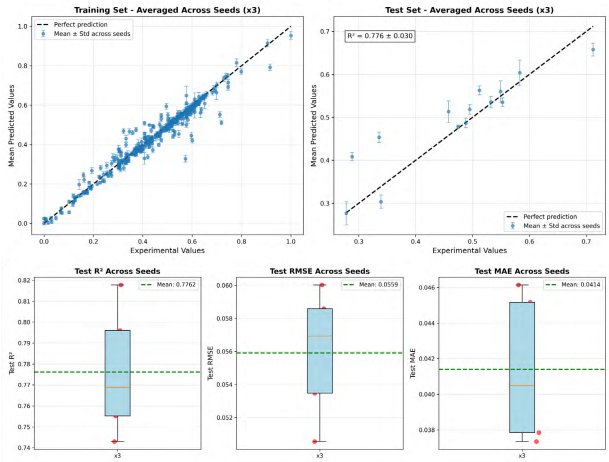

Figure 4. Reproducibility and uncertainty at the optimal augmentation level (×3). Top: parity plots for training and test sets using the mean prediction across five random seeds; error bars indicate ±1 standard deviation across seeds for each molecule. Bottom: seed-to-seed distributions of test R^2, RMSE, and MAE, demonstrating that the ×3 performance gain is robust to random initialization and training stochasticity.

At higher augmentation levels (×4–×5), performance stabilizes but does not surpass the ×3 optimum, indicating saturation or degradation from excessive reliance on pseudo-labeled data. This decline likely stems from error propagation in the oracle-based labeling process. The baseline GCN used for pseudo-labeling has inherent prediction uncertainty (test RMSE = 0.075), which compounds when generating labels for hundreds of synthetic molecules. As augmentation increases, the ratio of noisy pseudo-labels to ground-truth labels grows, eventually dominating the training signal and limiting model refinement. Additionally, the stratified generation strategy prioritizes distribution matching over tail-region accuracy, meaning molecules with extreme surface tension values are under-represented and potentially less reliable as training examples. This explains why augmentation benefits plateau rather than continue scaling linearly with dataset size.

The existence of an optimal intermediate augmentation level has practical implications for deployment: practitioners should systematically evaluate multiple augmentation ratios rather than maximizing synthetic data volume. Our multi-seed evaluation framework provides a template for identifying stable operating points where gains are reproducible across training stochasticity. Furthermore, the instability at intermediate levels suggests that augmentation strategies should either commit to minimal augmentation (×1.25–×1.5) for conservative noise control or move directly to higher levels (×3+) where increased sample diversity compensates for label uncertainty.

3.5 Practical Implications for Surfactant Discovery and Development

From a practical standpoint, achieving test R^2 = 0.78 with ×3 augmentation (375 total training samples: 125 real + 250 synthetic) represents performance equivalent to collecting approximately 250 additional experimental measurements; a substantial cost and time savings. Surface tension measurements via pendant drop or Wilhelmy plate methods typically require 2-4 hours per molecule including sample preparation, equilibration, and replicate measurements, translating to approximately 500-1000 person-hours or several months of laboratory work. The generative augmentation framework described here enables this performance gain using only computational resources (VAE training and GCN inference), which can be completed in hours on standard hardware.

This efficiency gain is particularly valuable in industrial surfactant development contexts where candidate molecules are continuously evaluated for specific applications. For example, formulation chemists might screen hundreds of potential surfactant structures in silico before selecting a subset for experimental validation. By training predictive models with augmented data, initial screening accuracy improves, reducing the number of false positives that proceed to costly synthesis and testing. The ×3 augmentation level identified in this work provides a practical guideline: for every 100 experimentally characterized molecules, generating 200 high-quality synthetic samples via stratified VAE sampling can enhance model performance by approximately 30% in R^2 while maintaining prediction stability.

However, several important limitations and boundary conditions constrain the applicability of this approach. First, augmentation effectiveness depends critically on the quality and diversity of the unlabeled structural data used for VAE pre-training. In our case, 461 unlabeled non-ionic surfactants provided sufficient coverage to learn valid molecular patterns, but domains with narrower structural diversity may require larger unlabeled pools or alternative generative architectures. Second, the approach assumes that the oracle model (baseline GCN) captures relevant structure-property relationships despite limited training data. If the baseline model has systematic biases or poor calibration in certain property regions, these errors will propagate through pseudo-labeling and potentially amplify during augmentation. Third, the framework is most effective for properties with smooth structure-property landscapes where interpolation is reliable; highly discontinuous properties or those governed by subtle molecular features may benefit less from density-based augmentation.

The stratified generation strategy is essential for maintaining distribution balance and should be adapted

to the specific property distribution of interest. For surfactant surface tension, which exhibits near-Gaussian behavior across our dataset, quartile-based stratification proved sufficient. Applications targeting skewed distributions or multimodal properties may require adaptive binning strategies or importance sampling to ensure adequate coverage of tail regions. Additionally, practitioners should validate that generated molecules satisfy domain-specific chemical rules—our surfactant-specific quality filter prevented generation of implausible structures, but similar heuristics would need to be developed for other molecular families.

Finally, while this work focused on single-property prediction (surface tension), the framework generalizes naturally to multi-property optimization scenarios common in formulation chemistry, where surfactants must simultaneously satisfy constraints on CMC, foaming behavior, biodegradability, and cost. By training the VAE on molecules with diverse property measurements and conditioning generation on multi-objective targets, the augmentation strategy could support more holistic molecular design workflows. This extension represents a promising direction for future research, particularly in specialty chemical domains where property trade-offs are critical and experimental data remains scarce.

CONCLUSION

We presented a semi-supervised generative augmentation framework for predicting surfactant surface tension from limited experimental data, demonstrating that intelligent data augmentation can overcome key bottlenecks in specialty chemical property prediction. By leveraging 461 unlabeled surfactant structures for VAE pre-training and 125 labeled molecules for fine-tuning, we generated 914 high-quality synthetic training samples (88.5% structural validity) that closely matched the target property distribution (Wasserstein distance = 0.030). Systematic evaluation across seven augmentation levels revealed optimal performance at ×3 augmentation, where test R^2 improved from 0.60 to 0.78 (+28% relative gain) with RMSE decreasing to 0.056, performance equivalent to tripling the experimental dataset.

Our multi-seed analysis confirmed that these gains are reproducible across random initialization (R^2 = 0.776 ± 0.030), establishing the stability of the approach. Chemical space analysis revealed that generated molecules primarily densify regions near existing training data rather than exploring entirely novel chemical space, which proves advantageous for interpolative prediction but limits extrapolation capabilities. The non-monotonic performance curve identified a critical trade-off between pseudo-label noise and sample diversity, with intermediate augmentation levels showing instability and diminishing returns at high augmentation ratios beyond ×3.

From a practical perspective, this framework provides a scalable and cost-effective approach to improve predictive models in data-scarce domains. Achieving high-quality predictions without collecting hundreds of additional experimental measurements addresses a fundamental challenge in specialty chemical development, where measurement costs and time constraints limit dataset sizes. The stratified generation strategy and multi-seed evaluation protocol offer a replicable template for applying generative augmentation to other molecular property prediction tasks where structural data is abundant but property labels are scarce.

Future work should focus on three key directions: (1) developing diversity-promoting generation strategies to expand chemical space coverage beyond interpolation, potentially through active learning or physics-guided sampling; (2) extending the framework to multi-property prediction scenarios where molecules must satisfy multiple constraints simultaneously; and (3) investigating transfer learning approaches where VAEs pre-trained on large unlabeled datasets can be rapidly adapted to new property targets with minimal labeled data. Additionally, integrating uncertainty quantification into the pseudo-labeling process could improve robustness by weighting synthetic samples based on oracle confidence, potentially addressing the performance degradation observed at high augmentation levels.

This work demonstrates that semi-supervised generative augmentation is a practical and effective strategy for molecular property prediction in small-data regimes, offering a promising pathway to accelerate discovery in specialty chemical domains where experimental data acquisition remains a significant bottleneck.

ACKNOWLEDGEMENTS

This work was supported by the U.S Department of Energy, Office of Science, under the Office of Basic Energy Science Separation Science program under Award No. DE-SC0022304. The authors gratefully acknowledge the computer time allotted by the high-performance computing center (HPC) at LSU and the Louisiana Network initiative.

AUTHOR IDENTIFIERS

Gabriela Theis Marchan ORCIDs: 0009-0000-9147-3943
Kyle Territo ORCIDs: 0009-0002-1133-915X
Jose Romagnoli ORCIDs: 0000-0003-3682-1305

REFERENCES

1. Rosen MJ, Kunjappu JT. Surfactants and interfacial phenomena. Wiley (2012). https://doi.org/10.1002/9781118228920

2. Scholz N, Behnke T, Resch-Genger U. Determination of the critical micelle concentration of neutral and ionic surfactants with fluorometry, conductometry, and surface tension—a method comparison. J Fluoresc 28:465-476 (2018). https://doi.org/10.1007/s10895-018-2209-4
3. Seddon D, Müller EA, Cabral JT. Machine learning hybrid approach for the prediction of surface tension profiles of hydrocarbon surfactants in aqueous solution. Journal of Colloid and Interface Science 625:328-339 (2022). https://doi.org/10.1016/j.jcis.2022.06.034
4. Chen J, Hou L, Nan J, Ni B, Dai W, Ge X. Prediction of critical micelle concentration (CMC) of surfactants based on structural differentiation using machine learning. Colloids and Surfaces A: Physicochemical and Engineering Aspects 703:135276 (2024). https://doi.org/10.1016/j.colsurfa.2024.135276
5. Marchan, G.T., et al., Harnessing Graph Learning for Surfactant Chemistry: PharmHGT, GCN, and GAT in LogCMC Prediction. 2025.
6. Liu K, Sun X, Jia L, Ma J, Xing H, Wu J, Gao H, Sun Y, Boulnois F, Fan J. Chemi-net: a molecular graph convolutional network for accurate drug property prediction. IJMS 20:3389 (2019). https://doi.org/10.3390/ijms20143389
7. Sanchez-Lengeling B, Aspuru-Guzik A. Inverse molecular design using machine learning: generative models for matter engineering. Science 361:360-365 (2018). https://doi.org/10.1126/science.aat2663
8. Wang H, Huang Z, Wu X, Xing E. Toward learning robust and invariant representations with alignment regularization and data augmentation. Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining :1846-1856 (2022). https://doi.org/10.1145/3534678.3539438
9. Jin, W., R. Barzilay, and T. Jaakkola. Junction tree variational autoencoder for molecular graph generation. in International conference on machine learning. 2018. PMLR.
10. Nnadili M, Okafor AN, Olayiwola T, Akinpelu D, Kumar R, Romagnoli JA. Surfactant-specific ai-driven molecular design: integrating generative models, predictive modeling, and reinforcement learning for tailored surfactant synthesis. Ind. Eng. Chem. Res. 63:6313-6324 (2024). https://doi.org/10.1021/acs.iecr.4c00401
11. Ricardo F, Ruiz-Puentes P, Reyes LH, Cruz JC, Alvarez O, Pradilla D. Estimation and prediction of the air–water interfacial tension in conventional and peptide surface-active agents by random forest regression. Chemical Engineering Science 265:118208 (2023). https://doi.org/10.1016/j.ces.2022.118208
12. Hödl SL, Hermans L, Dankloff PFJ, Piruska A, Huck WTS, Robinson WE. Surfpro – a curated database and predictive model of experimental properties of surfactants. Digital Discovery 4:1176-1187 (2025). https://doi.org/10.1039/d4dd00393d
13. Pedregosa, F., et al., *Scikit-learn: Machine learning in Python.* the Journal of machine Learning research, 2011. **12**: p. 2825-2830.
14. Landrum, G., RDKit: Open-source cheminformatics. 2006. Google Scholar, 2006.
15. Kipf, T.N. and M. Welling, *Semi-supervised classification with graph convolutional networks.* arXiv preprint arXiv:1609.02907, 2016.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A CFD Analysis of Dielectric Fluid Performance in Thermal Management of Li-ion Cells

Margarita G. Correa-Ibarra[a*], Jorge A. Alfaro-Ayala[a], Jose de J. Ramirez-Minguela[a], Zeferino Gamiño-Arroyo[a], and Agustin R. Uribe-Ramirez[a]

[a] University of Guanajuato, Department of Chemical Engineering, Mexico, Guanajuato, Guanajuato
* Corresponding Author: mg.correaibarra@ugto.mx.

ABSTRACT

The rapid electrification of the energy and transport sectors has increased the demand for efficient thermal management of Lithium-Ion Batteries (LIBs). LIBs operate optimally within a narrow temperature range (15–40 °C), effective Battery Thermal Management Systems (BTMS) are essential to prevent accelerated aging and performance degradation. Immersion cooling using dielectric fluids has emerged as a promising solution; however, a systematic comparison of their thermal and hydraulic performance remains limited. This study evaluates six dielectric fluids—Silicone Oil, Mineral Oil, Sunflower Oil, AmpCoolAC-100, E5-TM410, and Deionized Water—using Computational Fluid Dynamics (CFD) simulations of a lithium-ion battery pack composed of 32 fully immersed 18650 cells. The analysis focuses on key thermophysical properties, particularly the Prandtl number and thermal diffusivity, as criteria for coolant selection. The results demonstrate a strong correlation between cooling performance and the Prandtl number. Fluids with lower Prandtl numbers provide superior heat dissipation and lower hydraulic losses. Deionized water exhibited the best overall performance, maintaining maximum cell temperatures at 31.04 °C under a 2 LPM flow rate and achieving up to 31.33% better cooling than silicone oil. In contrast, high-viscosity oils showed poor thermal performance and significantly higher pressure drops. This work establishes the Prandtl number as a critical and reliable parameter for selecting dielectric fluids in immersion-based BTMS, highlighting deionized water as the most effective coolant among those evaluated.

Keywords: Battery Thermal Management., Computational Fluid Dynamics (CFD)., Dielectric Fluids., Prandtl Number., Immersion Cooling

INTRODUCTION

The urgent need for innovative solutions lowering the environmental impact of energy and transport sectors is leading to an unprecedentedly fast adoption rate of electrification [1]. Consequently, Lithium-Ion Batteries (LIBs) have been extensively adopted, whose choice is mainly driven by their high energy density (reaching 300 Wh/kg), efficiency (as high as 98%), lifetime (8-10 years) [2]. Lithium-ion batteries (LIBs) have established themselves as the primary energy storage technology for electric vehicles (EVs), grid-scale systems, and high-performance electronics due to their high power density, energy density, and extended cycle life. Despite these advantages, LIBs are highly sensitive to temperature, with their performance, safety, and reliability depending heavily on maintaining a narrow operational range, typically between 15°C and 40°C [3]. Operating outside these limits—specifically during high-current charging or rapid discharging—can result in accelerated aging, capacity degradation, and dangerous thermal runaway events [4].

To address these challenges, Battery Thermal Management Systems (BTMS) are employed to regulate peak temperatures and ensure uniform temperature distribution across individual cells and entire modules [3]. While conventional cooling strategies, such as forced air cooling and indirect liquid cooling, an emerging and highly effective alternative is direct liquid immersion cooling, where battery cells are submerged in an electrically non-conductive dielectric fluid [5].

While current literature has explored various categories of liquids—such as hydrofluoroethers, hydrocarbons, silicone oils, and esters—most studies focus on testing specific commercial products rather than establishing a systematic selection framework. For instance, research on deionized water (DELC-1) has shown temperature rise below 2.2 °C during 3C discharge with minimal parasitic power consumption [11]. Other studies highlight the "thermal buffer" created by low-boiling fluids like 3M Novec 7000, which utilizes latent heat to maintain cell temperatures near 34 °C even under aggressive 10C cycling [4].

This project aims to conduct a rigorous computational fluid dynamics (CFD) analysis of the thermal management of a lithium-ion battery pack. The study will focus on evaluating the cooling behavior of various dielectric fluids. The goal is to identify the best fluid that ensures the best cooling of the cells. Unlike other studies that analyze dielectric fluids without a robust foundation, this project adopts a strategic selection approach. It prioritizes comparative analysis of the thermophysical properties of different dielectric fluids to identify the most option for the thermal management of lithium-ion battery packs. This study is based on a previously reported cell pack geometry, which will allow us to validate existing simulation models or, alternatively, propose concrete improvements for cooling.

METHODOLOGY

Governing Equations

The numerical model of the battery pack cooling system is based on the fundamental principles of mass, momentum, and energy conservation, which are governed by the following equations. The fluid is considered to be an incompressible, Newtonian flow, and the thermal properties of the fluids are assumed to be constant.

The mass conservation equation, Eq. (1), describes how mass is conserved within the system.

$$\partial p/\partial t + \nabla \bullet (\rho v) = 0 \quad (1)$$

The momentum conservation equation, also known as the Navier-Stokes equation, Eq. (2), describes the motion of a fluid. It relates the forces acting on the fluid, such as pressure and viscosity, to changes in its velocity.

$$\rho(\partial v/\partial t + v \bullet \nabla v) = -\nabla P + \mu \nabla^2 v + \rho g \quad (2)$$

Finally, the energy conservation equation, Eq. (3) is utilized to model the thermal behavior of the system where S_T is heat generation of battery pack. This equation is crucial for understanding the heat transfer between the batteries and the dielectric fluid.

$$\rho Cp(\partial T/\partial t + v \bullet \nabla T) = -\nabla \bullet (k \nabla T) + S_T \quad (3)$$

Model-Geometry

The present study investigated the thermal performance of a Li-ion battery pack cooled by immersion in different dielectric fluids. A numerical model was developed and simulated using the commercial software ANSYS FLUENT. The model geometry, Figure 1, based on the work [6], consisted of a pack of 32 18650-type cells arranged with a 4 mm spacing in a total immersion cooling system. The manifold for the cooling system was designed with three inlets and three outlets, each featuring a diameter of 15 mm and a height of 25 mm.

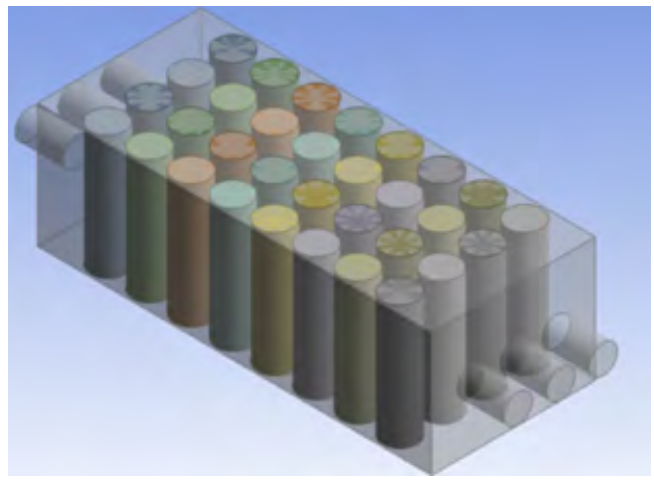

Figure 1. Geometry of pack Li-ion batteries analyzed.

Mesh Sensitivity Analysis

A mesh sensitivity analysis was conducted using the E5-TM410 dielectric fluid. The purpose of this analysis was to determine the best mesh density, balancing computational cost with solution accuracy. We focused on a range of mesh sizes between 1.55 mm and 3.75 mm, as this range consistently demonstrated stable and reliable behavior, thus justifying its selection for the rest of the study.

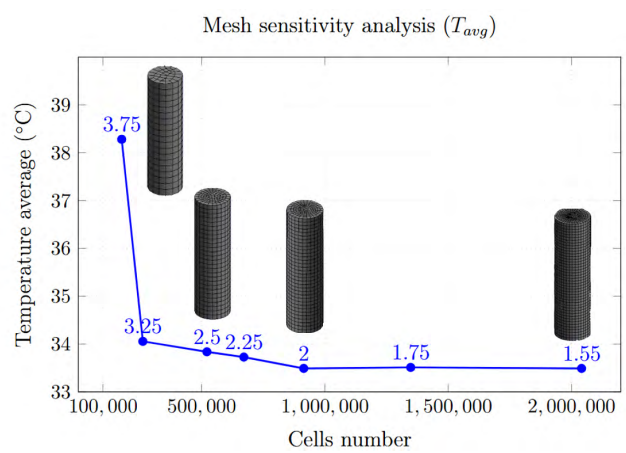

Figure 2. Mesh sensitivity analysis showing the battery temperature average as a function of mesh cell count.

Table 1. Thermal properties of dielectric fluids

Fluid	Viscosity (kg/m s)	Density (kg/m³)	Cp (J/kg K)	K (W/m K)	Reference
Deionized water (DW)	0.0009982	998.2	4182	0.6	[7]
Sunflower (SF)	0.035672	910	1800	0.175	[8]
AmpCool AC-100 (AC)	0.00799646	811	2203.2	0.1373	[7]
E5-TM410 (E5)	0.015617	805	2100	0.14	[6]
Mineral oil (MO)	0.05152	920	1900	0.13	[7]

As shown in Figure 2, the maximum battery temperature approaches an asymptotic value beyond a cell count of approximately 700, 000, indicating a stable, mesh-independent solution. Based on these findings, a mesh size of 2.0 mm was selected for the full study, as it provided an optimal balance. The corresponding mesh, consisting of 914, 140 cells, yielded accurate results without the excessive computational burden of finer meshes. The error analysis further confirmed this choice, showing a maximum error of 14.3% for the coarsest mesh (3.75 mm). Notably, the error stabilized at zero for both the 2 mm and 1.55 mm meshes.

Dielectric Fluids Selection

To ensure a rigorous selection of dielectric fluids, we began by analyzing their thermophysical properties to determine their capacity for heat removal from the battery cells. We evaluated the properties of six fluids reported in different studies like their optimal fluids see Table 1, focusing initially on viscosity. The initial hypothesis was that a fluid with high viscosity would impede heat removal from the batteries, even with a powerful pumping system, which would also be costly. This preliminary analysis led us to conclude that a low viscosity is a critical requirement. A similar analysis was performed for density. Subsequently, we focused on two others crucial thermophysical properties: heat capacity (Cp) and thermal conductivity (k). For a fluid to be a coolant, it must possess high values for both Cp and k, as these properties directly influence its ability to exchange heat with another body.

The relationship between these properties was further explored through two dimensionless parameters: the Prandtl number (*Pr*), as defined in Eq. (4), and thermal diffusivity (α), shown in equation Eq. (5). The analysis revealed that high thermal diffusivity is desirable for effective cooling. Conversely, a low Prandtl number is also indicative of a coolant. This is because a fluid with both high diffusivity and low viscosity will have a low Prandtl number, which correlates with superior cooling performance.

$$Pr = \frac{\mu}{\alpha} \tag{4}$$

$$\alpha = \frac{k}{Cp\rho} \tag{5}$$

Table 2 shows the Prandtl number (Pr) and thermal diffusivity values for the previously selected dielectric fluids. These values reveal a clear hierarchy, with deionized water exhibiting the lowest Pr number (Pr ≈ 7), and silicone oil presenting the highest (Pr ≈ 8808). This significant difference in Pr suggests a wide range of potential cooling performance. Based on this trend, we simulated the six fluids to validate the hypothesis. If the Prandtl number serves as a reliable indicator of cooling performance, we expect the cooling efficiency of the fluids to follow the same order as their Pr ranking.

Table 2. Prandtl number ranking of dielectrics fluids.

Fluid	α (m²/s)	Pr	Ranking
Deionized water	1.44×10^{-7}	6.96	1
AmpCool AC-100	7.68×10^{-8}	128.32	2
E5-TM410	8.28×10^{-8}	234.26	3
Sunflower	1.07×10^{-7}	366.91	4
Mineral oil	7.44×10^{-8}	752.98	5
Silicone oil	1.13×10^{-7}	8807.95	6

Boundary Conditions

The simulation was initialized under transient conditions with a time step of 720 seconds and was terminated once convergence was achieved. Heat generation was modeled using Eq. (6), within a numerical domain previously analyzed [6].

$$Q_b = 0.0013t^3 - 0.97t^2 + 143.12t + 293016 = 0 \tag{6}$$

The thermal properties of the batteries, presented in Table 3, were considered constant throughout the simulation. To accurately represent the thermal behavior of the cells, a thermal conductivity matrix was incorporated. This approach accounted for the anisotropic nature of

the material by specifying thermal conductivity in the axial (k_{zz}=19.03), radial (k_{yy}=1.07), and tangential (k_{xx}=19.03) directions, effectively modeling the thermal conductivity tensor.

Table 3. Properties of the Battery Pack Simulation

Thermal property	Value
Density (kg/m^3)	2055.1
Specific heat (J/kg K)	1129.95
Thermal conductivity (W/m K)	Anisotropic

Table 4. Boundary conditions of the fluids.

Boundary Conditions	Value
Inlet Temperature (°C)	25
Volumetric Flow Rates	2, 3, 4 and 5
Wall Slip Condition	Non-Slip

The boundary conditions for the fluids, as detailed in Table 4, were applied to the six fluids analyzed in this study. The investigation was conducted at four different volumetric flow rates (LPM) to evaluate the influence of fluid velocity on battery cooling. The thermophysical properties for each fluid are listed in Table 1.

RESULTS

Model Validation

As shown in Figure 3, the simulated maximum and average temperatures exhibit excellent agreement with the values reported in the reference article [6]. The maximum temperature data showed a maximum error of 2.25% at a flow rate of 5 LPM, while the minimum error was 0.49% at 3 LPM. Similarly, the average temperature data demonstrated a comparable trend, with a maximum error of 1.3% and a minimum error of 0.35%.

All percentage errors are well within the acceptable range for this type of analysis, these results successfully validate the numerical model against the published reference data.

Figure 4 presents the simulation results for the six dielectric fluids, clearly indicating deionized water delivered the best cooling performance. It maintained the maximum battery temperature within the optimal Li-ion operating range (15°C–40°C) across all volumetric flow rates, achieving a value of 30.04°C even at the lowest flow of 2 LPM. In contrast, silicone oil was the least effective, only maintaining a temperature within the safe range at the highest flow rate of 5 LPM (39.3°C).

As demonstrated, deionized water (DW) was the only coolant capable of maintaining the battery temperature within the desired operational range across all tested volumetric flow rates. In contrast, the other fluids—E5-TM410, AmpCoolAC-100, Sunflower oil, and mineral oil—only maintained temperatures within the acceptable range at higher volumetric flow rates of 4 and 5 LPM. Furthermore, even under these conditions, they operated at the upper thermal limit of the range, rendering them less reliable due to the heightened risk of temperature fluctuations exceeding the safe operational threshold.

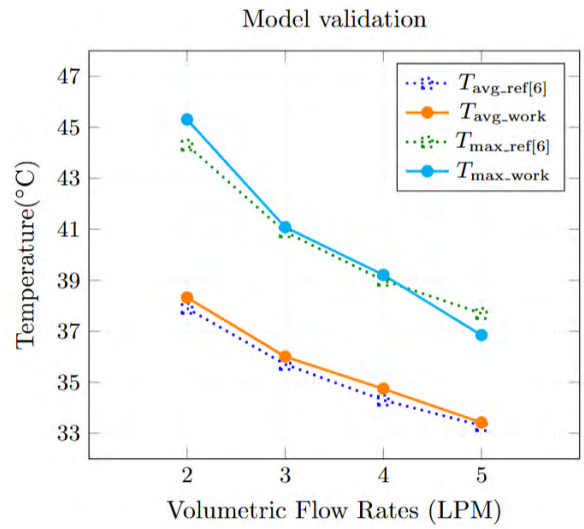

Figure 3. Thermal behavior of batteries using the E5-TM410 fluid for model validation.

Thermal Behavior of the Dielectric Fluids

At a reduced flow rate of 2 LPM, the performance difference among the coolants became more pronounced. Silicone oil exhibited the most significant thermal deviation, exceeding the maximum temperature limit by 17.7% (47.11°C), followed by mineral oil (14.66%, 45.86°C), sunflower oil (13.38%, 45.35°C), E5-TM410 (13.26%, 45.30°C), and AmpCoolAC-100 (11.76%, 44.71°C). Given that all tested fluids exceeded the safe operational threshold at this flow rate, it is deemed unsuitable for operation. In contrast, a flow rate of 5 LPM is recommended, as it maintained all fluids within the prescribed temperature range. Consequently, silicone oil was identified as the least effective coolant, maintaining the battery within the operational range only at the maximum flow rate of 5 LPM. These findings unequivocally designate silicone oil as unsuitable for this battery cooling application under most operational flow conditions.

The distribution of temperature within the battery pack provides direct insight into the effectiveness of heat transfer between the dielectric fluids and the batteries. As previously analyzed, deionized water provides superior cooling performance compared to silicone oil. These findings are further supported by the thermal behavior observed at a flow rate of 5 LPM, as illustrated in Figure 5a-c. The results compare the temperature profiles of three dielectric fluids selected to represent the top, median, and bottom rankings of the proposed methodology.

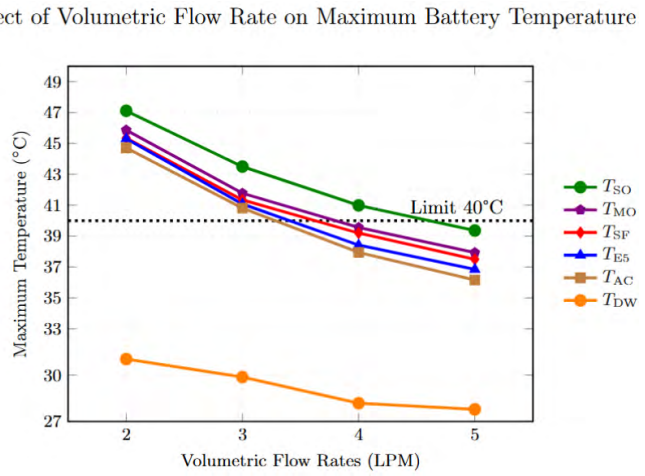

Figure 4. Relationship between the volumetric flow rate (LPM) and the maximum battery temperature for the different dielectric fluids.

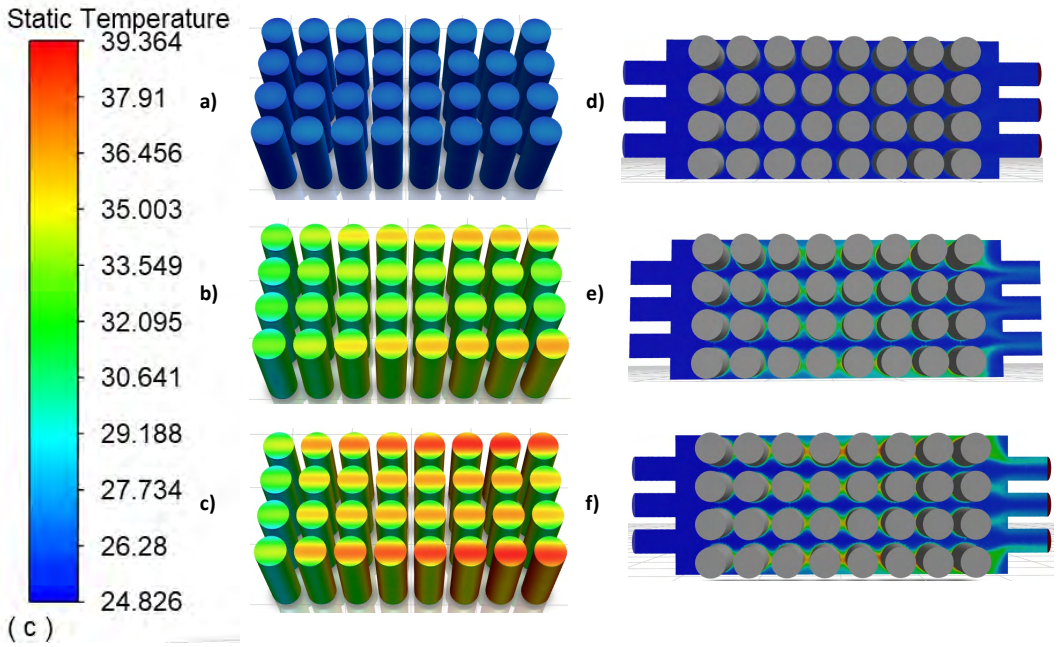

Figure 5. Comparison of the thermal distribution within the battery pack for three representative coolants selected through the proposed methodology: (a) deionized water (top-ranked), (b) E5-TM410 fluid (intermediate performance), and (c) silicone oil (lowest-ranked). Panels (a–c) show the temperature distribution across the battery cells, while panels (d–f) illustrate the corresponding temperature fields of the coolant flow within the pack for (d) deionized water, (e) E5-TM410, and (f) silicone oil.

Notably, the highest battery temperature was recorded when using silicone oil, confirming its lower performance relative to the other candidates. A closer inspection of an individual cell reveals evident anisotropic thermal diffusion. This phenomenon is a direct consequence of the simulated thermal conductivity varying with direction (radial, axial, and tangential).

The overall temperature distribution across the battery pack is non-uniform. Higher temperatures are consistently observed in the downstream and peripheral cells in the six fluids. Achieving a uniform temperature is critical for prolonging the cycle life of each cell. Thermal non-uniformity can lead to the premature degradation and eventual failure of individual cells, which necessitates the costly replacement of the entire battery pack.

To mitigate this issue, the implementation of a system that promotes a more uniform thermal distribution is essential. A potential solution is the use of turbulence promoters, which can disrupt the boundary layer and enhance heat transfer throughout the pack of batteries.

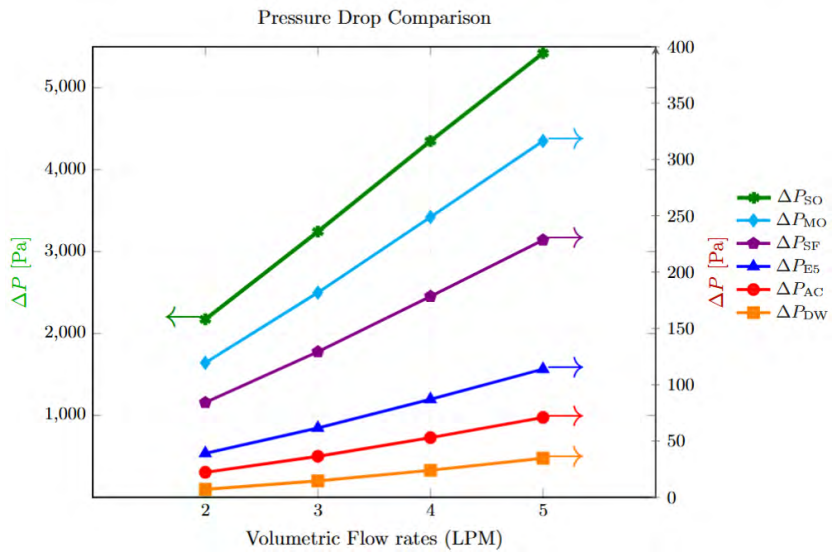

Figure 6. Pressure Drop comparison: Silicone Oil (left axis) vs. Others fluids (right axis).

These promoters could help transition the flow from a laminar to a turbulent regime, improving cooling efficiency at the implemented volumetric flow rates. The implementation of turbulence promoters is deemed necessary for this simulation. Promoters enhance fluid velocity and disrupt the boundary layer, which facilitates more effective heat removal from the batteries. This is crucial because, without them, the fluid's capacity to absorb heat diminishes as it flows through the pack, leading to higher temperatures in the downstream and peripheral cells.

This behavior is corroborated by the fluid temperature distribution shown in Figure 5d-f. The figure indicates that fluid temperature is significantly higher in the downstream cells compared to the inlet. This can be attributed to the fluid's initial momentum and impact with the first row of batteries, which promotes local turbulence and more effective cooling. As the fluid progresses, its momentum is dissipated and its heat absorption capacity diminishes, resulting in a less efficient heat exchange with the subsequent cells.

Further evidence for this behavior is provided by the fluid outlet temperatures at a flow rate of 5 LPM, the higher outlet temperatures observed in the more viscous fluids, such as silicone oil (26.34°C), mineral oil (26.01°C) and sunflower (26.02°C), particularly at low volumetric flow rates, indicate a less efficient heat exchange. This suggests that the longer residence time of these fluids in the battery pack allows for a greater increase in temperature as they absorb heat, which is less effective for overall thermal management.

The pressure drop across the battery pack, detailed in Figure 6, exhibits a direct correlation with fluid viscosity. As expected, the highly viscous silicone oil (μ= 0.96 kg/m·s) produced the most significant pressure drop, measuring 5422.29 Pa at 5 LPM—a value approximately 156 times greater than that of deionized water. Mineral oil (μ= 0.05 kg/ms), the second most viscous fluid, followed with a pressure drop of 331.42 Pa, which is about 9.12 times that of water. This substantial difference underscores the profound impact of even incremental viscosity increases on hydraulic resistance. The remaining fluids exhibited progressively lower pressure drops: sunflower oil at 228.4 Pa (6.58x DW), E5-TM410 at 113.83 Pa (3.2x DW), and AmpCoolAC-100 at 70.9 Pa (2.04x DW).

This behavior is a direct consequence of increased internal fluid friction and flow resistance at higher viscosities and volumetric flow rates. A high-pressure drop is a critical design drawback, as it necessitates a more powerful pumping system, thereby increasing the auxiliary energy consumption and reducing the net energy available for vehicle propulsion. While a turbulent flow regime—potentially enhanced by turbulence promoters—is desirable for improved cooling performance, its implementation would further exacerbate the pressure drop. Therefore, a strategic balance must be struck. It is suggested that turbulence promoters are most effectively deployed with low-viscosity fluids. This approach would enhance heat transfer while maintaining the associated pressure drop at a manageable level, avoiding the excessive energy costs incurred with high-viscosity coolants like silicone oil.

Hypothesis Validation

The simulation results validate the initial hypothesis that the Prandtl number (Pr) is a good parameter for selecting effective dielectric cooling fluids. As illustrated in Figure 4, the cooling performance of the fluids exhibits a strong inverse correlation with their Prandtl number. Deionized water, possessing the lowest Pr, delivered the best cooling performance, followed in descending order

of performance by AmpCoolAC-100, E5-TM410, Sunflower oil, mineral oil, and finally, silicone oil with the highest Pr. This trend is further corroborated by the data in Figure 7. At a flow rate of 5 LPM, silicone oil (Pr ≈ 8807.95) reached a maximum battery temperature of 39.3°C, approaching the upper safe operating limit, whereas deionized water (Pr ≈ 6.96) maintained temperatures well within the optimal range.

This analysis confirms that the Prandtl number effectively integrates the influence of key thermophysical properties, as it is defined as the ratio of momentum diffusivity (kinematic viscosity, μ) to thermal diffusivity (α), where α is itself a function of thermal conductivity (k), density (ρ), and specific heat (Cp). While high thermal conductivity and specific heat capacity are crucial for effective heat transfer, they are insufficient if not considered in conjunction with viscosity. Fluids with favorable k and Cp values but high viscosity will exhibit a high Prandtl number, which significantly impedes heat transfer by restricting fluid flow. This underscores the necessity of a balanced thermophysical profile, where high thermal diffusivity (enabling rapid heat propagation) is coupled with low viscosity (ensuring efficient flow) for superior thermal management.

Furthermore, a quantitative analysis was conducted to compare the cooling system of each dielectric fluid against the top-performing coolant, deionized water. This comparison was based on the average battery temperatures recorded at a flow rate of 5 LPM. The results demonstrate that deionized water provides superior cooling, outperforming silicone oil by 31.33%, mineral oil by 28.75%, sunflower oil by 25.45%, E5-TM410 by 24.73%, and AmpCoolAC-100 by 22.42%.

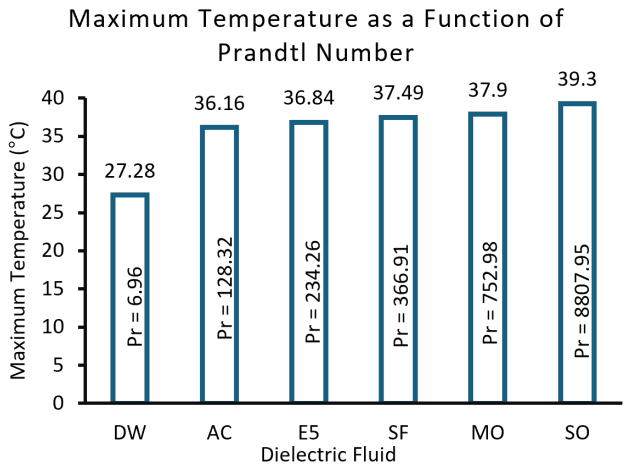

Figure 7. Comparison of Coolants: Maximum Temperature as a Function of Prandtl Number

Finally, the simulation results robustly support the hypothesis that a low Prandtl number and high thermal diffusivity are optimal for battery cooling. Deionized water is unequivocally the superior coolant among those tested, as it simultaneously excels in minimizing the maximum battery temperature, ensuring thermal uniformity, and operating with the highest hydraulic efficiency.

CONCLUSION

The analysis of the battery pack's thermal management through immersion in different dielectric fluids has conclusively validated the hypothesis that the thermophysical properties of the fluids, particularly the Prandtl number (Pr), are decisive for cooling. The simulation results confirmed that deionized water, with the lowest Prandtl number, provided the most superior cooling performance, maintaining maximum battery temperatures within the safe operational range of 20-40°C and outperforming silicone oil by 31.33%.

Furthermore, the analysis of pressure drop revealed a key disadvantage of the more viscous fluids. The high pressure drop observed in silicone oil implies a greater energy demand for the pumping system, which compromises the vehicle's overall energy efficiency. In contrast, deionized water achieved exceptional cooling with minimal pressure drop, positioning it as the most promising option for a high-performance and energy-efficient thermal management system.

The uniform thermal management of battery packs is critical for maintaining their performance and lifespan. In a further work, a key area for improvement lies in achieving a more uniform temperature distribution across the battery cells. We propose the implementation of turbulence promoters, cooling fins integrated into the battery cells, or a combined arrangement of both within the cooling channels. These promoters are designed to alter the fluid's flow patterns, transitioning from a laminar to a more turbulent regime. This change in flow behavior enhances the convective heat transfer coefficient, leading to a more effective heat exchange and, ultimately, a more uniform and efficient cooling of the batteries.

ACKNOWLEDGEMENTS

This work was supported by *Secretaría de Ciencia, Humanidades, Tecnología e Innovación* (SECIHTI) through the graduate scholarship awarded to author Correa [No. CVU: 2082437]. The authors express the support to the Sistema Nacional de Investigadoras e Investigadores (SECIHTI). The author also express their gratitude to Dr. Fernando Israel Gómez Castro for his technical guidance and for encouraging the presentation of this work at this international venue.

AUTHOR IDENTIFIERS

Author ORCIDs:
Correa Ibarra MG: 0009-0003-6487-0718

Alfaro Ayala JA: 0000-0003-3081-282X
Ramirez Minguela JJ: 0000-0003-1921-5864
Gamiño Arroyo Z: 0000-0002-6228-8586
Uribe Ramirez AR: 0000-0002-3768-5301

REFERENCES

1. Muratori M, Mai T. The shape of electrified transportation. Environ. Res. Lett. 16:011003 (2020). https://doi.org/10.1088/1748-9326/abcb38
2. Wen J, Zhao D, Zhang C. An overview of electricity powered vehicles: lithium-ion battery energy storage density and energy conversion efficiency. Renewable Energy 162:1629-1648 (2020). https://doi.org/10.1016/j.renene.2020.09.055
3. Venkateswarlu B, chavan S, Joo SW, Metwally ASM. Numerical study on cooling enhancement in lithium-ion battery modules using hybrid nanoliquids in flow channels. International Journal of Hydrogen Energy 112:231-242 (2025). https://doi.org/10.1016/j.ijhydene.2025.02.399
4. Giammichele L, D'Alessandro V, Falone M, Ricci R. Preliminary analysis of a novel battery thermal management system based on a low boiling dielectric fluid. IOP Conf. Ser.: Earth Environ. Sci. 1106:012017 (2022). https://doi.org/10.1088/1755-1315/1106/1/012017
5. Nitin M. *CFD Simulation of Thermal Management System (Immersion Cooling) of Lithium Ion Batteries in EVs* (2021). https://doi.org/10.13140/RG.2.2.19948.69767
6. Han JW, Garud KS, Kang EH, Lee MY. Numerical study on heat transfer characteristics of dielectric fluid immersion cooling with fin structures for lithium-ion batteries. Symmetry 15:92 (2022). https://doi.org/10.3390/sym15010092
7. Jithin KV, Rajesh PK. Numerical analysis of single-phase liquid immersion cooling for lithium-ion battery thermal management using different dielectric fluids. International Journal of Heat and Mass Transfer 188:122608 (2022). https://doi.org/10.1016/j.ijheatmasstransfer.2022.122608
8. Gutierrez CM, Fernandez AO, Renedo Estebanez CJ, Salas CO, Maina R. Understanding the ageing performance of alternative dielectric fluids. IEEE Access 11:9656-9671 (2023). https://doi.org/10.1109/access.2023.3239895
9. Roe C, Feng X, White G, Li R, Wang H, Rui X, Li C, Zhang F, Null V, Parkes M, Patel Y, Wang Y, Wang H, Ouyang M, Offer G, Wu B. Immersion cooling for lithium-ion batteries – a review. Journal of Power Sources 525:231094 (2022). https://doi.org/10.1016/j.jpowsour.2022.231094

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Simulation of Fixed-Bed Reactor System for Combined Ca–Cu Chemical Looping with Integrated Combustion and CO_2 Capture

Levente Biró, Norbert-Botond Mihály, Ana-Maria Cormoș*

Babeș – Bolyai University, Department of Chemical Engineering, Cluj-Napoca, Romania
* Corresponding Author: ana.cormos@ubbcluj.ro

ABSTRACT

As greenhouse gas emissions accelerate global warming, new capture and storage technologies are essential for reducing the industrial CO_2 concentration in the atmosphere. This study addresses the urgent need for greenhouse gas capture technologies by developing a detailed dynamic mathematical model for Chemical Looping Process with Integrated Combustion and CO_2 capture (CL-ICCC). In the CL-ICCC process configuration, CO_2 capture is integrated into the chemical looping combustion system, resulting in a higher-purity, more efficient process. In this work Cu/CuO oxygen carrier material and $CaO/CaCO_3$ sorbent materials were considered in a fixed bed reactor as solid phase to investigate Oxidation and Reduction/Calcination processes under different operating conditions. The simulation results were compared with experimental results from the literature. In case of the oxidation process, a sensitivity study was performed to investigate the behavior of the process for variation of different operating parameters under dynamic conditions: gas composition, initial gas temperature and solid phase mass composition. In the case of reduction/calcination, multiple fuel type and solid composition were investigated.

Keywords: Carbon Dioxide Capture, Chemical Looping, Fixed-Bed Reactor, Hydrogen generation, Sensitivity Study

INTRODUCTION

In the last years, the level of greenhouse gases (GHG) increased drastically, accelerating the process of global warming by raising the average temperature. [1] To create a pollution-free world, it is impossible to minimize all major pollution sources, such as industrial, transportation or agricultural GHG emissions. New GHG capture technologies, specifically CO_2 capture methods, that can have a negative, reducing effect on overall greenhouse gas levels are the solution. [2] Chemical Looping Combustion (CLC) is one of these new types of solutions and it can be adapted in multiple ways. The main purpose is to reduce GHG level in flue gas after the industrial processes. The working principle of the process is very simple: the process uses two different reactors: an Air Reactor (AR) where an oxidation process takes place between a metal or metal-oxide (known as Oxygen Carrier (OC)) and a Fuel Reactor (FR) where the combustion process can be performed in the absence of air. [3-5] An upgraded version of CLC process is the Chemical Looping with Integrated Combustion and CO_2 Capture (CL-ICCC) process. In this case the CO_2, which will be captured, is generated during the main process and that can increase the efficiency of the plant, also minimizing the climate change effect. [6] However, the products of the CL-ICCC process can be stored and used as raw materials for different processes, depending on the market and the plant operating profile. The main products of the CL-ICCC plant are pure H_2, CO_2 gas and electric power via steam. [7-9] The flowsheet of the presented, complex chemical process can be seen in Figure 1. All chemical processes take place in a single fixed-bed reactor. A complex valve system, illustrated in Figure 1, is included to switch between the different gaseous raw materials. In this study a copper-based OC, more specifically Cu/CuO material is used. For the CO_2 capture process the Ca-based compound is an optimal and efficient choice from an economic point of view. [9] The main objective

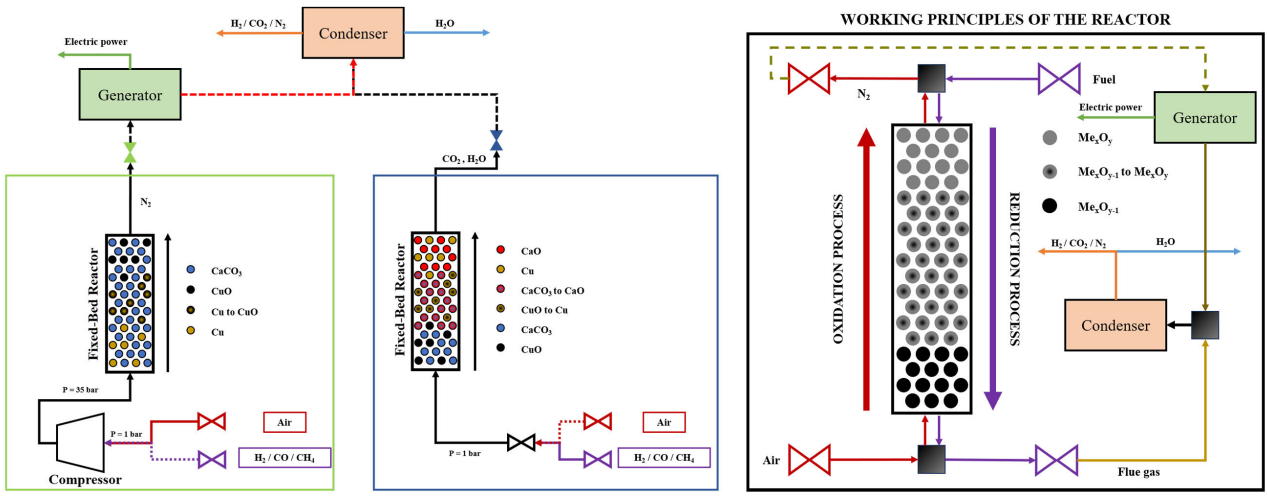

Figure 1. Schematic flowsheet of the presented CL-ICCC process (left side) and the working principles of the

of this work is to develop a dynamic mathematical model of the entire CL-ICCC process which can be divided into three different parts.

Carbonation process (A)

The H_2 synthesis occurs based on the reaction equations (1) – (3). At the same time the in-situ CO_2 capture takes place. As output of the system, the gaseous phase contains a significant quantity of H_2 gas, high-pressure steam, with fuel traces. The chemical process is performed in Fixed-Bed Reactors via Steam Enhanced Reforming (SER) and Water Gas Shift (WGS) processes [9, 10].

$$CH_4 + H_2O \rightarrow CO + 3\,H_2 \tag{1}$$

$$CO + H_2O \rightarrow CO_2 + H_2 \tag{2}$$

$$CH_4 + 2\,H_2O \rightarrow CO_2 + 4\,H_2 \tag{3}$$

$$CaO + CO_2 \rightarrow CaCO_3 \tag{4}$$

Oxidation (Regeneration) process (B)

The regeneration of the OC material (according to the equation 5) takes place in the oxidation process. The OC capability and type have a high influence on the efficiency of the carbonation process [11, 12]. During the chemical process, the oxidation stage generates a high amount of energy, and the excess energy can be used in steam generation, according to the scientific work of J. R. Fernández et al. [9].

$$2\,Cu + O_2 \rightarrow 2\,CuO \tag{5}$$

Calcination (reduction) process (C)

During the calcination process, the CO_2 gas is generated via two different ways: from the decomposition of the calcium carbonate, which was generated during the carbonation process and from the reaction between syngas and OC material [7, 11, 13]. The reactions described in equations:

$$4\,CuO + CH_4 \rightarrow 4\,Cu + CO_2 + 2\,H_2O \tag{6}$$

$$CuO + H_2 \rightarrow Cu + H_2O \tag{7}$$

$$CuO + CO \rightarrow Cu + CO_2 \tag{8}$$

$$CaCO_3 \rightarrow CaO + CO_2 \tag{9}$$

After the calcination process, the reactor remains at elevated temperature, due to the excess heat generated during the reduction stage and not consumed in the calcination. To improve overall efficiency, syngas synthesis can be performed immediately after the reduction, supplying fuel for subsequent or parallel reduction cycles.

MATERIALS AND METHODOLOGY OF THE RESEARCH

The aim of this study was to develop a dynamic mathematical model describing stages B and C, with emphasis on model validation and sensitivity analysis of key operating parameters. Although CO_2 capture and storage processes are well documented in the literature, differences arise in the choice of oxygen carrier and CO_2 sorbent, as well as in the level of model detail. J. R. Fernández et al. proposed several models for the same industrial process; however, their pseudo-homogeneous formulations simplify the temperature profile and provide only limited description of gas–solid interactions. [7-10, 12]

In stage B, a critical parameter is the stability of the solid phase under high-temperature and high-pressure conditions. Reactor temperatures may exceed 1000 °C and pressures approach 35 bar, emphasizing the need for appropriately selected support materials. To simulate the chemical process, the initial step involved validating

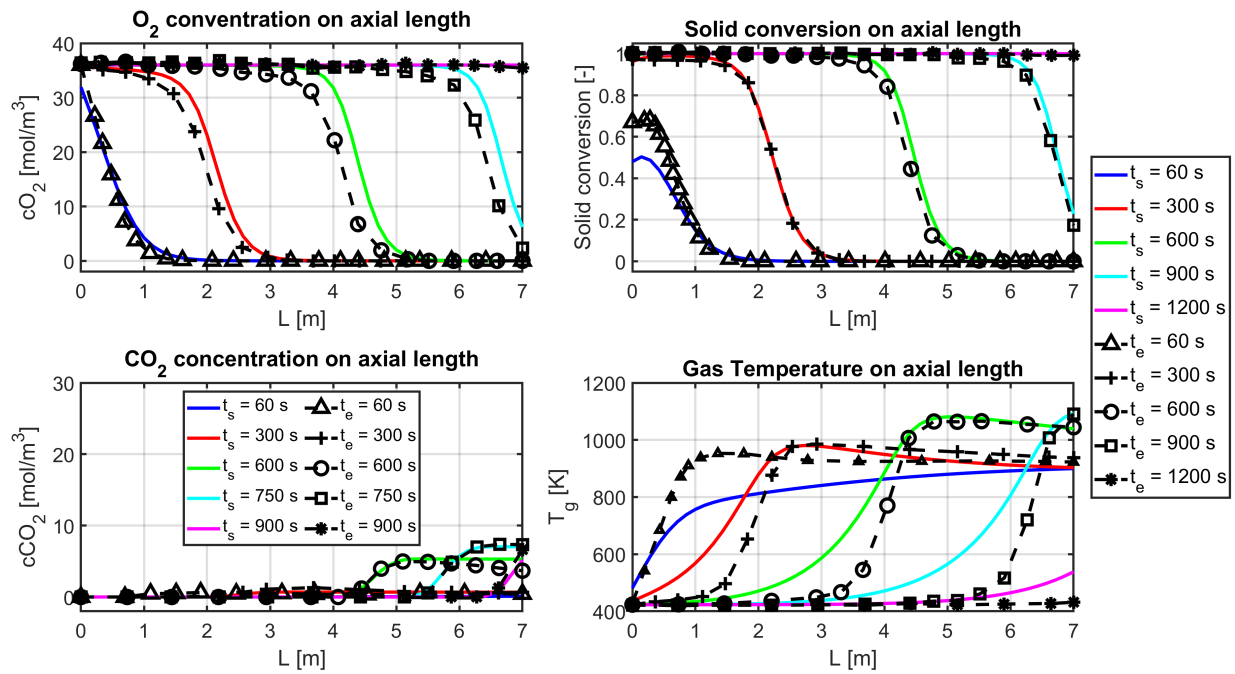

Figure 2. Validation diagrams of the stage B process (solid composition: Cu:Ca = 0.28:0.23) (simulated results vs.

the reaction kinetics under various conditions using experimental data from the literature. F. García-Labiano et al. conducted a comprehensive kinetic study on copper-based oxygen carriers, through first-order kinetic model, was chose for this work. [14]

Table 1. Kinetical parameters of Stage B and Stage C. [14, 15]

Parameters	U.M.	Values
Stage B		
$\rho_{m,Cu}$	kmol/m³	140.252
L_{Cu}	m	$4 \cdot 10^{-10}$
$k_{0,ox}$	m/s	$4.7 \cdot 10^{-6}$
$E_{a,ox}$	kJ/mol	15
Stage C		
$\rho_{m,CuO}$	kmol/m³	80.402
L_{CuO}	m	$2.3 \cdot 10^{-10}$
k_{0,H_2}	m/s	$1 \cdot 10^{-4}$
E_{a,H_2}	kJ/mol	33
$k_{0,CO}$	m/s	$5.9 \cdot 10^{-6}$
$E_{a,CO}$	kJ/mol	14
k_{0,CH_4}	m/s	$4.5 \cdot 10^{-4}$
E_{a,CH_4}	kJ/mol	60
$k_{0,red}$	kmol/(m³ • s)	252 015
$E_{a,red}$	kJ/mol	91.7

Although the previous studies provided a solid theoretical foundation, several literature sources indicate that the oxidation process is more accurately represented by a second-order kinetic model. In this formulation, key parameters include the oxygen-carrier layer thickness (L_{Cu}) and its density ($\rho_{m,Cu}$), as shown in equation (10) and (11):

$$\upsilon_{p_{ox}} = \frac{dX}{dt} = \frac{2}{\tau} \cdot (1 - X_{ox})^{0.5} \tag{10}$$

$$\tau = \frac{L_{Cu} \cdot \rho_{m,Cu}}{b \cdot k_{ox} \cdot C_g^n} \tag{11}$$

$$k_{ox} = k_{0,ox} \cdot e^{\left(-\frac{E_{a,ox}}{RT}\right)} \tag{12}$$

where X_{ox} is the conversion of the solid component, b is the stoichiometric factor of reaction, k_{ox} is the chemical reaction rate constant of the oxidation process [m/s], C_g is the gas component concentration in the reactor chamber [kmol/m³], n is the reaction order, τ is the necessary time to complete the solid conversion [s], R is the gas constant [J/mol•K]. The reaction rate constant is calculated by the Arrhenius equation (12).

For calcination, additional parameters have been defined in the literature to account for temperature effects, maximum conversion, and the equilibrium CO_2 concentration in the system, as expressed in equation (13) and (14):

$$\upsilon_{p_{cal}} = \frac{dX_{cal}}{dt} = k_{cal} \cdot (1 - X_{cal})^{\frac{2}{3}} \cdot \left(c_{CO_2,eq} - c_{CO_2}\right) \tag{13}$$

$$c_{CO_2,eq} = \frac{5.045 \cdot 10^8}{T} \cdot \exp\left(-\frac{20474}{T}\right) \tag{14}$$

$$k_{cal} = k_{0,cal} \cdot e^{\left(-\frac{E_{a,cal}}{RT}\right)} \tag{15}$$

where X_{cal} is the conversion of the solid component in the calcination process, k_{cal} is the chemical reaction rate constant [kmol/(m³•s)], c_{CO_2} is the CO_2 concentration in the reactor chamber [kmol/m³], $c_{CO_2,eq}$ is the equilibrium concentration of the CO_2 gas, calculated by equation (14), according to J. M. Alarcón et al. [13] The reaction rate constant is evaluated using the Arrhenius equation (15). The full set of kinetic parameters used in this work is summarized in Table 1.

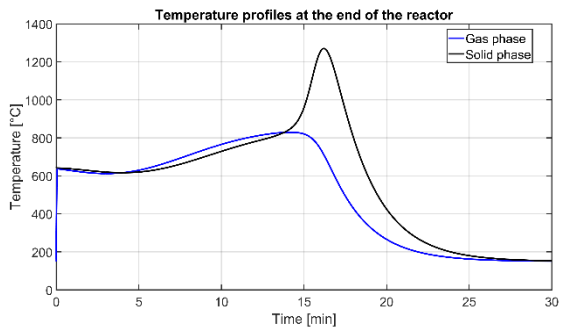

Figure 3. Temperature profiles at reactor exit are in a steady state regime.

For the model development, the governing equations were formulated using standard mass and energy balance relationships, consistent with established literature. In stage B, J. R. Fernandez et al. proposed two alternative modelling approaches, noting that the elevated temperatures reached during oxidation may induce partial calcination. [12] To examine this effect, they developed two different models: one model when the calcination reaction is not considered during the process and an other model when it is considered. In this study, the mathematical model was performed based on the second case. The components mass balance equations are presented via equations (16) – (22):

$$\varepsilon \frac{dm_{Cu}}{dt} = -(1-\varepsilon) \bullet \eta \bullet \upsilon_{p_{ox}} \bullet m_{Cu} \quad (16)$$

$$\varepsilon \frac{dm_{CuO}}{dt} = (1-\varepsilon) \bullet \eta \bullet \upsilon_{p_{ox}} \bullet m_{Cu} \bullet \frac{M_{CuO}}{M_{Cu}} \quad (17)$$

$$\varepsilon \frac{dm_{CaCO_3}}{dt} = -(1-\varepsilon) \bullet \upsilon_{p_{cal}} \bullet m_{CaCO_3} \quad (18)$$

$$\varepsilon \frac{dm_{CaO}}{dt} = -(1-\varepsilon) \bullet \upsilon_{p_{cal}} \bullet m_{CaO} \bullet \frac{M_{CaO}}{M_{CaCO_3}} \quad (19)$$

$$\varepsilon \frac{dc_{O_2}}{dt} = -w_g \bullet \frac{dc_{O_2}}{dz} - \frac{(1-\varepsilon)\bullet\eta\bullet\upsilon_{p_{ox}}\bullet m_{Cu}}{2\bullet M_{Cu}\bullet A\bullet dz} \quad (20)$$

$$\varepsilon \frac{dc_{N_2}}{dt} = -w_g \bullet \frac{dc_{N_2}}{dz} \quad (21)$$

$$\varepsilon \frac{dc_{CO_2}}{dt} = -w_g \bullet \frac{dc_{CO_2}}{dz} + \frac{(1-\varepsilon)\bullet \upsilon_{p_{cal}}\bullet m_{CaO}}{M_{CaCO_3}\bullet A\bullet dz} \quad (22)$$

where m_j is the mass of the *j* component in the solid bed (*j* = Cu, CuO, CaO, $CaCO_3$) [kg], η is an efficiency factor of the chemical reaction, ε is the solid bed porosity, c_i is the concentration of the *i* components in the reactor (*i* = O_2, N_2, CO_2) [kmol/m³], M_j is the molecular weight of the *j* component [kg/kmol], w_g is the velocity of the gas flow [m/s], A is the cross section of the reactor [m²], dz is the length of a section of the reactor (the reactor length was divide into 50 equal spaces).

The energy balance equations describe the conventional heat exchange between the solid and gaseous phases, according to the plug flow reactor operating principles. The balance equations are presented below:

$$\varepsilon \frac{dT_g}{dt} = -w_g \bullet \frac{dT_g}{dz} - \frac{K_T\bullet A_T\bullet(T_g-T_s)}{c_{p_g}\bullet\rho_g\bullet A\bullet dz} \quad (23)$$

$$\varepsilon \frac{dT_s}{dt} = \frac{K_T\bullet A_T\bullet(T_g-T_s)}{c_{p_s}\bullet\rho_s\bullet A\bullet dz} - \sum \frac{(1-\varepsilon)\bullet\eta\bullet\Delta H_x\bullet\upsilon_{p_x}\bullet m_j}{b\bullet M_j\bullet c_{p_s}\bullet\rho_s\bullet A\bullet dz} \quad (24)$$

where T_s is the temperature of the solid bed [K], T_g is the temperature of the gaseous phase [K], K_T is the heat transfer coefficient [W/m²K], A_T is the surface area where the heat exchange process takes place [m²], c_{p_s} and c_{p_g} are the specific heat capacity for the gaseous and solid phases [J/kgK], ρ_s and ρ_g are the density of the gaseous and solid phases [kg/m³], x is an indicator for the processes during the interpretation of the energy balance equations (x = oxidation or calcination).

Table 2. Process operating parameters.

Parameters	Values
Stage B	
Initial gas temperature [$T_{g,0}$]	423 K
Initial solid temperature [$T_{s,0}$]	923 K
Pressure [P]	35 bar
Inlet mass flow velocity [Q_s]	21 kg/m²s
Process effectiveness [η]	0.3
Inlet O_2 concentration [$c_{O_{2_{in}}}$]	3.6 %
Particle size [d_p]	10 mm
Porosity [ε]	0.5
Solid phase density [ρ_s]	3400 kg/m³
Cu weight fraction [ρ_{Cu}]	950 kg/m³
$CaCO_3$ weight fraction [ρ_{CaCO_3}]	750 kg/m³
Stage C	
Initial gas temperature [$T_{g,0}$]	973 K
Initial solid temperature [$T_{s,0}$]	973 K
Pressure [P]	1 atm
Inlet molar flow velocity [Q_s]	1 kmol/s
Process effectiveness [η]	0.8
Inlet fuel gas concentration [$c_{i_{in}}$]	75 %
Particle size [d_p]	5 mm
Porosity [ε]	0.5
Solid phase density [ρ_s]	2330 kg/m³

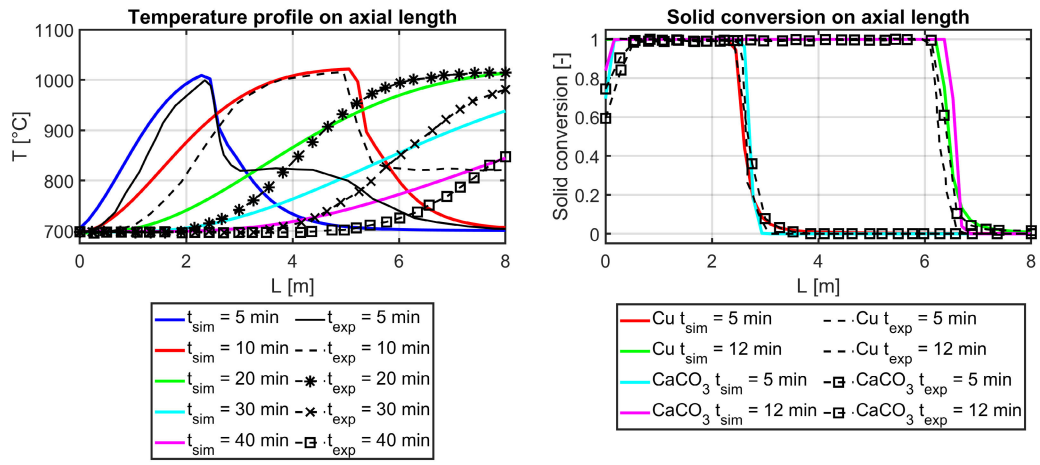

Figure 4. Validation diagrams of the stage C during H_2 process (solid composition: Cu:Ca = 0.3:0.15) (simulated results vs. experimental results)
From left to right: 1. Gas temperature profile on axial length; 2. Solid conversion on axial length;

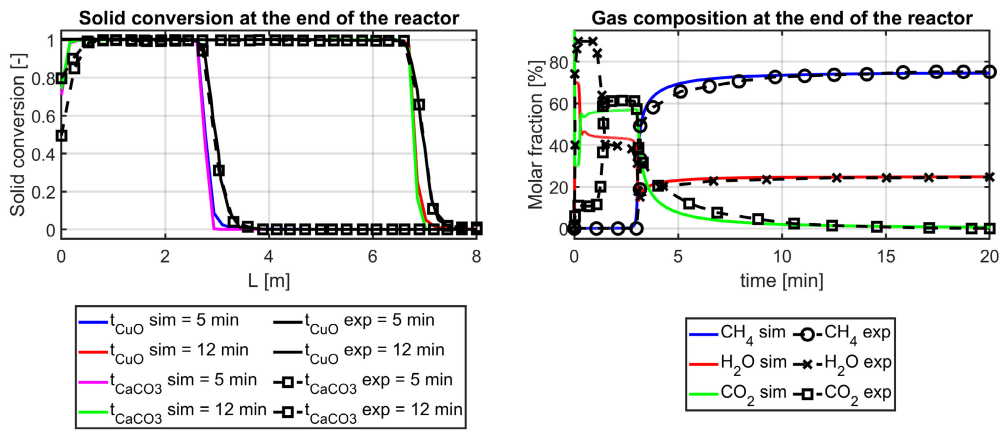

Figure 5. Validation diagrams of the stage C during CO and CH_4 process (solid composition: Cu:Ca = 0.3:0.21). (simulated results vs. experimental results)

For the calcination processes, the defined equations are the same, the differences come from the direction of the chemical reactions, from the raw material to the product, according to the reaction equations (6 - 9). The developed mathematical model is also operating in adiabatic conditions, where the energy for the calcination process is generated in situ via the reduction processes. The mass and energy balance equations are described below:

$$\varepsilon \frac{dm_j}{dt} = \pm (1-\varepsilon) \bullet \eta \bullet \upsilon_{p_{red}} \bullet m_j \tag{25}$$

$$\varepsilon \frac{dc_i}{dt} = -w_g \bullet \frac{dc_i}{dz} \pm \frac{(1-\varepsilon) \bullet \eta \bullet \upsilon_{p_{red}} \bullet m_j}{b \bullet M_j \bullet A \bullet dz} \tag{26}$$

$$\varepsilon \frac{dT_g}{dt} = -w_g \bullet \frac{dT_g}{dz} - \frac{K_T \bullet A_T \bullet (T_g - T_s)}{c_{p_g} \bullet \rho_g \bullet A \bullet dz} \tag{27}$$

$$\varepsilon \frac{dT_s}{dt} = \frac{K_T \bullet A_T \bullet (T_g - T_s)}{c_{p_s} \bullet \rho_s \bullet A \bullet dz} - \sum \frac{(1-\varepsilon) \bullet \eta \bullet \Delta H_x \bullet \upsilon_{p_x} \bullet m_j}{b \bullet M_j \bullet c_{p_s} \bullet \rho_s \bullet A \bullet dz} \tag{28}$$

where i = H_2, H_2O, CO, CO_2, CH_4.

The model simulation was performed in MATLAB/Simulink, the used solver was 'ode15s'. The operating parameters (Table 2) were chosen based on multiple literature sources. [8, 12, 13]

DYNAMIC SIMULATIONS' RESULTS AND DISCUSSION

Dynamic simulation and sensitivity study of stage B

The dynamic mathematic simulation was conducted to evaluate the systems' behavior under realistic operating conditions. The oxidation reaction concluded at approximately 15 minutes, after which the system began to cool due to heat transfer between the fresh gas stream and the oxidized solid phase. Owing to the elevated temperature, a minor extent of $CaCO_3$ decomposition occurred, by the calcination process. Since the minimum

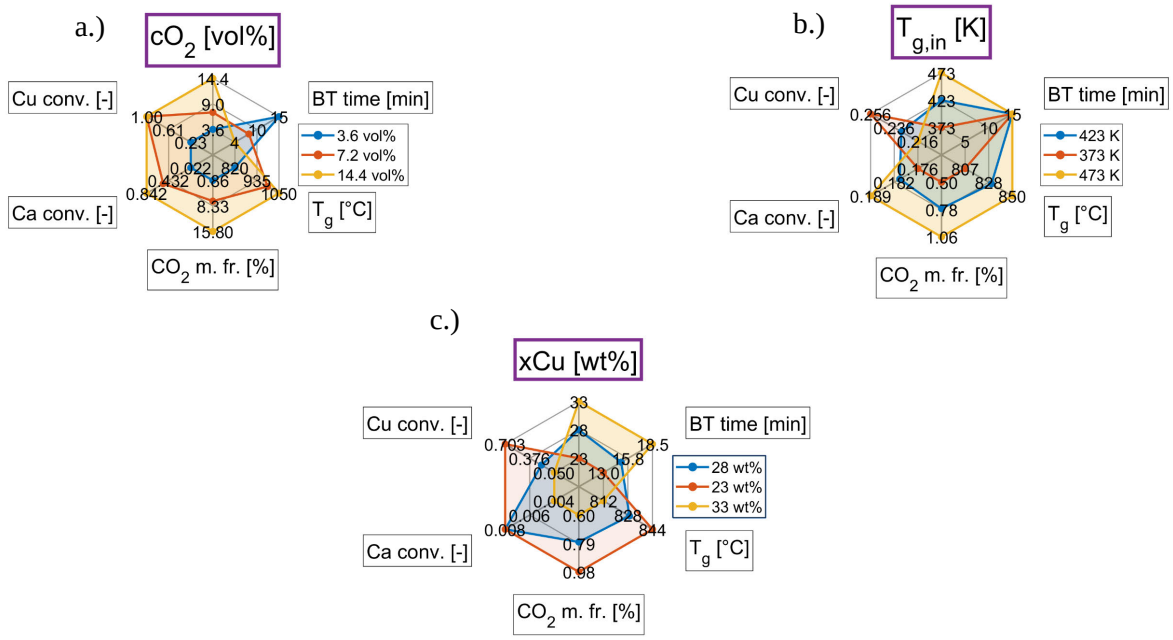

Figure 6. Spider web diagrams of the sensitivity study of the stage B (BT time is Breakthrough time): a.) Variation of gas composition [3.6 – 14.4 vol%]; b.) Variation of the initial gas temperature [373 – 473 K]; c.) Variation of the mass composition of the solid bed [23 – 33 wt% Cu];

temperature of the calcination exceeds 973 K, the peak temperature reached during stage B surpassed and threshold. The oxygen carrier achieved complete conversion at around 15 minutes, marking the breakthrough point. Beyond this point, the O_2 concentration returned to its initial level, signaling the end of the chemical reaction.

The steady state of the system within the reactor can be evaluated at two levels: the chemical composition and temperature. Figure 3 illustrates the temperature profiles of the solid and gas phases from the beginning to the breakthrough point. The temperature difference in the system is significant: the oxidation reaction generates a large amount of energy, and this energy is transferred to the gaseous phase to warm up the gas mixture to the operating temperature. After the breakthrough point, the temperature in the solid phase rises, caused by the excess heat from the oxidation process at the end of the undesired calcination process.

To investigate the system's response to variation in operating conditions, sensitivity analysis was performed. Three key parameters were examined: gas composition, initial gas temperature and mass difference in the case of Cu after 5 minutes on 2.5 m examination point. Based on the results, the hydrogen level can generate a non-favorable system behavior. The increase in O_2 concentration has a significant effect on system behavior, as illustrated in Figure 6. Higher hydrogen levels accelerate the reaction rate, resulting in shorter breakthrough times and an earlier appearance of the breakthrough curves. This intensified reaction also causes a significant temperature increase. At 14.4 vol% O_2, the peak gas-phase temperature is approximately 1, 050 °C, which is about 30% higher than under the reference operating conditions. At this point, the CO_2 concentration increases to 15.8 %, and the overall conversion of the Ca-based sorbent reaches 0.842 after 5 minutes at 2.5 m axial length. These conditions are undesirable, because the primary objective is to fully oxidize the oxygen carrier, while preserving the chemical integrity of the calcium-containing material. A lower initial gas temperature generates more desired results by higher Cu conversion and significantly lower Ca conversion. This phenomenon is caused by the temperature differences between the gas and the solid phase. The solid composition is also an important aspect of the system design and has significant impact on multiple evaluated parameters. A lower Cu composition in the solid bed causes a higher peak temperature, a shorter reaction time and a higher CO_2 level at the end of the reactor. The results indicate to estimate an optimized operating time, to decrease the appearance of the calcination in the oxidation process and maximize the Cu conversion.

Dynamic simulation of stage C

In the stage C, the chemical reaction is dictated by the composition of the fuel gas. Although the theoretical optimum operates with syngas ratio $CO:H_2$ = 1:3, trace alkanes (specifically CH_4) may be present due to valve malfunction or post-breakthrough operations. In order to develop a well-defined mathematical model, it is necessary to consider every case and perform three different

operating mods. Each operating mode was performed in isolated conditions, during the simulation steam was added to the pure fuel gas to create an inert atmosphere. The simulations were performed using the model equations (24-29), and the reactor was considered as an adiabatic system (perfectly isolated), which reduces the number of thermodynamic phenomena in the system.

The H_2 process was simulated for 80 minutes to evaluate the solid conversion, the gas composition and the temperature profiles. The breakthrough point occurred at 15 minutes when the solid phase reaches the maximum conversion level and the system reached its peak temperature. The process started with preheated gaseous phase at 973 K, which was also the temperature of the solid phase. After the fuel gas was introduced into the reactor chamber, an initial fluctuation of the gas composition occurred.

The CO process simulation duration was 30 minutes, and the breakthrough point was achieved in 12 minutes. During the process, significant amount of the CO_2 was detected, by the oxidation of the CO gas from syngas and from the calcination process for dissociation of the $CaCO_3$. In contrast with the H_2 solution, the reaction heat which comes from the CO oxidation process is much higher. The highest temperature obtained during the CO simulation was close to 1, 250°C, compared to 875°C obtained during the H_2 process.

With the CH_4 fuel gas, the simulation was performed under 20 minutes, and the end of the chemical reaction was nearly 3 minutes, which is almost 10 times shorter than the operating time of the H_2 simulation. The conversion of the $CaCO_3$ did not reach the maximum conversion level, which could be achieved by increasing the gas composition or the gas stream flow rate.

In the case of the H_2 model, the simulation results were validated for multiple mass composition by the experimental data provided by J. M. Alarcón et al. [13]

CONCLUSION

In this work, multiple mathematical models were developed successfully under dynamical operating conditions for the oxidation and calcination/reduction processes. The validation process was performed by experimental data presented in literature. During the validation process, the developed model accurately predicted the gas composition and the temperature profile at the end of the reactor, as well as the axial length and solid conversion. However, a sensitivity study was performed to examine the response of the oxidation process under dynamic operating conditions. Specifically, the elevated O_2 concentration (14.4 vol%) in the fuel gas stream indicates an undesirable calcination process. During the oxidation on an elevated peak temperature (1050°C) can produce the same unfavorable phenomenon: 0.189 Ca conversion after 5 minutes at 2.5 m axial length.

The calcination process is highly influenced by the composition and type of the fuel used. In the case of H_2 reduction, the peak temperature of the system was 875°C, while in the CO calcination, the value of the peak temperature was 1250°C. The CH_4 reduction process was performed under 3 minutes due to the large amount of energy generated via fuel gas. In practice, these three calcination types take place simultaniously in the reactor chamber in, complicates the modeling process.

ACKNOWLEDGEMENTS

This work was supported by a grant from Foundation „Sapientia Hungariae" (Hungary), through the „Collegium Talentum" Scholarship Program.

AUTHOR IDENTIFIERS

Author ORCIDs:
Biró L: 0009-0003-2971-3776
Mihály NB: 0000-0002-0126-2657
Cormoș AM: 0000-0003-0832-0722

REFERENCES

1. Beuttler C, Charles L, Wurzbacher J. The role of direct air capture in mitigation of anthropogenic greenhouse gas emissions. Front. Clim. 1: (2019). https://doi.org/10.3389/fclim.2019.00010
2. Europena Commision. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52018DC0773
3. Hossain MM, de Lasa HI. Chemical-looping combustion (CLC) for inherent 2 separations—a review. Chemical Engineering Science 63:4433-4451 (2008). https://doi.org/10.1016/j.ces.2008.05.028
4. Daneshmand-Jahromi S, Sedghkerdar MH, Mahinpey N. A review of chemical looping combustion technology: fundamentals, and development of natural, industrial waste, and synthetic oxygen carriers. Fuel 341:127626 (2023). https://doi.org/10.1016/j.fuel.2023.127626
5. Abuelgasim S, Wang W, Abdalazeez A. A brief review for chemical looping combustion as a promising CO2 capture technology: fundamentals and progress. Science of The Total Environment 764:142892 (2021). https://doi.org/10.1016/j.scitotenv.2020.142892
6. Liu G, Lisak G. Cu-based oxygen carriers for chemical looping processes: opportunities and challenges. Fuel 342:127828 (2023). https://doi.org/10.1016/j.fuel.2023.127828
7. Alarcón AM, Fernández JR, Abanades JC. Study of

a Cu-CuO chemical loop for the calcination of $CaCO_3$ in a fixed bed reactor. *Chem Eng J* 325:208-220 (2017) http://dx.doi.org/10.1016/j.cej.2017.05.070
8. Fernández JR, Abanades JC, Murillo R. Modeling of sorption enhanced steam methane reforming in an adiabatic fixed bed reactor. *Chem Eng Sci* 84:1-11 (2012) http://dx.doi.org/10.1016/j.ces.2012.07.039
9. Fernández JR, Abanades JC, Murillo R, Grasa G. Conceptual design of a hydrogen production process from natural gas with CO2 capture using a ca–cu chemical loop. International Journal of Greenhouse Gas Control 6:126-141 (2012). https://doi.org/10.1016/j.ijggc.2011.11.014
10. Fernández JR, Abanades JC, Grasa G. Modeling of sorption enhanced steam methane reforming – Part II: Simulation within a novel Ca/Cu chemical loop process for hydrogen production. *Chem Eng Sci* 84:12-20 (2012) http://dx.doi.org/10.1016/j.ces.2012.07.050
11. Noorman S, Gallucci F, van Sint Annaland M, Kuipers JAM. Experimental Investigation of chemical-looping combustion in packed beds: a parametric study. *Ind Eng Chem Res* 50:1968-1980 (2011) https://dx.doi.org/10.1021/ie1019112
12. Fernández JR, Abanades JC, Murillo R. Modeling of Cu oxidation in adiabatic fixed-bed reactor with N_2 recycling in a Ca/Cu chemical loop. *Chem Eng J* 232:442-452 (2013) http://dx.doi.org/10.1016/j.cej.2013.07.115
13. Alarcón JM, Fernández JR. $CaCO_3$ calcination by the simultaneous reduction of CuO in a Ca/Cu chemical looping process. *Chem Eng Sci* 137:254-267 (2015) http://dx.doi.org/10.1016/j.ces.2015.06.030
14. García-Labiano F, de Diego LF, Adánez J, Abad A, Gayán P. Reduction and oxidation kinetics of a copper-based oxygen carrier prepared by impregnation for chemical-looping combustion. Ind. Eng. Chem. Res. 43:8168-8177 (2004). https://doi.org/10.1021/ie0493311
15. Martínez I, Grasa G, Murillo R, Arias B, Abanades JC. Kinetics of calcination of partially carbonated particles in a ca-looping system for co₂ capture. Energy Fuels 26:1432-1440 (2012). https://doi.org/10.1021/ef201525k

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Comparison of Various Hydrogen Flux Trajectories in a Catalytic Membrane Reactor Operating Dehydrogenation of Ethylbenzene to Styrene

Nabeel S. Abo-Ghander[ab*]

[a] Department of Chemical Engineering, King Fahd University of Petroleum & Minerals, Dhahran 31261, Saudi Arabia.
[b] Interdisciplinary Center for Refining & Advanced Chemicals, King Fahd University of Petroleum & Minerals, Dhahran 31261, Saudi Arabia.
* Corresponding Author: naboghander@gmail.com.

ABSTRACT

Styrene is mainly produced by dehydrogenating ethylbenzene over an iron oxide-based catalyst. The reaction is endothermic and thermodynamically limited when operated in conventional catalytic fixed-bed reactors. This makes the styrene production process a conversion-selectivity trade-off, where different objectives must be compromised. In this work, a one-dimensional reactor model accounting for changes in molar flowrate, temperature, and pressure is used to predict the performance of a membrane reactor. Three main hydrogen flux profiles were assumed along the reactor axial direction: constant, linearly increasing, and linearly decreasing. It is found that the styrene yield and selectivity in a membrane reactor operated with a linearly decreasing hydrogen flux profile are higher than those with constant or linearly increasing hydrogen flux profiles in both isothermal and nonisothermal cases. It is also observed that the styrene yield and selectivity of the membrane reactor operated with a linearly decreasing hydrogen flux can be enhanced by improving the membrane's average hydrogen permeability. The influence of nonlinearity in the hydrogen flux equation is marginal and observed mainly at the reactor inlet. The improvement is due to enhanced hydrogen permeation, which is observed to be better for the linearly decreasing hydrogen profile than for the constant, linearly increasing hydrogen flux profile. When a permeation section is included in the reactor at the point where maximum hydrogen production is observed, the styrene yield is better than that of the completely permeated catalytic reactor.

Keywords: Reaction Engineering, Modelling and Simulations, Optimization, Hydrogen, Membranes, Hydrogen Flux

INTRODUCTION

Styrene is a crucial component for manufacturing high-performance polymers. It is mainly produced via the dehydrogenation reaction of ethylbenzene catalyzed by iron-oxide based catalyst. The reaction is endothermic and thermodynamically limited when operated in conventional fixed-bed reactors. This turns the process of styrene production into a conversion-selectivity problem where different objectives must be compromised.

Several researchers studied the kinetics of ethylbenzene dehydrogenation and its interactions with other chemicals. For instance, the influence of the presence of CO_2 in the reactor catalyzed by Fe_2O_3 supported on Al_2O_3 was addressed by Mimura and Saito, while the impact of the presence of steam in the reactor when ethylbenzene is dehydrogenated was studied by Zhu et al. via dynamic experiments in the presence of a potassium-promoted iron oxide-based catalyst [1, 2]. Rossetti et al. studied the causes of the deactivation of a commercial catalyst used for the dehydrogenation of ethylbenzene to styrene [3]. Lee and Froment studied the kinetics of the dehydrogenation of ethylbenzene to styrene catalyzed by a commercial iron oxide catalyst in a fundamental way, whereas Hossain M et al. studied the kinetics on a mesoporous alumina-supported iron catalyst [4, 5].

To improve styrene yield and selectivity, various reactor configurations have been proposed in the literature.

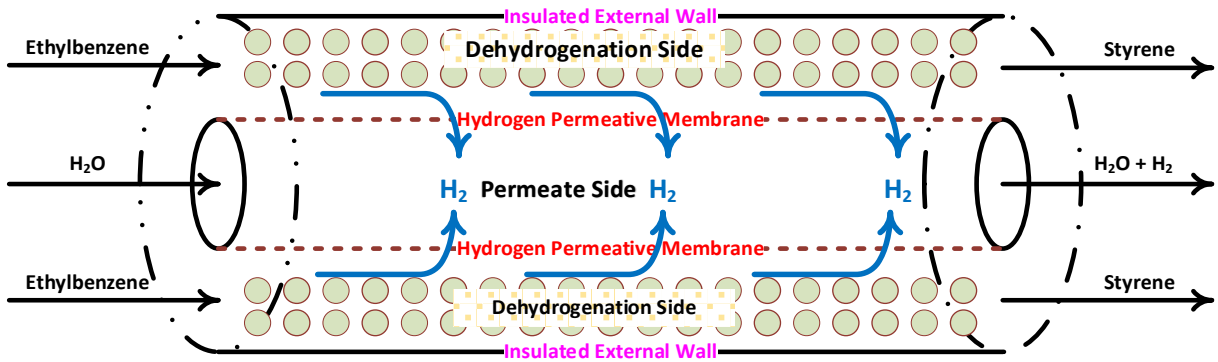

Figure 1: Cocurrent configuration of the catalytic fixed bed membrane reactor operating dehydrogenation of ethylbenzene.

For example, catalytic membrane reactors were used for the dehydrogenation of ethylbenzene to styrene, with the reaction integrated with either a sweep gas on the permeate side or with another hydrogenation reaction, such as the hydrogenation of nitrobenzene to aniline [6, 7]. Both 1D and 2D reactor mathematical models were used to describe and optimize the performance of the newly configured membrane reactors [8]. In catalytic membrane reactors, it is observed that styrene yield and selectivity were significantly improved compared to conventional fixed-bed reactors.

The performance of membrane reactors that enhance the selectivity and yield of a desired species is influenced by several parameters, among them the hydrogen permeation policy. So far, the hydrogen flux profile in membrane reactors has been estimated using Sievert's law, in which the driving force is the difference in hydrogen partial pressures between the reaction and permeate sides. The assessment of the various hydrogen permeation profiles and their impacts on reactor performance has not yet been addressed in the literature. The concept of nonuniform distribution has been already implemented in problems related to the catalyst activity within catalyst pellets and heat exchange in nonadiabatic catalytic reactors. Consequently, the main goal of this manuscript is to address the variation in hydrogen flux along the axial direction of a membrane reactor used to dehydrogenate ethylbenzene to styrene and to assess reactor performance. This will surely help reduce membrane reactor construction costs, as these costs are directly related to hydrogen permeability. Additionally, this study shed light on the problem of optimal hydrogen policy leading to the construction of optimally efficient catalytic membrane reactors.

KINETICS AND REACTION RATE LAWS

The reaction network of ethylbenzene dehydrogenation on a commercial iron oxide catalyst is composed of the following set of stoichiometric equations:

$$C_6H_5C_2H_5 \Leftrightarrow C_6H_5C_2H_3 + H_2 \quad [1]$$

$$C_6H_5C_2H_5 \rightarrow C_6H_6 + C_2H_4 \quad [2]$$

$$C_6H_5C_2H_5 + H_2 \rightarrow C_6H_5CH_3 + CH_4 \quad [3]$$

$$C_6H_5C_2H_3 + 2H_2 \rightarrow C_6H_5CH_3 + CH_4 \quad [4]$$

Two types of chemical reactions occur within the reactor: thermal and catalytic reactions, and neither can be neglected. The thermal reactions, following a radical mechanism, happen in the reactor free zone and are described by reactions [1], [2], and [3]. On the other hand, the catalytic reactions are described by reactions [1], [2], [3], and [4], which take place at the catalyst surface. The kinetic laws of thermal reactions based on molecular mechanisms and catalytic reactions based on intrinsic Hougen-Watson kinetic models are given by Lee and Froment [4].

REACTOR CONFIGURATION

The schematic diagram of the cocurrent configuration of a catalytic membrane reactor is depicted in Figure 1. The reactor is assumed to have a concentric configuration with an outer and an inner side. The outer side of the reactor is packed with a commercial iron oxide catalyst, catalyzing the dehydrogenation reaction of ethylbenzene to styrene. The reaction network is complex, leading to net endothermicity in conventional fixed-bed reactors. Benzene, toluene, and ethylene are produced as byproducts within the reactor. In conventional fixed bed reactors, hydrogen production and net endothermicity limit ethylbenzene conversion, styrene yield, and selectivity.

The external wall on the inner side of the reactor is fabricated to be hydrogen-permeable. It allows hydrogen produced on the dehydrogenation side to infiltrate to the permeate side, where pure steam is used to purge out the diffused hydrogen. Additionally, the permeate side may be used for in-situ heat exchange with the reaction side,

assisting in shifting the reaction's thermodynamic limitation. In-situ hydrogen permeation from the reaction side to the permeate side can significantly enhance the conversion and yield of ethylbenzene and styrene compared to conventional catalytic fixed-bed reactors. The external wall of the reactor on the reaction side is assumed to be perfectly insulated, preventing significant heat transfer between the reactor and its surroundings.

REACTOR MATHEMATICAL MODELLING

The mathematical model governing the performance of the catalytic membrane reactor is deduced considering the following assumptions:

1. Steady state operation
2. Ideal gas behavior on both sides of the reactor.
3. Ideal plug flow behavior.
4. Pseudo-homogeneous reactor model, i.e., neglect the transport effect between the reacting fluid and the catalyst pellets.
5. Neglecting the effect of catalyst deactivation.
6. A highly selective membrane reactor allowing only hydrogen to penetrate to the permeate side.
7. Adiabatic external wall of the membrane reactor
8. Pressure drop within the reaction side is correlated using Ergun's equation.

The reactor model equations on the reaction side of the catalytic membrane reactor may be mathematically presented as follows:

Mole balance equation on the reaction side:

$$\frac{dn_i}{dz} = \sum_{j=1}^{3} \upsilon_{ij}\, r_j^t A_c \varepsilon + \sum_{j=1}^{4} \upsilon_{ij}\, r_j^t \rho_s A_c (1-\varepsilon) - 2\pi R_3 \omega_i J_i(z) \quad [5]$$

Energy balance equation on the reaction side:

$$\frac{dT}{dz} = \frac{\begin{bmatrix} \sum_{j=1}^{3} \upsilon_{ij} r_j^t A_c \varepsilon [-\Delta H(T)]_j^t \\ + \sum_{j=1}^{4} \upsilon_{ij} r_j^t \rho_s A_c (1-\varepsilon)[-\Delta H(T)]_j \\ - \sum_{i=1}^{n} 2\pi R_3 \omega_i J_i(z) \int_{T'}^{T} C_{pi} dT + 2\pi R_3 U(T'-T) \end{bmatrix}}{\sum_{i=1}^{n} n_i C_{pi}} \quad [6]$$

Ergun's equation on the reaction side:

$$\frac{dP}{dz} = -\frac{G}{\rho g_c D_p}\left(\frac{1-\varepsilon}{\varepsilon^3}\right)\left[\frac{(1-\varepsilon)\mu}{D_p} 1.75G\right] \quad [7]$$

On the permeate side, the model equations may be stated as follows:

Mole balance equation on the permeate side:

$$\frac{dn_i'}{dz} = 2\pi R_3 \omega_i J_i(z) \quad [8]$$

Energy balance equation on the permeate side:

$$\frac{dT'}{dz} = \frac{\sum_{i=1}^{n} 2\pi R_3 \omega_i J_i(z) \int_{T'}^{T} C_{pi} dT - 2\pi R_3 U(T'-T)}{\sum_{i=1}^{n} n_i' C_{pi}'} \quad [9]$$

The boundary condition for the cocurrent flow configuration on both sides of the membrane reactor is defined at the reactor inlet as follows:

At $z = 0$:

$$n_i = n_{i0} \quad T = T_0 \quad P = P_0 \quad n_i' = n_{i0}' \quad T' = T_0' \quad [10]$$

In Eqs. [5], [6], [8], and [9], the Greek letter "ω_i" was introduced to make the derived model usable for both a conventional fixed bed reactor and a membrane catalytic reactor. If "$\omega_i = 0$" the model predicts the performance of the conventional fixed bed reactor, while "$\omega_i = 1$" the performance of the catalytic membrane reactor is predicted.

The set of reactor model equations forms an initial-value problem (IVP) that must be solved numerically using a suitable numerical method. In this work, a finite element orthogonal collocation technique is implemented using a predefined number of points per interval, thereby transforming the system of differential equations into a set of nonlinear algebraic equations. The set of algebraic equations has been solved using the "fsolve" MATLAB subroutine.

REACTOR PERFORMANCE MEASURE

The reactor performance indexes considered in this study are ethylbenzene conversion, styrene yield, and styrene selectivity. They are all estimated using the following equations:

$$x_{EB} = \frac{n_{EB,0} - n_{EB}}{n_{EB,0}} \quad [11]$$

$$Y_{ST} = \frac{n_{ST} - n_{ST,0}}{n_{EB,0}} \quad [12]$$

$$S_{ST} = \frac{n_{ST} - n_{ST,0}}{n_{EB,0} - n_{EB}} \quad [13]$$

These performance measures are generally conflicting; maximizing one leads to minimizing the others. As a result, multi-objective optimization problems that combine these measures might be considered at later stages of this study.

PROBLEM MATHEMATICAL FORMULATION

This study aims to investigate and assess the impact of implementing different hydrogen-permeation policies on the performance of a membrane reactor used for the dehydrogenation of ethylbenzene to styrene. The molar flux profile in the reactor model equation is given by the term $J_i(z)$ and it is defined only for hydrogen. The following mathematical model predicting the molar hydrogen flux as a function of axial distance is implemented in the reactor model equations:

$$J_{H_2}(z) = az^{\alpha} + b \quad [14]$$

The parameters "a" and "b" in the above equation

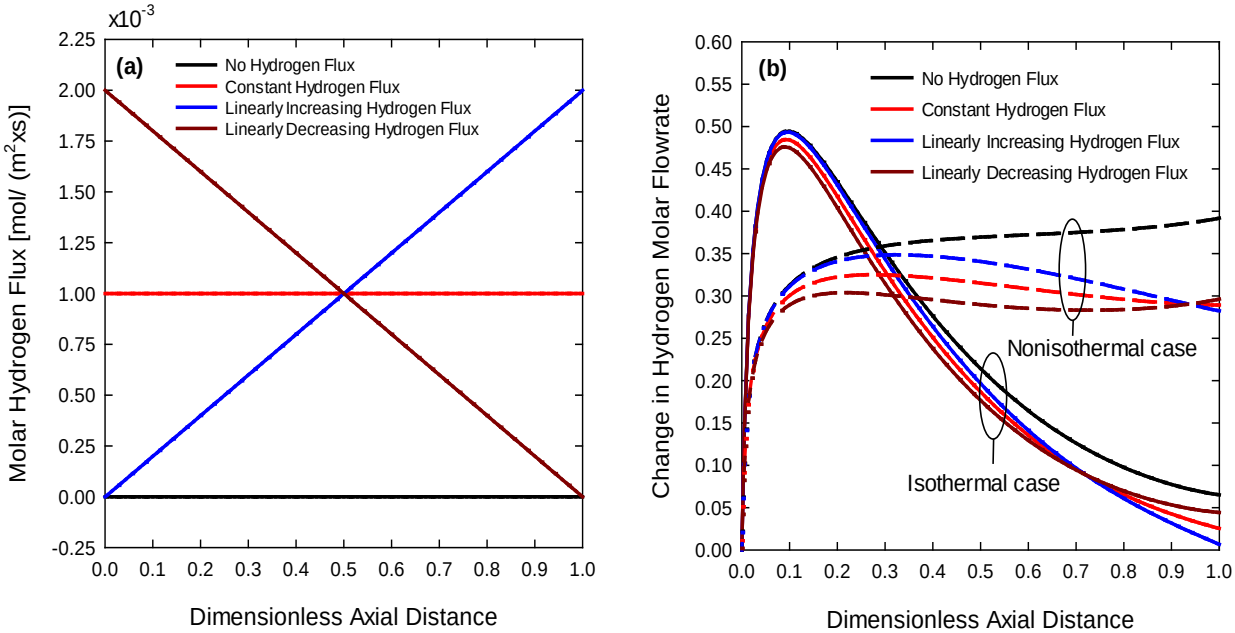

Figure 2: (a) hydrogen flux profiles implemented in the cocurrent configuration of the membrane reactor with an average permeability of 1.0×10^{-3} mol/($m^2\times s$). (b) change in hydrogen molar flowrates for both isothermal and nonisothermal cases.

have the units of mol/($m^{\alpha+2}\times s$) and mol/($m^2\times s$), respectively. Adjusting "α" in the above model provides a variety of linear and nonlinear distributions, i.e., α=0 is constant, α=1 is linear, and α≠1 is nonlinear. Three types of monotonic hydrogen fluxes may be defined: constant hydrogen flux, an increasing molar flux set to zero at the reactor feeding point, and a decreasing molar flux set to zero at the reactor exit. Consequently, the hydrogen molar flux is given for α≠1 as a function of the dimensionless distance $\zeta = \frac{z}{L}$ as follows:

Constant hydrogen flux:

$$J_{H_2}(z) = J_{ave} = \bar{J} \qquad [15]$$

Monotonically increasing hydrogen flux for $0 \le \zeta \le 1$ and $\zeta_1 < \zeta_2$:

$$J_{H_2}(\zeta) = \frac{(1+\alpha)(\zeta_2-\zeta_1)\bar{J}(\zeta^\alpha-\zeta_1^\alpha)}{\zeta_2(\zeta_2^\alpha-\zeta_1^\alpha)-\alpha\zeta_1^\alpha(\zeta_2-\zeta_1)} \qquad [16]$$

Monotonically decreasing hydrogen flux for $0 \le \zeta \le 1$ and $\zeta_1 < \zeta_2$:

$$J_{H_2}(\zeta) = \frac{(1+\alpha)(\zeta_2-\zeta_1)\bar{J}(\zeta^\alpha-\zeta_2^\alpha)}{\zeta_1(\zeta_2^\alpha-\zeta_1^\alpha)-\alpha\zeta_2^\alpha(\zeta_2-\zeta_1)} \qquad [17]$$

In the above molar flux models, the average membrane permeability in the above equations is defined as:

$$J_{ave} = \bar{J} = \frac{1}{L}\int_0^L J_{H_2}(z)dz = \int_0^1 J_{H_2}(\zeta)d\zeta \qquad [18]$$

In this study, the influence of average membrane permeability, nonlinearity, and hydrogen flux on the performance of the catalytic membrane reactor is examined primarily under isothermal conditions. The behavior of the membrane reactor under adiabatic, nonisothermal operating conditions is also presented to extend the conclusions. The focus is mainly on the isothermal condition, as establishing net endothermicity in the nonisothermal reactor will influence reaction kinetics and equilibrium conversion. Hence, the assessment of reactor performance will also be affected by temperature variations.

REACTOR OPERATING CONDITIONS

The geometric parameters and operating conditions used to simulate the performance of the catalytic membrane reactor are summarized in Table 1.

These parameters were kept constant while varying the trajectories of the molar hydrogen flux along the reactor axial distance and accounting for changes in related parameters, such as membrane permeability and the nonlinearity of the hydrogen flux profiles.

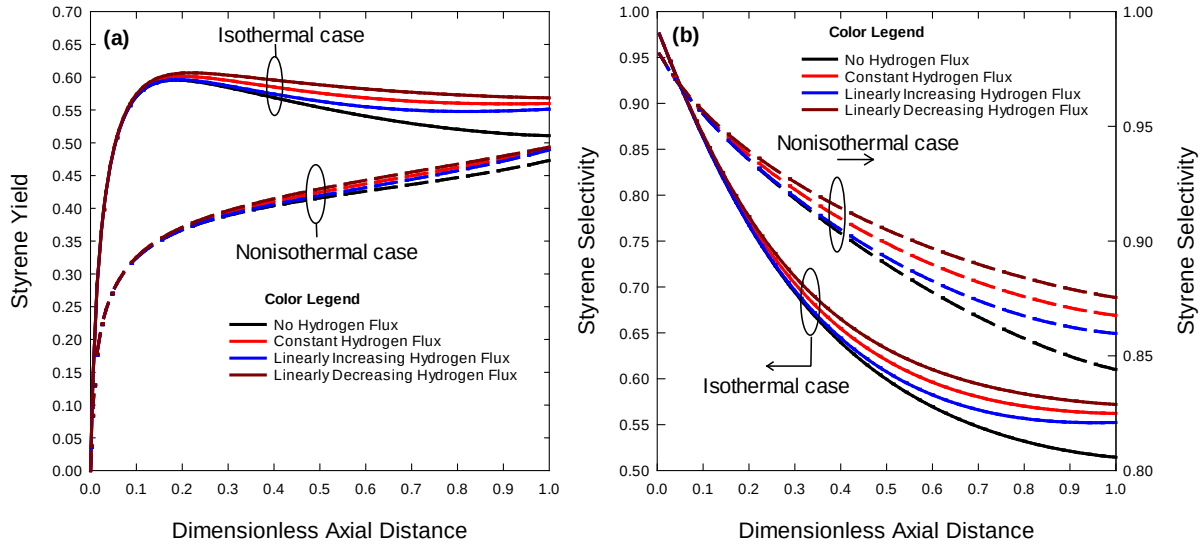

Figure 3: (a) styrene yield along the dimensionless axial distance of the reactor for both isothermal and nonisothermal cases. (b) selectivity to styrene along the dimensionless axial distance of the reactor for both isothermal and nonisothermal cases.

Table 1: List of operating and design parameters of the catalytic membrane reactor

Parameter	Value	Unit
Design Parameters		
Reactor length	0.22	m
Diameter of the permeate side	0.080	m
Equivalent diameter of the reaction side	0.055	m
Voidage	0.4312	--
Diameter of the catalyst particle	0.020	m
Operating Parameters		
Reaction side:		
Total input molar flowrate	1.0×10^{-2}	mol/s
Ethylbenzene mole fraction	0.08321	--
Styrene mole fraction	8.3612×10^{-4}	--
Benzene mole fraction	3.4485×10^{-5}	--
Toluene mole fraction	5.8472×10^{-4}	--
Steam mole fraction	0.9153	--
Feed temperature	886	K
Feed pressure	2.5	bar
Permeate side:		
Steam molar flowrate	0.01	mol/s
Feed temperature	1000	K
Feed pressure	1.0	bar

RESULTS AND DISCUSSION

In this study, the three basic molar hydrogen flux profiles considered, i.e., constant, linearly increasing, and linearly decreasing profiles, are shown in Figure 2a. For the constant molar hydrogen flux profile, the average molar hydrogen flux is maintained along the reactor, unlike the linearly increasing hydrogen flux profile, where the maximum permeability is observed at the reactor exit, or the linearly decreasing hydrogen flux profile, where the maximum permeability is observed at the reactor inlet.

In Figure 2b, the changes in the hydrogen molar flowrate defined as the difference in hydrogen molar flowrate to the ethylbenzene fed, are plotted along the axial distance of the reactor for the conventional fixed reactor, and the catalytic membrane reactor operated isothermally and nonisothermally. The maximum reduction in hydrogen on the reaction side of the membrane reactor occurs with a linearly decreasing hydrogen flux, followed by a linearly increasing hydrogen flux, and then a constant hydrogen flux profile. This observation is consistent whether the membrane reactor is operated isothermally or nonisothermally.

The influence of the application of the linear molar hydrogen flux profiles on the styrene yield and selectivity to styrene is depicted in Figure 3. In Figure 3a, the styrene yield is plotted as a function of the reactor axial distance for both the conventional catalytic fixed reactor and the catalytic membrane reactor, operated with three

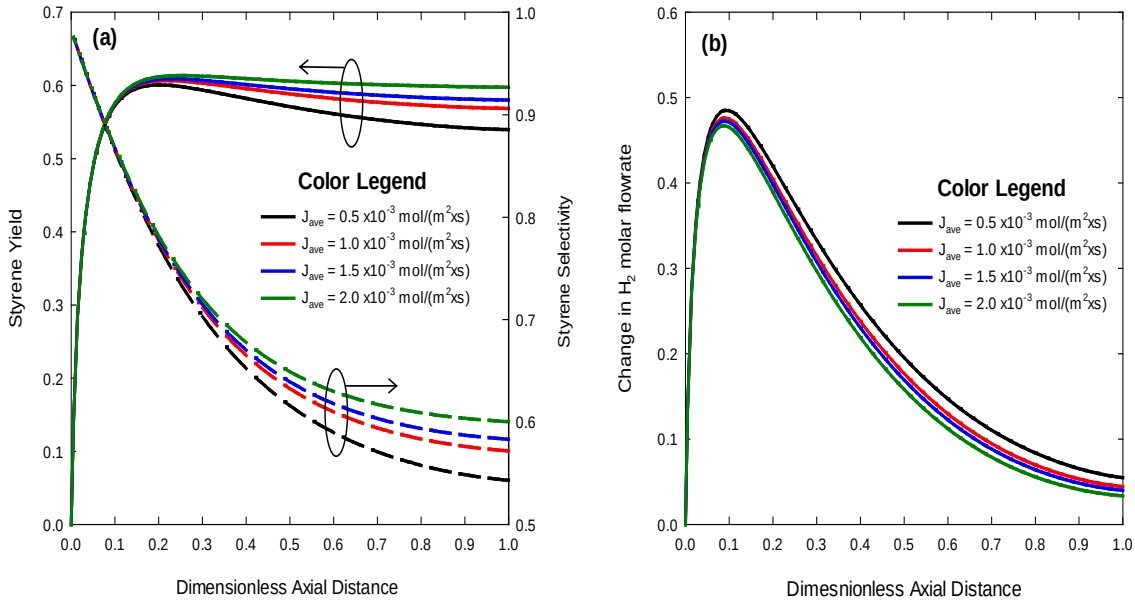

Figure 4: Influence of the average permeability of hydrogen on (a) styrene yield and selectivity to styrene and (b) change in hydrogen molar flowrate along the dimensionless axial distance of the reactor for the linearly decreasing hydrogen flux case.

different molar hydrogen flux profiles. It is found that the catalytic membrane reactor outperforms the conventional fixed-bed reactor, and the membrane reactor with a linearly decreasing hydrogen molar flux performs better than those with linearly increasing or constant hydrogen flux profiles.

This is because hydrogen permeation is significantly enhanced when linearly decreasing hydrogen flux profiles are used. The styrene selectivity along the reactor axial direction is plotted in Figure 3b. It is also observed that the performance of the catalytic membrane reactor with a linearly decreasing hydrogen molar profile exceeds that of conventional catalytic fixed reactors and of catalytic membrane reactors with linearly increasing or constant hydrogen flux profiles.

The effect of changing the average membrane permeability is assessed for a linearly decreasing hydrogen flux profile in Figure 4, where styrene yield, styrene selectivity, and change in hydrogen molar flowrate are plotted along the reactor axial distance. The values of 0.5×10^{-3}, 1.0×10^{-3}, 1.5×10^{-3}, and 2.0×10^{-3} mol/(m^2×s) were chosen to evaluate the performance of the membrane reactor under a linearly decreasing molar flux profile.

In Figure 4a, both the styrene yield and the styrene selectivity are plotted as a function of membrane reactor length for all four values of membrane permeability. It is found that as permeability increases, both the overall styrene yield and the styrene selectivity increase from ~0.53 to ~0.60. This refers to the permeation of hydrogen whose molar flow rate change is plotted along the axial distance of the membrane reactor in Figure 4b. It is shown that as membrane permeability increases, the hydrogen flowrate decreases on the reaction side, thereby enhancing both the styrene yield and selectivity to styrene

The effect of nonlinearity in the creasing flux equation is depicted in Figure 5. For α=0.3, α=1.0, and α=2.0, three nonlinearly decreasing hydrogen flux profiles are plotted along the reactor axial distance in Figure 5a. It is observed that increasing the value of "α" increases permeability at the reactor inlet, and it is higher for α=0.3 than for α=1.0 or α=2.0.

In Figure 5b, both the styrene yield and the change in hydrogen molar flowrate are plotted as a function of the reactor dimensionless distance for α=0.3, α=1.0, and α=2.0. It is found that as the value of "α" decreases, the styrene yield enhances at the reactor inlet, where the hydrogen flux is significant due to the decrease in the hydrogen molar flowrate. This suggests that when ethylbenzene is dehydrogenated in a catalytic membrane reactor, the maximum permeation should occur at the reactor inlet, where the hydrogen molar flow rate is high.

The impact of introducing a permeation section in an existing fixed bed reactor is also assessed in this study. Linearly decreasing hydrogen flux profiles with an average hydrogen molar flow rate of 3.0×10^{-3} mol/(m^2×s) are sketched along the reactor dimensionless axial distance

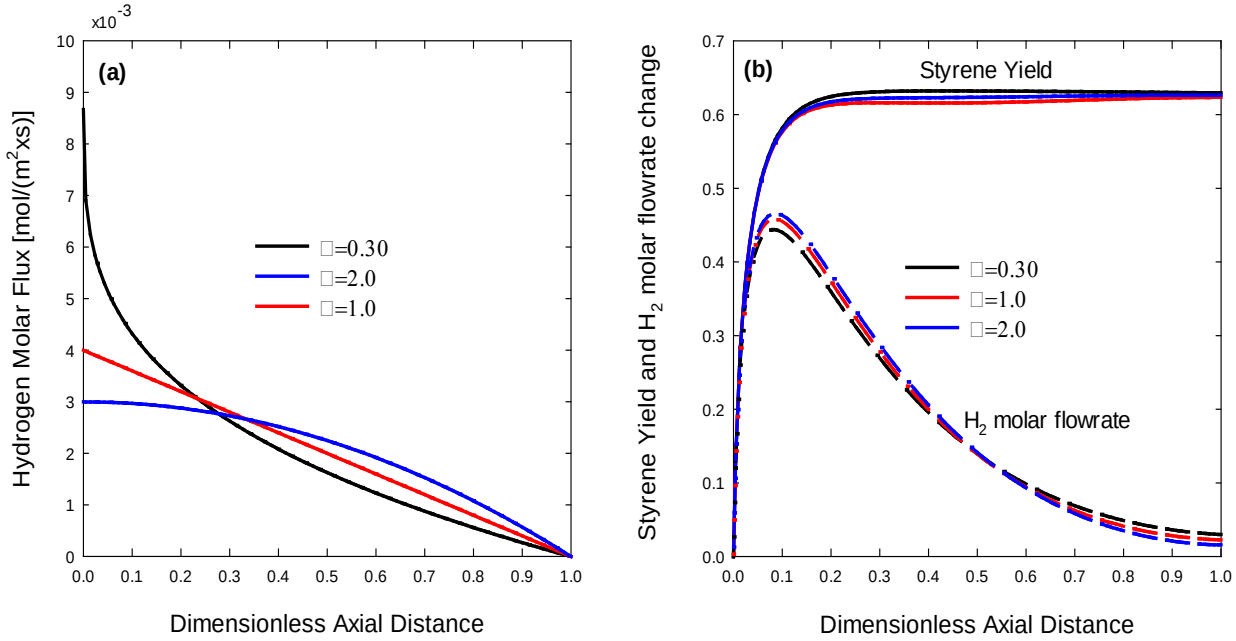

Figure 5: Influence of the nonlinearity in the hydrogen permeation model on (a) hydrogen flux, (b) styrene yield, and change in hydrogen molar flowrate along the dimensionless axial distance of the reactor for the linearly decreasing hydrogen flux case.

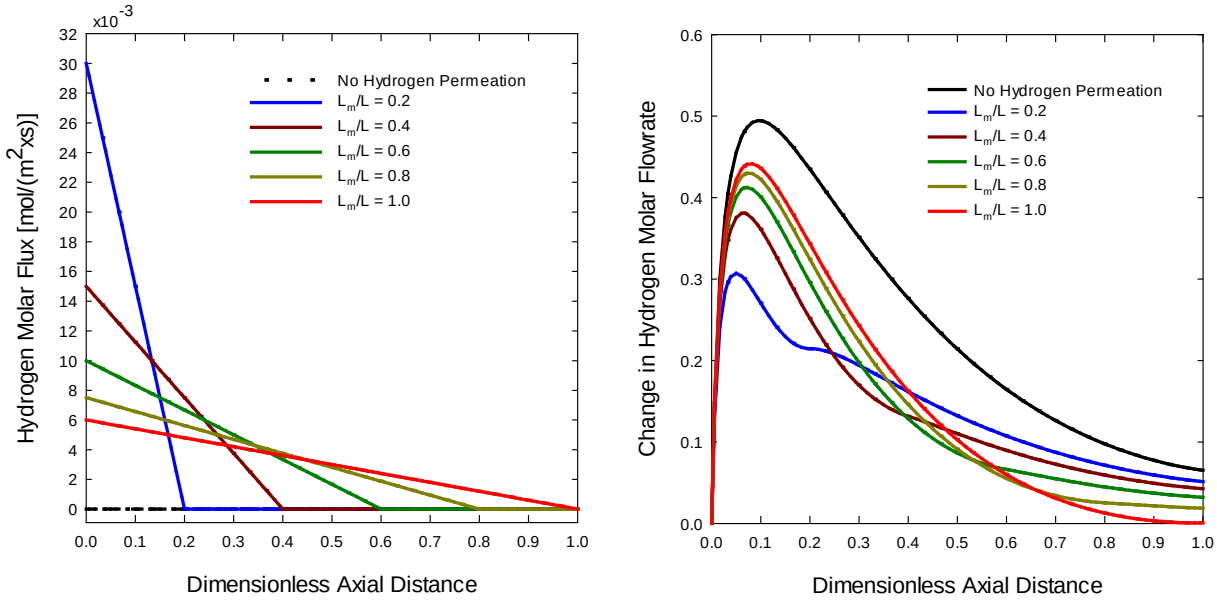

Figure 6: (a) decreasing hydrogen molar flux profiles applied at different fractions of the reactor length, (b) the corresponding change in hydrogen molar flowrate along the reactor axial distance.

in Figure 6a.

As the permeation fraction varies between 0.0 (non-permeable reactor) and 1.0 (Fully permeable reactor), a significant reduction in the hydrogen molar flow rate is observed, as shown in Figure 6b, compared to the conventional fixed bed reactor. The hydrogen molar flow rate increases at the beginning of the reactor due to the dehydrogenation of ethylbenzene, then falls after reaching a maximum at which the rate of hydrogen production is balanced by the rate of hydrogen permeation.

This shows that there is an optimal length when maximizing styrene production is the goal. The optimal

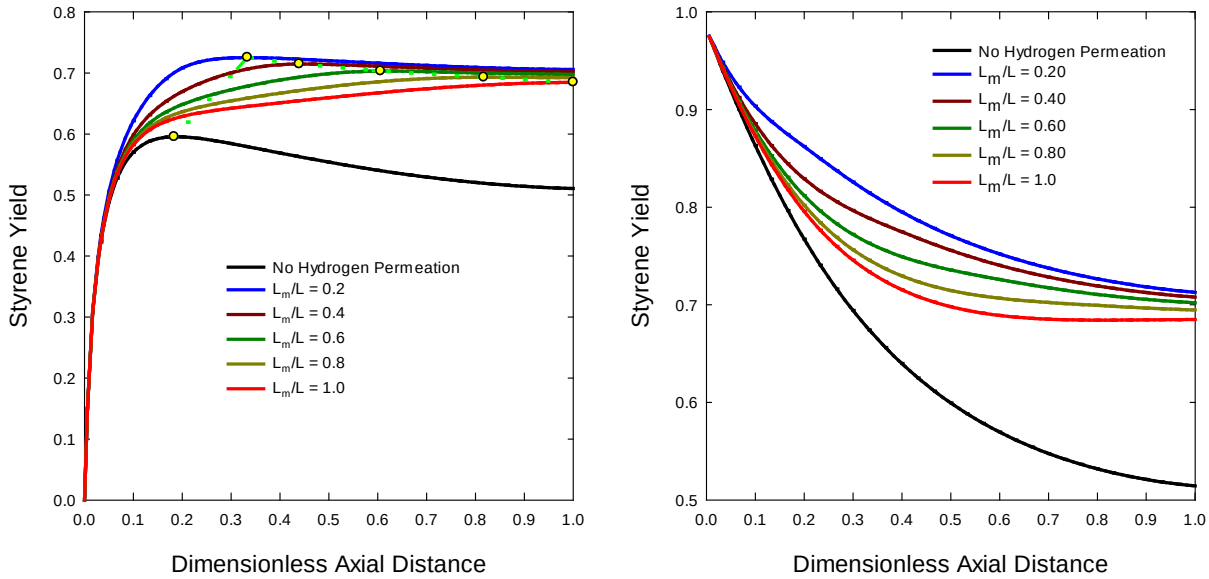

Figure 7: (a) styrene yield and (b) tyrene selectivity along the reactor axial distance for different fractions of the permeation section.

reactor fractional length is shown in Figure 7a by the light green curve, and a summary of the computed optimal reactor fractional length is given in Table 2. Based on the results summarized in Table 2, increasing the fraction of the permeation section within the reactor while maintaining the same average hydrogen permeation rate shifts the optimal reactor length to the right, improving yield. In general, the presence of a nonpermeating section in the reactor is detrimental, as it causes styrene to react with residual hydrogen to produce toluene. This unpleasant effect becomes more pronounced as the permeation section within the reactor is shortened. However, when a permeation section is placed around the point of maximum hydrogen production, the styrene yield in the reactor improves.

For instance, when the permeation section is placed within the first 20% of the reactor length, the predicted styrene yield is 0.73, the highest among all the trials considered. As far as styrene selectivity is concerned, it improves as the fraction of the reactor's permeation section increases.

Table 2: Summary of the optimal reactor length required to maximize styrene yield

Applied fraction of the permeation section	Optimal Reactor fractional length	Maximum styrene yield	Comment
0	0.18	0.60	No Permeation
0.20	0.33	0.73	--
0.40	0.44	0.71	--
0.60	0.61	0.70	--
0.80	0.82	0.69	--
1.0	1.00	0.68	Fully Permeated reactor

CONCLUSION AND RECOMMENDATIONS

The computational analysis of the impact of applying theoretical molar hydrogen profiles on the performance of a catalytic membrane reactor used for the dehydrogenation of ethylbenzene to styrene has been conducted using a one-dimensional reactor model that accounts for changes in molar flowrate, temperature, and

pressure on both the reaction and permeate sides. Three main molar hydrogen profiles are tested, namely constant, linearly increasing, and linearly decreasing profiles, and the performance of the catalytic membrane reactor is assessed via styrene yield and selectivity to styrene. It is found that the reactor should be designed to achieve a maximum hydrogen permeation rate at the inlet section, thereby enhancing styrene yield and selectivity. This can be achieved by applying linearly decreasing molar hydrogen profiles, which are more attractive than constant or linearly increasing hydrogen permeation profiles. Yet, with linearly decreasing hydrogen permeation profiles, the reactor's performance can be further enhanced by adopting a shell-and-tube configuration, which increases the available surface area for hydrogen permeation. Hydrogen membranes with linearly decreasing hydrogen profiles can be fabricated by correlating the palladium distribution with the hydrogen permeation rate. The steam feeding strategy on the permeate side might also be used to adjust the permeation rate of existing membranes to match the linear hydrogen flux observed in this work. The application of these results can significantly promote the use of hydrogen membranes with varying permeabilities in chemical reactors, thereby reducing construction costs. In addition, the results here shed light on the optimal hydrogen flux trajectory required to maximize the styrene yield and selectivity. Adding a permeation section around the point where maximum hydrogen is observed is preferable and yields higher conversion than a completely permeable reactor.

NOMENCLATURE:

A_c: reactor cross-sectional area, [m^2].
D_p: diameter of the catalyst pellet, [m].
G: mass velocity of the reacting mixture, [Kg/(m^2×s)].
$J_i(z)$: molar flux of species i along the axial distance of the reactor, [mol/(m^2×s)].
$J_{ave} or \bar{J}$: average molar flux of species i, [mol/(m^2×s)].
n_i: molar flowrate of species i in the reaction side of the reactor, [mol/s].
L: reactor length, [m].
P: Pressure in the reaction side of the reactor, [K].
r_j^c: catalytic reaction rate of j reaction, [mol/(Kg×s)].
r_j^t: thermal reaction rate of j reaction, [mol/(m^3×s)].
T: Temperature on the reaction side of the reactor, [K].
T': Temperature on the permeate side of the reactor, [K].
U: overall heat transfer coefficient, [J/(m^2×K×s)].
z: reactor axial distance, [m].
ζ: dimensionless reactor axial distance, [-].
μ: Viscosity of the reacting mixture, [Pa×s].
υ_{ij}: stoichiometric no. of species i in j reaction, [-].
ρ: density of the reacting mixture, [kg/m^3].
ρ_s: density of the catalyst particles, [kg/m^3].
ε: bed void fraction, [-].
ζ_1: the start of the membrane section, [-].
ζ_2: the end of the membrane section, [-].

ACKNOWLEDGEMENTS

The author would like to thank King Fahd University of Petroleum & Minerals (KFUPM) and the Interdisciplinary Center for Refining & Advanced Chemicals for the financial support.

DEDICATION

Dedicated to the soul of my great mentor, PhD advisor, and scholar, Prof. John R. Grace, who taught me scientific ethics, manners, and appreciation.

REFERENCES

1. Mimura N, Saito M. Dehydrogenation of ethylbenzene to styrene over fe2o3/al2o3 catalysts in the presence of carbon dioxide. Catalysis Today 55:173-178 (2000). https://doi.org/10.1016/s0920-5861(99)00236-9
2. Zhu XM, Schön M, Bartmann U, van Veen AC, Muhler M. The dehydrogenation of ethylbenzene to styrene over a potassium-promoted iron oxide-based catalyst: a transient kinetic study. Applied Catalysis A: General 266:99-108 (2004). https://doi.org/10.1016/j.apcata.2004.02.002
3. Rossetti I, Bencini E, Trentini L, Forni L. Study of the deactivation of a commercial catalyst for ethylbenzene dehydrogenation to styrene. Applied Catalysis A: General 292:118-123 (2005). https://doi.org/10.1016/j.apcata.2005.05.046
4. Jae Lee W, Froment GF. Ethylbenzene dehydrogenation into styrene: kinetic modeling and reactor simulation. Ind. Eng. Chem. Res. 47:9183-9194 (2008). https://doi.org/10.1021/ie071098u
5. Hossain MM, Atanda L, Al-Yassir N, Al-Khattaf S. Kinetics modeling of ethylbenzene dehydrogenation to styrene over a mesoporous alumina supported iron catalyst. Chemical Engineering Journal 207-208:308-321 (2012). https://doi.org/10.1016/j.cej.2012.06.108
6. Hermann C, Quicker P, Dittmeyer R. Mathematical simulation of catalytic dehydrogenation of ethylbenzene to styrene in a composite palladium membrane reactor. Journal of Membrane Science 136:161-172 (1997). https://doi.org/10.1016/s0376-7388(97)81990-4
7. Abo-Ghander NS, Grace JR, Elnashaie SSEH, Lim CJ. Modeling of a novel membrane reactor to integrate dehydrogenation of ethylbenzene to styrene with hydrogenation of nitrobenzene to

aniline. Chemical Engineering Science 63:1817-1826 (2008). https://doi.org/10.1016/j.ces.2007.12.007

8. Abo-Ghander NS, Logist F, Grace JR, Van Impe JFM, Elnashaie SSEH, Lim CJ. Optimal design of an autothermal membrane reactor coupling the dehydrogenation of ethylbenzene to styrene with the hydrogenation of nitrobenzene to aniline. Chemical Engineering Science 65:3113-3127 (2010). https://doi.org/10.1016/j.ces.2010.02.007

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Neural Model of Pinch-Based Multicomponent Distillation for Applications in Flowsheet Synthesis

Alexander B. Wolf[a], Mirko Skiborowski[b], and Jakob Burger[a]*

[a] Technical University of Munich, TUM Campus Straubing for Biotechnology and Sustainability, Laboratory of Chemical Process Engineering, Straubing, Germany
[b] Hamburg University of Technology, Institute of Process Systems Engineering, Hamburg, Germany
* Corresponding Author: burger@tum.de.

ABSTRACT

This work presents a data-driven surrogate modeling framework for predicting distillation behavior assuming an infinite number of stages and distillation limits informed by residue-curve topology and pinch-point feasibility analysis. The framework provides a direct mapping from feed composition and distillate-to-feed ratio (D/F) to distillate and bottom product compositions, making it suitable for flowsheet synthesis and optimization applications. The approach combines three components: a classifier that identifies feasible singular-point splits, a boundary regression model that predicts D/F limits separating pure- and mixed-product operating regimes, and a neural network that interpolates product compositions in the intermediate regime. The method is demonstrated for the ternary system ethanol, benzene, and water at 1 atm using data generated from rigorous vapor–liquid–liquid equilibrium analysis. Results show that the framework provides reliable predictions for pure splits while retaining smooth interpolation behavior in the mixed regime.

Keywords: Distillation, Surrogate Model, Machine Learning, Process Design, Modelling and Simulations

INTRODUCTION

Modeling distillation behavior of multicomponent systems is central to conceptual process design. Cumbersome process simulations based on equilibrium stage models are computationally demanding, prone to numerical failure, and require numerous specifications. They are therefore typically avoided in early design phases in favor of shortcut models that focus on the general feasibility of a split rather than detailed apparatus design. Shortcut models typically assume an infinite number of theoretical stages and target the minimum energy demand for pinched separations, e.g., the Underwood method for constant relative volatilities or the rectification body method for general nonideal mixtures [1]. Such shortcuts require prespecified products for which feasibility is determined. In many applications such as flowsheet synthesis, a mapping from feed and split specifications (e.g. distillate-to-feed ratio (D/F)) to product composition is required instead.

If such a mapping is required, one could rely on heuristics to suggest preferred column products [2] or wrap existing shortcut models inside iterative loops that adjust product specifications until material balances and split specifications are met [3]. Heuristics often restrict the design space, for example by excluding non-sharp splits [2], while iterative loops can be slow and may fail to converge, particularly in nonideal and azeotropic settings [3]. Ryll et al. [4] propose direct mapping from feed specifications and D/F to product compositions under infinite-stages and total reflux (∞/∞) assumptions; however, these methods are geometric, work only for discrete pressures, and are limited to four components [5] [6].

In this work, we address these limitations by proposing a data-driven surrogate model for distillation columns assuming an infinite number of stages and operating under general distillation limits defined by residue-curve behavior at total reflux and pinch point-based feasibility limits. The key objective is a deterministic, direct mapping from the molar feed composition and molar D/F to the distillate and bottom molar compositions for *a priori* defined systems of any vapor–liquid equilibrium (VLE) topology. Accordingly, these models serve as tailor-made surrogates for feasibility studies in conceptual process design, intended for flowsheet synthesis and machine learning (ML)-based process optimization, where a

large number of process alternatives have to be screened before optimizing the column parameters (e.g., number of stages, reflux ratio) in detail.

Data-driven surrogate models have been used successfully for distillation column design and optimization, including neural network-based [7-9] and Gaussian process-based [10, 11] models for different column configurations. Further, Panofen et al. [12] have recently proposed a surrogate model of distillation boundaries in azeotropic multicomponent mixtures. However, data-driven surrogates are always tailor-made for, and thus limited to one specific task. To the best of our knowledge, no surrogate currently exists for predicting distillation column products at the limiting cases of an infinite number of stages and distillation boundaries informed by residue-curve topology and pinch-point feasibility analysis. We close this gap with the present work.

We present results for ternary systems, although the methodology is in principle not limited and can be extended to multicomponent systems when sufficient computational power is available during training. The resulting surrogate model can be used in any flowsheet synthesis problem. Because the applied ML-based models evaluate quickly and robustly, it is particularly suited for optimization methods that identify flowsheet concepts, such as superstructure optimization and ML-based flowsheet design [3], where numerous flowsheet evaluations are needed.

METHODOLOGY

Overview

The objective is a direct and deterministic mapping from feed composition and D/F to the corresponding product compositions,

$$(\boldsymbol{z}, D/F) \rightarrow (\boldsymbol{x}, \boldsymbol{y}), \tag{1}$$

where $\boldsymbol{z} = (z_1, z_2, \dots, z_n)$ denotes the mole fractions of the feed an n-component mixture with $\sum_{i=1}^{n} z_i = 1$, $\boldsymbol{y} = (y_1, y_2, \dots, y_n)$ the distillate's mole fractions, and $\boldsymbol{x} = (x_1, x_2, \dots, x_n)$ the bottom product's mole fractions.

The training data for this mapping are generated using the Automated Process Synthesis (APS) Tool v1.0 [2], which performs thermodynamically rigorous VLE analysis of multicomponent mixtures. The APS Tool identifies azeotropes, distillation regions, and feasible singular-point splits, that is, separations involving pure components or azeotropes, using general distillation limits defined by residue-curve topology and continuation-based pinch-branch calculations. Residue curves characterize separation behavior in the total reflux limit, whereas pinch lines describe the limiting behavior of reversible columns operated at finite reflux. A split is considered feasible if it is admissible under either limiting case. Consequently, the feasible split set identified by the APS Tool represents the outer envelope of thermodynamically admissible separations for a given feed.

In the present work, these identified limiting singular-point splits are used as anchors for surrogate modeling. Once the limiting products are known, the separation behavior between them can be represented in terms of D/F, without explicitly distinguishing the underlying reflux mechanism.

A straightforward approach to learning the mapping in Eq. (1) would be to train a sufficiently large neural network directly on input—output data. While feasible, this approach was found to exhibit inconsistent accuracy in predicting azeotropic compositions and occasionally produced slightly impure product compositions in cases where a pure-product split is expected.

To address these limitations, we propose embedding the neural network into a structured hierarchical learning framework that leverages knowledge on distillation theory to guide the training process and enhance physical interpretability. The overall structure of the proposed framework is illustrated in Figure 1.

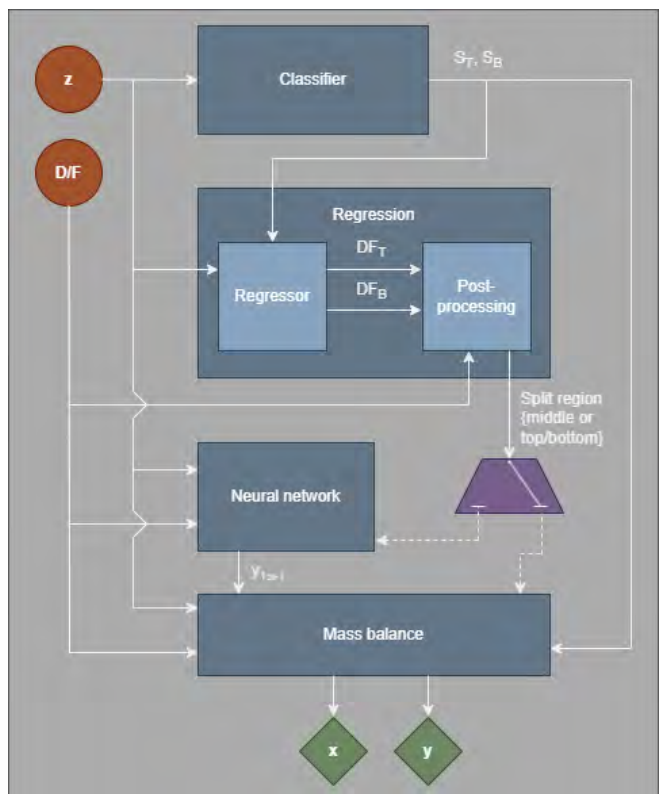

Figure 1. Hierarchical surrogate framework integrating data-driven models for predicting distillation behavior.

At the beginning of the pipeline, two learning tasks are introduced to extract physically interpretable information about the separation before predicting the compositions $(\boldsymbol{x}, \boldsymbol{y})$. The first task is a classification problem

that determines, based on $\boldsymbol{z}$, which singular points in the VLE act as unstable and stable nodes of the distillation lines governing the splits. The second task is a boundary regression that predicts the limiting D/F values associated with attainable singular-point separations for given feed composition $\boldsymbol{z}$ and singular points determined in the classifier.

In certain cases, the product compositions can be determined directly from the outcomes of these two tasks. For example, when separation to a pure singular-point product is feasible at a given D/F, the corresponding distillate or bottom compositions coincide with the composition of that singular point. Only for D/F values that do not admit a pure singular-point split, is the neural network used to predict the product compositions.

Both the classifier and the boundary regressor are entirely data-driven, with the required labels obtained and stored during the data-generation process. Overall, the proposed hierarchical framework effectively *distills* VLE knowledge into a data-driven surrogate model that is computationally efficient, numerically stable, and well-suited for flowsheet synthesis and ML-based process synthesis.

Data generation and representation

The training data are derived from rigorous VLE analyses performed with the MATLAB-based APS Tool. The proposed framework approximates these results while preserving physical interpretability and numerical robustness. Rather than relying on repeated evaluations with the APS Tool, physically meaningful separations are extracted from the tool's thermodynamic analyses and combined with material-balance calculations to determine the D/F value corresponding to a particular product split. This information can then be embedded into predictive models that accept D/F as an explicit input and can be evaluated efficiently in a Python-based environment.

First, a separation system is defined by specifying pure-component properties, such as Antoine equation parameters, together with a binary activity coefficient model. Based on this thermodynamic description, the singular points of the system are identified. The finite set of thermodynamically admissible singular points as top products is denoted as $\boldsymbol{S_T} = \{s_{T,1}, s_{T,2}, \dots, s_{T,K}, \varnothing_T\}$, and the corresponding set for bottom products by $\boldsymbol{S_B} = \{s_{B,1}, s_{B,2}, \dots, s_{B,K}, \varnothing_B\}$, where $\varnothing_T$ and $\varnothing_B$ denote the absence of a feasible pure distillate or pure bottom split, respectively.

Feed compositions $\boldsymbol{z}$ are then systematically sampled over the composition simplex. The evaluated feeds evenly cover the simplex, as studies on surrogate models show that such coverage reduces error and delivers large speedups when embedded later in flowsheet optimization [12]. For each feed composition, a one-hot encoding is assigned to indicate which elements of S_T and S_B correspond to feasible top and bottom splits, respectively. In addition, the corresponding product compositions $\boldsymbol{y}$ and $\boldsymbol{x}$ for the pure split at the maximum achievable yield are provided.

The APS Tool may identify multiple competing singular points for a given feed composition, for example near distillation boundaries where more than one pure top or bottom product is theoretically feasible. In such cases, expert judgment is applied to select the most appropriate split.

Once the singular-points splits have been identified, two critical values of D/F are defined for each feed composition: DF_T, the maximum D/F at which a pure distillate is achievable, and DF_B, the minimum D/F at which a pure bottom product is achievable. These values are calculated from a component-wise molar balance,

$$z_i = y_i(D/F) + x_i(1 - D/F). \tag{2}$$

The resulting values partition the operating space into three regimes:

Top: $0 < D/F \leq DF_T$,

Middle: $DF_T, < D/F < DF_B$

Bottom: $DF_B \leq D/F < 1$.

If no pure distillate split exists for a given feed, DF_T is set to zero, and if no pure bottom split DF_B is set to one. This information constitutes the complete dataset required to train the classifier and boundary regressor.

To train the neural network, additional data are generated to capture the relationship outlined in Eq. (1). For singular-point splits, product compositions are computed for pure-product yields Y_i ranging from Y_i =0.01 to the maximum feasible yield. An example of such data points for a single feed composition in a ternary mixture of ethanol, benzene, and water at 1 atm is illustrated in Figure 2.

The feed selected in Figure 2 corresponds to a separation governed by residue-curve topology, such that the intermediate operating points remain consistent with the distillation boundary associated with the total-reflux limit. Data points corresponding to a pure bottom split (orange circles) represent the bottom operating regime at high D/F values (D/F = 0.78-0.99), whereas points associated with a top split with a minimum ternary azeotrope (red squares) represent the top operating regime at low D/F values (D/F = 0.003-0.33).

Data for the middle operating regime are generated by interpolating between the limiting pure splits. When a feed admits both a pure distillate and bottom product, intermediate compositions are sampled along the discretized binary line connecting the corresponding pure-split product compositions, as indicated by the green diamonds in Figure 2. For each sampled composition, the corresponding D/F is computed using Eq. (2).

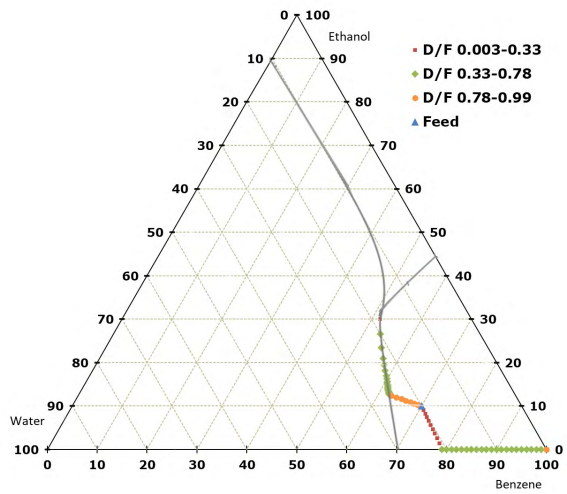

Figure 2. Example data points for ternary system collected using APS Tool [2]. Separation behavior falls within limits of residue-curve topology.

Figure 3 illustrates a case in which the selected feed admits feasible operating points associated with pinch-line feasibility. In this case, some middle-regime data points extend beyond the residue-curve-based distillation boundary, reflecting admissible separations under reversible-column limits. Both types of behavior are retained in the training dataset and treated as physically admissible by the surrogate framework.

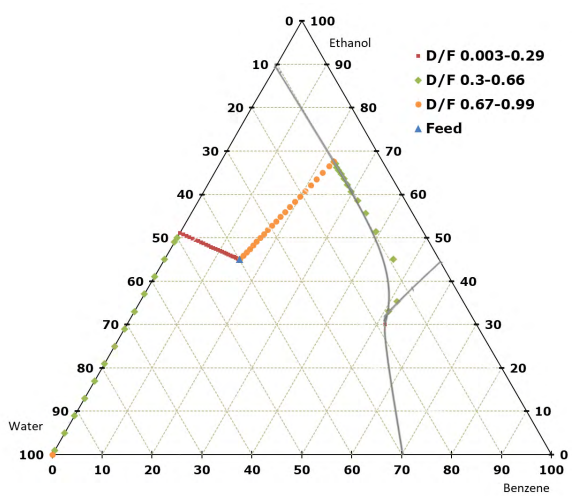

Figure 3. Example data points for ternary system collected using APS Tool [2]. Separation behavior falls within limits of pinch-line feasibility analysis.

Moreover, due to the geometry of the distillation boundaries and pinch-points limits, the resulting D/F values along these interpolation paths do not necessarily vary monotonically. This further motivates the use of a neural network to learn the relationship between D/F and product composition in the middle regime. In addition, numerical convergence issues may arise in rigorous VLE calculations for certain intermediate operating points, motivating the development of a surrogate model capable of delivering robust predictions across the full operating range.

Singular point split classification

The first ML task addresses the identification of singular-points for a given feed composition. The input to this task is the reduced feed composition vector $\tilde{\boldsymbol{z}} = (z_1, z_2, \dots, z_{n-1})$ for an n-component system. The output is a one-hot encoded vector over the predefined sets S_T.and S_B, indicating which singular points correspond to the top and bottom splits, respectively, for the given feed compositions.

The model employs a histogram-based gradient boosting classifier implemented in the *scikit-learn* library in Python. This method constructs an ensemble of decision trees sequentially, where each tree is trained to correct misclassifications of the current ensemble. The histogram-based discretization of the input space improves numerical robustness and computational efficiency while preserving the ability to represent nonlinear decision boundaries in the composition simplex.

The trained classifier yields a discrete and interpretable mapping from $\tilde{\boldsymbol{z}}$ to S_T.and S_B within the training domain.

Operating regime prediction via boundary regression

The regression task is formulated as a boundary regression problem. Two continuous regression models are trained to predict the functions $\widehat{DF}_T = f_T(\tilde{\boldsymbol{z}}, S_T, S_B)$ and $\widehat{DF}_B = f_B(\tilde{\boldsymbol{z}}, S_T, S_B)$. The operating regime for a queried D/F value is then obtained by comparison with these predicted boundaries.

Explicit constraints are enforced to ensure physically consistent predictions, including clamping predicted ratios to the interval [0, 1] and enforcing the condition $\widehat{DF}_T \geq \widehat{DF}_B$. The regression models are implemented using histogram-based gradient boosting regressors from the *scikit-learn* library, which efficiently capture nonlinear boundary surfaces over the composition simplex while remaining numerically robust.

Neural network for impure product splits

For operating points in the middle regime, product compositions are predicted using a feed-forward neural network implemented in Python with *PyTorch*. The network is trained as a regression model that maps $\boldsymbol{z}$ and D/F to the distillate compositions $\tilde{\boldsymbol{y}} = (y_1, y_2, \dots, y_{n-1})$. The remaining y_n and corresponding bottom compositions $\boldsymbol{x}$ are computed from an overall material balance.

Product composition reconstruction

The final product compositions are predicted using a hybrid reconstruction strategy that combines analytical material balances and the neural network prediction. In the top and bottom regimes, product compositions are computed using Eq. (2) consistent with the corresponding limiting pure split. For operating points in the middle regime, the neural network is used to predict product compositions.

To ensure physical feasibility, all predicted compositions are postprocessed to satisfy component-wise material constraints. For a given D/F, the distillate mole fractions must satisfy

$$y_i \leq \min\left(1, \frac{z_i}{D/F}\right), \tag{3}$$

reflecting that a distillate stream cannot be in excess of its availability in the feed. If necessary, the predicted distillate vector is projected onto the feasible composition simplex under these caps, ensuring nonnegativity and a unit sum. Once a feasible distillate composition is obtained, the bottom composition is computed from the overall material balance (see Eq. (2)).

EXAMPLE CASE STUDY

The proposed framework is demonstrated using the ternary mixture of ethanol, benzene, and water at 1 atm, modeled using the Non-Random Two-Liquid (NRTL) g^E-model [13] with parameters taken from Aspen Plus [14]. The mixture exhibits a liquid-liquid phase split and four minimum azeotropes (three binary, one ternary; two homogeneous, two heterogeneous). The set S_B comprises pure bottom splits yielding ethanol, benzene, or water, as well as the absence of a pure bottom split. Since all feasible compositions admit a top split with the ternary minimum azeotrope, the set S_T contains only a single option and yielding a trivial classification task.

Data for split classification and operating regime prediction are generated with approximately 124, 000 feed compositions samples across the VLE simplex. The samples are drawn from a structured grid with component mole fractions in the range 0.002 ≤ z_i ≤ 0.996. To avoid artificial grid-induced artifacts and improve generalization, small Gaussian perturbations with standard deviations of σ = 0.0003 are added to the sampled compositions. This stochastic smoothing prevents the classifier and regression models from memorizing grid-aligned decision boundaries and encourages learning of physically meaningful, continuous separation surfaces in the composition space.

A comparatively high sampling density is used for classification and boundary regression tasks because these models must resolve sharp transitions between distillation regions and accurately capture singular-point behavior near curved distillation boundaries. In addition, generating training data for pure splits is computationally efficient, as it requires only the identification of limiting separation behavior, rather than singular points across the entire operating range.

In contrast, training data for the neural network are generated using feed compositions in the range of 0.02 ≤ z_i ≤ 0.96. For middle-regime splits, product compositions are generated along the binary line connecting the pure bottom split and the bottom composition associated with the pure distillate split, as described above and outlined in Figures 2 and 3. These compositions are sampled in increments of Δz_i = ±0.01. In total, approximately 293, 000 data points are generated for neural network training.

Results

Classifier

The pure-split classification model was evaluated on 4, 851 validation feed compositions that were not used in the training dataset, with component mole fractions in the range 0.01 ≤ z_i ≤ 0.98. The results of the evaluation are shown in Figure 4, which shows the predicted pure bottom split for each feed composition $\boldsymbol{z}$ selected from the admissible set S_B.

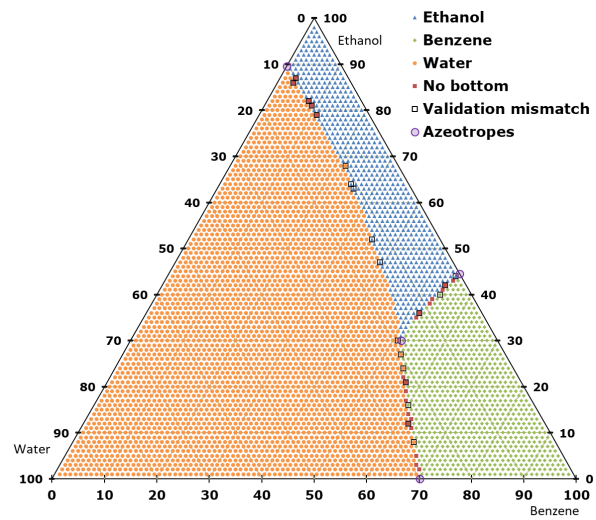

Figure 4. Validation run for singular-point split classification.

Of the analyzed feed compositions, only 21 showed mismatches with the expected splits. Notably, all of these cases occur at distillation boundaries separating feasible pure splits. In eight cases, the model predicted an incorrect pure bottom component, while in the remaining 13 cases it incorrectly classified whether a pure bottom split exists at the boundary point. Further improvements could be achieved by increasing the training density or applying additional weighting for feed compositions near distillation boundaries. In addition, incorporating confidence

measures based on the consistency of predictions in a local neighborhood of the composition space could enable the use of fallback strategies or alternative predictors when the model encounters uncertain splits.

Regression

The operating regime regression model was evaluated using the same validation set by comparing queried D/F values against the predicted DF_T and DF_B boundaries. The predicted DF_T values exhibited a mean absolute error of $|\varepsilon_T|$ = 0.0039 with a standard deviation of σ_T = 0.0056, with a maximum observed error of ε_T = 0.0178. The prediction error for DF_B was slightly greater, with a mean absolute error of $|\varepsilon_B|$ = 0.0062 and a standard deviation of σ_B = 0.0129. The two largest errors, ε_B = 0.277 and ε_B = 0.194, occurred for feed compositions where the bottom split was misclassified by the singular-point classifier, with ethanol predicted as the bottom product instead of water. Both cases lie near distillation boundaries, underscoring the sensitivity of the regression task to accurate split classification near these regions. While this suggests that additional fine-tuning of the classifier could reduce such errors, alternative strategies for improving the regression model could also be explored, including the use of different regression architectures or uncertainty-aware regression formulations.

Neural Network

The neural network is trained using mini-batch gradient descent with the Adam optimizer and a mean-squared-error (MSE) loss function. Key modeling choices, including data normalization, loss formulation, validation strategy, optimization settings, and stopping criteria, were systematically evaluated. A randomized train–validation split with a fixed random seed and early stopping based on validation loss were selected to balance accuracy and generalization. Among the tested configurations, a network with three hidden layers of 32 neurons each achieved the best validation performance, with an MSE of 2.36×10^{-5}.

Overall model

The full predictive framework was evaluated using representative feed compositions by first generating reference product compositions and the corresponding D/F values using the APS Tool. Figure 5 illustrates four representative cases. The first three feed compositions (cases (a)-(c)) were selected for their location within distinct distillation regions, for which a pure bottom split is expected. The fourth (case (d)) was chosen because it represents a composition for which the singular-point classifier produced a mismatch. The APS Tool-derived D/F values were provided as inputs to the hybrid surrogate and neural network models, and the resulting product compositions were compared against the APS reference data.

Although the hybrid surrogate and standalone neural network show similar overall agreement with the APS Tool data in cases (a)-(c), closer inspection reveals qualitative and quantitative differences in performance. In particular, for cases (a) and (b), the standalone network fails to correctly reproduce the expected pure-bottom product splits at high D/F values where thermodynamics dictates a pure split. The hybrid surrogate successfully predicts the compositions of these pure splits, as pure-split compositions are enforced by construction.

These discrepancies are reflected in the error statistics summarized in Table 1, which report the mean absolute error of the predicted product composition vectors relative to the APS Tool reference, evaluated across all components simultaneously and averaged over all considered D/F values for each split.

Table 1: Comparison of prediction errors between hybrid surrogate and neural network for representative feed compositions.

	Prediction error (mean ± std. dev.)	
Case study	**Surrogate model**	**Neural network**
(a)	4.17×10^{-4}	3.57×10^{-3}
	$\pm2.82\times10^{-3}$	$\pm2.54\times10^{-2}$
(b)	3.32×10^{-5}	1.69×10^{-3}
	$\pm1.58\times10^{-4}$	$\pm1.15\times10^{-2}$
(c)	8.56×10^{-5}	3.23×10^{-4}
	$\pm7.11\times10^{-4}$	$\pm7.59\times10^{-4}$
(d)	1.26×10^{-1}	1.10×10^{-2}
	$\pm2.83\times10^{-1}$	$\pm2.53\times10^{-2}$

For cases (a)-(c), the hybrid surrogate achieves lower mean errors, on the order of 10^{-5}-10^{-4}, compared with 10^{-4}-10^{-3} for the neural network, and shows substantially lower variability. In particular, the standard deviations of the surrogate predictions remain less than 3×10^{-3}, whereas the neural network exhibits standard deviations exceeding 1×10^{-2} for all splits. These differences also arise from the tendency of the neural network to produce inconsistent predictions of the ternary azeotrope composition, whereas the hybrid surrogate enforces exact singular-point compositions.

For the feed corresponding to the classifier mismatch (case (d)), errors show the opposite trend. The neural network exhibits a lower mean error than the hybrid surrogate, whose errors increase substantially due to the incorrect pure-split classification of ethanol instead of water, as seen in Figure 5. In this case, the surrogate predictions trend toward the misclassified bottom split, whereas the neural network captures a smooth transition toward the correct split. This behavior suggests that additional uncertainty-aware handling near classification boundaries could further improve the hybrid framework.

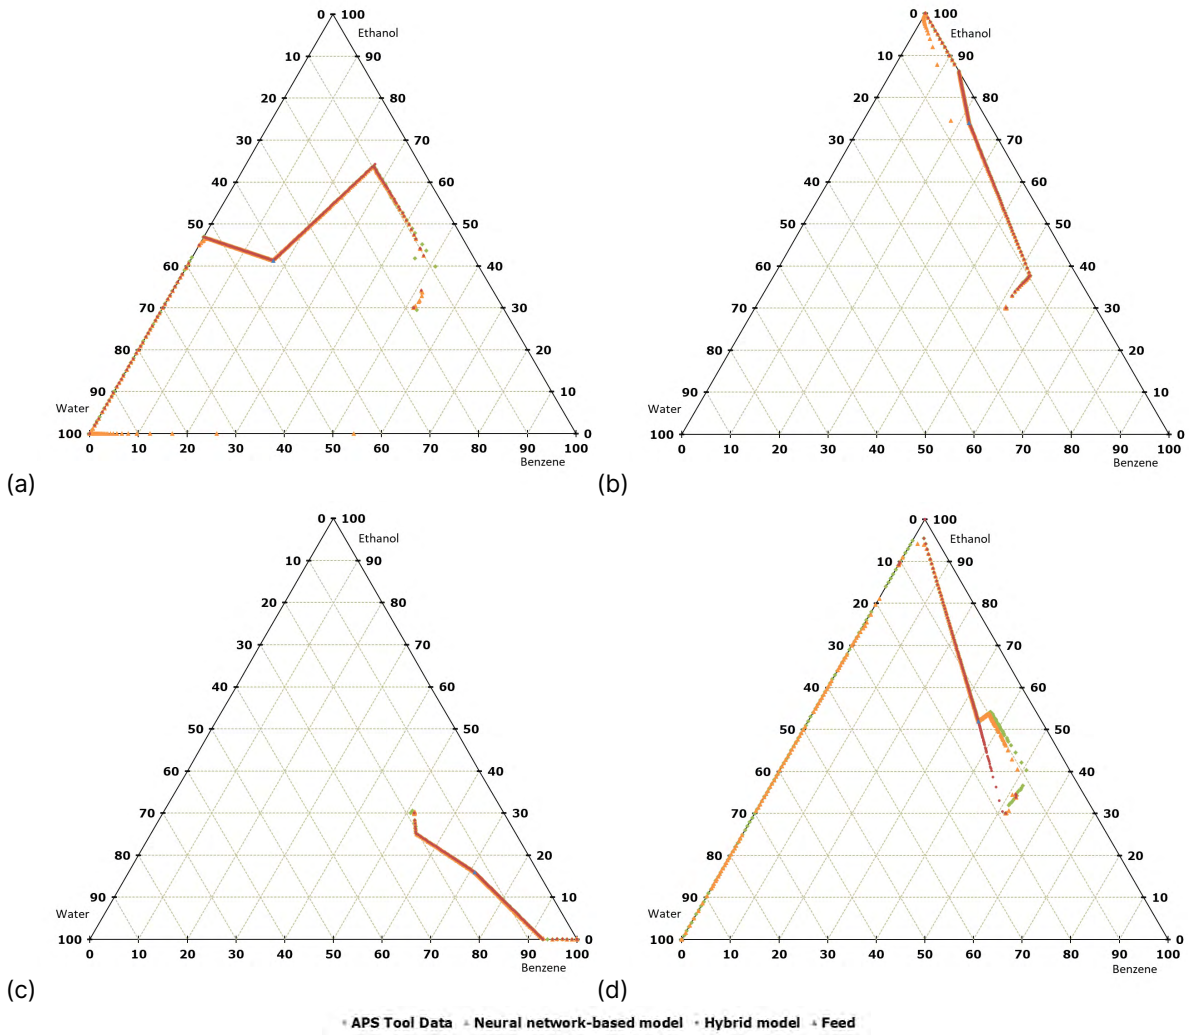

Figure 5. Hybrid surrogate vs. neural network-based product predictions for four feed compositions from different distillation regions

Nonetheless, due to its consistency and accuracy in predicting pure splits and its physically constrained behavior across operating regimes, the proposed hybrid surrogate provides a more reliable overall predictor than a purely neural network-based approach.

CONCLUSION

This work presents a structured, data-driven surrogate modeling framework for predicting distillation behavior under infinite-stage assumptions, operating within distillation limits informed by residue-curve behavior at total reflux and pinch point-based feasibility analysis. The framework provides a direct and deterministic mapping from feed composition and distillate-to-feed ratio to product compositions, making it well-suited for flowsheet synthesis and optimization applications in early conceptual design that require repeated and robust column evaluations. By decomposing the overall problem into a singular-point split classification task, an operating regime boundary regression, and a neural network-based predictor for impure splits, the approach combines physical interpretability with numerical robustness and computational efficiency.

The results demonstrate that the proposed framework reliably predicts feasible pure splits and operating regimes across the example ternary composition space of ethanol, benzene, and water. In particular, the combination of the classifier and regime regression outperforms a purely neural network-based predictor for pure splits, where exact reproduction of singular-point compositions is essential. The neural network, on the other hand, is well-suited for interpolating product compositions in the middle regime between limiting pure splits, as well as for feed compositions associated with greater uncertainty in the feasible bottom split. This hybrid

structure allows each modeling component to be applied where it is most effective, resulting in improved consistency and physical fidelity.

While the current implementation represents a significant step toward practical surrogate models for nonideal distillation systems, several opportunities for further improvement remain. The accuracy of the regime boundary predictions could be further refined, and the data generation workflow should be streamlined to reduce manual intervention and computational effort. In addition, the present framework is limited to fixed-pressure operation. Extending the model to include pressure as an additional input would require expanded training data and, in the case of azeotropic singular points, prediction of pressure-dependent azeotrope compositions. A longer-term objective is the development of a more general surrogate model that incorporates thermodynamic descriptors as inputs, enabling application to previously unseen chemical systems.

ACKNOWLEDGEMENTS

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) within the Priority Programme "SPP 2331: Machine Learning in Chemical Engineering."

AUTHOR IDENTIFIERS

Author ORCIDs:
Wolf AB: 0009-0007-0259-4518
Skiborowski M: 0000-0001-9694-963X
Burger J: 0000-0002-9376-9438

REFERENCES

1. Bausa J, Watzdorf RV, Marquardt W. Shortcut methods for nonideal multicomponent distillation: I. simple columns. AIChE Journal 44:2181-2198 (2004). https://doi.org/10.1002/aic.690441008
2. Sasi T, Skiborowski M. Automatic synthesis of distillation processes for the separation of homogeneous azeotropic multicomponent systems. Ind. Eng. Chem. Res. 59:20816-20835 (2020). https://doi.org/10.1021/acs.iecr.0c04555
3. Göttl Q, Pirnay J, Burger J, Grimm DG. Deep reinforcement learning enables conceptual design of processes for separating azeotropic mixtures without prior knowledge. Computers & Chemical Engineering 194:108975 (2025). https://doi.org/10.1016/j.compchemeng.2024.108975
4. Ryll O, Blagov S, Hasse H. ?/?-analysis of homogeneous distillation processes. Chemical Engineering Science 84:315-332 (2012). https://doi.org/10.1016/j.ces.2012.08.018
5. Bekiaris N, Morari M. Multiple steady states in distillation: ?/? predictions, extensions, and implications for design, synthesis, and simulation. Ind. Eng. Chem. Res. 35:4264-4280 (1996). https://doi.org/10.1021/ie950450d
6. O. Ryll, Thermodynamische Analyse gekoppelter Reaktions-Destillations-Prozesse: konzeptioneller Entwurf, Modellierung, Simulation und experimentelle Validierung, Ph.D. Thesis, University of Stuttgart 2009.
7. Gutiérrez-Antonio C. Multiobjective stochastic optimization of dividing-wall distillation columns using a surrogate model based on neural networks. Chem. Biochem. Eng. Q. 29:491-504 (2016). https://doi.org/10.15255/cabeq.2014.2132
8. Ibrahim D, Jobson M, Li J, Guillén-Gosálbez G. Optimization-based design of crude oil distillation units using surrogate column models and a support vector machine. Chemical Engineering Research and Design 134:212-225 (2018). https://doi.org/10.1016/j.cherd.2018.03.006
9. Wang Z, Zhou T, Sundmacher K. Data?driven integrated design of solvents and extractive distillation processes. AIChE Journal 69: (2023). https://doi.org/10.1002/aic.18236
10. Quirante N, Javaloyes J, Caballero JA. Rigorous design of distillation columns using surrogate models based on kriging interpolation. AIChE Journal 61:2169-2187 (2015). https://doi.org/10.1002/aic.14798
11. Keßler T, Kunde C, McBride K, Mertens N, Michaels D, Sundmacher K, Kienle A. Global optimization of distillation columns using explicit and implicit surrogate models. Chemical Engineering Science 197:235-245 (2019). https://doi.org/10.1016/j.ces.2018.12.002
12. Panofen M, Rolland T, Pela J, Skiborowski M. Surrogate models for distillation boundaries in azeotropic multicomponent mixtures. Ind. Eng. Chem. Res. 63:13723-13739 (2024). https://doi.org/10.1021/acs.iecr.4c00475
13. Prausnitz, J., Lichtenthaler, R., Azevedo, E. Molecular Thermodynamics of Fluid-Phase Equilibria, Prentice Hall PTR (1999).
14. Aspen Plus 8.8, Aspen Technology Inc.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Physics-informed Graph Neural Networks to Predict Thermodynamically Consistent Activity Coefficients in Multicomponent Mixtures

Lifeng Zhang[a], Benoît Chachuat[a], and Claire S. Adjiman[a]*

[a] Department of Chemical Engineering, The Sargent Centre for Process Systems Engineering, Imperial College London, London SW7 2AZ, United Kingdom
* Corresponding Author: c.adjiman@imperial.ac.uk.

ABSTRACT

Activity coefficients are key thermodynamic quantities for describing phase equilibria, but their experimental determination entails laborious and costly phase-equilibrium measurements, making predictive approaches highly desirable. The potential of machine learning for such predictions has received growing attention as an alternative to physics-based models that require experimental data or expensive calculations for parameterization. We propose a physics-informed edge-enhanced graph attention network (PEGAT) to predict activity coefficients in multicomponent mixtures, where each molecule is encoded as a graph in which the nodes correspond to atoms and the edges to chemical bonds. The excess Gibbs free energy of the mixture is predicted using the proposed model, including a nonlinear transformation in the final layer to ensure that the excess Gibbs free energy vanishes for pure components. To further enforce thermodynamic consistency, the relevant activity coefficients are obtained via the Gibbs–Duhem relation. Unlike machine-learning models developed primarily for binary systems, the proposed framework is directly applicable to arbitrary multicomponent mixtures. The PEGAT model is evaluated using a mixed dataset comprising both binary and ternary mixture data and demonstrates high predictive accuracy. Further validation on representative mixtures shows close agreement between predicted and reference activity coefficients. The results confirm that improved thermodynamic consistency can be achieved by embedding hard physical constraints into the graph neural network architecture. However, they also highlight that unphysical behaviors may still be predicted despite these constraints.

Keywords: Machine Learning, Graph Neural Network, Activity Coefficients, Physics-informed, Thermodynamic consistency

INTRODUCTION

Activity coefficients are fundamental thermodynamic quantities that can describe phase equilibria and reaction kinetics by accounting for non-ideal interactions through the chemical potentials of liquid mixtures [1]. The accurate prediction of activity coefficients is therefore essential for reliable modelling of vapor–liquid, liquid–liquid, and solid–liquid equilibria, as well as for assessing reaction feasibility in non-ideal liquid phases. Inaccurate activity-coefficient predictions can lead to erroneous estimates of phase behavior, such as spurious phase separation or miscibility gaps, which may ultimately result in infeasible process designs.

However, the experimental determination of activity coefficients is typically laborious, time-consuming, and costly, particularly for multicomponent mixtures where composition space grows combinatorially. Such experimental efforts are often impractical during early-stage molecular and/or process design, where the rapid evaluation of the properties of a large number of candidate mixtures and process conditions is required. Consequently, predictive thermodynamic models for activity coefficients play a central role in computer-aided mixture and blend design (CAM^{b}D) frameworks and in process design, enabling the efficient exploration of mixture

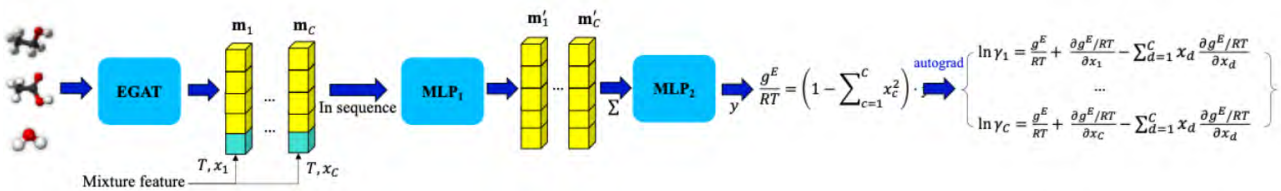

Figure 1. Schematic illustration of the proposed PEGAT model architecture. The edge-featured graph attention network (EGAT) converts molecules into latent space embeddings, followed by multi-layer perceptrons (MLPs) to capture the inherent nonlinearity. The derivatives with respect to composition are obtained through automatic differentiation (*Autograd* library).

space and/or informed optimization of process operating conditions [2].

A variety of thermodynamic models based on different physical principles have been developed for activity coefficient prediction, including local composition models such as NRTL [3], group-contribution approaches such as UNIFAC [4] and SAFT-γ Mie [5], and quantum chemical approaches exemplified by COSMO-RS [6]. Although these models have been successfully applied to a wide range of systems, their applicability is often limited by the availability of interaction parameters and the computational cost / effort required in modelling systems that fall outside of existing parameter sets. For instance, both UNIFAC and SAFT-γ Mie models rely on predefined group definitions and associated parameters, which restrict their use for molecules containing novel functional groups or lacking reliable parameterization. For COSMO-type models, costly quantum chemical calculations are required to generate the σ-profile data of new compounds. These limit the applicability of the methods for large-scale screening and early-stage process design. Overall, the transferability and generalization of these models to new mixtures, particularly those involving novel compounds or novel mixtures, remains limited.

Recently, supervised learning approaches have been widely leveraged to predict properties by learning structure–property relationships [7]. For example, Sanchez Medina et al. [8] trained a graph neural network (GNN) to predict infinite dilution activity coefficients at 298.15 K and achieved competitive performance compared to the UNIFAC and COSMO-RS methods. Winter et al. [9] employed a transformer model for binary activity coefficient prediction from SMILES for a temperature range from 273.15 to 598.15 K. Qin et al. [10] developed several SolvGNN models by adding hydrogen bonding features to capture intermolecular interactions and trained separate models for activity coefficients in binary and ternary mixtures, based on datasets generated from COSMO-RS at 298.15 K.

However, due to their black-box nature, machine learning models do not inherently guarantee the thermodynamic consistency of the predictions. To overcome this limitation, embedding physical knowledge is essential. Combining machine learning models with existing thermodynamic models has been attempted [11], but while effective, this approach is limited by the accuracy of the chosen thermodynamic model. Instead, Rittig et al. [12] developed a physics-informed GNN model for binary mixtures by imposing a Gibbs-Duhem penalty when training with a synthetic dataset at 298.15 K [10]. The model was later refined by embedding the Gibbs-Duhem relation as a hard constraint to obtain the activity coefficients in binary mixtures with the same dataset [13]. Specht et al. [14] proposed a hard-constraint neural network, called HANNA, trained on experimental data for binary mixtures at multiple temperatures. HANNA uses a pretrained model to convert a SMILES representation into a latent space and enforces thermodynamic consistency by incorporating the Gibbs-Duhem relation into the model as well as applying a nonlinear transformation to ensure that the predicted activity coefficients recover the correct behavior in the pure-component limit. HANNA was later improved by introducing a geometric projection method to represent the interactions between component pairs [14], thereby allowing the model to map the Gibbs excess free energy g^{E} trained from binary mixture data into multicomponent space. Nevertheless, such projections neglect the presence of higher-order interactions among the components when calculating the pairwise interactions.

Motivated by these earlier works, this paper proposes a physics-informed edge-enhanced graph attention network (PEGAT) model for thermodynamically consistent activity coefficient prediction. Rather than relying on a projection step, the PEGAT is designed to handle mixtures with an arbitrary number of components and to predict the overall Gibbs excess free energy directly. The corresponding activity coefficients are then computed using the Gibbs-Duhem relation. This formulation enables the model to be trained on heterogeneous, multi-source datasets that include not only binary mixtures but also higher-order multicomponent systems, with a potential to enhancing the model's generalisation capability.

Table 1: Overview of node features.

Index	Description
1-10	Atom type
11	Whether atom is aromatic
12	Whether atom is in a ring
13-23	Atom degree
24-34	Atom total degree
35-41	Explicit atom valence
42-48	Implicit atom valence
49-53	Hybridization type
54-58	Attached hydrogen numbers
59-63	Atom formal charge type
64-69	Atom radical electron
70-73	Atom chiral tag type
74-75	Atom chirality type
76	Atom mass
77	Whether atom is chiral center

Table 2: Overview of edge features.

Index	Description
1-4	Bond type
5	Whether bond is conjugated
6	Whether bond is aromatic
7	Whether bond is in a ring
8-13	Bond stereo type
14-15	Bond direction
16	Whether bond has bidirection

METHODOLOGY

The proposed PEGAT model architecture in Figure 1 implements an edge-featured graph attention network (EGAT) submodule to convert graphs into latent embeddings. These embeddings are further processed by two multi-layer perceptrons (MLPs) to predict the overall Gibbs excess free energy g^{E} and finally the activity coefficients as $\ln \gamma_c$. Details of the model can be found below.

Data Preparation

For the purpose of developing the model architecture, a synthetic dataset of activity coefficient values derived from COSMO-RS calculations is used [10], including 280, 000 binary and 160, 000 ternary mixture datapoints. The dataset is split into training (70%), validation (20%) and test (10%) subsets. Each molecule in a mixture is converted from a SMILES string into a graph $G = (V, E)$ using the *RDKit* package [15]. The nodes $v_i \in V$ represent the atoms in a molecule, and the edges $e_{ij} \in E$ describe the chemical bonds between atom i and j. With one-hot encoding, a 77-dimensional feature vector is generated for each node (Table 1) and a 16-dimensional feature vector for each edge (Table 2).

EGAT Module

Unlike other message-passing GNN models, the PEGAT model introduces an attention mechanism to update the node and edge features for obtaining a better representation. After a transformation through simple networks, each node $v_i \in V$ is associated with an initial atomic feature vector $\mathbf{h}_i^0 \in \mathbb{R}^h$ and each edge $e_{ij} \in E$ is associated with an initial bond feature vector $\mathbf{e}_{ij}^0 \in \mathbb{R}^e$.

At the lth layer and n^{th} attention head in EGAT, the unnormalised attention score between node i and node j, $s_{ij}^{l,n}$, is defined by incorporating the edge feature as follows:

$$s_{ij}^{l,n} = \mathrm{LeakyReLU}(\boldsymbol{\beta}^T[\mathbf{W}_h^n \mathbf{h}_i^l \parallel \mathbf{W}_h^n \mathbf{h}_j^l \parallel \mathbf{W}_e^n \mathbf{e}_{ij}^l]) \quad (1)$$

where $\mathbf{W}_h^n$ and $\mathbf{W}_e^n$ are learnable linear transformations for node and edge features at the nth attention head. The $\parallel$ refers to vector concatenation and $\boldsymbol{\beta}^T$ is the attention vector. The attention coefficients are obtained via softmax:

$$a_{ij}^{l,n} = \frac{\exp(s_{ij}^{l,n})}{\sum_{k \in \mathcal{N}_i} \exp(s_{ik}^{l,n})} \quad (2)$$

where $\mathcal{N}_i$ is the number of nodes that are connected to node i. The node feature is then updated via aggregating the neighbors in all the attention heads:

$$\mathbf{h}_i^{l+1} = \frac{1}{N} \sum_{n=1}^{N} \mathrm{LeakyReLU}\left(\sum_{j \in \mathcal{N}_i} a_{ij}^{l,n} \mathbf{W}_h^n \mathbf{h}_j^l\right) \quad (3)$$

where N is the total number of attention heads. Similarly, the edge feature vector for edge (i, j) at the $(l+1)^{\text{th}}$ layer is updated with the adjacent node features:

$$\mathbf{e}_{ij}^{l+1} = \mathrm{LeakyReLU}(\mathbf{W}_e^0[\mathbf{h}_i^l \parallel \mathbf{h}_j^l \parallel \mathbf{e}_{ij}^l]) \quad (4)$$

where $\mathbf{W}_e^0$ is a different linear transformation for updating the edge. Such EGAT transformation is applied to each molecule c in the mixture to obtain the final representation $h_{i,c}^L$ for each node at L^{th} layer, the embedding of the molecule c then formulated by summing up:

$$\mathbf{m}_c = \sum_{i \in I^c} h_{i,c}^L, c = 1, \ldots, C \quad (5)$$

where c represent the component in the mixture and I^c is the set of nodes in molecule c.

Thermodynamically Consistent Prediction

In addition to the structure feature vectors of all components, $\boldsymbol{m}_c, c = 1, \ldots, C$, the temperature and composition (mole fraction of each component, x_c) are added to the component input vectors which are individually converted via the first MLP, given by:

$$\mathbf{m}_c' = \mathrm{MLP}_1([\mathbf{m}_c \parallel T \parallel x_c]), c = 1, \ldots, C \quad (6)$$

The MLP_1 is defined with three layers whose dimensions

are $[513, 256, 64]$. The SiLU activation function is used to eliminate gradient vanishing in the model. Through MLP_1, the molecule embedding is now updated by capturing the temperature and respective composition information. The updated molecule representations are summed up to define the overall mixture before inputting into a second MLP, MLP_2, to obtain an initial scalar prediction y:

$$y = \mathrm{MLP}_2\left(\sum_{c=1}^{C} \mathbf{m}'_c\right) \tag{7}$$

Similar to MLP_1, the MLP_2 is defined as three layers of $[64, 256, 1]$. Also, SiLU is implemented as activation function here. The MLP_2 is converting the mixture representation into desired scalar output. While to ensure that g^{E} is zero for a pure component, a composition-dependent coefficient is introduced to obtain g^{E}:

$$\frac{g^{\mathrm{E}}}{RT} = \left(1 - \sum_{c=1}^{C} x_c^2\right) \cdot y \tag{8}$$

Lastly, the logarithm of the activity coefficient for each component c, $\ln \gamma_c$ is obtained through the GD relation:

$$\ln \gamma_c = \frac{g^{\mathrm{E}}}{RT} + \frac{\partial g^{\mathrm{E}}/RT}{\partial x_c} - \sum_{d=1}^{C} x_d \frac{\partial g^{\mathrm{E}}/RT}{\partial x_d}, c = 1, \ldots, C \tag{9}$$

Through these hard constraints, the GNN model now rigorously satisfies thermodynamic consistency, not only in terms of the Gibbs–Duhem relation, but also in terms of the limiting activity coefficients of pure components.

Model Training

The loss function $\mathcal{L}$ is taken as the mean squared error (MSE) for the logarithm of the activity coefficients:

$$\mathcal{L} = \frac{1}{N_p} \sum_{i=1}^{N_p} (\ln \gamma_i - \ln \tilde{\gamma}_i)^2 \tag{10}$$

where N_p refers to the total number of data points used in training. The PEGAT model is developed using the *DGL* package [16]. The Adam optimizer is used, with a batch size of 512, a learning rate of 0.001 and 200 epochs. The hyperparameters, including the learning rate, attention layers and number of neurons, are tuned with *optuna* [17]. With current settings, it takes around 3 hours to train this model on Windows11 using Intel(R) Xeon(R) Gold 6226R CPU @ 2.90GHz with 64GB RAM for data generation and NVIDIA RTX A4000 GPU with 16GB RAM for model training and evaluation.

Table 3: Performance evaluation of the PEGAT model on training, validation, and test datasets

Dataset	RMSE	MAE	R^2
Training	0.0665	0.0361	0.9922
Validation	0.0799	0.0388	0.9888
Test	0.0769	0.0387	0.9896
Binary mixtures	0.0723	0.0324	0.9924
Ternary mixtures	0.0683	0.0422	0.9891

RESULTS AND DISCUSSION

Overall Performance

The model is further evaluated on the binary dataset and the ternary dataset, respectively. The metric scores indicate that the model has comparable performance on these two datasets. For the ternary dataset, the RMSE is even smaller than those for binary mixtures. Given that the ternary dataset is smaller than the binary dataset, this indicates that the model enhances capturing the g^{E} function and activity coefficients in ternary mixtures from binary mixture data, contributing to the applicability of the model to ternary mixtures.

Performance on Binary Mixtures

Figure 2 compares the predicted $g^{\mathrm{E}}, \ln \gamma_1$, and $\ln \gamma_2$ values with available data [10] for two binary mixtures of acrolein-benzene and chloroform-hexane from the test dataset, at 298.15 K and 1 atm. The predicted g^{E} values for both mixtures in Figures 2(a, c) approach zero as the mixture composition approaches one of the two pure components, which is also consistent with the findings in [14]. Moreover, the imposed hard constraint (Eq. 7) guarantees the consistency of the behavior of the excess Gibbs free energy function, which is seen in Figures 2(b, d) where the respective $\ln \gamma_i$ becomes zero. Regarding the $\ln \gamma_i$ predictions, the model is in good agreement with the datapoints but is found to be of variable quality between points. Specifically, while the Gibbs–Duhem relation is satisfied at each value of x_1, the overall dependence on x_1 is sometimes unexpected. While the overall behavior of the curves for the acrolein-benzene mixture in Figure 2(a) is in agreement with the trend in the COSMO-RS data, including a maximum in $\ln \gamma_1$, multiple plateaus are observed between datapoints for the chloroform-hexane mixture in Figure 2(d). This latter behavior is unphysical and arises from the differentiation of the data-driven g^{E} model, as no constraints have been imposed on behavior of the derivatives during model training or in the model itself. This highlights one of the challenges in developing physically meaningful data-driven models.

Performance on Ternary Mixtures

Figure 3 compares the predicted g^{E} and $\ln \gamma_i$ values now for two ternary mixtures of cyclohexane-ethanol-methanol (CEM) and water-ethanol-benzene (WEB) from the test dataset, also at 298.15 K and 1 atm. For simplicity, the relative compositions of two of the components are set to $x_{\mathrm{ethanol}}{:}\,x_{\mathrm{methanol}} = 2{:}\,1$ in mixture CEM and $x_{\mathrm{ethanol}}{:}\,x_{\mathrm{benzene}} = 1{:}\,2$ in mixture WEB. Given that at most 5 datapoints are available for any given ternary mixture in the training dataset and the information relating to these specific ternary mixtures is limited, we use the COSMO-SAC [18] instead to generate the data for the

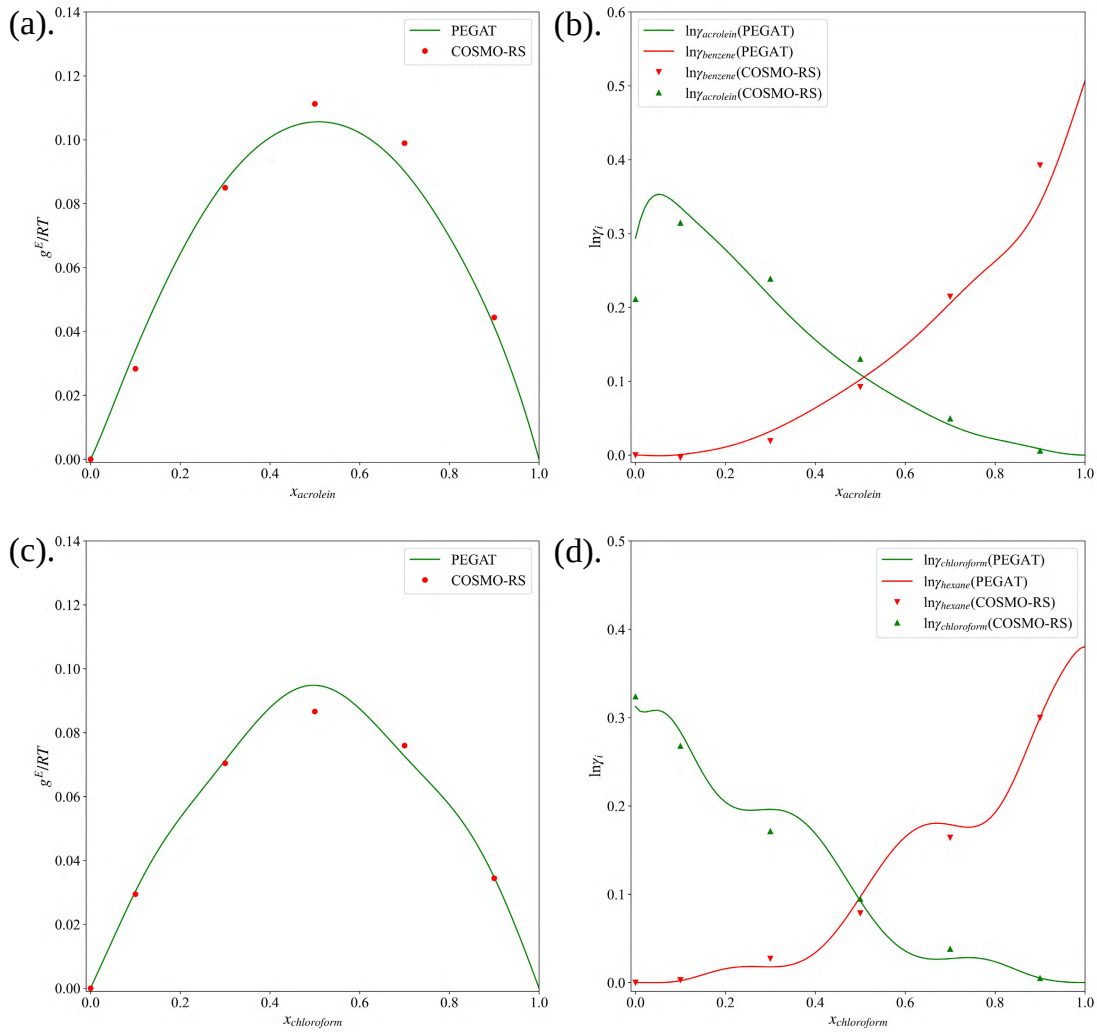

Figure 2: Calculated g^{E}, $\ln\gamma_1$ and $\ln\gamma_2$ values for two binary mixtures at $T = 298.15$ K and $P = 1$ atm as a function the mole fraction x_1. The curves are predictions with the proposed models and the symbols are values from the dataset of COSMO-RS calculations (a) g^{E} for acrolein (1)-benzene (2); (b) $\ln\gamma_1$ and $\ln\gamma_2$ values for acrolein (1)-benzene (2); (c) g^{E} for chloroform (1)-hexane (2); (d) $\ln\gamma_1$ and $\ln\gamma_2$ values for chloroform (1) -hexane (2).

plots as reference. To evaluate 100 data points for generating the diagrams, the PEGAT model takes only 0.05 seconds. While for the COSMO-SAC model, it can take up to 1.7 seconds if the σ-profiles have already existed; or otherwise, to hours if the σ-profiles need to be generated from beginning. This shows a significant reduction in computational expense with the proposed method. While it should be noted that the PEGAT model is trained based on the dataset from COSMO-RS and it is different from COSMO-SAC. Nevertheless, good quality fits are observed in Figures 3(a, c) for g^{E}, with the overall behavior following the trends in the data and good overall quantitative agreement, especially when $x_{\mathrm{cyclohexane}}$ and x_{water} approach 0 or 1. Furthermore, from Figures 3(b, d), the model demonstrates good agreement with the $\ln\gamma_i$ data at intermediate compositions, despite some deviations close to the 0 and 1 limits. Given the small amount of ternary mixture data used in training, this satisfactory performance confirms the value of combining binary and ternary data in model development. This can be helpful in enhancing performance for multicomponent mixtures, where generating a training dataset is challenging due to the large sampling space induced by the combinatorial explosion of possible components and compositions.

CONCLUSIONS AND FUTURE WORK

We have proposed a PEGAT model for the prediction of activity coefficients in multicomponent mixtures. The proposed model includes an EGAT module to convert molecular graphs into latent space embeddings. The edge-aware attention mechanism captures the interactions between atoms within the molecules. This is followed by a sequence of two MLPs that result in a

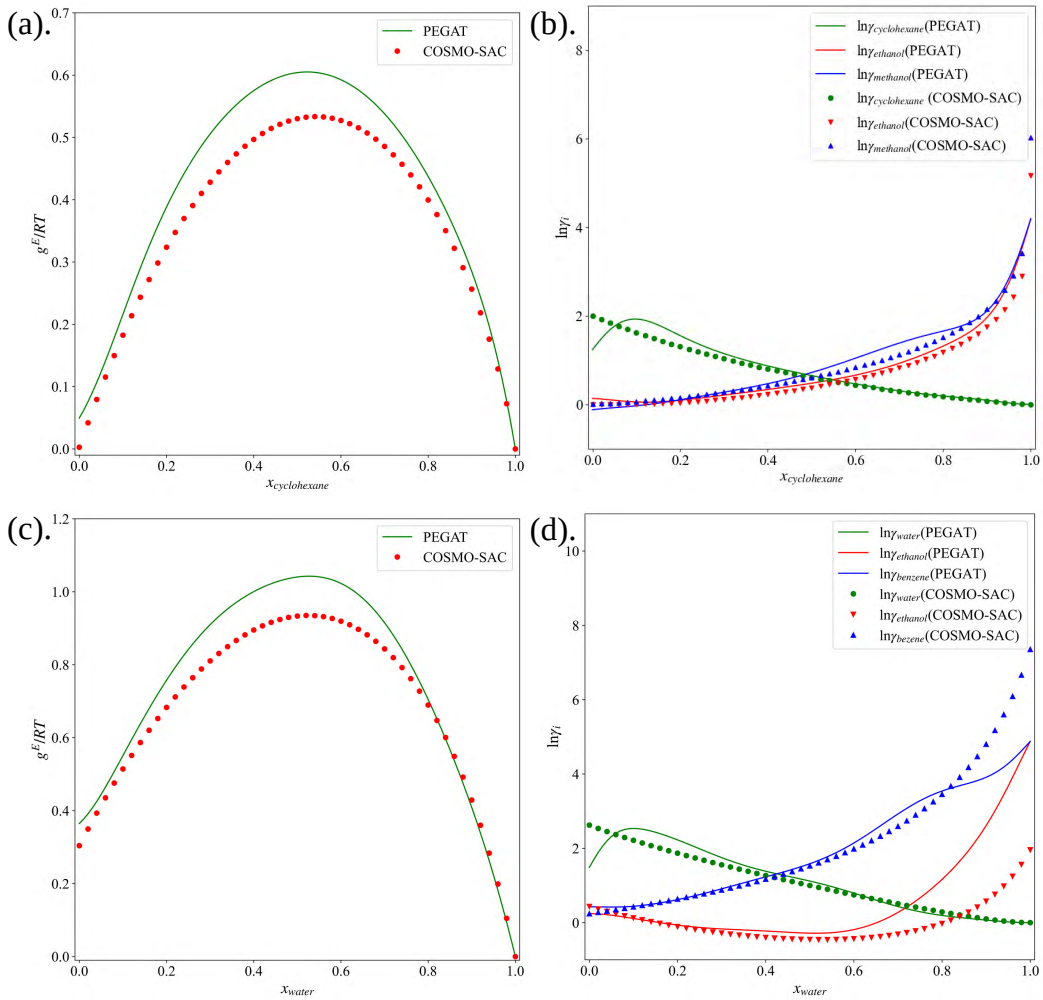

Figure 3. Values of g^{E} and $\ln\gamma_i$ for ternary mixtures at $T = 298.15$ K and $P = 1$ atm as a function of of the mole fraction of one components. The curves show the calculations with the proposed model and the symbols the COSMO-RS data (a) g^{E} and (b) $\ln\gamma_i$ for cyclohexane-ethanol-methanol with $x_{\mathrm{ethanol}}{:}\,x_{\mathrm{methanol}} = 2{:}\,1$; (c) g^{E} and (d) $\ln\gamma_i$ for water-ethanol-benzene with $x_{\mathrm{ethanol}}{:}\,x_{\mathrm{benzene}} = 1{:}\,2$.

prediction of the scalar excess Gibbs free energy. Thanks to a nonlinear transformation, g^{E} is guaranteed to approach zero for pure components, as is consistent with physical behavior. Then the activity coefficients are obtained via the Gibbs-Duhem relations by computing the derivatives of g^{E} with respect to compositions. These hard constraints ensure thermodynamic consistency in the predictions. Furthermore, the model is designed for mixtures with any number of components, which enables the use of a mixed dataset of binary and ternary mixture data for training. This architecture enhances the model's ability to learn from binary mixtures and extend to ternary mixtures, reducing the demand for expensive sampling. While good quality predictions are obtained, the interpolation behavior of the models can sometimes be unphysical due to use of derivatives for evaluation, highlighting the need to impose additional thermodynamic consistency requirements. In the future, more complex mixtures, such as quaternary mixtures, need to be furthermore validated to illustrate the potential of such model. The proposed model can be trained with experimental data and compared to existing thermodynamic models to demonstrate model's ability. Other alternative machine learning models, such as graph transformers, can be implemented to enhance learning both interactions within and between molecules for physical interpretability. The model could also be extended to include pressure as an input.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge funding from the UK Engineering and Physical Sciences Research Council (EPSRC) under grant EP/W003317/1.

REFERENCES

1. Fredenslund A., Jones R. L, Prausnitz J. M. Group-contribution estimation of activity coefficients in nonideal liquid mixtures. *AIChE Journal* **21**, 1086–1099 (1975).
2. Gani R. Computer-aided methods and tools for chemical product design. Chemical Engineering Research and Design 82:1494-1504 (2004). https://doi.org/10.1205/cerd.82.11.1494.52032
3. Vetere A. The NRTL equation as a predictive tool for vapor–liquid equilibria. Fluid Phase Equilibria 218:33-39 (2004). https://doi.org/10.1016/j.fluid.2003.10.013
4. Fredenslund A. UNIFAC and related group-contribution models for phase equilibria. Fluid Phase Equilibria 52:135-150 (1989). https://doi.org/10.1016/0378-3812(89)80320-6
5. Haslam AJ, González-Pérez A, Di Lecce S, Khalit SH, Perdomo FA, Kournopoulos S, Kohns M, Lindeboom T, Wehbe M, Febra S, Jackson G, Adjiman CS, Galindo A. Expanding the applications of the saft-? mie group-contribution equation of state: prediction of thermodynamic properties and phase behavior of mixtures. J. Chem. Eng. Data 65:5862-5890 (2020). https://doi.org/10.1021/acs.jced.0c00746
6. Scheffczyk J, Redepenning C, Jens CM, Winter B, Leonhard K, Marquardt W, Bardow A. Massive, automated solvent screening for minimum energy demand in hybrid extraction–distillation using COSMO-RS. Chemical Engineering Research and Design 115:433-442 (2016). https://doi.org/10.1016/j.cherd.2016.09.029
7. Wigh DS, Goodman JM, Lapkin AA. A review of molecular representation in the age of machine learning. WIREs Comput Mol Sci 12: (2022). https://doi.org/10.1002/wcms.1603
8. Sanchez Medina EI, Linke S, Stoll M, Sundmacher K. Graph neural networks for the prediction of infinite dilution activity coefficients. Digital Discovery 1:216-225 (2022). https://doi.org/10.1039/d1dd00037c
9. Winter B, Winter C, Schilling J, Bardow A. A smile is all you need: predicting limiting activity coefficients from SMILES with natural language processing. Digital Discovery 1:859-869 (2022). https://doi.org/10.1039/d2dd00058j
10. Qin S, Jiang S, Li J, Balaprakash P, Van Lehn RC, Zavala VM. Capturing molecular interactions in graph neural networks: a case study in multi-component phase equilibrium. Digital Discovery 2:138-151 (2023). https://doi.org/10.1039/d2dd00045h
11. Winter B, Winter C, Esper T, Schilling J, Bardow A. SPT-NRTL: a physics-guided machine learning model to predict thermodynamically consistent activity coefficients. Fluid Phase Equilibria 568:113731 (2023). https://doi.org/10.1016/j.fluid.2023.113731
12. Rittig JG, Felton KC, Lapkin AA, Mitsos A. Gibbs–duhem-informed neural networks for binary activity coefficient prediction. Digital Discovery 2:1752-1767 (2023). https://doi.org/10.1039/d3dd00103b
13. Rittig JG, Mitsos A. Thermodynamics-consistent graph neural networks. Chem. Sci. 15:18504-18512 (2024). https://doi.org/10.1039/d4sc04554h
14. Hoffmann M, Specht T, Göttl Q, Burger J, Mandt S, Hasse H, Jirasek F. Thermodynamically consistent machine learning model for excess gibbs energy. Nat Commun 17: (2026). https://doi.org/10.1038/s41467-026-71430-y
15. RDKit: Open-source cheminformatics. https://www.rdkit.org.
16. Wang M, Zheng D, Ye Z, Gan Q, Li M, Song X, Zhou J, Ma C, Yu L, Gai Y, Xiao T, He T, Karypis G, Li J, Zhang Z. Deep graph library: A graph-centric, highly-performant package for graph neural networks. *arXiv preprint* (2019). arXiv:1909.01315
17. Akiba T, Sano S, Yanase T, Ohta T, Koyama M. Optuna. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining :2623-2631 (2019). https://doi.org/10.1145/3292500.3330701
18. Antolovi? I, Stephan S, Vrabec J. High-throughput application and evaluation of the COSMO-SAC model for predictions of liquid–liquid equilibria. Digital Discovery 4:3191-3207 (2025). https://doi.org/10.1039/d5dd00259a

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

SMILE: Smell Maximisation In Low-cost Eau de parfum

Flora Esposito[a], Ulderico Di Caprio[a], Mattia Collu[b], Raffaele Graziano[b], Vincenzo Guida[b], Hasan Sildir[c], Idelfonso B.R. Nogueira[d], Florence Vermeire[e] and M. Enis Leblebici[a]*

[a] Process Engineering for Sustainable Systems, Department of Chemical Engineering, KU Leuven, Agoralaan Building B, 3590 Diepenbeek, Belgium
[b] Procter & Gamble, Brussels Innovation Center, Strombeek-Bever, Belgium
[c] Department of Chemical Engineering, Izmir Institute of Technology, Izmir 35430, Türkiye
[d] Chemical Engineering Department, Norwegian University of Science and Technology, Sem Sælandsvei 4, Kjemiblokk 5, Trondheim 793101, Norway
[e] KU Leuven, Department of Chemical Engineering, Celestijnenlaan 200F-bus 2424, Leuven 3001, Belgium
* Corresponding Author: muminenis.leblebici@kuleuven.be.

ABSTRACT

Despite the growing economic importance of the fragrance industry, perfume formulation remains largely guided by empirical knowledge and iterative trial-and-error approaches. The structured design of fragrances, typically organised into top, middle, and base notes through the blending of perfume raw materials, is therefore time-consuming, costly, and difficult to generalise. Existing computational approaches have begun to address this challenge, but are commonly limited to small ingredient sets or require extensive sensory data that are not always available. This work proposes a computer-aided optimisation framework for Eau de Parfum formulation that simultaneously maximises perceived olfactory intensity and minimises formulation cost. The resulting optimisation problem is formulated to preserve the structural balance of top, middle, and base notes inherent to the perfume pyramid. Application of the framework to a fruity–floral Eau de Parfum formulation demonstrates a substantial increase in perceived intensity alongside a reduction in total formulation cost of approximately 64%. The proposed approach is intended as a decision-support tool for perfumers, enabling rapid exploration of cost–intensity trade-offs once a candidate ingredient set has been defined. The framework offers a step towards more systematic, data-driven perfume engineering and provides a foundation for future extensions incorporating ingredient selection and sensory constraints.

Keywords: Optimisation, Industry 4.0, Perfume engineering, Formulation chemistry

INTRODUCTION

In 2017, the global Flavours and Fragrances industry was valued at over 26 billion US dollars, with projections indicating that it will exceed 36 billion US dollars by 2029 [1]. This industry spans a wide range of products, including fine fragrances, shampoos, candles, detergents, and fabric softeners [2]. Despite this sustained economic growth and the increasing interest in perfume engineering as a formal discipline, the literature still reports a notable lack of innovative, scientifically grounded tools for perfume formulation [3]. As a result, the design of fragrances remains largely dependent on empirical expertise and iterative trial-and-error procedures, which are time-consuming, costly, and difficult to generalise [4], [5]. Perfumes are complex mixtures of volatile compounds, typically dissolved in ethanol, water, or other carrier matrices, and are conventionally structured into top, middle, and base notes. In perfumery, the term notes refers to layers of scent that evaporate at different rates and collectively define the temporal evolution of a fragrance. Top notes consist of highly volatile compounds that dominate the initial sensory impression but dissipate rapidly. Middle notes, often described as the “heart” of the perfume, emerge shortly after application and establish its core olfactory character. Base notes comprise low-volatility compounds that persist for several hours and frequently act as fixatives by retarding the evaporation of more volatile constituents. The persistence and perceived intensity of each note are governed by

thermodynamic behaviour, mass-transfer phenomena, and psychophysical response mechanisms [2]. Carles conceptualised the perfumes structure as a pyramid divided into top, middle, and base notes, proposing typical proportions of 15–25%, 30–40%, and 45–55%, respectively [6]. These proportions were postulated to govern perfume tenacity, that is, the balanced evolution of odour intensity during evaporation.

The top, middle, and base notes described above are achieved by blending chemical ingredients known as perfume raw materials (PRMs), each contributing specific olfactory and volatilisation characteristics to the overall composition. Perfumers typically rely on experience-based heuristics to select and combine these PRMs, refining formulations through repeated sensory evaluation until the desired olfactory profile is achieved. This empirical methodology limits reproducibility and scalability, particularly given that modern cosmetic products often comprise more than 15 ingredients and may include up to 50 PRMs.

In recent years, several studies have tried to address these limitations through computational and data-driven approaches. Rodrigues et al. developed a machine learning-driven framework linking molecular structure to human olfactory perception, combining generative neural networks, odour-intensity prediction, thermodynamic modelling, and optimisation to reproduce target olfactory experiences from experimental data [3]. Viena Santana et al. proposed a systematic computational strategy for perfume formulation, employing deep learning trained on high-fidelity simulations alongside a meta-heuristic optimisation scheme to design compositions that achieve a desired odour spectrum while minimising undesirable notes [7]. Similarly, Zhang et al. introduced an optimisation-based framework for cosmetic formulation using a hierarchy of sensorial and product-property models, combined with iterative model substitution and generalised disjunctive programming, and demonstrated its applicability through a perfume design case study [8].

Despite their conceptual advances, these approaches suffer from several limitations. They typically focus on formulations containing a small number of PRMs (often three to four ingredients) or investigate the evaporation or performance of individual PRMs in isolation. Moreover, their application generally requires prior knowledge of odour detection thresholds for each ingredient, which are not always available or reliable. These constraints significantly limit their applicability to real-world cosmetic formulations, which are considerably more complex. Moreover, to the best of the authors' knowledge, the literature does not currently report any work aimed at mathematically optimising commercial perfume formulation.

In this work, an alternative optimisation framework for perfume formulation is proposed. Each PRM is characterised by its own odour intensity, reflecting differences in sensory experience between compounds. Consequently, two PRMs present at the same concentration will not necessarily be perceived at the same intensity level. The perceived odour strength is modelled explicitly in terms of olfactory intensity, based on a sigmoidal relationship between sensory response and the logarithm of stimulus concentration, consistent with established psychophysical models such as Hill's law, and the strongest-component model is employed to evaluate the overall intensity of the formulation [9], [10]. The approach integrates molecular property prediction, olfactory intensity modelling, and constrained numerical optimisation within a unified workflow. PRMs are characterised according to olfactory note classification, vapour–liquid equilibrium behaviour, intrinsic odour intensity, and economic cost, allowing for a quantitative assessment of their contribution to overall perfume performance. Olfactory intensities are computed from predicted gas-phase concentrations using established psychophysical models, while cost considerations are incorporated through a normalised economic metric. Crucially, the optimisation problem is formulated to preserve the structural balance of top, middle, and base notes inherent to the perfume pyramid, as illustrated in Figure 1, enabling the rational design of realistic Eau de Parfum formulations with industrial relevance.

METHOLOGY

The workflow of the current work is summarised in Figure 1. The proposed workflow begins with the selection of fragrance molecules, followed by the collection of key molecular characteristics, including olfactory note classification and economic information. These data were found through web scraping. In particular, molecules identifiers, molecular weights (MW_i), note type and character were found on Scent Tree, OQEMA Flavour & Fragrance Ingredients and The Good Scent Company. Commercial pricing data for each molecule were retrieved from Sigma-Aldrich, FluoroChem, Perfumer's Apprentice, and Pell Wall. To incorporate economic considerations, the cost per mole of each molecule was normalised between 10^{-3} and 1, generating a dimensionless cost metric suitable for inclusion in the optimisation problem. Moreover, parameters needed to characterise PRMs intensities were retrieved by Wakayama et al [9]. The selected molecules were then combined into fragrance formulations following a predefined Eau de Parfum composition, comprising 75 wt% ethanol, 20 wt% PRMs, and 5 wt% water. In this study, the initial composition was defined to mimic a realistic fragrance pyramid, wherein molecules are categorised as top, middle, or base notes. Each molecule was initially assigned a weight fraction corresponding to its contribution to the total perfume

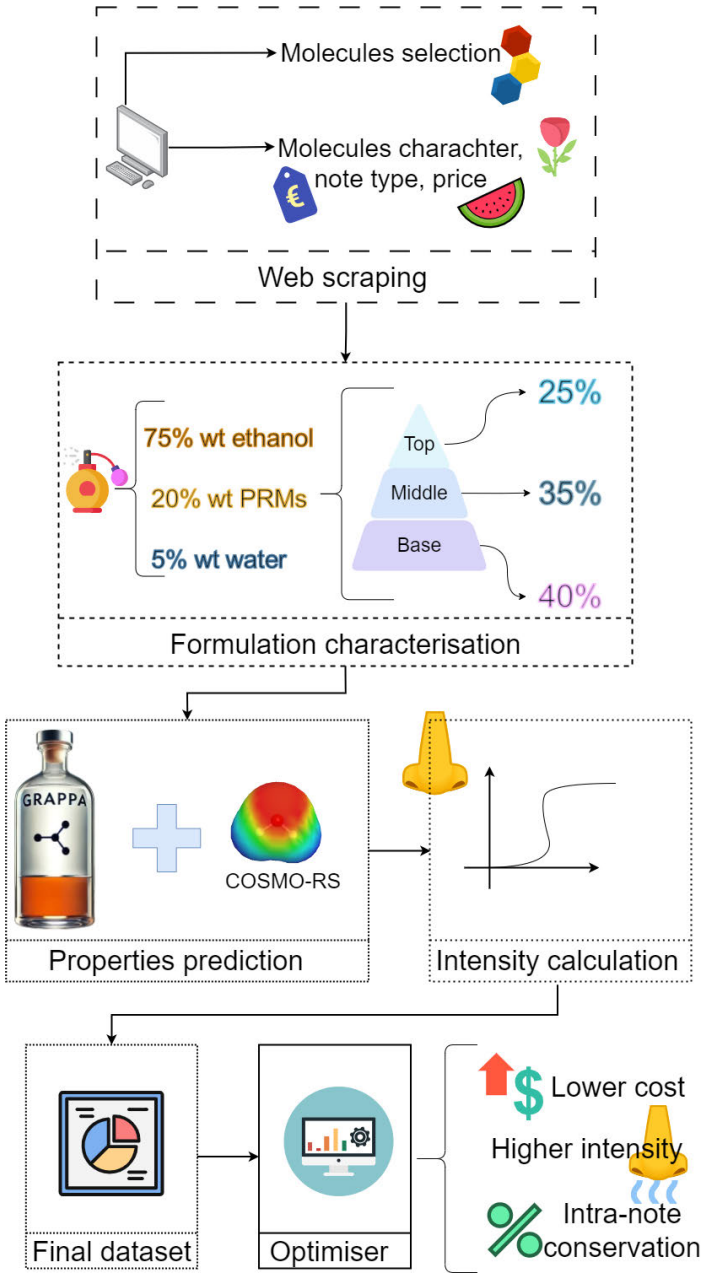

Figure 1. Schematic overview of the workflow adopted in this study. The process begins with web scraping to select and characterise PRMs. The selected molecules are then assembled into an Eau de Parfum formulation following predefined note-distribution constraints. Once the formulation is defined, relevant vapour–liquid equilibrium properties are predicted and used to calculate olfactory intensities. Finally, the molar fractions of the PRMs are optimised to maximise fragrance intensity while minimising formulation cost and preserving the relative proportions of top, middle, and base notes.

fraction. The relative distribution of weight fractions within each note category was determined to preserve intranote balance, reflecting the perfumer's approach of proportionally blending molecules within top, middle, and base notes to achieve the characteristic volatility and olfactory structure of the formulation (e.g. among 20%wt PRMs, 25% wt of top notes, 35% wt of middle notes and 40% wt of base notes)

Once the formulation is fully characterised, its vapour-liquid equilibrium properties were predicted. Activity coefficients (γ_i) at 36 °C (corresponding to human skin temperature) were calculated using COSMOTherm [11] (version 24.00), while vapour pressures (P_i^0) were estimated using GRAPPA [12]. COSMOTherm was

Table 1: Characterisation of the formulation developed for the current work, along with optimisation results. First, molecule identifiers and notes are reported. Then, activity coefficients, normalised costs, initial molar fraction and intensities are presented. Finally, optimised molar fractions and intensities are presented.

CAS	Name	Note	Note type	γ_i	Norm.cost	x_i^0	Ψ_i^0	x_i^{opt}	Ψ_i^{opt}
103-45-7	2-Phenylethyl Acetate	Floral	Middle	4.38	0.029	0.0027	0.0018	0.0099	0.0178
105-54-4	Ethyl Butanoate	Fruit	Top	2.85	0.157	0.0059	16.633	0.0083	20.754
6259-76-3	Hexyl Salicylate	Floral	Base	14.22	0.001	0.0023	0.0152	0.0094	0.1186
103-37-7	Benzyl Butyrate	Fruit	Middle	5.43	0.084	0.0045	$3.31 \cdot 10^{-6}$	$1 \cdot 10^{-6}$	$1.69 \cdot 10^{-16}$
10339-55-6	Ethyl Linalool	Floral	Middle	1.34	0.134	0.0018	0.0042	$1 \cdot 10^{-6}$	$8.26 \cdot 10^{-8}$
105-87-3	Geranyl Acetate	Floral	Middle	6.17	0.066	0.0011	$2.32 \cdot 10^{-6}$	0.0037	$7.11 \cdot 10^{-5}$
7452-79-1	Ethyl 2-Methylbutyrate	Fruit	Middle	2.90	0.088	0.0011	19.615	0.0064	39.246
103-60-6	Phenirat	Fruit	Middle	6.30	0.122	0.0010	$1.05 \cdot 10^{-13}$	$1 \cdot 10^{-6}$	$4.47 \cdot 10^{-28}$
118-58-1	Benzyl Salicylate	Floral	Base	9.36	0.058	0.0044	0.0044	$1 \cdot 10^{-6}$	$3.59 \cdot 10^{-9}$
122-40-7	Alpha-Amylcinnamaldehyde	Floral	Middle	6.31	0.123	0.0015	$1.72 \cdot 10^{-7}$	$1 \cdot 10^{-6}$	$8.41 \cdot 10^{-17}$
140-11-4	Benzyl Acetate	Floral	Middle	3.73	0.144	0.0019	0.0005	$1 \cdot 10^{-6}$	$1.09 \cdot 10^{-11}$
101-84-8	Diphenyloxide	Floral	Base	7.81	0.020	0.0030	0.0001	0.0099	0.0019
103-05-9	Benzyl tert-Butanol	Floral	Middle	1.02	1.000	0.0013	$1.39 \cdot 10^{-6}$	$1 \cdot 10^{-6}$	$2.60 \cdot 10^{-13}$
123-92-2	Isoamyl Acetate	Fruit	Top	3.07	0.259	0.0053	3.3286	$1 \cdot 10^{-6}$	0.0003
24851-98-7	Methyl Dihydrojasmonate	Floral	Middle	3.10	0.240	0.0016	$1.45 \cdot 10^{-9}$	$1 \cdot 10^{-6}$	$2.99 \cdot 10^{-18}$
106-25-2	Nerol	Floral	Middle	1.41	0.170	0.0029	$6.87 \cdot 10^{-6}$	$1 \cdot 10^{-6}$	$1.53 \cdot 10^{-12}$
120-72-9	Indole	Floral	Base	0.58	0.056	0.0033	$4.89 \cdot 10^{-5}$	$1 \cdot 10^{-6}$	$1.51 \cdot 10^{-11}$
57378-68-4	Delta-Damascone	Floral	Base	2.17	0.147	0.0039	0.0285	$1 \cdot 10^{-6}$	$2.86 \cdot 10^{-7}$
63500-71-0	Florosa Q	Floral	Base	0.58	0.032	0.0044	$6.16 \cdot 10^{-5}$	0.0020	$1.48 \cdot 10^{-5}$
78-70-6	Linalool	Floral	Top	1.01	0.011	0.0072	0.1569	0.01	0.2617

additionally employed to predict the pure-component densities of PRMs at ambient temperature, allowing all cost data to be expressed on a consistent molar basis, as supplier prices can be reported either by mass or by volume.

Optimisation problem formulation

The contribution of each molecule was quantified through its intensity Ψ_i:

$$\Psi_i = \frac{I_{max,i}}{1+e^{(-X_i - C_i/D_i)}} \quad (1)$$

Where X_i is the logarithm of the concentration of the ith PRM in the gas phase (µg/L air) and $I_{max,i}$, C_i, and D_i are parameters specific to each PRM: $I_{max,i}$ is the odour intensity at saturated gas concentration, C_i is the value of X_i when Ψ_i is $I_{max,i}/2$, and D_i is the steepness of the function. These parameters were retrieved by Wakayama et al [9]. Considering the ideal gas law, $X_i = {P_i}/{RT}$, where P_i (Pa) is the equilibrium partial pressure, R is the ideal gas constant (L Pa K^{-1} mol^{-1}), and T is the temperature (K). From Raoult's law $P_i = \gamma_i x_i P_i^0$, where x_i is the ith molecule molar fraction. Therefore, X_i is a function of the x_i as follows:

$$X_i = \log\left(\frac{\gamma_i x_i P_i^0}{RT} MW_i \, 10^3\right) \quad (2)$$

The objective function was defined as a weighted combination of total perceptual intensity and economic cost:

$$x_i^* = argmin\left(-\sum_{i=1}^{N} \Psi_i(x_i) + \lambda \sum_{i=1}^{N} c_i \, x_i\right) \quad (3)$$

Where c_i is the normalised cost of the molecule and λ is a tunable parameter balancing the importance of cost minimisation against intensity maximisation. In this case, $\lambda = 1$. The optimisation was subject to multiple equality constraints to ensure the structural integrity of the perfume formulation. Specifically, the sum of molar fractions within each note category (top, middle and base) was constrained to match the initial total fraction of that note type, reflecting the intranote balance inherent to the perfume pyramid:

$$\sum_{i \in Top} x_i = \sum_{i \in Top} x_i^0 = 0.0183, \quad (4)$$

$$\sum_{i \in Middle} x_i = \sum_{i \in Middle} x_i^0 = 0.0200 \quad (5)$$

$$\sum_{i \in Base} x_i = \sum_{i \in Base} x_i^0 = 0.0216 \quad (6)$$

The rationale for these constraints is as follows. The

initial molar fractions x_i^0 are derived from the target weight fractions of the perfume pyramid by converting each molecule's assigned weight fraction using its molecular weight and density. Consequently, the sums $\sum_{i\in Top} x_i^0$, $\sum_{i\in Middle} x_i^0$, and $\sum_{i\in Base} x_i^0$ encode the desired note-category proportions at the molar level. Enforcing Equations (4)–(6) as equality constraints during optimisation ensures that the total molar amount of each note category remains unchanged. The optimiser may therefore freely redistribute molar fractions among molecules within a note category, favouring those with higher intensity or lower cost, while the total molar amount of each note category, and hence the overall molar proportion of top, middle, and base notes in the Eau de Parfum, is held constant.

Additionally, target intensity constraints were imposed for specific olfactory categories, such as floral and fruity molecules, to ensure the final composition retained desired sensory characteristics:

$$\sum_{i\in Floral} \Psi_i(x_i) = \Psi_{Floral}^{target}, \quad (7)$$

$$\sum_{i\in Fruit} \Psi_i(x_i) = \Psi_{Fruit}^{target} \quad (8)$$

Where $\Psi_{Floral}^{target} = 0.4$ and $\Psi_{Fruity}^{target} = 60$. It is important to note that, from a psychophysical perspective, olfactory intensities are not strictly additive, and summing individual component intensities does not represent the true perceptual intensity of a fragrance. Nevertheless, the use of summed intensities is an easy and convenient mathematical construct within the optimisation framework, enabling the enforcement of compositional constraints on specific olfactory families. The optimisation problem was solved numerically using the Sequential Least Squares Programming (SLSQP) algorithm implemented in SciPy optimize minimize. x_i bounds ranged from 10^{-6} to 10^{-2}. Given the low concentration regime, infinite dilution conditions were assumed; accordingly, variations in X_i in Equation (2) arise solely from changes in x_i, while γ_i is kept constant.

Convergence was monitored using a 10^{-6} tolerance on the objective function, and the final solution was verified by computing the molar fractions of each note type to confirm they remained conserved within a 10^{-6} difference, thereby ensuring that the perfume pyramid structure was maintained.

RESULTS

Table 1 compares the initial and optimised compositions of the fragrance formulation in terms of molar fraction and olfactory intensity. A clear redistribution of material is observed following optimisation, reflecting the objective function, which maximises the sum of olfactory intensities while reducing economic cost. Molecules characterised by high initial intensities (Ψ_i^0) and relatively low normalised costs are preferentially selected and assigned increased molar fractions. In particular, Ethyl 2-methylbutyrate, Diphenyloxide, Linalool, and Hexyl salicylate exhibit substantial increases in molar fraction and corresponding olfactory intensity after optimisation. Among these, Ethyl 2-methylbutyrate emerges as a dominant contributor, with its optimised intensity ($\Psi_i^{opt} \approx 39.2$) significantly exceeding its initial value, while remaining economically efficient. Under the strongest component model, the perceived intensity of a perfume is governed primarily by the molecule with the highest individual intensity contribution. According to this criterion, the optimised formulation is strongly dominated by Ethyl 2-methylbutyrate, whose intensity far exceeds that of all other components. Compared to the initial formulation, where the maximum individual intensity is substantially lower, the optimisation therefore leads to a pronounced increase in overall perceived perfume intensity.

Conversely, several molecules are effectively excluded from the optimised formulation, with their molar fractions reduced to the imposed lower bound. This behaviour is particularly evident for compounds exhibiting extremely low initial intensities, such as Benzyl butyrate, Phenirat, Alpha-amylcinnamaldehyde, Ethyl linalool, Methyl dihydrojasmonate, and Benzyl salicylate. Despite some of these molecules having moderate or even low normalised costs, their negligible olfactory contributions render their inclusion unjustifiable within the optimisation framework. In addition, molecules associated with comparatively high costs and modest intensities, such as Benzyl tert-butanol and Isoamyl acetate, are strongly penalised and largely removed from the formulation. In addition to the objective function, explicit constraints were imposed on the summed olfactory intensities of fruity and floral components, as defined in Equations (7) and (8). The target floral and fruity intensities were set to Ψ_{Floral}^{target}= 0.4 and Ψ_{Fruit}^{target}= 60, respectively, compared to initial summed values of approximately 0.21 for floral notes and 39.6 for fruity notes. Following optimisation, both constraints are satisfied exactly, with optimised floral and fruity intensities of 0.400 and 60.0, respectively.

Under this formulation, the total summed intensity increases from 39.8 in the initial composition to 60.4 after optimisation. However, the actual perceived intensity of the perfume remains governed by the strongest component model and is therefore dominated by Ethyl 2-methylbutyrate.

Beyond sensory performance, the optimisation yields a substantial economic benefit. The total normalised cost of the formulation (calculated as $\sum_{i=1}^{N} c_i\, x_i$) is reduced from 0.00762 in the initial composition to 0.00279 after optimisation, corresponding to a cost reduction of approximately 64%. Importantly, this reduction is achieved simultaneously with an increase in overall perceived intensity, as dictated by the strongest component

model. Lastly, the influence of the cost–intensity trade-off parameter λ in the objective function was also examined. For the results presented above, λ = 1 was used, corresponding to equal weighting of intensity maximisation and cost minimisation. When λ was set to 0, such that only olfactory intensity was maximised, the optimiser selected a substantially larger number of molecules, including high-cost components. Although this intensity-only optimisation still reduced the total normalised cost from approximately 0.007 to 0.005, the reduction was markedly smaller than in the λ = 1 case. This comparison highlights that λ acts as a tunable formulation parameter, allowing the optimisation to prioritise either maximal intensity or economic efficiency depending on the desired design objective. The sensitivity of the optimisation to λ was further explored by increasing its value to 100, thereby strongly prioritising cost minimisation. In this specific case, no substantial changes in the optimised formulation were observed relative to the λ = 1 solution. This behaviour arises from the limited number of candidate components and, more importantly, from the presence of strict compositional constraints, including the perfume pyramid structure and the imposed floral and fruity intensity targets. These constraints must be satisfied simultaneously, which necessarily enforces the inclusion of certain molecules regardless of the cost penalty applied. This outcome reflects a limitation of the formulation space defined by the modeller rather than a failure of the optimisation strategy itself. In practical perfumery, expert knowledge would guide the initial selection and balancing of ingredients, preventing excessive dominance of individual components while ensuring that structural and sensory requirements are met without reliance on extreme parameter tuning.

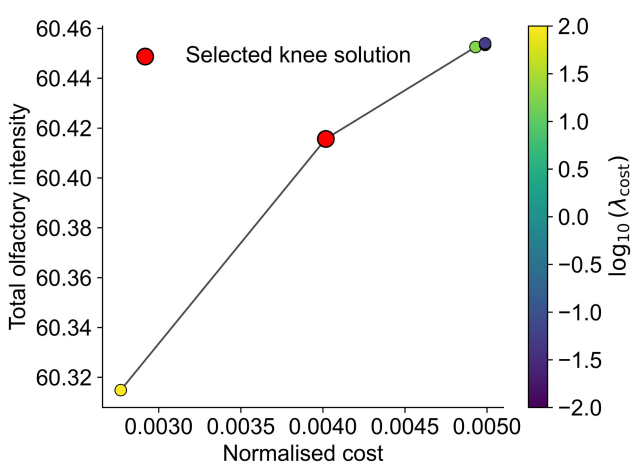

Figure 2. Pareto front: intensity–cost trade-off.

Pareto Front results

To further investigate the trade-off between olfactory performance and formulation cost, a Pareto-front analysis was performed by sweeping the cost regularisation parameter, λ_{cost}, over a logarithmically spaced grid and solving the constrained optimisation problem for each value.

In addition to the original note-fraction constraints (Equations (7)-(8)), soft constraints were introduced on the total floral and fruit intensities, allowing them to vary within 0.8–1.5 times their respective initial values. The resulting solutions define Pareto curves describing the compromise between total olfactory intensity and normalised cost.

An automatic knee-selection criterion based on the normalised intensity–cost trade-off was then applied to identify an optimal balance between performance and cost. The selected solution (λ_{cost} ≈ 43.3) yields a substantial increase in total intensity (from ≈ 39.8 to ≈ 60.4) alongside a marked reduction in normalised cost (from ≈ 0.0071 to ≈ 0.0040), while satisfying all constraints exactly. The floral and fruit intensities increase from 0.21 to 0.51 and from 39.6 to 59.9, respectively. This Pareto-selected formulation closely aligns with the previously reported optimisation case corresponding to λ = 100, both in terms of the selected molecules (seven components are retained with non-negligible contributions) and the overall composition profile, with most remaining components pushed towards their lower bounds. Moreover, inspection of the Pareto front indicates diminishing returns in total intensity for normalised costs above approximately 0.003–0.004, suggesting that further cost increases yield only marginal performance gains.

CONCLUSIONS AND PERSPECTIVES

In this work, a systematic optimisation framework for Eau de Parfum formulation has been proposed, aimed at simultaneously maximising olfactory intensity and minimising economic cost. By coupling physicochemical property prediction, intensity modelling, and constrained numerical optimisation, the methodology provides a quantitative tool capable of redistributing PRMs in a cost-efficient manner while preserving the structural integrity of the fragrance pyramid. The results demonstrate that significant gains in perceived intensity can be achieved alongside substantial reductions in formulation cost, highlighting the potential of optimisation-based approaches to support perfume engineering and formulation design. It should be emphasised that, in the present study, fragrance molecules were initially selected in a simplified manner, based primarily on their fruity and floral character, with olfactory intensities computed a posteriori. This led to a pronounced imbalance in intrinsic intensity across the candidate set, whereby a small number of molecules exhibited markedly higher intensity contributions than others. As a consequence, the optimiser was naturally driven towards selecting and amplifying

these dominant components, in accordance with both the objective function and the model of olfactory perception. While this behaviour is fully consistent with the mathematical formulation of the problem, it does not reflect the typical workflow of a professional perfumer, who would usually pre-select ingredients with broadly comparable olfactory strength in order to achieve a balanced sensory profile. From this perspective, the proposed framework is not intended to replace the creative expertise of the perfumer, but rather to act as a decision-support tool that accelerates and rationalises the formulation process. When applied to a candidate set already curated by a perfumer, based on experience, artistic intent, and desired olfactory balance, the optimisation tool can rapidly identify cost–intensity trade-offs and suggest optimal molar fractions that respect both sensory constraints and economic targets. In this sense, the methodology offers a means to formalise and streamline stages of formulation refinement that are traditionally addressed through iterative trial-and-error experimentation. Future work will focus on extending the optimisation framework to include a preliminary selection stage, in which candidate molecules are dynamically pooled and screened based on their intensity-to-cost ratios. Such an approach would enable the optimiser not only to determine optimal molar fractions, but also to identify the most promising subset of ingredients from a larger candidate library. The integration of molecule selection and composition optimisation within a unified framework would represent a significant step towards fully computer-aided perfume formulation, with the potential to further enhance efficiency, reduce development costs, and support innovation in fragrance design.

AUTHOR IDENTIFIERS

Author ORCIDs:
Esposito Flora: 0009-0008-9396-1922
Ulderico Di Caprio: 0000-0001-5194-8721
Mattia Collu: 0000-0003-2547-9057
Raffaele Graziano: 0000-0002-5767-3033
Vincenzo Guida: 0000-0003-0489-5596
Hasan Sildir: 0000-0003-1016-9865
Idelfonso B.R. Nogueira: 0000-0002-0963-6449
Florence Vermeire: 0000-0002-5607-8152
M. Enis Leblebici: 0000-0003-4599-9412

REFERENCES

1. "Flavors and Fragrances Market Share, Size, Report, 2024-2032, " 2024.
2. Rodrigues AE, Nogueira I, Faria RPV. Perfume and flavor engineering: a chemical engineering perspective. Molecules 26:3095 (2021). https://doi.org/10.3390/molecules26113095
3. Rodrigues BCL, Santana VV, Murins S, Nogueira IBR. Molecule generation and optimization for efficient fragrance creation. Ind. Eng. Chem. Res. 63:14480-14494 (2024). https://doi.org/10.1021/acs.iecr.4c00650
4. M. A. Teixeira, O. Rodríguez, G. Paula, V. Mata, and A. E. Rodrigues, *Perfume Engineering*. Oxford: Elsevier, 2013. doi: https://doi.org/10.1016/C2012-0-03624-1
5. Wolfson W. In the fragrance business, the right molecule smells like money. Chemistry & Biology 12:857-858 (2005). https://doi.org/10.1016/j.chembiol.2005.08.005
6. J. Carles, "A method of creation in perfumery, " *Soap Perfumery Cosmet*, vol. 36, pp. 328–335, 1962.
7. Santana VV, Martins MAF, Loureiro JM, Ribeiro AM, Rodrigues AE, Nogueira IBR. Optimal fragrances formulation using a deep learning neural network architecture: a novel systematic approach. Computers & Chemical Engineering 150:107344 (2021). https://doi.org/10.1016/j.compchemeng.2021.107344
8. Zhang X, Zhou T, Ng KM. Optimization?based cosmetic formulation: integration of mechanistic model, surrogate model, and heuristics. AIChE Journal 67: (2020). https://doi.org/10.1002/aic.17064
9. Wakayama H, Sakasai M, Yoshikawa K, Inoue M. Method for predicting odor intensity of perfumery raw materials using dose–response curve database. Ind. Eng. Chem. Res. 58:15036-15044 (2019). https://doi.org/10.1021/acs.iecr.9b01225
10. Lee C. Visualization of multidimensional database. SPIE Proceedings 6809:68090M (2008). https://doi.org/10.1117/12.768405
11. Dupeux T, Gaudin T, Marteau?Roussy C, Aubry J, Nardello?Rataj V. COSMO?RS as an effective tool for predicting the physicochemical properties of fragrance raw materials. Flavour & Fragrance J 37:106-120 (2022). https://doi.org/10.1002/ffj.3690
12. Hoffmann M, Hasse H, Jirasek F. GRAPPA—A hybrid graph neural network for predicting pure component vapor pressures. Chemical Engineering Journal Advances 22:100750 (2025). https://doi.org/10.1016/j.ceja.2025.100750

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Dynamic Operation of a Haber-Bosch Loop with Quench-Cooled Converter for Power-to-Ammonia Systems

José M. Pires[ab], Diogo A. C. Narciso[b], and Carla I. C. Pinheiro[a]*

[a] Centro de Química Estrutural, Institute of Molecular Sciences, Department of Chemical Engineering, Instituto Superior Técnico, Universidade de Lisboa, 1049-001, Lisboa, Portugal
[b] Centro de Recursos Naturais e Ambiente, Department of Chemical Engineering, Instituto Superior Técnico, Universidade de Lisboa, 1049-001, Lisboa, Portugal
* Corresponding Author: carla.pinheiro@tecnico.ulisboa.pt

ABSTRACT

This work presents a preliminary application of a developing methodology for assessing the operational flexibility of ammonia synthesis loops. Part of this methodology involves systematic dynamic testing in the synthesis loop. In the present case, a synthesis loop equipped with a two-bed, quench-cooled converter operating under variable feed conditions was considered. A high-fidelity model of the converter was developed in gPROMS Process using two-dimensional reactor models from its fixed-bed catalytic reactor library, and a synthesis loop configuration was modeled and designed in the same environment. A series of dynamic tests varied the make-up flow rate across four disturbance amplitudes (±25% and ±50%) and four disturbance durations (10 s, 600 s, 1800 s, 3600 s). The results showed that disturbances of ±50% magnitude led to the violation of one process operational constraint. These findings enable the construction of a preliminary operational map of this system, providing an initial determination of its process boundaries, and therefore its current degree of flexibility. Such knowledge is fundamental to understand how the synthesis loop can be made more flexible so that it suits Power-to-Ammonia applications. Overall, this study constitutes an initial step toward the development of a framework for enhancing the flexibility of ammonia synthesis loops.

Keywords: Dynamic Modelling, gProms, Process Design, Process Flexibility, Green Haber-Bosch Process

INTRODUCTION

In recent years, ammonia (NH_3) has attracted significant attention as a key piece on the strategic chessboard of energy transition, playing an important role in the value chain of green hydrogen. Ammonia possesses a higher volumetric energy density than hydrogen, making it easier to store. Furthermore, its handling and transport infrastructure is already well-established [1, 2]. Altogether, these features render ammonia an extraordinary carbon-free chemical energy carrier for hydrogen. Besides its potential novel role in the energy transition, ammonia is already one of the most important commodity chemicals worldwide. The expansion of its applications to include energy-related uses complements its traditional markets and is expected to further increase global demand. In 2021, the global ammonia market was valued at $US 67 billion, representing about 5% of the total chemical market value [3]. Looking ahead, ammonia demand is projected to increase to 223 million metric tons by 2030 [4, 5]. Despite these new applications, fertilizer production currently accounts for the majority of ammonia consumption, representing approximately 85% of the total demand, and remains its most significant use.

The Haber-Bosch (HB) process is the industrially relevant process when it comes to ammonia production, accounting for over 90% of global ammonia produced [1]. Conventional HB processes rely on an upstream operation to produce hydrogen (H_2) from fossil fuels, which is then mixed with dinitrogen to form the *synthesis gas* used for ammonia production. There are essentially two main process routes for producing hydrogen from fossil fuels, with the best available technology being steam methane reforming (SMR), which relies on a natural gas feedstock. The SMR section is responsible for much of the greenhouse gas (GHG) emissions associated with

ammonia production; in fact, for the SMR process route, it is estimated that 1.5 kilograms of CO_2-equivalent are emitted for every kilogram of ammonia produced [5]. As a result, about 1% of the global anthropogenic CO_2 emissions are attributed to the conventional HB process [5]. Shifting the hydrogen production route to electrolytically produced hydrogen greatly reduces the process dependence on fossil fuels, giving rise to the so-called green HB (gHB) process [1]. This shift entails a fundamental change in the nature of the HB process itself, which has traditionally been operated at a quasi-steady-state. The gHB process can therefore be viewed as a Power-to-Ammonia (P2A) process, which is a specific subclass of the broader Power-to-X (P2X) ones, wherein renewable energy is used to produce value-added commodity chemicals [1].

In a P2A process, green hydrogen is produced via a water electrolyser at a flow rate that is entirely dependent on the power supply to the electrolyser (*i.e.*, the availability of renewable energy). After being mixed with a nitrogen stream from an air separation unit (ASU), this stream is fed into the NH_3 synthesis loop [1]. The loop comprises an ammonia converter, a separator, and a recycle compressor, arranged as shown in **Figure 1**. Due to the low conversion-per-pass in the converter, unrecycled reactants must be recycled back to the converter, giving rise to the generic synthesis loop configuration. Different loop configurations exist depending on the quality of the make-up synthesis gas. For this reason, the gHB and the conventional HB process are expected to employ different loop configurations. In fact, even among commercial conventional HB process, several loop configurations are in use.

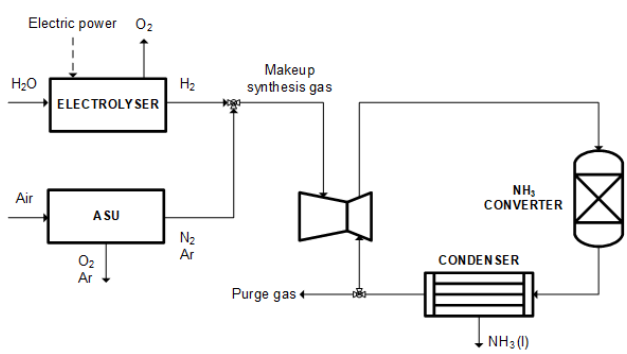

Figure 1. Generic Process Flow Diagram for the Green Haber-Bosch Process. Note that other loop configurations exist.

Under variable hydrogen supply (and hence variable make-up synthesis gas), the loop operates dynamically, requiring its unit operations to have a high degree of flexibility to accommodate different feed loads. This flexibility is currently missing in the conventional HB process.

In the context of Power-to-Ammonia systems, quantifying process flexibility through dynamic testing is certainly a most relevant task, because it enables the identification of current process boundaries, thereby shedding light on existing operational limitations. In other words, dynamic analysis helps clarify the conditions under which the process can resume operation without issues following a disturbance and ultimately provides guidance for improving process flexibility.

Existing works on the flexibility of synthesis loops have restricted their analysis to the converter [2, 6] under steady-state conditions or considered the loop, but again at steady-state [7]. Cheema and Krewer [7] quantified the operational flexibility of a three-bed converter inserted into a given synthesis loop, finding that, among all process variables, the inert gas fraction and the molar H_2:N_2 ratio were those whose variation enabled greater operational flexibility (*i.e.*, yielded more flexible designs). In more recent years, Fahr *et al.* [2] modelled Topsoe's proprietary S-300 converter and performed a steady-state parametric analysis to understand how that converter concept could be made more flexible. Across these studies, the central idea is that, for a given steady-state, the operating window of the converter can be widened by appropriately selecting the values of the operating variables. In all these works, however, the methodology did not involve dynamic tests. Conversely, Burrows and Bollas [8] employed dynamic testing but did not quantify flexibility, focusing instead on identifying faulty operating conditions. They investigated a conventional synthesis loop under dynamic feed conditions and conducted several dynamic tests at three different productivity scales. These tests subjected the process to a range of operating conditions, allowing process faults under such conditions to be identified [8]. Furthermore, Gottheil and Bremer [9] briefly addressed a systematic methodology for quantifying the flexibility of ammonia converter concepts by categorizing a given disturbance. However, their analysis was essentially limited to ramp-type disturbances, which could be characterised by a "ramp height" and a "ramp length". Moreover, this methodology was only applied to the ammonia converter.

Increasing the flexibility of the ammonia loop is enabled by increasing the flexibility of the individual unit operations. Of all existing unit operations within the loop, the ammonia converter is the most critical. The converter roughly consists of multiple catalyst beds and an iFEHE, all enclosed within a single pressurized vessel, with the iFEHE positioned downstream of the catalyst beds. There are three well-established types of converters on the market, depending on the cooling principle used to control the heat liberated by the ammonia synthesis reaction: the *Adiabatic Quench Cooled Reactor* (AQCR), the *Adiabatic Indirectly Cooled Reactor* (AICR), and the *Internally Cooled Multi-Tubular Reactor* (ICMR). Current industrial trends favour the AICR type, as all feed passes through

the entire catalyst bed (*i.e.*, no fraction of feed bypasses a given bed), allowing higher reactant conversion than the AQCR, even though the latter has found widespread acceptance in the industry for its simple mechanical design [10, 11].

When operating it in steady-state, the converter exhibits multiplicity of steady-state behavior when proceeding towards a steady-state after startup, owing to external recycle of heat [10, 12]. This extreme sensitivity ultimately implies that it can easily extinguish when transitioning between operating points. Consequently, its autothermal behavior is lost and productivity decreases sharply. When operating it dynamically, *i.e.*, when subjected to disturbances over a given time horizon, the converter exhibits several characteristic behaviors that have been extensively investigated in earlier studies [10, 13]. The first is the inverse response dynamics due to the packed bed. Catalyst overheating may be observed in consequence, which may lead to its deactivation, which is highly unwanted. Additionally, this response, combined with the feedback that the iFEHE introduces in the system, may evolve into (self) sustained oscillatory behavior (*i.e.*, limit cycles) [10, 13]. Such oscillations are also undesirable, as they induce pressure and temperature cycles, leading to material fatigue [1]. All these characteristics make the converter extremely sensitive to changes in process variables, hence particularly vulnerable for continuous dynamic operation.

If the converter design is sufficiently flexible, these operational issues may be avoided because the operating windows are wide enough that such disturbances do not push the process beyond its boundaries, allowing operation to be safely resumed. Therefore, assessing how converter designs are developed, quantifying their actual flexibility, and understanding how they can be made more flexible are critical questions in this context.

This work focuses on the initial steps in the development of a comprehensive framework for designing flexible synthesis loops. The authors focus on one specific component of this framework: systematic dynamic testing for the determination of process boundaries. Through dynamic tests, operational issues are intentionally induced in the system to identify its operating limits. For the purpose of systematic dynamic testing, the methodology developed by Gottheil and Brewer [9] has been adopted in this work and applied to a synthesis loop equipped with a commercially available AQCR-type converter (Topsoe's proprietary S-100 converter). Applying this methodology to the synthesis loop has not been explored in the literature and is therefore the main contribution of this study. The objectives of this study are twofold: (i) to assess the performance of this converter concept for Power-to-Ammonia applications, and (ii) to quantify the resulting flexibility of the synthesis loop when equipped with such a converter.

MODEL DEVELOPMENT

Reactor models belonging to the Advanced Model Library – Fixed Bed Catalytic Reactors (AML – FBCR) in gPROMS Process have been *composed* to assemble models of ammonia converters [14]. The model is a first-principles formulation, expressed as a system of differential-algebraic equations. In this model, variables are distributed over two spatial domains, namely the radial and axial directions, and separate conservation laws are written for the fluid and solid phases (heterogenous model). The catalyst consists of Fe-based spherical pellets, which have been modelled as a lumped system, considering the selected particle granulometry (industrial-grade particles typically range between 1.5 - 3 mm for radial-flow converters). The particle sizes employed guarantee there are no appreciable intraparticle gradients, which is to say there are practically no internal mass and heat transfer limitations. This methodological option results in a model that does not increase in dimensionality. Modified Temkin-Pyzhev kinetics by Dyson and Simon [15] have been selected, as follows in eq. (1):

$$r_{NH_3} = k\left[K_a^2\, a_{N_2}\left(\frac{a_{H_2}^3}{a_{NH_3}^2}\right)^{\alpha} - \left(\frac{a_{NH_3}^2}{a_{H_2}^3}\right)^{1-\alpha}\right] \quad (1)$$

where r is the rate of formation of ammonia (mol kgcat^{-1} s^{-1}), k is the kinetic constant (mol kgcat^{-1} s^{-1}) calculated using an Arrhenius-type equation with data provided in **Table 1**, K_a is the equilibrium constant computed using Gillespie and Beattie's equation [16], a_i the gas-phase activity of component i and α an empirical parameter which was set to 0.5.

Table 1: Kinetic data employed in the developed model.

Pre-exponential factor [kmol/kgcat h]	4.42×10^{11}
Activation energy [kcal/kmol]	4.08×10^{4}

All heat exchangers were modeled as black box units that satisfied mass and energy balances and a given thermal specification (e.g., temperature difference between outlet streams). As for the ammonia separator, a dynamic model for a flash separator has been used, using the gPROMS Separations Model Library. No compressors have been used to model this synthesis loop, since the pressure drop across the catalyst beds is negligible and the make-up stream is assumed to be available at the pressure of the synthesis loop. Simplifications in the compressors and heat exchangers design may influence the overall loop transient response. However, to keep the loop model simple yet functional and capture the physics of the system, these simplifying assumptions were adopted. Due to the packed-bed nature of the converter, the loop's response will be primarily determined by the converter's response, whose dynamics are much slower

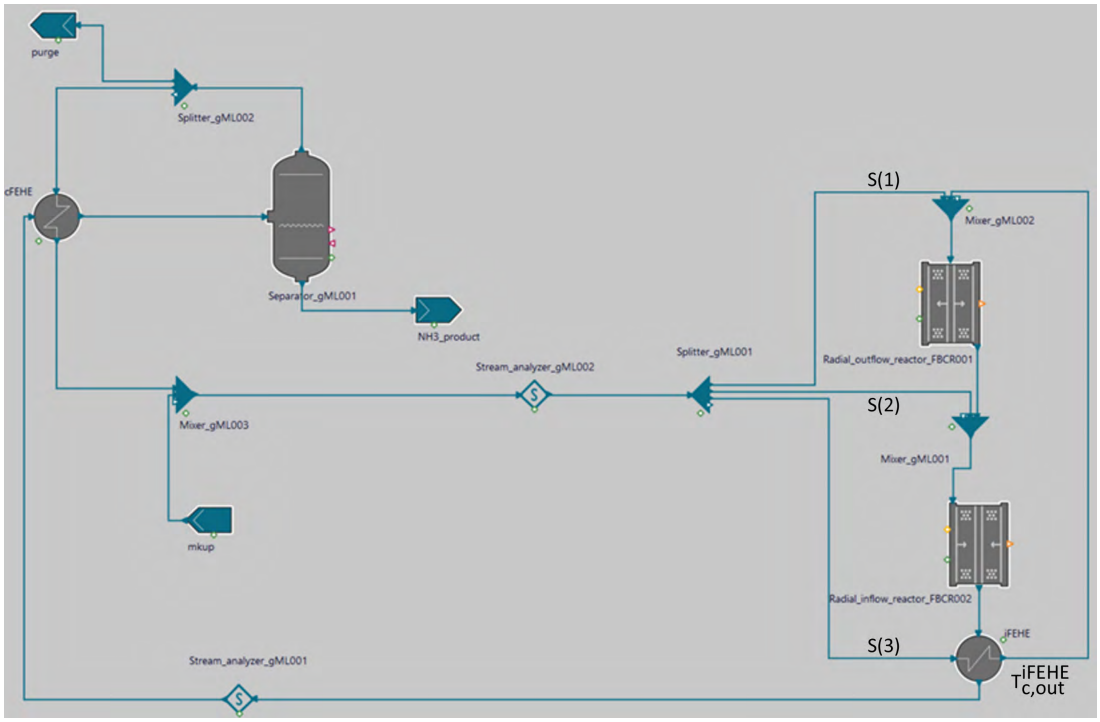

Figure 2. Flowsheet for the green ammonia loop modeled in gPROMS Process. In the loop, it is installed a 2-bed AQCR converter, reproducing Topsoe's S-100 proprietary converter. Selected key operating variables of the converter are shown in the figure.

than those of the compressors and heat exchangers.

These models were assembled to build a composite model of a synthesis loop, as illustrated in **Figure 2**. Overall, the mathematical model of the synthesis loop comprises approximately 14300 variables, and across all runs, simulation times reached a maximum of 7 seconds. Catalyst beds were discretized axially and radially, respectively, into 10 and 5 points.

METHODOLOGY

In this work, the authors detail a preliminary methodology for assessing the operational flexibility of the synthesis loop through systematic dynamic testing. Here, flexibility is evaluated by mapping the loop's operating window through a series of dynamic tests.

At this stage, the problem is formulated as a partial design problem, as it considers only performance. Economic considerations are not included in the present analysis.

Loop design and optimization

The loop was designed for a nominal ammonia production rate of 4 kt/year (approximating a small-scale ammonia production plant). In the selected loop configuration, the make-up synthesis gas is introduced downstream of the separator, since the existing trace amounts of water in the make-up gas do not pose a significant threat to the catalyst to the point of deactivating it. Iron-based catalysts, such as those employed in ammonia synthesis, are extremely sensitive to any oxygen-containing compounds, which can result in catalyst poisoning. Heuristic criteria were used to specify the separator operating temperature (-5 ºC) and purge rate (6 %). As the pressure drop across the catalyst beds was found to be negligible, a constant 200 bar pressure was assumed for all process units, and no recycle compressor was considered at this stage.

Starting from the flowsheet presented in **Figure 2**, the loop design was then optimized to achieve maximum conversion per pass. The design problem is formulated as a non-linear optimization problem, solved using gPROMS Optimization Entity, which is written as follows:

$$\max_D X_{N_2} \tag{2}$$

$$D = \left[T_{c,out}^{iFEHE}, S(1), S(2)\right] \tag{3}$$

s.t.

$$f(D, Y) = 0 \tag{4}$$

$$T_{bed\,i} \leq T_{upper} \tag{5}$$

$$T_{lower} \leq T_{bed\,i} \tag{6}$$

where D is the set of decision variables (*i.e.*, degrees of freedom), which, for an AQCR-type converter, are the quench fractions, $S(1)$ and $S(2)$, and the outlet

temperature of the feed stream (*i.e.*, the cold stream) of the feed-effluent heat exchanger, $T_{c,\,out}^{iFEHE}$.

$f(D, Y)$ represents the set of equations that define the model and must be satisfied at all times. Y denotes the remaining (dependent) model variables. Bed temperatures were constrained between 600 and 773 K, as literature reports 800 K to be the upper limit for catalyst operating temperature without risking deactivation [2, 17], while temperatures above 600 K are reported to yield appreciable reaction rates for the employed catalysts [18]. The results of this optimization study are presented in **Table 2** and reveal the optimal values for several key variables. Note that this represents an initial formulation, which at present considers only the reaction performance aspect of the problem. Process economic considerations and the extension of other loop variables to the design problem will be incorporated in future work.

Table 2: Results for the converter design before and after the optimization study.

Variable	Base design	Optimized design
Conversion [%]	22.52	23.61
$T_{c,\,out}^{iFEHE}$ [K]	690	669
Split to cold shot to bed 1, $S(1)$ [-]	0.200	0.229
Split to quench in between bed 1 and bed 2, $S(2)$ [-]	0.350	0.422

Dynamic tests planning and quantification of operational flexibility

As noted in the **Introduction**, disturbances applied to the converter at a given operating point may trigger one of the following phenomena: catalyst overheating (with subsequent deactivation), reaction extinction, or oscillatory behavior. The occurrence of any of these phenomena indicates that a process boundary has been crossed, rendering normal operation of the Power-to-Ammonia process infeasible. These phenomena, therefore, define the criteria for feasible operation of the synthesis loop.

Any dynamic tests performed on the modeled synthesis loop can thus be used to deliberately induce these phenomena and, whenever any of them arises, a given range of operating conditions is deemed unfeasible. From this information, the process boundaries can then be determined, and an operational map can be constructed from which the current loop flexibility can be inferred.

For the modeled loop, all dynamic tests were performed on the make-up gas feed flow rate. According to the methodology proposed in this work, all dynamic tests are characterized by (i) the *type* of disturbance (ramp, step, pulse, or sinusoidal); (ii) the *magnitude* or *amplitude* of disturbance, and (iii) the *duration* of disturbance. Specifically, the tests in this study consisted exclusively of linear ramps, applied with four different amplitudes (expressed as percentage deviations from the nominal value) and three disturbance durations. This resulted in a 4 × 3 test matrix, as summarized in **Table 3**.

Table 3: List of dynamic tests performed.

Type of disturbance	Linear ramp
Disturbed variable	Make-up gas flow rate
Magnitude of disturbances [%]	± 25 & ± 50
Duration of disturbances [s]	10, 600, 1800 & 3600

Therefore, the operational flexibility is defined by the maximum disturbance amplitude the system can withstand for a fixed disturbance time.

RESULTS AND DISCUSSION

In this section, the results for the dynamic tests conducted are presented, first focusing on the impact of these tests at the converter level and then expanding the study to the loop.

Table 4: Final steady-state values of relevant high-level converter variables for all dynamic tests considered.

Disturbance		X [%]	Θ_{conv} [s]	T_{out}^{conv} [K]
-50%	10 s	24.56	63.06	787.4
	600 s	24.56	63.09	787.5
	1800 s	24.59	63.19	787.3
	3600 s	24.61	63.22	787.2
-25%	10 s	24.54	42.52	787.0
	600 s	24.56	42.55	786.8
	1800 s	24.57	42.56	786.8
	3600 s	24.57	42.56	786.8
+25%	10 s	21.73	23.97	764.0
	600 s	21.71	23.96	763.8
	1800 s	21.70	23.95	763.6
	3600 s	21.69	23.94	763.5
+50%	10 s	4.580	8.33	624.2
	600 s	4.580	8.33	624.2
	1800 s	4.580	8.33	624.2
	3600 s	4.580	8.33	624.2

Converter performance

For all tests considered, the final steady-state values of the relevant high-level variables for the converter are summarized in **Table 4**. These include the conversion per pass, X, the overall average residence time, θ_{conv}, and the converter outlet temperature, T_{out}^{conv}.

From **Table 4**, it can be observed that decreasing

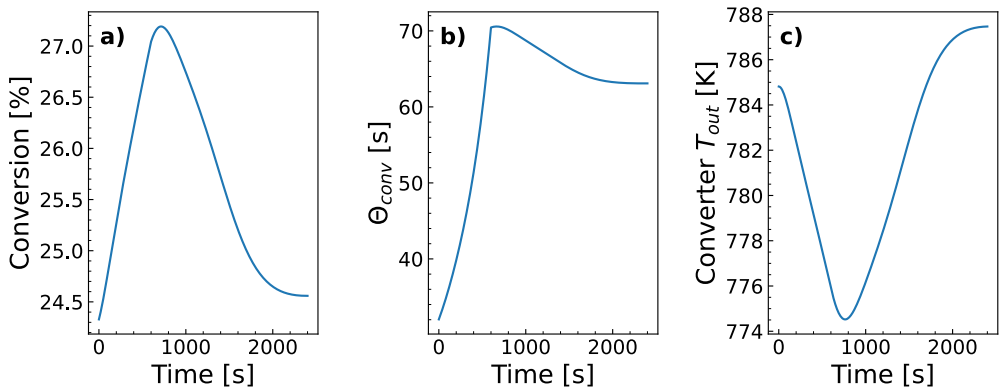

Figure 3: Profiles of percent conversion, overall average residence time, θ_{conv}, and outlet gas temperature, T_{out}, following a ramp disturbance of -50% on the make-up gas mass flow rate for 600 s.

feed flow rates lead to higher conversions (which are, in any case, limited by the equilibrium conversion). Conversely, admitting mass flow rates above the nominal to the converter reduces conversion and may even extinguish the converter if the superficial velocity becomes too high, substantially reducing the gas residence time.Furthermore, as expected, conversion and average residence time exhibit similar trends. That is, conversion is highest when the average residence time is greatest. Additionally, in cases where the conversion increased (*i.e.*, when the make-up flow rate decreased), the final converter outlet temperature (after the new steady-state was reached) was also higher, due to the enhanced occurrence of reaction. The opposite occurred when the conversion decreased. However, the transition between the two steady-state (initial and final) values of the converter outlet temperature is not as linear as it may seem. For that reason, the authors now restrict their analysis to the case in which the make-up gas flow was reduced by 50% over the course of 600 s. Along with the time variation of the make-up mass flow rate (and, by extension, the converter inlet mass flow rate), the profiles of the overall conversion, overall average residence time and converter outlet temperature are shown in **Figure 3**. The conversion profile exhibits a bell-shaped trend over time. As the inlet feed flow rate decreases, the gas superficial velocity reduces, and the gas spends more time in contact with the catalyst. Consequently, the increased residence time translates into higher conversion (see **Figure 3b**). Despite the increased conversion, which locally increases heat generation by the reaction, the solid bed temperature does not respond instantaneously due to its high thermal inertia, a consequence of to the densely packed catalyst bed. Additionally, the reduced gas mass flow rate the gas's weakens radial convective heat transport. As a result, the gas cannot absorb as much heat from the bed as before, so less energy is transported to the converter outlet. Consequently, the gas reaches the outlet at a lower temperature, explaining the decreasing trend observed in **Figure 3c**. Eventually, as the bed becomes progressively hotter by the enhanced occurrence of the reaction, part of this heat propagates by conduction, with the remainder carried by convection in the flowing gas. The gas temperature then rises, and when this temperature front reaches the converter outlet, the gas is seen as being hotter. This results in the outlet converter temperature profile shown in **Figure 3c**, where the outlet temperature initially decreases, goes through a minimum, and then increases again. Therefore, this sequence of a fast hydrodynamic effect followed by a slower thermal effect leads to the observed inverse dynamic response, characteristic of packed-bed reactors. However, as time goes by, the increase in the converter outlet temperature shifts the equilibrium conversion to lower values. Therefore, the maximum attainable conversion decreases, as does the actual conversion, giving rise to the decline observed in **Figure 3a**.

Loop performance

Table 5 presents the loop's key performance indicators. In this work, the following were considered: process yield and ammonia (loop) production rate.

When the make-up flow rate was increased, the loop production rate increased because more material circulated in the loop. Conversely, when the flow rate to the loop was decreased, the loop's productivity dropped. However, at the converter level, conversion per pass increased, which explains the rise in loop yield.

As more material circulated within the loop, the loop yield slightly increased reflecting the greater conversion per pass when the make-up flow rate is decreased. Moreover, for the disturbance with the largest absolute magnitude (± 50%), two different operational problems were encountered: increasing the make-up flow rate led to virtually no conversion (reaction extinction), whereas decreasing it caused only slight bed overheating on the

order of 5 K while still allowing production to continue (albeit at lower rates). Thus, the sign of the disturbance magnitude (*i.e.*, the direction of change) explained why the process was affected so differently in these cases.

Table 5: Key performance indicators for assessing loop performance.

Disturbance		X [%]	Loop yield [%]	Loop production rate [kt/y]
-50%	10 s	24.56	98.16	2.12
	600 s	24.56	98.15	2.16
	1800 s	24.59	98.19	2.16
	3600 s	24.61	98.19	2.16
-25%	10 s	24.54	98.15	3.18
	600 s	24.56	98.16	3.24
	1800 s	24.57	98.17	3.24
	3600 s	24.57	98.17	3.24
+25%	10 s	21.73	95.80	5.17
	600 s	21.71	95.78	5.27
	1800 s	21.70	95.77	5.27
	3600 s	21.69	95.76	5.27
+50%	10 s	4.580	55.39	3.59
	600 s	4.580	55.39	3.63
	1800 s	4.580	55.39	3.63
	3600 s	4.580	55.39	3.63
Nominal		82.46	97.98	4.32

Loop operability

So far, the dynamic behavior of both the converter and the loop has been analyzed in terms of relevant process indicators. This analysis revealed how the system responded whenever subjected to ramp-type disturbances. In addition, it allowed the feasibility of the studied loop configuration to be assessed when operating within dynamic conditions. Recall that for each dynamic test performed, the following conditions were investigated: (i) reaction extinction, (ii) oscillatory behavior in any of the output variables, and (iii) catalyst bed overheating. The occurrence of any of the previous rendered the corresponding transition unfeasible, but also the operation at this range of conditions.

With the +50% disturbance, reaction extinction was observed, resulting in a converter *cold stop*. In contrast, for the -50% disturbance, a bed overheating of about 5 K was detected. In the long term, bed overheating leads to catalyst deactivation and entails additional operating expenditures to replace the catalyst bed that has experienced a reduced service lifetime. Therefore, these cases were classified as unfeasible operating ranges.

This information is summarized in **Table 6** and forms the basis for an operational map. From the results in Table 6, it can be concluded that the modeled loop configuration is ± 25% flexible over the studied disturbance duration.

Table 6: Operational map of the dynamic tests conducted by varying the make-up gas mass flow rate.

Amplitude / Duration	-25%	-50%	+25%	+50%
10 s	OK[1]	CatH[2]	OK[1]	RxnEx[3]
600 s	OK[1]	CatH[2]	OK[1]	RxnEx[3]
1800 s	OK[1]	CatH[2]	OK[1]	RxnEx[3]
3600 s	OK[1]	CatH[2]	OK[1]	RxnEx[3]

[1] – denotes a final steady-state in which no process constraints are violated; [2] – denotes a final steady-state in which catalyst overheating is observed; [3] – denotes a final steady-state in which reaction extinction is identified.

CONCLUSION AND FUTURE WORK

This work represents an initial application of a systematic methodology to assess the flexibility of an ammonia synthesis loop equipped with a quench-cooled converter (AQCR) for Power-to-Ammonia applications. The methodology builds upon that previously proposed by Gottheil and Bremer [9]. The presented methodology should be regarded as part of a higher-level framework for designing flexible synthesis loops.

Dynamic tests were conducted by varying the make-up flow rate across four disturbance amplitudes and four disturbance durations. The disturbance space in this study considered only ramp disturbances, which effectively become step disturbances for short or very short disturbance times. Future work will address multivariable simultaneous disturbances and explicitly consider flow inputs from real renewable intermittency profiles, thus expanding the assessment beyond the present scope.

By compiling the results of the conducted transient simulations, it was possible to identify the effects of feed flow rate variations on overall loop performance (including converter performance) and, most importantly, to determine whether any operational limits were violated. If so, a process boundary is crossed, and the corresponding range of operating conditions must be considered infeasible. The present results therefore enable the construction of a preliminary operational map from which the current loop flexibility can be quantitively inferred.

Advancing such methodological solutions is essential for overcoming the limited flexibility of conventional ammonia synthesis loops, which remain one of the main bottlenecks to the industrial deployment of Power-to-Ammonia systems, alongside plant economics.

ACKNOWLEDGEMENTS

Project co-financed by the PRR – Recovery and

Resilience Plan by the European Union, in particular, the Mobilizing Agenda "Moving2Neutrality" (Project no. 32, with reference no. C644927397-00000038). The authors gratefully acknowledge the support of CQE (UID/00100/2025 and UID/PRR/100/2025) and IMS the Associated Laboratory funded by Project LA/P/0056/2020 (DOI 10.54499/LA/P/0056/2020), and of CERENA (FCT-UID/04028/2025).

AUTHOR IDENTIFIERS

Author ORCIDs:
Pires, José M.: 0009-0008-4598-1574
Narciso, Diogo A. C.: 0000-0002-2776-7477
Pinheiro, Carla I. C.: 0000-0002-2568-3974

REFERENCES

1. Narciso DAC, Pires JM, Fortunato J, Teixeira P, Castro PM, Pinheiro CIC, Matos HA. Design and operation of power-to-ammonia systems: a review. Energy Conversion and Management 327:119494 (2025). https://doi.org/10.1016/j.enconman.2025.119494
2. Fahr S, Schiedeck M, Schwarzhuber J, Rehfeldt S, Peschel A, Klein H. Design and thermodynamic analysis of a large-scale ammonia reactor for increased load flexibility. Chemical Engineering Journal 471:144612 (2023). https://doi.org/10.1016/j.cej.2023.144612
3. Lim J, Fernández CA, Lee SW, Hatzell MC. Ammonia and nitric acid demands for fertilizer use in 2050. ACS Energy Lett. 6:3676-3685 (2021). https://doi.org/10.1021/acsenergylett.1c01614
4. IRENA and AEA, Innovation Outlook: Renewable Ammonia (Report) (2022).
5. Smith C, Hill AK, Torrente-Murciano L. Current and future role of haber–bosch ammonia in a carbon-free energy landscape. Energy Environ. Sci. 13:331-344 (2020). https://doi.org/10.1039/c9ee02873k
6. Fahr S, Schiedeck M, Reinke M, Bohn JP, Rehfeldt S, Peschel A, Klein H. Simultaneous design and part-load optimization of an industrial ammonia synthesis reactor. Chemical Engineering Journal 480:148302 (2024). https://doi.org/10.1016/j.cej.2023.148302
7. Cheema II, Krewer U. Operating envelope of haber–bosch process design for power-to-ammonia. RSC Adv. 8:34926-34936 (2018). https://doi.org/10.1039/c8ra06821f
8. Burrows L, Bollas GM. Stability assessment of small-scale distributed ammonia production systems. Ind. Eng. Chem. Res. : (2022). https://doi.org/10.1021/acs.iecr.2c00631
9. Gottheil L, Bremer J. Shifting towards dynamic and load-flexible ammonia synthesis via polytropic fixed-bed reactors. Chemical Engineering Journal 526:170927 (2025). https://doi.org/10.1016/j.cej.2025.170927
10. Rosbo JW, Ritschel TKS, Hørsholt S, Huusom JK, Jørgensen JB. Flexible operation, optimisation and stabilising control of a quench cooled ammonia reactor for power-to-ammonia. Computers & Chemical Engineering 176:108316 (2023). https://doi.org/10.1016/j.compchemeng.2023.108316
11. Rosbo JW, Jensen AD, Jørgensen JB, Huusom JK. Comparison, operation and cooling design of three general reactor types for power-to-ammonia processes. Chemical Engineering Journal 496:153660 (2024). https://doi.org/10.1016/j.cej.2024.153660
12. Froment, GF, Bischoff, KB, and De Wilde, J Chemical Reactor Analysis and Design. John Wiley & Sons (2011).
13. Morud, J and Skogestad, S, Analysis of instability in an industrial ammonia reactor. *AIChE J.* 44:888-895 (1998).
14. Siemens. https://www.siemens.com/gproms-digital-process-design-and-operations.html
15. Dyson, D and Simon, J, Kinetic expression with diffusion correction for ammonia synthesis on industrial catalyst. *Ind. Eng. Chem. Fundamen.* 7:605-610 (1968).
16. Gillespie LJ, Beattie JA. The thermodynamic treatment of chemical equilibria in systems composed of real gases. I. an approximate equation for the mass action function applied to the existing data on the haber equilibrium. Phys. Rev. 36:743-753 (1930). https://doi.org/10.1103/physrev.36.743
17. Jennings JR and Ward SA. Ammonia Synthesis. In: Catalyst Handbook. Ed: Twigg MV. CRC Press (1989).
18. Appl M. Ammonia, 2. production processes. Ullmann's Encyclopedia of Industrial Chemistry : (2011). https://doi.org/10.1002/14356007.o02_o11

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Symbolic Regression-based approach for Modeling Fouling Resistance in Heat Exchangers

Fernando A. R. D. Lima[ab] Sergio G. de Oliveira[c], Antonioni B. Campos[c], Bruna Carla G. de Assis[c], Livia Pereira L. Costa[c], Fabio S. Liporace[c], Mauricio B. de Souza Jr.[ab] and Argimiro R. Secchi[ab]

[a] Chemical Engineering Program, PEQ/COPPE – Universidade Federal do Rio de Janeiro, Av. Horácio Macedo, 2030, CT, Bloco G, G115, 21941-914, Rio de Janeiro, RJ – Brazil
[b] EPQB, School of Chemistry, Universidade Federal do Rio de Janeiro, Av. Horácio Macedo, 2030, CT, Bloco E, 21941-914, Rio de Janeiro, RJ – Brazil
[c] CENPES, PETROBRAS, Rio de Janeiro, RJ, Brasil
* Corresponding Author: farrais@eq.ufrj.br.

ABSTRACT

Heat exchangers frequently suffer from fouling, which is the accumulation of unwanted deposits on heat-transfer surfaces. This issue reduces thermal performance, increases pressure drop, and raises energy use and operating costs. Predicting fouling resistance remains challenging in process engineering, yet it is important for monitoring, maintenance planning, and mitigation actions that reduce economic losses and environmental impacts. Symbolic regression (SR) is a machine learning approach that searches for an explicit mathematical expression that best represents the relationship between process inputs and a target output. Unlike many black-box models, SR can capture nonlinear behavior while producing compact, interpretable equations that are easier to deploy and analyze in industrial settings. In this work, a methodology to rapidly obtain algebraic models for fouling resistance in industrial heat exchangers using SR was proposed. Plant measurements of hot- and cold-side flow rates and inlet/outlet temperatures were transformed into dimensionless numbers, which were used as model inputs to predict fouling resistance. SR was applied to operational data from industrial heat exchangers, achieving coefficients of determination (R^2) above 0.95 and 0.87 for the training and validation datasets, respectively, with low prediction errors. A selected equation was then transferred to a second industrial heat exchanger by re-estimating a reduced set of parameters using plant data. Validation on data not used during parameter estimation showed that the adapted model maintained high predictive accuracy, with R^2 values of 0.96 and 0.91 for training and validation, respectively. Overall, the proposed approach provided an efficient and interpretable framework for fouling resistance modeling and supported practical deployment across similar units.

Keywords: Fouling resistance, Heat exchangers, Symbolic regression, Interpretable machine learning, Industrial process modeling

INTRODUCTION

Heat exchangers are key units in chemical processing, energy systems, and oil refining, where they enable efficient heat recovery and thermal integration. However, their performance is strongly affected by fouling, defined as the unwanted accumulation of solid or semi-solid material on heat transfer surfaces, which introduces an additional thermal resistance and progressively degrades heat transfer efficiency [1–3]. In refinery preheat trains, fouling leads to increased fuel consumption, higher greenhouse gas emissions, elevated pressure drops, and significant economic penalties due to throughput reduction and maintenance interventions [2, 4, 5]. Despite decades of research and industrial experience, fouling remains one of the most persistent and costly operational challenges in heat exchanger networks, with surveys indicating that more than 90% of industrial heat exchangers experience fouling during operation [1, 3, 6].

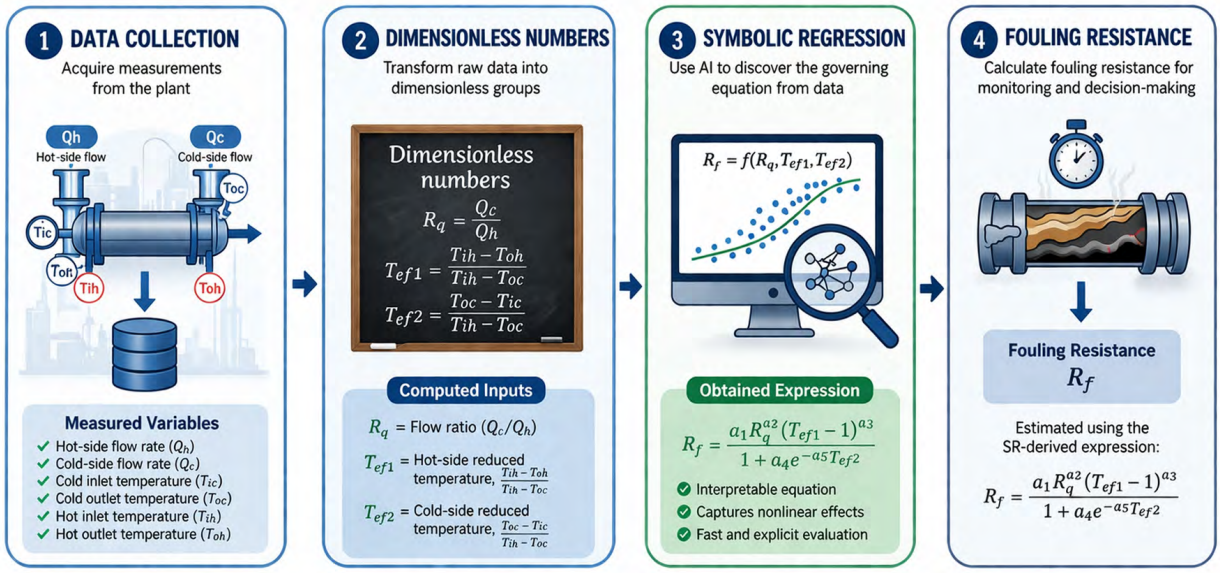

Figure 1: SR-based methodology proposed in this work to obtain expressions to calculate the fouling resistance.

The impact of fouling on heat exchanger performance is commonly quantified through the fouling resistance, which represents the additional thermal resistance introduced by the deposited layer and directly affects the overall heat transfer coefficient [1, 3, 7]. Fouling resistance is not constant over time and can exhibit markedly different temporal behaviors, including linear growth, asymptotic behavior, falling-rate trends, or oscillatory patterns, depending on the nature of the deposits, operating conditions, and dominant fouling mechanisms [1, 6, 8]. Multiple mechanisms may act simultaneously, including crystallization fouling, chemical reaction fouling, particulate deposition, corrosion, and ageing of deposits, resulting in highly complex and system-specific fouling dynamics [1, 5, 6]. This complexity makes fouling resistance difficult to predict reliably, particularly under industrial conditions where temperature, flow velocity, fluid composition, and shear stress vary along the length of the exchanger [4, 7, 9].

To model fouling resistance, a wide range of approaches has been proposed in the literature, spanning from empirical correlations to semi-empirical and physically motivated models [5, 7, 9]. Early models introduced the concept of competing deposition and removal mechanisms, expressing the net fouling rate as the difference between deposition and suppression terms [5, 7]. Building on this idea, threshold fouling models were developed to identify operating conditions under which deposition is balanced by removal, leading to negligible fouling rates [4, 5, 9]. These models have been widely adopted due to their relative simplicity and their ability to capture key trends observed in laboratory experiments. However, their predictive accuracy often deteriorates when applied to industrial systems, especially refineries, as the fitted parameters are strongly dependent on crude composition, exchanger geometry, and operating regime [5, 7, 10].

More advanced modeling approaches include deterministic and distributed models that explicitly account for local transport phenomena, reaction kinetics, and the dynamic interaction between fouling growth and heat transfer [2, 7, 9]. Such models have shown improved agreement with industrial data and can capture spatial variations of fouling resistance along the exchanger length [2, 7]. Computational fluid dynamics (CFD) has further enhanced mechanistic understanding by resolving local temperature, velocity, and concentration fields, offering insight into early stage fouling and induction periods [8, 9]. Nevertheless, these high-fidelity models are computationally intensive, require extensive parameterization, and are often impractical for routine monitoring, optimization, or real-time decision support in industrial environments [7–9].

As a result, fouling resistance modeling faces a fundamental trade-off between physical fidelity, computational tractability, and practical applicability. Empirical and semi-empirical models are easy to deploy but lack robustness and transferability, while detailed mechanistic models offer insight at the expense of complexity and computational cost [5, 7, 9, 10]. Moreover, many existing approaches rely on fouling resistance values inferred indirectly from heat balance calculations, embedding additional uncertainty associated with measurement noise and modeling assumptions [7, 9].

Machine learning (ML) has emerged as a powerful modeling paradigm in chemical and process engineering, particularly for systems where first-principles models are difficult to derive, computationally expensive, or affected by structural uncertainty. ML-based surrogate and data-driven models can learn complex nonlinear relationships directly from process data, enabling fast prediction, efficient optimization, and real-time decision support without explicit mechanistic formulations [11–13]. Compared to traditional empirical correlations, machine learning models can integrate large numbers of input variables, capture high-order interactions, and remain effective under noisy or partially observed conditions, which are common in industrial environments [11, 12]. As a result, ML approaches have been successfully applied to process modeling, optimization, monitoring, and fault detection across a wide range of chemical engineering applications, demonstrating significant gains in accuracy and computational efficiency over conventional modeling strategies [11–13].

In the context of heat exchangers, machine learning has shown promise for modeling fouling resistance. Artificial neural networks (ANNs) and deep learning models have been employed to predict fouling resistance directly from measurable process variables such as temperatures, flow rates, and pressure drops, achieving high predictive accuracy even under noisy conditions [14, 15]. These data-driven approaches bypass the need for detailed fouling mechanism identification and can be trained using either experimental data or high-fidelity simulation data, making them attractive for industrial deployment [14, 15]. Moreover, ML-based fouling models have demonstrated scalability across operating regimes and exchanger configurations, supporting their use in monitoring, maintenance planning, and performance optimization of heat exchanger systems [12, 14, 15]. However, many existing ML approaches rely on complex black-box models, which can limit interpretability and hinder trust and adoption in industrial decision-making.

Building on the broader adoption of ML in chemical engineering, symbolic regression (SR) has gained attention as a way to retain the predictive power of data-driven models while producing explicit, readable equations rather than opaque black-box predictors [16, 17]. SR searches over a space of mathematical expressions (e.g., combinations of process variables, constants, and operators) to identify compact formulas that balance accuracy and complexity, often using evolutionary or program-search strategies to evolve candidate models toward parsimonious structures [16, 17]. This is particularly attractive in process systems engineering because the resulting expressions are straightforward to audit, differentiate, and embed in simulation, optimization, and control workflows, and they can serve as interpretable surrogates or as corrective closures to improve mechanistic models [18]. Recent chemical-engineering studies highlight SR as a practical bridge between data and physics: it can recover compact relationships consistent with underlying transport/kinetic behavior and can be used to distill complex learned representations into deployable algebraic forms without sacrificing much accuracy [19, 20].

This work aims to address the limitations of existing fouling resistance modeling approaches by developing an accurate, interpretable, and transferable data-driven model suitable for industrial applications. Building on the strengths of ML while explicitly mitigating the drawbacks associated with black-box models, a SR–based framework was proposed to identify compact algebraic expressions that relate fouling resistance to operating conditions expressed through physically meaningful variables and dimensionless groups. The specific contributions of this study are: (i) the formulation of a symbolic regression methodology tailored to fouling resistance modeling using industrial heat exchanger data; (ii) the demonstration that the resulting symbolic models achieve high predictive accuracy while remaining transparent and easy to deploy; and (iii) the validation of the model's transferability by adapting it to a second industrial heat exchanger with minimal parameter re-estimation. By combining interpretability, computational efficiency, and practical relevance, the proposed approach provides a robust alternative for fouling monitoring, analysis, and decision support in heat exchanger operation.

METHODOLOGY

Data acquisition and preprocessing

The dataset used in this work was obtained from two industrial shell-and-tube heat exchangers operating under normal refinery conditions. For each unit, routinely measured operating variables were collected over extended periods of operation. The available measurements include the hot-side fluid flow rate (Q_h), cold-side fluid flow rate (Q_c), cold-fluid inlet temperature (T_{ic}), hot-fluid inlet temperature (T_{ih}), cold-fluid outlet temperature (T_{oc}), and hot-fluid outlet temperature (T_{oh}). These variables are typically available in industrial control systems and were selected to ensure that the proposed modeling framework relies exclusively on standard plant measurements.

Prior to model development, the datasets were screened to remove missing values and outliers associated with sensor faults or transient shutdown conditions. The resulting datasets were then divided into training and validation subsets, ensuring that both heat exchangers were represented in each subset. This strategy was adopted to assess not only predictive accuracy, but also the transferability of the identified models across different units.

Dimensionless variable formulation

To reduce scale effects and improve generalization, the measured variables were transformed into a set of dimensionless groups derived from energy balance considerations. In particular, the following dimensionless variables were defined:

$$R_q = \frac{Q_c}{Q_h} \quad (1)$$

$$T_{ef1} = \frac{T_{ih} - T_{oh}}{T_{ih} - T_{oc}} \quad (2)$$

$$T_{ef2} = \frac{T_{oc} - T_{ic}}{T_{ih} - T_{oc}} \quad (3)$$

These quantities capture the relative flow distribution and thermal effectiveness of the heat exchanger in a compact, normalized form. The use of dimensionless variables also facilitated comparison between exchangers of different sizes, operating ranges, and absolute flow rates. The fouling resistance, R_f, was treated as the target variable and was calculated from the model proposed by Liporace and Oliveira [21].

Symbolic regression framework

SR was used to discover an explicit algebraic relationship between fouling resistance R_f and the dimensionless variables (R_q, T_{ef1}, T_{ef2}). The workflow implemented in this study is summarized in Fig. 1, which outlines: (i) acquisition of plant measurements, (ii) computation of dimensionless groups, (iii) SR search for an analytical expression, and (iv) evaluation of R_f using the discovered expression. Symbolic regression was performed using the PySR library [22], an open-source framework for equation discovery based on a multi-population evolutionary search with an evolve–simplify–optimize loop for improving candidate expressions and refining numerical constants.

The PySR search was configured with 400 iterations and a population size of 30 candidate expressions. To promote interpretability and numerical stability, the operator set was restricted to common elementary functions. The allowed unary operators were exponential, logarithm, and square. The allowed binary operators were addition, subtraction, multiplication, and division. Candidate equations were evaluated by predictive performance and expression complexity, and the final model was selected to balance accuracy with parsimony, consistent with the goal of obtaining a compact and deployable equation for R_f.

RESULTS AND DISCUSSION

Application of Symbolic Regression

SR was first applied using data from the first industrial heat exchanger to identify an explicit algebraic relationship between the fouling resistance and the dimensionless variables defined in Methodology section. The resulting expression, denoted as Eq. 4, was selected based on the trade-off between predictive accuracy and model complexity. Table 1 summarizes the performance metrics obtained for this equation, showing coefficients of determination R^2 of 0.97 and 0.91 for the training and validation datasets, respectively, together with low mean squared error (MSE) values. These values indicate that the SR model can accurately capture the nonlinear relationship between operating conditions and fouling resistance for the first heat exchanger.

$$R_f = -0.192T_{ef1}R_q - 0.806T_{ef1} + \frac{0.806T_{ef1}}{T_{ef2}} + 0.806T_{ef2} - 0.943 \quad (4)$$

Table 1: Coefficient of determination (R^2) and mean squared error (MSE) of Eq. 4 to predict the training and validation sets for the first heat exchanger.

Dataset	R^2	MSE
Training	0.972	$2.48x10^{-8}$
Validation	0.910	$1.43x10^{-7}$

Figure 2 compares the measured fouling resistance values with the predictions obtained using Eq. 4. A good agreement between predicted and measured values was observed across the full range of fouling resistance, with no evident systematic bias. The dispersion around the parity line remains limited for both training and validation datasets, indicating that the identified symbolic expression generalizes well and does not suffer from overfitting despite its relatively compact analytical form. This result highlights one of the main advantages of SR, namely its ability to produce parsimonious and interpretable models while maintaining high predictive accuracy.

A second SR model was then identified using data exclusively from the second industrial heat exchanger, resulting in Eq. 5. The predictive performance of this model is reported in Table 2. The obtained R^2 values of 0.96 for the training dataset and 0.87 for the validation dataset confirm that SR is also capable of accurately modeling fouling resistance for a different unit operating under distinct conditions. Figure 3 further illustrates the agreement between measured and predicted fouling resistance values for this second heat exchanger, again showing a close alignment along the parity line and limited scatter.

$$R_f = -1.23R_q - 3.17T_{ef1} + \frac{1.65}{T_{ef2}} \quad (5)$$

The performance of Eqs. 4 and 5 shows that SR can successfully extract meaningful functional relationships between dimensionless operating variables and fouling resistance directly from industrial data. The identified

models achieved high predictive accuracy without relying on heat exchanger-specific mechanistic assumptions, supporting the applicability of the proposed approach to real industrial systems.

Table 2: Coefficient of determination (R^2) and mean squared error (MSE) of Eq. 5 to predict the training and validation sets for the second heat exchanger.

Dataset	R^2	MSE
Training	0.958	$3.56x10^{-8}$
Validation	0.871	$1.58x10^{-8}$

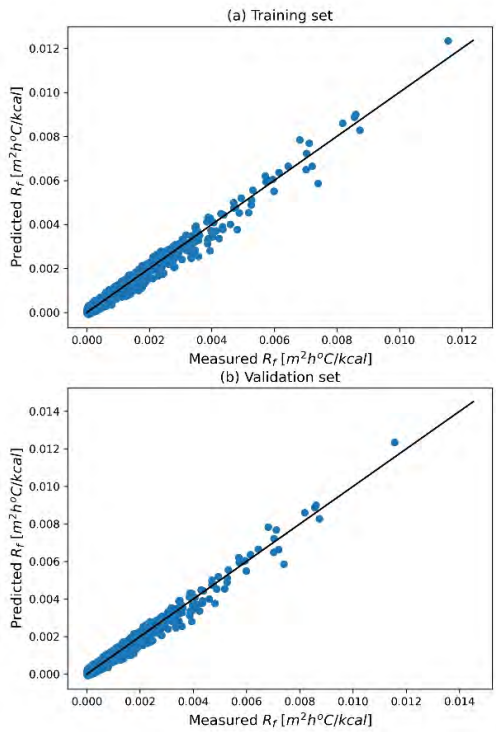

Figure 2. Comparison between the measured values of the fouling resistance and the predictions made by the SR model obtained with data from the first heat exchanger.

Generalization of the Symbolic Regression Equation

To assess the transferability of the SR model, the expression identified using data from the first heat exchanger (Eq. 4) was applied to the second heat exchanger by re-estimating its numerical parameters. Eq. 4 was chosen because of its better performance compared to Eq. 5, as presented in Tables 1 and 2. In this step, the functional structure of the symbolic equation was kept unchanged, and only a reduced set of coefficients was adjusted using plant data from the second unit. Parameter estimation was performed using particle swarm optimization (PSO) implemented with the *pymoo* Python library [23].

The predictive performance of the resulting model, denoted as Eq. 6, is summarized in Table 3. The re-parameterized equation achieved coefficients of determination R^2 of 0.96 and 0.91 for the training and validation datasets, respectively, with correspondingly low MSE values. These values are comparable to those obtained when fitting a symbolic regression model directly to data from the second heat exchanger (Table 2), indicating that the symbolic structure of the first heat exchanger was sufficiently flexible to describe fouling behavior across different units after limited recalibration.

$$R_f = -1.55T_{ef1}R_q - 6.79T_{ef1} + \frac{1.64T_{ef1}}{T_{ef2}} - 6.64T_{ef2} + 5.81 \quad (6)$$

.

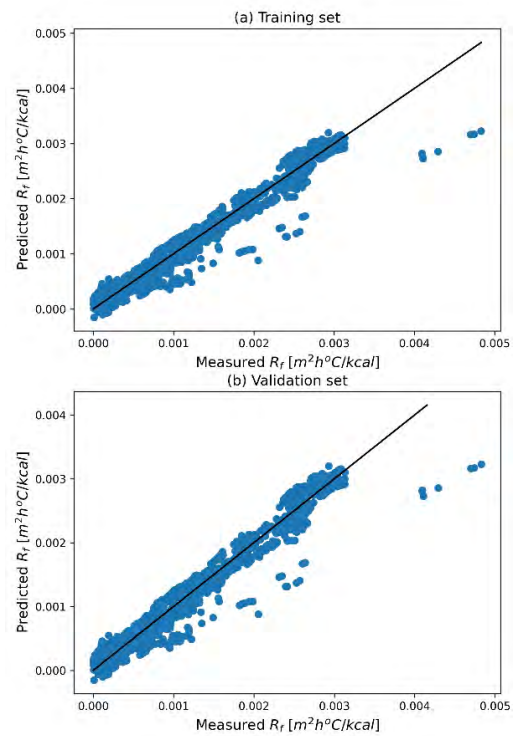

Figure 3. Comparison between the measured values of the fouling resistance and the predictions made by the SR model obtained with data from the second heat exchanger.

Figure 4 compares the measured fouling resistance values with the predictions obtained using Eq. 6. This figure presents a strong agreement between measured and predicted values over the full operating range, with no evident systematic deviation. The dispersion around the parity line remains similar to that observed for the heat exchanger-specific symbolic model of Eq. 5, confirming that re-estimating a small number of parameters is sufficient to adapt the model to a new heat exchanger

Table 3: Coefficient of determination (R^2) and mean squared error (MSE) of Eq. 6 to predict the training and validation sets for the second heat exchanger.

Dataset	R^2	MSE
Training	0.956	$3.72x10^{-8}$
Validation	0.910	$1.27x10^{-8}$

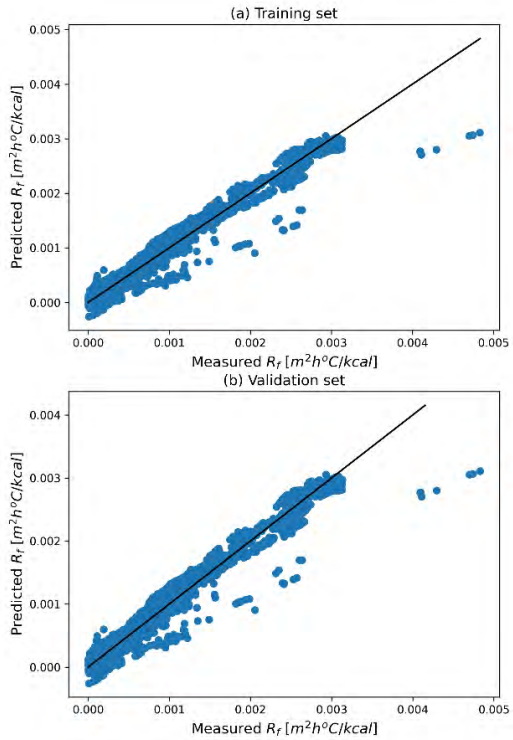

Figure 4. Comparison between the measured values of the fouling resistance and the predictions made by the SR model obtained with data from the first heat exchanger after estimating the model parameters with data from the second heat exchanger.

From an industrial perspective, the proposed approach presented a key advantage, which is the ability to obtain structurally transferable equations that can be adapted to new heat exchangers with minimal additional data. Rather than repeating the full SR procedure for each unit, plant engineers can reuse the learned functional form and perform a lightweight parameter estimation step, significantly reducing modeling effort and data requirements. This characteristic makes the proposed approach particularly attractive for large heat exchanger networks, where rapid and consistent fouling modeling across multiple units is essential for monitoring, optimization, and maintenance decision support.

CONCLUSION

In this work, a SR–based methodology was proposed for modeling fouling resistance in industrial heat exchangers using routinely available plant measurements. By expressing operating conditions through physically meaningful dimensionless variables, the approach enabled the identification of compact and interpretable algebraic equations that accurately described fouling behavior. The analysis of two industrial heat exchangers showed that SR can achieve high predictive accuracy for both training and validation datasets while avoiding the opacity commonly associated with black-box machine learning models. The explicit nature of the resulting expressions facilitates model inspection, interpretation, and direct integration into existing monitoring and analysis workflows.

Beyond predictive performance, the proposed framework was shown to provide transferability across heat exchangers. The symbolic equation identified for one heat exchanger could be successfully adapted to a second exchanger through limited parameter re-estimation, preserving high accuracy without repeating the full model discovery process. This capability highlights the potential of SR as a practical tool for fouling modeling at the plant and network levels, enabling rapid deployment, reduced data requirements, and consistent modeling across multiple units. Overall, the methodology offers a robust and interpretable alternative to conventional empirical correlations and black-box data-driven models, supporting improved fouling monitoring and decision-making in industrial heat exchanger operation.

ACKNOWLEDGEMENTS

This study was financed in part by CAPES - Finance Code 001 and by Petrobras S.A. (Cooperation term no. 0050.0125244.23.9). Professors MB de Souza Jr. and AR Secchi are grateful to financial support from CNPq (Grants No. 311153/2021-6 and 300744/2025-0).

AUTHOR IDENTIFIERS

Author ORCIDs:
Lima FARD: 0009-0004-5025-3804
Oliveira SGO: -
Campos AB: 0009-0000-1328-8339
Assis BCG: -
Costa LPL: -
Liporace FS: -
Souza MB Jr: 0000-0002-1090-8958
Secchi AR: 0000-0001-7297-3571

REFERENCES

1. Müller-Steinhagen H, Malayeri MR, Watkinson AP. Heat exchanger fouling: mitigation and cleaning strategies. *Heat Transf Eng.* 32(3–4):189–196

(2011). https://doi.org/10.1080/01457632.2010.505127
2. Coletti F, Macchietto S. A dynamic, distributed model of shell-and-tube heat exchangers undergoing crude oil fouling. Ind. Eng. Chem. Res. 50:4515-4533 (2011). https://doi.org/10.1021/ie901991g
3. Diaz?Bejarano E, Coletti F, Macchietto S. A new dynamic model of crude oil fouling deposits and its application to the simulation of fouling?cleaning cycles. AIChE Journal 62:90-107 (2015). https://doi.org/10.1002/aic.15036
4. Nebot E, Casanueva JF, Casanueva T, Sales D. Model for fouling deposition on power plant steam condensers cooled with seawater: effect of water velocity and tube material. International Journal of Heat and Mass Transfer 50:3351-3358 (2007). https://doi.org/10.1016/j.ijheatmasstransfer.2007.01.022
5. Mayer M, Bucko J, Benzinger W, Dittmeyer R, Augustin W, Scholl S. Modeling fouling factors for microscale heat exchangers. Experimental Heat Transfer 28:222-243 (2014). https://doi.org/10.1080/08916152.2013.854284
6. Palmer KA, Hale WT, Such KD, Shea BR, Bollas GM. Optimal design of tests for heat exchanger fouling identification. Applied Thermal Engineering 95:382-393 (2016). https://doi.org/10.1016/j.applthermaleng.2015.11.043
7. Qureshi BA, Zubair SM. Predicting the impact of heat exchanger fouling in power systems. Energy 107:595-602 (2016). https://doi.org/10.1016/j.energy.2016.04.032
8. Yang J. Computational fluid dynamics studies on the induction period of crude oil fouling in a heat exchanger tube. International Journal of Heat and Mass Transfer 159:120129 (2020). https://doi.org/10.1016/j.ijheatmasstransfer.2020.120129
9. Berce J, Zupan?i? M, Može M, Golobi? I. A review of crystallization fouling in heat exchangers. Processes 9:1356 (2021). https://doi.org/10.3390/pr9081356
10. Rehman OU, Ramasamy MG, Rozali NEM, Mahadzir S, Ghumman ASM, Qureshi AH. Modeling strategies for crude oil-induced fouling in heat exchangers: a review. Processes 11:1036 (2023). https://doi.org/10.3390/pr11041036
11. McBride K, Sundmacher K. Overview of surrogate modeling in chemical process engineering. Chemie Ingenieur Technik 91:228-239 (2019). https://doi.org/10.1002/cite.201800091
12. Özdemir P, Yıldırım R. The use of artificial intelligence for optimization and control of chemical processes: A review. *ChemBioEng Rev*. 12(4):e70012 (2025). https://doi.org/10.1002/cben.70012
13. Lima FARD, de Moraes MGF, Barreto AG Jr, Secchi AR, Grover MA, de Souza MB Jr. Applications of machine learning for modeling and advanced control of crystallization processes: developments and perspectives. Digital Chemical Engineering 14:100208 (2025). https://doi.org/10.1016/j.dche.2024.100208
14. Sundar S, Rajagopal MC, Zhao H, Kuntumalla G, Meng Y, Chang HC, Shao C, Ferreira P, Miljkovic N, Sinha S, Salapaka S. Fouling modeling and prediction approach for heat exchangers using deep learning. International Journal of Heat and Mass Transfer 159:120112 (2020). https://doi.org/10.1016/j.ijheatmasstransfer.2020.120112
15. Davoudi E, Vaferi B. Applying artificial neural networks for systematic estimation of degree of fouling in heat exchangers. Chemical Engineering Research and Design 130:138-153 (2018). https://doi.org/10.1016/j.cherd.2017.12.017
16. Wang Y, Wagner N, Rondinelli JM. Symbolic regression in materials science. MRS Communications 9:793-805 (2019). https://doi.org/10.1557/mrc.2019.85
17. Lima FARD, de Moraes MGF, Rebello CM, Barreto AG Jr, Secchi AR, de Souza MB Jr, Nogueira IBR. Interpretable and uncertainty-aware machine learning for trustworthy prediction in batch crystallization. Chemical Engineering and Processing - Process Intensification 215:110350 (2025). https://doi.org/10.1016/j.cep.2025.110350
18. Neumann P, Cao L, Russo D, Vassiliadis VS, Lapkin AA. A new formulation for symbolic regression to identify physico-chemical laws from experimental data. Chemical Engineering Journal 387:123412 (2020). https://doi.org/10.1016/j.cej.2019.123412
19. Rebello CM, Costa EA, Fontana M, Schnitman L, Nogueira IBR. Interpretable scientific machine learning approach for correcting phenomenological models: methodology validation on an ESP prototype. Ind. Eng. Chem. Res. 63:19030-19050 (2024). https://doi.org/10.1021/acs.iecr.4c02104
20. Santana VV, Costa E, Rebello CM, Ribeiro AM, Rackauckas C, Nogueira IBR. Efficient hybrid modeling and sorption model discovery for non-linear advection-diffusion-sorption systems: a systematic scientific machine learning approach. Chemical Engineering Science 282:119223 (2023). https://doi.org/10.1016/j.ces.2023.119223
21. Liporace FS, De Oliveira SG. Real time fouling diagnosis and heat exchanger performance. Heat Transfer Engineering 28:193-201 (2007).

https://doi.org/10.1080/01457630601064595
22. Cranmer M. Interpretable machine learning for science with PySR and SymbolicRegression.jl. *arXiv* arXiv:2305.01582 (2023) https://arxiv.org/abs/2305.01582
23. Blank J, Deb K. Pymoo: multi-objective optimization in python. IEEE Access 8:89497-89509 (2020). https://doi.org/10.1109/access.2020.2990567

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Energy Integration Via Heat Pump in a Simulated Fluidized TSA Column for CO_2 Capture from Biomass-Derived Flue Gases

Eduardo S. Funcia[a], Yuri S. Beleli[a], Enrique Vilarrasa-Garcia[b], Marcelo M. Seckler[a], José L. Paiva[a], Galo A. C. Le Roux[a]*

[a] University of Sao Paulo, Polytechnic School, Sao Paulo, Sao Paulo, Brazil
[b] Federal University of Ceara, Fortaleza, Ceara, Brazil
* Corresponding Author: galoroux@usp.br.

ABSTRACT

We present a steady-state, optimization-based techno-economic study of a continuous fluidized temperature-swing adsorption (TSA) system for post-combustion CO_2 capture from biomass-derived flue gas, using two adsorption stages and one desorption stage with integrated heat-pump thermal management. The GAMS/CONOPT4 model couples molar and energy balances, Toth adsorption equilibrium, fluidized-bed hydrodynamics and literature cost correlations. Optimization yields CO_2 purity of 96% v/v and 95.5% recovery at low, safe pressures with particle Reynolds numbers of 2-11, indicating near-minimum-fluidization operation. The nominal capture cost is 87 USD/tonCO_2 with an internal rate of return of 42%; utilities comprise 49% of annualized costs and the adsorption compressor dominates equipment capital. Disabling the heat pump increases modeled capture cost to 124 USD/tonCO_2, highlighting the heat pump's decisive role in reducing energy demand and costs. Adding adsorption stages lowers modeled cost further but produces impractically shallow beds, revealing a trade-off between mass-transfer performance and feasible bed geometry. A sensitivity analysis demonstrated a linear positive correlation between electricity cost and capture cost, with capture costs potentially as low as 50 USD/tonCO_2 and an internal rate of return as high as 64%. While results demonstrate the technical and economic plausibility of a heat-pump-assisted fluidized TSA under the stated assumptions, the study recommends including flue contaminants, refined heat-exchanger and pressure-drop models, and experimental validation for scale-up.

Keywords: Carbon Dioxide Capture, GAMS, Modelling and Simulations, Technoeconomic Analysis, Adsorption

INTRODUCTION

The 2015 Paris Agreement, ratified by 196 countries, established the ambitious objective of constraining global mean temperature rise to 1.5 °C through rapid reductions in greenhouse gas emissions — a target that demands substantial technological and policy effort [1]. Within the broader decarbonization agenda, carbon capture and storage (CCS) technologies are frequently identified as essential pathways to reduce emissions from hard-to-abate sectors [2]. In particular, bioenergy with carbon capture and storage (BECCS) has attracted attention because it can deliver net negative CO_2 emissions by capturing and permanently storing carbon that was originally sequestered from the atmosphere during biomass growth [3].

Despite this promise, most optimization studies addressing BECCS have concentrated on linear, system-level representations of spatial and temporal networks (plants, feedstock supply, and transport logistics), while paying limited attention to the internal operational dynamics of individual capture units [3, 4]. High-fidelity, process-level simulation and optimization can reveal critical interactions among process variables, energy flows, and equipment design that are necessary to bring emerging capture technologies from laboratory concepts to

industrially robust, cost-competitive systems. One such emerging technology is temperature-swing adsorption (TSA) implemented in a fluidized configuration, where a moving solid adsorbent circulates counter-current to an upward flow of CO_2-rich flue gas. Solid sorbents as binder-free zeolite 13X offer attractive properties for this application, including high CO_2 selectivity and reduced regeneration energy compared with many alternative materials [5, 6, 7]. Process-level optimization studies of TSA processes are comparatively rare in the literature and typically restricted to small-scale demonstrations that lack process cost estimates [8].

The authors' previous work on optimization of a simulated fluidized TSA process yielded interesting and potentially cost-saving insights while delivering a capture process cost estimate of about 84 USD/tonCO_2 [7]. However, that work relied on the heavy hypothesis of stage-wise mass-transfer equilibrium reached, and it lacked proper energy management between heat exchangers. Therefore, this study aims to improve on that previous work of an optimization-based process scale-up through developing a guidelines framework via in silico simulations.

This objective also aligns with the Sustainable Development Goals set by the United Nations for 2030 [9], including:

- 7: Affordable and Clean Energy
- 9: Industry, Innovation and Infrastructure
- 11: Sustainable Cities and Communities
- 12: Responsible Consumption and Production
- 13: Climate Action.

MATERIALS AND METHODS

The fluidized-bed TSA system simulated and optimized in this study was described in the authors' previous work [5, 6, 7]. It consists of two multi-stage fluidized columns – one for adsorption and another for desorption – with continuous adsorbent recirculation between them. The desorption column has a heat exchanger at the gas stream in its bottom, and each fluidized bed in each column also has a dedicated heat exchanger to optimize local temperature. The recirculating solid adsorbent also passes through a heat exchanger between the columns. Part of the gas desorbed in the desorption column is purged and recovered, and the rest is used as stripping gas recycled to the desorption column.

The number of adsorption stages was optimized based on CO_2 purity targets and design efficiency considerations. Our initial simulations indicated that only two adsorption stages and a single desorption stage could obtain CO_2 at over 90% v/v purity, a common benchmark in the field [10]. As will be discussed, adding more stages result in impractically wide and shallow adsorption beds.

A compressor is used at the bottom of each column to propel the gas upwards. A heat pump connects the solid-stream heat exchanger between columns to the desorption bed heat exchanger.

Table 1 presents the hypotheses for the process modeling, while Table 2 presents the values and units for several parameters. Figure 1 presents the process schematics, including adsorption and desorption stages, heat exchangers, compressors, the heat pump and main process streams.

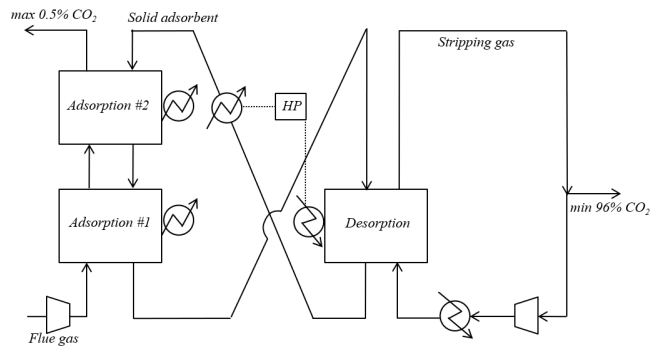

Figure 1. Simulated constinuous fluidized TSA process schematics.

Table 1: List of hypotheses for the fluidized TSA simulation.

No chemical reactions
No phase changes
No flue gas contaminants
Steady state
Ideal gases
No gradients within each stage
Spheric solid adsorbent particles
Air viscosity independent from pressure
Taxes and deprecation neglected
Labor, storage and transport costs neglected
Heat exchangers pressure drop neglected

The Marshall & Swift index, which adjusts equipment prices for inflation, was estimated as 2500 after extrapolation from values reported in the literature [12, 13, 14]. The monthly adsorbent loss was increased from 10% in our previous work [7] to 30% in this study to account for both mechanical losses and adsorption sites deactivation due to impurities. The flue gas molar flow was left as a variable to be fitted. The flue gas and the captured CO_2 molar compositions have also been changed from the authors' last work [7], which should be taken into consideration when comparing the cost estimates.

The process model couples molar and energy balances for each gaseous species in every stage of the adsorber with balances for the external heat-exchange loops, and it explicitly incorporates the thermodynamics of adsorption and the fluid dynamics characteristic of

fluidized beds. Adsorption equilibrium within each stage is represented by the Toth isotherm [19], using parameter values taken from prior experimental characterization [6]; the heterogeneity parameter for N_2 is fixed at unity and treated as temperature-independent for the present simulations. Mass and energy conservation equations are complemented by algebraic constraints representing thermodynamic relations, fluid dynamic closures, equipment efficiencies and cost correlations; together these expressions define the nonlinear constrained optimization problem.

Table 2: List of parameters for the fluidized TSA simulation.

Parameter description	Parameter value
Flue gas temperature	318.16 K
Flue gas pressure	1.01325 bar
Flue gas molar composition	10% v/v CO_2 and 90% v/v N_2
CO_2 max molar concentration after adsorption	0.5% v/v [5, 6, 7]
CO_2 min molar concentration after desorption	96% v/v
Zeolite 13XBF specific heat	920 J/kg.K [6]
Zeolite 13XBF bulk density	1287 kg/m^3 [6]
Zeolite 13XBF spheric particle diameter	0.5 mm [5, 6, 7]
Compressors efficiency	72% [11]
Overall transfer coefficient in each heat exchanger	200 W/m^2.K [11]
Marshall & Swift inflation index	2500 [12, 13, 14]
Work hours per year	8000 h/year [14]
Plant lifetime	20 years
Interest rate	10%
Cold water cost	0.506 USD/GJ [15, 16]
Steam cost	25.3 USD/GJ [15, 16]
Electricity cost	20.0 USD/GJ
Zeolite 13XBF cost	2 USD/kg
Adsorbent monthly losses	30%
Start-up cost	10% of fixed capital [14]
CO_2 market price	120 USD/tonCO_2
Carbon credit	60 USD/tonCO_2 [17]
Max column diameter	3.048 m [18]
Column height/diameter min ratio	1

The efficiency of CO_2 mass transfer between solid and gas, defined in Equation 1 (where q is the amount of gas adsorbed), was estimated by an empirical correlation found via in silico adsorption simulations, as shown in Equation 2 [20]. In Equation 2, u is the gas velocity, u_{mf} is the minimum fluidization velocity, S/G is the ratio between solid and gas mass flows, τ_S is the solid mean residence time and T is the temperature. The efficiency of N_2 mass transfer was left as a variable $\in [0, 1]$ to be fitted. The fixed capital investment (FCI) consisted in literature correlations for equipment cost for columns, heat exchangers, compressors and the heat pump [18, 21].

The heat pump is modeled as an auxiliary energy-recovery unit that redirects thermal energy from the solid-cooling heat exchanger between columns into useful heat for regeneration at the desorption stage [26]. Its required output heat duty is linked to the process heat demand, while the electrical power consumption is computed from a temperature-dependent coefficient of performance ($COP > 1$) as shown in Equation 3, where λ is the heat pump efficiency, Q_{out} is the heat output of the heat pump, W is the electric work supplied to the heat pump, T_{des} is the desorption stage temperature, and T_{Sout} is the outlet solid temperature at the heat exchanger between columns.

$$\eta_{i,j} = \frac{q_{i,j} - q_{i,j+1}}{q_{i,j}^{(Eq)} - q_{i,j+1}}, i = \mathrm{CO_2}, \mathrm{N_2}, j = \mathrm{stage} \quad (1)$$

$$\eta_{\mathrm{CO2}} = 1 - 0.1018 \left(\frac{u}{u_{mf}}\right)^{3.298} \left(\frac{S}{G}\right)^{0.1105} \left(\frac{\tau_S}{60\ \mathrm{s}}\right)^{-0.3055} \left(\frac{T}{300\ \mathrm{K}}\right)^{-0.7377} \quad (2)$$

$$COP = \frac{Q_{out}}{W} = \lambda\left(\frac{T_{des}+5}{(T_{des}+5)-(T_{Sout}-5)}\right) \quad (3)$$

The entire model (mass, energy, Toth equations, process economics and auxiliary relations) was encoded in GAMS (GAMS Software GmbH, release 52.5.0) and solved with the nonlinear programming solver CONOPT4 [22]. The formulation comprises 249 constraints (including inequality constraints) and 207 variables, with 58 definitional constraints – constraints which can be written as *a=f(b, c, d, ...)* and are solved separately after the problem is optimized. The optimization objective is the minimization of annualized process cost normalized by the annual mass of CO_2 captured.

Because the formulation is nonlinear, solver performance depended strongly on variable initializations and bounds [23]. To promote convergence we enforced physically meaningful upper and lower bounds and initial guesses on pressures, temperatures, gas properties and Toth parameters (set to their theoretical minima and maxima where applicable), and supplied strictly positive initial guesses for variables that would otherwise produce undefined expressions if equal to null.

RESULTS AND DISCUSSION

Tables 3 and 4 present the optimized values of main variables. The optimization resulted in no adsorption stage having a heat exchanger, with their thermal loads being automatically set to null. The external heat exchanger at the bottom of the desorption column was also

unnecessary to the optimized process.

Table 3: Optimized values of main global variables.

Variable	Optimized value
Flue gas flow	94 mol/s (7602 Nm^3/h)
Flue gas pressure	2.13 bar
Adsorbent flow	10.5 kg/s
Stripping gas flow	3.17 mol/s
Stripping gas temperature	465 K
Stripping gas pressure	1.11 bar
Recovered gas flow	9.4 mol/s
CO_2 recovery	95.5%
Internal rate of return	42%
Capture cost	87 USD/tonCO_2
Heat pump *COP*	2.0
Heat pump efficiency	62%
Heat pump electric work	663.4 kW

Table 4: Optimized values of variables per stage.

Variable	Adsorption bottom stage	Adsorption top stage	Desorption stage
Pressure	2.09 bar	1.94 bar	1.01 bar
Temperature	371 K	352 K	465 K
CO_2 efficiency	73%	60%	92%
Porosity	0.49 m^3/m^3	0.53 m^3/m^3	0.42 m^3/m^3
Minimum porosity	0.42 m^3/m^3	0.42 m^3/m^3	0.42 m^3/m^3
Velocity	0.19 m/s	0.25 m/s	0.10 m/s
Minimum velocity	0.12 m/s	0.12 m/s	0.10 m/s
Density	1.94 kg/m^3	1.86 kg/m^3	1.14 kg/m^3
Dynamic viscosity	0.0000217 Pa.s	0.0000209 Pa.s	0.0000257 Pa.s
Solid residence time	485 s	1651 s	189 s
Bed height	0.54 m	2.51 m	1.25 m
Bed diameter	3.048 m	2.62 m	1.25 m
Stage heat	0 W	0 W	+142 kW

The particle Reynolds numbers predicted for both adsorption and desorption stages remain low, ranging from approximately 2 to 11. These values confirm that the solids are readily fluidized and that the column operates close to minimum fluidization conditions, particularly in the desorption section. They also imply that viscous dissipation and inertial effects contribute comparably to the pressure losses associated with gas flow through the particulate phase. Under such conditions, a purely Stokesian drag description is inadequate, and inclusion of both viscous and kinetic terms in the Ergun formulation is therefore warranted [24].

The minimum fluidization porosity calculated from the semi-empirical Ergun correlation is approximately 0.42, in close agreement with the value of 0.43 reported in earlier work [8] and well within the typical 0.40–0.45 range expected for nearly spherical particles [24]. Superficial gas velocities fall within values commonly reported for industrial adsorption operations [25], and the optimized pressure levels across stages further indicate that TSA systems can be designed to operate safely at relatively low pressures.

The findings in this work further support claims that capturing CO_2 at purities greater than 90% is physically as well as economically feasible [10]. Low CO_2 mass transfer efficiencies in the adsorption stages indicate the optimal operating point is a compromise between the mass transfer process and economic parameters.

The objective function value, the CO_2 capture cost per mass unit, was optimized at 87 USD/tonCO_2, a value similar to that previously reported by the authors [7] – which shows that the economies derived from the efficient thermal management by the heat pump must be approximately similar to the increased costs of rejecting the strong stage-wise equilibrium hypothesis. Increasing the number of stages in adsorption resulted in decreased capture costs (to ≈76 USD/tonCO_2) while producing mathematically sound but practically unfeasible beds (of a few centimeters high or less).

In order to evaluate the improvement brought about by the energy management through the heat pump installation, the same process with the same constraints and same initial guess was optimized while setting variable Q_{out} to null. The resulting capture cost was 124 USD/tonCO_2, indicating that the heat pump reduced costs by 37 USD/tonCO_2 (or about 30%) and is therefore essential if the fluidized TSA technology is to find commercial success in the CCS field.

Figure 2 below shows the annualized costs of equipment, adsorbent and utilities. It is evident that utilities are still the main cost drivers for the fluidized TSA CO_2 capture system, comprising approximately half of the total annualized cost. However, that is down from 76% of our previous estimates [7]. Regarding equipment, the most expensive fixed capital investment is the adsorption compressor, nearly twice as expensive as its desorption counterpart.

A sensitivity analysis with respect to the effect of the electricity cost parameter on the capture cost yielded the correlation shown in Figure 3 below. The same initial guesses, parameters and constraints were used in each simulated process optimization. We observe that the overall capture cost increases linearly in a strong correlation with the electricity cost. It also noteworthy that if such electricity cost reduces by one order of magnitude

from the parameter value used in Table 2, the CO_2 capture cost may be reduced to 50 USD/tonCO_2, which suggests that this technology's economic feasibility may depend on regional utilities availability price.

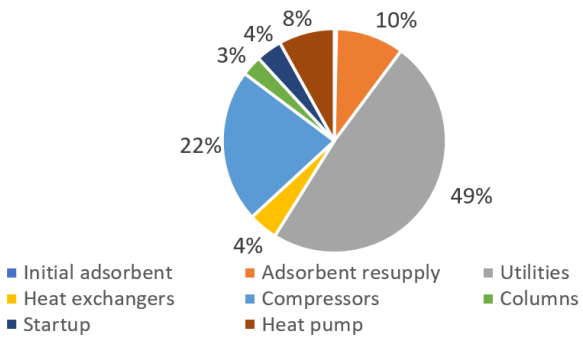

Figure 2. TSA annualized costs distribution.

The optimized capture costs obtained in this study, spanning approximately 50–87 USD/tonCO_2, lie within the range reported for several CCS technologies [2, 7] and are substantially lower than typical costs associated with direct air capture [2, 7]. Although these values remain higher than those reported for mature amine-based absorption processes [7], the corresponding economic performance is encouraging. In particular, the optimized annual internal rate of return reaches about 42% for the nominal case at 87 USD/tonCO_2 and increases to 64% at the lower-cost end of the electricity cost range. These results indicate that, despite its current developmental stage, fluidized TSA presents a level of economic attractiveness that justifies continued investigation, detailed modeling, and experimental validation.

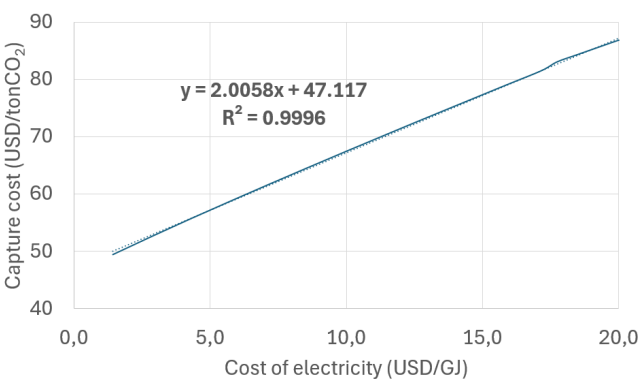

Figure 3. Sensitivity analysis of electricity cost on capture process cost.

CONCLUSION

A continuous fluidized TSA with two adsorption stages and one desorption stage can reach 96% CO_2 purity and 95.5% CO_2 recovery while operating at low, safe pressures and near minimum-fluidization hydrodynamics. The optimized nominal capture cost is about 87 USD/tonCO_2, with an internal rate of return of 42%, numbers that imply the concept is both physically feasible and economically promising. A sensitivity analysis resulted in a positive strongly linear relation between electricity and capture costs, with the optimized value potentially as low as 50 USD/tonCO_2.

A central technical and economic finding is that active thermal management via the heat pump is essential: removing the heat pump from the model raises the optimized capture cost from 87 to 124 USD/tonCO_2. The optimization also reveals trade-offs and practical constraints: mass-transfer efficiencies in adsorption stages are low, so the economic optimum is a compromise between increasing contact and the rising equipment and utility costs that follow. Increasing the number of adsorption stages can further reduce modeled capture costs (to ≈76 USD/tonCO_2) but the resulting optimal bed geometries become practically infeasible (very shallow beds).

In summary, this simulation study supports the physical and economic plausibility of a heat-pump-assisted, fluidized continuous TSA system for post-combustion CO_2 capture; further model refinement plus experimental validation is necessary before strong scale-up claims can be made.

Future developments in this work may include adding other costs, considering the effect of flue gas contamination, or enhancing heat exchangers modelling.

ACKNOWLEDGEMENTS

This study was financed by Brazil's National Council for Scientific and Technological Development – CNPq, scholarship 153526/2024-6, through the INCT project CAPICUA 406710/2022-8. This study was also financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001. We gratefully acknowledge support of the RCGI – Research Centre for Greenhouse Gas Innovation, hosted by the University of São Paulo (USP) and sponsored by FAPESP – São Paulo Research Foundation (2014/50279-4 and 2020/15230-5) and by Shell Brasil. We also thank the strategic and important support given by ANP (Brazil's National Oil, Natural Gas and Biofuels Agency) through the R&D levy regulation.

AUTHOR IDENTIFIERS

Author ORCIDs:
Funcia ES: 0000-0001-5408-2011
Beleli YS: 0009-0007-0967-0915
Vilarrasa-Garcia E: 0000-0003-4190-8763
Seckler MM: 0000-0001-9842-4896
Paiva JL: 0000-0001-9733-5744
Le Roux GAC: 0000-0001-7227-5271

REFERENCES

1. United Nations Framework Convention on Climate Change secretariat. https://unfccc.int/process-and-meetings/the-paris-agreement
2. International Energy Agency. https://www.iea.org/commentaries/is-carbon-capture-too-expensive
3. Sadr M, Esmaeili Aliabadi D, Jordan M, Thrän D. A bottom–up regional potential assessment of bioenergy with carbon capture and storage in germany. Environ. Res. Lett. 19:114047 (2024). https://doi.org/10.1088/1748-9326/ad7edd
4. Patange OS, Garg A, Jayaswal S. An integrated bottom-up optimization to investigate the role of BECCS in transitioning towards a net-zero energy system: a case study from gujarat, india. Energy 255:124508 (2022). https://doi.org/10.1016/j.energy.2022.124508
5. Beleli YS, Paiva JL, Seckler MM, Le Roux GAC. Optimization of a continuous multi-stage fluidized bed system for CO_2 capture utilizing temperature swing adsorption. *Comp Aided Chem Eng* 52:3233-3238 (2023). https://doi.org/10.1016/B978-0-443-15274-0.50515-1
6. Prado DS, Vilarrasa-García E, Sampronha E, Beleli YS, Moreira FS, Paiva JL, Roux GACL, Bastos-Neto M, Azevedo DCS, Silva ECN, Seckler MM. Multiple approaches for large-scale CO2 capture by adsorption with 13X zeolite in multi-stage fluidized beds assessment. Adsorption 30:429-455 (2023). https://doi.org/10.1007/s10450-023-00422-x
7. Funcia ES, Beleli YS, Vilarrasa-Garcia E, Seckler MM, Le Roux GAC. Design of a fluidized temperature swing adsorption process for biomass-derived flue gas carbon capture. Computers & Chemical Engineering 202:109280 (2025). https://doi.org/10.1016/j.compchemeng.2025.109280
8. Gao J, Hoshino Y, Inoue G. Honeycomb-carbon-fiber-supported amine-containing nanogel particles for CO2 capture using a rotating column TVSA. Chemical Engineering Journal 383:123123 (2020). https://doi.org/10.1016/j.cej.2019.123123
9. United Nations. https://sdgs.un.org/goals
10. Brandl P, Bui M, Hallett JP, Mac Dowell N. Beyond 90% capture: possible, but at what cost?. International Journal of Greenhouse Gas Control 105:103239 (2021). https://doi.org/10.1016/j.ijggc.2020.103239
11. Biegler LT, Grossmann IE, Westerberg AW. Systematic Methods of Chemical Process Design. Prentice Hall (1997).
12. Wright DG, Woods DR. Evaluation of capital cost data. part 7: liquid waste disposal with emphasis on physical treatment. Can J Chem Eng 71:575-590 (2009). https://doi.org/10.1002/cjce.5450710411
13. Camaraza-Medina Y, Sánchez-Escalona A, Retirado-Mediaceja Y, García-Morales O. Use of air cooled condenser in biomass power plants: a case study in cuba. IJHT 38:425-431 (2020). https://doi.org/10.18280/ijht.380218
14. Peters MS, Timmerhaus KD. Plant Design and Economics for Chemical Engineers. McGraw Hill (1991).
15. Turton R, Bailie RC, Whiting WB, Shaeiwitz JA, Bhattacharyya D. Analysis, Synthesis, and Design of Chemical Processes. Prentice Hall (2012).
16. Bureau of Labor Statistics, USA government. https://data.bls.gov/cgi-bin/cpicalc.pl
17. Congressional Research Service, USA government. https://www.congress.gov/crs-product/IF11455
18. Douglas JM. Conceptual Design of Chemical Processes. McGraw Hill (1988).
19. Do DD. Adsorption Analysis: Equilibria and Kinetics. Imperial College Press (1998).
20. Research Center for Greenhouse Gas Innovation. *Book of Abstracts: Energy Transition Research & Innovation Conference - ETRI 2024* 528:532 (2024). https://doi.org/10.5281/zenodo.17179238
21. Zuberi MJS, Hasanbeigi A, Morrow W. Techno-economic evaluation of industrial heat pump applications in US pulp and paper, textile, and automotive industries. Energy Efficiency 16: (2023). https://doi.org/10.1007/s12053-023-10089-6
22. GAMS Software GmbH. https://www.gams.com/46/docs/S_CONOPT4.html
23. GAMS Software GmbH. https://www.gams.com/latest/docs/UG_NLP_GoodFormulations.html
24. Tallmadge JA. Unit operations of chemical engineering. warren L. mccabe and julian C. smith. mcgraw?hill book company, inc., new yark (1956). 945 pages. $10.50. AIChE Journal 3: (2004). https://doi.org/10.1002/aic.690030126
25. Gabelman A. Adsorption Basics: Part 1. American Institute of Chemical Engineers (2017) https://www.aiche.org/sites/default/files/docs/pages/adsorption_basics_part_1.pdf

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Towards White-box Environmental and Economic Process Optimization: Tailoring Modelling Approaches to Multi-scale Simulations.

Thomas Hietala[a]*, Sonja Herres-Pawlis[b], Pedro S.F. Mendes[a]

[a] Centro de Química Estrutural - IMS, Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal
[b] Institute of Inorganic Chemistry, RWTH Aachen University, Aachen, Germany
* Corresponding Author: thomas.hietala@tecnico.ulisboa.pt.

ABSTRACT

Polylactic acid (PLA) is the most produced bioplastic, however, for it to compete with fossil-based plastics, maximum production efficiency is crucial. To achieve this, all scales, from catalyst scale to process scale, must be simultaneously considered. This challenge is particularly relevant for PLA production, where complex interactions of multiple phenomena occur in several unit operations and the development of active non-toxic catalysts is of major importance. For these reasons, developing a framework that allows a comprehensive understanding of the influence of design choices at different levels is of paramount importance. To address this, a multi-scale model was developed for the PLA production, coupling a custom devolatilizer reactor model to a process simulator that is subsequently linked to an environmental assessment tool via Python-based interfaces. With this model, sensitivity analysis was performed to assess the influence of operational variables on the most relevant KPI's. Results showed that changes of only 5.5 torr in the devolatilizer pressure can lead to total utility consumption and climate change impact variations of 13% and 2%, respectively. This demonstrates the critical role of the devolatilizers in the process, highlighting the importance of developing custom unit operation models coupled to the process simulator. Additionally, the results show that, with this framework, the influence of reactor scale variables on industrial environmental performance can be seamlessly assessed. Overall, this work demonstrates the feasibility and value of integrated multi-scale modeling for sustainable process design and provides a solid foundation for the future development of a multi-scale environmental and economic optimization framework.

Keywords: Sustainable process design, Multi-scale modeling, Custom unit operation model, Automated environmental assessment

INTRODUCTION

Polylactic acid (PLA) is the most produced bioplastic [1], but its carbon footprint, among other environmental metrics, still has margins that need to be improved, together with the production cost. To achieve utmost efficiency, all scales matter, from phenomena in a given unit operation to mass and energy integration, making multi-scale models necessary for the optimization of these processes. Currently, as the design of sustainable processes is mostly performed independently at different scales, the overall influence of design choices at different levels is not assessed in a seamless way, leading to a trial-and-error and inefficient design workflow.

In particular, in the case of PLA production, there are intricate interplay of mass-transfer, evaporation, and reaction phenomena in several unit operations [2]. Additionally, at the moment, stannous octanoate (tin octanoate, $Sn(Oct)_2$), is commonly used as a catalyst. Although classified as food safe, toxic tin can be released in the environment during PLA degradation [3]. Therefore, ongoing research is being made to develop nontoxic metal-based catalysts for this process [3]. To efficiently evaluate the impact of using different catalysts on the overall economic and environmental performance of the process, a multi-scale model is essential.

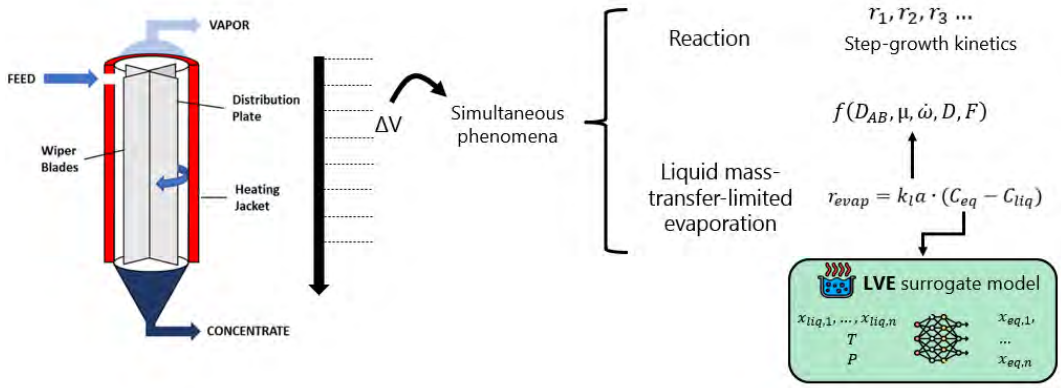

Figure 1. Schematic of the devolatilizer [6] used in the depolimerization reaction and the final PLA purification as well as a schematic for the custom reactor model developed.

To address this topic, in this work, a multi-scale model for PLA production was developed. As of today, the model covers the reactor, process, and environmental analysis scale. The process scale is simulated in Aspen Plus, to which a custom devolatilizer model is coupled using the Python Unit Operation tool [4].
For the environmental analysis, a Python interface that links the process model with OpenLCA is used. With this multi-scale model, sensitivity analysis was performed to evaluate the impact of different possible decision variables on relevant technical and environmental KPI's. The knowledge gained in this paper is crucial for the future development of a holistic optimization algorithm.

METHODOLOGY

Reactor-scale model via Python-implemented detailed devolatilizer model

The synthesis of PLA, starting from lactic acid, consists of three steps: oligomerization of lactic acid into a PLA pre-polymer with low molecular weight (MW ~1 kg/mol), depolymerization into lactide, and ring-opening polymerization (ROP) into high molecular weight PLA (100 kg/mol). In the depolymerization reaction, the low molecular weight PLA, with the aid of the catalyst, converts into lactide. However, if the lactide remains in the reaction mixture, the reaction will reach an equilibrium at low pre-polymer conversions. To favor the product side of the reaction, the lactide must be continuously removed from the liquid reaction mixture into the vapor phase. As the reactional mixture is quite viscous, due to the pre-polymer's molecular weight, the mass-transfer of the volatile components (water, lactic acid and lactide) from the liquid mixture to the liquid-vapor interface is limited, rendering the evaporation of the product and thus the progression of the reaction quite challenging. In the final high molecular weight PLA mixture purification, the viscosity of the mixture is even more relevant, turning the evaporation of the volatiles even more mass-transfer limited. For these reasons, in industry, both the depolymerization and the final PLA purification are performed in depolymerization equipments [5], schematized in Figure 1.

In these equipments, the liquid polymeric feed mixture travels in a thin liquid film, forced by the movement of wiper blades, reason why the equipment is also called "wiped-film evaporator". Being the mixture in a thin film, the liquid-vapor interfacial area is maximized, necessary to counteract the mass-transfer limitations in the liquid phase. This way, with high temperatures being maintained with heated jackets, the volatile components, such as lactide, will evaporate and leave the reactor in the vapor phase at the top of the equipment, while the polymeric mixture will leave the reactor at the bottom of the equipment. As Aspen Plus lacks a built-in model to adequately simulate all phenomena (reaction, mass transfer and evaporation) simultaneously, a custom model was developed to simulate the devolatilizers. The models were developed using the Python Unit Operation (PUO) [4], a CAPE-OPEN complying tool, that allows integration of Python code into process simulators, as can be seen in Figure 3, in which the devolatilizer models are identified as "*Python Unit Operation*" boxes.

As schematized in Figure 1, the reactor model discretizes the reactor volume into several control volumes. In each control volume, the reaction rates are calculated via step-growth kinetics implemented from literature [7]. Then, the evaporation rates of each component, $r_{evap,i}$, are computed using Equation 1, where $k_{l,i}$ is the mass-transfer coefficient of component *i* in the polymeric mixture, a is the interfacial area, $C_{eq,i}$ is the equilibrium concentration in the liquid phase, and $C_{liq,i}$ is the liquid phase concentration of component *i*.

$$r_{evap,i} = k_{l,i} a \bullet (C_{eq,i} - C_{liq,i}) \quad (1)$$

The mass transfer coefficient in devolatilizers is calculated through several thermodynamic and fluid

mechanic correlations from literature [8] and will depend of a number of different variables, such as the diffusivity of the volatile components in the in the polymeric mixtures (estimated with Stokes-Einstein equation), D_{AB}, the viscosity of the polymeric mixture, μ, the rotation speed of the wiper blades, $\dot{\omega}$, the diameter of the equipment, D, and the flowrate of the liquid mixture, F. The viscosity of the polymeric mixture is calculated as a function of the solute weight fraction, the molecular weight of the polymer, the temperature and the shear rate proportioned by the wiper blades through correlations given in the literature [9].

For the liquid-vapor equilibrium (LVE) concentrations calculation, a PUO-integrated function that connects to the process simulators thermodynamic model exists, however, as this would have to be done in every control volume, it would lead to high simulation times (up to 1 hour for a single run) and convergence issues that would render a future optimization framework very inefficient. Instead, neural networks (NN) were trained on the PolyNRTL thermodynamic data to calculate the LVE compositions at every control volume of the reactor. The NN model inputs are the input compositions, temperature, pressure, and molecular weight of the polymer, and the output data are the liquid-phase output compositions at equilibrium. In Figure 2, the parity plot for the testing of the NN used to calculate the lactide liquid equilibrium compositions with a 5% error cone is presented. In this case, the NN has 3 layers with 20, 16 and 12 nodes.

Process-scale model

This process model focuses on the transformation of lactic acid to PLA, as licensed by NatureWorks [5] (the largest PLA producer globally), based on an earlier implementation [7]. The simulations were done in Aspen Plus v14 software using its Polymer Plus module and PolyNRTL as thermodynamic method [7]. The process has three steps (see Figure 3): oligomerization of lactic acid into pre-polymer, its depolymerization into lactide, and, after a purification step, the ROP reaction into PLA, followed by the purification of the final PLA product. By default, detailed models were selected from the Aspen library apart from the devolatilizers, as mentioned previously.

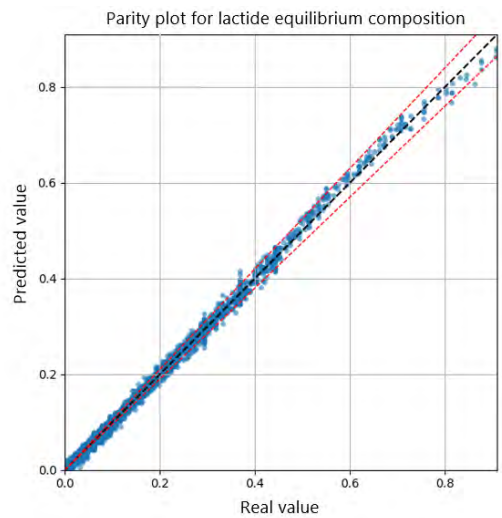

Figure 2. Parity plot of the neural network trained to calculate the lactide equilibrium liquid composition. The neural network uses 3 layers with 20, 16 and 12 nodes.

For the oligomerization step, lactic acid is converted into pre-polymer in a stirred-tank reactor connected to a distillation column to remove the produced water. The base case operating conditions of the reactor are 172ºC and 400 torr [7]. Next, the pre-polymer is depolymerized in a devolatilizer, i.e., an agitated thin-film evaporator, with simultaneous removal of volatiles, which will be further explained in the next section. The base case conditions for the reactor are an inlet temperature of 170ºC and an outlet temperature of 210ºC at 5 torr [7]. The non-reacted pre-polymer is recirculated into the reactor, while

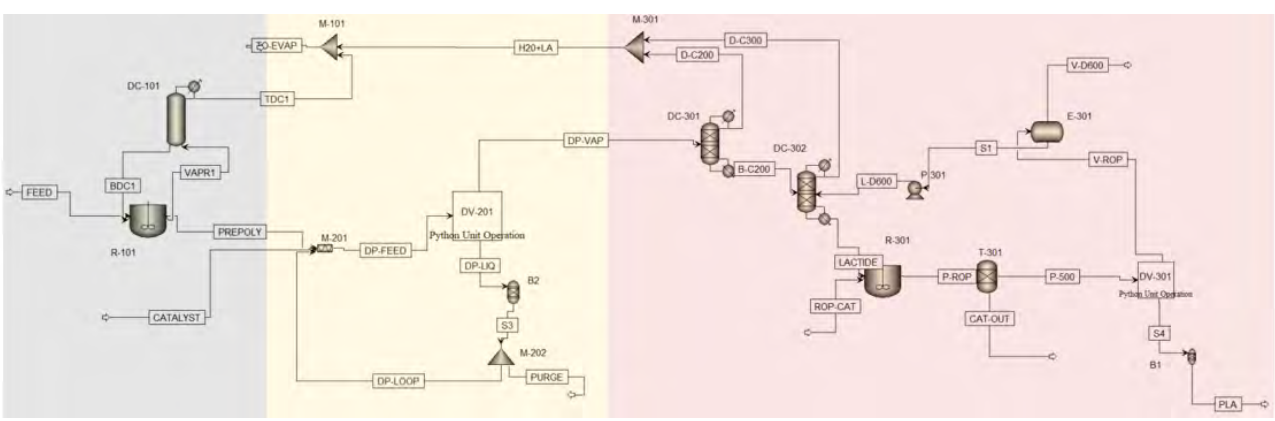

Figure 3. Aspen Plus model flowsheet developed in this work for the production process of PLA divided into three sections: Oligomerization in grey, depolymerization in yellow and ring-opening polymerization in red.

the produced lactide stream continues to the ROP section. This section starts with a purification section, constituted by two distillation columns. In the first column, the water from the lactide mixture is removed, while in the second column, the lactic acid is removed, both in the distillate, to render a purified lactide mixture in the base of the second column. The lactide stream is then converted, in a second stirred-tank reactor, into high MW PLA which undergoes a final purification in a devolatilizer to remove impurities. At the base case, the stirred tank reactor operates at 170ºC and 1 atm while the devolatilizer operates at 180ºC at the inlet, 200ºC at the outlet and 0.5 torr [7]. The crystallization process to form the PLA pellets was not simulated in this work.

To ensure that the modeled process produces the desired PLA production in terms of flowrate, molecular weight and purity, design specs were added to the Aspen Plus simulation. To fix the PLA flowrate at 8610 kg/hr (estimated value for a single production line of PLA), the lactic-acid feed mixture (88% wt.) flowrate is varied. The molecular weight of 100 kg/mol is ensured by varying the DC302 column distillate flowrate, as the molecular weight of the PLA produced will be mainly controlled by the purity of the lactide mixture that enters the ROP reactor. The higher the purity, the higher the molecular weight of the produced PLA will be. Finally, the purity of the final PLA mixture was kept at 98% wt. by varying the length of the DV301 devolatilizer, which can be translated into the number of devolatilizers needed in series. The values of the main KPI's for the base case operational variables values can be found in Table 1.

LCA model via Python interface

To perform an accurate and holistic environmental assessment of the modeled process, OpenLCA, an open-source LCA tool, was used. The LCA analysis was done with "*cradle-to-gate*" borders, meaning that the environmental impact includes the impact of all production processes from the extraction of the raw material up to the PLA product leaving the factory gate. This LCA analysis excludes the impact of the use of PLA and of its end-of-life. The Ecoinvent 3.11 was used as a database for the LCA analysis. The CML v4.8 2016 impact assessment method was used in the work. The LCA model consists of adding up the environmental impact of the necessary lactic acid production, for which the environmental impact was already defined in the database (corresponding to the production process of the manufacturer Cargill) and the environmental impact of the production of the different utilities used in the transformation of the lactic acid to PLA (such as steam, electricity and water). The climate change impact of 1 kg of lactic acid, through the Cargill production process, is 4.50 kg CO_2 eq.

To perform the LCA analysis, the necessary parameters to extract from the Aspen Plus process model are the lactic acid flowrate and the different utility usages. To extract these values, a Windows COM automation interface between Aspen Plus and Python was developed using the *win32com* library and Aspen's Variable Node Tree structure. Using this interface, the values from Aspen's simulation can be extracted, new parameter values can be set and new simulations can be run, which is critical for the sensitivity analysis performed in this work, and the future optimization framework. To run the LCA model with the values extracted, the *olca* library was used, which allows the coupling of the Python script with the OpenLCA model.

RESULTS AND DISCUSSION

Before implementing an optimization framework, it is crucial to identify the operational variables that influence the KPI's to be optimized, as these are best suited to be the decision variables. The operational variables tested were the ones related to the most relevant equipment. This leads to the following variables: the temperature and pressure of the reactors R101 and R301 and of the devolatilizers DV201 and DV301, and the blade rotation speeds of DV201 and DV301. Additionally, the influence of the length (proportional to the number of equipment in series) of DV201 will also be analyzed.

With the goal of understanding the impact on the process as a whole, KPIs from a technical, economic and environmental point of view will be selected. From the technical side, the process yield (PLA production flowrate divided by the lactic acid feed flowrate) will be analyzed. As the product flowrate is fixed by a design spec, in practice, variations in the process yield describe variations in the lactic acid feed flowrate. From the economic side, as a complete techno-economic analysis has not yet been implemented, the operational costs can be related to the necessary feed flowrate (already analysed by the process yield) and the total utility consumption, which will be used as a KPI. Regarding capital costs, its variations in these sensitivity analyses will only be due to changes in the necessary DV301 length necessary to obtain the desired PLA purity. For this reason, the DV301 length will be another KPI tracked. Finally, to assess the variations in environmental impact, the changes in the climate change impact will be analyzed, although with the developed LCA coupling, the influence on any other environmental KPI could also be analyzed.

The KPI values obtained for the base case operating conditions will first be presented, followed by the qualitative influence of each operating variable on the KPIs. Then, the main variables impacting each KPI will be discussed in more detail.

Base case results

At the base case, corresponding to the operating

Table 2: Influence of different operational variables ("T" – temperature, "P" – pressure, "$\dot{\omega}$" – blade rotation speed) on technical and environmental KPI's. Squares filled in green (or with "++") indicate strong influence of the variable on the KPI, yellow (or "+") indicate some or weak influence, red (or "0") indicate irrelevant or no influence and grey indicate variables that can't be changed or for which the simulation doesn't converge.

	Process yield			Total utility consumption		
	T	P	$\dot{\omega}$	T	P	$\dot{\omega}$
R101		+			+	
DV201	+	+	+	+	++	+
R301	0	0		+	0	
DV301	0	0	0	+	+	+
	DV301 length			**Climate change impact**		
	T	P	$\dot{\omega}$	T	P	$\dot{\omega}$
R101		++			+	
DV201	0	++	+	+	+	0
R301	+	0		0	0	
DV301	++	++	++	0	+	0

conditions mentioned in the process scale section, we obtain the following KPI values, presented in Table 1.

Table 1: KPI values obtained for the base case scenario.

KPI	Value
Process yield	64.52 %
Total utility consumption	13.91 MW
DV301 length	84.12 m
Climate change impact	6.869 kg CO_2 eq/kg PLA

The normalized utility consumption is equal to 5.82 MJ/kg PLA. Regarding the DV301 necessary length, it is considered that the devolatilizers have a maximum length of 6 meters and that all devolatilizers are placed in series (in order for the residence time to be inside the recommended operating range). This implies that 84.12 m is equivalent to having at least 14 devolatilizers in series. The number of devolatilizers, in both the case of the DV301 and the DV201, seems to be excessively high. This might be due to low diffusivity values calculated for the volatile components in the liquid mixture, which, for now, are calculated using the Stokes-Einstein equation. Future efforts will be made to obtain more accurate diffusivity values for this system, which will likely decrease the number of devolatilizer equipments necessary. Nonetheless, it has been noted that once the process is optimized, the number of necessary devolatilizers will be decreased significantly.

Impact on process yield

In Table 2, the impact of the different operating variables on the KPI's mentioned can be observed. Regarding the process yield, the most influential operational variables are related to the R101 and DV201 reactors, which is natural as these are the equipment
more upstream of the process, having more influence on the overall process than the more downstream equipment. For instance, increasing the operating temperature of DV201 to 175ºC increases the process yield to 64.96%, as it promotes a more effective product evaporation from the reactional mixture. The increase in the process yield is equivalent to the decrease in the necessary feed flowrate, which drops from 13340 to 13252 kg/hr.

Impact on total utility consumption

When it comes to total utility consumption, more operational variables have an impact. The most notable impact is noted with the DV-201 pressure. With lower pressure, such as 2 torr, the total consumption decreases to 13.24 MW, while at higher pressures, such as 7.5 torr, the total consumption increases to 15.25 MW. These relations are likely described by the increase and decrease in the process efficiency, which overcome the increase and decrease in vacuum costs, respectively, as lower pressures favor a more efficient lactide removal.

Impact on necessary DV301 length

Regarding the impact on the necessary DV301

length, which translates into the number of devolatilizers in series and is naturally related to the necessary capital investment cost, it is intuitive that the variables related to the same equipment have a more significant impact. In Figure 4, it can be seen that at higher DV301 inlet temperature, the necessary length decreases significantly. For a temperature of 187ºC, the necessary length is less than 60 m, corresponding to 10 devolatizers. In the same way, in Figure 5, it is seen that higher rotation speeds lead to lower lengths, with rotations of 250 rpm leading to a length of 74 m. Both relations are expected as the increase in both variables favor a more efficient volatile evaporation from the polymeric mixture, as increased rotations increase the mass transfer coefficient of the volatile components.

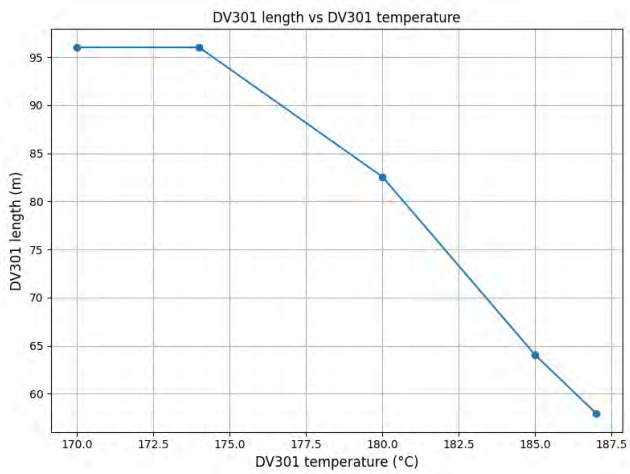

Figure 4. Impact of the DV301 temperature on the necessary DV301 length to perform the desired purification of the final PLA product.

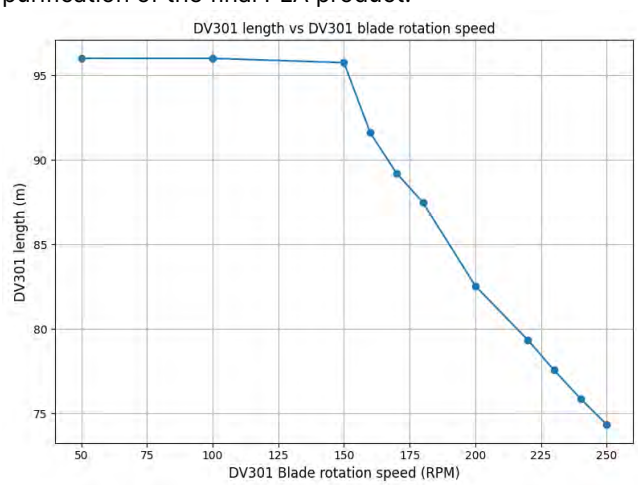

Figure 5. Impact of the DV301 blade rotation speed on the necessary DV301 length to perform the desired purification of the final PLA product.

Impact on climate change impact

In the case of climate change impact, in general, the more impactful variables are the ones that impact the feed flowrate and the energy consumption, as these are the main contributors to the environmental impact. For instance, the DV201 pressure has an influence on both the feed flowrate and the energy consumption, so naturally, as can be seen in Figure 6, it also has an influence on climate change impact. The influence is, however, not very significant, with the climate change impact decreasing to around 6.81 kg CO_2/kg PLA at 2 torr and increasing to 6.95 kg CO_2/kg PLA at 7.5 torr.

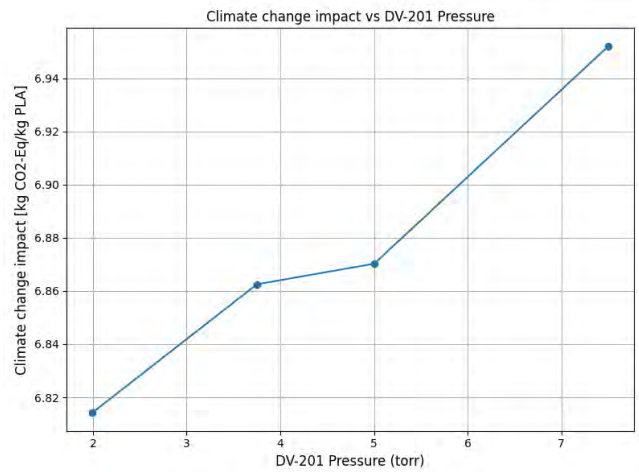

Figure 6. Impact of the DV201 pressure on the overall climate change impact.

In the case of the devolatilizers DV201 and DV301, variables such as temperature, pressure and blade rotation speed are all reactor-scale specific, as their impact couldn't be captured (at least not fully) without the custom devolatizer models developed. Thus, in Figure 6, we see the impact of a reactor scale variable directly on the environmental impact and in Figures 4 and 5 the impact, indirectly, on the capital investment cost. This is only possible with the coupling of the reactor, process, and environmental analysis scales as presented in this work.

DV201 length influence

In addition to the potential decision variables presented in Table 2, the impact of the length of DV201 (which in practice can be translated into devolatilizers in series) was analyzed. In the base case scenario, a length of 90 meters is used, which translates to 15 devolatilizers in series in two parallel production lines (in order to be inside the recommended residence time range), making a total of 30 devolatilizers. For the sensitivity analysis, the length was varied between 60 and 114 meters (20 and 38 devolatizers). The impact on the different KPI's can be seen in Table 3.

Table 3: Impact of the DV201 length on the different KPI's analyzed.

KPI	KPI Value for respective DV201 length	
	60 m	114 m
Process yield (%)	63.77	64.89
Total utility consumption (MW)	12.63	13.86
DV301 length (m)	96.00	83.19
Climate change impact (kg CO_2 eq/kg PLA)	6.875	6.831

As seen in Table 3, the DV201 length has an impact on all KPI's. Higher DV201 lengths favor lead to higher conversions, leading to more favorable KPI values. However, higher DV201 lengths also lead to higher capital cost investment and higher total utility consumption. This leads to an interesting trade-off between different KPI's, which could be explored in more detail with a future optimization framework.

The variables with the most significant impacts on the KPI's are the ones that should be chosen to be decision variables in an optimization framework. Having in account the analyses done, we conclude that the most important decision variables to include in an optimization framework are the following ones (in descending order): DV201 pressure, R101 pressure, DV201 temperature and blade rotation speed, DV301 pressure, temperature, and blade rotation speed, and finally, R301 temperature. Once the convergence of the process simulation is improved, the influence of the R101 temperature on the KPI's should be analyzed in more detail.

CONCLUSIONS AND FUTURE WORK

The first outcome of this work is how it shows the possibility of integrating models of different scales in a seamless manner and only using open-access tools. With the Python Unit Operation tool presented in this work, any complex unit operation at a smaller scale can be modeled in Python and coupled to a wide variety of process simulators (including Aspen Plus). Using a Python interface and the *olca* library, LCA analysis using OpenLCA can be performed, as long as the process simulator can be linked to a Python interface, which is the case with Aspen Plus.

With the sensitivity analysis performed, it was concluded that the variables related to DV201 and the temperature of R101 have a significant influence on practically all the KPIs considered, being the most relevant variables to be used in a future optimization algorithm as decision variables. The exception is the necessary DV301 length, for which the variables related to the DV301 devolatilizer have the largest influence. The crucial importance of both devolatilizers in the process highlights the paramount importance of having developed and coupled the detailed devolatilizer model, in order to understand the full influence of the devolatilizers on all KPI's with precision. With this coupling of the models, we were able to solve the common problem of assessing the impact of decision variables at smaller scales on industrially relevant environmental and, eventually, economic KPIs in an effortless way.

Future work will address, firstly, on improving the robustness and stability of the multi-scale model, so that the influence of all operational variables can be assessed at their full operational range. Additionally, future work will focus on coupling a complete techno-economic analysis tool in the same Python interface already used for the LCA analysis, so that the process can be accurately assessed not only environmentally, but also economically. Finally, the optimization algorithm can be developed in the Python interface, creating a multi-scale optimization framework, with which variables at all scales, potentially even catalyst-scale, can be optimized to improve the industrial economic and environmental performance of PLA production.

ACKNOWLEDGEMENTS

Centro de Química Estrutural is a Research Unit funded by Fundação para a Ciência e a Tecnologia through projects UIDB/00100/2020 (https://doi.org/10.54499/UIDB/00100/2020) and UIDP/00100/2020. Institute of Molecular Sciences is an Associate Laboratory funded by Fundação para a Ciência e a Tecnologia through project LA/P/0056/2020. Thomas Hietala PhD scholarship is funded by FCT (2023.05598.BD).

AUTHOR IDENTIFIERS

Author ORCIDs:
Hietala T: 0009-0009-3393-417X
Herres-Pawlis S: 0000-0002-4354-4353
Mendes S.F. P: 0000-0002-2768-353X

REFERENCES

1. European Bioplastics, https://www.european-bioplastics.org/bioplastics-market-development-update-2025/ visited on 02/02/2026
2. McNeeley A, Liu YA. Assessment of polylactic acid manufacturing from lactic acid for circular economy: process modeling and sustainable design. Ind. Eng. Chem. Res. 65:1679-1699 (2026). https://doi.org/10.1021/acs.iecr.5c04623
3. Conen N, Fuchs M, Hoffmann A, Herres?Pawlis S, Jupke A. Taking the next step—model?based analysis of robust and non?toxic zn catalysts for

the ring opening polymerization of lactide in the polymer melt. Advanced Sustainable Systems 7: (2022). https://doi.org/10.1002/adsu.202200359
4. CO-LaN, https://www.colan.org/process-modeling-component/python-cape-open-unit-operation/ visited on 02/02/2026
5. Vink ETH, Davies S, Kolstad JJ. ORIGINAL RESEARCH: the eco-profile for current ingeo ® polylactide production. Industrial Biotechnology 6:212-224 (2010). https://doi.org/10.1089/ind.2010.6.212
6. Lee BW, Yin K, Splaine K, Roesch B. Thin-film evaporator model for continuous active pharmaceutical ingredient manufacturing. Ind. Eng. Chem. Res. 59:3252-3260 (2020). https://doi.org/10.1021/acs.iecr.9b03974
7. Seavey KC, Liu YA. Step?growth polymerization process modeling and product design. Wiley (2009). https://doi.org/10.1002/9780470292488
8. Lopez-Toledo, Heat and Mass Transfer Characteristics of a Wiped Film Evaporator, University of Texas, Ph.D. dissertation (2006)
9. Witzke D.R., Introduction to Properties, Engineering, and Prospects of Polylactide Polymers, Michigan State University, Ph.D. dissertation (1997)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Designing $MgCl_2$-Based Ethanol Dehydration Systems: A Multi-Objective Approach with Open-Loop Controllability

Josué J. Herrera Velázquez[ac], J. Rafael Alcántara Avila[b], Salvador Hernández[a], and Julián Cabrera Ruiz[a]*

[a] Universidad de Guanajuato, Departamento de Ingeniería Química, Guanajuato, 36050, Mexico
[b] Pontificia Universidad Católica del Perú, Departamento de Ingeniería, Lima, 15074, Peru
[c] Instituto Tecnológico Superior de Guanajuato, Departamento de Ingeniería en Industrias Alimentarias, Guanajuato, 36262, Mexico
* Corresponding Author: j.cabreraruiz@ugto.mx.

ABSTRACT

Ethanol derived from biomass is a promising renewable fuel; however, its long-term use as a gasoline additive is becoming increasingly uncertain due to the rise of electric vehicles and alternative propulsion technologies. This trend motivates the exploration of higher-value applications for ethanol, particularly in the food and pharmaceutical sectors, where product safety is critical. A key challenge in ethanol purification is breaking the ethanol–water azeotrope, as conventional entrainers such as ethylene glycol or glycerol can leave residual traces that limit ethanol's use in sensitive markets. Magnesium chloride ($MgCl_2$) offers an effective alternative, enabling high-purity ethanol without introducing hazardous organic residues, while exhibiting favorable hygroscopic properties and operational reliability. Simulating this system is challenging due to strong non-ideal and electrolyte interactions in phase equilibrium. Conducting a rigorous controllability analysis is also difficult; therefore, within an Integrated Design and Control (IDC) framework, accurate approximations are essential for successful implementation. This work presents a multi-objective optimization of a $MgCl_2$-based ethanol dehydration column, simultaneously minimizing the Total Annual Cost (*TAC*) and the open-loop controllability criterion ($A_\gamma + \gamma_{sm}$). The optimization integrates Aspen Plus® with Python and employs validated surrogate cost models for the preconcentrator and salt concentrator processes. Dynamic boundaries ensure product purity by maintaining the ethanol/water feed below the salt feed. The *ASF* (Achievement Scalarizing Function) solution achieved a 64.33% improvement in controllability with a 13.37% increase in *TAC* relative to the reference case reported in the literature, illustrating the trade-off between economic and control objectives. This study demonstrates that incorporating simplified controllability metrics into multi-objective optimization enables the efficient and practical design of complex systems, providing a viable approach for managing cost–control trade-offs in processes with challenging dynamics.

Keywords: Magnesium Chloride, Ethanol Dehydration, Multi-objective Optimization, Surrogate Models, Aspen Plus - Python

INTRODUCTION

Ethanol produced from biomass has long been considered a promising renewable fuel, particularly as an additive in gasoline blends to reduce emissions and partially replace fossil fuels. However, the increasing adoption of electric and hydrogen mobility technologies, along with global decarbonization goals, may reduce the long-term demand for traditional liquid fuels, thereby affecting ethanol consumption [1]. Market studies indicate that ethanol demand is closely tied to gasoline and that, without supportive blending policies or mandates (such as E10, E15, or E85), its use could decline despite ongoing technological advancements [2]. Nevertheless, the future of ethanol is not predetermined; its role will continue to be shaped by regulatory frameworks, and significant opportunities exist for diversification as a chemical feedstock and within integrated biorefinery platforms, opening pathways to new applications in a low-carbon economy [3]. This transition underscores the need to explore

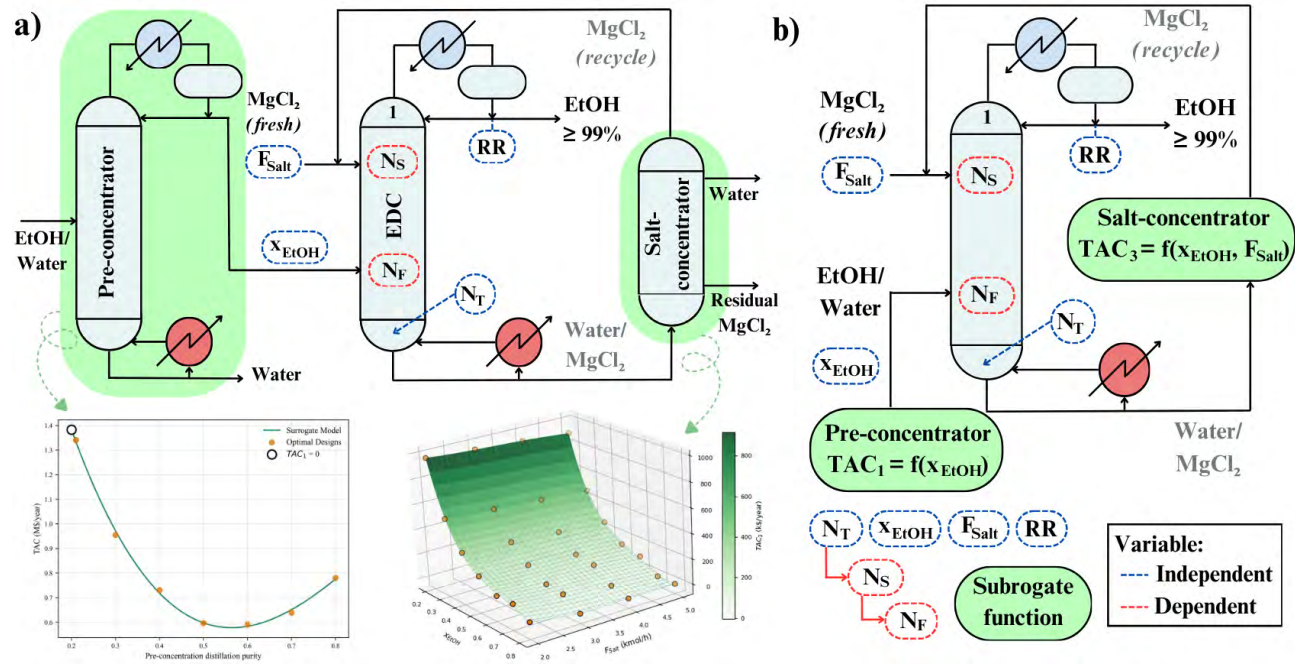

Figure 1: a) Actual extractive distillation process, including the mapping of the surrogate model variables used to derive Eqs. 2 and 3. b) Scheme to be optimized, in which the surrogate models replace the corresponding unit operations and variables for the optimization and their evaluation sequence.

alternative high-value uses of ethanol, particularly in the food and pharmaceutical industries, where stringent safety and quality requirements are paramount.

Ethanol with a purity exceeding 95 vol% is suitable for applications in the pharmaceutical, cosmetic, solvent, and chemical industries [4]. Conventional distillation is inherently limited to purities near the azeotropic composition, and achieving higher purities requires substantial energy input due to the reduced driving force as the azeotrope is approached. To address this limitation, several intensified separation strategies have been reported in the literature, including reactive distillation using ethylene oxide to remove water via ethylene glycol formation [5], and azeotropic distillation employing entrainers such as benzene [6]. Among these alternatives, extractive distillation has emerged as a well-established and widely implemented industrial process for producing anhydrous ethanol. This approach typically relies on high-boiling solvents (e.g. ethylene glycol and glycerol) to modify the relative volatility of the ethanol–water system, thereby enabling high-purity ethanol production [7]. However, the use of conventional extractive agents may result in residual traces in the final product, potentially limiting its suitability for sensitive applications such as the food and pharmaceutical industries.

In contrast, magnesium chloride ($MgCl_2$) has proven effective for producing high-purity ethanol without posing health risks, even at trace levels [8]. Additionally, $MgCl_2$ exhibits favorable hygroscopic properties and presents fewer operational challenges compared to alternative salts (*e.g.* $CaCl_2$) [9]. More recently, Herrera Velázquez et al. [10] optimized an $MgCl_2$-based extractive dehydration column by minimizing the Total Annual Cost (TAC), using surrogate models to represent the preconcentration and salt concentration processes. However, extractive distillation systems are inherently thermodynamically complex and involve more degrees of freedom than conventional distillation, which significantly complicates their control. In this context, previous studies have shown that improvements in controllability often conflict with economic performance metrics such as (TAC) [11], highlighting the need to explicitly incorporate control considerations during the process design stage.

Simulating these systems is challenging due to the strong non-ideal and electrolytic interactions present in phase equilibrium, which make dynamic-state simulations particularly difficult. In this context, obtaining controllability approximations from the steady-state conditions can be very useful for conducting such analyses. The Integrated Design and Control (IDC) propose designing the process and its controller simultaneously, thereby overcoming the traditional separation between process design and control [12]. This conventional separation can be suboptimal, as it often neglects the dynamic behavior of $u(t)$, where u is the dynamic model, and its controllability is unknown. Although controller design is not addressed in this work, understanding controllability behavior from the early design stages is fundamental.

Open-loop analysis provides a reliable approximation, since a well-designed process naturally facilitates control, whereas advanced control cannot fully compensate for an inadequate design.

Dynamic controllability analysis during stochastic optimization is computationally intensive because it involves nested optimization. As an alternative, open-loop controllability can be evaluated using first-order transfer function models and Singular Value Decomposition (SVD). The condition number γ^* is evaluated across a frequency range (10^2 rad/h), and system performance is quantified using the parameter $A_\gamma + \gamma_{sm}$, which represents the area under the condition number curve and penalizes noise amplification [13]. This approach is feasible because time constants can be estimated from steady-state data using the equations of Skogestad and Morari [14]. Employing approximate first-order models simplifies and generalizes the system analysis, allowing relevant controllability information to be obtained during the optimization process without requiring additional rigorous steps.

METHODOLOGY

Case Study

The extractive distillation scheme is shown in Figure 1. Figure 1a represents the real scheme, while Figure 1b shows the simplified scheme used in this work. The system feed is 500 $kmol/h$ with an ethanol ($EtOH$) concentration of $20\,\%$ mol. The feed to the extractive distillation column (EDC), however, depends on the value of x_{EtOH}, as the study evaluates whether water is removed in the preconcentrator or extracted in the EDC. The EDC is simulated in Aspen Plus® V12 using the APV140VLE-IG databank to estimate the dissociation reaction parameters (Eq. 1) required by the E-NRTL activity coefficient model. The column operates at a pressure of $1.2\,atm$, with a stage pressure drop of $0.05\,kPa$.

$$MgCl_2 \rightarrow Mg^{++} + 2Cl^- \tag{1}$$

The E-NRTL model parameters were fitted using equilibrium data reported by Zeng [8] and are summarized in Table 1.

Table 1: Electrolyte pair parameters.

Ethanol (1) - Water (2)	α_{m_1,m_2}	τ_{m_1,m_2}	τ_{m_2,m_1}
	0.3000	−0.1893	1.8224
Ethanol (1) - $\boldsymbol{MgCl_2}$ (ca)	α_{ca,m_1}	$\tau_{m_1,ca}$	τ_{ca,m_1}
	0.0600	24.5470	−12.5320
Water (2) - $\boldsymbol{MgCl_2}$ (ca)	α_{ca,m_2}	$\tau_{m_2,ca}$	τ_{ca,m_2}
	0.2000	7.8720	−4.1470

For the preconcentration and salt concentration process, surrogate models are used to provide the data required for the economic analysis. These units are represented with surrogate models because they correspond to well-established and optimized process configurations under the operating conditions considered, while also significantly reducing the computational complexity of the overall optimization problem. The surrogate model for the preconcentrator is defined in Eq. 2 and depends on the ethanol purity, as its primary purpose is to deliver the feed to the EDC at the specified ethanol composition, x_{EtOH}.

$$\begin{aligned} TAC_1[M\$/year] \approx f(x_{EtOH}) &= -4.1959x_{EtOH}^3 \\ &+ 11.6440x_{EtOH}^2 \\ &- 9.1315x_{EtOH} + 2.7769 \\ 0.2 < x_{EtOH} \leq 0.8 \end{aligned} \tag{2}$$

Eq. 2 is defined over the range $0.2 < x_{EtOH} \leq 0.8$, with the additional constraint that when $x_{EtOH} = 0.2$, the associated cost $TAC_1 = 0$. Under these conditions, a preconcentration unit is not required, since the ethanol composition already matches that of the feed to the separation system. This lower limit is based on the typical ethanol concentration found in the water/ethanol mixture at the outlet of a biomass fermentation reactor [16, 17].

Consequently, the maximum ethanol purity that can be fed to the extractive distillation column is $x_{EtOH} = 0.8$, which remains below the azeotropic composition. Eq. 2 depends solely on the purity of the $EtOH$, as it is a conventional distillation scheme. The surrogate model is adjusted based on optimized schemes for discretized ethanol purities (Figure 1a). Therefore, for any given $EtOH$ purity used to evaluate the preconcentrator cost, the model ensures that the design corresponds to the optimal configuration for that specific x_{EtOH} [10]. Eq. 3 presents the surrogate model for the salt concentrator located downstream of the EDC.

$$\begin{aligned} TAC_3[k\$/year] \approx f(x_{EtOH}, F_{Salt}) = \\ -8236x_{EtOH} + 534F_{Salt} + 9043x_{EtOH}^2 \\ -2374x_{EtOH}F_{Salt} + 9043F_{Salt}^2 + 71.53x_{EtOH}^3 \\ +7490x_{EtOH}^2F_{Salt} + 144.5x_{EtOH}F_{Salt}^2 \\ +77.52F_{Salt}^3 - 3382x_{EtOH}^4 - 9382x_{EtOH}^3F_{Salt} \\ -356.5x_{EtOH}^2F_{Salt}^2 + 1.927x_{EtOH}F_{Salt}^3 \\ -11.97F_{Salt}^4 + 4033x_{EtOH}^4F_{Salt} + 216x_{EtOH}^3F_{Salt}^2 \\ +8.697x_{EtOH}^2F_{Salt}^3 - 0.831x_{EtOH}F_{Salt}^4 \\ + 0.7047F_{Salt}^5 + 2158 \\ 0.2 \leq x_{EtOH} \leq 0.8;\ 2 < F_{Salt}[kmol/h] \leq 5 \end{aligned} \tag{3}$$

This model is a function of the x_{EtOH}, and the fresh salt flow rate, F_{Salt}. The dependence on x_{EtOH} arises because the amount of water to be removed is determined by the $EtOH$ concentration entering the system. Additionally, the salt flow rate directly affects both the

thermodynamic behavior of the unit and the required water removal, since a certain amount of water is needed to dissolve the solid salt (approximately $2\,\%\,mol$ relative to the total salt mixture). Both surrogate models are taken from Herrera Velázquez et al. [10]. The optimization variables for this problem are the ethanol fraction at the outlet of the preconcentrator and the inlet of the EDC (x_{EtOH}), the salt flow rate F_{Salt}, the total number of stages (N_T), the salt feed stage (N_S), the ethanol-water feed stage (N_F), and the reflux ratio (RR). The variables will be solved sequentially as shown in Figure 1b.

Mathematical Treatment

The problem is formulated as a Multi-Objective Optimization (MOO) problem where the objective functions are TAC and $A_\gamma + \gamma_{sm}$, as shown in Eq. 4.

$$\min(TAC, A_\gamma + \gamma_{sm}) = f(x_{EtOH}, F_{Salt}, N_T, N_s, N_F, RR)$$
$$s.t.\, Conv = 0;\; P_{EtOH} \geq 99\,\%\,mol \tag{4}$$

$Conv$ is a convergence constraint; if it is not met, a penalty value is assigned to each objective function. The distillate flow rate is set at $100\ kmol/h$, and the purity constraint is set to be greater than $99\,\%\,mol$ for the ethanol (P_{EtOH}); if this is not met, a penalty value is assigned. The optimization bounds are presented in Table 2.

Table 2: Optimization Bounds.

Variables	***Boundaries*** $[\boldsymbol{lower}, \boldsymbol{upper}]$
x_{EtOH}	$[0.2, 0.8]$
$F_{Salt}\ [kmol/h]$	$[2, 5]$
N_T	$[15, 50]$
N_S	$[2, (2/3)N_T]$
N_F	$[N_S + 2, N_T - 1]$
RR	$[0.2, 10]$

To ensure product purity, dynamic boundaries were applied to the feed stage positions, guaranteeing that the ethanol/water mixture is always fed below the salt. The variables were normalized following the method reported by Herrera Velázquez et al. [10] using Eq. 5.

$$VN_i = (B_{upper} - V_i)/(B_{upper} - B_{lower}) \tag{5}$$

Where VN_i is the normalized value of the variable V_i, and B_{upper} and B_{lower} are the upper and the lower bounds of variable i, respectively. The TAC_2 associated to the EDC is calculated as indicated in Eq. 6, and the overall TAC is calculated as shown in Eq. 7.

$$TAC_2\left[\frac{\$}{year}\right] = C_{Vapor} + C_{Cooling} + \frac{\left(\frac{CEPCI_{2022}}{CEPCI_{2001}}\right)(C_{Shell} + C_{Plates} + C_{Reboiler} + C_{Condenser})}{Recovery\ time\ (Years)} \tag{6}$$

$$TAC\left[\frac{\$}{year}\right] = (1x10^6)TAC_1 + TAC_2 + (1{,}000)TAC_3 \tag{7}$$

With respect to open-loop controllability, a sensitivity analysis indicates that the considered disturbances have a negligible effect on the gains of both the preconcentrator and the salt concentrator. Consequently, these units are excluded from the controllability objective function, which is therefore applied exclusively to the EDC.

For the calculation of $A_\gamma + \gamma_{sm}$, the transfer function matrix $G_j(s)$ is calculated as indicated in Eq. 8. $P_{EtOH,j}$ and $P_{W,j}$ (Water) represent the purities achieved for the design at iteration j. P_{EtOH,RR_j} and P_{W,RR_j} correspond to the purities obtained by applying a positive perturbation of 0.5% to the RR. Similarly, $P_{EtOH,Q_{H,j}}$ and $P_{W,Q_{H,j}}$ represent the purities obtained by applying a positive perturbation of 0.5% to the heat duty (Q_H). Each disturbance is evaluated through independent simulations, resulting in a total of three simulations per iteration j: (1) to verify that the design satisfies the constraints, (2) to apply a perturbation to RR and (3) to apply a perturbation to Q_H.

$$G_j(s) = \begin{bmatrix} \dfrac{\left(\frac{P_{EtOH,RR_j} - P_{EtOH,j}}{0.0025}\right)}{1 + \tau_{1,1}\, s} & \dfrac{\left(\frac{P_{EtOH,Q_{H,j}} - P_{EtOH,j}}{0.0025}\right)}{1 + \tau_{1,2}\, s} \\ \dfrac{\left(\frac{P_{W,RR_j} - P_{W,j}}{0.0025}\right)}{1 + \tau_{2,1}\, s} & \dfrac{\left(\frac{P_{W,Q_{H,j}} - P_{W,j}}{0.0025}\right)}{1 + \tau_{2,2}\, s} \end{bmatrix} \tag{8}$$

For τ, they are calculated as indicated in Eq. (9-12).

$$\tau_{1,1} = \frac{\sum_{k=2}^{N_T-1} H_{L,k}\left(x_{EtOH,RR_j,k} - x_{EtOH,k}\right)}{F_{D,RR_j}\left(P_{EtOH,RR_j} - P_{EtOH,j}\right) + F_{B,RR_j}\left(P_{W,RR_j} - P_{W,j}\right)} \tag{9}$$

$$\tau_{2,1} = \frac{\sum_{k=2}^{N_T-1} H_{L,k}\left(x_{W,RR_j,k} - x_{W,k}\right)}{F_{D,RR_j}\left(P_{EtOH,RR_j} - P_{EtOH,j}\right) + F_{B,RR_j}\left(P_{W,RR_j} - P_{W,j}\right)} \tag{10}$$

$$\tau_{1,2} = \frac{\sum_{k=2}^{N_T-1} H_{L,k}\left(x_{EtOH,Q_{H,j},k} - x_{EtOH,k}\right)}{F_{D,Q_{H,j}}\left(P_{EtOH,Q_{H,j}} - P_{EtOH,j}\right) + F_{B,Q_{H,j}}\left(P_{W,Q_{H,j}} - P_{W,j}\right)} \tag{11}$$

$$\tau_{2,2} = \frac{\sum_{k=2}^{N_T-1} H_{L,k}\left(x_{W,Q_{H,j},k} - x_{W,k}\right)}{F_{D,Q_{H,j}}\left(P_{EtOH,Q_{H,j}} - P_{EtOH,j}\right) + F_{B,Q_{H,j}}\left(P_{W,Q_{H,j}} - P_{W,j}\right)} \tag{12}$$

Where $x_{EtOH,k}$ and $x_{W,k}$ are the mole fractions of ethanol and water, respectively, at stage k. F_{D,RR_j} and F_{B,RR_j} are the distillate and bottom flow rates, respectively, for iteration j under the perturbation of RR. Similarly, $F_{D,Q_{H,j}}$ and $F_{B,Q_{H,j}}$ are the distillate and bottom flow rates for iteration j under the perturbation of Q_H. Finally, $H_{L,K}$ represents the liquid holdup at stage k. The $G_j(s)$ from Eq. (8) is then converted to the frequency domain ($s = \omega j$) and

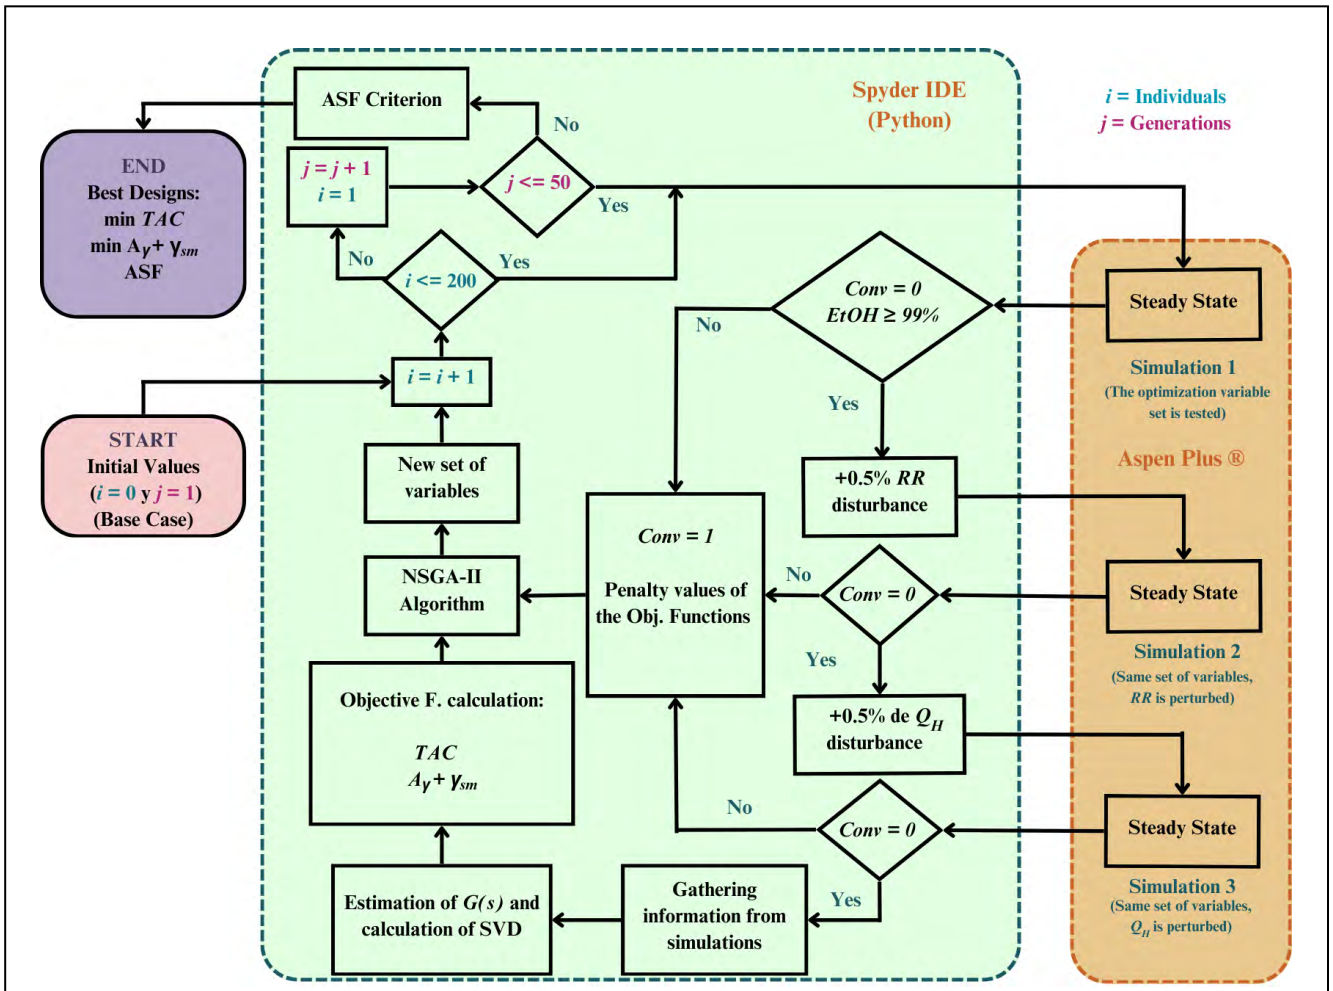

Figure 2. Flowchart of the mathematical model optimized for EDC with the NSGA-II algorithm.

the DVS of its γ^*, are evaluated over the frequency range from 10^{-4} rad/h to 10^{2} rad/h. Numerical integration is subsequently performed using Eq. 13 and the trapezoidal method. This equation is discussed in detail in Herrera Velázquez et al. [15].

$$A_\gamma + \gamma_{sm} = \frac{\int_{\log a}^{\log b} \log(\gamma^*)\, d(\log \omega)}{\log b - \log a} + \int_{\log a}^{\log b} \Delta log(\omega)\, |\gamma'_{i+1} - \gamma'_i|\, d(\log \omega)$$
$$\gamma'_{i+1} - \gamma'_i = \frac{log[\gamma^*(\omega_{i+1})] - log[\gamma^*(\omega_i)]}{log(\omega_{i+1}) - log(\omega_i)} - \frac{log[\gamma^*(\omega_i)] - log[\gamma^*(\omega_{i-1})]}{log(\omega_i) - log(\omega_{i-1})} \quad (13)$$

Where b and a are the upper and lower limits of the evaluated frequency range, respectively. The model was optimized using the NSGA-II algorithm [18], with 50 generations of 200 individuals each. Figure 2 shows the complete optimization process. The algorithm was implemented in Spyder IDE 5.2.2 and executed on an Intel® Xeon® W-3235 CPU at 3.30 GHz with 48 GB of RAM.

RESULTS

The execution time for the optimization was 9.79 hours. A total of $4,359$ designs $(43.59\ \%)$ satisfy the purity and convergence constraints. The summarized results are shown in Table 3. The criterion $A_\gamma + \gamma_{sm}$ is not reported for the reference case because, under the methodology described, the perturbation applied to the RR does not converge $(Conv = 1)$. Figure 3 shows the Pareto front obtained after the optimization process, highlighting the most relevant solutions $(\min TAC, \min A_\gamma + \gamma_{sm}, ASF)$. Figure 4 presents the distribution of partial costs $(TAC_1, TAC_2\ and\ TAC_3)$ within the overall TAC for all solutions, as well as for the reference design.

Figure 5 shows the values of condition number (γ^*) and the minimum singular values (σ_*) derived from the control estimates obtained using Eq. (8–12). Due to the difference in orders of magnitude, the trajectories of γ^* cannot be clearly observed; however, σ_* enables the identification of relevant differences in this analysis. The behavior of σ_* is consistent with expectations for each

solution, with the $\min TAC$ solution exhibiting the poorest performance. This trade-off highlights the inherently multi-objective nature of the problem.

Table 3: Summary of results of MOO.

$\boldsymbol{Variables}$	$\boldsymbol{Ref.}$ $\boldsymbol{Des.}$ [10]	$\min \boldsymbol{TAC}$	$\min \boldsymbol{A_\gamma} + \boldsymbol{\gamma_{sm}}$	$\boldsymbol{ASF}$
x_{EtOH}	0.598	0.700	0.796	0.699
F_{Salt} $[kmol/h]$	4.940	3.699	3.658	3.645
N_T	30	21	22	27
N_S	5	10	6	15
N_F	10	18	21	25
RR	0.370	0.424	1.475	0.832
P_{EtOH}	0.993	0.991	0.994	0.995
TAC $[\$/year]$	$1,119,007$	$1,161,517$	$1,534,817$	$1,279,075$
$A_\gamma + \gamma_{sm}$	–	0.999	0.033	0.315

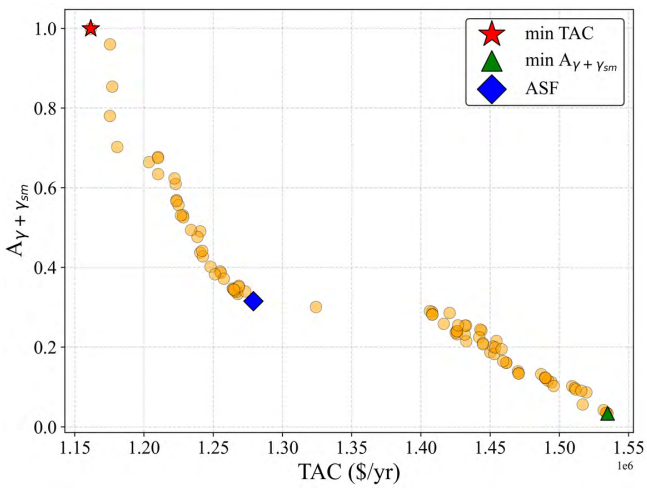

Figure 3. Pareto front solution from optimization.

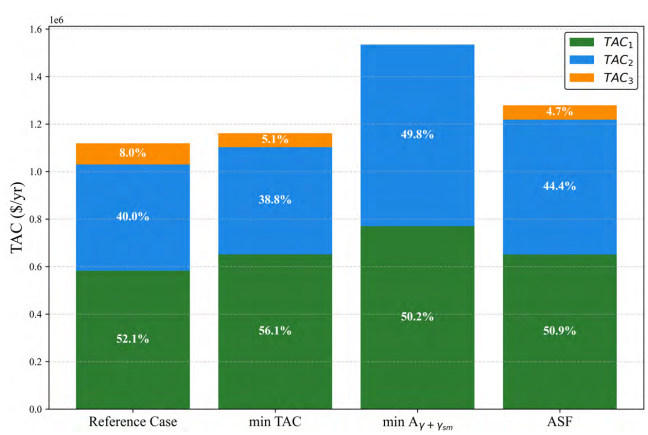

Figure 4. Distribution of partial costs within the TAC .

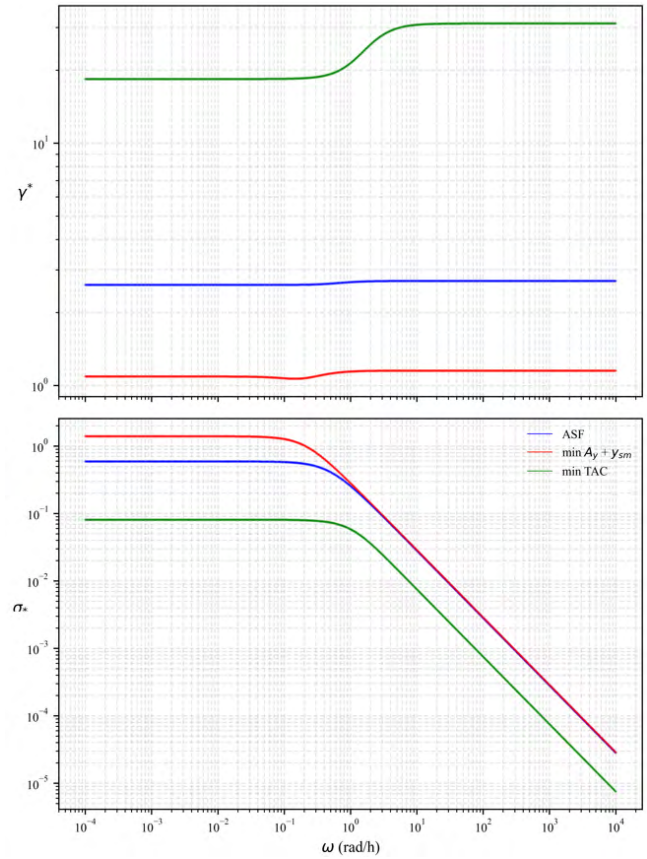

Figure 5. γ^* (above) and σ_* (below) of rigorous positive responses for each solution (EDC).

CONCLUSIONS

Compared to the reference design, the ASF solution (close to the utopian point) shows a 13.37% increase in TAC. However, from a controllability perspective, the reference case exhibits high sensitivity to perturbations in the reflux ratio (RR), as evidence by convergence difficulties in the simulation, suggesting unstable dynamic behavior. This result confirms the inherent trade-off between cost and controllability in the problem. A total of 89 non-dominated solutions were identified, forming the Pareto front and reflecting the multi-objective nature of the system. Most of these solutions, except for the $\min A_\gamma + \gamma_{sm}$ solution, use an EDC feed consisting of an ethanol-water mixture with 70% mol ethanol. This finding differs from what is commonly reported in the literature, where compositions close to the azeotrope are typically preferred. In this case study, however, feeding at a lower purity is more economical due the dominant contribution of the preconcentrator to the TAC, which amounts for more than 50% of the cost in all evaluated designs.

In contrast, the solution corresponding to the $\min A_\gamma + \gamma_{sm}$ proposes a feed with 80% mol ethanol, which coincides with the upper limit considered in the surrogate

model of the preconcetrator. This behavior can be explained by the relationship between feed composition and internal column flows, which directly influence the time constants used in controllability estimation. Higher ethanol purity reduces the amount of water entering the EDC, thereby modifying the system dynamics and improving controllability. For $\min A_{\gamma} + \gamma_{sm}$ solution, the salt concentrator has a negligible impact ($TAC_3 < 1,000\ \$/yr$), as salt recovery satisfies the water requirements for recirculation without incurring significant additional costs.

Finally, incorporating a controllability criterion based on simplified models within a MOO framework provides a practical approach for designing complex systems such as the $MgCl_2$-based EDC, particularly when full dynamic modeling is difficult to implement. Evaluating controllability using first-order transfer functions and open-loop SVD analysis yields quantitative indicators of dynamic behavior, enabling the identification of sensitive and robust configurations. In this study, an indirect relationship between controllability and TAC is observed, supporting a more comprehensive analysis of the trade-off between economic performance and operability. This approach enables the generation of Pareto fronts that facilitate the selection of robust and economically viable designs in systems with complex dynamics.

AUTHOR IDENTIFIERS

Author ORCIDs:
Herrera Velázquez JJ: 0000-0002-2975-360X
Alcántara Avila JR: 0000-0001-5974-2760
Hernández S: 0000-0002-4598-3392
Cabrera Ruiz J: 0000-0002-5144-5574

REFERENCES

1. Giehl A, Klanovicz N, Camargo AF, Albarello MLR, Treichel H, Alves SL Jr. Ethanol and electricity: fueling or fooling the future of road passenger transport?. Energy Nexus 12:100258 (2023). https://doi.org/10.1016/j.nexus.2023.100258
2. Curcio E. Accelerating Transportation Decarbonization: The Strategic Role of Ethanol Blends and Regulatory Incentives. arXiv (2025) https://doi.org/10.48550/arXiv.2504.06278
3. Susmozas A, Martín-Sampedro R, Ibarra D, Eugenio ME, Iglesias R, Manzanares P, Moreno AD. Process strategies for the transition of 1G to advanced bioethanol production. Processes 8:1310 (2020). https://doi.org/10.3390/pr8101310
4. Aslanzadeh S, Ishola MM, Richards T, Taherzadeh MJ. An overview of existing individual unit operations. Biorefineries :3-36 (2014). https://doi.org/10.1016/b978-0-444-59498-3.00001-4
5. Tavan Y, Hosseini SH. A novel integrated process to break the ethanol/water azeotrope using reactive distillation – part I: parametric study. Separation and Purification Technology 118:455-462 (2013). https://doi.org/10.1016/j.seppur.2013.07.036
6. Herrera Velázquez JJ, Cabrera Ruiz J, Alcántara Avila JR. Design of a Heterogeneous Azeotropic Distillation Processes Using Stochastic Optimization. *Comp Aided Chem Eng* 52:889-895 (2023) https://doi.org/10.1016/B978-0-443-15274-0.50142-6
7. García-Herreros P, Gómez JM, Gil ID, Rodríguez G. Optimization of the design and operation of an extractive distillation system for the production of fuel grade ethanol using glycerol as entrainer. Ind. Eng. Chem. Res. 50:3977-3985 (2011). https://doi.org/10.1021/ie101845j
8. Zeng L, Li Z. A new process for fuel ethanol dehydration based on modeling the phase equilibria of the anhydrous mgcl₂ + ethanol + water system. AIChE Journal 61:664-676 (2014). https://doi.org/10.1002/aic.14685
9. Llano-Restrepo M, Aguilar-Arias J. Modeling and simulation of saline extractive distillation columns for the production of absolute ethanol. Computers & Chemical Engineering 27:527-549 (2003). https://doi.org/10.1016/s0098-1354(02)00215-6
10. Herrera-Velázquez JJ, Cabrera-Ruiz J, Alcántara-Avila JR, Hernández-Hernández EJ, Perez-Gudiño MA, Luna-Perez R, Hernández S. Optimization of the bioethanol dehydration process using mgcl2 as mass separating agent. Chemical Engineering and Processing - Process Intensification 201:109809 (2024). https://doi.org/10.1016/j.cep.2024.109809
11. Alcántara-Avila JR, Cabrera-Ruiz J, Lee HY. Open-loop analysis for the stability of reactive distillation columns with intermediate condensers for the production of silane. IFAC-PapersOnLine 55:632-637 (2022). https://doi.org/10.1016/j.ifacol.2022.07.514
12. Vega P, Lamanna R, Revollar S, Francisco M. Integrated design and control of chemical processes – part II: an illustrative example. Computers & Chemical Engineering 71:618-635 (2014). https://doi.org/10.1016/j.compchemeng.2014.09.019
13. Cabrera-Ruiz J, Ramírez-Márquez C, Hasebe S, Hernández S, Alcántara Avila JR. Outlook of the dynamic behavior of closed-loop control through open-loop analysis for intensified separation processes. Ind. Eng. Chem. Res. 57:16795-16808 (2018). https://doi.org/10.1021/acs.iecr.8b04164
14. Skogestad S, Morari M. The dominant time

constant for distillation columns. Computers & Chemical Engineering 11:607-617 (1987). https://doi.org/10.1016/0098-1354(87)87006-0
15. Vel?zquez JJH, Avila JRA, Hern?ndez S, Ruiz JC. Evaluation of the controllability of distillation with multiple reactive stages. Systems and Control Transactions 4:1095-1100 (2025). https://doi.org/10.69997/sct.120398
16. Vane LM, Alvarez FR, Rosenblum L, Govindaswamy S. Efficient ethanol recovery from yeast fermentation broth with integrated distillation–membrane process. Ind. Eng. Chem. Res. 52:1033-1041 (2012). https://doi.org/10.1021/ie2024917
17. Tse TJ, Wiens DJ, Reaney MJT. Production of bioethanol—a review of factors affecting ethanol yield. Fermentation 7:268 (2021). https://doi.org/10.3390/fermentation7040268
18. Blank J, Deb K. Pymoo: multi-objective optimization in python. IEEE Access 8:89497-89509 (2020). https://doi.org/10.1109/access.2020.2990567

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Addressing Matrix Effects Through A Physical Prior-Informed Calibration Model For Quantitative Analysis

Onur C. Boy[a], Ulderico Di Caprio[a], Idelfonso Nogueira[b] and M. Enis Leblebici[a]*

[a] Center for Industrial Process Technology, Department of Chemical Engineering, KU Leuven, Agoralaan Building B, Diepenbeek, Belgium
[b] Department of Chemical Engineering, Norwegian University of Science and Technology, Gløshaugen, Trondheim, Norway
* Corresponding Author: muminenis.leblebici@kuleuven.be

ABSTRACT

Building a robust calibration curve is essential for accurate quantification of multicomponent mixtures. Matrix effects can distort the proportional relationship between the analyte concentration and instrumental response. In addition, classical machine learning models do not inherently incorporate simple physical constraints, such as the requirement that a zero response must correspond to a zero concentration, which can result in non-zero predictions. To address this limitations, prior-induced calibration models were proposed that inherently embed this physical constraint into the model architecture. A dataset was generated for ethanol electrooxidation products using headspace gas-chromatography-mass spectrometry (HS-GC-MS). Multiple linear regression (MLR), polynomial regression and artificial neural network (ANN) models were trained to investigate the effects of model complexity and the incorporation of physical information on predictive performance. Model selection and complexity were assessed using Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). The prior-induced polynomial regression model achieved improved predictive accuracy while maintaining low model complexity, whereas the ANN provided comparable accuracy but was strongly penalized due to its substantially higher complexity.

Keywords: Matrix effects, Calibration curve, Polynomial regression, Artificial neural network, Ethanol electrooxidation, Chemometrics

INTRODUCTION

Accurate quantification of individual species in multicomponent mixtures is essential for reliable chemical analysis, process monitoring and optimization of operating parameters on selectivity and conversion, which is critical for both batch studies and reactor design. Matrix effects and competitive partitioning among analytes complicate accurate quantification by altering the linear relationship between the analyte concentration and the response [1]. In such systems, the measured peak area is not solely proportional to the analyte concentration but is also influenced by the concentrations of the co-existing components, since each composition defines a new vapor–liquid equilibrium governed by non-ideal intermolecular interactions. To overcome this, an appropriate calibration model is required to accurately resolve the concentration of each component [2]. A calibration curve is a regression model between the known concentrations of analytes and its response obtained from an analytical instrument, which allows the prediction of the concentration of an unknown mixture [3].

Matrix effects have been addressed using calibration models, including linear and non-linear models such as MLR, principal component analysis, partial least squares (PLS), and ANNs [4]. Veiga-del-Baño et al. built LC–MS/MS and GC–MS/MS calibration models using linear, weighted linear, and polynomial regression, and reported that the weighted linear model had the best performance compared the other models [5]. Also, to identify adequate calibration models, Chen et al. investigated six different linear and non-linear calibration equations and reported that most real calibration curves exhibit measurable non-linearity, with non-linear models often outperforming linear models in terms of fitting and prediction [6]. Matrix effects were addressed by developing a calibration curve using PLS model to quantify volatile compounds in wine, where single internal standard

calibration failed to correct matrix effects [7]. Campmajó et al. used PLS to quantify peanut and hazelnut adulteration in almond products with prediction errors below 6.1% and no observable matrix effects [8]. Brusamarello et al. compared PLS and ANN models on UV–Vis spectra for soybean biodiesel quantification, finding that the ANN model achieved higher predictive accuracy for complex reaction mixtures [9]. Gholivand et al. addressed strong serum matrix effects in voltammetric analysis by calibrating PLS models in blank serum and applying baseline and alignment corrections, allowing reliable simultaneous determination of five electroactive alkaloids [10].

Classical machine learning algorithms cannot inherently incorporate simple physical constraints, such as the condition that in the absence of a given species, the corresponding peak area should be zero. As a result, such models may produce non-zero predictions even when the input peak area is zero. This issue can be addressed by applying a rule-based correction layer that forces the output to zero whenever the input value is zero or by manipulating the dataset by augmenting the training dataset with zero-concentration samples, however, this could introduce bias into the model. Alternatively, the model can be designed to inherently learn this constraint. In this case, the model cannot violate the physical boundary condition, as the physical information is enforced by design. Moreover, embedding this constraint reduces the effective model complexity and improves robustness near the detection limit.

This work aims to develop an data-driven calibration model for quantification of ethanol oxidation products while capturing the nonlinear relationships arising from matrix effects. The model also incorporates a simple physical prior: the principle that a zero peak area corresponds to zero concentration to increase prediction accuracy further. There are two models proposed: Prior-informed polynomial regression and prior-informed branched ANN model. The predictive performances of these models are compared with those of MLR, polynomial regression, and feed forward ANN. The model complexities are assessed using AIC and BIC.

METHODOLOGY

Chemicals

Ethanol (>99%) was purchased from Fisher Scientific (Hampton, NH, USA). Acetaldehyde (>99.5%), acetic acid (>99.9%), ethyl acetate (>99.7%), and 1-propanol (>99.9%) were obtained from Sigma-Aldrich (St. Louis, MO, USA). All chemicals were of analytical grade and used as received without further purification. Ultrapure water (resistivity ≥18.2 MΩ·cm) was used for all solution preparations

Method

Headspace gas chromatography–mass spectrometry (HS–GC–MS) was utilized to generate a dataset of standard solutions. Analyses were performed using an Agilent 8890 GC coupled to a 5977C MS and an 8697 HS. Chromatographic separation was achieved using an Agilent J&W HP-5 capillary column (30 m × 0.25 mm × 0.50 µm). Helium was used as the carrier gas at a constant flow rate of 2 mL min^{-1}. The GC inlet temperature was set to 250 °C, and a split injection mode with a split ratio of 100:1 was applied. The oven temperature was maintained isothermally at 35 °C for 5 min, which was sufficient to achieve complete elution and baseline separation of all target compounds. Headspace conditions were set to 70 °C for the oven, 80 °C for the loop, and 90 °C for the transfer line. Samples were equilibrated in the headspace oven for 10 min prior to injection, and the injection duration was 0.5 min. The mass spectrometer was operated in electron impact (EI) ionization mode at an ionization energy of 70 eV. The source and quadruple temperatures were set to 230 °C and 150 °C, respectively. Initial method development and peak identification were carried out in full-scan mode (m/z 30–80) to verify peak identity and chromatographic separation. Subsequently, selected ion monitoring (SIM) mode was employed to improve sensitivity and quantitative performance. The monitored ions and corresponding retention time windows for each analyte are summarized in Table 1:

Table 1: Retention time windows and quantifier m/z values used for SIM mode for each component.

Component	Time Segments	m/z
Acetaldehyde	-1.20	29
Ethanol	1.20-1.48	31
1-propanol	1.48-1.80	31
Ethyl acetate	1.80-	43

Sample preparation & data acquisition

Standard mixtures were prepared to generate calibration and simulated reaction datasets. All solutions were transferred into 20 mL headspace vials and salted with 3.20 g of sodium chloride per vial to enhance analyte partitioning into the headspace. The vials were immediately sealed with Teflon/silicone crimp caps prior to analysis. Standard mixtures containing 1-propanol as an internal standard (IS) were prepared, with individual analyte concentrations varied between 50 and 1000 mg L^{-1} across four concentration levels. A full factorial experimental design was employed to generate all possible combinations of species. For each combination, four calibration standards were prepared. The concentration range was divided into four equal intervals, and Latin hypercube sampling was used to assign one concentration value from each interval to every species present in a given mixture. The sampled concentration values were rounded for accurate and practical pipetting. Solutions

containing only acetic acid were excluded from the HS–GC–MS dataset. Due to its significantly lower volatility, acetic acid was quantified separately using ion chromatography (IC), as acetic acid is the only ionic species among the target analytes. Nevertheless, acetic acid was included in all standard mixtures, as its presence can influence vapor–liquid equilibrium and thereby affect the headspace partitioning behavior of the volatile components. For mixtures containing all analytes, simulated electrooxidation samples were generated by assuming a representative bioethanol feed concentration of approximately 2 M. Random selectivities toward products were applied, which yielded 12 simulated electrooxidation samples and a total of 64 calibration samples. The peak area of each analyte was extracted and normalized to the peak area of the internal standard. These normalized peak areas were used as input variables for model training, while the corresponding known concentrations served as output variables.

Model Development

To analyze the effects of increasing model complexity and prior-induced model architectures on prediction accuracy and model training, 5 different models were trained: MLR, polynomial regression, ANN, prior-induced polynomial regression, and prior-induced branched ANN. The MLR model was used as the baseline for comparison with the more complex models. Polynomial regression and ANN models were also included to enable direct comparison with the prior-induced models. All models were evaluated using 5-fold cross-validation, and MAE and RMSE were used as metrics to compare predictive performance. Model complexities were additionally assessed using the AIC an BIC.

Baseline calibration models

MLR model was employed as a baseline model due to its simplicity, interpretability, and widespread use in chemometrics. The MLR model is expressed as:

$$\hat{y} = \mathbf{X}\beta \tag{1}$$

Where $\hat{y}$ denotes the predicted output vector, $\mathbf{X}$ is the input matrix, β is the regression coefficients. Polynomial regression and ANN models were also developed to consider potential nonlinear relationships between normalized areas and analyte concentrations, and for comparison with the proposed prior-induced models. Polynomial model extends the linear model by applying a polynomial feature expansion $\mathbf{\Phi}(.)$ of a degree 2:

$$\hat{y} = \mathbf{\Phi}(\mathbf{X})^{\mathrm{T}}\beta \tag{2}$$

Despite the nonlinear transformation of the inputs, the model is still linear with respect to the parameters, β. ANNs were used as a flexible nonlinear models capable of capturing complex interactions among variables. The general expression for a feedforward ANN with L layers is given by:

$$\hat{y} = \boldsymbol{W}_L\sigma_{L-1}\left(\boldsymbol{W}_{L-1}\sigma_{L-2}\left(\ldots\sigma_1(\boldsymbol{W}_1\mathbf{X})\right)\right) \tag{3}$$

Where $\boldsymbol{W}_L$ represents the weight matrices (including bias terms) of layer L, σ_L denotes the activation function of the corresponding layer.

Linear and polynomial regression models were implemented using the scikit-learn library (version 1.7.2), while prior-induced polynomial regression and ANN models were developed using PyTorch (version 2.5.1) in Python. All models were trained and evaluated using 5-fold cross-validation to ensure robust performance assessment. For all models, normalized peak areas were used as input variables, yielding a total of 3 input features. Each model predicts 3 outputs, corresponding to the concentrations of the respective species. All input and output variables were normalized in the range 0-1. ANN training employed the mean squared error (MSE) as the loss function. Hyperparameters of the ANN models were selected using a manual exploration strategy to identify an optimized network architecture and training configuration considering dataset size. For this, a nested cross-validation strategy was adopted, where hyperparameter optimization was conducted using an additional 5-fold cross-validation within the training set of each k-fold. The activation functions evaluated included the rectified linear unit (ReLU) and the hyperbolic tangent (tanh). The number of hidden layers was varied between 1 and 3, and the number of neurons per layer ranged from 3 to 36. Model training was performed using the Adam optimizer, with learning rates manually explored in the range of 10^{-4} and 10^{-2}. Xavier weight initialization were applied to enhance training stability and convergence.

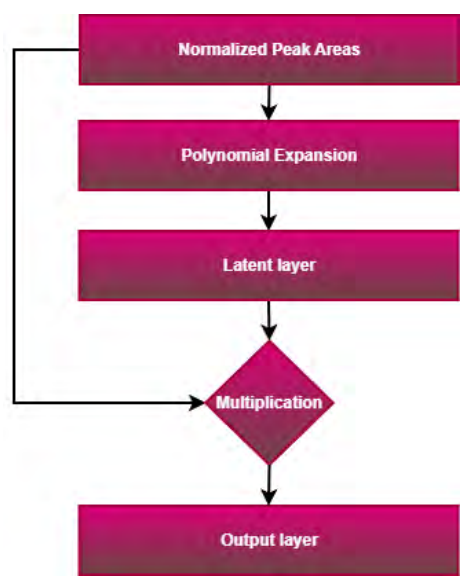

Figure 1. Physical prior-induced polynomial regression workflow.

Integration of physical prior

The prior-induced polynomial regression model was implemented in PyTorch using 3 input layers in the input layer. Polynomial feature expansion was applied to the inputs in the first layer. This layer was followed by the formation of intermediate terms. Then, these intermediate terms were multiplied with the input layer and linearly mapped to the output layer, yielding the final outputs. The resulting architecture is visualized in the Figure 1. Model training was performed using the MSE loss function and the Adam optimizer, with a fixed learning rate of 10^{-3}.

The ANN model was constructed following a similar prior-integration strategy, in which the original input variables were multiplied with the intermediate terms. To make ANN specialized for each of the output, the ANN architecture was partitioned into 3 separate sub-branches, each corresponding to an individual output. The model incorporates a shared base layer to capture common features across all outputs, followed by dedicated branch layers, which aims output-wise refinement and increased predictive performance. This branched ANN architecture, which balanced the shared information and output-specific modelling with an incorporation of physical prior is illustrated in the Figure 2.

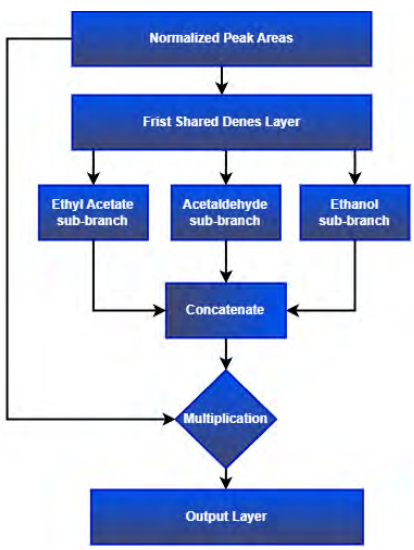

Figure 2. Physical prior-induced sub-branched ANN model.

Model Evaluation Criteria

AIC and BIC are widely used model selection criteria in machine learning that balances goodness of fit against model complexity. The two criteria differ in the way model complexity is penalized. In this study, both metrics were used to evaluate the proposed models in order to avoid selection of unnecessarily complex calibration models. For regression models, AIC is [11]:

$$AIC = n \bullet \ln\left(\frac{RSS}{n}\right) + 2k \quad (1)$$

Where n is the total number of data points used for fitting, and k is the number of parameters, which penalizes the model complexity. The first term in the equation indicates goodness of fit, and RSS is the residual sum of squares, and calculated across all components since there are 3 outputs:

$$RSS = \sum_{i=1}^{n} \sum_{j=1}^{3} \left(y_{ij} - \hat{y}_{ij}\right)^2 \quad (2)$$

In this case, AIC is calculated fold-wise and average of the 5-fold is reported to compare the models.

BIC is derived from Bayesian probability theory, and used for model selection [12]. It balances likelihood with a sample-size penalty. Like AIC, the first term evaluates the goodness of fit. BIC penalizes the number of parameters more strongly than AIC with the multiplication of $\ln(n)$ in the second term, which promotes simpler models.

$$BIC = n \bullet \ln\left(\frac{RSS}{n}\right) + k \bullet \ln(n) \quad (3)$$

For ANNs, BIC strictly limits the number of variables and favors linear models like polynomial regression with a low degree.

RESULTS & DISCUSSION

Figure 3 shows the relationship between concentration and the response behaviour of each analyte, where the primary y-axis is normalized area and the secondary y-axis corresponds to the raw peak area. The application of an internal standard has a very clear effect mitigating fluctuations between injections and instrumental variability, and improves stability and robustness significantly. However, Figure 3 also reveals that application of internal standard is not sufficient to explain relation with a perfectly linear trend, as there is still noticeable response variability.

Baseline Model Performance

Averaged over 5-fold cross-validation, MLR model yields a component-wise MAE and RMSE values of 29.26 and 35.40 for ethanol, 45.16 and 55.55 for acetaldehyde, and 49.75 and 58.00 for ethyl acetate. These results serve as a baseline for comparison with more flexible calibration models.

Parity plots for the test results of each k-fold obtained using the MLR model are shown in Figure 4 for each component separately. A clear model bias is observed in all plots, where predictions are above the parity line in the center of the range and below the parity line at the extremes. In addition, the MLR model produces non-zero predictions at zero concentration, this phenomenon is particularly evident for ethyl acetate predictions.

Standard polynomial regression with a second degree expansion results in a significant improvement in both metrics. The component-wise MAE and RMSE values are given in Table 2. The significant improvement in prediction accuracy indicates that the calibration model benefits from incorporating nonlinear relationships between the response and analyte concentrations.

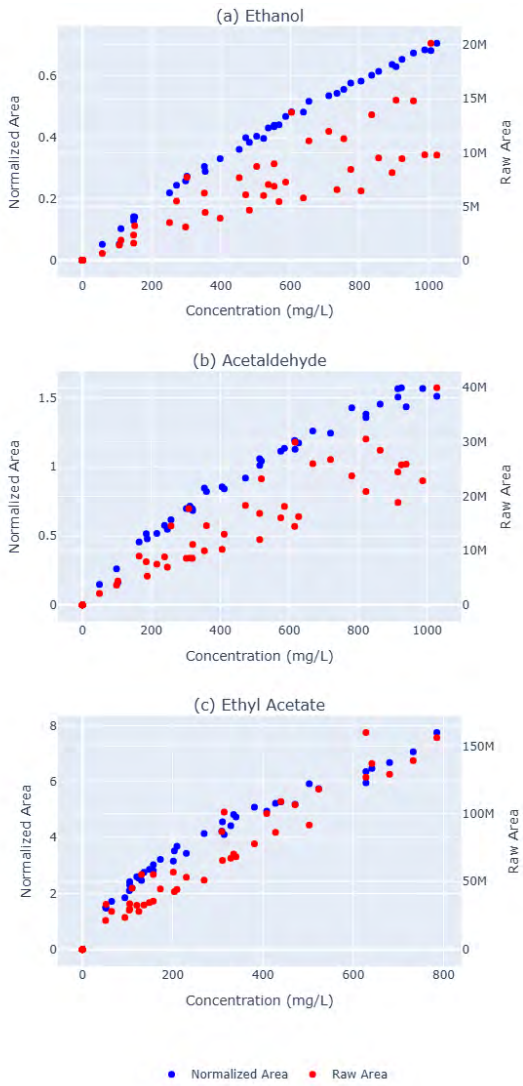

Figure 3. Standard concentrations and raw and internal standard normalized peak areas for (a) ethanol, (b) acetaldehyde, and (c) ethyl acetate.

The ANN model also exhibits predictive performance similar to polynomial regression. As shown in Table 3, the ANN yields a slightly higher MAE and RMSE for acetaldehyde, while achieving a nearly 10% decrease in MAE and RMSE for ethyl acetate. Overall, both polynomial regression and ANN outperform the MLR model, which indicates the importance of incorporating nonlinear relationships between the input variables and analyte concentrations.

Table 2: Component-wise MAE and RMSE for test set predictions, averaged over 5-fold cross-validation, obtained using polynomial regression.

Component	MAE	RMSE
Ethanol	10.11	15.03
Acetaldehyde	18.30	29.92
Ethyl acetate	15.17	23.89

Table 3: Component-wise MAE and RMSE for test set predictions, averaged over 5-fold cross-validation, obtained using ANN.

Component	MAE	RMSE
Ethanol	10.30	14.50
Acetaldehyde	19.08	30.46
Ethyl acetate	13.50	20.81

Table 4: Component-wise MAE and RMSE for test set predictions, averaged over 5-fold cross-validation, obtained using prior-induced polynomial regression.

Component	MAE	RMSE
Ethanol	7.50	12.55
Acetaldehyde	18.56	33.18
Ethyl acetate	12.04	20.50

Table 5: Component-wise MAE and RMSE for test set predictions, averaged over 5-fold cross-validation, obtained using polynomial regression prior-induced branched ANN model.

Component	MAE	RMSE
Ethanol	7.98	13.45
Acetaldehyde	17.37	29.85
Ethyl acetate	12.84	23.50

Table 6: Number of parameters, k, AIC, and BIC values for each model.

Model	k	AIC	BIC
MLR	12	332	352
Polynomial regression	30	302	352
Fully connected ANN	115	475	545
Prior-induced polynomial regression	30	316	370
Prior-induced branched ANN	155	665	803

Prior-Induced Models

The prior-induced polynomial regression and branched ANN models were evaluated using the same validation procedure as the baseline models. Incorporating prior information into the polynomial regression framework results in improved prediction accuracy for ethanol and ethyl acetate, while it yields comparable performance for acetaldehyde across both error metrics.

The component-wise MAE and RMSE values obtained using this model are summarized in Table 4.

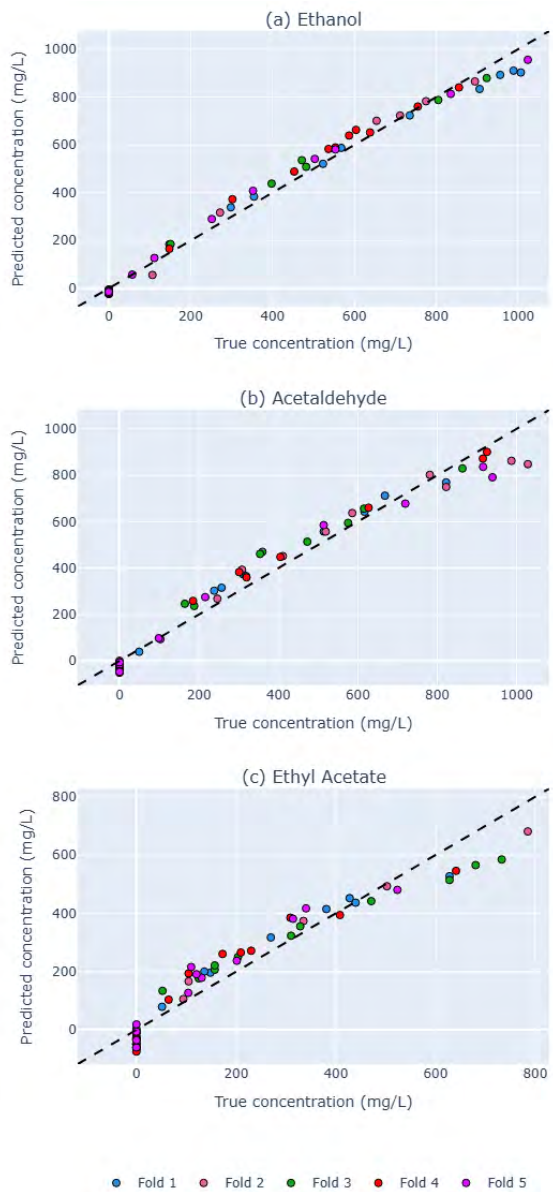

Figure 4. Parity plots for the test set of each k-fold using MLR for (a) ethanol, (b) acetaldehyde, and (c) ethyl acetate.

The prior-induced branched ANN model demonstrates consistent improvement across all metrics compared the fully connected ANN model, however, when it is compared with prior-induced polynomial regression model, branched ANN achieves slightly better prediction accuracy for acetaldehyde, while having marginally higher errors for ethanol and ethyl acetate. These results suggest that embedding prior knowledge into the network enhances predictive performance, while still having an efficient training process.

Model Complexity

Table 6 compares AIC and BIC values obtained for MLR, polynomial regression, fully connected ANN, and their prior-induced counterparts. Both AIC and BIC evaluate model performance by balancing goodness of fit against model complexity through an explicit penalty on the number of parameters. Among all evaluated models standard polynomial regression exhibits the lowest AIC and one of the lowest BIC values, which indicates that the improvement in predictive performance justifies the additional complexity relative to the MLR baseline. The fully connected ANN and prior-induced ANN results in substantially higher AIC and BIC, showing that large number of trainable parameters in ANNs are strongly penalized. The prior-induced polynomial regression has slightly increased AIC and BIC compared to standard polynomial regression. Since both models contain the same number of trainable parameters, this increase can be attributed to a higher RSS for the prior-induced formulation.

When all models are considered, ANN-based models do not yield proportional improvements in prediction accuracy despite their increased complexity, while the prior-induced polynomial regression model achieves predictive performance close to the ANN-based models, while having AIC and BIC comparable to MLR model, which makes it a good alternative in this specific use-case. However, both prior-induced polynomial regression model and branched-ANN model can provide efficient frameworks for incorporating physical prior information into calibration models for similar or more complex systems affected by matrix effects.

CONCLUSION

In conclusion, data-driven calibration models were developed for the accurate quantitative analysis of ethanol electrooxidation products using HS–GC–MS calibration data. MLR, polynomial regression, and fully-connected ANN were employed as baseline calibration models, while prior-induced polynomial regression and ANN models were developed by incorporating a simple physical constraint: The principle that a zero response corresponds to zero analyte concentration. Both polynomial regression and ANN models substantially outperformed the MLR baseline, achieving approximately 60–70% reductions in component-wise MAE and RMSE. Polynomial regression exhibited slightly better performance in acetaldehyde prediction, whereas ANN yielded marginally lower MAE and RMSE for ethyl acetate. Incorporation of prior information with polynomial regression further reduced MAE and RMSE by approximately 20–25% for ethanol and acetaldehyde. The prior-induced branched ANN achieved predictive accuracy comparable to that of the prior-induced polynomial regression model. However, both fully connected and branched ANN models

exhibited AIC and BIC values approximately twice those of the polynomial models, indicating that the ANN architectures are overly complex for the present problem. Overall, the proposed methodology demonstrates strong potential for application to more complex solution systems by effectively integrating prior physical knowledge and nonlinear relationships into model training.

ACKNOWLEDGEMENTS

Prof. Leblebici gratefully acknowledges the support of KU Leuven C2 project Scalable performant solar cell module for CO2 Reduction Coupled with Organic Synthesis (SPROUT) C2E/23/010

REFERENCES

1. St?ber M, Reemtsma T. Evaluation of three calibration methods to compensate matrix effects in environmental analysis with LC-ESI-MS. Analytical and Bioanalytical Chemistry 378:910-916 (2004). https://doi.org/10.1007/s00216-003-2442-8
2. Cheng WL, Markus C, Lim CY, Tan RZ, Sethi SK, Loh TP, . Calibration practices in clinical mass spectrometry: review and recommendations. Ann Lab Med 43:5-18 (2022). https://doi.org/10.3343/alm.2023.43.1.5
3. Moosavi SM, Ghassabian S. Linearity of calibration curves for analytical methods: a review of criteria for assessment of method reliability. Calibration and Validation of Analytical Methods - A Sampling of Current Approaches : (2018). https://doi.org/10.5772/intechopen.72932
4. Olivieri AC. Chemometrics and multivariate calibration. Introduction to Multivariate Calibration :1-26 (2024). https://doi.org/10.1007/978-3-031-64144-2_1
5. Veiga-del-Baño JM, Oliva J, Cámara MÁ, Andreo-Martínez P, Motas M. Matrix-matched calibration for the quantitative analysis of pesticides in pepper and wheat flour: selection of the best calibration model. Agriculture 14:1014 (2024). https://doi.org/10.3390/agriculture14071014
6. Chen HY, Chen C. Evaluation of calibration equations by using regression analysis: an example of chemical analysis. Sensors 22:447 (2022). https://doi.org/10.3390/s22020447
7. Ferreira V, Herrero P, Zapata J, Escudero A. Coping with matrix effects in headspace solid phase microextraction gas chromatography using multivariate calibration strategies. Journal of Chromatography A 1407:30-41 (2015). https://doi.org/10.1016/j.chroma.2015.06.058
8. Campmajó G, Saez-Vigo R, Saurina J, Núñez O. High-performance liquid chromatography with fluorescence detection fingerprinting combined with chemometrics for nut classification and the detection and quantitation of almond-based product adulterations. Food Control 114:107265 (2020). https://doi.org/10.1016/j.foodcont.2020.107265
9. Brusamarello CZ, , Di Domenico M, Da Silva C, de Castilhos F. A COMPARATIVE STUDY BETWEEN MULTIVARIATE CALIBRATION AND ARTIFICIAL NEURAL NETWORK IN QUANTIFICATION OF SOYBEAN BIODIESEL. Rev.Mex.Ing.Quim. 19:123-132 (2019). https://doi.org/10.24275/rmiq/bio579
10. Gholivand MB, Jalalvand AR, Goicoechea HC, Gargallo R, Skov T, Paimard G. Combination of electrochemistry with chemometrics to introduce an efficient analytical method for simultaneous quantification of five opium alkaloids in complex matrices. Talanta 131:26-37 (2015). https://doi.org/10.1016/j.talanta.2014.07.053
11. Cavanaugh JE, Neath AA. The akaike information criterion: background, derivation, properties, application, interpretation, and refinements. WIREs Computational Stats 11: (2019). https://doi.org/10.1002/wics.1460
12. Neath AA, Cavanaugh JE. The bayesian information criterion: background, derivation, and applications. WIREs Computational Stats 4:199-203 (2011). https://doi.org/10.1002/wics.199

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Multi-Objective CAPE Simulation of Agro-Industrial Systems Integrating High-Yield Sugarcane and the Inversion Process

Satoshi Ohara[a]*, Yoshifumi Terajima[b], Hiro Tabata[c], and Yasunori Kikuchi[cd]

[a] The University of Tokyo, Research Center for Advanced Science and Technology, Meguro-ku, Tokyo, Japan
[b] Japan International Research Center for Agricultural Sciences, Tropical Agriculture Research Front, Ishigaki, Okinawa, Japan
[c] The University of Tokyo, Presidential Endowed Chair for "Platinum Society", Bunkyo-ku, Tokyo, Japan
[d] The University of Tokyo, Institute for Future Initiatives, Bunkyo-ku, Tokyo, Japan
* Corresponding Author: s-ohara@g.ecc.u-tokyo.ac.jp

ABSTRACT

This study develops a multi-objective computer-aided process engineering (CAPE) framework to evaluate integrated sugarcane-based agro-industrial systems combining a high-yield cultivar, Haru-no-Ougi, and the "Inversion Process, " which reverses the conventional order of sugar crystallization and ethanol fermentation through selective fermentation of reducing sugars. The in-house CAPE tool SugaNol integrates agricultural, industrial, and environmental (life cycle assessment) models to simulate productivity, energy balance, greenhouse-gas (GHG) emissions, and relative economic performance on a per-hectare basis. Four scenarios were analyzed: NiF8–Conventional, KY01-2044–Conventional, Haru-no-Ougi–Conventional, and Haru-no-Ougi–Inversion Process. Simulation results showed that the combined Haru-no-Ougi and Inversion Process system increased total energy-equivalent productivity by approximately 40-45% compared with the baseline NiF8 system. Life cycle GHG emissions were reduced by 4–11%, while relative economic performance improved by approximately 35-40% due to enhanced ethanol and electricity co-production. These results demonstrate that CAPE-based multi-objective simulation enables holistic assessment and design of agro-industrial systems by explicitly linking crop traits and process design, providing a quantitative basis for sustainable sugar–ethanol–energy system design.

Keywords: Computer-Aided Process Engineering (CAPE), Life Cycle Assessment (LCA), Agro-Industrial Symbiosis, Bagasse Utilization, Sustainable Sugar–Ethanol–Energy Systems

INTRODUCTION

Sugarcane (*Saccharum* spp.) has historically been cultivated primarily for sugar production; however, in the era of global decarbonization and resource circularity, it is increasingly recognized as a promising feedstock for integrated agro-industrial systems producing sugar, ethanol, and electricity. Its unique combination of high biomass productivity and sugar accumulation makes it one of the few crops capable of providing both food and energy from the same land area [1-3]. Consequently, sugarcane-based industries play a central role in achieving carbon neutrality and the development of circular bioeconomies in tropical and subtropical regions.

Nevertheless, recent breeding strategies aimed at yield improvement have revealed a significant trade-off between total cane yield and sucrose content [3]. The conventional Japanese cultivar NiF8 exhibits a moderate cane yield (60-70 t ha^{-1}) with high sucrose content (13-14%), whereas the high-biomass cultivar KY01-2044 achieves a much higher cane yield (90-100 t ha^{-1}) but at the expense of lower sucrose (9-10%). As a result, despite increased biomass productivity, sugar yield per unit land area is reduced by on the order of 10% compared with conventional cultivars. Such compositional trade-offs necessitate process-system innovations that valorize both sugars and lignocellulosic residues to improve overall system efficiency.

A new high-yield Japanese cultivar, Haru-no-Ougi, was developed through breeding to overcome this

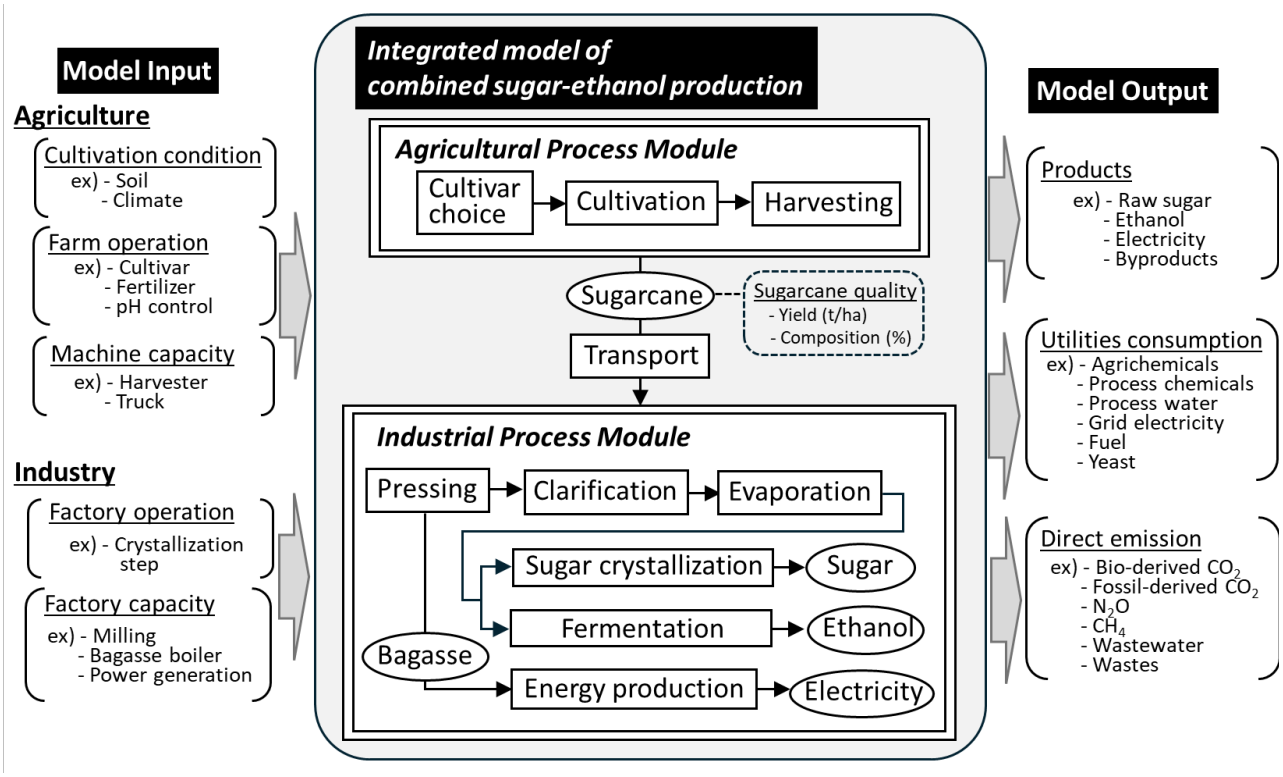

Figure 1. Integrated CAPE framework "SugaNol" linking agricultural, industrial, and environmental (LCA) models for sugarcane-based systems.

limitation [4]. It combines high cane yield (100-120 t ha^{-1}) with moderate-high sucrose content (11-12%) and increased fiber content (14-15%), enabling both efficient sugar crystallization and enhanced cogeneration from bagasse. To maximize the benefits of this cultivar, however, the industrial process configuration must be optimized in an integrated manner [2].

To address this challenge, the "Inversion Process" was introduced by [3, 5, 6] as a novel industrial scheme that reverses the conventional sugar–ethanol production sequence. Rather than crystallizing sucrose first and fermenting residual molasses, this approach selectively ferments reducing sugars (glucose and fructose) using invertase-defective yeast before crystallization. By removing inhibitory reducing sugars early, the process achieves higher sucrose purity, better crystallization efficiency, and greater ethanol yield, contributing to improved energy integration.

Despite these technological advances, a comprehensive evaluation integrating both agricultural and industrial performance remains limited [7]. Conventional life cycle assessment (LCA) and process optimization studies often focus on either cultivation or factory operation in isolation. Therefore, this study aims to establish an integrated Computer-Aided Process Engineering (CAPE) framework that couples agricultural, industrial, and environmental models to evaluate and optimize sugarcane-based agro-industrial systems.

The objective of this research is to apply the CAPE-based multi-objective simulation framework to the high-yield cultivar Haru-no-Ougi combined with the Inversion Process. The framework quantifies agricultural and industrial productivity, life cycle greenhouse-gas (GHG) emissions, and the associated economic performance on a per-hectare basis to identify synergistic solutions that simultaneously enhance food, fuel, and energy production toward carbon neutrality.

METHODOLOGY

Integrated CAPE Simulation Framework

An integrated CAPE framework was developed using the in-house Excel-based simulation tool *SugaNol*, originally introduced by Ouchida *et al.* [8]. The tool integrates three sub-models—agricultural, industrial, and environmental (LCA)—within a single platform to allow mass and energy balance consistency between field and factory operations. Figure 1 illustrates the overall structure of the integrated CAPE framework, linking agricultural production, industrial processing, and life cycle assessment within a unified simulation platform.

The simulation workflow begins with agricultural data (yield, composition) and propagates through process simulation to calculate product yields, energy use,

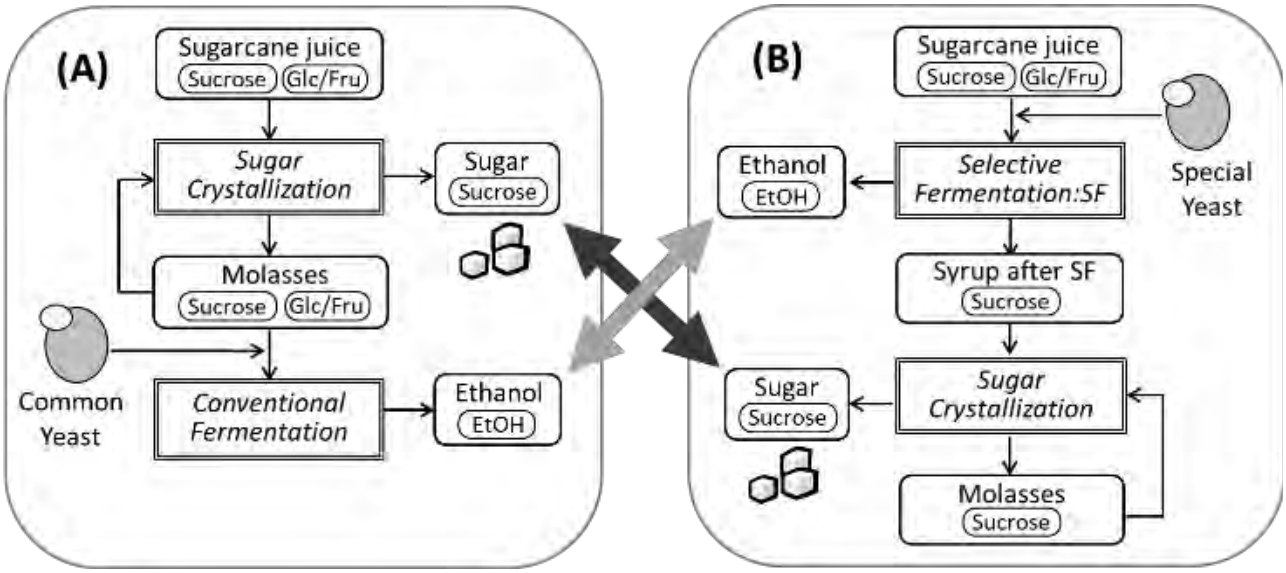

Figure 2. Process flow comparison between (A) the Conventional Process and (B) the Inversion Process for integrated sugar and ethanol production.

and emissions. The environmental model subsequently evaluates life cycle impacts and provides feedback for comparative multi-objective evaluation. All results are normalized to the functional unit of 1 ha of field.

Agricultural Model

The agricultural model represents field operations and compositional characteristics of different cultivars. It considers inputs of fertilizers (N, P, K), compost, pesticides, irrigation water, and fuel consumption for tillage, planting, harvesting, and transport. Outputs include cane yield and stalk composition (sucrose, reducing sugars, fiber, and moisture).

Three cultivars were simulated:
(1) NiF8 (baseline, conventional yield and high sucrose),
(2) KY01-2044 (high-yield, high-fiber), and
(3) Haru-no-Ougi (high-yield, moderate-high sucrose).

Agronomic input coefficients were derived from field trials in Kyushu, Japan, adjusted for energy and emissions factors based on the Japanese LCA database IDEA v3.5.1. [9]. Table 1 summarizes the cane yield and composition of the three cultivars used as primary inputs for the simulations.

Table 1: Yield and compositional characteristics of sugarcane cultivars (NiF8, KY01-2044, and Haru-no-Ougi).

	NiF8	KY01-2044	Haru-no-Ougi
	Conventional	High-yield	High-yield
Yield[t/ha]	67.1	93.5	107.3
Sucrose[%]	13.5	9.3	11.8
Glc & Fru[%]	0.6	1.2	0.4
Fiber[%]	12.3	14.0	14.8

The model assumes that all cultivars are cultivated under identical management practices to isolate genetic and compositional effects. Fertilizer rates were 150, 60, and 180 kg ha^{-1} for N, P_2O_5, and K_2O, respectively. Diesel fuel for agricultural machinery was set at 45 L ha^{-1}, and irrigation energy was assumed negligible under rain-fed conditions [7, 8].

Industrial Process Model

The industrial model simulates two process configurations: (a) Conventional Process – sugar crystallization followed by molasses fermentation; (b) Inversion Process – selective fermentation of reducing sugars before sucrose crystallization [3]. Mass and energy balances were solved across clarification, crystallization, fermentation, distillation, and cogeneration units. Steam and power demands were met through bagasse combustion in high-pressure boilers (boiler efficiency: 90%). Surplus electricity was exported to the grid. Fermentation efficiency was fixed at 85%, and sucrose crystallization recovery at 95%.

In the Inversion Process, glucose and fructose—natural inhibitors of sucrose crystallization—were selectively fermented to ethanol prior to crystallization using an invertase-defective yeast [5, 6]. This configuration improved energy balance and ethanol yield without sacrificing sugar quality. Figure 2 schematically compares the Conventional Process and the Inversion Process configurations considered in this study.

Life Cycle Assessment (LCA)

The environmental impact model calculates life cycle greenhouse-gas (LC-GHG) emissions on a cradle-to-gate basis, focusing on emissions from sugarcane

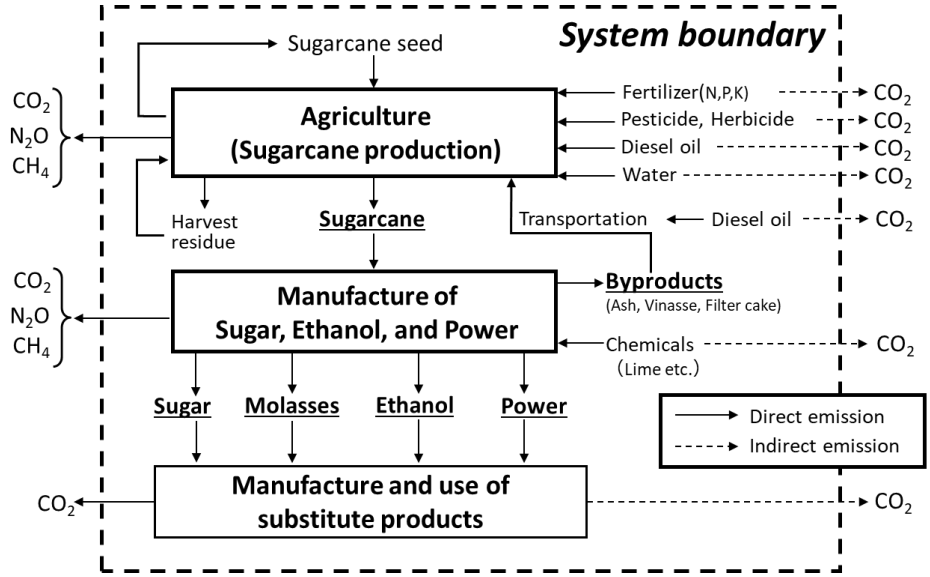

Figure 3. Life cycle assessment (LCA) system boundary for the integrated sugarcane-based agro-industrial system (cradle to gate).

cultivation through industrial processing. The system boundary encompasses agricultural production, transportation of sugarcane, and the production of sugar, ethanol, and electricity. The functional unit was defined as the "use of 1 ha of field." Figure 3 shows the system boundary of the life cycle assessment.

Input data included fertilizers, pesticides, compost, and fuel in agriculture, as well as materials, fuel, electricity, and auxiliary chemicals in factory operations. Both direct emissions from agricultural and industrial processes and indirect emissions associated with upstream production of inputs were accounted for. Background inventory data were obtained from the Japanese LCA database IDEA v3.5.1. [9]

Only the climate-change impact category was considered in this study. The greenhouse gases considered were CO_2, CH_4, and N_2O. Life cycle GHG emissions were calculated using 100-year global warming potentials (GWP100) from the IPCC Fourth Assessment Report (CO_2 = 1, CH_4 = 25, N_2O = 298) and expressed as LC-GHG emissions [kg CO_2 eq ha^{-1}] [10].

RESULTS AND DISCUSSION

Simulation Results on Productivity and Life Cycle GHG Emissions

Four simulation cases were analyzed:

(1) NiF8 – Conventional Process (baseline),

(2) KY01-2044 – Conventional Process,

(3) Haru-no-Ougi – Conventional Process, and

(4) Haru-no-Ougi – Inversion Process.

Figure 4 compares the production of sugar, ethanol, and electricity, as well as life cycle GHG emissions, across the four simulated scenarios.

For Case (1), the baseline configuration using NiF8 under the conventional process produced 8.3 t sugar ha^{-1} and 0.35 kL ethanol ha^{-1}, with 348 kWh ha^{-1} surplus power and LC-GHG emissions of 3.94 t CO_2 eq ha^{-1}.

In Case (2), where the high-biomass cultivar KY01-2044 was processed using the conventional process, ethanol production increased to 0.80 kL ha^{-1} (+126%) and electricity generation to 469 kWh ha^{-1}, while LC-GHG emissions were reduced by approximately 11%. However, the lower sucrose content of this cultivar resulted in a reduced sugar yield of 7.5 t ha^{-1}, confirming the yield–composition trade-off.

Case (3), corresponding to Haru-no-Ougi under the conventional process, achieved the highest sugar yield (11.4 t ha^{-1}) together with substantial ethanol co-production (0.43 kL ha^{-1}) and 669 kWh ha^{-1} surplus electricity, although total LC-GHG emissions were slightly higher (4.66 t CO_2 eq ha^{-1}) due to increased industrial energy demand.

In Case (4), which combines Haru-no-Ougi with the Inversion Process, the system delivered 11.6 t sugar ha^{-1} and 0.35 kL ethanol ha^{-1} with complete energy self-sufficiency and LC-GHG emissions of 4.47 t CO_2 eq ha^{-1} (–4% relative to Case (3)). When evaluated in energy-equivalent terms (sugar + ethanol), productivity per hectare increased by approximately 45% compared with the Case (1) baseline, demonstrating the synergistic benefit of integrating high-yield cultivars with process innovation.

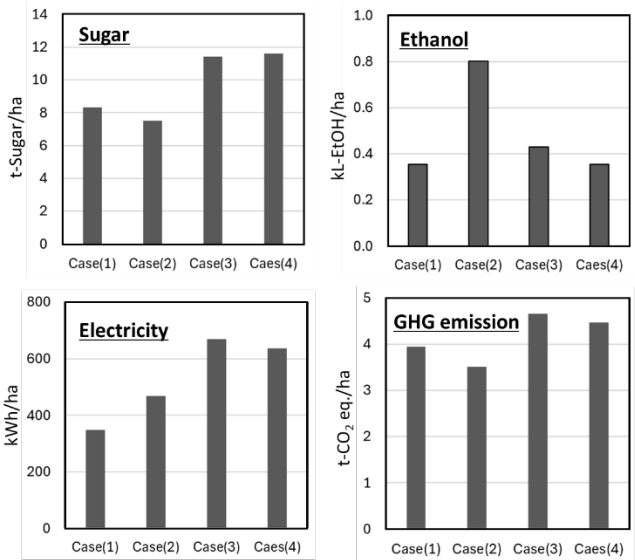

Figure 4. Comparison of sugar, ethanol, electricity production, and life cycle GHG emissions among four simulated scenarios.

Discussion: Environmental Impact and Bagasse Utilization

LC-GHG emissions ranged from 3.5 to 4.7 t CO_2 eq ha^{-1} among the four cases. Improvements in the Inversion Process stem primarily from enhanced ethanol yield and efficient steam integration. The higher emissions in Haru-no-Ougi–Conventional reflect the increased throughput and energy demand, but these are offset when normalized per unit of energy-equivalent output.

In this simulation, all surplus bagasse was assumed to be combusted in factory boilers for electricity generation. Alternative utilization routes, however, have been explored in previous studies using simulation-based approaches [11]. Corradini *et al.* clarified through process-system simulation that torrefaction—a mild pyrolysis process at 250–300 °C—can convert bagasse into a solid carbon-rich product with improved energy density and stability. Previous studies have also reported that biochar derived from thermochemical conversion can enhance soil fertility, improve water retention, and promote long-term carbon sequestration when applied to agricultural fields [12, 13]. Although detailed calculations for torrefaction-based bagasse recycling were not conducted in this study, these findings suggest that integrating such utilization pathways could further reduce LC-GHG emissions through additional carbon storage and reduced upstream inputs.

Overall, these results indicate that bagasse utilization strategies play a critical role in determining the environmental performance of sugarcane-based agro-industrial systems. While electricity generation via direct combustion represents a practical baseline configuration, simulation-based assessments highlight the potential of alternative pathways to support deeper decarbonization when integrated with agricultural and industrial system design.

Discussion: Economic Evaluation and Multi-Objective Assessment

A simplified economic evaluation was conducted on a relative basis using representative unit prices of 100 JPY kg^{-1} for sugar, 70 JPY L^{-1} for ethanol, and 15 JPY kWh^{-1} for electricity. These values were applied consistently across all cases to compare revenue structures on a per-hectare basis. The resulting relative economic performance of the four cases, normalized to the Case (1) baseline, is summarized in Table 2.

Table 2: Relative economic performance of the four simulated scenarios normalized to the Case (1) baseline.

Case	Cultivar & Process	Sugar revenue	Ethanol revenue	Electricity revenue	Total revenue
(1)	NiF8 & Conv.	100	100	100	100
(2)	KY01-2044 & Conv.	90	226	135	95
(3)	Haru-no-Ougi & Conv.	137	122	192	137
(4)	Haru-no-Ougi & Inversion process	140	100	182	139

In Case (2), based on the high-yield cultivar KY01-2044, sugar production decreased compared with the baseline, while ethanol and electricity outputs increased substantially. This case represents a configuration that sacrifices sugar production in favor of fuel and energy generation. As summarized in Table 2, the total economic performance of Case (2) was approximately 5% lower than that of the Case (1) baseline, indicating that increased fuel revenue did not fully compensate for reduced sugar output.

In contrast, Cases (3) and (4), based on the high-yield cultivar Haru-no-Ougi, achieved significantly higher sugar production together with increased electricity generation. These cases illustrate a configuration that balances sugar and power production rather than prioritizing fuel output. Table 2 shows that the total economic performance increased by approximately 35–40% relative to the baseline, reflecting the combined effects of higher cane yield, improved sucrose recovery, and enhanced bagasse utilization.

Case (4), which combines Haru-no-Ougi with the Inversion Process, further shifts the product balance toward sugar maximization. Although ethanol production was slightly reduced due to decreased molasses formation, this configuration achieved the highest overall economic performance among the four cases.

Rather than performing numerical optimization, the CAPE model was used to conduct a multi-objective assessment of trade-offs among productivity, life cycle GHG emissions, and economic performance. This comparative evaluation demonstrates that scenarios integrating Haru-no-Ougi with the Inversion Process achieve a balanced improvement across multiple performance indicators, without sacrificing sugar productivity or environmental performance.

The results further demonstrate that multi-objective evaluation can effectively identify system-level synergies even in the absence of formal optimization. By simultaneously considering agricultural yield, industrial output, environmental impact, and economic implications, the CAPE framework provides a quantitative basis for informed decision-making in agro-industrial system design. This approach is consistent with the "Design for Sustainability" concept proposed in the literature [14], which emphasizes the integrated evaluation of technical, economic, and environmental dimensions within a unified modeling framework.

CONCLUSIONS

This study developed a multi-objective CAPE-based simulation framework for integrated sugarcane-based agro-industrial systems. By coupling agricultural, industrial, and environmental models, the framework enables system-wide comparative evaluation of productivity, energy performance, life cycle GHG emissions, and relative economic performance.

Simulation results confirmed that combining the high-yield cultivar Haru-no-Ougi with the Inversion Process increased total energy-equivalent productivity by approximately 40–45% relative to the conventional NiF8 system. Life cycle GHG emissions were reduced by 4–11%, while relative economic performance improved by up to approximately 40%, depending on the balance between sugar, ethanol, and electricity production. The case-by-case comparison clarified distinct system roles, ranging from fuel-oriented configurations to designs prioritizing sugar maximization through improved sucrose recovery.

Rather than relying on numerical optimization, this study demonstrated that multi-objective CAPE-based assessment can identify system-level synergies and support design choices aligned with regional priorities. Overall, the proposed framework illustrates how modelling and simulation can bridge crop breeding and process design, providing a quantitative basis for data-driven strategies toward carbon-neutral and resource-efficient agro-industrial symbioses.

ACKNOWLEDGEMENTS

This work was supported by New Energy and Industrial Technology Development Organization (NEDO, Grant number JPNP20011), JSPS KAKENHI (Grant Number JP23K11521), JST COI-NEXT (Grant Number JPMJPF2003). The activities of the Presidential Endowed Chair for "Platinum Society" at the University of Tokyo are supported by Mitsui Fudosan Corporation, Sekisui House, Ltd., East Japan Railway Company, and Toyota Tsusho Corporation.

AUTHOR IDENTIFIERS

Author ORCIDs:
Satoshi Ohara: 0000-0002-5273-8969
Yoshifumi Terajima: 0000-0003-3201-3183
Hiro Tabata: 0000-0001-8137-9136
Yasunori Kikuchi: 0000-0003-0002-8233

REFERENCES

1. Macedo IC, Seabra JEA, Silva JEAR. Green house gases emissions in the production and use of ethanol from sugarcane in brazil: the 2005/2006 averages and a prediction for 2020. Biomass and Bioenergy 32:582-595 (2008). https://doi.org/10.1016/j.biombioe.2007.12.006
2. Ohara S, Hamada Y, Terajima Y, Ishida T, Kikuchi Y, Fukushima Y, Sugimoto A. Effect of single-boiling crystallization of high-yielding sugarcane KY01-2044 on sugar and molasses production. FSTR 30:47-56 (2024). https://doi.org/10.3136/fstr.fstr-d-23-00158
3. Ohara S, Fukushima Y, Sugimoto A, Terajima Y, Ishida T, Sakoda A. Rethinking the cane sugar mill by using selective fermentation of reducing sugars by saccharomyces dairenensis, prior to sugar crystallization. Biomass and Bioenergy 42:78-85 (2012). https://doi.org/10.1016/j.biombioe.2012.03.024
4. Hattori, T., Terajima, Y., Sakaigaichi, T., Terauchi, T., Tarumoto, Y., Adachi, K., Hayano, M., Tanaka, M., Ishikawa, S., Umeda, M., Matsuoka, M., Sugimoto, A. High ratoon yield sugarcane cultivar "Harunoogi" developed for Kumage region by using an interspecific hybrid between a commercial cultivar and *Saccharum spontaneum* L.
5. Bull NARO Agric. Food Res. Inst., 2. 21-44 (2019) https://doi.org/10.34503/naroj.2019.2_21
6. Ohara S, Kato T, Fukushima Y, Sakoda A. Selective ethanol production from reducing sugars in a saccharide mixture. Journal of Bioscience and Bioengineering 115:540-543 (2013). https://doi.org/10.1016/j.jbiosc.2012.11.010

7. Kato T, Ohara S, Fukushima Y, Sugimoto A, Masuda T, Yasuhara T, Yamagishi H. Development of flocculent saccharomyces cerevisiae strain GYK-10 for the selective fermentation of glucose/fructose in sugar mills. Journal of Bioscience and Bioengineering 122:58-63 (2016). https://doi.org/10.1016/j.jbiosc.2015.12.012
8. OHARA S, KIKUCHI Y, OUCHIDA K, SUGIMOTO A, HATTORI T, YASUHARA T, FUKUSHIMA Y. Reduction of greenhouse gas emissions in the introduction of inversion system to produce sugar and ethanol from sugarcane. Journal of Life Cycle Assessment, Japan 15:86-100 (2019). https://doi.org/10.3370/lca.15.86
9. Ouchida K, Fukushima Y, Ohara S, Sugimoto A, Hirao M, Kikuchi Y. Integrated design of agricultural and industrial processes: a case study of combined sugar and ethanol production. AIChE Journal 63:560-581 (2016). https://doi.org/10.1002/aic.15374
10. National Institute of Advanced Industrial Science and Technology (AIST), 2025. Inventory Database for Environmental Analysis (IDEA), Version 3.5.1, Tokyo, Japan .
11. Intergovernmental Panel on Climate Change (IPCC). Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the IPCC.
12. Corradini, L.L., Heiho, A., Kanematsu, Y., Shimono, R., Ohara, S., Kikuchi, Y. Simulation-based design of regional biomass thermochemical conversion system for improved environmental and socio-economic performance. *Comput Aided Chem Eng,* 52:2361-2366 (2023) https://doi.org/10.1016/B978-0-443-15274-0.50375-9
13. Bridgwater AV. Review of fast pyrolysis of biomass and product upgrading. Biomass and Bioenergy 38:68-94 (2012). https://doi.org/10.1016/j.biombioe.2011.01.048
14. Lehmann J, Rillig MC, Thies J, Masiello CA, Hockaday WC, Crowley D. Biochar effects on soil biota – a review. Soil Biology and Biochemistry 43:1812-1836 (2011). https://doi.org/10.1016/j.soilbio.2011.04.022
15. Gnansounou E, Dauriat A, Villegas J, Panichelli L. Life cycle assessment of biofuels: energy and greenhouse gas balances. Bioresource Technology 100:4919-4930 (2009). https://doi.org/10.1016/j.biortech.2009.05.067

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Towards the Decarbonization of a Conventional Ammonia Plant by the Gradual Incorporation of Green Hydrogen and Air-Separated Nitrogen

João Fortunato, Diogo A. C. Narciso*, and Henrique A. Matos

Centro de Recursos Naturais e Ambiente, Departamento de Engenharia Química, Instituto Superior Técnico, Universidade de Lisboa, Av. Rovisco Pais 1, 1049-001 Lisboa, Portugal
* diogo.narciso@tecnico.ulisboa.pt.

ABSTRACT

As economies advance towards decarbonization, industry follows suit. The ammonia (NH_3) sector heavily relies on the energy- and carbon-intensive Haber-Bosch (HB) process, which accounts for nearly 2% of global CO_2 emissions due to its reliance on fossil fuels. Emerging technologies are paving the way for fully renewable NH_3 production, although the most mature green process still relies on the HB process, entirely replacing fossil fuels with electrolytic green hydrogen (H_2). This work introduces the first developments towards the gradual incorporation of green H_2 and air-separated nitrogen (N_2) into a conventional NH_3 production plant. Using *Aspen Plus® V14* for modelling and simulation of Steam Methane Reforming (SMR), different scenarios incorporating 0 to 40 % of these alternative feedstocks were analyzed. An economic objective function is used in each scenario's optimization. To improve green H_2 incorporation and ensure operational constraints were met, the simulations used an adapted SMR section proposed in previous work. Preliminary results in the development of this methodology indicate that at 40 % green H_2 and 40 % air-separated N_2 incorporations, a 29.4 % reduction in carbon emissions was achieved, along with 30.1 % natural gas savings and a 35 % decreased in conventional operational costs (without green H_2 and air-separated N_2 costs) compared to the base case, where no alternative feedstocks were used. At current market prices, in the same scenario, incorporating alternative feedstocks can increase total operational costs by 52 %.

Keywords: Ammonia, Steam Methane Reforming, Green Hydrogen, Aspen Plus

INTRODUCTION

At the beginning of the 20th century, agriculture faced increasing pressure to secure a reliable source of nitrogen to enable higher food production for a rapidly growing population. Nitrogen is an essential nutrient for crop growth, and until then, it was obtained mainly from limited natural sources. In 1909, Fritz Haber achieved the fixation of atmospheric nitrogen (N_2) by synthesizing ammonia (NH_3) through its reaction with hydrogen (H_2) using an iron-based catalyst, Equation (1). Years later, in 1913, Carl Bosch successfully scaled up ammonia synthesis for industrial application, giving rise to what is now known as the Haber-Bosch (HB) process. This breakthrough is widely regarded as one of the most important technological developments in human history. [1]

$$N_2 + 3\,H_2 \rightleftharpoons 2\,NH_3 \tag{1}$$

Ammonia is the second-most-produced chemical globally, with 85 % of its output used in the nitrogen-based fertilizer industry, including urea and ammonium nitrate. It is estimated that almost half of the world's population relies directly on these fertilizers for food production. Besides agriculture, ammonia has diverse uses, including pharmaceuticals, textiles, explosives manufacturing, and deNOx technologies. [2, 3]

Given its global significance and extensive production scale, the Haber-Bosch (HB) process accounts for approximately 2 % of the world's total energy consumption and contributes between 1 % and 2 % to worldwide carbon dioxide (CO_2) emissions. [1] This considerable environmental footprint primarily stems from its substantial

reliance on fossil fuels, which serve as feedstocks for hydrogen production and process heat. Currently, around 75 % of global ammonia synthesis is derived from light hydrocarbons such as natural gas and naphtha, whereas the remaining 25 % originates from heavy hydrocarbons, including coal and Heavy Fuel Oil (HFO). In the European context, nearly all ammonia manufacturing facilities operate exclusively on natural gas, producing hydrogen via Steam Methane Reforming (SMR). [2, 3] This manufacturing route is regarded as the Best Available Technology, owing to its superior energy efficiencies, reduced emissions, and lower capital investment.

Driven by policies aimed at achieving carbon neutrality and energy independence, the ammonia industry is increasingly exploring decarbonization pathways. In the European Union (EU), the Renewable Energy Directive III stipulates that by 2030, 42 % of the hydrogen used in industrial processes must be derived from renewable sources. To align with these regulations and targets, the ammonia sector has intensified efforts to develop alternative synthesis technologies that transcend the conventional HB process, striving for a more decentralized, flexible, and energy-efficient production. The emerging routes under investigation are designed to facilitate the complete integration of renewable energy sources and encompass electrochemical, electrocatalytic, and photocatalytic methods for ammonia synthesis. Nevertheless, most of these technological approaches remain at early stages of development and are distant from commercial deployment. [4]

Currently, the most mature technology for green NH_3 production is based on electrifying the HB process, in which SMR is replaced by water electrolysis as the hydrogen source and an Air Separation Unit as the nitrogen source. By supplying renewable electricity to the process, green NH_3 can be produced. However, operating the HB process within the inherent variability of renewable energy sources, particularly solar and wind, remains a significant challenge. Numerous studies have addressed the issues related to fluctuating energy supply and, consequently, hydrogen production. Under such conditions, the operation of the HB process often requires hydrogen storage in buffer tanks and/or the use of energy storage systems such as batteries. Therefore, producing fully green NH_3 remains costly due to expensive infrastructure, high energy prices, and logistical challenges stemming from the underdeveloped electrolyser market. [4]

The NH_3 industry is consequently striving to balance economic competitiveness with the attainment of decarbonization targets, ideally leveraging existing capital investments in current infrastructure. As examples, Fertiberia and Yara are major players in the European ammonia industry, which has begun incorporating green hydrogen into its conventional plants. Inspired by the operating procedures at a real-life plant, a previous work addressed the incorporation of green hydrogen into a conventional, methane-fed NH_3 plant. [5] The authors pointed out that overheating in some equipment limits green H_2 integration and proposed a new SMR configuration to prevent it. With an adapted SMR, a maximum green hydrogen incorporation of 60 % can be achieved, although lower fractions may lead to issues with equipment load reduction. [5]

Building on the previous study, this work explores the simultaneous use of green H_2 and air-separated N_2 as alternative feedstocks in the SMR section of a conventional ammonia plant.

METHODOLOGY

As part of a conventional ammonia plant with an annual capacity of around 200 kton, an SMR section was modeled in *Aspen Plus® V14*, as shown in **Figure 1**.

SMR Section

The SMR section relies on natural gas, steam, and atmospheric air consumption to produce a synthesis gas mixture of H_2 and N_2. The process natural gas is first purified by hydrodesulfurization to remove sulphur compounds and prevent catalyst poisoning. This stream is then heated, mixed with process steam, and fed into a pre-reforming section, where the heavier hydrocarbons in natural gas are converted into methane (CH_4). Following this step, the reforming process begins in the primary reformer, Ref-I, where CH_4 reacts with steam to produce H_2, carbon monoxide (CO), and CO_2, Equations (2) and (3).

$$CH_4 + H_2O \rightleftharpoons CO + 3\,H_2 \tag{2}$$

$$CO + H_2O \rightleftharpoons CO_2 + H_2 \tag{3}$$

The gas mixture is mixed with compressed hot atmospheric air and fed into a secondary reformer, Ref-II, where CH_4 conversion is completed. The atmospheric air supplied in this unit must ensure the right amount of N_2 for ammonia synthesis, as well as the oxygen (O_2) needed for CH_4 partial and total oxidation, Equations (4) and (5).

$$CH_4 + 0.5\,O_2 \rightleftharpoons CO + 2\,H_2 \tag{4}$$

$$CH_4 + 2\,O_2 \rightleftharpoons CO_2 + 2\,H_2O \tag{5}$$

The mixture is then cooled, and CO is converted in a two-stage process: a high-temperature shift (HT) followed by a low-temperature shift (LT) according to the stoichiometry in Equation (3). CO_2 is removed from the process stream by scrubbing the mixture with an ammonia solution. To ensure that the syngas mixture produced does not contain oxygen-containing compounds, the unconverted CO fraction and the final traces of CO_2 are subjected to a methanation step, in which CH_4 and water

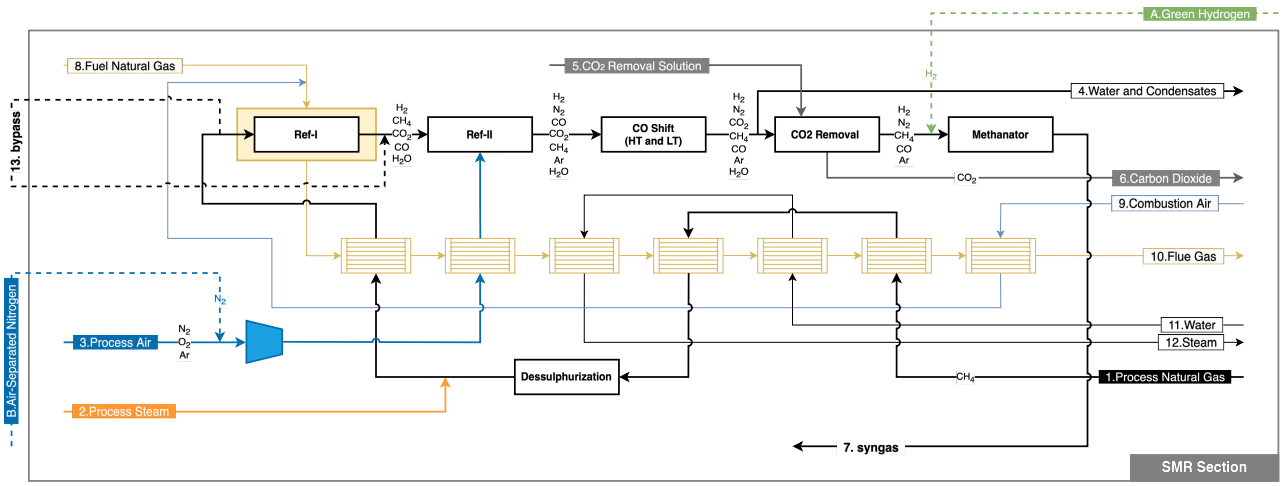

Figure 1: Simplified schematics of the SMR section model developed. The grey box contains the conventional SMR configuration; outside the box are represented the adapted SMR configuration (stream 13, a Ref-I Bypass proposed in [5,6]), stream A (green hydrogen incorporation), and stream B (air-separated nitrogen incorporation).

are produced.

The water is ultimately eliminated using molecular sieve adsorbers, allowing CH_4 and the Argon (Ar) introduced with atmospheric air to proceed to NH_3 synthesis as inert chemicals. The gas mixture obtained is then composed of H_2 and N_2 in a 3:1 molar ratio, with traces of CH_4 and Ar.

In addition to the main process line, the SMR section encompasses an auxiliary line that utilizes natural gas as fuel through combustion with air. The heat generated from this combustion process is employed to supply energy for reforming reactions (as described in Ref-I), to pre-heat various streams such as process natural gas and combustion air, and to generate steam.

Since the integration of alternative feedstocks is considered within the SMR section and the objective remains to produce equivalent flow rates of H_2 and N_2, the NH_3 synthesis process is not addressed in this analysis. Ultimately, incorporating green H_2 and air-separated N_2 could reduce the inert chemical fraction in the syngas stream, thereby decreasing inert component accumulation in the NH_3 synthesis loop and reducing purge fractions. Those additional advantages at NH_3 synthesis are not addressed in this work.

Green Hydrogen Incorporation

The incorporation of green H_2 was modeled as a pure stream (xH_2 = 1). The injection point into the process was positioned immediately prior to the methanation step, where the production of gry hydrogen (from CH_4 reforming) is completed. The study did not include modeling of the water electrolysis process used for green hydrogen production.

Green H_2 incorporation reduces grey H_2 production, thereby decreasing CH_4 reforming and helping achieve potential decarbonization. The extent of this alternative feedstock's incorporation is quantified by the Green Hydrogen Incorporation fraction (GHI), defined as the ratio of green H_2 in the total H_2 supply to NH3 synthesis (grey + green).

Air-Separated Nitrogen Incorporation

The incorporation of air-separated N_2 was treated similarly to green H_2, considering it as a pure stream (xN_2 = 1). This stream was intended to simulate the nitrogen-rich output typically produced by a well-established Air Separation Unit (ASU) technology, although the process itself was not explicitly modeled.

Air-separated N_2 incorporation decreases the requirement for atmospheric air, as the N_2 flow rate in syngas must stay constant to maintain consistent NH_3 production. The extent of incorporation of this alternative feedstock was quantified using the Air-Separated Nitrogen Incorporation fraction (SNI), which is defined as the ratio of N_2 incorporated relative to the total N_2 supplied for NH_3 synthesis (including both direct air supply and air separation).

Economic Optimization

Taking the green H_2 and air-separated N_2 incorporations as independent parameters in this framework, the main goal of this work is to explore how the SMR process can be operated to accommodate multiple combinations of these two key parameters. Building a thorough understanding of how SMR operation can be adapted in this context is vital to map the transition from the classic HB process towards partly or fully green NH_3 manufacture, while achieving the best economic performance. This

transition will only be considered in future work.

To achieve the first goal above, the process flowsheet model was configured and optimized as a function of green H_2 and air-separated N_2 incorporations using the Optimization tool in *Aspen Plus® V14* software. The optimization considers a set of manipulated variables listed in **Table 1**. These variables translate into different stream flow rates that may be adjusted by valves and their respective control, ensuring the optimal operation under different fractions of alternative feedstock incorporation.

Table 1: List of manipulated variables considered.

Variable	Lower Bound	Upper Bound	[Unit]
Process Natural Gas	0	19	ton/hr
Process Steam	0	60	ton/hr
Process Air	0	2 000	kmol/hr
Fuel Natural Gas	0	5	ton/hr
Combustion Air	0	140	ton/hr
CO_2 Removal Solution	0	150	ton/hr
Ref-I Bypass Fraction	0	1	-

Table 2: List of constraints considered.

Equality Constraint	Value	[Unit]
Steam/Carbon in Ref-I	3	mol/mol
H2 in Syngas Stream	3 000	kmol/hr
Combustion Air to Fuel	20	ton/ton
H2 to N2 at Syngas Stream	3	mol/mol
Inequality Constraint	**Value**	**[Unit]**
CO2 Content after Removal	≤ 0.01	%mol
Q (Ref-I) + Q (Fired-Box)	≤ 0	MW
Ref-II Outlet Temperature	≤ 921	ºC

In earlier works [5, 6], the authors introduced a Ref-I Bypass to enable the redistribution of CH_4 reforming reactions between primary and secondary reformers. They found that operating with the Ref-I Bypass can increase the fraction of green H_2 while keeping the Ref-II outlet temperature stable. The overheating problem in the Ref-II unit was identified in the same study when the Ref-I Bypass was not used and was observed to worsen as the GHI fraction rises. Due to its demonstrated effectiveness, the Ref-I Bypass Fraction is listed as a variable in **Table 1**.

The optimal solution must also satisfy the operational constraints listed in **Table 2**. A constraint is included to avoid overheating of Ref-IIThe problem is subsequently formulated to minimize the operational costs associated with the SMR section. For that purpose, a set of economic parameters (**Table 3**) is employed to estimate the costs of raw material streams (including alternative feedstocks), electricity consumption, thermal energy use and production, as well as taxes related to CO_2 emissions.

Table 3: Main economic parameters considered.

Economic Parameter	Value	[Unit]
Natural Gas	490	$/ton
Process Steam	37	$/ton
Electricity	77	$/MWh
CO_2 Emission Tax	83	$/ton
Green H_2	5 000	$/ton
Air-Separated N_2	50	$/ton

The objective function integrates all economic factors to minimize the Operational Net Cost (ONC), Equation (6).

$$\text{minimize } f(x) = ONC \tag{6}$$

The optimization problem was then solved, incorporating both alternative feedstocks from 0 to 40 % (0 < GHI < 40% and 0 < SNI < 40%) in intervals of 10%.

RESULTS

The Base Case

The base case is defined as the SMR operation with no alternative feedstock incorporated. This simulates a conventional NH_3 plant that operates entirely on natural gas and atmospheric air (GHI = SNI = 0%), with a non-adapted configuration (Ref-I Bypass Fraction = 0). **Table 4** lists some important consumption rates and indicators obtained for the base case.

Table 4: Results obtained for the base case (GHI = SNI = 0 %).

	Value	[Unit]
Natural Gas Consumption	20.8	ton/hr
Atmospheric Air Consumption	39.2	ton/hr
Ref-II Outlet Temperature	920	ºC
CO_2 Emissions	54.5	ton/hr
ONC	14.5	k$/hr
CO_2 Emission per kg of H_2	9.1	kg/kg
ONC per kg of H_2	2.4	$/kg

Alternative Feedstock Incorporation

From the base case, the incorporation of alternative feedstock was considered. Specifically, the incorporation of green H_2 raises concerns about overheating in the Ref-II unit [6]. **Figure 2** shows the outlet stream temperature

of this unit, demonstrating compliance with its constraint.

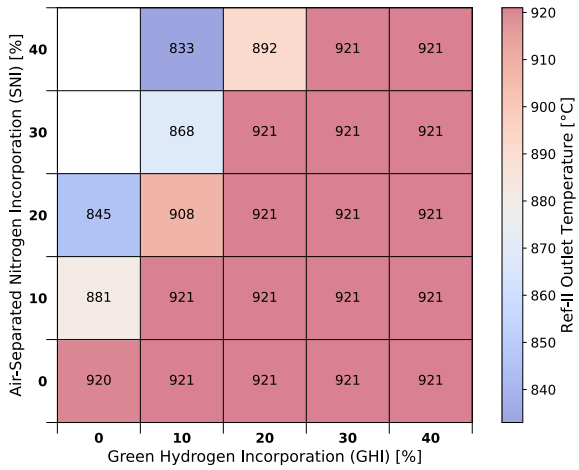

Figure 2: Matrix of Ref-II Outlet Temperature as a function of GHI and SNI (0 - 40 %; blank cells indicate no feasible solution was found).

As air-separated nitrogen incorporation increases, the results show that it has the opposite effect on Ref-II outlet temperature than hydrogen incorporation. This is especially evident when no green H_2 is considered (GHI = 0 %) and suggests that its incorporation may help control this temperature, thereby enabling higher green H_2 incorporation than previously reported [6]. Although the incorporation of green H_2 tends to increase the Ref-II outlet temperature, this temperature is kept constant at the base case value by activating the Ref-I Bypass. Therefore, the effect of increasing green H_2 is reflected not in overheating of this unit but in an increasing Ref-I Bypass fraction. **Figure 3** shows the Ref-I Bypass fraction required to comply with that limit.

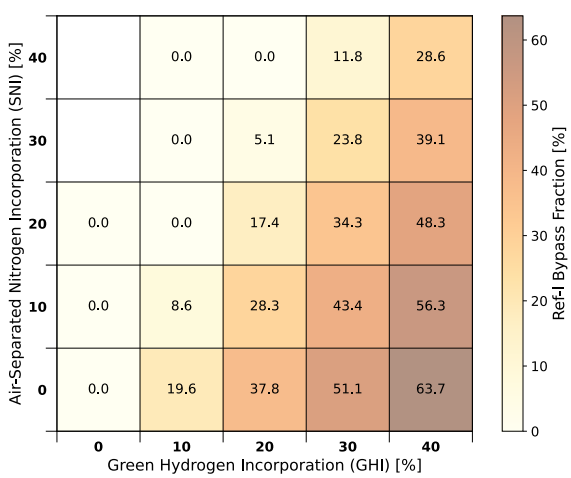

Figure 3: Matrix of Ref-I Bypass Fraction as a function of GHI and SNI (0 - 40 %; blank cells indicate no feasible solution was found).

Although the Ref-I Bypass operation is an important step toward greater green H_2 incorporation, it naturally reduces Ref-I load, as part of it is diverted by the bypass. A previous study considered a limit of 60 % on the reduction of Ref-I nominal capacity. Applying the same criteria in this work would make it impossible to achieve GHI values of 30 % and 40 % if no air-separated N_2 is considered. As SNI increases, the required Ref-I Bypass fraction is lower, which means not reporting such a sharp reduction in Ref-I load. This way, GHI of 30 % and 40 % are possible when SNI of 10 % and 40 %, respectively, are applied. Therefore, the results show that simultaneous incorporation of both alternative feedstocks can, in fact, enable higher green H_2 incorporation, contributing to greater decarbonization and deeper natural gas savings.

Figure 4 illustrates the Total Decarbonization extent achieved through both alternative feedstock incorporation. This indicator represents the relative difference in CO_2 emissions for each case compared to the previous base case. Included in CO_2 emissions are those captured in the CO_2 Removal Section, as well as emissions from flue gas generated by burning natural gas for energy production.

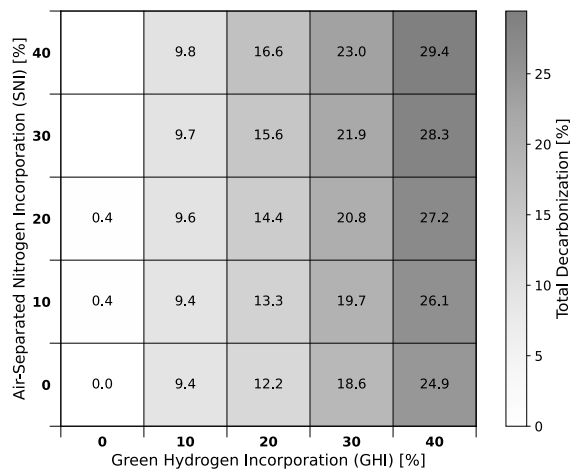

Figure 4: Matrix of Total Decarbonization as a function of GHI and SNI (0 - 40 %; blank cells indicate no feasible solution was found).

The decarbonization values achieved are notably significant, especially when green hydrogen is incorporated. These values arise from the reduced use of grey H_2 and the corresponding decrease in CH_4 required for both reforming and combustion, thereby substantiating the reported decarbonization levels. With respect to air-separated nitrogen, its influence is minor at low green hydrogen incorporation fractions (GHI < 10%), but it can contribute substantially to additional decarbonization at higher GHI values.

The reduction in natural consumption is quantified through Natural Gas Savings, which measures the

savings in natural gas (both as feedstock and fuel) reported relative to the base case. The results are presented in **Figure 5**.

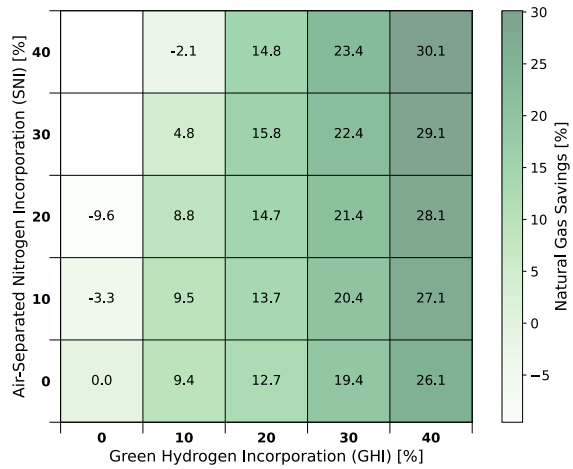

Figure 5: Matrix of Natural Gas Savings as a function of GHI and SNI (0 - 40 %; blank cells indicate no feasible solution was found).

Reducing natural gas consumption is a crucial consideration for the plant's economics, particularly given that, in the base case, natural gas costs account for approximately 70 % of total operational expenses. The economics obtained are reported in **Figure 6**.

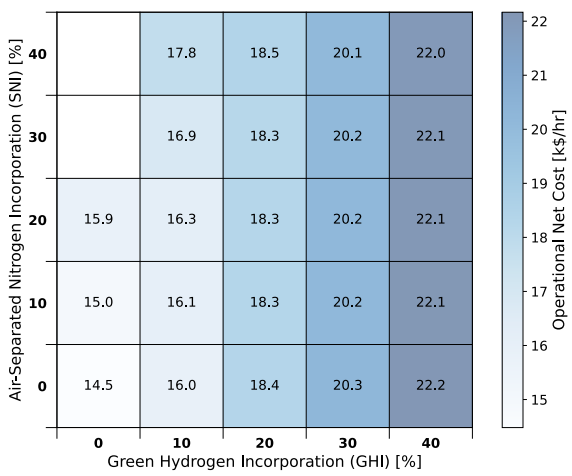

Figure 6: Matrix of ONC as a function of GHI and SNI (0 - 40 %; blank cells indicate no feasible solution was found).

Combining the effects of natural gas savings and the reduction in CO_2 emission taxes enabled by the extent of decarbonization, it would be expected that a cheaper operation would result as alternative feedstock incorporation is considered. Nevertheless, significant increases in operational costs are reported due to the cost of alternative feedstock considered. This is especially noted for green H_2, whose price is way higher than that of air-separated N_2. For instance, in the most decarbonized scenario, with GHI and SNI at 40 %, green H_2 expenses account for more than 50 % of total operating costs, while air-separation accounts for less than 3 %. Green H_2 may only be economically attractive once its production cost falls to around 2.4 $ per kilogram, at which point it becomes competitive with grey H_2 production.

CONCLUSIONS AND FUTURE WORK

The incorporation of green H_2 and air-separated N_2 into a conventional NH_3 plant was investigated using a model developed in *Aspen Plus® V14*. Incorporating green H_2 makes a greater contribution to natural gas savings and to the extent of SMR process decarbonization. Although its contribution is limited, air-separated N_2 can be significant in enabling broader control of the Ref-II outlet temperature.

Building on the previously proposed Ref-I Bypass operation, it was shown that higher GHI fractions can be achieved by considering a minimum fraction of air-separated N_2.

Despite technical viability, incorporating both alternative feedstocks (H_2 and N_2) increases SMR operational costs, mainly due to the cost of green H_2 considered. The threshold could be reached with a green H_2 cost of around 2.4 $ per kilogram, at which point it becomes competitive with grey H_2 production.

As these are preliminary results, future research may include exploring alternative feedstock compositions, varying injection points, or modifying the flowsheet to facilitate more flexible operation of the SMR process under the proposed conditions and to achieve greater decarbonization. Given the substantial costs involved, subsequent studies might also incorporate an economic analysis of green H_2 production costs to assess their impact on the economics of retrofitted NH_3 plants.

ACKNOWLEDGEMENTS

Project co-financed by the PRR – Recovery and Resilience Plan by the European Union, in particular, the Mobilizing Agenda “Moving2Neutrality” (Project no. 32, with reference no. C644927397-00000038), by the Fundação para a Ciência e a Tecnologia, I.P/; MCTES through national funds PIDDAC–CERENA (FCT-UIDB/04028/2025 and FCT-UIDP/04028/2025).

AUTHOR IDENTIFIERS

Author ORCIDs:
Fortunato, João: 0009-0002-6251-6427
Narciso, Diogo A.C.: 0000-0002-2776-7477
Matos, Henrique A.: 0000-0001-8128-7346

REFERENCES

1. Smith C, Hill AK, Torrente-Murciano L. Current and future role of haber–bosch ammonia in a carbon-free energy landscape. Energy Environ. Sci. 13:331-344 (2020). https://doi.org/10.1039/c9ee02873k
2. Rouwenhorst, K.; Castellanos, G.; International Renewable Energy Agency.; Ammonia Energy Association. *Innovation Outlook : Renewable Ammonia*; 2022; ISBN 9789292604233.
3. IEA International Energy Agency *Ammonia Technology Roadmap - Towards More Sustainable Nitrogen Fertiliser Production*; 2021.
4. Narciso DAC, Pires JM, Fortunato J, Teixeira P, Castro PM, Pinheiro CIC, Matos HA. Design and operation of power-to-ammonia systems: a review. Energy Conversion and Management 327:119494 (2025). https://doi.org/10.1016/j.enconman.2025.119494
5. Fortunato J, Castro PM, Narciso DAC, Matos HA. Towards the decarbonization of a conventional ammonia plant by the gradual incorporation of green hydrogen. Systems and Control Transactions 4:264-270 (2025). https://doi.org/10.69997/sct.104649
6. Fortunato, J.; Narciso, D.A.C.; Matos, H.A. Towards the Decarbonization of Ammonia Production through the Gradual Incorporation of Green Hydrogen (in Press). *Ind Eng Chem Res*.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Hybrid Modeling of Wastewater Treatment Dynamics Using Hammerstein-Wiener Structures

Arne Tirez[a], Niels Stevens[a], Dominik Bongartz[a], and José Matias Assumpcao[b]

[a] KU Leuven, Department of Chemical Engineering, Leuven, Flemish-Brabant, Belgium
[b] KU Leuven, Department of Chemical Engineering, Sint-Katelijne-Waver, Flemish-Brabant, Belgium

ABSTRACT

The zero-pollution ambition of the European Union requires improvements in wastewater treatment to meet increasingly stringent regulations at achievable cost. One promising approach consists in model-based optimal control. However, wastewater treatment plants involve highly nonlinear and time-varying processes, making existing mechanistic models such as the Benchmark Simulation Model no.1 (BSM1) challenging for direct use in online control. Therefore, this study explores a hybrid modeling approach using the Hammerstein-Wiener (HW) structure. The proposed model combines a mechanistic steady-state model, derived from BSM1, with a data-driven approximation of the system dynamics, incorporating low-order linear dynamic models. In this work, the HW model was used as a surrogate for BSM1. The HW surrogate model attained coefficients of determination (R^2) often exceeding 0.95 across key water quality indicators, such as total nitrogen and ammonium concentration. This accuracy was found to be substantially better than that of a First-Order-Plus Dead-Time model as conventional low-order linear model, demonstrating the usefulness of the HW structure in capturing key nonlinearities.

Keywords: Dynamic Modelling, Modelling, Modelling and Simulations, Wastewater, System Identification

INTRODUCTION

New regulation is imposing increasingly stringent requirements on wastewater treatment to tackle emerging sources of urban pollution and address societal challenges like the energy transition and climate neutrality [1]. One approach towards meeting these requirements consists in model-based methods, and in particular model-based optimal control. However, accurate and computationally tractable modeling of wastewater treatment dynamics remains a challenge.

In the 80s, the Activated Sludge Model no.1 (ASM1) became the first accepted mechanistic model across the industry, giving insights into the dynamic behavior of the carbon and nitrogen removal of activated sludge systems. Several versions have since been developed to target specific biological processes, such as phosphorus recovery and anaerobic digestion [2].

Due to the success of the ASM1 and its variants, mechanistic models are still used in wastewater treatment plants (WWTPs) because of the process insight and predictive capabilities they provide in plant design and operation optimization. The Benchmark Simulation Model No. 1 (BSM1) [3] is the industry standard for benchmarking and control application. Conversely, dynamic black-box models (a.k.a. data-driven models) are the most effective in the presence of a high volume of data. However, they have limited interpretability, large data requirements, and are unreliable in extrapolation [4].

A promising intermediate approach is hybrid modeling, which has already been successfully applied to certain aspects of wastewater treatment where some parameters, such as micropollutant concentration, colloid charge, and pathogen abundance, are hard or expensive to measure with sufficient precision [5]. By using data-driven surrogates for these complex variables, hybrid models offer a robust solution for real-time monitoring and process control. For instance, Gaublomme et al. [6] used hybrid models for solution-diffusion processes in membrane processes. Wang et al. [7] created an integrated framework to improve forecasting of salinity and dissolved oxygen with hybrid models. Nordstrand & Dutta [8] achieved a more robust and flexible control of capacitive-deionisation desalination processes by applying the

hybrid linear-state-space dynamic Langmuir model. Tomperi et al. [9] used a series-connected ARMAX model to estimate effluent COD concentration in full-scale industrial WWTPs.

Despite the fact that current hybrid models offer significant predictive capacity, they often face a trade-off between computational complexity and integration into real-time control and optimization frameworks. A critical gap remains in identifying hybrid structures that are complex enough to capture wastewater dynamics but simple enough to be used efficiently in mathematical optimization [10–12].

In this work, we address this by utilizing Hammerstein–Wiener (HW) models. These models constitute a class of block-structured dynamic models that represent systems as a linear dynamic block embedded between two nonlinear static transformations [13]. Specifically, we aim at using a mechanistic steady-state model as nonlinear static block, combined with a data-driven low-order dynamic block. We hypothesize that this hybrid model structure is well-suited for WWTP for two reasons: First, many dynamics observed in WWTP pilot plants can be effectively described as first-order models [2]. Second, by including a mechanistic model as static block, we can capture nonlinearities in the gains and represents steady states accurately. In this work, we use BSM1 both to derive the static block, and as a high-fidelity representation of the true system (cf. Figure 1).

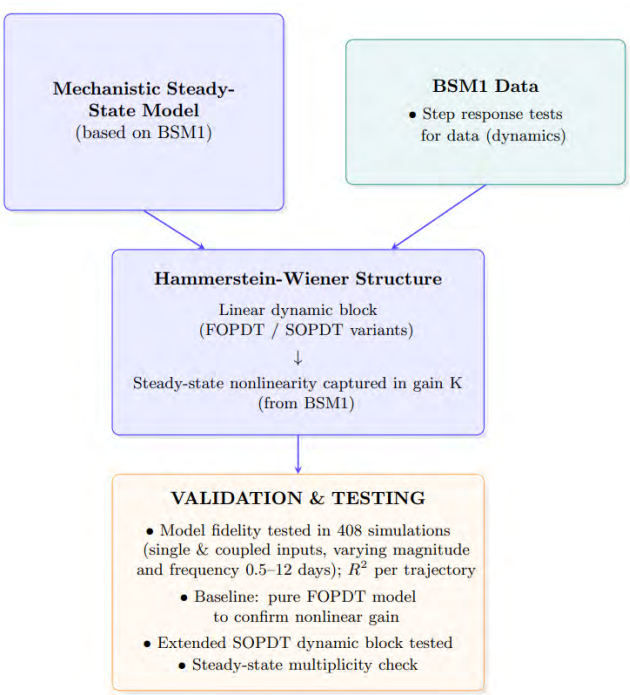

Figure 1. Simplified Flowsheet of Methodology

Our contribution lies in the development of a hybrid HW model that aims to strike a balance between accuracy and computational efficiency for real-time applications. In the following, we first describe the structure of the model as well as the experiments used for fitting the data-driven blocks. Then we evaluate the accuracy of the model in a number of validation experiments. Finally, we compare it to a first-order-plus-dead-time model as conventional, linear model to assess the benefit of retaining the nonlinear mechanistic part.

METHODOLOGY

BSM1 Model

The base model configuration of BSM1 consists of an activated sludge unit for biological wastewater treatment, followed by a circular secondary settler for separating treated water from sludge. The model consists of 205 differential equations and a minimum of 65 algebraic equations, depending on the specific implementation, and has 28 individual control handles [14]. The BSM1 reflects the standard design of many real-world secondary treatment stages in WWTPs. BSM1 provides standardized simulation environments to evaluate control strategies in SIMULINK.

In this work, we use the BSM1 in two ways: 1. We use a steady-state version of it within the HW structure as described below. 2. We use the full dynamic BSM1 as virtual plant, i.e., to generate data for the identification of the dynamic part of the HW structure, and as well as the ground truth for evaluating the accuracy of the HW model.

Hammerstein-Wiener Model

The Hammerstein-Wiener model is commonly used in nonlinear dynamic system identification. As illustrated in Figure 2, it is constructed from a linear block between two static nonlinear transformations.

Figure 2. Hammerstein-Wiener Structure

The formulation for a Hammerstein-Wiener model is presented in Equation (1)-(4) provided by Wills et al. [13]:

$$w(t) = f_H(u(t)), \forall t \in [t_0, t_f] \quad (1)$$

$$\dot{x}(t) = Ax(t) + Bw(t), \forall t \in [t_0, t_f] \quad (2)$$

$$z(t) = Cx(t) + Dw(t), \forall t \in [t_0, t_f] \quad (3)$$

$$y(t) = f_W(z(t)), \qquad \forall t \in [t_0, t_f] \quad (4)$$

Here, $f_H: \mathbb{R}^{n_u} \to \mathbb{R}^{n_w}$ represents the Hammerstein function for the input nonlinearity and $f_W: R^{n_z} \to R^{n_y}$ the

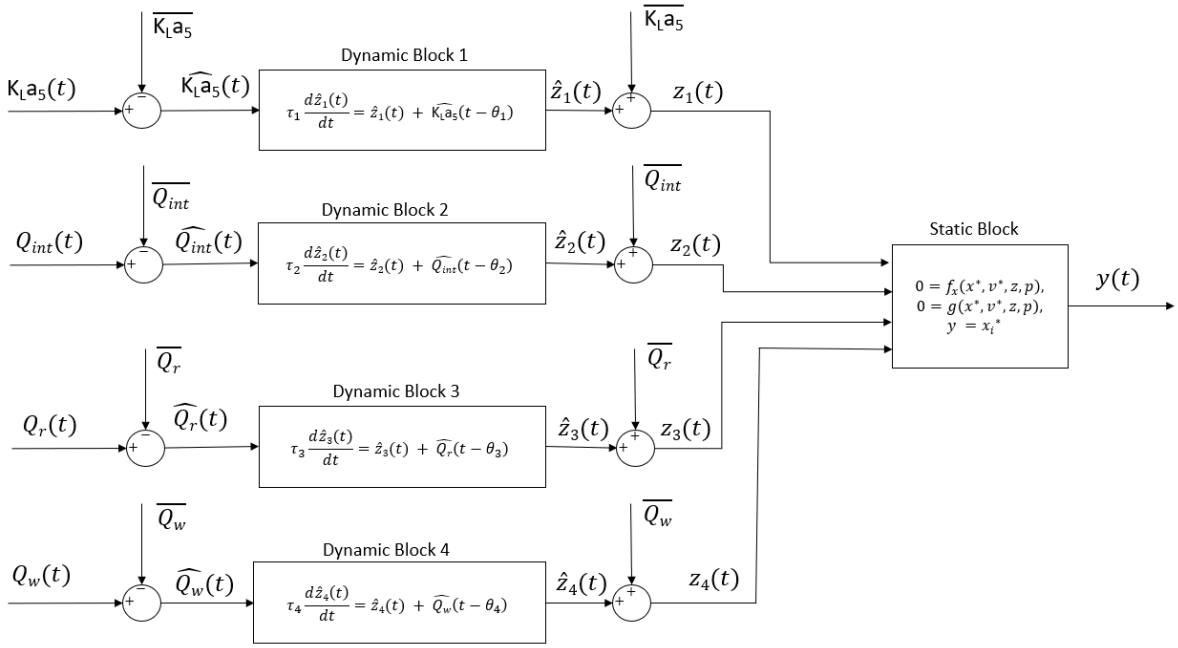

Figure 3. Hammerstein-Wiener surrogate model. $y(t)$ can be any of the 17 outputs in Table 1; the same model structure must be run separately for each output

Wiener function for the output nonlinearity. The inputs of the system are represented as $u: [t_0, t_f] \to \mathbb{R}^{n_u}$, while the inputs to the linear time-invariant (LTI) system are denoted by $w: [t_0, t_f] \to \mathbb{R}^{n_w}$. The system states are given by $x: [t_0, t_f] \to \mathbb{R}^{n_x}$, and the outputs of the LTI system are described by $z: [t_0, t_f] \to \mathbb{R}^{n_z}$. Finally, the overall system outputs are expressed as $y: [t_0, t_f] \to \mathbb{R}^{n_y}$.

Hammerstein-Wiener models can also be observed as two simpler variants: the Hammerstein model uses a nonlinear-then-linear (N-L) structure, and the Wiener model follows a linear-then-nonlinear (L-N) structure.

Hybrid modeling Approach

The HW model developed is inspired by the Wiener structure (linear-then-nonlinear) and takes 4 inputs: aeration intensity in the 5th aeration tank (K_La_5), internal recycle flow (Q_{int}), sludge recycle flow (Q_r), and waste sludge flow (Q_w). This selection was based on the most common control handles of a WWTP. Furthermore, aeration intensity and internal recycle flow have a built-in PID controller in BSM1, which provide a ground for control comparison. The modeled outputs are the effluent variables of secondary clarifier in BSM1, which can be found in Table 1. The resulting HW model architecture is shown in Figure 3.

Note that, to capture the distinct system dynamics of each variable, we developed a series of Multi-Input Single-Output (MISO) structures, where a dedicated HW model is trained for each specific effluent output. This is because the time constants τ_i and dead times θ_i of the FOPDT models in Equation (7) only captures the relation between a single input and output. Therefore, different outputs will require different parameter values. For example, ammonium concentrations and active heterotrophic biomass are affected differently by an increase in aeration intensity.

Parallel Linear Dynamic Blocks

The input entering each linear dynamic block i is defined as a deviation form by subtracting the nominal operating point from the actual input value:

$$\hat{u}_i(t) = u_i(t) - \overline{u}_i, (6)$$

where $\hat{u}_i(t)$ is the deviation input, the actual input $u_i(t)$, and nominal operating point $\overline{u}_i$. The linear block models only the time-dependence of the effect of input u_i on the output y, and consists of a FOPDT model with unit gain:

$$\tau_i \frac{d\hat{z}_i(t)}{dt} = -\hat{z}_i(t) + \hat{u}_i(t - \theta_i), \qquad (7)$$

where $\hat{z}_i(t)$ is the internal state (output) of the ith block, $\hat{u}_i(t)$ is the input, τ_i is the time constant, and θ_i is the dead time. The output $\hat{z}_i$ thus represents the same physical quantity as $\hat{u}_i$, but delayed in a way that represents the response *dynamics* of the overall output of the HW model (i.e., the output of the nonlinear block, y).

The parameters of each dynamic block i are only the time constant τ_i and time delay θ_i, which are determined from multiple step responses and validated against response signals from the ground truth system, where in each case the best fitting parameters are chosen based

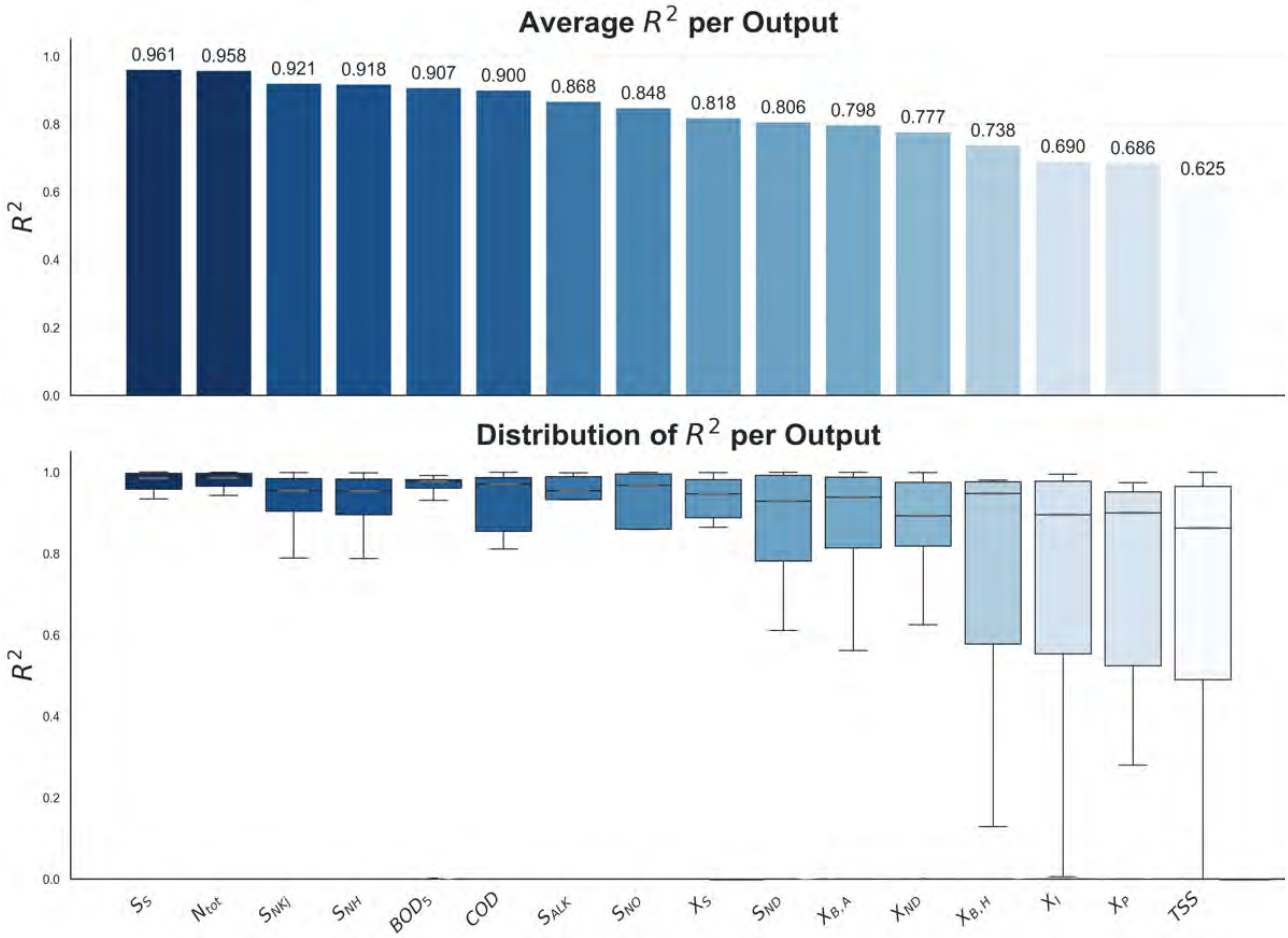

Figure 4: Predictive performance of the HW model across water quality outputs. Averages and distributions are computed over all time points of the 408 validation simulations.

on the coefficient of determination (R^2). The underlying assumption is that τ_i and θ_i stay rather constant across a wide operating range (i.e., that the main nonlinearity is in the gains). Moreover, it is assumed that the interaction between inputs in terms of dynamics can be ignored (superposition principle).

Finally, the operation point value is added to the dynamic block output to finalize the signal that flows into the nonlinear block:

$$z_i(t) = \hat{z}_i(t) + \overline{z}_i \quad (8)$$

where, $z_i(t)$ is the nominal term, $\hat{z}_i(t)$ is the deviation, and $\overline{z}_i = \overline{u}_i$ is the nominal operating point.

This modeling approach was used to enable dynamic prediction without introducing excessive model complexity, and much smaller data requirements than full black-box models. These smaller data requirements are enabled by the superposition principle of linear differential equations.

Nonlinear Static Block

This block consists of a steady-state model derived from the BSM1 equation set. Since the original BSM1 is implemented in Simulink, which limits the efficient computation of steady-states, we developed a standalone version in CasADi using the equations from [14]. This new model version allows for the calculation of steady states at a significantly lower computational cost. Note that, because the inputs are transformed based on their dynamic relationship with the output (captured via time constants and dead times), the steady-state block's solution represents a quasi-steady state, meaning the approximate response of the output to time-dependent input changes. The HW model is considered to be at a true steady state only when $u(t) = z(t)$, indicating the system has fully settled following input perturbations.

To compute the steady states, the BSM1, which is a nonlinear DAE, was converted into an algebraic system as follows. Let $x(t) \in \mathbb{R}^{n_x}$ denote the differential states, $v(t) \in \mathbb{R}^{n_v}$ the algebraic states, $z(t) \in \mathbb{R}^{n_z}$ the inputs, for which we here use the inputs transformed by the linear blocks rather than the actual external inputs to the system, and $p \in \mathbb{R}^{n_p}$ the model parameters. The BSM1 equations are

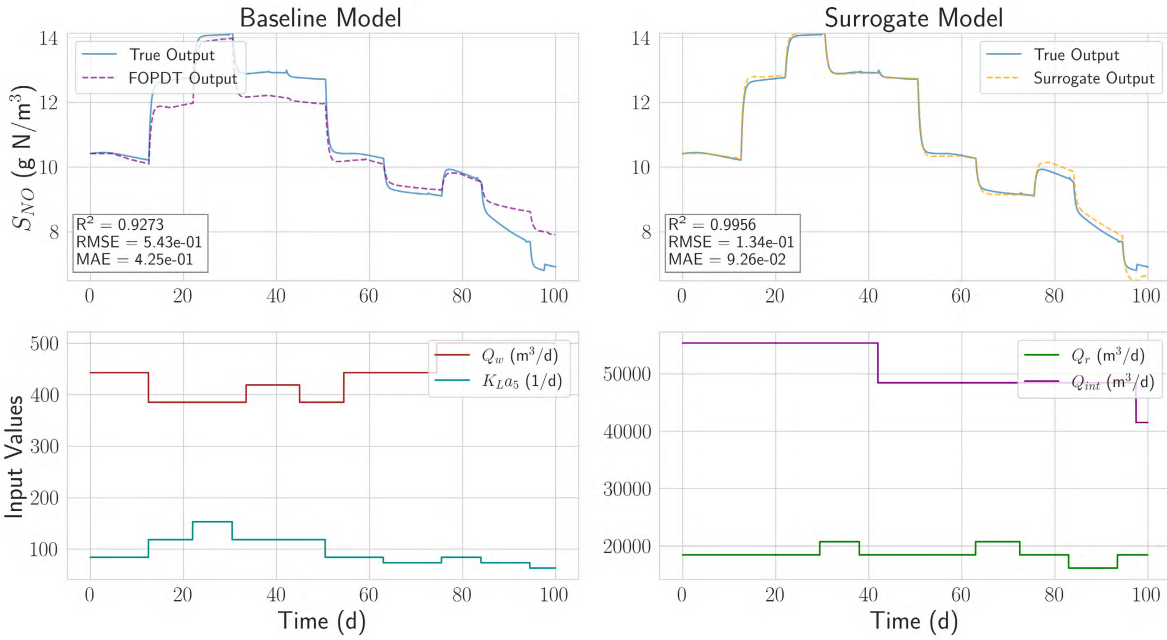

Figure 5. Comparison of a FOPDT baseline model (top left, dashed purple) against the proposed HW smodel (right, dashed orange) for nitrate and nitrite concentration, both in comparison to the BSM1 as ground truth (solid blue). The inputs are shown in the bottom left and right plots. Despite being shown in different figures, they were applied simultaneously to the models but were plotted on different graph because of the order of magnitude for the inputs.

$$\dot{x}(t) = f_x(x(t), v(t), z(t), p), (9)$$

$$0 = g(x(t), v(t), z(t), p), (10)$$

where f_x defines the differential equations (bioreactor, clarifier, and sensors), and g defines the algebraic equations (clarifier mappings and output connections). Computing the steady states reduces to solving the system

$$0 = f_x(x^*, v^*, z, p), (11)$$

$$0 = g(x^*, v^*, z, p), (12)$$

where $(x^*,\ v^*)$ represent the steady-state values for the differential and algebraic variables, respectively. The effluent concentration i can be extracted from x^*:

$$y = x_i{}^* \ (15)$$

The fact that only one output can be extracted for solving the steady state model (as described above) is a limitation of the proposed structure.

Model Training and Validation

All data was sampled at 15 minute intervals. This was an upper limit of sampling time for the regression to capture the fastest dynamic in the system. For the training data, step responses were applied to the ground truth BSM1, starting from different operating points. Given the linear dynamic models used, all inputs were varied separately to fit the respective blocks. For each input, 20 points evenly spaced across the feasible range of input values were used as initial operating points as well as final points of the input steps. From these 20 step responses per input, a FOPDT model was fitted to every input-output combination through least-squares regression, resulting in 20 parameters sets (K_i, τ_i, θ_i) per output (i.e., one per input for every output). The gains K_i were then discarded, since this they are not used in the HW model.

The validation dataset was generated utilizing a Pseudorandom Multilevel Sequence (PRMS) with Pseudorandom Frequency Modulation (PRFM) as input signal. The multilevel amplitudes varied across different experiments: some only covered the training space, while other extrapolated beyond the training space while remaining within the physical operating bounds of the plant. . The frequency modulation was varied from a $1/(0.5\ \text{days})$ to $1/(12\ \text{days})$ across experiments. This was designed to validate all aspects of the model, including extrapolation capabilities, accuracy under fast and slow dynamics, and ability to predict steady states. The data was generated with the same sampling rate as the training data. In total, 408 output trajectories were generated by applying the described inputs signals.

Model validation was performed in two steps; first the best among the 20 best parameter sets per input-output combination was selected based on the coefficient of determination (R^2), which was computed along

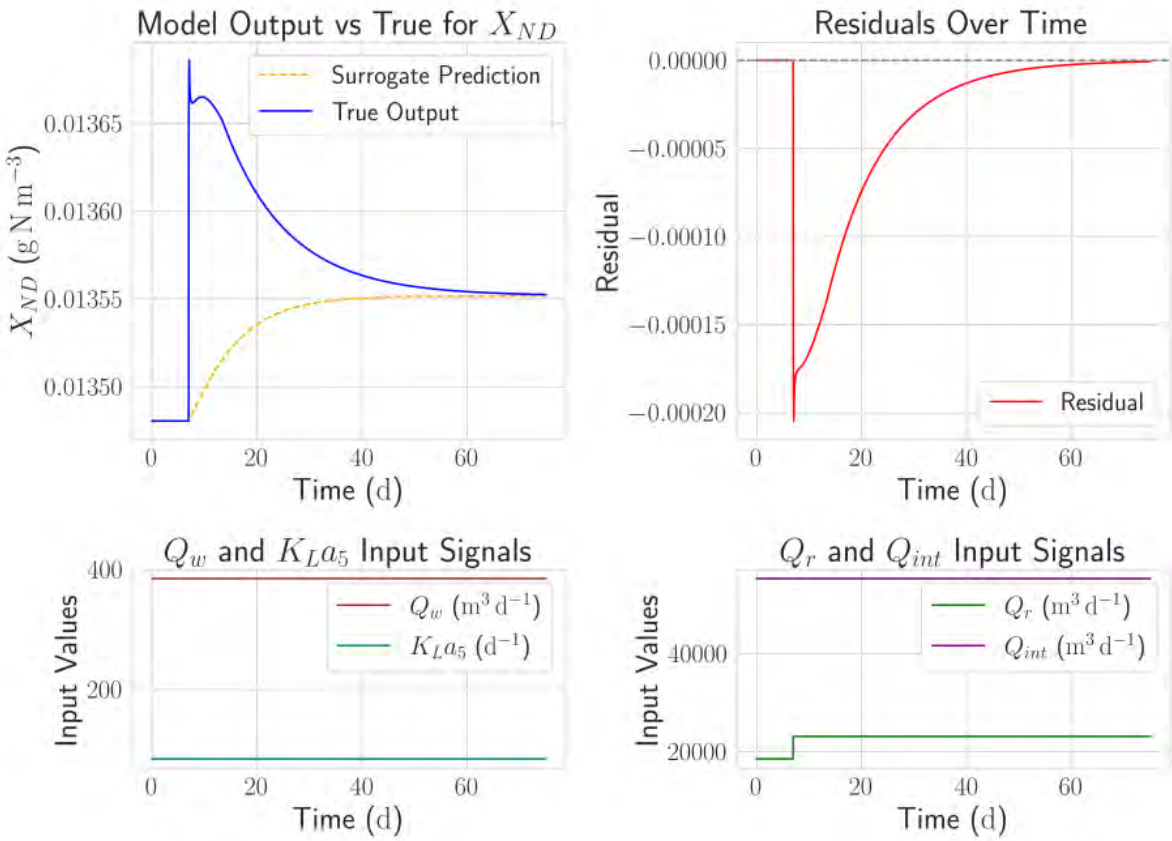

Figure 6. The particulate biodegradable organic nitrogen concentration cannot be captured by the proposed HW model with first-order dynamics because of its complex transient behavior.

the output trajectory. Afterwards, these best parameters were specified in the model.

In the second step of model validation (considered evaluation), the HW model's fidelity was tested. For each simulation the coefficient of determination (R^2) was captured for the whole trajectory of the simulation.

To evaluate the benefit of capturing the nonlinear gain through the mechanistic model in the HW structure, a pure FOPDT model was also developed as conventional low-order linear model. It was fitted to a step response where the final value was the upper bound of the training input range. As an extension, a HW model with a Second Order Plus Dead Time (SOPDT) function in its dynamic blocks was designed to consult test if an extra parameter in the structure leads to better predictive capabilities.

Although the BSM1 exhibits a unique steady state, the nonlinear blocks of the HW structure could introduce artificial steady-state multiplicity. To verify the reliability of the proposed model, we performed a multiplicity analysis using Latin Hypercube Sampling (LHS) to generate different initial guesses to the steady-state solver, ensuring the HW model maintains a unique, physically consistent mapping across the operational space.

RESULTS & DISCUSSION

Across all 408 validation simulations, the average R^2 for the outputs, as seen in Figure 4, achieved consistently high values, often exceeding 0.9, indicating accurate replication of both steady-state and dynamic behavior. The most important results from these experiments were the good predictive performance ($R^2 > 0.9$) across key water quality indicators such as readily biodegradable substrate (S_s), total nitrogen (N_{tot}), Kjeldahl nitrogen (S_{NKj}), and ammonium (S_{NH}), as well as Biochemical Oxygen Demand (BOD_5) and Chemical Oxygen Demand (COD). Lower-performing outputs were linked to specific limitations, particularly input-output relationships exhibiting complex nonlinear dynamics beyond FOPDT approximation. An example for such a case is shown in Figure 6.

The HW model achieved substantially better accuracy than the pure FOPDT model in the validation simulation: the HW model consistently attained higher

coefficients of determination (R^2) for the whole trajectory of a simulation, often exceeding 0.90 whereas the FOPDT model yielded noticeably lower R^2 values. An example is shown in Figure 5. Preliminary experiments indicated that extending the HW model's linear blocks to second-order (SOPDT) dynamics yielded negligible improvements in predictive accuracy over the first-order structure.

Steady-state multiplicity analysis confirmed the presence of multiple equilibria in roughly 10% of tested conditions. These multiplicities were concentrated in specific input regions, most notably when the wastage sludge flow Q_w dropped below $300\ \mathrm{m}^3 d^{-1}$. This confirms that in operating regimes where multiplicity is present the HW model is not reliable.

Comparing the HW model to traditional linear hybrid structures, such as ARMAX, highlights the potential advantages of the our approach. For instance, while an ARMAX model achieved an $R^2 \approx 0.67$ for effluent COD in a specific industrial WWTP [9], our results suggest that augmenting low-order dynamics with nonlinear static map can further enhance predictive capabilities by capturing the system's inherent nonlinearities more effectively.

CONCLUSION

This work demonstrates that a hybrid Hammerstein–Wiener model can effectively approximate the dynamics of key outputs in a wastewater treatment process by combining mechanistic steady-state modeling with data-driven linear dynamic blocks. This separation enables reasonably accurate dynamic prediction without introducing excessive model complexity, and much smaller data requirements than full black-box models. Moreover, the HW model retains the gain of the original rigorous model which is useful for control and optimization applications. The linear dynamics of the HW model can be an advantage also because they can be exploited algorithmically in optimization algorithms [15]. These properties make the Hammerstein–Wiener framework particularly well suited for integration into DRTO or NMPC and allow operation with online optimization crucial for WWTPs that operate under increasingly strict regulatory constraints.

Table 1: Inputs and outputs of Hammerstein-Wiener model

Symbol	Description	Unit
K_{La5}	Oxygen transfer coefficient in reactor	d^{-1}
Q_{int}	Internal recycle flow rate	$\mathrm{m}^3\,\mathrm{d}^{-1}$
Q_r	Return sludge flow rate	$\mathrm{m}^3\,\mathrm{d}^{-1}$
Q_w	Waste sludge flow rate	$\mathrm{m}^3\,\mathrm{d}^{-1}$
N_{tot}	Total nitrogen in effluent	$\mathrm{g\,m}^{-3}$
$\mathrm{COD}_{\mathrm{tot}}$	Total chemical oxygen demand in effluent	$\mathrm{g\,m}^{-3}$
S_{NH}	Ammonium nitrogen in effluent	$\mathrm{g\,m}^{-3}$
TSS	Total suspended solids	$\mathrm{g\,m}^{-3}$
BOD_5	5-day biochemical oxygen demand in effluent	$\mathrm{g\,m}^{-3}$
S_I	Soluble inert organic matter	$\mathrm{g\,m}^{-3}$
S_S	Readily biodegradable substrate	$\mathrm{g\,m}^{-3}$
X_I	Particulate inert organic matter	$\mathrm{g\,m}^{-3}$
X_S	Slowly biodegradable substrate	$\mathrm{g\,m}^{-3}$
$X_{B,H}$	Active heterotrophic biomass	$\mathrm{g\,m}^{-3}$
$X_{B,A}$	Active autotrophic biomass	$\mathrm{g\,m}^{-3}$
X_P	Particulate products from biomass decay	$\mathrm{g\,m}^{-3}$
S_O	Dissolved oxygen	$\mathrm{g\,m}^{-3}$
S_{NO}	Nitrate and nitrite nitrogen	$\mathrm{g\,m}^{-3}$
S_{ND}	Soluble biodegradable organic nitrogen	$\mathrm{g\,m}^{-3}$
X_{ND}	Particulate biodegradable organic nitrogen	$\mathrm{g\,m}^{-3}$
S_{ALK}	Alkalinity	$\mathrm{mmol\,L}^{-1}$

ACKNOWLEDGEMENTS

DB gratefully acknowledges funding through Internal Funds KU Leuven (STG/22/060). JM gratefully acknowledges funding through Internal Funds KU Leuven (STG/23/057).

AUTHOR IDENTIFIERS

Author ORCIDs:
Bongartz D: 0000-0003-1790-0235
Matias Assumpcao J: 0000-0002-1094-6145

REFERENCES

1. European Parliament and Council of the European Union. Directive (EU) 2024/3019 concerning urban wastewater treatment (recast). Dec. 2024. url: https://eurlex.europa.eu/eli/dir/2024/3019/oj.
2. Olsson G. ICA and me – a subjective review. Water Research 46:1585-1624 (2012). https://doi.org/10.1016/j.watres.2011.12.054
3. Petersen B, Gernaey K, Devisscher M, Dochain D, Vanrolleghem PA. A simplified method to assess structurally identifiable parameters in monod-based activated sludge models. Water Research 37:2893-2904 (2003). https://doi.org/10.1016/s0043-1354(03)00114-3
4. Alizadeh R, Allen JK, Mistree F. Managing computational complexity using surrogate models: a critical review. Res Eng Design 31:275-298 (2020). https://doi.org/10.1007/s00163-020-00336-7
5. Schneider MY, Quaghebeur W, Borzooei S, Froemelt A, Li F, Saagi R, Wade MJ, Zhu JJ, Torfs E. Hybrid modelling of water resource recovery facilities: status and opportunities. Water Science and Technology 85:2503-2524 (2022). https://doi.org/10.2166/wst.2022.115
6. Gaublomme D, Strubbe L, Vanoppen M, Torfs E, Mortier S, Cornelissen E, De Gusseme B, Verliefde A, Nopens I. A generic reverse osmosis model for full-scale operation. Desalination 490:114509 (2020). https://doi.org/10.1016/j.desal.2020.114509
7. Wang X, Zhang J, Babovic V. Improving real-time forecasting of water quality indicators with combination of process-based models and data assimilation technique. Ecological Indicators 66:428-439 (2016). https://doi.org/10.1016/j.ecolind.2016.02.016
8. Nordstrand J, Dutta J. Flexible modeling and control of capacitive-deionization processes through a linear-state-space dynamic langmuir model. npj Clean Water 4: (2021). https://doi.org/10.1038/s41545-020-00094-y
9. Tomperi J, Sorsa A, Ruuska J, Ruusunen M. Series-connected data-based model to estimate effluent chemical oxygen demand in industrial wastewater treatment process. Journal of Environmental Management 373:123680 (2025). https://doi.org/10.1016/j.jenvman.2024.123680
10. Iratni A, Chang NB. Advances in control technologies for wastewater treatment processes: status, challenges, and perspectives. IEEE/CAA J. Autom. Sinica 6:337-363 (2019). https://doi.org/10.1109/jas.2019.1911372
11. Han HG, Zhang Y, Sun HY, Liu Z, Qiao J. Data-knowledge-driven multiobjective adaptive optimal control for wastewater treatment processes under multiple operating conditions. IEEE Trans. Cybern. 55:1056-1069 (2025). https://doi.org/10.1109/tcyb.2025.3531514
12. Wu XL, Han WH, Yang HY, Li X, Han HG. Robust soft constrained model predictive control and its application in wastewater treatment processes. IEEE Trans. Automat. Sci. Eng. 22:13198-13211 (2025). https://doi.org/10.1109/tase.2025.3552421
13. Wills A, Schön TB, Ljung L, Ninness B. Identification of hammerstein–wiener models. Automatica 49:70-81 (2013). https://doi.org/10.1016/j.automatica.2012.09.018
14. J. Alex, L. Benedetti, J. Copp, K.V. Gernaey, Ulf Jeppsson, I. Nopens, M.-N. Pons, J.-P. Steyer, and P.A. Vanrolleghem. Benchmark Simulation Model no. 1 (BSM1). Vol. TEIE-7229. Lund University, 2008.
15. Kappatou CD, Bongartz D, Najman J, Sass S, Mitsos A. Global dynamic optimization with hammerstein–wiener models embedded. J Glob Optim 84:321-347 (2022). https://doi.org/10.1007/s10898-022-01145-z

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Re-parametrisation of NRTL model for C_1^+ organics and alcohols recovery from aqueous phase in pyrolysis oil production

Matteo Gilardi[a]*, Filippo Bisotti[a], Trung Trinh[a], Bernd Wittgens[a]

[a] SINTEF Industry, Department of Process Technology, Trondheim, Trøndelag, Norway
* Corresponding Author: matteo.gilardil@sintef.no.

ABSTRACT

Pyrolysis is an emerging green pathway to produce bulk chemicals and sustainable fuels. However, pyrolysis oil requires stabilisation via hydrotreatment, and this process generates an aqueous waste containing alcohols (mainly methanol and ethanol), carboxylic acids, and some ketones. To increase the economic sustainability of biofuels production, there is increasing interest in recovering these valuable chemicals from water. Reliable thermodynamics are necessary to address the separation and design of equipment to fractionate such complex mixtures, with multiple azeotropes and non-idealities. The Non-Random Two-Liquids (NRTL) models in both Aspen Plus V12.0 and COFE V3.7, a license-free software released by AmsterChem, do not accurately reproduce the equilibrium measurements of most of the binary and ternary mixtures involving water, a C_1-C_4 alcohol, and a light carboxylic acid. This work aims to retune the activity-based model to improve the NRTL model predictivity, using experimental data retrieved from NIST ThermoData Engine V10.1. The regression is performed in Aspen Plus V12 using the Maximum Likelihood algorithm. The resulting re-parametrised NRTL shows improved accuracy and reliability compared to the existing models. For all investigated binary and ternary mixtures, the Average Relative Deviation (ARD) is lower than 20%, and for 75% of them, ARD is lower than 10%. The ARD in the prediction of vapour phase composition drops from 16.0% and 12.1% of the original COFE and Aspen Plus models, respectively, down to 7.0%.

Keywords: pyrolysis, water phase valorisation, NRTL model, VLE model re-parametrisation, Aspen Plus, COCO-COFE

INTRODUCTION

Pyrolysis oil offers a valuable solution to the circular economy, as it allows valorisation of waste, including lignocellulosic biomass and plastics. It represents a renewable fuel for transport and power generation, and it can be used as a feedstock to produce valuable chemicals [1]. However, raw pyrolysis oils are chemically and physically unstable. A stabilisation is necessary to reduce acidity and reactive components (mainly carbonyls and carboxylic acids). The stabilisation process includes the following steps [2]:

1. Selective hydrogenation at mild temperatures (200-250°C). The obtained product is called Stabilised Pyrolysis Oil (SPO);
2. Deoxygenation at high temperature (350°C). The obtained product is called Stabilised Deoxygenated Pyrolysis Oil (SDPO);
3. Final hydrogenation (refinement step at 425°C) to produce Hydrogenated Pyrolysis Oil (HPO), which meets the specifications for blending with conventional fuels.

This stabilisation process generates a substantial amount of aqueous stream, which is currently treated as waste. This water stream contains methanol and ethanol in the order of 2 wt%-6 wt%, and minor concentrations of higher alcohols, C_1-C_4 carboxylic acids, and light esters.

In FUEL-UP [3], a research project funded by the European Commission, we aim to design a process to isolate the chemicals embodied in the aqueous phase and

exploit them as intermediates to generate biofuels, biochemicals, and pharmaceutical products.

Reliable thermodynamics are necessary to design, model, and optimise a process to fractionate such complex mixtures, characterised by multiple azeotropes. Among available Vapour-Liquid Equilibrium (VLE) models, the Non-Random Two-Liquid (NRTL) activity model is often used to characterise thermodynamic equilibrium in azeotropic systems. In this context, the reliability of existing models for multi-component mixtures can be checked by verifying their accuracy in predicting the VLE of binary and ternary mixture data available in the literature.

In this work, a statistical analysis is carried out to check the reliability of default NRTL models in both Aspen Plus V12.0 [4] and COCO-COFE V3.7 [5], a license-free process simulator, in predicting vapor-liquid equilibrium data for all possible combinations of mixtures involving the 14 compounds which are present in the water stream to be valorised in concentrations equal or higher than 0.1%. For the Aspen model, research on available interaction coefficients is enabled among all existing databanks. When available, APV121 VLE-IG database is selected, in accordance with the Aspen Plus default template.

NRTL combined with ideal gas law is considered instead of more accurate gas phase models where the gas phase is described by means of an equation of state, such as NRTL-RK. This choice is done for the sake of consistency in the model accuracy comparison, as this more advanced option to combine the NRTL model with an equation of state to describe the vapor phase is not available in COCO-COFE. Using NTRL in the liquid phase and ideal gas law in the gas phase can be justified considering that water-alcohol mixtures exhibit significantly stronger non-idealities in the liquid phase compared to the vapor phase and that the target separation process is to be carried out at atmospheric pressure.

After verifying the insufficient predictability of both existing Aspen Plus and COCO-COFE models, NRTL binary interaction coefficients are re-fitted in AspenPlus for all mixtures showing an Average Relative Deviation (ARD) > 10%. The model tuned in this work is used as the thermodynamic framework for the design and optimisation of a multi-step distillation process to recover alcohols, in particular methanol and ethanol, at the highest possible purity from pyrolysis oil aqueous waste. The process design and evaluation of Key Performance Indicators for this process are discussed in a separate contribution, entitled “Process design for the recovery of valuable organic compounds from pyrolysis oil aqueous phase”.

AQUEOUS STREAM CHARACTERISATION

10 litres of both SPO and SDPO have been collected and delivered to SINTEF by the Biomass Technology Group (BTG), a company specialised in converting biomass into sustainable biofuels located in the Netherlands. The composition of these streams has been analysed in SINTEF using the Gas Chromatography-Flame Ionisation Detector (GC-FID) technique. 14 different species were identified, and their concentration was measured after calibration following the Standard Operating Procedure (SOP). Peaks of unidentified species were combined with the peaks of the closest species in the chromatogram. For simplicity, alcohols heavier than C_5 were lumped into pentanol, and carboxylic acids higher than C_4 were lumped into butyric acid.

The composition of SPO and SDPO aqueous phases (SPO_{Aq} and $SDPO_{Aq}$, respectively) is shown in Table 1. Noteworthily, the measured concentrations of the compounds detected with the GC cover 85% and 95% of the mass balance for SPO_{Aq} and $SDPO_{Aq}$, respectively.

All organics are highly diluted in water (concentration < 5 wt%). The species present in the highest concentrations are alcohols, in particular methanol, ethanol, and pentanol. All other compounds, including carboxylic acids, acetates, ketones, and furans, have individual concentrations below 0.8 wt%.

Table 1: Composition of SPO and SDPO aqueous streams (wt %).

Component	Symbol	SPO_{Aq}	$SDPO_{Aq}$
Water	H2O	72.50%	89.90%
Methanol	MEOH	4.95%	0.33%
Ethanol	EtOH	2.70%	1.85%
1-Propanol	1-PrOH	0.32%	0.55%
Butanol	BuOH	0.13%	0.03%
Formic acid	FoAc	0.05%	0.05%
Acetic acid	AcAc	0.22%	0.00%
THF	THF	0.23%	0.19%
Methyl acetate	MeA	0.25%	0.00%
Ethyl acetate	EtA	0.59%	0.49%
Acetone	Acet	0.05%	0.08%
Isopropanol	2-PrOH	0.20%	0.10%
Propionic acid	PrAr	0.23%	0.47%
Butyric acid	BuAc	0.37%	0.50%
Pentanol	PeOH	1.91%	0.87%
TOTAL	-	84.70%	95.41%

METHODOLOGY

Experimental data overview

The reliability of the existing Aspen Plus V12.0 and COFE v3.7 models is checked by verifying their accuracy

in predicting the VLE of binary and ternary mixtures data available in the literature. Experimental data for model evaluation are retrieved from NIST ThermoData Engine version 10.1, considering all possible binary and ternary combinations among the compounds listed in Table 1.

Since the target process is to be operated at atmospheric pressure, the analysis considers all isobaric data collected at atmospheric pressure and isothermal data in a pressure range 80-120 kPa for the other mixtures. In the selection process, all data failing to pass the thermodynamic consistency test are disregarded. All collected data are then plotted in Excel in the form of a Temperature-composition chart to identify and remove any remaining outliers.

A statistical analysis is performed on the remaining data to obtain insight into how many independent measurements are available for each analysed mixture and into what share of the liquid composition range is covered by experiments. An overview of experimental data availability is provided in Table 2 and Table 3. Liquid composition ranges shown in the tables include all the components present in the mixture. Data availability is significantly larger for binary mixtures.

Table 2: Statistical analysis on experimental data availability on binary mixtures from NIST ThermoData Engine V10.1.

Indicator	Value
N° of mix. with available data	76 / 105
Covered composition range	
Liquid Composition <1% - >99% (mol)	21%
Liquid Composition <5% - >95% (mol)	68%
Liquid Composition <10% - >90% (mol)	90%
Number of experimental data	
N data > 100	21%
N data > 50	43%
N data > 20	75%
N data >10	88%

Table 3: Statistical analysis on experimental data availability on ternary mixtures from NIST ThermoData Engine V10.1.

Indicator	Value
N° of mix. with available data	31/455
Covered composition range	
Liquid Composition <1% - >99% (mol)	7%
Liquid Composition <5% - >95% (mol)	23%
Liquid Composition <10% - >90% (mol)	52%
Number of experimental data	
N data > 100	19%
N data > 50	52%
N data > 20	77%
N data >10	90%

Statistical analysis for model evaluation

Average Relative Deviation (ARD) is used as the statistical indicator to evaluate the model performance, as shown in expression (1), where *z* stands for a generic property (T, P, vapor composition), and *i* stands for the generic experimental measurement.

$$ARD = \frac{1}{N_{data}} \sum_{i=1}^{N_{data}} \left| \frac{z_{i,mod} - z_{i,exp}}{z_{i,exp}} \right| \quad (1)$$

For isothermal binary data, a VLE problem at fixed temperature and liquid composition is solved. Therefore, the ARD is calculated for pressure and vapor composition. Conversely, for isobaric data, pressure and liquid composition are fixed, and ARD is checked for temperature and vapor composition. For ternary data, the system has an additional degree of freedom according to the Gibbs phase rule, therefore both temperature and pressure are assigned, and the ARD is calculated for vapor phase composition only.

For a given mixture, the model predictability is considered satisfactory if the obtained average ARD is lower than 10%. These mixtures are coloured in green. Mixtures showing ARD higher than 10% but lower than 20% are coloured in yellow. This condition indicates that the model is still able to provide a reasonable description of the system equilibrium, but there is room for improvement through coefficient regression. Finally, mixtures showing ARD > 20% are coloured in red. This condition indicates that the current model is unsuitable for VLE prediction.

Regression procedure

NRTL binary interaction coefficients are regressed in AspenPlus V12.0 using the Aspen Plus regression tool for all mixtures showing ARD > 10% in both default models. Default values from the Aspen model are kept for all the other mixtures, as the Aspen model shows better predictability compared to the COFE one.

The generic expression of activity coefficient (γ_{ij}) in the NRTL model [6] is shown in expressions (2) to (3), where *j*, *k*, and *m* represent a generic component in a mixture, and *n* is the total number of components.

τ_{ij} is defined according to expressions (4) and (5) in Aspen Plus and COCO-COFE, respectively.

$$\ln(\gamma_i) = \frac{\sum_{j=1}^{n} x_j \cdot \tau_{ji} \cdot G_{ji}}{\sum_{k=1}^{n} x_k \cdot G_{ki}} + \sum_{j=1}^{n} \frac{x_j \cdot G_{ij}}{\sum_{k=1}^{n} x_k \cdot G_{kj}} \cdot \left(\tau_{ij} - \frac{\sum_{m=1}^{n} x_m \cdot \tau{:}mj \cdot G_{mj}}{\sum_{k=1}^{n} x_k \cdot G_{kj}} \right) \quad (2)$$

$$G_{ij} = \exp(-\alpha_{ij} \cdot \tau_{ij}) \quad (3)$$

$$\tau_{ij}(\text{Aspen}) = A_{ij} + \frac{B_{ij}}{T} \quad (4)$$

$$\tau_{ij}(\text{COFE}) = \frac{g_{ij} - g_{ij,T} \cdot \left(1 - \frac{T}{T_{ref}}\right)}{T} \quad (5)$$

COFE-compliant coefficients can be easily retrieved from the corresponding Aspen-compliant coefficients

using expressions (6) and (7), where Tref is set equal to 298.15 K.

$$g_{ij} = B_{ij} + A_{ij} \cdot T_{ref} \quad (6)$$

$$g_{ij,T} = \mathrm{A_{ij}} \cdot T_{ref} \quad (7)$$

The regression is performed following the same procedure adopted in Gilardi et al. [7]. Both temperature-independent (A_{ij}) and temperature-dependent (B_{ij}) coefficients (see expression (4)) are regressed, while the non-randomness factor (α_{ij}) is fixed at a value of 0.3. The regression is performed using the Modified-Barker algorithm, which means that the system solves a VLE at fixed temperature and liquid composition for isobaric data, and at fixed pressure and liquid composition for isothermal data. Each regression is repeated for 10 iterations to ensure the global minimum is reached, rather than just a local minimum. Overall, 108 parameters are re-fitted, improving the VLE prediction for 27 different mixtures.

RESULTS

The results of the COFE and Aspen Plus models evaluation for every single mixture are gathered in Figure 1 to Figure 2, and Figure 3 to Figure 4, respectively. A summary of the model reliability assessment results is shown in Table 4.

Table 4: Summary of Key Performance Indicators from the evaluation of both COFE and Aspen Plus default models.

KPI	Original COFE	Original Aspen
ARD-composition (%)	16.01	12.08
ARD-Temperature (%)	3.02	1.53
ARD – Pressure (%)	8.27	4.24
“Green” mixtures	36	62
“Yellow” mixtures	39	31
“Red” mixtures	32	14

The average ARD among all binary mixtures is 14.7% and 9.2% for the two models, respectively, while the average ARD among ternary mixtures is 19.1% and 12.5%. The results show that the Aspen Plus model outperforms the COFE model, but there is still a relevant number of mixtures for which the ARD is >20% (14), which makes it unsuitable for accurate process modelling of waste-water stream valorisation from pyrolysis oil stabilisation.

Mixture	Comp 1	Comp 2	ARD Y	ARD T	ARD P
BIN-1	AcAc	FoAc	11.01%	0.79%	not applicable
BIN-2	BuOH	AcAc	21.11%	2.34%	9.80%
BIN-3	BuOH	H2O	56.29%	6.24%	not applicable
BIN-4	EtOH	AcAc	16.69%	3.12%	2.67%
BIN-5	EtOH	BuOH	5.60%	0.44%	not applicable
BIN-6	EtOH	H2O	14.83%	2.51%	9.21%
BIN-7	EtOH	MeOH	2.47%	0.48%	5.46%
BIN-8	EtOH	1-PrOH	1.35%	0.34%	not applicable
BIN-9	H2O	AcAc	5.44%	1.75%	not applicable
BIN-10	H2O	FoAc	27.32%	3.97%	not applicable
BIN-11	H2O	1-PrOH	38.70%	9.60%	27.81%
BIN-12	MeOH	AcAc	13.21%	2.40%	not applicable
BIN-13	MeOH	BuOH	10.70%	6.52%	not applicable
BIN-14	MeOH	H2O	7.18%	1.20%	4.02%
BIN-15	MeOH	1-PrOH	6.89%	0.40%	not applicable
BIN-16	1-PrOH	AcAc	15.48%	3.10%	8.81%
BIN-17	1-PrOH	BuOH	5.11%	0.27%	not applicable
BIN-18	EtA	H2O	22.83%	not applicable	4.39%
BIN-19	EtA	THF	13.07%	0.06%	not applicable
BIN-20	MeA	H2O	72.44%	not applicable	34.49%
BIN-21	THF	H2O	33.59%	10.40%	not applicable
BIN-22	ACET	PROP	10.07%	not applicable	4.99%
BIN-23	ACET	AcAc	23.74%	not applicable	5.55%
BIN-24	ACET	FoAc	21.38%	11.08%	not applicable
BIN-25	THF	BuOH	4.83%	2.04%	not applicable
BIN-26	MeA	PROP	16.51%	not applicable	13.10%
BIN-27	MeA	BuOH	15.60%	5.68%	not applicable
BIN-28	MeA	FoAc	32.74%	1.37%	not applicable
BIN-29	MeA	AcAc	10.74%	0.89%	not applicable
BIN-30	EtA	PROP	14.30%	4.06%	not applicable
BIN-31	EtA	BuOH	13.84%	5.43%	not applicable
BIN-32	EtA	AcAc	22.04%	0.80%	not applicable
BIN-33	ACET	H2O	13.86%	not applicable	9.14%
BIN-34	ACET	MeOH	2.06%	not applicable	3.85%
BIN-35	ACET	EtOH	4.93%	not applicable	1.24%
BIN-36	THF	MeOH	2.69%	0.70%	not applicable
BIN-37	THF	EtOH	4.43%	0.61%	not applicable
BIN-38	THF	ACET	10.71%	not applicable	4.34%
BIN-39	MeA	MeOH	1.99%	not applicable	0.95%
BIN-40	MeA	EtOH	5.17%	1.88%	not applicable
BIN-41	MeA	ACET	3.01%	1.08%	not applicable
BIN-42	MeA	THF	3.88%	1.80%	not applicable
BIN-43	EtA	MeOH	2.21%	not applicable	1.06%
BIN-44	EtA	EtOH	5.75%	0.41%	not applicable
BIN-45	EtA	ACET	4.31%	1.38%	not applicable
BIN-46	EtA	MeA	2.85%	0.64%	not applicable
BIN-47	H2O	PrAc	22.81%	not applicable	16.26%
BIN-48	MeOH	PrAc	19.43%	3.66%	not applicable
BIN-49	EtOH	PrAc	21.67%	5.02%	not applicable
BIN-50	1-PrOH	PrAc	16.17%	1.33%	not applicable
BIN-51	BuOH	PrAc	18.05%	2.00%	not applicable
BIN-52	FoAc	PrAc	20.37%	not applicable	7.15%
BIN-53	AcAc	PrAc	11.76%	not applicable	1.85%
BIN-54	BuAc	PrAc	6.97%	0.72%	not applicable
BIN-55	MeA	PrAc	8.18%	2.46%	not applicable
BIN-56	H2O	BuAc	32.62%	8.81%	not applicable
BIN-57	FoAc	BuAc	21.25%	3.18%	not applicable
BIN-58	AcAc	BuAc	15.64%	0.61%	not applicable
BIN-59	2-PrOH	H2O	37.39%	10.21%	not applicable
BIN-60	2-PrOH	MeOH	4.18%	1.28%	not applicable
BIN-61	2-PrOH	EtOH	1.41%	0.23%	not applicable
BIN-62	2-PrOH	1-PrOH	2.30%	0.62%	not applicable
BIN-63	2-PrOH	BuOH	2.62%	0.10%	not applicable
BIN-64	2-PrOH	AcAc	21.23%	5.40%	not applicable
BIN-65	2-PrOH	PrAc	23.77%	5.27%	not applicable
BIN-66	2-PrOH	MeA	17.45%	5.99%	not applicable
BIN-67	2-PrOH	EtA	11.91%	3.72%	not applicable
BIN-68	2-PrOH	THF	7.73%	1.93%	not applicable
BIN-69	2-PrOH	ACET	12.97%	4.86%	not applicable
BIN-70	2-PrOH	BuAc	17.96%	not applicable	5.78%
BIN-71	PeOH	H2O	57.49%	14.63%	not applicable
BIN-72	PeOH	MeOH	12.92%	3.64%	not applicable
BIN-73	PeOH	EtOH	6.94%	1.25%	not applicable
BIN-74	PeOH	1-PrOH	4.84%	0.54%	not applicable
BIN-75	PeOH	BuOH	3.41%	0.21%	not applicable
BIN-76	PeOH	2-PrOH	4.10%	0.90%	not applicable

Figure 1. Results of the default COCO-COFE V7.3 NRTL model evaluation for binary mixtures: Average Relative Deviation on vapor composition (ARD Y), temperature (ARD T), and pressure (ARD P).

Mixture	Comp 1	Comp 2	Comp 3	Rel error Y
TER-1	ACET	H2O	MeOH	17.00%
TER-2	ACET	MeOH	EtOH	3.00%
TER-3	ACET	MeA	EtA	19.00%
TER-4	H2O	EtOH	MeA	19.00%
TER-5	H2O	EtOH	EtA	17.00%
TER-6	H2O	MeOH	MeA	35.00%
TER-7	H2O	MeOH	EtA	12.00%
TER-8	EtOH	MeOH	MeA	14.00%
TER-9	EtOH	MeOH	EtA	9.00%
TER-10	H2O	AcAc	ACET	17.00%
TER-11	H2O	AcAc	EtA	13.00%
TER-12	H2O	AcAc	MeA	31.00%
TER-13	MeOH	AcAc	MeA	15.00%
TER-14	BuAc	AcAc	PrAc	19.00%
TER-15	FoAc	AcAc	PrAc	26.00%
TER-16	2-PrOH	H2O	MeOH	32.00%
TER-17	2-PrOH	H2O	EtOH	21.00%
TER-18	2-PrOH	H2O	1-PrOH	23.00%
TER-19	2-PrOH	H2O	BuOH	23.00%
TER-20	2-PrOH	H2O	ACET	34.00%
TER-21	2-PrOH	MeOH	1-PrOH	18.00%
TER-22	2-PrOH	MeOH	ACET	14.00%
TER-23	2-PrOH	EtOH	EtA	5.00%
TER-24	EtOH	MeOH	1-PrOH	7.20%
TER-25	EtOH	H2O	BuOH	32.10%
TER-26	EtOH	MeOH	H2O	12.57%
TER-27	EtOH	1-PrOH	BuOH	25.90%
TER-28	EtOH	1-PrOH	H2O	15.80%
TER-29	H2O	AcAc	FoAc	9.43%
TER-31	H2O	1-PrOH	BuOH	22.30%
TER-32	MeOH	H2O	AcAc	31.70%

Figure 2. Results of the default COCO-COFE V7.3 NRTL model evaluation for ternary mixtures: Average Relative Deviation on vapor composition (ARD Y).

Mixture	Comp 1	Comp 2	ARD Y	ARD T	ARD P
BIN-1	AcAc	FoAc	11.78%	0.20%	not applicable
BIN-2	BuOH	AcAc	21.09%	2.37%	9.97%
BIN-3	BuOH	H2O	6.42%	0.83%	not applicable
BIN-4	EtOH	AcAc	18.85%	3.52%	5.01%
BIN-5	EtOH	BuOH	6.29%	0.41%	not applicable
BIN-6	EtOH	H2O	1.87%	0.26%	0.74%
BIN-7	EtOH	MeOH	2.67%	0.48%	4.33%
BIN-8	EtOH	1-PrOH	1.49%	0.29%	not applicable
BIN-9	H2O	AcAc	20.91%	2.01%	not applicable
BIN-10	H2O	FoAc	24.09%	3.51%	not applicable
BIN-11	H2O	1-PrOH	5.08%	0.55%	1.63%
BIN-12	MeOH	AcAc	15.62%	3.15%	not applicable
BIN-13	MeOH	BuOH	10.21%	6.65%	not applicable
BIN-14	MeOH	H2O	3.04%	0.32%	1.03%
BIN-15	MeOH	1-PrOH	6.89%	0.42%	not applicable
BIN-16	1-PrOH	AcAc	4.21%	0.98%	1.17%
BIN-17	1-PrOH	BuOH	5.12%	0.47%	not applicable
BIN-18	EtA	H2O	16.83%	not applicable	12.19%
BIN-19	EtA	THF	12.90%	0.67%	not applicable
BIN-20	MeA	H2O	10.26%	not applicable	4.15%
BIN-21	THF	H2O	4.93%	0.82%	not applicable
BIN-22	ACET	PROP	10.66%	not applicable	18.90%
BIN-23	ACET	AcAc	23.74%	not applicable	5.55%
BIN-24	ACET	FoAc	2.00%	0.30%	not applicable
BIN-25	THF	BuOH	2.46%	0.21%	not applicable
BIN-26	MeA	PROP	4.19%	not applicable	0.67%
BIN-27	MeA	BuOH	8.79%	2.24%	not applicable
BIN-28	MeA	FoAc	32.74%	1.37%	not applicable
BIN-29	MeA	AcAc	16.28%	1.57%	not applicable
BIN-30	EtA	PROP	2.56%	0.29%	not applicable
BIN-31	EtA	BuOH	3.35%	1.44%	not applicable
BIN-32	EtA	AcAc	28.73%	1.94%	not applicable
BIN-33	ACET	H2O	6.33%	not applicable	4.99%
BIN-34	ACET	MeOH	2.18%	not applicable	3.97%
BIN-35	ACET	EtOH	3.91%	not applicable	0.56%
BIN-36	THF	MeOH	2.13%	0.58%	not applicable
BIN-37	THF	EtOH	2.93%	0.34%	not applicable
BIN-38	THF	ACET	0.65%	not applicable	0.30%
BIN-39	MeA	MeOH	1.72%	not applicable	0.79%
BIN-40	MeA	EtOH	3.39%	0.88%	not applicable
BIN-41	MeA	ACET	1.31%	0.16%	not applicable
BIN-42	MeA	THF	1.16%	0.14%	not applicable
BIN-43	EtA	MeOH	2.23%	not applicable	0.96%
BIN-44	EtA	EtOH	5.80%	0.43%	not applicable
BIN-45	EtA	ACET	1.79%	0.20%	not applicable
BIN-46	EtA	MeA	2.81%	0.23%	not applicable
BIN-47	H2O	PrAc	11.80%	not applicable	5.66%
BIN-48	MeOH	PrAc	18.43%	2.51%	not applicable
BIN-49	EtOH	PrAc	16.63%	1.69%	not applicable
BIN-50	1-PrOH	PrAc	9.16%	1.48%	not applicable
BIN-51	BuOH	PrAc	12.69%	1.35%	not applicable
BIN-52	FoAc	PrAc	23.24%	not applicable	3.01%
BIN-53	AcAc	PrAc	11.88%	not applicable	1.87%
BIN-54	BuAc	PrAc	18.96%	2.47%	not applicable
BIN-55	MeA	PrAc	13.83%	0.97%	not applicable
BIN-56	H2O	BuAc	8.10%	0.80%	not applicable
BIN-57	FoAc	BuAc	35.40%	19.18%	not applicable
BIN-58	AcAc	BuAc	19.09%	2.23%	not applicable
BIN-59	2-PrOH	H2O	3.80%	0.41%	not applicable
BIN-60	2-PrOH	MeOH	6.40%	1.87%	not applicable
BIN-61	2-PrOH	EtOH	1.28%	0.19%	not applicable
BIN-62	2-PrOH	1-PrOH	2.18%	0.59%	not applicable
BIN-63	2-PrOH	BuOH	1.51%	0.35%	not applicable
BIN-64	2-PrOH	AcAc	21.23%	5.40%	not applicable
BIN-65	2-PrOH	PrAc	14.36%	3.28%	not applicable
BIN-66	2-PrOH	MeA	4.45%	0.84%	not applicable
BIN-67	2-PrOH	EtA	1.95%	0.55%	not applicable
BIN-68	2-PrOH	THF	1.02%	0.18%	not applicable
BIN-69	2-PrOH	ACET	4.91%	0.94%	not applicable
BIN-70	2-PrOH	BuAc	17.96%	not applicable	5.78%
BIN-71	PeOH	H2O	6.77%	1.72%	not applicable
BIN-72	PeOH	MeOH	14.72%	2.97%	not applicable
BIN-73	PeOH	EtOH	3.74%	0.24%	not applicable
BIN-74	PeOH	1-PrOH	4.78%	0.54%	not applicable
BIN-75	PeOH	BuOH	2.18%	0.14%	not applicable
BIN-76	PeOH	2-PrOH	2.64%	0.42%	not applicable

Figure 3. Results of the default Aspen Plus V12.0 NRTL model evaluation for binary mixtures: Average Relative Deviation on vapor composition (ARD Y), temperature (ARD T), and pressure (ARD P).

Mixture	Comp 1	Comp 2	Comp 3	Rel error Y
TER-1	ACET	H2O	MeOH	14.00%
TER-2	ACET	MeOH	EtOH	5.00%
TER-3	ACET	MeA	EtA	11.00%
TER-4	H2O	EtOH	MeA	7.00%
TER-5	H2O	EtOH	EtA	16.00%
TER-6	H2O	MeOH	MeA	11.00%
TER-7	H2O	MeOH	EtA	13.00%
TER-8	EtOH	MeOH	MeA	16.00%
TER-9	EtOH	MeOH	EtA	10.00%
TER-10	H2O	AcAc	ACET	18.00%
TER-11	H2O	AcAc	EtA	10.00%
TER-12	H2O	AcAc	MeA	20.00%
TER-13	MeOH	AcAc	MeA	14.00%
TER-14	BuAc	AcAc	PrAc	28.00%
TER-15	FoAc	AcAc	PrAc	26.00%
TER-16	2-PrOH	H2O	MeOH	8.00%
TER-17	2-PrOH	H2O	EtOH	5.00%
TER-18	2-PrOH	H2O	1-PrOH	9.00%
TER-19	2-PrOH	H2O	BuOH	11.00%
TER-20	2-PrOH	H2O	ACET	8.00%
TER-21	2-PrOH	MeOH	1-PrOH	9.00%
TER-22	2-PrOH	MeOH	ACET	10.00%
TER-23	2-PrOH	EtOH	EtA	4.00%
TER-24	EtOH	MeOH	1-PrOH	5.43%
TER-25	EtOH	H2O	BuOH	6.73%
TER-26	EtOH	MeOH	H2O	7.69%
TER-27	EtOH	1-PrOH	BuOH	20.29%
TER-28	EtOH	1-PrOH	H2O	8.40%
TER-29	H2O	AcAc	FoAc	8.76%
TER-31	H2O	1-PrOH	BuOH	11.03%
TER-32	MeOH	H2O	AcAc	35.45%

Figure 4. Results of the default Aspen Plus V12.0 NRTL model evaluation for ternary mixtures: Average Relative Deviation on vapor composition (ARD Y).

An improvement of the current Aspen Plus model is achieved in this work through interaction parameters re-fitting for all 45 mixtures coloured in red and yellow in Figure 3 and Figure 4. The list of the re-fitted coefficients is available in Table 5. The results of the re-fitted model evaluation are gathered in Figure 5 and Figure 6 for binary and ternary mixtures, respectively.

Mixture	Comp 1	Comp 2	ARD Y	ARD T	ARD P	Refitted?
BIN-1	AcAc	FoAc	4.02%	0.47%	not applicable	Yes
BIN-2	BuOH	AcAc	3.31%	0.41%	1.33%	Yes
BIN-3	BuOH	H2O	6.42%	0.83%	not applicable	No
BIN-4	EtOH	AcAc	13.48%	2.42%	2.82%	Yes
BIN-5	EtOH	BuOH	6.29%	0.41%	not applicable	No
BIN-6	EtOH	H2O	1.87%	0.26%	0.74%	No
BIN-7	EtOH	MeOH	2.67%	0.48%	4.33%	No
BIN-8	EtOH	1-PrOH	1.49%	0.29%	not applicable	No
BIN-9	H2O	AcAc	6.06%	0.42%	not applicable	Yes
BIN-10	H2O	FoAc	4.95%	0.37%	not applicable	Yes
BIN-11	H2O	1-PrOH	5.08%	0.55%	1.63%	No
BIN-12	MeOH	AcAc	7.48%	1.75%	not applicable	Yes
BIN-13	MeOH	BuOH	10.21%	6.65%	not applicable	No
BIN-14	MeOH	H2O	3.04%	0.32%	1.03%	No
BIN-15	MeOH	1-PrOH	6.89%	0.42%	not applicable	No
BIN-16	1-PrOH	AcAc	4.21%	0.98%	1.17%	No
BIN-17	1-PrOH	BuOH	5.12%	0.47%	not applicable	Yes
BIN-18	EtA	H2O	15.35%	not applicable	2.86%	Yes
BIN-19	EtA	THF	9.69%	0.46%	not applicable	Yes
BIN-20	MeA	H2O	10.26%	not applicable	4.15%	No
BIN-21	THF	H2O	4.93%	0.82%	not applicable	No
BIN-22	ACET	PROP	0.46%	not applicable	0.34%	Yes
BIN-23	ACET	AcAc	14.89%	not applicable	5.78%	Yes
BIN-24	ACET	FoAc	2.00%	0.30%	not applicable	Yes
BIN-25	THF	BuOH	2.46%	0.21%	not applicable	No
BIN-26	MeA	PROP	4.19%	not applicable	0.67%	No
BIN-27	MeA	BuOH	8.79%	2.24%	not applicable	No
BIN-28	MeA	FoAc	4.70%	0.73%	not applicable	Yes
BIN-29	MeA	AcAc	8.93%	1.02%	not applicable	Yes
BIN-30	EtA	PROP	2.56%	0.29%	not applicable	No
BIN-31	EtA	BuOH	3.35%	1.44%	not applicable	No
BIN-32	EtA	AcAc	11.81%	1.13%	not applicable	Yes
BIN-33	ACET	H2O	6.33%	not applicable	4.99%	No
BIN-34	ACET	MeOH	2.18%	not applicable	3.97%	No
BIN-35	ACET	EtOH	3.91%	not applicable	0.56%	No
BIN-36	THF	MeOH	2.13%	0.58%	not applicable	No
BIN-37	THF	EtOH	2.93%	0.34%	not applicable	No
BIN-38	THF	ACET	0.65%	not applicable	0.30%	No
BIN-39	MeA	MeOH	1.72%	not applicable	0.79%	No
BIN-40	MeA	EtOH	3.39%	0.88%	not applicable	No
BIN-41	MeA	ACET	1.31%	0.16%	not applicable	No
BIN-42	MeA	THF	1.16%	0.14%	not applicable	No
BIN-43	EtA	MeOH	2.23%	not applicable	0.96%	No
BIN-44	EtA	EtOH	5.80%	0.43%	not applicable	No
BIN-45	EtA	ACET	1.79%	0.20%	not applicable	No
BIN-46	EtA	MeA	2.81%	0.23%	not applicable	No
BIN-47	H2O	PrAc	7.39%	not applicable	2.45%	Yes
BIN-48	MeOH	PrAc	14.99%	2.98%	not applicable	Yes
BIN-49	EtOH	PrAc	13.65%	1.60%	not applicable	Yes
BIN-50	1-PrOH	PrAc	9.16%	1.48%	not applicable	No
BIN-51	BuOH	PrAc	6.26%	0.24%	not applicable	Yes
BIN-52	FoAc	PrAc	5.02%	not applicable	1.32%	Yes
BIN-53	AcAc	PrAc	5.48%	not applicable	1.17%	Yes
BIN-54	BuAc	PrAc	6.98%	0.70%	not applicable	Yes
BIN-55	MeA	PrAc	6.76%	0.81%	not applicable	Yes
BIN-56	H2O	BuAc	8.10%	0.80%	not applicable	No
BIN-57	FoAc	BuAc	0.27%	0.04%	not applicable	Yes
BIN-58	AcAc	BuAc	4.39%	0.61%	not applicable	Yes
BIN-59	2-PrOH	H2O	3.81%	0.41%	not applicable	No
BIN-60	2-PrOH	MeOH	6.40%	1.87%	not applicable	No
BIN-61	2-PrOH	EtOH	1.28%	0.19%	not applicable	No
BIN-62	2-PrOH	1-PrOH	2.18%	0.59%	not applicable	No
BIN-63	2-PrOH	BuOH	1.51%	0.35%	not applicable	No
BIN-64	2-PrOH	AcAc	7.90%	2.12%	not applicable	Yes
BIN-65	2-PrOH	PrAc	11.29%	1.59%	not applicable	Yes
BIN-66	2-PrOH	MeA	4.45%	0.84%	not applicable	No
BIN-67	2-PrOH	EtA	1.95%	0.55%	not applicable	No
BIN-68	2-PrOH	THF	1.03%	0.18%	not applicable	No
BIN-69	2-PrOH	ACET	4.91%	0.94%	not applicable	No
BIN-70	2-PrOH	BuAc	0.66%	not applicable	0.85%	Yes
BIN-71	PeOH	H2O	6.77%	1.72%	not applicable	No
BIN-72	PeOH	MeOH	14.72%	2.97%	not applicable	No
BIN-73	PeOH	EtOH	3.74%	0.24%	not applicable	No
BIN-74	PeOH	1-PrOH	4.78%	0.54%	not applicable	No
BIN-75	PeOH	BuOH	2.18%	0.14%	not applicable	No
BIN-76	PeOH	2-PrOH	2.64%	0.42%	not applicable	No

Figure 5. Results of the re-parametrised NRTL model evaluation for binary mixtures: Average Relative Deviation on vapor composition (ARD Y), temperature (ARD T), and pressure (ARD P). The column "refitted?" specifies which coefficients have been refitted in this work.

Mixture	Comp 1	Comp 2	Comp 3	ARD Y
TER-1	ACET	H2O	MeOH	14.44%
TER-2	ACET	MeOH	EtOH	4.54%
TER-3	ACET	MeA	EtA	11.10%
TER-4	H2O	EtOH	MeA	7.06%
TER-5	H2O	EtOH	EtA	15.66%
TER-6	H2O	MeOH	MeA	11.07%
TER-7	H2O	MeOH	EtA	15.04%
TER-8	EtOH	MeOH	MeA	16.10%
TER-9	EtOH	MeOH	EtA	9.69%
TER-10	H2O	AcAc	ACET	15.21%
TER-11	H2O	AcAc	EtA	15.48%
TER-12	H2O	AcAc	MeA	18.25%
TER-13	MeOH	AcAc	MeA	12.24%
TER-14	BuAc	AcAc	PrAc	12.94%
TER-15	FoAc	AcAc	PrAc	6.86%
TER-16	2-PrOH	H2O	MeOH	7.67%
TER-17	2-PrOH	H2O	EtOH	5.67%
TER-18	2-PrOH	H2O	1-PrOH	9.33%
TER-19	2-PrOH	H2O	BuOH	10.67%
TER-20	2-PrOH	H2O	ACET	8.33%
TER-21	2-PrOH	MeOH	1-PrOH	9.67%
TER-22	2-PrOH	MeOH	ACET	10.33%
TER-23	2-PrOH	EtOH	EtA	4.67%
TER-24	EtOH	MeOH	1-PrOH	5.43%
TER-25	EtOH	H2O	BuOH	6.73%
TER-26	EtOH	MeOH	H2O	7.69%
TER-27	EtOH	1-PrOH	BuOH	18.85%
TER-28	EtOH	1-PrOH	H2O	8.40%
TER-29	H2O	AcAc	FoAc	11.43%
TER-31	H2O	1-PrOH	BuOH	11.15%
TER-32	MeOH	H2O	AcAc	17.41%

Figure 6. Results of the re-parametrised NRTL model evaluation for ternary mixtures: Average Relative Deviation on vapor composition (ARD Y).

Table 5: Re-parametrised Aspen-compliant NRTL coefficients, including temperature-independent (Aij) and temperature-dependent (Bij) parameters (expression (4)). Acronyms are retrieved from Table 1.

i	j	Aij	Aji	Bij	Bji
MeOH	AcAc	-0.94	-1.75	760.71	157.89
BuOH	Acac	2.02	0.82	-268.78	-897.09
EtOH	AcAc	44.78	-4.22	-15450.2	940.49
H2O	FoAc	1.23	2.31	-922.19	-677.37
AcAc	FoAc	-19.49	-10.71	8365.84	3629.42
H2O	AcAc	-6.35	9.40	2202.34	-3173.0
THF	EtA	3.85	9.11	-1451.64	-2962.0
EtA	H2O	99.89	-31.06	-34457.7	12044.9
Acet	PrOH	-0.79	0.11	273.90	146.67
AcAc	BuAc	-11.82	10.81	4328.61	-2683.2
MeOH	PrAc	-1.92	686.6	585.61	1966.78
H2O	PrAc	-15.52	5.21	6519.08	-2027.7
EtOH	PrAc	281.2	-2.84	0.00	859.07
BuOH	PrAc	2.02	0.88	-330.78	-899.05
FoAc	PrAc	-20.68	-7.25	7676.38	4813.23
AcAc	PrAc	-17.89	-6.50	6640.57	4025.00
BuAc	PrAc	6.61	-6.94	-2775.23	2986.03
MeA	PrAc	21.09	-48.34	-8348.18	18479.3
FoAc	BuAc	-8.42	68.93	3088.05	-25002
Acet	FoAc	2.56	-4.97	-1050.83	1530.55
Acet	AcAc	-0.02	0.92	-314.35	-52.07
MeA	FoAc	23.93	-25.49	-8376.04	8924.19
MeA	AcAc	-14.05	10.47	5228.10	-3911.6
EtA	AcAc	28.98	-20.93	-10795.6	7836.64

The kPIs of the re-parametrised model are summarised in Table 6.

Table 6: Summary of Key Performance Indicators from the evaluation of both COFE and Aspen Plus default models.

KPI	Re-fitted model
ARD-composition (%)	6.96
ARD-Temperature (%)	0.89
ARD – Pressure (%)	2.01
“Green” mixtures	80
“Yellow” mixtures	27
“Red” mixtures	0

The re-fitted NRTL model shows a significantly improved capability to reproduce experimental VLE data. Indeed, for all investigated binary and ternary mixtures, the ARD is < 20%. Temperature versus composition plots (both boiling point and dew point curves) are drawn to check the model smoothness. For all binary mixtures of interest, a smooth trend is obtained throughout the whole composition range. As a conclusion. the re-fitted model can be considered sufficiently accurate for process design and optimisation purposes.

CONCLUSIONS

Using experimental data available in NIST ThermoData Engine collected at atmospheric pressure, a NRTL model has been re-parametrised to improve the VLE characterisation in strongly non-ideal mixtures containing water and short-chain alcohols, esters and carboxylic acids.

This model can be successfully used to improve the thermodynamic modelling of a wide range of both conventional and bio-based processes, including the recovery of diluted chemicals from fermentation broth, extraction and separation products for biochemicals upconcentration, Fischer esterification, vinegar production, and organics recovery and valorisation from wastewater in pyrolysis and hydrothermal liquefaction.

ACKNOWLEDGEMENTS

The present work received co-funding from FUEL-UP project, part of the European Union’s Horizon Europe Research and Innovation Programme under Grant Agreement N° 101136123.

The authors express their gratitude to BTG (Netherlands) for their precious contribution concerning feedstock supply and to Avecom (Belgium) for their support and cooperation in sample analysis.

The authors are also thankful to Rune Myrstad (SINTEF) and Camilla Otterlei (SINTEF) for their substantial contribution to sample analytics.

AUTHOR IDENTIFIERS

Author ORCIDs:
Gilardi Matteo: 0000-0003-2456-7549
Bisotti Filippo: 0000-0001-9106-1797
Bernd Wittgens: 0000-0001-5705-9255

REFERENCES

1. Barahmand Z, Wang L, Holm-Nielsen JB, Eikeland M. Significance of pyrolysis in the circular economy: an integrative review of technologies, potential chemicals, and separation techniques. Fuel 398:135539 (2025). https://doi.org/10.1016/j.fuel.2025.135539
2. BTG Biomass Technology Group BV. Production of sustainable transport biofuels via pyrolysis upgrading. Microsoft PowerPoint - IEA Webinar November 2024 - BTG presentation v1
3. FUEL-UP HEU project. Powering a greener future for aviation and marine transport. https://www.fuelup-project.eu/
4. AspenTech documentation. https://www.aspentech.com/en/products/engineering/aspen-plus
5. COCO. The CAPE-OPEN to CAPE-OPEN simulator. https://www.cocosimulator.org/
6. Renon, H., Prausnitz, J. M. Local compositions in thermodynamic excess functions for liquid mixtures. AIChE Journal, 14 (1): 135–144 (1968).
7. Gilardi M, Bisotti F, Tobiesen A, Knuutila HK, Bonalumi D. An approach for VLE model development, validation, and implementation in aspen plus for amine blends in CO2 capture: the HS3 solvent case study. International Journal of Greenhouse Gas Control 126:103911 (2023). https://doi.org/10.1016/j.ijggc.2023.103911

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Framework based on Population Balance Modeling for Predicting Li–O_2 Battery Discharge and Life Cycle Behavior

Nadia G. Khouri[a*], Jean F. Leal Silva[a], Letícia M. S. Barros[a], Viktor O. C. Concha[b], Rubens Maciel Filho[a]

[a] University of Campinas (UNICAMP), School of Chemical Engineering, Campinas, São Paulo, Brazil
[b] Federal University of São Paulo (UNIFESP), School of Chemical Engineering, Diadema, São Paulo, Brazil
* Corresponding Author: nadiagk@unicamp.br.

ABSTRACT

The growing integration of renewable energy sources such as solar and wind power has intensified the demand for advanced energy storage technologies. Lithium–air (Li–O_2) batteries are particularly attractive due to their exceptionally high theoretical specific energy, which surpasses that of the conventional lithium-ion system. However, their practical application is hindered by poor reversibility during discharge, primarily due to the formation and decomposition of lithium peroxide (Li_2O_2), which causes cathode passivation and capacity fading. Since the electrochemical performance of Li–O_2 batteries is strongly influenced by the morphology, size, and spatial distribution of Li_2O_2 crystals, understanding the mechanisms governing their nucleation and growth is critical. To address this challenge, this work proposes a computational framework based on population balance modeling (PBM) to describe Li_2O_2 crystallization dynamics during battery discharge. The framework integrates population, mass, and energy balances, allowing the coupled analysis of electrochemical kinetics, supersaturation effects, and the evolution of crystal size distributions. Compared with continuum-scale and phase-field models, the PBM approach offers reduced computational cost while naturally accounting for particle size distributions and linking microscopic crystallization phenomena to macroscopic battery performance and degradation. Although simplified assumptions were adopted in this initial formulation, the framework successfully captures the essential discharge behavior of Li–O_2 systems. As such, it provides a robust foundation for future model refinements incorporating additional physicochemical mechanisms and more detailed electrochemical and transport phenomena.

Keywords: Batteries, Dynamic Modelling, Energy Storage, Energy Systems, Simulation, Population Balance

INTRODUCTION

Renewable energy sources represent a promising alternative to mitigate environmental issues associated with conventional fossil-based systems. In addition to their role in reducing greenhouse gas emissions, renewables have become increasingly cost-competitive in electricity generation. According to the International Renewable Energy Agency [1], solar photovoltaic (PV) and onshore wind technologies are significantly more economical than fossil fuel–based alternatives with costs approximately 41% and 53% lower than conventional options, respectively [2]. Despite their advantages, sources such as solar and wind are inherently intermittent. Solar power generation is constrained to daylight hours, while wind energy availability is dependent on geographic and meteorological conditions [3]. This intermittency is a challenge to the reliable and continuous supply of electricity that motivates the development of integrated energy systems capable of storing excess energy and delivering it on demand. In this context, electrochemical energy storage technologies, particularly batteries, have emerged as a viable solution to enable the large-scale integration of renewables. Notably, the cost of battery technologies has decreased more than 90% since 2010, highlighting their growing viability [2].

Rechargeable lithium-ion batteries have become the dominant energy storage technology for mobile app-

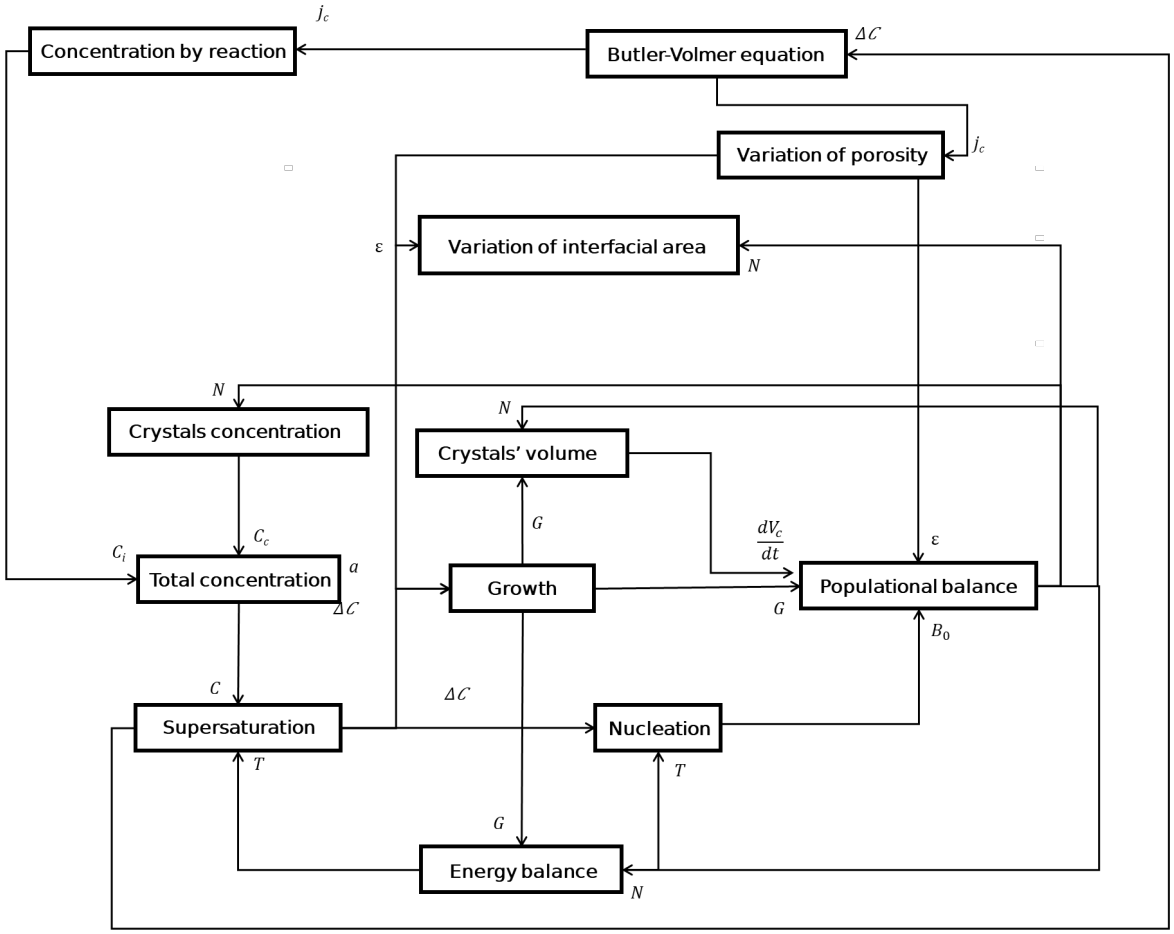

Figure 1. Information flow to model the discharge product involved in the Li_2-O_2 battery.

lications, largely due to their high energy density [4]. Despite these advantages, the energy density of lithium-ion batteries is approaching its theoretical limit and remains insufficient to fully meet the growing global demand for large-scale storage. In this context, lithium–air (Li–O_2) batteries have attracted attention because of their high theoretical specific energy, reported to reach approximately 3500 W.h/kg [5]. Nevertheless, this value is unlikely to be achieved in practical systems due to fundamental challenges [6].

In lithium–oxygen (Li–O_2) batteries, the overall reversible electrochemical reaction is governed by the formation and decomposition of lithium peroxide (Li_2O_2), in contrast to the intercalation-based mechanism of conventional lithium-ion batteries [7]. The morphology and accumulation of the Li_2O_2 discharge product are critical factors influencing the performance and reversibility of Li–O_2 cells. Specifically, Li_2O_2 formation and decomposition at the cathode directly affect round-trip efficiency, rate capability, and long-term cycling stability [8]. Studies have demonstrated that the insulating Li_2O_2 can deposit in diverse morphologies, including thin films and complex toroidal structures, depending on operating conditions [9]. Since the crystallization form of Li_2O_2 governs electron transport pathways and reaction accessibility during charging, controlling its morphology is essential. Consequently, a detailed understanding of how process variables influence Li_2O_2 nucleation, growth, and morphology is essential for optimizing battery capacity and achieving stable, reversible Li–O_2 operation.

Motivated by the influence of Li_2O_2 crystallization in the performance of Li–O_2 batteries, previous theoretical studies based on experimental observations have examined the formation of discharge products during battery operation [10-12]. In earlier work by the authors, a population balance–based model was applied to analyze Li_2O_2 deposition during discharge, and the resulting trends and performance-related implications were discussed [13]. The present study, however, addresses a distinct and complementary objective. Rather than focusing on the interpretation of modeling outcomes, this work provides a detailed exposition of the population balance modeling framework itself, including the governing equations, assumptions, numerical implementation, and coupling with

electrochemical variables. By systematically elucidating the structure and rationale of the model, this contribution establishes a reproducible simulation platform for Li_2O_2 crystallization, with direct applicability to other battery systems in which crystallization processes are crucial. As such, the novelty of this work lies in the formalization and clarification of the modeling methodology, enabling its reliable application and future expansion to more complex electrochemical energy storage systems.

LI-O2 CRYSTALLIZATION MODELLING

Crystallization produces solid particles with a distribution of sizes, commonly described by the crystal size distribution (CSD). In Li–O_2 batteries, the evolution of the CSD during discharge is governed by the coupled effects of electrochemical kinetics, mass and energy transport, and particle population dynamics. These phenomena are captured through an integrated modeling framework that accounts for Li_2O_2 saturation, nucleation, and crystal growth (Figure 1). Accordingly, the following sections detail the core components of the model, including the characterization of Li_2O_2 crystals, mass and energy balances coupled with electrochemical kinetics, the population balance formulation, and the expressions adopted for saturation, nucleation, and growth rates.

Characteristics of Li2O2 crystals

The external geometry of a crystal can be characterized by its characteristic length (L), the surface areas of its facets, and the angles between them. In modeling applications, these geometric attributes are conveniently represented through shape factors, which provide a compact description of crystal morphology. Shape factors are commonly defined with respect to crystal volume and surface area, as expressed by the volume-based and area-based shape factors (Equations 1 and 2), respectively [14]. These dimensionless parameters relate the characteristic length of a crystal to its corresponding volume and surface area. In the case of lithium peroxide, experimental observations indicate that the discharge product typically forms disk-like structures, such as toroids or oblate spheroids. Consequently, the associated shape factors are functions of the aspect ratio between the disk diameter and thickness, which governs the effective morphology of Li_2O_2 crystals [15].

$$V_c = k_V L^3 \tag{1}$$

$$A_c = k_a L^2 \tag{2}$$

Mass balance of Li_2O_2 crystals and kinetics

The material balance of Li_2O_2 crystals during battery discharge is formulated by accounting for the temporal accumulation of solid discharge products within the cathode. Equation 3 quantifies the total mass of Li_2O_2 as obtained from the distribution of insoluble crystals in the system and Equation 4 describes the evolution of crystal volume as a function of particle growth. To close the mass balance, the rate of Li_2O_2 formation must be linked to the underlying electrochemical processes. Lithium peroxide is generated through the cathodic oxygen reduction reaction, and its formation rate is evaluated using the Butler–Volmer formulation (Equation 5). This expression captures the dependence of the reaction rate on the applied overpotential as well as on the concentrations of the reactants (O_2 and Li^+) and the discharge product (Li_2O_2). Based on the calculated cathodic reaction rate, the resulting variations in Li_2O_2 concentration, electrode porosity, and electrochemically active interfacial area are subsequently determined (Equations 6–8). Finally, the effective suspension volume of the system (Vsusp) is estimated from the geometric dimensions of the cathode (radius and thickness) in combination with the evolving porosity (Equation 9) [16].

$$C_c(t) = k_V \rho_c \int_0^\infty N(L,t) L^3 dL \tag{3}$$

$$\frac{dV_c}{dt} = -3k_V \int_0^\infty G(L,t) N(L,t) L^2 dL \tag{4}$$

$$\frac{j_c}{nF} = k_{ano} \Delta C exp\left[\frac{(1-\beta)nF}{RT}\eta_c\right] - k_c\,(C_{Li})^2 C_{O_2} exp\left[\frac{-\beta nF}{RT}\eta_c\right] \tag{5}$$

$$\frac{\partial C_i}{\partial t} = -\frac{a j_c MM}{2F} \tag{6}$$

$$\frac{\partial \varepsilon}{\partial t} = \frac{a j_c MM}{2F\rho_c} \tag{7}$$

$$ae = ae_0\left[1 - \left(\frac{k_v \int_0^\infty N L^3 dL}{V_{susp}\varepsilon_0}\right)\right]^p \tag{8}$$

$$V_{susp} = \pi\, r_{el}{}^2 esp\, \varepsilon \tag{9}$$

Energy balance

The energy balance describes the heat transfer processes occurring within the system and is expressed by Equation 10. Two primary contributions are considered in this formulation. The first accounts for the thermal effects associated with crystallization and Li_2O_2 formation, represented by the crystallization enthalpy. The second term represents heat exchange with the surroundings. In the present model, heat dissipation to the environment is assumed to occur via forced air convection, characterized by an average overall heat transfer coefficient of 30 W m^{-2} K^{-1} [17].

$$\rho_e c_{p_e} \frac{dT}{dt} = \left(-3\Delta H_c \rho_c k_V \int_0^\infty G(L,t)N(L,t)L^2 dL\right) - UA_T(T - T_s) \quad (10)$$

Populational Balance

The general population balance equation (PBE) employed in this work was previously derived and reported by the authors in an earlier study (Equation 11) [18].

$$\frac{\partial N_\xi}{\partial t} + \frac{\partial}{\partial x} \bullet (U_o N_\xi) + \frac{\partial}{\partial \xi} \bullet (\dot{\xi} N_\xi) = h_\xi \quad (11)$$

Building on this formulation, several assumptions were adopted to tailor the population balance description to Li–O_2 battery systems. Crystal size is characterized using a single linear dimension, namely the characteristic length L, resulting in a univariate population balance (ξ=L). The system is modeled under transient operating conditions, reflecting the time-dependent nature of battery discharge. The control volume is defined as the cathode surface, which constitutes the primary site of Li_2O_2 nucleation and growth. Accordingly, the corresponding term in the PBE accounts for temporal changes in particle volume within the electrode domain. Particle transport within the phase volume is assumed to be governed solely by the linear crystal growth rate, while birth and death processes associated with discrete events are neglected (h_ξ =0). Under these assumptions, the population balance for Li_2O_2 crystallization reduces to the batch crystallizer formulation expressed in Equation 12, yielding a partial differential equation that depends on both particle size and time, N(L, t). Two boundary conditions are required to close the system: the first, N(0, t), represents the nucleation flux (Equation 13), while the second, N(L, 0), corresponds to the initial seed crystal size distribution. The resulting PBE is solved numerically using the method of classes, in which the size domain is discretized into finite intervals, transforming the PDE into a system of ordinary differential equations with established numerical solution procedures (Equations 14–16).

$$\frac{\partial N}{\partial t} + \frac{\partial (GN)}{\partial L} + N\frac{\partial V}{V\partial t} = 0 \quad (12)$$

$$N(0,t) = \frac{B_0(t)}{G(0,t)} \quad (13)$$

$$\frac{dN_1}{dt} + \frac{1}{V_{susp}}\frac{dV_c}{dt}N_1 + \frac{G(L_1)}{2\Delta c_1}N_1 + \frac{G(L_1)}{2\Delta c_2}N_2 = B_0 \quad (14)$$

$$\frac{dN_i}{dt} + \frac{1}{V_{susp}}\frac{dV_c}{dt}N_i + \frac{G(L_i)}{2\Delta c_{i+1}}N_{i+1} + \frac{G(L_i) - G(L_{i-1})}{2\Delta c_i}N_i - \frac{G(L_{i-1})}{2\Delta c_{i-1}}N_{i-1} = 0 \quad (15)$$

$$\frac{dN_N}{dt} + \frac{1}{V_{susp}}\frac{dV_c}{dt}N_N - \frac{G(L_{N-1})}{2\Delta c_N}N_N - \frac{G(L_{N-1})}{2\Delta c_{N-1}}N_{N-1} = 0 \quad (16)$$

Crystal saturation

A solution is defined as saturated when thermodynamic equilibrium is established between a dissolved species and its corresponding solid phase, indicating that the maximum solute concentration permitted under given conditions has been reached. Saturation is therefore an intrinsic property of the solute–solvent system and is strongly dependent on temperature [19]. In most cases, solubility increases with temperature; however, certain systems exhibit weak temperature dependence or even a decrease in solubility as temperature rises. In contrast, supersaturation arises when the solute concentration in the solution exceeds its equilibrium saturation value. The degree of supersaturation (ΔC) is quantified as the difference between the actual solute concentration (C) and the equilibrium saturation concentration (C∗), as defined in Equation 17. Supersaturated systems are thermodynamically unstable and provide the driving force for nucleation and crystal growth, making supersaturation important in crystallization processes.

$$\Delta C = C - C^* \quad (17)$$

Nucleation rate (B0)

Nucleation refers to the formation of stable crystal embryos from a supersaturated solution, marking the onset of a new solid phase. In this work, the nucleation rate of Li_2O_2 is modeled using classical nucleation theory [20]. For disk-shaped nuclei, the free energy of formation is evaluated by considering a critical nucleus size that ensures thermodynamic stability and sustained growth (Equation 18) [9]. The corresponding surface energy (γ) is estimated using the correlation given in Equation 19 [19]. On this basis, the nucleation rate expression (Equation 20) is formulated as a function of the Zeldovich factor (Equation 21), the diffusion coefficient of Li_2O_2 in the electrolyte (Equation 25), and the density of active nucleation sites (x0) The parameter x0 must be estimated for the system and represents one of the principal sources of uncertainty in the model. Accordingly, its influence on the predicted crystallization behavior is examined in detail in the previous work [13].

$$G_{crit} = \frac{\gamma^2 a^4}{2k_b T \ln(dC/C^*)} \quad (18)$$

$$\gamma = 0.414\, k_b T \ln(dC/C^*)\left(\frac{dC * Na}{MM_{Li_2O_2}}\right)^{2/3} \quad (19)$$

$$B_0 = D_l a_D{}^{-2} Z x_0 \exp\left(\frac{-G_{crit}}{k_b T}\right) \quad (20)$$

$$Z = \frac{[G_{crit}/(3\pi k_b T)]^{0.5}}{n_{crit}} \quad (21)$$

Growth rate (G)

The crystal growth rate, G, is defined as the temporal rate of change of the characteristic crystal length, dL/dt and is evaluated using the expression given in Equation 22. To account for the interplay between electrochemical reaction kinetics and mass transport limitations, an effectiveness factor for Li_2O_2 growth, ηr, is introduced and calculated according to Equation 23. This factor quantifies the relative contributions of reaction-controlled and diffusion-controlled growth regimes. Low values of effectiveness indicate dominance of the electrochemical reaction rate constant (k_c), whereas values approaching unity signify growth primarily limited by mass transfer, characterized by the mass transfer coefficient (kd). The coefficient kd is determined using Equation 24, as described in the literature [21-22]. The parameter θ represents the ratio of the operating current density to the limiting current density, with the latter corresponding to conditions under which reactant depletion in the electrolyte or membrane induces concentration polarization and constrains mass transport [23]. Finally, the diffusivity of disk-shaped Li_2O_2 crystals, D_l, is calculated using Equation 25 [24], which is derived from the linear regime observed in the early-time response of the curve obtained from intermittent potentiostatic titration experiments. The constant ksl corresponds to the slope of this linear relationship.

$$G = \frac{k_a k_c MM}{3\rho_c k_V} \eta_r (\Delta C)^{j_l} \quad (22)$$

$$\eta_r = \frac{1}{\frac{k_c}{k_d}(\Delta C) + 1} \quad (23)$$

$$k_d = \frac{I_c K_R RT}{8F^3 (C_{Li_2O_2}/MM)^2 D_l \; \theta \; ae \; esp} \quad (24)$$

$$D_l = \left(\frac{\pi}{36}\right)\left(\frac{k_{sl} L}{Q}\right)^2 \quad (25)$$

COMPUTATIONAL IMPLEMENTATION

The computational implementation of the proposed modeling framework follows the information flow illustrated in Figure 1, integrating all key components required to describe crystallization phenomena in battery systems. The system of ordinary differential equations resulting from the discretized population balance formulation was solved using the *solve_ivp* routine from the SciPy library, employing the implicit multistep Backward Differentiation Formula (BDF) method. This solver was selected for its robustness and numerical stability when handling stiff systems arising from the strong coupling between crystallization kinetics and electrochemical processes.

The discretization of the particle size domain required by the method of classes is inherently system-dependent and must be defined according to the characteristic crystal sizes, growth kinetics, and operating conditions of the electrochemical system under investigation. Both the extent of the size domain and the class width should therefore be chosen to adequately resolve the evolution of the crystal size distribution. In the present implementation, and consistent with the methodology established in the authors' previous work, the size domain was discretized using a uniformly spaced mesh spanning from 0 to 10 μm, with a class width of 0.05 μm. This discretization was shown to be appropriate for capturing the order of magnitude and temporal evolution of Li_2O_2 crystal growth during battery discharge. A uniform mesh was adopted because the population dynamics are governed exclusively by continuous nucleation and growth processes, with no contribution from discrete birth or death events.

The integral terms appearing in the mass and energy balance equations were evaluated using the *integrate.quad* function from the SciPy library, ensuring accurate numerical integration over the discretized size domain. Following implementation, the model outputs were systematically compared against selected response variables to verify numerical consistency and proper coupling among the governing equations, thereby validating the computational methodology employed in this study.

CONCLUSION

In this work, a population balance–based modeling framework was developed to describe the crystallization occurring during the discharge of Li–O_2 batteries. To enable an adequate formulation, several simplifying assumptions were adopted, which limit the quantitative accuracy of the model and its ability to capture highly complex operating scenarios, such as variable discharge currents, electrolyte compositions, or detailed electrode architectures. A key strength of the proposed approach lies in its ability to naturally account for crystal size distributions and morphological effects while maintaining a relatively low computational cost. By bridging microscopic crystallization mechanisms with macroscopic battery performance indicators, the framework provides physically meaningful insights without requiring excessively complex formulations. As such, it offers an accessible and flexible modeling platform for investigating crystallization-driven limitations in metal–air batteries. This study represents a step toward the systematic application of population balance modeling to Li–O_2 batteries, with a particular emphasis on formalizing and clarifying the

governing equations and numerical implementation. Future work may extend the present framework by incorporating additional mechanisms, such as side reactions, electrolyte-dependent effects, dynamic operating conditions, and alternative electrode designs. These developments would further enhance the predictive capability.

Symbology

a_e = interfacial area.
a_D = length scale of liquid Li_2O_2.
A_C = surface area of the crystals.
A_T = thermal exchange area.
B_0 (t)= nucleation function.
c_p = specific heat.
C = concentration.
C* = equilibrium saturation concentration
ΔC = supersaturation.
D_l= diffusivity constant of crystals.
esp = electrode thickness.
F = Faraday constant.
G^{crit} = critical formation energy of crystal nuclei.
G(L, t) = crystal growth function.
h_ξ = jump function (discrete events).
H_c = heat of formation.
I_c = current applied to the system.
j_c = interfacial transfer current density.
j_l = kinetic order of the growth law.
k_a = constant to obtain the crystal surface area.
k_{ano} = anodic kinetic constant.
k_b = Boltzmann constant.
k_c = cathodic kinetic constant.
k_d= mass transfer coefficient.
k_{sl}= diffusivity constant of the Li2O2 crystal.
k_v= constant to obtain the volume of the crystal.
K_r = electrolyte conductivity.
L = characteristic length.
MM = molar mass.
n = number of electrons in the reaction.
n_{crit} = critical nucleus size.
N = total number density.
N_A = Avogadro's number.
Q = battery capacity.
rel= electrode radius.
R = universal gas constant.
t = discharge time.
T = temperature.
U = overall heat transfer coefficient.
U_o = crystals velocity.
x_0 = nucleation site density.
V_c= volume of the crystals.
V_{susp}= effective volume of the system.
Z = Zeldovish factor.
ξ = internal coordinate vector of the PB.
ε = electrode porosity.
ρ_c = specific mass of Li2O2 crystals.
ρ_e= specific mass of the electrolyte.
η_r= effectiveness factor.
η_c= overpotential.
γ= surface energy.
θ = ratio of current density to limiting current density.

ACKNOWLEDGEMENTS

The authors acknowledge the financial support provided by the National Council for Scientific and Technological Development (CNPq, Brazil), the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES, Brazil; Finance Code 001), and the São Paulo Research Foundation (FAPESP, Brazil) through Grant Nos. 2018/16663-2, 2023/08594-9, and the Thematic Project Grant No. 2022/16379-8. The authors also gratefully acknowledge the Universidade Estadual de Campinas (UNICAMP) and the School of Chemical Engineering, Campinas, SP, Brazil, for the institutional support provided during the development of this work.

AUTHOR IDENTIFIERS

Author ORCIDs:
Khouri NG: 0000-0002-2954-4947
Leal Silva JF: 0000-0002-9915-5859
Barros LMS: 0000-0002-9786-8783
Concha VOC: 0000-0002-3520-4535
Maciel Filho R: 0000-0001-6511-7283

REFERENCES

1. International Renewable Energy Agency. *Renewable energy statistics 2025* https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2025/Jul/IRENA_DAT_RE_Statistics_2025.pdf
2. Reuters.https://www.reuters.com/business/energy/around-90-renewables-cheaper-than-fossil-fuels-worldwide-irena-says-2025-07-22
3. Conde HJC, Demition CM, Honra J. Storage is the new black: a review of energy storage system applications to resolve intermittency in renewable energy systems. Energies 18:354 (2025). https://doi.org/10.3390/en18020354
4. Adelhelm P, Hartmann P, Bender CL, Busche M, Eufinger C, Janek J. From lithium to sodium: cell chemistry of room temperature sodium–air and sodium–sulfur batteries. Beilstein J. Nanotechnol. 6:1016-1055 (2015). https://doi.org/10.3762/bjnano.6.105
5. Lampic G, Gotovac G, Geaney H, O'Dwyer C. Comparing the suitability of lithium ion, lithium sulfur and lithium air batteries for current and future vehicular applications. arXiv preprint

2016.06347 (2016) https://doi.org/10.48550/arXiv.1606.06347
6. Deng F, Zhang Y, Yu Y. Conductive metal–organic frameworks for rechargeable lithium batteries. Batteries 9:109 (2023). https://doi.org/10.3390/batteries9020109
7. Li S, Liu Y, Wu X, Wang K, Chen J. Tailoring the growth and morphology of lithium peroxide: nickel sulfide/nickel phosphate nanotubes with optimized electronic structure for lithium–oxygen batteries. Small 19: (2023). https://doi.org/10.1002/smll.202304435
8. Huang X, Li M, Gao Y, Park MG, Matsuda S, Amine K. Discharge rate?driven li ₂ o ₂ growth exhibits unconventional morphology trends in solid?state li?o ₂ batteries. Angew Chem Int Ed 64: (2025). https://doi.org/10.1002/anie.202507967
9. Horstmann B, Gallant B, Mitchell R, Bessler WG, Shao-Horn Y, Bazant MZ. Rate-dependent morphology of li₂o₂ growth in li–o₂ batteries. J. Phys. Chem. Lett. 4:4217-4222 (2013). https://doi.org/10.1021/jz401973c
10. Wang Y. Modeling discharge deposit formation and its effect on lithium-air battery performance. Electrochimica Acta 75:239-246 (2012). https://doi.org/10.1016/j.electacta.2012.04.137
11. Lau S, Archer LA. Nucleation and growth of lithium peroxide in the li–o₂ battery. Nano Lett. 15:5995-6002 (2015). https://doi.org/10.1021/acs.nanolett.5b02149
12. Yin Y, Gaya C, Torayev A, Thangavel V, Franco AA. Impact of li₂o₂ particle size on li–o₂ battery charge process: insights from a multiscale modeling perspective. J. Phys. Chem. Lett. 7:3897-3902 (2016). https://doi.org/10.1021/acs.jpclett.6b01823
13. Gagliardi Khouri N, Leal Silva JF, Sampaio Barros LM, Cárdenas Concha VO, Maciel Filho R. Lithium peroxide crystallization in lithium-air batteries: a population balance modeling approach to enhance energy storage. Journal of Energy Storage 140:118939 (2025). https://doi.org/10.1016/j.est.2025.118939
14. Costa CBB. Modelagem e controle ótimo do processo de cristalização do ácido adípico. PhD Thesis, Universidade Estadual de Campinas (2003) http://repositorio.unicamp.br/bitstream/REPOSIP/266500/1/Costa_CalianeBastosBorba_M.pdf
15. Endo Y, Chen R, Pui DYH. Theoretical consideration of permeation resistance of fluid through a particle packed layer. Powder Technology 124:119-126 (2002). https://doi.org/10.1016/s0032-5910(01)00479-x
16. Jiang K, Liu X, Lou G, Wen Z, Liu L. Parameter sensitivity analysis and cathode structure optimization of a non-aqueous li–o2 battery model. Journal of Power Sources 451:227821 (2020). https://doi.org/10.1016/j.jpowsour.2020.227821
17. Ling Z, Wang F, Fang X, Gao X, Zhang Z. A hybrid thermal management system for lithium ion batteries combining phase change materials with forced-air cooling. Applied Energy 148:403-409 (2015). https://doi.org/10.1016/j.apenergy.2015.03.080
18. Khouri NG, Bahú JO, Miranda NT, Batistella CB, Maciel MRW, Concha VOC, Maciel Filho R. Diesel upgrading: a modeling of its microemulsions. Fuel Processing Technology 239:107545 (2023). https://doi.org/10.1016/j.fuproc.2022.107545
19. Mersmann A. Calculation of interfacial tensions. Journal of Crystal Growth 102:841-847 (1990). https://doi.org/10.1016/0022-0248(90)90850-k
20. Kashchiev D, van Rosmalen GM. Review: nucleation in solutions revisited. Cryst. Res. Technol. 38:555-574 (2003). https://doi.org/10.1002/crat.200310070
21. Yuan H, Read JA, Wang Y. Capacity loss of non-aqueous li-air battery due to insoluble product formation: approximate solution and experimental validation. Materials Today Energy 14:100360 (2019). https://doi.org/10.1016/j.mtener.2019.100360
22. Barton JL, Milshtein JD, Hinricher JJ, Brushett FR. Quantifying the impact of viscosity on mass-transfer coefficients in redox flow batteries. Journal of Power Sources 399:133-143 (2018). https://doi.org/10.1016/j.jpowsour.2018.07.046
23. La Cerva M, Gurreri L, Tedesco M, Cipollina A, Ciofalo M, Tamburini A, Micale G. Determination of limiting current density and current efficiency in electrodialysis units. Desalination 445:138-148 (2018). https://doi.org/10.1016/j.desal.2018.07.028
24. Gallant BM, Kwabi DG, Mitchell RR, Zhou J, Thompson CV, Shao-Horn Y. Influence of li2o2 morphology on oxygen reduction and evolution kinetics in li–o2 batteries. Energy Environ. Sci. 6:2518 (2013). https://doi.org/10.1039/c3ee40998h

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Computational Framework for Simulation and Energy Evaluation of Sustainable Biodiesel Production Routes

Ian B. B. Batata[a], Emílio E X. Guimarães Filho[a], Victor H. S. Ramos[a], Maria R. Wolf Maciel[a], Nadia G. Khouri[a]*, Rubens Maciel Filho[a]

[a] University of Campinas (UNICAMP), School of Chemical Engineering, Campinas, São Paulo, Brazil.
* Corresponding Author: nadiagk@unicamp.br.

ABSTRACT

The growing global energy demand and the need to reduce dependence on fossil fuels have intensified efforts toward developing renewable alternatives. Among these, biodiesel and ethanol emerged as viable and sustainable fuel sources. In this context, the use of palm oil and ethanol as raw materials represents a promising production route, due to their availability, high productivity per unit of cultivated area, and renewable characteristics. However, ethanolic transesterification still faces challenges, such as lower productivity compared to methanolic processes and higher energy consumption due to reaction characteristics that impact the whole process. Bearing this in mind, this work aims to develop process simulations in Aspen Plus to optimize biodiesel production from palm oil and ethanol, coupled with an energy integration analysis. The process was divided into three main stages: (i) feed preparation, (ii) transesterification reaction, and (iii) separation and purification of biodiesel and glycerin. The key operational parameters evaluated were temperature, pressure, and ethanol:palm oil molar ratio, important factors to ensure reaction efficiency and biodiesel quality. Additionally, energy integration was performed using the Pinch Analysis method, which enabled the identification of major energy demands and potential opportunities for process improvement. Overall, the proposed simulation framework contributes to enhancing the competitiveness of biodiesel production through ethanolysis, reinforcing its role as a sustainable and viable alternative within the global energy mix.

Keywords: Aspen Plus, Biofuels, Ethanol, Methanol, Simulation

INTRODUCTION

Energy is essential in modern lifestyles and global demand has increased rapidly due to rising consumption. At present, fossil fuels such as petroleum, diesel, coal, and natural gas are widely used for power generation and mechanical energy [1]. However, the growing reliance on fossil fuels is a major concern because of their finite availability and significant environmental impacts [2]. Non-renewable energy sources dominate the global energy mix, and their combustion releases large quantities of greenhouse gases, including carbon dioxide and methane [3]. In addition, fossil fuel extraction and transportation often cause substantial environmental damage, such as habitat degradation, oil spills, and water pollution [4].

Biofuels have emerged as promising and sustainable alternatives that are derived from organic materials (plant biomass and agricultural residues) [2]. Among them, bioethanol and biodiesel show strong potential for integration into the global energy system [5]. Driven by rising oil prices and growing concerns over climate change, research into renewable biofuel technologies has intensified. Biofuels can reduce net carbon dioxide emissions because it is partially offset by the carbon dioxide absorbed by plants during growth, forming a near-closed carbon cycle. In addition to environmental benefits, biofuels reduce dependence on imported fossil fuels, enhance energy security, and create new opportunities for agricultural development and rural economic growth [6].

Palm oil and ethanol are promising biofuel feedstocks due to their high availability, renewable nature, and strong potential for sustainable large-scale energy production. Biodiesel, produced as monoalkyl esters from

vegetable oils or animal fats through transesterification with alcohol in the presence of a catalyst, can substitute conventional diesel fuels. The suitability of vegetable oil feedstocks for biodiesel production depends on regional soil and climatic conditions, with palm oil standing out for its high yields and efficiency [7]. Consequently, the palm oil sector supports food and biomass demands while ensuring economic decarbonization [8]. Ethanol is one of the most established and scalable biofuels globally, benefiting from mature production technologies [1]. Furthermore, the utilization of palm oil residues and ethanol-based pathways enhances sustainability by converting agricultural and processing wastes into biofuels, thereby improving resource efficiency and reducing environmental impacts.

Despite the advantages of ethanol as a renewable and less toxic alcohol derived from biomass, its application in biodiesel production remains technically challenging. Comparative studies indicate that ethanolic transesterification achieves lower conversions than methanolic routes, requiring higher reaction temperatures, longer reaction times, and higher alcohol-to-oil molar ratio [9-10]. These process characteristics increase overall energy consumption and reduce productivity, largely due to ethanol's higher polarity, steric hindrance, and reduced catalyst activity [11].

Therefore, this study focuses on the development of detailed process simulations using Aspen Plus to evaluate and optimize biodiesel production from palm oil via ethanolic transesterification, alongside a comprehensive energy integration analysis. The same framework is applied to the methanolic transesterification to compare the more sustainable process to a conventional one. Thereby, this work improves the technical feasibility of ethanol-based biodiesel production by systematically assessing process performance and energy efficiency, strengthening its potential as a sustainable and competitive fuel option.

PROCESS DESCRIPTION

This work models two biodiesel production routes in Aspen Plus V14, based on palm oil transesterification using sodium hydroxide as catalyst with short-chain alcohols: methanolysis and ethanolysis. In both cases, the plant was simulated with a production capacity of 1000 kg/h of biodiesel. The flowsheets were constructed to enable a consistent comparison of reaction performance, alcohol recovery, separation requirements, and energy demand. Both routes follow the same process structure, feed basis, and product specifications, with changes limited to the alcohol identity and dosage and the downstream separation steps required for alcohol recycling and product purification. Figure 1 shows the flowsheet of the biodiesel plant. The palm oil feed was represented using the triglyceride composition listed in Table 1. Phase equilibria and activity coefficients were calculated using the NRTL property method.

Table 1: Palm oil composition used to define the feed stream.

Fatty Acid	Palm Oil
Dyglyceride	5.20.
Triglyceride	94.80
Triunsaturated	
OOO	4.40
OOL	0.58
Total	4.98
Monosaturated	
PLO	9.68
POO	23.26
OOS	2.24
Total	35.18
Disaturated	
MPL	2.20
PPL	9.23
PPO	29.62
POS	4.90
Total	45.95
Trisaturated	
MMM	0.42
MMP	1.70
PPP	5.51
PPS	1.06
Total	8.69

The selected alcohol (methanol or ethanol) and palm oil feeds were combined and sent to a transesterification reactor operating with an alcohol-to-palm oil molar ratio of 6:1 and 1% NaOH catalyst (based on the oil feed) [12]. In the reactor, triglycerides were converted to fatty acid alkyl esters (FAAE) as the main product, namely FAME (fatty acid ethyl ester) in the methanolysis route and FAEE in the ethanolysis route, with glycerol formed as the byproduct. The transesterification reactor was modeled in Aspen Plus as a continuous stirred-tank reactor (CSTR). The reaction system and kinetic parameters were taken from Narváez et al. (2015) [13] and applied to both alcohols in the comparative assessment.

Conversion was calculated using Eq. (1), defined as the ratio of fatty acid alkyl esters (FAAE) to three times the initial triacylglyceride (TAG) amount on a molar basis. For ethanolysis, FAAE corresponds to FAEE, whereas for methanolysis it corresponds to FAME.

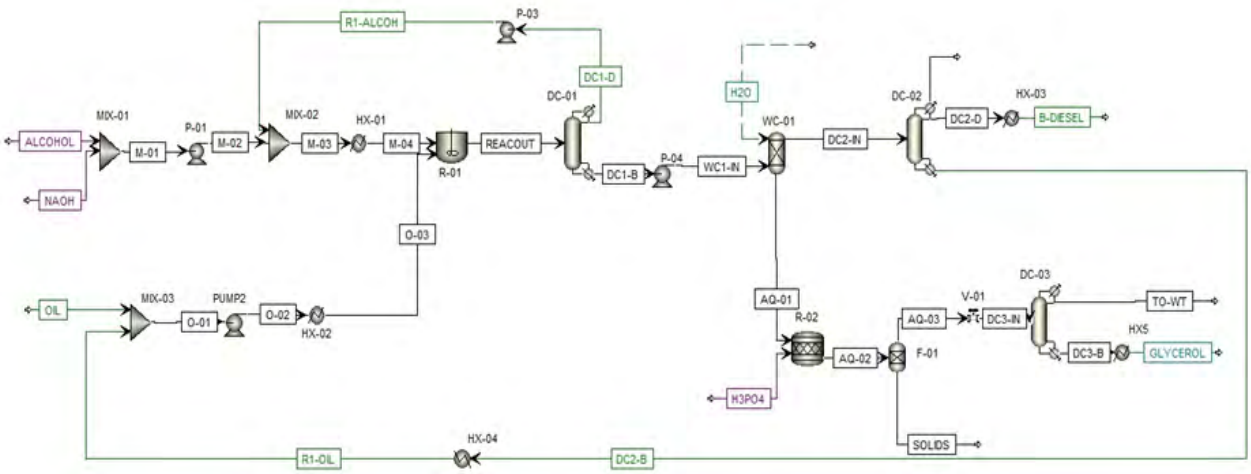

Figure 1. Process flowsheet used to evaluate methanolic and ethanolic transesterification routes for biodiesel production

$$Conversion\ (\%) = \frac{FAAE}{(3 * TAG)} . 100 \qquad 1)$$

The reactor effluent was then sent to the first separation unit, a distillation column (DC-01), which was specified to recover unreacted alcohol in the destilled stream. The alcohol-rich distillate was recycled to the reactor. The bottom stream from DC-01 was routed to the washing column (WC-01), where water was used as the washing medium to promote phase splitting into an aqueous phase and an organic phase. The aqueous phase was withdrawn as the bottom product and contained residual alcohol, NaOH, produced glycerol, and residual water. The organic phase was removed as the top stream, consisting mainly of biodiesel (FAAE) and unreacted glycerides (tri-, di-, and monoglycerides), and was sent to the subsequent separation column (DC-02). DC-02 was simulated for the purpose of purifying biodiesel in the distillate stream. The bottom stream, rich in unreacted glycerydes, was then recycled to the reactor.

The aqueous bottom stream from WC-01 was directed to a neutralization step to remove the homogeneous catalyst. Neutralization was modeled as a stoichiometric reactor (RSTOIC) considering the reaction between NaOH and phosphoric acid 20% w/w (H_3PO_4), producing inorganic salts. These solids were removed in a filtration unit (FILTER). The clarified aqueous phase was then fed to a distillation column for glycerol purification. Detailed operating specifications and performance indicators for each unit are reported in the Results and Discussion section.

RESULTS AND DISCUSSION

Process Results

This section compares the methanolysis and ethanolysis routes by adopting an excess-alcohol strategy as the starting point and then examining how the alcohol identity affects reaction performance, separation behavior, and overall process requirements. In addition to biodiesel formation, particular attention is given to the remaining unreacted glycerides in each route, since their accumulation directly reflects incomplete conversion and losses of valuable feedstock. The mass of unreacted glycerides is therefore quantified and discussed along with the role of recycling these streams back to the reactor, which increases overall feed utilization and supports higher biodiesel yields at the plant level. All stream data for the methanol and ethanol routes will be presented in an Excel spreadsheet, which will be made available as supplementary material.

A biodiesel production rate of 1000 kg/h was imposed as a design basis. Aspen Plus Design Spec was used to calculate the required fresh feeds of palm oil and pure alcohol so that, after mixing with the recycle streams, the overall alcohol-to-palm oil molar ratio entering the reactor remained at 6:1. In the methanolysis route, the reactor was operated at 50 °C with a residence time of 80 min [14], achieving a conversion of 96.35%. In the ethanolysis route, the reactor was operated at 70 °C with a residence time of 60 min [13], achieving a conversion of 96.42%.

The distillation columns were designed using a sequential modeling strategy in Aspen Plus. A shortcut model (DSTWU) was first employed to establish an initial design and provide starting estimates for key variables. The columns were then implemented in a rigorous framework using RadFrac, and separation targets were met through the Design Spec tool. Next, a sensitivity analysis was performed to identify the number of stages and the feed stage for each column, aiming to minimize reboiler

energy consumption while maintaining the required separation performance. Finally, hydraulic checks were carried out and column hydraulics were adjusted as needed to ensure that all stages operated within acceptable ranges. Table 2 and 3 summarizes the main design and operating parameters adopted for the distillation columns in the simulated flowsheets.

Table 2: Distillation column specifications used in the Aspen Plus V14 simulations for the methanolysis route.

	DC-01	DC-02	DC-03
Purpose	Alcohol recovery	Biodiesel purification	Glycerol purification
Number of stages	10	11	5
Feed stage	2	2	2
Reflux ratio (molar)	2.00	1.70	2.00
D:F (molar)	0.38	0.98	0.82
Operating Pressure (bar)	0.20	0.10	0.10

Table 3: Distillation column specifications used in the Aspen Plus V14 simulations for the ethanolysis route.

	DC-01	DC-02	DC-03
Purpose	Alcohol recovery	Biodiesel purification	Glycerol purification
Number of stages	23	20	5
Feed stage	2	2	2
Reflux ratio (molar)	2.53	0.38	2.00
D:F (molar)	0.44	0.96	0.92
Operating Pressure (bar)	0.20	0.10	0.10

The table reports, for each column, the molar reflux ratio, the distillate-to-feed ratio, the total number of equilibrium stages, and the feed stage location. These parameters are used to describe the separation strategy applied in each column and to support the comparison of column configurations between the methanolysis and ethanolysis routes. DC-01 specifications indicate that alcohol recovery is the main separation step distinguishing the two routes. Ethanolysis required a markedly higher column complexity, with 23 stages and a molar reflux ratio of 2.53, whereas methanolysis achieved the same task with 10 stages and a reflux ratio of 2.00, at the same operating pressure (0.20 bar). This higher separation effort is consistent with the ethanol–water azeotropic constraint in the presence of water, which limits alcohol dehydration by conventional distillation and increases the reflux and staging needed to meet recycle requirements. As a direct consequence, the ethanolysis configuration is expected to demand higher condenser and reboiler loads due to increased internal vapor and liquid traffic, which is reflected in the overall utility consumption of the separation section.

The bottom stream from the DC-01 distillation column was pumped to the washing column (WC-01). WC-01 was specified using an Aspen Plus Design Spec to determine the required water flow rate such that 90% of the residual alcohol was recovered in the aqueous phase withdrawn as the bottom product. The top stream from WC-01 was sent to DC-02, which was modeled with a partial condenser to achieve a biodiesel distillate with a molar purity of 99%. Notably, DC-02 adopted different stage and reflux configurations in the two routes, reflecting distinct separation strategies. The purified biodiesel was then cooled to 35 °C prior to storage and packaging.

In methanolysis, biodiesel purification was achieved with fewer stages (11) but a higher molar reflux ratio (1.70), whereas ethanolysis required more stages (20) but operated at a much lower reflux ratio (0.38). This contrast indicates a trade-off between column complexity and internal reflux demand: increasing the number of stages tends to reduce the reflux needed to meet a given distillate purity, but it increases column height and, consequently, capital requirements. Conversely, operating with fewer stages can be compensated by higher reflux, which increases internal vapor and liquid traffic and typically raises condenser and reboiler duties, increasing operating costs. These route-dependent designs suggest that the alcohol identity, along with the associated phase behavior in the separation section, shifts the balance between capital-lean and energy-lean configurations for achieving the same biodiesel specification.

The bottom stream from WC-01 was sent to the neutralization reactor. The required H_3PO_4 feed rate was calculated using the Aspen Plus calculator tool based on the NaOH molar flow rate in the WC-01 bottom stream, ensuring stoichiometric neutralization of the dissolved catalyst. Neutralization produces sodium phosphate salts, which were removed in a filtration unit prior to sending the clarified stream to the glycerol purification section. The resulting aqueous stream was then processed in DC-03 to concentrate and purify glycerol, completing the liquid-phase treatment train associated with catalyst removal and byproduct recovery. In contrast to DC-01 and DC-02, glycerol purification exhibited similar column requirements in both routes, indicating that the main process differences are concentrated in the alcohol recovery and biodiesel purification steps rather than in the final glycerol finishing section. With the separation tasks defined, the next section evaluates how these route-dependent column configurations translate into utility consumption and opportunities for heat integration.

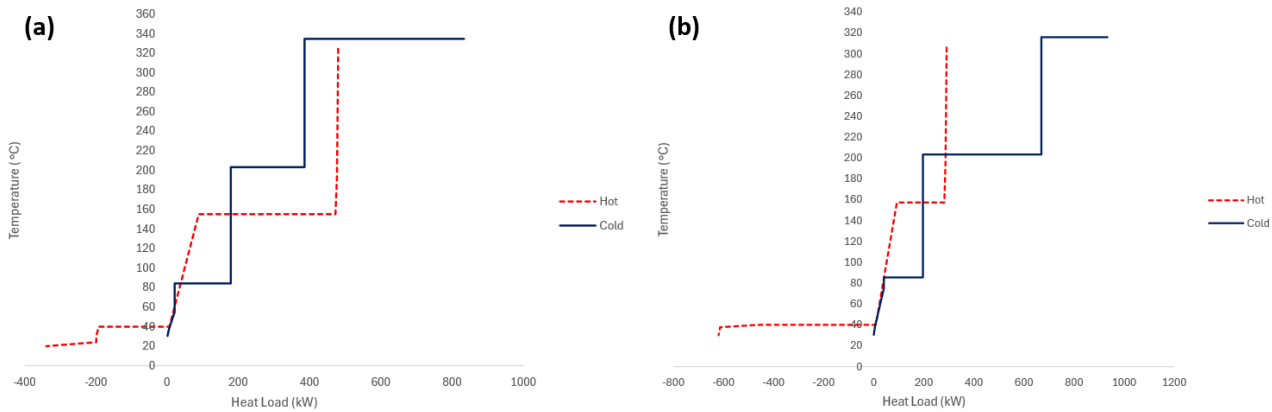

Figure 2. Composite curves for both evaluated scenarios. (a): methanolic route; (b): ethanolic route.

Energy Integration

Heat integration was evaluated in both scenarios (using methanol and ethanol) to investigate whether the choice of the primary alcohol in the transesterification reaction could provide energetic advantages. By applying the concepts of the pinch analysis, it is possible to identify opportunities for heat recovery and utilization between hot and cold utilities while respecting thermodynamic limits. A ΔT_{min} = 10 °C was adopted for initial insights. Focus was given to the existing heat exchangers and distillation columns, the main energy consumers in continuous biodiesel plants [15]. Also, it is assumed that there are no restrictions in energy exchange among all process streams.

Overall, the outcomes were similar in both cases. The pinch temperatures were determined as being 36.5 °C and 38.2 °C for the methanol and ethanol scenarios, respectively. It means that it is likely to have a similar profile in the heat cascade between them. Also, such low values of pinch temperatures indicate that the energy bottleneck occurs near the cold end of the process; in other words, the process is well integrated as the temperature intervals and heat capacities are similar throughout the process. In the methanol case, it was calculated a total heat recovery of 819 kW which comprehends the totality of the availability of the hot utilities; the same happens in the ethanol case, where the total heat recovery is 910.9 kW – the totality of the hot utilities of this specific case (Figure 2). The high energy consumption of the conventional biodiesel production and pinch points are similar to the findings in Plesu et al. (2015) [16].

As can also be observed, the ethanolic route has a higher energy demand. That's mainly because of two key factors: 1) the ethanol flow rate and molar proportion ethanol:oil in the feed stream is higher compared to the methanol; 2) the optimal temperature of the transesterification step for the ethanolic route was determined as 70 °C, which is 40% compared to the methanolic approach and thus promoting a slightly increase in utility demand.

CONCLUSION

This work developed a computational framework to simulate and compare methanolic and ethanolic biodiesel production routes from palm oil, with integrated energy analysis. Both routes achieved similarly high conversions (≈96%), indicating that ethanolysis is technically comparable to the conventional methanolysis process. The main differences between the routes were found in the separation section, particularly in alcohol recovery and biodiesel purification. While the ethanolic route required more energy input due to the challenging ethanol–water separation, these differences were subtle overall. The pinch analysis revealed similar thermal profiles, with effective utilization of available heat in each scenario.

These results indicate that the technical gap between methanolic and ethanolic biodiesel production is narrower than expected. Consequently, targeted improvements in separation design, alcohol recovery strategies, and process intensification will enhance the technical feasibility of ethanolysis. Given the renewable nature of ethanol, further investigation of the ethanolic route is essential to advance biodiesel production toward more sustainable configurations.

DIGITAL SUPPLEMENTARY MATERIAL

A spreadsheet containing the complete stream data for both the methanolic and ethanolic routes is provided as digital supplementary material and is accessible via hyperlink, as referenced in the Results and Discussion section: Stream data for the methanolic and ethanolic transesterification routes

ACKNOWLEDGEMENTS

The authors acknowledge the financial support provided by the National Council for Scientific and Technological Development (CNPq, Brazil) and the São Paulo Research Foundation (FAPESP, Brazil) through the Thematic Project Grant No. 2022/16379-8. This study was financed by the Human Resources Training Program of the Brazilian National Agency for Petroleum, Natural Gas and Biofuels – PRH/ANP, under the management of the São Paulo Research Foundation (FAPESP), Brazil, Process N. 2025/20397-0. The authors also gratefully acknowledge the Universidade Estadual de Campinas (UNICAMP) and the School of Chemical Engineering, Campinas, SP, Brazil, for the institutional support provided during the development of this work.

AUTHOR IDENTIFIERS

Author ORCIDs:
Batata IBB: 0009-0006-5764-2532
Guimarães Filho EEX: 0009-0007-5426-2187
Ramos VHS: 0000-0001-5121-7814
Wolf Maciel MR: 0000-0002-5922-1428
Khouri NG: 0000-0002-2954-4947
Maciel Filho R: 0000-0001-6511-7283

REFERENCES

1. Malik K, Capareda SC, Kamboj BR, Malik S, Singh K, Arya S, Bishnoi DK. Biofuels production: a review on sustainable alternatives to traditional fuels and energy sources. Fuels 5:157-175 (2024). https://doi.org/10.3390/fuels5020010
2. Kazmi A, Sultana T, Ali A, Nijabat A, Li G, Hou H. Innovations in bioethanol production: a comprehensive review of feedstock generations and technology advances. Energy Strategy Reviews 57:101634 (2025). https://doi.org/10.1016/j.esr.2024.101634
3. Cherwoo L, Gupta I, Flora G, Verma R, Kapil M, Arya SK, Ravindran B, Khoo KS, Bhatia SK, Chang SW, Ngamcharussrivichai C, Ashokkumar V. Biofuels an alternative to traditional fossil fuels: a comprehensive review. Sustainable Energy Technologies and Assessments 60:103503 (2023). https://doi.org/10.1016/j.seta.2023.103503
4. Das P, Ghosh B, Sarkar S. Biofuels: a sustainable alternative to fossil fuels. Microbial Menace and Miracle: The Double-Edged Sword of Infection and Immunity :317-322 (2025). https://doi.org/10.65525/svup.9788198471956.2025.317-322
5. Mazariegos I, Abdelfath-Aldayyat E, González-Rojo S, Gómez X. Reducing fossil fuel demand by using biofuels as an alternative hydrothermal liquefaction is a promising process for transforming biomass into drop-in fuels. RSC Sustainability 3:3228-3265 (2025). https://doi.org/10.1039/d5su00148j
6. Singh R, Gaur A, Soni P, Jain R, Pant G, Kumar D, Kumar G, Shamshuddin S, Mubarak NM, Dehghani MH, Suhas , Ansari K. A review of biofuels and bioenergy production as a sustainable alternative: opportunities, challenges and future perspectives. J Environ Health Sci Engineer 23: (2025). https://doi.org/10.1007/s40201-025-00946-0
7. Lamba BY, Jain S, Jha S. Biodiesel from palm vegetable oil. NR 15:51-60 (2024). https://doi.org/10.4236/nr.2024.152004
8. Munar-Flórez DA, Ramírez-Contreras NE, Albarracín-Arias JA, Arias-Camayo P, Rincón-Romero V, García-Núñez JA, Ardila-Badillo C, Cuéllar-Sánchez M. Decarbonizing aviation: the low-carbon footprint and strategic potential of colombian palm oil for sustainable aviation fuel. Energies 18:4978 (2025). https://doi.org/10.3390/en18184978
9. Rashid U, Ibrahim M, Ali S, Adil M, Hina S, Bukhari IH, Yunus R. Comparative study of the methanolysis and ethanolysis of maize oil using alkaline catalysts. Grasas aceites 63:35-43 (2012). https://doi.org/10.3989/gya.06891
10. García M A, Gonzaloa A, Sánchez L José, Arauzo J A, Simoes C. Methanolysis and ethanolysis of animal fats: a comparative study of the influence of alcohols. CI&CEQ 17:91-97 (2011). https://doi.org/10.2298/ciceq100224058g
11. Sánchez BS, Benitez B, Querini CA, Mendow G. Transesterification of sunflower oil with ethanol using sodium ethoxide as catalyst. effect of the reaction conditions. Fuel Processing Technology 131:29-35 (2015). https://doi.org/10.1016/j.fuproc.2014.10.043
12. Souza MF, Hirata GF, Batista EAC. Evaluation of kinetics and thermodynamic parameters for simulation of palm oil biodiesel production. Fluid Phase Equilibria 525:112792 (2020). https://doi.org/10.1016/j.fluid.2020.112792
13. Narváez PC, Noriega MA, Cadavid JG. Kinetics of palm oil ethanolysis. Energy 83:337-342 (2015). https://doi.org/10.1016/j.energy.2015.02.029
14. Narváez PC, Rincón SM, Sánchez FJ. Kinetics of palm oil methanolysis. J Americ Oil Chem Soc 84:971-977 (2007). https://doi.org/10.1007/s11746-007-1120-y
15. Boon-anuwat NN, Kiatkittipong W, Aiouache F, Assabumrungrat S. Process design of continuous biodiesel production by reactive distillation: comparison between homogeneous and heterogeneous catalysts. Chemical Engineering

and Processing: Process Intensification 92:33-44 (2015). https://doi.org/10.1016/j.cep.2015.03.025
16. Ple?u V, Subirana Puigcasas J, Benet Surroca G, Bonet J, Bonet Ruiz AE, Tuluc A, Llorens J. Process intensification in biodiesel production with energy reduction by pinch analysis. Energy 79:273-287 (2015). https://doi.org/10.1016/j.energy.2014.11.013

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Enhancing Pharmaceutical Supply Chain Robustness via Simulated Annealing

Nelson Chibeles-Martins[ab]*, Maria A. Monge[c] and Tânia Pinto-Varela[cd]

[a] Departamento de Matemática, Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, Caparica, Portugal
[b] NOVA MATH, Centro de Matemática e Aplicações, Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, Caparica, Portugal
[c] Centro de Estudos de Gestão, Instituto Superior Técnico, Universidade Lisboa, Lisboa, Portugal
* Corresponding Author: npm@fct.unl.pt.

ABSTRACT

The pharmaceutical sector is essential for ensuring universal access to medicines, demanding efficient supply chains that deliver drugs at optimal prices with minimal delays and shortages. Pharmaceutical supply chains (PSCs) face significant challenges, including strict quality controls, government regulations, drug perishability, high R&D costs, and complex transportation requirements. The sector is undergoing a shift, driven by the rise of pharmaceutical components in emerging markets, unpredictable demand, and reduced R&D investments by major companies, which struggle to compete with generic pharmaceutical brands. Post-pandemic challenges and geopolitical risks have further exposed vulnerabilities in PSCs, leading to frequent supply disruptions, product shortages, and unreliable transportation. The increasing focus on regionalization highlights the need for more resilient supply chains to manage disruptions effectively. PSCs must incorporate robustness to address uncertainties and disruptions, which can be internal, such as operational, financial, and quality risks, or external, including supply, demand, and environmental uncertainties. While recent research emphasizes the importance of integrating these uncertainties into decision-making tools, the field still lacks comprehensive solutions for PSC optimization. Previous studies have used heuristic and metaheuristic algorithms to address individual uncertainties and disruptions, providing feasible solutions with lower computational effort. However, a more holistic approach is required. This work introduces a new tool based on a Meta-heuristic algorithm, which evaluates the robustness of PSCs under multiple uncertainties and disruptions. The tool enables the development of what-if scenarios, helping decision-makers improve supply chain robustness and adaptability in the face of several risks

Keywords: Supply Chain, Algorithms, Modelling and Simulations, Optimization, Simulated Annealing, Uncertainty

INTRODUCTION

An essential role of pharmaceuticals is to provide stability and sustainability to any healthcare system, being one of the key areas of a country's economy. Ensuring demand satisfaction for pharmaceutical products and having a perfect service level with acceptable costs is crucial for every country. In top of that, pharmaceutical drugs gather special characteristics, like extensive and complex production processes, product perishability, intensive investment in R&D activities, and stringent transportation conditions that only create harder conditions for companies to achieve the desired service level.

For this reason, pharmaceutical supply chains (PSCs) need to guarantee a ceaseless flow of products throughout all stages and respond quickly and efficiently to possible changes and disruptions. Over the years, they became more complex, globalized, and subject to more strict safety procedures, quality controls, and regulations from governments and insurers. It is imperial to develop strategies that account for all these factors when managing and designing a network.

Furthermore, a good balance between all the factors and trade-offs affecting the performance of the PSC is

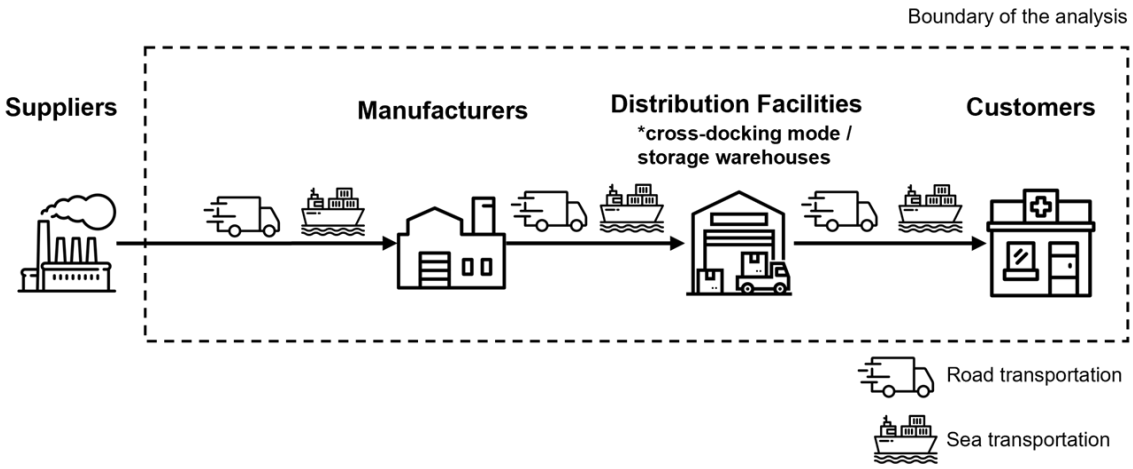

Figure 1. Pharmaceutical supply chain

essential, from resource allocation to inventory management and the location of entities.

Nowadays, the pharmaceutical sector is under a paradigm shift. The growing of pharmaceutical components in emerging countries and the unpredictability of markets is creating an additional layer for companies to account for and exposes them to more vulnerabilities. Moreover, the decrease of new pharmaceutical drugs by research-based companies due to reductions in R&D investment, price pressures, and the shortening of patent lives is making it difficult for these companies to maintain their performance and compete with generic pharmaceutical enterprises [1]. Drug shortages are also increasing due to supply, demand, and regulatory issues, which are also affecting the price and the production of its main components [2]. More recently, the impacts of the COVID-19 pandemic also affected this sector, as the world's main suppliers restricted production to domestic, and started to move to self-sufficiency and decentralized models [3].

All these challenges have a high degree of uncertainty associated with them, making the task of designing and operating an effective and responsive supply chain even more difficult.

While deterministic models have traditionally provided a baseline for pharmaceutical supply chain (PSC) design, they often fail to capture the inherent volatility of real-world operations, leading to vulnerable networks and ineffective strategies. In contrast, the metaheuristic approach proposed in this work allows for the integration of multiple uncertainties—such as raw material supply and transportation network disruptions—simultaneously. This provides decision-makers with the ability to develop "what-if" scenarios and assess the resistance of specific network entities, identifying a clear need for increased redundancy and decentralization as strategies to thrive under severe uncertainty.

Recent research shows that the field of optimizing PSC network design is still missing focus. The transition toward uncertainty-aware optimization is supported by a growing body of literature that identifies uncertainty quantification as a core requirement for process engineering. [1, 4] highlight that while the pharmaceutical sector has been extensively studied, most research still focuses on inventory models rather than the global optimization of network design under uncertainty.

In the last years, approaches based on Heuristics and Metaheuristics algorithms have been applied, since they allow the discovery of a good feasible solution with less computational effort than exact methods. Operational risks and their inherent uncertainties, encountered in customer demand, costs, and the capacity of facilities have been addressed in a large scale by the literature [5, 6, 7]. Particularly, [6, 7] have demonstrated the efficiency of hybrid metaheuristics in solving multi-objective PSC problems, proving that non-exact approaches can provide high-quality solutions for NP-hard network configurations. Furthermore, [8] emphasizes how deterministic models significantly underestimate risks in post-pandemic landscapes. Additionally, some authors have also been able to quantify the resilience of a network and model the effect of disruptions on production [4] and suppliers [9]. On top of that, disruptions in the transportation network were studied by Roshan [10], that developed a model for re-routing planning in the event of natural disasters.

By building on these foundations, this work addresses the existing research gap in simultaneous multi-uncertainty assessment through an adaptable Simulated Annealing algorithm. In the last years, approaches based on Heuristics and Metaheuristics algorithms have been applied, since they allow the discovery of a good and feasible solution for complex problems within a lower computational effort.

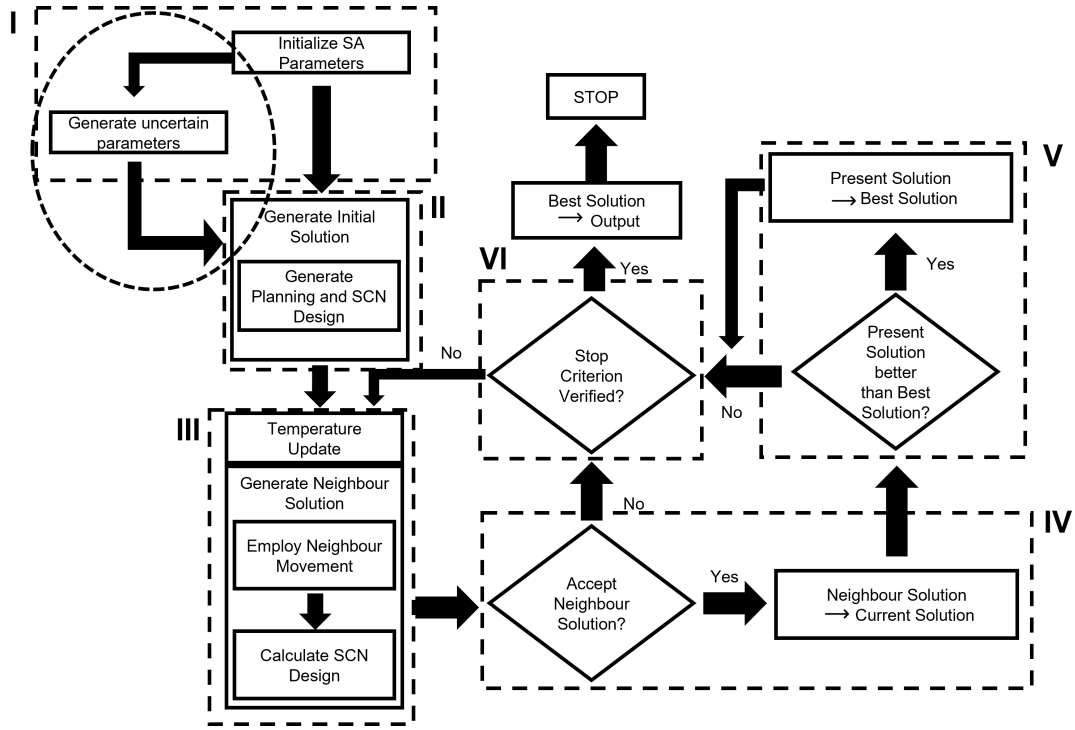

Figure 2. SA algorithm under uncertainty

PROBLEM CHARACTERIZATION

This work characterizes a design and planning of a global pharmaceutical supply chain (PSC) with 4-levels, shown in Figure 1.

The concerned PSC is defined by 4 stages: suppliers, manufacturers, distribution centers (DCs) and markets. Raw materials are delivered to factories by suppliers. After production, the final products are transported to distribution facilities, and from those to markets to satisfy demand. The distribution facilities can operate in a cross-docking mode, or function as typical warehouses and accommodate inventory of final products, that can also exist at the factories. The transportation between different levels is carried out using unimodal transportation, by road, or intermodal transportation, considering road and sea. In the latter, products are first transported from the origin entity to a seaport, where they are transshipped to a ship. Subsequently, they are transported to another seaport and finally to the end destination by road. Each entity has a seaport assigned, that works in a cross-docking mode. Supply, production levels, materials flow across entities and material storage handling are limited between specific bounds. Additionally, uncertainty is considered in some parameters of the problem, namely in production capacity. A possible disruption on the transportation network is also taken into account.

The aim of the problem is to define the PSC design over uncertainty alongside minimizing the total network cost. For that reason, the cost of factories and DCs installation, operational, storage, labor, transportation, and raw material are considered. Transportation costs include variable costs, arising from the distances travelled, contracted (fleet outsourced) and handling costs at the seaports.

From a strategic view, the entities are chosen from a set of potential locations. At a tactical level, the supply and production quantities, in addition to the materials flow within the network are determined.

At the same time, by studying this problem, it is possible to assess the severity of the impact of uncertainty across the multiple levels of the supply chain and understand the degree of robustness of the PSC that results from it.

METHODOLOGY

The Simulated Annealing (SA) algorithm is extended from its traditional form [11] to consider uncertainties in the design of the pharmaceutical supply chain. The SA algorithm under uncertainty considers 6 main steps, as shown in Fig. 2. The inclusion of uncertainty is defined at Procedure I, circled in dashes.

Procedure I is the key novelty in the algorithm, hence it is the only procedure highlighted and explained in more detail. This procedure is related with the initialization of the algorithm and determination of the parameters tainted with uncertainty. They are computed in the following way:

Uncertainty Generation Characterization. For each parameter that can be uncertain or disrupted, generate an arbitrary number $x \in [0,1]$. Considering the disruption (pd) or uncertainty probability (pu) of each of those parameters:

- If $x < pd/pu$, the parameter changes to a new value.
- If $x \geq pd/pu$, the value of the parameter remains the same.

After those parameters are settled, the algorithm is initiated, and moves to the next procedure, in order to find a compromise solution.

CASE STUDY

A real case study of a global supply chain from a pharmaceutical company that produces and delivers vaccines all over the world is explored.

The pharmaceutical company in question can install entities across 9 different locations. There are five possible manufacturing sites dedicated to the production of vaccines: U.S (Pennsylvania), Canada (Toronto), Europe (France), Asia (India) and Africa (Kenya). These manufacturers, when installed, receive the required raw materials from a supplier of the same region. Distribution centres can be installed in every location that contains a market, such as U.S (Pennsylvania), Canada (Toronto), Europe (France), Asia (India), Latin America (Brazil), Middle East (Israel), Eurasia (Russia) and Australia. In addition, each of these regions has a seaport that can be used for transportation purposes.

The aim of the case-study unfolds in two ways. 1) to determine which is the network design of the PSC capable of satisfying the total markets' demand with minimum cost, in terms of location, number of entities, and materials allocation. 2) analyse the impact of some uncertainties in the design and its costs, identifying which entities are robust. For that purpose, a baseline scenario and two additional ones are explored considering different types of uncertainty. The scenarios characterization is detailed in Table **1**, followed by the uncertainty or disruption probability associated to each of them, in Table 2.

An exploratory analysis was performed for the algorithm parametrization using: initial temperature (T_0), number of iterations at the same temperature (NIST) and cooling rate (α). The best values found were T_0 = 7500, NIST = 50 and α = 0.95. As a way of not falling in local minimums with higher network costs, an outer loop was added to the algorithm. For each run, the SA algorithm restarts its process 7 times, returning the solution with minimum costs encountered.

Table 1: Scenarios' characterization.

Scenario	Description
A: No Uncertainty	All the parameters of the PSC are considered deterministic.
B: Production Capacity	The maximum production capacity level of each manufacturer can decrease or maintain its original value.
C: Disruption on the transportation network	Some of the routes used for transportation between factories and DCs, and DCs and markets from the same region can suffer a disruption and become unavailable

Table 2: Scenarios' characterization.

Scenario	Disruption/ Uncertainty Probability (pd, pu)
A: No Uncertainty	Baseline
B: Production Capacity	$pu_{production} = 0.3$
C: Disruption on the transportation network	$pd_{transport} = 0.3$

RESULT ANALYSIS

The algorithm was performed 100 times for each scenario, and the solutions were analyzed considering their total network cost, general topology, and main characteristics of the supply chain regarding number, location of entities and possible reasons for the designs. The scenario **A** is deterministic and defined as the baseline. Therefore, all the remaining scenarios are compared with the baseline. The computational time required for an individual run in each of the scenarios addressed is displayed in Table 3.

Table 3: Time for each run in all scenarios.

Scenario	Time for each run (CPU s)
A: No Uncertainty	602
B: Production Capacity	802
C: Disruption on the transportation network	750

Scenario A – Baseline - Deterministic

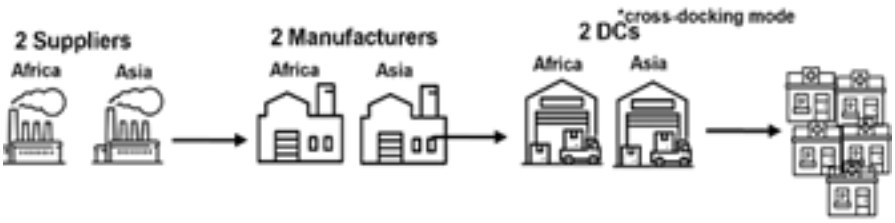

Figure 3. Baseline SCP configuration

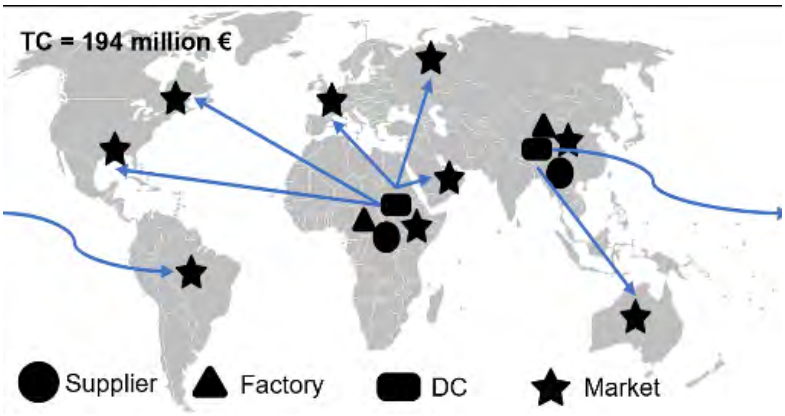

Figure 4. – Baseline SCP Locations

The PSC characterization of scenario **A** in terms of network configuration for the minimum cost solution is shown in Figure 3, corresponding to a total cost of 190 million €.

Regarding the network design most of the solutions shared the same overall structure, only differing in a small secondary part of the chain, such in some transfers of small quantities of vaccines between entities. Figure 3 illustrates the common topology for Baseline Scenario's solutions. The PSC is composed by two suppliers, two factories and two DCs functioning in a cross-docking mode, located both in Asia and Africa. Africa is the biggest supplier, producer, and distributor of vaccines, shipping them to its nearest markets, just like Asia, as shown in Figure 4.

These locations are among the ones with the lowest entity and transport related costs, not to mention that Asia is the region with the lowest overall costs. Their central location regarding other markets, and low costs justify the decision for the installation of entities.

Uncertainty behavior and impact on the PSC total cost and configuration

The results related with the cost, number, location and topology were explored and the most common topologies were characterized.

Scenario B

As can be seen in Figure 5, Scenario **B** requires an additional factory installed in Europe to satisfy the demand, due to capacity reduction of production in the factories used originally in Scenario **A**. As a result, the European supplier is also included in the PSC. The location for the factory comes from the fact that Europe is the third cheapest region to set up and operate a factory after Asia and Africa. Moreover, the PSC network for this scenario is composed by 6 DCs, resulting in a PSC with more entities. More DCs were necessary apart from the ones in Asia and Africa, for all solutions. The additional DCs were normally installed in Latin America, Middle East, Europe, or/and Eurasia. These results are justified by cheapest labour and seaport usage related costs of these 6 regions. Furthermore, these regions are among those with markets with higher demands. With the additional factory, and DCs the total cost of the PSC increased significantly to 220 million €.

Scenario C

The comparison with the baseline topology (Figure 3) shows that all solutions from Scenario C had similar topologies to the Baseline. The main difference is the instalation of two additional DCs. Given that Asia and Africa are the dominant regions for the installation of factories and DCs, it is logical that a blockage in one of the unimodal paths of these regions causes a change in the functioning of the chain since part of the vaccines need to be distributed by a different DC. Nevertheless, the solutions from this scenario had the least impact in terms of

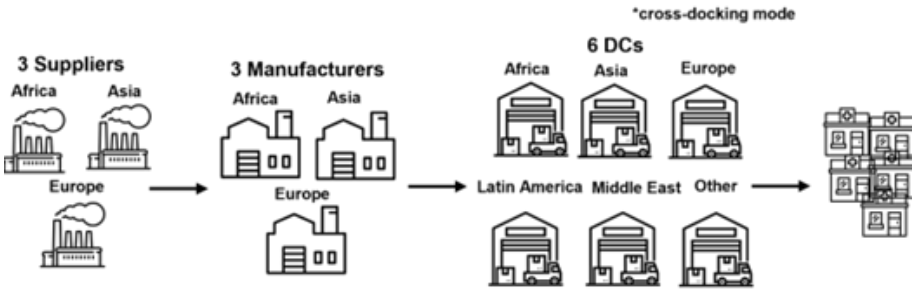

Figure 5. - Scenario B PSC configuration

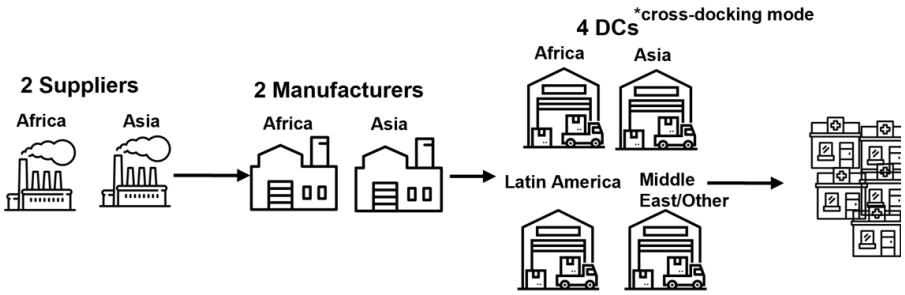

Figure 6. – Scenario **C** PSC configuration

costs. The least cost solution was comparable to the Baseline Scenario, implying that the instalation of additional DC higher costs is compensated by the supression of longer transportation and consequently lower transportation costs.

Some additional considerations

The results demonstrate that, as the severity of the uncertainty increases, costs also increase. However, this was not always the case, for Scenario **C**, indicating that the location where uncertainty is experienced may play an important role in the impact on the chain.

Given that Asia and Africa are the dominant regions for the installation of factories and DCs, it is logical that a decrease in the quantity of vaccines manufactured, or a blockage in one of the unimodal paths of these regions causes a major change in the functioning of the chain, and consequently higher costs, since part of the vaccines need to be produced in a new factory or even distributed by a different DC.

The effect of uncertainty on the PSC topology lead to more decentralized networks, requiring more entities, mainly DCs. This way, they get closer to the markets, with the DCs being able to satisfy their own local markets, instead of having two larger central DCs in Asia and Africa, as in the Baseline Scenario **A**, delivering vaccines to all markets. When comparing all the topologies, it is shown that a part of scenario **A** PSC is able to be robust and withstand both the production capacities uncertainty and the disruption. Asia and Africa continued to be the main regions for the supply and production of vaccines. Even with uncertainty, new entities emerge as common across the analysis, such as the use of a DC in Latin America and/or Middle East. The use of these entities brought additional costs, but a higher agility to the network, in case of uncertainty. Hence, it constitutes a trade-off that should be carefully studied.

CONCLUSIONS

Throughout this work, a type of uncertainty and disruption with consequences on a global PSC were explored in terms of their impact in the total network costs and its topology, in order to provide information for decision makers to create resilient and informed decisions when faced with uncertain future situations across the pharmaceutical sector.

To achieve this goal, a meta-heuristic (SA) tool able to consider different types of uncertainty and disruption was developed. To compare the severity of the uncertainty and the disruption, the algorithm explored a deterministic scenario (baseline), followed by two additional uncertain scenarios. A case study was provided for that purpose.

The results showed that the SA algorithm was able to address the inclusion of different disturbances on a PSC. It was concluded that uncertainty on the production capacity of manufacturers imposed the biggest impact on network costs, simultaneously with the larger changes on the PSC topology. The entities through which the majority of products were produced and distributed were able to withstand all the disturbances studied, being

considered the most robust. Furthermore, it was also noticeable a tendency for SC decentralization with increasing severity of uncertainty, as more local entities closer to markets tended to be installed.

As future works, deeper development of the algorithm should be made, by testing new rules for the search of the initial and neighbor solutions, in order to provide more efficient and effective solutions. Moreover, more types of uncertainties could be explored. Additional measures to quantify robustness and resilience, flexibility on the supply chain may be also explored in the future. Finally, the proposed problem could aim to incorporate sustainability and social issues, as their association with the economic part of a company is becoming inseparable.

ACKNOWLEDGEMENTS

The authors acknowledge the support provided by FCT under the project PTDC/EME-SIS/6019/2020 and UIDB/00097/2020, UIDP/00097/2020 (CEGIST).

This work is funded by national funds through the FCT – Fundação para a Ciência e a Tecnologia, I.P., under the scope of the projects UID/00297/2025 (https://doi.org/10.54499/UID/00297/2025) and UID/PRR/00297/2025 (https://doi.org/10.54499/UID/PRR/00297/2025) (Center for Mathematics and Applications – NOVA Math)

AUTHOR IDENTIFIERS

Author ORCIDs:
Chibeles-Martins N: 0000-0003-0411-6003
Pinto-Varela T: 0000-0002-4605-7083

REFERENCES

1. Marques CM, Moniz S, de Sousa JP, Barbosa-Povoa AP, Reklaitis G. Decision-support challenges in the chemical-pharmaceutical industry: findings and future research directions. Computers & Chemical Engineering 134:106672 (2020). https://doi.org/10.1016/j.compchemeng.2019.106672
2. Shukar S, Zahoor F, Hayat K, Saeed A, Gillani AH, Omer S, Hu S, Babar ZUD, Fang Y, Yang C. Drug shortage: causes, impact, and mitigation strategies. Front. Pharmacol. 12: (2021). https://doi.org/10.3389/fphar.2021.693426
3. Ayati N, Saiyarsarai P, Nikfar S. Short and long term impacts of COVID-19 on the pharmaceutical sector. DARU J Pharm Sci 28:799-805 (2020). https://doi.org/10.1007/s40199-020-00358-5
4. Franco C, Alfonso-Lizarazo E. A structured review of quantitative models of the pharmaceutical supply chain. Complexity 2017:1-13 (2017). https://doi.org/10.1155/2017/5297406
5. Zahiri B, Jula P, Tavakkoli-Moghaddam R. Design of a pharmaceutical supply chain network under uncertainty considering perishability and substitutability of products. Information Sciences 423:257-283 (2018). https://doi.org/10.1016/j.ins.2017.09.046
6. Goodarzian F, Kumar V, Ghasemi P. A set of efficient heuristics and meta-heuristics to solve a multi-objective pharmaceutical supply chain network. Computers & Industrial Engineering 158:107389 (2021). https://doi.org/10.1016/j.cie.2021.107389
7. Goodarzian F, Hoseini-Nasab H, Toloo M, Fakhrzad MB. Designing a new medicine supply chain network considering production technology policy using two novel heuristic algorithms. RAIRO-Oper. Res. 55:1015-1042 (2021). https://doi.org/10.1051/ro/2021031
8. Safaei F, Zarrinpoor N. Design of a pharmaceutical supply chain in uncertain conditions considering financial strategies and environmental concerns. Oper Manag Res : (2024). https://doi.org/10.1007/s12063-024-00440-0
9. Javadi Gargari F, Sayad M, Posht Mashhadi SA, Sadrnia A, Nedjati A, Yousefi Golafshani T. Five-echelon multiobjective health services supply chain modeling under disruption. Mathematical Problems in Engineering 2021:1-16 (2021). https://doi.org/10.1155/2021/5587392
10. Roshan M, Tavakkoli-Moghaddam R, Rahimi Y. A two-stage approach to agile pharmaceutical supply chain management with product substitutability in crises. Computers & Chemical Engineering 127:200-217 (2019). https://doi.org/10.1016/j.compchemeng.2019.05.014
11. Kirkpatrick S, Gelatt CD Jr, Vecchi MP. Optimization by simulated annealing. Science 220:671-680 (1983). https://doi.org/10.1126/science.220.4598.671

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Assessing the Impact of Thermodynamic and other Modelling Choices in MEA-based CO_2 Capture Simulations

Hassan Khaled Hassan Baabbad[a], Alberto Fernández[a], Fèlix Llovell[a], and Carlos Pozo[ab]*

[a] Universitat Rovira i Virgili, Departament d'Enginyeria Química, Tarragona, Spain
[b] ComSCIAM, Universitat Rovira i Virgili, 43007 Tarragona, Spain
* Corresponding Author: carlos.pozo@urv.cat.

ABSTRACT

Process modelling is essential to improve carbon capture unit design. However, depending on the modelling decisions made, such as the thermodynamic model and calculation method, the results obtained may vary significantly, hindering reliable process design. Nevertheless, studies that decouple the effects of thermodynamic packages and model approximations on simulation results are not well established. This contribution focuses on clarifying these effects and providing guidelines for the simulation of a CO_2 capture process with monoethanolamine (MEA) at different liquid-to-gas (L/G) ratios and CO_2 partial pressures in Aspen Plus. The calculations are validated with eight pilot campaign runs. The analysis reveals that the use of rate-based with kinetic reactions significantly improves the accuracy of the simulations. This approach, combined with the ENRTL-RK thermodynamic model, using Peng-Robinson for the vapor phase, provides the best performance, with average deviations below 3% in terms of CO_2 capture efficiency and energy requirements. On the other hand, strategies to improve convergence success in closed-loop simulations include replacing a single heat exchanger with two heater blocks, and a calculator block to automate mass-balance calculations and account for losses.

Keywords: Climate Change, Carbon Capture Utilization Storage (CCUS), Post Combustion Carbon Capture (PCC), Monoethanolamine (MEA), Aspen Plus Simulation

INTRODUCTION

The recent thirtieth conference of the parties (COP30) in Belém, Brazil, acknowledged the likelihood of overshooting the 1.5°C set during the Paris Agreement [1]. To limit global warming to 1.5°C requires rapid and deep emissions reductions of greenhouse gas (GHG) emissions such as carbon dioxide (CO_2) [2]. One of the key technologies that can curb these emissions, also reported in the Industrial Carbon Management (ICM) by the EU, is Carbon Capture, Utilization and Storage (CCUS) [3]. As the deployment of CCUS accelerates globally, rigorous process modelling becomes central to designing efficient and reliable CO_2 capture systems. Among different CO_2 capture technologies, the most mature and adaptable system at both pilot and industrial scale is post-combustion capture with monoethanolamine (MEA) [4].

Despite extensive literature on MEA capture systems, the bibliography on this topic remains fragmented and many results differ from one another. The main reason for these discrepancies is related to the choice of thermodynamic models, calculation methods (model complexity) and closed-loop convergence challenges. Most published works omit these details that matter in practice.

To address this gap, we investigate the independent effects and interactions between two different thermodynamic packages and calculation methods and document convergence tactics that enable stable solutions. The novelty of this study lies in the decoupling of thermodynamic models from calculation methods, along with the development of modelling strategies that facilitate convergence.

Specifically, the study compares the new unsymmetric electrolyte NRTL model (ENRTL-RK) for liquid with the Peng–Robinson Equation of State (EOS) for vapor phase, with ENRTL-RK combined with the Perturbed-Chain Statistical Association Fluid Theory (PC-SAFT). Furthermore, different degrees of model complexity,

such as equilibrium-stage, rate-based modelling, and chemical reactions are also revised to check their impact on Key Performance Indicators (KPIs). These KPIs include CO_2 capture efficiency (%) and thermal energy requirement (MJ/kg CO_2). Flowsheet convergence tips are also described to help in achieving closed-loop simulation convergence success.

The goal is to provide practical guidelines to follow when modelling post-combustion CO_2 capture processes using chemical absorption, especially at varying liquid-to-gas (L/G) ratios and CO_2 partial pressures.

PILOT PLANT SPECIFICATIONS

In this work, the pilot plant at the University of Kaiserslautern reported in [5] has been used for the validation of the process model developed in this work. The absorber and stripper column internals and the range of operating conditions are reported in Table 1. One of the parameters analyzed in this pilot plant was the variation in the L/G ratio while maintaining a constant gas flow rate. Two series of experiments were used for comparison: one with a low CO_2 partial pressure of 54 mbar (cases A1–A6) and another with a higher CO_2 partial pressure of 102 mbar (cases A7–12). In this work, eight of these experiments, covering both low and high CO_2 partial pressures and L/G ratios, were validated.

Key performance indicators

To validate the model using the experimental data from the pilot plant, two KPIs were defined, namely the CO_2 capture efficiency (η_{cap}) and thermal energy requirement (e_{req}). The following equations (1 – 3), based on the work of [6] were used to calculate these KPIs.

$$\dot{m}_{CO2,abs} = \dot{m}_{CO2,in} - \dot{m}_{CO2,out} \quad (1)$$

$$\eta_{cap} = \frac{\dot{m}_{CO2,abs}}{\dot{m}_{CO2,in}} \bullet 100 \quad (2)$$

$$e_{req} = \frac{\dot{Q}_{reb}}{\dot{m}_{CO2,abs}} \quad (3)$$

Here, $\dot{m}_{CO2,abs}$ is the absorbed CO_2 mass flow, $\dot{m}_{CO2,in}$ and $\dot{m}_{CO2,out}$ are the mass flow rates of CO_2 in the flue gas entering and exiting the absorber in kg/h, respectively, and $\dot{Q}_{reb}$ is the reboiler duty in MJ/h.

Model validation

The deviations between the experimental data and simulations results are given by the relative error (ε) and the percentage average absolute relative deviation ($\%AARD$). Equations (4 – 5) are used to calculate these deviations.

$$\varepsilon = \frac{x_{sim} - x_{exp}}{x_{exp}} \bullet 100 \quad (4)$$

$$\%AARD = \frac{\sum_1^n |\varepsilon|}{n} \quad (5)$$

where, x_{sim} is the simulation result, x_{exp} is the experimental value, and n is the number of data points.

METHODOLOGY

Process flowsheet

Aspen Plus v14 was used to build the flowsheet and simulate the MEA-based CO_2 capture process as depicted in Figure 1. The flue gas enters the bottom of the absorber while the lean solvent is introduced at the top. After absorption, the solvent is loaded with CO_2 and becomes the rich solvent. The rich solvent stream exits from the bottom of the absorber and is pumped through the heater block (half-heat exchanger), to finally enter the stripper, where the process is reversed. Following the stripper, the lean amine stream exits at the bottom and enters the second heater block (which acts as a cooler). This lean amine stream is mixed with the MEA and water makeups which are necessary to account for losses.

In the simulation, it is possible to use a single rigorous heat exchanger (HeatX) model to pre-heat the rich solvent and cool the lean solvent stream. However, for a closed-loop system, this approach creates two nested convergence loops: one for the HeatX block and the second one for the solvent recycle loop. This increases model complexity, making it prone to convergence issues, mostly related to temperature crossovers. To avoid this, two heater blocks connected using a heat stream are used instead of the single HeatX model [7]. In addition, for maintaining overall mass (material) balance, a calculator block is created in order to close the recycle solvent loop. This facilitates the convergence of the recycle stream by ensuring mass balances are accurate through the calculation of the required MEA and H_2O makeup streams.

Thermodynamic method and reactions

One of the most fundamental aspects to consider during process modelling is the choice of the thermodynamic method. This is because an incorrect decision of the thermodynamic package can lead to erroneous results in the simulation.

In this study, two slightly different thermodynamic methods are compared. The first uses the Electrolyte Non-Random Two-Liquid (ENRTL-RK) property method with the Peng–Robinson (PR) EOS to compute liquid and vapor properties, respectively [8]. This model was recently used by [9] and found to provide good predictions of the experimental data. The second thermodynamic package also uses the ENRTL-RK to calculate liquid phase properties but applies the PC-SAFT EOS for the calculation of the vapor phase fugacity coefficients [10].

For both property methods, the pure and binary parameters have been regressed or validated against available experimental data [8, 10]. Moreover, the following

chemical equilibrium reactions (equations (6 – 10)) were considered.

$$2H_2O \leftrightarrow OH^- + H_3O^+ \quad (6)$$

$$H_2O + HCO_3^- \leftrightarrow CO_3^{2-} + H_3O^+ \quad (7)$$

$$2H_2O + CO_2 \leftrightarrow HCO_3^- + H_3O^+ \quad (8)$$

$$H_2O + MEACOO^- \leftrightarrow MEA + HCO_3^- \quad (9)$$

$$H_2O + MEAH^+ \leftrightarrow MEA + H_3O^+ \quad (10)$$

Equilibrium constants of reactions (6)–(10) for both thermodynamic packages were calculated using the Gibbs free energy change described by equation (11).

$$\ln(K_j) = -\frac{\Delta G^o_{i,j}}{RT} \quad (11)$$

Here, K_j denotes the chemical equilibrium constant, $\Delta G^o_{i,j}$ is the standard free energy of formation of component i, R is the universal gas constant, and T is the temperature in Kelvin.

On the other hand, carbamate and hydrolysis reactions are typically expressed as kinetically controlled reactions, given by equations (12 – 13). Kinetic constants for these reactions can be found in [10].

$$MEA + CO_2 + H_2O \rightarrow MEACOO^- + H_3O^+ \quad (12)$$

$$MEACOO^- + H_3O^+ \rightarrow MEA + CO_2 + H_2O \quad (13)$$

Equilibrium-stage modelling

The Aspen Plus RadFrac block allows two calculation methods: equilibrium and rate-based modelling. Furthermore, in both methods it is possible to consider equilibrium or kinetic reactions in the liquid phase.

In this work, equilibrium-stage modelling was the first approach used, setting the number of stages to 20 for both columns. This approach assumes perfect mixing and does not take into account mass-transfer. Nonetheless, it is recommended as a first step because it facilitates the convergence of the closed-loop simulation, producing initial column estimates (such as temperature and composition) that are useful when moving to the more advanced rate-based modelling. Additionally, for the same equilibrium model, different model complexities regarding reaction information can still be applied, as reported in [11].

The first model used in this paper was the equilibrium-stage with chemical equilibrium reactions (EQ). Then, after successful convergence of the closed-loop simulation, the complexity can be increased to equilibrium-stage using kinetic reactions (EK). Nonetheless, when switching to EK model, data for liquid holdup or liquid residence time is required. In this contribution, the liquid holdup values reported by [12] were used. After successful convergence of the EK model, the columns can then be switched to rate-based modelling to account for mass transfer phenomena.

Rate-based modelling

In reality, the previous equilibrium-stage assumptions are never satisfied because of the highly non-ideal behavior of CO_2 capture systems. To address these limitations, the rate-based model is also tested. The rate-based model accounts for both mass and heat transfer. The mass transfer is explained by the two-film theory (Figure 2), which uses the Maxwell-Stefan approach [13].

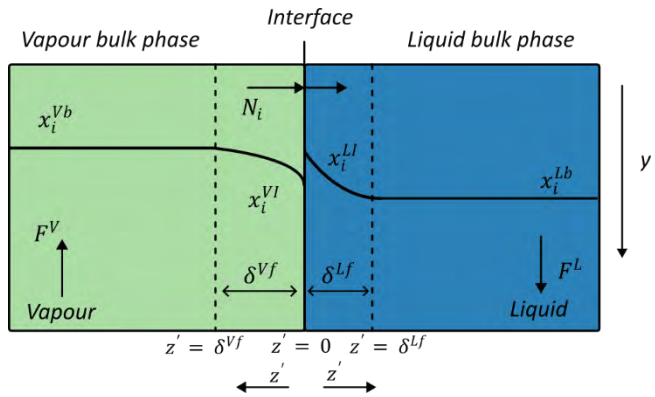

Figure 2. Two-film theory (model) for a packing segment adapted from [13].

Similar to equilibrium modelling, rate-based models also accept two approaches to define chemical reactions. The first one is rate-based model with chemical equilibrium reactions (RE), while the second one uses also the rate-based model but with kinetic (RK) and film reactions. This last model is considered the most complex.

In rate-based modelling, the column geometry must be specified for both absorber and stripper. To this end, data presented in Table 1 is utilized to replicate the pilot plant that will be used for validation. On the other hand, Table 2 reports the model parameters used for the rate-based model, which are common for the two thermodynamic models used in this work. For the liquid, vapor mass transfer, and interfacial area factor, the correlation by [14] is used. The heat transfer has been modeled using [15] and the liquid holdup was calculated by [16]. Additionally, the liquid film was discretized into 5 points and the reaction condition factor was set to 0.9 as suggested by [17].

Combining equilibrium-stage and rate-based modelling, along with equilibrium and rate-limiting kinetic reactions, and two thermodynamic packages, a total of eight different calculation models were tested.

Table 1: The main pilot plant specifications [5].

Parameter	Data	Unit
Flue gas source	Natural gas burner	(-)
CO_2 concentration in the flue gas	3–14	(vol%)
Liquid flow rate	20–350	(kg/h)
Flue gas flow rate	30–100	(kg/h)
Liquid load in the absorber	2–28.5	($m^3/(m^2h)$)
Electrical heat duty of evaporator	0–30	(kW)
Absorber pressure	Atmospheric pressure	(bar)
Stripper pressure	1.0–2.5	(bar)
	Absorber Geometry	
Total height of packing in absorber	4.25	(m)
Column internal diameter	0.125	(m)
Type of packing	BX 500	(-)
	Stripper Geometry	
Total height of packing in stripper	2.55	(m)
Column internal diameter	0.125	(m)
Type of packing	BX 500	(-)

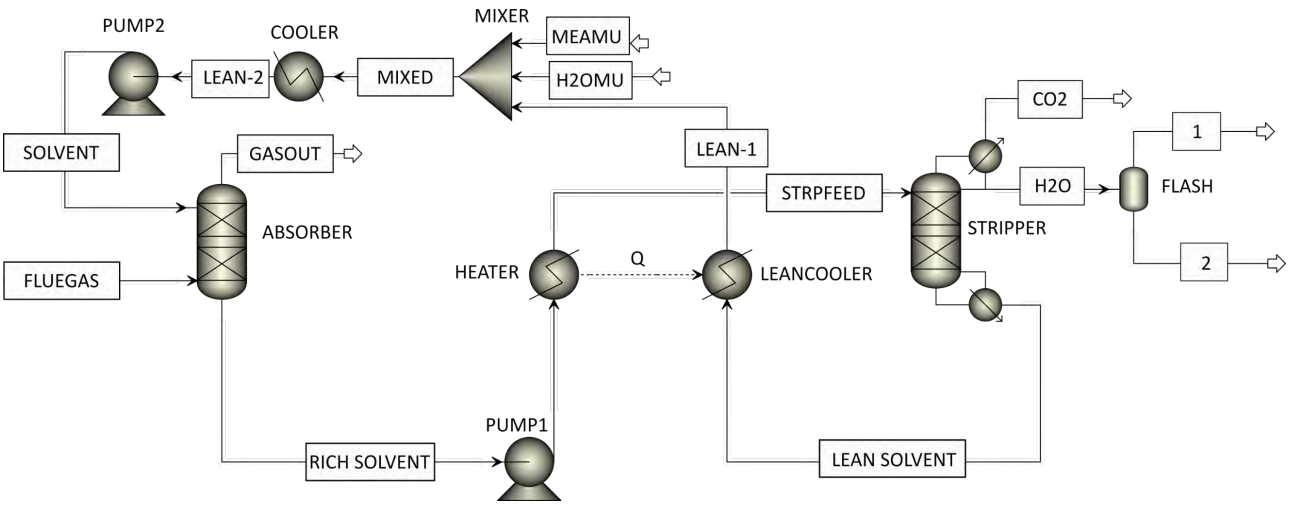

Figure 1. Process flowsheet of MEA CO_2 capture process.

Table 2: Rate-based model parameters used.

Parameter	Specification
Number of segments	20
Flow model	Mixed
Mass transfer coefficient method	Brf-85
Heat transfer coefficient method	Chilton and Colburn
Interfacial area method	Brf-85
Interfacial area factor	1
Liquid holdup correlation	Brf-92
Reaction condition factor	0.9
Film discretization ratio	5

RESULTS AND DISCUSSION

Effect of thermodynamic and model complexity

The process modelling results of the eight runs carried out at the University of Kaiserslautern are evaluated. The CO_2 capture efficiency (%) and the thermal energy requirement (MJ/kg CO_2) are the two KPIs evaluated and compared for the four different model complexities.

A summary of the %AARD in all KPIs for these models is given in Table 3. Both thermodynamic models predict the KPIs with an AARD of less than 10%. Among the KPIs, the thermal energy requirement is the one that can

be most accurately predicted. Nonetheless, some differences arise regarding the use of different model complexities. Including kinetically-controlled reactions significantly impact the results when using rate-based models. As an example, for the ENRTL-RK with Peng–Robinson, the model predicts the CO_2 capture efficiency and thermal energy with an AARD of 9.23% and 7.96%, respectively, when using rate-based model with equilibrium reactions (RE), but improves to 2.46% and 2.84%, respectively, when including kinetic reactions (RK). When using ENRTL-RK with PC-SAFT, the same pattern is noticed, but with minor improvements, especially for thermal energy requirement.

Lastly, the incorporation of mass-transfer kinetics also has a significant influence on the results, particularly when using ENRTL-RK with Peng–Robinson. Switching from equilibrium to rate-based model, when already considering kinetically controlled reactions (i.e., from EK to RK), improves predictions of CO_2 capture efficiency from 8.31% to 2.46%. Likewise, using PC-SAFT for the same KPI, the deviations decrease from 7.97% to 5.46%. For the thermal energy requirement, predictions improve by 61% (i.e., from 7.23% in EK to 2.84% in RK) when using Peng–Robinson and by 8% (from 7.05% to 6.47%) if using PC-SAFT.

These results highlight the fact that when thermodynamic binary parameters from both models have been properly validated against experimental data (heat capacity, density), they play a minor role compared with model complexity. In this case, where the same property method was used to calculate liquid properties (ENRTL-RK), the choice of different EOS for the vapor phase did not alter the results significantly, achieving minor differences in %AARD between Peng–Robinson and PC-SAFT for the same model complexity. However, for the same thermodynamic model, %AARD changed significantly when moving from classic equilibrium-stage model to the more advanced rate-based model with kinetic reactions.

Effect of L/G ratio and CO_2 partial pressure

After checking the average patterns from the two KPIs, this subsection examines these KPIs for each run analyzing the impact of the L/G ratio and CO_2 partial pressure. Figure 3 depicts the relative errors of the calculated CO_2 capture efficiency at low (A1, A4, A5, and A6) and high (A9, A10, A11, and A12) CO_2 partial pressure across various L/G ratios using the two thermodynamic methods.

It is apparent from Figure 3 that both thermodynamic methods perform well and are comparable to each other, although slight differences arise. At both, low and high CO_2 partial pressures and L/G ratios, the two thermodynamic models overestimate experimental results by 5–11% when using chemical equilibrium reactions (EQ and RE). However, when using kinetically-controlled reactions (RK, crimson-red) PC-SAFT underestimates experimental results by 1–11%, whereas, with Peng–Robinson, simulation results either underestimate or overestimate experimental values in the range of (-5%)–2%, showing excellent agreement with the data at high CO_2 partial pressures. Switching the mass-transfer model to EK (light blue), results in overestimation of the calculations by 0–11% for both thermodynamic packages.

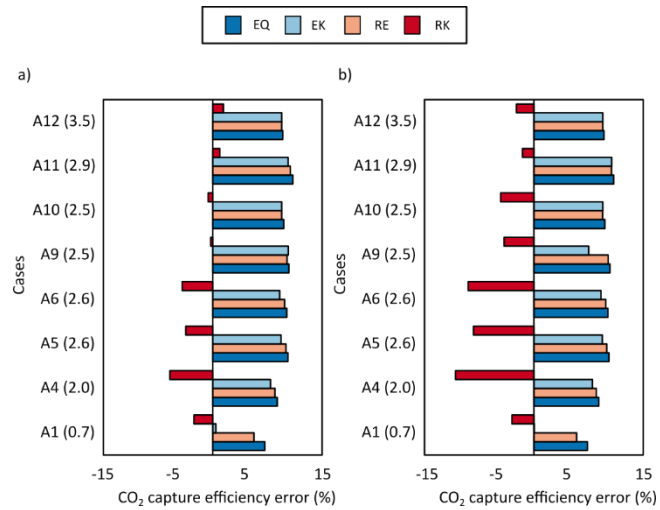

Figure 3. Relative error (%) for CO_2 capture efficiency using: (a) ENRTL-RK with Peng–Robinson and (b) ENRTL-RK with PC-SAFT. L/G ratio value in brackets.

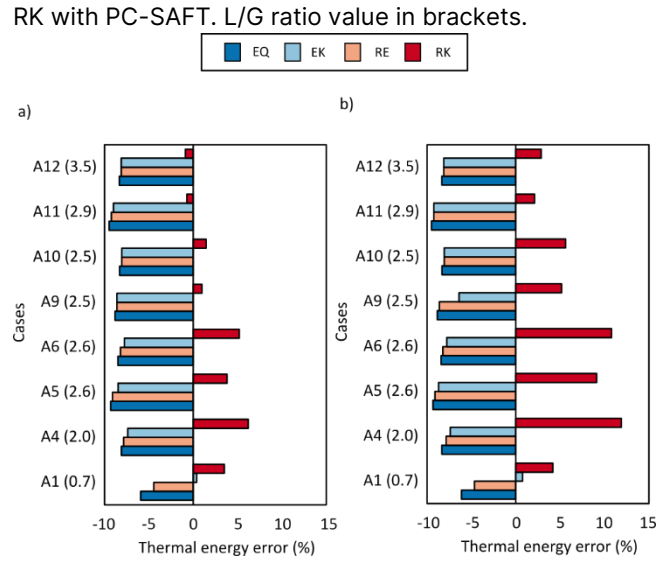

Figure 4. Relative error (%) for thermal energy requirement using: (a) ENRTL-RK with Peng–Robinson and (b) ENRTL-RK with PC-SAFT. L/G ratio value in brackets.

Figure 4 illustrates the relative errors of the calculated thermal energy requirement for both low and high CO_2 partial pressures and L/G ratios. As shown in the figure, the differences between the results using the two thermodynamic models are comparatively small, with the

Table 3: % AARD between the pilot plant data based on 8 runs and the process modelling results using different model complexities and two thermodynamic methods.

Thermodynamic method	Model complexity	CO_2 capture efficiency (%)	Thermal energy requirement ($MJ/kgCO_2$)
ENRTL-RK Peng–Robinson	EQ	9.66	8.35
	EK	8.31	7.23
	RE	9.23	7.96
	RK	2.46	2.84
ENRTL-RK PC-SAFT	EQ	9.68	8.39
	EK	7.97	7.05
	RE	9.25	7.98
	RK	5.46	6.47

exception of the RK model.

Interestingly, when using chemical equilibrium reactions (EQ and RE) the calculated values consistently underestimate the thermal energy requirement in the range of (-9%)–(-4%), irrespective of the mass-transfer method. Conversely, using kinetically-controlled reactions (RK, crimson-red) with PC-SAFT clearly over-predicts the experimental results in the range of 2–12%.

This also happens when using Peng–Robinson with kinetic reactions (crimson-red), but in these cases the thermal energy has more modest overestimations, with overall deviations between (-1%)–6%.

Additionally, it can be seen that the results using PC-SAFT for RK (crimson-red) overestimate the experimental values compared to Peng–Robinson at both low and high CO_2 partial pressure values (Figure 4). This result was also evident in Figure 3 when predicting CO_2 capture efficiency, yet, in that case, the deviations were causing an underprediction. The analysis shows that the differences between the two thermodynamic models is comparatively small at low CO_2 partial pressure and L/G ratio between 0.7 and 2.6.

However, at high CO_2 partial pressure and L/G >2.5 the choice of the EOS for the vapor phase (Peng–Robinson) affects the results, particularly in rate-based model with kinetic reactions.

CONCLUSIONS

The purpose of the present contribution is to provide a framework for accurately simulating a post-combustion CO_2 capture process using MEA by evaluating the effects of thermodynamic models and mass transfer methods, under different modelling decisions and chemical reactions assumptions.

During process modelling, two strategies were identified to facilitate convergence of closed-loop simulations. The first strategy consisted of replacing a rigorous heat exchanger with two heater blocks connected by a heat stream. The second strategy consisted of using a calculator block to automate the computation of solvent and co-solvent (water) makeup, thereby maintaining material balance in the system automatically.

The two thermodynamic models used were based on the ENRTL-RK method for liquid phase properties, but differing on how vapor phase properties are calculated (Peng–Robinson vs PC-SAFT).

The results indicated that both thermodynamic methods yielded an %AARD below 10% in all cases, showing comparable results in most mass transfer approaches. However, the use of the rigorous RK model revealed that ENRTL-RK with Peng-Robinson was superior to ENRTL-RK using PC-SAFT, reducing average deviation by approximately 50% for both KPIs.

For the computational cost, both equilibrium-stage and rate-based simulations exhibit acceptable performance for isolated runs, in the order of 1.0-2.0s (EQ; the simplest) and 11.0-22s (RK, the most complex) using a 13th Gen Intel(R) Core i7-1355U and 16.0 GB RAM machine. While this difference might seem small in practice, it can become relevant when the process model is embedded within an optimization framework potentially requiring the evaluation of thousands of solutions. Furthermore, the rate-based model provides greater accuracy, yet it requires additional data such as column internals (tray or packing specifications) and reliable mass transfer correlations, which might not always be available in practice. Hence, modellers may select simplified approaches in contexts where deviations on the order of 10% are deemed acceptable.

The current process model was validated using experimental data from a single pilot-plant operating under post-combustion conditions representative of gas-fired and coal-fired power plants. Although the validation considers a reasonably broad range of L/G ratios (0.7-3.5) and inlet CO_2 compositions (CO_2 partial pressures, 54-102 mbar), the applicability to other pilot plants or to op-

erating conditions outside this range cannot be guaranteed. Therefore, this should be regarded as a limitation of the present study.

Nevertheless, the results obtained here provide useful insight into the model performance within the tested domain. In particular, for the cases analysed, the Peng-Robinson equation of state for the gas phase showed better performance at L/G ratios above 2.5, suggesting that it may be a more suitable option for plants operating under similar conditions. The combination of this model within the rate-based with kinetic equations framework assures very good accuracy for the KPIs studied, particularly at higher CO_2 partial pressures. Still, extrapolation to other systems should be carried out with caution. For reliable application to different pilot plants or operating windows, the thermodynamic model should be assessed and, if necessary, fit using experimental data representative of the corresponding conditions. Even in that case, real-world systems may deviate from the assumptions made in the model. This is because the predictive accuracy depends on the quality and availability of data for parameters such as mass transfer coefficients and kinetic constants, which might not always be available in practice.

In conclusion, the results presented in this paper may be used as practical guidelines for the accurate design of post-combustion CO_2 capture systems using rate-based models with kinetic reactions, adaptable to various L/G ratios and CO_2 partial pressures.

ACKNOWLEDGEMENTS

This work is supported by grants PID2021-127713OA-I00 and PID2023-149713OB-I00, funded by MICIU/AEI/10.13039/501100011033 and "FEDER, UE", and grant CNS2023-144890 funded by MICIU/AEI/10.13039/501100011033 and the European Union NextGenerationEU/PRTR. Additional support from the Chair on Energy Transition URV-Fundación Repsol and from the Catalan Government through AGAUR as a Consolidated Research Group (SGR 2021-00738) are gratefully appreciated. H.K. Hassan Baabbad acknowledges grant 2023PMF-PIPF-06 from the Martí Franquès Research Fellowship Program.

AUTHOR IDENTIFIERS

Author ORCIDs:
Hassan Baabbad HK: 0009-0008-8974-9312
Fernández A: 0000-0002-1241-1646
Llovell F: 0000-0001-7109-6810
Pozo C: 0000-0002-5645-674X

REFERENCES

1. Al Amer N. COP30 Belém: Implications for carbon management implementation. *Global CCS Institute* (2025) https://www.globalccsinstitute.com/publications/cop30-belem-implications-for-carbon-management-implementation/
2. United Nations Environment Programme. Emissions Gap Report 2025: Off target continued collective inaction puts global temperature at risk. *UNEP* (2025) https://wedocs.unep.org/items/9f0bf855-2069-42a6-a856-4b389f740c5c
3. European Commission. *Towards an ambitious Industrial carbon management for the EU* (2024) https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52024DC0062
4. Cremona R, De Lena E, Conversano A, Spinelli M, Romano MC, Gatti M. Techno-economic assessment of high temperature heat pumps integrated in mea-based post-combustion CO2 capture for cement plant. Carbon Capture Science & Technology 16:100446 (2025). https://doi.org/10.1016/j.ccst.2025.100446
5. Mangalapally HP, Hasse H. Pilot plant study of post-combustion carbon dioxide capture by reactive absorption: methodology, comparison of different structured packings, and comprehensive results for monoethanolamine. Chemical Engineering Research and Design 89:1216-1228 (2011). https://doi.org/10.1016/j.cherd.2011.01.013
6. Notz R, Mangalapally HP, Hasse H. Post combustion CO2 capture by reactive absorption: pilot plant description and results of systematic studies with MEA. International Journal of Greenhouse Gas Control 6:84-112 (2012). https://doi.org/10.1016/j.ijggc.2011.11.004
7. Shaeiwitz J. Book review: learn aspen plus in 24 hours, 2nd edition by thomas A. adams, II. Chem. Eng. Ed. 57: (2023). https://doi.org/10.18260/2-1-370.660-132561
8. Aspen Technology, Inc. CO2 capture from coal power plant using MEA for Aspen Plus. Aspen Plus Documentation. Aspen Technology, Inc., Bedford, MA (2022).
9. Singh Rawat S, Gomes Fonseca F, María Isabel Roldán Serrano . Innovative high temperature heat pump concepts for an economic decarbonization of a carbon capture unit. Carbon Capture Science & Technology 17:100517 (2025). https://doi.org/10.1016/j.ccst.2025.100517
10. Aspen Technology, Inc. Rate-based model of the CO2 capture process by MEA using Aspen Plus. Aspen Plus Documentation. Aspen Technology, Inc., Bedford, MA (2014).

11. Kenig EY, Schneider R, Górak A. Reactive absorption: optimal process design via optimal modelling. Chemical Engineering Science 56:343-350 (2001). https://doi.org/10.1016/s0009-2509(00)00234-7
12. Dugas RE, *Pilot Plant Study of Carbon Dioxide Capture by Aqueous Monoethanolamine*, The university of Texas at Austin, TX, USA, (2006).
13. Lawal A, Wang M, Stephenson P, Yeung H. Dynamic modelling of CO2 absorption for post combustion capture in coal-fired power plants. Fuel 88:2455-2462 (2009). https://doi.org/10.1016/j.fuel.2008.11.009
14. Bravo JL, Rocha JA, Fair JR. Mass transfer in gauze packings. *Hydrocarb. process*. 64: 91–95 (1985)
15. Taylor R, Krishna R. Multicomponent Mass Transfer, New York (1993).
16. Bravo JL, Rocha JA, Fair JR. A comprehensive model for the performance of columns containing structured packings. *I. ChemE. Symp*. 1: A489–A507 (1992)
17. Zhang Y, Chen CC. Modeling CO2 absorption and desorption by aqueous monoethanolamine solution with aspen rate-based model. Energy Procedia 37:1584-1596 (2013). https://doi.org/10.1016/j.egypro.2013.06.034

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Experimental and Kinetic Study of Iron Oxide Reduction in a Fixed Bed Reactor using a Dynamic Shrinking Core Model

Emiliano Salucci[ab]*, Antonio D'Angelo.[ab], Vincenzo Russo[b], Henrik Grénman[a], and Henrik Saxén[a]

[a] Åbo Akademi University, Faculty of Science and Engineering, Henrikinkatu 2, Turku 20500, Finland
[b] University of Naples Federico II, Department of Chemical Sciences, Via Cintia 26, Napoli 80126, Italy
* Corresponding Author: emiliano.salucci@abo.fi.

ABSTRACT

The use of green hydrogen to reduce iron ore is a promising approach to drastically decrease CO_2 emissions in the steel industry. To enable the rapid adoption of this technology, it is essential to start from the fundamentals, namely understanding the intrinsic kinetics of iron oxide reduction. In this work, a kinetic investigation was conducted in a PBR-like system using both pure and commercial iron oxide powders under a wide range of operating conditions. The thermal conductivity of the outlet gas was measured and innovatively correlated with the extent of solid reduction through a rigorous mathematical procedure. To simulate the reduction process and determine the kinetic parameters, a deterministic axial dispersion model was developed in conjunction with a dynamic multistep shrinking core model. The model incorporates the particle size distribution of the solid into the mass balance and includes a reaction-front control mechanism to ensure physical consistency during kinetic parameter estimation. The model successfully reproduced all experimental data, confirming that the controlling mechanisms were predominantly chemical, internal diffusion, or mixed, depending on the type of solid material used. Furthermore, the model not only demonstrates excellent predictive capabilities through the validation of kinetic parameters but also provides new insights into the key phenomena governing iron oxide reduction.

Keywords: hydrogen, direct reduced iron, green ironmaking, TCD, SCM, deterministic kinetic modelling

INTRODUCTION

The steel industry is one of the most carbon-intensive sectors in the world, with estimated CO_2 emissions of 1.85 - 2.6 tons per ton of steel produced [1]. This is largely due to the dominance of the blast furnace–basic oxygen furnace route, which is based on the combustion of coke to generate carbon monoxide (CO), the most important reducing agent in the high-temperature chemical reduction of iron ores [2]. Various alternative processes have been developed in recent decades, e.g., direct reduction (DR), in which CO and hydrogen (H_2) gas mixtures are used to reduce iron ores at lower temperatures [3]. DR processes promise significant reductions in emissions, especially when hydrogen is used as the sole reducing agent [4]. However, the technological transition to hydrogen-based reduction faces challenges related to the availability of hydrogen, energy requirements and reactor customization. The reduction of iron oxides follows a multi-step reversible reaction sequence from hematite (Fe_2O_3) to magnetite (Fe_3O_4), wüstite (FeO) and finally metallic iron [5]. The reaction network is displayed in Table 1.

Table 1: Reactive network.

	Reactions
1)	$3Fe_2O_3 + H_2\,(CO) \rightleftharpoons 2Fe_3O_4 + H_2O\,(CO_2)$
2)	$Fe_3O_4 + H_2\,(CO) \rightleftharpoons 3FeO + H_2O\,(CO_2)$
3)	$FeO + H_2\,(CO) \rightleftharpoons Fe + H_2O\,(CO_2)$

The role of water vapor is particularly crucial during the second and third reduction steps in controlling the local thermodynamic equilibrium. A wide range of experimental methods have been used to investigate H_2-based reduction, including thermogravimetric analysis [6], gas

phase analyses [7] and morphological characterizations [8]. The current findings in this field remain largely inconsistent, primarily because several distinct rate-determining steps have been identified. Depending on the specific reduction stages and experimental conditions, these mechanisms range from chemical kinetics and intraparticle diffusion to hybrid regimes [9]. Consequently, the diversity of theoretical and practical methodologies has led to a significant divergence in reported kinetic parameters.

This discrepancy is mainly attributed to the varied experimental setups employed, as well as the fundamental characteristics of the reducing agents and solid materials. Literature reviews show that while some researchers estimated kinetic constants using samples with high gangue content [10] others utilized high-purity materials synthesized in situ [11]. Furthermore, fluctuations in particle geometry, scale, and physical attributes—such as density and porosity—further complicate the reliable assessment of published activation energies. As a result, deriving definitive intrinsic kinetic parameters solely from existing literature remains a significant challenge. These studies show a large variability in the reported kinetic parameters with activation energies between 18 kJ/mol and 246 kJ/mol [12], with results influenced by the solid materials, the experimental procedure as well as the mathematical approach taken to determine the parameters [13]. To overcome this inconsistency, a rigorous and innovative approach is followed in the present work. The reduction degree of the iron oxide sample in the experiment was determined using an innovative mathematical method by measuring the thermal conductivity of the reactor's outlet gases, which allows for an accurate tracking of the reaction progress. The experimental data were analyzed and interpreted through a detailed mathematical model based on an axial dispersion model coupled with a multi-stage dynamic shrinking core model (dSCM); furthermore, the model accounts for the particle size distribution of the powders [14].

MATERIALS AND METHOD

Experimental Setup

The iron oxide samples used in the reduction studies were high purity hematite powders (purity >96 %) from Sigma-Aldrich and commercial samples of pellets used in BFs, and DR shaft furnaces provided by steelmaking companies. Pure hematite has an average particle diameter below 5 µm, while the commercial materials were produced by crushing and sieving pellets to achieve a fine particle size distribution (< 50 µm). The chemical composition of the commercial material is reported in Table 2.

Table 2: BF and DRI pellets chemical compositions.

	Fe [%]	FeO [%]	Al_2O_3 [%]	CaO [%]	MgO [%]	MnO [%]	SiO_2 [%]
BF	64.9	0.15	0.34	1.04	0.39	0.15	5
DRI	66.9	0.69	0.43	0.99	0.28	0.05	1.7

The experiments were carried out under isothermal conditions using an Autochem 2910 device containing a thermal conductivity detector (TCD) for real-time analysis of the gas phase composition with high temporal resolution. The setup of the experimental system is shown in Figure 1. The reactor used was a straight glass tube with a length of approximately 190 mm and a radius of 11 mm, housed in a compact electric furnace that facilitated accurate thermal control (heating ramp: 10 °C/min). The powdered sample (30-200 mg) was positioned on a quartz disc supported by a layer of quartz wool.

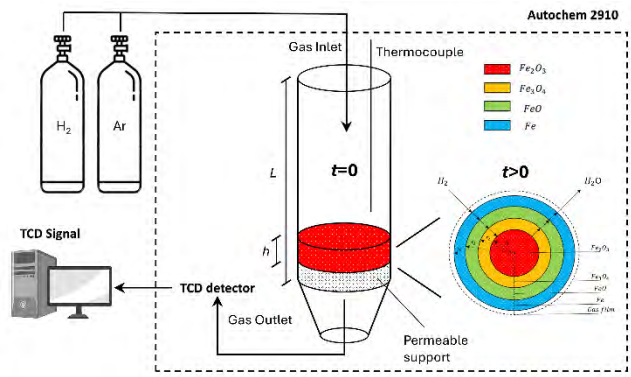

Figure 1. Schematic of the reactor experimental setup.

A wide range of operating parameters — including temperature, sample mass, total gas flow rate and hydrogen concentration — were systematically varied to investigate their effects on the reduction. The solid samples were characterized using several analytical techniques: X-ray diffraction (XRD) to identify chemical phases tracking the reduction process, light scattering to determine particle size distribution, nitrogen adsorption to assess porosity and specific surface area, and scanning electron microscopy (SEM) to evaluate surface morphology.

TCD Signal interpretation

Analyzing the thermal conductivity signal during iron oxide reduction is challenging because it reflects not only the conductivity difference between reference and exhaust gases but also the structural characteristics of the Wheatstone bridge. Despite this complexity, the high sensitivity of the data enables a detailed interpretation of the reduction process. Changes in signal slope are correlated with different reaction phases and indicate shifts in kinetic or thermodynamic regimes caused by the depletion of dominant solid species. An example of TCD results of pure hematite reduction is displayed in Figure 2.

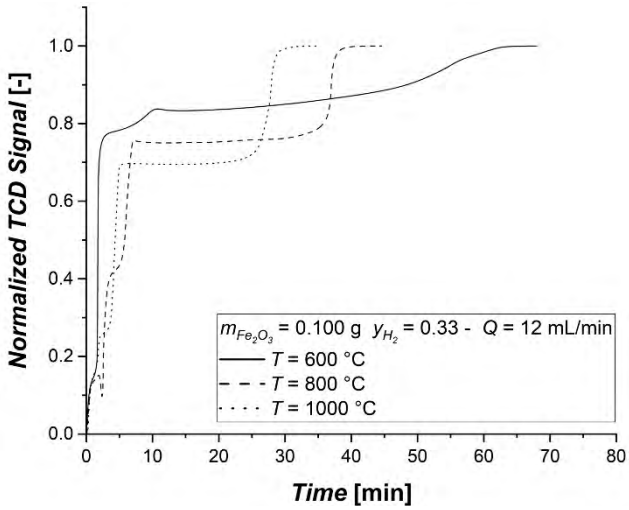

Figure 2. Examples of normalized TCD signals of pure hematite reduction at different operating temperatures.

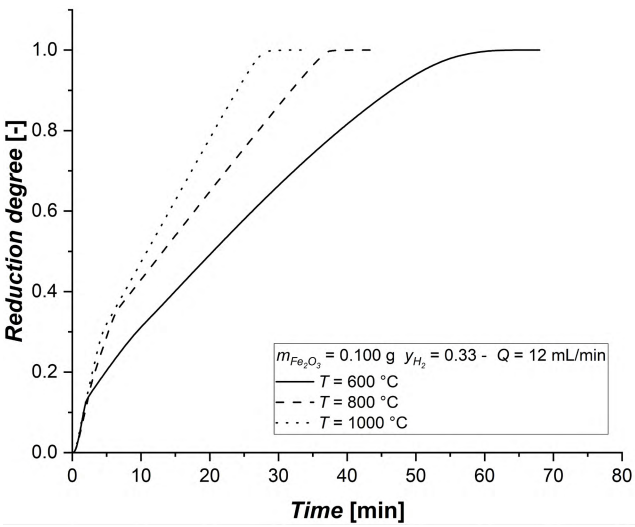

Figure 3. Examples of reduction degrees obtained from TCD signals of pure hematite reduction.

To validate these interpretations, stop-tests were performed at selected reaction times, and XRD analysis confirmed a direct correlation between the signal and the solid conversion. However, simulating the signal itself is complex because it is strongly affected by hydrogen concentration and reactor fluid dynamics. To overcome this, an alternative approach was developed to correlate the gas-derived signal with an indirect estimate of the overall reduction degree. The method involves deterministic mathematical operations: normalization, complement transformation, cumulative integration using the trapezoidal method, and correction by subtracting a blank signal obtained under identical conditions. This correction eliminates overestimation caused by axial dispersion and dilution effects, isolating the actual hydrogen consumption by the chemical reaction.

The processed signal can then be compared with the global reduction degree of the solid sample and used to estimate kinetic parameters for multistage iron oxide reduction. The key innovation lies in correlating highly accurate gas-phase measurements with solid conversion in heterogeneous reactions, enabling robust kinetic modeling and deeper physic-chemical interpretation of the reduction process. An example of reduction degrees obtained from TCD signals is displayed in Figure 3 from the same experiments used in Figure 2.

First principle mathematical model

To describe the experimental data, a mathematical model was developed that describes the gas–solid interaction during reduction. This model includes kinetic and thermodynamic evaluation as well as the fluid dynamics of the empty and non-empty reactor. It also includes a refined solid-phase mass balance that considers the particle size distribution and closely tracks the progression of the reaction from. The model couples an axial dispersion model with a dSCM, which is enhanced by a product layer. At each time step, the governing equations are selected dynamically depending on the number of active reaction fronts [13]. he temporal and spatial evolution of the molar content of the individual oxides is described along the axial coordinates z_1 and z_2 (i.e., the fixed-bed axial coordinate). The reactor is divided into a non-reactive inlet zone and a reactive fixed-bed region. The initial molar oxygen content n_0 is calculated based on the mass of the sample. The mass balance of the solid phase, which determines the degree of conversion X_j, is given by

$$\frac{\partial X_j(t,z_2)}{\partial t} = \frac{r_j(t,z_2)}{n_0 \cdot a_j} \quad (1)$$

where a_j is the stoichiometric oxygen content of the oxide j. The overall degree of reduction of the solid is therefore calculated by

$$X(t,z_2) = \sum_j^3 a_j \bullet X_j(t,z_2) \quad (2)$$

The local reduction rate r_j is expressed

$$r_j(t,z_2)\big|_f = \sum_{n=1}^{N} x_n \bullet \frac{P}{R_g T} \frac{3}{r_n^0} \frac{1}{A_j(t,z_2,r_n^0)\big|_f} \cdot R_j(t,z_2,r_n^0)\big|_f \quad (3)$$

$$R_j(t,z_2,r_n^0)\big|_f = \left(\frac{N_j(t,z_2,r_n^0)\big|_f}{W(t,z_2,r_n^0)\big|_f} - y_{\mathrm{H_2},j,\mathrm{eq}}(t,z_2) \right) \quad (4)$$

where P is the pressure, R_g is the universal gas constant and $y_{\mathrm{H_2},j,\mathrm{eq}}$ is the equilibrium molar fraction of hydrogen for the reaction j resulting from the local gas composition and the reaction-specific equilibrium constant K_j.

$$y_{\mathrm{H_2},j,\mathrm{eq}}(t,z_2) = \frac{1}{1+K_j} \cdot \left(y_{\mathrm{H_2},j}(t,z_2) + y_{\mathrm{H_2O},j}(t,z_2) \right) \quad (5)$$

The overall reaction rate is calculated by averaging the rates across the particle size classes, weighted by the initial radius r_n^0 and the mass fraction x_n. The key

parameters of the model include N_j and W, which correspond respectively to the kinetic resistance, the effective reaction rate and the total resistance. These terms evolve dynamically with the number of active reaction fronts, which is determined for each layer of the bed at each time step. A transition function f controls the number of active reaction fronts.

$$\begin{cases} X_1(t,z_2)=0 \implies f=1 \; j=1 \\ X_1(t,z_2)<1 \wedge X_2(t,z_2)=0 \implies f=2 \; j=1,2 \\ X_1(t,z_2)<1 \wedge X_2(t,z_2)<1 \wedge <0 \; X_3(t,z_2)<1 \implies f=3 \; j=1,2,3 \\ X_1(t,z_2)=1 \wedge X_2(t,z_2)<1 \wedge X_3(t,z_2)<1 \implies f=2 \; j=2,3 \\ X_1(t,z_2)=1 \wedge X_2(t,z_2)=1 \wedge X_3(t,z_2)<1 \implies f=1 \; j=3 \\ X_1(t,z_2)=1 \wedge X_2(t,z_2)=1 \wedge X_3(t,z_2)=1 \implies f=0 \end{cases} \quad (6)$$

N_j and W are calculated using the parameters A_j, B_j and F, which represent kinetic resistances, intraparticle diffusion resistances and external diffusion resistances, respectively. These correlations are described below in the case of three active fronts of reactions, two fronts of reactions and a single front.

$$\begin{aligned} N_1(t,z_2,r_n^0)|_3 &= (A_1A_2A_3)\cdot y_{\mathrm{H_2},1} + (A_2(B_3+F))\cdot y_{\mathrm{H_2},3,eq} \\ &+(A_1A_3B_2 + A_1A_3(B_3+F) + A_1B_2(B_3+F))\cdot y_{\mathrm{H_2},2,\mathrm{eq}} \\ &+(A_3B_1(A_2+B_2) + A_3B_1(A_2+B_2) + B_1(A_2+B_2)(B_3+F) \\ &+A_2A_3B_2 + A_2A_3(B_3+F) + A_2B_2(B_3+F))\cdot y_{\mathrm{H_2},1,\mathrm{eq}} \end{aligned} \quad (7)$$

$$\begin{aligned} N_2(t,z_2,r_n^0)|_3 &= (A_2A_3(A_1+B_1))\cdot y_{H_2,2} + (A_2(A_1+B_1) \\ &(B_3+F))\cdot y_{\mathrm{H_2},3,\mathrm{eq}} + (A_3(A_1+B_1)(B_3+F) + B_2(A_1+B_1) \\ &(B_3+F) + A_3B_2(A_1+B_1))\cdot y_{\mathrm{H_2},2,\mathrm{eq}} + (A_2A_3B_2 + A_2A_3B_2 \\ &+A_2B_2(B_3+F) + A_2A_3(B_3+F))\cdot y_{\mathrm{H_2},1,\mathrm{eq}} \end{aligned} \quad (8)$$

$$\begin{aligned} N_3(t,z_2,r_n^0)|_3 &= [A_3(A_1+B_1)(A_2+B_2) + A_2A_3B_2]\cdot y_{\mathrm{H_2},3} \\ &+[(A_1+B_1)(A_2+B_2)(B_3+F) + A_2B_2(B_3+F)]\cdot y_{\mathrm{H_2},3,\mathrm{eq}} \\ &+[A_3(A_1+B_1)(B_3+F)]\cdot y_{H_2,2,eq} + [A_2A_3(B_3+F)]\cdot y_{\mathrm{H_2},1,\mathrm{eq}} \end{aligned} \quad (9)$$

$$\begin{aligned} W(t,z_2,r_n^0)|_3 &= ((A_1+B_1)\cdot(A_3(A_2+B_2+B_3+F) \\ &+(A_2+B_2)\cdot(B_3+F)) + A_2(A_3(B_2+B_3+F) \\ &+B_2(B_3+F)))\cdot y_{\mathrm{H_2},1,\mathrm{eq}} \end{aligned} \quad (10)$$

$$\begin{aligned} N_2(t,z_2,r_n^0)|_2 &= (A_2A_3)\cdot y_{\mathrm{H_2},2} + [A_2(B_3+F)]\cdot y_{\mathrm{H_2},3,\mathrm{eq}} \\ &+[A_2B_1 + (A_3+B_2)(B_3+F)]\cdot y_{\mathrm{H_2},2,\mathrm{eq}} \end{aligned} \quad (11)$$

$$W(t,z_2,r_n^0)|_2 = A_3(A_2+B_2+B_3+F) + (A_2+B_2)\cdot(B_3+F) \quad (12)$$

$$N_3(t,z_2)|_1 = y_{\mathrm{H_2},3}(t,z_2) - y_{\mathrm{H_2},3,\mathrm{eq}}(t,z_2) \quad (13)$$

$$W(t,z_2,r_n^0)|_1 = \frac{A_3+B_3+F}{A_3} \quad (14)$$

The analytical formulations of A_j, B_j, and F are defined as follows

$$A_j(t,z_2)\big|_f = \frac{1}{\left(1-X_j(t,z_2)\right)^{\frac{2}{3}}}\cdot\frac{1}{k_j\cdot\left(1+\frac{1}{K_j}\right)} \quad (15)$$

$$B_j(t,z_2,r_n^0)\big|_f = \frac{\left(1-X_{j+1}(t,z_2)\right)^{\frac{1}{3}}-\left(1-X_j(t,z_2)\right)^{\frac{1}{3}}}{\left(1-X_j(t,z_2)\right)^{\frac{1}{3}}\cdot\left(1-X_{j+1}(t,z_2)\right)^{\frac{1}{3}}}\cdot\frac{r_n^0}{D_j} \quad (16)$$

$$F = \frac{1}{H} \quad (17)$$

where k_j is the kinetic constant, D_j is the intraparticle diffusion coefficient and H is the external mass transfer coefficient. The kinetic parameters were determined by fitting experimental reduction data using the linearized Arrhenius equation, while the thermodynamic and diffusion parameters were obtained from established theoretical and empirical expressions [13].

The gas phase dynamics are modelled by tracking changes in species concentrations over time. The equation for the gas species i is

$$\frac{\partial c_i(t,z_1)}{\partial t} = -\frac{u}{L-h}\cdot\frac{\partial c_i(t,z_1)}{\partial z_1} + \frac{D_z}{(L-h)^2}\cdot\frac{\partial^2 c_i(t,z_1)}{\partial {z_1}^2} \quad (18)$$

$$\frac{\partial c_i(t,z_2)}{\partial t} = -\frac{u}{h}\cdot\frac{\partial c_i(t,z_2)}{\partial z_2} + \frac{\partial^2 c_i(t,z_2)}{\partial {z_2}^2} + (1-\varepsilon)\cdot\sum_{j|_f}^{3}\nu_{i,j}\cdot r_j(t,z_2) \quad (19)$$

here, c_i is the concentration of the species i (H_2 or H_2O), $\nu_{i,j}$ are stoichiometric coefficients, and ε is the initial porosity of the bed, measured via nitrogen physisorption. The relevant parameters are the bed height h, superficial velocity u and the axial dispersion coefficient D_z. The bed height used corresponds to the theoretical height calculated based on the sample density and reactor geometry, while the axial dispersion coefficient was determined through fluid dynamics studies**.** The evaluation of the residence time distributions of the experimental setup enabled the evaluation of the Peclet number and of the axial dispersion coefficient D_z.

The boundary conditions of the PDE system are defined at the reactor inlet (z_1 = 0), at the outlet of the empty zone (z_1 = 1), at the inlet of the reactive bed (z_2 = 0) and at the system outlet (z_2 = 1).

$$c_i(t,z_1)|_{z_1=0} = c_{i,\mathrm{in}} \quad \frac{\partial c_i(t,z_1)}{\partial z_1}\bigg|_{z_1=1} = 0 \quad (20)$$

$$c_i(t,z_2)|_{z_2=0} = c_i(t,z_1)|_{z_1=1} \quad \frac{\partial c_i(t,z_2)}{\partial z_2}\bigg|_{z_2=1} = 0 \quad (21)$$

A detailed assessment was performed to identify which phenomena could be safely neglected in the kinetic modelling of hydrogen-based iron oxide reduction under controlled laboratory conditions. The analysis indicates that external mass transfer limitations can be neglected due to the very fine particle size used in the experiments. A detailed evaluation of the external mass transfer coefficient under different operating conditions confirmed values of H high enough to rule out significant resistance compared to other mechanisms in the dynamic shrinking core model. Similarly, thermal effects were assessed to verify the assumption of negligible temperature gradients within the particles. Despite the endothermic nature of the reduction reactions, the use of small beds of fine particles combined with strong external heating ensures nearly uniform temperature. Theoretical estimates of intraparticle temperature differences under worst-case conditions showed variations below 2 K for all oxides involved, supporting the assumption of isothermal behavior. These findings justify focusing on chemical and internal diffusion control while excluding external diffusion and thermal gradients from the kinetic model.

The mathematical model was implemented in Matlab 2025b to simulate the reduction process and estimate kinetic parameters. The system of partial differential equations was solved using the finite difference method, applying a backward scheme for first-order spatial derivatives and a centered scheme for second-order derivatives. The reactor's two axial coordinates were discretized into 20 nodes each, generating a system of ordinary differential equations solved with ode15s, a variable-step, variable-order solver based on numerical differentiation formulas (orders 1–5).

Parameter estimation was carried out using patternsearch, a deterministic, derivative-free optimization algorithm that explores the solution space through an adaptive grid (pattern). This approach ensures robustness in handling complex, nonlinear systems without requiring gradient information.

The objective function minimized the deviation between experimental data and the simulated average conversion of the thin layer over time ($\bar{X}_s$), expressed as:

$$O = \sum_s^S \left(\frac{X_{\text{meas},s} - \bar{X}_s}{X_{\text{meas},s}}\right)^2 \quad (21)$$

This formulation guarantees accurate fitting of experimental results, enabling reliable estimation of kinetic parameters for multistage iron oxide reduction.

RESULTS

The interpretation of the TCD response is inherently challenging. This difficulty stems not only from the conductivity contrast between the reference and outlet gas streams, but also from the specific physical layout and internal configuration of the detector. By applying a rigorous mathematical treatment, it was possible to link the processed signal directly to the conversion progress of the reacting solid. The operating conditions examined in this work include temperature, reducing-gas composition, volumetric flow rate, and the solid-sample loading in the reactor, as well as the type of solid employed, namely pure hematite and industrial samples used in BF and DRI processes. In total, 17 experiments were carried out for each type of solid sample, extensively varying all the different operating conditions.

Among all factors influencing the reduction of iron oxides, temperature is one of the most influential. The reason lies in the strong correlations that temperature has with both the reaction kinetics and the thermodynamics of the reduction process, as well as with diffusion-related phenomena. Figure 4 highlights how temperature affects the overall reduction degree over time. Panel (a) presents tests performed on pure hematite beds, whereas panel (b) and (c) displays results obtained from the reduction of compact BF and DRI powder respectively.

A key observation is the marked decrease in the total reduction time as temperature increases, reflecting the enhanced reaction rate. The data indicate that the conversion of hematite to magnetite occurs rapidly—within approximately 4 minutes—in all experiments. Some of them show a distinct change in the slope of the conversion curve at approximately 0.11 signals the completion of the hematite-to-magnetite transition. Further slope variations are more subtle, as the reductions of Fe_3O_4 and FeO overlap significantly. Consequently, a conversion value near 0.3 cannot be uniquely associated with the full reduction of magnetite. In these experiments, the height of the bed plays a crucial role in determining the reaction pathway. The water vapor formation introduces thermodynamic constraints deeper in the bed, observed as a noticeable slope break around a conversion of 0.2-0.3, identifying the approximation to the end of magnetite reduction and the initial stages of FeO formation.

The kinetic parameters for pure hematite were previously derived in an earlier study [14], whereas the parameters reported in Table 3 were obtained from the experiments performed on the commercial materials.

Table 3: Activation energy and pre-exponential factor estimated from BF and DRI powder.

Reactions	E_a [kJ/mol]	$k_{j,\text{ref}}$ [m/s]
$Fe_2O_3 \rightleftharpoons Fe_3O_4$	61 ± 3	$2.8 \cdot 10^{-1} \pm 6 \cdot 10^{-2}$
$Fe_3O_4 \rightleftharpoons FeO$	37 ± 2	$2.4 \cdot 10^{-3} \pm 9 \cdot 10^{-5}$
$FeO \rightleftharpoons Fe$	27 ± 1	$2.6 \cdot 10^{-3} \pm 6 \cdot 10^{-5}$

Figure 5 compares simulated and measured conversion values through reduction degree parity plots of all the different datasets: pure (a) BF (b) and DRI (c).

The plots show strong agreement between model predictions and experimental measurements, with most data points falling within a ±10% deviation. This consistency is further supported by the corresponding R^2 values. The robustness of the kinetic parameters was additionally confirmed by applying them to a broader set of reduction experiments involving hematite fixed beds.

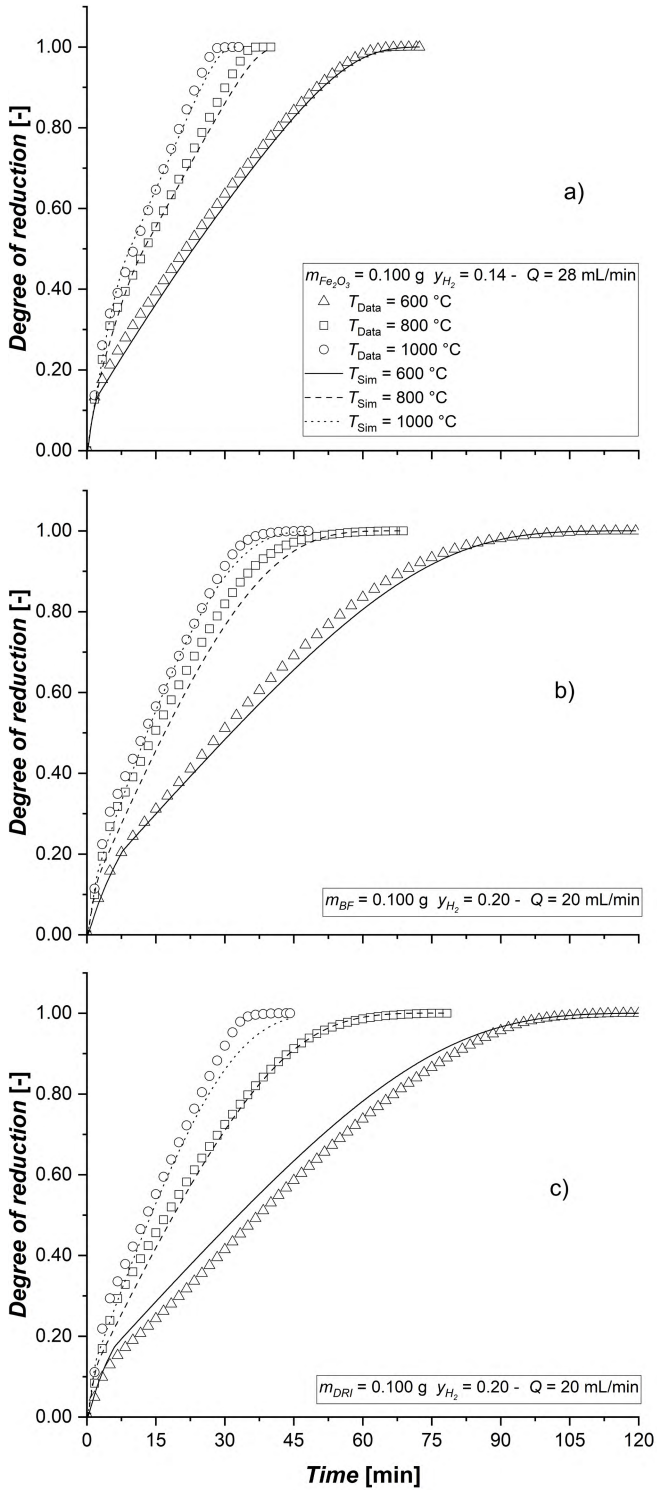

Figure 4. Influence of temperature on iron oxide reduction: a) pure hematite; b) BF; c) DRI.

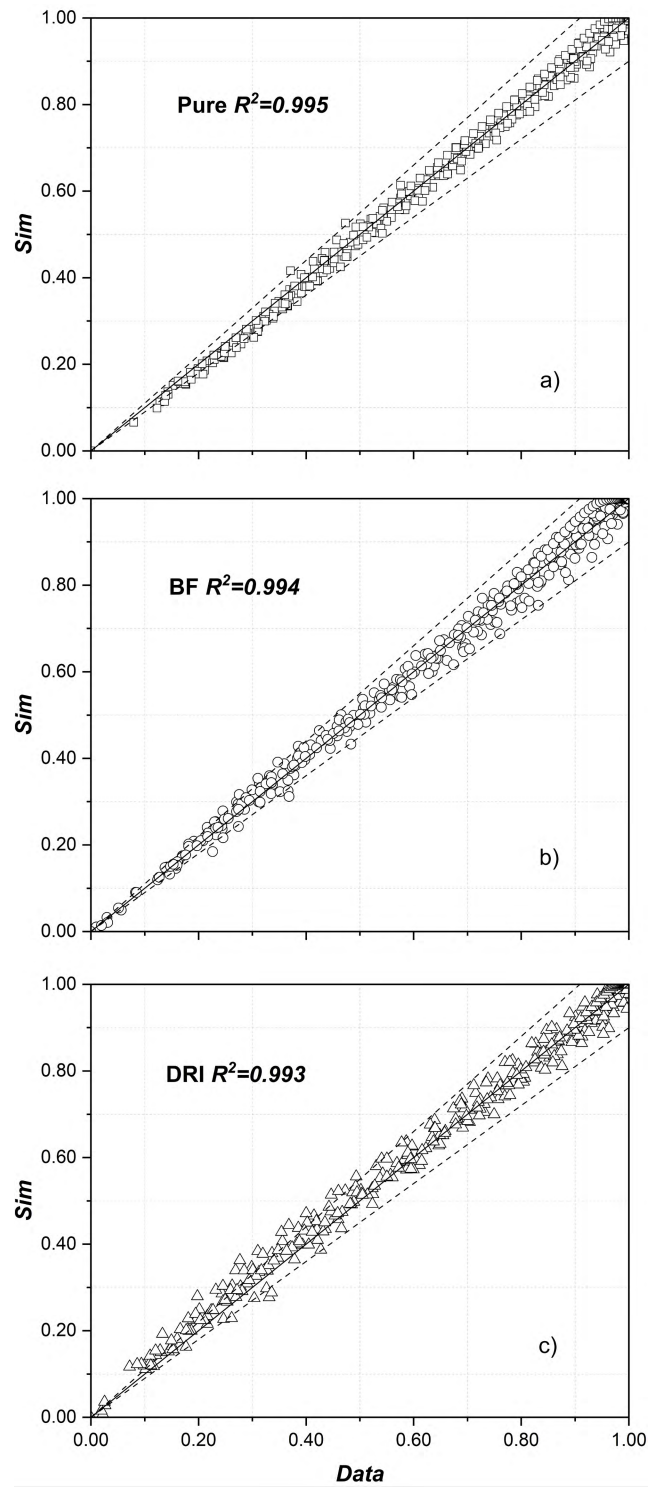

Figure 5. Reduction degree parity plots for a) pure hematite; b) BF; c) DRI.

CONCLUSIONS

This work has presented a comprehensive experimental and modelling investigation of the hydrogen-based reduction of iron oxide powders, with the aim of estimating the kinetic parameters of the reactions involved. In this study, the reduction of pure hematite, BF, and DRI powders in a fixed-bed reactor was examined under a wide range of operating conditions, including different temperatures, total volumetric flow rates, hydrogen contents in the feed gas, and solid sample loadings. The experiments were performed in a fixed-bed reactor equipped with a thermal conductivity detector, enabling fast and accurate gas-phase analysis.

The TCD signal was processed using a mathematical approach capable of correlating the thermal-conductivity response of the off-gas to the reduction degree of the sample. This indirect measurement strategy was rigorously validated through complementary analytical techniques, including morphological, chemical, structural, and granulometric characterization of the solids.

A novel mathematical modeling framework was developed, integrating a dynamic multi-front shrinking core formulation with a variable number of reaction interfaces and an axial dispersion description to accurately represent the behavior of the experimental system. The model dynamically tracks the sequential reduction reactions and their spatial evolution within the bed, enabling a detailed description of the reduction process under varying thermodynamic conditions. Numerical simulations showed excellent agreement with the experimental results over a wide range of operating conditions, demonstrating the predictive capability of the framework.

In parallel, a novel methodology based on thermal conductivity detector signal analysis was developed to determine the reduction degree directly from off-gas measurements. Through a deterministic signal-processing procedure, the evolution of the reaction could be continuously monitored, providing a robust and reproducible alternative to conventional gravimetric approaches for studying hydrogen-based reduction systems.

Importantly, the model proved capable of consistently describing the reduction behavior of different iron-bearing materials, including pure iron oxides, blast furnace materials, and direct reduced iron systems. Overall, the results demonstrate that the developed methodology—combining rigorous experimentation, advanced gas-phase diagnostics, solid-phase characterization, and physically grounded modeling—provides a robust and versatile framework for interpreting the dynamics of hydrogen-based iron oxide reduction.

ACKNOWLEDGEMENTS

We gratefully acknowledge financial support from the European Union within the projects Minimisation of CO_2 emissions from the BF by hydrogen containing injectants and use of DRI/HBI during transition to new ironmaking processes until 2030 (grant agreement 101057790) and Maximise H_2 Enrichment in Direct Reduction Shaft Furnaces (grant agreement 101058429). Views and opinions expressed in the paper are, however, those of the authors only and do not necessarily reflect those of the European Union, so the European Union cannot be held responsible for them.

AUTHOR IDENTIFIERS

Author ORCIDs: Salucci Emiliano 0009-0003-0673-2085

REFERENCES

1. Birat JP. Society, materials, and the environment: the case of steel. Metals 10:331 (2020). https://doi.org/10.3390/met10030331
2. Geerdes M, Chaigneau R, Lingiardi O, Molenaar R, van Opbergen R, Sha Y, Warren P. Modern blast furnace ironmaking. IOS Press (2020). https://doi.org/10.3233/stal9781643681238
3. Lüngen H B, Schmöle P, History, developments and processes of direct reduction of iron ores Online. 2022.
4. Number D, Eichberger H, Rollinger B, Circored fine ore direct reduction. 2006 stahl u.eisen, Vol. 126, no. 4.
5. El-Geassy AA. Gaseous reduction of fe2o3 compacts at 600 to 1050° c. J Mater Sci 21:3889-3900 (1986). https://doi.org/10.1007/bf02431626
6. Hessels CJM, Homan TAM, Deen NG, Tang Y. Reduction kinetics of combusted iron powder using hydrogen. Powder Technology 407:117540 (2022). https://doi.org/10.1016/j.powtec.2022.117540
7. He K, Zheng Z, Chen Z, Chen H, Hao W. Kinetics of hydrogen reduction of brazilian hematite in a micro-fluidized bed. International Journal of Hydrogen Energy 46:4592-4605 (2021). https://doi.org/10.1016/j.ijhydene.2020.10.263
8. D'Angelo A, Salucci E, Russo V, Grénman H, Saxén H. Hydrogen reduction of iron oxide powder in thin layers. Ind. Eng. Chem. Res. 64:158-170 (2024). https://doi.org/10.1021/acs.iecr.4c03138
9. Salucci E, D'Angelo A, Russo V, Grénman H, Saxén H. Review on the reduction kinetics of iron oxides with hydrogen-rich gas: experimental investigation and modeling approaches. Ind. Eng. Chem. Res. 64:1-35 (2024). https://doi.org/10.1021/acs.iecr.4c03136
10. Chen F, Mohassab Y, Jiang T, Sohn HY. Hydrogen reduction kinetics of hematite concentrate particles

relevant to a novel flash ironmaking process. Metall Mater Trans B 46:1133-1145 (2015). https://doi.org/10.1007/s11663-015-0332-z
11. Jozwiak WK, Kaczmarek E, Maniecki TP, Ignaczak W, Maniukiewicz W. Reduction behavior of iron oxides in hydrogen and carbon monoxide atmospheres. Appl Catal A: General 326:17–27 (2007) https://doi.org/10.1016/j.apcata.2007.03.021
12. Pineau A, Kanari N, Gaballah I. Kinetics of reduction of iron oxides by H2. Thermochimica Acta 447:89-100 (2006). https://doi.org/10.1016/j.tca.2005.10.004
13. Salucci E, D'Angelo A, Russo V, Grénman H, Saxén H. Modelling of iron oxide reduction with hydrogen in a small fixed bed. Chemical Engineering Science 292:119934 (2024). https://doi.org/10.1016/j.ces.2024.119934
14. Salucci E, D'Angelo A, Russo V, Grénman H, Saxén H. Kinetic modeling and parameter estimation of thin-layer iron oxide reduction with hydrogen in a fixed-bed reactor. Chemical Engineering Journal 529:172848 (2026). https://doi.org/10.1016/j.cej.2026.172848

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A General Framework for Model Recognition in Chemical Reactor Systems Using Artificial Neural Networks Classifiers

Emmanuel Agunloye, Asterios Gavriilidis, and Federico Galvanin*

Department of Chemical Engineering, University College London, London WC1E 7JE, UK
*Corresponding Author: f.galvanin@ucl.ac.uk

ABSTRACT

The identification of predictive mathematical model structures (i.e. set of model equations) is essential for the development of digital twin models of chemical reactor systems. Recent work demonstrated the use of artificial neural networks (ANNs) for kinetic model recognition in a conceptual batch reaction experimental system. In practical chemical processes, however, system behaviour is governed not only by reaction kinetics but also by reactor hydrodynamics and system thermodynamics. While a very recent study incorporated hydrodynamic effects, this work integrates the three aspects: reaction kinetics, reactor hydrodynamics, and system thermodynamics, to develop a general reactor modelling recognition framework. The framework, which comprises three modules: 1) model generator module; 2) data generation module; and 3) ANN classifier module, was applied to a case study of benzoic acid esterification in a Taylor vortex flow reactor system. Analysing the framework's sensitivity, results showed that ANN performance in classification deteriorates under increasing measurement noise but can be improved by increasing the number of simulated experiments. Further results show that extensive hyperparameter optimisation of ANN architectures provides no benefit over a fixed ANN architecture. This study highlights the potential of ANN-based frameworks for reactor model recognition while underscoring the dominant role of experimental design and data quality over network hyperparameter tuning.

Keywords: Modelling, Machine Learning, Optimization, Modelling and Simulations, Process Operations, Hybrid modelling, Artificial neural networks, Taylor vortex flow reactor

1.0 INTRODUCTION

Digital twin (DT) technology is transforming the way chemical and pharmaceutical processes are designed, monitored, and optimised [1 - 3]. A digital twin combines real-time data from a physical process with predictive computational models to enable autonomous operation and closed-loop optimisation [2]. In pharmaceutical manufacturing, DTs play a central role in accelerating process development, ensuring consistent product quality, and supporting regulatory compliance.

Despite the rapid rise of data-driven and machine learning (ML)-based digital twins, their adoption in safety-critical industries such as pharmaceuticals has been hindered by concerns about reliability, interpretability, and robustness [4]. Purely data-driven models can capture system trends efficiently but often fail to extrapolate beyond the training domain, leading to unsafe or suboptimal operation when process conditions change. To overcome these limitations, hybrid digital twins—which combine the computational efficiency of ML with the physical interpretability of first-principles models—have gained increasing attention [5-7].

Within this context, model identification remains a crucial step. Reliable digital twins require high-fidelity process models capable of representing reaction kinetics, reactor hydrodynamics, and thermodynamic interactions [8]. Recent research has demonstrated the use of artificial neural networks (ANNs) for rapid kinetic model recognition and optimal experimental design, particularly in batch reactor systems [9, 10]. However, these studies often neglected hydrodynamic and thermal effects whose representation is critical in nonisothermal flow reactors, which are increasingly central to pharmaceutical

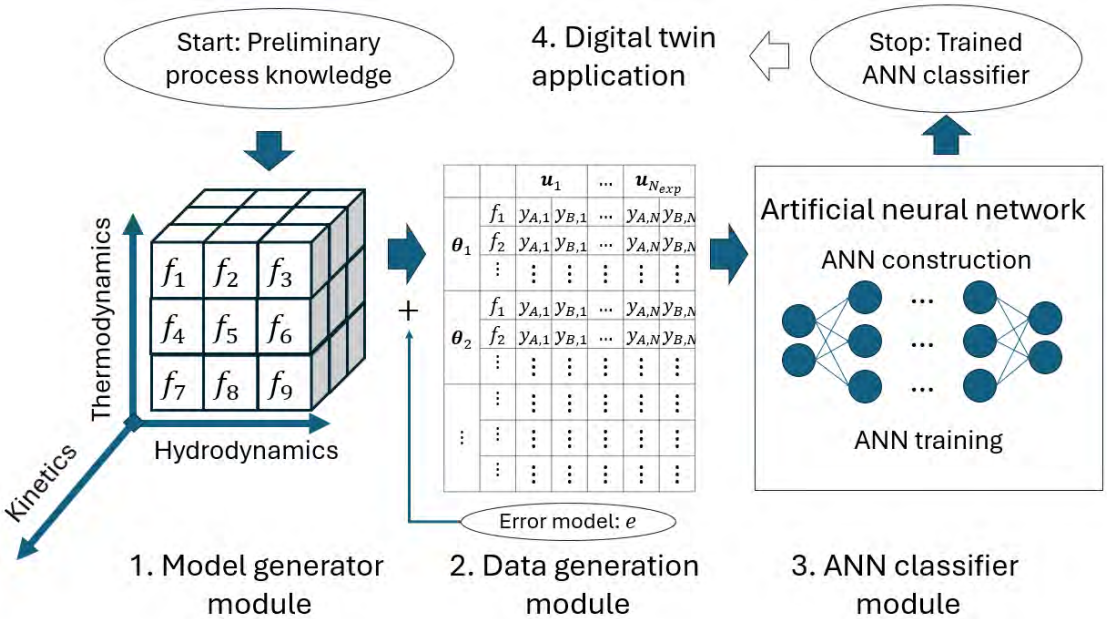

Figure 1: ANN hybrid modelling framework with 3-dimensional physics-based model generator

scale-up.

The present work extends previous ANN-based frameworks [9–11] by incorporating hydrodynamics and thermodynamics explicitly into the model recognition process. The extended framework is subsequently evaluated with respect to data quality and ANN architectural complexity to assess its scientific robustness and computational efficiency. The proposed framework identifies the most suitable reactor model structure among a library of possible configurations including kinetics, hydrodynamics, and energy interactions. This hybrid approach enables fast yet physically grounded model recognition — a key requirement for reliable DT implementation.

2.0 METHODOLOGY

The proposed hybrid framework comprises three interconnected modules (Fig. 1): 1. Model Generator Module, 2. Data Generation Module, and 3. ANN Classifier Module, for subsequent application as a digital twin.

2.1 Model Generator Module

The model generator defines the library of possible reactor models, organised as a three-dimensional tensor. Each axis of this tensor corresponds to a distinct modelling domain:

- Reaction Kinetics: Alternative rate-law formulations capturing alternative mechanistic hypotheses;
- Reactor Hydrodynamics: Different reactor flow models, considering in this work continuously stirred tank reactor (CSTR) and plug flow reactor (PFR) representations;
- System Thermodynamics: Alternative energy balance assumptions—adiabatic and isothermal operations.

Each model structure i is generally described by a system of differential and algebraic equations that can be written as:

$$\boldsymbol{f}_i(\boldsymbol{x}, \dot{\boldsymbol{x}}, \boldsymbol{u}, \theta, \tau) = \boldsymbol{0} \qquad \boldsymbol{y} = \boldsymbol{g}(\boldsymbol{x}) \tag{1}$$

where $\boldsymbol{x} \in \mathcal{X} \subseteq \mathbb{R}^{n_x}$: vector of state variables (e.g., concentrations and temperatures), $\dot{\boldsymbol{x}} \in \dot{\mathcal{X}} \subseteq \mathbb{R}^{n_x}$: vector of time derivatives of the state variables; $\boldsymbol{u} \in \mathcal{U} \subseteq \mathbb{R}^{n_u}$: vector of control/input variables (e.g., inlet temperature, stream flow rates and composition in flow reaction systems); $\theta \in \boldsymbol{\Theta} \subseteq \mathbb{R}^{n_\theta}$: vector of model parameters (e.g., reaction Arrhenius parameters); τ: space time variable; $\boldsymbol{y} \in \mathcal{Y} \subseteq \mathbb{R}^{n_y}$: vector of measured variables.

2.2 Data Generation Module

The data generation module produces synthetic datasets by simulating each model structure over the defined parameter and design spaces, following three main features:

- Parameter Space Sampling: Kinetic parameter sets $\{\dots, \theta_i, \dots\}$ are sampled from the feasible space of model parameters [9] using Latin hypercube

sampling (LHS) to ensure space-filling coverage.

- Experimental Design Space Sampling: Operating conditions $\boldsymbol{u}_i = [\boldsymbol{c}_0, T, \tau]$, denoting feed concentrations, temperature, and residence time, respectively are randomly designed using the Latin hypercube sampling technique from the experimental design space $\boldsymbol{\varphi}$ before employing the models to simulate the experimental conditions.
- Noise Modelling: to assess the impact of data quality on model identification Gaussian noise (e) with zero mean and relative variance ω has been varied between 0.01 and 0.1, and added to model outputs $\hat{\boldsymbol{y}}$ to mimic measurement uncertainty in real experiments, i.e.,

$$e = \mathcal{N}[0, \sigma^2 = (\omega y)^2] \tag{2}$$

The simulated data comprise model outputs and their labels, which serve as training and validation data set for ANN classifiers.

2.3 ANN Classifier Module

The third module performs model recognition using ANNs trained to classify which model structure best represents the test data. Two ANN architectures are considered:

- ANNfix: Fixed network structure comprising 3 hidden layers with 160 nodes per layer
- ANNopt: Hyperparameter-optimised network, varying the number of hidden layers (n_h = 1 – 5) and the number of nodes per layer (n = 32-256)

The ANN classification transformation can be described as:

$$p_i = \phi_{N_L}(\dots \phi_1(\boldsymbol{w}^T \boldsymbol{y} + \boldsymbol{b}) \dots); \forall i = [1, 2, \dots, N_m] \tag{3}$$

where p_i is the probability of ANN selection calculated for Model f_i; $\boldsymbol{w}$ and $\boldsymbol{b}$ are the weights and biases in ANN nodes respectively; ϕ_1 is the activation function of layer 1. For classification, SoftMax activation function is used in Layer N_L while ReLU activation function is used in all the preceding layers [12]. ANN parameters $\boldsymbol{w}$ and $\boldsymbol{b}$ are estimated from model training using standard backpropagation with categorical cross-entropy as the objective function. Classification accuracy $\mathrm{Acc}_{\Psi_{\mathrm{test}}}$—the proportion of correctly identified model structures in the test dataset, expressed in Eq. (4)—serves as the primary performance metric.

$$\mathrm{Acc}_{\Psi_{\mathrm{test}}} = \frac{N_{\Psi_{\mathrm{test}} \wedge \hat{l}_i = l_i}}{N_{\Psi_{\mathrm{test}}}} \bullet 100\% \tag{4}$$

where $N_{\Psi_{\mathrm{test}}}$ is the number of samples in the test dataset; $N_{\Psi_{\mathrm{test}} \wedge \hat{l}_i = l_i}$ is the number of test samples where the model label has been correctly classified.

3.0 CASE STUDY

The synthesis case study considered in this work is related to benzoic acid esterification [13] in a Taylor vortex reactor (TVR). In this esterification reaction, benzoic acid (BA) reacts with ethanol (EtOH) to produce water (W) and ethylbenzoate (EB) in a single reaction shown in Eq. 5. This reaction can be described using two rival reaction kinetics shown in equations M1 and M2, providing two options along the kinetic submodel dimension [13]

$$BA + EtOH \rightarrow W + EB \tag{5}$$

$$k_1 c_{BA} c_{EtOH} \tag{M1}$$

$$\frac{k_1 c_{BA} c_{EtOH}}{(1 + K_W c_W)^2} \tag{M2}$$

where: $k_1 = \exp\left(-KP_1 - \frac{KP_2 \times 10000}{RT}\left[\frac{1}{T} - \frac{1}{T_{ref}}\right]\right)$;

$$T_{ref} = 378.15\ K \tag{6}$$

Thus, $\boldsymbol{u} = [c_{BA}(0), T, F_{total}]$; $\hat{\boldsymbol{y}} = [c_{BA}, c_{EtOH}, c_W, c_{EB}, T]$ and $\theta = [KP_1, KP_2, K_W]$.

A TVR of 10 mL volume with 9.5 and $10.2\ mm$ in annulus inner and outer diameters is employed. Because of its rotating concentric cylinders (inner rotor and outer stator as shown in Fig. 2), the TVR offers mixing conditions that can vary from continuously stirred tank reactor (CSTR) (Eq. R1) to plug flow reactor (PFR) (Eq. R2) [14], so that a trained ANN within the hybrid framework will be able to recognise the prevailing hydrodynamics from experimental data.

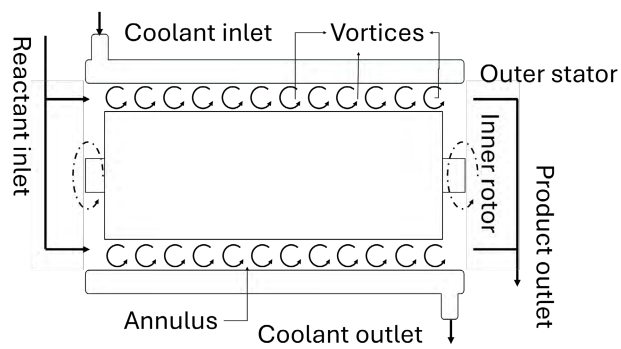

Figure 2. Axial cross section of a Taylor vortex reactor

Only this two hydrodynamics are considered in this work, yielding two options along the reactor submodel dimension.

$$c_i = c_{i0} + \frac{V}{F_{total}} \upsilon_i r; \tag{R1}$$

$$\frac{dc_i}{dV} = F_{total} \upsilon_i r; \tag{R2}$$

The reactor temperature profile expressed using the energy balance reported in Eq. (7) [15] can vary from an isothermal thermodynamic operation (Eq. D1), where the

reaction system temperature is controlled and under a constant temperature to an adiabatic thermodynamic operation (Eq. D2), where the reaction system temperature is uncontrolled and influenced by the reaction heat. These two operating modes provides the options along the thermodynamic submodel dimension.

$$dH = Q + W_s \tag{7}$$

where dH is the enthalpy change; Q is the heat transferred and W_s the shaft work impacted on the system.

$$T = constant \tag{D1}$$

$$Q = 0 \tag{D2}$$

Using these three submodel dimensions, 8 candidate models $(f_1, f_2, \dots, f_8)$ can be generated as shown in Table 1.

Table 1: The set of candidate models $\boldsymbol{f} = [f_1, f_2, \dots, f_8]$ obtained from the model generator.

	$D1$		$D2$	
	$R1$	$R2$	$R1$	$R2$
$M1$	$f_1 = (M1, R1, D1)$	$f_3 = (M1, R2, D1)$	$f_5 = (M1, R1, D2)$	$f_7 = (M1, R2, D2)$
$M2$	$f_2 = (M2, R1, D1)$	$f_4 = (M2, R2, D1)$	$f_6 = (M2, R1, D2)$	$f_8 = (M2, R2, D2)$

4.0 RESULTS AND DISCUSSION

Table 2 shows the experimental design for benzoic acid esterification where the reported kinetic models are valid.

Table 2: Benzoic acid experimental design space with three design variables and their lower and upper limits

Limits	$c_{BA}(0)$ (M)	T (K)	F_{total} $(\mu L/min)$
Lower	1.0	373.0	20.0
Upper	1.5	393.0	40.0

The impacts of three factors on the ANN classification performance are investigated: number of experiments, noise level and ANN architecture settings. In the number of experiments, five simulated experiment numbers were employed: 1 to 5, each experiment corresponding to five responses (= 4 exit chemical concentrations + 1 exit temperature). Fig. 3 shows the distribution of the 5 experiments within the experimental design space.

Two levels of constant relative noise were investigated: low (=0.01) and high (=0.1) noise level. In the ANN architecture settings, ANNfix and ANNopt were investigated. The Python optimiser package HyperOpt was employed for ANNopt. HyperOpt optimised the ANN layers and nodes along with their weights and biases during the training.

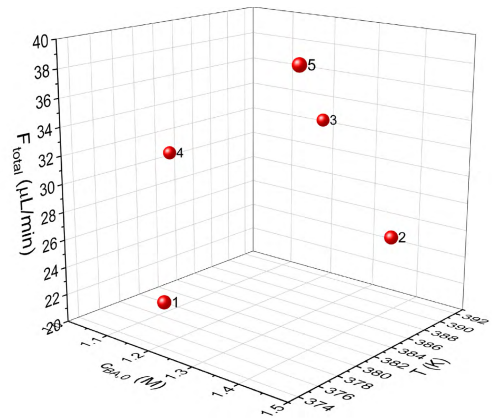

Figure 3. Distribution of 5 experiments designed using LHS in the experimental design space

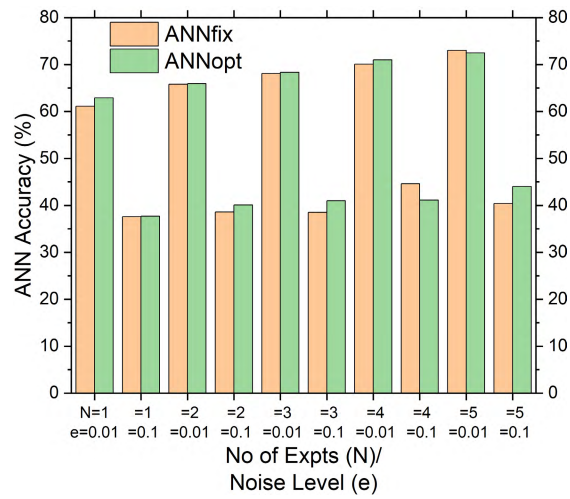

Figure 4. Comparison of the performance of the ANN network on the test dataset for the two different ANN architectures (ANNfix and ANNopt) for number of experiments between 1 and 5 and level of noise of 0.01 and 0.1.

Each of the 8 models in the model library was simulated using 1, 000 samples from the kinetic parameter space shown in Table 3. The model simulated responses in concentrations and temperature were subsequently corrupted randomly from a Gaussian distribution with a constant relative variance equalled to the level of noise.

Table 3: Benzoic acid esterification kinetic parameter space [13]

KP_1	KP_2	K_W
16.7 ± 0.02	5.93 ± 0.20	0.27 ± 0.10

The resulting in-silico data, split into training, validation and testing datasets in the ratio: 3:1:1, are fed into the ANN network. The ANN performance on the test dataset using the accuracy metric is shown in Fig. 4.

Results show that there is no significant difference between the performances of ANNfix and ANNopt in terms of global classification accuracy for the same number of experiments and level of noise. For example, when $N_{exp} = 5$ and noise level = 0.01, ANNfix achieved an accuracy of 73.0% while ANNopt achieved a slightly lower accuracy of 72.5%. At the lower number of experiments when $N_{exp} = 1$ and noise level = 0.01, ANNfix achieved an accuracy of 61.1% while ANNopt achieved only slightly a higher accuracy of 62.9%, a marginal increment not substantial enough to offset the required computational cost (86.0 seconds VS 6.7 seconds of ANNfix) and largely due to the stochastic shuffling of the datasets. Both architectures have the SoftMax activation function in the last layer required for the classification task. The input into the layer has been calculated using the trained prior network parameters and the last layer can still perform the classification, regardless of the network architectures. Further investigation confirmed that, provided the number of nodes per hidden layer in ANNfix was greater or equal to the input features, the hyperparameter optimisation employed in ANNopt did not confer a performance advantage over ANNfix. In the investigation at $N_{exp} = 20$ and 0.01 noise level, ANNfix achieved 86.3% while ANNopt achieved 86.9%. Similarly, a change in the design technique did not influence this trend. Using the full factorial design at 2 levels in the experimental design space, yielding 8 experiments from the corners in Fig. 3, ANNfix achieved 79.9% ANN classification accuracy but ANNopt performed about 2% worse at 78.1% accuracy.Different from the ANN architecture settings, other factors such as the noise level and number of experiments influence the network performance. For example, while ANNfix at a small level of noise (0.01) achieved an ANN performance of 61.1% with a single experiment, the same network architecture only achieved a performance of 37.6% at the higher noise level of 0.1. The same trend applies to ANNopt considering a single experiment, achieving 62.9% at 0.01 small noise level and 37.7% at 0.1 large noise level. By increasing the number of experiments to 2, both architectures improved in ANN performance to 65.8% and 65.9% at 0.01 noise level and to 38.6% and 40.1% at 0.1 noise level for ANNfix and ANNopt, respectively. The same trend can be observed by increasing the number of experiments further until, $N_{exp} = 5$, which has the highest ANN accuracy for both small and large noise levels in both ANN architectures. Increasing number of experiments increases the performance of the ANN while increasing noise level deteriorates the ANN performance.

The classification accuracy is illustrated graphically in Fig. 5 using the confusion matrices with two scenarios: Scenario 1 and Scenario 2. Scenario 1 corresponds to $N_{exp} = 5$ and 0.01 noise level where the achieved classification accuracy was 73.0% and Scenario 2 corresponds to $N_{exp} = 1$ and 0.1 noise level where the accuracy was 37.6 %, both considered using ANNfix architecture. The candidate models are labelled as in Table 4.

Table 4: The 8 candidate models: labels and components

Model Labels	Model components
Model 0	M1, R1, and D1
Model 1	M2, R1, and D1
Model 2	M1, R2, and D1
Model 3	M2, R2, and D1
Model 4	M1, R1, and D2
Model 5	M2, R1, and D2
Model 6	M1, R2, and D2
Model 7	M2, R2, and D2

The small noise in Scenario 1 only gives rise to classification confusion among alternative thermodynamic models (D1: isothermal and D2: adiabatic). For example, 76 out of 191 test datasets originating from Model 0 (kinetic M1, CSTR hydrodynamic R1 and isothermal thermodynamic D1) were wrongly classified as M2 (kinetic M1, CSTR hydrodynamic R1 and adiabatic thermodynamic D2). In a few instances data sets originating from a model with a CSTR hydrodynamic sub model were wrongly classified as PFR (Actual Model 2 and Predicted Model 7) or data sets originating from a model with kinetic M1 were wrongly classified as kinetic M2 (Actual Model 7 and Predicted Model 6).

In Scenario 2 with larger noise level, confusion in the ANN network classification extends beyond just one model label to multiple model labels. For example, test datasets originating from Model 0 are being wrongly classified as Model 1 (2), Model 2 (45), Model 3 (17), Model 4 (1), Model 5 (16) and Model 7 (11). Only Model 6 (kinetic M1, PFR hydrodynamic R2 and Adiabatic thermodynamic D2) was not wrongly predicted.

As a trend, the ANN performance decreases with increasing noise in the training data. At 0.01 noise variance and five simulated experiments, both ANNfix and ANNopt achieved comparable classification accuracies—73.0% and 72.5%, respectively. As noise increased to 0.1, accuracies fell below 60%, highlighting the detrimental impact of experimental uncertainty on model recognition.

Contrary to expectations [16, 17], extensive hyperparameter optimisation (ANNopt) did not improve classification performance over the fixed network (ANNfix). This suggests that data quality and diversity exert a stronger influence on model recognition accuracy than ANN complexity.

Increasing the number of simulated experiments substantially improved classification accuracy, even

under high-noise conditions. This reinforces the critical role of experimental design in ensuring model identifiability and supports the use of optimised designs in physical experimentation.

The framework provides a robust computational foundation for experimental studies in flow reactors, particularly TVRs, where mixing and thermal effects interact strongly. The ability to recognise between hydrodynamic and kinetic models rapidly can inform both reactor design and digital twin development for scale-up.

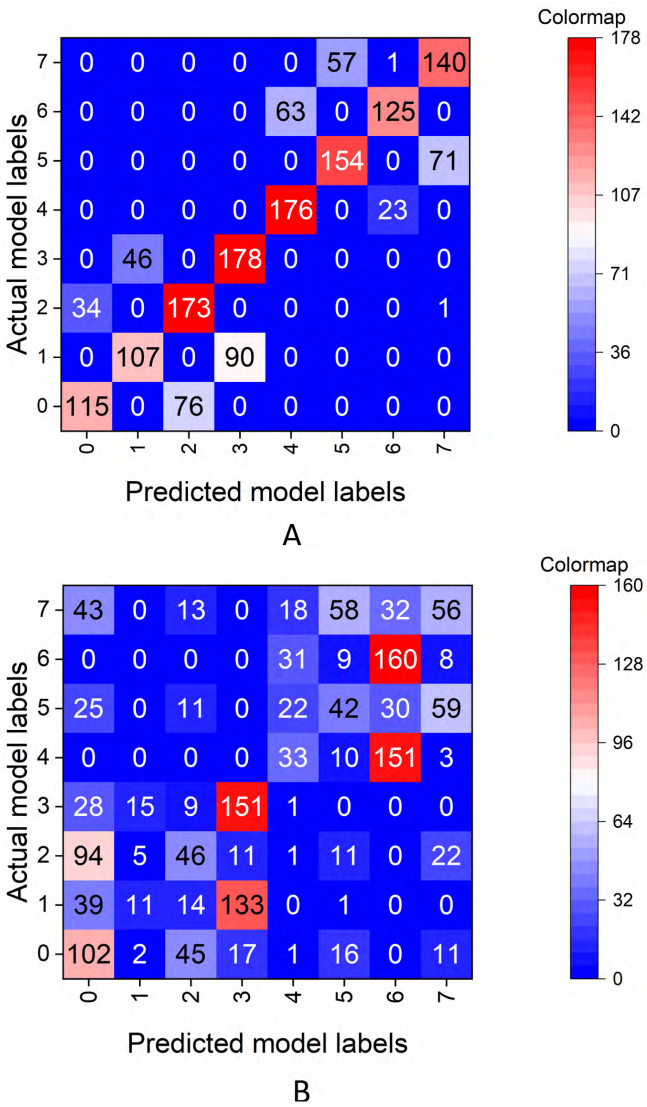

Figure 5. ANN test classification shown as confusion matrices. A) Scenario 1 ($N_{exp} = 5$, 0.01 noise level, 73.0 % accuracy); B) Scenario 2 ($N_{exp} = 1$, 0.1 noise level, 37.6 % accuracy; Model 0 (M1, R1, D1), Model 1 (M2, R1, D1), Model 2 (M1, R2, D1), Model 3 (M2, R2, D1), Model 4 (M1, R1, D2), Model 5 (M2, R1, D2), Model 6 (M1, R2, D2) and Model 6 (M2, R2, D2).

5.0 CONCLUSIONS

This study presents a hybrid digital twin framework for reactor model recognition integrating reaction kinetics, hydrodynamics, and thermodynamics. The framework combines first-principles modelling with ANN classification to achieve fast, data-driven yet physically meaningful model identification in a flow reactor system. Simulated results related to a case study of benzoic acid esterification in a Taylor vortex reactor show that ANN classification performance is highly sensitive to data noise but less dependent on network hyperparameters. Also, increasing the number of designed experiments mitigates noise effects and enhances model recognition accuracy. Hydrodynamic and thermal model inclusion improves the generality of the digital twin and its interpretability.

This work highlights the importance of combining physics-based and data-driven approaches in reactor model identification and provides a foundation for hybrid digital twin development in chemical and pharmaceutical manufacturing. Future work will therefore focus on the extension of the model library and the integration with the physical system.

DIGITAL SUPPLEMENTARY MATERIAL

Source Code: The code used in this work is available in GitHub at https://github.com/emmanagun/A-General-Artificial-Neural-Networks-Framework-for-Model-Recognition-in-Chemical-Reactor-Systems

ACKNOWLEDGEMENTS

This research was funded by EPSRC (grant number: EP/X024016/1).

AUTHOR IDENTIFIERS

Author ORCIDs:
Agunloye E: 0000-0002-1076-8385
Gavriilidis A: 0000-0003-3508-5043
Galvanin F: 0000-0002-3296-1581

REFERENCES

1. Galvanin F, Hartman RL, Kulkarni AA, Nieves-Remacha MJ. Introduction to the themed collection on digitalization in reaction engineering. React. Chem. Eng. 7:792-794 (2022). https://doi.org/10.1039/d2re90011d
2. Peterson L, Gosea IV, Benner P, Sundmacher K. Digital twins in process engineering: an overview on computational and numerical methods. Computers & Chemical Engineering

193:108917 (2025). https://doi.org/10.1016/j.compchemeng.2024.108917
3. Kasilingam S, Yang R, Singh SK, Farahani MA, Rai R, Wuest T. Physics-based and data-driven hybrid modeling in manufacturing: a review. Production & Manufacturing Research 12: (2024). https://doi.org/10.1080/21693277.2024.2305358
4. Destro F, Barolo M. A review on the modernization of pharmaceutical development and manufacturing – trends, perspectives, and the role of mathematical modeling. International Journal of Pharmaceutics 620:121715 (2022). https://doi.org/10.1016/j.ijpharm.2022.121715
5. Chen Y, Ierapetritou M. A framework of hybrid model development with identification of plant?model mismatch. AIChE Journal 66: (2020). https://doi.org/10.1002/aic.16996
6. Narayanan H, Luna M, Sokolov M, Arosio P, Butté A, Morbidelli M. Hybrid models based on machine learning and an increasing degree of process knowledge: application to capture chromatographic step. Ind. Eng. Chem. Res. 60:10466-10478 (2021). https://doi.org/10.1021/acs.iecr.1c01317
7. Geremia M, Marella T, Ardalani E, Beyramysoltan S, Chattoraj S, Facco P, Barolo M, Bezzo F. On the impact of mechanistic model quality and data availability in hybrid model development. Computers & Chemical Engineering 206:109536 (2026). https://doi.org/10.1016/j.compchemeng.2025.109536
8. Bird RB, Stewart WE and Lightfoot EN, Transport Phenomena (Second ed.). *John Wiley & Sons* (2001)
9. Quaglio M, Roberts L, Bin Jaapar MS, Fraga ES, Dua V, Galvanin F. An artificial neural network approach to recognise kinetic models from experimental data. Computers & Chemical Engineering 135:106759 (2020). https://doi.org/10.1016/j.compchemeng.2020.106759
10. Sangoi E, Quaglio M, Bezzo F, Galvanin F. An optimal experimental design framework for fast kinetic model identification based on artificial neural networks. Computers & Chemical Engineering 187:108752 (2024). https://doi.org/10.1016/j.compchemeng.2024.108752
11. Agunloye E, Gavriilidis A, Galvanin F. Hybrid modelling framework for reactor model discovery using artificial neural networks classifiers. SUSTENS 2025 :11 (2025). https://doi.org/10.3390/proceedings2025121011
12. Heaton J. Ian goodfellow, yoshua bengio, and aaron courville: deep learning. Genet Program Evolvable Mach 19:305-307 (2017). https://doi.org/10.1007/s10710-017-9314-z
13. Waldron C, Pankajakshan A, Quaglio M, Cao E, Galvanin F, Gavriilidis A. Closed-loop model-based design of experiments for kinetic model discrimination and parameter estimation: benzoic acid esterification on a heterogeneous catalyst. Ind. Eng. Chem. Res. 58:22165-22177 (2019). https://doi.org/10.1021/acs.iecr.9b04089
14. Fogler HS, Elements of Chemical Reaction Engineering, Prentice -Hall, Upper Saddle (2004) New Jersey
15. Sandler SI, Chemical, Biochemical, and Engineering Thermodynamics, fourth ed., John Wiley and Sons (2006) New Jersey.
16. Du X, Xu H, Zhu F. Understanding the effect of hyperparameter optimization on machine learning models for structure design problems. Computer-Aided Design 135:103013 (2021). https://doi.org/10.1016/j.cad.2021.103013
17. Sun L. Automated neural network architecture design and hyperparameter adjustment integrating distribution-based optimization strategy. ACE 147:181-186 (2025). https://doi.org/10.54254/2755-2721/2025.22721

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

PREDICTING FLOW REGIMES IN A WIPED FILM EVAPORATOR USING THE VOLUME OF FLUID METHOD

Gonçalo V.L. Pardal[a*] and Fernando P. Bernardo[a]

[a] CERES, University of Coimbra, Department of Chemical Engineering, Coimbra, Portugal
* Corresponding Author: gonalopardal@uc.pt.

ABSTRACT

To produce complex injectable formulations, it is imperative that the contents of the formulation are not damaged during processing. Some formulations require the removal of excess solvent often through evaporation. Wiped Film Evaporators can be applied for these scenarios, since they can operate at relatively low temperatures with mass transfer being promoted through a large surface area. This is created by wiping the liquid against the inner heated wall of the equipment, with the proper operation requiring, as much as possible, a stable and continuous liquid film. In some cases, however, depending on operating conditions or physical properties of the materials, film ruptures and wetting/de-wetting dynamics are observed, and in more severe situations the liquid is dispersed in isolated globules and drops. These are undesirable situations of operation, where it is not possible to guarantee an optimal and homogenous heat and mass transfer. In this work, Computational Fluid Dynamic (CFD) simulations are used to predict a map relating operating conditions to these different flow regimes. From this fundamental knowledge, it is possible to infer operatory windows for a given task and ultimately develop a reliable model for process scale-up.

Keywords: Computational Fluid Dynamics, Wiped Film Evaporator, Pharmaceutical Formulations

INTRODUCTION

Downstream processing of complex liquid pharmaceutical formulations requires a careful choice of equipment and operating conditions to avoid product damage. Such is the case of lipid nanoparticles and polymer nanoparticles, often produced as dilute suspensions, from which excess solvent must be removed under operating conditions that must not degrade the drug-charged particles [1-3]. The equipment here considered is a Wiped Film Evaporator (WFE), able to promote solvent evaporation at relatively low temperatures and shear rates.

In a WFE (Figure 1), the liquid is fed to a rotating plate that ejects it against the stationary heated wall of the main vessel. The liquid is then spread along the wall in a downwards helical pathway through spinning vertical blades (wipers), whose tip may have different geometries. The WFE under study has three wipers with tilted grooves that wipe the liquid upwards thus forcing it to rotate more in its helical path and have a longer residence time within the equipment. During operation, the evaporated solvent leaves the equipment upwards, while the concentrated liquid flows downwards and is collected at the bottom.

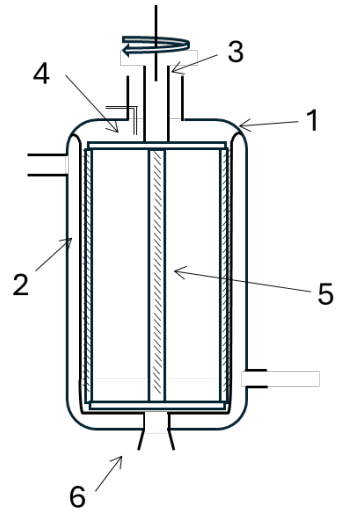

Figure 1. Schematic representation of a Wiped Film Evaporator. 1-main vessel; 2-heating jacket; 3-rotor; 4–feeding plate; 5- rotating wiper.

By adjusting inlet flowrate, temperature and/or rotating speed, one can in principle achieve an operating state with a reasonable evaporation rate, without exposing the mixture to excessive temperatures and wiping it gently enough, so that the suspended nanoparticles don't get damaged to the point of spilling their content. Achieving this desired state can be challenging. In some instances, other states apart from a stable wiped film are observed, such as the de-wetting of the inner wall or the formation of droplet streams instead of a constant film. These undesired flows are shown in Figure 2 and have already been reported in previous works regarding the study of WFE [4].

In this work, we present a CFD study aiming at identifying these different flow regimes using the Volume of Fluid model, well known by its ability to track complex gas-liquid interfaces. A simplified 2D version of the real flow is resolved, thus allowing to perform tens of simulations in the flowrate × rotation speed space.

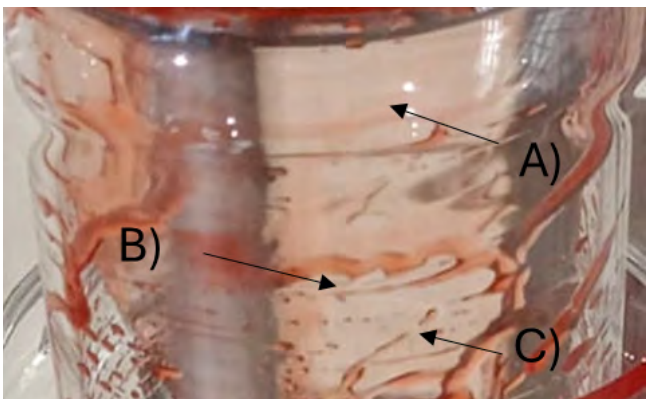

Figure 2. Image of our lab-scale WFE, highlighting the presence of 3 regimes: A–continuous film; B–film rupture into droplets; C–droplets.

CFD MODELLING

Model assumptions

The model used to simulate the film formation on the evaporator was a transient-state, laminar and multiphase model. To model the dynamic between liquid and gas, a Volume of Fluid (VoF) method was chosen. The VoF method is a commonly used technique that instead of solving momentum equations for each phase, it tracks the fraction of a phase in any given cell, in this case the volumetric fraction of liquid (α). The properties of the fluid are then calculated in each finite volume as a function of α, such as shown in Equation 1 for density. The sharp interface definition, which is of the highest interest in this case, is modeled through Equation 2, where F_σ represents the surface tension forces, k the curvature of the interface and σ the surface tension between both fluids. F_σ is added to the momentum balance and is what allows the VoF method to capture interface dynamics with great precision, given that the mesh chosen has enough finite volumes.

$$\rho = \alpha\rho_{liquid} + (1 - \alpha)\rho_{gas} \quad (1)$$

$$F_\sigma = \sigma k \nabla\alpha \quad (2)$$

Since the objective is to study the interface, a cross-section cut of the WFE was considered, thus the simulations carried out were 2D simulations. The cross section chosen is assumed to be around the middle of the equipment, where the fluid profiles should be well developed. This dimensional simplification drastically reduces computational time and can still give a good approximation of the regime in the WFE, although the lack of gravity is noticeable and will be discussed further in the next sections. The geometry used is present in Figure 3, along with the chosen boundaries.

During WFE operation, the rotating wipers can leave the rotor frame and move out, towards the inner wall, due to the centrifugal force [5]. Thus, since they are not fixed in any given position, to model the presence of the wipers, a porosity model was used. The porosity zone is presented in Figure 3 also. The chosen porosity model was the Darcy-Forchheimer equation, but only the viscous resistance was considered, since the Reynolds numbers of the simulations to study correspond mainly to laminar flows. The value used for viscous resistance coefficient was $4\times10^7\ m^{-2}$.

In this first approach, only the momentum transfer was considered, as for the time being it is only of interest to determine if the models can predict the scenarios of continuous film, film break or droplets. For the simulations considered, the liquid phase has the physical properties of water, while the gas was considered to have those of air. The model is solved with all velocities being relative to the spinning blades, as if blades were stationary and the inner wall was sliding (at the same speed of the blades and in the opposite direction). The relative velocities are reverted to absolute ones in the post processing of the simulation results.

In this 2D model it is not possible to directly manipulate the flow rates for the liquid entering the cavity, therefore the parameters to manipulate are the inlet velocity (U_{in}), inlet thickness of the fluid (ε), as well as the velocity of the sliding wall (U_{wall}). The values of U_{wall} are calculated based on the equivalent rotating wall velocity. To approximate to flowrate values, in the following sections, the product of ε and U_{in}, εU_{in}, will be considered for the global dynamic profile analysis.

Boundary Conditions

The inlet water velocity is a fixed value and although dependant on the scenario studied, its values are always related to those of U_{wall}. For scenario A the water enters the cavity fully wetting the sliding wall, with a parabolic

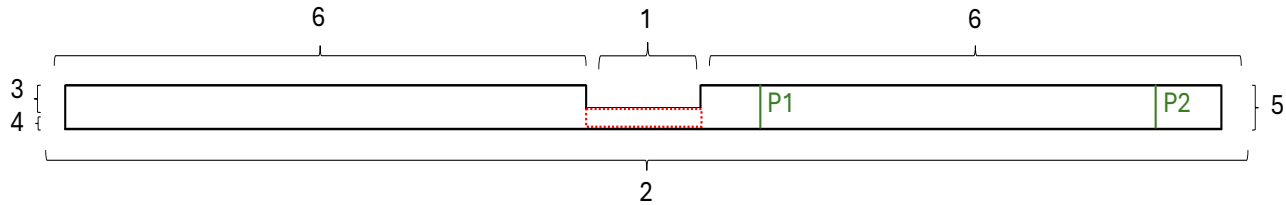

Figure 3. Cavity geometry and its boundaries. 1-resting blade; 2-sliding wall; 3-air inlet; 4-water inlet; 5-outlet; 6-back outlet. The dashed red line represents the porous media, and the green lines represent the P1 and P1 points.

velocity profile that increases the closer it is to the sliding wall, with U_{wall} being its maximum value. On the other hand, in scenario B it is assumed that the water enters diagonally relative to the sliding wall, approximately 2 *mm* above said wall, and with less kinetic energy than it. The velocity profile in this scenario B is a linear velocity, that is dependant on U_{wall}. Both scenarios were created to study: the water entering has an already established, well-defined profile for scenario A, while for scenario B studies the initial feed of liquid in the WFE, and how the difference of kinetic energy between feed and sliding wall may influence the flow pattern.

The air inlet and back outlet allow for air to leave but does not allow backflow. The outlet on the right side of the cavity has a similar boundary condition, allowing for matter to leave, but blocking backflows. The inner sliding wall has a fixed velocity, due to the exchange of reference with the rotating wiper. Thus, the wiper boundary has a no slip condition for velocity.

For the pressure, in the outlets the pressure was fixed as the atmospheric pressure, while the remaining boundaries were calculated depending on the flow on the cells near that boundary.

On the volumetric fraction side, it is important to note that for the inner sliding wall, a boundary condition that fixes the contact angle for water-air-glass was chosen, in order to better describe the interaction of the fluids with this wall, and to create more realistic and sharper droplets, in the cases where they appear.

For the face of the front and back of the cavity, a special boundary condition was chosen to not solve the momentum equation in the axial direction, thus making this a proper 2D simulation.

Meshing

The chosen mesh for this simulation corresponds to a rectangular cavity, with a height of 4 *mm*, and a width corresponding to a third of the perimeter of the evaporator, totaling 12 *cm*. The number of cells is around 33 000. From previous studies, it is imperative that in the gap between the resting blade and the inner wall, which is 2 *mm*, at least 16 cells are present, to capture the interface sharply [6]. In this study, a minimum of 28 cells were used to describe the gap, and for the cases where the feed height is smaller, more cells were used.

The simulations were conducted using the *openFOAM®* v2412 software, on a workstation equipped with an AMD Ryzen Threadripper™ PRO 7995WX, and 256 Gb of RAM. Post processing was carried out on the paraView software, along with using the Python™ pyVista package.

SIMULATION RESULTS AND DISCUSSION

As referred above, two scenarios where studied: A – a scenario where the inlet flow is parallel to the sliding wall, the inlet velocity was kept constant and ε varied between a defined range (Figure 4A and 4B); B – another one where ε has a fixed value, but the inlet velocity value varies, and its position is so that water enters diagonally to the sliding wall (Figure 4C). For scenario A, a range of U_{wall} and ε used were as follows: U_{wall} = [0.48; 0.64; 0.80; 0.96] *m/s*, and ε = [0.25; 0.5; 1.0 1.5; 2.0] *mm*, while for scenario B the range of U_{wall} used was the same as A, while the inlet thickness was a fixed value of ε = [0.25] *mm*. Furthermore, in scenario B the inlet velocity U_{in} varied based on the sliding mesh linear velocity by factors of [25; 50; 75; 100] %. The tests were conducted until the liquid thickness at positions P1 and P2 (see Figure 3), 2 *cm* after reaching the porous media and 1 *cm* before the main outlet, were considered constant, as well as analyzing the match between inlet and outlets flows. If both analyzes are constant, it is assumed that the simulation has achieved a steady state. The simulation time required to achieve a steady state varies according to the values of U_{wall} and ε. For the cases where droplets formed, the thickness was not evaluated, but rather the periodicity of the droplet formation. For cases where full film was formed, the convergence of the thickness analysed at P1 and P2 varies for each simulation condition, being as low as 0.17s (Scenario A, U_{wall} = 0.96 *m/s*; ε = 0.25 *mm*), and as high as 0.52s (Scenario A, U_{wall} = 0.64 *m/s*; ε = 0.25 *mm*), for example.

The CPU Clocktime for each simulation varied according to the value of ε, since smaller values of ε required finer meshes. The range of clocktimes was between [1650;2500] *s* per each second of results obtained.

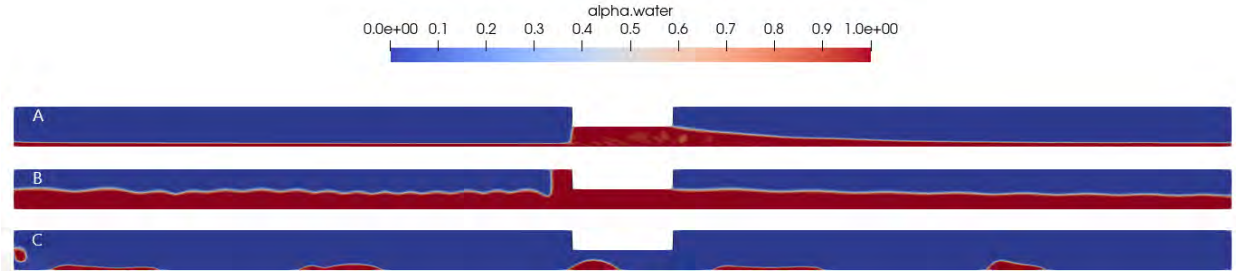

Figure 4. 2D representation of the fields of water volumetric fraction, showing examples of the different types of liquid film profiles identified during the simulations. A – Fully formed film (Scenario A, U_{wall} = 0.64 m/s, ε = 0.5 mm); B – Film with overflow (Scenario B, U_{wall} = 0.48 m/s, ε = 2.0 mm); C – Droplets (Scenario B, U_{wall} = 0.64 m/s, U_{in} = 0.32 m/s).

In each simulation the wetting behaviour of the liquid in the inner wall was analysed, to understand if it would fit one of the three categories at the main outlet: droplets, intermediate state from droplets to film, and finally a fully formed film. The simulations for scenario A showed full films and a different category, which was film overflowing form the cavity, while for scenario B the results varied between droplets and full film. An example of each of those cases are presented in Figure 4. To note that the characteristic "bow-wave" is formed in front of the blade, are visible in in those cases as well. The global analysis of the dynamic profile for the batch of simulations conducted is organized in the plots of Figures 5 and 6. Based on those results, it possible to identify some patterns for both scenarios. For scenario A, in the cases where εU_{in} are higher seem to lead more to the formation of liquid overflow, while for the cases where εU_{in} is smaller tendentiously lead to clear thin formed films. These phenomena can be interpreted as: the thinner films suffer less disruption from the porous media, meaning that they are less wiped by the rotating blade. Antithetical, the thicker films offer a larger liquid mass to be wiped by the blade, creating bigger bow waves in the front of the blade, and making so more liquid accumulates in that bow wave. This conclusion can also be corroborated by the α profiles shown in Figure 4. To overcome this, it is necessary to augment the values of U_{wall} for a given εU_{in}, similarly to what happened for cases like U_{wall} = [0.80; 0.96] *m/s*, and $\varepsilon U_{in} \approx$ [0.75:1] m^2/s, for example.

For Scenario B, meant to replicate the first passage of the liquid in the WFE, it is possible to see that for lower values of U_{in}, meaning lower inlet flow, it is not possible to develop a clear film right from the start. The droplets turn to complete film by augmenting the inlet flow, and it is possible to analyze that lower sliding wall velocities, like U_{wall} = 0.48 *m/s*, can not produce complete films from this diagonal feed but instead will produce pulses of droplets. This can be interpreted as follows: since the inlet liquid lacks kinetic energy compared to the scenario where it enters already wetting the wall and with a well-defined velocity profile, it is important to have a sliding wall that can provide enough kinetic energy to fully spread the liquid so that it forms a film.

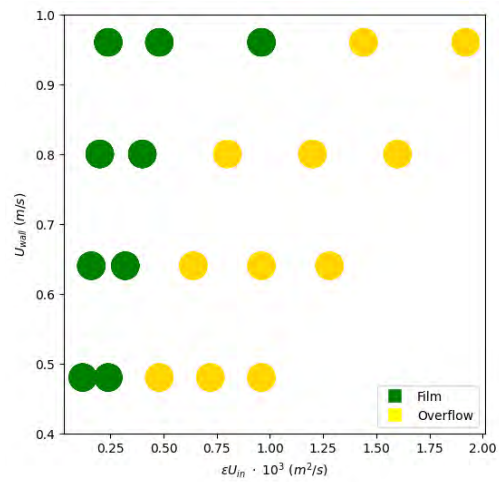

Figure 5. Flow regime map for the simulations of Scenario A.

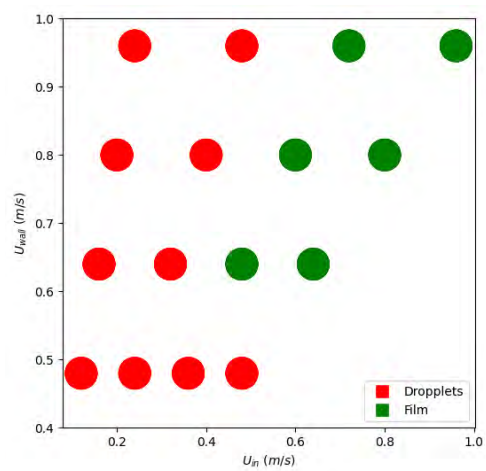

Figure 6. Flow regime map for the simulations of Scenario B.

Despite the promising results shown for the higher values of U_{wall}, working in a range of [0.80:0.96] *m/s*, or even above, may not be of interest to treat more sensitive pharmaceutical formulations. Furthermore, it is important to note that, the smaller the value of U_{wall}, the more noticeable the effects of gravity should be felt by the liquid in the real evaporator, since the gravitational force pulling downwards needs to be overcome by the centrifugal force the wiper combs spreading the fluid continuously. Without gravity, it is possible to identify the formation of droplets with considerable thickness and width, and the amount of liquid pushed in the bow wave also increased due to the lack of gravitational force, enabling it to accumulate. These phenomena can be seen on Figure 4B and 4C and should not be discard when analyzing the results. The larger droplets appear mostly in cases were the difference in kinetic velocity between U_{wall} and U_{in} is higher, and these cases were categorized as droplets. It was known that the initial assumptions would limit the model capability of replicating every aspect of the original equipment, nonetheless the computational speed reduction gained, as well as the sharp interface description demonstrated that there is value in these more elementary simulations.

To be able compare these results to a lab-scale evaporator, an auxiliary technology to determine film thickness inside the WFE would be necessary, such a high-speed image capturing equipment.

CONCLUSIONS AND FUTURE WORKS

From these simpler cases solved, it is possible to determine that the VoF method can describe the phenomena registered when operating a WFE for lower ranges of sliding wall velocity, initial film thickness and lower initial kinetic energy. The porosity model proved to be capable of retard the fluid movement enough to generate regimes that differ from the simple continuous film, and even creating the classic bow wave in front of the blade. The frame of reference exchange also proved viable, meaning that in future simulations it may be a useful alternative to the more conventional modelling approaches.

A range of operatory sliding wall velocities was also defined as best to be avoided, since there was fluid rupture and droplet formation for those cases. In spite of the simplifications used in the modeling step, this work proved that it is still possible to attain valuable information from this 2D simulations, and the creation of this first map of flow regimes will be most useful for future works, be it fully developed 3D models, or even for equipment scale-up.

Using this work as a base, the next step is to naturally transition to a more complex 3D model of the equipment, where gravity will play a crucial role in describing the dynamic behaviour of the film. It would also be an interesting complementary work to acquire high speed imaging inside a WFE to better identify the cases of film rupture, to then calibrate the 3D model. It would also be of interest to refine the porosity model or verify experimentally if the parameters can be optimized to match the effect that the wipers have on the fluid, be it interface definition or residence time.

ACKNOWLEDGEMENTS

The authors have the pleasure of acknowledging the support from “CiNTech - Technological Hub for Innovation, Translation and Industrialization of Complex Injectables Drugs” (reference:7131), in the scope of PRR – Recovery and Resilience Plan and by the Next Generation EU European Funds. Following NOTICE No.02/C05-i01/2022, Component 5 – Capitalization and Business Innovation – Mobilizing Agendas for Business Innovation. Also, the authors would like to extend their acknowledgment to the CERES - Chemical Engineering and Renewable Resources for Sustainability research center, for all the support provided.

AUTHOR IDENTIFIERS

Author ORCIDs:
Pardal, G.V.L.: 0000-0003-2764-4052
Bernardo, F.P.: 0000-0001-7339-0706

REFERENCES

1. Do JH, An J, Joun YS, Chung DJ, Kim JH. Cellular-uptake behavior of polymer nanoparticles into consideration of biosafety. Macromol. Res. 16:695-703 (2008). https://doi.org/10.1007/bf03218583
2. Xu Q, Hashimoto M, Dang TT, Hoare T, Kohane DS, Whitesides GM, Langer R, Anderson DG. Preparation of monodisperse biodegradable polymer microparticles using a microfluidic flow?focusing device for controlled drug delivery. Small 5:1575-1581 (2009). https://doi.org/10.1002/smll.200801855
3. John R, Monpara J, Swaminathan S, Kalhapure R. Chemistry and art of developing lipid nanoparticles for biologics delivery: focus on development and scale-up. Pharmaceutics 16:131 (2024). https://doi.org/10.3390/pharmaceutics16010131
4. Appelhaus D, Jasch K, Meyer H, Scholl S. Evaporation effects on residence time in wiped film evaporators with roller wipers. Chemical Engineering Research and Design 218:341-349 (2025). https://doi.org/10.1016/j.cherd.2025.04.036
5. Cvengroš J, Pollák Š, Micov M, Lutišan J. Film wiping in the molecular evaporator. Chemical

Engineering Journal 81:9-14 (2001). https://doi.org/10.1016/s1385-8947(00)00195-9
6. Pardal G, Bernardo FP, Multiphase CFD modelling of a Wiped Film Evaporator, oral presentation at ECCE/ECAB/CIBIQ 2025, 8-10 sep 2025, Lisbon.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Modelling of fouling dynamics in a falling-film evaporator

Johanne L. Christensen[a], Lukas S. Theisen[b] and Kevin Feldmann[b] and Jakob K. Huusom[a*]

[a] Department of Chemical and Biochemical Engineering, Technical University of Denmark, Kgs. Lyngby, Denmark
[b] GEA Process Engineering A/S, Søborg, Denmark
* Corresponding Author: jkh@kt.dtu.dk

ABSTRACT

Fouling is a persistent issue in industrial heat-transfer equipment, increasing energy demand and reducing efficiency. This is also true in the dairy industry where falling-film evaporators are central to powder production. Most dynamic models, however, neglect gradual fouling, limiting predictive accuracy during extended operation. As a result, model-based control can become unreliable when fouling becomes significant. The dynamic models by Bojnourd et al. [1] are widely used but assume clean-surface operation. While this captures short-term thermal behavior, it cannot represent the progressive decline in heat-transfer performance caused by fouling. Díaz-Ovalle et al. [2] introduced a fouling-layer model that explicitly describes the growth of a fouling deposit over time. Building on this concept, the present work incorporates a simple dynamic fouling model for falling-film evaporators and validates it using industrial data from a four-effect evaporator using thermal vapor recompression. Two fouling-modelling strategies are evaluated: a mechanistic fouling layer model and an empirical model describing the observed loss in heat-transfer efficiency. The fouling layer model substantially reduces prediction error, demonstrating that it captures the dominant fouling behavior. Although the empirical model fits all observed discrepancies, it lacks predictive capability and therefore cannot be applied for model-based prediction or control.

Keywords: Fouling dynamics, Falling-film evaporator, Dynamic modelling, Dairy industry

1. INTRODUCTION

The dairy industry uses drying to extend product shelf life and reduce transportation costs [3]. Dairy powder production relies on two main unit operations: falling-film evaporation and spray drying [4]. Spray drying is the most energy-intensive step, followed by falling-film evaporation. To lower overall operating expenditures, a falling-film evaporator is typically used beforehand to remove part of the water content [4]. In a falling-film evaporator, the liquid feed flows as a thin film through multiple vertical tubes, with hot steam applied on the other side. Water evaporates from the film, and a concentrated liquid exits the tubes, as illustrated in Figure 1. Optimizing the overall drying process requires maximizing the solids content at the evaporator outlet. Higher concentrations, however, increase the risk of fouling, a well-known challenge in evaporation systems. Fouling refers to the unwanted accumulation of deposits in process equipment, typically consisting of minerals, organic matter, and biological growth [5], as illustrated in Figure 2. Fouling reduces heat-transfer efficiency, which increases energy consumption and cleaning costs [2]. In thermal treatment of dairy products, fouling is caused by multiple interacting reactions, making the mechanisms difficult to capture [7]. The main reactions are protein agglomeration and salt precipitation. The initial protein agglomeration forms a monolayer on the heating surface in a slow "induction period," during which heat transfer is not significantly affected [8]. After the induction period, the fouling layer grows more rapidly on the protein layer surface.

This study uses industrial data from a four-effect evaporator with thermal vapor recompression (TVR) to develop and validate a simple model of fouling behavior. Validation includes estimating the fouling parameters incorporated in the model.

2. METHODOLOGY

2.1 Industrial case study

The available data for the 4-effect TVR evaporator is summarized in Table 1. Measurements are

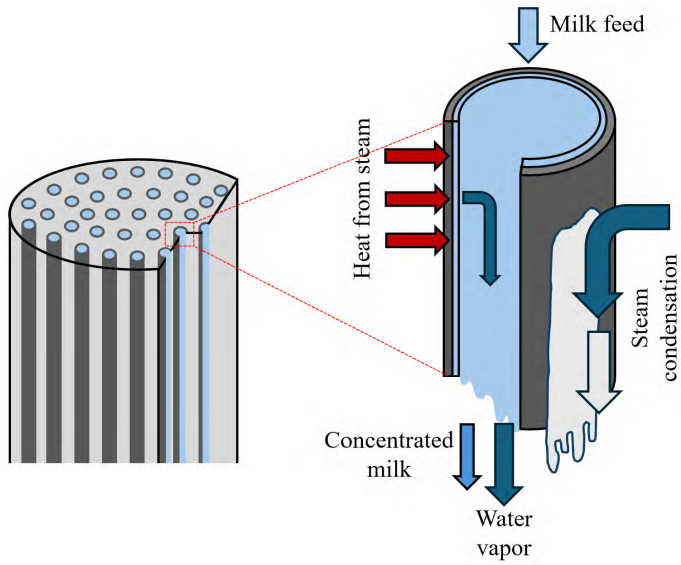

Figure 1. Basic principles of a falling-film evaporator.

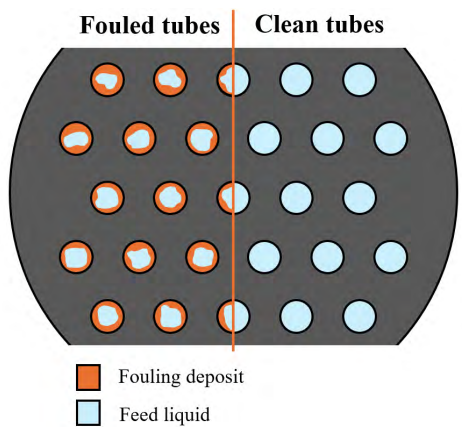

Figure 2. Illustration of fouling deposits in evaporator tubes.

recorded every 60 seconds. While the system's power consumption is also monitored, it is not required for simulation because both vapor and steam temperatures are measured, which are sufficient to model energy balances. Table 1 lists the measured variables, their descriptions, and units for both the industrial data and the simulation.

Table 1: Available data for a 4-effect TVR evaporator with data collected every 60 sec.

Symbol	Description	Unit
F_{in}	Feed flow	kg/h
x_{in}	Feed mass fraction	w/w%
T_{steam}	Steam temperature	°C
T_{vapor}	Vapor temperature	°C
F_{out}	Outlet flow	kg/h
x_{out}	Outlet mass fraction	w/w%

To ensure mass balance consistency, the industrial data are reconciled by accounting for a systematic bias in the feed concentration, x_{in}, and reducing measurement noise using a weighted least-squares method:

$$\min_{\hat{z}} \sum_{n=1}^{N} \sum_{k=0}^{K} \left(\frac{\hat{z}_{n,k} - z_{n,k}}{\sigma_n} \right)^2 \quad (1)$$

Subject to the mass balance constraint:

$$\hat{F}_{in} \cdot \hat{x}_{in} - \hat{F}_{out} \cdot \hat{x}_{out} = 0 \quad (2)$$

where $z = [F_{in}, x_{in}, F_{out}, x_{out}]$ are measured values, $\hat{z} = [\hat{F}_{in}, \hat{x}_{in}, \hat{F}_{out}, \hat{x}_{out}]$ are reconciled values and σ_n is the standard deviation for each variable in a steady state region. N and K are the number of variables and measurements, respectively. This reconciliation ensures a consistent mass balance even with sensor calibration errors or other systematic biases.

2.2 MATHEMATICAL MODELS

The evaporator dynamics are modeled using lumped models in both radial and axial directions, following Bojnourd et al. [1]. These models fit industrial data reasonably well when fouling is negligible. The mathematical description is divided into a base model, assuming no fouling, and a fouling layer model, incorporating fouling effects based on literature [2].

Base Model

The base model is governed by the following ODEs:

$$\frac{dF_{out}}{dt} = \frac{1}{\tau(1-\alpha)}(F_{in} - F_{out} - F_e) - \frac{\alpha}{1-\alpha}\frac{dF_{in}}{dt} \quad (3)$$

$$\frac{dx_{out}}{dt} = \frac{F_{in}(x_{in}-x_{av}) - F_{out}(x_{out}-x_{av}) + F_e x_{av}}{\tau(1-\alpha)F_{av}} - \frac{\alpha}{1-\alpha}\frac{dx_{in}}{dt} \quad (4)$$

Here, F_{out}, F_{in} and F_e are the outlet, feed and evaporated mass flows respectively. x_{in} and x_{out} are inlet and outlet solid mass fraction. t is time, τ is residence time and α is an averaging parameter used to define x_{av} and F_{av} in Equations (5) and (6).

$$F_{av} = \alpha F_{in} + (1-\alpha)F_{out} \quad (5)$$

$$x_{av} = \alpha x_{in} + (1-\alpha)x_{out} \quad (6)$$

The evaporated flow, F_e, is defined through the heat transfer from the steam to the liquid as shown by equation (7).

$$F_e = \frac{U \cdot A \cdot (T_{steam} - T_{vapor})}{H_{vap}(T_{vapor})} \quad (7)$$

Where U is the overall heat transfer coefficient across the tube, A is the surface area of the tubes, T_{steam} and T_{vapor} are the temperatures on the steam side of the tube and inside the tubes respectively. The heat of vaporization, $H_{vap}(T)$, is based on pure water and described using a fourth order polynomial based on experimental data.

The unknowns in the base model are U, A, τ. The value $U \cdot A$ can be estimated using a pseudo steady-state approximation over the first 20 minutes of operation:

$$UA_{base} = \text{mean}\left(\frac{H_{vap}(T_{vapor}[1:20]) \cdot (F_{in}[1:20] - F_{out}[1:20])}{T_{steam}[1:20] - T_{vapor}[1:20]}\right) \quad (8)$$

When no fouling is present this value is expected to fit the whole dataset. The final unknown parameter, the residence time τ, is theoretically expected to vary with changes in the flow within the system; however, it was observed through fitting τ that a value of 60 s was reasonable for all mass flows in the data. This constant value is used in all simulations.

Empirical Fouling Model

When fouling occurs, the heat transfer decreases because the fouling layer has lower thermal conductivity than the tube material. A pseudo steady-state approximation is used over 5-minute intervals to estimate the effective heat transfer coefficient. This is done to reduce noise.

$$UA = \frac{H_{vap}(T_{vapor}) \cdot (F_{in} - F_{out})}{T_{steam} - T_{vapor}} \quad (9)$$

This approach helps to quantify the reduction in heat transfer due to fouling.

Fouling Layer Model

A detailed fouling mechanism involves multiple reactions. The fouling layer model is based on the work by Díaz-Ovalle et al. [2], however the model is simplified heavily due to the many unknown parameters. The model is based on the fouling layer thickness, δ, growing over time, resulting in a decreased heat transfer coefficient.

$$\frac{d\delta}{dt} = \frac{k_d k_f}{U_C Re}\exp\left(-\frac{E}{R} \cdot \frac{k_f + \delta U_C}{k_f T_{steam} + T_{vapor} U_C}\right) - k_r \delta \quad (10)$$

Here the first term describes a deposition rate, dependent on the fouling parameters, flow regime and the temperatures in the system. The second term describes a removal rate, only dependent on the current fouling thickness and the parameter k_r, which is the rate of fouling layer removal. k_d describes the deposition rate, with k_f being the thermal conductivity of the fouling layer, Re is the Reynold number of the system, E is the activation energy of the overall fouling reaction and R is the universal gas constant. The corrected heat transfer coefficient, U_C, is defined in Equation (11).

$$U_C = \left(U_0^{-1} + \frac{\delta}{k_f}\right)^{-1} \quad (11)$$

Where $U_0 = UA_{base}/A$ is the initial heat transfer coefficient where fouling is not present. Only the ratio δ/k_f appears in U_C, resulting in δ and k_f are not being uniquely identifiable and they may scale together unless one of their magnitudes is known. For solid proteins the thermal conductivity is between 0.1-0.3 W/(m K) [9, 10] which is assumed for the fouling layer. The ODEs for the flow and solid mass fractions are identical to Equations (3) and (4) with the evaporated mass flow now being described with U_C as shown by Equation (13).

$$F_e = \frac{U_C \cdot A \cdot (T_{steam} - T_{vapor})}{H_{vap}(T_{vapor})} \quad (13)$$

Initial simulations indicated that, within the fouling layer model, both the temperatures and the corrected heat transfer coefficient U_C could be treated as constant, since their variation had a negligible effect on the results. The fouling layer model is simplified and defined through three constant parameters $K_{\text{pre}}, K_{\text{exp}}$, and k_r.

$$\frac{d\delta}{dt} = K_{pre} \cdot \exp(-K_{exp} \cdot \delta) - k_r \delta \quad (14)$$

$$K_{pre} = \frac{k_d k_f}{U_C Re} \cdot \exp\left(-\frac{E k_f}{R(k_f T_{steam} + T_{vapor} U_C)}\right) \quad (15)$$

$$K_{exp} = \frac{E \cdot U_C}{R(k_f T_{steam} + T_{vapor} U_C)} \quad (16)$$

The output data are used to determine these fouling parameters, along with the surface area, which is also unknown for the specific evaporator configuration. The fouling model is applied only when fouling is actively reducing heat transfer. Therefore, the simulation is divided into two segments: an induction period, during which the base model is used, and the post-induction period, where the fouling model is applied.

The induction period is determined by calculating the needed UA from Equation (9) and determining when the heat transfer decreases by 1% from the baseline in Equation (8). To ensure robustness, this decrease must persist for at least five consecutive measurements before the induction time is assigned.

3. SIMULATION RESULTS AND DISCUSSION

The initial conditions for the outlet flow and solid mass fraction are defined based on the first measurement. In all simulations, a value of $\alpha = 0.5$ is used, however, the exact value was found to have little influence on the simulation results [1]. Other constant parameters include a residence time of $\tau = 60\,\mathrm{s}$, and a thermal conductivity of the fouling layer of $k_f = 0.2\,\frac{\mathrm{W}}{\mathrm{m\,K}}$. All results are presented as normalized values, normalized to the maximum value of measured data in each subplot.

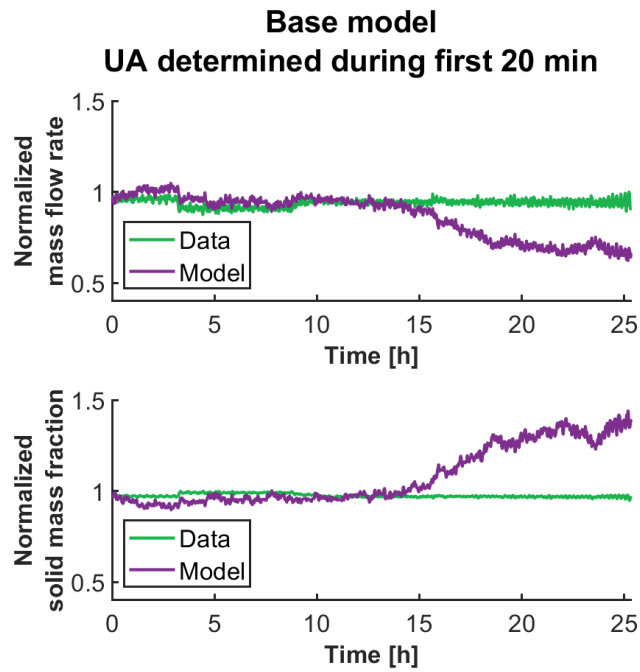

Figure 3. Simulation using base model.

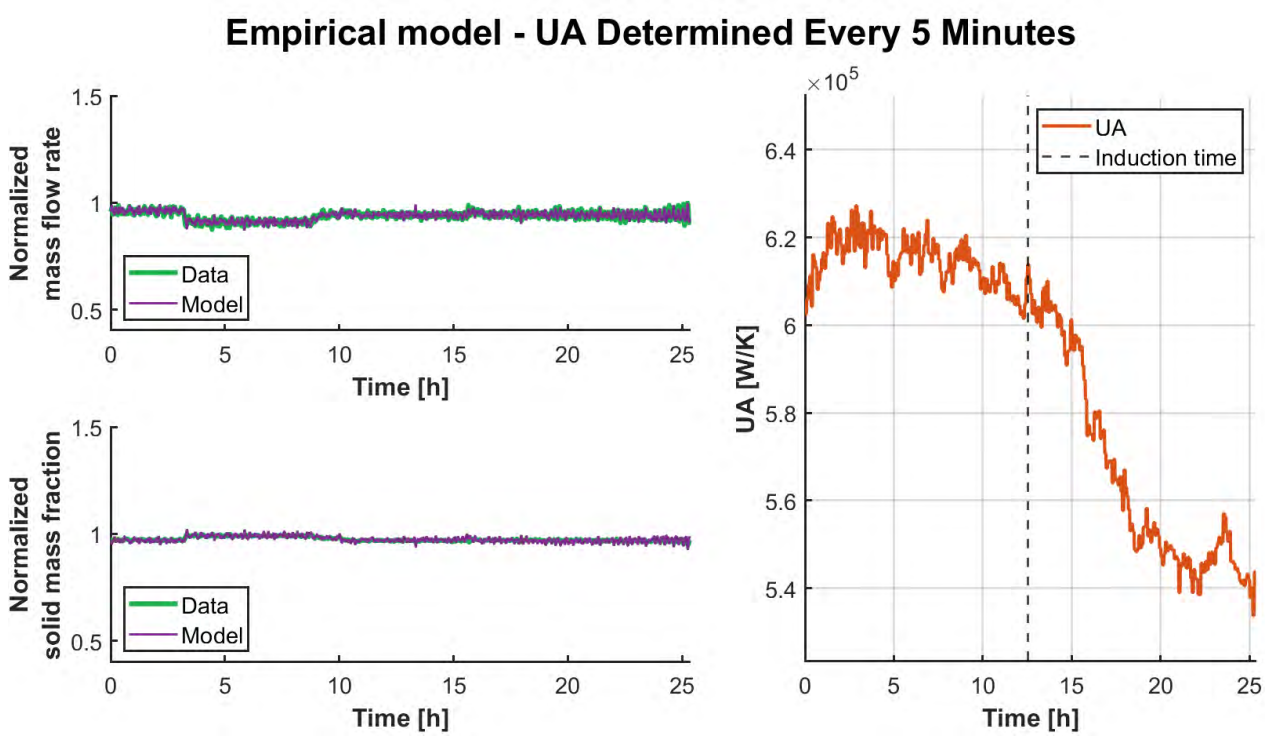

Figure 4. Simulation using empirical model, correcting the heat transfer coefficient UA.

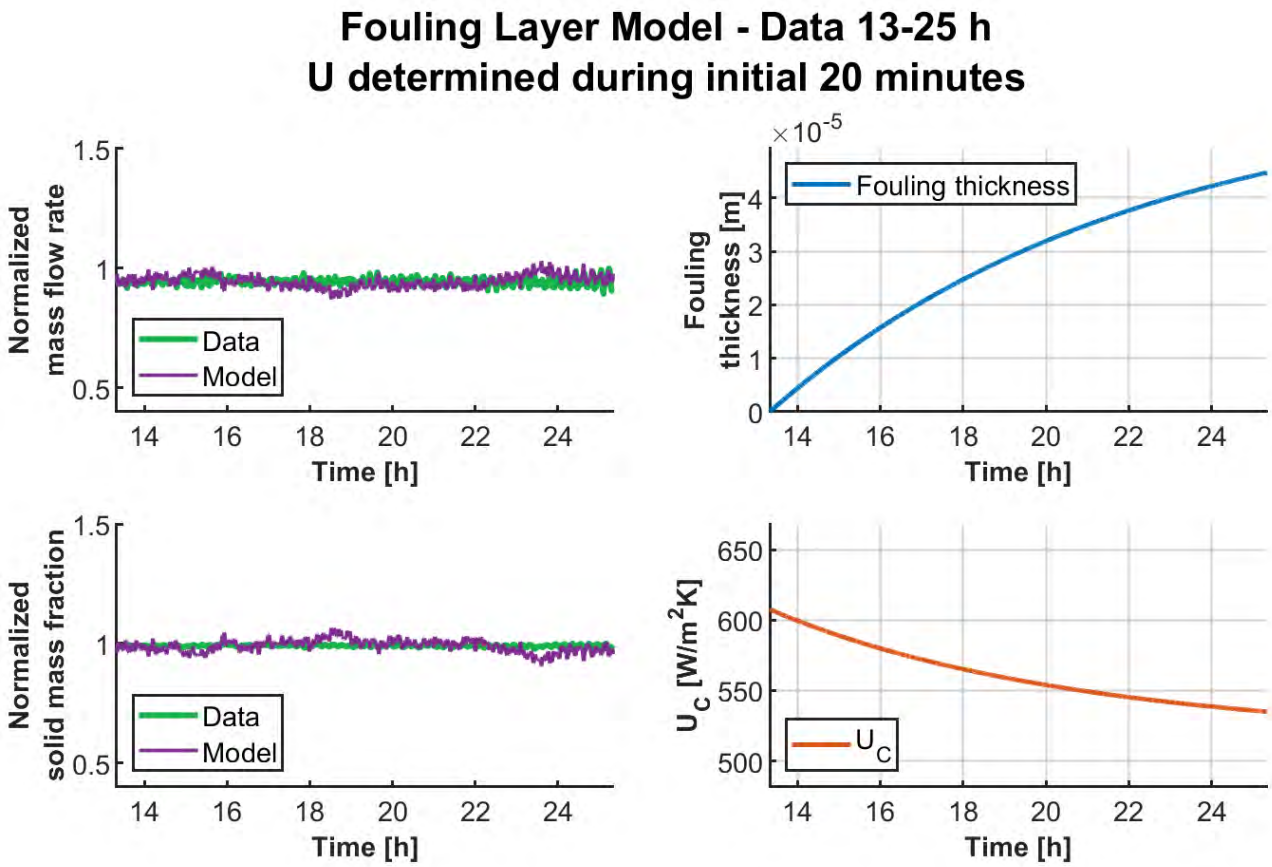

Figure 5. Simulation using fouling layer model on fouling data after the induction period.

The initial simulations using the base model, shown in Figure 3, illustrate the model's limited predictive capability. During the induction period, $t = [0, 13]$ h, the base model provides a reasonable fit. After approximately 12 hours, the model predictions increasingly deviate from the measured data. The empirical approach presented in Figure 4 shows a relatively constant overall heat transfer coefficient (UA) up to 12 hours, after which a noticeable decrease in heat transfer occurs. This decline correlates with the discrepancies observed in the base model predictions.

Table 2: 95% CI of the fouling parameters estimated for hour 13-25.

Parameter	Estimated value: 95% CI	Unit
A	1000 ± 294	m^2
K_{pre}	$[1.90 \pm 0.56] \cdot 10^{-9}$	m/s
K_{exp}	$[5.14 \pm 2.02] \cdot 10^{-2}$	1/m
k_r	$[3.16 \pm 0.27] \cdot 10^{-5}$	1/s

Table 3: Fouling parameters estimated for 8 runs, with the surface area fixed at 500 m^2.

Parameter	Estimated value: 95% CI	Unit
K_{pre}	$9.55\ (+0.46/-0.44) \cdot 10^{-10}$	m/s
K_{exp}	$4.12\ (+1.0/-0.81) \cdot 10^{4}$	1/m
k_r	$1.02\ (+17/-0.96) \cdot 10^{-6}$	1/s

The data from hours 13 to 25 are simulated in Figure 5 using the fouling model. The unknown parameter vector $p = [A, K_{pre}, K_{exp}, k_r]$ is fitted to the output data using a minimizing function. The 95% confidence intervals (CI) of the optimal parameters are presented in Table 2.The simulation results shown in Figure 5 indicate that the simplified fouling layer model, employing constant parameters, fits the experimental data well. The parameter estimation was performed using a statistical approach, yielding a cross-correlation matrix shown in Figure 6. This analysis reveals a strong correlation between the surface area A and the parameters K_{pre}.

Furthermore, the simulations demonstrated that different initial conditions could lead to different parameter sets while producing the same total error. To mitigate the

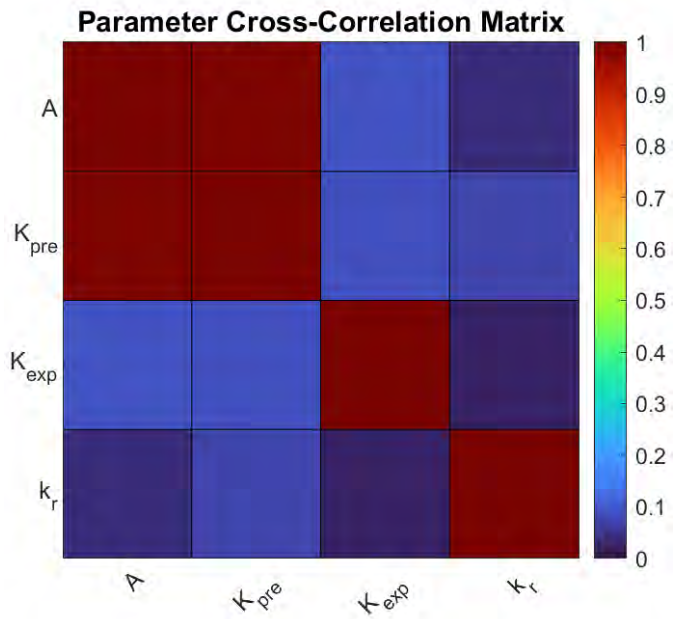

Figure 6. Parameter cross correlation matrix.

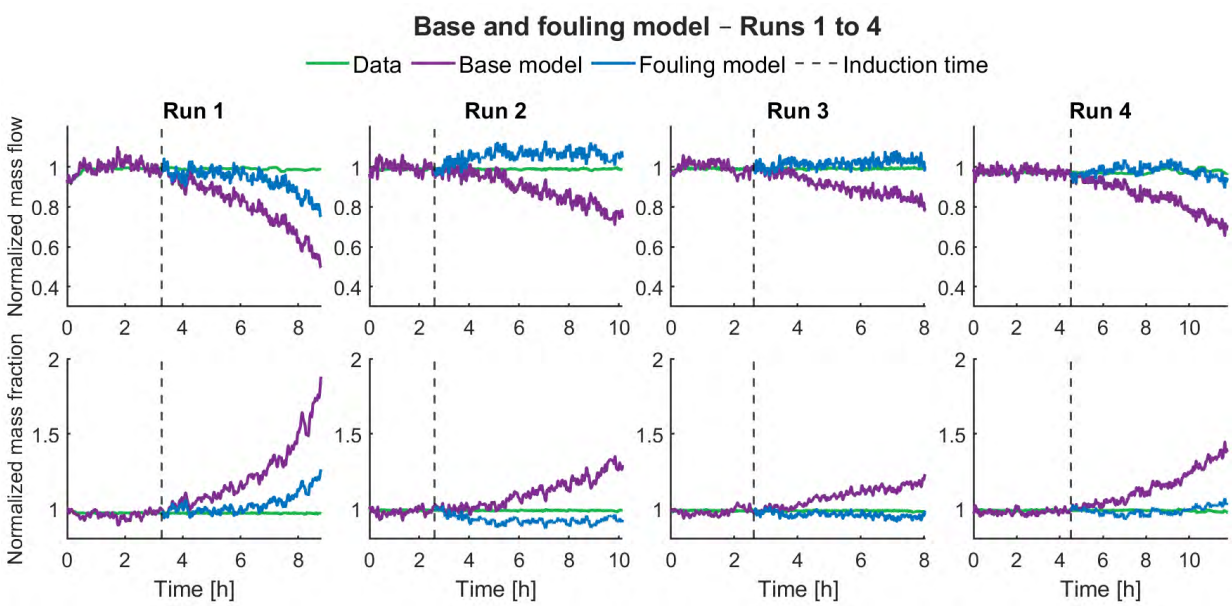

Figure 7. Simulation of run 1-4 using the optimal parameters in Table 3.

issue of multiple parameter solutions, the surface area is fixed at $A = 500\ \mathrm{m}^2$ for an extended parameter estimation.

A more global parameter estimation is performed using data from eight runs with the same feed type, each consisting of approximately 600 operational measurements. For all runs, the error between the model predictions and the measured data is included in a global optimization procedure. The data from the eight runs are distributed seasonally throughout the year and correspond to different durations of initial cleaning. This approach is adopted to evaluate the robustness of the estimated fouling parameters. The resulting parameter values are presented in Table 3, where the confidence intervals are asymmetric because the parameters were estimated on a logarithmic scale.

These parameters are applied in the simulation shown in Figure 7, where the fouling model becomes active after the induction period. The maximum relative error and the average relative error are calculated for all eight runs and are presented in Figure 8. The results clearly demonstrate that the use of the fouling model significantly reduces the prediction errors compared to the

base model, with the exception of run 5, for which both models perform equally well. using the fouling model compared with the base model. Runs 6-8 show behavior very similar to runs 1-4 and are therefore omitted.

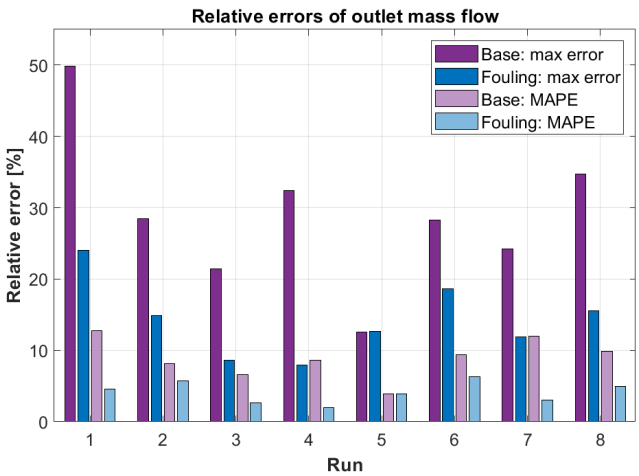

Figure 8. The maximum and average relative errors obtained using the fouling model compared with the base model.

The parameters estimated from individual runs generally exhibit substantial variability, particularly for the parameter K_{exp} which can span several orders of magnitude. In contrast, the value of K_{pre} and k_r show less variation. This variability in the fouling parameters may be attributed to changes in feed composition due to seasonal variations and/or differences in the operating conditions of the evaporator resulting from variations in cleaning procedures. The induction time was not observed to depend on the duration of the cleaning process.

FUTURE APPLICATION

The primary motivation for employing a simple model is to facilitate its use in process control of evaporator systems, where fouling significantly influences system behavior. Model predictive control (MPC) requires models that can reliably predict future system dynamics. A more accurate fouling model could improve control performance by enabling faster and more appropriate system responses.

The model is sensitive to parameter variations, particularly when only limited measurements are available, which can lead to uncertainty and non-unique solutions during early fouling, reducing control robustness. To address this, model parameters for different milk types could be systematically mapped and analyzed under varying feed compositions. This approach may provide a deeper physical understanding of fouling behavior and allow control sequences to anticipate fouling layer growth without estimating parameters for each individual run.

CONCLUSION

It has been shown that incorporating a simple fouling model in the dynamic operation of a four-effect falling-film evaporator leads to a substantial improvement in overall model performance.

The fouling layer model was shown to significantly improve predictive accuracy, reducing model errors by a factor of two or more in most cases. Using data from eight operational runs with the surface area fixed at 500 m², the fouling parameters were estimated as $K_{pre} = 9.55 \cdot 10^{-10}$ m/s, $K_{exp} = 4.12 \cdot 10^{4}$ 1/m and $k_r = 1.02 \cdot 10^{-6}$ 1/s, with k_r having a rather large confidence interval.

Although the model is sensitive to parameter variations, and the estimated parameter values can vary by several orders of magnitude due to seasonal changes in feed composition and differences in cleaning conditions, the inclusion of fouling dynamics consistently results in a substantially better fit than a model that neglects fouling altogether. These results demonstrate that even a simplified, physically motivated fouling model can provide robust improvements in predictive performance, supporting its applicability for both simulation and control-oriented applications in evaporator systems.

AUTHOR IDENTIFIERS

Author ORCIDs:
Christensen JL: 0009-0007-8698-7000
Theisen LS: 0000-0002-8019-5826
Huusom JK: 0000-0002-1367-7235

REFERENCES

1. Medhat Bojnourd F, Fanaei MA, Zohreie H. Mathematical modelling and dynamic simulation of multi-effect falling-film evaporator for milk powder production. Mathematical and Computer Modelling of Dynamical Systems 21:336-358 (2014). https://doi.org/10.1080/13873954.2014.980276
2. Díaz-Ovalle CO, Rico-Ramírez V, Ramos-Ojeda E. Control of falling-film evaporators with modeling of fouling. Food and Bioproducts Processing 132:68-82 (2022). https://doi.org/10.1016/j.fbp.2021.12.010
3. LAL Global FZE. https://lalglobalfze.com/industrial-applications-of-milk-powder-essential-across-sectors/
4. Zhang Y, Munir MT, Udugama I, Yu W, Young BR. Modelling of a milk powder falling film evaporator for predicting process trends and comparison of energy consumption. Journal of Food Engineering 225:26-33 (2018). https://doi.org/10.1016/j.jfoodeng.2018.01.016

5. Elliott Tool Technologies. https://www.elliott-tool.com/common-types-of-tube-failure-fouling-corrosion/
6. Ben-Mansour R, El-Ferik S, Al-Naser M, Qureshi BA, Eltoum MAM, Abuelyamen A, Al-Sunni F, Ben Mansour R. Experimental/numerical investigation and prediction of fouling in multiphase flow heat exchangers: a review. Energies 16:2812 (2023). https://doi.org/10.3390/en16062812
7. Huppertz T, Nieuwenhuijse H. Constituent fouling during heat treatment of milk: a review. International Dairy Journal 126:105236 (2022). https://doi.org/10.1016/j.idairyj.2021.105236
8. Fryer PJ. The uses of fouling models in the design of food process plant*. Int J of Dairy Tech 42:23-29 (2007). https://doi.org/10.1111/j.1471-0307.1989.tb01703.x
9. Xue Y, Lofland S, Hu X. Thermal conductivity of protein-based materials: a review. Polymers 11:456 (2019). https://doi.org/10.3390/polym11030456
10. Yamato T, Wang T, Sugiura W, Laprévote O, Katagiri T. Computational study on the thermal conductivity of a protein. J. Phys. Chem. B 126:3029-3036 (2022). https://doi.org/10.1021/acs.jpcb.2c00958

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Modeling Slug Flow Dynamics in Offshore Wells using Universal Differential Equations

Gustavo Luís Rodrigues Caldas[ad*], Giovani Gerevini[c], Fábio C. Diehl[c], Idelfonso B. R. Nogueira[d], Maurício B.de Souza Jr.[ab], Argimiro R. Secchi[a]

[a] Chemical Engineering Program, PEQ/COPPE, Universidade Federal do Rio de Janeiro, Rio de Janeiro, 21941-972, Brazil
[b] EPQB, Chemistry School, Universidade Federal do Rio de Janeiro, Rio de Janeiro, 21941-909, Brazil
[c] CENPES/Petrobras, Rio de Janeiro, 21941-915, Brazil
[d] Chemical Engineering Department, Norwegian University of Science and Technology, N-7491, Trondheim, Norway
* Corresponding Author: gcaldas@peq.coppe.ufrj.br.

ABSTRACT

Slug flow in multiphase production systems is a critical challenge in the oil and gas industry, characterized by complex and oscillatory dynamics, e.g., limit cycles. First-principle (FP) models often employ physics simplification, such as a virtual valve for the slug formation. To capture complex physics poorly modeled by FP models, hybrid models combine data-driven techniques and physical knowledge, such as the architecture known as universal differential equation (UDE). This work aims to employ a hybrid model based on neural networks to enhance the modeling of multiphase oil production systems. In the UDE model, a neural network is embedded within the structure of the FP differential equations. To demonstrate the feasibility of the methodology, the model was trained based on synthetic data, employing parameters estimated from OLGA simulations. Since the system faces oscillatory behavior, we trained the UDE in two stages: the first one employs smooth collocation on data to obtain an estimation of the derivatives. The results show that the two-stage strategy was successful in addressing the resolution of an oscillatory system. In the second stage, the algorithm reached a scaled MSE of 0.009 in 502 iterations. For the testing phase, using 20%, 40%, 60%, and 80% valve openings, the scaled MSE reported was 0.088. For the 100% valve opening, outside of the training part, the MSE reported was 0.039, showing a good extrapolation capacity.

Keywords: FOWM, hybrid model, scientific machine learning, oil and gas, neural networks

INTRODUCTION

After early works of machine learning (ML) in the 1990s in process systems engineering (PSE) [1], ML in the PSE field has experienced a growing interest due to the greater computational power available, simultaneously with the development of novel algorithms. This legacy was further employed to enhance first-principles (FP) models, which, in general, employ simplifying assumptions in their formulation. Therefore, research dedicated to incorporating physics into ML has attracted the attention of the scientific and engineering community in recent years, under the name physics-informed machine learning (PIML) [2], a class of hybrid models.

The advantages of hybrid models include improving the interpretability of neural network architectures [3]. Furthermore, ML can update the FP model based on new data. For instance, providing new terms in an equation if physics is partially known [4]. Machine learning is generally a good interpolator, but its results outside the data domain may fail and become unsatisfactory. On the other hand, when PIML is appropriately designed and trained, it can gain the capacity to extrapolate [5]. A third benefit of integrating physics into machine learning is promoting data parsimony: the data collection can be simple, but not simpler than required [5]. Therefore, the junction of the two fields can accelerate discovery in the field of science and engineering and enforce known physics.

The oil and gas industry faces various financial, operational, and environmental constraints. As a result, it is essential to develop dynamic models that account for complex behaviors to improve control, monitoring, and

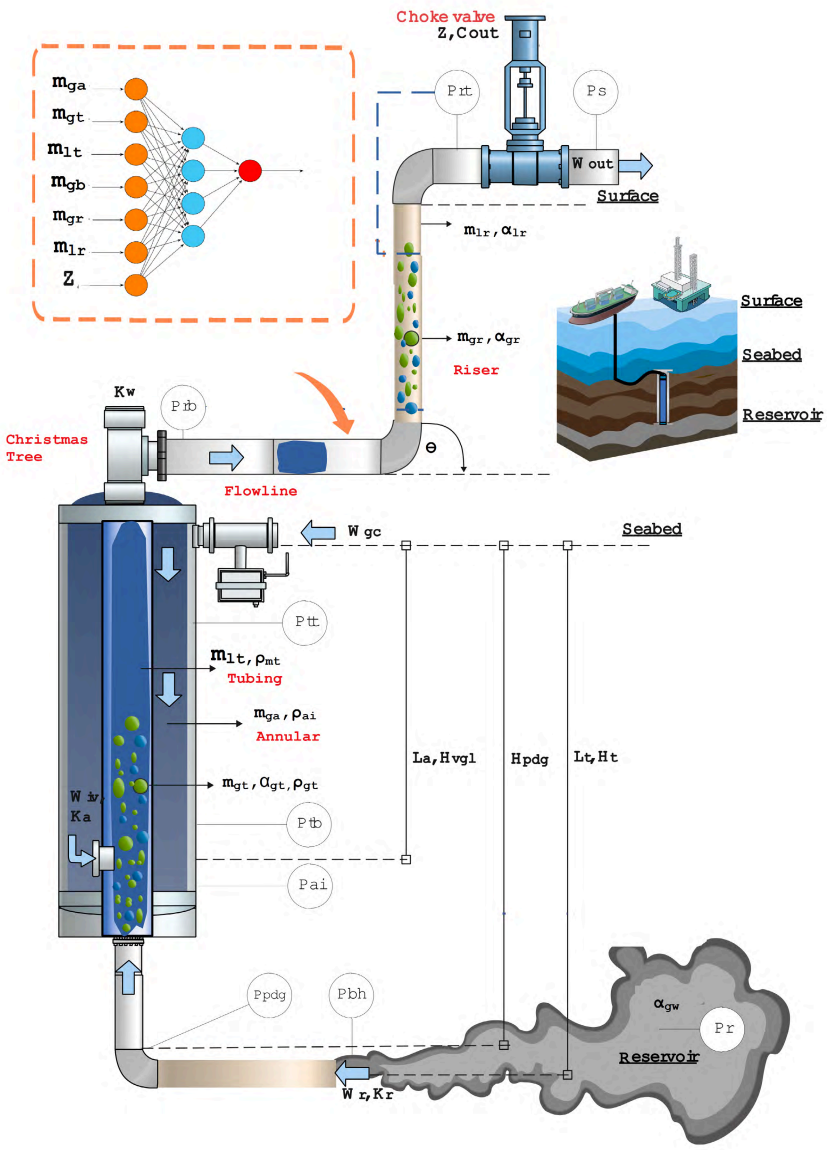

Figure 1. Neural Fast Offshore Model. Adapted from [11].

optimization. Neural networks can approximate the dynamic models, the so-called neural ordinary differential equations (NODEs). NODEs are designed by parametrizing the derivative of the hidden state using NNs. Then, it can be solved employing the adjoint method, which involves set of differential equations solved backward in time [6]. Physical constraints, also known as inductive bias, can be combined with the NODE in a universal differential equations (UDE) model [7]. In the oil and gas field, hybrid models combined with neural networks were performed by Kittelsen et al. [8], who employed physics-informed neural networks (PINNs) as a surrogate model for controlling gas-lift wells. However, PINNs need to consider the knowledge of the model in advance. In this sense, Faria et al. [9] developed a UDE model for gas-lift wells; however, the work aimed to replace algebraic equations with NNs for parameter estimation, without considering slug flow.

One significant challenge encountered by the industry is slug flow in multiphase flows. This phenomenon occurs when accumulated liquid obstructs the passage of gas. When the gas builds sufficient pressure, it can push the accumulated liquid, leading to sudden oscillations between liquid and gas flows. Additionally, the connections and angles between flowlines and risers can exacerbate this problem. Diehl et al. [10] developed a model that accounts for the reservoir, the gas-lift annulus, the production column, and the flowline-riser set. The authors employed a virtual valve to represent the point at which severe slug flow forms. The model has been used in several

works, for anti-slug PI control [11], for pressure estimation comparing Kalman filters and NN [12], and for transfer learning using recurrent networks [13].

Despite the model being validated in real production systems under some considerations, it is necessary to address a parameter estimation problem, which depends on the system conditions. Although simple, the model requires specific knowledge of the well to be tuned. Therefore, extending the model to other wells is a cumbersome task. Furthermore, the model assumes that the slug can be modeled based on a few parameters, considering that it occurs basically due to a virtual valve. For this reason, the long-term dynamics are particularly difficult to model, encouraging the use of more general tools, such as machine learning [13]. Therefore, we hypothesize that this virtual valve term, can be approximated by an NN embedded in a differential equation, with both first-principle and data-driven models working in tandem in the UDE scheme.

Apio et al. [12] compared the Kalman filter and NN for predicting downhole pressure measurements in slug flow. They conclude that NNs offer superior performance, but they failed outside the conditions of training. In that paper, just one choke condition was employed. In a recent work, Luz et al. [13] proposed a transfer learning approach to handle slug flows by using an NN trained over synthetic data. However, both works assume a pure black-box model to handle the system's complexity. In contrast, this work aims to employ universal differential equations to enhance the modeling of multiphase oil production systems under multiple choke conditions. Thus, we preserve the physical structure of the mass-balance model while letting the neural component capture unmodeled physics of the slugs.

This work is divided as follows: in the Methodology section, we describe the FP model used to generate the data and demonstrate the UDE model; in the next section, we present and discuss the UDE results. The final remarks and future works are presented in the Conclusions section.

METHODOLOGY

Description of the system

The multiphase production system model presented by Diehl et al. [10], referred to the Fast Offshore Wells Model (FOWM), was employed. The dynamic model is shown in Eqs. (1) to (6). The mass balance in the annulus, m_{ga}, is obtained from the mass flow rate of gas injected into the annulus (W_{gc}) and the gas flow entering the tubing, W_{iv}. In the tubing — Eq. (2), the gas balance also considers the production flow from the reservoir (W_{r}) and the gas flow at the wellhead (W_{whg}). The liquid flow rate in the tubing is derived in a similar manner — Eq. (3). α_{gw} is the gas fraction in the tubing, whereas E represents the mass fraction of gas that bypasses the slug in the subsea production line. The bubble formation through the slug — Eq. (4) — is modeled based on W_{g}, which corresponds to the flow through the virtual valve. In the riser, Eqs. (5) and (6), besides the upstream flows, the mass balance also accounts for W_{lout} and W_{gout}, which are, respectively, the liquid and gas flow rates through the surface choke valve.

$$\frac{\mathrm{d}m_{\mathrm{ga}}}{\mathrm{d}t} = W_{\mathrm{gc}} - W_{\mathrm{iv}} \tag{1}$$

$$\frac{\mathrm{d}m_{\mathrm{gt}}}{\mathrm{d}t} = W_{\mathrm{r}}\,\alpha_{\mathrm{gw}} + W_{\mathrm{iv}} - W_{\mathrm{whg}} \tag{2}$$

$$\frac{\mathrm{d}m_{\mathrm{lt}}}{\mathrm{d}t} = W_{\mathrm{r}}\left(1-\alpha_{\mathrm{gw}}\right) - W_{\mathrm{whl}} \tag{3}$$

$$\frac{\mathrm{d}m_{\mathrm{gb}}}{\mathrm{d}t} = (1-E)\,W_{\mathrm{whg}} - W_{\mathrm{g}} \tag{4}$$

$$\frac{\mathrm{d}m_{\mathrm{gr}}}{\mathrm{d}t} = E\,W_{\mathrm{whg}} + W_{\mathrm{g}} - W_{\mathrm{gout}} \tag{5}$$

$$\frac{\mathrm{d}m_{\mathrm{lr}}}{\mathrm{d}t} = W_{\mathrm{whl}} - W_{\mathrm{lout}} \tag{6}$$

The mass flow rates are modeled as shown in Eqs. (7) – (13). K_a is the flow coefficient between the annulus and the tubing in Eq. (7). K_r is a Vogel-type coefficient proportional to the reservoir contribution to the flow in the tubing in Eq. (8), whereas K_w represents the flow coefficient at the wellhead in Eqs. (9) and (10). The gas density in the annulus is denoted by ρ_{ai}. The gas mass fraction in the tubing is represented by α_{gt}, whereas α_{gr} and α_{lr}, denote, respectively, the gas and liquid mass fractions in the subsea production line. The constants C_g and C_{out} correspond to the virtual valve coefficient and the choke coefficient, respectively. The choke opening is defined by the fraction z. The pressures in the reservoir (bottomhole) and in the gravitational separator at the topside processing plant are represented by P_r and P_s, respectively.

$$W_{iv} = K_a\sqrt{\rho_{ai}(P_{ai}-P_{tb})} \tag{7}$$

$$W_{\mathrm{r}} = K_r\left[1 - 0.2\left(\frac{P_{bh}}{P_r}\right) - 0.8\left(\frac{P_{bh}}{P_r}\right)^2\right] \tag{8}$$

$$W_{\mathrm{whg}} = K_w\sqrt{\rho_L(P_{tt}-P_{rb})}\,\alpha_{gt} \tag{9}$$

$$W_{\mathrm{whl}} = K_w\sqrt{\rho_L(P_{tt}-P_{rb})}\left(1-\alpha_{gt}\right) \tag{10}$$

$$W_{\mathrm{g}} = C_g(P_{eb}-P_{rb}) \tag{11}$$

$$W_{\mathrm{gout}} = \alpha_{gr}\,C_{\mathrm{out}}\,z\sqrt{\rho_L(P_{rt}-P_s)} \tag{12}$$

$$W_{\mathrm{lout}} = \alpha_{lr}\,C_{\mathrm{out}}\,z\sqrt{\rho_L(P_{rt}-P_s)} \tag{13}$$

In addition to the density and gas fractions already described, the FOWM model also considers the following pressures: P_{ai}, pressure at the gas injection point from the annulus to the tubing; P_{tb}, pressure at the gas injection point on the tubing side; P_{bh}, bottomhole pressure; P_{pdg},

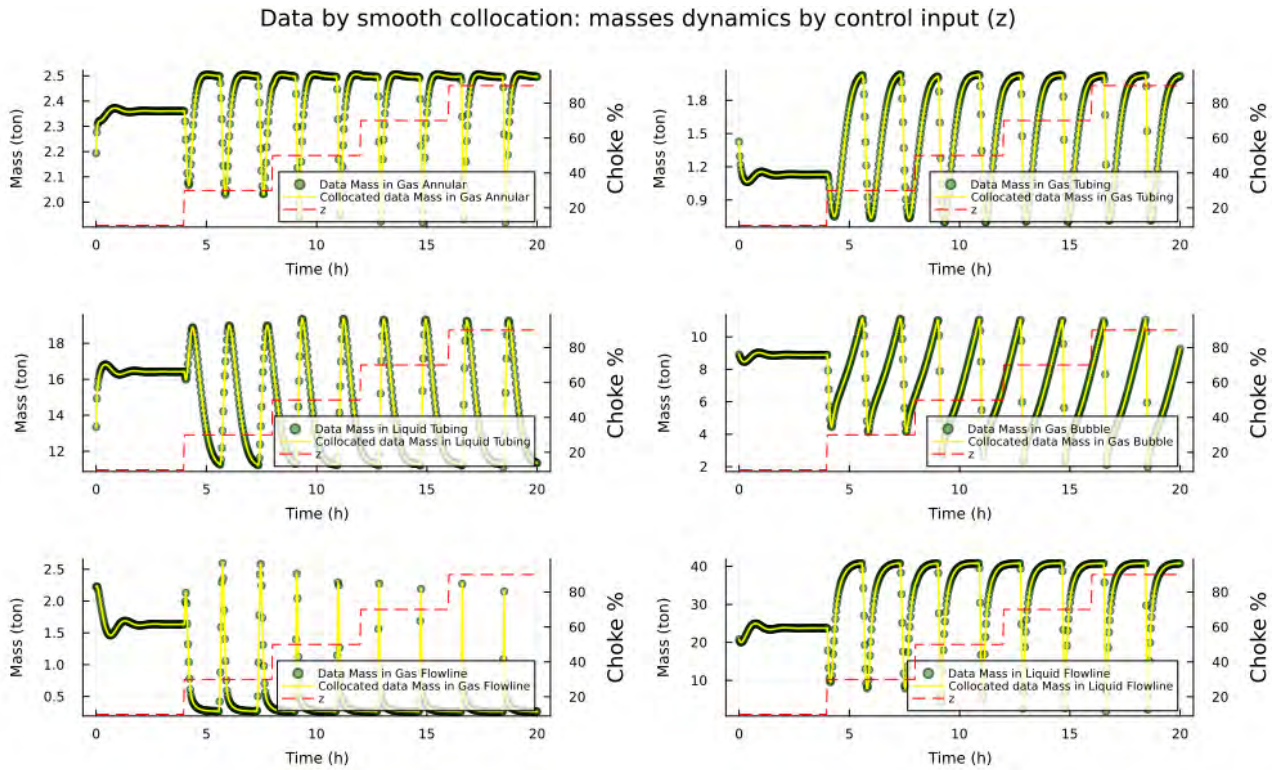

Figure 2. Results of interpolating data by cubic splines in the first-stage of training.

pressure at the PDG (Permanent Downhole Gauge); P_r, reservoir pressure near the bottom; P_{tt}, pressure at the tubing top; P_{rb}, pressure in the flowline upstream of the bubble point; P_{eb}, pressure at the bubble point; P_{rt}, pressure at the riser top. The described relations are represented in Eqs. (14) – (21):

$$P_{ai} = \left(\frac{RT}{V_a M} + \frac{gL_a}{V_a}\right) m_{ga} \tag{14}$$

$$P_{tb} = P_{tt} + \rho_{mt} g H_{vgl} \tag{15}$$

$$P_{bh} = P_{pdg} + \rho_{mres} g\left(H_t - H_{pdg}\right) \tag{16}$$

$$P_{pdg} = P_{tb} + \rho_{mres} g\left(H_{pdg} - H_{vgl}\right) \tag{17}$$

$$P_{tt} = \frac{\rho_{gt} RT}{M} \tag{18}$$

$$P_{rb} = P_{rt} + \frac{(m_{lr} + m_{L,still}) g \sin(\vartheta)}{A_{ss}} \tag{19}$$

$$P_{eb} = \frac{m_{gb} RT}{M V_{eb}} \tag{20}$$

$$P_{rt} = \frac{m_{gr} RT}{M}\left(\omega_u V_{ss} - \frac{m_{lr} + m_{L,still}}{\rho_L}\right) \tag{21}$$

The auxiliary variables used to compute gas density and mass fractions are given in Eqs. (22) – (25).

$$\rho_{ai} = \frac{M P_{ai}}{RT} \tag{22}$$

$$\alpha_{gt} = \frac{m_{gt}}{m_{gt} + m_{lt}} \tag{23}$$

$$\alpha_{gr} = \frac{m_{gr}}{m_{gr} + m_{lr}} \tag{24}$$

$$\alpha_{lr} = 1 - \alpha_{gr} \tag{25}$$

In the model, T denotes the average temperature along the production system, M represents the gas molecular weight, R is the universal gas constant, and g is the gravitational acceleration. The annulus is characterized by its volume V_a and length L_a. The mixture density in the tubing, ρ_{mt}, is computed from the fluid properties in this section, whereas the corresponding gas density, ρ_{gt}, is obtained separately. The mixture density in the reservoir, ρ_{mres}, is assumed constant. The vertical distances from the wellhead to key points of the system are: H_{vgl} to the gas-lift valve, H_{pdg} to the PDG sensor, and H_t to the bottomhole. The average riser inclination is denoted by ϑ.

The minimum liquid mass in the subsea line is represented by $m_{L,\,\mathrm{still}}$, and the gas bubble volume by V_{eb}. An auxiliary parameter, ω_u, is introduced to help locate the bubble. The cross-sectional area of the subsea line, A_{ss}, and its total volume are computed based on the pipeline geometry. Additionally, the volume of the gas-lift annulus is determined through standard geometric relationships. The production system parameters were tuned to

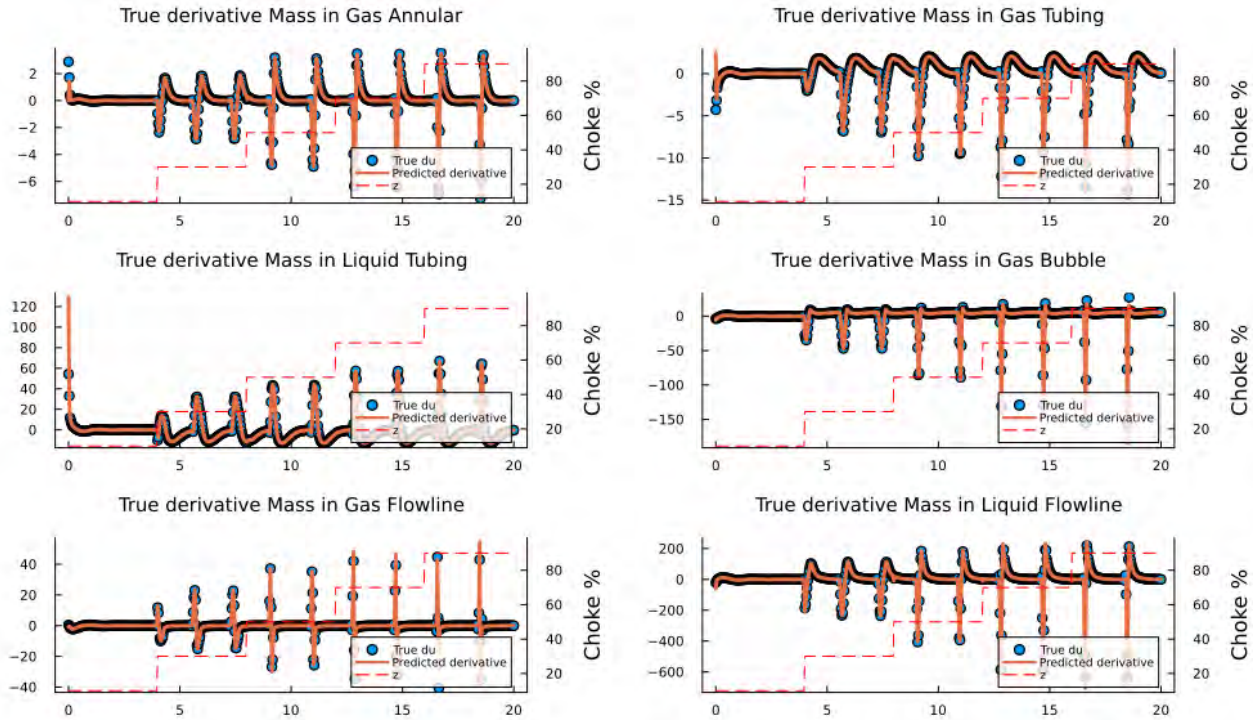

Figure 3. Predicted derivatives by the smooth colocation against data for each of the mass holdups.

reproduce real wells based on OLGA simulations, as listed in Table 1.

Table 1: Parameters from a real production well.

Parameter (Unit)	Unit
$\rho_L(kg/m^3)$	900
$P_r(bar)$	225
$P_s(bar)$	10
α_{gw}	0.0188
$\rho_{mres}\ (kg/m^3)$	892
$M\ (kg/kmol)$	18
$T\ (K)$	298
$L_r(m)$	1569
$L_{fl}(m)$	2928
$L_t(m)$	1639
$L_a(m)$	1118
$H_t(m)$	1279
$H_{pdg}(m)$	1117
$H_{vgl}(m)$	916
$D_{ss}(m)$	0.15
$D_t(m)$	0.15
$D_a(m)$	0.14
$m_{L,still}(kg)$	7.110×10^2
$C_g(ms)$	2.346×10^{-5}
$C_{out}(m^2)$	5.814×10^{-3}
$V_{eb}(m^3)$	9.016×10^1
E	3.582×10^{-2}
$K_w(m^2)$	1.021×10^{-3}
$K_a(m^2)$	1.767×10^{-4}
$K_r(kg/s)$	2.467×10^2
ω_u	2.780

The remaining geometric and mixture relations are given in Equations (26) – (32).

$$\rho_{mt} = \frac{m_{gt}+m_{lt}}{V_t} \tag{26}$$

$$\rho_{gt} = \frac{m_{gt}}{V_{gt}} \tag{27}$$

$$V_{gt} = V_t - \frac{m_{lt}}{\rho_L} \tag{28}$$

$$A_{ss} = \frac{\pi D_{ss}^2}{4} \tag{29}$$

$$V_{ss} = \frac{\pi D_{ss}^2}{4}\left(L_r + L_{fl}\right) \tag{30}$$

$$V_a = \frac{\pi D_a^2 L_a}{4} \tag{31}$$

$$V_t = \frac{\pi D_t^2 L_t}{4} \tag{32}$$

Modeling by Universal Differential Equations

This study uses a hybrid physics-informed machine learning approach to model and predict the dynamic

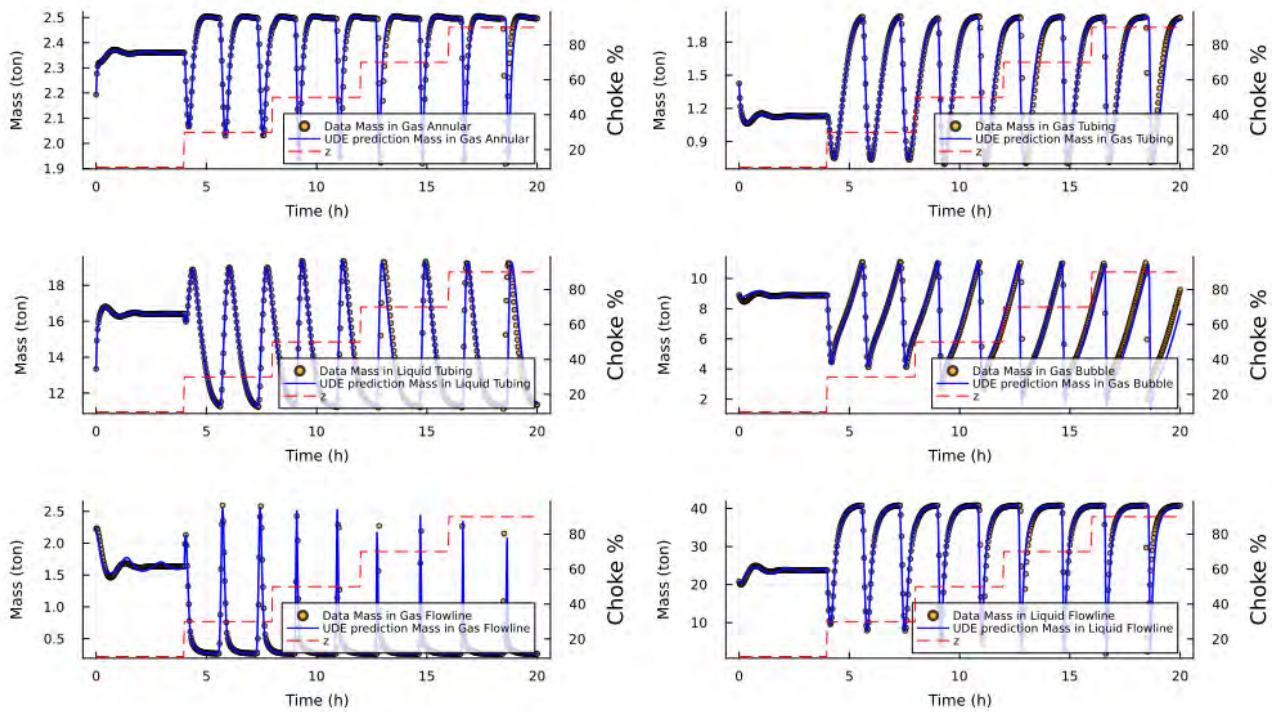

Figure 4. The initial UDE estimation of each of the mass holdups after the finishing the first-stage of training. The initial estimations are represented against data in different choke openings (z).

behavior of a multiphase oil production system. The methodology handles the complexity of simplified terms through a data-driven method. The data source comes from the simulation performed with the FOWM model described in the previous section. This work employs UDE, i.e., a neural network is embedded into the structure of the FOWM differential equations, replacing the virtual valve flow in Eqs. (4) and (5), as illustrated in Fig. 1. The neural network receives the six differential states, in addition to the choke valve opening (z), as inputs to a 3-layer hidden topology with 7 neurons each, using radial basis function (RBF) as the activation function. The virtual valve flow rate is considered as the output. The choke valve opening was varied from 10% to 90%, considering a 20% increment. The gas-lift flow rate was considered constant at 80, 000 m³/d. Training was performed using data from a 20-hour simulation of well operation, with a 1-minute time step. For the testing phase, the valve was considered with openings from 20% to 80%, using the same increment, to test generalization within the trained range, and at 100% to verify extrapolation capability.

Since the system exhibits oscillatory behavior, the UDE was trained in two stages: the first uses smooth data collocation to obtain an estimate of the derivatives, which will feed the UDE. In this work, we have access to the data of the states, but not to its derivatives. Therefore, we proceed to its estimation employing smooth collocation. There are several methods available to smooth the noisy data, including kernel-smoothing methods and piecewise interpolation. In this work, cubic splines were chosen instead of a kernel-based method because the data exhibit many oscillations. The second method would lead to discarding data when peaks are present. Once the data are interpolated, it is possible to compute the derivatives locally and then minimize the error between these derivatives and the UDE derivative outputs, as stated in Eq. (33). The θ represents the parameters of the NN, while the indexes i and j are, respectively, the states and the sample time. The superscript $\widehat{\mathrm{m}_i}$ represents the prediction of the state i by the UDE, the $\tilde{}$ represents the derivative estimation by the interpolation method, while the variables are marked accordingly to represent the data.

$$\mathrm{MSE}_{\mathrm{deriv}}(\theta) = \frac{1}{\mathrm{MN}} \sum_{i,j=1}^{M,N} \left\| \frac{\widehat{\mathrm{dm}_i(\theta,j)}}{\mathrm{dt}} - \frac{\mathrm{d}\tilde{\mathrm{m}}_i(j)}{\mathrm{dt}} \right\|^2 \quad (33)$$

This avoids a high computational cost in the initial phase, since it is not necessary to repeatedly solve the differential equation during the coarse tuning of the UDE parameters. The backpropagation process is employed to find an initial θ.

The incremental strategy provides a reasonable starting point for the second stage, the training of the complete UDE model, which involves backpropagation

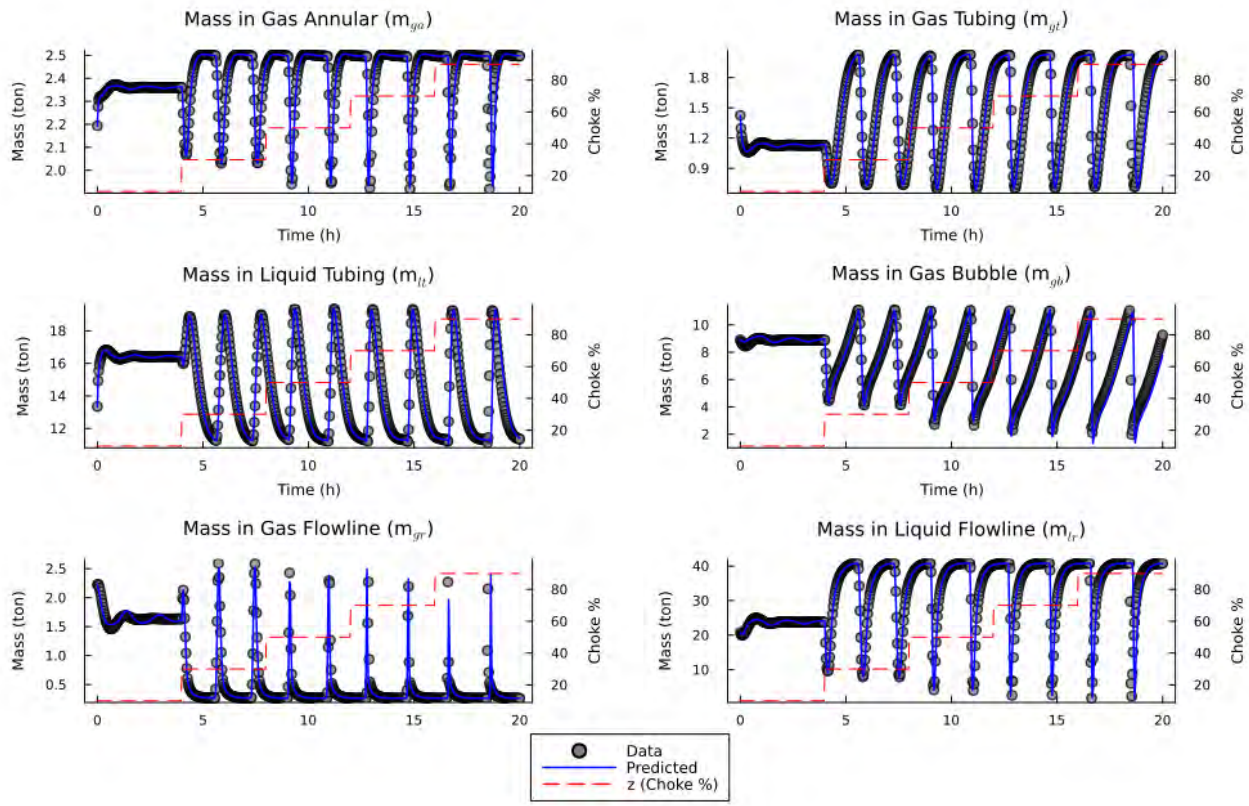

Figure 5. The UDE predictions of each of the mass holdups after the finishing the second-stage of training. The predictions are represented against data in different choke openings (z).

and solving the differential equation simultaneously. The objective function involves minimizing the mean squared error (MSE) between the predicted and expected (scaled) values of the states (mass holdup), Eq. (34).

$$\mathrm{MSE}_{\mathrm{state}}(\theta) = \frac{1}{\mathrm{MN}} \sum_{i,\mathrm{j}=1}^{\mathrm{M},N} \left\| \frac{\widehat{\mathrm{m}_{\mathrm{i}}(\theta,\mathrm{j})}}{\max(\mathrm{m}_i^{data}) - \min(\mathrm{m}_i^{data})} - \frac{\mathrm{m}_i^{data}(\mathrm{j})}{\max(\mathrm{m}_i^{data}) - \min(\mathrm{m}_i^{data})} \right\|^2 \tag{34}$$

The first phase was limited to a maximum of 10, 000 iterations, and the second to a maximum of 5, 000 iterations using the ADAM algorithm, followed by 1, 000 iterations with the LBFGS algorithm. Early stopping was used to interrupt training if no improvement was observed after 500 iterations. The code was developed using the Scientific Machine Learning package of the Julia language.

RESULTS AND DISCUSSION

The results of the first stage (smooth data collocation) with cubic splines interpolation are shown in Fig. 2 for the state variables and in Fig. 3 for their time derivatives. Since the data are noise-free simulation data, the piecewise polynomials pass through all data points. The results of the initially estimated UDE are shown in Fig. 4, where the derivatives minimization was able to reasonably predict the state data. The MSE of the derivatives were 205.39, while the scaled MSE was 0.0099. This enables fast convergence, in which the first stage only performs a coarse adjustment, while the full network training constitutes the fine-tuning, when the computational cost is higher. In the literature, this is considered a recommendation for complex systems. Huang et al. [14] trained a UDE for wastewater process employing smooth collocation followed by direct UDE training in a stiff system. As the authors stated, the stiffness may arise not only from the system itself but also from the introduction of the NN, leading to a numerical stiffness. By using this incremental strategy, it is possible to circumvent the instability hurdle at the beginning of training, providing smooth gradients to the second stage, in which the differential equation is solved simultaneously with the backpropagation process. Furthermore, this strategy can minimize the computational overhead.

Therefore, with a reasonably good initial estimate, the model can capture the dynamics present in the slugging system, as shown in Fig. 5. In this second stage, the

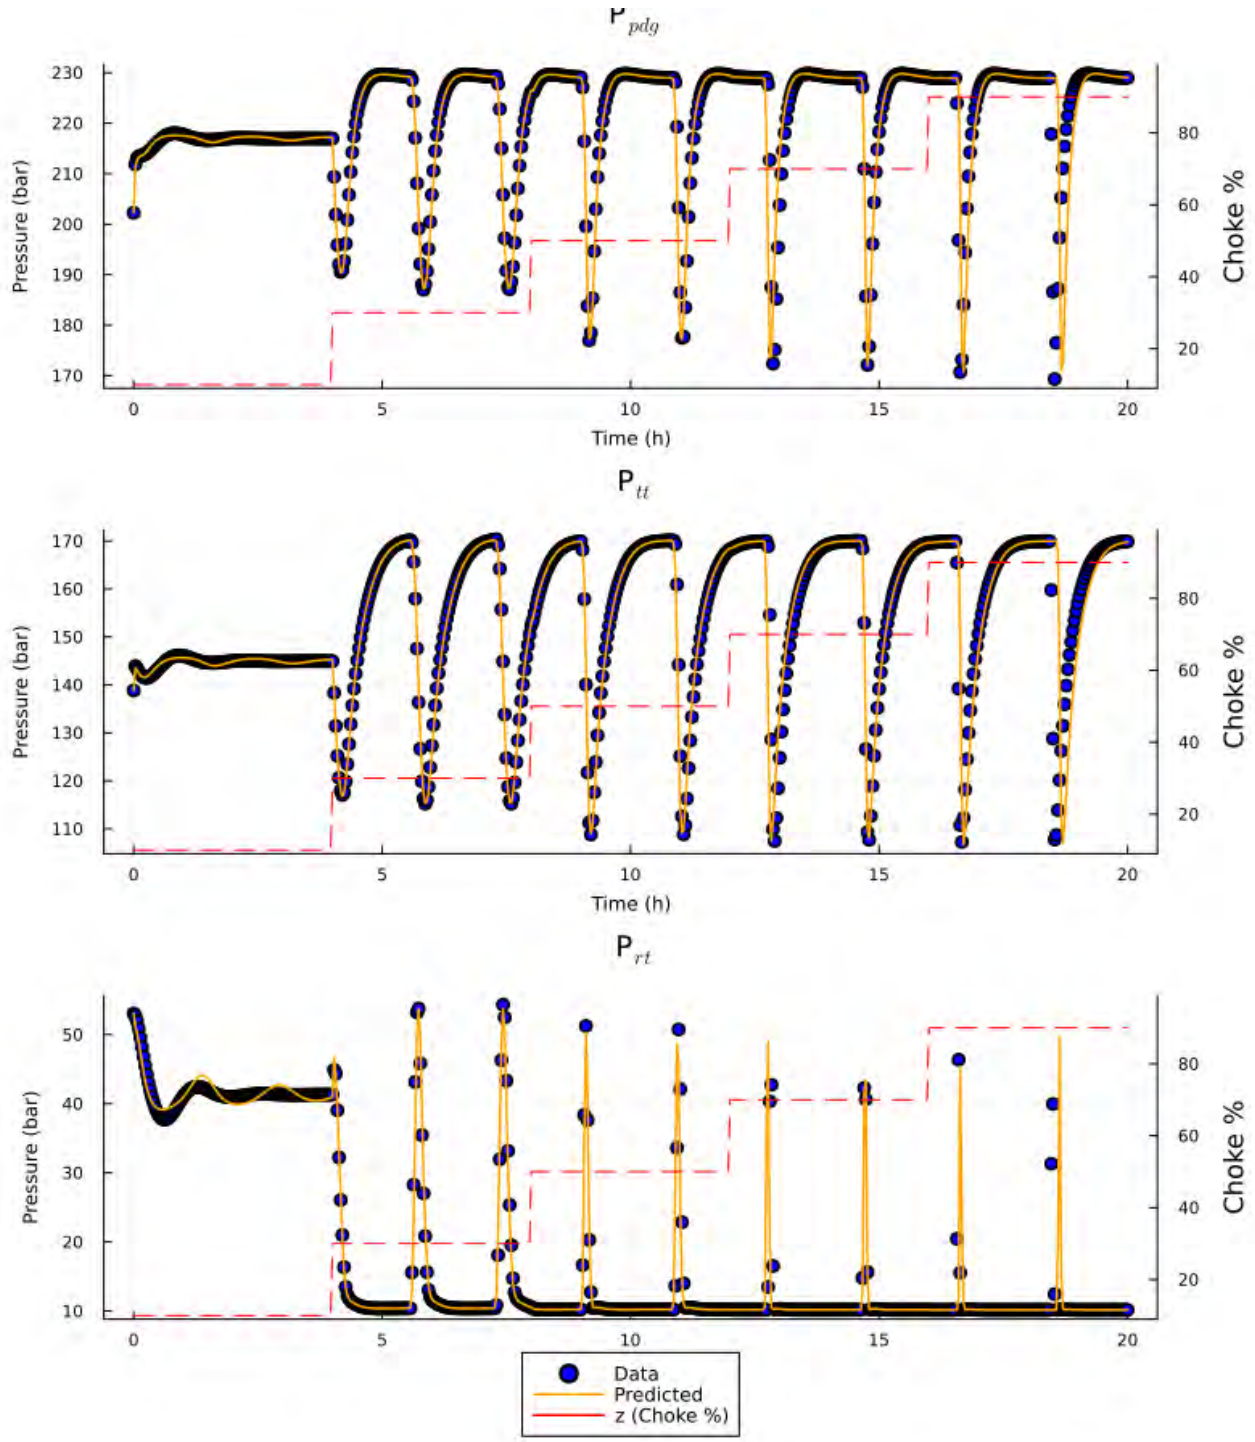

Figure 6. Pressure predictions by the UDE after the finishing the second-stage of training. The predictions are represented against data in different choke openings (z).

algorithm achieved a mean squared error of 0.0092 (in scaled units) after 505 iterations. During this phase, the algorithm explored regions in which the MSE increased, but it stopped after a small decrease of the MSE found based on the first stage. . It is worth noting that without the pre-training stage, network convergence is impossible due to the inherent numerical instability of hybrid models and the tendency to form local minima. We tried to bypass the pre-training step, but it failed to converge in all cases. It should be noted that the oscillations

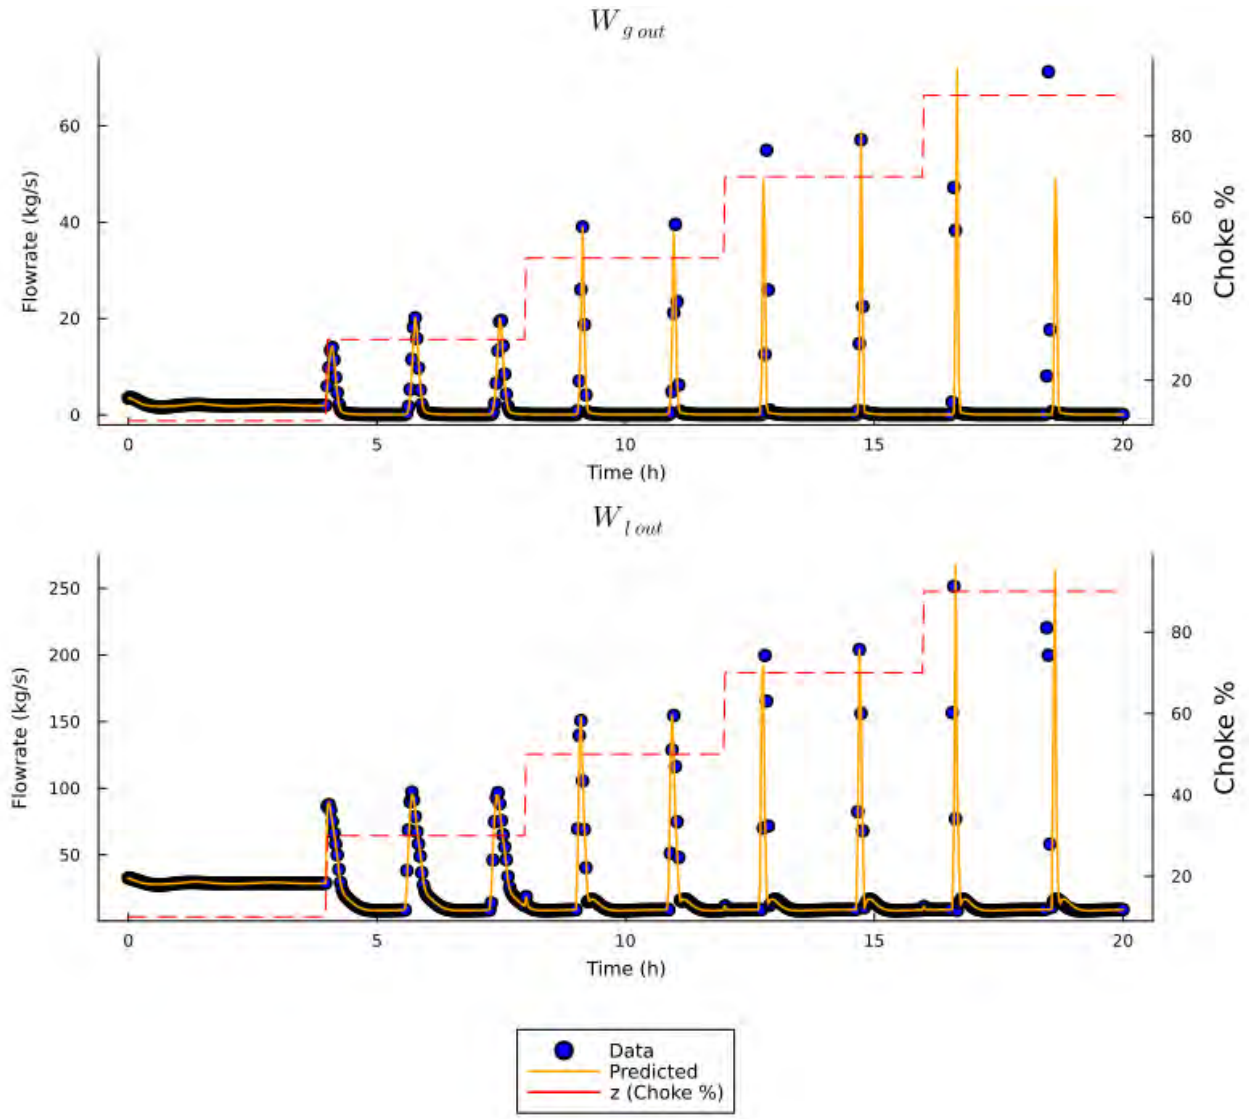

Figure 7. Flow rates predictions by the UDE after the finishing the second-stage of training. The predictions are represented against data in different choke openings (z).

increase as the choke is more opened. The NN can capture the trend as well, suggesting that it can generalize considering exogenous inputs.

Comparisons of pressures at the PDG location, at the tubing top, and across the choke are shown in Fig. 6. It can be observed that the pressures computed by the UDE model are accurate, as with a reasonable state estimate, they are calculated directly from Eqs. (17), (18), and (21). In real production systems, these are the measurements usually available, and they can be used as indicators of the model's quality. The model can be extended to consider production data from the plant.

Similarly, the predicted flow rates are shown in Fig. 7. The results suggest that universal differential equations can be applied to systems whose physics is poorly known. Compared to the transfer learning approach based on LSTM networks [12], the proposed UDE formulation does not rely on a purely black-box architecture. Instead, it also considers the governing mass-balance equations, thus, it enables extrapolation of the choke openings in the slug regime. Since the flow rates are usually not measured in the oil platforms, this has led to the development of virtual flow meters in recent works. With our model, it tackles the problem in a two-fold way: by predicting the pressures for anti-slug control and by acting as a virtual flow meter, which is vital for maximizing productivity. With a good prediction of the pressures, a control for suppressing slug regimes can be developed based on the UDE model, as demonstrated by using FP FOWM in [12].

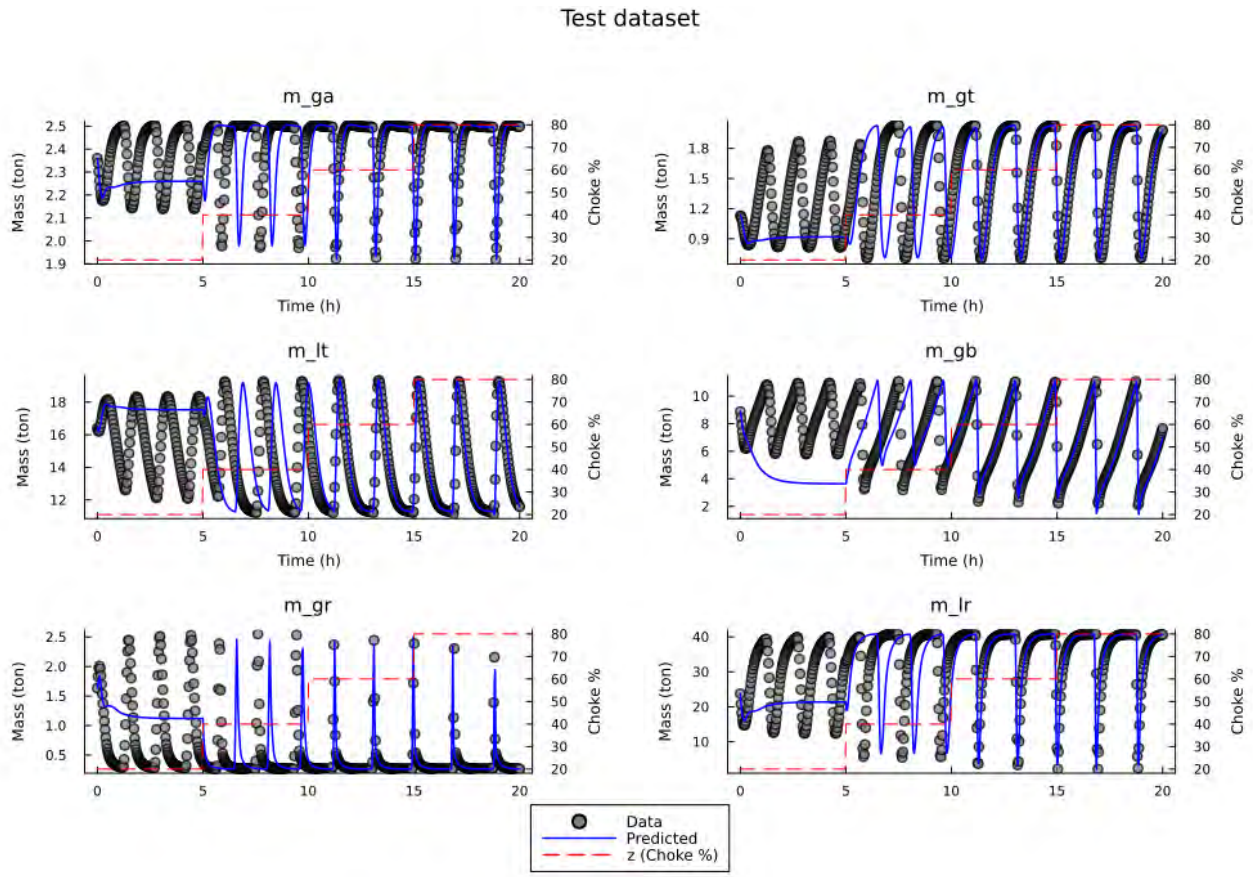

Figure 8. Test predictions by the UDE. The predictions are represented against data in different choke openings (z).

As discussed in previous work [12, 13], using a simple NN when changing choke scenarios can compromise its predictive capacity. This justifies the use of more complex ML strategies, such as transfer learning [13] or UDE modeling as in this work. However, previous works were based solely on pressure measurements, without accounting for the advantages of flow rate evaluation. Moreover, Luz et al. [13] conducted 1000 simulations to enable the NN to capture the underlying patterns. Results in the literature suggest that using physical priors can reduce the data requirement [2, 5], but more studies are needed to verify this claim in our context.

The results for the prediction on the test set are shown in Fig. 8. The test set considers interpolation between 20% and 80% choke opening, with a 20% increment. The reported scaled MSE was 0.088. The main limitation identified occurs in the prediction at 20% opening, in the transition region associated with a bifurcation. In this range, the errors are higher, and a local minimum is reached, inherited from the knowledge before the bifurcation. The results suggest that, for limit cycles, training over the transition region should be refined; once the oscillations are established, the model can capture the long-term dynamics.

On the other hand, for the extrapolation at 100% opening, the achieved scaled MSE was 0.039. As illustrated in Fig. 9, the UDE was able to adequately reproduce the trends of the states, suggesting that the model incorporated useful knowledge of the underlying physics even outside the exact training range.

As a main limitation of the work, we should note that, to apply to real data, only pressure measurements are available. In theory, one could use the MSE between the pressure measurements and the predictions in the UDE model, just changing Eq. (34). However, it should be noted that this would implying in skipping the first stage of training, as in Eq. (33). One suggestion would be to use the training of the adjusted FOWM data as warm up and apply transfer learning over the real data.

CONCLUSION

In this work, we proposed a universal differential equation (UDE) formulation for modeling slug-flow dynamics in offshore wells, using the Fast Offshore Wells

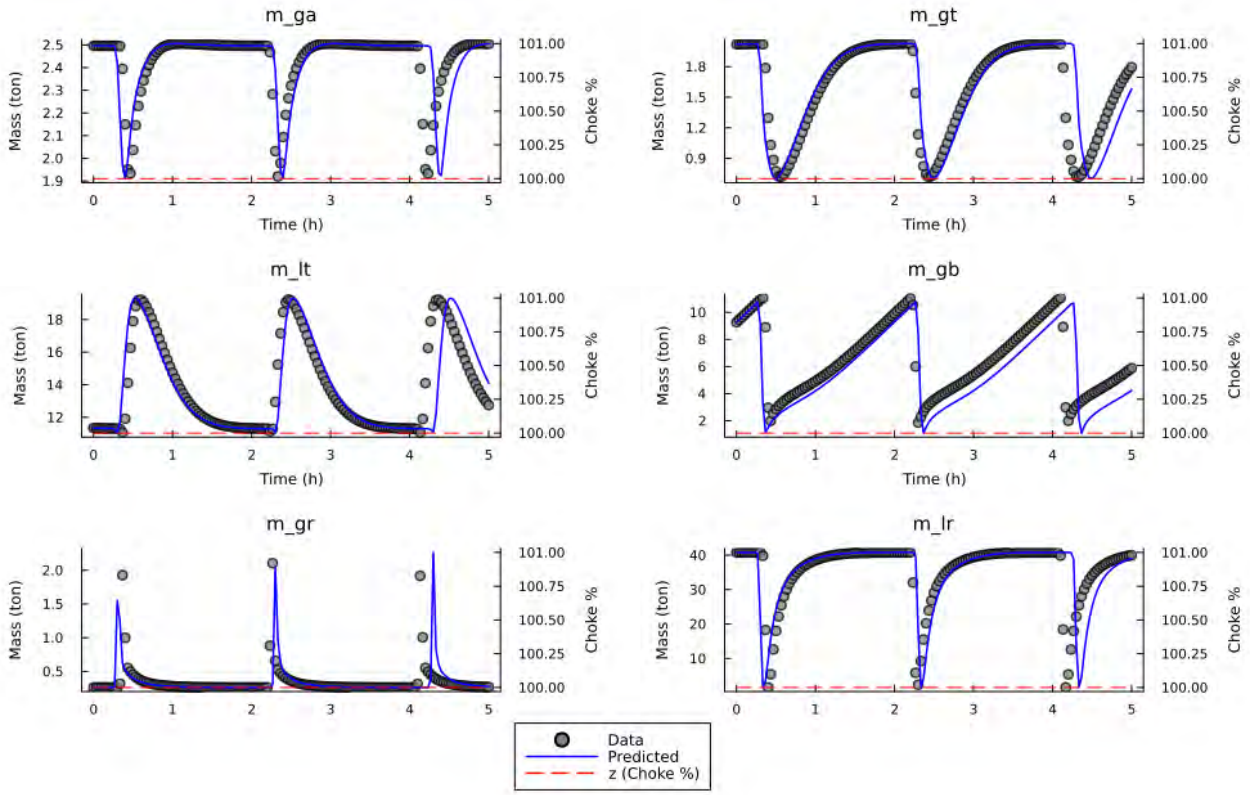

Figure 9. Extrapolation predictions by the UDE. The predictions are represented against data in different choke openings (z).

Model (FOWM) as the physical backbone. A neural network was embedded into the mass-balance equations by replacing the virtual valve flow term, allowing the model to capture unmodeled slugging physics while preserving the overall state-space structure.

A two-stage training strategy based on data collocation was employed to address numerical instability in the hybrid model. In the first stage, cubic splines were used to obtain local derivative estimates, which were used to get an initial estimate of the NN parameters. This step provided a suitable initialization for the second stage, in which the full UDE was trained by solving the differential equations and minimizing the mean squared error between the predicted and scaled state trajectories. The proposed approach achieved a scaled MSE of 0.0092 on the training scenarios and 0.088 on interpolation tests with choke openings between 20% and 80%. For the extrapolation case at 100% opening, the scaled MSE was 0.039, and the UDE was able to reproduce the mass holdup trends.

These results indicate that UDE-based hybrid models represent a feasible and promising possibility for modeling multiphase production systems. The methodology can be applied to other systems in the limit-cycle as well, especially when physics is partially known. Future work will focus on validating the methodology against high-fidelity simulators and field data, while incorporating the considerations of using only available measurements. Furthermore, future work should discuss in more detail the stability of the hybrid models.

ACKNOWLEDGEMENTS

This study was financed in part by CAPES - Finance Code 001 and by Petrobras S.A. (Cooperation term no. 0050.0124501.23.9). Professors Maurício B. de Souza Jr. and Argimiro R. Secchi are grateful to financial support from CNPq (Grants No. 311153/2021-6 and 300744/2025-0), and Prof. Maurício B. de Souza Jr. is also thankful to FAPERJ (Grant No. E-26/201.148/2022).

AUTHOR IDENTIFIERS

Author ORCIDs:
Caldas GLR: 0000-0003-2554-9561
Gerevini G: 0000-0002-1163-029X
Diehl FC: 0000-0002-7736-1395
Nogueira IBR: 0000-0002-0963-6449

Souza MB Jr: 0000-0002-1090-8958
Secchi AR: 0000-0001-7297-3571

REFERENCES

1. Psichogios DC, Ungar LH. A hybrid neural network?first principles approach to process modeling. AIChE Journal 38:1499-1511 (2004). https://doi.org/10.1002/aic.690381003
2. Bradley W, Kim J, Kilwein Z, Blakely L, Eydenberg M, Jalvin J, Laird C, Boukouvala F. Perspectives on the integration between first-principles and data-driven modeling. Computers & Chemical Engineering 166:107898 (2022). https://doi.org/10.1016/j.compchemeng.2022.107898
3. Gao L, Guan L. Interpretability of machine learning: recent advances and future prospects. IEEE MultiMedia 30:105-118 (2023). https://doi.org/10.1109/mmul.2023.3272513
4. Sahimi M. Physics-informed and data-driven discovery of governing equations for complex phenomena in heterogeneous media. Phys. Rev. E 109: (2024). https://doi.org/10.1103/physreve.109.041001
5. Zhu SP, Wang L, Luo C, Correia JAFO, De Jesus AMP, Berto F, Wang Q. Physics-informed machine learning and its structural integrity applications: state of the art. Phil. Trans. R. Soc. A. 381: (2023). https://doi.org/10.1098/rsta.2022.0406
6. Chen RTQ, Rubanova Y, Bettencourt J, Duvenaud D. Neural ordinary differential equations. Adv Neural Inf Process Syst 32:6572-6583 (2018)
7. Rackauckas C, et al. Universal differential equations for scientific machine learning. arXiv:2001.04385 (2020) https://doi.org/10.48550/arXiv.2001.04385
8. Kittelsen JE, Antonelo EA, Camponogara E, Imsland LS. Physics-informed neural networks with skip connections for modeling and control of gas-lifted oil wells. Applied Soft Computing 158:111603 (2024). https://doi.org/10.1016/j.asoc.2024.111603
9. Faria RR, Lima FARD, Secchi AR, de Souza MB. Enhanced hybrid model for gas-lifted oil production. IFAC-PapersOnLine 58:7-12 (2024). https://doi.org/10.1016/j.ifacol.2024.08.305
10. Diehl FC, Anzai TK, Almeida CS, Von Meien OF, Neto SS, Rosa VR, Campos MCMM, Reolon F, Gerevini G, Ranzan C, Farenzena M, Trierweiler JO. Fast offshore wells model (FOWM): a practical dynamic model for multiphase oil production systems in deepwater and ultra-deepwater scenarios. Computers & Chemical Engineering 99:304-313 (2017). https://doi.org/10.1016/j.compchemeng.2017.01.036
11. Clavijo NA, Diehl FC, Gerevini GG, Melo PA, Lemos T, Campos LFO, Pinto JC. Chaotic dynamics induced by anti-slug control in offshore oil production plants. Journal of Petroleum Science and Engineering 215:110716 (2022). https://doi.org/10.1016/j.petrol.2022.110716
12. Apio A, Dambros JWV, Diehl FC, Farenzena M, Trierweiler JO. PDG pressure estimation in offshore oil well: extended kalman filter vs. artificial neural networks. IFAC-PapersOnLine 52:508-513 (2019). https://doi.org/10.1016/j.ifacol.2019.06.113
13. Luz EML, Spina DE, Pinto JC. Application of transfer learning for modeling slugging in offshore oil production. Geoenergy Science and Engineering 255:214069 (2025). https://doi.org/10.1016/j.geoen.2025.214069
14. Huang X, Kandris K, Katsou E. Training stiff neural ordinary differential equations in data-driven wastewater process modelling. Journal of Environmental Management 373:123870 (2025). https://doi.org/10.1016/j.jenvman.2024.123870

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Estimation of Thermodynamic Properties for Cellulosic Biomass-Derived Compounds: Application to Heat and Work Balances in Process Simulation

Anthony D. Anastasi[a]*, Cornelius M. Masuku[b], Praveen Ravikumar[a], Shishir P.S. Chundawat[a], and Diane Hildebrandt[a]

[a] Rutgers, The State University of New Jersey, Department of Chemical and Biochemical Engineering, Piscataway, NJ, USA
[b] Purdue University, Davidson School of Chemical Engineering, West Lafayette, IN, USA
* Corresponding Author: dh985@soe.rutgers.edu.

ABSTRACT

Reliable data for the standard enthalpies and Gibbs free energies of formation, ΔH_f° and ΔG_f°, are essential for process synthesis, energy integration, and lost-work analysis. However, many biomass-derived compounds lack reliable thermodynamic property data, limiting optimization of energy and carbon utilization in biomass conversion processes. This study proposes a composition-based method to estimate ΔH_f° and ΔG_f° for compounds containing carbon, hydrogen, and oxygen. The method exploits widely available heats of combustion data and establishes a linear correlation between the enthalpy and Gibbs free energy of combustion, ΔH_C and ΔG_C, using tabulated organic compounds. The applicability of this relationship to biomass-derived compounds is tested using published data for cellulose, starch, and glucose. Thornton's correlation between heat of combustion and oxygen demand is then incorporated to derive simple expressions for estimating formation properties directly from elemental composition. The results show that the proposed framework provides sufficiently accurate estimates for conceptual thermodynamic analysis, enabling rapid evaluation of reaction enthalpies, minimum work requirements, equilibrium constants, and reaction driving forces. The method is therefore well suited for early-stage process synthesis and screening of biomass conversion pathways.

Keywords: Biomass, Biosystems, Energy Systems, Energy Efficiency, Modelling and Simulations

INTRODUCTION

The optimization of energy utilization in industrial processes involving reacting systems requires reliable values of the standard heats of formation, ΔH_f°, of all reacting components [1]. A complementary and more fundamental measure of process efficiency is the extent to which the chemical potential of the feed streams, together with process workflows, is conserved in the chemical potentials of the product streams. Any deviation from this ideal corresponds to lost work arising from irreversibilities associated with chemical reaction, separation, heat transfer, and mixing. Quantification and targeting of this lost work require accurate values of the standard Gibbs free energies of formation, ΔG_f°, for all compounds appearing in the feed and product streams [2].

Conversion of biomass feedstocks into chemicals and fuels represents a key pathway toward improving the sustainability of the chemical industry. From a thermodynamic perspective, knowledge of the Gibbs free energy of biomass feedstocks is essential for evaluating how efficiently chemical potential in the feed, together with supplied work, is transferred to the product streams. This information is required to quantify minimum work requirements, assess unavoidable irreversibilities, and determine how closely a process can approach its thermodynamically imposed limits on energy use and CO_2 emissions. A representative example is the conversion of lignocellulosic biomass to bioethanol, for which accurate Gibbs free energy data are necessary to distinguish fundamental thermodynamic constraints from inefficiencies arising from process structure and operation [3].

Such analyses and optimizations have been significantly hindered by the scarcity of reliable thermodynamic data for biomaterials, as well as by inconsistencies among reported values arising from differences in experimental methods, thermodynamic models, and sample characteristics [4-5]. In the case of cellulose, the wide range of reported thermodynamic properties may partly reflect its historically low economic value, as it was predominantly combusted rather than utilized as a chemical feedstock, reducing the incentive for detailed thermodynamic characterization. More recently, however, large quantities of underutilized lignocellulosic residues have become attractive as renewable feedstocks, motivating renewed interest in their thermodynamic properties [6-7].

Compounding this challenge, many biomass-derived compounds, including cellulose and xylan, exhibit polymorphism, with multiple allomorphs differing in crystallinity and molecular structure [5, 8]. These structural variations give rise to corresponding differences in thermodynamic properties and are further complicated by the inherently amorphous and heterogeneous nature of components such as lignin. Collectively, these gaps and uncertainties in thermodynamic property data constrain rigorous process optimization, CO_2 utilization studies, and the development of reliable life cycle and techno-economic analyses for biomass conversion processes. More critically, they preclude the application of lost-work and exergy-based targeting methods, which are central to early-stage process synthesis and the identification of fundamentally efficient process structures.

While values of ΔH_f° for biomass-derived compounds are generally absent from standard thermodynamic tables, heats of combustion, ΔH_C, have been reported for many biocompounds and bio-feedstocks (9-10). These reported values often exhibit substantial variability, which may be attributed to differences in elemental composition, crystallinity, and moisture content of the samples (10-11). If the elemental composition of a compound is known, ΔH_f° can be calculated from the measured value of ΔH_C using the appropriate combustion stoichiometry.

In contrast, the determination of ΔG_f° is considerably more challenging, as neither ΔG_f° nor the Gibbs free energy of combustion, ΔG_C, can be measured directly. In some cases, reaction equilibrium measurements may be used to estimate ΔG_f°; however, this approach is not feasible for many biomass-derived compounds. An alternative route is provided by the Third Law of Thermodynamics, which requires knowledge of the compound's heat capacity from 0 to 298 K. For biomass materials, such data are rarely available and typically rely on a combination of assumptions and indirect measurements.

In this paper, a method is proposed for estimating ΔG_C for a general organic compound of the form $C_xH_yO_z$ based on its heat of combustion, ΔH_C. The approach is based on the observation of a linear correlation between ΔH_C and ΔG_C across a wide range of organic compound families, including alkanes, alkenes, alcohols, ketones, aldehydes, and carboxylic acids. We hypothesize that this linear relationship, established for organic compounds composed of carbon, hydrogen, and oxygen, also holds for biomass-derived compounds. This hypothesis is tested using published values of $\Delta H_C, \Delta H_f^\circ$, and ΔG_f° for compounds such as cellulose, starch, xylan, and glucose. It is shown that the reported biomass data lie slightly above the correlation obtained for general organic compounds but nevertheless obey the same linear relationship to a good approximation. Consequently, for any compound of the form $C_xH_yO_z$, we proposed that the value of ΔG_C may be estimated from the measured ΔH_C, from which ΔG_f° can be calculated directly using reaction stoichiometry.

The method is further extended by incorporating the linear relationship between ΔH_c and oxygen demand proposed by Thornton, yielding a simple linear relationship between ΔG_f° and the elemental C:H:O ratio of the compound. This framework provides a practical basis for lost-work targeting in processes that convert one organic or biomass-derived compound into another, enabling early identification of minimum work requirements and unavoidable irreversibilities prior to detailed flowsheet development. Moreover, it enables estimation of reaction enthalpies, Gibbs free energies, equilibrium constants, and their temperature dependence for reactions involving organic and biomass compounds. These estimates offer high-level insight into minimum energy and work requirements and reaction driving forces, giving an understanding of the overall thermodynamic difficulty and energy operating costs of biomass conversion processes.

METHODS

Analysis of Enthalpy and Gibbs Free Energy of Tabulated Organic Compounds

To maintain a consistent thermodynamic basis, data for organic compounds were obtained from the NIST Chemistry WebBook [12]. For each compound, the database provides the standard heat of formation, ΔH_f°, in the appropriate standard state, together with the standard molar entropy, S°, at 298.15 K. These quantities were used to calculate the standard entropy of formation, ΔS_f°, relative to the elements in their standard states at 298.15 K and 1 atm, and subsequently the standard Gibbs free energy of formation, ΔG_f°, on the same basis.

The formation of a general organic compound of the form $C_xH_yO_z$ from its constituent elements is described

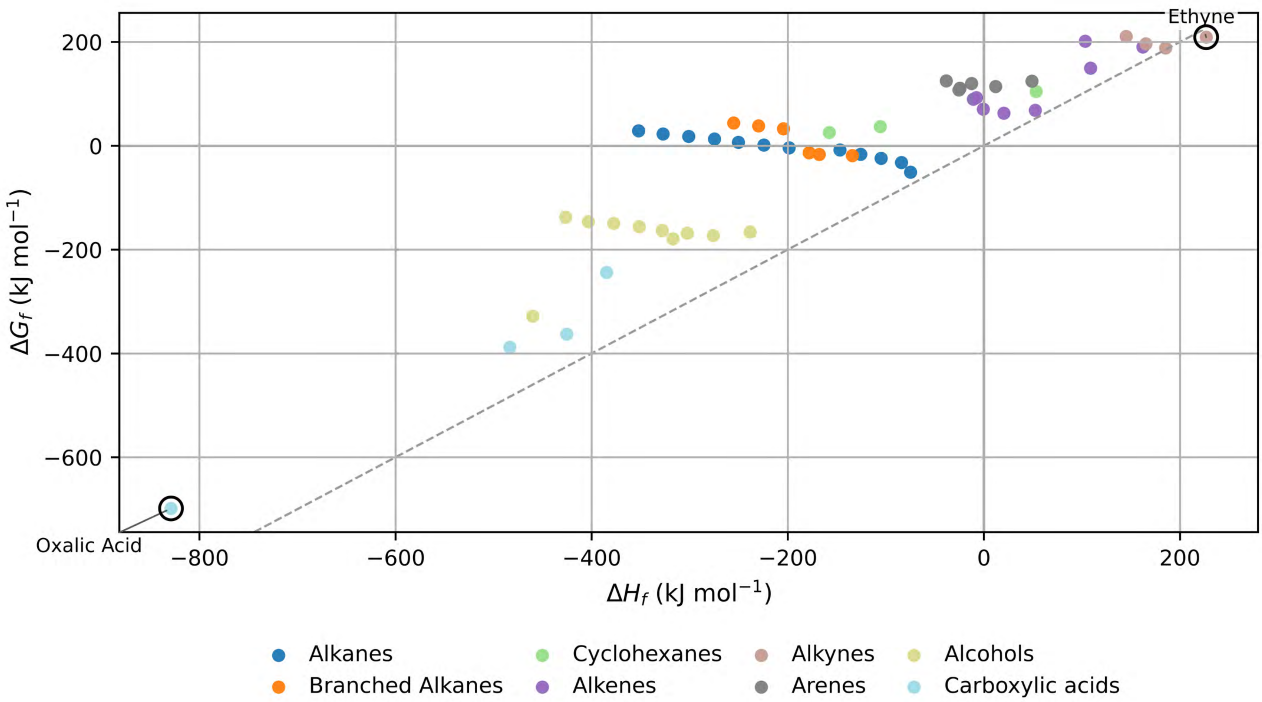

Figure 1. Standard Gibbs free energy of formation, ΔG_f°, versus standard heat of formation, ΔH_f°, for organic compounds grouped by chemical family (kJ mol^{-1}). The dotted line represents 45° line corresponding to $\Delta G_f^\circ = \Delta H_f^\circ$. Data extracted from NIST WebBook [12] and listed in Table S1.

by Eq.1.

$$xC(s) + \frac{y}{2}H_2(g) + \frac{z}{2}O_2(g) = C_xH_yO_z \quad (1)$$

The standard entropy of formation for each compound was calculated following Eq.2. Then standard Gibbs free energy of formation was obtained from Eq. (3), with $T = 298.15\ K$ and entropy expressed in kJ mol^{-1} K^{-1}.

$$\Delta S_F^0 = S_{C_xH_yO_z}^0 - xS_C^0 - \frac{y}{2}S_{H_2}^0 - \frac{z}{2}S_{O_2}^0 \quad (2)$$

$$\Delta G_F^0(C_xH_yO_z) = \Delta H_F^0(C_xH_yO_z) - T\Delta S_F^0 \quad (3)$$

Thermodynamic data for 53 organic compounds were extracted from the NIST WebBook. The compounds span a range of chemical families, including alkanes, alkenes, alcohols, and carboxylic acids. The values of $\Delta H_f^\circ, S^\circ$, and the calculated ΔG_f° are listed in Table S1 of the Supplementary Material. A plot of ΔG_f° versus ΔH_f° for these compounds is shown in Fig. 1.

Although the data exhibit scatter, clear trends are observed within individual chemical families. For example, oxalic acid exhibits one of the lowest values of ΔG_f°, indicating that its formation from the elements is both highly exothermic and strongly spontaneous. In contrast, ethyne exhibits a large positive ΔH_f° and a correspondingly large positive ΔG_f°, implying that its formation from the elements is highly endothermic and would require the input of approximately 200 kJ mol^{-1} of work.

Sempuga et al. [13] introduced a useful representation of thermodynamic feasibility by plotting ΔG versus ΔH, in which the axes correspond to reaction work and energy requirements, respectively. In this representation, compounds lying in the upper-right and lower-left quadrants are particularly attractive for thermocatalytic processes, as the temperature of the heat flow associated with the reaction may carry sufficient equivalent work to drive chemical conversion.

To evaluate combustion thermodynamics, the conversion of carbon dioxide and liquid water to a general organic compound $C_xH_yO_z$ and oxygen was considered:

$$xCO_2 + \frac{y}{2}H_2O(l) = C_xH_yO_z + \left(x + \frac{y}{4} - \frac{z}{2}\right)O_2 \quad (4)$$

This reaction corresponds to the photosynthetic production of organic matter and is equivalent to the reverse of the combustion reaction. When water is in the liquid phase, the heat of combustion corresponds to the higher heating value (HHV) of the compound.

The ΔG_C and ΔH_C were calculated from Eqs. (5) and (6).

$$-\Delta G_C = \Delta G_F(C_xH_yO_z) + \left(x + \frac{y}{4} - \frac{z}{2}\right)\Delta G_F(O_2) - x\Delta G_F(CO_2) - \frac{y}{2}\Delta G_F(H_2O) \quad (5)$$

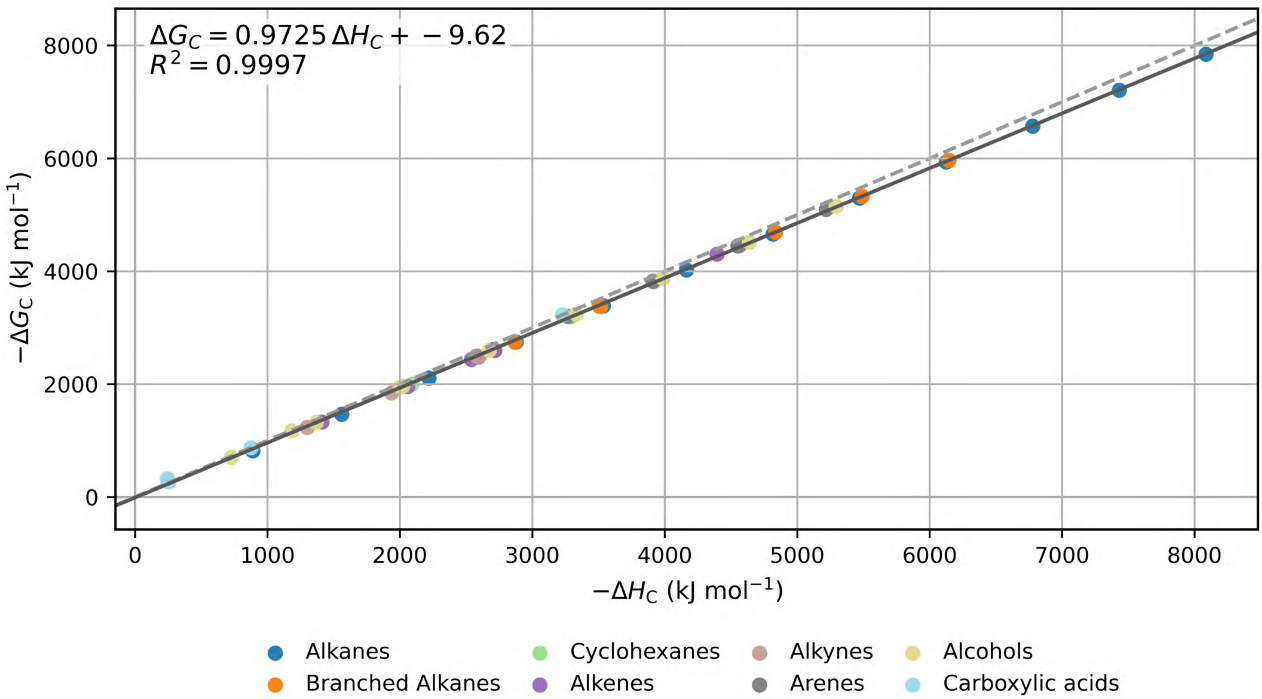

Figure 2. $-\Delta G_C$ versus $-\Delta H_C$.for organic compounds grouped by chemical family. The dotted line represents the 45° line corresponding to $\Delta G_C = \Delta H_C$; the solid line shows the linear regression given by Eq. (7).

$$-\Delta H_C = \Delta H_F(C_xH_yO_z) + \left(x + \frac{y}{4} - \frac{z}{2}\right)\Delta H_F(O_2) - x\Delta H_F(CO_2) - \frac{y}{2}\Delta H_F(H_2O) \quad (6)$$

Using the values of ΔG_f° and ΔH_f° obtained from Eqs. (2) and (3), the quantities $-\Delta G_C$ and $-\Delta H_C$ were calculated for the 53 organic compounds in the dataset, represented in Table S2. A plot of $-\Delta G_C$ and $-\Delta H_C$ is shown in Figure 2. The data collapse onto a linear relationship of the form in Eq. (7) where ΔG_C and ΔH_C are expressed in kJ mol^{-1}. The regression coefficient $R = 0.9997$ indicates the linear model in Eq. (7) gives a very good fit.

$$\Delta G_C = 0.9725\Delta H_C - 9.62 \quad (7)$$

From Fig. 2 it is evident that, for the set of organic compounds studied, which span a wide range of elemental C:H:O ratios, there exists a strong linear relationship between ΔG_C and ΔH_C. The slope is close to unity, and the intercept is small relative to the magnitude of the combustion energies, which typically lie in the hundreds to thousands of kJ mol^{-1}.

This simplified examination of organic compounds holds a mean average percent error, MAPE, of 1.71% with a mean average error, MAE, of 24.53 kJ/mol, proving there is very nominal error with this estimation.

Hypothesis

Based on the observed linear correlation in Fig. 2, it is hypothesized that that relationship described by Eq. (7) applies generally to compounds in the form $C_xH_yO_z$. Accordingly, the Gibbs free energy of combustion, ΔG_C, for any biomass-derived compound containing only carbon, hydrogen and oxygen may be estimated directly from its measured heat of combustion, ΔH_C. Given the elemental C:H:O ratios and estimated value of ΔG_C, the standard Gibbs free energy of formation, ΔG_f°, may then be calculated using Eq. (5), and the corresponding heat of formation, ΔH_f°, using Eq. (6).

Hypothesis Testing

To test the hypothesis that the linear relationship between the Gibbs free energy of combustion, ΔG_C, and the heat of combustion, ΔH_C, applies to biomass-derived compounds of the form $C_xH_yO_z$, published thermodynamic data for cellulose, starch, and glucose were examined. In particular, data for several allomorphic forms of cellulose with chemical formula $C_6H_{10}O_5$, reported by Ioelovich [14], were used. In this work, values of the standard heat of formation, ΔH_f°, were derived from measured heats of combustion, ΔH_C. In addition, published values of ΔH_f° and ΔG_f° for α-D-glucose and potato starch

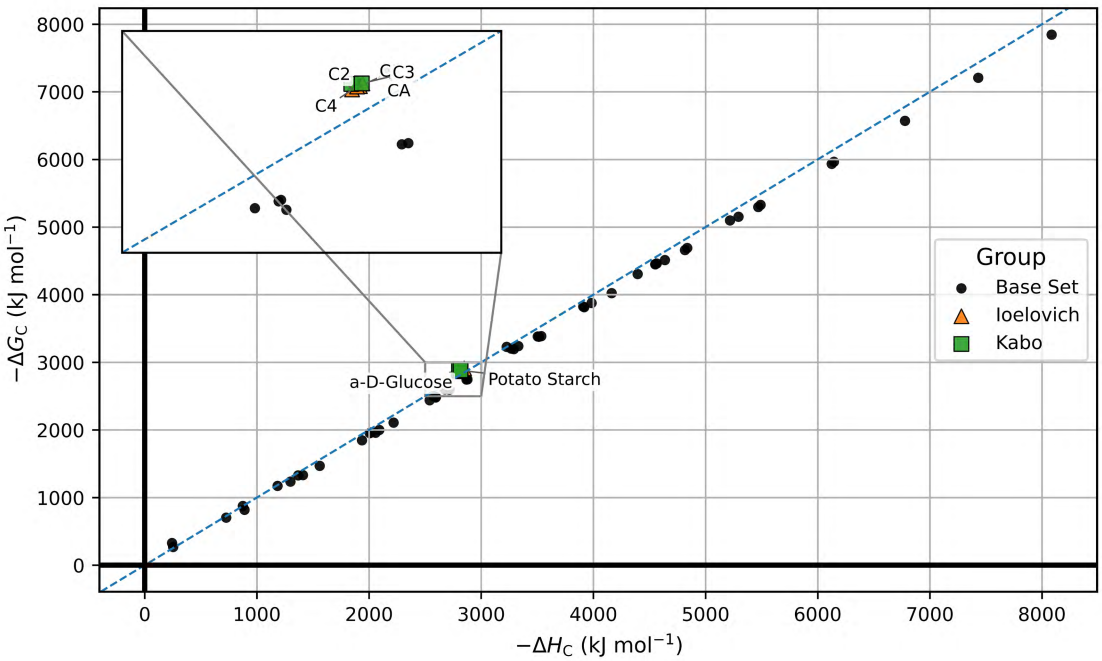

Figure 3. Comparison of $-\Delta G_C^*$ and $-\Delta H_C^*$ for biomass-derived compounds from references [14, 15] with $-\Delta G_C$ versus $-\Delta H_C$ for 53 organic compounds from the NIST database. The dotted line represents the 45° line corresponding to $\Delta G_C = \Delta H_C$.

reported by Kabo et al. [15] were included in the test. The thermodynamic data used in this analysis are summarized in Tables 1 and 2.

Table 1. ΔG_f° and ΔH_f° of formation of cellulose allomorphs as determined by Ioelovich [14].

Cellulose Allomorph	$-\Delta H_f^\circ$ (kJ mol^{-1})	$-\Delta G_f^\circ$ (kJ mol^{-1})
CI	979.6	676.4
CII	986.9	683.7
CIII	973.2	670.0
CIV	980.4	677.2
CA	942.4	642.6

Table 2. ΔG_f° and ΔH_f° of formation of glucose and starch as reported by Kabo *et al.* [15].

Compound	Formula	$-\Delta H_f^\circ$ (kJ mol^{-1})	$-\Delta G_f^\circ$ (kJ mol^{-1})
α-D-Glucose	$C_6H_{12}O_6$	1273.7	908.6
Potato Starch	$C_6H_{10}O_5$	974.1	670.4

Using the tabulated values of ΔH_f° and ΔG_f° from Tables 1 and 2, ΔH_C and ΔG_C were calculated using Eqs. (5) and (6). These calculated quantities are denoted as $-\Delta H_C^*$ and $-\Delta G_C^*$, respectively, and are listed in Table 3.

Table 3. $-\Delta H_C^*$ and $-\Delta G_C^*$ for cellulose allomorphs, glucose, and starch derived from published formation data [14, 15].

Compound	Formula	$-\Delta H_C^*$ (kJ mol^{-1})	$-\Delta G_C^*$ (kJ mol^{-1})
CI Cellulose	$C_6H_{10}O_5$	2813.6	2875.5
CII Cellulose	$C_6H_{10}O_5$	2803.3	2868.16
CIII Cellulose	$C_6H_{10}O_5$	2817.0	2881.9
CIV Cellulose	$C_6H_{10}O_5$	2809.8	2874.7
Amorphous Cellulose	$C_6H_{10}O_5$	2847.8	2909.3
α-D-Glucose	$C_6H_{12}O_6$	2802.34	2880.40
Potato Starch	$C_6H_{10}O_5$	2816.11	2881.46

The values of $-\Delta H_C^*$ and $-\Delta G_C^*$ for the biomass-derived compounds were combined with the 53 organic-compound data points obtained from the NIST WebBook and plotted together in Fig. 3. Both the biomass-derived data and the tabulated organic compounds lie close to the linear relationship described by Eq. (7); although all the biomass-derived data points lie above the trendline. The deviation from the line is generally less than 100 kJ mol^{-1}, which is small relative to the absolute magnitudes of the combustion energies.

This level of agreement indicates that, for the purposes of conceptual thermodynamic analysis and early-

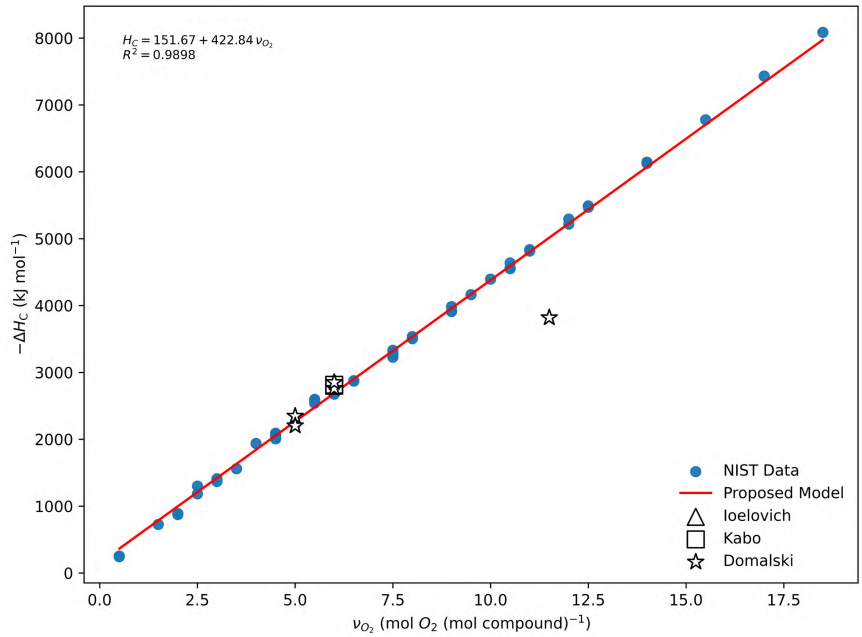

Figure 4. -Heat of combustion, -ΔH_C, versus oxygen demand, ν_{O_2}, for organic and biomass-derived compounds. Circles denote NIST data, triangles Ioelovich data [14], squares Kabo *et al.* data [15], and stars Domalski *et al* [17]. The solid line corresponds to the linear regression.

stage process synthesis, the proposed linear relationship provides sufficiently accurate estimates of ΔG_C from ΔH_C. On this basis, it can be concluded that the hypothesis holds: thus, the value of ΔG_C for biomass-derived compounds of the form $C_xH_yO_z$ can be calculated from ΔH_C using Eq. (7). It then follows that the value of ΔG_f° can be calculated using Eq. (5).

Thornton's Model

Thornton proposed that, for an organic compound of the form $C_xH_yO_z$, the heat of combustion is directly related to the amount of oxygen consumed during combustion [16]. He observed that the heat released per unit mass of oxygen consumed is approximately constant at about 13.1 MJ kg^{-1} O_2, leading to the linear relationship where $\nu_{O_2} = \left(x + \frac{y}{4} - \frac{z}{2}\right)$ is the stoichiometric oxygen demand for complete combustion, as defined in Eq. (4).

$$-\Delta H_C^T = \theta_0 + k_{O_2}\nu_{O_2} \tag{8}$$

The experimentally derived values of ΔH_C were plotted against oxygen demand for three datasets: organic compounds from NIST WebBook, cellulose allomorphs from Ioelovich [14] and starch and glucose from Kabo *et al.* [15]. These results are shown in Fig. 4. A linear regression by least squares was conducted on the data using Python's *sklearn.linear_model* package.

Regression of the combined dataset yields an excellent fit with $R^2 = 0.9898$ in the form of Eq. (9). The fitted slope of 422.84 kJ mol^{-1} O_2 agrees well with Thornton's estimate of approximately 460 kJ mol^{-1} O_2, confirming the applicability of the model to both organic and biomass-derived compounds.

$$-\Delta H_C^T = 151.67 + 422.84\left(x + \frac{y}{4} - \frac{z}{2}\right) \tag{9}$$

Eq. (9) implies that compounds with identical elemental composition will exhibit the same heat of combustion. Accordingly, all cellulose allomorphs examined by Ioelovich [14], which share the formula $C_6H_{10}O_5$, are predicted to have identical ΔH_C. Although starch and glucose differ chemically, their oxygen demand is identical to that of cellulose and therefore fall at the same point in Fig. 4.

Using Eq. (9), the predicted heat of combustion for all seven bio compounds is $-\Delta H_C^T = 2688.72$ kJ mol^{-1}. Comparison with the calculated values $-\Delta H_C^*$ derived from formation data (Table 3) is shown in Table 4. Deviations are on the order of 4-5%, indicating that Thornton's model provides sufficient accuracy for conceptual thermodynamic analysis.

The mean absolute percentage error corresponds to a mean absolute error of 78.82 kJ mol^{-1}. For compounds with large combustion energies, this results in small percentage errors. For compounds with lower ΔH_C, percentage errors increase. For example, oxalic acid ($(COOH)_2$) exhibits a 48.9% error, whereas 1-4-dimethylbenzene (C_8H_{10}) shows an error of only 0.85%. This indicates that

the proposed framework is most suitable for biomass-derived and other high-molecular-weight organic compounds.

Table 4. Comparison calculated heats of combustion, $-\Delta H_C^*$, with Thornton-model predictions $-\Delta H_C^T$.

Compound	$-\Delta H_C^*$ (kJ mol^{-1})	$-\Delta H_C^T$ (kJ mol^{-1})	Percent Error (%)
CI	2813.6	2688.72	4.44
CII	2803.3	2688.72	4.09
CIII	2817.0	2688.72	4.55
CIV	2809.8	2688.72	4.31
CA	2847.8	2688.72	5.59
α-D-Glucose	2802.34	2688.72	4.05
Potato Starch	2816.11	2688.72	4.52

Comparing HHV Data from Domalski *et al.* with Thornton Model

The applicability of Thornton's model to biomaterials was further examined using biomass compounds for which only higher heating values (HHV) are available. Domalski *et al.* [17] reported HHV values for several biomass-derived compounds (Table 5). These values were converted to heats of combustion using molecular weights derived from elemental composition, assuming pure components.

Table 5. HHV and chemical composition of select biomass compounds from Domalski *et al* [17].

Compound	Formula	HHV (BTU lb^{-1})
Cellulose	$C_6H_{10}O_5$	7325
Lignin	$C_{10}H_{12}O_3$	9111
Starch	$C_6H_{10}O_5$	7560
Xylan	$C_5H_8O_4$	7637
Hemicellulose	$C_5H_8O_4$	7170

When plotted in Fig. 4, all compounds except lignin follow the Thornton correlation closely. Using the nominal formula $C_{10}H_{12}O_3$, the Thornton model predicts $-\Delta H_C^T = 5014$ kJ mol^{-1}, whereas Domalski *et al.* report a value equivalent to 3819 kJ mol^{-1}. The deviation for lignin could be attributed to its structural complexity and compositional heterogeneity, leading to ν_{O_2} not being estimated correctly.

Relationship between $\Delta H_f^\circ, \Delta G_f^\circ$ and Oxygen Demand

By combining the linear relationship between ΔG_C and ΔH_C established in Eq. (7) with Thornton's oxygen-demand model Eq. (9), a simple framework is obtained for estimating the standard heats and Gibbs free energies of formation of organic and biomass-derived compounds of the form $C_xH_yO_z$.

Combining Eq. (6) and (9), the standard heat of formation may be expressed directly as a function of elemental composition:

$$\Delta H_F^0 = 151.67 + 29.33x - 37.21y - 211.42z \quad (10)$$

Similarly, combining Eqs. (7) and (9) yields an expression for the Gibbs free energy of combustion:

$$-\Delta G_C = 137.88 + 411.21x + 102.8y - 205.6z \quad (11)$$

Substitution of Eq. (11) into the combustion-free-energy balance, Eq. (5), allows the standard Gibbs free energy of formation to be estimated as:

$$\Delta G_F^0 = 147.5 + 16.65x - 15.77y - 205.6z \quad (12)$$

The units of enthalpy and Gibbs free energy are kJ mol^{-1}. Eqs (10) and (12) therefore provide a closed, composition-based method for estimating both ΔH_F^0 and ΔG_F^0 for any compound containing only carbon, hydrogen and oxygen.

Estimation of Enthalpy and Gibbs Free Energy of Reaction

Once formation properties are known, enthalpies of reaction, ΔH_{rxn} and Gibbs free energy of reaction, ΔG_{rxn} can be estimated for reactions between any organic or biomass-derived compounds. Consider the anaerobic conversion of reaction $C_{x1}H_{y1}O_{z1}$ to a product $C_{x2}H_{y2}O_{z2}$. Writing the reaction in terms of carbon dioxide and liquid water balance yields, in which $\nu_{O_{2,i}}$ is defined in Eq. (14):

$$C_{x1}H_{y1}O_{z1} + \left(\frac{\nu_{O_{2,1}}}{\nu_{O_{2,2}}}x_2 - x_1\right)CO_2 + \frac{1}{2}\left(\frac{\nu_{O_{2,1}}}{\nu_{O_{2,2}}}y_2 - y_1\right)H_2O = \frac{\nu_{O_{2,1}}}{\nu_{O_{2,2}}}C_{x2}H_{y2}O_{z2} \quad (13)$$

$$\nu_{O_{2,i}} = \left(x_i + \frac{y_i}{4} - \frac{z_i}{2}\right) \quad (14)$$

This can then be used to determine ΔH_{rxn} and ΔG_{rxn} by the following Eqs:

$$\Delta H_{rxn} = -\Delta H_{C,1} + \frac{z_1 - 2x_1 - \frac{y_1}{2}}{z_2 - 2x_2 - \frac{y_2}{2}}\Delta H_{C,2} \quad (15)$$

$$\Delta G_{rxn} = -\Delta G_{C,1} + \frac{z_1 - 2x_1 - \frac{y_1}{2}}{z_2 - 2x_2 - \frac{y_2}{2}}\Delta G_{C,2} \quad (16)$$

Connection to Process Synthesis and Estimation of Minimum Operating Costs

For a process exchanging a quantity of heat Q_0 at ambient temperature T_0 and work W with its surroundings, we can show that:

$$\Delta H_{rxn} = Q_0 + W \quad (17)$$

For a reversible process, the minimum reversible work, W_{rev}, requirement is given by Eq. (18) while for a real, irreversible process, Eq. (19) holds:

$$W_{rev} = \Delta G_{rxn} \quad (18)$$

$$W = W_{rev} + W_{lost} \quad (19)$$

Thus, estimates of ΔG_{rxn} provide direct insight into the minimum work requirements for a process and allow for estimation of irreversibilities if the work consumption of an actual process is known. In electrochemical and electrically driven processes, where operating costs are typically dominated by electricity consumption, the minimum operating cost can be calculated from the cost of a quantity of electricity ΔG_{rxn}. Large deviations between actual work input and ΔG_{rxn} indicate significant losses arising from reaction pathways, or equipment design and operation.

Chemical Reaction Equilibrium Analysis

The equilibrium constant at temperature T_0 is given by:

$$K_{eq}(T_0) = \exp\left(\frac{\Delta G_{rxn}(T_0)}{RT_0}\right) \quad (20)$$

Accordingly, Eq. (16) may be used to estimate equilibrium constants for reactions involving organic and biomass-derived compounds. The temperature dependence of $K_{eq}(T)$ can then be evaluated using the Gibbs-Helmholtz equation, with ΔH_{rxn} obtained from Eq. (15).

Scope and Use

The simplified composition-based estimation of ΔH_F^0 and ΔG_F^0 developed here is particularly valuable for early-stage process analysis and synthesis. It enables rapid assessment of reaction driving forces, equilibrium limitations, minimum work requirements and potential operating costs before detailed kinetic, separation, and flowsheet models are available. The only information that is required is the elemental composition of the components.

DISCUSSION

A central contribution of the proposed methodology is the establishment of a method for estimating ΔH_f^o and ΔG_f^o for biomass-derived compounds of the form $C_xH_yO_z$. At present, reported thermodynamic properties for biomass materials span wide ranges, reflecting differences in feedstock origin, crystallinity, polymorphism, moisture content, and experimental methodology. While such variability is important for detailed material characterization, it complicates early-stage techno-economic and thermodynamic analyses, where the objective is to assess feasibility, energy requirements, and fundamental limits rather than to resolve fine material distinctions. For many other biomass-derived compounds, reliable values of Gibbs free energy of formation are simply unavailable due to the difficulty of direct measurement. The framework proposed here provides a physically grounded reference point that is sufficiently accurate for conceptual analysis of heat and work requirements.

From a process systems perspective, the ability to estimate both ΔH_f° and ΔG_f° on a consistent basis is particularly valuable. Many process analysis tasks, including energy integration, lost-work analysis, reaction screening, and equilibrium assessment, require at least one of these properties. When different compounds within a flowsheet rely on inconsistent data sources or assumptions, the resulting calculations can be internally inconsistent and misleading. By expressing formation properties directly in terms of elemental composition, the proposed approach enables thermodynamically consistent calculations across all compounds in a system, facilitating reliable comparison of alternative process pathways.

The integration of the linear correlation between ΔH_C and ΔG_C with Thornton's oxygen-demand model provides a clear physical basis for the method. The strong linearity observed across a broad set of organic compounds, and its extension to biomass-derived materials, supports the hypothesis that combustion thermodynamics are dominated by elemental composition rather than molecular structure at a first-order level. The deviations observed for complex materials such as lignin highlight important practical considerations, most notably sample heterogeneity and hydration, but do not undermine the utility of the approach for early-stage analysis. Indeed, the ability to flag possible inconsistencies in experimental measurements could aid in identifying issues related to sample heterogeneity, incomplete drying, or other experimental issues.

From the standpoint of process synthesis and optimization, the principal value of the methodology lies in its connection to minimum work requirements and reaction driving forces. Estimates of ΔG_{rxn} provide direct insight into reversible work limits, equilibrium constants, and the magnitude of unavoidable irreversibilities. In electrically driven and electrochemical processes, where operating costs are dominated by electricity consumption, these estimates translate directly into lower bounds on operating cost. Conversely, large discrepancies between actual work input and estimated ΔG_{rxn} signal opportunities for improved reaction pathways, flowsheets or equipment design. In this sense, the proposed framework enables lost-work targeting for biomass systems that would otherwise be inaccessible due to missing thermodynamic data.

The methodology is subject to several limitations that define its intended scope of applicability. It is formulated for compounds containing only carbon, hydrogen, and oxygen, and does not explicitly account for nitrogen, sulphur, or other heteroatoms. The linear correlation observed between ΔG_C and ΔH_C in Fig. 2 is therefore expected to apply primarily to CHO compounds. In addition, the method does not distinguish between structural isomers or polymorphs, assigning identical thermodynamic properties to compounds with the same elemental

composition. While this limits its use for detailed materials characterization, it may be appropriate for early-stage process synthesis and conceptual thermodynamic analysis, where such distinctions typically have a limited impact on feasibility assessments and comparative process screening.

CONCLUSION

A composition-based methodology has been developed for estimating ΔH_f° and ΔG_f° of biomass-derived compounds containing C, H and O. By combining an observed linear relationship between the ΔH_C and ΔG_C with Thornton's oxygen-demand model, the approach enables consistent estimation of formation properties using widely available combustion data. Validation against published data for cellulose, starch, glucose, and related compounds indicates that the method provides sufficient accuracy for conceptual thermodynamic analysis.

The proposed framework supports early-stage process synthesis and screening by enabling rapid estimation of reaction enthalpies, minimum work requirements, equilibrium constants, and reaction driving forces in systems where reliable thermodynamic data are otherwise unavailable. While the method is limited to CHO compounds and does not account for structural or polymorphic effects, it provides a practical and internally consistent basis for evaluating the thermodynamic feasibility of biomass conversion and CO_2 utilization pathways prior to detailed flowsheet development.

DIGITAL SUPPLEMENTARY MATERIAL

Digital supplementary material, including model information, can be found at: https://psecommunity.org/LAPSE:2026.0036.

ACKNOWLEDGEMENTS

This material is based upon work supported by the National Institute of Food and Agriculture, U.S. Department of Agriculture (USDA), through the Northeast Sun Grant Initiative under award number #S005193-USDA.

AUTHOR IDENTIFIERS

Author ORCIDs:
Anastasi A: 0009-0006-2821-5523
Masuku CM: 0000-0003-0556-4588
Chundawat, SPS: 0000-0003-3677-6735
Hildebrandt D: 0000-0001-7873-8855

REFERENCES

1. Chou CC, Shih YS. A thermodynamic approach to the design and synthesis of plant utility systems. Ind. Eng. Chem. Res. 26:1100-1108 (2002). https://doi.org/10.1021/ie00066a009
2. Lopez-Arenas T, Mansouri SS, Sales-Cruz M, Gani R, Pérez-Cisneros ES. A gibbs energy-driving force method for the optimal design of non-reactive and reactive distillation columns. Computers & Chemical Engineering 128:53-68 (2019). https://doi.org/10.1016/j.compchemeng.2019.05.024
3. Hammond GP, Mansell RVM. A comparative thermodynamic evaluation of bioethanol processing from wheat straw. Applied Energy 224:136-146 (2018). https://doi.org/10.1016/j.apenergy.2018.04.123
4. Ioelovich, M. Application of Thermodynamic Methods to the Study of Plant Biomass and Its Components-A Review. *Appl Biosci*. 3:577-616 (2024) https://doi.org/10.3390/applbiosci3040036
5. Goldberg RN, Schliesser J, Mittal A, Decker SR, Santos AF, Freitas VLS, et al. A thermodynamic investigation of the cellulose allomorphs: Cellulose(am), cellulose Iβ(cr), cellulose II(cr), and cellulose III(cr). *J Chem Thermodyn* 81:184-226 (2015) https://doi.org/10.1016/j.jct.2014.09.006
6. Jeffri NI, Mohammad Rawi NF, Mohamad Kassim MH, Abdullah CK. Unlocking the potential: evolving role of technical lignin in diverse applications and overcoming challenges. International Journal of Biological Macromolecules 274:133506 (2024). https://doi.org/10.1016/j.ijbiomac.2024.133506
7. Yoo CG, Ragauskas AJ. Opportunities and challenges of lignin utilization. ACS Symposium Series :1-12 (2021). https://doi.org/10.1021/bk-2021-1377.ch001
8. Wang Y, Zhu Y, Li M, Ji B, Liu S, Liu W, Hu Z, Zhang H, Zhang W, Peng L, Kong Z. Distinct morphology and self-assembly behavior of xylan nanoparticles are correlated with arabinose side-chains in sorghum. Carbohydrate Polymers 366:123874 (2025). https://doi.org/10.1016/j.carbpol.2025.123874
9. Friedl A, Padouvas E, Rotter H, Varmuza K. Prediction of heating values of biomass fuel from elemental composition. Analytica Chimica Acta 544:191-198 (2005). https://doi.org/10.1016/j.aca.2005.01.041
10. Esteves B, Sen U, Pereira H. Influence of chemical composition on heating value of biomass: a review and bibliometric analysis. Energies 16:4226 (2023). https://doi.org/10.3390/en16104226
11. Blokhin AV, Voitkevich OV, Kabo GJ, Paulechka YU, Shishonok MV, Kabo AG, Simirsky VV. Thermodynamic properties of plant biomass components. heat capacity, combustion energy,

and gasification equilibria of cellulose. J. Chem. Eng. Data 56:3523-3531 (2011). https://doi.org/10.1021/je200270t
12. Linstrom, P.J., Mallard, W.G. NIST Chemistry WebBook. In: NIST Standard Reference Database Number 69. Ed: Linstrom, P.J., Mallard, P.J. National Institute of Technology.
13. Sempuga BC, Hildebrandt D, Patel B, Glasser D. Work to chemical processes: the relationship between heat, temperature, pressure, and process complexity. Ind. Eng. Chem. Res. 50:8603-8619 (2011). https://doi.org/10.1021/ie2004785
14. Ioelovich, M. Study of thermodynamic properties of various allomorphs of cellulose. *ChemXpress* 9(3):259-265 (2016).
15. Kabo GJ, Voitkevich OV, Blokhin AV, Kohut SV, Stepurko EN, Paulechka YU. Thermodynamic properties of starch and glucose. The Journal of Chemical Thermodynamics 59:87-93 (2013). https://doi.org/10.1016/j.jct.2012.11.031
16. Thornton WM. XV. the relation of oxygen to the heat of combustion of organic compounds. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 33:196-203 (2009). https://doi.org/10.1080/14786440208635627
17. Domalski, E.S., Jobe, T.L. Jr, Milne, T.A. Thermodynamic Data for Biomass Conversion and Waste Incineration. National Bureau of Standards (1986).

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Municipal Solid Waste Valorization into Chemical Solvents for Industrial Symbiosis: Techno-Economic and Environmental Assessment

Oktay Boztaş*[a], Daniel A. Flórez-Orrego[ab], Meire E. G. R. Domingos[a] and François Maréchal[a]

[a] IPESE EPFL, Federal Polytechnic School of Lausanne, Rue de L'Industrie 17, Valais/Wallis, Switzerland
[b] Faculty of Minas, National University of Colombia, Av. 80 #65 - 223 Medellín, Colombia
* Corresponding Author: oktay.boztas@epfl.ch.

ABSTRACT

Waste incineration with combined heat and power (CHP) supplies electricity and heat to cities and industrial clusters but remains a significant source of greenhouse gas emissions. This work develops an optimization-based, system-level decarbonization framework for an integrated waste-to-energy and chemical production cluster under operational, societal, and economic constraints. The framework is applied to a real-world case study including a municipal waste incineration plant, an energy utility system, and multiple chemical production units. A layered decarbonization strategy is implemented. First, energy efficiency is enhanced through waste heat valorization using heat pumps. Second, coordination between industrial actors is improved through solid waste storage management and operational alignment of heat and power supply with demand. Third, alternative feedstocks are introduced to reduce fossil-based inputs. Within the work material and heat-stream inventories are collected, and the required chemical synthesis pathways are modeled using Aspen Plus. The integrated system is optimized within the OSMOSE framework using a multi-period mixed-integer linear programming formulation that minimizes the annualized total cost while accounting for seasonal heat demand and electricity price variability. The results show that waste-heat upgrading alone enables fossil emission reductions of up to 22% while improving overall system profitability. More broadly, up to 80% of fossil-based Scope 1 and Scope 2 CO_2 emissions can be reduced through energy efficiency improvements, system integration, and feedstock substitution prior to carbon capture deployment. To address the remaining hard-to-avoid emissions and enable net-negative operation, post-combustion carbon capture options are evaluated in combination with solid oxide electrolysis cells (SOEC).

Keywords: Industrial symbiosis, Waste valorization, Waste heat, Carbon capture, Process integration

INTRODUCTION

In Switzerland, the most common practice for treating municipal solid waste is incineration. Approximately 30 waste incineration plants are currently in operation, treating about 4 Mt/yr of waste and accounting for roughly 5% of national CO_2 emissions [1]. A large share of these facilities is connected to district heating networks, while a smaller number directly supplies steam to nearby industrial zones. Although this configuration reduces additional fossil fuel combustion for steam generation and residential heating, current systems often fail to fully exploit the efficiency potential of the existing infrastructure. As a result, significant opportunities remain for optimization, particularly through improved valorization of industrial condensate using heat pumps and enhanced coordination between actors to further reduce fossil fuel use. In addition, waste can be used as a feedstock for the production of basic chemicals for nearby industries, extending industrial symbiosis beyond energy exchange to include material flows. This can further reduce the fossil dependency of the chemical sector, which is currently estimated to rely on fossil-based feedstocks for approximately 95% of its raw material input.

This work provides a system-level analysis involving multiple actors, combining process modeling, integration,

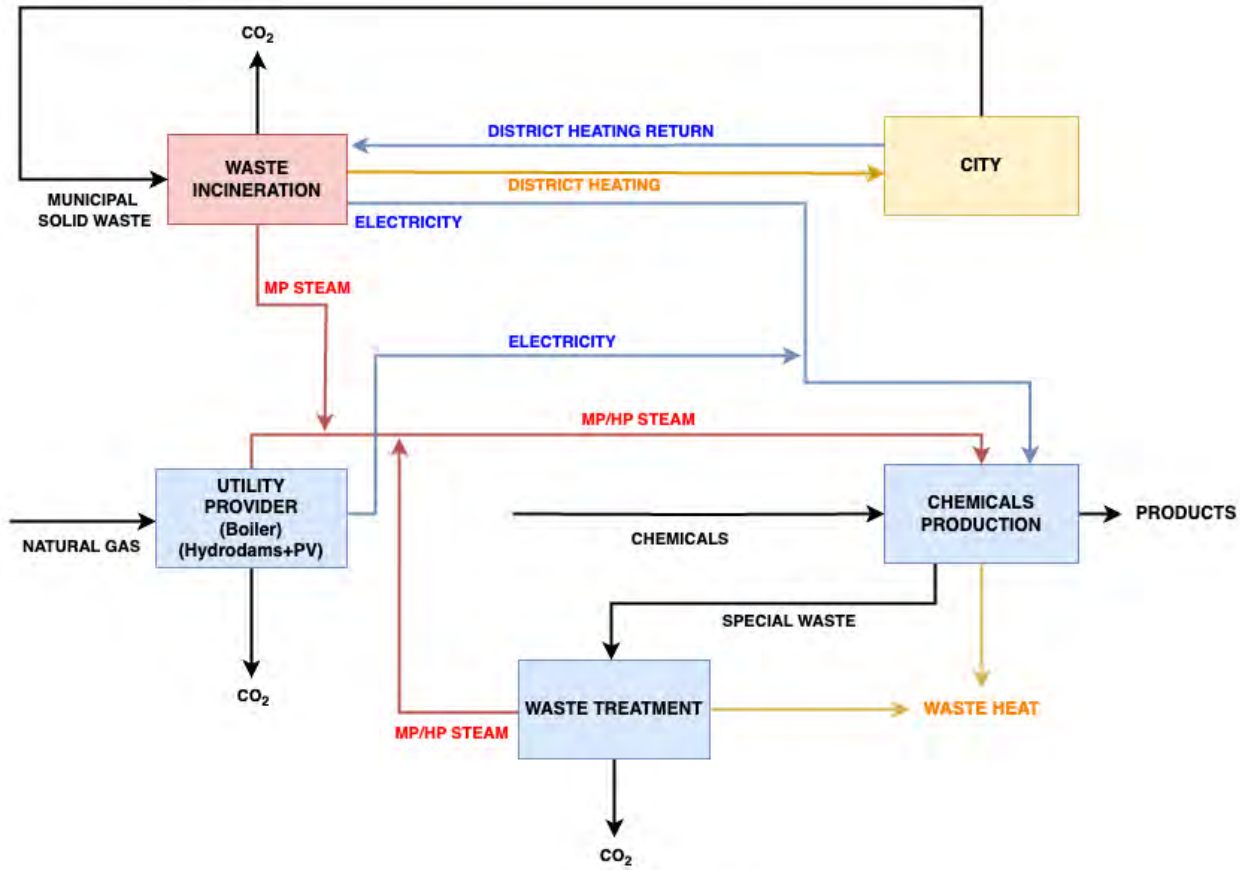

Figure 1. System boundary and energy–material interactions between system actors

and optimization to improve energy efficiency and reduce fossil-based emissions while explicitly accounting for economic constraints. A layered decarbonization approach is adopted, prioritizing measures that are easier to implement and adapt in industrial practice, thereby reducing reliance on and challenges associated with large-scale carbon capture deployment.

METHODS

System Boundary and Actors

Within the defined system boundaries shown in Figure 1, several industrial actors are considered. The first is the waste-to-energy (WtE) plant, which treats municipal solid waste (MSW) and supplies energy in the form of medium-pressure (MP) steam and electricity to the chemical processes, as well as hot water to the district heating network. The second actor is the utility company, which provides heat, power, and special waste treatment services to the chemical processes. Its available assets within the system include a natural gas boiler for the production of MP and high-pressure (HP) steam, waste incineration units capable of supplying MP and HP steam, and renewable power generation units, including hydropower dams and photovoltaic (PV) fields. The third actor consists of the chemical production units, which require electricity, steam, and basic chemical feedstocks, and produce high-value chemicals for the market. These units also benefit from cooling services provided by the utility company, generating low-grade waste heat that can be upgraded and valorized within the system. The final actor is the city, which requires heat for most of the year due to local climatic conditions, supplies municipal solid waste to the WtE plant, and acts as a net importer of electricity, with limited photovoltaic installations and hydropower capacity.

Although a certain degree of industrial symbiosis already exists between these actors, their current level of coordination remains far from optimal. As a result, significant opportunities for improving energy efficiency and reducing fossil-based carbon emissions remain unexploited. This work addresses these gaps by strengthening existing energy distribution and material exchange links and introducing new integration pathways within the system. The coordination of actors is optimized using a system-level approach to achieve the decarbonization objectives outlined above.

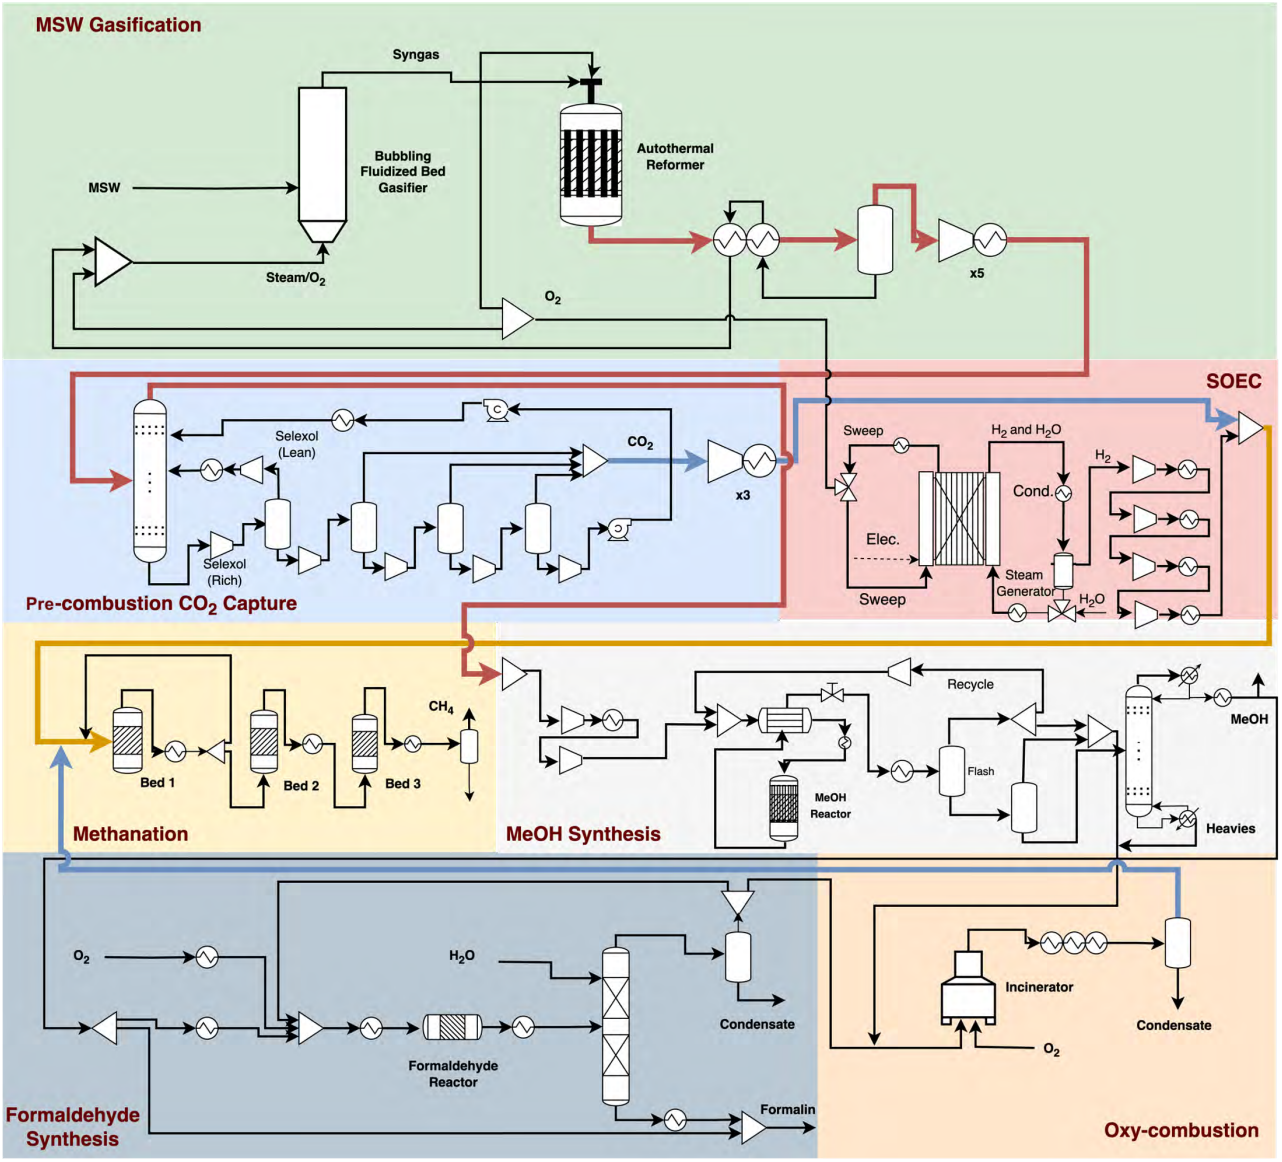

Figure 2. Superstructure of chemical synthesis pathways considered in the study, including municipal solid waste gasification for syngas production, pre-combustion carbon dioxide capture, H_2 production with SOEC, oxy-combustion for CO_2 capture, and downstream synthesis of methanol, formaldehyde, and synthetic natural gas.

Energy Demand and Supply in the System

In the current reference state of the system, approximately 67% of the total heat demand for the chemical processes is supplied as MP steam, while the remaining demand is covered by HP steam. Around half of the MP steam demand is met by the WtE facility, with the remainder supplied by the NG boiler operated by the utility company. For HP steam, the dominant share is provided by the incineration of special waste, accounting for approximately 80% of the supply, while the remaining share is supplied by the NG boiler. All heat demand of the city is covered by the WtE facility and represents approximately one fifth of the heat demand of the chemical processes.

In terms of power distribution, approximately 40% of the electricity consumed by the utility company and the chemical processes is supplied by the WtE facility. Around 8% is generated by PV installations, while the remaining share is provided by hydropower dams, with the possibility of electricity injection into the grid in the case of excess generation.

As a consequence of the current heat and power generation and distribution, approximately 65% of the total CO_2 emissions originate from the WtE facility, of which about half corresponds to biogenic carbon. Around 25% of the CO_2 emissions are associated with the incineration

of special waste, while the remaining 10% result from the combustion of natural gas for steam supply.

Process Modeling of Chemical Pathways

MSW can be valorized beyond direct incineration by converting it into syngas for the production of basic chemicals and by recovering high-grade heat for industrial applications. In this context, this study considers a superstructure of chemical synthesis pathways that enables the integrated valorization of material and energy streams within the overall system, as shown in Figure 2.

For the gasification of MSW, a model of bubbling fluidized gasifier using oxy-steam as the gasification agent, operated at 750 °C and 1 bar was used based on the study of Sebastiani et al. [2]. The chemical properties of the waste used is given in Table 1.

Table 1. Proximate and ultimate analysis of MSW (fuel as received)[2]

Composition		MSW
Proximate Analysis [%]	M	14.6
	FC	8.90
	VM	64.7
	Ash	11.8
Ultimate Analysis [%]	C	41.8
	H	5.05
	N	3.02
	O	23.4
Net calorific value [MJ/kg]		19.9

The produced syngas passes through an adiabatic autothermal reformer and is subsequently cooled to condense water. The resulting syngas is processed in an selexol-based carbon capture unit operating at 30 bar with 90% capture efficiency, adjusting the composition to the stoichiometric number (SN) of slightly above 2.0 required for methanol synthesis [3]. The resulting syngas after the upgrading has a molar composition of 68.3% H_2, 31.3% CO, 0.4% CO_2.

$$SN = (x_{H_2} - x_{CO_2})/(x_{CO} + x_{CO_2})$$

The methanol synthesis process is operated isothermally at 210 °C and 90 bar, based on Lurgi's Mega-Methanol® technology, achieving a mass yield of 95.8% [4]. The produced methanol can either be used as a base chemical within the system or converted to formaldehyde via the Formox® process, which is based on the catalytic oxidation of methanol [5]. The reaction is operated at 250 °C and 1 bar. Following formaldehyde production, the product is mixed with methanol and water to obtain a formalin solution with a composition of 38% formaldehyde, 14% methanol, and 48% water on a molar basis.

The CO_2 captured during syngas upgrading and from the oxy-combustion of purge streams is utilized in a methanation process coupled with SOEC, following the approach proposed by Flórez-Orrego et al [6].

Definition of Decarbonization Options

Following the layered decarbonization strategy, five system configurations are defined and analyzed. Each configuration progressively extends the level of integration, flexibility, and decarbonization options considered in the optimization framework.

- **Base**: The base configuration represents the reference system as defined in the previous sections, without additional flexibility or decarbonization measures.
- **A**: In this configuration, waste heat available from chemical synthesis processes and special waste incineration is valorized through the integration of decentralized and centralized steam-generating heat pumps. WtE is operated considering historical trends.
- **B**: Building on Configuration A, the waste-to-energy (WtE) facility is allowed for operational flexibility through the management of seasonal solid waste storage. The facility can adjust the allocation between steam, electricity, and district heating production in response to the demands of other system actors
- **C:** In addition to the features of Configuration B, waste heat from chemical synthesis and special waste treatment is further upgraded via heat pumps and allowed to supply the district heating network, thereby introducing competition with district heating generated by the WtE facility.
- **D**: The final configuration extends Configuration C by integrating a municipal solid waste (MSW) gasifier. The produced syngas is used as a feedstock for the synthesis of basic chemicals (i.e., methanol and formalin) at the chemical site. In addition, pre-combustion CO_2 captured from the gasification process is combined with hydrogen produced by a solid oxide electrolysis cell (SOEC) to increase chemical yields or enable the production of synthetic natural gas.

Each configuration is also evaluated with the integration of post-combustion carbon capture based on MEA for the capture of all the remaining CO_2 emissions within the system boundaries. Configurations including carbon capture are denoted by a prime symbol (') throughout this study.

Optimization Framework and Problem Formulation

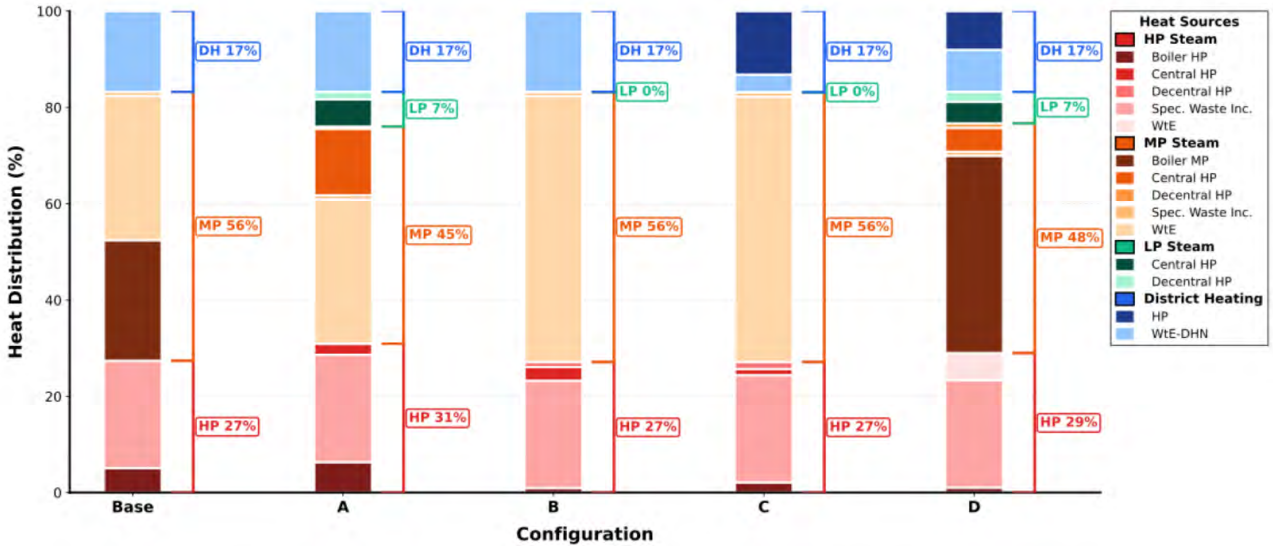

Figure 3a. Heat distribution by steam pressure level and heat source across the different decarbonization configurations.

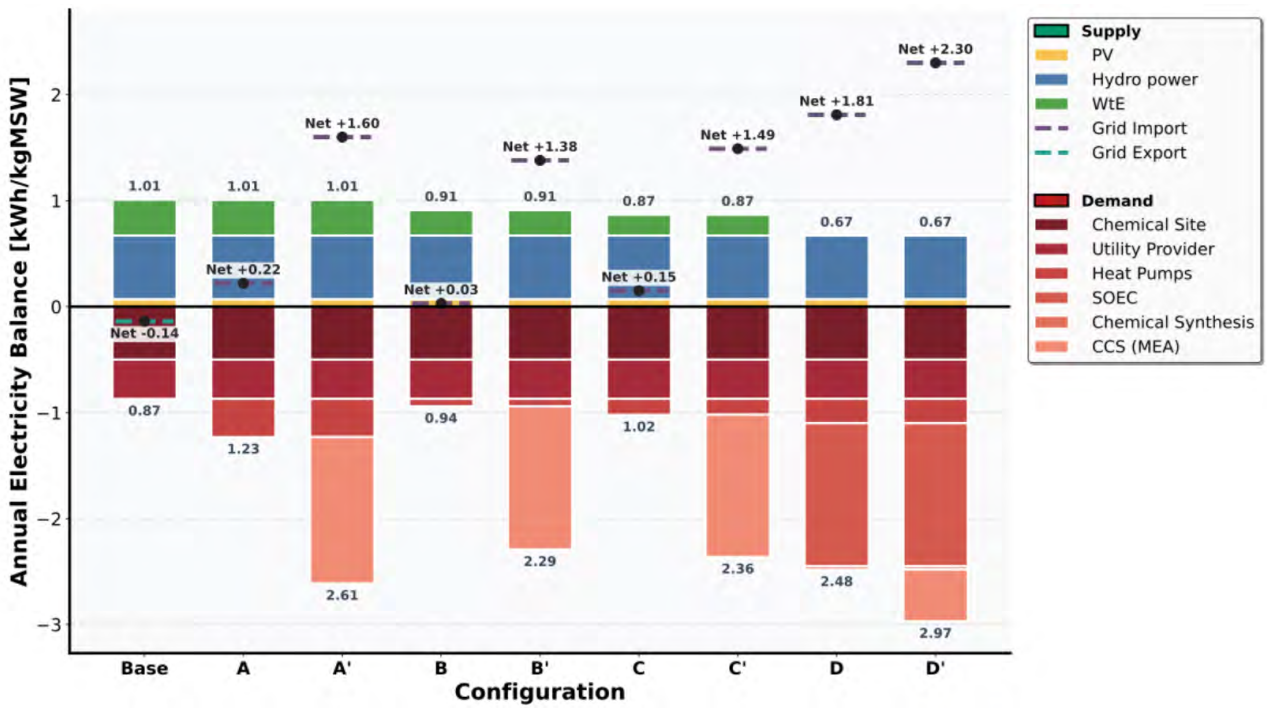

Figure 3b. Annual electricity balance of the integrated system across the different decarbonization configurations, showing electricity supply and demand contributions.

The proposed configurations are integrated and optimized within the OSMOSE Lua platform to identify cost-optimal operating conditions for the superstructure, as defined by the objective function presented below [7].

$$\min Cost_{tot} = \sum_{u \epsilon Units} \left[\sum_{t \epsilon Time} [(Cop1_u \bullet Y_{u,t} + Cop2_u \bullet M_{u,t})\Delta t + Cinv1_u \bullet Y_u + Cinv2_u \bullet M_u] \right]$$

- $Cop1_u$ [EUR/h]: fixed operating cost of unit u,
- $Cop2_u$ [EUR/kWh or EUR/kg]: variable operating cost of unit u,
- $Y_{u,t}$ [-]: integer variable standing for the activation of the unit u at a time t,
- $M_{u,t}$ [kg/h or kW]: continuous variable standing for the load of the unit u at a time t,

- $Cinv1_u$ [EUR/y]: fixed annualized cost of unit u,
- $Cinv2_u$ [EUR/kWh or EUR/kg]: variable annualized cost of unit u,
- Y_u [-]: integer variable for the activation of the unit u, maximum value of $Y_{u,t}$,
- M_u [kg/y or kWh/y]: maximum load of the unit.

The problem is solved using a multi-period formulation with 10 representative clusters, capturing steam demand of chemical processes, district heating requirements, MSW supply to WtE, and electricity/natural gas prices based on the Swiss spot market [8]. The electricity supply is considered to be fully decarbonized based on the real-case arrangements defined within the system boundaries. A Scope 2 emission factor of 0.05 $kgCO_2/kWh$ is applied to the natural gas supply.

RESULTS AND DISCUSSION

This section presents and discusses the results of the OSMOSE optimization framework for the different configurations, focusing on energy distribution, fossil CO_2 emissions, and the resulting economic trade-offs.

Energy Distribution and Electricity Demand

The heat distribution within the system for each configuration is shown in Figure 3a. In Configuration A, where only heat pump integration is introduced at temperature levels of 40 °C and 70 °C, the boiler supplying MP steam is no longer required. In contrast, the HP boiler remains in operation due to the higher temperature lift required. As a result, approximately 90% of the available waste heat is upgraded to LP and MP steam levels, while only about 10% is upgraded to HP steam, primarily during periods of low electricity spot prices.

In Configuration C, when the WtE facility is allowed to better align its operation with the other system actors and waste heat can be supplied to the district heating network, it becomes the sole provider of MP steam. At the same time, its direct contribution to district heating is reduced by approximately 80%. Given the high coefficient of performance of the heat pumps supplying the district heating network (COP = 4.3), the system favors cascading heat through the chemical processes while using the available waste heat primarily for meeting the city's heating demand. As a result, only about 18% of the available waste heat is utilized for steam generation.

Finally, in Configuration D, the system enables the production of synthetic natural gas, and the heat released from syngas cooling, exothermic methanol and formaldehyde synthesis, and purge stream combustion is recovered and reused for steam generation for the gasifier and the SOEC. As a consequence, the contribution of waste treatment to steam supply is reduced and the MP boiler is reactivated. In this configuration, approximately 60% of the upgraded waste heat is allocated to steam generation, while only about half of the city's heating demand is supplied via heat pumps.

As shown in Figure 3b, at its current electricity generation capacity the system is not able to remain fully self-sufficient and therefore requires electricity imports across all configurations. In Configuration A, the installation of heat pumps leads to an increase in electricity consumption of 0.36 $kWh_{EE}/kgMSW$. However, when improved alignment between system actors is introduced in Configuration B, the system reaches near self-sufficiency, with electricity imports limited to only 0.03 $kWh_{EE}/kgMSW$. The highest electricity consumption and imports are observed in Configuration D, as the SOEC installation requires 1.35 $kWh_{EE}/kgMSW$ and electricity generation from the WtE turbine is no longer available. As the heat demand of the MEA carbon capture unit is electrified, electricity consumption increases by up to 1.38 $kWh_{EE}/kgMSW$ in Configuration B′. In contrast, this increase is limited to 0.49 $kWh_{EE}/kgMSW$ in Configuration D′, since a larger fraction of the carbon contained in the MSW is retained within chemical products, thereby reducing the required size of the MEA carbon capture unit

CO_2 Emissions Performance

With the integration of heat pumps leading to the significant reduction in boiler usage, as illustrated in Figure 3a, net fossil emissions decrease by 17% in Configuration A, as shown in Figure 4a. Improved alignment between system actors and cascading of heat through the chemical processes before heating the city further extend the reduction in fossil emissions to 22% compared to the base case.

The remaining fossil CO_2 emissions consist of approximately 60% originating from MSW incineration and 40% from the incineration of special waste generated by the chemical processes; however, the system remains far from achieving carbon neutrality. In Configuration D, the valorization of waste as a feedstock for basic chemical production reduces net emissions to 0.23 $kgCO_2/kgMSW$. However, due to the combustion of synthetic natural gas, fossil CO_2 originating from the MSW is released again into the environment. Therefore, the integration of post-combustion carbon capture is required to achieve net negative emissions. With the deployment of MEA-based carbon capture and assuming that all remaining residual CO_2 is captured and permanently stored, the system can reach net negative emissions of −0.5 $kgCO_2/kgMSW$.

Economic Performance and Cost Trade-offs

The breakdown of marginal operating and capital expenditures for all configurations is shown in Figure 4b. Across Configurations A–C, optimization of the existing

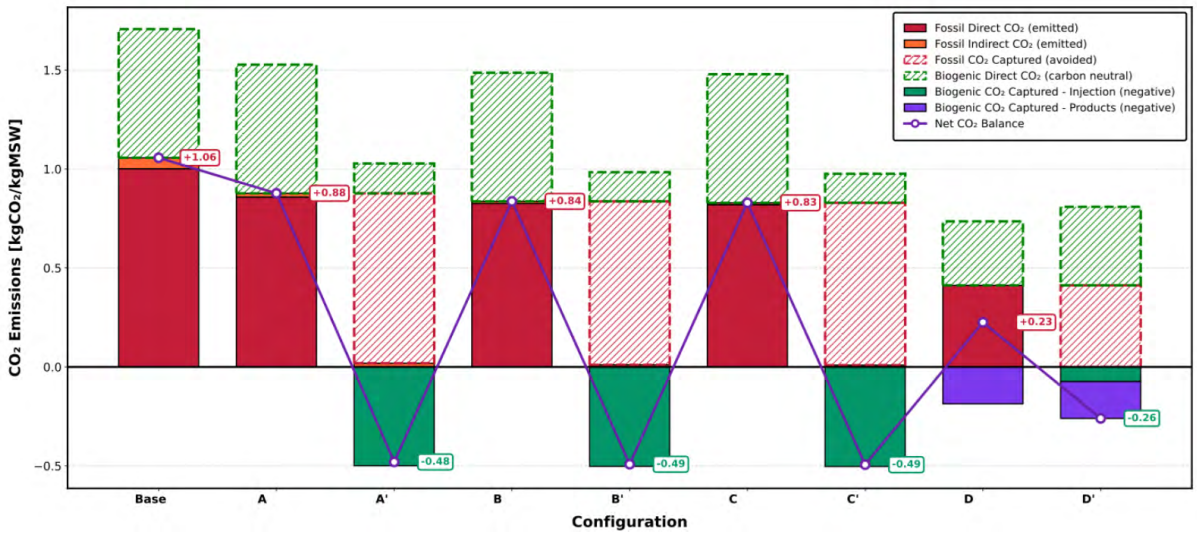

Figure 4.a. Breakdown of fossil and biogenic CO_2 emissions, captured CO_2, and net CO_2 balance across the different decarbonization configurations, normalized per kg MSW.

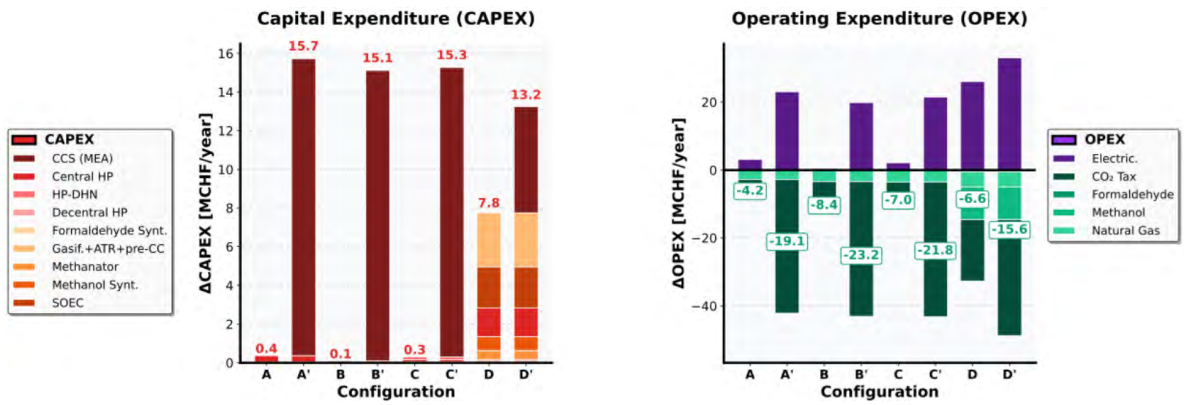

Figure 4.b. Breakdown of annualized capital expenditure (CAPEX) and operating expenditure (OPEX) across the different decarbonization configurations.

infrastructure combined with the installation of heat pump solutions requires limited capital investment and yields annual savings of up to 8.3 MCHF/year, with the maximum savings observed in Configuration B. In contrast, Configuration D results in a net loss of 2.2 MCHF/year. Although savings from avoided CO_2 taxes and additional revenues from methanol and synthetic natural gas exports exceed the required capital investments, the high cost of electricity imports ultimately leads to an overall net loss.

The integration of CCS technologies results in only a limited reduction in profitability for Configuration A'-C', amounting to approximately 0.2 MCHF/year, considering an average annual electricity spot price of 0.08 CHF/kWh and a conservative carbon price of 85 CHF/tCO_2 captured.

CONCLUSIONS

This work has demonstrated a systemic analysis of a real-world case study involving multiple actors within a symbiotic relationship, including a waste incineration facility, a utility company, a cluster of chemical industries, and the city. The analysis aimed to reduce fossil CO_2 emissions while concerning economic profitability through the combined use of process modeling, integration, and optimization methodologies.

The results showed that as a short-term solution, integration of heat pump technologies, optimizing the steam supply levels through means of heat cascading can lead to moderate reductions in fossil CO_2 emissions while providing significant economic benefits. However, these measures are followed by increase in electricity consumption, highlighting the need for adequate preparation of national electricity infrastructures in case of

mass adaptation.

While cost-effective and easily adoptable technologies are desirable, they fail on their own to achieve net-zero/negative operation. The findings of this work show the benefit of a layered decarbonization approach for a decarbonization pathway, in which energy efficiency and heat integration measures are prioritized before introducing capital intensive solutions. In the longer term, full defossilization requires a transition toward alternative renewable feedstocks and explicit consideration of carbon circularity to avoid burden shifting across sectors.

AUTHOR IDENTIFIERS

Author ORCIDs:
Boztaş Oktay : 0009-0003-8712-3432
Florez-Orrego, Daniel : 0000-0003-0103-7326
Domingos, Meire : 0000-0002-1699-8720
Maréchal, François : 0000-0003-1752-5690

REFERENCES

1. Swiss Federal Office for the Environment. Municipal waste incineration plants (MWI). 2025.
2. Sebastiani A, Macrì D, Gallucci K, Materazzi M. Steam - oxygen gasification of refuse derived fuel in fluidized beds: modelling and pilot plant testing. Fuel Processing Technology 216:106783 (2021). https://doi.org/10.1016/j.fuproc.2021.106783
3. Flórez-Orrego DA, Domingos MEGR, Nogueira Nakashima R, dos Santos MT, de Oliveira Junior S. Syngas purification by common solvents. Advances in Synthesis Gas : Methods, Technologies and Applications :27-72 (2023). https://doi.org/10.1016/b978-0-323-91877-0.00018-0
4. Ribeiro Domingos MEG, Flórez-Orrego D, Teles dos Santos M, de Oliveira S, Maréchal F. Techno-economic and environmental analysis of methanol and dimethyl ether production from syngas in a kraft pulp process. Computers & Chemical Engineering 163:107810 (2022). https://doi.org/10.1016/j.compchemeng.2022.107810
5. Franz AW, Kronemayer H, Pfeiffer D, Pilz RD, Reuss G, Disteldorf W, Gamer AO, Hilt A. Formaldehyde. Ullmann's Encyclopedia of Industrial Chemistry :1-34 (2016). https://doi.org/10.1002/14356007.a11_619.pub2
6. Flórez-Orrego D, Dardor D, Germanier R, Margni M, Maréchal F. A systemic study for decarbonizing secondary aluminium production via waste heat recovery, carbon management and renewable energy integration. Energy Conversion and Management 341:120021 (2025). https://doi.org/10.1016/j.enconman.2025.120021
7. Dardor D, Flórez-Orrego D, Terrier C, Ribeiro Domingos ME, Platteau C, da Silva JC, Lopez M, Maréchal F. ROSMOSE: a web-based decision support tool for the design and optimization of industrial and urban energy systems. Energy 304:132182 (2024). https://doi.org/10.1016/j.energy.2024.132182
8. Swiss Federal Office of Energy. Swiss Energy Dashboard. 2025.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Energy Baseline Surrogates for Modular Reactors from Generated Recipe-Based Process Data

Shreyas Parbat[a]*, Greeshmanth Rajanala[a], Isabell Viedt[a], and Leon Urbas[a]

[a] Technische Universität Dresden, Process-to-Order Group, Dresden, Germany
* Corresponding Author: shreyas.parbat@tu-dresden.de.

ABSTRACT

Energy Baselines (EnBs) provide a reference for evaluating Energy Key Performance Indicators (eKPIs), and their establishment is mandated under ISO 50001. Since eKPIs are typically defined per functional unit, such as product, recipe or recipe phase, EnBs should not be averaged across heterogeneous operating conditions but instead be defined in a context-specific manner. This requires detailed mechanistic models or sufficiently rich operational data for statistical approaches, both of which are often unavailable in highly flexible, semi-continuous production systems.This paper proposes a four-stage framework for the automated generation of surrogate EnBs to address this gap. In the first stage, the relevant training data space is defined, including non-influenceable variables (e.g., equipment deviations), design parameters (e.g., material properties), and adaptable recipe parameters (e.g., operating conditions and control actions). In the second stage, these parameters are systematically varied to automatically generate a labelled dataset representing process dynamics and energy behavior. The third stage trains and evaluates candidate surrogate models using the PRESTO recommendation framework, followed by a global sensitivity analysis based on Sobol indices to identify the dominant drivers of energy performance.The framework is demonstrated using a stirred tank reactor at laboratory (2 L) scale. The results show that Random Forest surrogates provide reliable performance across a wide operating range and that recipe parameters, particularly mixer rotational speed, dominate energy efficiency. Future work will address scalability to other equipment types and reactive processes.

Keywords: Energy Baseline, Surrogate Modeling, Recipes, Energy Efficiency, Energy Management, gProms

INTRODUCTION

Energy efficiency has become a central objective in the design and operation of chemical and biochemical production systems, driven by economic pressure, regulatory requirements, and sustainability targets [1]. Energy Management Systems (EnMS), standardized under ISO 50001, provide a structured framework for monitoring, evaluating, and improving energy performance [2]. Central to an EnMS is the definition of Energy Key Performance Indicators (eKPIs) together with appropriate Energy Baselines (EnBs), which serve as reference points for performance assessment, deviation detection, and verification of energy savings [2, 3].

In industrial practice, eKPIs are commonly expressed per functional unit, such as product [4], batch, recipe, or recipe phase [3]. This is particularly relevant for batch and semi-continuous processes, where production is governed by structured recipes as formalized in the ANSI/ISA-88 standard [5]. Consequently, meaningful EnBs must reflect the operational context in which the eKPIs are evaluated [3]. Averaging energy performance across heterogeneous recipes or operating modes can lead to misleading conclusions and conflicts ISO 50006 principles [6].

These challenges are particularly pronounced in flexible and modular production systems, where frequent reconfiguration, short production campaigns, and evolving system configurations limit the availability of sufficiently rich historical operating data. To address this, this work proposes a four-stage workflow for generating surrogate, context-specific EnBs based on systematically generated recipe-based process data and demonstrates its applicability using stirred-tank reactor case study.

FUNDAMENTALS

Energy Baselines and Energy Performance Indicators

EnBs and energy performance indicators (EnPIs) are fundamental elements of EnMS. ISO 50001 [2] defines EnBs as quantitative reference points against which changes in energy performance are assessed, while ISO 50006 provides guidance on their consistent establishment with respect to variables that significantly influence energy performance.

ISO 50006 distinguishes between influenceable parameters, defined as measurable factors whose values can change and that have a significant influence on energy performance [3, 6], and non-influenceable parameters, defined as measurable factors whose values are fixed or nearly fixed over the period under review but still significantly affect energy performance [4, 6].

Recipe-Based Operation

Batch and semi-continuous processes are typically operated using recipes, as formalized in the ANSI/ISA-88 [5]. Recipes describe the sequence of operations, setpoints, control actions, and phase transitions required to produce a given product [7]. In modular plants, recipes play an even more prominent role, as the same equipment modules may be reused across different products and operating modes.

Surrogate Modeling

Surrogate models are widely used to approximate complex process behavior when detailed mechanistic models are computationally expensive or difficult to maintain [8, 9]. In the context of energy systems, surrogate models have been applied for energy performance analysis, optimization, and decision support [4]. Common surrogate techniques include regression models, Gaussian processes, and ensemble methods such as Random Forest [10].

PROBLEM DEFINITION AND SCOPE

Conventional approaches to EnBs definition typically assume a stable operating context or rely on averaging energy performance across multiple operating conditions [6]. In recipe-driven batch systems, such averaging is inappropriate, as different recipes and recipe phases can exhibit fundamentally different energy behavior. Applying a single baseline across heterogeneous recipes may therefore mask inefficiencies and lead to misleading energy related performance assessments.

These limitations are further amplified in highly flexible and modular production systems, where frequent reconfiguration, and ongoing system evolution prevent the accumulation of sufficiently rich historical operating data. As a result, purely data-driven EnBs approaches become impractical, while the development and maintenance of detailed mechanistic energy models that capture non-influenceable effects are often costly. Despite the growing use of surrogate models for energy performance analysis, existing approaches rarely explicitly account for recipe-based operation and typically rely on historical plant data collected under fixed operating modes.

To address this gap, this work focuses on the generation of surrogate EnBs models using systematically generated simulation data from structured recipe execution. Deterministic process models are executed repeatedly under varying parameter realizations, with particular emphasis on recipe parameters.

FRAMEWORK FOR ENERGY BASELINES SURROGATE GENERATION

This section presents a four-stage framework for the generation of surrogate EnBs models using recipe-based simulation data. The framework is designed to address the lack of sufficiently rich operational data and the limited applicability of conventional EnBs approaches in highly flexible, semi-continuous production systems. Rather than replacing detailed mechanistic models, the proposed framework leverages such models as high-fidelity data generators to derive computationally efficient and interpretable surrogate EnBs representations.

I. Definition of the Training Data Space

The first stage focuses on the systematic identification and definition of the training data space for the EnBs surrogate model. As EnBs are inherently dependent on the operating context, it is essential to explicitly characterize the dimensions along which energy performance may vary. To this end, the parameter space [11] is structured into state variables describing the instantaneous dynamic condition of the system, model parameters representing physical and structural characteristics such as material properties and equipment geometry, recipe parameters capturing adaptable operational decisions including setpoints, phase durations, and control actions, input variables corresponding to external operational inputs, and output variables comprising energy-related quantities and derived eKPIs.

For each free parameter, appropriate bounds and probability distributions are defined, enabling a systematic exploration of feasible operating scenarios and ensuring that the generated dataset adequately represents the expected process dynamics and corresponding energy behavior. This structured definition of the training data space provides a consistent basis for generating representative simulation data to train reduced-order surrogate models for EnBs definition and analysis. While the framework includes all parameter categories as

potential dimensions of the training data space, individual studies may fix selected subsets depending on the analysis objective and system scope.

II. Automated Data Generation

In the second stage, the defined parameter space is systematically explored to generate labelled process data in an automated manner. All free parameters identified in Stage I are varied and combined according to their respective probability distributions, while the underlying process model is executed repeatedly using different parameter realizations. The simulations are carried out in gPROMS, with automated execution, parameter variation, and data extraction orchestrated via Python-based workflows. In particular, recipe parameters are explicitly varied, enabling the same recipe structure to be evaluated under a wide range of operating conditions.

Each simulation run yields a labelled data point comprising sampled input parameters and corresponding energy-related outputs. As the deterministic, noise-free process model is computationally expensive, the resulting dataset serves as a controlled reference for training surrogate models, enabling efficient EnBs evaluation, comparison across operating contexts, and sensitivity analysis without reliance on historical plant data.

A stratified sampling strategy is employed to ensure representative coverage of the feasible operating domain, resulting in a total of *N* simulation runs. The generated dataset represents idealized system behavior without measurement noise or structural model mismatch and is therefore intended for surrogate model development rather than direct representation of real plant variability.

While this provides a controlled benchmark for methodological development, it may lead to optimistic surrogate performance compared to real plant conditions. In practical applications, measurement noise, sensor uncertainty, and model–plant mismatch may affect predictive accuracy. Future work should therefore consider noise-augmented datasets and validation against measured plant data to ensure robustness.

III. Surrogate Model Training and Selection

In the third stage, the generated dataset is used to train candidate surrogate models for approximating the EnBs. The objective of the surrogate models is to capture the functional mapping between the selected input parameters and the resulting eKPIs. To enable a systematic and reproducible model selection, the PRESTO (Predictive REcommendation of Surrogate models To approximate and Optimize) [12] framework is employed. PRESTO assesses candidate models based on the characteristics of the dataset and the modeling task and recommends suitable surrogate structures accordingly.

The dataset is divided into training and test subsets using a fixed split to enable unbiased evaluation of predictive performance. All candidate surrogate models are trained using identical input parameter sets and evaluated on the same test data. Hyperparameters are not exhaustively optimized but are retained at default or PRESTO-recommended values to ensure a consistent and reproducible comparison. The objective of this stage is therefore comparative model screening rather than exhaustive surrogate tuning.

While this ensures a fair and reproducible comparison, it may disadvantage models sensitive to hyperparameter selection. Therefore, the reported performance reflects comparative behavior under standardized settings rather than maximum achievable accuracy.

IV. Global Sensitivity Analysis

The final stage focuses on improving the interpretability of the surrogate EnBs model and identifying the most influential parameters governing energy performance. To this end, a global sensitivity analysis [13] based on Sobol indices is conducted using the trained surrogate model. Both first-order (S_1) and total-order (S_T) Sobol indices are computed to quantify the individual contributions of each parameter as well as their combined interaction effects on the variability of the eKPIs. This analysis enables the identification of dominant parameters, the differentiation between influenceable and non-influenceable drivers of energy performance, and the assessment of the relative importance of recipe parameters in comparison to model parameters. Sobol indices are computed using surrogate model predictions rather than direct evaluations of the mechanistic model and may therefore be affected by approximation errors, particularly in regions with limited training data or lower predictive accuracy. While the high fidelity of the selected surrogate mitigates this effect, the results should be interpreted with consideration of potential surrogate-induced bias.

CASE STUDY: STIRRED TANK REACTOR

A stirred tank reactor (STR) having a total volume of 2 L is considered as the case study for process data generation. Figure 1 shows a schematic of the STR used in this study. The reactor is equipped with a thermal jacket for temperature control and inlet streams for reactants. Here, red and black colors highlight energy and mass streams, respectively, while the reactor dimensions are indicated in grey color.

The case study is intentionally selected as a structurally simple and well-understood system in order to isolate and demonstrate the proposed recipe-based EnBs surrogates generation framework without introducing additional complexity from reaction kinetics or phase behavior. In reactive systems, additional complexities are expected due to nonlinear coupling between mass and

energy balances, reaction heat effects, and concentration-dependent dynamics. These factors may increase model stiffness and introduce multiple operating regimes, requiring denser sampling and potentially more expressive surrogate structures. Nevertheless, the proposed framework remains applicable in principle and provides a foundation for such extensions. This allows the methodological aspects related to recipe execution, energy flows, and operating-context variability to be analyzed in a transparent manner.

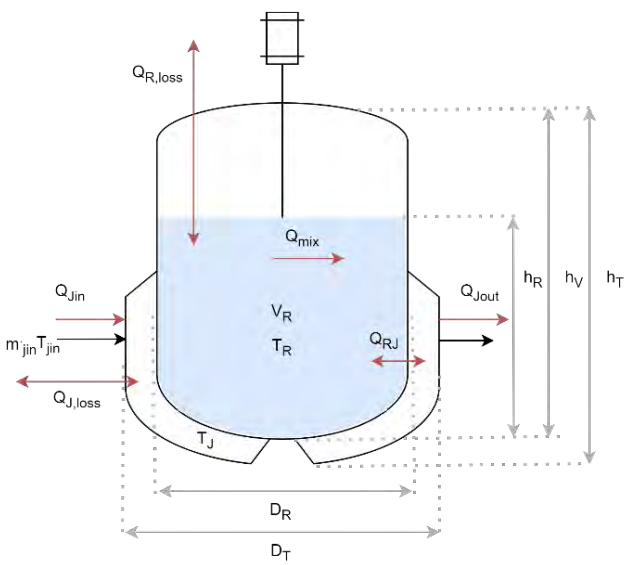

Figure 1. Schematic of the stirred tank reactor

The following assumptions are adopted to simplify the system modeling. Thermophysical fluid properties, such as density and heat capacity, are assumed to be constant and independent of temperature. The reactor contents are considered perfectly mixed, ensuring a spatially uniform temperature within each control volume. Reactor fouling and surface deposition effects are neglected, and the overall heat-transfer coefficient is therefore assumed to remain constant.

As the system contains only water, the reactor temperature dynamics are governed by the balance between heat transfer from the jacket, heat losses to the environment, and agitation-related heat input. The jacket temperature dynamics are described by the balance between the inlet enthalpy flow, heat exchange with the reactor, and heat losses to the environment:

$$\frac{dT_R}{dt} = \frac{1}{\rho_R V_R c_{pR}}(U_K A_{RJ}(T_J - T_R) + U_{RE} A_{RE}(T_{amb} - T_R) + Q_{mix}) \tag{1}$$

The reactor–jacket heat transfer area and reactor height are given by:

$$A_{RJ} = \pi D_R h_R + \frac{\pi D_R^2}{4} \tag{2}$$

$$h_R = \frac{4V_R}{\pi D_R^2} \tag{3}$$

$$A_{RE} = \frac{\pi D_R^2}{4} \tag{4}$$

The agitation-induced heat input is modeled via an empirical power correlation:

$$Q_{mix} = N_P \rho_R N^3 D_s^5 \tag{5}$$

$$N_P = \frac{3.26}{Re^{0.287}}\,; Re = \frac{\rho_R N D_s^2}{\mu_R} < 10^4 \tag{6}$$

$$N_P = \frac{0.484}{Re^{0.079}}\,; Re = \frac{\rho_R N D_s^2}{\mu_R} \geq 10^4 \tag{7}$$

The jacket temperature dynamics are described as and the jacket–environment heat transfer area is given by:

$$\frac{dT_J}{dt} = \frac{1}{\rho_J V_J c_{pJ}}(\dot{m}_{Jin} c_{pJ}(T_{Jin} - T_J) - U_K A_{RJ}(T_J - T_R) + U_{JE} A_{JE}(T_{amb} - T_J)) \tag{8}$$

$$A_{JE} = \pi D_T h_v + \frac{\pi D_T^2}{4} \tag{9}$$

A PI temperature control loop with sensor and actuator dynamics is implemented. The measured reactor temperature evolves according to:

$$\frac{dT_{Rmeas}}{dt} = \frac{(T_R - T_{Rmeas})}{\tau_s} \tag{10}$$

$$e = -(T_{sp} - T_{Rmeas}) \tag{11}$$

$$\frac{dI}{dt} = \frac{e}{t_i} \tag{12}$$

$$m_{Junsat} = K_c\,(e + I) \tag{13}$$

$$m_{Jclamped} = Max\,(m_{Jmin}, Min\,(m_{Junsat}, m_{Jmax})) \tag{14}$$

$$\frac{dm_J}{dt} = \frac{m_{Jclamped} - m_J}{\tau_v} \tag{15}$$

The reactor energy balance accounts for heat transfer between the reactor and the surrounding jacket, characterized by the overall heat-transfer coefficient U_K and the reactor–jacket heat-transfer area A_{RJ}, as well as heat losses to the environment and agitation-related heat input. The quantities $\rho_R, c_{p,R}$, and T_R denote the density, specific heat capacity, and temperature of the reactor contents, respectively, while $\rho_J, c_{p,J}, V_J$, and T_J describe the corresponding properties of the jacket. The thermal input to the jacket is represented by the inlet mass flow rate $\dot{m}_{J,in}$ and inlet temperature $T_{J,in}$, which together

determine the enthalpy flow driving the jacket temperature dynamics. Geometric parameters include the reactor volume V_R, reactor diameter D_R, impeller diameter D_s, and characteristic heights h_R and h_v, which define the relevant heat-transfer areas. Agitation is characterized by the impeller speed N, power number N_P, and Reynolds number Re, with μ_R denoting the dynamic viscosity of the reactor fluid. The PI controller is defined by the temperature setpoint T_{sp}, controller gain K_c, integral time constant t_i, and actuator time constant τ_v, with the jacket mass flow constrained between $m_{J,\min}$ and $m_{J,\max}$.

In this work, several eKPIs are considered, with a particular focus on thermal energy performance. The reactor heat exchange duty (Q_{RJ}) quantifies the useful heat transferred between the reactor and the jacket. Heat losses are captured by two additional indicators, the reactor heat loss ($Q_{R,loss}$), describing heat dissipated from the reactor to the environment, and the jacket heat loss ($Q_{J,loss}$), representing losses from the jacket to the surroundings. Based on these quantities, the thermal efficiency (η) is calculated as defined in Eq. (19), expressing the share of useful thermal energy relative to the total supplied heat, including unavoidable losses.

$$Q_{RJ} = U_K A_R (T_\mathrm{J} - T_\mathrm{R}) \tag{16}$$

$$Q_{R,loss} = U_\mathrm{RE} A_R (T_\mathrm{R} - T_{amb}) \tag{17}$$

$$Q_{J,loss} = U_\mathrm{JE} A_R (T_\mathrm{J} - T_{amb}) \tag{18}$$

$$\eta = \frac{Q_{RJ}}{Q_{RJ} + Q_{Jloss} + Q_{Rloss}} \cdot 100 \tag{19}$$

Here, the quantities A_{RJ}, A_{RE}, and A_{JE} correspond to the reactor–jacket, reactor–environment, and jacket–environment heat-transfer areas, respectively. U_K, U_{RE}, and U_{JE} are the associated heat-transfer coefficients.

While these eKPIs quantify instantaneous or batch-level energy performance, they are not directly used as EnBs. Instead, EnBs are defined in this work as expected reference values of the selected eKPIs under a specified operating context, including recipe parameters, system state, and relevant model parameters. To this end, the trained surrogate model provides a functional mapping from the operating context to the corresponding eKPIs, thereby serving as a context-dependent EnBs model. For this system, the parameter space is systematically explored. State, model, input, and recipe parameters are summarized in Table 1, while additional derived parameters include $N_P, Re, Q_{mix}, h_R, A_{RE}, A_{RJ}, A_{JE}, e, \dot{m}_{J,unsat}$, and $\dot{m}_{J,clamped}$. Parameters that directly affect the eKPIs and can be actively controlled ($T_{sp,preheat}, N_{mix,sp}$, etc.) are classified as influenceable. In contrast, design-specific parameters that influence the eKPIs only indirectly, through their impact on system properties or process dynamics, are considered non-influenceable. Based on the explored parameter space, simulation data are generated and used for training surrogates that provide operationally meaningful reference values for energy performance assessment.

The reactor operation is defined by a deterministic batch recipe composed of sequential phases:

1. Preheating: The jacket temperature is set to a predefined value to condition the reactor.
2. Mixing: The agitator operates at a specified rotational speed for a defined duration.
3. Heating: The reactor temperature is increased to a target setpoint via jacket control.

Table 1: Parameter Space Exploration.

Category	Parameters	Influenceability
State variables	T_R	Influenceable
	T_J	Influenceable
	m_J	Influenceable
	I	Influenceable
	T_{Rmeas}	Influenceable
Model parameters	ρ_R	Non-influenceable
	c_{pR}	Non-influenceable
	μ_R	Non-influenceable
	ρ_J	Non-influenceable
	c_{pJ}	Non-influenceable
	D_S	Non-influenceable
	T_{amb}	Non-influenceable
	T_{Jin}	Influenceable
	V_R	Non-influenceable
	V_J	Non-influenceable
	U_K	Non-influenceable
	U_{RE}	Non-influenceable
	U_{JE}	Non-influenceable
	D_R	Non-influenceable
	D_T	Non-influenceable
	h_v	Non-influenceable
	K_c	Influenceable
	t_i	Influenceable
	τ_v	Non-influenceable
	τ_s	Non-influenceable
	m_{Jmin}	Non-influenceable
	m_{Jmax}	Non-influenceable
Input variables	T_{sp}	Influenceable
	N	Influenceable
Recipe parameters	$T_{sp,preheat}$	Influenceable
	$t_{preheat}$	Influenceable
	$N_{mix,sp}$	Influenceable
	t_{mix}	Influenceable
	$T_{sp,heat}$	Influenceable
	t_{heat}	Influenceable
	t_{batch}	Influenceable

While the recipe structure and phase sequence remain unchanged across simulations, recipe parameters, including temperature setpoints, mixing speed, and phase durations, are treated as stochastic variables. The recipe is executed repeatedly while varying selected influenceable parameters, including temperature setpoints, mixing intensity, and phase durations. This design isolates the effect of recipe-level variability on eKPIs without introducing interference from structural influences.

Parameter bounds and probability distributions are defined a priori to ensure physically meaningful and operationally feasible scenarios. The process model is executed for each sampled parameter set, generating a diverse set of operating conditions.

A stratified sampling strategy ensures representative coverage of the operating domain. In this case study, 10 full batch simulations were performed. Each simulation produced approximately 662 time-step observations across four phases, resulting in 6, 621 input–output pairs. The dataset was split into training (80%) and test (20%) sets using a fixed random seed. Despite the limited number of batch simulations, the time-resolved data provides a sufficiently rich sample set for surrogate model training. The dataset is not used directly as an EnBs but to train a surrogate model that predicts expected energy performance under a given operating context. These predictions define context-specific EnBs, against which observed or simulated eKPIs are compared to assess performance deviations.

RESULTS

Multiple surrogate models, including artificial neural networks (ANN), extreme learning machines (ELM), Gaussian process regression (GPR), multivariate adaptive regression splines (MARS), radial basis function networks (RBFN), Random Forests (RF), and support vector regression (SVR), are evaluated using the PRESTO framework. The PRESTO variant proposed by Williams et al. [12] is excluded due to limitations associated with the high data volume required for reliable training. All remaining models are trained using identical input sets comprising all state, model, input, and recipe parameters, while excluding the output variables $T_R, T_J, Q_{RJ}, Q_{R,\mathrm{loss}}, Q_{J,\mathrm{loss}}$, and η.

Input features are standardized prior to training for models sensitive to feature magnitude (ELM, SVR, ANN, GPR, RBFN), using parameters derived from the training dataset and consistently applied to the test data. For RF and MARS, unscaled inputs are used, as these methods are invariant to monotonic transformations. Figure 2 compares the predictive performance of the trained models for the representative output Q_{Jloss}.

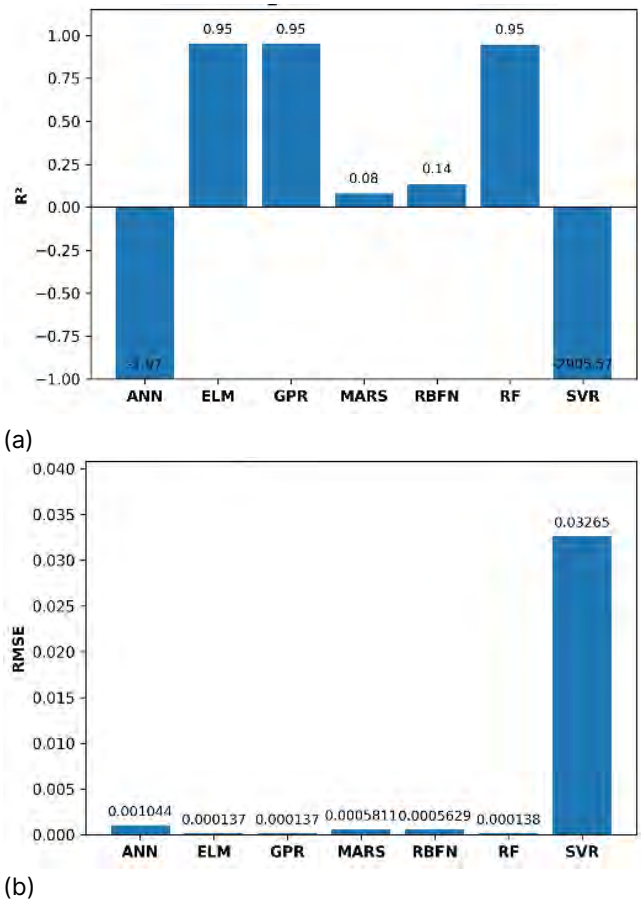

Figure 2 Performance comparison of candidate surrogate models (a) R^2, (b) RMSE

Among the evaluated candidates, the RF model consistently exhibits superior predictive accuracy and robustness and is therefore selected for all subsequent analyses. SVR exhibits poor predictive performance for outputs with small magnitudes (~10^{-1}), as the ε-insensitive loss (ε = 0.1) is comparable to the target scale, causing many deviations to fall within the ε-tube and thus not contribute to model learning. ANN similarly exhibit unstable convergence behavior despite input standardization, likely due to the high dimensionality of the input space and the absence of dedicated hyperparameter tuning, leading to ill-conditioned optimization and unstable gradient updates. No additional data normalization or architectural tuning was applied to ensure a consistent and unbiased comparison across surrogates.

Figure 3 presents parity plots comparing RF predictions with reference values obtained from process simulations for all considered outputs and eKPIs. Across all outputs, the surrogate model achieves coefficients of determination (R^2) close to unity and low root-mean-square errors (RMSE), indicating accurate representation of nonlinear relationships over the explored operating space. Minor deviations occur primarily near the boundaries of the parameter ranges or for outputs with limited

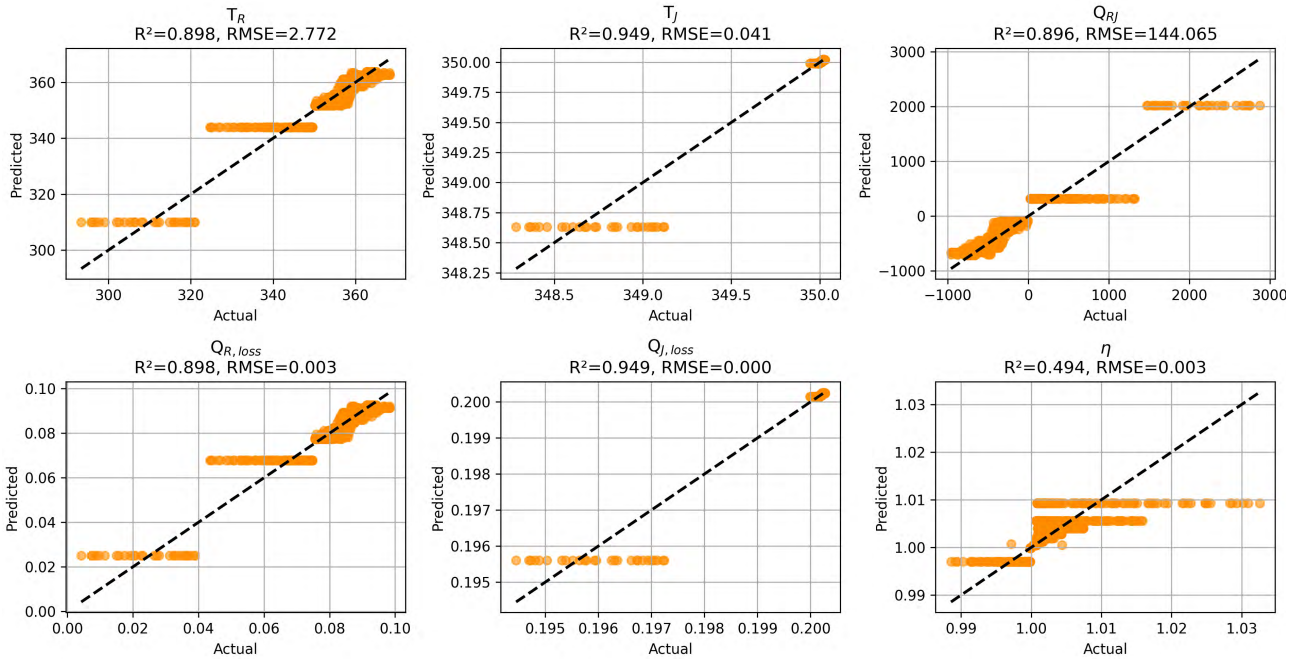

Figure 3. Parity plots of predicted versus simulated values for all output variables and eKPIs using the Random Forest surrogate model

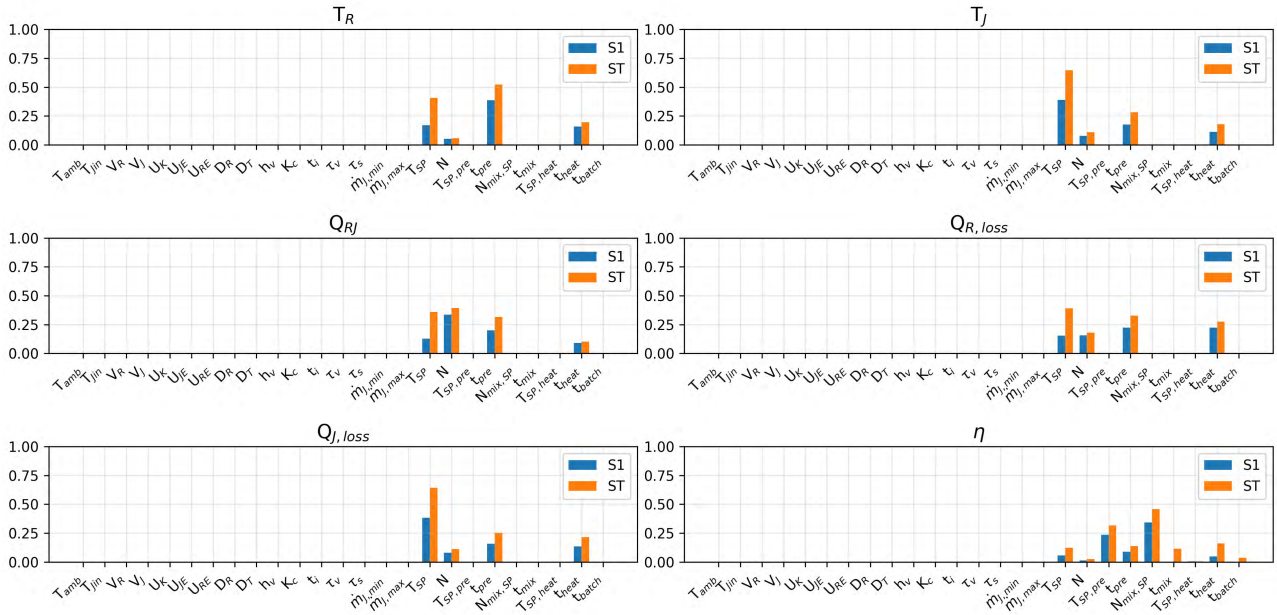

Figure 4. Global Sensitivity Analysis Using Sobol Indices

variability, where small absolute errors appear more pronounced. These results confirm that the trained RF surrogate is suitable for representing recipe-dependent EnBs and is subsequently used for global sensitivity analysis and performance assessment.

Figure 4 summarizes the first-order (S_1) and total-order (S_T) Sobol indices for the outputs $T_R, T_J, Q_{RJ}, Q_{R,loss}, Q_{J,loss}$, and η. While the first-order index quantifies the independent contribution of each parameter, the total-order index additionally captures interaction effects. Across all outputs, only a small number of parameters significantly influence the results, while most parameters have a negligible impact on energy-related behavior within the studied ranges.

For both reactor temperature (T_R) and jacket temperature (T_J), dominant sensitivities are associated with temperature setpoints, phase definitions, and reactor–jacket heat-transfer capability. This is physically con-

sistent, as temperature evolution in a non-reactive system is governed by thermal driving forces and heat transfer rates. The gap between total- and first-order indices indicates strong interaction effects, particularly between setpoints, heat-transfer capacity, and phase durations, shaping transient heating behavior.

The reactor–jacket heat transfer Q_{RJ} is primarily influenced by parameters determining the driving force and transfer rate, notably heat-transfer capability and recipe-dependent temperature differences and durations. The observation that S_T exceeds S_1 for key parameters confirms the presence of coupled effects, where both temperature response dynamics and phase durations govern heat transfer.

Loss-related outputs are most sensitive to temperature gradients, heat-loss coefficients, and phase durations, reflecting their dependence on both temperature differences and residence time. Interaction effects, particularly between setpoints and batch duration, further influence cumulative losses.

The mixing-related energy Q_{mix} exhibits sensitivity almost exclusively to recipe parameters associated with mixing intensity and mixing duration, while thermodynamic and structural parameters remain negligible. This behavior is consistent with physical expectations for a recipe-controlled agitation input and confirms that the surrogate-based sensitivity analysis preserves the intended separation between mechanical energy input and thermal exchange mechanisms in the adopted lumped, non-reactive formulation.

In contrast, the efficiency η displays the broadest sensitivity profile, with multiple contributing parameters and total-order indices that often exceed the corresponding first-order indices. As a KPI aggregating heat transfer and losses, it is strongly influenced by interactions between recipe parameters and system characteristics. Consequently, changes in operational decisions such as heating targets or phase durations can influence efficiency to a similar extent as variations in equipment-related parameters, underscoring the need for recipe-aware EnBs in batch and modular systems.

CONCLUSION AND OUTLOOK

This work demonstrates the feasibility of a four-stage framework for generating surrogate EnBs models from systematically generated recipe-based simulated data. By combining structured parameter space definition, automated data generation, surrogate modeling, and global sensitivity analysis, the framework enables EnBs definition without reliance on historical data.

Among the evaluated surrogate modeling approaches, Random Forests are found to be particularly well suited for the investigated case study, providing high predictive accuracy and compatibility with global sensitivity analysis. However, the optimal surrogate model choice is expected to depend on the characteristics of the underlying process, the dimensionality of the parameter space, and the structure of the available data.

The results show that recipe parameters, most notably mixing rotational speed, dominate the variability of energy-related outputs and eKPIs, while equipment-related parameters primarily influence system behavior through interaction effects. This provides quantitative justification for defining EnBs that are explicitly conditioned on recipe context, as aggregation across heterogeneous operating regimes can obscure meaningful performance differences. At the same time, the proposed approach should be understood as a model-based EnBs support tool whose validity depends on the adequacy of the underlying mechanistic model and requires validation against metered process data before deployment in operational EnMS.

Future work will focus on extending the framework to systems with increased modeling complexity, including reactive processes and coupled mass–energy interactions, as well as validating surrogate-based EnBs against measured process data to enable ISO-compliant energy performance assessment.

ACKNOWLEDGEMENTS

As part of the “Energy Efficiency in Industry”, research funding is provided by the Federal Ministry for Economic Affairs and Energy of Germany (BMWI) and Project Management Jülich (PtJ) under the title of project REUNION (FKZ 03EN2100C).

AUTHOR IDENTIFIERS

Author ORCIDs:
Shreyas Parbat: 0009-0000-3014-9859

REFERENCES

1. Ali M, Evans S. The impact of best practice on energy efficiency in industrial decarbonization policy. Energy Policy 204:114632 (2025). https://doi.org/10.1016/j.enpol.2025.114632
2. ISO 50001. ISO 50001: 2018(e) energy management systems — requirements with guidance for use. (2018)
3. Parbat S, Viedt I, Härtner S, Urbas L. Energy management system plug-&-monitor architecture with an asset administration shell. IFAC-PapersOnLine 59:59-64 (2025). https://doi.org/10.1016/j.ifacol.2025.11.925
4. Beisheim B, Rahimi-Adli K, Krämer S, Engell S. Energy performance analysis of continuous processes using surrogate models. Energy

183:776-787 (2019).
https://doi.org/10.1016/j.energy.2019.05.176
5. American National Standards Institute. ANSI-isa-88.00. 01-2010: batch control part 1: models and terminology. (2010)
6. ISO 50006. ISO 50006: 2023(e) energy management systems – measuring energy performance using energy baselines (enb) and energy performance indicators (enpi). (2023)
7. Winter M, Klausmann T, Kleinert T. Capability knowledge base query to allocate process resources for master recipe formulation. at - Automatisierungstechnik 72:906-918 (2024). https://doi.org/10.1515/auto-2023-0223
8. Kudela J, Matousek R. Recent advances and applications of surrogate models for finite element method computations: a review. Soft Comput 26:13709-13733 (2022). https://doi.org/10.1007/s00500-022-07362-8
9. Parbat S, Viedt I, Urbas L. A comparative evaluation of complexity in mechanistic and surrogate modeling approaches for digital twins. Systems and Control Transactions 4:166-172 (2025). https://doi.org/10.69997/sct.122855
10. Winz J, Nentwich C, Engell S. Surrogate modeling of thermodynamic equilibria: applications, sampling and optimization. Chemie Ingenieur Technik 93:1898-1906 (2021). https://doi.org/10.1002/cite.202100092
11. Potter C. PQLI application of science- and risk-based approaches (ICH Q8, Q9, and Q10) to existing products. J Pharm Innov 4:4-23 (2009). https://doi.org/10.1007/s12247-009-9051-9
12. Williams B, Otashu J, Leyland S, Eden MR, Cremaschi S. PRESTO: predictive recommendation of surrogate models to approximate and optimize. Chemical Engineering Science 249:117360 (2022). https://doi.org/10.1016/j.ces.2021.117360
13. looss B, Lemaître P. A review on global sensitivity analysis methods. Operations Research/Computer Science Interfaces Series :101-122 (2015). https://doi.org/10.1007/978-1-4899-7547-8_5

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Evaluation of dual pressure low-temperature distillation for LNG Production in CO_2-rich fields

Victor S. V. Mercado[a*], Dirceu Noriler[a], Laura Plazas Tovar[b], Radin Suhaib Salihuddin[c], Amiza Bt Surmi[c], Fadhli Hadana Rahman[c], and Jean F. Leal Silva[a]

[a] Universidade Estadual de Campinas (UNICAMP), Faculdade de Engenharia Química (FEQ), Campinas, SP, Brazil
[b] Universidade Federal de São Paulo (UNIFESP), Instituto de Ciências Ambientais, Químicas e Farmacêuticas (ICAQF), Diadema, SP, Brazil
[c] PETRONAS Research Sdn Bhd, Bangi, Malaysia
* Corresponding Author: v194994@dac.unicamp.br

ABSTRACT

Liquefied natural gas (LNG) plays a key role in the energy transition, but its production is often limited by the high CO_2 content of some reservoirs, which increases operating costs and solidification risks. This study evaluates the dual-pressure low-temperature cryogenic distillation process applied to a recently discovered gas field with a high CO_2 content (25%) for LNG production. The critical properties of the gas streams and the process operating conditions were analyzed using Aspen Plus v14. The results indicate that reducing the CO_2 concentration throughout the column is essential to prevent solid formation by maintaining a fluid composition with a freezing temperature below the operating temperature of the stages. It was also observed that the reflux affects LNG purity and freezing temperature in all stages. Furthermore, the adoption of a low-pressure separator upstream of the distillation proved crucial to producing condensates within commercial specifications. The work extends the understanding of the dual-pressure low-temperature cryogenic distillation to the 25% CO_2 range and highlights the potential of cryogenic distillation to enable LNG production in fields with high CO_2 content.

Keywords: LNG, Cryogenic distillation, Simulation, CO2, Gas field

INTRODUCTION

The energy transition from fossil fuels to renewable energy depends on the role of natural gas (NG) in mitigating the limitations of renewable energy. For example, the strong dependence on climate conditions, the weak structure to satisfy the high energy demand of cities during specific periods of the day, and the need to develop new technologies to store the energy supplied by the environment are among the key issues in this context [1]. NG alleviates these problems since it presents a competitive cost in comparison to oil [2] and has a lower greenhouse gas (GHG) emissions potential than coal [3]. In the past ten years, production of dry NG around the world has been increasing, supported by an increasing NG consumption in electricity generation [4]. As a result, the share of NG has practically remained constant (22% - 23%) despite the rise of renewable energy [5]. Together, these factors highlight the importance of NG in the energy mix.

One option for transporting NG is liquefaction, which reduces its volume by a factor of 600. This situation gives liquified natural gas (LNG) some notable advantages in NG supply. For instance, small quantities of LNG offer high NG availability and can be economically stored in tanks when a higher-than-needed supply is available. Additionally, LNG can be transported for long distances, which sometimes can be unfeasible with pipelines [6-8].

One of the main challenges in LNG processing is the high CO_2 content (15%-80% by mole) in some gas fields, which increases operating costs and poses a risk of solidification during processing, encompassing many technical challenges [6-8]. In the last ten years, companies have developed gas production projects in fields with high CO_2 content. For instance, gas production has been developed in the LaBarge field in Wyoming (USA) and in fields located in Southeast Asia (Malaysia, Indonesia, Vietnam, and the southern Gulf of Thailand) [7-12].

Similarly, some stakeholders have been evaluating the production potential in regions such as the Otway Basin (Southeast coast of Australia), the Kwanza Basin (Angola), and the Santos Basin (Brazil) [13-15]. This situation highlights the importance of innovation in producing LNG from those fields.

This work aims to evaluate the dual-pressure low-temperature distillation process for producing LNG from a recently discovered field with high CO_2 content. The process was simulated in Aspen Plus v14, and the results highlight the limitations and operational parameters to avoid solids formation and critical conditions.

MATERIALS AND METHODS

Data collection

The composition and operating conditions of a recently discovered gas field are shown in Table 1.

Table 1: Composition and operating conditions of the feed gas

Feed composition	Mole %
C_1	64.06
C_{2+}	10.33
CO_2	25
H_2S	0.01
N_2	0.60
Flowrate (MMscfd)	1
Pressure (bar)	70
Temperature (°C)	30

Process simulation

The process was simulated in Aspen Plus v14, to achieve an LNG production of 0.0096 MMTPA (million metric tons per annum), with a lower heating value (LHV) of 900.47 BTU/scf (British Thermal Unit per standard cubic foot of gas). To develop the simulation, the following aspects were considered: selection of the thermodynamic model, estimation of the pseudocritical temperatures of the streams, determination of the operating conditions of units, and sensitivity analysis.

Thermodynamic model

As the components can be considered all nonpolar, including H_2S because of its weak electronegativity, the most appropriate thermodynamic package to estimate the physical properties of the streams throughout the overall process corresponds to an Equation of State (EoS) [16-17]. The three main EoS considered in this study were: Peng-Robinson (PR), Soave-Redlich-Kwong (SRK), and Perturbed Chain Statistical Associating Fluid Theory (PC-SAFT). Results of equilibrium data for every EoS were compared to experimental data collected at the solid-vapor, liquid-vapor, and solid-liquid equilibria region for different alkane-CO_2 binary systems [18-25]. Finally, the PR EoS proved to be the thermodynamic model with the best fit to experimental data. This situation is explored in Figure 1, where the results of the selected EoS are compared with data reported by Wei *et al.* [25].

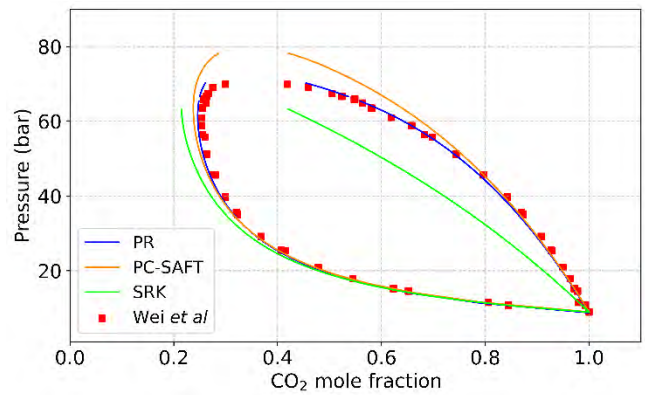

Figure 1. Validation of the Peng-Robinson Equation of State on the vapor-liquid equilibrium region of the C_1-CO_2 binary system at -43 °C

Pseudocritical temperatures

The critical point of a gas mixture is the temperature and pressure at which the distinction between the gas and liquid phases disappears. Under this condition, the dew point and bubble point curves converge to a single point in the pressure vs. temperature phase diagram [26].

Critical properties are not usually measured for hydrocarbon mixtures; instead, they are estimated as pseudocritical properties from correlations and depend strongly on gas composition [6, 26]. Once critical conditions are reached, measuring those properties in vapor and liquid becomes too difficult [27-28]. To ensure the accuracy of the results, the authors have defined that pseudocritical conditions constitute an operational limit in the simulation. This is illustrated in Figure 2, where the pseudocritical temperature of the process gas feed was estimated by varying the CO_2 content (25-65%) and the C_{2+} content (2-10%).

Operating conditions of units

The process is divided into two stages: condensate production (Figure 3) and dual-pressure, low-temperature distillation (Figure 4). First, the reservoir gas feeds the high-pressure separator HG-SEP, which operates at 70 bar and 30 °C. Then, the bottom product flows to the low-pressure separator LP-SEP, whose conditions can be optimized to achieve a Reid vapor pressure (RVP) between 0.69 and 0.82 bar at stream LKO-02 (based on a commercial specification for the liquid knock-out stream [29]). Both operations were simulated using a *Flash2* model.

CO_2 (mol %) \ C_{2+} (mol %)	2	4	6	8	10
25	-51.8	-49.4	-47.2	-45.1	-43.2
30	-46.1	-43.7	-41.5	-39.4	-37.4
35	-40.4	-38.0	-35.8	-33.7	-31.6
40	-34.7	-32.3	-30.0	-27.9	-25.9
45	-29.0	-26.5	-24.3	-22.1	-20.1
50	-23.3	-20.8	-18.5	-16.4	-14.3
55	-17.5	-15.1	-12.8	-10.6	-8.5
60	-11.8	-9.3	-7.0	-4.8	-2.7
65	-6.1	-3.6	-1.3	1.0	3.2

Pseudocritical temperature (°C): 3, -11, -24, -38, -52

Figure 2. Estimation of the pseudocritical temperature for different compositions of gas feed.

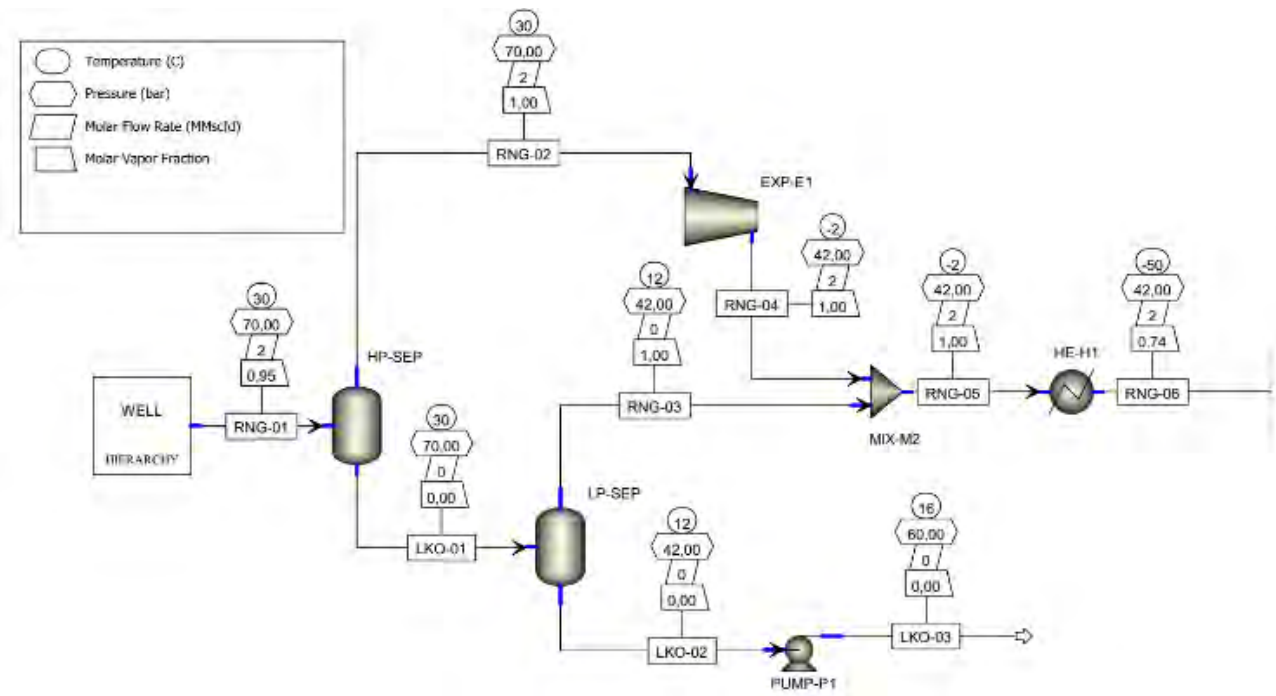

Figure 3. Setup of the condensate separation step.

At the same time, the top product of HP-SEP is directed towards the expander EXP-E1 (*Compr* model set as turbine), reducing its pressure to 42 bar. After the expansion, the stream is mixed with the top product of the LP-SEP and then cooled in the Heater model HE-H1, preparing it for the next step.

In the dual-pressure, low-temperature distillation stage (Figure 4), the gas stream RNG-06 feeds column COL-D1 (*RadFrac* model). This first column operates at 42 bar and a temperature range of -82 °C to -16 °C, producing an overhead product with approximately 80 mol% of methane. This distillate is then divided into two streams: one feeds the bottom of column COL-D2 (*RadFrac* model) directly, and the other (RNG-09) passes through the heater HE-H2 before heading to COL-D2. The second column operates at 39 bar and a temperature range of -89 °C to -84 °C. The bottom product is recycled to COL-D1, passing through *Pump* model PUMP-P2 to increase the stream pressure to 42 bar. As the CO_2 solid formation depends on the increase in CO_2 concentration [6], an additional outlet stream was considered in the column stage with a higher CO_2 concentration. As a result, a rich methane stream is produced at the top of COL-D2 (greater than 99 mol%). Likewise, a rich methane stream is produced at the bottom of COL-D1 (approximately 90 mol%).

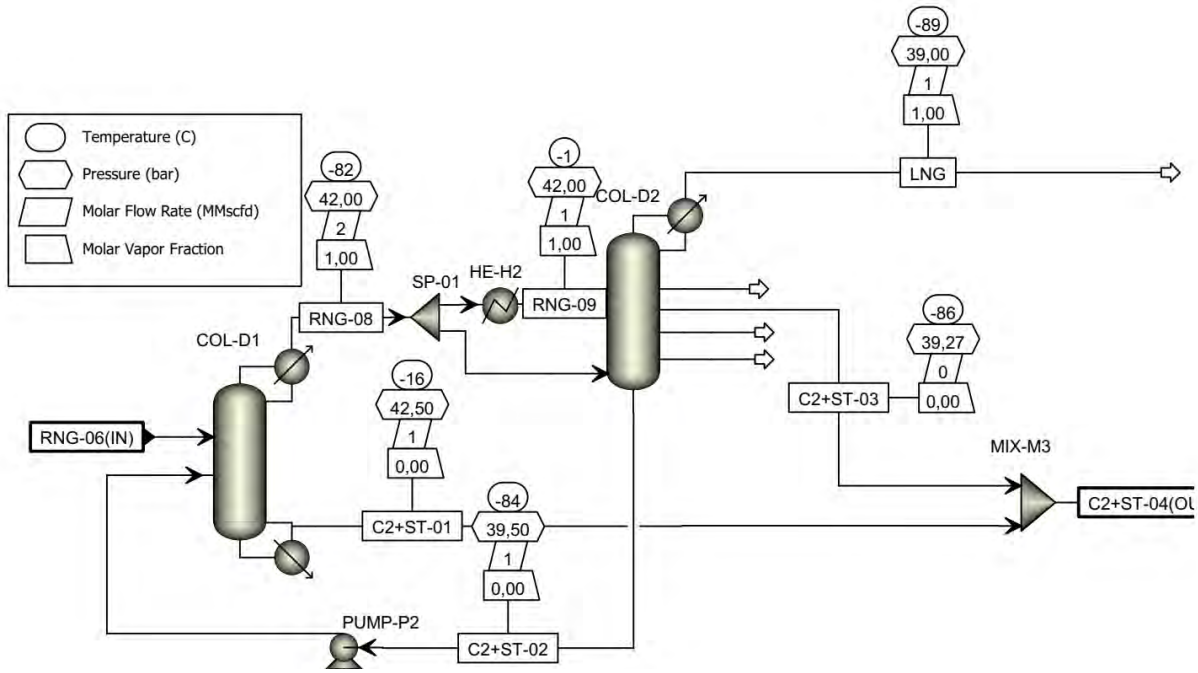

Figure 4. Setup of the dual-pressure, low-temperature distillation step.

Sensitivity analysis

The present study conducted two sensitivity analysis: one evaluating the operating conditions of the LP-SEP and another assessing those of COL-D2.

The project aims to produce a liquid condensate stream for sale to oil refineries, so the main objective of the first stage is to produce a liquid stream from the LP-SEP that meets commercial requirements. In this work, the Reid vapor pressure in LKO-03 (Figure 3) is used to measure its volatility. The Reid vapor pressure is defined as the vapor pressure of a liquid sample measured in a chamber at 37.8 °C, according to the D323 method of the American Society for Testing and Materials (ASTM) [30]. Usually, the Reid vapor pressure of a mixture is lower than its true vapor pressure, so crude stabilization systems can be designed to achieve a specified value to avoid the vaporization of volatile compounds [29].

On the other hand, the main variable used to assess LNG produced in the second stage of the process is CO_2 content, which is expected to be below 50 ppmv. This low concentration is only possible by operating COL-D2 at cryogenic temperatures ranging from -160 °C to -56 °C [7-8]. This means a challenge since the high concentrations of CO_2 in some stages of the column can lead to solids formation and to a proximity to the pseudocritical temperature [6, 17, 27, 28].

To evaluate operating conditions in COL-D2, this work examined the molar reflux ratio at the top and its potential impact on the freezing temperature in stages with higher CO_2 concentrations. Thus, the variable ΔT_F was defined as the difference between the stage temperature and the freezing temperature of the mixture, mathematically expressed as:

$$\Delta T_{Fi} = T_i - FT_i \tag{1}$$

where T_i is the temperature of the column at stage i, and FT_i corresponds to the mixture's freezing temperature at that stage. Additionally, the minimum reflux ratio is specified to achieve a CO_2 content below 50 ppmv, as LNG production is required to meet all specifications. In this context, ΔT_{Fi} works as an indicator of the safety margin between the freezing and operating temperature. This way, the research focuses on the operating parameters to avoid potential risks during the process.

RESULTS AND DISCUSSION

Condensate separation step

The operating conditions in LP-SEP were defined to achieve a commercial Reid vapor pressure between 0.76 and 0.80 bar. However, the separator pressure was also determined by the outlet pressure at EXP-E1, since a pressure difference between RNG-03 and RNG-04 could imply a problem before entering HE-H1. Similarly, the outlet pressure at EXP-E1 must support the LNG production in the next stage. Consequently, the operating pressure of LP-SEP was fixed, and only the effect of temperature on the Reid vapor pressure of LKO-03 was evaluated. The sensitivity of Reid vapor pressure to the temperature in LP-SEP is illustrated in Figure 5.

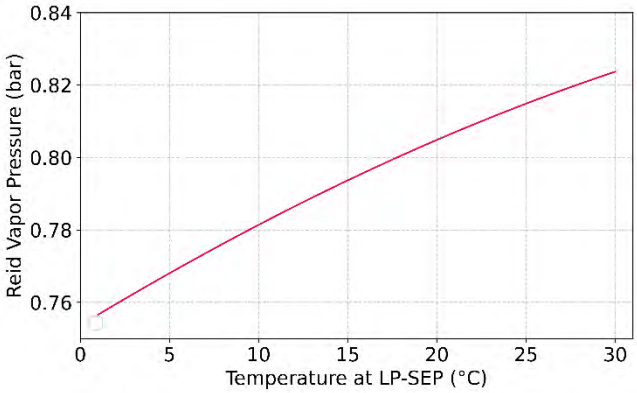

Figure 5. Reid vapor pressure of the condensate stream as a function of temperature at LP-SEP.

The increase in Reid vapor pressure in LKO-03, along with the corresponding temperature rise at LP-SEP, is due to two factors. First, the solubility of methane in light hydrocarbons (between C_4 and C_{10}), which increases strongly with the pressure rise and with the high content of such hydrocarbons [31]. Second, the solubility of CO_2 in hydrocarbon mixtures, which increases with temperature [17]. Both factors contribute to the presence of methane and CO_2 (volatile compounds) in LKO-03, reducing the boiling point of the stream and increasing its vapor pressure [32]. For example, the molar fraction of CO_2 plus the molar fraction of methane in LKO-03 ranged between 30 and 40% across the different temperatures.

Downstream, the top product of LP-SEP and the outlet stream RNG-04 from EXP-E1 are mixed in MIX-M2 (*Mixer* model), whose outlet constitutes stream RNG-05. This stream was then cooled in HE-H1, setting the outlet temperature between the freezing and pseudocritical points, -42 °C and -68 °C, respectively. This procedure avoids the solids formation without reaching the pseudocritical limit [6].

Dual-pressure, low-temperature distillation stage

The RNG-06 stream enters COL-D1, which at the same time is fed by a reflux stream. The composition profiles of methane and CO_2 are illustrated in Figure 6.

The top product in COL-D1 reached a molar fraction of approximately 95% of methane. This result is similar to that calculated by Anaugraha *et al.* [8], who estimated a methane molar fraction of 88% in the Natuna gas field, and 99% in the Bojonegoro gas field. In the same way, Pellegrini *et al.* [7] estimated a methane molar fraction between 85% and 92%. The reflux was determined to feed the column at the tenth stage, reducing the temperature and the CO_2 concentration. This configuration avoids the solids formation because the reflux stream has a methane molar fraction of approximately 90%. The temperature profile, the pseudocritical temperature, and the freezing temperature throughout each stage in COL-D1 are presented in Figure 7.

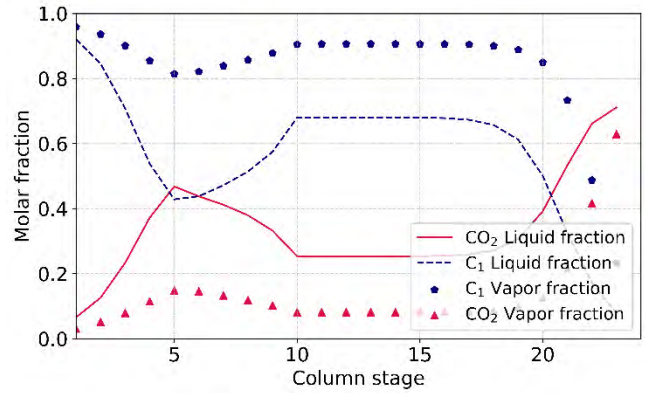

Figure 6. Compositions profile in COL-D1

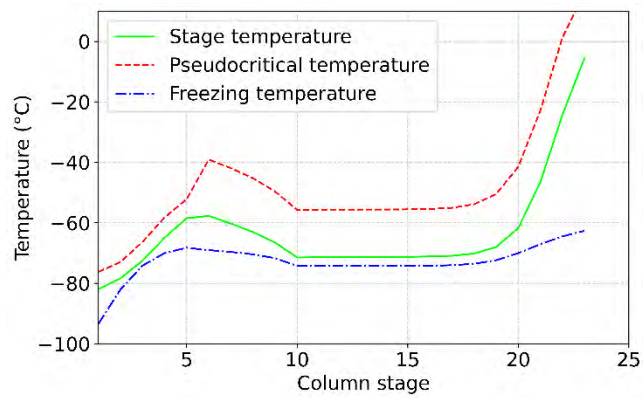

Figure 7. Temperature profile in COL-D1

The rich methane stream produced at the top of COL-D1 feeds COL-D2, whose methane and CO_2 composition profiles are shown in Figure 8.

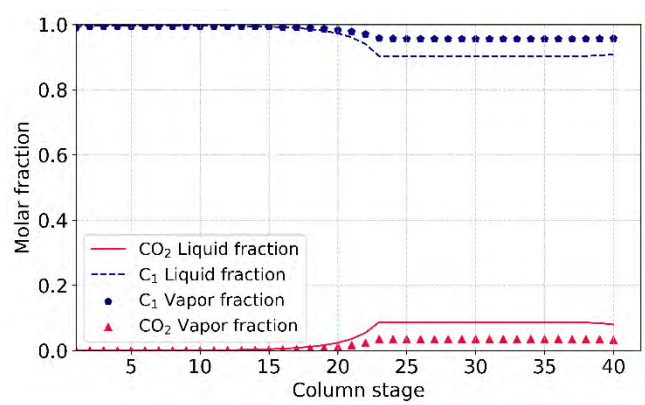

Figure 8. Compositions profile in COL-D2

As suggested by Anaugraha *et al.* [8] and Pellegrini *et al.* [7], the dual-pressure, low-temperature distillation presents the potential to produce natural gas with a minimum fraction of CO_2 content. However, the low temperatures and high operating pressure pose a risk of solid

formation. Figure 9 shows the freezing, pseudocritical and operating temperatures throughout the entire column COL-D2. This approach aligns with practices presented in GPSA [6] by considering solid formation and critical conditions as limits of the operational window.

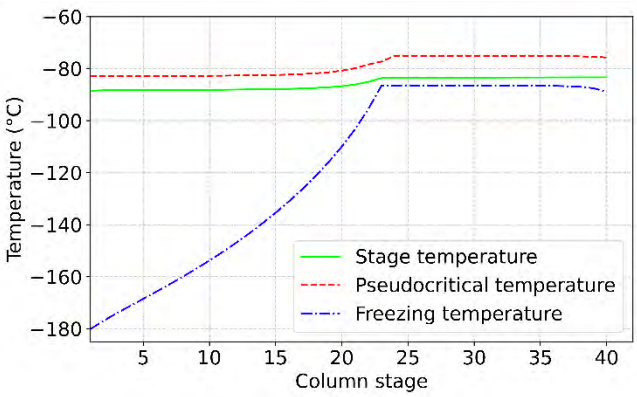

Figure 9. Temperature profile in COL-D2

It is important to notice in what stages the higher risk of solids formation can occur. In the case of this study the stages where the freezing temperature is close to the operating temperature are the 22nd, 23rd, and 24th. These stages have CO_2 and methane mole fractions of about 8.7% and 90.3%, respectively. In this context, the mixture at such stages behaves as a C_1-CO_2 binary system. These results align with the experimental data measured by ZareNezhad [21] and GPSA [33], reporting that the high increase in freezing temperature occurs with a CO_2 molar fraction between 6 and 10%.

Several authors report possible solutions for this cumbersome situation. Lastari [34] proposed an approach based on the Ryan/Holmes process by adding light hydrocarbons (C_{4+}), affecting the CO_2 solubility and concentrations. The GPSA [6], highlighted hydrocarbons addition led to a temperature rise in the overhead product, an increase in the relative volatility between methane and CO_2, and a rise in critical pressure, allowing higher-pressure operation. The Controlled-Freeze Zone (CFZ) developed by ExxonMobil promotes the CO_2 freezing in an open area of the column so that it can be dissolved and removed in the stripping section [35]. Pellegrini *et al.* [7] proposed the dual-pressure low-temperature distillation process to separate the methane from the CO_2. The high pressures in both columns and the recycling of methane-rich streams can prevent CO_2 solid formation despite the low temperatures.

In dual-pressure, low-temperature distillation, the low-pressure column COL-D2 serves as the final stage of methane purification. In this context, this work proposed an auxiliary outlet stream at stage 23 that can be mixed with the bottom product of COL-D1 to reduce CO_2 concentration. This scheme guarantees a methane recovery at distillate higher than 95%, which is similar to the results reported by Anaugraha *et al* [8] and Lastari [34], as shown in Table 2.

Table 2: Hydrocarbon recovery during LNG production

Hydrocarbon recovery (% mol/mol)	This work	Anaugraha	Lastari
C_1 (top)	95.8%	87.8% - 99.9%	99.9%
C_{2+} (top)	<0.1%	1.6%	0.22%
CO_2 (bottom)	99.9%	98.5% - 99.9%	99.9%

Furthermore, the research analyses the operating conditions in COL-D2. Thus, the impact of the molar reflux ratio that flows from the condenser on the ΔT_F variable on stages 22 and 24 is illustrated in Figure 10.

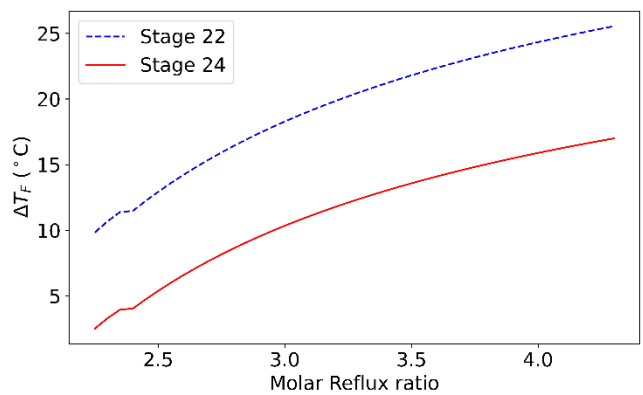

Figure 10. Molar reflux ratio effect on ΔT_F parameter on COL-D2

The safety margin between the operating temperature and freezing temperature increases as the molar reflux ratio in COL-D2 increases. This can be explained by the rise in methane concentration, driven by the higher reflux rate, which in turn reduces CO_2 concentration throughout the column. This affects solid-liquid and solid-vapor equilibria, reducing the freezing temperature at all stages.

CONCLUSIONS

The results show that the dual-pressure, low-temperature distillation process is a promising approach for LNG production in CO_2-rich gas fields. This product has a methane molar fraction of up to 99.1% and a CO_2 content lower than 50 ppmv. Reducing the CO_2 content at certain stages of the low-pressure column can prevent the formation of solids. This can be achieved by considering an outlet stream at the stage, increasing the column reflux rate, and adding light hydrocarbon streams. Additionally, the presence of a separator upstream of the distillation stage is key to producing a liquid stream of light hydrocarbons that meets the commercial requirement of an

Reid vapor pressure between 0.76 and 0.80 bar.

ACKNOWLEDGEMENTS

We gratefully acknowledge PETRONAS Petróleo Brasil Ltda. (PPBL) for their funding and support, as well as ANP (Brazil's National Oil, Natural Gas, and Biofuels Agency) for assistance through the R&D levy regulation (Project #25186-8). We also thank EQS – Sustainable Energy and Chemistry Laboratory, PQGe – Chemical Processes and Business Management Laboratory, the Center for Energy and Petroleum Studies (CEPETRO), and the School of Chemical Engineering (FEQ) at the University of Campinas (UNICAMP), Brazil. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

AUTHOR IDENTIFIERS

Author ORCIDs:
Victor S. V. Mercado: 0009-0002-1455-1623
Dirceu Noriler: 0000-0001-5281-5226
Laura Plazas Tovar: 0000-0003-1001-9250
Radin Suhaib Salihuddin: 0009-0009-6070-9429
Amiza Bt Surmi: 0009-0005-8071-3602
Fadhli Hadana Rahman: 0009-0008-7801-2198
Jean Felipe Leal Silva: 0000-0002-9915-5859

REFERENCES

1. Gürsan C, de Gooyert V. The systemic impact of a transition fuel: does natural gas help or hinder the energy transition?. Renewable and Sustainable Energy Reviews 138:110552 (2021). https://doi.org/10.1016/j.rser.2020.110552
2. EIA: Electric Power Monthly - December 2025., Washington (2025).
3. U.S. Environmental Protection Agency: GHG Emission Factors Hub | US EPA, https://www.epa.gov/climateleadership/ghg-emission-factors-hub, last accessed 2026/01/08.
4. EIA: International Energy Statistics, https://www.eia.gov/international/data/world/natural-gas/dry-natural-gas-production?pd=3002&p=00g&u=0&f=A&v=mapbubble&a=&i=none&vo=value&t=C&g=000000000000000000000000000000000000000000000000000001&l=249ruvvvvvfvtvnvv1vrvvvvfvvvvvvfvvvou20evvvvvvvvvvnvvvs0008&s=315532800000&e=1672531200000, last accessed 2026/01/15.
5. IEA: Electricity generation in World. (2025).
6. GPSA: GPSA Engineering Data Book. Gas Processors Suppliers Association (2004).
7. Pellegrini LA, De Guido G, Langè S, Moioli S, Picutti B, Vergani P, Franzoni G, Brignoli F. The potential of a new distillation process for the upgrading of acid gas. Abu Dhabi International Petroleum Exhibition & Conference : (2016). https://doi.org/10.2118/183154-ms
8. Anugraha RP, Pratiwi VD, Renanto R, Juwari J, Islami AN, Bakhtiar MY. Techno-economic analysis of CO2 cryogenic distillation from high CO2 content gas field: a case study in indonesia. Chemical Engineering Research and Design 202:226-234 (2024). https://doi.org/10.1016/j.cherd.2023.12.035
9. Masoudi R, Nayak S, Panting ATP, Hoesni J, Tiranda H, Samsuri MNB, Diah MABM. Geological controls on distribution of CO2 off the coast of sarawak, offshore malaysia: understandings from regional study and petroleum system modeling. Offshore Technology Conference Asia : (2024). https://doi.org/10.4043/34657-ms
10. Koronful N, Peters K, Ali MF, Skulsangjuntr J, Jiang L, Kleine A, Basu D, Bencomo J, Hernandez J, Brink G. Carbon dioxide in reservoir gases: new insights from basin and petroleum system modeling. SPE Asia Pacific Oil and Gas Conference and Exhibition : (2018). https://doi.org/10.2118/192011-ms
11. Burgers WFJ, Northrop PS, Kheshgi HS, Valencia JA. Worldwide development potential for sour gas. Energy Procedia 4:2178-2184 (2011). https://doi.org/10.1016/j.egypro.2011.02.104
12. Meyer, L., Walley, D., Sutton, J., Schapper, M.: Abstract: Ca Voi Xanh Field, Block 118, Offshore Da Nang, Vietnam: Exploration Highlights and Development Challenges. 10 (2017).
13. Watson NM, Zwingmann N, Lemon NM, Tingate PR. ONSHORE OTWAY BASIN CARBON DIOXIDE ACCUMULATIONS: CO2-INDUCED DIAGENESIS IN NATURAL ANALOGUES FOR UNDERGROUND STORAGE OF GREENHOUSE GAS. The APPEA Journal 43:637 (2003). https://doi.org/10.1071/aj02036
14. Liu J, Zhou H, Pujol M, Selby D, Li J, Tian H. The bitumen formation and re-os characteristics of a co2-rich pre-salt gas reservoir of the kwanza basin, offshore angola. Marine and Petroleum Geology 143:105786 (2022). https://doi.org/10.1016/j.marpetgeo.2022.105786
15. de Freitas VA, Vital JCS, Rodrigues BR, Rodrigues R. Source rock potential, main depocenters, and CO2 occurrence in the pre-salt section of santos basin, southeast brazil. Journal of South American Earth Sciences 115:103760 (2022). https://doi.org/10.1016/j.jsames.2022.103760
16. Carlson, E.C.: Don't Gamble With Physical Properties For Simulations. Chem. Eng. Prog. (1996).
17. De Guido G, Langè S, Moioli S, Pellegrini LA.

Thermodynamic method for the prediction of solid CO2 formation from multicomponent mixtures. Process Safety and Environmental Protection 92:70-79 (2014). https://doi.org/10.1016/j.psep.2013.08.001
18. NAGAHAMA K, KONISHI H, HOSHINO D, HIRATA M. BINARY VAPOR-LIQUID EQUILIBRIA OF CARBON DIOXIDE-LIGHT HYDROCARBONS AT LOW TEMPERATURE. J. Chem. Eng. Japan / JCEJ 7:323-328 (1974). https://doi.org/10.1252/jcej.7.323
19. Webster LA, Kidnay AJ. Vapor?liquid equilibria for the methane?propane?carbon dioxide systems at 230 k and 270 k. J. Chem. Eng. Data 46:759-764 (2001). https://doi.org/10.1021/je000307d
20. Ohgaki K, Katayama T. Isothermal vapor-liquid equilibrium data for the ethane—carbon dioxide system at high pressures. Fluid Phase Equilibria 1:27-32 (1977). https://doi.org/10.1016/0378-3812(77)80023-x
21. ZareNezhad B. Prediction of CO2 freezing points for the mixtures of CO2-CH4 at cryogenic conditions of NGL extraction plants. Korean J. Chem. Eng. 23:827-831 (2006). https://doi.org/10.1007/bf02705935
22. Brown TS, Sloan ED, Kidnay AJ. Vapor-liquid equilibria in the nitrogen + carbon dioxide + ethane system. Fluid Phase Equilibria 51:299-313 (1989). https://doi.org/10.1016/0378-3812(89)80372-3
23. Le TT, Trebble MA. Measurement of carbon dioxide freezing in mixtures of methane, ethane, and nitrogen in the solid?vapor equilibrium region. J. Chem. Eng. Data 52:683-686 (2007). https://doi.org/10.1021/je060194j
24. Neumann A, Walch W. Dampf/flüssigkeits?gleichgewicht co₂/ch₄ im bereich tiefer temperaturen und kleiner co₂?molenbrüche. Chemie Ingenieur Technik 40:241-244 (2004). https://doi.org/10.1002/cite.330400509
25. Wei MSW, Brown TS, Kidnay AJ, Sloan ED. Vapor + liquid equilibria for the ternary system methane + ethane + carbon dioxide at 230 k and its constituent binaries at temperatures from 207 to 270 k. J. Chem. Eng. Data 40:726-731 (2002). https://doi.org/10.1021/je00020a002
26. El-Banbi A, Alzahabi A, El-Maraghi A. Dry gases. PVT Property Correlations :29-63 (2018). https://doi.org/10.1016/b978-0-12-812572-4.00003-5
27. Xu X, Chen H, Liu C, Dang C. Prediction of the binary interaction parameter of carbon dioxide/alkanes mixtures in the pseudocritical region. ACS Omega 4:13279-13294 (2019). https://doi.org/10.1021/acsomega.9b01450
28. Chueh PL, Prausnitz JM. Vapor?liquid equilibria at high pressures: calculation of partial molar volumes in nonpolar liquid mixtures. AIChE Journal 13:1099-1107 (2004). https://doi.org/10.1002/aic.690130612
29. Stewart M, Arnold K. Crude stabilization. Emulsions and Oil Treating Equipment :81-106 (2009). https://doi.org/10.1016/b978-0-7506-8970-0.00002-5
30. ASTM: Standard Test Method for Vapor Pressure of Petroleum Products (Reid Method), ASTM D323 – 20a. (2020).
31. Pribylov AA, Skibitskaya NA, Zekel´ LA. Methane solubility in liquid hydrocarbons at high pressure. Russ Chem Bull 64:841-845 (2015). https://doi.org/10.1007/s11172-015-0942-2
32. Shashi Menon E. Physical properties. Transmission Pipeline Calculations and Simulations Manual :29-82 (2015). https://doi.org/10.1016/b978-1-85617-830-3.00003-1
33. GPSA: GPSA Engineering Data Book. Gas Processors Suppliers Association, Paris (1998).
34. Lastari, F.: Ryan-Holmes and Modified Ryan-Holmes Processes for LNG Production Fonny Lastari. (2009).
35. Haut RC, Denton RD, Thomas ER. Development and application of the controlled-freeze-zone process. SPE Production Engineering 4:265-271 (1989). https://doi.org/10.2118/17700-pa

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Multi-scale Metabolic Modeling and Simulation

Peter E. Carstensen[a], Teddy Groves[b], Lars K. Nielsen[b], Ulrich Krühne[c], Krist V. Gernaey[c], and John B. Jørgensen[a]*

[a] Technical University of Denmark, Department of Applied Mathematics and Computer Science, Kgs. Lyngby, Denmark
[b] Technical University of Denmark, Novo Nordisk Foundation Biotechnology Research Institute for the Green Transition, Kgs. Lyngby, Denmark
[c] Technical University of Denmark, Department of Chemical and Biochemical Engineering, Kgs. Lyngby, Denmark
* Corresponding Author: jbjo@dtu.dk.

ABSTRACT

Biological systems are governed by coupled interactions between intracellular metabolism and bioreactor operation that span multiple time scales. Constraint-based metabolic models are widely used to describe intracellular metabolism, but repeatedly solving the optimization problem at each time step in dynamic models introduces numerical challenges related to infeasibility and computational efficiency. This work presents a multi-scale modeling framework that integrates genome-scale, constraint-based metabolic models with dynamic bioreactor simulations. Intracellular metabolism is described using positive flux variables in a parsimonious flux balance analysis, and the resulting embedded optimization problem is replaced by a neural network surrogate. The surrogate provides a smooth approximation of the embedded optimization mapping and eliminates repeated linear program solves during simulation. The approach is demonstrated for fed-batch fermentation of *Escherichia coli*, in which the surrogate model yields intracellular fluxes under substrate-limited conditions, whereas the underlying linear program would otherwise be infeasible. The framework provides a continuous representation of intracellular metabolism suitable for dynamic simulation of genome-scale models in bioreactor configurations.

Keywords: Machine Learning, Multiscale Modelling, Modelling and Simulations, Surrogate Model, Dynamic Modelling

INTRODUCTION

Biological systems are governed by coupled interactions between cellular metabolism and process operation. These interactions span multiple time scales, from the rapid kinetics of intracellular reactions to the slower dynamics of bioreactor processes. Enzymatic reactions typically occur on the order of milliseconds, transcription and translation evolve over seconds to minutes, and bioreactor-scale phenomena develop over hours to days [1]. Modeling such systems requires a consistent representation of processes across these time scales [1, 2].

Constraint-based metabolic models are widely used to describe intracellular metabolism [3, 4] and are commonly embedded in dynamic reactor simulations through dynamic flux balance analysis (dFBA) [5-7]. This formulation links genome-scale metabolism with extracellular mass balances. Numerical challenges arise when the optimization problem approaches infeasible regions, for example, when maintenance requirements cannot be satisfied [5, 6, 8]. Changes in active constraints induce a piecewise-defined dependence of intracellular fluxes on extracellular conditions, potentially leading to infeasibility in substrate-limited regions [9].

In this work, a multi-scale modeling framework is presented that integrates genome-scale, constraint-based metabolic models with dynamic bioreactor simulations by replacing the embedded optimization problem with a neural network surrogate [10]. The framework consists of three coupled components. A reactor model describing flow, mixing, and volume dynamics. An extracellular reaction model defining mass balances for substrates, biomass, and products, and a kinetic cell model providing intracellular metabolic fluxes as functions of

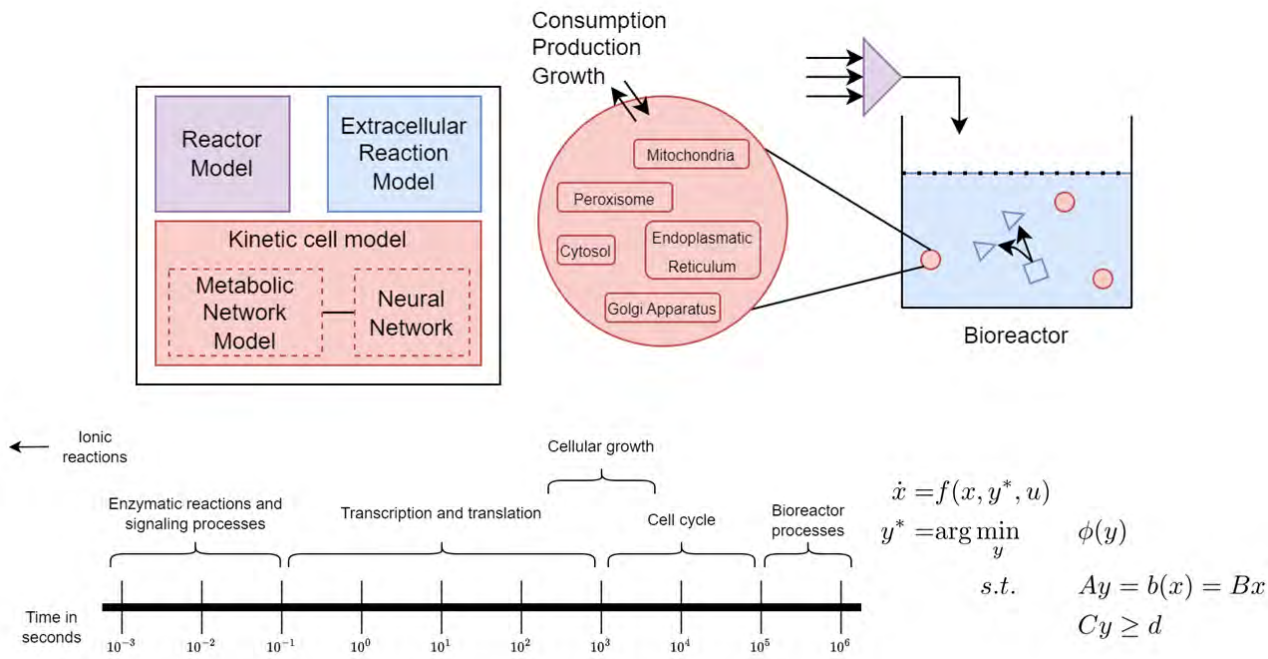

Figure 1. Overview of the multi-scale modelling framework linking intracellular metabolism to bioreactor dynamics. The upper panel shows the model structure, comprising the reactor model, the extracellular reaction model, and the kinetic cell model. The kinetic cell model describes intracellular metabolism using a constraint-based metabolic network formulation that a neural network surrogate can approximate [10,11]. The lower panel illustrates the wide separation of time scales in biological and process phenomena, ranging from picoseconds for ionic reactions to days for reactor operation [1]. Together, the three submodels capture the interaction between substrate consumption, product formation, and cellular growth across multiple time scales.

the extracellular conditions, as shown in Figure 1. Before surrogate replacement, the multi-scale model is expressed as a coupled dynamic optimization problem. Extracellular states evolve according to dynamic mass balances, whereas intracellular fluxes are obtained by solving an optimization problem subject to stoichiometric constraints, flux bounds, and a biological objective. Intracellular metabolism is assumed to operate in a quasi-steady state relative to the reactor dynamics, reflecting the separation of time scales between fast metabolic reactions and slower process-level phenomena [1, 5, 7].

The formulation is demonstrated using fed-batch fermentation of *Escherichia coli* with the genome-scale model iCH360 [3]. The coupled formulation links genome-scale metabolic models with reactor-scale dynamics within a single differentiable model.

The paper is structured as follows. The Method section presents the multi-scale modeling formulation, including the reactor model, extracellular mass balances, and kinetic cell model, as well as a description of the constraint-based metabolic formulation and the neural network surrogate model. The Results section presents simulation results for fed-batch fermentation of *E. coli* and a sensitivity analysis, followed by a discussion of model behavior and limitations, and a brief conclusion.

METHOD

Dynamic modeling

The multi-scale metabolic system is described by a set of ordinary differential equations coupled with optimal intracellular reaction rates. The extracellular dynamics depend on optimal rates obtained from an embedded optimization problem, or a surrogate model. The coupled system is written as,

$$\dot{x} = f(x, y^*, u) \tag{1a}$$
$$y^*(x) = y^* = \arg\min_y \quad \phi(y) \tag{1b}$$
$$s.t. \quad Ay = b(x) = Bx \tag{1c}$$
$$Cy \geq 0 \tag{1d}$$

x denotes the extracellular state variables, y^* is the optimal rates, as a function of the extracellular conditions, and u is the manipulated inputs. A surrogate representation may replace the embedded optimization problem in Equation (1b-1d) without altering the structure of the

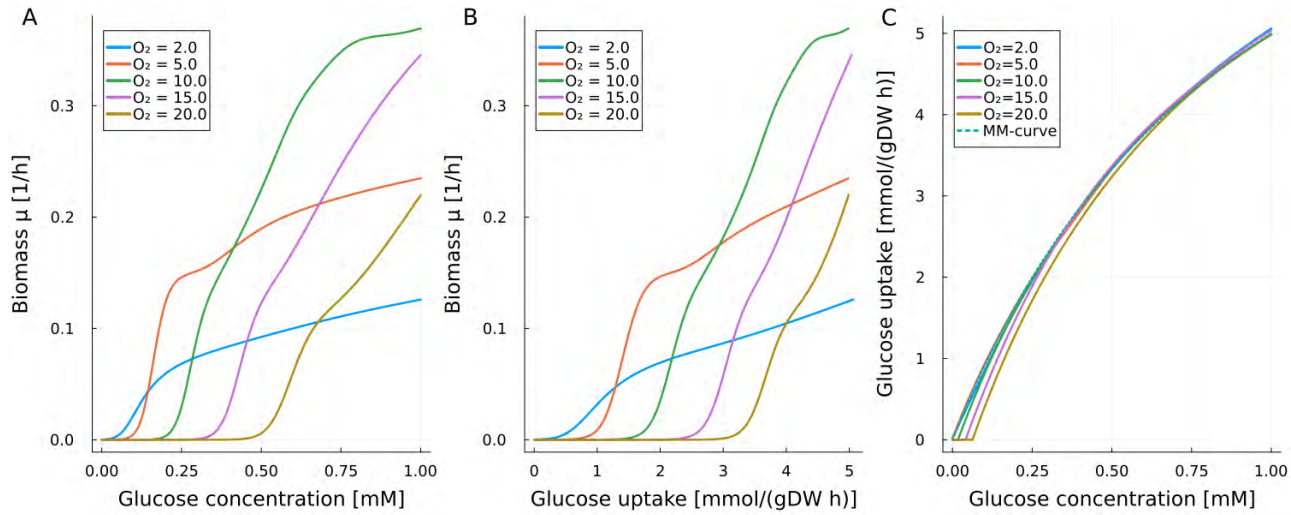

Figure 2. Surrogate model predictions of intracellular metabolism for the *E. coli* genome scale metabolic model. Biomass growth rate as a function of extracellular glucose concentration (A) for different fixed oxygen uptake rates, with growth decreasing as glucose availability is reduced. Biomass growth rate as a function of glucose uptake rate (B) at fixed oxygen uptake rate. Glucose uptake rate as a function of extracellular glucose concentration (C), with uptake decreasing towards zero at low glucose concentrations. The dotted line in (C) shows a Michaelis-Menten curve for comparison with $V_{max} = 10$ and $K_m = 1.0$.

remaining submodels.

Reactor modeling

The reactor model is modelled as a fed-batch system.

$$\dot{V} = e^T F_{in} \tag{2a}$$
$$\dot{n}_e = C_{in} F_{in} + R^*_{ext}(c_e)V, \quad c_e = \frac{n_e}{V} \tag{2b}$$
$$\dot{n}_X = R^*_X(c_e)V, \qquad c_X = \frac{n_X}{V} \tag{2c}$$
$$R^*_{ext}(c_e) = \nu^T_{ext} r^*(c_e) \tag{2d}$$
$$R^*_X(c_e) = \nu^T_X r^*(c_e) \tag{2e}$$
$$r^*(c_e) = \mu^*(c_e) c_X \tag{2f}$$

V is the volume, F_{in} is the inlet flow rate, n_e is the mole amount of extracellular species, C_{in} denotes inlet concentration. The vector $R^*_{ext}(c_e)$ represents the extracellular production and consumption rates, obtained from the intracellular model. n_X is the amount of biomass, and $R^*_X(c_e)$ is the production rate of biomass. ν_{ext} and ν_X are the stoichiometric matrices, $r^*(c_e)$ is the reaction rate vector, and $\mu^*(c_e)$ is the flux per unit of biomass and is a function of the extracellular concentration.

Constraint-based modeling

Flux balance analysis assumes that intracellular metabolism operates at steady-state and computes a flux distribution consistent with stoichiometric constraints and the specified objective [3, 12]. To account for changing extracellular conditions, uptake rates are fixed as functions of extracellular concentrations, denoted $\bar{\mu}_{in}(c_e)$. Embedding FBA within the reactor model yields a dFBA formulation [5, 12]. Intracellular reactions are assumed to equilibrate instantaneously relative to reactor-scale dynamics [5, 7]. The FBA problem is reformulated using positive forward and reverse fluxes to ensure directional consistency.

$$\mu^*(c_e) = \mu^* = \min_{\mu} \quad c^T(\mu^+ - \mu^-) \tag{3a}$$
$$s.t. \quad R_c = \nu_c^T(\mu^+ - \mu^-) = 0 \tag{3b}$$
$$0 \le \mu^+ \le \mu^+_{max} \tag{3c}$$
$$0 \le \mu^- \le \mu^-_{max} \tag{3d}$$
$$\mu^+_{in} - \mu^-_{in} = \bar{\mu}_{in}(c_e) \tag{3e}$$

μ is the vector of fluxes, decomposed into forward and reverse components, μ^+ and μ^-. The vector c defines the linear objective function, the matrix ν_c is the intracellular stoichiometric matrix, and the constraint $R_c = 0$ enforces steady-state balances for intracellular metabolites. The vectors μ^+_{max} and μ^-_{max} define the upper bounds for the forward and reverse fluxes. The FBA solution is generally non-unique, as the steady-state stoichiometric constraints define a high-dimensional feasible set [7, 13]. Linear optimization selects a single point from the optimal face, so parsimonious FBA (pFBA) is applied to identify flux distributions with minimal total flux, while maintaining the optimal growth rate z^*_{FBA} [14]. The pFBA optimization problem is defined as,

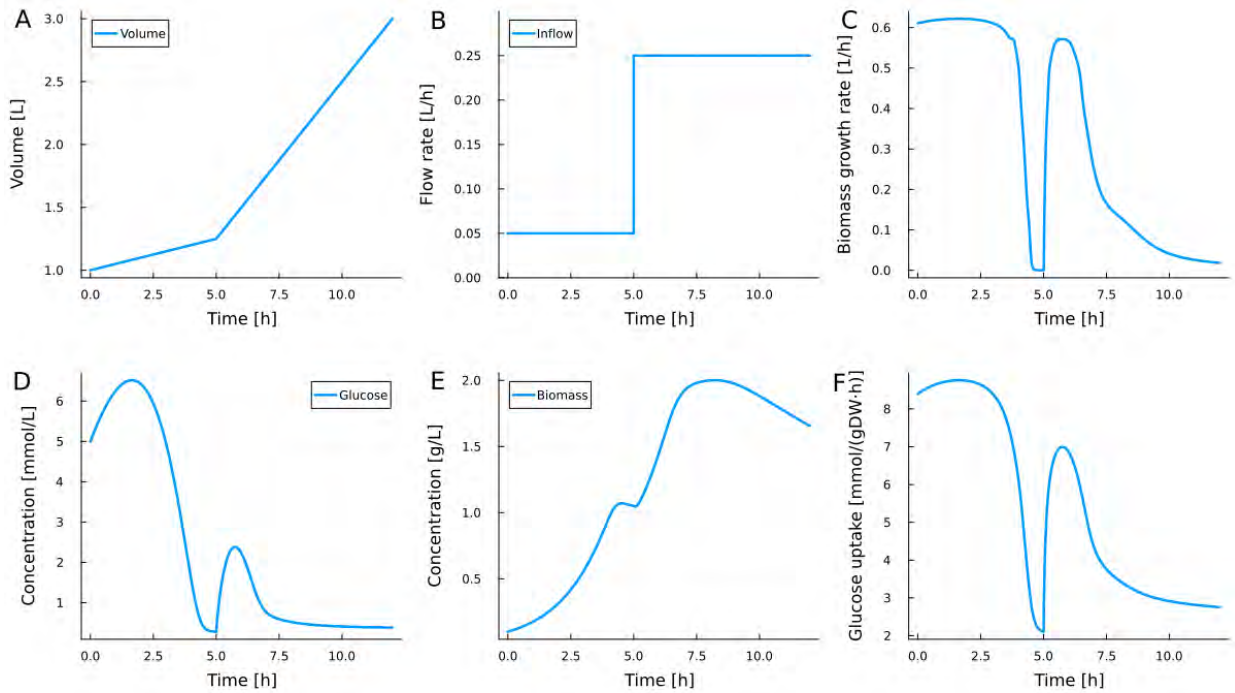

Figure 3. Dynamic fed-batch simulation of *E. coli* using the proposed multi-scale modeling framework with a neural network surrogate model. The reactor volume (A) increases over time in response to the feed profile (B). Biomass growth rates (C) are initially high and subsequently decrease as extracellular glucose (D) is consumed. This results in biomass accumulation (E), while glucose uptake rates (F) decrease but remain non-zero due to the continuous glucose feed. Glucose uptake follows a Michaelis-Menten dependence on extracellular concentration. Biomass growth ceases before complete substrate depletion when substrate uptake is insufficient to sustain growth. Oxygen uptake is kept constant at 15 mmol/(gDW h).

$$\mu^*(c_e) = \mu^* = \min_{\mu} \quad \sum_{i=1}^{n_\mu} \mu_i^+ + \mu_i^- \tag{4a}$$

$$s.t. \quad R_c = \nu_c^T(\mu^+ - \mu^-) = 0 \tag{4b}$$

$$0 \le \mu^+ \le \mu_{max}^+ \tag{4c}$$

$$0 \le \mu^- \le \mu_{max}^- \tag{4d}$$

$$\mu_{in}^+ - \mu_{in}^- = \bar{\mu}_{in}(c_e) \tag{4e}$$

$$\mu_X = z_{FBA}^* \tag{4f}$$

n_μ is the total number of reactions, and μ_X is the biomass growth flux.

Surrogate modeling

During dynamic simulation, substrate depletion may render the linear program infeasible or lead to discontinuous flux trajectories [8]. To obtain a smooth representation of intracellular fluxes, a neural network surrogate is used to approximate the pFBA solution map [9, 10]. The coupled system with the surrogate model can then be written as,

$$\dot{x} = f(x, y^*, u) \tag{5a}$$

$$y^*(x) = y^* = NN(\theta) \tag{5b}$$

$NN(\theta)$ is the trained neural network parameterized by θ. The surrogate model is trained on sampled uptake rates and the corresponding optimal flux distributions obtained from pFBA. In regions where the original pFBA formulation becomes infeasible, feasibility is maintained by introducing slack variables [15]. The surrogate is trained to approximate the resulting relaxed solution map, with fluxes decreasing smoothly towards zero. The surrogate model thus provides a smooth approximation of the optimal flux distribution within the sampled training region.

Training data consist of non-negative parsimonious fluxes. The surrogate is implemented in Julia [16] using Flux.jl [17, 18] and was trained using an NVIDIA H100 GPU and an AMD EPYC 9354 CPU. Training was performed on 40.401 pFBA solutions, each consisting of 349 intracellular fluxes. The pFBA solutions were generated from a uniform grid sampling over glucose and oxygen uptake rates from 0-20 mmol/(gDW h) with a step size of 0.1. The model was trained for approximately 15.000 epochs with early stopping.

Figure 2 shows surrogate model predictions of intracellular metabolism in *E. coli* under various conditions.

Panel A presents how the biomass growth rate changes as a function of the extracellular glucose concentration. In the low-substrate regions shown, growth decreases as glucose availability decreases and approaches zero. Higher oxygen uptake rates (15 and 20 mmol/(gDW h)) have not plateaued and can sustain higher biomass growth at higher glucose concentrations. Panel B shows the predicted biomass growth rate as a function of glucose uptake rate at a fixed oxygen uptake rate. Panel C shows the predicted glucose uptake rate as a function of extracellular glucose concentration. Uptake decreases smoothly towards zero as glucose concentration decreases and approximates the Michaelis-Menten kinetic equation (dotted line for comparison). When there is not enough glucose available for a fixed oxygen uptake, the glucose uptake rate smoothly approaches zero. Under fermentation conditions, neither the substrate nor oxygen uptake is fixed, but instead adapts to an optimum at the conditions the cell is in.

Figure 3 shows a dynamic fed-batch simulation of *E. coli* using the multi-scale modeling framework with the neural network surrogate. Panels A and B show the reactor volume and inflow, with a step increase in the inflow rate after 5 hours. Panel C shows the growth rate predicted by the surrogate model, while panels D and E show the glucose and biomass concentration in the bioreactor. Initially, the biomass concentration increases at a maximum rate under a fixed oxygen uptake. As the glucose concentration decreases, both the growth rate and glucose uptake shown in panel F decline. Following the step change in inflow, the reactor volume and glucose concentration increase, thereby increasing substrate availability. As a result, the glucose uptake, biomass growth rate, and biomass concentration increase. Once biomass accumulation is sufficient to take up more substrate than is supplied to the bioreactor, glucose concentration and biomass growth rate decrease. Here, the biomass concentration decreases due to dilution, while the total amount of biomass in the reactor continues to increase until the biomass growth rate is zero. At the step change, the non-surrogate model would have encountered an infeasible solution, and the simulation would have been terminated before additional feed [7].

Surrogate and pFBA flux predictions are compared in Figure 4 at glucose and oxygen uptake rates of 10 mmol/(gDW h). Panel A compares the total absolute intracellular flux, summed across all reactions, and the surrogate, which slightly underpredicts the sum of fluxes. Panel B shows the total flux related to glycolysis, specifically, the glycolysis reactions in Figure 5. In both cases, the surrogate reproduces the magnitude of the fluxes. Panels C and D show the parity plot comparing per-reaction flux predictions. Panel C includes all reactions, while Panel D consists of the glycolysis reactions in Figure 5. The low residuals and linear correlation between the surrogate and the pFBA flux predictions indicate a good fit for individual reactions.

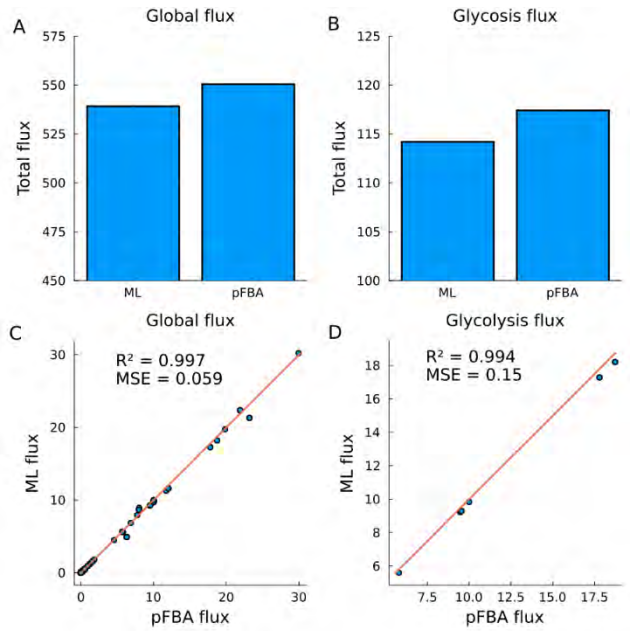

Figure 4. Comparison of surrogate and pFBA flux predictions using the *E. coli* genome-scale model. (A) Total absolute intracellular flux magnitude and (B) total glycolysis flux magnitude predicted by the surrogate and the pFBA at fixed glucose and oxygen uptake rates. (C-D) Per reaction flux comparison between surrogate and pFBA for the full network and the glycolysis pathway.

Results

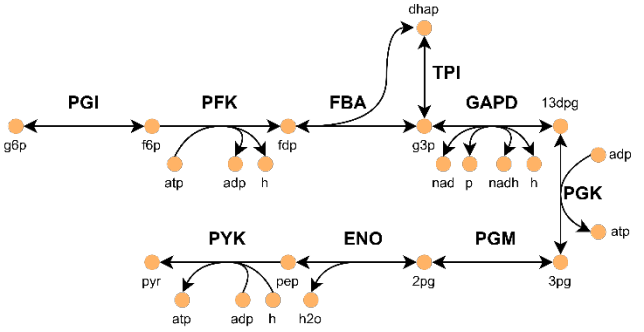

Figure 5. Graphical overview of the condensed glycolysis reactions used in Figure 4 (B, D). The enzyme abbreviations are in bold, and the yellow dots are intracellular species.

It should be noted that the uptake rates used in Figures 3 and 4 are included in the surrogate training data. The surrogate model should be applied only within the region of the input space covered by the training data, as it does not explicitly enforce mass-balance constraints or upper and lower bounds. Instead, the surrogate model reproduces empirical relationships observed during training.

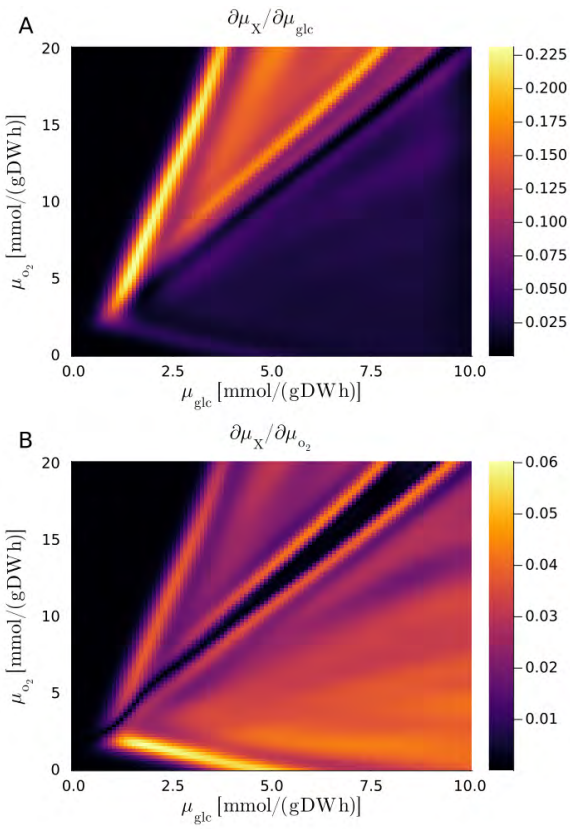

Figure 6. Sensitivity of the surrogate predicted biomass production rate μ_X with respect to glucose uptake, μ_{glc}, and oxygen uptake, μ_{o_2}.

Figure 6 shows sensitivity measures for the surrogate intracellular metabolism across the glucose and oxygen uptake range. Panel A shows the partial derivative of the biomass production rate with respect to glucose uptake, $\partial\mu_X/\partial\mu_{glc}$. High sensitivity is observed along a narrow diagonal region, where concurrent increases in glucose and oxygen uptake lead to increased biomass production. Sensitivity decreases outside this region and approaches zero in low-uptake or oxygen-limited regions. Panel B shows the partial derivative of the biomass production rate with respect to oxygen uptake, $\partial\mu_X/\partial\mu_{o_2}$. Similar regions are observed, however, the overall magnitude of the oxygen sensitivity is lower and more broadly distributed than that of glucose sensitivity.

DISCUSSION

In the iCH360 model, there is a minimum energetic requirement through the non-growth-associated maintenance reaction [3, 13]. At sufficiently low substrate uptake rates, this requirement cannot be met, and the pFBA problem becomes infeasible [8]. To have training samples, slack variables are introduced, and the surrogate constitutes a relaxed pFBA solution in these regions. The surrogate was trained using glucose and oxygen uptake rates as inputs, assuming Michaelis-Menten uptake kinetics without inhibition [1]. Mass-balance constraints and flux bounds are not explicitly enforced and are reflected only through the training data. Additional inputs, such as mineral uptake or alternative substrates, would require retraining [10, 19]. However, this retraining is performed offline and does not affect computational cost during simulation.

For the iCH360 model, solving the pFBA problem requires approximately 2.33 ms in Julia, where two linear programs are solved sequentially using Gurobi at fixed glucose and oxygen uptake rates. A classical FBA formulation requires approximately 120 μs. In contrast, the surrogate model predicts via a forward pass consisting of matrix-vector multiplications and nonlinear activation functions, in approximately 60 μs. This is 39 times faster than the pFBA-based formulation, and twice as fast as the classical FBA formulation.

The genome-scale model contains 349 intracellular fluxes, all of which were included in the training and validation datasets. Deviations between surrogate and pFBA predictions were generally small. The largest discrepancy between the pFBA solution and the surrogate prediction was observed for the hydrogen exchange reaction. Here, an absolute difference of 1.75 mmol/(gDW h), approximately 8%, was observed at a fixed glucose and oxygen uptake rate of 10 mmol/(gDW h). Increasing training coverage, adding more data points, or employing more advanced machine learning methods may reduce discrepancies [10, 19].

When implemented in Julia [16] using Flux.jl [17, 18], automatic differentiation is used to compute Jacobians and Hessians of intracellular fluxes with respect to states and parameters. State and parameter sensitivities are directly accessible through the formulation as well as for Julia's differential equation solvers. These sensitivities may be used in gradient-based state estimation and control [9].

CONCLUSION

This work extends dynamic FBA by replacing the embedded linear optimization problem with a surrogate model of intracellular metabolism. By approximating the parsimonious FBA solution, the surrogate enables dynamic simulations in operating regions, where the original optimization problem may become infeasible. The surrogate produces continuous intracellular flux trajectories and eliminates repeated optimization solves during simulation. Extension to additional substrates and operating conditions can be achieved through offline retraining. The surrogate provides a computationally efficient basis for integration into control frameworks to optimize fermentation processes.

ACKNOWLEDGEMENTS

Peter E. Carstensen is funded by The Novo Nordisk Foundation Center for Biosustainability, Technical University of Denmark, NNF20CC0035580.

AUTHOR IDENTIFIERS

Author ORCIDs:
Carstensen, PE: 0009-0008-9019-7607
Groves, T: 0000-0002-7109-3270
Nielsen, LK: 0000-0001-8191-3511
Krühne, U: 0000-0001-7774-7442
Gernaey, K: 0000-0002-0364-1773
Jørgensen, JB: 0000-0001-9799-2808

REFERENCES

1. Yasemi M, Jolicoeur M. Modelling cell metabolism: a review on constraint-based steady-state and kinetic approaches. Processes 9:322 (2021). https://doi.org/10.3390/pr9020322
2. Mahnert C, Oyarzún DA, Berrios J. Multiscale modelling of bioprocess dynamics and cellular growth. Microb Cell Fact 23: (2024). https://doi.org/10.1186/s12934-024-02581-0
3. Corrao M, He H, Liebermeister W, Noor E, Bar-Even A. A compact model of escherichia coli core and biosynthetic metabolism. PLoS Comput Biol 21:e1013564 (2025). https://doi.org/10.1371/journal.pcbi.1013564
4. Orth JD, Fleming RMT, Palsson BØ. Reconstruction and use of microbial metabolic networks: the core escherichia coli metabolic model as an educational guide. EcoSal Plus 4: (2010). https://doi.org/10.1128/ecosalplus.10.2.1
5. Mahadevan R, Edwards JS, Doyle FJ III. Dynamic flux balance analysis of diauxic growth in escherichia coli. Biophysical Journal 83:1331-1340 (2002). https://doi.org/10.1016/s0006-3495(02)73903-9
6. Gomez JA, Höffner K, Barton PI. Dfbalab: a fast and reliable MATLAB code for dynamic flux balance analysis. BMC Bioinformatics 15: (2014). https://doi.org/10.1186/s12859-014-0409-8
7. The Economic Cell Collective. Economic Principles in Cell Biology. (2026).
8. Klamt S, von Kamp A. Analyzing and resolving infeasibility in flux balance analysis of metabolic networks. Metabolites 12:585 (2022). https://doi.org/10.3390/metabo12070585
9. Scott F, Wilson P, Conejeros R, Vassiliadis VS. Simulation and optimization of dynamic flux balance analysis models using an interior point method reformulation. Computers & Chemical Engineering 119:152-170 (2018). https://doi.org/10.1016/j.compchemeng.2018.08.041
10. Faure L, Mollet B, Liebermeister W, Faulon JL. A neural-mechanistic hybrid approach improving the predictive power of genome-scale metabolic models. Nat Commun 14: (2023). https://doi.org/10.1038/s41467-023-40380-0
11. Sahu A, Blätke MA, Szyma?ski JJ, Töpfer N. Advances in flux balance analysis by integrating machine learning and mechanism-based models. Comput Struct Biotechnol J 19:4626-4640 (2021). https://doi.org/10.1016/j.csbj.2021.08.004
12. Varma A, Palsson BO. Stoichiometric flux balance models quantitatively predict growth and metabolic by-product secretion in wild-type escherichia coli W3110. Appl Environ Microbiol 60:3724-3731 (1994). https://doi.org/10.1128/aem.60.10.3724-3731.1994
13. Palsson BØ. Systems biology. Cambridge University Press (2015). https://doi.org/10.1017/cbo9781139854610
14. Lewis NE, Hixson KK, Conrad TM, Lerman JA, Charusanti P, Polpitiya AD, Adkins JN, Schramm G, Purvine SO, Lopez?Ferrer D, Weitz KK, Eils R, König R, Smith RD, Palsson BØ. Omic data from evolved e. coli are consistent with computed optimal growth from genome?scale models. Molecular Systems Biology 6: (2010). https://doi.org/10.1038/msb.2010.47
15. J. Nocedal and S. J. Wright, Numerical Optimization. Springer, New York, (2006).
16. Bezanson J, Edelman A, Karpinski S, Shah VB. Julia: a fresh approach to numerical computing. SIAM Rev. 59:65-98 (2017). https://doi.org/10.1137/141000671
17. Gandhi D, Innes M, Saba E, Fischer K, Shah V. Julia e flux: modernizando o aprendizado de máquina. Computação Brasil :41-45 (2019). https://doi.org/10.5753/compbr.2019.39.4527
18. Innes M. Flux: elegant machine learning with julia. JOSS 3:602 (2018). https://doi.org/10.21105/joss.00602
19. Merzbacher C, Mac Aodha O, Oyarzún DA. Modeling host–pathway dynamics at the genome scale with machine learning. Metabolic Engineering 91:480-491 (2025). https://doi.org/10.1016/j.ymben.2025.05.008

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Enhancing Parameter Identifiability in Capacitive Deionization: A Model-Based Design of Experiments Approach

Yuxuan Yang[a], Federico Glavanin[a]*

[a] University College London, Department of Chemical Engineering, London, United Kingdom
* Corresponding Author: f.galvanin@ucl.ac.uk

ABSTRACT

Capacitive Deionization (CDI) is an emerging electrochemical technology for energy-efficient brackish water desalination. However, the rigorous design and scale-up of CDI systems are frequently hindered by the complexity of validating predictive models. The coupling of electrochemical double-layer kinetics with macroscopic mass transport often leads to structural parameter correlations, where multiple combinations of kinetic rates yield indistinguishable effluent trajectories. This paper addresses these challenges by proposing a simulation-driven Model-Based Design of Experiments (MBDoE) framework. We develop and implement a reduced-order Dynamic Langmuir (DL) model within the gPROMS platform, designed to capture cyclic adsorption-desorption dynamics with high computational efficiency. Sensitivity analysis reveals that information content is highly transient, concentrated primarily in the short time windows following voltage switching, and that the effluent concentration is significantly more sensitive to desorption kinetics than adsorption. Fisher Information Matrix (FIM) analysis of baseline experimental data confirms a strong negative correlation between kinetic parameters, resulting in a poorly conditioned estimation problem. To resolve this, a large-scale in-silico screening of the experimental design space—spanning inlet concentration, flow rate, and cell volume—is conducted using a D-optimality criterion. The simulation results demonstrate that operating at low flow rates and large effective cell volumes maximizes parameter identifiability by enhancing the separation of dynamic signatures. This work illustrates the critical role of dynamic simulation in guiding experimental strategy, minimizing trial-and-error effort, and improving the robustness of process models.

Keywords: Desalination, Capacitive Deionization, Modelling and Simulation, System Identification, Design of Experiment

1. INTRODUCTION

Freshwater scarcity is a growing global constraint, driving the demand for reliable and energy-efficient desalination technologies. Capacitive Deionization (CDI) has emerged as a promising alternative to reverse osmosis, particularly for the treatment of low-to-moderate salinity brackish water [1]. The technology operates on the principle of electrosorption, where dissolved ions are removed from a feed stream by applying a low voltage across porous carbon electrodes, storing them in electrical double layers (EDLs) [2]. Unlike steady-state membrane processes, CDI is inherently dynamic and cyclic, alternating between a charging or purification phase and a discharging or regeneration phase.

The optimization of CDI cycles—balancing salt removal, energy consumption, and water recovery—requires robust predictive models. However, the modelling of CDI systems presents a significant fidelity-cost challenge [3]. High-fidelity models, such as those based on Poisson-Nernst-Planck (PNP) equations or modified Donnan theory, offer detailed resolution of pore-scale electrochemistry but are often computationally prohibitive for real-time control or extensive design screening [4]. Conversely, simple equivalent-circuit models (RC circuits) are computationally inexpensive but lack mechanistic insight into mass transfer limitations, making them poor tools for process scale-up.

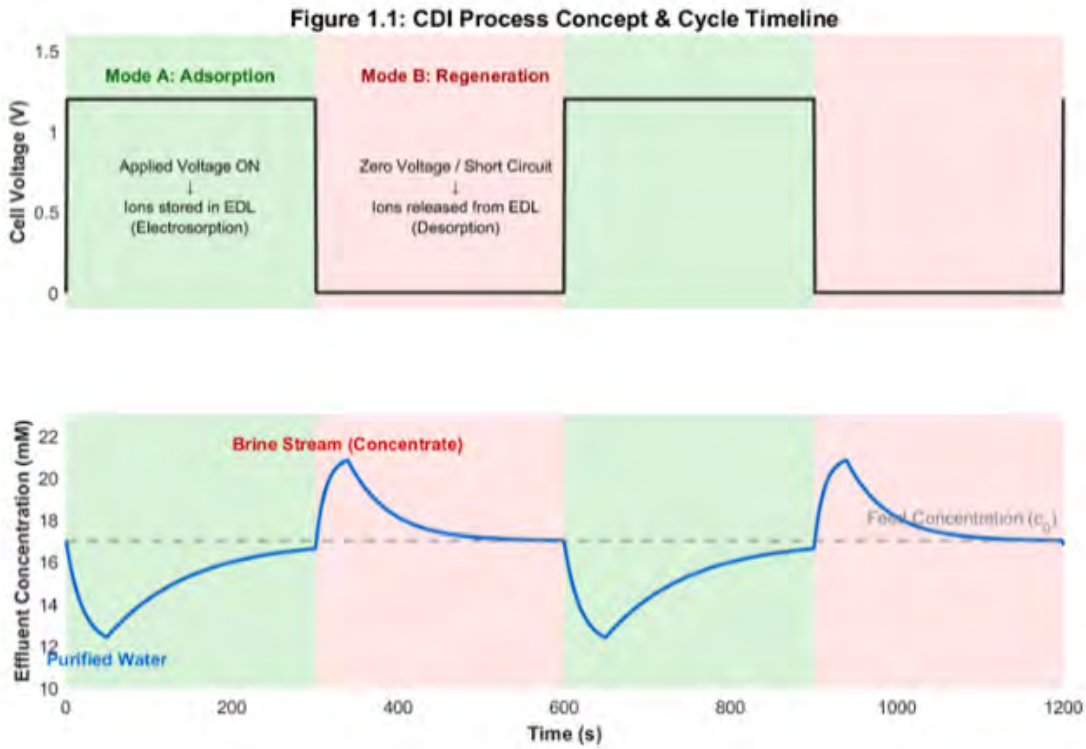

Figure 1. Ideal CDI process concept and timeline

An intermediate class of models, such as the Dynamic Langmuir (DL) model, offers a compelling compromise by describing the macroscopic adsorption dynamics with a compact set of parameters [5]. Yet, even simplified mechanistic models face the challenge of parameter identifiability. The coupling of reaction kinetics with fluid dynamics in a flow-through cell often leads to strong correlations between parameters. For instance, a faster adsorption rate might be mathematically indistinguishable from a different mass transfer coefficient under certain flow conditions. Standard experimental approaches, which often rely on normal one by one sweeps, is hard to generate data that can decouple these effects, leading to models with large predictive uncertainties [6].

Figure 1 shows the ideal CDI process concept and the adsorption regeneration cycle timeline. The top panel shows the applied cell voltage profile, which changes between a charging step (Mode A) where a constant voltage is applied to drive absorption, and a regeneration step (Mode B) where the voltage is removed or reversed. The bottom panel shows the corresponding effluent concentration response (blue line) relative to the constant feed concentration. During Mode A, ion removal causes the effluent concentration to drop, producing purified water. Conversely, in Mode B, the release of stored ions creates a concentration peak, resulting in a concentrated brine stream. These transient responses highlight why understanding dynamics and cycling is important to both system operation and model development.

Calibrating mechanistic models for CDI is challenging because different combinations of kinetic and transport parameters often yield indistinguishable effluent profiles, leading to poor parameter identifiability [7]. Traditional experimental design strategies, which typically rely on intuition or broad factorial sweeps, are often inefficient at resolving these correlations and can result in high experimental costs without reducing parameter uncertainty.

To address this, we employ Model-Based Design of Experiments (MBDoE) [8, 9]. Unlike classical design approaches that aim for broad empirical coverage, MBDoE uses the mechanistic model itself to quantitatively predict the information content of future experiments [10]. By analyzing dynamic sensitivities and the Fisher Information Matrix (FIM), MBDoE identifies specific operating windows, such as precise flow rates or cycle timings, that maximize the sensitivity of the measured outputs to the target parameters. This approach guides the selection of experimental conditions to decouple correlated parameters and minimize the variance of estimates with the least possible experimental effort [11]. Similar work in electrodialysis successfully minimise the number of trials and resources required for model identification [12].

This paper uses modelling and simulation as a tool

to solve this experimental design problem. We propose a MBDoE framework to systematically identify operating conditions that maximize the information content of collected data.

2. METHODOLOGY: MODELLING AND SIMULATION FRAMEWORK

2.1 The Dynamic Langmuir (DL) Model

To enable the rapid, repeated simulations required for extensive design screening ($N > 1000$ runs), we adopt a reduced version of the Dynamic Langmuir model. The original DL framework describes ion removal as a competitive adsorption process on voltage-induced sites.

The full kinetic description typically depends on the real-time free ion concentration inside the cell, $c(t)$. However, in this work, a key simplification to improve structural identifiability and computational speed was introduced [13]: It assume that for the purpose of the adsorption driving force, the cell concentration can be approximated by the constant inlet concentration, i.e., $c(t) \approx c_0$ in the kinetic rate term. This assumption transforms the nonlinear adsorption dynamics into a linear Ordinary Differential Equation (ODE) for the adsorbed concentration state, c_{ads} (mM)[14]:

$$\frac{dc_{ads}}{dt} = K_a - K_b c_{ads} \tag{1}$$

Here, K_a (mM/s) represents an effective adsorption rate that lumps together the site availability, voltage dependence, and feed concentration effects.K_b (s^{-1}) represents the effective desorption or relaxation rate constant.

The macroscopic mass balance for the cell is modelled as a Continuous Stirred Tank Reactor (CSTR). The change in free ion concentration c (mM) is driven by the net removal of ions to the electrode (adsorption sink) and the convective transport through the cell [14]:

$$\frac{dc}{dt} = -\frac{dc_{ads}}{dt} + \frac{Q}{V}(c_0 - c) \tag{2}$$

where Q is the volumetric flow rate (mL/s) and V is the effective free cell volume (mL). This structure retains the essential dynamic features of CDI, which is the rapid drop during charging and the concentration overshoot during regeneration, while reducing the estimation problem to two primary kinetic parameters, $\theta = [K_a, K_b]^T$.

2.2 Model-Based Design of Experiments (MBDoE) Formulation

The objective of the MBDoE framework is to determine the experimental design vector $\varphi = [c_0, Q, V]^T$ that minimizes the statistical uncertainty of the estimated parameters θ.

Sensitivity Analysis: The foundation of the analysis is the dynamic local sensitivity matrix $Q(t)$, which quantifies the influence of parameters on the measured output $y(t) = [c(t), c_{ads}(t)]^T$:

$$Q_{ij}(t) = \frac{\partial y_i(t)}{\partial \theta_j} \tag{3}$$

In this study, sensitivities were computed using a finite-difference perturbation method ($\delta = 1\%$) implemented directly within the simulator. Large sensitivity values indicate that the output is strongly responsive to the parameter, implying that the data contains significant information for estimation.

Fisher Information Matrix (FIM): To aggregate this information over the duration of a dynamic experiment (T), we compute the cumulative Fisher Information Matrix, $H(\theta)$. Assuming Gaussian measurement noise with variance σ^2, the FIM is defined as [8]:

$$H(\theta) = \frac{1}{\sigma^2}\int_0^T Q(t)^T Q(t) dt \tag{4}$$

The inverse of the FIM approximates the parameter covariance matrix ($Cov(\theta) \approx H^{-1}$). Therefore, "maximizing" the FIM corresponds to minimizing the variance and covariance of the parameters.

D-Optimality Criterion: Since the FIM is a matrix, a scalar metric is required for optimization. We employ the **D-optimality criterion**, which maximizes the determinant of the FIM [8]:

$$\Phi_D(\varphi) = \mathrm{logdet}\big(H(\varphi)\big) \tag{5}$$

Geometrically, maximizing the determinant minimizes the volume of the joint confidence ellipsoid in the parameter space. This is particularly critical in CDI modeling, where parameters are often highly correlated; D-optimality specifically penalizes designs that result in flattened, elongated confidence regions.

2.3 Implementation in gPROMS

The simulation framework was implemented in gPROMS ModelBuilder. The architecture was designed for scalability, consisting of two layers:

Unit Model: A rigorous dynamic model of a single CDI cell solving the mass balance and kinetic equations (Eqs. 1-2) under a cyclic voltage schedule].

Manager Model: A high-level wrapper that instantiates an array of unit models. This allows for the simultaneous execution of a full-factorial Design of Experiments (DoE) grid.

This architecture enabled the screening of $N_{exp} = 1000$ virtual experiments in a single simulation job, ensuring consistency across initial conditions and solver settings.

3. RESULTS AND DISCUSSION

3.1 Dynamic Behaviour and Sensitivity Analysis

The reduced model was first evaluated under baseline conditions ($c_0 = 17$ mM, $Q = 1$ mL/s, $V = 10$ mL) to establish the dynamic characteristics of the system. . The simulation successfully reproduces the classic CDI response: upon voltage application, the effluent concentration drops sharply as ions are electrosorbed, followed by a gradual saturation. During regeneration (voltage removal), a concentration peak (overshoot) is observed as stored ions are released back into the bulk

The agreement between the model predictions using the estimated parameters and the measured data is illustrated in Figure 2 The solid line represents the simulated effluent concentration, while the symbols with error bars denote the experimental measurements and their standard deviations based on the estimated variance model. The figure shows that the model captures both the rapid initial decrease in effluent concentration during adsorption and the subsequent recovery and slight overshoot during regeneration. All measured points lie within the ±1σ confidence bands, and no systematic deviation can be observed over time.

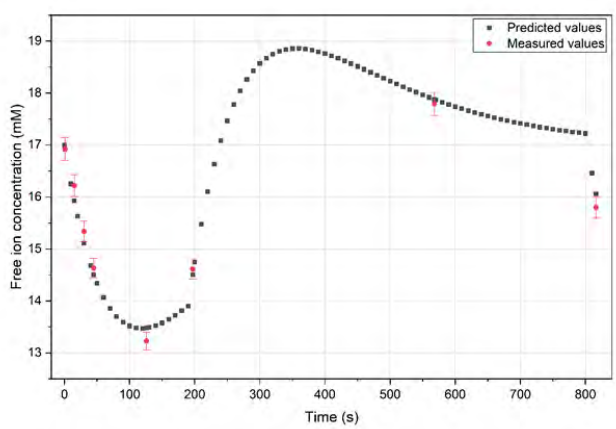

Figure 2. Comparison between measured and model-predicted effluent concentration c for the sample experiment

A critical finding from the sensitivity analysis is the disparity in parameter influence. The sensitivity of the effluent concentration to the desorption rate (K_b) is approximately an order of magnitude larger than that to the adsorption rate (K_a). Specifically, sensitivity to K_b peaks at values around 3×10^2, while sensitivity to K_a remains within ± 20.

Furthermore, the information content is highly localized in time. The sensitivity profiles exhibit sharp peaks immediately following the voltage switching events (charging and discharging steps) and decay to near zero during the steady-state phases. This temporal analysis suggests that measurements taken during the transient "switching windows" are exponentially more valuable than those taken at steady state. It confirms that dynamic operation is not just a feature of CDI but a necessity for accurate model calibration.

3.2 Parameter Estimation and Identifiability Issues

Using experimental data from a single charging cycle taken from literature, the kinetic parameters were estimated as $K_a \approx 0.0814$ mM/s and $K_b \approx 0.0071$ s^{-1}. While the model achieved a statistically valid fit (passing χ^2 and lack-of-fit tests), the FIM analysis reveals severe underlying identifiability issues.

The cumulative FIM at the end of the experiment showed a condition number of approximately 2.4×10^2, indicating a highly ill-conditioned estimation problem. The off-diagonal element of the covariance matrix revealed a correlation coefficient between K_a and K_b of $\rho \approx -0.83$. This indicates a strong negative correlation between Ka and Kb: an increase in one parameter can be partially compensated by a decrease in the other while still producing similar model responses. Geometrically, the confidence region in the (K_a, K_b)-plane is expected to be an elongated ellipse with a major axis tilted away from the coordinate axes.

Geometrically, this strong negative correlation results in a highly elongated joint confidence ellipse. This implies that K_a and K_b are coupled: an error in estimating the adsorption rate can be compensated by a proportional error in the desorption rate, leaving the model prediction largely unchanged. This "parameter compensation" effect is a primary cause of poor model predictive capability when operating conditions change.

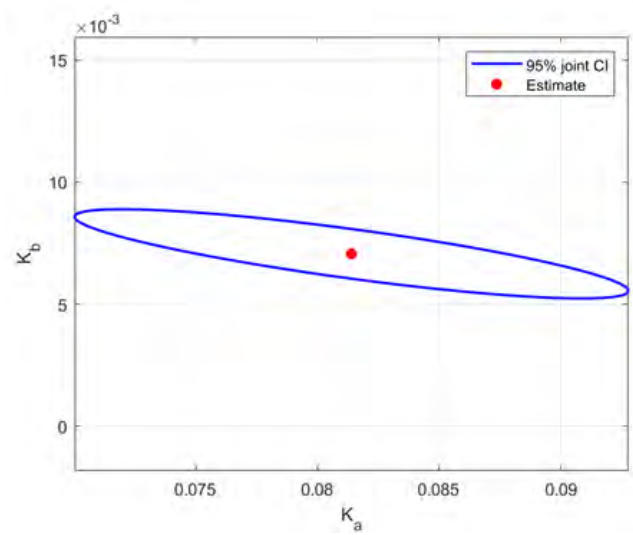

Figure 3. Joint 95% confidence region for the kinetic parameters $K_a$and K_b

Table 1. Estimated kinetic and variance-model parameters for the experiment data

Parameter	Final value	Initial guess	Std. dev.	95% CI lower	95% CI upper	t-value
K_a [units]	0.0814	0.0976	0.0046	0.059	0.119	6.82
K_b [s^{-1}]	0.00706	0.00880	0.00075	0.00530	0.00919	3.68
ω [–]	0.0130	0.0100	0.0049	0.0093	0.0185	1.10

Table 1 summarises the final parameter estimates, together with their standard deviations and 95% confidence intervals as reported by gPROMS. Starting from the literature-based initial guesses, the optimisation converged to slightly lower values for both kinetic parameters: K_a decreases from 0.0976 to approximately 0.081, while K_b decreases from 0.0088 to around 0.0071. These shifts indicate that the originally assumed rates moderately overestimated both the adsorption and desorption kinetics for the dataset considered.

3.3 Optimal Design Screening via Simulation

To break this parameter correlation, the simulation framework was used to screen the design space for conditions that maximize the D-optimality score. A full-factorial grid was evaluated:

- Inlet concentration $c_0 \in [0.1, 100]$ mM
- Flow rate $Q \in [0.05, 5.0]$ mL/s
- Cell volume $V \in [1, 100]$ mL.

The analysis of the resulting 1000 design points (Figure 4) yields three key insights for process engineering:

Invariance to Concentration: The D-optimality score showed negligible sensitivity to the inlet concentration c_0. This is consistent with the linear structure of the reduced model, suggesting that experimenters have flexibility in choosing c_0 based on safety or material constraints without compromising information quality.

Dominance of Flow Rate: The flow rate Q emerged as the most critical design factor, showing a strong negative correlation with the D-score ($r \approx -0.68$). Reducing the flow rate from 5.0 mL/s to 0.05 mL/s improved the determinant of the FIM by nearly a factor of three. Physically, lower flow rates increase the residence time of fluid in the cell. This allows the kinetic signature of the adsorption process to manifest more clearly against the background of convective transport, effectively "amplifying" the signal of the kinetic parameters.

Benefit of Cell Volume: Larger cell volumes were found to improve identifiability ($r \approx +0.31$). Similar to flow rate, a larger volume stretches the characteristic time of the mass balance, enhancing the separation of adsorption and desorption dynamic modes.

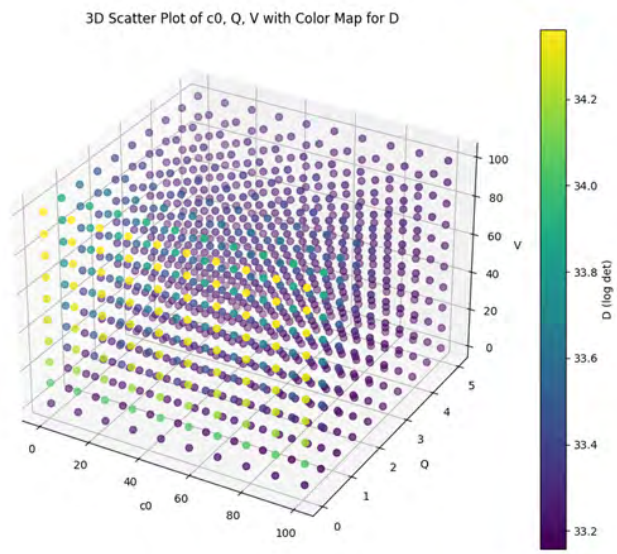

Figure 4. Three-dimensional scatter plot of the DOE design space (c_0, Q, V) coloured by the final D-criterion $D_{\log}(t_f)$, indicating the regions of highest and lowest information content for estimating K_a and K_b.

3.4 Recommended Design Strategy

Based on these simulation results, the optimal strategy for characterizing CDI kinetics is to operate at the **lowest feasible flow rate** and using a **larger effective cell volume**. The "worst-case" design was identified as a high-flow, small-volume configuration ($D_{log} \approx 33.16$), while the "best-case" design was a low-flow, large-volume configuration ($D_{log} \approx 34.36$). By shifting experiments to this optimal region, researchers can significantly reduce the volume of the parameter confidence ellipsoid, decoupling the kinetic rates and ensuring a more robust model calibration.

4. CONCLUSION

This study demonstrates the power of modelling and simulation as a core pillar of Computer Aided Process Engineering (CAPE) for emerging water technologies. By formulating a reduced-order dynamic model for Capacitive Deionization and integrating it with a rigorous MBDoE framework, we successfully quantified the information content of experimental data.

The simulation results highlight that "standard"

experimental conditions often yield poor parameter identifiability due to intrinsic correlations between adsorption and desorption kinetics. However, by using the model to navigate the design space, we identified that operating regimes with high residence times (low flow, large volume) maximize the D-optimality criterion. This approach moves beyond trial-and-error experimentation, providing a systematic workflow to minimize experimental effort while maximizing the reliability of process models. Future work will extend this framework to include multi-objective optimization, incorporating energy consumption and thermodynamic efficiency into the design criteria.

REFERENCES

1. Biesheuvel PM, Zhao R, Porada S, van der Wal A. Theory of membrane capacitive deionization including the effect of the electrode pore space. Journal of Colloid and Interface Science 360:239-248 (2011). https://doi.org/10.1016/j.jcis.2011.04.049
2. Suss ME, Porada S, Sun X, Biesheuvel PM, Yoon J, Presser V. Water desalination via capacitive deionization: what is it and what can we expect from it?. Energy Environ. Sci. 8:2296-2319 (2015). https://doi.org/10.1039/c5ee00519a
3. Porada S, Zhao R, van der Wal A, Presser V, Biesheuvel PM. Review on the science and technology of water desalination by capacitive deionization. Progress in Materials Science 58:1388-1442 (2013). https://doi.org/10.1016/j.pmatsci.2013.03.005
4. Hemmatifar A, Stadermann M, Santiago JG. Two-dimensional porous electrode model for capacitive deionization. J. Phys. Chem. C 119:24681-24694 (2015). https://doi.org/10.1021/acs.jpcc.5b05847
5. Nordstrand J, Dutta J. Dynamic langmuir model: a simpler approach to modeling capacitive deionization. J. Phys. Chem. C 123:16479-16485 (2019). https://doi.org/10.1021/acs.jpcc.9b04198
6. Lee S, Shim J, Kim HH, Yun N, Son M, Cho KH. Optimizing capacitive deionization operation using dynamic modeling and reinforcement learning. Desalination 602:118626 (2025). https://doi.org/10.1016/j.desal.2025.118626
7. Lam NN, Docherty PD, Murray R. Practical identifiability of parametrised models: a review of benefits and limitations of various approaches. Mathematics and Computers in Simulation 199:202-216 (2022). https://doi.org/10.1016/j.matcom.2022.03.020
8. Franceschini G, Macchietto S. Model-based design of experiments for parameter precision: state of the art. Chemical Engineering Science 63:4846-4872 (2008). https://doi.org/10.1016/j.ces.2007.11.034
9. Franceschini G, Macchietto S. Novel anticorrelation criteria for design of experiments: algorithm and application. AIChE Journal 54:3221-3238 (2008). https://doi.org/10.1002/aic.11607
10. Geremia M, Macchietto S, Bezzo F. A review on model-based design of experiments for parameter precision – open challenges, trends and future perspectives. Chemical Engineering Science 319:122347 (2026). https://doi.org/10.1016/j.ces.2025.122347
11. Maheshwari V, Rangaiah GP, Samavedham L. Multiobjective framework for model-based design of experiments to improve parameter precision and minimize parameter correlation. Ind. Eng. Chem. Res. 52:8289-8304 (2013). https://doi.org/10.1021/ie400133m
12. Galvanin F, Marchesini R, Barolo M, Bezzo F, Fidaleo M. Optimal design of experiments for parameter identification in electrodialysis models. Chemical Engineering Research and Design 105:107-119 (2016). https://doi.org/10.1016/j.cherd.2015.10.048
13. Nordstrand J, Dutta J. Flexible modeling and control of capacitive-deionization processes through a linear-state-space dynamic langmuir model. npj Clean Water 4: (2021). https://doi.org/10.1038/s41545-020-00094-y
14. Nordstrand J, Laxman K, Myint MTZ, Dutta J. An easy-to-use tool for modeling the dynamics of capacitive deionization. J. Phys. Chem. A 123:6628-6634 (2019). https://doi.org/10.1021/acs.jpca.9b05503

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Process Intensification for LNG Purification: Modeling CO_2 Separation in a Rotating Packed Bed

Alexander A. Zerwas[ab]*, Bruna L. V. Maia[a], Wilson Santos Neto[a], Radin Suhaib Salihuddin[c], Amiza Bt Surmi[c], Fadhli Hadana Rahman[c], Jean F. Leal Silva[a], Dirceu Noriler[a]

[a] University of Campinas, Department of Process Engineering of the Chemical Engineering Faculty, Campinas, Brazil
[b] University of São Paulo, Department of Chemical Engineering, São Paulo, Brazil
[c] PETRONAS Research Sdn Bhd, Bangi, Malaysia
* Corresponding Author: arzerwas@unicamp.br/alexander.zerwas@usp.br

ABSTRACT

Liquefied Natural Gas (LNG) plays a strategic role in the global energy transition, as it represents a less carbon-intensive alternative to coal. Separation of CO_2 from raw natural gas is a critical step for meeting LNG specifications and enabling Enhanced Oil Recovery (EOR) in offshore fields. However, high CO_2 concentrations and formation of a CO_2 ethane azeotrope increase the process complexity, often requiring extractive distillation with heavier hydrocarbons. Severe limitations are faced in offshore environments due to their weight, volume and high energy consumption. Due to that, Process Intensification (PI) seeks to enhance heat and mass transfer efficiency, potentially reducing equipment volume and weight. Rotating Packed Beds (RPB) have demonstrated significant potential for intensifying LNG purification by using centrifugal forces to drive liquid through a porous medium in contact with a gas stream. Experimental measurements of total pressure drop, and local liquid holdup are feasible in pilot-scale units, although typically requiring specialized equipment and instrumentation. Consequently, modeling approaches based on mass, momentum, and energy balances have been widely proposed in the literature to estimate the hydrodynamic behavior of RPB systems. While Computational Fluid Dynamics (CFD) models can provide detailed insights into equipment behavior, they are often computationally expensive for parametric studies and design optimization. In this context, this work proposes a one-dimensional model in cylindrical coordinates to analyze fluid flow in an RPB for natural gas/CO_2 separation process. The governing equations are formulated and validated against experimental data for pressure drop, interfacial velocity, and liquid holdup along the RPB radius.

Keywords: Fluid Dynamics, Modelling and Simulations, Natural Gas, Process Intensification, Distillation

INTRODUCTION

Liquefied Natural Gas (LNG) plays a strategic role in the global energy transition, as it is a less carbon-intensive alternative to coal. The growing need to reduce greenhouse gas emissions, coupled with the search for greater efficiency in oil exploration, has driven the application of Enhanced Oil Recovery (EOR) techniques in offshore fields. In these projects, the separation of CO_2 from raw natural gas is a critical step, both to meet LNG quality specifications and to enable the reinjection of CO_2 as a miscible displacement fluid.

However, many natural gas fields have high levels of carbon dioxide (CO_2), which makes commercial exploration unfeasible. One problem related to removal of CO_2 from raw natural gas is the significant hydrocarbon losses due to the formation of an azeotrope with ethane, which reduces overall efficiency of the process and the economic viability of LNG production. One way of reducing the loss of light carbon components due to the azeotrope formed may be to use a third component and subsequent fractionation using a low temperature distillation process, such as the use of compounds with 4 or more carbons, as reported in [1].

However, conventional separation methods, such as traditional cryogenic distillation columns, have significant limitations in offshore environments: high energy consumption, especially in mixtures with a high CO_2 content;

weight and volume of the equipment, restrictive factors on offshore platforms; operational complexity due to variations in load and gas composition.

One way of reducing the equipment volume is to intensify heat and mass transfer processes, with the aim of improving the efficiency of the mixing and transfer processes [2], which can reduce required volumes to a scale of 10 to 100 times and are interesting in environments with limited space.

For the case of intensifying the separation of CO_2 from the hydrocarbon stream, the application of a rotating packed bed (RPB) has demonstrated potential in the literature. In this method, a liquid stream is injected into the inner surface of the packing, while the packing is rotated, forcing the liquid to travel to the outer surface. A gas stream can be injected either at the outer casing (counter-current) or at the same location as the liquid inlet (co-current).

To understand and estimate separation efficiencies, experimental data or/and simulation and modeling of the transport phenomena inside the system have been carried out. Measurement of total pressure drop [3–5] and local liquid holdup [6, 7] are possible with pilot plant systems, although requiring specialized equipment. Due to that, modelling of the system based on mass, momentum and energy balances are proposed in the literature to estimate the system's behavior. For that, simplified models [8–12] can be solved to understand locally the process dynamics, which can include pressures, flow velocities and compositions across the packing. Moreover, Computational Fluid Dynamics (CFD) models can also be used to address different packing geometries [13–15], although often requiring more computational resources [16–18], but providing an insight in liquid distribution inside the packing, since it is difficult to take internal measurements during operation.

In this context, distillation using RPBs has emerged as a promising technology to overcome the economic barriers to LNG production in offshore wells and to understand the system behavior, a simpler 1D model in cylindrical coordinates are proposed to analyze the fluid flow, heat, and mass transfer in an RPB for the LNG/CO_2 separation process. Based on that, governing equations are described for modeling the process, and it is compared with available experimental data for pressure drop, interfacial velocity, and liquid holdup.

MATERIALS AND METHODS

Rotating packed bed

Rotating packed beds can be operated in co-current or counter-current, with typical rotations between 500 and 2500 RPM, with the packed bed being made of foams, spheres, wire meshes or other geometries and rotated vertically or horizontally. In the counter-current approach, liquid is inserted at the inner radius from the packing material and is driven to the outer radius due to the created gravitational field and gas flows in opposite direction. A schematic representation of the system is given in Figure 1. By increasing the liquid velocity, as compared to stationary packed beds, turbulence is intensified, which improves mass transfer coefficients inside the liquid. This effect is made possible by using packing elements that can simultaneously combine high surface area, promote low pressure drop and efficient liquid distribution [19].

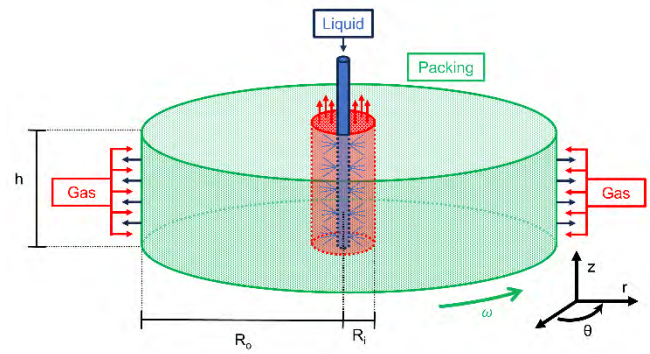

Figure 1. Schematic representation of a horizontal Rotating Packed Bed (RPB) system operating in a counter-current mode.

Hydrodynamic model

A counter-current gas-liquid RPB operating horizontally was assumed. For modelling the two-phase (gas, liquid) behavior, mass and r-component momentum balances are written in cylindrical coordinates, respectively, in Equation (1) and (2). Equation (3) was added to account for the presence of more phases, with all k-th volume fractions being equal to one. Gravitational force in the z-axis were disregarded, following similar approaches from the literature [20, 21].

$$\frac{1}{\alpha_k}\frac{\partial \alpha_k}{\partial r} + \frac{1}{u_{r,k}}\frac{\partial u_{r,k}}{\partial r} + \frac{1}{P}\frac{\partial P}{\partial r} = -\frac{1}{r} \tag{1}$$

$$\alpha_k \rho_k u_{r,k}\frac{\partial u_{r,k}}{\partial r} + \alpha_k \frac{\partial P}{\partial r} - 2\mu_k\left(\frac{1}{r}\frac{\partial u_{r,k}}{\partial r}\right) = \frac{\alpha_k \rho_k u_{\theta,k}^2}{r} + M_{r,k} \tag{2}$$

$$\sum_{k=0}^{n} \alpha_k = 1 \tag{3}$$

α_k represents the volume fraction with subindices k-th, representing the two phases (k – gas, liquid), P is the local pressure [Pa], which followed the ideal gas law, ρ_k is the density [kg/m^3], μ_k is the viscosity [Pa.s] $u_{r,k}$ is the area averaged radial velocity and $u_{\theta,k}$ is the tangential velocity of the k-th phase [m/s] and $M_{r,k}$ represents the momentum exchange between the phases. Assuming a solid body rotation, tangential velocities can be written in terms of radius and rotation velocity (ω), as $u_{\theta,k} = \omega r$.

In the modelling, simplifying hypotheses were assumed, such as steady-state operation, one-dimensional flow (velocity varies only in r component), no mass transfer occurs (only the hydrodynamic model is being developed) and excluded secondary velocity derivatives. The model was solved considering initial value problems, where conditions were set at the inner radius (R_i) and the equations were numerically integrated up to the outer radius of the packing material (R_o). Integrations were solved using the implicit Radau method, which is recommended for stiff-problems, available in Scipy library.

Interaction forces

During the flow of the gas and liquid phases through the voids of the solid packing material, the liquid often flows above the solid surface, partially wetting the surface and applying a force between at the solid boundary (f_{LS}). Gas flows in the center of the void section, forming an interface between liquid and solid. Due to velocities at the interface being the same, momentum is transferred at the interface due to the drag force between gas and liquid (f_{GL}). Also, gas momentum can be transferred directly to the solid (f_{GS}) through the liquid film, due to the tortuosity inside the packing material. As a result, momentum transfer to the liquid ($M_{r,L}$) and gas ($M_{r,G}$) are written, respectively, by Equations (4) and (5) [20].

$$M_{r,L} = f_{GL} - f_{LS} \tag{4}$$

$$M_{r,G} = -f_{GL} - f_{GS} \tag{5}$$

To calculate the interaction forces f_{LS}, f_{GS} and f_{GI}, different models are proposed in the literature. For studying the hydrodynamic behavior of a descending cocurrent gas-liquid flow through a trickle bed reactor, composed of packed spheres, Attou et al. (1999) [20] have used Equations (6), (7), (8) and (9), assuming that all solid surfaces were covered with liquid. For the model it was considered that the Ergun parameters are: $E_{\mu,LS} = E_{\mu,GS} = E_{\mu,GL} = 180$ and $E_{\rho,LS} = E_{\rho,GS} = E_{\rho,GL} = 1.8$

$$f_{LS} = E_{\mu,LS}\frac{(1-\epsilon)^2}{(1-\alpha_g)^3\epsilon^3 d_p^2}\mu_l w_l + E_{\rho,LS}\frac{1-\epsilon}{(1-\alpha_g)^3\epsilon^3 d_p}\rho_l w_l|w_l| \tag{6}$$

$$f_{GS} = \alpha_g\left(E_{\mu,GS}\frac{(1-\alpha_g\epsilon)^2}{\alpha_g{}^3\epsilon^3 d_p^2}\left(\frac{1-\epsilon}{1-\alpha_g\epsilon}\right)^{\frac{2}{3}}\mu_g w_g + E_{\rho,GS}\frac{(1-\alpha_g\epsilon)}{\alpha_g{}^3\epsilon^3 d_p}\left(\frac{1-\epsilon}{1-\alpha_g\epsilon}\right)^{\frac{1}{3}}\rho_g w_g|w_g|\right) \tag{7}$$

$$f_{GL} = \alpha_g\left(E_{\mu,GL}\frac{(1-\alpha_g\epsilon)^2}{\alpha_g{}^3\epsilon^3 d_p^2}\left(\frac{1-\epsilon}{1-\alpha_g\epsilon}\right)^{\frac{2}{3}}\mu_g w_r + E_{\rho,GL}\frac{(1-\alpha_g\epsilon)}{\alpha_g{}^3\epsilon^3 d_p}\left(\frac{1-\epsilon}{1-\alpha_g\epsilon}\right)^{\frac{1}{3}}\rho_g w_r|w_r|\right) \tag{8}$$

$$w_r = w_g - \left(\frac{\alpha_g}{1-\alpha_g}\right)w_l \tag{9}$$

Where ϵ is the packed bed porosity, d_p the diameter of the packing material [m], w is the superficial velocity [m/s], defined as $w_k = \epsilon u_k$. To account for a partial surface wettability (f_e), Lappalainen et al. (2008) [22] modified to include this effect in the interaction forces, as given by Equations (10) and (11).

$$M_{r,L} = f_e(f_{GL} - f_{LS}) \tag{10}$$

$$M_{r,G} = -f_e f_{GL} - (1-\mathrm{f_e})f_{GS} \tag{11}$$

In addition to that, Ergun parameters were changed to account for the different tortuosity in the packing. Ergun parameters are then modified according to Equations (12) and (13).

$$E_{\mu,i} = 72{T_i}^2 \tag{12}$$

$$E_{\rho,i} = 0.4554{T_i}^3 \tag{13}$$

With specific tortuosity (T_i) for the gas and liquid phase, being given respectively, by Equations (14) and (15), with $T_0 = 1.58$ being obtained for the tortuosity in an empty bed and the parameters $A = 3.592$ and $B = 1.14$ obtained by optimization.

$$T_G = \frac{T_o+1}{2} + \alpha_g\left(\frac{T_0+1}{2} - 1\right) \tag{14}$$

$$T_L = T_o A^{{\alpha_g}^B} \tag{15}$$

Lu et al. (2018) [23] modified the flow resistance model of Kolodziej et al. (2009) [24] to account for a packing bed made of wire meshes, resulting in interaction forces being given by Equations (16), (17) and (18).

$$f_{LS} = \alpha_l\epsilon\left(4(f_{app}+f_t)\frac{\rho_l|w_l|\epsilon\tau^3}{2d_w(\alpha_l\epsilon)^3\cos^3\theta}\right)w_l \tag{16}$$

$$f_{GS} = \alpha_g\epsilon\left(4(f_{app}+f_t)\frac{\rho_g|w_g|(1-\alpha_g\epsilon)\tau^3}{2d_w(\alpha_g\epsilon)^3\cos^3\theta}\right)w_g \tag{17}$$

$$f_{GL} = \alpha_g\epsilon\left(4(f_{app}+f_t)\frac{\rho_g|w_g-w_l|(1-\alpha_g\epsilon)\tau^3}{2d_w'(\alpha_g\epsilon)^3\cos^3\theta}\right)(w_g-w_l) \tag{18}$$

In which f_{app} is the corrected Fanning friction factor for non-developed laminar flows and f_t is the Fanning friction factor for turbulent flows, τ is the bed tortuosity factor, θ is the angle between the bed axis and flow direction and d_w is the diameter of the packed wires [m], which are given for each phase and detailed elsewhere [23].

Wetting efficiency

For the surface efficiency coverage (f_e), the model developed by Lu et al. (2018) [23] used the correlation provided by Onda et al. (1968) [25], as detailed in Equation (19).

$$f_e = 1 - e^{\left(-1.45\left(\frac{\sigma_c}{\sigma}\right)^{0.75} Re_l^{0.1} We_l^{0.2} Fr_l^{-0.05}\right)} \tag{19}$$

In which $Re_l = \rho_l u_l / a_s \mu_l$ is the Reynolds number of the liquid, $We_l = \rho_l u_l^2 / a_s \sigma$ is the Weber number, $Fr_l = a_s u_l^2 / g_c$ is the Froude number and centrifugal acceleration is given by $g_c = r\omega^2$. $\sigma = 0.072$ is the surface tension of the liquid, which was assumed as water [N/m] and $\sigma_c = 0.06$ is the critical surface tension in glass media. v_l is the liquid velocity [m/s] and a_s is the specific surface area of the dry packing media [1/m].

Lappalainen et al. (2008)[22] fitted a correlation from literature data for $Re_l, Ga_g = \rho_g^2 g d_p^3 \epsilon^3 / (\mu_g (1-\epsilon)^2)$ being the Galileo number, $Fr_g = u_g / \sqrt{g d_p}$ being the Froude number and $E\ddot{o} = \rho_l g d_p^2 \epsilon^2 / (\sigma (1-\epsilon)^2)$ being the modified Eötvos number, resulting in Equation (20).

$$f_e = 0.335 Re_l^{0.185} E\ddot{o}^{-0.188} Ga_g^{0.027} \left(1 + Fr_g\right)^{-0.014} \tag{20}$$

RESULTS AND DISCUSSION

For the model verification, experimental and simulation results obtained in the literature were compared with the current implementations using the three interaction forces of Attou et al. (1999) [20], Lappalainen et al. (2008) [22] and Lu et al. (2018) [23], with necessary parameters extracted from the sources. Attou et al. (1999) [20] and Lappalainen et al. (2008) [22] were developed for modelling the hydrodynamic behavior in a trickle bed reactor, composed of packed solid spheres, and Lu et al. (2018) [23] focused in modelling the flow inside a porous foam media. Comparisons were made in terms of pressure drop through the rotating media, liquid holdup, gas and liquid velocities. Experimental data from Sandilya et al. (2001), with parameters highlighted in Table 1, were used in the developed model, while varying two gas flow rates (2000 and 6000 cm³/s) and two rotations (950 and 1410). Another setup case from Lu et al. (2018) [23] was also tested using the different interaction forces.

Pressure drop

By comparing the pressure results obtained by the proposed model for gas flow rates in the range of 2000 cm³/s and 6000 cm³/s, as observed respectively in Figures 2 and 3. It is possible to note similar estimations for the pressure by using the interaction model of Lappalainen et al. (2008) [22] and Lu et al. (2018) [23], at the gas flow rate of 2000 cm³/s, which are somewhat close to the experimental data of Sandilya et al. (2001) [26], while Attou et al. (1999) [20] overpredicts pressure drop, with a relative difference of more than 72 % (950 rpm), which could be explained by considering that the whole solid surface is covered by liquid, resulting in a slightly higher pressure drop. When a higher gas flow rates of 6000 cm³/s are used, the three interaction force models predict higher values for pressure losses and the force models proposed in Lu et al. (2018) [23] and Lappalainen et al. (2008) [22] results in a value closer to the experimental data, with a maximum relative difference of 33 % (950 rpm). However, it is observed that higher rotation speeds increase pressure drop.

Table 1: Specification of the packing material used in the experimental setup and flow rates.

Var.	Description	[26]	[27]
R_i	Inner radius (mm)	30	35
R_o	Outer radius (mm)	155	160
h	Bed height (mm)	22	10
d_p	Diameter pore (mm)	0.25	7.6
a_s	Specific media area (1/m)	2196	786
ϵ	Porosity (-)	0.91	0.953
Q_g	Gas flow rate (cm³/s)	2000; 6 000	2.0
Q_l	Liquid flow rate (cm³/s)	14.3	17.5
ω	Rotation speed (rpm)	950; 1410	600; 917; 1200

By comparing the pressure results obtained by the proposed model for lower gas flow rates (2.0 cm³/s), as observed in Figure 4, it is possible to note that the three different interaction force models give a similar behavior for the different imposed rotations and result in differences lower than 0.2% with the correlation proposed in Sandilya et al. (2001) [26]. Moreover, it is observed that higher pressure losses as higher rotation speeds are employed and results are similar for the estimation from Equation (21), which is the correlation proposed by Sandilya et al. (2001) [26].

$$\begin{aligned} \Delta P_{RPB} = & \frac{150(1-\epsilon)^2 \mu}{{d_p}^2 \epsilon^3} \left(\frac{Q_g}{2\pi h}\right) \ln\left(\frac{R_o}{R_i}\right) \\ & + \frac{1.75(1-\epsilon)\rho_g}{d_p \epsilon^3} \left(\frac{Q_g}{2\pi h}\right)^2 \left(\frac{1}{R_i} - \frac{1}{R_o}\right) \\ & + \frac{1}{2} \rho_g \omega^2 (R_i^2 - R_o^2) \end{aligned} \tag{21}$$

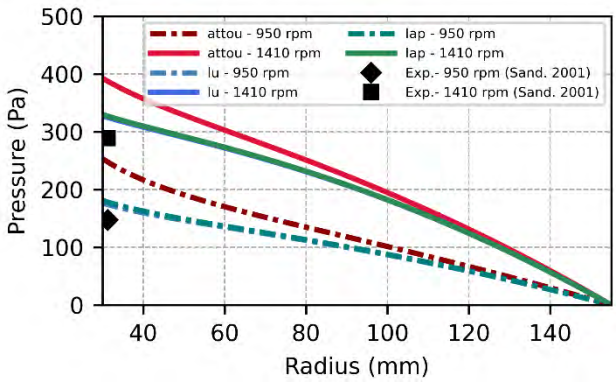

Figure 2. Relative pressures across the packing media, with atmospheric pressure condition being placed at the outer radius, for a gas flow rate of 2000 cm^3/s and liquid flow rate of 14.3 cm^3/s, and pressure obtained with the experimental data from Sandilya et al. (2001) [26].

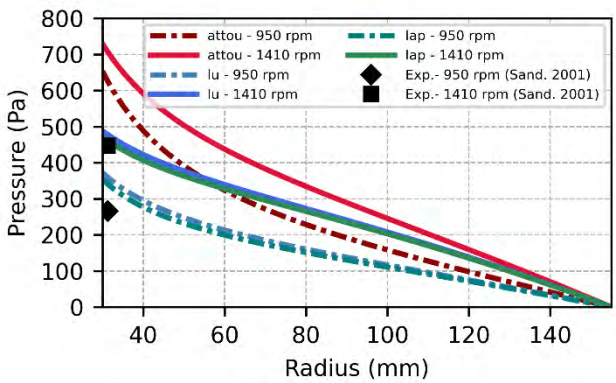

Figure 3. Relative pressures across the packing media, with atmospheric pressure condition being placed at the outer radius, for a gas flow rate of 6000 cm^3/s and liquid flow rate of 14.3 cm^3/s, and pressure obtained with the experimental data from Sandilya et al. (2001) [26].

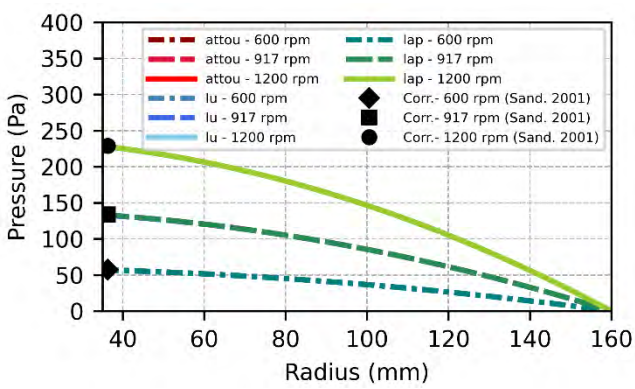

Figure 4. Relative pressures across the packing media, with atmospheric pressure condition being placed at the outer radius, for a gas flow rate of 2.0 cm^3/s and liquid flow rate of 17.5 cm^3/s, and pressure obtained with the correlation from Sandilya et al. (2001) [26].

Liquid holdup

Using a packed bed with inner radius of 35 mm and outer radius of 160 mm, with specific surface area of 786 1/m and porosity of 0.953, with rotations varying from 25 to 170 rad/s, Burns et al. (2000)[28] evaluated the liquid retained (ϵ_l) in the porous foam media and correlated that based on the rotation and the superficial liquid velocity, as observed in Equation (22).

$$\epsilon_l = 0.034\left(\frac{r\omega^2}{100}\right)^{-0,38}\left(\frac{w_l}{0.01}\right)^{0.62} \tag{22}$$

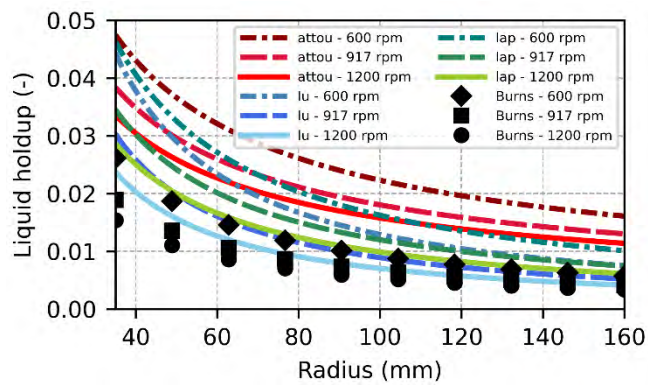

Figure 5. Comparison of liquid holdup obtained for three force interaction models across the packing media for the packed bed configuration used by Lu et al. (2018) [23], and experimental correlation obtained by Burns et al. (2000)[28].

By using similar packed bed configuration and process parameters, the three force models were compared with the correlation in Equation (22), as shown in Figure 5. It is possible to observe that for this case, that the interaction proposed by Attou et al. (1999) [20] resulted in higher liquid holdup, with a relative difference of 186% (1200 rpm), Lu et al. (2018) [23] predicted lower liquid holdups, with a relative difference of 81% (1200 rpm), and the model proposed by Lappalainen et al. (2008) [22] fell in the middle between the two models, with a relative difference of 33% (1200 rpm).

However, when higher gas flow rates are employed, such as the ones used in the experiment of Sandilya et al. (2001) with gas flow rate of 6000 cm^3/s, it is possible to observe lower liquid retentions obtained by using the model of Lu et al. (2018) [23], with close behavior than the other two interaction models, with a relative difference of 125% (950 rpm), as observed in Figure 6. Moreover, it is noticed that for all cases, higher rotations of the RPB ended up decreasing the liquid holdup and as the liquid traveled inside the packed media to the outer radius, lower local liquid retention is also observed, possibly due to the increase in surface area.

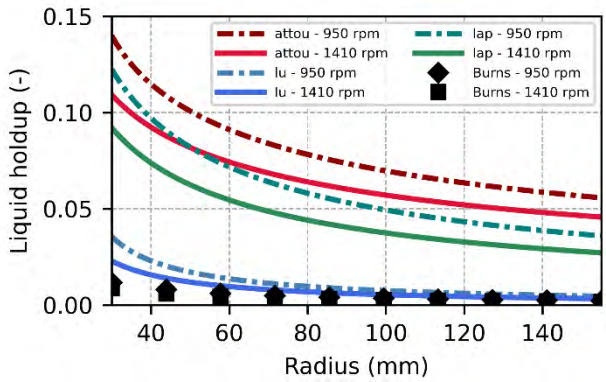

Figure 6. Comparison of liquid holdup obtained for three force interaction models across the packing media for the packed bed configuration used by Sandilya et al. (2001) [26] with gas flow rate of 6000 cm^3/s and liquid flow rate of 14.3 cm^3/s and experimental correlation obtained by Burns et al. (2000)[28].

Gas velocity

Radial gas velocities inside the packing bed were extracted for the configuration used by Sandilya et al. (2001) [26] with gas flow rate of 6000 cm^3/s with the three interaction force models, as shown in Figure 7. It is possible to observe an increase in gas velocity (u_g) as it travels to the center of the packed bed, an independent behavior of the rotation, as reported also in the literature.

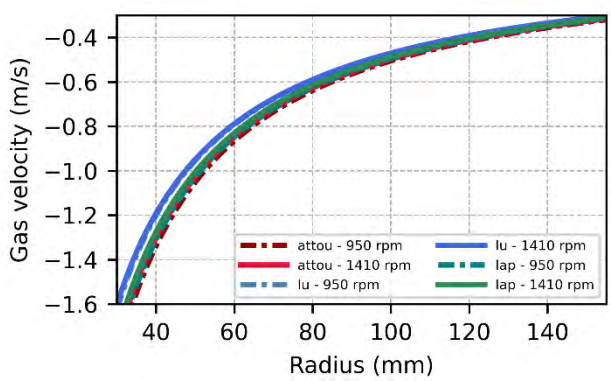

Figure 7. Comparison of radial gas velocities obtained for three force interaction models across the packing media for the packed bed configuration used by Sandilya et al. (2001) [26] with gas flow rate of 6000 cm^3/s and liquid flow rate of 14.3 cm^3/s.

Liquid velocity

Radial liquid velocities inside the packing bed were extracted for the three tested force models using the packed bed configuration of Lu et al. (2018) [23] with results shown in Figure 8. It is observed that the force model proposed by Attou et al. (1999) [20] predicts lower interstitial liquid velocities and tends to reduce the velocity as liquid travels to the outer radius of the RPB. For Lu et al. (2018) [23], higher velocities are observed, but with an increase in velocity as liquid travels the outward of the packing material. For Lappalainen et al. (2008) [22], lower and constant velocity were observed.

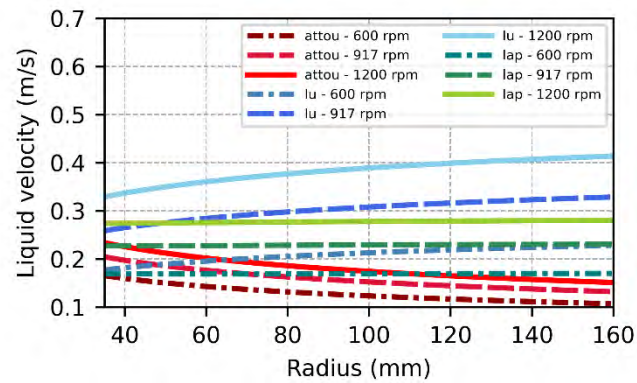

Figure 8. Comparison of radial liquid velocities obtained for three force interaction models across the packing media for the packed bed configuration used by Lu et al. (2018) [23].

When higher gas flow rates (6000 cm^3/s) are used, with configuration used by Sandilya et al. (2001) with the three interaction force models, as shown in Figure 9, liquid velocities for Attou et al. (1999) [20] and Lappalainen et al. (2008) [22] have similar values, whereas for Lu et al. (2018) [23], higher liquid velocities are observed, which still increases as liquid is driven outward from the packed bed.

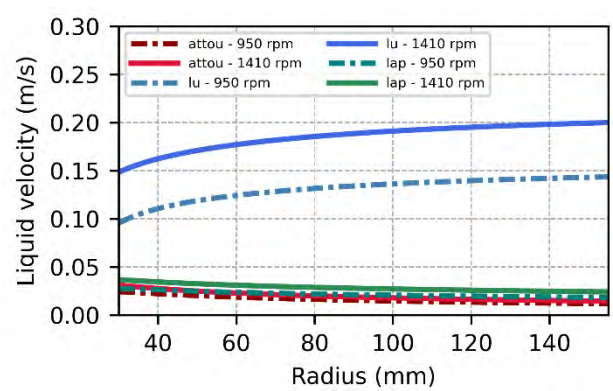

Figure 9. Comparison of radial liquid velocities obtained for three force interaction models across the packing media for the packed bed configuration used by Sandilya et al. (2001) [26] with gas flow rate of 6000 cm^3/s and liquid flow rate of 14.3 cm^3/s.

CONCLUSIONS AND OUTLOOKS

Based on the mass and momentum balance presented in this study, pressure, velocity and volume fractions were obtained for a gas-liquid RPB working at counter-current condition. To model the interaction between

the gas-liquid system with the solid packing material, three interaction force models were analyzed for different operating conditions and compared to available experimental data.

For the pressure losses across the bed material and operating at lower gas flow rates, similar predicted values were obtained for the different force models. However, for higher gas flow rates, differences were observed, with Lu et al. (2018) [23] and Lappalainen et al. (2008) [22] models obtaining closer values to the correlation.

For the liquid holdup estimation, it is observed that Lu et al. (2018) [23] model predicts closer values to the Burns et al. (2000)[28] correlation. When using the three force models, slightly differences were observed for gas velocities. Moreover, rotation did not have an effect in altering the gas velocities across the packed media, as reported also in the literature.

For liquid velocities, higher values were observed for the model proposed by Lu et al. (2018) [23], with lower values observed for the Attou et al. (1999) [20] model. When gas flow rates increase, similar behavior occurs for the liquid velocity.

Although the force model proposed by Attou et al. (1999) [20] and Lappalainen et al. (2008) [22] were focused in the flow behavior inside a sphere packed reactor, it could be observed good agreement with the experimental data for rotating packed beds with higher porosity medias, such as the one developed by Lu et al. (2018) [23], which used force models focused in flow around wire meshes.

Based on the presented study, it is possible to see that the hydrodynamic model for the RPB flow was successfully implemented, and further comparisons with experimental data and CFD simulations results are expected to occur. Moreover, it is possible to further develop the model to account for mass and heat transfer in the distillation of LNG/CO_2 streams, to obtain characteristic sizes and operational conditions for the equipment.

ACKNOWLEDGEMENTS

We gratefully acknowledge PETRONAS Petróleo Brasil Ltda. (PPBL) for their funding and support, as well as ANP (Brazil's National Oil, Natural Gas, and Biofuels Agency) for assistance through the R&D levy regulation (Project #25186-8). We also thank EQS – Sustainable Energy and Chemistry Laboratory, PQGe – Chemical Processes and Business Management Laboratory, the Center for Energy and Petroleum Studies (CEPETRO), and the School of Chemical Engineering (FEQ) at the University of Campinas (UNICAMP), Brazil. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001.

AUTHOR IDENTIFIERS

Zerwas AA: 0000-0002-0083-6196
Maia BLV: 0009-0006-9267-2905
Santos Neto W: 0009-0005-7915-0152
Salihuddin RS: 0009-0009-6070-9429
Surmi AB: 0009-0005-8071-3602
Rahman FH: 0009-0008-7801-2198
Leal Silva JF: 0000-0002-9915-5859
Noriler D: 0000-0001-5281-5226

REFERENCES

1. ZareNezhad B, Hosseinpour N. An extractive distillation technique for producing CO2 enriched injection gas in enhanced oil recovery (EOR) fields. Energy Conversion and Management 50:1491-1496 (2009). https://doi.org/10.1016/j.enconman.2009.02.016
2. Paul EL, Atiemo?Obeng VA, Kresta SM. Handbook of industrial mixing. Wiley (2004). https://doi.org/10.1002/0471451452
3. Neumann K, Hunold S, Skiborowski M, Górak A. Dry pressure drop in rotating packed beds—systematic experimental studies. Ind. Eng. Chem. Res. 56:12395-12405 (2017). https://doi.org/10.1021/acs.iecr.7b03203
4. Kelleher T, Fair JR. Distillation studies in a high-gravity contactor. Ind. Eng. Chem. Res. 35:4646-4655 (1996). https://doi.org/10.1021/ie950662a
5. Garba U, Rouzineau D, Meyer M. EXPERIMENTAL STUDY OF THE PRESSURE DROP AND FLOODING IN a ROTATING PACKED BED REACTOR. Malay. J. Catal. 7:42-48 (2023). https://doi.org/10.11113/mjcat.v7n2.174
6. Burns JR, Jamil JN, Ramshaw C. Process intensification: operating characteristics of rotating packed beds — determination of liquid hold-up for a high-voidage structured packing. Chemical Engineering Science 55:2401-2415 (2000). https://doi.org/10.1016/s0009-2509(99)00520-5
7. Yang Y, Xiang Y, Chu G, Zou H, Luo Y, Arowo M, Chen JF. A noninvasive x-ray technique for determination of liquid holdup in a rotating packed bed. Chemical Engineering Science 138:244-255 (2015). https://doi.org/10.1016/j.ces.2015.07.044
8. Joel AS, Wang M, Ramshaw C, Oko E. Modelling, simulation and analysis of intensified regenerator for solvent based carbon capture using rotating packed bed technology. Applied Energy 203:11-25 (2017). https://doi.org/10.1016/j.apenergy.2017.05.157
9. Guo F, Zheng C, Guo K, Feng Y, Gardner NC. Hydrodynamics and mass transfer in cross-flow rotating packed bed. Chemical Engineering Science 52:3853-3859 (1997). https://doi.org/10.1016/s0009-2509(97)00229-7

10. Lin CC, Chen YS, Liu HS. Prediction of liquid holdup in countercurrent-flow rotating packed bed. Chemical Engineering Research and Design 78:397-403 (2000). https://doi.org/10.1205/026387600527293
11. Iliuta I, Petre CF, Larachi F. Hydrodynamic continuum model for two-phase flow structured-packing-containing columns. Chemical Engineering Science 59:879-888 (2004). https://doi.org/10.1016/j.ces.2003.11.020
12. Chandra A, Goswami PS, Rao DP. Characteristics of flow in a rotating packed bed (HIGEE) with split packing. Ind. Eng. Chem. Res. 44:4051-4060 (2005). https://doi.org/10.1021/ie048815u
13. Qammar H, G?adyszewski K, Górak A, Skiborowski M. Towards the development of advanced packing design for distillation in rotating packed beds. Chemie Ingenieur Technik 91:1663-1673 (2019). https://doi.org/10.1002/cite.201900053
14. Wen ZN, Wu W, Luo Y, Zhang LL, Sun BC, Chu GW. Novel wire mesh packing with controllable cross-sectional area in a rotating packed bed: mass transfer studies. Ind. Eng. Chem. Res. 59:16043-16051 (2020). https://doi.org/10.1021/acs.iecr.0c01886
15. Zawadzki D, Majdzik M, Pi?tkowski M, Jaskulski M, Blatkiewicz M. Prototyping of rotating packed bed structural internals. Chemical Engineering and Processing - Process Intensification 196:109678 (2024). https://doi.org/10.1016/j.cep.2024.109678
16. Martínez EL, Jaimes R, Gomez JL, Filho RM (2012): CFD Simulation of Three-Dimensional Multiphase Flow in a Rotating Packed Bed. In: Computer Aided Chemical Engineering, Bogle IDL and Fairweather M, eds. Elsevier, pp. 1158–1162
17. Llerena-Chavez H, Larachi F. Analysis of flow in rotating packed beds via CFD simulations—dry pressure drop and gas flow maldistribution. Chemical Engineering Science 64:2113-2126 (2009). https://doi.org/10.1016/j.ces.2009.01.019
18. Lu X, Xie P, Ingham DB, Ma L, Pourkashanian M. Modelling of CO2 absorption in a rotating packed bed using an eulerian porous media approach. Chemical Engineering Science 199:302-318 (2019). https://doi.org/10.1016/j.ces.2019.01.029
19. Lukin I, G?adyszewski K, Skiborowski M, Górak A, Schembecker G. Aroma absorption in a rotating packed bed with a tailor-made archimedean spiral packing. Chemical Engineering Science 231:116334 (2021). https://doi.org/10.1016/j.ces.2020.116334
20. Attou A, Boyer C, Ferschneider G. Modelling of the hydrodynamics of the cocurrent gas–liquid trickle flow through a trickle-bed reactor. Chemical Engineering Science 54:785-802 (1999). https://doi.org/10.1016/s0009-2509(98)00285-1
21. Attou A, Ferschneider G. A two-fluid hydrodynamic model for the transition between trickle and pulse flow in a cocurrent gas–liquid packed-bed reactor. Chemical Engineering Science 55:491-511 (2000). https://doi.org/10.1016/s0009-2509(99)00344-9
22. Lappalainen K, Alopaeus V, Manninen M, Aittamaa J. Improved hydrodynamic model for wetting efficiency, pressure drop, and liquid holdup in trickle-bed reactors. Ind. Eng. Chem. Res. 47:8436-8444 (2008). https://doi.org/10.1021/ie8003754
23. Lu X, Xie P, Ingham DB, Ma L, Pourkashanian M. A porous media model for CFD simulations of gas-liquid two-phase flow in rotating packed beds. Chemical Engineering Science 189:123-134 (2018). https://doi.org/10.1016/j.ces.2018.04.074
24. Ko?odziej A, ?ojewska J. Experimental and modelling study on flow resistance of wire gauzes. Chemical Engineering and Processing: Process Intensification 48:816-822 (2009). https://doi.org/10.1016/j.cep.2008.10.009
25. Onda K, Takeuchi H, Okumoto Y. Mass Transfer Coefficients between Gas and Liquid Phases in Packed Columns. J Chem Eng Jpn 1:56-62 (1968). https://doi.org/10.1252/jcej.1.56
26. Sandilya P, Rao DP, Sharma A, Biswas G. Gas-phase mass transfer in a centrifugal contactor. Ind. Eng. Chem. Res. 40:384-392 (2000). https://doi.org/10.1021/ie0000818
27. Lu X, Xie P, Ingham DB, Ma L, Pourkashanian M. A porous media model for CFD simulations of gas-liquid two-phase flow in rotating packed beds. Chemical Engineering Science 189:123-134 (2018). https://doi.org/10.1016/j.ces.2018.04.074
28. Burns JR, Ramshaw C. Process intensification: visual study of liquid maldistribution in rotating packed beds. Chemical Engineering Science 51:1347-1352 (1996). https://doi.org/10.1016/0009-2509(95)00367-3

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Desing and optimization of a multi-objective plant to obtain the best furfural derivates

Melanie Coronel Muñoz[a], Brenda Huerta Rosas[a], Eduardo Sánchez Ramírez[a]*, Juan Gabriel Segovia Hernández[a] and Juan José Quiroz Ramírez[b].

[a] Universidad de Guanajuato, Department of Chemical Engineering, Guanajuato, Guanajuato, México.
[b] SECIHTI-CIATEC A. C. Center for Applied Innovation in Competitive Technologies, León 37545, Guanajuato, México.
* Corresponding Author: eduardo.sanchez@ugto.mx

ABSTRACT

The valorization of lignocellulosic biomass represents a key pathway toward sustainable chemical production, as it enables the development of circular economy products with reduced dependence on fossil resources. Among the platform molecules derived from biomass, furfural stands out as a versatile intermediate that can be transformed into several high-value chemicals, such as furfuryl alcohol, 2-methylfuran, tetrahydrofurfuryl alcohol, furan, tetrahydrofuran, and maleic anhydride. In this work, an integrated biorefinery scheme for producing the main furfural derivatives is proposed and evaluated through process simulation and sustainability metrics. The process was modeled in Aspen Plus V14, using furfural obtained from lignocellulosic biomass (corn stover) as raw material, following hydrogenation and oxidation routes. The process is multi-product, meaning that the main furfural derivatives are produced simultaneously within the same plant. This was achieved by implementing splitters that work at 50% capacity in the division of raw materials. The mass and energy balances from the simulation were used to evaluate environmental, economic, energy, and efficiency performance using a set of indicators. Economic data were obtained from Aspen Economics, yielding a total capital cost of 7.65 MUSD and a total annual cost of 1.93 MUSD/year. The environmental assessment resulted in an EI99 of approximately 2.0×10^{8} mPt/year, mainly influenced by hydrogen consumption and carbon monoxide emissions. From an efficiency perspective, the process shows favorable performance, with an MLI of 0.07 and an MCI of 69%, indicating high material efficiency despite the inherent complexity of multi-stage hydrogenation and oxidation pathways. Overall, the results demonstrate the technical and sustainability potential of an integrated furfural-based biorefinery, while highlighting opportunities for further improvement through process integration and hydrogen recovery.

Keywords: Furfural, Furfuryl alcohol, Biomass, Biorefinery, Performance Indexes

INTRODUCTION

Background

Sustainability, the need to generate circular economy products, furfural as a high-value product from biomass, a biorefinery scheme to produce all components for different applications, furfural as a versatile platform, explaining each of the products to be obtained.

The valorization of biomass as a sustainable alternative to its fossil counterpart, to create sustainable and profitable processes, is of great importance since lignocellulosic biomass is the only source that fixes carbon, is especially abundant, relatively inexpensive, and has great potential to be transformed into high-value bio-based chemicals. Due to this great economic advantage offered by using biomass as a starting point for the creation of a complete process, process design is relevant to identifying the appropriate kinetic and thermodynamic conditions.

The concept of biorefinery has become popular in recent years, as it involves the general concept of a plant that processes raw materials from biomass and converts them into a platform product, from which products with higher market value can subsequently be extracted, but

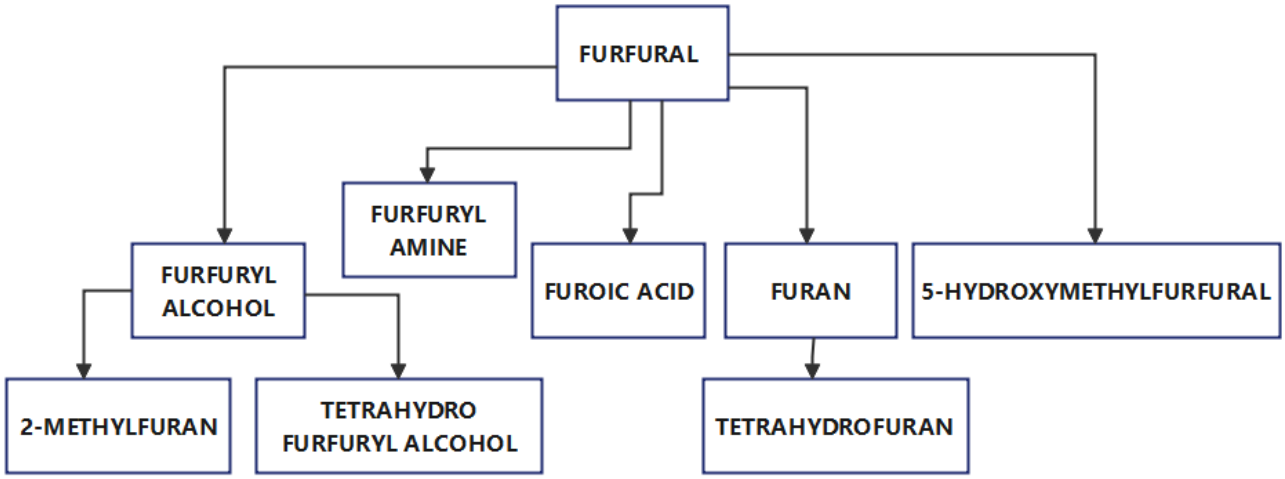

Figure 1. Primary products derived from the direct transformation of furfural [1].

which now have a more sustainable origin as they come from biomass.

This study considered integrated biorefinery processes by combining different processes to produce various products of interest. This was achieved by integrating mass and energy flows for the processing of furfural by-products, which can improve process efficiency in terms of yield, raw materials, and energy consumption, reducing waste generation and promoting economic and environmental performance. Since process integration provides a broader view of the use of products and by-products, through simple processes such as oxidation or hydrogenation, we can obtain a new product derived from the biorefinery. It is important to note that in previous cases, such as that carried out by Kabwe et al. (2025), scenarios are studied where production focuses only on some furfural derivatives. In this study, we carried out a multi-product process, meaning that all the main furfural derivatives are produced simultaneously within the same plant. This was achieved through the implementation of splitters that operate at 50% capacity and manage to divide the product flows that feed another sub-process.

The Chemistry of furfural and its derivatives

The chemistry of furfural and its derivatives is well developed. The main transformation of furfural is through hydrogenation and oxidation, resulting in primary families of derivatives that are more profitable than furfural, as most of them already have an established market, as can be seen in Figure 1. The different compounds obtained are widely used as solvents, transportation fuels, gasoline additives, lubricants, resins, addictive and decolorizing agents, and flavor enhancers for food and beverages.

Furfuryl alcohol

Among the main furfuryl derivatives, furfuryl alcohol is the most established, due to its broad market with applications in the chemical industry, specifically in pharmaceuticals, polymers, and resin manufacturing. It is also an important chemical intermediate for the manufacture of fragrances [1].

2-Methylfuran

In addition to the production of furfuryl alcohol, alternative furfural derivatives of interest that show market potential include 2-methylfuran (MF), which is currently the focus of research because it is an alternative to fossil fuels. Since MF has physicochemical characteristics like gasoline, its implementation in internal combustion engines is a potential field of application [2].

Tetrahydrofurfuryl

Tetrahydrofurfuryl alcohol is a promising source of organic chemicals that are currently produced from fossil fuels [3].

Furan

Furan is commercially produced from furfural by decarbonylation, where the loss of carbon monoxide directly yields furan. From a general chemical perspective, furan is a 5-membered heterocyclic unsaturated ring compound containing oxygen; therefore, it is the basic ring structure found in all kinds of industrially important products [4].

Tetrahydrofuran

Tetrahydrofuran has established itself as an excellent solvent, with high volatility, a low freezing point, and

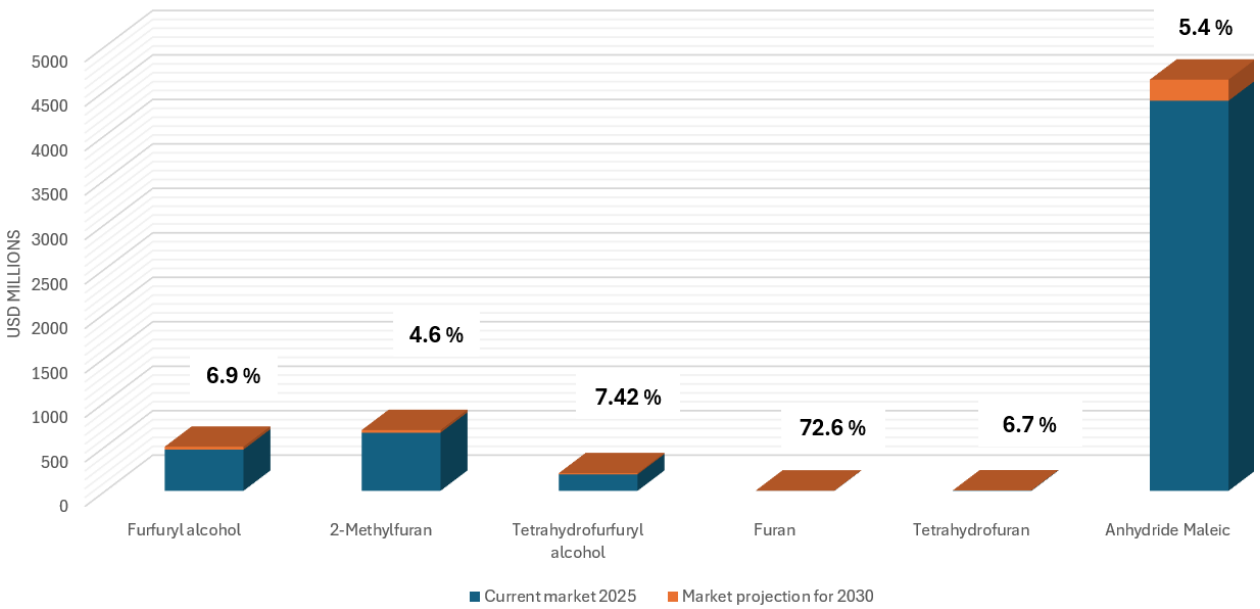

Figure 2. Growth projections for the main furfural derivatives for 2030 [8-13].

the ability to dissolve both polar and non-polar compounds. The applications of tetrahydrofuran are diverse, ranging from the production of polymer coatings and adhesives to pharmaceutical chemistry, where its presence reflects the use of a clean solvent that can be easily separated from the desired products [5]. Currently, THF is produced industrially from petroleum-derived intermediates, involving the hydrogenation of maleic anhydride or the dehydration of 1, 4-butanediol. Bio-based THF can be produced by direct catalytic hydrogenation of furfural or using furan as an intermediate, achieving typical yields of 52–100% [6].

Maleic anhydride

Finally, maleic anhydride is the anhydrous form of maleic acid and is currently produced on an industrial scale from fossil fuels, but there is another route for obtaining it from furan. Its main applications are in the pharmaceutical industry and in the manufacture of polyester resins. The largest percentage of maleic anhydride produced in the United States is used in the manufacture of polyester resins (52%). These applications are diverse, ranging from their use in reinforced plastics for automobiles, boats, and construction panels, among others [7]. However, the market for maleic anhydride resins is undoubtedly growing year after year, hence the importance of finding a new biosustainable production mechanism.

Market for the main furfural derivates

The market for furfural derivatives, as shown in Figure 2, has become even more important than that for furfural itself as a base product, because it is these transformed compounds that truly capture the greatest economic and technological value in industrial chains. While furfural is mainly used as a chemical platform and specialized solvent, its derivatives find direct and massive applications in resins, fuels, additives, polymers, coatings, and high-performance materials. This means that demand is driven more by the end markets for these derivatives than by the consumption of furfural itself, shifting the focus from simple intermediate production to the integration of higher value-added conversion processes. In this sense, furfural increasingly functions as a strategic “building block” within a biorefinery, while the real economic and growth engine of the sector lies in the diversification and expansion of its derivatives in larger-scale, higher-margin industrial markets.

Based on this, several researchers have focused their attention on conducting techno-economic analyses of furfural derivatives to compare production costs with prices established for different scenarios, such as Kabwe et al. (2025), who evaluated different scenarios for the conversion of furfural products, depending on the product of interest and the conditions established.

CASE STUDY

Furfural production

The background to furfural production is the design and optimization of a multi-product biorefinery aimed at obtaining furanic compounds proposed by Caceres-Barrera et al. (2024) from lignocellulosic biomass, using corn stover as the main raw material. The proposed process is based on a pretreatment that allows the biomass to be

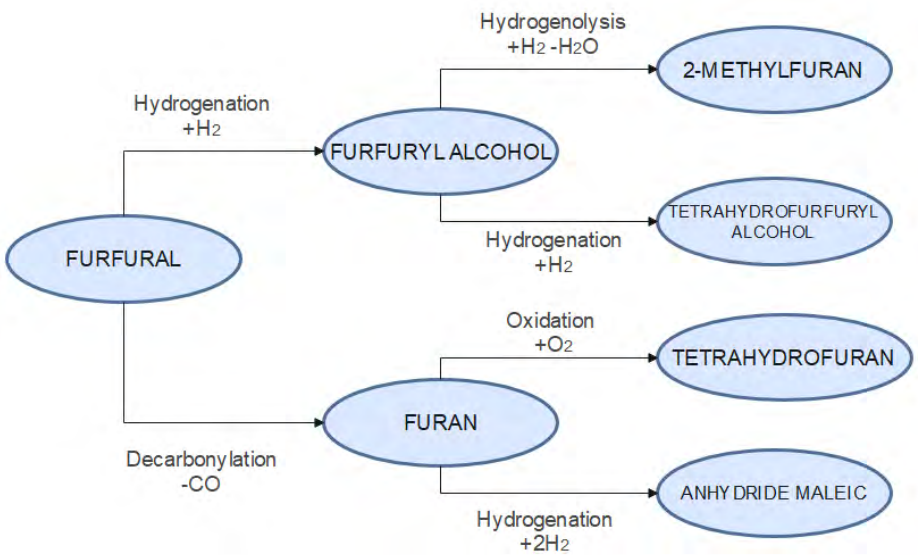

Figure 3. The main production routes for furfural derivatives.

broken down and the constituent sugars to be released, especially the pentoses from hemicellulose, which are subsequently converted into furfural through dehydration reactions. The design strategically integrates the reaction and separation stages to maximize the recovery of furfural and other valuable products such as HMF, levulinic acid, and γ-valerolactone. It is here that we segment the production route and take the final stream of the furfural product as feedstock for the biorefinery proposed in this study.

Process description

The production process was modeled in the Aspen Plus V14 process simulation tool using the diagram shown in Figure 3 to obtain mass and energy balances. The data retrieved from the simulation was used to perform multi-objective optimization using the mathematical model we constructed to estimate the quantity produced of each furfural derivative, maximizing profit. Assuming that the amount of furfural is fixed and predetermined (i.e., fixed plant size) retrieved from the article published by Cáceres-Barrera et al. (2024) [14].

The main production routes for furfural derivatives, as shown in Figure 3, are hydrogenation and oxidation. The reaction kinetics for each of the routes are shown in Table 1.

Finally, the process modeled in Aspen Plus follows the methodology shown in Figure 4. The process mainly consists of six reactive zones for obtaining the main products derived from furfural. The production of 2-Methylfuran includes an additional separation process through a distillation column for the purification of 2-Methylfuran. The maleic anhydride production process includes a purification process through an absorber to include a solvent that separates water from maleic anhydride, which is then obtained through a distillation column.

Table 1: The reaction kinetics.

Kinetics	Reference
$C_5H_4O_2 + H_2 \rightarrow C_5H_6O_2$	[15]
$C_5H_4O_2 + 2H_2 \rightarrow C_5H_6O + H_2O$	[15]
$C_5H_4O_2 + 4H_2 \rightarrow C_5H_{10}O_2$	[16]
$C_5H_4O_2 \rightarrow C_4H_4O + CO$	[17]
$C_4H_4O + 2H_2 \rightarrow C_4H_8O_2$	[18]
$C_4H_4O + \frac{3}{2}O_2 \rightarrow C_2H_2O_3 + H_2O$	[17]

Optimization

The production process for the main furfural derivatives presents a suitable case for optimization, with the analysis of the process's economic, environmental, and safety properties serving as the objective function. This can be explained through the design parameters involved in the process that directly impact economic, environmental, and safety indicators. It is important to note that to achieve a process with a positive environmental impact, it is necessary to consider various external variables, such as the chosen solvent. To achieve a design with high recovery of the main furfural derivatives, low environmental impact, low operating costs, and safety, it is necessary to have an appropriate design for the separation and purification units of the main products. Likewise, the choice of the solvent used to purify maleic

anhydride and the appropriate operating conditions play a significant role in the process.

In this project, the objective function involves indicators and metrics necessary for optimizing the production process of the main furfural products, which are described below.

Eco-indicator 99 (EI99)

The advantages of incorporating EI99 in sustainability evaluations are manifold. Firstly, it offers a thorough perspective on the environmental performance of a product or system, enabling decision-makers to prioritize actions and investments in areas with the most substantial impacts. EI99 covers a broad spectrum of environmental categories, aligning with the multidimensional approach of the UN's 2030 Agenda that seeks to tackle interconnected global challenges. Eco-indicator 99 is one of the most widely used environmental impact estimation methods and consists of a quantitative analysis of the life cycle evaluated from beginning to end. The calculation is performed using Eq. (1). Where ω represents the damage weight factor (Pts/kg), Ci represents the impact value for each of the categories i, and α is the value of subcategory j (kg/year). This methodology accounts for the origin of raw material in processing and degradation. It is based on standard ecological indicators, which are numbers that express the total environmental burden of a product or a process. The higher the value of the indicator, the greater the environmental impact. Therefore, the best scenario is the lowest possible EI99.

$$EI99 = \sum_i \sum_j \omega \bullet C_i \bullet \alpha_j \tag{1}$$

Total annual cost (TAC)

Assessing the overall yearly expenses of a process is crucial within a sustainability framework. The evaluation of total annual costs is essential for various compelling reasons. Firstly, by quantifying these costs, organizations and policymakers can comprehensively grasp the environmental, social, and economic implications of their actions. Sustainability goes beyond minimizing immediate expenditure; it requires an evaluation of long-term consequences. Considering the comprehensive cost, which includes direct and indirect financial expenses, environmental impacts, social implications, and potential externalities, allows stakeholders to make well-informed decisions considering the entirety of sustainability factors. This approach fosters transparency and accountability, as it enables a clear evaluation of trade-offs between short-term economic gains and long-term sustainability objectives. The TAC assumes the annualization of the investment cost of the main process equipment over a 10-year amortization period. To calculate it, Eq. (2) is used. Where $C_{TM,i}$ is the capital cost of the equipment in dollars ($), r represents the payback period in years, and $C_{ut,j}$ is the cost of cooling and heating services, in dollars per year ($/year). The objective function evaluates the lowest annual cost of the process while considering the utilities used by the plant and various units. The best scenario is a total annual cost as low as possible to recover the investment.

$$TAC = \frac{\sum_{i=1}^{n} C_{TM,i}}{r} + \sum_{j=1}^{n} C_{ut,j} \tag{2}$$

Individual risk (IR)

The Individual Risk (IR) metric is crucial for sustainability because it evaluates the likelihood of accidents or harmful events affecting individuals, thus supporting the design of safer industrial processes. By quantifying potential risks to human health and safety, IR ensures that processes prioritize safety, reducing the chance of incidents that could harm people or damage the environment [19]. This focus on minimizing risks aligns with the broader goals of sustainability, including regulatory compliance and social responsibility. Moreover, IR supports Sustainable Development Goals (SDGs) such as Good Health and Well-being (SDG 3) and Decent Work and Economic Growth (SDG 8) by promoting safe working conditions and protecting human health, which are essential components of sustainable development [19]. The IR identifies the risk that a person faces based on his position, including the likelihood of an accident resulting in death or serious injury. The IR is defined as shown in Eq. (3). Where, f_i represents the recurrence that one accident will occur, and $P_{x,y}$ is the likelihood that the accident will occur in a particular location.

$$IR = IR = \sum f_i P_{x,y} \tag{3}$$

In the methodology of this project, the instantaneous and continuous risk analysis was performed for each of the equipment involved in the CO_2 capture process that is in contact with the solvent of interest. Catastrophic events such as Boiling Liquid Expanding Vapor Explosion (BLEVE), Unconfined Vapor Cloud Explosion (UVCE), Jet Fire, Flash fire, and toxic explosion were evaluated using specific mathematical models for each case reported in the literature. The ideal scenario is the one with the lowest individual risk index (IR), as it would be a less risky process.

The objective function

In this study, the objective function encompasses three key areas: environmental, economic, and process safety. It is hoped that this work will contribute to selecting the optimal production route for the main furfural derivatives or to a specific analysis of a product of individual interest. Similar studies have been previously reported by Sanchez-Ramirez et al. [20] and Romero-Garcia et al.[21] for carbon capture cases, but a similar methodology is applied. Therefore, based on this prior information regarding the indicators, the general mathematical

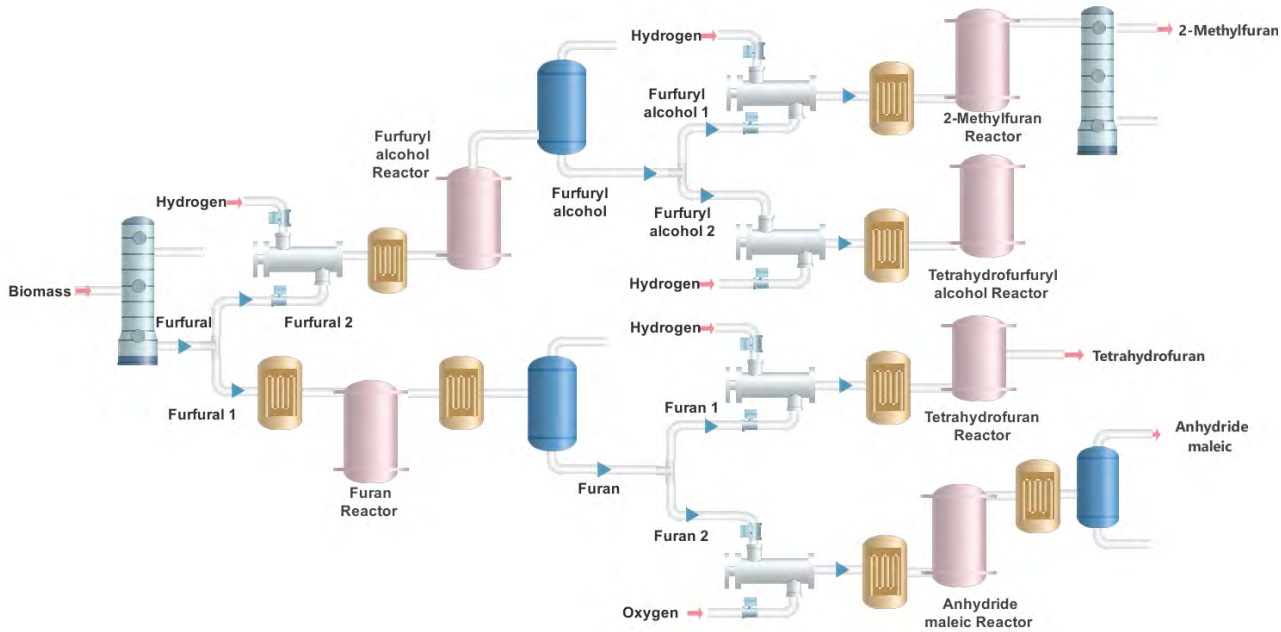

Figure 4. Process simulation diagram.

expression for the objective function, and the respective decision variables involved in the optimization process, are shown below in Equation (4).

$$Min\ [EI99, \mathrm{TAC}, \mathrm{IR}\ ] = f\left(N_{i,j,k}, NF_{i,j}, D_{i,j,k}, DBP, RR, Q\right) \quad (4)$$

For the variables related to the capture process $N_{i,j,k}$ represents the number of stages, $NF_{i,j}$ represents the feed stage, $D_{i,j,k}$ represents the diameter where the i is specified for distillation column, absorber j and de k corresponds to the desorber. DBP represents the solvent feed flow in the absorber, RR represents re reflux ratio and Q the reboiler duty, both in the desorber column.

The proposed model involves a significant number of variables that must be optimized to achieve an optimal design and operating conditions, with the goal of minimizing various performance metrics and maximizing the recovery of the main products of market interest. The implementation of the global optimization approach involved linking Aspen Plus with Microsoft Excel. To initiate the optimization process, the design variables were specified in Microsoft Excel using the stochastic optimization algorithm, Differential Evolution with Tabu List (DETL). During the optimization process, a vector of decision variables is created from Excel to Aspen; in Aspen, the process is evaluated under the parameters of the decision vector, then it is assessed whether the combination of variables is feasible; if not, the algorithm generates a new combination, the resulting data is returned from Aspen Plus and stored in Excel, to apply constraints based on the desired purities or flows, evaluating the previously mentioned indicators, to minimize the objective function and find the optimal process design.

The parameters considered for the DETL method are: population size of 120 individuals, 1, 000 generations, a tabu list size of 60, a tabu radius of 0.0001, a crossover rate of 0.9, and a mutation rate of 0.3. These parameters have been shown to work exceptionally well with nonlinear problems, yielding excellent results.

RESULTS

Analysis of the 3D Pareto Fronts

The scatter plot represents the system's solution space, highlighting the variability among the evaluated indicators. The blue points show all possible configurations, while the orange points represent the smoothed trend of the system's behavior. A nonlinear relationship is observed among the indicators, confirming the existence of a trade-off between them and the multi-objective nature of the problem.

Optimal Design

Table 2 provides a summary of the optimal design for the main purification units in the furfural production process; the optimal design is achieved when the objective functions are in equilibrium with one another and the trade-offs between the objectives are balanced.

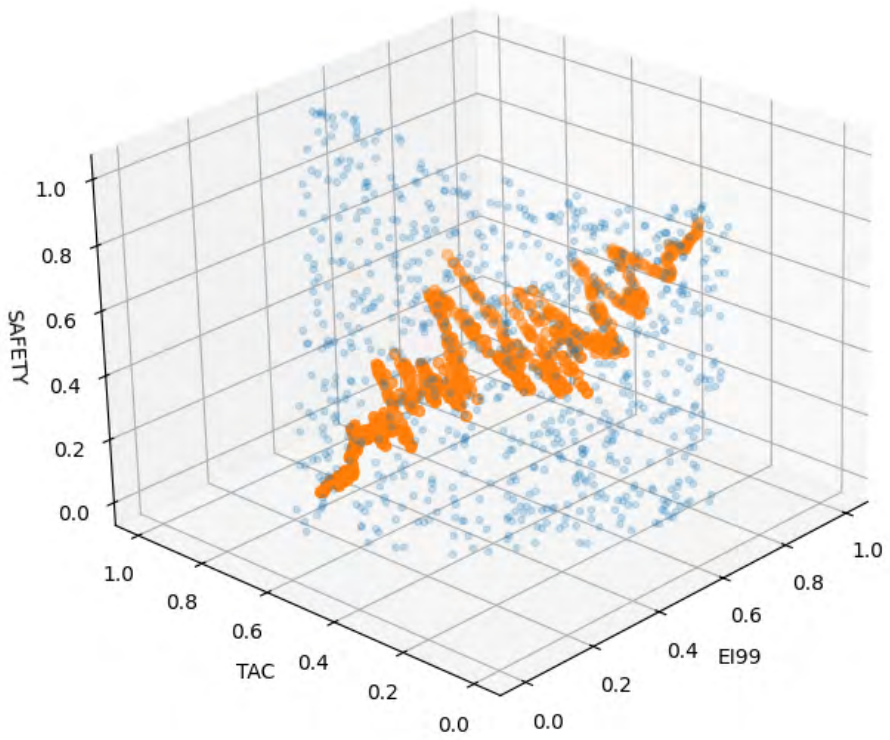

Figure 5. Pareto fronts between the Total annual cost (TAC), Ecoindicator 99 (EI99) and Safety.

Table 2: Design parameters for the best scenario.

Variables	Result
Column 1	
Number of stages	17
Feed stage	8
Diameter (m)	3.3163
Reflux ratio	3.1266
Reboiler duty (GJ/h)	0.6372
Adsorption	
Number of stages	17
Feed stage	1
Diameter (m)	1.9144
Desorption	
Number of stages	15
Feed stage	5
Diameter (m)	2.0566
Reflux ratio	2.6567
Reboiler duty (GJ/h)	1.4854
Objective function	
EI99	42896.24
TAC	4.3281E+10
SAFETY	3.2415E+09

It is possible to observe the impact of certain design parameters on process performance. In the case of the absorption unit, the solvent temperature plays an important role in achieving the adsorption of maleic anhydride and its subsequent separation, thus directly affecting the recovery efficiency of this main product. Furthermore, the reflux ratio is directly related to energy requirements. The higher the reflux ratio, the greater the energy requirements for capture. Regarding the analysis of sustainability indices, it is observed that the lowest EI99 is achieved when lower amounts of energy are used in the process, which is largely related to equipment sizes; this trend is related to the process topology.Energy efficiency

Figure 6 highlights an important perspective on efficiency by comparing the production rate (kg/h) with the energy consumption rate (MJ/h) for each component. It is evident that not all processes use energy with the same efficiency: for example, FURAN and FOL achieve relatively high production with low to moderate energy consumption, indicating efficient performance, while THF and THFA consume more energy without a proportional increase in production, making them less efficient. This contrast becomes more pronounced with MF and, especially, with MA, where energy consumption is significantly higher despite having production levels similar to those

of other components, revealing energy-intensive processes. Overall, the graph demonstrates that higher energy consumption does not necessarily translate to higher production, underscoring the importance of evaluating and optimizing energy efficiency in process design.

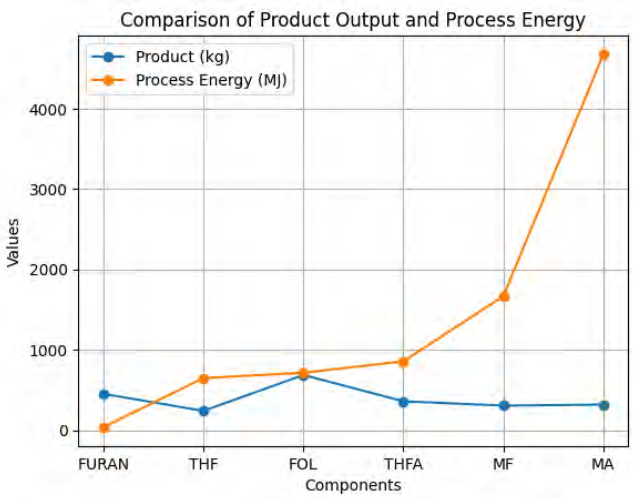

Figure 6. Energy efficiency.

CONCLUSIONS

This study proposes the design and optimization of a production plant for the main furfural derivatives. Simultaneous production is achieved through the implementation of flow separators for the main streams of first-generation furfural products, thereby enabling the simultaneous production of six of the main derivatives. The evaluation of metrics related to environmental impact, economic assessment, and process safety provides a broader perspective on how different variables can affect the sustainability of the process. Finally, the implementation of the multi-objective optimization algorithm provides a broader perspective on process control, with the objective function being to minimize these economic, environmental, and safety aspects of the process; this aims to achieve the optimal design that meets the established constraints and serves as the starting point for its implementation or scaling.

In conclusion, this work creates an opportunity for future research on the main furfural production routes. Considering the sustainability framework, there is significant potential for process improvement and adaptation, specifically regarding the implementation of new production routes based on this work or the implementation of a new solvent to recover maleic anhydride from the process.

ACKNOWLEDGEMENTS

The authors acknowledge the financial support provided by Universidad de Guanajuato and SECIHTI.

AUTHOR IDENTIFIERS

Author ORCIDs:
Coronel Muñoz Melanie: 0009-0004-4169-5758
Sánchez Ramírez Eduardo: 0000-0002-4326-4837
Segovia Hernández Juan G.: 0000-0003-4175-9809
Quiroz Ramírez Juan José: 0000-0002-9005-738X
Huerta Rosas Brenda: 0000-0002-6361-6094

REFERENCES

1. Corma A, Iborra S, Velty A. Chemical routes for the transformation of biomass into chemicals. Chem. Rev. 107:2411-2502 (2007). https://doi.org/10.1021/cr050989d
2. Tuan Hoang A, Viet Pham . 2-methylfuran (MF) as a potential biofuel: a thorough review on the production pathway from biomass, combustion progress, and application in engines. Renewable and Sustainable Energy Reviews 148:111265 (2021). https://doi.org/10.1016/j.rser.2021.111265
3. Nakagawa Y, Tomishige K. Production of 1,5-pentanediol from biomass via furfural and tetrahydrofurfuryl alcohol. Catalysis Today 195:136-143 (2012). https://doi.org/10.1016/j.cattod.2012.04.048
4. Kottke RH. Furan derivatives. Kirk-Othmer Encyclopedia of Chemical Technology : (2000). https://doi.org/10.1002/0471238961.0621180111152020.a01
5. Bowron DT, Finney JL, Soper AK. The structure of liquid tetrahydrofuran. J. Am. Chem. Soc. 128:5119-5126 (2006). https://doi.org/10.1021/ja0583057
6. Kabwe T, Louw J, Gorgens JF. A techno-economic comparison of chemical derivatives from furfural produced from bagasse at an existing sugar cane mill. Industrial Crops and Products 228:120890 (2025). https://doi.org/10.1016/j.indcrop.2025.120890
7. Maleic Anhydride. (s. f.). Google Books. https://books.google.com/books?hl=es&lr=&id=c8XeBwAAQBAJ&oi=fnd&pg=PA1&dq=Maleic+anhydride&ots=yDFpoHyM83&sig=utOgYrrP2G9tVpSrbrYg3t0Vmd0
8. Furfuryl Alcohol Market Size & Share | Industry Report, 2030. (s. f.). https://www.grandviewresearch.com/industry-analysis/furfuryl-alcohol-market
9. https://www.futuremarketinsights.com/reports/2-methylfuran-market

10. Research, I. M. (2025, 20 November). Tetrahydrofurfuryl Alcohol Market - latest advancement & future trends (2024-2032). Introspective Market Research. https://introspectivemarketresearch.com/reports/tetrahydrofurfuryl-alcohol-market/
11. Market Research Future. (2026, 13 January). Furan Resin Market Size, Share & Global Industry Report 2035. https://www.marketresearchfuture.com/es/reports/furan-resin-market-28311
12. Tetrahydrofuran Market Size, Share & Trends Report, 2030. (s. f.). https://www.grandviewresearch.com/industry-analysis/tetrahydrofuran-market-report
13. Maleic Anhydride Market Size, Share & Trends Report, 2030. (s. f.). https://www.grandviewresearch.com/industry-analysis/maleic-anhydride-market
14. Caceres-Barrera CR, Sánchez-Ramírez E, Juárez-García M, Alcocer-García H, Segovia-Hernández JG. Optimization of separation processes in multiproduct biorefinery design to produce furan-based compounds and their derivatives using performance indicators. Ind. Eng. Chem. Res. 63:20950-20962 (2024). https://doi.org/10.1021/acs.iecr.4c02646
15. Tseng YT, Wang WJ, Ward JD, Lee HY. Design and control of a process to produce furfuryl alcohol. Journal of the Taiwan Institute of Chemical Engineers 51:44-52 (2015). https://doi.org/10.1016/j.jtice.2015.01.001
16. Šivec R, Huš M, Likozar B, Grilc M. Furfural hydrogenation over cu, ni, pd, pt, re, rh and ru catalysts: ab initio modelling of adsorption, desorption and reaction micro-kinetics. Chemical Engineering Journal 436:135070 (2022). https://doi.org/10.1016/j.cej.2022.135070
17. Zeitsch, K. J. (2000). The Chemistry and Technology of Furfural and Its Many By-products. Elsevier Science Limited.
18. Godawa C, Gaset A, Kalck P, Maire Y. Mise en oeuvre d'un catalyseur actif pour l'hydrogenation selective du furanne en tetrahydrofuranne. Journal of Molecular Catalysis 34:199-212 (1986). https://doi.org/10.1016/0304-5102(86)87066-3
19. Mannan, S. Lees' Loss Prevention in the Process Industries: Hazard Identification, Assessment And Control: Fourth Edition. *Lees' Loss Prevention in the Process Industries: Hazard Identification, Assessment and Control: Fourth Edition* **2012**, *1–2*, 1–3642. https://doi.org/10.1016/C2009-0-24104-3
20. Sánchez-Ramírez E, Quiroz-Ramírez JJ, Hernández S, Segovia-Hernández JG, Kiss AA. Optimal hybrid separations for intensified downstream processing of biobutanol. Separation and Purification Technology 185:149-159 (2017). https://doi.org/10.1016/j.seppur.2017.05.011
21. Romero-García AG, Ramírez-Corona N, Sánchez-Ramírez E, Alcocer-García H, De Blasio C, Segovia-Hernández JG. Sustainability assessment in the CO2 capture process: multi-objective optimization. Chemical Engineering and Processing - Process Intensification 182:109207 (2022). https://doi.org/10.1016/j.cep.2022.109207

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Multi-Level Optimization of Crane Scheduling

Sophia Onyshkevych[a*], Bianca Springub[a], and Christos Galanopoulos[a]

[a] Aurubis AG, Data Analytics & Digital Twin, Hamburg, Germany
* Corresponding Author: s.onyshkevych@aurubis.com

ABSTRACT

Copper refining via electrolysis is a core metallurgical process that takes place in tankhouses, subject to strict temporal, spatial, and operational constraints. The efficiency and stability of this process depend critically on the coordinated scheduling of crane operations responsible for handling anodes, cathodes, and auxiliary tasks. In industrial practice, crane scheduling must simultaneously satisfy long-term production targets and short-term operational feasibility, while respecting process-dependent timing constraints imposed by electrochemical parameters. Inefficient or inconsistent schedules can lead to process delays, suboptimal resource utilization, and degraded electrolysis performance, ultimately affecting product quality and operational stability. This paper presents a modeling approach for optimizing tankhouse operations. The uniqueness of this case lies in the broad range of constraints, including human capacity, energy restrictions, metallurgical rules, and logistical specifications. The model operates on three levels, starting with generalized schedule optimization and culminating in the detailed optimization of crane movements and cathode stripping machine operations. Although the focus is on the specific use case of tankhouse operations within metallurgical production, the model can be adapted and fine-tuned for a wide range of scheduling applications. Additionally, we discuss the challenges encountered when integrating the model into existing production processes and outline directions for further work.

Keywords: Modelling, Optimization, Scheduling, Industry 4.0, Process Operations, Pyomo

INTRODUCTION

Every year, Aurubis produces 1.1–1.2 million metric tons of copper cathodes with an extremely high copper content, at a minimum of 99.99% purity [2].

To achieve this material purity, the anodes undergo an electrolysis process in the tankhouse as the final step [1]. The quality of the resulting cathodes depends on the duration and electricity levels during the electrolysis process. Therefore, timely processing of all groups is one of the main priorities.

The current setup includes several groups of different types of cathodes, some of which share electrical circuits, as well as cranes that transfer cathodes from the electrolyte bath to the copper stripping machines. These cranes are tasked with several operations, including lifting, transporting, and positioning both cathodes and anodes in and out of the electrolyte cells. These operations require high precision to avoid damaging the delicate metal deposits and to ensure the safe handling of heavy materials. The exchange process must be executed in a sequential manner, where each crane handles one anode or cathode at a time, often under strict timing constraints.

One of the main modeling challenges is the scalability of the model. In the production environment, operators aim to receive optimized results as quickly as possible, without sacrificing the operational details. Therefore, addressing the problem in the traditional MILP fashion might be quite tricky and lead to a rapid increase in variables [3]. However, other approaches often disregard essential operational rules, making them unsuitable for industrial settings.

In this paper, we describe a model that is feasible within our industrial setting, while keeping the parameters flexible to accommodate a broader range of applications. This balance between computational efficiency and operational accuracy is crucial for practical implementation in real-world production environments.

Problem setup

The key elements of the tankhouse setup are cathode and anode stripping machines and cranes which operate on groups of electrolyte cells. Each group contains

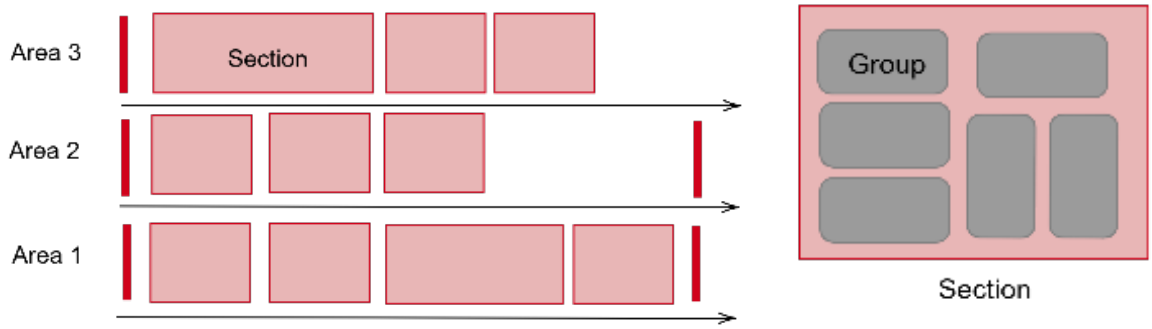

Figure 1. Schematic depiction of initial positions of cranes (in bold red) in the tankhouse layout (left) and a zoomed-in depiction of the example group layout in a section (right).

5-20 cells of cathodes/anodes. In addition, several groups are combined into section which share one electrical circuit. Further there are several tankhouse areas – lanes on which specific numbers of cranes operate. Some areas have one crane, whereas some have two, which causes the problem of tracking the distance between cranes and ensuring that the cranes do not intersect or collide.

The Cathode-Anode Exchange Process

During the copper refining process, metal ions are reduced and deposited onto a cathode in an electrolyte bath. The anode, typically made of a material like impure copper, supplies the metal ions that dissolve into the electrolyte. This process unfolds as follows: the anode gradually dissolves to provide metal ions to the electrolyte solution, which are then deposited on the cathode [1].

Therefore, cathodes must be exchanged regularly to remove the deposited metal, while a fresh cathode is placed in the cell. The anode also needs to be replaced, but the rate at which the anode is consumed, and the frequency of replacement are different from those of the cathode. In this paper, we assume that in one cycle (lifespan of one anode) there are two cathode changes followed by one cathode and anode change.

Typical operations in the tankhouse involve removing the spent anodes and replacing them with new ones while simultaneously removing the deposited cathodes. The sequence of operations consists of the following key steps:

- Lifting: The crane is responsible for lifting the spent anode and moving it from the electrolyte cell to the storage area.
- New Anodes Placement: Once the old anode is removed, the crane places a new anode into the electrolyte cell to continue the refining process.
- Lifting the Deposited Cathodes: The crane then lifts the cathode, which has the refined metal deposits, and moves it to a cathode stripping machine.
- Depositing the Refined Cathodes: After lifting, the crane transports the cathodes to a processing area for further refining or packaging.

Furthermore, the tankhouse operation is designed to be continuous and needs to maintain a constant flow of materials. Thus, it's important to create a schedule of operations that can be repeated indefinitely, accounting for more frequent change of the cathodes than the anodes as well as provide a model to support daily planning operations to optimize crane movements.

FIRST STAGE: CYCLIC OPERATION SCHEDULING MODEL

Problem Description

We consider a cyclic scheduling problem for repetitive operations on groups of cells. Each group consists of several cells and belongs to one of the areas (Figure 1). Over a finite planning horizon of D days, a schedule is created that is assumed to repeat indefinitely.

Each group must undergo three operations, denoted O_1, O_2 and O_3 (corresponding to two cathode changes and an anode change as described in previous section) which must always be executed in the fixed cyclic order:

$$O_1 \rightarrow O_2 \rightarrow O_3 \rightarrow O_1 \rightarrow \cdots$$

The time intervals between consecutive operations are fixed. Let δ_{ij} denote the number of days between O_i and O_j. These spacings, $\delta_{12}, \delta_{23}, \delta_{31}$, must be respected exactly, including across the boundary of the planning horizon.

Since the plan is cyclic, groups do not need to start

with O_1. Instead, each group may start at any operation in the sequence, which introduces a phase shift that is decided by the optimization model.

The main goal is to balance the daily workload, which is defined as the number of cells processed per day. Optionally, an additional balancing requirement can be imposed to ensure that the daily workload is also approximately balanced between the areas and groups (Figure 1). This ensures a more uniform utilization of resources and helps prevent operational bottlenecks across the planning horizon.

Model Formulation

The model uses the following sets:

- $\mathcal{G}$: set of groups
- $\mathcal{D} = \{1, \dots, D\}$: set of days in the planning horizon
- $\mathcal{O} = \{O_1, O_2, O_3\}$: set of operation types

The parameters of the model are:

- $c_g \in \mathbb{N}$: number of cells in group g
- $s_g \in \{A_1, A_2, A_3\}$: area of group g
- $\delta_{ij} \in \mathbb{N}$: fixed number of days between operation i and its successor j
- $D \in \mathbb{N}$: length of the cyclic planning horizon

Binary decision variables are defined as:

$$x(g,d,o) = \begin{cases} 1, \text{ if group } g \text{ performs operation } o \text{ on day } d \\ 0, \text{ else,} \end{cases} \forall g \in \mathcal{G}, d \in \mathcal{D}, o \in \mathcal{O}$$

To model the cyclic phase choice, additional binary variables are introduced:

$$y(g,o) = = \begin{cases} 1, \text{ if group } g \text{ starts its cycle with operation } o \\ 0, \text{ else,} \end{cases} g \in \mathcal{G}, o \in \mathcal{O}$$

The model is completed with the following constraints:

- Phase selection: Each group must select exactly one starting operation:

$$\sum_{o \in \mathcal{O}} y(g,o) = 1, \forall\, g \in \mathcal{G}.$$

- At most one operation per group per day:

$$\sum_{o \in \mathcal{O}} x(g,d,o) \le 1, \forall g \in \mathcal{G}, d \in \mathcal{D}.$$

- Cyclic sequence and spacing constraints: If a group performs operation i on day d, it must perform the successor operation j exactly δ_{ij} days later, where the day index is computed modulo D.This enforces the fixed sequence $O_1 \to O_2 \to O_3 \to O_1$ and ensures consistency under cyclic repetition. In the mixed-integer linear formulation, these implications are implemented using precomputed feasible day pairs or standard linearization techniques.

The total workload on day d is defined as the sum of the cells processed by all groups on that day:

$$W_d = \sum_{g \in \mathcal{G}, o \in \mathcal{O}} c_g x(g,d,o), \forall d \in \mathcal{D}.$$

Let $\bar{W}$ denote the average daily workload over the planning horizon. To obtain a balanced plan using linear objective, the model minimizes the sum of absolute deviations of the daily workload from this average.

This is achieved using linear constraints that relate W_d and $\bar{W}$ to the deviation variables z_d:

$$z_d \ge W_d - \bar{W}, \forall d \in \mathcal{D}.$$

$$z_d \ge -(W_d - \bar{W}), \forall d \in \mathcal{D}.$$

Then the objective function is given by

$$\min \sum_{\forall d \in \mathcal{D}} z_d.$$

If area balancing is activated, additional deviation terms are introduced to penalize differences between area workloads on each day. Analogously one can introduce deviation variables and constraints to add to the objective function balancing per section or specific group, if required.

Key Challenges and assumptions

The main challenges of the proposed model are:

- Representing cyclic time dependencies using modulo arithmetic.
- Allowing flexible starting phases while enforcing a strict operation sequence.
- Maintaining a linear formulation for solver efficiency.
- Ensuring scalability as the number of groups and planning days increases.

The proposed mixed-integer linear model generates cyclic operation schedules with fixed sequencing and exact inter-operation spacing. By allowing flexible phase shifts, the model achieves balanced daily workloads and can optionally balance workloads between spatial sections. The formulation is suitable for standard MILP solvers and can be extended to incorporate additional operational constraints.

In this setup, the time between each operation δ_{ij} is the same for each group. However, depending on the specific properties and composition of anodes, they can vary. This can be modelled by introducing $\Delta = \delta_{ij}(g)$, for

each group g. Consequently, the plan is executed with a periodicity equal to the Least Common Multiple (LCM) of all cycle durations. This extension increases the model's flexibility, while also introducing additional computational complexity due to the longer effective planning cycle.

SECOND STAGE: DAILY PROCESSING OPTIMIZATION

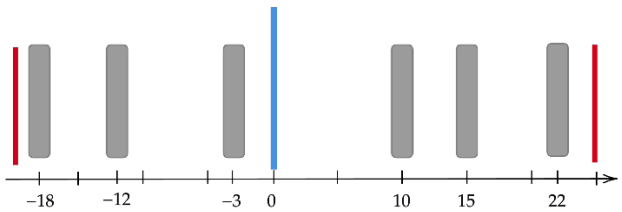

Figure 2. Schematic depiction of an area with two cranes (red), 6 groups to be processed (in grey) and a cathode machine (position 0, in blue).

Problem Description

We consider now a planning on a daily level. We aim to model the process given in section The Cathode-Anode Exchange Process. The problem can be summarized as follows: a set of groups must be processed using one or more overhead cranes and a single cathode stripping machine. Each group consists of multiple identical cells that must be transported individually from the group location to the machine, processed, and returned. The objective is to determine an optimal processing sequence that minimizes the total completion time (makespan) of all groups.

We set the following assumptions in our modeling setup:

- Each group contains between 5 and 20 cells.
- All cells within a group are assumed to be co-located at a fixed position for modeling simplicity.
- Groups are located at different distances to the left or right of the machine.
- Cranes move along a shared rail and cannot intersect or pass each other.
- Exact timing of individual movements is not critical; instead, the focus is on optimizing the processing sequence in a realistic manner.

Given these properties, the problem is modeled as a sequencing and assignment problem at the group level, where the processing time of each group is derived from the number of cells and the travel distance to the machine.

Model Formulation

The main bottleneck is efficient utilization of a cathode machine, since in our setup we consider only one cathode machine per plant, shared between cranes and areas which is used 3 times per cell every cycle. Following this we focus in our model on crane movement from group to cathode machine and back, excluding movement of the crane to the anode machine during anode operation (O_3) from the schedule.

Further, in this paper we do not focus on sections (Figure 1) which represent shared electrical circuit, since we do not aim to optimize shutdown time in our case. However, we remark that our model could be easily extended to this objective. Therefore, we schedule operation on group level considering only varying number of cells per group and ignoring to which section which group belongs (Figure 2).

The formulation employs a discrete-event-based approach to capture the interactions between cranes and jobs – operations performed in a specific group.

The cranes are modeled as parallel machines with sequencing decisions. Crane interference is handled by enforcing spatially consistent processing orders, reflecting the fact that cranes move on a single rail and cannot cross.

The model uses the following sets:

- $\mathcal{J}$: set of jobs
- $\mathcal{C}$: set of cranes
- $\mathcal{T}$: set of time slots
- The parameters of the model are:
- $P_InitialInventory(j)$: the initial number of cells in the group represented by job j
- $P_CathodeProcessingSpeed(t, j, c)$: processing speed of job j on crane c during time t. This parameter is 0 if job belongs to an area which is not covered by crane c.
- $P_Location(j)$: x-coordinate location of job j on the rail.

Binary decision variables are defined as:

$$BV_JobAssignCrane(j, c) = \begin{cases} 1, \text{if job j is assigned to crane } c \\ 0, \text{else,} \end{cases} \forall j \in \mathcal{J}, c \in \mathcal{C},$$

$$BV_JobCraneProcess(t, j, c) = \begin{cases} 1, \text{if job j is being processed at time t} \\ 0, \text{else,} \end{cases} \forall j \in \mathcal{J}, c \in \mathcal{C}, \text{t} \in \mathcal{T},$$

$$BV_RemainingCells(t, j) = \begin{cases} 1, \text{if job j has remaining cells at time t} \\ 0, \text{else,} \end{cases} \forall j \in \mathcal{J}, \text{t} \in \mathcal{T},$$

$$BV_BothActive(t,j,j') = \begin{cases} 1, \text{if job j and job j' are scheduled at time t} \\ 0, \text{else,} \end{cases} \forall j,j' \in \mathcal{J}, t \in \mathcal{T}$$

$$BV_JobLeft(t,j,j') = \begin{cases} 1, \text{if job } j \text{ is left to job } j' \text{ at time } t \\ 0, \text{else,} \end{cases} \forall j,j' \in \mathcal{J}, \text{t} \in \mathcal{T}$$

And the continuous decision variables:

- $V_Inventory(t,j)$: the inventory of cells for job j at time t
- $V_CellsProcessingAmount(t,j,c)$: the number of cells processed for job j by crane c at time t.
- The model aims to minimize an objective related to crane idle time or the makespan, i.e., the total time required to complete all jobs:

$$\min T_{max} = \sum_{\text{t}\in\mathcal{T}, j\in\mathcal{J}} t \bullet V_Inventory(t,j)$$

- The main constraints are given as follows:
- Each group is assigned to exactly one crane:

$$\sum_{c\in\mathcal{C}} BV_JobAssignCrane(j,c) = 1, \forall j \in \mathcal{J},$$

- At any given time, only one job can be processed on a crane:

$$\sum_{j\in\mathcal{J}} BV_JobCraneProcess(t,j,c) \leq 1, \forall\, t \in \mathcal{T}, c \in \mathcal{C}$$

- A job can only be processed on a crane if it is assigned to that crane:

$$BV_JobCraneProcess(t,j,c) \leq BV_JobAssignCrane(j,c), \forall\, t \in \mathcal{T}, c \in \mathcal{C}, j \in \mathcal{J}$$

- The mass balance given by:

$$V_Inventory(t,j) = V_Inventory(t-1,j) - \sum_{c\in\mathcal{C}} BV_JobCraneProcess(t,j,c), \forall\, t \in \mathcal{T}, c \in \mathcal{C}, j \in \mathcal{J}$$

- Job completeness:

$$\sum_{t\in\mathcal{T}, c\in\mathcal{C}} V_CellsProcessingAmount(t,j,c) = P_InitialInventory(j), j \in \mathcal{J}$$

- Continuous unloading across time slots:

$$BV_JobCraneProcess(t+1,j,c) \leq BV_JobCraneProcess(t,j,c) + 1, \forall\, t \in \mathcal{T}, c \in \mathcal{C}, j \in \mathcal{J}$$

$$BV_JobCraneProcess[t+1,j,c] \leq BV_RemainingCells(t,j), \forall\, t \in \mathcal{T}, c \in \mathcal{C}, j \in \mathcal{J}$$

Crane Non-Intersection Constraints

Another important constraint necessary is the crane non-intersection constraint. The hanging cranes operate on a single shared rail per area and therefore cannot cross or overtake one another during operation. Rather than explicitly modeling crane trajectories over time, the proposed formulation enforces non-intersection implicitly through sequencing constraints derived from the spatial ordering of groups.

Each group $j \in \mathcal{J}$ is associated with a fixed spatial coordinate $P_Location(j)$, representing its position along the crane rail. Without loss of generality, negative coordinates correspond to positions to the left of the machine, and positive coordinates to positions to the right.

For two groups $j, j' \in \mathcal{J}$, with $P_Location(j) < P_Location(j')$, any schedule in which crane c_1 completes processing of group j before crane c_2 completes processing of group j' would require the cranes to cross at some point in time. Such schedules are therefore physically infeasible. These constraints are encoded as logical sequencing restrictions and eliminate schedules that would require crane intersection as follows:

$$P_Location(j) + K \leq P_Location(j') + M \bullet (1 - BV_JobLeft(t,j,j')) + M \bullet (1 - BV_BothActive(t,j,j')) \forall j,j' \in \mathcal{J}, \text{t} \in \mathcal{T},$$

where $M > 0$ is a sufficiently large Big-M constant and K is a minimal safety distance between the cranes [4].

Possible Model Extensions

While this paper presents the simplest objective function, the proposed model can also be adapted to optimize additional factors, such as minimizing electricity usage (e.g., reducing shutdowns per section) or maximizing the throughput of the cathode machines.

Furthermore, in real-world setups, crane movement speeds may vary between areas, and certain groups may have higher processing priorities. These operational nuances can be incorporated into the model using penalty terms and scaling parameters.

The model can also be extended to add further realism, such as accommodating partial group processing, handling heterogeneous crane speeds, and introducing safety buffers for movement and waiting times at the groups and cathode machines.

RESULTS

Cyclic schedule

We present now the results of optimizing a cyclic schedule for operations for 24 days. In this example we consider around 3000 cells distributed across 10 groups.

We chose $\delta_{12} = 7, \delta_{23} = 8, \delta_{31} = 9$. We first try to find a schedule which balances the operations across 3 areas, without additional restriction on balancing the operations per group and per day. This results already in

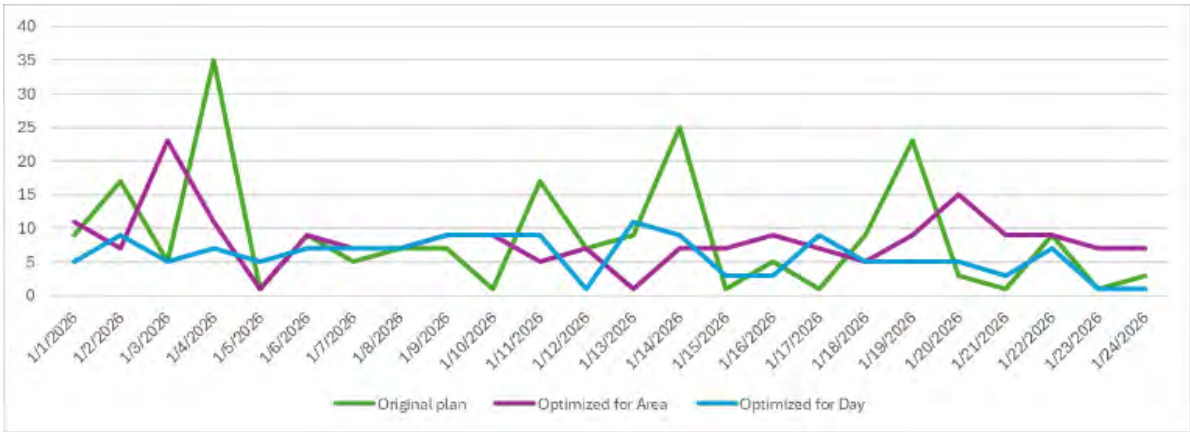

Figure 3. Deviation from processed avarage cell amount per day in the original plan (green), optimized plan with prioritized balanced distribution per area (purple), optimized plan with prioritized balancing per day (blue).

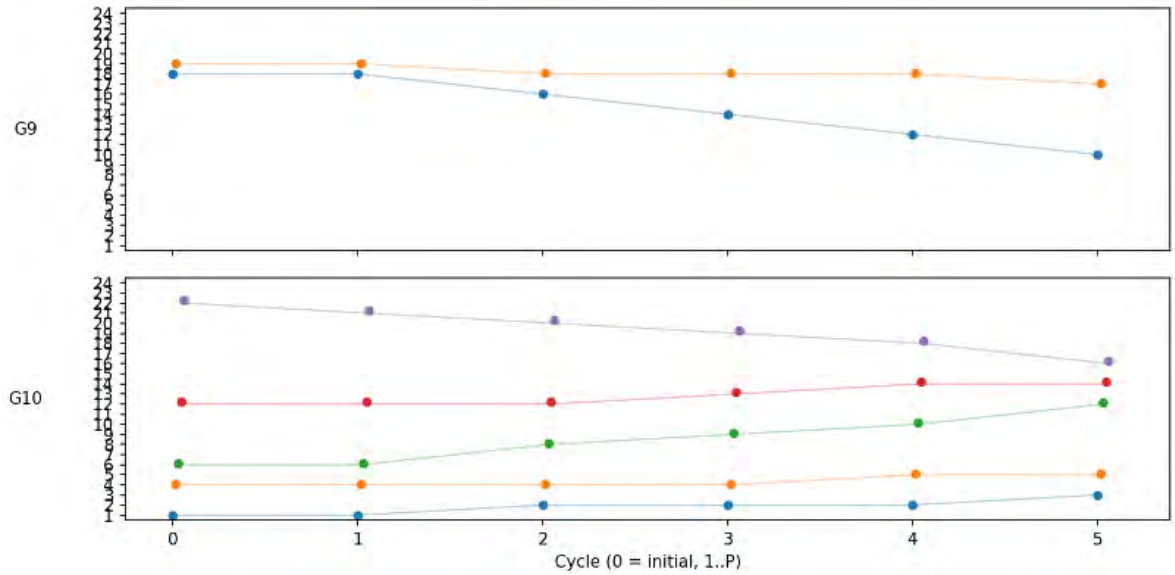

Figure 4. Optimized transition plan for group G9, G10 to be completed after P=5 cycles from the 24-day schedule to 21-day schedule.

improvement compared to the base scenario (Figure 3). By additionally minimizing the deviation per day for total amount of cells we achieve even more equal workload.

Daily schedule

We consider again the example showed on Figure 2: we are given a list of jobs to perform for each group. Each group is located at different position and can be accessed by one of two cranes. The details of each group are summarized in Table 1.

For simplicity we assume that a crane can transport only one cell per time (i.e. $V_CellsProcessingAmount$ is always 0 or 1), there is no additional time required for lifting or putting the cathode into the cell. We aim to omit crane intersections but allow multiple cranes to stay at the cathode machine (located at 0) to mimic the queuing. In practice, would add some crane movements to free up the space in post-processing.

Table 1: Groups to be scheduled with their location w.r.t. cathode machine (position 0) and number of cells.

Group	Location	# Cells
1	22	10
2	15	15
3	-12	12
4	10	16
5	-18	16
6	-3	16

The obtained results for this example are shown on Figures 5-7. The crane movement plot visually depicts the movement of the cranes across the rail system during the entire scheduling horizon. It can be observed that the crane movements are synchronized to avoid overlapping tasks and minimize the overall travel time. This is

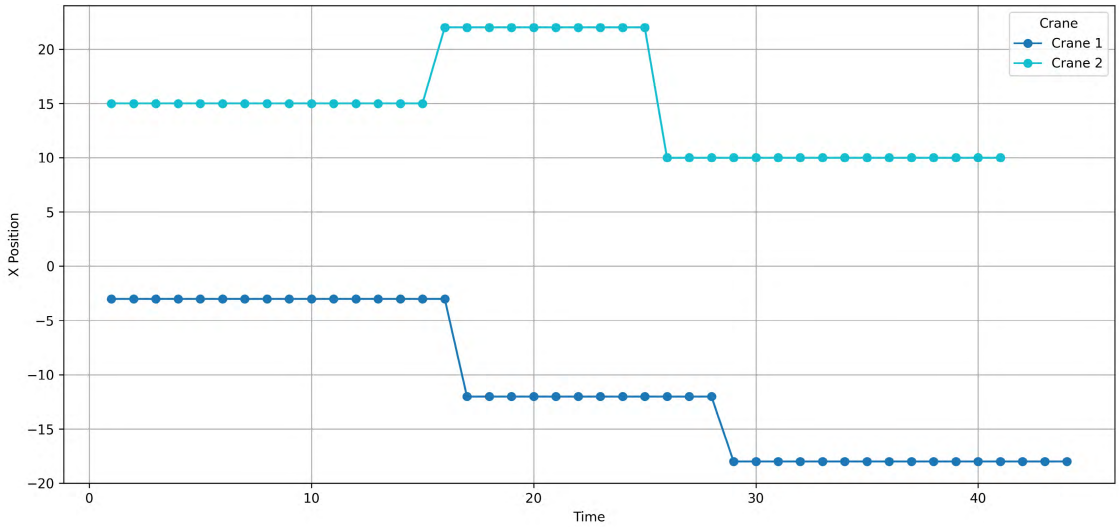

Figure 5. Crane movement with regards to the groups.

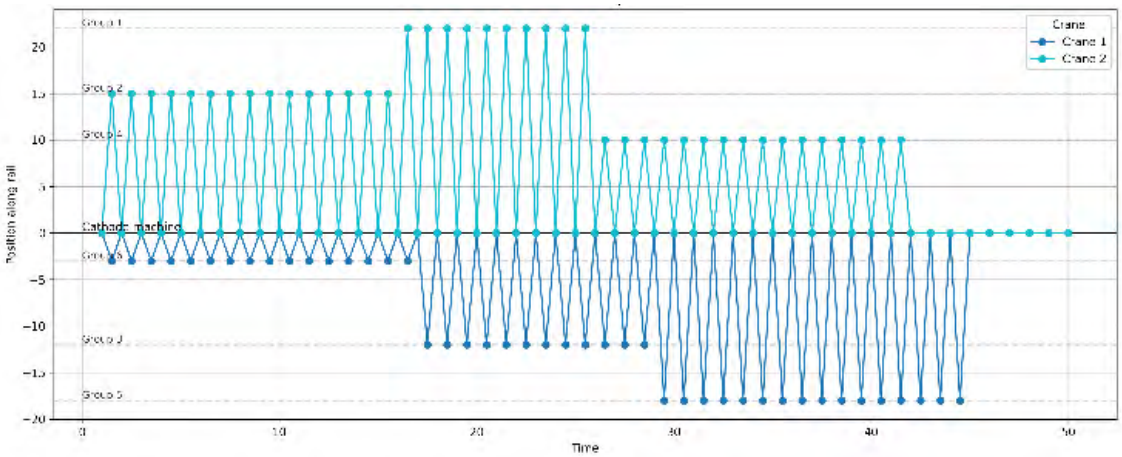

Figure 6. Crane movement with regards to the groups.

achieved by enforcing the non-interference constraints, which ensure that cranes do not cross paths or perform operations at the same location at the same time. The plot also shows how crane movements are efficiently scheduled, maximizing their utilization while respecting the available space and minimizing idle times.

On Figure 6 we show more detailed movement trajectory of each of the cranes to illustrate the active no-intersection condition and repeated travelling to cathode machine to return used and bring new cathodes.

Figure 7 provides a time-based visualization of the schedule, which shows the systematic distribution of job processing across time, ensuring that all jobs are completed within the required timeframe.

FUTURE WORK

While the results obtained from the optimization model discussed in previous sections are promising, an intriguing question arose during the project regarding the practical implementation of the optimized plan, particularly in terms of minimizing disruptions to an existing operational schedule. Specifically, the challenge lies in determining how to transition from the current, potentially suboptimal, schedule to the optimized plan while ensuring minimal disturbances to ongoing operations.

Further, depending on the anode availability and electricity prices the duration of the cycle can be increased or reduced from the modelled X-days. Therefore, the same question arises: how one would transition optimally (maybe within several cycles) from X-day plan to Y-day plan and from X day plan to the optimized X day plan.

A straightforward solution would be to employ linear interpolation, as illustrated in Figure 4. This approach would allow for a smooth transition between schedules

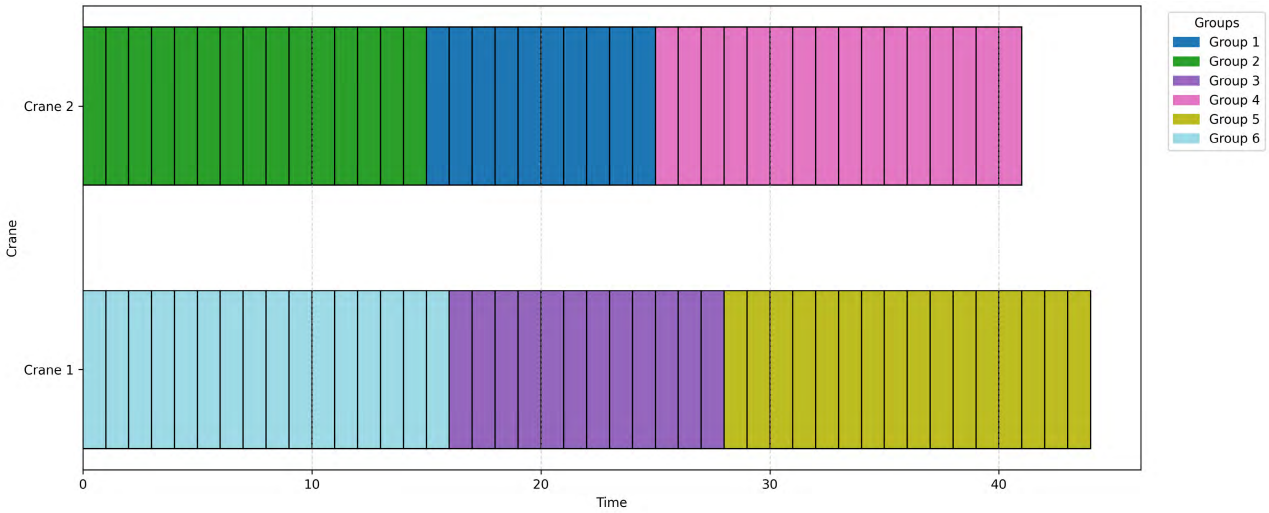

Figure 7. Job processing timeline for each of the cranes.

by gradually adjusting the start times and job assignments. However, while this method is simple, it may not fully capture the complexities of crane operations and job sequencing, especially under dynamic constraints.

In future work, we aim to explore more sophisticated techniques for these transitions. Possible approaches include the use of rolling horizon optimization, adaptive scheduling methods, or real-time re-optimization strategies, which would enable the system to adjust dynamically to changes in operational conditions, such as equipment availability and energy price fluctuations. These methods could potentially provide more robust solutions that balance operational efficiency with flexibility, ensuring smoother transitions between different scheduling plans.

CONCLUSION

This paper presents an optimization model for scheduling crane operations in copper refining tankhouses, addressing complex constraints such as metallurgical rules and operational feasibility. Despite the complexity of the problem, the model successfully integrates key operational parameters, ensuring minimal disruptions to existing processes while maximizing resource utilization and electrolysis efficiency. The results, demonstrated through optimized schedules, show potential for improving operational throughput and reducing downtime within the tankhouse.

ACKNOWLEDGEMENTS

The authors would like to express their gratitude to their colleagues from the production team at Aurubis for their valuable collaboration and insights throughout this project.

AUTHOR IDENTIFIERS

Author ORCIDs:
Onyshkevych S: 0009-0003-8735-4468
Springub B: 0009-0007-5036-8224
Galanopoulos C: 0000-0002-3353-3247

REFERENCES

1. Schlesinger M, King M, Davenport W. Extractive Metallurgy of Copper (5th Ed). Elsevier (2011) https://doi.org/10.1016/C2010-0-64841-3
2. Aurubis Website. https://www.aurubis.com/en/products/copper-cathodes
3. Sun Y, Onyshkevych et al. Optimal Crane Scheduling for Tankhouse Copper Processing. (preprint) http://dx.doi.org/10.2139/ssrn.6189940
4. Mondello L. Foodomics – advanced mass spectrometry in modern food science and nutritionalejandrocifuentes, ed. john wiley & sons, inc. 2013 ISBN 978?1?118?16945?2 hardcover, 580 pages € 118.40. J of Separation Science 36:1869-1869 (2013). https://doi.org/10.1002/jssc.201370103

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Modelling of carbon dioxide methanation in radial flow reactor

Salvatore Capasso[ab], Vincenzo Russo[b], Henrik Grénman[a*]

[a] Faculty of Science and Engineering, Johan Gadolin Process Chemistry Centre, Laboratory of Industrial Chemistry and Reaction Engineering, Åbo Akademi University, Turku/Åbo, Finland
[b] Department of Chemical Sciences, University of Naples Federico II, IT-80126 Naples, Italy
* Corresponding Author: henrik.grenman@abo.fi

ABSTRACT

Carbon dioxide hydrogenation to produce methane, as an energy carrier or raw material, has great potential for the chemical industry. Since methanation reaction is strongly exothermic and sensitive to diffusion, radial flow reactors represent a clear solution thanks to their low pressure drop and effective heat removal. A two-dimensional mixing cell network (MCN) approach to model the carbon dioxide methanation in a radial flow reactor is proposed. The reaction is catalyzed by a bifunctional Ni–Ce zeolite 13X supported catalyst, combining catalytic and adsorption functions. This contribution outlines the ongoing work, starting from a straightforward MCN pseudo-homogeneous approach comparing it with a tubular packed bed reactor. Both methanation kinetics and water adsorption have been successfully implemented in both models, setting feasibility for further improvements. Future developments will be necessary aiming to aid the design of units employing Ni–Ce/13X catalysts.

Keywords: Radial Flow Reactor, Methanation, Sorption-Enhanced, Mixing Cell Network, Carbon Dioxide Utilization

1. INTRODUCTION

The decarbonization of energy- and carbon-intensive sectors is pushing for the development of Power-to-Gas routes, in which renewable hydrogen is combined with carbon dioxide to produce synthetic natural gas (SNG). CO_2 methanation is an attractive pathway for the wide availability of catalytic materials and the production of grid-compatible fuel [1]. The reaction is reported below:

$$CO_2 + 4H_2 \rightleftarrows CH_4 + 2H_2O \quad (1)$$

$$\Delta_r H^0_{298K} = -165\ kJ/mol \quad (2)$$

As revealed, it occurs with a decrease in the number of moles and with the release of energy and is therefore thermodynamically favored by high pressures and low temperatures.

Many catalytic materials are available, differing from each other in the active phase, the support or the promoter [2]. Nonetheless, nickel is widely accepted to be fairly accessible balancing catalytic properties and application costs [3]. Furthermore, in recent research it was shown that using both physical mixture of catalyst and sorbent [4] and sorbent material as support for a catalyst [3] would improve the conversion of CO_2 by adsorbing water, exploiting Le Chatelier's principle. For these reasons, a bifunctional material made of Ni–Ce supported on zeolite 13X is implemented in the model discussed in this paper.

Different commercialized methanation technologies, like the TREMP Topsøe [5], are already common and employ packed bed reactors to produce SNG. However, as methanation is strongly exothermic an efficient heat removal is needed to avoid hot spot formation and to favor the thermodynamics of the reaction. Researchers suggested employing radial flow packed bed reactors, to enhance heat removal, to reduce the costs coming from using higher pressure to favor the thermodynamics [6]. In a radial reactor the catalyst is constrained between two perforated cylindrical metal screens enclosed in an outer shell which forms with the outer screen an annuls. The flow of the gas in a radial reactor can be centrifugal (CF) or centripetal (CP) and co-current (π-type) or counter-current (Z-type). There are four different

configurations of radial reactors, depending on how the flow is arranged, that are the CF flow Z-type, the CP flow Z-type, the CF flow π-type and the CP flow π-type [7]. In general, radial reactors are the favored solution when it comes to the scale up and when addressing the heat management and pressure drops [7].

Many strategies were adopted to model radial reactors, the majority of which were model based on the assumption that this type of reactor would behave similarly to a tubular packed bed reactor [8, 9]. However, in recent research, radial reactors were described as a matrix of interconnected CSTRs, each one representing a portion of the particle bed [10].

In the present paper the modelling of sorption enhanced methanation in a radial reactor is discussed. The radial reactor is modelled following a mixing cell network (MCN) approach, the kinetic is modelled for the bifunctional catalyst according to recent literature, accounting for the water adsorption. The results obtained from the radial reactor model will be compared to that of a tubular PBR.

2. MODELING

In this study the sorption enhanced methanation is modeled by developing and implementing a radial bed reactor adopting a pseudo-homogeneous mass balance. In the MCN the reactor is split into a matrix of interconnected tanks, each one representing the effect of a portion of the particulate bed. Figure 1 is a representation of the generic cell of the system. Figure 2 depicts the 2D MCN grid to model the radial reactor.

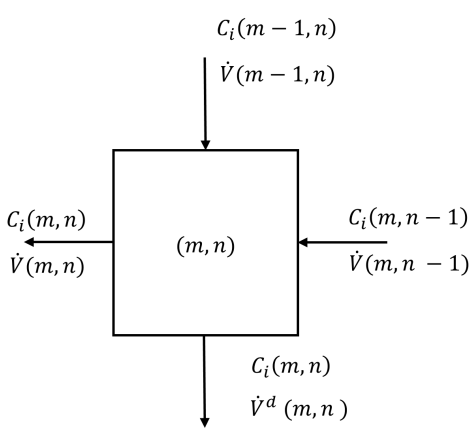

Figure 1. Schematic representation of the generic cell; m is the row index, n is the column index.

As a comparison the same reaction is also modelled in a packed bed reactor (PBR). The following assumptions were made:

- Ideal gas law applies.
- No side reactions are considered.
- Reactors are isothermal.
- A pseudo-homogenous mass balance was assumed.
- The model is written in dynamic form.
- Axial dispersion was assumed for PBR.
- Pressure drops for the radial reactor were neglected [11].
- A linear driving force (LDF) model is used to describe the adsorption rate [12].
- The reaction is assumed to be catalyzed by the bifunctional catalyst 5%Ni2.5%Ce13X [3].
- The outlet from each cell splits equally and the splitting is described by the ratio between the outlet volumetric flow rates from the cell.

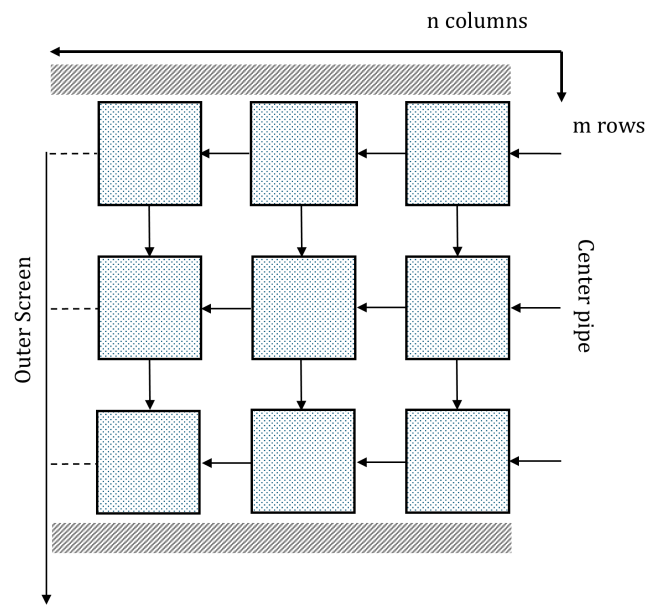

Figure 2. Schematic depiction of the 2D mixed cell network grid used for RFR.

Mass balance

For the MCN the mass balance equations are non-linear algebraic equations obtained by considering as a control volume the single cell. The general form of the dynamic mass balance in the gas phase for the generic (m, n) cell can be written as follows:

$$\overbrace{\varepsilon_B \frac{dC_i(m,n)}{dt} V_{cell}}^{Accumulation} =$$

$$\overbrace{C_i(m,n-1)\dot{V}(m,n-1) + C_i(m-1,n)\dot{V}(m-1,n)}^{Inlet} -$$

$$\underbrace{\left(C_i(m,n)\dot{V}(m,n) + C_i(m,n)\dot{V}^d(m,n)\right)}_{Outlet} + \underbrace{\rho_b \nu_i r_C(m,n)\, V_{cell}}_{Reaction} \quad (3)$$

Here, C_i is the concentration of the i-th component in the gas phase, ε_B is the bed void degree, V_{cell} is the volume of the single cell, V (dotted) is the volumetric flow rate, ρ_b is the catalyst bulk density, ν is the stoichiometric coefficient of the i-th component and r_c is the rate of the catalyzed reaction.

Kinetics and adsorption

For the methanation of CO_2 the kinetic was implemented grounding on previous studies [13, 14]. According to Koschany et al. for Ni-based catalyst, it was suggested that the rate of the catalyzed reaction was [14]:

$$r_c = \frac{k(C_{H_2} C_{CO_2})^{0.5}(R_{gas}\, T)\left(1 - \frac{C_{CH_4} C_{H_2O}^2}{C_{H_2}^4 C_{CO_2} K_{eq}*(R_{gas}\, T)^2}\right)}{\left(1 + (R_{gas}\, T)^{0.5}\left(K_{OH}\frac{C_{H_2O}}{C_{H_2}^{0.5}} + K_{H_2} C_{H_2}^{0.5} + K_{mix} C_{CO_2}^{0.5}\right)\right)^2} \quad (4)$$

Where k is the kinetic constant, R_{gas} is the ideal gas constant, T is the temperature, K_{eq} is the equilibrium constant, K_{OH}, K_{H2} and K_{mix} are adsorption constants. The kinetic constant is calculated from the modified Arrhenius law, written as follows:

$$k = k_0 exp\left(\frac{E_a}{R}\left(\frac{1}{T_{ref}} - \frac{1}{T}\right)\right) \quad (5)$$

The adsorption constants are computed as:

$$K_x = K_{ref,x} exp\left(\frac{\Delta H_x}{R}\left(\frac{1}{T_{ref}} - \frac{1}{T}\right)\right) \quad (6)$$

The equilibrium constant is calculated with the empirical formula:

$$K_{EQ} = 137\, T^{-3.994}\, exp\left(\frac{158.7\, kJ/mol}{R_{gas}\, T}\right) \quad (7)$$

The mass balance for the sorbent is described

adopting a linear driving force model [12]. Therefore, it can be written as follows:

$$\frac{\partial q}{\partial t} = K_{LDF}(q^* - q) \quad (8)$$

Where q is the amount of water adsorbed on zeolite 13X, q* is the equilibrium adsorption capacity of water on zeolite 13X and K_{LDF} is the linear driving force coefficient which takes into accounts internal mass transfer resistances during adsorption. The equilibrium adsorption capacity of water on the sorbent is defined as:

$$q^* = \frac{\rho(T)}{Mm_{H_2O}}\, W(T,P) \quad (9)$$

Where, Mm_{H2O} is the molecular mass of water and ρ(T) is the density of the adsorbed phase expressed as:

$$\rho(T) = \frac{\rho_{20°C}}{1+\beta_{20}(T-298.15K)} \quad (10)$$

Here, the parameter $\rho_{20°C}$ and β_{20} are the density and the thermal expansion coefficient of adsorbate at 20°C, respectively.

In eq. 9 W(T, P) is the volumetric adsorption capacity which is expressed according to previous studies [12]:

$$W = W_0\, exp\left[-\left(\frac{A}{E_{ads}}\right)^n\right] \quad (11)$$

$$A = \frac{R_{gas}}{Mm_{H_2O}}\, T\, ln\left(\frac{P_{sat}}{P_{H_2O}}\right) \quad (12)$$

In eq. 11 and 12 W is the adsorption volume, W_0 is the maximum adsorption volume, E_{ads} is the characteristic adsorption energy, A is the adsorption potential, P_{sat} is the saturation vapor pressure of water, P_{H2O} is the partial pressure of water and n is the heterogeneity factor of the micropore size distribution [12]. The linear driving force coefficient is computed as:

$$K_{LDF} = \frac{15}{0.5 d_p^2} D_{eff,H2O} \quad (13)$$

Where d_p is the particle diameter and $D_{eff,\, H2O}$ is the water effective diffusivity. Table 1 summarizes the geometrics of the reactor, the parameter set for the catalyst and the specification of the current fed to the reactor in the simulations. Table 2 gathers the values for all the parameters set for equation 4-13 which are kinetic, adsorption and thermodynamic data at reference temperature of 555K.

Table 1: Geometrics of the radial reactor, operation conditions and catalyst bed parameters.

Parameter	Value	Unit	Definition
R_o	0.02915	m	Outer radius
R_i	0.01	m	Inner radius
R_o- R_i	0.0192	m	Characteristic flow length
H	0.04	m	Bed height
V_{bed}	9.4213 10^{-5}	m^3	Bed volume
Operation conditions			
P	101325	Pa	Inlet pressure
T	523.15	K	Inlet temperature
yCO_2	0.04	-	Feed gas molar fractions
yH_2	0.16		
yCH_4	0.00		
yH_2O	0.00		
yHe	0.80		
V	100	mL/min	Vol. flow rate
R	0.5	-	Splitting ratio
Catalyst bed parameters			
ρ_b	500	kg/m^3	Solid bulk density
ε_B	0.40	-	Bed void degree

Table 2: Kinetic, adsorption and thermodynamic data at reference temperature of 555K.

Parameter	Value	Unit	Definition
Kinetic parameters			
k0	3.46 10^{-6}	mol/(Pa kg s)	Frequency factor
Ea	77.5 10^3	J/mol	Activation energy
Adsorption and thermodynamic parameters			
$K_{ref, OH}$	0.0016	Pa	Adsorption constant
$K_{ref, H2}$	0.0014		
$K_{ref, mix}$	0.0028		
ΔH_{OH}	22400	J/mol	Adsorption enthalpy
ΔH_{H2}	-6200		
ΔH_{mix}	-10000		
$\rho_{20°C}$	996	kg/m³	Density of adsorbate
β_{20}	2 10^{-3}	1/K	Thermal expansion coefficient
W_0	341 10^{-6}	m³/kg	Maximum adsorption volume
E_{ads}	1192.25 10^3	J/mol	Characteristic adsorption energy
d_p	2 10^{-4}	m	Particle diameter
$D_{eff, H2O}$	1 10^{-10}	m²/a	Effective diffusivity

RESULTS AND DISCUSSION

Radial reactor results

The simulations were carried out considering a 10x10 matrix of stirred tank reactors. The initial feeding was fed only to the first tank, and it subsequently split up in the following cells. The application of the mixed cell network led to obtain good first approximation results. Figure 3 reports the outlet gas molar fractions of the participant in the reaction. The breakthrough curve of water is depicted, showing the delayed outlet due to adsorption on zeolite 13X. Figure 4 shows the outlet molar flow rates, in particular the molar flow rate of carbon dioxide, hydrogen and methane variates, it increases for the reactants and decreases for the product. This happens at the water outlet, corresponding to the saturation of the adsorbent, and it is due to the shift of the reaction equilibrium to the reactants side. The variation is subtle due to the high equilibrium constant.

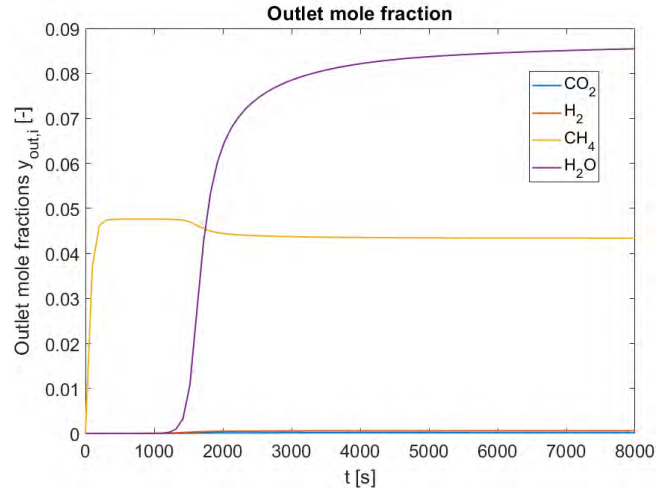

Figure 3. Radial flow reactor outlet gas molar fractions of the participant to the reaction versus time, the inert (helium) was excluded for reading pourpose.

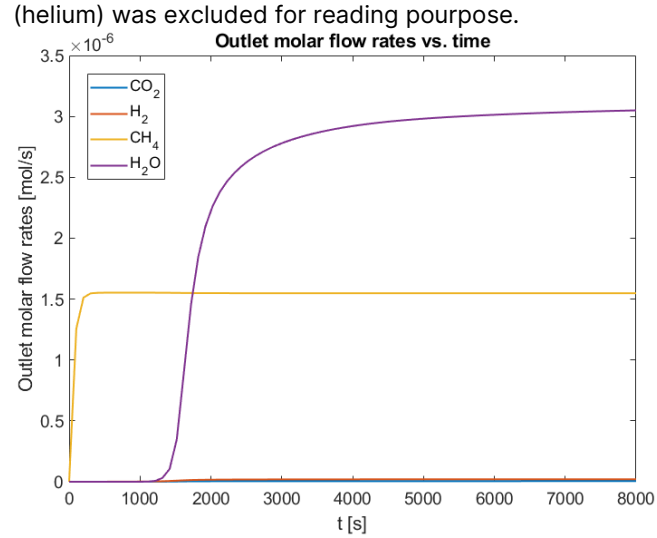

Figure 4. Radial flow reactor outlet molar flow rates of the participant to the reaction versus time.

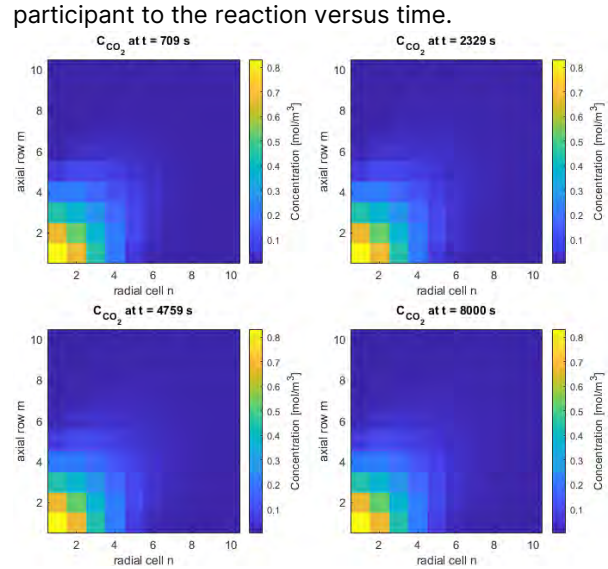

Figure 5. Carbon dioxide (limiting reactant) gas phase concentration [mol/m³] in each cell at fixed time instants.

Figure 5 shows the CO_2 gas phase concentration in each (m, n) cell at fixed time instants. Since the water is removed from the gas phase the equilibrium is shifted towards the products according to Le Chatelier's principle, therefore the reactants are consumed in the first part of the reactor. Figure 8 shows the development of the breakthrough curve of water. The water starts to appear in the reactor once the zeolite 13X has reached its maximum equilibrium capacity. When the sorbent can no longer retain water, the reaction is not enhanced anymore by adsorption.

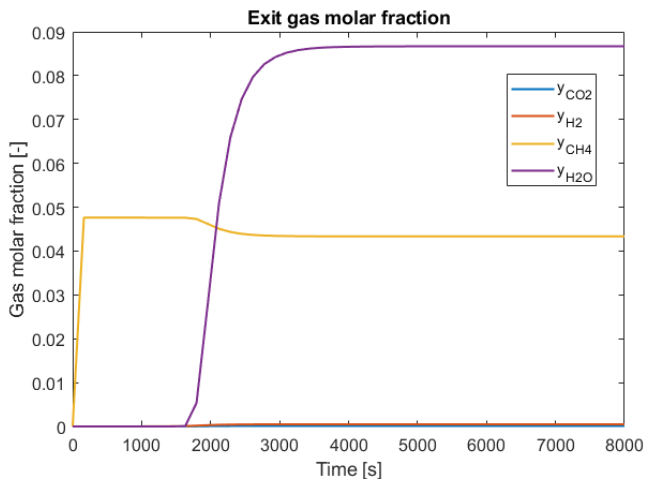

Figure 6. Tubular packed bed reactor outlet gas molar fractions versus time.

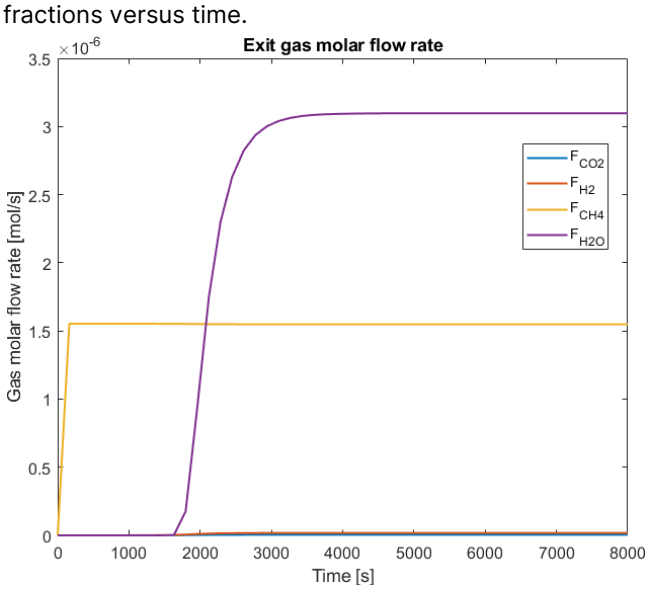

Figure 7. Tubular packed bed reactor outlet molar flow rates versus time.

Comparison with tubular PBR

To better inspect the behaviour of the MCN in a first stage the results can be compared to that of a tubular packed bed reactor of the same volume. This can be done because it is assumed that the reacting gas moves principally in the radial direction. The operation conditions and all the parameters used for the simulation of the radial reactor are the same for the tubular reactor. Figure 6 reports the outlet gas molar fractions for the tubular PBR.

It is possible to observe a slightly longer period of sorption enchantment, before zeolite 13X hits its maximum water uptake. In particular, the shorter period before the exit of water, in figure 3, in the radial reactor can be attributed to a distribution of residence time among the stirred cells due to the effect of the splitting parameter R. Figure 7 reports the outlet molar flow rates from the tubular reactor. Comparing the water breakthrough curves for both (fig. 3 and 6) reactors, it should be noted that in the region at longer times, the reaction should reach a new steady state, as the water is no longer being adsorbed by the zeolite 13X. This behaviour is correctly shown by the simulation in the tubular PBR, as can be seen in figure 6. Indeed, both the molar fraction of water and methane remain constant after the inflection point, where the sorbent has reached its maximum water uptake. However, in the radial reactor, deviations occur, and the steady state is delayed. Thus, the molar fraction of water and methane appear to increase and decrease, respectively (see figure 3). The delayed steady state is likely attributable to the onset of numerical diffusion due to the distribution of residence time among the mixed cells.

CRITICAL ASPECTS AND FUTURE DEVELOPMENT

Compared to other models available in literature for radial reactors, the mixed cell network avoids the widely spread approximation of assuming it to behave similarly to a tubular reactor. Therefore, the MCN has the potential to better describe a radial reactor with respect to classical plug flow models or axial dispersion models.

Pressure drops and velocity field handling are critical for the correct description of radial flow systems. This means combining the mass balance with a momentum balance and a continuity equation to describe the interaction of the flowing gas with the particle bed. In this study the fluid dynamic description was lumped in the parameter R which represents the splitting ratio between flow rates going out from each cell. Therefore, for future development the fluid dynamic will be addressed by solving the Ergun's equation for the radial coordinate and for the axial coordinate and combining it with continuity equation to obtain a velocity profile which can be coupled with the mass balance.

Another critical aspect is the appearance of internal diffusion limitations, which can be a bottleneck for the methanation reaction [13]. Therefore, an intraparticle mass balance can be implemented in the mixed cell network to account for this phenomenon.

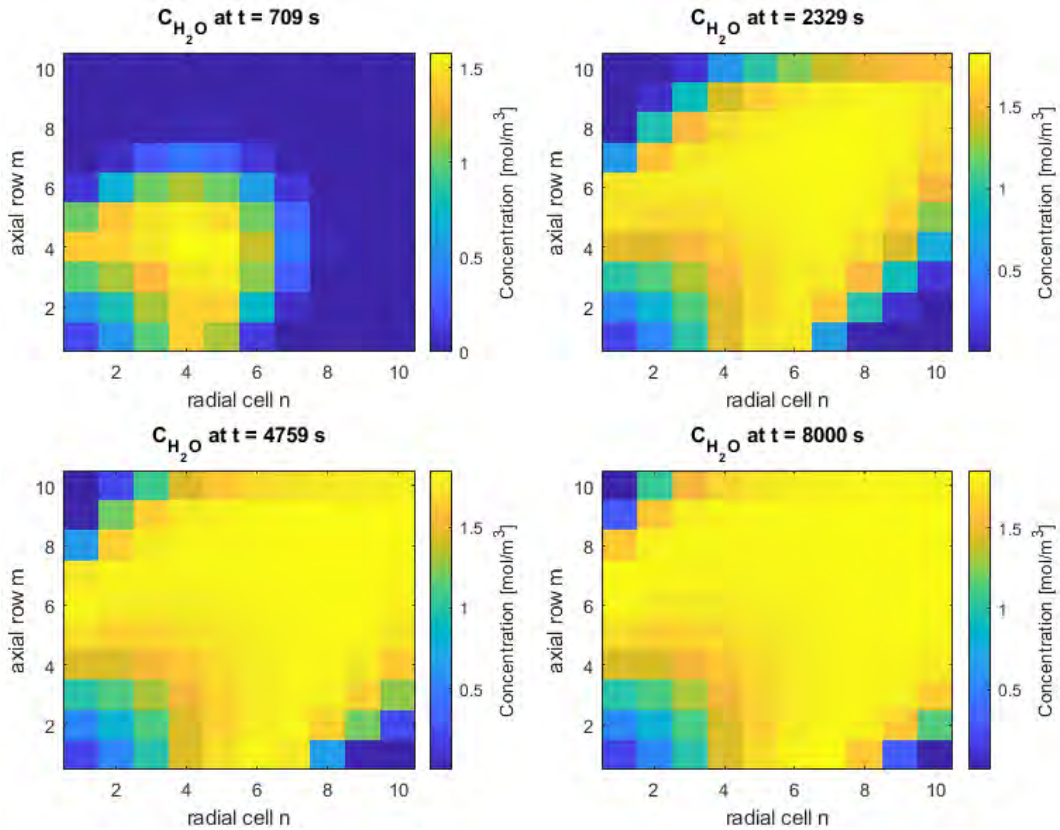

Figure 8. Water gas phase concentration [mol/m^3] in each cell at fixed time instants.

Methanation reaction is strongly exothermic, and this poses a design challenge when modelling radial reactors. An important future development is the implementation of the dynamic energy balance. Studying the transient of the temperature variation in the reactor is fundamental when trying to avoid hot spots formation as this can easily cause region of over-utilization of the bed and eventual runaways.

CONCLUSION

This study showed and discussed the modelling of the sorption enhanced methanation in a radial flow reactor system. The foundations of the MCN were presented. The mass balance on each cell and the implementation of the kinetics and adsorption was successfully, showing a good sorption enhancement period. The results fairly agree with that of a tubular PBR, and the difference are attributed to the handling of the fluid dynamics in the radial reactor. Critical aspects of the model and its future development have been highlighted showing good room for improvements implementing intraparticle balances, energy balance and pressure drop handling.

ACKNOWLEDGEMENTS

This work is a part of the activities of the Johan Gadolin Process Chemistry Centre at Åbo Akademi University. The authors gratefully acknowledge the financial support from the Research Council of Finland (Grant number 357454).

AUTHOR IDENTIFIERS

Author ORCIDs:
Salvatore Capasso: 0009-0009-2161-5665
Vincenzo Russo: 0000-0002-1867-739X
Henrik Grénman: 0000-0002-0129-0949

REFERENCES

1. Marchi E, Gangotena P, Frilund C, Salmi T, Simell P, Grénman H. Review on sorption-enhanced methanation of renewable hydrogen and carbon dioxide. International Journal of Hydrogen Energy 177:151628 (2025). https://doi.org/10.1016/j.ijhydene.2025.151628
2. Rönsch S, Schneider J, Matthischke S, Schlüter M, Götz M, Lefebvre J, Prabhakaran P, Bajohr S. Review on methanation – from fundamentals to current projects. Fuel 166:276-296 (2016). https://doi.org/10.1016/j.fuel.2015.10.111

3. Wei L, Azad H, Haije W, Grenman H, de Jong W. Pure methane from CO2 hydrogenation using a sorption enhanced process with catalyst/zeolite bifunctional materials. Applied Catalysis B: Environmental 297:120399 (2021). https://doi.org/10.1016/j.apcatb.2021.120399
4. Gómez L, Martínez I, Navarro MV, Murillo R. Selection and optimisation of a zeolite/catalyst mixture for sorption-enhanced CO2 methanation (SEM) process. Journal of CO2 Utilization 77:102611 (2023). https://doi.org/10.1016/j.jcou.2023.102611
5. Haldor Topsoe, 'From Solid Fuels to Substitue Natural Gas (SNG) using TREMP'. [Online]. Available: https://www.netl.doe.gov/sites/default/files/netl-file/tremp-2009.pdf
6. Zhang W, Lin Y, Norinaga K. Insights into structure-performance relationship in radial flow fixed bed reactor for CO2 methanation. International Journal of Hydrogen Energy 48:24594-24606 (2023). https://doi.org/10.1016/j.ijhydene.2023.03.343
7. J. C. H. Li, 'Radial-Flow Packed-Bed Reactors', Ullmanns Encycl. Ind. Chem. Wiley-VCH Verl. GmbH Co KGaA, 2012, doi: https://doi.org/10.1002/14356007.l22_l01
8. Farsi M. Mathematical modeling and optimization of a radial flow tubular reactor to produce methanol from syngas. Pet. Chem. 58:1091-1098 (2018). https://doi.org/10.1134/s0965544118100043
9. Iranshahi D, Karimi M, Amiri S, Jafari M, Rafiei R, Rahimpour MR. Modeling of naphtha reforming unit applying detailed description of kinetic in continuous catalytic regeneration process. Chemical Engineering Research and Design 92:1704-1727 (2014). https://doi.org/10.1016/j.cherd.2013.12.012
10. Yadav A, Roy S, Aijaz T. Modeling of three-phase radial flow reactor for diesel hydrotreating. Chemical Engineering Science 257:117713 (2022). https://doi.org/10.1016/j.ces.2022.117713
11. Hamedi N, Tohidian T, Rahimpour MR, Iranshahi D, Raeissi S. Conversion enhancement of heavy reformates into xylenes by optimal design of a novel radial flow packed bed reactor, applying a detailed kinetic model. Chemical Engineering Research and Design 95:317-336 (2015). https://doi.org/10.1016/j.cherd.2014.11.009
12. Mette B, Kerskes H, Drück H, Müller-Steinhagen H. Experimental and numerical investigations on the water vapor adsorption isotherms and kinetics of binderless zeolite 13X. International Journal of Heat and Mass Transfer 71:555-561 (2014). https://doi.org/10.1016/j.ijheatmasstransfer.2013.12.061
13. Mancusi E, Piso G, Shah HH, Pepe F, Tregambi C, Bareschino P. Modelling of a continuous sorption-enhanced methanation process in an adiabatic packed-bed reactor system. Chemical Engineering Science 301:120800 (2025). https://doi.org/10.1016/j.ces.2024.120800
14. Koschany F, Schlereth D, Hinrichsen O. On the kinetics of the methanation of carbon dioxide on coprecipitated nial(o). Applied Catalysis B: Environmental 181:504-516 (2016). https://doi.org/10.1016/j.apcatb.2015.07.026

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Dynamic Modeling of a Biomass Fluidized-Bed Gasifier

Jefferson D. C. Araujo[a]*, Fréderic Marias[a], and Sabine Sochard-Reneaume[a]

[a] University De Pau Et Des Pays De L'Adour, Thermal, Energy and Processes Laboratory, Pau, Pyrénées-Atlantiques, France
* Corresponding Author: jcdaraujo@univ-pau.fr

ABSTRACT

The climate crisis and dependence on fossil fuels make the transition to renewable energy sources imperative, with biomass standing out for promoting decarbonization and circular economy. In this context, fluidized bed gasification emerges as an efficient route for converting waste into syngas, applicable to power and hydrogen generation. Given the variability of real operating conditions, dynamic models are essential to represent coupled fluid dynamic, thermal, and kinetic phenomena over time. In this work, a dynamic phenomenological model was developed using a lumped 0D approach, in which the reactor is divided into two interacting zones represented as continuous stirred-tank reactors (CSTRs): a dense bed, where drying, devolatilization, and heterogeneous reactions occur, and a freeboard, dominated by homogeneous gas-phase reactions. The model was validated against experimental data from a bubbling fluidized bed gasifier, showing good agreement for major syngas species (CO and H_2, mean absolute error below 2%), with minor deviations for CH_4 due to simplified devolatilization and tar conversion kinetics. Transient simulations revealed strong sensitivity of process performance to particle size and biomass composition. Compared with the reference case (6 mm), reducing particle size to 1 mm enhanced hydrogen gas efficiency, cold gas efficiency, and carbon conversion due to intensified heterogeneous reactions and improved solid–gas heat transfer. Additionally, biomass switching simulations demonstrated that real-time air flow control effectively stabilizes syngas quality under feedstock variability. Overall, the model provides a computationally efficient tool for process optimization, control strategy development, and integration of biomass gasification into flexible and hybrid low-carbon energy systems.

Keywords: Gasification, Fluidized-bed, Dynamic modelling, Biomass, Syngas.

INTRODUCTION

The global climate crisis, intensified by persistent greenhouse gas (GHG) emissions, exposes the structural vulnerability of fossil fuel-based energy systems, whose geopolitical volatility and progressive depletion threaten global energy security. In this context, the Paris Agreement (2015) set ambitious targets to limit global warming to 1.5 °C, requiring a 45% reduction in emissions by 2030 and the achievement of carbon neutrality by 2050 [1].

Biomass emerges as a strategic vector in this transition, combining decarbonisation with circular economy principles by valorising agricultural, forestry and urban waste. The biomass is highly significant in the circular economy for both material products and energy provision, closing nutrient cycles and minimising waste [2].

According to the International Energy Agency's (IEA) Net Zero Emissions by 2050 Scenario, modern bioenergy is expected to triple its contribution, reaching approximately 100 EJ/year by 2050 (up from the current 60 EJ) to supply 15–20% of global energy demand [3].

Among thermochemical conversion routes, gasification stands out for its efficiency in producing syngas (CO + H_2), offering versatile applications in power generation, hydrogen production, and processes with the potential for negative emissions. Within this field, fluidized bed gasification is particularly notable due to its operational flexibility in processing heterogeneous biomass feedstocks. This reactor type provides efficient heat transfer and is well-suited for materials with compositional variability. Persistent challenges, though, including biomass moisture management and optimal operating pressure, require further development [4] .

In the case of real energy platforms, however, there

are load variations, fluctuations in fuel composition, and real-time control demands, requiring dynamic models that capture the coupled interactions between fluid dynamic, thermal, and kinetic phenomena over time [5]. While steady-state models are useful for point-in-time analysis, dynamic modeling is essential for design, operational optimization, and integration with hybrid systems. This approach is all the more necessary given the high complexity of thermochemical conversion processes, such as gasification, which involve non-linear parameters and high-dimensional conversion processes. Mathematical models therefore become indispensable tools for designing, optimizing, and intensifying these processes [6]. They make it possible to determine a system's optimal operating conditions, evaluate the impact of different variables on outcomes, and anticipate trade-offs between energy efficiency, thermal stability, and syngas quality. Furthermore, these models can be used for real-time control and optimization, as well as for predicting critical process performance parameters. These capabilities accelerate the technological maturity of processes toward industrial-scale deployment [7].

In this context, the present work aims to develop a numerical tool based on dynamic modeling capable of predicting the behavior of a biomass fluidized-bed gasifier when processing different types of biomass, each with distinct chemical composition, moisture content, and physical properties. The model seeks to capture, in a transient manner, the effects of these variations on the composition and yield of syngas, enabling the evaluation of their influence on process efficiency, operational stability, and the quality of the generated energy product.

METHODOLOGY

Model Description

A dynamic phenomenological model was developed to describe biomass gasification in a bubbling fluidized bed reactor, a schematic representation of the system and the boundaries of the domain studied is presented in the Figure 1. The objective of the model is to capture the main hydrodynamic, thermal, and chemical phenomena governing solid conversion and syngas composition under transient operating conditions.

To achieve this, a lumped-parameter approach (0D) was adopted, with the reactor divided into two distinct zones, both modeled as continuous stirred-tank reactors (CSTR). The reactor is divided into a dense bed zone and a freeboard zone, which are dynamically coupled through mass and energy exchanges.

The dense bed represents the fluidized region where gas and solid phases coexist. Biomass particles are introduced into an inert bed material and undergo drying, de-volatilization, char gasification, and partial oxidation re-actions. Gas–solid mass and heat transfer, as well as particle size reduction due to conversion, are explicitly considered in this zone.

The freeboard is modeled as a gas-phase region free of solid particles. Homogeneous gas-phase reactions and heat transfer with the reactor walls are considered. This zone contributes to the homogenization of the product gas composition before it exits the reactor.

The main modeling assumptions are summarized as follows: (i) Zone dense is a Multiphase CSTR (intra-phase homogeneity), (ii) Pressure is constant throughout each zone, (iii) Perfect gas behavior, (iv) Representation of biomass based on immediate and elemental analyses, (v) Layered model for biomass particles, (vi) Simplified thermal representation of walls and (vii) Particle shrinkage is considered, but average particle size is applied, while particle fragmentation is neglected.

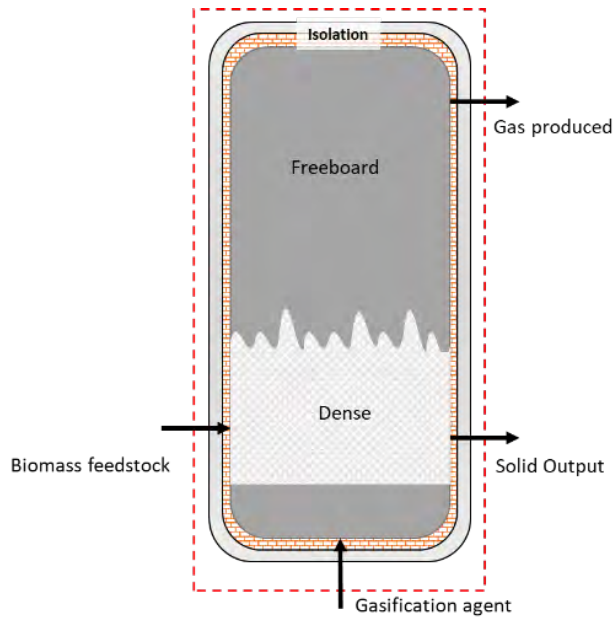

Figure 1. Schematic representation of the system – bubblubg fluidized bed reactor.

Phenomenological model construction

The proposed model considers the main stages of biomass conversion occurring in a bubbling fluidized bed gasifier. These stages include drying, devolatilization, heterogeneous reactions involving the solid char, and homogeneous gas-phase reactions. Each stage is modeled according to its dominant physical and chemical mechanisms and is dynamically coupled through mass and energy balances.

Drying

Drying describes the removal of moisture from biomass particles as their temperature increases. In the present model, moisture evaporation is driven by heat transfer between the gas and solid phases and occurs prior to devolatilization. The water evaporation rate is described

by an Arrhenius-type expression, dependent on the particle surface temperature and the amount of water contained in the particle, ensuring that the drying process progressively slows down as the available moisture is depleted [8].

Devolatilization

Devolatilization corresponds to the thermal decomposition of dry biomass, leading to the release of volatile species and the formation of a solid char residue. Devolatilisation is described using a global, lumped approach.

$$C_aH_bO_cN_dS_eCl_f \rightarrow \alpha C_{solide} + \sum \beta_i Gaz_i \quad (1)$$

For a given temperature, the devolatilisation products are represented by a limited number of gaseous and tar species, whose yields are determined by a set of empirical coefficients. Twelve coefficients are calculated: (i) one coefficient related to solid carbon, corresponding to the fixed carbon fraction, (ii) eight coefficients for gaseous species and (iii) three coefficients for tar compounds.

The system is closed using: (i) six atomic balance equations (C, H, O), (ii) three pyrolysis correlations from [9], [10] defining the ratios CO/CO_2, C_6H_6/C_6H_6O, and C_6H_6/C_7H_8, (iii) one gas-to-tar ratio correlation from [11] and (iv) one energy balance equation. This approach allows capturing the main trends of biomass devolatilisation while keeping the model computationally efficient and suitable for dynamic reactor simulations. Once the product distribution coefficients are established, the devolatilisation rate must be determined. This rate depends on the biomass type and is typically modeled using an Arrhenius-type kinetic expression [11].

Table 1: The kinetic reaction rates of reactions with carbon.

Reaction	k(m/s)
$C + H_2O \rightarrow CO + H_2$	$k_1 = 5.714 \bullet T \bullet \exp\left(\frac{-15600}{T}\right)$
$C + CO_2 \rightarrow 2CO$	$k_2 = 5.89 \bullet 10^2 \bullet T \bullet \exp\left(\frac{-26800}{T}\right)$
$C + \frac{1}{2}O_2 \rightarrow CO$	$k_3 = 2.30 \bullet T \bullet \exp\left(\frac{-11100}{T}\right)$
$C + O_2 \rightarrow CO_2$	$k_4 = k_3 \bullet \left(2512 \bullet \exp\left(\frac{-6420}{T}\right)\right)^{-1}$

Heterogeneous reactions

Heterogeneous reactions between char and gaseous species are modeled using a shrinking-core approach [12]. After devolatilization, the particle consists of a uniform reactive char layer surrounding an inert ash core. The entire external surface remains active throughout conversion, with ash contributing neither to reaction nor to internal diffusion resistance. Consequently, gaseous reactants encounter only gas-film resistance, and the particle diameter decreases progressively as char is consumed until only inert ash remains. All gasification and combustion reactions are assumed first-order in gaseous species and occur exclusively at the external char surface, with kinetic constants adopted from surface reaction data [13].

Homogeneous Reactions

Homogeneous gas-phase reactions play a central role in gasification by converting primary pyrolysis volatiles and heterogeneous reaction products into a more stable syngas composition. These reactions occur throughout the reactor, both in the dense bed and freeboard zones, and involve 16 key gaseous species, including major pyrolysis products (CO, CO_2, H_2, H_2O, CH_4), tar surrogates (benzene, toluene, phenol), light hydrocarbons (C_2H_2), particulate carbon, and oxidant gases (O_2, N_2). To balance accuracy and computational efficiency for dynamic reactor-scale simulation, a reduced kinetic scheme comprising 21 global reactions was adopted. The set captures the dominant chemical pathways: thermal cracking, steam and dry reforming, partial and complete oxidation, water–gas shift, and tar reforming/degradation [14-16]. Particular attention was given to tar conversion due to its dual energetic and environmental implications. The complete reaction list with stoichiometry and kinetic parameters is provided in the supplementary material.

Validation and Case Study

The predictive capability of the proposed model was evaluated by comparing simulation results with experimental data reported in the literature under similar operating conditions. Key process variables, including syngas composition, reactor temperature, and solid conversion, were analyzed. The operating parameters, such as gasifying agent composition, biomass characteristics, and reactor temperature, were adjusted to match the experimental setup.

Following validation, the model was applied as a predictive tool to investigate operational trade-offs associated with biomass gasification. The case study focused on two relevant transient operational scenarios: variation in particle diameter due to conversion and switching between different biomass feedstocks during reactor operation.

The impact of these parameters was evaluated in terms of syngas composition, reactor temperature evolution, and overall energy performance. To quantify process efficiency, several performance indicators were analyzed, including Hot Gas Efficiency (HGE), Cold Gas Efficiency (CGE), and Carbon Conversion Efficiency (CCE).

This combined validation and application approach highlights the capability of the model to support the analysis of transient operating strategies and to provide insights into the optimization of biomass gasification performance.

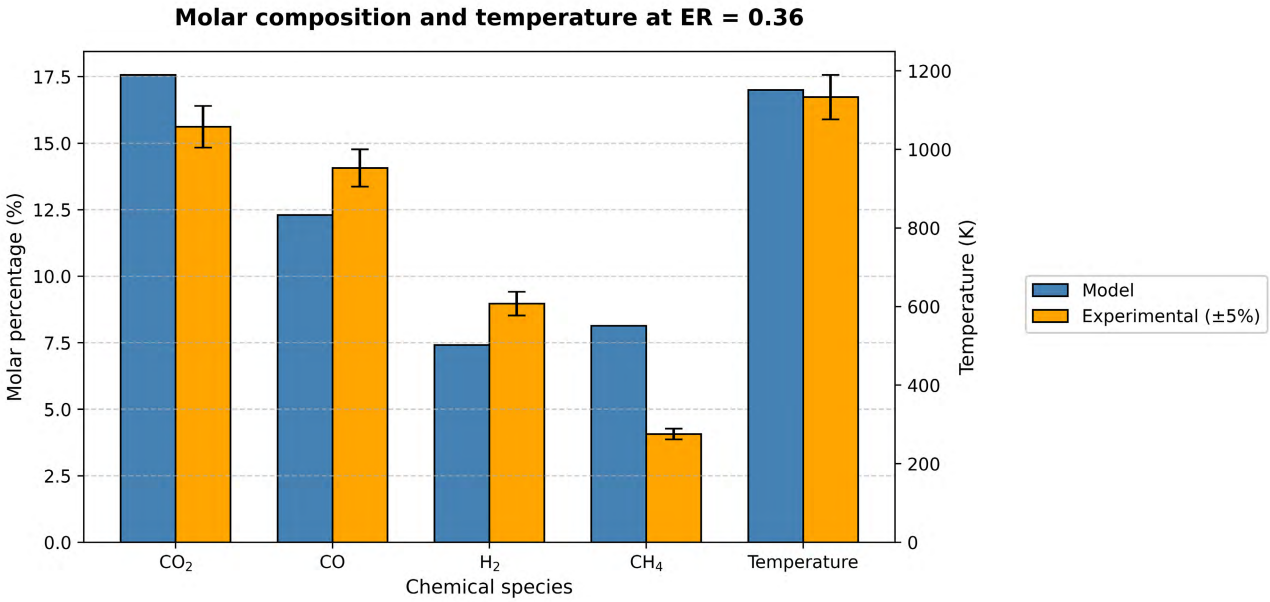

Figure 2. Comparison of Predicted and Experimental Gas Composition & Temperature (ER = 0.36)

RESULTS AND DISCUSSION

Validation

The model was validated against experimental data from a bubbling fluidized bed gasifier operating at 860 °C with Wood Pellets and air as the gasifying agent [17].Steady-state predictions of syngas composition (H_2, CO, CO_2, CH_4) and bed temperature were compared to measured values. As shown in Figure 2, the model reproduces the main trends of gas composition with a mean absolute error below 2% for H_2 and CO, the dominant species governing syngas heating value. A slight overprediction of CH_4 (≈4 % error) is observed, attributable to the simplified devolatilization and tar reforming kinetics. Overall, the results confirm the model's capability to predict the main features of biomass gasification steady-state conditions.

CASE STUDY

Reference case

Before the transient analysis, the model simulates a complete 48-hour startup sequence, from cold conditions (298 K) to steady-state operation, corresponding to the stabilization duration reported in [17]. The sequence begins with heating the bed to 380°C, followed by gradual feeding of biomass (wood pellets: d_p = 6.0 mm, moisture = 7.9% by weight, LHV = 16.95 MJ·kg^{-1}), followed by controlled adjustment of the air flow to achieve the target equivalence ratio (ER = 0.28). This controlled combustion phase drives the thermal increase of the reactor, favoring autothermal conditions, while air flow control will stabilize the composition of the synthesis gas. In 48 h, the system reaches steady-state conditions (CO = 5.03%, H_2 = 7.40%, dry basis; carbon conversion efficiency = 68.21%, HGE =59.68 % and CGE = 45.45%), with species variations below 1.0% in 1-hour windows (Figure 3).

The residence time was calculated based on the feed and withdrawal dynamics of the reactive solids in the reactor. In the reference scenario (base case), this parameter was set at 984.4 seconds, corresponding to approximately 16 minutes and 30 seconds. This validated baseline serves as the initial condition for the transient scenarios : (i) particle size (6.0 → 1.0 mm and 6.0 → 10.0 mm) and (ii) biomass switch (wood pellets →woodchips eucalyptus).

Building upon the validated steady-state baseline (48 h startup, d_p = 6.0 mm, residence time = 984.4 s), two transient scenarios were simulated to assess particle size sensitivity: (i) reduction to fine particles (d_p = 1.0 mm) and (ii) increase to coarse particles (d_p = 10.0 mm), as showed in Figure 4. To isolate the intrinsic effect of particle size on conversion kinetics, while maintaining equivalent hydraulic conditions, the solid outlet flow rate was dynamically recalculated at the instant of size change to preserve the reference residence time (984.4 s). This adjustment compensates for the altered bed inventory resulting from particle size variation, ensuring that observed transients reflect purely kinetic and heat/mass transfer effects rather than hydraulic residence time differences. The step-change in particle size was applied at t = 48 h, with biomass composition, flow rate and ER = 0.28 held

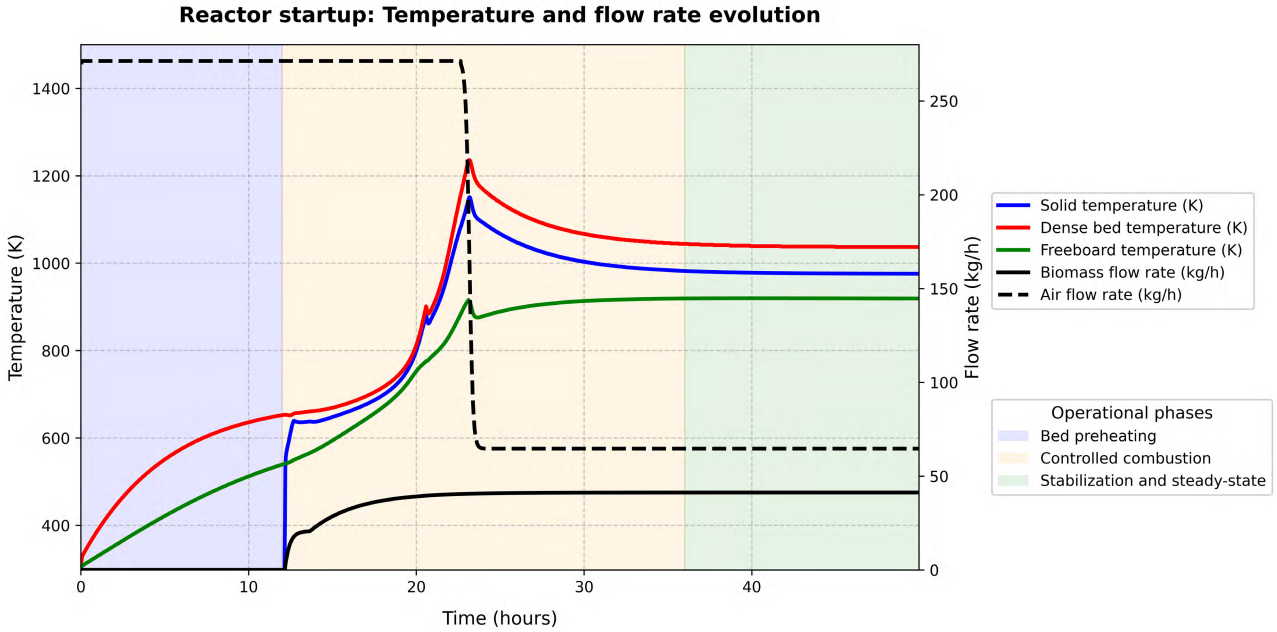

Figure 3. Reactor Startup : Temperature and flow rate evolution.

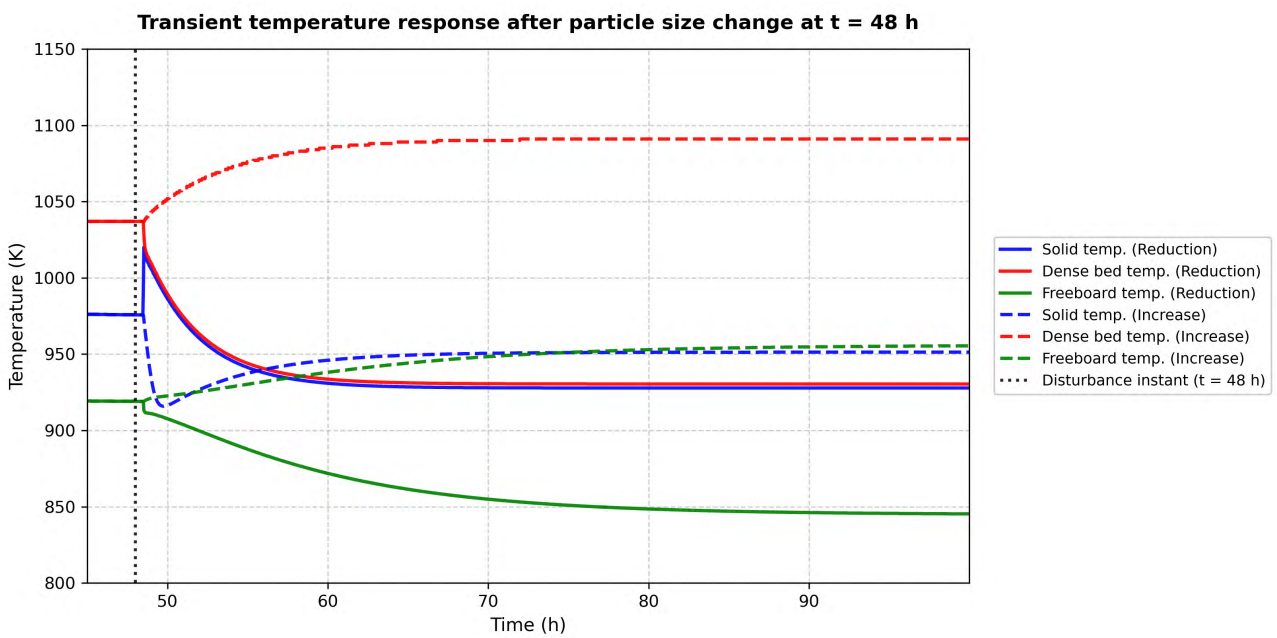

Figure 4. Transient temperature response after particle size change at t=48h.

constant throughout.

Particle size

Particle size strongly influences heat transfer, heterogeneous reaction rates, and overall gasification performance. The reference case (6 mm) represents a balance between reactive surface area and thermal behavior. Reducing particle size increases the specific surface area, enhancing solid–gas convective heat transfer and accelerating endothermic heterogeneous reactions such as steam and CO_2 gasification. Conversely, increasing particle size reduces the available reactive surface, limiting heterogeneous conversion and weakening thermal coupling between solid and gas phases.

Compared with the reference case, reducing particle size from 6 mm to 1 mm improves process performance, increasing hydrogen gas efficiency (HGE) from 59.68 % to 69.08 %, cold gas efficiency (CGE) from 45.45 % to 56.25 %, and carbon conversion efficiency (CCE) from 68.21 % to 77.01 %. This enhancement is associated

with intensified heterogeneous reactions (C + H_2O and C + CO_2), leading to higher char conversion and improved syngas quality, with CO and H_2 concentrations increasing from 5.03 % to 7.93 % and from 7.40 % to 11.24 %, respectively. In contrast, increasing particle size from 6 mm to 10 mm reduces HGE, CGE, and CCE to 54.54 %, 39.63 %, and 63.17 %, respectively, while decreasing CO and H_2 concentrations to 3.51 % and 5.24 %. These results confirm the dominant role of particle size in controlling thermal behavior and product distribution in fluidized bed gasifiers, highlighting the advantages of finer biomass particles, despite the additional milling energy requirement.

Biomass switch

To assess operational flexibility under feedstock variability, a transient scenario was simulated wherein the biomass type was switched from wood pellets to eucalyptus woodchips while maintaining a constant particle diameter (6 mm). The new biomass, characterized in [18], exhibits lower fixed carbon (16.21 wt.% vs. 20.21 wt.%) and same volatile content (79.81 wt.%), which inherently alters devolatilization intensity and char reactivity. To preserve syngas quality despite this compositional shift, the air flow rate was dynamically adjusted to maintain a constant CO/H_2 molar ratio, the targeted quality indicator, as showed in Figure 5.

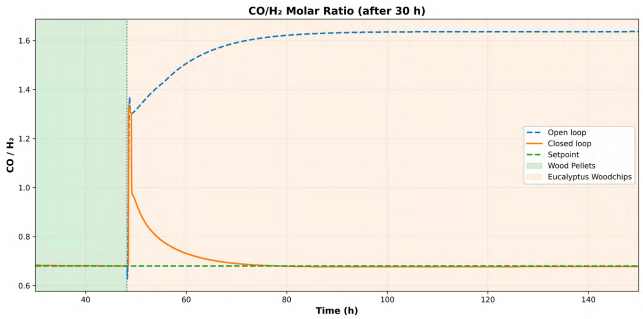

Figure 5. CO/H_2 molar ratio (after 30h).

Switching from wood pellets to eucalyptus woodchips necessitates dynamic air flow adjustment to maintain syngas quality, as demonstrated by the controlled case (Close loop), as showed in Figure 6. By modulating the air flow rate to stabilize the CO/H_2 molar ratio at 0.68, the system maintained consistent syngas composition (CO: 6%, H_2: 8.95%) and freeboard temperature (823 K) despite the feedstock change. In contrast, the uncontrolled case (Open loop) exhibited significant deviations, with the CO/H_2 ratio increasing to 1.6, CO concentration rising to 12%, H_2 to 7.33%, and freeboard temperature reaching 1057 K.

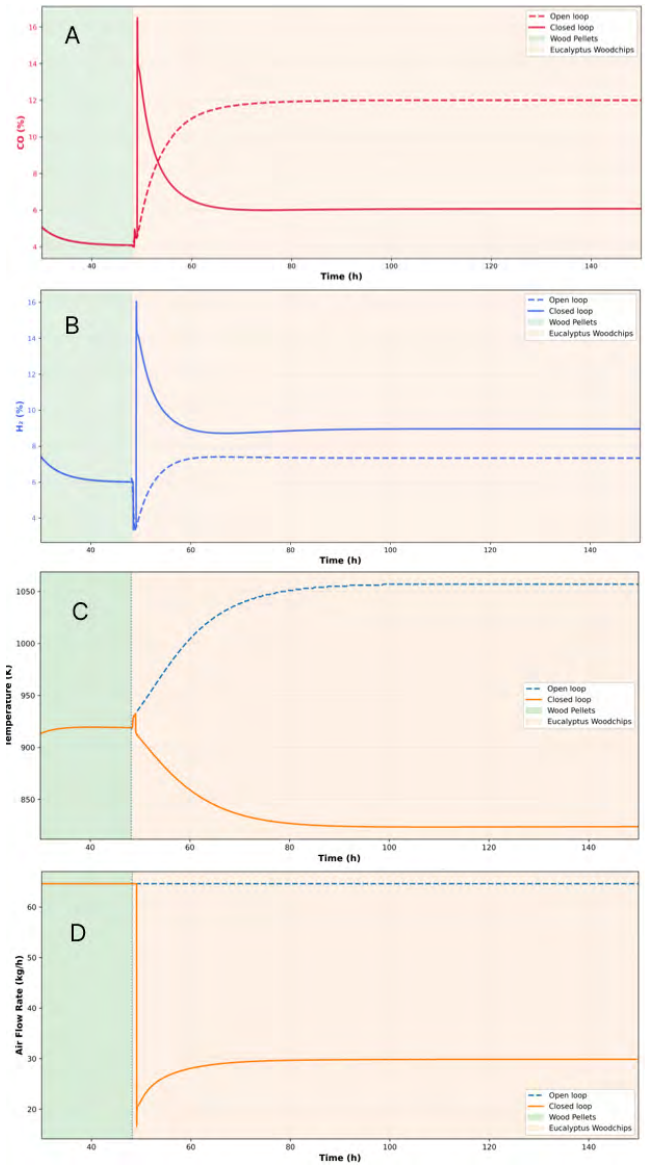

Figure 6. Dynamic response of syngas composition, thermal and flow rate behavior: (A) CO concentration, (B) H_2 concentration, (C) freeboard temperature and (D) air flow rate under open-loop and closed-loop operation during biomass transition.

This confirms that real-time air flow regulation effectively compensates for biomass-induced variations in reaction kinetics, ensuring operational stability during feedstock transitions. The comparison between Wood Pellets and Woodchips Eucalyptus highlights the strong influence of feedstock properties and process control on gasification performance. When Biomass 2 is operated under controlled conditions, the process shows improved gasification efficiency, with increases in HGC (from 59.7% to 63.61%) and CGC (from 45.46% to 55.37%) while maintaining a stable CO/H_2 ratio (0.68). These improvements are achieved with a significantly lower air flow rate, indicating a more efficient and optimized

conversion process. Conversely, when Biomass 2 is operated without control, the system exhibits substantially higher performance indicators (HGC = 85.43%, CGC = 67.44%, CCE = 99.06%), but this is accompanied by a strong increase in the CO/H_2 ratio (1.63) and higher air consumption. This suggests that although uncontrolled operation maximizes carbon conversion, it compromises gas quality and process selectivity, demonstrating the critical role of control strategies in balancing conversion efficiency, gas composition, and energy optimization.

CONCLUSION

This work presents a dynamic model for biomass gasification in bubbling fluidized bed reactors that accurately predicts both steady-state and transient behavior by integrating hydrodynamic, thermal, and kinetic phenomena within a unified framework. Validated against experimental data, the model successfully captures key aspects of biomass conversion, syngas composition, and thermal dynamics, while also simulating realistic operational disturbances such as particle size changes and biomass switching. Results highlight that feedstock properties significantly affect conversion efficiency and syngas quality, but real-time air flow control can effectively compensate for biomass variability, enabling stable operation with heterogeneous feedstocks, a common challenge in industrial settings.

The model's dynamic nature offers a strategic advantage for integrating gasification into hybrid, multi-product energy systems, as it enables analysis of load variations, start-up/shutdown sequences, and transient disturbances critical for coupling with downstream processes like methanation, hydrogen production, and Fischer–Tropsch synthesis. By predicting time-dependent syngas composition, it enhances compatibility with catalytic units, supports digital twin development, and facilitates advanced process control for improved efficiency and reduced environmental impact. Providing a computationally efficient yet physically representative tool, this work advances the design of flexible bioenergy platforms aligned with low-carbon energy transitions, with future extensions envisioned in areal-time control integration for industrial-scale deployment.

DIGITAL SUPPLEMENTARY MATERIAL

A supplementary material is available in LAPSE:2026.0024 with the complete reaction list with stoichiometry and kinetic parameters

ACKNOWLEDGEMENTS

This study was funded by the Communauté de communes de Lacq-Orthez and the Institut Carnot ISIFOR

REFERENCES

1. U. Nations, “Net Zero Coalition, ” United Nations. Accessed: Jan. 04, 2026. [Online]. Available: https://www.un.org/en/climatechange/net-zero-coalition
2. Sherwood J. The significance of biomass in a circular economy. Bioresource Technology 300:122755 (2020). https://doi.org/10.1016/j.biortech.2020.122755
3. “WS30: Bioenergy in a Net Zero Future | Bioenergy.” Accessed: Jan. 06, 2026. [Online]. Available: https://www.ieabioenergy.com/blog/ieaevent/ws30-bioenergy-in-a-net-zero-future/
4. Motta IL, Miranda NT, Maciel Filho R, Wolf Maciel MR. Biomass gasification in fluidized beds: a review of biomass moisture content and operating pressure effects. Renewable and Sustainable Energy Reviews 94:998-1023 (2018). https://doi.org/10.1016/j.rser.2018.06.042
5. Boujjat H, Yuki Junior GM, Rodat S, Abanades S. Dynamic simulation and control of solar biomass gasification for hydrogen-rich syngas production during allothermal and hybrid solar/autothermal operation. International Journal of Hydrogen Energy 45:25827-25837 (2020). https://doi.org/10.1016/j.ijhydene.2020.01.072
6. Hussain M, Ali O, Raza N, Zabiri H, Ahmed A, Ali I. Recent advances in dynamic modeling and control studies of biomass gasification for production of hydrogen rich syngas. RSC Adv. 13:23796-23811 (2023). https://doi.org/10.1039/d3ra01219k
7. Rey JRC, Longo A, Rijo B, Pedrero CM, Tarelho LAC, Brito PSD, Nobre C. A review of cleaning technologies for biomass-derived syngas. Fuel 377:132776 (2024). https://doi.org/10.1016/j.fuel.2024.132776
8. Benkoussas B, Consalvi JL, Porterie B, Sardoy N, Loraud JC. Modelling thermal degradation of woody fuel particles. International Journal of Thermal Sciences 46:319-327 (2007). https://doi.org/10.1016/j.ijthermalsci.2006.06.016
9. Dufour A, Masson E, Girods P, Rogaume Y, Zoulalian A. Evolution of aromatic tar composition in relation to methane and ethylene from biomass pyrolysis-gasification. Energy Fuels 25:4182-4189 (2011). https://doi.org/10.1021/ef200846g
10. Abdelouahed L, Authier O, Mauviel G, Corriou JP, Verdier G, Dufour A. Detailed modeling of biomass gasification in dual fluidized bed reactors under aspen plus. Energy Fuels 26:3840-3855 (2012). https://doi.org/10.1021/ef300411k
11. Van de Velden M, Baeyens J, Boukis I. Modeling CFB biomass pyrolysis reactors. Biomass and

Bioenergy 32:128-139 (2008). https://doi.org/10.1016/j.biombioe.2007.08.001
12. M. De Souza-Santos, Solid Fuels Combustion and Gasification: Modeling, Simulation, and Equipment Operations. 2004. doi: 10.1201/9780203027295.
13. Wu Y, Zhang Q, Yang W, Blasiak W. Two-dimensional computational fluid dynamics simulation of biomass gasification in a downdraft fixed-bed gasifier with highly preheated air and steam. Energy Fuels 27:3274-3282 (2013). https://doi.org/10.1021/ef4003704
14. R. Demarthon, “Modélisation et simulation d’un étage haute température pour la purification d’un gaz chargé en goudrons et en particules carbonées par assistance plasma, ” PhD Thesis, 2013. [Online]. Available: http://www.theses.fr/2013PAUU3001/document
15. C. Oliveira Martins, “Reformage de gaz sidérurgiques par assistance plasma, ” PhD Thesis, 2022. [Online]. Available: http://www.theses.fr/2022PAUU3010/document
16. A. Fourcault, “Modélisation d’un étage haute température alimenté par une torche à plasma pour la dégradation des goudrons présents dans les gaz de synthèse, ” PhD Thesis, 2010. [Online]. Available: http://www.theses.fr/2010PAUU3015
17. Badu PO, Debal M, Girods P, Aubert S, Rogaume Y. Biomass gasification in an autothermal semi-industrial fluidized bed gasifier: syngas characterization and energy balance. Biomass and Bioenergy 197:107831 (2025). https://doi.org/10.1016/j.biombioe.2025.107831
18. Parrillo F, Ardolino F, Calì G, Marotto D, Pettinau A, Arena U. Fluidized bed gasification of eucalyptus chips: axial profiles of syngas composition in a pilot scale reactor. Energy 219:119604 (2021). https://doi.org/10.1016/j.energy.2020.119604

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Comprehensive Framework for Model Discovery and Discrimination Based on Symbolic Regression and Structural Identifiability - Application to a Partially Observed Chemical Reaction System

Xuming Yuan[a], Brahim Benyahia[a]

[a] Loughborough University, Department of Chemical Engineering, Loughborough, Leicestershire, UK
* Corresponding Author: B.Benyahia@lboro.ac.uk.

ABSTRACT

Traditional approaches for mechanistic modelling require in-depth understanding of the underlying chemical and physical phenomena to construct reliable and predictive models. However, at early stages of development, limited experimental data, incomplete expert knowledge, and non-observable states often hinder a full understanding of the underlying mechanisms. Symbolic regression (SR) enables systematic model discovery and offers a practical route to addressing these challenges by automating the identification of interpretable model structures and the estimation of associated parameters from available data. However, structural identifiability and observability (SIO), a critical property of such models, is often overlooked in SR, thereby limiting its broader adoption and effective deployment. To address these limitations, this study proposes a comprehensive framework, which leverages scarce prior knowledge in SR and incorporates SIO analysis, offering a potential solution to capture the effects of all state variables and ensure rigorous structure of the resulting models. The proposed strategy is demonstrated on a partially observed sulfide oxidation system, demonstrating how an SIO-assured model can be systematically identified from limited data and incomplete system knowledge. Overall, this work presents a potential pathway for extending model discovery to partially observed systems and enhancing model robustness, thereby positioning SR as an alternative tool with respect to traditional kinetic modeling strategies.

Keywords: Modelling and Simulations, Symbolic Regression, Partially Observed Systems, Structural Identifiability & Observability Analysis, Systematic Model Development

1. INTRODUCTION

The availability of robust mathematical models and digital twins constitutes a cornerstone of digital transformation across multiple industries [1, 2]. Mathematical models are essential for both off-line and real-time decision-making and underpin process design, optimization, scale-up, control, operation, and integration [3, 4]. The recent emergence of digitally driven paradigms for process and product quality and environmental sustainability, such as Green-by-Design and Quality-by-Digital-Design frameworks [5-7], has further reinforced the central role of mathematical models, particularly within the pharmaceutical industry. These recent shifts coincide with the broader adoption of continuous manufacturing technologies, collectively driving increased demand for high-quality mathematical models that support design space identification and exploration, start-up and shut down, and control strategies [8, 9]

Traditionally, the development of a mathematical model starts by postulating the inputs and outputs, then constructing the model equations which involve the unknown model parameters to be calibrated using experimental data. Often the underlying phenomena to be modelled are intricate, subject to limited measurements, non-observables, and expert bias. These challenges result in multiple model candidates which may require further benchmarking and discrimination procedures. However, the identification of affective model through the traditional route may demand extensive prior knowledge of

the system and domain expertise, which can be scarce especially for early-stage studies. More importantly, the optimality of the chosen kinetic expression is probably questionable since discrimination is normally restricted to a handful of manually crafted alternatives, while an exhaustive search remains prohibitively labor-intensive. On the contrary, data-driven models (e.g., response surface models, artificial neural networks) dispense with mechanistic assumptions and exhibit superior adaptability to the data supplied, offering accurate predictions within the training space. Nonetheless, they are known to be data-hungry and commonly difficult to extrapolate. In addition, they are operated as opaque black boxes with poor interpretability, which potentially hinders the comprehension of the underlying natural phenomenon.

The emergence of automated model discovery approaches is possibly delivering a way out of this impasse. Through automating the exploration and identification of the most appropriate model over a large combinatorial search space, the demand of human data-science expertise is minimized, thereby dramatically reduces the effort for the determination of the optimal set of model equations. In particular, symbolic regression (SR), a process that discovers parsimonious mathematical expressions from datasets, is a group of the most attractive approaches in automated model discovery. SR refers to the practice not only estimating the numerical coefficients of a pre-defined mathematical expression but also exploring the most appropriate functional forms describing the data [14], which fuses applied mathematics and machine learning to provide deep insights of the underlying relationship between the input, output and state variables with limited prerequisites of the model structure. Therefore, it is greatly beneficial to unveil the unknown governing equations of the systems that manifest complex or nonlinear behaviors. Hence, SR is rapidly becoming a catalyst for advanced model development. For instance, two automated model discovery frameworks (ADoK-S and ADoK-W) were developed for chemical reactions to automatically yield process models utilizing a Python-based deep SR toolbox PySR [11]. SR is also applied to discover the nucleation and growth expressions for a batch crystallization case of potassium sulphate, yielding highly resembling equations as the ones of standard method of moments, confirming the interpretability of the model and the effectiveness of the methodology [10]. Another SR implementation on crystallization process [15] leverages SINDy to obtain the nucleation and growth governing equations, which fitted well with the 1st and 2nd moments data but were not very satisfactory for 0th and 3rd moments, probably due to the invalidity of the inherent assumption of linear combination of the terms in SINDy. In a biological context, the kinetic model of a microbial fermentation was derived through implementing SR, which was then embedded in a multi-objective optimization to maximize cell biomass, protein titre and minimize batch time [10]. Experimental validation showed excellent agreement with model predictions, demonstrating the practical utility of SR in real-world bioprocesses.

SR is an extraordinary tool to uncover concise, predictive equations, yet most studies stop at goodness-of-fit and ignore structural identifiability and observability (SIO). Structural identifiability asks whether the parameters can be uniquely recovered from the observables, while observability determines whether unmeasured states can be uniquely inferred from the observables. Neglecting either property incurs ill-posed models that fit the training data yet yield unreliable predictions [17].

To bridge this gap, we couple PySR with the MATLAB-based toolbox STRIKE-GOLDD and its extension AutoRepar [16]. PySR is selected as the algorithm for model generation in this study due to greater flexibility in the functional forms and better compatibility to unmeasured states compared with SINDy. After PySR proposes a parsimonious model, STRIKE-GOLDD investigates it SIO, and the equations will be reparametrized by AutoRepar if any structural defect is pinpointed. The loop repeats until a structurally sound, data-efficient model is returned.

The resulting workflow furnishes an end-to-end pipeline from raw data to trustworthy equations, requiring minimal prior mechanistic insight and human intervention. We demonstrate its value on a partially observed sulfide-oxidation system [12], which starts from sparse measurements and recovers a compact, identifiable and observable high-fidelity model. By embedding SIO analysis inside SR, the approach offers a rigorous new paradigm for early-stage process development where data are scarce and mechanistic knowledge is limited.

2. METHODOLOGY

The proposed holistic framework which integrates the automated model discovery and SIO analysis for chemical reaction systems as shown in Figure 1. Initially, PySR is used to perform SR based on the measured data of the observables ($\boldsymbol{y}$), generating a large set of plausible concentration-time equations $\boldsymbol{y}(t)$. The $\boldsymbol{y}(t)$ with the lowest Akaike information criterion (AIC) is chosen, approximating the concentration evolution of the observed species. Thereafter, the concentration-time expressions for the entire set of state variables $\boldsymbol{x}(t)$ are derived from $\boldsymbol{y}(t)$ applying mass balance and stoichiometry, which allows the numerical computation of derivative values $\dot{\boldsymbol{x}}(t)$. Next, the evaluated $\dot{\boldsymbol{x}}(t)$ values are leveraged in couple with the measured data of $\boldsymbol{y}$ in the SR exploring $\dot{\boldsymbol{y}}(\boldsymbol{x})$ expressions using PySR. It is worth highlighting that $\dot{\boldsymbol{y}}(\boldsymbol{x})$ is the kinetic equations set describing the chemical reaction system. Afterwards, an AIC test is performed to find the top N

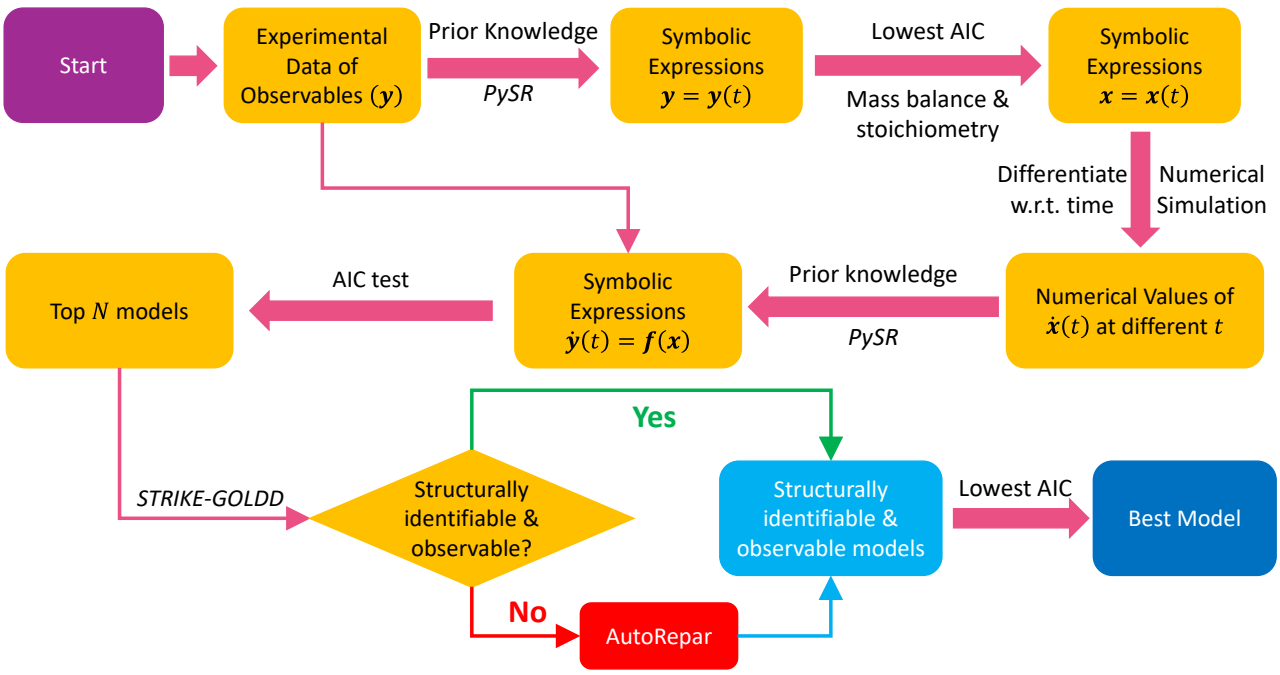

Figure 1. Proposed framework which integrates the automated model discovery and SIO analysis for chemical reaction systems.

models best fitting the data. The SIO of these models are scrutinized using STRIKE-GOLDD, while the models with structural defects will be reparametrized by AutoRepar. Eventually, the SIO-assured model with the lowest AIC is selected as the best model. This pipeline delivers models that both fit the data and exhibit rigorous structure, minimizing the requirement of prior mechanistic knowledge and domain expertise, thereby reducing human input to problem specification and data upload.

3. CASE STUDY

A sulfide oxidation reaction system [12] is selected as the case study, which involves four species: raw material methyl phenyl sulfide (R), desired product methyl phenyl sulfoxide (D), side-product methyl phenyl sulfone (U) and reactant hydrogen peroxide ($\mathrm{H_2O_2}$). The main reaction and the side reaction are presented in Equations 1 and 2, respectively. An experimental study was conducted in [12], providing the concentration-time data of R, D and U, as depicted in Figure 2, while the initial concentration of $\mathrm{H_2O_2}$ is 2.5 times as the initial concentration of R. The process involves four state variables $\boldsymbol{x} = [\mathrm{R}, \mathrm{D}, \mathrm{U}, \mathrm{H_2O_2}]$ and three observables $\boldsymbol{y} = [\mathrm{R}, \mathrm{D}, \mathrm{U}]$. The goal is to find the underlying rate equations of $\mathrm{R}, \mathrm{D}, \mathrm{U}$ and $\mathrm{H_2O_2}$ based on the data provided.

$$\mathrm{R} + \mathrm{H_2O_2} \rightarrow \mathrm{D} + \mathrm{H_2O} \tag{1}$$

$$\mathrm{D} + \mathrm{H_2O_2} \rightarrow \mathrm{U} + \mathrm{H_2O} \tag{2}$$

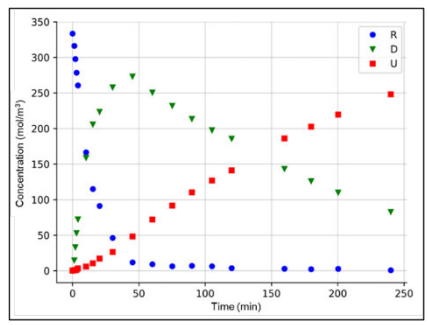

Figure 2. Experimental concentration-time data of methyl phenyl sulfide oxidation [12]. R: Methyl phenyl sulfide. D: Methyl phenyl sulfoxide. U: Methyl phenyl sulfone.

3.1. Obtain $\boldsymbol{y}(t), \boldsymbol{x}(t)$ and $\dot{\boldsymbol{x}}(t)$

The first step in the workflow is to obtain a concentration-time model for the observables denoted by $\boldsymbol{y}(t)$. Before running SR, it should be noticed that the atomic balance of sulfur can be given by Equation 3, where $\mathrm{R_{ini}}$ the initial concentration of R. Since $\mathrm{R_{ini}}$ is a known constant, U can be treated as a dependent variable, thereby it is only necessary to symbolically regress $\mathrm{R}(t)$ and $\mathrm{D}(t)$. The PySR settings in this step are summarized in Table 1.

$$\mathrm{R}(t) + \mathrm{D}(t) + \mathrm{U}(t) = \mathrm{R_{ini}} \tag{3}$$

By running SR, a series of mathematical expressions of $\mathrm{R}(t)$ and $\mathrm{D}(t)$ are obtained. In this case, 19 candidates for $\mathrm{R}(t)$ and 17 candidates for $\mathrm{D}(t)$ are identified by PySR, among which the expressions corresponding to the

Table 1: PySR settings for the discovery of concentration-time relationships

Setting	Value	Description
niterations	100	Maximum number of iterations in symbolic regression
binary_operators	["+", "-", "*", "/"]	Allowed computational operations between two terms
unary_operators	["exp"]	Allowed computational operations on one term
loss	loss(x, y) = (x - y)^2	Fitting objective function: mean squared error
population_size	100	Population size for seeking the mathematical expression
maxsize	30	Complexity limit of the mathematical expression
timeout_in_seconds	60	Time limit of symbolic regression search
parsimony	10^{-3}	Parsimony coefficient determining the sparsity of the mathematical expression

Complexity	Loss	Equation
1	88.72346	-1.5131524
3	0.65134954	-0.063507065 * R
5	0.64904344	(-0.06375289 * R) + 0.06231785
7	0.42687345	R * (-0.061005794 + (-5.4023287e-5 * D))
9	0.14379378	(-0.057983927 + ((D * R) * -5.498032e-7)) * R
11	0.08519997	(-0.04215632 + (R * ((D * -8.229583e-7) + -4.6016456e-5))) * R
13	0.0780262	((-0.057921443 - ((R * 8.5146423e-10) * (D * H2O2))) * R) - -0.19413921
15	0.07224689	(R * (-0.057804443 - (((R - U) * 8.545725e-10) * (D * H2O2)))) - -0.13526249
17	0.055178992	(27.877422 + R) * (-0.05317548 + (((D * (-76.895355 + R)) - (H2O2 * U)) * -8.234853e-7))
19	0.055041444	(27.877422 + R) * (-0.05316254 + ((((D * (R + -76.895355)) - (H2O2 * U)) + U) * -8.2493364e-7))

(a)

Complexity	Loss	Equation
1	78.04503	1.3129557
3	1.8908302	R * 0.058173776
5	1.2114038	-1.0696173 + (R * 0.062388875)
7	1.2113937	((R * 1.0623888) - 1.0695953) - R
9	0.5922489	(R * ((D * 9.781079e-5) + 0.059605878)) - 1.5127325
11	0.11398397	((0.05509739 + (D * (8.210762e-7 * R))) * R) - 1.3124856
13	0.0565364	((0.060073934 + ((8.721774e-7 * R) * D)) * R) - (H2O2 * 0.0036195528)
15	0.0357263	(R * (0.05330363 + ((3.2462417e-6 * ((R * 0.2775878) - U)) * D))) - 0.83941805
17	0.03287169	((0.05330363 + ((3.7689076e-6 * ((R * 0.22400005) - (U + -3.4678288))) * D)) * R) - 0.8074039
19	0.031657755	(R * (0.05332232 + ((4.102189e-6 * ((0.194399 * R) - (U + (-2.7117753 - 2.7252011)))) * D))) - 0.780967

(b)

Figure 3. Symbolically regressed reaction rate expressions of: (a) Methyl phenyl sulphide (R), (b) Methyl phenyl sulfoxide (D).

lowest AIC values are selected as the best concentration-time equations (Equations 4 and 5 below).

$$\mathrm{R}(t) = \frac{\exp(8.678)}{t+17.577+0.0732(\exp(0.192t)+t^2)} + \exp(1.469) \tag{4}$$

$$\mathrm{D}(t) = \exp\left(\exp\left(0.645 + \frac{t+114.992}{t+\frac{212.903}{t+1.676}+95.674}\right)\right) - 0.559t - 24.232 \tag{5}$$

With $\mathrm{R}(t)$ and $\mathrm{D}(t)$ determined, the mathematical expression of $\mathrm{U}(t)$ can be readily calculated using Equation 3. To derive $\mathrm{H_2O_2}(t)$, the mass balance can be implemented, deducing that the consumption of $\mathrm{H_2O_2}$ is equal to the sum of the consumption of R and the generation of U, resulting in Equation 6. As the initial concentrations of R ($\mathrm{R_{ini}}$) and $\mathrm{H_2O_2}$ ($\mathrm{H_2O_{2ini}}$) are known, Equation 6 enables the calculation of $\mathrm{H_2O_2}$ value at any given time.

$$\mathrm{H_2O_2}(t) = \mathrm{H_2O_{2ini}} - \left(\mathrm{R_{ini}} - \mathrm{R}(t)\right) - \mathrm{U}(t) \tag{6}$$

Therefore, $\boldsymbol{x}(t) = [\mathrm{R}(t), \mathrm{D}(t), \mathrm{U}(t), \mathrm{H_2O_2}(t)]$ is obtained, which allows the numerical computation of the $\dot{\boldsymbol{x}}(t)$ values at the sampling points, approximating the reaction rates

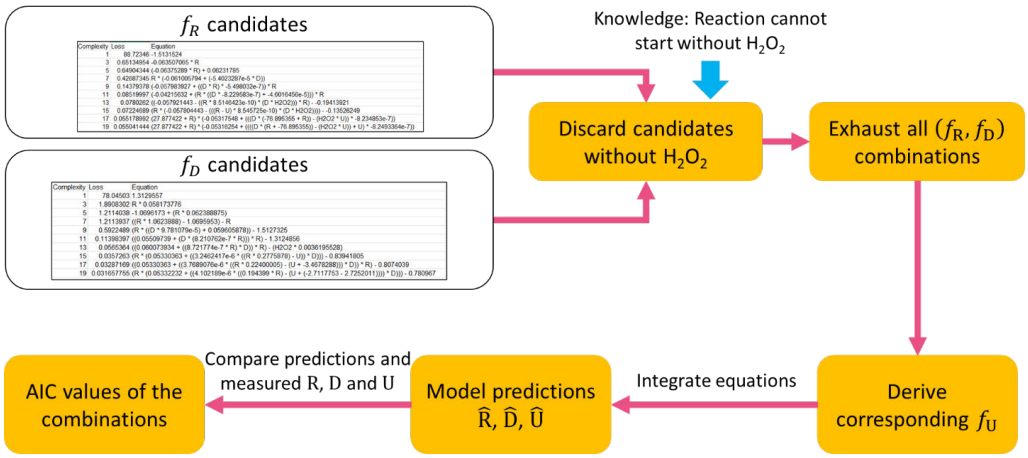

Figure 4. Procedure to construct valid rate model candidates and calculate the AIC values.

for each species in the experiment. The measured of $\mathrm{R}, \mathrm{D}, \mathrm{U}$ data, the simulated $\mathrm{H_2O_2}$ values, together with the simulated $\dot{\boldsymbol{x}}(t)$ values at the sampling points, are subsequently used in the next step discovering the $\dot{\boldsymbol{y}}(\boldsymbol{x})$ mathematical expressions.

3.2. Derive $\dot{\boldsymbol{y}}(\boldsymbol{x})$

As emphasized in section 3.1, U and $\mathrm{H_2O_2}$ are dependent on R and D, thus it is only necessary to attain the mathematical expressions $\dot{\mathrm{R}} = f_{\mathrm{R}}(\mathrm{R}, \mathrm{D}, \mathrm{U}, \mathrm{H_2O_2})$ and $\dot{\mathrm{D}} = f_{\mathrm{D}}(\mathrm{R}, \mathrm{D}, \mathrm{U}, \mathrm{H_2O_2})$ via SR, which are the reaction rate equations sought in this study. Compared with the settings for seeking the concentration-time relationships (Table 1), the settings for the SR of the reaction rate equations (f_{R} and f_{D}) apply a deeper search, which raises the maximum number of iterations from 100 to 150 and the population size from 100 to 200. Furthermore, the division operator and the exponential operators are disabled, as they are rarely seen in rate laws. Moreover, the complexity cap is tightened from 30 to 20, reflecting the parsimony of natural kinetics.

Figure 3 presents the resulting f_{R} and f_{D} candidates given by SR. Some mathematical expressions do not contain $\mathrm{H_2O_2}$ in the terms, violating the knowledge that the reaction cannot start without its existence. Thus, the expressions without $\mathrm{H_2O_2}$ are discarded.

Thereafter, an exhaustive mapping is carried out to construct all possible $(f_{\mathrm{R}}, f_{\mathrm{D}})$ combinations from expressions that pass the screening and derive the corresponding f_{U}. The resulting model candidates, consisting ordinary differential equations, are integrated with respect to time, giving the predicted $\hat{\mathrm{R}}, \hat{\mathrm{D}}$ and $\hat{\mathrm{U}}$ at the sampling times. Afterwards, the AIC values are calculated by leveraging the measured data of R, D and U. This procedure is graphically presented in Figure 4.

3.3. SIO Analysis & Best Model Selection

Utilizing the MATLAB-based tool STRIKE-GOLDD, the SIO of the 10 model candidates with the lowest AIC values are analyzed, where the models with structural defects will be reparametrized by AutoRepar to fix this issue. Note that the AIC of the reparametrized models are re-computed, since the number of model parameters have changed. Eventually, the best kinetic model is selected among these SIO-assured models based on the updated AIC values. The rate equations of the best model are given by Equations 7 to 10, with the parameter values presented in Table 2. The SIO of this model is guaranteed, as evidenced by the SIO analysis records presented in Figure 5. Moreover, this model has an outstanding prediction performance as demonstrated in Figure 6. Hence, a high-fidelity and structurally rigorous reaction rate model is obtained in this case study.

$$\dot{\mathrm{R}} = p_1 + (p_2 - p_3 \bullet \mathrm{R} \bullet \mathrm{D} \bullet \mathrm{H_2O_2}) \bullet \mathrm{R} \tag{7}$$

$$\dot{\mathrm{D}} = (p_4 + p_5 \bullet \mathrm{R} \bullet \mathrm{D}) \bullet \mathrm{R} - p_6 \bullet \mathrm{H_2O_2} \tag{8}$$

$$\dot{\mathrm{U}} = -(p_1 + (p_2 - p_3 \bullet \mathrm{R} \bullet \mathrm{D} \bullet \mathrm{H_2O_2}) \bullet \mathrm{R}) - ((p_4 + p_5 \bullet \mathrm{R} \bullet \mathrm{D}) \bullet \mathrm{R} - p_6 \bullet \mathrm{H_2O_2}) \tag{9}$$

$$\dot{\mathrm{H_2O_2}} = 2(p_1 + (p_2 - p_3 \bullet \mathrm{R} \bullet \mathrm{D} \bullet \mathrm{H_2O_2}) \bullet \mathrm{R}) + ((p_4 + p_5 \bullet \mathrm{R} \bullet \mathrm{D}) \bullet \mathrm{R} - p_6 \bullet \mathrm{H_2O_2}) \tag{10}$$

Table 2: Parameter estimates of the best rate model

Parameter	Value	Unit
p_1	0.194	$\mathrm{mol \bullet m^{-3} \bullet min^{-1}}$
p_2	-5.79×10^{-2}	$\mathrm{min^{-1}}$
p_3	-8.51×10^{-10}	$\mathrm{mol^{-3} \bullet m^{-9} \bullet min^{-1}}$
p_4	6.01×10^{-2}	$\mathrm{min^{-1}}$
p_5	8.72×10^{-7}	$\mathrm{mol^{-2} \bullet m^{-6} \bullet min^{-1}}$
p_6	-3.62×10^{-3}	$\mathrm{min^{-1}}$

4. CONCLUSION

An integrated workflow coupling SR with joint SIO analysis is proposed and demonstrated through a partially observed sulfide oxidation case study. By embedding scarce prior knowledge and limited experimental

```
--------------------------------
>>> STRIKE-GOLDD toolbox 3.0
--------------------------------

Analyzing the TempCombo_k0_0_162105 model...

>>> The model contains:
4 states:
[R; D; U; H2O2]
3 outputs:
[R; D; U]
0 known inputs:

0 unknown inputs:

6 parameters:
[p1; p2; p3; p4; p5; p6]

>>> Building the observability-identifiability matrix requires at least 3 Lie derivatives
    Calculating derivatives: 1 2 3
>>> Observability-Identifiability matrix built with 3 Lie derivatives
    (calculated in 2.425598e-01 seconds)
>>> Calculating rank...
    Rank = 9 (calculated in 6.448850e-02 seconds)
>>> Observability-Identifiability matrix built with 4 Lie derivatives
    (calculated in 3.031976e-01 seconds)
>>> Calculating rank...
    Rank = 10 (calculated in 1.766829e-01 seconds)

------------------------
>>> RESULTS SUMMARY:
------------------------

>>> The model is Fully Input-State-Parameter Observable (FISPO):
    All its states are observable.
    All its parameters are locally structurally identifiable.
Total execution time: 1.485672e+00


----------------------------------------
The model is already FISPO.
 There is no need for reparameterization.
----------------------------------------
```

Figure 5. SIO analysis record of the best model given by STRIKE-GOLDD.

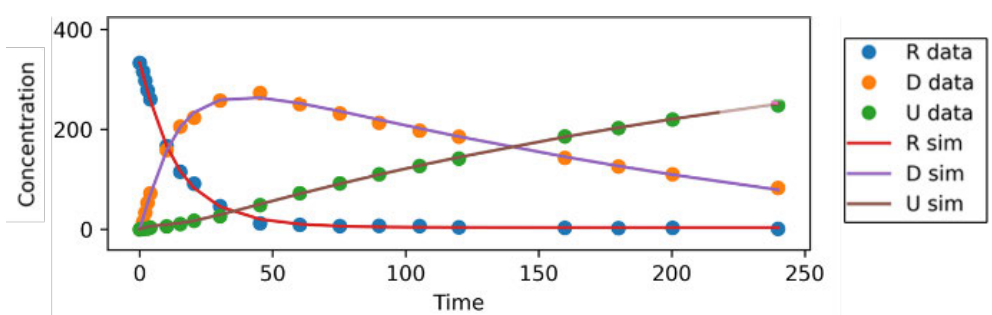

Figure 6. Fitting performance of the best model.

data, the approach returns a model that simultaneously fits the data and preserves rigorous structure. In contrast to purely data-driven SR pipelines, which often yield non-identifiable or structurally inconsistent models, the proposed framework systematically filters candidates based on identifiability and observability, leading to a substantially reduced and more interpretable model space. Relative to conventional mechanistic model

development, the workflow achieves comparable structural guarantees while requiring markedly less a priori specification, thereby improving robustness and credibility in early-stage process modelling where data and mechanistic insight are both limited.

ACKNOWLEDGEMENTS

The authors acknowledge funding from the UK Engineering and Physical Sciences Research Council (EPSRC), for Made Smarter Innovation – Digital Medicines Manufacturing Research Centre (DM2), EP/V062077/1.

AUTHOR IDENTIFIERS

Author ORCIDs:
Yuan X: 0009-0006-2211-4879
Benyahia B: 0000-0002-5397-1498

REFERENCES

1. Ramin, P, Mansouri, SS, Udugama, IA, Benyahia, B, Gernaey, KV (2018) Modelling continuous pharmaceutical and bio-based processes at plant-wide level: A roadmap towards efficient decision-making, *Chemistry Today*, 36(2), 26-30.
2. Ma Y, Yuan X, Yewale A, Benyahia B. Recent advances in the digital transformation of crystallization process development and operation: synergy between model- and ai-driven strategies. Particuology 110:223-253 (2026). https://doi.org/10.1016/j.partic.2026.01.014
3. Campbell, TJS, Rielly, CD, Benyahia, B (2024) Sustainability and Quality-by-Digital Design of an Integrated End-to-End Continuous Pharmaceutical Process. *Computer Aided Chemical Engineering*, 53, 343-348
4. Benyahia, B. (2018) Applications of a plant-wide dynamic model of an integrated continuous pharmaceutical plant: Design of the recycle in the case of multiple impurities. In Singh, R and Yuan, Z (ed) *Process Systems Engineering for Pharmaceutical Manufacturing: From Product Design to Enterprise-Wide Decisions*, Elsevier, pp.141-157, DOI: 10.1016/B978-0-444-63963-9.00006-3.
5. Mustoe CL, Turner AJ, Urwin SJ, Houson I, Feilden H, Markl D, Al Qaraghuli MM, Chong MWS, Robertson M, Nordon A, Johnston BF, Brown CJ, Robertson J, Adjiman C, Batchelor H, Benyahia B, Bresciani M, Burcham CL, Cardona J, Cottini C, Dunn AS, Fradet D, Halbert GW, Henson M, Hidber P, Langston M, Lee YS, Li W, Mantanus J, McGinty J, Mehta B, Naz T, Ottoboni S, Prasad E, Quist PO, Reynolds GK, Rielly C, Rowland M, Schlindwein W, Schroeder SLM, Sefcik J, Settanni E, Siddique H, Smith K, Smith R, Srai JS, Thorat AA, Vassileiou A, Florence AJ. Quality by digital design to accelerate sustainable medicines development. International Journal of Pharmaceutics 681:125625 (2025). https://doi.org/10.1016/j.ijpharm.2025.125625
6. Campbell TJS, Rielly CD, Benyahia B. Digital design and optimization of an integrated reaction–extraction–crystallization–filtration continuous pharmaceutical process. Computer Aided Chemical Engineering 51:775–780 (2022) https://doi.org/10.1016/B978-0-323-95879-0.50130-2
7. Gao S, Benyahia B. Robust techno-economic analysis, life cycle assessment, and quality and sustainability by digital design of three alternative continuous pharmaceutical tablet manufacturing processes. Systems and Control Transactions 4:2568-2573 (2025). https://doi.org/10.69997/sct.104102
8. Liu, J., Benyahia, B. Systematic model-based dynamic optimization of a combined cooling and antisolvent multistage continuous crystallization process. Computer Aided Chemical Engineering 50, 1221-1227.
9. Benyahia B, Anandan PD, Rielly C. Robust model-based reinforcement learning control of a batch crystallization process. 2021 9th International Conference on Systems and Control (ICSC) :89-94 (2021). https://doi.org/10.1109/icsc50472.2021.9666494
10. Alam MN, Bhat S, Kodamana H, Rathore AS. A data-driven symbolic regression framework for modeling and multiobjective optimization of a microbial fermentation system. Ind. Eng. Chem. Res. 64:23220-23233 (2025). https://doi.org/10.1021/acs.iecr.5c02567
11. de Carvalho Servia MÁ, Sandoval IO, Hii KK(, Hellgardt K, Zhang D, Antonio del Rio Chanona E. The automated discovery of kinetic rate models – methodological frameworks. Digital Discovery 3:954-968 (2024). https://doi.org/10.1039/d3dd00212h
12. Doherty S, Knight JG, Carroll MA, Clemmet AR, Ellison JR, Backhouse T, Holmes N, Thompson LA, Bourne RA. Efficient and selective oxidation of sulfides in batch and continuous flow using styrene-based polymer immobilised ionic liquid phase supported peroxotungstates. RSC Adv. 6:73118-73131 (2016). https://doi.org/10.1039/c6ra11157b
13. Javadi AB, Pong P. A review on symbolic regression in power systems: methods, applications, and future directions. Renewable and Sustainable Energy Reviews 224:116075 (2025). https://doi.org/10.1016/j.rser.2025.116075
14. Koza J. Genetic programming as a means for

programming computers by natural selection. Stat Comput 4: (1994). https://doi.org/10.1007/bf00175355
15. Lima FARD, de Moraes MGF, Rebello CM, Barreto AG Jr, Secchi AR, de Souza MB Jr, Nogueira IBR. Interpretable and uncertainty-aware machine learning for trustworthy prediction in batch crystallization. Chemical Engineering and Processing - Process Intensification 215:110350 (2025). https://doi.org/10.1016/j.cep.2025.110350
16. Massonis G, Banga JR, Villaverde AF. Autorepar: a method to obtain identifiable and observable reparameterizations of dynamic models with mechanistic insights. Intl J Robust & Nonlinear 33:5039-5057 (2021). https://doi.org/10.1002/rnc.5887
17. Yuan X, Nagy ZK, Benyahia B. A comprehensive framework for model evaluation and refinement using mbdoe estimability and structural identifiability: application to a crystallization process. Chemical Engineering Journal Advances 23:100823 (2025). https://doi.org/10.1016/j.ceja.2025.100823

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

SEMPRE-BIO project: comparison of three innovative scaled up and optimised technologies for biomethane production and its purification

Filippo Bisotti[a]*, Matteo Gilardi[a]*, and Bernd Wittgens[a]*

[a] SINTEF Industry – Process Technology – P.O. Box 4760 Torgarden, 7465, Trondheim, Norway
* Corresponding Authors: filippo.bisotti@sintef.no, matteo.gilardi@sintef.no, bernd.wittgens@sintef.no

ABSTRACT

This work presents the scale-up and detailed analysis, including comparison of relevant key performance indicators (KPIs) and energy analysis, of three innovative technologies for producing biomethane and its subsequent upgrading. The SEMPRE-BIO project tested and validated three different innovative technologies and pilots within the Horizon Europe framework. The three pilots are in Spain (biomethanation of biogas from wastewater treatment fermentation), France (biomethanation of syngas from biomass gasification), and Belgium (purification of biogas from manure anaerobic digestion). The three case studies will be presented and discussed. The analysis will dive into the layout of the optimised process layout of the scale-up plants as well as an exhaustive comparison, presenting advantages, shortcomings and bottlenecks of each technology, accounting for the main KPIs, i.e., electricity and steam demand, consumables including cooling water and others specific to the technologies.

Keywords: Biogas, Biomethanation, Biomethane, Gasification, Waste valorisation, Carbon circularity, Energy transition

INTRODUCTION

Biomethane (bio-CH_4) is gaining attention for its application as fuel and feedstock in chemical synthesis [1-4]. Currently, replacing natural gas with bio-CH_4 is part of the EU's decarbonisation strategy. The implementation of a large and integrated biogas network should increase the EU's energy independence and resilience to market shocks and shortages [5]. Biogas is produced via anaerobic fermentation. It contains mainly wet bio-CH_4 and biogenic CO_2 (bio-CO_2), but, in minor content, also pollutants, such as NH_3 and H_2S [5, 6]. As for natural gas, biogas needs purification to meet specifications for use and/or transport. The purification of biogas, i.e., dehydration, cleaning, and removal of bio-CO_2, is the most straightforward strategy [7]. Direct biomethanation of biogas [8] or (bio)syngas [9] is becoming appealing to boost bio-CH_4 yield and site productivity, but challenges are to be fixed [10]. SEMPRE-BIO project, funded under the EU's Horizon Europe program, aims to secure domestic production of cost-effective bio-CH_4. For this reason, pilots for three technologies have been run at different locations using different feedstocks. The general scope is increasing the European portfolio, adding more options for bio-CH_4 production and unlocking the biogas economy. The proposed technologies covered either biomethanation with downstream cleaning or biogas purification. For all case studies, the focus was on validating the concepts and debottlenecking their deployment. This work presents the three case studies (CSs), the results from process simulations, and compares the technologies regarding to the most relevant key performance indicators (KPIs).

MATERIALS AND METHODS

Case studies presentation

This section presents the case studies (CSs) and the corresponding processes. For each demo site, the block flow diagram (BFD) and model assumptions are reported. All process simulations have been performed in COCO COFE V3.10 using both conventional units and custom-based solutions.

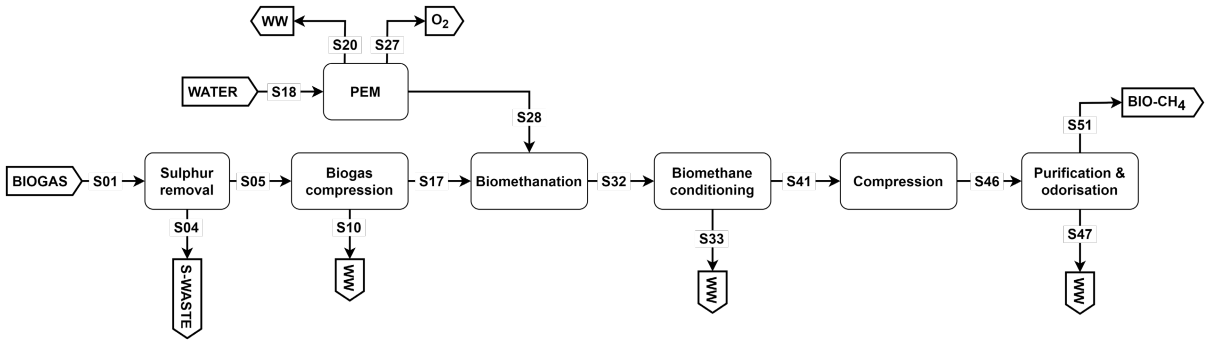

Figure 1. CS1 (Spain) - BFD

Case Study 1 (CS1) - Spain

Process description

Figure 1 depicts the BFD for the demo site in Spain to validate the direct biomethanation of cleaned biogas. The produced bio-CH_4 is expected to meet quality for fuelling public transport vehicles. The raw biogas comes from the fermentation of wastewater treatment sludge. The resulting raw biogas (S01) is desulfurised and the cleaned biogas (S05) compressed to 7 bar(a) and partially dehydrated (S17). The compressed biogas is sent to the biomethanation reactor, where H_2 from a proton-exchange membrane (PEM) module (S28) is fed and consumed. The relative ratio between H_2 and CO_2 is sensitive. The resulting gas mixture (S32), rich in bio-CH_4, is post-treated and conditioned to reduce its moisture content below the target value. Bio-CH_4 (S41) is compressed to the desired pressure based on storage and final user requirements. For this case, we considered 250 bar(a). The compression further refines the water content in the final product (S51).

Model assumptions

For the simulation of this case study, the following assumptions have been adopted:

- Redlich-Kwong-Soave (RKS) cubic equation of state has been used to model both the liquid and gas phases.
- The desulphurisation unit has been designed as a splitter with assigned efficiency, whose value has been assigned based on the average value from literature and providers [6, 11].
- Compressors and pumps have been simulated with an assigned total efficiency of 75% (accounting for both mechanical and polytropic factors).
- The PEM electrolyser has been modelled as a simplified unit where cathode and anode reactions occur in a reactor at a fixed conversion of water. The conversion is a function of the operating temperature of the electrolyser. In addition, reaction efficiency has been considered based on the suggested value from the provider (ProPuls). This approach was already adopted in other studies [12, 13].
- Biomethanation reactor has been simulated as a reactor with fixed conversion for the limiting reactant (bio-CO_2). The average CO_2 conversion for thermophilic bacteria consortia has been retrieved from experimental evidence and the literature [14].

Case Study 2 (CS2) - France

Process description

Figure 2 shows the BFD for the demo site in France, for the biomethanation of biosyngas from woody biomass gasification. The chips (S01) are fed to the gasification. Steam is not supplied as the solid presents a high content of moisture, and steam is produced during the heating. The gasification chamber is run at 700-750°C under autothermal conditions. The solid (char) is discharged from the chamber, while the hot raw gas leaves the gasifier (S15). The sensible heat of S15 is recovered to sustain the gasification. The recovered heat is insufficient as gasification is highly endothermic. The remaining energy demand is covered with the combustion of a fraction of the produced biosyngas (S31). The raw biosyngas (S20) is conveyed to a cleaning unit to refine the tars and pollutants content. The cleaned syngas (S22) is split into two streams. One fraction (S31) is pre-mixed with air (S32) and burnt. The remaining (S23) is fed to a biomethanation reactor. Conversely to CS1 (biogas as feed), bio syngas contains H_2. Hence, two alternatives for the biomethanation have been considered: with co-feed of extra H_2 from an electrolyser (CS2-A) and without supplying outsourced H_2 (CS2-B). The resulting gas mixture (S28) is purified to deliver the best purity for bio-CH_4 (S30). During the cooling and purification condenses are generated. These streams are sent to wastewater treatment, not been considered in our simulations.

Model assumptions

For the simulation of this case study, the following assumptions have been adopted:

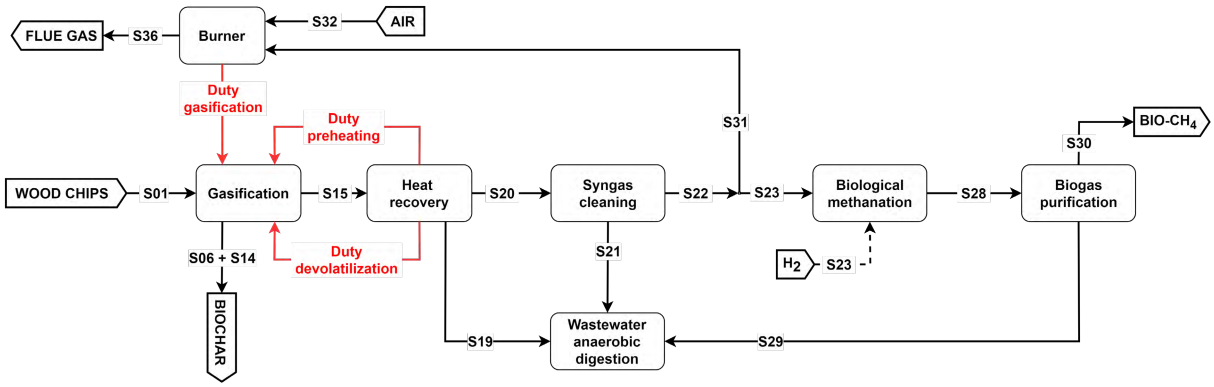

Figure 2. CS2 (France) - BFD

- The same equation of state has been used.
- Elemental and proximate analyses are considered to characterise the feedstock (sensitive).
- Gasification involves a complex chemistry with hundred thousand of species and reactions. The thermochemical process has been simulated considering the different steps of gasification, i.e., heating, devolatilization, pyrolysis, tar cracking, tar reforming, and secondary gas-phase lumped reaction. Details on model and preliminary validation are available [15].
- Syngas cleaning and biomethanation reactor have been simulated as for CS1.
- The fraction of burnt biosyngas is adjusted to reach autothermal conditions in the gasifier. Air is the comburent, in 10% excess. The flue gas leaves the burner at 850°C to mimic flame temperature and ensure that heat is supplied to the gasification chamber (operated at 700-750°C).

Case Study 3 (CS3) - Belgium

Process description

The process and the refrigeration strategies are described with more details in the paper *"Energy recovery from process purges: steam turbine integration and operation optimisation in biogas upgrading within SEMPRE-BIO project"*, which is the complement of this work. Figure 3 depicts the three process layouts considered in the analysis. The proposed BFDs differ for the "cold box". The remaining part of the process is kept unchanged. Manure (S1-01) is digested, resulting in the slurry digestate (S1-04) and raw biogas (S1-07). The produced biogas is purified from H_2S and NH_3, then dehydrated and compressed (S1-19). Its residual moisture is reduced in triethylene glycol (TEG) absorption below 3 ppm vol. Dry biogas (S3-01) leaves the absorber, and a bio-CH_4-rich stream, the top of the stripper of TEG regeneration. This stream is post-treated in a membrane to recover part of the bio-CH_4. Pure bio-CH_4 is recirculated to the stripper. Injecting hot methane as stripping gas in TEG regeneration is an industrial practice to improve TEG purity from 98 to >99.9 wt%, without outsourcing heat. Unrecovered bio-CH_4 (S2-34) is burnt to generate steam for the turbines. The dry biogas (S3-01) is cooled and S3-02 conveyed to the cryogenic distillation: part of the bio-CO_2 is recovered as pure liquid (>99.99 vol%) from the bottom, while bio-CH_4-enriched stream (S3-11 with CO_2 ~80 vol%) is sent to a membrane cascade to purified until outlet gas (S4-04) reaches a residue CO_2 content below 50 ppm (vol), respecting the specifications for grid injection. The membranes work at ambient temperature. The permeated gas is enriched in CO_2 but contains bio-CH_4 as CO_2 permeation is favoured. Part of the permeate is recycled internally, while the purge (S4-12) is cooled and recycled to the cryogenic distillation. Pure bio-CH_4 is cooled and conveyed to the CH_4 liquefaction process (S5-01), where the expansion (Joule-Thomson effect) generates the frigories to liquefy part of the bio-CH_4 (S5-03). The not liquefied fraction (S5-20) is purged. The purges, S2-34 (from TEG dehydration) and S5-20 (from bio-CH_4 liquefaction), are burnt. The combustion of these streams produces heat for steam and power production. Part of the steam is used for the regeneration of TEG solvent (topic of the complementary paper).

Refrigeration strategies

Figure 3 shows three layouts depicting three alternatives to manage the refrigeration and keep the cold box of the system:

- Option A (single loop): the total cooling demand is covered using a hydrocarbon (HC) mixture as working fluid. The HC blend contains C_2-C_3 hydrocarbons. The relative ratio of C_2-C_3 has been tuned to keep the cold box minimum temperature (sensitive).
- Option B (double loop): the same mixture is used

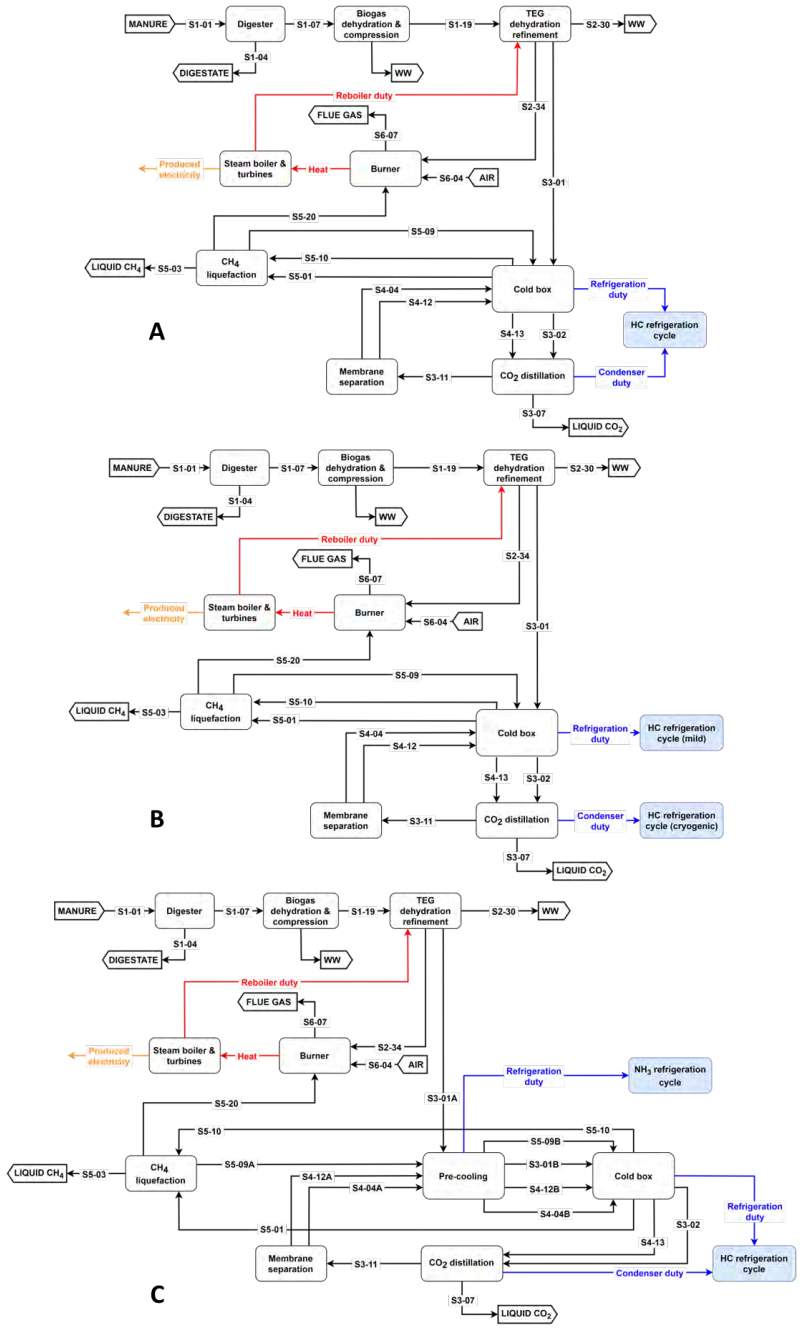

Figure 3. CS3 (Belgium) - BFDs

as working fluid but split into two independent mild refrigeration loops. In the first refrigeration loop (named as loop), all warm streams are precooled from ambient temperature (25-30°C) to reach the same intermediate temperature of the gas leaving the top of the cryogenic column. This temperature is sensitive and is not reported. In the second refrigeration loop (named as cryo loop), all the streams are cooled to the final desired temperature (sensitive).

- Option C (double loop with two refrigerants): is like Option B, but all warm streams are pre-cooled to the same intermediate temperature using ammonia for the mild loop. The temperature is reduced to the final target value in a second refrigeration cycle, using the same HC mixture as for options A and B.

Model assumptions

The following assumptions are considered:

- Biogas flow and volumetric composition are assigned (sensitive).
- The same efficiency for pumps and compressors, H_2S and NH_3 removal as for CS1 and CS2.
- The template for TEG has been retrieved from examples and ready-to-use flowsheets available in COCO COFE webpage [16].
- The membranes are modelled as a Python-based custom module integrated within the flowsheet. Permeance and selectivity for an industrial membrane are retrieved from the literature [17].
- Other model assumptions related to the refrigeration cycles and the steam turbine are available in the mentioned work.

Technology assessment: definition of key performance indicators (KPIs)

The three technologies have been assessed both individually and in comparison. The specific KPIs for each case study (CS) will be discussed in the dedicated section to contrast our results with pieces of evidence from prior literature. This analysis has been carried out to show the consistency of the implemented models. Regarding the comparison of the three technologies, the assessment has been structured to fairly compare them in terms of net production and other relevant indices. The demo sites validate different concepts with different solids starting feedstock, capacity, scale, and target. Hence, their direct comparison through extensive properties would be misleading. In addition, CS1 and CS2 include biomethanation, whereas purification of biogas is the core of CS3. For this reason, the analysis focuses on the net (dry) bio-CH_4 production per unit of both dry solid feedstock and volume of dry biogas or biosyngas. Here, biosyngas is considered because CS2 consider gasification of woody biomass and not biogas from solid waste fermentation. Remarkably, both bio-CH_4 and the reference solid/gas feedstock are considered on a dry basis. The bio-CH_4 purity, electricity demand, and cooling water consumption are other relevant parameters. Remarkably, the PEM electrolyser is considered for both CS1 and CS2A to supply green H_2. For these cases, the electrolyser demand is set from the simulation. The average demand for electrolysis is ~56 kWh/kg_{H2}, in line with literature and technical reports [18-20]. Regarding cooling water demand, the source is assumed at 20°C and the maximum allowable discharge temperature is set 30°C.

RESULTS AND DISCUSSION

Simulation results

This section will show and comment on the results from the simulations. In the first part, each case study will be discussed. Finally, an overview and comparison based on KPIs of the different case studies is addressed.

CS1 results

The most relevant property of the compressed bio-CH_4 is its dew point (T_{dew}) at 250 bar(a). This is the storage pressure of the product for end-user usage (fuel for public transport). The dew point is directly proportional to the purity of delivered bio-CH_4 and the content of impurities. CS1 can deliver compressed bio-CH_4 with a purity of 98.45 vol%. The dew point of the gas mixture has been calculated based on the Redlich-Kwong-Soave model. It corresponds to -60.5°C, meaning that the level of other impurities (moisture, H_2, and CO_2) is kept below the threshold value. The estimated dew temperature respects the European standard EN 16723-2:2017 for compressed natural gas for the use of Bio-CNG (compressed natural gas) as an automotive fuel stored at 200–250 bar.

CS2 results

Table 1 summarises the results related to the steam gasification of the woody biomass. The main KPIs are the gas yield, i.e, the amount of dry gas produced per unit of mass of dry solid, the volumetric composition of the syngas (still on a dry basis) limited to H_2 and CH_4, and the corresponding H_2/CO_x, defining the excess of H_2 compared to the sum of CO_2 and CO. In addition, the analysis considers the tar yield, the energy input, defined as the total amount of heat supply to the gasifier per unit of dry solid, and the heating values (both high and low). The results have been compared to experimental works and review papers referring to steam gasification [20-21]. The mentioned references confirm that the obtained values are aligned with experimental values and ranges for softwood biomass.

Table 1: Overview of CS2: KPIs for biomass pyrolysis (all values refer to dry product and feedstock)

KPI	Value	References
Gas yield	1.23 Nm^3_{syngas}/kg_{solid}	[20, 21]
H_2 content	50.9 vol%	[20, 21]
CH_4 content	6.10 vol%	[20, 21]
H_2/CO_x	1.18 vol/vol	[20, 21]
Char yield	0.25 kg_{char}/kg_{solid}	[20, 21]
Tar yield	17.7 g_{tar}/kg_{solid} 14.4 g_{tar}/Nm^3_{syngas}	[20, 21]
Energy input/ demand	1.02 kWh_{th}/kg_{solid} or 3.66 MJ_{th}/kg_{solid}	[20, 21]
HHW	12.35 MJ/Nm^3_{syngas}	[20, 21]
LHW	10.60 MJ/Nm^3_{syngas}	[20, 21]

Finally, Table 2 gathers the main performance of the technology, including the biomethanation process. For CS2, two different configurations have been considered: with (CS2A) and without (CS2B) extra green H_2 supplied

to the biomethanation.

Remarkably, the results show that the H_2 in the biosyngas is not sufficient to run the biomethanation to full depletion of CO_x. Hence, extra H_2 boosts the productivity of the system; thus, the purity and concentration of bio-CH_4 in the product mixture. The moisture content is still large for transport and use. Therefore, a deep dehydration step should be considered. Beyond this consideration, the analysis highlights the need to combine H_2 production with woody biomass gasification.

Table 2: Overview of KPIs of CS2 (all values refer to dry bio-CH_4 and solid feedstock)

KPI	CS2A	CS2B
Productivity	0.32 Nm^3_{CH4}/kg_{solid}	0.15 Nm^3_{CH4}/kg_{solid}
H_2 addition	1.91 Nm^3_{H2}/Nm^3_{CH4}	Not available
CH_4 purity	95.2 vol%	45.0 vol%
Impurities		
CO_2	0.03 vol%	30.1 vol%
CO	0.07 vol%	20.6 vol%
H_2	0.34 vol%	< 100 ppmv
H_2O	4.34 vol%	4.34 vol%

CS3 results

CS3 is the only demo site where direct biogas purification is performed. Compared to alternative processes such as amines [22], the technology validated in CS3 enables the delivery of both pure bio-CH_4 and bio-CO_2 (purity > 99.99 vol%). This is the main advantage of this new cryogenic system. Other relevant results refer to the integration of boiler and steam turbines, and how the selection of refrigerant and the refrigeration cycle layout affect the energy balance of the system. This paragraph will briefly summarise the main results. More details are reported in *"Energy recovery from process purges: steam turbine integration and operation optimisation in biogas up-grading within SEMPRE-BIO project"*. The analysis showed that keeping the cold box through a refrigeration cycle is energy-demanding. Compression of refrigerants accounts for 34 to 37% of the total electricity demand. The implementation of two independent refrigeration loops (Option B and C) lowers the electricity demand. Using ammonia as a refrigerant in the mild loop (Option C) further drops the electricity demand and valorising waste streams is needed to mitigate the energy demand of the overall process. The simulations highlight that the steam cycle allows for recovering up to 10% of the total electricity demand of the upgrading plant.

Processes comparison

Table 3 highlights the main results of the global KPIs of the CSs, including all process configurations. Looking at the results, it is possible to draw the following conclusions:

- Woody biomass has the highest gas yield, meaning that it is the feedstock that releases the largest amount of gas (biosyngas) per unit of solid. Gasification indeed is an efficient technology to convert solids into syngas. Unfortunately, the same information is not available for wastewater sludge due to missing data related to the efficiency of the sludge fermentation.
- Apparently, biomethanation of biogas from sludge fermentation is the most productive route. It is worth remarking that for CS1 (but also CS2A), the total productivity is boosted due to the addition of green H_2 (in excess to the stoichiometric amount) to run the biomethanation. Conversely, for CS3, the productivity is the net productivity of the system because bio-CH_4 is recovered from biogas directly.
- The electricity consumption of the three technologies is completely different. Whenever biomethanation is part of the process, most of the electricity demand is concentrated in the green H_2 production. For CS3, the energy demand covers the compression of biogas and the energy for the refrigeration cycles (refrigerant compression). If PEM energy demand is disregarded, CS3 consumes almost 10 times more energy than other options.
- Directly from the previous comment, it is evident that biomethanation is an effective solution to produce bio-CH_4 from both biogas and biosyngas. The technology enables a substantial conversion of COx into CH4, thereby breaking the thermodynamic equilibrium and operating the system at thermophilic conditions (milder temperatures, e.g., ~50-55°C). However, despite this advantage, H_2 demand is still relevant and represents one of the bottlenecks of the technology.
- Cooling water demand is in the same order of magnitude for all the considered processes. This parameter is less relevant to the present analysis as it highly depends on the operative condition. Anyway, there are no significant differences in electricity demand.

Regarding CS3, it is worth underlining that the electricity demand covers both the refrigeration cycle and the compression of biogas; more in general, the overall electricity demand to deliver both pure pressurised bio-CH4 and bio-CO2 (99.99+ vol% purity). Refrigeration covers from 34 to 37% of the total energy depending on refrigerant and cycle configuration. The process up to bio-CH4 purification, i.e., disregarding liquefaction of bio-CH4, consumes 1.70 kWh_{el}/Nm^3_{biogas} and 0.15 kWh_{th}/Nm^3_{biogas}.

Table 3: Overview of KPIs of all CSs

	Solid feed-stock	Biogas/ Biosyngas pro-duced	Bio-CH_4 production[1]	Bio-CH_4 purity (vol%)	Electricity de-mand (w/o PEM)[2, 3] [kWh_{el}/Nm^3_{CH4}]	Cooling Water (CW) demand [t_{CW}/Nm^3_{CH4}]
CS1	Sludge	Not Available	0.709 Nm^3_{CH4}/Nm^3_{biogas}	98.5	11.9 (0.30)	0.174
CS2A (extra H_2)	Woody biomass	1.23 Nm^3_{syngas}/kg_{solid}	0.260 $Nm^3_{CH4}/Nm^3_{biosyngas}$ 0.320 Nm^3_{CH4}/kg_{solid}	95.2	10.1 (0)	0.523
CS2B (no extra H_2)	Woody biomass	As CS2A	0.122 $Nm^3_{CH4}/Nm^3_{biosyngas}$ 0.150 Nm^3_{CH4}/kg_{solid}	45.0	0	0.804
CS3 Option A	Manure	3.80 Nm^3_{biogas}/t_{solid}	0.591 Nm^3_{CH4}/Nm^3_{biogas} 0.09 Nm^3_{CH4}/kg_{solid}	99.99+	2.21	0.374
Option B	Manure	As Option A	As Option A	As above	2.15	0.207
Option C	Manure	As Option A	As Option A	As above	2.10	0.205

[1] Referred to dry biogas/biosyngas
[2] Assuming PEM consuming 56 kWh_{el}/kg_{H2}
[3] Energy recovered from turbine is included (only CS3)

Recent literature investigated several amine systems for upgrading biogas [22].

This work reports on average 0.15 kWh_{el}/Nm^3_{biogas} and 0.28 kWh_{th}/Nm^3_{biogas} but neglects the compression and purification of CO_2 and bio-CH_4. The purity of bio-CH_4 is poor, ranging 90-95 vol%; hence, further refinement is needed. This comment holds also for bio-CO_2 which can contain acid pollutants (H_2S) and degradation products. The purification and compression of CO_2 is expensive: values reported in the literature does not represent the actual energy cost for delivering bio-CH_4 matching market specification. The analysis is limited to the bulk separation of CO_2 from biogas. Moreover, the technology tested in the SEMPRE-BIO project (Cryo Inox's patent) allows for delivering both CH_4 and CO_2 pure and liquid to ease their transport. CO_2 liquefaction is always neglected in articles dealing with the biogas upgrading.

CONCLUSION AND FUTURE WORK

This work presented the results of the process simulation and comparison for three technologies for bio-CH_4 production. The analysis showed that biomethanation of both biogas and biosyngas is a valuable opportunity to convert bio-CO_2 into bio-CH_4. However, this option is still under investigation and constrained to the availability of cheap green-H_2. Despite this potential limitation, it is worth noting that biomethanation can increase overall production of the site. Biogas upgrading is a well-established technology to purify bio-CH_4. The SEMPRE-BIO project validates a new option with cryogenic distillation. The analysis showed more efforts should focus on reducing the energy demand. This work is limited to the technical analysis of the three technologies. However, it is important to note that the analysis should be more holistic, accounting for integrated techno-economic and life-cycle assessments to address the economic and environmental aspects. Only a complete analysis encompassing all these aspects allows for leveraging advantage and identifying bottlenecks. This holistic analysis will be finalised later within the project.

ACKNOWLEDGEMENTS

This work received funding from the SEMPRE-BIO project (Link to project webpage and Link to official EU portal page), funded within the EU's Horizon Europe programme under grant agreement Nº 101084297. Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or CINEA. Neither the European Union nor the granting authority can be held responsible for them.

REFERENCES

1. IEA. An introduction to biogas and biomethane. https://www.iea.org/reports/outlook-for-biogas-and-biomethane-prospects-for-organic-growth/an-introduction-to-biogas-and-biomethane
2. IEA. Outlook for biogas and biomethane. https://www.iea.org/reports/outlook-for-biogas-and-biomethane

3. Kougias PG, Angelidaki I. Biogas and its opportunities—a review. Front. Environ. Sci. Eng. 12: (2018). https://doi.org/10.1007/s11783-018-1037-8
4. Deng R, Wu J, Huang Z, Feng Z, Hu W, Tang Y, Tan H, Zhang H, Zairov R, Pan Z. Biogas to chemicals: a review of the state-of-the-art conversion processes. Biomass Conv. Bioref. 15:14653-14673 (2024). https://doi.org/10.1007/s13399-024-06343-1
5. EBA. Statistical Report 2025. https://www.europeanbiogas.eu/news/eba-statistical-report-2025/
6. Werkneh AA. Biogas impurities: environmental and health implications, removal technologies and future perspectives. Heliyon 8:e10929 (2022). https://doi.org/10.1016/j.heliyon.2022.e10929
7. Swinbourn R, Li C, Wang F. A comprehensive review on biomethane production from biogas separation and its techno?economic assessments. ChemSusChem 17: (2024). https://doi.org/10.1002/cssc.202400779
8. Jønson BD, Tsapekos P, Tahir Ashraf M, Jeppesen M, Ejbye Schmidt J, Bastidas-Oyanedel JR. Pilot-scale study of biomethanation in biological trickle bed reactors converting impure CO2 from a full-scale biogas plant. Bioresource Technology 365:128160 (2022). https://doi.org/10.1016/j.biortech.2022.128160
9. Grimalt-Alemany A, Skiadas IV, Gavala HN. Syngas biomethanation: state-of-the-art review and perspectives. Biofuels Bioprod Bioref 12:139-158 (2017). https://doi.org/10.1002/bbb.1826
10. Krishania M, Kumar V, Vijay VK, Malik A. Analysis of different techniques used for improvement of biomethanation process: A review. Fuel 106:1-9 (2013). https://doi.org/10.1016/j.fuel.2012.12.007
11. Jusoh N, Adibah Tengku Hassan TN, Suhaimi NH, Mubashir M. Hydrogen sulfide removal from biogas: an overview of technologies emphasizing membrane separation. Separation and Purification Technology 373:133466 (2025). https://doi.org/10.1016/j.seppur.2025.133466
12. Ancona MA, Antonioni G, Branchini L, De Pascale A, Melino F, Orlandini V, Antonucci V, Ferraro M. Renewable energy storage system based on a power-to-gas conversion process. Energy Procedia 101:854-861 (2016). https://doi.org/10.1016/j.egypro.2016.11.108
13. Abad D, Vega F, Navarrete B, Delgado A, Nieto E. Modeling and simulation of an integrated power-to-methanol approach via high temperature electrolysis and partial oxy-combustion technology. International Journal of Hydrogen Energy 46:34128-34147 (2021). https://doi.org/10.1016/j.ijhydene.2021.04.029
14. Li C, Zhu X, Angelidaki I. Carbon monoxide conversion and syngas biomethanation mediated by different microbial consortia. Bioresource Technology 314:123739 (2020). https://doi.org/10.1016/j.biortech.2020.123739
15. Bisotti F, Gilardi M, De Carli P, Mercier Y, Wittgens B. A Novel Semi-rigorous Model for Biomass Pyrolysis and Gasification, Chem. Eng. Trans., 119: 241-246 (2025). https://doi.org/10.3303/CET25119041
16. AmsterCHEM. https://www.cocosimulator.org
17. Sinaei Nobandegani M, Yu L, Hedlund J. Zeolite membrane process for industrial CO2/CH4 separation. Chemical Engineering Journal 446:137223 (2022). https://doi.org/10.1016/j.cej.2022.137223
18. Berrada A, Sanjari MJ, El Mrabet R. Techno-economic assessment of hydrogen production: comparative analysis of electrolyser technologies in a hybrid pv/wind system. International Journal of Hydrogen Energy 141:193-211 (2025). https://doi.org/10.1016/j.ijhydene.2025.05.384
19. Elrhoul D, Naveiro M, Romero Gómez M. Thermo-economic comparison between three different electrolysis technologies powered by a conventional organic rankine cycle for the green hydrogen production onboard liquefied natural gas carriers. JMSE 12:1287 (2024). https://doi.org/10.3390/jmse12081287
20. Cerone N, Zimbardi F, Contuzzi L, Stri?gas N, Eimontas J, Zito GD. Production of hydrogen-rich syngas from biomass gasification by double step steam catalytic tar reforming. International Journal of Hydrogen Energy 95:1215-1221 (2024). https://doi.org/10.1016/j.ijhydene.2024.07.290
21. Daugaard DE, Brown RC. Enthalpy for pyrolysis for several types of biomass. Energy Fuels 17:934-939 (2003). https://doi.org/10.1021/ef020260x
22. Jabraeelzadeh A, Gharehghani A, Ahbabi Saray J, Mahmoudzadeh Andwari A, Borhani TN. Techno-economic analysis of biogas upgrading through amine scrubbing: a comparative study of different single amines. Fuel 381:133662 (2025). https://doi.org/10.1016/j.fuel.2024.133662

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Modeling and Simulation of Nitrogen Generation by Pressure Swing Adsorption for Power-to-Ammonia

Marcus J. Schytt[a], Lorenz T. Biegler[b], and John B. Jørgensen[a]*

[a] Technical University of Denmark, DTU Compute, Kongens Lyngby, Capital Region, Denmark
[b] Carnegie Mellon University, Department of Chemical Engineering, Pittsburgh, Pennsylvania, USA
* Corresponding Author: jbjo@dtu.dk.

ABSTRACT

Power-to-ammonia (P2A) provides a carbon-free alternative to conventional ammonia production by replacing fossil-based feedstocks with electrolytic hydrogen and nitrogen from air separation. For decentralized P2A systems, pressure swing adsorption (PSA) offers a flexible alternative to cryogenic air separation. However, its industrial implementations are largely proprietary, and open, first-principles models capable of simulating its cyclic, nonlinear transport are scarce in literature. This work presents a first-principles, dynamic, one-dimensional model of a PSA superstructure for nitrogen generation, formulated with thermodynamically consistent equations of state, coupling multicomponent mass, energy, and momentum balances with kinetically limited adsorption on carbon molecular sieves. The resulting system of partial differential-algebraic equations is semi-discretized using the finite volume method, integrated using diagonally implicit Runge-Kutta methods, and cyclic steady states (CSS) are computed via shooting-based solution methods. The framework is implemented in Julia, combining analytical derivatives with automatic differentiation and utilizing sparse linear algebra for efficient solution of the arising large nonlinear systems. The framework is demonstrated on a two-bed PSA cycle for air separation, comparing spatial and temporal discretization strategies, CSS solution methods, and the effects of ideal versus real-gas thermodynamics on predicted nitrogen purity and recovery. The proposed framework establishes an extensible basis for PSA simulation and optimization.

Keywords: Adsorption, Alternative Fuels, Dynamic Modelling, Numerical Methods, Simulation

INTRODUCTION

Conventional ammonia production relies on the famous Haber–Bosch process, combining nitrogen with hydrogen derived from fossil-based syngas. The resulting grey ammonia, of which approximately 80% is used for fertilizer production, accounts for about 4.5% of global industrial CO_2 emissions [1]. To mitigate this environmental impact, power-to-ammonia (P2A) has emerged as a carbon-free alternative, replacing fossil-based hydrogen with hydrogen produced by water electrolysis and sourcing nitrogen from atmospheric air. When powered by renewable electricity, P2A enables the production of green ammonia, which can serve both as a carbon-free hydrogen carrier and as a potential maritime fuel.

Nitrogen is traditionally produced by large-scale cryogenic air separation. However, for the smaller and decentralized production units envisioned in P2A systems, pressure swing adsorption (PSA) offers a flexible alternative. PSA achieves high-purity (95%+) nitrogen through cyclic adsorption and desorption of air components on carbon molecular sieves (CMS). Despite its widespread industrial adoption, PSA process design remains largely proprietary, and the open literature, often relying on commercial simulators, offers limited first-principles models capable of simulating its inherently cyclic operation. This limitation is exacerbated by strong multiscale coupling between transport, adsorption kinetics, and thermodynamics, making systematic computation of cyclic steady states (CSS) especially challenging.

Motivated by these gaps, this work develops a first-principles adsorption column model and superstructure for PSA-based nitrogen generation, which, using well-established numerical discretization methods, enables systematic computation of thermodynamically consistent cyclic operation.

ADSORPTION COLUMN MODEL

We model adsorption columns with fixed beds composed of carbon molecular sieve pellets, specifically the commercial Takeda 3A CMS adsorbent. Parameters describing the adsorbent are given in the digital supplementary material. We model the adsorption column as a cylinder with length L on the axial interval domain $\Omega = [0, L]$. We study a set of components $\mathcal{C} = \{\mathrm{N}_2, \mathrm{O}_2, \mathrm{Ar}\}$ representative of the composition of air (78%, 21%, 1%). To model the adsorption process, the column is treated heterogeneously and divided into a mobile phase and a stationary phase with porosity ε. Accordingly, the column volume V holds the gaseous volume $V_{\text{gas}} = \varepsilon V$, and the ratio between the gaseous and stationary volumes is given by $\varphi = \varepsilon/(1-\varepsilon)$. We denote mobile-phase concentrations by $c = [c_\alpha]_{\alpha\in\mathcal{C}}$, and stationary-phase concentrations by $q = [q_\alpha]_{\alpha\in\mathcal{C}}$. Moreover, we let u denote the pseudo-homogeneous total internal energy density, accounting for both the mobile and stationary phases. The mass and energy balances of the adsorption column are then summarized by a system of one-dimensional transport equations subject to thermodynamic constraints

$$\partial_t c = -\partial_z N + R, \tag{1a}$$

$$\partial_t u = -\partial_z E + Q. \tag{1b}$$

Here, $N = [N_\alpha]_{\alpha\in\mathcal{C}}$ denotes the molar fluxes, $R = [R_\alpha]_{\alpha\in\mathcal{C}}$ the adsorption rates, E the heat flux, and Q the heat transfer rate.

Thermodynamics

We introduce thermodynamic functions to evaluate bulk gas-phase properties. Assuming local thermodynamic equilibrium, the thermodynamic state is characterized by the local temperature, pressure, and composition of the system. While the composition is distributed between the mobile and stationary phases, the temperature and pressure are common to both phases. A gas phase with temperature T, pressure P, and moles $n = [n_\alpha]_{\alpha\in\mathcal{C}}$, can be described by thermodynamic functions

- Volume: $\mathcal{V}(T, P, n)$
- Enthalpy: $\mathcal{H}(T, P, n)$
- Internal energy: $\mathcal{U}(T, P, n)$

provided by an equation of state (EOS). The thermodynamic functions satisfy standard identities. In particular, they are related by

$$\mathcal{U}(T, P, n) = \mathcal{H}(T, P, n) - P\,\mathcal{V}(T, P, n). \tag{2}$$

Moreover, the thermodynamic functions are first-order homogeneous in the component moles n, i.e.,

$$\mathcal{V}(T, P, \lambda n) = \lambda \mathcal{V}(T, P, n), \qquad \lambda > 0, \tag{3}$$

with analogous relations holding for $\mathcal{H}$ and $\mathcal{U}$. The EOS may be specified by ideal-gas correlations, with or without cubic real-gas modifications.

Transport

The molar fluxes account for mobile-phase convection and are given by

$$N = vc, \tag{4}$$

assuming negligible diffusive contributions. Here, v denotes the interstitial axial velocity. The velocity is related to the pressure gradient through a momentum balance. We employ Ergun's equation

$$-\partial_z P = 150\mu \left(\frac{\varphi}{d_p}\right)^2 v + 1.75\rho_{\text{gas}} \frac{\varphi}{d_p} v|v|. \tag{5}$$

Here, μ denotes the mobile-phase dynamic viscosity, d_p the effective particle diameter, and

$$\rho_{\text{gas}} = \textstyle\sum_{\alpha\in\mathcal{C}} c_\alpha M_\alpha, \tag{6}$$

the mobile-phase density, which is defined by molar masses $M = [M_\alpha]_{\alpha\in\mathcal{C}}$. We fix the dynamic viscosity to be that of air at ambient conditions and invert (6) to yield an explicit expression for the velocity $v = v(\partial_z P, c)$.

The heat flux accounts for mobile-phase convection and is expressed in terms of the enthalpy flux

$$E = \varepsilon \mathcal{H}(T, P, N), \tag{7}$$

neglecting gas- and solid-phase conduction. Here, first-order homogeneity is employed to express the enthalpy flux in terms of the enthalpy function and molar fluxes.

Adsorption and heat transfer

The adsorption rate R is defined by the stationary-phase mass balances

$$R = -\varphi\, \partial_t q, \tag{8}$$

where the volume ratio φ ensures consistent mobile- and stationary-phase volume bases. The stationary-phase mass balances are modeled using the linear driving force (LDF) approximation

$$\partial_t q = \omega \odot (q^* - q), \tag{9}$$

With rate constants $\omega = [\omega_\alpha]_{\alpha\in\mathcal{C}}$ and equilibrium concentrations $q^* = [q^*_\alpha]_{\alpha\in\mathcal{C}}$. They are described by correlations, with parameters provided in the supplementary material.

Adsorption equilibria

Following [2], we consider the multi-component Langmuir isotherm, which for a given component $\alpha \in \mathcal{C}$ is defined by

$$q^*_\alpha = q_{\max,\alpha} \frac{b_\alpha P_\alpha}{1 + \sum_{\beta\in\mathcal{C}} b_\beta P_\beta}. \tag{10}$$

Here, $q^{\max}_\alpha$ denotes the component's saturation capacity,

b_α its Langmuir parameter, and P_α its partial pressure. The partial pressure is given by $P_\alpha = Pc_\alpha / \sum_{\beta\in\mathcal{C}} c_\beta$. The remaining quantities are specified by temperature-dependent correlations

$$q_{\max,\alpha} = k_{1,\alpha} - k_{2,\alpha} T, \tag{11a}$$

$$b_\alpha = k_{3,\alpha} exp\left(\frac{k_{4,\alpha}}{T}\right), \tag{11b}$$

with component-wise constants $k_{i,\alpha}$ for $i = 1, 2, 3, 4$.

Adsorption rates

For each component $\alpha \in \mathcal{C}$, the adsorption rate constant is defined as

$$\omega_\alpha = \Omega \frac{D_\alpha}{r_p^2}, \tag{12}$$

where Ω is a geometric factor, D_α denotes the component's effective pore diffusivity, and r_p is the effective particle radius. Assuming effectively spherical particles, we take $\Omega = 15$. The apparent time constant D_α / r_p^2 is modeled as a pressure-dependent quantity using the supercritical structural Langmuir model [2]

$$\frac{D_\alpha}{r_p^2} = k_\alpha \sqrt{P_{r,\alpha}} (1 + b_\alpha P_\alpha)^2, \tag{13}$$

where k_α is a kinetic parameter and $P_{r,\alpha} = P_\alpha / P_{c,\alpha}$ denotes the reduced pressure defined by the component's critical pressure $P_{c,\alpha}$. The Langmuir parameter b_α is consistent with the adsorption equilibrium model.

Heat of adsorption

Neglecting wall effects, the heat transfer rate accounts for the heat of adsorption

$$Q = -(1-\varepsilon) \sum_{\alpha\in\mathcal{C}} \Delta H_\alpha \, \partial_t q_\alpha, \tag{14}$$

where $\Delta H = [\Delta H_\alpha]_{\alpha\in\mathcal{C}}$ denotes the component-wise heats of adsorption. Note that the heat of adsorption is taken to be negative for exothermic adsorption. Component-wise heats of adsorption are obtained from Clausius-Clapeyron analysis of the equilibrium isotherm [2].

Thermodynamic constraints

The thermodynamic EOS is enforced through closures that recover the pseudo-homogeneous temperature and pressure fields. We provide constraints on the internal energy and volume.

Internal energy constraint

The total internal energy density is constrained by enforcing the decomposition into mobile- and stationary-phase contributions

$$u = \varepsilon u_{\text{gas}} + (1-\varepsilon) u_{\text{stationary}}. \tag{15}$$

Exploiting first-order homogeneity, the mobile-phase internal energy density is given by

$$u_{\text{gas}} = \mathcal{U}(T, P, c). \tag{16}$$

For the stationary phase, the internal energy density is split into its pore-phase and solid-phase contributions $u_{\text{stationary}} = u_{\text{pore}} + u_{\text{solid}}$. The pore-phase internal energy density is similarly given by

$$u_{\text{pore}} = \mathcal{U}(T, P, q). \tag{17}$$

Moreover, the solid-phase internal energy density is defined as

$$u_{\text{solid}} = \rho_{\text{solid}} c_{p,\,\text{solid}} (T - T_{\text{ref}}), \tag{18}$$

where ρ_{solid} denotes the solid-phase density, $c_{p,\,\text{solid}}$ the heat capacity, and T_{ref} a reference temperature.

Volume constraint

The total gas volume comprises the mobile-phase and pore gas volumes. Due to sorption effects, the pore gas volume is not well-described by the EOS. Instead, a volume constraint is enforced on the mobile-phase gas

$$\mathcal{V}(T, P, c) = 1, \tag{19}$$

which, by first-order homogeneity, ensures that the relation $\mathcal{V}(T, P, n) = V_{\text{gas}}$ holds.

Initial and boundary conditions

The differential mass and energy balances are closed by a set of initial and boundary conditions.

Initial conditions

Initial conditions are prescribed for the differential states (c, q, u) at time $t = 0$ across the axial domain $z \in \Omega$

$$c(0, z) = c_0(z), \tag{20a}$$

$$q(0, z) = q_0(z), \tag{20b}$$

$$u(0, z) = u_0(z). \tag{20c}$$

From these, consistent initial profiles of the algebraic states (T, P) are subsequently obtained from the thermodynamic constraints (15) and (19).

Boundary conditions

Boundary conditions are prescribed directly on the boundary fluxes

$$N(t, 0) = N_{\text{in}}(t), \tag{21a}$$

$$E(t, 0) = E_{\text{in}}(t), \tag{21b}$$

$$N(t, L) = N_{\text{out}}(t), \tag{21c}$$

$$E(t, L) = E_{\text{out}}(t). \tag{21d}$$

We employ flow-driven boundary conditions, specifying the upstream flows and downstream pressure. Consequently, the inlet fluxes are defined by

$$N_{\text{in}}(t) = \frac{f_{\text{in}}(t)}{S_{\text{gas}}}, \tag{22a}$$

$$E_{\text{in}}(t) = \varepsilon \frac{h_{\text{in}}(t)}{S_{\text{gas}}}, \tag{22b}$$

where $f_{\text{in}} = [f_{\text{in},\alpha}]_{\alpha \in \mathcal{C}}$ denotes the inlet molar flow rates and $h_{\text{in}} = \mathcal{H}(T_{\text{in}}, P_{\text{in}}, f_{\text{in}})$, the corresponding inlet enthalpy flow, specified by the inlet temperature T_{in} and pressure P_{in}. Moreover, $S_{\text{gas}} = V_{\text{gas}}/L$ denotes the cross-sectional gas area. At the outlet, the prescribed pressure $P_{\text{out}}(t)$ determines the outlet pressure gradient $\partial_z P(t, L)$ and thus the outlet velocity $v_{\text{out}}(t)$ via the momentum balance (5). To prevent backflow, we adopt a check-valve formulation and determine

$$v_{\text{out}}(t) = \max_{\delta}\{0, v\big(\partial_z P(t, L), c(t, L)\big)\}, \tag{23}$$

where the smoothed maximum operator is given by $\max_{\delta}\{a, b\} = \left(a + b + \sqrt{(a-b)^2 + \delta}\right)/2$ with $\delta = 10^{-4}$. The outlet fluxes are then defined by

$$N_{\text{out}}(t) = v_{\text{out}}(t) c(t, L), \tag{24a}$$

$$E_{\text{out}}(t) = \varepsilon \mathcal{H}\big(T(t, L), P(t, L), N_{\text{out}}(t)\big). \tag{24b}$$

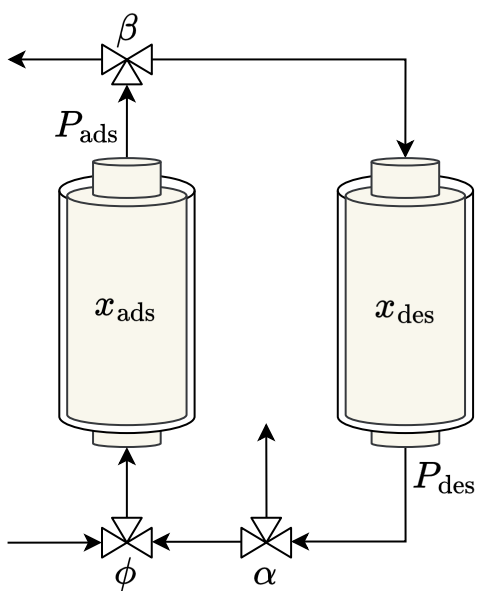

Figure 1. Diagram of the two-bed PSA superstructure.

PSA SUPERSTRUCTURE

The PSA process is simulated by periodically sequencing a set of operating steps (pressurization, adsorption, blowdown, desorption, purge, and pressure equalization) across multiple adsorption beds. In a conventional two-bed PSA configuration, the operating steps are symmetric and may be divided into two complementary half-cycles: an adsorption half-cycle and a desorption half-cycle. While one bed undergoes the adsorption half-cycle, the other simultaneously undergoes the desorption half-cycle. Under this symmetry, it is sufficient to simulate the two-bed interactions across a single half-cycle, with the remaining half-cycle obtained by swapping beds. This motivates the PSA superstructure in [3], as shown in Figure 1. The superstructure defines the interconnections between the adsorption beds through idealized valves and imposed adsorption and desorption pressure levels. We omit the modeling of auxiliary equipment, and set pressure levels via boundary conditions.

Mathematical formulation

We formulate the PSA superstructure over a half-cycle time horizon $t \in [0, \tau]$. Each bed is represented by an instance of the adsorption column model described previously. For each bed, we denote its differential states by $x = [c; q; u]$ and its algebraic states by $y = [T, P]$. The two beds are distinguished by subscripts indicating whether the bed undergoes adsorption or desorption. The left-most bed in Figure 1 undergoes the adsorption half-cycle, while the right-most bed undergoes the desorption half-cycle.

Boundary conditions

In the left-most bed, a fraction $\phi(t)$ of the constant feed stream $(f_{\text{in}}, h_{\text{in}})$ enters from the bottom inlet, together with a fraction $\alpha(t)$ of the right-most bed outlet

$$N_{\text{in, ads}}(t) = \phi(t) \frac{f_{\text{in}}}{S_{\text{gas}}} + \alpha(t) N_{\text{out, des}}(t), \tag{25a}$$

$$E_{\text{in, ads}}(t) = \varepsilon \phi(t) \frac{h_{\text{in}}}{S_{\text{gas}}} + \alpha(t) E_{\text{out, des}}(t). \tag{25b}$$

Similarly, a fraction $\beta(t)$ of the left-most bed outlet enters the top inlet of the right-most bed

$$N_{\text{in, des}}(t) = \beta(t) N_{\text{out, ads}}(t), \tag{26a}$$

$$E_{\text{in, des}}(t) = \beta(t) E_{\text{out, ads}}(t). \tag{26b}$$

Here, the outlet fluxes are defined by (24), via the prescribed adsorption pressure $P_{\text{ads}}(t)$ and desorption pressure $P_{\text{des}}(t)$. The superstructure enforces unidirectional flow through the check-valve formulation (23), allowing the half-cycle to be simulated without flow reversal.

Cyclic steady state

A consistent PSA cycle is obtained from the simulated half-cycle at CSS. Exploiting the symmetry of the two-bed superstructure, the full PSA cycle is recovered by imposing periodic boundary conditions, such that the terminal state of one bed at the end of a half-cycle provides the initial condition for the other bed at the beginning of the subsequent half-cycle, and vice versa. Accordingly, the differential states of the superstructure are required to solve the boundary value problem (BVP)

$$x_{\text{des}}(\tau, L - z) = x_{\text{ads}}(0, z), \tag{27a}$$

$$x_{\text{ads}}(\tau, L - z) = x_{\text{des}}(0, z). \tag{27b}$$

for all $z \in \Omega$, where the axial reversal reflects the flow

configuration of the superstructure. Consistent periodic profiles of the algebraic states are recovered from the thermodynamic constraints as part of the CSS solution.

NUMERICAL SIMULATION

We simulate the PSA superstructure by discretizing the governing system of PDAEs in space using the method of lines. The resulting semi-discrete system of DAEs is then integrated in time.

Finite volume discretization

We discretize the axial domain Ω into K uniform volumes $\Omega_h^k = [z_h^{k-1/2}, z_h^{k+1/2}]$ of length $h = L/K$, and consider a cell-centered finite volume approximation $w_h = [w_h^k]_{k=1}^K$ of a given quantity w. For a volume Ω_h^k, we define its midpoint $z_h^k = (k-1/2)h$ and interfaces $z_h^{k\pm1/2} = z_h^k \pm h/2$, and approximate the volume-averaged quantity

$$w_h^k(t) \approx \langle w \rangle_h^k(t) = \frac{1}{h}\int_{z_h^{k-1/2}}^{z_h^{k+1/2}} w(t,z)\,\mathrm{d}z, \tag{28}$$

by midpoint evaluation $w_h^k(t) = w(t, z_h^k)$. Applying the fundamental theorem of calculus and the same midpoint approximation to the volume-averaged mass and energy balances (1), we find that

$$\partial_t c_h^k = -\frac{1}{h}\left(N_h^{k+1/2} - N_h^{k-1/2}\right) + R_h^k, \tag{29a}$$

$$\partial_t u_h^k = -\frac{1}{h}\left(E_h^{k+1/2} - E_h^{k-1/2}\right) + Q_h^k, \tag{29b}$$

where $N_h^{k\pm1/2}$ and $E_h^{k\pm1/2}$ denote the interface molar and heat fluxes, respectively, and where R_h^k and Q_h^k denote the midpoint approximations of (8) and (14)

$$R_h^k = -\varphi\,\partial_t q_h^k, \tag{30a}$$

$$Q_h^k = -(1-\varepsilon)\sum_{\alpha\in\mathcal{C}} \Delta H_\alpha (\partial_t q_\alpha)_h^k. \tag{30b}$$

Here, $\partial_t q_h^k = \left[(\partial_t q_\alpha)_h^k\right]_{\alpha\in\mathcal{C}}$ is defined by the midpoint approximations of the LDF approximation

$$\partial_t q_h^k = \omega_h^k \odot \left((q^*)_h^k - q_h^k\right), \tag{31}$$

with $(q^*)_h^k = q^*(T_h^k, P_h^k, c_h^k)$ and $\omega_h^k = \omega(T_h^k, P_h^k, c_h^k)$.

Numerical fluxes

Since the superstructure enforces unidirectional flow, stability is ensured by prescribing upwind numerical fluxes at the volume interfaces. They are defined as

$$N_h^{k+1/2} = v(\partial_z P_h^{k+1/2}, c_h^k) c_h^k, \tag{32a}$$

$$E_h^{k+1/2} = \varepsilon\mathcal{H}(T_h^k, P_h^k, N_h^{k+1/2}), \tag{32b}$$

where the interface pressure gradient is approximated by $\partial_z P_h^{k+1/2} = (P_h^{k+1} - P_h^k)/h$. This results in a monotone first-order accurate spatial discretization.

Thermodynamic constraints

The thermodynamic constraints are imposed locally in each finite volume Ω_h^k. Applying the midpoint approximation to the internal energy constraint (15) yields

$$u_h^k = \varepsilon(u_{\text{gas}})_h^k + (1-\varepsilon)(u_{\text{stationary}})_h^k, \tag{33}$$

with $(u_{\text{stationary}})_h^k = (u_{\text{solid}})_h^k + (u_{\text{pore}})_h^k$. Here, the corresponding phase contributions are given by midpoint evaluation of (16)-(18) as

$$(u_{\text{gas}})_h^k = \mathcal{U}(T_h^k, P_h^k, c_h^k), \tag{34a}$$

$$(u_{\text{pore}})_h^k = \mathcal{U}(T_h^k, P_h^k, q_h^k), \tag{34b}$$

$$(u_{\text{solid}})_h^k = \rho_{\text{solid}} c_{p,\,\text{solid}}(T_h^k - T_{\text{ref}}). \tag{34c}$$

Similarly, the volume constraint (19) is also enforced by midpoint evaluation

$$\mathcal{V}(T_h^k, P_h^k, c_h^k) = 1. \tag{35}$$

Initial and boundary conditions

The initial conditions are also discretized by midpoint evaluation

$$c_h^k(0) = c_0(z_h^k), \tag{36a}$$

$$q_h^k(0) = q_0(z_h^k), \tag{36b}$$

$$u_h^k(0) = u_0(z_h^k), \tag{36c}$$

from which the initial algebraic states are recovered by the discrete thermodynamic constraints. The inlet boundary conditions are simply given by the inlet fluxes

$$N_h^{1/2}(t) = N_{\text{in}}(t), \tag{37a}$$

$$E_h^{1/2}(t) = E_{\text{in}}(t), \tag{37b}$$

and the outlet boundary conditions are given by the upwind fluxes

$$N_h^{K+1/2}(t) = v\left(\partial_z P_h^{K+1/2}(t), c_h^K(t)\right) c_h^K(t), \tag{38a}$$

$$E_h^{K+1/2}(t) = \varepsilon\mathcal{H}\left(T_h^K(t), P_h^K(t), N_h^{K+1/2}(t)\right), \tag{38b}$$

Here, the outlet pressure gradient is approximated by $\partial_z P_h^{K+1/2}(t) = \left(P_{\text{out}}(t) - P_h^K(t)\right)/(h/2)$ using the prescribed outlet pressure $P_{\text{out}}(t)$.

Semi-discrete system

We let $x_h = [c_h; q_h; u_h]$ and $y_h = [T_h; P_h]$ denote the finite-volume discretizations of the differential and algebraic states, respectively. For the two-bed PSA superstructure, the corresponding discretized states of both beds are collected into global differential and algebraic state vectors $\mathbf{x}$ and $\mathbf{y}$. The spatial discretization of the superstructure then yields a semi-discrete system of semi-explicit index-1 DAEs

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{y}, p), \mathbf{x}(0) = \mathbf{x}_0, \tag{39a}$$

$$0 = g(\mathbf{x}, \mathbf{y}, p), \tag{39b}$$

where $\mathbf{x}_0$ is defined by the discretized initial conditions, and p the prescribed boundary conditions.

Initialization

Before integration, we solve for consistent initial algebraic states $\mathbf{y}(0) = \mathbf{y}_0$ satisfying

$$g(\mathbf{x}_0, \mathbf{y}_0, p) = 0. \tag{40}$$

To this end, we employ Newton's method with backtracking line search, which restricts trial iterates to the feasible domain of the model functions and prevents evaluation at physically inadmissible states.

Diagonally implicit Runge-Kutta methods

Time integration of (39) is performed using a diagonally implicit Runge-Kutta (DIRK) method [5]. Given states $(\mathbf{x}_n, \mathbf{y}_n)$ at time $t = t_n$, an s-stage DIRK method advances to the states $(\mathbf{x}_{n+1}, \mathbf{y}_{n+1})$ at time $t = t_{n+1}$, with step size $h_n = t_{n+1} - t_n$, by computing stage values $(\mathbf{X}_{n,i}, \mathbf{Y}_{n,i})$ satisfying the implicit stage equations

$$\mathbf{X}_{n,i} = \mathbf{x}_n + h_n \sum_{j=1}^{s} a_{i,j} f(\mathbf{X}_{n,j}, \mathbf{Y}_{n,j}, p), \tag{41a}$$

$$0 = g(\mathbf{X}_{n,i}, \mathbf{Y}_{n,i}, p), \tag{41b}$$

defined at stage times $T_{n,i} = t_n + c_i h$, for $i = 1, \dots, s$. The advancing states are found via the stage values using the quadrature rule

$$\mathbf{x}_{n+1} = \mathbf{x}_n + h_n \sum_{i=1}^{s} b_i f(\mathbf{X}_{n,i}, \mathbf{Y}_{n,i}, p), \tag{42a}$$

$$0 = g(\mathbf{x}_{n+1}, \mathbf{y}_{n+1}). \tag{42b}$$

The DIRK method is characterized by a Butcher tableau with coefficients $a = [a_{i,j}]_{i,j=1}^{s}, b = [b_i]_{j=1}^{s}$, and $c = [c_i]_{i=1}^{s}$, where $a_{i,j} = 0$ for $i < j$. In this work, we employ singly-diagonally implicit Runge-Kutta (SDIRK) methods, and in particular explicit first-stage SDIRK (ESDIRK) methods, with constant diagonal coefficients $a_{i,i} = \gamma$ and $a_{1,1} = 0$. We further adopt stiffly accurate schemes where $b_i = a_{s,i}$, such that the final stage directly gives the advancing state. This combination of properties renders stiffly accurate ESDIRK methods particularly efficient. The lower-triangular structure enables sequential solution of the stage equations via Newton's method, while the constant diagonal allows reuse of a single Jacobian factorization across all stages within a time step, making quasi-Newton schemes especially effective. The remaining coefficients are then chosen to ensure high-order accuracy and stability. Step sizes are chosen adaptively, to control the temporal error and to handle transient stiffness arising from changes in boundary conditions. We use PID-based step-size control via embedded error estimates, as described in [5].

CYCLIC STEADY-STATE SIMULATION

While the previous section describes numerical simulation of a single half-cycle of the PSA superstructure, a physically consistent PSA cycle is obtained only when the half-cycle satisfies the CSS conditions defined by the BVP (27). Using the method-of-lines discretization, the CSS BVP is solved numerically using shooting methods.

Single shooting

Single shooting aims to solve the BVP by iteratively integrating the semi-discrete DAE system and updating the initial state until CSS is achieved. Using the ESDIRK integrators, we introduce the half-cycle propagator Φ, which, for given initial differential states $\mathbf{x}_0$ and boundary conditions p, returns the terminal states $\Phi(\mathbf{x}_0, p)$ at the end of the half cycle via numerical integration over $t \in [0, \tau]$. Additionally, we define the operator Π, which implements the permutation of the superstructure states that swaps the two beds and applies the axial reversal. The discretized CSS BVP reduces to the fixed-point problem

$$\Pi\Phi(\mathbf{x}, p) = \mathbf{x}, \tag{43}$$

whose solution defines initial states for a consistent PSA cycle. Alternatively, we adopt the residual formulation

$$R(\mathbf{x}) \quad = \Pi\Phi(\mathbf{x}, p) - \mathbf{x}, \tag{44}$$

and seek a solution to $R(\mathbf{x}) = 0$. The fixed-point formulation of single shooting lends itself to simple fixed-point (Picard) iterations and acceleration techniques, such as Anderson acceleration (AA). In contrast, Newton's method may be used to solve the residual formulation.

IMPLEMENTATION DETAILS

The simulation framework is implemented in the Julia 1.10 language and executed on an Intel Xeon E5-2650 (2.2 GHz) workstation. Model derivatives are computed by combining analytical derivatives of the finite-volume discretization with forward-mode automatic differentiation of the thermodynamic functions provided by `Clapeyron.jl` [6]. Dense linear algebra is performed with Intel MKL-accelerated BLAS, while sparse model Jacobians are assembled explicitly and factorized using sparse direct linear algebra within the Newton schemes used for consistent initialization and ESDIRK time integration. Sensitivities of the half-cycle propagators are computed via staggered direct forward sensitivity analysis of the ESDIRK method and are used in the Newton-based single shooting solver.

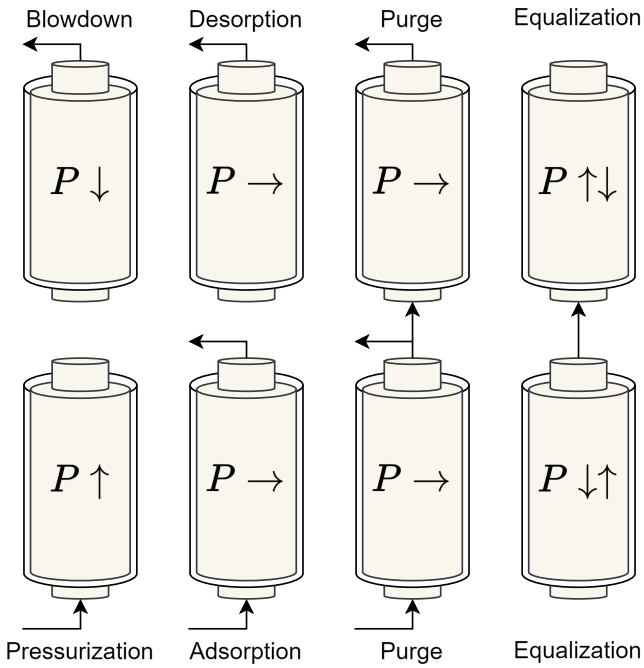

Figure 2. Diagram of the 8-step PSA cycle.

Table 1. Process conditions of the 4-step PSA half-cycle.

Step	**PR**	**ADS**	**PUR**	**EQ**
Time [s]	20	20	45	5
P_{ads} [bar]	8	8	8	4.5
P_{des} [bar]	1	1	1	8
α [-]	0	0	0	0
β [-]	0	0	0.5	1
φ [-]	1	1	1	0

Table 2. Column and feed parameters.

L [m]	D [m]	T_{in} [K]	P_{in} [bar]	v_{in} [cm/s]
2	1	298	8	5

Table 3. Simulated cycle performance (Tol=10^{-8}).

Statistic	EOS	K=10	K=20	K=40	K=80
Purity [%]	Ideal	95.4%	95.8%	96.0%	96.2%
	Real	95.3%	95.8%	96.1%	96.1%
Recov. [%]	Ideal	45.9%	45.9%	45.9%	45.9%
	Real	45.7%	45.7%	45.7%	45.8%

NUMERICAL EXPERIMENTS

We consider simulation of the example two-bed PSA cycle described in Table 1 and shown in Figure 2. The columns are fed dry air at ambient temperature, supplied at constant pressure and flow rate. Column and feed parameters are summarized in Table 2. Studying a converged CSS, we first assess how spatial resolution and thermodynamics affect simulated cycle performance. Using tight integration tolerances, purity and recovery are compared for ideal and real-fluid thermodynamics across increasing finite-volume resolutions. As shown in Table 3, the predicted product purity increases with spatial resolution, whereas the recovery remains essentially unchanged. Differences between ideal and real-fluid thermodynamics are negligible, on the order of 1–2‰, indicating that spatial discretization effects dominate.

The dominant computational task is governed by time integration, and in particular thermodynamic property evaluation. Table 4 compares several ESDIRK schemes from [5] for CSS simulation across error tolerances. While higher-order methods require fewer time steps, their larger stage counts make lower-order schemes slightly faster at loose tolerances. However, higher-order schemes are more efficient at stricter tolerances. It is worth noting that simulation scales linearly in K, since we employ sparse linear algebra. Accordingly, since model evaluation is dominated by the thermodynamics, ideal conditions are roughly 30-50% faster to simulate across all spatial resolutions.

Beyond the cost of individual cycle simulations, the computation of the CSS hinges on how efficiently the CSS BVP is solved. Using the 4th order ESDIRK method, Table 5 compares Picard iterations, Anderson acceleration, and Newton-based solution. Newton's method converges in only three iterations with quadratic convergence, achieving a tighter solution tolerance in minutes, whereas Picard and Anderson-accelerated iterations require hundreds of iterations and substantially longer runtimes. The dominant eigenvalue of the propagation operator lies within the integration tolerance of unity, making fixed-point convergence noisy, necessitating tight integration tolerances.

Table 4: CSS simulation statistics (EOS=Real, K=80).

Method	Statistic	Tol=10^{-2}	Tol=10^{-4}	Tol=10^{-6}	Tol=10^{-8}
ESDIRK2(1)3L[2]SA	Steps [#]	178	266	1589	14542
2nd order 3 stages	Time [min:sec]	**00:04**	**00:06**	00:31	04:43
ESDIRK3(2)5L[2]SA	Steps [#]	143	163	282	725
3rd order 5 stages	Time [min:sec]	00:05	**00:06**	**00:10**	00:27
ESDIRK4(3)7L[2]SA	Steps [#]	**129**	**144**	**233**	**377**
4th order 7 stages	Time [min:sec]	00:06	00:07	00:11	**00:20**

Table 5. Solver statistics (EOS=Real, K=40, Tol=10^{-6}).

Statistic	**Picard**	**AA**	**Newton**
Iterations [#]	1123	416	**3**
Time [hr:min]	01:19	00:23	**00:02**

Finally, the two half-cycles are stitched together to recover the full periodic operation and simulated spatio-temporal composition profiles are shown in Figure 3. The resulting product flow rate of approximately 450 [m^3/h] at ambient conditions falls well within the range of small-scale PSA-based air separation units.

CONCLUSION

This work presented a first-principles, dynamic modeling framework for PSA-based nitrogen generation. Thermodynamically consistent column models, combined with finite-volume discretization and implicit time integration, enable systematic simulation and reliable CSS computation via shooting methods. Numerical experiments show that solver choice strongly affects computational efficiency, with Newton-based methods converging substantially faster than fixed-point approaches. The framework provides a systematic basis for optimization of PSA cycles for nitrogen generation, following [3].

DIGITAL SUPPLEMENTARY MATERIAL

Adsorbent and adsorption rate parameters are available in the LAPSE at LAPSE:2026.0016.

ACKNOWLEDGEMENTS

This project has been funded by the MissionGreen-Fuels project DynFlex under The Innovation Fund Denmark project no. 1150-00001B.

AUTHOR IDENTIFIERS

Author ORCIDs:
Schytt MJ: 0009-0001-8143-1878
Biegler LT: 0000-0003-3875-4441
Jørgensen JB: 0000-0001-9799-2808

REFERENCES

1. Noordende HV, Tubben G, Rouwenhorst KH, Fruytier M. Clean Ammonia Roadmap. ISPT (2024).
2. Jee JG, Kim MB, Lee CH. Pressure swing adsorption processes to purify oxygen using a carbon molecular sieve. Chemical Engineering Science 60:869-882 (2005). https://doi.org/10.1016/j.ces.2004.09.050
3. Dowling AW, Vetukuri SRR, Biegler LT. Large?scale optimization strategies for pressure swing adsorption cycle synthesis. AIChE Journal 58:3777-3791 (2012). https://doi.org/10.1002/aic.13928
4. Hesthaven JS. Numerical methods for conservation laws. Society for Industrial and Applied Mathematics (2018). https://doi.org/10.1137/1.9781611975109
5. Kennedy CA, Carpenter MH. Diagonally Implicit Runge-Kutta Methods for Ordinary Differential

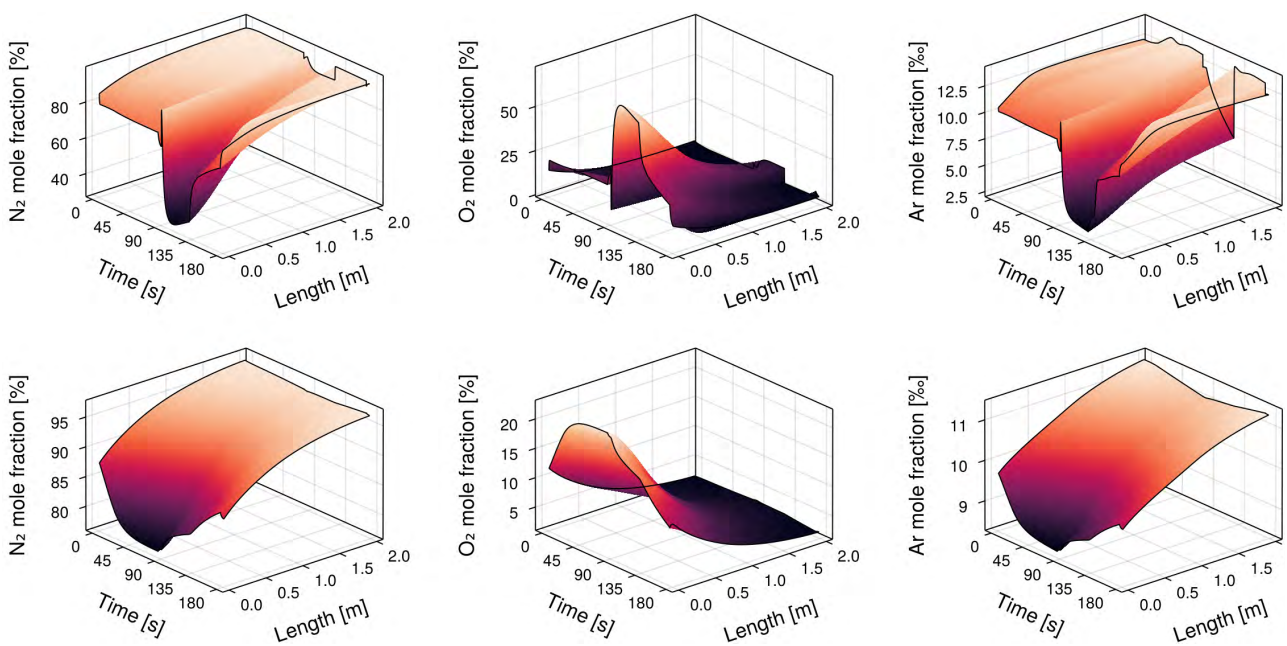

Figure 3. Fraction profiles throughout the 8-step PSA cycle at CSS (top: mobile, bottom: stationary).

Equations: A Review. NASA (2016).
6. Walker PJ, Yew HW, Riedemann A. Clapeyron.jl: an extensible, open-source fluid thermodynamics toolkit. Ind. Eng. Chem. Res. 61:7130-7153 (2022). https://doi.org/10.1021/acs.iecr.2c00326

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Energy recovery from process purges: steam turbine integration and operation optimisation in biogas upgrading within SEMPRE-BIO project

Filippo Bisotti[a]*, Matteo Gilardi[a]*, and Bernd Wittgens[a]*

[a] SINTEF Industry – Process Technology – P.O. Box 4760 Torgarden, 7465, Trondheim, Norway
* Corresponding Authors: filippo.bisotti@sintef.no, matteo.gilardi@sintef.no, bernd.wittgens@sintef.no

ABSTRACT

The SEMPRE-BIO project tested and validated three different innovative technologies and pilots within the Horizon Europe framework. One of the pilots is commissioned in Belgium. The proposed technology purifies biogas from manure anaerobic digestion and delivers simultaneously pure biomethane and food-grade CO_2, conversely to conventional purification technologies such as absorption and adsorption. Due to the severe cryogenic conditions, energy recovery for purge and waste streams becomes relevant to improve the energy demand of the process. The present work will show an effective solution to reduce the electricity demand of the process. Biomethane slip and other purge stream are valorised in a steam boiler and a two-pressure-level steam turbine to deliver both middle pressure steam as utility in distillation reboilers and produce electricity. The analysis will propose a simple, but rigorous methodology to maximise the steam turbine loop and the net power. The present work is based on solid and concrete assumptions for the design of the steam turbine. The analysis shows that the process allows for recovering up to 12% of the total electricity and around 55% of the total energy demand for the refrigeration loop to preserve the cold box of the bio-CH_4/CO_2 distillation.

Keywords: Turbines, Energy Efficiency, Combustion Heat and Power (CHP), Energy recovery, Biogas upgrade

INTRODUCTION

Biogas appears as a promising feedstock for combustion heat and power (CHP) applications [1, 2] as well as for replacing fossil methane in chemical production [3]. Biogas is a wet mixture of biomethane (bio-CH_4) and biogenic CO_2 (bio-CO_2) with a variable ratio of the two main components. Generally, bio-CH_4 ranges from 45 vol% to 75 vol% [4]. The remaining is mainly wet bio-CO_2. Prior to any application and transport, biogas undergoes deep dehydration and purification/upgrading by removing bio-CO_2 [5]. In addition, bacteria releases H_2S and NH_3 in traces during the anaerobic digestion of biological wastes (food residue, sludges, wastewater, crops, and manure) [4, 5]. These species are harmful and corrosive. Hence, biogas purification is energy intensive [6]. The bulk capture of CO_2 and successive dehydration of bio-CH_4 consumes heat, cooling water, and electricity [7]. This work will introduce a novel concept for biogas upgrading patented by Cryo Inox S.L. and tested within SEMPRE-BIO, an EU project under the EU's Horizon Europe Program. The site capacity and biogas productivity are sensitive. The analysis shows that valorising side streams covers part of the electricity energy and satisfy the total demand for thermal energy of the plant.

MATERIALS AND METHODS

Process block flow diagram (BFD)

Figure 1 depicts the three different process layouts considered in the analysis. The proposed BFDs differ for the "cold box" and how refrigeration is performed. The remaining of the biogas upgrading is kept unchanged. Regardless of the process configuration, manure (S1-01) is digested resulting in the slurry digestate (S1-04) and raw biogas (S1-07). The produced biogas is purified (H_2S and NH_3 removal) while dehydration and compression are carried out. The compressed biogas (S1-19) still contains

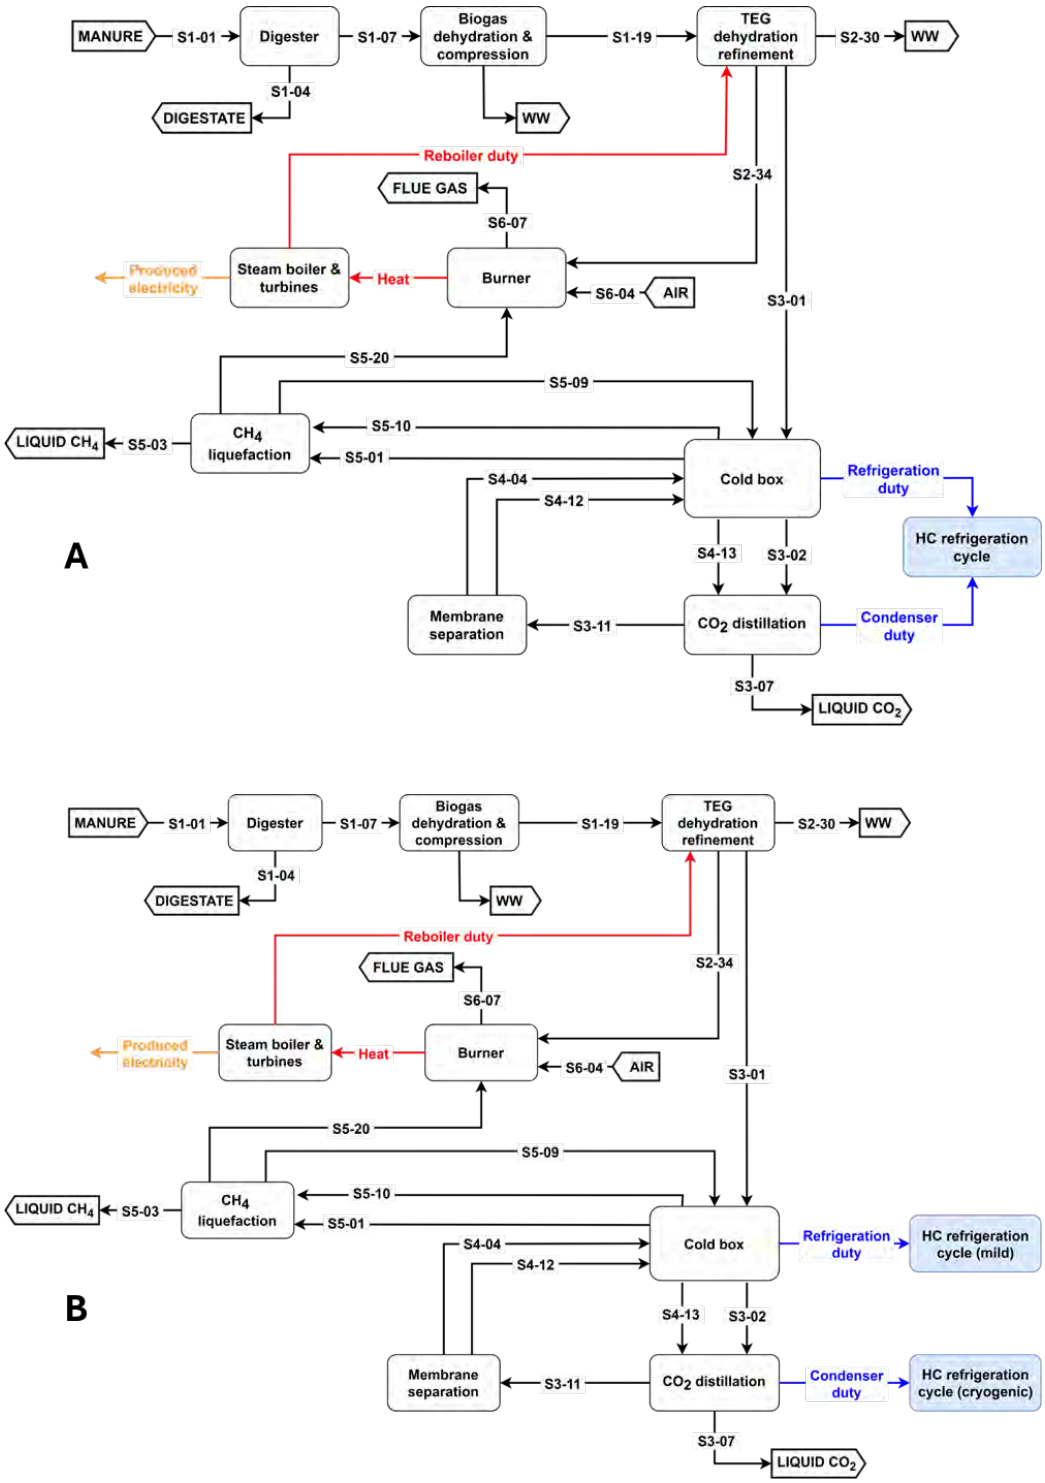

Figure 1: Process BFDs for three different configurations of the analysed process. The processes differ for the cold box configuration and how cooling (blue arrows) is managed within the process. The remaining part of the process remains the same. WW stands for wastewater.

a relevant residual moisture. The moisture content is refined through a conventional triethylene glycol (TEG) absorption process to keep the residue water content below 3 ppm (vol). The dry biogas (S3-01) leaves the absorber. As the absorber is run at high pressure and low temperature to abate residue moisture, bio-CH_4 and bio-CO_2 accumulates in the solvent [8, 9], which are easily released as heat is supplied. TEG regeneration is run

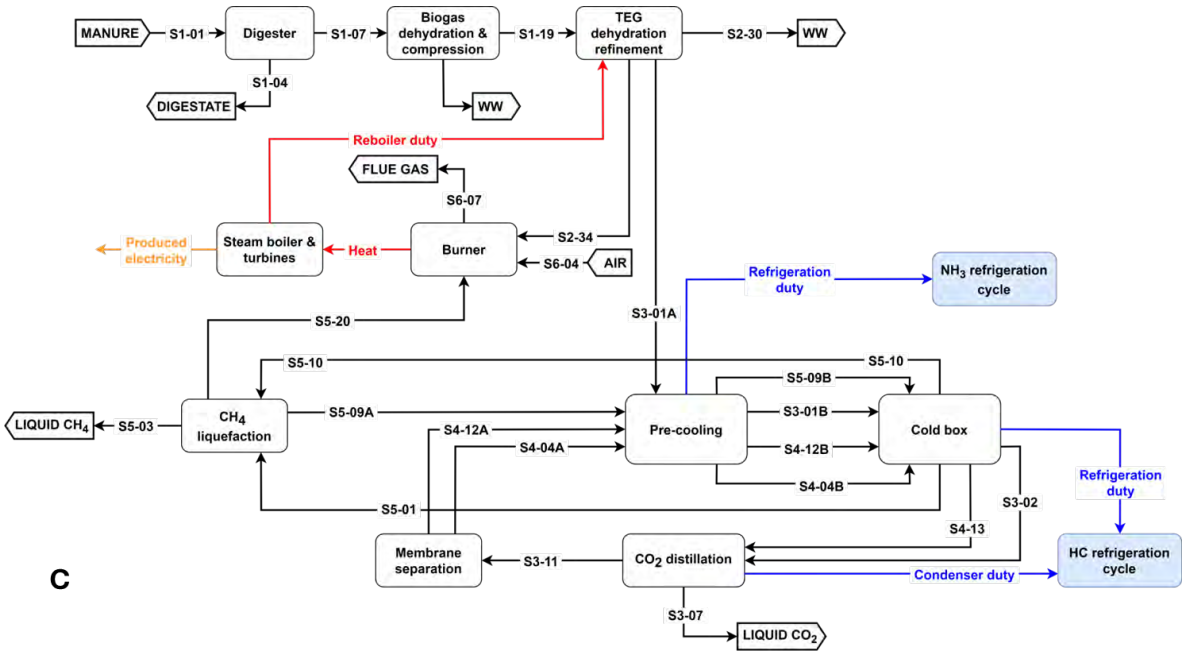

Figure 1. (continues):

thermally, and the process require steams. However, the temperature in the reboiler of the TEG regeneration is kept below 207°C to prevent it from degrading. To keep a minimum temperature approach ~10°C and avoid formation of hotspots in the reboiler, steam must be supplied around 219-221°C. A bio-CH_4-rich stream is released from the top of the stripper of the TEG regeneration. This stream is post-treated in a membrane to recover part of bio-CH_4. Pure methane is recycled back to the stripper after recompression. Injecting hot methane as stripping gas in TEG regeneration column is an industrial practice to increase the TEG purification from 98 to 99.9+ wt%, while lowering the thermal demand for the regeneration and reducing the temperature at the bottom [10, 11]. This is reflected in a minimum TEG degradation and liquid-to-gas ratio in the absorber to reduce the moisture content in the gas phase below 3 ppm (vol). The unrecovered bio-CH_4 (S2-34) is a purge and valorised in the burned to deliver steam for the steam turbine.

The dry biogas (S3-01) is cooled, and the resulting cold biogas (S3-02) is conveyed to the cryogenic distillation, whose operating conditions, namely, temperature and pressure, are sensitive. Part of the bio-CO_2 is recovered as pure liquid (>99.99 vol%) from the bottom, while bio-CH_4-enriched gas stream (S3-11 with CO_2 ~80 vol%) is sent to a membrane separation system where residue CO_2 selectively permeates and the gas stream is progressively purified until the outlet gas (S4-04) reaches a residue CO_2 content below 50 ppm (vol), which respects the specifications for grid injection. The membrane system works at ambient temperature. The permeated gas is enriched in CO_2 but still contains bio-CH_4 as membrane preferably permeate CO_2, but the separation is not sharp. Part of the permeated gas is recycled internally, while the purge (S4-12) is cooled and recycled to the cryogenic distillation to maximise both bio-CH_4 recovery and liquid bio-CO_2 production. Pure bio-CH_4 is cooled and conveyed to CH_4 liquefaction process (S5-01) where it is compressed and partially expanded to exploit Joule-Thomson effect and generate the frigories to liquefy part of the bio-CH_4 (S5-03). The not liquefied fraction (S5-20) is purged. The purged streams, namely S2-34 (from TEG dehydration) and S5-20 (from bio-CH_4 liquefaction), are transferred to a burned where excess air is co-fed. The combustion of purge streams is used to produce heat (steam) and power. Part of the produced steam is used for the regeneration of TEG solvent. The implementation and optimisation of combined steam boiler and turbine is described in the next sections.

Refrigeration strategies

Figure 1 shows three different layouts depicting three alternatives to manage the refrigeration and keep the cold box of the system. The basic process is described in the previous section, and it is kept equal regardless of the refrigeration strategy and configuration. The cold box satisfies the frigories demand in the cryogenic distillation (bio-CH_4 purification through bio-CO_2 liquefaction) and liquefaction of purified bio-CH_4. For the first, pre-cooling of the dry biogas and reaching cryogenic temperature in the top condenser of the distillation column (condenser duty in Figure 1) are needed. For the latter, cryogenic conditions are required to reduce the inefficiency. Indeed, if frigories are recovered exclusively

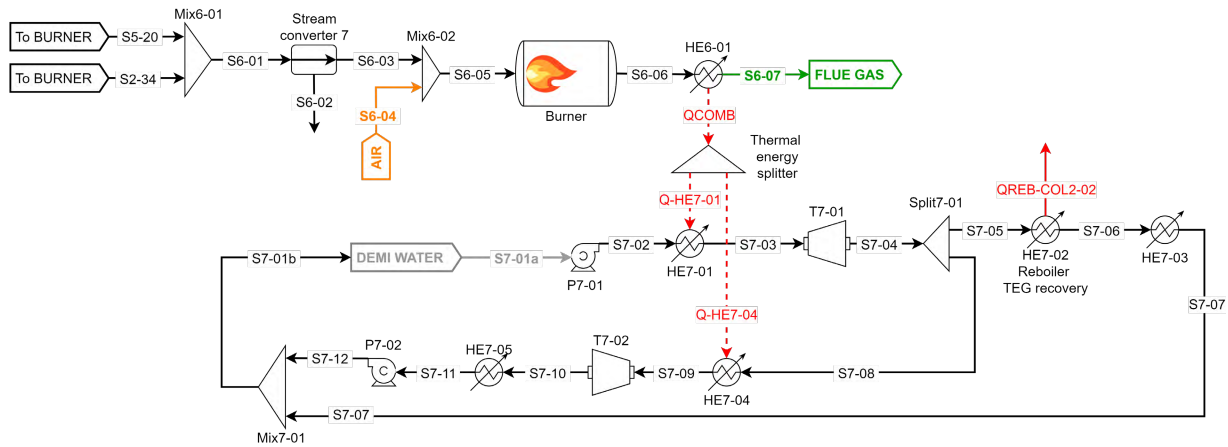

Figure 2. Combustion chamber and steam boiler with corresponding dual-pressure steam loop

using the liquefied bio-CH_4 without outsourced cooling duty, the net production of liquid bio-CH_4 would be reduced. Figure 1 distinguishes between the duty for the condenser of the cryogenic distillation and all the pre-cooling demand distributed around the plant. This distinction will be clarified later. Based on these considerations, three different layouts have been identified:

- Option A (single loop): the total cooling demand is covered using a hydrocarbon (HC) mixture as working fluid. The HC blend contains C_2-C_3 hydrocarbons. The relative ratio of C_2-C_3 has been tuned to keep the cold box minimum temperature (sensitive).
- Option B (double loop): the same mixture is used as working fluid but split into two independent refrigeration loops. In the first refrigeration loop (named as mild loop), all warm streams are precooled from ambient temperature (25-30°C) to reach the same intermediate temperature of the gas leaving the top of the cryogenic column. This temperature is sensitive and here not reported. In the second refrigeration loop (named as cryo loop) all the streams are cooled to final desired temperature (sensitive). This option is justified by thermodynamic constraints. Warmer streams allow for having higher outlet temperature of the cooling fluid leaving the evaporator. Conversely, colder streams are more challenging to further refrigerate because of the minimum temperature approach in the condenser heat exchanger. Then, the milder loop would treat a lower amount of gas because the gas from the cryogenic distillation is treated and cooled only in the cryo loop.
- Option C (double loop with two refrigerants): is like Option B, but all warm streams are pre-cooled to the same intermediate temperature using ammonia refrigerant for the mild loop. The temperature is reduced to the final target value in a second refrigeration cycle still using the same HC mixture as for options A and B.

Steam boiler and turbine

Process description

Figure 2 depicts the configuration of the burner and a conventional dual-pressure steam turbine system. Cold water (S7-01a) is pumped (P7-01), superheated (HE7-01) and conveyed to the first steam turbine T7-01. Expanded steam leaves the unit and it is split into two flows: S7-05 and S7-08. S7-05 is the steam supplied to the reboiler of the TEG regeneration. The resulting condensed steam (S7-06) is further cooled (HE7-03) and recycled back. The other stream from the split (S7-08) is heated a second time and sent to a second expander (T7-02). The low-pressure steam is then fully condensed (HE7-05), pumped (P7-02) and mixed with S7-07. The total condensed water is then recycled to the main pump (P7-01). This work aims to optimising the steam cycle to maximise the net work recovered from the combustion of bio-CH_4 rich purges defined as

$$W_{Net} = W_{T7-01} + W_{T7-02} - W_{P7-01} - W_{P7-02} \quad (1)$$

where W stands for the power generated (for turbines) and consumed (for pump) and consumed (for pump).

Input data and Degrees of Freedom (DoF) of the optimisation

The combustion of bio-CH_4 purges releases 96.04 kW_{th}. This is the total energy input to the steam cycle. The steam cycle is designed to deliver 39.12 kW_{th} to the reboiler of the TEG regeneration. As aforementioned, steam should be supplied saturated at 220°C, i.e., at 23.5 bar (absolute). The steam quality is stringent due to TEG thermal stability. The degrees of freedom of the system are: (1) the total flow of water/steam circulating in the cycle (S7-01a), (2) the upper pressure of the steam cycle,

(3) the thermal duty (Q-HE07-01) supplied to the first superheater (HE07-01), and (4) the discharge pressure of the second steam turbine (T7-02). These variables correspond to the optimisation parameters.

The system has been simulated in Aspen Hysys V12 using ASME steam thermodynamic model. The optimisation has been carried out using the sensitivity analysis tool in the simulation environment. For pump and turbines, a global efficiency, i.e., product of polytropic and mechanical efficiency, has been assumed. The following ranges have been investigated for the optimisation parameters:

Total mass flow: 95 – 115 kg/h with a step of 5 kg/h. Our investigation included 98, 107, and 112 kg/h to refine the domain and further optimise the process.

(1) Maximum/upper pressure of the steam cycle: ranges between 60 – 120 bar (absolute), typical for high-efficiency system for heat recovery and electricity production [12, 13]. 10 bar is the considered discretization step. The lower and upper limited has been identified based on steam turbine producers' suggestions and standard, i.e., not customised, units.

(2) Superheater Q-HE7-01 duty: ranges 90 – 96 kW_{th} and energy balances imposes that

$$Q_{HE7-04} = 96.04 - Q_{HE7-01} \quad (2)$$

(3) Discharge pressure of the second turbine T7-02: as been limited to 0.3 – 1.2 bar (absolute) with a step of 0.3 bar. The bottom and upper limits are retrieved from technical reports and producers [14]. It is unlikely to lower the discharge pressure below 30 kPa. It is possible to run the turbine above 1.2 bar absolute, however, this is normally done in district heating and not in steam boiler and turbine loops where the main target is the energy recover.

Constraints and criteria

The results have been screened and filtered based on the following criteria (checkpoint list) and constraints:

1. The maximum inlet temperature of superheated steam to turbines (S7-03) is 600°C to avoid thermal stress [12, 15], i.e., periodic maintenance, and expensive construction materials [15]. This constraint links the duty to the first superheater (HE07-01) and the total water mass flow circulating in the loop. If the inlet temperature exceeds this threshold, i.e., T_{S7-03} > 600°C, the operating point is unfeasible. The same constraint holds for the steam conveyed to the second steam turbine (T7-02), hence T_{S7-09} < 600°C must be fulfilled.

2. S7-05 is the saturated steam for the reboiler of TEG regeneration. The maximum temperature of the steam is assumed 220°C to keep a minimum temperature approach of ~10°C in the reboiler and the avoid hotspots in the unit as TEG degrades above 207°C. Setting the steam pressure implies that the discharge pressure of T7-01 is 23.5 bar. Moreover, the reboiler duty for the TEG regeneration (QREB-COL02-02) is assigned; indeed

$$m_{S7-04} = \frac{Q_{REB-COL2-02}}{\Delta H_{EV}(T_{sat}, P_{sat}(T_{sat}))} \quad (3.1)$$

$$m_{S7-04} = \frac{39.12}{\Delta H_{EV}(220°C, 23.5\ bar)} \quad (3.2)$$

3. From the previous point it follows that the temperature of the superheated steam leaving the turbine T7-01 (S7-04) could be warmer than the one selected to deliver steam to the TEG reboiler. Hence, the optimal operational point should deliver steam as much as possible to the saturation point to avoid a kind of transfer line heat exchanger to remove the "excess" heat (Q_{excess}). This device would increase the inefficiency of the system.

4. The lower pressure of the loop, i.e., discharge pressure of the second turbine T7-02, is an optimisation variable and DoF of the system. The discharge pressure is constrained to the physical state of the steam leaving the unit. If steam starts condensing, the expander cannot be operated. Hence, the discharge pressure is selected as the lowest pressure at which only steam leave the unit, i.e., vapour fraction α = 1. This condition preserves the operation of the turbine while ensuring to recover the maximum work. It is straightforward, that the same condition must be satisfied also at the outlet of first turbine (T7-01). This condition is less stringent, unless to large water mass flow is processed in the cycle.

RESULTS DISCUSSION

Steam cycle optimisation

Table 1 gathers the results from the sensitivity analysis performed in Aspen Plus. Several feasible points have been identified; only the optimal one per each tested total mass flow of steam/water in the loop are reported. The optimal points are selected based on the maximum net work recovered from the loop. The trend of net work recovered from the steam cycle is depicted in Figure 3. The results show that the steam cycle can be operated only within a limited domain for the steam/water mass flows in the loop: between 98 kg/h (minimum) and 112 kg/h (maximum). Indeed, below 98 kg/h, the total fluid is too little and the inlet temperature of superheated steam either to the first or to second turbine violates the threshold value of 600°C. Under this condition, the combustion releases an excessive heat. Conversely, above 112 kg/h, the total fluid is too large: the first superheater requires a substantial fraction of the heat released from the combustion, but the remaining heat is insufficient to reheat the steam after the splitter.

Table 1: Optimal points for the steam cycle

Steam flow (kg/h)	P max (bar abs)	P min (bar abs)	Q_{HE7-01} (kW)	Q_{excess} (kW)	Inlet T (°C) T7-01 / T7-02	Net work (kW)
95	-	-	-	-	-	-
98	120	0.4	92.5	6.07	588 / 588	19.11
100	120	0.4	92.0	4.80	554 / 573	19.10
105	**120**	**0.4**	**90.0**	**1.46**	**468 / 566**	**19.04**
107	120	0.6	90.0	0.60	448 / 526	18.34
110	100	0.9	90.0	0.05	407 / 487	16.64
112	80	1.0	90.0	0.05	374 / 472	15.26
115	-	-	-	-	-	-

Table 2: Comparison of the refrigeration strategies

Option	P_{tot} (kW)	$P_{tot-net}$ (kW)	Compression refrigerant (kW)	$Compr\%_{tot}$	Electricity Saving (ES)	ΔE_{tot}	ΔE_{net}
A	191.2	172.2	70.8	37.0%	-9.94%	Reference	Reference
B	186.2	167.2	66.1 (total)	35.5%	-10.2%	-2.6%	-2.9%
C	181.9	162.9	61.9 (total)	34.0%	-10.4%	-4.9%	-5.4%

This results in incipient condensation of the steam even above the atmospheric pressure (1.2 bar absolute has been considered as the upper limit). In other words, the system experiences a lack of thermal energy. This explanation justifies also the reduction of the upper pressure of the steam loop, which progressively reduces from 120 bar (maximum) to 80 bar. This is the result of the combination of saturation temperature (higher as pressure increases) and latent of vaporization (lower as pressure increases). Figure 3 displays an almost flat profile up 105 kg/h and then a sharp drop of the net work recovered. The average net work recovery is around 19 kW for the flat area. The reduction of the net work is substantial beyond 105 kg/h. Indeed, for an increment of the mass flow from 105 to 112 kg/h (+6.25%), the work recovery falls from 19 kW to 15.3 kW (-19%). The physics behind the system is straightforward and already commented above, when describing the boundaries of the feasible domain to run the steam cycle. The optimal operational point has been selected by globally accounting for the criteria listed and described in the dedicated section. In particular, the steam cycle run with 105 kg/h is considered the best option because of the limited excess heat, meaning that the steam withdrawn after the first expansion is almost close to saturation and no quench is required. In addition, in that configuration, superheated steam is fed to both turbines at milder temperatures, i.e., lower than the maximum allowable one (600°C), lowering the risk of mechanical failure.

Effect of refrigeration strategies and steam cycle integration in the biogas upgrading

In addition, the system has been analysed in terms of refrigerant demand and total energy demand to identify the best cold box configuration and the benefit associated with the addition of a steam turbine system in the base process. Based on the simulation of the whole process under the different options (Figure 1). The total cooling demand is around 49 kW_{th}, split into 19.3 kW_{th} for the condenser at the cryogenic distillation and 29.7 kW_{th} to provide frigories in other sections of the process.

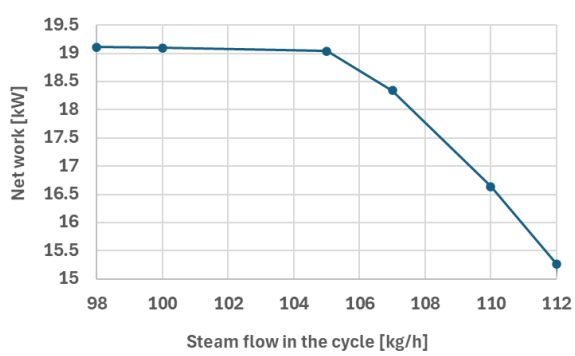

Figure 3. Locus of optimal points

Table 2 gathers the results from the simulation. The results show a progressive reduction in the fluid demand and the corresponding compression costs. Hence, the possibility to split the refrigeration into two independent loops is a benefit. Moreover, the possibility to run the mild loop with ammonia as refrigerant improves further the performance. Indeed, the pre-cooling with ammonia and reduces the total electricity demand to operate the refrigeration cycle. At this point, it is worth contextualising the role of the energy recovery and cold box optimisation in the plant for biogas upgrading and liquefaction.

Additional data are necessary to carry out this last assessment. Table 3 resumes the results for all the simulations. The total electricity, in Eq (4), power demand (P_{tot}) accounts for compressors and pumps and it has been estimated based on the simulations; while the net power ($P_{tot-net}$) is defined as the difference between the total power (P_{tot}) and the net total power/work recovered from the steam boiler and turbine system (W_{net}), as defined in Eq (1). All entities are expressed in kilowatt (kW).

$$P_{tot-net} = P_{tot} - W_{net} \quad (4)$$

It is important to highlight that the previous analysis concluded that the maximum net power (W_{net}) recovered from the turbine loop is around 19 kW (Table 1). The analysis considers the following key performance indicators (KPIs):

- Percentage of energy demand of compression of refrigerant(s) on the total electricity demand:

$$Compr\%_{tot} = \frac{W_{refrigerant\ compression}}{P_{tot}} \quad (5)$$

where the numerator is provided in Tables 2 and 3.

- The electricity saving defined as:

$$ES = \frac{W_{net}}{P_{tot}} \quad (6)$$

- Energy reduction factors (both using as reference the total and net power of Option A) calculated as:

$$\Delta E_{tot} = \frac{P_{tot} - P_{tot}(option\ A)}{P_{tot}(option\ A)} \quad (7)$$

$$\Delta E_{net} = \frac{P_{tot-net} - P_{tot-net}(option\ A)}{P_{tot-net}(option\ A)} \quad (8)$$

Table 3 shows that energy demand for the compression of refrigerant(s) represents one of the most substantial energy consumptions for running the purification of biogas. The implementation of a steam turbine enables to cover almost 10% of the total electricity of the system. In addition, Options B and C introduce additional improvements to the benchmark (Option A). These two configurations further reduce the total energy demand thanks to a better layout of the refrigeration loop (Option B) and selection of the refrigerants (Option C). The energy reduction with regards to the benchmark is still limited, but the electricity power demand drops up to -4.9% and –
5.4% if you consider the either the total or the net energy demand as baseline, respectively. Though the energy savings could sound limited, but the steam cycle delivers 100% of the steam demand for the entire process. The end user does not have to burn extra methane or outsource fuel, i.e., biomass waste, to integrate the amount of steam. Beyond the technical aspect and comments related to the energy analysis of the system, it is important to stress that also economic aspects must be considered to identify the best process configuration. However, this is out of the scope of the present paper.

CONCLUSIONS AND FUTURE WORKS

The paper presented a novel technology for the upgrading of the biogas. Remarkably, the system allows for purifying bio-CH_4, while producing bio-CO_2 with purity of both products 99.99+ vol%, meeting specs for injection in pipeline (methane) and carbon use, including food and beverage, or permanent storage (CO_2). The proposed technology, patented by Cryo Inox, represents an example of full electrification of biogas upgrading. The system includes the bio-CH_4 liquefaction. Liquefaction can be considered for small clusters of bio-CH_4 producers. Indeed, liquid bio-CH_4 can be transported. Hence, the proposed process is ideal whenever natural gas pipelines are far from the biogas site and/or connecting to existing infrastructure is not feasible. The work showed that the recovery and valorisation of waste stream coming from the process becomes relevant to improve the key performance of the system. For instance, the valorisation of purges containing bio-CH_4 can be exploited to generate steam and heat. The energy analysis highlighted that the integration of the combustion chamber with a steam boiler and a steam cycle can benefit to the upgrading process. The optimization of the fully integrated steam cycle enables to recover up to 10% of the total electrical energy demand of the system and satisfy 100% of the total thermal energy demand (steam). In addition, the paper also showed the optimisation of the refrigeration strategy, and its configuration have an impact on the overall

Table 3: Energy analysis of the three options.

Option	HC mixture (kg/h)	NH_3 (kg/h)	Refrigerant reduction	HC mixture compression (kW)	NH_3 fluid Compression (kW)	Compression reduction
A	800	NA	Reference	70.8	NA	Reference
B	392.5 (mild) 326.7 (cryo) 755.2 (total)	NA	-5.6%	37.5 (mild loop) 22.6 (cryo loop) 66.1 (total)	NA	-6.2%
C	637.0 (cryo)	34.0 (mild)	-16.1%	yo loop)	5.90 (mild loop)	-13.0%

indicators. Three configurations to supply frigories have been investigated: the selection of both the process layout and the proper refrigeration affects the performance. Beyond these technical aspects, it is worth underlining that the analysis should not be limited to "engineering" aspect, i.e., energy recovery and process optimisation, but integrate the economics assessment to identify the best options among the candidates and account for the impact of process configuration on the profitability. This is not included the present work, but it will be investigated within the SEMPRE-BIO project.

ACKNOWLEDGEMENT

This work received funding from the SEMPRE-BIO project (Link to project webpage and Link to EU portal page), funded within the EU's Horizon Europe programme under grant agreement Nº 101084297. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or CINEA. Neither the European Union nor the granting authority can be held responsible for them.

REFERENCES

1. Kabeyi MJB, Olanrewaju OA. Biogas production and applications in the sustainable energy transition. Journal of Energy 2022:1-43 (2022). https://doi.org/10.1155/2022/8750221
2. Kougias PG, Angelidaki I. Biogas and its opportunities—a review. Front. Environ. Sci. Eng. 12: (2018). https://doi.org/10.1007/s11783-018-1037-8
3. Deng R, Wu J, Huang Z, Feng Z, Hu W, Tang Y, Tan H, Zhang H, Zairov R, Pan Z. Biogas to chemicals: a review of the state-of-the-art conversion processes. Biomass Conv. Bioref. 15:14653-14673 (2024). https://doi.org/10.1007/s13399-024-06343-1
4. European Biogas Association. EBA Statistical Report 2025. https://www.europeanbiogas.eu/news/eba-statistical-report-2025/
5. International Energy Agency (IEA). An introduction to biogas and biomethane. https://www.iea.org/reports/outlook-for-biogas-and-biomethane-prospects-for-organic-growth/an-introduction-to-biogas-and-biomethane
6. Swinbourn R, Li C, Wang F. A comprehensive review on biomethane production from biogas separation and its techno?economic assessments. ChemSusChem 17: (2024). https://doi.org/10.1002/cssc.202400779
7. Yentekakis IV, Goula G. Biogas management: advanced utilization for production of renewable energy and added-value chemicals. Front. Environ. Sci. 5: (2017). https://doi.org/10.3389/fenvs.2017.00007
8. Rasoolzadeh A, Raeissi S, Shariati A, Peters CJ. Experimental Measurement and Thermodynamic Modeling of Methane Solubility in Triethylene Glycol within the Temperature Range of 343.16–444.95 K. *J. Chem. Eng. Data*, 65, 8, 3866–3874 (2020) https://pubs.acs.org/doi/10.1021/acs.jced.0c00198
9. Soroodan E, El-Okazy MA, Macario E, Kentish SE. Triethylene glycol solubility in high-pressure methane: new experimental data at 40 to 60 °C. Ind. Eng. Chem. Res. 63:2895-2900 (2024). https://doi.org/10.1021/acs.iecr.3c04002
10. Stewart M, Arnold K. Dehydration considerations. Gas Dehydration Field Manual :55-168 (2011). https://doi.org/10.1016/b978-1-85617-980-5.00002-1
11. Aimikhe VJ, Okologume CW. Environmental challenges of natural gas dehydration technologies. Advances in Natural Gas: Formation, Processing, and Applications :53-75 (2024). https://doi.org/10.1016/b978-0-443-19221-0.00011-9
12. Ohji A, Haraguchi M. Steam turbine cycles and cycle design optimization: the rankine cycle, thermal power cycles, and integrated gasification-combined cycle power plants. Advances in Steam Turbines for Modern Power Plants :11-40 (2022). https://doi.org/10.1016/b978-0-12-824359-6.00020-2
13. Hall S. Refrigeration. Branan's Rules of Thumb for Chemical Engineers :190-202 (2012). https://doi.org/10.1016/b978-0-12-387785-7.00011-6
14. U.S. Environmental Protection Agency (EPA) - Combined Heat and Power Partnership. Section 4. Technology Characterization - Steam Turbine. March 2025. https://www.epa.gov/sites/default/files/2015-07/documents/catalog_of_chp_technologies_section_4._technology_characterization_-_steam_turbines.pdf
15. Siemens Energy. https://www.siemens-energy.com/global/en/home/products-services/product-offerings/steam-turbines.html

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

ProcessSimulator.jl: A Symbolic-Numeric Open-Source Framework for Process Simulation in Julia Language

Vinicius V. Santana[a]*, Christopher V. Rackauckas[bc], Idelfonso Nogueira[a]*

[a] Norwegian University of Science and Technology, Department of Chemical Engineering, Trondheim, Norway
[b] Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory, JuliaLab, Boston, USA
[c] JuliaHub, Boston, USA
* Corresponding Authors: vinicius.santana@ntnu.no, Idelfonso.b.d.r.nogueira@ntnu.no .

ABSTRACT

This paper presents ProcessSimulator.jl, an open-source process simulation framework built in Julia that combines acausal, equation-oriented modelling with seamless integration of procedural code. The framework leverages ModelingToolkit.jl to enable modular construction of unit operations using symbolic–numeric representations, facilitating the extension of models with advanced thermodynamics, kinetics, and data-driven components. Inspired by the ModuSim concept [3], ProcessSimulator.jl introduces an extensible control-volume abstraction and connector-based composition at the unit-operation level. A steady-state CSTR case study is presented and compared against Aspen Plus, showing good agreement in key variables. The results demonstrate the feasibility of a flexible, open, and composable process simulation paradigm for research and education.

Keywords: Process Simulation, Acausal Modeling, Modularisation, Open Source Software, Julia Language

INTRODUCTION

Process simulation is a central component of process systems engineering, enabling the coupling of thermodynamics, transport, and reaction kinetics across interconnected unit operations for design, optimisation, and control of complex chemical systems. Current simulators have enabled remarkable advances in systems engineering since the 60s, thanks to prebuilt, configurable unit operations and thermophysical models. While prebuilt components are powerful, innovating and designing new processes requires going beyond them. In this sense, commercial simulators enable users to create unconventional unit operations through custom modelling languages. ASPEN Custom Modelling Language and gPROMS Model Builder are powerful tools for achieving this goal. One drawback, however, is that unconventional unit operations often require an interface between the modelling language (declarative), where the conservation laws are written, and a procedural language, where the thermodynamics, kinetics, and other calculation routines, such as those for Deep Learning and modern data-driven tools, are written. This interplay is often far from simple in commercial tooling, since such tooling is often designed to be self-contained.

We believe that reducing the gap and friction in this interplay is critical to the easy integration of emerging, complex process models. One way to achieve this is to run the simulator in a fully featured language that supports seamless interplay between declarative and procedural programming. Addressing these limitations requires a new simulation paradigm grounded in open, high-level, and composable programming frameworks that bridge symbolic and numeric representations at high performance.

Another highly desirable feature for process systems engineering simulators is modularisation at the unit operation level, as it readily enables modelling of non-standard problems. The idea of composing submodels into a unit operation that, when combined, forms a flowsheet has been part of process engineering research and practice since the 90s [1], [2], [3]. The ModuSim concept [3], for example, was proposed in the 90s along these lines. Concretely, the ideas presented in ModuSim are very amenable to the equation-oriented modelling paradigm, which partitions a system or unit operation into small sub-models and assembles them using acausal connectors. The connectors represent boundary

conditions spanning convective and diffusive transport across ports. This can be done at the unit operation level up to the flowsheet level. This is a highly flexible and powerful modelling concept for chemical engineering, especially when easily integrated into procedural code, where thermophysical and transport-property correlations are typically written (calculations that require complex code constructions – for, while, conditionals, data structures).

There exists open-source software for process simulation written in modern, fully featured programming languages built on different modelling concepts. IDAEs PSE [4] is a Python-based process modelling framework and model library for process simulation and optimisation. The software core stack has at its bottom the Pyomo Algebraic Modeling Language. On top, it contains a library of models for common process unit operations and thermophysical properties that can be used to rapidly create process systems. Modularity of unit operations is enabled by the object-oriented nature of the Python language through inheritance and method overloading, but it is not part of the language's inherent modelling philosophy, as it is in the Modelica environment, for example. NeqSim is a Java-based library for process simulation and fluid thermodynamics. It hosts a comprehensive library of unit operations and thermodynamic models, but it does not mention any constructs for extending models from basic elements. OMCHEMSim is an open-source library in OpenModelica supporting implementations of unit operations and thermodynamic models (within the Modelica language). At the moment, however, no open-source software is available for structured acausal network modelling with smooth integration with procedural coding for process and chemical engineering.

In this setting, we present ProcessSimulator.jl, an open-source package that aims to be a modern library for modularising process systems engineering models at the unit-operation level using acausal equation-oriented modelling, with a smooth interface to procedural code. The technologies that enable it are the open-source Julia programming language and the ModelingToolkit.jl (MTK.jl) package [5]. ModelingToolkit.jl is a symbolic-numerical package in Julia that enables the application of the concept of Physical Network Acausal modelling to build complex models by assembling components with connectors. The symbolic constructs of the components enable ModelingToolkit.jl to perform symbolic manipulations, such as automatic index reduction, variable transformation, and problem simplification, to provide the optimal form for a numerical solver.

ModelingToolkit.jl is inspired by the Modelica language and shares the same principles of acausal modelling; however, Modelica is not a fully featured programming language, which makes it difficult to express complex code constructs for efficient thermophysical property calculations. Consequently, attempts to construct a process simulation environment require rewriting thermodynamic calculations in Modelica [6] or using workarounds to access external packages written in other languages. Moreover, because ModelingToolkit is a Julia modelling language that generates Julia code (functions) using metaprogramming, it can be composed efficiently and smoothly with other libraries available in the language. Examples of such compositions are illustrated using the Clapeyron.jl package, which provides libraries for thermodynamic and transport properties.

Built-in compositions in ModelingToolkit.jl are also used to solve the resulting numerical problem, with DifferentialEquations.jl for dynamic simulations and NonlinearSolve.jl for steady-state problems. As in Modelica, there is a separation of concerns – the system is represented in a modelling language and further simplified by an algorithm to best suit the numerical solvers, leaving the modeller with the efforts to do what they know – to model systems.

It is worth noting that ProcessSimulator.jl is still in its early stages, but this publication aims to communicate its philosophy. We believe that with a well-structured philosophy, advancing the software becomes easier, since many other contributors can materialise the philosophy in code, a task that is cumbersome for a single person.

The benefits of creating a process simulator in a language such as Julia using the acausal modelling paradigm are several-fold. It is not only useful for custom and advanced modelling but also for educational purposes. The state of open-source scientific computing is advanced enough to simulate complex models in process and chemical engineering, especially in Julia. Using the freely available latent infrastructure to advance and expand computational tools in chemical engineering is therefore worth pursuing.

On the educational side, acausal-equation-oriented modelling prompts reflection on many important aspects that shape a good engineer - every part of the system must be placed through a rational, conscious decision: a system's degrees of freedom, what to include and what to exclude when modelling it. This is somewhat lost when engineers spend time configuring numerous boxes without a direct understanding of their impact on the model.

METHODOLOGY

Previous modularisation methodologies - The ModuSim

The constructs presented here aim to reflect many of the concepts and ideas in the ModuSim concept, using equation-oriented modelling. Both ModuSim and ModelingToolkit use different wording for similar concepts. Similarly, in ProcessSimulator.jl, we introduce additional vocabulary that shares meanings with many concepts in

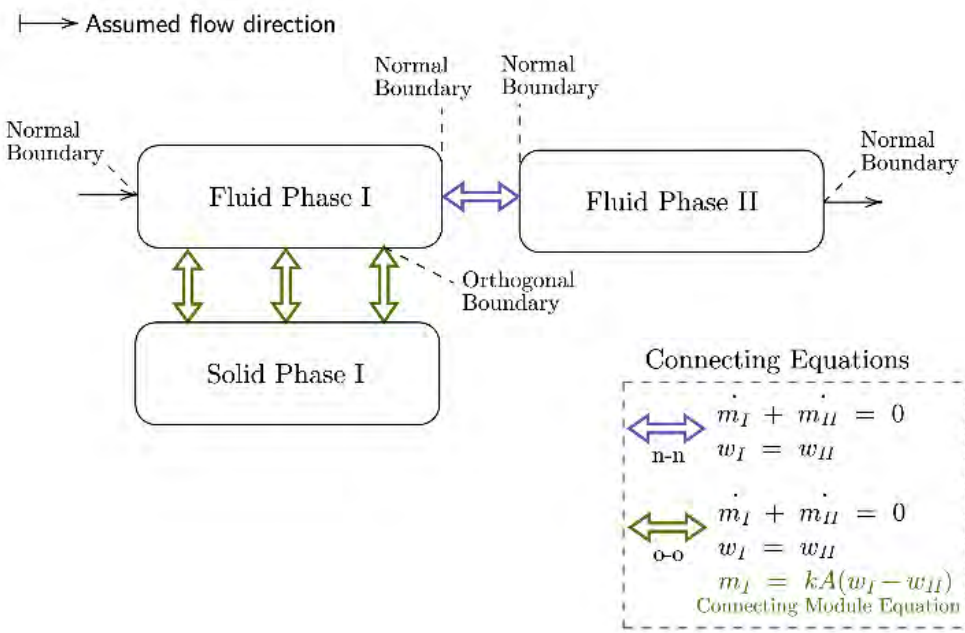

Figure 1. The ModuSim modularisation methodology. Adapted from [3].

ModuSim.

In ModuSim notation, there are two classes of main modelling objects: devices and connections. A device is a model object that may have a holdup of extensive quantities and that holds the relationships used to calculate state variables such as temperature, composition, and density. A connection creates equations that, when solved, calculate unknown fluxes from known variables at device boundaries. The ModuSim concept distinguishes between two types of connectors: normal and orthogonal boundaries. Normal boundaries point in the direction of the coordinate system (the direction of the convective flow), and orthogonal boundaries have no specified direction (depends on the potential for transfer in the system). The connection between two devices can only take place through a connector.

The acausal equation-oriented equivalent of ModuSim in ProcessSimulator.jl

As previously mentioned, the core modelling language in ProcessSimulator.jl is ModelingToolkit.jl (MTK.jl). Therefore, the equivalent ModuSim construction in ProcessSimulator.jl is implemented using MTK.jl.

While ModuSim is chemical-engineering-oriented, the ModelingToolkit.jl notation is generic. Therefore, effort is necessary to translate ModuSim concepts into MTK.jl. MTK.jl also has classes of modelling objects. A component in MTK.jl is a data structure that holds variables, parameters, equations, a name, and potentially other inner components, nested within it. In some sense, a connector in MTK.jl is one special kind of component that also holds variables, parameters and a name. The only difference is that connectors automatically generate equations depending on the type of connection one assigns to their variables. ProcessSimulator.jl uses two types of connector variables from MTK.jl: flow, regular connectors. By tagging a connector variable as flow, whenever two connectors are connected, it creates the equation:

$$w_I + w_{II} = 0 \tag{1}$$

Moreover, the regular connector tag creates the following equation:

$$v_I - v_{II} = 0 \tag{2}$$

Where w_I, w_{II}, v_I, v_{II} are the variables tagged as flow and regular variables from the connectors I and II.

It is possible to see that the ModuSim normal-normal (n-n) connecting equations and part of the orthogonal-orthogonal (o-o) are displayed in Figure 1 can be well represented with connector variables through equations (1) and (2). The connecting module equation seen in green on Figure 1 It depends on which types of interaction arise from the two phases in contact. In the Figure, it can, for example, be an adsorbent in contact with a carrier gas during an adsorption separation process. In these lines, the connecting equation can be a linear driving-force mass-transfer equation.

The ModuSim paper has not provided details on how to modularise devices, i.e., what is the minimum set of equations and variables that can be later extended to become conventional or advanced unit operations –

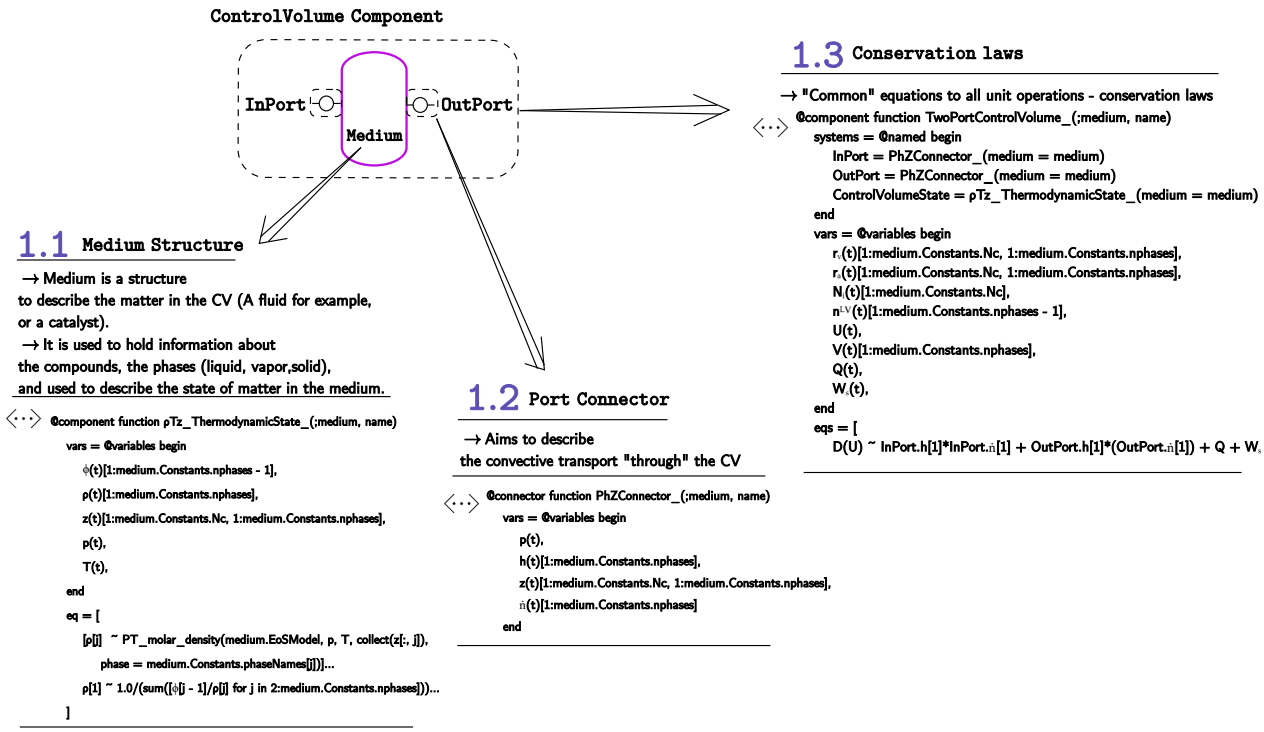

Figure 2. The ControlVolume modularisation strategy.

reactors, separators, etc. In ProcessSimulator.jl, we explored this aspect of the modularisation with the ControlVolume component. This is presented in the next section.

The ControlVolume component

The idea behind the control volume component (in MTK.jl notation) is to provide an extensible component that can be easily transformed into a specific unit operation. To achieve this goal, some reasoning on what the equations are that are common to almost all "conventional" unit operations is required.

To keep it simple, we start by analysing perfectly mixed, multicomponent, non-distributed systems and their governing conservation equations (mole component balance and energy):

$$\frac{d(N_i)}{dt} = \dot{m}_{in} z_{in,i} + \dot{m}_{out} z_{out,i} + \mathrm{r}_{i,V} V + r_{i,A} A \tag{3}$$

$$\frac{d(\mathrm{U})}{dt} = \dot{m}_{in} h_{in} + \dot{m}_{out} h_{out} + \dot{\mathrm{Q}} + \dot{W} \tag{4}$$

V is the control volume's volume, A is the control volume's area, U is the control volume internal energy, h_{in} and h_{out} are the inlet and outlet streams' molar enthalpies, $\dot{m}_{in,}$ and $\dot{m}_{out}$ are the overall inlet and outlet molar flow rates, $\dot{\mathrm{Q}}$ is the rate of heat to/from the control volume, $\dot{W}$ is the work rate done on the control volume by the environment. Additionally, N_i is the molar holdup, $z_{in,i}$ and $z_{out,i}$ are the inlet and outlet streams' mole fraction, $\mathrm{r}_{i,V}$ is the volumetric molar rate of generation, all with respect to the ith component.

In addition to the differential equations representing conservation laws, we need algebraic expressions for each of the terms of the equations, except N_i which is set by an initial condition.

By default, equations for $\mathrm{r}_{i,V}, r_{i,A}, \dot{\mathrm{Q}}, \dot{W}$ are not specified and are left to the user to specify according to the unit operation being modelled. For example, a heated stirred tank can be extended by including the expression $\mathrm{r}_{i,V} = 0, r_{i,A} = 0, \dot{\mathrm{Q}} = Q_0, \dot{W} = W_0$ where Q_0, W_0 are real numbers. A heated stirred tank reactor can be extended by replacing zero generation rate by $\mathrm{r}_{i,V} = kinetics(N_i, V, T)$.

All the equations associated with variables marked with the in Subscripts are derived from the control volume's connection to the nearest upstream components. These equations are defined in the ControlVolume using the InPort connector. InPort is of type PhZnConnector because it connects pressure, enthalpy, composition, and flow rate information. Pressure, compositions and enthalpy are regular connector variables – create equations like Equation (2). Flow rate is a flow variable; it yields equations like Equation (1).

When it comes to outlet variables (out subscripts),

only the outlet flow rate ($\dot{m}_{out}$) information depends on downstream equipment. h_{out} is calculated from the control volume's state and "transmitted" downstream. The ControlVolume has an OutPort of type PhZnConnector to transmit information downstream.

While rate laws are unit-operation-specific, thermophysical properties are not. Therefore, in ProcessSimulator.jl, the ControlVolume component encapsulates the "StateComponent". Internally, it calls functions which are used to evaluate its thermophysical properties, i.e.; $\mathrm{U} = internal_{energy(medium,T,\rho,N_i)}, \rho = density(medium, T, P, N_i)$ $h_{out} = enthalpy(medium, T, \rho, N_i)$

Where $medium$ is a structure that holds the information of the thermophysical model and $internal\ energy, density, enthalpy$ are functions to calculate the properties of interest. For density, for example, it typically involves a phase stability check followed by a sequence of other calculations depending on the $medium$ type. All these concepts are illustrated in the Figure 2.

When composing thermodynamic calculations with MTK.jl, direct calls from within those calculations can cause problems. MTK.jl makes extensive use of symbolic manipulation, and tracing symbolic variables through complex code structures leads to a symbolic explosion. To stop tracing, functions of this kind have to be registered. Registration informs MTK.jl that particular calculations should be treated as a single block rather than decomposed. Registering a function in MTK.jl is very simple, allowing users to treat any complex block in the modelling language as an isolated calculation block. If the registered functions are Julia-native, the Julia compiler will perform its optimisations, as it does in any other piece of Julia code, after code generation. This makes the modelling environment extremely flexible, powerful and performant, as discussed before.

Another challenge in chemical and process systems is that phase splitting can occur at any time during numerical simulation; therefore, it must be checked during the simulation. In practice, it means that h_{out}, U depends on the number of phases present, each phase's composition, and the fraction of each phase. If only one phase exists, it means that the fraction of the others is zero, and their compositions can be assigned to an arbitrary value. Stability checks must be run during simulation and used to detect phase splitting when it occurs.

RESULTS AND DISCUSSION

To illustrate the modularisation strategy implemented in ProcessSimulator.jl in this work, a steady-state, fixed-volume CSTR reactor simulation is presented.

The reactor has $1\ \mathrm{m}^3$ volume and carry out the reaction between ethylene oxide and water to produce ethylene glycol with all stoichiometric coefficients equal to 1, and the reaction rate given by a power law with a first-order dependency on the concentration of ethylene oxide. The kinetics are expressed in code using the structure `PowerLawReaction` developed in ProcessSimulator.jl:

```
rxn1 = PowerLawReaction(components = compo-
nents, stoichiometry = Dict("ethylene oxide"
=> -1.0, "water" => -1.0, "ethylene glycol"
=> 1.0), order = Dict("ethylene oxide" =>
1.0, "water" => 0.0), A = 0.5, Eₐ = 0.0)
```

The inlet temperature is set to 350.15 K and 5 bar, with a flow rate of 100 mol/s, and an excess of water (60 mol/s of water and 40 mol/s of ethylene oxide). The reactor pressure is set to match the inlet pressure. The inlet specifications are encapsulated in the `Boundary_pTzn` component, which is equivalent to an infinite source of matter that drains at a fixed pressure, temperature, composition, and flow rate. It was not discussed in the methodology, but anyone familiar with common simulators is familiar with the material stream on those. It is the exact same construction. In terms of code, it is wrriten as:

```
@named S1 = Boundary_pTzn(medium = medium, p
= 5*101325.0, T = 350.15, z = [0.4, 0.6,
0.0], flowrate = 100.0, flowbasis = :molar)
```

Medium is created using the Clapeyron.jl [7] extension, which is built as part of ProcessSimulator.jl. The details of the thermodynamic model are beyond the scope of this work, but it is correlation-based and based on established literature.

```
components = ["ethylene oxide", "water",
"ethylene glycol"]
idealmodel = CompositeModel(components,
 liquid = RackettLiquid,
 gas = ReidIdeal(components, reference_state
= :formation),
 saturation = DIPPR101Sat,
 hvap = DIPPR106HVap)
activity = NRTL(components)
model = CompositeModel(components, fluid =
idealmodel, liquid = activity)
medium = EoSBased(components = components,
eosmodel = model)
```

To create the fixed-volume reactor, we use an extension of the control volume from lines 150 to 174 in https://github.com/SciML/ProcessSimulator.jl/blob/sigle-phase-approach/src/reactors/CSTR.jl. These are the code implementations of the additional equations that extend the control volume to a reactor.

With this set of specifications, it is easy to demonstrate that no additional specifications are required to fulfil the degrees of freedom, and the outlet flow rates and temperature are then determined by solving the conservation laws, kinetics and additional constraints.

The results for the outlet stream were compared with those obtained in ASPEN V12.1 using the NRTL non-ideal model and the default ideal models. The comparison is shown in Table 1. The largest difference is in the ethylene oxide (EO) mole fraction (about 60%). This is attributed to solver configurations rather than implementation errors, since the other variables show discrepancies smaller than 10% and orders of magnitude much higher than the mole fraction.

Table 1: Outlet stream results comparison – ProcessSimulator.jl and ASPEN.

Property	Unit	Aspen (SI)	ProcessSimulator (SI)	% diff (PS vs Aspen)
Temperature	K	350.15	350.15	0.0%
Pressure	Pa	5.07×10^5	5.07×10^5	0.0%
Liquid molar flow	mol/s	60.2	60.34	0.2%
Molar enthalpy	J/mol	-3.94×10^5	-3.55×10^5	9.7%
Total mass flow	kg/s	2.84	2.84	0.0%
Mole frac — EO	—	0.00355	0.00562	58.5%
Mole frac — H_2O	—	0.336	0.337	0.4%
Mole frac — EG	—	0.661	0.657	−0.5%
Heat duty	W	-4.00×10^5	-3.81×10^5	4.8%

CONCLUSIONS

ProcessSimulator.jl demonstrates that acausal, symbolic–numeric modelling in Julia provides a flexible and extensible foundation for process simulation. By combining modular unit-operation design with seamless access to procedural thermodynamic and kinetic models, the framework reduces modelling friction and supports advanced custom applications. Although still in early development, the presented philosophy and case study indicate strong potential for future research, industrial prototyping, and education in process systems engineering.

DIGITAL SUPPLEMENTARY MATERIAL

All codes are available at https://github.com/SciML/ProcessSimulator.jl/blob/sigle-phase-approach/test/reactors/ss_cstr_test.jl.

ACKNOWLEDGEMENTS

This publication has been produced with support from the HYDROGENi Research Centre (hydrogeni.no), performed under the Norwegian research program FMETEKN. The authors acknowledge the industry partners in HYDROGENi for their contributions and the Research Council of Norway (333118).

REFERENCES

1. Preisig HA. Constructing and maintaining proper process models. Computers & Chemical Engineering 34:1543-1555 (2010). https://doi.org/10.1016/j.compchemeng.2010.02.023
2. Mansouri SS, Sivaram A, Savoie CJ, Gani R. Models, modeling and model-based systems in the era of computers, machine learning and AI. Computers & Chemical Engineering 194:108957 (2025). https://doi.org/10.1016/j.compchemeng.2024.108957
3. Dieterich EE, Eigenberger G. The modusim concept for modular modeling and simulation in chemical engineering. Computers & Chemical Engineering 21:S805-S809 (1997). https://doi.org/10.1016/s0098-1354(97)87601-6
4. Lee A, Ghouse JH, Eslick JC, Laird CD, Siirola JD, Zamarripa MA, Gunter D, Shinn JH, Dowling AW, Bhattacharyya D, Biegler LT, Burgard AP, Miller DC. The idaes process modeling framework and model library—flexibility for process simulation and optimization. J Adv Manuf & Process 3: (2021). https://doi.org/10.1002/amp2.10095
5. Rackauckas C. Modelingtoolkit, modelica, and modia: the composable modeling future in julia. The Winnower : (None). https://doi.org/10.15200/winn.162133.39054
6. Nayak P, Dalve P, Sai RA, Jain R, Moudgalya KM, Naren PR, Fritzson P, Wagner D. Chemical process simulation using openmodelica. Ind. Eng. Chem. Res. 58:11164-11174 (2019). https://doi.org/10.1021/acs.iecr.9b00104
7. Walker PJ, Yew HW, Riedemann A. Clapeyron.jl: an extensible, open-source fluid thermodynamics toolkit. Ind. Eng. Chem. Res. 61:7130-7153 (2022). https://doi.org/10.1021/acs.iecr.2c00326

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Nanoparticle Nucleation and Growth Model Exploration with Perturbative Analysis

Stephen T. King[a], Antonios Armaou[ab#], Themis Matsoukas[a], Griffin A. Canning[a], and Robert M. Rioux[ac*]

[a] The Pennsylvania State University, Department of Chemical Engineering, University Park, PA, 16802, USA
[b] University of Patras, Department of Chemical Engineering, Patras, 26504, Greece
[c] The Pennsylvania State University, Department of Chemistry, University Park, PA, 16802, USA
* Corresponding Author 1: rmr189@psu.edu.
# Corresponding Author 2: armaou@upatras.gr.

ABSTRACT

Nanoparticle (NP) synthesis has been extensively studied since the mid-1800s and are utilized across numerous fields due to their unique microscopic properties that collectively yield macroscopic benefits. Of particular interest are silver (Ag) NPs, whose controllable size and morphology impart distinct catalytic, electronic, and optical properties advantageous for environmental and energy-related applications. The theoretical understanding of NP nucleation and growth has advanced considerably starting with classical nucleation theory, evolving into the LaMer model centering on burst nucleation and diffusion-limited growth and resulted in near monodispersed hydrosols. Finke and Watzky later introduced the autocatalytic model considering a slow and continuous nucleation and autocatalytic surface growth not limited by monomer diffusion. However, the precise mechanisms remain the subject of active debate for the different homogeneous and heterogenous nucleation systems. In this study, simulation models for NP population balances are developed using both material balanced and constant number Monte Carlo methods to describe NP formation and growth under diffusion-limited and autocatalytic conditions. Perturbative analysis of the model provides insight into whether NP formation is diffusion-limited or autocatalytic and how addition(s) or subtraction(s) of precursor influence the NP growth mechanism. Perturbative analysis reveals the mean particle size and variance can be tuned by adding or removing precursor molecules during the reaction, although the overall size distribution trend remains consistent. The present study together with new experimental capabilities lays the foundation for a model-based design of experiments to explore the nucleation and growth mechanism of metal nanoparticles.

Keywords: Simulation, Modelling, Materials, Catalysis, Nanoparticles

INTRODUCTION

The development of nanoparticle (NP) synthesis models began with classical nucleation theory (CNT), progressed through the LaMer diffusion-limited model [1], and later evolved into the autocatalytic Finke–Watzky framework [2]. Early experimental observations, starting in the mid-19th century, established that NPs could form under controlled chemical conditions and exhibit unique properties distinct from bulk materials [3]. As experimental techniques improved, these foundational ideas were extended into population balance modeling (PBM), which captures the dynamic evolution of particle size distributions by accounting for nucleation, growth, and aggregation processes [4].

CNT emerged from pioneering work by Faraday, Ostwald, Zsigmondy, and later Becker and Döring, who formalized nucleation kinetics using population-based equations [5]. CNT successfully explains the energetic barrier to nucleus formation and distinguishes between homogeneous and heterogeneous nucleation. However, it treats nuclei as bulk-like entities and provides limited mechanistic insight into post-nucleation growth and final particle size distributions, leaving important gaps in understanding nanoparticle growth behavior [5].

To address growth and size uniformity, LaMer

introduced a diffusion-limited model describing three distinct stages: monomer accumulation, burst nucleation once a critical concentration is reached, and subsequent growth limited by monomer diffusion. This model explained the formation of near-monodisperse particles and became widely cited as a general framework for nanoparticle synthesis [1]. Nonetheless, because it was developed for sulfur sols and extrapolated to other systems, its applicability, particularly to transition metals, has been questioned, as many systems do not exhibit instantaneous burst nucleation or purely diffusion-limited growth [6].

The autocatalytic Finke–Watzky model challenged the LaMer paradigm by proposing slow, continuous nucleation at low monomer concentrations coupled with rapid autocatalytic surface growth. In this framework, smaller particles grow faster than larger ones, naturally leading to near-monodisperse (or sized focused) distributions without burst nucleation [2]. Subsequent refinements added diffusion driven agglomeration and secondary autocatalytic steps, supported by kinetic studies, microscopy, and simulations, further improving the model's ability to capture experimentally observed nanoparticle growth pathways [7].

Modern approaches integrate the autocatalytic framework with mechanism-enabled PBM to explicitly track particle dispersity throughout nucleation and growth [8]. While this approach has been successful for certain transition metals, evidence from silver nanoparticle and nanowire systems shows that diffusion processes remain important alongside autocatalytic growth and LaMer growth may still be important [7, 9, 14]. Consequently, accurate modeling of silver NP formation requires a hybrid framework that incorporates both autocatalytic surface reactions and diffusion-limited transport, motivating the development of tailored kinetic and population balance models for silver NPs (AgNPs). These computational studies support our experimental efforts to understand the mechanism of Ag NP growth by the polyol process [10-13].

THEORETICAL OBJECTIVES

The objective of this research is to develop a mechanistic and kinetic understanding of NP nucleation, growth, and particle size distribution (PSD), with particular emphasis on how transient perturbations in reactant concentration can differentiate formation mechanisms. By elucidating not only how NPs form but also how their size and PSD can be controlled, this work aims to enable the rational design of AgNPs for applications in medicine, catalysis, sensing, and related commercial fields [14]. A testable, mechanistic hypothesis is proposed in which nucleation follows first-order kinetics and growth proceeds via second-order autocatalytic kinetics, represented through reaction extents for nucleation (ξ_0) and growth (ξ_{n+1}, n > 0). This framework (Table 1) allows systematic investigation of how changes in reaction conditions, precursor concentration(s), and additives can be used to manipulate AgNP size and dispersity, providing a foundation for model refinement using empirical data.

Table 1. Hypothesized reaction mechanism for NP nucleation and growth.

Step	Reaction	Rate
Nucleation	$A \xrightarrow{k_0} C_1$	$R_0 = \frac{d\xi_0}{dt} = k_0[A]$
Growth	$A + C_n \xrightarrow{k_n} C_{n+1}$	$R_n = \frac{d\xi_n}{dt} = k_n[A][C_n], \forall\, n \geq 1$

The corresponding material balance equations for the different species generating NPs are given in Table 2, assuming a batch reactor system. These predict how the NP concentration changes over time in relation to the precursor and nucleate.

Table 2. Species concentration rate over time under reaction mechanism of Table 1.

Species	Rate
Precursor	$\frac{d[A]}{dt} = -k_0[A] - \sum_{n=1}^{\infty} k_n[A][C_n]$
NP Nucleate	$\frac{d[C_1]}{dt} = k_0[A] - k_1[A][C_1]$
NP	$\frac{d[C_n]}{dt} = k_{n-1}[A][C_{n-1}] - k_n[A][C_n], \forall n \geq 2$

Table 3. Normalized NP species concentrations over time.

Species	Normalized Rate
Precursor	$\frac{dc_A}{d\tau} = -\kappa_0 c_A - c_A \sum_{n=1}^{\infty} \kappa_n c_n$
NP Nucleate	$\frac{dc_1}{d\tau} = \kappa_0 c_A - \kappa_1 c_A c_1$
NP	$\frac{dc_n}{d\tau} = \kappa_{n-1} c_A c_{n-1} - \kappa_n c_A c_n, \forall n \geq 2$

Species concentrations are normalized to the initial precursor concentration $[A]_0$, yielding the dimensionless quantities c_A for the precursor and c_n for the nanoparticles. Characteristic times are defined based on the reaction rate constants: the nucleation time τ_N and the growth time τ_G, normalized to the nucleation and growth rate

constants, respectively, as summarized in Table 3. Normalizations for autocatalytic growth are also provided by:

$$\tau = \frac{t}{\tau_G} \quad , \kappa_0 = \frac{\tau_G}{\tau_N} = \frac{k_0}{[A]_0 k_n} \tag{1}$$

These equations form the foundation of the NP nucleation and growth mechanisms simulated in the two models developed here. To enable versatility across both diffusion-limited and autocatalytic regimes, a diffusion component is incorporated to account for the transport of monomer to the NP surface. The Smoluchowski equation, which characterizes diffusion-limited aggregation [15-18], was originally formulated to describe the collision rate between two colloidal particles. A simplified form of this expression is used here by assuming the precursor is a single Ag atom of unit radius r_{Ag}, the NP radius follows a 1/3 power-law dependence on the number of Ag atoms per NP (or n), and the diffusion coefficient (D):

$$k_{D,n} = 4\pi D r_{Ag} n^{1/3} \tag{2}$$

Here, the effective NP growth rate constant is expressed as the inverse sum of the autocatalytic and diffusion-limited rate constants per Collins-Kimball [16-17].

$$k_{eff,n}^{-1} = k_{D,n}^{-1} + k_{A,n}^{-1} \tag{3}$$

This formulation allows the model to represent purely diffusion-limited growth ($k_{A,n} = 0$), purely autocatalytic growth ($k_{D,n} = 0$), or a combination of both mechanisms. Accordingly, the reaction rates listed in Table 1 can be extended to incorporate diffusion effects by using Equation 3. The rate of precursor diffusion to the NP surface relative to autocatalytic growth rate can be summarized by the Damköhler number (Da).

$$Da_n = \frac{k_{A,n}}{k_{D,n}}, for\ n > 0 \tag{4}$$

Larger Damköhler numbers yield higher autocatalytic growth of NPs relative to diffusive flux of precursor atoms to the NP surface [18]. For this study, we can consider the ratio of the rate constant for the first growth event to the diffusion coefficient, $k_{A,1}/D$, for characterization of varying degrees of diffusion to autocatalytic growth.

MODELING METHODOLOGY

Two NP nucleation and growth simulation models were developed, a material balanced (or constant volume) model and a constant Number Monte Carlo Model. Both utilize the Metropolis-Hastings Markov Chain Monte Carlo (MCMC) method to nucleate and grow NPs [19-20].

Material Balanced Monte Carlo Model

A stochastic model for NP formation was implemented in Python using the Metropolis–Hastings MCMC method. The NP formation system can be considered a Markov Chain since the NP reaction rates are predicated on the current concentration of the NPs and precursor [20]. The material balanced NP formation algorithm proceeds in two stages.

For the first stage, nucleation is represented as the initial element, while subsequent elements correspond to autocatalytic growth. Reaction rate constants follow an Arrhenius-like decay reflecting the lower surface area to volume as NP size increases. Alternatively, a power law decay could also be used but Arrhenius-like decay was used for this work.

$$k_{A,n} = k_{A,1} e^{-\lambda n} \tag{5}$$

First, a uniform random number, $N_{R,1}$, between 0 and 1 is selected based on a cumulative normalized vector of reaction rates defined by Equation 6 below:

$$X_n = \frac{\sum_0^i R_n}{\sum_0^T R_n} \tag{6}$$

where T is the maximum size of the NP reaction rate vector. Within the normalized vector of reaction rated X_n, an NP size is identified such that:

$$X_{n-1} < N_{R,1} \leq X_n \tag{7}$$

The normalized cumulative vector enables selection of nucleation or NP growth where nucleation dominates at the initial stage, while growth probabilities become more dispersed as particles enlarge.

Second, the accept–reject algorithm determines whether the chosen event occurs. Reaction rates are normalized to the maximum rate per Equation 8:

$$Y_n = \frac{R_n}{\max(R_n)} \tag{8}$$

a second uniform random number, $N_{R,2}$, is drawn and, if it is less than or equal to the normalized rate, the event is accepted. If it is greater, the event is rejected. This ensures particles with higher reaction rates are more likely to grow, while those with lower rates stagnate.

Ultimately, this framework provides mechanistic insight into NP formation and growth. By training the model with experimental data, particle size distributions can be mapped over time using heat maps, contour plots, or other 3-D visualizations. This enables model-based design of experiments to compare classical LaMer and Finke–Watzky autocatalytic models, offering a pathway to control NP size and distribution through mechanistic understanding of nucleation and growth kinetics.

Adaptability of NP Formation

The material balanced model is extended to incorporate monomer additions and subtractions at defined times during nucleation and growth. Resulting changes in reaction rate constants, surface free energy, and precursor reduction are evaluated to assess their impact on NP formation. NP evolution is tracked through changes in mean particle size and the full-size distribution. Key perturbations include adding precursor atoms during

growth, introducing structure-directing agents (H_2), and accounting for ambient oxygen. Incorporating these variations enables the simulation to quantify how each perturbation influences NP formation and growth dynamics, yielding a clearer picture of the underlying mechanism.

Although precursor concentrations are difficult to control experimentally, perturbations observed in experimental data can be used to validate the simulation. Canning et al have used perturbation analysis that nucleation is fast relative to growth [10]. Minute changes in precursor concentration are challenging to monitor in real time, but the model enables controlled virtual experiments that replicate these conditions. By comparing simulation results with measurable experimental perturbations, the robustness and predictive capability of the model can be validated.

The material balanced model, while effective in simulating precursor concentrations, becomes computationally expensive and less practical for large NP sizes, as all precursor and NPs must be explicitly accounted for. To address this limitation, a second approach was developed: the constant Number Monte Carlo model. This framework tracks a fixed number of particles that may nucleate, grow, or leave the system based on probabilities dictated by reaction rates. This approach reduces computational demands and extends functionality, making it more suitable for simulating NP nucleation and growth at larger scales.

Constant Number Monte Carlo (cNMC) Model

In the constant Number Monte Carlo (cNMC) method, an array of NPs is tracked for nucleation and growth within a simulation box. This approach is analogous to transmission electron microscopy (TEM), where a representative sample slide reveals the NP distribution [21]. In the cNMC model, the number of NPs in the simulation box (c_N) remains constant, while the NP mass concentration (c_M) increases over time as NPs nucleate and grow. The average volume $\langle \upsilon \rangle$ is defined as:

$$\langle \upsilon \rangle = \frac{c_M}{c_N} \tag{9}$$

Thus, NP distributions and concentrations must be tracked individually. Assuming normalized precursor and NP concentrations, the initial conditions are defined as:

$$c_{0,(t=0)} = 1, c_{n,(t=0)} = 0 \tag{10}$$

$$c_N = \sum_1^\infty c_n = 0 \tag{11}$$

The NP autocatalytic growth rate and mean growth rate constant are expressed as:

$$r_A = c_0 c_N \frac{\sum_{n=1}^\infty k_n c_n}{c_N} \tag{12}$$

Experimentally, TEM measurements of NP size and number yield a similar mean growth rate constant:

$$\langle k_N \rangle = \frac{\sum_{n=1}^\infty k_n c_n}{c_N} = \frac{\sum_{n=1}^N k_{n(i)}}{N} \tag{13}$$

with nucleation and autocatalytic growth rates defined as

$$r_N = k_0 c_0 \tag{14}$$

$$r_A = c_0 c_N \langle k_N \rangle \tag{15}$$

The simulation box contains N particles representing normalized concentrations $c_N(t)$, with N constant while c_N and c_M evolve. Each Monte Carlo step updates the system by either: (1) nucleation adds an NP of mass 1 to the box and ejects a random NP from the box to maintain a constant N or (2) growth adds one unit of mass ($\delta_\upsilon = +1$) to an existing NP in the box. The time increment and NP mass concentration change are given by:

$$\delta t = \frac{1}{r_N + r_G} \tag{16}$$

$$\delta c_M = \frac{c_m \delta \upsilon}{N \langle \upsilon \rangle} \tag{17}$$

with initial conditions:

$$t = 0, c_0 = 1, c_M = 0, c_N = 0, \langle \upsilon \rangle = 0$$

The first event is always nucleation and therefore Equations 14-15 reduce to:

$$\delta t = \frac{1}{k_0} \tag{18}$$

$$\delta c_M = \frac{1}{N} \tag{19}$$

Subsequent events update the mean NP size and growth rate constant:

$$\langle k_n \rangle = \frac{\sum_{i=1}^N k_i}{N}, \langle \upsilon \rangle = \frac{\sum_{i=1}^N \upsilon_i}{N} \tag{20}$$

At each step, nucleation or growth is chosen probabilistically. The probabilities are:

$$P_N = \frac{r_N}{r_N + r_G} \tag{21}$$

$$P_G = \frac{r_G}{r_N + r_G} \tag{22}$$

$$P_N + P_G = 1 \tag{23}$$

If growth occurs, an NP is selected based on its relative reaction rate constant. The chosen NP increases in mass by one unit, and system variables are updated in Equations 14-15. This algorithm is repeated until the desired number of iterations is reached, enabling simulation of NP nucleation and growth dynamics under constant particle number constraints.

RESULTS

Two MCMC models, the material balanced model and the cNMC model, were developed and validated to simulate NP formation and growth. These models generate data to examine how NPs nucleate, grow, and how their size distributions can be manipulated and controlled over time. Both frameworks can be adapted to represent either diffusion-limited LaMer growth, the autocatalytic

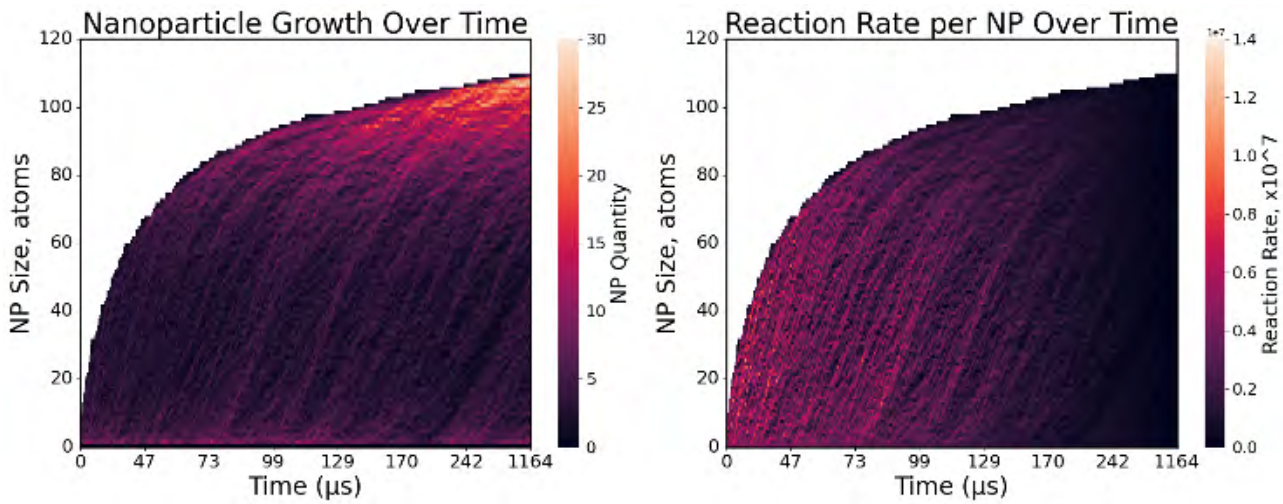

Figure 1. Material balanced model simulation of NP formation and growth and reaction rate over time. NP size (left) and reaction rate (right) show convergence of larger NPs at longer time scales (size-focusing) and decreasing reaction rate as precursor atoms are consumed. This simulation was executed assuming a nucleation rate, $k_0 = 100\ \frac{1}{s}$ initial growth rate, $k_1 = 1,000 \frac{1}{Ms}$, successive k_n values following equation 5, $D = 1.0\ \frac{l^2}{s}$, λ = 0.05, and an initial silver concentration population of 50, 000 atoms resulting in $\frac{k_{A,1}}{D} = 1,000\ \frac{1}{l^2 M}$ ($\overline{Da}$=2.51). The simulation time was calculated using the inverse sum of the reaction rates for each step.

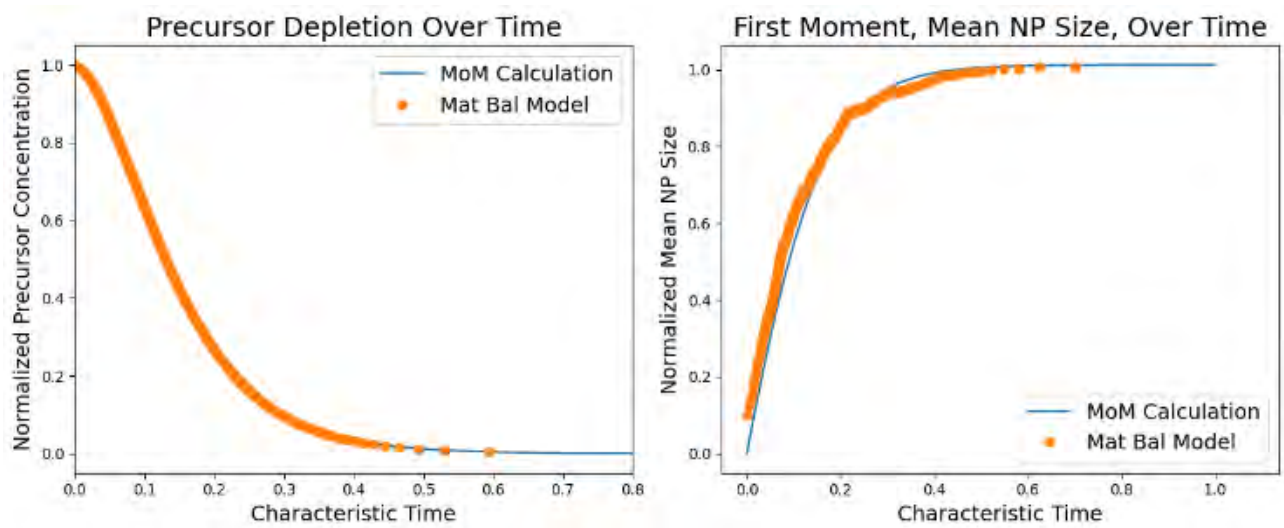

Figure 2. MoM validation of the material-balanced model. The material balanced model exhibits alignment with the MoM calculations for both precursor concentration (left) and the first-moment NP mean size (right) over time, demonstrating consistent predictive behavior across the two approaches.

Finke–Watzky mechanism of NP formation, or both. The simulations were performed on Google Colab using a T4 GPU. NP size distributions were tracked in pandas arrays, and three-dimensional heat maps were generated with the seaborn package in Python.

Material Balanced Monte Carlo Model

The material balanced (constant volume) model was implemented following the Methodology section of the referenced paper and used to simulate nanoparticle formation. This framework is well suited for systems with a finite precursor supply, enabling perturbative analyses in which precursor is added or removed and allowing calculation of NP formation yield. Consistent with literature consensus, reaction rate constants decrease as NP size increases [22]. Figure 1 shows the evolution of particle size distribution and reaction rates over time; the heat-map representation highlights their coupled,

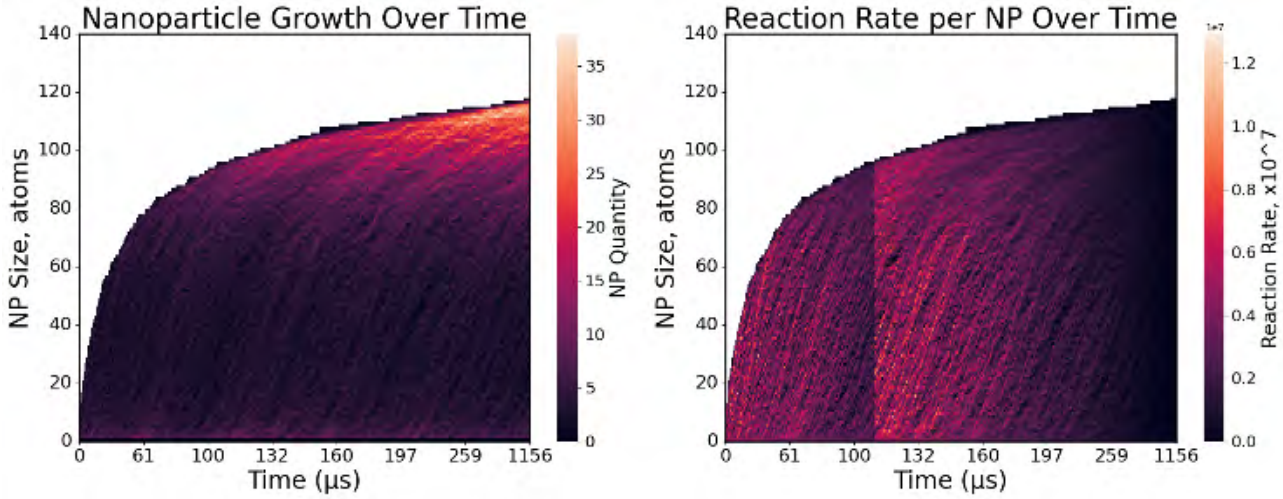

Figure 3. Perturbative analysis of NP growth where 25, 000 precursor molecules are added once 50% of the 50, 000 initial precursor molecules have reacted with conditions described in Figure 1. A sudden increase in reaction rate is noted as precursor is injected leading to a larger NP mean size.

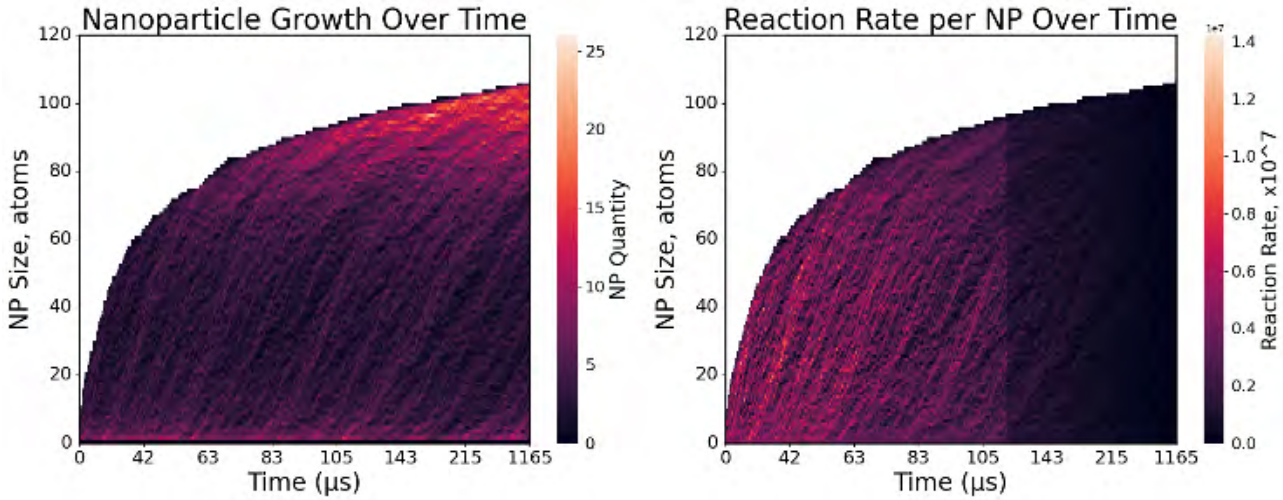

Figure 4. Perturbative analysis of NP growth where 10, 000 precursor atoms are subtracted once 50% of the 50, 000 initial precursor molecules have reacted with conditions described in Figure 1.

three-dimensional progression. As illustrated, the population size-focuses toward a mean of 73.75 ± 32.96 Ag atoms.

The material balanced model was validated using a Method of Moments (MoM) simulation for NP generation and results are shown in Figure 2. The normalized precursor concentration exhibits a sigmoidal decay profile where nucleation dominates early, growth accelerates precursor consumption, and depletion slows once precursor is consumed. The NP mean size follows the expected trend as well, increasing as NPs form and then tapering as precursor is exhausted which is in alignment with the MoM simulated mean.

Development and validation of the material balanced model provides the foundation for simulating perturbations to assess their effects on NP mean size and dispersity, which offers insight into the nanoparticle formation and growth mechanism.

Adaptability of NP Formation

The effect of precursor additions and changes in reaction conditions on NP growth was examined. In one scenario, an additional 50% of precursor was injected once half of the initial precursor had been consumed. As shown in Figure 3, starting with 50, 000 precursor atoms, an extra 25, 000 molecules were introduced after 50% precursor conversion. The reaction rate then increased sharply, evident from the intensified color on the heat map. This behavior occurs because the sudden availability of precursor accelerates growth across all NP sizes

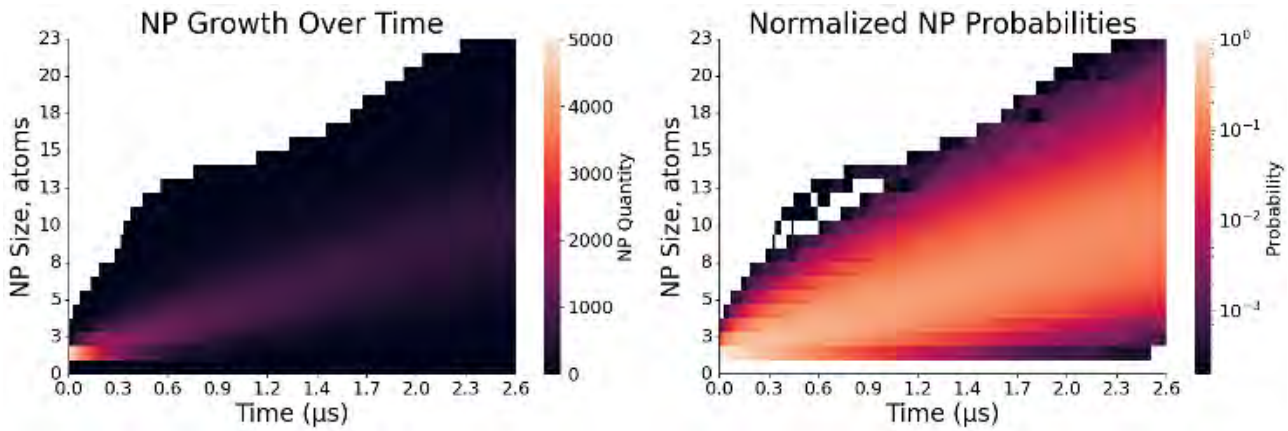

Figure 5. cNMC simulation using a 5000-particle box with k_0 is 1x10^{-12} and k_n is based on Arrhenius-like Decay (λ = 0.05). NP growth for 50, 000 iterations allows NPs to grow up to ~90 atoms.

simultaneously. In this case, the NPs size focus toward a mean of 82.33 ± 34.39 Ag atoms.

Conversely, we looked at subtracting precursor molecules when 50% of precursor atoms have reacted. Experimentally, this could be achieved by adding a deactivation reagent to the solution where suddenly a fraction of precursors are deactivated and unable to nucleate or add to existing NPs. Figure 4 shows this simulation and, in comparison with Figure 3, shows the reaction rates suddenly decrease once the fraction of precursor molecules are added to the system. In this case, the NPs size focus toward a mean of 67.80 ± 32.50 Ag atoms.

As such, perturbations in the addition or subtraction of precursor atoms during the reaction have a strong influence on the mean and the variance of the final NP size distribution. Adding precursor atoms increases the mean size and variance of the NP population whereas subtraction of the precursor molecules during the reaction lowers the mean size and variance.

Expanding on this point, Table 4 summarizes the relative variation in NP size increases as the diffusion coefficient decreases. In other words, greater resistance to diffusion leads to a more dispersed, or less size-focused, NP population. This increase in dispersity becomes even more pronounced when the system experiences perturbations that add or subtract precursor molecules.

The results in Table 4 show that both diffusion coefficient and precursor perturbations systematically shape NP size evolution. Lower diffusion coefficients produce larger mean NP sizes because restricted mass transport limits precursor delivery to NP surfaces, promoting continued growth of existing particles. At the same time, the normalized standard deviation and polydispersity increase as diffusion slows, yielding broader, less size-focused populations. Additive perturbations suppress dispersity, whereas precursor removal amplifies it, especially at low diffusion. Systems with restricted diffusion are therefore more sensitive to precursor fluctuations, resulting in greater heterogeneity in the final NP size distribution.

Table 4. Final NP size and dispersity parameters after precursor perturbations at three diffusion coefficients.

Diffusion Coeff	**0.010**	**0.100**	**1.000**
$k_{A,1}/D$	100,000	10 000	1, 000
No Perturbation			
Mean NP size	104.38	90.42	73.75
Std Dev (SD)	62.46	48.19	32.96
Norm NP Size	1.00	1.00	1.00
Norm SD	0.60	0.53	0.45
Polydispersity	1.36	1.28	1.20
+25, 000 Perturbation			
Mean NP size	120.58	104.89	82.33
Std Dev (SD)	63.70	49.38	34.39
Norm NP Size	1.16	1.16	1.12
Norm SD	0.61	0.55	0.47
Polydispersity	1.28	1.22	1.17
-10, 000 Perturbation			
Mean NP size	91.12	81.14	67.80
Std Dev (SD)	60.46	47.30	32.50
Norm NP Size	0.87	0.90	0.92
Norm SD	0.58	0.52	0.44
Polydispersity	1.44	1.34	1.23

Constant Number Monte Carlo (cNMC) Model

The cNMC model has the advantage that it can grow nanoparticles (NPs) to very large sizes and run for an effectively indefinite simulation time. Since the algorithm

does not impose a finite precursor pool, additional precursor atoms can be introduced into the simulation box as NPs grow. Under the conditions shown here in Figure 5, the mean NP size converges towards a mean of 10.99 ± 2.50 Ag atoms.

Event probabilities indicate nucleation is initially far more likely than growth; however, once NPs form, the probability rapidly shifts toward growth. This behavior is consistent with both the LaMer burst-nucleation picture and the autocatalytic Finke–Watzky mechanism. As NPs enlarge, their total mass increases while the particle concentration (c_n) remains constant within the simulation box. The cNMC framework, therefore, provides an alternative approach to simulating NP nucleation and growth that is not constrained by a fixed particle number. Additionally, it is more representative of TEM experiments, where a sample of the reaction mixture is analyzed as a distribution within a finite observation volume.

Table 5. Summary of NP size and dispersity parameters obtained from the Constant Number Monte Carlo (cNMC) simulations at varying diffusion coefficients. Increasing diffusion leads to reduced polydispersity and a systematic increase in mean NP size.

Diffusion Coeff	**0.010**	**0.100**	**1.000**
$k_{A,1}/D$	100,000	10,000	1, 000
Result			
Mean NP size	5.87	9.32	10.76
Std Dev (SD)	4.47	4.53	3.09
Norm SD	0.76	0.49	0.29
Polydispersity	1.58	1.24	1.08

As observed in the material balanced model, NP polydispersity decreases with increasing diffusion coefficient. However, in contrast to the material balanced model, the mean NP size increases with greater precursor diffusivity. This difference arises because, in the cNMC model, the simulation box effectively expands as precursor atoms are added during growth. Faster diffusion therefore increases the rate at which precursors reach the NP surface, enabling larger average NP sizes while still narrowing the size distribution. In the material balanced model, by contrast, the simulation box is fixed by the initially supplied precursor, making growth increasingly supply limited as NPs grow and precursors depleate. Together, these frameworks clarify how diffusion-limited and supply limited regimes shape NP growth. Both the material balanced and cNMC models will be used for experimental studies in which simulated size distributions can be directly compared with TEM-derived NP populations to validate mechanistic assumptions about NP formation under various kinetic regimes.

DISCUSSION

A nanoparticle growth model was proposed (Table 1) incorporating LaMer diffusion-limited growth via the Smoluchowski equation with the Collins–Kimball framework, and implemented along with two formation and growth models using a Metropolis–Hastings MCMC algorithm [19-20]. Both the material balanced and cNMC models were developed to study NP formation under experimental conditions and to accommodate diffusion-limited LaMer [1] and autocatalytic Finke–Watzky [2] mechanisms. Diffusion-limited growth was modeled with Smoluchowski diffusion, while autocatalytic growth followed an Arrhenius-like decay in rate constants, consistent with observations that larger NPs grow more slowly [22]. Although the original autocatalytic framework was designed for iridium, its applicability to diffusion-limited silver systems is uncertain; diffusion remains a critical parameter for AgNP growth [7, 14] and is explicitly included here. Increasing diffusional resistance (i.e., decreasing diffusion coefficient) produces more disperse NPs in both the material balanced and cNMC models.

A perturbative analysis was performed in which precursor molecules were added or removed once 50% of the initial precursor had been consumed. This analysis revealed that subtracting precursor decreased the mean and increased the dispersity of the final NP distribution, while adding precursor caused the reverse. Despite these perturbations, the overall particle size distribution remained relatively consistent, demonstrating the robustness of the modeling framework.

A constant Number Monte Carlo (cNMC) NP growth model was also developed to incorporate both autocatalytic and diffusive components. Event probabilities show that nucleation dominates first, but once NPs begin to form, the likelihood rapidly shifts toward growth. This transition mirrors the mechanistic behavior described by both the LaMer model and the autocatalytic Finke–Watzky framework. As in the material balanced model, increasing diffusion decreases NP dispersity. However, in contrast to the material balanced model, the mean NP size increases with increasing diffusion. This difference arises because the cNMC framework does not impose a finite precursor pool and is not supply limited. Precursor can be continuously added into the simulation box and are therefore not limited by the initial concentration.

Development of these models, with the use of controlled perturbations, provides a powerful framework for understanding how NPs grow and achieve size focusing. These perturbation studies reveal how fluctuations in precursor availability can reshape both the mean and dispersity of the final NP population, offering mechanistic insight into the sensitivity of NP growth pathways. Future work will focus on aligning experimental data with both material balanced and cNMC models to determine how

nucleation and early-stage growth proceed at scales below typical instrumental detection limits and define fundamental NP formation mechanisms. The simulations presented here demonstrate precursor perturbations exert measurable and systematic effects on NP size evolution, underscoring the value of this modeling approach for probing nanoscale processes that are otherwise experimentally inaccessible.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge Honeywell International, Inc. for STK's financial support. RMR acknowledges funding from the National Science Foundation through NSF Award #2404221.

AUTHOR IDENTIFIERS

Author ORCIDs:
King ST: 0009-0003-0600-9604
Armaou A: 0000-0002-8592-7934
Matsoukas T: 0000-0002-0097-2673
Canning GA: 0000-0002-8448-4174
Rioux RM: 0000-0002-6019-0032

REFERENCES

1. LaMer VK, Dinegar RH. Theory, Production, and Mechanism of Formation of Monodispersed Hydrosols. *J. Am. Chem. Soc.* 72:4847-4854 (1950)
2. Watzky MA, Finke RG. Transition metal nanocluster formation kinetic and mechanistic studies. a new mechanism when hydrogen is the reductant: slow, continuous nucleation and fast autocatalytic surface growth. J. Am. Chem. Soc. 119:10382-10400 (1997). https://doi.org/10.1021/ja9705102
3. Whitehead CB, Özkar S, Finke RG. Lamer's 1950 model of particle formation: a review and critical analysis of its classical nucleation and fluctuation theory basis, of competing models and mechanisms for phase-changes and particle formation, and then of its application to silver halide, semiconductor, metal, and metal-oxide nanoparticles. Mater. Adv. 2:186-235 (2021). https://doi.org/10.1039/d0ma00439a
4. Ramkrishna D, Singh MR. Population balance modeling: current status and future prospects. Annu. Rev. Chem. Biomol. Eng. 5:123-146 (2014). https://doi.org/10.1146/annurev-chembioeng-060713-040241
5. Polte J. Fundamental growth principles of colloidal metal nanoparticles – a new perspective. CrystEngComm 17:6809-6830 (2015). https://doi.org/10.1039/c5ce01014d
6. Whitehead CB, Özkar S, Finke RG. Lamer's 1950 model for particle formation of instantaneous nucleation and diffusion-controlled growth: a historical look at the model's origins, assumptions, equations, and underlying sulfur sol formation kinetics data. Chem. Mater. 31:7116-7132 (2019). https://doi.org/10.1021/acs.chemmater.9b01273
7. Hornstein BJ, Finke RG. Transition-metal nanocluster kinetic and mechanistic studies emphasizing nanocluster agglomeration: demonstration of a kinetic method that allows monitoring of all three phases of nanocluster formation and aging. Chem. Mater. 16:139-150 (2003). https://doi.org/10.1021/cm034585i
8. Handwerk DR, Shipman PD, Whitehead CB, Özkar S, Finke RG. Particle size distributions via mechanism-enabled population balance modeling. J. Phys. Chem. C 124:4852-4880 (2020). https://doi.org/10.1021/acs.jpcc.9b11239
9. Qi X, Chen Z, Yan T, Fichthorn KA. Growth mechanism of five-fold twinned ag nanowires from multiscale theory and simulations. ACS Nano 13:4647-4656 (2019). https://doi.org/10.1021/acsnano.9b00820
10. Canning GA, Ogozaly S, Pekarcik M, Lieb CA, Jharimune S, Pfukwa R, Klumperman B, Rioux RM. Complementary reducing agents are responsible for temporally distinct nucleation and growth phases during the polyol synthesis of ag nanocubes. Angew Chem Int Ed 64: (2025). https://doi.org/10.1002/anie.202418611
11. Jharimune S, Pfukwa R, Chen Z, Anderson J, Klumperman B, Rioux RM. Chemical identity of poly(n-vinylpyrrolidone) end groups impact shape evolution during the synthesis of ag nanostructures. J. Am. Chem. Soc. 143:184-195 (2020). https://doi.org/10.1021/jacs.0c08528
12. Chen Z, Balankura T, Fichthorn KA, Rioux RM. Revisiting the polyol synthesis of silver nanostructures: role of chloride in nanocube formation. ACS Nano : (2019). https://doi.org/10.1021/acsnano.8b08019
13. Chen Z, Chang JW, Balasanthiran C, Milner ST, Rioux RM. Anisotropic growth of silver nanoparticles is kinetically controlled by polyvinylpyrrolidone binding. J. Am. Chem. Soc. 141:4328-4337 (2019). https://doi.org/10.1021/jacs.8b11295
14. Ben Moshe A, Markovich G. Synthesis of single crystal hollow silver nanoparticles in a fast reaction-diffusion process. Chem. Mater. 23:1239-1245 (2011). https://doi.org/10.1021/cm102991z
15. Eggers J, Fontelos M. Anomalous dimensions of the smoluchowski coagulation equation. Phys. Rev. E 108: (2023).

https://doi.org/10.1103/physreve.108.064110
16. Collins FC, Kimball GE. Diffusion-controlled reaction rates. Journal of Colloid Science 4:425-437 (1949). https://doi.org/10.1016/0095-8522(49)90023-9
17. Agmon N, Szabo A. Theory of reversible diffusion-influenced reactions. The Journal of Chemical Physics 92:5270-5284 (1990). https://doi.org/10.1063/1.458533
18. Deen WM, Analysis of Transport Phenomena, 1st ed. *Oxford Univ. Press* (1998)
19. Beichl I, Sullivan F. The metropolis algorithm. Comput. Sci. Eng. 2:65-69 (2000). https://doi.org/10.1109/5992.814660
20. Hastings WK. Monte carlo sampling methods using markov chains and their applications. Biometrika 57:97 (1970). https://doi.org/10.2307/2334940
21. Smith M, Matsoukas T. Constant-number monte carlo simulation of population balances. Chemical Engineering Science 53:1777-1786 (1998). https://doi.org/10.1016/s0009-2509(98)00045-1
22. Whitehead CB, Watzky MA, Finke RG. “burst nucleation” vs autocatalytic, “burst” growth in near-monodisperse particle-formation reactions. J. Phys. Chem. C 124:24543-24554 (2020). https://doi.org/10.1021/acs.jpcc.0c06875

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Sustainable and Optimized Biorefinery Design for the Production of High-Value Catechol Derivatives from Lignin

Alden Paul Rangel-López[a]*, Eduardo Sánchez-Ramírez[a], Maricruz Juarez-García[a], Jesús Manuel Núñez-López[a], and Juan Gabriel Segovia-Hernández[a]

[a] Departamento de ingeniería química, Universidad de Guanajuato, Campus Guanajuato, Guanajuato, 36050, México
* Corresponding Author: ap.rangellopez@ugto.mx

ABSTRACT

This work presents the conceptual design, rigorous simulation, and multi-objective optimization of an integrated multiproduct biorefinery focused on the sustainable valorization of lignin into catechol and a portfolio of high-value aromatic derivatives, including guaiacol, vanillin, and vanillic acid. The proposed methodology combines a critical review of synthesis pathways and kinetic data with detailed process modeling in Aspen Plus® V14. The process framework begins with lignocellulosic biomass pretreatment via acid hydrolysis, followed by catalytic lignin depolymerization to obtain catechol as a platform intermediate. Subsequently, dedicated reaction and separation sections are designed for the selective conversion of catechol into value-added derivatives, incorporating energy-efficient and thermally integrated separation schemes. A stochastic multi-objective optimization strategy based on Differential Evolution with a Tabu List (DETL) is implemented to simultaneously assess economic and environmental performance indicators. The results demonstrate that lignin can be effectively upgraded into commercially relevant fine chemicals, confirming its strategic role within next-generation biorefineries and reinforcing the transition toward circular, low-emission chemical processes.

Keywords: Biorefinery, Catechol, Catechol Derivatives, Process Design, DETL Optimization

INTRODUCTION

Currently, approximately 80% of global primary energy consumption and nearly 90% of chemical products are derived from fossil-based resources. The extensive exploitation of these resources has led to significant environmental, economic, and social impacts that are increasingly incompatible with the commitments established in the United Nations 2030 Agenda for Sustainable Development [1]. In this context, the development of sustainable chemical production routes capable of reducing greenhouse gas emissions and fossil dependency has become a strategic priority.

Biorefineries have emerged as a promising technological paradigm for converting renewable biomass into fuels, energy carriers, and high-value chemicals. Among the main biomass constituents, lignin represents the most abundant renewable source of aromatic structures [2].

Historically, lignin has been significantly underutilized within the lignocellulosic biorefinery framework, predominantly burned for low-grade heat and power generation [3]. Early techno-economic and life cycle assessments established the fundamental viability of shifting this paradigm. For instance, Montazeri and Eckelman [4] demonstrated the environmental advantages of lignin-derived catechol over its petrochemical counterpart, while Mabrouk et al. [5] confirmed its market competitiveness using olive tree pruning as a primary feedstock.

However, the field has rapidly evolved beyond single-product feasibility. Recent comprehensive reviews, such as John et al. [6], confirm that the current state-of-the-art is shifting toward increasingly integrated and efficient multi-product systems. These studies encompass a diverse range of feedstocks showing that product diversification is fundamental to mitigating economic risks and improving the overall sustainability of the system.

Despite these advances, a critical limitation remains: most studies focus exclusively on the production of a

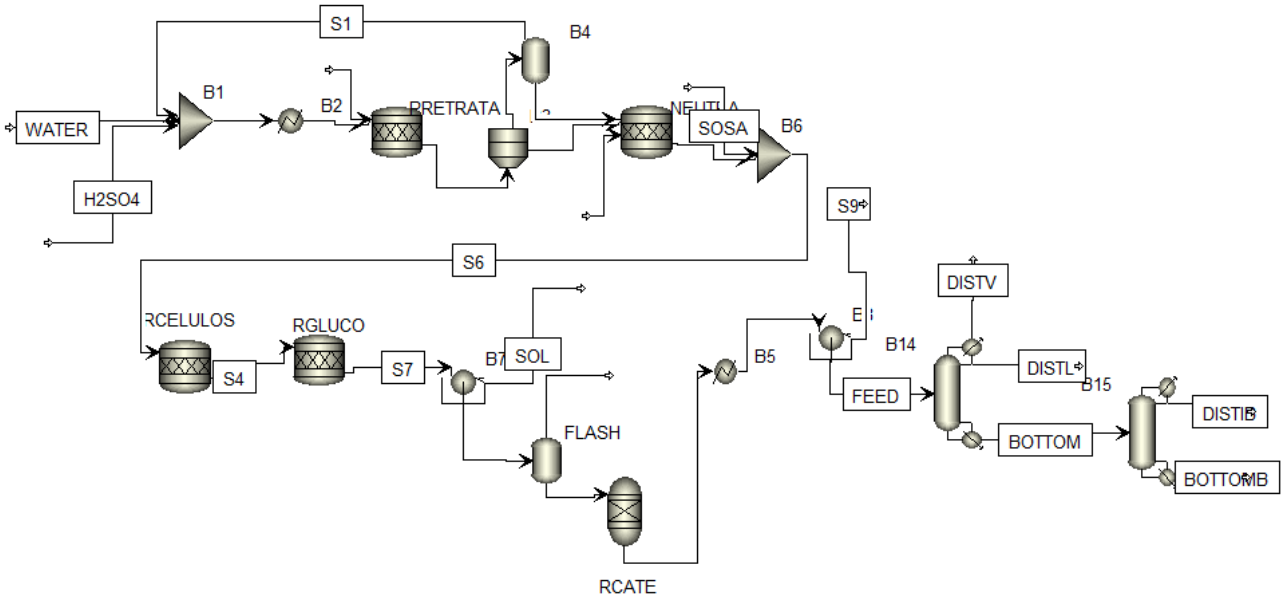

Figure 1. Diagram Process of Catechol Production.

single platform molecule. This work addresses this gap by proposing a multiproduct biorefinery where catechol serves as a strategic intermediate for the synthesis of higher value-added products, specifically focusing on the guaiacol pathway and its conversion into vanillin and vanillic acid. The novelty of this contribution lies in the integrated optimization of the entire value chain—from lignocellulosic biomass to purified fine chemicals—under circular economy principles.

By combining rigorous simulation with advanced stochastic optimization, this work provides a robust and scalable framework aligned with Sustainable Development Goals 7 (Affordable and Clean Energy), 9 (Industry, Innovation, and Infrastructure), 12 (Responsible Consumption and Production), and 13 (Climate Action), offering a viable pathway toward a bio-based and low-carbon chemical industry [7].

METHODOLOGY

The proposed methodology comprises the conceptual design, rigorous simulation, kinetic modeling, and multi-objective optimization of a multiproduct biorefinery aimed at lignin valorization. All process units were modeled using Aspen Plus® V14, ensuring thermodynamic consistency and reliable mass and energy balances.

Simulation of Catechol Production

A thermodynamic framework based on the Non-Random Two-Liquid (NRTL) activity coefficient model was selected to accurately represent the highly non-ideal behavior of the multicomponent liquid mixtures involved. The process begins with the fractionation of lignocellulosic biomass through dilute-acid hydrolysis, modeled as a continuous reactor operating at temperatures between 140 and 200 °C using sulfuric acid as catalyst. This pretreatment selectively solubilizes hemicellulose into a xylose-rich aqueous stream, while producing a solid fraction primarily composed of cellulose and lignin.

The solid fraction is subsequently processed in a catalytic depolymerization section to break down the lignin macromolecular structure. Cellulose conversion to glucose is simultaneously considered. Due to the limited availability of detailed kinetic data for lignin depolymerization, a yield-based reactor model (RYield) was employed to represent the conversion of solubilized lignin into catechol under high-temperature conditions (300 °C). The resulting product stream is purified through a dedicated separation train consisting of two distillation columns designed to recover catechol with commercial-grade purity.

Simulation of Guaiacol Pathway

Catechol is further valorized via vapor-phase O-methylation with methanol over an alumina-based catalyst. This reaction system was modeled using a tubular fixed-bed reactor (RPlug), incorporating a kinetic scheme that accounts for both guaiacol formation and the parallel production of veratrole as an undesired byproduct. Reactor operating conditions were selected to maximize selectivity toward guaiacol.

Given the formation of azeotropic mixtures, the

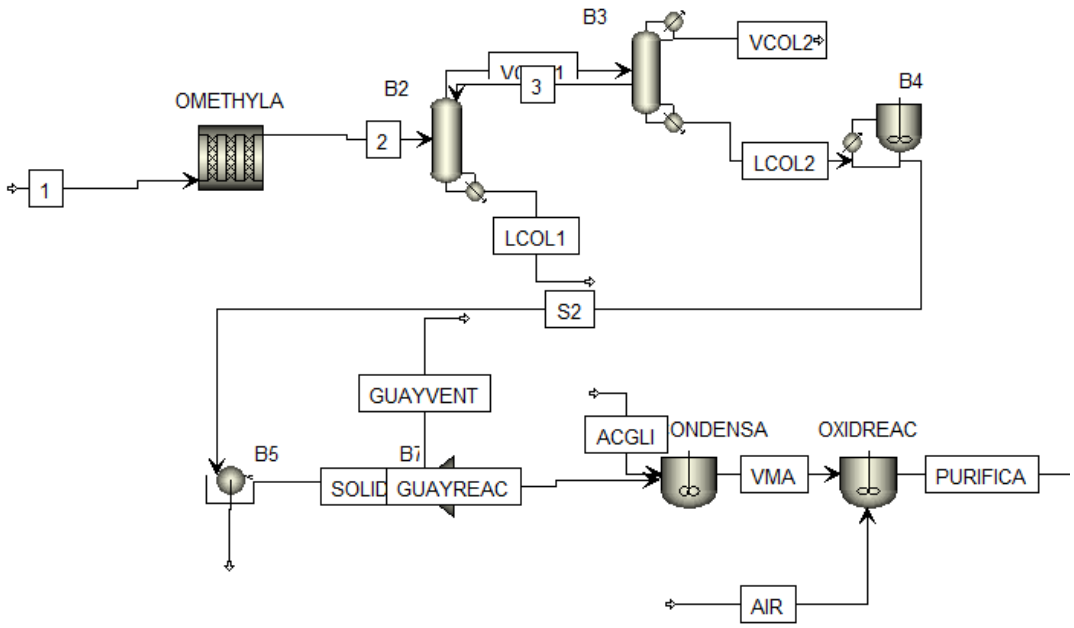

Figure 2. Diagram Process of Guaiacol Pathway.

downstream separation was modeled using a hybrid separation strategy combining distillation and melt crystallization within a thermally coupled indirect distillation sequence. This approach significantly reduces energy consumption compared to conventional standalone distillation systems.

For vanillin synthesis, the purified guaiacol stream was conditioned and fed to a liquid-phase reaction section composed of two continuous stirred-tank reactors (CSTRs) arranged in series. In the first reactor, the condensation reaction between guaiacol and glyoxylic acid was carried out in an aqueous basic medium, producing the key intermediate vanillylmandelic acid (VMA).

Finally, the conversion of VMA into the target product was achieved through oxidative decarboxylation in the second CSTR. A continuous bubbled air feed stream was introduced to supply the required oxygen. From a kinetic standpoint, this stage was defined as a system of competitive parallel reactions of zero order with respect to the reactants, favoring vanillin formation (approximately 89%) over over-oxidation to vanillic acid (approximately 11%).

Optimization of Kinetic Parameters

For the O-methylation reaction system, a first-order kinetic model with respect to each reactant was formulated. The kinetic parameters, including pre-exponential factors and activation energies, were estimated through a stochastic optimization approach using the Differential Evolution with Tabu List (DETL) algorithm. The objective function minimized the global error between simulated and experimental molar flow rates of guaiacol, veratrole, and water, ensuring robust predictive capability.

Process-level optimization was subsequently carried out using the same DETL framework, treating the overall flowsheet as a black-box system. Aspen Plus® acted as the calculation engine, while Microsoft Excel was used to evaluate economic and environmental objective functions. Decision variables included operating temperatures, pressures, flow rates, reflux ratios, and the number of theoretical stages in separation units.

$$C_6H_6O_2 + CH_3OH \rightarrow C_7H_8O_2 + H_2O \quad (1)$$

$$r_1 = k_1 C_A C_B \quad (2)$$

Where k_1 is the first-order kinetic constant, C_A is Catechol's molar concentration and C_B is the Methanol's molar concentration.

Optimization Process

Process optimization is essential due to the inherent complexity of the models, which integrate highly nonlinear and non-convex equations to describe the behavior of kinetic reactors and rigorous separation units. In this work, a stochastic black-box optimization strategy was implemented using the Differential Evolution with a Tabu List (DETL) algorithm, which interacts directly with the simulator without requiring prior knowledge of the internal equations or precise system initialization.

For the execution of this phase, an integrated framework was developed to dynamically link Aspen Plus V14 with Microsoft Excel. In this framework, the simulator

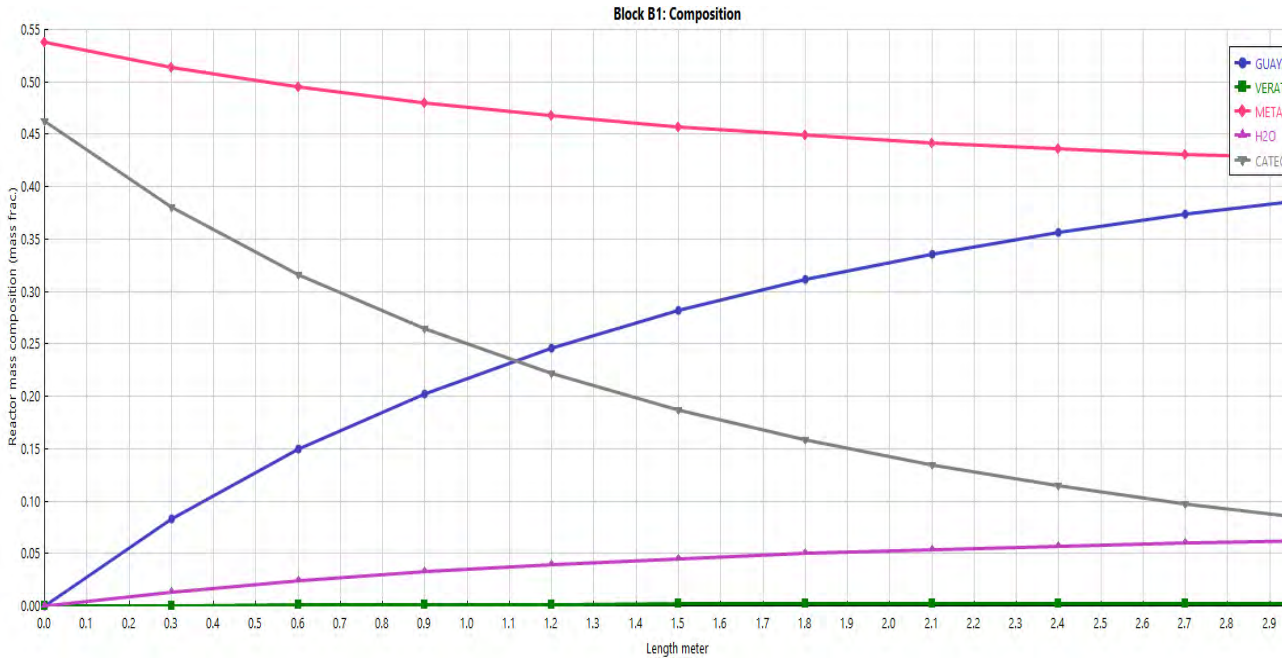

Figure 3. O-methylation reactor composition profiles

acts as the computational engine for mass and energy balance calculations for a given set of operating variables, while the spreadsheet centralizes the results to compute the objective functions. The algorithm manages the exploration of the solution space by manipulating key variables such as stream flow rates, operating temperatures and pressures of the equipment, the number of stages, reflux ratios, among others, with the aim of minimizing the objective functions.

RESULTS

The multi-objective optimization identified operating configurations that simultaneously enhance technical performance, economic viability, and environmental sustainability. Key improvements include significant reductions in energy consumption through optimized reflux ratios and thermally integrated separation schemes, increased selectivity toward high-value derivatives, and a substantial decrease in carbon footprint due to reduced utility demand and improved material efficiency.

The multi-objective optimization ensured a technically and environmentally efficient catechol production stage, establishing a robust foundation prior to its conversion into high-value derivatives. As detailed in the Table 1, the refined design of the separation trains significantly reduces energy consumption and operating costs by configuring 54 stages for Column 1 and 30 stages for Column 2. The implementation of optimized reflux ratios (21.54 and 72.68, respectively) facilitates the recovery of a platform intermediate with commercial-grade purity while minimizing the overall carbon footprint. These results supported the economic and sustainable viability of catechol as a precursor within the proposed biorefinery scheme.

Table 1: Catechol Purification Results

Variable	Column 1	Column 2
Pressure	1.01 bar	1.01 bar
Number of Stages	54	30
Feed stage	50	13
Distillate rate	1925.74 kmol/h	0.682033 kmol/h
Bottom Rate	140.736 kmol/h	140.054 kmol/h
Reflux Ratio	21.5	72.7
Condenser Duty	-4.7878×10^6 cal/sec	-175740 cal/sec
Reboiler Duty	4.988×10^6 cal/sec	176079 cal/sec

Table 2 demonstrates a highly efficient material balance for the catechol production stage. The separation train processes a substantial distillate rate in Column 1 to isolate key components, eventually yielding a steady bottom rate of 140.05 kmol/h in Column 2, which corresponds to the purified catechol stream. This optimized flow distribution minimizes waste and ensures that catechol is recovered with the high purity required for its subsequent conversion into high-value derivatives.

Table 2: Catechol Production Streams

Compound	Flow
Biomass	15000 kg/h
Xilose	4602.25 kg/h
Glucose	824.998 kg/h
Soluble Lignin	1375.5 kg/h
Methanol	5928.53 kg/h
Catechol	200 kg/h

Through the implementation of the DETL stochastic optimization algorithm, the global error objective function was successfully minimized, yielding a set of kinetic parameters that faithfully reproduce the experimental flow profiles. The optimized values for the pre-exponential factors and activation energies regarding the proposed reaction scheme are summarized in Table 3.

Table 3: Optimization of Kinetic Parameters Results

Variable	Value
Pre-exponential factor 1	5.96×10^{9}
Activation Energy 1	22931.2 cal/mol
Pre-exponential factor 2	4.08×10^{7}
Activation Energy 2	20052.9 cal/mol
Pre-exponential factor 3	5.50×10^{7}
Activation Energy 3	23500.5 cal/mol

Table 4: Guaiacol Purification Results

Variable	Column 1	Column 2
Pressure	1.01 bar	1.01 bar
Number of Stages	58	20
Feed stage	15	10
Distillate rate	8.6263 kmol/h	7.2653 kmol/h
Bottom Rate	0.4553 kmol/h	1.3609 kmol/h
Reflux Ratio	16.4	9.0
Condenser Duty	-554, 992 cal/sec	-178, 723 cal/sec
Reboiler Duty	547, 441 cal/sec	156, 341 cal/sec

The integration of thermally coupled demonstrates high technical and economic efficiency for the separation of higher-value derivatives. As shown in the process flowsheet, the effluent from the O-methylation reactor is processed through a sequence of columns where thermal coupling minimizes energy demand. The optimized design, detailed in the Table 4, achieves commercial-grade purity while significantly decreasing operating costs through equipment size reduction and improved heat integration.

Based on the mass balance results presented in the Table 5, a catechol feed of 200 kg/hr is successfully processed to yield 166.66 kg/hr of guaiacol as a primary intermediate. Subsequent condensation and oxidation stages further valorize this stream, resulting in the production of 60.911 kg/hr of vanillin and 7.5284 kg/hr of vanillic acid. The precise control of side-reactions is evidenced by the minimal formation of by-products such as veratrole (1.0286 kg/h), ensuring high selectivity and product purity throughout the integrated biorefinery scheme. These flow rates confirm the robustness of the optimized chemical pathways in maximizing the economic revalorization of lignin-derived platform molecules.

Table 5: Guaiacol Pathway streams.

Compound	Flow
Catechol	200 kg/hr
Guaiacol	166.66 kg/hr
Veratrole	1.0286 kg/h
VMA	68.44 kg/hr
Vanillin	60.911 kg/h
Vanillic Acid	7.5284 kg/h

These results consolidate the proposed scheme as a viable alternative for next-generation biorefineries focused on the sustainable valorization of lignin, contributing to Sustainable Development Goals (SDGs) 7, 9, 12, and 13 related to clean energy, innovative industry, responsible production, and climate action.

CONCLUSION

This study demonstrates the technical feasibility of an integrated biorefinery, achieving a high-purity separation of aromatic derivatives through an optimized distillation sequence. Simulation results in Aspen Plus show that the process can recover catechol with a molar purity of 99.99% and a recovery rate of 200 kg/hr, and production of 166.66 kg/hr of Guaiacol, 60.911 kg/hr of vanillin, and 7.52 kg/hr of vanillic acid from the lignin-derived stream.

The energy performance of the process was significantly enhanced through the integrated design, where the total thermal duty of the distillation reboilers was minimized to ensure an efficient resource allocation within the circular economy framework. By combining rigorous simulation with advanced stochastic optimization, this work establishes a robust numerical foundation for further life cycle assessment and process intensification. These findings reinforce the strategic role of biomass-derived aromatics in the transition toward a sustainable and low-carbon chemical industry.

ACKNOWLEDGEMENTS

Authors acknowledge the support provided by ADRESS Group, SECIHTI, and Universidad de Guanajuato

AUTHOR IDENTIFIERS

Author ORCIDs:
Alden Paul Rangel-López: 0009-0004-9316-0157
Eduardo Sánchez- Ramírez 0000-0002-4326-4837
Maricruz Juarez-García: 0000-0002-6613-9377
Jesús Manuel Núñez-López: 0009-0003-3080-4325
Juan Gabriel Segovia-Hernández: 0000-0003-4175-9809

REFERENCES

1. Takkellapati S, Li T, Gonzalez MA. An overview of biorefinery-derived platform chemicals from a cellulose and hemicellulose biorefinery. Clean Techn Environ Policy 20:1615-1630 (2018). https://doi.org/10.1007/s10098-018-1568-5
2. Sharara MA, Clausen EC, Carrier DJ. An overview of biorefinery technology. Biorefinery Co?Products :1-18 (2012). https://doi.org/10.1002/9780470976692.ch1
3. A. Giuliano, M. Poletto, and D. Barletta, “Process Design of a Multi-Product Lignocellulosic Biorefinery, ” in Computer Aided Chemical Engineering, vol. 37, Elsevier B.V., 2015, pp. 1313–1318. doi: 10.1016/B978-0-444-63577-8.50064-4.
4. Montazeri M, Eckelman MJ. Life cycle assessment of catechols from lignin depolymerization. ACS Sustainable Chem. Eng. 4:708-718 (2016). https://doi.org/10.1021/acssuschemeng.5b00550
5. Mabrouk A, Erdocia X, Alriols MG, Labidi J. Economic analysis of a biorefinery process for catechol production from lignin. Journal of Cleaner Production 198:133-142 (2018). https://doi.org/10.1016/j.jclepro.2018.06.294
6. A. B. John et al., "Trends in integrated multi-product biorefinery systems: A comprehensive review, " Renew. Sust. Energ. Rev., vol. 2, pp. XXX-XXX, 2025, doi: https://doi.org/10.20935/AcadEnergy7605
7. United Nations. https://www.un.org/sustainabledevelopment/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Integrated Data-Driven Optimisation of LNG Hot Section for Energy Efficiency and Decarbonization

Aisha Al-Hammadi[a]*, Dr Tareq Al-Ansari[a], Dr Ahmed AlNouss[a], Abdul Aziz Shaikh[a]

[a] College of Science and Engineering, Hamad Bin Khalifa University (HBKU), Qatar Foundation, Doha, Qatar
* Corresponding Author: aial34241@hbku.edu.qa.

ABSTRACT

In today's competitive LNG market, reducing energy consumption is critical for enhancing both profitability and sustainability. The hot section of the LNG processing, which includes inlet receivers, acid gas removal, and dehydration units, is the most thermally demanding. Previous optimisation methods targeted each unit separately. On the other hand, this work details the development of a data-driven optimisation framework to minimise energy across these interdependent units. Preliminary application of the framework has yielded encouraging results. Utilising HYSYS process simulation data, the study successfully identifies critical operating variables—such as reboiler duty, amine circulation rate, and air-to-furnace stoichiometry—that drive production efficiency and energy consumption. Results indicate that a baseline condensate mass flow of 2, 048.71 kg/h is achieved at a stripper bottom temperature of 137.74 °C, while the AGRU produces sweet gas with 0.18 ppm H_2S. Optimisation using Pareto frontier analysis reveals a "knee" point in the SRU at 96.9% recovery efficiency, balancing elemental sulfur production (7, 622.24 kg/h) against heating loads. The findings provide mathematical coefficients for plant control to maximise yield while maintaining strict environmental stack limits.

Keywords: Natural Gas, Liquified Natural Gas, Optimization, LNG Optimization, Data-driven optimization, Hot section, Energy Efficiency.

INTRODUCTION

The liquified natural gas (LNG) industry faces pressure to reduce energy intensity as global markets demand both economic competitiveness and environmental sustainability. The liquefaction process is highly energy-intensive, consuming a significant portion of the feed gas's energy, mainly to power the large compression systems that drive refrigeration cycles. This high energy consumption results in both high operating costs and significant CO2 emissions. For a typical LNG facility, even a small reduction in specific energy consumption results in significant annual savings and reduces the carbon footprint.

LNG production involves two process sections (hot section and cold section). The hot section operates at ambient temperature to pre-treat the feed gas, removing impurities that would otherwise freeze in cryogenic equipment. As shown in Figure 1, the hot section consists of an inlet receiver, a condensate stabilisation unit, an acid gas removal unit (AGRU), and molecular sieves dehydration. The treated gas enters the cold section for liquefaction through different cooling cycles, such as propane pre-cooling and mixed refrigerant cycles, reaching -162 ºC. The hot section accounts for a substantial portion of total plant energy consumption through thermal duties in amine regeneration reboilers, condensate stabilisation columns, and molecular sieve regeneration systems [1].

While research has focused on optimising the cold section (cryogenic liquefaction processes), the hot section has received limited attention. A critical, overlooked characteristic is the high interdependence among the hot-section units. Condensate stabiliser overhead composition affects molecular sieve loading, acid gas removal efficiency influences downstream dehydration requirements, and heat integration opportunities create coupling through the thermal utility system. Despite these interdependencies, traditional optimisation approaches have mainly focused on individual unit operations in isolation, leading to suboptimal overall plant performance.

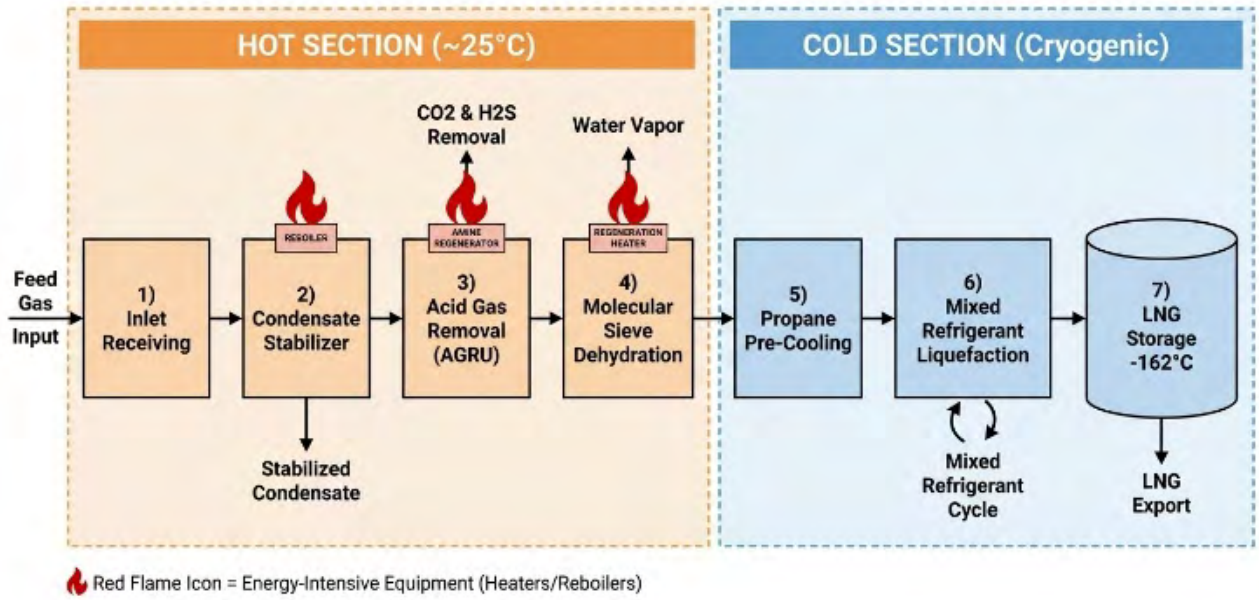

Figure 1. LNG plant hot and cold sections configuration.

This work addresses the gap between traditional unit-specific optimisation and the operational reality of interdependent systems. Rather than employing conventional model-based optimisation that requires extensive first-principles modelling, this study develops a data-driven optimisation framework that uses simulation data to identify optimal operating strategies.

LITERATURE REVIEW

Hot Section Energy Intensity

The hot section consists of several thermally intensive units. The acid gas removal unit (AGRU), employing amine-based absorption-regeneration cycles, is the most significant energy consumer, with large reboiler duties for solvent regeneration [2]. Condensate stabilisation separates light hydrocarbons from heavier liquids that require significant thermal energy, while molecular sieve dehydration requires periodic regeneration with hot gas at elevated temperatures [1].

These units show significant interdependencies. Feed composition variations affect all units; for instance, changes in CO_2 content alter amine circulation requirements and reboiler duty, while variations in heavy hydrocarbon content affect both condensate stabilisation loads and downstream molecular sieve performance. Heat integration opportunities further couple the units through the thermal utility system [3].

Conventional Optimisation Approaches

LNG process optimisation has historically relied on employing detailed first-principles models in platforms. Chew et al. (2022) developed an AGRU model and optimised operating conditions using a hybrid genetic algorithm and sequential quadratic programming, achieving a 22% reduction in reboiler duty. Shaikh et al. (2022) applied pinch analysis, demonstrating that waste heat recovery could supply a substantial portion of molecular sieve regeneration energy requirements.

First-principles approaches face limitations, such as the burden of extensive model development and calibration, computational expense when optimising large flowsheets, difficulty in translating theoretical optimal solutions into practical operational changes, and a static nature that provides limited guidance for real-time adaptation to changing conditions [4]. Industrial practice often relies on heuristic optimisation based on operational practice, which has limited exploration, difficulty handling variable interactions, and a tendency to underestimate trade-offs.

Research Contribution

This work addresses critical gaps: (1) limited data-driven optimisation in LNG hot section despite the successful applications in other industries, (2) inadequate consideration of units' interdependencies with most studies addressing individual units separately, and (3) disconnect between academic optimisation and industrial application due to computational complexity and model validation challenges. The primary challenge lies in the competitive nature of the optimisation objectives. For instance, increasing condensate stabilisation to meet

Reid Vapour Pressure (RVP) limits inherently reduces mass yield by vaporising light hydrocarbons. Similarly, the AGRU must balance solvent circulation against regeneration energy, while the SRU requires precise stoichiometric control to prevent H_2S breakthrough. This report quantifies these trade-offs through rigorous sensitivity mapping and optimisation formulation.

The study is conducted using a steady-state HYSYS simulation framework. The methodology followed a four-step process:

1. Baseline Validation: Current plant data for condensate stabilisation, AGRU, and SRU are used to establish the baseline.
2. Framework Analysis: The simulation utilised the Sour-PR and Sulfinol-D fluid packages for accurate phase behaviour and acid gas kinetics.
3. Sensitivity Studies: Operating parameters are varied by ±10% to ±40% to observe the impact on Objective Functions (Mass Flow, Duty, Purity).
4. Optimisation Formulation: Competitive objectives were plotted to identify the Pareto Frontier.

The simulation is based on a high-pressure feed gas arriving at **7, 800 kPa**. Key baseline values extracted from the simulation workbooks are summarised below:

Table 1: Key baseline values from simulation

Parameter	Value
Condensate Mass Flowrate	2, 048.71 kg/h
Condensate Stripper Duty	1, 766, 960 kJ/h
Stripper Bottom Temperature	137.74 °C
Sweet Gas Mass Flowrate from AGRU	412, 312.66 kg/h
AGRU Solvent Circulation Rate	1, 319, 223.78 kg/h
AGRU Regenerator Reboiler Duty	1.81×10^9 kJ/h
AGRU Overall Duty	2.49×10^9 kJ/h
CO_2 Content in AGRU Sweet Gas	81.5 ppm
H_2S Content in AGRU Sweet Gas	0.18 ppm
Lean Solvent Temperature to Absorber	55 °C
H_2S mol% Entering SRU (Acid Gas)	48.7 mol%
H_2S ppm Content from SRU (Stack)	4.26 ppm
SRU Overall Duty	1.22×10^9 kJ/h
Elemental Sulphur Production	7, 622.24 kg/h

Sensitivity Analysis

Sensitivity curves depict the competitive relationships among production capacity, energy consumption, and product purity across the different units of the LNG hot section.

Figure 2 represents the trade-off between maximising mass flow rate and meeting stabilisation requirements. Reboiler Duty varied within +-20% range to study Condensate Mass Flow and RVP. The curves show a decreasing trend of condensate flow and RVP with increasing reboiler duty.

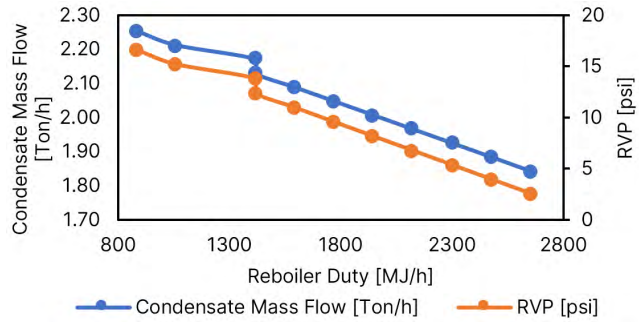

Figure 2. Sensitivity analysis of condensate stripper reboiler duty against condensate production and RVP.

Figure 3 shows the effect of decreasing reboiler duty on stripper bottom temperature and its dependence on feed temperature.

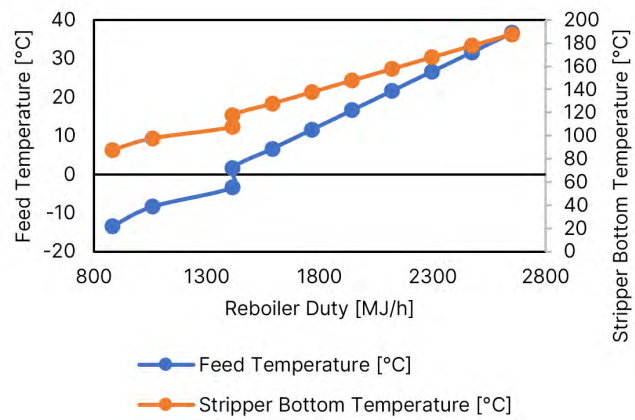

Figure 3. Sensitivity analysis of condensate stripper reboiler duty against feed and stripper bottom temperatures.

Sensitivity analysis shows that mass flow is highly sensitive to the stripper bottom temperature. A 10 °C increase in bottom temperature results in a ~2% reduction in mass yield due to the vaporisation of butane into the overhead.

Figures 4 and 5 show the competitive nature of amine circulation; more solvent improves purity but significantly increases pumping and regeneration duty. Amine rate varied within +-20% range to study CO2 content in sweet natural gas, regen, and pump duties.

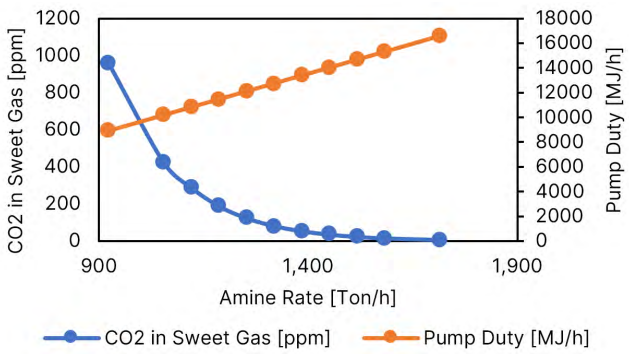

Figure 4. Sensitivity analysis of amine circulation against amine pumping duty and CO2 content in sweet natural gas.

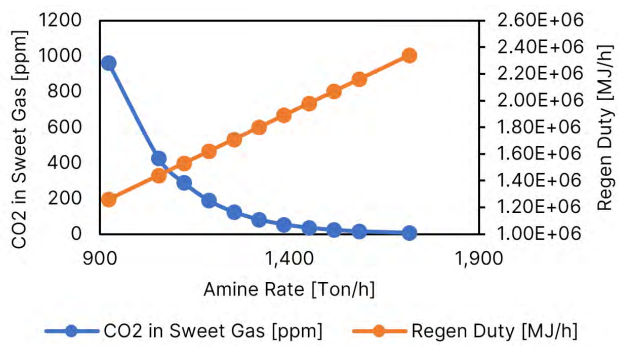

Figure 5. Sensitivity analysis of amine circulation against regeneration duty and CO2 content in sweet natural gas.

Similarly, increasing the amine rate enhances H2S removal from natural gas but reduces sweet natural gas production, as illustrated in Figure 6.

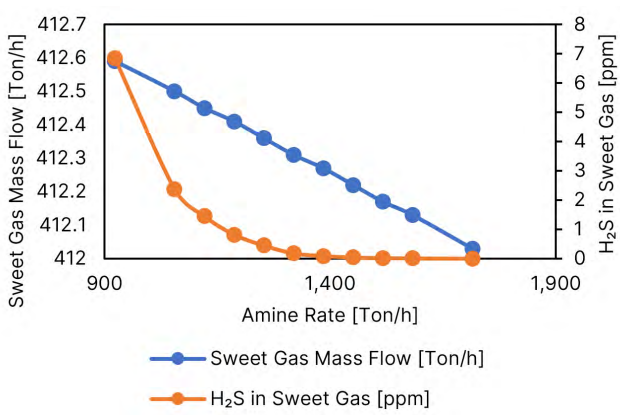

Figure 6. Sensitivity analysis of amine circulation against sweet gas flow and CO2 content in sweet natural gas.

The amine circulation rate shows diminishing returns in terms of gas purity. Increasing the rate from 1, 319 to 1, 583 Ton/h improves H_2S removal from 0.18 ppm to 0.01 ppm but increases pumping duty by 20%.

Figure 7 shows the most critical "competitive" parameter in sulfur recovery: the Air/Acid Gas ratio must be precisely balanced to maintain the 2:1 H2S: SO2 stoichiometric ratio. Air flow varied within ±20 % to study H2S content in the stack and sulfur production.

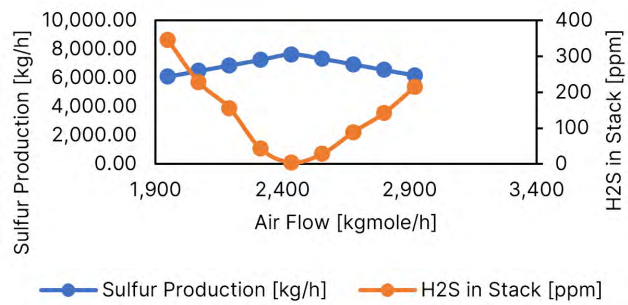

Figure 7. Sensitivity analysis of air flow against sulfur production and H2S content in the stack.

Figure 8 represents another critical parameter in sulfur recovery, where recovery efficiency and net heating duty are measured against air flow variation. Recovery efficiency in SRU is calculated as per Eq 1.

$$\eta_{SRU} = \frac{m_{S-produced}}{m_{H_2S-inlet}} \quad \text{Eq 1}$$

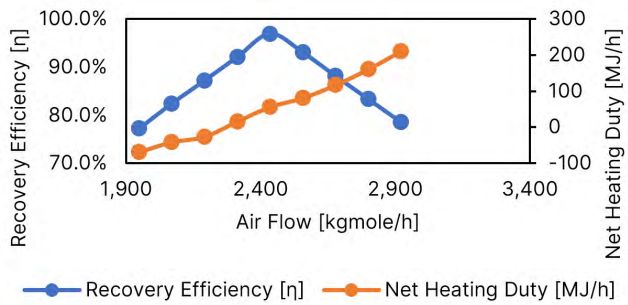

Figure 8. Sensitivity analysis of air flow against recovery efficiency and net heating duty.

Optimization Formulation

The following linear and non-linear optimisation formulations are derived from the sensitivity analysis of simulation data points to achieve the desired optimisation goals.

1. Maximise Field Condensate Mass Flowrate

Objective Function:

$$\boldsymbol{Max\ m_c = C_o + a_1(Q) + a_2(RVP)} \quad \text{Eq 2}$$

Where,

Co = 2457.4

a1 = -0.23

a2 = -0.00313

Subject to:

$$RVP \leq 10\ psi \quad \text{Eq 3}$$

$$T_{bottom} \leq Thermal\ stability\ limit\ of\ hydrocarbons$$

2. Maximise Sweet Gas Mass Flowrate from AGRU

Objective Function:

$$\boldsymbol{Max\ m_{NG} = C_1 + b_1(m_{amine}) + b_2(Duty_{pump}) + b_3(Duty_{regen})}$$

Eq 4

Where,

C1 = 413.24

b1 = 0.0121

b2 = -0.00077

B3 = -3.9x10-6

Subject to:

$H_2S \leq 4\ ppm$ Eq 5

$CO_2 \leq 50\ ppm$ Eq 6

3. Maximise SRU Efficiency

Objective Function:

$$\boldsymbol{Max\ \eta_{SRU} = \frac{m_{S-produced}}{m_{H_2S-inlet}} = C_2 + d_1(m_{Air}) + d_2(Duty_{net})}$$

Eq 7

Where,

C2 = 5.171

d1 = -0.00232

d2 = 6.59x10-6

Subject to:

$H_2S\ in\ stack \leq 270\ ppm$ Eq 8

OPTIMIZATION RESULTS

The optimisation problems represent critical competitive challenges in the LNG hot section.

For the condensate stabilisation section, the multi-objective optimisation yielded a Pareto Frontier between the condensate production rate and the required reboiler duty, as illustrated in Figure 9. Increasing reboiler duty stabilises the product (lowers RVP) but vaporises more hydrocarbons into the off-gas.

For the acid gas removal section, the multi-objective optimisation yielded a Pareto Frontier between sweet natural gas production and required solvent regeneration duty as illustrated in Figure 10. Increasing regeneration duty purifies the product (lowering CO2 and H2S content) but decreases sweet gas production due to greater absorption.

The most critical competitive optimisation is the SRU recovery efficiency versus external heating load. At a sulfur production rate of 7, 622 kg/h, the unit achieves 96.9% efficiency, as illustrated in Figure 11. Attempting to reach 97.8% recovery requires an additional 16% increase in heater duty, exceeding the energy provided by the Waste Heat Exchanger (1.64E+08 kJ/h).

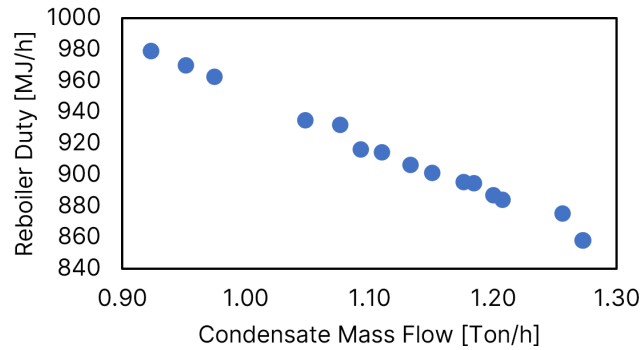

Figure 9. Pareto Frontier of condensate production and reboiler duty.

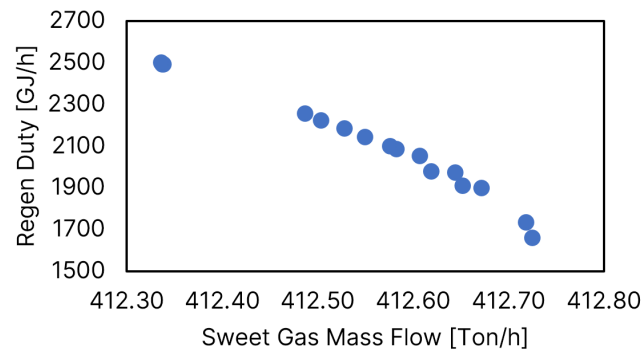

Figure 10. Pareto Frontier of condensate production and reboiler duty.

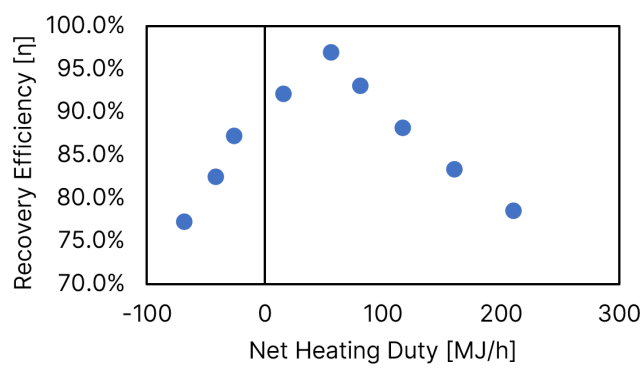

Figure 11. Optimisation curve of SRU recovery efficiency and net heating load.

Comparison with Literature

The achieved 0.18 ppm H_2S in sweet gas is significantly better than the typical 4.0 ppm pipeline limit, indicating the AGRU is conservatively designed. In addition, the 96.9% SRU efficiency is standard for a two-stage Claus unit. To achieve higher rates (99%+), the literature suggests adding a Tail Gas Treatment Unit (TGTU).

CONCLUSION

Optimising the integrated facility highlights the baseline operating point as the Pareto-optimal knee. Specifically, the Condensate Unit should maintain a bottoms temperature of 137.7 °C to maintain a balance between a 2, 048 kg/h flow rate and a 9.6 psi RVP. AGRU should operate at 1, 319 Ton/h amine rate; further increases lead to excessive reboiler duty without meaningful purity gains. The SRU must be strictly controlled at 2, 435.34 kgmole/h air flow to maintain stack emissions at 4.26 ppm H_2S.

REFERENCES

1. Bouabidi Z, Katebah MA, Hussein MM, Shazed AR, Al-musleh EI. Towards improved and multi-scale liquefied natural gas supply chains: thermodynamic analysis. Computers & Chemical Engineering 151:107359 (2021). https://doi.org/10.1016/j.compchemeng.2021.107359
2. Chew YE, Putra ZA, Foo DCY. Process simulation and optimisation for acid gas removal system in natural gas processing. Journal of Natural Gas Science and Engineering 107:104764 (2022). https://doi.org/10.1016/j.jngse.2022.104764
3. Shaikh AA, AlNouss A, Al-Ansari T. A heat integration case study for the dehydration and condensate stabilization units in LNG plants for economic and energy savings. Computers & Chemical Engineering 168:108062 (2022). https://doi.org/10.1016/j.compchemeng.2022.108062
4. Katebah MA, Hussein MM, Shazed A, Bouabidi Z, Al-musleh EI. Rigorous simulation, energy and environmental analysis of an actual baseload LNG supply chain. Computers & Chemical Engineering 141:106993 (2020). https://doi.org/10.1016/j.compchemeng.2020.106993
5. Zhu J, Zhang W, Liu S, Li Y, Liu M, Yin Q, Xie B, Yu X, Li Y. Experiment and dynamic simulation study on propane pre-cooling double nitrogen-expander liquefaction process for medium-pilot LNG plant. Applied Thermal Engineering 170:114994 (2020). https://doi.org/10.1016/j.applthermaleng.2020.114994
6. Lee JH, Shin J, Realff MJ. Machine learning: overview of the recent progresses and implications for the process systems engineering field. Computers & Chemical Engineering 114:111-121 (2018). https://doi.org/10.1016/j.compchemeng.2017.10.008

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Modeling and Optimization of Sonochemical Reactors through simulations

Nikolaos I. Vittas[a], Antonios Armaou[ab*]

[a] Chemical Engineering Department, University of Patras, Patras 26504, Greece
[b] Chemical Engineering Department, The Pennsylvania State University, University Park 16802, PA, USA
* Corresponding Author: armaou@upatras.gr .

ABSTRACT

Sonochemical reactors are a promising technology in process intensification, offering a sustainable and energy-efficient means of enhancing chemical reactions. By harnessing acoustic cavitation – the formation, oscillation and violent collapse of bubbles in a liquid medium – these systems generate local hotspots that can accelerate reaction kinetics. Despite its potential, efficient design and scale-up of sonochemical reactors remain major challenges, mostly because the cavitation phenomena take place close to the ultrasonic transducer. This work presents a simulation-based framework for the optimization of sonochemical batch reactors by coupling microscopic-level bubble behavior with macroscopic-level reactor performance, focusing on the placement of transducers to maximize reaction activity.

Keywords: Batch Process, Modelling and Simulations, Optimization, Sonochemistry, Acoustic Cavitation

INTRODUCTION

Sonochemistry relies on the phenomenon of acoustic cavitation, i.e. the type of cavitation induced in a liquid medium by ultrasound or acoustic waves [1]. When such a wave propagates through a liquid it causes cavities (bubbles) to form, oscillate and finally collapse [2, 3]. During the rarefaction phase of the wave, cavities form and preexisting bubbles expand, while during the compression phase the bubbles contract. This oscillatory pattern continues until the prevailing conditions drive the bubble into a violent collapse. The result of this process is the creation of localized hotspots, characterized by extreme pressure and temperature, which in turn enhances chemical reactions.

For accurate modeling of sonochemical reactors, several physical phenomena must be considered in combination: wave propagation inside the liquid, bubble dynamics under the ultrasound, heat diffusion from the collapsing bubbles and reaction kinetics. The bubble dynamics were captured using the Keller–Miksis model, accounting for liquid compressibility, while the Helmholtz equation described wave propagation within the reactor, incorporating viscous, thermal, and acoustic damping. Reaction kinetics for the Fenton assisted degradation of cresols were modelled as a pseudo–first-order process governed by Arrhenius temperature dependence.

BUBBLE DYNAMICS

The dynamic oscillations of a single bubble in a liquid medium under the influence of applied ultrasound can be described based on the fundamental principles of force and energy equilibriums, which are employed to derive the Rayleigh-Plesset equation. Following this, the refinement suggested by Keller and Miksis will be presented to account for more complex phenomena and to provide a more accurate representation of the oscillations and collapse. Finally, the effect of multiple acoustic waves on single-bubble dynamics will be examined.

Force Equilibrium

A typical cavitation bubble, assumed to pre-exist in the liquid medium, consists of both vapor and gas. Consequently, the total pressure inside the bubble (p_{in}) is the sum of the vapor pressure (p_v) and the gas pressure (p_g).

$$p_{in} = p_v + p_g \tag{1}$$

The vapor pressure is assumed to remain constant, whereas the gas pressure undergoes significant changes and can be approximated as:

$$p_g\, V_b{}^k = constant \tag{2}$$

where $V_b = {}^4/_3\, \pi R^3$ is the volume and R the radius of the bubble. The exponent k represents the heat capacity ratio and the specific value depends on the model approach and the material. In the case of an adiabatic process involving air, the ratio can approximately be:

$$k = \frac{C_p}{C_v} \cong 1.4 \tag{3}$$

For the bubble to expand by dR, the work required is $\sigma(8\pi R\, dR)$, where σ is the surface tension. Using the above definitions, the force equilibrium on the bubble wall can be expressed as:

$$p_{in}\,(4\pi R^2) = p_b(2\pi R^2) + \sigma(8\pi R) \tag{4}$$

where p_b is the pressure on the bubble wall.

It is assumed that the bubble is oscillating from an equilibrium point with radius R_0 and the static ambient pressure is $p_\infty = p_0$. Combining the above equations provides an expression for the pressure on the bubble wall.

$$p_b = \left(p_\infty + \frac{2\sigma}{R_0} - p_v\right)\left(\frac{R}{R_0}\right)^{3k} + p_v - \frac{2\sigma}{R} \tag{5}$$

The effect of the acoustic wave, represented as time dependent variable $p_s(t)$, is incorporated into the ambient pressure term

$$p_\infty = p_0 + p_s(t) \tag{6}$$

In most typical cases the pressure wave is sinusoidal with pressure amplitude of P_a and frequency f, superimposed on hydrodynamic pressure p_0:

$$p_s(t) = p_0 - P_a\ sin(2\pi f t) \tag{7}$$

Energy Equilibrium

The conservation of energy can be described as: the work of the oscillating bubble W_b, either positive for expansion or negative for contraction and collapse, must be equal to the work W_l plus kinetic energy E_k required by the liquid volume surrounding the cavity.

$$W_b = W_l + E_k \tag{8}$$

Each of these terms is calculated separately:

$$W_b = p_b \frac{4}{3}\pi\left(R^3(t) - R_0{}^3\right) \tag{9}$$

$$W_l = p_\infty \frac{4}{3}\pi\left(R^3(t) - R_0{}^3\right) \tag{10}$$

$$E_k = 2\pi\rho_L R^3(t)\left(\frac{dR(t)}{dt}\right)^2 \tag{11}$$

Rayleigh-Plesset

The combination of equations (5)-(11) yields the Rayleigh-Plesset equation, describing the radius of the bubble in respect to time.

$$\frac{3}{2}\left(\frac{dR(t)}{dt}\right)^2 + R(t)\frac{d^2R(t)}{dt^2} =$$
$$= \frac{1}{\rho_L}\left(p_g + p_v - \frac{2\sigma}{R(t)} - 4\mu\frac{1}{R(t)}\frac{dR(t)}{dt} - p_0 - p_s(t)\right) \tag{12}$$

Keller-Miksis

The Rayleigh – Plesset equation is foundational for any work done in acoustic cavitation. However, the main assumptions made, that the liquid is incompressible and the bubble remains spherical, are less than accurate for such high-speed phenomenon. To address these limitations, the suggested modifications by Keller and Miksis [4, 5] account mainly for the compressibility of the liquid, as well as effects of viscosity, surface tension and acoustic radiation

$$\left[1 - \frac{1}{3}(3\Lambda + 1)\frac{1}{c_L}\frac{dR(t)}{dt}\right]\frac{3}{2}\left(\frac{dR(t)}{dt}\right)^2 +$$
$$+\left[1 - (\Lambda + 1)\frac{1}{c_L}\frac{dR(t)}{dt}\right]R(t)\frac{d^2R(t)}{dt^2} =$$
$$= \left[1 + (1 - \Lambda)\frac{1}{c_L}\frac{dR(t)}{dt} + \frac{R}{c_L}\frac{d}{dt}\right]\frac{p_b - p_\infty}{\rho_L} \tag{12}$$

In the present work the parameter $\Lambda = 0$, referring to the Keller form of the equation.

Collapse

The analysis of the transient collapse of the oscillating gas filled cavity [6] assumes that the bubble reaches a maximum radius R_{max} just before the violent collapse to a minimum radius $R_{min} \to 0$. The collapse is characterized by the bubble wall reaching its maximum velocity under the assumption of adiabatic compression of the internal gas. As a result, the gas reaches exceedingly high pressures and temperatures during the collapse. The maximum temperature reached during the process can be determined, using the initial temperature T_0, by the following relationship [7].

$$T_{max} = T_0\left(\frac{R_{max}}{R_{min}}\right)^{3(\gamma-1)} \tag{13}$$

Different Wave Sources

A preliminary approach to analyze the effect of multiple wave sources on the single bubble dynamics can be made via simple modifications to the existing model. The equation (7) can be modified to account, for instance, two distinct waves characterized by amplitudes and frequences P_{a1}, f_1 and P_{a2}, f_2 respectively [8, 9].

$$p_s(t) = p_0 - P_{a1} sin(2\pi f_1 t) - P_{a2}\ sin(2\pi f_2 t) \tag{14}$$

WAVE PROPAGATION

Analyzing wave propagation in a liquid medium containing dispersed bubble population requires a comprehensive approach. One of the most important parameters

in the cavitation theory is the void fraction β. It accounts for the total volume occupied by cavities inside the control volume V:

$$\beta = \frac{4}{3}\pi R_0^3 n \tag{15}$$

where $n = {}^{N}/_{V}$ is bubble number density and R_0 is the equilibrium radius that all bubbles are assumed to have [10].

Mass Conservation

The conservation of mass for a mixture of liquid and dispersed bubbles is expressed as:

$$\frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \bullet (\rho \boldsymbol{v}) = 0 \tag{16}$$

where $\boldsymbol{v}$ is the velocity and ρ is the density of the mixture, based on the void fraction

$$\rho = \rho_L(1-\beta) + \rho_g\beta \cong \rho_L(1-\beta) \tag{17}$$

It is apparent that the mixture density approaches the liquid density ρ_L, since the gas density ρ_g is negligibly small.

Assuming (a) negligible bubble coalescence and fragmentation, (b) bubbles move at the same velocity as the liquid and (c) the void fraction is of the order of 10^{-4} or smaller, the equation for conservation of mass is:

$$\frac{1}{\rho_L}\frac{D\rho_L}{Dt} + \boldsymbol{\nabla} \bullet v = 4\pi R^2 \frac{DR}{Dt} \tag{18}$$

Energy Conservation

The momentum conservation is ensured via Cauchy relationship:

$$\rho_L\left[\frac{\partial \boldsymbol{v}}{\partial t} + \boldsymbol{v} \bullet \boldsymbol{\nabla}\boldsymbol{v}\right] = \boldsymbol{\nabla} \bullet \bar{\bar{\boldsymbol{\sigma}}} + \boldsymbol{f} \tag{19}$$

Neglecting viscous terms expressed via the stress tensor $\bar{\bar{\boldsymbol{\sigma}}}$ and the body forces given via $\boldsymbol{f}$ term, thereby considering only for pressure forces, the momentum conservation equation is simplified to:

$$\rho_L\left[\frac{\partial \boldsymbol{v}}{\partial t} + \boldsymbol{v} \bullet \boldsymbol{\nabla}\boldsymbol{v}\right] - \boldsymbol{\nabla} p = \boldsymbol{0} \tag{20}$$

Wave equation

The non-homogeneous wave equation can be derived through the nondimensionalization of the governing mass and momentum conservation equations. Following the work done by Calfisch and al. [11] the non-homogeneous wave equation can be obtained:

$$\frac{1}{{c_L}^2}\frac{\partial^2 p}{\partial t^2} - \nabla \bullet \nabla p - \frac{\partial^2 \beta}{\partial t^2} = 0 \tag{21}$$

It can be proven that the wave propagation in the bubbly liquid is following the Helmholtz's equation

$$\Delta P(t,\boldsymbol{r}) + {k_m}^2 P(t,\boldsymbol{r}) = 0 \tag{22}$$

where the complex wavenumber, a monodisperse bubble population with equilibrium radius R_0, for is defined by the relationship:

$${k_m}^2 = \frac{\omega^2}{{c_L}^2}\left(1 + 4\pi\,\frac{{c_L}^2\frac{N}{V}R_0}{\omega_0 - \omega^2 + 2ib\omega}\right) \tag{22}$$

The wave number incorporates variables like resonance frequency (ω_0), damping constant (b) that accounts for viscous, thermal and acoustic effects and velocity potential (Φ) given by the following relationships:

$$\omega_0 = \frac{p_0}{\rho_L {R_0}^2}\left[Re(\Phi) - \frac{2\sigma}{R_0 p_0}\right] \tag{23}$$

$$b = \frac{2\mu}{\rho_L {R_0}^2} + \frac{p_0}{2\rho_L {R_0}^2} Im(\Phi) + \frac{\omega^2 R_0}{2c_L} \tag{24}$$

$$\Phi = \frac{3\gamma}{1 - 3(\gamma-1)i\frac{D_g}{\omega {R_0}^2}\left[\sqrt{i\frac{\omega {R_0}^2}{D_g}}\,coth\left(i\frac{\omega {R_0}^2}{D_g}\right) - 1\right]} \tag{25}$$

REACTION

Fenton oxidation is a promising method for the degradation of pollutants like phenols. In Fenton process, hydroxyl radicals OH^- are produced from hydrogen peroxide H_2O_2 catalyzed by iron $Fe(II)$.

$$Fe^{2+} + H_2O_2 \rightarrow Fe^{3+} + OH^- + OH\bullet \tag{26}$$

The generated radicals can react effectively with a variety of organic compounds RH.

$$RH + OH\bullet \rightarrow H_2O + R\bullet \rightarrow \textit{further oxidation} \tag{27}$$

$$R\bullet + H_2O_2 \rightarrow ROH + OH\bullet \tag{28}$$

Batch Reactor

Considering a well-mixed mixture, two areas of interest can be assumed. The first one is where the cavitation takes place, which is essentially the part of the liquid that has elevated temperature due to a bubble collapse, thus an enhanced reaction rate. The second one is the rest of the bulk liquid having an ambient temperature, thus a base reaction rate. The areas are symbolized with c and B respectively. The mass equilibrium for a product with concentration, C, may then be written as:

$$\frac{dC}{dt} = \beta r_c + (1-\beta) r_B \tag{29}$$

where β is the void fraction and r the rate of the reaction.

Effective Rate

Although this theory may be applied to different reactions, this work focuses on the degradation of cresols, due to the reaction being well documented under a variety of conditions.

$$C_7H_8O + 17H_2O_2 \rightarrow 7CO_2 + 21H_2O \tag{30}$$

Assuming abundance of H_2O_2, a pseudo-first-order reaction rate for the organic compound is given by:

$$r = -kC \tag{31}$$

where the constant k has an Arrhenius dependency of temperature:

$$k = A\exp\left(\frac{E_A}{RT}\right) \tag{32}$$

In the bulk liquid with an ambient temperature T_0, the rate constant would be:

$$k_B = A\exp\left(\frac{E_A}{R_g T_0}\right) \tag{33}$$

While in the area affected by cavitation phenomena, heat diffusion in the surrounding liquid must be considered. Thus, the rate constant would be given by the following equation:

$$k_{eff} = \frac{\int_t \int_V k(t,\boldsymbol{r})dVdt}{\int_t \int_V dVdt} \tag{34}$$

By integrating Eq. (29) a more convenient variable for comparison can be derived, hereafter referred to as effective rate r_{eff}:

$$-ln\left[\frac{C}{C_0}\right] = \left[\beta k_{eff} + (1-\beta)k_B\right]t \tag{35}$$

where:

$$r_{eff} = \beta k_{eff} + (1-\beta)k_B \tag{36}$$

SIMULATION DETAILS

Geometry

The sonochemical reactor is a cylindrical batch reactor, open at the top, with a cylindrical transducer submerged from the top. The ultrasound is introduced into the system from the tip of the transducer.

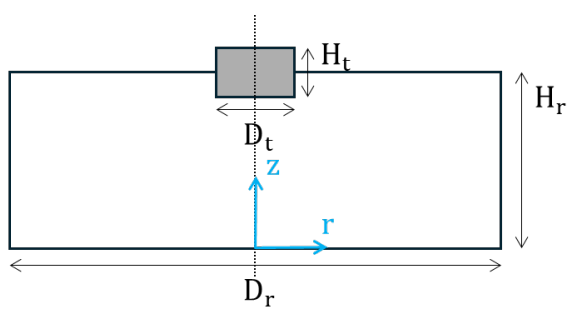

Figure 1. Geometric characteristics of the reactor.

The reactor is filled with an aquatic solution. The geometry of the reactor is described in Figure 1 and Table 1

Table 1. Geometrical characteristics of the reactor.

Symbol	Value
D_r	0.5 m
D_t	0.08 m
H_r	0.2 m
H_t	0.1 m

Parameters

The liquid medium is water at standard temperature and pressure, thus $T_0 = 298.15$ K and $p_0 = 1.01325 \cdot 10^5$ Pa. The physical properties of water are given in Table 2.

Table 2. Physical properties of water.

Symbol	Value
Γ	4/3
Σ	$0.0725\ \mathrm{N \cdot m^{-1}}$
M	$1.0016 \cdot 10^{-3}\ \mathrm{Pa \cdot s}$
ρ_L	$1000\ \mathrm{kg \cdot m^{-3}}$

The acoustic wave travels at the frequency of $f = 20$ kHz and with an amplitude that is calculated by the following relationship [8, 12]:

$$P_a = \sqrt{2\rho_L c\frac{P}{A_t}} \tag{37}$$

where P is the power of the transducer set to 40 W and $A_t = 0.005\ \mathrm{m^2}$ the area of the tip of the transducer. Thus, the amplitude of the wave is set to be $P_a = 1.53 \cdot 10^5$ Pa.

In terms of the reaction, for p-cresols the Arrhenius constants are $A = 0.1419\ \mathrm{s^{-1}}$ and $E_a = 14.25 \cdot 10^3\ \mathrm{J\ mol^{-1}}$.

It should be noted that equilibrium radius of the single bubble is considered to be $R_0 = 5\ \mu\mathrm{m}$ and the void fraction is calculated via the following empirical relationship [13]:

$$\beta = N \cdot |p| \tag{38}$$

where $N = 2 \cdot 10^{-9}$.

RESULTS

The results of the simulation can be divided into two categories: the single-bubble environment, which accounts for the enhancement of the reaction through an increase of the rate constant and wave propagation inside the batch reactor, which reflects the intensity of cavitation phenomena through an increase of the void fraction.

Single-Bubble Environment

The effect of one wave and two coherent waves on the single bubble have been simulated. Figures 2 and 3 demonstrate the radius changes for each case.

The bubble grows 15 and 50 times its equilibrium size respectively and then it rapidly collapses before the completion of on cycle. The temperature reached by the bubble during the collapse, based on the maximum radius the bubble reaches, is 4780 K for one wave and 13987 K for two waves. Applying a simple conduction model, the area effected by cavitation is 2 times the equilibrium radius and this effect is sustained for 10 periods for both cases.

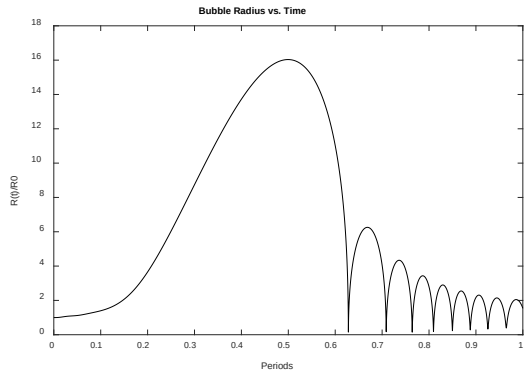

Figure 2. Bubble radius as a function of time under one acoustic wave.

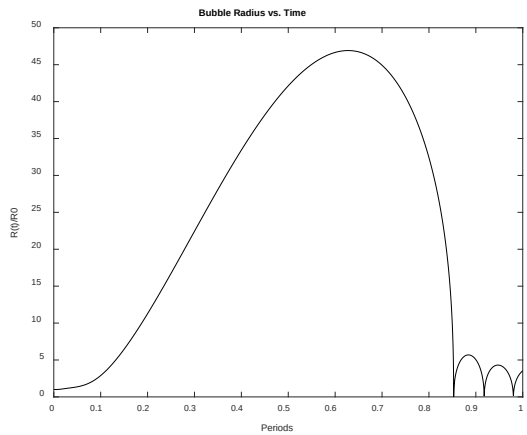

Figure 3. Bubble radius as a function of time under two identical acoustic waves

The rate constant in these conditions can be calculated to be $0.2788\ \mathrm{s}^{-1}$ and $0.7733\ \mathrm{s}^{-1}$ respectively, while in bulk conditions the constant is $4.5213 \cdot 10^{-4}\ \mathrm{s}^{-1}$. This accounts for an increase of 615% and 1700% for the rate constant in the cavitation zone compared to the bulk of the liquid.

One transducer

The first simulation examines a transducer placed at the center of the top of the reactor as shown in Figure 4.

The propagation of the wave in the liquid medium was described by the pressure distribution inside the reactor, while the void fraction demonstrated where cavitation phenomena occur. The YZ cut-plane of the reactor is presented for both parameters in Figures 5 and 6.

The pressure wave appears to have an effect mostly close to the transducer, although small pressure variations exist everywhere inside the reactor. The higher void fraction, in the order of $9 \cdot 10^{-4}$, exists just under the transducer. The mean void fraction was calculated to be $5.32 \cdot 10^{-5}$.

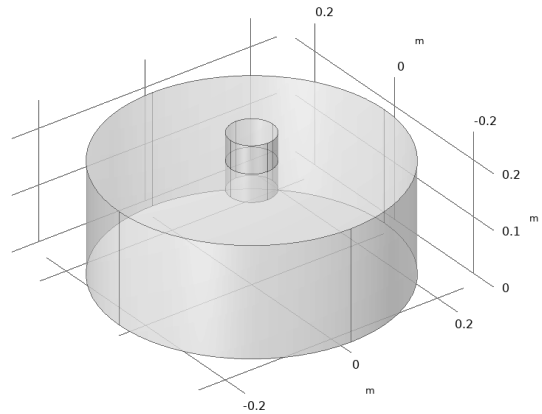

Figure 4. Geometry of one transducer at the center of the top of the reactor.

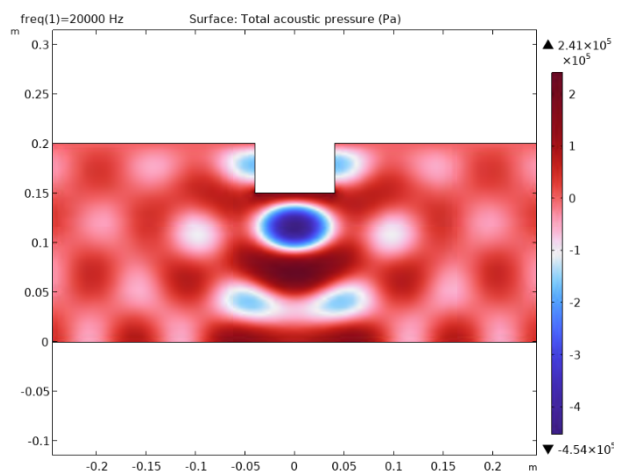

Figure 5. Pressure distribution for one transducer.

Given this result and the rate constant for one wave, the effective rate was calculated to be $4.648 \cdot 10^{-4}\ s^{-1}$. This is an increase of 2.8% in the reactions rate. Literature data suggest that the rate is around $4.5 \cdot 10^{-4} s^{-1}$, depending on the exact conditions [14, 15].

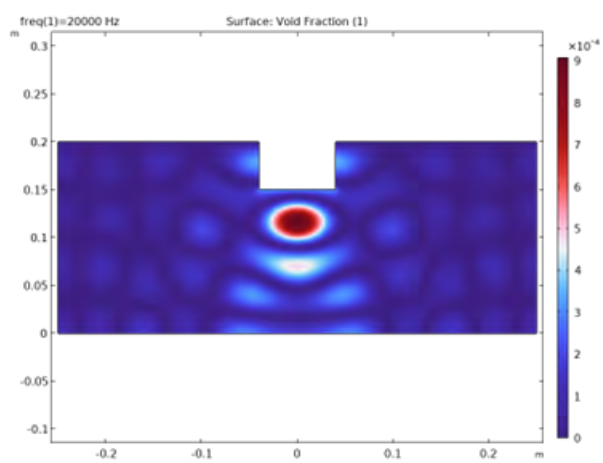

Figure 6. Void fraction for one transducer.

Two transducers

The simplest geometry to simulate the wave propagation from two different transducers is one at the center top of the reactor (as before) and the other at the center bottom of the reactor. Note that for practical reasons the tip of the bottom transducer is placed on the same plane as the bottom of the reactor.

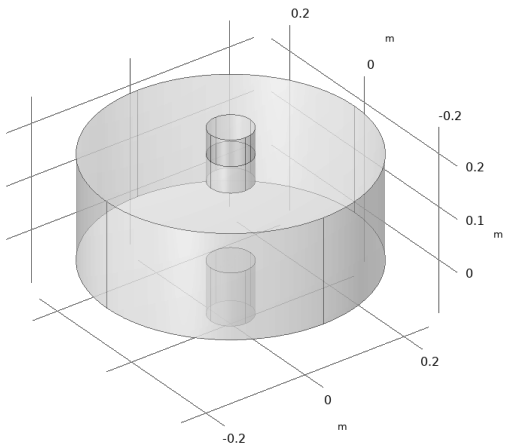

Figure 7. Geometry for two transducers.

The same YZ-plane plots for the pressure and void fraction distributions are presented in Figures 8 and 9 respectively. Both figures show little to no changes for the addition of another transducer, just a slightly increased intensity of the variations in the central axis. The mean void fraction was calculated as $4.82 \cdot 10^{-5}$, lower than previously.

Despite that, the two waves introduced into the system affect the single bubble dynamics differently, resulting in more extreme phenomena. Thus, the effective rate was calculated to be $4.873 \cdot 10^{-4} s^{-1}$, representing a 4.8% increase compared to the one transducer simulation.

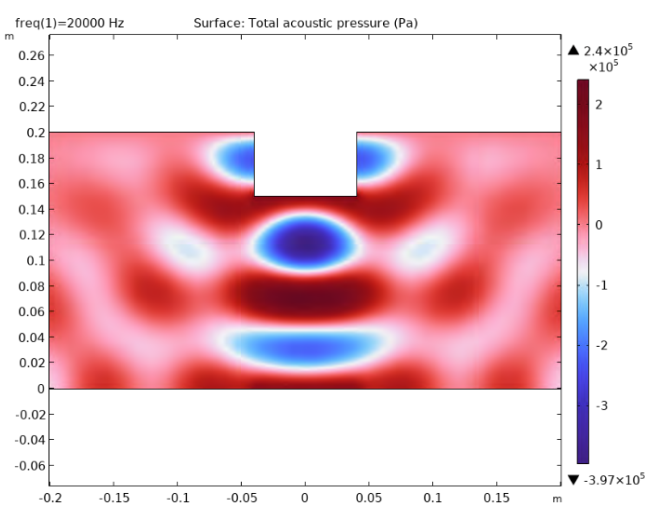

Figure 8. Pressure distribution for two transducers

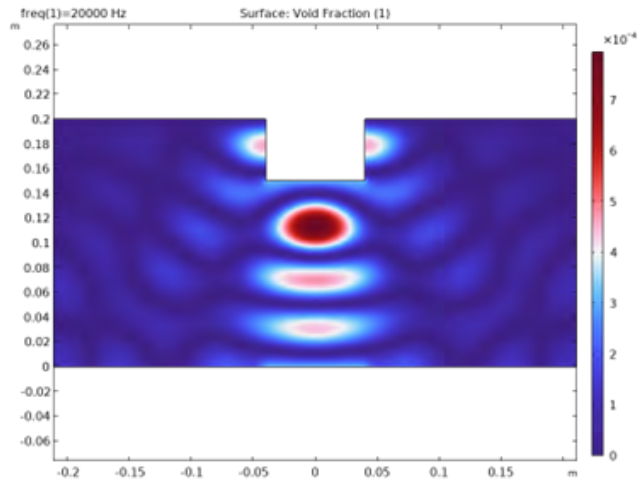

Figure 9. Void fraction for two transducers.

Optimization

To maximize the effective rate we investigated changing the position of the top transducer in all directions and the bottom transducer in x and y directions.

The optimum configuration was found to involve both transducers positioned at the center of the reactor, with the top transducer submerged as deeply as possible. The maximum effective rate was calculated to be $5.640 \cdot 10^{-4} s^{-1}$, which accounts for an increase of 15.7% of the rate calculated in the simple simulation.

In this case, the results were verified by the wave propagation as well. As shown in Figure 10, the pressure distribution changed drastically, when the transducer was well placed. On the central axis, the combination of the two transducers resulted in higher pressure amplitude. Due to more efficient reflection on the walls, the phenomenon was spreading to the rest of the bulk liquid, since pressure variations are observed in all the reactor.

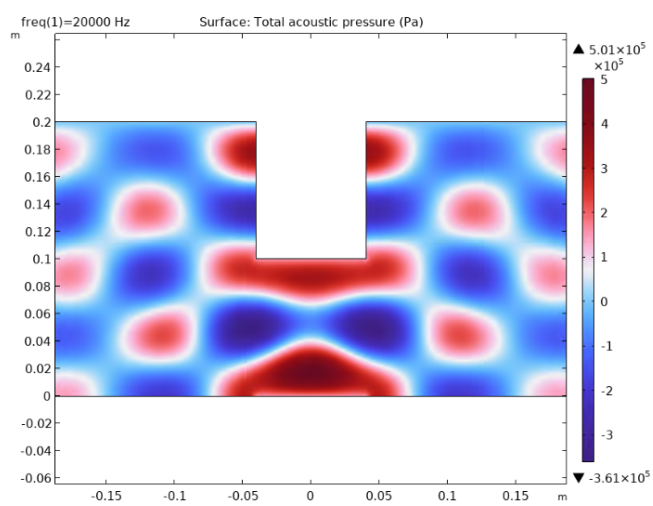

Figure 10. Pressure distribution for optimal position for the transducers.

The Void fraction mapping presented in Figure 11 further clarifies the result. Not only the areas with high void fraction were increased between the transducers, but the value of void fraction within these areas was increased as well. The mean void fraction was calculated to be $1.475 \cdot 10^{-4}$, an increase of 206%.

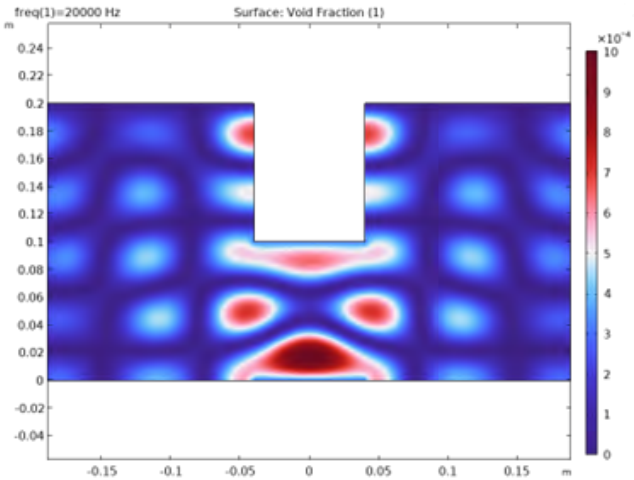

Figure 11. Void fraction for optimally positioned transducers.

It should be noted that the void fraction depends on the total volume of the liquid. Submerging the transducer more into the liquid, lowers the liquid volume, resulting in an increase of the void fraction. This plays an important role in the increase and should be expected. Furthermore, the fact that both transducers were optimally placed along the central axis rather than near the walls indicates that, for the present geometry characteristics, the most effective wave propagation occurs at the center, where the reflections constructively reinforce the acoustic field. In this arrangement, the pressure pulsations were symmetrically distributed throughout the reactor volume, ensuring that cavitation activity was present across most of the liquid domain.

Side transducers

Considering the optimization simulations performed up to this stage, the simulation of the side transducer was omitted, since the behavior of the effective rate was already established. The subsequent step was to determine the optimal placement along the z direction.

The optimal position was found to be in about ¾ the height of the reactor. The mean void fraction and the effective rate of the reaction were calculated to be $8.828 \cdot 10^{-5}$ and $4.746 \cdot 10^{-4}\ s^{-1}$ respectively. This was a slight improvement over only one transducer present.

The intuitive position for the transducer would probably be at half of the height of the reactor for symmetry or as low as possible to enable strong reflection from the walls. However, the optimum position was not even close to these, proving the complexity of wave propagation inside the reactor.

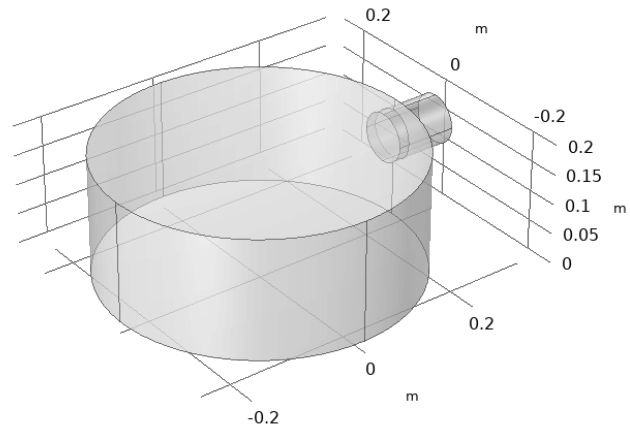

Figure 12. Optimal position for one side transducer.

Since the optimal height for one transducer had been calculated, the next optimization was focused on a second side transducer and its relative position to the first one. The optimal position is shown in Figure 13.

The mean void fraction was found to be $1.216 \cdot 10^{-4}$, while the rate of the reaction was $5.440 \cdot 10^{-4} s^{-1}$, which accounted for a 14.6% increase relative to just one transducer.

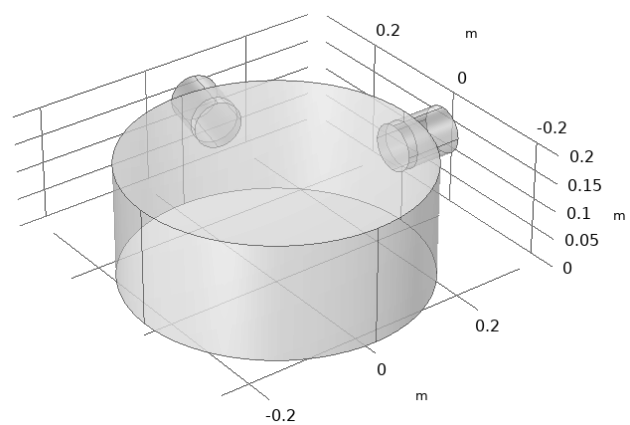

Figure 13. Optimal position for two side transducers.

It was expected that the optimum placement for the two side transducers would also be at the opposite sides of the reactor (like the vertical positioning problem solution). However, the optimization results indicates that this is not the case. This outcome can be attributed either to the horizontal orientation of the transducers that change drastically the wave propagation, or the fact that in the opposite-side configuration their separation distance is larger than in the vertical arrangement, increasing the wave damping throughout the liquid.

CONCLUSION

The present study examined the enhancement of batch reactor performance through the optimal placement of one or more ultrasonic transducers to increase chemical reaction rates. This objective was addressed through an optimization formulation incorporating multiscale modeling, integrating cavitation theory, wave propagation, heat diffusion, and reaction kinetics into a batch reactor model equipped with one or two transducers. Particular emphasis was placed on capturing the effect of transducer placement on reactor performance.

At the single-bubble scale, simulations showed that a bubble exposed to a single ultrasonic wave reached a collapse temperature of ~4800 K, whereas the interaction of two identical waves could increase the collapse temperature to 15, 000 K. This increase in temperature produced a substantial kinetic enhancement, with the reaction rate within the cavitation-affected volume increasing by 600% in the single-wave case and by 1700% in the two-wave case relative to nominal conditions.

At the reactor scale, a configuration with a single transducer submerged from the open top was used to establish baseline conditions. Cavitation activity, quantified by the void fraction, was concentrated beneath the transducer and provided reference values for both the mean void fraction and the effective reaction rate. The introduction of a second transducer increased the effective reaction rate by 5%, despite a slight decrease in the mean void fraction. This result indicates that reactor performance is governed not only by macroscopic wave propagation, but also by microscopic bubble dynamics.

The optimal placement of the two transducers was found to be along the central axis on opposite sides of the reactor, with the upper transducer submerged as deeply as possible. Under these conditions, the effective reaction rate was maximized, reaching a value 16% higher than that of the baseline configuration. Placing a single transducer on the reactor sidewall, had no significant performance enhancement. However, the addition of a second sidewall transducer at an optimized location increased the effective reaction rate by approximately 15% relative to the single-side configuration.

These findings demonstrate that simulation-based optimization can significantly improve the efficiency of sonochemical reactors by guiding the spatial arrangement of transducers. The proposed study highlights the potential of computational approaches for the rational design and scale-up of high-performance ultrasonic systems for chemical engineering applications.

AUTHOR IDENTIFIERS

Author ORCIDs:
Nikolaos Vittas:
Antonios Armaou: 000-0002-8592-7934

REFERENCES

1. Shah YT, Pandit AB, Moholkar VS. Cavitation reactors. The Plenum Chemical Engineering Series :193-245 (1999). https://doi.org/10.1007/978-1-4615-4787-7_6
2. Suslick KS. The chemical effects of ultrasound. Sci Am 260:80-86 (1989). https://doi.org/10.1038/scientificamerican0289-80
3. Kentish S, Ashokkumar M. Ultrasonic membrane processing. Food Engineering Series :583-598 (2010). https://doi.org/10.1007/978-1-4419-7472-3_23
4. Keller JB, Miksis M. Bubble oscillations of large amplitude. The Journal of the Acoustical Society of America 68:628-633 (1980). https://doi.org/10.1121/1.384720
5. A. Prosperetti and A. Lezzi, "Bubble dynamic in a compressible liquid. Part 1. First-order theory, " Fluid Mech., vol. 168, pp. 457-478, 1986.
6. Neppiras EA. Acoustic cavitation. Physics Reports 61:159-251 (1980). https://doi.org/10.1016/0370-1573(80)90115-5
7. Merouani S, Hamdaoui O, Rezgui Y, Guemini M. Theoretical estimation of the temperature and pressure within collapsing acoustical bubbles. Ultrasonics Sonochemistry 21:53-59 (2014). https://doi.org/10.1016/j.ultsonch.2013.05.008
8. Tatake PA, Pandit AB. Modelling and experimental investigation into cavity dynamics and cavitational yield: influence of dual frequency ultrasound sources. Chemical Engineering Science 57:4987-4995 (2002). https://doi.org/10.1016/s0009-2509(02)00271-3
9. Prabhu AV, Gogate PR, Pandit AB. Optimization of multiple-frequency sonochemical reactors. Chemical Engineering Science 59:4991-4998 (2004). https://doi.org/10.1016/j.ces.2004.09.033
10. Commander KW, Prosperetti A. Linear pressure waves in bubbly liquids: comparison between theory and experiments. The Journal of the Acoustical Society of America 85:732-746 (1989). https://doi.org/10.1121/1.397599
11. R. Caflish, M. Miksis, G. Papanicolaou and L. Ting, "Effective equations for wave propagation in bubbly liquids, " Fluid Mech., vol. 153, pp. 259-273, 1985.
12. Kanthale PM, Gogate PR, Pandit AB. Modeling aspects of dual frequency sonochemical reactors. Chemical Engineering Journal 127:71-79 (2007). https://doi.org/10.1016/j.cej.2006.09.023
13. Tiong TJ, Chu JK, Tan KW. Advancements in acoustic cavitation modelling: progress, challenges,

and future directions in sonochemical reactor design. Ultrasonics Sonochemistry 112:107163 (2025). https://doi.org/10.1016/j.ultsonch.2024.107163
14. Kavitha V, Palanivelu K. Destruction of cresols by fenton oxidation process. Water Research 39:3062-3072 (2005). https://doi.org/10.1016/j.watres.2005.05.011
15. Okouchi S, Nojima O, Arai T. Cavitation-induced degradation of phenol by ultrasound. Water Science and Technology 26:2053-2056 (1992). https://doi.org/10.2166/wst.1992.0659

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

MCSGP dynamic simulation for peptides separation using Aspen Chromatography

Ivan Chóez-Guaranda[a], Emmanuel Appiah-Danquah[a], Bogdan Dorneanu[a] and Harvey Arellano-García[a]

[a] Brandenburg University of Technology (BTU) Cottbus-Senftenberg, Department of Process and Plant Technology, Cottbus, Brandenburg, Germany
* Corresponding Author: arellano@b-tu.de.

ABSTRACT

The purification of therapeutic peptides represents a major bottleneck in biopharmaceutical downstream processing due to the structural similarity between target products and closely related impurities. In this study, a shortcut dynamic simulation model of a two-column Multi-Column Countercurrent Solvent Gradient Purification (MCSGP) process is implemented in Aspen Chromatography for peptide separation. Each column is described using a one-dimensional axial dispersion model coupled with mass transfer kinetics and a modulated Langmuir adsorption equilibrium, while time-dependent boundary conditions are applied to represent solvent gradient elution. The model explicitly incorporates internal recycle streams between columns using the cycle organizer approach, capturing the defining operational features of MCSGP. This enables a unified representation of chromatographic transport phenomena, gradient operation, and discrete recycle logic within a single flowsheet-based framework. The novelty lies in treating MCSGP as a hybrid dynamic process rather than a purely chromatographic unit, providing a transferable modelling framework suitable for operability analysis, process comparison, and future optimization of continuous peptide purification systems. Dynamic simulation results for a model mixture of closely eluting peptides demonstrate stable cyclic operation and a substantial improvement in mass recovery compared with conventional batch chromatography under identical transport assumptions. In particular, systematic recycling of intermediate fractions leads to an order-of-magnitude increase in recovered target product.

Keywords: Downstream processing, Modelling, Peptides, Purification, Preparative chromatography

1. INTRODUCTION

Therapeutic peptides represent an increasingly important class of biopharmaceutical products, yet their purification remains one of the most challenging and cost-intensive stages of production. During solid-phase synthesis, numerous closely related impurities are formed, often differing from the target molecule by only a few amino acids [1, 2]. These impurities exhibit very similar physicochemical properties, leading to overlapping retention behaviour during chromatographic separation and making selective purification particularly demanding [3].

Preparative batch chromatography is widely used for peptide purification but typically requires multiple sequential steps and large solvent volumes to achieve the desired purity. This results in long processing times, limited product recovery, and high operational costs. As a consequence, intensified chromatographic processes have gained increasing attention to improve efficiency, sustainability, and scalability in downstream processing [4].

Multi-Column Countercurrent Solvent Gradient Purification (MCSGP) effectively combines the high separation selectivity of gradient elution with the process efficiency offered by countercurrent operation to enhance separation performance. By internally recycling overlapping fractions between columns, MCSGP simultaneously allows for higher product yield and purity levels, while significantly reducing solvent consumption compared to conventional batch chromatography. In the MCSGP process, the chromatographic columns operate cyclically with alternating roles. One column performs gradient elution, during which overlapping fractions containing weak

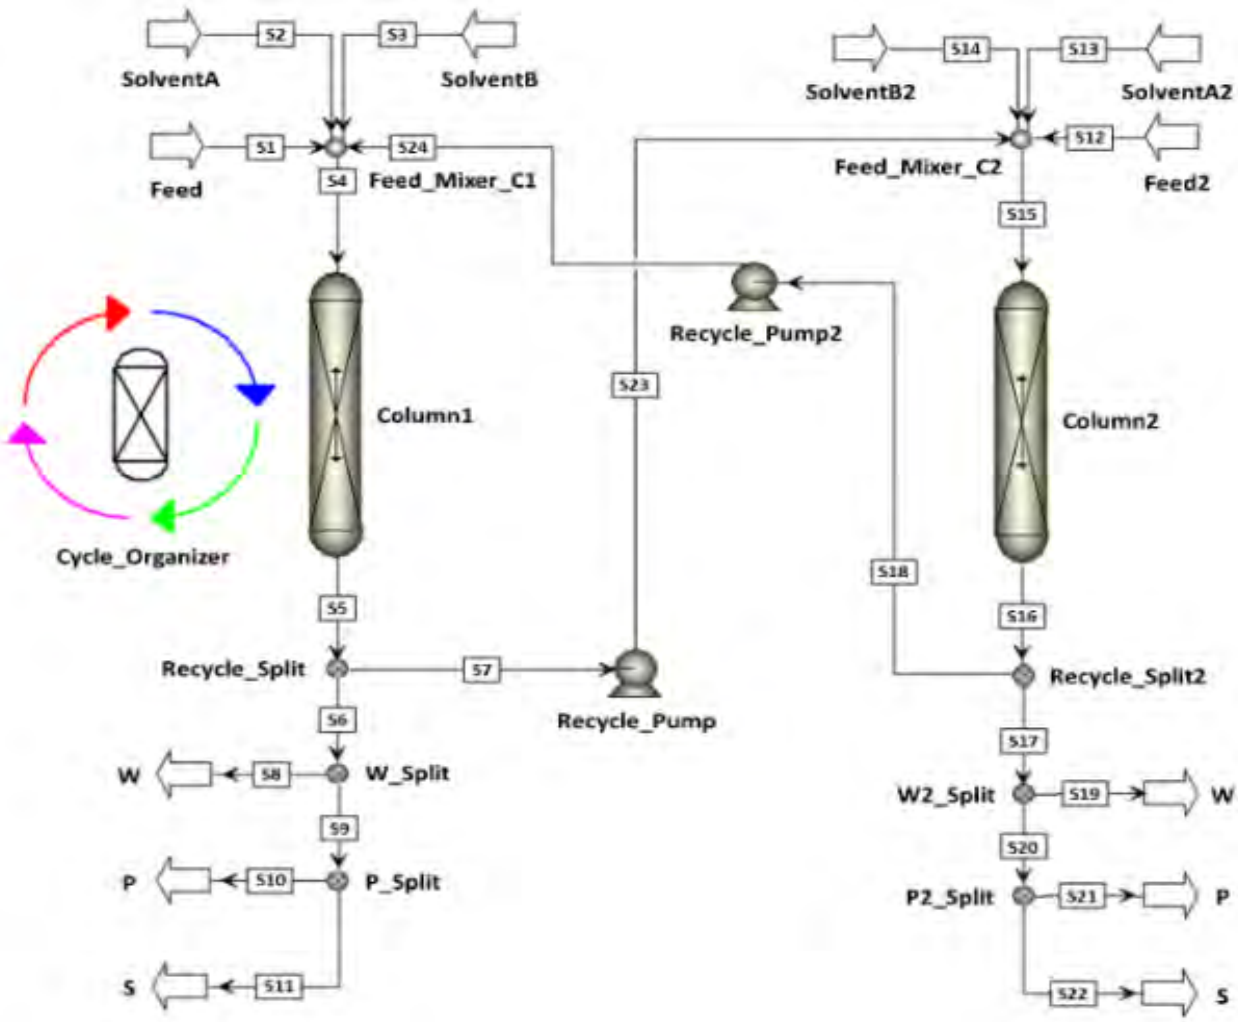

Figure 1. Two-column MCSGP process simulation diagram for the continuous separation of peptides developed in Aspen Chromatography.

adsorbed-product and product-strong adsorbed mixtures are internally recycled, while the pure target product fraction is collected. Meanwhile, the second column receives the recycled streams and fresh feed. After regeneration steps, the columns exchange their functions. Following two switches, the system returns to its initial configuration, completing one operating cycle [5–8]. While experimental and modelling studies have demonstrated these advantages, most dynamic MCSGP models for peptides have been developed using in-house numerical solvers, focusing primarily on separation physics rather than process operability [9–11]. From a process systems engineering perspective, MCSGP is inherently a hybrid dynamic process, characterized by continuous transport phenomena, time-dependent solvent gradients, discrete switching events, and internal cycle streams. Capturing these interactions in a transparent and reusable manner is essential for process comparison, operability analysis, and future integration with optimization control strategies.

In this context, the present work develops a dynamic, cycle-resolved simulation for peptide purification within Aspen Chromatography. The objective is not to introduce new separation mechanisms, but to demonstrate that the essential hybrid dynamic features of MCSGP can be represented transparently at the process level, enabling systematic analysis and reuse beyond custom numerical implementations.

2. METHODOLOGY

2.1. Process representation and simulation strategy

The MCSGP process is represented as a two-column system consisting of identical reversed-phase

chromatographic columns interconnected by internal recycle streams. The overall configuration, shown in Figure 1, reflects the cyclic nature of MCSGP operation, where columns alternate between loading, elution, and regeneration roles.

Each column is modelled using a one-dimensional axial dispersion formulation accounting for convection and axial mixing in the mobile phase. Intraparticle mass transfer resistance is described using the particle mass balance (PMB) kinetic model, capturing pore diffusion and external film resistance. Adsorption equilibrium is represented using a modulated Langmuir isotherm, allowing solvent composition effects during gradient elution to be explicitly incorporated. Dynamic simulations are conducted for a model peptide mixture characterized by closely eluting retention times, representative of realistic purification challenges. The solvent gradient is segmented to reproduce reversed-phase elution conditions commonly employed in peptide purification. Time-dependent boundary conditions are applied to represent both gradient elution and cyclic column switching.

A key methodological aspect of this work is the explicit implementation of column switching, gradient ramps, product withdrawal, and internal recycle streams using the cycle organizer approach in Aspen Chromatography. This avoids custom scripting and enables the MCSGP process to be treated as a hybrid dynamic system within a flowsheet environment, where both continuous dynamics and discrete operational logic are represented consistently.

2.2. Mathematical model formulation

The governing equations correspond to a conventional chromatographic transport model extended to include time-dependent boundary conditions and internal recycling. The axial mass balance accounts for accumulation, convection, dispersion, and mass transfer between mobile and stationary phases (Equation 1):

$$\frac{\partial C_i}{\partial t} + \frac{(1-\varepsilon_i)}{\varepsilon_i}\frac{\partial q_i}{\partial t} + \nu\frac{\partial C_i}{\partial z} = \frac{\nu.L}{Pe.\varepsilon_i}\frac{\partial^2 C_i}{\partial z^2} \quad (1)$$

where C_i is the concentration of the compound i in the mobile phase, t is the time, ε_i is the inter-particle voidage, q_i is the concentration of the compound i in the stationary phase, ν is the superficial velocity of the mobile phase, z is the axial distance through the length of the column, L is the column length, and Pe is the particle Péclet number.

Intraparticle mass transfer is described using the PMB kinetic model (Equation 2):

$$D_{pi}\left(\frac{\partial C_{pi}}{\partial r}\right) = k_i(C_i - C_{pi}) \quad (2)$$

where D_p is the pore/effective diffusion coefficient, C_p is the macropore concentration, r is the particle radius, k is the external film mass transfer coefficient, while local equilibrium between phases is defined by Equation 3:

$$\frac{dq_i}{dt} = \sum_{j=1}^{N}\left(\frac{dq_i}{dc_j}\frac{dc_j}{dt}\right) \quad (3)$$

Adsorption is modelled using a modulated Langmuir isotherm (Equation 4), which captures nonlinear adsorption behaviour and the influence of solvent strength during gradient elution. This formulation has been shown to be suitable for reversed-phase separation of peptides and related biomolecules.

$$q_i = \frac{IP_{1i}\, 10^{IP_{2i} - IP_{3i}\log_{10}(C_m) + IP_{4i}(C_m)}\, C_i}{1 + \sum\left(IP_{1i}\, 10^{IP_{2i} - IP_{3i}\log_{10}(C_m) + IP_{4i}(C_m)}\right) C_i} \quad (4)$$

where IP is the isotherm parameter, and C_m is the concentration of the mobile phase.

The numerical solution of the resulting system of differential–algebraic equations is performed in Aspen Chromatography v12 using the fourth-order orthogonal collocation finite element method (OCFE4) with one node per column, combined with an implicit Euler integration scheme suitable for stiff systems.

Table 1: Column specifications and mass transfer parameters used for the two-column MCSGP model.

Parameter	Value/Unit
Column length	10.00 cm
Column internal diameter	1.00 cm
Inter-particle voidage	0.4 m^3 void/m^3 bed
Intra-particle voidage	0.4 m^3 void/m^3 bead
Particle radius	60.00 µm
Péclet number for dispersion/ACN	400.00
Péclet number for dispersion / P	350.00
Péclet number for dispersion / S	400.00
Péclet number for dispersion / W	350.00
Diffusion coefficient / ACN	2.292 x10^{-3} cm^2/min
Diffusion coefficient / P	9.167x10^{-4} cm^2/min
Diffusion coefficient / S	1.375 x10^{-3} cm^2/min
Diffusion coefficient / W	9.167x10^{-4} cm^2/min
Isotherm parameter 1 / P	0.10
Isotherm parameter 1 / S	0.10
Isotherm parameter 1 / W	0.10
Isotherm parameter 2 / P	3.00
Isotherm parameter 2 / S	4.50
Isotherm parameter 2 / W	1.80
Isotherm parameter 3 / P	3.00
Isotherm parameter 3 / S	3.50
Isotherm parameter 3 / W	2.60
Isotherm parameter 4 / P	0.0032
Isotherm parameter 4 / S	0.0032
Isotherm parameter 4 / W	0.0032

2.3. Model simulation parameters

The column geometry and mass transfer parameters used in the simulations are summarized in Table 1. Identical parameters are assumed for both columns to isolate the effect of process configuration and operation. Feed

and mobile phase specifications are provided in Table 2, representing a mixture of target peptide (P), a weakly adsorbed impurity (W), and a strongly adsorbed impurity (S).

The solvent gradient strategy employed in the simulations is detailed in Table 3. A sequence of isocratic and linear gradient steps is defined for each column using the cycle organizer. Precisely, post feed loading (step 1) for the first elution of column 2 which receives the overlapped fractions from column 1, the elution of W (step 2), elution of P (step 4), and elution of S (step 6) are achieved during the different established linear ramps. In the meantime, the elution of the overlapping fractions containing the mixture W/P (step 3), and the mixture P/S (step 5) are also completed and internally recycled to column 1. Then, the post-elution time (step 7) allows the complete separation and collection of the strongly adsorbed compound, prior to column equilibration. After two column switches, the system returns to its initial configuration, completing one full MCSGP cycle [12].

Internal recycle streams corresponding to overlapping W/P and P/S fractions are treated as free variables, reflecting their dynamic nature during the process operation.

Table 2: Inlet/outlet stream specifications used for the two-column MCSGP model.

Stream	Value/Unit	Specification
Feed		
Flowrate	0.20 mL/min	Fixed
Concentration (ACN)	1.50 g/L	Fixed
Concentration (P)	0.05 g/L	Fixed
Concentration (S)	0.01 g/L	Fixed
Concentration (W)	0.015 g/L	Fixed
Pressure	20.00 bar	Fixed
Mobile Phase		
Solvent A (ACN)		
Flowrate	1.00 mL/min	Fixed
Concentration	1.50 g/L	Fixed
Pressure	20.00 bar	Fixed
Solvent B (ACN)		
Concentration	45.00 g/L	Fixed
Pressure	20.00 bar	Fixed
Internal recycles (W/P, P/S)		
Flowrate	5 x10^{-4} L/min	Free
Pressure	10.00 bar	Free
Products (P, S, W)		
Flowrate	5 x10^{-4} L/min	Free
Pressure	10.00 bar	Free

Table 3: Mobile phase gradient elution conditions employed for the two-column MCSGP model and dynamic simulation.

Step	Time (min)	Mobile phase Solvent A (mL/min)	Mobile phase Solvent B (mL/min)
1	1.00	1.00	0.00
2	7.74	1.00-0.94	0.00-0.06
3	1.00	0.94-0.93	0.06-0.07
4	6.86	0.93-0.50	0.07-0.40
5	2.20	0.50-0.18	0.40-0.79
6	1.20	0.18-0.00	0.79-1.00
7	2.50	0.00	1.00

3. RESULTS AND DISCUSSION

Dynamic simulations are performed for one complete MCSGP cycle with a total duration of approximately 90 minutes, which is sufficient to capture the transition from initial batch-like operation to fully cyclic steady behaviour. Figure 2 shows the outlet concentration profiles of the weakly adsorbed impurity (W), target peptide (P), and strongly adsorbed impurity (S) for both columns. The concentration profiles are reported in absolute concentration units, enabling direct comparison between batch and continuous operation.

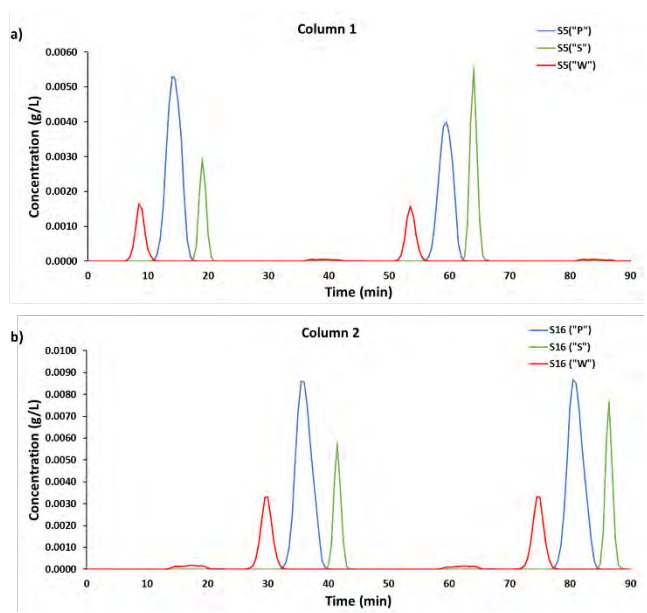

Figure 2. Concentration profiles of the MCSGP outlet streams (weakly adsorbed W, red line; target product P, blue line; strongly adsorbed S, green line) obtained after the dynamic simulation of the one cycle for: a) Column 1, and b) Column 2.

Distinct and recurrent elution peaks are observed

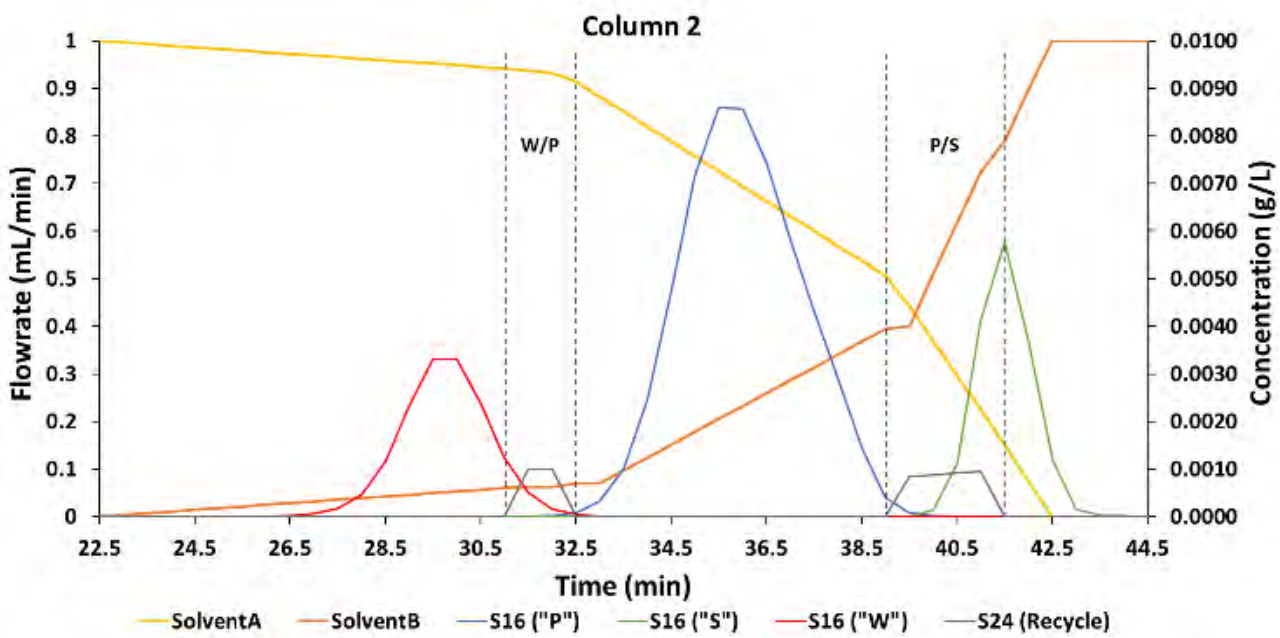

Figure 3. Mobile phase flowrates (solvent A, orange line; solvent B, yellow line), recycle flowrates (grey line), and concentration profiles of the outlet streams (weakly adsorbed W, red line; target product P, blue line; strongly adsorbed S, green line) obtained during the first elution of Column 2.

for all components, indicating stable cyclic operation after the initial transient. During the first elution of Column 1, the system operates effectively in batch mode, as no internal recycle streams are yet active. This is reflected in comparatively lower peak concentrations and a single-pass utilization of the injected feed. Once internal recycle streams are activated, subsequent elution exhibits a recurrent and stable pattern, indicating that MCSGP process has reached a cyclic operating regime. In particular, the concentration profiles obtained for Column 2 show noticeably higher peak intensities for the target peptide P, demonstrating the cumulative effect of recycling overlapping W/P and P/S fractions from Column 1.

Figure 2 further shows that peak shapes remain well defined over successive cycles, suggesting that axial dispersion and intraparticle mass transfer resistance are adequately captured by the adopted modelling framework. Limited peak broadening is observed, despite the repeated recycling of intermediate fractions, indicating that the chosen gradient-design and switching strategy preserve separation efficiency while enabling enhanced recovery.

The coupling between solvent gradient elution and component separation is examined in more detail in Figure 3, which reports the mobile phase composition alongside outlet profiles during the first elution of Column 2.

The weakly adsorbed impurity W elutes at low solvent strength during the early gradient segment, while the target peptide elutes at intermediate solvent composition. The strongly adsorbed impurity S is released only at higher solvent strengths during the final gradient steps. Quantitatively, the elution windows of each component correspond closely to the segment gradient ramps defined in Table 3, confirming the cycle organizer approach allows precise synchronization between gradient evolution and recycle timing.

Of particular relevance is the treatment of overlapping fractions. As shown in Figure 3, the elution intervals where W/P and P/S overlap are deliberately short but non-negligible, reflecting realistic separation challenges for closely eluting peptides. Instead of being discarded, these fractions are redirected through internal recycle streams and reintroduced into the opposite column. This process-level mechanism is central to the performance gains observed in MCSGP and cannot be achieved in conventional batch chromatography without additional purification steps.

A quantitative comparison between conventional batch chromatography and two-column MCSGP is presented in Figure 4. Under identical transport and adsorption assumptions, the recovered mass of the target peptide increases from approximately 0.067 g in batch operation to about 0.674 g in MCSGP, corresponding to an order of magnitude improvement in recovery. Similarly, recovery of the strongly adsorbed impurity increases from roughly 0.017 g to 0.598 g, confirming the effectiveness of internal recycling for late-eluting components. The weakly adsorbed fraction shows a more modest increase, from about 0.016 g in batch mode to 0.078 g in MCSGP, consistent with its early elution and

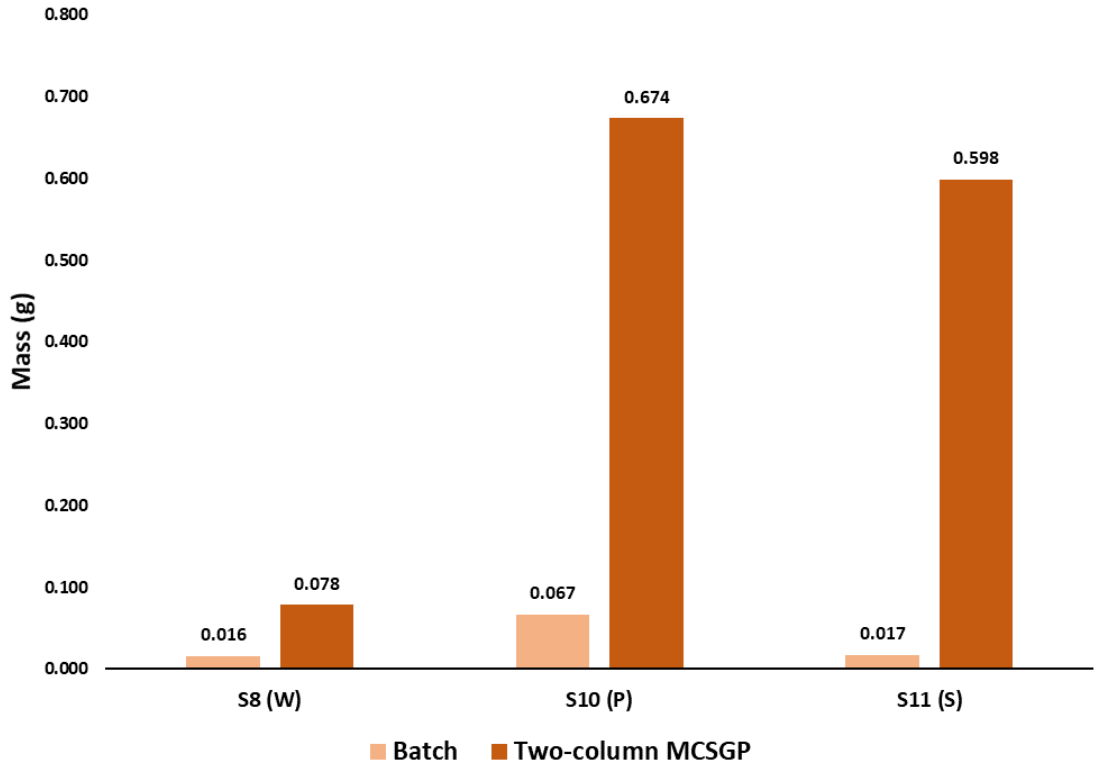

Figure 4. Comparative analysis of mass recovery after dynamic peptide separation using conventional batch chromatography and a two-column MCSGP process.

limited overlap.

These quantitative results demonstrate that the performance gains achieved by MCSGP are not due to changes in column design or adsorption parameters, but arise directly from the process configuration and operating strategy. The coordinated interaction between solvent gradient design, recycle timing, and column switching enables the systematic reuse of overlapping fractions that would otherwise be discarded in batch operation.

The observed trends are consistent with previously reported MCSGP modelling studies for peptides and nucleotides [9, 10], particularly with respect to enhanced recovery and selectivity. However, the present work differs in its explicit treatment of MCSGP as a hybrid dynamic process within a commercial simulation environment. By embedding chromatographic transport, gradient elution, and recycle logic in a unified flowsheet-based framework, the model facilitates process-level interpretation and provides a practical basis for future operability, optimization, and control-oriented studies.

4. CONCLUSIONS

A dynamic, cycle-resolved simulation model of a two-column MCSGP process for peptide purification has been developed and implemented in Aspen Chromatography. The preliminary results confirm that MCSGP operation substantially improves product recovery compared with conventional batch chromatography through the coordinated application of solvent gradient elution, column switching, and internal recycling of overlapping fractions. Under identical transport and adsorption assumptions, the simulated MCSGP process achieves an order of magnitude increase in recovered target peptide, clearly demonstrating the process-level benefits of continuous operations.

Beyond confirming known separation advantages, the key contribution of this work lies in framing MCSGP as a hybrid dynamic process and demonstrating its essential features can be captured in a transparent and reusable manner within a commercial flowsheet-based simulation environment rather than relying on custom numerical implementations. The proposed model serves as a transferable template for systematic process comparison, operability analysis, and early-stage decision-support in biopharmaceutical downstream processing.

Future work will focus on strengthening the predictive capability and practical relevance of the proposed

framework. This includes validation of the model against experimental peptide purification data, as well as parameter estimation and sensitivity analysis to better quantify uncertainty in mass transfer and adsorption parameters. Extending the model to multi-column configurations and incorporating additional performance indicators, such as solvent consumption, productivity, and robustness to operating disturbances, will further enhance its applicability. Finally, integration of the dynamic MCSGP model with optimization and control strategies represents a promising direction toward digital twin-based design and operation of continuous chromatographic processes.

ACKNOWLEDGEMENTS

The authors sincerely express their gratitude to the Open Access Funding provided by Brandenburg University of Technology (BTU) Cottbus-Senftenberg.

AUTHOR IDENTIFIERS

Chóez-Guaranda I: 0000-0003-2603-9491
Appiah-Danquah E: 0000-0002-3826-2634
Dorneanu B: 0000-0003-3553-6625
Arellano-García H: 0000-0002-8297-0232

REFERENCES

1. De Luca C, Lievore G, Bozza D, Buratti A, Cavazzini A, Ricci A, Macis M, Cabri W, Felletti S, Catani M. Downstream processing of therapeutic peptides by means of preparative liquid chromatography. Molecules 26:4688 (2021). https://doi.org/10.3390/molecules26154688
2. Weldon R, Müller-Späth T. Enrichment and purification of peptide impurities using twin-column continuous chromatography. Journal of Chromatography A 1667:462894 (2022). https://doi.org/10.1016/j.chroma.2022.462894
3. Catani, M., De Luca, C.: Chapter three - application of Multicolumn Countercurrent Solvent Gradient Purification to the polishing of therapeutic proteins. In: Moscatelli, D. and Sponchioni, C.E. (eds.) Process Intensification in the Manufacturing of Biotherapeutics. pp. 69–100. Academic Press (2022)
4. Luca C, Felletti S, Lievore G, Buratti A, Vogg S, Morbidelli M, Cavazzini A, Catani M, Macis M, Ricci A, Cabri W. From batch to continuous chromatographic purification of a therapeutic peptide through multicolumn countercurrent solvent gradient purification. Journal of Chromatography A 1625:461304 (2020). https://doi.org/10.1016/j.chroma.2020.461304
5. De Luca C, Felletti S, Lievore G, Chenet T, Morbidelli M, Sponchioni M, Cavazzini A, Catani M. Modern trends in downstream processing of biotherapeutics through continuous chromatography: the potential of multicolumn countercurrent solvent gradient purification. TrAC Trends in Analytical Chemistry 132:116051 (2020). https://doi.org/10.1016/j.trac.2020.116051
6. Krättli M, Müller?Späth T, Morbidelli M. Multifraction separation in countercurrent chromatography (MCSGP). Biotech & Bioengineering 110:2436-2444 (2013). https://doi.org/10.1002/bit.24901
7. Aumann L, Morbidelli M. A continuous multicolumn countercurrent solvent gradient purification (MCSGP) process. Biotech & Bioengineering 98:1043-1055 (2007). https://doi.org/10.1002/bit.21527
8. Fioretti I, Müller?Späth T, Weldon R, Vogg S, Morbidelli M, Sponchioni M. Continuous countercurrent chromatographic twin?column purification of oligonucleotides: the role of the displacement effect. Biotech & Bioengineering 119:1861-1872 (2022). https://doi.org/10.1002/bit.28093
9. Vetter-Joss L, Sellberg A, Hansen EB. Model-supported design of multicolumn countercurrent solvent gradient purification processes for rapid implementation. Ind. Eng. Chem. Res. 63:17612-17624 (2024). https://doi.org/10.1021/acs.iecr.4c02092
10. Taguado Menza S, Prestia R, Fioretti I, Sponchioni M. Model-based optimization strategy for intensification in the chromatographic purification of oligonucleotides. Journal of Chromatography A 1736:465321 (2024). https://doi.org/10.1016/j.chroma.2024.465321
11. Grossmann C, Ströhlein G, Morari M, Morbidelli M. Optimizing model predictive control of the chromatographic multi-column solvent gradient purification (MCSGP) process. Journal of Process Control 20:618-629 (2010). https://doi.org/10.1016/j.jprocont.2010.02.013
12. Baur, D.: Design, modeling and optimization of multi-column chromatographic processes, https://doi.org/https://doi.org/10.3929/ethz-a-010881404, (2017)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Multisectorial Energy Integration of Low-Temperature Brewery Process, Manufacturing Industry and District Heating Network

Pullah Bhatnagar[a]*, Daniel Florez-Orrego[a], Oktay Boztas[a], Meire Ribeiro Domingos[a], Manuele Margni[b], Francois Marechal[a]

[a] EPFL-Valais, École polytechnic federal de Lausanne, Sion, Switzerland
[b] HES-SO Valais-Wallis, Sion, Switzerland
* Corresponding Author: pullah.bhatnagar@epfl.ch.

ABSTRACT

Low-temperature industrial processes release substantial amounts of waste heat, representing a largely untapped renewable energy resource. This study focuses on the brewery sector, encompassing both beer and whiskey production, along with its integration with manufacturing and city. The brewery industry generates approximately 0.061 kWh of waste heat per liter of beer, while whiskey production releases around 2.2 kWh per liter, with most of this waste heat available at temperatures close to 95 °C. Such low-grade heat is well suited to meet heating demands in manufacturing industries and urban district heating networks, where temperature requirements typically remain below 80 °C. Multiple technological options for meeting process heat requirements and recovering waste heat are evaluated using the OSMOSE energy integration framework. The study assesses the technical performance and economic viability of these options under varying assumptions for electricity prices, natural gas prices, and carbon pricing. Depending on market conditions, overall external utilities consumption reductions ranging from 22% to 63% are achieved. Price of utilities also has an effect on reduction of emission, as it is ranging from 25% to 90%.The results highlight the significant potential for cross-sectoral heat integration between breweries, manufacturing industries, and urban energy systems. A comparison is drawn between a case in which 5 th generation system (case 1) is activated and a case in which 5 th generation is not active (case 2). It was observed in most of the market conditions, case 1 is energetically, environmentally and economically more sustainable compared to case 2. In one of the market condition it was possible for both the cases to be energetically a carbon negative solution as it was economically feasible to capture biogenic CO_2. The work helps in validating that strategic waste heat recovery and utilization can substantially enhance energy efficiency, support industrial decarbonization pathways, and improve the resilience of urban energy infrastructures.

Keywords: Cross-sectoral Integration, Waste heat utilization, Brewery, Decarbonization, Renewable energy

INTRODUCTION

In 2023 the energy consumption for all the industrial activities in Europe (majorly EU and UK) was about 2750 TWh/year [1-2]. Meanwhile, the total energy consumption for residential sector is 2650 TWh/year [1-2]. The current work focuses on two industrial sectors: 1) Brewery and Whiskey production and 2) Manufacturing industry and their symbiosis with a typical city. Brewery industry act as source of low-grade heat, which is well suited in meeting the heat demand for manufacturing industry and also, urban district heating networks, where temperature requirements remain typically below 80 °C.

METHODOLOGY

The section first presents, the brewery and manufacturing process description, along with the considered

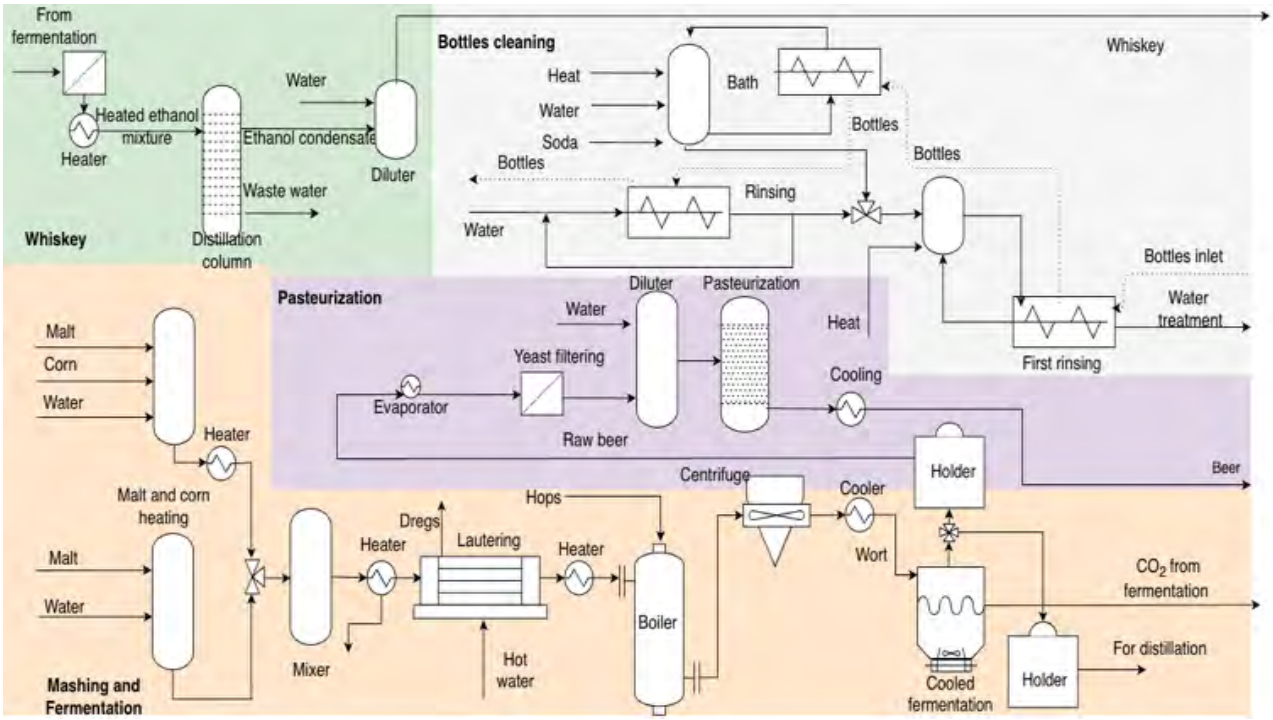

Figure 1. Process flowsheet of the beer and whiskey production with cleaning in place of bottles.

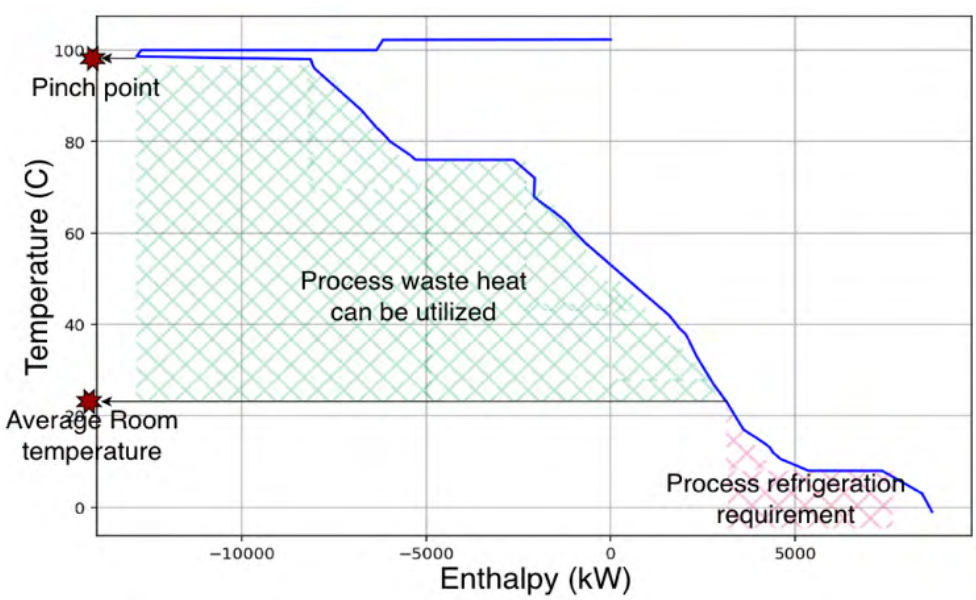

Figure 2. Integrated composite curve of the beer and whiskey production. Green shaded region identifies the process waste heat. Red shaded region identifies refrigeration need.

city demands. Next, the cluster analysis is discussed, followed by explaining in brief the various technological options accounted to meet the industrial utility demands. Finally, discussing about the different approaches for waste heat recovery and its subsequent utilization.

Brewery Industry

The process description is shown in Figure 1. Primary stages in the production of beer and whiskey include malting and fermentation along with pasteurization for beer [3], [4], and distillation for whiskey [5]. The final products are then bottled. In the current analysis the capacity from beer production Is around 49 m^3/h, while the capacity of whiskey production is around 2.9 m^3/h.

Energy requirement for the processes was

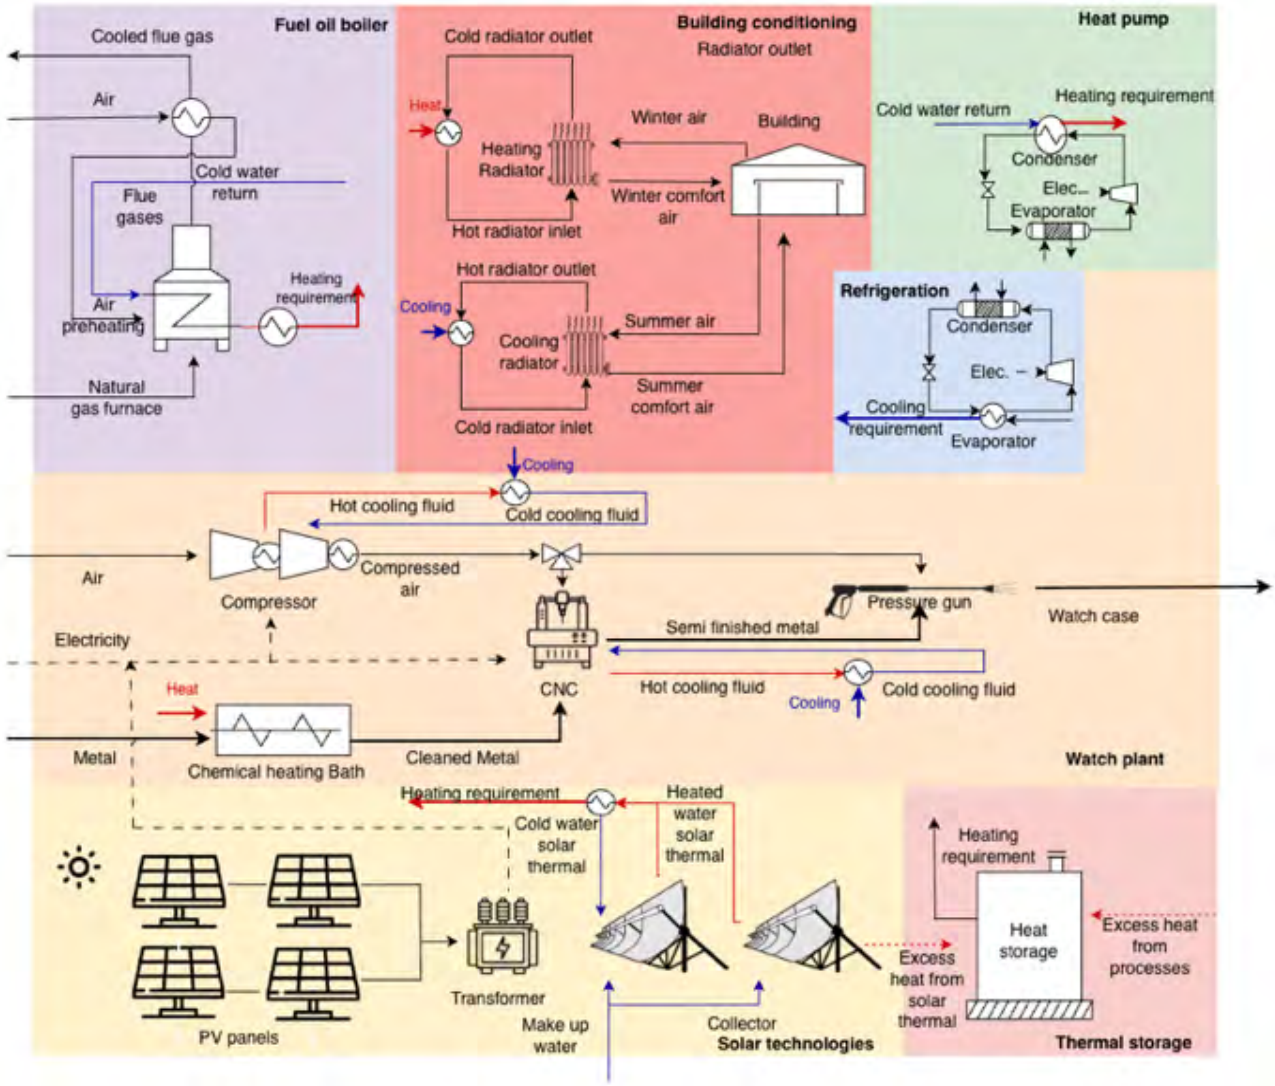

Figure 3. Superstructure of manufacturing industry.

simulated using the opensource modelling tool DWSIM. Process integration was performed using OSMOSE [8]. The waste heat available as well as its amounts are shown in Figure 2. Energy requirement remains constant for the whole year as in this analysis the industry is considered to be a largescale industry.

Manufacturing Industry

The major part of the energy requirement for the manufacturing industry is electricity. Heat is generally required for heating of offices and incoming metal. ([7]). Figure 3 helps in identifying what are the different process involved in machining industry. Energy requirement is varying throughout the year, the variation in energy is depicted in Figure 4.

City

Energy requirement in city generally corresponds to the energy requirement for residential purpose. The majority of energy consumption is in the form of space heating, water heating, lighting and electrical appliances and cooking.

Energy requirement varies depending upon the season. Variation in energy demand for a city is shown in figure 5.

Cluster analysis

Accurately capturing temporal variations in industrial demand is essential for system-level energy analysis. However, representing demand at an hourly resolution over an entire year (8, 760 hours) leads to a substantial computational burden. To address this challenge, a K-means clustering algorithm is employed to aggregate hourly demand profiles into 15 representative operating points distributed across the year. A silhouette score of 0.73 indicates good cluster separation, while the elbow

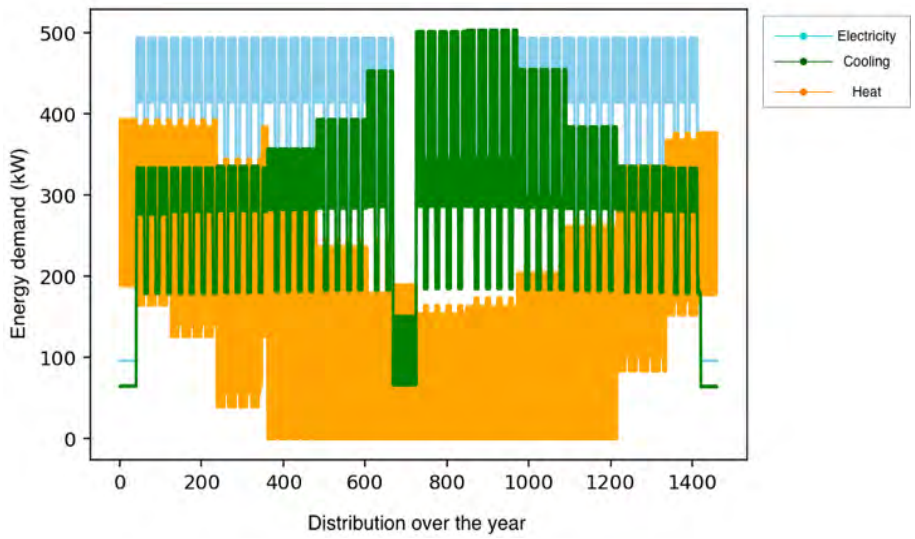

Figure 4. Variation in energy demand across the year in manufacturing unit. Each time step represents a typical 6 hours of working.

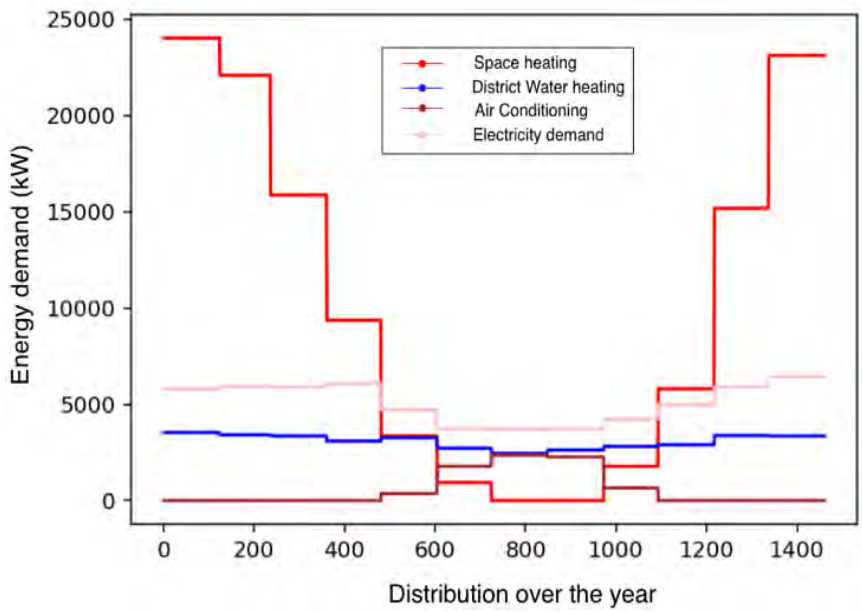

Figure 5. Variation in energy demand across the year in a city. In the figure each time step represents a typical 6 hours of working.

method shows that the within-cluster sum of squares (WCSS) approaches a minimum beyond 15 clusters. The resulting representative points preserve both seasonal and operational variability while significantly reducing computational complexity.

Utilities superstructure

Figure 6 presents the utility supply and waste valorization options available for the brewery industry. A share of the heating and cooling demand can be supplied by heat pumps operating at different temperature levels. The brewery generates by-products such as dregs and releases waste streams including wastewater and biogenic CO_2. These streams offer multiple valorization pathways. Dregs can be converted into synthetic natural gas (SNG), which can be used either in gas engines for electricity generation or in furnaces for heat production. Wastewater can be treated to produce sludge, which is subsequently converted into SNG via anaerobic digestion. Biogenic CO_2 released during fermentation can be converted into SNG through methanation using hydrogen produced by electrolysis [6].

For the manufacturing (machining) sector, part of

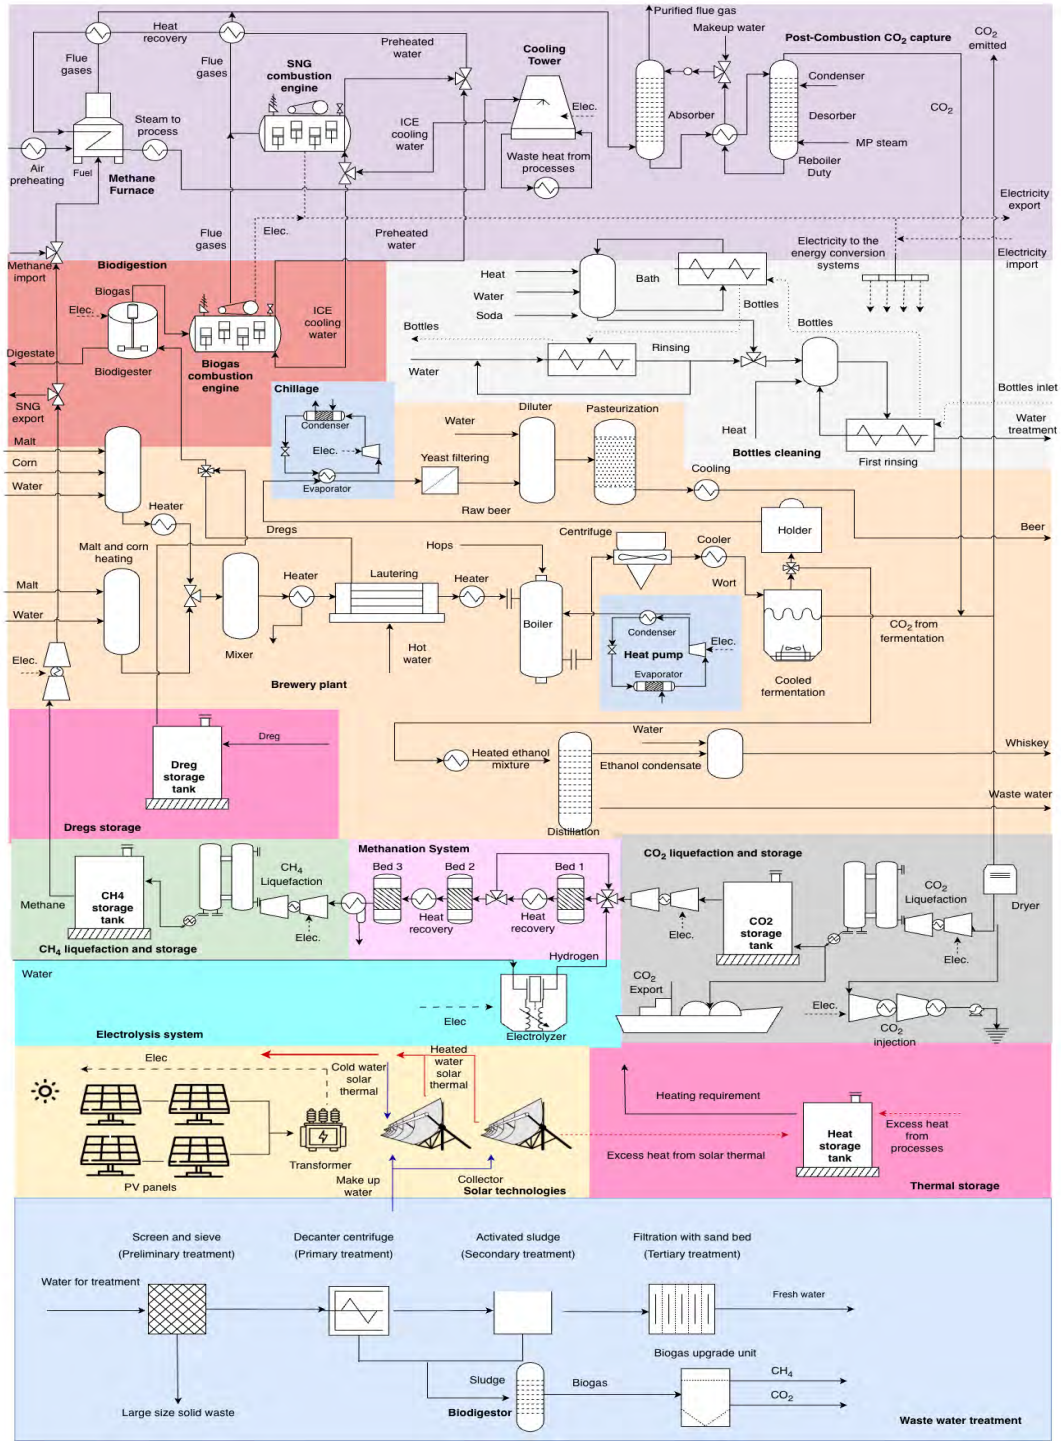

Figure 6: Utilities and waste valorization

the heating and cooling demand can similarly be met using heat pumps, while solar thermal energy is considered as an additional option for supplying low-temperature heat. The corresponding utility options are shown in Figure 3 [7].

For the urban sector, heating is supplied through a district heating system in which heat is generated centrally and distributed to residential buildings. Centralized heat supply can be provided by either large-scale heat pumps or NG furnaces, while decentralized NG furnaces can alternatively be replaced by decentralized heat pumps [4].

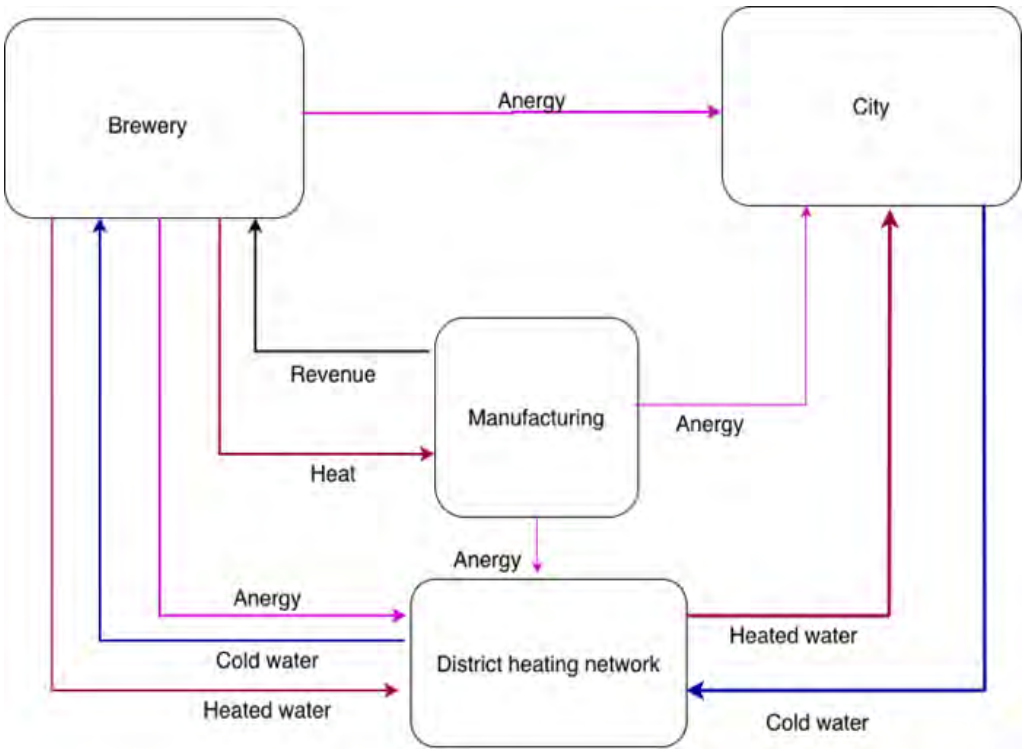

Figure 7. 5th generation heat recovery and utiliztaion

Waste heat recovery.

The brewery industry generates approximately 0.061 kWh of waste heat per liter of beer produced, while whiskey production releases around 2.2 kWh per liter, with most of the waste heat available at temperatures close to 95 °C. For the studied system, this corresponds to a total waste heat potential of approximately 12.76 MW.

Cold outlet from the houses can be heated up to required temperature with the process waste heat released in brewery. In this case some section of district heating network is provided with the help of process waste heat. In brewery and manufacturing industry as depicted in figure. In brewery process as identified from figure 2, there is need of refrigeration system to provide cooling. These refrigeration system act as a source of anergy (low grade heat which cannot be directly used in direct heating). The anergy act as a source of heat for the evaporator section of decentralized or centralized heat pump and provide heating to the house directly or to the district heating network. All these possible options are depicted in figure 7.

Sensitivity analysis.

Given the multiple utility and waste valorization options shown in Figures 3, 6 and 7, the optimal system configuration is determined using the optimization framework OSMOSE [8]. Technology selection is driven by the chosen objective function. In this study, total cost (TOTEX), defined as the sum of annualized investment and operating costs, is used as the optimization criteria.

Key parameters influencing technology selection include electricity price, natural gas price, and carbon tax level. A sensitivity analysis is therefore performed by varying these parameters to evaluate their impact on system configuration, energy self-reliance, and emissions.

RESULT

The analysis is carried out for two distinct system configurations. For the System Configauration 1 (SC 1), all three sectors are fully integrated and allowed to exchange mass and energy flows. In System Configuaration 2 (SC 2), brewery, manufacturing sector, and the city operate independently, with no exchange of heat, anergy, or mass flows among them.

For both scenarios, a sensitivity analysis is performed and the results are evaluated based on two criteria: (i) improvements in energy efficiency, (ii) reductions in carbon emissions. In all the comparison, results are evaluated against the base system in which the heating requirement is met with the help of natural gas furnace. The cooling requirement is met with the help of cooling tower and refrigeration system. In all the further discussion this scenario will be considered as base system.

Self-reliance of system

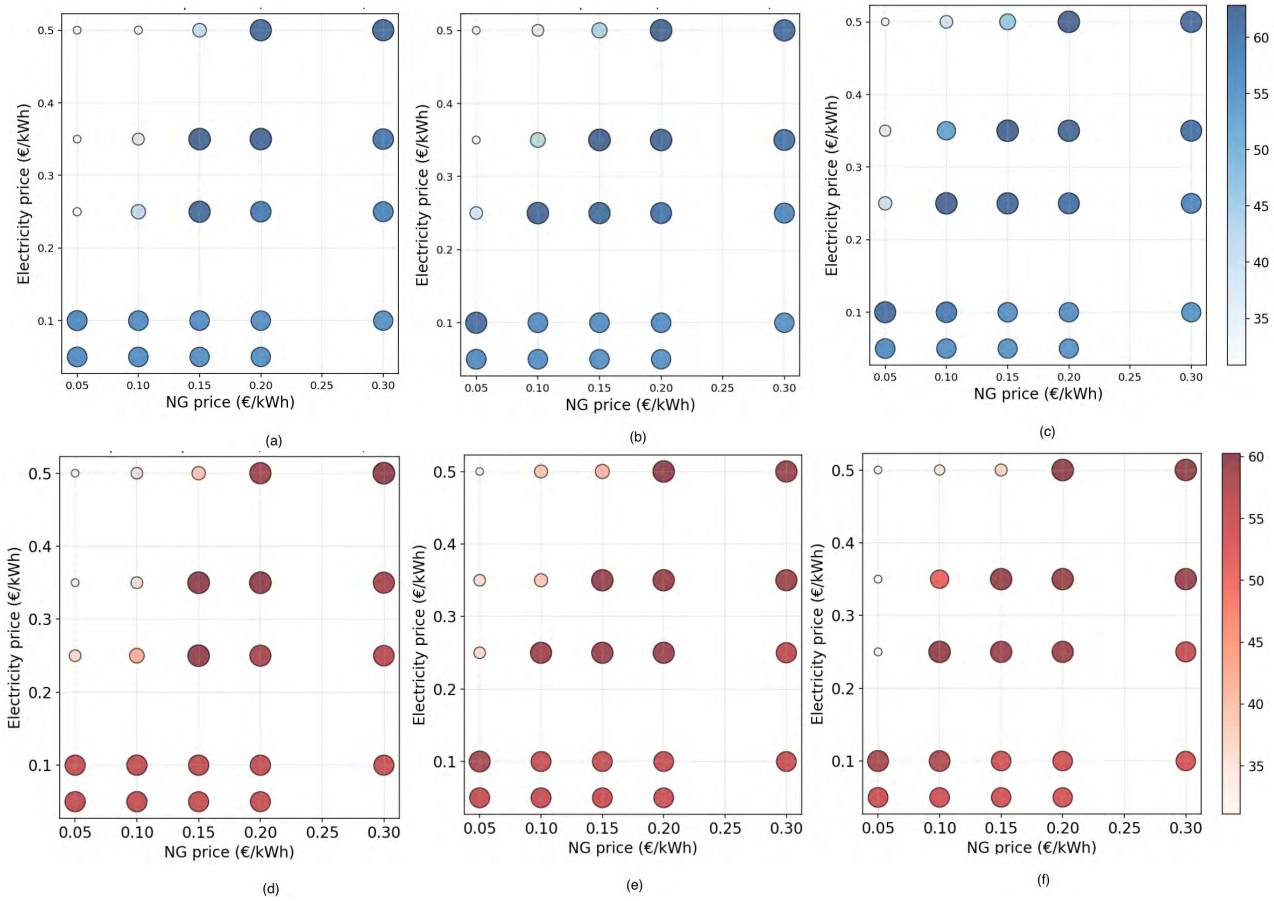

Figure 8: Percentage increase in system self-reliance compared to the base system cofiguration, across different electricity prices, natural gas prices, and carbon taxes in both system conifgurations Panels (a), (b), and (c) correspond to the SC1 while panels (d), (e), and (f) correspond to the SC 2. Graph (a) and (d), corresponds to a carbon tax of 0 €/ton, Graph (b) and (e), value of carbon tax is 80 €/ton; and in Graphs (c) and (f), it is 200 €/ton. Darker and larger points indicate solutions that are more energy independent.

In the base system, total natural gas consumption is approximately 216 TWh/year, while total electricity consumption is around 63 TWh/year. In SC 1 and SC 2, consumption of natural gas and electricity varies as a function of electricity price, natural gas price, and carbon tax. Here, energy self-reliance is defined as the reduction in externally purchased electricity and natural gas relative to the base case. Figure 8 illustrates improvements in energy self-reliance by summing purchased electricity and natural gas in each sensitivity case and comparing them to the base scenario.

Across the sensitivity analysis, energy self-reliance improvement has a very wide range of 61–23% for the SC 1, and also 60–25% for SC 2. The reason for the wide range is due to shift of energy vector. As the price of electricity increases, the optimizer chooses to shift from eletrical based utilities to natural gas based utilities. This shift decreases the self sufficieny of the system. The reason is because, heat pumps (electrically driven utility) operate across the system's pinch point and can internally redistribute energy to meet both heating and cooling demands, thereby reducing external energy purchases.

Comparing (Figure 8 (a) and 8(d)) with (Figures 8 (b) and 8(e)) and, (Figure 8(c) and 8 (f)) respectively, increasing ca rbon taxes chooses more energy-independent utilities option. The reason is, electrically driven utilities generally exhibit lower carbon emissions than natural gas-based alternatives. As carbon tax increases, the economic penalty associated with CO_2 emissions rises, making low-carbon, electricity-based solutions more attractive despite higher electricity prices. This highlights the role of carbon pricing in incentivizing integrated, low-carbon energy systems, even in contexts with relatively high electricity prices.

Figure 9. compares the energy self-reliance of the SC 1 from SC 2, for the value where carbon tax of zero. Overall, the SC 1 exhibits higher self-sufficiency compared to the SC 2, except in scenarios where electricity prices are high and natural gas prices are low. Under these conditions, the SC 1 consumes more natural gas to

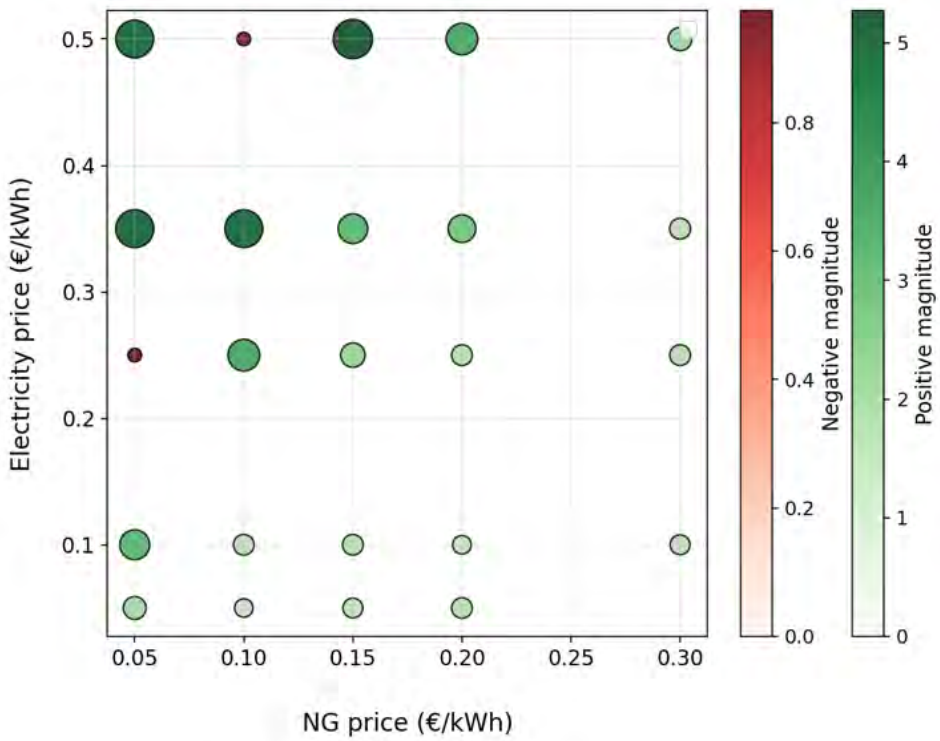

Figure 9. Comparing self sufficiency among different system configuartions, across variable electricity and natural gas prices. Green points indicate choices of optmizer where SC 1 results in higher self-sufficiency than the SC 2. Conversely, red point indicate the choices where SC 2 has higher self sufficiency than SC 1. Higher the size and intensity of the color, larger the difference among both the system configuartions.

satisfy energy demand, as the absence of a carbon penalty favors cost-optimal but fossil-intensive solution.

Environmental analysis

In the base system configuration, the total amount of CO_2 released into the atmosphere is approximately 14, 490 kton/year, of which emissions from fermentation account for around 14, 440 kton/year. In the present analysis, only emissions associated with utility choices are considered for comparison. Although fermentation dominates total CO_2 emissions, the analysis focuses on utility-related emissions, as these are directly influenced by system configuration and energy price assumptions.

Figure 10 compares the differences in emission reductions among both the system configuration under varying prices of electricity, natural gas, and carbon tax. When the cost of natural gas is higher than or equal to the cost of electricity, SC 1 exhibits slightly lower emissions than the SC 2. However, the difference is relatively small, ranging between 0.2 and 0.3 kton/year. This can be attributed to anergy utilization in the SC 1, which reduces the electricity demand for residential heating. In the same price regime, residential heating in the SC 2 is supplied by decentralized heat pumps operating at higher temperature levels, resulting in lower coefficients of performance (COP) and higher electricity consumption.

When natural gas is cheaper than electricity, SC 2 shows lower emissions than the SC 1 in some cases. The maximum observed difference is approximately 2.2 kton/year, occurring at a carbon tax of 80 €/ton, an electricity price of 0.35 €/kWh, and a natural gas price of 0.05 €/kWh. Under these conditions, the SC 1 consumes more natural gas due to increased availability of waste heat. For system configuration 1, when a heat pump is selected, it upgrades heat below the pinch point to temperatures above the pinch point for use in brewery processes. This reduces the amount of waste heat available for exchange within the proposed heat-exchange network, thereby increasing reliance on natural gas. This also explains the presence of two red points in Figure 10(a) when the carbon tax is zero, since higher natural gas consumption does not incur any economical penalty for the case.

Scenario analysis

The sensitivity analysis revealed an interesting scenario under specific energy price conditions. When the

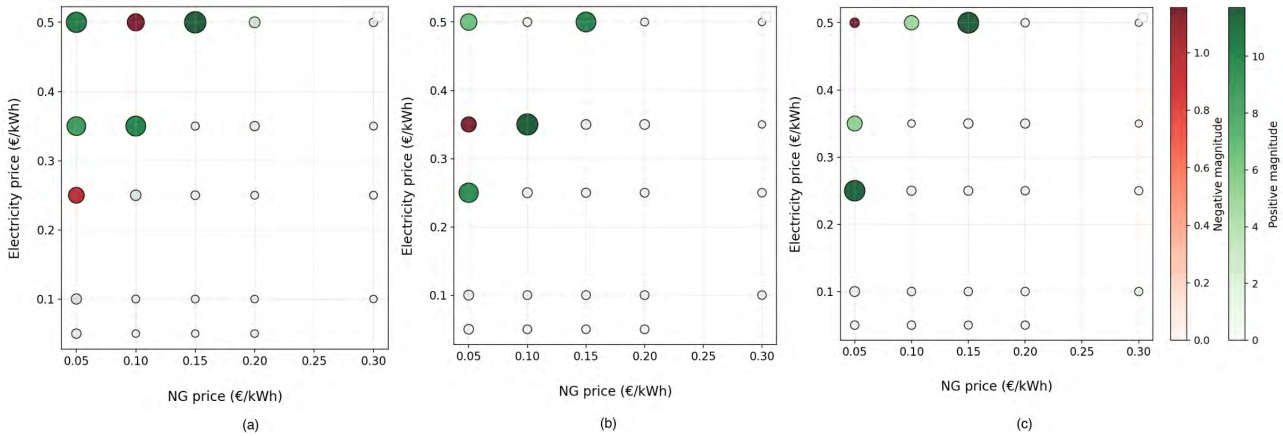

Figure 10: Variation in emission reductions across both the system configuration in varible electricity and natural gas prices. Graphs (a), (b), and (c) represent emission reductions for carbon tax values of 0, 80, and 200 €/tCO_2, respectively.Green points indicate decisions in which SC 1 results in lower emissions than SC 2. Size of green points reflects the absolute difference in emission reduction between the system configurations. Conversely, red points indicate the decision where SC 2 has lower emissions than the SC 1, with larger red points representing greater differences between the two systems.

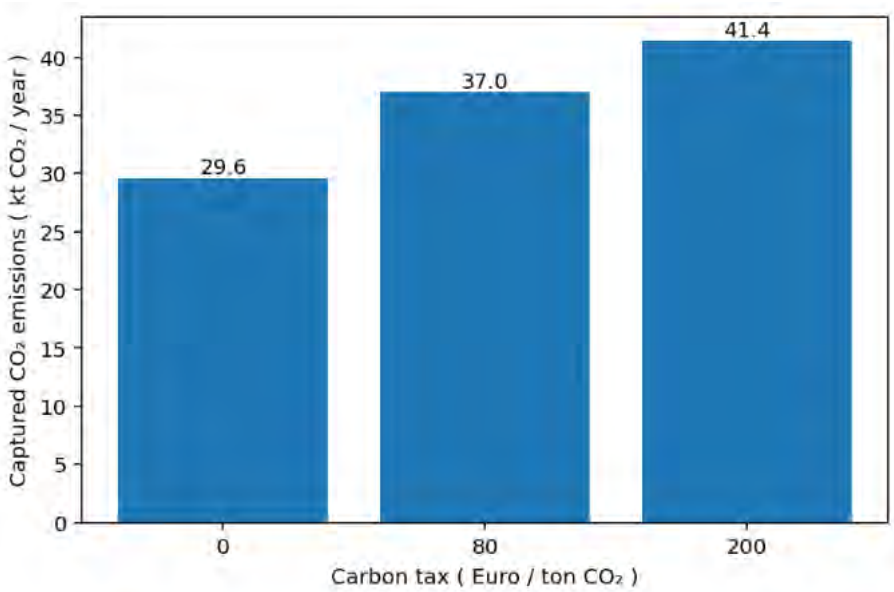

Figure 11. Net amount of CO_2 captured for SC 1 scenario as a function of carbon tax level.

electricity price is approximately 0.05 €/kWh and the natural gas price is around 0.3 €/kWh, the optimizer captures biogenic CO_2 and it is conversted into e-methane via methanation.

The capture and utilization of biogenic carbon enables the brewery system to operate in an energetically carbon-negative mode. In the zero carbon tax scenario (0 €/tCO_2), electricity consumption amounts to 691.7 GWh yr^{-1} for the SC 1 and 1, 055 GWh yr^{-1} for the SC 2, while e-methane production reaches 335 GWh yr^{-1} and 520 GWh yr^{-1}, respectively. Assuming a grid emission factor of 0.054 kg CO_2 kWh^{-1}, the CO_2 emissions associated with electricity consumption remain lower than the amount of biogenic CO_2 captured from fermentation, resulting in net negative emissions in both cases.

In both the system configurations a consistent trend was observed: increasing carbon tax levels leads to more capture of biogenic CO_2: increasing the net amount of CO_2 catured, thereby enhancing the carbon-negative performance of the system (Figures 11 and 12).

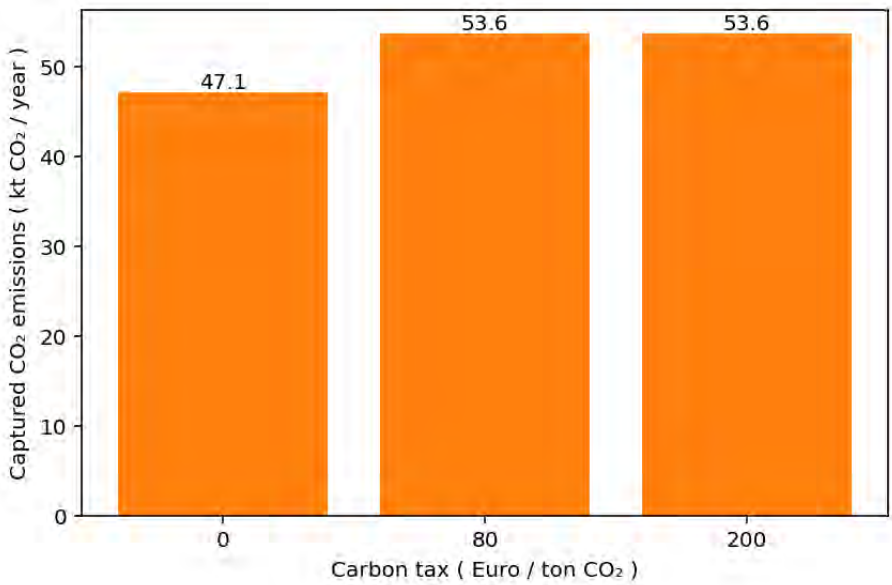

Figure 12. Net amount of CO_2 captured SC 2 as a function of carbon tax level.

Methanation and electrolysis are exothermic, and the resulting heat generation constrains the maximum CO_2 capture through cooling capacity limits. For SC 1, these limits are determined by the combined refrigeration and heat-storage capacities across the brewery, city, and manufacturing sectors enabled by the fifth-generation heat network, whereas in the SC 2 system they are governed by the brewery refrigeration system and cooling tower capacity.

CONCLUSION

This study demonstrates how integrated industrial–urban energy systems can significantly improve energy self-reliance and environmental performance through coordinated utility selection, waste valorization, and heat recovery. By combining temporal demand clustering with cost-optimal optimization using OSMOSE, the analysis highlights the strong influence of energy prices, carbon taxation, and system integration on technology choices.

The results show that SC 1 generally achieve higher self-sufficiency and lower emissions than SC 2 configurations, particularly under higher carbon tax levels. Notably, under favorable energy price conditions, the system enables the capture and conversion of biogenic CO_2 into e-methane, resulting in energetically carbon-negative operation. These findings underline the potential of integrated, low-carbon energy systems and policy instruments such as carbon pricing to support deep industrial decarbonization.

ACKNOWLEDGEMENTS

I would like to thank CIRAIG which is helping me in performing such project and analysis.

REFERENCES

1. Eurostat. "Final Energy Consumption in Industry—Detailed Statistics." (2023).
2. OSR, UK ENERGY IN BRIEF 2024. Department for Energy Security & Net Zero (2024).
3. Okon JE, Oboh IO. Modelling and optimization of a brewery plant from starch sources using aspen plus. J. Eng. Res. Rep. 26:82-93 (2024). https://doi.org/10.9734/jerr/2024/v26i81231
4. Marechal F, Sachan AK, Salgueiro L. Application of process integration methodologies in the brewing industry. Handbook of Process Integration (PI) :820-863 (2013). https://doi.org/10.1533/9780857097255.5.820
5. Gaiser M, Bell GM, Lim AW, Roberts NA, Faraday DBF, Schulz RA, Grob R. Computer simulation of a continuous whisky still. Journal of Food Engineering 51:27-31 (2002). https://doi.org/10.1016/s0260-8774(01)00033-4
6. Flórez-Orrego D, Domingos MER, and Maréchal F. "Heat pumping and renewable energy integration for decarbonizing brewery

industry and district heating." In *Computer Aided Chemical Engineering*, vol. 52, pp. 3177-3182. Elsevier, 2023.

7. Bhatnagar P, Florez-Orrego D, Baibhav V, Margni M and Maréchal F. "Role of process integration and green technologies utilization for the decarbonization of the watchmaking sector." In *Proceedings of the 35th European Symposium on Computer Aided Process Engineering (ESCAPE 35)*, pp. 252-257. PSE Press, 2025.Webpage author. http://hyper.link/deep/link.html
8. Dardor D, Flórez-Orrego D, Terrier C, Ribeiro Domingos ME, Platteau C, da Silva JC, Lopez M, Maréchal F. ROSMOSE: a web-based decision support tool for the design and optimization of industrial and urban energy systems. Energy 304:132182 (2024). https://doi.org/10.1016/j.energy.2024.132182

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Process-Intensified Oscillatory Opposed-Jet Mixers: Mixing Quantification and Operational Guidelines

Sofia P. Brandão[a], Ricardo J. Santos[a], Madalena M. Dias[a], José C. Lopes[ab], and Margarida S. C. A. Brito[a]*

[a] LSRE-LCM, ALiCE, Faculty of Engineering, University of Porto, Rua Dr. Roberto Frias, 4200-465 Porto, Portugal
[b] STAR Institute – Science & Technology Applied Research, Zona Industrial do Salgueiro, Lugar do Pombal, 3530-259 Mangualde, Portugal.
* Corresponding Author: mbrito@fe.up.pt

ABSTRACT

This work presents guidelines for controlling and intensifying mixing in oscillatory opposed-jet mixers, focusing on Confined Impinging Jets (CIJs) as a model system where flow behavior is primarily governed by oscillatory parameters, decoupled from geometric complexity. Computational Fluid Dynamics (CFD) simulations were used to investigate the effects of oscillation amplitude and frequency on mixing. The results show that at high amplitudes, mixing is robust across a broad frequency range, as energy injection is sufficient to promote vortex formation and their propagation to the reactor's outlet. At low amplitudes, mixing is highly sensitive to the oscillation frequency and occurs only near the resonance frequency, the specific frequency at which the flow's response to the applied oscillation is maximized. At low amplitude, lower frequencies fail to inject sufficient energy, while higher frequencies promote flow segregation. Remarkably, effective vortex propagation and mixing were achieved even at very low net Reynolds numbers ($\mathrm{Re_{net}} = 25$), increasing the residence time from 0.042 s at $\mathrm{Re_{net}} = 300$ to 0.5 s at $\mathrm{Re_{net}} = 25$. Two quantitative mixing metrics corroborate the qualitative flow observations and were instrumental in defining practical guidelines for process intensification. These findings provide new insights into how oscillatory forcing can be tuned to control mixing scales, enabling efficient mixing under a wide range of operating conditions, including shear-sensitive processes. The study offers a foundation for the model-based design and optimization of intensified static mixers, linking oscillation parameters, Reynolds number, and residence time to application-specific mixing performance.

Keywords: Computational Fluid Dynamics, Fluid Dynamics, Mixing, Process Intensification, Process Design

INTRODUCTION

Static mixers are characterized by the absence of external moving parts, offering efficient mixing control and low energy consumption, in contrast to mechanically agitated devices. The mixing performance of these devices is strongly dependent on the operating Reynolds number, which should be above a critical threshold, usually $\mathrm{Re_c}$ > 120 [1].

Under these specific conditions, the flow evolves from a steady state to a chaotic regime, characterized by the formation of dynamic vortices that enhance heat and mass transfer [2]. Yet, these advantages are restricted to operations above $\mathrm{Re_c}$, and are ineffective when dealing with applications that require long residence times.

Oscillatory flow is an alternative to improve mixing stability in static mixers across a broader Reynolds number range. It consists of imposing an oscillatory flow rate, defined by an oscillation frequency and amplitude.

This concept is already used in Oscillatory Flow Reactors (OFRs), characterized by a baffled tubular reactor with a superimposed oscillatory stream [3], [4]. Mixing in conventional OFRs has mainly evolved from the interplay between geometric design and oscillation conditions, making it difficult to isolate or generalize the influence of oscillation alone. In both OFRs and the more recent O-

NETmix configuration [5], this interdependence arises from the repeated contractions and expansions in the reactors, which modulate the flow field.

Promising alternatives are systems in which oscillation can act as the primary driver of flow dynamics, such as opposed-jet mixers, which allow the development of predictive models that describe flow behaviour as a function of oscillatory parameters alone, decoupled from geometric complexity and net flow rate.

Moreover, opposed jets mixers have been widely implemented in pharmaceutical applications processes, such as the production of lipid nanoparticles for mRNA encapsulation – e.g., mRNA vaccines - [6], where precise control over micromixing scales is critical for product quality. Mixing is achieved by injecting two streams through opposed jets into a confined chamber, where they evolve into larger vortices. These vortices promote the engulfment of the streams, leading to a phenomenon called the inverse energy cascade, typical of 2D turbulent flows, where the energy is injected from smaller scales, approximately with the size of feed flow streams ($\approx d$), growing to feed larger eddies, with roughly the dimension of the reactor diameter or width ($\approx D$) [7]. The dominant mechanism of energy dissipation is the viscous stress field associated with vortices that evolve throughout the reactor, sustained by the inlet jet streams. Previous studies on oscillatory flow in opposed-jet mixers have shown promising results in controlling the 2D turbulent energy spectrum in response to an external stimulus [8].

However, existing numerical studies remain limited: studies on Oscillatory-CIJs (OCIJs) are at an early proof-of-concept stage, lack quantification of mixing, and cover only a small range of oscillation frequencies and Reynolds numbers, constraining model generalization [8].

The present study aims to establish guidelines for implementing oscillatory flow for efficient mixing in chemical reactors, with a demonstrative application to the CIJs case. It also intends to use oscillatory flow to operate at low Reynolds numbers, achieving longer residence times with comparable mixing efficiencies to those in cases with higher Reynolds numbers.

METHODOLOGY

Computational Fluid Dynamics (CFD)

The flow was simulated using the CFD package ANSYS/Fluent 2021 R2, considering a 2D model of the CIJs geometry (Figure 1). The geometry consists of a rectangular chamber of width D and length H, and two injectors of size d, positioned at a distance h from the top wall.

The continuity, Navier-Stokes, and transport of species equations are solved for an isothermal flow and a Newtonian, incompressible fluid,

$$\nabla \cdot \vec{v} = 0 \tag{1}$$

$$\rho\left(\frac{\partial \vec{v}}{\partial t} + \vec{v} \cdot \nabla \vec{v}\right) = -\nabla p + \mu\, \nabla^2 \vec{v} \tag{2}$$

$$\frac{\partial Y_A}{\partial t} + \nabla \cdot \left(\vec{v} Y_A\right) = D_A\, \nabla^2 Y_A \tag{3}$$

where $\vec{v}$ is the velocity vector field, p the static pressure, Y_A is fluid A mass fraction, and $D_A = 10^{-9}\ \mathrm{m^2\,s^{-1}}$ is the mass diffusion coefficient of fluid A in fluid B. The working fluid consists of a mixture of two viscous fluids (fluids A and B) both with identical physical properties (ρ = 1000 kg/m^3 and μ = 0.02 Pa·s). Conditions of no gravity were considered in the simulations. These properties were considered to mimic the fluids used in CIJs industrial applications.

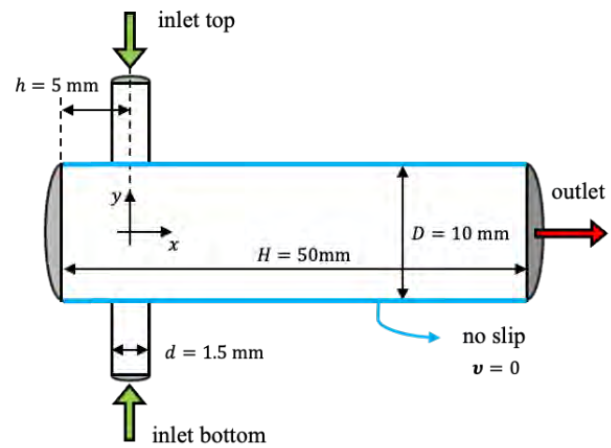

Figure 1. 2D CIJs geometry.

The velocities of the liquid stream have two components: (i) a net flow part characterized by a net Reynolds number, and (ii) an oscillatory part described mathematically by a sine wave.

The net Reynolds number is defined at each inlet by

$$\mathrm{Re_{net}} = \frac{\rho d v_{inj}}{\mu} \tag{4}$$

where ρ and μ are the fluid's density and viscosity, respectively, d is the inlet's width, and v_{inj} is the inlet stream velocity.

The velocity profiles at the inlets are defined by Equation 5 (bottom inlet) and Equation 6 (top inlet)

$$v_{Fluid\,B}(x,t) = \frac{3}{2}\, v_{inj}\left[1 - \frac{4x^2}{d^2}\right]\left[1 + \Psi \sin(2\pi f_o t + \varphi)\right] \tag{5}$$

$$v_{Fluid\,A}(x,t) = \frac{3}{2}\, v_{inj}\left[1 - \frac{4x^2}{d^2}\right]\left[1 + \Psi \sin(2\pi f_o t)\right] \tag{6}$$

where $\Psi = \mathrm{Re_o}/\mathrm{Re_{net}}$ is the amplitude (or velocity ratio), f_o is the oscillatory frequency, and φ is the phase-shift between the two inlet streams. In this study, it was considered $\varphi = \pi$, which means that the liquid streams are injected in the confined chamber in out-of-phase.

The average residence time in the reactor can be computed from

$$\tau = \frac{HD}{2v_{inj}d} \quad (7)$$

Simulations employed a pressure-based laminar solver using the Coupled scheme. Discretization utilized Second Order Implicit time integration, Least Squares Cell-Based gradients, PRESTO! for pressure, and Second Order Upwind for momentum.

The CFD mesh considered was validated in previous works [9], consisting of a fine mesh with an element size of Δx = 50 μm.

Although the adopted grid does not resolve the Batchelor scale associated with the high Schmidt number, capturing these micro-scales is unnecessary for this macroscopic study. The mesh has already been validated for hydrodynamics and the resolution of large-scale transport. Furthermore, because this resolution effect is systematic across all simulated cases, the applied metrics remain a strong indicator for comparison of mixing quantification.

The initial solution for the dynamic simulations was a steady-state solution that considered only the net Reynolds number for each tested condition. The reactor was initially filled with fluid A. The simulations ran for 300 000 time-steps, corresponding to 20 times the residence time (20τ), considering a time-step of $\Delta t = 8.3 \times 10^{-6}$ s, with a convergence criterion of 10^{-4}.

The different conditions tested in the CFD simulations are summarized in Table 1.

Table 1: Conditions used for inlet velocity profiles in CFD

Strategy	τ (s)	Re_{net}	Re_o	Ψ	f_o (Hz)
No oscillation	0.25	50	-	-	-
	0.042	300	-	-	-
Low amplitude	0.25	50	25	0.5	6
					13
					80
High amplitude	0.25	50	300	6	6
					13
					80
Low Reynolds	0.50	25	25	1	6

Mixing of species

Two different approaches were used to quantify mixing and analyse its mechanisms in the CIJs simulations of the binary mixture.

First, the length of the interface between Fluids A and B was computed, which represents the generation of interfacial area in a two-dimensional configuration. The interface can be identified from the simulated flow by extracting an `Iso-Line` corresponding to a fluid's mass fraction of 0.5. The total length of this `Iso-Line`, $s_{interface}$, is then computed.

One of the key characteristics of CIJ reactors is that optimal mixing performance is achieved under chaotic flow conditions, for which the interfacial length grows exponentially along the flow direction [10]. At this regime, mixing can be quantified by the stretching of a material line defining the interface between the two fluids. Consequently, the interface length provides a direct measure of mixing performance: an increase in $s_{interface}$ indicates enhanced stretching and folding, resulting in a larger number of locations where the mass fraction approaches 0.5.

The interfacial area normalized by the inlet's width, d, is defined by

$$\alpha = \frac{s_{\mathrm{interface}}}{d} \quad (8)$$

A second metric was computed from the CFD simulations to quantify the species' mixing area. This metric accounts for the fraction of the reactor area in which the fluid's mass fraction lies within a prescribed tolerance of the completely mixed value of 0.5. A tolerance of $\pm$ 20% was adopted, such that regions with mass fraction values between 0.45 and 0.55 are considered mixed. An `Iso-Clip` was applied in CFD-Post to identify these regions, and the mixing area fraction was computed as

$$\beta = \frac{A_{\mathrm{mixed}}}{A_{\mathrm{reactor}}} \quad (9)$$

where the $A_{reactor} = H \times D = 50 \times 10 = 500$ mm^2. Together, these metrics provide complementary insights into both mixing performance and the mechanisms underlying it.

The α and β values were computed at the last time-step simulated, corresponding to $t = 20\tau$.

RESULTS AND DISCUSSION

Non-oscillatory benchmark

Non-oscillatory flow was simulated at Reynolds numbers of Re = 50 and Re = 300, to serve as a benchmark for the oscillatory cases. Figure 2 shows the interface (black line) and mixing area (green) contour map for these two reference cases without oscillation.

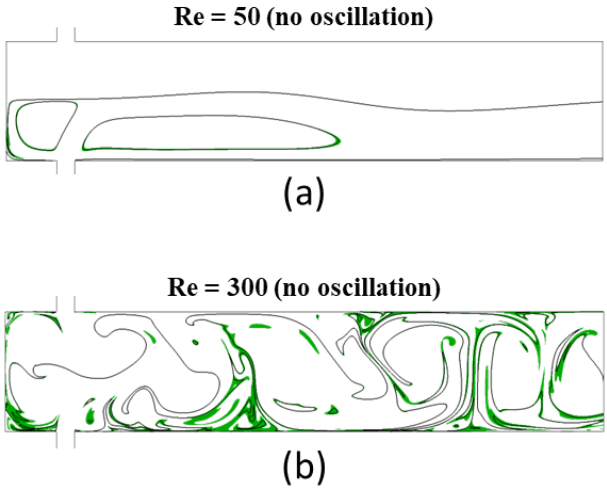

Figure 2. Interface line and mixing area contour maps (green) for non-oscillatory flow at: (a) Re = 50 and (b) Re = 300.

The flow regime at Re = 50 is characterized by steady-state flow, in which the two fluids flow parallel to each other in the CIJs chamber [11]. No convective mixing occurs between the two halves of the CIJs chamber, which is shown by the straight interface line formed between the two fluids. The normalized interfacial area achieved a final value of α = 110. This indicates that the interface between fluids A and B increased to 110 times the initial inlet width, $d = 1.5$ mm. The final mixing area is limited to β = 2%, meaning that only 2% of the reactor achieved a well-mixed state, defined as having mass fractions between 0.45 and 0.55.

In contrast, at Re = 300, the flow regime is characterized by a self-sustained chaotic flow. The two streams from each jet impact at the center of the chamber, engulfing each other and forming large vortices that propagate throughout the main chamber [11]. The interface line delineates the vortices circular shape, achieving a final value of $\alpha = 312$. The chaotic advection occurring at this regime significantly enhances homogeneity. The well-mixed region expanded to cover a mixing area of β = 12 %. These results confirm the strong dependence of the CIJs' mixing performance on the inlet Reynolds number. Operating at Re = 300 produced a threefold increase in interfacial area (α) and a sixfold increase in mixing area (β) compared to the Re = 50 case.

Oscillatory Flow

Previous work suggests that flow dynamics improve when the oscillation frequency resonates with the frequency of vortex formation [5], [8]. Based on this, a frequency screening was performed at a net Reynolds number of Re = 50 with amplitudes of Ψ = 0.5 and Ψ = 6. The frequencies tested (6 Hz, 13 Hz, and 80 Hz) were specifically chosen to match vortex formation frequencies at Re = 25, 50, and 300. The frequencies of vortex formation (f_{vortex}) were computed using the following equation [7], [8]

$$f_{vortex} = \frac{4\mu\,\mathrm{Re}}{\pi\rho\,D^2} \tag{10}$$

Additionally, a lower net Reynolds number, Re = 25, imposing the resonance frequency of 6 Hz and an amplitude Ψ = 1 was also tested.

Flow dynamics and contour plots

Figures 3 to 5 show the interface (black line) and mixing area (green) for all oscillatory conditions tested.

At low amplitude (Re50_Ψ0.5), the hydrodynamic behaviour strongly depends on the frequency used. With an oscillatory Reynolds number of $\mathrm{Re_o}$ = 25, significantly lower than the critical value of $\mathrm{Re_c}$ = 120, the flow regime is dominated by frequency, with amplitude playing a secondary role. This is clearly observed from the interface lines in Figure3b: at f = 13 Hz, with a resonance frequency of Re = 50, there is effective vortex shedding inside the CIJs chamber. Although at f = 6 Hz (Figure 3a), large vortices are formed near the impact point, they tend to lose their shape at the reactor's outlet. At f = 80 Hz (Figure 3c), the oscillation becomes too fast for the low amplitude, resulting in the formation of a waving interface rather than large vortices.

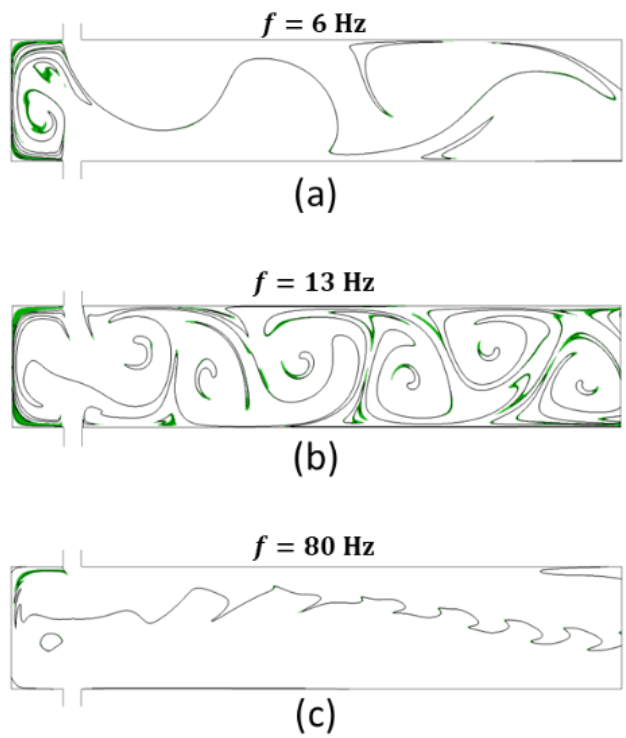

Figure 3. Interface line and mixing area contour maps (green) for oscillatory flow with low amplitude ($\Psi = 0.5$) for: (a) $f = 6$ Hz, (b) $f = 13$ Hz and $f = 80$ Hz.

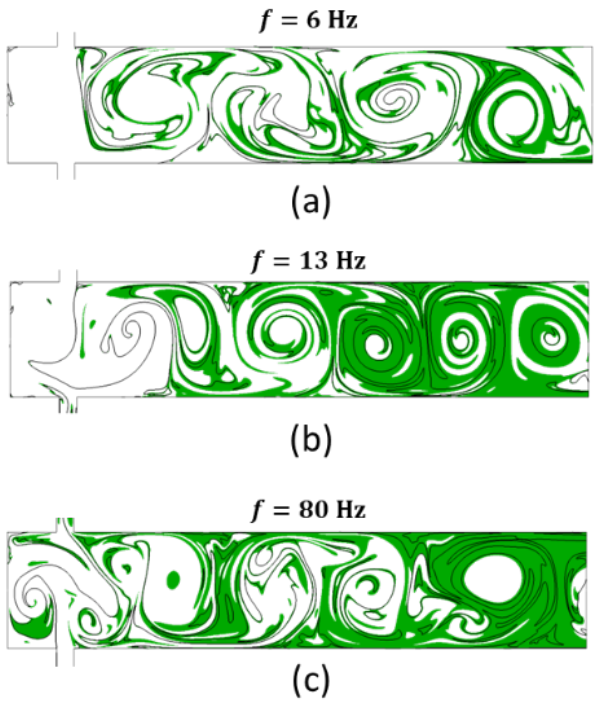

Figure 4. Interface line and mixing area contour maps (green) for oscillatory flow with high amplitude ($\Psi = 6$) for: (a) $f = 6$ Hz, (b) $f = 13$ Hz and $f = 80$ Hz.

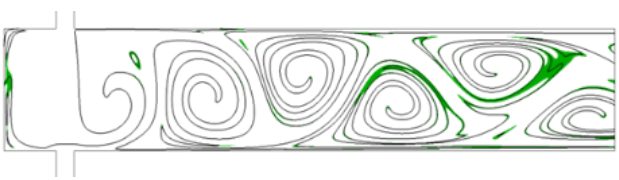

Figure 5. Interface line and mixing area contour maps (green) for oscillatory flow with low Reynolds ($\mathrm{Re}25_\Psi1_f_o6$Hz).

In contrast, at high amplitude ($\mathrm{Re}50_\Psi6$), the effect of the oscillation is very similar across the three frequencies studied (6 Hz to 80 Hz). Under these conditions, the oscillatory Reynolds achieves $\mathrm{Re_o}$ = 300, which is above the critical Reynolds. At f_o = 6 Hz (Figure 4a), the vortices formed have an elongated, ellipsoidal shape, which becomes more circular and organized when the frequency is increased to f_o = 13 Hz (Figure 4b). The resonance effect leads to the formation of more coherent circular vortices of size D along the entire reactor. In the 80 Hz case (Figure 4c), the vortices degenerate into smaller structures. A progressive increase in the mixing area (green zone) with increasing frequency is observed. A larger mixing area is achieved for f_o = 80 Hz, at high amplitudes.

However, improved control of the mixing scales, defined by the size of the resulting vortex structures, is achieved at low oscillation frequencies, most notably at the resonance frequency of f_o = 13 Hz.

For the low Reynolds condition (Figure 5), a frequency of 6 Hz was applied in the oscillation, matching the frequency of vortex formation at Re = 25. Under these conditions, successive circular vortices develop along the chamber, during which the fluid becomes striated, and lamellar structures are formed. Although the mixing area is smaller than in the high-amplitude cases (Figure 4), the oscillatory Reynolds number is $\mathrm{Re}_o = 25$, which is twelve times lower than the $\mathrm{Re}_o = 300$ used in the high-amplitude case. Consequently, the mixing mechanism is non-chaotic and more closely resembles the striation observed in the low-amplitude cases shown in Figure 3.

In summary, these qualitative observations demonstrate that the mixing phenomenon in oscillatory flow is governed by the interplay between the imposed oscillation frequency and amplitude. The oscillation amplitude acts as the main mechanism for energy injection, providing the necessary momentum to transport vortices from the impact point to the outlet of the chamber.

The energy injection is sufficient to exceed the critical limit when the amplitude considered satisfies $\Psi > \mathrm{Re_c}/\mathrm{Re_{net}}$, a condition that directly drives the oscillatory Reynolds number above the critical value ($\mathrm{Re_o} > \mathrm{Re_c}$). Here, vortex formation becomes widespread and independent of the imposed frequency, even while the net Reynolds number remains below the critical value.

However, when the oscillation frequency is tuned to the natural vortex shedding frequency (defined by $\mathrm{Re_{net}}$), a resonance effect is generated. This resonance allows the instabilities formed at the impingement point to propagate along the length of the reactor, enabling efficient mixing even at low amplitudes.

Mixing quantification

Figure 6 shows the normalized interfacial area α and mixing area β for the low amplitude (orange triangle), high amplitude (black square), and low net Reynolds (yellow circle) conditions.

Both plots, α and β, confirm a strong frequency dependence in the low-amplitude case (orange). At f_o = 13 Hz, α surpasses the value obtained for the non-oscillatory flow at Re = 300 (green dashed line), while at f_o = 80 Hz, α is below the Re = 50 benchmark (red dashed line). This implies that the combination of low amplitude and high frequency results in a quasi-segregated flow regime, where the two inlet streams barely interact.

Despite the high α at 13 Hz, the mixing area (β) never exceeds the Re = 300 benchmark. This indicates that low-amplitude oscillatory mixing is driven by laminar stretching and folding rather than the chaotic advection associated with higher Reynolds numbers.

On the other hand, α and β are always above the Re = 300 values, when high amplitude is applied (black). This efficient performance indicates that chaotic advection is occurring under these oscillatory conditions.

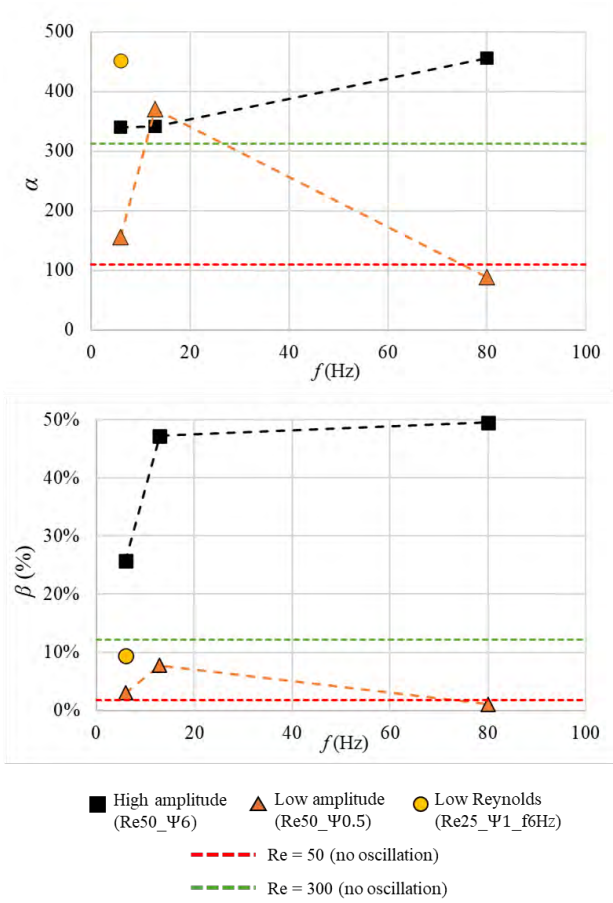

Figure 6. Mixing metrics for oscillatory flow in CIJs and non-oscillatory references: (a) normalized interfacial area, and (b) normalized mixing area.

Although the interfacial area values are significantly higher at f_o = 80 Hz, this behavior is not reflected in the normalized mixing area, which shows similar values at $\beta \approx$ 50%, for f_o = 13 Hz and 80 Hz. This means that approximately 50% of the CIJs area is well-mixed, according to the defined criterion.

Finally, for the low net Reynolds condition (yellow), the normalized interfacial area generated is high, and comparable to the high amplitude case with f_o = 80 Hz. This highlights a synergetic effect of frequency and amplitude, since in this case the frequency chosen resonates with both the net Reynolds and the oscillatory Reynolds because $\Psi = 1 \Rightarrow \mathrm{Re_{net}} = \mathrm{Re_o}$. However, as expected, the mixing mechanism is similar to the low amplitude case, laminar stretching and folding, leading to β values below the Re = 300 benchmark.

Guidelines for oscillatory mixing in CIJs

Based on these results, optimal mixing is obtained by operating at the resonance frequency of the net Reynolds number and using high oscillation amplitudes.

However, the choice of oscillatory conditions strongly depends on the intended application. Operating at low amplitudes, for which the oscillatory Reynolds number Re_o does not exceed the critical Reynolds threshold, is an advantage for shear-sensitive processes. In this regime, the flow remains at low Reynolds numbers, thereby limiting shear stresses and strong inertial forces. Moreover, at low amplitudes, the flow dynamics are entirely governed by the imposed sinusoidal inlet flow rate, in addition to the associated increase in residence time.

In contrast, high amplitudes are particularly advantageous for processes characterized by low net Reynolds numbers that require fast, efficient mixing, if shear sensitivity is not a limiting factor.

Figure 7 shows a schematic overview of the proposed guidelines for process intensification in oscillatory opposed-jet mixers.

The primary objective of intensifying mixing in these devices is to increase the average residence time without compromising mixing quality. Therefore, the first step is to determine the required average residence time, τ, based on the specific application (e.g., crystallization or cell growth). If Equations 4 and 7 are combined, the net Reynolds number can be computed from the residence time as $\mathrm{Re_{net}} = \frac{\rho H D}{2\mu\tau}$.

It is important to note that this procedure is only valid for $\mathrm{Re_{net}} < \mathrm{Re_c}$. This constraint must therefore be considered when selecting the reactor residence time.

The second step concerns the computation of the oscillatory frequency using Equation 8. Then, the oscillation amplitude should be selected according to the nature of the application: (i) for shear-sensitive processes, a low amplitude is recommended; (ii) for non-shear-sensitive processes, a high amplitude may be employed. The amplitude is classified as low or high depending on whether $\mathrm{Re_o} < \mathrm{Re_c}$ (i.e., $\Psi < \mathrm{Re_c}/\mathrm{Re_{net}}$), or $\mathrm{Re_o} > \mathrm{Re_c}$ (i.e., $\Psi > \mathrm{Re_c}/\mathrm{Re_{net}}$).

Finally, at low oscillation amplitudes, laminar mixing dominates, leading to the formation of an ordered vortex street and predictable fluid striation dictated by the sinusoidal inlet conditions. In contrast, high-amplitude oscillations promote chaotic advection, resulting in a higher degree of fluid homogeneity than in the non-oscillatory chaotic flow regime used as a benchmark.

CONCLUSIONS

This work represents the first quantification of mixing in oscillatory CIJs using these metrics, interfacial area and mixing area. These metrics were fundamental in defining practical guidelines for process intensification and corroborating the qualitative flow observations. This study also provides new insights into the combined effects of amplitude and frequency and demonstrates effective mixing at net Reynolds numbers way below the critical value.

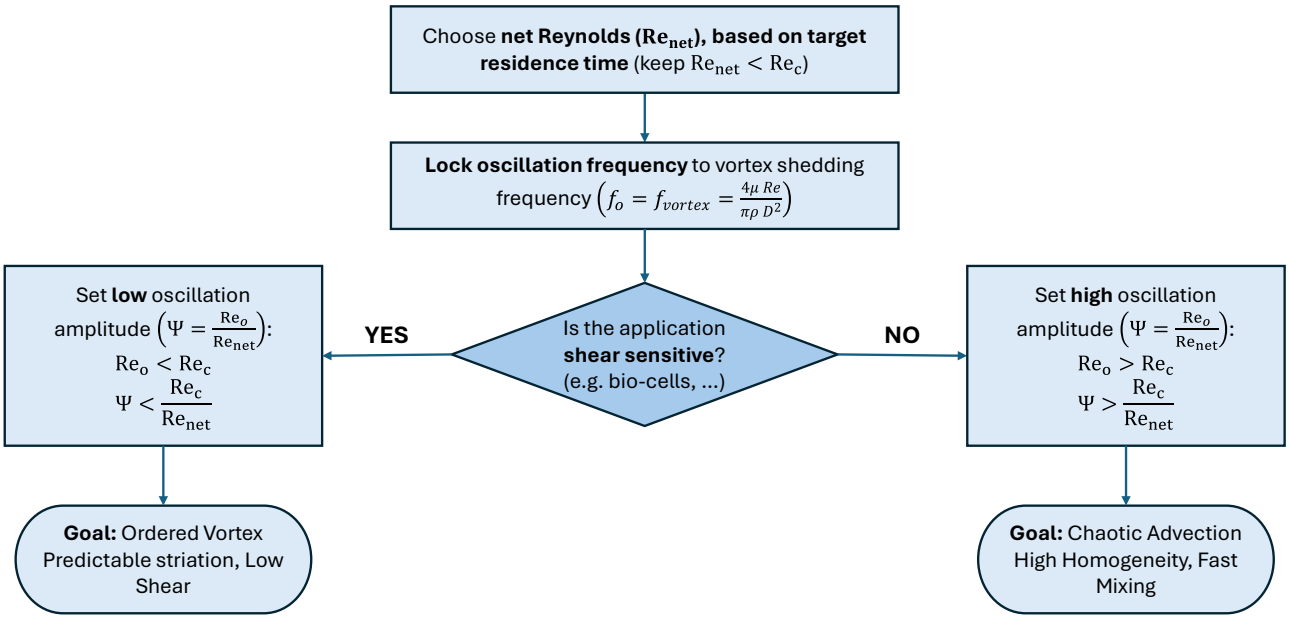

Figure 7. Guidelines for mixing in opposed jets mixers

Optimal mixing in oscillatory opposed-jet mixers is achieved at the resonance frequency of the net Reynolds number with high oscillation amplitudes. Nevertheless, the choice of oscillatory conditions depends on the intended application: low amplitudes favor shear-sensitive processes by maintaining laminar mixing, while high amplitudes enhance mixing efficiency and homogeneity in non-shear-sensitive applications.

These findings offer a framework for setting mixing efficiency and dependent product properties from control of the residence time, oscillation frequency, and amplitude based on application requirements, while ensuring operation within the appropriate Reynolds number regime.

ACKNOWLEDGEMENTS

This work was financially supported by Fundação para a Ciência e a Tecnologia, I.P. /MCTES through national funds: LSRE-LCM, UID/50020/2025 (DOI: 10.54499/UID/50020/2025); and ALiCE, LA/P/0045/2020 (DOI: 10.54499/LA/P/0045/2020). This work was also supported by the project ComputReactors: Advanced computational design of novel chemical reactors, 2023.10740.CPCA.A2, (DOI: 10.54499/2023.10740.CPCA.A2). Margarida Brito acknowledges the FCT Individual Call to Scientific Employment Stimulus O-MIX - Mixing Principles in Modulated Flow Reactors, 2023.06416.CEECIND (DOI:10.54499/2023.06416.CEECIND/CP2834/CT0008). Sofia Brandão acknowledges her FCT PhD scholarship 2023.00452.BD (DOI: https://doi.org/10.54499/2023.00452.BD).

AUTHOR IDENTIFIERS

Author ORCIDs:
Brandão S: 0000-0002-0267-4100
Santos R: 0000-0002-9133-2187
Dias MM: 0000-0003-1208-1496
Lopes JCB: 0000-0001-6127-6340
Abrunhosa Brito M: 0000-0003-1134-1743

REFERENCES

1. Santos RJ, Erkoç E, Dias MM, Teixeira AM, Lopes JCB. Hydrodynamics of the mixing chamber in RIM: PIV flow?field characterization. AIChE Journal 54:1153-1163 (2008). https://doi.org/10.1002/aic.11472
2. Santos RJ, Sultan MA. State of the art of mini/micro jet reactors. Chem Eng & Technol 36:937-949 (2013). https://doi.org/10.1002/ceat.201200678
3. Avila M, Kawas B, Fletcher DF, Poux M, Xuereb C, Aubin J. Design, performance characterization and applications of continuous oscillatory baffled reactors. Chemical Engineering and Processing - Process Intensification 180:108718 (2022). https://doi.org/10.1016/j.cep.2021.108718
4. Ferreira A, Rocha F, Mota A, Teixeira JA. Nonmechanically agitated bioreactors. Current

Developments in Biotechnology and Bioengineering :217-233 (2017). https://doi.org/10.1016/b978-0-444-63663-8.00008-2
5. Fernandes IS, Matos J, Dias MM, Lopes JCB, Santos RJ. A new operating mode of netmix for active mixing under batch operation. Chemical Engineering Science 313:121707 (2025). https://doi.org/10.1016/j.ces.2025.121707
6. Devos C, Mukherjee S, Inguva P, Singh S, Wei Y, Mondal S, Yu H, Barbastathis G, Stelzer T, Braatz RD, Myerson AS. Impinging jet mixers: a review of their mixing characteristics, performance considerations, and applications. AIChE Journal 71: (2024). https://doi.org/10.1002/aic.18595
7. Gonçalves ND, Salvador HM, Fonte CP, Dias MM, Lopes JCB, Santos RJ. On the 2D nature of flow dynamics in opposed jets mixers. AIChE Journal 63:2335-2347 (2016). https://doi.org/10.1002/aic.15566
8. Brito MSCA, Brandão SP, Gonçalves ND, Fonte CP, Dias MM, Lopes JCB, Santos RJ. On the control of 2D turbulence by active mixing in opposed jets mixers. Chemical Engineering Research and Design 195:272-285 (2023). https://doi.org/10.1016/j.cherd.2023.05.045
9. C. P. Fonte, R. J. Santos, M. M. Dias, and J. C. B. Lopes, “Quantification of mixing in RIM using a non-diffusive two-phase flow numerical model, ” International Journal of Chemical Reactor Engineering, vol. 9, no. 1, Dec. 2011, doi: 10.2202/1542-6580.2564/MACHINEREADABLECITATION/RIS.
10. J. M. Ottino, The Kinematics of Mixing: Stretching, Chaos, and Transport. in Cambridge Texts in Applied Mathematics. Cambridge University Press, 1989. [Online]. Available: https://books.google.pt/books?id=8OLVcbRoNSgC
11. Fonte CP, Sultan MA, Santos RJ, Dias MM, Lopes JCB. Flow imbalance and reynolds number impact on mixing in confined impinging jets. Chemical Engineering Journal 260:316-330 (2015). https://doi.org/10.1016/j.cej.2014.08.090

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Section 6

Concepts, Methods and Tools

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Multimodal Framework Integrating Procedural Texts and Visual Perception for Laboratory Safety Monitoring

Shuo Xu[ab], Jinsong Zhao[ab]*

[a] Tsinghua University, Department of Chemical Engineering, Beijing, China
[b] State Key Laboratory of Chemical Engineering and Low-Carbon Technology, Tsinghua University
* Corresponding Author: jinsongzhao@tsinghua.edu.cn.

ABSTRACT

Laboratory safety is procedure-dependent: required personal protective equipment (PPE) and permissible actions vary across experiments and across experimental steps, yet most vision-based monitoring remains appearance-driven and often produces generic warnings without reliable procedural context. We propose a multimodal framework for step-aware safety monitoring in laboratory videos. The framework first localizes procedural context through clip-level step prediction and protocol alignment to identify the experiment and current step. Given this context, it retrieves step-specific safety constraints, extracts evidence of step-relevant equipment and interactions using an equipment database, and prompts a video-capable vision-language model (VLM) to generate structured (JSON) monitoring reports supported by retrieved constraints and visual evidence. Experiments on protocol-annotated molecular biology lab videos show that our approach improves the mean score from 0.4352 to 0.6430 and reduces the missing rate from 65.00% to 33.75% relative to a video-only baseline, demonstrating more faithful and step-specific safety judgments.

Keywords: Artificial Intelligence, Laboratory Safety Monitoring, Vision-Language Model

1. INTRODUCTION

Laboratory research—especially in chemistry and molecular biology—often involves hazardous chemicals, biohazards, pressurized systems, and high-risk instruments. Safe operation therefore requires step-specific compliance: the required personal protective equipment (PPE), permissible actions, and risk controls can change across experiments and even across steps within the same experiment. However, much laboratory safety practice is still evaluated using coarse, self-reported proxies rather than direct, task-level evidence of compliance [1]. Accordingly, modern safety principles emphasize assessing hazards for the specific task and revisiting controls as conditions evolve across steps [2]. This motivates methods that can continuously observe laboratory work and anchor safety judgments in the current procedural context, enabling step-aware checks rather than generic, appearance-driven warnings.

In busy laboratories, day-to-day compliance is difficult to ensure despite training and written protocols. Manual supervision is costly and intermittent, while violations (e.g., missing eye protection, unsafe reagent handling, or incorrect waste practices) may occur abruptly and unpredictably. Computer vision (CV) has therefore been explored for automated inspection—most notably PPE compliance monitoring—because it can, in principle, provide scalable and continuous observation [3, 4]. Yet many existing pipelines remain largely appearance-driven: they often focus on detecting visible PPE items or generic unsafe behaviors, and do not reliably condition judgments on which experiment is being performed and which step is underway. This is a critical limitation because identical motions or object interactions can be compliant in one step but unsafe in another.

A key missing piece is reliable scenario recognition—identifying the ongoing experiment and localizing the current step—so that monitoring can retrieve and apply the correct step-level requirements. Recent procedure-aware video understanding research shows that long activities can be modeled as temporally ordered steps and aligned to protocol text, providing a natural mechanism to connect video observations with procedural context [5-9]. For laboratory environments, ProBio

offers protocol-guided, fine-grained hierarchical annotations for molecular biology experiments, making step-level modeling in realistic lab settings practically feasible [10]. These advances suggest that step localization plus protocol alignment can anchor downstream safety reasoning to the correct procedural context.

In parallel, large vision-language models (VLMs) have advanced toward open-ended multimodal reasoning and instruction following [11, 12], which is attractive for safety monitoring where rules are textual and evidence is visual. However, open-ended generation can produce hallucinated or weakly evidence-supported claims when models lack access to task-specific rules or when evidence is ambiguous [13]. Retrieval-augmented generation (RAG) provides a principled way to inject verifiable external knowledge by conditioning outputs on retrieved documents, improving controllability and evidence linkage [14-16]. For laboratory safety, this suggests combining step-aware retrieval of safety constraints with evidence-supported interpretation of video evidence.

In this work, we propose a procedure-aware laboratory safety monitoring framework that explicitly conditions safety reasoning on the current procedural context. Our key idea is to decouple low-cost procedural localization from higher-cost safety reasoning: we first infer the experiment identity and step context from long, untrimmed videos, and then use step-conditioned retrieval and evidence linkage to produce step-aware, structured safety monitoring reports. Experiments on protocol-annotated molecular biology laboratory videos show that incorporating procedural context substantially improves both faithfulness and coverage compared with a video-only VLM baseline.

Our contributions are threefold:

- **Protocol-aware scenario recognition.** We formulate scenario recognition as clip-level experimental step identification and protocol-based experiment linkage, enabling step-level alignment in long laboratory videos.
- **Evidence-linked multimodal monitoring.** We integrate retrieved step-level constraints and structured equipment evidence with video observations to improve faithfulness and reduce generic warnings and poorly evidence-supported claims.
- **End-to-end video-to-safety framework.** We combine procedure-aware video understanding with VLM-based reasoning to move beyond isolated PPE checks toward context-sensitive, step-specific safety monitoring.

The remainder of this paper is organized as follows. Section 2 presents our methodology, including scenario recognition and the proposed multimodal safety monitoring pipeline. Section 3 reports experimental results and discussion. Finally, Section 4 concludes and outlines limitations and future directions.

2. METHODOLOGY

2.1 Problem Definition and Framework Overview

We study procedure-aware laboratory safety monitoring, which maps a long laboratory video to step-aware, structured safety issues. As illustrated in **Figure 1**, our framework consists of two stages: (i) scenario recognition, which localizes the procedural context by clip-level experimental step identification and experiment linkage, and (ii) multimodal safety monitoring, which retrieves step-level safety constraints, extracts equipment evidence, and prompts a video-capable VLM to produce structured issues.

Given a laboratory video V that contains a single experiment, we split it into T 1-second clips:

$$V = \{x_t\}_{t=1}^{T} \tag{1}$$

We assume three data sources: a candidate experiment set E, where each experiment $e \in E$ has an ordered procedure $P_e = \langle s_1^e, \dots, s_{L_e}^e \rangle$; a step-level safety analysis report library R; and an equipment database D that stores equipment names and visual references.

We predict the experimental step for each clip:

$$\widehat{s_t} = \mathrm{StepID}(x_t) \tag{2}$$

Let $\hat{S} = \langle \widehat{s_1}, \dots, \widehat{s_T} \rangle$ denote the predicted step sequence. We perform experiment linkage for the full video by aligning $\hat{S}$ with candidate protocols:

$$\hat{e} = \mathrm{argmax}_{e \in E}\ \mathrm{Align}(\hat{S}, P_e) \tag{3}$$

To reduce the number of expensive VLM calls while retaining temporal context, we monitor variable-length segments with a consistent predicted step label. Specifically, we partition the clip sequence into K maximal contiguous segments $\{z_k\}_{k=1}^{K}$, such that all clips within a segment share the same predicted experimental step label:

$$z_k = \{x_t\}_{t=\tau_k}^{\tau_{k+1}-1} \tag{4}$$

$$\widehat{s_t} = \widetilde{s_k} \quad \forall t \in [\tau_k, \tau_{k+1} - 1] \tag{5}$$

Where $\tau_1 = 1, \tau_{K+1} = T + 1$, and $\widetilde{s_k}$ is the shared step label of segment z_k. The segment length is therefore adaptive and determined by the duration of the predicted experimental step run, i.e., $W_k = \tau_{k+1} - \tau_k$.

Given $(\hat{e}, \tilde{s}_k)$, we retrieve the corresponding step-level safety analysis report and its equipment list:

$$r_k = \mathrm{Retrieve}(R, \hat{e}, \widetilde{s_k}) \tag{6}$$

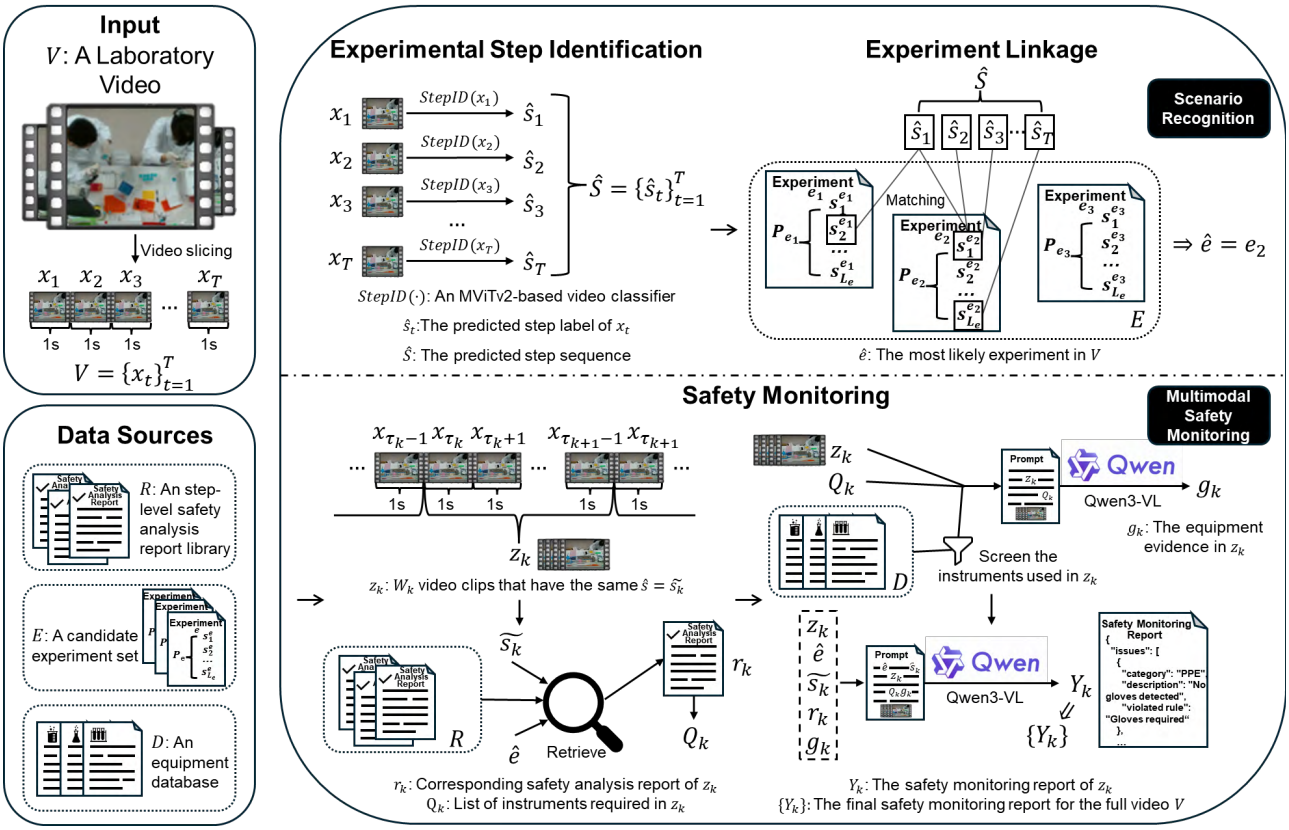

Figure 1. Overall framework.

$$Q_k = \mathrm{EquipList}(r_k) \tag{7}$$

We then extract equipment evidence in the segment and prompt the VLM to output structured safety issues:

$$g_k = \mathrm{ExtractEquipEvidence}(z_k, Q_k, D) \tag{8}$$

$$Y_k = \mathrm{Monitor}(z_k, \hat{e}, \widetilde{s_k}, r_k, g_k) \tag{9}$$

Aggregating $\{Y_k\}$ over all segments yields the final safety monitoring report for the full video. This formulation separates dense, low-cost procedural localization from sparser, higher-cost safety reasoning, while ensuring each monitored segment is associated with a single experimental step.

2.2 Data Sources

Our framework relies on three complementary data sources to connect procedural text with visual evidence. We first use ProBio molecular biology laboratory recordings with protocol and step annotations to train and evaluate scenario recognition [10]. Each video is split into 1-second clips to produce step predictions over time, and the resulting step sequence is later used for experiment linkage by aligning it to candidate protocols.

Once the experiment and current step are localized, we anchor monitoring in a step-level safety analysis report library that is aligned with experimental procedures. Each entry corresponds to a specific experiment step and summarizes step intent, potential hazards, required PPE, cautions, and forbidden actions. Importantly, the same entry also specifies the equipment involved in that step, providing an explicit, procedure-derived prior for downstream equipment understanding.

To support reliable equipment evidence extraction in video, we further maintain an equipment database that maps equipment names to prompt-ready appearance descriptions and visual references when available. The database covers common laboratory instruments and items such as centrifuges, pipettes, 96-well plates, EP tubes, and frequently handled reagents stored in tubes or vials. Conditioned on the report-provided equipment list, a VLM uses these appearance cues to extract equipment location evidence and human-equipment interaction evidence from the video. This equipment-level evidence is then passed to the final safety monitoring stage to enable more specific and better-supported safety judgments.

2.3 Scenario Recognition

To retrieve the correct step-level safety constraints, the system must first localize procedural context in a long, untrimmed video. Scenario recognition provides this context by predicting a step label for each 1-second clip and linking the entire video to an experiment via protocol alignment. These outputs serve as the conditioning signal

for safety monitoring in the next stage.

2.3.1 Experimental step identification

Given a laboratory video $V = \{x_t\}_{t=1}^{T}$, we infer the current experimental step using a video classifier. Specifically, the clip-level step prediction is defined in **Eq. (2)**.

We instantiate $StepID(\cdot)$ with an MViTv2-based classifier [17] finetuned on ProBio step annotations [10]. In our implementation, we finetune and evaluate the classifier on step-annotated video segments whose durations range from a few seconds to tens of seconds. Each segment is represented by uniformly sampled frames resized to 224 × 224, and we adopt a fixed train / validation / test split after filtering extremely rare step categories. Step identification is evaluated at the segment level using top-1 accuracy and weighted F1 to account for class imbalance.

Separately, during inference on an untrimmed video V, we split it into 1-second clips $\{x_t\}_{t=1}^{T}$ and apply $StepID(x_t)$ to obtain a step sequence $\hat{S} = \langle \hat{s}_1, \dots, \hat{s}_T \rangle$. For robustness in downstream linkage, we retain the top-3 predicted steps (by posterior probability) in addition to the top-1 label $\hat{s}_t$.

2.3.2 Experiment linkage

Given the predicted step sequence $\hat{S} = \langle \hat{s}_1, \dots, \hat{s}_T \rangle$, we link the long video to the most likely experiment by aligning it to each candidate protocol using **Eq. (3)**.

We implement $Align(\cdot)$ as a soft matching score based on the step posterior from $StepID(\cdot)$. Let $p_t(s)$ be the predicted probability of step s for clip t. For each experiment, we accumulate probability mass on protocol steps and average the log score over supported clips:

$$\mathrm{Align}(\hat{S}, P_e) = \frac{1}{|\Omega_e|} \sum_{t \in \Omega_e} log\Big(\sum_{s \in P_e} p_t(s)\Big) \tag{10}$$

$$\Omega_e = \Big\{ t \;\Big|\; \sum_{s \in P_e} p_t(s) > 0 \Big\} \tag{11}$$

In practice, we approximate $p_t(\cdot)$ by retaining only the top-3 predicted steps for each clip and renormalizing their scores to form a sparse distribution, while setting the probabilities of all remaining steps to zero. This reduces noise from low-confidence classes and improves robustness in long, untrimmed videos.

2.4 Multimodal Safety Monitoring

After scenario recognition, we perform procedure-aware safety monitoring on multi-second segments z_k that stay within a single predicted step (**Eq. (4)-(5)**). For each segment, we use the linked experiment $\hat{e}$ and step $\widetilde{s_k}$ to retrieve the corresponding step-level safety entry, as defined in **Eq. (6)**, which summarizes step intent, hazards, required PPE, cautions, forbidden actions, and a structured set of monitoring rules (e.g., PPE, required objects, environment constraints, forbidden actions).

To improve evidence support, we extract a step-specific equipment list Q_k from r_k and combine it with the equipment database D to obtain equipment evidence, following **Eq. (7)-(8)**.

We implement $ExtractEquipEvidence(\cdot)$ as a constrained VLM call for evidence extraction rather than safety judgment. Given z_k and the step-specific equipment list Q_k (from r_k), we augment each candidate with appearance cues from D and ask the VLM to verify which items are visible and how they are being used. The output g_k is a structured equipment-evidence summary (e.g., present/absent/uncertain with brief visual and interaction cues) that is fed into $Monitor(\cdot)$ to provide stronger evidential support and reduce hallucinations.

Finally, we prompt a video-capable VLM (Qwen3-VL) with (i) the segment z_k, (ii) metadata $(\hat{e}, \widetilde{s_k})$, (iii) retrieved constraints r_k, and (iv) verified evidence g_k, and ask it to produce structured safety issues as defined in **Eq. (9)**. The output is formatted as a list of issues (category, description, violated rule, supporting evidence). Aggregating $\{Y_k\}$ over all segments yields the final video-level safety monitoring report, naturally organized by experimental step.

3. RESULTS AND DISCUSSION

3.1 Scenario Recognition

Scenario recognition provides the procedural context required for step-level safety retrieval; we therefore evaluate both experimental step identification and protocol-based experiment linkage on long, untrimmed laboratory videos. In this section, we report results for the two components: Experimental step identification and Experiment linkage. Overall, $StepID(\cdot)$ achieves strong step recognition, and protocol alignment yields reliable experiment linkage, especially when videos contain multiple distinct experimental steps.

3.1.1 Experimental step identification

We finetune an MViTv2-based classifier (initialized from Kinetics-400 pretraining) to recognize experimental steps in ProBio [10, 17]. After filtering extremely rare step categories (fewer than 5 samples), the final label space contains 77 step classes. The dataset is split into 2, 232 / 1, 116 / 372 samples for training / validation / testing (0.6 / 0.3 / 0.1). As shown in **Table 1**, the best checkpoint is selected by validation accuracy (epoch 36, 0.8952) and achieves 0.8952 top-1 test accuracy and 0.8905 weighted-F1, indicating strong overall performance under class imbalance.

Figure 2 shows that the confusion matrix is strongly diagonal, with remaining errors concentrated in a few visually similar step pairs. Most frequent confusions occur between near-duplicate primitives with almost identical motions and objects (e.g., different “incubate on ice” targets, and “mix/invert/resuspend” variants), as well as

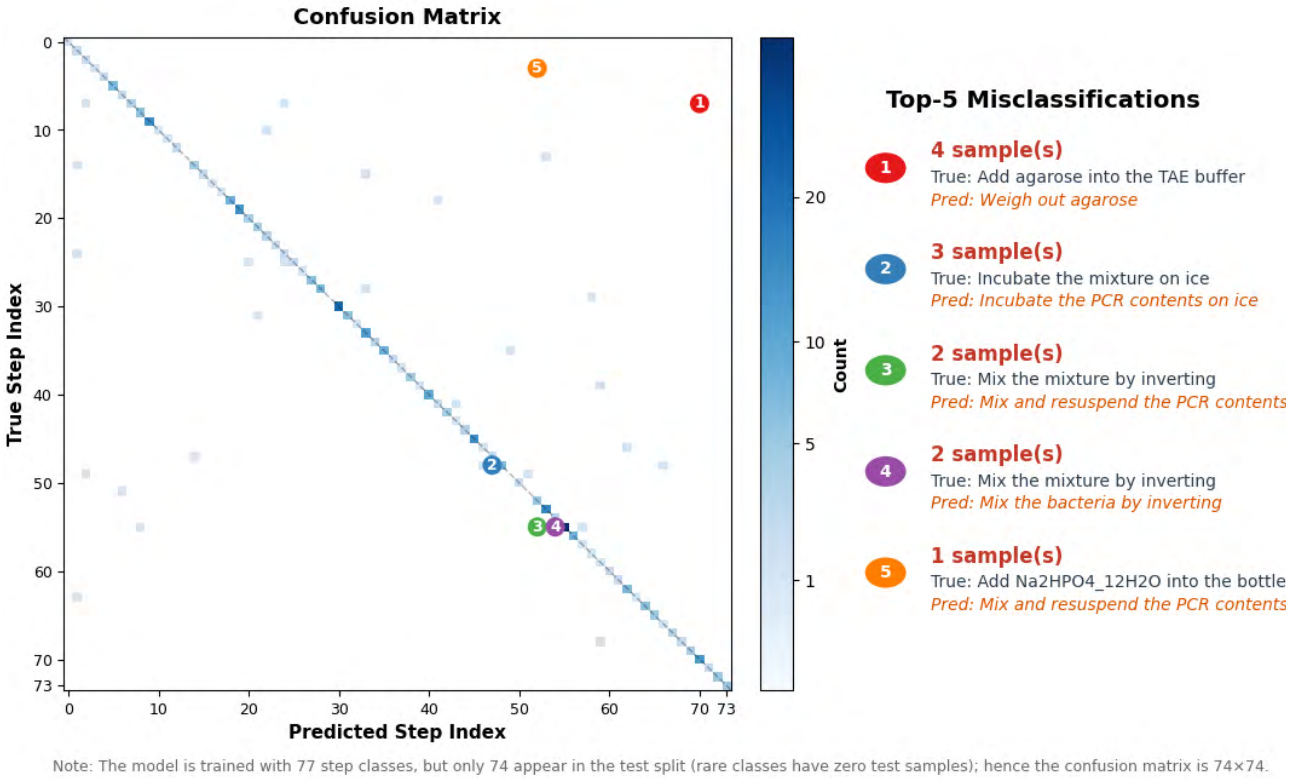

Figure 2. Step identification confusion matrix with top misclassifications.

temporally adjacent operations sharing materials (e.g., "weigh agarose" vs. "add agarose into buffer"). The high top-3 test accuracy (0.9731) indicates that misclassifications are typically near-misses rather than gross failures. Accordingly, we propagate the top-3 step hypotheses to the downstream modules to better handle local visual ambiguity during protocol alignment and step-conditioned safety monitoring.

Table 1: Experimental step identification performance.

	Item	Value
Model	Backbone	MViTv2
Dataset	Step classes	77
	Train / Val / Test samples	2232 / 1116 / 372
Metrics	Best Val Acc (epoch)	0.8952 (36)
	Test Top-1 Accuracy	0.8952
	Test Top-3 Accuracy	0.9731
	Test Weighted-F1	0.8905

3.1.2 Experiment linkage

Using the 1-second step predictions produced by $StepID(\cdot)$, we infer the experiment identity by aligning the predicted step sequence $\hat{S}$ with each candidate protocol P_e via the soft matching score in equation (10)-(11). In practice, we approximate $p_t(\cdot)$ using each clip's top-3 step predictions to improve robustness to locally ambiguous top-1 outputs. We evaluate linkage over 12 experiment classes. Because some videos contain only a few distinct step types (thus providing weak protocol evidence), we report results on subsets defined by the number of unique top-1 step labels $\boldsymbol{n}$ observed across clips. Note that the subset split is computed from the per-clip top-1 labels to count the number of unique predicted steps $\boldsymbol{n}$, whereas the protocol alignment score is still computed by aggregating the top-3 predictions for each clip.

Table 2. Experiment linkage performance under different procedural diversity.

Subset	Videos	Top-1 Acc	Top-3 Acc	Weighted-F1
$n \geq 1$	260	0.8346	0.9654	0.8272
$n \geq 2$	127	0.9213	0.9843	0.9156
$n \geq 3$	14	1.0000	1.0000	1.0000

As shown in **Table 2**, linkage accuracy improves monotonically with procedural diversity: when n is small, multiple experiments that share common subroutines can obtain similar alignment scores; when n increases, the

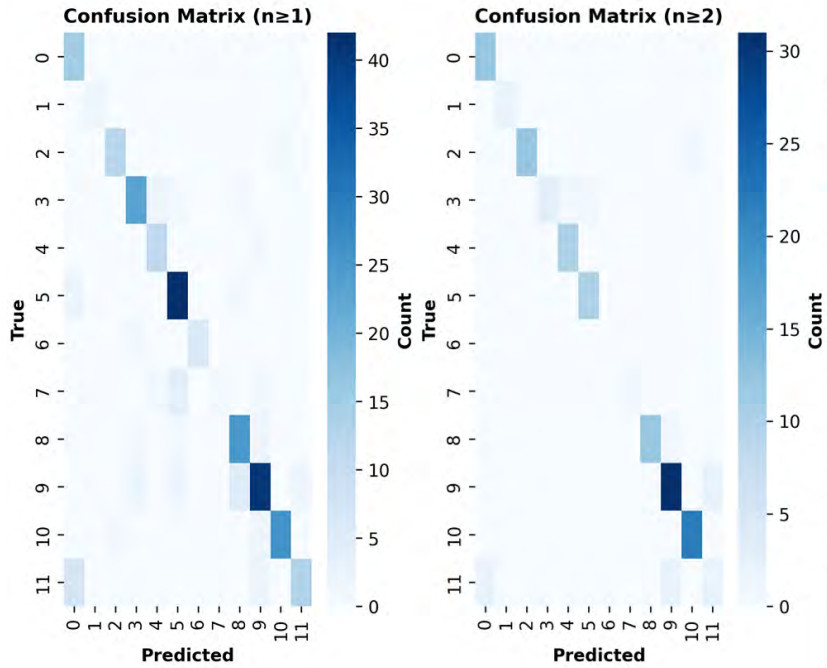

Index Mapping

ID	Experiment Name
0	E.coli transformation
1	ELISA
2	Glue production
3	Glue recycling
4	LB_solid_medium preparation
5	PCR
6	PCR fragment recycling
7	PCR products digest using DpnI enzyme
8	Plasmid Maxiprep
9	Plasmid Maxiprep II
10	SDCAA_medium_solution preparation
11	Yeast transformation

Figure 3. Experiment linkage confusion matrices under different procedural diversity (n≥1 vs. n≥2).

protocol signature becomes more discriminative and top-1 accuracy rises while top-3 accuracy nearly saturates.

Figure 3 further illustrates this trend with confusion matrices for n≥1 and n≥2: the n≥2 subset is more diagonal with fewer off-diagonal confusions. The accompanying index mapping links matrix indices to experiment names for inspecting confused pairs.

In practice, ambiguity can often be detected from the margin between the best and second-best alignment scores. When the margin is small, we can defer downstream safety monitoring to a top-k candidate set, consistent with the strong top-3 accuracy of experiment linkage. We leave this ambiguity-aware inference to future work. Further improvements may also come from stronger recognition of protocol-distinctive steps and lightweight temporal smoothing to reduce spurious step switches in $\hat{S}$.

3.2 Multimodal Safety Monitoring

After scenario recognition provides the procedural context $(\hat{e}, \widetilde{s_k})$, we perform procedure-aware safety monitoring on each segment z_k. The monitoring prompt is constructed by integrating four components: (i) the video segment z_k; (ii) experiment and step metadata $(\hat{e}, \widetilde{s_k})$; (iii) the retrieved step-level safety entry r_k (including required PPE, cautions, and forbidden actions); and (iv) structured equipment evidence g_k extracted via a constrained "equipment presence and interaction" call. The final VLM (Qwen3-VL) outputs a strict JSON report with a list of issues (category, description, violated rule, and supporting evidence). We compare against a video-only baseline that prompts the same VLM using only the video segment, without retrieval of step safety entry and without equipment evidence.

3.2.1 Quantitative Results

We evaluate $N = 80$ monitored segments. For each monitored segment z_k, an expert provides a reference issue set $\mathcal{R}_k = \{I_{k,i}\}_{i=1}^{|\mathcal{R}_k|}$. The model outputs a set of issues O_k. The expert aligns model outputs to reference issues at a consistent granularity. For each reference issue $I_{k,i}$, we assign a coverage score $c_{k,i} \in \{1.0,\ 0.5,\ 0.0\}$, where 1.0 denotes full coverage, 0.5 partial coverage, and 0.0 not covered. Any model-reported issue that cannot be aligned to any reference issue is treated as an additional issue. Let $A_k \subseteq O_k$ denote the set of such additional issues. The expert verifies each additional issue against the video and assigns a correctness score $v_{k,j} \in \{1.0,\ 0.5,\ 0.0\}, j = 1, \ldots, |A_k|$, where 1.0 means verified correct, 0.5 partially correct/plausible, and 0.0 incorrect. We compute a single segment score by averaging over reference issues + additional issues:

$$S_k = \frac{\sum_{i=1}^{|\mathcal{R}_k|} c_{k,i} + \sum_{j=1}^{|A_k|} v_{k,j}}{|\mathcal{R}_k| + |A_k|} \tag{12}$$

This penalizes over-generation: adding many incorrect extra issues increases the denominator while contributing little to the numerator. Finally, the **Mean Score** is the average segment score over all monitored

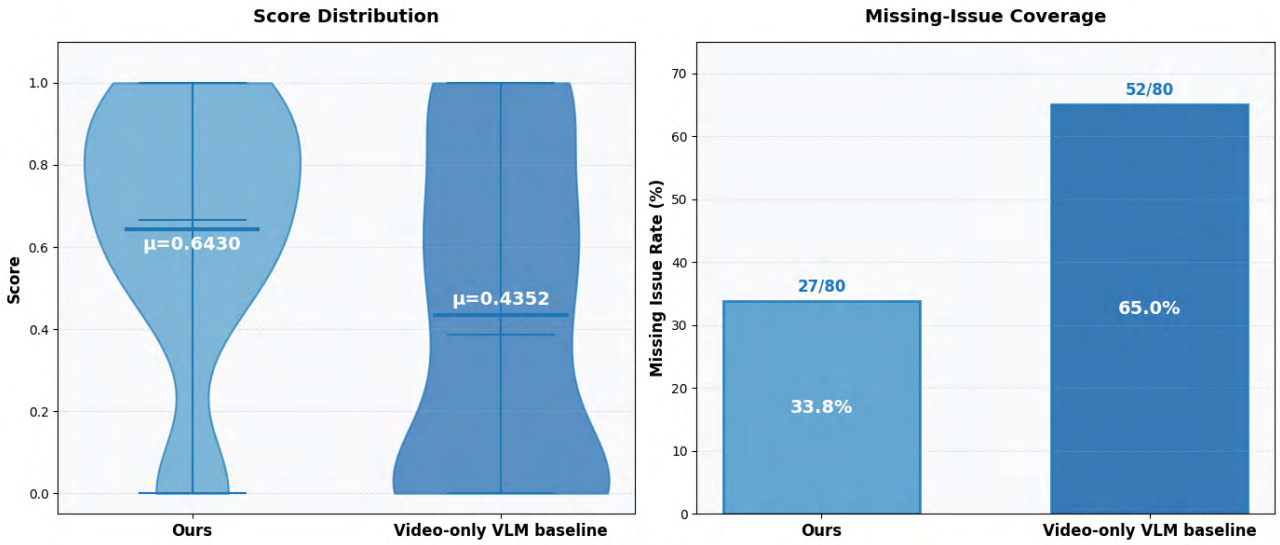

Figure 4. Score distribution and missing rate: ours vs. video-only VLM baseline.

segments:

$$MeanScore = \frac{1}{N}\sum_{k=1}^{N} S_k \tag{13}$$

We additionally report a segment-level **Missing Rate**. A segment is marked as missing if the reference issue set $\mathcal{R}_k$ is not fully covered, i.e., there exists at least one reference issue with $c_{k,i} < 1$. Let $m_k \in \{0,1\}$ denote this missing indicator.

$$m_k = [1 - \min_i c_{k,i}] \tag{14}$$

$$MissingRate = \frac{1}{N}\sum_{k=1}^{N} m_k \tag{15}$$

Table 3 summarizes the quantitative comparison. Procedure-aware monitoring substantially improves the mean score and reduces the missing rate relative to the video-only baseline.

Table 3. Multimodal safety monitoring performance.

Method	Mean Score ↑	Missing Rate ↓
Video-only VLM baseline	0.4352	65.00%
Ours	0.6430	33.75%

Figure 4 visualizes these gains. The left panel shows that our method shifts the per-segment score distribution upward, indicating more valid and better-supported issues across segments. The right panel shows improved coverage: our method reduces missing segments to 27/80 (33.75%), compared with 52/80 (65.0%) for the video-only baseline. Overall, injecting step-specific constraints and providing explicit equipment interaction evidence leads to more evidence-supported, actionable safety judgments while substantially reducing missed violations.

3.2.2 Qualitative Analysis

We present a representative case to illustrate how procedural anchoring changes VLM behavior in practice. The example is from Plasmid Maxiprep, step *add_plasmid_into_EP_tube*, using the 0-65s segment of *0802_0_21.mp4*. In this step, the operator pipettes plasmid into EP tubes, and the retrieved step-level safety entry emphasizes PPE compliance and contamination/traceability control (e.g., minimizing open-tube exposure and ensuring proper tube identification before addition).

As shown in **Figure 5**, manual inspection identifies five issues: ① a nearby observer lacks basic PPE (no lab coat or gloves), ② the main operator is missing safety goggles, ③ the EP tubes show no clearly visible labeling before plasmid addition (a traceability risk), ④ EP tubes are left open on the bench during pipetting (increasing contamination/aerosol exposure risk), and ⑤ an experimentalist briefly uses a mobile phone in the laboratory. Our method identifies four issues (①-④). Notably, it goes beyond the most salient PPE cues and also flags step-relevant risks (③-④) that are explicitly emphasized in the retrieved step-level safety entry for this operation, such as contamination control and traceability. In contrast, the video-only VLM baseline identifies three issues, focusing mainly on the more visually obvious PPE-related problems (①-②), and it additionally reports a generic proximity-related concern (standing too close to the work area) that is not part of the manual reference in this example. Both methods miss the phone-use issue ⑤, plausibly because it appears in only a few frames and can be easily overlooked in segment-level monitoring.

Ours	Manual	Video-only VLM baseline
The observer (female) is not wearing a lab coat or gloves.	① An experimentalist is not wearing an experimental suit and is not wearing gloves.	The observer (female) is not wearing a lab coat or gloves, exposing her to potential chemical or biological hazards.
The experimenter is not wearing safety goggles.	② The experimenters are not wearing protective glasses.	The lab technician appears to be wearing gloves but may not be wearing safety goggles, which are essential during pipetting to prevent accidental splashes.
Potential cross-contamination risk due to lack of visible labeling on EP tubes before addition of plasmid.	③ The EP tube does not appear to be marked.	/
EP tubes are left open on the bench during pipetting, increasing risk of contamination and aerosol exposure.	④ During the pipetting process, the EP tube remains open for an extended period of time.	/
/	⑤ An experimentalist takes out a mobile phone in the laboratory.	/
/	/	The observer is standing too close to the active work area, potentially causing distraction or contamination risk.

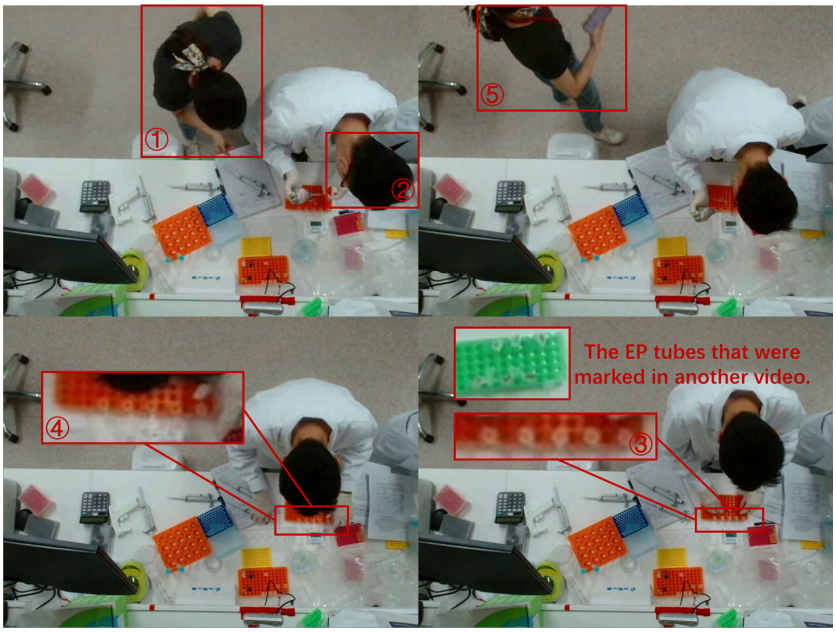

Figure 5. Qualitative comparison on plasmid maxiprep. Top: issue lists produced by our method, a manual baseline, and a video-only VLM baseline. Bottom: representative frames with visual evidence; numbered red boxes correspond to the reported issues.

Overall, this case illustrates that procedural anchoring leads to richer and more step-aware safety reports by anchoring generation to step-specific constraints known for the current operation, thereby surfacing context-dependent risks that are less likely to be produced by video-only prompting.

4. CONCLUSION

This paper presents a multimodal framework that integrates procedural texts with visual perception for procedure-aware laboratory safety monitoring. Motivated by the experiment- and step-dependent nature of laboratory safety requirements, we decompose the task into two stages: scenario recognition and step-conditioned monitoring. In the first stage, we localize procedural context through clip-level step identification and protocol-based experiment linkage, enabling the system to infer both the most likely experiment and the current step in long, untrimmed laboratory videos. In the second stage, we retrieve the corresponding step-level safety analysis entry, extract a step-specific equipment list, extract equipment presence and interaction evidence via an equipment database, and prompt a video-capable VLM to produce strict JSON-formatted safety issues supported by step-specific constraints and video evidence.

Experiments on a protocol-annotated molecular biology laboratory video dataset show that procedural anchoring improves both faithfulness and coverage. Compared with a video-only VLM baseline, our approach increases the mean score from 0.4352 to 0.6430 and reduces the missing rate from 65.00% to 33.75%. Qualitative results further indicate that conditioning generation on step-level constraints yields richer, more step-aware findings—not only salient PPE violations but also step-dependent risks such as traceability issues and open-tube exposure that the baseline often overlooks.

Despite these gains, several limitations remain. First, fleeting violations that occur in only a few frames can be missed, suggesting the need for finer temporal sensitivity or event-triggered mechanisms. Second, the framework depends on the coverage and quality of the safety report library and the equipment database; incomplete or mismatched entries may reduce retrieval effectiveness and evidence-alignment accuracy. Third, although segment-level processing reduces the number of VLM calls, cost and latency still pose practical challenges for real-world deployment, and privacy considerations must be addressed when monitoring human behavior in laboratory settings.

Future work will focus on robustness and deployment readiness. Promising directions include: (i) stronger temporal modeling and smoothing to reduce spurious step switches and improve detection of fleeting events; (ii) uncertainty-aware retrieval and monitoring that can defer judgment or expand to top-k candidate experiments/steps when linkage margins are small; (iii) scalable approaches to expand and maintain safety-report and equipment knowledge bases, potentially via semi-automated extraction from institutional SOPs and safety documents; and (iv) broader evaluations across more diverse laboratory environments and experiment types.

Overall, our framework provides a practical blueprint for context-sensitive, procedure-aware safety monitoring, moving beyond appearance-driven checks toward step-specific compliance reasoning supported by both text and video evidence.

REFERENCES

1. Ménard AD, Trant JF. A review and critique of academic lab safety research. Nat. Chem. 12:17-25 (2019). https://doi.org/10.1038/s41557-019-0375-x
2. Hill RH Jr. Recognizing and understanding hazards — the key first step to safety. J. Chem. Health Saf. 26:5-10 (2019). https://doi.org/10.1016/j.jchas.2018.11.005
3. Vukicevic AM, Petrovic M, Milosevic P, Peulic A, Jovanovic K, Novakovic A. A systematic review of computer vision-based personal protective equipment compliance in industry practice: advancements, challenges and future directions. Artif Intell Rev 57: (2024). https://doi.org/10.1007/s10462-024-10978-x
4. Luo H, Liu W, Xu P, Zhang L, Li L. Recognition algorithm for laboratory protective equipment based on improved yolov7. Heliyon 10:e36264 (2024). https://doi.org/10.1016/j.heliyon.2024.e36264
5. Tang Y, Ding D, Rao Y, Zheng Y, Zhang D, Zhao L, Lu J, Zhou J. COIN: a large-scale dataset for comprehensive instructional video analysis. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) :1207-1216 (2019). https://doi.org/10.1109/cvpr.2019.00130
6. Miech A, Zhukov D, Alayrac JB, Tapaswi M, Laptev I, Sivic J. Howto100m: learning a text-video embedding by watching hundred million narrated video clips. 2019 IEEE/CVF International Conference on Computer Vision (ICCV) :2630-2640 (2019). https://doi.org/10.1109/iccv.2019.00272
7. Zhukov D, Alayrac JB, Cinbis RG, Fouhey D, Laptev I, Sivic J. Cross-task weakly supervised learning from instructional videos. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) :3532-3540 (2019). https://doi.org/10.1109/cvpr.2019.00365
8. Afouras T, Mavroudi E, Nagarajan T, Torresani L, Wang H. Ht-step: aligning instructional articles with how-to videos. Advances in Neural Information Processing Systems 36 :50310-50326 (2023). https://doi.org/10.52202/075280-2188
9. Zhou H, Martín-Martín R, Kapadia M, Savarese S, Niebles JC. Procedure-aware pretraining for instructional video understanding. 2023 IEEE/CVF Conference on Computer Vision and Pattern

Recognition (CVPR) :10727-10738 (2023). https://doi.org/10.1109/cvpr52729.2023.01033
10. Cui J, Gong Z, Huang S, Jia B, Ma J, Zheng Z, Zhu Y. Probio: a protocol-guided multimodal dataset for molecular biology lab. Advances in Neural Information Processing Systems 36 :41543-41571 (2023). https://doi.org/10.52202/075280-1799
11. Zhang J, Huang J, Jin S, Lu S. Vision-language models for vision tasks: a survey. IEEE Trans. Pattern Anal. Mach. Intell. 46:5625-5644 (2024). https://doi.org/10.1109/tpami.2024.3369699
12. Bai S, Cai Y, Chen R, Chen K, Chen X, Cheng Z, et al. Qwen3-VL technical report. arXiv preprint arXiv:2511.21631 (2025)
13. Huang L, Yu W, Ma W, Zhong W, Feng Z, Wang H, Chen Q, Peng W, Feng X, Qin B, Liu T. A survey on hallucination in large language models: principles, taxonomy, challenges, and open questions. ACM Trans. Inf. Syst. 43:1-55 (2025). https://doi.org/10.1145/3703155
14. Petroni F, Piktus A, Fan A, Lewis P, Yazdani M, De Cao N, Thorne J, Jernite Y, Karpukhin V, Maillard J, Plachouras V, Rocktäschel T, Riedel S. KILT: a benchmark for knowledge intensive language tasks. Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies :2523-2544 (2021). https://doi.org/10.18653/v1/2021.naacl-main.200
15. Abootorabi MM, Zobeiri A, Dehghani M, Mohammadkhani MA, Mohammadi B, Ghahroodi O, Soleymani Baghshah M, Asgari E. Ask in any modality: A comprehensive survey on multimodal retrieval-augmented generation. arXiv preprint arXiv:2502.08826 (2025)
16. Guo KX, Wong PKY, Cheng JCP, Chan CF, Leung PH, Tao X. Enhancing visual-llm for construction site safety compliance via prompt engineering and bi-stage retrieval-augmented generation. Automation in Construction 179:106490 (2025). https://doi.org/10.1016/j.autcon.2025.106490
17. Li Y, Wu CY, Fan H, Mangalam K, Xiong B, Malik J, Feichtenhofer C. Mvitv2: improved multiscale vision transformers for classification and detection. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) :4794-4804 (2022). https://doi.org/10.1109/cvpr52688.2022.00476

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimal Stopping of Batch Processes with Stochastic Dynamics – A Study of When to Act under Uncertainty

Rafif S. Ramadhan[ab], Luca Grebe[c], Maximilian Maschmeier[c], Johannes Pastoors[c], Eike Cramer[bd]*

[a] Davidson School of Chemical Engineering, Purdue University, West Lafayette, IN 47907, USA
[b] RWTH Aachen University, Process Systems Engineering (AVT.SVT), Aachen 52074, Germany
[c] RWTH Aachen University, Biochemical Process Engineering (AVT.BioVT), Aachen 52074, Germany
[d] University College London, Sargent Centre for Process Systems Engineering, Department of Chemical Engineering, London WC1E 7JE, United Kingdom
* Corresponding Author: e.cramer@ucl.ac.uk

ABSTRACT

Mathematical models in process systems engineering (PSE) are widely used to support decision-making in design and operation, but they are mostly limited to deterministic models. For biochemical systems, the biological variability can give rise to stochastic dynamics. This work addresses the question of when to act in such processes, as the stochastic dynamics affect the timing of important events. We consider the case of batch production of malic acid using *Ustilago trichophora*. The goal is to predict when the substrate concentration falls below a predefined threshold. We extend an existing deterministic model of the process to a stochastic differential equation (SDE) formulation by introducing a Monod-like noise term. Simulations of the SDE model reveal a distribution of substrate depletion times and a deviation between the mean of the stochastic trajectory and the deterministic solution due to nonlinear effects. To determine optimal intervention times under uncertainty, we formulate a finite-horizon optimal stopping problem and solve it using the Longstaff–Schwartz algorithm, also known as the Least-Squares Monte Carlo (LSMC). The resulting distribution of optimal stopping times from the LSMC algorithm is shown to closely match the actual first threshold hitting times obtained from *a posteriori* analysis of the stochastic simulations, as confirmed by a two-sample Kolmogorov–Smirnov test. The results demonstrate that optimal stopping provides a framework for decision-making in stochastic biochemical processes, enabling risk-aware operational strategies beyond deterministic optimisation.

Keywords: Stochastic differential equations (SDEs), optimal stopping, decision-making under uncertainty

INTRODUCTION

Biochemical processes play an increasingly important role in the transition toward more sustainable chemical production. By exploiting the metabolic activity of microorganisms, bioprocesses enable the conversion of renewable feedstocks into fuels and platform chemicals that are essential for downstream manufacturing [1]. As these processes move from laboratory to industrial scale, the need for reliable modeling and decision-support tools becomes increasingly critical.

In process systems engineering, mathematical models are routinely used to support design, optimization, and operational decision-making [2]. For biochemical systems, such models are traditionally formulated as ordinary differential equations (ODEs). While such deterministic models are valuable for capturing average system behavior, biological processes are inherently variable due to intracellular heterogeneity, environmental fluctuations, and nonlinear reaction kinetics [3]. As a result, deterministic predictions may fail to adequately represent the range of possible process trajectories, particularly when decisions depend on threshold-based events such as substrate depletion or intervention timing.

Although uncertainty is often addressed in process control frameworks, stochastic effects are rarely incorporated explicitly into optimization and planning problems for biochemical batch processes [2]. Existing studies on biological noise primarily focus on stochastic simulation and analysis [3], while the implications of

stochastic dynamics for decision-making and optimization remain underexplored. In particular, there is a lack of systematic approaches for determining intervention times in stochastic bioprocesses that account for the full distribution of possible system trajectories rather than a single deterministic prediction.

Stochastic decision-making problems have been extensively studied in financial mathematics, where uncertainty is intrinsic to system dynamics [4, 5]. Here, optimal stopping theory provides a framework for determining when to take an action in a stochastic environment, e.g., to maximize an expected reward. The Longstaff–Schwartz algorithm, also known as Least-Squares Monte Carlo (LSMC), enables the solution of optimal stopping problems using simulated trajectories and regression-based approximation of conditional expectations [6].

In this work, we formulate a stochastic model of a biochemical batch reactor using stochastic differential equations (SDEs) and embed this model within an optimal stopping framework. The stochastic formulation explicitly captures process variability through state-dependent noise, while the optimal stopping problem is defined in terms of substrate depletion. The Longstaff–Schwartz algorithm is then applied to learn an optimal intervention time directly from simulated stochastic trajectories.

In this work, we formulate a stochastic model of a batch bioreactor to produce malic acid via *Ustilago trichophora* [7]. We extend an existing ODE model to an SDE formulation to represent the stochastic process. Based on the stochastic model, we analyse an optimal stopping problem using the Longstaff-Schwartz algorithm. The results from the algorithm show that the stochastic system exhibits behaviours that lead to different intervention strategies compared to the deterministic system. The findings of this work highlight the importance of stochastic modelling for decision-making problems.

STOCHASTIC PROCESS MODEL

This work considers the batch production of the platform chemical malic acid using *U. trichophora* as biomass. For this process, the substrates sucrose, fructose, glucose, and ammonium (NH_4^+), as the source of nitrogen, are supplied. During the initial growth phase, sucrose is enzymatically hydrolysed to the monosaccharides glucose and fructose and utilised mainly for biomass formation. As nitrogen becomes depleted, the culture shifts from growth to production. Under these conditions, malic acid is formed as the predominant product. The deterministic ODE model is taken from Maschmeier [7].

$$\frac{d}{dt}\begin{bmatrix} X_a \\ X_i \\ Suc \\ FruGlu \\ N \\ P \end{bmatrix} = \begin{bmatrix} \mu_a \\ \mu_i \\ -q_{split} \\ q_{split} - \frac{\mu_a}{Y_{XS,a}} - \frac{\mu_i}{Y_{XS,i}} - \frac{q_P}{Y_{PS}} \\ -\frac{\mu_a}{Y_{XN}} \\ q_P \end{bmatrix} X_a \tag{1}$$

Here, the specific growth and product formation rates are:

$$\mu_a = \mu_{a,max} \frac{FruGlu}{FruGlu+K_{FG}} \frac{N}{N+K_N} \tag{2}$$

$$\mu_i = \mu_{i,max} \frac{FruGlu}{FruGlu+K_{FG,2}} \frac{K_{IN}}{N+K_{IN}} \left(1 - e^{\frac{X_i/X_a - \phi_X}{\chi_{acc}}}\right) \tag{3}$$

$$q_{split} = q_{split,max} \frac{Suc}{Suc+K_S} \tag{4}$$

$$q_P = q_{P,max} \frac{FruGlu}{FruGlu+K_{PFG}} \frac{K_{IN}}{N+K_{IN}} \frac{K_{IP}}{K_{IP}+\frac{w_N N}{X_a+X_i}} \tag{5}$$

For the definition of the variables in the deterministic model, see Table 1.

Table 1: State variables in the model.

Symbol	Description
X_a	Active biomass concentration
X_i	Inactive biomass concentration
Suc	Sucrose concentration
$FruGlu$	Sum of Fructose and Glucose concentrations
N	Nitrogen concentration
P	Malic acid concentration

We extend the deterministic model to describe the stochastic dynamics by adding a differential noise term to the ODE, turning the model equation into an SDE. In general, SDEs have the form

$$dX_t = \mu(X_t)\,dt + \sigma(X_t)\,dB_t, \tag{6}$$

where $\mu(X_t)$ is the drift term, $\sigma(X_t)$ the diffusion term, and B_t the Brownian motion (standard Wiener process) [8]. The deterministic model Eq. 1 serves as the drift term. In the following, we will define the noise term.

In biochemical processes, biomass activities are dependent on the availability of substrates. Therefore, the random fluctuations are modelled to follow the same phenomenon. The biomass activities exhibit minimal fluctuations when the substrate is depleted and reach a saturation point when the substrate is abundant. To capture the substrate dependence, this work proposes modelling the noise term σ_i in component *i* using a Monod-like equation. The Monod model is chosen because it relates microbial activities to substrate availability [9]. We note that Eq. 7 does not itself constitute the stochastic noise term. Rather, it defines the substrate-dependent

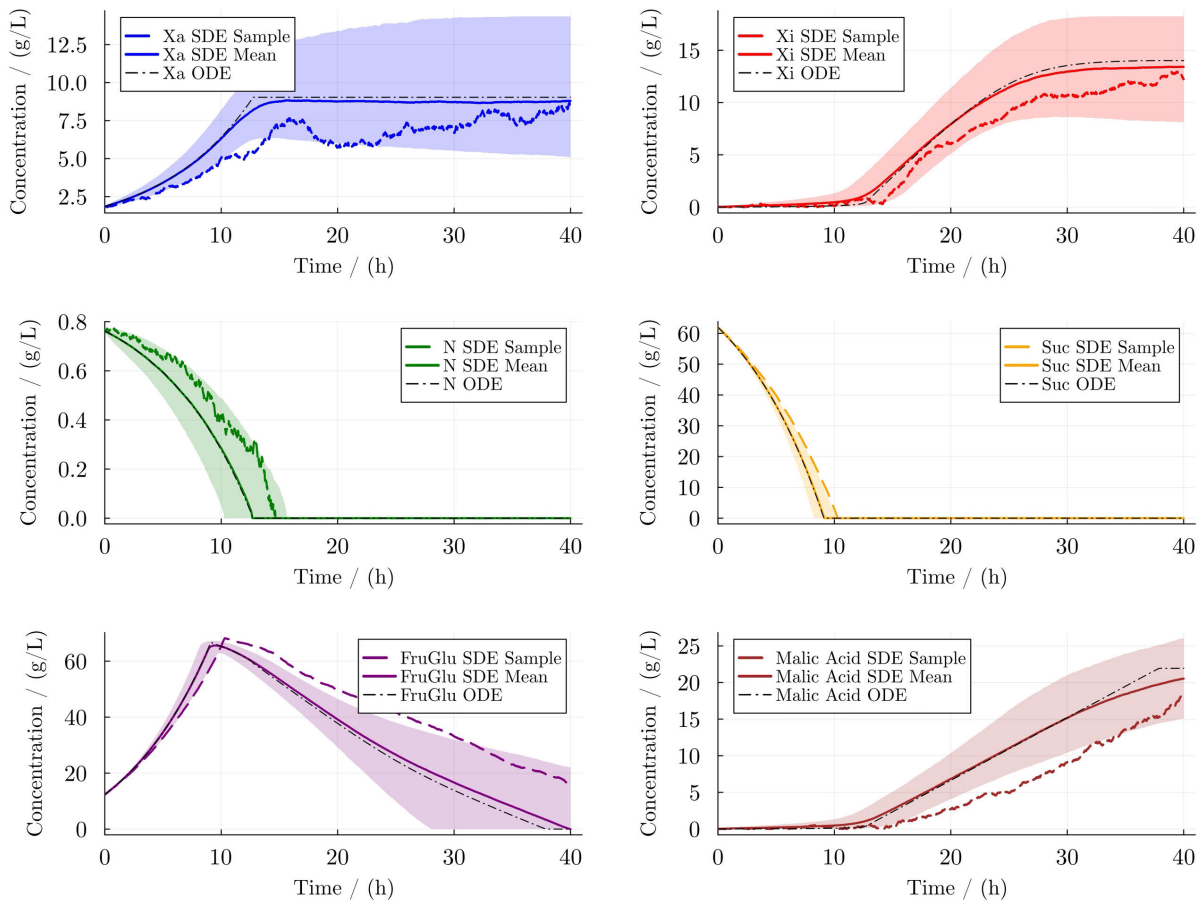

Figure 1. Sample stochastic scenario (dashed), deterministic profile (dash-dot), and SDE summary of several simulated scenarios (Set 1) consisting of the averaged concentration profile (solid line) and the 5th to 95th percentile of the scenario distribution (shaded region). From top to bottom, left to right: active biomass, inactive biomass, nitrogen, sucrose, fructose and glucose, malic acid.

diffusion coefficient σ_i, which modulates the stochastic fluctuations in a physically meaningful way. The noise itself enters the SDE through the Wiener increment dB_t.

$$\sigma_i = \sigma_{i,max} \frac{S}{K_{\sigma_i}+S} \tag{7}$$

Here, $\sigma_{i,max}$ is the maximum noise intensity and K_{σ_i} is the associated half-saturation constant. The variable S represents the substrate, which can be taken up by either $Suc, FruGlu,$ or N, depending on which equation the noise term is in. The full noise equation is given in Eq. 8. Each row represents the diffusion coefficient of the corresponding state variable. Thus, together with Equation (1), the mole balance of each species is given as a stochastic differential equation in vector form, with the drift term given by Equation (1) and the diffusion term by Equation (8). The half-saturation constant in the noise term K_{σ_i} is set to the value of the reaction kinetics.

This work aims to discuss the optimal stopping algorithm, and no measurements of the variance parameters are available. Instead, the parameter values in the noise terms are selected arbitrarily to be $\sigma_{i,max} = 0.05\ 1/\sqrt{\mathrm{h}}$ for all components. The validation of this proposed noise model and its parameters based on experimental data falls outside the scope of this work and, therefore, can be conducted as a separate study, where other noise models can also be considered.

$$\sigma(X_t) = \begin{bmatrix} \sigma_{X_a,max} \frac{FruGlu}{K_{FG}+FruGlu} \\ \sigma_{X_i,max} \frac{FruGlu}{K_{FG,2}+FruGlu} \\ \sigma_{Suc,max} \frac{Suc}{K_S+Suc} \\ \sigma_{FG,max} \frac{FruGlu}{K_{FG}+FruGlu} \\ \sigma_{N,max} \frac{N}{K_N+N} \\ \sigma_{P,max} \frac{FruGlu}{K_{PFG}+FruGlu} \end{bmatrix} X_a \tag{8}$$

In this work, we simulated two sets of stochastic sample paths, which will be referred to as Sets 1 and 2 in the following. Each set contains five thousand scenarios. Figure 1 shows a sample stochastic scenario from Set 1

and the summary of Set 1 obtained via the Euler-Maruyama scheme. The Euler-Maruyama method is the stochastic extension of Euler's method to numerically solve SDEs [8]. The ensemble summary consists of an averaged profile and the 95% confidence interval at every time point. The initial values are provided in Table 2, and the model parameters in Table 3.

Table 2: Initial values

Variable	Concentration (g/L)
X_a	1.85
X_i	0
Suc	61.96
$FruGlu$	12.3
N	0.762
P	0

Figure 1 shows that the average concentration profile of the SDE does not align with the deterministic ODE model (shown in dashed lines) and that the concentrations from the stochastic SDE model are distributed over a range of values. This observation suggests that different results may be obtained when the models are applied to further analysis, such as optimisation or process control. The difference between the stochastic scenarios and the deterministic model can be attributed to the interplay between the model's nonlinearities and the noise. This interplay can be described using Jensen's inequality, which states that the mean of a nonlinear function of a random variable is not equal to the function of the expected value itself [8, 10].

OPTIMAL STOPPING

Optimal stopping is a problem commonly encountered in derivatives finance, specifically in the context of American Options. An optimal stopping problem poses the question: "When should one exercise an American Option to maximise the payoff?" For more background information on derivatives finance, the reader is referred to [4, 5].

Least-Squares Monte Carlo

LSMC was originally developed for the valuation of American Options [6]**.** However, the algorithm itself only requires simulated trajectories and a reward function. Thus, the algorithm is model-agnostic, and the method is applicable to problems outside of finance if the underlying model is a stochastic process. This work investigates the application of LSMC to solve an optimal stopping problem, where the underlying stochastic process is the stochastic reactor model.

Consider the SDE given in Eq. 6. The objective of an optimal stopping problem is to determine a time τ that maximises the expected value of a certain reward function G_t. The reward function encodes the performance criterion associated with stopping the process at time t.

$$V^* = \sup_{\tau \in [0,T]} \mathbb{E}[G_\tau] \tag{9}$$

Table 3: Model parameters

Symbol	Description	Value	Unit
$\mu_{a,max}$	Max. growth rate of active biomass	0.125	1/h
$\mu_{i,max}$	Max. growth rate of inactive biomass	0.125	1/h
$q_{split,max}$	Max. split rate	1.985	1/h
$q_{P,max}$	Max. product formation rate	0.095	1/h
K_S	Half-saturation constant of sucrose	0.00321	g/L
K_{FG}	Half-saturation constant of active biomass	0.147	g/L
$K_{FG,2}$	Half-saturation constant of inactive biomass	3.277	g/L
K_N	Half-saturation constant of ammonia	3.8E-5	g/L
K_{PFG}	Half-saturation constant of product	0.0175	g/L
K_{IN}	Nitrogen inhibition constant	0.0147	g/L
K_{IP}	Product inhibition constant	0.147	g/L
w_N	Internal nitrogen mass fraction	0.08	-
χ_{acc}	Acceleration factor	0.3	-
ϕ_X	Max. ratio of inactive to active biomass	1.56	-
$Y_{XS,a}$	Yield factor FruGlu to active biomass	0.531	-
$Y_{XS,i}$	Yield factor FruGlu to inactive biomass	0.799	-
Y_{PS}	Yield factor FruGlu to product	0.508	-
Y_{XN}	Yield factor ammonia to biomass	9.428	-

This formulation constitutes a finite-horizon optimal stopping problem, i.e., the process must be stopped before or at time T. Furthermore, the stopping decision must be adapted to the filtration generated by the stochastic process $\{\mathcal{F}_t\}$. In other words, the stopping decision must be made based on the observed process up to time t. The first step of the algorithm is to approximate the continuous time horizon on a discrete time grid $\{t_n\}_{n=0}^N$. Here, G_n is the reward from stopping at time t_n. The algorithm proceeds with backward induction to determine the optimal stopping time. The backward induction corresponds to the dynamic programming recursion,

Eq. 10, with the terminal condition $S_N = G_N$.

$$S_n = \max\{G_n, \mathbb{E}[S_{n+1}|\mathcal{F}_n]\} \tag{10}$$

At each time step t_n, the decision is obtained by comparing the immediate reward G_n with the continuation value, which is defined as the conditional expectation of the future rewards given the current information $\mathbb{E}[S_{n+1}|\mathcal{F}_n]$. The optimal stopping time is thus the earliest time at which stopping is preferable to continuation, i.e., $G_n > \mathbb{E}[S_{n+1}|\mathcal{F}_n]$.

The main component of the LSMC algorithm is to approximate the continuation value, which is essential in the backward induction step. Longstaff and Schwartz proposed approximating the continuation value via simulation and regression [6]. First, a number of sample paths of the stochastic process are generated via Monte Carlo simulation. At each time step, the continuation value is approximated by a regression of realised future rewards onto a set of basis functions ϕ_i:

$$\mathbb{E}[\, G_{n+1} \mid \mathcal{F}_n \,] \approx \sum_{i=0}^{K} \beta_i \, \phi_i(X_n) \tag{11}$$

Here, the regression coefficients β_i are obtained via least-squares minimisation over simulated paths where stopping is admissible. The original method uses the first three weighted Laguerre polynomials. However, numerical tests indicate that the choice of basis polynomials, e.g., unweighted Laguerre, Hermite, Legendre, or Chebyshev, is trivial [6].

Optimal Stopping of Stochastic Batch Reactor

We now formulate a finite-horizon optimal stopping problem based on the stochastic batch reactor model. The objective is to determine the intervention time at which the substrate concentration approaches a predefined depletion threshold ϵ. We define the reward function to be

$$G_t = -(FruGlu(t) - \epsilon)^2, \tag{12}$$

which is maximised when $FruGlu(t) = \epsilon$. The depletion threshold is defined as 5% of the initial concentration of sucrose, fructose, and glucose $\epsilon = 0.05 \cdot \big(FruGlu(t_0) + Suc(t_0)\big)$. The optimal stopping problem is given by Eq. 9, and the continuation value approximation is

$$\mathbb{E}[\, G_{n+1} \mid \mathcal{F}_n \,] \approx \sum_{i=0}^{3} \beta_i \, \phi_i\big(FruGlu(t_n)\big). \tag{13}$$

In this work, we use the first four unweighted Laguerre polynomials as basis functions:

$$\phi_n(x) = \frac{e^x}{n!}\frac{d^n}{dx^n}(x^n e^{-x}) \text{ for } n \in \{0, 1, 2, 3\}. \tag{14}$$

RESULTS AND DISCUSSION

The Longstaff-Schwartz algorithm is applied to the formulated optimal stopping problem to obtain a distribution of optimal intervention time. Figure 2 shows the distribution of stopping times obtained from the LSMC algorithm along with a kernel density estimate (KDE) of the distribution. The resulting distribution is unimodal with a mean of 32.91 h and a mode at 33.46 h, which is significantly earlier than the time the deterministic model reaches the depletion threshold, which is at 35.68 h.

To assess whether the LSMC algorithm correctly identifies the relevant stopping point, the distribution of the optimal stopping times is compared to the distribution of the first hitting times. We define the first hitting time as the time when the stochastic path first crosses the depletion threshold, which is obtained from *a posteriori* analysis of the generated stochastic sample paths. The distribution of the first hitting times is shown in Figure 3. We found the average first hitting time to be 32.91 h and the mode of the distribution is 33.47 h, which matches the statistics of the optimal stopping times.

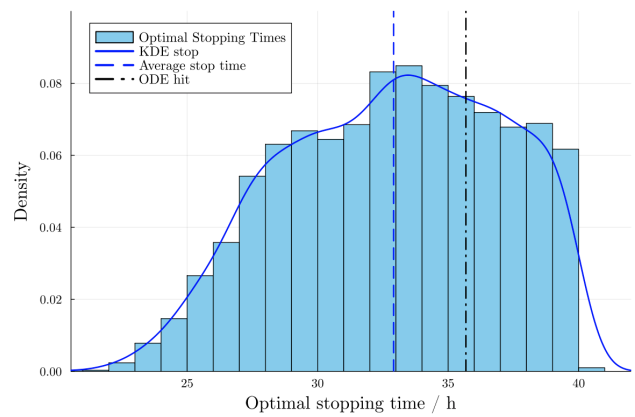

Figure 2. Distribution of the optimal stopping times from LSMC and the corresponding KDE.

The similarity between the two distributions is further quantified using a two-sample Kolmogorov-Smirnov test. The test fails to reject the null hypothesis that the two samples are drawn from the same distribution, with a negligible maximum deviation between the empirical cumulative distribution functions obtained from the time distributions, which is shown in Figure 4. This statistical agreement supports the conclusion that the LSMC algorithm yields consistent results.

To assess the robustness of the LSMC method, the analysis was repeated on Simulation Set 2. Set 2 differs in that the generated scenarios now exhibit increased variability and multimodality. However, the average optimal stopping time of 32.96 h differs only slightly from the average time found in the first simulation set (Set 1). This similarity is shown in Figure 5. The Kolmogorov-Smirnov test again indicates statistical similarity between the distributions of the optimal stopping times and the actual first hitting times for Set 2. This finding suggests that the

LSMC algorithm can reliably identify stopping times across different stochastic realisations, considering the nonlinearities of the stochastic process.

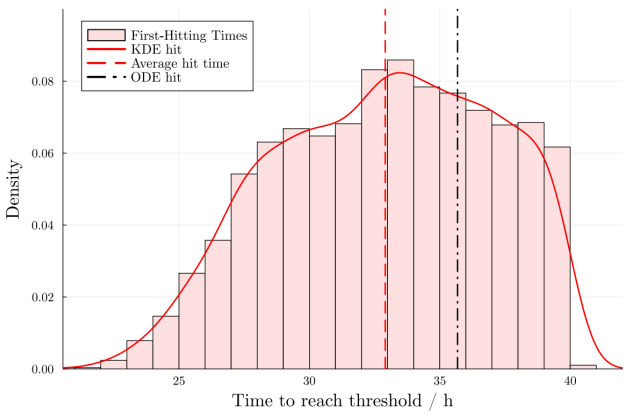

Figure 3. Distribution of the first hitting times (actual times when the substrate reaches the threshold) and the corresponding KDE.

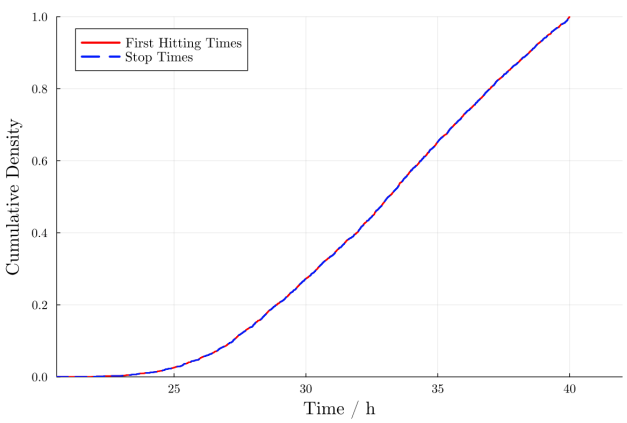

Figure 4. Cumulative distribution function (CDF) of the first hitting times and the optimal stopping times

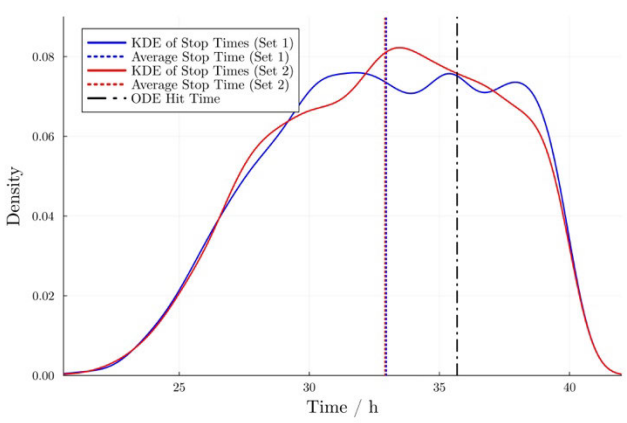

Figure 5. Optimal stopping time distributions from Simulation Sets 1 (red) and 2 (blue).

CONCLUSION

This work demonstrates that stochastic optimal stopping provides a viable and interpretable framework for intervention timing in biochemical batch processes, e.g., by adding new substrate or stopping the production process. By extending a deterministic bioprocess model to a stochastic formulation and coupling it with a regression-based optimal stopping algorithm, the study shows that uncertainty-aware decision-making can lead to systematically earlier interventions than deterministic predictions.

The results demonstrate that stochastic dynamics have a measurable impact on the timing of interventions in biochemical batch processes. Across both simulation sets, the obtained optimal stopping times consistently precede the time the deterministic model reaches the depletion threshold. This shift reflects the influence of stochastic fluctuations on process trajectories and highlights the limitations of deterministic models in decision-making problems under uncertainty.

We have also shown through the close agreement between the optimal stopping time distribution and the distribution of the first-hitting times that the LSMC algorithm accurately captures the stochastic dynamics of the system with little sensitivity across different sets of stochastic realisations. This finding provides empirical validation of the optimal stopping formulation and highlights the broader applicability of methods originating in financial mathematics to process systems engineering problems involving stochastic dynamics. The effectiveness of the optimal stopping framework depends critically on the fidelity of the underlying stochastic model. Future work should therefore prioritise validation and refinement of stochastic models using experimental data, including the exploration of species-specific or correlated stochastic perturbations.

A notable downside of the SDE model in this work is that it allows for physically unrealistic behaviour, e.g., the small increase in nitrogen concentration observed in Figure 1. Future work will focus on modelling with alternative stochastic models with an emphasis on physical constraints, specifically mass and energy balances as well as nonequilibrium reactions. Further extensions may include the application of the proposed framework to fed-batch or continuous systems, where real-time measurements enable dynamic intervention strategies. Integrating stochastic optimal stopping with techno-economic analyses also represents a promising direction, allowing intervention decisions to be evaluated directly in terms of operational cost or economic performance. Together, these developments would move stochastic optimal stopping from a methodological concept toward a practical decision-support tool for industrial bioprocess operations.

AUTHOR IDENTIFIERS

Author ORCIDs:
Ramadhan, R.S.: 0009-0007-6861-0066
Grebe, L.: 0009-0009-5532-5797
Pastoors, J.: 0000-0003-2654-0593
Cramer, E.: 0000-0002-6882-5469

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open-access article under the Creative Commons CC-BY-SA licensing terms. Credit must be given to the creator, and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

REFERENCES

1. Almquist J, Cvijovic M, Hatzimanikatis V, Nielsen J, Jirstrand M. Kinetic models in industrial biotechnology – improving cell factory performance. Metabolic Engineering 24:38-60 (2014). https://doi.org/10.1016/j.ymben.2014.03.007
2. M. D. Graham and J. B. Rawlings. Modeling and Analysis Principles for Chemical and Biological Engineers. Santa Barbara: Nob Hill Publishing, 2022. isbn:9780975937761.
3. Pischel D, Sundmacher K, Flassig RJ. Efficient simulation of intrinsic, extrinsic and external noise in biochemical systems. Bioinformatics 33:i319-i324 (2017). https://doi.org/10.1093/bioinformatics/btx253
4. Privault N. Introduction to stochastic finance with market examples. Chapman and Hall/CRC (2022). https://doi.org/10.1201/9781003298670
5. Glasserman P. Monte carlo methods in financial engineering. Springer New York (2003). https://doi.org/10.1007/978-0-387-21617-1
6. Longstaff FA, Schwartz ES. Valuing american options by simulation: a simple least-squares approach. Rev. Financ. Stud. 14:113-147 (2001). https://doi.org/10.1093/rfs/14.1.113
7. , Diaz F, Sommerfeld M, , Hovestadt G, , Latacz D, , Friedrich B, . Lessons learned from attempts at minimising CO2 emissions in process metallurgy – pyrolysed secondary raw materials, bio-coke, and hydrogen as alternative reducing agents. 12th International Conference of Molten Slags, Fluxes and Salts (MOLTEN 2024) Proceedings :1519-1532 (2024). https://doi.org/10.62053/zkqg2762
8. Thygesen UH. Stochastic differential equations for science and engineering. Chapman and Hall/CRC (2023). https://doi.org/10.1201/9781003277569
9. Monod J. THE GROWTH OF BACTERIAL CULTURES. Annu. Rev. Microbiol. 3:371-394 (1949). https://doi.org/10.1146/annurev.mi.03.100149.002103
10. A. Nobel. Convex Functions and Jensen's Inequality. Sept. 2024. url: https://nobel.web.unc.edu/teaching/stor- 555-theoretical- statistics/stor- 555-lecture-notes/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Multi-objective Experimental Design Framework Leveraging Hybrid Modelling and Gaussian Process Optimization

Michael Aku[a], Solomon Gajere Bawa[a], Ye Seol Lee[a], Federico Galvanin[a]*

[a] University College London, Department of Chemical Engineering, London, United Kingdom
* Corresponding Author: f.galvanin@ucl.ac.uk

ABSTRACT

Digitalization, artificial intelligence, and autonomous experimentation are reshaping chemical process development by enabling data-driven system identification and model-based optimization. Despite these advances, mechanistic models remain a cornerstone for predicting chemical reaction behavior and supporting optimization. However, purely mechanistic models often exhibit limited predictive accuracy when key phenomena affecting kinetics, mass and energy transfer are not fully captured. To address limitations on kinetic modelling, a hybrid modelling framework is proposed in this work that integrates a lumped power-law kinetic model with a Gaussian Process (GP) residual model to predict the reaction rate across the experimental design space while quantifying the uncertainty of the predicted rate. The hybrid model is then coupled with multi-objective Bayesian optimization (MOBO) by employing a weighted-sum approach and an upper confidence bound acquisition function to guide experimental design by simultaneously maximizing reaction rate and minimizing the associated uncertainty on prediction. The approach is tested on a case study related to catalytic methane oxidation in an autonomous microreactor system, where the proposed hybrid model performs comparably to the ground truth model, identifying the same optimal operating conditions when validated against the original experimental data, with an average prediction uncertainty of approximately 2.7%. The results highlight the potential of hybrid modelling and uncertainty-aware optimization for guiding autonomous catalytic experiments in flow systems.

Keywords: Modelling and Simulations, System Identification, Machine Learning, Bayesian Optimization

1. INTRODUCTION

System identification is reshaping autonomous experimentation by playing a key role in the use of mathematical models and input-output data to describe process behaviour alongside its underlying physics [1]. Identifying mathematical models that adequately represent a process system can be challenging [2], time-consuming and resource-intensive, and constrained by limited budgets, experimental scope, and strict project timelines [3]. Mechanistic models that describe a system can either be (i) simplified (low-fidelity) models, which include only limited mechanistic information and serve as approximate representations of an "assumed" true model or (ii) high-fidelity models that describe the system, capture observable phenomena but require extensive domain knowledge, computational efforts, and detailed satisfactory values of the model parameters.

Mechanistic models (simplified or high-fidelity) can provide valuable physical insights, but their predictive accuracy depends on how completely key phenomena such as kinetics, mass and heat transfer or catalyst deactivation are fully represented in the model [4]. Inaccurate representation of these phenomena in the model leads to a mismatch between the model's prediction and observable reaction rates.

This limitation motivates the development of a systematic hybrid modelling framework [5, 6] that retains the interpretability of mechanistic model while correcting for kinetic mismatch identified from experimental data. The

framework is consistent with recent hybrid Gaussian-process-based multi-objective Bayesian Optimization (MOBO) methods for reaction optimization under uncertainty [7]. This proposed framework offers a new direction in experimental design approaches in that it combines a lumped power-law model (simplified from a Langmuir-Hinshelwood (LH) adsorption mechanism equation) with a GP correction model (which provides data-driven correction for the portions of the reaction behaviour not described by the mechanistic model). The resulting hybrid model was then used in a MOBO procedure to guide optimal experimental design. The overall objective of the design is to identify experimental conditions that: i) maximize reaction rate; and ii) minimize predictive uncertainty in the rate predictions. The framework was demonstrated on a case study related to the catalytic oxidation of methane over Pd/Al_2O_3, a system previously investigated in an automated [8] and autonomous microreactor platform [9].

2. METHODOLOGY

2.1 System Description

The methane oxidation reaction system (see section 2.2 for further details) is modelled as a plug flow reactor (PFR) using differential-algebraic equations (DAEs) describing state variables, inputs, and measured outputs:

$$\mathbf{f}(\dot{\mathbf{x}}, \mathbf{x}, \mathbf{u}, \theta, t) = 0 \quad (1) \qquad \mathbf{y} = g(\mathbf{x}, \mathbf{u}, \mathbf{v}) \quad (2)$$

In Eq. (1), and Eq. (2) $\mathbf{x} \in \mathbb{R}^{n_x}, \dot{\mathbf{x}}, \mathbf{u} \in \mathbb{R}^{n_u}, \mathbf{y} \in \mathbb{R}^{n_y}, \theta \in \mathbb{R}^{n_\theta}$, and t denote the vectors for state variables of the system, their time derivative, the manipulated control variables (considered as time-independent inputs in this work), the N_y -dimensional vector of response variables that are measurable, the N_θ -dimensional vector of model parameters, and time variable, respectively.

2.2 Reactor Model for Methane Oxidation

Although the general system formulation in Equation (1) is expressed in terms of time t, the PFR model applied here (Eq. 3) is defined with respect to catalyst mass w as the independent variable. The physical system had some noisy observation described as $\mathbf{v} \in \mathbb{R}^{n_v}$ which represents random experimental errors in the measured outlet mole fractions where $\mathbb{R}$ denotes the set of real numbers and $n_v$is the dimension of the noise vector.

$$\frac{dx_1}{dw} = \frac{R.u_1}{u_2 P_{avg}} . (-r) \; x_1(0) = u_4$$

$$\frac{dx_2}{dw} = \frac{R.u_1}{u_2 P_{avg}} . (-2r) \; x_2(0) = u_3 . u_4$$

$$\frac{dx_3}{dw} = \frac{R.u_1}{u_2 P_{avg}} . (r) \qquad (3)$$

$$\frac{dx_4}{dw} = \frac{R.u_1}{u_2 P_{avg}} . (2r) .$$

In Eq. 3 and following the notation introduced in Eqs. 1-2, the $\mathbf{x}$ vector of state variables x_1, x_2, x_3 and x_4 represents the mole fractions of methane, oxygen, carbon dioxide and water respectively in the reactor outlet stream, and R $[J\,\mathrm{mol}^{-1}K^{-1}]$ is the universal gas constant, r $[mol\,\mathrm{g}^{-1}min^{-1}]$ represents the reaction rate according to a postulated kinetic model. The set of control variables $\mathbf{u}$ includes the reaction temperature T[°C], flow rate of the feed $[N\mathrm{ml\ min}^{-1}]$, oxygen to methane mole fraction ratio $[\mathrm{mol\ mol}^{-1}]$ and inlet methane mole fraction $[\mathrm{mol\ mol}^{-1}]$ which are respectively denoted as u_1, u_2, u_3, u_4 [9]. These controls form the design vector φ which is bounded within the experimental design domain as shown in Table 1. The kinetic parameter vector θ = $[k_0$ and $E_a]$ consists of the pre-exponential factor and activation energy of the Arrhenius rate expression.

Table 1: Upper and Lower Bounds on Control Variables Defining the Experimental Design Space

Control Variables	Range	Units
Temperature (T)	250-350	°C
Feed flow rate ($\dot{m}$)	20-30	$\mathrm{Nml\ min}^{-1}$
Oxygen to methane mole ratio (R_{O_2/CH_4})	2-4	$\mathrm{mol\ mol}^{-1}$
Inlet conc. of methane (CH_4)	0.005-0.025	$\mathrm{mol\ mol}^{-1}$

2.3 Hybrid (Power-Law & GP) Framework

A schematic of the proposed hybrid workflow combining the lumped power law model and the GP is shown in Figure 1 and is articulated in the steps below:

1. Lumped Power-law model formulation: a power-law model is a simplified approximation of the LH model for the adsorption of methane. The simplification of the power law model is discussed in Section 2.4.
2. Experiment Execution: the experimental design space (see Table 1) was sampled using Latin Hypercube Sampling (LHS) technique, and physical experiments were conducted in the laboratory generating a set of preliminary experimental data.
3. Parameter Estimation: the lumped power law model parameters were then estimated using the experimental data generated in Step 2. The parameter estimation technique is discussed in Section 2.5
4. Stopping criteria: the process stops when the t-test is satisfied for the model parameters or the experimental budget is terminated. If the stopping criteria are not met new experiments from DoE need to be executed.
5. In silico Data Generation from the identified Power-

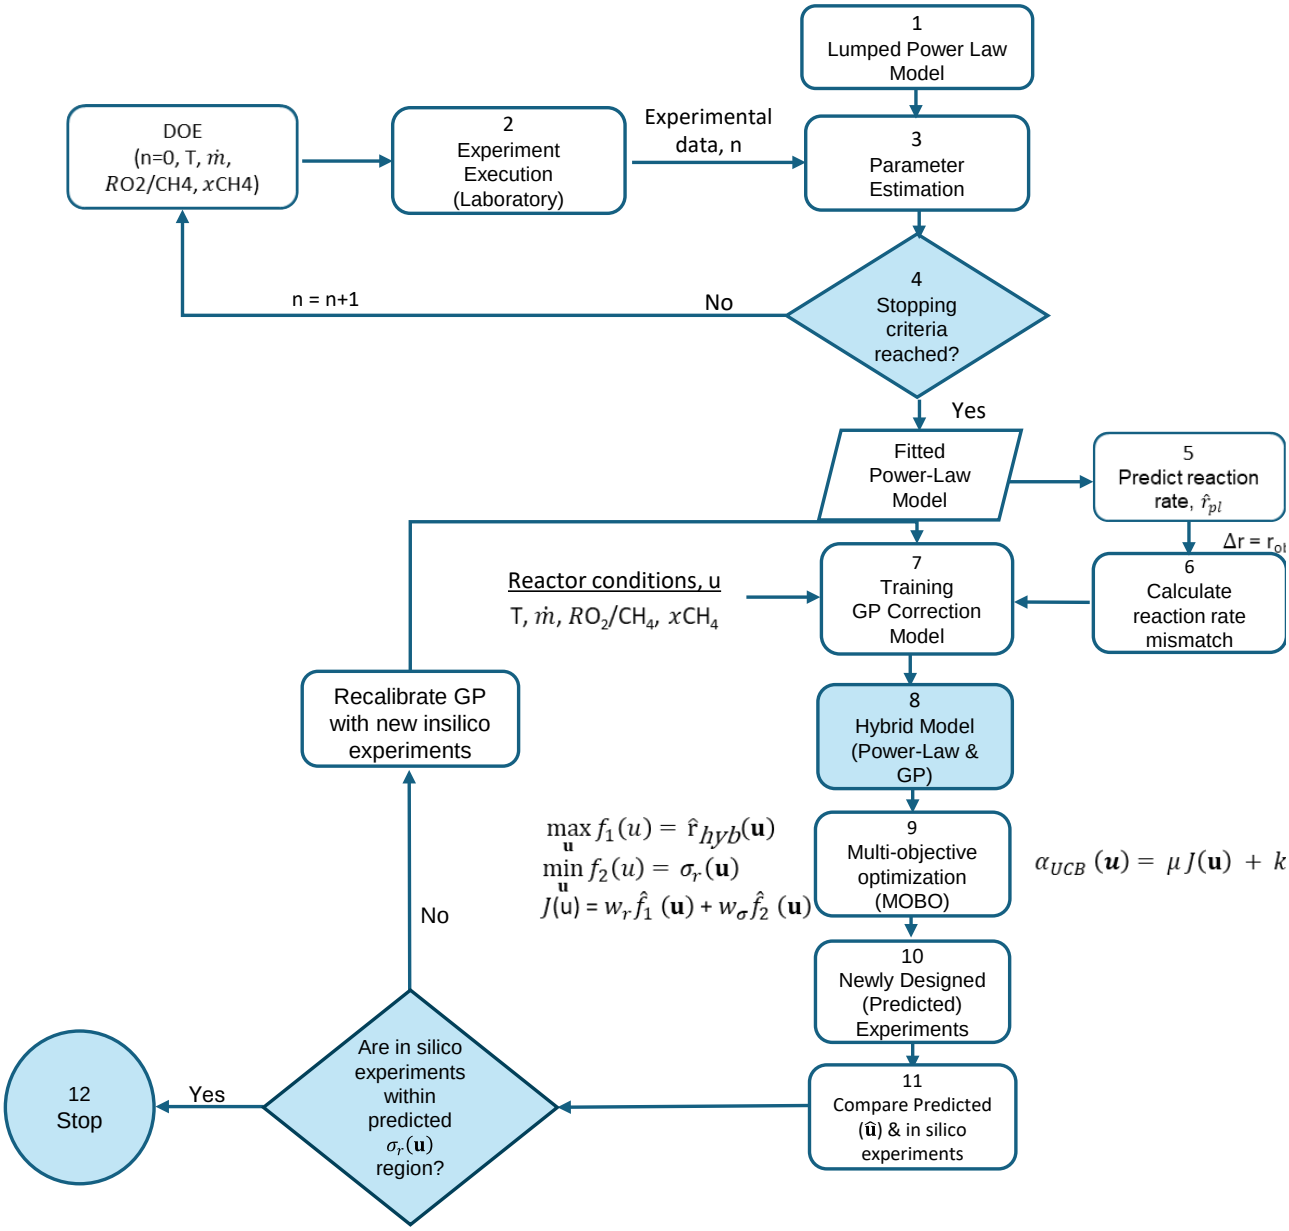

Figure 1. Procedural framework for hybrid (Power-Law & GP) model identification and the MOBO workflow.

Law model: after the *t*-test is satisfied, the corresponding parameter values are then applied to the-power-law model, and the reactor model is used to simulate the reaction rates at the LHS generated conditions.

6. Calculate Reaction Rate Mismatch: the rates predicted by the model are then compared with the experimentally calculated rates. The rate mismatch calculated across experiments is used to train a GP model at the corresponding input conditions of the experiments.
7. GP Modelling: the GP model is trained using experimental data to learn the relationship between the input conditions and the corresponding system rate mismatch [10]. The dataset includes 30 experiments, with 21 used for training and 9 reserved for testing. The GP is built with a composite kernel function Eq. 4

$$k(\mathbf{u}, \mathbf{u}') = C.RBF\,(\mathbf{u}, \mathbf{u}') \tag{4}$$

The kernel consists of a constant term (C = 1.0, bounds 10^{-3}–10^{3}) and an *RBF* with length scale 1.0 (bounds 10^{-2}–10^{2}). The hyperparameters of the kernel are optimized during model training by maximizing the log marginal likelihood using gradient-based

optimization technique. In practice, the model automatically tunes its settings during training using a gradient-based method (L-BFGS-B), while respecting predefined bounds.

8. Finalized Hybrid Model: the general form for the reaction rate predicted by the hybrid model is represented in Eq. 5 as:

$$\hat{r}_{\text{hyb}}(\mathbf{u}) = r_{PL}(\mathbf{u};\theta)\exp\left(\boldsymbol{\mu}_{GP}(\mathbf{u})\right) \quad (5)$$

where $r_{PL}(\mathbf{u};\theta)$ is the mechanistic power-law model prediction and $\boldsymbol{\mu}_{GP}(\mathbf{u})$ is the posterior mean of the GP trained on the log-residuals between experimental and mechanistic reaction rates.

9. Multi-objective Bayesian Optimization (MOBO): the multi objectives formulation involve maximizing reaction rate and minimizing uncertainty; this is discussed in details in section 2.6.

10. Newly Designed (Predicted) Experiment: the hybrid-MOBO is then used to design new experiments. These experiments are a pareto sets of optimal experiments (in this case we selected top three (3) experiments) which gives: (i) the predicted reactor conditions ($\hat{\mathbf{u}}$); (ii) the corresponding maximum reaction rate $\hat{r}(\mathbf{u})$; and (iii) minimum uncertainty $\hat{\sigma}_r(\mathbf{u})$ at these conditions.

11. Comparison (Predicted Experiments) With Insilico Experiments: in silico experiments generated using a ground truth model are then compared to the predicted experiments by the hybrid-MOBO framework, where the agreement is evaluated using the GP uncertainty bounds, i.e, the predicted mean (from the ground truth model) reaction rate $\hat{r}_{\text{true}}(\mathbf{u})$ must lie within the confidence region of the hybrid-MOBO. If the insilico results falls inside the uncertainty region, the hybrid model is considered reliable and the workflow terminates. Otherwise, the discrepancy is used to recalibrate the GP correction model, and the optimization-prediction loop is repeated.

12. Stop: the final workflow is terminated when the insilico experimental results consistently fall within the predicted uncertainty bounds of the hybrid-MOBO predictions. i. e., for all newly proposed conditions $\hat{\mathbf{u}}$ from the MOBO step. This indicates that the hybrid power-law-GP model has achieved sufficient predictive accuracy and reliability, and no further model corrections are required.

2.4 Lumped Power-Law model formulation

LH mechanism has been proposed as a plausible model by other researchers [8, 11] for catalytic methane oxidation over Pd/Al_2O_3. The reaction mechanism occurs with methane and oxygen been adsorbed on the catalyst surface prior to the reaction. The LH rate expression for the system [11] is defined in Eq. 6 as:

$$r_{CH_4} = \frac{k_r K_{CH_4} p_{CH_4}\sqrt{(K_{O_2}p_{O_2})}}{\left(1+K_{CH_4}p_{CH_4}+\sqrt{(K_{O_2}p_{O_2})}\right)^2} \quad (6)$$

where k_r is the intrinsic surface-reaction rate constant, K_{CH_4} K_{O_2}, p_{CH_4} and p_{O_2} are the adsorption equilibrium constants and partial pressures of methane and oxygen respectively. Methane oxidation occurs at low surface coverage with excess oxygen, reflecting adsorption–reaction balance [12, 13].

Under these conditions, $K_{CH_4}p_{CH_4}$ <<< 1, $K_{O_2}p_{O_2}$ >>>1, this approximation reduces the LH rate law to a lumped power-law form defined in Eq. 7 as:

$$r_{CH_4} = k^{'}\ p_{CH_4}{}^{a} p_{O_2}{}^{b} \quad (7)$$

where k' is the lumped rate constant and a, b are the reaction orders. In the presence of excess oxygen methane adsorption was identified as the rate controlling step [13, 15], thus giving first-order methane and zero-order oxygen dependence; thus, reaction orders were fixed at a = 1 and b = 0, and only k_0 and E_a were estimated. The rate constant is expressed in Eq. 8 as:

$$k(T) = exp\left(\log k(T_{ref}) - \frac{E_a}{R}\left(\frac{1}{T} - \frac{1}{T_{ref}}\right)\right) \quad (8)$$

where $k(T)$ is the reaction rate constant at temperature T, and T_{ref} = 320 °C and E_a is the activation energy in kJ/mol

The resulting lumped power-law model provides a simplified mechanistic representation of the dominant kinetic behaviour derived from the Langmuir–Hinshelwood mechanism under the experimentally relevant conditions. In the proposed hybrid modelling framework, this mechanistic model serves as the baseline predictive model, preserving physically meaningful kinetic parameters such as the activation energy $E_a$and pre-exponential factor k_0. The Gaussian Process component is subsequently trained on the residual error between the mechanistic model predictions and the experimental observations, allowing the model to capture discrepancies arising from unmodelled effects such as transport limitations or catalyst heterogeneity while retaining the physical interpretability of the kinetic parameters.

2.5 Parameter Estimation

Parameter estimates (indicated as $\hat{\theta}$) were obtained by minimizing the sum of squared residuals between the logarithm of the experimental rates and the logarithm of the model predictions. Assuming an independent Gaussian measurement errors with constant variance $\sigma^2{}_y$, the negative log-likelihood for Nexperiments is expressed in Eq. 9 as:

$$L(\hat{\theta}) = \frac{N}{2}\log(2\pi\sigma^2{}_y) + \frac{1}{2\sigma^2{}_y}\sum_{i=1}^{N}(\log r_i{}^{exp} - \log r_i{}^{model}(\theta))^2 \quad (9)$$

where r_i^{exp} is the observed reaction rate for experiment $i, r_i^{model}(\theta)$ is the power-law model prediction, N is the number of experiments, and $\sigma^2{}_y$ is the constant variance of the measurement noise.

2.6 Multi-objective Bayesian Optimization

Bayesian optimization (BO) guided the selection of subsequent experiments using probabilistic predictions from the hybrid power-law–GP model to maximize methane oxidation rate while minimizing uncertainty.

2.6.1 Hybrid model predictions

For a given set of experimental conditions $\mathbf{u}$, The hybrid model predicts the methane reaction rate as a posterior distribution $\mathrm{r}(\mathbf{u})$, whose predictive mean is given by $\hat{\mathrm{r}}_{\mathrm{hyb}}(\mathbf{u})$, and expressed in Eq. 10 as:

$$r(\mathbf{u}) \sim N\left(\hat{\mathrm{r}}_{\mathrm{hyb}}(\mathbf{u}), \sigma_r^2(\mathbf{u})\right) \tag{10}$$

where $\sigma_r^2(\mathbf{u})$ is the total prediction uncertainty of the methane reaction rate obtained by combining the GP posterior variance with mechanistic parametric uncertainty propagated through the power-law model under the assumption that all uncertainty contributions are assumed independent, is expressed in Eq. 11 as:

$$\sigma_r(\mathbf{u}) = \sqrt{(\sigma_{mech}^2(\mathbf{u}) + \sigma_{r,GP}^2(\mathbf{u})} \tag{11}$$

where $\sigma_{mech}^2(\mathbf{u})$ is the mechanistic model uncertainty propagated from parameter estimates, and $\sigma_{r,GP}^2(\mathbf{u})$ is the GP-derived predictive uncertainty of methane reaction rate mapped to rate space.

2.6.2 Multi objective optimization problem

The resulting multi-objective optimization problem is defined as (i) maximizing the reaction rate of methane; and (ii) minimizing the uncertainty on the predicted rate, and is mathematically formulated in Eq. 12 as:

$$\max_{\mathbf{u}} f_1(u) = \hat{\mathrm{r}}_{\mathrm{hyb}}(\mathbf{u}) \quad \min_{\mathbf{u}} f_2(u) = \sigma_r(\mathbf{u}) \tag{12}$$

To enable practical optimization, the two competing objectives are combined using a weighted sum scalarization shown in Eq 13 as:

$$J(\mathrm{u}) = w_r \hat{f}_1(\mathbf{u}) + w_\sigma \hat{f}_2(\mathbf{u}) \tag{13}$$

where $\hat{f}_1$, and $\hat{f}_2$ denote normalized forms of the rate and uncertainty objectives respectively, and w_r and w_σ are non-negative weighting factors satisfying $w_r + w_\sigma = 1$. The weights reflect a user-defined trade-off choice, here prioritizing reaction performance (w_r = 0.6) over uncertainty (w_σ= 0.4).

2.6.3 BO with Upper Confidence Bound (UCB)

The scalarized objective function $J(\mathbf{u})$ defined in Eq. 13 is optimized using BO was implemented in Python using a typed function called "scikit-learn". The Gaussian process surrogate hyperparameters are learned by maximizing the log marginal likelihood using the L-BFGS-B gradient-based optimizer, while candidate operating conditions are selected by maximizing a UCB acquisition function defined in Eq. 14 (which balances exploitation of conditions with high predicted objective values, and exploration of regions of high uncertainty) over a discretized design space.

$$\alpha_{UCB}(\boldsymbol{u}) = \mu\, J(\mathbf{u}) + k\sigma\, J(\mathbf{u}) \tag{14}$$

where $\mu_j(\boldsymbol{u})$ and $\sigma_j(\boldsymbol{u})$ are the predictive mean and standard deviation of the scalarized objective, and k whose value is 0.1 controls the exploration-exploitation trade-off. At each iteration, the next experiment is selected using the expression in Eq. 15:

$$\boldsymbol{u}^* = \underset{\mathbf{u}}{\operatorname{argmax}}\, \alpha_{UCB}(\boldsymbol{u}) \tag{15}$$

UCB was chosen over alternative acquisition function such as Expected Improvement (EI) due to its explicit control of uncertainty-driven exploration, which is particularly suited for autonomous experimentation [14].

Thirty experiments were generated via Latin hypercube sampling, with 21 used for training and 9 for testing. The Mars Van Krevelen (MVK) model reaction rate expression (Eq. 16) was adopted as a "ground truth" model in this study to generate in-silico data in the MOBO step and was experimentally validated in [8] where the authors calibrated the MVK model from flow reactor data.

$$r_{ch4} = \frac{k_1 k_2 P_{CH_4} P_{O_2}}{k_1 P_{O_2} + 2\, k_2 P_{CH_4} + \frac{k_1 k_2}{k_3}(P_{O_2} P_{CH_4})} \tag{16}$$

The MVK model was also adopted to validate the predictive performance of the hybrid model, with the corresponding MVK parameter vector θ = [5.77, 6.72, 5.87, 9.51, 10.17, 7.98] as reported in [9].

2.7 Reliability Map of the Identified Hybrid Model

To analyse the contribution of each component in the hybrid model, two normalized indices were defined in Eq. 17 based on the hybrid prediction ŷ, which is taken as the reference (ground truth):

$$\hat{I}_{PL} = \frac{\hat{y}_{PL}}{\hat{y}} \quad \hat{I}_{GP} = \frac{\hat{y}_{GP}}{\hat{y}} \tag{17}$$

where $\hat{y}_{PL}$ is the rate prediction from the mechanistic power-law model across the design space, and $\hat{y}_{GP}$ is the GP correction term, and the total hybrid prediction is given by:

$$\hat{y} = \hat{y}_{PL} + \hat{y}_{GP} \tag{18}$$

These indices quantify the relative contribution of each component to the final prediction. Specifically, $\hat{I}_{PL}$ indicates how much of the prediction is explained by the mechanistic model, while $\hat{I}_{GP}$ quantifies the contribution of the data-driven correction.

3. RESULTS & DISCUSSION

This section discusses: (i) the predictive performance of the fitted lumped power-law model for observed reaction rate; (ii) the hybrid (power-law–GP) model performance relative to the reference (MVK) model across the design space; and (iii) the reactor operating conditions suggested by the optimization that simultaneously maximize reaction rate and minimize predictive uncertainty.

3.1 Performance of the lumped power-law model

As shown in Figure 2, the fitted lumped power-law model captures the general trend of methane oxidation rates across the experimental dataset but exhibits systematic deviations at higher reaction rates.

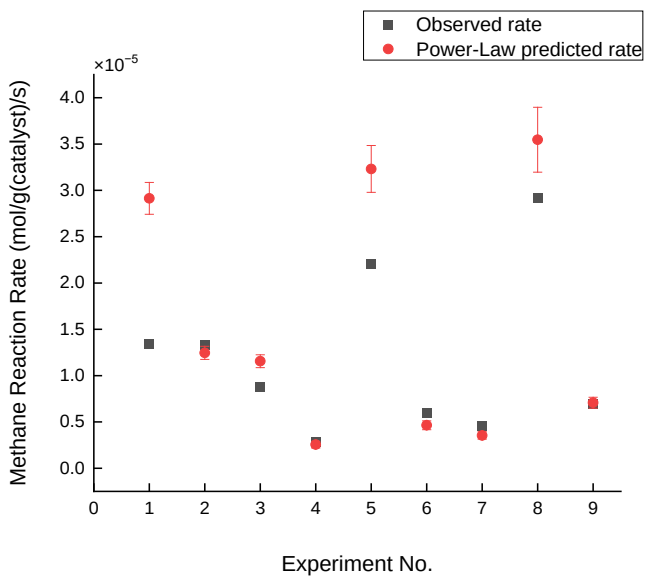

Figure 2. Observed methane oxidation reaction rates and Power-Law predictions for the test dataset.

This reflects the breakdown of its simplified kinetic assumptions. Under high reactor rate conditions (i.e. conditions close to the following: T = 350 ℃, $\dot{m}$ = 30 Nml min^{-1}, RO_2/CH_4 = 4 mol mol^{-1}, xCH4 = 0.025 mol mol^{-1}) methane oxidation is likely influenced by surface coverage and adsorption saturation effects, which are not captured by the reduced power-law form.

3.2 Hybrid Model performance: Power-Law model with Gaussian Process Correction

As shown in Figure 3, the hybrid power-law–GP model reproduces the experimental methane oxidation rates with good accuracy across the test dataset. Predicted rates closely match observed reaction rate values, for example 1.43×10^{-5} vs. 1.35×10^{-5} mol g^{-1} s^{-1} (Exp.1) and 8.84×10^{-6} vs. 8.74×10^{-6} mol g^{-1} s^{-1} (Exp.3). At higher temperature (T=342℃, Exp. 8; Table 2), the hybrid model predicts 3.20×10^{-5} mol g^{-1} s^{-1}, compared to an experimental value of 2.92×10^{-5} mol g^{-1} s^{-1}, capturing the trend without the systematic overprediction observed for the power-law model. Experiments 1, 3, and 8 were selected as representative cases covering the test dataset operating range (Table 2).

Table 2: Selected Experiments from Test Datasets

Exp	Temp ℃	Feed flow Nml min^{-1}	O_2/CH_4 mol mol^{-1}	CH_4 conc mol mol^{-1}
1	311	23.8	3.6	0.025
3	290	23.5	3.2	0.021
8	342	25.8	2.9	0.022

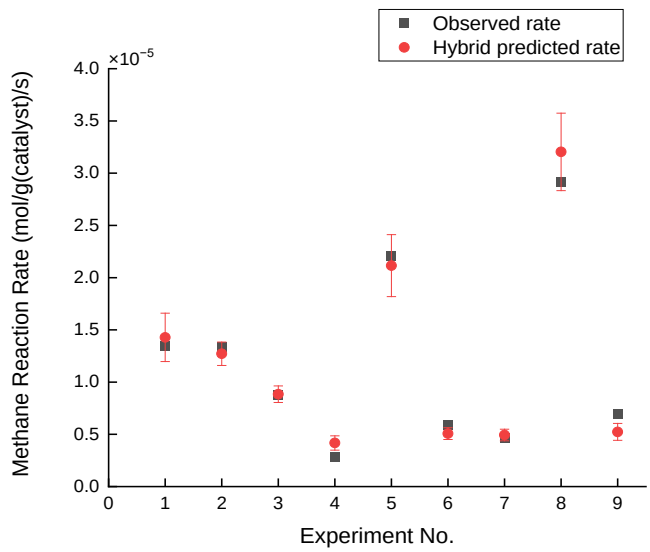

Figure 3. Observed methane oxidation reaction rates and hybrid model predictions for the test dataset.

3.3 Hybrid Model performance against MVK Prediction

The hybrid model exhibits close agreement with the observed reaction rates across the investigated operating conditions (test dataset), while consistently following the same overall trends as the mechanistic MVK model as observed in Figure 4. At moderate temperature (e.g. Exp. 3, $r_{exp} \approx 8.7 \times 10^{-6}$ mol g^{-1} s^{-1}), both the MVK and hybrid predictions are in close agreement with the experimental data, indicating that the MVK formulation adequately captures the dominant reaction pathway in this regime. At higher rates (e.g. Exp. 5 and 8, $r_{exp} \approx (2.2–2.9) \times 10^{-5}$ mol g^{-1} s^{-1}), the hybrid model remains well aligned with the measurements, whereas the MVK model slightly underpredicts the observed rates.

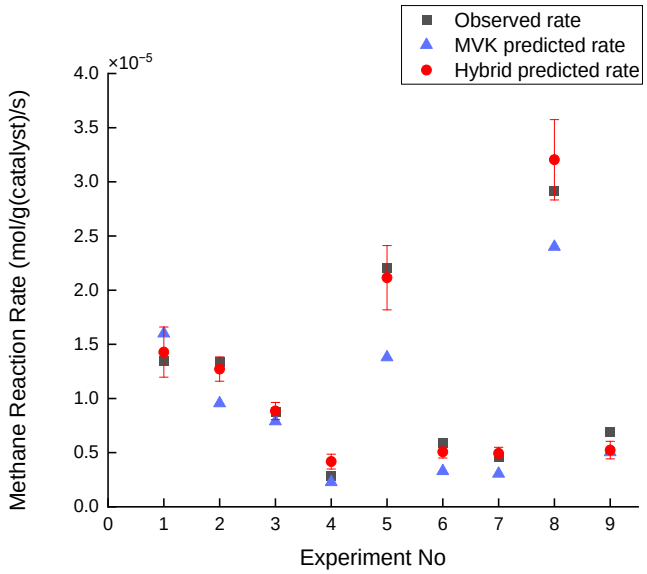

Figure 4. Observed MVK, and Hybrid Model Predictions of Methane Oxidation Rates for test dataset.

3.4 Experimental Design using Hybrid (Power-law & GP) MOBO

The Hybrid-MOBO strategy identifies the Pareto set which consists of three high-ranking experiments at 350 °C and 30 Nml min^{-1} by balancing predicted reaction rate and uncertainty (Table 3). For O_2/CH_4 ratios of 4.0, 3.5, and 3.0, and a constant inlet conc. of methane at 0.005 mol/mol. Hybrid-MOBO predicts reaction rates between 4.71 × 10^{-5} mol g^{-1} s^{-1} and 4.03 × 10^{-5} mol g^{-1} s^{-1}.

Table 3: Experimental conditions for the top three experimental candidates from MOBO achieving maximum reaction rate and minimum uncertainty.

Exp	Temp °C	Feed flow $Nml\ min^{-1}$	O_2/CH_4 mol mol^{-1}	CH_4 rate (x 10^{-5}) mol g^{-1} s^{-1}
1	350	30	4	4.7± 2.9
2	350	30	3.5	4.4± 2.6
3	350	30	3	4.0± 2.7

When the mechanistic MVK model is evaluated at these same conditions, comparable rates (4.3 × 10^{-5} mol g^{-1} s^{-1} to 4.1 × 10^{-5} mol g^{-1} s^{-1}) are obtained, as shown in Figure 5, confirming that the Hybrid-MOBO selected operating conditions align with mechanistic trends while explicitly accounting for predictive uncertainty.

3.5 Reliability Maps of the Identified Hybrid Model

Figures 6 and 7 show the spatial distribution of $\hat{I}_{PL}$ and $\hat{I}_{GP}$, respectively, across the design space. The results indicate that the power-law model dominates most of the domain, particularly in the mid-temperature region (~270–320 °C), where $\hat{I}_{PL}$ approaches unity and the GP contribution is minimal. This suggests that the mechanistic model alone adequately captures the system behaviour in this region.

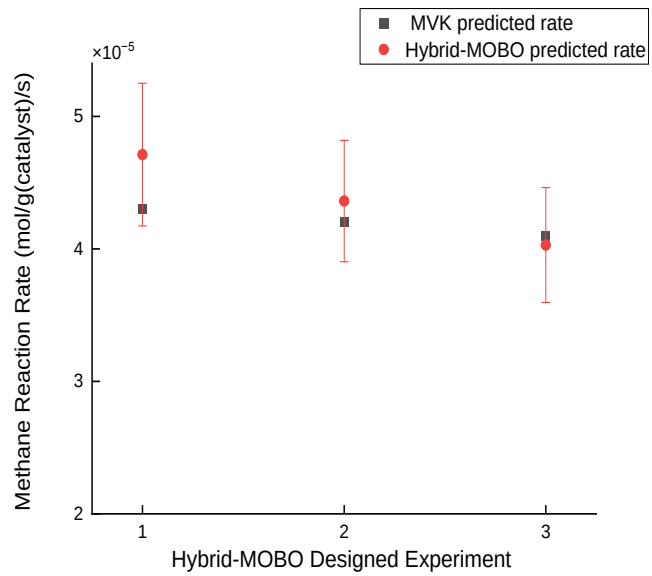

Figure 5. Hybrid-MOBO and MVK reaction rate predictions for the top three experiments selected by multi-objective Bayesian optimization.

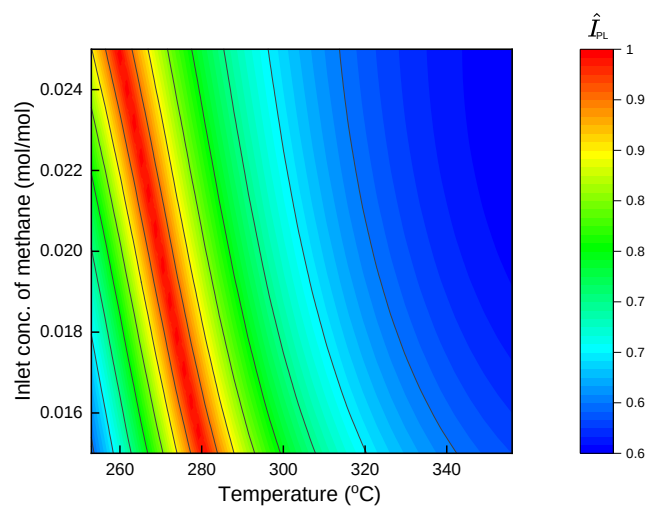

Figure 6. Relative local contribution of the mechanistic power-law model across the design space.

At higher temperatures (above ~320 °C), the GP contribution increases (Figure 7), reflecting the need for data-driven corrections to account for deviations from the simplified kinetic model. Overall, these maps highlight regions where the mechanistic model is sufficient and where the GP component becomes essential for accurate predictions.

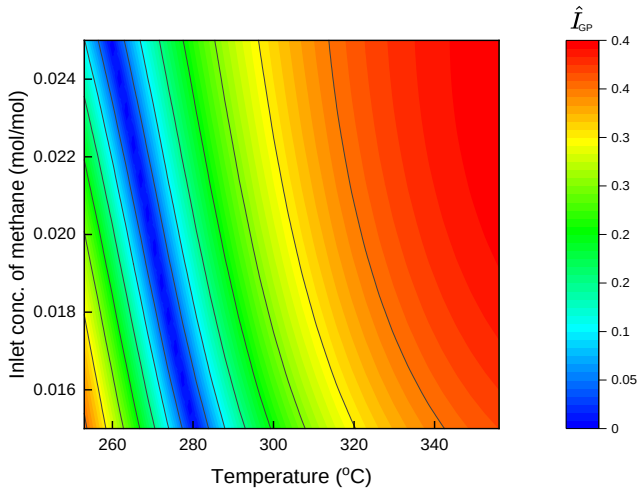

Figure 7. Relative local contribution of the GP correction across the design space.

4. CONCLUSION

A hybrid modelling MOBO framework was proposed in this work where a lumped power-law kinetic model with a GP residual model was used to improve reaction rate prediction while minimising the uncertainty of the predicted rate. The hybrid MOBO framework identified operating conditions providing high reaction rates with minimum uncertainty. The relative percentage in reaction rate prediction uncertainty of the three top-ranked MOBO designed experiments averaged 2.7% and was in close agreement with the optimal value determined by the MVK ground truth model. Future work will extend the optimization framework to alternative objective functions based on key performance indicators such as cost of operations.

ACKNOWLEDGEMENTS

The authors are grateful to the Department of Chemical Engineering at University College London (UCL), the Galvanin System Identification Group (GSIG) and the Petroleum Trust Development Fund (PTDF) for their support and sponsorship of the research.

REFERENCES

1. Ljung L. Perspectives on system identification. Annual Reviews in Control 34:1-12 (2010). https://doi.org/10.1016/j.arcontrol.2009.12.001
2. Van den Hof, P. M. (2020). System identification-Data-driven modelling of dynamic systems. *Lecture Notes*, 305.
3. Taylor CJ, Booth M, Manson JA, Willis MJ, Clemens G, Taylor BA, Chamberlain TW, Bourne RA. Rapid, automated determination of reaction models and kinetic parameters. Chemical Engineering Journal 413:127017 (2021). https://doi.org/10.1016/j.cej.2020.127017
4. Stalidzans E, Zanin M, Tieri P, Castiglione F, Polster A, Scheiner S, Pahle J, Stres B, List M, Baumbach J, Lautizi M, Van Steen K, Schmidt HHHW. Mechanistic modeling and multiscale applications for precision medicine: theory and practice. Network and Systems Medicine 3:36-56 (2020). https://doi.org/10.1089/nsm.2020.0002
5. Bradley W, Kim J, Kilwein Z, Blakely L, Eydenberg M, Jalvin J, Laird C, Boukouvala F. Perspectives on the integration between first-principles and data-driven modeling. Computers & Chemical Engineering 166:107898 (2022). https://doi.org/10.1016/j.compchemeng.2022.107898
6. Chen Y, Ierapetritou M. A framework of hybrid model development with identification of plant?model mismatch. AIChE Journal 66: (2020). https://doi.org/10.1002/aic.16996
7. Zhang J, Semochkina D, Sugisawa N, Woods DC, Lapkin AA. Multi-objective reaction optimization under uncertainties using expected quantile improvement. Computers & Chemical Engineering 194:108983 (2025). https://doi.org/10.1016/j.compchemeng.2024.108983
8. Bawa SG, Pankajakshan A, Waldron C, Cao E, Galvanin F, Gavriilidis A. Rapid screening of kinetic models for methane total oxidation using an automated gas phase catalytic microreactor platform. Chemistry Methods 3: (2022). https://doi.org/10.1002/cmtd.202200049
9. Pankajakshan A, Bawa SG, Gavriilidis A, Galvanin F. Autonomous kinetic model identification using optimal experimental design and retrospective data analysis: methane complete oxidation as a case study. React. Chem. Eng. 8:3000-3017 (2023). https://doi.org/10.1039/d3re00156c
10. Cicirello A, Giunta F. Machine learning based optimization for interval uncertainty propagation. Mechanical Systems and Signal Processing 170:108619 (2022). https://doi.org/10.1016/j.ymssp.2021.108619
11. Hurtado P, Ordóñez S, Sastre H, D??ez FV. Development of a kinetic model for the oxidation of methane over pd/al2o3 at dry and wet conditions. Applied Catalysis B: Environmental 51:229-238 (2004). https://doi.org/10.1016/j.apcatb.2004.03.006
12. Froment, G. F., Bischoff, K. B., & De Wilde, J. (2011). *Chemical reactor analysis and design* (3rd ed.). John Wiley & Sons.
13. Kumar RS, Mmbaga JP, Semagina N, Hayes RE. Methane combustion kinetics over palladium-

based catalysts: review and modelling guidelines. Catalysts 14:319 (2024). https://doi.org/10.3390/catal14050319

14. Munyebvu N, Dunn S, Howes PD. Multiobjective platform for autonomous property targeting and optimization of colloidal lead halide perovskite quantum dots. Chem. Mater. 37:6629-6641 (2025). https://doi.org/10.1021/acs.chemmater.5c01153
15. Escandón LS, Ordóñez S, Vega A, D??ez FV. Oxidation of methane over palladium catalysts: effect of the support. Chemosphere 58:9-17 (2005). https://doi.org/10.1016/j.chemosphere.2004.09.012

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Universal Framework for Automated Reaction Network Identification and Interpretable Rate Model Generation

Harry Kay[a], Alexander Rogers[a], and Dongda Zhang[a]*

[a] The University of Manchester, Department of Chemical Engineering, Manchester, M13 9PL, UK
* Corresponding Author: dongda.zhang@manchester.ac.uk

ABSTRACT

Mathematical models are paramount to the field of reaction engineering, facilitating reaction mechanism discovery, process optimisation, and informed decision making in academic and industrial settings. Nevertheless, the development of precise mechanistic reaction rate models remains experimentally intensive, requires expert knowledge, and is susceptible to the introduction of structural bias. Similarly, the identification of a suitable reaction network that depicts all chemical transformations remains a non-trivial task, with existing techniques often being ill-suited for large and complex systems, hence limiting their scalability and implementation within chemical and biochemical applications. This work develops a two-stage autonomous framework that exploits non-linear sparse optimisation to identify the minimum size global reaction network representative of the system under study, and subsequently proposes and discriminates between interpretable rate equations developed through symbolic regression (SR). The generated SR expressions are constrained to a mechanistically meaningful form through the use of a novel substructure decomposition strategy, largely reducing the search space that must be explored and increasing model interpretability. The framework is evaluated on two case studies; the first, depicting a catalytic methanol synthesis reaction network, and the second, an enzymatic kinetic resolution network. The methodology exhibited high levels of accuracy, scalability, data efficiency, and robust network identification consistent with known physics. Lastly, the potential of augmented intelligence, is discussed as a method to enhance fidelity. Therefore, this work represents a key step toward autonomous process modelling and digitalisation in reaction engineering, providing a foundation for accelerated design and development of chemical and biochemical processes.

Keywords: Reaction network identification, Symbolic regression, Interpretable model construction, Model based design of experiments, Augmented intelligence

INTRODUCTION

The growing global emphasis on sustainability, coupled with mounting economic pressures, has driven the chemical and biochemical industries to seek more efficient processes with reduced waste, energy consumption and raw material costs. A central aspect to enabling the advancement of (bio)chemical processes is the accurate characterisation of reaction mechanisms through reliable modelling techniques, which underpins process optimisation, scale-up, and control.

Conventional modelling practices are typically based on first-principle [1-3] or empirical kinetics such as power-law, Langmuir–Hinshelwood, or Michaelis–Menten expressions [4-7] due to their ability to impart kinetic understanding through their inherent interpretability. However, for complex reaction networks that are less well understood, identifying a representative model is difficult [8], and accounting for all physical phenomena may drive model complexity beyond means which available data can sustain [9]. In order to limit complexity, expert knowledge is exploited, which introduces inductive bias and may reduce model generalisability.

In order to eliminate introducing inductive bias for reaction rate model construction, recent literature has focused on developing efficient automated chemical reaction network (CRN) identification techniques. By formulating the CRN as a mixed-integer linear programming

problem [10, 11], all mass balanced chemical transformations can be found, and rate expressions fitted to experimental data assuming first-order or categorical power law dependencies. However, if the reaction under study is defined by complex kinetics, such as typical within catalytic systems, simple power law dependencies are not sufficient for capturing the underlying physics. Additionally, for reaction networks described by larger numbers of species, the number of mass balanced models that must be evaluated increases exponentially. As a result, existing modelling approaches struggle with accuracy and scalability, particularly for complex systems which are commonplace in industrial and academic research.

To address these challenges, this work introduces a universal data-driven framework for automated reaction network identification and generation of interpretable rate expressions that resemble mechanistic models. We develop a sparse regression based optimisation strategy to effectively identify the underlying reaction network, and evolve an iterative symbolic regression (SR) procedure integrated with model-based design of experiment (MBDoE) to maximise information gain, thus facilitating rapid discovery of mechanistic knowledge for complex reaction systems. The framework was validated on two case studies: (i) a heterogeneous catalytic methanol synthesis reaction network defined by 3 complex Langmuir-Hinshelwood type rate expressions representative of an industrial reaction, and (ii) a larger kinetic resolution network defined by 9 states and 5 net reactions.

We thoroughly investigate how the different attributes of a reaction system influence the identification of characteristic and balanced reaction networks, the achievable model accuracy of any generated SR expressions, and the prospective mechanistic insight that can be gained. Finally, the prospect of augmented intelligence is examined, integrating available domain knowledge with data-driven interpretable modelling in order to accelerate the model development process and promoting the discovery of kinetic knowledge that is physically meaningful. Consequently, this study highlights the potential of the developed framework towards autonomous reaction rate modelling and discovery of process knowledge for chemical and biochemical engineering applications.

CASE STUDIES

In this work, two in-silico case studies were implemented based on real world scenarios. The first delineates the heterogeneous catalytic hydrogenation of carbon dioxide [12] to methanol and was selected due to the complex nature of the rate expressions that describe the underpinning physics. The second study is an enzyme catalysed kinetic resolution reaction network [13], which was chosen to represent a larger reaction system, hence testing the scalability of the developed framework.

Methanol Synthesis from CO_2 Hydrogenation

The first case study examines the catalytic hydrogenation of CO_2 over an unsupported In_2O_3 catalyst. Experiments were performed in an isothermal fixed-bed reactor operated under isobaric conditions, with both internal and external mass transfer limitations assumed to be negligible. A comprehensive description of the experimental configuration and the corresponding operating conditions is provided in [12]. Within the operating regions that were investigated, three reversible reactions (R1) were identified that govern the systems behaviour as follows:

$$\begin{aligned} CO_2 + 3H_2 &\rightleftharpoons CH_3OH + H_2O \;\; \Delta G(25°C) = -0.47kJ\ mol^{-1} \\ CO_2 + H_2 &\rightleftharpoons CO + H_2O \;\; \Delta G(25°C) = 23.4kJ\ mol^{-1} \\ CO_2 + 4H_2 &\rightleftharpoons CH_4 + 2H_2O \;\; \Delta G(25°C) = -130.8kJ\ mol^{-1} \end{aligned} \tag{R1}$$

where the first reaction corresponds to the primary methanol formation pathway via CO_2 hydrogenation, and the second and third represent competing side reactions, namely the water–gas shift and methanation reactions. To enable reaction network identification within this in-silico study, four initial experiments were generated, yielding 16 data points across 6 state variables. All generated data was simulated using the literature proposed kinetic model at a temperature of 493 K and weight hourly space velocities (WHSV) ranging from 6000 to 16000 mL g^{-1}_{cat} h^{-1}.

The kinetic expressions that describe this system are highly non-linear and have been shown characteristic of standard industrial reactions. Notably, the application of an automated rate model generation technique to a case study high kinetic complexity has not previously been reported in the literature [14]. The ground truth expressions adopted to simulate the in-silico methanol synthesis system are provided in Equation 1.

$$\begin{aligned} r_1 &= \frac{k_1\left(P_{CO_2}P_{H_2}^3 - \frac{P_{CH_3OH}P_{H_2O}}{K_{eq,1}}\right)}{P_{H_2}^2\left(1 + K_{CO_2}P_{CO_2} + \sqrt{K_{H_2}P_{H_2}}\right)^2} \\ r_2 &= \frac{k_2\left(P_{CO_2}P_{H_2} - \frac{P_{CO}P_{H_2O}}{K_{eq,2}}\right)}{\sqrt{P_{H_2}}\left(1 + K_{CO_2}P_{CO_2} + \sqrt{K_{H_2}P_{H_2}}\right)^2} \\ r_3 &= \frac{k_3\sqrt{P_{CO_2}}\sqrt{P_{H_2}}\left(1 - \frac{P_{CH_4}P_{H_2O}^2}{K_{eq,3}P_{CO_2}P_{H_2}^4}\right)}{\left(1 + K_{CO_2}P_{CO_2} + \sqrt{K_{H_2}P_{H_2}}\right)^2} \end{aligned} \tag{1}$$

Here, r_1, r_2 and r_3 refer to the global reaction rates for methanol synthesis via CO_2, the water gas shift, and the methanation reaction, respectively. K_j and P_j are the adsorption equilibrium constants and partial pressures of

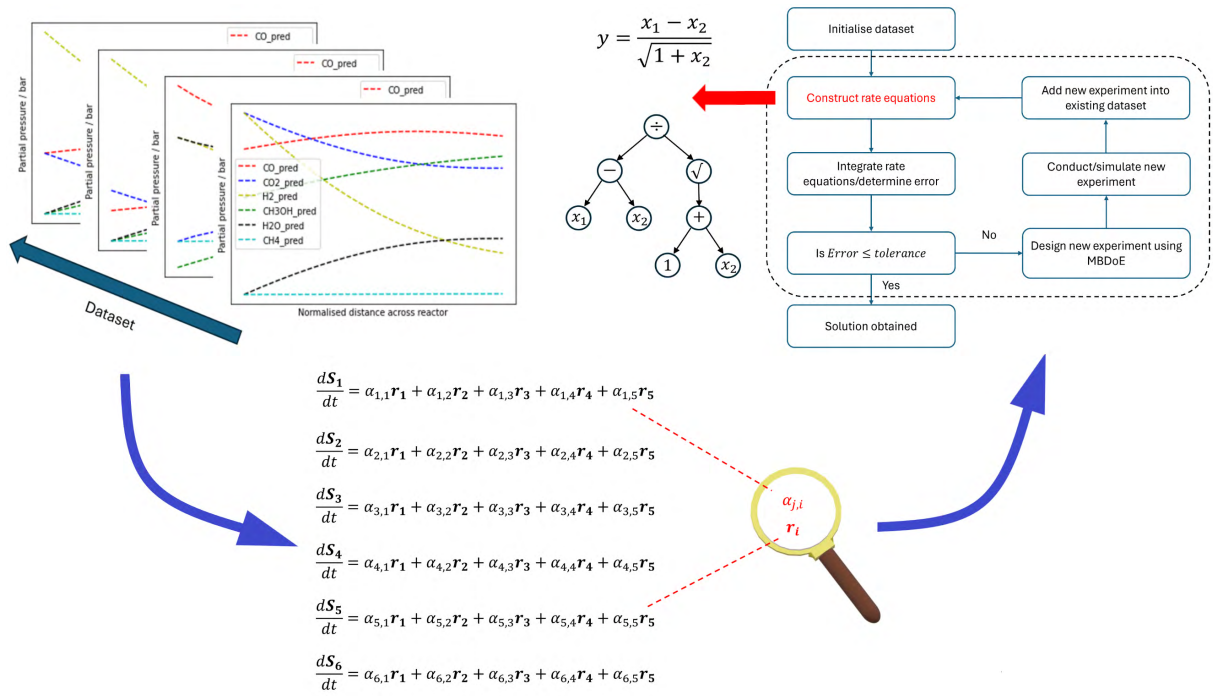

Figure 1: Schematic of the general devised framework for the identification of reaction networks and construction of representative rate equations

species j. k_i and $K_{eq,i}$ represent the reaction rate constant and equilibrium constant for a given reaction i. All values used to parameterise Equation 1 can be found in [12].

Enzyme Catalysed Kinetic Resolution

The second case study investigates the construction of kinetic models for the kinetic resolution of (+-)-trans-1, 2-cyclohexanediol in supercritical CO_2. The reactions take place in a batch reactor containing immobilised enzymes which is operated under isothermal and isobaric conditions. All developed literature rate models were formulated under the assumptions that: all reaction steps are elementary, the concentration of non-activated enzyme remains in excess, the catalysed reactions proceed irreversibly, and perfect mixing is maintained. Additional details surrounding the experimental setup and modelling assumptions can be found in [13].

For the purpose of the present study, the formation of the diacetate product S_3 was not ignored in the network to investigate the influence of minor side reactions on the accuracy of the proposed network identification and model generation framework. The reaction network and corresponding rate expressions used as the ground truth for generating the in-silico data are presented in Reaction Network (R2) and Equation (2).

$$VA \rightleftharpoons \mathrm{AE} + AA$$

$$AE + S_1 \rightarrow S_2$$

$$AE + S_2 \rightarrow S_3 \qquad (R2)$$

$$AE + R_1 \rightarrow R_2$$

$$AE + R_2 \rightarrow R_3$$

$$\begin{aligned} r_1 &= k_1 C_{VA} - k_{-1} C_{AE} C_{AA} \\ r_2 &= k_2 C_{AE} C_{S_1} \\ r_3 &= k_3 C_{AE} C_{S_2} \\ r_4 &= k_4 C_{AE} C_{R_1} \\ r_5 &= k_5 C_{AE} C_{R_2} \end{aligned} \qquad (2)$$

Species VA, AE and AA denote vinyl acetate, the active enzyme, and the by-product acetaldehyde, respectively. S_1 and R_1 correspond to cyclohexanediols, S_2 and R_2 to the monoacetate products, and S_3 and R_3 to the diacetate products, with R and S indicating the respective enantiomers. The terms r_i and k_i, in Equation 2, are the global reaction rate and rate constant of reaction i, while C_j denotes the concentration of species j. All parameter values used to parameterise the rate are reported in [13].

Similar to the previous case study, four in-silico experiments were generated, however since there are 9 state variables and 5 overall reactions present, more datapoints are required for network identification; consequently, 21 datapoints were collected per experiment.

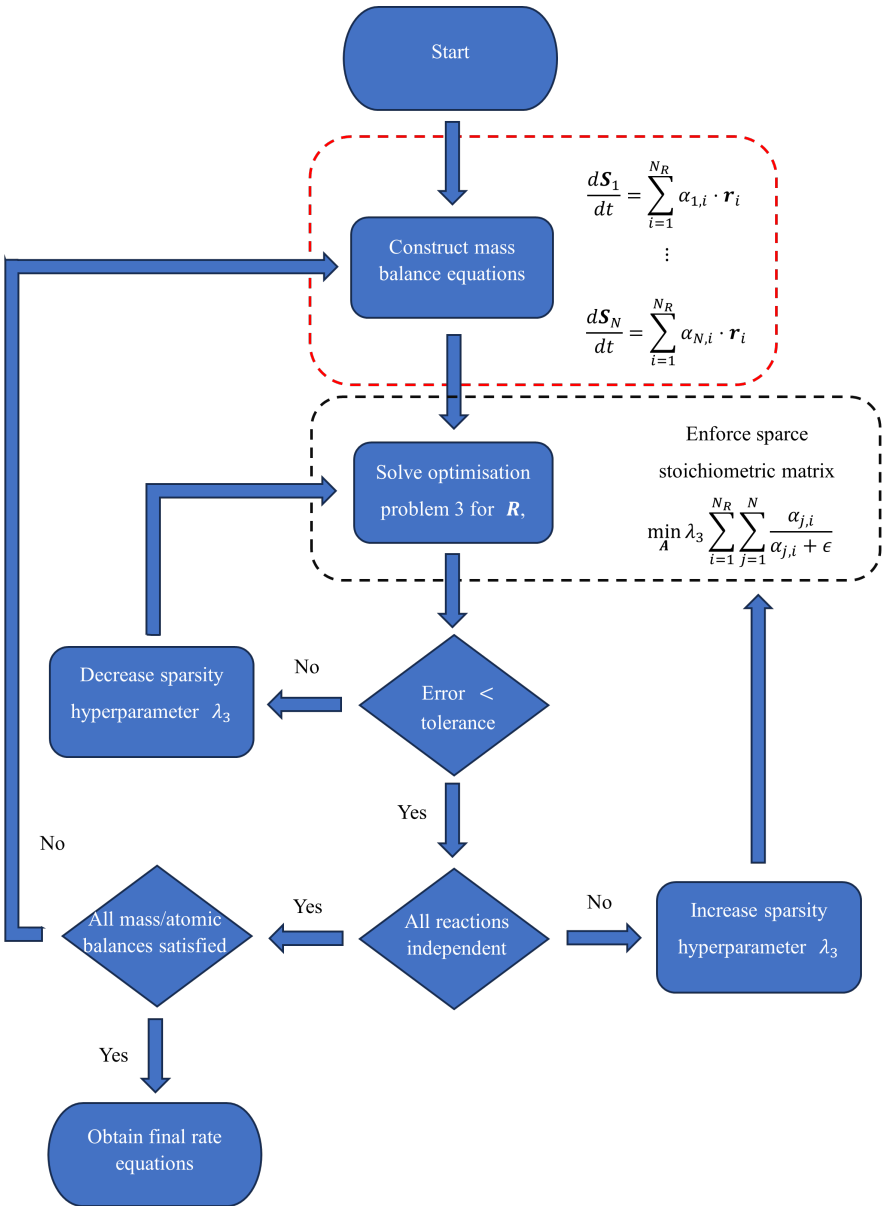

Figure 2: Flowchart representation of the methodology for identifying the minimum size global reaction network

METHODOLOGY

This section will summarise the development of the data-driven framework for autonomous reaction network identification and reaction rate model generation. Explicitly, we will first detail the method used for ascertaining the minimum number of global reactions in addition to the active species involved, then we will show the iterative symbolic regression (SR) strategy employed for the construction and refinement of representative rate equations as illustrated in Figure 1.

To reduce the associated computational expense, we propose a 2-stage workflow. Firstly, we exploit the experimental data to formulate a non-linear programming problem, subject to mass balance constraints, that yields the active reactions, their net reaction rates, $\boldsymbol{R} = \{r_1, r_2 \dots\}$, and the stoichiometric matrix, $\boldsymbol{A}$. Next SR is employed in conjunction with MBDoE in an iterative fashion to produce informative experiments that can facilitate model discrimination and improve the SR results. The iterations continue until no improvements are observed in the pareto front of proposed model candidate or until a threshold criterion is met.

Reaction Network Identification

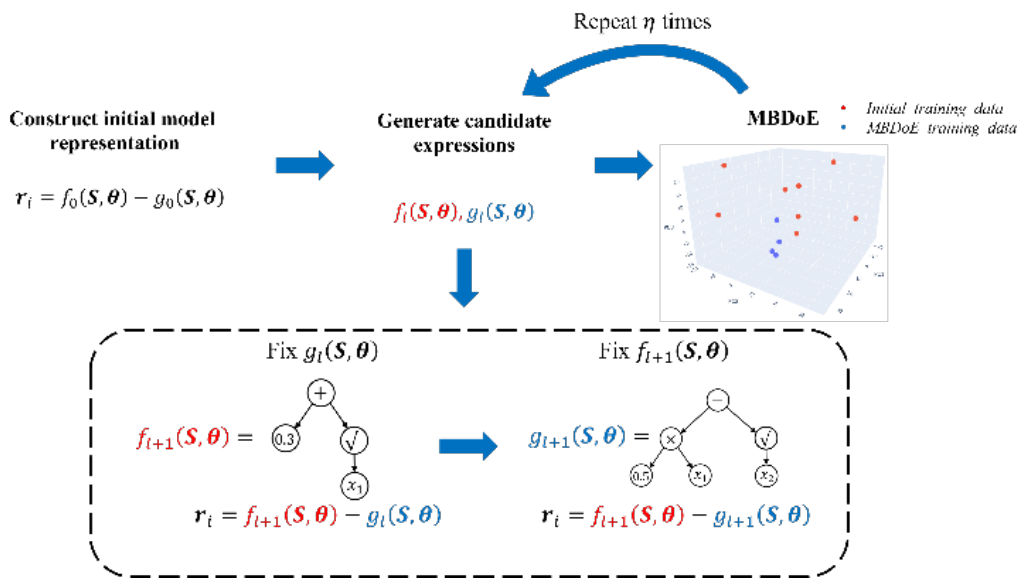

Figure 3. Diagram of the iterative approach used for generating rate expressions from data across η global iterations of the framework at local iteration, l, of tuning the SR model structure.

A chemical reaction network (CRN) is described by a set of characteristics, being the number of global reactions, N_R, the total number of active species, N, and the stoichiometric matrix, $\boldsymbol{A}$, which captures the interactions among all species within the network. Furthermore, the CRN must adhere to linear algebra constraints such that $N_R \leq N$ [15] because there cannot exist more independent reactions than active species. Using this information, we can derive an optimisation problem of which the solution details the mass balance of species, $\boldsymbol{S}$, as a function of net reactions, $\boldsymbol{R}$, across N_e experiments. The formulation used, given a single experiment n, is shown in Equation 3

$$\begin{aligned} \frac{dS_1}{dt} &= \sum_{i=1}^{N_R} \alpha_{1,i} r_i \\ &\vdots \\ \frac{dS_N}{dt} &= \sum_{i=1}^{N_R} \alpha_{N,i} r_i \end{aligned} \tag{3}$$

where $\alpha_{N,i}$ is the stoichiometric coefficient for an individual species within net reaction, i.

Once the reaction rates are determined, it is imperative to ensure they are independent from one another. Many techniques exist to this purpose; however, we employ the use of sparse optimisation such that we can simultaneously minimise the total number of identified net reactions and find the corresponding rate profiles. The methodology for this procedure is outlined in Figure 2.

In order to find the minimum dimension reaction network that well fits experimental data, we conduct parameter estimation using the modified least squares objective function in Equation 4, subject to mass and atomic balance constraints.

$$\min_{\boldsymbol{R}(t),\boldsymbol{A}} \sum_{n=1}^{N_e} (\boldsymbol{S}_n - \boldsymbol{S'}_n(\boldsymbol{R}(t),\boldsymbol{A}))^2 + \lambda \sum_{i=1}^{N_R} \sum_{j=1}^{N} \frac{\alpha_{j,i}}{\alpha_{j,i} + \epsilon} \tag{4}$$

Here $\boldsymbol{S}$ and $\boldsymbol{S'}$ are the measured and simulated species concentrations, λ is a weighting parameter, and the term $\frac{\alpha_{j,i}}{\alpha_{j,i}+\epsilon}$ encourages model sparsity by minimising the rank of the stoichiometric matrix.

Reaction Rate Model Generation

Once the global reaction rates have been identified, a SR based approach is implemented to automatically generate representative reaction rate equations directly from the estimated rate profiles, using available literature to exclude terms with no physical relevance to the system under study. We overcome the limitation of SR for complex systems by introducing a substructure decomposition strategy that breaks down the global reaction rates into various terms, each with physical meaning (e.g., forward and reverse reaction rates), as shown in Figure 3.

As illustrated in Figure 3, after the initial candidate expression has been suggested by SR, the individual substructures are then iteratively updated and fine tuned and, once this local iteration completes, MBDoE is performed by maximising the variance between the predictions of the pareto front of SR candidate functions.

RESULTS AND DISCUSSION

Results of the Methanol Synthesis Reaction Network

As described in the case study section, the methanol synthesis via CO_2 reaction network consists of 6 species in a closed system, meaning that there is a maximum of 5 global reactions. By deriving the problem formulation in Equation 3 and using the sparse optimisation methodology in Figure 2 and Equation 4 to solve the parameter estimation framework, we obtain a set of independent reaction rates, $\boldsymbol{R}$, and stoichiometric coefficients, $\boldsymbol{A}$. Upon rescaling the values of the stoichiometric coefficients to ensure they are all of integer values, we are left with the reduced reaction network in Equation 5

$$\begin{aligned} \frac{dF_{CO}}{d\tau} &= r_2 \\ \frac{dF_{CO_2}}{d\tau} &= -r_1 - r_2 - r_3 \\ \frac{dF_{H_2}}{d\tau} &= -3r_1 - r_2 - 4r_3 \\ \frac{dF_{CH_3OH}}{d\tau} &= r_1 \\ \frac{dF_{H_2O}}{d\tau} &= r_1 + r_2 + 2r_3 \\ \frac{dF_{CH_4}}{d\tau} &= r_3 \end{aligned} \tag{5}$$

where F_j is the molar flow rate of species j and τ is the residence time.

When equated to the ground truth in reaction network R1, it can be seen that they are identical. The use of mass and atomic balance constraints allowed for side reactions with minimal influence to be discovered, as is true of the Sabatier reaction, r_3, which holds reaction rates 10-100 times less that of r_1 and r_2.

With the establishment of a valid reaction network, the second stage of the workflow can proceed, enabling the creation of reaction rate models from data. As symbolic regression leverages genetic algorithms to design candidate expressions, it is necessary to introduce approaches to reduce the search space that must be explored to improve solution quality. This can be achieved in several manners, namely the use of structural constraints and imposing restrictions on the list of operators, $\varphi = \{+, -, \times, \div, x_a\}$, and inputs. Within the literature surrounding heterogenous catalysis, it is commonplace to use Langmuir-Hinshelwood type structures to define the kinetics present and, as such, we can enforce SR to abide by the structural form as shown Equation 6.

$$r_i = \frac{f_i(\boldsymbol{S}, \theta) - g_i(\boldsymbol{S}, \theta)}{(1 + h_i(\boldsymbol{S}, \theta))^n} \tag{6}$$

In this manner, the proposed equations will be more interpretable, and we can decompose the LH models based on their definition, where the numerator relates to the thermodynamic driving force and the denominator represents the kinetic resistance. A further advantage stems from stage 1 of the workflow. Since a complete reaction network has been attained, the input features can be curated based on these results, as we have knowledge of which species participate within each reaction, including within the different substructures (forward rate, reverse rate and kinetic inhibition).

Following the iterative SR procedure outlined in Figure 3, 8 further experiments were designed and incorporated into the algorithm, with the best performing models of the final iteration listed in Table 1.

Table 1: Top ranked discovered Langmuir-Hinshelwood type equations from the iterative SR procedure in Figure 3. $a_{j,i}$ is a parameter of model i.

Rate	Equation
r_1	$\dfrac{\alpha_{1,1}P_{CO_2}P_{H_2}^2 - \alpha_{1,2}P_{CH_3OH}P_{H_2O}}{P_{H_2}(1 + \alpha_{1,3}P_{CO_2}P_{H_2} + \alpha_{1,4}P_{H_2} + \alpha_{1,5}P_{CO_2} + \alpha_{1,6}P_{CO_2}^2)}$
r_2	$\dfrac{\alpha_{2,1}P_{CO_2}P_{H_2} - \alpha_{2,2}P_{CO}P_{H_2O}}{(\alpha_{2,3} + \alpha_{2,4}P_{CO_2} + \alpha_{2,5}P_{H_2})(\alpha_{2,6} + \alpha_{2,7}P_{CO_2} + \alpha_{2,8}P_{H_2})}$
r_3	$\dfrac{\alpha_{3,1}P_{H_2}(1 - \frac{\alpha_{3,2}P_{CH_4}P_{H_2O}^2}{P_{H_3}^3})}{(\alpha_{3,3} + \alpha_{3,4}P_{CO_2} + \alpha_{3,5}P_{H_2})(\alpha_{3,6} + \alpha_{3,7}P_{CO_2} + \alpha_{3,8}P_{H_2})}$

From an analysis of Table 1, there exists considerable equivalence to the ground truth structures in Equation 1, indicating that the methodology was capable of learning crucial insight from the data regarding the kinetics. Specifically, the algorithm discerned that, not only do H_2 and CO_2 individually affect the kinetic resistance, but their interactions also have a substantial contribution. Additionally, in the methanol synthesis (r_1) and water gas shift reactions (r_2), the influence of CO_2 the on the forward rate and CH_3OH, CO, and H_2O on the reverse rates were correctly identified as first order. Although, for r_1, the order of P_{H_2} was determined to be 2 instead of 3, the dominant impact of H_2 in the forward reaction rate was uncovered correctly.

As well as providing deep insight into the kinetics of the reaction, the parameterisation of the models in Table 1 yielded low average fitting percentage errors of 3.3% as shown in Figure 4. In order to assess the predictive capabilities across the entire design space, an experiment was designed to discriminate between the generated SR models and the ground truth. The maximum

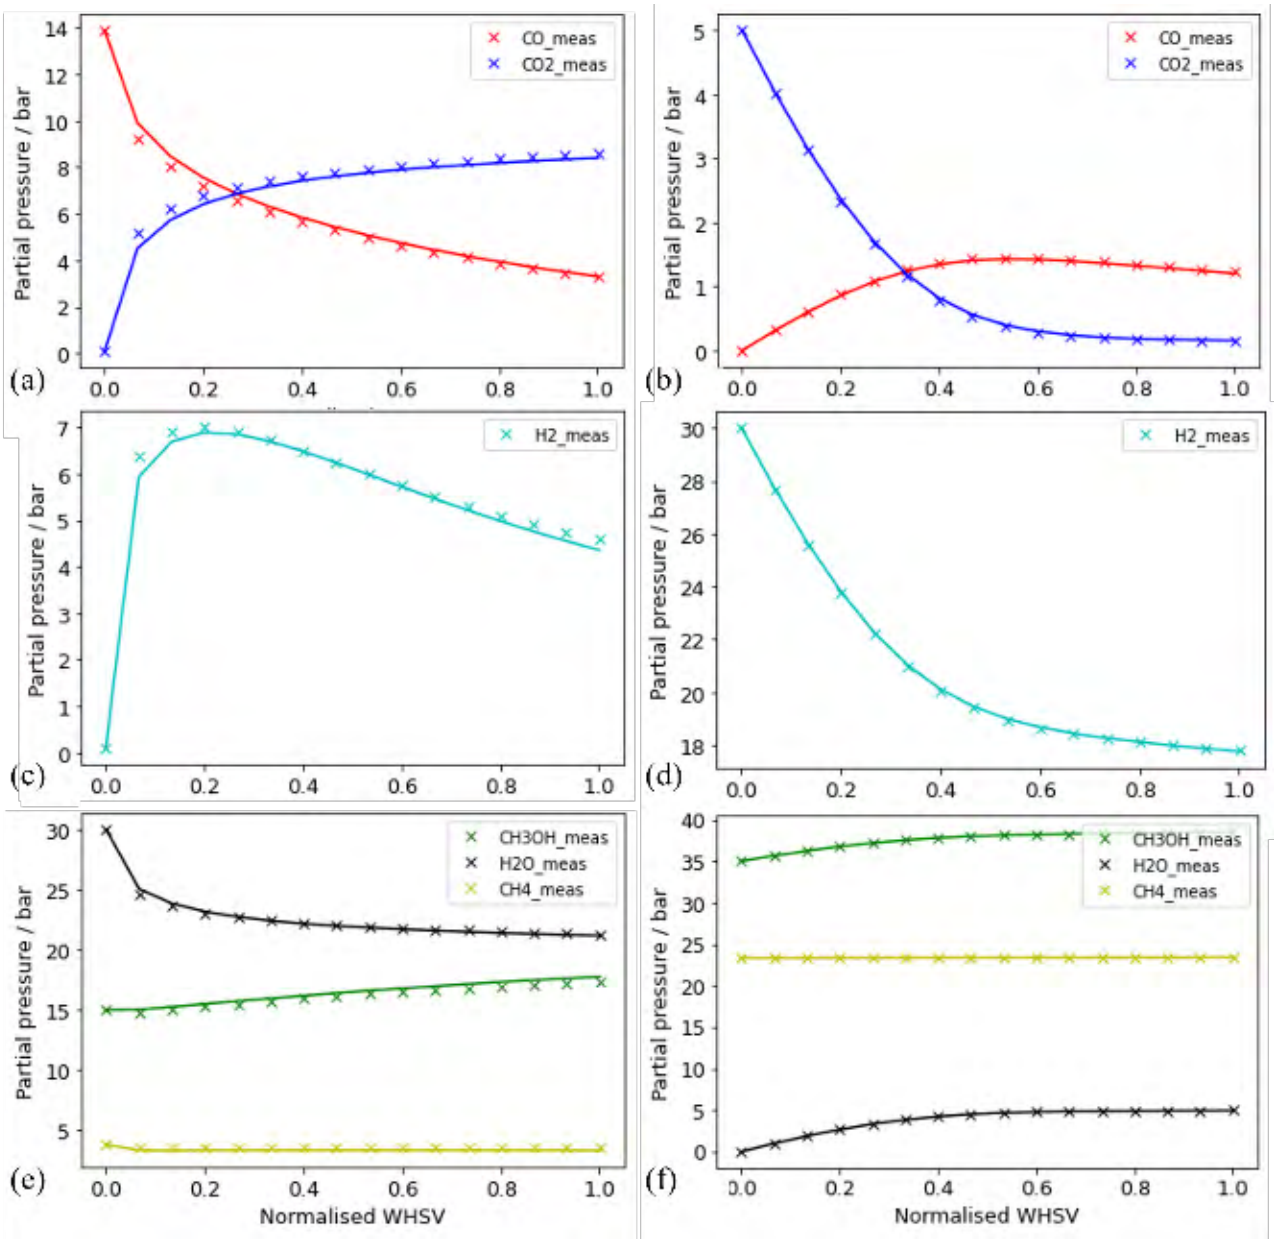

Figure 4. Fitting results of the top ranked SR generated models on the MBDoE experiments. (a), (c), and (e) refer to a single experiment and (b), (d), and (f) to another.

average percentage error across the entire operational region was 16.9%, demonstrating that the framework is robust and data efficient, being able to capture highly complex underlying kinetic information with a limited experimental budget.

Results of the Kinetic Resolution Network

In a similar manner to the methanol synthesis reaction network, we can derive the problem formulation in Equation 3 using 9 active species and therefore a maximum of 8 independent reactions. After solving the problem formulation, it was discovered that only 5 global reactions are required to fit the experimental data, with the final reaction network presented in Equation 7.

$$
\begin{aligned}
r_1&: VA \rightleftharpoons \mathrm{AE} + AA \\
r_2&: AE + S_1 \rightarrow S_2 \\
r_3&: VA + S_2 \rightarrow AA + S_3 \\
r_4&: AE + R_1 \rightarrow R_2 \\
r_5&: AE + R_2 \rightarrow R_3
\end{aligned}
\tag{7}
$$

Through comparing the results in Equation 7 to the ground truth (R2), it is possible to observe that all the

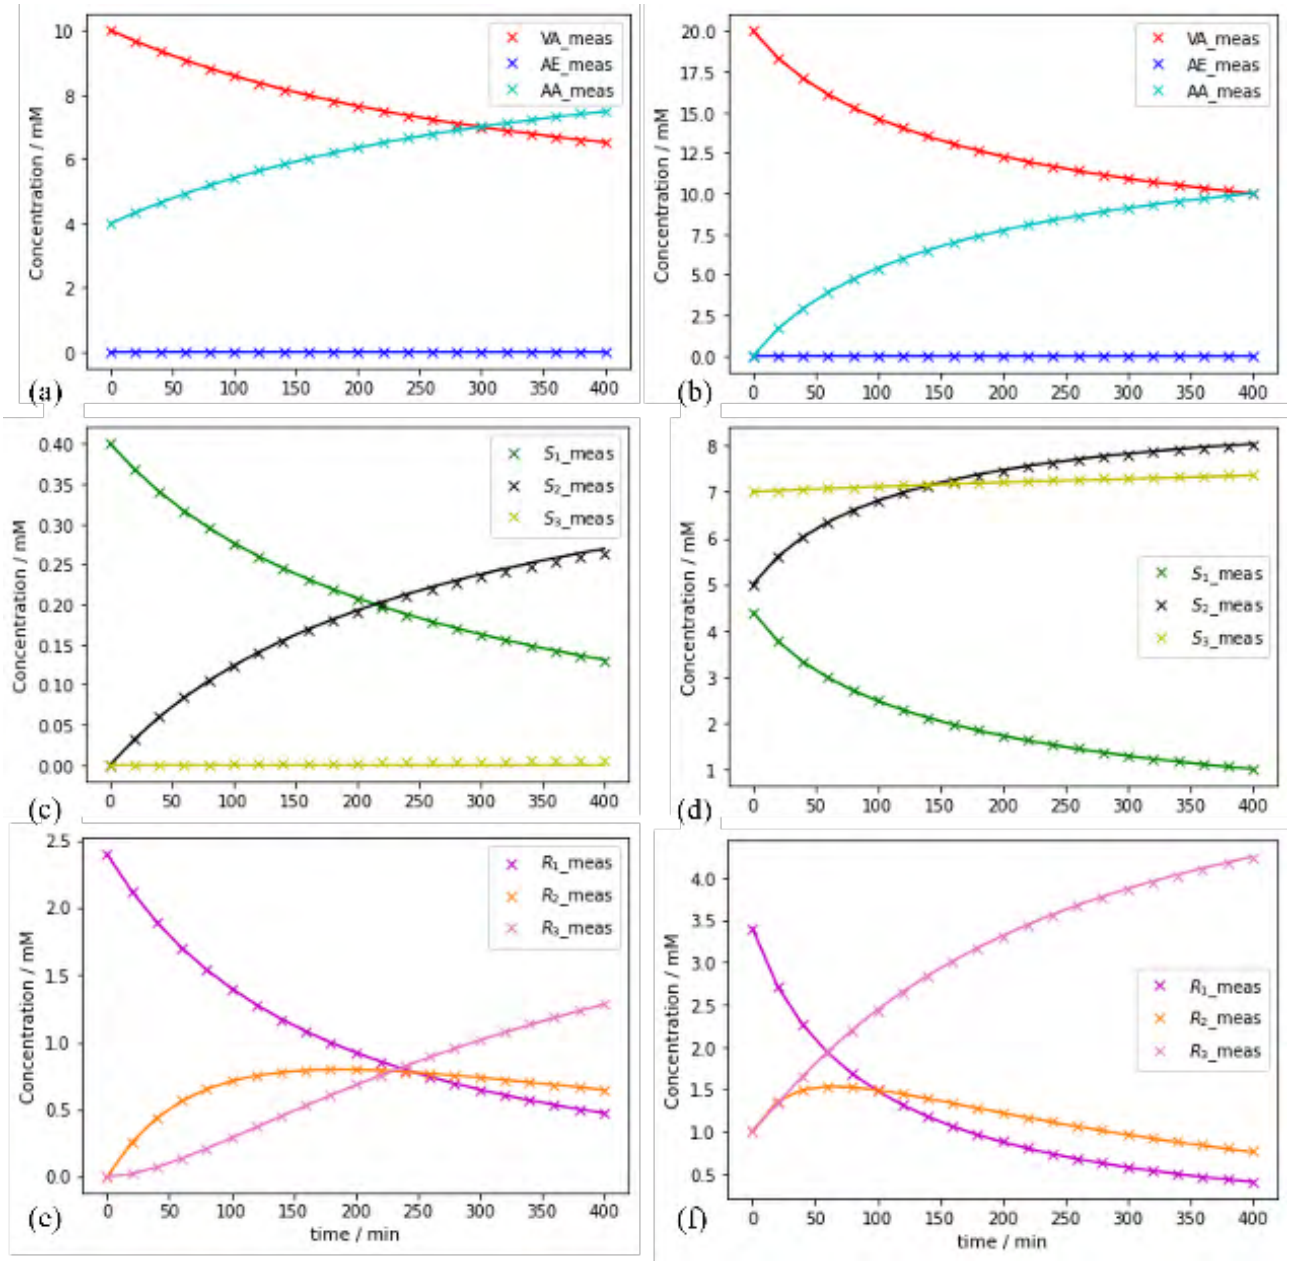

Figure 5. Fitting results for the parameterised SR models in Table 2. (a), (c), and (e) refer to a single experiment and (b), (d), and (f) to another.

species that are involved within the reactions as well as their respective stoichiometric coefficients were found with the sole inconsistency being in r_3, where VA and AA were included in place of AE. Notably, this is a lumped representation of the r_1 and r_3 in the ground truth reactions and still satisfies all mass balances. Therefore, the result is physically valid. It is hypothesised that the lumped representation was found due to the high dimensional nature of the optimisation problem requiring over 240 parameters to be solved for. This issue is further complicated by the low relative magnitudes of reaction rates for r_3 being 10-1000 times smaller than the reaction rates of all other reactions.

The accurate determination of four global reaction rates provides affirmation of the accuracy and reliability of the developed methodology. Applying the parameterised reaction network in Equation 7 to the initial 4 experiments used for network identification yields minimal percentage errors ($\approx 0\%$), both proving the effectiveness of the strategy, yet also a limitation where recovered reactions may, in actuality, be a lumped representation of the several ground truth reactions.

With the construction of a suitable reaction network, SR can now be applied as dictated in Figure 3. Similar to the previous case study, the literature can be exploited to add suitable constraints to the inputs, operators $\boldsymbol{\varphi} = \{+, -, \times, \div, x_a\}$ and the equations structure. Succeeding the first iteration of running the framework, it was surmised that no further MBDoE was needed due to the correct rate models being recovered immediately for all reactions except r_3, as shown in Table 2, with a fitting percentage error of 3.0%.

Table 2: Top ranked rate models discovered for the kinetic resolution case study.

Rate	Equation
r_1	$\alpha_{1,1}C_{VA} - \alpha_{1,2}C_{AA}C_{AE}$
r_2	$\alpha_{2,1}C_{S_1}C_{AE}$
r_3	$\alpha_{3,1}C_{S_3}C_{VA}$
r_4	$\alpha_{4,1}C_{R_1}C_{AE}$
r_5	$\alpha_{5,1}C_{R_2}C_{AE}$

From Table 2, the interpretability of the generated model structures is apparent, indicating first order dependencies with respect to any species involved, and that all reaction rates possess irreversible behaviour apart from r_1. Finally, in both case studies, none of the proposed reaction rate models exhibited complexities beyond that of the ground truth kinetics, providing strong evidence that the framework is capable of yielding highly interpretable and precise representations that describe the underlying reaction networks.

CONCLUSION

The case studies presented demonstrate that the proposed two-stage framework is capable of accurately recovering underlying reaction networks and automatically constructing interpretable kinetic rate expressions consistent with known physics, operating effectively under a limited experimental budget. Furthermore, the novel iterative application of symbolic regression mitigates the risk overly complex model forms, increasing the practical applicability of SR to real-world reaction systems whilst preserving high levels of interpretability.

Importantly however, the present study is conducted assuming complete observability of all reacting species and noise-free measurements, which are rarely fully satisfied in industrial practice. Future research should therefore assess the robustness of the framework when applied to noisy and sparsely sampled experimental datasets. Nevertheless, when combined with human-in-the-loop guidance, the methodology shows considerable promise for gaining mechanistic insight to support expert analysis. By incorporating augmented intelligence within the workflow, the framework offers a pathway towards accelerated knowledge generation and automated kinetic modelling for complex reaction engineering applications.

REFERENCES

1. Becker C. From langmuir to ertl: the “nobel” history of the surface science approach to heterogeneous catalysis. Encyclopedia of Interfacial Chemistry :99-106 (2018). https://doi.org/10.1016/b978-0-12-409547-2.13527-9
2. Marin GB, Yablonsky GS, Constales D. Kinetics of chemical reactions. Wiley (2018). https://doi.org/10.1002/9783527808397
3. Motagamwala AH, Dumesic JA. Microkinetic modeling: a tool for rational catalyst design. Chem. Rev. 121:1049-1076 (2020). https://doi.org/10.1021/acs.chemrev.0c00394
4. Nakajima EA, Oliveira LG, Gasparrini LJ, Souza GEQ, Ignacio AA, Alves HJ, Borba CE. Kinetics of dry reforming of methane catalyzed by ni/si-mcm-41. International Journal of Hydrogen Energy 48:32331-32341 (2023). https://doi.org/10.1016/j.ijhydene.2023.05.010
5. Armenise, S.; García-Bordejé, E.; Valverde, J.; Romeo, E.; Monzón, A. A Langmuir–Hinshelwood approach to the kinetic modelling of catalytic ammonia decomposition in an integral reactor. Phys.
6. Ali SH. Kinetics of catalytic esterification of propionic acid with different alcohols over amberlyst 15. Int J of Chemical Kinetics 41:432-448 (2009). https://doi.org/10.1002/kin.20416
7. Mendoza-Covarrubias C, Romero CE, Hernandez-Rosales F, Agarwal H. N₂o formation in selective non-catalytic no_x reduction processes. JEP 02:1095-1100 (2011). https://doi.org/10.4236/jep.2011.28126
8. FISHTIK I, DATTA R. A NEW APPROACH FOR THE CLASSIFICATION AND ENUMERATION OF UNIQUE REACTION ROUTES AND UNIQUE OVERALL REACTIONS IN MULTIPLE CHEMICAL REACTION SYSTEMS. Chemical Engineering Communications 191:373-397 (2004). https://doi.org/10.1080/00986440490272555
9. Miskovic L, Béal J, Moret M, Hatzimanikatis V. Uncertainty reduction in biochemical kinetic models: enforcing desired model properties. PLoS Comput Biol 15:e1007242 (2019).

https://doi.org/10.1371/journal.pcbi.1007242
10. Taylor CJ, Seki H, Dannheim FM, Willis MJ, Clemens G, Taylor BA, Chamberlain TW, Bourne RA. An automated computational approach to kinetic model discrimination and parameter estimation. React. Chem. Eng. 6:1404-1411 (2021). https://doi.org/10.1039/d1re00098e
11. Taylor CJ, Booth M, Manson JA, Willis MJ, Clemens G, Taylor BA, Chamberlain TW, Bourne RA. Rapid, automated determination of reaction models and kinetic parameters. Chemical Engineering Journal 413:127017 (2021). https://doi.org/10.1016/j.cej.2020.127017
12. Ghosh S, Sebastian J, Olsson L, Creaser D. Experimental and kinetic modeling studies of methanol synthesis from CO2 hydrogenation using in2o3 catalyst. Chemical Engineering Journal 416:129120 (2021). https://doi.org/10.1016/j.cej.2021.129120
13. Gustavo MF, Székely E, Tóth J. Kinetic modeling of a consecutive enzyme-catalyzed enantioselective reaction in supercritical media. ACS Omega 5:26795-26806 (2020). https://doi.org/10.1021/acsomega.0c02405
14. Dong J, Zhong J. Recent advances in symbolic regression. ACM Comput. Surv. 57:1-37 (2025). https://doi.org/10.1145/3735634
15. Feinberg M. Species-reaction graph foundations. Applied Mathematical Sciences :419-440 (2019). https://doi.org/10.1007/978-3-030-03858-8_19

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Unveiling Reaction Patterns in Thermal and Catalytic Biomass Pyrolysis Using PCA and Multivariate Analysis

Martín Rodríguez-Fragoso[a], Sandro González-Arias[a], Octavio Elizalde-Solis[a], and Edgar Ramírez-Jiménez[a]*

[a] Department of Chemical Petroleum Engineering, ESIQIE, Instituto Politécnico Nacional, Mexico City, 07738, Mexico
* Corresponding Author: eramirezj@ipn.mx.

ABSTRACT

Understanding the relationships between operating conditions and product formation pathways in biomass pyrolysis remains challenging due to the complex interactions among temperature, catalytic effects, and feedstock composition. In this work, principal component analysis (PCA) was applied to investigate the combined influence of temperature and catalyst-to-biomass ratio on the pyrolysis of sugarcane bagasse and Salicornia. To preserve mechanistic interpretability, two complementary analyses were performed: one considering only catalytic experiments and a second integrating both thermal and catalytic conditions. Separate PCA were conducted for product yields, gas and liquid compositions, and solid-phase FTIR features. The results reveal that thermal conditions promote severe cracking and solid carbonization, whereas catalytic operation favors secondary pathways associated with controlled dehydration and partial stabilization of liquid products. Distinct patterns between the two feedstocks were also identified, reflecting their intrinsic compositional differences. Based on the multivariate trends identified, a conceptual reaction scheme is proposed to describe the interplay between thermal and catalytic pathways.

Keywords: Biomass, Multivariate Statistics, Reaction, Reaction Patterns

INTRODUCTION

Biomass pyrolysis is a promising thermochemical route to produce biofuels, value-added chemicals, and carbonaceous materials from renewable resources. Lignocellulosic biomass, due to its abundance and carbon-neutral nature, has attracted significant attention as a feedstock for sustainable energy systems. Nevertheless, the intrinsic complexity of biomass and the large number of simultaneous reactions occurring during thermal conversion hinder a detailed understanding of the reaction mechanisms and the control of product distribution [1-2].

From a mechanistic standpoint, biomass pyrolysis is generally described as the superposition of three main primary pathways: char formation, depolymerization, and fragmentation of the macromolecular constituents. These reactions are accompanied by secondary processes such as cracking, recombination, and recondensation of volatile species, which may lead to the formation of low- and high-molecular-weight products as well as secondary char. The relative contribution of these pathways strongly depends on temperature and reactor conditions, resulting in a highly interconnected and nonlinear reaction network [1-2].

The pyrolysis behavior of lignocellulosic biomass is governed by the relative proportions and interactions of cellulose, hemicellulose, and lignin. However, several studies have shown that the pyrolysis of whole biomass cannot be accurately described as a simple superposition of its individual components. Interactions among the main constituents significantly affect product yields and reaction pathways, and their detailed mechanisms remain not fully understood [3].

As an alternative, data-driven methodologies have gained increasing attention for the analysis of complex pyrolysis datasets. Principal Component Analysis (PCA) has been successfully employed to identify dominant patterns, correlate product families, and distinguish between primary and secondary reactions, supporting the proposal of reaction or kinetic schemes without imposing *a priori* assumptions [4].

In this work, an integrated PCA-based approach is

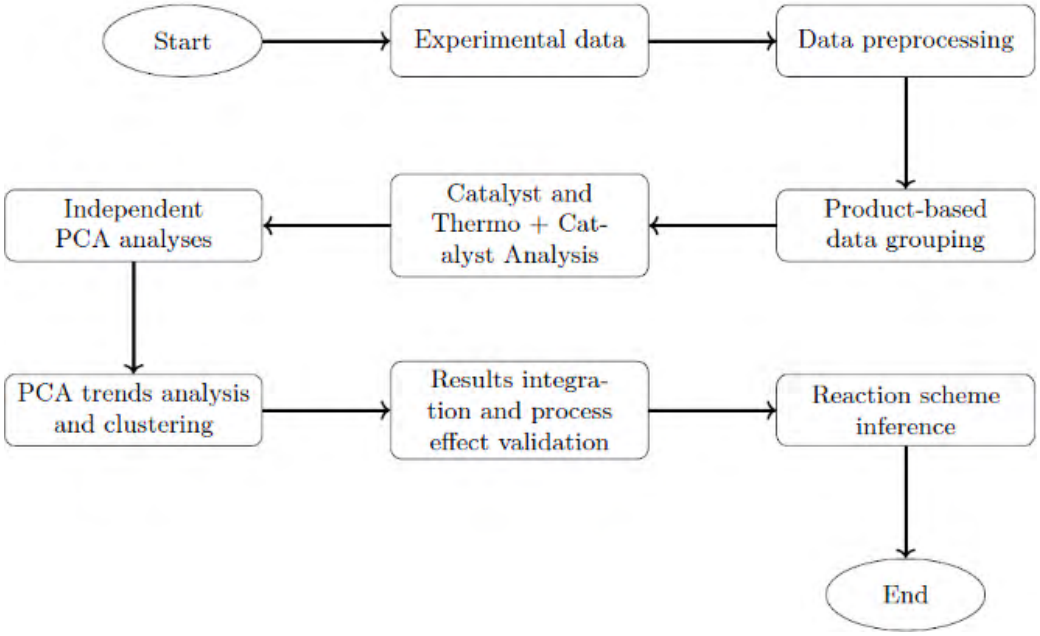

Figure 1. Flowchart of PCA methodology.

applied to investigate the catalytic and thermal pyrolysis of lignocellulosic biomass by combining complementary experimental data sets, including global product yields, gas chromatography, and FTIR analysis of solid residues. Independent PCA analyses are performed and subsequently interpreted as a function of temperature and C/B ratio, together with clustering techniques to identify distinct reaction regimes.

EXPERIMENTAL SETUP

Two biomasses were investigated: sugarcane bagasse and Salicornia sp. Sugarcane bagasse (SCB) was obtained from the Tierra Blanca Sugar Mill (Veracruz, Mexico), while Salicornia (SAA) was collected from cultivated fields in Baja California Sur, Mexico. Both materials were used as feedstock for thermal and catalytic pyrolysis experiments.

Pyrolysis reactions were carried out in a semi-continuous fixed-bed tubular reactor made of titanium, using 1 g of biomass per experiment. Thermal pyrolysis was conducted in the absence of catalyst, whereas catalytic pyrolysis experiments were performed using catalyst-to-biomass (C/B) mass ratios of 1:1, 1:2, and 1:3. For both biomasses, experiments were conducted at reaction temperatures ranging from 400 to 600 °C, while the residence time was kept constant. Nitrogen was continuously fed from the bottom of the reactor to ensure an inert atmosphere.

The experimental setup consisted of a heating furnace, the tubular reactor, a condensation column with a recirculating cooling system, a round-bottom flask for liquid recovery, and a gas collection system. Condensable volatiles were recovered using ethylene glycol as cooling medium at 5 °C, while non-condensable gases were collected in inert multilayer foil sampling bags.

The experimental campaign generated four types of data: (i) overall product yields, (ii) gas-phase composition determined by gas chromatography, (iii) liquid-phase composition analyzed by gas chromatography, and (iv) solid-phase FTIR spectra. Gaseous products were analyzed using a GC-7890B gas chromatograph equipped with a Molecular Sieve 13X column and argon as carrier gas. Liquid products were analyzed using a Clarus 500 gas chromatograph. Solid residues were characterized by FTIR spectroscopy using a PerkinElmer Precisely Spectrum 100 spectrometer.

METHODOLOGY

The methodology implemented in this work follows a product-based multivariate analysis strategy, as summarized in Fig. 1. The approach integrates experimental data from different product phases with structured data preprocessing and Principal Component Analysis (PCA) to identify dominant trends and infer reaction pathways during biomass pyrolysis.

Data Preprocessing

All data processing was performed using `Python`, employing standard scientific libraries (`NumPy`, `Pandas`, and `scikit-learn`)[5-7].

Overall product yields, gas-phase compositions, liquid-phase grouped variables, and solid-phase FTIR

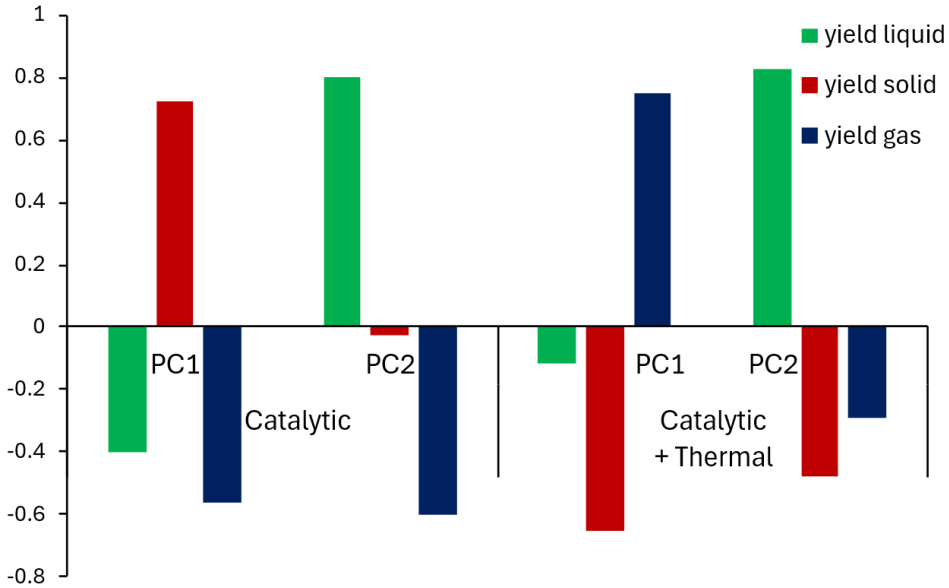

Figure 2. Bar plots of yield loadings for each data group.

descriptors were scaled using z-score normalization to ensure comparable variable variance prior to multivariate analysis.

Gas-phase data consisted of the molar fractions of the main species (CH_4, CO_2, C_2H_6, C_3H_8, C_4H_{10}, H_2, and CO), which were directly used as input variables. Liquid-phase chromatographic data were first grouped into chemical families (alcohols, ketones, methylphenols, nitrophenols, and chlorophenols) to reduce dimensionality and improve chemical interpretability.

For solid-phase analysis, FTIR spectra were divided into predefined wavenumber intervals (2700–2350, 2350–2100, 2100–1800, 1800–1500, and 1300–850 cm^{-1}). The area under each interval was calculated using numerical integration, providing quantitative descriptors associated with major functional groups in the solid residues.

Product-based PCA strategy

PCA was applied independently to each product-based dataset using the `scikit-learn` implementation. In all cases, the first two principal components were retained, as they captured the dominant variance and allowed visualization of process trends. To enhance discrimination of reaction regimes, the data were separated into two subsets: catalytic-only and combined thermal–catalytic experiments. Independent PCA models were constructed for each subset, the analysis resulted in a total of eight PCA models: four corresponding to catalytic experiments and four corresponding to the thermal–catalytic dataset.

For each model, the loading patterns of PC1 and PC2 were analyzed to assign physical and chemical meaning to the principal components, while score plots were used to evaluate the evolution of these behaviors as a function of temperature, C/B ratio, and biomass type. Clustering in the score space was further assessed using the `KMeans` algorithm to identify groups exhibiting similar process behavior.

RESULTS AND DISCUSSION

Table 1 summarizes the explained variance of the first two principal components for yields, gas, liquid, and solid (FTIR) datasets under catalytic and combined catalytic–thermal conditions.

Table 1: Explained Variance for each PCA analysis.

	Explained Variance			
	Catalytic		Catalytic + Thermal	
	PC1	PC2	PC1	PC2
Yields	0.639	0.361	0.523	0.477
Gas	0.331	0.237	0.331	0.237
Liquid	0.480	0.303	0.492	0.317
Solid FTIR	0.800	0.175	0.747	0.227

For product yields, the catalytic system is predominantly described by PC1, indicating a dominant reaction pathway. When thermal experiments are included, the variance is more evenly distributed between PC1 and PC2, revealing the coexistence of temperature-driven and catalyst-assisted mechanisms.

Gas-phase data display a similar variance

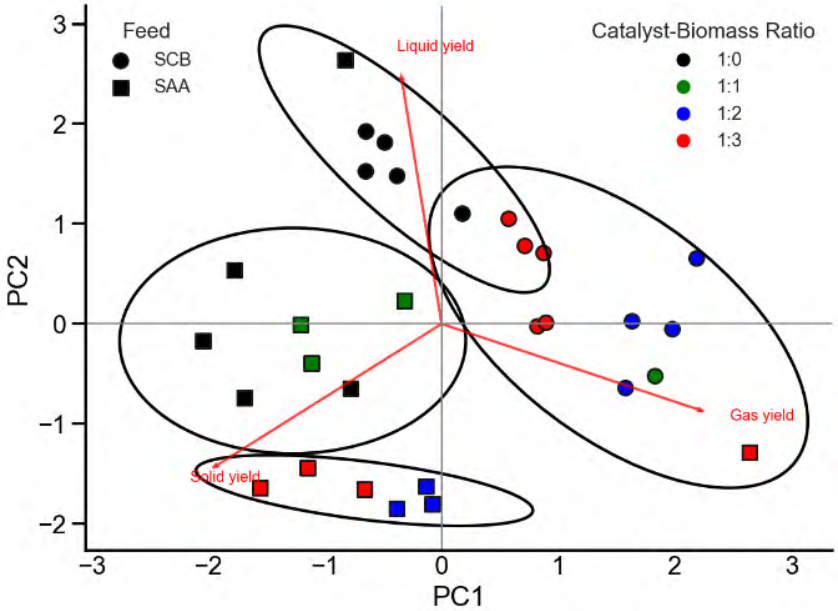

Figure 3. Biplot of thermal + catalytic yields scores.

distribution in both cases, suggesting that gas formation is primarily governed by temperature, with limited sensitivity to catalytic effects. In contrast, liquid-phase data show a more balanced contribution of PC1 and PC2, reflecting the superposition of thermal fragmentation and catalytic stabilization pathways. Solid residues characterized by FTIR are largely dominated by PC1 in both scenarios, indicating a robust primary mechanism associated with char formation and structural evolution. The increased contribution of PC2 upon inclusion of thermal experiments suggests secondary temperature-induced modifications without altering the dominant catalytic signature.

PCA of Product Yields

Loading Analysis

The loadings bar plots for product yields (see Fig. 2) highlight clear differences between the PCA performed using only catalytic experiments and the one including both catalytic and thermal conditions.

For the catalytic dataset, PC1 accounts for most of the variance and reflects the overall redistribution of products driven by secondary catalytic reactions, separating gas-dominated regimes from those favoring liquids and solids. This indicates that, in the presence of a catalyst, product yields are mainly controlled by the efficiency of catalytic cracking, dehydration, and aromatization pathways.

PC2 captures a secondary contrast associated with the balance between liquid stabilization and solid residue formation, characteristic of intermediate catalytic severities.

When thermal experiments are included, the variance structure becomes more evenly distributed between PC1 and PC2. This redistribution indicates that an additional, non-catalytic source of variability is introduced. In this case, PC1 represents a more general conversion severity axis, combining both catalytic effects and temperature-driven primary decomposition. PC2 gains importance as a discriminator between thermally dominated pathways, characterized by early gas formation and poor liquid stabilization, and catalytically assisted regimes, where yields are redistributed toward upgraded liquids and structured solids.

Biplot Analysis

Figs. 3 and 4 shows the PCA biplots of product yields for thermo-catalytic and purely catalytic experiments, with samples colored according to the C/B ratio. These biplots allow visualization of the combined effect of reaction regime and catalyst loading on the distribution of products.

For the thermo-catalytic experiments, most SCB samples are in the (+PC1, +PC2) region, while thermal bagasse experiments tend to shift toward the (−PC1, +PC2) quadrant. In contrast, SAA samples predominantly occupy the (−PC1, −PC2) region, with some thermo-catalytic conditions extending toward (−PC1, +PC2). A consistent downward trend along PC2 is observed with increasing C/B ratio for SAA, whereas for SCB the effect of C/B is more pronounced along PC1, with lower PC1 scores associated with lower catalyst loadings.

When considering only catalytic experiments (see

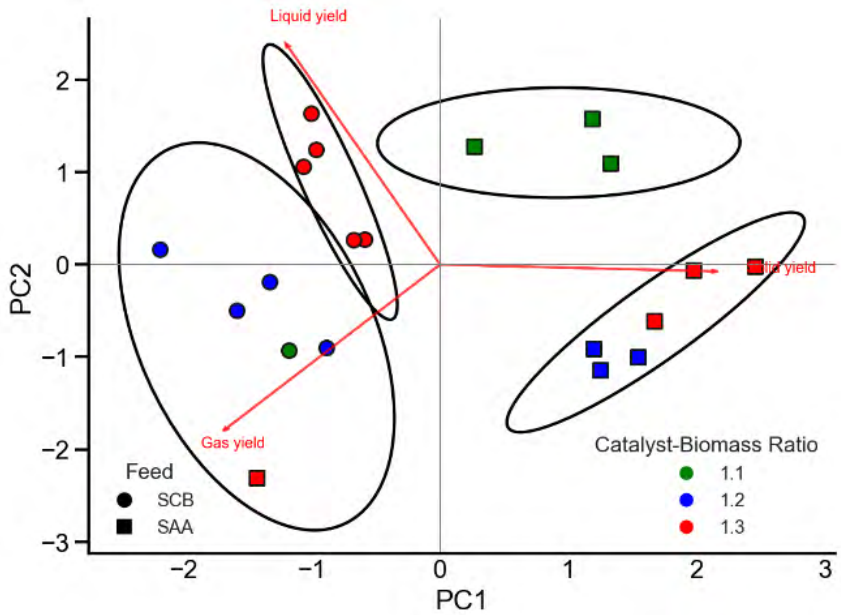

Figure 4. Biplot of only catalytic yields scores.

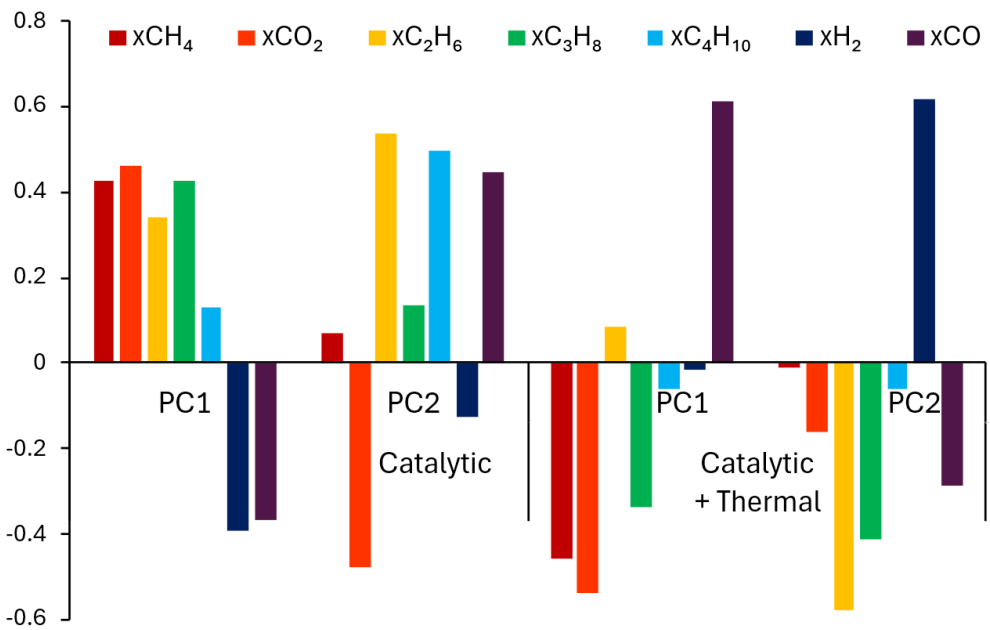

Figure 5. Bar plots of gas loadings for each data group.

Fig. 4), the separation between biomasses and catalyst loadings becomes more defined. For SCB, higher C/B ratios are associated with an upward displacement along PC2, indicating a stronger catalytic influence on yield redistribution. In the case of SAA, samples with lower C/B ratios cluster in the (+PC1, +PC2) region, while higher catalyst loadings shift the samples toward (+PC1, −PC2), with a simultaneous increase along PC2 at the highest C/B values.

Overall, these biplots indicate that the C/B ratio governs the main directions of variance in product yields, but its influence depends strongly on both the biomass type and the reaction regime.

Gas-Phase PCA

Loadings Analysis

The bar plot of the gas-phase PCA loadings (see Fig. 5) reveals a consistent structure between the catalytic

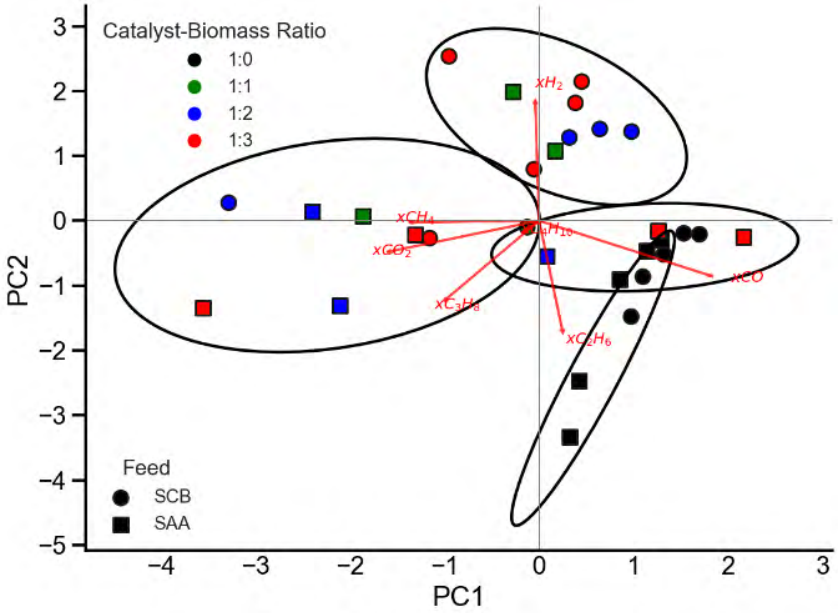

Figure 6. Biplot of thermal + catalytic gas scores.

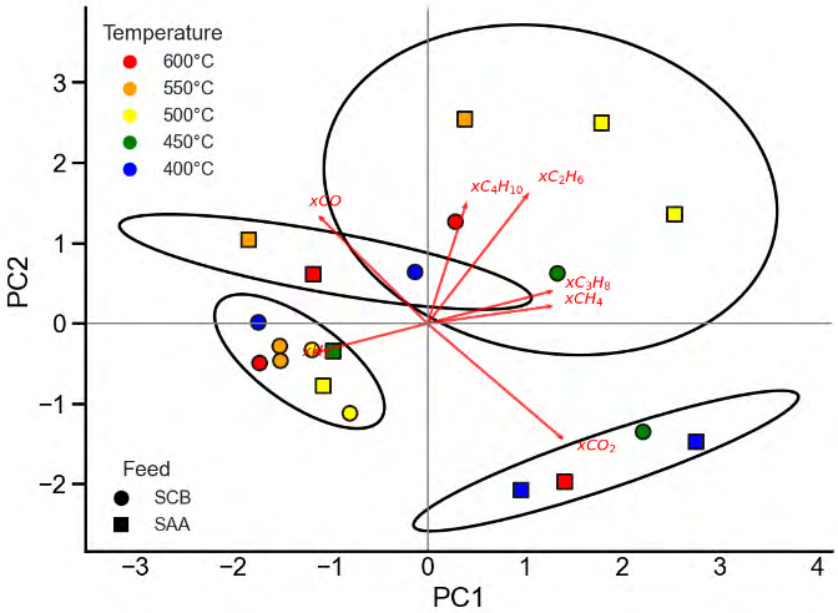

Figure 7. Biplot of only catalytic gas scores.

dataset and the combined catalytic–thermal dataset, indicating that the main gas-forming pathways are intrinsically robust and largely governed by fundamental decomposition and reforming reactions.

In both cases, PC1 is dominated by the major permanent gases, reflecting an axis associated with the extent of gasification and secondary cracking severity. Positive PC1 loadings are linked to increased formation of light gases. This confirms that PC1 represents a global severity coordinate.

PC2 shows a more selective contribution, separating gas species associated with primary thermal devolatilization from those arising from secondary catalytic transformations. In the purely catalytic PCA, PC2 captures subtle shifts in gas selectivity driven by catalyst-assisted pathways, whereas in the combined dataset its relative importance increases, highlighting the growing role of temperature-controlled primary gas release.

The similarity of the loadings pattern indicates that gas-phase composition is less sensitive to the coexistence of thermal and catalytic routes. Instead, gas

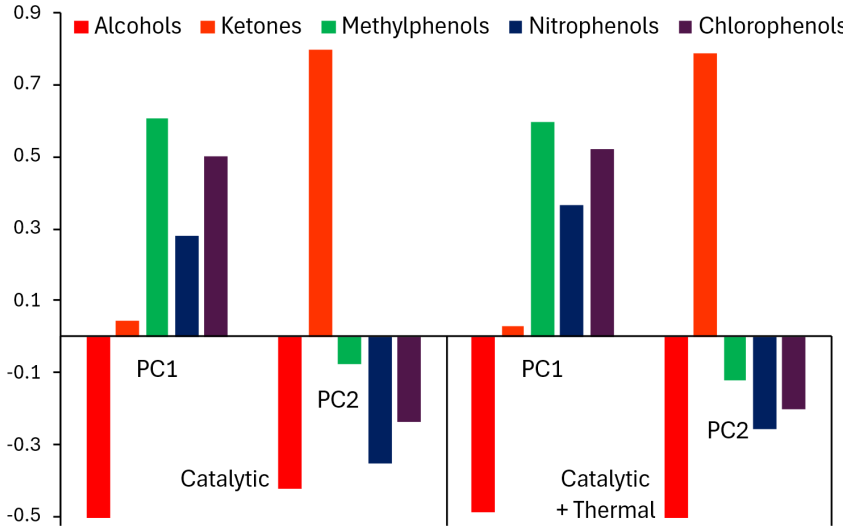

Figure 8. Bar plots of liquid loadings in both groups.

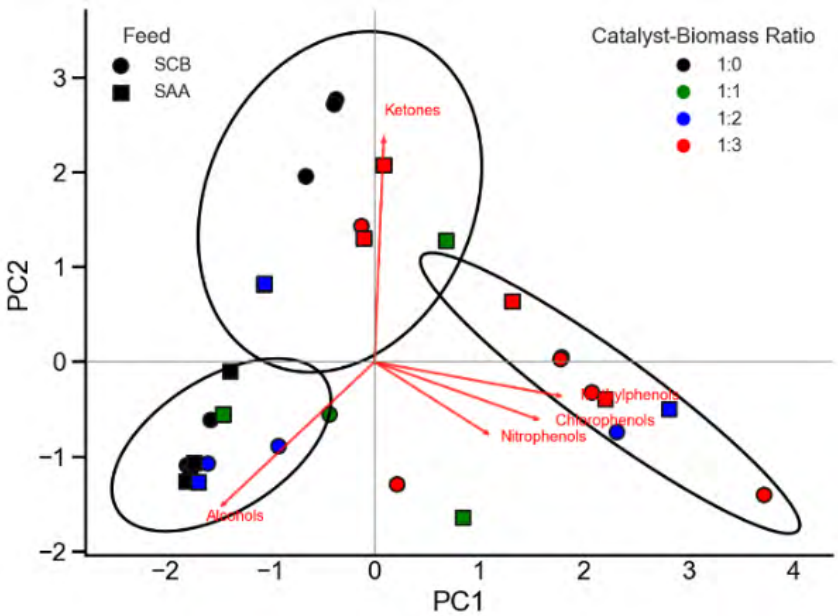

Figure 9. Biplot of thermal + catalytic s liquid scores.

formation responds primarily to the overall intensity of bond scission and reforming reactions, where temperature is the dominant driver and catalyst modulating selectivity.

Biplot Analysis

For the thermo-catalytic dataset, the biplot (see Fig. 6) reveals that thermal experiments are mainly located in the (+PC1, −PC2) region, showing a broader dispersion associated with feedstock type. SCB experiments tend to occupy higher PC2 values, whereas SAA samples are preferentially located at lower PC2 values, indicating distinct gas-phase response patterns between both biomasses under thermal conditions.

In contrast, the catalytic experiments exhibit a clearer organization in terms of the C/B ratio. For SCB, most points are in the (+PC1, +PC2) quadrant, with a systematic upward shift along PC2 as the catalyst loading increases, suggesting an enhanced catalytic contribution to gas-phase transformations. For SAA, although data points remain close to the PC1 axis, an increase in C/B ratio leads to a noticeable displacement toward higher PC1 values, indicating a catalyst-dependent modulation of gas composition.

When considering only catalytic experiments and classifying the data by temperature (see Fig. 7), a strong thermal effect emerges. For SCB, low-temperature experiments cluster in the (+PC1, −PC2) region, while

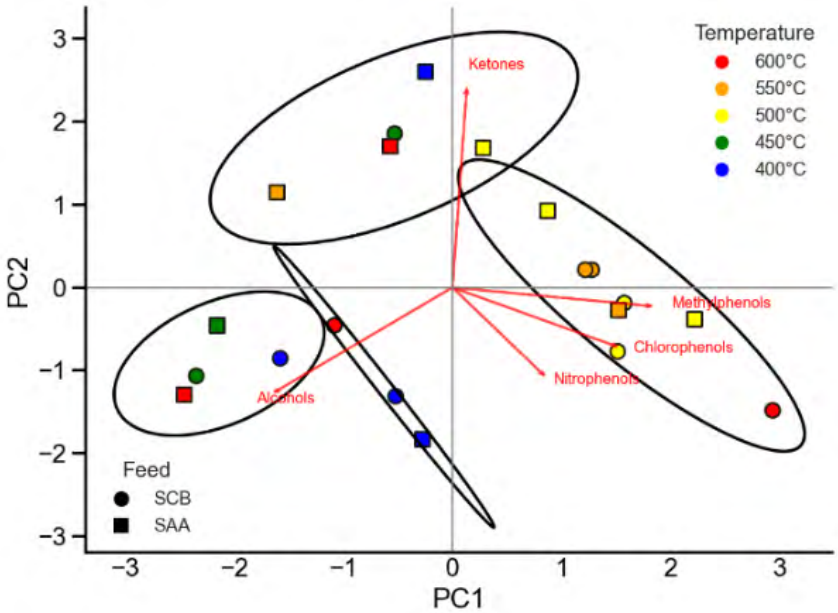

Figure 10. Biplot of only catalytic liquid scores.

increasing temperature drives the data upward along PC2 and progressively toward lower PC1 values. This trend suggests a temperature-driven shift in dominant gas-phase pathways. Conversely, SAA exhibits an opposite behavior along PC2: increasing temperature results in a downward displacement, although a consistent shift toward lower PC1 values is also observed, indicating a shared thermal influence on gas evolution despite biomass-specific differences.

Liquid-Phase PCA

Loadings Analysis

The loadings bar plot for the liquid-phase PCA (see Fig. 8) captures the dominant chemical differentiation within the bio-oil fraction with respect to the first two principal components.

PC1 is primarily governed by the opposition between alcohols and substituted phenolic compounds, defining an axis associated with the structural origin of liquid products. Negative PC1 loadings are dominated by alcohols, which are characteristic of primary depolymerization products from cellulose and hemicellulose. In contrast, positive PC1 loadings are associated with methyl-, chloro-, and nitrophenols, reflecting lignin-derived compounds and secondary aromatic stabilization pathways. This confirms PC1 as a descriptor of increasing aromaticity and secondary processing severity in the liquid phase.

PC2 is strongly controlled by ketones, which exhibit the largest positive loadings in both PCA configurations. This component represents the prevalence of carbonyl-forming fragmentation routes, typically linked to carbohydrate decomposition. Negative PC2 contributions correspond to alcohols and phenolic species, indicating a shift away from simple fragmentation toward more stabilized or condensed liquid products.

When thermal experiments are incorporated, only minor variations in loading magnitudes are observed, while the overall structure of PC1 and PC2 remains unchanged.

This indicates that the chemical partitioning of liquid products is fundamentally preserved, with temperature primarily affecting the distribution of experiments in score space rather than redefining the underlying liquid-phase reaction pathways.

Biplot Analysis

For the thermo-catalytic dataset, the biplot (see Fig. 9) shows that thermal experiments are exclusively located in the negative PC1 region, indicating the absence of liquid-phase stabilization associated with positive PC1 scores.

A clear separation by feedstock is observed along PC2: SCB thermal experiments are mainly located at positive PC2 values, aligned with the ketone-related vector, whereas SAA experiments cluster at negative PC2 values, following the alcohol-related direction.

When analyzing the thermo-catalytic data according to the C/B ratio, most experiments with higher catalyst loading for both biomasses shift toward the (+PC1, −PC2) quadrant. This behavior suggests a catalytic-driven transition toward liquid compositions characterized by enhanced dehydration and secondary transformation pathways.

For the catalytic-only dataset (see Fig. 10), the

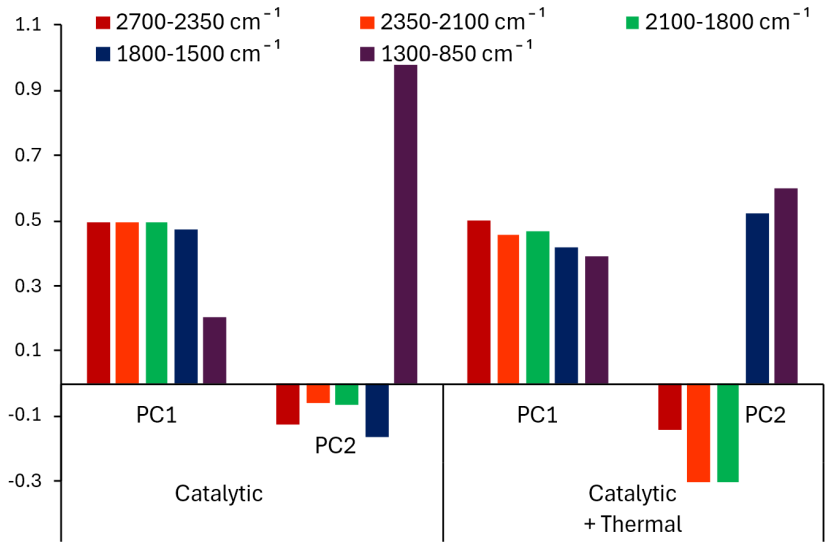

Figure 11. Bar plots, solid FTIR loadings in both groups.

temperature-colored biplot highlights a distinct thermal effect. In SCB, increasing temperature promotes a displacement toward positive PC1 values with intermediate PC2 scores, although all bagasse experiments remain in the negative PC2 region. In contrast, SAA exhibits a different high-temperature response: experiments at 600 °C shift toward lower PC1 values, indicating a tendency toward alcohol- and ketone-rich liquid compositions. Despite these differences, both biomasses show a temperature-dependent reorganization of liquid-phase chemistry under catalytic conditions.

FTIR-Based PCA

Loadings Analysis

The FTIR-based PCA of solid residues (see Fig. 11) exhibits a highly structured behavior, with the first principal component explaining the majority of the variance in both datasets. This confirms that the chemical evolution of the solid phase is governed by a limited number of dominant transformation pathways rather than by random spectral variations.

PC1 is characterized by uniformly positive loadings across all FTIR spectral regions, indicating that it represents the overall intensity and preservation of functional groups in the solid residue. High PC1 values correspond to solids retaining a broad range of oxygenated and aliphatic functionalities, whereas low PC1 values are associated with extensive devolatilization and carbonization, resulting in weak FTIR signals. The slight reduction in PC1 explained variance upon inclusion of thermal experiments suggests the emergence of additional structural diversity in the solid phase at high temperatures.

PC2 provides chemical discrimination among functional groups, separating spectral regions associated with aliphatic and carbohydrate-related vibrations from those linked to more condensed aromatic structures. Positive PC2 loadings are dominated by low-wavenumber regions (1300–850 cm^{-1} and 1800–1500 cm^{-1}), typically attributed to C–O, C–C, and aromatic skeletal vibrations, while negative PC2 contributions are associated with higher-wavenumber regions linked to less condensed structures.

Compared to the catalytic-only PCA, the catalytic–thermal PCA shows a redistribution of PC2 loadings, indicating that temperature enhances structural differentiation within the solid, promoting the coexistence of partially functionalized aromatic char and more condensed carbonaceous domains. Importantly, the interpretation of PC1 as a severity axis and PC2 as a structural-type axis remains robust across both datasets.

Biplots Analysis

For the thermo-catalytic dataset, the biplot (see Fig. 12) reveals a clear and consistent separation between thermal and catalytic regimes along PC1. All thermal experiments are in the negative PC1 region, whereas catalytic experiments shift toward positive PC1 values, indicating the development of more structured and chemically evolved solid residues under catalytic conditions. Along PC2, feedstock-dependent separation is observed: SCB thermal samples are mainly located at negative PC2 values, while SAA samples cluster toward positive PC2.

Within the catalytic experiments, a systematic trend with C/B ratio is evident. For SAA, increasing catalyst loading induces a displacement along a negative slope, with both PC1 and PC2 values decreasing, suggesting progressive modification of surface functionalities and aromatic structures. In contrast, SCB exhibits a more

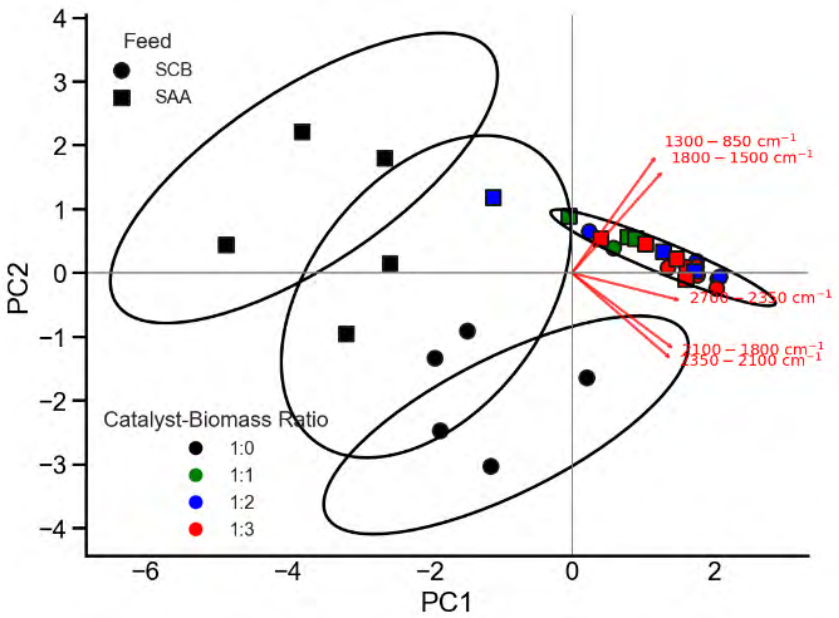

Figure 12. Biplot of thermal+catalytic Solid-FTIR scores.

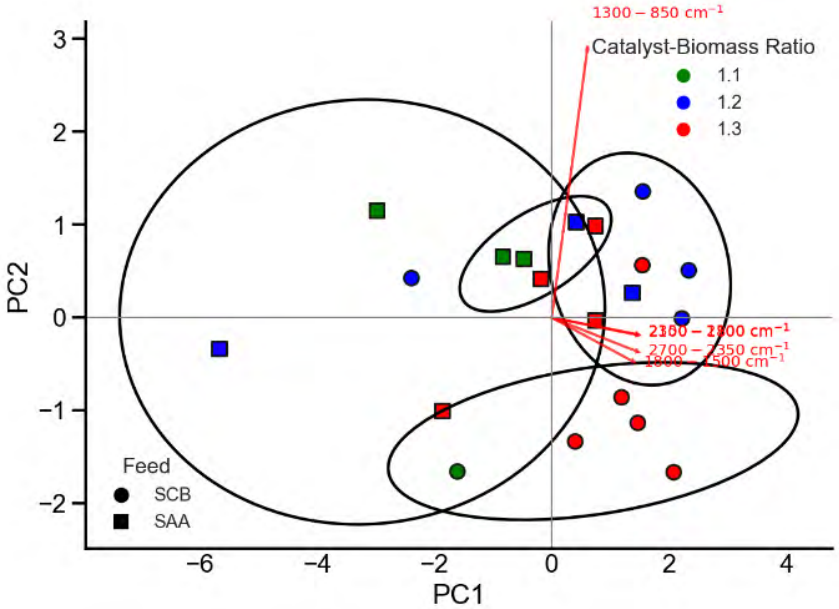

Figure 13. Biplot of only catalytic solid-FTIR scores.

pronounced displacement along PC2, with comparatively minor changes along PC1, indicating a different catalytic restructuring pathway of the solid phase.

For the catalytic-only dataset (see Fig. 13), a distinct feedstock-dependent response is observed. In SAA, increasing C/B ratio promotes a shift toward positive PC1 values, consistent with enhanced catalytic control over char aromatization and functional group evolution. Conversely, SCB samples primarily shift downward along PC2, suggesting progressive loss or transformation of specific surface functionalities rather than changes in overall aromaticity.

Reaction Trends and Proposed Mechanistic Pathways

The integrated interpretation of the PCA results obtained from product yields, gas-phase composition, liquid-phase composition, and solid-phase FTIR provides a

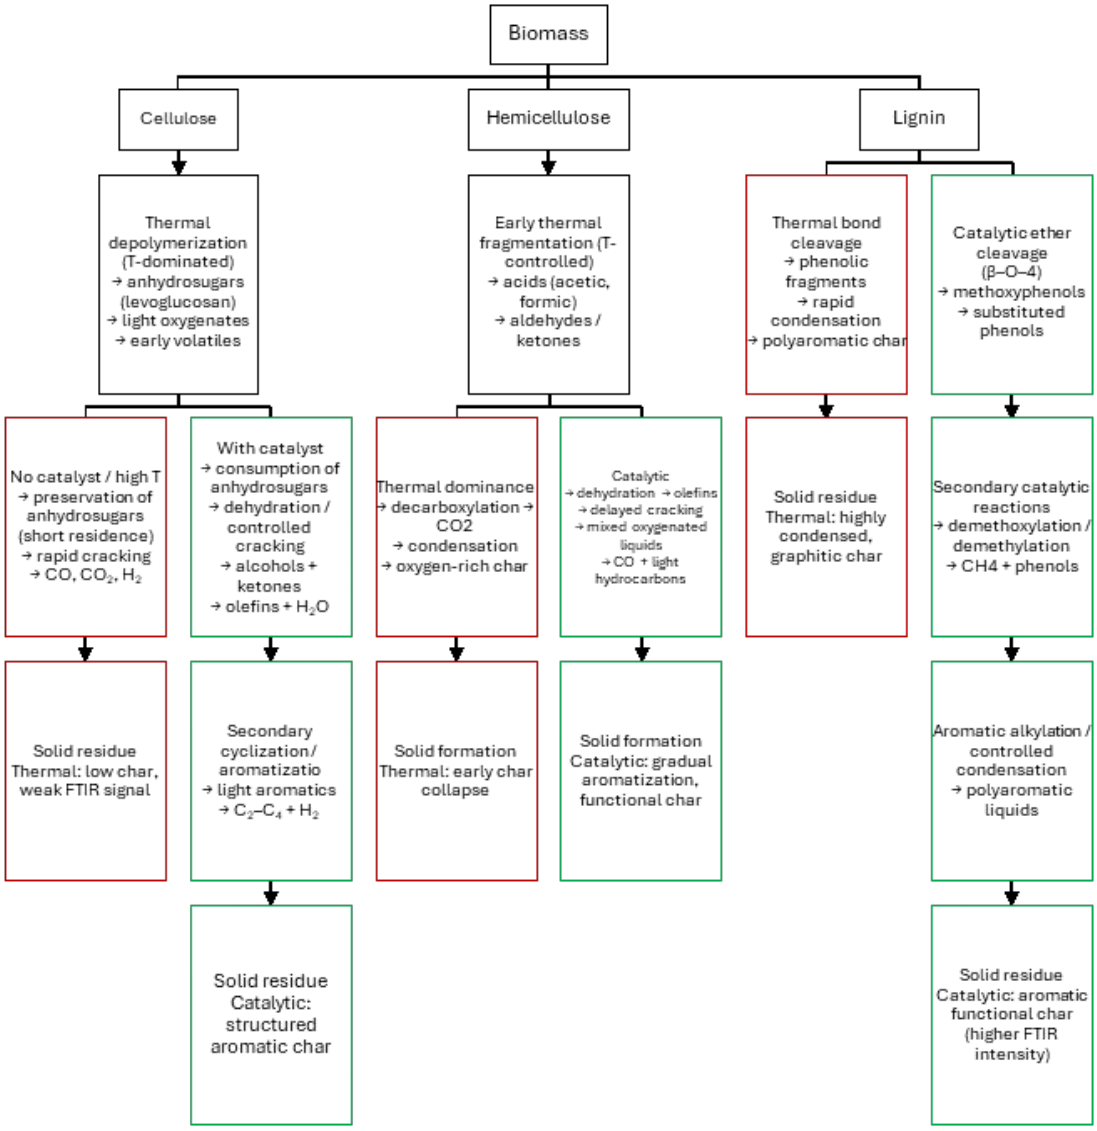

Figure 14: Conceptual reaction scheme for biomass pyrolysis highlighting the main transformation pathways of cellulose, hemicellulose, and lignin into gas, liquid, and solid products of temperature-dominated (red boxes) and catalyst-assisted (green boxes) routes.

consistent basis for proposing the reaction pathways summarized in Fig. 14. The multivariate trends indicate that biomass pyrolysis proceeds through the coexistence of temperature-dominated (red boxes) and catalyst-assisted (green boxes) routes, whose relative contributions depend strongly on feedstock composition.

For SCB, thermal effects primarily govern the initial depolymerization of cellulose and hemicellulose, leading to rapid volatilization and gas formation, as reflected by the clustering of thermal experiments in PCA regions associated with negative PC1 or PC2 values in the yields and gas-phase analyses. As depicted in Fig. 14, the presence of catalyst promotes secondary dehydration, cracking, and aromatization reactions, which are evidenced by systematic displacements along PC1 in yields and gas-phase PCA and along PC2 in FTIR. These trends indicate enhanced formation of light hydrocarbons and aromatics, together with the development of a more structured and aromatic solid residue. The C/B ratio therefore modulates reaction severity and solid-phase restructuring rather than fundamentally altering the primary decomposition pathways of bagasse.

In contrast, SAA exhibits a stronger coupling between catalytic effects and reaction selectivity. The biplots for yields, gases, and liquids show clear and systematic shifts with increasing C/B ratio, particularly along

PC2, indicating a catalyst-driven redistribution among oxygenated liquids, permanent gases, and stabilized solids. This behavior is consistent with the reaction scheme in Fig. 14, where catalytic dehydration, decarboxylation, and controlled condensation pathways become increasingly dominant. The FTIR PCA further supports this interpretation, as catalytic experiments display concurrent displacements along PC1 and PC2, reflecting simultaneous changes in aromaticity and surface functionality of the resulting char.

Overall, the PCA trends across all product phases converge toward a unified mechanistic framework, illustrated in Fig. 14, in which temperature controls primary bond cleavage and volatilization, while the catalyst selectively enhances secondary reactions such as dehydration, decarboxylation, aromatization, and controlled condensation. The balance between these pathways is feedstock-dependent, leading to distinct reaction trends and product distributions for SCB and SAA. This figure-assisted mechanistic interpretation captures the essential reaction behavior inferred from the multivariate analysis without requiring an explicit global PCA.

These observations are also consistent with previously reported descriptions of biomass pyrolysis [1, 8].

CONCLUSIONS

The PCA results demonstrate that separating catalytic and thermal–catalytic datasets is essential to disentangle the dominant reaction regimes during biomass pyrolysis. Temperature primarily governs severe cracking and carbonization pathways, while the catalyst introduces selectivity by promoting secondary reactions and moderating product distribution. The consistent displacement of experimental conditions in the PCA space highlights distinct behaviors for SCB and SAA, linked to their different biochemical compositions. The reaction scheme proposed in this work summarizes these trends and provides a coherent mechanistic framework connecting multivariate patterns with underlying transformation pathways.

ACKNOWLEDGEMENTS

The authors thank the Secretary of Science, Humanities, Technology and Innovation (SECIHTI) of Mexico for the scholarship granted to Martín Rodríguez-Fragoso.

AUTHOR IDENTIFIERS

Author ORCIDs:
Martín Rodríguez-Fragoso: 0009-0003-4299-9641,
Sandro González-Arias: 0000-0002-7395-1566,
Octavio Elizalde-Solis: 0000-0002-7282-3554,
Edgar Ramirez-Jimenez: 0000-0002-3509-0101.

REFERENCES

1. Collard FX, Blin J. A review on pyrolysis of biomass constituents: mechanisms and composition of the products obtained from the conversion of cellulose, hemicelluloses and lignin. Renewable and Sustainable Energy Reviews 38:594-608 (2014). https://doi.org/10.1016/j.rser.2014.06.013
2. Vuppaladadiyam AK, Varsha Vuppaladadiyam SS, Sikarwar VS, Ahmad E, Pant KK, S M, Pandey A, Bhattacharya S, Sarmah A, Leu SY. A critical review on biomass pyrolysis: reaction mechanisms, process modeling and potential challenges. Journal of the Energy Institute 108:101236 (2023). https://doi.org/10.1016/j.joei.2023.101236
3. Wang S, Dai G, Yang H, Luo Z. Lignocellulosic biomass pyrolysis mechanism: a state-of-the-art review. Progress in Energy and Combustion Science 62:33-86 (2017). https://doi.org/10.1016/j.pecs.2017.05.004
4. Aguado R, Elordi G, Arrizabalaga A, Artetxe M, Bilbao J, Olazar M. Principal component analysis for kinetic scheme proposal in the thermal pyrolysis of waste HDPE plastics. Chemical Engineering Journal 254:357-364 (2014). https://doi.org/10.1016/j.cej.2014.05.131
5. Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, Taylor J, Berg S, Smith NJ, Kern R, Picus M, Hoyer S, van Kerkwijk MH, Brett M, Haldane A, del Río JF, Wiebe M, Peterson P, Gérard-Marchant P, Sheppard K, Reddy T, Weckesser W, Abbasi H, Gohlke C, Oliphant TE. Array programming with numpy. Nature 585:357-362 (2020). https://doi.org/10.1038/s41586-020-2649-2
6. McKinney W. Data structures for statistical computing in python. Proceedings of the Python in Science Conference :56-61 (2010). https://doi.org/10.25080/majora-92bf1922-00a
7. Pedregosa F, Varoquaux, G.Scikit-learn: Machine learning in python. *JMLR* 12:2825-2830 (2011)
8. Ranzi E, Cuoci A, Faravelli T, Frassoldati A, Migliavacca G, Pierucci S, Sommariva S. Chemical kinetics of biomass pyrolysis. Energy Fuels 22:4292-4300 (2008). https://doi.org/10.1021/ef800551t

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimisation of Synthetic Natural Gas Production via Direct Air Capture and Utilisation using Reduced Models under a Novel Trust-Region Funnel Method

Gul Hameed[a], Tao Chen[a], Antonio del Rio Chanona[b], Lorenz T Biegler[c], and Michael Short[a]*

[a] School of Chemistry and Chemical Engineering, University of Surrey, Guildford GU2 7XH, UK
[b] Department of Chemical Engineering, Imperial College London, United Kingdom
[c] Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA
* Corresponding Author: m.short@surrey.ac.uk.

ABSTRACT

In this study, we propose a novel trust-region funnel (TRF) optimisation framework for process systems that integrate external black-box models, such as rigorous models, within equation-oriented (EO) formulations. The framework is applied to optimise a synthetic natural gas production process combining direct air capture and catalytic CO_2 conversion using dual-function material (DFM) technology, with the objective of minimising the total annualised cost. The problem is formulated in Pyomo and solved using IPOPT, treating the DFM reactor as an external black-box model. The TRF method achieves substantial improvements compared to published mixed-integer nonlinear programming and direct nonlinear programming approaches, reducing capture cost from 460 USD to 426 USD per tonne of CO_2. Key design improvements include reducing the number of DFM units per train by one-third and achieving a 22% reduction in DFM capital costs. These results highlight the TRF framework's ability to overcome numerical challenges and unlock economically superior designs for next-generation carbon capture and utilisation systems.

Keywords: Algorithms, Derivative Free Optimization, Catalysis, Carbon Capture, Reaction Engineering

INTRODUCTION

Modern challenges in process systems engineering (PSE) increasingly involve multi-scale optimisation problems aimed at rigorously assessing the performance and scalability of emerging technologies, such as carbon capture and utilisation (CCU) systems [1]. These problems are inherently hybrid, combining explicit equation-oriented (EO) glass-box process models with external, implicitly defined components, commonly referred to as black-box models. Black-box components typically originate from rigorous high-fidelity models (e.g., systems of computationally challenging nonlinear algebraic equations or ordinary and partial differential equations) or simulation packages, whose analytical structure is often unavailable or impractical to embed directly within EO formulations. When integrated into EO optimisation frameworks, rigorous high-fidelity models frequently introduce numerical instability and cause convergence issues in traditional nonlinear programming (NLP) solvers, while simulation packages cannot be incorporated as NLPs due to their model form [2].

Achieving convergence to first-order optimality in hybrid glass-box/black-box optimisation problems requires specialised solution strategies that explicitly account for the presence of external models. The classical trust-region filter method is a promising approach for solving such hybrid problems, but it faces significant challenges, including excessive numbers of expensive model evaluations and strong sensitivity to initialisation and algorithmic tuning parameters [3]. To address the limitations of NLP solvers and the trust-region filter method in handling rigorous high-fidelity models, this work proposes a trust-region funnel (TRF) optimisation framework for embedding rigorous external models within nonlinear process flowsheets. The TRF method extends the trust-region filter approach [4] by replacing the filter globalisation strategy with a funnel-based method [5]. This modification improves robustness, preserves convergence guarantees, and reduces reliance on

heuristic parameter tuning, enabling the optimisation of tightly coupled hybrid systems. A key feature of the proposed TRF strategy is its ability to efficiently utilise derivative information whenever available, thereby reducing the computational cost associated with external model evaluations. When analytic derivatives of black-box models are unavailable, the method remains applicable by employing derivative approximation at the cost of additional model evaluations. This flexibility makes the TRF method particularly well-suited for large-scale hybrid optimisation problems involving cyclic operation, implicit system behaviour, or computationally challenging and derivative-free sub-models.

Growing concerns regarding the long-term environmental risks and uncertainties associated with CO_2 storage have intensified interest in CO_2 utilisation pathways [5]. Among these, catalytic conversion of captured CO_2 into synthetic natural gas (SNG) using renewable hydrogen, commonly referred to as power-to-gas (PtG), has emerged as a promising option. This study demonstrates the application of the proposed TRF framework to optimise an integrated SNG production process based on direct air capture (DAC) and catalytic conversion using dual-function material (DFM) technology. The DFM reactor integrates CO_2 capture and methanation within a single unit [6]. In addition to the DFM unit, the process flowsheet includes upstream air intake and high-temperature solid oxide electrolyser (SOEL) units, and downstream heat-integration and separation systems.

The resulting problem is characterised by a large number of equality constraints and strong nonlinear coupling across multiple process units, making it computationally challenging for standard NLP solvers. To address these challenges, previous works [6][7] imposed restrictive assumptions, such as fixing key design and operating variables (e.g., adsorption time and DFM mass). While effective for numerical convergence, these restrictions significantly reduce the feasible design space, preventing identification of the true economic optimum.

In this work, we overcome these limitations by introducing additional degrees of freedom, allowing all process variables, including adsorption time and DFM catalyst weight, to vary during the optimisation process. The DFM reactor is treated as an external multi-output black-box component, embedded within an EO formulation, and the resulting hybrid optimisation problem is solved using the proposed TRF method. The results demonstrate that the TRF approach achieves robust and systematic optimisation of process systems that are otherwise difficult to optimise using traditional NLP techniques.

TRUST-REGION FUNNEL METHOD

The general form of the hybrid optimisation problem considered in this work can be expressed as:

$$\min_{w,z} \; f(w,y,z) \tag{1.1}$$

$$s.t. \; h(w,y,z) = 0 \tag{1.2}$$

$$g(w,y,z) \leq 0 \tag{1.3}$$

$$y = t(w) \tag{1.4}$$

where f denotes the objective function, h and g represent equality and inequality constraints, and $t(w)$ is the rigorous (truth) model describing external black-box components. The functions f, g and h are assumed twice continuously differentiable. The variables w correspond to the inputs of the truth model, y are the associated external model outputs, and z denotes the remaining variables.

Trust-Region Subproblem

Trust-region methods iteratively solve a trust-region subproblem (TRSP) by replacing the truth model with a reduced model (RM). The RM is constructed to be sufficiently accurate within a neighbourhood of the current iterate, satisfying the κ-fully linear property as defined in [8]. The TRSP at iteration k is formulated as:

$$\min_{w,z} \; f(x) \tag{2.1}$$

$$s.t. \; h(x) = 0 \tag{2.2}$$

$$g(x) \leq 0 \tag{2.3}$$

$$y = r_k(w) \tag{2.4}$$

$$\|u - u_k\| \leq \Delta_k \tag{2.5}$$

where $x = [w,y,z]^T, r_k(w)$ denotes the RM at iteration k, u is a selected sub-vector from the variable vector $v = [w,z]^T$ that represents independent variables, u_k is the current base point, and Δ_k is the trust-region radius. The trust-region constraint restricts the step size in u to ensure that the RM provides a sufficiently accurate approximation of the truth model within the trust-region.

Reduced Model Construction

To enforce local consistency between the reduced and truth models, zero- and first-order correction terms are introduced whenever derivative information from the truth model is available or can be approximated. In this work, the RM is constructed as:

$$r_k(w) = \tilde{r}(w) + \left(t(w_k) - \tilde{r}(w_k)\right) + (\nabla t(w_k) - \nabla \tilde{r}(w_k))^T (w - w_k) \tag{3}$$

where $\tilde{r}(w)$ is a basis (or base surrogate model). In the present case study, $\tilde{r}(w) = 0$, such that the RM corresponds to a first-order Taylor approximation of the truth model about the current base point w_k. Although convergence guarantees hold for arbitrary choices of $\tilde{r}(w)$, algorithmic performance improves significantly as the accuracy of the $\tilde{r}(w)$ increases.

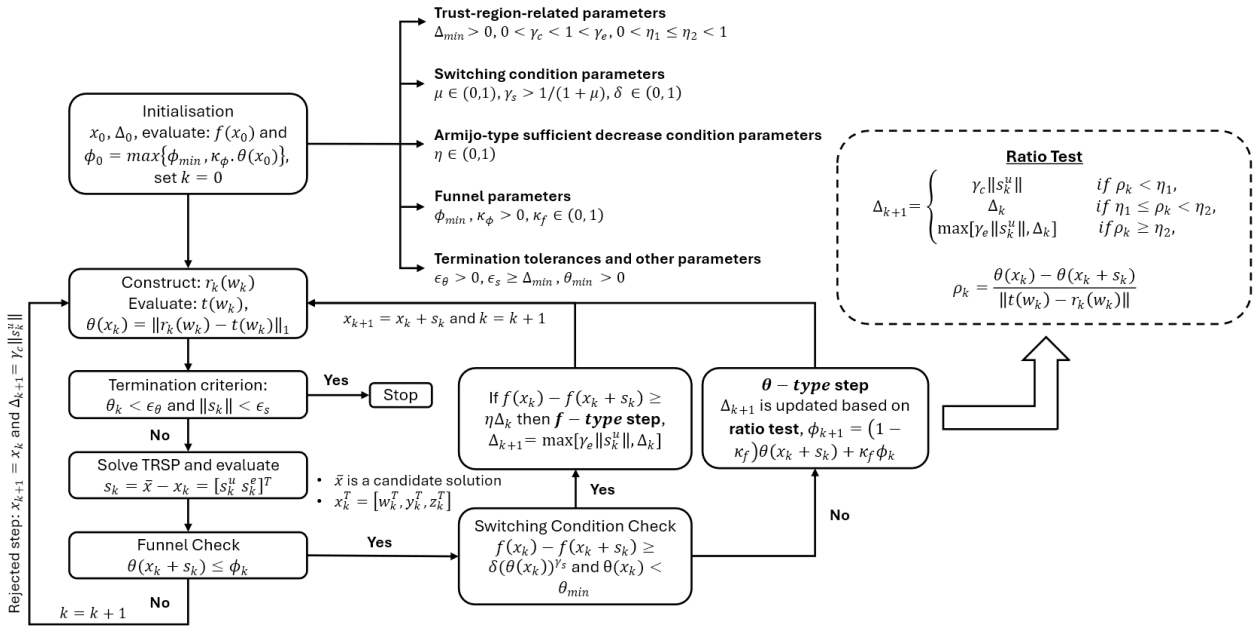

Figure 1: Trust-region funnel algorithm.

Funnel Globalisation Strategy

After solving the TRSP to obtain candidate solutions w_k^* and z_k^*, a funnel-based globalisation strategy is applied. The funnel enforces a progressively tightening upper bound ϕ_k on the RM approximation error θ, defined as:

$$\theta(w) = \|r(w) - t(w)\|_1. \tag{4}$$

Based on the agreement between predicted and actual reductions in the approximation error, quantified through a ratio test (see Figure 1 for details), and pre-defined funnel acceptance criteria, the trust-region radius is updated, and the candidate iterate is either accepted or rejected.

While the solution of each TRSP primarily targets improvement in the objective function, the funnel systematically drives successive iterates toward feasibility, enabling robust convergence to a locally optimal solution that satisfies first-order optimality conditions [9].

Algorithm Overview

The TRF algorithm proceeds iteratively as summarised in Figure 1. At initialisation, a base point x_0, initial trust-region radius Δ_0, and initial funnel bound $\phi_0 = \max\{\phi_{\min}, \kappa_\phi \cdot \theta(x_0)\}$ are set. At each iteration k, the reduced model $r_k(w_k)$ is constructed as per Equation (3) and the truth/rigorous model $t(w_k)$ is evaluated to compute the approximation error $\theta(x_k) = \| r_k(w_k) - t(w_k) \|_1$. If the termination criteria $\theta_k < \varepsilon_\theta$ and $\| s_k \| < \varepsilon_s$ are satisfied, the algorithm stops. Otherwise, the TRSP defined in Equations (2.1)-(2.5) is solved to compute a candidate step $s_k = [s_k^u, s_k^z]^T$, where superscripts u and z denote subvectors of the decision variables and the remaining variables, respectively.

The candidate iterate is then assessed through various sequential checks. First, a funnel check determines whether $\theta(x_k + s_k) \leq \phi_k$. If satisfied, a switching condition is evaluated: if sufficient objective reduction is achieved relative to the approximation error (i.e., switching condition) and trust-region size (i.e., Armijo-type sufficient decrease condition), an $f - type$ step is accepted and the trust-region radius is expanded. Otherwise, a $\theta - type$ step is taken, during which the funnel bound is updated as $\phi_{k+1} = (1 - \kappa_f)\theta(x_k + s_k) + \kappa_f \phi_k$, progressively tightening the feasibility envelope, and the trust-region radius is adjusted via the ratio test (as given in Figure 1). If the funnel check fails, the step is rejected and the trust-region radius is contracted.

This structure shares the switching condition and step-type logic of the classical trust-region filter method, but differs fundamentally in the globalisation strategy. The filter maintains an iteratively updated list of pairs (f, θ) and requires heuristic tuning of various problem-dependent tuning parameters. By contrast, the funnel can be interpreted as a filter with a single entry, i.e., the funnel width ϕ_k, updated only during $\theta - type$ iterations and initialised from $\theta(x_0)$ using a problem-independent parameter κ_ϕ. This eliminates the need to maintain non-dominated (f, θ) pairs or tune sensitive filter parameters, offering a key practical advantage. Importantly, this structural similarity enables convergence guarantees for the TRF method to be derived directly from trust-region filter theory [10].

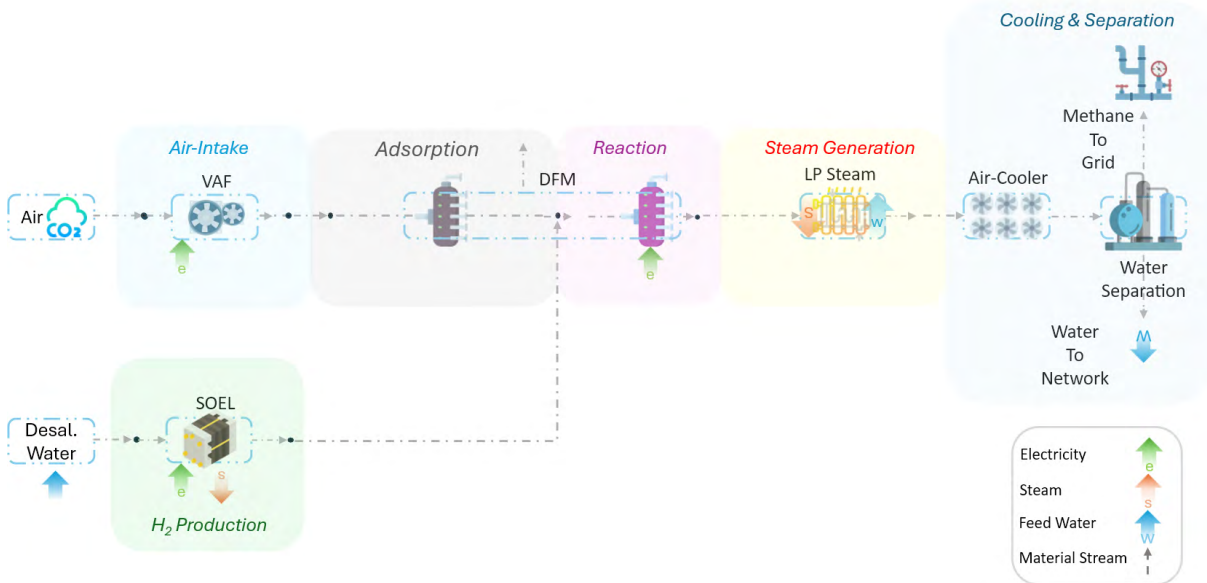

Figure 2: DFM DAC-to-SNG process flowsheet.

Convergence Guarantees

Under standard regularity assumptions, the proposed TRF method guarantees convergence to a first-order optimal point of the original hybrid optimisation problem, even in the presence of inexact or approximated derivative information from the external model [9][4].

PROCESS MODELLING

The optimisation objective is to minimise the total annualised cost (TAC) of the integrated process flowsheet shown in Figure 2. TAC accounts for capital and operating expenditures, including equipment sizing, material requirements such as water and DFM, and utility costs. The DAC process is assumed to be located within an industrial plant with access to electricity, cooling water, steam, and a boiler feed water network, ensuring the availability of required utilities. The mathematical models for the process units are adapted from [6][7].

Process Description

As illustrated in Figure 2, ambient air is introduced into the CO_2 capture unit via a vane axial fan (VAF) and directed into columns containing large tubes packed with honeycomb-shaped monoliths. These monolithic structures are coated with DFM, which facilitates CO_2 adsorption during the adsorption phase as the air stream passes through the monolith channels. The DFM formulation used in this study, $Ni^{(15\%)}Ru^{(1\%)}Na_2O^{(10\%)}/CeO_2\text{–}\gamma\ Al_2O_3^{(74\%)}$, is based on published works [11][7]. During this cyclic process, CO_2 is captured on the DFM surface, followed by a short nitrogen purge to remove residual gases. Subsequently, a hot hydrogen stream at 350 °C, supplied by SOELs, is introduced to hydrogenate the adsorbed CO_2 in situ, producing methane. A final purge completes the cycle before repetition. All electrolysers operate at atmospheric pressure with 100% stoichiometric conversion of desalinated water to H_2 and O_2, simplifying the model by omitting downstream separation.

The hot gas generated during the cyclic operation is utilised in steam generators to produce low-pressure (LP) steam for export or internal heat integration. After heat recovery, the process stream is trim cooled using an air-cooler and directed to a flash drum for water removal. The resulting gas mixture primarily consists of methane, with residual amounts of water vapour, CO_2, and hydrogen. During high-conversion operation, unreacted hydrogen is minimal; however, when excess hydrogen is employed for temperature control or enhanced heat recovery, pressure swing adsorption (PSA) may be used for hydrogen recovery and recycle.

Reactor Modelling

The DFM reactor is modelled as a multi-output external black-box component within the nonlinear optimisation formulation of the flowsheet. Although the DFM model is equation-based, consisting of a coupled system of nonlinear algebraic equations that describe adsorption, hydrogenation, reactor geometry, energy balances, velocity, and pressure drop, its direct inclusion in the EO flowsheet results in non-convergence of classical NLP solvers, even with near-optimal feasible initialisation.

To ensure numerical tractability while preserving model fidelity, the DFM reactor is externalised as a black-box mapping, which returns product compositions and energy requirements as functions of the operating

conditions. The surrounding process flowsheet is formulated in Pyomo, and the TRSPs are solved using IPOPT.

RESULTS AND DISCUSSION

The case study considers a techno-economic optimisation of a DAC-DFM flowsheet designed to capture 10, 000 tonnes of CO_2 per year, consistent with recently reported industrial-scale benchmarks [6][7]. The proposed TRF framework successfully converged to a first-order stationary point in 38 iterations, as shown in Figure 3, which summarises the solver's performance.

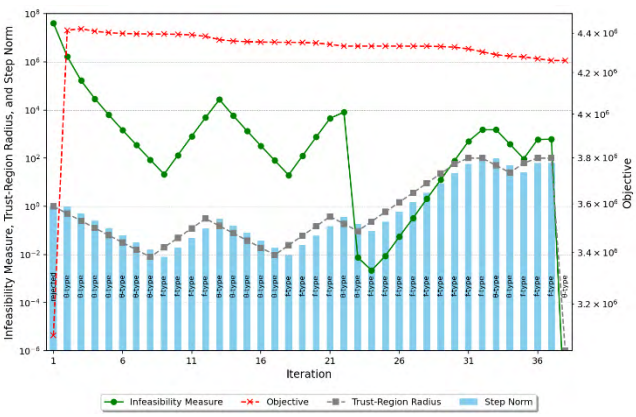

Figure 3. Performance metrics for the TRF solver.

Comparison of Optimisation Approaches

Key process metrics are summarised in Table 1, with an aggregated cost comparison presented in Figure 4. The study evaluates three optimisation approaches.

- The first approach uses a mixed-integer nonlinear programming (MINLP) formulation, where the DFM route, adsorbent mass, and adsorption time were fixed a priori [7].
- The second (direct NLP) optimisation approach converts the published MINLP flowsheet [7] to an NLP formulation while retaining fixed DFM weight and adsorption time to ensure solver convergence.

In the third (TRF-based grey-box) optimisation approach, the DFM reactor is treated as an external multi-output model, allowing both adsorbent weight and adsorption time to vary as decision variables.

As illustrated in Figure 4, the TRF framework achieves reduction in TAC compared to both the published MINLP and direct NLP solutions. While the published and direct NLP cases yield nearly identical results due to their fixed DFM design assumptions, the TRF solution reduces TAC by ~8%, demonstrating the economic benefits of relaxing restrictive design assumptions. The lower TAC is due to lower capital costs, while operational expenditures and revenues are observed to be similar for all optimisation approaches.

The observed improvement primarily stems from the optimisation of the adsorption–reaction unit design. Under the TRF framework, the number of DFM units per train is reduced from 16 to 6, while individual unit size and channel velocity are increased by ~2.9 and ~3.5 times, respectively. Adsorption time is also reduced from 1000 seconds to 626 seconds. These changes lead to a ~6% reduction in adsorbent weight and a corresponding ~23% decrease in DFM capital cost, without compromising capture targets or pressure-drop constraints. Note that the expensive catalyst sites in DFM reactor, particularly Ruthenium, result in significant capital expenditure in all optimisation approaches [7].

These results highlight a fundamentally different trade-off between reactor parallelisation (i.e., monolithic design), adsorption time, and adsorbent weight, systematically identified by the TRF algorithm. Unlike traditional approaches that rely on conservative over-design with multiple small parallel units, the TRF optimiser favours fewer, larger units operating at higher velocities and shorter cycle times. This design strategy reduces both equipment count and adsorbent inventory while maintaining overall process performance.

Economic and Practical Implications

This study illustrates that the algorithmic structure decisively influences not only whether a solution can be obtained but also its qualitative characteristics. By removing restrictive design and operational assumptions, imposed purely for numerical convenience, TRF method enables a faithful exploration of techno-economic search space and unlocks economically better designs, which is critical for emerging CCU and PtG systems to remain competitive. These systems inherently involve cyclic operation, strong nonlinear coupling, and implicit model formulations.

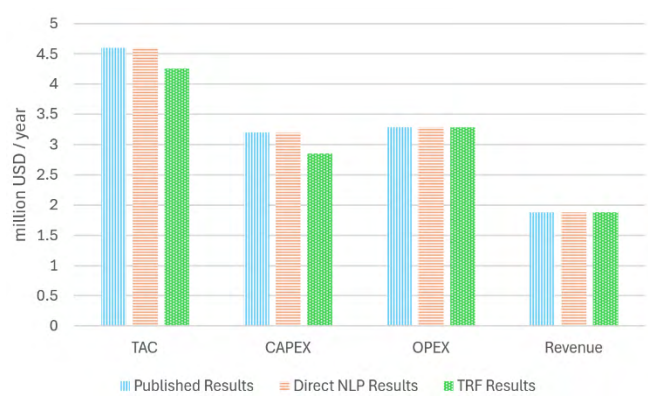

Figure 4. Economic comparison with published results.

Table 1: Optimisation results for the DFM-based process flowsheet (VAF: vane axial fan, SOEL: solid oxide electrolyser, DFM: dual-function material, LP: low pressure).

Variable Name	Unit	Published Results	Direct NLP Results	TRF Results
Air-Intake System: VAF				
Total capital cost	USD	705,100	705,030	705,030
Operational cost	USD/year	166,500	166,300	166,300
Number of fans	-	6	6	6
Fan power	kW	430	420	420
Hydrogen Production System: SOEL				
Total capital cost	USD	32,747,900	32,747,867	32,747,867
Operational cost	USD/year	3,115,800	3,115,738	3,115,738
Adsorption-Reaction Unit: DFM				
Total capital cost	USD	30,364,200	30,361,756	23,512,996
Operational cost	USD/year	0	0	0
Number of units in each train	-	16	16	6
Unit size (each)	m^3	22	22	63
Air velocity inside channels	m/s	2	2	7
Unit diameter	m	4	4	4
Unit height/length	m	2	2	5
Pressure drop in each unit	Pa	530	522	522
Total adsorbent weight	kg	28,350	28,343	26,678
Net CO_2 uptake	kmol/kg	0.000255	0.000255	0.000135
Steam Generation Level: LP				
Total capital cost	USD	54,600	54,555	54,555
Boiler heat transfer area	m^2	65	65	65
Steam production revenue	USD/year	26,000	25,950	25,950
Downstream Units				
Air-cooler total capital cost	USD	47,300	47,229	47,229
Air-cooler operational costs	USD/year	1,900	1,836	1,836
Separation vessel capital cost	USD	24,900	24,900	24,900
Separation vessel size	kg	340	339	339
Cost per tonne of CO^2 captured	USD/t_{CO^2}	460	460	426

CONCLUSION

This study demonstrates the effectiveness of the proposed TRF framework for optimising hybrid glass-box/black-box models in large-scale CCU processes, applied to a SNG production process integrating DAC and catalytic CO_2 conversion via DFM technology. The TRF method reduced TAC from 4.6 million USD to 4.26 million USD, achieving an ~8% decrease in the cost per tonne of captured CO_2. Key design improvements include reducing the number of DFM units per train (16 to 6), increasing unit size (22 m^3 to 63 m^3), and decreasing total adsorbent weight (28, 350 kg to 26, 678 kg), resulting in a 22% reduction in DFM capital costs. These results highlight the TRF framework's ability to overcome numerical challenges and deliver economically superior designs compared to traditional NLP/MINLP approaches. Future research will focus on embedding DFM reactor simulations based on partial differential equations within EO flowsheets to explore reactor dynamics and their interactions with the process. Moreover, integrating PSA as a downstream black-box unit can enable hydrogen recovery and recycling for enhanced process efficiency.

AUTHOR IDENTIFIERS

Author ORCIDs:
Hameed G: 0000-0002-2069-5132
Chen T: 0000-0002-3941-1603
Chanona AR: 0000-0003-0274-2852
Biegler LT: 0000-0003-3875-4441
Short M: 0000-0002-9719-3368

REFERENCES

1. Beykal B, Pistikopoulos EN. Data-driven optimization algorithms. Artificial Intelligence in Manufacturing :135-180 (2024). https://doi.org/10.1016/b978-0-323-99134-6.00006-2

2. Biegler LT. Multi-level optimization strategies for large-scale nonlinear process systems. Computers & Chemical Engineering 185:108657 (2024). https://doi.org/10.1016/j.compchemeng.2024.108657
3. Biegler LT. The trust region filter strategy: survey of a rigorous approach for optimization with surrogate models. Digital Chemical Engineering 13:100197 (2024). https://doi.org/10.1016/j.dche.2024.100197
4. Yoshio N, Biegler LT. Demand?based optimization of a chlorobenzene process with high?fidelity and surrogate reactor models under trust region strategies. AIChE Journal 67: (2020). https://doi.org/10.1002/aic.17054
5. Hasan MMF, Zantye MS, Kazi MK. Challenges and opportunities in carbon capture, utilization and storage: a process systems engineering perspective. Computers & Chemical Engineering 166:107925 (2022). https://doi.org/10.1016/j.compchemeng.2022.107925
6. M. Dolat, M. S. Duyar, and M. Short, "Techno-economic evaluation of novel dual-function reactor for direct air capture applying superstructure optimisation, " in *Computer Aided Chemical Engineering*, vol. 53, F. Manenti and G. V Reklaitis, Eds., Elsevier, 2024, pp. 1213–1218. doi: https://doi.org/10.1016/B978-0-443-28824-1.50203-9
7. Dolat M, Keynejad K, Duyar MS, Short M. Superstructure optimisation of direct air capture integrated with synthetic natural gas production. Applied Energy 384:125413 (2025). https://doi.org/10.1016/j.apenergy.2025.125413
8. Conn AR, Scheinberg K, Vicente LN. Introduction to derivative-free optimization. Society for Industrial and Applied Mathematics (2012). https://doi.org/10.1137/1.9780898718768
9. Hameed G, Chen T, del Rio Chanona A, Biegler LT, Short M. A trust?region funnel algorithm for gray?box optimization. AIChE Journal 72: (2026). https://doi.org/10.1002/aic.70258
10. Hameed G, Chen T, del Rio Chanona A, Biegler LT, Short M. Trust?region filter algorithms utilizing hessian information for gray?box optimization. AIChE Journal 72: (2026). https://doi.org/10.1002/aic.70236
11. Merkouri LP, Ramirez Reina T, Duyar MS. Feasibility of switchable dual function materials as a flexible technology for co₂ capture and utilisation and evidence of passive direct air capture. Nanoscale 14:12620-12637 (2022). https://doi.org/10.1039/d2nr02688k

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Control-Guided Reinforcement Learning for Cooperative Energy Management

Isabela Fons Moreno-Palancas[a]*, Raquel Salcedo Díaz[a], Rubén Ruiz Femenía[a], José A. Caballero[a], Antonio del Río Chanona[b]

[a] University of Alicante, Department of Chemical Engineering, Alicante, Comunidad Valenciana, Spain
[b] Imperial College London, Sargent Centre for Process Systems Engineering, Department of Chemical Engineering, London, United Kingdom
* Corresponding Author: isabela.fons@ua.es.

ABSTRACT

Addressing the urgent transition to low-carbon energy systems requires microgrids capable of locally coordinating electricity generation, storage, and flexible consumption. Their efficient integration calls for control strategies that are scalable, privacy-preserving, and robust to uncertainty. To address such a challenging control problem, this work proposes a decentralised Multi-Agent Reinforcement Learning (MARL) approach based on the Cross-Entropy Method (CEM) for the coordination of prosumers, equipped with renewable generation and vehicle-to-grid capabilities. To improve sample efficiency and robustness, the policy is warm-started using Behaviour Cloning (BC) from a classical Proportional-Integral-Derivative (PID) controller, resulting in a hybrid BC–CEM framework. The proposed method is evaluated in a realistic microgrid simulation with stochastic demand and real weather and generation profiles. Results show that BC–CEM accelerates convergence and achieves lower energy costs compared to both PID control and randomly initialized CEM, without sacrificing comfort or mobility requirements. The findings highlight the effectiveness of combining derivative-free optimization with imitation learning in complex MARL tasks, such as energy flexibility coordination.

Keywords: Reinforcement Learning, Behavioral Cloning, Derivative-Free Optimization, Machine Learning, Energy Management

INTRODUCTION

Distributed Energy Resources (DERs) such as rooftop solar panels or batteries empower households to act as prosumers, simultaneously generating and consuming electricity within the energy system. By producing their own electricity from renewable sources close to the point of use, DERs contribute to decarbonisation, improve grid resilience to disturbances, and reduce energy costs for end users [1]. Still, their efficient integration requires advanced control systems, capable of handling the stochastic, high-dimensional, and dynamic nature of these environments.

Although classical mathematical programming techniques provide reliable solutions with optimality guarantees, their applicability in real-life settings is limited, as their centralized formulations may raise privacy concerns, are computationally expensive and require decisions to be computed over the whole planning horizon upfront [2]. The dynamic adaptability of model predictive control (MPC) has motivated its application to DER optimization, but its performance is strongly dependent on the accuracy of the system model and shares computational intractability of classical optimization methods [3].

In this context, Reinforcement Learning (RL) has emerged as a powerful data-driven alternative to traditional control strategies. MARL in particular provides a promising framework for distributed autonomous energy management, as it allows multiple agents to cooperate and facilitates decentralized learning through experience-based optimization [4].

A wide range of approaches have been investigated in the literature. Value-based methods have been proven effective in microgrid coordination [5], but struggle with large continuous action spaces. Policy-based deep RL

algorithms, overcome this limitation and have broadened the applicability of RL to energy management problems [6, 7]. More recently, attention mechanisms have been incorporated into MARL architectures to further improve scalability in complex energy systems [8].

Despite these advances, exiting literature relies mostly on gradient-based optimisation techniques, which typically require extensive exploration and are highly sensitive to hyperparameter tuning. Ji et al. address these challenges by integrating evolutionary strategies in the MARL framework [9], but few studies explore the integration of expert knowledge in this context, despite their success in other control domains [10].

These gaps motivate the proposed hybrid Behaviour Cloning and Cross-Entropy Method (BC–CEM) approach for energy coordination in residential microgrids.

BACKGROUND

Reinforcement Learning

RL addresses sequential decision-making problems, modelled as Markov Decision Processes (MDPs), due to its inherent trial-and-error approach. In RL, an agent is defined as a system which takes a specific action based on information about the environment state that results in a transition of the environment to a new state. The goal is to find the policy, i.e., a mapping of states to actions, that maximises the expected cumulative reward over an episode.

In scenarios where multiple entities interact concurrently within a shared environment, agents are typically constrained in their capacity to observe the state of the environment, thereby restricting their ability to make informed decisions, which are based solely on local observations. MARL formalizes such problems as Decentralised Partially Observable MDPs (Dec-POMDPs), enabling each agent to learn autonomously, whilst adapting to other agents' behaviour. Additional challenges include non-stationarity induced by simultaneously learning agents, scalability with increasing agents, and the difficulty of credit assignment when global rewards are used.

Despite the decentralised nature of policy execution, the training process can follow centralized or independent paradigms. In centralized training, agents learn decentralized policies that are assessed by a centralized critic, typically a learned value function, which evaluates agents' actions based on global information about the environment state. Conversely, independent learning treats other agents as part of the environment, with each agent learning to make individual decisions that collectively maximize the global objective.

In energy microgrids, where households must make local control decisions under uncertainty and with limited information sharing, MARL provides a natural solution approach.

Evolutionary strategies in RL

Evolutionary Strategies (ES) are a class of derivative-free optimization algorithms inspired by biological evolution. Policy parameters are optimized by iteratively sampling a population of policies and recombining the most successful individuals to form the subsequent generation of candidate solutions.

Recent studies have evidenced ES can be competitive with gradient-based RL methods in complex control tasks [9, 10]. Moreover, their advantages are particularly attractive in MARL applications for energy systems given their simple implementation (no backpropagation necessary), resilience to highly stochastic or sparse rewards, and improved credit allocation for long-term horizons or when reliable value estimates are unavailable [11].

SYSTEM DESCRIPTION

Environment

The environment considered in this work models a microgrid of prosumers, equipped with photovoltaic panels (PV) and vehicle-to-grid (V2G) technologies, as depicted in **Figure 1**. Though adapted from Charbonnier et al. [5, 7], our implementation differs on action handling and data sampling procedures.

Each prosumer is considered an autonomous agent in a multi-agent framework. The environment operates with hourly time steps $t \in T$ over a 24-hour horizon. PV generation and electricity demands are stochastic, reflecting realistic weather conditions and user behavior. The goal is that each prosumer p learns a cost optimal energy management policy π_θ whilst ensuring comfort and mobility needs.

Only the core equations are reported in this section. For the full mathematical description, the reader is referred to the original formulation [5, 7].

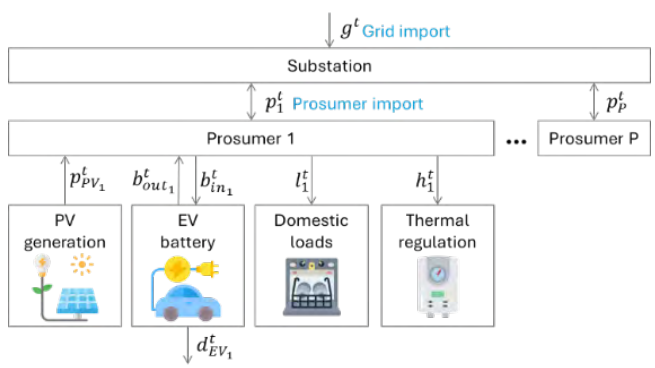

Figure 1. Microgrid of prosumers

Objective

The task is formulated as a multi-objective optimization problem. The primary objective is to minimize total

operational costs c^t according to Eq. 1:

$$c^t = C_g^t(g^t + \epsilon_g^t) + C_d \sum_p max(0, -p_p^t) + C_s \sum_p (b_{in_p}^t + b_{out_p}^t) \quad (1)$$

C_g^t, C_d, and $C_s$denote cost coefficients for energy imports, exports, and battery depreciation. Battery charge and discharge flows are represented by $b_{in}$and b_{out}, while p and g denote household-level and aggregate energy imports, with ϵ_g accounting for grid losses.

The maintenance of thermal comfort and sufficient battery levels to ensure mobility are imposed as soft constraints or penalties added to the economic objective.

Constraints

Home-level constraints determine the energy to be imported and exported, as well as the evolution of the system state. Thus, power, battery state of charge (SoC) and demand balances are enforced in the environment as hard constraints. Energy flows are expressed in kWh and temperatures in °C.

Power balance

$$p_p^t = l_p^t + h_p^t + b_{in_p}^t \eta_{ch} - \eta_{dis} b_{out_p}^t - p_{PV_p}^t \quad (2)$$

where l_p^t and h_p^t are domestic and heating loads, η_{ch}, η_{dis} are battery efficiencies, and $p_{p,t}^{PV}$ is PV generation.

Battery energy balance

$$E_p^{t+1} = E_p^t + b_{in_p}^t - b_{out_p}^t - d_{EV_p}^t \quad (3)$$

$$\mu_p^t \underline{E} \leq E_p^t \leq \mu_p^t \bar{E} \quad (4)$$

with EV demand $d_{EV_p}^t$, battery SoC bounds $\underline{E}$ and $\bar{E}$, and μ_p^t being a binary variable that equals 1 if EV is plugged and 0 otherwise.

Battery charge/discharge constraints

$$b_{in_p}^t \leq \mu_p^t \overline{b_{in}} \; ; \; b_{out_p}^t \leq \mu_p^t \overline{b_{out}} \quad (5)$$

where $\overline{b_{in}}$ and $\overline{b_{out}}$ are maximum charge and discharge rates.

Thermal model

$$\begin{bmatrix} BMT_p^{t+1} \\ IAT_p^{t+1} \end{bmatrix} = \kappa[1, BMT_p^t, IAT_p^t, \phi_p^t, h_p^t] \quad (6)$$

where BMT_p^t and IAT_p^t denote building mass and indoor temperatures, ϕ_p^t is the solar radiation and κ is a coefficient matrix, as defined in [5].

Consumption flexibility

$$\sum_{t_C \in T} \hat{l}_{p,k,t_C,t_D} f_{p,k,t_C,t_D} = d_{p,k}^{t_D} \quad (7)$$

$$\sum_{t_D \in N_{flex}} \hat{l}_{p,k,t_C,t_D} f_{p,k,t_C,t_D} = l_{p,k}^{t_C} \quad (8)$$

where $\hat{l}_{p,k,t_C,t_D}$ is the fraction of domestic loads $d_{p,k}^{t_D}$ of type k (flexible or fixed) requested at t_D and served at t_C. Flexible loads can be delayed up to n_{flex} time steps, and $f_{p,k,t_C,t_D} = 1$ if $t_D \leq t_C \leq t_D + n_{flex}$.

Action translation

At each time step, prosumers select three actions:

1. **Battery action** $a_b \in [-1, 1]$ controls actual battery charge b_{in} and discharge b_{out} flows according to Eq. 9, where maximum feasible amounts of charge ΔE^+ and discharge ΔE^- and are derived from Eq. 3 to 5.

$$\begin{aligned} b_{in_p}^t &= a_{b_p}^t \Delta E^+ \; ; \; b_{out_p}^t = 0 \quad a_{b_p}^t < 0 \\ b_{out_p}^t &= a_{b_p}^t \Delta E^- \; ; \; b_{in_p}^t = 0 \quad a_{b_p}^t > 0 \end{aligned} \quad (9)$$

2. **Thermal action** $a_h \in [-1, 1]$ regulates the electricity consumed by a heat pump h_p^t for heating/cooling, relative to its rated power P_{max} and coefficient of performance COP.

$$h_p^t = |a_{h_p}^t P_{max}| / COP \quad (10)$$

3. **Consumption action** $a_c \in [0, 1]$ determines the fraction of deferrable loads l_{flex}, within feasible bounds $[\underline{l}_{flex}, \overline{l}_{flex}]$ obtained from Eq. 7 and 8, to be satisfied at each time step t.

$$l_{p,flex}^t = \underline{l}_{p,flex}^t + a_{c_p}^t (\overline{l}_{p,flex}^t - \underline{l}_{p,flex}^t) \quad (11)$$

Policy network

The policy network outputs the mean and standard deviation of Gaussian distributions based on local observations, consisting of a window of Δt time steps of battery level, temperature measurements, PV generation and flexible loads, along with temporal features. As outlined in **Figure 2**, observations are normalized before being fed to the network.

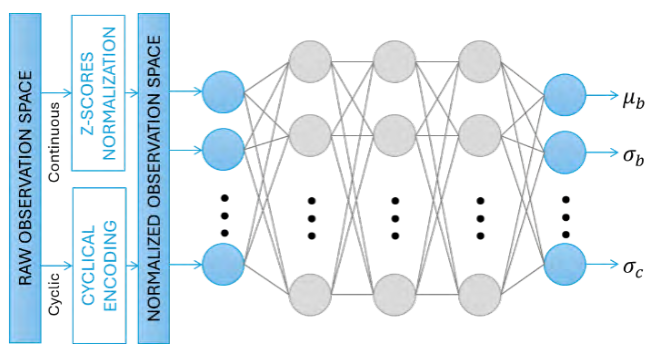

Figure 2. Forward pass of the policy network

METHODOLOGY

Cross-Entropy Method (CEM) is a stochastic, derivative-free optimization algorithm used to train a policy network π_θ by maximising the expected episode return

$J(\theta)$ [12], evaluated running n_r Monte Carlo rollouts:

$$J(\theta) = E_{\tau\sim\pi_\theta}\left[\sum_{t=0}^{T-1} r^t\right] \approx \frac{1}{n_r}\sum_{m=1}^{n_r}\sum_{t=0}^{T-1} r_m^t \tag{12}$$

Each iteration k starts by sampling a random population of n_c candidate parameter vectors $\Theta = \{\theta_i\}_{i=1}^{n_c}$:

$$\theta_i \sim \mathcal{N}\left(\boldsymbol{\mu}_k, \mathrm{diag}\left(\boldsymbol{\sigma}_k^2\right)\right), i = 1, \dots, n_c \tag{13}$$

These candidates are ranked according to their returns, and the top $(1-\rho)$-quantile is selected as the elite set $\mathcal{E}_\rho$, defined by the return threshold γ:

$$\mathcal{E}_\rho = \{\theta \in \Theta | \, J(\theta) \geq \gamma\} \tag{14}$$

The sampling distribution is updated to increase the likelihood of producing the elite following Eq. 15 and 16.

$$\boldsymbol{\mu}_{k+1} = \alpha_\mu \boldsymbol{\mu}_k + (1-\alpha_\mu)\boldsymbol{\mu}_{\mathbf{elite}} \tag{15}$$

$$\boldsymbol{\sigma}_{k+1} = \alpha_\sigma \boldsymbol{\sigma}_k + (1-\alpha_\sigma)\boldsymbol{\sigma}_{\mathbf{elite}} \tag{16}$$

where $\boldsymbol{\mu}_{\mathbf{elite}}$ and $\boldsymbol{\sigma}_{\mathbf{elite}}$ denote the per-parameter mean and standard deviation of the elite set, and $\alpha_\mu, \alpha_\sigma \in [0,1]$ are smoothing factors.

These steps are iterated and eventually yield an improved distribution over the parameter space. The best-performing candidate encountered during optimisation is retained.

While CEM evaluates a population of policies per iteration, its high degree of parallelization and absence of gradient estimation make it computationally attractive for stochastic, non-stationary MARL problems. Notwithstanding, in high-dimensional parameter spaces, vanilla CEM may require extensive exploration and is susceptible to premature convergence to local optima [13], motivating its combination with imitation learning.

Behavior Cloning (BC) is a supervised learning approach where agents learn to mimic expert demonstrations [14]. In this work, a PID controller is rolled out n_{BC} episodes to collect a dataset $\mathcal{D} = \left\{\left(s^{(i)}, a^{(i)}\right)\right\}_{i=1}^{n_{BC}}$. The policy network is then pretrained for a single agent and the optimal weights returned by Eq. 17 are used to initialize the policies of all agents during CEM execution in a favourable region of the parameter space, thus, reducing exploration requirements during CEM phase.

$$\theta_0 \in \underset{\theta}{\mathrm{argmin}} \frac{1}{|\mathcal{D}|}\sum_{(s,a)\in\mathcal{D}} \|\pi_\theta(\boldsymbol{s}) - \boldsymbol{a}\|_2^2 \tag{17}$$

The pseudocode of the BC-CEM strategy is given in **Table 1**. It is worth noting that we opted for an independent learning approach, i.e., CEM is applied per-agent, so that each one learns its own optimal parameter vector $\theta_{\mathbf{p}}^*$. However, agent returns are computed considering the joint action.

Table 1: BC-CEM algorithm

Input: $\mathcal{D}$: expert trajectory dataset; n_c: number of candidates; ρ: elite quantile; $\alpha_\mu, \alpha_\sigma$: mean and standard deviation update smoothing factors; σ_0: initial standard deviation; K: iterations
Output: Optimal multi-agent policy parameters θ^*
Initialization:
$\boldsymbol{\mu}_0^p \leftarrow$ Eq. 17; $\boldsymbol{\sigma}_0^p \leftarrow \boldsymbol{\sigma}_0$
$J_p{}^* \leftarrow -\infty$; $\theta_p{}^* \leftarrow \boldsymbol{\mu}_0$
for k in $\{1, \dots, K\}$ **do**
  for p in $1, \dots, n_p$ **do**
    for i in $1, \dots, n_c$ **do**
      $\theta_i^p \leftarrow$ Eq. 13
    end
  end
  for i in $1, \dots, n_c$ **do**
    $J\left(\{\theta_i^p\}_{p=1}^P\right) \leftarrow$ Eq. 12
  end
  for p in $1, \dots, P$ **do**
    $\mathcal{E}_\rho^p \leftarrow$ Eq. 14
    $\boldsymbol{\mu}_{\mathbf{elite}}^p \leftarrow \frac{1}{|\mathcal{E}_\rho^p|}\sum_{i\,\epsilon\,\mathcal{E}_\rho^p}\theta_i^p$; $\boldsymbol{\sigma}_{\mathbf{elite}}^p \leftarrow \sqrt{\frac{\sum_{i\,\epsilon\,\mathcal{E}_\rho^p}\left(\theta_i^p - \boldsymbol{\mu}_{\mathbf{elite}}^p\right)^2}{|\mathcal{E}_\rho^p|}}$
    $\boldsymbol{\mu}_{k+1}^p \leftarrow$ Eq. 15; $\boldsymbol{\sigma}_{k+1}^p \leftarrow$ Eq. 16
    if $J\left(\boldsymbol{\mu}_{\mathbf{elite}}^p\right) > J_p^*$ **then**:
      $J_p^* \leftarrow J\left(\boldsymbol{\mu}_{\mathbf{elite}}^p\right)$; $\theta_p^* \leftarrow \boldsymbol{\mu}_{\mathbf{elite}}^p$
    end
  end
end

Table 2: Hyperparameter set-up for BC-CEM algorithm

Hyerparameter	Value
n_c	100
ρ	0.1
α_μ	0.7
α_σ	0.5
σ_0	0.01
n_r	50

RESULTS AND ANALYSIS

Computational implementation

The training was conducted on a Windows 64-bit computer with an Intel i7-14700K CPU at 3.4 GHz and an NVIDIA GeForce RTX 5080 GPU, requiring 215 minutes (wall-clock) to complete $K = 250$ iterations and coordinate $P = 10$ prosumers. The hyperparameters used to run the training process (**Table 2**) were tuned through Bayesian optimization to maximize performance.

All agents share a common policy architecture, modelled as a three-layer feed-forward network with

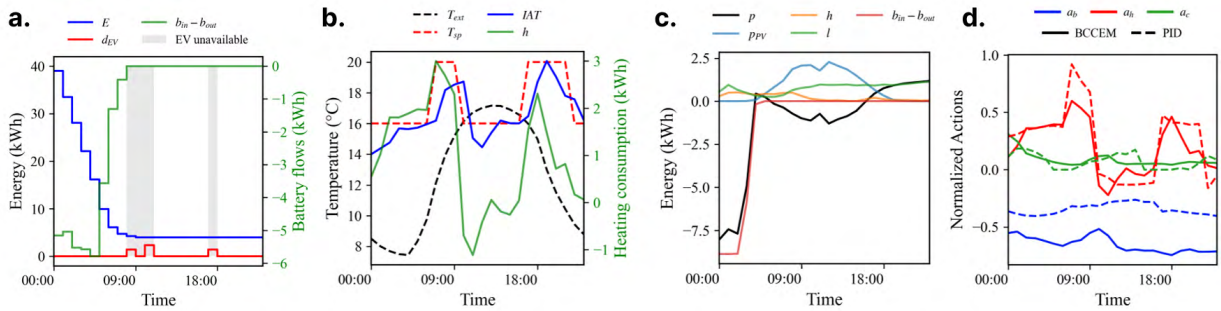

Figure 3. Example day representation of energy flows (a), battery SoC evolution (b), temperature evolution (c) and

structure [128, 128, 64] and ELU activations. Inputs concatenate $\Delta t = 3$ time steps of local observation history.Training

Figure 3 illustrates the evolution of the aggregate episode return $J(\mu_{elite})$ across agents during training, confirming the capacity of CEM to learn effective control policies despite the high stochasticity of the energy system environment.

Plus, it highlights the benefits of BC warm-starts. Agents initialized with expert demonstrations require fewer samples to reach high-performing solutions and consistently achieve higher returns for the same training budget compared to random initialisation.

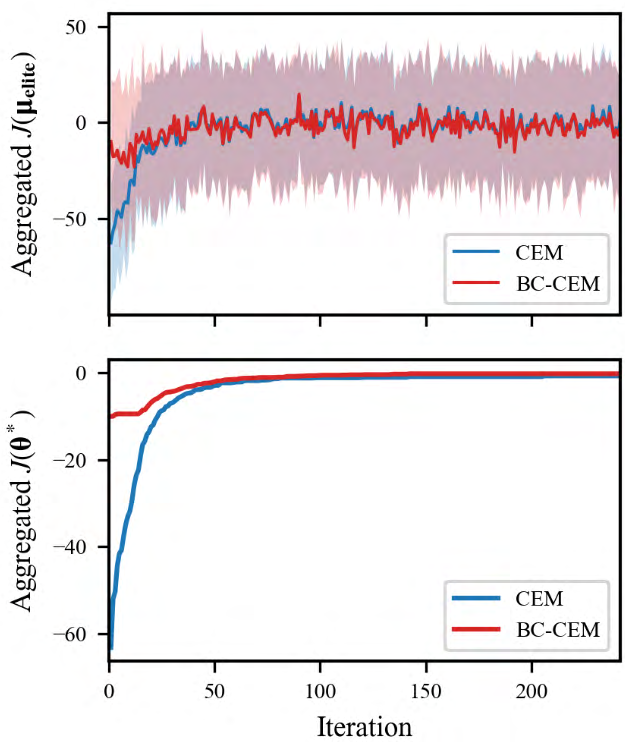

Figure 4. Learning curves for CEM and BC-CEM

Performance assessment

The learned policies are tested on a set of 200 episodes not seen during training, and the distribution of undiscounted episode returns is presented in **Figure 4**.

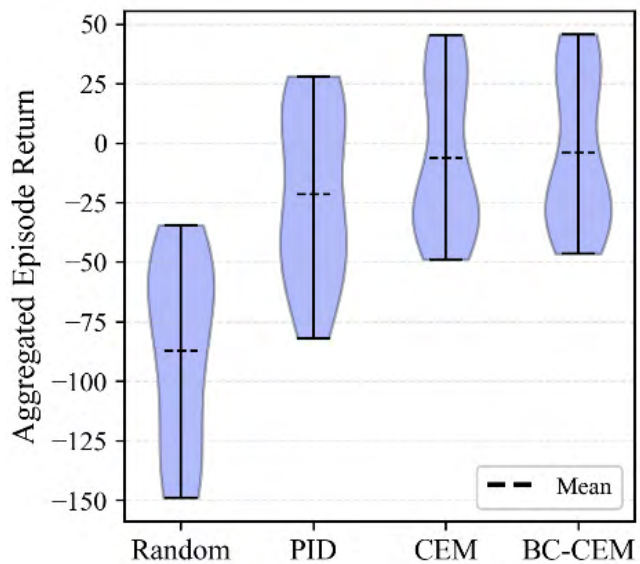

Figure 5. Violin plots of episode returns

Our approach outperforms the PID controller, as evidenced by a higher mean return. Quantitative results reported in **Table 3** further corroborate this result. BC–CEM achieves an average cost reduction of 180% relative to the PID baseline, while maintaining comparable levels of mobility feasibility. Although RL policies prioritise economic performance over thermal comfort, deviations remain limited, corresponding to an average hourly temperature mismatch of approximately 0.5 °C per prosumer

Table 3: Mean daily results for a 10-prosumer microgrid

	PID	CEM	BC-CEM
Total daily cost (£)	17.2	-12.4	-13.9
Comfort deviation (°C)	21.4	134	133
Battery SoC deficit (kWh)	1.41	1.58	1.42

Finally, a representative example day is illustrated in **Figure 5a-c**, whereas **Figure 5d** compares the actions suggested by the PID controller with those generated by a BC-CEM agent. While the fine-tuned policies preserve

the smoothness and stability of the PID trajectories, they drive the system towards more cost-optimal behaviour by exploiting flexibility more effectively.

CONCLUSION

This paper presented a hybrid Behaviour Cloning and Cross-Entropy Method (BC–CEM) approach for decentralised coordination of energy flexibility in residential microgrids. By combining expert demonstrations from classical PID control with derivative-free evolutionary optimisation, this work aggresses training instability under highly stochastic environments and reduces reliance on complex optimization pipelines and hyperparameter tuning. Numerical results in a microgrid environment reflecting realistic consumption patters and leveraging real climate data demonstrate that BC–CEM accelerates convergence and achieves lower operational costs compared to both PID control and vanilla CEM, while maintaining comparable comfort and mobility constraints. These findings highlight the benefits of integrating strong control priors with evolutionary strategies for scalable and robust energy management.

Future work will investigate scalability to larger agent populations, the incorporation of recurrent architectures to better exploit temporal dependencies, the impact of feature engineering on learning performance, the application of constrained reinforcement learning to eliminate residual constraint violations or the use of centralized training with decentralized execution MARL paradigms.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge financial support from the Spanish Ministry of Science, Innovation, and Universities via the doctoral grant for training university teachers FPU22/01276 and the R&D and Innovation Project PID2021-124139NB-C21.

AUTHOR IDENTIFIERS

Fons Moreno-Palancas I: 0009-0000-3503-9165
Salcedo Díaz R: 0000-0001-8290-0595
Ruiz Femenia R: 0000-0003-0652-5177
Caballero Suárez JA: 0000-0001-6470-2907
Del Río Chanona EA: 0000-0003-0274-2852

REFERENCES

1. IEA. Unlocking the Potential of Distributed Energy Resources. IEA (2022) Licence: CC by 4.0 https://www.iea.org/reports/unlocking-the-potential-of-distributed-energy-resources
2. Mohammadi P, Darshi R, Shamaghdari S, Siano P. Comparative analysis of control strategies for microgrid energy management with a focus on reinforcement learning. IEEE Access 12:171368-171395 (2024). https://doi.org/10.1109/access.2024.3495032
3. Zhang H, Seal S, Wu , Bouffard F, Boulet B. Building energy management with reinforcement learning and model predictive control: a survey. IEEE Access 10:27853-27862 (2022). https://doi.org/10.1109/access.2022.3156581
4. Zhang B, Hu W, MYM Ghias A, Xu X, Chen Z. Multi-agent deep reinforcement learning based distributed control architecture for interconnected multi-energy microgrid energy management and optimization. Energy Conversion and Management 277:116647 (2023). https://doi.org/10.1016/j.enconman.2022.116647
5. Charbonnier F, Morstyn T, McCulloch MD. Scalable multi-agent reinforcement learning for distributed control of residential energy flexibility. Applied Energy 314:118825 (2022). https://doi.org/10.1016/j.apenergy.2022.118825
6. Han Y, Wu J, Chen H, Si F, Cao Z, Zhao Q. Enhancing grid-interactive buildings demand response: sequential update-based multiagent deep reinforcement learning approach. IEEE Internet Things J. 11:24439-24451 (2024). https://doi.org/10.1109/jiot.2024.3357109
7. Charbonnier F, Peng B, Vienne J, Stai E, Morstyn T, McCulloch M. Centralised rehearsal of decentralised cooperation: multi-agent reinforcement learning for the scalable coordination of residential energy flexibility. Applied Energy 377:124406 (2025). https://doi.org/10.1016/j.apenergy.2024.124406
8. Savino S, Minella T, Nagy Z, Capozzoli A. A scalable demand-side energy management control strategy for large residential districts based on an attention-driven multi-agent DRL approach. Applied Energy 393:125993 (2025). https://doi.org/10.1016/j.apenergy.2025.125993
9. Ji C, Xiao H, Pei W, Wang X, Pu X. Coordinated voltage control for distribution network and multi-microgrids based on improved EM-RACE multi-agent reinforcement learning. *Int. J. Electr. Power Energy Syst*. 172:111315 (2025) https://doi.org/10.1016/j.ijepes.2025.111315
10. Bloor M, Ahmed A, Kotecha N, Mercangöz M, Tsay C, del Río-Chanona EA. Control-informed reinforcement learning for chemical processes. Ind. Eng. Chem. Res. 64:4966-4978 (2025). https://doi.org/10.1021/acs.iecr.4c03233
11. Salimans T, Ho J, Chen X, Sidor S, Sutskever I. Evolution Strategies as a Scalable Alternative to Reinforcement Learning. *arXiv* (2017).

https://arxiv.org/abs/1703.03864
12. de Boer PT, Kroese DP, Mannor S, Rubinstein RY. A tutorial on the cross-entropy method. Ann Oper Res 134:19-67 (2005). https://doi.org/10.1007/s10479-005-5724-z
13. Zhang Z, Jin J, Jagersand M, Luo J, Schuurmans D. A Simple Decentralized Cross-Entropy Method. *arXiv* (2022) https://doi.org/10.48550/arXiv.2212.08235
14. Levine S, Kumar A, Tucker G, Fu J. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv* (2020) https://doi.org/10.48550/arXiv.2005.01643

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

New tools, new thinking: Biomimetic Process Design through Parametric Modelling and Simulation

Alix Saury[abc], Thibaut Houette[a], Pierre-Emmanuel Fayemi[c], Jean-Matthieu Cousin[a], Jérôme Fortin[c], Arnaud Dujany[c]*

a Ceebios, Centre d'études et d'expertise en biomimétisme, 62 rue du Faubourg Saint-Martin, 60300 Senlis, France Provincial University, Department of Chemical Engineering, City, Province, Country
b IkosLab, 155 Rue Anatole France, 92300, Levallois-Perret, France
c Institut Polytechnique UniLaSalle, GeNumEr, UP 2018.C100, 14 Quai de la Somme 80082 Amiens, France
* Corresponding Author: arnaud.dujany@unilasalle.fr

ABSTRACT

This paper examines the mutually beneficial relationship between biomimetics and modelling and simulation tools, showing how each can enhance the other. Through a literature review and a detailed use case on anaerobic digestion, the study highlights how the complexity, multiscale organisation, and functional richness of biological systems challenge current modelling capabilities. By analysing the contributions of modelling and simulation to product development, such as early performance validation, rapid and low-cost iteration, and multicriteria evaluation, the paper questions whether integrating modelling and simulation tools to biomimetics would bring similar benefits to the design process. Several hypotheses are formulated regarding the potential contributions of modelling and simulation to biomimetics, particularly the improvement of biological system understanding through advanced visualisation and the assessment of functional viability using parametric modelling. Integrating such tools into biomimetics is envisioned as a means to reinforce the existing methodology, support more reliable principle transfer, and facilitate the wider adoption of biomimetics as a problem-solving design approach, ultimately promoting the emergence of biomimetic innovations across diverse applications.

Keywords: Biosystems, Modelling and Simulations, Multiscale Modelling, Process Design, Natural Gas

INTRODUCTION

The integration of modelling and simulation tools into product design and development processes has been widely recognized as a key enabler of reduced time-to-market, increased product complexity, and early validation of performance, safety, and fault tolerance [1]. Simulations allow engineers to rapidly and accurately understand physical behaviours before building physical prototypes, supporting a "learn-before-doing" approach that enhances design efficiency [2]. Beyond accelerating learning, simulation-based approaches enable designers to consider technological, economic, and logistic criteria in parallel, helping identify feasible design alternatives and suitable process plans [3]. Such tools support the design of complex, multidisciplinary systems, facilitate experimentation with technology integration, reduce physical prototyping costs, and promote collaboration among engineering teams [2, 4].

Recent decades have witnessed exponential growth in both numerical computational capabilities and advanced manufacturing technologies. Indeed, the improvements in computing power led to strong progression in the capacity to simulate, model, and forecast intricate systems and processes [5]. Manufacturing technologies have evolved from primarily subtractive processes to include sophisticated additive techniques capable of producing geometries with internal features, gradient materials, and hierarchical structures, previously impossible to fabricate [6]. This convergence of computational and manufacturing capabilities seems to have induced a shift in engineering design, from "what can do" to "what should we do" [7]. In other words, from "what can we feasibly produce given our constraints?" to "what should we create to best meet our objectives?". Technological ceilings have been pushed further, thus opening

the door to design spaces that were previously inaccessible. A next step could be in the remoulding of design approaches from the current top-down approach to a bottom-up one, so true hierarchical structural design of multifunctional materials and products could be achieved. The bottom-up approach is notably observed in nature and living systems. There is a design process, named biomimetics, that is the practice of mimicking biological strategies to fix practical problems in human activities (ISO 18458:2015). Biomimetics offers a compelling framework for the development and improvement of modelling and simulation tools, as biological organisms embody complex, multiscale, and multifunctional systems, difficult to replicate or to model accurately [5, 8].

Drawing on both a literature review and the analysis of a use case on anaerobic digestion, this work aims to highlight the mutually beneficial relationship between modelling and simulation tools and biomimetics. The objective is twofold: first, to explain how biomimetics can drive advancements in modelling and simulation; and second, to explore how these tools, in turn, can enhance the biomimetic design process. To address this objective, the first part outlines the biomimetic approach and summarises its relevance supported by literature for modelling and simulation tools. The second part introduces the anaerobic digestion use case, which examines the relevance in a practical context.

STATE OF THE ART

Biomimetics, a branch of design

Biomimetics, or bio-inspired design (BID) is a subset of design sciences, from which it borrows its procedural structure. Two methodological paradigms predominate in BID. The first, known as 'problem-driven', begins with the identification of opportunities and needs. The aim is to search for solutions to a specific technical problem in the living world. The second, known as 'solution-based', starts from an already identified biological mechanism for which potential technical applications are sought [9].

To illustrate this methodological dichotomy, the example of aerodynamic devices at the tips of aeroplane wings (winglets) can be considered. Inspired by the flexible feathers of large birds of prey, these elements significantly reduce the aerodynamic drag of the aircraft, thereby optimising its performance and reducing its energy consumption. This innovation perfectly exemplifies the 'problem-driven' approach, where a specific technical issue has guided research towards an analogous biological solution. Conversely, the mechanism by which burdock seeds attach themselves illustrates the 'solution-based' approach. This plant, commonly found in meadows, uses zoochory as a dispersal strategy, exploiting the passage of animals to passively transport its seeds to new territories. The characteristic hooks on these seeds allow them to adhere effectively to the fur of wildlife, while detaching easily during grooming activities. Georges de Mestral, the inventor of Velcro was inspired by this simple and repeatable hook-and-loop principle to develop a revolutionary fastening system, initially applied to the footwear industry, which has since become a universally adopted product.

These two design pathways are formalised in BID, they answer different stakes, but in themselves are not superior to one another [10].

Methodological approaches developed

The scientific community has developed various methodological approaches to guide and facilitate the biomimetic process. Zhang identifies three distinct phases in the evolution of these solutions [11]. The first focused on developing generic tools to facilitate the transition between biological and technical fields. The second phase focused on the development of processes and guidelines encompassing the entire biomimetic approach. A third, more contemporary phase focuses on the evolution towards integrated data platforms, with or without the incorporation of artificial intelligence technologies.

The growing interest in integrating advanced digital technologies into biomimetics (digital imaging, parametric design, artificial intelligence...) stems from their ability to facilitate the analysis, emulation, and optimisation of biological structures [12]. The incorporation of these technologies allows iterative feedback loops to be established between the different phases of biomimetic design, simultaneously enriching the design process and the definition of manufacturing requirements [12].

Biomimetics as a North star

As technological constraints diminish and design possibilities expand, engineering requires guiding principles to direct innovation efforts productively [7]. Biomimetics offers a compelling framework for this purpose.

Biological organisms represent the outcomes of billions of years of evolutionary optimization and adaptability under real-world fluctuating constraints including material availability, energy efficiency, and environmental compatibility [5]. Living systems exhibit remarkable properties achieved through hierarchical structuring, multiscale optimization, and efficient use of available materials [12]. Spider silk, for example, achieves extraordinary tensile strength and toughness through a hierarchical protein structure spanning from the molecular to the fibre scale, with mechanical properties rivalling synthetic high-performance materials despite being produced at ambient temperature and pressure using aqueous processing [13]. This biological example presents both inspiration and technical challenges for engineers. The complexity, multiscale nature, and multifunctionality inherent in biological systems make them difficult to

model accurately with computational tools and challenging to replicate with existing manufacturing processes [8, 14]. However, this difficulty aligns well with the need for ambitious targets that push technological capabilities forward. Biomimetics could thus provide direction for innovation efforts through modelling, simulation and manufacturing, which would stimulate the development of new tools and upgrades [11].

By essence, biomimetics is an interdisciplinary approach. Its integration to modelling and simulation tools would foster the crossing and collaboration of domains, which has been identified as a success factor [15]. Integrating biomimetics to modelling and simulation tools is an opportunity to broaden the problem-solving mindset of engineers [9]. Living principles differ from engineering ones, and so, looking at a problem through a biomimetics lens can sometimes shift perspective and possible solutions on said problem [16].

Another addition to the relevance of biomimetics as a north star for future development in modeling and simulation, is that through biomimetic design, the notions of sustainability, resource efficiency, and environmental compatibility can be accounted for [17]. This is encompassed in the broader term of biomimicry. By essence, a biomimetic product is not intentionally sustainable, but the design process is highly compatible with the inclusion of such factors [18]. This inclusion integrates biomimetic design in the demands set by societal objectives for the environmental transition.

Therefore, biomimetics is an opportunity for the modelling and simulation community to profit from the newly acquired capacities of their tools, and challenge them with the ability to model complex, multiscale, multifunctional systems. Biomimetics also brings an input of interdisciplinarity, and fits well with the consideration of sustainability factors, both essential elements in product design and development.

MATERIALS AND METHODS

Building on the arguments put forward by the literature in favour of biomimetics for modelling and simulation tools, a use case was set up. The following section presents the materials and methods used in a modelling and simulation-assisted biomimetic approach on the topic of anaerobic digestion (AD).

Use case presentation

Anaerobic digestion is a biological process in which microorganisms break down organic matter to produce biogas, a renewable energy source [19]. Current digesters offer satisfactory performances, but remain optimisable, as they encounter problems of preferential pathways and dead volume amongst others. On average, 30% of the introduced matter is not recycled in valuable elements [20]. The best digesters (with the strongest yield) are in the living world, in the rumen of bovines. The rumen is a large, pouchlike fermentation chamber, divided into dorsal and ventral sacs by muscular folds. A student project was built around the topic of optimizing the anaerobic digestion process through the development of a biomimetic digester, inspired by the rumen of cows. Students were Master 1 engineering students from UniLaSalle with specialisations raging from geosciences to nutrition, to life-sciences and agronomy. The purpose of the project was to model and run simulations on both the biological model, a rumen, and a biomimetic version of an anaerobic digester, built with the learnings they gathered from the rumen model.

Use case set up

Students had two weeks to build the models, run the simulations, and come up with a biomimetic digester. Their progression was guided by the design thinking framework. To narrow down the scope of potential upgrades, students were asked to focus on the hydrodynamics, and so mostly worked with Computational Fluid Dynamics (CFD). Modelling was conducted using the Xgenerative design software in the CATIA suite. Simulations were run through Xflow, using the Large Eddy Simulation (LES) approach. The latter was selected for its high accuracy, despite its demand for substantial computational resources. The entry parameter was set to a constant speed, and the exit parameter to atmospheric pressure.

RESULTS

Biological model

In the modelling process, students aimed to remain as close as possible to the biological system. The overall shape and architecture were replicated: one big pouch, split into two pockets, dorsal and ventral. Both are visible on the figure below, the dorsal pouch is the top one, and the ventral pouch, the bottom one.

The rumen is submitted to different forces through muscle contraction and movement that aim to move the content of the rumen around (grass and other elements eaten by the animal) so it can be digested by the bacterial mass. These movements also contribute to the physiological regurgitation of nutrition, which can then be chewed and salivated again, an action known as rumination. To model the rumination process, a circulation system was set up, which ensured proper breakdown of solid matter, and prevented congestion in potential dead volumes. Movement is induced through a pump, and reinjection is done by a pipe going from the ventral pouch to the dorsal pouch. The dorsal pouch is where solid matter is broken down and digested. To make sure that the solid matter stayed as long as digested in the dorsal

pouch, and that no amount of it would fall and clog the papillae or the recirculation system in the ventral pouch, a narrow grid was built in. It is only once reduced to particles by the digestate, that former solid matter can fall into the lower pouch.

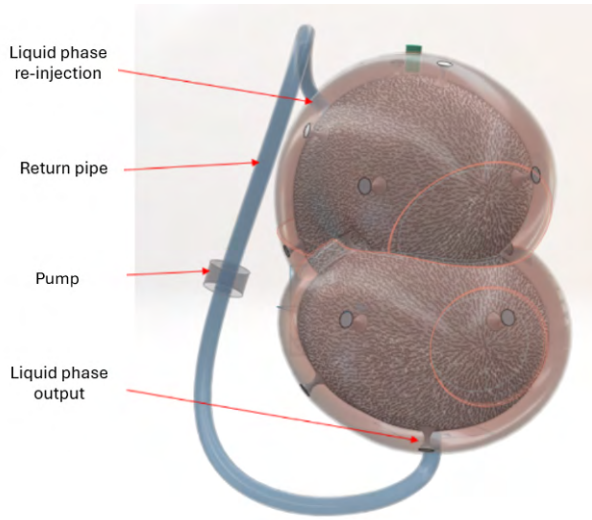

Figure 1a Biological model – a rumen

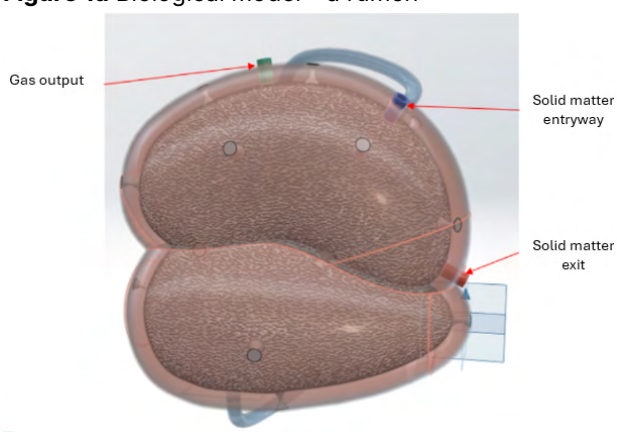

Figure 1b Biological model – a rumen

Rumen walls are lined on the inside with shaggy, tongue-like papillae that give it a textured, velvety appearance. Within the mindset of representing all components of the system, to be as reliable as possible in the visual but also functional representation, students integrated inner, smaller-scale structuration into their rendition.

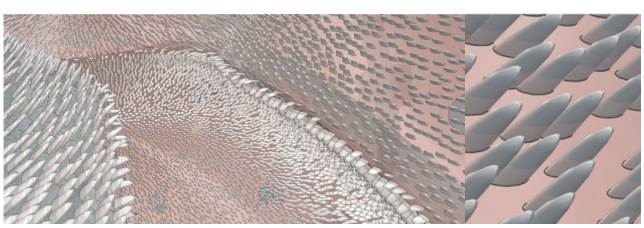

Figure 2 Inner wall layering in biological model

Biomimetic model

Anaerobic digestion can follow two main pathways depending on whether the handled feedstock is considered as solid or liquid. Solid-state digestion is built to manage unpumpable, and dry feedstock, where liquid-state digestion handles pumpable, slurry-like feedstock. Students observed CSTR solid and liquid-state digesters as reference systems to design the biomimetic versions. Both are of cylindrical shape, with the addition of a mixer for the liquid-state. Simulations were run to observe matter movement inside of traditional digesters. Below are the simulation results for the solid-state.

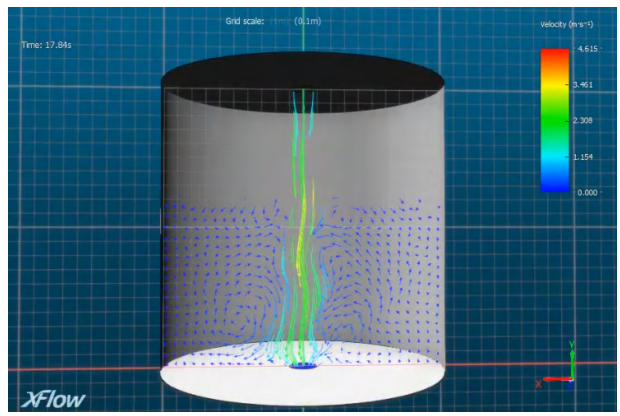

Figure 3 Solid-state digester fluid dynamic simulation Color gradient shows fluid velocity in $m.s^{-1}$

Fig. 3 simulates matter movement inside the solid-state digester. The first observation is that there is a definite preferential path to matter movement: a straight line from entry to exit points. The existence of a preferential path in the digester is not a desired effect. Moreover, the path visible here is the most direct one, and carries matter at a high speed (1.15 to 3.40 $m.s^{-1}$) which is not intended. Indeed, the objective is to have sufficient retention time in the digester to allow solid matter exposure with methanogenic bacteria. If the matter tends to go straight from entry to exist, in the fastest way through, retention time is not optimized. Also, if matter follows a preferential path, it means that the matter outside of this path gets stuck, creating dead weight, also an unwanted phenomenon.

Dead volume is the second observation that can be taken from this simulation result. The matter on the left and right of the preferential path moves very little, to not at all. Arrows on the graph do not indicate movement, and if so, only minimal and slow.

As displayed in Fig.4 below, the liquid-state digester also shows subpar fluid dynamics.

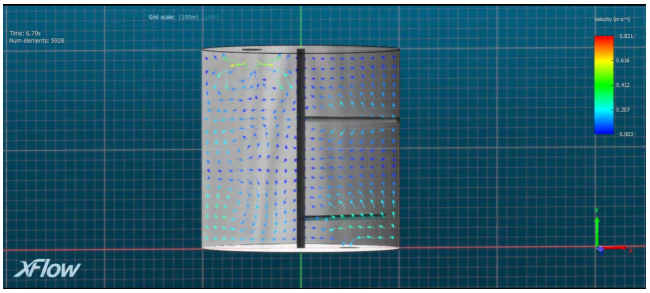

Figure 4 Liquid-state digester fluid dynamic simulation
Color gradient shows fluid velocity in m.s^{-1}

The action of the mixer ensures matter movement around the digester. It appears that matter is pushed around and sped up by the mixer into circle-like pathways. An observation that is worth underlining is that, compared to the solid-state digester, velocity around entry and exit point is not significantly bigger than in the rest of the digester, specifically at the exit point. This leads to wonder if matter is going through. Fig. leaves the question on the recirculation of matter in the digester. In the solid-state digester, the seemingly present issue was of too little retention time, but here, it might be of too big retention time. If feedstock doesn't get though, the bacteria could exhaust the needed molecules for methanogenesis, which would put a halt to production.

Two focus points can be extracted from the analysis of the non-biomimetic digesters: limiting dead volume, and preferential pathways and optimizing retention time.

The goal was then to build a biomimetic digester with improved hydrodynamics to deals with these issues. Students built two different systems for solid-state and liquid-state pathways and ran simulations of internal hydrodynamic flow. For both systems, the shape was designed as rumen-like, with two pockets (dorsal and ventral/upper and lower), and an oval architecture to avoid dead volumes, and allow the entire volume of the digester to be used effectively. Below are the simulation results for the solid-state biomimetic digester.

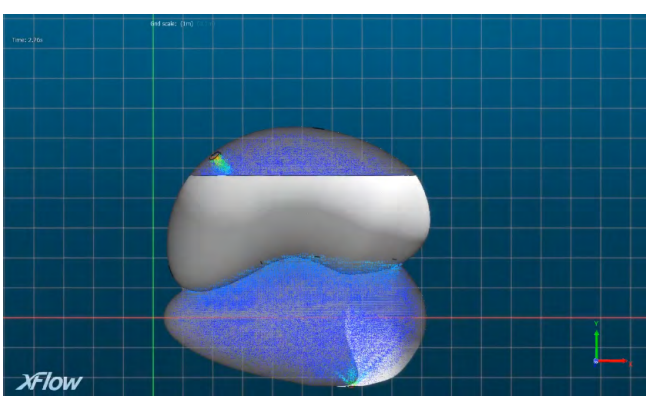

Figure 5a. Biomimetic solid-state digester fluid dynamics
Color gradient shows fluid velocity in m.s^{-1}

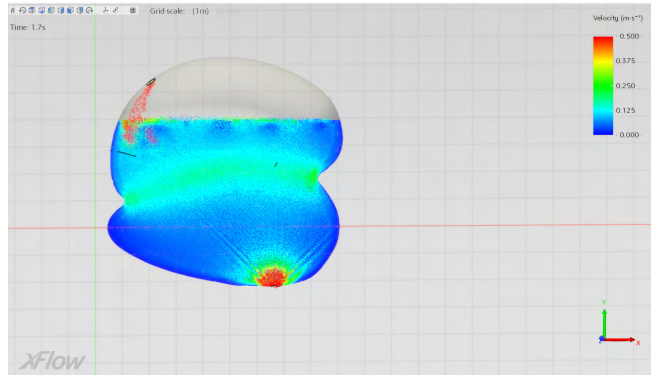

Figure 5b. Biomimetic solid-state digester fluid dynamics
Color gradient shows fluid velocity in m.s^{-1}

Fluid dynamic simulation in the biomimetic solid-state digester shows varying speed of matter inside of the system. The separation between dorsal and ventral pouches is a "fast-moving" spot with matter circulating around 0, 25 m.s^{-1}. It is worth noting that the acceleration occurs through the entire length of the separation and not specifically on one spot. It seems that matter is injected at high speed, slows down in the dorsal pouch, gradually speeds up going through the separating grid, and gradually slows down in the ventral pouch before being pushed out by overflow.

The following results were obtained for the liquid-state biomimetic digester.

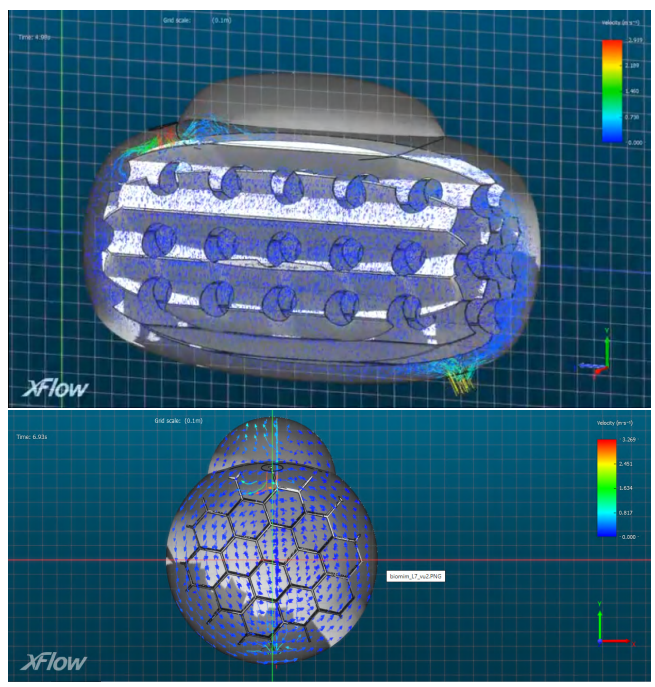

Figure 6. Biomimetic liquid-state digester fluid dynamics
Top figure : side view ; Bottom figure : front view
Color gradient shows fluid velocity in m.s^{-1}

The main takeaway from these results is that the

redesign of the mixer and the overall shape of the digester are very beneficial to matter movement. Gas production is stored up in the superior pouch, and fluid simulations show the entire volume of matter circulating around all the available space. The light blue strands in Fig. indicate that once matter is introduced in the digester, it does not remain at top or bottom parts but goes through the hive-like mixer. Overall speed is relatively slow (0 to 0, 75 m/s), but movement is homogeneous, which is wished for.

To sum up, a model of a biological system, a rumen, was designed and underwent simulation testing to gather knowledge on its interior fluid dynamics. Fluid dynamics were also observed in non-biomimetic digesters in two pathways for anaerobic digestion: solid-state and liquid state. The analysis highlighted issues of preferential pathways and dead volume in the digester current design. The insights on the inner workings of the rumen, which stands at the top of the yield pyramid for methane production, were used to design biomimetic digesters, with the specific focus of hydrodynamic improvement. Initial simulations of the biomimetic systems show an upgrade in fluid circulation, and so a reduced tendency for preferential pathways and a minimization of dead volume.

Results are still preliminary and superficial and stand more as a qualitative validation of the overall improvement of digesters through a biomimetic approach. Additional data and testing are needed to push the biomimetic systems further. For example, it would be interesting to calculate retention time and biogas production in the biomimetic systems and compare them to the performances of current systems (non-biomimetic digesters). Another opportunity would be to dig deeper on the papillae structures that maximize absorption surface. It was not studied here as only hydrodynamics were of concern, but they remain an addition proposed by the biological system for the technical one and probably play a substantial role.

DISCUSSION

The central point argued in this paper is that biomimetics and modelling and simulation tools have the potential to work in a mutually reinforcing relationship. Recent works support and put forward arguments showcasing this synergy [5, 8, 14, 21].

On the one hand, the literature and the studied use case support that biomimetics can contribute to the development and improvement of modelling and simulation tools. Following a biomimetic approach indicates the study of biological systems, which are known as complex, multiscale, and multifunctional. This degree of intricate, hierarchal geometry stands on the current edge of the modelling and simulation capacities, despite recent improvements. It is only through these improvements that trying to model biological systems is conceivable, still, not yet easily feasible. The rumen use case discussed previously led to conclusions that support the positioning of biological models on the limit of what tools are currently able to model. Indeed, the design of the rumen pushed the software to its computational limits, specifically, in attempting to replicate the lining of rumen walls with papillae. Due to the sheer scale and number of required features, students were forced to model the papillae at a larger-than-natural scale to achieve convergence. This observation underlines that modelling and simulation tools need to improve in some way to compute and navigate such complexity (ie. multiscale structure). Ability to do so is essential as in biological systems, multiscale structure is often crucial to the performance of a specific function. So, biomimetics can contribute to the development of modelling and simulation tools by challenging them with the innate complexity of biological systems. For example, biology works with both hard and soft components to profit of their respective advantages (resistance, flexibility....), where engineering processes are focused on “hard materials” (steel, concrete...). Modelling biological systems could push the development and expand the versatility of Fluid-Structure Interaction (FSI) capabilities in software. In the presented use case, it would have been more accurate to represent rumen wall deformability and observe the influence of such capacity on fluid dynamics, but it was not feasible.

On the other hand, modelling and simulation tools could also bring beneficial elements to biomimetics.
Their integration to product development has increased product complexity, enabled early validation of performance, as well as quick and cheap iteration loops to test design options, and multicriteria analysis to identify the most suitable design alternative overall (finances, manufacturing capabilities, level of performance...). The biomimetic process would benefit from such additions as it faces struggles, notably related to the disruptiveness of the systems it leads to. This hypothesis opens the door to future work on the validation of mentioned benefits in the case of biomimetic design.

An essential part of the biomimetic process is transferring the principle of the biological function to the chosen field of application. If done poorly, it is likely that the resulting system will not deliver to expected levels. Proper transfer is deeply related to biological system understanding. Visual representations have always been used to describe, and comprehend the way systems, including living ones, function. Modelling and simulation tools lead to advanced visuals, gathering more information of the modelled system with the addition of one dimension (2D to 3D). A double hypothesis stands that through the standard use of modelling and simulation tools in biomimetic design, biological system

understanding would be improved, and resulting biomimetic systems would meet expectations.

In the use case, without modelling and simulation tools, getting to a deep and intricate enough level of understanding of the rumen, and the fluids dynamics it holds, would have been extremely complicated, and vastly more expensive. Using a porthole cow would not have necessarily brought the needed data, on top of being an ethically questionable approach. Perhaps, required data could have been obtained through the ingestion and monitoring of movements, of a traceable sensor by a cow, but it remains a demanding experiment in terms of resources.

Going further than improving the understanding of systems of interest, models and simulations could help figure out the viability frame of biological strategies. As seen in the use case with papillae structures, biology often works at small scales (nano and micro). Being able to assess if the desired function is governed by scale, or other parameters, would be a game-changing ability. In the example, knowing if absorption improvements brought by structures like papillae are only present if said structures are of nano or microscale, is an extremely important fact. Indeed, it indicates that nano or micro texturation would be an essential element to the biomimetic system for it to be functional. Information that signals early-on future manufacturing constraints. A hypothesis is here that models and simulations, specifically parametric models, could assist in the determination of the viability frame of a function, thus informing on its usability by a specific actor, depending on the field of application and manufacturing capacities.

All hypotheses made above must be tested through rigorous experimentation. They stem from the authors reflexion around the relevance of biomimetics and modelling tools together but are not yet verified statements.

CONCLUSION

The aim of this article was to showcase the mutually beneficial relationship between biomimetics and modeling and simulation tools. The difficulty to model biological systems challenges modelling and simulation tools and stimulates their development and improvement to be able to cope with that level of complexity. Modelling and simulation tools are believed to be relevant to assist in a core step of biomimetics, which is the transfer of the biological function to the chosen field of application. This assistance would stem from the improvement of the understanding of biological systems through informative visuals. Parametric modelling could play a break-through role in determining the viability frame of biological functions, thus advising on biomimetic system feasibility. it is conceivable

Future work could focus on the experimental validation of the hypothesis around modelling and simulation tools for biomimetics, as well as on building the methodological steps needed to standardise the use of such tools in biomimetic design.

The end goal of such work is to facilitate the wider adoption of biomimetics as a problem-solving design approach, ultimately promoting the emergence of biomimetic innovations across diverse applications.

ACKNOWLEDGEMENTS

The authors wish to express their gratitude to the students of UniLaSalle for their work on the use case.

Acknowledgements to IKOS Consulting and Ceebios for their financial contributions to the PhD work underpinning this research.

AUTHOR IDENTIFIERS

Author ORCID: 0009-0008-0273-6819

REFERENCES

1. Van der Auweraer H, Anthonis J, De Bruyne S, Leuridan J. Virtual engineering at work: the challenges for designing mechatronic products. Engineering with Computers. 2013 Jul 1;29(3):389–408. https://dl.acm.org/doi/10.1145/3351095.3372871
2. Dodgson M, Gann DM, Salter A. The impact of modelling and simulation technology on engineering problem solving. Technology Analysis & Strategic Management 19:471-489 (2007). https://doi.org/10.1080/09537320701403425
3. Nguyen D, Martin P. Product design-process selection-process planning integration based on modeling and simulation. Int J Adv Manuf Technol 77:187-201 (2014). https://doi.org/10.1007/s00170-014-6446-7
4. Sinha R, Paredis CJJ, Liang VC, Khosla PK. Modeling and simulation methods for design of engineering systems. Journal of Computing and Information Science in Engineering 1:84-91 (2000). https://doi.org/10.1115/1.1344877
5. Fattepur G, Patil AY, Kumar P, Kumar A, Hegde C, Siddhalingeshwar IG, Kumar R, Khan TMY. Bio-inspired designs: leveraging biological brilliance in mechanical engineering—an overview. 3 Biotech 14: (2024). https://doi.org/10.1007/s13205-024-04153-w
6. Luo GJ, Dyson RJ, Thomas-Seale LEJ. A Bio-Inspired Mathematical Approach to Design for Additive Manufacturing. In: Leveraging Transdisciplinary Engineering in a Changing and Connected World [Internet]. IOS Press; 2023 [cited

2026 Jan 24]. p. 92–101. Available from: https://ebooks.iospress.nl/doi/10.3233/ATDE230601
7. Abebe R, Barocas S, Kleinberg J, Levy K, Raghavan M, Robinson DG. Roles for computing in social change. In: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency [Internet]. New York, NY, USA: Association for Computing Machinery; 2020 [cited 2026 Jan 26]. p. 252–60. (FAT* '20). Available from: https://dl.acm.org/doi/10.1145/3351095.3372871
8. Singh B, Basri AA, Yidris N, Pai R, Ahmad KA. Introduction to biomimetics, modelling and analysis. Series in BioEngineering :1-20 (2024). https://doi.org/10.1007/978-981-97-1017-1_1
9. Helms M, Vattam SS, Goel AK. Biologically inspired design: process and products. Design Studies 30:606-622 (2009). https://doi.org/10.1016/j.destud.2009.04.003
10. Hashemi Farzaneh H, Helms MK, Muenzberg C, Lindemann U. Technology-pull and biology-push approaches in bio-inspired design - Comparing results from student teams. In Dubronik, Croatia; 2016.
11. Zhang J, Kestem L, Wommer K, Wanieck K. Biomimetic tools: insights and implications of a comprehensive analysis and classification. Bioinspir. Biomim. 20:026014 (2025). https://doi.org/10.1088/1748-3190/adaff6
12. Efstathiadis A, Symeonidou I, Tzimtzimis EK, Avtzis D, Tsongas K, Tzetzis D. Parametric design and mechanical characterization of a selective laser sintering additively manufactured biomimetic ribbed dome inspired by the chorion of lepidopteran eggs. Biomimetics 10:1 (2024). https://doi.org/10.3390/biomimetics10010001
13. Eisoldt L, Smith A, Scheibel T. Decoding the secrets of spider silk. Materials Today 14:80-86 (2011). https://doi.org/10.1016/s1369-7021(11)70057-8
14. Singh B, Ahmad KA, Pai R. Biomimetic modeling and analysis using modern architecture frameworks like CUDA. Series in BioEngineering :223-239 (2024). https://doi.org/10.1007/978-981-97-1017-1_10
15. Graeff E, Letard A, Raskin K, Maranzana N, Aoussat A. Biomimetics from practical feedback to an interdisciplinary process. Res Eng Design 32:349-375 (2021). https://doi.org/10.1007/s00163-021-00356-x
16. Nagel JKS, Nagel RL, Stone RB. Abstracting biology for engineering design. IJDE 4:23 (2011). https://doi.org/10.1504/ijde.2011.041407
17. Ilieva L, Ursano I, Traista L, Hoffmann B, Dahy H. Biomimicry as a sustainable design methodology—introducing the 'biomimicry for sustainability' framework. Biomimetics 7:37 (2022). https://doi.org/10.3390/biomimetics7020037
18. Jatsch AS, Jacobs S, Wommer K, Wanieck K. Biomimetics for sustainable developments—a literature overview of trends. Biomimetics 8:304 (2023). https://doi.org/10.3390/biomimetics8030304
19. Meegoda JN, Li B, Patel K, Wang LB. A review of the processes, parameters, and optimization of anaerobic digestion. IJERPH 15:2224 (2018). https://doi.org/10.3390/ijerph15102224
20. Rocha-Meneses L, Zannerni R, Inayat A, Abdallah M, Shanableh A, Ghenai C, et al. Current progress in anaerobic digestion reactors and parameters optimization. Biomass Conv Bioref. 2025 Dec 1;15(23):29573–96.
21. Houette T, Dibia M, Mahabadi N, King H. Pullout resistance of biomimetic root-inspired foundation systems. Acta Geotech. 19:1191-1210 (2023). https://doi.org/10.1007/s11440-023-02118-6

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Comparative Analysis of Sequential Active Learning Approaches: Statistical Design of Experiments versus Bayesian Optimisation

Daniel V. Batista[a] and Marco S. Reis[a]*

[a] University of Coimbra, CERES, Department of Chemical Engineering, Coimbra, Portugal
* Corresponding Author: marco@eq.uc.pt.

ABSTRACT

As chemical processes become increasingly complex and costs of experimentation increase, understanding the practical effectiveness of Active Learning methodologies is essential. In this regard, an ongoing debate is occurring within the research community about the use of Design of Experiments (DOE) and Bayesian Optimisation (BO). However, this debate is limited by the scarcity of systematic comparative studies. Therefore, this work provides a comparative analysis of two widely adopted data-driven optimisation approaches: DOE and BO. The comparison is conducted across two distinct case studies reflecting different levels of complexity, regarding the quantity and variety of input variables involved. The first case study represents a realistic in silico experimental scenario, with multiple decision variables of different types (continuous, categorical and mixture), and two distinct single-objective optimisation goals, while the second one considers a simpler, well-known benchmark model with just two input continuous variables. Both studies were designed and analysed, while acknowledging the inherent conceptual and operational differences between DOE and BO. The results showed that a sequential DOE strategy, adapted here for optimisation purposes, consistently outperformed classical BO under the tested conditions. Particularly in terms of convergence towards the target response, robustness in identifying the optimal region, and overall experimental budget efficiency. Rather than asserting the superiority of one methodology over the other, this work highlights the need for case-specific adaptation when applying data-driven optimisation strategies in Chemical Engineering. The findings emphasise that theoretical guarantees alone under relatively strict assumptions are insufficient and must be complemented by problem-driven evaluation to support informed decision-making.

Keywords: Statistical Design of Experiments (DOE), Bayesian Optimisation (BO), Optimisation, Active Learning (AL) approaches

INTRODUCTION

Systematic experimentation and mathematical optimisation have played a pivotal role in the advancement of Chemical Engineering Science, as evidenced by the development of new products and processes and their effective transfer and scaling in industrial applications.

Over the years, various optimisation strategies have been developed to address distinct objectives while taking the specific case study into account as much as possible. However, identifying an optimal strategy can be a challenging, time-consuming, and costly endeavour, particularly when the objective function and process model are not explicitly known and must be induced through experimentation.

This situation can prove even more challenging when multiple objectives must be considered and balanced simultaneously (e.g., yield, robustness, cost) and when variables of different types are present (e.g., continuous, multi-level categorical, mixture). Therefore, the primary challenge may lie in establishing an effective data-driven optimisation methodology that can efficiently optimise the systems towards a desired goal, within a reasonable budget of experiments.

Until recently, this activity has been supported by the extensive framework of Statistical Design of

Experiments (DOE) [1-3], widely applied in both research and industry, a methodology that traditionally enables careful planning of experiments followed by a separate optimisation stage. However, a new category of methods has recently attracted considerable interest for its conceptually intuitive formulation, flexibility, and reduced reliance on strong problem-specific assumptions. In particular, Bayesian Optimisation (BO)[4-6] offers a unified approach that incorporates both modelling and optimisation within a single framework.

Furthermore, authors tend to dismiss DOE because it typically relies on a pre-determined sampling pattern, which can hinder adaptation to features that appear during the course of experimentation [7]. Additionally, BO uses spatial correlation, avoiding the matrix singularity problems that can complicate DOE, namely when dealing with mixture variables.

Currently, it is unclear which of the two alternatives is the better choice, and when. The decision-making process can be challenging, particularly due to a lack of comparative evidence on the advantages and disadvantages of each strategy in different practical situations. This ambiguity may hinder engineers and analysts in selecting the most appropriate Active Learning (AL) approach for their specific case, at the cost of being suboptimal in terms of data efficiency and overall experimental costs.

To address such current information bottleneck, we present a comparative analysis of the two approaches (DOE and BO) in two *in silico* scenarios reflecting different application realities, where sequential experimentation is carried out with the goal of optimisation.

METHODHOLOGY

Active Learning approaches

In the comparative analysis, two AL approaches were considered (DOE and BO). Regarding DOE, we tested an adaptation of the base methodology that implements a sequential AL approach. This approach initially explores an initial sampling space, followed by sequential exploitation, and it was adapted to each case study scenario, as described in the following sections.

As for BO, the classical methodology was adopted by building a Gaussian process-based probabilistic surrogate model from observed data. This model captures mean predictions and uncertainties, which are fed to the acquisition function, to guide the selection of informative points for sampling, while balancing exploration and exploitation. In this work, we have selected Expected Improvement (EI) for the acquisition function. The model is updated successively, with new points sampled to maximise the acquisition function until the predefined optimisation budget is reached.

Although the statistical foundations underlying DOE and BO are quite different, the study was carefully designed to provide a fair and realistic comparison, minimising potential biases towards either approach. At the same time, the analysis acknowledged the inherent advantages of each method, such as the reduced number of experiments required for starting BO, owing to its sample-efficient nature [5]. Both methodologies were implemented in Python and customised for each case study. For BO, the Ax Python library was used to conduct the study.

Case studies

The practical comparison was conducted on two distinct case study scenarios, with different levels of complexity regarding the number and typology of input variables involved.

The first case study considers a realistic simulated experimental scenario with two single-optimisation opposing goals (maximisation and minimisation) and several variable types, including continuous, mixture, and multi-categorical variables. The scenario also considers a predetermined experimental budget and target response value.

The second case study represents a typical maximisation problem based on a simple and well-established mathematical function with a known optimum point. In this scenario, the study flexibly covers different options for the initial sample space used for both strategies. For simplicity, contrary to the first case study, only continuous variables are used, and the system is low-dimensional (2 input variables). The scenario also considers various target experimental budgets and a predefined bounded space for analysis.

CASE STUDY 1: IN SILICO OPTIMISATION OF LNP SYNTHESIS

Comparison Study

For the first case study, an *in silico* system was adopted, based on a model previously developed from experimental data collected [8]. This model was originally discussed and analysed in earlier work [9] and was adapted for the present study.

The *in silico* system represents the synthesis of lipid nanoparticles (LNPs) by microfluidics and acts as a computational surrogate of a real experimental setting. Due to the high structural complexity of the underlying simulation model, it was treated as unknown for the purposes of this study, effectively constituting a black-box system.

According to the simulation model's specifications, several types of input variables were considered, and their ranges were selected while respecting the feasibility limits of the simulation model, as shown in **Table 1**.

Table 1: Input variable ranges and types in simulation model.

Input Variables	Ranges	Type
Ionizable Lipid	[0.3;0.6]	Mixture
Structural Lipid	[0.185;0.555]	Mixture
Helper Lipid	[0.1;0.28]	Mixture
PEG Lipid	[0.01;0.045]	Mixture
N/P ratio	[6, 14]	Continuous
Total Flow Rate	[1, 3] mL/min	Continuous
Ionizable Lipid Type	H101 or H102 or H103	Categorical

The model also included two response variables: *in vivo* potency and the average size of the LNP. These responses were linked to separate single-objective optimisation goals: potency should be maximised and particle size, minimised. The goals and predefined targets to meet are shown in **Table 2**.

Table 2: Response targets and goals.

Responses	Goal	Target
Potency (%)	Maximize	98 %
Size (nm)	Minimize	90 nm

For the DOE approach, an initial sample space was generated using a D-optimal design strategy for two different model structures with distinct levels of complexity (M2 and M3). Model M2 represents a lower-complexity structure and includes main effects of the mixture components, nonlinear blending terms between pairs of mixture components, and second-order interaction terms among all the remaining input variables. Model M3 extends M2 by incorporating additional Scheffé cubic mixture terms, such as nonlinear blending effects among three mixture components. As the number of model terms increased with model complexity, the required number of experiments also increased. To enable a fair comparison, a minimum initial sample size was imposed for both models, corresponding to 50 experiments, which met the minimum requirements of the more complex model.

Therefore, for the initial DOE sample space, we considered different levels of exploration, using 50, 55, and 60 initial experiments, and a pre-established target budget of 55, 60, and 70 experiments for each case, respectively, during optimisation.

Once the initial sample space was generated, the models were fitted using least-squares estimation. The corresponding optimal input conditions were then determined through an adapted sequential least-squares programming (SLSQP) nonlinear optimisation procedure for categorical variables.

The newly identified optimal input conditions were evaluated using the *in silico* model, after which the existing dataset was appended, and the data-driven model was re-estimated. This sequential process of model fitting and optimisation was repeated iteratively until the predefined experimental budget was exhausted. A schematic representation of this workflow is provided in **Figure 1**.

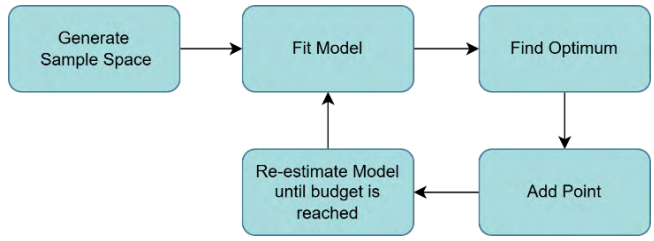

Figure 1. DOE framework in Case Study 1.

In the case of Bayesian optimisation, a classical BO framework was adopted, following the previously described reasoning, until the same target budget established for the DOE approach was reached. Given the reduced information requirements of BO, a smaller initial sample space (N_{init})was adopted using a practical rule-of-thumb, based on the number of variables (N_{var}), as shown in **Equation (1)**. The sample space was generated using SOBOL sampling, a quasi-random (low-discrepancy) space-filling initialisation method, within the established ranges in **Table 1.**

$$N_{init} = N_{var} + 1 \quad (1)$$

Results and Discussion

For both active learning approaches, 1, 000 simulation runs per approach were conducted for each scenario to characterize and take into account stochastic variability in the results, while maintaining the same initial sample space across simulations. In addition to the baseline noise level inherent to each simulation model, an alternative scenario with a 20% increase over the original noise level was also investigated. This additional scenario was included to assess the robustness of both approaches under more challenging, noise-dominated conditions.

For each simulation and each AL approach, the best response value achieved (maximum *for* potency or minimum for LNP size) was recorded along with the corresponding experiment location (sequential number of the run where it was obtained).

Additionally, the computational time required to complete each optimisation in DOE and BO were recorded for both responses. The results showed that on a 13th Gen Intel i7-1365U processor, the optimisation time for DOE was approximately between [1.34, 8.90] s for potency and [1.40, 7.32] s for size. In contrast, BO required longer computational time, approximately between [200, 453.70] s and [188.79, 269.53] s for potency and size, respectively.

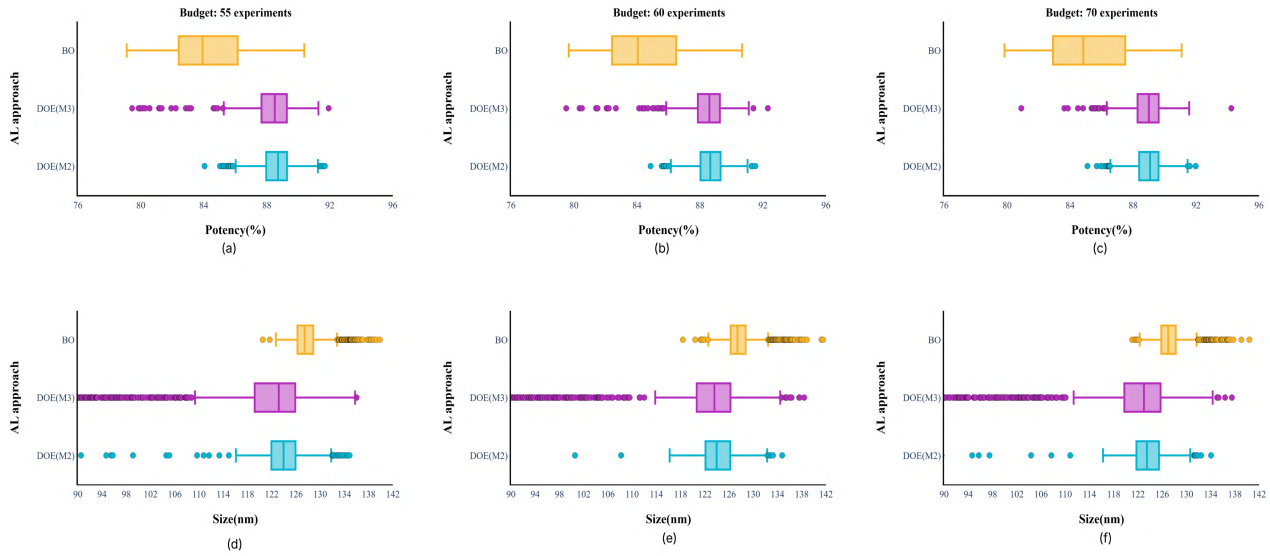

Figure 2. Distribution of the responses' best values obtained from 1,000 simulations per AL approach (BO, DOE(M3) and DOE(M2)) at original noise level. (a)-(c) Potency' best values (maximum); (d)-(f) Size's best values (minimum).

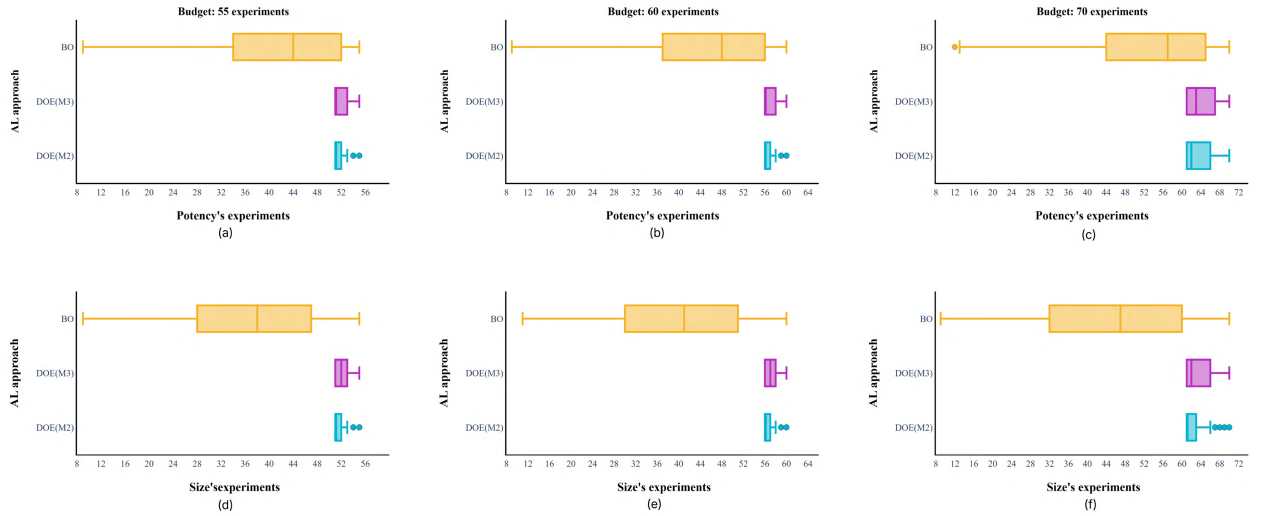

Figure 3: Distribution of the location of the experiment's best value, obtained from 1,000 simulations per AL approach (BO, DOE(M3) and DOE(M2)) at original noise level. (a)-(c) Potency; (d)-(f) Size.

The distributions of the optimality outcomes were subsequently visualised using box plots (**Figure 2** and **Figure 3**) to facilitate a comparative assessment of the optimisation performance.

The results in **Figure 2** show that, in both responses, DOE outperformed BO across all established target experiment budgets, as indicated by the spread of box-plot values in each response and by the closeness to target response values (**Table 2**). In particular, for size (**Figure 2 (d-f)**), BO failed to achieve values sufficiently close to the target value (90 nm).

The disparity in performance becomes particularly pronounced when examining the median values depicted in the box plot. These values clearly illustrate a central tendency in DOE that aligns more closely with the target value than BO. Furthermore, in the case of DOE, comparing both DOE strategies, DOE(M2) and DOE(M3), reveals a similar median value, suggesting that, regardless of model complexity, both approaches generally perform similarly in optimisation for this case study.

Lastly, as shown in **Figure 3**, BO exhibits greater variability in pinpointing the optimum location for both

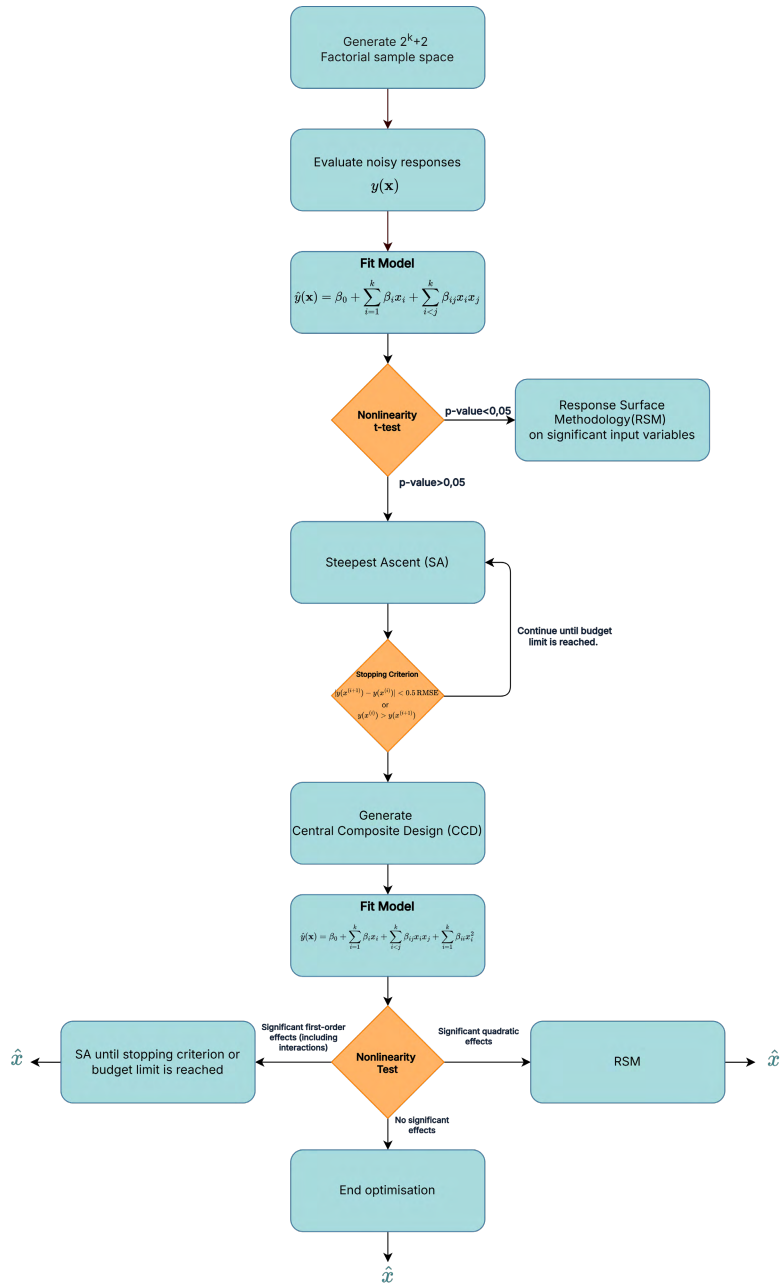

Figure 4: DOE framework in Case Study 2.

responses than DOE. This increased variability may be attributed to the exploratory nature of the BO process during the earlier stages of the optimisation phase.

It is important to note that, although not presented in this section, a consistent trend of results was observed regardless of the increased noise level, as shown in the Supplementary Material, **Figure A.1** and **Figure A.2**.

CASE STUDY 2: MATHEMATICAL BENCHMARK FUNCTION

Comparison Study

To complement the experimentally driven Case Study 1, this second case study evaluates and benchmarks the performance of active learning approaches in a more controlled benchmark scenario, independent of

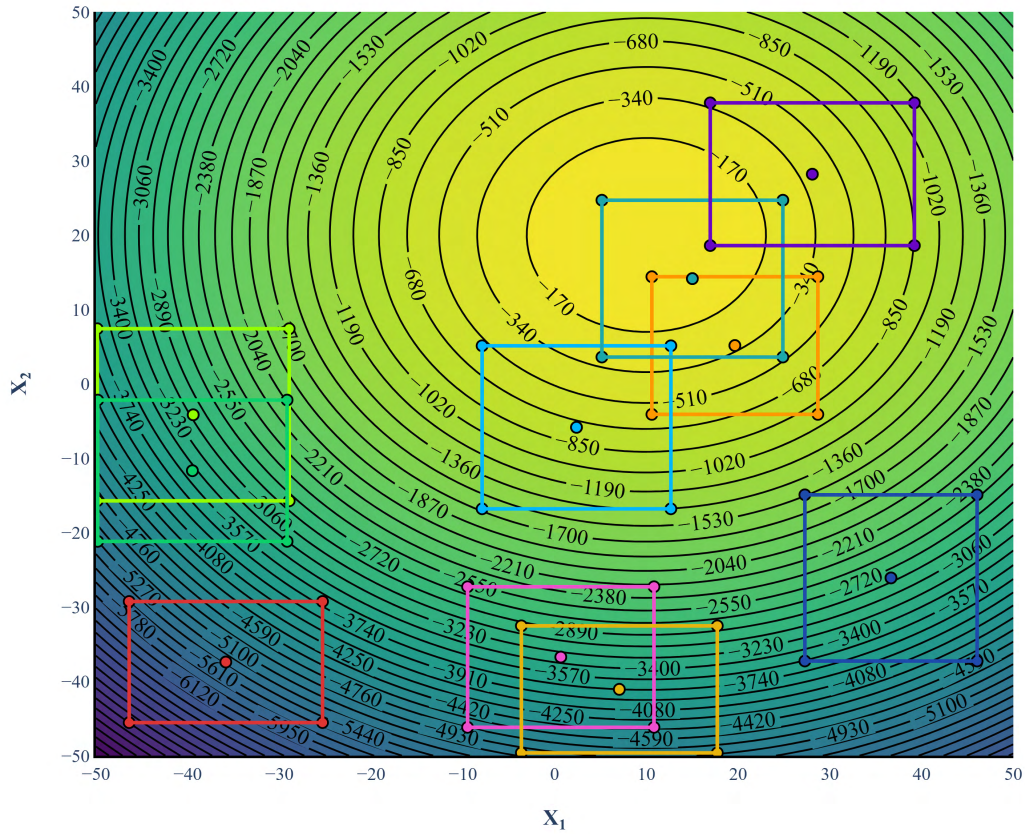

Figure 5: Design "Box" factorial sizes from ten simulations distributed across the response surface.

the complexities and uncertainties inherent in experimental systems, where the ground truth is often unavailable or costly to obtain. Therefore, for this case study, we considered a classical maximisation problem based on a well-known two-variable concave function, $f(x)$, shown in **Equation (2)**, which possesses a known global maximum at (10, 20) with an optimum value of 0.

$$f(x) = -(x_1 - 10)^2 - (x_2 - 20)^2 \quad (2)$$

Although the underlying function is known a priori for validation, it is treated as unknown during the optimisation stage for both approaches, with synthetic noise added to the responses to simulate realistic conditions. To reflect stochastic variability, responses were generated from **Equation (2)** with additive Gaussian noise, $y(x) = f(x) + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. The standard deviation was set as $\sigma = 1.6525$.

Following the methodology from the first case study, the DOE approach was tailored to fit the current problem, namely the different postulated model structure that implies a distinct strategy for its optimisation. However, the sequential active learning principle of exploration followed by exploitation is preserved. The DOE framework was developed in Python, following the workflow shown in **Figure 4.**

The DOE began by generating an initial two-level full factorial design with two replicated central points ($2^k + 2, k = 2$), while adhering to the bounded range for both variables of [-50, 50].

The central point (x_{centre}) was randomly selected from a uniform distribution within this range ($x_{\text{centre}} \sim U(-50, 50)$). The high and low levels of both input variables (X1 and X2) were determined by adding and subtracting, respectively, a predefined distance (l) sampled from a normal distribution, $l \sim N(10, 1)$, thereby allowing for multiple plausible initial sample configurations and design "box" sizes across various simulations (ten in this case study), as shown in **Figure 5**.

Once the initial dataset was established, a least-squares model was fitted to the simulated noisy responses and tested for nonlinearity using a statistical hypothesis test [3]. Sequential optimisation then followed two possible pathways: (i) Response Surface Methodology (RSM) via a central composite design (CCD) on significant inputs if significant nonlinearity was detected, or (ii) Steepest Ascent (SA) starting from the x_{centre} if no evidence of nonlinearity was observed.

Note that the SA here conducted, considers the

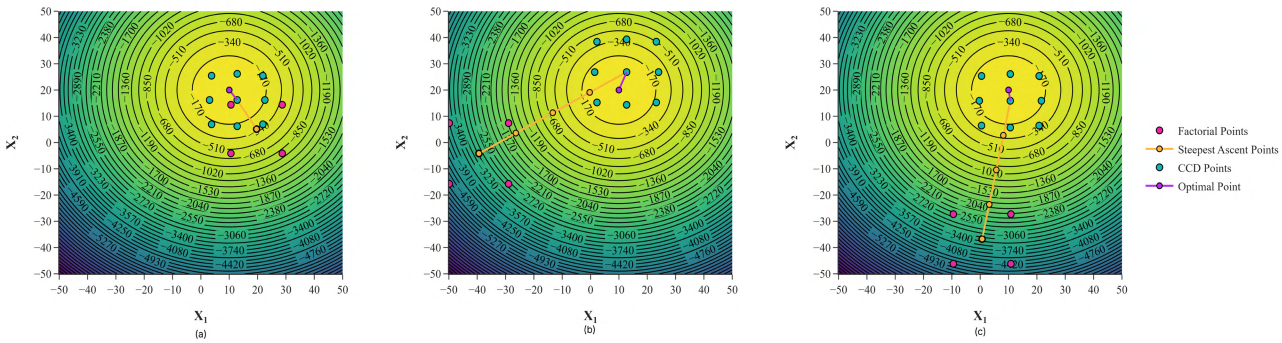

Figure 6. DOE optimisation path across the response surface. (a)-(c) First three simulations (Remaining simulations, please consult the Supplementary Material, **Figure B.1**).

fitted model with second-order interactions and recalculates the gradient at each new optimization iteration, as specified by **Equation (3)**.

$$x^{(i+1)} = x^{(i)} + \lambda \frac{\nabla \hat{y}(x^{(i)})}{\left\| \nabla \hat{y}(x^{(i)}) \right\|_2}, \lambda = \sqrt{2} \quad (3)$$

Additionally, the SA process continued until either a set number of experiments (50 in this case) was exhausted, or the stopping criterion was reached. Note that if the stopping criterion was not met, the optimisation could either continue using SA for additional iterations or be terminated.

If the stopping criterion was met during SA, a CCD was generated from the last point, and the model was refitted in the new sample space. Lastly, RSM was applied when significant quadratic effects were identified; otherwise, the SA continued until convergence or the last point in the previous SA was taken as the optimal point.

For the sake of illustration, some results derived from the sequential DOE approach during optimisation are presented in **Figure 6**; for more detailed information please refer to the Supplementary Material, **Figure B.1**.

It is important to note that, as shown in **Figure 6** and **Figure B.1**, all ten simulations during optimisation chose an SA approach, followed by RSM and stopped before reaching the experiment budget of 50.

The total number of experiments conducted was recorded for each simulation, along with the estimated optimal input ($\hat{x}$) and the corresponding noiseless response, $f(\hat{x})$, using Equation (2).

Additionally, across ten simulations, the Euclidean distance from $\hat{x}$ to the true optimal point of $f(x)$, (10, 20) was calculated using the **Equation (4)**.

$$d(\hat{x}) = \sqrt{(\hat{x}_1 - 10)^2 + (\hat{x}_2 - 20)^2} \quad (4)$$

In the case of BO, the classic BO approach was evaluated across three variants to assess the influence of initial knowledge: BO-N_{min}, BO-N_{DOE}, and BO-N_{max}.

For BO-N_{min}, a minimal sample space of three points, as defined by **Equation (1)**, was generated using SOBOL across ten independent runs. BO-N_{DOE} was initialised with the same initial sample space knowledge as the DOE approach, consisting of six points, across ten simulations. Lastly, BO-N_{max} employed a SOBOL design of the same size as the DOE initial sample space (six points) across ten independent runs. Note that for BO, a target budget of 50 experiments was used, matching DOE, and a smaller budget of 25 experiments was also considered for comparison.

The optimal point (maximum) at each strategy and simulation was recorded, along with the associated metrics, $d(\hat{x})$ and $f(\hat{x})$.

Results and Discussion

Figure 7 presents a summary of results for both AL approach, based on the selected metrics ($d(\hat{x})$ and $f(\hat{x})$). The numerical data are available for consultation in Supplementary Material, **Table B.1**.

The results indicate that, overall, the DOE method, which had lower budget requirements for optimisation, consistently outperformed BO by predicting input values ($\hat{x}$) that were nearer to the true optimal point (10, 20) under noisy conditions, **Figure 7 (a)**.

Additionally, in terms of the optimal value, DOE results were generally closer to the reference optimum value (0), **Figure 7 (b)**.

Although BO employed higher budgets, it did not consistently achieve similar proximity to the optimal values. Nevertheless, it is important to highlight the differences between the BO strategies used, as shown in **Figure 8**.

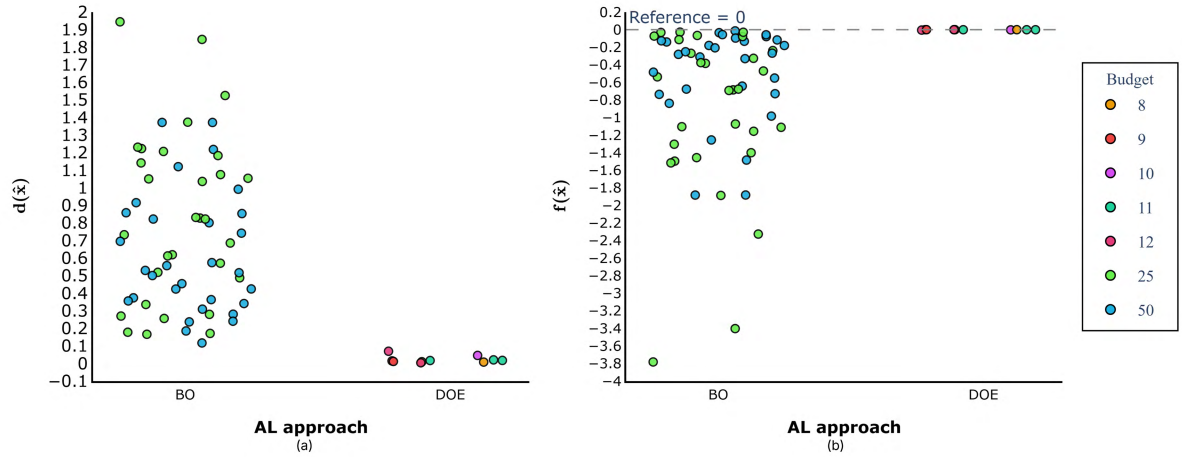

Figure 7. Results of $d(\hat{x})$ and $f(\hat{x})$ per the AL approach, coloured by the experiments' budget. (a) $d(\hat{x})$ values; (b) $f(\hat{x})$ values.

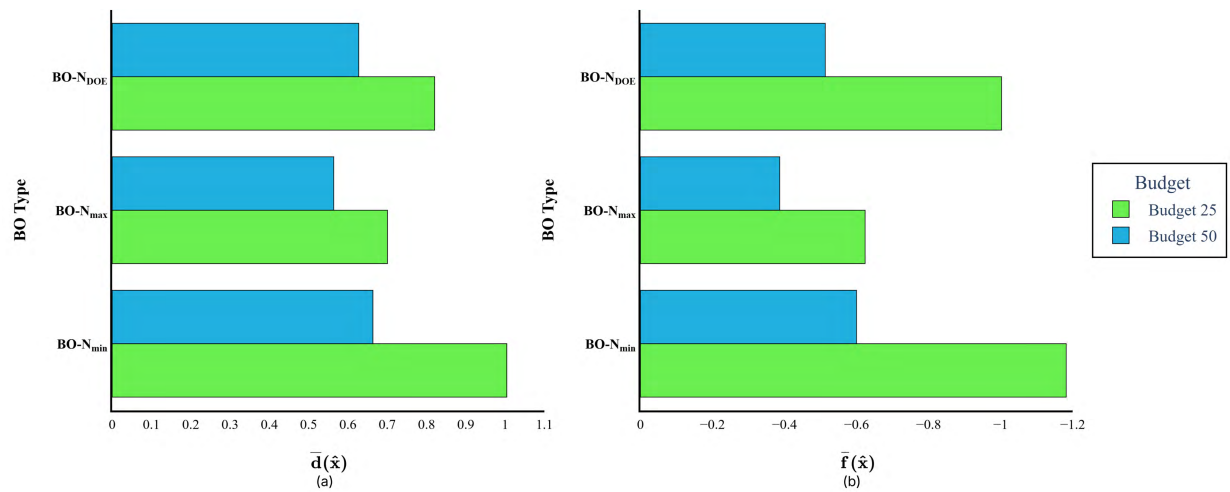

Figure 8. Results of average $d(\hat{x})$ and $f(\hat{x})$ per the BO Type, coloured by the two experiments' budget (25 and 50). (a) $\bar{d}(\hat{x})$ values; (b) $\bar{f}(\hat{x})$ values.

According to these findings, in the case of BO for both strategies, a lower experimental budget (25) on average performed worse, as indicated by the metric scores. For both budgets, BO-N_{min} showed the worst performance, whereas BO-N_{max} showed superior average performance, followed by BO-N_{DOE}.

These results suggest that, for this scenario, a larger Sobol-generated space (BO-N_{max}) may generally enhance the effectiveness of BO optimisation, as opposed to a factorially structured initial design.

Lastly, for this case study, the optimisation times on a 13th Gen Intel i7-1365U processor were recorded once again. For DOE, the optimisation took approximately between [0.22, 0.59] s, whereas BO required between [18.79, 119.96] s.

CONCLUSIONS

The practical comparison study yielded insights across multiple scenario conditions, ranging from a complex, unknown experimental setting with categorical. continuous and mixture variables to a simpler, well-known system, just composed of continuous variables.

Across these studies, the findings consistently

indicated that an adapted sequential DOE approach consistently outperforms classical BO in several aspects: achieving closer convergence to the target response and the optimal point, decreasing variability in locating the optimum, and requiring a lower experimental budget.

It is important to note that these results do not suggest that BO lacks the potential to excel in other case scenarios. As highlighted by the No Free Lunch theorem, no single algorithm can be expected to outperform all others across every possible problem [10].

As shown here, BO remains promising when carefully tailored, as demonstrated by the size and design space generation discussed in the Case Study 2. Additionally, several aspects remain to be explored in future studies (e.g., the implementation of stopping criteria in BO).

In conclusion, the insights from this analysis indicate that, despite the well-established theoretical foundations of both DOE and BO, focused practical studies remain critical to fully understand their real-world performance and applicability, rather than treating them as competing approaches. Future research will further expand this analysis to tackle the current scarcity in practical comparisons, such as exploring their robustness under model sparsity and evaluating common benchmark optimisation cases (e.g., Hartman6D and Hartman3D).

DIGITAL SUPPLEMENTARY MATERIAL

To access the Supplementary material mentioned here, please refer to LAPSE ID: LAPSE:2026.0009

ACKNOWLEDGEMENTS

This work was supported by the CInTech project - Technological Hub for Innovation, Translation and Indus trialisation of Complex Injectable Drugs -, under refer ence no. C644865576-00000005, co-financed by Componente C5 - Capitalização e Inovação Empresarial integrada na Dimensão Resiliência do Plano de Recuperação e Resiliência (PRR), through the NextGenerationEU fund. Authors acknowledge support from CERES – Chemical Engineering and Renewable Resources for Sustainability Research Center, funded by FCT – Fundação para a Ciência e Tecnologia (UID/00102/2025) and PRR – Recovery and Resilience Program, of the Portuguese Republic (UID/PRR/00102/2025).

AUTHOR IDENTIFIERS

Author ORCIDs:
Batista DV: 0009-0006-0856-2739
Reis MS: 0000-0002-4997-8865

REFERENCES

1. Wu CFJ, Hamada MS, Experiments - Planning, Analysis, and Optimization. John Wiley & Sons (2009).
2. Santos CP, Rato TJ, Reis MS. Design of experiments: a comparison study from the non?expert user's perspective. Journal of Chemometrics 33: (2018). https://doi.org/10.1002/cem.3087
3. Montgomery DC. Design and Analysis of Experiments.John Wiley & Sons(2013)
4. Eggensperger K, Feurer M, Hutter F, Bergsta J, et al. Towards an empirical foundation for assessing bayesian optimization of hyperparameters. NIPS workshop on Bayesian Optimization in Theory and Practice (2013)
5. Snoek J, Larochelle H, Adams RP.Practical Bayesian Optimization of Machine Learning Algorithms, Adv Neural Inf Process Syst 25(2012)
6. Brochu E., Cora VM, Freitas N. A Tutorial on Bayesian Optimisation of Expensive Cost Functions, with Application to Active User Modelling and Hierarchical ReinforcementLearning.arXiv preprint arXiv:1012.2599 (2010)
7. Greenhill S, Rana S, Gupta S, Vellanki P, Venkatesh S. Bayesian optimization for adaptive experimental design: a review. IEEE Access 8:13937-13948 (2020). https://doi.org/10.1109/access.2020.2966228
8. Karl AT, Essex S, Wisnowski J, Rushing H. A workflow for lipid nanoparticle (LNP) formulation optimization using designed mixture-process experiments and self-validated ensemble models (SVEM). JoVE : (2023). https://doi.org/10.3791/65200
9. Batista DV, Reis MS. Balancing modelling complexity and experimental effort for conducting qbd on lipid nanoparticles (lnps) systems. Systems and Control Transactions 4:2548-2553 (2025). https://doi.org/10.69997/sct.163183
10. Wolpert DH, Macready WG. No free lunch theorems for optimization. IEEE Trans. Evol. Computat. 1:67-82 (1997). https://doi.org/10.1109/4235.585893

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Advancing Industrial Fermentation across scales: Model Development, Cost Analysis, and Predictive Control

Marc Lemperle[a]*, Pedram Ramin[a], Julian Kager[a], Benny Cassells[b], Stuart Stocks[b], Krist V. Gernaey[a]

[a] Process and Systems Engineering Center (PROSYS), department of Chemical and Biochemical Engineering, Technical University of Denmark, Building 228A, 28000 Kgs. Lyngby
[b] Novonesis, Fermentation Pilot Plant, Krogshoejvej 36, 2880 Bagsværd, Denmark
* Corresponding Author: marlem@dtu.dk

ABSTRACT

The bioprocess industry is actively exploring technologies associated with the fourth industrial revolution, with modeling offering considerable potential for process optimization. Nevertheless, model adoption in industry remains limited. This is partly because model development continues to depend heavily on offline sampling, and because relatively few industrial applications convincingly demonstrate their practical value. This study therefore first examines the benefits of online rheology and online biomass measurements for model development and demonstrates, among other aspects, that online biomass significantly improves model fidelity. The second part examines how electricity prices affect process conditions, a key factor in production, and finds that, contrary to common practice, maximizing all operating parameters is not the most cost-effective strategy. Finally, an in-silico framework for model predictive control, applied to a reactor end-fill scenario, demonstrates that oxygen-controlled processes can be dynamically optimized, highlighting the strong potential of bioprocess models for industrial usage.

Keywords: Bioprocess Modelling, Cost Analysis, Model Predictive Control

INTRODUCTION

Fermentation processes are well-established biochemical operations that utilize microorganisms to produce specific compounds, ranging from monoclonal antibodies to chemicals and industrial enzymes [1]. Although individual processes differ substantially, all industrial fermentations rely on a complex interplay between the underlying biological system and its physical reactor environment [2], [3].

Accurate modeling of fermentation processes is considered a crucial step forward in the ongoing optimization, monitoring, control and scaling of such processes [4]. Nevertheless, the application and usage of such mathematical bioprocess models in industry remains limited, largely because the model development is highly resource-intensive and time-consuming [5].

A major contributor to this challenge is the reliance on offline measurements for rheological characterization and cell-weight–based analyses [6].

However, online measurement technologies are constantly emerging within industrial settings, raising the question of to what extent they can support model development and improve model accuracy from an industrial perspective [7].

For instance, rheological offline measurements are laborious and additionally require an estimate of the apparent viscosity based on numerous assumptions. This can be substituted by online measurements of the dynamic viscosity using a vibrational viscometer operating in the upper Newtonian region, in which the fluid behaves shear-independently. Dynamic viscosity is referred to in the following as online viscosity [8]. Additionally, online biomass sensors that quantify viable cells based on permittivity can replace the cumbersome and, under current industrial conditions, error-prone measurements of cell dry weight (CDW) [9], [10], [11].

This work first aims to investigate the extent to which the integration of those two online measurement technologies improves model development and enhances model accuracy for industrial fungal fermentation processes at pilot and production scale.

This is done by comparing three modeling approaches, leading to three different models, developed

using pilot plant data and validated through pilot and large-scale simulations: (i) a mechanistic model based on offline CDW and offline rheology measurements, (ii) a hybrid model combining offline CDW and online viscosity, and (iii) a hybrid model relying on online viable cell concentration and online viscosity. The complexity of the model is addressed in more detail in the "Methodology" section.

Alongside the effort required for development, the usage of models in industry may also be slowed down by a lack of practical examples that clearly demonstrate their benefits [12], [13], [14]. Therefore, the second part of this work presents two model applications tailored to an established dissolved oxygen (DO)-controlled enzyme-production process. First, a cost analysis under varying electricity prices is conducted to provide economic insights into fermentation operations. Second, an in-silico framework for predictive control is implemented to dynamically optimize DO-controlled fermentations, while ensuring a defined reactor fill [15], [16].

This study aims to bridge the gap between existing modeling approaches and their application in industry. This is done by exploring the potential of online measurements for model development, on the one hand, and by presenting model applications specifically tailored to enzyme production processes, on the other hand.

METHODS AND MODELS

Three bioprocess models are used to discuss and explore the advantages of integrating novel online measurements into model development. All models are designed for the same industrial bioprocess, assuming substrate as the primary limiting factor for cell growth and death. Each model incorporates mass balances for key components, including reactor mass, biomass, substrate, product and dissolved oxygen. Due to confidential reasons, model parameters for each of the models cannot be disclosed.

A pure mechanistic model (model i) exhibits the highest dependency on offline data, requiring both CDW and offline rheology measurements. Essentially the model structure can be seen in Albeak et al., 2012 [6], with the difference being that yields have been linked to the growth rate, rather than being constant as well as the integration of a Monod growth kinetic.

Cell death is not modelled explicitly in model i. This decision reflects the limited precision of CDW measurements, which prevents robust estimation of death kinetics. The linear yields correlation as well as μ_{max} has been determined experimentally, whereas the remaining parameters were fitted using data from six fermentations covering the relevant operational space. Additionally, enzyme production is modelled based on a constant rate and includes a product inhibition term. The applied k_La correlation is presented in equation 1.

$$k_La = C * \left(\frac{P}{V}\right)^a * v_s^b * \mu_{app}^c \qquad 1$$

The apparent viscosity is described using the Henzler and Otto model in which the power-law parameters are fitted to rheological measurements, as presented in Albeak et al., 2012 [6].

The only distinction between the mechanistic model (model i) and the second model, model ii, is that the apparent viscosity is replaced by online viscosity. In this approach, the k_La correlation, shown in equation 2, includes online viscosity that is predicted by a supervised LGBM machine learning (ML) model, resulting in a sequential hybrid model. The ML model, trained with 27 historical fermentation processes from pilot plant and production, utilizing different fungal strains, predicts the online viscosity based on mechanistic modelled input features, as described in more detail in [17].

$$k_La = C * \left(\frac{P}{V}\right)^a * v_s^b * \mu_{online}^c \qquad 2$$

Therefore, model ii advances model i by increasing the use of online data (i.e. viscosity), while model iii builds upon model ii by integrating additional online measurements (i.e. capacitance), consequently reducing the need for offline sampling. The availability of online viable cell concentration measurements enables model iii to explicitly incorporate cell-growth and cell-death kinetics, both formulated as functions of substrate concentration and dissolved oxygen levels [18].

In addition, model iii integrates an extended ML–based viscosity model from model ii that includes capacitance as an input feature. However, this dataset consists of only two pilot-scale fermentation processes conducted with the same production strain. The restricted size of this dataset reflects the current limited availability of fermentation batches, equipped with online capacitance probes.

Consequently, in contrast to model i and model ii, which model biomass growth with a Monod model, model iii, predicts the viable cell concentration with a Contois-type growth and O'Brien-type death kinetic [19], [20]. As this extension increases the number of kinetic parameters in model iii, a differential evolution algorithm is employed to ensure sufficient exploration of the enlarged parameter space. The optimization applied the '`best1bin`' strategy with a population size of 6 and a maximum of 30 generations as well as a relaxed tolerance of 0.01 to enable early termination when improvement falls below 1% [21].

Due to limitations in the availability of offline and online biomass and rheology measurements, it was not possible to obtain a sufficiently comprehensive data set to calibrate all models on the same dataset, resulting in the use of the two data sets. Model i and ii are calibrated

on six 550L fermentation processes designed to explore the operational space of the system. In contrast, model iii is based on two 550L fermentations with different designs, one targeting three specific growth rates, while the other examines cell behavior under varying k_La conditions. The presented results are validated on additional fermentation processes at pilot and large scale that were not used for calibrating the models.

Cost optimization

In-silico cost optimization of a fermentation process using a comprehensive bioprocess model represents a compelling application of modeling for industrial purposes [22].

The first step of a cost optimization involves linking the relevant cost factors, as listed in Table 1, to the simulation. For the purposes of this study, and due to the limited cost information available in the literature, we assumed a fixed baseline operating cost, denoted as fixed overhead costs. These represent expenses that occur regardless of how the fermentation is run. The values are based on internal company data and cannot be disclosed.

Other costs are associated with either the reactor start or end filling, the inherent process operating conditions as well as the fermentation time.

The energy prices in the analysis are based on historical data from the Danish energy market and reflect the minimum, maximum and average price over the last 10 years [23]. The other costs are based on literature values selected to best approximate the present case; however, the availability of detailed information in the literature is rather limited [16], [24], [25].

Table 1: Cost factors of a fermentation process

Fixed Costs	fixed overhead costs	confidential
Variable Costs	Electricity	0.8, 1.8, 2.8 [23] [DKK/kwh]
	Carbon source	4.23 [16] [DKK/kg]
	Ammonia	1.94 [16] [DKK/kg]
	Media	0.26 [24] [DKK/kg]
	Recovery Cost	4.49 [25] [DKK/kg]

In the second step, a realistic manufacturing scenario is embedded within the optimization problem. The problem is formulated as a nonlinear, simulation-based, single-objective constrained optimization problem. The decision variables $\mathbf{x} = [\mathrm{p}, \mathrm{a}, \mathrm{g}, \mathrm{t_f}]^{\mathrm{T}}$, represents pressure, aeration, agitation, and fermentation time. The objective is to minimize the total fermentation cost $\mathrm{C}(\mathbf{x})$, subject to meet a required enzyme production $\mathrm{P}(\mathbf{x}) \geq \mathrm{P_{target}}$ and respecting operational bounds $\mathrm{x_i^{min}} \leq \mathrm{x_i} \leq \mathrm{x_i^{max}}$, represented in the following formulation.

$$\min_{x} C(x),$$

$$s.t.\ P(\boldsymbol{x}) \geq P_{target},$$

$$x_i^{min} \leq x_i \leq x_i^{max}, i \in \{p, a, g, t_f\}$$

A nested root-finding procedure adjusts the initial reactor fill to ensure identical final fill levels across all simulations. Because the model exhibits non-linear, non-convex behavior without analytical gradients, the optimization is solved using a Genetic Algorithm with a population size of 50, up to 25 generations, using tournament selection, two-point crossover (probability 0.5), and Gaussian mutation (probability 0.5). The optimization process is initiated by generating 50 initial points within the design space using Latin Hypercube Sampling (LHS) [26].

Advanced control for the reactor end fill

Many industrial enzyme production processes have already been extensively optimized, including by adjusting the feed rates to maintain a stable oxygen balance. However, as this normally leads to varying reactor end fills, a subsequent objective might be to reliably hit a defined end fill. Reactor end fills are linked to the total product produced and beyond that enhances manufacturing planning reliability by reducing process deviations. Model predictive control (MPC) is widely regarded as a powerful approach to enable such advanced control [6].

To reduce batch-to-batch variability and achieve this in DO-controlled fungal fermentation, a modular in-silico workflow that combines biased initial planning, real-time parameter adjustment, and constraint-based input optimization, is presented. First, a realistic model-process mismatch is created by distorting a calibrated "source" process model: the dynamic substrate and oxygen yields are influenced by a constant factor that is randomly determined in the range of ±30% of the yield's value. Using this deliberately distorted model, a Brent root-finding optimization is employed to determine the initial reactor fill that would result in the fermentation reaching the predefined final reactor fill. That start fill is applied to the unperturbed source (or truth) model to generate synthetic batch data, representing the realized process.

During the run, the algorithm corrects the distorted model at specified control points by finding correction factors for the dynamic yields, using the available online data of viable cell concentration, oxygen uptake rate (OUR), and feed rate by minimizing a normalized mean-squared error over the time to the checkpoint with Nelder-Mead under conservative tolerances. Normalization guards against misalignment and scale differences and regularization with constraints keeps parameters physically plausible. The corrected model is then rolled forward to forecast the rest of the batch. If the predicted

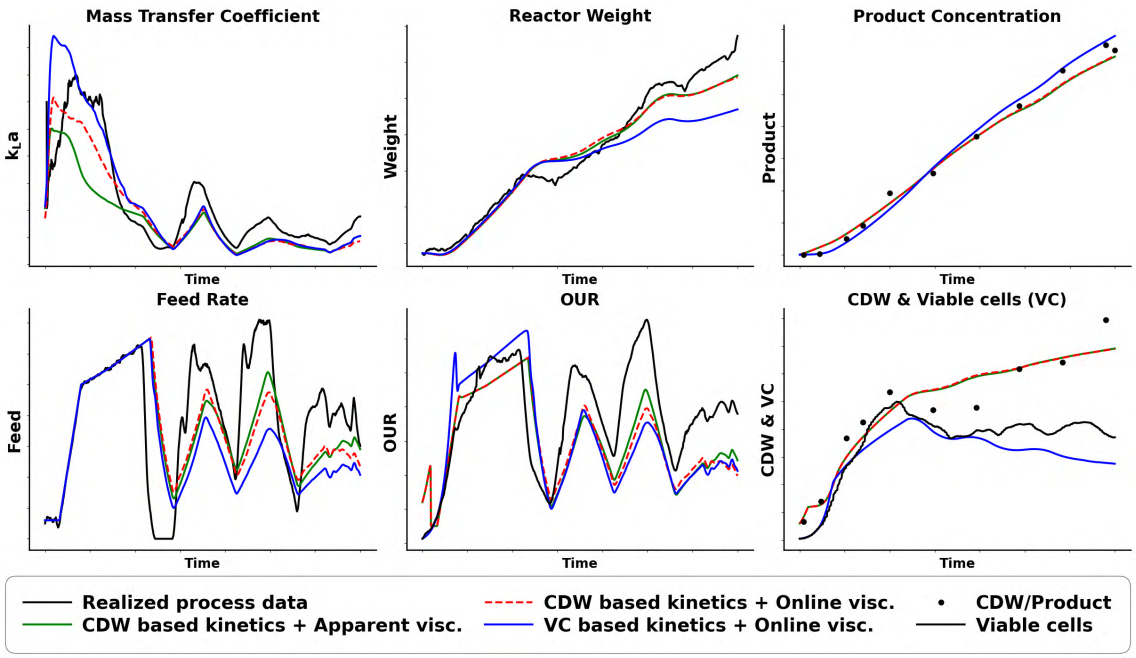

Figure 1. Pilot-scale predictions for the three models.

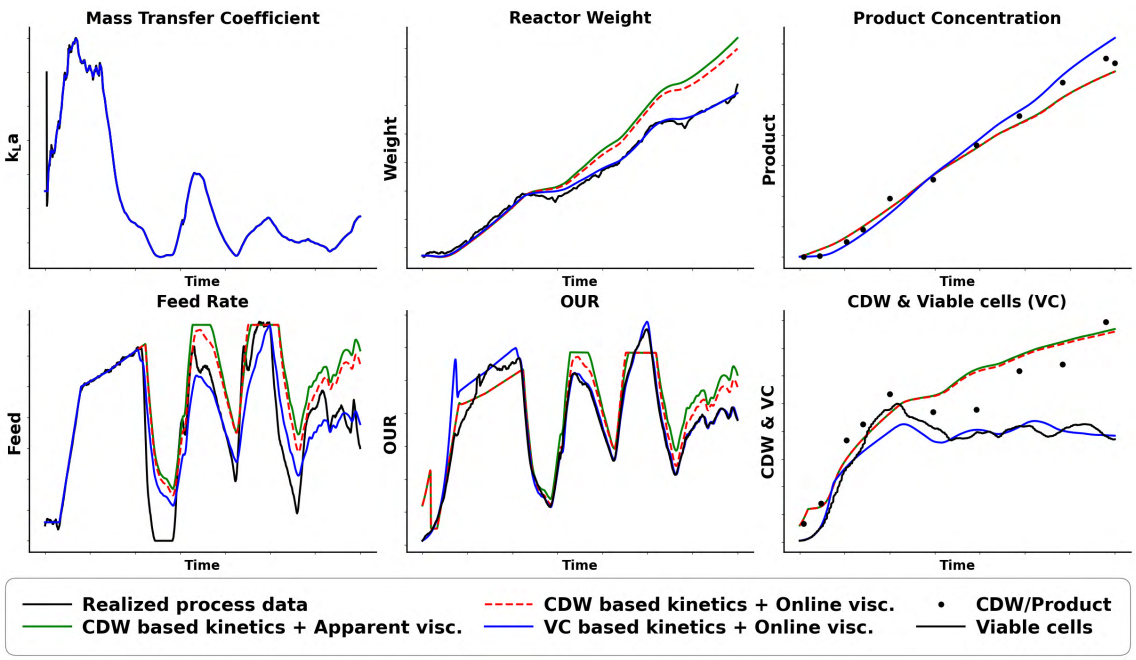

Figure 2. Pilot-scale predictions for the three models with realized k_La.

end fill based on the new model deviates from the target beyond a threshold of e.g. ±1%, a bounded L-BFGS-B adjusts the agitator speed and aeration. This approach is possible, as with a DO-controlled process, there is a direct correlation between the OTR and the final reactor fill of a process.

$$Agitation\ and\ Aeration \uparrow \rightarrow OTR \uparrow \rightarrow DO \uparrow \rightarrow feed\ rate \uparrow \rightarrow end\ fill \uparrow$$

By adjusting the agitator speed and aeration under DO-controlled settings, a DO concentration close to the set point as well as a final fill that corresponds to the desired final fill of the process is ensured. This can be achieved by optimizing both the agitation and the

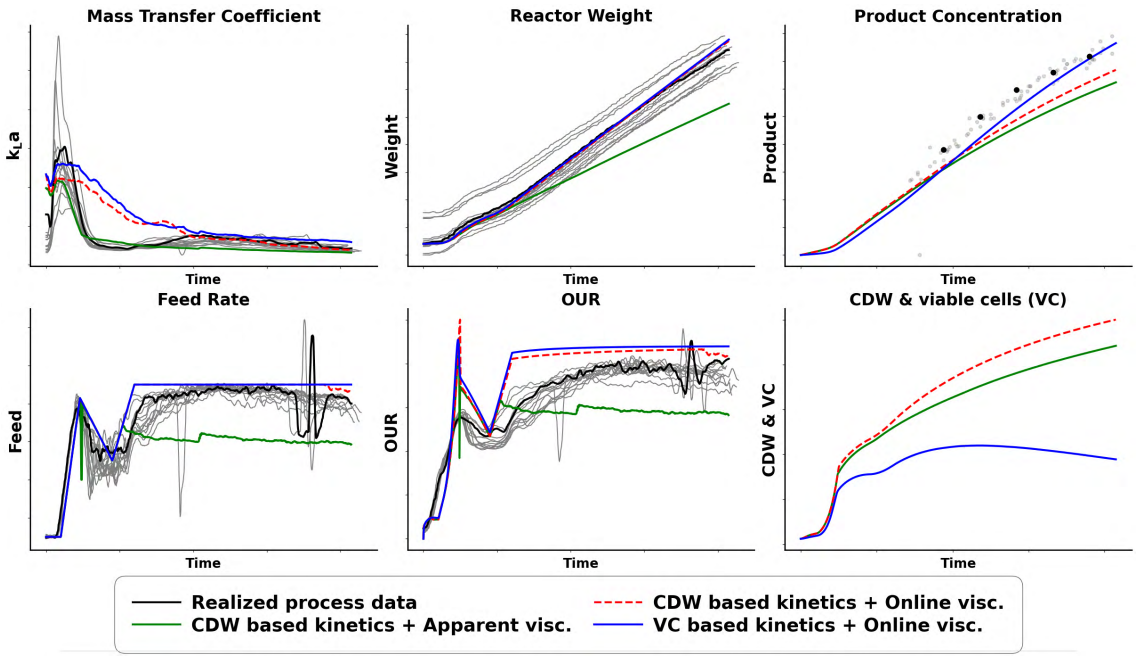

Figure 3. Large-scale predictions for the three models.

aeration based on the squared error between the final reactor fill prediction of the corrected model and the target fill using an L-BFGS-B optimization algorithm with moderate tolerances.

After executing this framework at each checkpoint, all actions are logged and saved, as they define the process until the next checkpoint. This fit-predict-control loop repeats to batch completion, maintaining robust DO control while steering the process to the target end fill under initial model inaccuracy. To reduce the computational cost, the framework is only activated at those time points, in which the modelled predictions deviate from the actual realized process data.

RESULTS

Model simulation for pilot and large scale

Figure 1 presents the closed-loop simulation results for models i, ii, and iii. The empirical k_La models, each depending on their respective biological and physical sub-models, show reasonable agreement with the measured data throughout most of the fermentation process. Deviations in the predicted k_La values occur primarily during the initial hours and can largely be attributed to viscosity effects. These early-phase differences arise either from the use of different viscosity models or from discrepancies in the biological sub-models implicitly embedded in the k_La estimation.

Because the process is intentionally operated at its maximum feed rate during the initial hours, the early differences in k_La do not translate into additional feed addition. As the predicted k_La converges for the remainder of the fermentation, and because model i and ii share similar underlying biological sub-models, the resulting state predictions overlap for most variables during the closed-loop pilot-scale validation run.

Model iii, which was developed using only two fermentation datasets and viable cell measurements, underestimates the reactor fill compared to the observed data. Further investigation is required to understand the root causes of the discrepancies among the models. One possible diagnostic approach is to rerun the simulations using the realized k_La profile, effectively decoupling the physical mass-transfer model. This allows a focused assessment of the biological sub-models' accuracy, as illustrated in Figure 2.

Figure 2 shows that the biological sub-models in models i and ii systematically result in an overprediction of the reactor mass when the realized $\mathrm{k_La}$ profile is imposed. In contrast the biological sub-model in model iii lead to an accurate prediction of the reactor mass. This indicates that, for models i and ii, inaccuracies in the physical $\mathrm{k_La}$ model and the biological growth model partially cancel each other out, leading to an apparently improved mass prediction in the fully coupled simulation. In contrast, the biological sub-model of model iii seems to be more accurate but does not benefit from this compensatory effect. Therefore, model iii suffers more directly from the limitations of the $\mathrm{k_La}$ model, when it's simulated.

Figure 3 illustrates the performance of the same models at production scale, compared against multiple manufacturing datasets shown in grey, with the black

trajectory representing the specific batch, which operational settings are used as input for the simulation.

Unlike the pilot-scale scenario, the k_La predictions at large scale do not align well with the observed data. This mismatch indicates the limitations of the empirical k_La, developed at pilot scale and employed at production scale. Because of these challenges, there might be the need for either re-calibration at production scale or a theoretically grounded k_La model that can scale effectively from pilot to production.

Since the hybrid models generally overpredict k_La, the process operates at its maximum feed ramp due to the DO-control, whereas the purely mechanistic model underpredicts k_La, resulting in a reduced feed rate and deviations in both reactor fill and product predictions. No biomass measurements are available due to limited availability at production scale.

Overall, the results demonstrate that capacitance measurements, enables the development of a more accurate and realistic biological model at pilot scale. In contrast, integrating online viscosity data does not lead to a noticeable improvement in prediction accuracy for this case, but it also does not negatively affect the physical k_La model. Combined with the advantage that both, capacitance and online viscosity reduce the need for frequent offline sampling, these findings highlight the value utilizing online data for model development.

Cost optimization

Cost optimization was carried out using model ii at a pilot scale, as it is known from previous work that this model provides so far, the most robust and broadly applicable predictions [17]. This is essential for such a cost optimization analysis in which a wide range of operational conditions are evaluated. The primary objective is to minimize the overall campaign cost under specific constraints as elaborated in the method section.

Due to confidentiality constraints, the results in these sections are presented in relative terms. Accordingly, the individual cost factors of a fermentation process in Table 2 are shown as percentages of the total fermentation cost. This allows identification of each factor's contribution under different electricity price scenarios. The analysis of Table 2 reveals that as electricity prices increase, the share of electricity costs rise substantially, from 12.4% for 0.8 DKK/kWh to 24.9% for 2.8 DKK/kWh, which is expected. Consequently, the relative contribution of other cost factors decreases. This trend aligns well with the general understanding of cost distribution in fermentation processes: at lower electricity prices, recovery, substrate and fixed costs dominate, whereas at higher prices, electricity contributes more than substrate alongside recovery and fixed costs.

Table 2: Individual fermentation cost factors at each electricity price, scaled by the corresponding total fermentation cost.

Electricity [DKK/kWh]	**0.8**	**1.8**	**2.8**
Starting costs [%]	2.2	2.2	2
Electricity [%]	12.4	18	24.9
Substrate [%]	17.3	13.9	12.7
Ammonia [%]	1.9	1.5	1.4
Recovery [%]	56.6	55	50.4
Fixed costs [%]	9.6	9.3	8.6

Table 3 lists the process decision variables x across the three electricity price scenarios. The decision variables x in Table 3 are normalized by their respective upper bound $x_i^{\max}$, providing insights into how intensively each decision variable is utilized. Interestingly, fermentation time remains nearly constant across all cases. Pressure is found at approximately 93% of its maximum value for the low electricity scenario and increases to almost 99% under higher electricity prices, which is consistent with its low-cost impact on the process. Higher headspace pressure reduces evaporation by lowering the vapor liquid driving force, thereby influencing the water balance and indirectly affecting the mass-based calculations [15].

In contrast, energy-intensive parameters such as agitation and aeration are significantly reduced. For example, aeration starts at about 100% for 0.8 DKK/kWh but drops to around 60% for 1.8 and 2.8 DKK/kWh. Agitation already starts with 60% utilization of $x_i^{\max}$ for the low energy cost scenario and drops to 50% for the higher energy costs, which also aligns with the overall expectations, as those decision variables are directly linked to higher electricity costs.

Table 3: Optimized decision variables x, scaled by their corresponding upper limit $x_i^{\max}$.

Electricity [DKK/kWh]	**0.8**	**1.8**	**2.8**
Aeration [%]	100	61.6	61.4
Agitation [%]	60.7	52.8	51.8
Pressure [%]	92.7	95.5	98.8
Fermentation time [%]	81.8	81.8	82.1

Finally, the results in Table 4 and Figure 4 are scaled relative to the cost associated with a reasonable current industrial standard fermentation, defined as operating all decision variables x at their respective upper limits $x_i^{\max}$, as implied in [2], [3]. Table 4 summarizes the scenario-specific total campaign costs. The results indicate that rising electricity prices amplify the cost advantage of the optimized campaign relative to a strategy in which all operational variables are maintained at their maximum setpoints. E.g. at a price of 0.8 DKK/kWh, the optimized campaign costs only 80.5% of the cost of the campaign that would result from a typical current standard strategy.

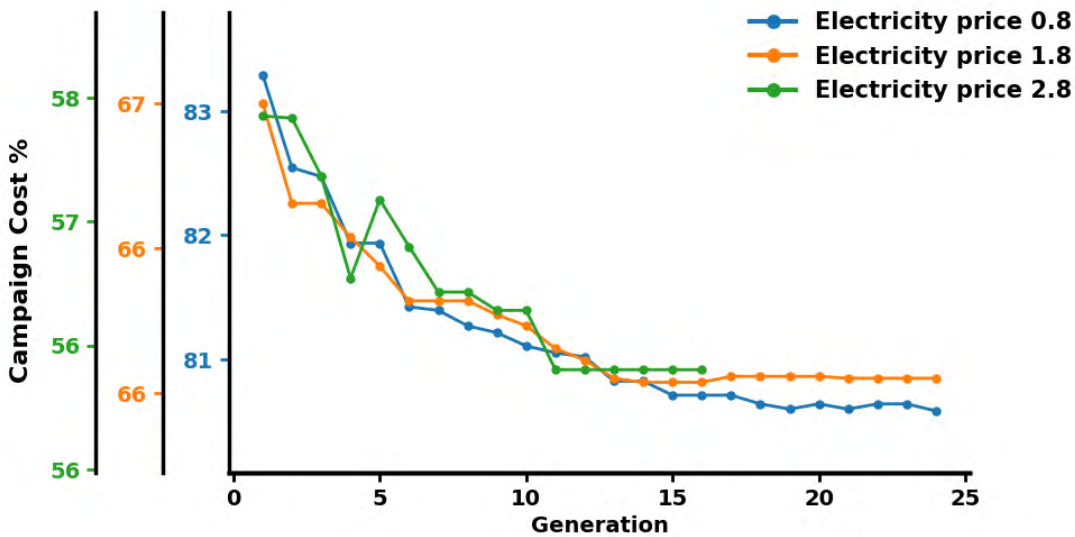

Figure 4. Campaign costs in percentage, scaled in the same manner as in Table 4.

Table 4: Total campaign costs for each electricity cost scenario, scaled by cost of a state-of-the-art operational strategy ($x_i = x_i^{\max}$)

Electricity [DKK/kWh]	0.8	1.8	2.8
Total campaign costs [%]	80.5	65.6	56.3

Figure 4 illustrates the campaign cost progression across generations of the differential evolution algorithm, scaled in the same manner as the values in Table 4. The visualization reveals that for each electricity price scenario, the optimizer convergences and successfully identifies optimal process conditions within the given constraints. It also reveals that the optimization finds a very good first condition in the first generation of the optimization, which is likely caused by the broad initial LHS screening.

The study provides an example of how a bioprocess model can be used to evaluate the cost implications for different electricity price scenarios. The results reveal that operating all parameters at 100% ($x_i = x_i^{\max}$) is not the most cost-efficient, which becomes increasingly evident as energy prices rise. Also, the results indicate that increasing energy prices primarily affect energy-intensive operations such as agitation and aeration, which should be scaled down to reduce costs, while pressure can remain near its maximum due to its negligible cost impact [16]. Interestingly, the comparison between 1.8 and 2.8 DKK/kWh shows minimal differences in the process conditions, suggesting a potential plateau in cost sensitivity beyond a certain threshold.

The optimization, performed at pilot scale, assumes constant values for agitation, aeration, and pressure throughout the fermentation process, which may not reflect the most efficient strategy in practice, but should be efficient enough for this purpose of energy cost evaluation. Furthermore, the results are considered to be highly dependent on the underlying model, scale and cost assumptions, limiting their generalizability to other processes. Future work could address time-dependent operational adjustments and include a broader range of electricity price scenarios to better capture system dynamics and identify potential tipping points for operational changes. Similar analysis at large scale might also be very interesting but first requires a similar broadly accurate process model . Finally, many more industrial relevant business cases can be approached using the same framework and subsequently customized for the specific application.

Advanced control for the reactor end fill

Advanced control was implemented using the viable cell model (model iii), which offers the advantage of incorporating additional online measurements into the control strategy. To imitate a realistic scenario for MPC of an industrial fermentation process, the first step is to define a target reactor fill and perform an optimization to determine the required initial fill, which has been done with the same procedure as for the cost analysis. This analysis was executed using a distorted process model to create a discrepancy between the model and the process, as might occur in a real case.

The mimicked fermentation process, implemented by using the correct model and initiated with the estimated start fill, is represented by the black trajectories in

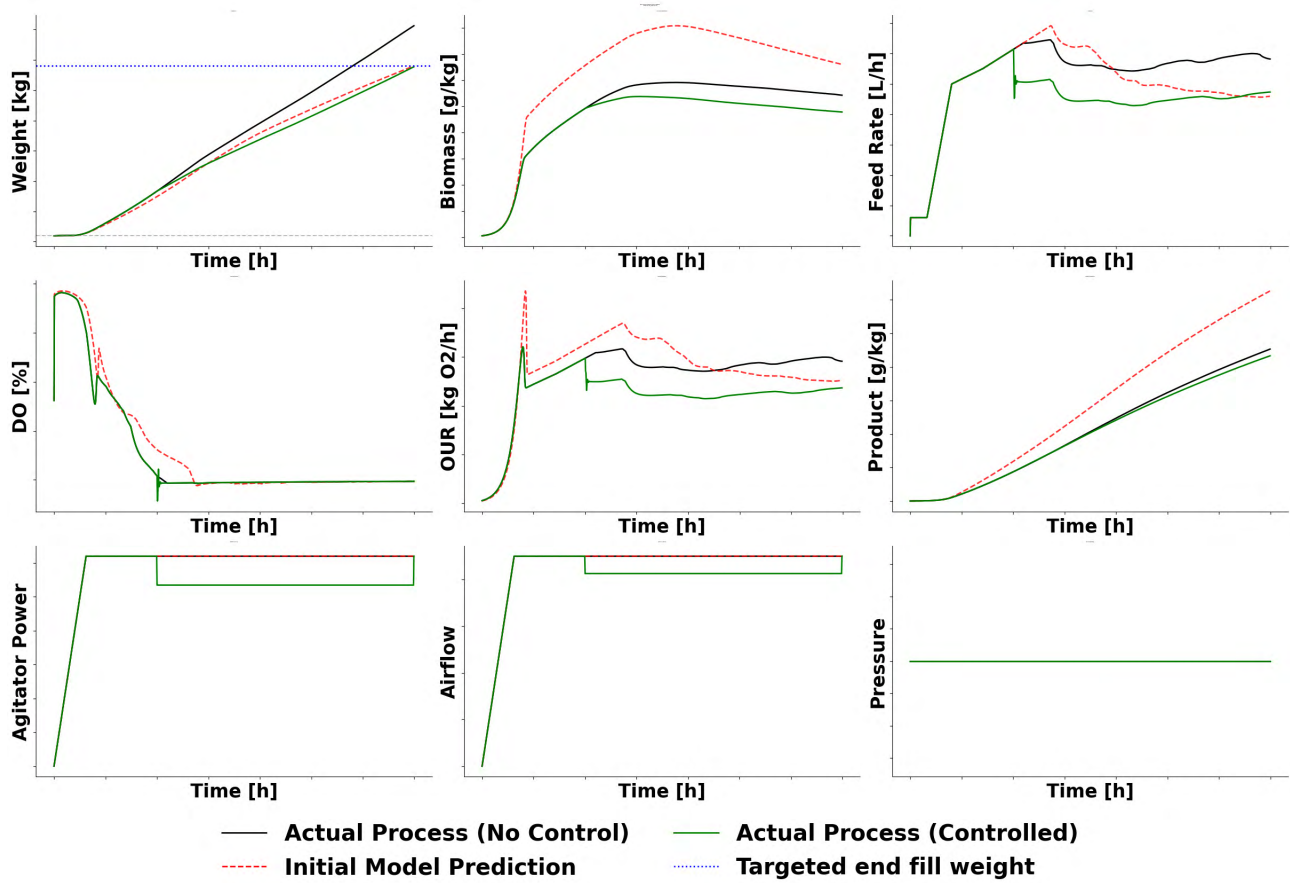

Figure 5. Full Fermentation process for the initial model and the synthetic process with and without advanced control.

Figure 5. As expected, without a corrective control mechanism, this leads to a reactor end fill that deviates from the defined end fill, underscoring the need for a corrective control strategy.

In contrast, the controlled fermentation process successfully reaches the target end fill. The optimizer adjusts agitation and aeration by lowering their setpoints, consequently reducing the DO-controlled feed rate and thus the final fill. The agitation is tuned a bit more than the aeration, which can be explained by the greater influence on the oxygen mass transfer. In addition, less aeration not only reduces the feed rate, but also evaporation, which lowers the total effect on the final weight.

Those results demonstrate the effectiveness of the control strategy described in the *Methods* section. Enabling sequential adjustments in which agitation and aeration are modified to achieve the desired final fill, while keeping DO close to its setpoint.

Under control, the final product concentration is slightly lower compared to the uncontrolled scenario and the initial model, reflecting the reduced feed addition and corresponding decrease in biomass concentration. Nevertheless, the reactor does not overfill, which is set as the main objective. This scenario illustrates that, even with an inaccurate initial model, the proposed framework can iteratively correct the model and control the process by fine-tuning agitation and aeration to ensure the process meets the end fill target without violating DO constraints.

Future work should focus on implementing and validating this approach in a real-world setting, which requires careful consideration. System boundary conditions, easily overlooked in simulations, could lead to operational failures if not properly addressed. These constraints must be incorporated into the simulation and thoroughly evaluated before deploying such a closed-loop model predictive control strategy in a pilot plant environment.

Additionally, attention must be given to the time intervals for evaluating control actions. These intervals should align with process dynamics and ideally exceed the time required by the computational framework to update the model, and adjust agitation, aeration and feed rate settings before the next evaluation point.

CONCLUSION

This study first discusses the benefits of utilizing online rheology and online biomass measurements for fermentation models, emphasizing that especially the integration of capacitance data seems to improve the biological sub-model accuracy. Combined with the reduced labor associated with offline measurements, this supports the usage of online data to tackle the resource-intensive and time-consuming model development.

Secondly, the study demonstrates the applicability of bioprocess models tailored for industrial enzyme production, showing that electricity prices substantially influence the optimal operating strategy. The results clearly indicate that maximizing all decision variables, such as agitation, aeration, or fermentation time, is not cost-efficient under the given optimization problem and varying electricity prices. Finally, the in-silico MPC case illustrates that an initial model–process mismatch can be corrected in real time, enabling dynamic adjustment of agitation and aeration to ensure that the reactor achieves its target end-fill while maintaining DO control throughout the process.

ACKNOWLEDGEMENTS

Special thanks to the entire team at the Novonesis fermentation pilot plant for their valuable involvement in this project

AUTHOR IDENTIFIERS

Author ORCIDs:
ML: 0009-0006-6389-6563

REFERENCES

1. P. Stanbury, A. Whitaker, and S. Hall, *Principles of fermentation technology*. Elsevier, 2013.
2. Villadsen J, Nielsen J, Lidén G. Bioreaction engineering principles. Springer US (2011). https://doi.org/10.1007/978-1-4419-9688-6
3. Doran PM. Reactor engineering. Bioprocess Engineering Principles :333-391 (1995). https://doi.org/10.1016/b978-012220855-3/50013-4
4. Rydal T, Frandsen J, Nadal?Rey G, Albæk MO, Ramin P. Bringing a scalable adaptive hybrid modeling framework closer to industrial use: application on a multiscale fungal fermentation. Biotech & Bioengineering 121:1609-1625 (2024). https://doi.org/10.1002/bit.28670
5. Albino M, Gargalo CL, Nadal-Rey G, Albæk MO, Krühne U, Gernaey KV. Hybrid modeling for on-line fermentation optimization and scale-up: a review. Processes 12:1635 (2024). https://doi.org/10.3390/pr12081635
6. Albaek MO, Gernaey KV, Hansen MS, Stocks SM. Modeling enzyme production with aspergillus oryzae in pilot scale vessels with different agitation, aeration, and agitator types. Biotech & Bioengineering 108:1828-1840 (2011). https://doi.org/10.1002/bit.23121
7. Kager J, Herwig C, Stelzer IV. State estimation for a penicillin fed-batch process combining particle filtering methods with online and time delayed offline measurements. Chemical Engineering Science 177:234-244 (2018). https://doi.org/10.1016/j.ces.2017.11.049
8. Pu Y, Chaudhry S, Parikh M, Berry J. Application of in-line viscometer for in-process monitoring of microcrystalline cellulose-carboxymethylcellulose hydrogel formation during batch manufacturing. Drug Development and Industrial Pharmacy 41:28-34 (2014). https://doi.org/10.3109/03639045.2013.845837
9. Bergin A, Carvell J, Butler M. Applications of bio-capacitance to cell culture manufacturing. Biotechnology Advances 61:108048 (2022). https://doi.org/10.1016/j.biotechadv.2022.108048
10. Fehrenbach R, Comberbach M, Pêtre JO. On-line biomass monitoring by capacitance measurement. Journal of Biotechnology 23:303-314 (1992). https://doi.org/10.1016/0168-1656(92)90077-m
11. Konakovsky V, Yagtu A, Clemens C, Müller M, Berger M, Schlatter S, Herwig C. Universal capacitance model for real-time biomass in cell culture. Sensors 15:22128-22150 (2015). https://doi.org/10.3390/s150922128
12. Herwig C, Pörtner R, Möller J. Digital twins. Springer International Publishing (2021). https://doi.org/10.1007/978-3-030-71660-8
13. Udugama IA, Lopez PC, Gargalo CL, Li X, Bayer C, Gernaey KV. Digital twin in biomanufacturing: challenges and opportunities towards its implementation. Syst Microbiol and Biomanuf 1:257-274 (2021). https://doi.org/10.1007/s43393-021-00024-0
14. Mears L, Stocks SM, Albaek MO, Sin G, Gernaey KV. Mechanistic fermentation models for process design, monitoring, and control. Trends in Biotechnology 35:914-924 (2017). https://doi.org/10.1016/j.tibtech.2017.07.002
15. Mears L, Stocks SM, Albaek MO, Cassells B, Sin G, Gernaey KV. A novel model?based control strategy for aerobic filamentous fungal fed?batch fermentation processes. Biotech & Bioengineering 114:1459-1468 (2017). https://doi.org/10.1002/bit.26274
16. Ferreira RG, Azzoni AR, Freitas S. Techno-economic analysis of the industrial production of a low-cost enzyme using E. coli: the case of

recombinant ?-glucosidase. Biotechnol Biofuels 11: (2018). https://doi.org/10.1186/s13068-018-1077-0
17. M. Lemperle, P. Ramin, J. Kager, B. Cassells, S. Stocks, and K. V. Gernaey, “Hybrid bioprocess model towards the development of a digital twin for an industrial fermentation process, ” *ESCAPE 35: 35th European Symposium on Computer Aided Process Engineering 2025*, pp. 199–200, 2025.
18. Esener AA, Veerman T, Roels JA, Kossen NWF. Modeling of bacterial growth; formulation and evaluation of a structured model. Biotech & Bioengineering 24:1749-1764 (2004). https://doi.org/10.1002/bit.260240803
19. CONTOIS DE. Kinetics of bacterial growth: relationship between population density and specific growth rate of continuous cultures. Journal of General Microbiology 21:40-50 (1959). https://doi.org/10.1099/00221287-21-1-40
20. O'Brien CM, Zhang Q, Daoutidis P, Hu WS. A hybrid mechanistic-empirical model for in silico mammalian cell bioprocess simulation. Metabolic Engineering 66:31-40 (2021). https://doi.org/10.1016/j.ymben.2021.03.016
21. Abimbola M Enitan . Food processing optimization using evolutionary algorithms. Afr. J. Biotechnol. 10: (2011). https://doi.org/10.5897/ajb11.410
22. Wagner N, Bosshart A, Wahler S, Failmezger J, Panke S, Bechtold M. Model-based cost optimization of a reaction–separation integrated process for the enzymatic production of the rare sugar d-psicose at elevated temperatures. Chemical Engineering Science 137:423-435 (2015). https://doi.org/10.1016/j.ces.2015.05.058
23. Statista, “Denmark: Monthly wholesale electricity price.” Accessed: Jan. 13, 2026. [Online]. Available: https://www.statista.com/statistics/1271525/denmark-monthly-wholesale-electricity-price/
24. R. Davis *et al.*, “Process Design and Economics for the Conversion of Lignocellulosic Biomass to Hydrocarbons: Dilute-Acid and Enzymatic Deconstruction of Biomass to Sugars and Catalytic Conversion of Sugars to Hydrocarbons, ” Mar. 2015.
25. J. Wimble, R. Ashizawa, and E. W. Swartz, “Analysis of the economic viability and environmental impacts of a conceptual process model for the recovery of lactic acid from spent media in cultivated meat production, ” *Biotechnol. Prog.*, 2025, doi: 10.1002/btpr.70094.
26. R. Storn and K. Price, “Differential Evolution - A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces, ” *Journal of Global Optimization*, vol. 11, no. 4, pp. 341–359, 1997, doi: 10.1023/A:1008202821328/METRICS.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Predicting Ecotoxicity (HC50) Values Using Symbolic Regression for Transparent Life Cycle Assessment

Abdulhakeem Ahmed[a], Nitya Kasera[a], and Ana I. Torres[ab]*

[a] Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA, 15213, USA
[b] Wilton E. Scott Institute for Energy Innovation, Carnegie Mellon University, Pittsburgh, PA, 15213, USA
* Corresponding author: aitorres@cmu.edu

ABSTRACT

Accurate life cycle assessment (LCA) depends on robust characterization factors (CFs), which quantify impacts such as ecotoxicity through the integration of fate (FF), exposure (XF), and effect (EF) factors. While databases such as USEtox and Ecoinvent provide essential CFs, significant data gaps remain, particularly in ecotoxicity endpoints like hazardous concentration 50% (HC_{50}), which directly inform effect factor calculations. Existing machine learning models can predict such values, but they often lack interpretability, which limits trust and transparency in environmental modeling. To address this, a machine learning framework is applied that utilizes symbolic regression (SR) and genetic programming (GP) to predict missing HC_{50} values from physicochemical descriptors. A dataset with 14 descriptors was used to train SR models capable of generating interpretable mathematical expressions that link chemical properties to HC_{50} values. SR models were benchmarked against prominent black-box models such as random forest (RF) and neural network (NN) models. SR performance was consistent across multiple parameter configurations, while revealing recurring patterns in variable selection. These results demonstrate that symbolic regression can both predict ecotoxicity values with comparative accuracy and provide transparent functional relationships, thereby filling existing data gaps in characterization factors needed for a complete LCA.

Keywords: Life Cycle Assessment, Symbolic Regression, Machine Learning

INTRODUCTION

Life Cycle Assessment is a cornerstone of environmental modeling, enabling quantification of the potential impacts of chemical substances across their life cycles. Among the various impact categories, ecotoxicity assessment is particularly important for evaluating risks to aquatic and terrestrial ecosystems. Ecotoxicity characterization within LCA depends on effect factors such as the *Hazardous Concentration for 50% of species* (HC_{50}), which represents the concentration at which half of the tested species experience a measurable toxic effect. Reliable HC_{50} data for chemicals are essential for computing characterization factors used in assessment frameworks such as USEtox and TRACI [1, 2]. However, measured HC_{50} values are missing for many chemicals, limiting the completeness of LCA inventories and impact models.

Machine learning (ML) models are powerful tools for estimating missing ecotoxicity values based on physicochemical properties of substances. While black-box algorithms such as random forests or neural networks provide accurate predictions, their limited interpretability can hinder informed decision-making. Symbolic regression (SR) offers a promising alternative, deriving analytical expressions that directly relate input descriptors to target variables. The use of explainable data-driven models in chemical engineering is well established, with a review of past efforts, developments, and future directions presented in [3]. Most interpretable SR methods rely on genetic programming, as discussed in [4], where it was applied to mechanism discovery. Researchers have also explored optimization-based formulations; for example, in [5], a mixed-integer nonlinear programming (MINLP) model is proposed for SR to compete with traditional, stochastic genetic programming methods. The benefit of optimization is stronger optimality guarantees, unlike

genetic algorithms, which can be computationally less expensive but may yield lower solution quality [6].

Despite growing interest in data-driven methods, few studies have applied ML models to predict chemical properties in the context of life-cycle assessment or environmental impact evaluation. One such example is [7], where the authors employed ML models to estimate environmental endpoint indicators such as human health, resource utilization, and climate change, though these efforts relied exclusively on black-box models.

The present study applies symbolic regression to predict HC_{50} values using a dataset of 14 physicochemical descriptors (e.g., molecular weight, partition coefficients, solubility, and bioconcentration factor). The overall goal is to derive an interpretable model that connects the property to the descriptors. The applied SR workflow encompasses data preprocessing, model training via evolutionary optimization, and performance evaluation across multiple configurations of population size and generation count. Model interpretability is investigated through the analysis of recurrent variable patterns and expression structures.

By comparing predicted HC_{50} values against baseline models also assessed in [8] for the same dataset, we evaluate whether symbolic regression can achieve comparative predictive performance while preserving explicit analytical expressions. These equations can help bridge existing data gaps in effect factor estimation, thereby improving the robustness of LCA-based ecotoxicity characterization. The remainder of this paper details the symbolic regression framework, presents the results and interpretability analysis, and examines implications for model generalization, descriptor roles, and ecotoxicity characterization.

METHODS

Dataset

The ecotoxicity dataset used in this study originates from [8] and contains 2, 122 chemicals with experimentally determined values for HC_{50} and 14 descriptors. These descriptors include molecular weight, boiling point, and water solubility and are listed in Table 1. The descriptors were compiled into a numerical matrix and HC_{50} values were log-transformed ($\log(\mathrm{HC}_{50})$) to form the target vector.

Table 1: Description of molecular features

Variable	Descriptor	Abbreviation
X_0	Molecular Weight	MW
X_1	Atmospheric Hydroxylation Rate	AOH
X_2	Bioconcentration Factor	BCF
X_3	Biodegradation half-life	BioHL
X_4	Boiling point	BP
X_5	Henry's law constant	HL
X_6	Fish biotransformation half-life	KM
X_7	Octanol/air partition coefficient	KOA
X_8	Octanol/water partition coefficient	LogP
X_9	Melting point	MP
X_{10}	Soil adsorption coefficient	KOC
X_{11}	Vapor pressure	VP
X_{12}	Water solubility	WS
X_{13}	Mode of action	MoA

Baseline Models

Two common ML models are used to benchmark the symbolic regression approach: a Random Forest (RF) regressor and a Neural Network (NN). The RF model consists of an ensemble of decision trees constructed using bootstrap samples of the training data. Each tree is built by recursively partitioning the descriptor space using impurity-minimizing splits, while random feature selection at each node promotes decorrelation between trees. Final predictions are obtained by averaging the outputs of all trees. This structure enables RF to capture nonlinear interactions between descriptors while maintaining robustness to noise and outliers [9].

The NN model is a fully connected feed-forward network trained using backpropagation. The network maps the standardized descriptor vector to the $\log(\mathrm{HC}_{50})$ target through a sequence of nonlinear transformations defined by hidden layers and activation functions. By adjusting its weights during training, the NN model can capture complex continuous relationships. Although it requires careful regularization to avoid overfitting on medium-sized datasets [10].

Symbolic Regression Framework

Symbolic regression was implemented using a genetic programming (GP) framework developed in [11]. SR searches for explicit analytical expressions that relate the input descriptors to the target $\log(\mathrm{HC}_{50})$. The approach evolves populations of candidate mathematical expressions using operations such as mutation, crossover, and selection.

The GP search was configured to explore functional structures while preventing unnecessary algebraic complexity. The operator set was restricted to the four arithmetic operations {+, −, ×, ÷}. Candidate expressions were initialized as shallow trees with depths between 0 and 3, and the maximum allowable depth during evolution was

constrained to 10. Population sizes ranged from 4 to 12 individuals, and each evolutionary run proceeded for 100–2, 000 generations.

At every generation, each candidate expression was evaluated using the root-mean-square error (RMSE) between predicted and measured $\log(\mathrm{HC}_{50})$ values. To balance accuracy and interpretability, adjusted R^2 was incorporated into the fitness function to penalize unnecessarily complex expressions. After the evolutionary search identified a set of high-performing symbolic structures, the coefficients associated with each expression were re-optimized using a post-evolution linear regression step.

The final SR model takes the general form shown in Equation (1), where each term $f_i(X)$ represents an evolved symbolic transformation of the descriptor set $\{X_0, X_1, \ldots, X_{13}\}$, and the coefficients β_i correspond to optimized weights.

$$\hat{y} = \sum_i^k \beta_i f_i(X) \qquad (1)$$

Model development

All baseline models were trained using a three-way data partition consisting of training (70%), validation (10%), and testing subsets (20%). The training set was used for model fitting, the validation set for hyperparameter tuning, and the test set for final out-of-sample evaluation. The SR development strategy described in [11] is applied here to develop the SR models. The strategy involves training and validation on multiple data splits and GP parameters. To address the probabilistic nature of the algorithm, multiple runs are made for each data split and configuration. For the study, 100 expressions are generated per configuration and data split. Five data splits are evaluated for six parameter configurations varying in population size and number of generations, yielding a total of 3000 expressions.

Population size and number of generations control how thoroughly the genetic program explores the search space. Higher values for both increase the computational budget, and larger populations increase the set of candidate solutions maintained and evolved in each generation. By assessing various configurations, we can evaluate the effect of these parameters on predictive performance. This is done by performing statistical tests on the validation set to compare different data splits and model configurations. The Friedman and Wilcoxon Rank tests are employed to measure statistical significance between configurations and data splits. The Friedman test determines whether configurations differ significantly within each split. The Wilcoxon test evaluates statistical significance in performance differences between configurations. All tests are performed for a 99% confidence level.

The development of the RF and NN baselines is streamlined by the scikit-learn [12] and hyperopt [13] Python packages for model construction and hyperparameter optimization, respectively. Model accuracy was evaluated using the standard regression metrics (RMSE and R^2). The coefficient of determination (R^2) quantifies the proportion of variance in the measured HC_{50} values explained by the model, while the root-mean-square error (RMSE) measures the average magnitude of prediction error in the units of the log-transformed target variable. These metrics were computed for the training, validation, and test subsets to evaluate model fit and generalization behavior.

Feature Importance and Structural Analysis

The contribution of each feature to the predicted target value for all models is evaluated using SHAP (SHapley Additive exPlanations) [14]. SHAP is an approach for explaining individual predictions of an ML model by attributing a contribution value to each input feature. It builds on Shapley values from cooperative game theory, which fairly distribute total gains or costs among collaborating players. In this context, each feature acts as a player, and SHAP fairly allocates the model's prediction difference (from a baseline) across them.

In addition to understanding how individual features contribute to model predictions, we examine the structural diversity of the expressions generated via SR. This structural analysis provides complementary insights into the models and their underlying patterns. To characterize the diversity of the discovered expressions, we perform structural pattern mining by extracting abstract syntax tree (AST)-based features for each unique mathematical expression. Each expression is parsed into its AST representation, from which quantitative features, including operation counts, tree structural properties (depth, total nodes), and variable usage patterns, are extracted. Two feature extraction modes are considered: a variable-agnostic mode that treats all variables as generic placeholders (focusing purely on the expression skeletal structure) and a variable-specific mode that encodes which molecular descriptors are used with binary indicator features. AST features are compiled into a feature matrix with rows representing unique expressions and columns representing the extracted structural characteristics. Once these feature vectors are obtained, a systematic comparison of the structural characteristics of the different expressions follows.

To quantify structural similarity between expressions, we compute pairwise Euclidean similarity matrices from the standardized feature vectors. Euclidean similarity measures the distance between feature vectors with values ranging from 0 (completely dissimilar) to 1 (identical) after taking the normalized inverse distance. By computing separate similarity metrics for both modes, we assess whether expressions share similar structural motifs and/or molecular descriptors.

To illustrate this, consider the analytical expressions

in Equation (2). In the variable-agnostic case, we evaluate the number of variables used, the types of operations, and the overall structure. The variables, X_0,X_1, and X_2 are treated as interchangeable placeholders, as are X_5,X_7, and X_9. Consequently, $f_1$and f_2 are structurally similar despite their different variable names, resulting in a Euclidean similarity score of 1. In contrast, in the variable-specific case we also consider exact variable usages. Since $f_1$and f_2 use different variables, the measured similarity score is low, with a score of 0.143.

$$\begin{aligned} f_1 &= X_0 + X_1X_2 \\ f_2 &= X_5 + X_7X_9 \end{aligned} \tag{2}$$

RESULTS & DISCUSSION

SR Hyperparameter Tuning

Symbolic regression experiments were performed for six configurations of generation and population size pairs; performance was measured by the RMSE values on the validation set. Figure 1 depicts the RMSE distribution for each configuration across all data splits. Model statistics averaged across the splits for the validation set are recorded in Table 2.

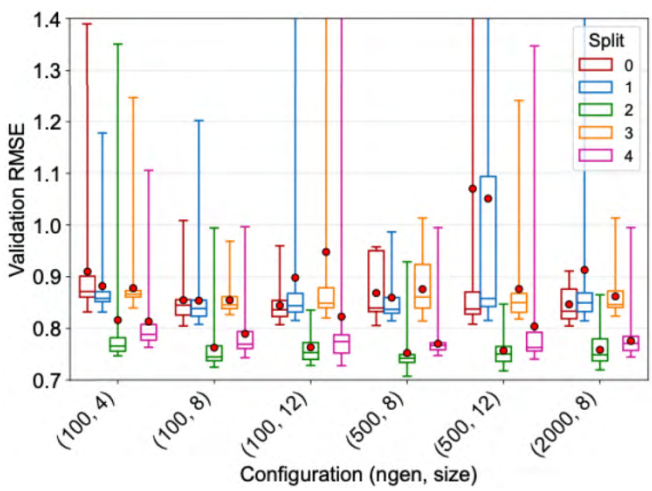

Figure 1: RMSE distribution across evolutionary parameter configurations and data splits. The red dot is the mean. In the box, the middle line is the median, the bottom line is the 1st quartile, and the top line is the 3rd quartile. The tails are the minimum and maximum values.

Friedman and Wilcoxon tests were conducted across data splits and configurations. In all splits, the Friedman test yielded p-values ≤ 0.01, indicating statistically significant differences among configurations. Building on these results, of the 75 Wilcoxon pairwise comparisons, 32 were significant, with the parameter set (ngen = 100, size = 8) showing statistical dominance. Notably, this configuration also exhibits a stable RMSE distribution and the lowest mean validation RMSE (see Table 2); therefore, it was selected as the optimal SR model parameter set. Consequently, all subsequent SR model results and baseline comparisons are based on this configuration, and use Split 1 to ensure consistent training, validation, and testing sets. The SR model with the lowest validation RMSE under these specifications is presented in Equation (3).

Table 2: RMSE validation statistics averaged across split sets.

Config.	Mean	Median	Std
(100, 4)	0.860	0.856	0.0985
(100, 8)	0.823	0.826	0.0652
(100, 12)	0.855	0.831	0.3044
(500, 8)	0.825	0.829	0.0699
(500, 12)	0.912	0.828	0.9603
(2000, 8)	0.831	0.827	0.1161

$$\begin{aligned} \log(\mathrm{HC}_{50}) = \; & 0.241X_{12} + 0.279X_{13} - 0.259\mathrm{X}_2 \\ & -0.113X_9 - 0.000413\frac{X_8^2}{X_2X_9} \\ & + 0.106\frac{X_8}{X_2X_5} + 0.145\frac{-X_1 + \frac{X_{11}}{X_4} - X_{13}}{X_0} \\ & + 27.1\frac{-X_1 + \frac{X_{11}}{X_4} - X_{13}}{X_1X_{13}(X_1 - X_7)} \end{aligned} \tag{3}$$

As shown, Equation (3) is a compact interpretable model with high variable diversity, using 11 out of the 14 available descriptors; several appear multiple times. The optimal weights (β_i) are the coefficients preceding each subexpression. Prominent monomials include water solubility (X_{12}), mode of action (X_{13}), bioconcentration factor (X_2), and melting point (X_9). Based on the signs of each variable, water solubility and mode of action appear to be positively correlated with $\log(\mathrm{HC}_{50})$, while bioconcentration factor and melting point are inversely correlated with $\log(\mathrm{HC}_{50})$. However, this is not the full picture, considering the presence of polynomial terms. These relationships are made explicit via feature importance analysis in the subsequent section.

Comparative Performance of SR, RF, and NN

All models were evaluated using the same train/validation/test sets and assessed using the coefficient of determination (R^2) and the root-mean-square error (RMSE). The performance results for training and testing are summarized in Table 3. Furthermore, parity plots (Figure 2) provide a visual assessment of model generalization, comparing SR, RF, and NN models.

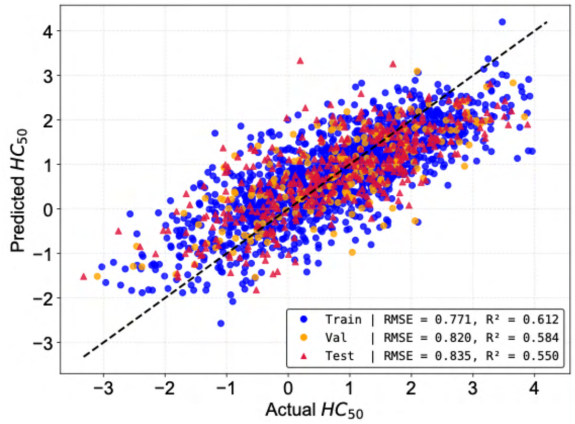

(a) SR model

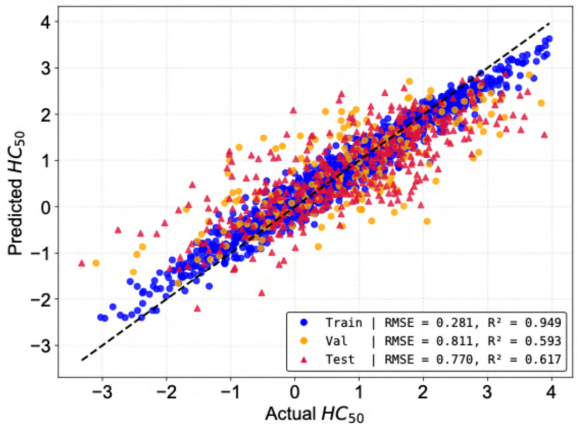

(b) RF model

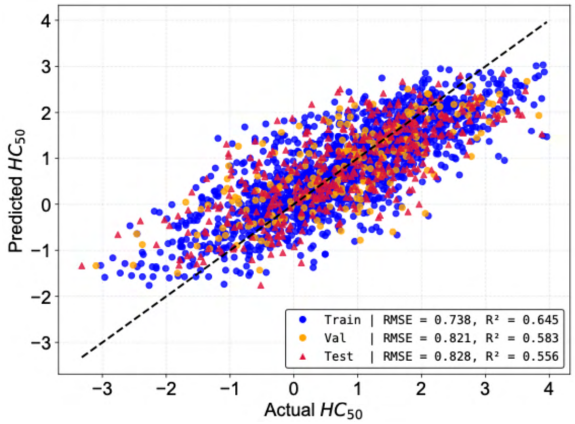

(c) NN model

Figure 2: Parity plots comparing predicted and actual $\log(\mathrm{HC}_{50})$ values for the SR, RF, and NN models on split set 1.

As shown in Table 3, the RF model achieves the highest performance on the training set, but its accuracy drops considerably on the test set- suggesting overfitting. This is evident from the large differences in both RMSE and R^2 between the training and test results. The NN model displays lower performance than the RF model in training and testing, while the SR model has the lowest training scores but performs about equally to the NN model on the test set. While RF achieves the strongest apparent fit on the training data, this advantage does not translate to unseen samples. On average, all three models perform roughly equally on the test set across all data splits.

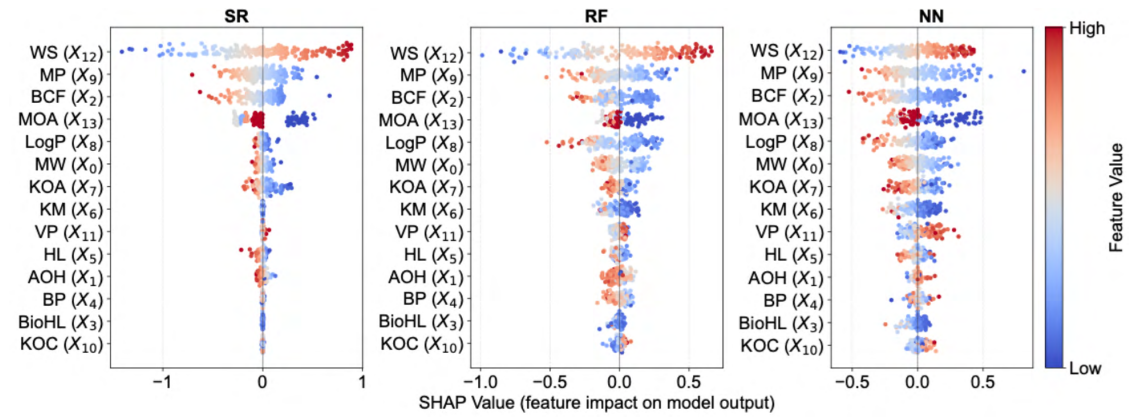

Figure 3. SHAP-based feature importances for SR, RF, and NN models. Particles are samples representing the impact distribution of each feature on the target output, indicating the direction that each feature value drives the predicted $\log(\mathrm{HC}_{50})$ output. High positive values are positively correlated with the target, and high negative values are negatively correlated.

Table 3: Average model performance metrics for RF, NN, and SR models on 5 data splits. SR corresponds to configuration ($ngen$=100, $size$=8). Standard deviations in parentheses.

Model	R^2 (Train)	RMSE (Train)	R^2 (Test)	RMSE (Test)
RF	0.930 (0.0446)	0.322 (0.0874)	0.609 (0.0136)	0.766 (0.0142)
NN	0.719 (0.0598)	0.659 (0.0707)	0.573 (0.0269)	0.800 (0.0225)
SR	0.606 (0.004)	0.785 (0.012)	0.573 (0.019)	0.801 (0.030)

Feature Importance Across Models

The feature importances for the three models, obtained using the SHAP methodology, are shown in Figures 3 and 4. Figure 3 describes the effect of each feature on $\log(\mathrm{HC}_{50})$ prediction, while Figure 4 shows the cumulative mean absolute SHAP value for all features of the compared models. The two most important features across all models are water solubility (X_{12}) and melting point (X_9). The next prominent features include bioconcentration factor (X_2), mode of action (X_{13}), and the octanol/water partition coefficient (X_8).

Figure 3 reveals that for the top features, higher water solubility increases the predicted $\log(\mathrm{HC}_{50})$ value, while higher melting point reduces the predicted value. A negative correlation is also observable for bioconcentration factor, mode of action, and octanol/water partition coefficient with higher feature values reducing the predicted $\log(\mathrm{HC}_{50})$ value.

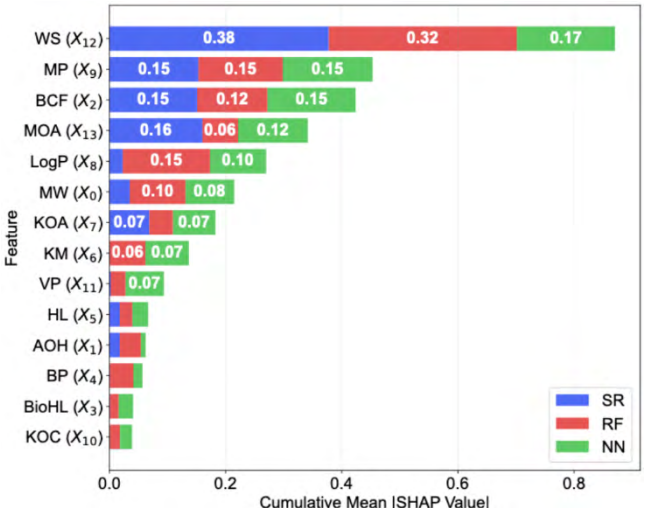

Figure 4. Mean absolute Shapley values for SR, RF, and NN models.

Overall, the findings of the feature importance analysis agree with the visual inspection of the expression in Equation (3). It clarifies the contributions of the octanol/water partition coefficient (X_8) and mode of action (X_{13}), which are ambiguous in Equation (3) due to their repeated occurrences. The analysis also reveals that vapor pressure (X_{11}) and boiling point (X_4) primarily act as structural constructors in Equation (2), as they do not exert a direct influence on the predicted value, as shown by Figure 4. While feature importances can be used to effectively extract variable information from black box models, an explicit expression remains desirable. It can be inspected against domain knowledge to verify if the discovered relationships are plausible and consistent with established theory or intuition. Furthermore, when an expression satisfies these criteria, it may uncover new

mechanistic insights.

Structural Expression Analysis

Having established how the individual features in each model contribute to predicting the target, we next examine the structural characteristics of the expressions obtained by symbolic regression. The goal is to see if a common canonical form emerges among the generated expressions. From this data-driven assessment, we seek to uncover mechanistic insights and identify the elements of these expressions that drive their convergence.

Across all evolutionary configurations, 3000 symbolic expressions were generated, of which 999 are unique and used to quantify structural similarity. Figure 5 summarizes the distribution of these expressions. The most frequent expression (Equation (4)) occurs 16 times, while roughly 300 expressions appear only once, indicating that most derived expressions are recurrent.

$$\begin{aligned}\log(\mathrm{HC}_{50}) = \quad & 0.161X_1 - 0.000167X_{10} \\ & - 0.265X_{12} - 0.170X_2 \\ & - 0.00133X_8 - 0.161X_9 + \frac{0.654}{X_{13}}\end{aligned} \tag{4}$$

Comparing the expression with the most frequent appearance against Equation (3), there is a clear divergence in the structures. Equation (4) is more compact and uses fewer features than Equation (3). Except for the soil adsorption coefficient (X_{10}), which is accompanied with a low coefficient value, all variables present in (4) also appear in (3).

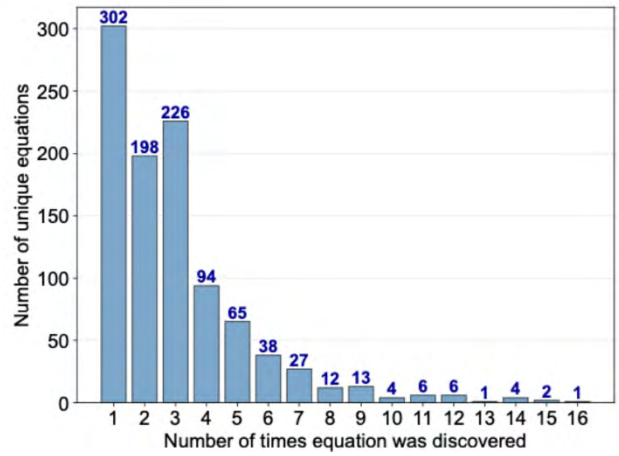

Figure 5. Frequency of unique expressions from all SR models. Labels above each bar represent the number of unique expressions for the corresponding appearance count on the x-axis.

After identifying the unique expressions, we

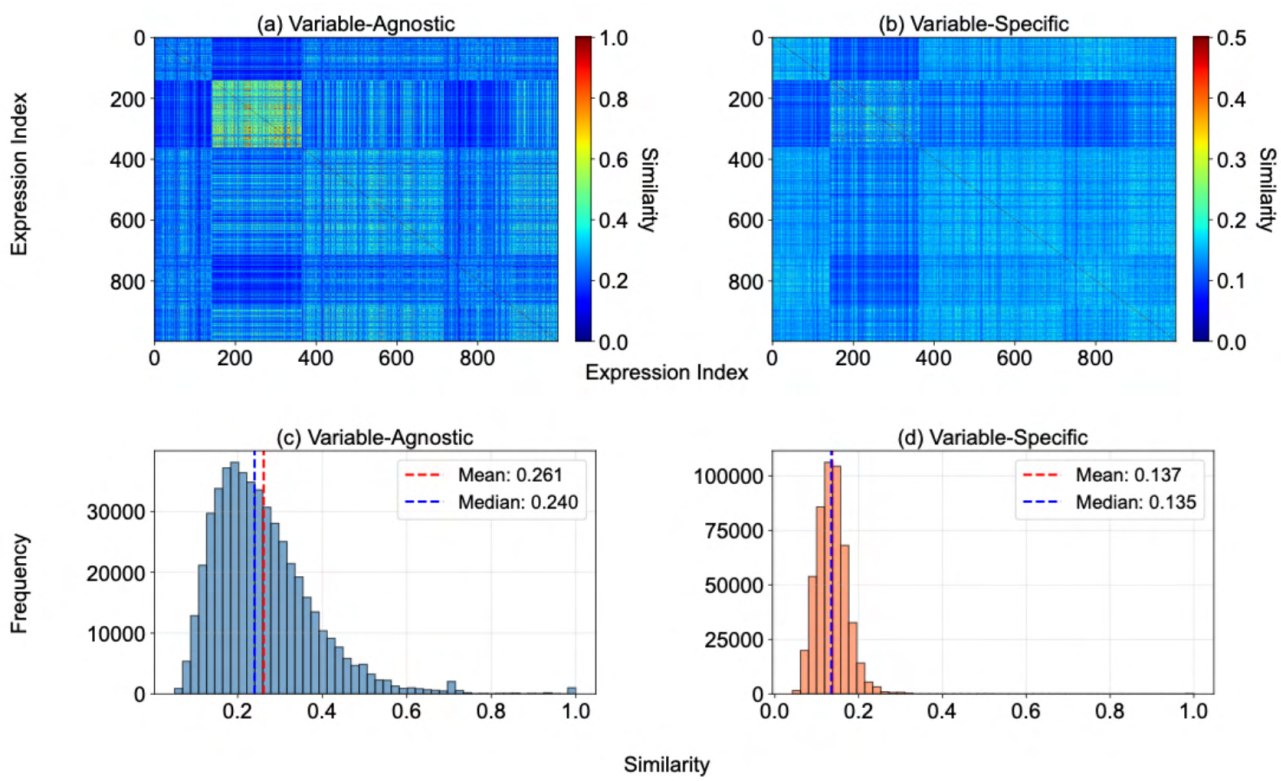

Figure 6. Characterization of structural diversity of SR expressions for variable-agnostic (left) and variable-specific (right) structural signatures. Frequencies in (c) and (d) are the number pairs appears with the similarity score on the x-axis.

proceed to characterize the structural diversity of the SR expressions for both variable-agnostic and variable-specific evaluations by measuring the similarity between expression pairs. The findings of this analysis are presented in Figure 6. Figure 6 (a and c) shows that for the variable-agnostic case (variables not explicitly considered), most expressions are not structurally similar, with a mean similarity score of 0.261. This means that expressions are extremely diverse in their use of number/type of operators, as well as in expression lengths. If we consider Equations (3) and (4), for example, the similarity score is 0.138, which falls below the mean.

The variable-specific case shown in Figure 6 (b and d) also indicates high feature diversity with a mean similarity score of 0.137. Again, comparing Equations (3) and (4), the similarity score in this case is 0.110. The near-zero mean and the distribution of the similarity scores shown in Figure 6 suggest that most expressions have little to no overlap in the overall structure. Yet, performance remains consistent across all models.

CONCLUSION

This study demonstrates that symbolic regression is a strong alternative to conventional black-box models for predicting ecotoxicity effect factors used within life cycle assessment frameworks. The symbolic regression models showed stable performance across multiple evolutionary configurations. When benchmarked against random forest and neural network models, symbolic regression achieved comparable predictive accuracy while offering an explicit analytical expression. This is particularly valuable for LCA, where traceability and scientific justification of characterization factors are paramount. Furthermore, feature importances obtained with SHAP were consistent across models, reinforcing SR's ability to compete with traditional black-box methods. Lastly, the structural analysis of the derived SR expressions revealed substantial variation, indicating that caution is needed when assigning causality. At the same time, this variation can be viewed positively, as it provides multiple pathways to accurately compute a property based on the data available. In summary, our results highlight symbolic regression as a promising method for reducing data gaps in CF estimation.

ACKNOWLEDGEMENTS

Funding received under the GEM Fellowship is gratefully acknowledged.

REFERENCES

1. Rosenbaum RK, Bachmann TM, Gold LS, Huijbregts MAJ, Jolliet O, Juraske R, Koehler A, Larsen HF, MacLeod M, Margni M, McKone TE, Payet J, Schuhmacher M, van de Meent D, Hauschild MZ. Usetox—the UNEP-SETAC toxicity model: recommended characterisation factors for human toxicity and freshwater ecotoxicity in life cycle impact assessment. Int J Life Cycle Assess 13:532-546 (2008). https://doi.org/10.1007/s11367-008-0038-4
2. Bare J. TRACI 2.0: the tool for the reduction and assessment of chemical and other environmental impacts 2.0. Clean Techn Environ Policy 13:687-696 (2011). https://doi.org/10.1007/s10098-010-0338-9
3. Venkatasubramanian V. The promise of artificial intelligence in chemical engineering: is it here, finally?. AIChE Journal 65:466-478 (2018). https://doi.org/10.1002/aic.16489
4. Chakraborty A, Sivaram A, Venkatasubramanian V. AI-DARWIN: a first principles-based model discovery engine using machine learning. Computers & Chemical Engineering 154:107470 (2021). https://doi.org/10.1016/j.compchemeng.2021.107470
5. Cozad A, Sahinidis NV. A global MINLP approach to symbolic regression. Math. Program. 170:97-119 (2018). https://doi.org/10.1007/s10107-018-1289-x
6. Talbi E. Metaheuristics. Wiley (2009). https://doi.org/10.1002/9780470496916
7. Aboagye E, Longo J, Conway M, Pazik J, Barkow M, McMahon M, Weil B, Appiah HD, Hesketh R, Yenkie KM. Leveraging machine learning algorithms to predict life cycle inventory assessments (LCIA) to facilitate sustainable process design. Computers & Chemical Engineering 201:109217 (2025). https://doi.org/10.1016/j.compchemeng.2025.109217
8. Hou P, Jolliet O, Zhu J, Xu M. Estimate ecotoxicity characterization factors for chemicals in life cycle assessment using machine learning models. Environment International 135:105393 (2020). https://doi.org/10.1016/j.envint.2019.105393
9. Breiman L. Random forests. Machine Learning 45:5-32 (2001). https://doi.org/10.1023/a:1010933404324
10. Bishop CM. Neural networks for pattern recognition. Oxford University PressOxford (1995). https://doi.org/10.1093/oso/9780198538493.001.0001
11. Ferreira J, Pedemonte M, Torres AI. Development of a machine learning-based soft sensor for an oil refinery's distillation column. Computers & Chemical Engineering 161:107756 (2022). https://doi.org/10.1016/j.compchemeng.2022.107756

12. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, É. Scikit-learn: machine learning in Python. *J. Mach. Learn Res.* 12:2825-2830 (2011).
13. Bergstra J, Yamins D, Cox D. Making a science of model search: hyperparameter optimization in hundreds of dimensions for vision architectures. *Proc. Int. Conf. Mach. Learn. PMLR* 28:115-123 (2013). https://doi.org/10.48550/arXiv.1209.5111
14. Lundberg SM, Lee SI. A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Sys. 30* (2017). https://doi.org/10.48550/arXiv.1705.07874

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Adaptive soft sensor to estimate alite fraction in clinker production through quasi-ensemble PLS modelling

Mihnea Stefan[ab], Wilson R. Leal da Silva[b], Fabrizio Bezzo[a], Pierantonio Facco[a]*

[a] *CAPE*-Lab, Computer-Aided Process Engineering Laboratory, University of Padova, Department of Industrial Engineering, Via Marzolo 9, Padova, 35131, Italy
[b] Fuller Technologies, Scandiagade 8, 4, 2450, Copenhagen, Denmark
* Corresponding Author: pierantonio.facco@unipd.it.

ABSTRACT

Cement is regarded as the most widely used construction material worldwide; however, its production is also recognized as a major contributor to global CO_2 emissions. Strict control of cement quality is therefore required to prevent excessive consumption of raw materials and energy, which would otherwise increase the process environmental footprint. Cement quality is largely governed by clinker quality, which is primarily characterized by two quality control parameters: free-lime content and alite fraction. At present, these are characterized by costly and time-consuming laboratory analyses that are not optimal for real time process control and optimization. Hence, in this work, a soft sensor for the real-time estimation of the clinker alite fraction is proposed. The developed soft sensor is designed to adapt to process drifts and operating condition changes, capture nonlinear and dynamic behavior, and retain interpretability through a Partial Least Squares (PLS) modelling framework. To this end, a novel recursively adaptive local dynamic soft-sensing strategy (ALD-PLS) is introduced and implemented within a multi-model ensemble structure referred to as Quasi-Ensemble PLS (QE-PLS). Unlike conventional ensemble approaches, where model diversity is generated through data resampling or training–testing partitioning, the proposed framework constructs multiple sub-models using combinations of model hyperparameters, evaluated on the same evolving dataset. As a result, improved predictive accuracy and robustness are achieved, while estimation uncertainty is quantified. The proposed QE-PLS soft sensor is shown to outperform, in terms of R^2 and $RMSE$, PLS-based and single-instance implementations ALD-PLS for a similar task. The novel methodology is validated against data collected from an industrial cement production plant.

Keywords: Modelling, adaptive modelling, cement industry, ensemble modelling, PLS, soft sensing

1. INTRODUCTION

Cement is regarded as the most widely used construction material worldwide due to its scalability, affordability, and widespread availability. Nevertheless, cement production is recognized as a major contributor to global carbon emissions, accounting for ca. 8% of total anthropogenic CO_2 emissions [1]. Strict quality control is therefore required, since deviations in product quality may result in unnecessary consumption of raw materials and energy, increasing the overall environmental footprint of the process and product. Cement quality is largely governed by clinker quality, which is primarily controlled by two key indices, i.e. free-lime content and alite fraction, to mention a few [2]. At present, these are determined by laboratory analyses, which are costly and time-consuming, thus not optimal for real-time process monitoring, control, and optimization.

An alternative to laboratory measurements is provided by soft sensors, through which difficult-to-measure quality variables are estimated using data-driven models built from readily available process measurements, e.g. temperatures, pressures, and material flow rates [3]. In clinker production, real-time estimation of critical quality indices by soft sensors can support advanced control strategies and improved operational efficiency. However, research on soft sensing in the cement industry has remained limited, with most existing studies

focusing on free lime and few addressing the alite fraction. Existing approaches can be grouped into three categories. First, black-box models, e.g. neural networks, have been associated with high predictive accuracy, while limited process interpretability has been provided [4]. Second, white-box Partial Least Squares (PLS)-based models have been employed to retain interpretability; however, performance has often been constrained by difficulties in capturing nonlinearity, dynamics, and process drifts typical of modern cement plants [5]. Third, mechanistic first-principles models have been characterized by both accuracy and interpretability, but high computational demand has frequently rendered their use impractical for real-time control and optimization [6].

In this work, the focus is on the real-time estimation of clinker alite fraction through an interpretable, adaptive soft-sensing strategy. In particular, PLS modelling has been adopted as the modelling strategy due to its robustness in handling collinearity and high-dimensional process data. A recursively adaptive PLS-based soft sensor has been proposed based on the established incremental PLS (I-PLS) strategy [7]. Process dynamics have been explicitly incorporated using lagged variables [8], while local modelling strategies have been employed to tailor the model to the prevailing data environment; improving predictive performance under nonlinear and non-stationary conditions. Such adjustments are needed since conventional PLS-based soft sensors do not adapt to evolving industrial conditions. In clinker production, non-stationarity, seasonal drifts, abrupt operating changes, and dynamic behavior have frequently been observed, often in association with multiple operating modes, thereby challenging static, global modelling approaches.

Note, however, that the proposed models' predictive performance was found to depend on a relatively large number of interdependent hyperparameters, including the number of latent variables and the dynamic lag structure, whose optimal values vary over time and operating conditions. Thus, the models' hyperparameters have been allowed to vary within physically and statistically meaningful bounds - instead of being fixed a priori. As a result, a new methodology based on a multi-model structure that is inspired by ensemble learning principles has been developed, resulting in an ensemble over model structures, not calibration datasets. This novel strategy is referred to as Quasi-Ensemble PLS modelling. To prove its predictive superiority the performance of QE-PLS, in terms of estimation R^2 and $RMSE$, is compared to that of its components, I-PLS and ALD-PLS.

2. MATERIALS AND METHODS

The proposed novel soft-sensing methodology was developed and validated using industrial data collected from a cement production plant. In this section, the clinker production process is briefly described, followed by an overview of the dataset and preprocessing steps, and a detailed presentation of the modelling framework adopted in this work.

2.1 Clinker production process

Cement is produced by grinding clinker together with supplementary cementitious materials and gypsum. Clinker itself is obtained by thermally processing a raw meal composed mainly of limestone and clay. The raw meal is preheated in a multistage cyclone tower, partially calcined in a precalciner, and subsequently fed into a rotary kiln where clinker formation occurs. The hot clinker is then rapidly cooled to preserve the formed alite phase [10]. A schematic representation of the clinker production process is shown in Figure 1.

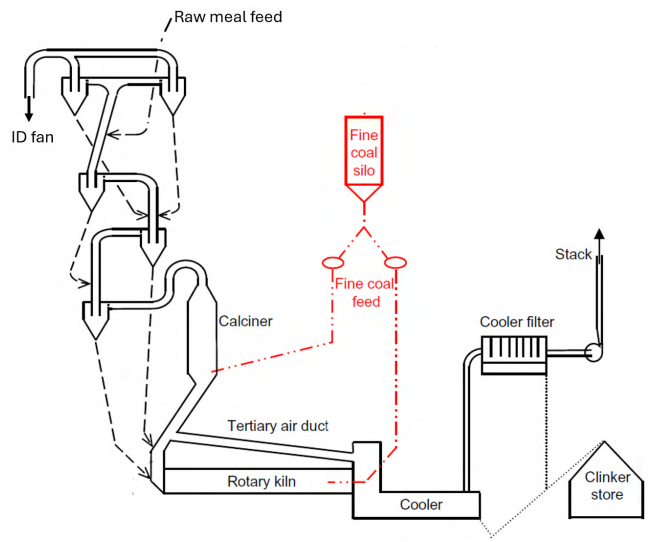

Figure 1. Clinker production process; adapted from [10].

2.2 Dataset preprocessing

The dataset used in this study comprises 74 process variables and the corresponding product quality measurements in terms of alite fraction. For confidentiality reasons, the dataset origin (cement plant) and the description of individual variables cannot be disclosed. All process variables were recorded at a nominal sampling interval of 1 min. While process variables were continuously measured, the alite fraction was obtained through laboratory analysis at a lower frequency (ca. once per hour). Hereafter, the process variables and the alite fraction are referred to as the **X**-block and the **y**-vector.

Data preprocessing steps include: (i) removal of missing values, sensor decoupling periods, and unphysical measurements; and (ii) outlier detection and removal based on Q (SPE) residuals and Hotelling's T^2 statistics [11], exceeding twice their respective 95% confidence limits in PLS modelling. This resulted in a final dataset spanning ca. two months of operation, where $\mathbf{X}$ has the

dimension of $[N_{observations} \times V_{variables}] = [826 \times 74]$ and **y** of $[N_{observations} \times 1] = [826 \times 1]$.

2.3 Partial Least Squares (PLS)

Partial Least Squares (PLS) regression was selected due to its effectiveness in handling highly correlated, high-dimensional industrial data. All PLS models were built using the NIPALS algorithm [12]. Equations 1-4 represent the model formulation and estimation formulae.

$$\mathbf{X} = \mathbf{T}\mathbf{P}^{\mathbf{T}} + \mathbf{E} \tag{1}$$

$$\mathbf{Y} = \mathbf{U}\mathbf{Q}^{\mathbf{T}} + \mathbf{F} \tag{2}$$

$$\mathbf{U} = \mathbf{T}\mathbf{B} \tag{3}$$

$$\hat{\mathbf{y}}_{\text{new}} = \mathbf{x}_{\text{new}}\mathbf{W}(\mathbf{P}^{\text{T}}\mathbf{W})^{-1}\mathbf{B}\mathbf{Q}^{\text{T}} \tag{4}$$

where, $\mathbf{X}, \mathbf{Y}$ refer to the predictor and response matrices, $\mathbf{T}$ and $\mathbf{U}$ the scores matrices, $\mathbf{P}$ and $\mathbf{Q}$ the loadings matrices, and $\mathbf{E}$ and $\mathbf{F}$ the residual matrices resulted from the decomposition of the predictor and response blocks. Besides, $\mathbf{W}$ represents the weights matrix, $\mathbf{B}$ the regression coefficients vector, and $\mathbf{x}_{\text{new}}$ the new observation vector while $\hat{\mathbf{y}}_{\text{new}}$ the estimation for the corresponding response.

However, an adjustment to classical PLS is needed, since it is inherently linear and static. In particular, models are calibrated once and, subsequently, applied without adaptation, rendering them incapable of accounting for the frequent process drifts and operating mode changes – typical of industrial clinker production. Hence, an adaptive modelling strategy – described as follows – is required to overcome these limitations.

2.4 Incremental update strategy PLS

The incremental recursive updating strategy for PLS [7] is adopted to enable model adaptation over time. Specifically, at each time instant, a PLS model is recalibrated using all available historical data, and the current process observation is projected to estimate the corresponding alite fraction. Once the true measurement becomes available, it is incorporated into the historical dataset and the procedure is repeated.

While this recursive approach allows for tracking process changes, it suffers from two drawbacks: (i) as the dataset size grows, diminishing influence on the recalibrated model is exerted by the new observations, reducing adaptivity; and (ii) increasing computational cost due to repeated recalibration on an ever-growing dataset. These limitations motivate the proposal of a local adaptive dynamic PLS framework (see Section 2.6), where a similarity-based selection of historical samples is used to maintain model relevance and estimation accuracy.

2.5 Lagged variables PLS

The standard PLS framework is extended by lagged-variables PLS through augmentation of the **X**-block with time-shifted versions of the process variables, enabling the model to capture process dynamics and temporal autocorrelation [8]. In classical dynamic PLS, the augmented **X**-block is formed by concatenating copies of the original variables shifted by predefined time lags.

In this work, a modified lag-construction strategy is adopted for the adaptive local dynamic PLS framework (Section 2.6). Specifically, instead of generating lagged variables from the main modelling dataset, lagged information is extracted from the higher-resolution 1-minute dataset (Section 2.2). This helps incorporate the dynamic information present at finer time scales, while the main **X–y** modelling dataset remains aligned with the lower sampling frequency of the alite measurements (ca. one hour).

As for lagged variables, they are constructed by selecting a lag spacing W, defining the time shift between consecutive lags. In the main dataset, for each sample at time t, process variable observations at times $t - W, t - 2W, \ldots, t - R$ are retrieved from the 1-minute dataset and horizontally concatenated to form the augmented **X**-block. Thus, the total number of lags added in the dynamic X-block reads R/W.

Altogether, this approach enables the incorporation of dynamic process information without increasing the temporal resolution of the primary dataset and forms the basis for the adaptive local modelling strategy described as follows.

2.6 Proposed methodology: adaptive local dynamic PLS (ALD-PLS)

The proposed methodology integrates I-PLS, dynamic lagged-variable methods, and local modelling into a unified adaptive local dynamic PLS (ALD-PLS) strategy. The objective is to retain the adaptability and interpretability of PLS while overcoming the limitations of static global models when applied to non-stationary, dynamic industrial processes. A schematic overview of the methodology is provided in Figure 2.

At each time t, the modelling procedure is carried out according to the following update paradigm:

1. **Step 1:** data from time 1 to $t - 1$ are used to construct a global dynamic PLS model, where the **X**-block consists of the lag-augmented process variables and the **y**-vector the corresponding alite fraction measurements. The construction of the dynamic **X**-block requires the selection of a proper lag W and the maximum lag range R (Section 2.5). The global model is built using A_g latent variables.

2. **Step 2:** the current observation of the $V = 74$ online process variables at time t, represented in its dynamic lagged-variable form, is projected into the global PLS

score space.

3. **Step 3:** Euclidean distances between the score vector of the current point and those of all historical samples are computed, and the L closest historical points are selected. These samples are considered the most relevant local neighborhood for modelling the current process state.
4. **Step 4:** using the selected L samples, a local dynamic PLS model is constructed with A latent variables.
5. **Step 5:** Variable Importance in Projection (VIP) indices [13] are calculated for each predictor in the dataset. Predictors with VIP < 1 are subsequently removed [13], and the local dynamic PLS model is recalibrated using the reduced predictor set.
6. **Step 6:** the current observation at time t is projected onto the score space of the refined local PLS model, and the alite fraction is estimated.

The abovementioned local modelling strategy ensures that estimations are based on the most relevant historical data while maintaining robustness against irrelevant or degenerate predictors.

2.7 Quasi-Ensemble PLS Modelling

The predictive performance of the proposed methodology depends on a set of critical hyperparameters (Section 2.6) summarized in Table 1. Note that the choice of hyperparameters is interdependent.

Table 1. Hyperparameters of the proposed ALD-PLS.

Hyperparameter	Description
L	size of similarity neighborhood
A	number of LVs of the local PLS model
A_g	number of LVs of the global PLS model
W	lag frequency (time shift of the variables between two consecutive lags)
R	maximum time shift considered

Instead of selecting a single optimal hyperparameter configuration, a multi-model strategy inspired by ensemble learning [9] was adopted. In contrast to classical ensemble approaches, where diversity is generated through resampling or partitioning of the dataset, the proposed framework constructs diversity through different combinations of model hyperparameters evaluated on the same evolving dataset. This approach is referred to as Quasi-Ensemble PLS modelling (QE-PLS).

All possible combinations of hyperparameters values in reasonable ranges are selected to build the sub-models of the Quasi-Ensemble (QE). The selected values of the hyperparameters, selected through heuristics, are shown in Table 2. For each sub-model, the ALD-PLS is executed, yielding an independent estimation of the alite fraction at each time point.

The QE estimation is obtained by averaging the sub-model outputs, while the associated estimation uncertainty is quantified through the standard deviation in the estimation. Specifically, for the i-th observation, the ensemble mean and standard deviation are given by Equations (5-6):

$$\hat{y}_i = \bar{y}_i = \frac{1}{M}\sum_{j=1}^{M} \hat{y}_i^{(j)} \tag{5}$$

$$\sigma_i = \sqrt{\frac{1}{M}\sum_{j=1}^{M}\left(\hat{y}_i^{(j)} - \bar{y}_i\right)^2} \tag{6}$$

where, $\hat{y}_i^{(j)}$ denotes the estimation of the i-th sample produced by the j-th sub-model, and M is the total number of sub-models in the QE; M =3840 in the present case.

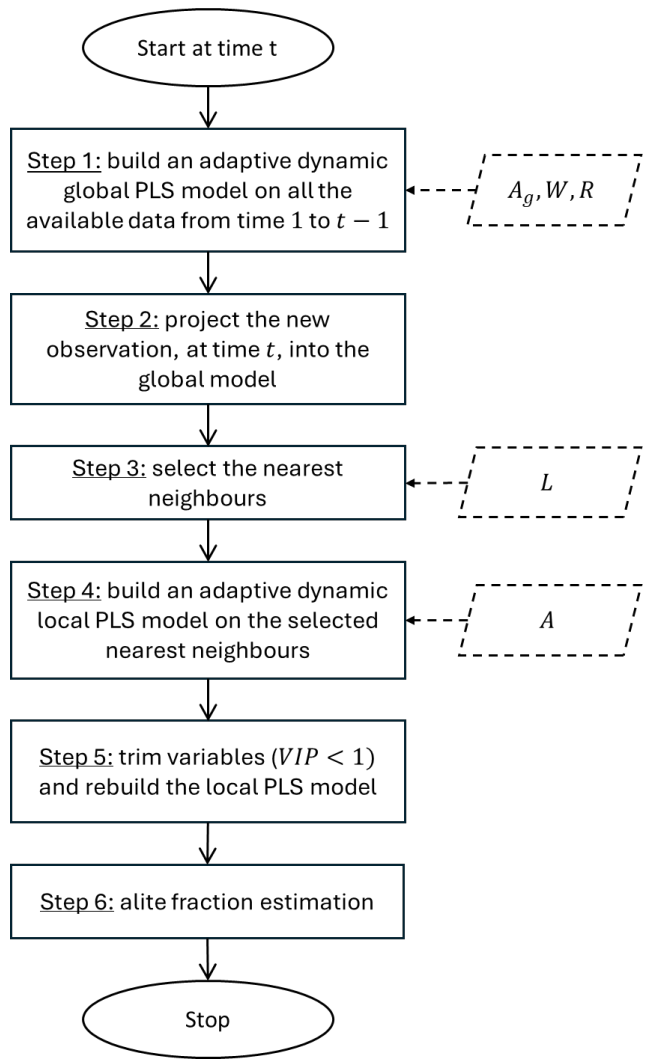

Figure 2. Schematic of the ALD-PLS at time t.

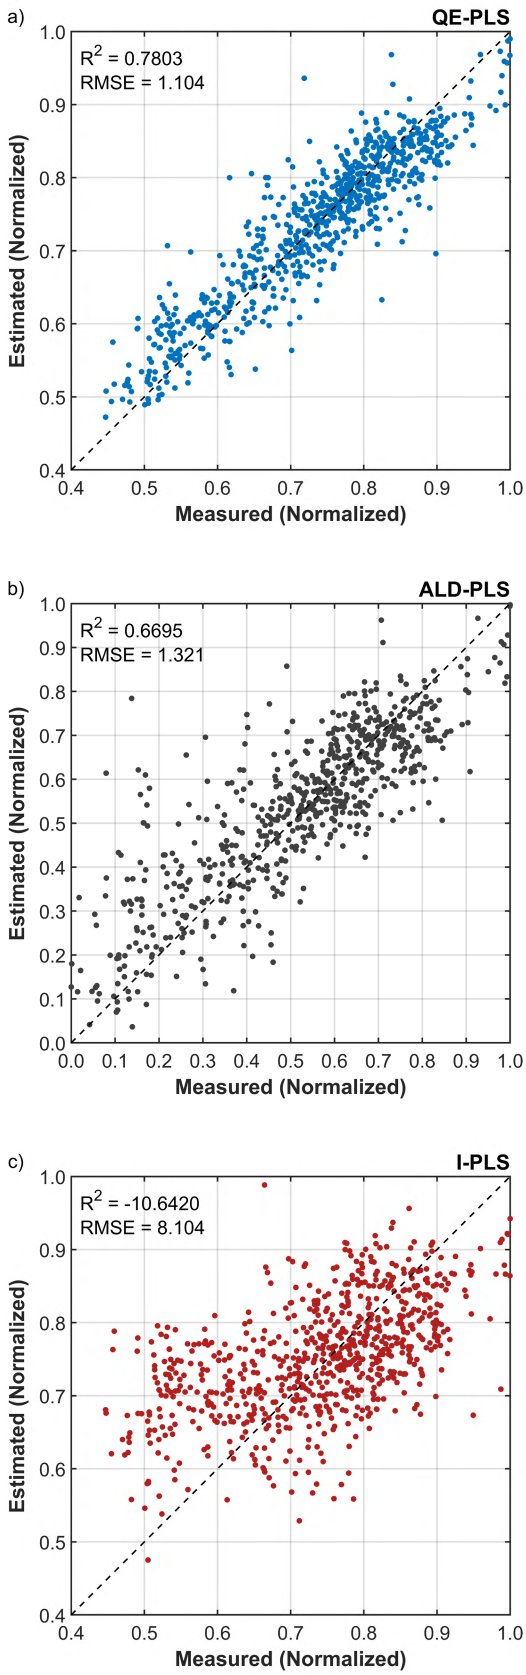

Figure 3. Parity plot of predicted alite fraction vs measured alite fraction: a) QE-PLS, b) ALD-PLS and c) I-PLS.

Based on σ_i, approximate 95% confidence intervals for the final estimations are computed as $\bar{y}_i \pm 2\sigma_i$. The QE-PLS allows the proposed soft sensor to remain robust to hyperparameter uncertainty and process variability, while simultaneously providing predictive uncertainty estimates without requiring explicit probabilistic modelling.

Table 2. Selected hyperparameters ranges.

Hyperparameter	Range
L	30:10:60
A	1:1:8
A_g	1:1:10
W	30 minutes
R	2:2:24 hours

2.8 Performance Metrics

In this work the performance metrics used to assess the quality of the proposed soft sensor are the determination coefficient, R^2, and the root mean square error, $RMSE$; i.e.:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} \tag{7}$$

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{8}$$

where, y_i is the i-th elements of the **y** measurements, $\hat{y}_i$, the estimation for that element, $\bar{y}$ the mean alite fraction data, and n the number of samples in the **y** dataset.

3. RESULTS

The proposed methodology was applied to the preprocessed **X** and **y** datasets (Section 2.2). The resulting alite fraction estimations are shown in the parity plot of Figure 3 a). The QE-PLS soft sensor achieved an R^2 of 0.7803 and a $RMSE$ of 1.1037. Considering the intrinsic variability of the alite fraction measurements, this indicates an accurate and reliable predictive performance.

To further assess the proposed approach, a comparative study was conducted against two state-of-the-art PLS-based soft-sensing strategies, i.e.: (i) a standard incremental update PLS model, i.e. I-PLS (Section 2.4), and (ii) a single instance of the adaptive local dynamic PLS model, i.e. ALD-PLS, corresponding to one sub-model of the QE. The corresponding parity plots for a single instance of the ALD-PLS and the I-PLS are presented in Figure 3 b) and 3 c), respectively.

The comparative study results are listed in Table 3. They indicate that the QE-PLS model outperforms both I-PLS and ALD-PLS models on R^2 and $RMSE$. Besides, QE-PLS is shown to improve the performance of ALD-PLS sub-models, demonstrating the benefit of aggregating

estimations across multiple hyperparameter configurations. Such improvement highlights the ability of the QE framework to enhance robustness and adaptability when modelling the nonlinear and dynamic behavior of alite fraction data in industrial clinker production.

Table 3. Comparison of the estimation performance of tested models on predicting alite fraction.

Model	R^2	*RMSE*
QE-PLS	0.7803	1.1037
ALD-PLS	0.6695	1.3212
I-PLS	-10.6420	8.1038

The tested models were evaluated on a laptop equipped with a 13th Gen Intel Core i9-13900H processor (2.60 GHz), 32 GB RAM, and MATLAB R2024b. Under this configuration, I-PLS showed the lowest computational cost, requiring milliseconds for the estimation of one observation, compared with deciseconds for ALD-PLS and a few minutes for QE-PLS. Thus, I-PLS was much faster than ALD-PLS and QE-PLS. This trend is consistent with the increasing complexity of the models; notably, QE-PLS comprises 3840 ALD-PLS sub-models, which explains its substantially higher runtime. Although this computational burden is a limitation of the proposed soft sensor, the method remains suitable for quality control, since the current laboratory-based reference measurements are acquired only once per hour.

4. CONCLUSIONS & FUTURE WORK

In this work, a soft-sensing framework for the real-time estimation of clinker alite fraction was proposed. The results obtained for the industrial case study (clinker production) indicate that the proposed adaptive PLS-based soft sensor is capable of effectively capturing the nonlinear and dynamic behavior of the alite fraction under realistic operating conditions.

In addition, a novel modelling strategy, i.e. Quasi-Ensemble PLS (QE-PLS), was introduced. The proposed strategy constructs multiple sub-models using different combinations of model hyperparameters evaluated on the same evolving dataset. The QE-PLS framework was shown to outperform both classical incremental PLS (I-PLS) models and single-instance adaptive local dynamic PLS (ALD-PLS) models, demonstrating improved predictive accuracy and robustness.

Despite the achieved favorable performance, the proposed QE-PLS framework entails a relatively high computational cost due to the evaluation of several sub-models. Hence, future work shall focus on reducing computational requirements through optimization of the hyperparameter combination space, as well as through the investigation of more advanced parallelization and asynchronous execution strategies. Besides, further validation will allow for assessing the generalizability and robustness of the proposed approach. Finally, the inherent interpretability of the PLS-based framework will be exploited to support process understanding, anomaly detection, and explainability in industrial process monitoring and control applications.

5. ACKNOWLEDGEMENT

The authors would like to acknowledge the support from the project “Data Enabling Transformation and Optimization towards Concrete Sustainability” (DETOCS) that is funded by the European Union’s Horizon Europe research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 101119929. Funded by the European Union - views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency (REA). Neither the European Union nor the REA can be held responsible for them.

6. AUTHOR IDENTIFIERS

Author ORCIDs:
Stefan M.: 0009-0006-2898-4023
Leal da Silva W. R.: 0000-0001-9895-5128
Bezzo F.: 0000-0003-1561-0584
Facco P.: 0000-0001-8383-6783

REFERENCES

1. Lehne J, Preston F. Chatham House Report Making Concrete Change Innovation in Low-carbon Cement and Concrete #ConcreteChange [Internet]. www.chathamhouse.org
2. Lyu X, Chu D, Lu X, Mu J, Zhang Z, Yun D. Quality control of cement clinker through operating condition classification and free calcium oxide content prediction. Applied Sciences 14:1119 (2024). https://doi.org/10.3390/app14031119
3. Souza FAA, Araújo R, Mendes J. Review of soft sensor methods for regression applications. Chemometrics and Intelligent Laboratory Systems 152:69-79 (2016). https://doi.org/10.1016/j.chemolab.2015.12.011
4. Pani AK, Vadlamudi VK, Mohanta HK. Development and comparison of neural network based soft sensors for online estimation of cement clinker quality. ISA Transactions 52:19-29 (2013). https://doi.org/10.1016/j.isatra.2012.07.004
5. Lin B, Recke B, Knudsen JKH, Jørgensen SB. A systematic approach for soft sensor development. Computers & Chemical Engineering 31:419-425

(2007). https://doi.org/10.1016/j.compchemeng.2006.05.030

6. Fayaz SJ, Montiel-Bohórquez N, Bishnoi S, Romano M, Gatti M, Krishnan NMA. Industrial-scale prediction of cement clinker phases using machine learning. Commun Eng 4: (2025). https://doi.org/10.1038/s44172-025-00432-3
7. Helland K, Berntsen HE, Borgen OS, Martens H. Recursive algorithm for partial least squares regression. Chemometrics and Intelligent Laboratory Systems 14:129-137 (1992). https://doi.org/10.1016/0169-7439(92)80098-o
8. Dong Y, Qin SJ. Regression on dynamic PLS structures for supervised learning of dynamic data. Journal of Process Control 68:64-72 (2018). https://doi.org/10.1016/j.jprocont.2018.04.006
9. Cao D, Deng Z, Zhu M, Yao Z, Dong J, Zhao R. Ensemble partial least squares regression for descriptor selection, outlier detection, applicability domain assessment, and ensemble modeling in QSAR/QSPR modeling. Journal of Chemometrics 31: (2017). https://doi.org/10.1002/cem.2922
10. Hewlett PC, Liska M. Preface. Lea's Chemistry of Cement and Concrete :xxxiii-xxxiv (2019). https://doi.org/10.1016/b978-0-08-100773-0.09998-6
11. Odgers J, Kappatou C, Misener R, García Muñoz S, Filippi S. Probabilistic predictions for partial least squares using bootstrap. AIChE Journal 69: (2023). https://doi.org/10.1002/aic.18071
12. Geladi P, Kowalski BR. Partial least-squares regression: a tutorial. Analytica Chimica Acta 185:1-17 (1986). https://doi.org/10.1016/0003-2670(86)80028-9
13. Zheng R, Chen Z, Guan Z, Zhao C, Cui H, Shang H. Variable importance for projection (VIP) scores for analyzing the contribution of risk factors in severe adverse events to xiyanping injection. Chin Med 18: (2023). https://doi.org/10.1186/s13020-023-00718-8

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Semantic PEA Datasheets for digitalised modular plant documentation

Sascha Lamm[a]*, Sebastian Tecl[a], Ingo Dietrich[b], Sissy Sommer[b], Markus Heinbücher[a], Peter F. Pelz[a]

[a] Technical University of Darmstadt, Chair of Fluid Systems, Darmstadt, Germany
[b] Industrial Science GmbH, Darmstadt, Germany
* Corresponding Author: sascha.lamm@tu-darmstadt.de.

ABSTRACT

Modular plants emerged as the key solution for reducing time-to-market and increasing flexibility in the process industry by combining different modules known as Process Equipment Assemblies (PEAs). While PEA automation is standardised through the Module Type Package (MTP), comparable tools for their documentation remain absent. This work presents the Semantic PEA Datasheet (SPEAD) ontology, which represents PEA documentation as a machine-readable knowledge graph that adheres to the FAIR principles. SPEAD integrates established standards such as DEXPI and the VDI 2776 guidelines and ensures data quality through comprehensive annotations and constraint-based validation. The ontology was evaluated against twelve competency questions derived from a representative use case as well as competency questions from the literature using a continuous stirred-tank reactor PEA as well as a dosing PEA as example systems. SPEAD successfully covers operational and design parameters as well as interface definitions, and provides provenance tracking through rich metadata. Comparison with existing information models OntoCAPE and DEXPI shows that SPEAD provides enhanced provenance tracking essential for trust in engineering documentation. Future work includes enriching capability and safety descriptions and exploring integration with the Asset Administration Shell (AAS).

Keywords: Information Management, Modelling, Industry 4.0, Documentation, Modular Plants, Ontology, Knowledge Graphs

INTRODUCTION

In recent years, modular process plants have emerged as a key solution to tackle the growing need for reduced time-to-market and increasing flexibility in the process industry [1]. Modular plants are built by combining several modules, known as Process Equipment Assemblies (PEAs) [2]. PEAs perform single process steps, such as dosing or heating [2]. By combining and re-combining different PEAs, the Owner/Operator (O/O) of the plant can adapt their plant to fit the requirements of the product to be produced [2]. This process is called orchestration and replaces detailed equipment engineering with module selection [2].

Guidelines for the design of modules, their automation both on PEA and plant level, and the documentation of the modules are already established. The design of a PEA depends on the vendor, the operator, and general design requirements and is thus not unified across the industry [3]. The automation of a module is performed in a standardised manner using the Module Type Package (MTP) [4]. However, no comparable standard tool for the documentation of modules exists to date. In the guideline VDI 2776-1 [2], minimum requirements for the content of the documentation are given. The guideline further recommends the provision of documentation in a modular, machine-readable format. Guideline VDI 2770 provides such a format [5], which defines the minimum requirements for the exchange of digital manufacturer documentation. Neither guideline specifies how these requirements should be fulfilled, nor how information should be represented. This may lead to inconsistent data, metadata and contradictory interpretation of data. In order for the documentation to be usable across the lifecycle of a PEA, it must contain information on all aspects of the PEA and provide clear, unambiguous definitions of critical terminology and data.

Research Objective

To address this gap, we previously proposed the Semantic PEA Datasheet (SPEAD) [6]. SPEAD utilises semantic web technologies to represent the documentation of a PEA as a knowledge graph.

In this work, we present the SPEAD ontology, focusing primarily on the modelling of physical properties. This paper is structured as follows: First, we present current documentation standards and guidelines, semantic web technologies, and related models and standards. Next, we present our modelling methodology and the resulting ontology. Afterwards, we evaluate our model against competency questions developed during the modelling phase using a continuous stirred-tank reactor PEA and a dosing PEA as examples. Finally, the results are discussed and a conclusion is given.

Documentation Standards and Guidelines

Guideline VDI 2776-1 recommends that the documentation of a modular plant is hierarchically structured in the same way as its modular structure [2]. Documentation for PEAs, their contained functional equipment assemblies (FEAs), and additional components should be stored in a database that enables cross-referencing between corresponding entities. Updates performed on any element are thus inherited by higher level elements.

Guideline VDI 2776-2 recommends that PEA documentation contains at least (1) an overview of implemented process functions, (2) an overview of implemented services, (3) the interface requirements both towards adjacent PEAs as well as the (non-modular) environment, and (4) a presentation of the intended area of use. The PEA documentation should reflect its life cycle [3]. Services should be described according to the MTP [4], and a P&ID should be provided in an open data format such as DEXPI. Furthermore, it is recommended that documentation be supplied as modular digital information with an open data format, and that the exchange format described in guideline VDI 2770-1 be used. In this guideline, documents are classified according to their purpose and stored alongside relevant metadata such as ID and provenance in containers [5].

Lastly, Regulation 2023/1230 of the European Parliament and of the Council (commonly “Machinery Directive”) demands certain information to be present in the technical documentation in accordance with harmonised standards or common specifications, such as design and manufacturing drawings [7]. Where these standards and specifications have not been or have only been partially applied, descriptions of the utilised standards or specifications must be provided [7].

Semantic Web Technologies

Engineering documentation for modular plants should adhere to the FAIR principles to ensure data is Findable, Accessible, Interoperable, and Reusable throughout the plant lifecycle and across organisational boundaries. Semantic web technologies provide the standardised framework necessary to achieve these goals through structured, machine-readable knowledge representation. In the following paragraph we introduce the core W3C standards that form the technical foundation of our ontology.

The Resource Description Framework (RDF) is the core technology for the semantic web. It allows information about resources, e.g., documents, physical objects, etc., to be expressed in the form of triples [8]. Triples always have the form subject-predicate-object, where the subject and object are the resources being related, and the predicate is their relationship. A collection of these triples forms a directed labelled graph in which the subject and object are nodes and relationships are edges. RDF enables the creation of vocabularies and ontologies that provide reusable definitions of common terminology. The creation of ontologies is aided by the Web Ontology Language (OWL). OWL itself is an RDF vocabulary and supports the publication and sharing of ontologies by providing abstract constructs for the definition of classes, properties, restrictions, and logical axioms [9]. To ensure data quality and consistency, RDF graphs can be validated against a set of constraints using the Shapes Constraint Language (SHACL). Shapes are defined to specify how targeted class instances and properties can be constructed. These constraints include but are not limited to data type, cardinality, and logical constraints [10]. In this work we use the Terse RDF Triple Language (Turtle) syntax for serialising RDF data due to its human-readable format and widespread support.

Related vendor-independent standards

Koch et al. [11] evaluated existing vendor-independent standards in information modelling for process engineering, identifying three key standards: OntoCAPE, DEXPI, and IEC 61512.

1. OntoCAPE [12] is an ontology covering the domain of computer aided process engineering. It consists of multiple modules spread across six layers, each covering different topics across the process engineering domain.
2. The DEXPI (Data Exchange in the Process Industry) specifications [13] provide two complementing data models, DEXPI Plant and DEXPI Process. DEXPI Plant is a standardised data model for P&IDs, whereas DEXPI Process is a data model for process engineering information. Both specifications aim to enable consistent data exchange in the process industry. Both were recently unified into the DEXPI 2.0 model.
3. The IEC 61512 standard [14] deals with batch control and defines three models: a process model, a physical

model, and a procedural model for linking automation, operations and procedures to the other models. It defines three types of control and distinguishes between different recipe types.

Based on their evaluation, OntoCAPE and the DEXPI specifications were identified as the most suitable information models for modular plants [11]. However, neither model should be considered as a standalone solution. Instead, Koch et al. suggest the integration and extension of these models towards an ontology for modular plants [11]. The goal of this work is to lay a foundation for such an ontology, with the focus of applying the model to the documentation of modular plants.

METHODOLOGY

Development Approach

The SPEAD ontology was modelled using the use case- and competency question-based method by Kendall and McGuinness [15]. Competency questions (CQ) are derived from the use case and represent requirements for the knowledge graph that must be fulfilled by it. A representative, iteratively developed use case for the SPEAD ontology is given in the following.

Use case: Exchange of PEAs in a modular plant

A service technician is charged with replacing the reactor PEA from an existing plant with a new one from a pool of available reactor PEAs. To do this, they require information about the operating conditions, structure, and available services of all PEAs in the pool.

The derived competency questions CQ1 – CQ8 are listed in Table 1. These CQs focus on the content of the information model. Based on the CQs, lists are created which contain the necessary terminology to answer the CQs. For example, to answer CQ1, terms such as “PEA” and “operating pressure” must be defined. Further, a representation for quantitative properties in the format name/symbol = value-unit, e.g., vessel diameter D = 300 mm, must be established. To further organise the domain knowledge, we group the information into five categories: (i) operating range, (ii) construction, (iii) interfaces, (iv) services, and (v) safety. In this work we focus on (i) and (ii) and only briefly introduce our concept for (iii), (iv), and (v).

In addition, we derived questions regarding the provenance and terminology of the model, which is shown in Table 2. The provenance-related questions enable the model to capture information that is critical for regulatory compliance, change management, and establishing accountability in collaborative engineering environments. The terminology-related questions ensure that domain concepts are consistently defined, thus enabling unambiguous communication between stakeholders and interoperability between systems. Together these CQs address the metadata layer essential for trust, traceability, and quality assurance in engineering documents.

Table 1: Competency questions and the expected answers regarding the content of the information model.

CQ	Competency Question	Expected answer
1	What is the maximum operating pressure of this PEA?	The maximum operating pressure of this PEA is X bara.
2	List the boundaries of the operating temperature of this PEA.	A table with the desired boundaries.
3	Find all PEAs with a maximum operating pressure larger than 20 bar and a minimum operating temperature below 10°C.	A list of IDs of PEAs that fulfil the criteria.
4	What is the total volume of the reactor vessel of this PEA, and how much of it is heated?	The total volume and the heated volume of the reactor vessel.
5	Which FEAs and components does this PEA contain?	A list of FEAs and components of this PEA.
6	How many inlet interfaces does this PEA have that can be used for educts or flushing?	The count of inlet interfaces with the specified function.
7	What is the IP rating of this PEA?	The IP rating of this PEA.
8	Which MTP services are available at this PEA?	A list of services this PEA provides.

Table 2: Competency questions and the expected answers regarding the terminology and provenance of information inside the model.

CQ	Competency Question	Expected answer
9	Which document is the source of this value?	Reference to the source document.
10	When was this value added to the documentation of this PEA?	The date on which the value was added.
11	What is the definition of the class *spead:ProcessEquipmentAssembly*?	The definition text from the ontology.
12	List all physical properties of this PEA whose *rdfs:label* property contains the substring “flow”.	A list of physical properties with labels such as “volume flow” or “flow rate”.

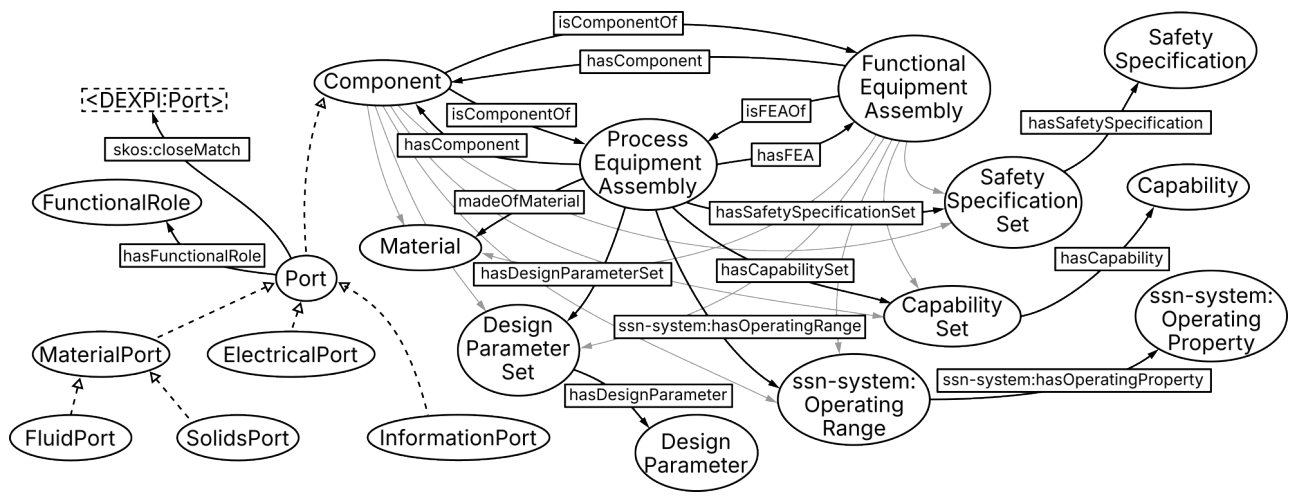

Figure 1. Core classes and properties from the SPEAD ontology. The thin grey lines depict relationships between the Component class and FEA class and other core classes that are only labeld between the PEA class and other core classes. Dashed lines indicate that the subject is a subclass of the object. The *spead:Port* class closely matches the definition of a port from the DEXPI 2.0 specification.

SPEAD Ontology

The core classes and properties are shown in Figure 1. Unless otherwise stated, all definitions shown are our own, and labelled with the prefix *spead:*. Components, FEAs, and PEAs are specialisations of OntoCAPE's plant item. Their aggregation relationships are derived from the VDI 2776-1 guideline [2]. We introduce intermediate grouping classes (e.g., *spead:SafetySpecificationSet* and *spead:DesignParameterSet*) to aggregate related property specifications from four of our five categories, drawing on the system capabilities module of the Semantic Sensor Network (SSN) ontology (prefix *ssn-system:*) [16] for the operating range. While these add verbosity, the grouping allows logically related specifications to be directly referenced. Further, sets can carry their own metadata and can be shared across different PEA instances, which simplifies the instance graph when a PEA pool contains multiple copies of the same type of PEA. Components and FEAs may also carry these sets, as indicated by the thin grey lines in Figure 1.

For interfaces, we define the abstract *spead:Port* class, which is specialised further into material, electrical, and information ports. This split was derived from the DEXPI 2.0 process model [13], and is mapped to the respective classes using *skos:closeMatch* as shown in Figure 1.

We define the abstract class *spead:QuantifiableProperty* for all equipment-related properties which are typically expressed as a name, a value, and a unit. In some cases, these kinds of properties occur as single values, e.g., the volume of a tank vessel. In other cases, values are a range or set, e.g., the maximum and minimum operating pressure of a tank vessel. These properties are also categorised differently, e.g., design parameters for attributes such as tank volume or procedure parameter for properties related to an MTP service [4]. In the instance graph, these properties can also be instances of the *spead:QuantifiableProperty* class, providing them with a shared structure. We choose this approach over subclassing to allow qualitative properties (categorisations, colours, etc.) to be described using the same classes as quantitative properties.

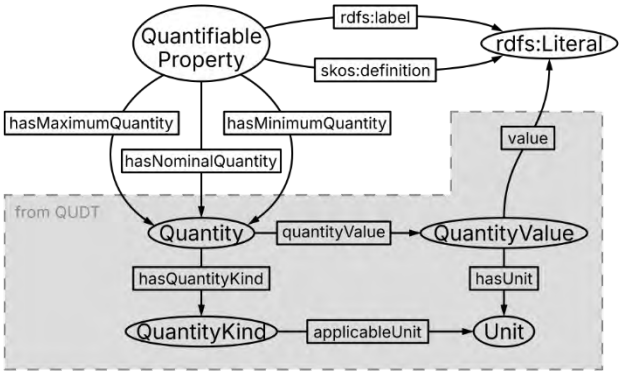

Figure 2. Schema for describing quantifiable properties. Classes and properties outside the gray box are own definitions, excluding SKOS and RDFS terms.

A detailed schema for the *spead:QuantifiableProperty* class is shown in Figure 2. We utilise the structure of the Quantities, Units, Dimensions and Data Types Ontologies (QUDT), which provides a schema for the description of quantities (prefixed with *qudt:*) as well as vocabularies that define a large number of quantity kinds (prefix *quantitykind:*), units (*prefix unit:*), dimension vectors, and constants [17]. Using these kinds of vocabularies ensures

that every stakeholder in the PEA lifecycle has the same, unambiguous understanding of the described quantities. This is relevant in cases where vendors and operators use different internal vocabulary. For example, one party might call the rate at which a fluid flows through a dosing module "flow rate", while the other calls it "volume flow". In our model, both parties link to the QUDT instance of *quantitykind:VolumeFlowRate*, maintaining their internal naming schema whilst preserving the interoperability of the data. We deviate from the structure of physical quantities in OntoCAPE, since QUDT provides an actively maintained, cross-domain, standardised structure for the representation of quantities, dimensions and units. Additionally, it contains an extensive vocabulary of kinds of quantities and units. These include 2896 unit descriptions, whereas OntoCAPE's 41 unit descriptions miss crucial units for process systems, such as those for volume flow. By leveraging QUDT, we eliminate the need of creating additional unit and quantity descriptions, thus removing a potential source of inconsistencies.

For provenance tracking, we transform the data model from the VDI 2770 guideline [5] into an RDF format. As a base we use the PROV Ontology due to its extensive vocabulary for tracking entities, their derivations, involved agents, and activities [18]. Figure 3 shows a small part of the transformed data model and how it relates to the components, FEAs, and PEAs.

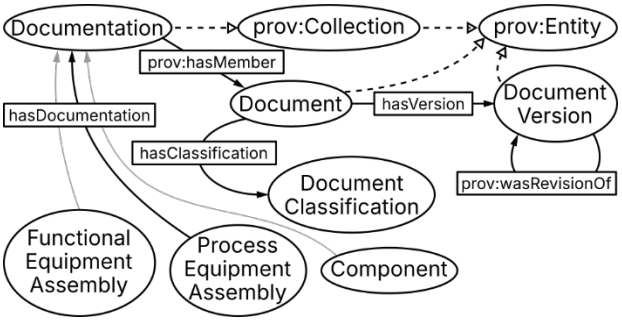

Figure 3. Classes and properties for tracking provenance. The thin grey lines depict the hasDocumentation relationship between FEAs and Components and their documentation.

Data Quality Assurance

Data quality is ensured at two levels. Our OWL ontology is intentionally kept lightweight, defining only class hierarchy and property semantics and avoiding complex class restrictions. Structural validation is instead delegated to SHACL shapes. In an industrial context, the Closed World Assumption SHACL operates under is important: unlike OWL's Open World Assumption, which treats missing information as merely unknown, SHACL treats absence as a validation error.

At the ontology level, we utilise ontologies for metadata and provenance descriptions such as Simple Knowledge Organization System (SKOS, prefix *skos:*) [19] and DCMI Metadata Terms (DC Terms, prefix *dcterms:*) [20]. All class and property definitions were annotated with labels, definitions and examples, as well as common alternative labels and sources where appropriate in both English and German. External resources (e.g., definitions from DEXPI) were related using SKOS mapping properties. Finally, we utilise the Basic Formal Ontology (BFO, prefix *bfo:*), as a foundational ontology. Foundational ontologies supply generic categorisation of the kind of things that an ontology describes, domain-independent relations, and means to represent attributes of the described things. The use of foundational ontologies provides a framework of fundamental distinctions, and facilitates the integration of different domain-specific ontologies based on the same foundational ontology. In our ontology, we use the OWL implementation of BFO, and declare our classes and properties as specialisations of the appropriate BFO terms. For example, a PEA is an object aggregate, i.e., a material object that consists of multiple parts that work together as a functional entity. We document the alignment rationale with the *skos:scopeNote* property. The serialisation of this example is as follows, where the object aggregate term is called *bfo:BFO_0000027*:

```
spead:ProcessEquipmentAssembly a owl:Class ;
  rdfs:subClassOf bfo:BFO_0000027 ;
  skos:scopeNote "Aligned to BFO 'object
aggregate', as PEAs are aggregations of FEAs
and components that together form a
unit."@en .
```

At the instance level, SHACL (prefix *sh:*) shapes are used to validate structural conformance. Amongst others, these include shapes that validate the aggregation of components and FEAs to PEAs, which is comprised of validating the domain and range of the properties shown in Figure 1, as well as ensuring that components and FEAs are only part of one PEA, excluding transitive relationships (e.g., a component that is part of a FEA which is part of a PEA is by the transitive characteristic of their relationships also part of that PEA). Further, we introduced shapes to ensure that all instances have an identifier, expressed as *dcterms:identifier*, and a source expressed as *dcterms:source.*

Since we retained the structure of QUDT, the SHACL shapes provided by QUDT can be reused. We expand on them, adding cardinality restrictions and further structural requirements. For example, QUDT allows for a quantity to have zero or more quantity kinds. Since the data in a documentation must be clear and leave as little room for faulty interpretation as possible, we enforce that every quantity has exactly one quantity kind, as shown in the following.

```
spead:QuantityShape a sh:NodeShape ;
```

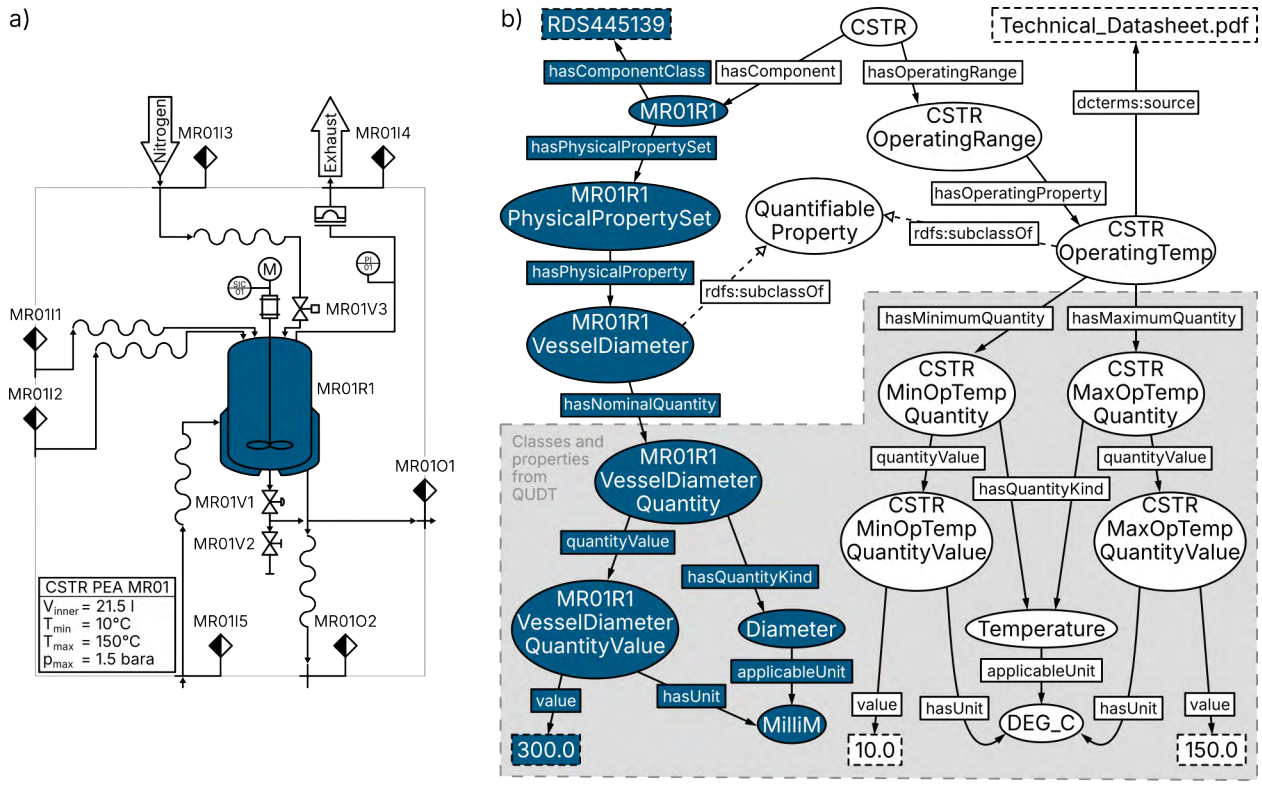

Figure 4. a) Simplified P&ID of a Continuous Stirred-Tank Reactor (CSTR) PEA developed within the project REUNION. The reactor component is highlighted in blue. b) Excerpt from the SPEAD representing the CSTR PEA. Highlighted in blue are nodes and properties relating to the reactor component in a). Boxes with dashed outlines indicate literal values, and the properties they are the subject of are datatype properties. Unless otherwise stated, all classes and properties are own definitions.

```
sh:targetClass qudt:Quantity ;
sh:property [
   sh:path qudt:hasQuantityKind ;
   sh:class qudt:QuantityKind ;
   sh:minCount 1 ;
   sh:maxCount 1 ;
] .
```

Further, the minimum and maximum quantity of a property must link to the same quantity kind using the *qudt:hasQuantityKind* property. The quantity values of those quantities must also be associated with at least one unit that is applicable to that specific quantity kind.

EVALUATION

Case study: CSTR PEA & Dosing PEA

Two PEAs are used as case studies for the evaluation of our ontology: a continuous stirred-tank reactor (CSTR) PEA as well as a dosing PEA. Figure 4a) shows a simplified P&ID of the CSTR PEA, which is developed as part of the REUNION project. The reactor component is highlighted in blue. The documentation of both PEAs was evaluated and information relevant to the competency questions was transformed into an instance graph of the SPEAD ontology, an excerpt of which is shown in Figure 4b) for the CSTR PEA. Nodes, relationships and values that directly relate to the reactor component are highlighted in blue. Figure 4b) shows how the diameter of the reactor vessel as well as the operating temperature of the CSTR PEA are modelled using our ontology. The vessel component is identified as a vessel-class component through a reference to the same Reference Data Library (RDL) used by DEXPI via the datatype property *spead:hasComponentClass*. Provenance of the operating temperature property is given with the *dcterms:source* property. In total, the SPEAD of the CSTR PEA contains 358 triples, while the SPEAD of the dosing PEA includes 214 triples, excluding any inferred triples.

We validated the ontology against the CQs shown in Table 1 and Table 2 through SPARQL queries. As an example, the following query was used to validate CQ1:

```
SELECT ?val ?unitLbl
WHERE {
   ?pea dcterms:identifier “MR01” .
   ?pea ssn-system:hasOperatingRange/
ssn-system:hasOperatingProperty ?prop .
```

```
  ?prop spead:hasMaximumQuantity/qudt:hasQuantityKind quantitykind:Pressure .
  ?prop spead:hasMaximumQuantity/qudt:quantityValue ?valNode .
  ?valNode qudt:value ?val .
  ?valNode qudt:unit/rdfs:label ?unitLbl .
}
```

For the content-related questions, the instance graph of both PEAs successfully covers operating parameters (CQs 1-3, see also Figure 4b)), as well as design parameters and PEA structure details (CQs 4-6). Certification and service details (CQs 7 and 8) are technically answerable, though the current implementation has limited expressiveness compared to the structure of quantifiable properties, which is discussed below. Integration of metadata properties from VDI 2770 enables the instance graph to answer provenance-related questions (CQs 9 and 10). CQ11 is covered through our strict policy to include labels, alternative labels, definitions, and sources for every newly introduced class and property. By validating the instance graph with our SHACL shapes we ensured that CQ12 was also successfully met.

In addition to our own CQs, we evaluated the ontology against 8 of the 11 competency questions derived by Koch et al. based on the selection of PEAs for a neutralisation process [11]. The remaining 3 CQs cover the entire modular plant, which exceeds the scope of our ontology. SPEAD performs well in answering their questions related to asset specification. For example, the question "How many centrifugal pumps does the PEA module have?" can be answered by counting the number of component nodes directly or indirectly linked to the PEA node through the *spead:hasComponent* property which are the subject of the property *spead:hasComponentClass* with the value set to the appropriate component class from the DEXPI RDL. However, the questions regarding the functional design cannot be answered, since SPEAD does not include mathematical models for the process design. This could be addressed in the future by further integrating the relevant submodels of OntoCAPE.

DISCUSSION

The SPEAD ontology sufficiently addresses CQs 1-6 by modelling the structure of PEAs and their quantifiable properties in a consistent, verifiable manner. However, in its current state, the implementation of the *spead:Capability* and *spead:SafetySpecification* classes limit the quality of the stored information. This stems from a deliberate design choice. In automation terms, the capabilities of a PEA are reflected through the implemented MTP services. While enriching the SPEAD ontology with service descriptions would improve human readability, the operational information already exists in the MTP files. The current implementation only adds metadata, avoiding the duplication of information.

We did not focus on the safety specifications, only implementing and validating them in a string-based format. In the future, especially in cases such as ATEX certifications or the IP rating of a PEA, small controlled vocabularies could be used to express safety related properties of a PEA instead.

Compared to OntoCAPE and DEXPI, SPEAD defines terminology definitions more precisely (*rdfs:comment* vs. *skos:definition* and *dcterms:source*), though all three models are capable of answering CQ11. However, neither OntoCAPE nor DEXPI can answer CQ9 and CQ10, which are crucial for traceability and trust, especially in documentation. These questions are addressed by SPEAD through its utilisation of PROV-O and its implementation of metadata properties from VDI 2770.

The evaluated CSTR and dosing PEAs in our study consisted of 358 and 214 triples respectively. Information was primarily drawn from P&IDs and technical data sheets, which were the only available documents due to the ongoing development of the PEAs. We expect a single SPEAD to be of the order of 10^4 triples, and a modular plant to be of the order of 10^5 triples.

At this scale, both reasoning and validation remain tractable. The OWL constructs used in SPEAD are intentionally conservative, avoiding complex role chains and nominals. This keeps reasoning complexity polynomial with respect to graph size, meaning that a tenfold increase in graph size does not lead to a disproportionate increase in reasoning time. SHACL validation performance is similarly predictable: most shapes are node- or property-scoped and scale linearly with the number of instances. One exception is the path-based constraint excluding transitive aggregation relationships, which could become costly for deeply nested structures. In practice, nesting depth is bounded by the four-level component-FEA-PEA-MP hierarchy defined in VDI 2776-1.

CONCLUSION & OUTLOOK

We introduced the SPEAD ontology for the documentation of Process Equipment Assemblies that collects and categorises information on PEAs and enriches this information with documentation metadata. We evaluated the ontology against competency questions that cover both content of the information model as well as terminology and provenance-related information, which we derived from an exemplary use case. We also evaluated the ontology against competency questions from the literature. The results show that the ontology works well for the considered use case of exchanging PEAs in an existing plant. Compared to existing data models, it improves the tracking of provenance related metadata.

The next steps in the development of SPEAD include

the improvement of the description of module capabilities and safety-related properties. Further, additional use cases for SPEAD will be investigated. For example, we aim to investigate how SPEAD can be used in the matching of PEAs during the plant engineering phase. SPEAD could also be used as a source of static information for process simulation models, e.g., stirrer speed or wall thickness. Further, we aim to extend the ontology to also cover the non-modular environment. This is based on the VDI 2776-1 guideline, which states that the non-modular environment can be considered as a PEA in the context of automation [2]. We extend this assumption to the documentation. The resulting "Infrastructure SPEAD" can be used to perform matching queries with the SPEADs of a pool of PEAs to select PEAs based on their compatibility with the infrastructure.

Since the creation of instance graphs can be a time-consuming task, an automated pipeline to convert existing documentation such as technical data sheets into SPEADs should also be explored.

Lastly, due to the emergence of the Asset Administration Shell (AAS) as a key digital twin technology, integration of SPEAD either as a source for semantic identifiers or as its own submodel should also be considered.

ACKNOWLEDGEMENTS

This work is conducted within the project REUNION. The authors thank all project partners and the German Federal Ministry of Economic Affairs and Energy for the funding (funding codes #03EN2100A & #03EN2100J)

AUTHOR IDENTIFIERS

Author ORCIDs:
Lamm S: 0000-0001-5912-1948
Tecl S: 0009-0004-3194-3199
Dietrich I: 0009-0008-9691-972X
Sommer S: 0009-0001-5287-2682
Heinbücher M: 0009-0003-5097-791X
Pelz PF: 0000-0002-0195-627X

REFERENCES

1. DECHEMA e.V.. Modular Plants: Flexible chemical production by modularization and standardization – status quo and future trends. (2016) https://dechema.de/dechema_media/modularplants.pdf
2. Verein Deutscher Ingenieure e.V. (VDI). VDI 2776-1, Process engineering plants – Modular plants – Fundamentals and planning modular plants. DIN Media (2020)
3. Verein Deutscher Ingenieure e.V. (VDI). VDI 2776-2, Process engineering plants – Modular plants – Designing modular plants. DIN Media (2024)
4. Verein Deutscher Ingenieure e.V. (VDI). VDI/VDE/NAMUR 2658-1, Automation engineering of modular systems in the process industry – General concept and interfaces. DIN Media (2022)
5. Verein Deutscher Ingenieure e.V. (VDI). VDI 2770, Operation of process engineering plants – Minimum requirements for digital manufacturer information for the process industry – Fundamentals. DIN Media (2020)
6. Lamm S, Wetterich P, Dietrich I, Sommer S and Pelz PF. Semantic PEA Data sheet (SPEAD) – enhancing modular plant documentation with the Asset Administration Shell (AAS). Zenodo (2025) https://doi.org/10.5281/zenodo.15187664
7. European Parliament and Council. Regulation (EU) 2023/1230 of the European Parliament and of the Council of 14 June 2023 on machinery and repealing Directive 2006/42/EC of the European Parliament and of the Council and Council Directive 73/361/EEC (Text with EEA relevance). Official Journal of the European Union (2023) http://data.europa.eu/eli/reg/2023/1230/oj
8. World Wide Web Consortium (W3C). RDF 1.1 Concepts and abstract syntax. W3C (2014) https://www.w3.org/TR/rdf11-concepts/
9. World Wide Web Consortium (W3C). OWL 2 Web Ontology Language Document Overview (Second Edition). W3C (2012) https://www.w3.org/TR/owl2-overview/
10. World Wide Web Consortium (W3C). Shapes Constraint Language (SHACL). W3C (2017) https://www.w3.org/TR/shacl/
11. Koch A, Hamedi N, Furtner L, Kock T, Klose A, Mädler J. Standards for information models considering knowledge distribution in modular plants. 2023 IEEE 21st International Conference on Industrial Informatics (INDIN) :1-7 (2023). https://doi.org/10.1109/indin51400.2023.10218218
12. Marquardt W, Morbach J, Wiesner A, Yang A. OntoCAPE – A Re-Useable ontology for Chemical Process Engineering. Springer Berlin Heidelberg (2010) https://doi.org/10.1007/978-3-642-04655-1
13. DEXPI e.V.. DEXPI Specification 2.0.0 (2025) https://dexpi.gitlab.io/-/Specification/-/jobs/11676485644/artifacts/src/.build/html/html/index.html
14. International Electrotechnical Commission (IEC). IEC 61512-1, Batch control – Part 1: Models and terminology. IEC (1997)
15. Kendall EF and McGuinness DL. Ontology Engineering. Springer International Publishing (2019) https://doi.org/10.1007/978-3-031-79486-5
16. World Wide Web Consortium (W3C). Semantic Sensor Network Ontology. W3C (2017)

https://www.w3.org/TR/vocab-ssn/
17. QUDT.org. Quantities, Units, Dimensions and Types (2014) https://doi.org/10.25504/FAIRsharing.d3pqw7
18. World Wide Web Consortium (W3C). PROV-O: The PROV Ontology. W3C (2013) https://www.w3.org/TR/prov-o/
19. World Wide Web Consortium (W3C). SKOS Simple Knowledge Organization System Reference. W3C (2009) https://www.w3.org/TR/skos-reference/
20. Dublin Core Metadata Initiative (DCMI). DCMI Metadata Terms. DCMI (2020) https://www.dublincore.org/specifications/dublin-core/dcmi-terms/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

libDIPS: An Open-Source Platform for Global Optimization of Hierarchical Optimization Problems

Adrian W. Lipow[b], Daniel Jungen[b], Aron Zingler[b], Hatim Djelassi[b] and Alexander Mitsos[abc]*

[a] JARA-CSD, 52062 Aachen, Germany
[b] Process Systems Engineering (AVT.SVT), RWTH Aachen University, 52074 Aachen, Germany
[c] Institute of Climate and Energy Systems: Energy Systems Engineering (ICE-1), Forschungszentrum Jülich GmbH, 52425 Jülich, Germany
* Corresponding Author: amitsos@alum.mit.edu.

ABSTRACT

Hierarchical optimization problems such as (generalized) semi-infinite optimization problems and bilevel problems appear in various disciplines of process systems engineering, such as flexibility analysis or parameter estimation. Adaptive discretization-based algorithms are a family of methods to solve these problems. In these methods, the original problem is decomposed into subproblems, which are solved with a standard optimization solver and then refined iteratively. Several related algorithms have been published. Until recently, computational studies have typically been performed using publication-specific implementations and benchmark problems. We recently published a software package – libDIPS – comprising existing adaptive discretization-based algorithms and a library of test problems for comparison. Several of the algorithms implemented in libDIPS exhibit strong parallelization potential in their algorithmic steps: In the algorithms of Mitsos [Optimization 60:1291-1308 (2011)] and Mitsos and Tsoukalas [J Glob Optim 61:1-17 (2015)], the upper- and lower-bounding procedures are independent. In the algorithm of Tsoukalas and Rustem [Optim Lett 5:705-716 (2011)], several objective function values can be probed independently for feasibility. However, algorithm statements in the literature typically only consider serial execution of the algorithm steps. While parallelization of the used subsolver is already possible in libDIPS, leveraging parallelization in the algorithmic steps has not been investigated until now. In this contribution, we present an overview of libDIPS and introduce our parallelized implementations of several existing algorithms. We investigate the effect of parallelization at the algorithm-level compared to parallelization within the subsolver. Furthermore, we propose two improvements to speed up the solution process of the algorithms compared to their original publications. Results are presented for the existing library of test problems collected from literature.

Keywords: Optimization, Parallelization, Numerical Methods, Semi-Infinite Programming

INTRODUCTION

Hierarchical optimization problems occur in various applications in process engineering, such as flexibility analysis [12] or thermodynamics [11]. Broadly speaking, hierarchical optimization problems include (generalized) semi-infinite optimization problems ((G)SIPs), existence-constrained semi-infinite optimization problems (ESIPs), bilevel problems and minmax problems. In hierarchical optimization problems, an upper-level problem is constrained by the solution of an embedded lower-level problem. Solving these optimization problems to global optimality is challenging especially when their lower-level problems are nonconvex, since even establishing feasibility of a point requires global optimization of the lower-level problem.

Several approaches for global optimization of hierarchical problems have been developed over the past decades, among them adaptive discretization-based methods, which decompose the original problem into several subproblems and iteratively refine them. A major advantage of adaptive discretization methods is that standard optimization solvers can be used to solve the subproblems. However, until recently, computational

studies to evaluate different methods' performance were typically done using publication-specific prototype implementations and benchmark problems, thereby limiting the possibility for fair comparisons between different methods. In previous work, we published libDIPS [7], a software of unified implementations of existing adaptive discretization-based algorithms [1-6, 8-10] and a large library of test problems collected from literature, enabling benchmarking and comparisons of the different algorithms.

In this contribution, we present several recent additions to libDIPS aimed at improving the computational performance of several solvers. We focus on solvers for SIPs and GSIPs and exploit parallelization potentials at the algorithmic level (as opposed to the subproblem level, which is already possible with appropriate subproblem solvers). Furthermore, we present two heuristics that improve some of the solvers' performance compared to their original publications.

We consider SIPs (1) and GSIPs (2):

$$\min_{\boldsymbol{x}\in X} f(\boldsymbol{x}) \tag{1a}$$
$$\text{s.t.}\ g(\boldsymbol{x},\boldsymbol{y}) \leq 0\ \forall \boldsymbol{y}\in Y \tag{1b}$$

$$\min_{\boldsymbol{x}\in X} f(\boldsymbol{x}) \tag{2a}$$
$$\text{s.t.}\ g(\boldsymbol{x},\boldsymbol{y}) \leq 0\ \forall \boldsymbol{y}\in Y(\boldsymbol{x}) \tag{2b}$$

We call $\boldsymbol{x}$ the upper-level variables and $\boldsymbol{y}$ the lower-level variables. We assume continuity of f on X and g on $X\times Y$. Further, we assume that the sets X, Y, and $Y(\boldsymbol{x})$ are compact and defined by continuous functions. These assumptions are made by the algorithms considered in the following [8, 9, 13] and are also typically required by the subsolvers. We only consider a single semi-infinite constraint ((1b) and (2b), respectively) for ease of notation, although the algorithms and our implementations can handle multiple semi-infinite constraints.

A SOFTWARE FRAMEWORK FOR SOLVING HIERARCHICAL OPTIMIZATION PROBLEMS

We briefly describe libDIPS from a user's perspective, specifically focusing on aspects of modeling and solving hierarchical optimization problems. For a detailed description of all features and implementation details, we refer to [7] and the libDIPS repository (see Digital Supplementary Material).

To model a problem in libDIPS, users must specify the subproblems as text files using the domain-specific language provided by libALE [18], which allows writing optimization problems in an intuitive fashion. There are two options for modeling the subproblems: users can formulate them in an opinionated way or use the provided template files to automatically derive the subproblems from the overall problem description. While the latter approach is arguably simpler to follow, the former allows users to exploit problem characteristics to potentially construct subproblem formulations tailored for the used subsolver. We illustrate this using an example workflow below.

As mentioned above, one advantage of adaptive discretization algorithms is that their subproblems can be solved by state-of-the-art MINLP optimization solvers. libDIPS provides interfaces to several commercial and open-source solvers, among them CPLEX [15], Gurobi [16], and our in-house solver MAiNGO [17]. For a continuously updated list of supported solvers, we refer to the libDIPS repository.

Example workflow

At a high level, the workflow for solving problems with libDIPS consists of writing the subproblems into text files, parsing these and passing the parsed problem objects to one of the solvers implemented in libDIPS, which then solves the problem. Output is written to the console and subsolver logs are written to files, enabling the user to track the overall solution progress.

Modeling in libDIPS

Assume we want to solve the following SIP using the *B&F* algorithm [1] (to be described):

$$\min_{\boldsymbol{x}\in[-2,2]^2} x_1^2 + 2x_2^3 \tag{3a}$$
$$\text{s.t.}\ \min\{x_1-y_1, x_2-y_2\} \leq 0, \forall \boldsymbol{y}\in[-1,1]^2 \tag{3b}$$

The required input files are given in Listings 1-4. Note that we show two formulations of the lower-level problem: Listing 2 uses the provided template and Listing 3 formulates the problem in an opinionated way. Besides defining the subproblems `lbp` and `llp`, a definitions file is required to mark special symbols in the subproblems, such as those related to the discretization or algorithm-specific ones.

Solving problems using libDIPS

As the name suggests, libDIPS is a library of solvers, which users can utilize in their own software tools. Furthermore, we provide standalone demos for all solvers, which demonstrate how to call the solvers from a C++ workflow. The demo executables can also be used to solve any problem defined in the format specified above. The default subsolver for the subproblems is MAiNGO; users can easily change this to their preferred solver in the build configuration file.

IMPROVING THE PERFORMANCE OF (G)SIP ALGORITHMS

The original publications of all algorithms

```
definitions:
real[2] x in [-2,2];
set{index} lbp_k_disc := {};
real[0,2] lbp_y_disc := 0;
real f(real[2] x) := x[1]^2+2*x[2]^3;
real g(real[2] x, real[2] y) := min(x[1]-y[1], x[2]-y[2]);
objective:
f(x);
constraints:
forall k in lbp_k_disc:
    g(x, lbp_y_disc[k]) <= 0;
```

Listing 1: The file `lbp.txt` for the example problem (3). Function symbols `f` and `g` are used to enable templated derivation of the subproblems.

```
definitions:
real[2] y in [-1,1];
objective:
-g(x,y);
```

Listing 2: The file `llp.txt` for the lower-level problem of the example problem (3), using the provided template. For this problem, this results in a nonlinear optimization problem due to the min-operator. Note that subproblems are always written as minimization problems.

```
definitions:
real[2] y in [-1,1];
real t in [-3,3];
objective:
-t
constraints:
x[1]-y[1] >= t;
x[2]-y[2] >= t;
```

Listing 3: The file `llp.txt` for the lower-level problem of the example problem (3), using an opinionated formulation. For this problem, this results in a linear optimization problem.

```
programdefinitions("lbp"):
set_disc(1) = lbp_k_disc;
set_disc_parameter_src(1, lbp_y_disc) = y;
```

Listing 4: The file `def.txt` for the example problem (3), defining `lbp_k_disc` as the set of discretization indices and `lbp_y_disc` as the discretized set Y^{LBP}.

implemented in libDIPS only consider serial execution of the algorithmic steps. Multi-core architectures can still be exploited by parallelism in the subsolver called to solve the subproblems. However, several of the algorithms offer potential for parallelization on their own. Here, we consider the *RRHS* algorithm [8], the *GSIP-RRHS* algorithm [9] and the *Oracle* algorithm [13]. Furthermore, we propose changes to the algorithmic steps to improve the algorithms' performance.

Parallel bounding in *RRHS* and *GSIP-RRHS*

RRHS [8] and *GSIP-RRHS* [9] combine the lower-bounding procedure from the *B&F* algorithm [1] with an upper-bounding procedure based on restricting the

right-hand-side of the semi-infinite constraint. We briefly state the main ideas of the algorithm here and refer to the original publications for full details. The two subproblems are given for an SIP in the following:

$$\min_{\boldsymbol{x}\in X} f(\boldsymbol{x}) \tag{4a}$$
$$\text{s.t.}\, g(\boldsymbol{x}, \boldsymbol{y}^k) \leq 0 \;\forall \boldsymbol{y}^k \in Y^{\text{LBD}} \tag{4b}$$

$$\min_{\boldsymbol{x}\in X} f(\boldsymbol{x}) \tag{5a}$$
$$\text{s.t.}\, g(\boldsymbol{x}, \boldsymbol{y}^k) \leq -\varepsilon \;\forall \boldsymbol{y}^k \in Y^{\text{UBD}}, \tag{5b}$$

Where $\varepsilon > 0$, and $Y^{\text{UBD}}, Y^{\text{LBD}}$ constitute sets of finite cardinality, i.e., discrete subsets of the infinite set Y. Problem (4) is the lower-bounding problem. Problem (5) is the upper-bounding problem, which is generally neither a relaxation nor a restriction and may not yield an upper bound in every iteration. Note that the two problems use separate discretizations, Y^{LBD} and Y^{UBD} (a shared discretization is possible, but this leads to larger subproblems and may offset any performance gains seen by reducing the number of iterations). Thus, while the original publications alternate between solving the upper- and lower-bounding subproblems, these can be solved in parallel, as subsequent iterations are not dependent on information from the respective other procedure.

Candidate points $\bar{\boldsymbol{x}}$ computed from problems (4) and (5) are checked for feasibility by solving the lower-level problem:

$$g^* = \max_{\boldsymbol{y}\in Y} g(\bar{\boldsymbol{x}}, \boldsymbol{y}) \tag{6}$$

If g^* is non-positive, the candidate point $\bar{\boldsymbol{x}}$ is SIP-feasible, otherwise, it is infeasible. In the latter case, the associated value $\boldsymbol{y}^*$ is used to refine the discretization of the corresponding upper-level problem (i.e., if $\bar{\boldsymbol{x}}$ was generated by (4), Y^{LBD} would be updated to $Y^{\text{LBD}} \cup \{\boldsymbol{y}^*\}$). Note that problems (4) and (6) are the subproblems used in the *B&F* algorithm, which only computes lower bounds.

Besides the obvious benefit of potentially being able to solve more subproblems in a fixed amount of time, decoupling the upper- and lower-bounding procedures has another advantage: It enables the algorithm to progress faster when one procedure has already found a "good" bound. The sequential implementation of *RRHS* and *GSIP-RRHS* alternates between solving problems (4) and (5). If one of the procedures quickly finds a "good" bound and stops progressing, subsequent iterations of that procedure are still performed but do not contribute to convergence of the algorithm. By decoupling the two procedures, the optimality gap may be closed faster because neither procedure needs to wait for an iteration of the other to complete before continuing. We test this claim using the library of test problems in libDIPS.

Early refining of the upper-bounding discretization in *RRHS* and *GSIP-RRHS*

For *RRHS*, we already presented a guard heuristic in [7], in which the upper-bounding procedure is skipped until the lower-bounding procedure produces ε^a-SIP-feasible points with $\varepsilon^a > 0$. Whenever the upper-bounding procedure is skipped, the upper-bounding discretization is populated with the discretization point obtained from the lower-bounding procedure. The aim of the guard is to quickly refine the upper-bounding discretization in early iterations, when the upper-bounding procedure is unlikely to find a feasible point. For *GSIP-RRHS*, this approach is analogously applicable. However, the guard heuristic only makes sense when the algorithms are executed in serial fashion. In the parallel implementations, such a guard cannot be directly applied, as the upper- and lower-bounding procedures are run in parallel, so skipping the upper-bounding procedure does not make sense.

In the parallel versions of *RRHS* and *GSIP-RRHS*, we propose to achieve this aim in a similar fashion. Whenever the lower-bounding procedure furnishes a new discretization point, we check if a valid upper bound has been found. If this is not the case, we add the discretization point to the upper-bounding discretization.

Parallel objective probing in *Oracle*

The *Oracle* algorithm [13] adapts the algorithm from [1] to obtain both lower bounds and feasible points. It does so by solving the oracle problem (7) to determine whether an objective function value f^{target} is attainable.

$$\min_{\boldsymbol{x}\in X} \max\{f(\boldsymbol{x}) - f^{\text{target}}, \max_{\boldsymbol{y}^k\in Y^{\text{ORA}}} g(\boldsymbol{x}, \boldsymbol{y}^k)\}, \tag{7}$$

Where Y^{ORA} is a set of finite cardinality If f^{target} is not attainable with a given discretization Y^{ORA}, it is a lower bound to the optimal objective value. On the other hand, if the target is attainable in the discretized subproblem and the lower-level problem proves the candidate point SIP-feasible, an upper bound is obtained. In the latter case, the lower-level problem may also declare the candidate point SIP-infeasible, in which case the discretization is refined, and the oracle problem is re-solved for the same target value. After successfully finding a new upper or lower bound, the target value is updated (in the original publication by bisection) and the oracle problem is solved again, starting with the discretization of the last iteration.

The *Oracle* algorithm offers great potential for parallelization similarly to runahead computing [14], where additional threads operate a few steps ahead of the main thread to speculate on the future state. In *Oracle*, this means that multiple objective value targets can be probed in parallel. For example, rather than just probing the midpoint between the current lower and upper bound, two additional threads can be used to probe the

upper and lower quarter of the interval as well, thereby effectively computing the next iteration of the algorithm in parallel. This idea can be generalized to arbitrarily many targets, assuming sufficient computing capabilities. Whenever a thread computes a new bound, any threads probing targets outside the remaining interval are aborted and assigned new valid targets.

Using multiple threads to solve multiple instances of problem (7) for different values of f^{target} raises the question of how to deal with the discretization Y^{ORA}. On one hand, Y^{ORA} should not grow unnecessarily large to avoid generating overly large subproblems. On the other hand, discretization points generated by one thread for a particular value of f^{target} may also be relevant for other target values. Thus, we present three strategies for handling the discretization. First, it is possible for each thread to have its own discretization, which is never synchronized with other threads (referred to as strategy "Off" in our numerical experiments later). This results in discretizations of minimal size in the sense that every thread only includes those points in its discretization which have been proven to be necessary in at least one iteration. However, this approach may lead to multiple threads computing the same discretization points, potentially increasing the total number of subproblems solved.

Second, it is possible for the threads to share a single common discretization (strategy "All"). Whenever the solution of a lower-level problem in a thread adds a discretization point, all other threads also use that point in subsequent iterations. While this leads to the subproblems becoming tighter and potentially requiring fewer iterations, it may also yield unnecessarily fine discretizations, increasing the difficulty of solving individual subproblems.

A third option is to combine the previous two approaches. Each thread uses its own discretization, but whenever a thread produces an update to the lower or upper bound, all other discretizations are updated to include the discretization points of the successful thread (strategy "Opt"). This approach attempts to combine the best aspects of the other two approaches: sharing discretization points that have proved useful in computing tighter bounds without producing unnecessarily fine discretizations. We test all three approaches using the library of test problems in libDIPS to confirm our claims.

A secant-based heuristic for updating the target value of *Oracle*

In the original publication of *Oracle* [13], new objective function target values are selected through bisection of the objective space. While this approach is simple to implement, it may lead to situations where a very good upper or lower bound is found quickly, but many iterations are required to prove (approximate) optimality or find a "good" feasible point.

To improve the convergence of the algorithm, we attempt to improve on the simple bisection of the objective space when selecting a new target objective value. In particular, we want to exploit iterates that lie near the boundary of the semi-infinite constraint. Once the algorithm has furnished an upper or lower bound with a sufficiently small absolute value of the semi-infinite constraint, we propose to probe the objective space in the vicinity of this bound as follows:

$$f^{\text{target}} = LBD + \frac{-g^{\text{LBD}}}{g^{\text{UBD}} - g^{\text{LBD}}}(UBD - LBD) \quad , \tag{8}$$

Where $g^{\text{LBD}} > 0$ and $g^{\text{UBD}} < 0$ are the violations of the semi-infinite constraint at the solutions generating the lower and upper bound, respectively, and LBD and UBD refer to the current lower and upper bounds on the objective value. The intuition behind (8) is to estimate a root of the semi-infinite constraint using the secant method, and that this approximation provides a valid proxy for the optimal objective function value, i.e., we expect the ratio $\frac{-g^{\text{UBD}}}{g^{\text{LBD}}}$ to be a good approximation of $\frac{UBD - f^*}{f^* - LBD}$, where f^* denotes the optimal objective value. In a neighborhood around the optimum, this is typically a valid assumption and can lead to significant reductions in the remaining objective space compared to bisection. We only apply the heuristic once a threshold g^{t} for the absolute value of the semi-infinite constraint is reached, i.e., $\min\{g^{\text{LBD}}, -g^{\text{UBD}}\} < g^{\text{t}}$, as we do not necessarily expect the linear approximation from the constraint space to the objective space to be sufficiently accurate far away from the minimum due to non-convexity. We present numerical experiments to determine this threshold below.

NUMERICAL RESULTS

We performed numerical experiments using the library of test problems collected in the libDIPS repository, which contains 336 SIPs and 86 GSIPs. Further details on the test set can be found in [7].

All computations were performed on the RWTH High-Performance Computing cluster CLAIX-2023. All subproblems were solved with MAiNGO v0.9.1, configured to use CPLEX v22.1.2 to solve (mixed-integer) linear or quadratic optimization problems. Additionally, multi-starting in MAiNGO was performed using SLSQP, implemented as part of NLopt [19]. In some of the following experiments, we configured MAiNGO to use multiple threads. In those runs, MAiNGO set the number of threads used by CPLEX correspondingly. All other settings were left unchanged compared to those used for benchmarks in [7]. As in [7], we rounded all runtimes below 1 second (s) to 1 s, as we compare the solver configurations relative to each other and cannot fairly differentiate times below 1 s. Due to the increased memory requirement of the multithreaded implementation, 18 large

SIP instances were not tractable (despite being tractable but unsolved in [7]) and are omitted in the results below. This omission increases the relative percentage of problems solved compared to [7], however, the absolute numbers of our serial implementation are consistent with the results reported there.

Parallelization of *RRHS* and *GSIP-RRHS*

We first investigated the performance of our parallelized implementations of *RRHS* and *GSIP-RRHS* using several parameter configurations. The results in [7] already showed that the guard heuristic performs favorably for a serial implementation of the algorithm. We now compare this to our parallel implementation, both with and without the heuristic enabled. For the parallel implementation, we ran the upper- and lower-bounding procedures using a single thread each. We also re-ran the serial implementation using two threads for the subsolver to enable a fair comparison and demonstrate the difference between parallelizing the algorithmic steps and parallelizing in the subsolver.

Figure 1 shows the performance profile for all configurations of *RRHS* on the SIP problem set. The multithreaded implementation without the early refinement of the upper-bounding problem's discretization performs best, solving problems faster than the other configurations and solving a larger fraction of problems overall. Surprisingly, the early refinement does not carry its advantage from the serial case to the parallel case. We suspect that this is due to the extra discretization points not being as useful in the parallel case, as the upper-bounding problem is solved from the start. Thus, the upper-bounding problem's discretization may become unnecessarily dense in this case, leading to harder subproblems.

For unsolved instances, we define the remaining optimality gap as $\min\{UBD - LBD, \frac{UBD-LBD}{UBD}\}$. In Figure 2, we plot the percentage of unsolved problems over the remaining optimality gap for each configuration. The multithreaded implementation achieves better remaining gaps than the serial implementation, indicating that parallelizing the algorithmic steps of *RRHS* is indeed more efficient than using all available threads for the subsolver.

We show the performance profile for GSIP-*RRHS* applied to the GSIP problem set in Figure 3. Again, parallelism at the algorithmic level pays off, as the multithreaded implementations are faster and solve a slightly higher percentage of the problem set. In contrast to the SIP case, extra refinement of the upper-bounding procedure does not seem to impact the performance of the multithreaded implementation but harms the serial one.

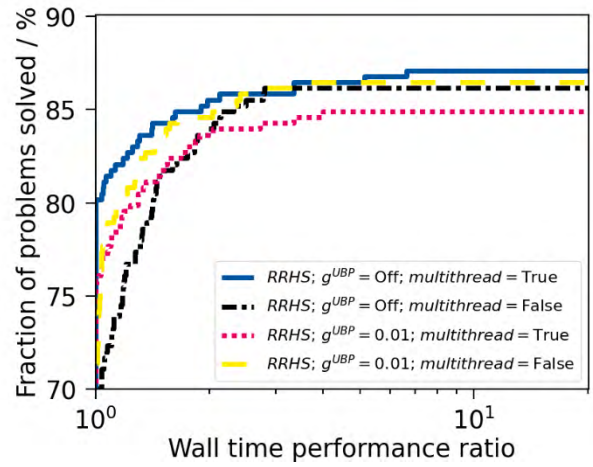

Figure 1. Dolan-Moré plot of wall time performance ratio for *RRHS* on the SIP problem test set.

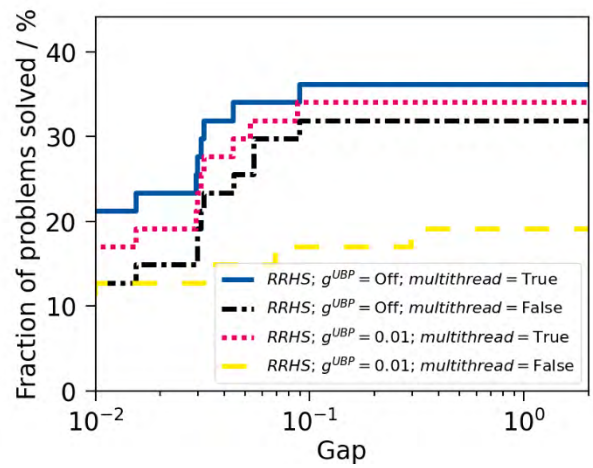

Figure 2. Fraction of SIP instances that are not solved by one of the *RRHS* configurations but would have been considered solved with a larger optimality tolerance.

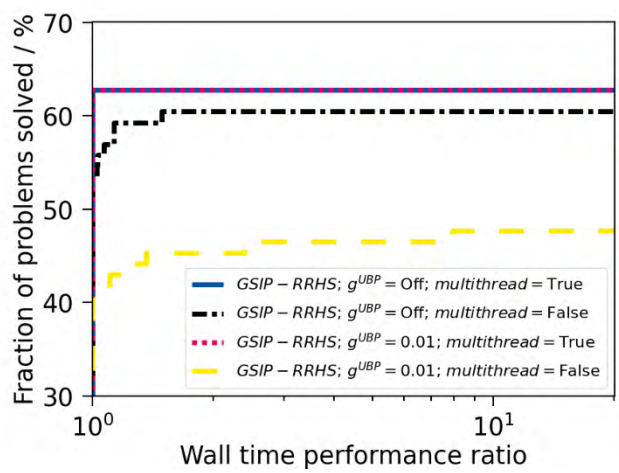

Figure 3. Dolan-Moré plot of wall time performance ratio for GSIP-*RRHS* on the GSIP problem test set.

Secant-based heuristic in *Oracle*

To determine a good value for the threshold hyperparameter g^t, we solved the benchmark set of SIPs using the serial implementation of *Oracle* while varying $g^t \in \{0, 0.01, 0.1, 0.5, 1, 1000\}$. Note that setting $g^t = 0$ disables the heuristic and setting $g^t = 1000$ effectively enables it as soon as a feasible upper bound is found. We evaluated each setting based on (rounded) runtime (RT) and the total number of subproblems (NSP) solved. We present the results in Table 1. Note that in many cases, multiple settings yielded equal numbers of subproblems or (rounded) runtimes. The number of unique winners is lower for both metrics.

The results indicate that always using the secant heuristic for target selection is superior in terms of the number of subproblems solved in many cases, while there are some problems in the test set that benefit from setting the threshold g^t to a small value or completely disabling the heuristic. In terms of runtime, all settings performed similarly due to most of the problems being solvable in under 1 s. Nevertheless, the number of unique winners exhibits qualitatively similar behavior for runtime as for the number of subproblems.

We found that enabling the heuristic resulted in up to six problems (depending on the value of g^t) of the test set not being solved that were solved without the heuristic. There were two reasons for this: in some cases, the heuristic was exact and proposed a target value equal to the minimum, in which case finite termination is not guaranteed [13]. In other cases, the initially computed upper bound was feasible and triggered the heuristic, but the ratio $\frac{-g^{\mathrm{UBD}}}{g^{\mathrm{LBD}}}$ was significantly smaller than $\frac{UBD-f^*}{f^*-LBD}$, leading to target updates close to the upper bound. In combination with a very large upper bound, this resulted in the algorithm failing to converge within the iteration limit. On the other hand, there were two instances that were solved with some settings of the heuristic but were unsolved without it.

Table 1. Statistics for the secant heuristic over the entire benchmark set. The number of unique winners is given in parentheses.

Threshold g^t	# lowest NSP	# lowest RT
0	99 (11)	217 (11)
0.01	125 (15)	228 (19)
0.1	143 (19)	220 (12)
0.5	132 (0)	215 (6)
1	139 (3)	224 (23)
1000	214 (95)	228 (28)

Parallelization of *Oracle*

We investigated multiple configurations of *Oracle* to determine the benefit of parallelism. Based on the results of the previous section, we always enabled the secant heuristic as it usually required fewer subproblems to be solved. As a baseline, we used the serial implementation, giving 16 threads to the subsolver. We then compared this to two levels of parallelism: probing four targets in parallel with four threads per subsolver and probing 16 targets in parallel with one thread per subsolver. For both configurations involving parallelism at the algorithmic level, we investigated all three strategies for the discretization as mentioned earlier.

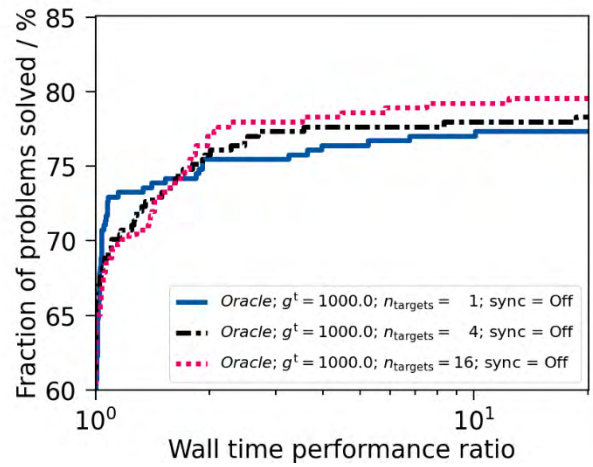

Figure 4. Dolan-Moré plot of wall time performance ratio for different parallelization strategies in *Oracle*.

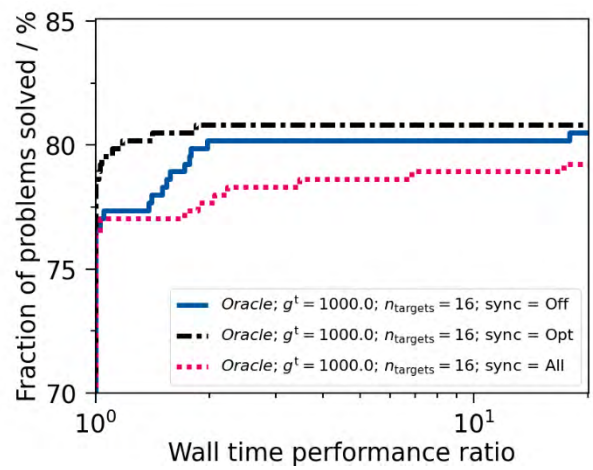

Figure 5. Wall time performance ratio for different discretization synchronization strategies in *Oracle*.

We plot the wall time performance ratio for all three parallelization strategies with no synchronization of discretizations in Figure 4. As for *RRHS*, parallelization at the algorithmic level leads to a larger percentage of the test set being solved. At small time ratios, the plot shows that the serial strategy outperforms the parallel strategies. We attribute this to the lack of a proper method to abort the optimization within MAiNGO. Thus, the parallel implementations may achieve optimality faster than their reported runtimes, as they must wait for all running

MAiNGO instances to complete before terminating. In some rare instances, this led to an estimated performance degradation by a factor of 10 or more. We aim to add an abort-functionality to MAiNGO in a future release and test this hypothesis.

In Figure 5, we plot the performance ratios of different discretization synchronization strategies when probing 16 targets in parallel. Strategy "Opt" performs best, resulting in faster solutions and solving the largest fraction of the problem set. Synchronizing every discretization point (strategy "All") performs worst, indicating that the extra information gained from more discretization points is not worth the extra computational effort involved with solving larger subproblems. Letting each thread use its own discretization with no synchronization (strategy "Off") performs in between the other two strategies, achieving similar performance to strategy "All" for fast problems (likely due to computing the same discretization points across multiple threads), but solving nearly the same fraction of the problem set overall as strategy "Opt" (likely due to the better performance of later subproblems with smaller discretizations).

CONCLUSION

We presented improved implementations of some of the algorithms for hierarchical optimization problems in libDIPS. We exploited parallelism in the algorithmic steps and demonstrated that this is a more efficient use of available threads than parallelizing only at the subsolver level, enabling faster solutions and increasing the percentage of solved problems. This observation holds for all investigated algorithms. However, parallelism increases memory usage, which led to some instances becoming intractable. Additionally, we presented a promising strategy for synchronizing the discretization used in the multithreaded implementation of the *Oracle* algorithm. The open-source nature of libDIPS enables others to implement and benchmark their algorithms.

We plan to extend parallelization in libDIPS in future releases. This includes investigating parallelization in the steps of other implemented algorithms as well as improving the existing parallelization, such as also enabling the solution of multiple lower-level problems in parallel.

DIGITAL SUPPLEMENTARY MATERIAL

libDIPS is freely available at https://git.rwth-aachen.de/avt-svt/public/libdips.

ACKNOWLEDGEMENTS

Computations were performed with computing resources granted by RWTH Aachen University. This research is funded by Réseau de transport d'electricité (RTE, France) through the project "Hierarchical Optimization for Worst-Case Analysis of Power Grids."

AUTHOR IDENTIFIERS

Author ORCIDs:
Lipow AW: 0000-0002-5860-5323
Jungen D: 0000-0002-1802-1884
Zingler A: 0000-0003-1135-0236
Djelassi H: N/A
Mitsos A: 0000-0003-0335-6566

REFERENCES

1. Blankenship JW, Falk JE. Infinitely constrained optimization problems. J Optim Theory Appl 19:261-281 (1976). https://doi.org/10.1007/bf00934096
2. Djelassi H, Mitsos A. A hybrid discretization algorithm with guaranteed feasibility for the global solution of semi-infinite programs. J Glob Optim 68:227-253 (2016). https://doi.org/10.1007/s10898-016-0476-7
3. Djelassi H, Mitsos A. Global solution of semi-infinite programs with existence constraints. J Optim Theory Appl 188:863-881 (2021). https://doi.org/10.1007/s10957-021-01813-2
4. Djelassi H, Glass M, Mitsos A. Discretization-based algorithms for generalized semi-infinite and bilevel programs with coupling equality constraints. J Glob Optim 75:341-392 (2019). https://doi.org/10.1007/s10898-019-00764-3
5. Falk JE, Hoffman K. A nonconvex max?min problem. Naval Research Logistics 24:441-450 (2006). https://doi.org/10.1002/nav.3800240307
6. Jungen D, Mitsos A. An improved oracle adaption for bilevel programs. In: Computer Aided Chemical Engineering. Ed: Manenti F, Reklaitis GV. Elsevier (2024).
7. Jungen D, Zingler A, Djelassi H, Mitsos A. libDIPS – discretization-based semi-infinite and bilevel programming solvers. *Math Prog Comp* (in press).
8. Mitsos A. Global optimization of semi-infinite programs via restriction of the right-hand side. Optimization 60:1291-1308 (2011). https://doi.org/10.1080/02331934.2010.527970
9. Mitsos A, Tsoukalas A. Global optimization of generalized semi-infinite programs via restriction of the right hand side. J Glob Optim 61:1-17 (2014). https://doi.org/10.1007/s10898-014-0146-6
10. Mitsos A, Lemonidis P, Barton PI. Global solution of bilevel programs with a nonconvex inner program. J Glob Optim 42:475-513 (2007). https://doi.org/10.1007/s10898-007-9260-z

11. Mitsos A, Bollas GM, Barton PI. Bilevel optimization formulation for parameter estimation in liquid–liquid phase equilibrium problems. Chemical Engineering Science 64:548-559 (2009). https://doi.org/10.1016/j.ces.2008.09.034
12. Swaney RE, Grossmann IE. An index for operational flexibility in chemical process design. part I: formulation and theory. AIChE Journal 31:621-630 (2004). https://doi.org/10.1002/aic.690310412
13. Tsoukalas A, Rustem B. A feasible point adaptation of the blankenship and falk algorithm for semi-infinite programming. Optim Lett 5:705-716 (2010). https://doi.org/10.1007/s11590-010-0236-4
14. Bakshalipour M, Sarbazi-Azad H. Parallelizing bisection root-finding: a case for accelerating serial algorithms in multicore substrates. https://arxiv.org/abs/1805.07269
15. International Business Machines Corporation. IBM ILOG CPLEX v 22.1.2.
16. Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual. https://www.gurobi.com
17. Bongartz D, Najman J, Sass S, Mitsos A. MAiNGO – McCormick-based Algorithm for mixed-integer Nonlinear Global Optimization. http://permalink.avt.rwth-aachen.de/?id=729717
18. Zingler A, Jungen D, Djelassi H, Mitsos A. libALE – a library for algebraic-logical expression trees. https://git.rwth-aachen.de/avt-svt/public/libale
19. Johnson SG. The NLopt nonlinear-optimization package. https://github.com/stevengj/nlopt

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Model verification and Uncertainty Quantification methods using the CCSI simulation model for CO_2 capture

Jessica V. Scheffer[ab*], Serena Delgado[a], Olivier Authier[a], Valentin Loubière[a], Franchine Ni[a], Christophe Castel[b] and Jean-Marc Commenge[b]

[a] EDF R&D, Chatou, France
[b] Université de Lorraine, CNRS, LRGP, F-54000 Nancy, France
* Corresponding Author: jessica.vaz-scheffer@edf.fr.

ABSTRACT

This work aims at verifying the CO_2 absorption capture model using monoethanolamine (MEA) solvent developed by the U.S. DOE's Carbon Capture Simulation Initiative (CCSI) and performing uncertainty propagation of mass transfer, liquid hold-up and reaction kinetics properties in the complete model, which includes absorber and stripper columns. The verification of the Aspen Plus CCSI model, based on pilot plant data from the National Carbon Capture Center (NCCC) for a CO_2 flue gas concentration between 7 and 11% (mol) allowed uncertainty quantification (UQ) analysis for four different selected operational points using Monte Carlo Simulation (MCS), where low liquid mass transfer parameters exhibited an impact on calculation convergence. Gaussian Processes (GP) surrogate model was implemented, followed by a sensitivity analysis in order to correlate the most sensitive parameters with studied outputs.

Keywords: Carbon Capture, Simulation, Aspen Plus, Absorption, Surrogate Model

INTRODUCTION

Carbon capture, utilization and storage (CCUS) is emerging as a technology to support the energy transition toward net-zero emissions by 2050, to limit the rise in global temperatures as close as possible to 1.5°C. According to the International Energy Agency (IEA) Net Zero Emission (NZE) scenario, CCUS technologies play some critical roles in the transition to net zero, especially considering the difficulty of eliminating emissions in hard-to-abate industrial sectors, aviation and maritime shipping. Since carbon capture technologies can also be applied to already existing energy assets, it has been considered as part of the energy transition strategies, with a post-combustion application of CCS technologies (targeting 95% capture) from fossil-fuel thermal power plants.

The most mature technique for CO_2 capture is chemical absorption with amine-based aqueous solutions as absorbents. A steady-state open-source CO_2 capture model with 30 wt.% aqueous monoethanolamine (MEA) as the solvent was developed and validated by the U.S. DOE's Carbon Capture Simulation Initiative (CCSI) based on a pilot plant at National Carbon Capture Center (NCCC) [3, 4]. The flowsheet (v3.2.1) of the NCCC pilot plant MEA Steady State Model developed by the U.S. DOE's CCSI and modelled in the process modelling software Aspen Plus is represented in Figure 1. The open-source code is available on the release page [5].

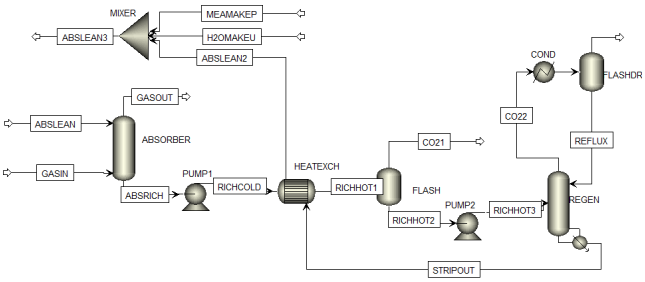

Figure 1. CCSI model flowsheet based on NCCC pilot plant, available in [5].

CO_2 is captured through the absorber column and recovered in the flash unit and stripper column. A design specification is used in Aspen Plus (rate-based columns using RateSep) to adjust the reboiler duty based on lean CO_2 loading setpoint. An economiser heat exchanger is

also included between the lean solvent stripper output and rich solvent absorber output for energy integration.

Uncertainty Quantification (UQ) is a method used to determine the confidence level of a numerical model, outlining the most sensitive sources (process variables or model parameters) and predict outlet variabilities, considering the uncertainty carried by those sources. UQ analysis requires having first verified the model for physical consistency. To allow the UQ analysis in the CCSI CO_2 capture model, a set of operational data from NCCC pilot plant from 2014 and 2017 was used for model verification. The %CO_2 capture, rich CO_2 loading, reboiler duty and rich solvent flow rate were compared to data points to attest model accuracy. Once the model is verified, uncertainty quantification implementation requires the definition of: (1) uncertain sources and their respective distributions and (2) output variables and their quantities of interest (e.g., mean, median, dispersion parameters).

The CCSI initiative developed a complete CO_2 absorption process model [3, 4] using UQ analysis while other recent studies such as [8] compared different UQ analysis techniques for an absorber model under one operational condition. This work complements the CCSI uncertainty propagation over larger ranges of liquid and gas flow rates to obtain %CO_2 capture and reboiler duty probability distributions. Mass-transfer, hold-up and reaction kinetics properties are used as uncertain sources to obtain %CO_2 capture and reboiler duty probability distributions. Then, sensitivity analysis is performed to correlate the most sensitive parameters with studied outputs. While the CCSI initiative obtained the input probability distributions by parameter regression, in this work they are built based on literature data.

METHODOLOGY

Uncertainty Quantification Techniques

In general, UQ consists of identifying uncertain parameters and propagating their uncertainty to the numerical model. Different routes exist for UQ. The most straightforward way uses a rigorous model (in the case of this work, the Aspen Plus model) to propagate uncertainties. Initially, a sample from the probability distributions of the chosen uncertain sources is generated. For this work the chosen sampling method is the Monte Carlo Simulation (MCS) where samples are randomly generated according to the defined set of sources probability distributions. Because MCS (as well as other sampling methods) requires a large number of simulations, UQ can be propagated through surrogate models, which approximate the rigorous models with lower complexity and computation cost, but do not represent chemical or physical phenomena. MCS and other computations can then be performed at a lower cost with the surrogate model instead of the rigorous model, for example to perform sensitivity analysis. In this work, Gaussian Processes (GP) surrogate model was implemented in order to obtain 1st and total Sobol indices.

Uncertain properties and selected outputs distributions

Mass-transfer properties in both absorber and stripper columns (including mass-transfer coefficients and interfacial area), liquid hold-up and reaction kinetics are chosen as uncertainty sources considering previous works on UQ analysis of CO_2 chemical absorption process models [3, 4, 8]. The probability distributions were defined according to a literature review of the variability of selected properties or parameters.

Because experimental data are primarily based on measured properties rather than model parameters (except for the reaction rate, for which dedicated experiments were conducted to estimate the rate constant) this work focuses on propagating uncertainties associated with properties rather than parameters. However, in order to perform an uncertainty analysis under different operating conditions, parameters that are directly proportional to these properties were selected to represent the property uncertainty.

The selected uncertain sources, their relative Absolute Average Deviation (AAD), relative standard deviation (σ) or coefficient of variation (CV), the name of the parameters used to propagate property uncertainty, their mean value and standard deviation are presented in Table 1. The relative AAD is presented in Equation 1 where N is the number of experimental points. The relative AAD is defined as the average of normalised absolute deviations of literature models compared with experimental values. The CV is the normalised standard deviation of the data points with respect to the mean value.

$$relative\ AAD(\%) = \frac{1}{N}\sum\frac{|y_{exp}-y_{model}|}{y_{exp}} \quad (1)$$

For the liquid and gas mass-transfer coefficients and interfacial area, the relative standard deviation is obtained from the relative AAD considering a normal distribution and the Mean-Absolute Deviation model, as discussed in [15]. For liquid hold-up, the AAD and standard deviation were directly obtained using experimental points from [10].

Regarding reaction kinetics, although the rate constants were directly measured, they differ from those implemented in the CCSI model. Usually, reaction rate equations are measured as functions of components concentrations. However, the CCSI model calculates kinetics from components activities, which changes rate constant dimensions. To maintain consistency, the coefficient of variation obtained from the literature rate constants is used to calculate the standard deviation of the CCSI model rate constants.

Table 1: Sources of uncertainty and probabilities for uncertainty analysis on the CCSI model

Measured Source	Source	Relative AAD (%)	Relative σ (%)	CV	Parameter	Mean	Standard deviation
Liquid mass-transfer coefficient (k_L)	[9]	24	30.08	-	C_L	0.2033	0.0612
Gas mass-transfer coefficient(k_G)	[9]	21	26.32	-	C_G	0.3522	0.0927
Interfacial area (a_e)	[9]	8.9	11.15	-	A_1	1.4225	0.1586
Hold-up (h_L)	[10]	14.3	17.22	-	H_{L1}	11.45	1.9725
Kinetic coefficient for Reaction 1 (k_{F1})	[11,12]	-	-	0.22	k_{F1}	$2.51\ 10^9$	$5.54\ 10^8$
Kinetic coefficient for Reaction 2 (k_{F2})	[13,14]	-	-	0.17	k_{F2}	$2.29\ 10^4$	$3.85\ 10^3$

Reactions 1 and 2 presented below are used in the CCSI model [3, 4]. Reaction 2 can also be written as a combination of Reactions 3 and 4. Due to a lack of experimental data for Reaction 2, Reactions 3 and 4 were used for literature review. Since Reaction 4 is considered to be at equilibrium, Reaction 3 is therefore rate-limiting. Its coefficient of variation is used to predict the standard variation of k_{F2}.

$$2MEA + CO_2 \leftrightarrow MEAH^+ + MEACOO^- \quad (1)$$

$$MEA + CO_2 + H_2O \leftrightarrow MEAH^+ + HCO_3^- \quad (2)$$

$$CO_2 + OH^- \leftrightarrow HCO_3^- \quad (3)$$

$$MEA + H_2O \leftrightarrow MEAH^+ + OH^- \quad (4)$$

The mean and standard deviation of obtained probability distributions are then reported in Table 1. The standard deviations are the product of the AAD or the CV (in the case of rate constants) and the parameters values from the CCSI model [4], used as the mean value. All probability distributions are considered as gaussian distributions.

RESULTS AND DISCUSSION

Model verification

For the purpose of uncertainty analysis, the CCSI model was verified using NCCC pilot plant data. The variation range of NCCC data is presented in Table 2. The open-source data were reported in [6, 7]. Figure 2 compares pilot plant data and CCSI model simulations for different outputs over the operational range reported in Table 2. The %CO_2 capture is obtained through the gas mass balance in the absorber, which is equivalent to the CO_2 produced in the flash and at the top of the stripper. The data was classified by the use of intercooling at two different heights of the absorber and the number of activated absorber sections, which changes packing height in the simulation model. Out of 52 simulation runs, 44 successfully met the convergence criteria, providing a statistically significant dataset for this study.

Table 2: Variation range of NCCC pilot plant data

Inlet conditions	Range
Lean solvent flow rate (kg/h)	3016-11794
Lean loading (mol CO_2/ mol MEA)	0.065-0.419
Flue gas flow rate (kg/hr)	1000-2997
%CO_2 vol in flue Gas	0.067-0.115
Absorber active sections	1-3

A good agreement between model predictions and reference data provides an indication of the quality of the overall model and, by extension, of the mass transfer model. Nevertheless, higher deviation is observed at lower CO_2 capture rates.

Since the lean solvent loading is an input variable and the CO_2 capture rate is well represented, it is expected that the rich solvent loading should also be accurately reproduced. As, in addition, solvent flow rate does not vary significantly throughout the absorber and the flash unit, the rich solvent flow rate presents a good model prediction for a wide range of flow rates.

Among the output variables, the reboiler duty is the less accurately represented. This variable is more strongly influenced by thermodynamic modelling within the stripper, where only equilibrium reactions are considered. After verifying the model with the operational conditions in Table 2, the uncertain parameters in Table 1 were propagated to obtain the %CO_2 capture and reboiler duty probability distributions.

Uncertainty Quantification analysis

For the Monte Carlo Simulation (MCS), four different sets of operating conditions, representative of conditions reported in [6, 7] and showed in Table 3, were chosen to propagate the mass-transfer, reaction kinetics and hold-up probability distributions through the simulation model.

The selected cases, K1 and K13, differ primarily from each other due to the presence of intercooling, as well as

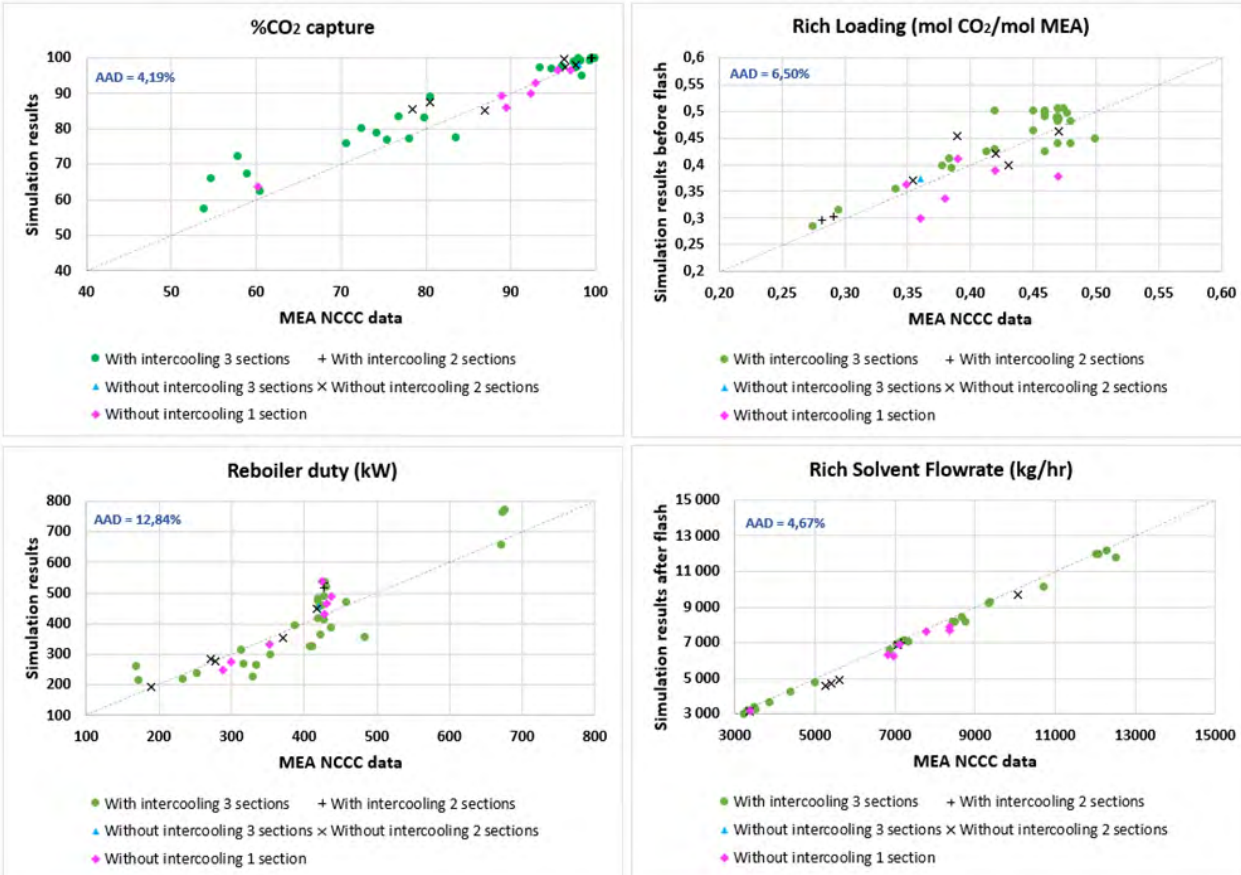

Figure 2: CCSI model verification using NCCC data pilot plant with respective relative AAD for the following outputs: %CO_2 capture, rich loading, reboiler duty and rich solvent flowrate.

slight variations in other operational variables. In contrast, the last two cases, 12A and 2D, differ mainly from the first pair because of the lower ratio of lean solvent flow rate to flue gas flow rate. Additional distinctions include minor variations in other operational variables, the presence or absence of intercooling, and, in the case of 2D, a different number of active sections.

A sample consisting of 300 points was generated for each operational condition using Persalys, a graphic interface of OpenTURNS, an open-source software developed by EDF R&D and Phimeca Engineering to facilitate uncertainty analysis in numerical simulation. The %CO_2 capture and reboiler duty probability distributions were analysed with the same sample for the four operational conditions. A Gaussian Processes (GP) metamodel was then constructed for each operational condition for Sobol indices computation.

Monte Carlo Simulation

Convergence domains

The simulation framework demonstrated high numerical consistency, achieving a success rate of 95.7% across 1200 runs (300 per operational point). 1148 calculations were completed successfully, while the remainder resulted in warnings (21), errors (26) or non-convergence (5).

Table 3: Operational conditions chosen for MCS

Cases	K1	K13	12A	2D
Lean solvent flow rate (kg/h)	6804	6804	4166	4600
Lean loading (mol CO_2/ mol MEA)	0.152	0.172	0.144	0.200
Flue Gas flow rate (kg/hr)	2248	2238	2976	1998
%CO_2 vol on flue gas	0.115	0.093	0.097	0.107
Absorber active sections	3	3	3	2
Intercooling	Yes	No	Yes	No
%CO_2 capture of deterministic model	***99.95***	***98.37***	***76.81***	***87.4***

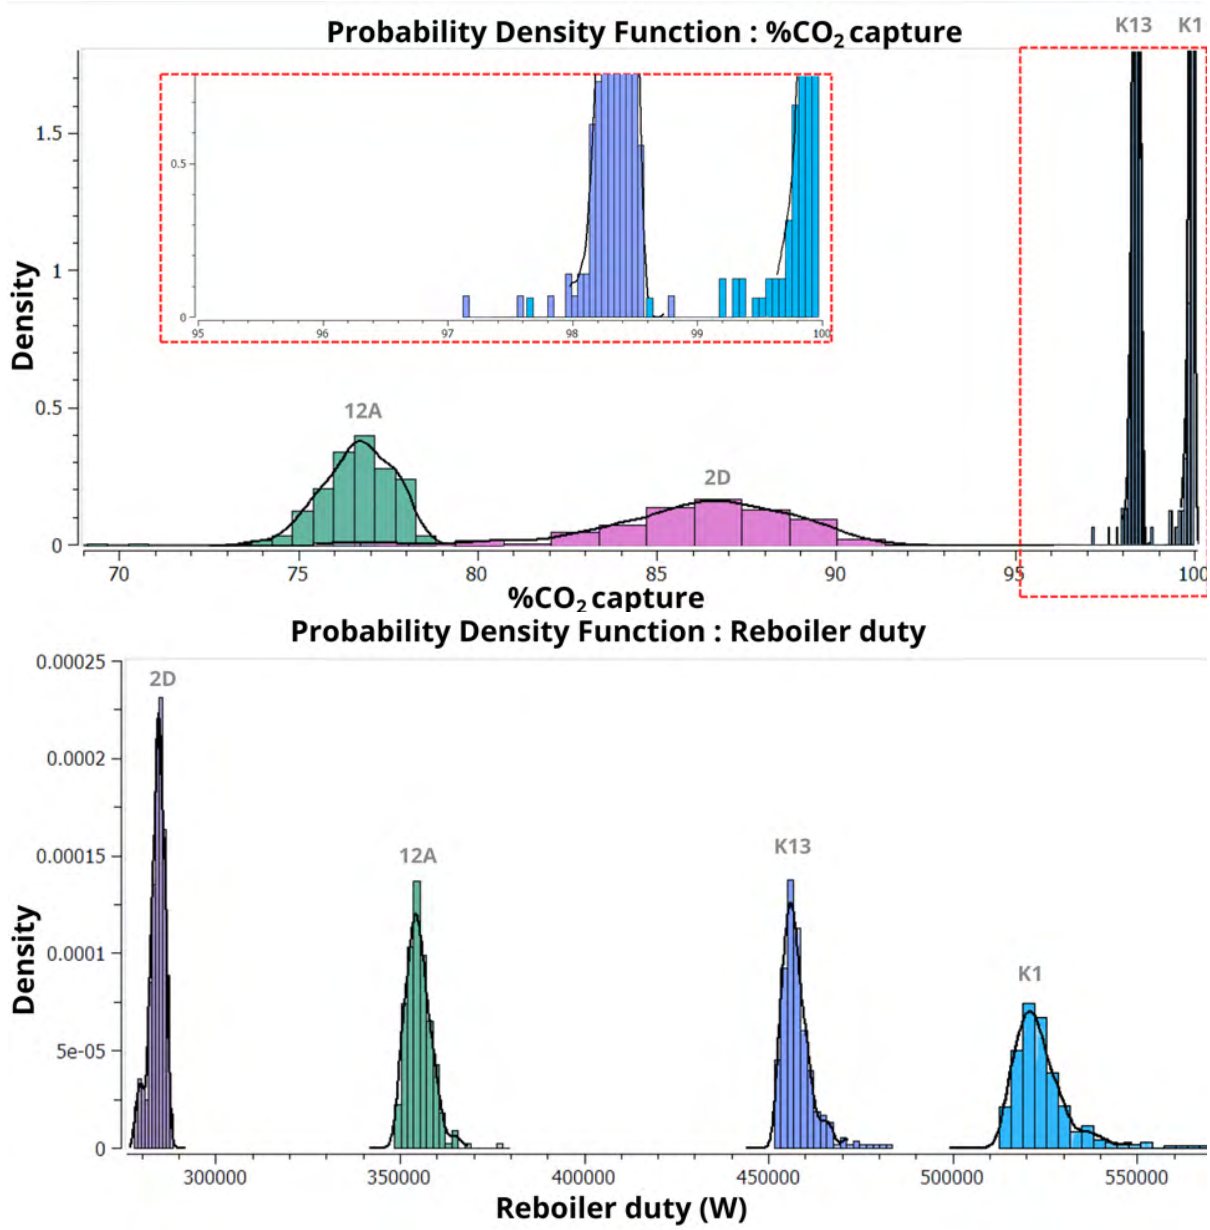

Figure 3. Probability distribution of $\%CO_2$ capture (%) and reboiler duty (W) using CCSI model for mass transfer and hold-up uncertain parameters listed in Table 1 and operational points listed in Table 3 : Blue (K1), purple (K13), green (12A) and pink (2D).

To maintain the independence of the results and prevent error propagation, the system was reinitialised following any failure to converge. Discrepancies were primarily confined to cases involving low liquid mass transfer C_L coefficient (< 0.2) and/or lower lean flow rates.

RESULTS

Figure 3 reports the probability distribution of $\%CO_2$ capture and reboiler duty for the conditions reported in Table 3. Table 4 reports the probabilistic moments, mean (μ) and standard deviation (σ), for a wider range of outputs for the studied operational conditions.

For the cases K1 and 12A, the most divergent $\%CO_2$ capture result of 45.5% and 32.4% respectively, and reboiler duty of 652 kW and 465 kW respectively, were removed from the figure in order to better display the results (even if the simulations were completed normally; the probabilistic moments presented in Table 4 still include those outlier values). This divergence is explained by the lower C_L value (0.01). The other two cases resulted

Table 4: Probability moments for UQ analysis of CCSI model using MCS sampling

Cases	$\mu_{CO2\ capture}$ (%)	$\sigma_{CO2\ capture}$ (%)	$\mu_{reboiler\ duty}$ (kW)	$\sigma_{reboiler\ duty}$ (kW)	$\mu_{CO2\ loading}$ (mol CO_2/ mol MEA)	$\sigma_{CO2\ loading}$ (mol CO_2/ mol MEA)	$\mu_{rich\ solvent\ flowrate}$ (kg/h)	$\sigma_{rich\ solvent\ flowrate}$ (kg/h)
K1	99.71	3.15	524.60	11.28	0.411	0.0084	7125	12.88
K13	98.34	0.14	458.10	4.90	0.374	0.0003	6905	4.57
12A	76.43	2.79	355.60	7.73	0.493	0.0139	4200	5.33
2D	86.26	3.14	284.16	2.03	0.461	0.0057	4585	0.88

in an error for this parameter value. The other outlier %CO_2 capture points on the PDF were also obtained with low C_L values.

Surrogate Models

To complete UQ analysis of CCSI model obtaining the Sobol indices, the GP surrogate model was implemented and tested using Persalys. The two mentioned outlier points in cases K1 and 12A (%CO_2 capture result of 45.5% and 32.4%) were removed from the database for surrogate model construction in order to improve the model accuracy. Because GP implementation for K1 was not successful using the complete sample (model not validated for the test set), two additional outlier points were removed (with %CO_2 capture result of 97.7% and 98.61%).

Table 5 gives the residue (ε) and the agreement with the test set (Q^2) using the leave-one-out (LOO) analytic cross-validation method for each output for the Gaussian Process (with Matern method). LOO analytic method details can be consulted in [17].

Table 5: Residual error and Q^2 LOO for Gaussian Process Metamodel

Cases	K1	K13	12A	2D
$\varepsilon_{CO_2\ capture}$ (%)	0.002	0.003	0.017	0.015
$Q^2_{CO_2\ capture}$	0.87	0.88	0.94	0.99
$\varepsilon_{reboiler\ duty}$ (kW)	0.046	0.031	0.046	0.008
$Q^2_{reboiler\ duty}$	0.99	0.99	0.96	0.99
$\varepsilon_{CO_2\ loading}$ (mol CO_2/ mol MEA)	6.4 10^{-6}	5.7 10^{-6}	7.7 10^{-5}	4.5 10^{-5}
$Q^2_{CO_2\ loading}$	0.87	0.89	0.95	0.98
$\varepsilon_{rich\ solvent\ flowrate}$ (kg/h)	0.05	0.03	0.02	0.01
$Q^2_{rich\ solvent\ flowrate}$	0.32	0.99	0.99	0.97

The surrogate model of K1 is less accurate and does not present a good agreement with the original model for rich solvent flow rate (rich CO_2 loading and %CO_2 capture are also less accurate compared with the other GP surrogate models).

Figure 4 reports the 1st and total Sobol indices obtained by the GP surrogate model presented in Table 5 for %CO_2 capture and reboiler duty. As a method based on variance decomposition, the 1st order Sobol index is a global measurement of the output sensitivity to each parameter input [16], while higher-order Sobol indices also show parameter interactions. The total Sobol indices represent the global contribution of the parameters to the variance of the outputs. For cases with lower solvent flow rate (12A and 2D), C_L present-ed the stronger impact on the %CO_2 capture, followed by A_1. For cases with higher solvent flow rate (K1 and K13), the parameters A_1 and K_{F1} were more influent for %CO_2 capture, with a contribution of C_L for K_1 and C_G for K_{13}. Those cases also have a bigger deviation between the 1st and total Sobol indices for C_L, A_1 and K_{F1} for %CO_2 capture results. C_L presented an impact on the reboiler duty for all operational conditions. Since the gas flow rate is in a similar range for the analysed cases, higher solvent flow rates reduce the contact time between the liquid and gas phases. With less time to reach equilibrium, the interfacial area and reaction kinetic properties become limiting factors in the CO_2 absorption process. However, the reboiler duty does not follow the same phenomena as CO_2 capture because it is connected with the stripper column thermodynamics, which is less dependent on the contact area between the liquid and gas phases and where no kinetic reactions are considered. Therefore, the most influential property for reboiler duty analysis is the liquid mass-transfer coefficient.

CONCLUSION

The model validation was sufficiently accurate to allow the uncertainty quantification analysis for different points across operational condition ranges. For the analysed uncertain properties (mass-transfer, liquid hold-up and reaction kinetics), uncertainty propagation using

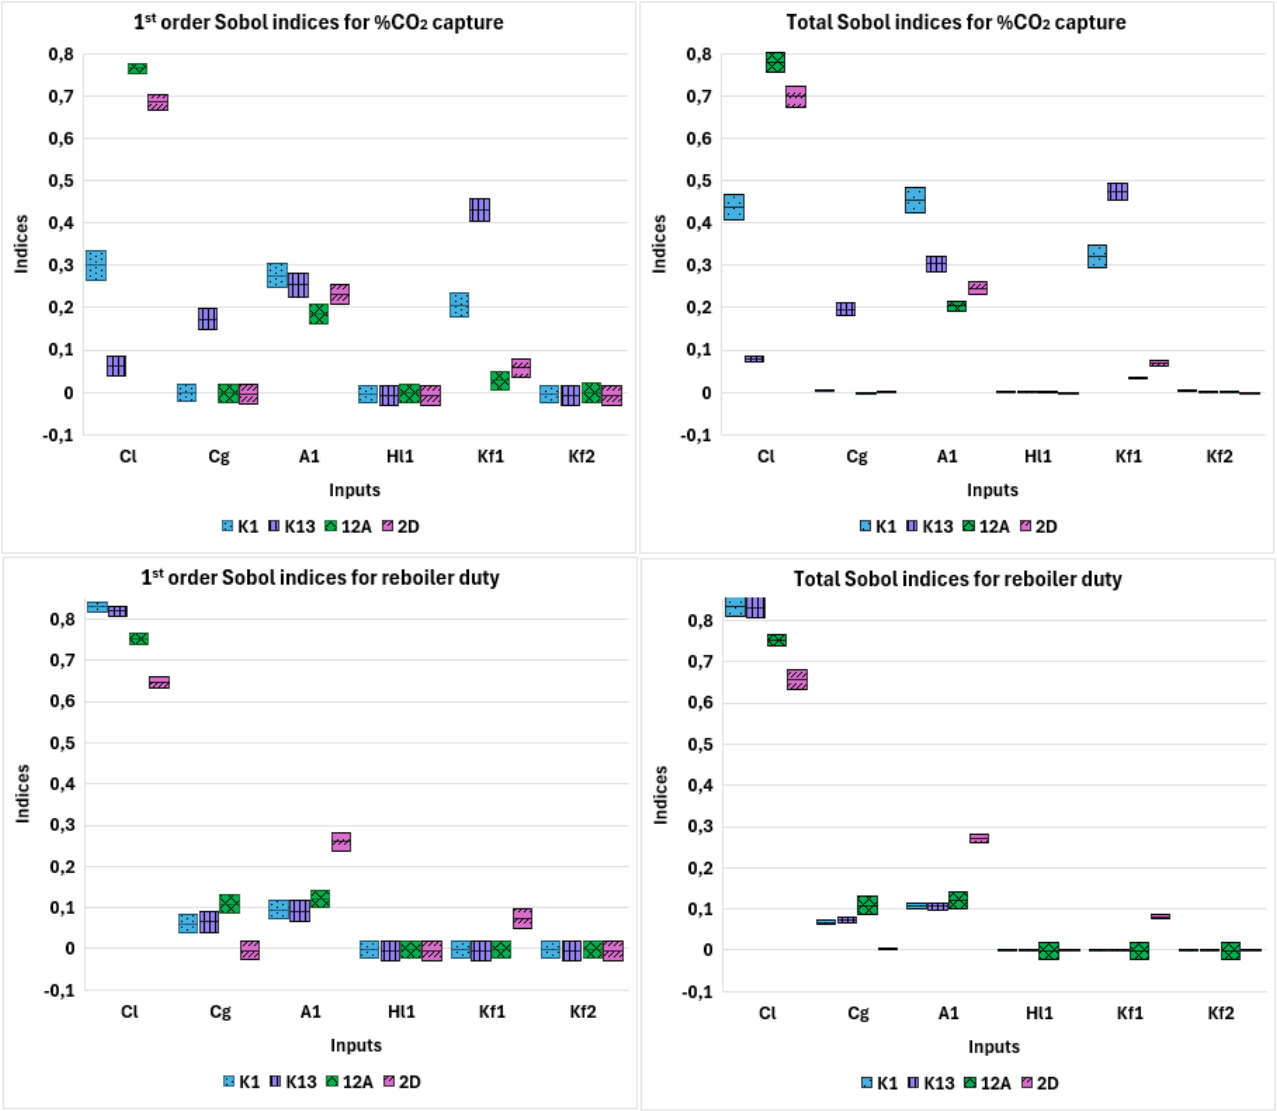

Figure 4: 1st and total Sobol indices with 95% confidence interval obtained by the GP surrogate model: Blue (K1), purple (K13), green (12A) and pink (2D).

MCS sampling demonstrated a reasonably low deviation from average values for the analysed outputs, especially for %CO_2 capture which presented a maximum standard deviation of 3.15% for the four operational conditions investigated (considering outlier points). Therefore, the GP surrogate model presented satisfying agreement for K13, 12A and 2A cases and presented some limitations for K1 (for rich solvent flow rate property). The sensitivity analysis of the complete model shows different impacts on the output distributions according to operational conditions. Among analysed properties, liquid mass-transfer coefficient shows the highest impact on the convergence of the samples. This is consistent with sensitivity analysis that demonstrated the strong impact of the C_L parameter, especially for lower solvent flow rate, indicating the liquid mass-transfer property influence. Interfacial area and reaction kinetic properties become more influent for greater liquid-gas ratio.

A more thorough analysis is expected for further works: (1) to consider other uncertain properties in the UQ analysis, and (2) other methods for sensitivity analysis. Surrogate models could also be built to replace the rigorous model, with varying operational conditions, which would require a deeper investigation on model convergence (e.g., initialisation) but allow wider UQ analysis techniques with larger samples.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the Agence Nationale de la Recherche et de la Technologie (ANRT) for their financial support (Cifre 2024/0837).

Additionally, the authors would like to thank Bertrand looss for UQ- and Persalys-related discussions and Joshua Morgan for very open data and information-sharing on the CCSI model and NCCC experimental campaigns.

REFERENCES

1. International Energy Agency (IEA). (2021). Net Zero by 2050: A Roadmap for the Global Energy Sector. IEA, Paris. Available at: https://www.iea.org/reports/net-zero-by-2050
2. International Energy Agency (IEA). (2020). CCUS in Clean Energy Transitions. IEA, Paris. Available at: https://www.iea.org/reports/ccus-in-clean-energy-transitions
3. Morgan, Joshua Cole, "Physical Property Modeling of Solvent-Based Carbon Capture Processes with Uncertainty Quantification and Validation with Pilot Plant Data" (2017). Graduate Theses, Dissertations, and Problem Reports. https://researchrepository.wvu.edu/etd/6262/
4. Soares Chinen, Anderson, "Development of Hydraulic Models, Mass Transfer Models, and Dynamic Models of Solvent-Based Carbon Capture Processes with Uncertainty Quantification and Validation with Pilot Plant Data" (2018). Graduate Theses, Dissertations, and Problem Reports. https://researchrepository.wvu.edu/etd/6685/
5. Carbon Capture Simulation Initiative (CCSI), *MEA Steady State Model*, version 3.2.1, GitHub, 2022. [Online]. Available: https://github.com/CCSI-Toolset/MEA_ssm/releases/tag/3.2.1
6. Morgan, J.C. et al. (2018) 'Development of a Rigorous Modeling Framework for Solvent-Based CO2 Capture. Part 2 : Steady-State Validation and Uncertainty Quantification with Pilot Plant Data', Industrial & Engineering Chemistry Research, 57(31), p. 10464-10481. https://doi.org/10.1021/acs.iecr.8b01472
7. Morgan JC, Chinen AS, Anderson-Cook C, Tong C, Carroll J, Saha C, Omell B, Bhattacharyya D, Matuszewski M, Bhat KS, Miller DC. Development of a framework for sequential bayesian design of experiments: application to a pilot-scale solvent-based CO2 capture process. Applied Energy 262:114533 (2020). https://doi.org/10.1016/j.apenergy.2020.114533
8. Kuncheekanna VN, Jakobsen JP. Exploring global uncertainty quantification and sensitivity analysis methodologies: CO2 capture absorber model with MEA solvent as a test case. Chemical Engineering Research and Design 200:107-133 (2023). https://doi.org/10.1016/j.cherd.2023.09.030
9. Song , Seibert AF, Rochelle GT. Mass transfer parameters for packings: effect of viscosity. Ind. Eng. Chem. Res. 57:718-729 (2018). https://doi.org/10.1021/acs.iecr.7b04396
10. Tsai RE, Seibert AF, Eldridge RB, Rochelle GT. A dimensionless model for predicting the mass?transfer area of structured packing. AIChE Journal 57:1173-1184 (2011). https://doi.org/10.1002/aic.12345
11. Aboudheir A, Tontiwachwuthikul P, Chakma A, Idem R. Kinetics of the reactive absorption of carbon dioxide in high co2-loaded, concentrated aqueous monoethanolamine solutions. Chemical Engineering Science 58:5195-5210 (2003). https://doi.org/10.1016/j.ces.2003.08.014
12. Vaidya PD, Kenig EY. Co₂?alkanolamine reaction kinetics: a review of recent studies. Chem Eng & Technol 30:1467-1474 (2007). https://doi.org/10.1002/ceat.200700268
13. Knuutila H, Juliussen O, Svendsen HF. Kinetics of the reaction of carbon dioxide with aqueous sodium and potassium carbonate solutions. Chemical Engineering Science 65:6077-6088 (2010). https://doi.org/10.1016/j.ces.2010.07.018
14. Gondal S, Svendsen HF, Knuutila HK. Activity based kinetics of CO2–OH? systems with li+, na+ and K+ counter ions. Chemical Engineering Science 151:1-6 (2016). https://doi.org/10.1016/j.ces.2016.05.009
15. Konno H, Koshizuka T. Mean-absolute deviation model. IIE Transactions 37:893-900 (2005). https://doi.org/10.1080/07408170591007786
16. Marrel A, looss B, Laurent B, Roustant O. Calculations of sobol indices for the gaussian process metamodel. Reliability Engineering & System Safety 94:742-751 (2009). https://doi.org/10.1016/j.ress.2008.07.008
17. Beaume, G. (2009). Modélisation et simulation numérique de l'écoulement d'un fluide complexe (Doctoral thesis). HAL Open Science. https://theses.hal.science/tel-00416435

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

An End-to-End Pure Component Property Prediction Framework Based on a Hierarchical Molecular Fragmentation Method

Jianfeng Jiao[a] and Jie Li[a]*

[a] The University of Manchester, Department of Chemical Engineering, Manchester, UK
* Corresponding Author: jie.li-2@manchester.ac.uk

ABSTRACT

The accurate prediction of pure component properties has consistently been a critical issue in fields such as chemical engineering, biomedicine, and environmental science. In recent years, end-to-end deep learning methods have shown significant improvement over traditional machine learning approaches. This is due to their ability to automatically learn task-relevant representations from raw molecular data. In addition to accurate property prediction, researchers have increasingly focused on how specific fragment structures influence molecular properties. However, existing fragmentation methods based on predefined rules and group libraries struggle to capture novel molecular structures, which hampers the development of new materials and drugs. To address these challenges, this work proposes a hierarchical molecular fragmentation method. This method can automatically segment molecules into multiple fragments containing key functional groups. Then a three-branch graph attention network was constructed to achieve multi-level representation. Finally, a multi-layer perceptron is employed to establish the mapping relationship between molecular features and physical property values. Twenty datasets were used for validation, which can be grouped into four categories: Thermodynamic Properties, Pharmacokinetics, Toxicological Properties, and Industrial Safety. The results show that the best performance is achieved, with the average error reduced by 6.8% compared to existing research.

Keywords: Machine Learning, Artificial Intelligence, Algorithms, Multiscale Modelling, Property Prediction

INTRODUCTION

Accurate prediction of molecular physicochemical properties relies on effective molecular representations. Traditional descriptors like MACCS keys and Extended-Connectivity Fingerprints (ECFP) have been widely used to encode structural information into vectors [1]. However, these fixed-length representations suffer from limitations such as hash collisions, information loss in many-to-one mappings, and a lack of direct chemical interpretability. In parallel, Group-Contribution (GC) methods estimate properties by summing the contributions of predefined structural groups [2]. While recent integrations with machine learning have improved their nonlinear expression ability [3]. However, GC methods inherently ignore topological connectivity. Furthermore, their reliance on fixed libraries results in high-dimensional, sparse feature vectors that can lead to model overfitting and poor generalization to molecules containing novel substructures [4].

With the rise of deep learning, Graph Neural Networks (GNNs) have achieved success by treating molecules as graphs and updating atomic representations via message passing [5], meanwhile attention mechanisms have further improved accuracy [6]. However, chemists typically reason via specific structural motifs rather than treating the molecule as an indivisible whole. Hence, researchers have employed the predefined BRICS scheme to decompose molecules into fragments, subsequently treating each fragment as a subgraph and representing it using GNNs. Similarly, predefined group libraries have been used to fragment molecules, where each group is regarded as a subgraph. These approaches have been shown to effectively improve predictive accuracy,

because molecular properties are often governed by specific functional groups; for example, nitro, amino, and quinone groups often increase molecular toxicity, whereas carboxyl groups are widely recognised as detoxifying motifs. The resulting fragments can capture common functional groups and establish mappings to target properties via attention mechanisms. However, owing to the limited coverage of predefined rules, prediction accuracy often deteriorates when extrapolating to out-of-distribution (OOD) molecules. To address the aforementioned issues, Li et al. [7] proposed SMILES Pair Encoding (SPE) method, which learns frequent SMILES substrings from a large dataset, then segments molecules into tokens based on learned SMILES substrings. Benefiting from the large scale of the dataset, the learnable chemical space is substantially expanded, thereby alleviating the out-of-distribution (OOD) issue. However, SPE completely loses topological information. The model only knows that the molecule contains these three fragments. It does not know how these fragments are connected.

To address these challenges, based on a novel hierarchical fragmentation method, a molecular property prediction framework is proposed. Then, a three-branch GNNs was developed to enable representations at both the molecular-graph level and the fragment level. Twenty datasets were used for validation, which can be grouped into four categories: Thermodynamic Properties, Pharmacokinetics, Toxicological Properties, and Industrial Safety.

METHODOLOGY

Hierarchical molecular fragmentation method

The quality of molecular fragments depends on whether they contain functional groups that determine molecular properties or the core scaffolds that host these functional groups. Such motifs are typically highly cohesive in topology; they exhibit dense internal bonding and are connected to the rest of the molecule via only a small number of external bonds. In this subsection, we propose a novel hierarchical molecular fragmentation method. It systematically searches for highly cohesive topological substructures within molecules and enables explicit modelling of molecular fragments.

We treat each molecule as an undirected molecular graph $G = (V, E)$, where V are non-hydrogen nodes and $E$are bonds. Fragmentation is performed by a greedy-and-recursive procedure that repeatedly extracts one dominant local substructure and then applies the same procedure to the remaining subgraph.

Step 1: State initialization and candidate generation.

We first define the core C as a set of k directly connected atoms in graph G. We set a maximum core order k (in this work $k = 4$) and define $S_k(G)$ denote the collection of all such connected k -atom cores $S_k(G) = \{C \subseteq V \mid |C| = k, G[C] \text{ is connected}\}$.

For each core $C \in S_k(G)$, we define its environment as the core together with all atoms directly bonded to any atom in the core: $\Omega(C) = C \cup \{u \in V \mid \exists v \in C, (u, v) \in E\}$. Cores that yield the same environment are grouped, producing a candidate environment set $\mathcal{M}_{ca}$.

Step 2: Weight evaluation and greedy selection.

For each candidate environment $M \in \mathcal{M}_{ca}$, we assign a weight based on how many cores map to it across orders:

$$W(M) = \sum_{k=1}^{K} |\{C \in S_k(G) \mid \Omega(C) = M\}|$$

We then select the environment M^* with the largest weight as the dominant fragment, $M^* = \arg\max_{M \in \mathcal{M}_{ca}} W(M)$.

Step 3: State update and recursive splitting.

After selecting M^*, we remove M^* from the graph G to obtain the residual subgraph. At this stage, we examine whether the residual subgraph contains more than two non-hydrogen nodes. If so, the residual subgraph is treated as an intact unit, and Steps 1–2 are repeated to further fragment the residual subgraph until the stopping criterion is satisfied.

The final fragmentation results for the molecule O=C(N)Nc1ccc(C)cc1 are shown in **Fig. 1**. As can be seen, the proposed method successfully identifies the urea moiety (–NH–C(=O)–NH2) and isolates it as an individual fragment (blue), while the remaining part corresponds to a tolyl aromatic ring (red). For this molecule, the urea group governs its polarity and hydrogen-bonding capability and is therefore one of the most critical structural units for property characterization.

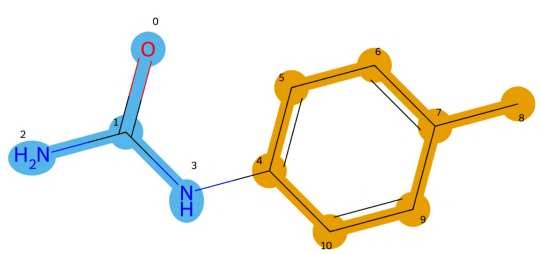

Figure 1. The fragmentation results of O=C(N)Nc1ccc(C)cc1 based on the proposed hierarchical molecular fragmentation method.

In addition, we fragmented the molecule using a predefined group-contribution (GC) library [8] and the BRICS rules [9], as shown in **Figs. 2–3**. Under the GC scheme, the phenyl ring is decomposed into multiple

aromatic carbon atoms, and the methyl group becomes disconnected from the ring, preventing the formation of an intact fragment. Under BRICS, the complete urea structure is not preserved; instead, it is split into a carbamoyl fragment (–NH2C(=O)–) and an –NH fragment.

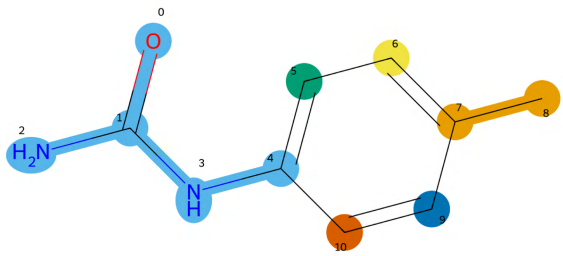

Figure 2. The fragmentation results of O=C(N)Nc1ccc(C)cc1 based on the group contribution method.

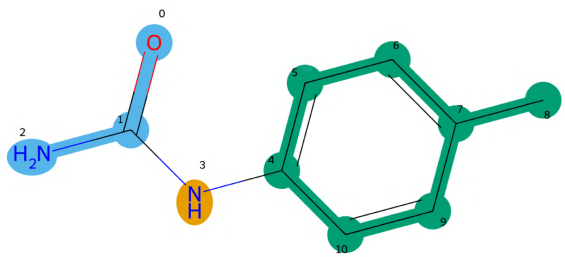

Figure 3. The fragmentation results of O=C(N)Nc1ccc(C)cc1 based on the BRICS method.

HierAttnGNN

The HierAttnGNN contains three branches. The first branch learns representations at the whole molecule level. The second branch encodes molecular fragments. The third branch is a junction tree that represents the connectivity relationships between fragments, as shown in **Fig. 4.**

Molecule level representation

This branch learns a global embedding for the complete molecular graph. Given the molecular graph defined above, let x_v and $x_{e(v,u)}$ denote the node and edge features, as shown in **Table 1-2** respectively.

Table 1 Atom information used to featurize the molecular graph

No.	Feature Name	Description	Encoding	Dim
1	Atom Symbol	Atom type (C, N, O, S, F, P, Cl, Br, I)	One-hot	9
2	Atom Degree	Number of covalent bonds (0–5)	One-hot	6
3	Total Num Hs	Total number of hydrogen atoms (0–4)	One-hot	5
4	Implicit Valence	Implicit valence electrons (0–5)	One-hot	6
5	Hybridization	Orbital hybridization type (SP, SP2, SP3, SP3D, SP3D2)	One-hot	5
6	Is Aromatic	True (1) / False (0)	Boolean (Numeric)	1
7	Chirality Possible	Potential for chirality (RDKit property)	Boolean (Numeric)	1
8	Chirality Type (CIP Code)	CIP code classification (R, S)	One-hot	2
9	Formal Charge	Electrical charge of the atom	Direct Numeric	1
Total				36

Table 2 Bond information used to featurize the molecular graph

No.	Feature Name	Description	Encoding	Dim
1	Is Single Bond	Single, Double, Triple, or Aromatic	Boolean (0/1)	1
2	Is Double Bond	Conjugation status of the bond	Boolean (0/1)	1
3	Is Triple Bond	Checks if the bond type is Triple	Boolean (0/1)	1
4	Is Aromatic Bond	Checks if the bond type is Aromatic	Boolean (0/1)	1
5	Is Conjugated	Checks if the bond is conjugated	Boolean (0/1)	1
6	Is In Ring	Ring membership status	Boolean (0/1)	1
Total				6

We first project the atom and bond feature vectors into a shared hidden space, yielding the initial hidden states of atoms and bonds, $h_v^{(0)}$ and $h_{e(v,u)}^{(0)}$, as in Eqs. (1)–(2).

$$h_v^{(0)} = \mathrm{ReLU}(\mathbf{W}_{node} x_v + b_n) \quad (1)$$

$$h_{e(v,u)}^{(0)} = \mathbf{W}_{edge} x_{e(v,u)} + b_e \quad (2)$$

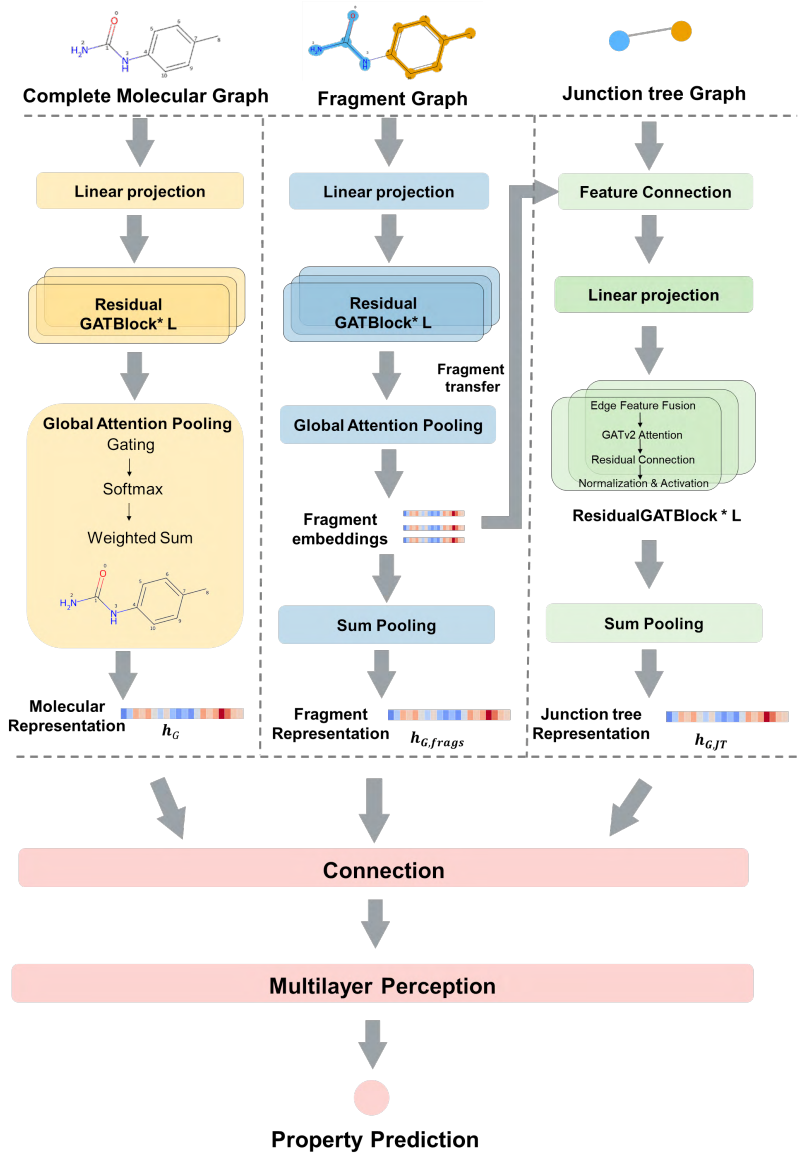

Figure 4. The architecture of the proposed HierAttnGNN framework, consisting of molecule-level, fragment-level, and junction-tree branches

We then stack L ResidualGATBlock modules to update node representations. Each block is built on a GATv2 layer, which performs neighbourhood aggregation with learnable attention. At layer l, edge features incident to node v are first fused into the node state, so that subsequent attention can depend on both node and bond information, as in Eq. (3).

$$h'_v = h_v^{(l)} + \sum_{u \in \mathcal{N}(v)} h_{e(v,u)}^{(l)} \tag{3}$$

Next, for each neighbour u of v, GATv2 computes an unnormalised attention score a_{vu} using a dynamic attention function (Eq. (4)), which is more expressive than the original GAT and can capture more complex structural dependencies [10]. The scores are normalised with a softmax to obtain the importance weights α_{vu} (Eq. (5)).

$$a_{vu} = w_{attn}^T \bullet \mathrm{LeakyReLU}(\mathbf{W}_L h'_v + \mathbf{W}_R h'_u) \tag{4}$$

$$\alpha_{vu} = \mathrm{softmax}_u(a_{vu}) = \frac{\exp(a_{vu})}{\sum_{k \in \mathcal{N}(v) \cup \{v\}} \exp(a_{vk})} \tag{5}$$

where, $\mathbf{W}_L$ and $\mathbf{W}_R$ are learnable weight matrices. Specifically, $W_L, W_R \in \mathbb{R}^{d_h \times d_h}$ are learnable linear transformation matrices. The vector $w_{\mathrm{attn}} \in \mathbb{R}^{d_h}$ is a learnable attention weight vector.

Using α_{vu}, the model forms a weighted sum of neighbour features to update each node. A multi-head

mechanism is applied for stability and capacity, and the outputs of K heads are concatenated, as shown in Eq. (6).

$$h_v^{(l+1)} = \|_{k=1}^{K} \sum_{u \in \mathcal{N}(v) \cup \{v\}} \alpha_{vu}^{(k)} \mathbf{W}^{(k)} h'_u \quad (6)$$

where, $\|$ denotes concatenation.

To mitigate over-smoothing and to stabilise optimisation, each block combines residual connections with batch normalisation, giving the complete update in Eq. (7).

$$h_{v,\text{out}}^{(l+1)} = \text{BatchNorm}(\text{ELU}(h_v^{(l+1)}) + h_v^{(l)}) \quad (7)$$

After L GATv2 layers, we obtain the final node embeddings $\{h_v^{(L)}\}$. A Global Attention Pooling readout is then applied to produce the graph-level representation h_G, as in Eq. (8), where W_{gate} learns the gating (importance) score of each node and $h_G$is taken as the final molecule-level representation.

$$h_{\mathcal{G}} = \sum_{v \in V} \left(\frac{\exp(\mathbf{W}_{gate} h_v^{(L)})}{\sum_{u \in V} \exp(\mathbf{W}_{gate} h_u^{(L)})} \right) h_v^{(L)} \quad (8)$$

where, $h_v^{(L)}$ is the final feature vector of node $v$after L GNN layers. $\mathbf{W}_{gate}$ is a learnable linear transformation used to compute the gating or importance score of each node. The aggregated vector $h_{\mathcal{G}}$ is taken as the final representation of the molecular graph.

Fragment level representation

The second branch encodes molecular fragments produced by the hierarchical fragmentation. Each fragment is treated as an independent subgraph and is processed using the same representation pipeline as the whole-molecule branch, i.e., feature initialization, GATv2-based message passing, and Global Attention Pooling. This yields a fragment embedding $h_{\mathcal{G}_{frag}}$ that summarizes the corresponding fragment subgraph.

To aggregate fragment-level information back to the parent molecule, we apply sum pooling over all fragment embeddings that originate from the same molecular graph G. The resulting vector $h_{G,frags}$ is defined as:

$$h_{\mathcal{G},frags} = \sum_{g_{frag} \in \mathcal{F}(\mathcal{G})} h_{\mathcal{G}_{frag}} \quad (9)$$

where, $F(\mathcal{G})$ denotes the set of fragments derived from the original molecular graph $\mathcal{G}$.

Junction tree

In the third branch, we explicitly model fragment connectivity by constructing a junction tree, where each molecular fragment is treated as a super node and bonds between fragments are treated as edges. Message passing on this coarse grained graph enables the model to capture spatial and chemical interaction patterns among fragments. The initial feature vector of each junction tree node, $x_{v,jt\text{-}aug}$, is formed by concatenating two parts: (i) static fragment descriptors and (ii) a dynamic fragment embedding. In particular, the scalar W_i incorporates the fragment weight, allowing the network to recognise more stable structural cores and assign them higher importance during message passing, as shown in Eq. (10).

$$x_{v,jt-aug} = [c_i;\ W_i; h_{\mathcal{G},_{frag}}] \quad (10)$$

where, $c_i$is the number of atoms contained in fragment i, W_i is the statistical weight of this fragment, and $h_{\mathcal{G},\text{frag}}$ is the graph level embedding of this fragment obtained from the second branch.

Next, the augmented feature vector is passed through a linear projection layer and mapped into the hidden space of this branch, yielding the initial hidden state of each junction tree node, as shown in Eq. (11).

$$h_{v,jt}^{(0)} = \mathbf{W}_{jt} x_{v,jt-aug} + b_{jt} \quad (11)$$

As in the first two branches, node features on the junction tree graph are then updated by L ResidualGATBlock modules. Information is propagated between nodes along edges, and the updated hidden states $h_{v,jt}^{(L)}$ represent each fragment in the context of its fragment neighbourhood. Finally, to obtain the junction tree view representation of the whole molecule, we aggregate the final hidden states of all junction tree nodes using sum pooling, as shown in Eq. (12).

$$h_{\mathcal{G},jt} = \sum_{v \in V_{JT}(\mathcal{G})} h_{v,jt}^{(L)} \quad (12)$$

where, $V_{JT}(\mathcal{G})$ denotes the set of nodes in the junction tree corresponding to molecule $\mathcal{G}$.

DATASETS AND MODEL TRAINING

This work uses 20 datasets released by Alshehri et al. [3, 11], covering four property categories: thermophysical, pharmacokinetic, toxicological, and industrial safety properties, summarised in **Table 3**. Thermophysical property prediction is directly relevant to process design and operation; for example, the normal boiling point T_b is a basic design parameter that affects equipment selection and operating conditions, while the critical constants (T_c, P_c) determine the critical point and are widely used in equations of state and phase-equilibrium calculations, including supercritical-fluid applications. Accurate prediction of safety- and environment-related properties is likewise essential for chemical safety management and environmental risk evaluation. As a representative example, auto-ignition temperature is important for fire and explosion prevention and risk assessment, and reliable prediction of solubility is critical for toxicological evaluation, drug discovery, and human-health risk assessment.

Table 3. Dataset categories and descriptions

Classification	Property	Sample number	Description
Thermophysical properties	T_b	5276	Normal boiling point (K)
	T_m	9248	Normal melting point (K)
	T_c	776	Critical temperature (K)
	P_c	774	Critical pressure (bar)
	V_c	773	Critical volume (cc/mol)
	H_v	425	Enthalpy of vaporization at 298 K (kJ/mol)
	H_{fus}	748	Normal enthalpy of fusion (kJ/mol)
	H_f	1059	Standard enthalpy of formation (kJ/mol)
	G_f	756	Gibbs energy of formation at 298 K (kJ/mol)
	L_{mv}	1055	Liquid molar volume at 298 K (cc/mol)
	HSolP	1368	Hildebrandt solubility parameter at 298 K) (MPa1/2)
Pharmacokinetics properties	BCF	589	Bioconcentration factor
	LogP	12184	Octanol–water partition coefficient
	$LogW_s$	2563	Aqueous solubility (mol/L)
	PCO	606	Photochemical oxidation potential
	pKa	1632	Acid dissociation constant
Toxicological Properties	$LC50_{(FM)}$	705	Fathead Minnow 96-H LC50 (mol/L)
	LD50	4780	Toxicity (oral rat) (mol/kg)
Industrial Safety	OSHA-TWA	422	Permissible exposure limit (mol/m3)
	AiT	570	Auto ignition temperature (K)

A six-fold cross-validation is used for dataset splitting. Specifically, 10% of the samples are first held out as an independent test set. This test set is strictly excluded from model training, hyperparameter tuning, and validation, and is used only once for the final evaluation of the selected model. The remaining 90% of the data are then randomly and evenly divided into six folds. For each run, five folds are used for training and the remaining fold is used for validation, and the procedure is repeated six times such that every fold serves as the validation set exactly once. To prevent data leakage, data standardisation is performed only after the test set has been separated.

Training, evaluation metrics and hyperparameters

Model parameters were optimised using AdamW with decoupled weight decay, which applies weight decay independently from the gradient-based Adam update and thus provides more effective regularisation than coupling L2 penalties into the gradient. The training objective was to minimise the mean squared error (MSE). To improve training stability and avoid gradient explosion, we applied gradient clipping: after backpropagatio, the global L2 norm of all parameter gradients $\| \mathbf{g} \|_2$was computed, and when it exceeded a threshold g_{clip}, gradients were rescaled. Following common practice, we set $g_{\text{clip}} = 1$. To mitigate overfitting, two complementary strategies were used. First, early stopping was adopted by monitoring the validation root mean squared error (RMSE); if no improvement was observed for 25 consecutive epochs, training was terminated and the checkpoint with the lowest validation RMSE was retained. Second, a learning-rate scheduler was employed: when validation performance plateaued, the learning rate was multiplied by 0.5 until reaching a minimum of 1×10^{-5}, enabling finer optimisation in later training stages.

Generalisation performance was evaluated using

RMSE and the coefficient of determination R^2. Under six-fold cross-validation, model selection and reporting followed a consistent protocol: within each fold, the checkpoint achieving the lowest validation RMSE was saved as the best model for that fold. After all six folds, the six resulting models were evaluated on the independent held-out test set, and the reported test metrics were computed as the average across these six models rather than from a single run.

Because different fragmentation schemes produce fragments with substantially different sizes and topological complexity, a single fixed hyperparameter configuration can unfairly constrain some models and obscure their attainable performance. We therefore applied Bayesian hyperparameter optimisation using Optuna [12]. A unified search space was defined, and for each GNNs variant constructed under a specific fragmentation method, 50 Optuna trials were executed independently. To avoid information leakage, the dataset splits used during hyperparameter search were generated with random seeds independent of those used in final evaluation. Validation RMSE was used as the optimisation objective; the best hyperparameter set was then fixed and used for the subsequent six-fold cross-validation. The complete search space and ranges are reported in **Table 4**.

Table 4 The detailed search space and ranges of all hyperparameters

Hyperparameter	Search space and ranges
Learning Rate	$[1\times10^{-4}, 1\times10^{-2}]$
Hidden Dimension	{128, 256, 512}
Attention Heads	{4, 8}
Layers	[2, 11]
Dropout	[0.0, 0.5]
Weight Decay	$[1\times10^{-6}, 1\times10^{-3}]$

RESULTS

The final predictive results are reported in **Table 5-8**. The results of the GP-WP model were taken directly from the literature [13]. The BRICS-GNN, SPE-GNN and GC-GNN models were built under different fragmentation methods. For each dataset, molecular fragments were generated for every molecule, a three-branch architecture was constructed, and property prediction was performed. HiGNN [14] and AttentiveFP [6] model were evaluated using publicly available molecular encoding pipelines, and public implementations were used for prediction and hyperparameter optimisation. To further validate the effectiveness of incorporating fragment weight features in the junction tree branch, an ablation model, Hier-AttnGNN-Ablation, was also constructed. All results are reported as the average performance on the test set over six-fold cross validation.

Table 5 Predictive performance of different models across Pharmacokinetics property datasets

	BCF	LogP	$LogW_s$	PCO	pKa
HierAttnGNN	0.72	0.39	0.76	0.14	1.67
HierAttnGNN-Ablation	0.91	0.403	0.77	0.15	1.76
GPR-WP	0.97	0.408	0.79	0.16	1.78
BRICS-GNN	0.91	0.403	0.79	0.18	1.72
SPE-GNN	0.9	0.43	0.92	0.21	1.95
GC-GNN	0.73	0.4	0.77	0.18	1.87
HiGNN	0.79	0.412	0.78	0.18	1.77
Attention FP	0.74	0.405	0.77	0.16	1.89

Table 6 Predictive performance of different models across Toxicological and Industrial Safety property datasets

	AiT	$LC50_{(FM)}$	LD50	OSHA-TWA
HierAttnGNN	76.09	0.67	0.43	0.66
Hier-AttnGNN-Ablation	78.3	0.703	0.45	0.74
GPR-WP	94.64	0.76	0.43	1.3
BRICS-GNN	89.97	0.96	0.44	0.78
SPE-GNN	83.47	0.89	0.47	1.08
GC-GNN	81.22	0.78	0.46	0.88
HiGNN	81.44	0.69	0.46	0.87
Attention FP	83.01	0.68	0.45	0.77

Table 7 Predictive performance of different models across Thermophysical property datasets (Part 1)

	T_b	V_c	T_c	P_c	G_f
Hier-AttnGNN	19.19	27.31	21.06	2.43	22.45
Hier-AttnGNN-Ablation	19.97	29.57	25.76	3.15	25.39
GPR-WP	25.39	33.43	86.8	2.77	44.37
BRICS-GNN	23.06	32.73	27.5	2.53	32.24
SPE-GNN	27.86	75.12	43.24	4.02	36.11
GC-GNN	19.4	29.88	21.37	2.47	24.66
HiGNN	19.27	34.62	23.21	2.84	24.06
Attention FP	20.77	30.44	22.62	2.99	26.14

Table 8 Predictive performance of different models across Thermophysical property datasets (Part 2)

	H_f	H_{fus}	HSolP	H_v	L_{mv}
Hier-AttnGNN	20.95	3.74	1.66	4.82	0.0068
Hier-AttnGNN-Ablation	24.25	3.88	1.68	4.95	0.0103
GPR-WP	37.42	4.099	2.31	5.72	0.007
BRICS-GNN	28.56	4.44	1.78	5.97	0.019
SPE-GNN	45.72	4.53	1.72	5.95	0.02
GC-GNN	24.9	5.3	1.71	5.01	0.0105
HiGNN	27.83	4.04	1.7	5.18	0.007
Attention FP	26.69	3.76	1.72	4.98	0.0074

HierAttnGNN achieved the best performance across all 20 prediction tasks, with an average RMSE of 12.44, representing a 6.8% reduction relative to GC-GNN at 13.36. Among the models sharing the same three-branch architecture, SPE-GNN consistently produces the highest errors, confirming that discarding topological connectivity between fragments severely impairs predictive accuracy. BRICS-GNN generally outperforms SPE-GNN but suffers notable degradation where its predefined rules fail to preserve intact functional groups, as evidenced by RMSE values of 0.96 on $LC50_{(FM)}$ and 89.97 on AiT, compared with 0.67 and 76.09 for HierAttnGNN. GC-GNN performs comparably to HierAttnGNN on classical thermophysical properties where its predefined library provides adequate coverage, yet it yields the worst RMSE of 5.3 on H_{fus}, illustrating how fine-grained fragmentation of out-of-library substructures leads to sparse feature vectors and poor generalisation. AttentiveFP achieves competitive results on properties governed by local atomic environments, attaining RMSE values of 3.76 on H_{fus} and 0.68 on $LC50_{(FM)}$, but falls behind on tasks requiring multi-level structural reasoning, such as Hf and pKa. HiGNN ranks among the top three on many datasets but inherits the functional-group integrity limitations of BRICS, leading to higher errors on toxicological and structurally diverse thermophysical datasets.

CONCLUSION

A hierarchical molecular fragmentation method is proposed in this work. Without relying on predefined rules or a static vocabulary, the proposed method can automatically partition a molecule into fragments that contain key functional groups. On this basis, a three-branch GNNs is constructed to learn hierarchical representations from the atom level to the fragment level and further to the junction tree level. Evaluation on 20 benchmark datasets shows that HierAttnGNN significantly outperforms existing mainstream model. The average RMSE is reduced by 6.8% relative to the second-best model.

ACKNOWLEDGEMENTS

Authors appreciate financial support from China Scholarship Council (202406440073)

AUTHOR IDENTIFIERS

Author ORCIDs:
Jianfeng Jiao: 0009-0009-8246-3861
Jie Li: 0000-0001-5196-2136

REFERENCES

1. Rogers D, Hahn M. Extended-connectivity fingerprints. J. Chem. Inf. Model. 50:742-754 (2010). https://doi.org/10.1021/ci100050t
2. Gani R. Group contribution-based property estimation methods: advances and perspectives. Current Opinion in Chemical Engineering 23:184-196 (2019). https://doi.org/10.1016/j.coche.2019.04.007
3. Alshehri AS, Tula AK, You F, Gani R. Next generation pure component property estimation models: with and without machine learning techniques. AIChE Journal 68: (2021). https://doi.org/10.1002/aic.17469
4. Aouichaoui ARN, Fan F, Abildskov J, Sin G. Application of interpretable group-embedded graph neural networks for pure compound properties. Computers & Chemical Engineering 176:108291 (2023). https://doi.org/10.1016/j.compchemeng.2023.108291
5. Gilmer J, Schoenholz S S, Riley P F, Vinyals O, Dahl G E. Neural message passing for quantum chemistry. in *International conference on machine learning* 1263–1272 (2017). https://doi.org/proceedings.mlr.press/v70/gilmer17a
6. Xiong Z, Wang D, Liu X, Zhong F, Wan X, Li X, Li Z, Luo X, Chen K, Jiang H, Zheng M. Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism. J. Med. Chem. 63:8749-8760 (2019). https://doi.org/10.1021/acs.jmedchem.9b00959
7. Li X, Fourches D. SMILES pair encoding: a data-driven substructure tokenization algorithm for deep learning. J. Chem. Inf. Model. 61:1560-1569 (2021). https://doi.org/10.1021/acs.jcim.0c01127
8. Hukkerikar AS, Sarup B, Ten Kate A, Abildskov J,

Sin G, Gani R. Group-contribution+ (GC+) based estimation of properties of pure components: improved property estimation and uncertainty analysis. Fluid Phase Equilibria 321:25-43 (2012). https://doi.org/10.1016/j.fluid.2012.02.010

9. Wang J, Wang Y. Brics-based generation and ai-assisted screening of ionic liquids with mechanistic insights into lithium transport in electrolytes. J. Chem. Inf. Model. 65:10961-10976 (2025). https://doi.org/10.1021/acs.jcim.5c01824
10. Brody S, Alon U, Yahav E. How attentive are graph attention networks? *arXiv Prepr. arXiv2105.14491* (2021) https://doi.org/10.48550/arXiv.2105.14491
11. Jiao J, Gao X, Li J. Pure component property estimation framework using explainable machine learning methods. Chinese Journal of Chemical Engineering 84:158-178 (2025). https://doi.org/10.1016/j.cjche.2025.05.011
12. Akiba T, Sano S, Yanase T, Ohta T, Koyama M. Optuna. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining :2623-2631 (2019). https://doi.org/10.1145/3292500.3330701
13. Cao X, Gong M, Tula A, Chen X, Gani R, Venkatasubramanian V. An improved machine learning model for pure component property estimation. Engineering 39:61-73 (2024). https://doi.org/10.1016/j.eng.2023.08.024
14. Zhu W, Zhang Y, Zhao D, Xu J, Wang L. Hignn: a hierarchical informative graph neural network for molecular property prediction equipped with feature-wise attention. J. Chem. Inf. Model. 63:43-55 (2022). https://doi.org/10.1021/acs.jcim.2c01099

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Global Optimization of a Hydrodealkylation Flowsheet through Spatial Decomposition with SNoGloDe

Madeline Leppla[a], Georgia Stinchfield[a], Norman Tran[a], and Carl D. Laird[a]*

[a] Carnegie Mellon University, Department of Chemical Engineering, Pittsburgh, Pennsylvania, USA
* Corresponding Author: claird@andrew.cmu.edu

ABSTRACT

Global optimization of industrial-scale chemical process flowsheets remains challenging due to nonlinearity, nonconvexity, and large problem scale. While equation-oriented modeling frameworks enable high-fidelity representation of industrial processes, obtaining globally optimal solutions is often computationally intractable for off-the-shelf solvers. In this work, we present a decomposition-based global optimization strategy that solves a high-fidelity flowsheet model from the IDAES framework with the Structured Nonlinear Global Decomposition (SNoGloDe) framework. The proposed approach exploits spatial decomposability by partitioning the flowsheet into coupled subproblems linked through a small set of complicating variables and solving them within a prioritized spatial branch-and-bound framework. The methodology is demonstrated on a hydrodealkylation (HDA) process for benzene production, a nonconvex and industrially relevant case study. The flowsheet is decomposed into reactor and separation subproblems, enabling efficient computation of valid global bounds while preserving convergence guarantees. SNoGloDe successfully identifies a (verifiably) globally optimal solution. These results illustrate the potential of decomposition-based global optimization for complex process systems and highlight the advantages of integrating rigorous process modeling with flexible, customizable global optimization algorithms.

Keywords: Optimization, Algorithms, Parallelization, Process Operations

INTRODUCTION

Optimizing chemical process flowsheets presents a central challenge in process systems engineering due to the inherent nonlinearity, nonconvexity, and scale of industrially relevant processes. Traditional nonlinear programming techniques for large-scale process optimization are limited to locally optimal optima. Global solutions to these flowsheets are frequently intractable for general off-the-shelf solvers. On the other hand, global solvers, such as BARON [6], have promising capabilities, but many realistic flowsheets remain computationally intractable at scale [1, 4]. We aim to develop a decomposition-based framework for tackling process flowsheet optimization problems, enabling its use within a global optimization framework.

BACKGROUND

Global optimization of process flowsheets with high-fidelity models remains a major open challenge in process systems engineering: such problems are typically large-scale, nonconvex, and computationally intractable for general-purpose solvers. In this work, we demonstrate a framework for globally optimizing complex process flowsheets by decomposing them spatially into more manageable subproblems. We solve the decomposed problem using a prioritized spatial-branch-and-bound algorithm implemented within the Structured Nonlinear Global Decomposition (SNoGloDe) framework [5]. The Institute for the Design of Advanced Energy Systems (IDAES) provides high-fidelity equation-oriented flowsheet models [3]. Utilizing the extensive process modeling framework in IDAES, in conjunction with the solution strategies developed in SNoGloDe, demonstrates the potential to obtain globally optimal solutions to critical process systems problems that are intractable with widely available off-the-shelf solvers, such as BARON [6]

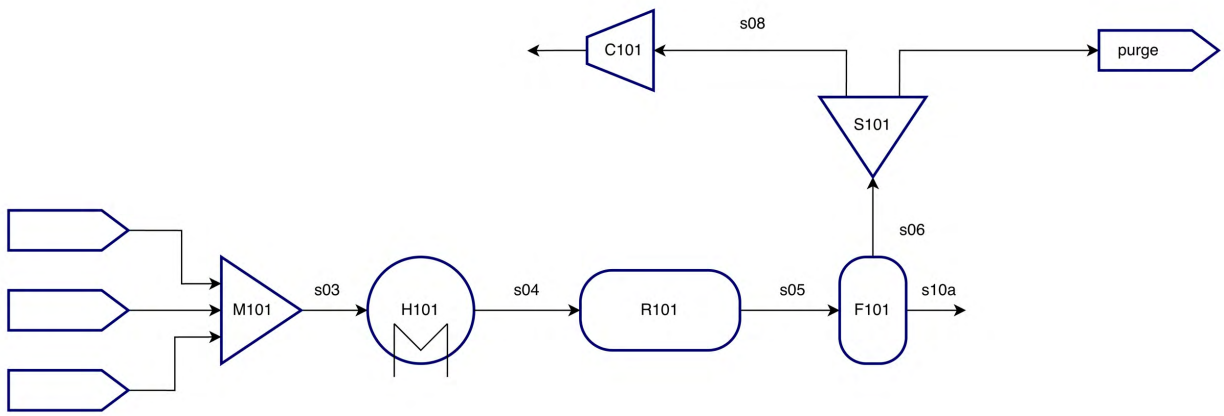

Figure 1: Complete flowsheet for hydrodealkylation process [3]. Three feed streams (toluene feed, hydrogen feed, and combined vapor feed) are combined in a mixer (M101), then heated in a heat exchanger (H101). Reactor (R101) is a CSTR, and its product is fed to a flash (F101). The resulting liquid stream (s10a) is the product, and the remaining vapor stream (s06) is fed to a splitter (S101) sent to purge or to a compressor (C101).

or Gurobi [2]. To evaluate this approach on an industrially relevant and economically significant problem, we select the hydrodealkylation (HDA) process for benzene production as our case study.

Benzene is a highly versatile aromatic raw material used in the production of ethylbenzene and cumene, which are feedstocks for producing styrene and phenol [7], cyclohexene [8], and chlorobenzene derivatives [9]. The major industrial sectors that use these products include the chemical, automotive, pharmaceutical, consumer products, and renewable energy sectors [10-12]. The demand for benzene and its derivatives is projected to grow rapidly between 2024 and 2030, driven by growth in the industry sectors [13]. According to the market research report, the market size was valued at 39.30 billion USD in 2023, with a compound annual growth rate of 6.5%, and is expected to reach 61.07 billion USD by 2030.

The primary process for producing benzene is hydrodealkylation, in which toluene is reacted with hydrogen to produce the simpler aromatic hydrocarbon benzene (without functional groups) and methane as a byproduct [14]. Recent work by Yalçın (2025) compared HDA with catalytic reactions – hydrocracking, isomerization, dehydrocyclization, and dehydrogenation - versus the thermic reaction [15]. While catalytic reactions can produce commercially relevant benzene and its derivatives, the process's operational conditions, energy intensity, and catalyst degradation outweigh the economic viability. As a result, recent optimization advances aimed at enhancing heat and energy integration of the HDA thermic reaction have shown promising results in achieving yield, energy, and cost reductions [14-16]. A model of this process serves as the case study for our global optimization approach based on spatial decomposition.

HDA FLOWSHEET MODEL

For our case study, we adapt a hydrodealkylation flowsheet from the Institute for Design of Advanced Energy Systems (IDEAS) example library [3] and modify its objective function to maximize revenue. This process converts toluene and hydrogen into benzene; the reaction is described in Equation 1.

$$C_6H_5CH_3 + H_2 \longrightarrow C_6H_6 + CH_4 \tag{1}$$

Figure 1 shows the complete flowsheet, where three feed streams (hydrogen feed, toluene feed, and a mixed vapor stream) are combined in a mixer (M101). The resulting stream is fed into a heater (H101) and then reacted in a continuous stirred tank reactor (CSTR) (R101). The reactor outlet is flashed in F101, yielding a liquid product stream primarily composed of benzene and a vapor stream containing the remaining hydrogen, toluene, methane, and some benzene. The vapor stream is split (S101) and either purged or sent to a compressor (C101). In future work, the compressor outlet will be recirculated to the mixer (M101); in this study, the compressor outlet is treated as an additional vapor waste stream.

The following subsections summarize the features of each unit operation.

Mixer (M101)

Three feed streams, fresh toluene, hydrogen, and a vapor recycle stream, are fed to the mixer unit and combined into a single outlet stream. As no recycle is implemented in this case study, the composition of the vapor recycle feed stream is fixed. Mixer units have zero degrees of freedom, as all outlet conditions are fully determined by the inlet streams and balances [3].

The mixer maintains all mass and energy balances. Energy conservation is satisfied by calculating the outlet stream temperature such that enthalpy is conserved among inlet streams that may enter at different temperatures. Vapor stream enthalpies are calculated using Equation 2.

$$h_{\mathrm{Vap},i} = \Delta H_{\mathrm{vap},i} + \frac{c_{p,5,\,\mathrm{Vap},i}}{5}\left(T^5 - T_{\mathrm{ref}}^5\right) + \frac{c_{p,4,\,\mathrm{Vap},i}}{4}\left(T^4 - T_{\mathrm{ref}}^4\right) + \frac{c_{p,3,\,\mathrm{Vap},i}}{3}\left(T^3 - T_{\mathrm{ref}}^3\right) + \frac{c_{p,2,\,\mathrm{Vap},i}}{2}\left(T^2 - T_{\mathrm{ref}}^2\right) + c_{p,1,\,\mathrm{Vap},i}\left(T - T_{\mathrm{ref}}\right) \quad (2)$$

Where $h_{Vap,i}$ is the component enthalpy for species i, the c_p's are specific heat coefficient values, $\Delta H_{\mathrm{vap},i}$ is the heat of vaporization for component i, T_{ref} is a reference temperature (298.15 K), and T is the temperature of the stream. Liquid-stream enthalpies are calculated similarly but do not include the heat of vaporization term.

Heater (H101)

The heater raises the temperature of the mixer outlet stream to the level required for the downstream reaction. The heater model in IDAES has one degree of freedom, typically specified by the outlet temperature or the heat duty [3]. All mass and energy balances are enforced, and there is no pressure drop across the heater.

Continuous Stirred Tank Reactor (R101)

The reactor in this flowsheet is modeled as a continuous stirred-tank reactor. The unit has one degree of freedom, namely the reactor volume and conversion, which impact capital cost (see Reactor Subproblem section) [3]. The rate constant of the reaction is governed by the Arrhenius equation. The reactor enforces all mass and energy balances. Since the reaction between toluene and hydrogen is exothermic, the reactor outlet stream is at a higher temperature than its inlet.

Flash Unit (F101)

The flash unit performs a vapor-liquid equilibrium separation of the reactor outlet stream, enabling recovery of benzene from the remaining components. In general, the flash unit has two degrees of freedom, typically either the heat duty or the pressure drop across the unit [3]. The flash is isobaric in this case study, leaving the heat duty as a degree of freedom. The phase equilibrium is determined using Antoine's equation, along with Raoult's law and the energy balance. In this model, toluene and benzene are treated as condensable species that may enter the liquid phase, while hydrogen and methane are assumed to be non-condensable and remain entirely in the vapor phase after the flash.

Splitter (S101)

The splitter divides the vapor outlet stream of the flash unit into two streams: a purge stream and a vapor-recycle stream. The IDAES splitter model has one degree of freedom, defined by the split fraction between the outlet streams [3]. Although the outlet stream is labeled vapor recycle, it is not returned to an upstream unit in this case study. The splitter enforces overall mass balance, with each stream receiving a fixed fraction of the inlet flow. Outlet streams inherit the same temperature, pressure, and composition as the inlet stream.

Compressor (C101)

The compressor is included in the flowsheet to account for potential pressure changes in the vapor-recycle stream. This flowsheet uses the Pressure Changer model from the IDAES library, which has one degree of freedom in this application, the outlet pressure.

SOLUTION METHODOLOGY

SNoGloDe is a framework (built in Python relying on the algebraic modeling language Pyomo [7]) that exploits the decomposable structure of an optimization problem and integrates it into a priority spatial branch-and-bound algorithm to efficiently solve large-scale optimization problems to global optimality [5]. To use SNoGloDe, the optimization formulation must contain variables that couple otherwise separable subproblems, which are first reformulated as complicating constraints to enable decomposition. This transformation allows the overall problem to be expressed as a collection of subproblems linked through a shared set of complicating variables. The resulting structure is particularly amenable to algorithms that exploit this decomposability to solve larger problems by decomposing them into significantly smaller, more manageable subproblems. SNoGloDe exploits this structure via a prioritized branching strategy and strategically constructed lower and upper bounds.

Algorithm

The underlying solution strategy is based on spatial branch-and-bound, a global optimization approach that explores the feasible region by recursively partitioning variable domains. Each node in the search tree corresponds to a restricted domain for the decision variables and is associated with both a lower bound, obtained from a relaxation, and an upper bound, obtained from a feasible solution. Nodes are pruned if they are infeasible or if their lower bound exceeds the best known upper bound. As the algorithm progresses, global bounds tighten, and the optimality gap decreases until a specified tolerance

is met. This process guarantees convergence to a globally optimal solution under standard assumptions [5].

To exploit the decomposable structure of the reformulated problem, a priority spatial branch-and-bound scheme is used. Rather than branching on *all* optimization variables, the search tree is restricted to the *complicating* variables only, while subproblem-specific variables are handled implicitly within each node. At a given node, the lower bound is computed by dropping the complicating constraints and solving each subproblem independently to global optimality with respect to the current bounds on the complicating variables, and then aggregating the resulting objective values. This ensures valid global lower bounds without requiring additional relaxations, although such enhancements can be incorporated if desired. Importantly, global convergence is preserved because all variables are ultimately optimized to global optimality, either through the branch-and-bound tree or within the subproblem solvers.

Upper bounds are generated by fixing a candidate assignment of the complicating variables and solving each subproblem independently to obtain a feasible solution. If all subproblems are feasible, their combined objective value provides a valid upper bound for the node. When a node cannot be pruned, branching is performed by selecting a complicating variable and splitting its domain, creating two child nodes that are added to the search queue. Throughout the algorithm, node-level computations are separable across subproblems and can therefore be parallelized. The method terminates when the optimality gap is sufficiently small or no nodes remain, at which point the best feasible solution found is returned.

Algorithm 1: Priority Spatial Branch-and-Bound (SNoGloDe)

In: Subproblems $\{P_1, \ldots, P_k\}$, complicating variables $z \in Z^0$, tolerance ε

Out: Globally optimal solution z^*, x^*

1. Initialize: $UB \leftarrow +\infty$, $LB \leftarrow -\infty$, queue $Q \leftarrow \{\text{root node}\}$
2. **while** Q not empty and $(UB - LB) > \varepsilon$ **do**
   3. Select node N from Q *[customizable: node selection]*
   4. Compute lower bound by relaxing complicating constraints and solving each subproblem P_k independently: $LB_n \leftarrow \sum_k \min f_k(x_k, z)$ s.t. $x_k \in X_k$, $z \in Z_n$ *[customizable: relaxation refinement]*
   5. **if** $LB_n \geq UB$ **then** prune N and continue
   6. Select candidate solution, fix values for complicating variables z *[customizable: incumbent generation]*
   7. Compute upper bound by solving each P_k with z fixed: $UB_n \leftarrow \sum_k \min f_k(x_k, \hat{z})$ s.t. $x_k \in X_k$ (if all P_k feasible)
   8. **if** $UB_n < UB$ **then** $UB \leftarrow UB_n$; store best solution
   9. Select $z_i \in z$ to branch on and split point s *[customizable: branching variable & split point]*
   10. Add child nodes N^+ ($z_i \leq s$) and N^- ($z_i \geq s$) to Q
11. **end while**
12. **return** best feasible solution

Algorithm Customization via SNoGloDe

The algorithmic framework SNoGloDe implements the priority spatial branch-and-bound approach described above. Within this setting, the proposed approach is intentionally designed to be highly customizable. Rather than prescribing a fixed set of heuristics, the package exposes multiple decision points throughout the branch-and-bound process at which problem-specific knowledge can be incorporated. These include strategies for node selection, refinement of lower-bound relaxations, generation of incumbent solutions, selection of branching variables, and determination of split points. Each of these choices directly influences how the search tree evolves, affecting both solution quality and computational efficiency.

The flexibility of this design is critical, as the performance of prioritized branch-and-bound methods is driven primarily by heuristics tailored to the specific problem. Intelligent node selection can dramatically reduce the size of the search tree by prioritizing regions with higher potential for bound improvement. Similarly, tighter relaxations enable earlier pruning, whereas effective incumbent generation accelerates convergence by improving global upper bounds more quickly. Branching and partitioning strategies further shape the combinatorial and numerical properties of the tree. By allowing these components to be tailored to the structure of a given application, the framework supports a wide spectrum of problem-aware strategies, transforming a generic decomposition into a tailored algorithm.

Allowing models to be expressed in Pyomo [7] ensures broad compatibility with nonlinear, mixed-integer, and nonconvex formulations and enables seamless interaction with global and local solvers. The implementation further supports parallel execution via an MPI-based workflow, enabling efficient distribution of subproblem solves and bound computations. To facilitate customization, SNoGloDe provides base classes with virtual methods that map directly to the key algorithmic decision points, making it straightforward for users to experiment

with and deploy customized heuristics.

HDA FLOWSHEET DECOMPOSITION

To solve using SNoGloDe, the hydrodealkylation flowsheet was partitioned into two subproblems. The first subproblem included the feed streams, the mixer (M101), the heat exchanger (H101), and the reactor (R101). The second subproblem included the flash (F101), the splitter (S101), and the compressor (C101). The first subproblem was called the reactor subproblem, and the second subproblem was called the flash subproblem. The reactor outlet temperature and the component molar flow rates were treated as complicating variables, meaning that their equality was enforced when generating the upper bound and candidate solution, but relaxed when determining the lower bound for a given node. This resulted in nine complicating variables (4 components, 2 phases, and 1 temperature). This decomposition scheme was chosen to place the two most complex unit operations, R101 and F101, in separate subproblems. Each subproblem has its own objective function and can be solved to global optimality using Baron [6] or Gurobi [2].

Reactor Subproblem

The model for the reactor subproblem included all equations and constraints that govern the unit operations and flows among the feeds, mixer, and reactor. The objective was to minimize the capital and operating costs. The costing equations were adopted from [3]. The heat exchanger outlet temperature (affecting the heat duty) and the reactor volume are both free variables. The annualized operating cost of the heat exchanger was found using Equation 3.

$$\text{Operating Cost} = (3600 \times 24 \times 365) \times \text{Heating Cost} \quad (3)$$

The heating cost for the heat exchanger was determined using Equation 4.

$$\text{Heating Cost} = 2.2 \times 10^{-7} \times Q_{\text{heater}} \quad (4)$$

Where Q_{heater} is the heat duty of the H101. The capital cost for the CSTR was calculated using Equation 5.

$$\text{Capital Cost} = 1.0 \times 10^{5} \times V \quad (5)$$

Where V is the reactor volume in cubic meters. The objective function for the reactor subproblem is found in Equation 6.

$$\text{min.} \quad (\text{Operating Cost} + \text{Capital Cost}) \times 10^{-6} \quad (6)$$

The objective function was scaled by 10^{-6}.

Flash Subproblem

The subproblem model for the flash included all equations and constraints that govern the unit operations and flows among the flash, splitter, and compressor. The objective for the flash subproblem was to maximize profit by minimizing negative profit. The costing equations were adopted from [3]. The free variable was the flash outlet temperature, which impacts the heat duty of the flash and the amount of liquid benzene generated. Annual revenue was calculated using Equation 7.

$$\text{Revenue} = F^{\text{out}}_{\text{Liq, benzene}} \times (3600 \times 24 \times 365) \times 78.11 \times 0.82 \times 10^{-3} \quad (7)$$

The annualized operating cost for the flash was found using Equation 8.

$$\text{Operating Cost} = (3600 \times 24 \times 365) \times \text{Cooling Cost} \quad 8)$$

The cooling cost for the flash was determined using Equation 9.

$$\text{Cooling Cost} = 0.25 \times 10^{-7} \times (-Q_{Flash}) \quad (9)$$

Where Q_{Flash} is the heat duty of the F101. It was negated because the flash is cooled. The objective function for the reactor subproblem is reported in Equation 10.

$$\text{min} \quad -(\text{Revenue} - \text{Operating Cost}) \times 10^{-6} \quad (10)$$

Note that capital cost is not included for the flash unit since the vessel cost is assumed to be much less than the reactor cost. Similarly to the reactor subproblem, the objective was scaled by 10^{-6}.

RESULTS

As highlighted in the Methodology section, a significant advantage of SNoGloDe is its ability to customize multiple elements. For this problem, we developed a warm-start capability at the lower bound (first solving a reduced square problem with IPOPT to obtain strong initial values for process variables and then solving globally with Gurobi [2]). This approach improved the convergence reliability for Gurobi on this NLP. To produce upper bounds, we used the built-in option to solve the monolithic problem *locally* with IPOPT and used the solutions for the complicating variables as the candidate solution for that upper-bound iteration. For branching, we selected variables using a Pseudocost methodology (also built in) and a worst-bound node queue. The results for the flowsheet optimization are summarized in Table 1 and shown in Figure 2.

SNoGloDe reached the globally optimal objective value with a 0.6% gap. Note that the monolithic model was also solved using Baron to confirm global optimality; both approaches yielded the same solution. This approach demonstrates a method for decomposing a nonlinear, nonconvex process flowsheet.

This model and solution are available as an example problem in the SNoGloDe code base.

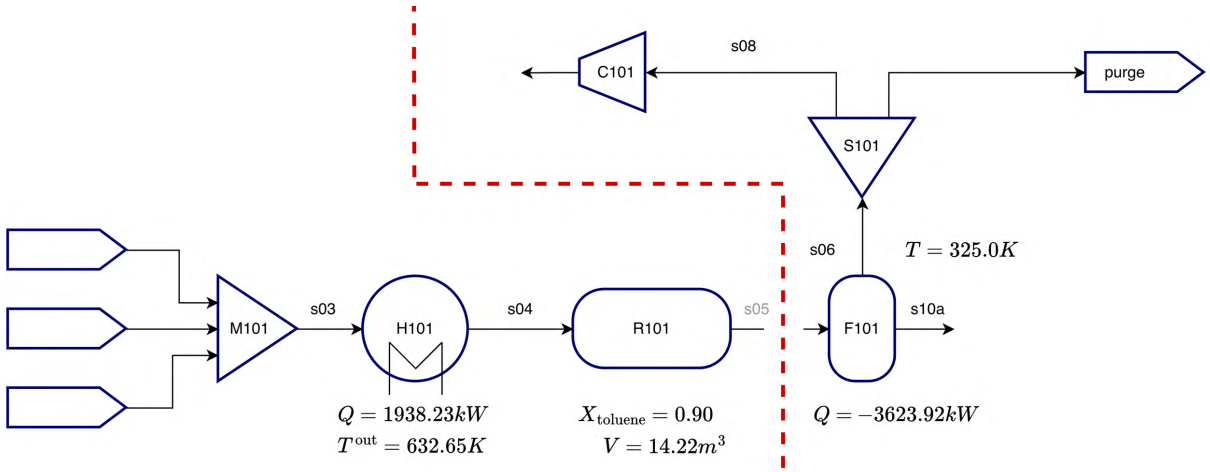

Figure 2. Summary of solution for the decomposed flowsheet. Values are summarized in Table 1. The first subproblem is the feed streams, mixer (M101), heat exchanger (H101), and the reactor (R101). The second subproblem is the flash (F101), splitter (S101), and compressor (C101). The molar flowrate for each component and the temperature of the reactor outlet are taken as complicating variables.

Table 1: SNoGloDe results for HDA flowsheet

Variable	Value	Units
Profit	9.96	M$/year
H101 Outlet Temperature	632	K
H101 Heat Duty	1, 940	kW
Reactor Volume	14.2	m^3
Toluene Conversion	0.90	
F101 Outlet Temperature	298	K
F101 Heat Duty	-3, 620	kW

The resulting costs for the flowsheet are reported in Table 2.

Table 2: Costing Results for HDA Process

Item	Value	Units
H101 Operating Cost	13, 447, 000	$/yr
F101 Operating Cost	2, 857, 000	$/yr
R101 Capital Cost	1, 421, 000	$
Annual Revenue	27, 687, 000	$/yr
Net Profit	9, 961, 000	$/yr

CONCLUSIONS

Our approach and case study illustrate how decomposition-based global optimization can provide results for complex, integrated flowsheet models. Beyond this example, integrating SNoGloDe with IDAES enables the extension of these methods to large-scale industrial challenges. By combining rigorous process modeling with a solution framework that yields globally optimal solutions, we present an alternative approach to solving process-system flowsheet optimization problems. Future work includes connecting the recycle between the compressor outlet and the mixer inlet.

AUTHOR IDENTIFIERS

Leppla M: 0009-0001-1295-5571
Stinchfield G: 0009-0004-5130-5661
Tran N: 0009-0007-2932-7786
Laird CD: 0000-0001-8430-1561

REFERENCES

1. Dowling AW, Biegler LT. A framework for efficient large scale equation-oriented flowsheet optimization. Computers & Chemical Engineering 72:3-20 (2015). https://doi.org/10.1016/j.compchemeng.2014.05.013
2. Gurobi Optimization LLC. Gurobi Optimizer reference manual. (2024).
3. Lee A, Ghouse JH, Eslick JC, Laird CD, Siirola JD, Zamarripa MA, Gunter D, Shinn JH, Dowling AW, Bhattacharyya D, Biegler LT, Burgard AP, Miller DC. The idaes process modeling framework and model library—flexibility for process simulation and optimization. J Adv Manuf & Process 3: (2021). https://doi.org/10.1002/amp2.10095
4. Ryoo HS, Sahinidis NV. Global optimization of nonconvex nlps and minlps with applications in process design. Computers & Chemical Engineering 19:551-566 (1995). https://doi.org/10.1016/0098-1354(94)00097-2
5. Stinchfield G, Walsh M, Cao Y, Laird CD. SNoGloDe: a flexible tool for decomposition and global

optimization. **Stockholm Glob Optim Workshop (STOGO)** (2025).
6. Zhang Y, Sahinidis NV. Solving continuous and discrete nonlinear programs with BARON. Comput Optim Appl 92:1123-1161 (2024). https://doi.org/10.1007/s10589-024-00633-0
7. Franck HG, Stadelhofer JW. Production and uses of benzene derivatives. Industrial Aromatic Chemistry :132-235 (1988). https://doi.org/10.1007/978-3-642-73432-8_5
8. Smith JA, Wilson KB, Sonstrom RE, Kelleher PJ, Welch KD, Pert EK, Westendorff KS, Dickie DA, Wang X, Pate BH, Harman WD. Preparation of cyclohexene isotopologues and stereoisotopomers from benzene. Nature 581:288-293 (2020). https://doi.org/10.1038/s41586-020-2268-y
9. Kulprathipanja S, Rekoske JE, Wei D, Slone RV, Pham T, Liu C. Modern petrochemical technology. Wiley (2021). https://doi.org/10.1002/9783527818167
10. Gentry JC. Benzene production and economics: a review. Asia-Pacific J Chem Eng 2:272-277 (2007). https://doi.org/10.1002/apj.18
11. Shanmugaraj K, Bustamante TM, Campos CH, Torres CC. Liquid phase hydrogenation of pharmaceutical interest nitroarenes over gold-supported alumina nanowires catalysts. Materials 13:925 (2020). https://doi.org/10.3390/ma13040925
12. Shewchuk SR, Mukherjee A, Dalai AK. Selective carbon-based adsorbents for carbon dioxide capture from mixed gas streams and catalytic hydrogenation of CO2 into renewable energy source: a review. Chemical Engineering Science 243:116735 (2021). https://doi.org/10.1016/j.ces.2021.116735
13. Benzene and its derivatives market—global industry analysis and forecast (2024–2030). **Maximize Market Research** (2024).
14. Rani J, Thakur P, Majumder S. Energy integration of hydrodealkylation (HDA) of toluene. Materials Today: Proceedings 111:60-68 (2024). https://doi.org/10.1016/j.matpr.2023.11.044
15. Yalç?n A. The techno?economic analysis and description of the hydrodealkylation process for benzene production from toluene. International Journal of Energy Research 2025: (2025). https://doi.org/10.1155/er/2393398
16. Dani S, Atqiyani AAT, Isna AA, Berliana Y, Maulana VR. Process intensification of hydrodealkylation (HDA) for benzene production through heat integration and gas recycle optimization. J. Chem. Eng. Res. Prog. 2:184-192 (2025). https://doi.org/10.9767/jcerp.20416

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Superstructure Framework for Feasibility and Flexibility Analysis Methods in Modular Plant Design

Julian Pamperin[a], Jonathan Mädler[b], Amy Koch[a], Isabel Viedth[a], and Leon Urbas[c]

[a] Dresden University of Technology, Process-to-Order Group, Dresden, Saxony, Germany
[b] FLEXCELERATE Solutions GmbH, Dresden, Saxony, Germany
[c] Dresden University of Technology, Chair of Process Control Systems, Dresden, Saxony, Germany
* Corresponding Author: julian.pamperin@mailbox.tu-dresden.de.

ABSTRACT

Modular plant design requires assessing whether independently characterized process requirements and module capabilities are compatible—a challenge that established methods address incompletely. Feasibility and flexibility analysis, as well as Quality by Design, typically assume integrated single-domain models where all variables belong to one coherent description, yet modular design involves domains that originate from different sources, evolve independently, and connect through interface variables. This work proposes Quantified Constraint Satisfaction Problems (QCSPs) as a formulation for interface-level suitability assessment: universal quantification encodes properties that must hold across their entire admissible range (e.g., physical properties, uncertain or environment-dependent characteristics requiring robustness), while existential quantification encodes variables where at least one feasible value must exist (e.g., critical process parameters, control inputs, configuration options). By operating over domain interfaces rather than requiring integrated system models, the formulation accommodates heterogeneous information sources within a common problem specification, with uncertainty handled through probabilistic intervals governed by a global confidence threshold. The framework supports both finite-domain solving on pre-evaluated datasets and numeric solving through on-demand model evaluation. A bioreactor suitability case study provides proof-of-concept, illustrating how the framework accommodates realistic problem complexity and offering initial observations on method selection trade-offs.

Keywords: Process Design, Design Under Uncertainty, Information Management, Interdisciplinary, Modelling and Simulations, Optimization

INTRODUCTION

Modular plant design has emerged as a paradigm for flexible, reconfigurable production systems in process industries facing volatile markets or rapid scale-up needs [1, 2]. By composing standardized, pre-engineered equipment units rather than designing monolithic facilities [3, 4], modular approaches have the potential to reduce engineering effort and support faster time-to-market [5]. However, this introduces a fundamental challenge: assessing whether a given module is suitable for a specific process before committing to detailed engineering [6, 7].

Two established methods offer relevant foundations. Feasibility and flexibility analysis evaluates whether process designs can maintain feasible operation under uncertainty [8, 9]. A key aspect within Quality by Design is the design space: the region of input variables and process parameters within which product quality is reliably met [10–12]. Despite their complementary strengths, neither directly addresses the module selection problem: determining whether independently characterized process requirements and module capabilities are compatible. Both typically assume integrated single-domain models, yet modular design involves domains that originate from different sources.

Bridging this gap requires confronting several challenges. First, process requirements and module capabilities may be characterized using different model types, information sources, and uncertainty representations.

Second, the interface between domains involves shared variables where semantic alignment introduces uncertainty. Third, practical suitability assessment must accommodate varying levels of detail: coarse screening across large module libraries demands different methods than detailed verification of selected candidates.

This work proposes a framework that addresses these challenges by reformulating suitability assessment as a Quantified Constraint Satisfaction Problem (QCSP) operating over domain interfaces. The framework provides a unified problem description capable of integrating diverse information types, systematic handling of uncertainties through interval specifications, and scalable assessment procedures. The following sections describe the theoretical foundations, present solution procedures, and demonstrate the framework through a case study in bioprocess equipment selection.

CONCEPTUAL FOUNDATIONS

Assessing module suitability requires clarity on what is being compared and at what level of abstraction. This section establishes the conceptual structure: domains as bodies of knowledge, their diverse characterizations, and approaches for comparison.

Domains and Abstractions

A domain denotes a bounded, abstract description space in which variables, constraints, and relations are defined—a consistent viewpoint independent of physical realization. In module-based plant design, a separation is typically made between the (equipment-) module domain and the process domain [13, 14].

Each domain may be characterized through diverse means—mechanistic models, experimental data, datasheets, simulation results, or expert knowledge—varying in fidelity, formalism, and the phenomena they capture. Crucially, characterizations from different domains are often not directly comparable: a detailed dynamic model of a bioreactor and a Quality-by-Design design space for a cultivation process do not share a common language.

Two strategies exist for enabling evaluation across domains. The first connects domain characterizations directly—for instance, coupling dynamic process and equipment models—and derives abstractions that evaluate the joint system's suitability for a given goal. This integrated approach permits analytical depth but introduces complexity: proprietary interfaces must propagate uncertainties across subsystems, and metrics must address the combined system. The second strategy, emphasized here, derives abstractions from individual domain characterizations independently, then assesses compatibility at the interface level. From process characterizations, requirements are derived (e.g., operating ranges, quality specifications); from module characterizations, capability metrics are extracted (e.g., operational envelopes). These transformations require analytical effort but yield quantities that are comparable without coupling the underlying characterizations, sacrificing analytical depth for manageability and accepting conservatism in exchange for modularity.

In practice, process and equipment characterizations are rarely created independently: experimental data emerges from specific equipment configurations, while equipment specifications are refined based on target processes. The framework accommodates such entanglement through uncertainty specifications on interface variables, allowing imperfect separation to be represented explicitly rather than assumed away.

Assessment Scope

Adopted from the concept of intra-modular safety in modular safety engineering [15, 16], this work introduces *intra-modular suitability*: the suitability of an individual module for a given process, assessed via boundary variables that characterize its interface with the environment. A module is considered intra-modularly suitable if its capability envelope covers the corresponding process requirements across all relevant boundary variables. If intra-modular suitability is established for all modules within a configuration, system-wide suitability can be inferred without further analysis.

CSP FRAMEWORK FOR SUITABILITY ASSESSMENT

The preceding conceptual foundations established the domains and the scope of the suitability assessment considered in this work. This section introduces a formal framework that operationalizes these concepts using QCSPs [17–19]. The following section relates the framework to feasibility analysis and Quality by Design, presents the advantages of the CSP-based formulation, provides the formal QCSP definition with uncertainty handling, and defines satisfaction semantics.

Relation to Feasibility Analysis and Quality by Design

The suitability assessment problem shares foundational objectives with feasibility and flexibility analysis as well as Quality by Design: ensuring constraint satisfaction under uncertainty through appropriate use of degrees of freedom. Classical feasibility and flexibility analysis evaluates whether feasible operation can be maintained for all uncertain parameters θ within a specified set T through appropriate adjustment of control variables z given constraints $g(d, z, x, \theta) \leq 0$ [8, 20, 21]. Quality by Design similarly distinguishes between material attributes and process parameters that define a design space

within which product quality is assured [10, 11]. Both frameworks implicitly encode a quantifier structure analogous to that of QCSPs: in flexibility analysis, uncertain parameters θ correspond to universal quantification $\forall$ (feasibility must hold across their entire range), while control variables z correspond to existential quantification $\exists$ (finding any feasible adjustment suffices). In Quality by Design, Critical Material Attributes (CMAs) typically relate to universal quantification as properties that must be tolerated across their specified ranges, while Quality Target Product Profiles (QTPPs), Critical Quality Attributes (CQAs), and Critical Process Parameters (CPPs) relate to existential quantification as targets where achieving any value within acceptable bounds suffices.

However, neither method directly addresses the matching problem in module-based plant design. Both typically require an integrated model relating all variables within one coherent description. Performing a feasibility test for a given process-module combination therefore requires semantic alignment: the meaning encoded in design space characterizations must be transferred into a problem structure compatible with flexibility analysis. This translation is non-trivial and may need to be repeated for each combination, as these approaches are commonly implemented ad hoc rather than integrated into information model standards. While recently established standards such as VDI/VDE/NAMUR 2658 (Module Type Package) [22] or DEXPI [23] provide structured representations of equipment and process data, they do not prescribe a formal problem structure for suitability assessment. An interface language is therefore needed that encodes the semantics of feasibility and flexibility analysis and Quality by Design into a common problem specification.

Advantages of CSP-Based Formulation

QCSPs address this by providing a *common interface language* for independently characterized domain constraints. Each domain translates its characterization into QCSP-compliant interval specifications; matching then operates on these specifications without requiring knowledge of underlying models.

This interface-language property derives from several characteristics of the CSP formalism. First, CSP formulations treat all variables uniformly without requiring explicit differentiation between inputs, outputs, or other dependency structures. This uniformity shifts modeling complexity into the domain descriptions themselves, decoupling problem specification from solution strategy. The framework thus accommodates heterogeneous information sources—mechanistic models, empirical correlations, experimental datasets, manufacturer specifications—without prescribing how these sources internally relate variables.

Second, while the quantifier structure underlying feasibility and flexibility analysis and Quality by Design remains implicit in their problem formulations, the QCSP formalism makes this structure explicit and systematic. Variables are declared with their quantifiers as part of the problem specification, enabling direct expression of which conditions must be tolerated versus which represent degrees of freedom. This explicitness facilitates communication across domain boundaries and supports automated reasoning about suitability.

Third, the formulation accommodates expressiveness beyond classical approaches. Quantifier semantics may depend on assessment context: a critical process parameter such as cultivation temperature is existentially quantified for single-product assessment (one feasible setpoint suffices) but universally quantified for multi-product flexibility (all temperatures in the range must be achievable). The QCSP formulation handles such context shifts through quantifier specification rather than problem reformulation. Furthermore, discrete variables representing categorical choices (equipment variants, operating modes) integrate naturally alongside continuous variables—a capability that flexibility analysis handles less directly.

Fourth, integrating uncertainty directly into constraint specifications enables cleaner problem descriptions and more generic uncertainty handling than approaches common in flexibility analysis or Quality by Design. The formulation presented below embeds probability distributions as first-class constraint elements, with a global confidence threshold governing the overall satisfaction guarantee.

The approach localizes transformation effort to domain boundaries: while defining the translation from domain-specific characterizations to QCSP specifications requires non-trivial effort, such effort is inherent whenever knowledge from different sources must be integrated. The formalization provides a uniform interface that, once established, simplifies subsequent usage across different assessment contexts. Beyond binary satisfaction, CSP-based analysis can yield information analogous to flexibility analysis, including identification of binding constraints, feasible region characterization, and flexibility quantification through systematic constraint bound variation [24].

Quantified Constraint Satisfaction Problem Formulation

The QCSP is defined as a tuple $\langle X, D, C \rangle$ where:

- $X = \{X_1, \dots, X_n\}$ is a set of variables
- $D = \{D_1, \dots, D_n\}$ where each D_i is the domain for variable X_i, which may be continuous ($D_i \subseteq \mathbb{R}$) or discrete ($D_i = \{v_1, \dots, v_k\}$)
- $C = \{C_I, C_F, C_D, C_\alpha\}$ is a structured set of constraints

Each constraint type serves a distinct role in the problem specification: The set C_I contains interval constraints that bound individual variables:

$$C_I = \{(Q_i, X_i, B_i) : i \in \mathcal{I}\} \tag{1}$$

where $Q_i \in \{\exists, \forall\}$ denotes the quantifier specifying whether satisfaction is required for *some* (∃) or *all* (∀) values of the variable, $X_i \in X$ is the constrained variable, and B_i specifies the constraint bound. The bound B_i may take the form of either a *deterministic interval* $[l_i, u_i]$ with $l_i, u_i \in \mathbb{R}$, or a *probabilistic specification* P_i, i.e., a probability distribution from which operational bounds are derived during solving. Deterministic intervals directly restrict a variable to a known range. Probabilistic specifications, by contrast, encode uncertain knowledge about the constraint—for example, a process viscosity requirement subject to measurement or modeling uncertainty, or an equipment limit known only within confidence bounds. The translation from a probability distribution to an effective interval depends on the global confidence threshold and the solving procedure, allowing the same problem specification to induce different operational constraints under different confidence requirements.

The set C_D restricts variables to specific discrete values:

$$C_D = \{(X_k, V_k) : k \in \mathcal{K}\} \tag{2}$$

where $V_k \subset D_k$ is a finite set of allowed values for variable X_k. These constraints represent categorical choices such as equipment variants, operating modes, or material selections.

The set C_F captures relational constraints among multiple variables:

$$C_F = \{(t_j, R_j) : j \in \mathcal{J}\} \tag{3}$$

where $t_j \subseteq X$ denotes the scope of the constraint, i.e., the subset of variables involved, and $R_j \subseteq \prod_{X_k \in t_j} D_k$ defines the relation specifying the allowed combinations of values over that scope. The relation R_j may be specified explicitly (e.g., as a lookup table or pre-evaluated dataset) or implicitly (e.g., through algebraic equations, differential equations, or trained surrogate models). This flexibility accommodates the diverse information types encountered in modular plant design.

For any variable $X_i \in X$, an optional associated distribution P_{X_i} may be specified, indicating that values of X_i are known only probabilistically. The impact of such variable-level uncertainty on constraint satisfaction is governed by the global confidence threshold.

The constraint C_α specifies the required confidence level for the overall problem:

$$C_\alpha = \alpha_{\text{global}} \in [0, 1) \tag{4}$$

This threshold defines an upper bound on the probability with which a declared solution may violate one or more constraints. A high confidence requirement (e.g., $\alpha_{\text{global}} = 0.01$) demands that uncertainties be treated conservatively, yielding solutions that are highly trustworthy but potentially restrictive. Larger values of α_{global} permit more optimistic interpretations of uncertain constraints, expanding the solution space. The global confidence constraint affects all probabilistic elements in the problem: interval constraints specified as distributions ($B_i = P_i$) and variables with associated uncertainty (P_{X_i}). Deterministic constraints remain unaffected. The specific mechanism by which α_{global} propagates to individual uncertain elements—for instance, through Bonferroni correction (as described in the case study) or more sophisticated allocation schemes—is a solver concern rather than a problem specification concern. This separation ensures that the problem formulation remains independent of the solution method.

Satisfaction Semantics

The QCSP is satisfied if a strategy σ exists for existentially quantified variables such that, for any assignment to universally quantified variables, all constraints are satisfied with significance at least α_{global}.

Formally, let $X_\forall = \{X_i : \exists (Q_i, X_i, B_i) \in C_I \text{ with } Q_i = \forall\}$ denote the universally quantified variables, and $X_\exists = X \backslash X_\forall$ the existentially quantified variables. The QCSP is satisfied if and only if:

$$\mathbb{P}\left[\exists \sigma : \prod_{X_j \in X_\exists} D_j \text{ s.t. } \forall \mathbf{x}_\forall \in \prod_{X_i \in X_\forall} \tilde{B}_i : (\mathbf{x}_\forall, \sigma(\mathbf{x}_\forall)) \vDash C\right] \geq 1 - \alpha_{\text{global}} \tag{5}$$

where $\tilde{B}_i$ denotes the effective interval for variable X_i (either the deterministic interval $[l_i, u_i]$ or an interval derived from P_i consistent with α_{global}), and $\vDash C$ indicates satisfaction of all constraints in C. When $\alpha_{\text{global}} = 0$ and all constraints are deterministic, this reduces to classical QCSP semantics. When only existential quantifiers are present ($X_\forall = \varnothing$), the problem simplifies to a standard CSP.

CASE STUDY

This section demonstrates the framework through a proof-of-concept assessment of bioreactor suitability for cell cultivation processes. First, the solving procedures are presented; then, these procedures are applied to four representative scenarios.

Solving Procedures

The QCSP formulation admits various solving strategies; the methods employed here prioritize simplicity to illustrate the connection between the problem formulation and established techniques. More sophisticated solvers from flexibility analysis or constraint programming can be substituted as problem complexity

increases, while the fundamental structure of QCSP solving remains unchanged.

As noted above, constraint relations may be specified through pre-evaluated datasets (finite-domain solving via direct lookup) or model structures requiring computation (numeric solving via on-demand evaluation). This distinction is decisive: these information sources entail fundamentally different solver strategies and different potentials for characterizing domain properties.

For uncertainty handling, a spectrum of approaches exists between repeated sampling from probability distributions followed by solving CSP instances, and conservative treatment through worst-case bounds on confidence intervals. The latter approach, employed here via Bonferroni correction, offers computational efficiency at the cost of conservatism. When multiple uncertain constraints must be satisfied jointly, the union bound controls overall failure probability: if α_{global} is the acceptable global significance level and N constraints carry uncertainty, allocating a budget α_i to each constraint such that $\sum_{i=1}^{N} \alpha_i = \alpha_{\text{global}}$ guarantees that the global failure probability remains below α_{global}. The simplest allocation divides α_{global} uniformly; more sophisticated schemes can assign stricter thresholds to critical constraints, enabling optimized prioritization. For a constraint with probabilistic specification (e.g., Gaussian with mean μ_i and standard deviation σ_i), the effective interval becomes $[\mu_i - z_{\alpha_i/2}\sigma_i, \mu_i + z_{\alpha_i/2}\sigma_i]$, where $z_{\alpha_i/2}$ is the critical value corresponding to the allocated local significance level. This transformation converts probabilistic constraints into deterministic intervals suitable for standard constraint checking, while preserving statistical guarantees on the overall solution.

For standard CSPs (only existential quantifiers), satisfaction requires finding at least one assignment satisfying all constraints. The solving procedure applies a search strategy appropriate to the information type: for finite-domain problems, constraint propagation successively filters candidate points, or backtracking examines candidates sequentially with early termination upon finding a valid assignment; for numeric problems, sampling-based search (e.g., random or structured sampling) generates input configurations for on-demand model evaluation. In both cases, uncertain constraints are checked against intervals derived from their associated distributions using the allocated local significance thresholds.

For QCSPs containing universal quantifiers ∀, direct verification would require checking infinitely many value combinations within ∀-intervals—an intractable task. The approach adopted here decomposes universal constraints into boundary evaluations: each ∀-constraint on interval $[l_i, u_i]$ generates sub-problems at the boundary values $\{l_i, u_i\}$, yielding $2^{|X_\forall|}$ standard CSP sub-problems. Under a continuity assumption—that feasible regions form convex sets such that boundary feasibility implies interior feasibility—satisfaction of all boundary sub-problems guarantee satisfaction of the original QCSP. This assumption holds for many physical systems but may yield false positives when solution spaces exhibit discontinuities or non-convexities; such cases require explicit mode enumeration or conservative interior sampling.

Algorithm 1 is a unified solving procedure. The algorithm accepts a QCSP specification and returns a satisfaction status. The SEARCH subroutine encapsulates information-type-specific strategies: for finite-domain problems, it implements constraint propagation or backtracking over pre-evaluated datasets; for numeric problems, it implements sampling-based search with on-demand model evaluation. The DECOMPOSE subroutine generates boundary sub-problems from universal constraints.

Algorithm 1: QCSP Solver

Input: QCSP $\langle X, D, C\rangle$ with $C = \{C_I, C_F, C_D, C_\alpha\}$
Output: Satisfaction status
$X_\forall \quad \{X_i : (Q_i, X_i, B_i) \in C_I \wedge Q_i = \forall\}$
$\mathcal{S} \leftarrow$ Decompose$(C_I, X_\forall)$ `// 2^|X∀|` `sub-problems (1 if` $X_\forall = \emptyset$`)`
Compute α_{local} from α_{global} $|\mathcal{S}|$, and uncertain constraint count;
Derive effective intervals $\tilde{B}_i$ from probabilistic specifications using α_{local}
foreach *sub-problem* $s \in \mathcal{S}$ **do**
  if SEARCH$(X, D, C, \tilde{B}^{(s)}) = \mathit{unsatisfied}$ **then**
    return `unsatisfied`
return `satisfied`

Bioreactor Suitability Assessment

The suitability assessment integrates a process domain characterized by Quality-by-Design-style cultivation requirements with a module domain described by a steady-state simulation model. Four representative scenarios, selected from a broader set of 18 assessment cases, illustrate the trade-offs between finite-domain and numeric solving while validating the framework's capability to handle problems of typical complexity in this domain. The remaining cases are omitted for brevity.

The bioreactor model represents a 10–200 L stirred-tank reactor typical of pilot-scale cell culture systems. The model accepts 12 input variables spanning thermal properties, equipment parameters, operational settings, and medium physical properties. Each output incorporates Gaussian uncertainty distributions with mechanistically derived mean values and standard deviations, enabling uncertainty propagation through the assessment. Despite the 12-dimensional input space, gradient-based principal component analysis revealed an effective dimensionality of approximately 5.4 dimensions, indicating significant redundancy due to variable correlations—a

Table 1: QCSP formulation for Case 1: CHO high cell density cultivation (20–30 × 10^6 cells/mL) with high tempering capability of the bioreactor.

Variable	Unit	Interval	Quant.
Process Requirements			
Medium Temperature	°C	[36.5, 37.0]	∀
Medium Volume	m³	[0.15, 0.17]	∀
Dynamic Viscosity	Pa·s	[0.001, 0.008]	∀
Medium Density	kg/m³	[1010, 1040]	∀
Specific Heat Capacity	J/(kg·K)	[3800, 3900]	∀
Thermal Conductivity	W/(m·K)	[0.58, 0.62]	∀
Eff. Shear Stress	Pa	[0.0, 1.4]	∃
Mixing Time	s	[0.0, 30]	∃
Min. Spec. Heat Flow	W/kg	[−732, −2.25]	∃
Equipment Specification			
Min. Jacket Temp.	°C	[15, 25]	∃
Max. Jacket Temp.	°C	[50, 70]	∃
Wall Therm. Resistance	m²·K/W	[0.0002, 0.0006]	∃
Jacket Heat Trans. Coeff.	W/(m²·K)	[400, 800]	∃
Fouling Resistance	m²·K/W	[0.00015, 0.0005]	∃
Max. Spec. Heat Flow	W/kg	[−203, 557]	∃
Agitation Speed	rpm	[20, 400]	∃

characteristic that influences data requirements for explicit capability characterization.

The cultivation scenarios were selected to represent diverse assessment outcomes: Case 1 with requirements well within equipment capabilities; Case 2 requiring significantly more data due to proximity to capability boundaries; Case 3 where requirements approach the equipment limits; and Case 4 representing genuine incompatibility. Table 1 shows the complete QCSP formulation for Case 1 as an exemplary problem specification, showing how cultivation requirements and equipment specifications are transformed into interval constraints with appropriate quantifiers. The six universally quantified variables—representing process conditions that must be satisfied across their entire specified range—yield $2^6 = 64$ subproblems per assessment. A global confidence threshold of $\alpha_{\text{global}} = 0.05$ governed uncertainty propagation via Bonferroni correction.

For finite-domain solving, a dataset of 5×10^6 uniformly random samples were generated across the 12-dimensional input domain. The solver employed hybrid constraint propagation and backtracking with early termination. The numeric solver implemented uniform random search with on-demand model evaluation. Both solvers were implemented in Julia and executed single-threaded on a Windows system with 64 GB RAM and Intel vPRO EVO Edition CPU.

The assessment outcomes are summarized in Table 2. Both methods correctly identified Case 1 and Case 2 as satisfiable and Case 4 as infeasible. However, for Case 3, the finite-domain approach produced a false negative: it failed to find satisfying assignments despite the numeric solver confirming that feasible operational configurations exist. This discrepancy arises because the requirements in Case 3 approach capability boundaries where the feasible region becomes too narrow for the pre-computed dataset to capture. The finite-domain method cannot detect feasible regions between sparse datapoints, regardless of total dataset size—a fundamental precision limitation when requirements lie near equipment limits.

The performance comparison for Cases 1 and 2 (see Table 3) reveals the trade-off structure between methods. Case 1, with requirements well within capabilities, achieved satisfaction with 5×10^5 datapoints in 5.2 s, while the numeric solver required only 64 model evaluations. Case 2, despite both methods succeeding, demanded 5×10^6 datapoints and 36.9 s for finite-domain solving—a tenfold increase in data requirements due to proximity to capability boundaries—whereas numeric solving required only 69 evaluations. The breakeven simulation time, at which both methods require equal computational effort, shifts from 0.081 s per evaluation (Case 1) to 0.535 s (Case 2). For models with evaluation times above these thresholds, pre-computed datasets offer superior performance for repeated assessments; below them, on-demand evaluation becomes more efficient. This breakeven analysis, combined with the precision limitations observed in Case 3, suggests a practical strategy: finite-domain methods provide efficient initial

Table 2: Suitability assessment results for four representative scenarios comparing finite-domain (pre-evaluated dataset) and numeric (on-demand evaluation) solving approaches. Temp. indicates the bioreactor's tempering capability configuration (high: extended jacket temperature range; low: restricted range).

Case	Cultivation	Temp.	Numeric	Finite-D.
1	CHO high cell dens.	high	satisfied	satisfied
2	E. coli low-temp expr.	high	satisfied	satisfied
3	E. coli low-temp expr.	low	satisfied	unsatisfied
4	P. pastoris methanol	low	unsatisfied	unsatisfied

Table 3: Computational performance comparison for cases where both methods found satisfaction. Datapoints: minimum dataset size required for finite-domain satisfaction; FD Time: finite-domain solving duration; Num. Evals: model evaluations required by numeric solver; Breakeven: simulation time per evaluation at which both methods require equal total computation time.

Case	Datapoints	FD Time (s)	Num. Evals	Breakeven (s)
1	5×10^5	5.20	64	0.081
2	5×10^6	36.92	69	0.535

screening when requirements lie safely within anticipated capabilities, while numeric methods become essential for boundary exploration. The framework's unified problem formulation enables seamless transition between these approaches as assessment demands evolve.

CONCLUSION

This work addressed a methodological gap in modular plant design: while feasibility and flexibility analysis provide rigorous frameworks for evaluating process robustness, and Quality by Design establishes systematic design space characterization, both traditions typically assume integrated single-domain models. Modular plant design, however, requires assessing compatibility between process requirements and module capabilities that originate from different sources, evolve independently, and connect through shared interface variables.

The proposed QCSP formulation addresses this gap primarily as a *formal interface language*. Rather than requiring integrated models that connect diverse domains manually—which demands detailed knowledge of both domain-specific model structures—the QCSP provides a common specification through which independently characterized domains can express their constraints. Each domain translates its characterization into QCSP-compliant interval specifications; the matching process then operates on these specifications without requiring knowledge of underlying model structures.

The framework accommodates heterogeneous information sources—mechanistic models, empirical datasets, expert-specified bounds—through uniform variable treatment, and supports both finite-domain solving on pre-evaluated datasets and numeric solving through on-demand model evaluation. The bioreactor case study provided proof-of-concept that the formulation can represent realistic suitability problems and that established solving techniques remain applicable.

Several limitations warrant emphasis. The case study employed an unvalidated first-principles model; industrial validation remains necessary. The solving procedures presented are illustrative rather than contributory—more sophisticated solvers from flexibility analysis or constraint programming literature can and should be substituted when problem complexity demands. Furthermore, while the framework provides mechanisms for uncertainty handling, systematic study of how uncertainty specifications affect assessment outcomes in practice was out of scope.

The breakeven thresholds (0.08–0.54 s per evaluation) illustrate trade-offs critical when assessing large module libraries. Separating problem specification from solution method enables staged workflows where inexpensive screening precedes costly verification. Capability characterizations generated during assessment can be stored and reused across future queries, amortizing effort and improving assessment efficiency over time.

OUTLOOK

The primary need is empirical validation: applying the framework to industrial module selection problems would test whether the formalization provides practical value beyond conceptual clarity. Integration with

industrial information model standards could enable automated translation from existing domain characterizations to QCSP specifications, reducing the manual effort currently required; relevant standards include the Module Type Package (VDI/VDE/NAMUR 2658) for modular automation [22], DEXPI (IEC 62424) for process equipment data [23], IEC 61512 for batch control models [25], and ICH Q8–Q12 for pharmaceutical design spaces [10]. Development of accessible tooling would facilitate transition from methodological contribution to practical engineering workflow. Finally, comparative studies against established flexibility analysis tools would help quantify whether the interface-language benefits outweigh the translation overhead in realistic engineering contexts.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the Process2Order Lab at TU Dresden for institutional support. This work is based on a diploma thesis funded and supported by FLEXCELERATE Solutions, whose input significantly contributed to its development.

REFERENCES

1. Baldea M, Edgar TF, Stanley BL, Kiss AA. Modular manufacturing processes: status, challenges, and opportunities. AIChE Journal 63:4262-4272 (2017). https://doi.org/10.1002/aic.15872
2. Becker T, Lier S, Werners B. Value of modular production concepts in future chemical industry production networks. European Journal of Operational Research 276:957-970 (2019). https://doi.org/10.1016/j.ejor.2019.01.066
3. Hohmann L, Kössl K, Kockmann N, Schembecker G, Bramsiepe C. Modules in process industry ? a life cycle definition. Chemical Engineering and Processing: Process Intensification 111:115-126 (2017). https://doi.org/10.1016/j.cep.2016.09.017
4. Lier S, Paul S, Ferdinand D, Grünewald M. Modulare verfahrenstechnik: apparateentwicklung für wandlungsfähige produktionssysteme. Chemie Ingenieur Technik 88:1444-1454 (2016). https://doi.org/10.1002/cite.201600015
5. O'Connor JT, Kowall CP, Haapala KR, Agrawal NV, Paul BK. Specialty chemicals production case study: economic analysis of modular chemical process intensification versus conventional stick?built approaches. J Adv Manuf & Process 3: (2021). https://doi.org/10.1002/amp2.10102
6. Schindel , Polyakova , Harding , Weinhold , Stenger , Grünewald , Bramsiepe . General approach for technology and process equipment assembly (PEA) selection in process design. Chemical Engineering and Processing - Process Intensification 159:108223 (2021). https://doi.org/10.1016/j.cep.2020.108223
7. Harding D, Polyakova M, Nowara D, Rech S, Grünewald M, Bramsiepe C. New process function-based selection and configuration methodology for process equipment assemblies (peas) exemplified on the unit operation distillation. Chemical Engineering and Processing - Process Intensification 168:108531 (2021). https://doi.org/10.1016/j.cep.2021.108531
8. Swaney RE, Grossmann IE. An index for operational flexibility in chemical process design. part I: formulation and theory. AIChE Journal 31:621-630 (2004). https://doi.org/10.1002/aic.690310412
9. Bansal V, Perkins JD, Pistikopoulos EN. Flexibility analysis and design using a parametric programming framework. AIChE Journal 48:2851-2868 (2004). https://doi.org/10.1002/aic.690481213
10. International Council for Harmonisation, "ICH Quality Guidelines Q8–Q12, " 2005–2019. Available: https://www.ich.org
11. Yu LX, Amidon G, Khan MA, Hoag SW, Polli J, Raju GK, Woodcock J. Understanding pharmaceutical quality by design. AAPS J 16:771-783 (2014). https://doi.org/10.1208/s12248-014-9598-3
12. Holden JM. Many lives, many masters, by brian L. weiss. J Near-Death Stud 11:57-61 (1992). https://doi.org/10.1007/bf01082738
13. Maedler, J., I. Viedt, J. Lorenz, and L. Urbas, "Requirements to a digital twin-centered concept for smart manufacturing in modular plants considering distributed knowledge, " in Computer Aided Chemical Engineering, vol. 49, Elsevier, 2022, pp. 1507–1512. doi: 10.1016/B978-0-323-85159-6.50251-7
14. Koch A, Hamedi N, Furtner L, Kock T, Klose A, Mädler J. Standards for information models considering knowledge distribution in modular plants. 2023 IEEE 21st International Conference on Industrial Informatics (INDIN) :1-7 (2023). https://doi.org/10.1109/indin51400.2023.10218218
15. Klose A, Merkelbach S, Menschner A, Hensel S, Heinze S, Bittorf L, Kockmann N, Schäfer C, Szmais S, Eckert M, Rüde T, Scherwietes T, da Silva Santos P, Stenger F, Holm T, Welscher W, Krink N, Schenk T, Stutz A, Maurmaier M, Stark K, Hoernicke M, Unland S, Erben S, Kessler F, Apitz F, Urbas L. Orchestration requirements for modular process plants in chemical and pharmaceutical industries. Chem Eng & Technol 42:2282-2291 (2019). https://doi.org/10.1002/ceat.201900298
16. VDI 2776-3:2024, Process engineering plants – Modular plants – Safety of modular plants, Jan. 2024.

17. Gent, I. P., P. Nightingale, and K. Stergiou, "QCSP-Solve: A Solver for Quantified Constraint Satisfaction Problems, " 2005.
18. Mamoulis N, Stergiou K. Algorithms for quantified constraint satisfaction problems. Lecture Notes in Computer Science :752-756 (2004). https://doi.org/10.1007/978-3-540-30201-8_60
19. Qureshi AJ, Dantan JY, Bruyere J, Bigot R. Set based robust design of mechanical systems using the quantifier constraint satisfaction algorithm. Engineering Applications of Artificial Intelligence 23:1173-1186 (2010). https://doi.org/10.1016/j.engappai.2010.02.003
20. Grossmann IE, Floudas CA. Active constraint strategy for flexibility analysis in chemical processes. Computers & Chemical Engineering 11:675-693 (1987). https://doi.org/10.1016/0098-1354(87)87011-4
21. Floudas CA, Gümü? ZH, Marianthi G Ierapetritou . Global optimization in design under uncertainty: feasibility test and flexibility index problems. Ind. Eng. Chem. Res. 40:4267-4282 (2001). https://doi.org/10.1021/ie001014g
22. VDI/VDE/NAMUR 2658-4:2022, Automation engineering of modular systems in the process industry – Modellierung von Moduldiensten, Oct. 2022.
23. Tolksdorf G, Cameron DB, Theißen M. DEXPI 2.0: synergistic integration of PFD and p&id in a unified digital model. Chemie Ingenieur Technik 97:1065-1069 (2025). https://doi.org/10.1002/cite.70009
24. Hu J, Aminzadeh M, Wang Y. Searching feasible design space by solving quantified constraint satisfaction problems. Journal of Mechanical Design 136: (2013). https://doi.org/10.1115/1.4026027
25. IEC 61512-1, Batch control – Part 1: Models and terminology, 1997

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Process Flowsheet Synthesis via Quantum Reinforcement Learning with Improved Scalability

Austin Braniff[a], Fengqi You[b], and Yuhe Tian[a]*

[a] Department of Chemical and Biomedical Engineering, West Virginia University, Morgantown, WV, United States
[b] Smith School of Chemical and Biomolecular Engineering, Cornell University, Ithaca, NY, United States
* Corresponding Author: yuhe.tian@mail.wvu.edu.

ABSTRACT

In this work, we present quantum reinforcement learning algorithms for process flowsheet synthesis. Particularly, we discuss the implementation of encoding strategies to improve the algorithmic scalability. Reinforcement learning (RL)-driven flowsheet synthesis techniques provide a promising approach for conceptual process design, in addition to traditional optimization-based methods. These RL-based strategies identify the optimal flowsheet configurations from a maximum set of available processing units, without requiring to pre-postulate an interconnected superstructure. However, the resulting combinatorial design space for RL can scale extensively with the increased number of available processing units, which can render the algorithms to be computationally intensive or even intractable. To address this challenge, our prior work has introduced a quantum-enhanced approach to RL-driven process synthesis. However, this algorithm was limited in its capacity to solve larger flowsheeting problems as a large number of qubits were needed. To reduce qubit requirements and improve scalability, this work presents two different encoding strategies to embed information into the gates of quantum circuits instead of embedding it into qubits alone. We demonstrate the efficacy of these algorithms on three different scenarios of a flowsheet synthesis problem and interpret the results obtained from the IonQ quantum computing platform.

Keywords: Process Synthesis, Process Design, Reinforcement Learning, Quantum Computing, Machine Learning

1. INTRODUCTION

There has been increasing interest in RL-driven process synthesis methods [1-3], which strive to determine the optimal design solutions while minimizing the impact of prior expert knowledge and the solution of large-scale combinatorial optimization algorithms (e.g., generalized disjunctive programming, mixed integer nonlinear programming). However, a significant challenge for these strategies is the large process design space which can be highly computationally intensive or even intractable for RL to identify the truly optimal design solution.

RL-driven process synthesis involves training an agent to build optimal flowsheet designs for a given process application. This is posed as a discrete decision-making process. The RL agent is usually represented via a neural network (NN). This parameterization allows a deep RL algorithm to adjust tunable parameters within the NN-based agent. In this way, the agent iteratively becomes better at developing process flowsheets. Several techniques have been developed by the research community in recent years [1]. These include graph-based [2] or matrix-based [3] flowsheet representations, leveraging transfer learning models instead of expensive simulators [4], and simultaneously addressing design and control decisions [5].

To potentially speed up the RL-driven process synthesis, our prior work [4] presented the use of Quantum computing to leverage the theoretical advantages and continuous improvements of quantum hardware [5-6]. However, this quantum-accelerated methodology was constrained to very small flowsheet synthesis problems due to the poor scaling relationship between the required number of qubits and the problem size (e.g., the number of unit operations considered for process design). In this work, we improve this methodology by applying new

information encoding strategies to reduce the number of qubits needed in the quantum algorithm. This allows larger process design problems to be solved.

The remaining sections are organized as follows: Section 2 provides an overview of RL-driven process synthesis and covers the fundamentals of quantum RL. Section 3 introduces the re-uploading parameterized quantum circuit (PQC) approach used in our prior work and the proposed two new encoding strategies for qubit reduction. Section 4 showcases the new algorithms performance in solving three different scenarios of a process flowsheet synthesis problem. Section 5 discusses the outlook for next-step implementation on current-scale quantum hardware. Section 6 gives the concluding remarks and directions of ongoing research.

2. BACKGROUND

In this section, we review the basics of: (i) the RL-driven process synthesis approach adopted in this work, and (ii) quantum machine learning (QML).

2.1 RL-driven Process Synthesis

In this work, we adopt the matrix-based method developed by Wang et al. [3], which was also used in our prior works [4, 7].Fig. 1 provides a conceptual overview of this approach [3]. The flowsheet is represented by a stream matrix which represents the interconnectivity between potential unit operations using binary (0-1) values. It is important to note that not all available unit operations must be used. This stream matrix representation is then used as the input to an NN-based agent. The agent serves as a Q-function approximator by assigning a Q-value to each potential action. In the context of process synthesis, the action involves adding, changing, or removing a connecting stream within a process flowsheet. By taking the action that corresponds to the highest Q-value, the agent has made the theoretically optimal decision.

In this algorithm, the agent is limited to making only one action per step (i.e., only one stream connection can be changed). The RL agent can thus determine all the mixed-integer variables through iterations with the environment, which define the flowsheet structure. Afterwards, the resulting flowsheet is simulated using the Institute for the Design of Advanced Energy Systems (IDAES) platform and the continuous process variables are optimized using nonlinear optimization solvers (e.g., IPOPT) [8-9].

This nonlinear optimization is conducted by maximizing a reward or objective function defined by the user. Examples of the reward functions include maximizing economic profits, productivity, sustainability, etc. This function is also used within the RL algorithm to update the tunable parameters of the agent. At each step of the algorithm, a gradient-based constraint-free optimization technique is used to improve the agents decision-making capability. The RL method used here is the Deep Q-Network algorithm.

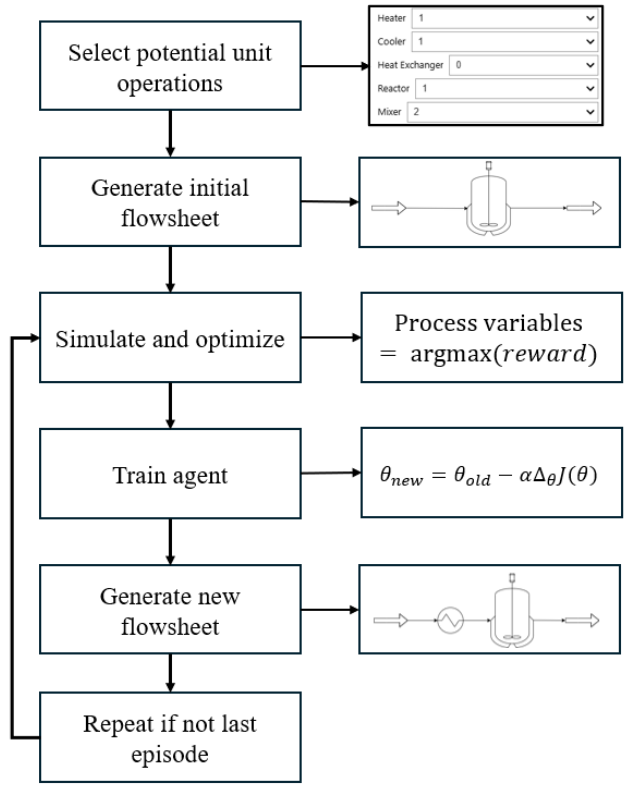

Figure 1. An overview of RL-driven process synthesis.

2.2 Quantum Reinforcement Learning

Quantum RL relies on QML which implements ML algorithms on quantum computing platforms and manipulates information encoded in qubits. The fundamental component of QML is the use of parameterized quantum circuits (PQCs), which can be described as the quantum analog of classical NNs [10].

PQCs consist of quantum circuits that take in classical or quantum data typically as quantum states and produce outputs via measurement combinations known as observables. The quantum gates have adjustable parameters that are trainable. Figure 2 illustrates an example of PQC with three qubits. Each qubit is sequentially acted upon by single-qubit rotation gates R_X, R_Y, and R_Z, followed by an entangling architecture composed of controlled-Z (CZ) gates. The rotation gates implement transformations about the x, y, and z axes of the Bloch sphere representation of the qubit's quantum state. The rotation angles are defined by the trainable parameters θ_i. The CZ gates introduce conditional phase shifts that depend on the state of another qubit, thereby generating quantum entanglement among the qubits. In this work, we will refer to a sequence of x, y, z rotation gates followed by CZ

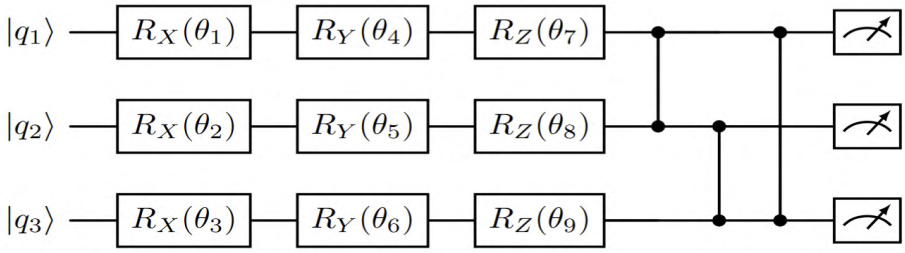

Figure 2. PQC with measurments: a 3-qubit example.

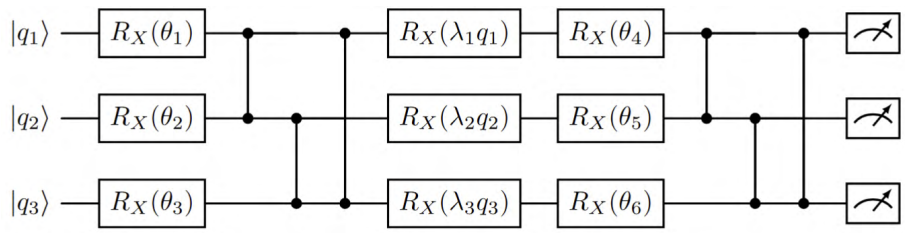

Figure 3. 3-qubit re-uploading PQC with measurements.

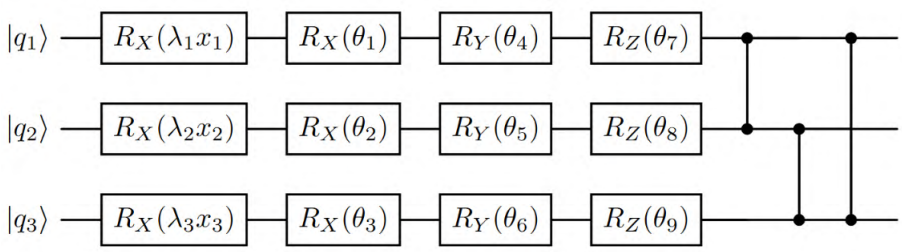

Figure 4. 3-qubit PQC using "X-only" encoding: an example of one PQC layer.

gate entanglement as a "layer" of a PQC.

For quantum RL, these PQCs can replace NNs in deep RL algorithms, as that developed in our prior work [4]. These would act as quantum agents to be trained. In addition, different with the data fed to NNs, the data fed to PQCs may need to be encoded as a quantum state. All other components of the classical RL algorithms may remain the same for quantum RL. Essentially any modern RL algorithm (e.g., DQN, DDPG) can be modified to use quantum agents by replacing NNs with PQCs [10].

3. QUANTUM RL-DRIVEN SYTHESIS

In this section, we first provide a brief review of the re-uploading PQC algorithm for the quantum RL-driven process synthesis approach presented in our prior work. We then present two new encoding approaches for qubit reduction and scalability improvement.

3.1 Quantum RL with Re-Uploading PQC

Our prior work utilized a special type of PQC, called a re-uploading PQC as shown in Fig. 3 [4]. This type of PQC increases expressibility by reinjecting the input data into the circuit after sets of CZ gates. This reinjection is accompanied with its own set of tunable parameters, which is represented by λ in Fig. 3. The data re-encoding step can be repeated many times, thus forming a quantum circuit with several layers. Despite the advantage, re-uploading PQC does require that the number of qubits in the circuit equals the size of the input data. The many layers of entanglement may also be expensive to compute.

For the RL-driven process synthesis algorithm, the input dimension is determined by the number of available unit operations. This is not an issue for classical RL, since NNs can handle extremely large input dimensions. However, for PQCs the circuit complexity increases vastly with more qubits. This necessitates the need for new quantum-based approaches which can reduce the number of qubits needed and therefore solve larger process synthesis problems.

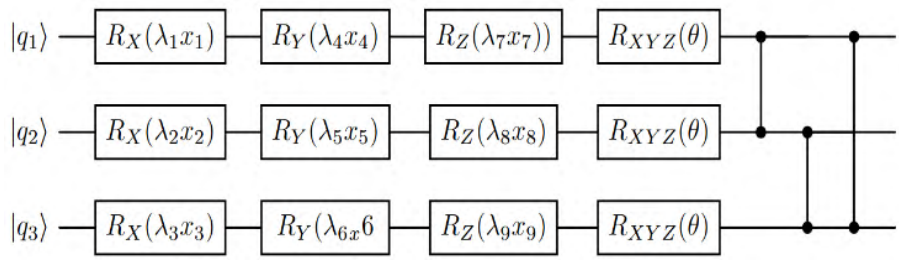

Figure 5. 3-qubit PQC using "XYZ" encoding: an example of one PQC layer.

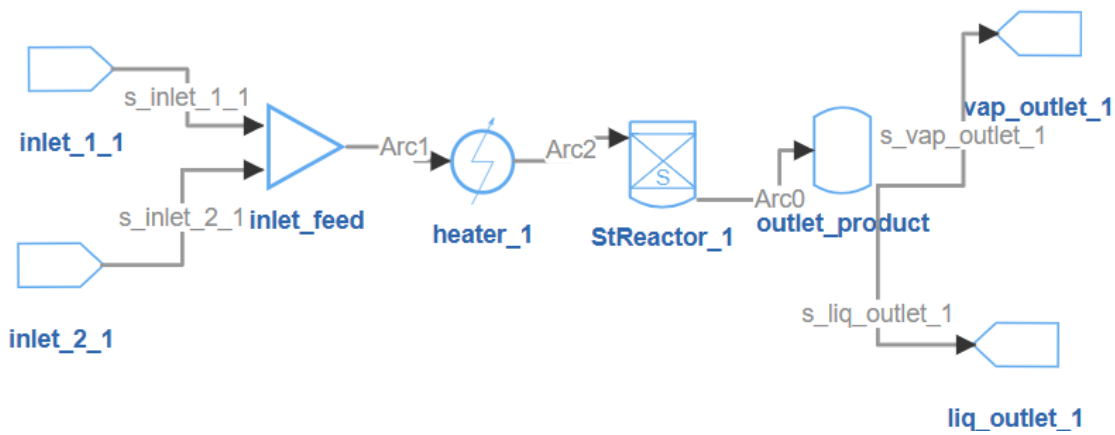

Figure 6. Optimal flowsheet corresponding to the maximum acheivable reward.

3.2 Improved Algorithm with Encoding

Motivated by the re-uploading PQC and other works in the field of QML [11-12], our proposed approaches below adopt new encoding strategies. These strategies take advantage of the fact that input data can be encoded into the angle of rotation gates. This differs with the previous algorithm which only encodes the input information to the qubits. In this way, larger problems can be solved using fewer qubits and less complex circuits.

3.2.1 X-only Encoding

In the "X-only" encoding strategy, the input features are injected directly into rotational gate angles with an additional trainable parameter λ used as a multiplier (Fig. 4). The input to each qubit of the circuit first enters an x rotation gate. It is then processed by a PQC layer, i.e. a sequence of x, y, z rotation gates followed by CZ gate entanglement. Since each qubit can only encode one extra input per layer, the number of qubits used in the circuit plus the number of layers needed should be equal to the size of the input. In other words, the number of qubits used or available determines how many layers should be added.

3.2.2 XYZ Encoding

The "XYZ" encoding strategy is similar to the "X-only" but adds additional y and z rotation gates before the qubit reaches the remaining parts of the PQC layer (Fig. 5). This allows more information to be encoded into each layer which can drastically save computational time. Again, the number of qubits used determines the number of layers needed since there is an encoding limit on each qubit per layer. This algorithm theoretically saturates the amount of information that can be embedded per layer using rotational gates, since x, y, and z are the only possible rotations.

4. CASE STUDIES

The section applied the improved encoding strategies on a case study, solving three process synthesis problem with increasing sizes to investigate algorithmic scalability. We also compare the performance of the quantum RL algorithm with classical RL.

4.1 Problem Description

Consistent with our prior works, this case study focuses on the flowsheet design of a hydrodealkylation (HDA) process which converts toluene ($C_6H_5CH_3$) into the product benzene (C_6H_6). This reaction is carried out in the presence of hydrogen (H_2) as described by Eq. 1. The reaction is solely active in the vapor phase and is

considered to occur without producing any byproducts.

$$H_2 + C_6H_5CH_3 \rightarrow H_2 + C_6H_6, \Delta H = -1.08 \times 10^5 \frac{\text{J}}{\text{mol}} \quad (1)$$

The HDA process feed conditions are summarized in Table 1 and consist of two inlet streams: (i) a liquid-phase stream with hydrogen and methane, and (ii) a vapor-phase stream composed entirely of toluene. A range of process units is available for flowsheet development, including a heater, cooler, heat exchanger, reactor, mixer, flash separation, splitter, compressor, expander, and a turbine. Each unit involves their own operational constraints, such as maintaining outlet temperatures in the range of 500-600 K. All phase separations within the process are assumed to obey ideal vapor-liquid equilibrium. The final product of the process must be a vapor-phase stream with a benzene mole fraction greater than 0.55.

In the following scenarios, the amount of available process units varies. The primary design goal is to maximize the vapor product stream flowrate and purity, forming a productivity-based reward function. The optimal flowsheet for this problem only uses a reactor and a heater with no other processing units, as shown in Fig. 6. The goal of each scenario is to build this optimal flowsheet given a set of potential process unit operations.

4.2 Scenario Analyses

4.2.1 Scenario 1

This scenario matches with the case study presented in our previous work [4], except the amount of training episodes has been extended from 1, 500 to 30, 000. The extended number of episodes is made possible by the newly proposed algorithms which only requires 4 qubits. The prior approach required 12 qubits. There are only two available unit operations: (i) reactor, and (ii) heater. This set represents the smallest possible subset of unit operations where the optimal flowsheet design can be found. The results of each algorithm for this scenario are provided in Table 2. Note that "classical" refers to the classical RL algorithm without quantum computing. "X-only" and "XYZ" refer to the quantum RL algorithms respectively with X-only encoding and XYZ encoding.

4.2.2 Scenario 2

In this scenario, there are three potential unit operations: (i) reactor, (ii) heat exchanger, and (iii) heater. Each algorithm is run for a total of 30, 000 episodes. In our previous algorithm with re-uploading PQC, this scenario would have required twenty qubits which is approaching the limit for simulation capability. The newly proposed encoding algorithms allow to solve this problem via quantum RL, only needing 5 qubits. The results of this scenario are highlighted in Table 3.

Table 1: Compositions and inlet conditions for feed streams.

	Liquid Feed	Vapor Feed
Temperature (K)	303.2	303.2
Pressure (kPa)	350	350
Flowrate (mol/s)		
Hydrogen	0.30	0.00
Toluene	0.00	0.30
Methane	0.02	0.00
Benzene	0.00	0.00

Table 2. Algorithm results for Scenario 1.

	Classical	X-only	XYZ
Max Reward Solutions	>100	>100	>100
Unique Flowsheets	3	3	3
Feasible Designs	1	1	1
Tunable Parameters	2, 374	64	40
Qubits Required	n/a	4	4
Number of Layers	2	3	1
Total Time (hr)	0.056	1.029	0.499

Table 3. Algorithm results for Scenario 2.

	Classical	X-only	XYZ
Max Reward Solutions	20	12	11
Unique Flowsheets	7	11	11
Feasible Designs	2	1	1
Tunable Parameters	2, 374	100	80
Qubits Required	n/a	5	5
Number of Layers	2	5	2
Total Time	0.0599	1.433	0.959

Table 4. Algorithm results for Scenario 3.

	Classical	X-only	XYZ
Max Reward Solutions	6	1	0
Unique Flowsheets	31	19	19
Feasible Designs	1	1	0
Tunable Parameters	2, 374	144	96
Qubits Required	n/a	6	6
Number of Layers	2	5	2
Total Time	0.086	3.040	2.315

4.2.3 Scenario 3

For this scenario, there are four potential unit operations: (i) reactor, (ii) cooler, (iii) heat exchanger, and (iv) heater. The number of episodes is extended to 60, 000, because no algorithm (including the classical RL algorithm) can find the optimal solution within 30, 000 episodes. For our previous algorithm with re-uploading PQC, this study would have required multiple simulations of a 30-qubit circuit which is not feasible. The results of this scenario are given in Table 4.

4.3 Results and Discussion

The ability of the quantum algorithms to use fewer parameters but achieve similar performance to classical algorithms is evident through the three scenarios. The quantum algorithms use an order of magnitude less tunable parameters. All three algorithms perform exceptionally well for Scenario 1 and find the optimal flowsheet over one-hundred times within the total training episodes (Table 1). In Scenario 2, an interesting result occurs in which the quantum algorithms propose more unique flowsheet designs than the classical algorithm, yet they find the maximum reward solution at only about half of the rate (Table 2). Additionally, the classical algorithm find two feasible designs while the quantum algorithms only find one each. For Scenario 3, the classical algorithm outperforms the quantum algorithms by discovering the optimal flowsheet six times compared to one and zero times for the quantum approaches, while also proposing significantly more unique flowsheets (Table 3).

Of the two newly proposed algorithms, the "X-only" encoding approach significantly outperforms the "XYZ" approach. This is particularly evident in Scenario 3 when the "X-only" method finds the optimal flowsheet while the "XYZ" method does not. This could be explained by the difference in trainable parameters. In these studies, the minimum number of layers are used. By design, the "X-only" encoding scheme must use more layers than "XYZ" encoding. It is possible that the superior performance is indicative on only the increased circuit complexity but not the different encoding schemes.

The classical RL algorithm appears to perform better than the quantum approaches especially when considering the computational time taken. However, there are two important aspects to consider: (i) the classical algorithm does use significantly more tunable parameters, and (ii) the computational time of the quantum algorithms are for simulated circuits run on a classical computer. Adding additional trainable parameters in the quantum circuits via increasing qubit or layer amounts may allow the quantum agents to become more expressive and therefore potentially as competitive as the classical agent. The computational time taken for the quantum approaches may be reduced when moving to quantum hardware.

5 CURRENT-SCALE QUANTUM HARDWARE OUTLOOK

The full end-to-end RL algorithm cannot currently run on quantum hardware due to prohibitive cost restrictions. It is technically possible, but the sampling costs of current-scale quantum circuits render this approach infeasible because of the large amount of total training episodes needed. However, it is still critical to understand whether the trained quantum agents could be realized via real-world quantum machines. To do this, we ran our trained circuits on IonQ's quantum computing simulator which mathematically represents their quantum hardware. These results are provided in Figs. 7 and 8 for the "X-only" algorithm in Scenario 1.

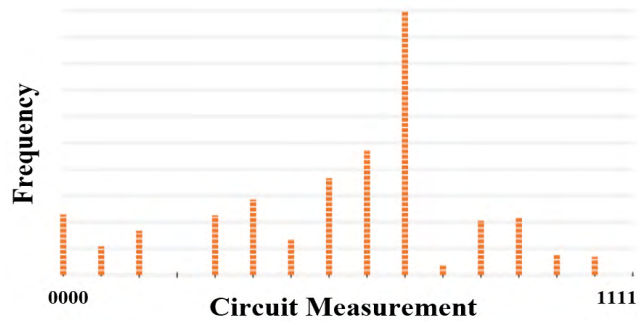

Figure 7. Scenario 1 X-only agent before training.

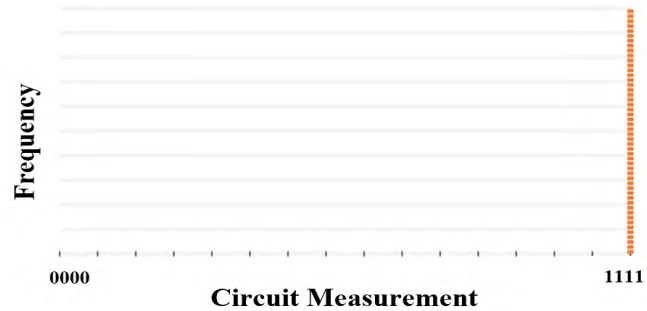

Figure 8. Scenario 1 X-only agent after training.

The results from the IonQ simulator give a probability distribution over the expected circuit measurement. These results are for the same random flowsheet before and after training. Before training the agent is random in its decision-making on which action to take while after training the probability distribution collapses to a single measurement. This shows that the agent's decision-making becomes deterministic in selecting the optimal action.

6 CONCLUDING REMARKS

In this work we presented two new quantum-accelerated RL-driven process synthesis algorithms. These algorithms were introduced to overcome a limitation of our previous approach which scaled poorly with problem size. The algorithms were demonstrated through three different scenarios of varying problem sizes. It was shown that the quantum algorithms can solve these problems while needing fewer parameters than classical RL algorithms without quantum computing. Ongoing efforts will include trying to reduce the algorithm complexity further to shorten the computational time needed on both simulated and real-world quantum hardware.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the financial

support from NSF GRFP DGE-1102689, NSF EPSCoR RII Track-4 OIA-2327303, and West Virginia University.

AUTHOR IDENTIFIERS

Author ORCIDs:
Braniff A: 0009-0006-6835-320X
You F: 0000-0001-9609-4299
Tian Y: 0000-0001-7937-5785

REFERENCES

1. Gao Q, Schweidtmann AM. Deep reinforcement learning for process design: review and perspective. Current Opinion in Chemical Engineering 44:101012 (2024). https://doi.org/10.1016/j.coche.2024.101012
2. Gao Q, Yang H, Theisen MF, Schweidtmann AM. Accelerating process synthesis with reinforcement learning: transfer learning from multi-fidelity simulations and variational autoencoders. Computers & Chemical Engineering 201:109192 (2025). https://doi.org/10.1016/j.compchemeng.2025.109192
3. Wang D, Bao J, Zamarripa-Perez MA, Paul B, Chen Y, Gao P, Ma T, Noring AA, Iyengar AKS, Schwartz DT, Eggleton EE, He Q, Liu A, Marina OA, Koeppel B, Xu Z. A coupled reinforcement learning and IDAES process modeling framework for automated conceptual design of energy and chemical systems. Energy Adv. 2:1735-1751 (2023). https://doi.org/10.1039/d3ya00310h
4. Stops L, Leenhouts R, Gao Q, Schweidtmann AM. Flowsheet generation through hierarchical reinforcement learning and graph neural networks. AIChE Journal 69: (2022). https://doi.org/10.1002/aic.17938
5. Reynoso-Donzelli S, Ricardez-Sandoval LA. An integrated reinforcement learning framework for simultaneous generation, design, and control of chemical process flowsheets. Computers & Chemical Engineering 194:108988 (2025). https://doi.org/10.1016/j.compchemeng.2024.108988
6. Braniff A, You F, Tian Y. Enhanced reinforcement learning-driven process design via quantum machine learning. Systems and Control Transactions 4:1403-1408 (2025). https://doi.org/10.69997/sct.149501
7. Ajagekar A, You F. New frontiers of quantum computing in chemical engineering. Korean J. Chem. Eng. 39:811-820 (2022). https://doi.org/10.1007/s11814-021-1027-6
8. Bernal DE, Ajagekar A, Harwood SM, Stober ST, Trenev D, You F. Perspectives of quantum computing for chemical engineering. AIChE Journal 68: (2022). https://doi.org/10.1002/aic.17651
9. Tian Y, Akintola A, Jiang Y, Wang D, Bao J, Zamarripa MA, Paul B, Chen Y, Gao P, Noring A, Iyengar A, Liu A, Marina O, Koeppel B, Xu Z. Reinforcement learning-driven process design: a hydrodealkylation example. Systems and Control Transactions 3:387-393 (2024). https://doi.org/10.69997/sct.119603
10. Lee A, Ghouse JH, Eslick JC, Laird CD, Siirola JD, Zamarripa MA, Gunter D, Shinn JH, Dowling AW, Bhattacharyya D, Biegler LT, Burgard AP, Miller DC. The idaes process modeling framework and model library—flexibility for process simulation and optimization. J Adv Manuf & Process 3: (2021). https://doi.org/10.1002/amp2.10095
11. Biegler LT, Zavala VM. Large-scale nonlinear programming using IPOPT: an integrating framework for enterprise-wide dynamic optimization. Computers & Chemical Engineering 33:575-582 (2009). https://doi.org/10.1016/j.compchemeng.2008.08.006
12. Skolik A, Jerbi S, Dunjko V. Quantum agents in the gym: a variational quantum algorithm for deep q-learning. Quantum 6:720 (2022). https://doi.org/10.22331/q-2022-05-24-720
13. Jerbi, S., Gyurik, C., Marshall, S. C., Briegel, H. J., & Dunjko, V. (2021). Parametrized quantum policies for reinforcement learning (arXiv:2103.05577).
14. Chen, S. Y.-C., Yang, C.-H. H., Qi, J., Chen, P.-Y., Ma, X., & Goan, H.-S. (2019, June 30). *Variational Quantum Circuits for Deep Reinforcement Learning*. arXiv.Org.
15. Pérez-Salinas, A., Cervera-Lierta, A., Gil-Fuster, E., & Latorre, J. I. (2019, July 3). *Data re-uploading for a universal quantum classifier*. arXiv.Org.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Unified Multi-Scale TCN Framework for Batch Manufacturing Soft Sensing and Monitoring

Yee Hung Hong[a], Zhao Jinsong*

[a] Tsinghua University, Department of Chemical engineering, Beijing, China
* Corresponding Author: jinsongzhao@mail.tsinghua.edu.cn.

ABSTRACT

Batch manufacturing is central to fine chemicals, pharmaceuticals, and bioprocessing. Its operation evolves across phases and recipes, which yields high-dimensional trajectories and strong batch-to-batch variability. Meanwhile, key quality-indicative variables are often measured offline and cannot be used as online model inputs. This work presents an integrated deep learning framework that unifies soft sensing and process monitoring in a single module using only process variables as inputs. A multi-scale Temporal Convolutional Network with multiple kernel sizes extracts complementary dynamic features from sliding windows. These features are concatenated and pooled into a compact representation that feeds two task branches. A variational autoencoder branch reconstructs the input window and provides fault monitoring signals via reconstruction deviation while regularizing the latent space through KL divergence. In parallel, a prediction branch estimates the quality-indicative variable directly from the pooled temporal features without using the variational latent sample. This separation preserves a stable quality mapping while retaining a probabilistic reconstruction model for monitoring. During inference, reconstruction error and prediction error are fused into a joint state score that more comprehensively reflects system state changes than either deviation alone. Diagnostic heatmaps are produced from residual maps and optional SHAP attributions to highlight contributing variables and time segments. The framework is validated on the Industrial Penicillin Simulation, an industrial-scale penicillin fermentation benchmark. Results show stable convergence of reconstruction, prediction, and KL terms, clear fault-set monitoring trajectories, and interpretable heatmaps that support actionable diagnosis.

Keywords: Batch Process, Fermentation, Artificial Intelligence, Process Monitoring, Machine Learning

INTRODUCTION

Modern batch manufacturing increasingly operates in an environment that is both data-rich and information-poor. Hundreds of process variables are measured online, but the variables that truly define economic value and product quality are often unavailable in real time. Many quality-indicative variables require laboratory assays, which are expensive and delayed. Operators therefore face a practical dilemma: the plant generates high-frequency trajectories, yet the most critical quality information arrives late. In this work, we use the term **quality-indicative variable (QIV)** to denote a product-quality-related variable that is economically important but typically unavailable online and only measured through delayed laboratory assays.

This reality is particularly severe in dynamic batch systems. A batch is not a stationary time series. It is a sequence of phases with changing kinetics, changing control objectives, and changing degrees of freedom. Recipes evolve. Controllers are tuned. Disturbances occur. These factors produce distribution shifts across batches and time, which is now recognized as a central scientific and engineering challenge in process monitoring [1].

Two tasks are indispensable for safe and profitable operation. The first task is soft sensing, which aims to estimate quality-indicative variables from process measurements. Deep learning has become a major driver for soft sensors because it can model nonlinear and dynamic relationships and it can exploit large-scale historical data [2]. The second task is fault monitoring, which aims to

detect abnormal states early, before they propagate into off-spec product or unsafe operation. Variational autoencoders provide a principled approach to fault monitoring because reconstruction deviation can indicate departures from a learned normal operating manifold [4].

However, in many practical deployments, soft sensing and monitoring are still implemented as separate modules. This separation is convenient but it is scientifically limiting in evolving batch environments. It forces independent feature extraction mechanisms, reduces feature efficiency, and increases maintenance cost under shift. Life-cycle studies in industrial fermentation emphasize that drift handling and model maintenance become bottlenecks once a soft sensor is deployed [3]. At the same time, recent ESCAPE-related work on IndPenSim demonstrates that even predictive monitoring frameworks must address frequent operational shifts and high retraining demand, which motivates architectures that can reuse learned temporal representations [18].

This paper addresses a core scientific question: **How can we build a unified architecture that learns a shared temporal representation for batch trajectories, performs online soft sensing and process monitoring simultaneously, does not require online quality variables as inputs, and yields interpretable diagnostic outputs that operators can trust?**

Bearing the core scientific question in mind, we propose this work which makes four contributions.

1. It formulates joint soft sensing and monitoring as a unified representation learning problem under the industrial constraint that QIV is not an online input stream.
2. It proposes a multi-scale multi-kernel TCN backbone that extracts complementary temporal patterns from batch windows and supports two monitoring tasks through a shared feature space.
3. It designs a dual-branch head in which a VAE branch provides probabilistic reconstruction learning with KL divergence for monitoring, while a separate prediction branch provides a stable quality mapping that does not depend on latent sampling.
4. It introduces a joint state score that fuses reconstruction deviation and prediction deviation to better capture system state changes, and it supports diagnosis with residual heatmaps and optional SHAP explanations [12].

RELATED WORK

Batch Process Monitoring under Non-stationary and Operational Diversity

Batch processes exhibit multi-phase dynamics and recipe dependent policies. These properties yield strong non-stationarity. They also yield batch to batch variability. In such settings, monitoring must remain reliable when the operating distribution shifts. Zhao discussed nonstationary monitoring in the era of industrial AI and summarized how distribution shift challenges classical condition monitoring assumptions [2].

Non-stationarity is amplified by practical realities. Operating strategies evolve, controllers are retuned, raw materials vary, equipment ages, and so on. These factors create gradual drift and abrupt change. Monitoring systems must therefore emphasize robustness and maintainability. A practical monitoring framework should reuse representations and reduce fragmentation. This is especially important for batch operations where each run is a finite trajectory that must be interpreted in context.

Soft Sensing with Scarce and Delayed Quality Measurements

Soft sensing is essential when quality is hard to measure online. Deep learning soft sensors can learn nonlinear mappings from multivariate trajectories to quality outcomes. **Sun et al.** surveyed deep learning for soft sensors and emphasized feature learning for complex dynamics [1].

A persistent barrier is label scarcity. Quality labels can be sparse and delayed. In many plants, QIV is not available as an online input stream. Semi supervised learning provides one route. It allows the model to learn from abundant process trajectories while using QIV only as supervision when it exists. **Lee et al.** developed a semi supervised latent dynamic variational framework to improve quality prediction under limited quality labels [7].

Beyond temporal modeling, some works model variable relations explicitly. Graph convolutional soft sensors embed interactions among variables to improve prediction [8]. Temporal graph approaches integrate temporal dependence with structured relationships and add transparency modules for interpretability [9]. Transfer and adversarial learning have also been studied to improve generalization across grades and modes when operating conditions change [10]. These directions highlight the need for representations that remain useful across shifting batch environments.

Fault Detection and Monitoring with Reconstruction based Generative Models

Reconstruction based monitoring is widely used in process monitoring. It learns normal behavior and detects deviation through residuals. Variational autoencoders provide a principled generative formulation. They pair a reconstruction term with KL divergence regularization of the latent distribution. Kingma and Welling introduced the core variational framework and the reparameterization method that enables stable training.

VAE based monitoring has been extended in

chemical processes. Interpretable variants embed attention to support fault identification and variable contribution analysis. **Bi et al.** proposed an orthogonal self-attentive VAE for interpretable fault detection and identification [4]. Distributed VAE monitoring has also been proposed for large scale systems and variable group structures [5].

A limitation of purely reconstructive monitoring is that reconstruction deviation does not always align with quality relevance. A deviation may be detectable in variables yet have limited effect on quality. A quality critical deviation may also be subtle in measured variables. This motivates quality related monitoring and combined monitoring statistics. **Tang et al.** proposed a nonlinear quality related method that combines variational information bottleneck ideas with VAE learning [6]. This line of work supports monitoring metrics that reflect both operating deviation and quality impact.

Unified Learning across Soft Sensing and Monitoring

Soft sensing and monitoring are often implemented separately. Yet both tasks aim to infer a latent process state from trajectories. Separate modeling duplicates feature extraction and increases lifecycle burden. It also creates inconsistencies across modules when the operating distribution shifts.

Lifecycle and deployment studies in industrial fermentation highlight this problem. **Metcalfe et al.** discussed MLOps oriented soft sensors and showed that drift handling and maintenance can dominate the cost of deployment [3]. This motivates unified architectures that share temporal representations and reduce retraining effort.

Recent work on IndPenSim also reflects community interest in transferable monitoring. **Hung HY et al.** reported predictive monitoring enhanced by transfer learning on IndPenSim [13]. This underscores the importance of representation reuse under operational shift. A remaining gap is a unified architecture that performs both online quality estimation and monitoring while respecting the constraint of no online quality inputs.

Explainability and Actionable Diagnosis

In industrial operation, an alarm alone is not sufficient. Operators require diagnosis that supports corrective actions. Explainability is therefore critical.

SHapley Additive exPlanations (SHAP) provides an additive attribution framework grounded in Shapley values from cooperative game theory. It offers a consistent approach for interpreting model outputs through feature contributions. Explainable diagnosis in chemical processes increasingly combines data driven learning with domain structure. **Wu et al.** proposed a topology aware diagnosis model that integrates process topology knowledge with self attention mechanisms [11]. Comprehensive fault detection and diagnosis studies continue to highlight interpretability as a key requirement for industrial adoption.

Positioning of this work: This paper proposes a unified batch monitoring module. It uses process variables only as inputs. It learns one shared temporal representation. It supports soft sensing and monitoring simultaneously. It provides interpretable diagnosis artifacts that can be used in practice.

PRELIMINARIES

Variable Screening with Lagged Mutual Information and Redundancy Control

Batch trajectories are often high dimensional. Many variables can be weakly related to quality. Many variables can also be redundant. Feeding all variables into a deep model can increase noise sensitivity and reduce interpretability. We therefore perform variable screening using nominal batches only.

We quantify nonlinear relevance to QIV using lagged mutual information. For each candidate variable x_i, we compute mutual information between $x_i(t-k)$ and $y(t)$ for lags $k = 0, \dots, K$. We take the maximum over lags as the relevance score. We also record the lag that achieves the maximum. This captures delayed effects of process variables on quality.

We control redundancy using an mRMR style greedy selection. Each candidate is penalized by its redundancy with already selected variables. Redundancy is approximated through correlation. The procedure yields a compact variable subset that preserves relevance and reduces collinearity.

Pseudocode is available in Algorithm 1.

Temporal Convolutional Networks for Causal Multi Scale Dynamics

Batch monitoring requires models that can learn long range temporal dependence. Training must remain stable. We adopt Temporal Convolutional Networks based on causal dilated convolutions and residual blocks. Causal convolution enforces that the representation at time t depends only on present and past samples within a window. Dilation expands the receptive field without recurrent computation. Residual connections stabilize optimization. We implement a multi scale TCN using parallel branches with different kernel sizes. Each branch captures temporal patterns at a different scale. Branch outputs are concatenated and pooled. This yields a compact shared temporal representation that is used by both downstream tasks.

Inputs: Nominal batches$D_0 = \{(X, y) \mid Fault = 0\}$, Candidate variables$X = \{x_1, \dots, x_d\}$, QIV y, Maximum lag K, Redundancy weight λ Number of selected variables k.

Outputs: Selected variable set S, Relevance scores $Rel(x_i)$ and optimal lags k_i^*.

Step 1: Compute lagged MI relevance

for each $x_i \in X$ do
for $k = 0$ to K do
Construct paired samples $\{(x_i(t-k), y(t))\}_t$ using nominal batches D_0
$I_{i,k} \leftarrow MI(x_i(t-k), y(t))$
end for
$Rel(x_i) \leftarrow max_{k\in[0,K]} I_{i,k}$
$k_i^* \leftarrow argmax_{k\in[0,K]} I_{i,k}$
end for

Step 2: Greedy mRMR-style selection

$S \leftarrow \varnothing$
while $|S| < k\ do$
for each $x_i \in X \backslash S\ do$
if $S = \varnothing\ then\ R_i \leftarrow 0$
else
$R_i \leftarrow \frac{1}{|S|} \sum_{x_j \in S} |corr(x_i, x_j)|$
end if
$Score(x_i) \leftarrow Rel(x_i) - \lambda R_i$
end for
$x^* \leftarrow argmax_{x_i \in X\backslash S} Score(x_i)$
$S \leftarrow S \cup \{x^*\}$
end while
$return\ S, \{Rel(x_i), k_i^*\}$

Algorithm 1: Variable Screening with Lagged Mutual Information and mRMR-style selection

Semi Supervised Training with Sparse Quality Labels

In realistic plants, QIV labels are sparse and delayed. QIV is not used as an online input channel. It is used only as a training label when it exists. We adopt a semi supervised training protocol.

A reconstruction objective is trained on nominal windows. It uses only process variables. This objective is always available. A prediction objective is trained only when a QIV label exists. We use a label mask $m_{b,t} \in \{0, 1\}$ so that unlabeled windows do not contribute to the prediction loss. This enables joint training on mixed labeled and unlabeled data.

PROBLEM FORMULATION

Let $x_{b,t} \in \mathbb{R}^D$denote process variables at time t for batch b. We use a sliding window

$$X_{b,t} = [x_{b,t-W+1}, \dots, x_{b,t}] \in \mathbb{R}^{W\times D}$$

Let $y_{b,t}$ denote a quality indicative variable aligned with the window end. In the target setting, $y_{b,t}$ is sparse and delayed. It is not an online input channel.

We aim to build one module that:

1 takes $X_{b,t}$ as input.
2 outputs a reconstructed window $\hat{X}_{b,t}$ for monitoring.
3 outputs a quality estimate $\hat{y}_{b,t}$ for soft sensing.
4 computes a joint state score by fusing monitoring deviation and quality deviation.
5 supports interpretable diagnosis through heatmaps and optional SHAP attributions.

METHODOLOGY

Architecture Overview

The framework consists of a shared temporal backbone and two task branches.

1. Shared backbone. A multi scale TCN uses multiple kernel sizes in parallel.

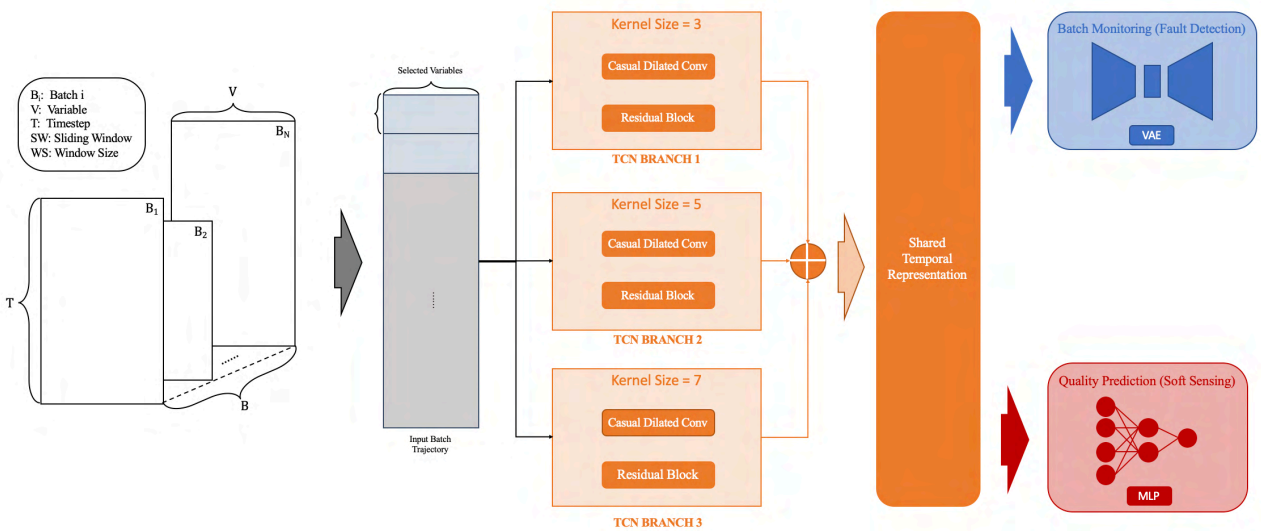

Figure 1. Multi scale multi kernel TCN backbone with two branches

2. Monitoring branch. A VAE reconstructs the input window. Monitoring is based on reconstruction deviation.
3. Soft sensing branch. A predictor estimates QIV from pooled temporal features. It does not use the latent sample.

This design enforces controlled sharing. Both tasks reuse temporal features. The predictor avoids stochastic sampling to preserve stable quality mapping.

Multi Scale TCN Backbone with Multiple Kernel Sizes

Batch trajectories contain patterns at different time scales. Early phases can exhibit fast transients. Later phases can exhibit slow drifts. A single kernel size can under represent this diversity. We therefore build parallel TCN branches with kernel sizes $k_1, k_2, \dots \dots$.

Each branch uses causal dilated convolutions. Dilation increases receptive field efficiently. Residual connections improve stability for deep temporal stacks. Let $f_{b,t}^{(k)}$ denote features from kernel branch k. We concatenate:

$$F_{b,t} = Concat(f_{b,t}^{(k_1)}, f_{b,t}^{(k_2)}, \dots \dots)$$

We pool across time to obtain compact shared features:

$$g_{b,t} = Pool(F_{b,t})$$

VAE Monitoring Branch with KL Divergence Regularization

The VAE branch maps $g_{b,t}$ into a latent distribution $q(z_{b,t} \mid g_{b,t})$ with parameters $\mu_{b,t}$ and $\sigma_{b,t}$. We sample using reparameterization:

$$z_{b,t} = \mu_{b,t} + \sigma_{b,t} \odot \epsilon, \qquad \epsilon \sim \mathcal{N}(0, I)$$

A decoder reconstructs the window and outputs $\hat{X}_{b,t}$ Reconstruction loss:

$$\mathcal{L}_{rec} = MSE(X_{b,t}, \hat{X}_{b,t})$$

KL divergence:

$$\mathcal{L}_{KL} = D_{KL}(q(z_{b,t} \mid g_{b,t}) \,||\, \mathcal{N}(0, I))$$

Monitoring deviation:

$$e_{rec} = MSE(X_{b,t}, \hat{X}_{b,t})$$

Prediction Branch for Soft Sensing without Latent Sampling

The prediction branch estimates QIV from pooled temporal features:

$$\hat{y}_{b,t} = f_{pred}(g_{b,t})$$

The prediction branch does not take $z_{b,t}$ as input. This avoids sampling noise in the quality mapping. QIV is not an input stream. It is only used as a label when available.

We use a robust prediction loss:

$$\mathcal{L}_{pred} = m_{b,t} \bullet Huber(y_{b,t}, \hat{y}_{b,t})$$

where $m_{b,t}$ is the label mask defined in the Semi-Supervised Training section.

Joint Multi Task Objective

The overall objective is:

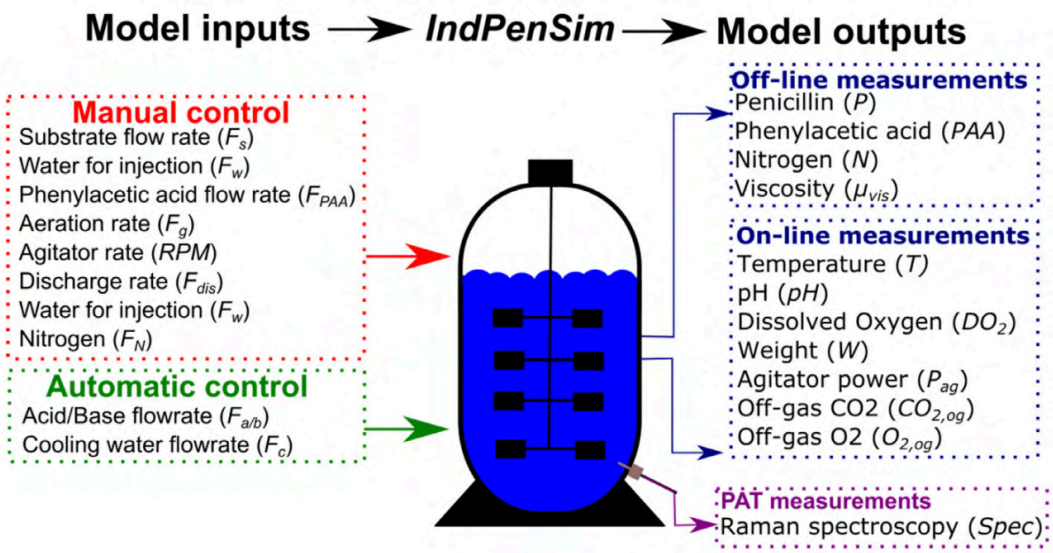

Figure 2. Summary of Inputs and Outputs of IndPenSim

$$\mathcal{L} = \lambda_{rec}\mathcal{L}_{rec} + \lambda_{pred}\mathcal{L}_{pred} + \beta\mathcal{L}_{KL}$$

This objective forces the backbone to learn shared temporal features. It also maintains a regularized latent space for stable reconstruction learning.

Joint State Score by Fusing Reconstruction and Prediction Deviation

At inference, the model provides reconstruction deviation:

$$e_{rec}(t) = MSE(X_{b,t}, \hat{X}_{b,t})$$

The model also provides a quality estimate $\hat{y}_{b,t}$. When a lab assay arrives, prediction deviation can be computed:

$$e_{pred}(t) = |y_{b,t} - \hat{y}_{b,t}|$$

We define a joint state score:

$$S(t) = \alpha \bullet norm\big(e_{rec}(t)\big) + (1-\alpha) \bullet norm\big(e_{pred}(t)\big)$$

When $y_{b,t}$ is not yet available, the system runs in early warning mode with $S(t) = norm\big(e_{rec}(t)\big)$. This supports online monitoring under delayed quality measurement.

Interpretable Inference with SHAP Heatmaps

For diagnosis, we compute the residual map:

$$R_{b,t} = (X_{b,t} - \hat{X}_{b,t})^2$$

We visualize $R_{b,t}$ as a variable time heatmap for alarm windows. This highlights which variables and which time segments contribute most to reconstruction deviation.

If attribution based interpretation is required, we apply SHAP to the prediction output or to the joint score. This yields variable contribution values that support human interpretation.

CASE STUDY: INDPENSIM BATCH PENICILLIN FERMENTATION

We validate the framework on IndPenSim. IndPenSim simulates an industrial scale fed batch penicillin fermentation process and includes realistic variability, disturbances, and sensor noise.[16] It is widely used as a benchmark for monitoring and soft sensing studies. Unlike the traditional PenSim benchmark [15], IndPenSim models a 100, 000 L bioreactor, bringing it closer to real-world industrial conditions and validated by historical industrial data. Detailed process descriptions and operational specifics can be found at the IndPenSim website and in the original publications [14]. **Figure 2** illustrates the summary of IndPenSim.

Recent ESCAPE community work also used IndPenSim to study predictive monitoring with transfer learning under changing control strategies [13]. This supports the view that transferable temporal representations are important. Our focus differs in one key constraint. We enforce that QIV is not an online input. We unify soft sensing and monitoring within one module and fuse reconstruction and prediction deviations for system state tracking.

RESULTS AND DISCUSSION

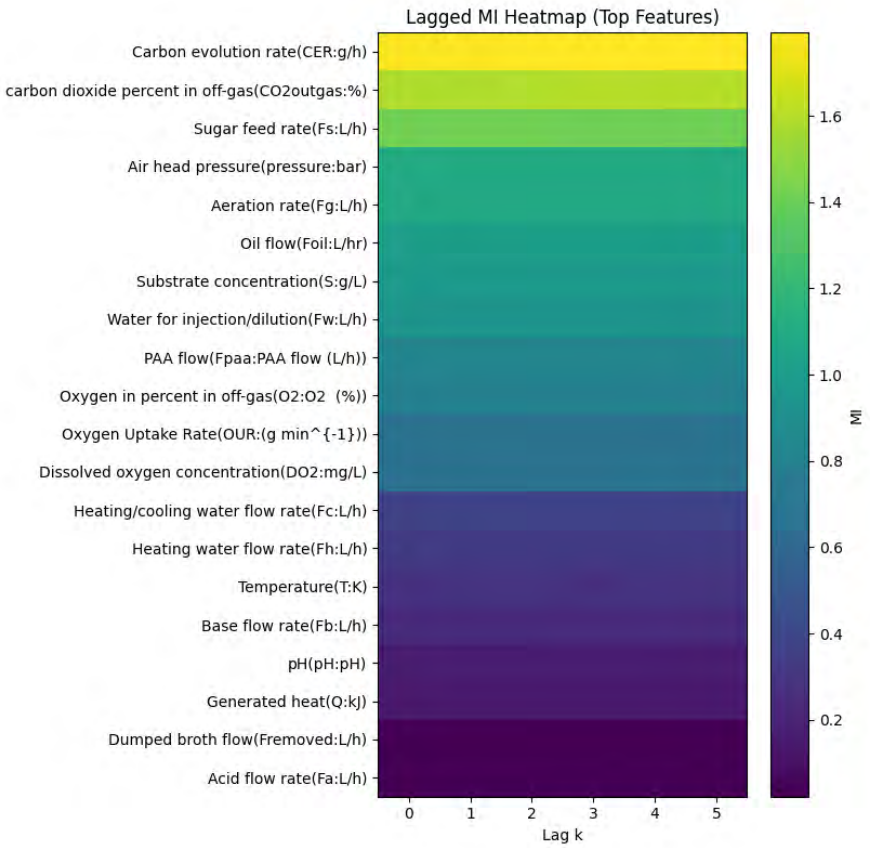

Figure 3. Lagged MI heatmap of top variables

Variable Screening Results

To reduce input dimensionality while retaining quality relevance, we performed variable screening using **nominal batches only**. For each candidate variable x_i, we computed **lagged mutual information (MI)** between $x_i(t-k)$ and the quality-indicative variable $y(t)$ over lags $k = 0, \dots, 5$. The **maximum** MI across lags was used as the relevance score $Rel(x_i) = max_k I(x_i(t-k), y(t))$, and the corresponding **optimal lag** k_i^* was recorded.

Figure 3 above visualizes the lag-dependent MI patterns for the most informative variables. Clear delayed effects can be observed: several key variables achieve their peak relevance at $k^* = 4{\sim}5$, including **CER** ($k^* = 4$), **CO**_2 in off-gas ($k^* = 4$), **substrate concentration** ($k^* = 5$), **OUR** ($k^* = 5$), and **temperature** ($k^* = 5$). This is consistent with fed-batch fermentation where operational actions (e.g., feeding/aeration) and physiological responses propagate to quality with a time delay.

Figure 4 summarizes the **max-MI ranking**, where variables related to **metabolic activity and gas exchange** (CER and off-gas CO_2) dominate the relevance ranking, followed by **feeding** and **aeration/pressure** variables. These variables are physically meaningful: gas-out measurements and CER reflect respiratory activity and carbon balance, which are closely tied to the process state and product formation, while feed and aeration govern substrate availability and oxygen transfer.

Because high-dimensional trajectories often contain strong collinearity, we further applied an **mRMR-style greedy selection** to control redundancy. Each candidate variable was penalized by its average absolute correlation with already selected variables, and the selection score was computed as

$$Score(x_i) = Rel(x_i) - \lambda Red(x_i)$$

This step prevents the final input set from being dominated by highly correlated variables. For example, although **CO**_2 in off-gas has the second highest max-MI, it is selected after **CER** and **sugar feed rate**, reflecting redundancy suppression against already selected strong predictors.

The final selected subset contains **18 variables**, reported in **Table 1**, which are used as the input channels for the subsequent multi-kernel TCN backbone. Overall, the screening results indicate that the quality-relevant information is concentrated in variables representing **respiration/gas exchange**, **feeding**, and **oxygen-transfer-related operations**, while low-MI variables (e.g., control

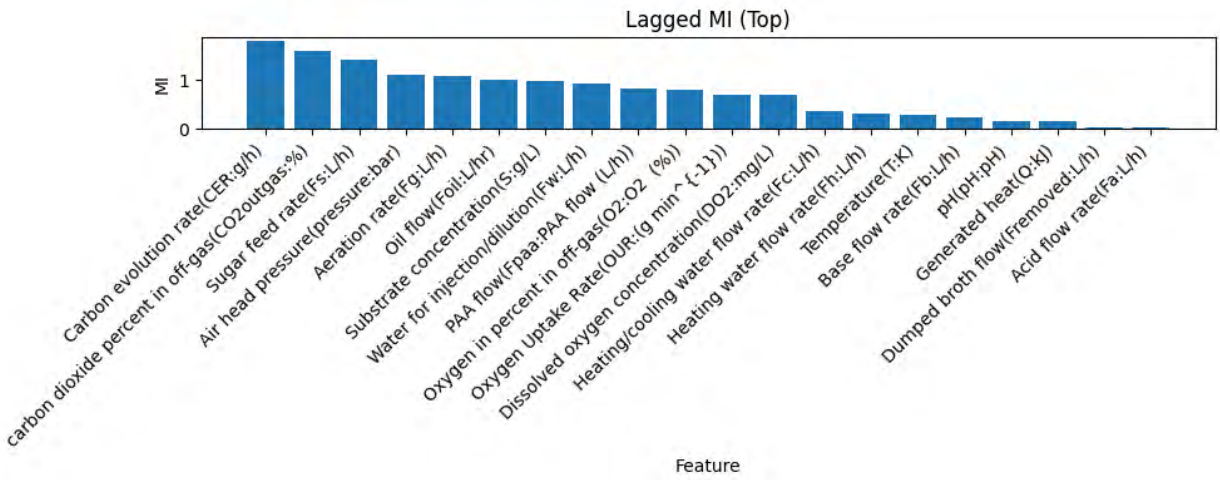

Figure 4. Max lagged MI ranking of candidate variables

Table 1. Selected variables from lagged MI screening and redundancy-controlled selection (k=18)

Rank	Variable	Max lagged MI	Best lag k^*	mRMR score
1	Carbon evolution rate(CER:g/h)	1.793	4	1.793
2	Sugar feed rate(Fs:L/h)	1.417	0	1.343
3	carbon dioxide percent in off-gas(CO2outgas:%)	1.597	4	1.312
4	Air head pressure(pressure:bar)	1.107	0	0.962
5	Aeration rate(Fg:L/h)	1.091	1	0.889
6	Substrate concentration(S:g/L)	0.982	5	0.875
7	Water for injection/dilution(Fw:L/h)	0.940	0	0.854
8	Oil flow(Foil:L/hr)	1.017	1	0.804
9	Oxygen in percent in off-gas(O2:O2 (%))	0.797	1	0.701
10	PAA flow(Fpaa:PAA flow (L/h))	0.829	3	0.647
11	Dissolved oxygen concentration(DO2:mg/L)	0.694	4	0.648
12	Oxygen Uptake Rate(OUR:(g min^{-1}))	0.695	5	0.520
13	Heating water flow rate(Fh:L/h)	0.326	0	0.274
14	Heating/cooling water flow rate(Fc:L/h)	0.376	4	0.263
15	Temperature(T:K)	0.288	5	0.231
16	Base flow rate(Fb:L/h)	0.244	4	0.181
17	pH(pH:pH)	0.174	0	0.153
18	Generated heat(Q:kJ)	0.150	1	0.128

reference flags) are naturally deprioritized.

Baseline Comparison with the Proposed Framework

This section compares the proposed unified predictive monitoring framework with two recurrent baselines. The baselines are **Baseline-1 (LSTM)** and **Baseline-2 (GRU)**, both following the same predictive monitoring pipeline with a fixed alarm threshold. Here, **LSTM** denotes **Long Short-Term Memory**, and **GRU** denotes **Gated Recurrent Unit**. Both are recurrent neural network architectures for sequential modeling, while GRU uses a simpler gating mechanism with fewer parameters than LSTM.

The proposed method combines a shared temporal representation with a reconstruction-based monitoring

branch and a deterministic quality prediction branch, and triggers alarms using a unified monitoring score.

To ensure a fair comparison, all methods were trained and evaluated under an identical experimental specification. The dataset was split into 80%/10%/10% train/validation/test batches. All models used sliding windows of length $W = 10$ with stride 1. The same feature screening procedure was applied: lagged mutual information was computed up to max lag $K = 5$, followed by redundancy control with weight $\lambda = 0.5$, and the **top** $k = 18$ variables were retained as input channels for every model.

All networks were optimized with **Adam** using a learning rate of 1×10^{-3} and weight decay 1×10^{-4}, with batch size 64 and up to 50 epochs. Early stopping was applied with patience = 6 on the validation set. Input variables were scaled using robust scaling (clipping at 5.0), and the target was scaled during training (scale_y = True). For the joint learning objective, we used the same loss weights across runs:$\lambda_{rec} = 1.0, \lambda_{pred} = 1.0$, and $\lambda_{pred} = 0.5$. During inference, the alarm threshold was computed consistently as the 0.99 quantile of nominal scores for all methods. We note that this choice is not claimed to be a universal or golden standard. Rather, it is a practical empirical calibration rule that uses the upper tail of the nominal score distribution to control false alarms. In this study, we selected a high nominal quantile because batch trajectories contain strong phase evolution and transient variability, and a conservative threshold helps suppress nuisance alarms during nominal operation. More importantly, using the same thresholding principle for all methods ensures a fair comparison, so that the observed differences reflect the quality of the learned monitoring statistic rather than method-specific threshold tuning.

We evaluated all methods on eight fault batches generated by simulator. For each batch, the result sheet reports the fault interval $[t_{fault_start}, t_{fault_end}]$, the first alarm time $t_{alarm,}$ (FAT, the first time index where the score exceeds the threshold), and the detection delay $\Delta t = t_{alarm,} - , t_{fault_start}$. If a method never crosses the threshold during the entire batch, then $t_{alarm,}$ is undefined and the table records "-", which also makes Δt undefined.

According to **Table 2**, across the eight fault batches, the LSTM baseline does not produce valid alarms and therefore has 0/8 detection. At the same time, it generates frequent nuisance alarms before the fault (average 19.75 pre-fault false alarms), which indicates that its monitoring statistic is strongly affected by early-batch transients or phase evolution. The GRU baseline achieves 8/8 detection, but the detection delay is highly variable, with a large worst case delay (maximum 631). In contrast, the proposed framework achieves 8/8 detection with zero pre-fault false alarms, and yields substantially smaller delays (mean 17.5, median 6.5, maximum 54). Overall, the proposed method improves both reliability (consistent detection) and timeliness (lower delay) while maintaining clean monitoring behavior.

Table 2. Overall comparison

Method	Detection rate	Mean delay	Median delay	Max delay	Avg. false alarms (pre-fault)
Baseline-1 (LSTM)	0/8 (0%)	–	–	–	19.75
Baseline-2 (GRU)	8/8 (100%)	99.75	29.0	631	0.125
Proposed	8/8 (100%)	17.5	6.5	54	0.0

The stark difference between GRU and LSTM suggests that, under the present data scale and training budget, the simpler recurrent parameterization of GRU is easier to optimize than LSTM. GRU uses a lighter gating structure and may therefore adapt more effectively to the nonstationary batch trajectories and limited fault examples considered here. By contrast, the LSTM monitoring score appears to be more strongly affected by early-batch transients and phase evolution, which leads to poor separation between nominal and faulty score distributions. This interpretation is consistent with the high pre-fault false alarm count of LSTM (19.75 on average) and its failure to produce any valid alarm in the actual fault intervals.

Table 3. Batch Wise Result

batch_id	fault_start_t	fault_end_t	LSTM FAT	LSTM delay	GRU FAT	GRU delay	Proposed FAT	Proposed delay
1	379	459	–	–	407	28	386	7
2	349	449	–	–	417	68	403	54
3	499	549	–	–	503	4	503	4
4	379	459	–	–	409	30	387	8
5	379	459	–	–	409	30	385	6
6	499	549	–	–	503	4	504	5
7	499	549	–	–	502	3	501	2
8	349	449	–	–	980	631	403	54

The batch-wise results clarify where the baselines fail. For the LSTM baseline, the first alarm time and delay are missing ("-") because the score never exceeds the threshold; hence, detection cannot be declared. This is not a reporting issue but a direct consequence of the

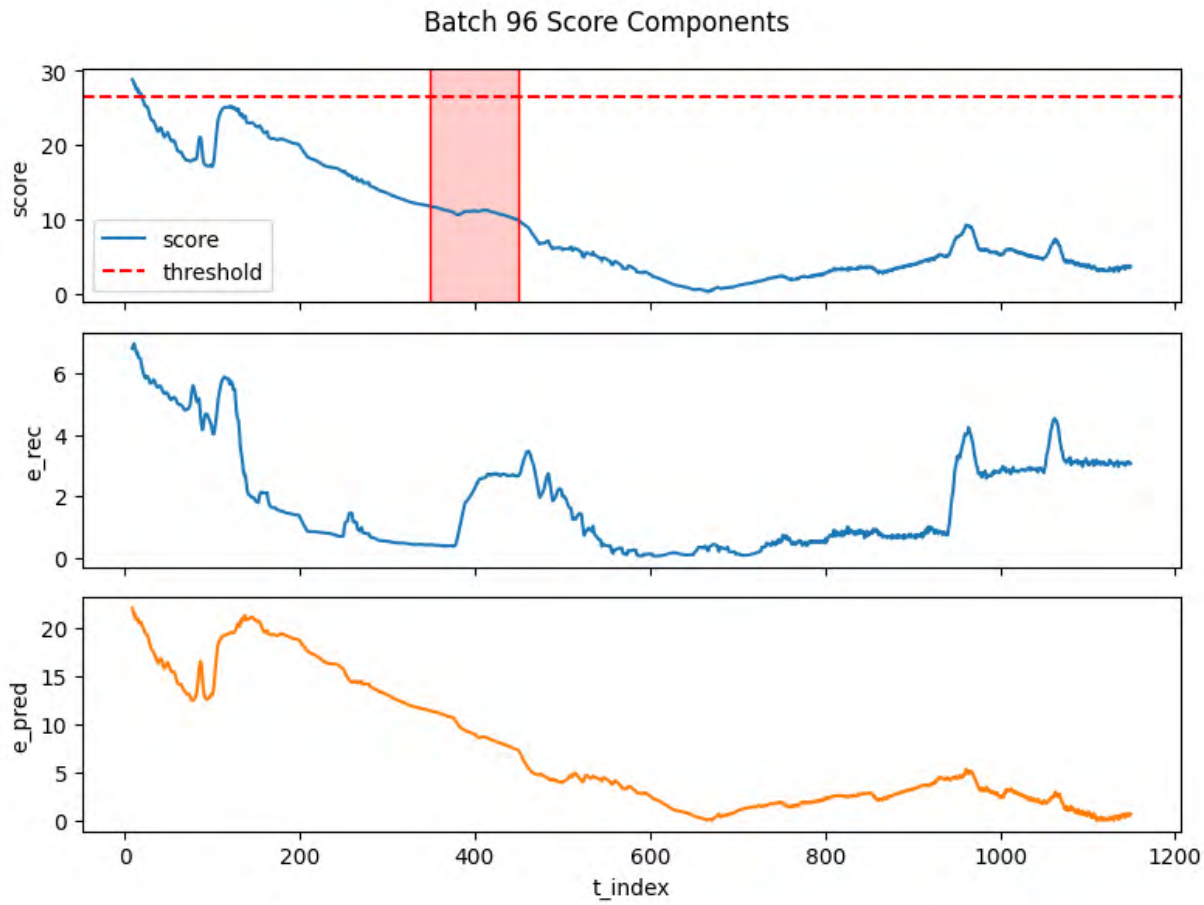

Figure 5. Score components for Batch 8 under the LSTM baseline. The shaded region denotes the fault interval and the dashed line denotes the fixed alarm threshold.

alarm rule. Batch 8 provides a representative example and is used to illustrate the mechanism behind these missing entries.

Table 4. Fault Batch 8 case study result

Batch ID	Fault Type	FDS	FAR	EDT
8	Coolant-Flowrate Fault	1	0	54

Figure 5 shows the monitoring score and its components for Batch 8. The score (top) stays below the fixed threshold throughout the highlighted fault interval, so the LSTM baseline never triggers an alarm; therefore t_{alarm} and the detection delay are recorded as "-". This indicates that the LSTM score is strongly influenced by early-stage variability and phase transitions, which raises the calibrated threshold and makes mid-batch faults difficult to detect. Therefore, the underperformance of LSTM in this study is not merely a consequence of one difficult batch, but reflects a systematic calibration issue: nominal early-phase variability inflates the threshold-relevant score distribution and reduces sensitivity to mid-batch faults.

The reduced delay of the proposed framework can be attributed to two design properties. First, the multi-scale TCN backbone extracts temporal patterns at different receptive fields, which is beneficial for batch processes where short transients and slow drifts coexist. Second, the proposed monitoring statistic fuses reconstruction deviation with prediction deviation, so the model can respond not only to trajectory abnormality but also to quality-relevant inconsistency. This dual evidence appears particularly useful in difficult cases such as Batch 8, where a single recurrent monitoring pathway may detect the fault only after a substantial lag.

Overall, under identical training and inference settings, the proposed framework delivers more actionable

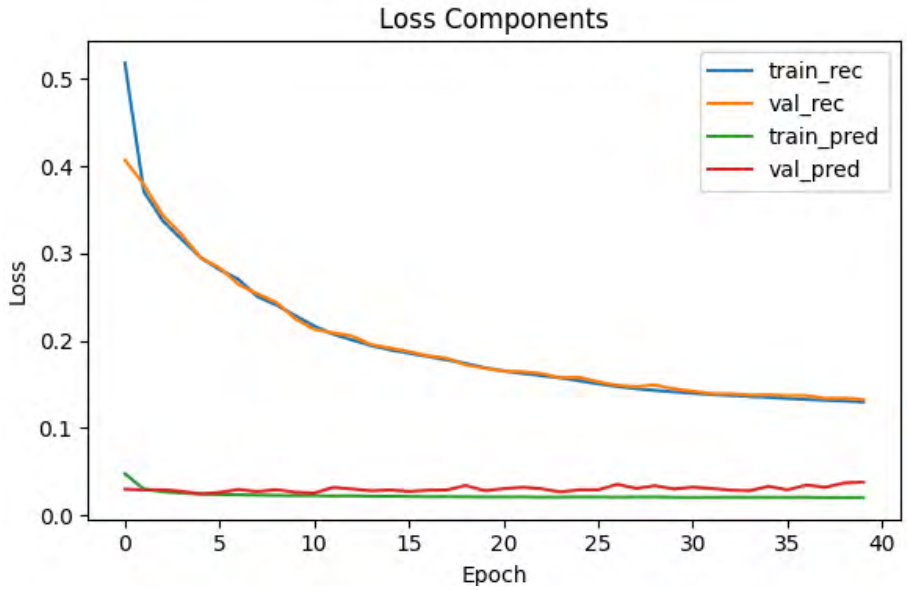

Figure 6. Training Loss Component

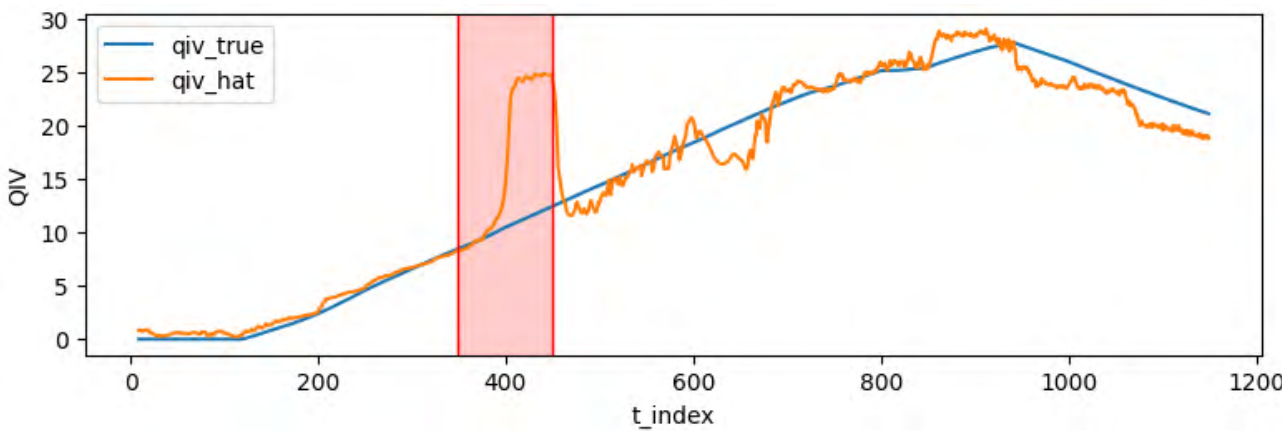

Figure 7. QIV prediction vs. true QIV for Batch 8

monitoring than the recurrent baselines: LSTM suffers from frequent pre-fault false alarms and missed detections, while GRU can detect faults but may alarm very late on hard cases (e.g., Batch 8). The proposed method achieves reliable detection with lower delays and cleaner alarm trajectories.

Case Study of Batch 8: Monitoring, Soft-Sensing and SHAP Diagnosis

This subsection presents a representative fault-batch case study to illustrate how the proposed framework behaves during inference and how the diagnostic module supports interpretation. In our experimental outputs, the plots run corresponds to fault batch_id = 8 in the evaluation list. For this batch, the fault-metric summary reports Fault Detect Success (FDS) = 1, False Alarm Rate (FAR) = 0, and Event/early detection delay (EDT) = 54 time steps, indicating that the fault is detected successfully with no pre-fault alarms and a moderate detection delay.

We first verify that the joint learning objective is optimized stably. The overall loss curve (**Figure 6**) decreases smoothly for both training and validation, and the gap remains small, suggesting stable convergence without severe overfitting. The loss component plot further shows that the reconstruction term dominates the objective, while the prediction term remains comparatively small yet consistent across epochs. This behavior is expected under semi-supervised operation: reconstruction learning is always available from process trajectories, while prediction learning is constrained by sparse quality labels. Together, these curves support that the shared temporal representation can be trained reliably and that both task branches are optimized simultaneously.

We then examine the online soft-sensing behavior for this batch. **Figure 7** plots the ground-truth QIV (y) and

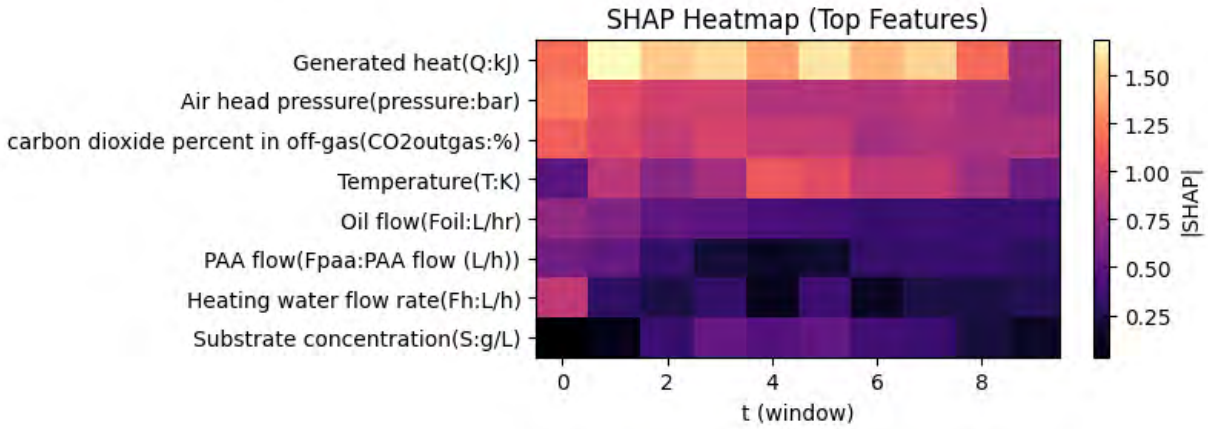

Figure 8. SHAP heatmap for the alarm window of Batch 8. Warmer colors indicate larger positive contributions to the score.

the model estimate ($\hat{y}$) over time, with the fault interval highlighted. Outside the fault window, $\hat{y}$ generally tracks the smooth trend of y, indicating that the deterministic prediction head provides a stable quality mapping from the pooled temporal features. In the highlighted fault interval, the predicted QIV shows a pronounced deviation from the true trajectory, which is consistent with the idea that faults can manifest as quality-relevant inconsistencies even when process trajectories are still within a plausible operating range. This mismatch later becomes a key contributor to the monitoring statistic through the prediction deviation component.

Finally, we demonstrate interpretability using SHAP attributions on an alarm window. **Figure 8** visualizes a SHAP heatmap for the top contributing variables within the local inference window (the horizontal axis corresponds to time positions inside the sliding window). The heatmap indicates that the score at this alarm moment is dominated by Generated heat (Q), Heating/cooling-related signals, Temperature, and closely coupled variables such as Air head pressure and CO_2 in off-gas, with additional contributions from Oil flow, PAA flow, and Heating water flow rate. Importantly, this attribution pattern is not only interpretable but also diagnostically specific: the concentration of contributions on thermal-balance and heat-removal pathways is consistent with the known fault label for Batch 8, i.e., a Coolant-Flowrate Fault. A reduction or perturbation in coolant flow directly weakens heat removal, which manifests as abnormal thermal behavior (Q and temperature), and subsequently propagates to gas-exchange and pressure dynamics through shifts in metabolic state and mass-transfer conditions. Therefore, the SHAP heatmap provides an actionable diagnosis that can localize the event to a cooling/coolant subsystem disturbance, rather than suggesting a feed-driven or aeration-driven root cause, effectively aligning the model's explanation with the ground-truth fault type (Coolant-Flowrate Fault) for Batch 8.

In summary, Batch 8 illustrates the full workflow of the proposed framework: (i) stable joint training with simultaneous reconstruction and prediction learning; (ii) online quality tracking with fault-induced prediction deviation during the fault window; and (iii) SHAP-based explanation that localizes the alarm to a small set of physically interpretable variables, supporting actionable diagnosis.

CONCLUSIONS

This work advances batch analytics beyond isolated "soft sensor + monitor" toolchains by introducing a unified, explainable framework that learns a reusable process-state representation for nonstationary operations while respecting the industrial reality that quality measurements are delayed and unavailable online. Looking forward, this architecture provides a strong foundation for the next step in industrial AI: integrating topology or physics constraints for sharper root-cause localization, adopting uncertainty-aware and adaptive fusion to handle evolving label availability, and enabling transfer/continual learning across recipes and plants to minimize retraining. Together, these directions point toward monitoring systems that are not just accurate, but trustworthy, maintainable, and decision-shaping in order to close the loop from data to diagnosis to intervention in modern batch manufacturing.

ACKNOWLEDGEMENTS

This work was supported by the National Key R&D

Program of China under Grant 2024YFC3013601.

REFERENCES

1. Sun Q, Ge Z. A survey on deep learning for data-driven soft sensors. IEEE Trans. Ind. Inf. 17:5853-5866 (2021). https://doi.org/10.1109/tii.2021.3053128
2. Zhao Z, Lee JM. Soft sensing and fault detection of batch processes using a multi-task deep learning framework. Comput Chem Eng 163:107827 (2022) https://doi.org/10.1016/j.compchemeng.2022.107827
3. Metcalfe B, Acosta-Pavas JC, Robles-Rodriguez CE, Georgakilas GK, Dalamagas T, Aceves-Lara CA, Daboussi F, Koehorst JJ, Corrales DC. Towards a machine learning operations (mlops) soft sensor for real-time predictions in industrial-scale fed-batch fermentation. Computers & Chemical Engineering 194:108991 (2025). https://doi.org/10.1016/j.compchemeng.2024.108991
4. Zhao C. Perspectives on nonstationary process monitoring in the era of industrial artificial intelligence. Journal of Process Control 116:255-272 (2022). https://doi.org/10.1016/j.jprocont.2022.06.011
5. Bi X, Zhao J. A novel orthogonal self-attentive variational autoencoder method for interpretable chemical process fault detection and identification. Process Safety and Environmental Protection 156:581-597 (2021). https://doi.org/10.1016/j.psep.2021.10.036
6. Tang P, Peng K, Dong J. Nonlinear quality-related fault detection using combined deep variational information bottleneck and variational autoencoder. ISA Transactions 114:444-454 (2021). https://doi.org/10.1016/j.isatra.2021.01.002
7. Lee YS, Chen J. Developing semi-supervised latent dynamic variational autoencoders to enhance prediction performance of product quality. Chemical Engineering Science 265:118192 (2023). https://doi.org/10.1016/j.ces.2022.118192
8. Jia M, Xu D, Yang T, Liu Y, Yao Y. Graph convolutional network soft sensor for process quality prediction. Journal of Process Control 123:12-25 (2023). https://doi.org/10.1016/j.jprocont.2023.01.010
9. Guo W, Zhu J, Yu X, Jia M, Liu Y. Temporal graph convolutional network soft sensor for molecular weight distribution prediction. Chemometrics and Intelligent Laboratory Systems 252:105196 (2024). https://doi.org/10.1016/j.chemolab.2024.105196
10. Dai Y, Yang C, Zhu J, Liu Y. Adversarial transferred data-assisted soft sensor for enhanced multigrade quality prediction. ACS Omega 8:19900-19911 (2023). https://doi.org/10.1021/acsomega.3c01832
11. Wu D, Bi X, Zhao J. Protopormer: toward understandable fault diagnosis combining process topology for chemical processes. Ind. Eng. Chem. Res. 62:8350-8361 (2023). https://doi.org/10.1021/acs.iecr.3c00206
12. Rao S, Wang J. A comprehensive fault detection and diagnosis method for chemical processes. Chemical Engineering Science 300:120565 (2024). https://doi.org/10.1016/j.ces.2024.120565
13. Hung HY, Jinsong Z. Enhancing batch chemical manufacturing via development of deep learning based predictive monitoring with transfer learning. Systems and Control Transactions 4:1548-1554 (2025). https://doi.org/10.69997/sct.188760
14. Goldrick S, Ştefan A, Lovett D, Montague G, Lennox B. The development of an industrial-scale fed-batch fermentation simulation. Journal of Biotechnology 193:70-82 (2015). https://doi.org/10.1016/j.jbiotec.2014.10.029
15. Birol G, Ündey C, Çinar A. A modular simulation package for fed-batch fermentation: penicillin production. Computers & Chemical Engineering 26:1553-1565 (2002). https://doi.org/10.1016/s0098-1354(02)00127-8

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimizing the Solubility of Organic Molecules in Mixed Solvents Using Bayesian Optimization and Multicomponent Directed-Message Passing Neural Networks

Simona Buzzi[a], Ulderico Di Caprio[a], Dominik Bongartz[a], and Florence Vermeire[a]*

[a] KU Leuven, Department of Chemical Engineering, Celestijnenlaan 200F, Leuven 3001, Belgium.
* Corresponding Author: florence.vermeire@kuleuven.be

ABSTRACT

Accurate prediction of solubility limits of organic compounds in mixed solvents is critical for the design and optimization of chemical and pharmaceutical processes. Recent advances in machine learning have enabled fast and reliable prediction of physicochemical properties of molecules, including solubility. In this work, we present a Bayesian Optimization framework to identify optimal solvent combinations, compositions, and temperatures that maximize the solubility of active pharmaceutical ingredients. The optimization strategy leverages a multicomponent directed-edge message passing neural network trained on solvent mixtures to predict solubility in ternary systems consisting of a solute and two solvents. To enable efficient Bayesian optimization, we represented the solvents in a continuous space and compare three different strategies: integer enumeration, numerical descriptors, and deep embeddings. The proposed approach was tested on a dataset comprising 14 299 points in solvent mixtures. The results indicate that integer enumeration and embedding-based representations achieve the best performance in identifying solvent mixtures that maximize solubility for each compound at a fixed number of model evaluations. While demonstrated on ternary systems, the framework is readily extensible to mixtures with a larger number of components. With this work we aim to identify which chemical representation is most suitable for similar optimization problems, providing a solid foundation for further and broader exploration, for example in the optimization of crystallization processes.

Keywords: Bayesian Optimization, Deep Learning, Mixed Solvents, Solubility

INTRODUCTION

Predicting solubility limits of organic compounds in mixed solvents is crucial in various fields, including environmental science, chemical engineering and pharmaceuticals. In the pharmaceutical industry, for instance, the synthesis of Active Pharmaceutical Ingredients (APIs) is gradually moving from traditional batch processes to processes in continuous flow, due to their enhanced safety and sustainability [1, 2]. However, one of the main hurdles of continuous flow reactors is the precipitation of solids within the reactor, leading to equipment clogging and severe safety risks. Therefore, controlling the solubility of reaction intermediate species and products is crucial for the design and optimization of such processes. Moreover, information on solubility limits is equally important in the stages following the synthesis, such as crystallization, as well as in telescoped reaction sequences.

Traditionally, solubility is calculated using thermodynamic cycles of fusion or sublimation. When considering a fusion cycle, the solubility can be calculated using the definition of ideal solubility which is expressed in terms of the fusion properties of the solute, namely the melting temperature and the enthalpy of fusion [3]. The non-ideality is accounted through activity coefficients (γ). Several approaches are available to estimate γ, including the Wilson equation of state [4] and the Non-Random Two-

Liquid (NRTL) model [5]. However, these models rely on interaction parameters that must be regressed from vapor–liquid equilibrium data, which are often unavailable for many systems of practical interest. In the absence of such parameters, group contribution methods, including UNIFAC [6] can be employed. Methods from quantum chemistry, including both implicit and explicit solvent models, can also be used to compute activity coefficients. Nevertheless, these approaches are computationally demanding, as they require extensive geometry optimization. Moreover, such methods often do not account for the presence of a solid phase, limiting their applicability for accurate solubility predictions. More recently, purely data-driven techniques, including quantitative structure-property relationship (QSPR) models [7] and machine learning (ML) [8] approaches have been promising in predicting molecular properties. In particular, ML has emerged as a valuable alternative to traditional modeling approaches in chemical engineering applications, enabling fast and accurate predictions of complex physicochemical properties.

Algorithms for solvent optimization are closely tied to the accuracy of the predictive models they rely on. Although several studies exist, most of them employ group-contribution approaches as predictive models [9]. In our recent work, we developed robust deep learning models for solubility prediction in mixtures of organic solvents [10]; thanks to these accurate and predictive models, it is now possible to develop optimization strategies to identify solvent mixtures and experimental conditions that optimize solubility. Such tools enable a more efficient exploration of solvent space, reduce the experimental costs and improve the process efficiency. The main limitation of using these ML models is the computational cost associated with their evaluation; direct evaluation of all possible solvent combinations with predictive deep learning models is often computationally prohibitive due to the combinatorial explosion of possible mixtures, especially when considering multiple solvents beyond binary systems.

Bayesian Optimization (BO) [11] is an effective approach for optimizing costly objective functions when their analytical form is unknown or difficult to evaluate, as is the case with graph neural networks. In BO, a surrogate model, often a Gaussian process (GP), is constructed to approximate the expensive objective function. This surrogate is then used to guide the selection of new evaluation points through an acquisition function, which balances exploration and exploitation. By leveraging the acquisition function, BO can efficiently identify maxima or minima with a limited number of function evaluations. A key strength of this approach is its ability to handle both continuous and discrete variables.

In this work, we present an optimization tool based on BO to identify optimal solvent combinations, compositions, and temperatures that maximizes the solubility of APIs. The method leverages a multicomponent directed-edge message passing neural network (D-MPNN) [12] to predict the solubility in ternary systems, i.e., an API dissolved in two solvents [10]. BO is employed to efficiently explore the solvent design space by constructing a GP surrogate model of the mixture D-MPNN, using a radial basis function kernel defined over continuous input variables. Since the molecular identifiers for the solvents cannot be directly incorporated into the Gaussian process kernel, three representation strategies were adopted: integer labels, molecular descriptors, and deep embeddings. This study primarily investigates the selection of chemical representations for the optimization algorithm, aiming to develop a fully automated digital framework that allows users, given a target molecule and a set of solvents, to determine the optimal conditions for pharmaceutical substance design, from chemical synthesis to purification.

BACKGROUND

Directed-Message Passing Neural Networks

D-MPNNs are a variant of message passing neural networks (MPNNs) [13] in which, in contrast to standard MPNNs, each bond is represented as two directed edges, allowing messages to flow explicitly from a source atom to a target atom. For each directed edge, a message is computed in a message passing function M_t (eq. 1) based on the features of the source atom, neighboring atoms, and the connecting bond. These messages update the hidden states of the edges through learned transformations and non-linear activations (e.g., ReLU) (eq.2). The initial hidden state for each edge is created by concatenating the features of the source atom and the directed bond. This combined vector is transformed using a learned weight matrix and an activation function (eq.3). The final hidden state for a certain atom, h_v, is computed by using the final message, m_v, obtained by summing all the messages from the previous steps (eq.4). The final message is then used together with the initial atom features, x_v, to give a representation of the atom. Then all the atom hidden states are aggregated, e.g. summed, and a feature vector for the molecule is obtained (eq.5).

$$m_{vw}^{t+1} = \sum_{k \in N\{(v)\setminus w\}} M_t\,(x_v, x_k, e_{vw}) \quad (1)$$

$$h_{vw}^{t+1} = U_t(h_{vw}^{t}, m_{vw}^{t+1}) = \tau(h_{vw}^{0}, + W_m\, m_{vw}^{t+1}) \quad (2)$$

$$h_{vw}^{0} = \tau(W_i\, cat\,(x_v, e_{vw})) \quad (3)$$

$$m_v = \sum_{k \in N(v)} h_{kv}^{T} \quad (4)$$

$$h = \sum_{v \in G} h_v \quad (5)$$

Finally, the molecular vector is passed through a neural

network to predict the properties of interest.
The D-MPNN architecture is implemented in Chemprop [12], one of the most widely used software tools in chemistry.

Bayesian Optimization

Bayesian Optimization (BO) is a probabilistic framework for the global optimization of functions $f(\mathbf{x})$ that are costly to evaluate or lack analytical expressions. The primary goal of BO is to find a value of $\mathbf{x}^*$, which is the globally optimal value of $\boldsymbol{x}$, that maximizes the objective function $f(\mathbf{x})$ over a predefined domain $\mathcal{X} \subset \mathbb{R}^d$

$$\mathbf{x}^* = \arg\max_{\mathbf{x}\in\mathcal{X}} f(\mathbf{x}) \tag{6}$$

In some cases, evaluating $f(\mathbf{x})$ requires substantial computational resources or experimental effort. Moreover, BO is particularly useful when the objective function is treated as a black box, meaning that only function evaluations are available while derivatives are either inaccessible or difficult to compute (e.g., D-MPNN models). Bayesian optimization addresses this challenge by constructing a probabilistic surrogate model that approximates the objective function, often a Gaussian Process, while quantifying uncertainty. A GP is defined by a mean function $\mu(\mathbf{x})$ and a covariance kernel $k(\mathbf{x}, \mathbf{x}')$

$$f(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \tag{7}$$

Given a set of observations $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, where $y_i = f(\mathbf{x}_i) + \epsilon$ and $\epsilon \sim \mathcal{N}(0, \sigma_n^2)$ represents observational noise, the GP posterior at a new input $\mathbf{x}_*$ is Gaussian with mean and variance

$$\begin{aligned}\mu_n(\mathbf{x}_*) &= \mathbf{k}_*^\top(\mathbf{K} + \sigma_n^2\mathbf{I})^{-1}\mathbf{y}, \sigma_n^2(\mathbf{x}_*) \\ &= k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top(\mathbf{K} + \sigma_n^2\mathbf{I})^{-1}\mathbf{k}_*\end{aligned} \tag{8}$$

where $\mathbf{K}$ is the covariance matrix of training points and $\mathbf{k}_*$ is the covariance vector between $\mathbf{x}_*$ and the training points. This probabilistic prediction allows BO to identify regions of high uncertainty and guide exploration efficiently. The next point to evaluate is chosen by maximizing an acquisition function $\alpha(\mathbf{x}; \mathcal{D}_n)$, which balances exploration and exploitation. Common acquisition functions include Expected Improvement (EI), Probability of Improvement (PI), and Upper Confidence Bound (UCB), each combining the GP mean and variance in different ways to guide the search. The BO procedure is iterative: the surrogate is updated with new observations, the acquisition function is optimized to propose the next evaluation point, and the objective function is sampled. Finally, this loop continues until convergence, or a maximum number of evaluations is reached. By efficiently leveraging both the predicted mean and uncertainty, Bayesian Optimization reduces the number of expensive function evaluations, making it particularly suited for applications where computational or experimental costs are significant.

METHODOLOGY

Solvent Feature Encoding

In this section we describe the three molecular representations employed for the BO kernel: integer enumeration, numerical descriptors and deep embeddings. For the integer enumeration, all unique solvents and solutes were assigned discrete numerical identifiers. Specifically, the solvent SMILES were extracted, mapped to integer values starting from one and stored in a dictionary, where the keys correspond to the SMILES strings and the values to the assigned integers. The second encoding method employs molecular descriptors. Numerical descriptors were computed using RDKit (version 2024.09.6). Each solvent SMILES was converted into an RDKit mol object, from which a set of physicochemical and structural descriptors was calculated. These included molecular weight, LogP, topological polar surface area, numbers of hydrogen bond donors and acceptors, number of rotatable bonds, ring count and fraction of sp^3-hybridized carbons. The descriptors were manually selected to capture relevant chemical properties for solubility. Finally, solvent embeddings were generated using the multicomponent directed edge graph neural network trained on single solvent systems and mixtures. To ensure consistency, embeddings were extracted using the same pre-trained model employed as the objective function in the optimization framework (see "Solubility predictor"). The latent representation of the solvents was extracted from the solvent-specific message-passing block of the D-MPNN. For dimensionality reduction and visualization, principal component analysis (PCA) from scikit learn (version 1.7.1) was applied, retaining the first 5, 10, 20, and 30 principal components for comparison. The reduced representations were subsequently used both for visualization and exported in JSON format for downstream analysis.

Solubility Predictor

We employ a D-MPNN extended to multicomponent solvent systems implemented in Chemprop (version 2.1.2) as objective function. Molecular graphs of the solute and solvents are encoded using a shared bond-centered D-MPNN in which messages are defined on directed bonds and propagated for three message-passing steps while excluding reverse-bond information. Initial bond messages are constructed from default atom and bond features (86-dimensional) and projected into a latent space of dimension 300. Message updates are performed using learned linear transformations with ReLU activations, resulting in 300-dimensional directed bond embeddings. After 3 message passing steps, bond representations are aggregated to atom-level embeddings

and pooled using mean aggregation to obtain fixed-size 300-dimensional molecular embeddings for each component. Solvent embeddings are then combined using a mixture aggregation module that incorporates solvent mole fractions in a permutation-invariant manner, yielding a 300-dimensional mixture embedding. The solute and mixture embeddings are concatenated (602 dimensions in total, including global features) and passed to a feed-forward neural network with a single hidden layer of size 300 to predict solubility.

An ensemble of four models was trained on logarithmic solubility data (logS [mol/L]) in binary and ternary systems from the datasets BiggerSolDB [14] and MixtureSolDB [15] respectively. The combined dataset comprises 227 998 datapoints, 9 690 unique solutes, and 132 unique solvents in the range of 243.15-592.9K. Training was performed for 100 epochs with a batch size of 256 for training and 2048 for validation and testing. The mean squared error (MSE) was used as the loss function. The dataset was randomly split into training, validation, and test sets with ratios of 0.89, 0.10, and 0.01, respectively. The Chemprop default optimizer (Adam) and learning rate were used, and the predictive performance of the model was evaluated on a held-out test set. Figure 1 shows a schematic representation of the multicomponent D-MPNN architecture.

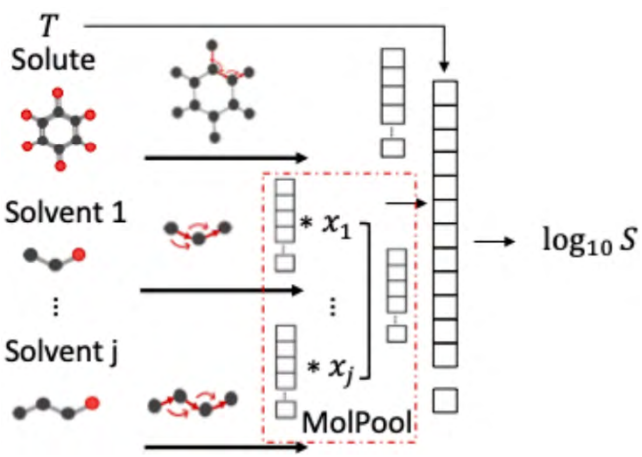

Figure 1. schematic representation of the D-MPNN architecture employed in this work as the objective function within the optimization framework.

Optimization algorithm

As the optimization algorithm, gp_minimize from the scikit-optimize library (version 0.10.2) was employed. Regarding the search space, different bounds were adopted for the categorical variables depending on the solvent representation (i.e., integers, descriptors, or deep embeddings). In the case of integer enumeration, each solvent variable was represented as a discrete integer corresponding to a unique solvent in the dataset. The search space was therefore discrete, ranging from 1 to the total number of unique solvents.

For descriptors and deep embeddings, the search space was instead constructed directly in the continuous representation space of the solvents. Optimized continuous solutions in the descriptor or embedding space were subsequently mapped back to real solvents using a nearest-neighbor projection, selecting the solvent whose representation is closest in Euclidean distance to the optimized point, without modifying the kernel.

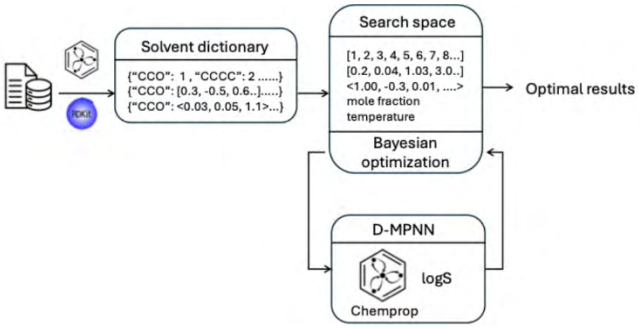

Figure 2. Overview of the optimization workflow: Solvents are extracted from the dataset and encoded as random integers, deep embeddings, or numerical descriptors, which are stored in dictionaries. Bayesian optimization iteratively searches for the best combinations, which are evaluated by the predictor. This loop continues until convergence or a fixed number of iterations is reached.

For all three optimization approaches, the mixture composition was encoded as a continuous variable. The temperature was also modeled as a continuous variable, defined over a real-valued domain bounded by the minimum and maximum temperatures observed in the dataset. The boiling temperature of each pure solvent was estimated using the Joback group contribution method [16]. These boiling temperatures were encoded as numerical values and used together with the solvent representations (integer encoding, molecular descriptors, or embeddings). These values were implemented as a hard-coded constraint in the optimization algorithm, ensuring that solvent mixtures were not proposed at temperatures exceeding the boiling point of the individual pure solvents. This constraint represents an approximation, as the pure component boiling point does not necessary represent the one of the mixtures. Future work could improve this aspect by incorporating predictive models for bubble and dew points. Additionally, to prevent the optimization algorithm from proposing solvent combinations that are implausible from a miscibility standpoint, a constraint based on the water–ethanol partition coefficient was introduced. Future work will focus on extending this framework by incorporating more accurate thermodynamic models for miscibility. Finally, the BO procedure was initialized with 10 random samples and employed

expected improvement as acquisition function. Figure 2 shows an overview of the optimization tool.

RESULTS

The performance of the optimization framework was evaluated on a dataset consisting of 14 299 data points, 66 solutes and 38 solvents in the temperature range 263-367 K. Before reporting the results of the optimization algorithm, we evaluate the performance of the D-MPNN predictive model for the solubility of these solutes. The performance of the model to predict solubility in mixtures of organic solvents for a broader set of solutes and solvents was already extensively evaluated in prior work. In Figure 3, a parity plot shows the performance on the solubility of the 66 tested solutes in the experimentally reported solvents mixtures. Also, for these solutes, a good performance is achieved, indicating that the predicted solubility in new, optimized solvents can be expected to be reliable. Nevertheless, further experimental validation is needed to support these conclusions.

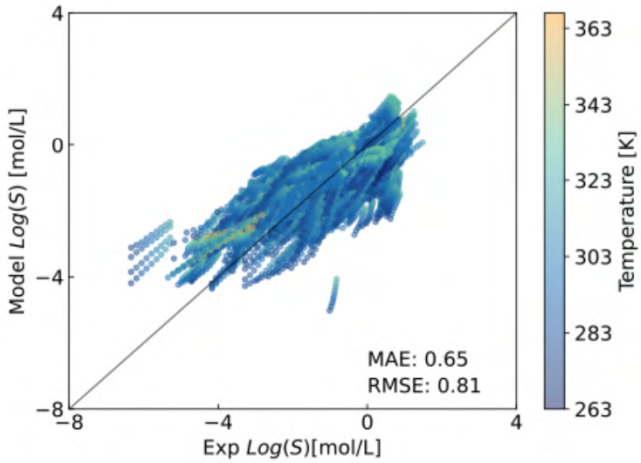

Figure 3. performance of the D-MPNN evaluated on the dataset used in this work, which comprises 66 solutes, 38 solvents, and a total of 14 299 datapoints.

As an example of the performance of the optimization algorithm, we report mixtures proposed by the tool for the three encoding strategies that yield high predicted solubilities. For the enumeration approach, the tool reaches a predicted solubility of 1.4950 $\log_{10}$ mol·L^{-1} for the compound octadecane in a mixture of 4-methylpentan-2-one (0.1 mole fraction) and hexane (0.9) at 333.26 K. This solution is chemically feasible, as both solvents are known to be suitable for octadecane, and the temperature is below their boiling points. The dataset used in this study contains experimentally measured solubilities for only a limited number of solvent mixtures. For octadecane, available experimental data are limited to mixtures of ethanol and cyclohexane. The highest experimentally measured solubility for octadecane was obtained in a mixture of ethanol (0.4) and cyclohexane (0.6) at 297.45 K, with a value of 0.8506 $\log_{10}$ mol·L^{-1}. For the descriptors-based approach, we report the example of propylparaben in a mixture of ethyl acetate (0.2) and dimethylformamide (0.8) at 318.39 K, which was found to have the highest solubility among all solute–solvent combinations. This mixture is feasible because the polar solvents are compatible with the polarity of propylparaben. The optimization algorithm predicted a solubility of 0.8487 $\log_{10}$ mol·L^{-1} for the compound, exceeding the highest experimentally measured solubility in the dataset (0.6249 $\log_{10}$ mol·L^{-1} in pure methanol at 298 K). This demonstrates that the use of solvent mixtures can enhance solubility compared to single-solvent systems. Using the embeddings approach, the maximum predicted solubility (1.6623 $\log_{10}$ mol·L^{-1}) was found for octadecane in a mixture of hexane (0.25) and n-heptane (0.75) at 338.7 K, using an embedding size of 5; since this embedding size yielded the best performance, only results for this setting are reported. The temperature is below the boiling points of both solvents, making the mixture chemically appropriate for octadecane, a paraffin. These results highlight that the embeddings-based approach successfully identified the optimal solvent combination for octadecane, outperforming both other encoding schemes (the descriptor-based model reaches 0.5333 $\log_{10}$ mol·L^{-1}) and the experimental data.

In Figure 4, we present two examples, fumaric acid and 5-nitro-8-hydroxyquinoline, for which we compare the performance of the three representations. In the first case, the enumeration approach achieves higher predicted solubility than both the experimental value and the two encoding schemes. In contrast, for 5-nitro-8-hydroxyquinoline, the chemistry-informed representations (deep embeddings and descriptors) yield better results. For fumaric acid, the enumeration procedure predicts the highest solubility (0.8273 $\log_{10}$ mol·L^{-1}) in a mixture of ethyl acetate (0.1) and dimethylformamide (0.9) at 324.24 K. Descriptors suggest best solubility (-0.0355 $\log_{10}$ mol·L^{-1}) to be in a mixture of diethylene glycol monoethyl ether (0.6) and n-methylpyrrolidone (0.4) at 327.1 K. Finally, for the deep embeddings, the optimal mixture corresponds to diglyme (0.1) and n-methylpyrrolidone (0.9) at 336.97 K (0.0625 $\log_{10}$ mol·L^{-1}). The highest experimentally measured solubility for this molecule, fumaric acid, was obtained in a methanol (0.8) and water (0.2) mixture at 333 K, 0.0898 $\log_{10}$ mol·L^{-1}.

For the other compound, 5-nitro-8-hydroxyquinoline, the best mixture predicted by the enumeration procedure consists of acetone (0.2) and cyclohexanone (0.8) at 300.94 K, yielding −0.2633 $\log_{10}$ mol·L^{-1}. In contrast, the deep embeddings suggest 2-butanone (0.05) and cyclopentanone (0.95) at 324.2 K, with a predicted solubility of 0.1172 $\log_{10}$ mol·L^{-1}. Finally, the descriptor-based model predicts the highest solubility (−0.1440

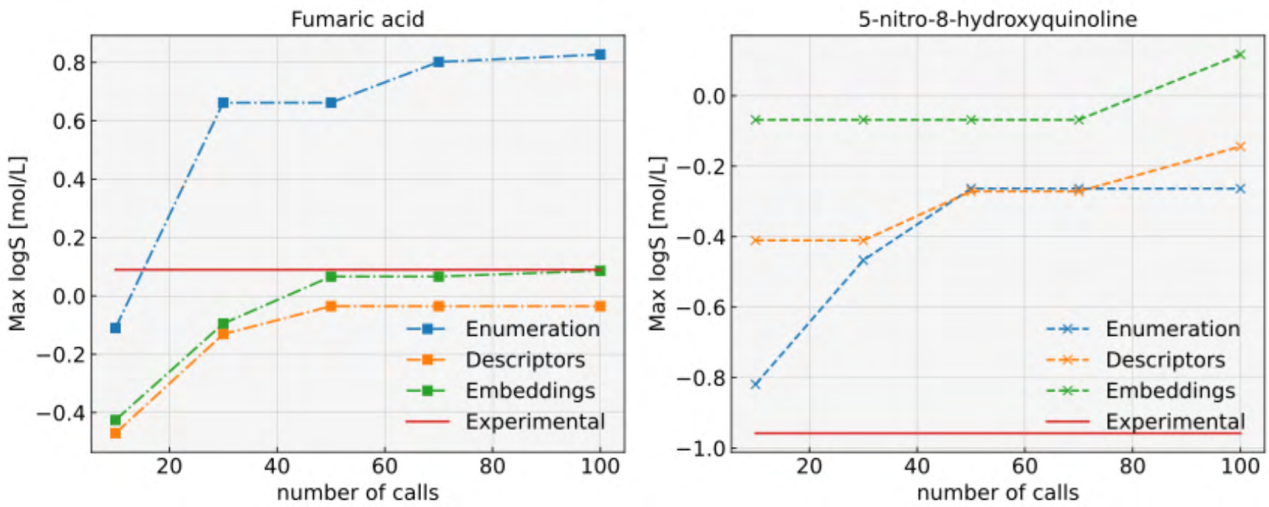

Figure 4. Maximum solubility as a function of the number of optimization calls for the three representations, shown for fumaric acid and 5-nitro-8-hydroxyquinoline. For fumaric acid, integer enumeration and embeddings yield the best performance, whereas for 5-nitro-8-hydroxyquinoline all three representations achieve solubility values higher than the experimental reference. The red continuous line indicates the highest experimentally measured solubility for fumaric acid and 5-nitro-8-hydroxyquinoline.

$\log_{10}$ mol·L^{-1}) in a mixture of pyridine (0.35) and propionic acid (0.65) at 307.5 K. The highest experimental value was -0.9578 $\log_{10}$ mol·L^{-1} found for the mixture ethyl acetate (0.77) and methanol (0.33) at the temperature of 313.15K.

We further analyzed the maximum solubility achieved for each compound in the original curated experimental dataset and compared it with the maxima predicted by the optimization algorithm. Under the enumeration scheme, the algorithm proposed improved solvent mixtures for 36 out of 66 solutes, indicating that for more than half of the compounds it can suggest solvent mixtures that improve solubility as compared to the best experimental measurement. For the embeddings-based approach, 33 solutes exhibited higher predicted solubility than in the original experimental dataset, whereas for the descriptors-based approach, 21 out of 66 solutes showed improved solubility (see table 1). Table 1 summarizes the results for the three enumeration schemes in terms of improvement with respect to the experimental data, mean gain, and mean loss.

These results highlight several important points. First, the enumeration and embedding approaches are the most effective at identifying solvent mixtures that maximize solubility. Second, the proposed method is computationally efficient and improve solubility for a substantial subset of solutes, demonstrating their practical utility for experimental planning. In practice, this approach can significantly reduce the cost associated with testing large numbers of solvent mixtures experimentally. Third, all predicted solutions are chemically reasonable, as they account for solute–solvent, solvent-solvent compatibility and safe temperature ranges.

Table 1: Summary of the three molecular representations in terms of improvement relative to experimental data.

Representation	Better than experiments	Worse than experiments	Mean gain	Mean loss
Enumeration	36	30	0.4881	-0.654
Descriptors	21	45	0.5628	-0.644
Embeddings	33	33	0.5554	-0.599

CONCLUSIONS

In this work, we developed an optimization workflow based on directed message passing neural networks to identify, for each solute, the optimal combination of solvents, mole fractions, and temperatures that maximize solubility. Three different solvent representations were investigated: integer enumeration, physicochemical descriptors, and learned deep embeddings. Our results show that, for a curated dataset with a limited number of solvents and solutes, enumeration and embedding-based representations are the most effective at identifying solvent mixtures that maximize solubility. Both the enumeration strategy and the embeddings-based approach were able to propose improved solvent mixtures for 36 out of 66 solutes and 33 out of 66 solutes whereas the descriptors-based approach led to improved

solubility for a smaller subset of compounds (21 out of 66). To ensure physically meaningful predictions, solvent constraints were imposed to avoid phase separation and to restrict solutions to chemically reasonable temperature ranges. These constraints are necessary because current predictive models do not account for phase stability and lack inherent temperature constraints. However, they should be viewed as approximations that may ultimately be replaced by models that directly enforce thermodynamic constraints within the architecture. Ultimately, the proposed approach lays the groundwork for a fully automated optimization pipeline capable of identifying optimal experimental conditions for organic molecules, spanning the entire workflow from synthesis to crystallization. Future directions include modifying the kernel to directly incorporate physicochemical information, developing improved models for boiling point estimation, enabling the use of optimization algorithms that leverage the structure of the ANN model (e.g., gradient-based methods or mixed-integer linear programming), and experimental validation of the proposed improvements.

ACKNOWLEDGEMENTS

S.B acknowledges the Fonds Wetenschappelijk Onderzoek (FWO) for funding (G021924N). U.D.C acknowledges the Fonds Wetenschappelijk Onderzoek (FWO) for funding (12A4D26N). Resources and services used in this work were provided by the VSC (Flemish Supercomputer Center), funded by the Research Foundation - Flanders (FWO) and the Flemish Government.

AUTHOR IDENTIFIERS

Author ORCIDs:
Simona Buzzi: 0009-0007-9521-7921
Ulderico Di Caprio: 0000-0001-5194-8721
Dominik Bongartz: 0000-0003-1790-0235
Florence Vermeire: 0000-0002-5607-8152

REFERENCES

1. Plutschack MB, Pieber B, Gilmore K, Seeberger PH. The hitchhiker's guide to flow chemistry. Chem. Rev. 117:11796-11893 (2017). https://doi.org/10.1021/acs.chemrev.7b00183
2. Ince MC, Benyahia B, Vilé G. Sustainability and techno-economic assessment of batch and flow chemistry in seven industrial pharmaceutical processes. ACS Sustainable Chem. Eng. 13:2864-2874 (2025). https://doi.org/10.1021/acssuschemeng.4c09289
3. Nordström FL, Rasmuson ÅC. Prediction of solubility curves and melting properties of organic and pharmaceutical compounds. European Journal of Pharmaceutical Sciences 36:330-344 (2009). https://doi.org/10.1016/j.ejps.2008.10.009
4. Sadeghi M, Rasmuson ÅC. On the estimation of crystallization driving forces. CrystEngComm 21:5164-5173 (2019). https://doi.org/10.1039/c9ce00747d
5. Wilson, G. M. (1964). A new expression for the excess free energy of mixing. J. Am. Chem. Soc., 86, 127.
6. Renon, H., & Prausnitz, J. M. (1969). Estimation of parameters for the NRTL equation for excess Gibbs energies of strongly nonideal liquid mixtures. Industrial & Engineering Chemistry Process Design and Development, 8(3), 413-419.
7. Fredenslund, A., Jones, R. L., & Prausnitz, J. M. (1975). Group-contribution estimation of activity coefficients in nonideal liquid mixtures. AIChE Journal, 21(6), 1086-1099.
8. Katritzky AR, Lobanov VS, Karelson M. QSPR: the correlation and quantitative prediction of chemical and physical properties from structure. Chem. Soc. Rev. 24:279 (1995). https://doi.org/10.1039/cs9952400279
9. Leenhouts RJ, Morgan N, Al Ibrahim E, Green WH, Vermeire FH. Pooling solvent mixtures for solvation free energy predictions. Chemical Engineering Journal 513:162232 (2025). https://doi.org/10.1016/j.cej.2025.162232
10. Karia T, Chaparro G, Chachuat B, Adjiman CS. Classifier surrogates to ensure phase stability in optimisation-based design of solvent mixtures. Digital Chemical Engineering 14:100200 (2025). https://doi.org/10.1016/j.dche.2024.100200
11. Buzzi, S., Vermeire, F., Al Ibrahim, E., & Green, W. (2025, September). Predicting the Solubility of Organic Compounds in Mixtures with Thermodynamic Cycles and Machine Learning. In 8th RSC-CICAG/RSC-BMCS Artificial Intelligence in Chemistry 2025, Date: 2025/09/22-2025/09/24, Location: Cambridge (UK).
12. Frazier, P. I. (2018). A tutorial on Bayesian optimization. arXiv preprint arXiv:1807.02811.
13. Yang K, Swanson K, Jin W, Coley C, Eiden P, Gao H, Guzman-Perez A, Hopper T, Kelley B, Mathea M, Palmer A, Settels V, Jaakkola T, Jensen K, Barzilay R. Analyzing learned molecular representations for property prediction. J. Chem. Inf. Model. 59:3370-3388 (2019). https://doi.org/10.1021/acs.jcim.9b00237
14. Gilmer J, Schoenholz SS, Riley PF, Vinyals O, Dahl GE. Message passing neural networks. Lecture Notes in Physics :199-214 (2020). https://doi.org/10.1007/978-3-030-40245-7_10
15. Al Ibrahim E, Morgan N, Müller S, Motati S, Green

WH. Accurately predicting solubility curves via a thermodynamic cycle, machine learning, and solvent ensembles. J. Am. Chem. Soc. 147:45057-45069 (2025). https://doi.org/10.1021/jacs.5c13746

16. Malikov, D., Krasnov, L., Kiseleva, M., Meshcheriakova, E., Kuznetsov, F., Elistratov, V., ... & Bezzubov, S. (2025). MixtureSolDB, dataset of solubility values for organic compounds in binary mixtures of solvents at various temperatures.
17. Devotta, S., & Pendyala, V. R. (1992). Modified Joback group contribution method for normal boiling point of aliphatic halogenated compounds. Industrial & engineering chemistry research, 31(8), 2042-2046.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Coupling Analytical Derivatives with Adjoint Automatic Differentiation in a Modular Process Simulator

Andrés Piña-Martinez[a*] and Jean-Marc Commenge[a]

[a] Université de Lorraine, CNRS, LRGP, F-54000 Nancy, France
* Corresponding Authors: andres-david.pina-martinez@univ-lorraine.fr

ABSTRACT

Modular process simulators are widely used in industry due to their robust and detailed unit operation models. However, their application to gradient-based process optimization remains challenging, as these simulators are typically treated as black boxes, limiting access to internal equations and derivatives. As a result, finite difference methods are commonly employed for gradient estimation, despite their sensitivity to numerical noise and poor scalability. While previous studies have demonstrated the benefits of analytical derivatives in modular simulators, these approaches have largely relied on tangent differentiation modes. This work proposes a non-intrusive methodology that couples analytical derivatives with the adjoint mode of automatic differentiation to efficiently compute gradients for process optimization in modular simulators. The approach preserves the robustness of existing simulation tools by performing simulations normally to convergence, followed by external adjoint-based derivative evaluation using analytically available sensitivities. Theoretical advantages of adjoint differentiation are thus exploited, particularly for optimization problems where the number of decision variables exceeds the number of objective outputs. The proposed framework is evaluated through a case study based on a combined heat and power cycle, using three benchmark objective functions. Gradient computation cost, robustness, and overall optimization performance are compared against forward and centered finite difference methods. Results demonstrate that the analytical adjoint approach significantly reduces gradient computation time while maintaining high robustness, achieving up to 56% reduction in computational cost compared to finite differences. These findings highlight the potential of analytical adjoint differentiation as an efficient and reliable alternative for optimization in modular process simulators.

Keywords: Modelling and Simulations, Optimization, Process Design, Simulation, Energy Systems

INTRODUCTION

Effective decision-making to address societal challenges, such as how energy, fuels, chemicals, and other products should be produced and stored, requires advanced tools and methodologies [1]. In the context of chemical and biochemical processes, Computer-Aided Process Engineering (CAPE) is the field related to the development and application of computational models, methods, and software tools to support their design, simulation, optimization, control, and operation [2]. Most of the CAPE methodologies or tools require the calculation of numerical values for a mathematical function (e.g., process efficiency, product purity, etc), often computed by running the simulation of a process, and its derivatives as well. In particular, derivatives play a central role in product-and-process design optimization, operating parameter sensitivity analysis, or even data reconciliation.

In the case of a steady-state process design, the aforementioned simulation contains material and energy balances, along with physical and thermodynamic property estimations. The simulation is performed by process simulators, which offer a rigorous modelling environment that rely on extensive thermodynamic properties databanks and tailored flowsheeting algorithms. These simulators act as black-box tools and generally rely on sequential-modular or equation-oriented (EO) approaches [3]. The former solves one unit operation at a time,

determining output streams from input streams and equipment parameters. In contrast, the EO approach gathers all the equations of all the units included in the flowsheet resulting in a system of equations which must be solved simultaneously [3].

Concerning derivative calculation, there exist four alternatives: (1) analytical, (2) finite, (3) symbolic, and (4) automatic differentiation. Analytical differentiation consists of determining the exact derivative of a function using calculus rules and formulas. Finite differentiation corresponds to the approximation of derivatives using finite differences formulas. Symbolic differentiation computes derivatives using algebraic manipulations by computer algebra systems. Finally, automatic differentiation applies the chain rule systematically to computer programs to compute derivatives with machine precision.

In the context of process simulation, various of the differentiation techniques presented above have been employed. EO simulators compute derivatives analytically [4], which gives virtually direct access to Jacobians, gradients and sensitivities with no numerical error. On the contrary, modular ones, due to their black-box nature, use finite differentiation for derivative calculation, which may result in inaccurate derivatives due to truncation errors and an elevated computational cost.

Several studies have been focused on analytical derivative calculation in modular process simulators. Early work by Chen and Stadtherr [5] highlighted the benefits of analytical derivatives in flowsheet calculations, while Chan and Prince [6] showed improved accuracy over methods using finite differentiation. Wolbert et al. [4] demonstrated the benefits from exact derivatives for flowsheet optimization. Nevertheless, all of these methods effectively employed a tangent, AD-like differentiation mode, propagating derivatives from input variables to output variables [7]. The objective of this work is to explore the capabilities of coupling analytical derivatives with the adjoint mode of AD, since, in theory, it offers better performances for problems such as optimization, in which the number of input variables is larger than the number of output variables [8].

The paper is structured as follows: background information is provided in Section 2, the methodology is described in Section 3, the case study and benchmarks are presented in Section 4, results are discussed in Section 5, and conclusions are drawn in Section 6.

BACKGROUND

This section provides the necessary background for introducing the proposed methodology. It first outlines the structure and numerical operation of modular process simulators, highlighting their modular organization and coupling with external tools. Next, existing strategies for derivative calculation in such simulators are discussed, with particular attention to intrusive and non-intrusive approaches. Finally, the fundamentals of automatic differentiation are introduced, focusing on the tangent and adjoint modes that underpin the derivative propagation framework used in this work.

Structure of a modular process simulator

A module in a modular process simulator determines the output stream values of a unit (η_k), given its input stream values (ξ_k) and the unit parameter values (p_k). This is achieved by solving a set of equations representing the model of that unit, i.e., mass and energy balances, phase equilibrium [3].

In modular process simulators, once the flowsheet units are connected through material and energy balance equations, the simulator automatically establishes a calculation sequence and applies internal numerical algorithms to solve the resulting system. These simulators typically embed robust numerical methods for tasks such as root-finding, convergence of recycle loops, and local optimization of unit or flowsheet variables. However, their built-in capabilities do not preclude the use of external tools. In many advanced applications, the process simulator can be coupled with external optimization solvers, allowing users to exploit more specialized or large-scale optimization algorithms, integrate custom objective functions and constraints, or link the simulation to broader decision-making frameworks such as process design, control, or economic optimization.

Derivative calculation with modular process simulators

Two strategies can be considered for derivative calculation with a modular process simulator. The first one consists of "opening" the simulator's source code and implementing derivative expressions for function relationships within each unit module. Then, derivatives can be evaluated as part of the simulation process [4]. The second strategy does not modify the simulator's source code. In contrast, the simulation is performed first, followed by a derivative calculation procedure that receives the calculated output stream values and unit parameter values.

The first strategy provides exact gradient information that can be used by both internal or external numerical methods. For instance, if a flowsheet contains a recycle stream, it is necessary to implement a root-finding algorithm that may require gradient information. Similarly, the second strategy is also able to provide exact gradient information. The difference lies in the fact that such gradients are obtained a posteriori, based on the converged simulation results, rather than being computed and propagated internally during the solution of the flowsheet equations.

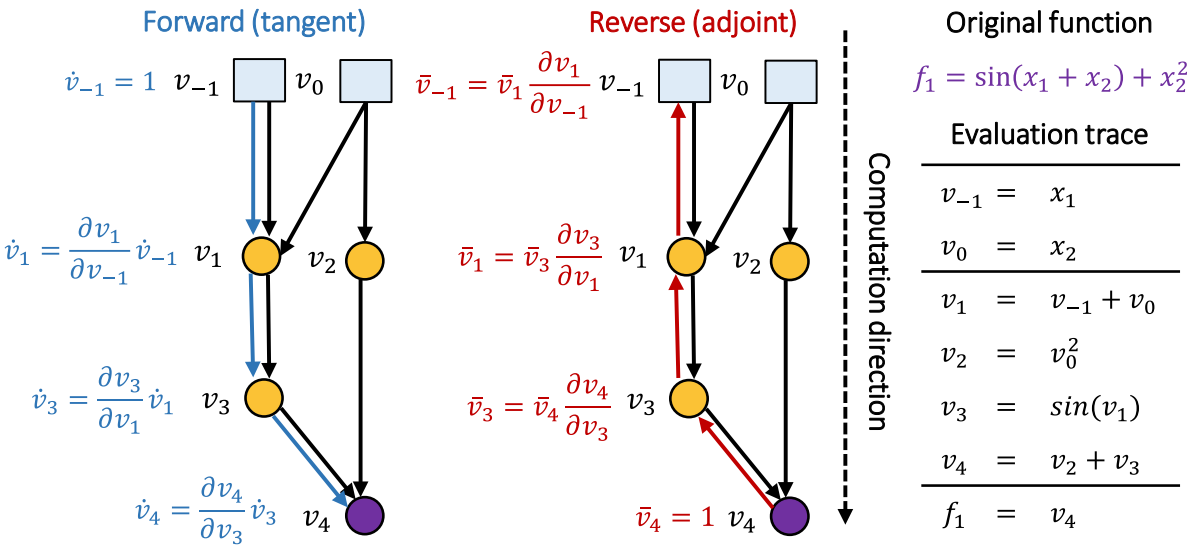

Figure 1. Sketch of the computational graph for the computation of f_1 and its derivatives.

Automatic differentiation

Automatic Differentiation (AD) is based on the observation that a sequence of code statements can be regarded as the composition of the functions representing the computation performed by each statement. This composition is represented by a computational graph (CG). In Figure 1, vertices (v_i) represent each sub-function composing the original function (f_1), and an arc runs from v_j to v_i when v_i depends on v_j ($j \prec i$). The roots of the graph (blue) represent the independent variables (x), and the leaves (purple) represent the dependent ones (f). Then, AD works by propagating gradients in the computational graph using the chain rule. Two basic modes exist: the tangent (forward) and adjoint (reverse) mode. Forward mode propagates tangents values ($\dot{v}_i$) in the same direction as function evaluation, while reverse mode propagates adjoint values ($\bar{v}_j$) backward [8].

Automatic differentiation can be implemented through: (1) operator overloading, where standard arithmetic operations are redefined on custom data types that store both values and derivatives, enabling derivatives to be computed as the program runs. (2) Source code transformation, which analyses the original program and generates new code that simultaneously computes the function and its derivatives, often yielding highly optimized performance. (3) Tape-based execution records operations during the forward pass and replays them in reverse to efficiently propagate gradients, a technique widely used in modern deep learning frameworks [8].

For the following sections, consider a vector of input variables as $\boldsymbol{x} \in \boldsymbol{X} \subset \mathbb{R}^n$, a vector function (or vector valued) $F: \boldsymbol{X} \longmapsto \mathbb{R}^m$, which is the composition of a sequence of l once continuously differentiable elemental functions $(\varphi_i)_{i=1,\dots,l}$, representing each internal node. The Jacobian is a well-defined matrix-valued function $F': \boldsymbol{X} \longmapsto \mathbb{R}^{m \times n}$.

Tangent mode

The first derivative vector in tangent mode depends on x and a certain seed direction $\dot{x}$. The desired output is $\dot{y}$, defined as follows:

$$\dot{y} = F'(x)\dot{x} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1}\dot{x}_1 + \cdots + \frac{\partial f_1}{\partial x_n}\dot{x}_n \\ \vdots \\ \frac{\partial f_m}{\partial x_1}\dot{x}_1 + \cdots + \frac{\partial f_m}{\partial x_n}\dot{x}_n \end{bmatrix} \tag{1}$$

One might be tempted to evaluate the full Jacobian $F'(x)$ and then multiplying it with the direction $\dot{x}$. In turn, the Jacobian is evaluated by tangent propagation of each element φ_i composing the function F. In consequence, the first step is to define the seed direction $\dot{x}$ for which $\dot{y}$ will be calculated. Then, all the nodes $\dot{v}_i$ are calculated as follows:

$$\dot{v}_i = \sum_{j \prec i} \frac{\partial v_i}{\partial v_j} \dot{v}_j \tag{2}$$

where $j \prec i$ means that node j precedes node i. Consider the case with $\dot{x}_j = 1$ and $\dot{x}_{i \neq j} = 0$. Then, $\dot{y}$ corresponds to the jth column of $F'(x)$:

$$\dot{y} = \begin{bmatrix} \frac{\partial f_1}{\partial x_j} & \cdots & \frac{\partial f_m}{\partial x_j} \end{bmatrix}^T \tag{3}$$

It is then possible to compute a full Jacobian matrix ($F': \boldsymbol{X} \longmapsto \mathbb{R}^{m \times n}$) in n forward sweeps. In the example of Figure 1, $n = 2$ and $m = 1$. It is thus necessary to set $\dot{x} = (1, 0)$ and $\dot{x} = (0, 1)$ to obtain $\frac{\partial f_1}{\partial x_1}$ and $\frac{\partial f_1}{\partial x_2}$, respectively. Note

that if $m > 1$, the number of required forward sweeps does not increase.

Adjoint mode

In the adjoint mode, the objective is to compute the gradient of a function defined as the weighted sum of the components (f_k) of a vector-valued function $F(x)$, namely,

$$s = \sum_{k=1}^{N} \bar{y}_k f_k \tag{4}$$

where $\bar{y}$ correspond to the weight vector. The desired output is $\bar{x}^T = \nabla s = \bar{y}^T \cdot F'(x)$:

$$\bar{x}^T = \left[\bar{y}_1 \frac{\partial f_1}{\partial x_1} + \cdots + \bar{y}_m \frac{\partial f_m}{\partial x_1} \quad \ldots \quad \bar{y}_1 \frac{\partial f_1}{\partial x_n} + \cdots + \bar{y}_m \frac{\partial f_m}{\partial x_n}\right] \tag{5}$$

In order to compute the gradient of s, the first step is to define the weight vector. Then, all the nodes $\bar{v}_j$ are calculated as follows:

$$\bar{v}_j = \sum_{i \succ j} \bar{v}_i \frac{\partial v_i}{\partial v_j} \tag{6}$$

where $i \succ j$ means that node i follows node j. Consider the case with $\bar{y}_i = 1$ and $\bar{y}_{j \neq i} = 0$. Then, $\bar{x}^T$ corresponds to the ith row of $F'(x)$:

$$\bar{x}^T = \left[\frac{\partial f_i}{\partial x_1} \quad \ldots \quad \frac{\partial f_i}{\partial x_n}\right] \tag{7}$$

It is then possible to compute a full Jacobian matrix $(F' : \boldsymbol{X} \longmapsto \mathbb{R}^{m \times n})$ in m reverse sweeps. In the example of Figure 1, it is thus necessary to set $\bar{y} = 1$ to simultaneously obtain $\frac{\partial f_1}{\partial x_1}$ and $\frac{\partial f_1}{\partial x_2}$. Note that if $n > 2$, the number of required backward sweeps does not increase.

METHODOLOGY

In this work, we propose a non-intrusive methodology that couples analytical derivatives with the adjoint mode of automatic differentiation to perform gradient computation in modular process simulators. The methodology proposed in this work is composed by four main steps:

1. Definition of the flowsheet, specifications, objective functions and action variables.
2. Definition of the computational graph of the process model.
3. Implementation of analytical derivatives at the unit-operation level.
4. Generation of the computation procedure for derivatives using the adjoint mode of AD.

Implementation strategy

Once the flowsheet, specifications, objective functions and action variables are defined in the simulator. The computational graph can be automatically extracted from the simulator or manually defined by the user by following the unit operation calculation sequence. At each new equation of a given unit operation module, new nodes, if they do not exist yet, are included in the graph. In this work, the computational graph was manually defined.

At the unit-operation level, analytical derivatives involving mass and energy balances are hand-coded in a user-defined module or custom scripting environments depending on the simulator native features. If the derivative of a physical property, such as enthalpies or entropies, is required, they are obtained via simulator APIs, such as Simulis Thermodynamics, Aspen Properties, or CoolProp.

Finally, all the derivatives required for the computation of adjoints for each node of the computational graph are linked manually following Equation 6.

Implementation example

The gradient calculation procedure is presented through an example based on the Rankine cycle described in Figure 2. Consider a function equal to the net power of the cycle: $\dot{W}_{net}$, whose calculation is described by the set of equations S1 and performed by a process simulator The variables involved in the calculation are the turbine and pump powers $(\dot{W}_T, \dot{W}_P)$, their isentropic efficiencies (η_T, η_P), enthalpies (h_i) calculated at temperature and pressure conditions (T_i, P_i), enthalpies (h_i^{is}) calculated at isentropic conditions and pressure P_i. $\dot{m}$ represents the mass flow rate. $T_1, P_1, P_2, T_3, P_3, P_4, \eta_T, \eta_P$ and $\dot{m}$ are known.

The cycle flowsheet is created by the user, and the model will be generated and solved automatically by the process simulator. Then, the computational graph is manually generated and analytical derivatives are coded following the set of equations S1. At this point, it is necessary to define which variables from the simulation will be reused in the gradient calculation and which not. For instance, the enthalpies or entropies will be reused, while enthalpy or entropy derivatives will be calculated in the AD interface. This approach preserves the robustness of the existing simulator while leveraging adjoint efficiency.

In execution time, if $\frac{\partial \dot{W}_{net}}{\partial T_1}$ is required, the simulation will be executed first. Then, the calculated output stream values and unit parameter values that will be reused in the AD procedure are stored. This is followed by the calculation starting with $\bar{\dot{W}}_{net}$ and finishing with $\bar{T}_1$.

CASE STUDIES AND BENCHMARK

To evaluate the proposed framework for AD adjoint mode with analytical derivatives within modular process simulators, a series of optimization tests, based on a Combined Heat and Power (CHP) cycle, were conducted. The case study considers the CHP cycle presented in Figure 3. The working fluid in a liquid state is compressed

by a pump (1-2) from a low pressure (LP) to an intermediate pressure level (IP) corresponding to $P_{LP} = P^{sat}(T_A)$ and $P_{IP} = P^{sat}(T_M)$, respectively. It is then mixed with a bleeding stream at P_{IP}. The resulting liquid stream is compressed by a pump (3-4) to a high-pressure level (HP) equal to $P_{HP} = P^{sat}(T_H)$. Then, the fluid is preheated, evaporated at constant pressure (4-5). At the expansion stage, part of the fluid is expanded to the IP level and cooled at constant pressure in a condenser in order to provide the required end-user heating power ($\dot{Q}_{user}$) at T_M. The remaining stream is condensed at constant pressure (8-1) to the initial state. The working fluid is neopentane.

Table 1. CHP cycle specifications.

Variable	Value	Units
T_A	298	K
$\dot{m}$	1	kg.s^{-1}
η_T / η_P	0.75 / 0.80	-
x_{v1} / x_{v7}	0.0 / 0.0	-

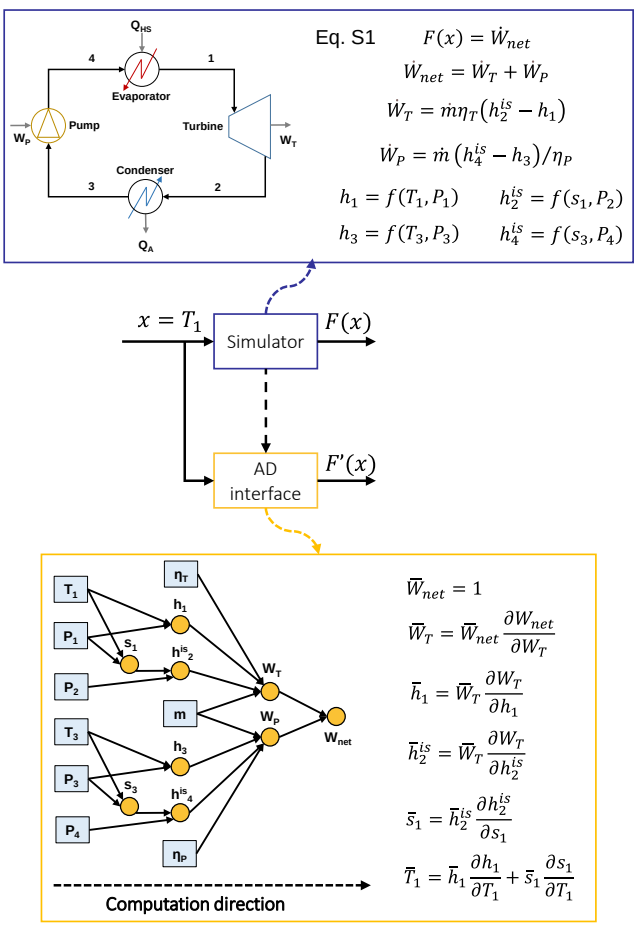

Figure 2. Representation of the communication between the automatic-differentiation based interface and the process simulator for gradient calculation.

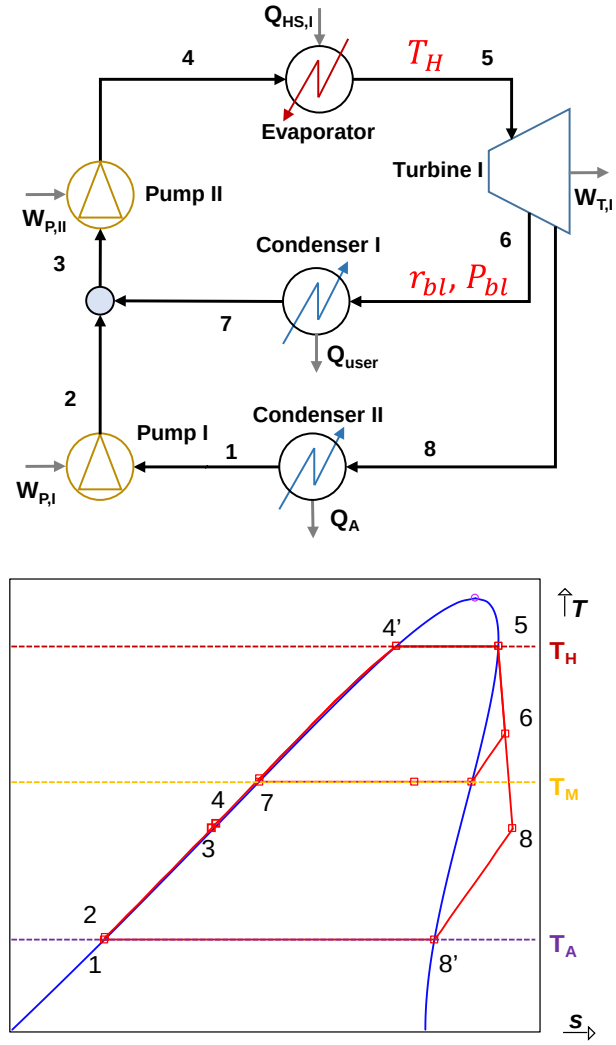

Figure 3. Top: process flow diagram of the studied CHP cycle. Bottom: Temperature-entropy diagram.

For this cycle, nine degrees of freedom (DoF) are initially identified: vapor fractions at states 1 and 7 (x_{v1}, x_{v7}), turbine and pump isentropic efficiencies (η_T, η_P), mass flow rate ($\dot{m}$), low-level temperature (T_A), bleeding ratio (r_{bl}), bleeding pressure (P_{bl}) and outlet temperature at the evaporator (T_H). For benchmarking, only three DoFs are accepted: r_{bl}, P_{bl}, T_H. The specifications of the remaining DoFs are listed in Table 1.

Problem formulation

Given CHP cycle specifications, action variables bounds and an objective function, an optimization problem is formulated to determine optimal operating conditions that minimizes the objective function.

Action variables

Although three DoFs and different combinations are possible, only three cases were considered. If one, two or three action variables are considered, then the

Table 4: Values of a_i, b_i, c_i for the Rosenbrock, Spheric and Exponential functions.

Function	Index	a	b	c
Rosenbrock	0	$\pi - 2$	-	-
	1	250	1	-0.4556
	2	5	1	-0.3684
	3	-	-	-0.5360
Sphere	0	$\pi - 2$	-	-
	1	100	1	0.5443
	2	1	1	0.6316
	3	1	1/2	0.4640
Exponential	1	-100	-20	0.5443
	2	-100	-20	0.6316
	3	-100	-20	0.4640

decision vector corresponds to $x = (r_{bl}), x = (r_{bl}, P_{bl})^T$ or $x = (r_{bl}, P_{bl}, T_H)^T$, respectively. The bounds, or specification if a given variable is not considered in the action variable vector, are listed in Table 2.

Table 2. Action variables defined in this study.

Variable	Bounds (Specification)	Units
r_{bl}	0.05 – 0.95 (NA)	-
P_{bl}	3 – 9 (5.4)	bar
T_H	343.15 – 413.15 (390.65)	K

Objective functions

A common choice for the objective function is the Energy Utilization Factor (EUF): $EUF = -\frac{(\dot{W}_{net} + \dot{Q}_{user})}{\dot{Q}_{HS}}$, where $\dot{W}_{net}, \dot{Q}_{HS}$ and $\dot{Q}_{user}$ represent, respectively, the net mechanical power, the required heat-source duty, and the heating power provided to the end user. One might be tempted to adopt the EUF or other efficiency-based indicator as a benchmark objective, given its clear physical meaning and widespread use in cycle optimization. However, the resulting landscape is largely determined by the coupled physics and constraints of the cycle, which can make features such as smoothness, conditioning, and multimodality difficult to assess independently. Efficiency-based objectives may therefore reflect characteristics of the full thermodynamic formulation, potentially limiting the ability to adjust features such as curvature, scaling, or the presence of local optima.

To address this, the process model is first used to map physical decision variables to $\dot{W}_{net}, \dot{Q}_{HS}$ and $\dot{Q}_{user}$, which are then normalized and used as coordinates of a benchmark test function, creating a controlled optimization landscape over the physically admissible operating space. The benchmark functions implemented in this work are:

Rosenbrock:

$$F_{obj} = \sum_{i=1}^{N-1} \left[a_i \cdot \left(\phi_{i+1} - \phi_i^2\right)^2 + b_i \cdot (1 - \phi_i)^2 \right] + a_0 \quad (9)$$

Spheric:

$$F_{obj} = \sum_{i=1}^{N} \left[a_i \cdot \left(\frac{\phi_i}{b_i}\right)^2 \right] + a_0 \quad (10)$$

Exponential:

$$F_{obj} = \sum_{i=1}^{N} \left[a_i \cdot exp(b_i \cdot \phi_i^2) \right] \quad (11)$$

with $N = 3$, corresponding to the number of physical decision variables ($\dot{W}_{net}, \dot{Q}_{HS}$ and $\dot{Q}_{user}$), and ϕ_i defined as follows:

$$\phi_i = \frac{\alpha_i - \alpha_{i,min}}{\alpha_{i,max} - \alpha_{i,min}} - c_i \quad (12)$$

with $\alpha_1 = \dot{W}_{net}, \alpha_2 = \dot{Q}_{HS}$ and $\alpha_3 = \dot{Q}_{user}$. The values of $\alpha_{i,min}$ and $\alpha_{i,max}$ were obtained via systematic sampling of the admissible operating space for r_{bl}, P_{bl} and T_H, and are reported in Table 3. The values of a_i, b_i and c_i, are constant and reported in Table 4 for each benchmark function.

Table 3. Admissible bounds for $\dot{W}_{net}, \dot{Q}_{HS}$ and $\dot{Q}_{user}$.

Variable	α_{min}	α_{max}	Units
$\dot{W}_{net}$	7197	63362	W
$\dot{Q}_{HS}$	255973	428123	
$\dot{Q}_{user}$	12247	349282	

Computational graph

The computational graph linking the input variables with the objective function is represented in Figure 4. For the sake of readability, the graph is represented by three layers representing the input nodes (9) corresponding to the maximum DoFs, the internal nodes (39) and the output node (1). It is worth recalling that, at most, three of the input nodes represent action variables of the optimization problems. The resulting computational graph is not

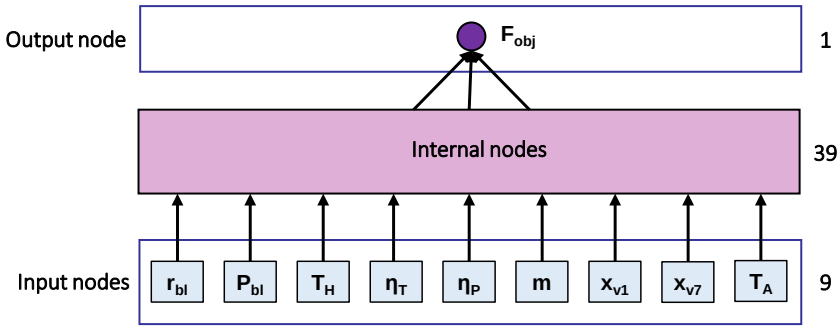

Figure 4. Sketch of the computational graph for the computation of F_{obj} and its derivatives.

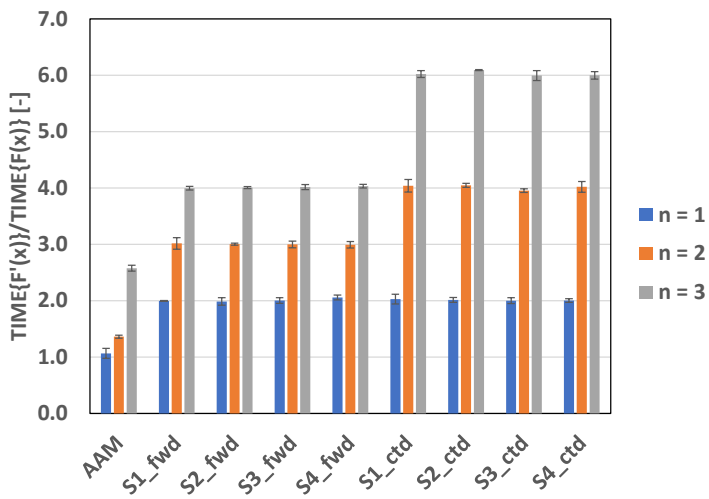

Figure 5. Average relative runtime of the gradient calculation compared to a single function evaluation over all the runs and objective functions. Abbreviations: AAM: analytical adjoint mode. fwd: forward finite difference. ctd: centered finite difference.

a complete graph, i.e., where every pair of distinct nodes is connected by an edge. In consequence, the number of nodes connecting two different input nodes to the output node is not necessarily the same.

Benchmark procedure

Each benchmark test consisted of 30 independent runs, corresponding to 30 different initializations of the design variables. For each run, gradients were computed using both automatic differentiation and finite difference methods. Finite difference gradients were calculated using four perturbation step sizes, defined as: $S_1 = 10^{-2}, S_2 = 10^{-3}, S_3 = 10^{-4}, S_4 = 10^{-5}$. A solution was declared optimal if it satisfied the following $f(x^*) \leq \max(f(z^*) \cdot \epsilon, f(z^*) + \epsilon)$, where $f(x^*)$ denotes the objective function evaluated at the computed optimal solution for a given run, $f(z^*)$ corresponds to the analytically known optimal value, and ϵ represents the tolerance. It was set to 10^{-2} as reference [9].

Computational details

The optimizations were performed on an Intel(R) Core(R) Ultra 7, 3.80 GHz, 32 GB RAM, Windows 11 64-bit operating system. The optimization solver is based on a Sequential Least Squares Programming (SLSQP) algorithm, implemented via pySLSQP, while all simulations were carried out using ProSimPlus 3.7.9.

RESULTS

In this section, gradient computation cost, robustness, and overall optimization performance of the analytical adjoint mode (AAM) are compared against forward and centered finite difference methods. For clarity in the presentation of results, finite-difference gradients computed using centered difference schemes are denoted as S_i_ctd, while those obtained using forward difference schemes are denoted as S_i_fwd, with $i = 1, \ldots, 4$ indicating the corresponding perturbation sizes S_1-S_4.

Gradient computation cost

Figure 5 illustrates the average relative runtime of the gradient calculation with respect to a single function evaluation ($TIME\{F'(x)\}/TIME\{F(x)\}$), computed over all

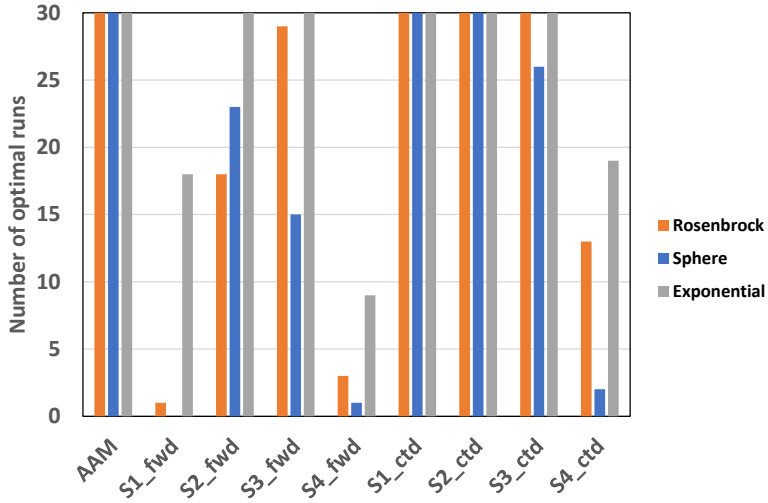

Figure 6. Number of optimal runs. Abbreviations: AAM: analytical adjoint mode. fwd: forward finite difference. ctd: centered finite difference.

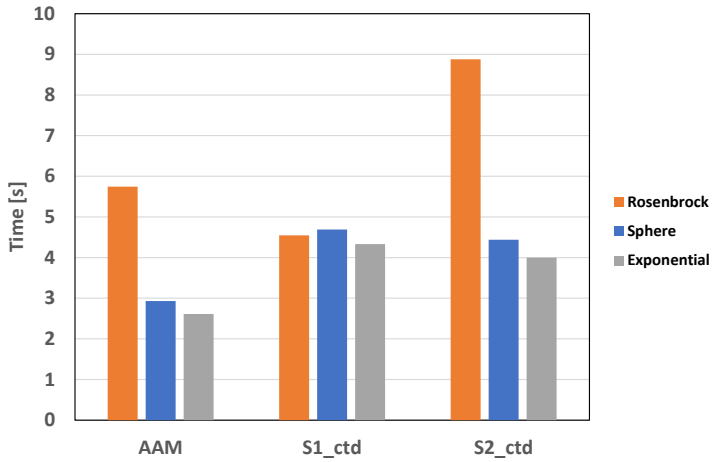

Figure 7. Optimization time of analytical adjoint mode (AAM), and centered differences with a perturbation step of $S_1 = 10^{-2}$ and $S_2 = 10^{-3}$.

runs and objective functions. From this figure, it is possible to verify that the computational cost of the gradient calculation ($F'(x)$) for a function $F(x)$ with n inputs and m outputs using finite differentiation corresponds, as expected, to $(n+1)$ and $(2n)$ function evaluations for forward and centered differences, respectively**.** Adjoint mode of AD with analytical derivatives offers better scaling, requiring, at most, $(3m)$ function evaluations [5]. For the case with $n = 3$, it is around 2.6 times the cost of a function evaluation, which results in a reduction of 35% and 56% compared to forward and centered finite differentiation, respectively.

It is possible to observe that the ratio for AAM is not the same for the three cases. This difference can be explained, in part, by the fact that the computational graph is not a complete graph. For instance, the number of internal nodes involved in the cases with $n = 1$ is 17, $n = 2$ is 26 and $n = 3$ is 37. Nevertheless, it is observed that the ratio does not increase linearly. Although the difference of internal nodes ($\Delta v_{2\to3} = 11$) between cases 2 and 3 is almost the same between cases 1 and to ($\Delta v_{1\to2} = 9$), the change in the ratio is not the same. It would have meant that each node has the same computational cost. In contrast, when $n = 3$, the outlet temperature at the evaporator (T_H) is an action variable and involves nodes in which derivatives of enthalpy, entropy or vapor pressure at a dew point are calculated.

It is also important to highlight that by considering $x = (r_{bl}, P_{bl}, T_H)^T$ as action variables vector among the 9 potential input nodes (Figure 4), only two internal nodes are not visited during the gradient calculation. Such nodes would be visited if $\mathrm{T_A}$ was included as an action variable. Consequently, assuming $n = 3$ as the baseline case, adding one additional action variable ($n = 4$) or up to five additional variables ($n = 9$) should not significantly affect the computational cost of the gradient calculation.

Robustness and runtime

In this section, it is decided to discuss the case with $n = 3$ since it included the highest number of internal nodes for gradient calculation in AAM. In terms of robustness, in Figure 6 it is observed that only three approaches were able to reach 30 optimal runs for each benchmark function: analytical adjoint mode, and centered differences with a perturbation step of 10^{-2} and 10^{-3}.

Optimization times for these three approaches are presented in Figure 7. At a first glance, it is observed that for the Sphere and Exponential functions, the computation time of analytical adjoint mode is reduced by 35%. In the case of the Rosenbrock function, centered differences with a perturbation step of 10^{-2} outperforms by 26% analytical adjoint mode and by 49% centered differences with a perturbation step of 10^{-3}.

It is worth noting that even when using AAM, the optimization time may not follow the trends presented in Figure 5. This is due to the fact that the SLSQP algorithm may require substantially more evaluations of the objective function than of its gradient. This disparity arises because, at each major iteration, SLSQP typically computes the gradient only once to form the quadratic programming subproblem, whereas the objective function may be evaluated multiple times during the subsequent line search and step acceptance procedures. In particular, candidate step lengths along the computed search direction are repeatedly probed to enforce sufficient decrease in a merit function that combines the objective value with constraint or bound violation measures.

CONCLUSION

This work demonstrates the feasibility and benefits of coupling analytical derivatives with the adjoint mode of automatic differentiation in modular process simulators. The proposed non-intrusive methodology enables efficient and robust gradient computation without modifying the internal structure of existing simulation tools. Through a combined heat and power cycle case study, the analytical adjoint approach was shown to outperform finite difference methods in terms of computational efficiency and robustness, particularly for problems with multiple decision variables. While finite difference schemes can occasionally achieve competitive performance, they require careful tuning of perturbation steps, which increases overall optimization effort. In contrast, the analytical adjoint method provides consistent accuracy and favorable scaling properties without additional sensitivity analysis. These results confirm the strong potential of adjoint-based analytical differentiation as a practical solution for gradient-based optimization in industrial modular process simulation environments.

AUTHOR IDENTIFIERS

Author ORCIDs:
Piña-Martinez Andrés: 0000-0003-0691-0075
Commenge Jean-Marc: 0000-0003-2792-9357

REFERENCES

1. Adjiman CS, Sahinidis NV, Vlachos DG, Bakshi B, Maravelias CT, Georgakis C. Process systems engineering perspective on the design of materials and molecules. Ind. Eng. Chem. Res. 60:5194-5206 (2021). https://doi.org/10.1021/acs.iecr.0c05399
2. E.N. Pistikopoulos, A. Barbosa-Povoa, J.H. Lee, R. Misener, A. Mitsos, G.V. Reklaitis, V. Venkatasubramanian, F. You, R. Gani, Process systems engineering – The generation next?, *Comput. Chem. Eng.* 107252-147 (2021) https://doi.org/10.1016/j.compchemeng.2021.10725
3. A.W. Westerberg, H.P. Hutchinson, R.L. Motard, P. Winter. Process Flowsheeting. Cambridge (1979).
4. Wolbert D, Joulia X, Koehret B, Biegler LT. Flowsheet optimization and optimal sensitivity analysis using analytical derivatives. Computers & Chemical Engineering 18:1083-1095 (1994). https://doi.org/10.1016/0098-1354(94)e0020-n
5. Chen H, Stadtherr MA. A simultaneous?modular approach to process flowsheeting and optimization. part I: theory and implementation. AIChE Journal 31:1843-1856 (2004). https://doi.org/10.1002/aic.690311110
6. Chan WK, Prince RGH. Application of the chain rule of differentiation to sequential modular flowsheet optimization. Computers & Chemical Engineering 10:223-240 (1986). https://doi.org/10.1016/0098-1354(86)85004-9
7. Mischler C, Joulia X, Hassold E, Galligo A, Esposito R. Automatic differentiation applications to computer aided process engineering. Computers & Chemical Engineering 19:779-784 (1995). https://doi.org/10.1016/0098-1354(95)87129-2
8. A. Griewank, A. Walther. Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation. SIAM (2008).
9. Rios LM, Sahinidis NV. Derivative-free optimization: a review of algorithms and comparison of software implementations. J Glob Optim 56:1247-1293 (2012). https://doi.org/10.1007/s10898-012-9951-y

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Hand-crafted Feature Fusion for Deep Learning-Based Instance Segmentation in Microfluidics

Wenle Xu[ab], Lin Sheng[ab], Qichen Shang[ab], Mengqi Liu[ab], Tong Qiu[ab*], Kai Wang[ab], Guangsheng Luo[ab]

[a] Department of Chemical Engineering, Tsinghua University, Beijing, China
[b] State Key Laboratory of Chemical Engineering and Low-carbon Technology, Tsinghua University, Beijing, China
* Corresponding Author: qiutong@tsinghua.edu.cn.

ABSTRACT

High-throughput analysis of microfluidic droplets and bubbles is essential for chemical engineering but remains challenging due to the inherent loss of high-frequency details in standard deep learning models. This study proposes a novel Hand-crafted Feature Fusion framework that explicitly integrates physical priors, specifically Local Binary Patterns and Discrete Wavelet Transform, into a two-stage instance segmentation network. We design an adaptive attention-based fusion module embedded within both the Feature Pyramid Network and Region Proposal Network to synergize explicit texture cues with implicit semantic features. Validated on a large-scale dataset comprising over 64000 instances, our method achieves a test mAP of 0.808, significantly outperforming state-of-the-art architectures. Crucially, the framework effectively resolves the detection bottleneck for minute targets and elevates the small-object accuracy to 0.764, representing an improvement of nearly 20% over the baseline. This work demonstrates that incorporating physical priors offers a superior strategy for precise scientific image analysis compared to generic data-driven models.

Keywords: Microfluidics, Computer Vision, Instance Segmentation, Hand-crafted Features

1. INTRODUCTION

Microfluidics involves manipulating fluids at the micrometer scale (typically 10–100 μm), creating a high surface-area-to-volume ratio. This configuration enables superior mass and heat transfer and a narrow residence time distribution (RTD) [1]. Consequently, microfluidics has become a critical platform in chemical engineering, biomedicine, and materials synthesis. The properties of the dispersed phase, such as size, flow pattern, and velocity, are key determinants of mechanistic understanding and product quality in both research and industry. Therefore, high-speed cameras coupled with microscopy serve as a primary means of capturing this essential information.

This high-speed imaging inherently generates a prohibitive volume of visual data, typically yielding thousands of frames per second. Conventionally, quantitative analysis relied on manual measurement using software tools such as ImageJ. While effective for small datasets, this laborious workflow is time-consuming and prone to subjective bias, creating a significant bottleneck for high-throughput analysis. To overcome the limitations of manual measurement, the field initially adopted traditional computer vision algorithms such as Canny edge detection and Watershed segmentation. These methods rely on explicitly engineered hand-crafted features and operate by detecting low-level attributes like pixel-intensity gradients or geometric contours. Although these rule-based algorithms offer mathematical interpretability, they inherently lack the adaptability required to distinguish targets effectively within complex or noisy microfluidic backgrounds.

The subsequent emergence of deep learning precipitated a fundamental shift from manual feature engineering to automatic representation learning. Convolutional Neural Networks (CNNs), such as YOLO and Mask R-CNN, surpassed earlier methods by extracting hierarchical features directly from raw data. More recently, this evolution has extended to large-scale Transformer-based architectures, such as the Segment Anything Model (SAM), which leverage self-attention mechanisms to capture global context. However, while these data-driven approaches provide enhanced robustness and

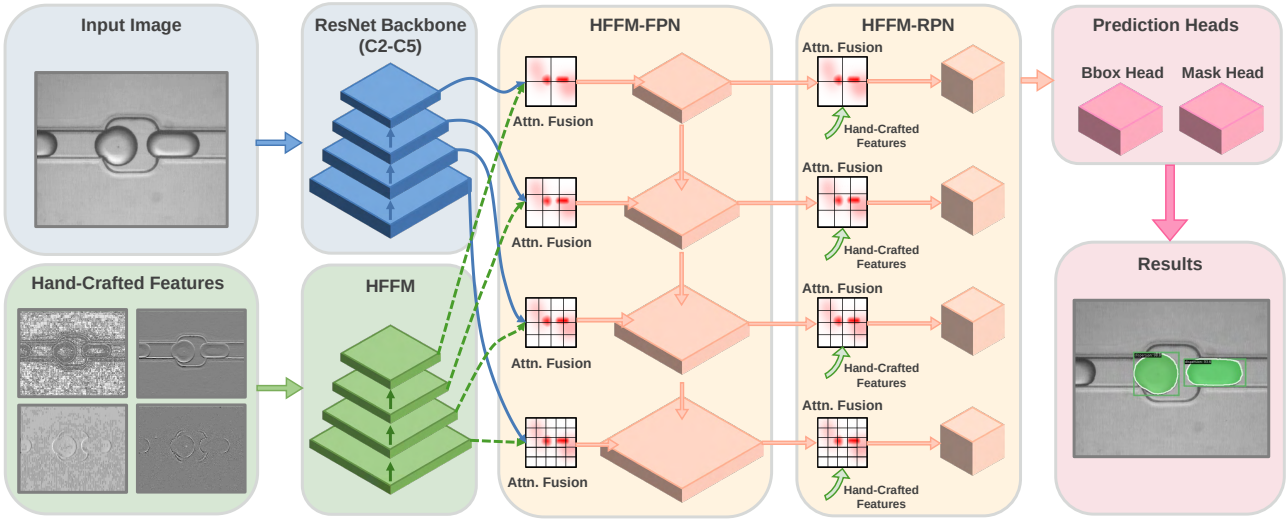

Figure 1: Overview of our proposed HFFM-Intergrated model.

end-to-end automation, they often obscure the precise physical priors explicitly preserved by traditional mathematical operators. This limitation suggests that relying solely on implicit deep features may be insufficient for high-precision microfluidic analysis.

To address these challenges, this study aims to explicitly integrate hand-crafted physical priors into a deep learning architecture, thereby enhancing model robustness and interpretability in complex microfluidic scenarios. The primary contributions of this work are summarized as follows:

1. HFFM-Integrated model: We propose a deep learning-based instance segmentation framework that explicitly integrates hand-crafted features. Through the proposed Hand-crafted Feature Fusion Module (HFFM), the model effectively fuses explicit physical prior knowledge with implicit deep representations. This synergy enhances robustness and interpretability in complex microfluidic scenarios without altering the core backbone of the detection architecture.

2. Attention-based embedding mechanism: We design a specialized attention-based weighting scheme to govern the feature fusion process. This mechanism allows the model to adaptively assess the relevance of hand-crafted inputs and effectively bridges the semantic gap between explicit physical descriptors and implicit deep features.

3. Generalized dual-stage integration strategy: We introduce a hierarchical fusion framework tailored for two-stage detectors. This strategy incorporates the fusion module into the Feature Pyramid Network for multi-scale feature alignment and restructures the Region Proposal Network with a guided anchoring mechanism, strictly enforcing physical constraints during proposal generation.

2. METHODS

The overall architecture of the proposed framework is illustrated in Figure 1. While adhering to the fundamental paradigm of two-stage instance segmentation methods, our approach distinguishes itself from conventional models that rely solely on implicit deep learning for feature extraction. To complement these deep representations, we introduce an auxiliary hand-crafted feature extraction branch explicitly designed to capture texture and edge information. To synergize these dual streams, we propose the hand-crafted feature fusion module, which performs dimensional alignment and adaptive fusion of the heterogeneous features. This mechanism is strategically integrated into the Feature Pyramid Network and the Region Proposal Network to strictly enforce physical constraints. A detailed description of these components follows.

2.1 Overall Architecture

Current methodologies in instance segmentation can be primarily categorized into three distinct paradigms comprising one-stage detectors, query-based transformers, and two-stage frameworks. One-stage detectors typified by SOLO and PolarMask prioritize computational efficiency by formulating detection as a dense regression problem in a single pass. However, this speed often comes at the cost of localization precision particularly for minute targets. Recently, query-based approaches such as DETR have emerged by leveraging self-attention mechanisms to capture global context. Despite their potential, these transformer-based architectures typically lack the inductive biases inherent in convolutional neural networks resulting in high data

dependency and slow convergence which constitute significant drawbacks for microfluidic applications where annotated data is often scarce. Conversely, two-stage frameworks operate on a robust coarse-to-fine philosophy where region proposals are first generated and subsequently refined.

We adopt the two-stage Mask R-CNN [2] as our baseline framework, as it offers a superior balance of stability, data efficiency, and segmentation accuracy. The workflow initiates with a convolutional backbone, such as ResNet, which extracts hierarchical features. These features are further enhanced by a Feature Pyramid Network to accommodate the multi-scale nature of microfluidic droplets and bubbles. Subsequently, the multi-scale features are then processed by the Region Proposal Network to generate a sparse set of candidate regions of interest effectively filtering out background noise. To ensure pixel-level precision, the architecture incorporates the RoIAlign layer which is a critical mechanism that eliminates the quantization errors associated with traditional pooling operations by preserving exact spatial alignment. Finally, the aligned features are fed into parallel multi-task heads to simultaneously predict class labels, regress bounding box coordinates, and generate binary segmentation masks.

2.2 Hand-crafted Features

Although deep learning architectures excel at extracting semantic representations, they inherently suffer from the erosion of high-frequency details due to successive down-sampling and pooling operations [3]. This information loss is particularly detrimental in microfluidics, where the precise delineation of phase boundaries and internal morphology is critical. To mitigate this limitation, we explicitly incorporate hand-crafted feature extraction to capture texture and frequency-domain attributes, serving as immutable physical priors. Specifically, we employ Local Binary Patterns (LBP) and the Discrete Wavelet Transform (DWT) as primary descriptors to complement deep features.

2.2.1 LBP Feature

Local Binary Patterns serve as a robust texture descriptor that characterizes the local spatial structure of the images [4]. The LBP code is derived by evaluating the local difference between the center pixel and its neighbors, formulated as:

$$LBP_{P,R} = \sum_{p=0}^{P-1} s(g_p - g_c) \cdot 2^p \tag{1}$$

Where g_c denote the gray-level value of the center pixel and g_p $(p = 0, \dots, P-1)$ represent the gray values of P equally spaced sampling points on a circle of radius R. And $s(x)$ is the sign function, defined as:

$$s(x) = \begin{cases} 1, if\ x \geq 0 \\ 0, if\ x < 0 \end{cases} \tag{2}$$

In this calculation, the term $s(g_p - g_c)$ effectively captures the local gradient direction. By assigning a binomial weight 2^p to each position, the operator transforms the topological relationship of the neighborhood into a unique integer value.

A concrete instance of this calculation process is illustrated in Figure 2(a). For a standard 3×3 neighborhood, taking the central pixel with an intensity of 6 as the threshold, neighbors are compared against this center. Sequentially collecting these bits produces the binary pattern "11110001", which corresponds to the decimal value 241.In this study, we employ a configuration of $(P = 8, R = 1)$ to capture fine-grained details. The practical efficacy of this approach is demonstrated in Figure 2(b). While the raw microfluidic image exhibits subtle phase boundaries, the LBP-transformed image significantly enhances these high-frequency edge features while suppressing low-frequency background variations, thereby facilitating more precise segmentation.

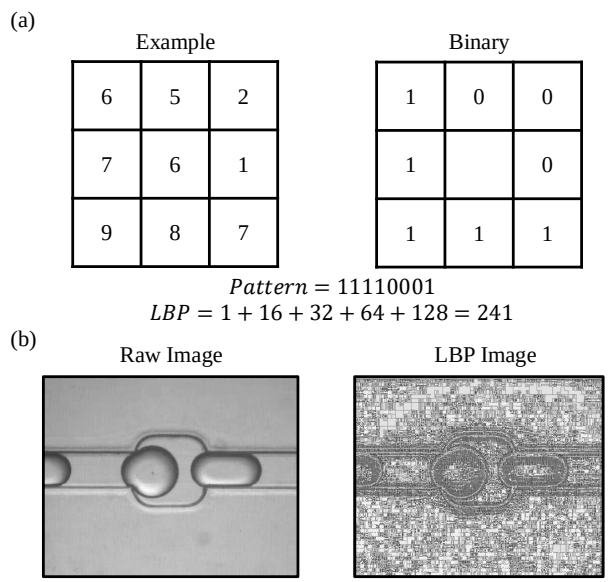

Figure 2: Illustration of the Local Binary Pattern operator. (a) Schematic representation of the LBP calculation process within a 3×3 neighborhood. (b) Visual comparison showing the efficacy of LBP in enhancing high-frequency edge features against the original image.

2.2.2 DWT Features

Discrete Wavelet Transform is utilized to decompose the input image into orthogonal frequency subbands, allowing for the separation of high-frequency edge information from low-frequency background [5]. Mathematically, the 2D-DWT operation on an input image I can be succinctly expressed as:

$$\{LL, LH, HL, HH\} = \mathrm{DWT}(I) \tag{3}$$

In this decomposition, the LL component retains the low-frequency global topology, while the detail components LH, HL, and HH capture high-frequency features corresponding to horizontal, vertical, and diagonal edges respectively. Figure 3 visually underscores this distinct separation where the low-frequency component is dominated by smooth background intensity while the high-frequency sub-bands collectively isolate the sharp intensity gradients characteristic of phase boundaries. Consequently, we explicitly prioritize these high-frequency details to provide the network with a robust structural guide for precise edge delineation, effectively discarding the redundant background information present in the low-frequency approximation.

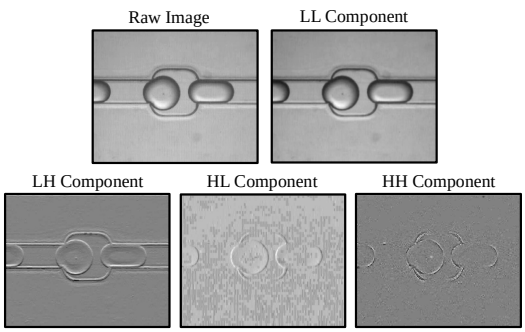

Figure 3: Visualization of the 2D-DWT decomposition.

2.3 Hand-crafted Feature Fusion Module

The core objective of the proposed Hand-crafted Feature Fusion Module is to synergistically integrate the explicit physical priors extracted by traditional operators with the implicit representations learned by the deep neural network. Within the standard convolutional backbone, the generated feature hierarchy inherently exhibits a trade-off between resolution and semantic abstraction where deeper feature maps possess reduced spatial dimensions but enriched semantic information while shallower feature maps retain high-resolution spatial details yet lack high-level contextual understanding.

To address this hierarchical discrepancy, the HFFM first implements a multi-level alignment strategy utilizing simple convolutional layers to independently project the raw hand-crafted features onto the varying spatial resolutions and channel dimensions corresponding to different stages of the feature pyramid. Following this dense alignment, we employ a specialized attention-based fusion mechanism rather than naive element-wise addition to integrate these heterogeneous features. This attention mechanism dynamically generates spatial weight maps to selectively emphasize regions rich in high-frequency textural and structural details effectively filtering out irrelevant background noise. By guiding the network to focus specifically on the informative boundary cues provided by the physical priors, this strategy ensures that the hand-crafted features are robustly incorporated into the deep representation space at every scale of the network.

2.4 Dual-Stage Fusion Strategy

To maximize the efficacy of the proposed physical priors, we deploy a dual-stage integration strategy that embeds the HFFM into both the feature enhancement and region proposal phases as shown in Figure 1. In the first stage, the integration is performed within the Feature Pyramid Network where the hierarchical feature maps generated by the backbone are fed into the HFFM alongside the hand-crafted priors. This early-stage fusion serves to fundamentally enrich the foundational feature representation to ensure that the multi-scale pyramidal features are imbued with explicit structural cues before reaching any downstream functional heads.

In the second stage, we restructure the Region Proposal Network to further exploit these enhanced features by replacing the standard anchor generation mechanism with the Guided Anchoring strategy [6]. This modification is essential because the conventional fixed heuristic is often ill-suited for the irregular and deformable shapes of microfluidic droplets and bubbles. Unlike rigid sliding windows, this mechanism dynamically predicts the probability of object presence and regresses the optimal anchor shapes based on the underlying image semantics. Crucially, we substitute the standard input features of the anchor generation branch with our physically-enhanced fused feature maps to leverage the sharpened edge information and texture details. This integration directs the Guided Anchoring module to generate proposals that are not only spatially adaptive but also strictly aligned with the actual physical boundaries of the dispersed phase thereby improving the quality of the initial candidate regions.

3. RESULTS AND DISCUSSIONS

3.1 Dataset Construction and Statistical Overview

To rigorously validate the efficacy of the proposed method in real-world scenarios, we curate a comprehensive dataset specifically focused on microfluidic bubbles and droplets. This dataset encompasses over 2200 high-resolution images aggregating to more than 64000 individual targets. To ensure consistent evaluation standards, all collected instances are meticulously annotated manually and unified under the single category of "dispersion". For experimental purposes, the dataset is randomly partitioned into training, validation, and testing subsets following a ratio of 60%, 20%, and 20%

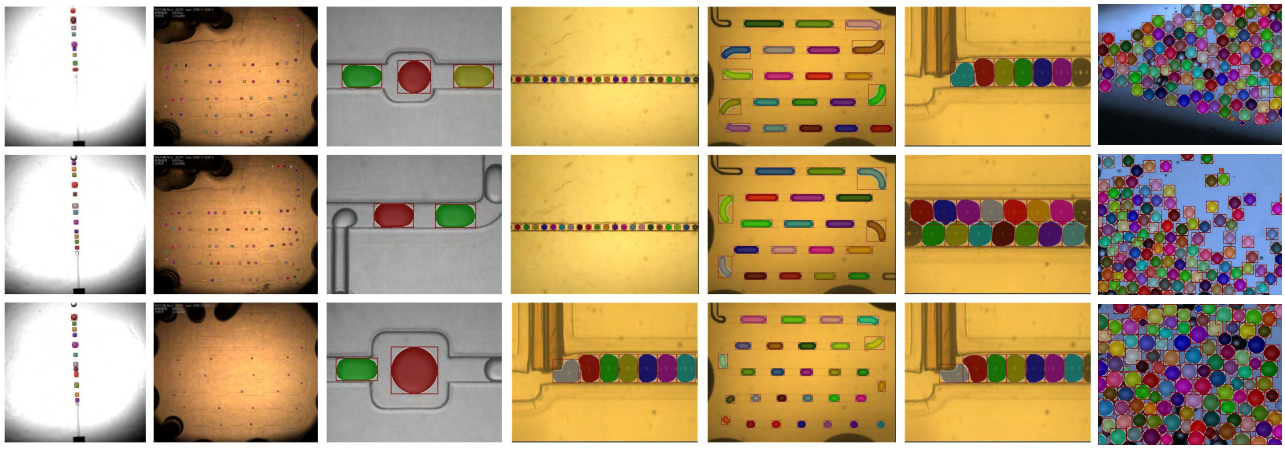

Figure 4: Overview of the constructed microfluidic dataset.

respectively.

As illustrated in Figure 5, we analyze the distribution of critical attributes across these subsets, including instance counts, object area ratios, and aspect ratios. The visualization confirms that there is no significant statistical divergence among the three subsets, indicating a homogeneous data distribution. This consistency ensures that the subsequent evaluation results accurately reflect the model's true generalization capability rather than artifacts of dataset bias.

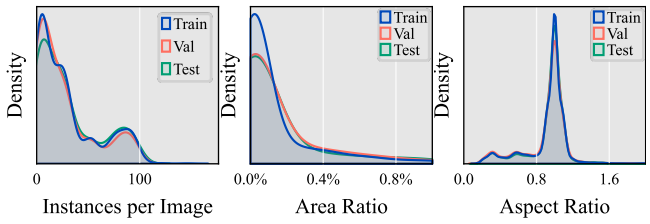

Figure 5: Statistical overview of the constructed microfluidic dataset.

3.2 Training and Evaluation Metrics

To strictly benchmark the performance of the proposed HFFM-Integrated framework, we conduct comparative experiments against a diverse suite of state-of-the-art (SOTA) instance segmentation models. To ensure a unified and fair evaluation, all architectures are implemented within the MMDetection framework and are trained under identical environmental settings. Furthermore, to accelerate model convergence and bolster generalization, the backbone networks are initialized with weights pre-trained on the COCO dataset [7].

To quantitatively assess the performance of the proposed framework, we employ the standard evaluation protocols defined by the Microsoft COCO dataset. The primary metrics utilized are Mean Average Precision (mAP) and Mean Average Recall (mAR), which serve as rigorous indicators of segmentation accuracy. mAP reflects the model's ability to correctly identify and precisely delineate target boundaries without generating false positives, while mAR indicates the model's success in capturing all existing dispersed phases within a frame.

These metrics are fundamentally derived from the Intersection-over-Union (IoU), defined as the ratio of the intersection area to the union area between the predicted mask and the ground truth:

$$IoU = \frac{\text{Area}(\text{Pred} \cap \text{GT})}{\text{Area}\ (\text{Pred} \cup \text{GT})} \tag{4}$$

Based on the IoU thresholds, the predictions can be classified into True Positives (TP), False Positives (FP), and False Negatives (FN). The Average Precision (AP_{IoU}) and Average Recall (AR_{IoU}) are subsequently calculated as:

$$AP_{IoU} = \frac{TP}{TP + FP} \tag{5}$$

$$AR_{IoU} = \frac{TP}{TP + FN} \tag{6}$$

And the mAP and mAR are calculated by averaging the AP over 10 distinct IoU thresholds ranging from 0.50 to 0.95 with a step size of 0.05:

$$mAP = \frac{1}{10}\sum_{i=0}^{9} AP_{IoU=0.50+0.05\times i} \tag{7}$$

$$mAR = \frac{1}{10}\sum_{i=0}^{9} AR_{IoU=0.50+0.05\times i} \tag{8}$$

Furthermore, given the significant scale variations inherent in microfluidic dispersed phases, we conduct a stratified analysis based on object size. We report AP_S, AP_M, and AP_L for objects with pixel areas less than 32^2, between 32^2 and 96^2, and greater than 96^2, respectively. This granular breakdown allows for a specific assessment of the model's robustness in detecting minute

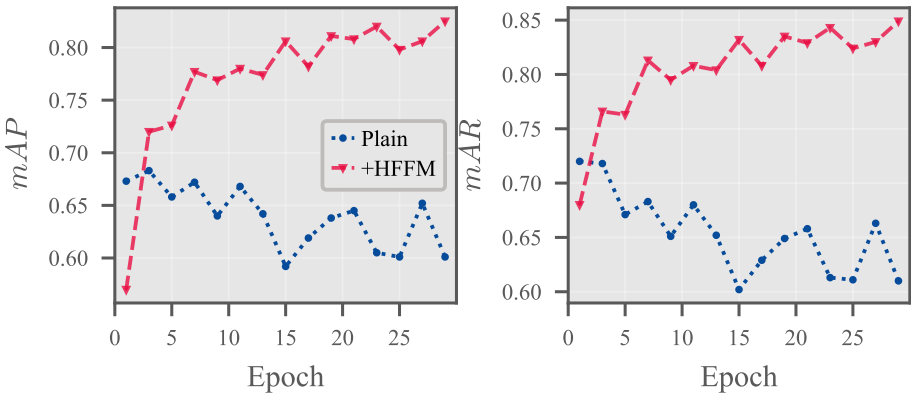

Figure 6: Comparative evaluation of detection accuracy.

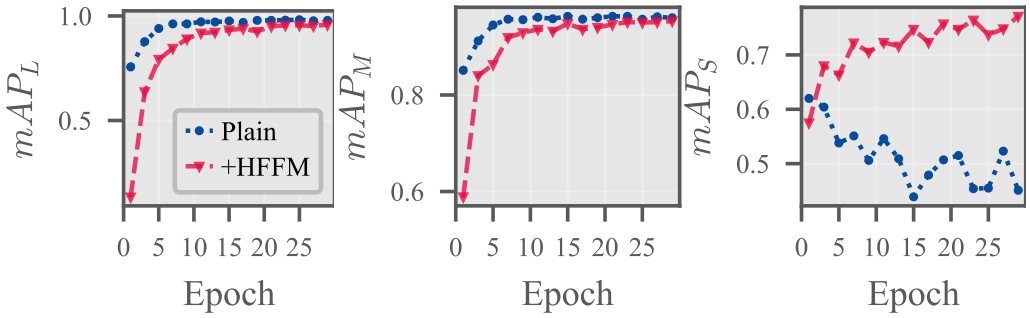

Figure 7: Comparative evaluation of detection accuracy stratified by object size.

targets.

3.3 Training Curves

The training progression of our proposed framework versus the baseline is visualized in Figure 6. As evidenced by the curves, the baseline model (denoted as 'Plain') exhibits a tendency to plateau early, consistently fluctuating within the 0.60–0.70 range for both mAP and mAR without showing further convergence. In contrast, the HFFM-Integrated model catalyzes a significant performance leap. Despite a brief initial adjustment period, the proposed model establishes a robust and monotonic upward trajectory, rapidly surpassing the baseline around the 3rd epoch and steadily converging to high-performance values exceeding 0.80. This substantial margin validates that the explicit injection of physical priors effectively prevents the network from stagnating in local optima, resulting in superior global precision and recall. To rigorously deconstruct the source of these performance gains and strictly verify whether the hand-crafted features successfully enhance the detection different scales of targets as hypothesized, it is imperative to conduct a stratified evaluation based on object size metrics.

To investigate the underlying drivers of the performance improvement, we further deconstruct the evaluation metrics based on object scale as illustrated in Figure 7. The comparative curves for large (mAP_L) and medium (mAP_M) targets reveal that both the baseline and the proposed framework achieve high saturation levels, with detection accuracy exceeding 0.95. In these categories, the baseline model already exhibits robust performance, indicating that deep semantic features are sufficiently distinct for identifying prominent dispersed phases. However, a striking disparity emerges in the small-object regime (mAP_S). As clearly depicted in the rightmost subplot, the baseline model struggles significantly to resolve minute targets, suffering from severe instability and failing to surpass 0.60 after the initial training phase. This degradation corroborates the limitation of standard CNN architectures, where successive downsampling erodes the spatial cues necessary for detecting small-scale features. In sharp contrast, the HFFM-Integrated model maintains a vigorous upward trajectory for mAP_S, steadily converging to approximately 0.78. This substantial margin, representing an improvement of nearly 20% over the baseline, provides empirical evidence that the fused hand-crafted physical priors successfully preserve high-frequency edge details. Consequently, the proposed method effectively recovers the detection capability for micro-scale targets that are otherwise lost in the deep feature hierarchy.

Table 1: Comparison with SOTA models and ablation study on test set.

Methods		Val				Test			
		mAP	mAR	mAP_S	mAR_S	mAP	mAR	mAP_S	mAR_S
Mask R-CNN Based	Plain	0.741	0.788	0.693	0.744	0.613	0.649	0.55	0.585
	+HFFM FPN	0.811	0.837	0.752	0.787	0.751	0.773	0.686	0.717
	+HFFM RPN	0.819	0.845	0.769	0.799	0.798	0.824	0.750	0.782
	+HFFM (Both)	**0.825**	**0.849**	**0.771**	**0.802**	**0.808**	**0.834**	**0.764**	**0.795**
SOTA Models	Mask2Former	0.574	0.626	0.451	0.523	0.458	0.515	0.329	0.405
	CondInst	0.617	0.670	0.510	0.569	0.390	0.417	0.244	0.268

3.4 Comparison with SOTA Models and Ablation Study

To validate the superiority of the proposed framework within the specific domain of microfluidics, we conduct a comparative analysis against representative state-of-the-art architectures, including the transformer-based Mask2Former and the CNN-based CondInst. As detailed in Table 1, the quantitative results reveal a substantial performance gap favoring our approach. While Mask2Former utilizes powerful self-attention mechanisms to capture global context, it achieves a test mAP of only 0.458 and a remarkably low mAP_S of 0.329. In the context of microfluidics where training data is relatively scarce and targets are minute, the transformer struggles to converge on precise boundaries. Similarly, CondInst yields the lowest performance with a test mAP of 0.390. In contrast, our proposed method achieves a test mAP of 0.808 and a test mAP_S of 0.764. This significant margin demonstrates that explicitly embedding hand-crafted physical priors into a stable two-stage baseline is a far more effective strategy for scientific image analysis than merely applying complex generic deep learning models.

To rigorously decouple the individual contributions of the Hand-crafted Feature Fusion Module at different stages of the detection pipeline, we perform an ablation study as summarized in the upper section of Table 1. The baseline Mask R-CNN, denoted as Plain, establishes a starting performance of 0.613 for test mAP. The first experimental configuration introduces the HFFM solely into the Feature Pyramid Network. This integration results in a significant performance boost, raising the test mAP to 0.751. This improvement verifies that injecting texture and edge cues at the feature extraction level fundamentally enriches the representational power of the network. Subsequently, we evaluate the impact of integrating the HFFM exclusively into the Region Proposal Network to guide anchor generation. This configuration yields a test mAP of 0.798, indicating that physically-guided proposals are critical for reducing false negatives and accurately localizing small targets. Finally, the complete architecture which synergistically combines HFFM in both the FPN and RPN achieves the highest performance metrics across all categories with a test mAP of 0.808 and a test mAP_S of 0.764. These results confirm that the dual-stage fusion strategy offers a complementary gain where the FPN enhances feature discriminability while the RPN ensures precise spatial localization.

3.5 Qualitive Analysis

To provide an intuitive understanding of how the hand-crafted features enhances the network's representational capability, we visualize the qualitative detection results and the corresponding intermediate feature maps in Figure 8. The bottom row compares the feature maps at FPN Level 0, which is the critical stage for retaining high-resolution spatial details. As observed in the bottom-right subfigure, the feature map of the plain model appears diffuse and lacks structural definition, indicating that the subtle boundaries of the targets are eroded during the feature extraction process. In sharp contrast, the HFFM-Integrated model (bottom-left) exhibits distinct, high-contrast activations that precisely delineate the serpentine channel structure and the individual droplets. This visual evidence confirms that the attention-based fusion successfully injects the hand-crafted physical priors into the deep feature hierarchy. Consequently, this enhanced feature discriminability directly translates to the superior detection performance shown in the top row. While the Plain model misses minute targets or generates ambiguous bounding boxes in low-contrast regions, the HFFM-Integrated model demonstrates robust localization capabilities, accurately capturing the targets with high fidelity.

3.6 Computational Efficiency

To evaluate the impact of the introduced modules on inference speed, we conduct a comprehensive benchmark comparing the baseline Mask R-CNN with our HFFM-Integrated model. For a fair comparison, all input images are resized to 512×512 pixels, and testing is performed on an NVIDIA RTX 5880 Ada GPU. The stage-wise results are summarized in Table 2.

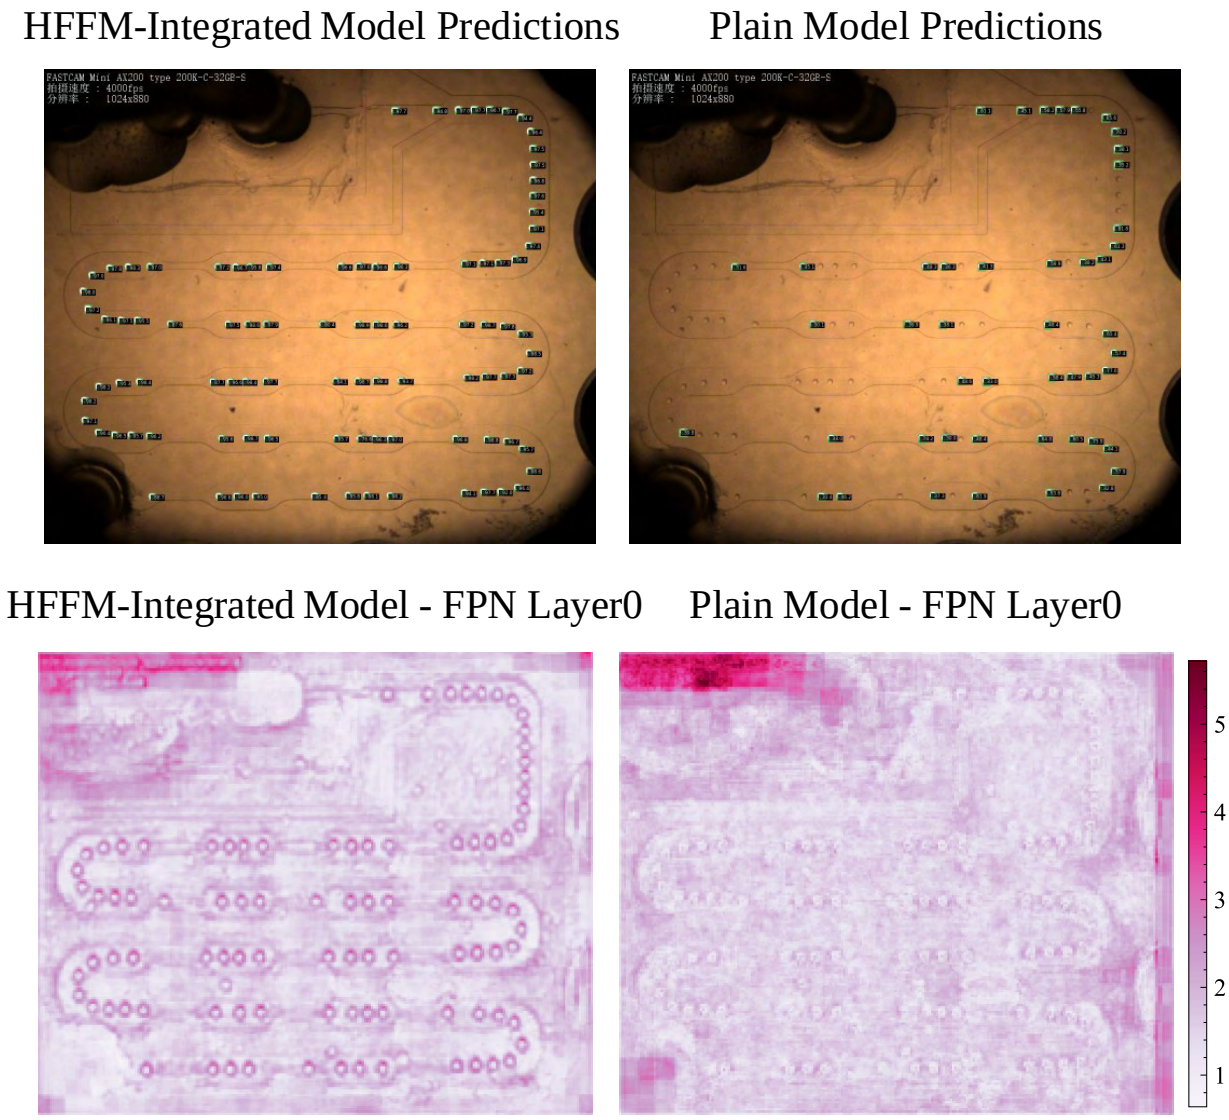

Figure 8: Visualization of detection results and intermediate outputs.

Table 2: Comparison of inference latency between models (ms).

Stage (ms)	Plain	+HFFM
Image Read	3.46	3.58
Preprocess	4.41	8.53
Backbone	4.92	5.09
FPN	0.73	2.81
RPN	0.80	8.46
Prediction Heads	7.00	9.95
Total	21.32	38.41

The results indicate that the integration of handcrafted features and the Guided Anchoring mechanism introduces a computational overhead of approximately 17 ms per image. This latency represents a strategic trade-off. While the plain model is slightly faster, our proposed framework ensures a more stable training process and significantly faster convergence.

4. CONCLUSIONS

In this study, we propose a novel HFFM-Integrated instance segmentation framework that synergistically integrates explicit physical priors with implicit deep representations. By embedding Local Binary Patterns and Discrete Wavelet Transform into the deep learning pipeline, our approach effectively compensates for the erosion of texture and boundary details caused by down-sampling operations. Comprehensive evaluations on our constructed dataset comprising over 64000 instances

demonstrate the framework's superiority. Our method achieves a test mAP of 0.808, significantly outperforming state-of-the-art architectures like Mask2Former and CondInst. Most notably, the proposed strategy effectively resolves the detection bottleneck for small-scale targets, elevating the mAP_S to 0.764. These findings confirm that injecting domain-specific physical constraints is a highly effective strategy for scientific image analysis, superior to relying solely on generic deep models. Ultimately, this study provides a robust, high-precision tool for automated microfluidic monitoring, paving the way for more accurate mechanistic studies in microreaction.

While the current performance is substantial, there is still progress needed to reach an ideal accuracy of 0.99. This gap primarily stems from the inherent complexity of the microfluidic dataset, which encompasses targets with vast disparities in size, geometric morphology, and experimental conditions. Achieving absolute saturation in precision remains a persistent challenge even in mainstream computer vision benchmarks like COCO. Consequently, future research focuses on developing even more robust, multi-scale deep learning architectures to capture these diverse dispersed phases more effectively across various flow regimes. Furthermore, although the integration of physical priors introduces a manageable computational overhead of approximately 17 ms per image, optimizing inference efficiency remains a priority for ultra-high-speed monitoring. Future efforts involve transitioning the current CPU-bound preprocessing operations, specifically the LBP and DWT transformations to high-performance GPU CUDA kernels to maximize real-time processing capability for online chemical process analysis.

AUTHOR IDENTIFIERS

Author ORCIDs:
Wenle Xu: 0009-0001-9228-0956

REFERENCES

1. Jensen KF. Flow chemistry—microreaction technology comes of age. AIChE Journal 63:858-869 (2017). https://doi.org/10.1002/aic.15642
2. He K, Gkioxari G, Dollar P, Girshick R. Mask R-CNN. 2017 IEEE International Conference on Computer Vision (ICCV) :2980-2988 (2017). https://doi.org/10.1109/iccv.2017.322
3. K. Xu, M. Qin, F. Sun, Y. Wang, Y.-K. Chen, F. Ren, Learning in the Frequency Domain, (2020). https://doi.org/10.48550/arXiv.2002.12416
4. Ojala T, Pietikainen M, Maenpaa T. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. IEEE Trans. Pattern Anal. Machine Intell. 24:971-987 (2002). https://doi.org/10.1109/tpami.2002.1017623
5. Mallat SG. A theory for multiresolution signal decomposition: the wavelet representation. IEEE Trans. Pattern Anal. Machine Intell. 11:674-693 (1989). https://doi.org/10.1109/34.192463
6. Wang J, Chen K, Yang S, Loy CC, Lin D. Region proposal by guided anchoring. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) :2960-2969 (2019). https://doi.org/10.1109/cvpr.2019.00308
7. Zhang S, Li H, Wang K, Qiu T. Accelerating intelligent microfluidic image processing with transfer deep learning: a microchannel droplet/bubble breakup case study. Separation and Purification Technology 315:123703 (2023). https://doi.org/10.1016/j.seppur.2023.123703

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Methodology to assess the integrity of Water and Energy Integration Systems (WEIS) models using the ThermWatt computational tool

Miguel Castro Oliveira[ab]*, Rita Castro Oliveira[c] and Henrique A. Matos[b]

[a] Research, Development and Innovation, Instituto de Soldadura e Qualidade, 2740-120 Porto Salvo, Portugal
[b] Centro de Recursos Naturais e Ambiente (CERENA), Department of Chemical Engineering, Instituto Superior Técnico, Avenida Rovisco Pais 1, 1049-001 Lisboa, Portugal
[c] Department of Computer Science and Engineering, Instituto Superior Técnico, Avenida Rovisco Pais 1, 1049-001 Lisboa, Portugal
* Corresponding Author: dmoliveira@isq.pt, miguelmcoliveira@tecnico.ulisboa.pt.

ABSTRACT

Type your abstract text here. This work presents an essential methodological framework oriented to the implementation of sustainability promotion measures in process industries. It makes use of a previously developed paradigm, designated as Water and Energy Integration Systems (WEIS), which are fundamentally conceptual systems based on the implementation of several technologies implemented with the end to minimize water use, energy use and related environmental burdens. The primarily conceptual nature of these systems is significant that these have not been significantly implemented in real-life, and that these have been essentially implemented in the virtual basis of digital twin-based computational models. This work extensively presents a methodology developed for the assessment of the integrity of WEIS models, which have been developed using the capacities of a customised computational tool designated as ThermWatt. Two previously approached case-studies have been considered to perform a proof-of-concept in respect to the applicability of the delineated methodology. The delineated methodology proved to be adequate for the whole process of the elaboration of the WEIS Engineering projects, paving way to the implementation of these in real-life in a future perspective.

Keywords: Water and energy integration systems, simulation, optimisation, model integrity, sustainability promotion

INTRODUCTION

This is Header 2Water and Energy Integration Systems (WEIS) subsist in complex installations to be installed in end-use sector facilities set to perform the maximum use of the recirculation of waste heat and water streams for the purpose of causing an overall sustainability promotion. These systems use a set of technologies that may be used in the context of the exploitation of the interdependencies of energy and water resources, as encompassed by the concept of water-energy nexus [1].

The WEIS paradigm was introduced and extensively developed in previous works developed by the authors [2], along with the tool created for the purpose of modelling, the ThermWatt computational tool [3]. This tool subsists in several simulation and optimisation models created in the Modelica and Python languages, which fundamentally serve as digital twins of real-life installations.

The duality between the WEIS paradigm and the ThermWatt tool has been making possible the existence of a complex framework in which the viability associated to sustainability promotion technologies-based systems is set to be firstly assessed in a virtual basis. In this sense, it is possible to develop proposals of systems that promote overall sustainability in a site with minimized requirements of real-life testing of technologies to be implemented. Owing to the primarily conceptual nature of the WEIS paradigm, the assurance of the integrity of computational models of this type of systems is essential

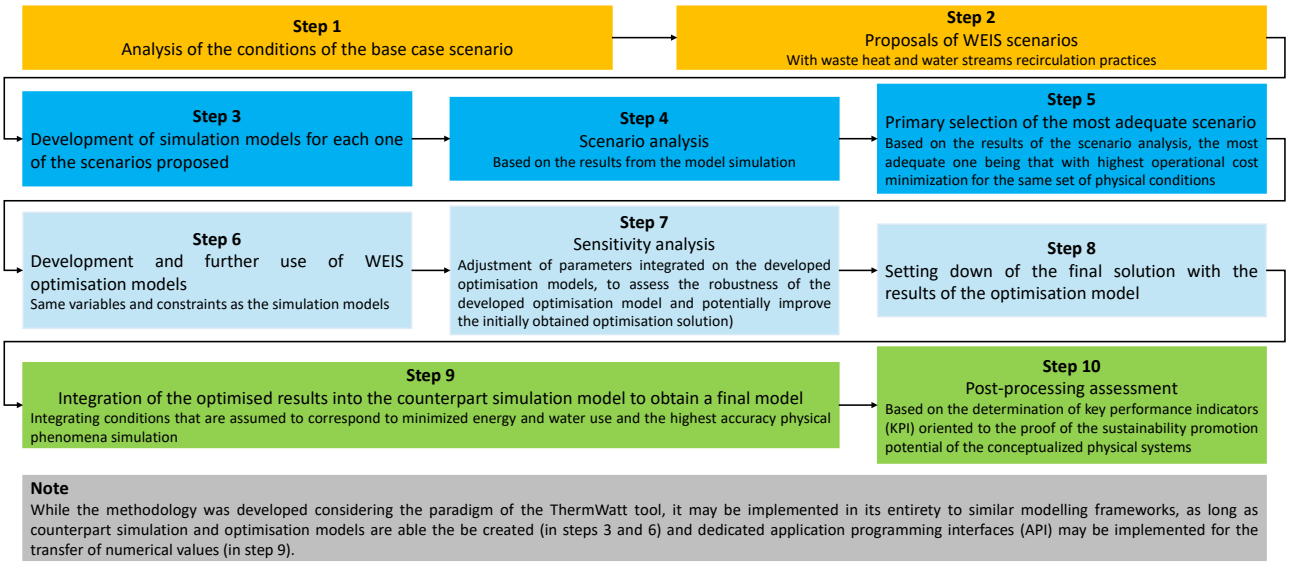

Figure 1. Methodology for model integrity assessment

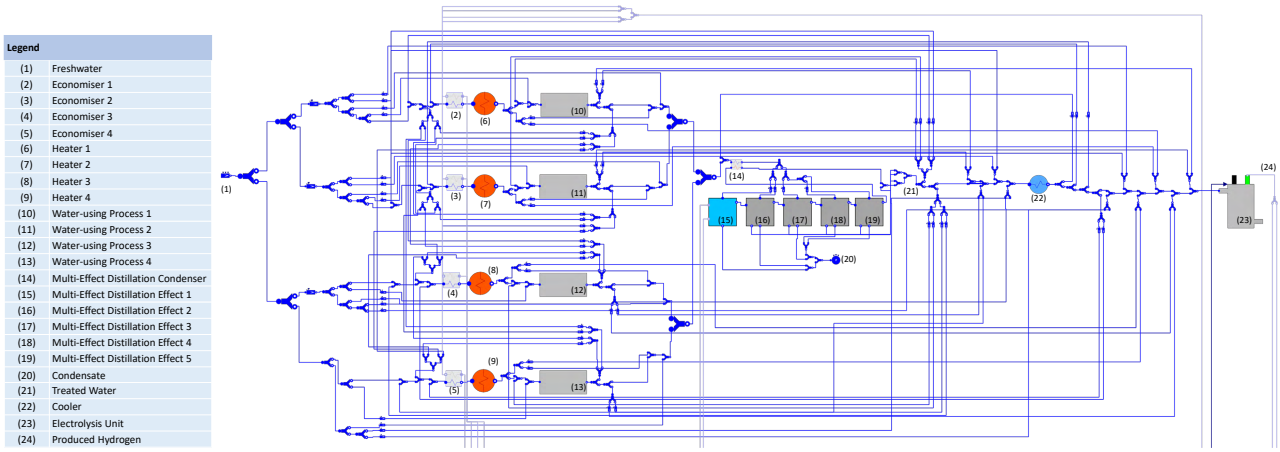

Figure 2. Water system side of both WEIS Scenarios 1 and 2 for Case-study 1 (flowsheet-based model created using the capabilities of the ThermWatt tool)

to proceed with the most accurate analysis of the physical phenomena inherent to these systems and the assessment of the sustainability promotion potential of these.

In this work, the methodology developed for the purpose of assessing the integrity of WEIS models is extensively presented. A set of results related to two case-studies within the Portuguese process industry (namely a ceramic plant) are presented and furtherly used to prove the applicability of the developed methodology.

METHODOLOGY

The assessment of the potential implementation of Water and Energy Integration Systems (WEIS) must subsist on a methodology that is set to allow the whole process from the initial analysis of a case-study (which is commonly a process industry plant and its whole system comprising energy sources, energy supply, energy demand and energy-using processes) to the obtainment of indicators reflecting the sustainability promotion potential brought by the implementation of such systems. In the context of the integration of computational model (which has been the main form of study of the implementation of WEIS, owing to the mostly still conceptual nature of these systems), the implementation of such methodology requires the development of validated and useful models. As such, the delineation of such methodology requires that this is not only one oriented to the assessment of the potential implementation of WEIS Engineering projects but also to the evaluation of the integrity of computational models of these systems.

In Figure 1, the methodology inherent to the process of WEIS model integrity evaluation (and, in this prospect,

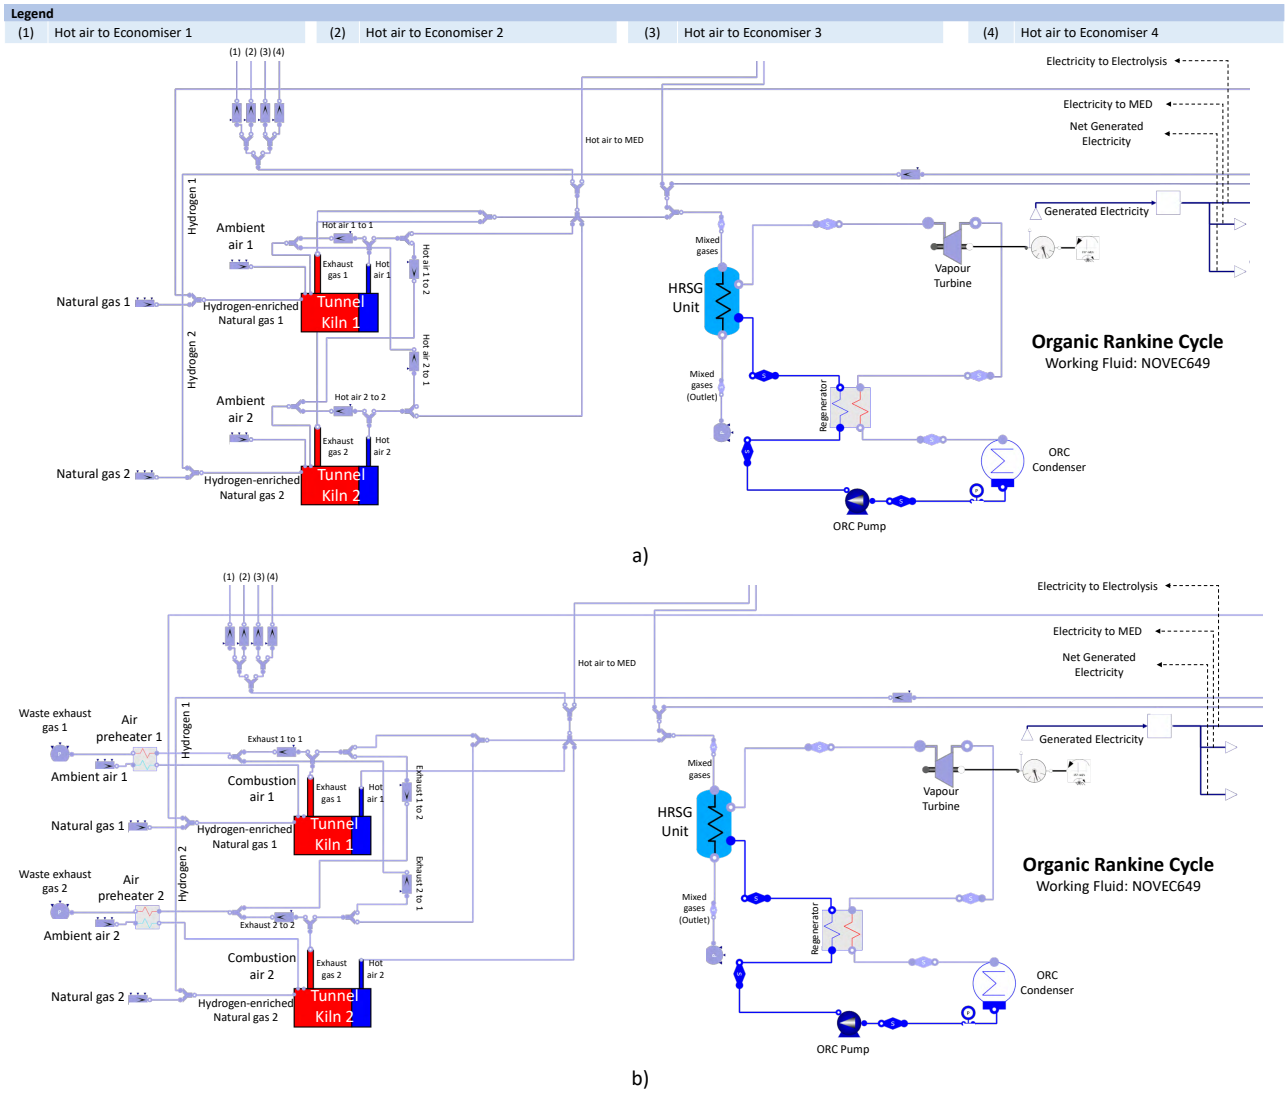

Figure 3: WEIS (Thermal system side) a) Scenario 1 and b) Scenario 2 for Case-study 1 (flowsheet-based model created using the capabilities of the ThermWatt tool)

the assessment of the potential implementation of WEIS within specific case-studies) is extensively detailed.

The development of all the previous works within the overall topic of Water and Energy Integration Systems [2, 3] have subsisted on the application of the afore identified methodology in a subtle manner. The delineation of this methodology which serves as the main subject of this work serves as a manner to turn accessible the implementation of the WEIS paradigm by interested parts of scientific communities, as well as interested parts from end-use sectors representatives.

In the prospect to assess model integrity, the methodology delineated in Figure 1 must be implemented in an aggregated manner, in which the application of each one of the delineated steps converges in the proof of integrity of a developed modelling framework. The quality of a model to be integrate subsists on the verification of the following aspects:

- Simulation-based scenarios must be able to generate most favourable results in comparison to the base case scenarios;
- Optimisation results must be overall more favourable in comparison to the results generated by running of the corresponding simulation models;
- Models are valid, according to the comparison between real-life and simulation data;
- The sustainability promotion capacity (economic, environmental and social pillars) associated to WEIS implementation must not be compromised by the application of the overall modelling framework.

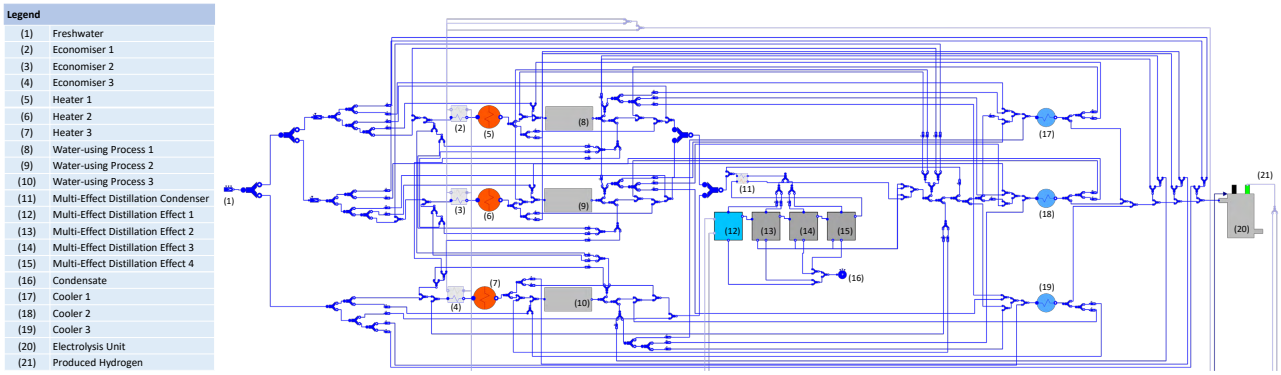

Figure 4. Water system side of both WEIS Scenarios 1 and 2 for Case-study 2 (flowsheet-based model created using the capabilities of the ThermWatt tool)

BACKGROUND

In the prospect to perform a proof of concept of the WEIS paradigm, a set of two process industry-oriented case-studies have been approached. Both these case-studies are set in a Portuguese ceramic plant, has been exploited through the development and further use of simulation and optimisation models (integrated in the ThermWatt computational tool), followed by a complex post-processing assessment based on the evaluation of the viability associated to proposed WEIS.

The afore proposed methodology has been implemented for both case-studies (up until step 3). For both case-studies, the conceptualized WEIS consider the implementation of:

- An Organic Rankine cycle (ORC) to generate electricity;
- A multi-effect distillation (MED) unit for removal of contaminants (salts) from wastewater streams (and thus to produce treated water streams with negligible concentration of contaminants);
- An Electrolysis unit to produce hydrogen streams from treated water streams, which are then blend with the natural gas fuel streams at the inlet of the combustion-based processes.

For case-study 1, the conceptualized WEIS specifically considers:

- A steady-state perspective (owing to the inexistence of combustion-based processes operating in batch);
- Two tunnel kilns as combustion-based processes;
- Four water-using lines (each one encompassing a heater and a water-using process), which are mixed in single water-using line which is then cooled down in a cooler;
- The recirculation of the waste heat streams (either exhaust gas or hot air) to constitute additional energy inputs in combustion-based processes with direct mixture with inlet combustion air (Scenario 1) or by the implementation of air preheaters (Scenario 2).
- For case-study 2, the conceptualized WEIS specifically considers:
- A transient-state/ dynamic perspective (owing to the existence of batch combustion-based processes);
- Four combustion-based processes, two tunnel kilns and two intermittent kilns;
- Three water-using lines (each one encompassing a heater, and a water-using process and a cooler);
- The recirculation of the waste heat streams (either exhaust gas or hot air) to constitute additional energy inputs in combustion-based processes with direct mixture with inlet combustion air;
- The aforementioned streams are also recirculated to a phase change material (PCM)-based thermal energy storage (TES) unit, which is heated up by the passage of the exhaust gases (Scenario 1) or hot air (Scenario 2) streams (while the intermittent kilns are not operating) and from which thermal energy is discharged by the passage of an ambient air stream which is then distributed to each one of the four combustion-based processes.

Each one of the WEIS scenarios developed for case-studies 1 are pictorially presented in the sequence of Figures 2 and 3. Each one of the WEIS scenarios developed for case-studies 1 are pictorially presented in the sequence of Figures 4 and 5.

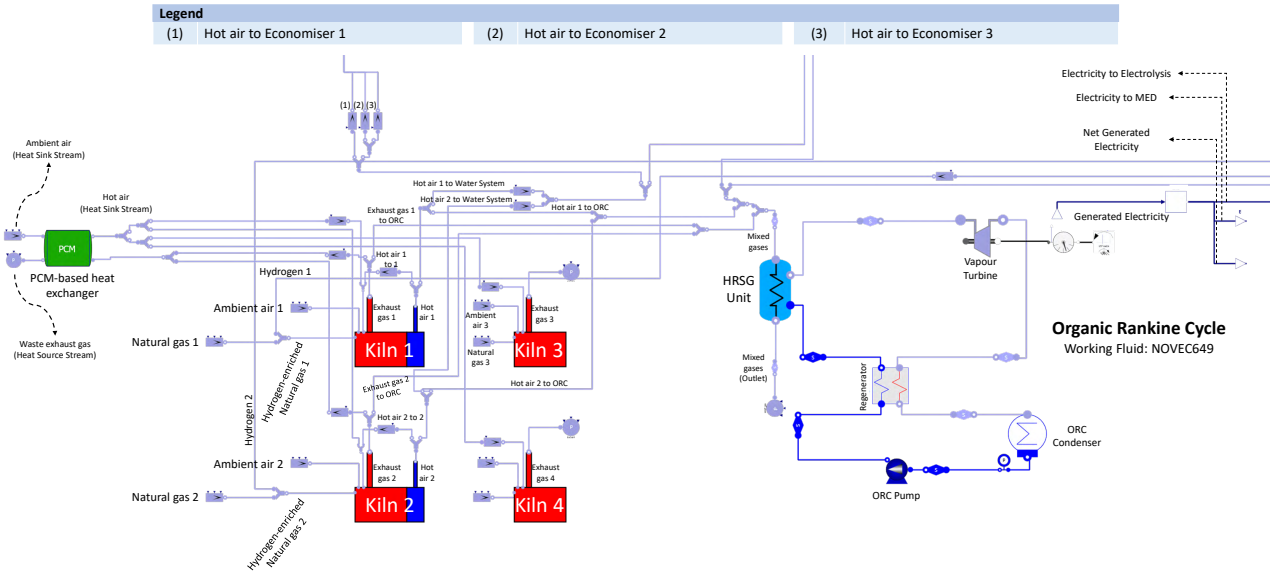

a)

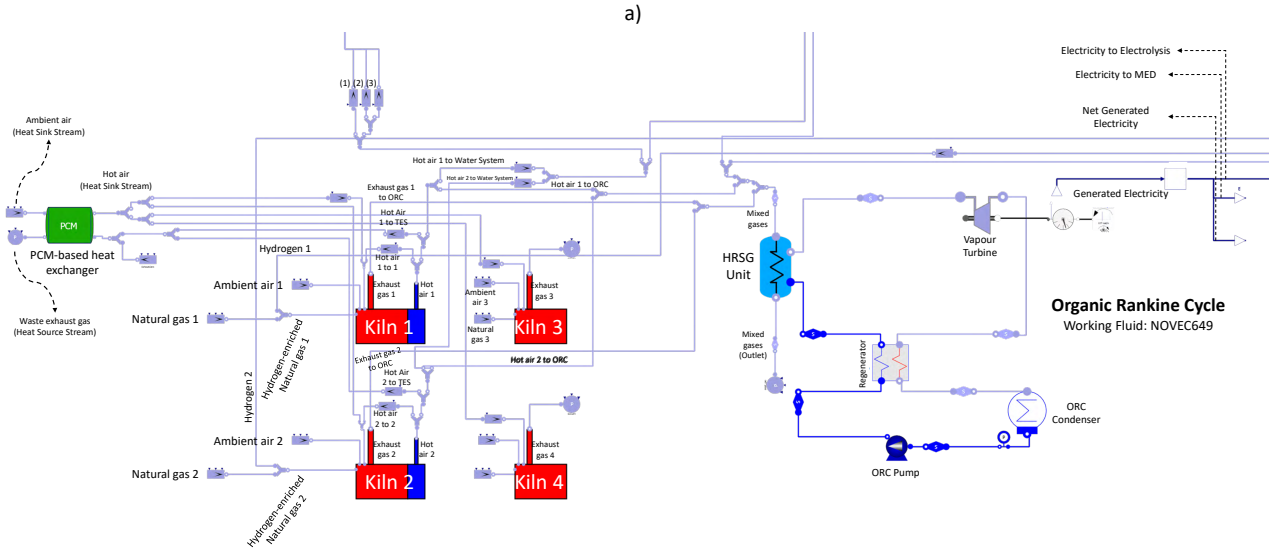

b)

Figure 5. WEIS (Thermal system side) a) Scenario 1 and b) Scenario 2 for Case-study 2 (flowsheet-based model created using the capabilities of the ThermWatt tool)

MODEL INTEGRITY ASSESSMENT

The assessment of model integrity associated to each one of the approached case-studies may be performed by the analysis of aggregated results. In this case, the results which may be analysed in aggregated manner and thus allow the attainment of the specific objectives of the implementation of the afore presented methodology are the energy use or production associated to each one of the energy sources of interest (natural gas use in combustion-based processes, net electricity generation, freshwater consumption and hot/ cold utility consumption), as well as the total operational costs attributed to the overall inputs of energy and water. In the context of the integration of developed computational models, the following sets of results are set to be considered for analysis:

- Base cases;
- Simulation results for each WEIS scenario;
- Initial Optimisation results obtained by the running of the developed optimisation models before sensitivity analysis;
- Optimisation results obtained by the running of the

Table 3: Model integrity assessment for each one of the approached case-studies

Case-study 1	Case-study 2
• Scenario 1 is favoured in relation to scenario 2, owing to the relatively lower operational costs; • It is possible to verify that the obtained values for the hot utility consumption in heaters 1, 2 and 3 are higher in the context of the final model, being slightly higher than the null values obtained by the previous two sets of optimisation results. Such is attributed to the different temperature of the water streams at the inlet of the mentioned three heaters on the final model, which results from the mixture of several recirculated water streams within the whole water system within the overall WEIS; • These temperatures values are slightly lower on the final model (developed in the Modelica language) in relation to the optimisation models (developed with the Python languages), owing to the different modelling packages used to simulate water properties in both categories of models. Such slight decrease of temperature at the inlet of the heaters are responsible for a slight requirement of hot utility supply to achieve the objective temperature values at the outlet of the heaters.	• Scenario 2 is favoured in relation to scenario 1, owing to the relatively lower operational costs (as may be verified by the operational costs associated to the thermal process systems standalone); • The results associated to each one of the parcels associated to energy use (namely the natural gas consumption in each one of the four kilns) are generally lower in the respective optimisation-based scenarios in comparison to the counterpart simulation-based scenarios (which is according to the expected verifications); • The respective energy use levels are relatively higher in the final model (which also causes the total operation costs to be relatively higher). Such relative increase may be attributed to the respective methods implemented for the modelling of the PCM-TES unit in the optimisation model developed for Case-study 2 and in the equipment-level model implemented within the case-study's system-level model. The latter uses a much more accurate method, based on the finite difference method (FDM) for the modelling of heat transfer phenomena; • The afore delineated verification does not considerably affect overall model integrity. The disparity between the results from the counterpart optimisation and final models may be considered a negligible difference, not affecting the balance between modelling accuracy and attainment of the highest sustainability promotion level for the approached case-study.
Overall	
• The implementation of all scenarios for both case-studies is set to generate significant results in terms of reduced energy and water-related operational costs in relation the base case scenario; • All the sets of optimised scenarios generate more favourable results in comparison the ones obtained by simulation models (proving the primary effectiveness of the developed optimisation models); • The results obtained after the sensitivity analysis correspond to a point in which the thermal-to-electric efficiency considered to assess potential electricity generation on the optimisation models is reduced in order to ensure that this parameter is certainly higher in the context of the running of the final simulation model; • The aforementioned verification ensures that the definitive solution is the one associated to the minimum possible level of total operational costs, while still considering the most accurate modelling of physical phenomena. In this sense, the net generated electricity (and thus operational costs savings) is relatively lower after the performance of the sensitivity analysis; • The final results (optimisation results integrated in the simulation model) correspond to significant operational costs savings, evaluated using a model proved to be accurate in terms of the simulation of physical phenomena.	

Consumption (MJ/h)	**Cooler 2**	133.00					
	Cooler 3	17.00					
Overall							
Total Operational Cost (M€/year)		2.53	2.04	2.04	2.01	1.87	1.89

developed optimisation models after the performance of sensitivity analysis;

- Final results obtained by the aggregated model.

In Tables 1 and 2, the aggregated results for each one of the WEIS scenarios of interest are presented. In Table 3, the aspects related to the model integrity assessment process for each one of the approached case-studies are presented.

Overall, model integrity associated to the whole modelling framework developed for each one of the case-studies may be considered to have been secured, attending that:

- All sets of simulation results are considered to be more favourable in comparison to base case scenarios (significant reduction of total operational costs);
- All sets of optimisation results are considered to be more favourable in comparison to the

counterpart simulation scenarios (considerable relative reduction of total operational costs);

- Models are valid;
- The overall modelling framework allows the attainment of final WEIS configurations that have a secured sustainability promotion capacity.

In the framework of the present work, the implementation of methodology presented in Figure 1 may be considered to have been fully fulfilled by section 3 (up until step 3) and the present section 4 (up until step 9). The sustainability assessment process (corresponding to step 10) is considered to also have been completed in the scope of the overall study inherent to this work, although such process have been mostly presented and extensively exploited in the previous related works developed by the authors [4]. In this prospect, the presentation of a significantly higher detail related to this step is considered to be unnecessary in a standalone perspective of the focus of this work. An extensive model validation procedure have been also presented in previous works (the models approached in this work are considered to be a priori valid in terms of the representation of physical phenomena) [5].

An extensive delineation of the sustainability promotion capacity associated to the implementation of the conceptualized WEIS (in respect to the proof of the attainment of benefits associated to the economic, environmental and social pillars of sustainability) is presented in the digital supplementary material (with information adapted from [5]).

CONCLUSIONS

This work extensively presents the methodology developed for the purpose of assessing the integrity of Water and Energy Integration Systems (WEIS) computational models. Although the WEIS paradigm has been introduced and extensively exploited in previous works, the delineation of the approached methodology is fundamental for the scientific communities and end-use sectors representative to replicate similar sustainability promotion results for case-studies other than the ones that have been already exploited. The existence of such a methodology is important due to the primarily conceptual nature of the WEIS paradigm. As such, the assurance of the integrity of models of this type of systems is essential to proceed with the most accurate analysis of the physical phenomena inherent to these systems and the determination of several sustainability promotion indicators (which is the foremost final objective of these of these models).

From the analysis of aggregated results obtained by the use of computational model created for two Portuguese process industry case-studies, it is possible to affirm that:

- The running of initial simulation models for several WEIS scenarios for a same case-study allow the attainment of significant savings at the level of energy use, water use and total operational costs;
- The running of counterpart optimisation models allows the attainment of significantly more favourable results in comparison to the simulation models, with such attainment process being later refined by the consideration of results obtained after the performance of sensitivity analyses to these models;
- The integrity of the final models (through the integration of optimisation results into the counterpart simulation models, considering a compromise between the accuracy of physical phenomena modelling and optimised results) is not compromised by the allocation of numerical outputs from the optimisation models as inputs on the simulation models.

The proposed methodology may be considered adequate for the whole process of the elaboration of the Engineering projects of a WEIS, from the first step of analysing base case data to the final proof of sustainability promotion. Such is supported by the successful implementation of the delineated methodology in two real-life case-studies, having been proved that the integrity of the developed overall modelling frameworks is secured for both cases.

DIGITAL SUPPLEMENTARY MATERIAL

Digital supplementary material is available at https://psecommunity.org/LAPSE:2026.0021.

ACKNOWLEDGEMENTS

This work was funded through the base funding component of the Center for Technology and Innovation—ISQ, under the terms defined in AAC nº 03/C05i02/2022, and CERENA, UIDB/04028/2020_UIDP/04028/2020.

AUTHOR IDENTIFIERS

Author ORCIDs:
Castro Oliveira M: 0000-0002-3649-8838
Castro Oliveira R: 0009-0003-3114-7857
Matos HA: 0000-0001-8128-7346

REFERENCES

1. Walsh BP, Murray SN, O'Sullivan DTJ. The water

energy nexus, an ISO50001 water case study and the need for a water value system. Water Resources and Industry 10:15-28 (2015). https://doi.org/10.1016/j.wri.2015.02.001

2. Oliveira MC, Oliveira RC, Castro PM, Matos HA. The paradigm of water and energy integration systems (WEIS): methodology and performance indicators. Systems and Control Transactions 4:1288-1293 (2025). https://doi.org/10.69997/sct.130880
3. Castro Oliveira M., Castro Oliveira R., Matos H.A. Dynamic simulation and optimisation of water and energy consumption in a ceramic plant: Application of the customised ThermWatt computational tool. *Comput. Aided Chem. Eng.* 53: 2377–2382 (2024) https://doi.org/10.1016/B978-0-443-28824-1.50397-5
4. Oliveira MC, Matos HA. Sustainability and strategic assessment of water and energy integration systems: case studies of the process industry in portugal. Energies 17:195 (2023). https://doi.org/10.3390/en17010195
5. Castro Oliveira M. Simulation and Optimisation of Water and Energy Integration Systems (WEIS): An Innovative Approach for Process Industries. *PhD Thesis, IST* (2023) https://scholar.tecnico.ulisboa.pt/records/KwfdB28dqObhPEiL7RtB3VJ6YxviyXFEhWnt

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Targeted Olfactory Molecule Generation for Vanilla Scents Using Generative Flow Networks

Bruno C. L. Rodrigues[a]*, Paul J. Groening[a], Laura Sisson[b], Mumin Enis Leblebici[c] and Idelfonso B. R. Nogueira[a]

[a] Department of Chemical Engineering; Norwegian University of Science and Technology, Gløshaugen, Trondheim, 7034, Norway
[b] Patina, New York, New York, United States of America
[c] Department of Chemical Engineering, Center for Industrial Process Technology, KU Leuven, Agoralaan Building B, Diepenbeek 3590, Belgium
* Corresponding Author: bruno.rodrigues@ntnu.no.

ABSTRACT

This work explores Generative Flow Networks (GFlowNets) as a computational approach for sustainable fragrance design, focusing on generating novel molecules that reproduce the scent profile of vanillin while reducing reliance on resource-intensive synthesis and environmentally vulnerable natural sources. An integrated pipeline couples a GFlowNet generator with a fragrance note predictor, which guides learning toward a target odor by rewarding molecules predicted to be aromatically similar to vanillin. Chemical validity and realism are enforced through chemistry filters that penalize unstable or implausible structures and through an odorless-vs-odorant classifier, so only chemically and olfactorily plausible candidates are selected. The agent is trained in a hybrid offline–online regime, implementing reinforcement-based exploration, with hyperparameters tuned via Bayesian optimization. As an independent validation layer, an olfactory receptor docking model estimates binding affinities to receptors associated with vanillin, and several generated candidates show comparable or higher predicted affinities than vanillin while retaining structural novelty. Overall, the results suggest that GFlowNets can capture multidimensional olfactory patterns and support CAPE-style exploration of vast chemical spaces under uncertainty, enabling more efficient exploration of olfactory ingredients.

Keywords: GFlowNet, molecular generation, odorant design, sustainability, green chemistry, CAPE, fragrance engineering, generative AI, vanillin

INTRODUCTION

Odors play a central role in human perception, influencing emotion, memory, and cultural experience [1-3]. Odors are central to the global fragrance and flavor industry, underpinning applications ranging from perfumes and cosmetics to food flavorings and household products [4]. The market was valued at approximately 30.61 billion USD in 2023 and is projected to grow to 44.6 billion USD by 2030 [4]. Many commercially relevant odorants rely on resource-intensive laboratory experimentation or extraction from rare and environmentally vulnerable natural sources, raising concerns related to sustainability, scalability, and long-term supply security [3].

Recent advances in artificial intelligence, particularly in generative modeling, offer new opportunities for in silico molecular design [5]. By learning from existing datasets of odorants and their associated properties, generative models can propose candidate molecules with desired sensory characteristics, potentially reducing reliance on trial-and-error laboratory synthesis. From a sustainability perspective, such approaches align with green chemistry principles by enabling the design of synthetic alternatives to naturally sourced odorants, thereby reducing environmental pressure while maintaining sensory fidelity.

Among generative approaches, Generative Flow Networks (GFlowNets) [6] have emerged as a promising framework for sampling diverse candidates from complex target distributions. They provide an alternative to conventional reinforcement learning (RL) and Markov Chain Monte Carlo methods, which often struggle in drug

discovery settings to identify diverse molecules efficiently [7]. GFlowNets are trained so that the probability of sampling a terminal molecule is proportional to its assigned reward, thereby aligning generation with the desired target distribution. Training proceeds in an active learning setting similar to RL, where an agent, typically parameterized by a graph neural network, constructs molecules through sequential actions that define trajectories [6]. This formulation supports exploration of large combinatorial chemical spaces while preserving diversity among high-reward solutions, which is particularly important in molecular design where multiple distinct structures can satisfy the same functional objective.

This work investigates the use of GFlowNets for targeted odorant generation, focusing on reproducing the characteristic vanilla scent associated with vanillin. Vanilla is among the most widely preferred odor profiles across cultures [8], and the interactions between vanillin and multiple olfactory receptors have been well studied in prior work, making it a well-suited reference target for this project [9]. We present a GFlowNet pipeline that integrates learned olfactory representations, explicit chemical validity constraints, and an independent bio-informed validation step, enabling the generation of structurally novel candidates consistent with the desired sensory profile.

BACKGROUND AND RELATED WORK

Computational approaches to odor modeling have gained increasing attention. Graph neural network-based models have demonstrated the ability to embed molecules into high-dimensional representations that capture perceptual odor similarities across multiple scent descriptors [10]. These models enable data-driven prediction of odor categories and provide a quantitative basis for guiding molecular design [3, 10].

Generative molecular design has relied on architectures such as Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), and Reinforcement Learning (RL). While effective in specific scenarios, these methods face distinct challenges in navigating discrete chemical spaces. As highlighted in comprehensive reviews, VAEs and GANs often struggle to maintain chemical validity or suffer from mode collapse, failing to capture the full diversity of the target distribution [11]. Similarly, standard optimization strategies can have difficulty balancing exploration and exploitation, frequently leading to premature convergence on narrow regions of chemical space rather than generating a diverse set of high-performing candidates [11].

GFlowNets address this challenge by learning a policy that samples entire trajectories through a construction space such that the probability of reaching a terminal state is proportional to its assigned reward [6]. This formulation enables controlled exploration of multiple high-reward regions and is particularly attractive for problems where the reward function is sparse, noisy, or expensive to evaluate. Previous work has shown that GFlowNets are well suited for molecular graph construction [12], where molecules are built sequentially through valid chemical transformations.

Despite these advantages, the application of GFlowNets to olfactory molecule generation remains relatively unexplored, particularly in settings that require explicit chemical plausibility constraints and evaluation beyond internal reward-model scores.

METHODOLOGY

Problem Formulation

Building on the objective of generating novel molecular structures with olfactory characteristics similar to vanillin, molecules are represented as graphs and constructed step-by-step through a sequence of chemically valid actions. The generative process is framed as a trajectory sampling problem in which each completed molecule corresponds to a terminal state and is assigned a scalar reward reflecting its suitability as a vanilla-like odorant.

Formally, a state represents the current partial molecular graph, while an action applies a valid graph modification, such as adding an atom or bond, or a terminal stop action that completes the molecule. The GFlowNet agent is parameterized by a policy that assigns probabilities to valid actions given the current state. Applying an action transforms the current state into a new partial graph. A sequence of such state transitions defines a trajectory that terminates in a complete molecule. Each terminal molecule is assigned a reward based on its predicted similarity to the vanilla odor profile, together with chemical validity and odorant feasibility constraints. Training is performed in a hybrid offline–online regime, where offline samples are drawn from a dataset of known fragrant molecules [13] and online samples are generated by the current policy. Both sets of samples are used jointly to compute the training loss and update the model parameters.

The reward function is designed to favor molecules that (i) exhibit a high predicted similarity to the vanilla odor profile, (ii) fall within the distribution of known odorants, and (iii) satisfy chemical validity and safety constraints. This multi-component formulation ensures that high-reward samples are not only perceptually relevant but also chemically plausible. Figure 1 summarizes the workflow.

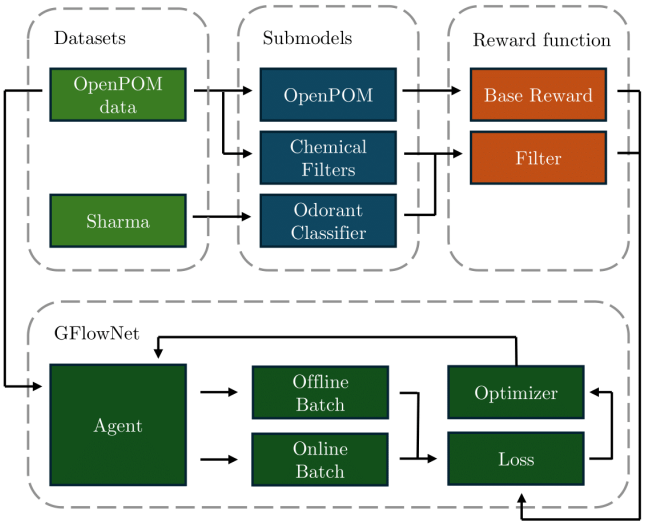

Figure 1. – Flowchart summarizing the methodology.

GFlowNet Architecture

A GFlowNet agent is employed to sample molecular construction trajectories. At each step, the agent selects an action that modifies the current molecular graph, such as adding an atom or bond, subject to chemical valency rules. A special stop action terminates the trajectory and yields a complete molecule.

The agent's policy is trained to distribute probability mass across terminal molecules in proportion to their reward values. Training follows the standard GFlowNet paradigm, where batches of trajectories are sampled, rewards are computed for terminal states, and network parameters are updated to minimize the discrepancy between inflow and outflow of probability mass along each trajectory.

Olfactory Reward Modeling

Olfactory similarity is quantified using OpenPOM [10, 13], a message-passing neural network trained to predict odor categories across a wide range of fragrance notes. For each generated molecule, OpenPOM outputs a probability vector over odor descriptors, including vanilla-related dimensions such as sweet, phenolic, smoky, spicy, floral, and vanilla.

The base reward component is computed by comparing cosine similarity between the predicted odor profile of a generated molecule to that of vanillin, encouraging alignment across multiple odor dimensions rather than a single scalar label. This multidimensional formulation allows the model to capture the complex perceptual structure of vanilla scent.

The reward formulation is intentionally defined over the full predicted odor probability vector rather than relying solely on the vanilla odor probability. This design choice addresses the inherent sparsity of vanilla-labeled molecules in the dataset and reflects the perceptual reality that vanilla scent is characterized by a combination of correlated odor dimensions rather than a single categorical label. By rewarding similarity across all odor dimensions, molecules with moderate vanilla probability but high perceptual coherence are retained, resulting in a denser and more informative reward landscape during training. This approach was found to be essential for maintaining learning signal in early and intermediate training stages.

Chemical and Odorant Constraints

To prevent the generation of chemically implausible or unsafe molecules, several constraint mechanisms are incorporated. Rule-based chemical filters penalize structures containing reactive or undesirable functional groups, excessive ring strain, or unstable motifs. These filters are derived from properties observed in the training data and from established chemical safety considerations. The filters considered and their ranges are shown in Table 1. Here, LogP denotes the logarithm of the octanol–water partition coefficient, TPSA the topological polar surface area, and NRB the number of rotatable bonds.

Table 1: Chemical filters considered and their ranges

Parameter	Min	Max
LogP	0.06	4.91
Molecular weight	100.12	282.46
Hydrogen bond acceptors + donors	0	5
TPSA	0	63.32
NRB	0	11
NO_2	0	0
Heteroatoms	0	5
Aromatic rings	0	1
Atoms	0	20
Triple bonds	0	0
Radical electrons	0	0
Oxygen chain length	0	1

In addition, a random forest classifier trained on molecular fingerprints is used to distinguish odorant molecules from non-odorants. This classifier serves as an auxiliary filter to ensure that generated candidates are likely to be perceptible at all, rather than chemically valid but odorless. The dataset used for training this model is the Sharma dataset [14].

Only molecules passing both chemical and odorant plausibility checks are assigned nonzero rewards and propagated during training, reducing noise in the learning signal and improving convergence, though they are still quantified for loss metrics.

Training Strategy and Hyperparameter Optimization

The GFlowNet is trained in a hybrid offline–online

regime. Offline training batches are drawn from the OpenPOM dataset, providing stable supervision during early learning phases by exposing the model to chemically valid and odor-relevant molecular structures. Online batches are generated by the current policy, enabling exploration of novel molecular graphs beyond the training data and progressively refining the sampling distribution toward high-reward regions of chemical space.

Training can be formulated using the Trajectory Balance (TB) objective, which enforces consistency between the probability of sampling a complete trajectory and the reward assigned to its terminal state. In this formulation, the product of forward transition probabilities along a trajectory is matched to the terminal reward, ensuring that the probability of sampling a molecule is proportional to its reward. The TB objective provides a principled mechanism for learning a target distribution over terminal states, but it is known to be sensitive to long trajectories, since trajectory probabilities are obtained as products of many terms and can therefore become extremely small, leading to large loss values when paired with mismatched rewards.

To mitigate this issue, the Subtrajectory Balance (SubTB) objective is employed. Instead of enforcing balance only at the level of complete trajectories, SubTB decomposes each trajectory into overlapping subtrajectories and applies the balance condition to each of them. The final loss is computed as a weighted aggregation over all subtrajectories. This formulation preserves the theoretical objective of matching the target distribution while reducing the impact of extreme probability values associated with long molecular construction paths. In practice, SubTB provides improved numerical stability and smoother training dynamics, particularly in molecular generation tasks where trajectory lengths can vary substantially.

In this work, SubTB is preferred over the standard TB objective due to the presence of long and heterogeneous construction trajectories and the additional noise introduced by hybrid offline–online training. While SubTB does not fully eliminate instability, it significantly reduces loss explosions and improves convergence behavior compared to TB, making it more suitable for the present setting.

Model training follows the standard GFlowNet procedure. At each training iteration, a batch of trajectories is constructed by sampling actions from the current policy. Each trajectory starts from an initial empty molecular graph and is extended through sequential graph-modifying actions until a terminal stop action is selected, yielding a complete molecule. Rewards are computed only for terminal states and reflect olfactory similarity, chemical validity, and odorant plausibility. Model parameters are then updated by minimizing the SubTB loss computed over both offline and online trajectories.

Hyperparameters governing learning rates, batch composition, reward scaling, and gradient clipping are optimized using Bayesian optimization with Optuna [15]. The tuned parameters and the two final configurations retained for evaluation are reported explicitly in Table 2. The sampling tau controls the rate at which the sampling policy is updated toward the latest trained parameters, trading off exploration stability against responsiveness. The Z learning rate sets the step size for updating the GFlowNet normalization constant Z, which stabilizes the global flow matching between trajectory probabilities and rewards. The learning rate controls the update magnitude for the main network parameters, whereas momentum accelerates optimization by smoothing gradients across iterations. Weight decay applies L_2 regularization to limit parameter growth and reduce overfitting. Finally, gradient clipping caps the gradient norm to prevent occasional large updates, which is particularly important under long trajectories and noisy offline–online mixing. Based on validation performance and training stability, we retain two optimized configurations (Table 2) for the final experiments.

Table 2: Final hyperparameter values

Parameter	Search space	First	Second
Sampling tau	0.8 - 1	0.95864	0.98403
Z learning rate	10^{-8} - 10^{-1}	9.1725 x 10^{-4}	3.1573 x 10^{-5}
Gradient clipping	0.1 – 100	9.9068	32.278
Momentum	0.5 – 0.9999	0.65129	0.65129
Learning rate	10^{-8} - 10^{-1}	3.3456 × 10^{-4}	2.5660 × 10^{-4}
Weight decay	10^{-10} - 10^{-4}	3.9731 × 10^{-7}	3.9731 × 10^{-7}

VALIDATION AND EVALUATION STRATEGY

Model validation was conducted using multiple complementary criteria. Offline validation was performed by computing loss values on a held-out set of unseen molecules from the OpenPOM dataset, while online validation considers the molecules being generated by the GFlowNet model. However, because the OpenPOM model is used both during training and reward computation, these metrics alone are susceptible to circular validation. Olfactory receptor docking was therefore introduced as an independent, physically motivated validation layer, selected for its interpretability and relevance to

molecular-level odor perception.

A receptor docking model is constructed using AutoDock Vina [16] to simulate binding interactions between generated molecules and ten human olfactory receptors known to exhibit high affinity for vanillin. Protein structures are obtained from AlphaFold [17], and active site bounding boxes are defined manually based on literature-reported amino acid residues. Docking exhaustiveness is set to a moderate value to balance computational cost and accuracy.

Prior to evaluating generated molecules, the docking model is benchmarked by repeatedly simulating vanillin–receptor interactions. The resulting binding energy distributions exhibit low variance and reasonable agreement with literature values, indicating that the model provides a suitable independent metric for comparative evaluation. After training, generated molecules can be sampled and checked on existing databases such as PubChem [18] to check if the they are already catalogued, or novel. Figure 2 summarizes the validation strategy.

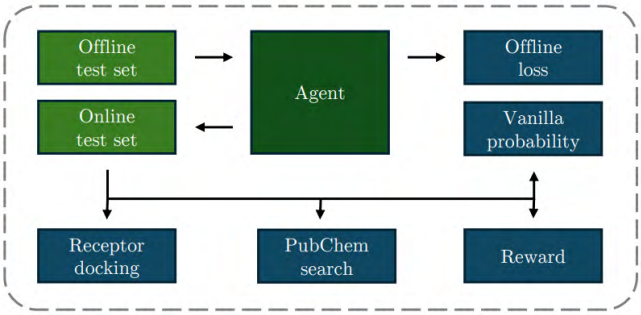

Figure 2. – Scheme of the validation strategy.

RESULTS AND DISCUSSION

The model was trained with two different hyperparameter configurations. For each training regimen, 500 molecules were sampled and analyzed at two checkpoints: an intermediate stage (3000 iterations) and the end of training (6000 iterations). Across both checkpoints, the first hyperparameter configuration consistently outperformed the second.

The best-performing configuration achieved median reward values above 0.5, with improved reward concentration at the final checkpoint compared to the intermediate stage, indicating progressive refinement of the sampling distribution. Median docking energies across vanillin-sensitive receptors were comparable to those obtained for vanillin itself, with reduced variance relative to intermediate checkpoints. In contrast, the second hyperparameter configuration exhibited both lower median rewards and broader docking energy distributions, indicating less consistent receptor engagement. These results highlight the sensitivity of GFlowNet performance to hyperparameter selection and training stability.

Qualitative inspection of high-reward molecules reveals structural diversity, including candidates with and without ring systems. While certain high-reward structures exhibit features that may pose synthetic challenges, others remain relatively simple and chemically feasible. Radar plots of predicted odor profiles show that the best-performing molecules align closely with the six dominant odor dimensions associated with vanillin. This qualitative diversity is more pronounced for the first hyperparameter configuration, which consistently samples a broader range of structurally distinct candidates at comparable reward levels. Figure 3 shows four of the best molecules sampled in the first hyperparameter configuration.

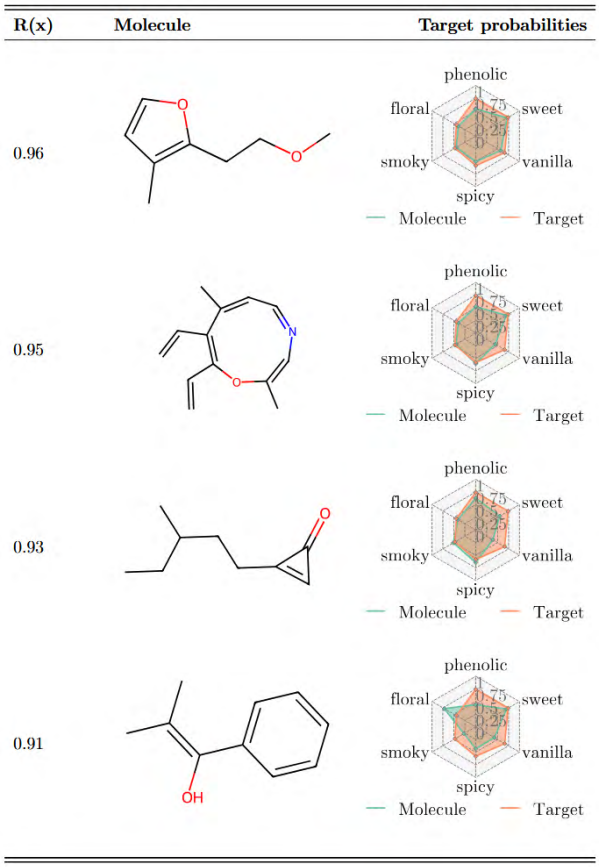

Figure 3. – Four best molecules sampled from the best configuration of hyperparameters.

Docking-based validation indicates that many generated molecules exhibit binding affinities to vanillin-sensitive receptors that are comparable to, and in some cases stronger than, those predicted for vanillin itself. Although median binding energies are generally slightly weaker than vanillin, the distributions suggest meaningful receptor interaction rather than random binding. Importantly, the best-performing model shows a reduction in high-energy outliers relative to intermediate

checkpoints, indicating improved consistency. Figure 4 shows the predicted binding energies for the molecules sampled using the first configuration for several receptors known to bind with vanillin. Additionally, the estimated values for vanillin using the AutoDock Vina and and the values from literature [9] are given for comparison.

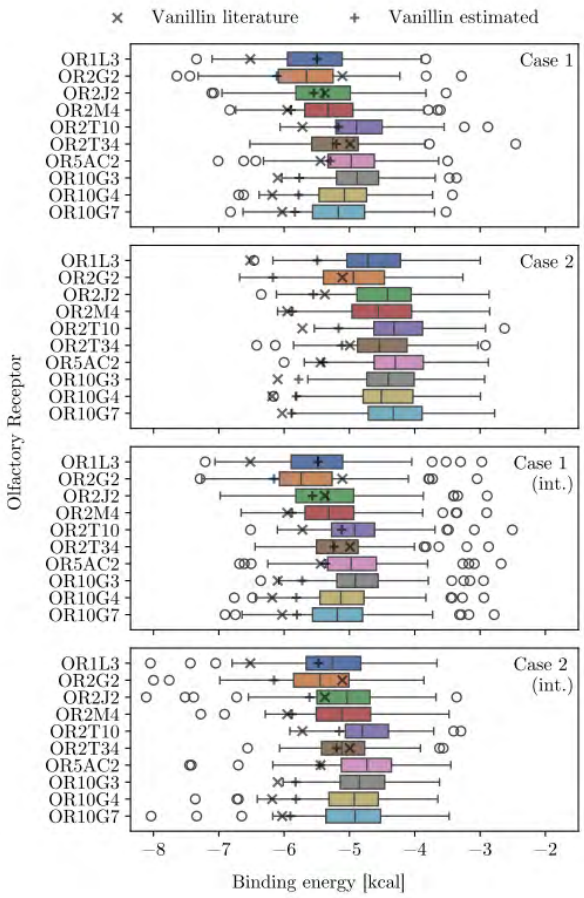

Figure 4. - Docking binding energies for the molecules sampled from the first hyperparameter configuration run.

The second hyperparameter configuration yields less favorable results, with lower reward values and weaker receptor binding. Training dynamics for this configuration exhibit stronger oscillations, suggesting instability and misalignment between offline and online objectives. These findings highlight the sensitivity of GFlowNet training to hyperparameter choices and reward sparsity.

Despite the overall improvement in sampled molecule quality over training, learning dynamics presented instabilities. In particular, oscillatory reward trajectories and sporadic spikes in offline loss were observed. These effects are attributed to a misalignment between offline and online training objectives, since offline molecules bypass several validity and plausibility filters that are applied during online sampling. As a result, offline trajectories can induce large loss gradients that do not reflect the effective online sampling distribution, leading to noisy parameter updates. As anticipated, the use of the Sub-Trajectory Balance (SubTB) objective reduced the severity of loss explosions associated with long trajectories, although oscillatory dynamics were still observed. Figure 5 and 6 show the loss metrics during training for both configurations. Note that reward curves shown during training include unfiltered offline and online trajectories, which is lower than the reward of chemically valid, odorant-positive molecules only, reported earlier. Figure 7 shows the percentage of molecules generated by the model that could be found in the PubChem [18] database; molecules were sampled at 3000 and 6000 iterations, as previously described. For this figure, 500 molecules were sampled without applying chemistry filters, presented under “All”; the second row shows the results considering only chemically valid molecules, and predicted to be odoriferous.

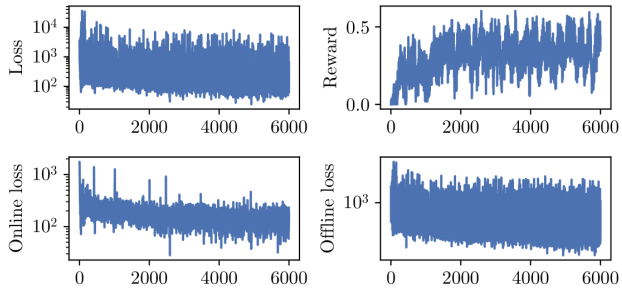

Figure 5. – Loss metrics for the first hyperparameter configuration.

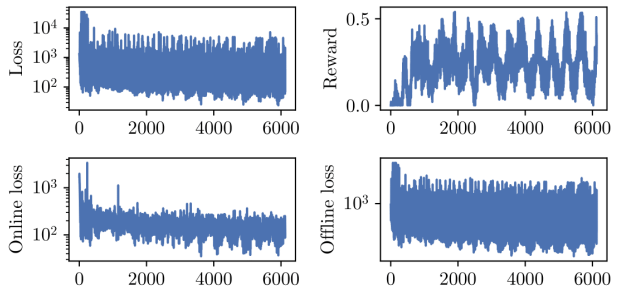

Figure 6. – Loss metrics for the second hyperparameter configuration.

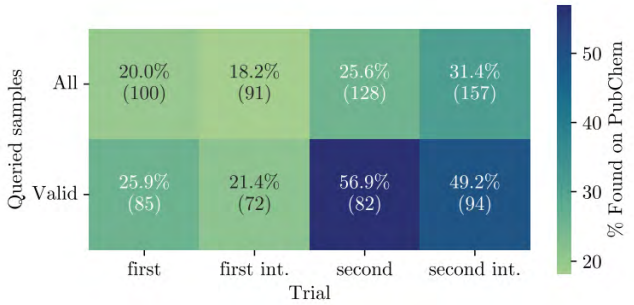

Figure 7. – Percentage of sampled molecules found in Pubchem.

Taken together, these results indicate that model performance is governed not only by the expressiveness of the reward function, but also by the interaction between offline supervision, online exploration, and trajectory-level credit assignment. While the Sub-Trajectory Balance objective reduces the impact of extreme loss values, instability persists when offline trajectories induce gradients that are inconsistent with the effective online sampling distribution. This misalignment manifests as oscillatory reward dynamics and sporadic spikes in offline loss, particularly in less favorable hyperparameter configurations.

From a modeling perspective, these findings suggest that training stability, rather than representational capacity, currently constitutes the primary bottleneck for applying GFlowNets to chemically constrained generative tasks. Improving the alignment between offline and online objectives, for example through staged training or uncertainty-aware weighting of offline samples, is therefore expected to yield more consistent convergence and improved sampling efficiency.

CONCLUSION

This work demonstrates that Generative Flow Networks can be applied to targeted olfactory molecule generation, producing structurally novel candidates with predicted sensory characteristics similar to vanillin. Despite training instabilities and reward noise, the proposed framework consistently generates high-reward molecules and exhibits meaningful agreement with independent receptor-based validation metrics.

At the same time, the study highlights important limitations. Oscillatory training dynamics and misalignment between offline and online objectives affect loss stability and complicate hyperparameter optimization. In addition, high reward or strong predicted receptor affinity does not guarantee synthetic accessibility, and receptor docking remains an approximate proxy for olfactory perception. Future work should therefore focus on improving reward smoothness, incorporating synthesis-aware and uncertainty-aware objectives, and refining training strategies to better balance exploration and constraint enforcement. Experimental validation and synthesis-informed modeling represent critical next steps toward practical deployment.

From a broader CAPE perspective, the proposed approach can be interpreted as a feasible-region exploration problem in a vast molecular design space under uncertainty. By enabling the generation of synthetic alternatives to naturally sourced odorants, the methodology contributes to sustainability and supply-chain resilience while extending established CAPE concepts from process-level design to molecular-level innovation.

ACKNOWLEDGEMENTS

BCLR and IBRN acknowledge funding from SCENTINEL project (UV101923100).

AUTHOR IDENTIFIERS

Author ORCIDs:
Rodrigues BCL: 0009-0006-0513-3699
Leblebici ME: 0000-0003-4599-9412
Nogueira IBR: 0000-0002-0963-6449

REFERENCES

1. Herz RS, Schankler C, Beland S. Olfaction, emotion and associative learning: effects on motivated behavior. Motiv Emot 28:363-383 (2004). https://doi.org/10.1007/s11031-004-2389-x
2. Bembibre C, Strli? M. Smell of heritage: a framework for the identification, analysis and archival of historic odours. Herit Sci 5: (2017). https://doi.org/10.1186/s40494-016-0114-1
3. Rodrigues BCL, Santana VV, Murins S, Nogueira IBR. Molecule generation and optimization for efficient fragrance creation. Ind. Eng. Chem. Res. 63:14480-14494 (2024). https://doi.org/10.1021/acs.iecr.4c00650
4. Grand View Research. Flavors and fragrances market size & share report (2026 – 2033). https://www.grandviewresearch.com/industry-analysis/flavors-fragrances-market
5. Elton DC, Boukouvalas Z, Fuge MD, Chung PW. Deep learning for molecular design—a review of the state of the art. Mol. Syst. Des. Eng. 4:828-849 (2019). https://doi.org/10.1039/c9me00039a
6. Bengio Y, Lahlou S, Deleu T, Hu EJ et al. GFlowNet foundations. arXiv (2026) https://doi.org/10.48550/arXiv.2111.09266
7. Bengio E, Jain M, Korablyov M, Precup D et al. Flow network based generative models for non-iterative diverse candidate generation. arXiv (2021) https://doi.org/10.48550/arXiv.2106.04399
8. Spence C. Odour hedonics and the ubiquitous appeal of vanilla. Nat Food 3:837-846 (2022). https://doi.org/10.1038/s43016-022-00611-x
9. Zeng S, Zhang L, Zheng R, Li P, Fu Y, Xi H, Wang D, Pu D, Mao J, Sun B, Sun S, Zhang Y. Molecular recognition mechanisms of vanillin and high-throughput screening of its analogs based on olfactory receptors. LWT 202:116305 (2024). https://doi.org/10.1016/j.lwt.2024.116305
10. Lee BK, Mayhew EJ, Sanchez-Lengeling B, Wei JN, Qian WW, Little KA, Andres M, Nguyen BB, Moloy T, Yasonik J, Parker JK, Gerkin RC, Mainland JD, Wiltschko AB. A principal odor map unifies diverse

tasks in olfactory perception. Science 381:999-1006 (2023). https://doi.org/10.1126/science.ade4401
11. Sanchez-Lengeling B, Aspuru-Guzik A. Inverse molecular design using machine learning: generative models for matter engineering. Science 361:360-365 (2018). https://doi.org/10.1126/science.aat2663
12. Madan K, Rector-Brooks J, Korablyov M, Bengio E et al. Learning GFlowNets from partial episodes for improved convergence and stability. arXiv (2023) https://doi.org/10.48550/arXiv.2209.12782
13. Barsainyan AA, Kumar R, Saha P, Schmuker M. https://github.com/BioMachineLearning/openpom
14. Sharma A, Saha BK, Kumar R, Varadwaj PK. Olfactionbase: a repository to explore odors, odorants, olfactory receptors and odorant–receptor interactions. Nucleic Acids Research 50:D678-D686 (2021). https://doi.org/10.1093/nar/gkab763
15. Akiba T, Sano S, Yanase T, Ohta T, Koyama M. Optuna. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining :2623-2631 (2019). https://doi.org/10.1145/3292500.3330701
16. Trott O, Olson AJ. Autodock vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. J Comput Chem 31:455-461 (2009). https://doi.org/10.1002/jcc.21334
17. Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, Tunyasuvunakool K, Bates R, Žídek A, Potapenko A, Bridgland A, Meyer C, Kohl SAA, Ballard AJ, Cowie A, Romera-Paredes B, Nikolov S, Jain R, Adler J, Back T, Petersen S, Reiman D, Clancy E, Zielinski M, Steinegger M, Pacholska M, Berghammer T, Bodenstein S, Silver D, Vinyals O, Senior AW, Kavukcuoglu K, Kohli P, Hassabis D. Highly accurate protein structure prediction with alphafold. Nature 596:583-589 (2021). https://doi.org/10.1038/s41586-021-03819-2
18. PubChem. https://pubchem.ncbi.nlm.nih.gov/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Strategy for Limiting the Effects of Nonconvexities in Mixed-Integer Nonlinear Programming Reformulation of Nonconvex Generalized Disjunctive Programs

Miloš Bogataj[a*], Chiara Železnik[a], and Zdravko Kravanja[a]

[a] University of Maribor, Faculty of Chemistry and Chemical Engineering, Maribor, Slovenia
* Corresponding Author: milos.bogataj@um.si.

ABSTRACT

Nonconvex generalized disjunctive programs (GDPs) frequently arise in chemical engineering applications and are commonly reformulated as mixed-integer nonlinear programs (MINLPs). However, nonconvexities in these reformulations often lead to numerical difficulties, sensitivity to initialization, and degraded solution quality when solved with general-purpose MINLP solvers. This work proposes a two-phase strategy to mitigate these effects by generating improved initial points through the solution of a sequence of relaxed MINLPs, which are subsequently used to initialize the original formulation. The approach is evaluated on a family of purely disjunctive benchmark problems, referred to as the Crescent problems, with sizes ranging from 60 to 1000 binary variables. Numerical experiments using the DICOPT and SBB solvers assess performance in terms of objective value distributions, the percentage of feasible initial points, and average constraint violation. The results indicate that the proposed strategy improves solution quality, increases the likelihood of feasibility, and reduces the magnitude of constraint violations across all problem sizes.

Keywords: nonconvex optimization, mixed-integer nonlinear programming, generalized disjunctive programming, global optimization, local optimization

INTRODUCTION

Optimization problems in chemical engineering are predominantly nonlinear because the relationships describing reaction kinetics, thermodynamic equilibria, and transport phenomena are inherently governed by nonlinear dependencies among variables. Additionally, constraints such as energy and mass balances, as well as equipment design and sizing relationships, are often described by nonlinear – generally nonconvex – equations. In the context of optimization, this poses two problems: first, the need to provide good initial points, and second, the potential existence of multiple local optima.

In the past two decades, significant advances have been made in the field of global optimization. Numerous global optimization strategies have been reported in the literature [1-4], some designed for very specific problems and others more general. Additionally, there has been substantial improvement in the computational efficiency of commercially available deterministic general-purpose global optimization solvers (BARON, ANTIGONE, LINDO-Global). However, when nonconvex nonlinear optimization problems include discrete decisions – namely, mixed-integer nonlinear problems (MINLP)—the ability to find the global optimum arguably decreases exponentially with problem size. Therefore, research into approaches that provide good locally optimal solutions using general MINLP solvers is relevant.

PROBLEM STATEMENT

The "Crescent" problem considered in this work is stated as a Generalized Disjunctive Program (1). It is a purely disjunctive problem consisting of a nonconvex objective function, D disjunctions and T_d terms per disjunction d.

Each term t includes a Boolean variable Y_{td}, two nonlinear constraints (one of which is nonconvex), and a "charge" parameter γ_{td}. The logical propositions

$\Omega(Y_{td})$ ensure that only one Boolean variable in each disjunction is $True$.

$$\min Z = (x_1 - 3)^2 - (x_2 - 3)^2 + \sum_{d \in D} c_d$$

s.t.

$$\bigvee_{t \in T_d} \begin{bmatrix} Y_{td} \\ (x_1 - a_1)^2 + (x_2 - b_1)^2 - r_1^2 \le 0 \\ (x_1 - a_2)^2 + (x_2 - b_2)^2 - r_2^2 \ge 0 \\ c_d = \gamma_{td} \end{bmatrix}, \quad d \in D$$

$$\Omega(Y) = True$$

$$\boldsymbol{x}^{\mathrm{LO}} \le \boldsymbol{x} \in \mathbb{R}^2 \le \boldsymbol{x}^{\mathrm{UP}}, c_d \in \mathbb{R}^1,$$

$$Y_{td} \in \{True, False\}, t \in T_d, d \in D \qquad (1)$$

The disjunctive nonconvex feasible region of a fifteen-term disjunction is shown schematically in Figure 1.

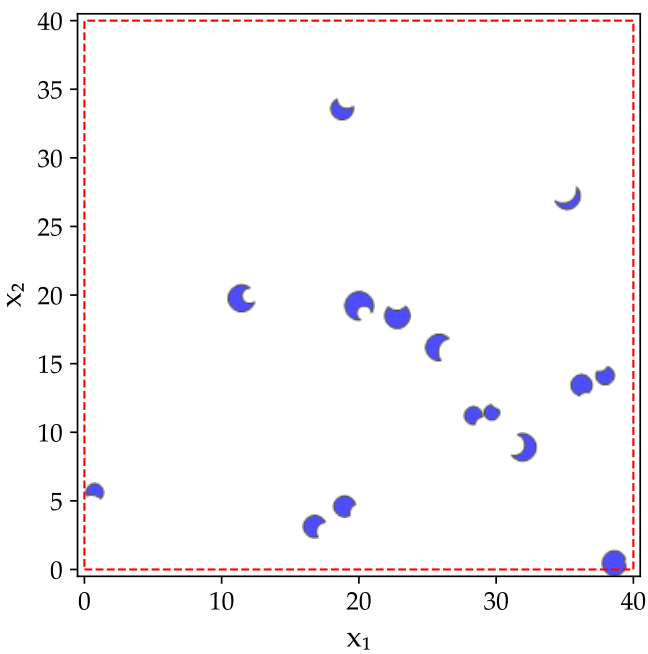

Figure 1. Projection of a disjunctive feasible region onto x_1-x_2 plane.

GDP-to-MINLP reformulation

The proposed GDP-to-MINLP reformulation is a variation of the method proposed by Bogataj and Kravanja [5] and can be viewed as a hybrid between the Hull Reformulation [6-7] and the Big-M reformulation.

The nonlinear constraints in the terms of the disjunctions $g(\boldsymbol{x}) \le 0$ are formulated in terms od disaggregated variables $\boldsymbol{x}^{\mathrm{D}}$as $g(\boldsymbol{x}^{\mathrm{D}}) - (1-y)g(\boldsymbol{x}^{\mathrm{F}}) \le 0$, where $y \in \{0,1\}$, and $\boldsymbol{x}^{\mathrm{D,LO}} \le \boldsymbol{x}^{\mathrm{F}} \le \boldsymbol{x}^{\mathrm{D,UP}}$ is a fixed parameter. The disaggregated variables are bounded by $\boldsymbol{x}^{\mathrm{LO}}\boldsymbol{y} + (\boldsymbol{1} - \boldsymbol{y})\boldsymbol{x}^{\mathrm{F}} \le \boldsymbol{x}^{\mathrm{D}} \le \boldsymbol{x}^{\mathrm{UP}}\boldsymbol{y} + (\boldsymbol{1} - \boldsymbol{y})\boldsymbol{x}^{\mathrm{F}}$. The latter constraint enforces the disaggregated variable to satisfy $x^{\mathrm{LO}} \le x^{\mathrm{F}} \le x^{\mathrm{UP}}$ if $y = 1$, and $x^{\mathrm{D}} = x^{\mathrm{F}}$ if $y = 0$. Finally, the constraint $x = \sum(x^{\mathrm{D}} - (1-y)x^{\mathrm{F}})$ ensures the regularity of the relationship between the original variable x and the disaggregated variables $\boldsymbol{x}^{\mathrm{D}}$. For more information on the reformulation, we refer the reader to Bogataj and Kravanja [5].

Algorithm

The algorithm (2) consists of two phases. In the first phase, the continuous relaxation is tightened by solving a sequence of relaxed MINLPs using a local NLP solver (e.g., CONOPT). In the second phase, the MINLP formulation is solved from the initial points determined in the first phase. Regarding first-phase initialization, it should be noted that if the initial points $(\boldsymbol{x}_0, \boldsymbol{x}_0^{\mathrm{D}})$ are selected randomly, most are expected to be infeasible with respect to the constraints in the original GDP. However, it should also be emphasized that any initial point $\boldsymbol{x}_0^{\mathrm{D}}$ for which the corresponding constraints $g(\boldsymbol{x}) \le 0$ in the GDP are defined is a feasible point from the perspective of the proposed MINLP formulation.

$$\begin{aligned}
&\mathit{set}\ Z = -\infty \\
&\mathit{set}\ \boldsymbol{x}_0, \boldsymbol{x}_0^D, \boldsymbol{x}^F = \boldsymbol{x}_0^D \\
&\mathit{loop}\ (i = 1, \ldots, I \\
&\quad \mathit{solve}\ \text{R-MINLP} \quad \min(Z\ ) \\
&\quad \mathit{get}\ \ Z_i^*, \boldsymbol{x}_i^*, \boldsymbol{x}_i^{D*} \\
&\qquad \text{if}\ |Z - Z_i^*| \le \varepsilon\ \ \text{break} \\
&\qquad \mathit{else}\ Z = Z_i^*, \boldsymbol{x}^F = \boldsymbol{x}_i^{D*}) \\
&\mathit{set}\ \boldsymbol{x}_0 = \boldsymbol{x}_i^*, \boldsymbol{x}_0^D = \boldsymbol{x}_i^{D*}, \boldsymbol{x}^F = \boldsymbol{x}_i^{D*} \\
&\mathit{solve}\ \text{MINLP}\ \min(Z)
\end{aligned} \qquad (2)$$

Table 1: Crescent problem instances.

Problem	**Crescent 60**
Instance	1-60, 2-30, 3-20, 4-15, 5-12,
$D-T$	60-1, 30-20, 20-3, 15-4, 12-5
Problem	**Crescent 200**
Instance	2-100, 4-50, 5-40, 10-20, 25-8,
$D-T$	100-2, 50-4, 40-5, 20-10, 8-25
Problem	**Crescent 1000**
Instance	5-200, 8-125, 10-100, 20-50, 25-40,
$D-T$	200-5, 125-8, 100-10, 50-20, 40-25

NUMERICAL EXPERIMENTS

The proposed approach was evaluated on three groups of Crescent problem instances containing 60, 200, and 1000 binary variables, respectively; the full set of instances is listed in Table 1. Each instance, defined by the parameter pair (D, T), was solved $N_s = 200$ times from randomly generated initial points, with a time limit of 200 seconds per run. Two MINLP solvers were considered: the Outer Approximation/Equality Relaxation algorithm (DICOPT with CONOPT4 as NLP and CPLEX as MIP solver) and the Branch-and-Bound algorithm (SBB with CONOPT4 as NLP solver). To ensure the solvers do not terminate before the set time limit per run if the default

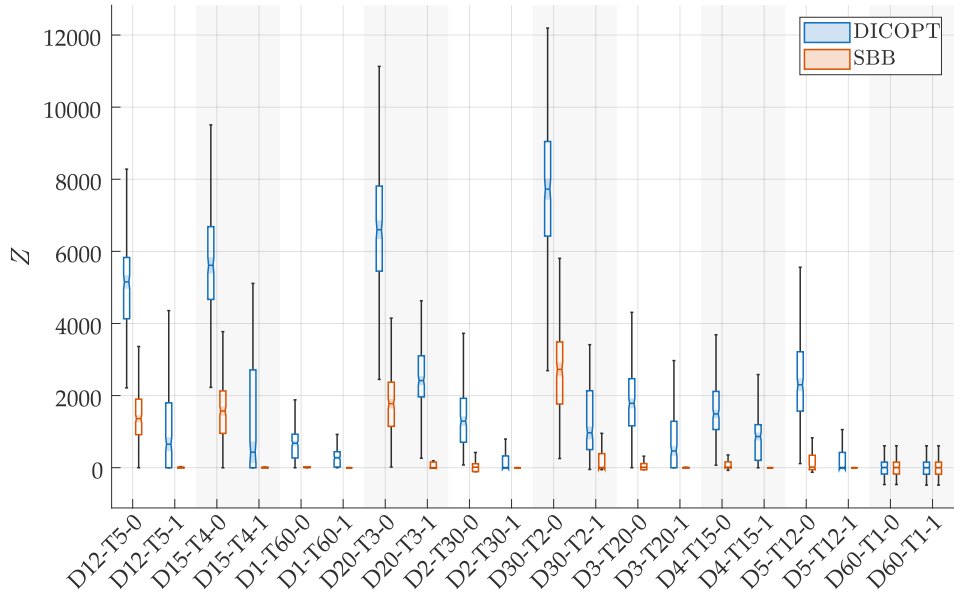

Figure 2. Distribution of objective function values (Crescent 60).

optimality criteria are not reached, the following changes were made to the default algorithmic options: DICOPT – convex = 0, maxcycles = 200, stop = 1; SBB – memnodes = 10000. For each solver, instances were solved both without the initial phase of algorithm (2) and using the full algorithm (2). For each instance and solver, both algorithmic variants were initialized using the same randomly generated starting points.

RESULTS

The goal of the numerical experiments was to assess the efficiency of the proposed methodology in limiting the effects of nonconvexities. The numerical experiments were statistically analyzed. Box-plots were used to characterize the distribution of objective functions. To further assess the efficiency of the algorithm, average constraint violation and the percentage of feasible initial points were utilized. The average constraint violation per instance $\bar{v}$ was computed as $\bar{v} = \frac{1}{2K}\sum_{k=1}^{K}\sum_{j=1}^{J} v_{j,k}$, where $v_{j,k} = \max\left\{0, g_j\left(\boldsymbol{x}_{0,k}^{\mathrm{D}}\right)\right\}, K = N_{\mathrm{s}} \bullet D \bullet T$, and $J = 2$. The percentage of initial feasible points was determined by $P_{\mathrm{f}} = N_{\mathrm{f}}/K$, where N_{f} denotes the number of initial points $\boldsymbol{x}_{0,k}^{\mathrm{D}}$ that satisfy both constraints in disjunctive terms.

Distributions of objective function values

Figures 2–4 summarize the distributions of the objective function values. For clarity, the box plots in each figure are best interpreted in groups of four, with each group corresponding to a single problem instance defined by the pair (D, T).

For each instance, results are shown for two solvers—DICOPT (blue) and SBB (red)—and for two algorithmic variants. In the first variant, the MINLP formulation is solved directly from randomly generated initial points, without the initial phase of the algorithm; in the second variant, the full algorithm (2), including the initial phase, is applied. In the Figures, these two variants are distinguished by the suffixes "–0" and "–1, " respectively, appended to the instance label.

To account for instance-dependent variability in objective function values, the results were normalized independently for each problem instance. For a given instance (D, T), four sets of objective values are therefore available, corresponding to the two solvers and the two algorithmic variants. The normalization constant was defined as the minimum of the four corresponding medians, and this value was subtracted from all objective function values associated with that instance. Note that although the normalization alters the absolute magnitude of the objective function values, it leaves the shape of the distributions and the relative performance of solvers within each instance unchanged.

Analysis of the box plots in Figure 2 (problem Crescent 60) clearly shows that better solutions are obtained when the full algorithm (2) is used. The solutions are improved not only in terms of objective values, but also in the reduced spread around the median. This is indicated by the shorter heights of the boxes and the shorter lengths of the whiskers. An exception to this general observation is instance (D60-T1), which is a somewhat

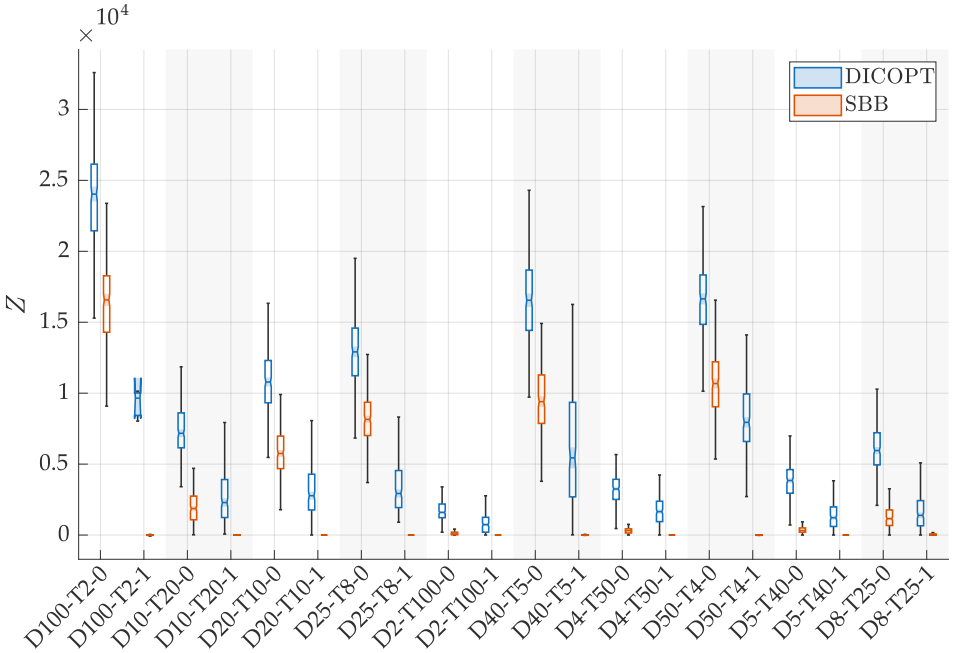

Figure 3. Distribution of objective function values (Crescent 200).

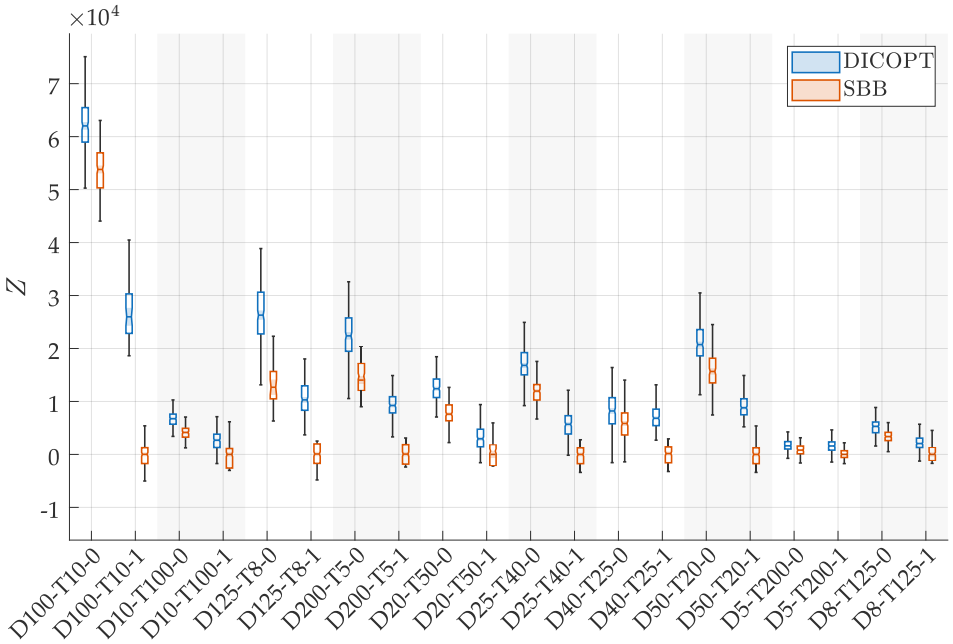

Figure 4. Distribution of objective function values (Crescent 1000).

trivial case because each disjunction consists of a single term. An integer solution is guaranteed in the relaxed MINLP for this instance; therefore, the impact of the first phase of the algorithm is negligible.

Similar observations can be made regarding the results shown in Figure 3 (problem Crescent 200) and Figure 4 (problem Crescent 1000). In both problems the improvement of objective function values is clearly visible when full algorithm (2) is applied. Perhaps the most noticeable among all is the considerably tightened distribution of objective function values in problem Crescent 200 obtained using SBB solver.

Percentage of feasible initial points

When local optimization solvers are applied to nonconvex optimization problems, the choice of initial points is critical for obtaining high-quality locally optimal solutions. Poor initial points can lead the solver to inferior local optima, infeasible solutions, or premature termination. In practice, effective initialization often relies on a deeper

understanding of the optimization problem, including valid bounds on the variables and, more importantly, the geometric structure of the feasible region. However, such information is difficult to obtain, especially for high-dimensional problems, which motivates the need for systematic and robust initialization procedures.

Figures 5–7 report the percentage of feasible initial points. In total, 12, 000 initial points were generated per instance for Crescent 60, 40, 000 for Crescent 200, and 200, 000 for Crescent 1000. As expected, given the random generation of starting points, only a small fraction of the initial points was feasible with respect to both constraints in (1) for each problem instance (see Table 2).

Table 2: Number of feasible randomly generated points.

Problem	**Crescent 60**
Instance (N_f)	1-60(**0**), 2-30(**0**), 3-20(**0**), 4-15(**0**), 5-12(**0**), 60-1(**0**), 30-20(**0**), 20-3(**0**), 15-4(**0**), 12-5(**0**)
Problem	**Crescent 200**
Instance (N_f)	2-100(**5**), 4-50(**1**), 5-40(**0**), 10-20(**0**), 25-8(**0**), 100-2(**0**), 50-4(**0**), 40-5(**0**), 20-10(**0**), 8-25(**0**)
Problem	**Crescent 1000**
Instance (N_f)	5-200(**3**), 8-125(**1**), 10-100(**0**), 20-50(**0**), 25-40(**0**), 200-5(**0**), 125-8(**0**), 100-10(**0**), 50-20(**0**), 40-25(**0**)

When the full algorithm (2) is applied, however, a substantial increase in the number of feasible points is observed. The highest percentage of feasible points, approaching 50 %, occurs for instance D200–T5–1 (Crescent 1000), whereas the lowest percentage, approximately 2 %, is observed for instance D2–T100–1 (Crescent 200).

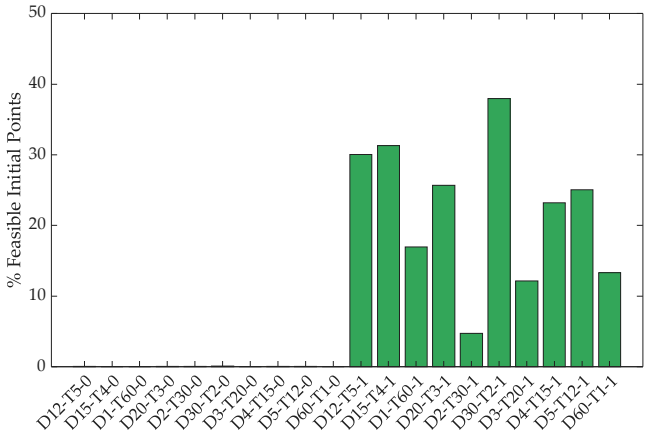

Figure 5. Percent of feasible initial points (Crescent 60).

Average constraint violation

Average constraint violation complements the percentage of feasible initial points by capturing the severity of constraint violations when infeasibility occurs. Lower average violations indicate that a solver consistently produces solutions close to the feasible region, reflecting potentially greater numerical robustness.

The results for the average constraint violation show consistent trends across all three problem sizes (see Figures 8-10). The magnitudes of average constraint violations for randomly generated initial points, shown in orange, are almost uniform across the three problems. In contrast, the configurations using the full algorithm (2), shown in green, yield lower average violations, often close to zero, which indicates convergence of initial points toward disjunctive feasible regions.

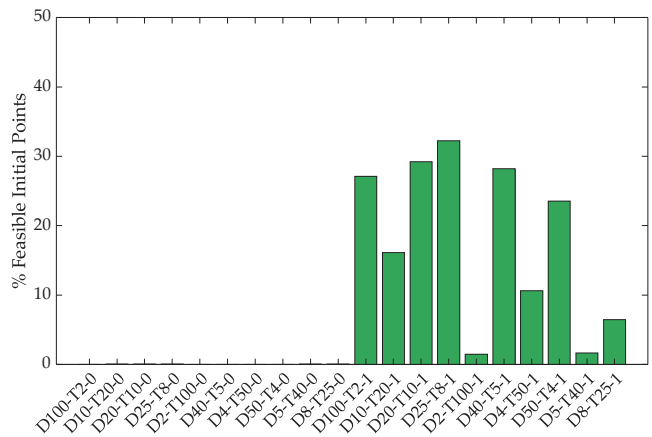

Figure 6. Percent of feasible initial points (Crescent 200).

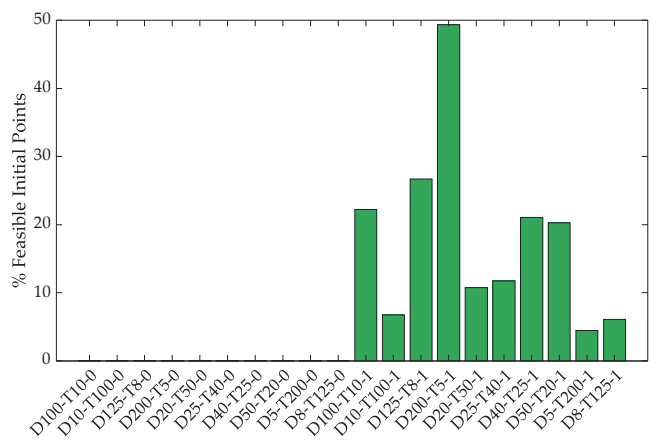

Figure 7. Percent of feasible initial points (Crescent 1000).

Computational efficiency

Finally, we briefly discuss the computational efficiency of the first phase of the algorithm. Across all

instances, the first phase required on average five R-MINLP iterations, with a minimum of two and a maximum of ten. Since these iterations involve solving NLP problems, the computational effort of the first phase is relatively modest compared to solving the full MINLP. In numerical terms, the CPU time of the first phase corresponded to approximately 1–10% of the imposed computational time limit. Although the computational cost of the first phase increases with problem size, the cost of solving the full MINLP increases as well.

CONCLUSIONS

This work investigated a two-phase strategy for limiting the effects of nonconvexities in MINLP reformulations of nonconvex generalized disjunctive programs. The approach combines a preliminary phase, in which a sequence of relaxed MINLPs is solved to improve initial points, with a subsequent solution of the original MINLP formulation. The strategy was assessed on the Crescent benchmark problems of increasing size using two general-purpose MINLP solvers, DICOPT and SBB, and evaluated through statistical measures of objective values, feasibility, and constraint violations.

The numerical results demonstrate that the proposed approach improves solver performance. Across all problem sizes, the use of the full algorithm leads to improved objective function values and reduced variability, as evidenced by tighter distributions in the box-plot analysis. In addition, the percentage of feasible initial points is consistently higher when the first phase is applied, highlighting its effectiveness in steering solutions toward disjunctive feasible regions. Complementary to this measure, the analysis of average constraint violation shows that the proposed strategy not only increases feasibility but also reduces the severity of constraint violations when infeasibility persists.

While the improvements are problem-dependent, the observed trends are consistent across solvers and problem sizes, suggesting that the proposed approach enhances numerical robustness without relying on problem-specific tuning. The results indicate that systematic initialization procedures can play a significant role in mitigating the challenges associated with nonconvex MINLP reformulations. Future work will focus on extending the approach to more general GDP structures and investigating its interaction with global optimization techniques.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of the Slovenian Research and Innovation Agency ARIS (programme P2-0414 and project J2-60044).

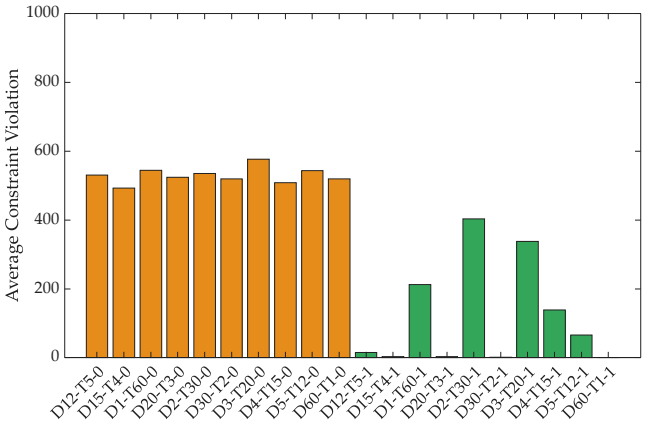

Figure 8. Average constraint violation (Crescent 60).

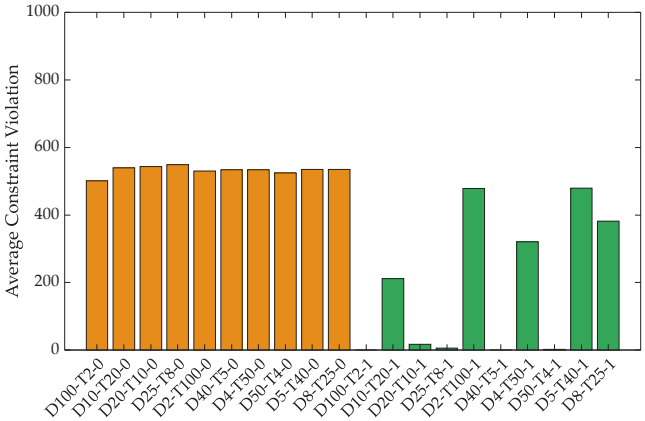

Figure 9. Average constraint violation (Crescent 200).

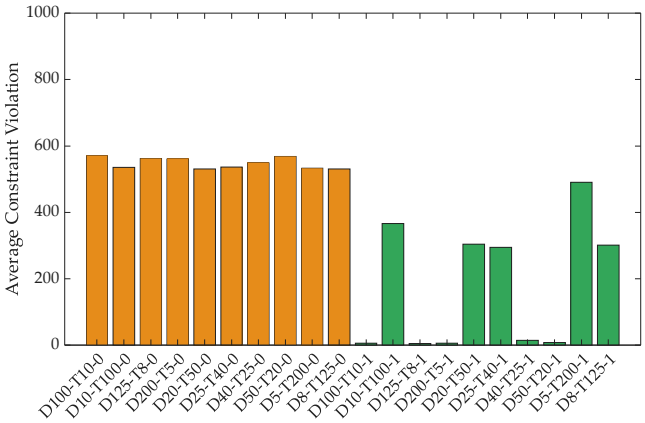

Figure 10. Average constraint violation (Crescent 1000).

REFERENCES

1. Zamora JM, Grossmann IE. Global optimization of heat exchanger networks. Encyclopedia of Optimization :1337-1342 (2008). https://doi.org/10.1007/978-0-387-74759-0_235
2. Kolodziej S, Castro PM, Grossmann IE. Global optimization of bilinear programs with a multiparametric disaggregation technique. J Glob

Optim 57:1039-1063 (2013). https://doi.org/10.1007/s10898-012-0022-1
3. Bongartz D, Mitsos A. Deterministic global optimization of process flowsheets in a reduced space using mccormick relaxations. J Glob Optim 69:761-796 (2017). https://doi.org/10.1007/s10898-017-0547-4
4. Montez C, Sanjeevi S, Sundar K. Global optimization algorithm for mixed-integer nonlinear programs with trigonometric functions. J Glob Optim 93:605-650 (2025). https://doi.org/10.1007/s10898-025-01547-9
5. Bogataj M, Kravanja Z. On numerical stability of relaxations of generalized disjunctive programs. *Comput-Aided Chem En* 3469–3474 (2024)
6. Lee S, Grossmann IE. New algorithms for nonlinear generalized disjunctive programming. Computers & Chemical Engineering 24:2125-2141 (2000). https://doi.org/10.1016/s0098-1354(00)00581-0
7. Sawaya NW, Grossmann IE. Computational implementation of non-linear convex hull reformulation. Computers & Chemical Engineering 31:856-866 (2007). https://doi.org/10.1016/j.compchemeng.2006.08.002

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Beyond Tennessee Eastman: Benchmarking Deep Anomaly Detection on Real-World Pilot-Scale Continuous Distillation Data

Fabian Hartung[ac†], Aparna Muraleedharan[b†], Marius Kloft[a] and Jakob Burger[b*]

[a] RPTU Kaiserslautern, Department of Machine Learning, Kaiserslautern, Germany
[b] Technical University of Munich, Laboratory for Chemical Process Engineering, Straubing, Germany
[c] BASF, Gas Treatment, Monheim am Rhein, Germany
* Corresponding Author: burger@tum.de.

ABSTRACT

Anomaly detection is essential for ensuring the safe and efficient operation of chemical plants. Although many deep-learning-based methods have been proposed in recent years, their evaluation remains largely limited to synthetic benchmarks such as the Tennessee Eastman Process (TEP) [1]. While these simulators enable controlled and reproducible comparisons, they fail to capture the noise characteristics, operational complexity, and irregular fault dynamics of real industrial plants, leaving the practical generalizability of many methods unclear. In this work, we extend our earlier ESCAPE study [2] beyond water-based systems to industrially relevant chemical processes. We analyze multivariate time-series data from two continuously operated pilot-plant scenarios at the Technical University of Munich, namely *n*-butanol/water heteroazeotropic distillation and poly(oxymethylene) ether purification, whose datasets were recently published at NeurIPS 2025 [3]. Using the open-source TimeSeAD library [4], we benchmark 30 anomaly detection methods, including 26 deep-learning-based and 4 classical approaches, under a unified preprocessing, model-selection, and evaluation pipeline. Performance is assessed using the F1-score and the area under the precision–recall curve (AUPRC). Our results show a substantial performance drop when moving from synthetic to real process data, with average scores far below those commonly reported for TEP. No single method performs consistently best across all datasets, and the ranking depends strongly on the chosen metric and process scenario. These findings highlight the limitations of synthetic benchmarks and underscore the need for more realistic industrial datasets, process-aware methods, and evaluation practices that better reflect real operating conditions.

Keywords: Machine Learning, Anomaly Detection, Continuous Distillation, Pilot Plant Data, Heteroazeotropic distillation, Tennessee Eastman Process Data

1. INTRODUCTION

Continuous distillation is the most widely used separation technique in the chemical industry, enabling steady-state operation through continuous feed input and simultaneous product withdrawal [5, 6]. Anomaly detection (AD) is therefore critical, as accidents in chemical plants—although rare—can result from leaks, fires, or explosions and may harm personnel, equipment, or the environment. Over the past two decades, numerous data-driven AD techniques have been developed, ranging from classical multivariate statistical to principal component analysis, and, more recently, unsupervised machine-learning-based approaches [7, 8].

Most existing AD methods have been developed and evaluated on simulated benchmarks such as the Tennessee Eastman Process (TEP) [1]. Despite their strong reported performance, it remains unclear to what extent models trained and validated on simulated data generalize to real industrial processes operating under

† These authors contributed equally to this work.

non-ideal and evolving conditions. Cheng et al. [7] report that only 12% of the processing methods in their analysis have been applied to real industrial data, primarily due to data confidentiality and limited public availability. This highlights a key challenge: while simulated data represent idealized scenarios with a limited set of fault cases, real industrial data are far more complex, influenced by factors such as minor operational disturbances, equipment aging, and complex chemical mixtures. Industrial process data are typically proprietary, as they contain sensitive information on instrumentation, operating conditions, process variables, and control strategies, and are therefore rarely accessible in the public domain.

To address this gap, we recently published process data from three scenarios of a continuous distillation mini-plant at the TUM Campus in Straubing [3]. This publicly available dataset [9] is used to train and evaluate AD methods. By benchmarking a broad range of anomaly detection approaches from the literature on our data, this work provides a more realistic assessment of model performance under practical industrial operating conditions.

2. RELATED WORK

The detection of anomalous behavior in industrial systems has a long tradition in process engineering. Early work framed fault diagnosis as a multivariate statistical process control problem, relying on dimension-reduction techniques such as PCA, PLS, or SIMCA to construct control charts that identify deviations from regular operation [10, 11]. Subsequent approaches explored probabilistic models, including Gaussian mixture models [12], and kernel-based one-class methods [13]. While these “shallow” models have proven effective at small to medium scales, their ability to capture the complex, nonlinear dynamics of chemical plants is inherently limited. As a result, recent advances in deep learning have been transferred to the process domain. Ruff et al. [14] introduced deep one-class classification, followed by numerous studies proposing autoencoder-, GAN-, and VAE-based architectures tailored to multivariate time series from simulators such as the TEP [15, 16].

Comprehensive surveys [17, 18] categorize AD techniques into probabilistic, distance-based, domain-based, reconstruction-based, forecasting-based, and classification-based approaches. Our previous work presented the first large-scale comparison of 27 deep AD algorithms on the TEP and demonstrated the strong performance of reconstruction-based models [15]. While these results are promising, it remains unclear whether models that perform well on synthetic benchmarks generalize to real industrial processes. The present paper addresses this gap by evaluating the same methodological spectrum on three newly released real plant datasets [9]. In practice, anomalies in real plants rarely manifest as abrupt, well-defined events; instead, they often evolve gradually, resemble normal process variability, or are partially masked by control actions.

A key challenge in AD benchmarking is choosing evaluation protocols that account for the temporal nature of anomalies. Point-wise metrics such as the F1-score and AUPRC are still most commonly reported, as they summarize precision–recall trade-offs in a single value [19, 20]. However, they ignore event duration and detection latency. Consequently, different metrics can yield substantially different conclusions about model performance. To mitigate this issue, the literature has proposed window-based scores (NAB, NAB-modified) [21], time-to-detect measures [22], and cost-sensitive utility functions [23]. In this work, we report both F1 and AUPRC in time-series-adapted variants.

The continued popularity of the TEP is driven by its realistic flowsheet, stochastic disturbances, and public availability. However, several limitations warrant caution: (a) faults are introduced as persistent step changes over entire test runs, amplifying anomaly signals and biasing results toward high recall; (b) only a single operating point is represented, whereas real plants undergo grade changes and load variations; (c) the simplified control structure lacks advanced layers such as MPC or APC; and (d) hyperparameter optimization on a small set of standard faults increases the risk of overfitting and score inflation. Recent studies have shown that even trivial baselines, such as an untrained LSTM autoencoder, can achieve F1-scores above 0.93 on TEP [24]. These limitations suggest that results obtained on synthetic benchmarks should be interpreted with care when assessing real-world applicability, underscoring the need for more rigorous and realistic benchmarks—an objective addressed by the datasets used in this study.

3. DATASETS

We use recently published steady-state multivariate time-series data [9] generated on a pilot-scale continuous distillation mini-plant with a capacity of 5 t/a located at the TUM Campus Straubing. The setup comprises two distillation columns and a decanter, instrumented with 34 sensors and actuators. Three operating scenarios of increasing industrial relevance are considered: water runs, a heteroazeotropic *n*-butanol/water system, and a reactive poly(oxymethylene) ether (OME) system. Water runs were performed on a single column, yielding approximately 520 h of steady-state operation. Experiments with the *n*-butanol/water system used two columns and a decanter, with campaigns averaging 10 h and 80 h of steady-state data. The OME system was operated in a single-column configuration (30h of data).

All process variables were recorded at a 30s sampling interval, sufficient to capture relevant steady-state

dynamics. Measurements include column and reboiler temperatures, mass flow rates, column head pressure, pump power, and differential pressure signals. A detailed description of the plant setup, process, and instrumentation diagram, sensors and actuators, control strategies, data preparation, and occurring anomalies is provided in the related work [9]. Startup and shutdown phases were removed prior to analysis, leaving only steady-state operation. During routine operation, a range of typical industrial anomalies was observed, including clogging, reboiler level instabilities, overheating, instrument malfunctions, pressure control deviations, and human operating errors. When specific anomaly types did not occur naturally, they were deliberately induced under controlled conditions to ensure a representative yet realistic range of fault behaviours, cf. [9] for details.

Anomalies were labeled following the classification commonly used in anomaly detection studies [15]. For each anomaly, metadata describing its type and cause, affected sensors, and timestamps are provided in the related work [9]. We additionally document the time of fault introduction, its first observable manifestation in the data, and the return to regular operation. Detailed dataset statistics are given in [3]; for the continuous water, butanol, and OME datasets, anomalous samples account for 20%/25%, 24%/41%, and 4%/36% of train/test time steps, respectively. Table 1 summarizes the key characteristics distinguishing our pilot-scale datasets from the TEP benchmark, namely: (1) real sensor noise levels 5–20× higher than in simulation, (2) closed-loop control that actively responds to and masks faults, and (3) gradual fault development rather than step changes.

Table 1: Process comparison of TEP data [1] vs data available from the pilot plant [9]

Aspect	TEP	Butanol Data	OME Data
Fault Behavior	Step Changes	Gradual Evolution	Very Gradual
Control Mode	Often open-loop	Closed-loop automatic	Closed-loop automatic
Sensor Noise	Roughly 0.1%	0.5-2%	0.5-2%
Time Constants	Fast	Medium 10-30 min	Slow 30-60 min

The Supporting Information compares faulty and normal runs from TEP and our plant, illustrating that anomalies in our data are more gradual and subtle. Additional examples in [9] further show that behavior appearing anomalous to a neutral observer may in fact correspond to normal and expected plant operation.

4. COMPUTATIONAL EXPERIMENTS

All experiments are conducted using **TimeSeAD**, our open-source benchmark library for time-series anomaly detection [4]. The evaluated methods were benchmarked using their implementations in the open-source TimeSeAD framework and trained by the authors on the considered datasets under a unified preprocessing, hyperparameter-selection, and evaluation protocol. The framework handles data ingestion, stratified splitting, normalization, hyperparameter selection, and metric aggregation, ensuring full reproducibility across datasets and methods. Training data are split into folds for model training and validation across different hyperparameter configurations. To ensure a fair comparison, the total runtime is limited to 24 hours per method. The best-performing configuration is then applied to compute anomaly scores on the test data. For each dataset, sensor channels with constant readings are removed, and the initial transient phase is discarded. The remaining channels are standardized to zero mean and unit variance within each dataset to stabilize optimization [25].

We benchmark 30 state-of-the-art anomaly detection methods, grouped into the following categories:

- **Reconstruction-based (8 methods):** Models are trained to reconstruct normal data accurately; anomalies are identified when reconstruction errors exceed a threshold.
- **Prediction-based (5 methods):** Temporal or probabilistic dynamics are learned to predict future values; large deviations between predictions and observations indicate anomalies.
- **GAN-based generative (3 methods):** A generator and discriminator are trained adversarially; anomalies are poorly reconstructed or classified as unlikely by the discriminator.
- **VAE-based generative (6 methods):** Models learn a latent probabilistic representation while optimizing reconstruction fidelity; anomalies correspond to low likelihood or high reconstruction error.
- **Hybrid (4 methods):** Multiple paradigms are combined to exploit complementary strengths, typically integrating representation learning with explicit anomaly scoring.
- **Shallow baselines (4 methods):** Classical, non-deep approaches based on statistical or simple machine-learning principles, included as interpretable and computationally efficient reference methods.

Evaluation Metrics

Model performance is assessed using the F1-score and the area under the precision–recall curve (AUPRC), both derived from true positives (TP), false positives (FP), false negatives (FN), and true negatives (TN). Precision measures the fraction of correctly detected anomalies among all detections, while recall quantifies the fraction of actual anomalies that are successfully identified.

$$Precision = \frac{TP}{TP+FP} \quad Recall = \frac{TP}{TP+FN}$$

The F1-score is defined as the harmonic mean of precision and recall:

$$F1 = 2.\frac{Precision.Recall}{Precision + Recall}$$

In addition, we report the AUPRC, which balances precision and recall across all detection thresholds and is particularly well suited to the highly imbalanced nature of anomaly detection datasets. However, conventional point-wise definitions of these metrics are ill-suited for time-series data, where anomalies typically occur as contiguous temporal segments rather than isolated points.

To address this, we use the time-series-aware versions of F1 and AUPRC implemented in the TimeSeAD library [26]. These formulations extend the classical TP/FP/FN definitions by evaluating the temporal overlap between predicted and ground-truth anomaly intervals instead of comparing individual time points. Each predicted segment contributes to precision and recall in proportion to its overlap with one or more true anomaly segments. Partial overlaps are weighted by their intersection length and penalized when a predicted segment covers multiple true segments, or vice versa. The resulting time-series-aware precision and recall are obtained by aggregating these weighted overlaps across all predicted and true intervals, and the F1-score and AUPRC are then computed analogously to their classical definitions.

5. RESULTS

Detailed results for all 30 evaluated methods on the three pilot-plant datasets, including F1-score and AUPRC, are provided in the supplementary material (Tables S8–S10). Table S11 reports the mean performance across all three datasets, while Table S12 lists the corresponding results on the Tennessee Eastman Process (TEP) for comparison. To give a compact overview of score distributions, Table 2 summarizes the average, minimum, and maximum F1-scores achieved per dataset.

A clear performance gap emerges when comparing simulated TEP data with real pilot-plant data. While TEP achieves high average scores (F1 = 0.859, AUPRC = 0.906), performance on the pilot-plant datasets is substantially lower across all cases (F1 = 0.532, AUPRC = 0.361). No method exceeds an F1-score of 0.718 or an AUPRC of 0.597, indicating that none of the evaluated approaches approaches the benchmark performance observed on synthetic data.

Despite this overall performance degradation, deep-learning-based methods consistently outperform shallow baselines on the real-world datasets. Multiple deep learning approaches achieve substantially higher F1-scores than the best-performing shallow baseline across all settings, reaffirming the advantage of deep representation learning even under realistic industrial conditions (Table S11 – results for AUPRC in supporting information). Among the deep-learning-based approaches, the strongest performers vary by dataset: LSTM-2S2-P achieves the best F1-score on Butanol and Water, whereas generative VAE methods, particularly LSTM-VAE and LSTM-DVAE, perform best on OME. At the method-type level, hybrid approaches achieve the highest mean F1-score across the three pilot datasets, whereas, for deep methods, reconstruction-based methods rank first in mean AUPRC, highlighting strong dataset- and metric-dependence.

Table 2: Dataset-level F1-score statistics across evaluated AD methods showing average, minimal, and maximal F1-Score, and the span (AUPRC results in appendix)

F1	average	min	max	Δ
Butanol	0.538	0.291	0.624	0.333
OME	0.629	0.215	0.718	0.503
Water	0.428	0.104	0.531	0.427
Average	0.532	0.267	0.590	0.323
TEP	0.859	0.497	0.911	0.414

Table 3: F1-scores of the best-performing deep-learning and shallow methods, and the mean across all three datasets

	n-butanol	OME	water	mean
Best Deep	0.624	**0.718**	0.531	**0.624**
Best Shallow	0.588	0.666	0.433	0.562

6. DISCUSSION

Our earlier benchmark study on the TEP dataset [11–15] showed that most deep anomaly-detection methods achieve near-perfect performance when anomalies appear as persistent fault states. In contrast, our evaluation

of real pilot-plant distillation data reveals a markedly different picture. When anomalies are transient, gradual, or intermittent, as is typical in industrial practice, the same methods achieve at best an F1-score of 0.718. This pronounced performance degradation highlights a key limitation of current approaches: their strong reliance on temporally persistent and clearly separable fault signatures, which are common in synthetic benchmarks but rare in real plants. Importantly, this performance gap cannot be attributed to label uncertainty. In our datasets, the onset and resolution times of anomalies are known, allowing us to rule out annotation noise as a confounding factor. Instead, the observed differences directly reflect model behavior under realistic operating conditions.

Across the three datasets and both evaluation metrics, no single method or method family consistently dominates. However, across all datasets, several deep-learning-based methods outperform all shallow baselines in terms of F1-score, and in two out of three datasets also with respect to AUPRC. While GenAD [27] and AnomalyTransfer [28] achieve comparatively low average F1-scores, they exhibit stronger AUPRC performance on selected datasets. Hybrid approaches show a slight average advantage but also pronounced dataset-specific weaknesses, underscoring the lack of robust generalization across fault types and evaluation metrics.

The root cause of this discrepancy lies in the development of industrial faults. Real process anomalies typically evolve gradually due to fouling, clogging, drift, or degradation rather than abrupt step changes. The fault scenarios in our datasets [3, 9] were explicitly designed to reflect this behavior. Gradual faults are inherently more difficult to detect: early deviations often remain within normal variability, fault signatures evolve slowly, and closed-loop control systems may partially compensate disturbances, thereby masking their effects. Controllers suppress deviations in controlled variables while redistributing the effects of anomalies into manipulated variables, such as flow rates or reboiler duties. Consequently, faults may be physically present but only indirectly observable through control actions, in stark contrast to TEP, where faults are introduced as step disturbances with clear onsets and pronounced responses.

These differences lead to a systematic performance gap, with F1-scores of approximately 0.80–0.90 on TEP compared to 0.55–0.72 on real plant data. Crucially, this gap does not primarily indicate insufficient model complexity or flawed algorithmic design. Instead, it reflects fundamental process-engineering challenges: (i) gradual fault evolution, (ii) active feedback control that masks faults in controlled variables, and (iii) sensor noise levels that are 5–20 times higher than in synthetic benchmarks.

From a deployment perspective, reduced detection performance on real industrial data should therefore be expected rather than interpreted as model failure. Anomalies often become observable first through changes in control actions rather than deviations in process outputs. Effective anomaly-detection systems must explicitly incorporate manipulated variables and control behavior and align detection window sizes with process time scales. For continuous distillation, fault dynamics on the order of 10–60 minutes imply an unavoidable trade-off between detection reliability and alarm latency.

Robust methods must therefore demonstrate stable performance across diverse datasets rather than excelling only on highly structured synthetic benchmarks. Deep-learning-based methods not only outperform shallow approaches under realistic conditions but also offer greater potential for further methodological advances, making them a promising path toward closing the current performance gap between synthetic and real-world benchmarks.

Finally, this work provides the first publicly available pilot-scale distillation benchmark featuring realistic faults, industrial noise characteristics, and closed-loop control interactions. We hope it fosters the development of process-aware anomaly-detection methods that explicitly account for plant physics and control behavior and are validated under realistic industrial conditions, rather than treating chemical plants as generic time-series systems.

7. LIMITATIONS AND FUTURE WORK

While this benchmark focuses on distillation-related process data, several limitations should be acknowledged. Although the datasets and benchmarks presented in this work provide a more realistic assessment of anomaly detection methods than commonly used synthetic benchmarks, they cannot exhaustively represent all anomalies encountered in industrial distillation processes. Rare or safety-critical events, such as severe equipment failures or cascading faults, are inherently difficult or unsafe to reproduce in a pilot-scale plant. Future data collection campaigns could therefore expand the diversity of fault types, include additional separation processes, and capture longer-term degradation phenomena as well as more subtle control-related anomalies. In addition to collecting further real-world data, data augmentation and generation offer substantial potential to mitigate current data limitations [58].

Second, the benchmarking framework relies on the current capabilities of the TimeSeAD library. While TimeSeAD enables reproducible preprocessing, training, and evaluation across a wide range of anomaly detection methods, it does not yet cover all approaches or evaluation protocols proposed in the literature. Future work will extend TimeSeAD with additional models, improved hyperparameter optimization strategies, and alternative evaluation metrics that more explicitly account for

temporal aspects such as detection delay and fault duration. Given the still considerable development potential of deep-learning-based methods, compared to the largely saturated performance of shallow approaches, the integration of transformer-based models, in particular the recently proposed NeuTraL AD method [56], into the library and our benchmark datasets is of high interest.

Third, while F1-score and AUPRC provide complementary perspectives on model performance, they can lead to different rankings and interpretations. To achieve a more comprehensive and task-specific evaluation, future work should investigate additional metrics tailored to anomaly detection, such as the recently proposed ALARM metric [57], which may improve comparability across datasets and operating conditions.

Finally, although the presented pilot-scale and industrial datasets capture realistic process dynamics, noise characteristics, and operational variability, they cannot be assumed to be universally representative of all large-scale industrial systems. Differences in plant design, control strategies, sensor quality, and operating regimes may substantially influence anomaly characteristics and model performance. Consequently, future studies should incorporate data from additional plants and process types to further assess the robustness and generalizability of anomaly detection methods across industrial settings.

8. SUMMARY OF CONTRIBUTIONS

The main contributions of this work are:

- We benchmark 30 anomaly detection methods, including deep-learning-based and classical approaches, on steady-state multivariate time-series data from pilot-scale and industrial continuous distillation processes.
- We show that methods achieving strong performance on synthetic benchmarks, such as the Tennessee Eastman Process, experience a pronounced drop in performance when evaluated on real process data.
- We demonstrate that no single method consistently performs well across datasets and operating conditions, while several deep-learning-based methods outperform shallow baselines under realistic process conditions.
- We highlight the strong dependence of benchmarking outcomes on the choice of evaluation metrics and emphasize the need for time-series–aware evaluation in industrial anomaly detection.

DIGITAL SUPPLEMENTARY MATERIAL

Comparison of TEP data vs data from our pilot plant. Details on the performance of the different AD methods on the Butanol, OME, Water, and TEP datasets. Mean performance of the AD Methods. Supporting Information available at:
https://psecommunity.org/LAPSE:2026.0038

ACKNOWLEDGEMENTS

This work was funded by the German Research Foundation (DFG) and conducted in collaboration with the FOR5359 Research Unit on Deep Learning on Sparse Chemical Process Data.

AUTHOR IDENTIFIERS

Author ORCIDs:
Hartung: 0009-0004-5093-3287
Muraleedharan: 0009-0006-7232-1550
Kloft: 0000-0001-6829-3725
Burger: 0000-0002-9376-9438

REFERENCES

1. Downs JJ, Vogel EF. A plant-wide industrial process control problem. Computers & Chemical Engineering 17:245-255 (1993). https://doi.org/10.1016/0098-1354(93)80018-i
2. Muraleedharan A, et al. Benchmarking deep anomaly detection on real process data of a continuous distillation process. ESCAPE-34 / PSE Conf (2024).
3. Wagner D, et al. NoBOOM: Chemical process datasets for industrial anomaly detection. Proc NeurIPS Datasets Benchmarks Track (2025). https://openreview.net/forum?id=qiLboR0ocm
4. Wagner D, et al. TimeSeAD: Benchmarking deep multivariate time-series anomaly detection. Trans Mach Learn Res (2023). https://openreview.net/forum?id=iMmsCI0JsS
5. Liu J, Ren J, Yang Y, Liu X, Sun L. Effective semicontinuous distillation design for separating normal alkanes via multi-objective optimization and control. Chemical Engineering Research and Design 168:340-356 (2021). https://doi.org/10.1016/j.cherd.2021.02.018
6. Safrit BT, Westerberg AW, Diwekar U, Wahnschafft OM. Extending continuous conventional and extractive distillation feasibility insights to batch distillation. Ind. Eng. Chem. Res. 34:3257-3264 (2002). https://doi.org/10.1021/ie00037a012
7. Ji C, Sun W. A review on data-driven process monitoring methods: characterization and mining of industrial data. Processes 10:335 (2022). https://doi.org/10.3390/pr10020335

8. Siegel B. Industrial anomaly detection: a comparison of unsupervised neural network architectures. IEEE Sens. Lett. 4:1-4 (2020). https://doi.org/10.1109/lsens.2020.3007880
9. Muraleedharan A, et al. Experimental time series data with and without anomalies from a continuous distillation mini-plant for development of machine learning anomaly detection methods. engrXiv (2025). https://doi.org/10.31224/5631
10. Ku W, Storer RH, Georgakis C. Disturbance detection and isolation by dynamic principal component analysis. Chemometrics and Intelligent Laboratory Systems 30:179-196 (1995). https://doi.org/10.1016/0169-7439(95)00076-3
11. Kresta JV, Macgregor JF, Marlin TE. Multivariate statistical monitoring of process operating performance. Can J Chem Eng 69:35-47 (2009). https://doi.org/10.1002/cjce.5450690105
12. Yin S, Ding SX, Xie X, Luo H. A review on basic data-driven approaches for industrial process monitoring. IEEE Trans. Ind. Electron. 61:6418-6428 (2014). https://doi.org/10.1109/tie.2014.2301773
13. Chiang LH, Russell EL, Braatz RD. Fault diagnosis in chemical processes using fisher discriminant analysis, discriminant partial least squares, and principal component analysis. Chemometrics and Intelligent Laboratory Systems 50:243-252 (2000). https://doi.org/10.1016/s0169-7439(99)00061-1
14. Ruff L, et al. Deep one-class classification. Proc Int Conf Mach Learn 80:4393–4402 (2018).
15. Hartung F, Franks BJ, Michels T, Wagner D, Liznerski P, Reithermann S, Fellenz S, Jirasek F, Rudolph M, Neider D, Leitte H, Song C, Kloepper B, Mandt S, Bortz M, Burger J, Hasse H, Kloft M. Deep anomaly detection on tennessee eastman process data. Chemie Ingenieur Technik 95:1077-1082 (2023). https://doi.org/10.1002/cite.202200238
16. Li X, Wang J, Qin SJ. Temporal convolutional autoencoder for fault detection: A case study on the Tennessee Eastman process. Comput Chem Eng 165:107998 (2022).
17. Chandola V, Banerjee A, Kumar V. Anomaly detection. ACM Comput. Surv. 41:1-58 (2009). https://doi.org/10.1145/1541880.1541882
18. Ruff L, Kauffmann JR, Vandermeulen RA, Montavon G, Samek W, Kloft M, Dietterich TG, Muller KR. A unifying review of deep and shallow anomaly detection. Proc. IEEE 109:756-795 (2021). https://doi.org/10.1109/jproc.2021.3052449
19. Tatbul N, et al. Precision and recall for time series. Adv Neural Inf Process Syst 31:1920–1930 (2018).
20. Hanselmann M, Strauss T, Dormann K, Ulmer H. Canet: an unsupervised intrusion detection system for high dimensional CAN bus data. IEEE Access 8:58194-58205 (2020). https://doi.org/10.1109/access.2020.2982544
21. Lavin A, Ahmad S. Evaluating real-time anomaly detection algorithms -- the numenta anomaly benchmark. 2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA) :38-44 (2015). https://doi.org/10.1109/icmla.2015.141
22. Görnitz N, et al. Toward supervised anomaly detection. J Artif Intell Res 46:235–262 (2013). https://doi.org/10.1613/jair.3965
23. Siffer A, Fouque PA, Termier A, Largouet C. Anomaly detection in streams with extreme value theory. Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining :1067-1075 (2017). https://doi.org/10.1145/3097983.3098144
24. Fährmann D, Damer N, Kirchbuchner F, Kuijper A. Lightweight long short-term memory variational auto-encoder for multivariate time series anomaly detection in industrial control systems. Sensors 22:2886 (2022). https://doi.org/10.3390/s22082886
25. Heaton J. Ian goodfellow, yoshua bengio, and aaron courville: deep learning. Genet Program Evolvable Mach 19:305-307 (2017). https://doi.org/10.1007/s10710-017-9314-z
26. Hua X, Zhu L, Zhang S, Li Z, Wang S, Deng C, Feng J, Zhang Z, Wu W. Genad: general unsupervised anomaly detection using multivariate time series for large?scale wireless base stations. Electronics Letters 59: (2022). https://doi.org/10.1049/ell2.12683
27. Xu J, et al. Anomaly transformer: Time series anomaly detection with association discrepancy. arXiv 2110.02642 (2022).
28. Abdi H, Williams LJ. Principal component analysis. WIREs Computational Stats 2:433-459 (2010). https://doi.org/10.1002/wics.101
29. Kumar A, Tripathi AR, Satapathy SC, Zhang YD. Sars-net: COVID-19 detection from chest x-rays by combining graph convolutional network and convolutional neural network. Pattern Recognition 122:108255 (2022). https://doi.org/10.1016/j.patcog.2021.108255
30. Tatbul N, et al. Precision and recall for time series. Adv Neural Inf Process Syst 31:1920–1930 (2018).
31. Hanselmann M, Strauss T, Dormann K, Ulmer H. Canet: an unsupervised intrusion detection system for high dimensional CAN bus data. IEEE Access 8:58194-58205 (2020). https://doi.org/10.1109/access.2020.2982544
32. Lavin A, Ahmad S. Evaluating real-time anomaly detection algorithms -- the numenta anomaly benchmark. 2015 IEEE 14th International Conference on Machine Learning and Applications

(ICMLA) :38-44 (2015). https://doi.org/10.1109/icmla.2015.141
33. Görnitz N, et al. Toward supervised anomaly detection. J Artif Intell Res 46:235–262 (2013). https://doi.org/10.1613/jair.3965
34. Siffer A, Fouque PA, Termier A, Largouet C. Anomaly detection in streams with extreme value theory. Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining :1067-1075 (2017). https://doi.org/10.1145/3097983.3098144
35. Fährmann D, Damer N, Kirchbuchner F, Kuijper A. Lightweight long short-term memory variational auto-encoder for multivariate time series anomaly detection in industrial control systems. Sensors 22:2886 (2022). https://doi.org/10.3390/s22082886
36. Thill M, Konen W, Bäck T. Time series encodings with temporal convolutional networks. Lecture Notes in Computer Science :161-173 (2020). https://doi.org/10.1007/978-3-030-63710-1_13
37. Hinton GE, Salakhutdinov RR. Reducing the dimensionality of data with neural networks. Science 313:504-507 (2006). https://doi.org/10.1126/science.1127647
38. Malhotra P, et al. LSTM-based encoder-decoder for multi-sensor anomaly detection. arXiv 1607.00148 (2016). https://doi.org/10.48550/arXiv.1607.00148
39. Malhotra P, et al. Long short-term memory networks for anomaly detection in time series. Proc ESANN :89 (2015).
40. Xu H, Feng Y, Chen J, Wang Z, Qiao H, Chen W, Zhao N, Li Z, Bu J, Li Z, Liu Y, Zhao Y, Pei D. Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications. Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18 :187-196 (2018). https://doi.org/10.1145/3178876.3185996
41. Sölch M, et al. Variational inference for on-line anomaly detection in high-dimensional time series. Stat 1050:23 (2016). https://doi.org/10.48550/arXiv.1602.07109
42. Su Y, Zhao Y, Niu C, Liu R, Sun W, Pei D. Robust anomaly detection for multivariate time series through stochastic recurrent neural network. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining :2828-2837 (2019). https://doi.org/10.1145/3292500.3330672
43. Li L, Yan J, Wang H, Jin Y. Anomaly detection of time series with smoothness-inducing sequential variational auto-encoder. IEEE Trans. Neural Netw. Learning Syst. 32:1177-1191 (2021). https://doi.org/10.1109/tnnls.2020.2980749
44. Park D, Hoshi Y, Kemp CC. A multimodal anomaly detector for robot-assisted feeding using an lstm-based variational autoencoder. IEEE Robot. Autom. Lett. 3:1544-1551 (2018). https://doi.org/10.1109/lra.2018.2801475
45. Audibert J, Michiardi P, Guyard F, Marti S, Zuluaga MA. USAD. Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining :3395-3404 (2020). https://doi.org/10.1145/3394486.3403392
46. Guo Y, et al. Multidimensional time series anomaly detection: A GRU-based Gaussian mixture variational autoencoder approach. Proc Asian Conf Mach Learn :97–112 (2018).
47. He Y, Zhao J. Temporal convolutional networks for anomaly detection in time series. J. Phys.: Conf. Ser. 1213:042050 (2019). https://doi.org/10.1088/1742-6596/1213/4/042050
48. Mirza AH, Cosan S. Computer network intrusion detection using sequential LSTM neural networks autoencoders. 2018 26th Signal Processing and Communications Applications Conference (SIU) :1-4 (2018). https://doi.org/10.1109/siu.2018.8404689
49. Said Elsayed M, Le-Khac NA, Dev S, Jurcut AD. Network anomaly detection using LSTM based autoencoder. Proceedings of the 16th ACM Symposium on QoS and Security for Wireless and Mobile Networks :37-45 (2020). https://doi.org/10.1145/3416013.3426457
50. Geiger A, Liu D, Alnegheimish S, Cuesta-Infante A, Veeramachaneni K. Tadgan: time series anomaly detection using generative adversarial networks. 2020 IEEE International Conference on Big Data (Big Data) :33-43 (2020). https://doi.org/10.1109/bigdata50022.2020.9378139
51. Zhan J, Wang S, Ma X, Wu C, Yang C, Zeng D, Wang S. Stgat-mad : spatial-temporal graph attention network for multivariate time series anomaly detection. ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) :3568-3572 (2022). https://doi.org/10.1109/icassp43922.2022.9747274
52. Li D, Chen D, Jin B, Shi L, Goh J, Ng SK. MAD-GAN: multivariate anomaly detection for time series data with generative adversarial networks. Lecture Notes in Computer Science :703-716 (2019). https://doi.org/10.1007/978-3-030-30490-4_56
53. Zhao H, Wang Y, Duan J, Huang C, Cao D, Tong Y, Xu B, Bai J, Tong J, Zhang Q. Multivariate time-series anomaly detection via graph attention network. 2020 IEEE International Conference on Data Mining (ICDM) :841-850 (2020). https://doi.org/10.1109/icdm50108.2020.00093

54. Munir M, Siddiqui SA, Dengel A, Ahmed S. Deepant: a deep learning approach for unsupervised anomaly detection in time series. IEEE Access 7:1991-2005 (2019). https://doi.org/10.1109/access.2018.2886457
55. Deng A, Hooi B. Graph neural network-based anomaly detection in multivariate time series. AAAI 35:4027-4035 (2021). https://doi.org/10.1609/aaai.v35i5.16523
56. Qiu C, et al. Self-supervised anomaly detection with neural transformation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021). https://doi.org/10.1109/TPAMI.2024.3519543
57. Wagner, D. et al. Formally Exploring Time-Series Anomaly Detection Evaluation Metrics. In Proceedings of the AISTATS (to appear in 2026).
58. Manduchi, L., et al. (2024). On the Challenges and Opportunities in Generative AI. arXiv. https://doi.org/10.48550/arXiv.2403.00025

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

From P&ID Drawings to Process Graphs: A Multimodal Language Model Approach

Baikai Zhu[a], Samuel Duong[a], Javal Vyas[a*], Mehmet Mercangöz[a]

[a] Imperial College London, Department of Chemical Engineering, London, Greater London, United Kingdom
* Corresponding Author: j.vyas24@imperial.ac.uk

ABSTRACT

Piping and instrumentation diagrams (P&IDs) encode the functional structure of process plants and are a critical yet underutilised source of engineering knowledge for digital twins and intelligent decision support. However, digitising legacy P&IDs remains challenging due to heterogeneous drawing standards and the reliance of existing methods on brittle symbol recognition and rule-based connectivity reconstruction. This work reframes P&ID digitization as the extraction of equipment tags and inference of process topology, rather than graphical reproduction. We propose a two-stage workflow based on multimodal large language models, in which visual extraction and topology reconstruction are treated as distinct reasoning stages guided by chemical engineering process knowledge. The approach is evaluated on two ANSI-standard P&ID case studies of increasing complexity. Results show that decomposing visual extraction and topology reasoning yields more accurate and structurally consistent process representations than end-to-end digitization, highlighting the potential of language-model-based, knowledge-guided workflows for scalable and semantically reliable P&ID digitization.

Keywords: P&ID digitisation, Multimodal large language models, Graph reconstruction

INTRODUCTION

Piping and instrumentation diagrams (P&IDs) constitute the structural backbone of a process plant, capturing how equipment, instrumentation, and control elements are interconnected to realise the intended process operation. Beyond providing spatial or locational information, P&IDs encode the functional logic of the plant, describing how units interact, how material and signals propagate, and how control strategies are implemented. As such, they represent an authoritative source of process understanding that spans design intent, operational constraints, and safety considerations.

Despite their central role, the information embedded in P&IDs remains largely underexploited in industrial practice. In many plants, P&IDs continue to be treated primarily as static reference documents rather than as structured, machine-interpretable sources of knowledge. This limits their integration into advanced engineering workflows, where explicit representations of process topology could support automated consistency checking, decision support, and interaction with digital twins. As the process industries move toward smart factories and data-driven operation, the ability to interrogate and reason over P&IDs becomes increasingly important.

A major barrier to such “smart” use of P&IDs is the legacy nature of existing documentation. A substantial fraction of industrial P&IDs still exists as paper drawings or unstructured raster images, often produced over decades using heterogeneous drafting standards and conventions. Converting these legacy artefacts into structured, machine-readable representations remains labour-intensive and typically requires extensive manual interpretation by domain experts. This manual effort does not scale and represents a significant obstacle to unlocking the full value of P&IDs within modern digital infrastructures, including retrieval-augmented generation systems and interactive engineering tools.

From a technical standpoint, P&ID digitization is challenging due to variability in symbol rendering, annotation practices, and drawing quality across companies and engineering contractors. While human engineers resolve such variability using contextual and domain knowledge, automated systems must operate on

inconsistent and ambiguous graphical representations. Conventional digitization pipelines therefore rely on deep learning models for symbol and text recognition combined with rule-based methods to reconstruct process connectivity [1–3]. Although effective in constrained settings, these approaches exhibit limited generalisation across drawing standards [1, 5], require extensive manual engineering of reconstruction rules [4], and propagate errors from early recognition stages into the reconstructed topology. Consequently, many existing methods emphasise graphical detection while falling short of producing semantically reliable representations of process structure.

Recent advances in vision-language models (VLMs) and multimodal large language models (MLLMs) provide an opportunity to reconsider this paradigm. By jointly grounding visual and semantic information, such models can reduce dependence on specific symbol appearances and bespoke training for individual drawing standards. When combined with language-based reasoning, they further enable inference over extracted engineering artefacts in a manner that more closely resembles expert interpretation of P&IDs.

In this work, we investigate the use of multimodal language models and multi-agent workflows to transform legacy P&IDs into structured, semantically meaningful representations of process topology. By decoupling visual extraction from topology reasoning, the proposed approach aims to reduce manual digitization effort and enable P&IDs to function as active, machine-interpretable sources of process knowledge, supporting intelligent interaction and decision-making in future smart manufacturing environments.

LITERATURE REVIEW

Existing research on P&ID digitization has primarily focused on automating the extraction of symbols, text, and connectivity from legacy diagrams using computer vision and machine learning techniques. The following review summarises the main methodological directions, highlighting their strengths and limitations with respect to robustness, scalability, and semantic correctness.

Classical image processing and rule-based pipelines

Early approaches to P&ID digitization predominantly relied on classical image processing, optical character recognition (OCR), and engineering heuristics. These systems typically detect graphical primitives such as symbols, lines, and text using template matching and rule-based algorithms, followed by proximity-based association to identify equipment tags and infer connectivity. Kang et al. [6] demonstrated the feasibility of converting raster P&IDs into structured databases using such techniques, while maintaining compatibility with downstream engineering tools. However, these methods are highly sensitive to drawing quality, symbol variability, and vendor-specific drafting conventions, which limits their robustness when applied to heterogeneous industrial datasets [7].

Deep learning for visual recognition

To address the limitations of purely rule-based methods, subsequent research has integrated deep learning for visual recognition tasks. Convolutional neural networks and modern text detection architectures have been applied to identify equipment symbols, instrumentation, and tag text in P&IDs [3]. Industrially oriented pipelines combine deep learning–based detection of symbols, text, and lines with heuristic association rules to reconstruct connectivity and export structured representations aligned with engineering information models [8]. While these approaches significantly improve robustness in component recognition, topology reconstruction remains challenging, as rule-based connectivity inference is sensitive to broken lines, ambiguous junctions, and overlapping annotations commonly found in real plant drawings.

Graph-based formulations of topology reconstruction

Recognising that the primary engineering value of a P&ID lies in its process topology, several studies explicitly model digitization as a graph reconstruction problem. In this formulation, equipment and instruments form graph nodes, while piping and signal connections define edges. Mani et al. (2020) [9] combine deep learning–based detection with graph search techniques to traverse line structures and infer connectivity. Similarly, Digitize-PID introduces a structured pipeline that incorporates domain-based validation to correct implausible connections and improve graph consistency [1]. More recently, relation-based and transformer-style models have been proposed to directly infer connectivity from diagram images, framing P&ID understanding as an image-to-graph task [9]. Although these methods show improved structural reasoning, they typically require specialised training data with explicit edge annotations and may struggle to generalise across diverse drafting styles.

Data scarcity and generalisation challenges

A practical challenge in P&ID digitization is the scarcity of annotated datasets, as industrial drawings are often proprietary. To mitigate this, recent work explores semi-supervised and weakly supervised learning strategies. Gupta et al. (2024) [11] propose a semi-supervised symbol detection approach that reduces annotation effort while improving generalisation across symbol variants. Complementary methods exploit auxiliary

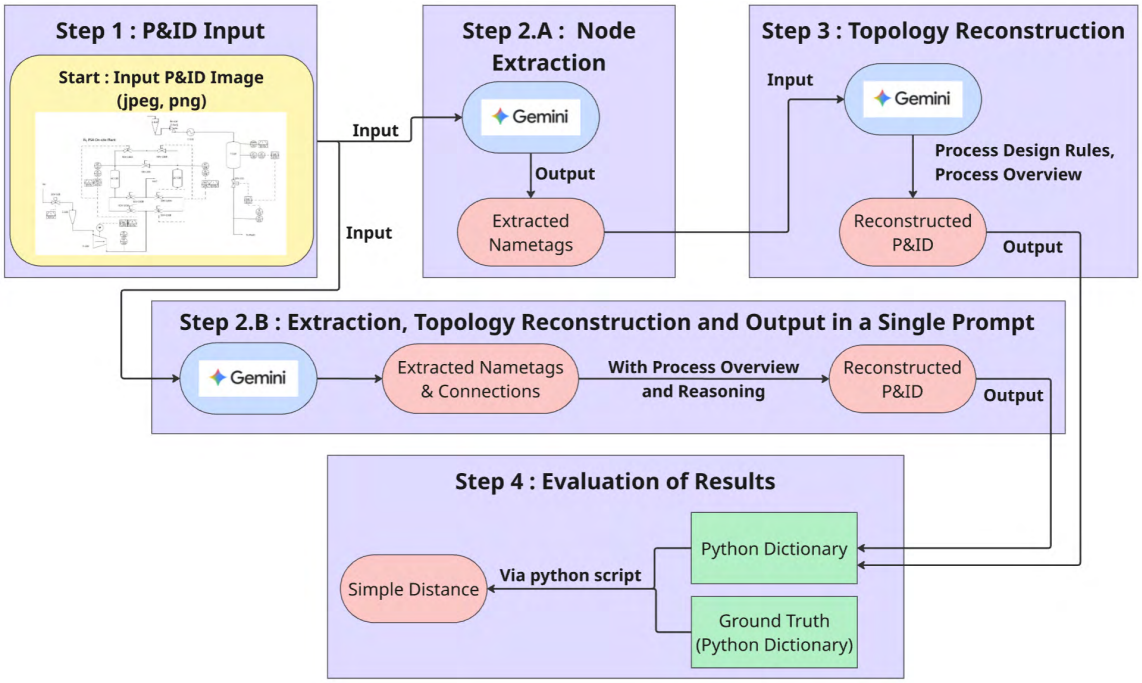

Figure 1. Overview of the proposed P&ID digitisation workflows

engineering artefacts such as legend sheets to infer symbol semantics with minimal manual labelling [12]. While these approaches improve robustness in symbol recognition, they do not directly address the inference of process connectivity under uncertainty.

Engineering knowledge and validation

Across the literature, it is widely acknowledged that purely visual or data-driven methods are insufficient to guarantee engineering correctness. Equipment type, tag nomenclature, and process conventions impose strong constraints on allowable connectivity—for example, pumps require suction and discharge connections, vessels connect via nozzles, and instruments interface through signal rather than process lines. Several studies therefore incorporate post-processing validation stages based on engineering rules to filter implausible structures and enforce consistency with expected process behaviour [8, 9]. However, such rules are typically applied locally and procedurally, limiting their ability to resolve global ambiguities in complex diagrams.

Language-based reasoning for P&ID understanding

Recent advances in large language models (LLMs) and vision-language models (VLMs) have motivated exploratory studies on their application to engineering diagrams. Initial investigations indicate that direct application of LLMs to P&ID digitization without grounding or constraint enforcement leads to inconsistent or non-physical interpretations [13]. Nevertheless, these models demonstrate strong capabilities in reasoning over structured engineering information when provided with explicit inputs and constraints, suggesting their potential role in topology inference when combined with reliable visual extraction and engineering knowledge.

METHODOLOGY

Problem formulation

We adopt a process-centric formulation of the P&ID digitization problem. Rather than aiming to reproduce graphical details of the diagram, the primary objective is to extract (i) identifiable process units and instrumentation, represented by their tags, and (ii) the interconnections between these units that define the process topology. Under this formulation, a digitised P&ID is represented as a graph, where nodes correspond to equipment or instrumentation tags and edges represent material, signal, or control connections between them.

Conventional digitization pipelines typically address this task through a combination of deep learning–based visual recognition and rule-based connectivity reconstruction. In contrast, we investigate whether modern multimodal large language models (MLLMs), combined with language-based reasoning, can directly infer structured process representations from P&ID images, and whether decomposing the task into extraction and reasoning stages improves robustness and accuracy.

Overview of evaluation pathways

We assess language-model-based P&ID digitization using two complementary workflows, illustrated schematically in Figure 1. The first workflow performs extraction and topology reconstruction in a single step using an MLLM. The second workflow decomposes the task into two stages—node extraction and topology reconstruction—handled by separate language model instances, with the reconstruction stage augmented by explicit chemical engineering process knowledge. In both cases, the final output is a structured graph representation of the P&ID, which is evaluated against a manually constructed ground truth.

Pathway I: End-to-end digitization with an MLLM

In the first workflow (Figure 1, Step 2.B), the P&ID image is divided into eight spatial slices to mitigate resolution and context-length limitations of the multimodal model. These image slices, together with a short high-level textual description of the diagram, are provided as input to an MLLM. The model is prompted to digitise the P&ID directly into a structured representation expressed as a Python dictionary, encoding process units as nodes and their interconnections as edges.

To assess the quality of the generated representation, a ground-truth graph of the same P&ID is manually constructed using the same dictionary format. The predicted and ground-truth graphs are compared using three metrics. Node identification accuracy measures the proportion of correctly identified tags, while edge recognition accuracy measures the proportion of correctly inferred connections. In addition, overall structural similarity is quantified using a simple distance metric. The simple distance is defined as the total number of node- and edge-level discrepancies between the predicted and ground-truth graphs, including missing nodes, extra nodes, missing edges, extra edges, and edge-type mismatches.

This workflow evaluates the feasibility of treating P&ID digitization as a single-step multimodal reasoning task.

Pathway II: Decomposed extraction and topology reconstruction

The second workflow (Figure 1, Steps 2.A and 3) explicitly separates visual extraction from topology inference, reflecting the hypothesis that reasoning over process structure benefits from abstraction away from raw visual input.

Node extraction

In the extraction stage (Figure 1, Step 2.A), an MLLM is provided only with the P&ID image slices and tasked with identifying the set of process units and instrumentation present in the diagram, essentially functioning as an advanced Optical Character Recognition (OCR) tool. The model outputs a list of equipment and instrument tags without attempting to infer connectivity, thereby isolating the visual recognition task.

Topology reconstruction with engineering knowledge

In the reconstruction stage (Figure 1, Step 3), the output of the extraction step is passed to a second language model. This model has no access to the original P&ID images and relies exclusively on symbolic inputs. In addition to the extracted node list, the model is provided with (i) a high-level process overview and (ii) a document describing general chemical engineering process design rules. It is important to note that the inclusion of such chemical engineering context does not involve fine-tuning or training the underlying model. Instead, this study utilises an off-the-shelf MLLM (Gemini-3.0-Pro), where domain-specific knowledge is introduced entirely via prompt engineering, supplied to the language model in the form of an accompanying text document.

The process design rules are organised into three categories. The first category defines common tag nomenclature and abbreviations for equipment, valves, and instrumentation, capturing typical conventions used to encode process streams, destinations, and unit sequencing. The second category describes common physical connectivity patterns between process units, reflecting standard design and maintenance practices. The third category addresses control and instrumentation structure, guiding the model in inferring control loops and signal hierarchies, including the directionality of signals between sensors, controllers, and actuators.

Using these inputs, the reconstruction model infers the connectivity between nodes and outputs a structured graph representation in the same format as Workflow I.

Evaluation Protocol

For both workflows, the reconstructed P&ID graph is evaluated against the same ground-truth representation using identical metrics, as shown in Figure 1 (Step 4). This enables a direct comparison between end-to-end multimodal digitization and the decomposed extraction–reasoning approach, isolating the effect of task decomposition and explicit process knowledge on topology reconstruction accuracy.

CASE STUDY

To evaluate the proposed P&ID digitization workflows, two representative process diagrams were selected as case studies. Both P&IDs were designed in accordance with the ANSI standard and use consistent naming conventions for equipment and instrumentation

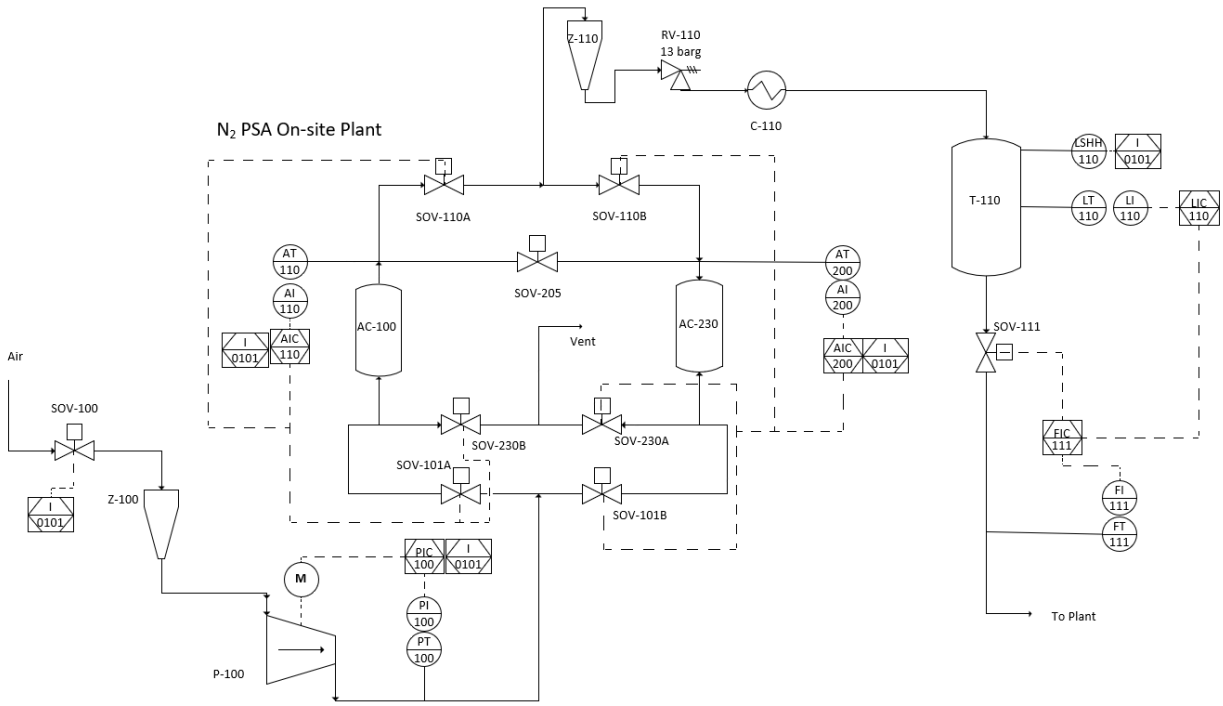

Figure 2. Nitrogen Pressure Swing Adsorption P&ID

to minimise variability arising from tag nomenclature rather than diagram structure.

The first case study, shown in **Figure 2**, corresponds to a nitrogen pressure-swing adsorption (PSA) process. The second case study, shown in **Figure 3**, represents a wet flue gas desulfurization (FGD) process. These two systems were chosen to reflect different levels of process and topological complexity while remaining representative of industrial process flows commonly encountered in practice.

P&ID complexity

We quantified the complexity of each P&ID by comparing the total number of nodes and edges, where a higher number of nodes and edges represents a more complex case study, as shown in Table [1].

Table 1. P&ID Complexity Metrics

P&ID	Complexity
Nitrogen Pressure Swing Adsorption	Total nodes: 38 Total edges: 53 Total combined: 91
Wet Flue Gas Desulfurization	Total nodes: 73 Total edges: 101 Total combined: 174

Experimental setup

At the start of each case study, the P&ID image was divided into eight equal-resolution slices to mitigate resolution and context-length limitations of the multimodal language models. In addition, a short, high-level textual description of the process was prepared for each P&ID to provide contextual information where applicable.

For the end-to-end workflow, the P&ID slices and process overview were provided to Gemini-3.0-pro in a single prompt. The model was instructed to output a Python dictionary representing the P&ID, with equipment and instrumentation tags as nodes and their interconnections as edges.

For the decomposed workflow, digitization was split into two stages using a multi-agent setup implemented in Dify, an agentic workflow framework. In the first stage, the language model was provided only with the P&ID image slices and tasked with extracting equipment and instrumentation tags without inferring connectivity. The model is prompted to solely extract equipment tags and numbering while ignoring the graphical symbols entirely. In the second stage, the extracted node list was passed to a separate instance of the language model, together with the process overview and a document describing general chemical engineering process design rules. This second instance inferred the interconnections between nodes using language-based reasoning on the linguistic meaning of the extracted tags, guided by chemical engineering domain knowledge regarding standard configurations of process design. Since no graphical symbols are used as an input to this reconstruction stage, the

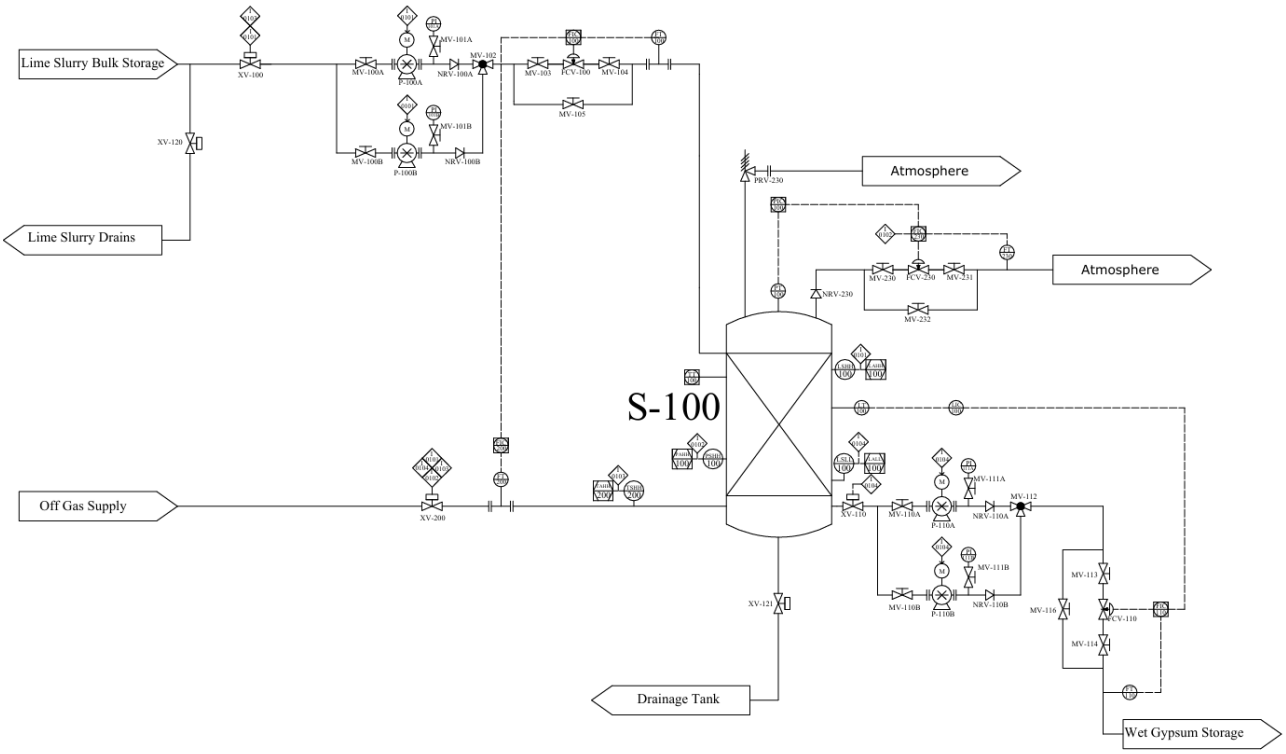

Figure 3. Wet Flue Gas Desulfurization P&ID

traditional challenge of spatially mapping text annotations to corresponding visual symbols is bypassed. Importantly, the reconstruction model had no access to the original P&ID images and shared no memory with the extraction model beyond the list of identified nodes.

Although Gemini-3.0-pro was used for both extraction and reconstruction in this case study, the workflows are model-agnostic. Any vision-language model can be substituted in the extraction stage, and any reasoning-capable language model can be used in the reconstruction stage.

Evaluation procedure

For each P&ID case study, a ground-truth representation was manually constructed in the same Python dictionary format as the model outputs, with equipment and instrumentation tags represented as nodes and their interconnections represented as edges. Performance was evaluated by comparing the predicted graph against the ground truth using accuracy-based metrics and a structural discrepancy metric.

To account for stochastic variability in language model outputs, each workflow was executed ten times using identical prompts. All reported results correspond to the mean and standard deviation across these runs.

Comparative Analysis

The experimental design enables several comparative analyses. First, the performance of end-to-end digitisation with an MLLM (E2E-MLLM) is compared against decomposed extraction and topology reconstruction (DETR) to assess the effect of task decomposition. Second, the end-to-end workflow is evaluated with and without a high-level process overview to examine the impact of contextual information. Finally, within the decomposed workflow, topology reconstruction with and without explicit chemical engineering design rules is compared to quantify the contribution of domain-specific knowledge to connectivity inference.

Together, these comparisons isolate the roles of contextual information, task decomposition, and engineering knowledge in language-model-based P&ID digitisation.

RESULTS

The results reported in Tables 2.A and 2.B are evaluated using three accuracy-based metrics and one structural discrepancy metric. Total accuracy measures the proportion of correctly reconstructed nodes and edges relative to the ground truth, while node accuracy and edge accuracy quantify the correctness of equipment identification and connectivity inference, respectively.

In addition to these local accuracy measures, global structural similarity between the predicted and ground-truth graphs is quantified using a simple distance metric. This metric counts the total number of discrete modifications required to reconcile the predicted graph with the ground truth, including missing nodes, extra nodes, missing edges, extra edges, and edge-type mismatches. A lower simple distance therefore indicates that fewer structural changes are required for the reconstructed P&ID to match the reference topology, providing a direct and interpretable measure of topological correctness. All reported values are expressed as mean ± standard deviation over ten independent runs.

Formally, the simple distance is defined as:

$$D_s = | V_{miss} | + | V_{extra} | + | E_{miss} | + | E_{extra} | + | E_{type} | \quad (1)$$

where V_{miss} and V_{extra} denote missing and extra nodes, E_{miss} and E_{extra} denote missing and extra typed edges, and E_{type}captures edge-type mismatches between the predicted and ground-truth graphs.

The accuracy metrics are computed as follows:

$$\boldsymbol{Total\ Accuracy} = \frac{\sum(\boldsymbol{Correct\ Nodes + Correct\ Edges})}{\sum(\boldsymbol{Total\ Nodes + Total\ Edges})} \quad (2)$$

$$\boldsymbol{Node\ Accuracy} = \frac{\sum \boldsymbol{Correct\ Nodes}}{\sum \boldsymbol{Total\ Nodes}} \quad (3)$$

$$\boldsymbol{Edge\ Accuracy} = \frac{\sum \boldsymbol{Correct\ Edges}}{\sum \boldsymbol{Total\ Edges}} \quad (4)$$

Nitrogen PSA plant

Results for the nitrogen pressure-swing adsorption (PSA) P&ID are summarised in Table 2.A. Using end-to-end digitization with a multimodal large language model (E2E-MLLM) as the baseline configuration, the model achieves a total accuracy of 68.56%, with relatively high node accuracy (86.32%) but substantially lower edge accuracy (55.58%). This disparity indicates that while equipment identification is largely successful, connectivity inference remains challenging when extraction and reasoning are performed jointly.

Augmenting the end-to-end prompt with a process overview yields a modest improvement across all metrics, increasing total accuracy to 70.67% and reducing the simple distance from 45.6 to 41.0. This suggests that high-level contextual information helps resolve some connectivity ambiguities; however, structural discrepancies remain pronounced, as reflected by the still relatively high simple distance.

A more substantial improvement is observed when switching to the decomposed extraction and topology reconstruction (DETR) workflow augmented with both a process overview and explicit process design rules. In this configuration, DETR achieves a markedly higher total accuracy of 88.78%, with node accuracy reaching 100% and edge accuracy improving to 80.58%. This improvement is accompanied by a significant reduction in simple distance to 24.0, indicating that substantially fewer node and edge modifications are required to align the reconstructed graph with the ground truth.

Slime and Off gas Plant

Results for the slime and off-gas plant P&ID, which represents a more complex process topology, are shown in Table 2.B. In this case, the limitations of end-to-end digitisation are more pronounced. The baseline E2E-MLLM configuration achieves a total accuracy of 67.31% and exhibits a high simple distance (68.20), indicating substantial structural deviation from the reference topology.

Providing a process overview in the end-to-end workflow improves total accuracy to 71.81% and reduces the simple distance to 59.70, highlighting the increasing importance of contextual information as diagram complexity grows. Nevertheless, a significant number of structural discrepancies persist, particularly in connectivity inference.

Transitioning to the decomposed extraction and topology reconstruction (DETR) workflow further clarifies the role of domain knowledge. While decomposed extraction and reconstruction without explicit process knowledge does not consistently outperform end-to-end digitisation, the highest performance is achieved when DETR is augmented with both a process overview and chemical engineering design rules. In this configuration, the model attains a total accuracy of 85.21%, near-perfect node accuracy (99.73%), and a substantially reduced simple distance of 48.30, indicating a closer match to the ground-truth process topology.

Table 2A. Nitrogen PSA Plant Results

	E2E	E2E + Knowledge	DETR	DETR + Knowledge
Total Acc.(%)	68.56 ± 0.07	70.67 ±0.06	71.22 ± 0.03	88.78 ± 0.02
Node Acc. (%)	86.32 ± 0.07	87.37 ± 0.05	100.00± 0.00	100.00 ± 0.00
Edge Acc. (%)	55.58 ± 0.08	59.98 ± 0.09	50.19 ± 0.05	80.58 ± 0.04
Simple Distance	45.60 ± 6.52	41.00 ± 6.70	51.80 ± 3.99	24.00 ± 4.67

Table 2B. Slime & Offgas Plant Results

	E2E	E2E + Knowledge	DETR	DETR + Knowledge
Total Acc. (%)	67.31 ± 0.07	71.81 ± 0.06	64.74 ± 0.02	85.21 ± 0.03
Node Acc. (%)	78.38 ± 0.08	78.20 ± 0.09	100.00± 0.00	99.73 ± 0.01
Edge Acc. (%)	58.87± 0.07	66.63 ± 0.07	37.84± 0.04	74.12 ± 0.05
Simple Distance	68.20 ± 12.64	59.70± 10.90	116.60 ± 8.22	48.30 ± 10.23

Across both case studies, several consistent trends emerge. First, DETR augmented with process knowledge consistently outperforms end-to-end digitisation across all metrics, with the largest gains observed in edge accuracy and simple distance. Second, providing high-level process context improves performance for both workflows but does not fully resolve structural inconsistencies when extraction and reasoning are coupled. Third, incorporating explicit chemical engineering design rules during topology reconstruction yields the most structurally faithful representations, as evidenced by the lowest simple distances.

Taken together, these results show that accurate P&ID digitisation requires not only reliable visual extraction but also explicit reasoning over process structure, and that separating these concerns—when guided by domain knowledge—leads to reconstructions that are both more accurate and more closely aligned with engineering intent.

CONCLUSION

This work investigated the use of multimodal large language models and decomposed, agent-based workflows for the semantic digitisation of piping and instrumentation diagrams. By reframing P&ID digitisation as a problem of extracting equipment tags and inferring process topology, we demonstrated that separating visual extraction from topology reconstruction leads to more accurate and structurally consistent representations than end-to-end digitisation with a single prompt. Across two industrially representative case studies, the proposed decomposed extraction and topology reconstruction workflow consistently outperformed end-to-end approaches, particularly in edge accuracy and structural consistency.

The results show that incorporating high-level process context and explicit chemical engineering design rules during topology reconstruction substantially improves performance, especially for more complex diagrams. This highlights the importance of embedding domain knowledge directly into the digitisation workflow, allowing chemical engineers to actively contribute to the definition of process rules rather than relying on extensive manual refinement of rule-based or data-driven computer vision systems. The resulting structured representations align naturally with graph-based standards such as DEXPI, supporting downstream integration with digital twins and intelligent engineering tools.

A practical limitation of the proposed approach is the use of proprietary language models, which raises confidentiality concerns when handling sensitive industrial P&IDs. However, the workflow itself is model-agnostic and can be readily adapted to open-source multimodal language models, enabling deployment in environments with strict data governance requirements. A systematic comparison between proprietary and open-source models was beyond the scope of this study and remains an important direction for future work.

Future research will focus on closing the loop between reconstruction and validation by introducing feedback mechanisms that iteratively refine the inferred topology based on consistency checks and engineering constraints. Such feedback-driven refinement has the potential to further reduce structural discrepancies and enable fully automated, trustworthy P&ID digitisation pipelines based on multimodal language models.

REFERENCES

1. Paliwal S, Sharma M, Vig L. OSSR-PID: one-shot symbol recognition in p&id sheets using path sampling and GCN. 2021 International Joint Conference on Neural Networks (IJCNN) :1-8 (2021). https://doi.org/10.1109/ijcnn52387.2021.9534122
2. Kim H, Lee W, Kim M, Moon Y, Lee T, Cho M, Mun D. Deep-learning-based recognition of symbols and texts at an industrially applicable level from images of high-density piping and instrumentation diagrams. Expert Systems with Applications 183:115337 (2021). https://doi.org/10.1016/j.eswa.2021.115337
3. Yu ES, Cha JM, Lee T, Kim J, Mun D. Features recognition from piping and instrumentation diagrams in image format using a deep learning network. Energies 12:4425 (2019). https://doi.org/10.3390/en12234425
4. Schulze Balhorn, L., Seijsener, N., Dao, K., Kim, M., Goldstein, D. P., Driessen, G. H., & Schweidtmann, A. M. (2025). Rule-based autocorrection of Piping and Instrumentation Diagrams (P&IDs) on graphs. *arXiv e-prints*, arXiv-2502.
5. Jamieson L, Elyan E, Moreno-García CF. Few-shot symbol detection in engineering drawings. Applied Artificial Intelligence 38: (2024). https://doi.org/10.1080/08839514.2024.2406712
6. Kang SO, Lee EB, Baek HK. A digitization and conversion tool for imaged drawings to intelligent piping and instrumentation diagrams (p&id). Energies 12:2593 (2019). https://doi.org/10.3390/en12132593
7. Choi J, An D, Kang D. Object recognition scheme for digital transformation in marine science and engineering. JMSE 11:1914 (2023). https://doi.org/10.3390/jmse11101914
8. Kim BC, Kim H, Moon Y, Lee G, Mun D. End-to-end digitization of image format piping and instrumentation diagrams at an industrially applicable level. Journal of Computational

Design and Engineering 9:1298-1326 (2022). https://doi.org/10.1093/jcde/qwac056
9. Mani S, Haddad MA, Constantini D, Douhard W, Li Q, Poirier L. Automatic digitization of engineering diagrams using deep learning and graph search. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) :673-679 (2020). https://doi.org/10.1109/cvprw50498.2020.00096
10. Stürmer JM, Graumann M, Koch T. From engineering diagrams to graphs: digitizing p&ids with transformers. 2025 IEEE 12th International Conference on Data Science and Advanced Analytics (DSAA) :1-11 (2025). https://doi.org/10.1109/dsaa65442.2025.11248012
11. Gupta M, Wei C, Czerniawski T. Semi-supervised symbol detection for piping and instrumentation drawings. Automation in Construction 159:105260 (2024). https://doi.org/10.1016/j.autcon.2023.105260
12. Hain A, Gölzhäuser S, Réhault N, Brox T, Demant M. Legend-informed symbol recognition in engineering diagrams with self-supervised learning. Lecture Notes in Computer Science :403-421 (2025). https://doi.org/10.1007/978-3-662-72243-5_23
13. Lu C, Walker S, Struck C, Itard L, Saelens D. Poster abstract: p&id-to-graph: llm-assisted digitalization of HVAC diagrams. Proceedings of the 12th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation :296-297 (2025). https://doi.org/10.1145/3736425.3772104

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Task-Conditioned Hierarchical Representations for Controllable AI-Assisted Process Synthesis

Ali Tarik Karagoz[a], Omar Alqusair[ab], and Jie Li[a]*

[a] The University of Manchester, Department of Chemical Engineering, Manchester, UK
[b] King Saud University, Department of Chemical Engineering, Riyadh, Saudi Arabia
* Corresponding Author: jie.li-2@manchester.ac.uk.

ABSTRACT

Machine learning (ML) has attracted growing interest in process systems engineering for its potential in process design, synthesis, and optimization. By learning complex patterns from data, ML methods complement traditional first-principles modelling and heuristic approaches, particularly for conceptual process design and the exploration of alternatives. Although current text-based representations capture unit-level connectivity, they lack a holistic view of process intent, equipment hierarchy, and contextual information to guide learning and inference. Consequently, models trained on such linear token sequences tend to reproduce syntactic structure rather than underlying process reasoning, thus limiting interpretability and user control. In this work, we introduce a contextual framework for representing process flowsheet information in ML models that embeds process engineering logic directly into the model inputs. The approach combines a structured, text-based representation of process topology with context descriptors that define the process type and synthesis task. The task context, such as simplified block flow design, enables engineers to steer learning and generation towards intended designs. This shifts control from post-training filtering to context-driven learning, allowing models to act as flexible, engineer-guided synthesis tools rather than passive pattern generators. Transformer models were trained on 100, 000 synthetic ethylene glycol process flowsheets, with and without context features. To emulate user-guided synthesis, the models were asked to generate different design configurations involving critical equipment (reactors and distillation columns). Models trained without contextual input achieved an average success rate of 24.6% on average (59.6% in the best case), whereas context-aware models achieved 98.8% on average.

Keywords: Artificial Intelligence, Process Synthesis, Machine Learning, Process Design, Intelligent Systems

INTRODUCTION

Machine learning (ML) and artificial intelligence (AI) have a long history in both computer science and chemical engineering research. Since the 1980s, early trials sought to integrate AI methods into chemical engineering toolboxes [1]. Although research activity was intermittent, progress was made in property prediction, process control, and surrogate modelling. There were some trials, even in process synthesis and design as early as expert systems [2], but the exploration was not continuous until recently. With the progress of ML research since the 2010s, a new wave of AI-assisted or ML-integrated process synthesis approaches has emerged. Recent advances in generative AI have shown promise for producing candidate flowsheets, completing partial designs, and accelerating conceptual process design studies [3–7]. Different methodologies and models have been applied, ranging from power cycles [8, 9] to chemical production processes of ethylene oxide [10], ethyl tert-butyl ether [11], dimethyl ether [12], and methyl acetate [13]. These developments raise the prospect of interactive, data-driven tools that can support engineers in exploring wide design spaces during early-stage process development.

One of the central challenges lies in how to represent process flowsheet information for ML. Depending on the chosen models and approach, text sequences [9, 14], graphs [13], or matrices [15] can be used to carry the topological information. These representations capture unit-

level connectivity and basic chemical engineering principles and may include control elements [14] or operational information [16]. Nevertheless, when training ML models, most approaches treat process flowsheets as pure autoregressive problems, learning how to predict the next token from previous information. As a result, higher-level engineering concepts, such as equipment hierarchy or relative importance, are not explicitly communicated to the model. It is expected that the model should infer such abstractions from sufficient examples.

This limitation becomes particularly apparent when generative models are used to explore alternatives or to produce a certain type of process. Without explicit information about the desired process identity or design objective, models tend to reproduce biases occurring in the training data. In the work of Nabil et al. [9], although the model can learn higher-complexity structures syntactically, these gradually disappear from the outputs due to computational limitations, finite sample sizes, and complex structures failing filters more frequently. This creates an undesired and uneven exploration-exploitation trade-off that is dominated by simpler designs [17].

Probable solutions include inference-time masking strategies to promote exploration or information-gain objectives to increase solution-space coverage. Although the latter have not yet been explored in the process synthesis literature, there are masking examples to increase model accuracy and success rates [12]. However, its primary role is to simplify decision-making and accelerate convergence rather than to increase structural diversity. Although such strategies can improve many performance metrics, they effectively force outputs toward prescribed patterns rather than enabling the model to internalize nuances or design principles.

These limitations motivate incorporating controllability into the representation itself, rather than relying on inference-time interventions. Moreover, aggressive masking intended to enforce user-specified topologies may push models outside their learned distributions, leading to degraded quality or unrealistic design actions. In contrast, human engineers naturally reason at multiple levels of abstraction and information before engaging with detailed process structures.

In this work, we argue that embedding these engineering abstractions directly into ML inputs is essential for controllable and interpretable AI-assisted process synthesis. We introduce a contextual framework that augments flowsheet representations with task descriptors specifying the number and ordering of critical units with an emphasis on unit hierarchies. This formulation enables the model to associate each training example with an explicit design intent, such as simplified block-flow design, rather than inferring such objectives implicitly from topology alone. To complement this approach, we introduce a new representation that linearizes process flowsheets by focusing on the main branch, which is constructed from major feed to major product to emphasize the primary design objective.

The proposed approach is demonstrated using the ethylene glycol process. Synthetic flowsheet examples were generated for two datasets of increased structural complexity. Transformer models trained with and without task context are compared to evaluate syntactic learning, predictive performance, and the ability to generate user-specified process identities.

METHODS

Task Context

The task context is not the same as the context used in NLP applications, but rather an additional feature to further describe the flowsheet topology based on the design intent. Nevertheless, they both serve the purpose of reducing the ambiguity in learning and generating token sequences. Task context is not inherently present in the process topology data. It is an engineering perspective to create an elevated level of abstraction analogous to the ideas in PSE, such as block flow diagrams (BFD) and process flowsheet diagrams (PFD). Like BFDs and PFDs, task context is aimed at pointing to the same process flowsheet information as the topology at a different level without any contradiction.

The importance hierarchy of process elements is another core idea. From the earliest examples of systematic process synthesis approaches to the contemporary research efforts in specific units, the criticality of the equipment is not related to the frequency of its presence. While the reaction and separation approaches in the process shape the process identity, the rest of the equipment is there to create a feasible foundation or enhance the process performance most of the time. However, this idea is not present when applied to process topology alone. These critical units that define the characteristics of the process can be reduced to a small group of tokens that are necessary for successful completion but not critical.

Combining these two ideas of information abstraction and importance hierarchy for conceptual process synthesis gives us a task context that relates to the number and the sequence of the critical equipment that is desired to be present in the process flowsheet. It signals the model what the process identity or design intent is when it is learning or generating the process flowsheets with desired characteristics. It does not provide any information about how the rest of the process equipment should be constructed or constrain certain outputs directly through the manipulation of probability distributions.

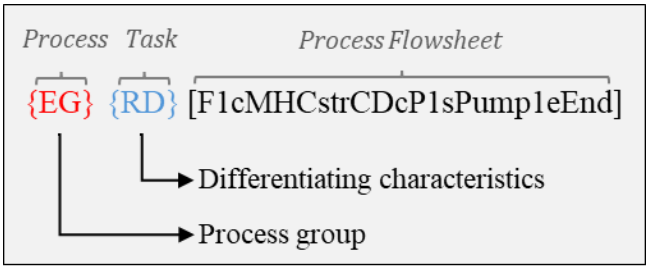

Figure 1. Task context and corresponding text-based flowsheet representation for an EG process. The {R, D} task context specifies the number and ordering of critical units.

A simple example is illustrated in **Figure 1** using a simple ethylene glycol production process based on ethylene oxide hydrolysis. The {R, D} task context indicates that the following process flowsheet has exactly 1 reactor equipment that precedes 1 distillation equipment in the process flow. It does not tell where exactly they are placed or any restrictions regarding the ancillary equipment.

Text Representation

The task context approach is an independent abstraction feature that can be modified to work with any kind of process representation. However, since the approach is based on better distinction between the process topology characteristics, a new style of text-based representation was developed for process flowsheet information to aid this purpose. The representation is similar to SFILES 2.0 [14] in some parts, but it is combined with case sensitivity and numerical tokens [8, 9].

Building on this foundation, three critical aspects were further improved. Unlike the current local graph traversal approaches, this representation tries to focus on the process flowsheet holistically to create a linearized global connection. The assumption comes from the fact that most processes have raw materials that are either consumed heavily or a critical stream, which constitutes the major feed as the starting point to reach the major product as the end. The shortest path of streams and equipment between the major feed and the major product defines the main branch, which contains the most critical equipment for achieving the desired production.

Additionally, all the equipment is assumed to have a major outlet, which is signified by the subtype of representing token in the case of multiple outlets. The connections and detachments created by minor feeds and products, recycles, and other equipment outlets creates subbranches that are represented uniquely. This procedure yields a canonical textual form for each flowsheet topology, reducing alternative linearizations when the main feed and product are fixed. Such invariance is expected to improve learning efficiency for limited datasets similar to the graph representation [18].

Lastly, human readability was also prioritized. All representations are aimed at transforming the process flowsheet information into a machine-readable and easy-to-store format. Nevertheless, as the machine itself just sees them as vectors or tensors, the text side should aim for humans to bridge the gap between human and machine understanding. Units and connectivity tokens with their representations are shown in **Table 1**. Capital letters denote the start of unit tokens, lowercase characters for remaining characters or specify sub-types, and numbers indicate the subbranches with lowercase characters to signal start, connection, and end of the subbranch.

Table 1: Equipment and connectivity used in the proposed text-based representation.

Equipment	Text Representation
Major Feed	F
Minor Feed	Fa, Fb, ...
Major Product	P
Minor Product	Pa, Pb, ...
Subbranch Starter	1s, 2s, 3s
Subbranch Connector	1c, 2c, 3c
Subbranch End	1e, 2e, 3e
Mixer	M
Heater	H
Cooler	C
Pump	Pump
CSTR	Cstr
PFR	Pfr
Distillation C. (Top P.)	Dc
Distillation C. (Bottom P.)	Dcb
End	End

The representation is still under active development for further grouping based on unit purpose, inclusion of more unit types, and rules for complex process topologies. Regardless, it is fully functional for the scale of this example.

Dataset

The datasets for the tasks were composed of synthetic process examples of EG production processes via EO hydrolysis. All samples are unique and passed through the syntax filter to ensure a valid process flowsheet representation. To ensure variety and an adequate representation of process topology domain, random sampling from uniform distributions was used for the total token number, task contexts, and equipment placement.

Using bounded random generation and syntactic filtration, we obtain D1 and D2 datasets, each containing 100, 000 unique process flowsheet representations. To illustrate the procedure, more than 400, 000 examples were generated and filtered to create the D2 dataset. The datasets differ in the total number of reactors and

distillation columns per flowsheet (3 for D1 and 4 for D2) and their task representations. It results in increased topological complexity and length of representations in D2. Distribution of data with respect to the task context is reported in **Table 2**. Uniform distribution across task context classes was not enforced, as a skewed distribution is more realistic and helps the analysis of the model performance in underrepresented configurations.

Table 2: Distribution of task context classes in the training datasets of D1 and D2

Task Context	D1 (%)	D2 (%)
RD	22.9	6.4
RDD	21.5	7.0
RDR	39.6	13.8
RRD	16.1	4.9
RDDD	-	7.1
RDDR	-	12.7
RDRD	-	7.9
RDRR	-	18.6
RRDD	-	6.1
RRDR	-	11.8
RRRD	-	3.7

Although it is out of scope for this experiment, a batch of 1000 designs was simulated in Aspen Plus to see if generated data and syntax filters would lead to real process examples. All designs were successfully built, simulated, and the unit design parameters were explored. Over 50% of the designs converged to feasible operating conditions with a limited computational budget.

Models, Training, and Inference

As the representation and the problems are formulated as symbolic sequences like NLP applications, a standard transformer architecture with ~1M parameters and 4 attention heads was used. This is a considerably smaller model compared to both applications in NLP and process synthesis [3]. The resulting token to parameter-ratio is on the lower side of the acceptable bounds for this kind of model training, but higher than contemporary works in chemical engineering applications.

All models were trained using an 80/20 train–validation split for ten epochs with an autoregressive objective, cross-entropy loss, and the Adam optimizer. Identical datasets and training protocols were used for both context-aware and no-context models; in the latter case, task tokens were removed from the inputs to replicate the standard autoregressive model behavior.

To illustrate the effectiveness of task context, all models were asked to generate 10, 000 process flowsheets for every trial at the inference stage. In the first trial, models are just generating examples without any restrictions to show their syntax learning capabilities to establish the baseline. Then, models are tasked to generate processes with a desired identity for which similar examples were already present in their respective training datasets. This trial is the main experiment that imitates the model control by an engineer. The generation requests are picked from a uniform random distribution to symbolize user input, e.g., "Provide me an EG hydrolysis process example with 1 reactor and 1 distillation column". Lastly, a dynamic inference masking was introduced to help models generate designs with the chosen task context. The mask tries to control the model outputs by blocking the probability distribution of unwanted critical equipment until the necessary ones are picked according to the desired sequence.

RESULTS AND DISCUSSION

Training with task context influences the evolution of the loss function, as shown in **Figure 2**. Introducing task descriptors provides an additional conditioning signal that helps the model differentiate equipment choices associated with varying process structures. Although the primary objective of task context is improved controllability during generation, it also improves predictive performance during training by better capturing variations based on task context in process topology. This trend is reflected in the final perplexity values reported in **Table 3**, where context-aware models achieve consistently lower scores across both datasets.

Table 3: Training and validation perplexity for transformer models trained with and without task context.

Model	PP_{tr}	PP_{val}
No-context (D1)	3.05	3.05
Context (D1)	2.77	2.77
No-context (D2)	3.20	3.21
Context (D2)	2.77	2.77

The magnitude of the improvement remains modest, indicating that the model capacity and dataset size are already sufficient for learning syntactic structure. Nevertheless, the consistent reduction across datasets suggests that task context provides a useful inductive bias in the training.

To further establish the model capabilities, the first trial is conducted without any restrictions on process identity, focusing only on syntactic production and process representation. All models are fully capable of producing text representations that pass our syntax filters with almost perfect scores. As it is shown in **Table 4**, the numerical differences are insignificant. We used nucleus sampling ($p = 0.9$) to generate these results; however, they did not show any noteworthy change when multinomial sampling was used as well. For the successful synthesis rate, not just validity but uniqueness was enforced, as the model that answers a single process configuration

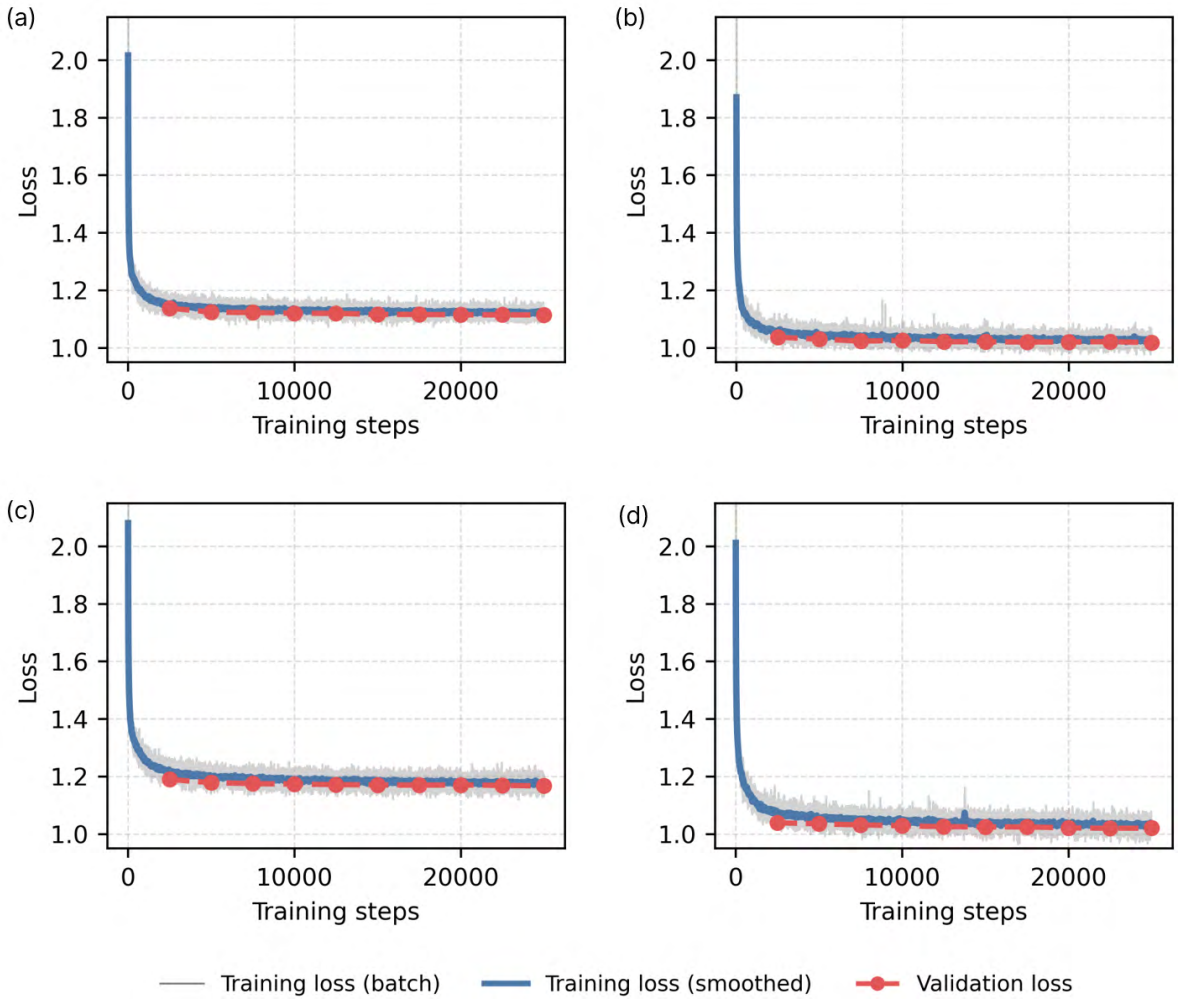

Figure 2. Training and validation loss for the four experimental configurations: (a) no-context training on D1, (b) context-conditioned training on D1, (c) no-context training on D2, and (d) context-conditioned training on D2. Grey lines denote batch level training loss, blue lines show the smoothed trajectory, and red dashed lines correspond to validation loss.

for 10, 000 requests would not be a desired final model.

Table 4: Successful synthesis rates (valid and unique) for unrestricted generation across datasets and model variants.

Model	Successful Synthesis (%)
No-context (D1)	98.94
Context (D1)	98.88
No-context (D2)	99.10
Context (D2)	98.65

These results indicate that the structure and the amount of data are appropriate for the complexity of the task. Reducing the dataset size from 100, 000 to 1, 000 leads to a substantial performance drop in both cases. Success rates dropped as low as 86.9% for D1 and 79.6% for D2, reflecting the increased complexity of designs in D2. In heavily restricted data regimes, context-conditioned models initially perform worse than their counterparts because they must learn an additional conditioning signal. Notably, this difference diminished as structural complexity increased, reducing from 7% to 0.6% between D1 and D2.

As the generative competence of models is established, the next step was to measure their performance under user-specified process synthesis based on critical equipment number and sequence. As no-context models cannot receive explicit information, they have no way of distinguishing the input and the process that is generated. The outcomes basically reflect the model bias from training data towards certain types.

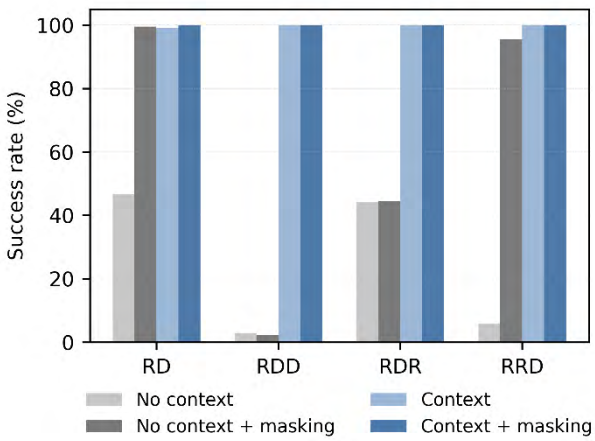

Figure 3. Distribution of generated process identities for no-context and context-aware models under uniform user requests, with and without inference-time masking, for dataset D1.

As shown in **Figure 3**, no-context models favor designs with a reactor and a distillation column (RD) and the addition of another reactor (RDR) as its output. These two categories account for more than 90% of the model answer. However, since the requests are sampled from a uniform distribution, the average correct answer rate drops to 24.6%. Even with the dynamic inference mask, this number only reaches 60%. The mask can be improved further by disabling end token selection until the desired sequence is produced to increase the success rate, but it risks pushing the model outside of the learned distribution. Thus, the accuracy and quality of generated designs can degrade. On the other hand, task context conditioned models perform almost perfectly with or without the mask, as their average performance never drops below 98.6% for both datasets.

The effect becomes even more pronounced in the more complex D2 dataset, where the number of possible task categories increases substantially. Here, the no-context model's average success rate drops from 24.6% to 8.72%. Category-level analysis in **Figure 4** shows that four configurations (RD, RDR, RDRR, and RRDR) account for 76% of generated designs, further highlighting strong structural bias. Masking again improves performance, but only to 40.3%, whereas task-conditioned models sustain an average success rate of 98.8%.

Overall, these results demonstrate that task context embeds high-level design intent directly into the learning

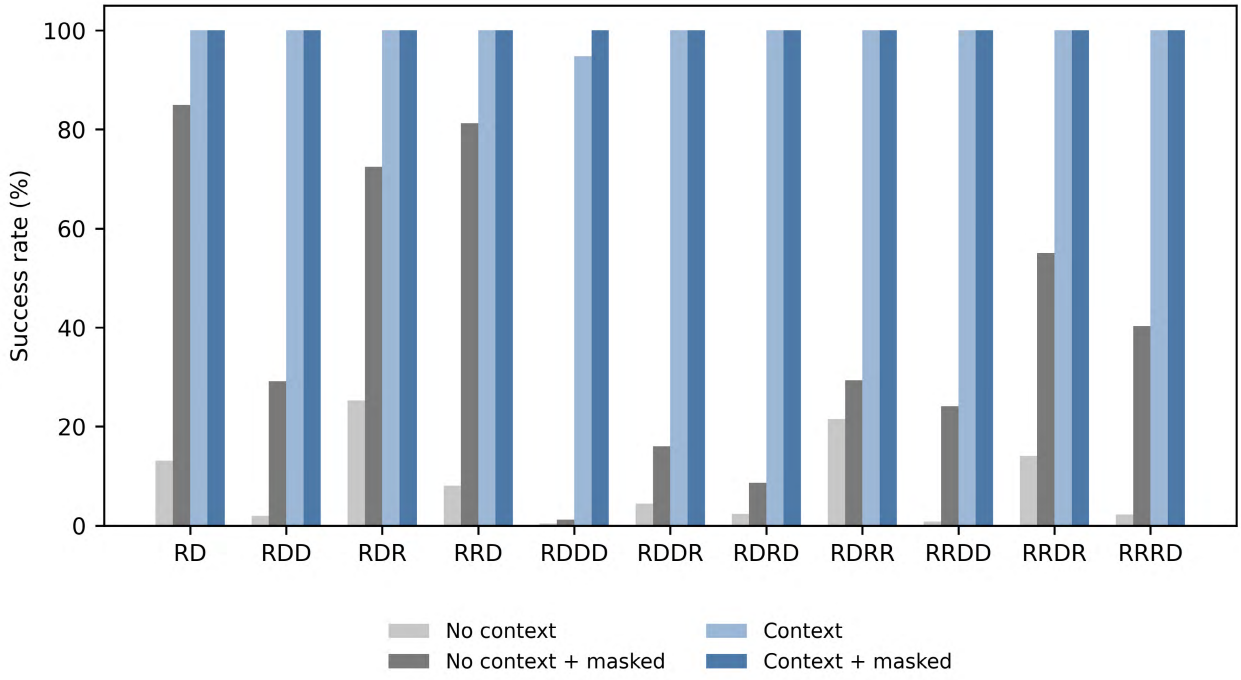

Figure 4. Success rates for generating valid and unique flowsheet structures under four settings for D2 dataset for user control trial. Bars compare models trained without contextual conditioning ("No context"), with masking only ("No context + masked"), with task context conditioning ("Context"), and with both conditioning and masking ("Context + masked"). Each group corresponds to a target structure class that was present in the training dataset. Contextual models achieve near-perfect compliance across all specifications, while masking alone improves but does not close the gap relative to explicit contextual conditioning.

problem, improving not only controllability during synthesis but also predictive performance during training. When deployed in AI-assisted conceptual process design, contextual conditioning provides a practical mechanism for interpreting model biases, steering generation toward desired structures, and reducing reliance on brittle post-processing rules.

CONCLUSION

In this work, we introduced a task-context feature for AI-assisted process synthesis to increase the interpretability of model outputs and user control. By embedding a higher-level abstraction of process topology into the training data through the number and ordering of critical units, machine-learning models move beyond purely syntactic generation. These models are therefore able to capture relationships between design examples and can be guided by an engineer during both learning and generation.

A new text-based representation emphasizing a main process branch and global traversal was developed to complement this objective. Together with task descriptors, this representation provides the model with explicit structural cues while retaining flexibility in the construction of ancillary equipment and subbranches. The resulting models achieve near-perfect syntactic validity across datasets of increasing complexity, while modestly improving perplexity relative to their no-context counterparts. More importantly, the inclusion of task context enables reliable controlled synthesis. Models trained without context largely reproduce dataset and model biases and perform poorly when asked to generate specific process identities, even when supported by dynamic inference masking. Context-aware models consistently satisfy user-specified design objectives across both simple and complex datasets. These results highlight that conditioning during training is substantially more effective than post-hoc probability manipulation at inference.

Beyond quantitative performance, the proposed framework offers practical advantages for AI-assisted process design. Task context provides an interpretable interface between human engineers and generative models and supports interactive exploration of alternative process concepts. In this sense, the approach shifts generative models from passive pattern reproducers toward controllable design assistants. Although demonstrated here for ethylene glycol production, the proposed framework is process-agnostic and intended to extend to other synthesis domains.

Future work will focus on expanding the vocabulary of task context and developing extensions for additional process design objectives, such as process integration and process intensification. Within process synthesis, the framework will be extended to accommodate more complex flowsheets and cross-process training to alleviate data limitations in further applications. Integrating optimization and evaluation algorithms with this framework will further enable efficient exploration of design alternatives and identification of high-quality solutions within the feasible design space. Overall, this study demonstrates that embedding engineering intent directly into representations is a promising direction for scalable and controllable AI-assisted process synthesis.

ACKNOWLEDGEMENTS

ATK would like to appreciate financial support from The University of Manchester. ATK is a participant of the Study Abroad Programme from the Republic of Turkey which he is grateful for the continuous support from the Ministry of Education and the Secretariat of Defence Industries. OA would like to appreciate financial support from King Saud University.

AUTHOR IDENTIFIERS

Author ORCIDs:
Karagoz AT: 0009-0002-4105-880X
Alqusair O: 0009-0006-6461-4388
Li J: 0000-0001-5196-2136

REFERENCES

1. Venkatasubramanian V. The promise of artificial intelligence in chemical engineering: is it here, finally?. AIChE Journal 65:466-478 (2018). https://doi.org/10.1002/aic.16489
2. Stephanopoulos G, Johnston J, Kriticos T, Lakshmanan R, Mavrovouniotis M, Siletti C. Designkit: an object-oriented environment for process engineering. Computers & Chemical Engineering 11:655-674 (1987). https://doi.org/10.1016/0098-1354(87)87010-2
3. Vogel G, Schulze Balhorn L, Schweidtmann AM. Learning from flowsheets: a generative transformer model for autocompletion of flowsheets. Computers & Chemical Engineering 171:108162 (2023). https://doi.org/10.1016/j.compchemeng.2023.108162
4. Schweidtmann AM, Esche E, Fischer A, Kloft M, Repke J, Sager S, Mitsos A. Machine learning in chemical engineering: a perspective. Chemie Ingenieur Technik 93:2029-2039 (2021). https://doi.org/10.1002/cite.202100083
5. Nabil T, Commenge JM, Neveux T. Generative Approaches for the Synthesis of Process Structures. In: Computer Aided Chemical Engineering. Elsevier; 2022. p. 289–94. (14

International Symposium on Process Systems Engineering; vol. 49).
6. Daoutidis P, Lee JH, Rangarajan S, Chiang L, Gopaluni B, Schweidtmann AM, Harjunkoski I, Mercangöz M, Mesbah A, Boukouvala F, Lima FV, del Rio Chanona A, Georgakis C. Machine learning in process systems engineering: challenges and opportunities. Computers & Chemical Engineering 181:108523 (2024). https://doi.org/10.1016/j.compchemeng.2023.108523
7. Karagoz AT, Alqusair O, Liu C, Li J. Advances in conceptual process design: from conventional strategies to ai-assisted methods. Chinese Journal of Chemical Engineering 84:60-76 (2025). https://doi.org/10.1016/j.cjche.2025.05.014
8. Zhou P, Zhang J, Le Moullec Y. Dynamic modeling and transient analysis of a recompression supercritical CO2 brayton cycle. AIP Conference Proceedings 2303:130011 (2020). https://doi.org/10.1063/5.0029260
9. Nabil T, Noaman M, Morosuk T. Data-driven structural synthesis of supercritical CO2 power cycles. Front. Chem. Eng. 5: (2023). https://doi.org/10.3389/fceng.2023.1144115
10. Khan AA, Lapkin AA. Designing the process designer: hierarchical reinforcement learning for optimisation-based process design. Chemical Engineering and Processing - Process Intensification 180:108885 (2022). https://doi.org/10.1016/j.cep.2022.108885
11. Göttl Q, Tönges Y, Grimm DG, Burger J. Automated flowsheet synthesis using hierarchical reinforcement learning: proof of concept. Chemie Ingenieur Technik 93:2010-2018 (2021). https://doi.org/10.1002/cite.202100086
12. Reynoso?Donzelli S, Ricardez?Sandoval LA. A reinforcement learning approach with masked agents for chemical process flowsheet design. AIChE Journal 71: (2024). https://doi.org/10.1002/aic.18584
13. Stops L, Leenhouts R, Gao Q, Schweidtmann AM. Flowsheet generation through hierarchical reinforcement learning and graph neural networks. AIChE Journal 69: (2022). https://doi.org/10.1002/aic.17938
14. Vogel G, Hirtreiter E, Schulze Balhorn L, Schweidtmann AM. SFILES 2.0: an extended text-based flowsheet representation. Optim Eng 24:2911-2933 (2023). https://doi.org/10.1007/s11081-023-09798-9
15. Göttl Q, Grimm DG, Burger J. Automated synthesis of steady-state continuous processes using reinforcement learning. Front. Chem. Sci. Eng. 16:288-302 (2021). https://doi.org/10.1007/s11705-021-2055-9
16. Mann V, Sales-Cruz M, Gani R, Venkatasubramanian V. Esfiles: intelligent process flowsheet synthesis using process knowledge, symbolic AI, and machine learning. Computers & Chemical Engineering 181:108505 (2024). https://doi.org/10.1016/j.compchemeng.2023.108505
17. Lee JH, Shin J, Realff MJ. Machine learning: overview of the recent progresses and implications for the process systems engineering field. Computers & Chemical Engineering 114:111-121 (2018). https://doi.org/10.1016/j.compchemeng.2017.10.008
18. Schulze Balhorn L, Degens K, Schweidtmann AM. Graph-to-sfiles: control structure prediction from process topologies using generative artificial intelligence. Computers & Chemical Engineering 199:109121 (2025). https://doi.org/10.1016/j.compchemeng.2025.109121

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Recommendation System for Prediction of Adsorption Properties using Kernelized Probabilistic Matrix Factorization

Gnaneshwar Sampathirao[a], Sasidhar Gumma[a], Nabil Magbool Jan[a*]

[a] Indian Institute of Technology Tirupati, Department of Chemical Engineering, Tirupati, Andhra Pradesh, India
* Corresponding Author: nabil@iittp.ac.in

ABSTRACT

Porous materials such as Metal-Organic Frameworks and Covalent Organic Frameworks are emerging adsorbent materials with tunable structures and chemistry, making them useful for applications such as carbon capture, drug delivery, gas separations, and storage. This work aims to design and develop a systematic approach to build a data-driven recommendation system that leverages the historical experimental data or simulation data to assist process engineers in identifying the most suitable adsorbents from a large candidate space. In general, only some of the adsorption properties are available for porous materials owing to limited experimental data. In this scenario, this problem can be formulated as a matrix completion problem, which aims to impute the missing data by exploiting the underlying pattern in the available data. To this end, we propose a parameterization of the kernelized probabilistic matrix factorization framework, which aims to determine the nonlinear latent factors that are parameterized. The resulting bi-convex Maximum a Posteriori objective with reproducing kernel Hilbert space penalties can be solved using an alternating minimization approach. We demonstrate this approach on a publicly-available COF dataset. Of the 16 adsorption properties considered in this study, 15 of them can be predicted very accurately for a missing percentage of up to 60%. Further, the proposed approach preserves the ranking of COF candidates, which helps in accurate screening of best COFs for a given task.

Keywords: Matrix factorization, Alternating Minimization, Kernels, COFs

INTRODUCTION

Porous materials, particularly Metal-Organic Frameworks (MOFs) and Covalent Organic Frameworks (COFs), have emerged as materials characterized by their exceptional tunability in structure, and chemical functionality [1]. COFs are built from elements connected through covalent bonds, yielding materials with permanent porosity, high surface areas, and controllable pore environments [2]. These attributes have positioned these porous candidates as leading for a diverse applications including carbon capture, hydrogen and methane storage, selective gas separations, and catalysis. The vastness of the COF chemical space emphasizes the need of computational screening to identify promising candidates for specific applications, yet this screening is fundamentally constrained by the availability of property data. Experimental measurement of adsorption isotherms, selectivity, and Henry coefficients is time-consuming and resource-intensive, often requiring specialized equipment and expertise. Molecular simulations, can predict these properties computationally, but remain expensive for large-scale screening across diverse conditions [3, 4]. This scenario gives rise to a matrix completion problem.

To formally define the problem, consider an incomplete matrix $\mathrm{X} \in \mathrm{R}^{\mathrm{m}\times\mathrm{p}}$ where the rows correspond to m COF candidates, and the columns correspond to p adsorption properties. Only a subset of entries $\Omega \subseteq \{1, \dots, \mathrm{m}\} \times \{1, \dots, \mathrm{p}\}$ is observed, and the goal of the recommendation engine is to accurately impute the missing entries to allow complete screening and ranking of COF candidates for a given task. This problem of imputing missing values in materials property matrices has been

approached by Sturluson et al. where they used linear matrix factorization with l_2 regularization [5]. However, the relationships among material properties are nonlinear, and thus kernel methods offer an attractive approach to incorporate the similarity between materials and properties while guaranteeing the computational tractability of the proposed model [6].

In this work, we develop a parameterization of the kernelized probabilistic matrix factorization approach for imputing missing entries to provide reliable predictions for the recommendation system. Our approach extends the Kernelized Probabilistic Matrix Factorization (KPMF) by parameterizing both row and column latent factors through kernel matrices. The resulting inference problem is formulated as Maximum a Posteriori (MAP) estimation problem with RKHS-induced penalties on the latent factors, yielding a bi-convex objective that we optimize efficiently via alternating minimization. We demonstrate the efficacy of our approach on a benchmark adsorption property dataset [3, 4].

The main contributions of this work includes: (a) a parameterization of the kernelized matrix factorization model that uses latent factors as $\mathrm{Z} = \mathrm{K_M A}$ and $W = K_P B$, enabling the incorporation of side information on both materials and properties, (b) formulate the corresponding learning problem as MAP estimation with RKHS-induced penalties, (c) demonstrate the imputation performance of the proposed approach in a publicly available COF dataset for varying missing percentages, and (d) finally, we present screening of adsorbent materials to identify the best COF candidates for targeted adsorption applications.

The remainder of this paper is organized as follows. Firstly, we present the problem statement for developing the data-driven recommendation system, which involved accurate imputation and screening of materials. Next, we propose a data imputation technique using kernelized probabilistic matrix factorization followed by a parametrization strategy. Next, an alternating minimization algorithm for solving the proposed formulation is presented. Later, the screening of adsorbent material is presented. In the subsequent section, we demonstrate the effectiveness of the proposed recommendation system in a COF dataset. Finally, conclusions are presented.

PROBLEM STATEMENT

The objective of this work is to design and develop a systematic approach to build a data-driven recommendation system that leverages the historical experimental data or simulation data to assist process engineers in identifying the most suitable adsorbents from a large candidate space. The recommendation system should be capable of handling sparse and noisy data to capture the nonlinear relationship to accelerate screening of adsorbents and improve the decision quality. The computational screening of porous materials generates vast datasets of potential material candidates. However, the cost of molecular simulations – particularly Grand Canonical Monte Carlo (GCMC) calculations for generating the adsorption isotherms – scales up significantly, making comprehensive characterization of all candidates computationally infeasible. On the other hand, performing experiments to determine the adsorption properties of all adsorbent candidates is prohibitive due to the high cost. This naturally leads to datasets with substantial missing entries, thus solving a matrix completion problem is a crucial step in the development of a recommendation system.

In this context, given partial observations $\{\mathrm{X_{mp}}: (\mathrm{m}, \mathrm{p}) \in \Omega\}$ of a materials property matrix, our goal is to determine the missing observations $\{\mathrm{X_{mp}}: (\mathrm{m}, \mathrm{p}) \notin \Omega\}$ of the complete data matrix $\hat{\mathrm{X}}$ such that, $\hat{X}_{\mathrm{mp}} \approx \mathrm{X_{mp}}$ for all $(\mathrm{m}, \mathrm{p}) \in \Omega$ and the complete data matrix exhibits low-rank structure or lies on a low-dimensional manifold. This low-rank assumption is physically motivated, as in practice, material properties are governed by a smaller set of underlying factors that collectively determine macroscopic adsorption behavior. This latent structure enables imputation of unobserved entries by exploiting correlations among materials and properties. Consider the observation model of the form

$$X_{mp} = \sum_{l=1}^{r} f_l(m) g_l(p) + \epsilon_{mp} \quad \epsilon_{mp} \sim \mathcal{N}(0, \sigma^2) \quad (1)$$

where $\mathrm{f_l}(.)$ and $\mathrm{g_l}(.)$ denote the lth latent function of materials (rows) and properties (columns), respectively. r denotes the latent dimensions. The measurement errors ϵ_{mp} follow a univariate Gaussian distribution with mean zero and variance σ^2. Let $\mathrm{X_{mis}}$ and $\mathrm{X_\Omega}$ denote the missing and observed entries of a complete data matrix X. Then, the standard imputation problem is often cast as a maximum-a-posteriori (MAP) inference problem, and it can be stated in a probabilistic setting as follows

$$\hat{\mathrm{X}}_{miss} = \arg\max P(X_{mis} \mid X_\Omega) \quad (2)$$

The proposed formulation belongs to a class of probabilistic methods in a collaborative filtering framework. This approach enables one to reconstruct the full data matrix, which is useful for providing the ranking of adsorbent materials for a decision-focused objective of a recommendation system.

PROPOSED METHODOLOGY

This section develops a parametrized kernelized probabilistic matrix factorization approach for imputing missing entries. Firstly, we present the vanilla version of a non-parameterized kernelized probabilistic matrix factorization, followed by the parameterized version for

accurate imputation of adsorption properties [6]. To this end, we formalize the problem with kernel-encoded side information, derive the MAP formulation. Next, we develop the alternating minimization technique to solve the probabilistic imputation problem. Finally, we present the screening of adsorbent materials for potential applications.

Kernelized Probabilistic Matrix Factorization

In this section, we present the kernelized probabilistic matrix factorization approach to solve the matrix completion problem that characterizes the nonlinear latent dimensions, following the work of Zhou et al. [6]. In a standard matrix factorization problem, all materials' specific factors are considered to be independent (i.e., the isotropic Gaussian prior), regardless of any structure or similarity. However, in materials science applications, we often have side information indicating that certain materials/properties are similar---for instance, MOFs sharing amine functional groups are likely to exhibit good CO_2 adsorption properties [7], Here, we use kernels to encode prior knowledge about similarity among materials or properties. In this method, each latent dimension follows a Gaussian process prior. In other words, each latent dimension is assumed to be a random function drawn from a Gaussian process.

$$f_l(.) \sim \mathcal{GP}(0, k_m) \quad g_l(.) \sim \mathcal{GP}(0, k_p) \qquad l = 1, \cdots, r$$

where k_m, k_p are covariance functions which signify the similarity between two COFs (rows) or two properties (columns), respectively. Denoting $z_{ml} = f_l(m)$ and $w_{pl} = g_l(p)$, the observation model can be expressed as

$$X_{mp} = \sum_{l=1}^{r} f_l(m) g_l(p) + \epsilon_{mp} \qquad (3)$$

$$X_{mp} = \sum_{l=1}^{r} z_{ml} w_{pl} + \epsilon_{mp} \qquad (4)$$

Notice that it is a product of two Gaussian processes which makes the model non-Gaussian over X. It is not a Gaussian process over matrix entries. Now, the MAP estimation problem is presented as follows:

$$\{f_l^*, g_l^*\} = \arg\max P(X_{mis} \mid X_\Omega) = \arg\max P(\{f_l, g_l\} \mid X_\Omega) \qquad (5)$$

We can equivalently express it as,

$$\mathcal{L}_{\mathrm{KPMF}}(Z, W) = \sum_{(m,p)\in\Omega} \left(X_{mp} - \sum_{l=1}^{r} f_l(m) g_l(p) \right)^2 + \lambda_M \sum_{l=1}^{r} |f_l(m)|_{\mathcal{H}_m} + \lambda_P \sum_{l=1}^{r} |g_l(p)|_{\mathcal{H}_p} \qquad (6)$$

This is a variational problem in infinite-dimensional Reproducing Kernel Hilbert Spaces. This defines the matrix completion problem entirely in a function space. As this is computationally not-tractable, we consider evaluating the Gaussian process at discrete indices (say n) gives

$$f_l(.) = [f_{l1}(.), \cdots, f_{ln}(.)]^\top \sim \mathcal{N}\left(0, \frac{1}{\lambda_M} K_M\right) \qquad (7)$$

$$g_l(.) = [g_{l1}(.), \cdots, g_{ln}(.)]^\top \sim \mathcal{N}\left(0, \frac{1}{\lambda_P} K_P\right) \qquad (8)$$

where K_M and K_P are Gram matrices obtained by evaluating the Gaussian process covariance functions on the row and column index sets. Let $\mathrm{Z} = [\mathrm{f}_1(.)\ \mathrm{f}_2(.)\ \cdots \mathrm{f}_r(.)] \in \mathrm{R}^{\mathrm{M}\times\mathrm{r}}$ and $\mathrm{W} = [\mathrm{g}_1(.)\ \mathrm{g}_2(.)\ \cdots \mathrm{g}_r(.)] \in \mathrm{R}^{\mathrm{P}\times\mathrm{r}}$. Each column is a GP draw evaluated at all row indices, and columns are independent across latent dimensions l. Thus, the prior over the whole matrix Z is given by

$$p(Z) = \Pi_{i=1}^{r} \mathcal{N}(f_i | 0, K_M) \propto exp\left(-\frac{1}{2} \sum_{l=1}^{r} f_l^\top K_M f_l \right) \qquad (9)$$

Similarly, the prior over the whole matrix W is given by

$$p(W) = \Pi_{i=1}^{r} \mathcal{N}(g_i | 0, K_P) \propto exp\left(-\frac{1}{2} \sum_{l=1}^{r} g_l^\top K_P g_l \right) \qquad (10)$$

Now, the MAP objective with the GP prior on Z and W becomes [6]:

$$\mathcal{L}_{\mathrm{KPMF}}(Z, W) = \sum_{(m,p)\in\Omega} \left(X_{mp} - z_m^\top w_p \right)^2 + \lambda_M tr\left(Z^\top K_M^{-1} Z \right) + \lambda_P tr\left(W^\top K_P^{-1} W \right) \qquad (11)$$

It can be observed that (eq 11) is a bi-convex function in (Z, W), enabling the solution to the matrix completion problem using the alternating minimization technique. While KPMF can successfully incorporate side information on materials and properties, it has limitations such as the regularization term requires computation of K_M^{-1}, which is computationally expensive for large M and can be numerically unstable for ill-conditioned kernels, and optimization over Z and W cannot fully exploit the kernel structure.

Parametrization of Kernelized Probabilistic Matrix Factorization

We now present the Parameterization of Kernelized Probabilistic Matrix Factorization (P-KPMF) model, to addresses the limitations of KPMF through: (i) kernel regularization on both materials and properties, and (ii) a reparameterization that avoids explicit kernel matrix inversion. Inspired by the representer theorem from kernel methods, we parameterize the latent factors as linear combinations of kernels

$$Z = K_M A, \qquad W = K_P B \qquad (12)$$

where $\mathrm{A} \in \mathrm{R}^{\mathrm{M}\times\mathrm{r}}$ and $\mathrm{B} \in \mathrm{R}^{\mathrm{P}\times\mathrm{r}}$ are coefficient matrices to be learned and we place Gaussian priors on the columns of A and B:

$$a_l \sim \mathcal{N}\left(0, \frac{1}{\lambda_M} K_M^{-1}\right), \quad b_l \sim \mathcal{N}\left(0, \frac{1}{\lambda_P} K_P^{-1}\right), \quad l = 1, \ldots, r \qquad (13)$$

Notice that each column of matrices Z and W are

realizations of a zero-mean Gaussian Process with covariance functions k_m and k_p, respectively. The resulting MAP objective is:

$$\mathcal{L}_{\text{P-KPMF}}(A,B) = \sum_{(m,p)\in\Omega} \left(X_{mp} - k_m^\top A B^\top k_p\right)^2 + \lambda_M tr\left(A^\top K_M A\right) + \lambda_P tr\left(B^\top K_P B\right) \quad (14)$$

Define the residual for j^{th} latent as $r_{mp}^{(j)} = X_{mp} - \sum_{l\neq j}(k_p^\top b_l)(k_m^\top a_l)$

$$\mathcal{L}_{\text{P-KPMF}}(a_j,b_j) = \sum_{(m,p)\in\Omega} \left(r_{mp}^{(j)} - (k_p^\top b_j)(k_m^\top a_j)\right)^2 + \lambda_M tr\left(A^\top K_M A\right) + \lambda_P tr\left(B^\top K_P B\right) \quad (15)$$

Similar to (eq 11), the above objective function (eq 15) is also bi-convex; however, it is expressed in terms of Gram matrices K_M and K_P without requiring inverse computation. Owing to the bi-convexity, we can use alternating minimization to solve for the parameterization matrices A and B. Thus, the first derivative of the objective function with respect to a_l and b_l yields

$$\left[\sum_{(m,p)\in\Omega} \left(k_p^\top b_j\right)^2 k_m k_m^\top + \lambda_M K_M\right] a_j = \sum_{(m,p)\in\Omega} r_{mp}^{(j)} \left(k_p^\top b_j\right) k_m \quad (16)$$

$$\left[\sum_{(m,p)\in\Omega} \left(k_m^\top a_j\right)^2 k_p k_p^\top + \lambda_P K_P\right] b_j = \sum_{(m,p)\in\Omega} r_{mp}^{(j)} \left(k_m^\top a_j\right) k_p \quad (17)$$

Now, we present the alternating minimization-based solution methodology for imputing the missing observations using the proposed parameterized kernelized probabilistic matrix factorization. The detailed methodology is presented in Algorithm 1. This iterative approach involves two steps: in the first step, we fix b_j and solve for a_j using (16) and in the second step, we fix a_j and solve for b_j using (17). Notice that fixing one variable, result in a convex optimization problem. Thus, the alternating minimization enables one to solve the imputation problem.

Algorithm 1: Alternating Minimization for P-KPMF

Initialize a_l, b_l, K_M, K_p
Set hyperparameters r, λ_M, λ_P
Set tolerance ϵ, maximum iterations T
for $t = 1, \ldots, T$ **do**
for $j = 1, \ldots, r$ **do**
  (A-step) Fix b_j
  Solve for a_j using (16)
  (B-step) Fix a_j
  Solve for b_j using (17)
end for
Compute $\mathcal{L}_{P-KPMF}^{t}$ via (15)
if $|\mathcal{L}_{P-KPMF}^{t-1} - \mathcal{L}_{P-KPMF}^{t}|/|\mathcal{L}_{P-KPMF}^{t}| < \epsilon$ **then**
  break
end if
end for

Screening of Adsorbent Candidates

The primary objective of the proposed framework in a recommendation system is to rank candidate materials. Therefore, the model performance will also be evaluated using rank-based metrics in addition to regression error measures. To this end, we use Spearman's rank correlation coefficient (ρ_s), which quantifies the monotonic agreement between the imputed and true material properties. The Spearman rank correlation coefficient is defined as:

$$\rho_s = 1 - \frac{6\sum_{i=1}^{n} d_i^2}{n(n^2-1)} \quad (18)$$

where d_i denotes the difference between the ranks of the corresponding imputed and observed values for the ith COF, and n denotes the total number of COFs. A higher value of Spearman correlation indicates the true ranking is preserved by the proposed imputation model, and a lower value of Spearman correlation indicates that the ordering of COFs obtained by the imputation model is significantly different from the true order of COFs. Thus, high values of Spearman correlation signify that the recommendation system can provide a reliable ranking of COFs for screening.

RESULTS AND DISCUSSION

In this section, we demonstrate the proposed parameterization of the kernelized probabilistic matrix factorization approach to learn the nonlinear surface from the adsorption properties of COF materials when there are missing data. In this work, we use the adsorption data of 573 COF candidates taken from AiiDA platform [3] for 16 features which includes uptake values and Henry Coefficients for CH_4, CO_2, H_2, H_2O, H_2S, Kr, N_2, O_2, Xe gas species (same as in the work of [5]). We performed preprocessing of the data using logarithmic transformation for features that have values in the range (10^{-5} to 10^{-7}).

For fair comparison, we consider the same missing data points for performing imputation using the three models: Linear Matrix Factorization (LMF) - studied by [5], Kernelized Probabilistic Matrix Factorization (KPMF), and the proposed parameterized form of Kernelized Probabilistic Matrix Factorization (P-KPMF). To systematically evaluate different models, we consider imputation performance across a range of missingness fractions between 10% to 90%. For each missingness, we generate 20 independent Missing Completely At Random (MCAR) experiments, yielding a total of 180 experimental runs per method. These experiments enable statistical

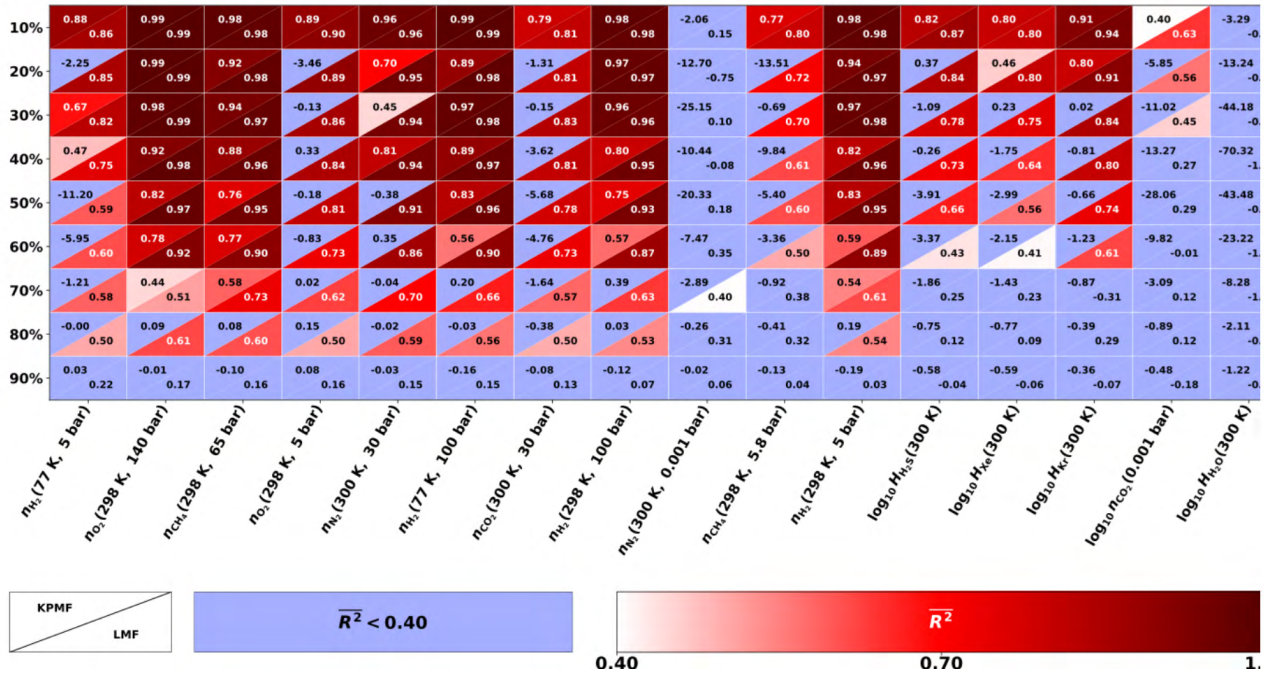

Figure 1: Heatmap showing the comparison of the imputation performance of all 16 adsorption properties for varying percentages of missingness using LMF (bottom triangle) and KPMF (top triangle), respectively.

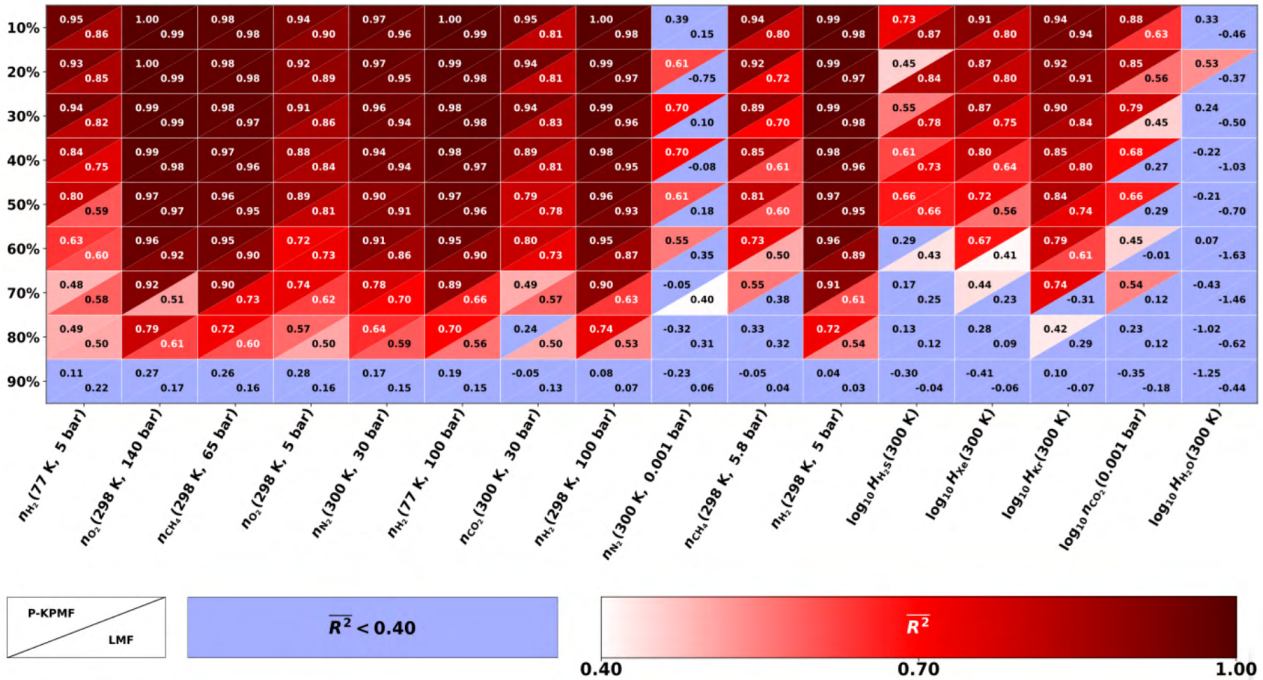

Figure 2. Heatmap showing the comparision of the imputation performance of all 16 adsorption properties for varying percentage of missingness using LMF (bottom traingle) and P-KPMF (upper traingle), respectively.

characterization of performance variability attributable to the increase in the percentage of missingness for various imputation models under consideration. Hyperparameters are optimized using grid search, and we use radial basis function kernels in this study.

Figure 1 presents the heatmap showing the comparison of imputation performance (that is, coefficient of determination R^2) of all 16 adsorption properties over 20

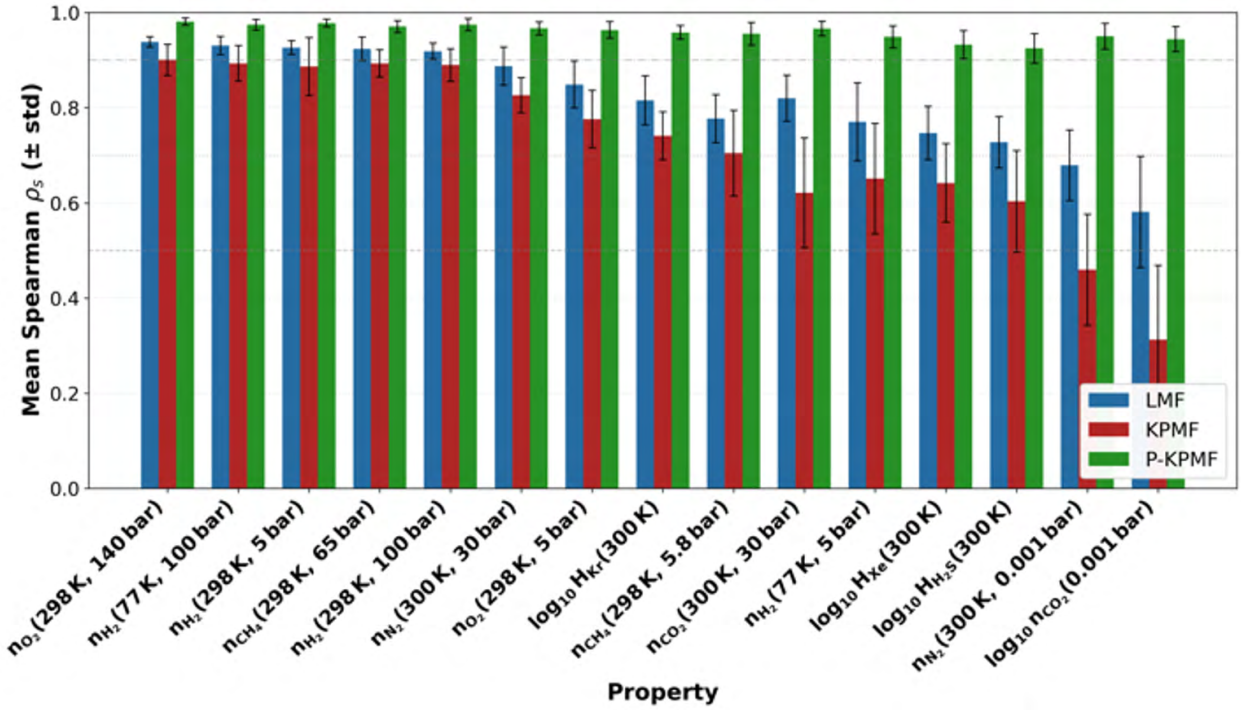

Figure 3. Bar plot showing the comparison of mean Spearman rank correlation obtained using LMF, KPMF, and P-KPMF models.

MCAR runs for varying percentages of missingness using LMF (bottom triangle) and KPMF (top triangle), respectively. The imputation performance of the LMF model shows good predictions in 13 of the 16 adsorption properties. The adsorption properties such as uptake values of H_2 and CO_2 (at 300 K and 0.001 bar), and Henry coefficient of H_2O at 300 K cannot be predicted using both models. It can be observed that the imputation performance of the KPMF model is very poor, as can be seen from the blue colors in the top triangles. This can be attributed to the numerical instability in determining the inverse of the Gram matrices K_P and K_M. Thus, it can be observed LMF performs better than the non-parameterized KPMF model.

Figure 2 presents the heatmap showing the comparison of imputation performance (that is, coefficient of determination R^2) of all 16 adsorption properties over 20 MCAR runs for varying percentages of missingness using LMF (bottom triangle) and P-KPMF (top triangle), respectively. Notice that the predictive performance of 15 out of 16 adsorption properties can be determined accurately by the proposed parameterization of the KPMF model. More specifically, the adsorption properties such as uptake values of H_2 and CO_2 (at 300 K and 0.001 bar) can be accurately predicted up to a missing percentage of 60%. Recall that these properties were poorly predicted by the LMF model. Further, it can be observed that the overall prediction accuracy of the remaining 15 properties shows improved performance using the P-KPMF model over the LMF model. However, Henry coefficient of water cannot predict accurately using both models. This improvement in performance can be attributed to the numerical stability of the proposed parameterization of PKMF, and use of kernel methods to capture the nonlinear relationship among different adsorption properties. Thus, the proposed P-KPMF model can provide accurate imputations even when there is a missing percentage of 60%.

Figure 3 presents the bar plot showing the comparison of mean Spearman rank correlation obtained for 60% missing data using LMF, KPMF, and P-KPMF models. Standard deviation values in predictions across 20 MCAR runs for the 15 properties are also shown in the bar chart. As the predictions for the Henry coefficient of H_2O are poor using all methods, we have omitted it. Among all models, our proposed parametrization of kernelized probabilistic matrix factorization yields a Spearman rank correlation of 0.9, which indicates that the predicted ranks are close to the true ranks. This highlights that our proposed P-KPMF-based recommendation system will yield reliable COF candidates for a given task. To illustrate, we have considered the screening of best-performing COFs that achieves higher uptake values for CO_2. Table presents the COF IDs as obtained from Materials Cloud repositry of the best 3 COFs selected for CO_2 uptake at 30 bar and 300 K using LMF, KPMF, and P-KPMF

models. True indicate the COF IDs with three highest CO_2 uptake values at 30 bar and 300 K reported in the Materials Cloud repositry. The proposed P-KPMF-based recommendation system can identify all three best COFs showing superior performance over the LMF and KPMF models.

Table 1: Top 3 COF IDs for n_{CO_2}(300 K, 30 bar).

Methods	COF ID
LMF	20116N3, **18143N3, 20430N3**
KPMF	20116N3, **18143N3, 20430N3**
P-KPMF	**18142N3, 20430N3, 18143N3**
True	18142N3, 18143N3, 20430N3

CONCLUSIONS

In this work, we presented a systematic approach to develop a recommendation system that is based on the parameterization of kernelized probabilistic matrix factorization. This parameterization provided the accurate determination of the material properties by eliminating the inverse of computation of Gram matrices, which improved numerical stability. Also, the use of Gaussian process-based latent functions effectively captured the nonlinear relationship. The resulting bi-convex MAP objective with regularization on the latent function has been efficiently solved using an alternating minimization approach. We demonstrated that this proposed approach showed superior predictive performance of adsorption properties over the linear matrix factorization technique, and also preserves the ordering for the effective screening of COF candidates up to 60% missing data. This highlights the significance of the proposed P-KPMF-based recommendation system can predict the adsorption properties and screening the candidates under a data-scarce scenario. Future work will explore the multi-kernel or deep kernel learning strategies that integrate heterogeneous material descriptors, such as geometric, chemical, and topological features, within a unified latent representation.

REFERENCES

1. Furukawa H, Cordova KE, O'Keeffe M, Yaghi OM. The chemistry and applications of metal-organic frameworks. Science 341: (2013). https://doi.org/10.1126/science.1230444
2. Co?te? AP, Benin AI, Ockwig NW, O'Keeffe M, Matzger AJ, Yaghi OM. Porous, crystalline, covalent organic frameworks. Science 310:1166-1170 (2005). https://doi.org/10.1126/science.1120411
3. Ongari D, Yakutovich AV, Talirz L, Smit B. Building a consistent and reproducible database for adsorption evaluation in covalent–organic frameworks. ACS Cent. Sci. 5:1663-1675 (2019). https://doi.org/10.1021/acscentsci.9b00619
4. Talirz L, Kumbhar S, Passaro E, Yakutovich AV, Granata V, Gargiulo F, Borelli M, Uhrin M, Huber SP, Zoupanos S, Adorf CS, Andersen CW, Schütt O, Pignedoli CA, Passerone D, VandeVondele J, Schulthess TC, Smit B, Pizzi G, Marzari N. Materials cloud, a platform for open computational science. Sci Data 7: (2020). https://doi.org/10.1038/s41597-020-00637-5
5. Sturluson A, Raza A, McConachie GD, Siderius DW, Fern XZ, Simon CM. Recommendation system to predict missing adsorption properties of nanoporous materials. Chem. Mater. 33:7203-7216 (2021). https://doi.org/10.1021/acs.chemmater.1c01201
6. Zhou T, Shan H, Banerjee A, Sapiro G. Kernelized probabilistic matrix factorization: exploiting graphs and side information. Proceedings of the 2012 SIAM International Conference on Data Mining :403-414 (2012). https://doi.org/10.1137/1.9781611972825.35
7. Soleimani B, Niknam Shahrak M, Walton KS. The influence of different functional groups on enhancing CO2 capture in metal-organic framework adsorbents. Journal of the Taiwan Institute of Chemical Engineers 163:105638 (2024). https://doi.org/10.1016/j.jtice.2024.105638

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimizing MIP-Heuristics: Generic Formulation and Code

Sophie Hildebrandt[ab], Meik Franke[b], Edwin Zondervan[b] and Guido Sand[a]*

[a] Pforzheim University of Applied Science, School of Engineering and Technology, Pforzheim, Germany
[b] University of Twente, Faculty of Science and Technology, Enschede, The Netherlands
* Corresponding Author: guido.sand@hs-pforzheim.de

ABSTRACT

Large-scale mixed-integer programs (MIPs) typically cannot be solved by standard solvers with reasonable computational cost. MIP-heuristics decompose large-scale monolithic mixed-integer programs into polylithic programs such that they can be solved with reasonable computational cost at the price of loosing their optimality certificate. The decomposition is steered by hyperparameters that impact the solution quality and the computational cost diametrically. The proper selection of the hyperparameter values is a black-box optimization problem which is mostly solved by grid search or random search. In previous publications the authors proposed a novel hyperparameter optimization method based on Bayesian optimization and studied a use case from the PSE domain. Computational studies showed that the BO-based algorithm is superior for objective functions with few optimal solutions.This contribution generalizes the description of the MIP-Heuristic Optimization Problem (MIP-HOP) and the computer implementation of the solution-method. It supports the transferability of the method to other use cases. The general mathematical description of the MIP-HOP comprises eight expressions covering the monolithic and the polylithic MIP-model, the heuristic decomposition and composition operator, and the MIP-HOP. The formulation is explained and linked to the computer implementation of the solution method. The Python-code integrates state-of-the-art tools for mathematical programming (Pyomo and CPLEX) and Bayesian optimization (SMAC3) and will be published on GitHub.

Keywords: Algorithms, Derivative Free Optimization, Machine Learning, Surrogate Model, MIP-Heuristics

INTRODUCTION

Mixed-integer programming (MIP) is a methodology to model and solve optimization problems with mixed (integer and continuous) degrees of freedom subject to complex constraints. Exact MIP-algorithms provide proven optimal solutions or feasible solutions along with conservative optimality certificates ("optimality gaps"). However, many real-world optimization problems are NP-complete such that their computational costs scale poorly with their sizes. A common approach to deal with large-scale problem instances is to heuristically decompose a monolithic MIP-model into a polylithic MIP-model consisting of a set of partial MIP-models of smaller size. By sequentially solving the set of partial MIP-models by exact MIP-algorithms the polylithic MIP-model can be solved faster than the monolithic MIP-model at the price of losing its optimality certificate.

Schemes that decompose monolithic MIP-models into polylithic MIP-models based on problem insights are a special type of matheuristics named MIP-heuristics. They typically comprise hyperparameters that steer the decomposition and have an impact on both, the computational cost for solving the polylithic MIP-model and the objective value of its solution. MIP-heuristics as well as MINLP-heuristics are applied to a wide range of optimization problems, such as batch scheduling, maintenance planning, project scheduling, supply network planning, production routing, bin packing, locomotive assignment, healthcare planning, flowsheet optimization and lot sizing. The MIP/MINLP-heuristics comprise different decomposition approaches like temporal, area, activity, scenario and geographic decomposition. State-of-the art methods for tuning the hyperparameters of the MIP-heuristics are limited to grid search, random search and heuristic search; these methods suffer from their computational costs for large hyperparameter spaces (grid search and random search) or explore only a heuristically reduced hyperparameter space (heuristic search).

In previous publications the authors proposed a

novel hyperparameter optimization method based on Bayesian optimization (BO) and studied a use case from the PSE domain [7, 8]. The use case is a two-step moving horizon heuristic for a hoist scheduling problem comprising a construction step and an improvement step. The six hyperparameters are the window size and the step size in each step, the initial job sequence and the accepted optimality gap of the solutions of the polylithic MIP-models. The contribution at hand generalizes the description of the MIP-Heuristic Optimization Problem (MIP-HOP) and the computer implementation of the solution-method.

The remainder of the paper is structured as follows: The following section "Generic Formulation" describes and explains the MIP-HOP in detail. The section "Related Work" summarizes existing research on tuning of MIP-heuristics, hyperparameter tuning methods and the authors' own work in this field. The section "Software Architecture" describes the implementation of the BO-based approach for solving the MIP-HOP with references to the generic formulation. "Conclusions" and a link to "Digital Supplementary Material" close the paper.

GENERIC FORMULATION

The following large scale monolithic MIP-model M (describing a hoist scheduling problem for instance, see above) is given:

$$M: z = \min_{x \in \mathbb{R}^{n1} \times \mathbb{N}^{n2}} cx \mid Ax \leq d \tag{1}$$

In (1) z denotes the scalar objective value, x denotes the vector of variables, $n1$ and $n2$ denote the number of continuous and integer variables, c denotes the parameter-vector of the linear objective function cx, and A denotes the parameter-matrix and d denotes the parameter-vector of the linear constraints $Ax \leq d$. A feasible or optimal solution of M is denoted as $\hat{x}$.

The decomposition scheme of the monolithic MIP-model M into a polylithic MIP-model $P(\lambda)$ according to a MIP-heuristic is also given and described by a decomposition operator $\mathcal{D}$:

$$P(\lambda) = \mathcal{D}(M, \lambda) \tag{2}$$

In (2) the λ denotes the vector of hyperparameters of a MIP-heuristic. For the hoist scheduling example, the hyperparameters "window size" control the size of the partial MIP-models and the hyperparameters "step size" control the number of the partial MIP-models. It is assumed that the hyperparameters can be represented by real or integer values:

$$\lambda \in \mathbb{R}^{h1} \times \mathbb{N}^{h2} \tag{3}$$

In (3) $h1$ and $h2$ denote the number of real-valued and integer-valued hyperparameters, respectively. Their values may be subject to a vector of (non-linear) constraints:

$$g(\lambda) \leq 0 \tag{4}$$

For the hoist scheduling example, the step size must not be greater than the window size.

The polylithic MIP-model $P(\lambda)$ consist of a set of $n(\lambda)$ partial MIP-models P_i^λ:

$$P(\lambda) = \{P_1^\lambda, P_2^\lambda, \dots, P_i^\lambda, \dots, P_{n(\lambda)}^\lambda\} \tag{5}$$

Each partial MIP-model P_i^λ is defined as shown below and typically much smaller than the large-scale monolithic MIP-model M:

$$P_i^\lambda: z_i^\lambda = \min_{x_i^\lambda \in \mathbb{R}^{n1_i^\lambda} \times \mathbb{N}^{n2_i^\lambda}} c_i^\lambda x_i^\lambda \mid A_i^\lambda x_i^\lambda \leq d_i^\lambda \tag{6}$$

In (6) z_i^λ denotes the scalar objective value, x_i^λ denotes the vector of variables, $n1_i^\lambda$ and $n2_i^\lambda$ denote the number of continuous and integer variables, c_i^λ denotes the parameter-vector of the linear objective function $c_i^\lambda x_i^\lambda$, and A_i^λ denotes the parameter-matrix and d_i^λ notes the parameter-vector of the linear constraints $A_i^\lambda x_i^\lambda \leq d_i^\lambda$.

The composition of the partial solutions $\hat{x}_i^\lambda$ according to the given MIP-heuristic to a solution $\hat{x}^\lambda$ is described by a composition operator $\mathcal{C}$:

$$\hat{x}^\lambda = \mathcal{C}(M, \lambda, \hat{x}_1^\lambda, \hat{x}_2^\lambda, \dots) \tag{7}$$

Finally, the MIP-heuristic optimization problem (MIP-HOP) can be stated as the following multi-objective optimization problem with mixed degrees of freedom subject to three types of constraints:

$$\begin{aligned} &MIP-HOP: \hat{\lambda}(f_1, f_2, \dots) = \\ &\underset{\lambda}{argmin}(f_1(\lambda), f_2(\lambda), \dots) \mid g(\lambda) \leq 0, P(\lambda) = \\ &\mathcal{D}(M, \lambda), \hat{x}^\lambda = \mathcal{C}(M, \lambda, \hat{x}_1^\lambda, \dots, \hat{x}_{n(\lambda)}^\lambda) \end{aligned} \tag{8}$$

In (8) $f_1(\lambda), f_2(\lambda), \dots$ denote general scalar objective functions and $f_1, f_2, \dots$ denote the corresponding objective values.

The MIP-HOP typically exhibits the following properties:

- The exact values of the general objective functions $f_1(\lambda), f_2(\lambda), \dots$ for a certain value of λ can only be observed a posteriori, i.e. after a $P(\lambda)$ is solved.
- Two natural objective functions $f_1^n(\lambda)$ and $f_2^n(\lambda)$ of the *MIP-HOP* are the objective function of the monolithic MIP-model M and the computational cost for solving $P(\lambda)$, respectively.
- The optimal objective value of M is a lower bound of the exact $f_1^n(\lambda)\ \forall\lambda$.
- Exactly solving M (and possibly $P(\lambda)$ for some values of λ) is prohibited by too high

computational costs.

Exact MIP-algorithms can solve the monolithic MIP-model M and the partial MIP-models P_i^λ exactly (feasible solution with optimality gap 0), or they can approximate their solutions (feasible solutions with optimality gap >0). In this paper, a solution of a polylithic MIP-model $P(\lambda)$ is called "exact" if all partial MIP-models P_i^λ are solved to optimality. Exact MIP-algorithms are steered by special hyperparameters (e.g. an accepted optimality gap); these special hyperparameters are included in λ.

The MIP-model M and the partial MIP-models P_i^λ are "white-box" optimization problems since the mathematical formulation is fully accessible to the solver at any time of the solution process. Due to the algorithmic nature of the MIP-heuristic, in general no such formulation exists for $\mathcal{C}$ and $\mathcal{D}$, such that the MIP-HOP belongs to the class of "black-box" optimization problems.

RELATED WORK

Tuning of MIP-Heuristics

A search in Google scholar for "MIP-heuristic" leads to around 700 publications from the years 1965 to 2025; some 100 of them were published since 2024. To gain a deeper insight into the state of the art of tuning of MIP-heuristics 30 publications were selected and analyzed. The publications comprise MIP-based as well as MINLP-based decomposition heuristics for a wide range of optimization problems, such as batch scheduling, maintenance planning, project scheduling, supply network planning, production routing, bin packing, locomotive assignment, healthcare planning, flowsheet optimization and lot sizing. The MIP/MINLP-based heuristics comprise different decomposition approaches like temporal, area, activity, scenario and geographic decomposition. The selected papers were published between 2002 and 2025, and include 26 journal articles, 3 conference papers, and 1 book chapter.

A majority of 25 of the 30 reviewed publications (83%) exhibits MIP/MINLP-based heuristics that include at least one hyperparameter. These hyperparameters are not in all papers explicitly mentioned in the text but were identified during the review by closely analyzing the proposed heuristics. Only 12 publications (40%) report some form of tuning: 9 use grid search while 3 rely on hyperparameter values derived from random or heuristic sampling (based on pre-studies). The remaining 13 publications with hyperparameters do not describe any tuning procedures while 2 of them devise the tuning to the user. In the remaining 5 publications no hyperparameter could be identified. A full list of the reviewed papers along with the detailed analysis will be provided in [10].

The result from this literature review shows that most MIP-heuristics exhibit tunable hyperparameters and roughly half of the authors identified the tuning of hyperparameters as an optimization problem. However, the applied tuning methods still leave space for methodological improvement: Grid search as the predominant tuning method requires the continuous hyperparameters to be discretized. It finds the global optimum in the discretized hyperparameters space, but its computational cost suffers from its sample-inefficiency in higher dimensional hyperparameter-spaces. Random and heuristic sampling typically use less samples but do not explore the hyperparameter space effectively.

Hyperparameter Tuning Methods

The application of hyperparameter tuning and algorithm selection methods to mixed-integer programming was recently reviewed by the authors in [9]; it can be summarized as follows:

Numerous studies have demonstrated that proper hyperparameter tuning can enhance machine learning model accuracy, generalization and performance, as well as prevent overfitting (see e.g. [1-3]). Beyond traditional machine learning models, hyperparameter tuning and algorithm selection techniques have also been applied to mathematical programming (see e.g. [4-6]). However, there is no research on its use in MIP-heuristics known to the authors. A key challenge in MIP-heuristic tuning is the high computational cost associated with evaluating each hyperparameter value, as it involves solving one or more MIPs. Consequently, exhaustive methods such as grid search, random search, or even population-based approaches like evolutionary algorithms or particle swarm optimization often become impractical due to their high computational cost and low sample efficiency. To address this issue, more efficient methods such as successive halving, hyperband, and BO have been proposed. Both bandit-based methods (like successive halving and hyperband) and BO-based methods are promising approaches for the hyperparameter tuning of matheuristics.

Own Work

The authors proposed a novel BO-based algorithm based to solve the MIP-HOP (see (1)-(8)) and studied it for a hoist scheduling problem [7-10]. The description of the MIP-heuristic and hyperparameters can be found in [9]. The architecture of the algorithmic is shown in Figure 1.

In each iteration, the BO proposes a new point in the space of hyperparameters λ using a surrogate model for the objective functions $f_1(\lambda), f_2(\lambda), \ldots$ and an acquisition function such that the hyperparameter constraints (4) are met. The decomposition operator $\mathcal{D}$ and the composition operator $\mathcal{C}$ represent the MIP-heuristic. The decomposition operator $\mathcal{D}$ uses the λ-values and the formulation of

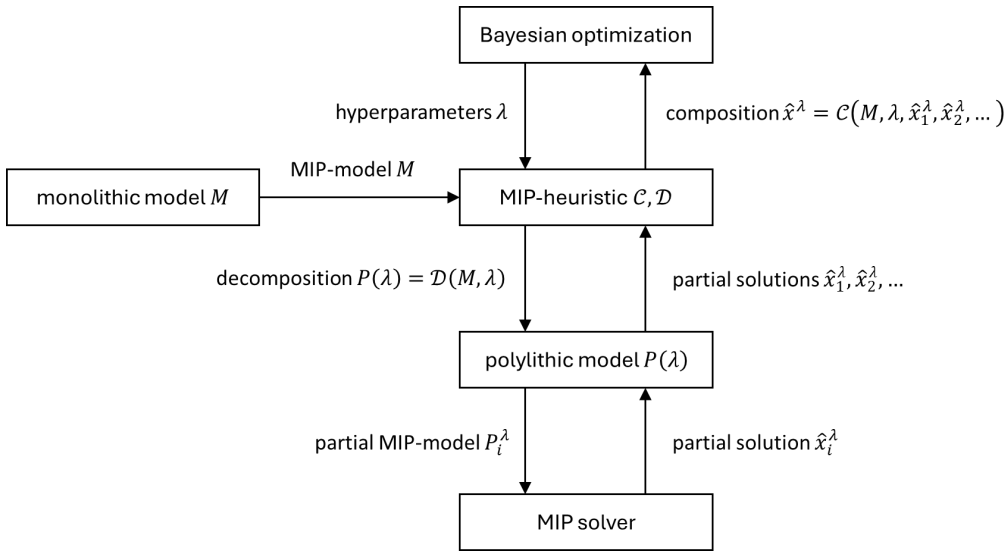

Figure 1. Algorithmic architecture

the monolithic MIP-model M as input to generate the polylithic MIP-model $P(\lambda)$ according to (2). The polylithic MIP-model $P(\lambda)$ consists of a set of interconnected partial MIP-models P_i^λ; they are solved sequentially by a MIP standard solver. The MIP standard solver computes a solution $\hat{x}_i^\lambda$ of the partial MIP-model P_i^λ which is used to parametrize the next partial MIP-model P_{i+1}^λ. The composition operator $\mathcal{C}$ uses the λ-values, the formulation of the monolithic MIP-model M and the solutions $\hat{x}_1^\lambda, \hat{x}_2^\lambda, \dots$ as input to generate a solution $\hat{x}^\lambda$ of M. The objective functions $f_1(\lambda), f_2(\lambda), \dots$ are evaluated and the objective values $f_1, f_2, \dots$ are returned to the BO. The BO updates its surrogate model and calculates new promising values for the hyperparameters λ using an acquisition function. This is done until a termination criterion for the BO, such as the number of iterations, a computational time limit or the number of retries to propose new λ-values is met.

SOFTWARE ARCHITECTURE

The proposed approach is implemented in python within a modular software architecture that separates MIP-heuristic execution $\mathcal{D}, \mathcal{C}$, MIP-model definition M, and hyperparameter tuning of the MIP-heuristic (MIP-HOP). This separation enables the individual components e.g. the MIP-model to be exchanged or extended independently.

The software architecture supports multiple solving strategies for the underlying MIP-model: Monolithic solving with adapted solver parameters λ like the target optimality gap and polylithic solving based on MIP-heuristics $\mathcal{D}, \mathcal{C}$ that decompose the monolithic model M into a sequence of partial MIP-models P_i^λ are implemented. In the polylithic case, the MIP-heuristic is parameterized by a set of hyperparameters λ that control, among others, the number $n(\lambda)$ and size of the generated partial MIP-models and thus define the trade-off between the two natural, competing objectives of the MIP-HOP, the solution quality $f_1^n(\lambda)$ and computational cost $f_2^n(\lambda)$. The hyperparameters λ can either be evaluated using a single fixed hyperparameter configuration or optimized automatically using hyperparameter tuning.

Several hyperparameter tuning methods are implemented, including grid search and random search as baseline methods, as well as Bayesian optimization for more efficient exploration of the non-linear hyperparameter spaces. All tuning methods interact with the MIP-heuristic through evaluations of $P(\lambda)$ and the resulting objective values $f_1(\lambda), f_2(\lambda), \dots$, and thus treat the MIP-heuristic as a black-box objective function.

The implementation integrates several external optimization tools. The MIP-model M and all partial MIP-models P_i^λ are formulated using the algebraic modeling language Pyomo, while all MIPs are solved using the commercial MIP-solver CPLEX. Bayesian optimization for solving the MIP-HOP is realized using SMAC3 [11].

Figure 2 illustrates the main components of the software architecture (`experiment setup`, `experiment controller`, `heuristic execution`, `model definition` and external optimization tools) as well as the control flow between them.

Figure 2 illustrates the main components of the software architecture (`experiment setup`, `experiment controller`, `heuristic execution`, `model definition` and `external optimization tools`) as well as the control flow between them.

The `experiment setup` specifies the hyperparameters λ and selects the solving strategy, such as

monolithic solving or polylithic solving combined with random search. The `experiment controller` then orchestrates the execution by invoking either the monolithic solving strategy or the MIP-heuristic and, if selected, executes the hyperparameter tuning procedure to solve the MIP-HOP. The `heuristic execution` and `model definition` implement the problem-specific logic of the underlying application, such as hoist scheduling. The following subsections describe these components in detail, following the order in which they interact during an experimental run.

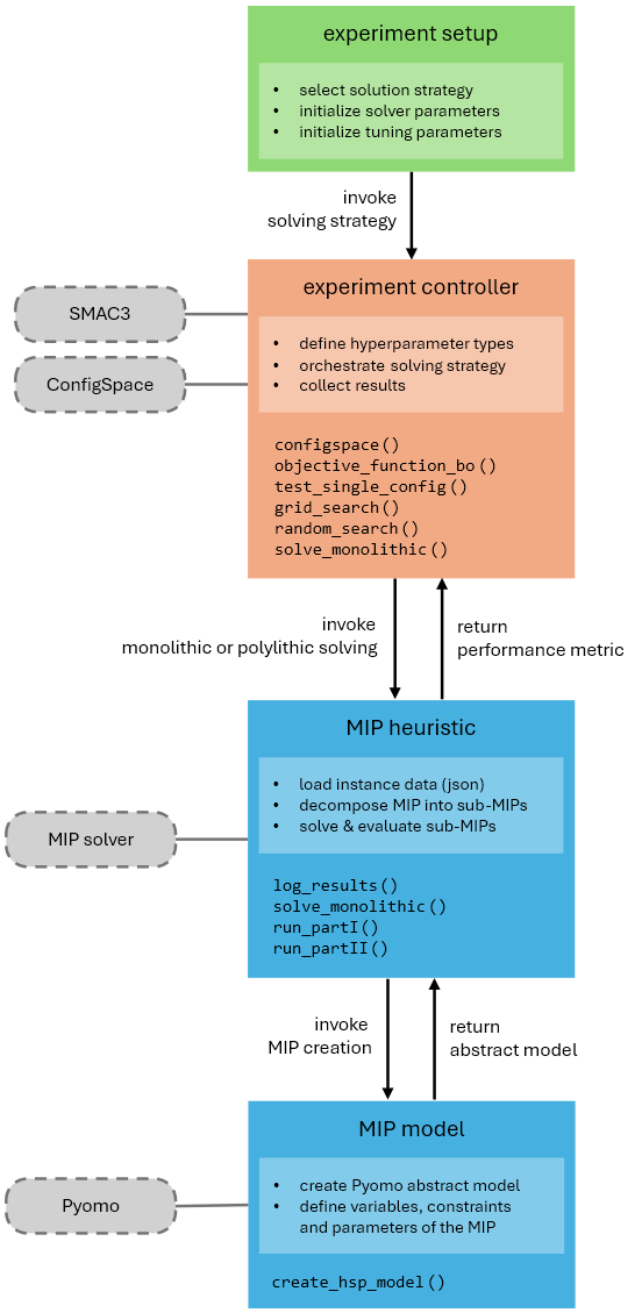

Figure 2. Software architecture

Experiment setup

The `experiment setup` is implemented in `main.py` and serves as the central entry point of the software architecture. Its primary responsibility is to configure and initiate experimental runs and delegate their execution to the `experiment controller`.

At the beginning of an experimental run the `experiment setup` loads the problem instance data from a JSON file and forwards the instance-specific information to the `experiment controller`. These data remain fixed throughout an experimental run. In addition, `experiment setup` initializes output structures for storing results and metadata, supporting reproducibility and systematic comparison of experiments.

`Experiment setup` defines solver-related parameters. These include parameters such as time limits, number of threads, clock type, and target optimality gap. The parameters are passed to the `experiment controller` and subsequently applied during solution of the MIP-model.

The solving strategy selected in `experiment setup` determines whether the `experiment controller` evaluates a monolithic or polylithic MIP-model. In the monolithic case, decomposition-specific parameters are inactive, while parameters, such as the initial job sequence (see [9]), can still be provided explicitly. For polylithic solving, `experiment setup` specifies whether a single hyperparameter configuration λ is evaluated or whether automated hyperparameter tuning is applied. In the latter case, `experiment setup` selects the tuning strategy (grid search, random search, or Bayesian optimization) and defines termination criteria such as a time budget or a maximum number of evaluations. When Bayesian optimization is selected, `experiment setup` additionally specifies hyperparameters of the Bayesian optimization, such as the choice of acquisition function and surrogate model, and delegates the execution of the MIP-heuristic hyperparameter tuning to the `experiment controller`.

Experiment controller

The `experiment controller`, implemented in `experiment_controller.py`, is responsible for orchestrating the execution of the selected solving strategy. It acts as the central coordination component that links experiment setup, heuristic execution, model solving, and hyperparameter tuning.

For monolithic solving, the `experiment controller` delegates execution to the monolithic solving routine where the full MIP-model M is instantiated and solved. The resulting solution quality $f_1^n(\lambda)$, computational cost $f_2^n(\lambda)$, and solver status are collected and passed to

the logging routines.

For heuristic-based solving, the `experiment controller` instantiates and executes the MIP-heuristic $\mathcal{D}$. Depending on the selected hyperparameter tuning method, it either evaluates a single fixed hyperparameter configuration or systematically explores multiple configurations using grid search, random search, or Bayesian optimization.

For Bayesian optimization the type of the hyperparameters to be tuned and the constraints between them are defined using `ConfigSpace`. Based on the defined configuration space the `experiment controller` employs the SMAC3 framework to iteratively select and evaluate hyperparameter configurations. Each hyperparameter configuration proposed by SMAC3 is evaluated by executing the MIP-heuristic $\mathcal{D}$, which is treated as black-box objective function returning performance metrics such as solution quality $f_1^n(\lambda)$ and computational cost $f_2^n(\lambda)$.

For heuristic execution of the hoist scheduling problem, the experiment controller manages the two-phase structure of the rolling horizon approach described in [7]-[9]. It first triggers the initial construction phase (Part I). If this phase completes successfully and the overall time limit is not exhausted, the `experiment controller` initiates the improvement phase (Part II). If a termination criterion is reached within the construction phase, a composition $\mathcal{C}$ is required to ensure a feasibility of the resulting schedule. In this case the remaining unscheduled jobs are appended by the `experiment controller`.

By design, the `experiment controller` interacts with the heuristic exclusively through its external interface and does not modify its internal logic or the structure of the MIP model. Its role is limited to configuration, orchestration, and result collection, ensuring a clean separation between control flow and algorithmic implementation.

MIP model and heuristic execution

The mathematical model and the heuristic execution are implemented in HSP_MIP_model.py and HSP_MIP_heuristic.py, respectively.

The MIP-model is defined as a Pyomo abstract model, encapsulating decision variables, constraints, and the objective function independently of any specific data instance. A factory function returns this abstract model, which is instantiated at runtime using instance data provided via JSON files. This design allows the same model definition to be reused for monolithic solving of M as well as for all heuristic subproblems P_i^λ.

The `HeuristicSolving` class provides methods for:

Solving the full monolithic MIP-model

Executing Part I of the rolling horizon heuristic, which incrementally constructs an initial schedule by solving a sequence of partial MIP-models, and

Executing Part II of the rolling horizon heuristic, which iteratively improves the initial solution by rescheduling subsets of jobs

The `heuristic execution` receives all hyperparameters of the MIP-heuristic, solver settings, and instance references from the `experiment controller`. Based on the hyperparameters provided, it dynamically constructs the corresponding partial MIP-models P_i^λ by fixing or relaxing subsets of decision variables while reusing the same underlying Pyomo abstract model. Solver execution, result extraction and status handling are encapsulated within the heuristic implementation.

Result logging is handled directly within the class `HeuristicSolving` and includes both human-readable text files and structured CSV output capturing solver parameters, heuristic settings, solution quality $f_1^n(\lambda)$, computational cost $f_2^n(\lambda)$, and termination status. This ensures consistent reporting across monolithic and polylithic runs.

RESULTS AND CONCLUSIONS

In previous work, computational studies with the hoist scheduling use case were performed to compare the BO-based algorithm with random search. The MIP-HOP was solved with two different single objective functions: The makespan (objective of the large-scale MIP model M) with many optimal solutions (see Figure 3); and the computational cost (measured by the CPU-time) with few optimal solutions (see Figure 4).

The results showed that the BO-based algorithm is slower than the random search but finds better (or at least equally good) solutions than the random search [10].

This paper introduced a generic formulation of the MIP-heuristic optimization problem (MIP-HOP) and a corresponding modular software architecture. By abstracting from problem-specific characteristics, the formulation unifies the description of monolithic and polylithic MIP models, heuristic decomposition and composition, and hyperparameter optimization of MIP-heuristics within a single architecture. The proposed software reflects this formulation and supports reproducible experimentation by separating experiment setup, experiment control, heuristic execution, and model definition. The integration of established tools for mathematical programming and Bayesian optimization, together with the public release of the code, facilitates reuse and transfer of the approach to other application domains.

Future work will build on the presented formulation and software to conduct an extensive empirical

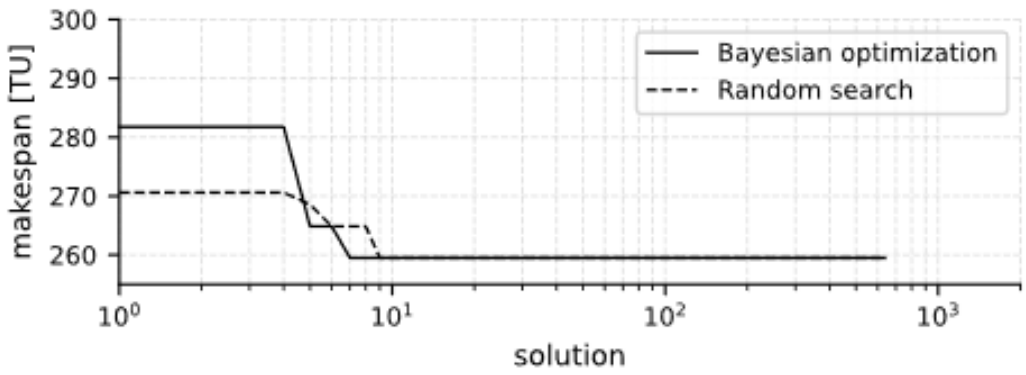

Figure 3. Convergence of the makespan (in time units) over the number of solutions λ in the hyperparameter space

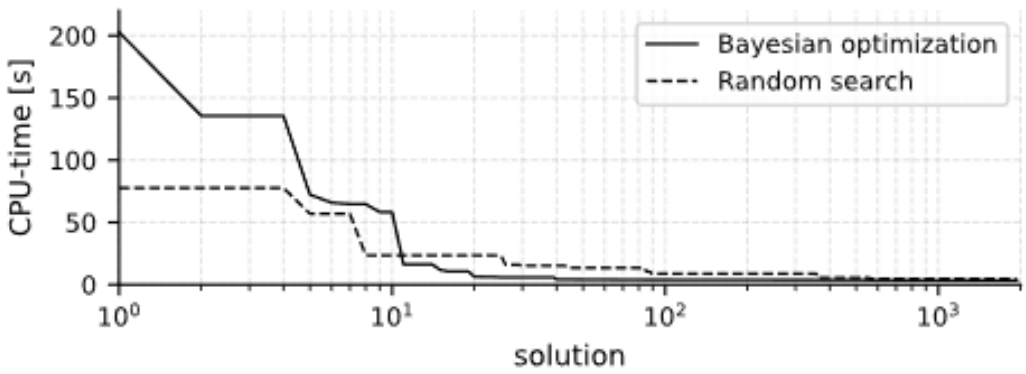

Figure 4. Convergence of the CPU-time (in seconds) over the number of solutions λ in the hyperparameter space

evaluation. The behavior of the two-phase rolling horizon heuristic will be analyzed on small- and large-scale hoist scheduling instances within two- and five-dimensional hyperparameter spaces. This evaluation aims to quantify how MIP-heuristic hyperparameters affect solution quality and computational cost, depending on instance characteristics and dimensionality of the hyperparameter space. Different hyperparameter tuning strategies, including grid search, random search, and Bayesian optimization, will be compared to derive recommendations on the choice of the tuning strategy with respect to the competing objectives solution quality and computational cost.

DIGITAL SUPPLEMENTARY MATERIAL

The digital supplementary material accompanying this paper is publicly available at:

https://github.com/SoHildebrandt/ES-CAPE36_MIP-HOP_GenericFormulationAndCode

The repository contains the complete implementation described in the Software Architecture section, including:

- the experiment setup and control logic (main.py, experiment_controller.py),
- the mixed-integer programming model of the hoist scheduling problem (`HSP_MIP_model.py`),
- the two-phase rolling horizon MIP-heuristic (HSP_MIP_heuristic.py), and
- a set of small-scale and large-scale hoist scheduling instances provided in JSON format.

The provided code realizes the modular software architecture described in this paper, including the interaction between experiment setup, experiment control, heuristic execution, and mathematical model definition. The included input instances serve as illustrative examples and enable readers to inspect, execute, and extend the implementation.

AUTHOR IDENTIFIERS

Author ORCIDs:
Hildebrandt S: 0000-0001-9606-3695
Franke M: 0000-0002-0382-0480
Zondervan E: 0000-0003-2659-1622
Sand G: 0009-0005-3162-4726

REFERENCES

1. A llemobayo J, Durodola O, Alade O, J Awotunde O, T Olanrewaju A, Falana O, Ogungbire A, Osinuga A, Ogunbiyi D, Ifeanyi A, E Odezuligbo I, E Edu O. Hyperparameter tuning in machine learning: a comprehensive review. J. Eng. Res. Rep. 26:388-395 (2024). https://doi.org/10.9734/jerr/2024/v26i61188
2. Hutter F, Kotthoff L, Vanschoren J. Automated machine learning. Springer International Publishing (2019). https://doi.org/10.1007/978-3-030-05318-5
3. Bischl B, Binder M, Lang M, Pielok T, Richter J, Coors S, Thomas J, Ullmann T, Becker M, Boulesteix A, Deng D, Lindauer M. Hyperparameter optimization: foundations, algorithms, best practices, and open challenges. WIREs Data Min & Knowl 13: (2023). https://doi.org/10.1002/widm.1484
4. Basso S, Ceselli A, Tettamanzi A. Random sampling and machine learning to understand good decompositions. Ann Oper Res 284:501-526 (2018). https://doi.org/10.1007/s10479-018-3067-9
5. Mitrai I, Daoutidis P. Taking the human out of decomposition-based optimization via artificial intelligence, part I: learning when to decompose. Computers & Chemical Engineering 186:108688 (2024). https://doi.org/10.1016/j.compchemeng.2024.108688
6. Mitrai I, Daoutidis P. Taking the human out of decomposition-based optimization via artificial intelligence, part II: learning to initialize. Computers & Chemical Engineering 186:108686 (2024). https://doi.org/10.1016/j.compchemeng.2024.108686
7. S. Hildebrandt, G. Sand. Hyperparameter optimization of matheuristics for hoist scheduling. ESCAPE34/PSE24. Ed: Manenti F, Reklaitis GV. Elsevier (2024) https://doi.org/10.1016/B978-0-443-28824-1.50488-9
8. Sand G, Hildebrandt S, Nunes S, Yip CO, Franke M. Tune decomposition schemes for large-scale mixed-integer programs by bayesian optimization. Systems and Control Transactions 4:1468-1473 (2025). https://doi.org/10.69997/sct.120655
9. Hildebrandt S, Nunes S, Franke M, Sand G. Studies on a bayesian optimization based approach to tune hyperparameters of matheuristics. Lecture Notes in Computer Science :203-217 (2026). https://doi.org/10.1007/978-3-032-09192-5_14
10. S. Hildebrandt et al.. Bayesian Optimization of Matheuristics for Large-Scale Mixed-Integer Programs. OR spectrum (work in progress)
11. M. Lindauer et al.. SMAC3: A versatile Bayesian optimization package for hyperparameter optimization. Journal of Machine Learning Research 23:1–9 (2022)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Large Language Model Enhanced Fault Diagnosis Framework for Chemical Processes

Jingkang Liang[a], Gürkan Sin[a]*

[a] Process and Systems Engineering Center (PROSYS), Department of Chemical and Biochemical Engineering, Technical University of Denmark (DTU), 2800 Kgs. Lyngby, Denmark
* Corresponding Author: gsi@kt.dtu.dk

ABSTRACT

Fault diagnosis is essential for ensuring safety and efficiency in chemical process industries. Conventional diagnostic systems often generate raw numerical outputs that require extensive human interpretation, increasing the operator's workload and slowing decision-making during abnormal events. To overcome these limitations, this work introduces a model context protocol (MCP)-integrated fault diagnosis framework, where a Large Language Model (LLM) functions as the MCP client, coordinating multiple diagnostic tools through a unified protocol. Within the proposed framework, the LLM interacts with specialized diagnostic tools, including a convolutional neural network-based fault diagnosis model and an ensemble-based variant for uncertainty-aware analysis. The LLM synthesizes the outputs of these tools and generates operator-oriented natural-language reports that summarize diagnostic results and explicitly communicate uncertainty, thereby supporting more transparent and efficient decision-making.A benchmarking study based on the Tennessee Eastman (TE) process is conducted to evaluate the framework using multiple LLMs under identical settings. The evaluation assesses diagnostic accuracy, tool adherence, and operator report quality using an LLM-as-a-judge methodology. Experimental results show that LLM performance varies significantly across models, and that reliable tool calling and uncertainty-aware reasoning are more critical than model size for effective fault diagnosis assistance.Overall, the proposed framework demonstrates strong potential for enhancing fault diagnosis workflows in chemical processes by improving usability, interpretability, and decision-making efficiency without modifying existing diagnostic algorithms.

Keywords: Artificial Intelligence, Fault Detection, Large Language Model

1. INTRODUCTION

Fault diagnosis plays an important role in chemical process industries, where undetected or misdiagnosed faults can lead to severe safety incidents, environmental damage, and significant economic losses [1]. To address these challenges, machine learning and deep learning techniques have been widely applied for fault detection and diagnosis [2], enabling earlier identification of abnormal operating conditions and improving overall process safety and reliability [3, 4]. Despite their performance, most existing fault diagnosis systems remain algorithm-centric, providing outputs in the form of numerical scores or fault probabilities [5]. Such results often lack intuitive, human-friendly explanations, increasing the cognitive workload of operators and potentially slowing down decision-making during abnormal events.

In parallel, recent advances in natural language processing have led to the rapid development of large language models (LLMs) based on transformer architectures [6]. These models demonstrate strong capabilities in reasoning, contextual understanding, and natural-language generation [7]. Beyond standalone text generation, LLMs have evolved into agentic systems, where an LLM can autonomously select and invoke external tools, interact with real-world applications, and maintain contextual memory to support more adaptive and personalized interactions [8]. In chemical engineering, LLMs have begun to emerge as promising tools for different tasks. Caccavale et al. [9] introduced ChatGMP, an LLM-based chatbot to support student learning in chemical engineering education. Liang et al. [10] leveraged LLM agent in

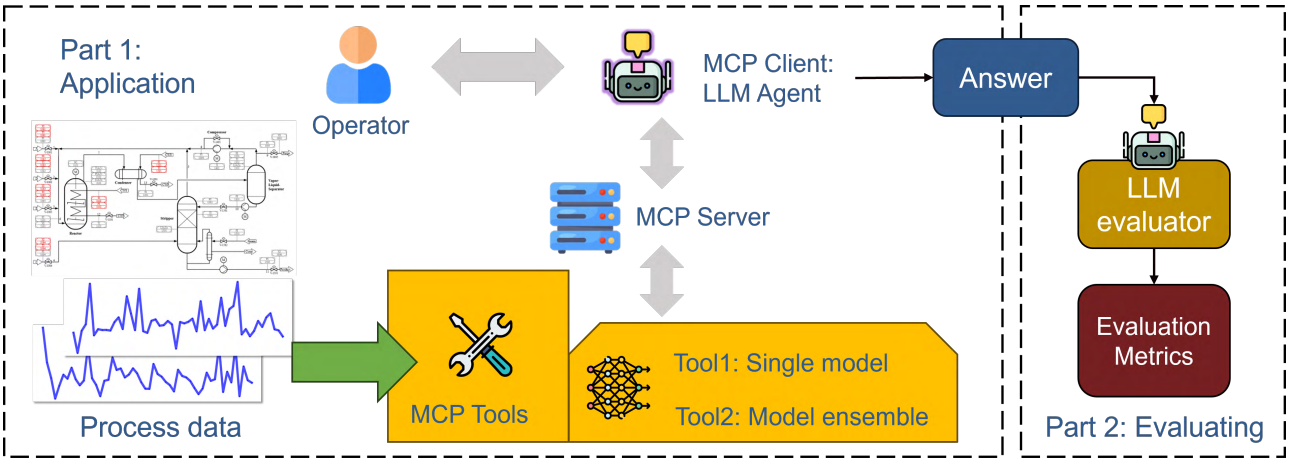

Figure 1. The proposed LLM-enhanced fault diagnosis framework with Model Context Protocol (MCP)

chemical simulations for simulation flowsheet analysis and synthesis. Zeng et al. [11] demonstrate an LLM-guided multi-agent optimization system that outperforms traditional methods on hydrodealkylation optimization . In fault diagnosis, FaultExplainer [12] demonstrates the potential of large language models to enhance fault diagnosis systems by generating interpretable, human-readable explanations alongside data-driven diagnostic results.

Motivated by the reasoning and tool-interaction capabilities of LLM agents, this work proposes a framework that integrates LLMs into chemical process fault diagnosis to enhance usability and decision support. Operators can query the system using natural language, while the LLM agent autonomously selects and invokes appropriate diagnostic tools using model context protocol (MCP), evaluates confidence in the results, and synthesizes the outputs into interpretable, operator-oriented explanations. By combining the reliability of established fault diagnosis algorithms with the flexibility and expressiveness of LLM agents, the proposed framework aims to reduce operator workload, improve situational awareness, and support faster and more informed decision-making during abnormal operating conditions.

In addition to proposing the framework, this work conducts a systematic benchmarking study using a dedicated evaluation dataset to compare the performance of different LLMs in terms of diagnostic accuracy, tool adherence and report quality. This study provides practical insights into LLM selection for fault diagnosis tasks that require both tool calling and operator-oriented report generation.

The main contributions of this paper are summarized as follows:

1) An MCP-based fault diagnosis framework is proposed, in which an LLM agent selects diagnostic tools based on prediction uncertainty and generate user-friendly report.

2) A structured evaluation experiment for LLM agents in fault diagnosis is carried out, enabling a comparative study of different LLMs under fair settings.

The remainder of this paper is organized as follows. Section 2 introduces the methodologies used in the proposed framework, including the CNN-based fault diagnosis algorithm, the MCP tools, and the LLM-as-a-judge evaluation strategy. Section 3 presents the experimental study of the proposed framework under different language models, using the Tennessee Eastman Process as a benchmarking dataset. Section 4 concludes the paper and outlines directions for future work.

2. METHODOLOGIES

2.1 Proposed workflow

As shown in Figure 1, the proposed framework defines an interaction between the operator, the LLM agent, the MCP server, diagnostic tools, and the chemical process. The operator interacts with the system by submitting natural-language queries to the LLM agent, which acts as an MCP client. The chemical process generates process data that are accessed by fault diagnosis algorithms exposed as MCP tools. These tools are managed by an MCP server, which provides standardized interfaces for tool discovery and invocation to the LLM. Upon receiving a query, the LLM agent interprets the user intent, selects appropriate diagnostic tool(s) via the MCP server, and invokes them using the relevant process data. The MCP tools perform fault diagnosis and return structured results to the LLM agent. The LLM agent aggregates and interprets the tool outputs and generates a user-facing fault diagnosis response for the operator.

In addition, an evaluation module is incorporated, where the generated responses are assessed using an LLM-as-a-judge [13] approach to compute multiple evaluation metrics. This evaluation provides insights into the usefulness, reliability, and decision-support quality of the LLM agent.

The following sections describe the diagnostic tools, the LLM tool-calling strategy, and the evaluation methodology, which together form the key components of this work.

2.2 Fault diagnosis algorithms

In this work, a multi-scale convolutional neural network (MSCNN) [14] is adopted as the core fault diagnosis algorithm due to its demonstrated effectiveness in previous fault diagnosis studies. Two variants are employed: a single MSCNN model and an ensemble-based MSCNN approach designed to provide more robust predictions.

The single MSCNN model is trained using the training dataset of the Tennessee Eastman Process. For the ensemble approach, multiple MSCNN models with identical network architectures and hyperparameters are trained independently. Model diversity is introduced through random initialization of model parameters and by training each model on different subsets of the training data. The final ensemble prediction is obtained by aggregating the outputs of all individual models, resulting in improved robustness against uncertainty in individual model predictions.

Based on prior experimental results, the ensemble-based MSCNN achieves an overall diagnostic accuracy of 73.67% in difficult faults, outperforming the single-model MSCNN, which attains an accuracy of 68.04%. These results demonstrate the advantage of the ensemble approach in providing more reliable fault diagnosis outcomes for subsequent interpretation by the LLM agent.

2.3 MCP tools

By using the Model Context Protocol (MCP) [15], a standardized protocol designed to expose external tools and their interfaces to large language models, the LLM is provided with structured contextual information about the available tools, including their functionality, required input parameters, and expected outputs. The tools are defined in a MCP server, the LLM agent, is a MCP client that can connect to the MCP server and get the available tools. MCP enables LLMs to interact with domain-specific software in a controlled and interpretable manner without requiring direct access to the underlying code or execution environment.

With this information, the LLM can reason about when and how to invoke a tool to solve tasks that require interaction with external software or applications. In this way, MCP serves as a standardized interface between the LLM and domain-specific tools, significantly simplifying tool invocation and integration. In this work, two diagnostic tools are exposed to the LLM through MCP, along with their corresponding descriptions:

Predict_sample_single_model(): Predicts the fault type and associated probabilities for a given sample index using a single diagnostic model.

Predict_sample_ensemble_with_plot(): Predicts fault types and probabilities using a model ensemble and generates a probability distribution plot across fault classes.

By invoking these tools, the LLM gains access to the fault diagnosis results produced by the underlying algorithms. The LLM then interprets, summarizes, and presents these results in a human-readable form, enabling operators to understand both the predicted fault and the associated uncertainty.

2.4 Agent Reasoning and Response Strategy

A rule-based tool selection logic is defined to guide the reasoning behavior of the LLM agent. Let $p_{\max}$ denote the maximum fault probability predicted by the single-model fault diagnosis algorithm. The tool selection strategy is formulated as follows:

- If $p_{\max} \geq 0.95$, the LLM agent directly uses the output of the single-model diagnosis.
- Otherwise, the LLM agent invokes the ensemble-based diagnosis tool and analyzes alternative faults, as well as plotting the result to the operator.

The rationale behind this strategy is to balance efficiency and robustness in operator-facing fault diagnosis. When the prediction confidence is high, the LLM agent is expected to provide a concise and confident report that focuses on the primary fault, thereby minimizing unnecessary computation and cognitive load. In contrast, when the prediction confidence is low, the LLM agent prioritizes robustness and transparency by incorporating ensemble results into the report. In such cases, the agent presents the most likely fault together with plausible alternative faults and their corresponding vote counts, explicitly communicating diagnostic uncertainty to the operator. This design supports efficient yet transparent decision-making in uncertain fault conditions.

2.5 LLM-as-a-judge

To evaluate the quality of the responses generated by the LLM agent, human feedback is generally regarded as the most reliable approach, since the goal of the system is to support human operators. However, for large-scale and repeated evaluations, relying on human assessment can be time-consuming, costly, and inconsistent. As an alternative, this work adopts an automated evaluation strategy known as LLM-as-a-judge [13], where a well-defined LLM is used to assess the outputs of another LLM.

The evaluation workflow using LLM-as-a-judge is illustrated as follows. First, a user submits a query to the LLM agent, which then performs reasoning and tool invocation to complete the fault diagnosis task and generate a final response. All intermediate information, including

the user query, tool calls, and the final answer, is logged and provided as input to the LLM evaluator. In addition, auxiliary information from the evaluation dataset—such as the true fault label and the predefined uncertainty level of the sample—is included. A system prompt is also supplied to guide the judge LLM on the evaluation criteria and scoring procedure.

Based on the provided information, the judge LLM performs its own reasoning and assigns scores according to the specified metrics. This process closely resembles human evaluation, while offering improved efficiency and consistency across repeated experiments. The specific evaluation metrics used in this study are detailed in the following section.

3. EXPERIMENTAL STUDIES

3.1 Tennessee Eastman Process Dataset

The Tennessee Eastman process is used as a case study for the proposed LLM-enhanced fault diagnosis task. It is used to simulate fault scenarios in a chemical process where an operator queries an LLM agent to analyze process data, invoke diagnostic algorithms, and generate fault diagnosis reports. The TE process is a widely used benchmark in the literature for fault detection and fault diagnosis in chemical processes. In this work, its usage is extended to evaluate how LLMs can enhance fault diagnosis workflows through tool calling and natural-language interaction.

The dataset consists of simulation running with 21 different fault types [16]. In total 52 process variables are input to the algorithm. The output of the algorithm is typically the fault number with the highest probability, which is usually regarded as the predicted fault. However, in the uncertainty aware method, we take the probability distribution of all faults, so that we can see how confident it is for the model to predict the primary fault. The training, validation and test sets are separated by the ratio of 6:2:2, and the test set will contribute to the LLM evaluation dataset in section 3.2.

3.2 LLM-fault diagnosis evaluation dataset

To evaluate the behavior of the LLM agent in a controlled and quantitative manner, an evaluation dataset with ground-truth labels is constructed. This dataset enables systematic assessment of the LLM's reasoning, tool usage, and response quality under different levels of diagnostic uncertainty. A total of 100 samples are selected, consisting of 50 low-uncertainty samples and 50 high-uncertainty samples. All samples are drawn from the test set of the Tennessee Eastman process dataset, ensuring that they are not seen by the diagnostic models during training.

Each data instance in the evaluation dataset is formulated as a natural-language query and is accompanied by additional labels used for evaluation. The structure of the dataset is illustrated in Table 1.

Table 1: Example of the LLM Evaluation Dataset

Question	True label	Uncertainty level
What is the fault type of sample 761?	13	Low
...	...	...
What is the fault type of sample 4201?	21	High

It is important to emphasize that this work focuses on evaluating the role of LLMs in enhancing fault diagnosis workflows, rather than improving the predictive accuracy of the underlying diagnostic algorithms. Accordingly, this study does not aim to develop or benchmark new fault diagnosis models. Instead, the emphasis is on exploring an LLM-centered framework that improves usability, transparency, and decision support in fault diagnosis by integrating existing diagnostic models with LLM agents.

3.3 Evaluation Metrics

To comprehensively evaluate the performance of the proposed LLM-enhanced fault diagnosis framework, three complementary metrics are employed. These metrics assess technical correctness, rule compliance, and operator-oriented response quality.

1) Diagnostic Accuracy

Diagnostic accuracy is evaluated using an LLM-as-a-judge approach with a tiered scoring scheme of $\{0, 0.5, 1\}$, based on the true fault label and the response generated by the LLM agent. A score of 1 is assigned if the primary fault reported by the agent matches the true label. If the true fault is mentioned as an alternative fault, a score of 0.5 is assigned. Otherwise, a score of 0 is given. This tiered scheme accounts for uncertainty-aware responses and avoids penalizing the agent for explicitly acknowledging diagnostic ambiguity.

2) Tool Adherence

Tool adherence is evaluated using a rule-based evaluator that checks whether the LLM agent follows the tool selection strategy defined in Section 2.4. The evaluation is performed based on the labeled uncertainty level of each sample and the corresponding expected tool invocation. An LLM agent is considered compliant if it consistently invokes the appropriate diagnostic tool (single model or ensemble) according to the predefined uncertainty threshold.

3) Operator Report Quality

Operator report quality is assessed using an LLM-

as-a-judge approach, with scores ranging from 1 to 5. A system prompt is provided to guide the judge LLM in scoring the responses based on multiple criteria, including alignment with the labeled uncertainty level, clarity and directness of the explanation, effectiveness in supporting operator decision-making, and the actionability of the generated report.

To help readers better understand the LLM's reasoning in the fault diagnosis task, Table 2 presents an example output generated by *gemini-2.5-flash* for sample 4201.

Table 2: Answer from gemini-2.5-flash of sample 4201

Question: What is the fault type of sample 4201?
Answer: For sample 4201, the primary fault type is 15. However, the confidence in this prediction is low (around 30%), indicating significant uncertainty. Alternative fault types to consider are: ▪ **Fault type 21** (with a mean probability of 29.17%) ▪ **Fault type 9** (with a mean probability of 11.84%) To distinguish between these, the operator should check for symptoms related to both fault type 15 and fault type 21, as they have very similar probabilities. Additionally, a review of indicators for fault type 9 would be prudent.

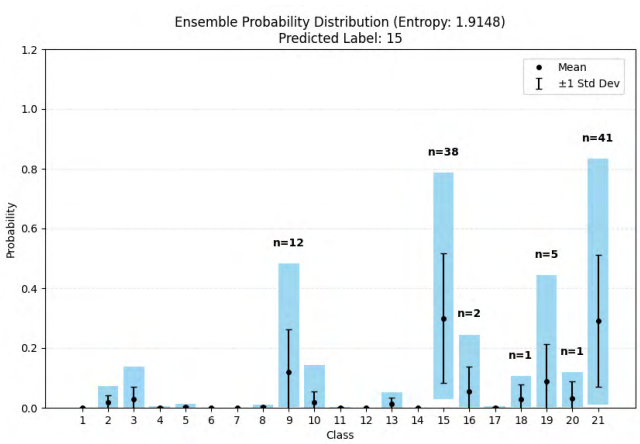

Figure 2. Diagnostic result visualization of sample 4201 by calling tool Predict_sample_ensemble_with_plot()

For this sample, the generated response achieves a fault accuracy score of 0.5, a report quality score of 5, and a tool adherence score of 1. These scores reflect that, although the prediction is uncertain, the LLM correctly follows the intended diagnostic strategy. Specifically, the LLM invokes the ensemble-based diagnostic tool, which produces a visualization of the fault probability distribution and exposes the uncertainty across fault classes. The LLM then incorporates this information into a structured and cautious explanation, discussing multiple plausible fault hypotheses instead of over-committing to a single class.

Overall, this example illustrates how the LLM leverages tool outputs and uncertainty information to generate an interpretable, decision support-oriented diagnosis, highlighting its potential role in enhancing fault diagnosis beyond raw classification results.

3.4 Results and Discussion

To evaluate the effectiveness of different LLMs within the proposed MCP-based fault diagnosis framework, experiments were conducted using multiple LLMs under identical settings. The evaluated models span multiple providers, architectures, and deployment types, including both open-source and closed-source models with varying parameter sizes. All models were evaluated under the same environment with identical prompts and tools, ensuring a fair comparison. The experimental results are summarized in Table 3.

Table 3: Experimental results from the LLM-enhance fault diagnosis report

Model name	Fault Accuracy	Tool Adherence	Report Quality
openai/gpt-4o-mini	0.83	1.00	5.00
x-ai/grok-4-fast	0.82	1.00	5.00
anthropic/claude-3.5-haiku	0.82	0.99	5.00
google/gemini-2.5-flash	0.82	1.00	4.94
openai/gpt-4o	0.82	1.00	5.00
deepseek/deepseek-chat-v3.1	0.79	0.97	4.85
meta-llama/llama-3.1-8b-instruct	0.75	0.85	4.60
google/gemini-2.5-flash-lite	0.70	0.69	4.89
qwen/qwen3-235b-a22b-2507	0.62	0.86	4.47
meta-llama/llama-3.3-70b-instruct	0.60	0.64	3.20

From these results, several conclusions can be drawn.

First, models from major commercial providers consistently demonstrate strong performance across all three metrics, with relatively small performance differences among them. This indicates that advanced LLMs are generally capable of strictly following the proposed tool selection strategy and generating high-quality, operator-oriented fault diagnosis reports. These results highlight the effectiveness of leveraging LLMs to enhance

fault diagnosis workflows in chemical processes.

Second, differences between lightweight and full-scale models are model-dependent rather than universal. For example, Gemini-2.5-Flash-Lite performs noticeably worse than Gemini-2.5-Flash, suggesting that model compression may negatively impact tool adherence and diagnostic reasoning. In contrast, gpt-4o-mini achieves performance comparable to its full-scale counterpart (gpt-4o), demonstrating that it is possible to deploy lightweight models while still achieving high diagnostic accuracy and report quality when tool-calling behavior is well aligned with the task requirements.

Third, open-source models are generally less competitive in this setting, exhibiting lower fault accuracy, weaker tool adherence, and reduced report quality. Among them, deepseek-chat-v3.1 performs the best, while llama-3.1-8b, qwen3-235b, and llama-3.3-70b show significantly poorer performance. Notably, the wide variation in parameter sizes among these models further indicates that larger model size does not necessarily lead to better performance in MCP-based fault diagnosis tasks.

Finally, it can be observed that the upper bound of fault diagnosis accuracy across all evaluated models is approximately 0.83, even when tool adherence reaches 100%. This suggests that the accuracy of the LLM agent is fundamentally constrained by the performance of the underlying diagnostic tools. This behavior is intuitive: by mitigating hallucinations through strict tool usage, the LLM's outputs are limited to the information provided by the invoked algorithms.

Overall, these results demonstrate that LLMs can effectively assist fault diagnosis in chemical processes by leveraging tool-calling capabilities and natural-language report generation. The findings further indicate that LLM performance varies across models, and that selecting models with strong tool-calling reliability and reasoning ability is more important than choosing models solely based on parameter size. Rather than establishing a permanent ranking of LLMs, these results highlight the importance of a systematic benchmarking process for evaluating LLM-based agents in fault diagnosis tasks.

4. CONCLUSION AND FUTURE WORK

This work proposes an LLM-enhanced fault diagnosis framework that integrates large language models with fault diagnosis algorithms through structured tool calling. By enabling the LLM to invoke diagnostic tools and generate operator-oriented reports, the proposed framework improves the usability and efficiency of fault diagnosis in chemical processes. Experimental results across multiple LLMs demonstrate that many models are capable of correctly following the proposed strategy, while the achievable diagnostic accuracy is fundamentally constrained by the performance of the underlying fault diagnosis algorithms. Advanced LLMs, however, are effective in interpreting diagnostic results and communicating them to operators in a clear and actionable manner.

Overall, the results highlight the potential of applying LLMs to enhance fault diagnosis workflows in chemical engineering, particularly by improving user experience, transparency, and decision-making efficiency rather than by replacing existing diagnostic models.

Future work will focus on extending the proposed framework toward more advanced multi-agent systems and scaling up the validation to a broader range of industrial process systems. In addition, the observed performance gap between open-source and commercial LLMs motivates further research into improved system design and agent strategies to narrow this gap. Advancing the capability of open-source LLM-based systems is particularly important for addressing data privacy and deployment constraints in industrial environments.

Specifically, many critical chemical facilities operate on strictly air-gapped control networks without internet access to mitigate cybersecurity risks and protect sensitive intellectual property. Cloud-based commercial LLMs cannot be deployed in these settings, making locally deployed open-source models a strict industrial necessity. Furthermore, because chemical process monitoring is a continuous task, relying on commercial APIs would incur significant ongoing token costs, whereas locally deployed open-source models incur zero variable query costs once established. Therefore, enhancing the performance of local, open-source LLMs remains a critical priority for practical industrial deployment.

ACKNOWLEDGEMENTS

This research was funded by the European Union Horizon Europe 2022 Research and Innovation Programme under the Marie Sklodowska-Curie Grant Agreement No. 101119358 (ProSafe).

AUTHOR IDENTIFIERS

Author ORCIDs:
Liang Jingkang: 0000-0003-0815-3135
Sin Gürkan: 0000-0003-0513-4502

REFERENCES

1. Bi X, Qin R, Wu D, Zheng S, Zhao J. One step forward for smart chemical process fault detection and diagnosis. Computers & Chemical Engineering 164:107884 (2022). https://doi.org/10.1016/j.compchemeng.2022.107884

2. Taqvi SAA, Zabiri H, Tufa LD, Uddin F, Fatima SA, Maulud AS. A review on data?driven learning approaches for fault detection and diagnosis in chemical processes. ChemBioEng Reviews 8:239-259 (2021). https://doi.org/10.1002/cben.202000027
3. Deng X, Zhang Z. Nonlinear chemical process fault diagnosis using ensemble deep support vector data description. Sensors 20:4599 (2020). https://doi.org/10.3390/s20164599
4. Lomov I, Lyubimov M, Makarov I, Zhukov LE. Fault detection in tennessee eastman process with temporal deep learning models. Journal of Industrial Information Integration 23:100216 (2021). https://doi.org/10.1016/j.jii.2021.100216
5. Shi T, Yang A, Jin Y, Ren J, Shen W, Dong L, Man Y. Artificial intelligence in process systems engineering. Applications of Artificial Intelligence in Process Systems Engineering :1-10 (2021). https://doi.org/10.1016/b978-0-12-821092-5.00010-3
6. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention Is All You Need 2023. https://doi.org/10.48550/arXiv.1706.03762
7. Cao Y, Griffiths T, Narasimhan K, Shafran I, Yao S, Yu D, Zhao J. Tree of thoughts: deliberate problem solving with large language models. Advances in Neural Information Processing Systems 36 :11809-11822 (2023). https://doi.org/10.52202/075280-0517
8. Ferrag MA, Tihanyi N, Debbah M. From LLM Reasoning to Autonomous AI Agents: A Comprehensive Review 2025. https://doi.org/10.48550/arXiv.2504.19678
9. Caccavale F, Gargalo CL, Kager J, Larsen S, Gernaey KV, Krühne U. Chatgmp: a case of AI chatbots in chemical engineering education towards the automation of repetitive tasks. Computers and Education: Artificial Intelligence 8:100354 (2025). https://doi.org/10.1016/j.caeai.2024.100354
10. Liang J, Groll N, Sin G. Large Language Model Agent for User-friendly Chemical Process Simulations 2026. https://doi.org/10.48550/arXiv.2601.11650
11. Zeng T, Badrinarayanan S, Ock J, Lai C-K, Farimani AB. LLM-guided Chemical Process Optimization with a Multi-Agent Approach 2025. https://doi.org/10.48550/ARXIV.2506.20921
12. Khan A, Nahar R, Chen H, Constante Flores GE, Li C. Faultexplainer: leveraging large language models for interpretable fault detection and diagnosis. Computers & Chemical Engineering 199:109152 (2025). https://doi.org/10.1016/j.compchemeng.2025.109152
13. Gu J, Jiang X, Shi Z, Tan H, Zhai X, Xu C, Li W, Shen Y, Ma S, Liu H, Wang S, Zhang K, Lin Z, Zhang B, Ni L, Gao W, Wang Y, Guo J. A survey on llm-as-a-judge. The Innovation :101253 (2026). https://doi.org/10.1016/j.xinn.2025.101253
14. Liang J, Sin G. Enhancing fault diagnosis for chemical processes via MSCNN with hyperparameters optimization. Systems and Control Transactions 4:1712-1717 (2025). https://doi.org/10.69997/sct.181373
15. What is the Model Context Protocol (MCP)? Model Context Protoc n.d. https://modelcontextprotocol.io/docs/getting-started/intro (accessed December 3, 2025).
16. Melo A, Câmara MM, Clavijo N, Pinto JC. Open benchmarks for assessment of process monitoring and fault diagnosis techniques: a review and critical analysis. Computers & Chemical Engineering 165:107964 (2022). https://doi.org/10.1016/j.compchemeng.2022.107964

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Unified Python/JAX Framework for Thermodynamic Modeling, Nonlinear Solvers, and DAE Solution of Hydrocarbon Systems

Carlos C. Sanz[a] and Galo Le Roux[a]*

[a] Universidade de São Paulo, Chemical Engineering Department, São Paulo, Brazil
* Corresponding Author: galoroux@usp.br

ABSTRACT

Dynamic simulation of distillation columns and chemical reactors remains essential for plant design, controllability analysis, and economic optimization. High-purity separations of close-boiling mixtures present significant computational challenges due to nonlinear thermodynamic behavior and stiff differential-algebraic equation (DAE) systems. This work presents a unified Python/JAX framework integrating four computational modules: (1) Peng-Robinson thermodynamics with complex-step differentiation, (2) nonlinear solvers (Newton, Broyden, Newton-Krylov) with automatic Curtis-Reid scaling, (3) DAE solver with Radau IIA collocation and intelligent auto-selection, and (4) constrained optimization using the Augmented Lagrangian Method with JAX automatic differentiation. The framework leverages JAX's just-in-time compilation (JIT), vectorization (vmap), and automatic differentiation (AD) to achieve near-compiled-language performance. Validation includes: nonlinear solver benchmarks with Newton-Krylov achieving 100% success across seven problems (n=2 to 5000), Williams-Otto reactor optimization with 0.06% deviation from published literature and 1325× real-time speedup, and a 180-stage propylene-propane splitter with 18× real-time performance under three concurrent disturbances. The framework shows that is possible to have an open-source alternative suitable for real-time optimization, operator training, and advanced process control applications.

Keywords: Python, JAX, Process Simulation, Nonlinear Solvers, DAE Systems, Optimization, Distillation

1. INTRODUCTION

Dynamic modeling of chemical processes supports applications from conceptual plant design to real-time optimization and operator training. However, several technical challenges complicate simulation of hydrocarbon separation processes. High-purity product specifications (>99%) demand rigorous thermodynamic models that accurately capture vapor-liquid equilibrium. The resulting mathematical models are differential-algebraic equations (DAE) systems in which algebraic constraints (phase equilibrium, summation) couple with differential mass balances. Large industrial distillation columns (50-200 stages) produce systems with thousands of coupled equations, necessitating efficient sparse linear algebra.

This work addresses these challenges using Python with JAX [1, 2], providing: (1) automatic differentiation for exact Jacobians and gradients, (2) just-in-time compilation via XLA [2], achieving near-C++ performance, (3) vectorization through vmap for batch processing, and (4) functional purity ensuring reliable transformations.

The framework comprises four modules: **thermodynamics** (Peng-Robinson with complex-step differentiation), **nonlinear solvers** (Newton/Broyden/Newton-Krylov with Curtis-Reid scaling), **DAE solver** (Radau IIA with auto-selection), and **optimization** (Augmented Lagrangian with JAX AD).

Sections 2–5 present methods, Section 6 validates via benchmarks and applications, and Section 7 concludes.

Table 1: JAX Commands Reference.

Command	Description	Application	Example
@jit	**JIT compilation to machine code via XLA**	**All computational kernels**	**@jit def f(x): return x**2**
vmap(f, in_axes)	Vectorize function over batch dimension	Multi-stage VLE, collocation	vmap(stage_func)(stages)
jacfwd(f)(x)	Forward-mode AD: Jacobian via n forward passes	Newton solver, constraints (n_out ≥ n_in)	J = jax.jacfwd(F)(x)
grad(f)(x)	Reverse-mode AD: gradient via 1 forward + 1 backward	Scalar optimization objectives	g = jax.grad(f)(x)
jvp(f, (x,), (v,))	Jacobian-vector product J·v	Newton-Krylov matrix-free solver	_, Jv = jax.jvp(F, (x,), (v,))
jnp.*	JAX NumPy functions (GPU/TPU compatible)	All array operations	jnp.sum(x), jnp.linalg.solve(A, b)

2. THERMODYNAMIC MODEL

2.1 Peng-Robinson Equation of State

The Peng-Robinson equation of state [3] relates pressure P, temperature T, and molar volume V. Parameters a(T) and b are computed from critical properties (Tc, Pc) and acentric factor ω [4]. For mixtures, van der Waals mixing rules employ binary interaction parameters kij fitted to experimental VLE data. Fugacity coefficients φ_i enable K-value calculations: $K_i = \varphi_i^L / \varphi_i^V$. Complete derivation in [3]; property calculations following [4].

2.2 Complex-Step Differentiation

Computing derivatives for Newton-Raphson VLE solvers requires careful consideration of JAX transformation composition. Using JAX automatic differentiation inside PREOS functions would create nested JIT/vmap contexts when called from DAE solvers, causing exponential memory growth and compilation overhead.

Complex-step differentiation [12] computes derivatives via $f'(x) = \mathrm{Im}[f(x + ih)]/h$ with $h \approx 10^{-20}$, achieving machine precision (10^{-16} accuracy). This approach operates as pure arithmetic without additional JAX transformations, avoiding nested issues while maintaining exact derivatives. Care was taken to contain complex dtype within thermodynamic kernels through dtype-agnostic implementations. Complete mathematical treatment in [12].

2.3 JAX Implementation

All thermodynamic kernels use @jit for compilation (10–50× speedup). Vectorization via vmap enables batch VLE calculations across stages (8–15× speedup). Precision enforced at 64-bit: jax.config.update("jax_enable_x64", True).

3. NONLINEAR EQUATION SOLVERS

3.1 Problem Formulation

General nonlinear system: $F(x) = 0$ where $F: \mathbb{R}^n \to \mathbb{R}^n$. Jacobian computed via forward-mode AD: J = jax.jacfwd(F)(x) evaluates n forward passes for exact derivatives.

3.2 Newton Method

Newton iteration: $J(x_k)\Delta x_k = -F(x_k)$, $x_{k+1} = x_k + \alpha_k \Delta x_k$. Three variants implemented: pure (α=1), line search (backtracking Armijo condition), dogleg (trust-region [10]). Cost: $O(n^3)$ due to LU decomposition.

3.3 Broyden Quasi-Newton

Broyden's method [10] uses rank-1 Jacobian approximation where $s_k = x_{k+1} - x_k$, $y_k = F(x_{k+1}) - F(x_k)$. Periodic Jacobian refresh (every 10–20 iterations) improves convergence. Cost: $O(n^2)$ per iteration.

3.4 Newton-Krylov Method

For large sparse systems, Newton-Krylov [6] solves the Newton equation using GMRES (Generalized Minimal Residual method for linear systems [6]) without forming the explicit Jacobian. Matrix-vector products computed efficiently via JVP. For n=5000, jax.jvp evaluates J·v in O(n) time with single forward pass, versus $O(n^2)$ to form explicit Jacobian. Eisenstat-Walker forcing terms [7] adapt GMRES tolerance dynamically.

Parameter sensitivity: Requires tuning GMRES restart (10–40), linear solver tolerance (1e-4 to 1e-6), and preconditioner choice. Framework uses conservative defaults (restart=40, tolerance=1e-4, no preconditioner) ensuring reliability without manual tuning. Cost: O(n) per iteration for sparse systems.

3.5 Curtis-Reid Scaling

Ill-conditioned systems with variables spanning multiple orders of magnitude cause convergence failure. Curtis-Reid scaling [5] addresses this through diagonal scaling. Activated when conditioning ratio $\kappa = \max(Dx)/\min(Dx) > 1000$. Scaled system: $\tilde{F}(\tilde{x}) = F(Dx \cdot \tilde{x})/Df = 0$ is solved iteratively.

4. DAE SOLVER

4.1 DAE Formulation

We solve an index-1 DAE in residual form

$$\mathbf{F}(t,\ \mathbf{y},\ \dot{\mathbf{y}},\ \mathbf{z}) = \mathbf{0},$$

with differential states $\mathbf{y} \in \mathbb{R}^{n_y}$ and algebraic variables $\mathbf{z} \in \mathbb{R}^{n_z}$. For example, for a distillation column, $\mathbf{y}(t)$ stacks stage component holdups (and controller integrators, if any), while $\mathbf{z}(t)$ stacks stage temperatures and algebraic quantities needed to enforce VLE + summation constraints. The model is semi-rigorous: it advances component material balances dynamically and enforces phase equilibrium algebraically, without explicit tray-by-tray energy balances. Temperatures are therefore algebraic unknowns solved from equilibrium closure.

4.2 Radau IIA Collocation

Time integration uses a stiffly accurate Radau IIA implicit collocation method with $K = 3$ nodes

$$\tau_1 = \frac{4-\sqrt{6}}{10}, \qquad \tau_2 = \frac{4+\sqrt{6}}{10}, \qquad \tau_3 = 1,$$

chosen for L-stability (critical for stiff DAEs).

For each time step $[t_n, t_{n+1}]$ with $h = t_{n+1} - t_n$, introduce collocation unknowns

$$\mathbf{Y}_i \approx \mathbf{y}(t_n + h\tau_i), \qquad \mathbf{Z}_i \approx \mathbf{z}(t_n + h\tau_i), \qquad i = 1, \dots, K,$$

with $\mathbf{Y}_0 = \mathbf{y}(t_n)$ known.

Differentiation matrix.

Collocation reconstructs $\dot{\mathbf{y}}_i$ from nodal values using a differentiation matrix $\mathbf{A}$:

$$\dot{\mathbf{Y}}_i \ = \ \sum_{j=0}^{K} A_{ij}\, \mathbf{Y}_j, \qquad i = 1, \dots, K.$$

In the implementation, $\mathbf{A}$ is computed once from the Lagrange basis polynomials on $\{0, \tau_1, \tau_2, \tau_3\}$, and its derivatives are obtained via JAX AD (no numerical differencing for $\mathbf{A}$).

Collocation residuals.

At each collocation node, enforce the DAE residual:

$$\mathbf{F}\left(t_n + h\tau_i,\ \mathbf{Y}_i,\ \sum_{j=0}^{K} A_{ij}\, \mathbf{Y}_j,\ \mathbf{Z}_i\right) = \mathbf{0}, \qquad i = 1, \dots, K.$$

This produces a single nonlinear system per step with

$$K(n_y + n_z) \text{ unknowns} \quad \leftrightarrow \quad K(n_y + n_z) \text{ equations},$$

over the stacked vector

$$\mathbf{w} = \begin{pmatrix} \mathbf{Y}_1 \\ \vdots \\ \mathbf{Y}_K \\ \mathbf{Z}_1 \\ \vdots \\ \mathbf{Z}_K \end{pmatrix}.$$

So, "time integration" is dominated by repeatedly solving $\mathbf{R}(\mathbf{w}) = \mathbf{0}$ for each step.

4.3 Residual Structure and Vectorization

The DAE residual $\mathbf{R}(\mathbf{w})$ decomposes naturally into:

Differential residuals: stage component material balance mismatches, using $\dot{\mathbf{Y}}_i$ from the collocation reconstruction.

Algebraic residuals: VLE relations, e.g.,

$$y_j = K_j(T, P, \mathbf{x})\, x_j,$$

and summation constraints ($\sum_j x_j = 1,\ \sum_j y_j = 1$, etc.).

Residual evaluation is organized stage-wise and then vectorized across stages via a single vmap call. This avoids Python loops over trays and exposes SIMD parallelism to XLA. Practically, one batched evaluation replaces $\mathcal{O}(10^2)$ sequential stage evaluations, which is a major speed driver in stiff dynamic runs.

4.4 Dense Newton vs. Matrix-Free Newton–Krylov

Each step solves $\mathbf{R}(\mathbf{w}) = \mathbf{0}$. Two solvers are supported:

Dense Newton (explicit Jacobian).

Compute the Jacobian $\mathbf{J} = \partial\mathbf{R}/\,\partial\mathbf{w}$ with JAX forward-mode AD (jacfwd) and solve the Newton system

$$\mathbf{J}\Delta\mathbf{w} = -\mathbf{R}(\mathbf{w})$$

via dense linear algebra (LU). This is robust and low-tuning, but the linear solve cost grows rapidly with n.

Newton–Krylov (matrix-free).

Avoid forming $\mathbf{J}$ explicitly. Use JVP to evaluate Jacobian–vector products $\mathbf{Jv}$ inside GMRES, i.e.,

$$\mathbf{Jv} = \mathrm{jvp}(\mathbf{R},\ \mathbf{w};\ \mathbf{v}).$$

This reduces memory pressure and can be faster for large n, especially when the effective Jacobian is sparse/structured.

4.5 Auto-Selection: Size/Sparsity-Aware Method Choice

Because the "best" method depends on system size and effective sparsity, the solver estimates Jacobian density ρ by sampling a small number of columns using JVP and computing

$$\rho = \frac{\mathrm{nnz}}{n^2}.$$

A simple cost model compares:

Dense Newton: Jacobian build $\mathcal{O}(n^2)$ + factorization $\mathcal{O}(n^3)$

Newton–Krylov: residual $\mathcal{O}(n)$ + GMRES $\mathcal{O}(k \cdot n)$ per Newton step

Decision rules:

- *Very large systems:* force matrix-free (memory-bound)
- *Small systems:* prefer dense (robust, minimal tuning)
- *Intermediate:* choose lower predicted cost based on ρ and n

When Newton–Krylov is chosen, conservative defaults (tolerances/restarts) are used to preserve convergence without case-specific parameter tuning. However, for ill-conditioned or highly nonlinear systems, some tuning may still be necessary.

5. CONSTRAINED OPTIMIZATION

5.1 Augmented Lagrangian Method

General constrained problem: minimize f(x) subject to equality constraints h(x) = 0 and inequality constraints g(x) ≤ 0. Augmented Lagrangian formulation [11] uses penalty parameter μ > 0 and Lagrange multiplier estimates λh, λg. Algorithm alternates between solving unconstrained subproblem for x and updating multipliers. Penalty parameter μ increases if constraint violations persist, gradually enforcing feasibility.

5.2 Gradient Computation via JAX AD

JAX automatic differentiation provides exact derivatives with minimal computational overhead. For Williams-Otto reactor (n=2 decision variables), jax.grad returns exact 2D gradient in one forward plus one backward pass. Forward-mode AD is optimal when n_outputs > n_inputs [10]. The framework includes complex-step and finite-difference options for comparison.

BFGS quasi-Newton [10] solves unconstrained subproblems using rank-2 Hessian approximation updates with line search satisfying Wolfe conditions.

6. RESULTS AND VALIDATION

6.1 Nonlinear Solver Benchmarks

Seven Moré-Garbow-Hillstrom benchmark problems [13] spanning dimensions n=2 (Rosenbrock) to n=5000 (Brusselator 2D). Convergence tolerance 10^{-6}, maximum 100 iterations, default parameters.

Table 2: Success Rates on Standard Benchmarks.

Solver	Successful	Failed	Success Rate
Newton-Pure	6/7	1/7	85.7%
Newton-LineSearch	5/7	2/7	71.4%
Broyden-Pure	6/7	1/7	85.7%
Newton-Krylov-Pure	7/7	0/7	100.0%
Newton-Krylov-LineSearch	6/7	1/7	85.7%

Newton-Pure and Broyden-Pure achieved 85.7% success using fixed default parameters. Newton-Krylov-Pure achieved 100% success across all seven problems, notably solving ill-conditioned Brown Almost Linear (n=100) where all other methods failed due to near-singular Jacobians. This demonstrates the fundamental advantage of iterative linear solvers for poorly conditioned systems [6].

For Brusselator 2D (n=5000): Newton-Pure 15 iterations/19.5s versus Newton-Krylov-Pure 16 iterations/18.6s, validating theoretical O(n) scaling advantage of matrix-free methods.

Parameter sensitivity analysis: Newton-Krylov required problem-specific tuning of GMRES restart (10–40), linear tolerance (1e-4 to 1e-6), and preconditioner to achieve optimal performance. Newton-Pure converged reliably with fixed defaults, representing the fundamental trade-off between asymptotic complexity and configuration effort [6, 7].

6.2 Williams-Otto Reactor

Reaction System

The Williams-Otto reactor [14] consists of three series-parallel reactions occurring in a continuous stirred-tank reactor (CSTR) with constant holdup W = 2105 kg. Six components participate: reactants A and B, intermediate C, desired products P and E, and by-product G.

Reaction Scheme

$$A + B \rightarrow C, r_1 = k_1 C_A C_B$$

$$B + C \rightarrow P + E, r_2 = k_2 C_B C_C$$

$$C + P \rightarrow G, r_3 = k_3 C_C C_P$$

where Ci denotes the mass fraction of component i in the reactor (dimensionless), and ri is the reaction rate (s^{-1}).

Arrhenius Kinetics

The temperature-dependent rate constants follow the Arrhenius law:

$$k_i = k_i^0 \, exp\left(-\frac{E_{a,i}}{RT_R}\right), i = 1,2,3$$

where k_i^0 is the pre-exponential factor, $E_{a,i}$ is the activation energy, R = 8.314 J/(mol·K) is the universal gas constant, and T_R is the reactor temperature (K).

Table 3. Arrhenius kinetic parameters for the Williams-Otto reactor [14].

Reaction	Stoichiometry	k_0 (s^{-1})	*Ea*/R (K)	*Ea* (kJ/mol)
1	A + B → C	5.9755×10^9	12000	99.77
2	B + C → P + E	2.5962×10^{12}	15000	124.71
3	C + P → G	9.6283×10^{15}	20000	166.28

Component Mass Balances

For a CSTR with constant holdup W (kg), the dynamic mass balances for the six components are:

$$\frac{dC_A}{dt} = \frac{F_A C_{A,in} - FC_A}{W} - r_1$$

$$\frac{dC_B}{dt} = \frac{F_B C_{B,in} - FC_B}{W} - r_1 - r_2$$

$$\frac{dC_C}{dt} = -\frac{FC_C}{W} + r_1 - r_2 - r_3$$

$$\frac{dC_P}{dt} = -\frac{FC_P}{W} + r_2 - r_3$$

$$\frac{dC_E}{dt} = -\frac{FC_E}{W} + r_2$$

$$\frac{dC_G}{dt} = -\frac{FC_G}{W} + r_3$$

where F = $F_A + F_B$ is the total inlet mass flow rate (kg/s), F_A = 1.8275 kg/s is the fixed feed of pure A ($C_{A,in}$ = 1), and F_B is the adjustable feed of pure B ($C_{B,in}$ = 1). No C, P, E, or G enter with the feed.

Steady-State Condition

At steady state, all time derivatives vanish. The system reduces to six nonlinear algebraic equations:

$$F(C, T_R, F_B) = 0, where\ C = (C_A, C_B, C_C, C_P, C_E, C_G)$$

Given the reactor temperature TR and feed flow FB, the steady-state concentrations C are implicitly determined by solving the nonlinear system (11) via Newton's method.

Operating Conditions and Economic Parameters

Table 4. Reactor operating parameters and economic data.

Parameter	Symbol	Value	Units
Reactor holdup	W	2105	Kg
Feed flow of A (fixed)	F_a	1.8275	kg/s
Feed flow of B (adjustable)	F_E	4.0–6.0	kg/s
Reactor temperature range	T_r	70–100	°C
Parameter	**Symbol**	**Value**	**Units**
Base case temperature	T_r^0	85.0	°C
Base case feed flow B	F_E^0	4.50	kg/s
Price of product P	p_P	80	$/kg
Price of product E	p_E	20	$/kg
Cost of reactant A	c_A	40	$/kg
Cost of reactant B	$c_Æ$	25	$/kg

Constrained Optimization Problem

The economic optimization seeks to maximize the instantaneous profit by selecting the reactor temperature and feed flow of B. The problem is formulated as a constrained Nonlinear Programming (NLP) problem, where the decision variables are $x = (T_R, F_B)^T \in \mathbb{R}^2$

Objective Function

The profit function Π ($/h) accounts for product revenues minus raw material costs: The complete optimization problem is defined as:

$$max_{T_R,F_B}\, \Pi = (p_P X_P + p_E X_E)\, F - (c_A F_A + c_B F_B)$$

subject to: $F(C, T_R, F_B) = 0$

$$T_R \le T_R^{max} = 100\,°C$$

$$F_B \le F_B^{max} = 6\,kg/s$$

$$T_R^{min} \le T_R \le T_R^{max}, F_B^{min} \le F_B \le F_B^{max}$$

Inequality constraints (operational limits):

$$g_1(x): T_R \le 100\,°C\ (thermal\ safety\ limit)$$

$$g_2(x): F_B \le 6\,kg/s\ (pump\ capacity\ limit)$$

Equality constraints (steady-state process model):

$$h(x): F(C, T_R, F_B) = 0\ (6\ nonlinear\ algebraic\ equations$$

Variable bounds:

$$70\,°C \le T_R \le 100\,°C, 2\,kg/s \le F_B \le 6\,kg/s$$

where X_P and X_E are the mass fractions of products P and E at steady state (functions of x through the implicit steady-state solution), and F = $F_A + F_B$ is the total outlet flow.

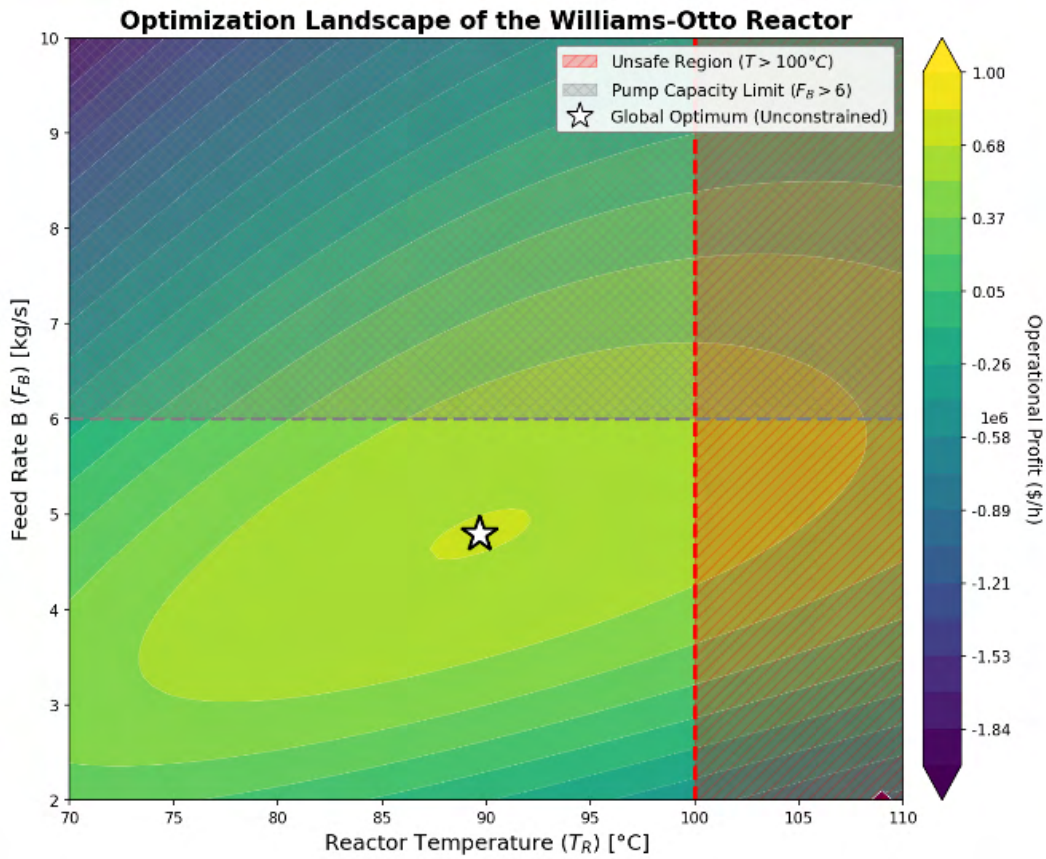

Figure 1. Optimization landscape over $T^R \in [70, 110]$°C and $F^B \in [2, 10]$ kg/s. Constraint regions indicated by red hatching (T>100°C) and gray hatching (F^B>6 kg/s). White star marks constrained global optimum.

The steady-state equations couple the decision variables to the component mass fractions, making the objective function implicitly dependent on the process model solution. The problem is solved using the Augmented Lagrangian Method with JAX automatic differentiation for exact gradient computation.

This is a two-dimensional NLP with 6 implicit equality constraints (steady-state model), 2 inequality constraints (safety and capacity), and box bounds on the decision variables. The steady-state equations couple the decision variables to the component mass fractions through the nonlinear kinetic model, making the objective function implicitly dependent on the process model solution.

The problem is solved using the Augmented Lagrangian Method with JAX automatic differentiation for exact gradient computation (see Section 5).

Steady-state calculation at base operating point converged in 5 Newton iterations to machine precision (residual 9.96×10^{-16}). Dynamic simulation over 1500s horizon using K=3 Radau IIA collocation completed in 1.132s wall-clock time, demonstrating 1325× real-time capability suitable for operator training simulators.

Constrained optimization via Augmented Lagrangian Method with JAX automatic differentiation converged in 0.894s (12 outer iterations) to global optimum: profit increased 1.9% from base case ($687, 529/h versus $674, 580/h). Both constraints remain inactive at optimum (temperature slack 10.3°C, flow slack 1.21 kg/s). Solution exhibits strong sensitivity to temperature as evidenced by tight contour spacing near optimum, justifying gradient-based optimization.

Comparison with Roberts (1979) published results [15] shows excellent agreement: temperature deviation +0.057°C (0.06%), flow deviation +0.004 kg/s (0.08%), profit deviation <0.01%. This validates all four framework modules against established literature benchmarks.

6.3 Propylene-Propane Splitter

High-purity separation of propylene from propane represents a major challenge in petrochemical industry due to close boiling points. At typical operating conditions (15–16 bar), these C3 hydrocarbons exhibit relative volatility α≈1.1 as computed from rigorous thermodynamics, requiring tall columns (150–200 stages) to achieve commercial purity specifications.

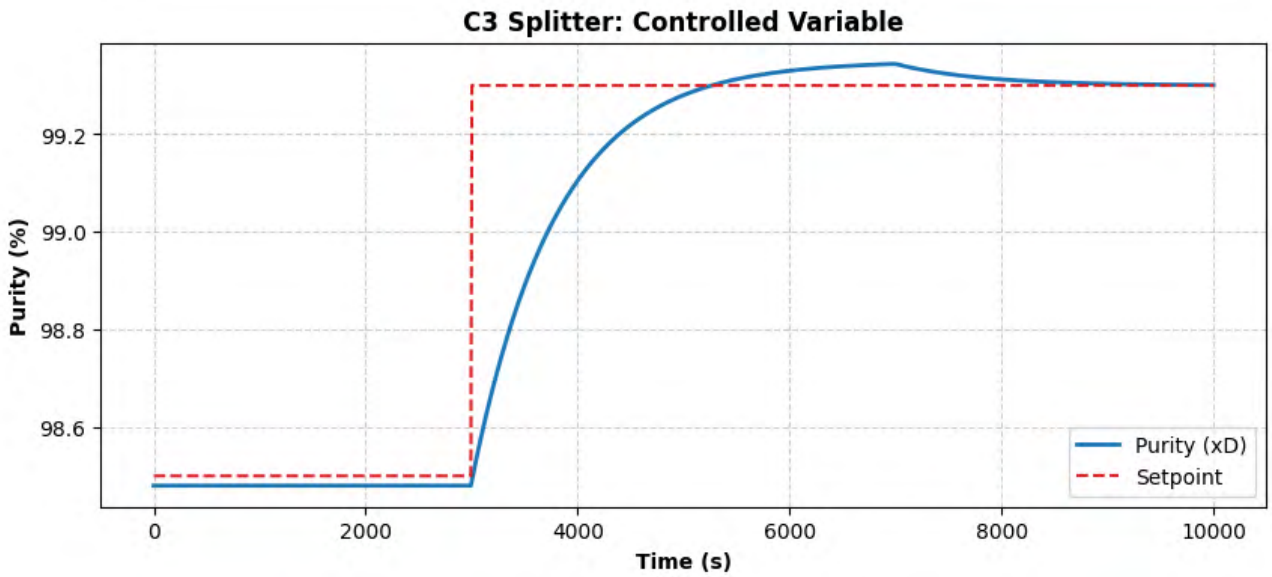

Figure 2. Dynamic response: propylene purity tracking

Table 5: Williams-Otto Comprehensive Performance.

Parameter	Base Case	Optimum	Units
Reactor Temperature TR	**85.00**	**89.70**	**°C**
Feed Rate FB	4.50	4.79	kg/s
Profit Π	674, 580	687, 529	$/h
Profit Increase	—	+1.9	%
Product P mass fraction XP	0.1098	0.1095	—
Product E mass fraction XE	0.2840	0.2906	—
Newton iterations	5	8	—
Residual norm	9.96×10^{-16}	1.2×10^{-15}	—
Optimization time	—	0.894	s
Dynamic simulation time (1500s)	1.132	—	s
Real-time speedup	1325×	—	—
Temperature deviation	—	+0.057	°C
Flow deviation	—	+0.004	kg/s
Profit deviation	—	<0.01	%

Table 6: C3 Splitter Complete Specification and Results.

Parameter	Value	Units
Number of stages	180	—
Feed stage	90	—
Top pressure (Ptop)	15.91	bar
Stage pressure drop (ΔP)	50	Pa
Feed flow rate (F)	100	mol/s
Feed temperature (Tfeed)	315	K
Propylene composition (z_C3H6)	0.55	mol frac
Binary interaction (k_{12})	0.0063	—
Controller gain (Kc)	85	—
Integral time constant (τI)	150	s
t=3000s: Setpoint step	98.5→99.3	%
Parameter	**Value**	**Units**
t=3000s: Feed flow ramp	100→120	mol/s
t=6000s: Composition ramp	0.55→0.60	—
Differential states (holdups)	360	—
Algebraic states (T, compositions)	180	—
Total DAE variables	540	—
Physical time simulated	10, 000	s
Wall-clock computation time	544.2	s
Real-time speedup factor	18.4×	—
Auto-selected DAE method	Dense Newton	—

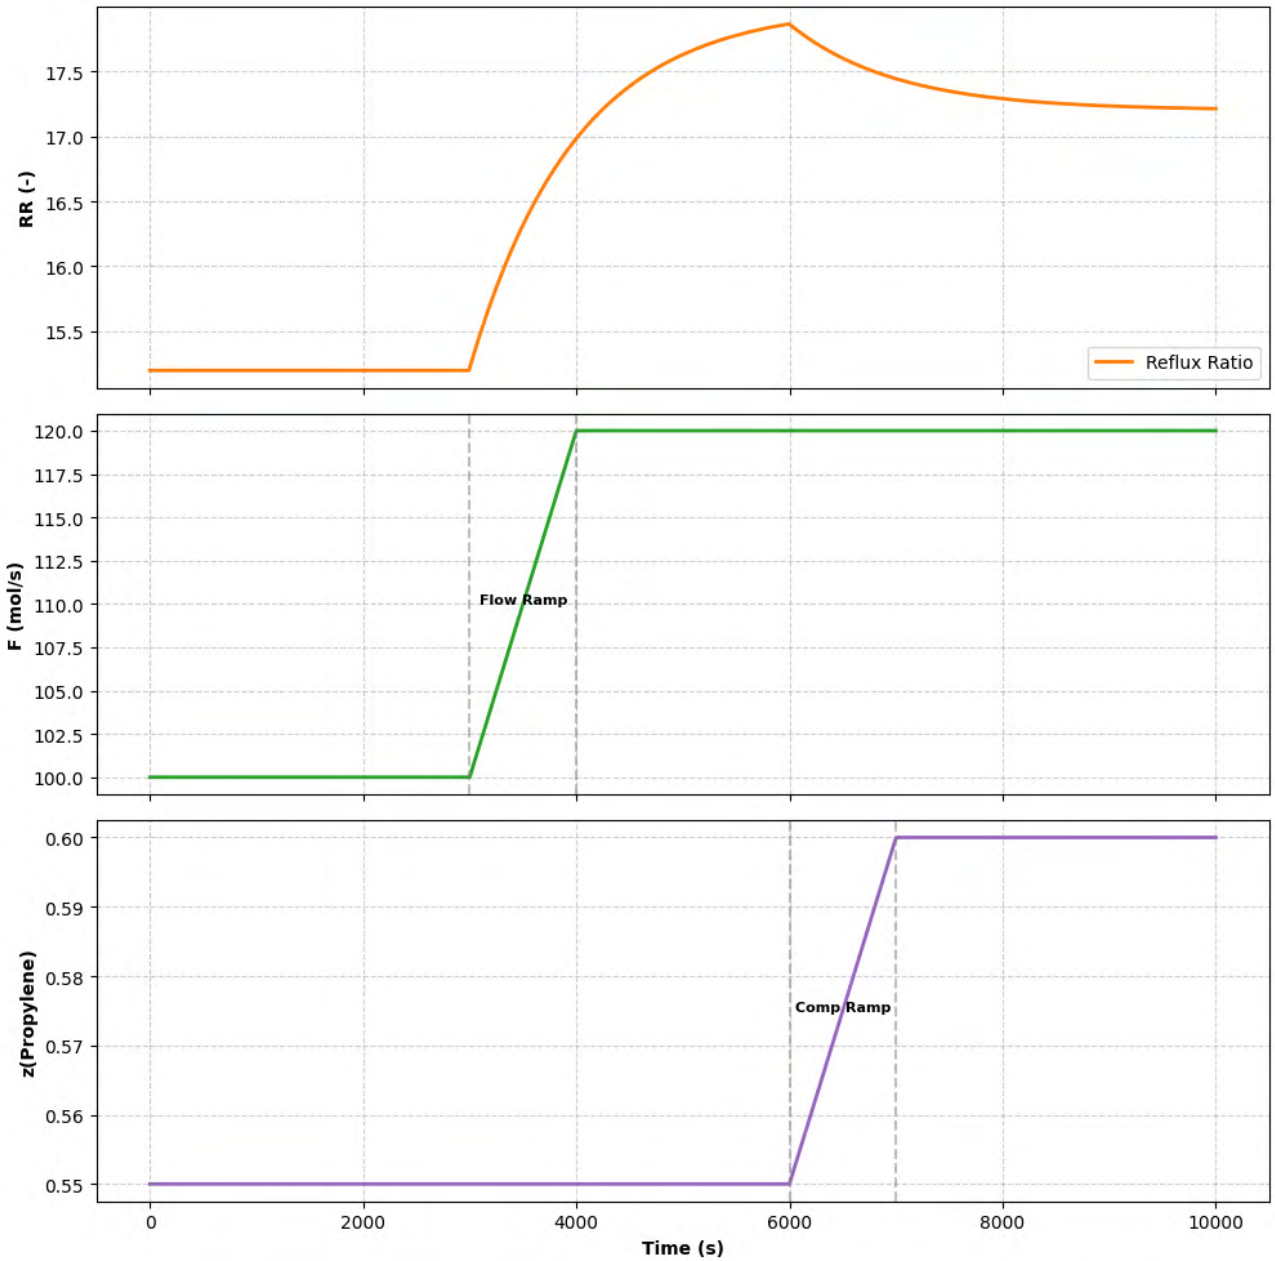

Figure 3 Dynamic response: top plot, reflux ratio manipulation (15.3 to peak 19.2); middle plot, feed flow ramp (+20%); bottom plot, feed composition evolution. Vertical dashed lines mark disturbance events.

The framework employs stage-specific rigorous VLE calculations rather than simplified constant relative volatility correlations. Pressure profile construction: Given top pressure Ptop = 15.91 bar and stage pressure drop ΔP = 50 Pa/stage, yielding pressures from 1591 kPa (top) to 1600 kPa (bottom). Rigorous VLE via Peng-Robinson EOS at each stage with binary interaction parameter k_{12}=0.0063 fitted to experimental C3–C3 VLE data. This rigorous approach captures pressure-dependent thermodynamic behavior during large flow disturbances.

Three sequential disturbances stress-test the controller: (1) t=3000s: concurrent setpoint step from 98.5% to 99.3% propylene purity synchronized with severe feed flow ramp from 100 to 120 mol/s over 1000s (+20% throughput increase); (2) t=6000s: feed composition ramp from 55% to 60% propylene over 1000s (+9% relative increase).

The controller successfully tracked the 99.3% setpoint with 0.4% overshoot and ~600s rise time. During the concurrent feed flow ramp, the controller increased reflux ratio from 15.3 to 19.2 (+25%) to maintain separation, demonstrating excellent disturbance rejection with

a maximum purity deviation of only −0.1%. The delayed composition disturbance at t=6000s was smoothly compensated, with the system stabilizing by t=8000s.

The DAE solver auto-selected dense Newton method based on system size (540 collocation variables < 800 threshold), completing 10, 000s physical time in 544.2s wall-clock time (18.4× real-time capability). With 540, 000 rigorous VLE calculations using Peng-Robinson thermodynamics with complex-step differentiation, this performance demonstrates suitability for control system design, operator training simulators, and multi-scenario optimization studies.

Additional validation: A 15-stage benzene-toluene-o-xylene (BTX) ternary column was simulated to validate multi-component capabilities beyond binary mixtures. The BTX column achieved 49× real-time performance under multiple disturbances, confirming that the framework's complex-step differentiation strategy and JAX vectorization scale effectively to systems with three or more components.

7. CONCLUSIONS

This work presents a Python/JAX framework for hydrocarbon process simulation validated through three applications: (1) nonlinear solver benchmarks with Newton-Krylov achieving 100% success on seven standard problems (n=2 to 5000), (2) Williams-Otto reactor optimization showing 0.06% deviation from published literature with 1325× real-time dynamic simulation capability, (3) 180-stage propylene-propane splitter achieving 18× real-time performance with 540, 000 rigorous VLE calculations under three concurrent disturbances.

Key implementations: Peng-Robinson thermodynamics with complex-step differentiation [12], avoiding nested JAX transformation issues, Newton/Broyden/Newton-Krylov solvers with Curtis-Reid auto-scaling [5], Radau IIA DAE solver with auto-selection between dense and matrix-free methods [8, 9], Augmented Lagrangian optimization with JAX automatic differentiation [11].

Current limitations: Semi-rigorous distillation formulation without tray-by-tray energy equations, Peng-Robinson equation of state limited to nonpolar hydrocarbon systems, Newton-Krylov methods may require parameter tuning for optimal performance.

Why a JAX-friendly DAE solver? While several Python libraries exist for DAE systems (Assimulo [17], GEKKO [18], CasADi [19]), they typically lack one or more of: (1) native automatic differentiation for exact Jacobians, (2) JIT compilation for near-C++ performance, (3) vectorization primitives (vmap) for batch stage-wise calculations, or (4) transparent pure-Python implementations enabling straightforward debugging and customization.

Future directions: (1) Rigorous distillation models implementing MESH equations and/or the DDPA model of Matias et al (2018) [16], (2) extended thermodynamics including CPA and NRTL/UNIQUAC activity models, (3) state estimation integration (EKF, UKF, MHE) for real-time applications, (4) GPU acceleration evaluation for large-scale simulations, (5) Sparse matrix handling JAX-friendly, as an alternative for Newton-Krylov

ACKNOWLEDGEMENTS

The authors acknowledge the support of the Chemical Engineering Department at Universidade de São Paulo.

AUTHOR IDENTIFIERS

Author ORCIDs:
Sanz CC: 0009-0008-4267-3410
Le Roux G: 0000-0001-7227-5271

REFERENCES

1. J. Bradbury et al., "JAX: composable transformations of Python+NumPy programs, " 2018. http://github.com/google/jax
2. OpenXLA Project, "XLA: Optimizing Compiler for Machine Learning." https://openxla.org/xla
3. Peng D-Y, Robinson DB. A New Two-Constant Equation of State. Ind Eng Chem Fundam 15(1):59-64 (1976)
4. Cox KR, Chapman WG. The properties of gases and liquids, 5th edition by bruce E. poling (university of toledo), john M. prausnitz (university of california at berkeley), and john P. o'connell (university of virginia). mcgraw-hill: new york. 2001. 768 pp. $115.00. ISBN 0-07-011682-2.. J. Am. Chem. Soc. 123:6745-6745 (2001). https://doi.org/10.1021/ja0048634
5. CURTIS AR, REID JK. On the automatic scaling of matrices for gaussian elimination. IMA J Appl Math 10:118-124 (1972). https://doi.org/10.1093/imamat/10.1.118
6. Knoll DA, Keyes DE. Jacobian-free newton–krylov methods: a survey of approaches and applications. Journal of Computational Physics 193:357-397 (2004). https://doi.org/10.1016/j.jcp.2003.08.010
7. Eisenstat SC, Walker HF. Choosing the forcing terms in an inexact newton method. SIAM J. Sci. Comput. 17:16-32 (1996). https://doi.org/10.1137/0917003
8. Hairer E, Wanner G. Order conditions for index 2 DAE. Springer Series in Computational Mathematics :506-518 (1996). https://doi.org/10.1007/978-3-642-05221-7_35

9. Ascher UM, Petzold LR. Computer Methods for ODE and DAE, SIAM (1998)
10. Nocedal J, Wright SJ. Numerical Optimization, 2nd Ed., Springer (2006)
11. Conn AR, Gould NIM, Toint PhL. Trust-Region Methods, SIAM (2000)
12. Martins JRRA, Sturdza P, Alonso JJ. The complex-step derivative approximation. ACM Trans. Math. Softw. 29:245-262 (2003). https://doi.org/10.1145/838250.838251
13. Moré JJ, Garbow BS, Hillstrom KE. Testing unconstrained optimization software. ACM Trans. Math. Softw. 7:17-41 (1981). https://doi.org/10.1145/355934.355936
14. Williams TJ, Otto RE. A generalized chemical processing model for the investigation of computer control. Trans. AIEE, Part I: Comm. Electron. 79:458-473 (1960). https://doi.org/10.1109/tce.1960.6367296
15. Roberts PD. An algorithm for steady-state system optimization. Int J Syst Sci 10(7):719-734 (1979)
16. Matias JOA, Ebrahimpour M, Le Roux GAC. Development and validation of a reduced-index dynamic model of an industrial high-purity column. Ind. Eng. Chem. Res. 57:1531-1546 (2018). https://doi.org/10.1021/acs.iecr.7b02789
17. Andersson C, Führer C, Åkesson J. Assimulo: a unified framework for ODE solvers. Mathematics and Computers in Simulation 116:26-43 (2015). https://doi.org/10.1016/j.matcom.2015.04.007
18. Beal LDR, Hill DC, Martin RA, Hedengren JD. GEKKO optimization suite. Processes 6:106 (2018). https://doi.org/10.3390/pr6080106
19. Andersson JAE, Gillis J, Horn G, Rawlings JB, Diehl M. Casadi: a software framework for nonlinear optimization and optimal control. Math. Prog. Comp. 11:1-36 (2018). https://doi.org/10.1007/s12532-018-0139-4

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Digital Twin Supported FAIR Electronic Lab Notebooks for Simulated Experiments

Amy Koch[a*], Isabell Viedt[a], Leon Urbas[a]

[a] Technische Universität Dresden, Process-to-Order Group
* Corresponding Author: amy.koch@tu-dresden.de

ABSTRACT

The use of equipment digital twins of standardized, multi-purpose units can accelerate process development and reduce experimental effort. Experimental data are essential not only for identifying critical process parameters and enabling model-based methods within a Quality by Design framework, but also for constructing and validating the simulation models that describe digital twin behavior. To achieve high-fidelity and robust predictive models, structured concepts are required to manage metadata and process-, product-, and resource-specific information exchanged between physical and digital twins. Electronic lab notebooks (ELNs), which contextualize experimental data, must therefore be structured and standardized to ensure interoperability and seamless data exchange. For integration into digital twin workflows and process transfer between equipment instances of the same category, ELNs must comply with FAIR (Findable, Accessible, Interoperable, Reusable) data principles. This work proposes a Semantic Integration Pattern for FAIR-compliant ELNs to document simulated experiments using domain-specific, vendor-independent standards such as DEXPI 2.0 and ISA 88/IEC 61512-1. DEXPI 2.0 supports standardized descriptions of equipment and process steps, while ISA 88/IEC 61512-1 structures experimental recipes. The resulting FAIR ELN Semantic Integration Pattern enables standardized data exchange between ELNs and digital twin tools. Its applicability is demonstrated through an illustrative example documenting simulated saponification experiments in gPROMS® using simple batch reactor model. The results will support subsequent real experiments conducted in a 2L stirred tank unit.

Keywords: digital twins, electronic lab notebooks, gProms, Simulation, Process Design

MOTIVATION

Digital twins can be utilized to support process development, optimization, and scale-up. In modular and multi-purpose process environments, application of digital twins of the equipment enables the reuse of simulation models across different process applications and equipment instances, thereby reducing experimental effort and time-to-process. However, the predictive quality of digital twins critically depends on the availability of high-quality experimental and simulation data together with rich contextual metadata.

Electronic lab notebooks (ELNs) play a central role in capturing experimental data and associated metadata. In current practice, ELNs are often used as flexible documentation tools with limited structure, leading to heterogeneous data representations, incomplete metadata, and poor interoperability. As a result, experimental data rapidly lose their value for reuse, model validation, and transfer between projects, tools, and organizations.

To address these challenges, the FAIR data principles [1] provide a widely accepted framework for sustainable scientific data management. Applying FAIR principles to ELNs is particularly important in digital twin workflows [2], where experimental data from both the laboratory and simulation models must be exchanged between modeling tools, automation systems, and laboratory or pilot-scale equipment. This paper focuses on the documentation of experimental data from simulation models (simulated experiments), which are a key component of model-based process development and Quality by Design (QbD) methodologies.

In this paper, first the relevant background is presented and key domains specific ontology, metadata

models, and data models presents. Next, a concept for a structured Semantic Integration Pattern for FAIR-compliant ELNs for simulated experiments based on established, vendor-independent domain specific metadata guidelines are proposed. This Semantic Integration Pattern is explained for a saponification reaction in a 2L batch process using a gPROMS® Process simulation.

METHODOLOGY

Digital Twins in Modular Plants

Digital twins for modular plants have been discussed in terms of a single PEA [3] and for process transfer [2]. Here, the common approach is to use a Product-Process-Resource-Model to structure the digital twin [4] and leverage the knowledge distribution between the equipment vendor and the owner/operator [3] as well as for process transfer between modules [2]. A similar approach is used in the domain of discrete manufacturing with a focus on standardized description of tasks and skills of machinery, e.g. [6], [7], [8]. For modular plants, these equipment-based digital twins can be used to facilitate faster time-to-process. These equipment -based digital twins are integrated into digital twin workflows for process development which typically span multiple asset lifecycle phases, including process concept design, model development, parameter identification, optimization, and scale transfer. In modular plant concepts, digital twins are often associated with standardized equipment classes rather than single physical assets. This enables the transfer of process knowledge between different equipment instances of the same category.

Within such workflows, simulated experiments are used extensively to explore design spaces, analyze process dynamics, and support model-based design of experiments. The results of these simulations must be traceable, reproducible, and interpretable to be reused for decision-making and subsequent validation with physical experiments.

FAIR DATA

A critical consideration for the effective utilization and reuse of experimental data is adherence to the FAIR data principles [1]. Alignment with these principles ensures that metadata is sufficiently enriched with descriptive information that captures the context in which the data were collected, including data quality, collection conditions, and other relevant characteristics. In accordance with FAIR requirements metadata must “be richly described with a plurality of accurate and relevant attributes.” Within the context of ELNs and simulated experiments, these principles imply:

- **Findable:** Experiments and datasets can be located based on product, process, resource, purpose, or author.
- **Accessible:** Data and metadata can be retrieved via well-defined access mechanisms, including access restrictions where required.
- **Interoperable:** Data and metadata use shared vocabularies, formats, and semantics to enable interpretation by humans and machines.
- **Reusable:** Data are sufficiently well-described to allow repetition of experiments and reuse in new contexts.

Applying FAIR principles requires explicit information models and controlled vocabularies rather than ad-hoc documentation.

Electronic Lab Notebooks

The increasing shift towards digitalization in the laboratory extends beyond the use of electronic devices to encompass the networking of experimental workflows and the management of research data in accordance with the FAIR principles. Electronic laboratory notebooks (ELNs) support structured, machine- and human-readable metadata capture, enable direct integration of measurement data, and facilitate data traceability, reuse, and long-term archiving.

NFDI4Chem [9] provides guidelines for managing research data in chemistry and chemical domain specific guidelines for the following aspects: type of experiment, planning of experiments, documentation of experiments, data producing methods, data analysis, and publication.

Simulation-Supported DoE

Model-based Design of Experiments (MBDoE) [10] integrates simulation models, typically formulated as differential–algebraic equations, with classical Design of Experiments to systematically plan informative experiments. Simulation models incorporating prior knowledge of process structure and parameter estimates are used to identify experimental variables with high sensitivity to model parameters, thereby maximizing information gain. This model-driven approach enables efficient identification of optimal process conditions while reducing experimental effort, making it particularly valuable in time- and resource-constrained process development.

Information Models and Standards

In addition to the FAIR Data requirements [1], other ontology and data model standards exist to describe the high-level metadata.

Dublin Core Metadata Terms

The **Dublin Core Metadata Element Set** [11] is a widely adopted standard for describing digital and physical resources to support discovery, interoperability, and long-term access. It consists of 15 core metadata

elements, such as *title, creator, subject, description, publisher,* and *date*, which are designed to be simple, generic, and domain independent. Dublin Core is formally standardized as ISO 15836. The standard supports extensibility through qualified elements and encoding schemes, allowing more precise descriptions when required by specific communities.

ProMetaS

ProMetaS [12] is a comprehensive metadata schema designed to document and describe data generated in chemical process engineering and industry to support interoperability and reuse. The schema was developed within the **KEEN research project** to address the diversity and complexity of industrial and research data, improving findability and understanding. In line with FAIR principles ProMetaS builds on international metadata standards (e.g., Dublin Core) and was iteratively refined with domain experts to include only relevant metadata elements. The applicability of this schema has been demonstrated for real use cases in process data collection and management in the KEEN project.

The information model divides the data in the process industry into categories: plant info, process info, chemical, sensor, image control valve, time series, and machine learning.

DEXPI 2.0 for Process and Equipment Descriptions

DEXPI (Data Exchange in the Process Industry) is an initiative of the process industry to create a universal data model for the information described in flow diagrams along the asset lifecycle. DEXPI Plant [13] provides a standardized, vendor-independent representation of process equipment and piping and instrumentation diagrams. DEXPI Plant can be used to describe equipment classes and topology. DEXPI Process serves to describe the block and process flow diagram type information which can be mapped to simulated data, particularly for flowsheet simulation. The DEXPI 2.0 und data model serves to provide a consistent reference between simulation models and physical assets.

ISA-88 / IEC 61512-1 for Recipe Descriptions

In addition to the process description, it is necessary to describe the experimental procedure, or "recipe." ISA-88 which is formally standardized in Europe with the IEC 61512 [14] specifies a hierarchical recipe model consisting of procedures, unit procedures, operations, and phases. This structure is well suited for documenting experimental workflows in both simulation and laboratory contexts. In the FAIR ELNs, ISA-88 is used to: (1) structure the experimental procedure of simulated experiments, (2) define input parameters, control actions, and execution order, and (3) support repeatability and comparison between experiments.

PEA Harmonization for Modular Plants

Additionally, for modularization further requirements with respect to the documentation are needed. Here, it is important to be able to describe both process and equipment in a modular way. To address this challenge, the Process-to-Order Group has introduced a PEA Harmonization [17] which aligns the hierarchical modular terms from process engineering [16] and automation engineering [17] and links it to the process steps in the DEXPI Process specification [18]. This harmonization enables a direct link from the domain specific terms in modularization to domain specific process steps used in the process industry to describe process phenomena in different levels of abstraction.

REQUIREMENTS

Based on experience from experimental and simulation studies, several challenges motivate this work. First, the same process step can be carried out in models as well as different apparatus configurations, leading to ambiguous documentation. This is an issue for both real experiments as well as simulation runs. A second issue is that equipment of the same category may exist in multiple instances with different geometries, sensors, or performance characteristics. This is important because depending on the purpose of the simulation, the same simulation models can be applied to different equipment instances, but with different parameterization. Another issue is sensors. Sensors with different accuracies and variances are often used without explicit documentation of uncertainty models, however, uncertainty models and estimates play a significant role on the predicted outcome of a simulation. Another challenge is the data itself, experimental data may originate from both online (time-series) and offline (sample-based) sources, which are rarely integrated consistently. These challenges highlight the need for structured, semantically rich ELNs.

Another issue to address is the separation of the domain specific data. Here, it is necessary to have a separation of data and metadata: Raw simulation results, derived results, and metadata must be clearly distinguished. Next, it is necessary to structure the domain specific data. As in the ProMetaS [12], it is necessary to have structured description for the different domains. For ELNs specifically, this means separate descriptions of process, product, and resources. Therefore, equipment topology, process steps, and material information must be described using standardized models.

In addition to the above listed challenges, it is further necessary to delineate between the type of simulation which is being carried out. Here, the follow categories are suggested:

- Static simulations for design space identification,

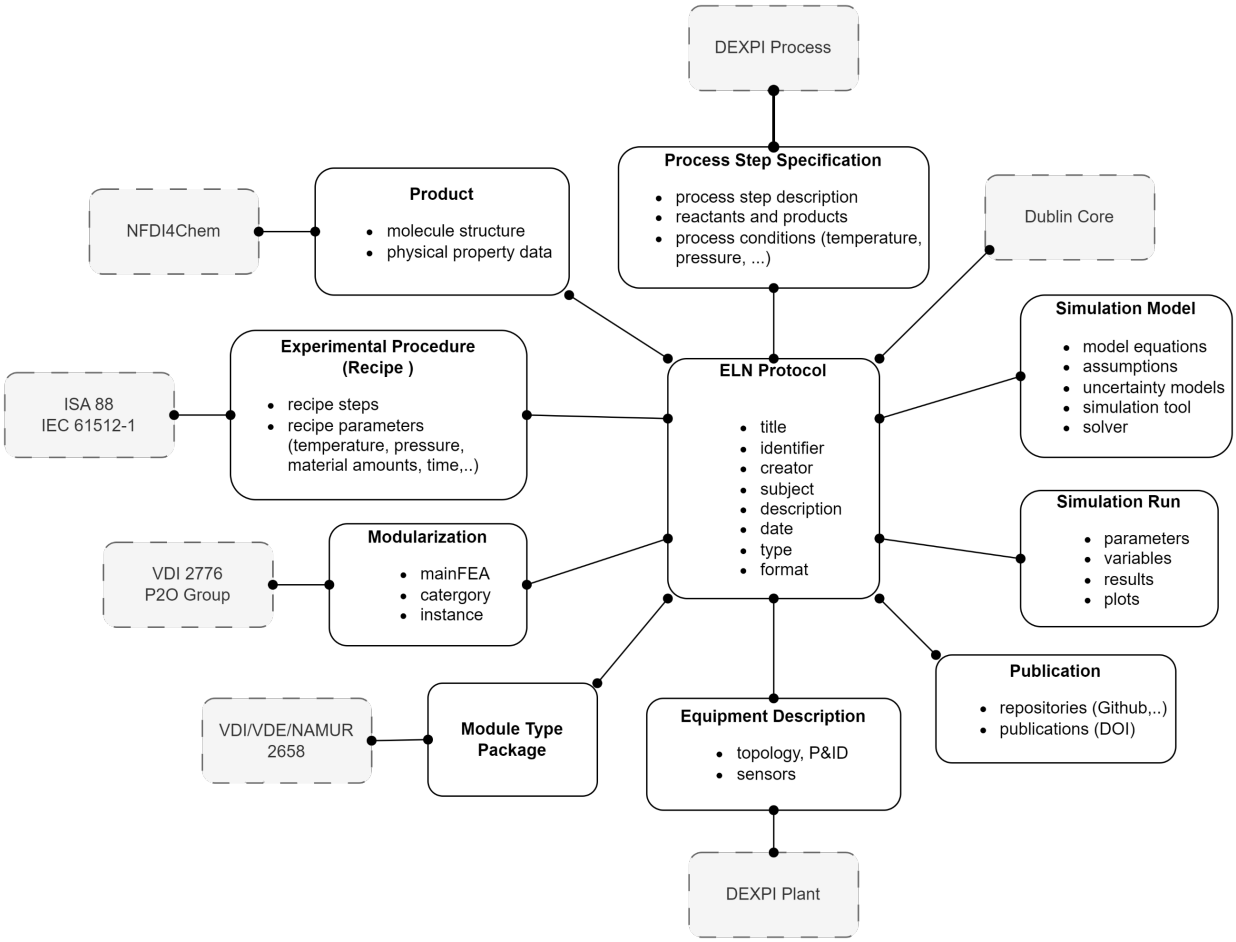

Figure 1. ELN Semantic Integration Pattern for documentation of simulated experiments.

parameter estimation, process design, and optimization

- Dynamic simulations for system behavior analysis, model discrimination, and control loop design
- Stochastic simulations for model-based experimental design and uncertainty analysis/reduction

ELN SEMANTIC INTEGRATION PATTERN FOR SIMULATED EXPERIMENTS

To address the above-mentioned challenges, an ELN semantic integration pattern is proposed for documentation of simulation runs for application in digital twin workflows for modular plants. Here, the Dublin Core [11] was utilized for the ELN Protocol, ISA 88 for the experimental procedure, and ProMetas for domain specific categories. The following domains for ELNs documenting simulated experiments are identified as:

- **ELN Protocol** for the explicit identification of experiments. Each ELN entry must have a unique and persistent identifier and should include metadata according to the Dublin Core [11]. Specifically, experiment title, unique identifier, creator, subject (purpose of experiment), description, data, type of simulation, format of the simulation, and the source.
- The **Product** must include a description of the chemicals. Here, information such as molecule structure, physical property data, physical property models, etc. is included. The NFDI4Chem [9] provides domain specific guidelines depending on the chemical domain.
- The **Process Step Specification** describes the type of process step which is considered and utilized terms from the DEXPI Process specification
- The **Equipment Description** to describe the topology, P&ID and available sensors in alignment with the DEXPI Plant [13] .
- The **Experimental Procedure** must be documented

in a reproducible, machine-readable form. Definition of the experimental procedure using ISA-88/IEC 61512-1 [14] structures to include recipe steps and parameters.

- **Modularization:** Functional Equipment Assemblies (FEAs), module category, and instance as outlined in the PEA Harmonization.
- **Module Type Package:** Link to the data in the real plant along with the Human Machine Interface (HMI).
- **The Simulation Model** to document versions, solvers, and assumptions must be recorded: Documentation of assumptions, observations, and derived quantities.
- The **Simulation Runs** to document the parameters, variables, results, and plots unique to the simulation run.
- **Publication:** Storage of ELN entries in a repository enabling search, access control, and reuse.

This ELN Semantic Integration Pattern ensures that simulated experiments are documented consistently and can be FAIRly [1] integrated into digital twins for modular plants.

ILLUSTRATIVE EXAMPLE

To demonstrate the utilization of the ELN Semantic Integration Pattern, a small illustrative example is provided. Here, a simple batch model created in gPROMS® Process was used to validate kinetics of a saponification reaction at $T = 313\,K$ according to the following equation:

$$CH_3COOC_2H_5 + NaOH \rightarrow CH_3COONa + C_2H_5OH$$

The simulation model employs three equations: Equation 1 describes the concentration profile of each reaction component i over time with the assumption of constant volume V, Equation 2 describes the Arrhenius rate constant kinetics, and Equation 3 describes the reaction rate.

$$\frac{\mathrm{dCi}}{\mathrm{dt}} = -\mathrm{r} * \mathrm{V} \tag{1}$$

$$k = k_{i,\infty} \exp\left(\frac{-E_{A,i}}{R\,T}\right) \tag{2}$$

$$r = kC(1)C(2) \tag{3}$$

Next, for the simulation run values are assumed for the parameters of the Arrhenius rate expressions:

$$k_\infty = 3.5e^7 \frac{L}{mol\ s}$$

$$E_A = 48000\ \frac{J}{mol}$$

Additionally, it assumed that an equimolar ratio of the reactants, 1 = ethyl acetate ($CH_3COOC_2H_5$) and 2 = sodium hydroxide ($NaOH$) is present in the reactor at the start of the reaction time. The results and plots from the simulation run are stored in a gPROCESS Case File.

Figure 2 shows a delimited ELN Protocol in XML utilizing the ELN Semantic Integration Pattern. The ELN Protocol highlights the following domains from the ELN Semantic Integration Pattern from Figure 1: ELN Protocol, Modularization, Simulation Model, and Simulation Run. Here, an example of the general meta data in accordance with the Dublin Core [11] is provided in the ELN Protocol section. The **Modulization** description of M22 which is an instance of a 2L Glass Stirred Tank Unit whose main Functional Equipment Assembly is a reactor. Next, the **Simulation Model** is described according to equations (1) to (3), the modelling assumption of constant volume, along with a list of parameters and units, as well is the solver and any uncertainty assumptions. Lastly, for the **Simulation Run** domain estimates of the Arrhenius parameters is provided as well as the initial amount of moles of the reaction constitutes in the vessel at the start of the reaction. The name of the Case File where the simulation results are stored is also provided in the XML code.

CONCLUSION AND OUTLOOK

The proposed FAIR ELN Semantic Integration Pattern addresses key challenges in documenting simulated experiments for digital twin workflows in modular plants. By combining domain-specific standards with general metadata vocabularies, both semantic richness and broad interoperability are achieved. The integration of DEXPI Plant [13], DEXPI Process [18], ISA-88/IEC 61512-1 [14], and Dublin Core [11] provides a structured, vendor-independent approach to experimental documentation which can be extended to modularization and simulation. However, it should be noted that this work focuses on the concept, thus the illustrative example in Figure 2 is underspecified. This example is not intended to be exhaustive in scope; rather, it aims to describe the ELN documentation method for a simulation run, not the subsequent utilization of the simulated results within the digital twin.

Remaining challenges comprise the effort required for initial model alignment, the establishment of controlled semantic vocabularies, and the development of integration tools. However, the long-term benefits in terms of data reuse, reproducibility, and model fidelity outweigh these initial limitations. Future work will therefore focus on the implementation of tool-supported approaches, the automation of metadata capture, and the extension of the framework to facilitate digital twin–supported information exchange based on documented simulation runs.

```
<ELN Protocol>
        <Title = "SimpleBatchRunSaponification001" Identifier = "BatchSaponifcation001" Creator = "Amy Koch" Subject
        "Saponification Kinetic Validation"     Description = "Validation of Reaction Kinetics for Ethyl Acetate Saponification
        Reaction" Date = "25.11.2025" Type = "Simulated Experiment/>
<Modularization>
        <mainFEA = "Reactor" Category = "Stirred Tank Unit" Instance = "2L Glass Stirred Tank Unit"/>
</Modularization>
<Simulation Model>
        <simulation tool = "gPROMS Process Academic Research Version 2025.2.0/>
        <Model ="Simple Batch Model"/>
                <Equation>
                        <!-- Species balances -->
                        <Equation variable="C(1)">
                        d(Ci * V)/dt = -r * V
                 </Equation>
                <Kinetics>
                        <!-- Arrhenius rate constant -->
                        <RateConstant name="k">
                                 k0 * exp(-Ea / (R * T))
                        </RateConstant>
                         <!-- Reaction rate: CH3COOC2H5 (1) + NaOH (2) → CH3COONa (3) + C2H5OH (4) -->
                        <ReactionRate name="r" units="mol/L/s">
                         k * C(1) * C(2)
                         </ReactionRate>
                </Kinetics>
                <Parameters>
                        <Parameter name="R" units="J/mol/K"/>
                        <Parameter name="kinf" units="L/mol/s"/>
                        <Parameter name="Ea" units="J/mol" >
                        <Parameter name="T" units="K"/>
                </Parameters>
                <Variables>
                        <Variable name="M22.N(1)" units="mol/L"/>
                        <Variable name="M22.N(2)" units="mol/L"/>
                        <Variable name="M22.N(3)" units="mol/L"/>
                        <Variable name="M22.N(4)" units="mol/L"/>
                        <Variable name="Volume"  units="L"/>
                </Variables>
        <Assumptions = "constant volume"/>
                        <!-- Volume balance (constant volume) -->
                        <Equation variable="V">
                                dV/dt = 0
        </Assumptions>
        <Uncertainty = "none"/>
        <Solver = "DAEBDF"/>
  <Simulation Run>
        <Parameters>
                <Parameter name="R" units="J/mol/K" value="8.314"/>
                <Parameter name="kinf" units="L/mol/s" value="3.57e7"/>
                <Parameter name="Ea" units="J/mol" value="48000"/>
                <Parameter name="T" units="K" value="313.15"/>
        </Parameters>
        <Variables>
                <Variable name="M22.N(1)" value="0.020043"/>
                <Variable name="M22.N(2)" value="0.020043"/>ee/
                <Variable name="M22.N(3)" value="0"/>
                <Variable name="M22.N(4)" value="0"/>
                <Variable name="Volume" value="0.22"/>
        </Variables>
        <Results = "Simulation_Batch_Recipe_20251125_132208.gSCP-PB"/>
</Simulation Model>
</ELN Protocol>
```

Figure 2. Example XML documentation of the ELN Protocol for a simulated exeriment.

ACKNOWLEDGEMENTS

The authors would like to thank our partners Merck Electronics KGaA and Siemens AG in the industrial research project Lab2Plant for the fruitful discussions as well as P2O-Lab Learning Factory at TU Dresden for information about their modules.

AUTHOR IDENTIFIERS

Author ORCIDs:
Koch A: 0009-0009-4740-0590
Viedt I: 0000-0002-5814-5128
Urbas L: 0000-0001-5165-4459

REFERENCES

1. M. D. Wilkinson *et al.*, “The FAIR Guiding Principles for scientific data management and stewardship, ” *Scientific Data*, vol. 3, no. 1, p. 160018, Mar. 2016, doi: 10.1038/sdata.2016.18.
2. Koch A, Viedt I, Bamberg A, Urbas L. Applying quality by design to digital twin supported scale?up in modular plants. Chemie Ingenieur Technik 97:1127-1137 (2025). https://doi.org/10.1002/cite.70016
3. J. Mädler, I. Viedt, J. Lorenz, and L. Urbas, “Requirements to a digital twin-centered concept for smart manufacturing in modular plants considering distributed knowledge, ” in *Computer Aided Chemical Engineering*, vol. 49, Y. Yamashita and M. Kano, Eds., in 14 International Symposium on Process Systems Engineering, vol. 49., Elsevier, 2022, pp. 1507–1512. doi: 10.1016/B978-0-323-85159-6.50251-7.
4. Bamberg A, Urbas L, Bröcker S, Bortz M, Kockmann N. The digital twin – your ingenious companion for process engineering and smart production. Chem Eng & Technol 44:954-961 (2021). https://doi.org/10.1002/ceat.202000562
5. A. Koch, J. Mädler, A. Bamberg, and L. Urbas, “Towards PEA Matching from Simulation as Part of a Digital Twin Concept for Scale-Up in Modular Plants, ” in *Computer Aided Chemical Engineering*, vol. 53, Elsevier, 2024, pp. 595–600. doi: 10.1016/B978-0-443-28824-1.50100-9.
6. Schleipen M, Lüder A, Sauer O, Flatt H, Jasperneite J. Requirements and concept for plug-and-work. at - Automatisierungstechnik 63:801-820 (2015). https://doi.org/10.1515/auto-2015-0015
7. Pfrommer J, Stogl D, Aleksandrov K, Escaida Navarro S, Hein B, Beyerer J. Plug & produce by modelling skills and service-oriented orchestration of reconfigurable manufacturing systems. at - Automatisierungstechnik 63:790-800 (2015). https://doi.org/10.1515/auto-2014-1157
8. Pfrommer J, Klein JF, Wurster M, Rapp S, Grauberger P, Lanza G, Albers A, Matthiesen S, Beyerer J. An ontology for remanufacturing systems. at - Automatisierungstechnik 70:534-541 (2022). https://doi.org/10.1515/auto-2021-0156
9. Steinbeck C, Koepler O, Herres-Pawlis S, Bach F, Jung N, Razum M, Liermann JC, Neumann S. Nfdi4chem—a research data network for international chemistry. Chemistry International 45:8-13 (2023). https://doi.org/10.1515/ci-2023-0103
10. Franceschini G, Macchietto S. Model-based design of experiments for parameter precision: state of the art. Chemical Engineering Science 63:4846-4872 (2008). https://doi.org/10.1016/j.ces.2007.11.034
11. ISO 15836-1:2017: Information and documentation — The Dublin Core metadata element set, May 2017.
12. Sherpa L, Müller-Pfefferkorn R, Tolksdorf G, Khaydarov V, Wiedau M, Urbas L. Prometas – a metadata schema for process engineering and industry. Chemie Ingenieur Technik 95:1041-1048 (2023). https://doi.org/10.1002/cite.202200225
13. Wiedau M, von Wedel L, Temmen H, Welke R, Papakonstantinou N. ENPRO data integration: extending DEXPI towards the asset lifecycle. Chemie Ingenieur Technik 91:240-255 (2019). https://doi.org/10.1002/cite.201800112
14. Hughes RC. Preparation by IEC and CIGRE for IEC publication: “IEC 60-2: 1994. high voltage test techniques part 2. measuring systems”. IEE Colloquium on NAMAS Accreditation for High Voltage Measurements and Tests 1997:1-1 (1997). https://doi.org/10.1049/ic:19970280
15. Urbas L. 2nd p2o-symposium 2025 : process-to-order. Tagungsband 2nd P2O- Symposium 2025 : Process-to-Order (P2O) Group : (2025). https://doi.org/10.25368/2025.292
16. VDI 2776-2:2024, Verfahrenstechnische Anlagen - Modulare Anlagen - Design modularer Anlagen, Richtlinie 2, Jan. 2024.
17. Maurmaier ASM. Dezentralisierte kommunikationsarchitektur für direkte modul-modul-interaktion auf basis von standardisierten konstrukten der vdi/vde/namur 2658. Automation 2019 :343-360 (2019). https://doi.org/10.51202/9783181023518-343
18. Cameron DB, Otten W, Temmen H, Hole M, Tolksdorf G. DEXPI process: standardizing interoperable information for process design and analysis. Computers & Chemical Engineering 182:108564 (2024). https://doi.org/10.1016/j.compchemeng.2023.108564

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Nonconvex Robust Optimization for Process Design with Artificial Neural Networks Embedded

Diego Izquierdo González[†], Basit Adeogun[†], Yuhui Yin[a], Vassilis M. Charitopoulos*

Department of Chemical Engineering, The Sargent Centre for Process Systems Engineering, University College London (UCL), Torrington Place, WC1E 7JE, UK
[†] Equally contributing authors.
* Corresponding Author: v.charitopoulos@ucl.ac.uk

ABSTRACT

Artificial neural networks (ANNs) have emerged as powerful surrogate models in process design and optimisation, capable of capturing complex nonlinear process behaviour while significantly reducing computational cost compared to detailed first-principles simulations. However, ANN prediction errors in safety-critical applications can lead to suboptimal or vulnerable designs, necessitating rigorous treatment of approximation uncertainties. While probabilistic approaches exist for surrogate-based decision making, risk-averse contexts that require formal robustness guarantees face a fundamental challenge: the nonconvex nature of ANN-embedded models hinders the employment of standard robust optimisation methods. To this end, in this work we explore the global robust optimisation of process design problems with embedded ANNs. A robust spatial branch-and-bound (RsBB) algorithm to achieve global optimality is proposed while enforcing constraint satisfaction across all uncertainty realisations. This approach integrates spatial branch-and-bound with adversarial cuts generation, concurrently searching for the global optimum while ensuring robust feasibility despite ANN prediction uncertainties. We test the proposed algorithm on the ANN-embedded optimisation of a cumene manufacturing process, across a range of conservatism levels. The algorithm computes solutions within the prescribed optimality tolerance while quantifying the robustness-performance trade-off.

Keywords: Robust optimisation, Process design, Hybrid modelling, Global optimisation, Machine learning-based optimisation

INTRODUCTION

Recent developments in Process Systems Engineering (PSE) have accelerated the use of optimisation to support decision-making in process design and operations. However, many process design problems remain computationally challenging [1]. High-fidelity, first-principles models can be expensive to evaluate and strongly nonlinear, which makes repeated model evaluations during optimisation impractical. This has motivated approaches that retain sufficient model fidelity while achieving more tractable optimisation problems [2, 3]. In this context, data-driven surrogate models, particularly machine-learning (ML) models such as artificial neural networks (ANNs), have become attractive because they can approximate complex input-output relationships learned from simulation or plant data while enabling fast, differentiable evaluations suitable for embedding in optimisation problems [4, 5].

Despite these advantages, surrogate-based optimisation introduces a critical challenge, as uncertainty is intrinsic to data-driven models. In practice, ANN surrogates carry inherent uncertainty due to prediction error, limited or biased training data and variability in training procedures and model architecture choices [2]. Consequently, designs that appear optimal under nominal surrogate predictions may perform poorly, or even violate constraints, when evaluated on the real system, especially near tight feasibility boundaries or when extrapolating beyond the training region. Robust optimisation provides a way to address this by explicitly accounting for uncertainty during optimisation, aiming to identify designs that remain feasible and perform well across a prescribed uncertainty set [6, 7]. This is particularly relevant

in process engineering settings where hard feasibility constraints must be satisfied, and probabilistic characterisation of uncertainty may not be available [6-8].

At the same time, embedding ANNs into process optimisation models can significantly increase nonlinearity and introduce nonconvexities due to the nature of their activation functions. This often yields large-scale nonconvex problems, motivating the use of global optimisation methods [1]. Spatial Branch-and-Bound (sBB) is a standard methodology for deterministic global optimisation of nonconvex problems that relies on partitioning the variable domain, solving convex relaxations to obtain lower bounds, and pruning regions that cannot contain a globally optimal solution, but its performance depends strongly on the tightness of the relaxations and variable bounds [9, 10]. For ANN-embedded formulations, relaxations based on conventional McCormick envelopes can be weak, particularly for convexo-concave activation functions such as sigmoid and hyperbolic tangent, leading to poor relaxations, excessive branching and poor algorithmic scalability [11-13]. This has motivated the development of more tight envelopes for activation functions to reduce the size of the search tree in ANN-embedded global optimisation. Recent research efforts in this direction include the software packages MeLOn[12] and the recent extensions to EAGO.jl [11].

Robust optimisation of nonconvex problems adds another layer of difficulty, since standard tractable robust reformulations rely on convexity and duality assumptions that often do not hold in surrogate-embedded optimisation formulations [6]. Cutting-set methods provide an iterative alternative by alternating between optimisation and worst-case "pessimisation" checks [14], but can become expensive when the underlying nonconvex problem is already challenging. In [15], Wolfe duality and robust chance constraints were proposed to tackle uncertainty in problems with warped Gaussian Processes. Recently, Marousi and Charitopoulos proposed a Robust spatial Branch-and-Bound (RsBB) algorithm that combines robust infeasibility checks at each node of the branch-and-bound tree with robust cuts inside the global search for solving nonconvex QCQPs [7].

Nonetheless, a key gap remains: a robust global optimisation framework for surrogate-based process design problems in which (i) an ANN surrogate is embedded within an algebraic optimisation model, (ii) uncertainty in the surrogate predictions is explicitly accounted for, and (iii) tight relaxations tailored to ANN activation functions are used to solve the resulting nonconvex problem. This work addresses that gap using a cumene process design case study [12, 13]. High-fidelity data is generated from Aspen Plus and used to train an ANN surrogate, which is then embedded in an optimisation model to couple the surrogate consistently with process constraints and decision variables. The resulting nonconvex optimisation problems are solved using an RsBB framework that incorporates tailored activation function envelopes [11] and Optimisation-Based Bound-Tightening (OBBT), and robustness is enforced under prescribed uncertainty sets on the ANN outputs. Preliminary computational results compare deterministic and robust variants and quantify the trade-off between conservativeness of the solution and objective performance.

METHODOLOGY

Robust spatial Branch-and-Bound Algorithm

Our approach builds on the RsBB algorithm of Marousi and Charitopoulos, which integrates a spatial branch-and-bound (sBB) global optimisation with a robust cutting-set method to enforce feasibility under uncertainty. A schematic overview of the resulting workflow is provided in Figure 1. In robust optimisation, feasibility is required for all uncertainty realisations in a prescribed set $\mathcal{U}$, resulting in a semi-infinite program of the form

$$\min_{x \in \mathcal{X}} f(x)$$

$$\text{s.t.}\, g(x,u) \leq 0 \; \forall u \in \mathcal{U}$$

The aim of RsBB is to avoid spending large effort optimising solutions that later turn out to be robust-infeasible, by performing robustness checks during the tree search [5]. At each node n of the tree, RsBB maintains two optimisation models defined on the node's variable bounds: (i) a nonconvex "sampled" problem that provides an upper bound and (ii) a convex relaxation that provides a lower bound used for fathoming. Robustness is enforced using a cutting-set method where, instead of requiring $g(x,u) \leq 0$ for all $u \in \mathcal{U}$ directly, the algorithm maintains a finite sampled set $\hat{\mathcal{U}} = \{u^0, \dots, u^K\} \subseteq U$ and solves the sampled problem with constraints enforced only for $u \in \hat{\mathcal{U}}$. The robustness of any candidate solution x^* is then evaluated by an infeasibility test that searches for the worst-case uncertainty realisation $u^* \in \mathcal{U}$ that maximises constraint violation, and if a violation is found, u^* is appended to $\hat{\mathcal{U}}$ and the node is re-solved. This is repeated until no further violations are detected, at which point the node solution is certified robust-feasible with respect to $\mathcal{U}$ [7].

The global-search component of the algorithm follows standard sBB logic [16]. Nodes are selected using a best-first strategy based on the lower bound provided by the relaxation, prioritising the most promising subdomains in the waiting set. At a selected node, the branching variable is selected as the pre-activation variable with the largest relaxation gap score, measured as the deviation between the post-activation variable and the activation function evaluated at the relaxed pre-activation variable value. If all candidate gaps fall below a small tolerance, the algorithm reverts to branching on the variable

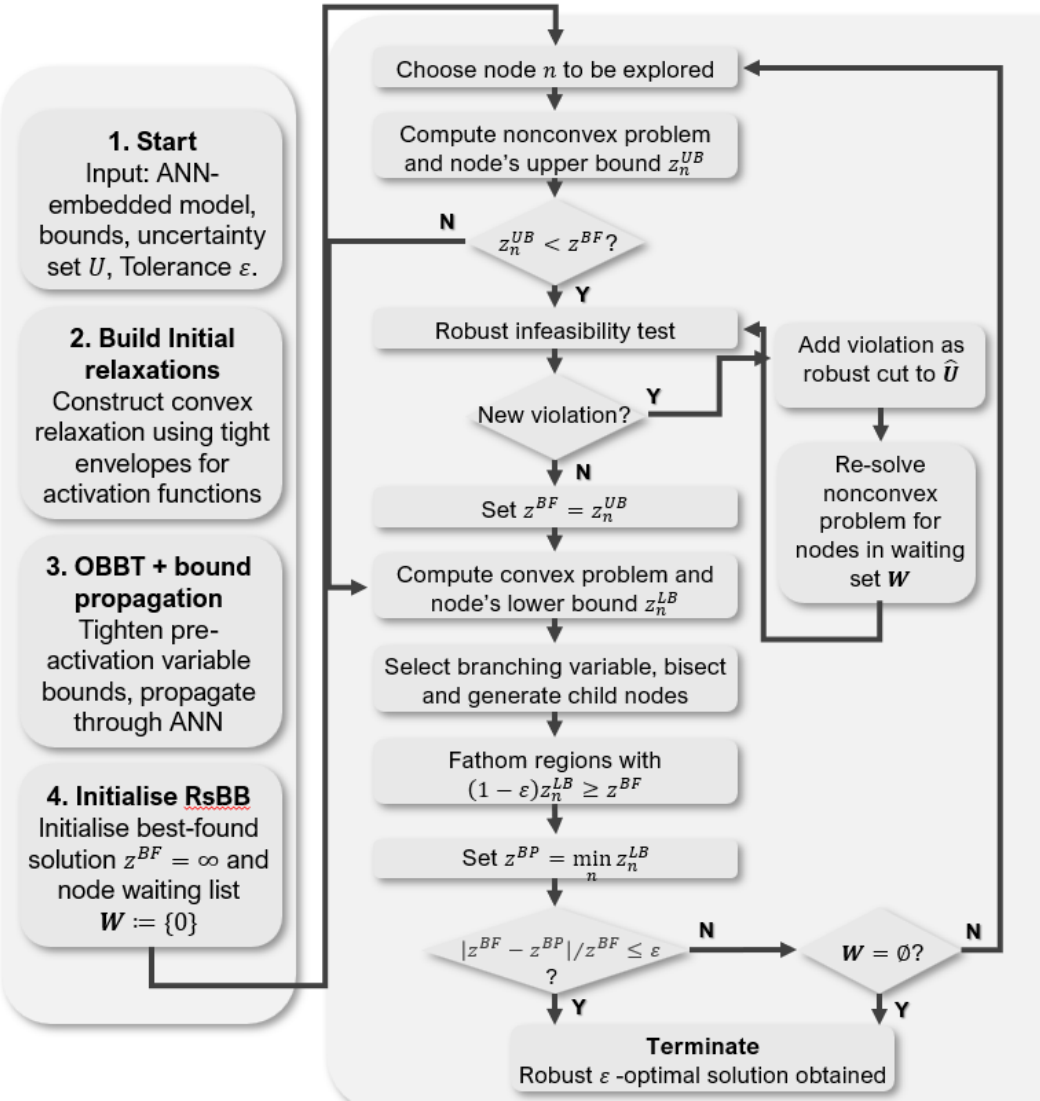

Figure 1. Algorithmic flowchart of RsBB with ANNs-embedded.

with the largest bound interval width. Once the branching variable is selected, bisection branching is applied. Nodes are fathomed whenever their relaxation bound cannot improve upon the incumbent objective within the predefined tolerance.

Relaxation Envelopes

To obtain the convex relaxations required for the RsBB algorithm, the ANN activation functions $z = f(x)$ are relaxed using tight envelopes for convexo-concave activation functions such as hyperbolic tangent and sigmoid. For a pre-activation variable $x \in [x^L, x^U]$ and post-activation variable z, we define the secant on $[a, b]$ and the tangent at $c \in [a, b]$ as:

$$sct_{[a,b]}(x) = f(a) + \frac{f(b) - f(a)}{b - a}(x - a)$$

$$tan_c(x) = f(c) + f'(c)(x - c)$$

The activation function is then relaxed via $z \geq f^{cv}(x)$ and $z \leq f^{cc}(x)$, where f^{cv} and f^{cc} are the convex under-estimator and concave over-estimator, respectively, given by the following expressions [9]:

$$f^{cv}(x) = \begin{cases} f(x), x^U \leq 0 \\ sct_{[x^L,x^U]}(x), x^L \geq 0 \\ \max\{tan_{x^L}(x), tan_{\xi^L}(x), sct_{[x^{cv},x^U]}(x)\}, x^L < 0 < x^U \end{cases}$$

$$f^{cc}(x) = \begin{cases} sct_{[x^L,x^U]}(x), x^U \leq 0 \\ f(x), x^L \geq 0 \\ \min\{tan_{x^U}(x), tan_{\xi^U}(x), sct_{[x^L,x^{cc}]}(x)\}, x^L < 0 < x^U \end{cases}$$

The tie points x^{cv} and x^{cc} are defined as the points where the tangent to $f(x)$ matches the secant connecting the tie point to x^U and x^L, respectively [11]:

$$f'(x^{cv}) = \frac{f(x^U) - f(x^{cv})}{x^U - x^{cv}};\ f'(x^{cc}) = \frac{f(x^{cc}) - f(x^L)}{x^{cc} - x^L}$$

Additionally, the points ξ^U and ξ^L are defined as:

$$\xi^U = \frac{x^U + x^{cc}}{2};\ \ \xi^L = \frac{x^L + x^{cv}}{2}$$

Compared to McCormick relaxations, the envelopes presented here are significantly tighter for convexo-concave functions, avoiding the dependency problem whereby repeated occurrences of the same variable are relaxed independently, leading to overestimation and loose bounds [9].

Optimisation-Based Bound Tightening

Since the tightness of the envelopes depends strongly on the bounds $[x^L, x^U]$, we apply Optimisation-Based Bound Tightening (OBBT) for domain reduction and to improve relaxations. In our implementation, OBBT is performed as a preprocessing step by solving a pair of auxiliary optimisation problems for each pre-activation and input variable: one problem minimises the variable and one maximises it, subject to all constraints of the current relaxed model. The resulting optimal values provide tighter feasible bounds, which are accepted only if they improve the current bounds and do not invert the interval.

Because this procedure is computationally expensive, we run it iteratively with a pre-specified budget and reduction improvement tolerance. After each OBBT iteration, we perform bound propagation through the ANN to improve the tightening without additional optimisation problems: bounds are propagated through affine layer relations using interval arithmetic and through the activation functions. Then, the updated bounds are passed on to the relaxation envelopes so that they are updated consistently with the new bounds.

CASE STUDY

Process and Decision Problem Description

We test the proposed algorithm on the optimisation of the cumene process, which synthesises cumene via Friedl-Crafts alkylation of benzene with propene over a solid catalyst [17]. The process exhibits a classic trade-off between reactor design and separation efficiency, where lower reactor temperatures improve selectivity to cumene but require a larger vessel and higher recycle rates, while higher temperatures permit smaller reactors but produce more side products. The flowsheet comprises of three main sections: preheating of fresh and recycled feeds, a cooled packed-bed reactor, and a distillation train that separates light inerts, recovers unreacted benzene for recycle and delivers the final cumene product [17, 18]. This work develops a neural network surrogate model to predict cumene product purity and flow rate, which are the two primary output variables controlling revenue, as functions of seven input variables. The input variables comprise of the propylene and benzene feed flow rates (F_{C3} and F_{BZ}), the reactor inlet temperature (T_R), the reboiler duty of the two distillation columns (Q_{C1} and Q_{C2}) along with their associated reflux ratios (RR_{C1} and RR_{C2}) [17]. The process economics are driven by maximising profit, which is simplified as the revenue generated from cumene sales minus the cost of feeds. The surrogate ANN model enables the evaluation of alternative operating strategies to balance the competing objectives such as purity vs feed efficiency vs throughput, along with constraints imposed on the minimum product purity and flow rate. The combination of nonlinearity and multiple active constraints makes the cumene process an appropriate case study for surrogate-based optimisation.

Optimisation Problem Formulation

We seek the design and operating decisions that maximise profit while guaranteeing feasibility for all realisations of bounded uncertainty on the ANN outputs (product purity and flow rate). The robust problem is formulated as:

$$\min_{\boldsymbol{x}\in\mathcal{X},\boldsymbol{y}\in\mathcal{Y}} -\Pi(\boldsymbol{x},\boldsymbol{y})$$

$$s.t.\ \boldsymbol{y} = f^{ANN}(\boldsymbol{x},\boldsymbol{u})$$

$$\boldsymbol{\mu}(\boldsymbol{x},\boldsymbol{y},\boldsymbol{w}) \le 0$$

$$y_P^{(u)} \ge 0.99\ \ \forall u_P \in \mathcal{U}_P = [-\Delta P, \Delta P]$$

$$y_F^{(u)} \ge 70\ \ \forall u_F \in \mathcal{U}_F = [-\Delta F, \Delta F]$$

$$\boldsymbol{x} \in \mathcal{X}, \boldsymbol{y} \in \mathcal{Y}, \boldsymbol{u} \in \mathcal{U}$$

where $\Pi(\boldsymbol{x},\boldsymbol{y})$ is the profit (thousand £/year), $\boldsymbol{y} = [y_P, y_F]$ are the ANN outputs, $\boldsymbol{x} \in \mathcal{X}$ are the process decision variables, $f^{ANN}(\cdot)$ represents the trained ANN, $\boldsymbol{\mu}(\boldsymbol{x},\boldsymbol{y},\boldsymbol{w}) \le 0$ denotes the activation function envelope constraints (with auxiliary variables $\boldsymbol{w}$), $y_P^{(u)}$ and $y_F^{(u)}$ denote the uncertain purity and flow rate outputs associated with uncertainty realisation $\boldsymbol{u} \in \mathcal{U}$, and $\mathcal{U} = \mathcal{U}_P \times \mathcal{U}_F$ is the box uncertainty on $(\Delta P, \Delta F)$.

Data Generation and ANN Training Workflow

Data were generated using Aspen Plus V14, using the flowsheet proposed by Luyben [17] and Schultz [18]. Python 3.10.19 was then used to automate the Aspen Plus simulations by utilising Aspen's COM interface to programmatically set inputs, run simulations and extract results, which were stored in a CSV file. Within the Python code used for sampling, the Aspen Plus stream variable bounds were set along with a fixed number of input samples using Sobol sampling. A varying number of samples were used for ANN validation, ranging from 1, 000 to 10,

000 samples. The sampled data were then cleaned to ensure only physically viable results were used in the ANN training, e.g. non-convergent runs and infeasible solutions (e.g., negative flow rates, unbounded temperatures) were excluded, retaining 99.95% of generated samples.

ANN training was completed using PyTorch. After training the ANN in PyTorch, a set of weights and biases for each neuron were generated. Within PyTorch, the data were split into training and testing sets using an 80/20 training-testing split and min-max scaling was selected. The ANN architecture consisted of 7 input nodes (representing feed composition and operating parameters), five hidden layers (16→32→32→32→16 neurons) with tanh activation functions, and 2 output nodes (representing product purity and flow rate) with linear activation. This architecture was selected to balance the nonlinear complexity of the Aspen flowsheet while minimising the computational burden for the RsBB optimisation algorithm. Batch normalisation was applied after each linear layer (16, 32, 32, 32, 16 nodes) to stabilise training and accelerate convergence. Dropout with a rate of 30% was applied after the activation functions to reduce overfitting. This ordering ensures that batch normalisation statistics are not corrupted by dropout-induced sparsity

The trained ANN was then exported to ONNX for it to be in a format compatible with the Optimization and Machine Learning Toolkit (OMLT). OMLT v1.2.2 was used to embed the trained ANN into Pyomo [19]. The ANN constraints were built using FullSpaceSmoothNNFormulation for tanh activations and the final ANN block was then stored in order to be inserted into the RsBB algorithm.

Validation and Errors

The trained ANN model was evaluated on 10, 000 process samples to predict the output product purity and flow rate. The model demonstrated strong predictive performance with an overall R^2 of 0.772, with an R^2 of 0.739 for purity and 0.804 for flow rate. Error analysis reveals that purity predictions were exceptionally consistent with a mean signed error of 0.000378 and a standard deviation ±0.0186, indicating a minimal systematic bias and tight clustering around the true values. The model exhibits a slight bias towards underestimation, as 48.31% of purity predictions fall within the -0.05 to 0 unit range, while overestimations only account for 1.29% of predictions. Only 0.31% of purity predictions are classified as outliers (>3 standard deviations) and these outliers are all negative, suggesting the model rarely produces overestimations.

Flow rate predictions demonstrate broader variability, with a mean signed error of -0.000593 and a standard deviation of ±4.575. The flow rate error distribution spans from -56.63 to +39.72, with approximately 52% of predictions concentrated within ±3 unit range. An outlier rate of 0.89% was observed and the extremes were distributed relatively symmetrically between overestimations and underestimations, indicating no directional bias.

RESULTS

All computations were performed on a workstation with an Intel Core i9-13900K (13th Gen) CPU (24 cores, 32 logical processors, 3.00 GHz base clock) and 128 GB RAM.

Deterministic Optimisation

The deterministic optimisation problem (i.e., without uncertainty on ANN outputs) was solved globally using the proposed RsBB framework with the tailored relaxations for activation functions and a preprocessing OBBT step. At the root node, the baseline relaxation gap was 3.47%, with a best-found objective of -16, 299.70 and a lower bound of -16, 865.76. Applying OBBT reduces the gap to 2.02% after seven OBBT iterations.

The globally optimal nominal solution returned by RsBB corresponds to a profit of 16, 299.70 thousand £/year. The optimal decision vector is: $F_{C3} = 105.00\ kmol\ h^{-1}, F_{BZ} = 95.00\ kmol\ h^{-1}, T_R = 360.00\ ^\circ\mathrm{C}, Q_{C1} = 1.3100\ Gcal\ h^{-1}, Q_{C2} = 1.5758\ Gcal\ h^{-1}, RR_{C1} = 0.9999,$ and $RR_{C2} = 1.1998$. At this optimum, the ANN predicts a product purity of 0.9900 mol/mol and a flow rate of 86.90 kmol/h. The total solution time was 299.75 s, of which 204.66 s were spent in OBBT and 95.09 s in the RsBB tree search, exploring a total of 202 nodes.

For benchmarking, the same deterministic model was solved with BARON (v24.10.10 using GAMS 48.1.0) under a 5, 400 s time limit. BARON recovered an identical best-found solution with the same decision variables; however, terminating with a reported gap of 8.356%. A summary comparison of RsBB and BARON [20] on the deterministic problem is provided in Table 1.

Robust Optimisation Under Box Uncertainty

Robust optimisation was performed by enforcing feasibility for all uncertainty realisations within a box uncertainty set defined on the ANN outputs: product purity uncertainty with radius ΔP (mol/mol) and product flow rate uncertainty with radius ΔF (kmol/h). Results are reported for several uncertainty levels, from the deterministic case up to $(\Delta P, \Delta F) = (0.004, 10)$. A detailed solver comparison with RsBB vs BARON for the nominal case, plus a representative robust run at $(\Delta P, \Delta F) = (0.002, 5)$, is given in Table 1, while Figure 2 summarises the profit-conservativeness trade-off across all uncertainty levels tested.

For the representative robust case $(\Delta P, \Delta F) = (0.002, 5)$, RsBB returns a globally robust solution with a profit of 15, 850 10^3£/year, a 2.75% reduction with respect to the deterministic case. Robust feasibility is

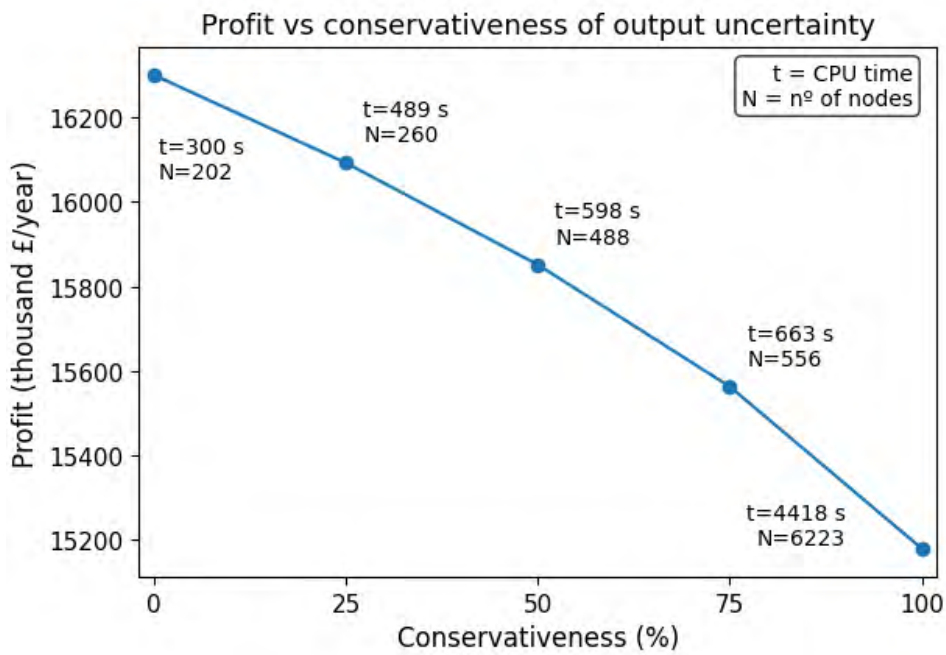

Figure 2. Profit-conservativeness trade-off under box uncertainty on ANN outputs. Conservativeness is scaled relative to the maximum uncertainty and annotations report CPU time and number of nodes explored.

enforced through the cutting-set mechanism embedded within the sBB search. In this representative case, a single robust cut was generated (at node 5) to enforce the purity requirement at the worst-case realisation, after which the solution was certified robust-feasible. Relative to the nominal optimum, the robust solution changes through the reflux ratio of the first column and the reboiler duty of the second column: RR_{C1} decreases from 0.9999 to 0.6107 and Q_{C2} decreases from 1.5758 to 1.5373, while other inputs remain unchanged. The robust optimum corresponds to ANN outputs of product purity 0.9920 mol/mol and flow rate 86.00 kmol/h, reflecting the need to maintain feasibility under the prescribed box uncertainty.

Increasing the conservativeness of the uncertainty set imposes a clear penalty in the objective value. As shown in Figure 2, profit decreases monotonically as the uncertainty scale increases from 0% (nominal) to 100% $\big((\Delta P, \Delta F) = (0.004, 10)\big)$, quantifying the robustness-performance trade-off as shown in Figure 2. These results highlight both (i) the practical cost of robustness in terms of objective value and solve time, and (ii) the ability of the proposed RsBB framework with tailored envelopes and preprocessing OBBT to provide globally optimal and robust solutions for ANN-embedded process optimisation problems.

CONCLUSIONS

This work presented a robust global optimisation framework for ANN-embedded process optimisation under uncertainty by integrating RsBB with tight envelope relaxations for convexo-concave activation functions and a preprocessing OBBT step. Using a cumene process case study, an ANN surrogate trained on Aspen Plus simulation data was embedded into an algebraic optimisation model, and robustness was enforced via a cutting-set strategy under box uncertainty on ANN outputs (product purity and flow rate).

For the nominal (deterministic) case, RsBB obtained a globally optimal solution after exploring 202 nodes in 299.75 s total CPU time. In contrast, BARON recovered an identical best found solution but did not close the optimality gap within the 5, 400 s time limit, and therefore did not provide a global optimality certificate for this ANN-embedded problem. These results highlight the value of tailored activation function relaxations, which improve envelope tightness and allow for effective pruning in the RsBB algorithm for ANN-embedded global optimisation.

Robust optimisation introduced the expected performance-conservativeness trade-off. For the representative uncertainty level $(\Delta P, \Delta F) = (0.002, 5)$, the robust

objective worsened by 2.75% relative to the nominal optimum, reflecting the cost of enforcing feasibility across the uncertainty set. As conservativeness increased, the objective worsened monotonically, reaching a 6.87% degradation at the maximum uncertainty considered $(0.004, 10)$. This higher conservativeness also increased computational effort considerably, as stricter robust feasibility requirements and associated robust cuts increased the search required to meet the same optimality tolerance. Overall, the proposed RsBB approach provides a practical framework to obtaining certified globally optimal robust solutions for surrogate-based process optimisation under bounded uncertainty.

ACKNOWLEDGEMENTS

Financial support from EPSRC grants EP/T022930/1, EP/W003317/1 and EP/V051008/1 is gratefully acknowledged.

AUTHOR IDENTIFIERS

Author ORCIDs:
Charitopoulos VM: 0000-0001-9051-917X
Yuhui Y: 0009-0000-5360-720X

REFERENCES

1. Huster WR, Schweidtmann AM, Lüthje JT, Mitsos A. Deterministic global superstructure-based optimization of an organic rankine cycle. Computers & Chemical Engineering 141:106996 (2020). https://doi.org/10.1016/j.compchemeng.2020.106996
2. Hüllen G, Zhai J, Kim SH, Sinha A, Realff MJ, Boukouvala F. Managing uncertainty in data-driven simulation-based optimization. Computers & Chemical Engineering 136:106519 (2020). https://doi.org/10.1016/j.compchemeng.2019.106519
3. McBride K, Sundmacher K. Overview of surrogate modeling in chemical process engineering. Chemie Ingenieur Technik 91:228-239 (2019). https://doi.org/10.1002/cite.201800091
4. Luo Y, Ierapetritou M. Multifeedstock and multiproduct process design using neural network surrogate flexibility constraints. Ind. Eng. Chem. Res. 62:2067-2079 (2023). https://doi.org/10.1021/acs.iecr.2c02968
5. Beykal B, Boukouvala F, Floudas CA, Sorek N, Zalavadia H, Gildin E. Global optimization of grey-box computational systems using surrogate functions and application to highly constrained oil-field operations. Computers & Chemical Engineering 114:99-110 (2018). https://doi.org/10.1016/j.compchemeng.2018.01.005
6. Ben-Tal A, Nemirovski A. Robust solutions of uncertain linear programs. Operations Research Letters 25:1-13 (1999). https://doi.org/10.1016/s0167-6377(99)00016-4
7. Marousi, A., & Charitopoulos, V. M. (2025). Global and Robust Optimisation for Non-Convex Quadratic Programs. arXiv preprint arXiv:2503.07310. doi: https://doi.org/10.48550/arxiv.2503.07310
8. Gorissen BL, Yan?ko?lu ?, den Hertog D. A practical guide to robust optimization. Omega 53:124-137 (2015). https://doi.org/10.1016/j.omega.2014.12.006
9. Ryoo HS, Sahinidis NV. Global optimization of nonconvex nlps and minlps with applications in process design. Computers & Chemical Engineering 19:551-566 (1995). https://doi.org/10.1016/0098-1354(94)00097-2
10. Liberti L, Maculan N. Global optimization. Springer US (2006). https://doi.org/10.1007/0-387-30528-9
11. Wilhelm ME, Wang C, Stuber MD. Convex and concave envelopes of artificial neural network activation functions for deterministic global optimization. J Glob Optim 85:569-594 (2022). https://doi.org/10.1007/s10898-022-01228-x
12. Schweidtmann AM, Mitsos A. Deterministic global optimization with artificial neural networks embedded. J Optim Theory Appl 180:925-948 (2018). https://doi.org/10.1007/s10957-018-1396-0
13. Anderson BG, Ma Z, Li J, Sojoudi S. Tightened convex relaxations for neural network robustness certification. 2020 59th IEEE Conference on Decision and Control (CDC) :2190-2197 (2020). https://doi.org/10.1109/cdc42340.2020.9303750
14. Mutapcic A, Boyd S. Cutting-set methods for robust convex optimization with pessimizing oracles. Optimization Methods and Software 24:381-406 (2009). https://doi.org/10.1080/10556780802712889
15. Wiebe J, Cecílio I, Dunlop J, Misener R. A robust approach to warped gaussian process-constrained optimization. Math. Program. 196:805-839 (2022). https://doi.org/10.1007/s10107-021-01762-8
16. Smith EMB, Pantelides CC. Global optimisation of nonconvex minlps. Computers & Chemical Engineering 21:S791-S796 (1997). https://doi.org/10.1016/s0098-1354(97)87599-0
17. Luyben WL. Design and control of the cumene process. Ind. Eng. Chem. Res. 49:719-734 (2009). https://doi.org/10.1021/ie9011535
18. Schultz ES, Trierweiler JO, Farenzena M. The importance of nominal operating point selection in self-optimizing control. Ind. Eng. Chem. Res.

55:7381-7393 (2016). https://doi.org/10.1021/acs.iecr.5b02044
19. Ammari B, Ceccon F, Jalving J, Haddad J, Thebelt A, Tsay C, Bynum M, Laird C, Misener R. Machine learning surrogates with OMLT and IDAES for improved design and analysis of energy systems. 2022 INFORMS Annual Meeting - Indianapolis, Indiana, United States of America - October - 2022 : (2022). https://doi.org/10.2172/2432248
20. Zhang Y, Sahinidis NV. Solving continuous and discrete nonlinear programs with BARON. Comput Optim Appl 92:1123-1161 (2024). https://doi.org/10.1007/s10589-024-00633-0

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Using Active Learning to Efficiently Calibrate Foundation Models on Raman Spectra in Upstream Bioprocess Fermentations

Christoph Lange[a*], Ernesto Martínez[ab], Peter Neubauer[a] and Mariano Nicolas Cruz Bournazou[a]

[a] Technische Universität Berlin, Faculty III Process Sciences, Institute of Biotechnology, Chair of Bioprocess Engineering, Straße des 17. Juni 135, 10623 Berlin, Germany
[b] INGAR, (CONICET - UTN), Avellaneda 3657, Santa Fe, Argentina
* Corresponding Author: christoph.lange@tu-berlin.de

ABSTRACT

Real-time monitoring of metabolite concentrations is critical for optimising bioprocess performance. While Raman spectroscopy offers a non-invasive solution, translating spectra into metabolite concentration estimates requires robust machine learning models. Foundation models such as TabPFN demonstrate exceptional predictive performance but suffer from high inference complexity when trained on large calibration datasets, hindering their use in real-time laboratory settings. This study proposes a batch Active Learning (AL) strategy to efficiently calibrate TabPFN using a minimal subset of data. We employ a weighted K-means clustering strategy that balances model uncertainty and dataset diversity to select the most informative calibration samples. We evaluated this method on a dataset of nearly 7, 000 Raman spectra covering eight substances. Our AL strategy achieved a mean R^2 score greater than 0.95 with approximately 1, 000 samples, significantly outperforming random sampling. Notably, the method matched the accuracy of a model trained on the full dataset using only 20% of the data. This reduction lowers computational complexity by a factor of 25, enabling millisecond-scale inference times suitable for high-throughput bioprocess monitoring.

Keywords: Algorithms, Artificial Intelligence, Fermentation, Machine Learning, Modelling

INTRODUCTION

In bioprocess engineering, the ability to monitor metabolite concentrations in real time during microbial or cell cultivations is critical for understanding, controlling, and optimising biotechnological processes with respect to yield and product quality. Traditionally, metabolite analysis is based on offline sampling and subsequent analytical techniques such as High-performance liquid chromatography or Gas chromatography–mass spectrometry, which are time-consuming, labour-intensive, and often unsuitable for high-throughput experimentation.

As an alternative, Raman spectroscopy has emerged as a promising Process Analytical Technology (PAT) for non-invasive monitoring of metabolite concentrations as it is fast, label-free, and capable of simultaneously quantifying multiple compounds in complex mixtures. Recent advancements enable the integration of Raman spectrometers into high-throughput systems, significantly enhancing the volume and quality of data generated during parallel cultivations [1], [2].

However, a major challenge remains the translation of raw Raman spectra into metabolite concentration estimates. This requires a machine learning model that maps Raman spectra to estimations of concentration values. Traditionally, this has been achieved using multivariate statistical methods such as Partial Least Squares (PLS) [2] or Deep Neural Networks [4].

However, training deep learning models typically requires large, diverse calibration datasets, which are costly and time-consuming to generate. A recent study

showed that a tabular foundation model TabPFN [5] delivers exceptional prediction performance when provided with a large calibration dataset [6]. This model is different from classic regression models as it uses in-context learning. When predicting concentrations for an unknown Raman spectrum, the model computes similarities between all calibration data points and the unknown spectrum using a Transformer architecture [7]. Unfortunately, the computational complexity of this step increases quadratically with the number of calibration points. This results in slow prediction speeds, particularly in laboratory contexts where modern GPUs are often unavailable. On a CPU, the average inference time per spectrum is approximately 10 seconds [6].

To address this limitation we reduce the size of the calibration dataset using Active Learning (AL) algorithms which select the most informative data points from the calibration set [8]. While many AL algorithms suggest one sample at a time this procedure is not efficient when each step requires retraining a whole neural network. A naive approach is to select the most informative data points based solely on a specific metric. However, this approach might lead to a batch of similar data points such that the model would be better off with labelling more diverse data points that yield better a generalisation error. This diversity-informativeness trade-off is addressed by methods that optimise the informativeness of the whole batch of data points [9], [10], [11].

MATERIAL AND METHODS

Dataset

In this study we use a dataset of nearly 7000 Raman spectra with annotations for 8 substances [6] that are relevant in high-throughput bioprocess development [12]. As the different substances have been added to the dataset over time the concentration labels of some substances are imbalanced which makes the dataset an interesting subject to study for AL.

Model

As a model we use version 2.5 of TabPFN [5], [13], [14] which is a foundation model for tabular datasets that uses in-context learning on the whole training set when predicting new labels. This model has been used in version 2 on this dataset before and showed exceptional performance in particular for low concentrations [6].
For regression problem like ours the model discretises the concentration space and actually solves a classification problem under the hood. Therefore, it yields probabilistic predictions of the concentration labels [15] that widen the range of possible information criteria that can be used during AL.

Setup

To select Raman spectrum–concentration pairs that most improve the predictive performance of the model we perceive the problem as pool-based active learning [8]. We join the training dataset with the validation dataset to a pool of more than 6000 spectra to choose from. We use the test dataset for evaluating the model in each step.

The reason why we can merge the training and the validation set in our evaluation setup is that we chose to not do any hyperparameter tuning. This decision is motivated by two considerations. First an foremost, TabPFN is a foundation model that uses in-context learning [14]. Therefore we did not do classic model training on the spectral datasets where we can pick hyperparameters that shape TabPFN's architecture like the encoding dimension, the number of transformer encoder layers, the dropout rate etc. Second, for the hyperparameters that we could have tuned like the number of estimators or the kernel size for the average pooling step, we use the same values as in [6].

Typically, active learning algorithms suggest labeling one sample at a time. However, that procedure is very time-consuming for a datasets size above 6000 when each suggestion requires retraining the whole model. Therefore, we choose an AL procedure that suggests batches of data points.

Algorithm

The batch suggested by the AL algorithm should be optimised in terms of two objectives. On the one hand, we need to quantify the informativeness of unlabelled data points. On the other hand we need a mechanism that ensures diversity within the batch.

Informativeness

As most machine learning algorithms do not deliver probabilistic predictions the classic active learning approach is to work with an ensemble of regressors and judge the informativeness of an unlabelled data point by the discrepancy of their predictions. Since we use TabPFN, we obtain a predictive distribution over discretised bins of the concentration space [13]. Consequently, we can use information criteria that are directly related to the posterior distribution$P(y|x)$and which are usually used for classification models. As the discretisation consists of 5000 bins the typical least confident approach [8] or margin sampling [16] neglect most of the distribution. Therefore, we use the entropy [17] of the predictive distribution to measure the uncertainty of the model

$$U(x_i) = \sum_{j=1}^{5000} -P(y_j|x_i) log P(y_j|x_i) \quad (1)$$

in its own prediction. We use this uncertainty as a metric for the informativeness of each spectrum.

Algorithm 1 Batch Active Learning with Weighted K-Means

Require: Pool of training spectra X_{train},
Test set X_{test},
batch sizes $\{n_0, n_1, \ldots, n_T\}$,
number of independent runs $N = 10$,
Ensure: R^2 scores per step for Active Learning

Preprocess all spectra: average pooling + water-peak intensity filtering
for run $= 1$ **to** N **do**
  $X_u \leftarrow X_{\text{train}}$
  $\mathcal{L} \leftarrow \text{RandomSample}(X_u,\ n_0)$
  for $t = 1$ **to** T **do**
    Fit TabPFN on $\mathcal{L}$; Evaluate TabPFN on X_{test}
    for $s \in \mathcal{S}$ **do**
      Fit TabPFN on $\{(x_i, y_i^s) : x_i \in \mathcal{L},\ y_i^s \neq \text{NaN}\}$
      Compute uncertainty for all $x_i \in X_u \setminus \mathcal{L}$:

$$U^{(s)}(x_i) = \sum_{j=1}^{5000} -P(y_j^s \mid x_i) \log P(y_j^s \mid x_i)$$

    end for
    $U(x_i) \leftarrow \frac{1}{8} \sum_s U^{(s)}(x_i)$
    Select candidate pool: $\mathcal{C} \leftarrow \text{top}_{5n_t}(\{U(x_i)\}_{i \in X_u \setminus \mathcal{L}})$
    Solve weighted K-means on $\mathcal{C}$ with $K = n_t$ clusters:

$$\min_{\mu_k,\, z_{ik}} \sum_{x_i \in \mathcal{C}} \sum_{k=1}^{K} U(x_i)\, z_{ik}\, \|x_i - \mu_k\|$$

    $\mathcal{S} \leftarrow$ spectra in $\mathcal{C}$ closest to each cluster centre μ_k
    Resolve duplicates using the highest-U unselected candidate
    $\mathcal{L} \leftarrow \mathcal{L} \cup \mathcal{S}$
  end for
end for

Algorithm 1. Uncertainty-guided batch active learning for spectral calibration. Candidate spectra are scored by averaged predictive entropy across substances and clustered via weighted K-means to balance informativeness with diversity in each acquired batch.
Active Learning Method

Diversity

We use a method that is suggested in [11] and incorporate minor changes that are inspired from [9]. So let X_u be the set of unlabelled spectra and let $S \subseteq X_u$ be the set of Kspectra that we want to label. This subset should be most representative of the set of unlabelled spectra which we quantify via:

$$\sum_{x \in X_u} min_{s \in S} ||x - s|| \tag{2}$$

However, this problem is known to be NP-hard which is why we solve a similar problem that is known as K-means clustering

$$\sum_{x_i \in X_u} \sum_{k=1}^{K} z_{ik} ||x - \mu_k|| \tag{3}$$

where we would like to find both the cluster centres $\mu_k$and the binary assignment$z_{ik}$for each spectrum$x_i$to the clusters$k$. Even though the two optimisation problems look very similar the crucial difference is that for K-means the cluster centres$\mu_k$do not need to be part of the set of unlabelled spectraX_u. So instead of the cluster centres we choose the spectra that are closest to the cluster centres of K-means. In case there is one spectrum that is closest to two or more cluster centres we pick the spectrum with the highest uncertainty that has not been picked before.

Moreover, we do not cluster the whole set of unlabelled spectra, but choose thespectra that are most informative according to their entropy. This is inspired from [9] and increases the fraction of highly uncertain spectra

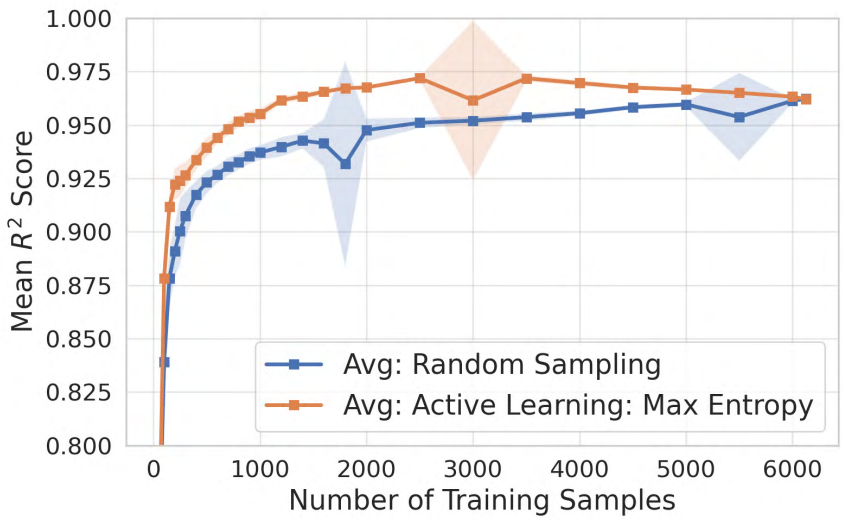

Figure 1. Overall model performance comparison, averaging R^2 scores across all eight substances. This plot summaries the aggregate performance of the active learning and random sampling strategies. The x-axis is the number of training samples, and the y-axis is the mean R^2 score. Lines and shaded areas represent the mean and standard deviation over 10 runs, respectively

within the batch.

Joint Objective

To consider both aspects of informativeness and diversity we alter the K-mean clustering problem (Equation 3) by introducing the informativeness criterion as a weight for each spectrum. So we solve the weighted K-means clustering problem

$$\sum_{x_i \in X_u} \sum_{k=1}^{K} U(x_i) z_{ik} ||x - \mu_k|| \quad (4)$$

that gives more importance to unlabelled spectra based on the model's uncertainty $U(x_i)$.
To evaluate the effectiveness of that approach we compare it against Random Sampling (RS). As the first batch of data is randomly selected for both methods we conduct 10 independent experiments while we use the same first batch for both AL as well as RS. More details on the exact procedure are given in Algorithm 1.

RESULTS

We compare the AL method described previously against RS.

Overall Predictive Performance

The predictive performance of both methods, RS and AL, follows a characteristic learning curve, with R^2 scores improving rapidly before reaching a plateau. When averaged across eight substances and ten experimental runs, the AL strategy consistently outperformed RS (Figure 1). While RS shows a slight advantage on five substances in low-data scenarios, AL makes significant gains on the remaining substances (Figure 2). This results in a higher mean R^2 score throughout training sizes.

Substance-Specific Sensitivity

A breakdown of performance by substance reveals that the effectiveness of AL is highly dependent on the specific substance (Figure 2) and can be grouped into three distinct categories:

AL-Dominant Profiles (EnPump and Yeast Extract): AL provides a substantial advantage here. For EnPump, RS fails to exceed an R^2 of 0.85 even with 6, 000 samples, whereas AL reaches near-perfect prediction almost immediately. This suggests these substances are defined by features sparsely distributed in the parameter space, which benefit significantly from targeted sampling.

Rapid **Convergence Profiles (Glucose, Glycerol, Acetate, Nitrate, and Total Phosphate):** Both strategies perform similarly, reaching R^2 values above 0.90 within the first 500 samples. The underlying relationships for these substances appear easy to capture regardless of the sampling method.

RS-Dominant **Profile (Total Sulfate):** RS outperforms AL until the training set reaches 3, 000 samples. The high variability in the AL results suggests that model uncertainty for Sulfate did not effectively guide the selection of data points, leading to a less balanced training set.

Convergence and Significance

As the training set size approaches 6, 000 samples, the gap between AL and RS narrows for most substances. However, the early-stage superiority of the AL

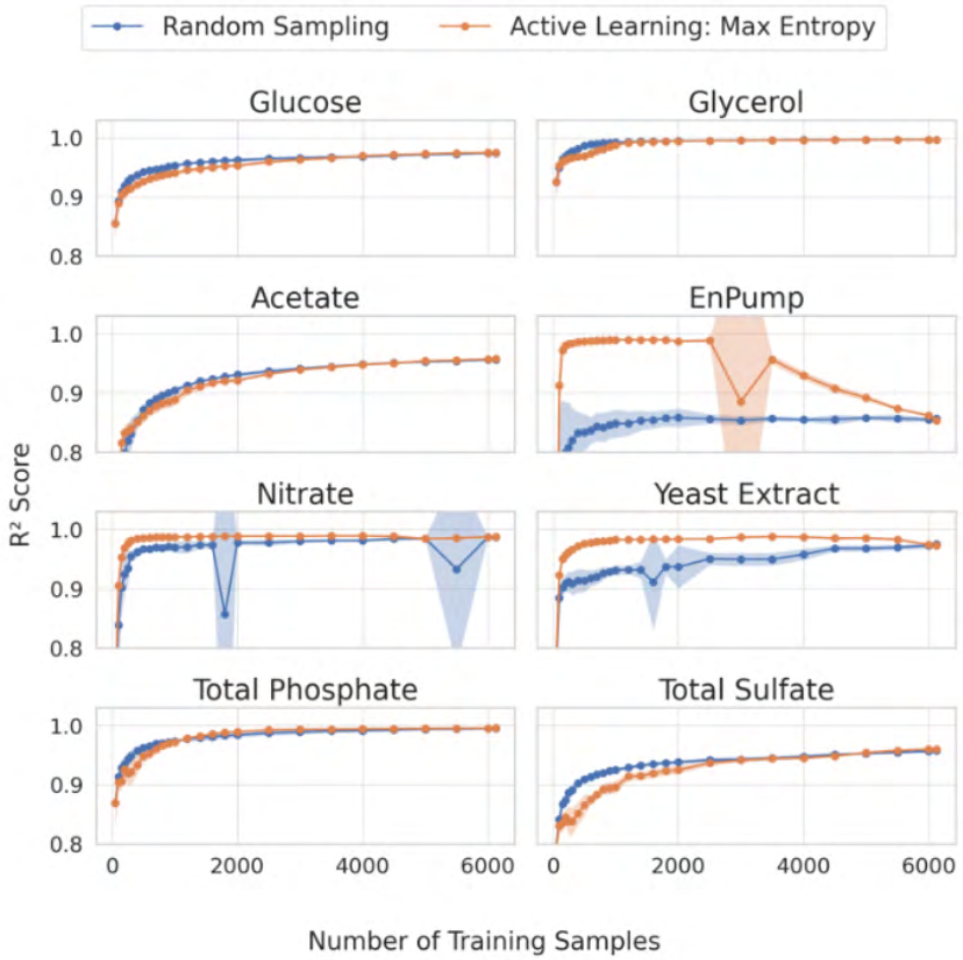

Figure 2: Comparison of model performance (R^2 score) for predicting individual substance concentrations using active learning versus random sampling. Each subplot corresponds to a specific substance. The x-axis represents the number of samples used for training the model, while the y-axis shows the resulting R^2 score on a held-out test set. The solid lines indicate the mean performance across 10 independent experimental runs, and the shaded areas represent the standard deviation.

Data Efficiency and Peak Performance

approach remains significant. A paired t-test to compare the two strategies via p-values (Figure 3) confirms that the advantage of AL is statistically significant across the majority of training set sizes.

The AL algorithm demonstrated superior data efficiency, reaching high accuracy levels much faster than RS. Specifically, AL achieved a mean R^2 score of 0.95 with approximately 1, 000 samples, whereas RS required nearly 2, 500 samples to reach the same threshold (Figure 1).

Notably, the AL model reached the same accuracy as the final model (trained on the full dataset) with only 1, 200 samples. Performance continued to climb until it peaked at 2, 500 samples, where the AL-selected subset actually achieved a higher accuracy than the model trained on the complete dataset. Paired t-tests (Figure 4) indicate that this peak performance at 2, 500 samples is statistically significant. This suggests that by the 2, 500-sample mark, the AL algorithm had successfully identified the most informative data points, and the slight improvements for some substances cannot make up for the deterioration of EnPump and Yeast Extract.

DISCUSSION

First, the use of an AL algorithm based on entropy as a selection criterion to discriminate between different Raman spectra is quite new as using accurate probabilistic regression models on Raman spectra is a very recent phenomenon. However, TabPFN delivers very high accuracies in a probabilistic form that enables the AL algorithm to select training samples that improve model performance compared to RS.

Therefore, the results discussed here highly depend

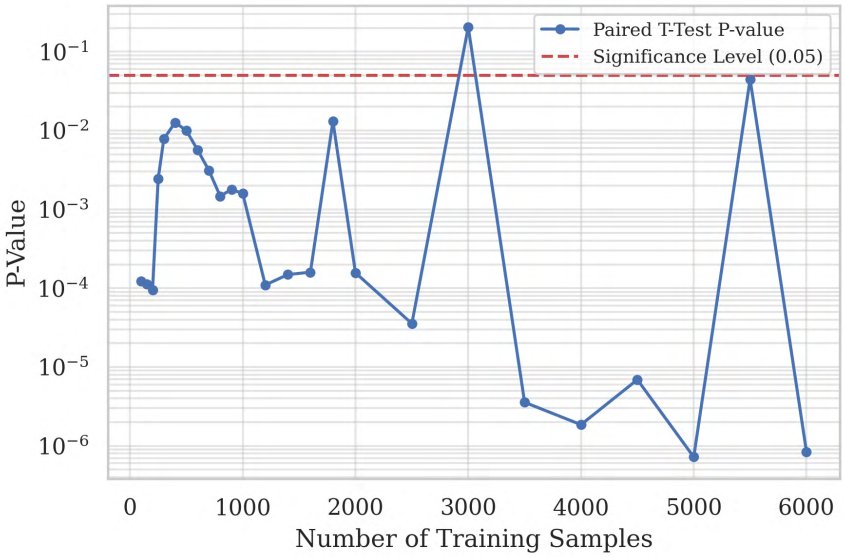

Figure 3. Statistical significance of the performance improvement of Active Learning over Random Sampling. The plot shows the p-value from a paired t-test performed on the R^2 scores across all predicted substances at different training set sizes and across 10 trials.

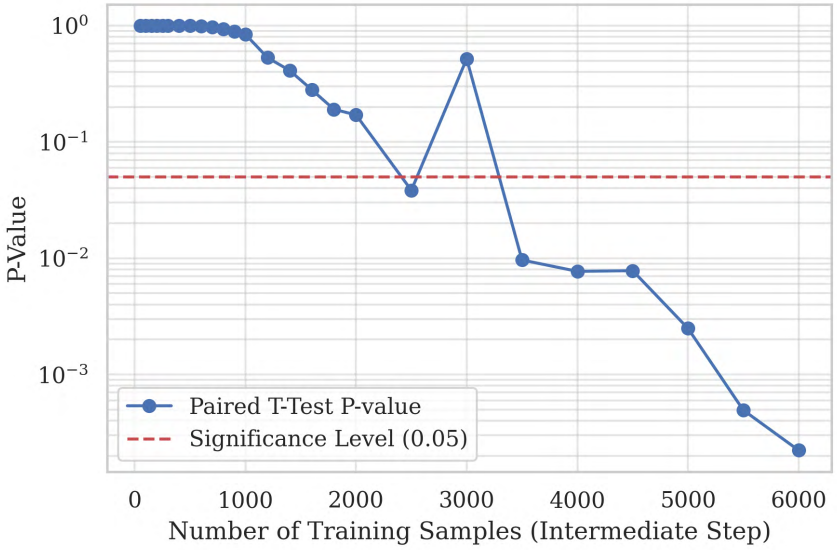

Figure 4. The plot displays the p-values from a one-sided t-test comparing the model's performance at intermediate training stages against its final performance. The x-axis represents the number of training samples, while the y-axis shows the corresponding p-value.

on the model type TabPFN. In particular, the structure of TabPFN with its in-context learning paradigm enables evaluations like this one, that would otherwise require a whole hyperparameter tuning run for each number of training samples.

The data selection strategy of the AL algorithm is helpful from an efficiency perspective. In particular, we can reach the level of accuracy of the model with the full dataset with already 1200 data points which is roughly 20 percent of the whole training dataset. Therefore, the computational complexity for training and inference is reduced by a factor of 25. This allows inference times in the realm of milliseconds compared to 10 seconds before. This represents an important speed improvement, particularly for high-throughput bioprocess development.

When there is a powerful GPU available that enables fast computations we can even use AL as a filter on the training dataset that picks the most informative points as long as the validation performance increases. So in this case we would use the 2500 data points for training that achieved the highest overall accuracy.

Model performance varies substantially across substances. While for most substances the accuracy increases with the number of data points, for EnPump, Nitrate and Yeast Extract the accuracy of both models decreases for dataset sizes over 2500. Looking at the data distributions in [6] we observe that these are the three most imbalanced substances in the dataset and mostly contain zero concentrations. It is likely that the AL algorithm predominantly selects zero-concentration samples, exacerbating the label imbalance that is harder to handle for the model. This also highlights the importance of carefully designing one's calibration dataset with methods like Design of Experiments.

For the future it would be interesting to compare other batch AL methods to the current approach, for instance [9], [10]. Or one could use the Core-Set approach from [18] and combine it with an informativeness criterion. Moreover, evaluating the benefit of AL approaches on other imbalanced Raman datasets in particular from fermentations might be an approach to increase the model accuracy on the small datasets present in biotechnological applications.

ACKNOWLEDGEMENTS

This research is funded by DFG's project MEDEA (Grant No. 561489561).

AUTHOR IDENTIFIERS

Author ORCIDs:
Lange Ch: 0009-0008-0100-8596
Martinez E: 0000-0002-2622-1579
Neubauer P: 0000-0002-1214-9713
Cruz Bournazou MN: 0000-0001-9461-4414

REFERENCES

1. Lange C, Seidel S, Altmann M, Stors D, Kemmer A, Cai L, Born S, Neubauer P, Bournazou MNC. A setup for automatic raman measurements in high?throughput experimentation. Biotech & Bioengineering 122:2751-2769 (2025). https://doi.org/10.1002/bit.70006
2. R. C. Rowland-Jones *et al.*, "Spectroscopy integration to miniature bioreactors and large scale production bioreactors–Increasing current capabilities and model transfer, " *Biotechnol. Prog.*, vol. 37, no. 1, p. e3074, Jan. 2021, doi: 10.1002/btpr.3074.
3. Wold S, Sjöström M, Eriksson L. Pls-regression: a basic tool of chemometrics. Chemometrics and Intelligent Laboratory Systems 58:109-130 (2001). https://doi.org/10.1016/s0169-7439(01)00155-1
4. Lange C, Borisyak M, Kögler M, Born S, Ziehe A, Neubauer P, Bournazou MNC. Comparing machine learning methods on raman spectra from eight different spectrometers. Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy 334:125861 (2025). https://doi.org/10.1016/j.saa.2025.125861
5. Hollmann N, Müller S, Purucker L, Krishnakumar A, Körfer M, Hoo SB, Schirrmeister RT, Hutter F. Accurate predictions on small data with a tabular foundation model. Nature 637:319-326 (2025). https://doi.org/10.1038/s41586-024-08328-6
6. Lange C, Altmann M, Stors D, Seidel S, Moynahan K, Cai L, Born S, Neubauer P, Bournazou MNC. Deep learning for raman spectroscopy: benchmarking models for upstream bioprocess monitoring. Measurement 258:118884 (2026). https://doi.org/10.1016/j.measurement.2025.118884
7. A. Vaswani *et al.*, "Attention Is All You Need, " Aug. 02, 2023, *arXiv*: arXiv:1706.03762. doi: 10.48550/arXiv.1706.03762.
8. B. Settles, "Active learning literature survey, " University of Wisconsin-Madison Department of Computer Sciences, 2009. [Online]. Available: https://minds.wisconsin.edu/handle/1793/60660
9. G. Citovsky *et al.*, "Batch Active Learning at Scale, " July 29, 2021, *arXiv*: arXiv:2107.14263. doi: 10.48550/arXiv.2107.14263.
10. J. T. Ash, C. Zhang, A. Krishnamurthy, J. Langford, and A. Agarwal, "Deep Batch Active Learning by Diverse, Uncertain Gradient Lower Bounds, " Feb. 24, 2020, *arXiv*: arXiv:1906.03671. doi: 10.48550/arXiv.1906.03671.
11. F. Zhdanov, "Diverse mini-batch Active Learning, " Jan. 17, 2019, *arXiv*: arXiv:1901.05954. doi: 10.48550/arXiv.1901.05954.
12. Bensch M, Schulze Wierling P, von Lieres E, Hubbuch J. High throughput screening of chromatographic phases for rapid process development. Chem Eng & Technol 28:1274-1284 (2005). https://doi.org/10.1002/ceat.200500153
13. L. Grinsztajn *et al.*, "TabPFN-2.5: Advancing the State of the Art in Tabular Foundation Models, " Nov. 11, 2025, *arXiv*: arXiv:2511.08667. doi: 10.48550/arXiv.2511.08667.
14. N. Hollmann, S. Müller, K. Eggensperger, and F. Hutter, "TabPFN: A Transformer That Solves Small Tabular Classification Problems in a Second, " Sept. 16, 2023, *arXiv*: arXiv:2207.01848. doi: 10.48550/arXiv.2207.01848.
15. S. Müller, N. Hollmann, S. P. Arango, J. Grabocka, and F. Hutter, "Transformers Can Do Bayesian Inference, " Aug. 13, 2024, *arXiv*: arXiv:2112.10510. doi: 10.48550/arXiv.2112.10510.
16. Scheffer T, Decomain C, Wrobel S. Active hidden markov models for information extraction. Lecture Notes in Computer Science :309-318 (2001). https://doi.org/10.1007/3-540-44816-0_31

17. Shannon CE. A mathematical theory of communication. Bell System Technical Journal 27:379-423 (2013). https://doi.org/10.1002/j.1538-7305.1948.tb01338.x
18. O. Sener and S. Savarese, “Active Learning for Convolutional Neural Networks: A Core-Set Approach, ” June 01, 2018, *arXiv*: arXiv:1708.00489. doi: 10.48550/arXiv.1708.00489.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Utilizing Machine Learning for Phenomena-based Synthesis of Intensified Process Flowsheets

Omar Alqusair[ab], and Jie Li[a]*

[a] Centre for Process Integration, Department of Chemical Engineering, The University of Manchester, Manchester, M13 9PL, UK
[b] Department of Chemical Engineering, College of Engineering, King Saud University, P.O. Box 800, Riyadh 11421, Saudi Arabia
* Corresponding Author: jie.li-2@manchester.ac.uk

ABSTRACT

The increasing demand for energy, water, and chemical products signals the need for more sustainable and efficient process design methodologies. Traditional methods for conceptual process design constrains the exploration of novel and intensified process alternatives, as they rely on prior knowledge in defining the design space. Previous studies employing bottom-up approaches, such as phenomena building blocks (PBBs), suggest that the synthesis of complex bottom-up flowsheets remains computationally challenging and is thus limited to the synthesis of individual units of operation. This work proposes a bottom-up, data-driven framework for process synthesis and intensification based on phenomena building blocks (PBBs), in which process flowsheets are constructed from their underlying physical and chemical phenomena rather than conventional units of operation. The proposed framework introduces a phenomena-based text representation and data collection module. Furthermore, a sequence training and generation module is developed, which learns patterns governing PBB interactions and placement within flowsheets. A case study on ethylene glycol production is presented, in which nine fundamental phenomena yield 49 distinct PBB tokens. Results show that the trained model successfully reproduces established flowsheet structures and generates flowsheets using novel PBB combinations, highlighting its potential to support sustainable and intensified process designs.

Keywords: Process Design, Process Synthesis, Process Intensification, Artificial Intelligence, Machine Learning

INTRODUCTION

In conceptual process design, finding the optimal flowsheet for a process often involves the predefinition of units of operation and their possible configuration [1]. This paradigm limits the discovery of novel process alternatives involving unconventional units of operation and configurations, as it relies on expert knowledge to define the process's design space [2]. Rather than defining existing units of operation, phenomena-based methods offer an innovative solution to process design, as they rely on defining existing chemical components and their physicochemical interactions. This led to the introduction of phenomena building blocks (PBBs) [3], where abstract blocks are used to model chemical species and the phenomena that govern them. Then a process flowsheet may be constructed using a combination of these PBBs.

Demirel et al. [4] proposed a phenomena-based optimization framework, where superstructures of PBBs are formulated as a mixed integer non-linear programming problem (MINLP), that are solved to obtain the optimal allocation of PBBs in a 2D grid, leading to the synthesis of intensified process flowsheet. Then, Demirel et al. [5] modified the design framework by explicitly defining the grid structure, introducing an initializing strategy, and including a two-block formulation for vapor-liquid equilibria, showcasing the framework's ability in intensifying an ethylene glycol production process. Moreover, Monjur and Hasan [6] built upon this modification to develop a design framework for the intensification of natural gas to

methanol process that includes membrane reactors.

Krone et. al. [7] proposed an optimization framework employing mixed-integer non-linear programming (MINLP) to synthesize distillation columns using superstructures of VLE building blocks. To address the large design space of the problem, which often prevents MINLP solvers from converging, they introduced a structural screening procedure to eliminate infeasible process alternatives prior to optimization. On the other hand, Cortés-Peña and Zavala [8] explored the application of phenomena-based representation for the simulation of process flowsheets formulated as graphs consisting of equation and variable nodes.

Nonetheless, using mathematical programming to solve superstructures of PBBs remains challenging. One challenge is the computational complexity of handling a large combinatorial design space. Another challenge is that achieving optimal design alternatives depends heavily on the formulation and may require the inclusion of disjunctive programming, which increases the model complexity. Moreover, solving these complex models requires good initial conditions, an effective initialization strategy must be developed to complement formulation.

In data-driven approaches, a design framework often trains a machine-learning (ML) model on process data to learn structural relationships and decision-making rules directly from data rather than prescribing them using expert knowledge. Because such frameworks enable the simultaneous generation and evaluation of process alternatives, they can also navigate large design spaces to discover new alternatives [9].

Inspired by how the Simplified Molecular Input Line Entry System [10] (SMILES) can represent molecular structures using text, d'Anterroches [11] proposed the Simplified Flowsheet Input Line Entry System (SFILES). Moreover, Zhang et. al. [12] proposed a pattern recognition framework that can detect patterns for flowsheets represented using SFILES, and discovered patterns that match those described by heuristic rules in the literature. Nabil et. al. [13] employed some elements of the SFILES representation in developing vocabulary and grammar rules tailored for the synthesis of a supercritical CO_2 power cycle, and trained a long short-term memory (LSTM) model to predict the syntax of a text-based process flowsheet.

With recent interest in data-driven approaches, further developments on the SFILES notation have been made. SFILES 2.0 [14] expanded the notation to capture more interactions between units of operation and represent control modules within flowsheets [15]. While eSFILES [16] built on the original representation to establish hierarchical grammar trees and process ontology that can be utilized for machine learning. Li [17] demonstrated that PBBs can be employed for the synthesis of a hybrid extraction and distillation unit using phenomena-based string representation. Using a ML model on a dataset of vapor-liquid equilibrium (VLE), liquid-liquid equilibrium (LLE), and mixing building blocks. The results highlighted the framework's potential in generating novel extraction-distillation configurations with lower separation costs.

Nevertheless, recent works remain constrained to the representation of phenomena across single units of operations, as extending these approaches to entire process flowsheets will result in large and complex problems that are difficult to solve and require expert knowledge.

This work adopts a phenomena-based, data-driven approach for conceptual process design, enabling the automatic completion of process flowsheets. The proposed framework aims to represent complex process flowsheets as structured sequences of PBB tokens and train an ML model to learn relationships between PBBs and their placement within process flowsheets, providing a computationally lighter alternative for exploring large design spaces compared to superstructure-based methods. A case study for the production of ethylene glycol from ethylene oxide and water is highlighted.

METHODOLOGY

The proposed methodology follows a two-step approach: (1) text-representation and rule-consistent data preparation; (2) LSTM model training, finetuning, and generation of string-based flowsheets.

Data Preparation

An existing text-based representation paradigm, the SFILES 2.0 [14], is adopted for encoding process flowsheet information. In this notation, strings are read left to right, with parentheses-enclosed words representing a unit of operation token, and the tokens appearing next to one another have an implicit connection from left to right. However, we modify the notation to be based on PBBs rather than units of operation.

Each block token can represent one or more phenomena and connects to other blocks via directed, branch, and recycle/bypass streams to form the process topology. SFILES 2.0 encodes a flowsheet as a single linearization of a process graph, augmented with special symbols to mark branches. The same augmentation style is retained, but the vocabulary and grammar rules are modified to suit a phenomenon-based representation.

A concise vocabulary is defined to align with the PBBs representation; consequently, tokens fall into three classes: (a) input/output blocks, (b) phenomena tokens, and (c) topology control tokens. Input/output blocks are process-specific and rely on a prior definition of available raw materials and possible products. With input blocks being denoted as (ik) where i refers to an input, and k refers to the chemical species. Similarly, for output blocks, (ok) is used to denote an outlet block, where o

type	1-phenomenon tokens	subtype
Mass Change	(m): mixing	Mix
	(s): splitting	Split (no comp. change)
	(f): flash	Split (1-stage comp. change)
	(t): distillation	Split (n-stage comp. change)
Energy Change	(c): cooling	Thermal
	(h): heating	Thermal
	(p): pressurizing	Mechanical
	(d): depressurizing	Mechanical
Reaction	(r): reaction	Homogenous Reaction

(a)

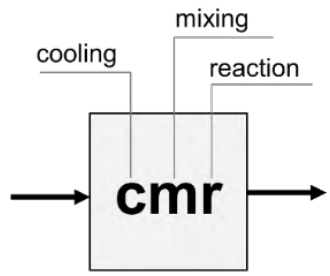

(b)

{...(cmr)...}

(c)

Figure 1. Text-based representation of PBBs: **(a)** one-phenomenon tokens, **(b)** a PBB, and **(c)** its equivilant string.

refers to an outlet, and k refers to the chemical species.

Furthermore, phenomenon tokens are derived from nine basic phenomena across three types, with each phenomenon being denoted using a unique letter, as shown in **Figure 1.a**. From these basic phenomena, 49 tokens are defined. These tokens were defined to allow one or more phenomena to exist in a block, as illustrated in **Figure 1.b**, the tokens follow simple thermodynamic rules (e.g., cooling and heating cannot occur in the same block). In a string format, input/output tokens and phenomena tokens are enclosed by parentheses to clearly distinguish them, as shown in **Figure 1.c**.

On the other hand, topology control tokens use symbols to denote a special action on the topology of the process, such as a branch entering or exiting, or a recycle or bypass stream from one block to another. **Table 1** shows the full list of the topology control tokens. Moreover, recycle/bypass streams come in pairs where each inlet must have an outlet, where '#' is an integer number that appears in both the stream's inlet and outlet. For a process with 2 input blocks and 2 output blocks, and one allowed recycle/bypass stream, this makes up a vocabulary size of 59, which consists of 49 phenomena tokens, 6 topology control tokens, and 4 input/output tokens. This vocabulary is not only expressive enough to support complex flowsheet patterns, but also to encode the grammar and thermodynamic rules that are used to generate the initial dataset.

Table 1: Process Topology Tokens.

Token	Description
<&	Start of a nest for a mixing branch.
<&|	End of a nest for a mixing branch.
[	Start of a nest for a splitting branch.
]	End of a nest for a splitting branch.
<#	Recycle/bypass stream inlet
#	Recycle/bypass stream outlet

The initial dataset is randomly generated using a rule-guided procedure, where two sets of rules govern how a string-based dataset of process flowsheets is generated. The set of heuristic rules is shown the supplementary material. In **Table S.1**; these rules consist of grammar rules that describe how text-based representation is achieved, and thermodynamic rules that ensure that a design-informed string is generated. On the other hand, the sampling logic rules establish string targets as shown in **Table S.2**. Some of these rules follow probabilistic controls to randomly sample the next token in the string.

Also, these rules can be modified to control the complexity of the generated string-based flowsheet, such as modifying the maximum length of strings or the size of nests for branches. Although the sampling logic rules can be considered heuristic, they enforce constraints related to the string generation logic that are not grammar or thermodynamic rules.

Finally, an algorithm that combines the heuristic and sampling logic rules is used to sequentially build up strings, following a token-to-token approach. For each position in the string, a list of eligible tokens is updated, so that the algorithm picks a token that follows the rules. Moreover, for each step, the algorithm conducts checks related to which tokens have appeared in the string, their count, and whether sampling targets have been met. Nonetheless, the algorithm only requires the definition of the number of strings to be generated, and can output the full list of generated strings, including unique ones.

Model Training and Flowsheet Generation

A machine-learning pipeline is employed to build a data-driven model capable of autocompleting string-based flowsheets and generating new valid strings. The pipeline begins by retrieving the initial dataset from the data preparation step, preprocessing it into a suitable format, training an LSTM model, and finally generating

new strings. The LSTM architecture is chosen because it can handle short and long-range interactions in sequences, which is essential in capturing various token dependencies in string-based flowsheets.

Each flowsheet string is tokenized and mapped to integer IDs in $[1, v]$, where v is the vocabulary size; with ID =0 reserved for padding. The IDs are to be used by an embedding layer in the LSTM model, which maps integers to d-dimensional vectors, which are then fed to an LSTM cell. Label encoding is used to assign integer IDs to tokens, which can then be used as the embedding layer input. While the integers themselves carry no physical meaning, encoding strings into integers serves as a reference for the learning.

To ensure that strings have uniform sequence lengths during training, all sequences are padded to a global maximum length L_{max}, corresponding to the longest string in the dataset. Nonetheless, this padding token is masked when calculating the loss and accuracy values. Given a dataset of sequences as input, output token sequences are created by shifting input sequences by one position to the left and removing the first token; tokens in any given position in the shifted sequence are output tokens (ground truths) to their corresponding tokens in the original sequence.

A unidirectional LSTM is adopted with input embedding and linear projection to the vocabulary to predict the next token in a sequence. The model is trained on a synthetically generated dataset, capturing PBB strings for the production of ethylene glycol from ethylene oxide and water. The training uses class-weighted pad-masked cross-entropy, averaging the loss across all non-pad positions to handle variable-length sequences and class imbalance.

The dataset is split into 80/10/10% for training, validation, and testing respectively, with the training dataset being further divided into batches of equal lengths. The embedded token is passed to an LSTM module consisting of l-layers of hidden size h, then projected to g logits per position that convert them to probabilities representing next-token placement.

The training commences with the first LSTM consuming the embedded sequence one token at a timestep, as shown in **Figure 2**, carrying hidden states forward. Initial states are set to zero per sequence. At each timestep, the LSTM's layer output is passed through a linear projection to produce logits over the vocabulary size. Softmax converts these logits to a next-token probabilities that are compared with the ground truth of the next-token to calculate the loss function. For timestep t, the loss function is calculated using equation (1), where L_t is the class-weighted cross-entropy function, m is the masking term with 0 for padding tokens and 1 for other tokens, w_y is the weight of the class, and p_y is the probability of picking that class as the next token.

$$\mathcal{L}_t = -m w_y \log p_y \qquad (1)$$

At the end of an epoch, the loss function across all batches is averaged to produce a class-weighted cross-entropy for that epoch, which is then used in backpropagation to update the hidden states and embedding table for the next epoch. Moreover, accuracy calculation can be performed to better understand how the predictions are. For a single timestep, the accuracy is calculated by taking the highest probability of non-pad tokens in the next-token prediction and checking whether it matches the ground truth, assigning an accuracy of 1 if it does, and 0 if it does not. This is then averaged across all batches in an epoch to get that epoch's accuracy.

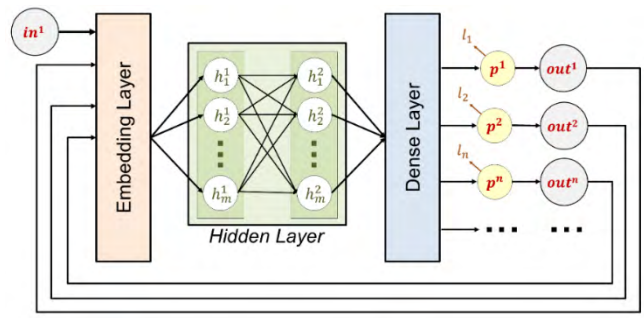

Figure 2. Unidirecotional LSTM module used in training.

As the aim of the training is to get a model that minimizes the cross-entropy loss function, selecting appropriate hyperparameter values is crucial for this stage. Consequently, a grid search will be used where the combination of selected hyperparameter values will be used to get their initial estimates. Furthermore, Adam [18] optimizer is to be employed to dynamically change the learning rate across batches, rather than using a global learning rate. Finally, Optuna [19] optimizer is to be used for a final round of hyperparameters finetuning.

After a trained model is obtained, it is used to generate new phenomena-based flowsheets using text-based representation. The generation procedure employs a top-k probabilistic decoding strategy, where the next token in the sequence is sampled from the k most likely predictions for the output token. First, a sequence is initiated by randomly sampling an input token in the first timestep. The sequence is then autocompleted until a maximum string length is reached, or the stopping criteria are met, namely, the presence of all output tokens and balanced process topology symbols. Then, each generated sequence is checked for adherence to the original heuristic rules and flagged as a valid sequence.

RESULTS

The first step in data preparation is to establish the vocabulary set representing tokens that can appear in

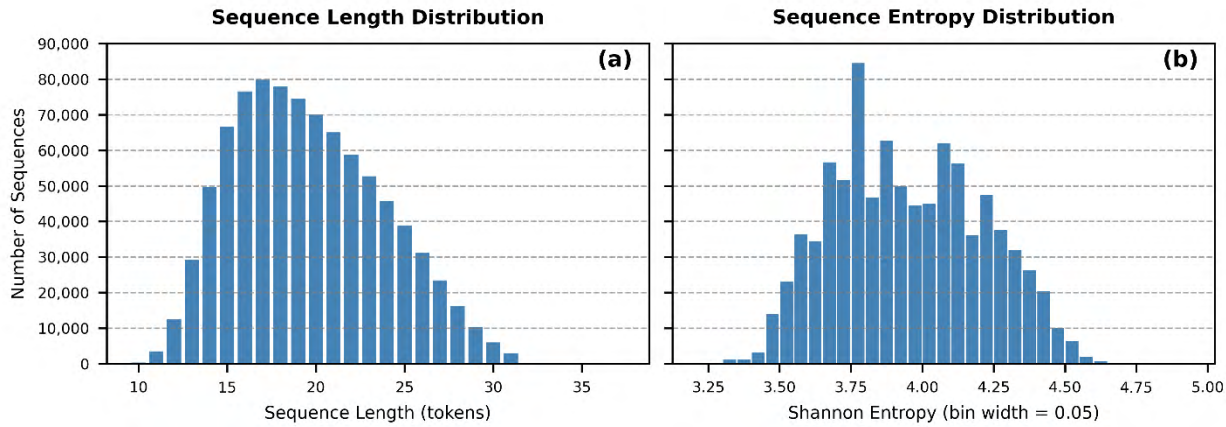

Figure 3. Initial dataset analysis: **(a)** sequence length distribution, and **(b)** entropy distribution.

string-based flowsheets. The proposed framework incorporates all nine basic phenomena to enable a more comprehensive flowsheet representation for a case study on ethylene glycol production from ethylene oxide and water. Consequently, many multi-phenomena tokens introduced during the token definition stage were not previously examined within the context of ethylene glycol production. Furthermore, two input tokens, (iW) and (iEO), were defined to represent the water and ethylene oxide feed streams respectively. Similarly, two output tokens. (oW) and (oEG), were defined to represent the water and ethylene glycol product streams, respectively, with (oEG) designated as the desired product output token.

A maximum of two recycle/bypass streams was allowed in the process. Based on these definitions, a vocabulary set containing 61 tokens was constructed, comprising 49 PBB tokens, 4 input/output tokens, and 8 process topology tokens. Consequently, a total of 2, 000, 000 strings were prepared, yielding 900, 425 valid strings that complied with the heuristic rules. Among these, 892, 768 were unique, non-repeating strings. A maximum sequence string length of 30 is used for the generation procedure; however, the observed maximum sequence length is 37. This is due to the post-processing step in the procedure, which enforces the sampling logic rules, such as balancing recycle and branching symbols, or ensuring that the desired output appears in the string.

As a result, the sequence length distribution shown in **Figure 3.a** exhibits a unimodal, centred around sequence length 17, reflecting the constrained randomness of the generated sequences. Additionally, the sequence entropy distribution presented in **Figure 3.b** highlights entropy values ranging from 3.3-4.7 bits, corresponding to 10-26 tokens respectively. This indicates that the generated sequences have sufficient randomness and token diversity to represent the design space of the process.

The LSTM model is trained through a multi-stage procedure, which starts by pretraining using an initial set of hyperparameters to obtain baseline learning curves and assess the model's ability to learn. Then, to find the minimum cross-entropy loss function, these hyperparameters were optimized using grid-search exploration, Adam optimizer, and Optuna optimizer. The final stage is to generate novel strings and validate their structure and adherence to heuristic rules. All training was done using the PyTorch library, utilizing a CUDA-compatible GPU (Nvidia 4070 Super).

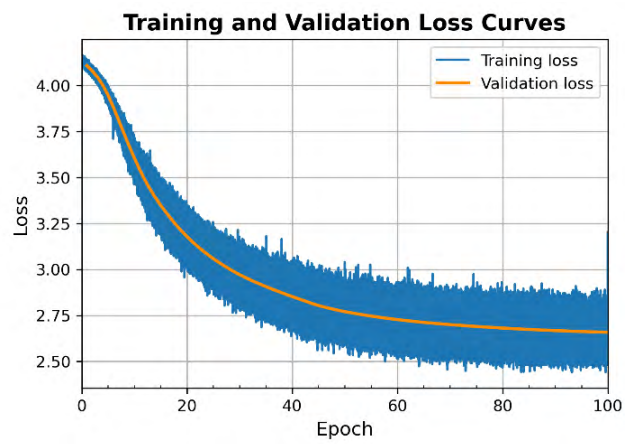

Figure 4. Training/validation curves in pretraining.

The pretraining stage was carried out using initial hyperparameter values similar to those used by Hirtreiter et. al. [16] in their SFILES 2.0 representation. Using a batch size of 32, hidden size and embedding dimension of 64, 1 LSTM layer, and a learning rate of 3e-4. Stochastic Gradient Descent (SGD) is initially used to calculate the gradient in backpropagation. **Figure 4** shows the evolution of training and validation loss over 100 epochs, where the training curve is composed of per-batch loss evaluations to demonstrate how the model is changing the LSTM states for each batch, while the validation curve is evaluated at the end of each epoch. The curve

shows a gradual decrease in the cross-entropy loss function, converging to a loss of 2.63.

Although these results confirm that a baseline model is capable of learning sequential ordering of how tokens appear in the strings, the slow convergence indicates the chosen hyperparameters are not optimal. And any reruns using the same -or different- datasets will be computationally expensive. Moreover, early stopping is to be implemented in the subsequent stages, where training terminates when the validation losses across the last 5 epochs are within 0.001. The grid search provided coarse insight into the suitability of certain hyperparameter ranges, as shown in **Figure S.1** and **Figure S.2** in the supplementary material, highlighting how the training loss responds to variations in batch size, learning rate, hidden dimensions, and embedding dimensions over the first ten epochs.

Consequently, an Optuna study consisting of 50 trials was performed using continuous sampling of the learning rate and discrete sampling for all other hyperparameters. An additional training run using a batch size of 512 and the optimized hyperparameters produced the loss curve shown in **Figure 5**. Compared with the pretraining results, the optimized model exhibits faster convergence, with the training loss stabilizing within the first epoch and converging to a loss of around 2.58.

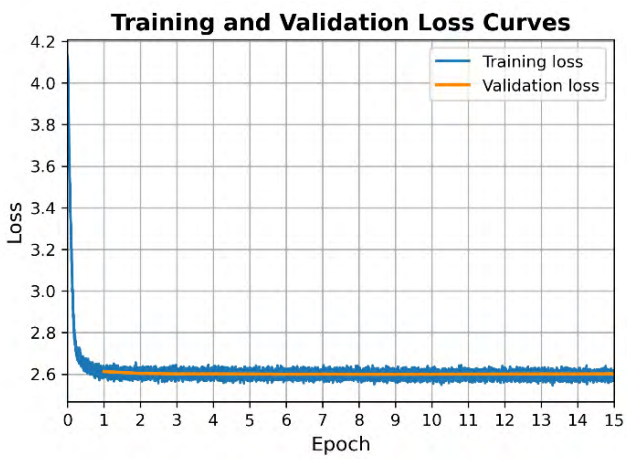

Figure 5. Training/validation curves after finetuning.

The rapid convergence of the optimized model indicates that the selected hyperparameters not only offer a slight improvement in accuracy but also enhance the efficiency and smoothness of the training process. To evaluate the generative capabilities of the trained LSTM model, a generation module was constructed. It starts by retrieving the best LSTM model after the last training cycle and implements basic decoding rules. These rules include balancing or process topology tokens, the presence of raw materials and desired products, and termination conditions. Using top-k sampling (with k=20), the module is tasked with generating a batch of 10, 000 flowsheet strings without any explicit rules or structural guidance.

Across the full batch, the module outputs 9, 760 valid sequences, with valid sequences defined as those that do not violate the heuristic rules used in generating the initial dataset, demonstrating that the learned sequence distribution is highly consistent with the rules embedded in the training data. This high validity is notable given that the generator applies no grammar rules, thermodynamic rules, or sampling logic rules. Furthermore, only 170 valid sequences (about 1.7% of the generated dataset) appeared in the initial dataset, indicating that the model is not simply memorizing the training dataset. This means that the trained model can synthesize novel, previously unseen string-based flowsheet configuration that do not violate any of the heuristic rules.

The model was also tested for the autocompletion of strings, where the first couple of tokens are provided and the model predicts a valid continuation of the string. The module first checks whether the given initial sequence of tokens is probable under the model, with structurally inconsistent combinations being labelled as an invalid input sequence. When a probable input sequence is provided, the model extends it into a full flowsheet string. Examples of generated and autocompleted strings are presented in **Table 2**. The examples demonstrate the model's ability to generate alternatives that include no recycle streams, one or multiple recycle streams, alternatives with or without side products, and complex reaction-separation sequences.

Table 2: Examples of Input strings and their corresponding autocompleted strings.

Input String	Autocompleted String
(iW)	(iW)(m)<&|(iEO)(h)(d)&|(dr)(rt)[(d)(oEG)](h)(d)(oW)
(iEO)	(iEO)(hm)<1(dmr)<&|(iW)(d)&|(pt)[(h)(d)(oEG)](df)1
(iW)(p)(m)	(iW)(p)(m)<&|(iEO)(d)&|(dr)(fr)[(d)(h)(oEG)](oW)
(iW)(cr)	invalid input

A parsing module is also developed to translate each string into a graph representation that can be easily interpreted. An example of this is provided in **Figure 6**, showing a string being parsed into a graph-representation, with input PBB blocks highlighted in light red, output PBB block in light green, red arrows for directed edges on the main branch, grey arrows for directed edges on side branches, and green arrow for the recycle stream.

CONCLUSIONS

The reliance on predefined units of operation in traditional process design constrains the exploration of

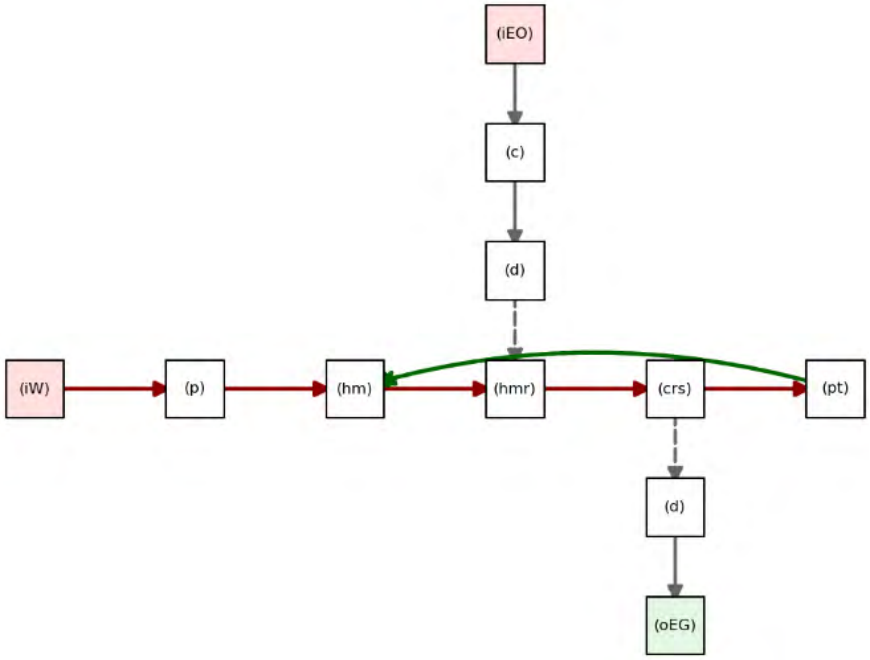

Figure 6. Process Graph of the Generated String '(iW)(p)(hm)<1(hmr)<&|(iEO)(c)(d)&|(crs)[(d)(oEG)](pt)1'.

novel and intensified process alternatives. Recent phenomena-based studies have highlighted the advantages of representing processes through fundamental physicochemical interactions, although their extension to complex, complete flowsheet synthesis remains challenging.

The proposed framework demonstrates the potential of a text-based data-driven approach for conceptual process design, and its ability in generating novel and large design spaces for chemical processes. Using an initial dataset of about 900, 000 alternatives, the model underwent multiple training stages, including pretraining with initial hyperparameters, followed by finetuning. The optimized model was used to generate 10, 000 alternatives, of which 97.6% implicitly follow the heuristic rules. Furthermore, among the generated alternatives, 98.3% were novel and had not appeared in the initial dataset, indicating that the model successfully learned phenomena building block relationships using the modified SFILES 2.0 representation.

Nonetheless, several limitations arise when implementing text-based representation for PBB flowsheets. The current vocabulary and grammar rules do not provide a true phenomenon representation for separation. This is due to the inherent restriction associated with the adopted SFILES 2.0 representation, which does not allow mixing and splitting to occur simultaneously in the same token. Similarly, mixing and splitting tokens are restricted to a single branch each, further limiting the representation of separation phenomena.

Although the trained model generated novel strings, their thermodynamic feasibility is to be investigated. The heuristic and sampling logic rules ensure structural consistency and basic PBB logic. However, process conditions need to be incorporated, and the process alternatives should be evaluated, simulated, and optimized. As process conditions were not considered in the current framework, certain thermodynamic interactions may be infeasible during evaluation, such as using a simple distillation for azeotropic mixtures. Another limitation is that the LSTM model was trained only on synthetically generated data, without incorporating real process data or simulation results. Future work should address these limitations to achieve a more robust phenomena-based design methodology.

DIGITAL SUPPLEMENTARY MATERIAL

The heuristic and sampling logic rules, and hyperparameters grid search results can be accessed via the link: https://psecommunity.org/LAPSE:2026.0019

ACKNOWLEDGEMENTS

Omar Alqusair would like to appreciate financial support from King Saud University.

AUTHOR IDENTIFIERS

Author ORCIDs:
Alqusair O: 0009-0006-6461-4388
Li J: 0000-0001-5196-2136

REFERENCES

1. Pistikopoulos EN, Tian Y. Advanced modeling and optimization strategies for process synthesis. Annual Review of Chemical and Biomolecular Engineering 15:81-103 (2024). https://doi.org/10.1146/annurev-chembioeng-100522-112139
2. Karagoz AT, Alqusair O, Liu C, Li J. Advances in conceptual process design: from conventional strategies to ai-assisted methods. Chinese Journal of Chemical Engineering 84:60-76 (2025). https://doi.org/10.1016/j.cjche.2025.05.014
3. Lutze P, Babi DK, Woodley JM, Gani R. Phenomena based methodology for process synthesis incorporating process intensification. Ind. Eng. Chem. Res. 52:7127-7144 (2013). https://doi.org/10.1021/ie302513y
4. Demirel SE, Li J, Hasan MMF. Systematic process intensification using building blocks. Computers & Chemical Engineering 105:2-38 (2017). https://doi.org/10.1016/j.compchemeng.2017.01.044
5. Demirel SE, Li J, El-Halwagi M, Hasan MMF. Sustainable process intensification using building blocks. ACS Sustainable Chem. Eng. 8:17664-17679 (2020). https://doi.org/10.1021/acssuschemeng.0c04590
6. Monjur MS, Hasan MMF. computer?aided process intensification of natural gas to methanol process. AIChE Journal 68: (2022). https://doi.org/10.1002/aic.17622
7. Krone D, Esche E, Skiborowski M, Repke JU. Optimization-based process synthesis by phenomena-based building blocks and an MINLP framework featuring structural screening. Computers & Chemical Engineering 194:108955 (2025). https://doi.org/10.1016/j.compchemeng.2024.108955
8. Y. Cortes-Pena and V. Zavala, "Phenomena-Based Graph Representations and Applications to Chemical Process Simulation, " Sep. 11, 2025, ChemRxiv. doi: 10.26434/chemrxiv-2025-86jhv
9. A. Khan, "A framework for autonomous process design: towards data-driven and knowledge-driven systems", University of Cambridge, 2022, doi: https://doi.org/10.17863/CAM.93910
10. Weininger D. SMILES, a chemical language and information system. 1. introduction to methodology and encoding rules. J. Chem. Inf. Comput. Sci. 28:31-36 (2002). https://doi.org/10.1021/ci00057a005
11. Ludvigsen E, , Telecka A, , Li T, , Ndoni S, , Taboryski R, , Li T, . Self-assembly of block-copolymer structures forin-situ nano-lithography. JSAME 6:1-1 (2018). https://doi.org/10.13052/jsame2245-4551.2018014
12. Zhang T, Sahinidis NV, Siirola JJ. Pattern recognition in chemical process flowsheets. AIChE Journal 65:592-603 (2018). https://doi.org/10.1002/aic.16443
13. Nabil T, Noaman M, Morosuk T. Data-driven structural synthesis of supercritical CO2 power cycles. Front. Chem. Eng. 5: (2023). https://doi.org/10.3389/fceng.2023.1144115
14. Vogel G, Hirtreiter E, Schulze Balhorn L, Schweidtmann AM. SFILES 2.0: an extended text-based flowsheet representation. Optim Eng 24:2911-2933 (2023). https://doi.org/10.1007/s11081-023-09798-9
15. Schulze Balhorn L, Degens K, Schweidtmann AM. Graph-to-sfiles: control structure prediction from process topologies using generative artificial intelligence. Computers & Chemical Engineering 199:109121 (2025). https://doi.org/10.1016/j.compchemeng.2025.109121
16. Mann V, Sales-Cruz M, Gani R, Venkatasubramanian V. Esfiles: intelligent process flowsheet synthesis using process knowledge, symbolic AI, and machine learning. Computers & Chemical Engineering 181:108505 (2024). https://doi.org/10.1016/j.compchemeng.2023.108505
17. Li J. Learning hybrid extraction and distillation using phenomena-based string representation. Systems and Control Transactions 3:300-307 (2024). https://doi.org/10.69997/sct.171879
18. D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization." 2017. [Online]. Available: https://arxiv.org/abs/1412.6980
19. Akiba T, Sano S, Yanase T, Ohta T, Koyama M. Optuna. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining :2623-2631 (2019). https://doi.org/10.1145/3292500.3330701
20. E. Hirtreiter, L. S. Balhorn, and A. M. Schweidtmann, "Towards automatic generation of piping and instrumentation diagrams (P&IDs) with artificial intelligence, " arXiv preprint arXiv:2211.05583, 2022

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Separation of Concern Capabilities of Information Model Candidates for Modular Plant System Engineering Lines

Tobias Kock[a]*, Isabell Viedt[a], Amy Koch[a], Leon Urbas[a]

[a] TU Dresden University of Technology, Chair of Process Control Systems & Process Systems Engineering Group, 01062 Dresden, Germany
* Corresponding Author: tobias.kock@tu-dresden.de.

ABSTRACT

Pharmaceutical and fine chemical industries face strong pressure to shorten time-to-market while maintaining compliance with complex regulatory frameworks. These conflicting demands require rapid process design, validation, and scale-up. Modular production plants standardized in VDI 2776 and VDI/VDE/NAMUR 2658 have emerged as a promising strategy to shorten engineering and validation efforts. The Product–Process–Resource (PPR) philosophy represents a key approach to efficient data management in modular plant engineering. It enables the separation of different flexibility dimensions into distinct, relevant aspects that can ideally be exchanged or modified independently. To realize this principle in practical applications, formalized information models and ontologies serve as a key enabler for structuring and managing semantic data. This work investigates several information models and ontologies for the process engineering domain regarding their suitability to support separation according to the dimensions defined by the PPR approach. As a metric for separability, coupling measures from Design Mapping Matrix (DMM) research are applied. This work provides the foundation for improving computer-aided engineering tools to efficiently handle modular processes while considering the reshaped information distribution.

Keywords: Process Design, Industry 4.0, AAS, DEXPI, Modular Plants

INTRODUCTION

The chemical and fine chemical industry has an increased need for shorter times to market [1]. One approach to address this challenge is the utilization of modular plants [2]. In this context, the chemical process is divided into sub-functions, each provided by a Process Equipment Assembly (PEA). PEAs may be engineered prior to or in parallel with process development. As detailed PEA engineering is performed independently by the manufacturer, process-specific requirements are reduced to selecting an appropriate PEA topology.

This approach results in a significant knowledge shift compared to traditional monolithic plants. In monolithic plants, most of the knowledge about the chemical process, its operating conditions, recipes, operation, and detailed plant engineering is centralized within a single organization, often supported by service providers [3]. In modular plants, this knowledge is distributed among different parties, each specialized in a specific domain. [4] introduced a philosophy that represents this changed knowledge distribution and its functional implications as a key approach for efficient data management in modular plants. This philosophy is known as the Product-Process-Resource (PPR) model. Building on this concept, the model has since then been extended with additional requirements[5][4], [5], [6], [7].

Because the different domains represented in the PPR model must interact and exchange information, efficient and standardized information exchange is a key factor in realizing the advantages promised by the modular plant approach [8], [9]. For instance, VDI/VDE/NAMUR 2658 [10] introduced the Module Type Package (MTP) as a standardized, machine-readable exchange format that enables rapid integration of PEAs into the Process Orchestration Layer (POL) of a modular plant. The MTP was specifically designed to address a particular information exchange requirement.

In addition, further standards exist that were not specifically developed for modular plants but can

nevertheless contribute to the need for standardized, machine-readable information exchange. [11] demonstrated the application of the Data Exchange in the Process Industry (DEXPI Plant & Process) standard [12] in the context of modular plants. Furthermore, it showed that a piping and instrumentation diagram (P&ID) modeled using DEXPI Plant can be successfully transferred to an MTP representation [11]. Furthermore, [9] demonstrated that interlinked information between DEXPI Plant and OntoCAPE [13] can support modular plant design within the context of the PPR model.

This work aims to provide a structured overview of how existing information models contribute to defined requirements derived from the PPR model. The overview supports the future development of machine-readable information exchange throughout the engineering and operation of modular plants. To enable systematic mapping, only information models that provide a sufficiently high level of abstraction to allow allocation within the PPR model are considered. At the same time, the evaluated information models are required to offer a technical specification that is detailed enough to enable matching at the class level.

PRODUCT PROCESS RESOURCE MODEL

In [4] the authors describe the structure of a digital twin for the process industry from a user-oriented perspective. They identify several data models that play an important role throughout the lifecycle of a chemical process plant. The authors argue that a plant-independent description level is a prerequisite for the effective use of digital twins and for the efficient and effective design of new production methods and processes [4]. To address this, the digital twin is structured into partial models that can be created and maintained by different disciplines, departments, and organizations. Overall, the digital twin is subdivided into six models:

- **Product model:** Contains information about the intended product.
- **Process model**: Contains information about thermodynamic, chemical, physical, and control-engineering relationships.
- **Resource model:** Contains information about the mechanical functionality and operational attributes of the plant and equipment structure.
- **Transformation model:** Contains product-related information about requirements for process control as well as individual transformation and separation steps. It is derived by combining the product and process models.
- **Capabilities model:** Contains plant-related information about the capability of equipment to perform specific process-engineering operations, including static and dynamic parameters. It is derived by combining the process and the resource models.
- **Operation model**: Contains information about optimized operating parameters and recipes to enable optimized production scenarios. It is derived by combining the transformation and capabilities models.

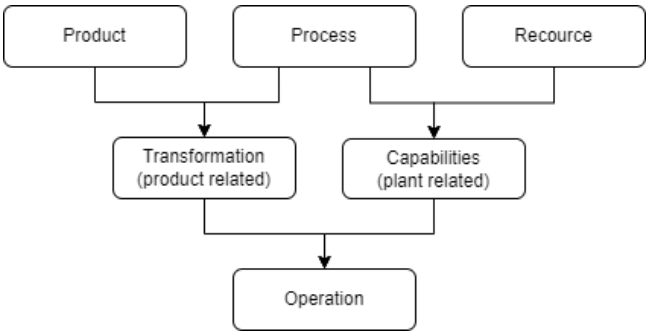

Figure 1. Visualization of the proposed Digital Twin structure according to the PPR-Model by [4].

In [5] several general challenges arising from this structure are highlighted, including:

- quality assurance and specification of simulation models,
- open and standardized interfaces,
- simulation and optimization algorithms that consider intellectual property (IP) issues, and
- model structuring and semantic description as key research and development challenges.

In the following section, requirements resulting from the different disciplines utilizing the model are summarized.

Product Model

The product model encompasses information related to the properties and specifications of chemical components and products. It contains component- and substance-specific properties, including physical properties of component mixtures and detailed product specifications [4] [5]. In addition, the product model incorporates specifications and data sheets [14] and provides information on product and component storability, safety- and transport-related requirements, as well as application and sales aspects [4].

Process Model

The process model contains information on thermodynamic, chemical, and physical engineering relationships, such as material transport coefficients and

reaction kinetics [4]. It further provides the information required for use in process simulation environments and includes the equations employed for process modeling and analysis [4], [5].

Resources Model

The resource model represents equipment-related information of the production system. It models attributes and characteristics of equipment, including properties specific to mobile equipment, and provides a structured description of equipment-related characteristics relevant for plant operation and engineering [4], [5]. In addition, the resource model includes a comprehensive list of equipment elements forming the plant structure [8].

Transformation Model

The transformation model contains equipment-independent information describing the production process. This includes laboratory measurement data, equipment-independent simulation models, and conceptual design knowledge related to the process [5], [8]. Furthermore, the transformation model captures key process characteristics such as residence times, throughputs, and risk potentials [4]. It also defines the capability requirements that must be fulfilled by the corresponding capabilities model [9].

Capabilities Model

The capabilities model describes plant-related information characterizing the ability of equipment to perform specific process-engineering operations. It includes stationary and dynamic parameters of a largely product-independent process plant, as well as capabilities resulting from the boundary conditions of the employed equipment [4]. In addition, the model captures material property parameters of the equipment and provides operational and behavioral information, including aspects of automation [5] [14]. Technical operating limits, such as maximum and minimum values for temperature, pressure, and viscosity, are also represented within the capabilities model [14].

Operation Model

The operation model addresses the execution and optimization of plant operation. It includes the modeling of recipes, as well as the monitoring and execution of operational procedures [15]. Furthermore, the model supports plant and PEA configuration and captures optimized process parameters to enable efficient and stable operation of the production system [4], [15].

SELECTED INFORMATION MODELS

The evaluated information models are selected based on the stated criteria. Specifically, the selected models:

- provide a detailed and formally defined specification,
- exhibit a level of abstraction that allows concrete allocation to aspects of the PPR model,
- support technical implementation through machine-readable or formally specified structures, and
- classes are excluded that serve only internal or structural purposes of the model (e.g., datatype definitions), as these do not contribute directly to fulfilling PPR submodel requirements.

For example, the Asset Administration Shell (AAS) submodel *Handover Documentation* is capable of describing common document types and can therefore be applied across multiple PPR submodels. In contrast, the AAS submodel *Carbon Footprint* models calculation methods in a descriptive form (e.g., as strings) and can thus be concretely allocated to a specific PPR submodel, such as the *Product* submodel.

DEXPI Plant & DEXPI Process

DEXPI aims to provide a general method for data exchange, interoperability, and data integration covering all lifecycle phases of (petro-)chemical plants, from functional specification to operation. It seeks to avoid format conversions and data loss across CAE system boundaries, enable efficient project handover, reduce barriers between different CAE systems, support long-term storage in CAE-independent formats, and facilitate the coexistence of multiple engineering systems within an organization.[16]

The DEXPI Specification includes the DEXPI Plant and the DEXPI Process Model. DEXPI Plant represents plant engineering information and associated graphical representations e.g. Piping and Instrumentation Diagrams (P&ID). DEXPI Process represents chemical processes and associated graphical representations e.g. Block Flow Diagrams (BFD).

AAS Submodels

The Carbon Footprint submodel template supports the exchange of an asset's carbon footprint along the value chain to improve interoperability between stakeholders involved in documenting, evaluating, or optimizing environmental impacts. While it may contribute to initiatives such as the Digital Product Passport or Product Environmental Footprint. [17]

The Digital Nameplate submodel template provides interoperable asset nameplate information for Asset Administration Shells by defining standardized properties, ideally aligned with dictionaries such as ECLASS and the IEC Common Data Dictionary. It establishes a

standardized digital nameplate structure for process industry and factory automation equipment, while allowing application-specific extensions through supplementary submodel templates. [18]

The Functional Safety submodel defines a structured engineering data model for safety-relevant devices to support functional safety and reliability assessments in the design of safety control systems and related applications. Its intended benefits include reducing data-mapping effort, optimizing information exchange with engineering tools, supporting safety-oriented product selection, and enabling globally accessible, consistent product data independent of country, language, or culture. [19]

Module Type Package

The Module Type Package (MTP) provides a vendor-independent description interface for the automation interface of a Process Equipment Assembly (PEA). It includes descriptions of the human–machine interface (HMI), available services, procedures, and configuration parameters, as well as IP addresses and OPC UA nodes and their relations to data objects defined in VDI/VDE/NAMUR 2658 [10].

OntoCape

OntoCAPE is a large-scale ontology for the domain of Computer-Aided Process Engineering (CAPE), represented in a formal, machine-interpretable ontology language. It captures consensual process-engineering knowledge in a generic and reusable form, enabling knowledge sharing across software systems and supporting advanced engineering applications such as simulation model management, design data integration, electronic procurement, and mathematical modeling. [20]

Metadata4Ing Ontology

One of the main goals of the Metadata4Ing ontology is to enable documentation of data and material provenance in research processes. It is based on a generalized process model centered on the *processing step* class, in which data and material objects are represented as outputs, while methods, tools, and other contextual information are linked via dedicated classes, allowing complex research processes to be described through compositions of processing steps. [21]

METHODOLOGY

To achieve a transparent evaluation of various technically implemented standards with respect to an abstract philosophy representing general knowledge distribution, a multi-step methodology is applied. First, a literature review is conducted to identify suitable information models. This review focuses on standards that (i) provide a formal technical specification, (ii) exhibit an active community or documented update history, and (iii) are sufficiently mature to support practical application.

To operationalize the PPR model, a second literature review is performed to identify requirements specific to the individual PPR submodels. Based on these requirements, a mapping matrix is constructed that relates each identified PPR submodel requirement to every specified class within the selected information models. This fine-grained level of detail enables a direct connection between the technical implementation of an information model and the abstract structure of the PPR model.

The generation and evaluation of the mapping matrices are performed automatically. The assessment of whether a given class contributes to a specific requirement is carried out manually using a binary (true/false) decision. To simplify the evaluation process, the assumption is made that if a superclass within an information model satisfies a requirement, all its' subordinate classes do so as well. The resulting matrices are then evaluated automatically. To account for differing numbers of classes across PPR submodels, the results are normalized accordingly. This normalization enables a comparative allocation of each standard to the respective PPR submodels.

The Python scripts used for generating and evaluating the matrices, as well as the completed matrices themselves, are provided in the accompanying GitHub repository https://github.com/p2o-lab/-ESCAPE-36-Domain-Mapping-Matrices-supplementary-information. Finally, the results are analyzed and discussed in detail.

DMM INTRODUCTION

To evaluate technically implemented classes against PPR submodel requirements, the Design Matrix Methodology (DMM) is adapted [22] [23]. To provide a structured overview of the suitability of the investigated information models for fulfilling PPR submodel requirements, a matrix-based evaluation is conducted for each combination of PPR submodel and information model.

In each matrix, the PPR submodel requirements are represented by the rows, while the classes of the respective information model form the columns. Each requirement–class combination is assigned a value of 0 if the class description does not contribute to fulfilling the requirement, or a value of 1 if it does. If no meaningful assignment is possible, the corresponding cell is left empty and excluded from the evaluation. This reflects the assumption that the implementation of a requirement does not necessitate the use of all available classes within an information model.

For each requirement, the assigned values are aggregated and normalized by the number of evaluated entries. This results in a representation in which, for each PPR submodel, a table is obtained where the rows

correspond to the requirements and the columns represent the normalized degree to which each information model satisfies the respective requirement. In the subsequent step, the mean value across all requirements of a PPR submodel is calculated. This yields a single value between 0 and 1 that describes the overall suitability of an information model to fulfill the requirements of that PPR submodel. A visualization is provided in Figure 2.

RESULTS

Figure 3 presents the outcome of the cross-dimensional coverage analysis as an overall heatmap. A higher value represents a better matching between the specification and the PPR-Submodel requirements. The rows correspond to the considered PPR-related concepts (Product, Process, Resource, Transformation, Capability, Operation), while the columns represent the analyzed standards and information models (OntoCAPE, AAS Carbon Footprint, DEXPI Process, MTP, DEXPI Plant, AAS Digital Nameplate, AAS Functional Safety, metadata ing classes). All values are normalized to the interval [0, 1] and express the relative degree to which a given model explicitly addresses a specific concept.

OntoCAPE achieves the highest overall coverage across all evaluated dimensions. In particular, strong scores are observed for Capability (1.000), as well as consistently high values for Product, Resource, and Operation (each 0.750) and Transformation (0.714). This indicates a broad and balanced semantic scope with respect to the PPR model.

The AAS Carbon Footprint submodel shows a pronounced focus on the Process dimension (1.000), with moderate coverage of Product and Transformation (0.500 each). All remaining dimensions exhibit negligible or zero coverage.

DEXPI Process demonstrates high alignment with Process (0.800) and Transformation (0.667). Product and Resource are only partially covered, while Capability and Operation are not explicitly represented.

The Module Type Package (MTP) model exhibits strong coverage for Capability (0.818) and Operation (0.750). Other PPR dimensions, including Product, Process, and Resource, are largely not addressed.

DEXPI Plant shows its highest relevance for Transformation (0.667) and Capability (0.600), with lower coverage for Operation (0.250). Product- and process-oriented aspects are only marginally represented.

The AAS Digital Nameplate submodel is strongly associated with Resource (1.000) and shows moderate relevance for Capability (0.400). All other dimensions show no explicit coverage.

The AAS Functional Safety submodel exhibits selective coverage, mainly for Capability (0.600) and Operation (0.250), while other PPR dimensions remain unaddressed.

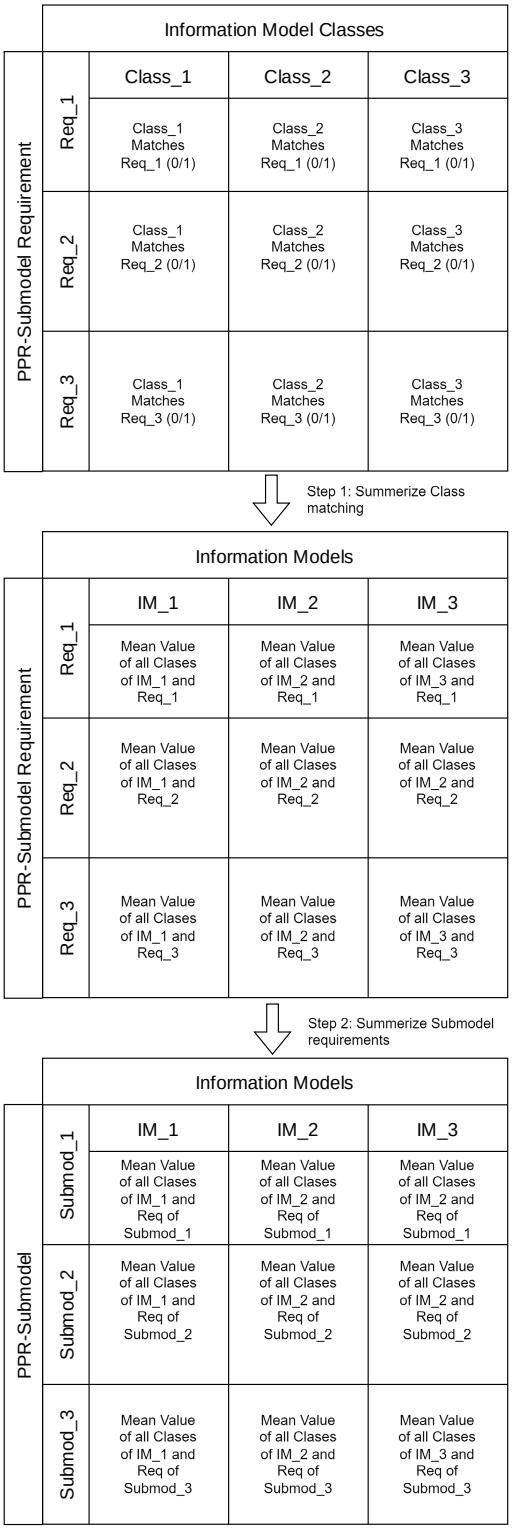

Figure 2. Description of the three step process to evaluate the information matching results. The corresponding data and code is available under https://github.com/p2o-lab/-ESCAPE-36-Domain-Mapping-Matrices-supplementary-information

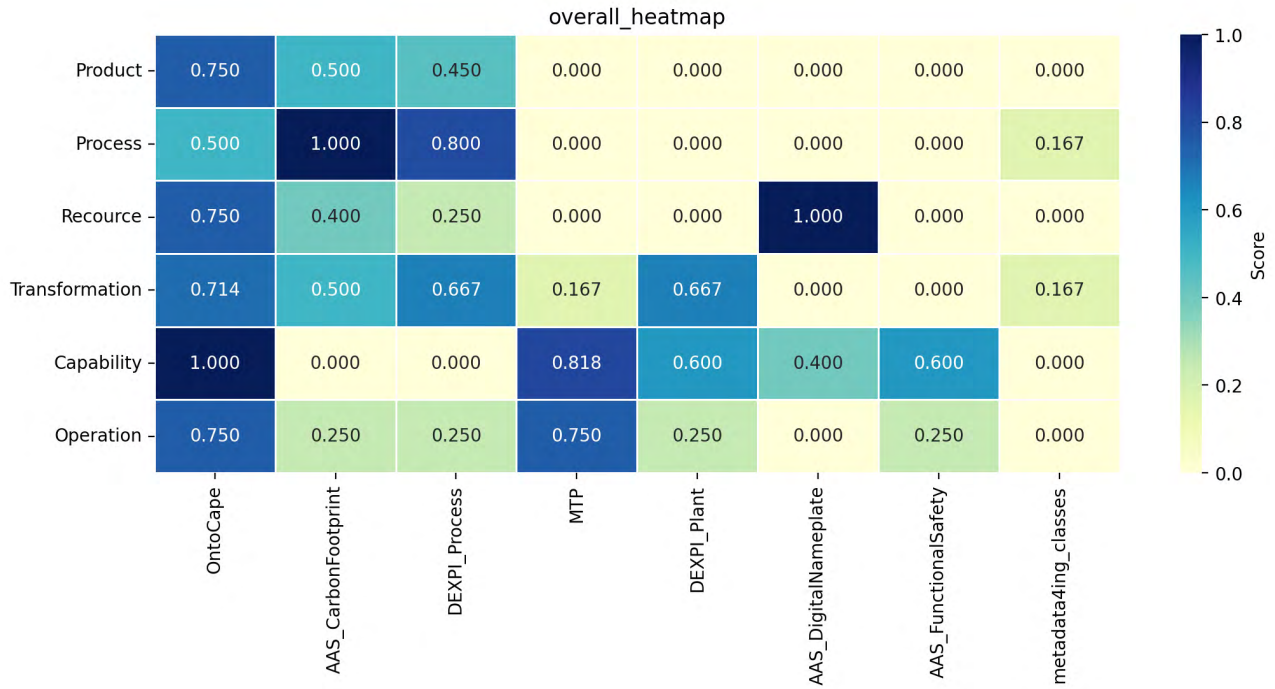

Figure 3. Heatmap of the cross dimensional analysis of the specified standards to the PPR-Model. A higher value represents a better matching between the specification and the PPR-Submodel requirements

Finally, the Metadata4Ing classes show only minimal coverage, primarily for Process and Transformation (0.167 each), indicating a limited contribution to explicit PPR modeling.

DISCUSSION & OUTLOOK

The results highlight a clear differentiation between conceptually comprehensive ontologies and practically oriented engineering standards. OntoCAPE stands out with the most complete and balanced coverage across all PPR dimensions, underscoring its strength as a reference ontology from a semantic modeling perspective. However, this result must be interpreted with caution: OntoCAPE has not undergone a major specification update since 2009 and, to date, lacks documented large-scale industrial or production-level deployments. Its high scores therefore reflect conceptual expressiveness rather than validated applicability in current modular plant engineering practices.

In contrast, standards such as DEXPI, MTP, and the various AAS submodels exhibit narrower but more focused coverage profiles. DEXPI Process and DEXPI Plant emphasize process execution and plant structure, respectively, while MTP concentrates on capabilities and operations required for modular automation and orchestration. The AAS submodels further describe asset-centric, safety-related, or sustainability-related concerns. Although none of these standards individually achieve comprehensive PPR coverage, they are actively maintained and have demonstrable relevance in industrial contexts.

The limited coverage of the Metadata4ing classes reinforce their role as generic framework and metadata structure rather than as domain-specific engineering model. They are well suited to complement other standards by providing traceability and documentation capabilities, but they cannot substitute detailed PPR semantics.

Overall, the analysis reveals a trade-off between semantic completeness and practical adoption. While OntoCAPE provides a rich conceptual foundation, its limited real-world uptake constrains its immediate applicability. Conversely, industrial standards offer pragmatic, implementation-ready solutions at the cost of reduced semantic breadth. These findings support the need for a combined modeling approach in modular plant engineering, in which actively used standards are combined for expressiveness, maintainability, and industrial relevance.

In the next steps, interfaces between the individual models can be evaluated according to their allocation within the PPR model, and the foundation for computer-aided tools suitable for the modular plant approach can be further substantiated.

ACKNOWLEDGEMENTS

Funded by the Federal Republic of Germany FKZ

03EI3115A Funding body: Federal Ministry for Economic Affairs and Energy based on a resolution of the German Bundestag

AUTHOR IDENTIFIERS

Author ORCIDs:
Kock T: 0009-0002-0419-5549

REFERENCES

1. Schindel, Polyakova, Harding, Weinhold, Stenger, Grünewald, Bramsiepe. General approach for technology and process equipment assembly (PEA) selection in process design. Chemical Engineering and Processing - Process Intensification 159:108223 (2021). https://doi.org/10.1016/j.cep.2020.108223
2. VDI 2776-1:2020, *Verfahrenstechnische Anlagen - Modulare Anlagen - Grundlagen und Planung modularer Anlagen*, Richtlinie 1, Berlin., Nov. 2020. Accessed: Jan. 17, 2022. [Online]. Available: https://www.vdi.de/richtlinien/details/vdi-2776-blatt-1-verfahrenstechnische-anlagen-modulare-anlagen-grundlagen-und-planung-modularer-anlagen
3. K. H. Weber, *Engineering verfahrenstechnischer Anlagen*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016. doi: 10.1007/978-3-662-52897-6.
4. Bamberg A, Urbas L, Bröcker S, Bortz M, Kockmann N. The digital twin – your ingenious companion for process engineering and smart production. Chem Eng & Technol 44:954-961 (2021). https://doi.org/10.1002/ceat.202000562
5. J. Mädler, I. Viedt, J. Lorenz, and L. Urbas, "Requirements to a digital twin-centered concept for smart manufacturing in modular plants considering distributed knowledge, " in *Computer Aided Chemical Engineering*, vol. 49, Y. Yamashita and M. Kano, Eds., in 14 International Symposium on Process Systems Engineering, vol. 49., Elsevier, 2022, pp. 1507–1512. doi: 10.1016/B978-0-323-85159-6.50251-7.
6. Schleipen M, Lüder A, Sauer O, Flatt H, Jasperneite J. Requirements and concept for plug-and-work. at - Automatisierungstechnik 63:801-820 (2015). https://doi.org/10.1515/auto-2015-0015
7. Pfrommer J, Stogl D, Aleksandrov K, Escaida Navarro S, Hein B, Beyerer J. Plug & produce by modelling skills and service-oriented orchestration of reconfigurable manufacturing systems. at - Automatisierungstechnik 63:790-800 (2015). https://doi.org/10.1515/auto-2014-1157
8. Koch A, Hamedi N, Furtner L, Kock T, Klose A, Mädler J. Standards for information models considering knowledge distribution in modular plants. 2023 IEEE 21st International Conference on Industrial Informatics (INDIN) :1-7 (2023). https://doi.org/10.1109/indin51400.2023.10218218
9. A. Koch, F. Kunkel, L. Theophile, V. De La Vega Hernandez, J. Madler, and L. Urbas, "Aiding Modular Plant Design by Linking Capability and Transformation Models, " in *Computer Aided Chemical Engineering*, vol. 53, Elsevier, 2024, pp. 583–588. doi: 10.1016/B978-0-443-28824-1.50098-3.
10. VDI/VDE/NAMUR 2658 Blatt 1 - Automatisierungstechnisches Engineering modularer Anlagen in der Prozessindustrie - Allgemeines Konzept und Schnittstellen. 2019. Accessed: Jan. 21, 2026. [Online]. Available: https://www.vdi.de/mitgliedschaft/vdi-richtlinien/details/vdivdenamur-2658-blatt-1-automatisierungstechnisches-engineering-modularer-anlagen-in-der-prozessindustrie-allgemeines-konzept-und-schnittstellen
11. Rahm J, Theißen M, Klose A, Soemers M, Temmen H, Schäfer C, Urbas L. Efficient automation engineering of modular process equipment assemblies using the digital twin. Chemie Ingenieur Technik 93:2081-2091 (2021). https://doi.org/10.1002/cite.202100100
12. Tolksdorf G, Cameron DB, Theißen M. DEXPI 2.0: synergistic integration of PFD and p&id in a unified digital model. Chemie Ingenieur Technik 97:1065-1069 (2025). https://doi.org/10.1002/cite.70009
13. Marquardt W, Morbach J, Wiesner A, Yang A. Overview on ontocape. RWTHedition :35-56 (2009). https://doi.org/10.1007/978-3-642-04655-1_3
14. A. Koch, J. Mädler, A. Bamberg, and L. Urbas, "Digital Twins for Scale-Up in Modular Plants: Requirements, Concept, and Roadmap, " in *Computer Aided Chemical Engineering*, A. C. Kokossis, M. C. Georgiadis, and E. Pistikopoulos, Eds., Elsevier, Jan. 2023, pp. 2063–2068. doi: 10.1016/B978-0-443-15274-0.50328-0.
15. Klose A, Merkelbach S, Menschner A, Hensel S, Heinze S, Bittorf L, Kockmann N, Schäfer C, Szmais S, Eckert M, Rüde T, Scherwietes T, da Silva Santos P, Stenger F, Holm T, Welscher W, Krink N, Schenk T, Stutz A, Maurmaier M, Stark K, Hoernicke M, Unland S, Erben S, Kessler F, Apitz F, Urbas L. Orchestration requirements for modular process plants in chemical and pharmaceutical industries. Chem Eng & Technol 42:2282-2291 (2019). https://doi.org/10.1002/ceat.201900298
16. "Specifications – DEXPI." Accessed: Jan. 27, 2026. [Online]. Available: https://dexpi.org/specifications/

17. admin-shell-io, “submodel-templates/published/Carbon Footprint/1/0 at main · admin-shell-io/submodel-templates, ” GitHub. Accessed: Jan. 27, 2026. [Online]. Available: https://github.com/admin-shell-io/submodel-templates/tree/main/published/Carbon%20Footprint/1/0
18. admin-shell-io, “submodel-templates/published/Digital nameplate/3/0/1 at main · admin-shell-io/submodel-templates, ” GitHub. Accessed: Jan. 27, 2026. [Online]. Available: https://github.com/admin-shell-io/submodel-templates/tree/main/published/Digital%20nameplate/3/0/1
19. admin-shell-io, “submodel-templates/published/Functional Safety/1/0 at main · admin-shell-io/submodel-templates, ” GitHub. Accessed: Jan. 27, 2026. [Online]. Available: https://github.com/admin-shell-io/submodel-templates/tree/main/published/Functional%20Safety/1/0
20. Morbach J, Wiesner A, Marquardt W. Ontocape—a (re)usable ontology for computer-aided process engineering. Computers & Chemical Engineering 33:1546-1556 (2009). https://doi.org/10.1016/j.compchemeng.2009.01.019
21. “NFDI4ING / metadata4ing / metadata4ing · GitLab, ” GitLab. Accessed: Jan. 27, 2026. [Online]. Available: https://git.rwth-aachen.de/nfdi4ing/metadata4ing/metadata4ing
22. Danilovic M, Browning TR. Managing complex product development projects with design structure matrices and domain mapping matrices. International Journal of Project Management 25:300-314 (2007). https://doi.org/10.1016/j.ijproman.2006.11.003
23. Browning TR. Applying the design structure matrix to system decomposition and integration problems: a review and new directions. IEEE Trans. Eng. Manage. 48:292-306 (2001). https://doi.org/10.1109/17.946528

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Framework for Flexible Start/Stop Operation of Electrified Chemical Processes

Samuel Mercer[a] and Michael Baldea[ab]*

[a] McKetta Department of Chemical Engineering, The University of Texas at Austin, Austin, TX, USA
[b] Oden Institute for Computational Science and Engineering, The University of Texas at Austin, Austin, TX, USA
* Corresponding Author: mbaldea@che.utexas.edu.

ABSTRACT

A flexible start-stop operating policy that involves full shut-down and start-up may be beneficial for electrified plants under certain grid conditions, such as dispatchable demand response. This paper introduces a multi-period Hamilton-Jacobi reachability framework to explore the space of state trajectories for plant shut-down and start-up. Shut-down is defined in terms of operations leading to a stand-by state with no material flows or energy inputs, and variables within safety constraints. Candidate stand-by states are identified by constructing backwards reachability tubes from the desired steady-state operating point. The candidate shut-down/stand-by state is partitioned in fast and slow regions. Admissible control input trajectories are determined for the fast region, from which the minimum time trajectory is selected as optimal for fast start-up. A proof-of-concept simulation using a reaction/separation/recycle plant is presented.

Keywords: Plant Start-up, Process Electrification, Hamilton-Jacobi Reachability, Optimal Control

INTRODUCTION

Electrified processing and, in particular, electrified process heating offer new opportunities for the precise delivery of process energy, enabling faster, smaller, safer, and more efficient processes [1]. Operational flexibility is critically important for electrified processes. With a growing share of electricity generation from renewable sources, characterized by inherent temporal variability, and rapidly increasing demand from large electricity users such as data centers, plants with large electricity loads should have the capability to modulate power consumption in response to grid conditions and (ideally) benefit economically from doing so [1-2]. Modulation typically consists of temporary decreases or increases in production rates when electricity prices are high or low, respectively, as in price-driven demand response; flexible start-stop operation extends this idea, potentially allowing plants to take further advantage of incentive-driven demand side management such as interruptible power [3].

The start-stop operation of a chemical plant is a transient mode of operation that uses start-up and shut-down procedures to temporarily interrupt and (eventually) restart continuous production [3, 4, 5, 6]. Shut-down involves transitioning the plant from the nominal operating state to a passive stand-by state without process (material and energy) flows (and likely without control input). Start-up consists of transitions from the stand-by state back to the nominal state under safety and operability constraints by restarting process flow and control [5, 6]. For continuous processes, such events are typically rare and restricted to maintenance or emergency situations. The loss of production over the start-stop period, high energy requirements to restart process units (e.g., staged units such as distillation columns) and long start-up durations are costly. Nevertheless, fast start-stop operations (where the cycle occurs over a time interval that is relevant to daily fluctuations in electricity prices) could have substantial benefits [1, 2].

Motivated by the above, the contribution of this work is to develop an analysis framework for start-stop operations based on Hamilton-Jacobi reachability. We are specifically interested in describing the "stopped/stand-by" state that meets the above conditions and results in the fastest startup and restoration of the nominal operating state. A case study on the start-up of a simple reactor/separation/recycle process illustrates

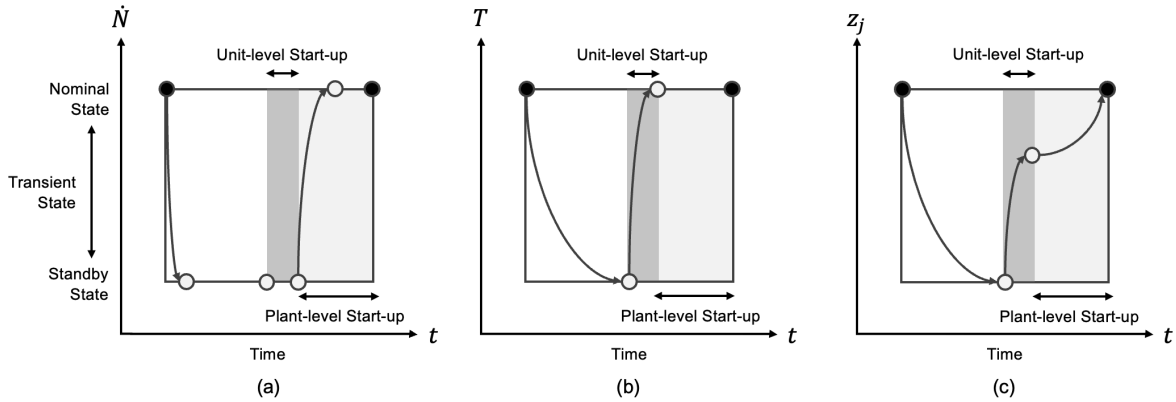

Figure 1. Schematic of the state trajectories for the (a) flow rate, (b) temperature, (c) composition for component *j* over a full shut-down and start-up cycle. The white region signifies the shut-down and stand-by period, the dark gray region the feed ramp-up start-up period, and the light gray region the composition control start-up period. The temperature is restored to its nominal operating state quickly followed by the flow rate, while the composition is the slowest.

the framework.

PLANT START-STOP OPERATIONS

Plant *start-stop* operations define a full shut-down to start-up cycle where the plant temporarily stops material and energy flow. To guide the discussion, the state trajectories of the flow, temperature, and composition variables are shown in Figure 1 for the full cycle. We define the following cycle periods:

Shut-down: starting from nominal steady-state operating conditions, the plant stops production by closing feed and product streams. The closed plant gradually shuts interconnections between individual units (or sub-systems).

Stand-by: individual units operate as autonomous closed systems without control action. The state derivatives are near zero and the holdup is at a near constant value. Control inputs are disabled and units are not interconnected. Material holdup is redistributed due to pressure and/or temperature gradients, resulting in changes in internal flow traffic such as draining in staged systems, pressure losses, or backmixing due to condensation. Energy is lost in the form of passive cooling from heat dissipation. The longer the standby period, the farther each sub-system drifts from the nominal steady-state operating point.

Unit-level Start-up: each unit starts up individually, similar to a batch operation, as their control inputs are reactivated. Their state progresses towards an intermediary (local) operating point that precedes plant-wide interconnection. This period ends when the slowest unit reaches its local operating point.

Plant-level Start-up: interconnections are re-established between local sub-systems to form the global system. Feed and product streams are re-introduced as well as feedback interconnections, such as recycle loops. The plant returns to the nominal steady-state driven by the manipulated variables, as dictated by a control system and/or operator inputs. The full shut-down-startup cycle is complete once the system returns to its nominal steady-state. It is assumed that the steps taken in the start-up and shut-down procedures are known (i.e., we are not attempting to devise new procedures).

CONCEPT: REACHABILITY FRAMEWORK

Efficient and timely start-stop operations depend on finding a (set of) plant state(s) that meet the description of the stand-by state, and for which feasible trajectories exist to and from the nominal operating state.

Hamilton-Jacobi (HJ) reachability analysis [8, 9, 10] is a powerful global method to explore all possible paths between two operating points, that satisfy system constraints. It comprises of formulating a forward or backward problem that derives feasible state trajectories over a finite time horizon, subject to an admissible control input trajectory. The forward problem maps all reachable final states from an initial state. The backward problem starts from the target set and maps the set of all initial states that reach the target, which is termed as the attraction basin. For start-stop operations, we formulate the backward reachability problem as the target operating state (i.e., the nominal state), which is known *a priori*.

On a related note, HJ backward reachability is closely related to minimum-time optimal control. There may exist an initial state and some admissible control

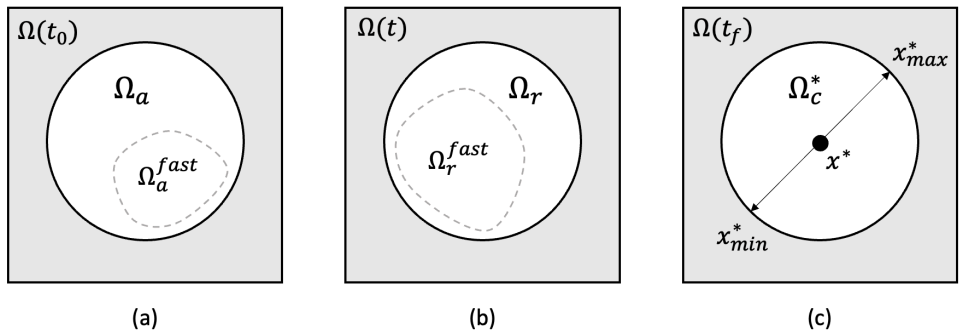

Figure 2. Plant abstract operating space Ω cross-sections for (a) the attraction basin Ω_a with fast region Ω_a^{fast} at initial time t_0, (b) the reachable set Ω_r with fast region Ω_r^{fast} at some intermediate time t , and (c) the capture basin Ω_c^* at terminal time t_f containing the target set x^*.

trajectory that reaches the target set the fastest. For plant-wide start-up, we aim to identify this path.

MATHEMATICAL FORMULATION

Plant Network Dynamics

Consider a plant with $i = 1, \dots, N_{units}$ process units, where each unit has state variables $x \in \mathbb{R}^m$ in the vector $\mathbf{x}$ and control inputs $u \in \mathbb{R}^p$ in the vector $\mathbf{u}$ given by:

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}, \mathbf{u}, t) \tag{1}$$

over the finite time horizon $t: [t_0, t_f]$, where t_0 and t_f represent the initial and final times, respectively. x and u have operability constraints $x_{min} \le x \le x_{max}$ and $u_{min} \le u(t) \le u_{max}$, which are valid for all time [11, 12]. The maximum allowable time horizon τ satisfies $\tau \ge t_f - t_0$. The state trajectory ϕ is defined as:

$$x(t) = \phi(t, t_0, x_0; u(t)) \tag{2}$$

while u has an admissible control trajectory ϕ^{u}:

$$u(t) = \phi^u(t, t_0, u_0; x(t)) \tag{3}$$

The plant network is a collection of process units with a set of state variables $\Omega = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_{units}}]^T \in \mathbb{R}^{N_{units} \times m}$ and set of control inputs $\Omega^{\mathrm{u}} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_{N_{units}}]^T \in \mathbb{R}^{N_{units} \times p}$:

$$\dot{\Omega}(t) = (I \bigotimes A + \gamma(t) \bigotimes \zeta(t))\Omega(t) + (I \bigotimes B)\Omega^{\mathrm{u}}(t) \tag{4}$$

where $A \in \mathbb{R}^{m \times m}$ is a state matrix, $\gamma(t) \in \mathbb{R}^{N_{units} \times N_{units}}$ is a graph Laplacian matrix defining the connectivity of the network topology, $\zeta(t) \in \mathbb{R}^{m \times m}$ is a coupling matrix between the state variables of connected process units, B is a control input matrix, and I is the identity matrix, all of appropriate dimensions.

Equivalently, the plant state trajectory $\Phi \in \mathbb{R}^{N_{units} \times m}$ is defined as:

$$\Omega(t) = \Phi(t, t_0, \Omega(t_0); \Omega^{\mathrm{u}}) \tag{5}$$

where $\Omega(t_0)$ is the set of initial states for all state variables. Ω^{u} has an admissible plant control trajectory Φ^{u}:

$$\Omega^u(t) = \Phi^u(t, t_0, \Omega^u(t_0); \Omega(t)) \tag{6}$$

Backward Reachability

Ω^* is a target steady-state operating region that consists of a subset of target state variables x^* subject to production and quality constraints $x^*_{min} \le x^* \le x^*_{max} \in \mathbb{R}^m$ (Figure 2c). The capture basin Ω_c^* is the region of state space enclosing x^* such that $x^* \in \Omega_c^*$. It is implied that Ω_c^* is a terminal condition that is reached at $t = t_f$ such that $\Omega_c^* \subset \Omega(t_f)$. It is also assumed that x^* exists on a stable equilibrium manifold.

A state $\Omega(t_0)$ is backward reachable from Ω_c^* if there exists an admissible control trajectory Φ^{u} such that the state trajectory Φ reaches Ω_c^* within the finite time horizon satisfying $\tau \ge t_f - t_0$. The attraction basin $\Omega_a(t_0, \Omega(t_0))$ (Figure 2a) consists of the set of initial states that are backward reachable.

Ω_a can be extended to become a backward reachability tube $\mathcal{B}$ (Figure 3), which is defined as the collection of Ω_a for some intermediate time on the interval $[t_0, t_f]$ that reaches x^*:

$$\mathcal{B} = \bigcup_{t \in [t_0, t_f]} \Omega_{\mathrm{a}}(t, \Omega^*) \times \{t\} \tag{7}$$

For each $\Omega_a \in \mathcal{B}$, the mapping from Ω_a to Ω_c^* contains a set of admissible state and control trajectories that span a range of reachable states $\Omega_r(t, t_0, \mathcal{B})$ (Figure 2b) within the operable state space. The reachable space Ω_r over τ is therefore:

$$\Omega_r(\Omega(t_0)) = \bigcup_{t \in [t_0, t_f]} \mathcal{B}(t, t_0) \tag{8}$$

which implies $\Omega_r \subseteq \Omega$.

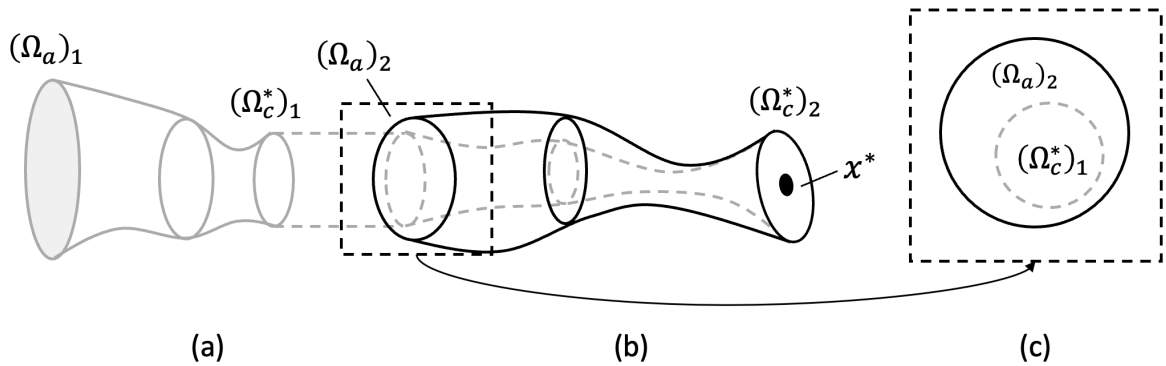

Figure 3. The backward reachability tube $\mathcal{B}$ for (a) $\mathcal{B}_1$ feed ramp-up start-up period (indexed by 1 and colored in gray solid lines) and (b) $\mathcal{B}_2$ composition control start-up period (indexed by 2 and black solid lines). The attraction and capture basins for each period are visible at an intermediary operating point, where the intersection between $(\Omega_a)_2$ and $(\Omega_c^*)_1$ (Equation 12), as shown in (c), forms a reachable multi-period tube (traced by the gray tube in solid lines for $\mathcal{B}_1$ and dashed lines for $\mathcal{B}_2$).

Minimum Time Optimal Control

The backward reachability tube (BRT) is computed using Hamilton-Jacobi reachability theory. The first step is to define an implicit level set function Ψ to define the proximity of the current state $x(t)$ from x^* for all $x^* \in \Omega^*$:

$$\Psi = \left| x - \frac{x^*_{min} + x^*_{max}}{2} \right| - \frac{x^*_{min} - x^*_{max}}{2} \tag{9}$$

Equation 9 suggests that if $\Psi \leq 0$, then x is within Ω_c^*. This implies that Ψ encodes $\mathcal{B}$; it is also equivalent to the value function V in the constrained Hamilton-Jacobi-Bellman partial differential equation (HJB-PDE, Equation 10a) with terminal condition (Equation 10b), which computes the manipulated input trajectory that reaches the target set the fastest:

$$\frac{\partial V}{\partial t}(t,x) + \min_u \left(\nabla V(x(t),t) \cdot f(x,u,t) \right) = -1 \tag{10a}$$

$$V(x(t_f), t_f) = \Psi(t_f) \tag{10b}$$

The optimal admissible control trajectory for each possible initial state $x(t_0)$ can be derived from Equation 10 for minimum time:

$$u^*(t,x) = \underset{u}{\operatorname{argmin}} \left(\nabla V(x(t),t) \cdot f(x,u,t) \right) \tag{11}$$

which is measured by the reachability time τ^* for each initial state given u^*, where $\tau^* \leq \tau$ and is defined as the time interval for a trajectory to reach Ω_c^* from some initial state. The u^* for each process unit constitutes the set Ω^u.

After V is computed, the BRT may be partitioned between fast and slow regions based on the reachability time τ^*. The smallest τ^* corresponds to the fastest start-up time. As such, the initial state $\Omega(t_0)$ lying within the fast region of Ω_a, denoted by Ω_a^{fast}, and its corresponding minimum-time control trajectory Ω^u with the fastest state trajectory Ω is equivalent to the minimum-time optimal control solution.

Multi-Period Reachability

The above approach for one time period can be extended to multiple time periods. Discrete behavior for start-stop operations, such as alterations to the process topology (e.g. turning on specific units or enabling/disabling streams), can be split into $k = 1, \ldots, N_{period}$ periods. The connection between periods are intermediary operating points that intersect the attraction basin of the kth period $(\Omega_a^*)_k$ and capture basin of the k-1 period $(\Omega_c^*)_{k-1}$, as visualized in Figure 3c.

Remark: the operating spaces in Figure 2 show one step in the backward reachability for each time period τ_k, as reflected in Figure 3.

$$(\Omega_c^*)_{k-1} \cap (\Omega_a^*)_k \tag{12}$$

Thus, a reachable multi-period tube assumes that the intersection in Equation 12 is non-empty, otherwise the multi-period tube does not exist.

Start-up Reachability

Plant start-up reachability is represented by the possible backward paths from the nominal state to the stand-by state. The following framework is proposed to map reachable and optimal trajectories:

(1) Identify the region of possible stand-by states.

(2) From the nominal steady state, identify each step in the operating sequence and their intermediary operating points based on where discrete behavior arises (Figure 3b). Examples are feed ramp-up or turning on control valves.

(3) Validate the connectivity between time periods based on Equation 12 (Figure 3c) to fully trace the $\mathcal{B}$ from the desired/target state. Work backwards until reaching the initial start-up state, which corresponds to the final stand-by state. The state derivatives are near zero and the holdup is at a constant value. Control inputs are disabled and units are not interconnected.

(4) Determine where the region of possible stand-by states in Step (1) intersect with $(\Omega_a)_1$, as depicted in in Figure 3a. This is the region in the state space that is the most favorable stand-by region for fast start/stop operation where there exists fast and slow regions.

(5) Identify the fast and slow regions based on the total time to reach steady-state from the grid of all feasible initial conditions (Figure 2a) for each time period. The full start-up time is given by the sum of τ_k for each period.

CASE STUDY

Consider a prototype reactor-recycle-separation system with purge (shown in Figure 4), similar to systems found in [13, 14]. Assume the following: the holdup is constant for all units (and, therefore, the collective plant holdup), the temperature is constant, and the energy balance is neglected. The authors recognize that this may not be fully representative of electrified processes, where it is likely that the largest use of electricity will be in the form of process heat. Nevertheless, the features highlighted here are broadly applicable.

The process has two components: A and B, where B is the desired product. The composition variables are subject to constitutive relations:

$$1 = x_{A1} + x_{B1} \tag{13}$$

$$1 = x_{A2} + x_{B2} \tag{14}$$

x_{A1} and x_{A2} are differential variables of component balances:

$$\frac{d}{dt}\begin{bmatrix} x_{A1} \\ x_{A2} \end{bmatrix} = \begin{bmatrix} -\frac{F_1}{N_1} - k_1 & \frac{\alpha_A R}{\beta N_1} \\ \frac{F_1}{N_2} & -\frac{F_2}{N_2} - \frac{\alpha_A F_3}{\beta N_2} \end{bmatrix} \begin{bmatrix} x_{A1} \\ x_{A2} \end{bmatrix} + \begin{bmatrix} \frac{F_0 x_{A0}}{N_1} \\ 0 \end{bmatrix} \tag{15}$$

The overall plant material balance is between the feed F_0, product F_2, and purge P streams

$$F_0 = F_2 + P \tag{16}$$

$$P = \lambda F_3 \tag{17}$$

where P is determined by the split ratio λ between the vapor stream F_3 and the recycle stream R

$$R = (1 - \lambda) F_3 \tag{18}$$

F_0 is pure A ($x_{A0} = 1$) and mixes with R (whose composition is x_{A3} and x_{B3}). The combined stream $F_0 + R$ enters a liquid phase reactor (indexed as Unit 1) where A converts to B in an irreversible forward reaction $A \rightarrow B$.

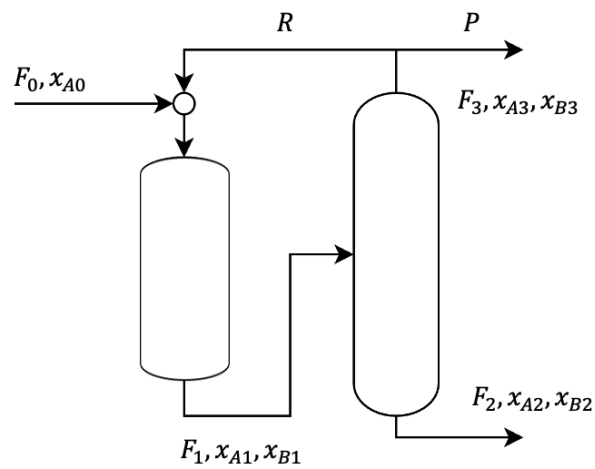

Figure 4. Chemical plant with a reactor, a flash drum, a recycle stream, and a purge stream.

$$\frac{d}{dt}(N_1) = 0 = F_0 + R - F_1 \tag{19}$$

The reactor outlet stream F_1 has composition x_{A1} and x_{B1}, and feeds into a single stage separator (indexed as Unit 2) that splits F_1 into a liquid product stream F_2, which concentrates the desired product B, and a vapor stream F_3 that concentrates the unreacted A.

$$\frac{d}{dt}(N_2) = 0 = F_1 - F_2 - F_3 \tag{20}$$

The flash calculation in the separator is computed using a linearized Rachford-Rice equation, where α_A and α_B represent the relative volatility of components A and B while β represents a phase transfer coefficient.

$$0 = \frac{x_{A2}}{(1-\alpha_A\beta^{-1})^{-1} - F_3 F_1^{-1}} + \frac{x_{B2}}{(1-\alpha_B\beta^{-1})^{-1} - F_3 F_1^{-1}} \tag{21}$$

$$\beta = \alpha_A x_{A2} + \alpha_B x_{B2} \tag{22a}$$

$$x_{A3} = \frac{\alpha_A}{\beta} \tag{22b}$$

$$x_{B3} = \frac{\alpha_B}{\beta} \tag{22c}$$

The plant start-up comprises of manipulating the split ratio λ to control the purity of the separator product stream during the feed ramp-up and composition control periods. The goal is to identify the stand-by region that allows startup, as well as the subregion for the fastest startup. The model parameters are available in Table 1.

Table 1: Model parameters.

Symbol	Parameter	Value
α_A	Relative volatility of A	10.5
α_B	Relative volatility of B	3
F_0^{nom}	Nominal feed flow rate	25 kmol/hr
N_1	Reactor holdup	50 kmol
N_2	Drum holdup	45 kmol
k_1	First-order rate constant	10 s^{-1}
x_{A1}	A mole fraction (reactor)	0.5
x_{A2}	B mole fraction (drum)	0.8
t_f	Time horizon	6 hr

Numerical Implementation

A dynamic programming approach is adopted to efficiently solve the finite horizon HJB equation. The state and time variables of Equation 10 are discretized where x_{A1} in $m = 0, \dots, N_{x_{A1}}$ grid points and x_{A2} in $n = 0, \dots, N_{x_{A2}}$ grid points. The new discrete variables are x_{A1}^m and x_{A2}^n. The control variable λ is discretized in $l = 0, \dots, N_\lambda$ grid points, to form the discrete control set λ^l. Piecewise-constant controls are used, updated on an interval of Δt_c.

The value function is defined as follows:

$$V_{nm}(t_f) = \begin{cases} 0, x \in \Omega_c^* \\ \infty, x \notin \Omega_c^* \end{cases} \tag{23}$$

There are two distinct time discretization schemes. An explicit backward Euler integration scheme for Δt steps is applied to discretize the process dynamics:

$$\frac{dx_{A1}}{dt} = \frac{x_{A1}^{t+1} - x_{A1}^t}{\Delta t} \tag{24}$$

$$\frac{dx_{A2}}{dt} = \frac{x_{A2}^{t+1} - x_{A2}^t}{\Delta t} \tag{25}$$

and a semi-Lagrangian backward propagation scheme of the HJB equation for Δt_c steps over the control trajectory [15].

$$V_{nm}^{t_c+1} = \min_{\lambda^l}\left[\Delta t_c + V_{nm}^{t_c}\right] \tag{26}$$

The full λ^l space is recursively searched for a time-optimal control λ_{nm}^{l*}:

$$\lambda_{nm}^{l*} = \operatorname*{argmin}_{\lambda^l}\left[\Delta t_c + V_{nm}^{t_c}\right] \tag{27}$$

The numerical settings are specified in Table 2. All computations were performed using Python 3.13.5 and NumPy 2.1.3. The runtime was 6 minutes.

Table 2: Numerical settings.

Symbol	Parameter	Value
$N_{x_{A1}}$	x_{A1} grid resolution	120
$N_{x_{A2}}$	x_{A2} grid resolution	120
Δt	Euler time step size	0.05
Δt_c	Control time step size	0.2

Results

The purity set-point is 0.95 mol% for x_{B2}. Figure 5 shows the shape of the BRT (as anticipated in Figure 3a and 3b) and the intersection of the tubes for each period, tracing a full feasible BRT for start-up. Figure 6 compresses the BRT onto a 2D grid for each period.

In Figures 6a-i and 6a-ii, the attraction basin expands as the time horizon increases, implying that more initial states are reachable for longer times. Considering that the holdup is constant, the feedback from the recycle stream shows how the initial holdup composition in the reactor and separator affects the start-up length. The boundary of the attraction basin expands to more initial x_{A2} states at higher initial x_{A1} for a given time. While this is intuitive, we find the converse perspective more useful: the state space region that allows for fast start-up becomes narrower as the desired start-up time gets shorter.

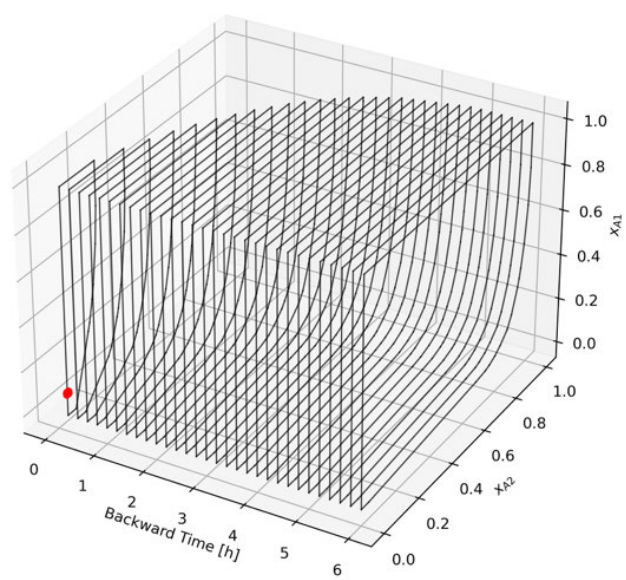

Figure 5. Backwards reachability tube for the start-up of the reactor-recycle-separation system. The red dot denotes the target set x^*.

Figures 6b-i and 6b-ii illustrate the optimal time-averaged split ratio for each initial state. Higher x_A initial reactor concentrations, which are far from the nominal state, taper towards $\lambda = 0$ and favor recycle-only operation. Purge-only operation $\lambda = 1$ is favored when x_{A1} is near the target operating state. Alternative split ratios are favored over the interval $x_{A1} = [0.2, 0.5]$.

Remark: In Figures 5 and 6, the discretization scheme is a coarse-grained approximation of the attraction basin and set of recycle ratios, which is why the profiles have streaks along the basin boundary. A

finer approximation with more grid points would generate a smoother profile but would require longer computational time.

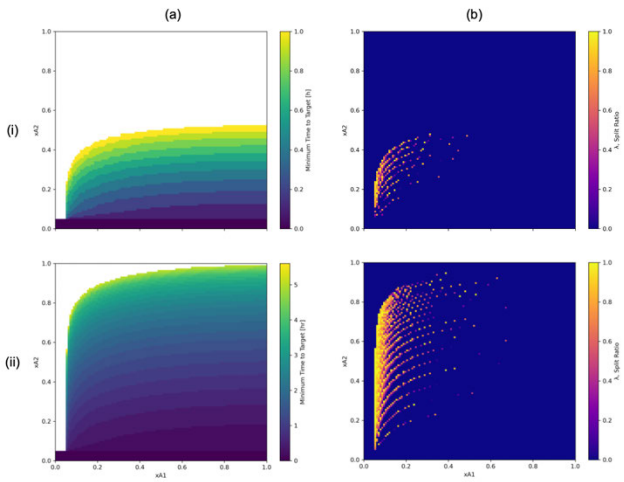

Figure 6. Set of (a) attraction basins and (b) optimal split ratios for (i) feed ramp-up period and (ii) composition control period.

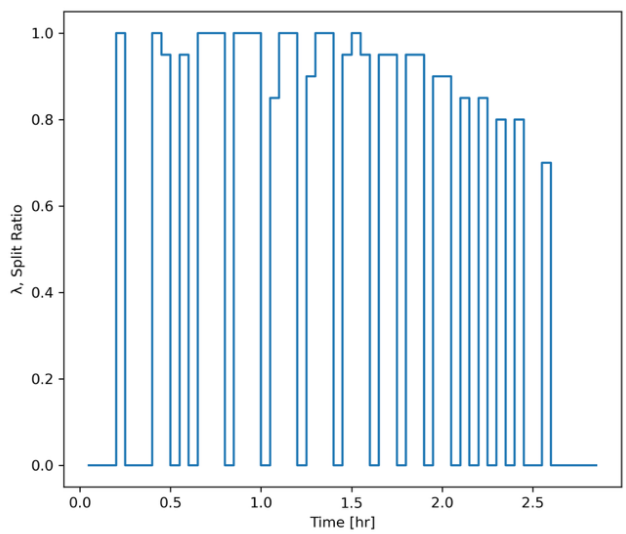

Figure 7. Optimal control trajectory of the split ratio during the composition control period.

From the results in Figures 5 and 6, let us consider an initial state of $[x_{A1,0} = 0.9, x_{A2,0} = 0.9]$ whose total start-up time is 2.6 hr. An optimal control trajectory for the split ratio is presented in Figure 7, where there is initially large variation until the system reaches the nominal steady state. The optimal trajectory is reflected in the black-colored trajectory on the phase portrait shown in Figure 8. Notably, state trajectories from all initial conditions converge to the same path to the capture basin. The proposed framework provides a full view of the feasible operating space under operability and finite time constraints. It is shown how the feasible state and control trajectories are systematically determined by a global exploration of the control space. Provided the shape of the BRT and its associated attraction basins, the initial state (which is, in fact, the stand-by state of the plant) has a substantial impact on the start-up speed. The ability to confine the process unit state to the fast region of an attraction basin during shut-down and stand-by has a great effect on start-up, suggesting that an interruptible operation is governed as much by its available control inputs as the ability to maintain the system state near the target conditions.

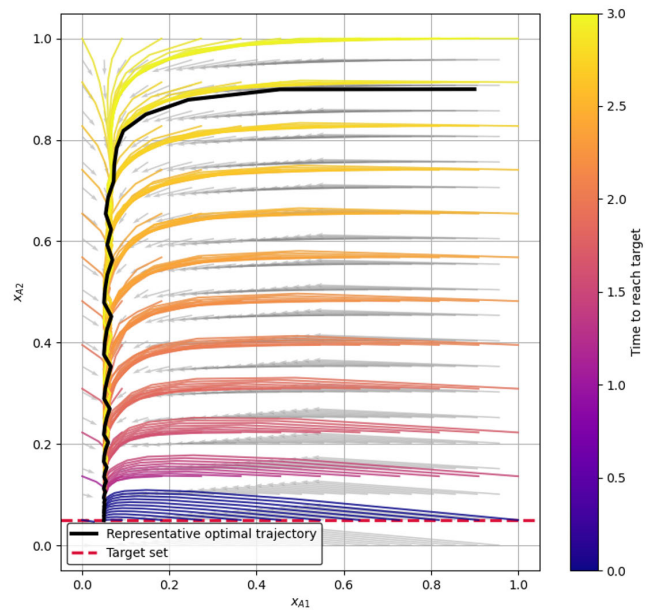

Figure 8. Phase portrait of control trajectories during composition control period. The black trajectory is an optimal path given the initial state (x_{A1} = 0.9, x_{A2} = 0.9) from stand-by.

CONCLUSIONS AND FUTURE WORK

In this paper, a conceptual framework for analyzing start-stop operations was developed based on Hamilton-Jacobi reachability. Attraction basins were found by deriving backward reachability tubes from a target set, from which the fast region of the attraction basin was identified. Within this region, the optimal admissible control trajectory corresponding to the fastest state evolution is found. A plant start-up case study for reactor/separation/recycle focusing solely on composition dynamics was demonstrated as a proof-of-concept. The next step is to validate the proposed framework on an industrial system with higher dimensionality.

ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation Graduate Research Fellowship Program

(GRFP) and the Provost's Graduate Excellence Fellowship at the University of Texas at Austin. This material is based upon work supported by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under the Industrial Technologies Office (ITO), Award Number DE-EE0010725. The views expressed herein do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

AUTHOR IDENTIFIERS

Author ORCIDs:
Mercer SP: 0000-0003-3224-5670
Baldea M: 0000-0001-2345-6789

REFERENCES

1. Baldea M, Endler EE, Hale E, Maravelias CT, Barolo M, Harjunkoski I, Mercangoz M, Shah SL, Soroush M, Young BR, Zhang Q. Transforming the process industries through electrification: challenges and opportunities. Ind. Eng. Chem. Res. 64:16466-16478 (2025). https://doi.org/10.1021/acs.iecr.5c01256
2. Harjunkoski I, Vaara S. Demand-side opportunities for process industry participation in electricity markets. Ind. Eng. Chem. Res. 64:24675-24688 (2025). https://doi.org/10.1021/acs.iecr.5c03604
3. Zhang Q, Morari MF, Grossmann IE, Sundaramoorthy A, Pinto JM. An adjustable robust optimization approach to scheduling of continuous industrial processes providing interruptible load. Computers & Chemical Engineering 86:106-119 (2016). https://doi.org/10.1016/j.compchemeng.2015.12.018
4. Chong, Z. Dynamic optimization formulations for plant operation under partial shutdown conditions. McMaster University PhD Dissertation (2012).
5. Durkin, A, Stolte, J, Mercangöz, M. Hoflon: hybrid offline learning and online optimization for process start-up and grade-transition control. *arXiv preprint* (2025).
6. Srinivasan R, Viswanathan P, Vedam H, Nochur A. A framework for managing transitions in chemical plants. Computers & Chemical Engineering 29:305-322 (2005). https://doi.org/10.1016/j.compchemeng.2004.09.024
7. Verwijs JW, van den Berg H, Westerterp KR. Startup strategy design and safeguarding of industrial adiabatic tubular reactor systems. AIChE Journal 42:503-515 (2004). https://doi.org/10.1002/aic.690420219
8. Kanavalau A, Lall S. Thermal runaway avoidance using hamilton–jacobi reachability and model predictive control. Computers & Chemical Engineering 157:107605 (2022). https://doi.org/10.1016/j.compchemeng.2021.107605
9. Kurzhanski AB, Varaiya P. Trajectory tubes state-constrained feedback control. Systems & Control: Foundations & Applications :311-339 (2014). https://doi.org/10.1007/978-3-319-10277-1_8
10. Yin H, Arcak M, Packard A, Seiler P. Backward reachability for polynomial systems on a finite horizon. IEEE Trans. Automat. Contr. 66:6025-6032 (2021). https://doi.org/10.1109/tac.2021.3056611
11. Dinh S, Lima FV. Dynamic operability analysis for process design and control of modular natural gas utilization systems. Ind. Eng. Chem. Res. 62:2052-2066 (2023). https://doi.org/10.1021/acs.iecr.2c03543
12. Dinh S, Lima FV. Linear dynamic operability analysis with state-space projection for the online construction of achievable output funnels. Computers & Chemical Engineering 205:109428 (2026). https://doi.org/10.1016/j.compchemeng.2025.109428
13. Falcone M, Ferretti R. Semi-lagrangian approximation schemes for linear and hamilton—jacobi equations. Society for Industrial and Applied Mathematics (2013). https://doi.org/10.1137/1.9781611973051
14. Baldea M, Daoutidis P. Dynamics and nonlinear control of integrated process systems. Cambridge University Press (2012). https://doi.org/10.1017/cbo9780511978760
15. Zhang J, Liu J. Distributed moving horizon state estimation for nonlinear systems with bounded uncertainties. Journal of Process Control 23:1281-1295 (2013). https://doi.org/10.1016/j.jprocont.2013.08.005

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PhoSim V.0 - Towards A Digital Twin for an Industrial Wet-Process Phosphoric Acid Production

Ilias Bouchkira[a], Sanae Elmisaoui[c], Abderrazak Latifi[bc*]

[a] RWTH Aachen University, Process Systems Engineering (AVT.SVT), Aachen 52074, Germany
b Université de Lorraine, LRGP, CNRS-ENSIC, Nancy 54000, France
c Mohammed VI Polytechnic University, Lot 990 Hay Moulay Rachid, Benguerir, Morocco
* Corresponding Author: abderrazak.latifi@univ-lorraine.fr

ABSTRACT

In this paper, PhoSim V.0, a physics-based digital twin core of a wet-process phosphoric acid (WPPA) digestion reactor, is presented as an initial step toward a high-fidelity virtual representation of industrial WPPA plants. The simulator integrates dissolution, reaction, and crystallization phenomena under strongly non-ideal electrolyte conditions. The reactor is modeled as a batch reactor, where the dissolution of tricalcium phosphate is described using a shrinking-core approach, while gypsum precipitation is represented by a one-dimensional population balance equation capturing particle size distribution evolution. Supersaturation ratios governing nucleation and crystal growth are computed from non-ideal activities using a Pitzer-based thermodynamic model, ensuring consistent coupling between dissolution, reaction stoichiometry, and crystallization kinetics. The resulting stiff system of equations is solved to predict key process indicators such as speciation, supersaturation ratio, and characteristic crystal size. Implemented within a user-friendly simulation environment, PhoSim V.0 provides a modular and extensible foundation that will be progressively refined to mimic industrial WPPA digestion units and supports future monitoring, optimization, and control-oriented applications.

Keywords: Digital twin, Wet-process phosphoric acid, Shrinking-core model for dissolution, Population balance modeling, Gypsum crystallization, Non-ideal thermodynamics

INTRODUCTION

Wet-process phosphoric acid (WPPA) production is a cornerstone of the phosphate and fertilizer industries, yet it remains a highly complex process involving the strong coupling between phosphate-rock dissolution, reactions kinetics, gypsum crystallization, and thermodynamics of highly non-ideal electrolytes. The digestion reactor, typically operated as a continuous stirred-tank reactor (CSTR) at the unit scale, plays a decisive role in determining phosphoric acid yield, gypsum crystal properties, fouling behavior, and downstream filtration performance.

Reliable modeling of this unit therefore represents a critical prerequisite for process analysis, optimization, and future digitalization efforts.

Over the past years, the authors of this paper have developed an integrated modeling framework addressing the main physico-chemical phenomena governing industrial WPPA digestion.

This body of work includes first-principles process modeling and multi-criteria optimization of industrial digestion tanks [1–6], rigorous thermodynamic modeling of sulfuric–phosphoric acid systems with impurities using Pitzer-type model [7], and thermodynamic-based analysis of fouling and crystallization phenomena in WPPA [8]. In parallel, detailed gypsum crystallization models based on population balance equations have been developed and applied to industrial digestion conditions [9-11], providing quantitative information on the evolution of particle size distribution (PSD) and its operational implications.

In addition to crystallization, phosphate rock dissolution has been investigated through shrinking-core approaches, with particular emphasis on the influence of particle size distribution and dissolution mechanisms under industrial conditions [9-12]. These dissolution models have been further complemented by surrogate and

reduced-order representations to support sensitivity analysis, parameter estimability, and model-based experimental design [13-15]. Together, these contributions form a unified and internally consistent modeling foundation for WPPA digestion, with consistent thermodynamic closure and kinetically grounded sub-models.

The present work builds exclusively upon these established developments and consolidates them into PhoSim V.0, a first-generation digital twin core of the WPPA digestion reactor. Rather than introducing new sub-models, this contribution focuses on their integration within a single simulation environment, enabling time-resolved prediction of speciation, supersaturation, nucleation and growth rates, and gypsum PSD indicators. This version intentionally targets the core digestion phenomena and serves as a foundational step toward a high-fidelity digital twin capable of mimicking industrial WPPA plants.

Future developments will extend this framework toward plant-scale integration, data-driven calibration, and control-oriented applications.

Process Description & Digital Twin Scope

Wet-Process Phosphoric Acid Digestion

Wet-process phosphoric acid production is based on the acidulation of phosphate rock by sulfuric acid in an agitated reactor, producing phosphoric acid and calcium sulfate dihydrate (gypsum) as the principal solid by-product. The digestion step is typically operated as a continuous stirred-tank reactor (CSTR) to ensure adequate mixing between the solid phosphate particles and the acidic liquid phase. For modeling purposes, the wet-process route is represented by a two-step reaction scheme. In the first step, tricalcium phosphate (TCP) is attacked by phosphoric acid to form soluble monocalcium phosphate (MCP). In the second step, the formed MCP reacts with sulfuric acid to regenerate phosphoric acid and precipitate gypsum (Figure 1).During digestion, several strongly coupled phenomena take place simultaneously: dissolution of phosphate rock, homogeneous and heterogeneous chemical reactions in the liquid phase, and nucleation and growth of gypsum crystals. The dissolution of TCP is often governed by diffusion and reaction resistances within the solid particle, which can be described using a shrinking-core model. As dissolution progresses, calcium and phosphate species are released into the liquid phase, modifying solution speciation and supersaturation conditions, which in turn control gypsum crystallization kinetics. The resulting gypsum particle size distribution (PSD) has a direct impact on downstream filtration performance and acid losses.
The digestion environment is characterized by high ionic strength and strongly non-ideal electrolyte behavior, making activity-based thermodynamic descriptions necessary for reliable prediction of solution equilibria and supersaturation. These characteristics motivate the use of rigorous thermodynamic and process models when representing the digestion reactor.

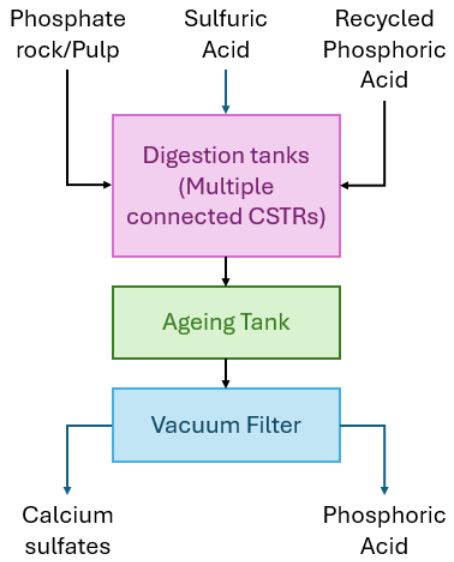

Figure 1. Simplified representation of the wet-process phosphoric process.

Digital Twin Architecture

The digital twin developed in this work targets the dynamic behavior of a single Wet Phosphoric Acid (WPPA) digestion reactor operated in batch mode. The scope is deliberately restricted to the core unit operation to establish a robust, transparent, and extensible modeling framework before addressing plant-wide integration or advanced control layers. The reactor is assumed to be perfectly mixed; such that spatial gradients are neglected, while temporal variations in composition, supersaturation, and crystal population are resolved.

Figure 2 presents a high-level and simplified architecture of the digital twin, intended to illustrate the conceptual structure and data flow of the simulation code rather than its numerical or software-level implementation. At this level, the model integrates three coupled components: a shrinking-core dissolution model, a thermodynamic calculation module, and a population balance equation (PBE) for gypsum crystallization.

The simulation starts with an Inputs and Initialization block that specifies operating conditions and initial states, including the initial amounts of TCP, sulfuric acid, phosphoric acid, water, gypsum seed crystals, and temperature. These inputs define the solution composition, the initial gypsum particle size distribution (PSD), and the thermodynamic database required for activity-based calculations.

The dissolution model is based on a shrinking-core formulation and describes the progressive dissolution of

solid TCP into soluble calcium and phosphate species. At each time step, the module updates the dissolved fraction of TCP, thereby governing the availability of Monocalcium phosphate (MCP) for subsequent gypsum crystallization. The dissolution rate is influenced by the solution state, whereas feedback from crystallization on dissolution is neglected.

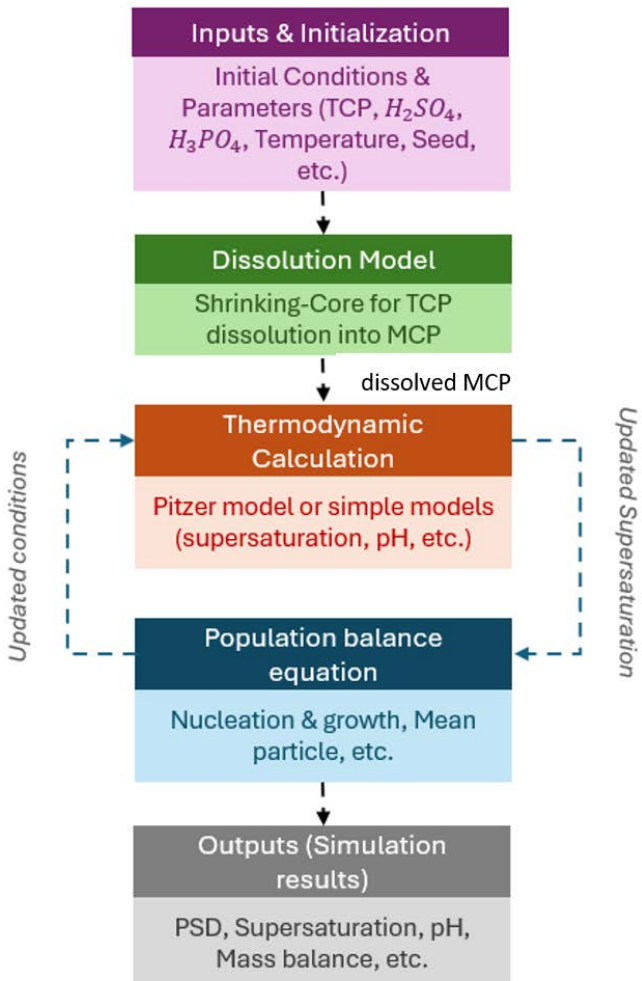

Figure 2. Schematic representation of the digital twin core architecture, showing the coupling between TCP dissolution, thermodynamic calculations, and gypsum crystallization via a population balance model.

The thermodynamic calculation module receives the updated solution composition and computes non-ideal solution properties, including ionic speciation, pH, and gypsum supersaturation ratio. Depending on the simulation settings, this module may rely on a rigorous Pitzer-based model or on simplified thermodynamic models. The resulting supersaturation ratio constitutes the key driving force for crystallization kinetics and is communicated to the PBE module.

The PBE describes gypsum nucleation and growth as a function of supersaturation. Solving the PBE yields the temporal evolution of the crystal size distribution (CSD), its moments (e.g. mean particle size), and the cumulative mass of precipitated gypsum. The crystallization of gypsum consumes calcium and sulfate from the liquid solution; this consumption is fed back to the thermodynamic modules to update the operating conditions.

Finally, the output block gathers both directly observable quantities and internal process states. Typical outputs include PSD metrics, supersaturation, nucleation and growth rates, speciation, and cumulative gypsum production. The digital twin core is implemented in a user-friendly simulation environment, allowing systematic variation of operating conditions, initial states, and kinetic parameters, which makes it well suited for scenario analysis, sensitivity studies, and future extensions toward optimization and control.

MODEL FORMULATION

Dissolution Reaction

The dissolution of TCP constitutes the first and governing step of the wet-process phosphoric acid route, as it directly determines the amount of calcium released into the liquid phase and thus available for subsequent gypsum precipitation and phosphoric acid formation. The digestion reaction considered is the acid attack of TCP by phosphoric acid, leading to the formation of MCP according to the following reaction:

$$\mathrm{Ca_3(PO_4)_2 + 4H_3PO_4 \rightarrow 3Ca(H_2PO_4)_2} \quad (1)$$

The dissolution kinetics of TCP is described using a shrinking core model, in which the solid particle progressively shrinks while maintaining a spherical shape. The model resolves mass transfer and reaction across three regions: the solid core, the surrounding liquid film, and the well-mixed liquid bulk [10-14]. The dissolution rate is expressed in terms of the temporal decrease in particle radius (or equivalently particle mass), written as:

$$\frac{dX}{dt} = \frac{3M_{\mathrm{TCP}}D_M}{x_{\mathrm{TCP}}R_0\delta\rho_s}(1-X)^{2/3}\left(C_M|_{r=R} - C'_M\right) \quad (2)$$

with $X(0) = 0, R_0$ the initial particle radius, X the TCP conversion, D_M the diffusion coefficient, δ the liquid film thickness, C_M and C'_M the concentration of MCP at the particle surface and in the bulk, respectively. For a spherical shrinking particle, the conversion $X$is related to the instantaneous particle radius R through the shrinking-core relation:

$$R = R_0(1-X)^{\frac{1}{3}} \quad (3)$$

Diffusion of species i (phosphoric acid or MCP) through the liquid film is described by the 2nd Fick's law in spherical coordinates:

$$\frac{\partial C_i}{\partial t} = \frac{1}{r^2}\frac{\partial}{\partial r}\left(D_i r^2 \frac{\partial C_i}{\partial r}\right) \quad (4)$$

with $C_i(0) = C_i^0$ and boundary conditions at the solid surface ($r = R$):

$$-D_i \frac{\partial C_i}{\partial r}\Big|_{r=R} = \nu_i k_r C_{H_3PO_4}\big|_{r=R} \quad (5)$$

and at the film–bulk interface ($r = R + \delta$):

$$C_i|_{r=R+\delta} = C_i' \quad (6)$$

The liquid-film thickness is linked to hydrodynamic conditions through a mass-transfer correlation. In our previous papers, the film thickness was derived from the Sherwood correlation, and then combined with film theory to express the evolution of δ as follows [10-12]:

$$\delta = R\left[(1+\alpha)\left(\frac{R}{R_0}\right)^{8/9} D_M^{-1/3} - 1\right] \quad (7)$$

Finally, liquid bulk concentrations evolve according to overall mass balances, which couple TCP dissolution to calcium availability in solution. The solution of this system provides the time-dependent TCP conversion and, equivalently, the amount of calcium released into the liquid phase. This calcium dissolution constitutes the direct input for gypsum precipitation and phosphoric acid formation in the subsequent modeling stages.

Population Balance Model for Gypsum Crystallization

For a population of crystals undergoing temporal evolution, the population balance framework describes the number density function n as a function of time t and a characteristic crystal size x. In addition to the dissolving phosphate particles, a second solid phase is considered and corresponds to precipitated gypsum, formed through the following reaction between aqueous calcium and sulfate species to produce calcium sulfate dihydrate.

$$\mathrm{Ca(H_2PO_4)_2 + H_2SO_4 + 2H_2O \rightarrow CaSO_4.2H_2O + 2H_3PO_4} \quad (8)$$

Assuming an isothermal, perfectly mixed batch reactor with a constant volume, and size-independent crystal growth, the population balance equation (PBE) used within the digital twin core is written as:

$$\frac{\partial n(t,x)}{\partial t} = -G(t)\frac{\partial n(t,x)}{\partial x} + R_N(t)\delta(x - x_0) \quad (9)$$

where R_N denotes the total nucleation rate, accounting for both primary and secondary nucleation mechanisms, G represents the crystal growth rate, and δ is the Dirac delta function used to introduce nuclei at a critical size. The expressions for the nucleation and crystal growth rates are given by the following relations:

$$R_N(t) = A.\exp\left(\frac{-B}{\ln\ S(t)^2}\right) + k_s.[\ln S(t)]^s.\mu_3(t) \quad (10)$$

$$G(t) = k_g[\ln\ S(t)]^g \quad (11)$$

where A, B, k_s, s, k_g, and g are unknown kinetic parameters to be identified from experimental data, S denotes the supersaturation ratio, primarily governed by the speciation of the reacting medium, and μ_3 is the third-order moment of the crystal number density function n. To solve the population balance equation efficiently, the method of moments is employed as a model reduction technique. This approach converts the original partial differential equation into a system of ordinary differential equations (ODEs) whose state variables are independent of the size coordinate x. The method relies on the definition and computation of the kth-order moment μ_k, which is expressed as follows:

$$\mu_k(t) = \int_0^\infty x^k n(t,x)dx \quad (12)$$

By multiplying Eq. (9) by $x^k$and integrating over the entire size domain, from zero to infinity, a closed system of ordinary differential equations governing the temporal evolution of the moments is obtained:

$$\frac{d\mu_0(t)}{dt} = R_N(t) \quad (13)$$

$$\frac{d\mu_k(t)}{dt} = kG(t)\mu_{k-1}(t), k \geq 1 \quad (14)$$

$$\mu_k(0) = \int_0^\infty x^k n_0(x)dx \quad (15)$$

where n_0 denotes the CSD at $t = 0$. The third moment of the crystal size distribution can be used to quantify the amount of gypsum formed during the process. Specifically, the cumulative gypsum production is expressed as:

$$C_{\Delta\text{gyp}}(t) = \frac{k_v \rho_s}{M_c}(\mu_3(t) - \mu_3(0)) \quad (16)$$

where $k_v$is the volume shape factor, $\rho_s$is the crystal density, and M_c is the molecular weight of the precipitated solid. Based on this definition, the overall material balance for the batch reactor can be written as follows:

$$C_i(t) = C_{i,0} + \nu_i C_{\Delta\text{gyp}}(t) \quad (17)$$

where $C_{i,0}$ denotes the initial concentration of component i, and $\nu_i$is its stoichiometric coefficient in the reactions from Eq. (1) and Eq. (7).

Thermodynamic Framework and Supersaturation Calculation

Solving the material balance subsequently enables the calculation of the supersaturation ratio, which is expressed as follows:

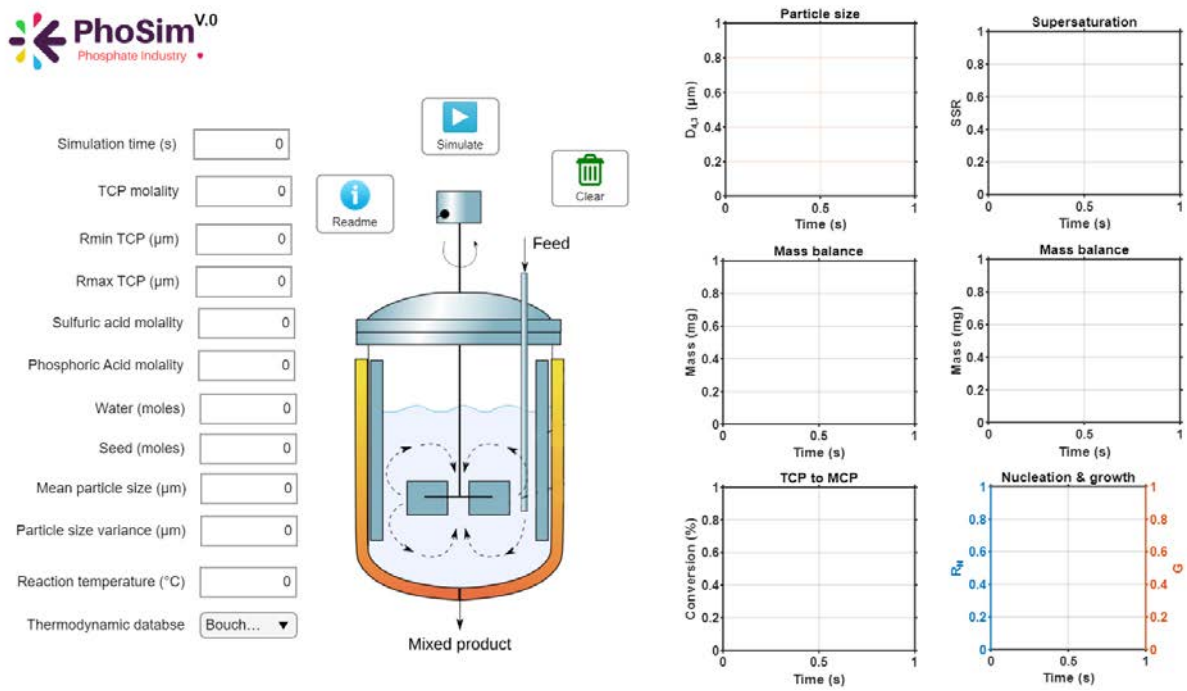

Figure 3. Graphical user interface of the PhoSim V.0 digital twin for batch wet-process phosphoric acid digestion and gypsum crystallization, showing input parameters, reactor schematic, and real-time simulation outputs.

$$S = \left(\frac{a_{Ca^{2+}}\, a_{SO_4^{2-}} a_w^2}{K_{sp}}\right); a_i = m_i \cdot \gamma_i \tag{18}$$

where a_i, γ_i, and m_i denote the activity, activity coefficient, and molality of species i, respectively, and $K_{sp}$is the solubility product of gypsum. The activity coefficients are evaluated using the Pitzer model, as described below [7]:

$$\ln(\gamma_i) = \frac{z_i^2}{2} f + 2\sum_j m_j \lambda_{i,j}(I) + z_i^2 \sum_j m_j m_k \lambda'_{j,k}(I) + \cdots \tag{19}$$

where z_i denotes the electrical charge of species i, f is the Debye–Hückel term, and I is the ionic strength of the solution. The parameters $\lambda_{i,j}$ and $\psi_{i,j,k}$ correspond to short-range binary and ternary interaction coefficients, respectively, whose values are identified from experimental equilibrium measurements.

DIGITAL TWIN CORE

Figure 3 illustrates the graphical user interface (GUI) of PhoSim V.0, the digital twin core developed in this work for the simulation of the wet-process phosphoric acid (WPPA) digestion and gypsum crystallization in a batch mode. The interface is designed to ensure transparency between user-defined operating conditions, underlying physicochemical models, and the resulting simulation outputs.

The GUI is organized in three main functional areas: input and configuration, process visualization, and output monitoring.

Input and configuration panel

The left-hand panel gathers all parameters required to initialize and configure a simulation run.

Simulation time defines the total integration horizon of the dynamic model and sets the temporal window over which all state variables are computed and displayed. TCP specifies the initial TCP amount introduced in the reactor. This parameter directly governs calcium availability for phosphoric acid production and gypsum precipitation. Minimum TCP radius (Rmin TCP) and maximum TCP radius (Rmax TCP) define the size bounds associated with the initial TCP particle population, which condition the dissolution kinetics through the shrinking-core representation. Sulfuric acid sets the initial sulfuric acid charge, controlling the extent and rate of TCP digestion and strongly influencing solution supersaturation with respect to calcium sulfate. Phosphoric acid represents the initial phosphoric acid concentration in the liquid phase and affects solution composition and thermodynamic speciation. Water defines the dilution level of the reacting medium and impacts concentrations, activities, and supersaturation behavior. calcium sulfate dihydrate (gypsum) corresponds to an initial gypsum mass, representing either pre-existing solids or deliberate seeding, which

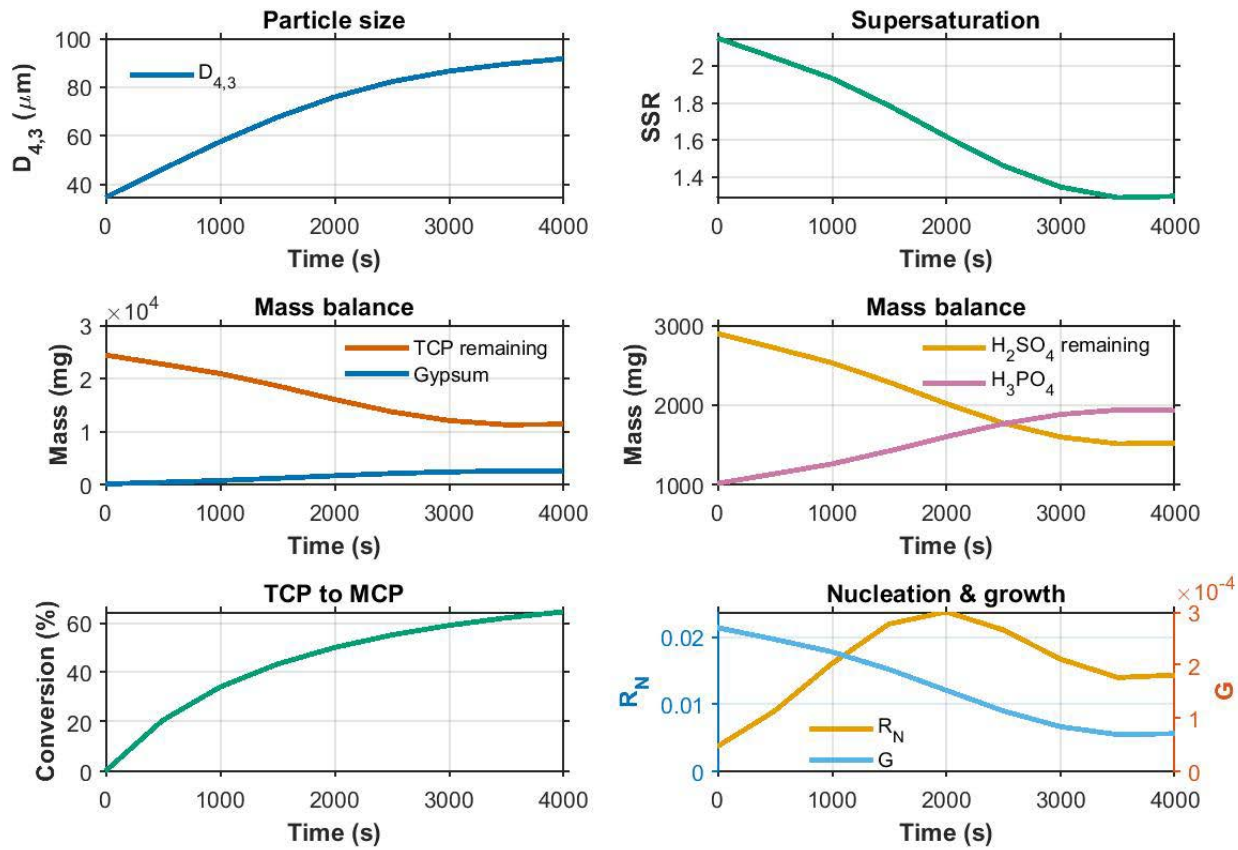

Figure 4. Dynamic simulation results for the case study: evolution of gypsum mean particle size D_{43}, supersaturation ratio (SSR), TCP and gypsum mass balances, sulfuric and phosphoric acid mass balances, TCP conversion, and nucleation and growth rates over the 4000 s batch operation.

can significantly influence nucleation dynamics and crystal size evolution. The mean particle size and gypsum particle size variance are used to define the initial gypsum particle size distribution (PSD), providing control over the initial crystal surface area available for growth. Reaction temperature sets the operating temperature used consistently across kinetic, dissolution, and thermodynamic sub models. Thermodynamic database selector allows the user to choose the electrolyte property database used for activity coefficients, speciation, and saturation index calculations, ensuring flexibility and extensibility of the digital twin.

Simulation control and information access

Run (Simulate) button launches the simulation by initializing the digital twin, solving the coupled mass balance, dissolution, and population balance equations, and updating all output plots. Readme / information button opens a compact pop-up window providing a concise description of the model scope, key assumptions, bibliographic references, and author contact information, supporting traceability and reproducibility. Clear button removes all previously plotted curves while preserving axis limits, labels, and formatting, enabling direct visual comparison between successive simulation scenarios.

Digital twin core graphical interface

The right-hand side of the interface contains six synchronized plots that display key model outputs in real time. Particle size evolution shows the time evolution of a representative volume-based mean particle size (e.g., D_{43}), summarizing gypsum crystal growth dynamics. The supersaturation (SSR) panel displays the supersaturation ratio as a function of time, representing the thermodynamic driving force for nucleation and crystal growth. Additional plots report mass balances for TCP and gypsum, sulfuric- and phosphoric-acid consumption and production, TCP conversion, and the nucleation and growth rates using a dual-axis representation. Together, these dashboards provide a comprehensive and physically interpretable overview of reaction progress, crystallization mechanisms, and particle population dynamics. The GUI enables rapid scenario testing while maintaining a clear correspondence between user inputs, model structure, and process-relevant outputs, which is a key requirement for digital twins intended for analysis, optimization, and

future control-oriented applications.

CASE TEST DEMONSTRATION

To illustrate the capabilities of the proposed digital twin core, a representative case study is conducted using a batch digestion–crystallization scenario typical of the WPPA route. The simulation conditions are summarized below and correspond to a demonstrative, non-optimized operating point.

The batch reactor is initialized with 0.08 mol/kgw of TCP, 0.03 mol/kgw of sulfuric acid, 0.01 mol/kgw of phosphoric acid, and 40 mol of water, together with an initial gypsum amount of 5×10^{-4} mol/kgw, representing either residual solids or deliberate seeding. The reactor temperature is fixed at 80 °C, and thermodynamic properties (speciation, activity coefficients, and saturation indices) are evaluated using the Bouchkira_et_al_Database.dat electrolyte database. The simulation time horizon is set to 4000 s.

The solid phase is described using a population balance framework. TCP particle sizes are bounded between R_min = 5×10^{-4} m and R_max = 1×10^{-3} m, while the initial gypsum particle size distribution is characterized by a mean size parameter $\bar{x}$ = 30 μm and a variance σ = 70 μm.

Figure 4 presents the main dynamic outputs of the digital twin core. The evolution of the volume-based mean particle size D_{43} shows a continuous increase throughout the batch, indicating growth-dominated gypsum crystallization. This behavior is consistent with the progressive decrease in supersaturation ratio (SSR), which initially exceeds unity and gradually relaxes as calcium sulfate precipitates and the system approaches equilibrium.

Mass balance results confirm the internal consistency of the model. TCP mass decreases monotonically due to acid digestion, while gypsum mass increases as a result of calcium sulfate precipitation. In parallel, sulfuric acid is consumed and phosphoric acid is produced, in agreement with the stoichiometry of the digestion reaction. The TCP conversion increases steadily and reaches approximately 65% at the end of the simulation, highlighting the combined effect of reaction kinetics and evolving thermodynamic driving forces.

The nucleation and growth kinetics further clarify the crystallization mechanism. The nucleation rate exhibits an initial rise associated with high supersaturation, followed by a decline as supersaturation decreases. Conversely, the growth rate remains significant over the entire batch, confirming that crystal growth dominates over nucleation kinetics.

This case test demonstrates the ability of the proposed digital twin core to coherently capture the coupled dynamics of reaction, thermodynamics, and particle population evolution, while providing physically interpretable outputs suitable for process analysis, sensitivity studies, and future optimization or control applications.

CONCLUSIONS AND PERSPECTIVES

A physics-based digital twin core (PhoSim V.0) for WPPA digestion and gypsum crystallization has been developed and demonstrated. The framework consistently couples reaction kinetics, thermodynamic equilibrium, and population balance modeling within a dynamic simulation environment supported by a dedicated graphical interface.

The case study confirms that the model captures the main process trends in a physically coherent manner, including TCP digestion, acid consumption and production, supersaturation relaxation, gypsum precipitation, and crystal size growth. The simulated transition from nucleation-dominated to growth-dominated crystallization regimes and the associated mass balance closure illustrate the internal consistency and interpretability of the approach.

The proposed digital twin core provides a flexible tool for process analysis and scenario testing, enabling systematic investigation of the impact of operating conditions and particle properties on both macroscopic performance indicators and crystallization mechanisms.

Future work will focus on experimental validation, extension to continuous reactor configurations, integration of inline measurements, and coupling with optimization and model-based control strategies, with the objective of supporting advanced monitoring and decision-making in industrial phosphoric acid production.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the financial support of OCP Group and Mohammed VI Polytechnic University (UM6P) over the period 2017–2023, during which most of the models presented in this work were developed.

REFERENCES

1. Bouchkira I, Latifi AM, Khamar L, Benjelloun S. Modeling and multi-objective optimization of the digestion tank of an industrial process for manufacturing phosphoric acid by wet process. Computers & Chemical Engineering 156:107536 (2022). https://doi.org/10.1016/j.compchemeng.2021.107536
2. Bouchkira I, Latifi AM, Khamar L, Benjelloun S. Multi-objective optimization of the digestion tank of an industrial phosphoric acid manufacturing

process. 2020 6th IEEE Congress on Information Science and Technology (CiSt) :389-394 (2020). https://doi.org/10.1109/cist49399.2021.9357206

3. Bouchkira, I., Latifi, A. M., Khamar, L., & Benjelloun, S. (2021). Process modeling and multi-criteria optimization of an industrial phosphoric acid wet-process. In *Computer Aided Chemical Engineering* (Vol. 50, pp. 499-504). Elsevier.
4. Bouchkira, I., Latifi, A. M., Khamar, L., & Benjelloun, S. (2022). Modeling of the crystallization of gypsum produced in the digestion tank of an industrial phosphoric acid manufacturing process. In *Computer Aided Chemical Engineering* (Vol. 51, pp. 445-450). Elsevier.
5. Bouchkira, I. (2022). Modélisation thermodynamique des solutions d'acides sulfurique et phosphorique en présence du minerai de phosphate: applications à l'encrassement, à la cristallisation et à l'optimisation multicritère d'une unité industrielle de production d'acide phosphorique. *Université de Lorraine*.
6. Bouchkira, I., Latifi, A. M., Khamar, L., & Benjelloun, S. (2021). Process modeling and multi-criteria optimization of an industrial phosphoric acid wet-process. In *Computer Aided Chemical Engineering* (Vol. 50, pp. 499-504). Elsevier.
7. Bouchkira I, Benjelloun S, Khamar L, Latifi AM. Thermodynamic modeling and parameter estimability analysis of a wet phosphoric acid process with impurities. Fluid Phase Equilibria 564:113594 (2023). https://doi.org/10.1016/j.fluid.2022.113594
8. Bouchkira, I., Benjelloun, S., Khamar, L., & Latifi, A. M. (2021). Thermodynamic-based model for the prediction of the fouling phenomena in a wet phosphoric acid process. *Chemical Engineering Transactions*, *86*, 1273-1278.
9. Elmisaoui, S., Latifi, A. M., Khamar, L., & Salouhi, M. (2021). Shrinking core approach in the modelling and simulation of phosphate ore acidulation. *Chemical Engineering Transactions*, *86*, 871-876.
10. Elmisaoui S, Latifi AM, Khamar L. Analysis of the dissolution of phosphate ore particles in phosphoric acid: influence of particle size distribution. Hydrometallurgy 223:106197 (2024). https://doi.org/10.1016/j.hydromet.2023.106197
11. Elmisaoui, S., Latifi, A. M., Farghi, F., Elmisaoui, S., & Khamar, L. (2024). Modeling of the coupling of dissolution and crystallization in the digestion tank of a wet phosphoric acid manufacturing process. In *Computer Aided Chemical Engineering* (Vol. 53, pp. 991-996). Elsevier.
12. Elmisaoui, S., Latifi, A. M., Khamar, L., & Salouhi, M. (2021). Analysis of the dissolution mechanism in the phosphoric acid manufacturing process: modelling and simulation. In *Computer Aided Chemical Engineering* (Vol. 50, pp. 891-897). Elsevier.
13. Elmisaoui S, Benjelloun S, Chkifa MA, Latifi AM. Surrogate model based on hierarchical sparse polynomial interpolation for the phosphate ore dissolution. Computers & Chemical Engineering 173:108174 (2023). https://doi.org/10.1016/j.compchemeng.2023.108174
14. Elmisaoui, S., Benjelloun, S., Chkifa, A., & Latifi, A. M. (2022). A sparse polynomial surrogate model for the shrinking core model in phosphate ore digestion. In *Computer Aided Chemical Engineering* (Vol. 51, pp. 1291-1296). Elsevier.
15. Bouchkira I, Latifi AM, Benyahia B. ESTAN—A toolbox for standardized and effective global sensitivity-based estimability analysis. Computers & Chemical Engineering 186:108690 (2024). https://doi.org/10.1016/j.compchemeng.2024.108690

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

MatStudio: A Human-in-the-Loop Framework for Microstructure Segmentation with SAM-Guided Refinement

Yao Xue[a], Yanhu Wang[b], Antonios Armaou[cd]*

a The Pennsylvania State University, Department of Mechanical Engineering, University Park, PA 16802, USA
b Wenzhou University, College of Mechanical and Electrical Engineering, Wenzhou, 325035, Zhejiang, China
c University of Patras, Department of Chemical Engineering, Rio, 26504, Greece
d The Pennsylvania State University, Department of Chemical Engineering, University Park, PA 16802, USA
* Corresponding Author: armaou@psu.edu; armaou@upatras.gr

ABSTRACT

Microstructure segmentation is essential for quantitative materials analysis; however, supervised deep learning demands substantial annotation, whereas general-purpose foundation models such as the Segment Anything Model (SAM) offer limited domain-specific semantic control. This paper presents MatStudio, a human-in-the-loop framework for microstructure segmentation that is proposed and implemented end to end in this work. MatStudio couples an interactive workflow for batchwise micrograph annotation and model adaptation with a dual-head convolutional architecture and SAM-guided boundary refinement. The loop combines sparse supervision with SAM-assisted labeling, task-specific training, and iterative batch-level correction, typically converging within two to three cycles.. The network comprises a shared encoder initialized from a pretrained backbone and two decoders: a UNet-style segmentation head that jointly predicts class labels and pixelwise uncertainty, and a prototype branch that measures texture similarity against a memory bank of learnable prototypes. An uncertainty-aware training objective couples label estimation with confidence; uncertainty weighting emphasizes prototype learning in ambiguous regions, and purity filtering restricts prototype updates to phase-interior features to mitigate boundary contamination. At inference, decoder outputs are fused in a spatially adaptive manner according to local confidence. We further introduce SAMFusion, which treats SAM as a morphology-aware proposal source and accepts proposals that are consistent with the fused prediction under a hierarchical small-to-large activation criterion, yielding improved boundary fidelity without case-specific tuning of SAM prompts. Experimental evaluation on Ni-alloy micrographs and public benchmarks shows consistent gains relative to interactive scribble-based segmentation and SAM-centric reference methods.

Keywords: Machine Learning, Artificial Intelligence, Materials, Microstructure segmentation, Human-in-the-loop, Uncertainty quantification, Prototype learning, Segment Anything Model

INTRODUCTION

Microstructure governs the mechanical, thermal, and chemical properties of engineering materials, making quantitative microstructural characterization essential for process optimization and structure-property relationship discovery [1, 2]. The spatial distribution of phases, precipitates, grain boundaries, and defects directly influences macroscopic performance metrics such as strength, ductility, and corrosion resistance. As materials informatics advances toward high-throughput analysis and autonomous experimentation, accurate and efficient microstructure segmentation has become a critical bottleneck.

Traditional image analysis approaches relied on threshold-based methods, morphological operations,

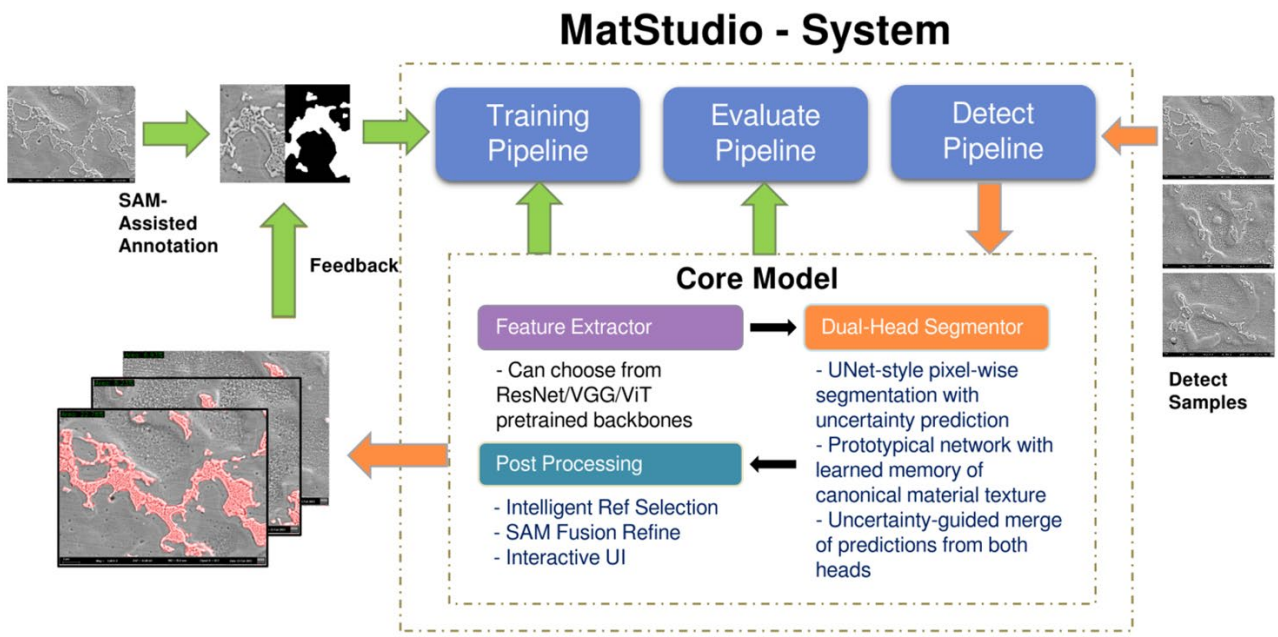

Figure 1. System-level overview of MatStudio. Training, evaluation, and detect interface with a shared core model (backbone, dual-head segmentor, and post-processing); feedback closes the loop between predictions and SAM-assisted annotation.

and hand-crafted features, which require extensive manual parameter tuning and exhibit poor generalization across imaging conditions. The advent of deep learning has substantially advanced automation: U-Net [3] and fully convolutional architectures demonstrated remarkable performance on biomedical and materials images, with subsequent work applying these methods to steel microstructures [4] and precipitate quantification. Pretrained encoders on large microscopy datasets have further improved generalization [5]. However, standard encoder-decoder architectures optimize for pixel-wise classification and may struggle to maintain consistent predictions across texturally homogeneous regions, particularly when training data is limited [6]. This limitation motivates complementary approaches based on metric learning.

Siamese networks [7] address few-shot recognition by learning embeddings where similar samples cluster together, enabling classification via distance comparison to reference examples. Prototypical networks [8] extend this paradigm by representing each class as a single prototype vector—the mean embedding of support examples—and classifying queries based on proximity to these prototypes. For semantic segmentation, PANet [9] demonstrated that prototype alignment enables effective few-shot segmentation by matching query features to class-specific prototypes derived from support masks. These metric-learning approaches excel at capturing texture consistency: once a canonical texture representation is established, regions sharing that texture can be reliably identified regardless of their spatial distribution. However, prototype-based methods typically operate at reduced resolution and may produce coarse boundaries, while their performance depends heavily on the quality and purity of the reference examples used to construct prototypes.

More recently, foundation models trained on massive datasets have demonstrated remarkable zero-shot generalization capabilities. The Segment Anything Model (SAM) [10], trained on over one billion masks, has attracted particular attention for materials microscopy. MatSAM [11] employs automated prompt generation to leverage SAM's zero-shot segmentation for microstructure extraction, SAM-I-Am [12] uses geometric and textural features for semantic boosting in electron microscopy, and SAMRefiner [13] adapts SAM for universal mask refinement. In parallel, Human-in-the-Loop (HITL) frameworks such as scribble-based segmentation [14] have effectively mitigated data scarcity through iterative expert feedback.

Despite these advances, each paradigm exhibits inherent trade-offs. Standard segmentation networks excel at boundary delineation but may produce inconsistent predictions in material interiors where local context alone is insufficient. Prototype-based methods enforce texture consistency but sacrifice spatial precision and require

careful reference selection to avoid prototype pollution from boundary features. SAM-based zero-shot approaches achieve impressive results without training, but directly applying SAM to material microscopy poses challenges due to the dense and dispersed nature of microstructural features; moreover, these methods create dependency on prompt quality where suboptimal prompts directly degrade results. This analysis suggests an opportunity: rather than relying on any single approach, robust microstructure segmentation may be achieved by combining texture-consistent prototype matching with pixel-precise segmentation, while leveraging SAM's morphological sensitivity for boundary refinement.

Building on this insight, we present MatStudio, a unified framework that integrates an uncertainty-guided dual-head architecture with SAMFusion—a novel proposal-guided refinement strategy. The primary segmentation head produces pixel-wise predictions along with learned uncertainty estimates [15], which guide both training and inference of an auxiliary prototype memory bank. During training, high-uncertainty regions receive amplified supervision for the prototype branch, forcing it to specialize on cases where the segmentation head struggles; at inference, predictions are dynamically fused based on local confidence. SAMFusion further refines boundaries by treating SAM purely as a morphology-aware proposal generator, activating proposals only when corroborated by the fused prediction. This design inherits the boundary precision of segmentation networks, the texture consistency of prototype matching, and SAM's morphological sensitivity—while avoiding the failure modes of each approach used in isolation.

The remainder of this article is organized as follows. Section 2 presents the MatStudio methodology. Figure 1 gives a system-level overview; the text then details the interactive workflow, the dual-head architecture with prototype-based texture modeling, uncertainty-guided training and memory updates, spatially adaptive fusion at inference, and SAMFusion for proposal-gated boundary refinement. Section 3 reports comparative experiments on Ni-alloy micrographs and public benchmarks against ScribbleSeg and MatSAM, including the qualitative assessment in Figure 2. Section 4 summarizes the principal findings and discusses implications for interactive microstructure segmentation.

METHODOLOGY

Figure 1 sketches MatStudio at architecture level: training, evaluation, and detect pipelines connect to a shared core model, with outputs feeding back to annotation. The operational workflow follows a human-in-the-loop paradigm for narrow batches of micrographs under consistent imaging conditions. Users begin by loading a batch into an interactive annotation interface and providing sparse labels—point clicks or bounding boxes—for the target microstructural features. These annotations, optionally accelerated by SAM-assisted region proposals, bootstrap a task-specific segmentation model. After an initial training pass, the model generates predictions across the batch; the user then reviews these outputs and provides corrective feedback on under-segmented or mis-segmented regions, which refines the supervision signal for the next iteration. This iterative refinement cycle continues until the segmentation quality converges, typically within two to three rounds.

Dual-Head Architecture with Prototype Memory

The core segmentation model employs a dual-head architecture where a primary segmentation decoder is complemented by an auxiliary prototype-based similarity head. A shared convolutional encoder, initialized from a pretrained backbone, extracts hierarchical feature representations that feed two parallel branches: (1) a UNet-style segmentation head producing dense pixel-wise predictions along with a learned uncertainty signal, and (2) a prototype head that computes texture similarity against a dynamically maintained memory bank of canonical material textures.

The segmentation head is designed to output not only class predictions but also a per-pixel confidence estimate, inspired by heteroscedastic uncertainty modeling in Bayesian deep learning [15]. This enables the network to explicitly signal regions where it lacks confidence—typically phase boundaries, texturally ambiguous zones, and imaging artifacts. The auxiliary prototype head maintains a compact set of learnable prototype vectors representing characteristic material textures. For each spatial location, the head performs texture similarity matching against all prototypes and outputs a dense similarity map that captures texture consistency across the image.

2.2 Uncertainty-Guided Learning

The key insight of our training strategy is that the segmentation head's uncertainty provides a natural supervisory signal for the auxiliary branch. The segmentation head is trained with an uncertainty-aware loss formulation that jointly learns predictions and confidence: when prediction error is high, the network learns to increase its uncertainty estimate, effectively admitting difficult regions. At equilibrium, the uncertainty map captures precisely those regions where the segmentation head consistently struggles.

This learned uncertainty is then used to modulate supervision for the prototype head. Regions with high uncertainty receive amplified training signal, directing the prototype branch to specialize on cases the segmentation head cannot handle reliably. This uncertainty-

weighted training ensures complementary specialization: the segmentation head handles straightforward regions while the prototype head focuses its learning capacity on challenging textures and ambiguous boundaries. The total training objective combines both components with appropriate weighting.

To maintain prototype quality, we employ a conservative update strategy with strict purity filtering. Before updating prototypes, we identify regions that are unambiguously inside material phases—far from boundaries—and use only these "pure" features to update the memory bank. This prevents prototype pollution from boundary features, which would otherwise cause boundary blurring and increased false positives.

Inference and SAMFusion Refinement

At inference time, the learned uncertainty directly governs how predictions from both heads are combined. In confident regions, the segmentation head dominates; in uncertain regions, the prototype-based similarity takes over. This dynamic, spatially-adaptive fusion requires no additional hyperparameters and automatically adapts to local image content.

To further enhance boundary quality, we introduce SAMFusion, a proposal-guided refinement strategy that leverages SAM's exceptional morphological sensitivity without inheriting its failure modes. Unlike approaches that treat SAM outputs as authoritative segmentations—and consequently suffer when SAM misinterprets complex microstructures—SAMFusion treats SAM purely as a morphology-aware proposal generator. We invoke SAM to obtain an exhaustive set of candidate regions, then selectively activate proposals through an agreement-gated mechanism: proposals are accepted only when they exhibit sufficient overlap with the network's fused prediction. This hierarchical activation process, which evaluates proposals from small to large, ensures that SAM's boundary precision refines the segmentation without introducing spurious detections in texturally ambiguous regions. Crucially, this design eliminates the need for case-by-case hyperparameter tuning of SAM's prompt strategy.

RESULTS

We benchmarked MatStudio against ScribbleSeg [14] and MatSAM [11] on four segmentation tasks spanning self-collected Ni-alloy micrographs (cases a–b) and public datasets—UHCS spheroidite [16] and NBS-2 eutectic carbides [14]—as shown in Figure 2.

Cases (a) and (b) expose complementary failure modes of the baselines. In (a), dendritic arms spanning the full field-of-view defeat MatSAM's edge-based prompting, which cannot reliably distinguish foreground from background when boundaries lack clear termination; ScribbleSeg under-segments due to textural ambiguity between dendrites and intermediate grey substructures in the matrix. In (b), off-axis detector gradients and secondary-electron edge halos generate spurious boundaries that MatSAM interprets as phase interfaces, producing extensive false positives; the same gradient also causes ScribbleSeg to fail.

Case (c) presents well-defined spheroidite particles that all methods largely detect; however, surface cracks partition the matrix into coarse enclosed domains that, despite containing no target particles, exhibit closed-contour morphology indistinguishable from valid regions under edge-based prompting, causing MatSAM to register these large spurious areas as false positives. Case (d) constitutes the most demanding scenario: eutectic carbides with amorphous, interconnected morphologies cause catastrophic failure in MatSAM regardless of hyperparameter tuning, while ScribbleSeg partially succeeds but cannot resolve regions where carbide and matrix textures overlap. We note that published metrics for ScribbleSeg and MatSAM on these benchmarks could not be reproduced under our experimental protocol; moreover, manual inspection revealed annotation inconsistencies in the ground-truth masks—several target particles remain unlabeled while certain background regions are erroneously marked—rendering absolute metric comparisons unreliable. Despite these caveats, IoU and Dice reported as overlay annotations on each prediction panel in Figure 2 still permit relative comparison among methods on the same exemplars and under our evaluation settings.

The IoU and Dice coefficients printed on those panels summarize performance on the displayed crops. MatStudio attains the highest IoU and Dice in every row: in (a), 0.92 and 0.96 versus 0.44 / 0.62 (ScribbleSeg) and 0.41 / 0.58 (MatSAM); in (b), 0.74 / 0.85 versus 0.15 / 0.26 and 0.56 / 0.72; in (c), 0.66 / 0.80 versus 0.48 / 0.65 and 0.57 / 0.72; in (d), 0.95 / 0.97 versus 0.64 / 0.78 and 0.05 / 0.10. These rankings are consistent with the qualitative failure modes discussed above and should be read together with the protocol and labeling limitations noted earlier.

Across all four cases, and in agreement with both the visual evidence and the panel-reported scores in Figure 2, MatStudio maintains robust performance by exploiting the complementarity between learned texture representations and SAM's morphological sensitivity. The agreement-gated activation mechanism filters out proposals that, while geometrically plausible, lack corroboration from the semantic predictor—effectively suppressing edge artifacts and spurious enclosed regions. Crucially, by treating SAM outputs as candidate refinements rather than authoritative segmentations, the framework inherits SAM's boundary-delineation strengths without suffering its failure modes when edge-

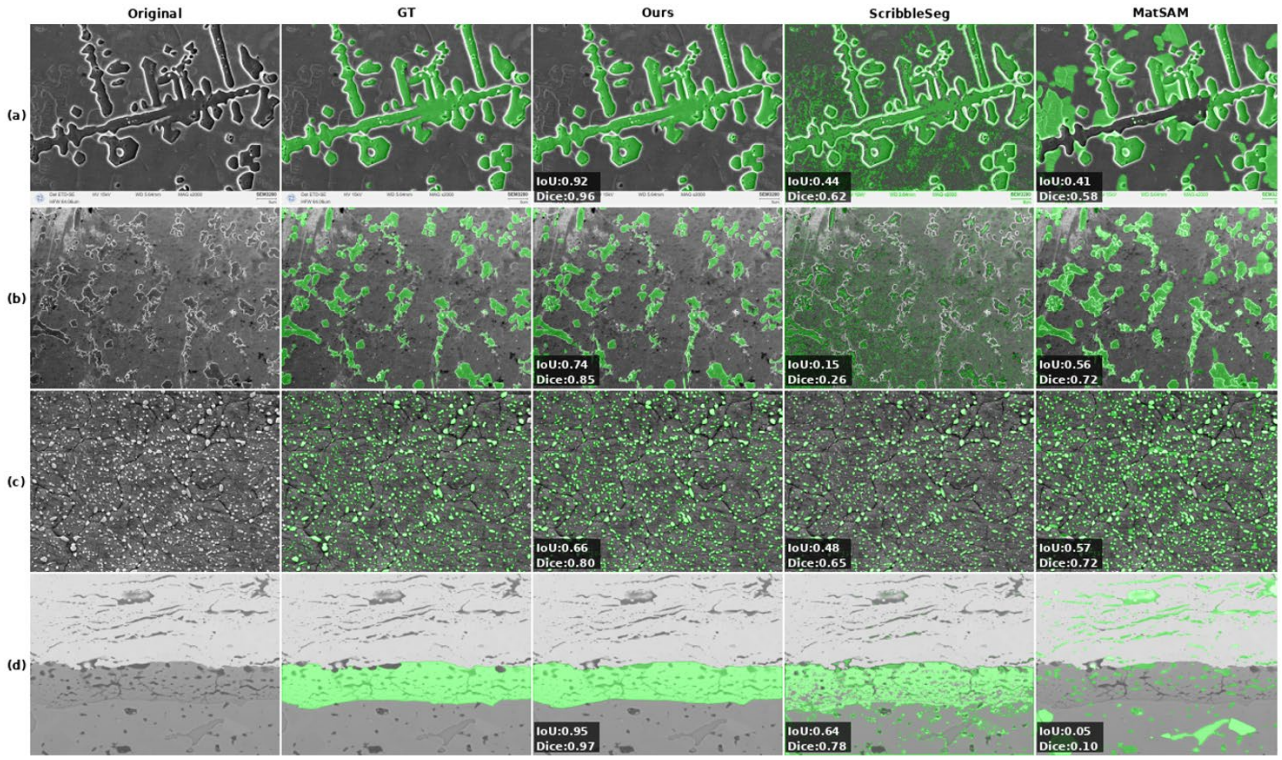

Figure 2. Qualitative comparison across four material systems. Columns show original image, ground truth, and predictions from MatStudio (ours), ScribbleSeg, and MatSAM. IoU and Dice scores are presented in overlay annotations. We observe that scribble-seg fails when there is no clear textural difference between the target and background. Similarly, MatSAM fails when boundary of the target area is complex or arbitrary. Our MatStudio is capable of delivering consistent and robust performance across difference scenarios and outperforms its counterparts in most cases.

based prompting encounters ambiguous or artifact-laden microstructures.

CONCLUSION

We presented MatStudio, a human-in-the-loop framework for microstructure segmentation that combines an uncertainty-guided dual-head model with SAMFusion for proposal-gated boundary refinement. Across the heterogeneous microscopy settings considered in this work, the comparison in results section illustrates a recurring pattern: methods that rely predominantly on edge- or prompt-driven segmentation exhibit substantial fragility when microstructural morphology, contrast, or imaging artifacts violate the assumptions those pipelines implicitly encode, whereas MatStudio remains comparatively stable by coupling texture-aware predictions with conservative use of SAM as a geometric proposal source. The dual-head design uses estimated uncertainty to allocate responsibility between dense segmentation and prototype-driven texture reasoning at inference, and SAMFusion retains SAM refinements only when they agree with the fused model, reducing spurious activations that arise when proposals are treated as definitive masks. We report exemplar-level overlays under our training and evaluation protocol together with the limitations discussed in the Results section; accordingly, we frame the findings as consistent relative evidence for the proposed integration rather than as universal rankings or as claims of proportional improvement over prior published numbers.

AUTHOR IDENTIFIERS

Author ORCIDs:
Xue Y: 0000-0002-2944-0005
Wang Y: 0009-0001-1200-6157
Armaou A: 0000-0002-8592-7934

REFERENCES

1. Cecen A, Dai H, Yabansu YC, Kalidindi SR, Song . Material structure-property linkages using three-dimensional convolutional neural networks. Acta Materialia 146:76-84 (2018).

https://doi.org/10.1016/j.actamat.2017.11.053
2. Shen C, Wang C, Wei X, Li Y, van der Zwaag S, Xu W. Physical metallurgy-guided machine learning and artificial neural network for high-strength alloy design. Acta Materialia 2021;205:116601.
3. Ronneberger O, Fischer P, Brox T. U-net: convolutional networks for biomedical image segmentation. Lecture Notes in Computer Science :234-241 (2015). https://doi.org/10.1007/978-3-319-24574-4_28
4. Azimi SM, Britz D, Engstler M, Fritz M, Mücklich F. Advanced steel microstructural classification by deep learning methods. Sci Rep 8: (2018). https://doi.org/10.1038/s41598-018-20037-5
5. Stuckner J, Harder B, Smith TM. Microstructure segmentation with deep learning encoders pre-trained on a large microscopy dataset. npj Comput Mater 8: (2022). https://doi.org/10.1038/s41524-022-00878-5
6. Yang L, Zhang Y, Chen J, Zhang S, Chen DZ. Suggestive annotation: a deep active learning framework for biomedical image segmentation. Lecture Notes in Computer Science :399-407 (2017). https://doi.org/10.1007/978-3-319-66179-7_46
7. Koch G, Zemel R, Salakhutdinov R. Siamese Neural Networks for One-shot Image Recognition. In: ICML Deep Learning Workshop 2015.
8. Snell J, Swersky K, Zemel R. Prototypical Networks for Few-shot Learning. In: NeurIPS 2017. p. 4077-4087.
9. Wang K, Liew JH, Zou Y, Zhou D, Feng J. Panet: few-shot image semantic segmentation with prototype alignment. 2019 IEEE/CVF International Conference on Computer Vision (ICCV) :9196-9205 (2019). https://doi.org/10.1109/iccv.2019.00929
10. Kirillov A, Mintun E, Ravi N, Mao H, Rolland C, Gustafson L, Xiao T, Whitehead S, Berg AC, Lo WY, Dollár P, Girshick R. Segment anything. 2023 IEEE/CVF International Conference on Computer Vision (ICCV) :3992-4003 (2023). https://doi.org/10.1109/iccv51070.2023.00371
11. Li C, Han X, Yao C, Guo Y, Li Z, Jiang L, Liu W, Huang H, Fu H, Ban X. A novel training-free approach to efficiently extracting material microstructures via visual large model. Acta Materialia 290:120962 (2025). https://doi.org/10.1016/j.actamat.2025.120962
12. Abebe W, Strube J, Guo L, Tallent NR, Bel O, Spurgeon S, Doty C, Jannesari A. Sam-i-am: semantic boosting for zero-shot atomic-scale electron micrograph segmentation. Computational Materials Science 246:113400 (2025). https://doi.org/10.1016/j.commatsci.2024.113400
13. Deng X, Wu H, Zeng R, Qin J. Memsam: taming segment anything model for echocardiography video segmentation. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) :9622-9631 (2024). https://doi.org/10.1109/cvpr52733.2024.00919
14. Na J, Kim SJ, Kim H, Kang SH, Lee S. A unified microstructure segmentation approach via human-in-the-loop machine learning. Acta Materialia 255:119086 (2023). https://doi.org/10.1016/j.actamat.2023.119086
15. Kendall A, Gal Y. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In: NeurIPS 2017. p. 5574-5584.
16. DeCost BL, Lei B, Francis T, Holm EA. High throughput quantitative metallography for complex microstructures using deep learning: a case study in ultrahigh carbon steel. Microsc Microanal 25:21-29 (2019). https://doi.org/10.1017/s1431927618015635

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and ad-aptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Hybrid Multi-Task Learning for Sustainability-Aware Pharmaceutical Molecular Design

Yiming Ma[a], Shang Gao[a], and Brahim Benyahia[a]*

[a] Department of Chemical Engineering, Loughborough University, Epinal Way, Loughborough, LE11 3TU, United Kingdom
* Corresponding Author: B.Benyahia@lboro.ac.uk.

ABSTRACT

Environmental sustainability is increasingly recognized as a critical consideration in pharmaceutical development, yet it is rarely incorporated at the scale of molecular-level design. This study introduces a strategy to predict cradle-to-gate indicators that can be flexibly incorporated into multiple early-stage molecular prioritization scenarios. A dataset of 150 pharmaceutical-relevant molecules was compiled, with each molecule described by structural descriptors, thermophysical properties, and ReCiPe endpoint indicators representing human health, ecosystem quality, and resource scarcity. A dual-branch multi-task model combining graph-based and descriptor-based representations was trained to predict these three endpoint indicators. Model performance was evaluated through validation metrics, local sensitivity analysis, and SHAP-based interpretability. A case study with solubility-based feasibility constraints was then used to illustrate how different sustainability weighting schemes affect molecular ranking and to demonstrate the potential for incorporating sustainability assessment into early-stage molecular prioritization. The results indicate that sustainability preferences can lead to distinct prioritization patterns, while some candidates remain comparatively favourable across scenarios.

Keywords: Machine Learning, Modelling and Simulations, SimaPro, Life Cycle Analysis, Surrogate Model

1. INTRODUCTION

Drug molecular design constitutes a critical interface between microscopic chemical structure and macroscopic product performance [1-3]. Small-molecule therapeutics continue to dominate the pharmaceutical industry due to their well-established development pipelines, scalable manufacturing routes, and broad therapeutic applicability [4, 5]. In parallel, rapid advances in computational chemistry, machine learning (ML), and optimization algorithms have positioned computer-aided molecular design as a powerful tool for navigating the vast chemical space and enabling rational, data-driven discovery beyond traditional trial-and-error approaches [6, 7].

Environmental sustainability is increasingly being incorporated into chemical research and development, not merely to satisfy policy and regulatory requirements, but also as an important consideration in the design and manufacture of chemical products across their life cycles [8, 9]. In current pharmaceutical development practice, however, sustainability considerations are typically introduced only after target molecules have been selected. At this stage, environmental impacts are primarily addressed through downstream, process-level interventions, such as route optimization, solvent substitution, or constraint-based decision-making [10-12]. Within this established paradigm, sustainability is primarily optimized with respect to process choices, whereas molecular structure itself is rarely examined as an explicit dimension for shaping cradle-to-gate sustainability outcomes. In the drug discovery workflow, although extensive combinatorial chemistry and structure-based design efforts are routinely conducted and a range of toxicity- and safety-related indicators may be evaluated, sustainability information remains fragmented, indirect, and rarely integrated into molecular-level candidate assessment [13, 14]. As a result, cradle-to-gate sustainability risks may be implicitly embedded in early design choices and only become apparent during scale-up and manufacturing stages, where opportunities for more fundamental improvement are already limited.

Recognizing this gap, primary approaches to linking chemicals with sustainability-related attributes have primarily relied on empirical heuristics and rule-based methods. In these methods, QSPR models are often employed as surrogate mappings from molecular structure to sustainability-relevant endpoints [15]. A typical workflow involves decomposing chemical structures into molecular descriptors, functional-group counts, and/or thermophysical properties, which are then combined with life-cycle impact indicators obtained from simulation tools or experimental databases. These relationships are commonly formalized using regression-based or other empirical modeling approaches to link readily computed chemical attributes with sustainability metrics[16, 17]. Such QSPR-enabled predictive LCA strategies are particularly attractive for early-stage decision-making, as they provide fast, structure-informed estimates enriched with basic physicochemical information. However, these classical empirical approaches typically rely on extensive, system-specific datasets for calibration and tend to exhibit limited extrapolation capability when applied to novel or sparsely explored chemical spaces. ML methods have rapidly emerged as an alternative paradigm to address these challenges. By learning complex structure-property relationships, ML models can achieve improved predictive performance across diverse molecular structures [18, 19]. For example, Song et al. [20] developed descriptor-based models to predict six life-cycle impact categories of chemicals, while Karka et al. [21] proposed a standardized ML workflow to digitize sustainable process development and enable a transition from ex post to ex ante life-cycle assessment. Despite these advances, the role of sustainability in early-stage molecular-level decision-making remains insufficiently explored. In pharmaceutical settings, the predictive models used to support such decisions are expected not only to provide reasonable accuracy, but also to be accompanied by interpretability analyses that promote trust, transparency, and informed decision-making [22, 23].

Therefore, this work presents a strategy for predicting cradle-to-gate sustainability indicators and examining their use in early-stage molecular prioritization. The proposed approach integrates structural descriptors, thermophysical properties, and graph-based molecular information within a hybrid representation, and employs a multi-task learning strategy to jointly predict multiple sustainability endpoints from a curated pharmaceutical-relevant dataset. A case study with solubility-based feasibility constraints is further used to illustrate how predicted sustainability indicators can affect molecular ranking under different weighting schemes.

2. METHOD

2.1 Database

A unified molecular database was constructed to support early-stage sustainability screening in pharmaceutical applications. A focused set of 150 candidate molecules relevant to pharmaceutical manufacturing were selected. For each molecule *i*, a vector of *M* sustainability-related outputs was defined, where M corresponds to multiple environmental impact indicators. In this study, three ReCiPe-based endpoint categories aligned with the areas of protection were considered: human health, ecosystem, and resources. Environmental impacts were quantitatively evaluated using SimaPro (v9.5) with the ReCiPe 2016 v1.1 method, based on the Ecoinvent 3 database and conducted in accordance with ISO 14040–14043 standards. Endpoint indicators were adopted to provide a comprehensive assessment of environmental performance, capturing aggregated long-term consequences across the three protection areas. Human health impacts were quantified using disability-adjusted life years, representing years of life lost or lived with disability. Ecosystem quality was measured as time-integrated local species loss (species·year), while resource scarcity was expressed in monetary terms (USD), reflecting the additional future costs associated with mineral and fossil resource extraction.

2.2 Functional Group Representation and Thermophysical Properties

2.2.1 Functional-group-based representation

Molecular structures were decomposed into a predefined dictionary of structural motifs using SMARTS-based substructure matching implemented in RDKit. The dictionary comprised several categories, including acid derivatives, carbonyl motifs, O-containing motifs, N-containing motifs, S/P-containing motifs, halogen motifs, heteroaromatic motifs, and auxiliary descriptors. Representative examples are summarized in **Table 1**. A hierarchical and non-overlapping matching rule set was applied, in which more specific motifs were matched before more general ones. For atom-consuming rules, matched atoms were assigned exclusively and excluded from subsequent atom-consuming matches. Non-consuming auxiliary descriptors were evaluated separately and retained in the final representation.

For a molecule *i*, the resulting representation was expressed as a vector of motif occurrence counts, where each element denotes the number of occurrences of motif *k* in molecule *i*. Binary presence indicators were generated for the same motif set. Additional discrete descriptors were included to encode stereochemical and topological information, including the number of chiral centers, double-bond configurations, and scaffold-level discriminators reflecting pharmaceutically relevant isomeric patterns.

Prior to feature extraction, all molecules were standardized using a uniform preprocessing pipeline

Table 1: Main categories of the predefined structural motif dictionary and representative examples.

Motif category	Representative examples
Acid derivatives	lactone_5, lactam_6, ester, carbamate
Carbonyl motifs	benzaldehyde, enone, dialkyl_ketone, carboxylate
O-containing motifs	phenol, alcohol, aryl_ether, epoxide
N-containing motifs	aniline, amine_primary, imine, amidine, guanidine
S/P-containing motifs	sulfone, sulfonamide, phosphate_triester, boronic_acid
Halogen motifs	aryl_halide, alkyl_halide, benzylic_halide
Heteroaromatic motifs	pyridine, oxazole, thiazole, quinoline
Auxiliary descriptors	aromatic_N, aromatic_O, sp2_C, sp3_C

comprising fragment removal, charge neutralization, stereochemical assignment, and tautomer canonicalization under controlled iteration limits. In total, the structure-based representation comprised 422 discrete variables, including 211 integer-valued motif counts and 211 corresponding binary indicators.

2.2.2 Thermophysical property descriptors

In addition to structure-based features, 38 continuous thermophysical descriptors were included in the molecular representation. These descriptors included molecular weight, critical properties, heat capacity, vapor pressure, viscosity, surface tension, enthalpy of vaporization, and related thermodynamic variables. Thermophysical data were compiled from open-source chemical information resources and property libraries, including PubChem and the *thermo* package, using SMILES, compound names, CAS numbers, and CIDs as identifiers. Missing entries were supplemented using estimation routines implemented in *thermo* and chemicals. Property values were collected or calculated at 101325 Pa and at 298.15, 323.15, and 373.15 K.

All continuous descriptors were standardized using z-score scaling. Categorical and binary variables were retained in their original form. Descriptors with zero or near-zero variance were removed prior to model training.

2.2.3 Graph-level attributes and reversibility

In parallel with the tabular descriptor representation, each molecule was encoded as a graph using atom-level and bond-level structural matrices. Graph data were stored in JSON-formatted atom and bond matrices and converted into fixed-size graph tensors prior to model training.

The atom matrix was parsed into a node-feature array, and the bond matrix was parsed into a symmetric weighted adjacency matrix. Bond entries retained bond-order information, including fractional values for aromatic bonds. Self-loops were added to valid nodes. Graph tensors were generated with a maximum size of 64 nodes, consisting of node features, adjacency matrix, and node mask. Molecules with fewer than 64 atoms were zero-padded. For disconnected graphs, the largest connected component was retained. These tensors were used as inputs to the graph branch of the model.

2.3 Multi-Task Model Training

A dual-branch multi-task neural network was constructed to jointly predict three sustainability endpoints. The model combined a tabular branch operating on descriptor-based inputs and a graph branch operating on molecular graph tensors. Three predictions were generated during training: a fused output, a graph-branch output, and a tabular-branch output.

The tabular input was formed by combining discrete structural descriptors with continuous thermophysical descriptors when available. Continuous descriptors were imputed using training-set means, standardized, and then filtered by variance thresholding and pairwise correlation analysis. The dataset was divided into training, validation, and test subsets with a ratio of 0.7/0.2/0.1. The three target variables were processed independently using column-specific screening rules. For each endpoint, a binary mask was generated to indicate retained samples, and retained values were scaled to the interval [1, 2] using training-set minimum and maximum values. The training target for each sample therefore consisted of the scaled response vector together with the corresponding output mask. The graph branch was implemented using a two-layer graph neural network operating on node features, weighted adjacency matrices, and node masks. The graph encoder included linear node projection, adjacency normalization, graph aggregation, layer normalization, PairNorm, dropout, gated residual connection, and combined mean-pooling and attention-pooling readout. The tabular branch consisted of residual fully connected layers with layer normalization and dropout. The fused prediction was defined as the sum of the graph-branch and tabular-branch predictions. Model training used masked mean squared error with endpoint-specific weights, and masked mean absolute error was used as an additional evaluation metric. Hyperparameters were optimized using Bayesian optimization. The tabular branch and graph branch were pretrained separately, followed by joint training of the full model. Adam optimization, gradient clipping, early stopping,

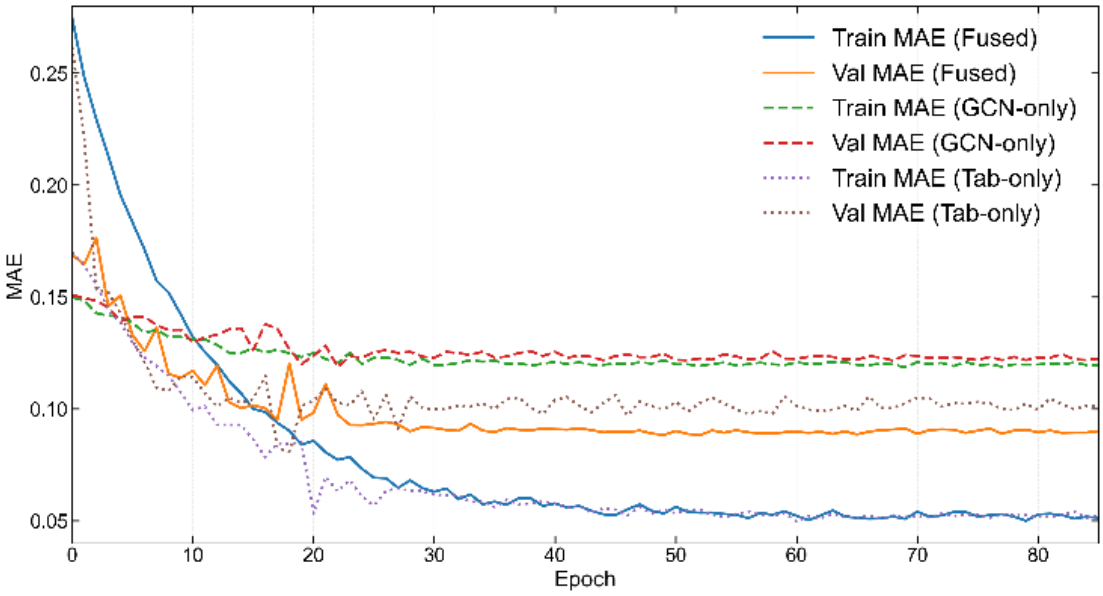

Figure 1. Training and Validation of MAE Across Model Heads

checkpointing, and learning-rate scheduling were applied during training. Model performance was evaluated on the held-out test set using MSE and MAE.

Model interpretability was assessed using SHAP-based feature attribution applied to the tabular descriptor branch. This analysis quantified the contributions of individual structural and thermophysical features to each sustainability endpoint, enabling chemically meaningful interpretation of learned structure–sustainability relationships and supporting transparent decision-making in early-stage molecular screening.

2.4 Probabilistic Molecular Design

The trained predictive model served as a surrogate evaluator within a genetic-algorithm-based search procedure. Candidate molecules were encoded using the same descriptor set as in model training and were iteratively evaluated by the surrogate to obtain performance scores. A distribution-based solubility feasibility criterion was applied, in which the feasibility of each candidate was assessed probabilistically to reflect model and property uncertainty. The detailed formulation of the optimization problem and the associated design rationale are presented in the application section.

3. RESULTS AND DISCUSSION

3.1 Dataset Description

3.1.1 Dataset Construction

The 150 molecular entries span a broad and representative range of chemical space relevant to pharmaceutical manufacturing. From a molecular weight perspective, the dataset covers compounds ranging from small, commonly encountered synthesis intermediates such as hydroxylamine (molecular weight 33 g·mol^{-1}, crystalline at ambient conditions) to formally regulated pharmaceutical chemicals and ligands such as diethylenetriaminepentaacetic acid (DTPA, molecular weight 393 g·mol^{-1}, crystalline at ambient conditions).

In addition to solid-state pharmaceutical compounds and intermediates, the dataset also includes a subset of widely used organic solvents, such as alcohols and esters, which are routinely involved in pharmaceutical synthesis and downstream processing. The inclusion of these solvent-type molecules further extends the chemical diversity of the dataset and ensures coverage of molecular classes for which extensive physicochemical and sustainability-related data are available.

3.1.2 Statistical Characteristics of Molecular Descriptors

The statistical characteristics of the descriptors were examined to assess diversity, redundancy, and representational conditioning. The pairwise binary Jaccard similarities of the discrete molecular fingerprints showed substantial heterogeneity across the 150 molecules, suggesting that the fingerprint representation preserved structural diversity and did not collapse into a highly redundant encoding.

For the continuous thermophysical property descriptors, principal component analysis indicated that variance was distributed across multiple principal components, with the first five components explaining approximately 80% of the total variance. Moreover, the property loadings were unevenly distributed across

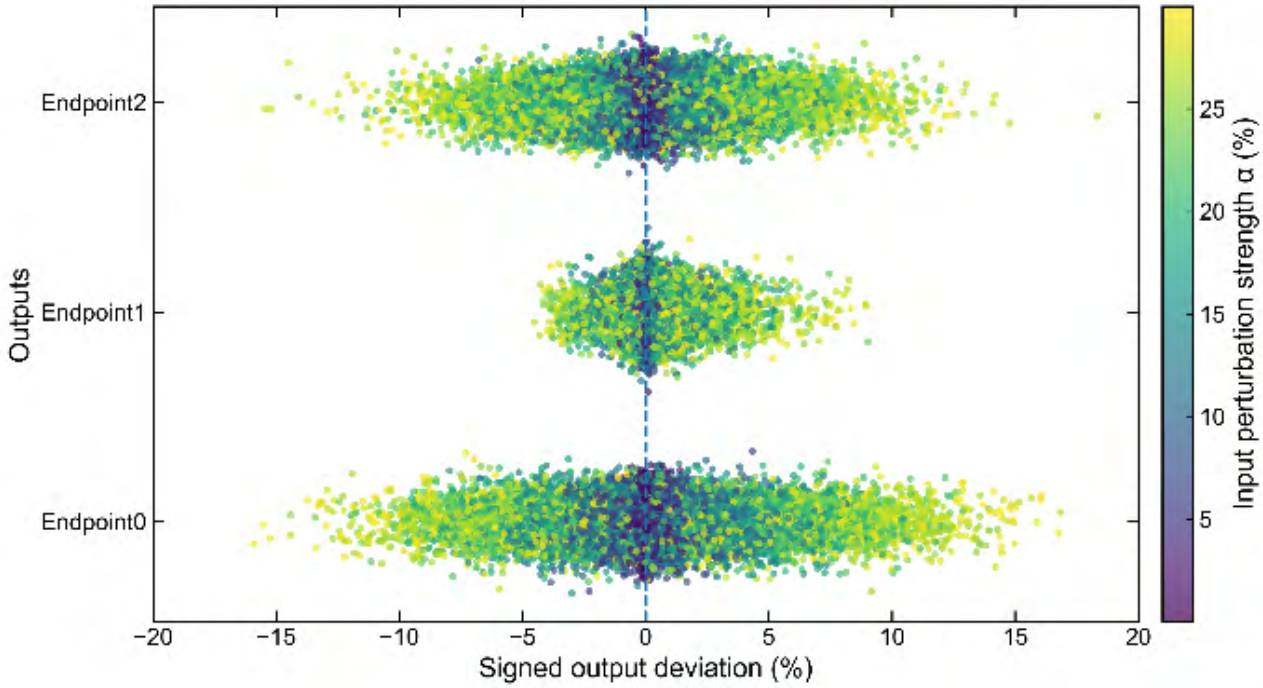

Figure 2. Monte Carlo Beeswarm of Output Deviations under Input Perturbations**.** For each Monte Carlo trial, a validation sample was randomly selected and its tabular input features were perturbed by a random multiplicative factor within ±α (α ∈ [0, 30%]), while the molecular graph inputs were kept unchanged. The horizontal axis shows the signed relative deviation of the predicted output from its unperturbed baseline, $(\hat{y}' - \hat{y})/(|\hat{y}| + \epsilon) \times 100\%$, centred at 0 (blue dashed line). Each point corresponds to one perturbed evaluation, and the colour encodes the input perturbation strength α (%). Narrower, more concentrated clouds around zero indicate higher robustness of the endpoint prediction; wider spreads and heavy tails indicate stronger sensitivity.

these leading components, implying that different subsets of descriptors contributed to different latent directions. Together, these observations suggest that the continuous descriptor set contained both shared and complementary information rather than complete redundancy. Furthermore, all 150 molecules were successfully reconstructed under the same fingerprint definitions, corresponding to a reconstruction success rate of 100%.

3.2 Model Training and Validation

This section evaluates the model's training behaviour and reliability from three perspectives: (i) the functional roles of the prediction heads and whether they operate as intended, (ii) potential overfitting under small-sample, mixed-variable conditions, and (iii) generalization performance on the validation set. Training dynamics are shown in **Figure 1** and **Table 1**. All heads show steadily decreasing training MAE with stable convergence, and the fused output exhibits consistently lower validation error than either individual branch, confirming the complementary roles of graph and descriptor representations.

Given the dataset size of 150 samples and the high-dimensional, mixed discrete–continuous feature space, overfitting risk was examined carefully. Most input variables are discrete counts without learned embeddings, and only a limited set of continuous descriptors contributes to interpolation. As shown in the training–validation curves, no divergence indicative of uncontrolled overfitting is observed.

To further assess model robustness, a local sensitivity analysis was performed by applying random multiplicative perturbations of up to ±30% to the continuous input descriptors, while keeping discrete features and molecular graph representations fixed. Across Monte Carlo perturbations, the resulting predictions remained centred around their unperturbed baselines and exhibited smoothly varying, bounded deviations (**Figure 2**). Importantly, the overall response patterns were preserved across all output endpoints, with no evidence of abrupt jumps or erratic behaviour. These observations indicate that the learned input–output mapping is locally smooth and internally consistent, supporting stable generalization rather than sensitivity to small, high-frequency input fluctuations.

Finally, the generalization performance of the proposed fused multi-task model was evaluated on the

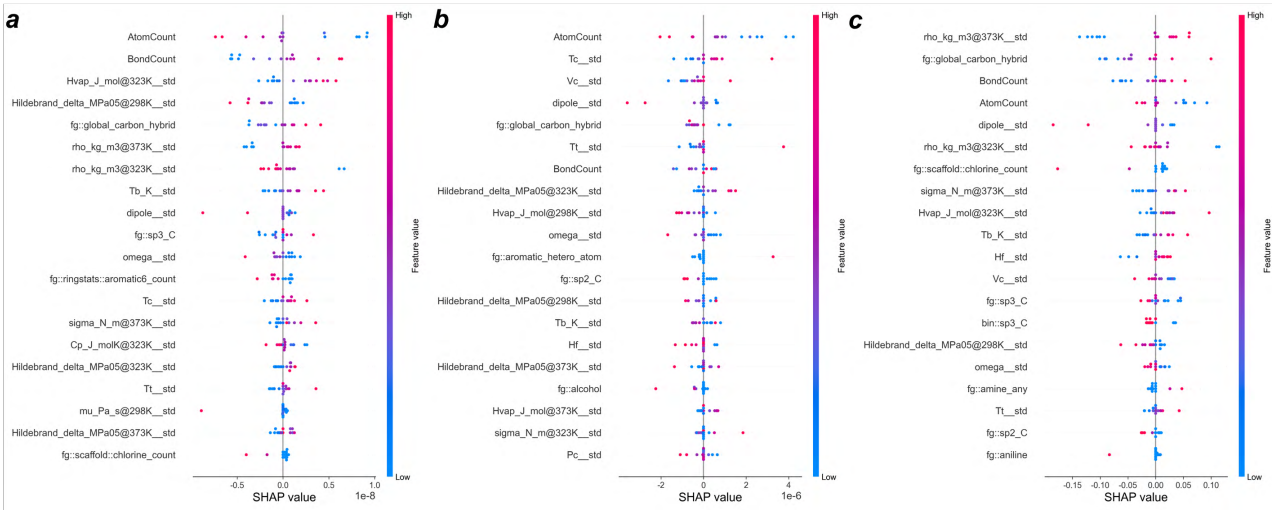

Figure 3. SHAP-based feature attribution on test set for three endpoints. SHAP beeswarm plots are shown for each endpoint, with features ordered by their mean absolute contribution. The x-axis denotes SHAP values reflecting the impact on model predictions, while the colour gradient (blue to red) indicates low to high feature values across the dataset.

validation set (**Table 2**). To provide a scale-aware assessment of predictive accuracy across outputs with different magnitudes, the symmetric mean absolute percentage error (sMAPE) was employed. The resulting sMAPE value of 23.82% indicates that, on average, prediction deviations remain bounded within a practically acceptable range under relative scaling. When considered together with the bounded absolute error metrics, this result suggests stable generalization behaviour on unseen data, rather than sensitivity to specific training instances.

Table 2: Prediction performance on the validation set.

Model	RMSE	MAE	MSE	sMAPE*
Fused-Multi-Task Model	0.271	0.190	0.0736	23.82

*sMAPE: Symmetric Mean Absolute Percentage Error (%)

3.3 SHAP

Figure 3 presents SHAP Beeswarm plots for the three predictive endpoints on test set. For endpoint human health (a), global molecular size descriptors and thermophysical properties, such as atom and bond counts, vaporization enthalpy, and Hildebrand solubility parameters, dominate the model contributions. Endpoint ecosystem quality (b) is primarily governed by standardised thermodynamic and state variables (e.g., critical temperature and volume), indicating a stronger reliance on continuous physicochemical attributes. In contrast, endpoint resource scarcity (c) shows increased sensitivity to structural and functional-group descriptors, suggesting that local chemical features play a more prominent role in shaping the model response.

The observed SHAP attribution patterns provide additional validation of the adopted hybrid descriptor strategy. Discrete structural features and continuous thermophysical properties are both actively utilised by the model, with their relative contributions varying coherently across endpoints. Such structured and endpoint-dependent feature usage suggests that the descriptor selection captures relevant chemical and physical information, rather than relying on arbitrary or redundant inputs.

3.4 Application

3.4.1 Problem Setting

The proposed molecular design strategy can be coupled with appropriate engineering and chemical constraints across multiple unit operations in pharmaceutical manufacturing, including synthesis, separation, crystallization, and downstream filtration–washing–drying steps. As a representative application, this study focuses on crystallization. When water is regarded as the most desirable and environmentally preferred crystallization solvent, the design task reduces to identifying the "cradle-to-gate greenest" molecular candidates whose crystallization yield is thermodynamically favored. Assuming that yield is governed solely by solubility thermodynamics and that solubility follows the General Solubility Equation (GSE), the problem becomes one of selecting the optimal molecule under a solubility-driven feasibility framework. The optimization problem is expressed as:

$$\min_{x \in \mathcal{X}} F(x) = \lambda_{\text{sol}}(1 - P_{\text{sol}}(x)) + \lambda_{\text{soft}}\,\Phi(x) \tag{1}$$

$$s.t. \quad x \text{ canonical SMILES, } \text{Elements}(x) \subseteq \{C, H, N, O, S, F\},$$

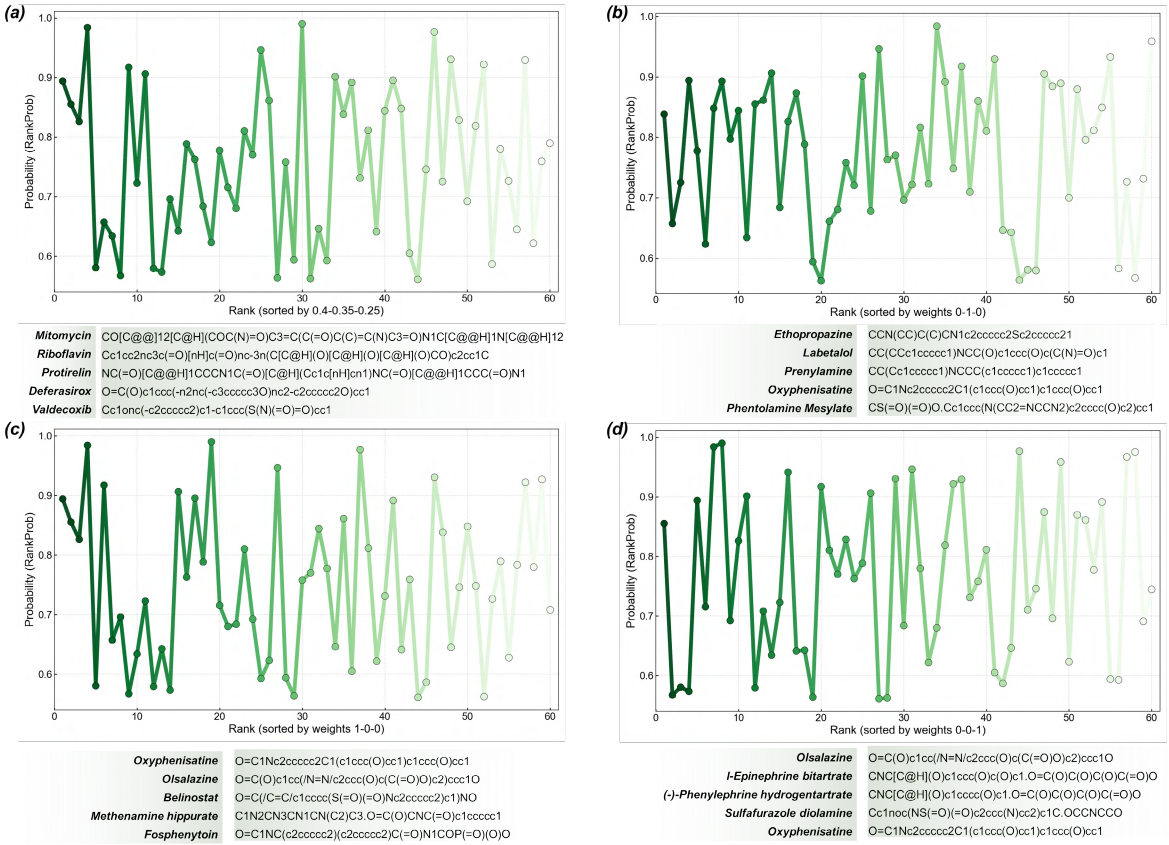

Figure 4: Candidates ranking profiles under four normalized endpoint weighting schemes: (a) 0.4–0.35–0.25, (b) 0–1–0, (c) 1–0–0, (d) 0–0–1, corresponding to emphasis on human health, ecosystem quality, and resource depletion, respectively.

$$\mathrm{Atoms}(x) \le 80,\ 300 \le MW(x) \le 400,\ \ \mathrm{Charge}(x) = 0$$

$$P_{\mathrm{sol}}(x) = \Pr(\log S(x) \ge \log S_{\min} = 10^{-2}),$$

$$\mathrm{Log}\, S(x) = a + b\,(T_m(x) - 25) + c \log P(x) + \epsilon,$$

$$a \sim N(0.5, 0.1^2), b \sim N(-0.01, 0.002^2),$$

$$c \sim N(-1.0, 0.10^2), \epsilon \sim N(0, 0.30^2),$$

$$\hat{\mathbf{y}}(x) = \mathcal{M}_{\mathrm{Fused(ML),\ 3\text{-}out}}(x) = \big(\hat{y}_1(x), \hat{y}_2(x), \hat{y}_3(x)\big),$$

$$0 < \hat{y}_k(x) < z, k = 1, 2, 3.$$

Within this formulation, molecular design is cast as an optimization problem in which each candidate SMILES must satisfy structural and com-positional feasibility constraints. The fitness function consists of a solubility-feasibility term and a soft structural penalty, where feasibility is defined as the probability that the predicted solubility exceeds a minimum threshold. This probability is obtained through Monte-Carlo evaluation of an uncertainty-aware GSE model, where a melting-point surrogate and calculated lipophilicity serve as inputs and the regression coefficients and residual variability are represented as random variables to capture intrinsic model uncertainty and propagated input noise. Three indicators predicted by the multi-task fused ML model prediction results simultaneously from molecular structure, where lower values correspond to greener and more sustainable process outcomes. Different weight combinations were applied to flexibly adjust the relative emphasis on sustainability objectives.

3.4.2 Solubility-Constrained Molecular Optimization

To benchmark the design outcomes against real pharmaceutical space, the generated candidates were compared with Phase IV approved drug substances curated from ChEMBL. These approved APIs were screened under the same structural and solubility constraints, and their sustainability profiles under different weighting schemes are summarized in **Figure 4**.

Under the human health dominated configuration (1 0 0), the top ranked molecules exhibit a consistent set of structural features, including moderate polarity, low halogen content, and stable aromatic or semi rigid scaffolds. When ecosystem quality receives full weight (0 1 0), the ranking shifts toward molecules with fewer persistent functional motifs, reduced halogenation, and thermophysical attributes that correspond to faster

environmental degradation. In contrast, the resource scarcity dominated configuration (0 0 1) systematically promotes structurally simpler molecules that contain fewer heteroatoms and less complex functional groups, reflecting the cost-based nature of the resource metric.

Across the different weighting schemes, a subset of molecules consistently appears in the top ranked region. These repeatedly selected structures share general characteristics such as medium polarity, low halogenation, and the absence of resource intensive functional groups. Their recurrent presence indicates that certain structural motifs provide balanced sustainability performance across human health, ecosystem, and resource endpoints. This suggests the existence of a structural region that jointly satisfies the requirements of the three metrics.

Finally, it is important to clarify that the ranking behavior in **Figure 4** is determined entirely by ReCiPe endpoint indicators. The GSE based solubility probability serves a separate purpose within the optimization framework. It acts solely as a feasibility constraint, where a higher probability reflects a greater likelihood of meeting crystallization requirements rather than an improvement in sustainability performance.

4. CONCULSION

This work presents an approach for integrating environmental sustainability considerations into early-stage molecular design of pharmaceuticals. By combining structural descriptors, thermophysical attributes, and graph-level information through a hybrid representation, a multi-task model was developed to predict three ReCiPe endpoint indicators. The model exhibits stable training behaviour, controlled sensitivity to input perturbations, and interpretable feature usage, supporting its suitability for downstream design tasks.

Incorporation of the predictive model into a solubility-constrained optimization framework enables the evaluation of structurally feasible molecular candidates under uncertainty. Comparative analysis with Phase IV approved drug substances illustrates how endpoint-specific weighting schemes influence molecular prioritization. The human-health metric favours moderately polar, low-halogen rigid scaffolds; the ecosystem metric favours molecules with fewer persistent structural motifs; and the resource metric favours structurally simpler and heteroatom-sparse designs. Several molecules repeatedly appear in top-ranked regions across different weighting configurations, suggesting the presence of a structural domain that offers balanced performance across all three sustainability endpoints.

Overall, the strategy provides a systematic route for incorporating sustainability metrics at the molecular-design stage, supporting well-informed decision-making prior to downstream process development.

ACKNOWLEDGEMENTS

This project has received funding under the call HORIZON-HLTH-2021-IND-07 (SusPharma, grant agreement no. 101057430) and under the UKRI Horizon Europe Guarantee scheme (SusPharma, project reference no. 10038378).

AUTHOR IDENTIFIERS

Author ORCIDs:
Yiming Ma: 0000-0001-6937-1924
Shang Gao: 0009-0001-0541-1015
Brahim Benyahia: 0000-0002-5397-1498

REFERENCES

1. Rarey M. Molecular design. concepts and applications. von gisbert schneider und karl?heinz baringhaus.. Angewandte Chemie 121:1746-1747 (2009). https://doi.org/10.1002/ange.200900047
2. I.D. Kuntz, E.C. Meng, B.K. Shoichet, Structure-based molecular design, Accounts of Chemical research 27(5) (1994) 117-123.
3. R.B. Silverman, M.W. Holladay, The organic chemistry of drug design and drug action, Academic press2014.
4. Wu K, Karapetyan E, Schloss J, Vadgama J, Wu Y. Advancements in small molecule drug design: a structural perspective. Drug Discovery Today 28:103730 (2023). https://doi.org/10.1016/j.drudis.2023.103730
5. Cummings MD, Sekharan S. Structure-based macrocycle design in small-molecule drug discovery and simple metrics to identify opportunities for macrocyclization of small-molecule ligands. J. Med. Chem. 62:6843-6853 (2019). https://doi.org/10.1021/acs.jmedchem.8b01985
6. Song CM, Lim SJ, Tong JC. Recent advances in computer-aided drug design. Briefings in Bioinformatics 10:579-591 (2009). https://doi.org/10.1093/bib/bbp023
7. Singh K, Bhushan B, Singh B. Advances in drug discovery and design using computer-aided molecular modeling. CAD 20:697-710 (2024). https://doi.org/10.2174/1573409920666230914123005
8. Jairoun AA, Al-Hemyari SS, Shahwan M, Alkhoujah S, El-Dahiyat F, Jaber AAS, Zyoud SH. Towards eco-friendly pharmaceuticals: regulatory and policy approaches for sustainable medicines use. Exploratory Research in Clinical and Social

Pharmacy 17:100576 (2025). https://doi.org/10.1016/j.rcsop.2025.100576
9. De Spiegeleer B, Wynendaele E. Sustainability of drug discovery, development and use as embedded in european pharmaceutical policies. Current Opinion in Green and Sustainable Chemistry 53:101028 (2025). https://doi.org/10.1016/j.cogsc.2025.101028
10. Ma Y, Sun M, Liu Y, Chen M, Wu S, Wang M, Wang L, Gao Z, Han D, Liu L, Wang J, Gong J. Design of spherical crystallization of active pharmaceutical ingredients via a highly efficient strategy: from screening to preparation. ACS Sustainable Chem. Eng. 9:9018-9032 (2021). https://doi.org/10.1021/acssuschemeng.1c01973
11. Y. Ma, S. Gao, N. Mehta, Q. Fu, W. Li, B. Benyahia, SolECOs: a data-driven platform for sustainable and comprehensive solvent selection in pharmaceutical manufacturing, Green Chemistry (2025).
12. Ma Y, Li W, Liu J, Shang G, Yang H, Gong J, Nagy ZK, Benyahia B. Digital design and optimization of the integrated synthesis and crystallization process using data?driven approaches. AIChE Journal 71: (2025). https://doi.org/10.1002/aic.18931
13. Yang H, Sun L, Li W, Liu G, Tang Y. In silico prediction of chemical toxicity for drug design using machine learning methods and structural alerts. Front. Chem. 6: (2018). https://doi.org/10.3389/fchem.2018.00030
14. Guengerich FP. Mechanisms of drug toxicity and relevance to pharmaceutical development. Drug Metabolism and Pharmacokinetics 26:3-14 (2011). https://doi.org/10.2133/dmpk.dmpk-10-rv-062
15. P. Baxevanidis, S. Papadokonstantakis, A. Kokossis, E. Marcoulaki, Group contribution-based LCA models to enable screening for environmentally benign novel chemicals in CAMD applications, AIChE Journal 68(3) (2022) e17544.
16. Calvo-Serrano R, González-Miquel M, Papadokonstantakis S, Guillén-Gosálbez G. Predicting the cradle-to-gate environmental impact of chemicals from molecular descriptors and thermodynamic properties via mixed-integer programming. Computers & Chemical Engineering 108:179-193 (2018). https://doi.org/10.1016/j.compchemeng.2017.09.010
17. Calvo-Serrano R, González-Miquel M, Guillén-Gosálbez G. Integrating cosmo-based ?-profiles with molecular and thermodynamic attributes to predict the life cycle environmental impact of chemicals. ACS Sustainable Chem. Eng. 7:3575-3583 (2018). https://doi.org/10.1021/acssuschemeng.8b06032
18. Zhao K, Peng X, Hu S, Wang X. Machine learning models for rapid prediction of chemicals' life-cycle environmental impacts: current status, challenges, and future directions. Green Carbon 3:433-442 (2025). https://doi.org/10.1016/j.greenca.2025.05.004
19. Kleinekorte J, Kleppich J, Fleitmann L, Beckert V, Blodau L, Bardow A. APPROPRIATE life cycle assessment: a process-specific, predictive impact assessment method for emerging chemical processes. ACS Sustainable Chem. Eng. 11:9303-9319 (2023). https://doi.org/10.1021/acssuschemeng.2c07682
20. Song R, Keller AA, Suh S. Rapid life-cycle impact screening using artificial neural networks. Environ. Sci. Technol. 51:10777-10785 (2017). https://doi.org/10.1021/acs.est.7b02862
21. Karka P, Papadokonstantakis S, Kokossis A. Digitizing sustainable process development: from ex-post to ex-ante LCA using machine-learning to evaluate bio-based process technologies ahead of detailed design. Chemical Engineering Science 250:117339 (2022). https://doi.org/10.1016/j.ces.2021.117339
22. Carter A, Imtiaz S, Naterer GF. Review of interpretable machine learning for process industries. Process Safety and Environmental Protection 170:647-659 (2023). https://doi.org/10.1016/j.psep.2022.12.018
23. Linardatos P, Papastefanopoulos V, Kotsiantis S. Explainable AI: a review of machine learning interpretability methods. Entropy 23:18 (2020). https://doi.org/10.3390/e23010018

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Machine Learning-Assisted Multi-PAT Data Fusion for Physics Consistent Crystallization Monitoring

Yiming Ma[a], Xuming Yuan[a], and Brahim Benyahia[a]*

[a] Department of Chemical Engineering, Loughborough University, Leicestershire, LE113TU, United Kingdom
* Corresponding Author: B.Benyahia@lboro.ac.uk.

ABSTRACT

Reliable multimodal monitoring in crystallization processes remains challenging due to heterogeneous PAT signal quality, sensor drift, asynchronous sampling and nonstationary noise. This work presents a machine-learning-assisted fusion framework that integrates multimodal PAT alignment, estimation and physics-guided regularisation to generate coherent concentration and particle-size trajectories. A mechanistically informed simulation platform is developed to produce synthetic Raman, FTIR, FBRM and image-based crystal size data with realistically simulated drift, heteroscedastic noise, dropouts and distortion patterns. Sensor reliability is inferred through a Random Forest model trained on variance-normalised discrepancies and quality metrics, which allows the dynamic adjustment of channel contributions. Across modalities, the Random Forest achieves MAE values of 0.03-0.20 for probability-type indicators and shows stable explanatory power for variance-inflation factors on particle-size channels (R^2 = 0.81-0.93). Two representative PAT scenarios illustrate the performance of the proposed framework under different cross-sensor discrepancy structures, demonstrating improved robustness, reduced local variance and enhanced physical coherence compared with individual signals. Overall, the results highlight the potential of multimodal information fusion as a foundation for trustworthy online monitoring and modelling in crystallization.

Keywords: Machine Learning, Process Monitoring, Modelling and Simulations, Surrogate Model

1. INTRODUCTION

With the rapid progression of continuous manufacturing, digital twins and process intensification, achieving reliable online observability has become central to ensuring product quality and supporting high-fidelity process modeling [1-4]. This need is particularly acute in crystallization-centered systems, where the time dependent evolution of solute concentration and crystal size governs supersaturation control, population-balance dynamics and critical quality attributes [5, 6]. Accurate and high-resolution online measurements that capture the underlying physicochemical evolution are therefore essential for model development, and model-based optimization [7, 8].

Industrial crystallization is typically monitored using process analytical technology (PAT) tools such as Raman and FTIR spectroscopy for concentration estimation, and FBRM or in situ imaging for particle size characterization [9-11]. Although these sensors provide complementary information, their performance in practice is affected by temperature- and solids-dependent spectral drift [12, 13], the structural mismatch between chord length and true particle sizes [14, 15], and imaging artefacts such as defocus or occlusion [16]. Additional challenges arise from asynchronous sampling between fast and slow PAT instruments, which leads to temporal misalignment and apparent discontinuities [17]. While various single sensor enhancement strategies have been proposed, inconsistencies between PAT channels remain common when individual sensors experience degradation or temporary downtime, complicating unified trajectory reconstruction for kinetic modeling, PBM simulation and soft sensing [18, 19].

To obtain more robust estimates of concentration and particle size in particulate processes, researchers have explored signal level fusion strategies, including basic weighting schemes such as inverse-variance

weighting [14, 20], rule-based gating heuristics [21, 22], and classical filtering or smoothing [23, 24] to combine information from heterogeneous PAT channels. While these approaches improve interpretability under relatively stable conditions, their performance remains largely confined to laboratory settings, with limited consideration of realistic industrial signal situations. In response, researchers have increasingly moved towards model level fusion frameworks that explicitly encode multi-sensor data structures, mechanistic constraints, and probabilistic noise models [25, 26]. These approaches treat heterogeneous PAT channels as jointly informative measurements embedded within unified state-space or hybrid mechanistic–statistical formulations. For example, nonlinear state estimators such as EKF, UKF and moving-horizon observers have been applied to reconstruct solute concentration and crystal size distributions from multiple process measurements in seeded and industrial crystallization systems [27, 28]. In parallel, multi-PAT sensor arrays that combine spectroscopic and particle focused modalities have demonstrated the value of structurally modelling cross channel complementarity for polymorphism and particle size monitoring [29]. Nonetheless, these methods are often built upon assumptions of stable noise characteristics and consistently reliable sensor availability, whereas industrial crystallization environments exhibit dynamic hydrodynamics, varying solids loading and fluctuating PAT signal quality. Achieving robustness, physical consistency and interpretability simultaneously therefore remains a critical open challenge.

To address these challenges, a system level framework is developed that integrates multimodal PAT data simulation, dynamic sensor weight learning, and robust fusion. First, a multi-sensor data generation platform is established. This platform emulates industrial PAT behaviours with drift, offset, heteroscedastic noise, local dropouts and asynchronous sampling, aiming to generate complex and highly perturbed signals that capture the multi-factor variability observed under realistic industrial operating conditions. Next, an ML-based adaptive weighting mechanism is introduced to adjust sensor credibility in real time according to signal quality. Finally, robust statistical methods combined with temporal and physicochemical constraints ensure that the fused concentration and particle size trajectories remain smooth, consistent and physically coherent even under challenging noise conditions.

2. METHODOLOGY

2.1 PAT Data Structure and Feature Construction

The analysis considers four PAT channels: Raman and FTIR for concentration, and image-based sensors (IMA, e.g., PVM) and FBRM-proxy channels for particle size. Temperature is taken as the reference timeline, and all PAT streams are merged using an as-of alignment within a fixed tolerance so that each time point contains a consistent set of observations and associated time-step information. The aligned dataset is defined in **Eq. (1)**.

$$\mathcal{D}(t) = \{C_{\mathrm{Raman}}(t),\ C_{\mathrm{FTIR}}(t),\ L_{\mathrm{PVM}}(t),\ L_{\mathrm{FBRM}}(t),\ T(t)\} \quad (1)$$

Concentration-related feature vectors $X_{conc}(t)$ are formed from the instantaneous Raman and FTIR values, first-order temporal differences computed over a 7-point window, local variance estimates, normalised inter-channel discrepancies, and the time step statistics inherited from the alignment procedure. Spectral quality metadata including correlation coefficients and wavelength-shift indicators are included as direct numerical features when present in the raw PAT output.

Particle size feature vectors $X_{size}(t)$ comprise the instantaneous size estimates obtained from IMA, the calibrated FBRM-proxy signals, as well as derived descriptors including local temporal gradients, rolling-window variability, and the normalised discrepancy between the image-based and FBRM-proxy channels. To represent channel specific noise behaviour, image-quality indicators and count variance metrics associated with the construction of the FBRM-proxy are also incorporated.

For general applications, the underlying dataset must provide synchronised multi-sensor time series with explicit timestamps, channel resolved variance estimates, and relevant quality metadata. In the present study, all PAT datasets used for model development and evaluation are synthetically generated cases designed to strictly conform to this structural specification.

2.2 Pseudo-Label Construction and Model Training

For each case, Raman–FTIR and IMA–FBRM pairs are aligned on the common time axis. At each time index t, the two sensor readings $y_1(t)$ and $y_2(t)$, together with their nominal standard deviations $s_1(t)$ and $s_2(t)$, are combined to obtain a robust reference value. The reference is defined as the minimiser in **Eq. (2).**

$$x(t) = \arg\min_{x} \left[\rho\left(\frac{x - y_1(t)}{s_1(t)}\right) + \rho\left(\frac{x - y_2(t)}{s_2(t)}\right)\right] \quad (2)$$

where $\rho(\cdot)$ is the Huber loss. When only one measurement is available, the reference is taken directly from that channel. Residuals $r_i(t)$ are computed, and the normalised deviation is calculated in **Eq. (3)**, which is used to characterise the discrepancy between each sensor reading and the reference, relative to its nominal noise level. Based on this deviation, a variance inflation factor (VIF) (**Eq. (4)**) is assigned to each sensor and subsequently clipped within a fixed interval to suppress unrealistically large inflation values.

$$z_i(t) = \frac{|\, r_i(t) \,|}{s_i(t)} \tag{3}$$

$$\alpha_i(t) = \frac{r_i(t)^2}{s_i(t)^2} \tag{4}$$

A probability type failure indicator is constructed by combining $z_i(t)$ with channel specific quality metrics, including spectral correlation and peak-shift descriptors for Raman/FTIR, and image correlation and count-variance descriptors for IMA/FBRM. The final probability is obtained in **Eq. (5)**.

$$p_{\mathrm{bad},i}(t) = \sigma(z_i(t) + \gamma\, q_i(t)) \tag{5}$$

where $q_i(t)$ is the aggregated quality score and γ is a constant scale factor.

The input feature matrices are those constructed in Section 2.1. Two independent multi-output Random Forest regressors are trained for concentration and size separately. Each model outputs a four-dimensional vector ($p_{bad,\,i}(t)$, $\alpha_i(t)$). The regressors consist of 500 trees, use bootstrap sampling, employ mean-squared error as the splitting criterion, and restrict leaf size to a minimum of three samples to stabilise the estimates. Maximum tree depth is not constrained, allowing the ensemble to approximate nonlinear patterns. Feature subsets at each split are selected using sqrt-d sampling, where dis the number of feature dimensions.

The full dataset is randomly divided into an 80/20 train–test split with a fixed random seed to ensure reproducibility. Each Random Forest model is trained on the training portion and evaluated on the test portion using mean absolute error, root-mean-square deviation, coefficient of determination, and signed bias. After prediction, outputs corresponding to probability-type indicators are clipped to [0, 1], and VIFs are clipped to the prescribed interval. These trained regressors serve as the reliability estimators for the sensor weighting scheme used in the fusion procedure of Section 2.3.

2.3 Sensor Fusion with Constraints and Regularisation

The trained models provide estimates of the probability of sensor failure and the corresponding VIF. These outputs jointly modulate the effective contribution of each channel. A channel is considered unreliable when $p_{bad,\,i}(t)$ exceeds a high-risk threshold, in which case it is temporarily removed from fusion; when only one channel remains valid, the fused estimate simply follows that channel. For all remaining channels, the effective measurement uncertainty is computed in **Eq. (6)**. The precision score can be defined in **Eq. (7)**.

$$\tilde{v}_i(t) = \alpha_i(t)\, v_i(t) \tag{6}$$

$$\mathrm{prec}_i(t) = \frac{(1 - p_{\mathrm{bad},i}(t))^{\gamma}}{\tilde{v}_i(t)} \tag{7}$$

where γ controls the sharpness of the penalty applied to unreliable sensors. The contribution weights are then obtained by normalisation (Eq. (8)), after which they are smoothed in time to prevent abrupt switching caused by local noise or transient prediction variability. The fused estimate and its propagated variance follow the standard precision weighted averaging rules (**Eqs. (9-10)**).

$$w_i(t) = \frac{\mathrm{prec}_i(t)}{\sum_j \mathrm{prec}_j(t)} \tag{8}$$

$$y_{\mathrm{fused}}(t) = \sum_i w_i(t)\, y_i(t) \tag{9}$$

$$v_{\mathrm{fused}}(t) = \sum_i w_i^2(t)\, v_i(t) \tag{10}$$

Short gaps caused by sensor dropouts are filled by linear interpolation with aggressive variance inflation, ensuring that reconstructed segments carry low confidence and are naturally down weighted in later steps. To maintain consistency with crystallisation physics, the fused trajectories are further regularised by two physical constraints. The first ensures that the fused concentration remains consistent with the solubility relation at the corresponding temperature. The second enforces growth-dominated behaviour, where particle number is approximately conserved and the evolution of the characteristic size follows the prescribed baseline kinetics.

3. RESULTS AND DISCUSSION

3.1 Model Performance

The behaviour of each output in **Table 1** is consistent with the way pseudo-labels are generated in the training pipeline. In the code, both $p_{\mathrm{bad},i}(t)$ and $\alpha_i(t)$ are computed from residuals relative to a robust fused reference obtained via Huber IRLS. Channels whose simulated noise structure is tightly controlled (FBRM counts and IMA imaging) produce pseudo-labels with low intrinsic variability, resulting in smoother targets and correspondingly higher R^2 values for both p_{bad} and α. In contrast, the Raman and FTIR channels incorporate larger fluctuations in their synthetic noise terms, including peak-shift, correlation degradation and variable spectral perturbations. These lead to a wider spread in the pseudo-labels and therefore lower R^2 for the corresponding regression tasks.

For the probability type indicators $p_{\mathrm{bad},i}$, the mapping is dominated by monotonic functions of the normalised deviation $z_i(t)$ and channel specific quality metrics. These relationships are relatively smooth and piecewise monotonic, which aligns well with the structure of the RF

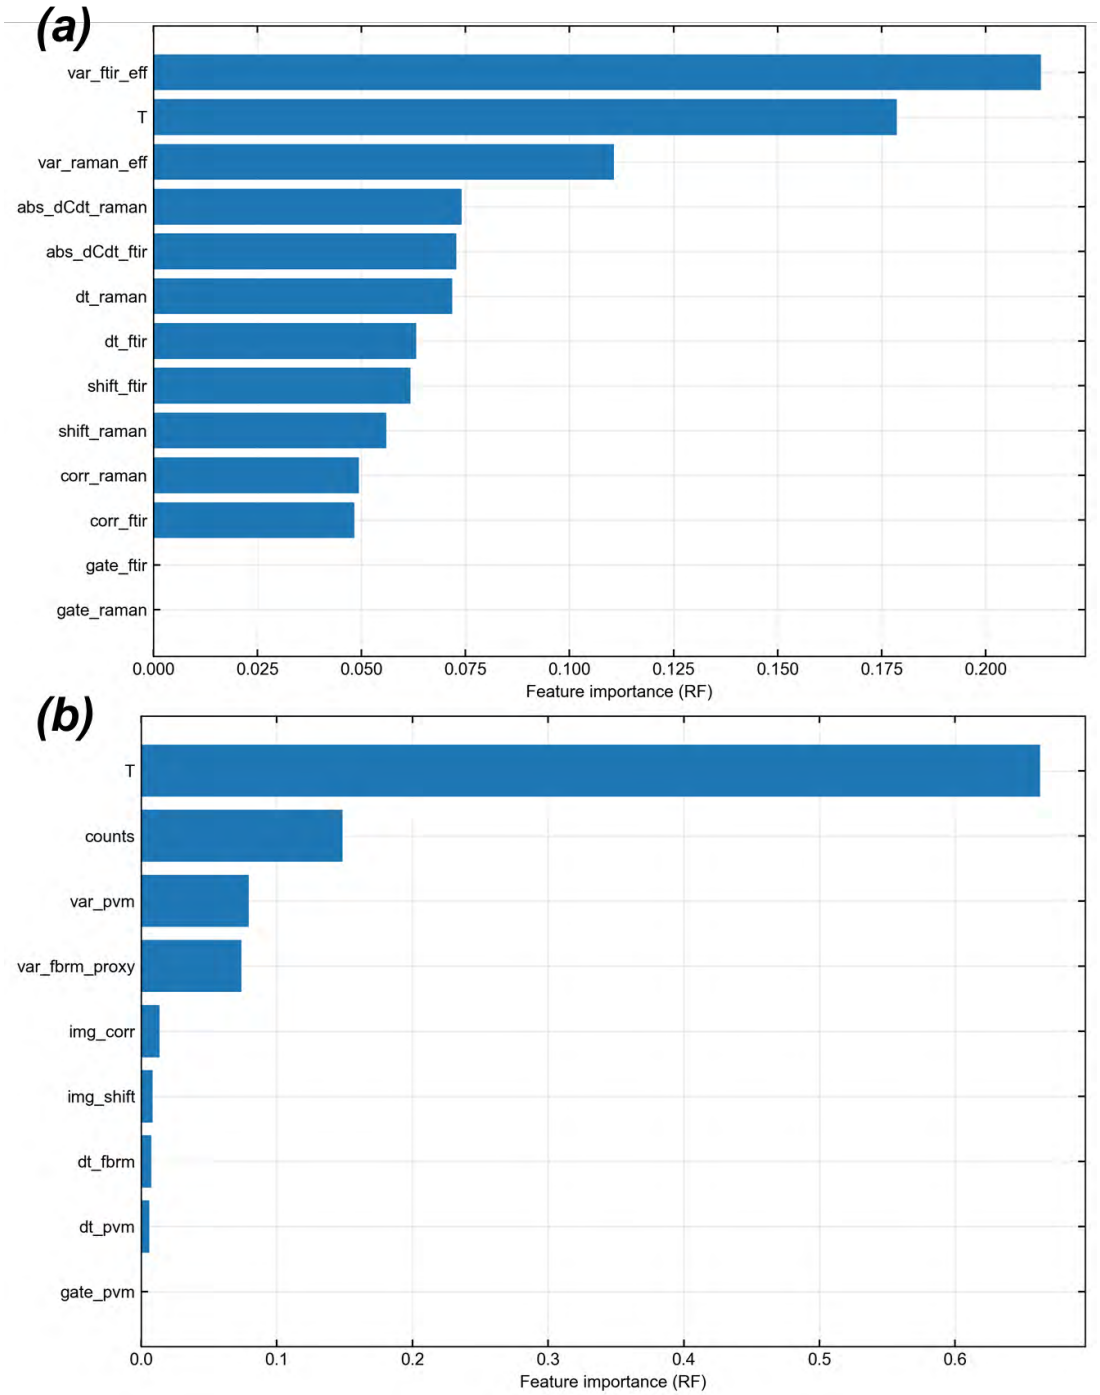

Figure 1. Feature importance of the RF models for (a) concentration and (b) particle size signals.

model. Consequently, all PAT channels yield low prediction errors and stable R^2, including the Raman and FTIR channels.

The VIFs α_i depend quadratically on residuals, and therefore amplify local variability more strongly. In the synthetic dataset, the Raman and FTIR channels possess higher perturbation amplitude and more irregular local fluctuations. This produces labels with higher dispersion and less predictable fine-scale structure, reducing the achievable R^2 even when MAE and bias remain small. The lower R^2 for α_{Raman} and α_{FTIR} therefore reflects the variability present in the labels rather than a systematic model deficiency, and does not affect the downstream fusion

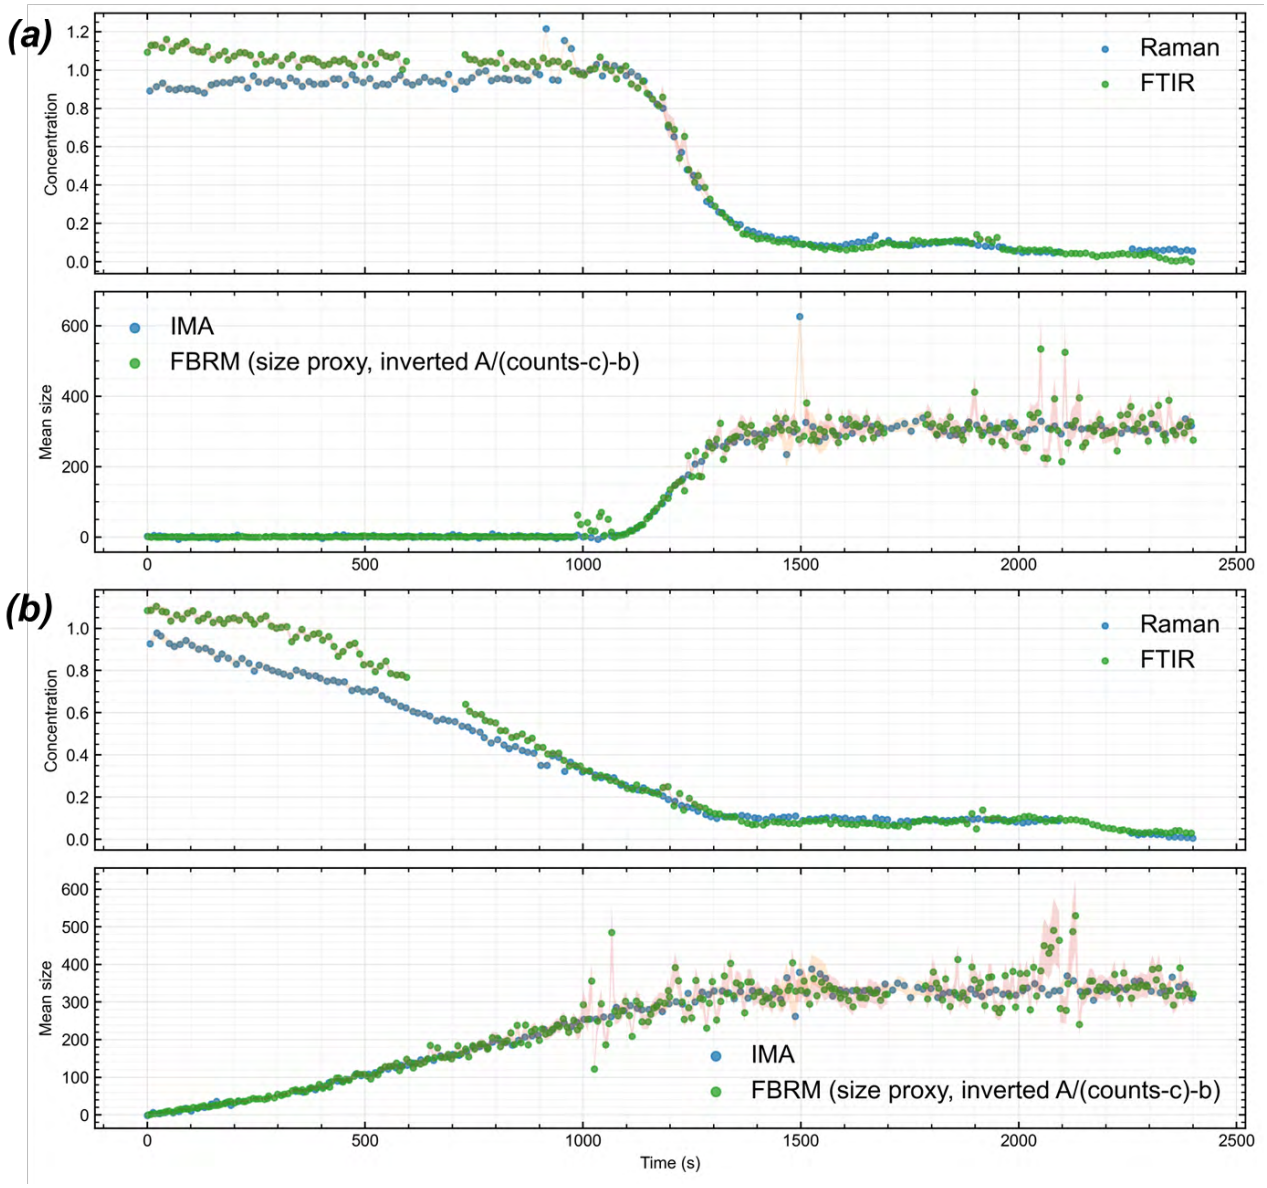

Figure 2. Synthetic scenarios used to evaluate fusion robustness cases (a) and (b).

stage where post clamped values are used only to modulate channel wise uncertainties.

Overall, the contrasting performance across channels directly follows from the pseudo label construction mechanism in the simulation code. The lower $R^2$for the Raman/FTIR αoutputs is expected under the imposed noise model and does not compromise the reliability estimates used in the subsequent fusion step.

The feature importance results reflect how the random forest allocates predictive weight across the input descriptors within the training dataset (**Figure 1**). For the concentration RF model, the highest ranked features are the effective variance terms (var_ftir_eff and var_raman_eff), followed by temperature and a set of short-term derivative descriptors (|dC/dt|, **Figure 1(a)**). These descriptors dominate the model's internal decision rules because they exhibit the strongest association with the pseudo-labelled reliability targets generated during training. Additional spectral structure descriptors, including peak-shift and correlation metrics, also receive non-negligible importance, consistent with their role in characterising signal perturbations within the dataset.

For the size RF model, temperature again appears with the highest importance values (**Figure 1(b)**). Other influential descriptors include the chord-count signal, variance terms derived from both IMA and FBRM-proxy channels, and image correlation and image shift measures. These variables receive higher weights

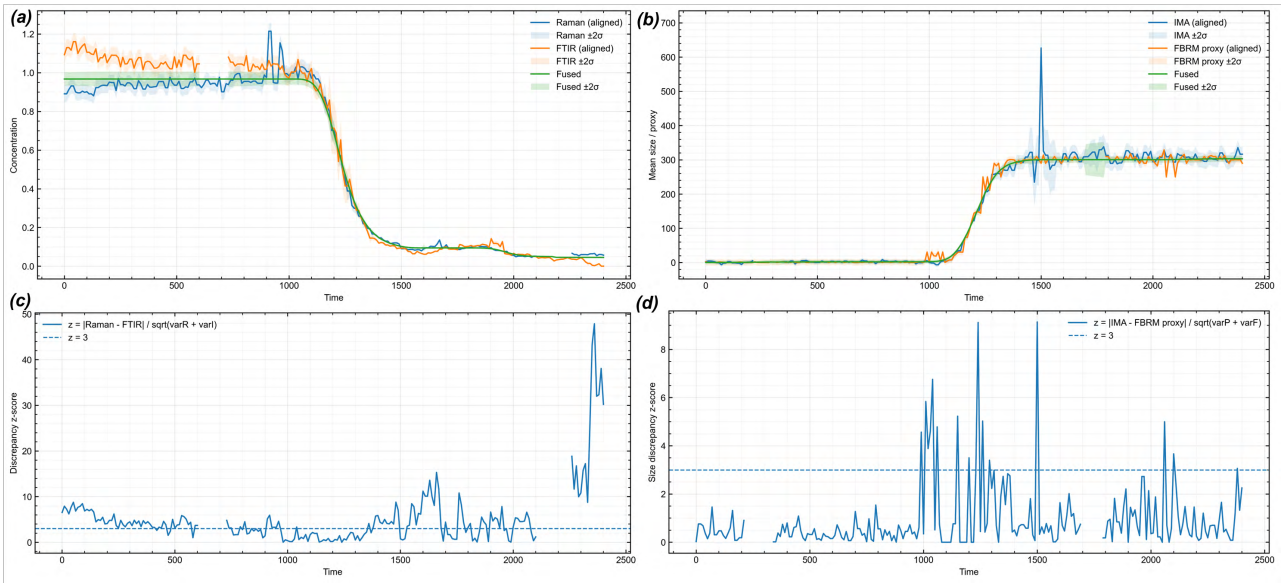

Figure 3. Multimodal PAT fusion results for Case 1. (**a**) aligned Raman and FTIR concentration measurements with fused estimates and uncertainty bands, (**b**) aligned IMA and FBRM-proxy mean size trajectories with fused estimates and uncertainty bands, and (**c–d**) discrepancy z-scores quantifying cross-sensor consistency over time.

because they co-vary with the reliability pseudo-labels across the synthetic cases used for training. Time-alignment and gating flags also contribute, although with lower importance relative to variance- and correlation-based descriptors.

Overall, the models emphasise descriptors that vary most consistently with the reliability pseudo-labels in the generated dataset, particularly those reflecting signal variance, short-term fluctuations, and structural perturbations in the PAT measurements. These importance patterns should be interpreted strictly as reflections of the models' internal learning behaviour under the specific data-generation and pseudo-labelling scheme adopted here, and do not imply generalisable diagnostic value beyond this study's design.

Table 1: Prediction performance of the multi-output Random Forest model for each PAT channel.

Output	MAE	RMSE	R^2	Bias
$P_{bad\ IMA}$	0.0605	0.1081	0.7220	-0.0041
α_{IMA}	0.7789	2.0620	0.8118	-0.0865
$P_{bad\ FBRM}$	0.0312	0.0586	0.9001	-0.0016
α_{FBRM}	0.2796	1.1892	0.9317	0.0017
$P_{bad\ Raman}$	0.1596	0.2115	0.5093	-0.0035
α_{Raman}	3.8445	5.6881	0.4896	-0.1377
$P_{bad\ FTIR}$	0.2047	0.2558	0.4654	0.0011
α_{FTIR}	6.1362	7.5885	0.4111	-0.0160

3.2 Case Studies

In this section two representative multimodal PAT cases are examined to illustrate the primary challenges in crystallization monitoring (**Figure 2**). Case a exhibits relatively stable concentration measurements but highly dynamic particle size evolution with an initial induction period, a rapid growth phase and sensor degradation episodes including imaging defocus, chord-length distortion and temporary dropouts. Case b features stronger baseline drift and temperature-dependent bias in the concentration channels, while the size measurements display continuous growth with local anomalies and inversion-related inconsistencies.

The fusion results for case 1 are available in **Figures 3 (a-d)**. The concentration trajectories in **Figure 3 (a)** show that Raman and FTIR follow the same overall evolution, although they display a noticeable baseline offset in the early stage. When the system enters the dominant crystallization transition between approximately 1100 and 1400 seconds, the two sensors respond coherently and the fused trajectory captures a physically smooth decline while presenting a narrower uncertainty envelope than either individual channel. This behaviour indicates effective fusion: the model is able to combine the complementary information from both sensors, reduce high-frequency noise and suppress misalignment-related variability. In the final segment of the experiment, roughly after 2200 seconds, the two concentration channels diverge substantially. The discrepancy z-score in **Figure 3 (c)** rises to values far beyond the expected statistical range, signalling that the disagreement between sensors is too large to be explained by their combined variance. This pattern is characteristic of sensor drift, intermittent artefacts or reliability degradation. The fused estimate remains stable during these periods and avoids

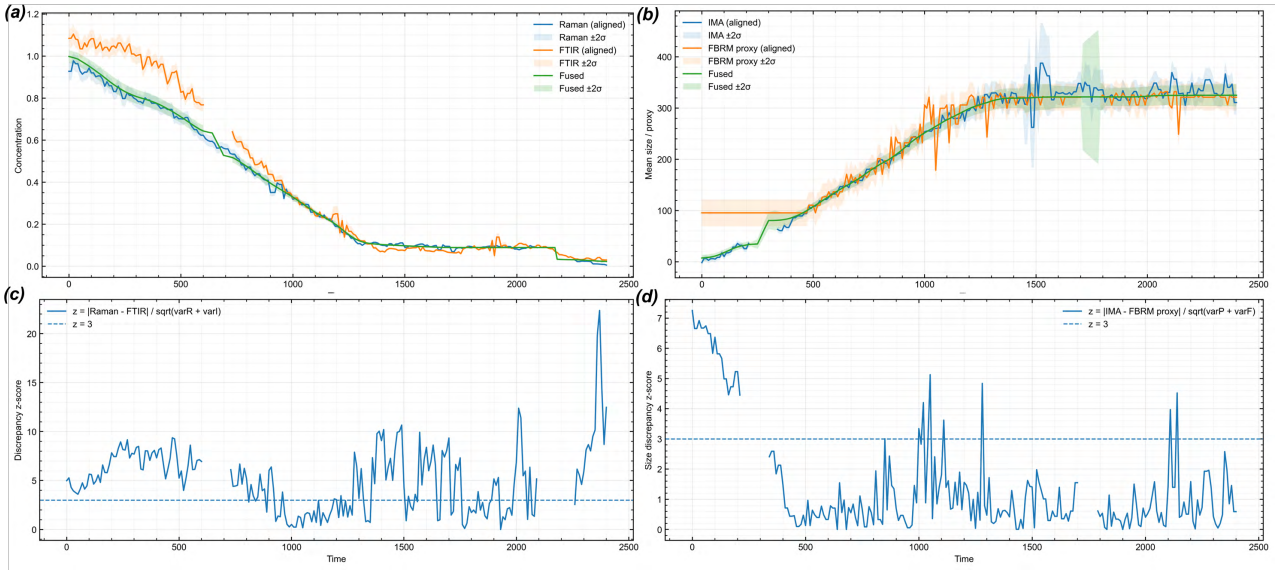

Figure 4. Multimodal PAT fusion results for Case 2. (**a**) aligned Raman and FTIR concentration measurements with fused estimates and uncertainty bands, (**b**) aligned IMA and FBRM-proxy mean size trajectories with fused estimates and uncertainty bands, and (**c–d**) discrepancy z-scores quantifying cross-sensor consistency over time.

propagating spurious fluctuations.

A similar behaviour is observed for the size signals in **Figure 3 (b)**. IMA and the FBRM-derived proxy exhibit consistent growth dynamics during the nucleation and growth regime and converge to a comparable plateau. IMA shows isolated spikes, most prominently near 1500 seconds. The fused estimate does not follow these irregular excursions and instead yields a smooth, monotonic trajectory that aligns with the expected crystallization trend. The discrepancy z-scores in **Figure 3 (d)** are small during the steady and slowly evolving periods, and markedly larger during transitional and disturbed regions. This indicates that the sensors agree well under stable conditions and that disagreements emerge mainly during rapid changes or localised noise events. Overall, the fused concentration and size trajectories are smoother, more physically interpretable and more robust to sensor artefacts than any individual signal.

In Case 2, the concentration signals display a stronger systematic discrepancy than in Case 1, with FTIR exhibiting an elevated baseline and larger drift throughout the early and mid stages (**Figures 4 (a-d)**). Despite the offset, both sensors follow similar downward kinetics, allowing the fused concentration to track the global trend while moderating channel specific bias and reducing local variance. The discrepancy z-scores remain consistently above the statistical threshold, indicating that the two sensors disagree more persistently than in Case 1 and signalling limited redundancy within the concentration pair.

For the particle size, IMA and the FBRM proxy show broadly similar growth behaviour during the early stage but diverge substantially after the main transition. IMA provides smoother and more continuous growth, while the FBRM proxy introduces abrupt fluctuations and intermittent artefacts, resulting in elevated size discrepancy z-scores across a wider time window. The fused size estimate mitigates these inconsistencies by suppressing FBRM-induced spikes and maintaining a physically coherent trajectory. Compared with Case 1, this case demonstrates higher cross-sensor disagreement and places greater emphasis on the fusion mechanism's ability to reject outliers and stabilise the evolution of both concentration and size.

3.3 Limitations and Future Perspectives

Although the present framework demonstrates robustness across heterogeneous PAT signals, several limitations remain. The analysis relies on synthetic scenarios and may not fully capture the complexity of real sensor behaviour, particularly long-term drift, chemometric model degradation, and process-dependent nonlinearities that can reshape the noise structure. In addition, the current fusion mechanism operates primarily at the signal level and does not explicitly integrate mechanistic constraints or population-balance consistency, which limits its ability to enforce global physical plausibility under severe sensor failure.

Future developments could incorporate hybrid physical statistical priors, adaptive drift modelling, and online reliability learning from historical process batches. Extending the framework to multivariate spectral domains, image-based morphological descriptors, and higher-frequency transient events would further enhance

its applicability to industrial crystallization.

4. CONCLUSIONS

This study developed a machine learning-assisted fusion framework for achieving reliable and physically coherent multimodal PAT monitoring in crystallization processes. By combining mechanistically informed data simulation, variance normalised discrepancy analysis and Random Forest-based reliability estimation, the method dynamically adjusts sensor contributions and mitigates the impact of drift, noise bursts, dropouts and cross-sensor inconsistencies. The physics-guided regularisation step further constrains the fused concentration and particle size trajectories to follow thermodynamic relations and growth-dominated crystallization behaviour.

Across two representative multimodal PAT scenarios, the framework consistently produced smoother, more interpretable and more robust trajectories than any individual sensor channel. In cases with mild disagreement, fusion effectively reduced noise and improved local consistency, whereas in scenarios with strong baseline drift or conflicting size measurements, the reliability model successfully down-weighted degraded sensors and preserved coherent kinetic trends. These results demonstrate that reliability-aware fusion can substantially enhance the integrity of multimodal crystallization monitoring under realistic industrial signal challenges.

ACKNOWLEDGEMENTS

This project has received funding under the call HORIZON-HLTH-2021-IND-07 (SusPharma, grant agreement no. 101057430) and under the UKRI Horizon Europe Guarantee scheme (SusPharma, project reference no. 10038378).

AUTHOR IDENTIFIERS

Author ORCIDs:
Yiming Ma: 0000-0001-6937-1924
Xuming Yuan: 0009-0006-2211-4879
Brahim Benyahia: 0000-0002-5397-1498

REFERENCES

1. Gao Y, Zhang T, Ma Y, Xue F, Gao Z, Hou B, Gong J. Application of pat-based feedback control approaches in pharmaceutical crystallization. Crystals 11:221 (2021). https://doi.org/10.3390/cryst11030221
2. Ma Y, Yuan X, Yewale A, Benyahia B. Recent advances in the digital transformation of crystallization process development and operation: synergy between model- and ai-driven strategies. Particuology 110:223-253 (2026). https://doi.org/10.1016/j.partic.2026.01.014
3. Ma Y, Wu S, Macaringue EGJ, Zhang T, Gong J, Wang J. Recent progress in continuous crystallization of pharmaceutical products: precise preparation and control. Org. Process Res. Dev. 24:1785-1801 (2020). https://doi.org/10.1021/acs.oprd.9b00362
4. Houson I, Siddique H, Chong MWS, Robertson M, Turner AJ, Nordon A, Osman A, Galindo A, Dunn AS, Johnston B, Benyahia B, Brown CJ, Mustoe CL, Price CJ, Reilly CD, Adjiman C, Jackson G, Prasad E, Halbert G, Feilden H, Sefcik J, McGinty J, Robertson J, Smith K, Al-Attili M, McGowan M, Siddique M, Pathan M, Rajoub N, Hamilton N, Feeney R, Brown S, Urwin SJ, Bernet T, Pickles T, Li W, Florence AJ. Quality by digital design in action: a workflow for crystallisation and isolation. International Journal of Pharmaceutics 686:126343 (2025). https://doi.org/10.1016/j.ijpharm.2025.126343
5. Li W, Hatcher LE, Wilson CC, Rielly CD, Benyahia B. A framework for systematic crystal shape tuning – case of lovastatin's needle-shaped crystals. Chemical Engineering Research and Design 202:126-146 (2024). https://doi.org/10.1016/j.cherd.2023.12.016
6. Houson I, Siddique H, Chong MWS, Robertson M, Turner AJ, Nordon A, Osman A, Galindo A, Dunn AS, Johnston B, Benyahia B, Brown CJ, Mustoe CL, Price CJ, Reilly CD, Adjiman C, Jackson G, Prasad E, Halbert G, Feilden H, Sefcik J, McGinty J, Robertson J, Smith K, Al-Attili M, McGowan M, Siddique M, Pathan M, Rajoub N, Hamilton N, Feeney R, Brown S, Urwin SJ, Bernet T, Pickles T, Li W, Florence AJ. Quality by digital design in action: a workflow for crystallisation and isolation. International Journal of Pharmaceutics 686:126343 (2025). https://doi.org/10.1016/j.ijpharm.2025.126343
7. Liu J, Benyahia B. Single and multiobjective shutdown optimization of a multistage continuous crystallizer. Ind. Eng. Chem. Res. : (2024). https://doi.org/10.1021/acs.iecr.3c03441
8. Yuan X, Nagy ZK, Benyahia B. A comprehensive framework for model evaluation and refinement using mbdoe estimability and structural identifiability: application to a crystallization process. Chemical Engineering Journal Advances 23:100823 (2025). https://doi.org/10.1016/j.ceja.2025.100823
9. Togkalidou T, Fujiwara M, Patel S, Braatz RD. Solute concentration prediction using chemometrics and ATR-FTIR spectroscopy. Journal of Crystal Growth 231:534-543 (2001).

https://doi.org/10.1016/s0022-0248(01)01518-4
10. Nagy ZK, Fujiwara M, Braatz RD. Monitoring and advanced control of crystallization processes. Handbook of Industrial Crystallization :313-345 (2019). https://doi.org/10.1017/9781139026949.011
11. Li W, Rielly CD, Benyahia B. Targeting desired particle size for improved dissolution and manufacturability of mefenamic acid. The 3rd International Online Conference on Crystals :21 (2022). https://doi.org/10.3390/iocc_2022-12165
12. Chong MWS, McGlone T, Chai CY, Briggs NEB, Brown CJ, Perciballi F, Dunn J, Parrott AJ, Dallin P, Andrews J, Nordon A, Florence AJ. Temperature correction of spectra to improve solute concentration monitoring by in situ ultraviolet and mid-infrared spectrometries toward isothermal local model performance. Org. Process Res. Dev. 26:3096-3105 (2022). https://doi.org/10.1021/acs.oprd.2c00238
13. Zhang F, Liu T, Wang XZ, Liu J, Jiang X. Comparative study on ATR-FTIR calibration models for monitoring solution concentration in cooling crystallization. Journal of Crystal Growth 459:50-55 (2017). https://doi.org/10.1016/j.jcrysgro.2016.11.064
14. Wynn EJW. Relationship between particle-size and chord-length distributions in focused beam reflectance measurement: stability of direct inversion and weighting. Powder Technology 133:125-133 (2003). https://doi.org/10.1016/s0032-5910(03)00084-6
15. Pandit AV, Ranade VV. Chord length distribution to particle size distribution. AIChE Journal 62:4215-4228 (2016). https://doi.org/10.1002/aic.15338
16. Cardona J, Ferreira C, McGinty J, Hamilton A, Agimelen OS, Cleary A, Atkinson R, Michie C, Marshall S, Chen YC, Sefcik J, Andonovic I, Tachtatzis C. Image analysis framework with focus evaluation for in situ characterisation of particle size and shape attributes. Chemical Engineering Science 191:208-231 (2018). https://doi.org/10.1016/j.ces.2018.06.067
17. Qu H, Alatalo H, Hatakka H, Kohonen J, Louhi-Kultanen M, Reinikainen SP, Kallas J. Raman and ATR FTIR spectroscopy in reactive crystallization: simultaneous monitoring of solute concentration and polymorphic state of the crystals. Journal of Crystal Growth 311:3466-3475 (2009). https://doi.org/10.1016/j.jcrysgro.2009.04.018
18. Zhao Y, Yuan J, Ji Z, Wang J, Rohani S. Combined application of in situ FBRM, ATR-FTIR, and raman on polymorphism transformation monitoring during the cooling crystallization. Ind. Eng. Chem. Res. :120912161452001 (2012). https://doi.org/10.1021/ie301241h
19. Orehek J, Tesli? D, Likozar B. Continuous crystallization processes in pharmaceutical manufacturing: a review. Org. Process Res. Dev. 25:16-42 (2020). https://doi.org/10.1021/acs.oprd.0c00398
20. Hansen TB, Simone E, Nagy Z, Qu H. Process analytical tools to control polymorphism and particle size in batch crystallization processes. Org. Process Res. Dev. 21:855-865 (2017). https://doi.org/10.1021/acs.oprd.7b00087
21. Zhao J, Wang M, Dong B, Feng Q, Xu C. Monitoring the polymorphic transformation of imidacloprid using in situ FBRM and PVM. Org. Process Res. Dev. 17:375-381 (2013). https://doi.org/10.1021/op300320a
22. Agimelen OS, Svoboda V, Ahmed B, Cardona J, Dziewierz J, Brown CJ, McGlone T, Cleary A, Tachtatzis C, Michie C, Florence AJ, Andonovic I, Mulholland AJ, Sefcik J. Multi-sensor inline measurements of crystal size and shape distributions during high shear wet milling of crystal slurries. Advanced Powder Technology 29:2987-2995 (2018). https://doi.org/10.1016/j.apt.2018.09.003
23. Fujiwara M, Chow PS, Ma DL, Braatz RD. Paracetamol crystallization using laser backscattering and ATR-FTIR spectroscopy: metastability, agglomeration, and control. Crystal Growth & Design 2:363-370 (2002). https://doi.org/10.1021/cg0200098
24. Trifkovic M, Sheikhzadeh M, Rohani S. Multivariable real?time optimal control of a cooling and antisolvent semibatch crystallization process. AIChE Journal 55:2591-2602 (2009). https://doi.org/10.1002/aic.11868
25. Wu WL, Mills MM, Schacht U, Rabinowitz C, Vlachos V, Nagy ZK. Sensor fusion and calibration-based adaptive image analysis procedure for in situ crystal size measurement. Crystal Growth & Design 23:7076-7089 (2023). https://doi.org/10.1021/acs.cgd.3c00273
26. Jul-Jørgensen I, Facco P, Gernaey KV, Barolo M, Hundahl CA. Data fusion of raman spectra in MSPC for fault detection and diagnosis in pharmaceutical manufacturing. Computers & Chemical Engineering 184:108647 (2024). https://doi.org/10.1016/j.compchemeng.2024.108647
27. Mesbah A, Huesman AEM, Kramer HJM, Van den Hof PMJ. A comparison of nonlinear observers for output feedback model-based control of seeded batch crystallization processes. Journal of Process Control 21:652-666 (2011). https://doi.org/10.1016/j.jprocont.2010.11.013
28. Porru M, Özkan L. Monitoring of batch industrial

crystallization with growth, nucleation, and agglomeration. part 2: structure design for state estimation with secondary measurements. Ind. Eng. Chem. Res. 56:9578-9592 (2017). https://doi.org/10.1021/acs.iecr.7b00243

29. Simone E, Saleemi AN, Nagy ZK. In situ monitoring of polymorphic transformations using a composite sensor array of raman, NIR, and atr-uv/vis spectroscopy, FBRM, and PVM for an intelligent decision support system. Org. Process Res. Dev. 19:167-177 (2014). https://doi.org/10.1021/op5000122

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Analysis of Ultrasound-Assisted Transesterification for Sustainable Biodiesel Production via Inline Raman spectroscopy

Ilias Bouchkira, Adel Mhamdi*

RWTH Aachen University, Process Systems Engineering (AVT.SVT), Aachen 52074, Germany
* Corresponding Author: Adel.Mhamdi@avt.rwth-aachen.de

ABSTRACT

We investigate ultrasound-assisted transesterification for biodiesel production. We use inline Raman spectroscopy to quantify its impact on reaction kinetics, catalyst reduction, and temperature sensitivity. We perform a systematic experimental study at different temperatures (50, 55, and 60 °C), different catalyst loadings, with and without ultrasound. The results show that ultrasound significantly accelerates early reaction kinetics at all temperatures, with the strongest effect observed at 55 °C, where both Fatty Acid Methyl Ester (FAME) formation rate and final conversion increase by up to 7 wt%. Under reduced catalyst conditions, ultrasound restores high conversion levels, leading to up to 20 wt% higher final FAME compared to operation without ultrasound and achieving performances comparable to, or exceeding, those obtained with (normal) catalyst without ultrasound. This is mainly because ultrasound primarily enhances mass transfer and phase contact, thereby reducing the system's sensitivity to catalyst loading. These findings demonstrate that ultrasound enables catalyst reduction while maintaining high biodiesel yields, enabling more sustainable and intensified transesterification processes.

Keywords: Biodiesel production, transesterification, ultrasound-assisted process intensification, catalyst optimization, Raman spectroscopy, process monitoring, clean energy systems

INTRODUCTION

Biodiesel is typically produced via the transesterification of triglycerides with short-chain alcohols. It offers advantages such as biodegradability, lower particulate emissions, and compatibility with existing diesel engines [1–3]. Despite its industrial maturity, conventional biodiesel production still faces several challenges, including high catalyst consumption, sensitivity to operating conditions, and mass-transfer limitations arising from the immiscibility between oil and alcohol phases [2, 3]. These limitations directly affect reaction kinetics, energy efficiency, and downstream purification costs, motivating the development of intensified and more sustainable processing routes.

Among the various process-intensification strategies proposed for biodiesel synthesis, ultrasound-assisted transesterification has attracted increasing attention over the past two decades [4, 5]. The application of ultrasound generates acoustic cavitation, leading to localized high temperatures and pressures, micro-jets, and intense micro-mixing at phase interfaces [4–6]. These phenomena enhance interfacial area and mass transfer between reactants, thereby accelerating reaction rates and potentially reducing the dependence on high catalyst loadings or elevated temperatures. Several studies have reported shorter reaction times and improved biodiesel yields when ultrasound is applied, particularly under kinetically or mass-transfer-limited conditions [7–8]. Nevertheless, most existing works focus on final conversion values or batch end-point analyses, providing limited insight into reaction evolution and the temperature-dependent role of ultrasound. Systematic studies quantifying the effect of ultrasound on biodiesel synthesis under different catalyst loadings remain limited.

Moreover, the potential of ultrasound to compensate for reduced catalyst loading, which is critical for improving the sustainability and economics of biodiesel production, has not been comprehensively quantified

under dynamic, inline-monitored conditions. Catalyst reduction is highly desirable, as it directly lowers raw material costs, reduces soap formation, and simplifies downstream separation. Yet, reducing catalyst loading often leads to slower kinetics and incomplete conversion, especially at lower temperatures or in the absence of efficient mixing. Understanding whether ultrasound can mitigate these drawbacks across different temperature regimes is therefore of both scientific and industrial relevance.

Raman spectroscopy has proven particularly suitable for quantifying process conditions due to its non-invasive nature, chemical specificity, and ability to operate under harsh reaction conditions. We demonstrated recently that inline Raman spectroscopy measurements can be employed successfully to monitor triglyceride consumption, fatty acid methyl ester (FAME) formation, and glycerol production during transesterification, offering unprecedented access to reaction kinetics and transient process behavior [9]. However, the combined use of ultrasound-assisted transesterification and inline Raman spectroscopy remains scarcely explored, and systematic studies coupling these two approaches are still lacking.

In this context, the objective of this work is to provide a systematic and quantitative analysis of ultrasound-assisted transesterification using inline Raman spectroscopy. To study the individual and combined effects of different parameters, we employ an experimental design combining three temperatures (50, 55, and 60 °C), two catalyst loadings (normal stoichiometry and 30% reduced), and the presence or absence of ultrasound. By continuously monitoring the evolution of key species using Raman spectroscopy, we aim to (i) study the temperature-dependent impact of ultrasound on reaction kinetics, (ii) assess the ability of ultrasound to compensate for reduced catalyst loading, and (iii) identify operating regimes where ultrasound-based process intensification delivers the greatest benefit. The insights gained contribute to the development of more energy-efficient, catalyst-lean, and digitally monitored biodiesel production processes.

MATERIALS AND METHODS

Materials

Refined vegetable oil is used as the triglyceride feedstock for all experiments. served as the alcohol reactant, while sodium methylate, a homogeneous alkaline catalyst, is employed to promote the transesterification reaction. Two catalyst loadings were investigated: a reference “normal” catalyst amount representative of conventional operating conditions, and a reduced catalyst loading corresponding to a 30% decrease relative to the reference case. This reduction level is selected to remain within an industrially relevant range while still inducing measurable kinetic limitations in the absence of process intensification.

All experiments were carried out using a fixed methanol-to-oil molar ratio, chosen sufficiently high to avoid equilibrium limitations and ensure comparability between runs. The reaction temperature is varied between 50 and and 60 °C (lower and higher temperatures were intentionally avoided to reach enough yield and minimize methanol evaporation).

Experimental Setup

For the semi-batch transesterification of vegetable oil into biodiesel, we use the experimental setup illustrated in Fig. 1. The core of the system is a continuously stirred lab-scale reactor (1) with a nominal working volume of 500 mL. The reactor is jacketed and connected to a thermostatically controlled heat circulator (2) (Julabo GmbH, Germany), allowing precise regulation of the reaction temperature. The temperature inside the reactor is continuously monitored using an immersed PT100 temperature sensor (3).

The feedstock solutions, refined vegetable oil and a pre-mixed methanol/sodium methylate catalyst solution, are stored in separate supply tanks (6 and 7). Each stream is delivered to the reactor via independent high-precision metering pumps (4 and 5) (Knauer GmbH, Germany), enabling accurate control of the injection strategy and the methanol-to-oil molar ratio during semi-batch operation. The reaction mixture is agitated using an overhead mechanical stirrer (8) (Heidolph GmbH, Germany) to ensure effective macroscopic mixing and homogeneous temperature and composition throughout the liquid phase.

Ultrasound irradiation is provided by an external ultrasonic device (11) mechanically coupled to the reactor vessel. The system operates at a fixed low ultrasonic frequency characteristic of cavitation-dominated regimes and is activated continuously during ultrasound-assisted experiments. For reference experiments, the ultrasound unit is completely switched off, while mechanical stirring is maintained in all cases to provide a consistent baseline mixing condition and enable direct comparison between mechanically stirred and ultrasound-enhanced systems.
Real-time process monitoring is achieved using an inline Raman spectroscopy system consisting of a Raman spectrometer (12) (tec5 AG, Germany) equipped with a 785 nm laser source and an immersion Raman probe (9). The probe is mounted directly through the reactor headplate, allowing continuous, non-invasive measurement of the reacting mixture without the need for sampling or offline analysis. The probe position is selected to avoid direct exposure to intense cavitation zones or interfaces while ensuring representative sampling of the liquid phase.

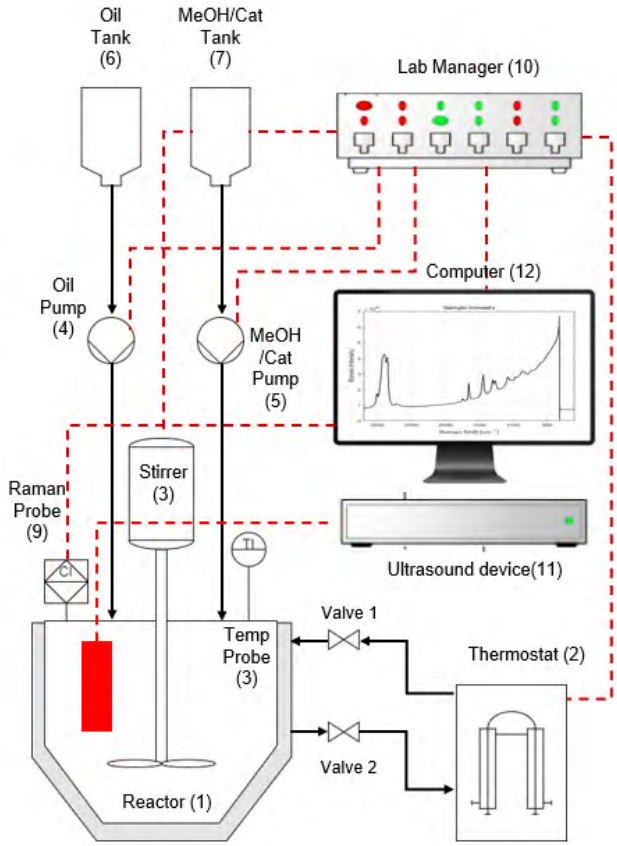

Figure 1: Schematic of the experimental setup showing the semi-batch reactor, mechanical stirring, ultrasound transducer, temperature control, and inline Raman.

All system components, including pumps, temperature control, and Raman data acquisition, are supervised by a custom LabManager control platform (10). This software coordinates sensor data acquisition, actuator control, and real-time visualization, ensuring reproducible operating conditions and synchronized process monitoring throughout each experiment.

Design of Experiments

We adopt an experimental design to investigate systematically the effects of ultrasound, catalyst loading, and temperature on transesterification performance. Three temperatures (50, 55, and 60 °C), two catalyst loadings (normal and 30% reduced), and two operating modes (with and without ultrasound) were combined, resulting in a comprehensive set of operating conditions (Table 1). For clarity, the experiments are categorized into four main cases:

(A) Ultrasound applied with normal catalyst loading.

(B) Ultrasound applied with reduced catalyst loading.

(C) No ultrasound with normal catalyst loading.

(D) No ultrasound with reduced catalyst loading.

This classification enables interesting pairwise comparisons, such as: (i) the effect of ultrasound at constant catalyst loading, (ii) the effect of catalyst reduction under ultrasound, (iii) the role of catalyst loading without ultrasound, and (iv) the ability of ultrasound to compensate for reduced catalyst loading.

Table 1: Summary of experimental conditions with temperature, catalyst loading, ultrasound condition, and case designation (A–D)

Case	T (°C)	Catalyst Loading	Ultra-sound
A	50, 55, 60	Normal	Yes
B	50, 55, 60	Reduced (−30%)	Yes
C	50, 55, 60	Normal	No
D	50, 55, 60	Reduced (−30%)	No

Raman Spectroscopy and Data Processing

Raman spectra were collected continuously throughout each experiment with a constant acquisition time, providing high temporal resolution of the transesterification process. Characteristic Raman bands associated with triglycerides, fatty acid methyl esters (FAME), methanol, and glycerol were used to quantify the evolution of key species during the reaction, following established methodologies explained in our previous work [9]. In summary, a multivariate calibration model, developed and validated using reference samples covering the expected composition range, is employed to convert Raman spectra into quantitative concentration profiles. This approach enabled direct tracking of triglyceride consumption and biodiesel (FAME) formation without the need for offline sampling.

The resulting time-resolved composition profiles were used to analyze reaction kinetics, compare conversion rates across operating conditions, and evaluate final biodiesel yields. All Raman-based measurements are synchronized with temperature and ultrasound operation logs to ensure consistent interpretation of dynamic effects.

Data Analysis and Performance Metrics

Reaction performance is evaluated using both kinetic and conversion-based indicators. The rate of FAME formation and triglyceride consumption during the initial reaction period is used as a measure of kinetic enhancement, while final FAME yield and residual oil content were used to assess overall conversion efficiency.

Comparisons between experimental cases were performed at identical temperatures to isolate the influence of ultrasound and catalyst loading. Attention is

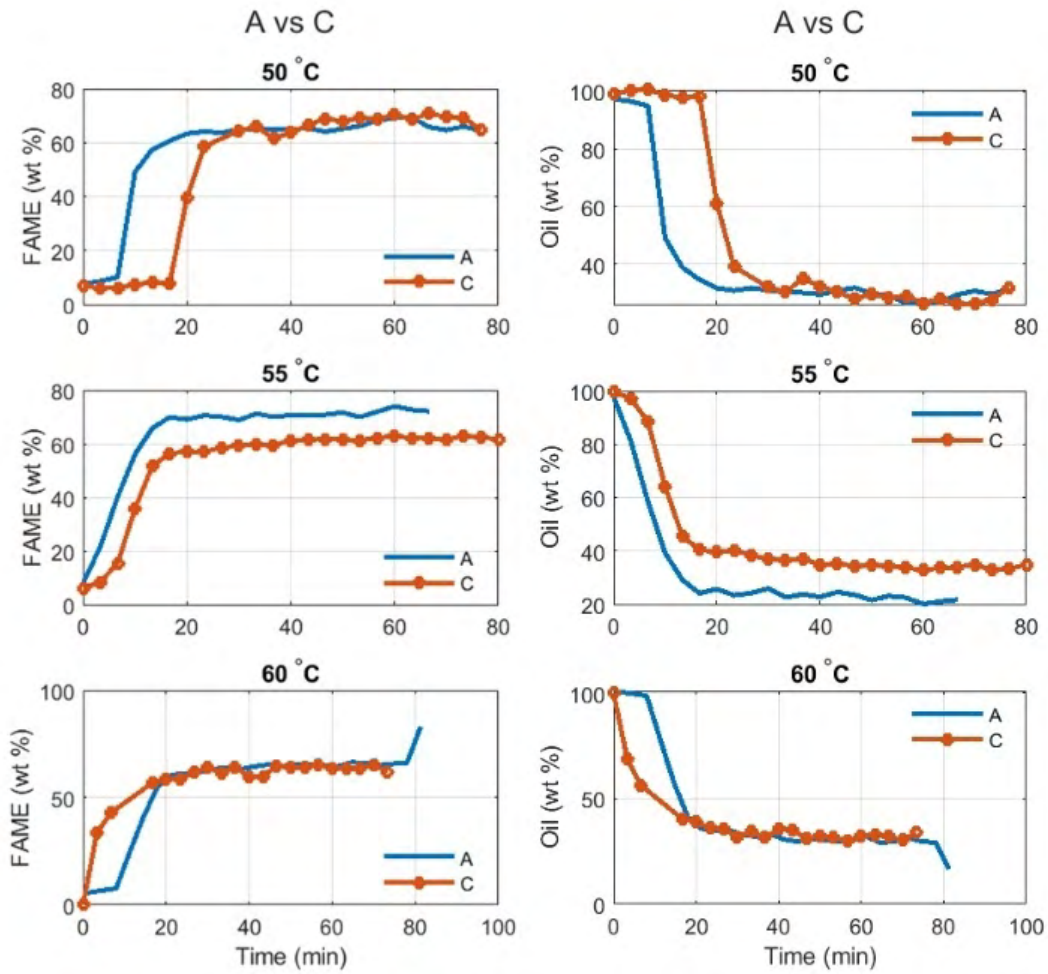

Figure 2. Effect of ultrasound at reduced catalyst

given to identifying temperature regimes where ultrasound-induced process intensification is most pronounced and to assessing whether ultrasound could compensate for the loss of catalytic activity under reduced catalyst conditions.

RESULTS AND DISCUSSION

Effect of Ultrasound at Normal Catalyst Loading

The effect of ultrasound at normal catalyst loading is assessed by comparing experiments performed with and without ultrasound at identical temperatures (Cases A and C), based on FAME formation and oil consumption (Figure 2). At 60 °C, ultrasound accelerates the initial reaction rate, but final FAME yields and residual oil contents remain similar, indicating that thermal activation already dominates the kinetics. At 55 °C, ultrasound has a pronounced and sustained effect, significantly enhancing both the rate and extent of FAME formation by overcoming mass-transfer limitations. At 50 °C, ultrasound mainly improves the early reaction stage, with faster initial FAME formation and oil consumption, while final conversions are comparable.

Effect of Ultrasound at Reduced Catalyst Loading

The influence of ultrasound under reduced catalyst loading is evaluated by comparing experiments performed with and without ultrasound at identical temperatures (Cases B and D), using FAME formation and oil consumption as performance indicators (Figure 3). At 60 °C, ultrasound leads to a substantial improvement in reaction performance, with much faster FAME formation, higher final yields, and more oil conversion, effectively compensating for the reduced catalytic activity observed without ultrasound. At 55 °C, ultrasound continues to enhance kinetics and conversion, although the performance gap is slightly smaller than at 60 °C. At 50 °C, the effect of ultrasound is particularly pronounced: while the non-ultrasound case remains strongly kinetically limited,

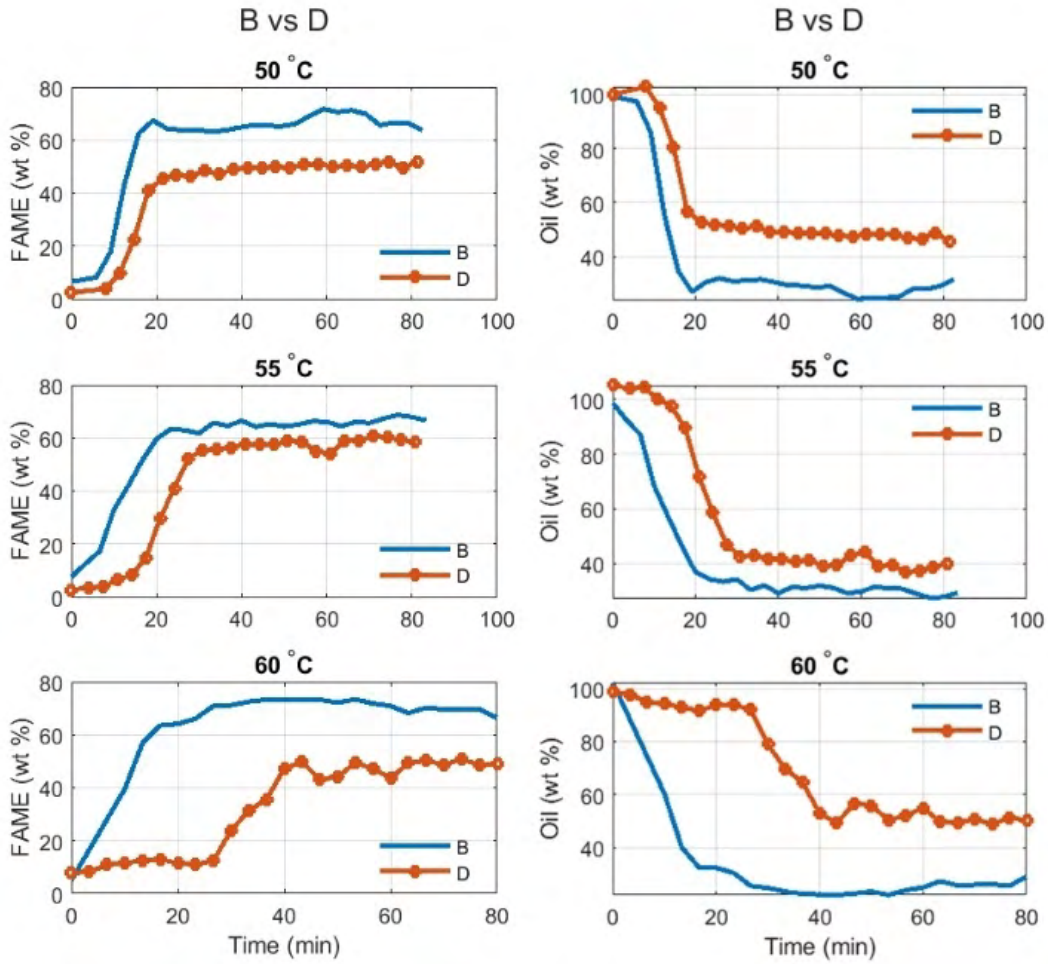

Figure 3. Effect of ultrasound at reduced catalyst

ultrasound-assisted operation achieves rapid oil consumption and high final FAME yields.

Effect of Catalyst Amount under Ultrasound

The effect of catalyst loading under ultrasound-assisted conditions is assessed by comparing experiments performed with normal and reduced catalyst amounts while ultrasound is applied (Cases A and B), based on FAME formation and oil consumption (Figure 4). At 60 °C, normal catalyst loading results in slightly faster kinetics and deeper oil conversion, although the performance gap is much smaller than in the absence of ultrasound, indicating reduced sensitivity to catalyst concentration. At 55 °C, the influence of catalyst loading becomes moderate, with both cases converging toward similar final FAME yields and residual oil contents. At 50 °C, catalyst reduction has almost no effect, as FAME formation and oil consumption profiles are nearly identical. The results show thus, that under ultrasound-assisted operation, catalyst loading mainly affects reaction rate rather than final conversion, particularly at intermediate and low temperatures.

Effect of Catalyst Amount without Ultrasound

The influence of catalyst loading in the absence of ultrasound is evaluated by comparing experiments performed with normal and reduced catalyst amounts under identical thermal and mixing conditions (Cases C and D), using FAME formation and oil consumption as indicators (Figure 5). At 60 °C, normal catalyst loading results in markedly faster kinetics, deeper oil conversion, and higher final FAME yields, whereas reduced catalyst loading leads to slow reaction rates and incomplete conversion. At 55 °C, the effect of catalyst loading remains evident, although the performance gap between the two cases is slightly reduced. At 50 °C, catalyst reduction has a particularly severe impact, with strongly kinetically limited behavior and poor oil consumption compared to the normal catalyst case.

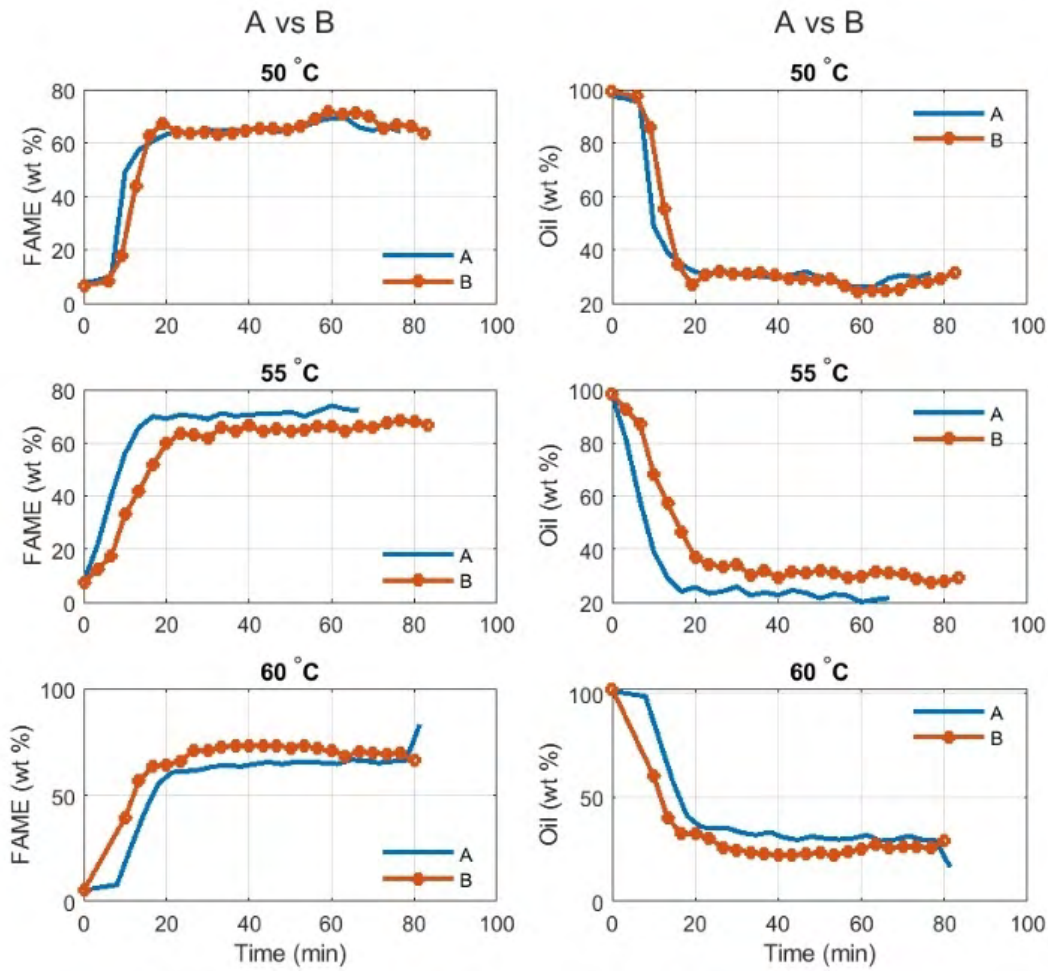

Figure 4. Effect of catalyst amount under Ultrasound

Can Ultrasound Compensate for Reduced Catalyst Loading?

The ability of ultrasound to compensate for reduced catalyst loading is assessed by comparing experiments conducted with reduced catalyst under ultrasound to those performed with normal catalyst without ultrasound (Cases B and C), based on FAME formation and oil consumption (Figure 6). At 60 °C, both cases reach very similar final FAME yields and residual oil contents, while the ultrasound-assisted reduced-catalyst case exhibits faster initial kinetics, indicating effective compensation of lower catalytic activity. At 55 °C, the compensation effect is even more pronounced, as the reduced-catalyst case with ultrasound slightly outperforms the normal-catalyst case without ultrasound in terms of final FAME yield and oil conversion. At 50 °C, both scenarios converge toward similar final conversion levels, although ultrasound consistently accelerates early-stage FAME formation and oil consumption.

Best and Worst Operating Scenarios

The combined influence of ultrasound and catalyst loading is evaluated by comparing the best operating condition—normal catalyst loading with ultrasound (Case A)—to the worst condition—reduced catalyst loading without ultrasound (Case D), based on FAME formation and oil consumption (Figure 7). At 60 °C, the performance gap is substantial: the best-case scenario shows much faster FAME formation and deeper oil conversion, while the worst-case scenario remains strongly kinetically limited and reaches significantly lower final conversion. At 55 °C, this disparity remains pronounced, with higher FAME yields and lower residual oil content under ultrasound and normal catalyst loading. At 50 °C, the contrast becomes even more striking, as reduced catalyst without ultrasound leads to severe kinetic limitations, whereas ultrasound-assisted operation with normal catalyst maintains high conversion efficiency.

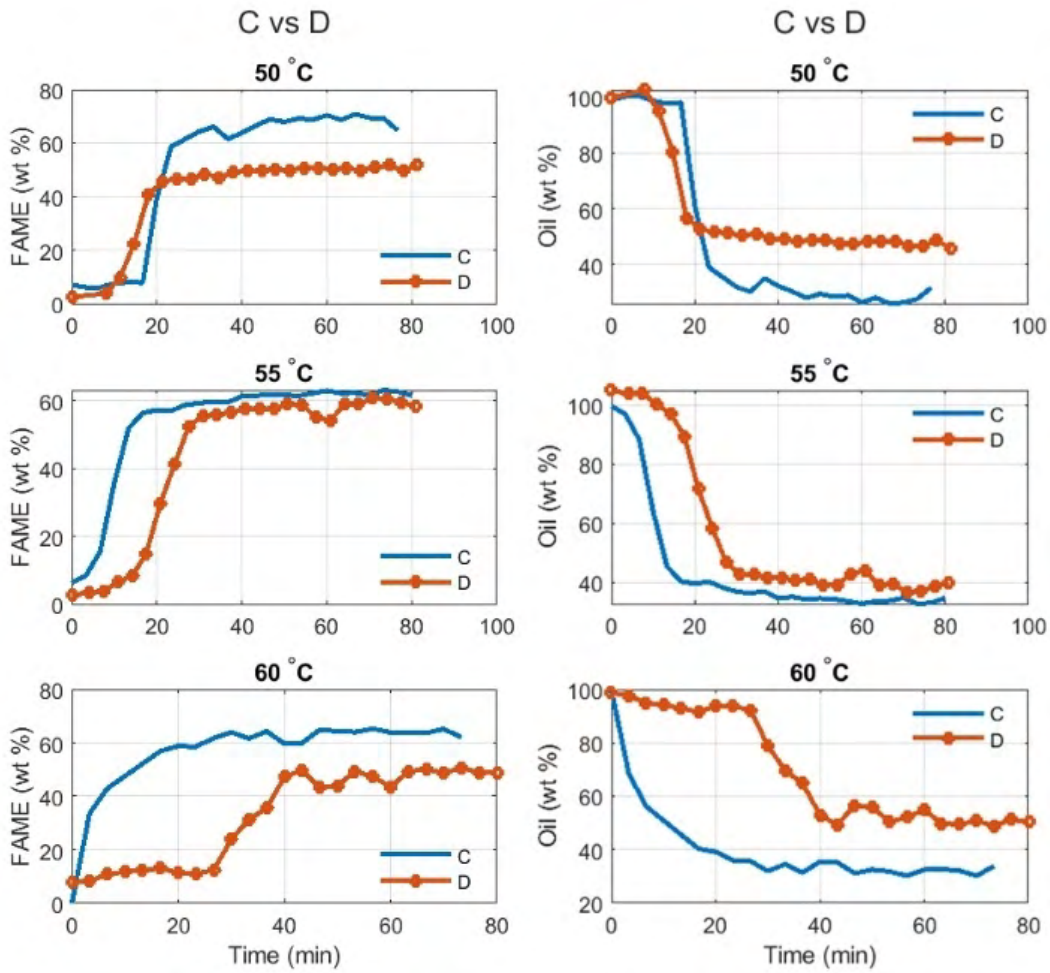

Figure 5: Eeffect of catalyst amount without ultrasound

CONCLUSIONS AND PERSPECTIVES

This work provides a systematic analysis of ultrasound-assisted transesterification for biodiesel production using inline Raman spectroscopy to monitor the dynamic evolution of FAME formation and triglyceride consumption under different temperatures and catalyst loadings. The results demonstrate that ultrasound acts as an effective process-intensification strategy, accelerating reaction kinetics and significantly reducing the sensitivity of the process to catalyst loading. At normal catalyst conditions, ultrasound primarily enhances the early stages of the reaction, with the strongest and most sustained impact observed at intermediate temperature (55 °C). At higher temperature (60 °C), the additional benefit of ultrasound becomes more limited due to dominant thermal activation, while at lower temperature (50 °C) ultrasound mainly reduces the time required to reach comparable final conversion levels.

Under reduced catalyst loading, ultrasound plays a particularly important role. Without ultrasound, catalyst reduction leads to strong kinetic limitations and incomplete oil conversion across all temperatures. In contrast, ultrasound-assisted operation consistently restores high conversion levels, achieving faster oil consumption and higher FAME formation. In several cases, reduced catalyst combined with ultrasound achieves performances comparable to, or even exceeding, those obtained with normal catalyst loading without ultrasound. These findings confirm that ultrasound can effectively compensate for catalyst reduction while maintaining high biodiesel production efficiency.

From a Process Systems Engineering perspective, the combination of ultrasound intensification and inline Raman monitoring provides a valuable framework for data-rich process understanding and model development. The time-resolved composition data generated in this work constitute an important basis for developing phenomenological kinetic models capable of explicitly

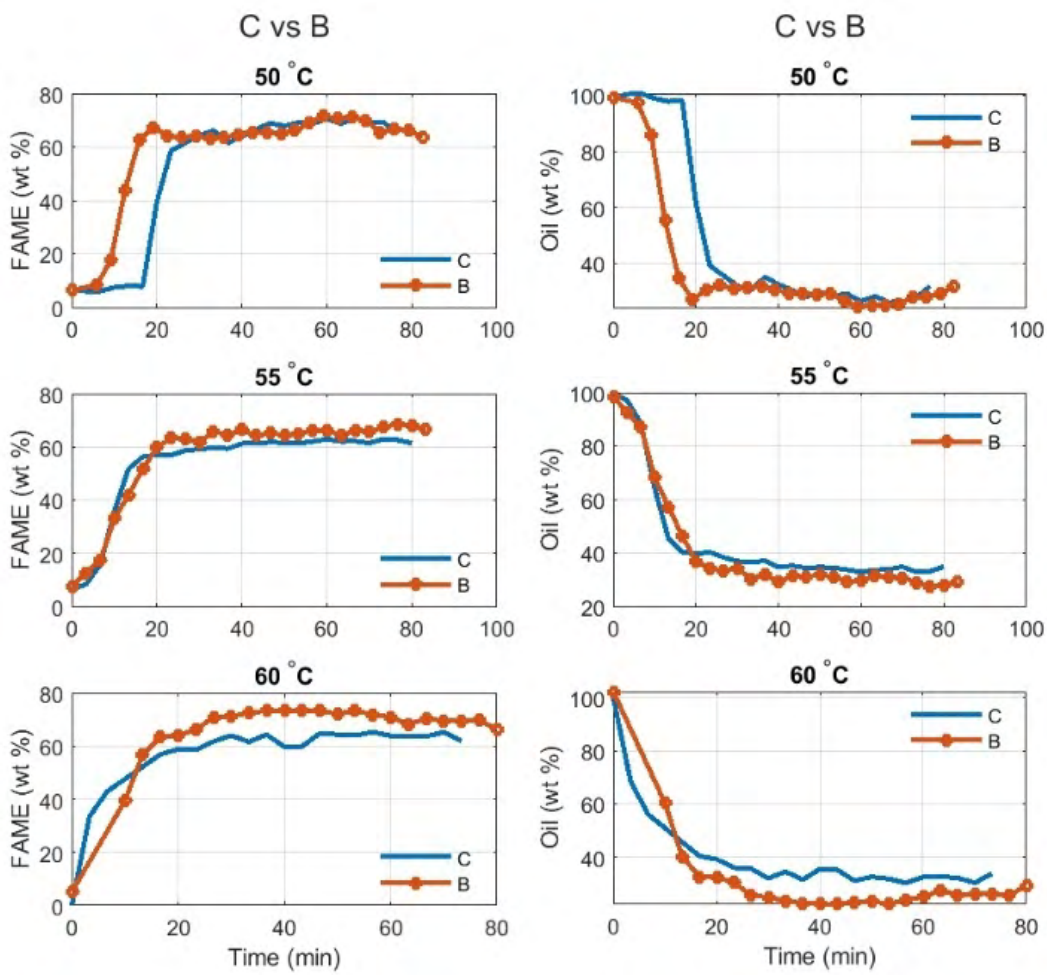

Figure 6. Can ultrasound compensate reduced catalyst?

capturing the influence of ultrasound-induced mass-transfer enhancement on transesterification dynamics. Such models can then be integrated into process simulation and optimization frameworks to identify optimal operating conditions balancing catalyst usage, temperature, and energy consumption.

Moreover, the use of inline Raman spectroscopy enables continuous measurement of key species, which opens the possibility of implementing advanced monitoring and model-based control strategies. Integrating ultrasound-assisted transesterification within digitalized process architectures, including soft sensors, real-time optimization, or digital twins, could further enhance process efficiency and robustness.

Future work will therefore focus on extending the current experimental insights toward first-principles modeling, model-based optimization, and control of ultrasound-intensified biodiesel reactors, as well as evaluating the scalability and energy efficiency of ultrasound-assisted systems under continuous industrial operation. These developments would contribute to the design of intensified, catalyst-lean, and digitally enabled biodiesel production processes.

ACKNOWLEDGEMENTS

We thank Martin Kleber and Wolfgang Hoffmann (Cargill GmbH, Frankfurt) for the fruitful technical discussions on the operating scenarios and for providing the chemicals used in the experimental campaigns. We also thank the financial support of the Kopernikus project SynErgie by the Federal Ministry of Research, Technology and Space (BMFTR) and the project supervision by the project management organization Projektträger Jülich (PtJ).

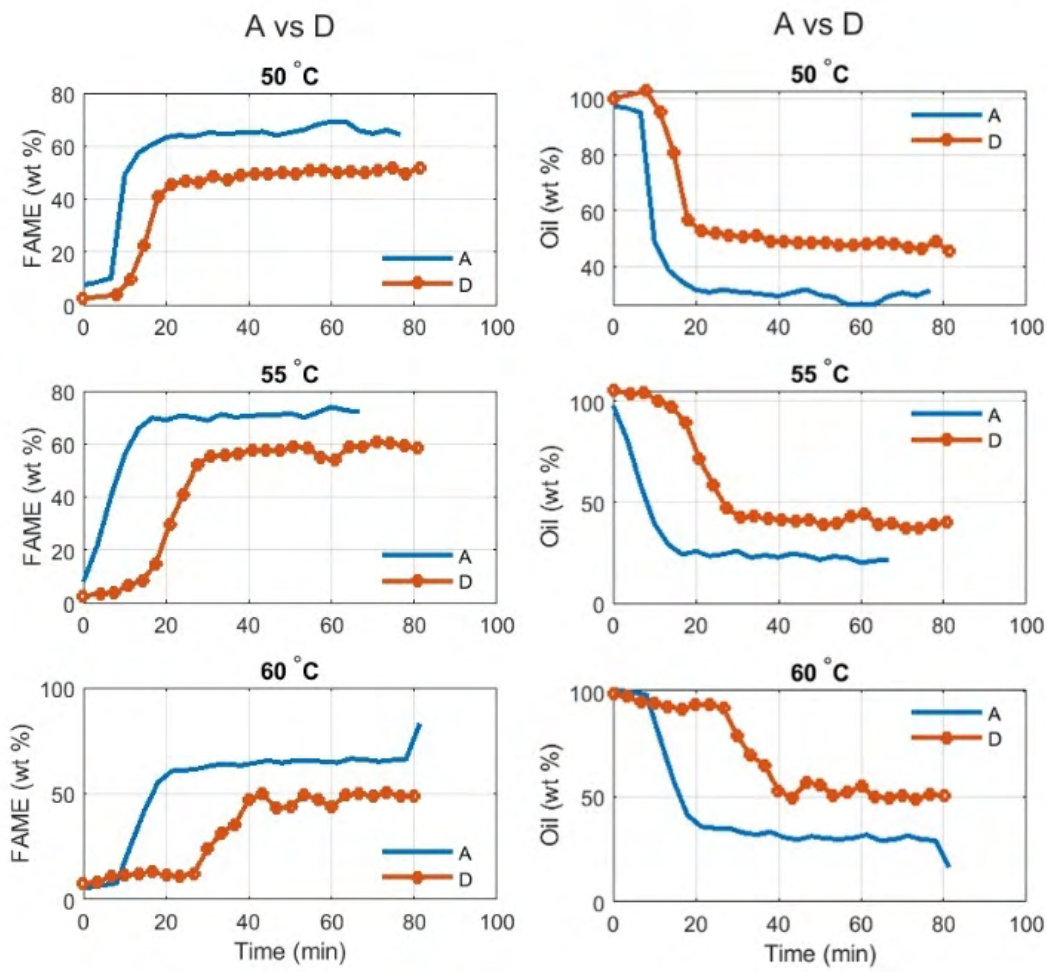

Figure 7: Best and worst operating scenarios

REFERENCES

1. Negm NA, Abou Kana MTH, Youssif MA, Mohamed MY. Biofuels from vegetable oils as alternative fuels: advantages and disadvantages. Surfactants in Tribology, Volume 5 :289-367 (2017). https://doi.org/10.1201/9781315120829-18
2. Mathew GM, Raina D, Narisetty V, Kumar V, Saran S, Pugazhendi A, Sindhu R, Pandey A, Binod P. Recent advances in biodiesel production: challenges and solutions. Science of The Total Environment 794:148751 (2021). https://doi.org/10.1016/j.scitotenv.2021.148751
3. Ambat I, Srivastava V, Sillanpää M. Recent advancement in biodiesel production methodologies using various feedstock: a review. Renewable and Sustainable Energy Reviews 90:356-369 (2018). https://doi.org/10.1016/j.rser.2018.03.069
4. Tan SX, Lim S, Ong HC, Pang YL. State of the art review on development of ultrasound-assisted catalytic transesterification process for biodiesel production. Fuel 235:886-907 (2019). https://doi.org/10.1016/j.fuel.2018.08.021
5. Ho WWS, Ng HK, Gan S. Advances in ultrasound-assisted transesterification for biodiesel production. Applied Thermal Engineering 100:553-563 (2016). https://doi.org/10.1016/j.applthermaleng.2016.02.058
6. Tan SX, Lim S, Ong HC, Pang YL. State of the art review on development of ultrasound-assisted catalytic transesterification process for biodiesel production. Fuel 235:886-907 (2019). https://doi.org/10.1016/j.fuel.2018.08.021
7. Mostafaei M, Javadikia H, Naderloo L. Modeling the effects of ultrasound power and reactor dimension on the biodiesel production yield: comparison of prediction abilities between response surface methodology (RSM) and adaptive neuro-fuzzy

inference system (ANFIS). Energy 115:626-636 (2016). https://doi.org/10.1016/j.energy.2016.09.028

8. Prakash Maran J, Priya B. Modeling of ultrasound assisted intensification of biodiesel production from neem (azadirachta indica) oil using response surface methodology and artificial neural network. Fuel 143:262-267 (2015). https://doi.org/10.1016/j.fuel.2014.11.058
9. Bouchkira I, El Wajeh M, Mhamdi A. Modeling of biodiesel production via transesterification using inline raman spectroscopy. IFAC-PapersOnLine 59:600-605 (2025). https://doi.org/10.1016/j.ifacol.2025.08.011

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Evaluating and adapting modelling strategies for data-driven prediction of solvent effects on reaction barriers

Daeun Shin[ab], Lingfeng Gui[c], Jonggeol Na[d], Won Bo Lee[a], and Lauren Ye Seol Lee[b]*

[a] Department of Chemical and Biological Engineering, Seoul National University, Seoul 08826, Republic of Korea
[b] Department of Chemical Engineering, University College London, Torrington Place, London, WC1E 7JE, United Kingdom
[c] School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, Scotland EH14 4AS, United Kingdom
[d] Department of Chemical Engineering and Materials Science, Ewha Womans University, Seoul 03760, Republic of Korea
* Corresponding Author: lauren.lee@ucl.ac.uk.

ABSTRACT

Predicting solvent effects on reaction activation barriers is central to understanding chemical reactivity and reaction kinetics, and guiding solvent selection. The solvent-induced change in activation free energy ($\Delta\Delta G^{\ddagger}_{\mathrm{solv}}$) provides a quantitative descriptor of this effect, but remains costly to evaluate across vast reaction-solvent spaces, using quantum mechanical methods. Recent data-driven models have enabled prediction of solvent effects. However, most typically rely on two-dimensional representation of reactions and do not explicitly encode sufficient reaction context, such as transition-state information, or three-dimensional structural changes along the reaction, resulting in limited generalizability and predictive accuracy. In this study, systematic evaluation is presented of modelling strategies for predicting $\Delta\Delta G^{\ddagger}_{\mathrm{solv}}$, with a focus on the role of reaction-state representation, input-geometry fidelity, and input modality. Using a large reaction-solvent dataset, models based on two-dimensional condensed reaction graphs are compared with models incorporating three-dimensional geometries of reactants, transition states, and products. The sensitivity of geometry-based models to structural accuracy is assessed by replacing quantum-chemically optimized transition states with structures predicted by a generative model. In addition, a dual-modality architecture combining two-dimensional graph-based and three-dimensional geometry-based representations is examined. The results show that explicit inclusion of both reactant and transition-state geometries leads to improved prediction accuracy relative to representations based on reaction endpoints or transition states alone. However, model performance depends strongly on the fidelity of the input geometries, with substantial degradation observed when low-quality structures are used. The dual-modality approach partially mitigates this sensitivity by adaptively reweighting two-dimensional and three-dimensional information, leading to performance recovery under low-fidelity conditions.

Keywords: Solvent effect, Solvation free energy of reaction, Transition state, 3D geometry, Multi-modality

1. INTRODUCTION

Solvents play a central role in chemical reactions because preferential stabilization of transition state (TS) relative to the reactant (R) can shift reaction energy barriers and alter reaction rate and selectivity. This solvent effect can be quantified by the solvation contribution to the activation free energy, defined as $\Delta\Delta G^{\ddagger}_{\mathrm{solv}} = \Delta G^{\mathrm{TS}}_{\mathrm{solv}} - \Delta G^{\mathrm{R}}_{\mathrm{solv}}$. The solvation free energy of reaction, $\Delta\Delta G^{\ddagger}_{\mathrm{solv}}$ can serves as a key performance descriptor, making it possible to compare solvent effects across reaction-solvent combinations and therefore useful for solvent selection. Traditionally, $\Delta\Delta G^{\ddagger}_{\mathrm{solv}}$ has been calculated via quantum mechanics/molecular mechanics (QM/MM) simulations or implicit solvation models, including such as the Solvation Model based on Density (SMD) and the COnductor-like Screening Model for Real Solvent (COSMO-RS). While these methods can provide reliable estimates, their high computational cost often limits scalability to a large and diverse reaction-solvent space [1]. In particular,

locating and optimising TS structures represents a major bottleneck, making systematic evaluation across many reactions and solvents impractical [2, 3].

In recent years, data-driven approaches have emerged as efficient alternatives by enabling rapid estimation of solvent effects at substantially reduced computational cost. Early work in this area has largely focused on predicting solvation energy of individual molecules [4-6] and extension to reaction kinetics have been more limited. Existing approaches for modelling solvent effect on activation energies broadly fall into two classes. The first class consider a single fixed reaction, treating solvent as the sole variable [7, 8]. While this yields meaningful results within specific reaction classes, it faces a fundamental limitation in generalizability across diverse reaction-solvent combinations. The second class aims to model general reaction-solvent combinations by using both reaction and solvent as inputs, and hence covering a broader chemical space [9, 10].

Recent progress in the latter direction has been driven by the effectiveness of two-dimensional (2D) reaction representations for modelling reaction properties, such as activation energies and reaction yields [10-12]. These representations include condensed reaction graph (CGR) [13], which is obtained by superposing 2D molecular graphs of R and product (P), and reaction fingerprint, which is generated by open-source software RDKit [14, 15] or deep learning models [16, 17]. Such representations have been used as input and achieved high predictive accuracy by effectively encoding structural changes resulting from reactions. This early success presents a promising modelling strategy of solvent effect on activation energy by combining 2D reaction representations with molecular representations of solvent. For example, Chung and Green [9] developed a data-driven model that can predict solvent-induced changes in activation energy utilizing CGR and 2D molecular graph of solvent. By integrating reaction and solvent embeddings independently encoded by directed message passing neural network (D-MPNN) encoders, the model achieved state-of-the-art performance in both $\Delta\Delta G^{\ddagger}_{\mathrm{solv}}$ calculated by COSMO-RS and experimental relative rate constant (k_{rel}) prediction. However, these representations primarily capture topological differences between R and P and do not explicitly encode information along the reaction coordinate or structural features of the TS. Moreover, 2D representations cannot uniquely determine three-dimensional (3D) properties, such as charge redistribution and spatial arrangement of functional groups, which are known to influence solvation and activation energetics. As a result, current models may struggle to generalize across reaction types and capture key physical factors governing solvent effects on reaction barriers. While employing TS geometries as input can provide richer information of reaction context, this has been considered impractical due to the difficulties in computation, as we described. To address this limitation, generative models that predict TS geometries have recently been actively proposed as a practical alternative [18-21]. These models generate TS geometries from relatively accessible inputs, e.g. CGR and geometries of R and P, at a fraction of the computational cost of QM methods. However, to the best of our knowledge, the effectiveness of explicit incorporation of the low-cost generated TS geometries for modelling solvent effect on activation energy has yet to be explored.

In this study, we develop systematic modelling strategies for $\Delta\Delta G^{\ddagger}_{\mathrm{solv}}$ prediction that explicitly address these limitations. Rather than proposing a new architecture, we analyse how predictive performance depends on three controlled aspects of model input design: (i) the reaction-state(s) used to represent the reaction (R, TS, and/or P), (ii) the fidelity of the associated 3D geometries, and (iii) the representation modality, including 2D graph-based, 3D geometry-based, and combined dual-modality inputs. Using a large reaction-solvent dataset derived from COSMO-RS calculations, we compare $\Delta\Delta G^{\ddagger}_{\mathrm{solv}}$prediction performance on unseen reactions of a conventional 2D model using CGR with models that explicitly incorporate geometries of various input configurations for reactants, transition states, and products. By replacing quantum-chemically optimised structures with geometries locally optimised by force field or predicted from a generative model, the sensitivity of the model to geometry quality is assessed. Dual-modality approaches combining 2D graph-based and 3D geometry-based representations are implemented to further explore the robustness of the models depending on fidelity of input geometry.

2 METHOD

2.1 Molecular and reaction representation

To investigate how input design influences model performance, we varied three factors systematically: (i) the state(s) of the reaction path encoded as input (R, TS, and/or P), (ii) the fidelity of the associated 3D geometries, and (iii) the modality of the representation i.e., whether the model used 2D graph-based inputs, 3D geometry-based inputs, or both.

For 2D representations, reactions were encoded as CGR [13], which represent the transformation between reactants and products within a single graph. CGRs were generated from atom-mapped SMILES using Chemprop [22]. Solvents were represented as standard 2D molecular graphs. Initial atom and bond features for both CGRs and solvent graphs were constructed using standard RDKit descriptors, including atomic number, bond order, formal charge, and aromaticity.

To incorporate explicit structural information along the reaction coordinate, 3D molecular geometries were

used to construct reaction representations based on different states of the reaction path. Three reaction-path configurations were considered (Figure 1**a**): (1) an R-P representation, which encodes the reaction using the 3D geometries of reactants (R) and products (P); (2) a TS-only representation, which encodes the reaction using only the 3D geometry of the transition state (TS); and (3) R-TS representation, which encodes the reaction using the 3D geometries of both reactants (R) and transition states (TS). These configurations enable a controlled comparison of how explicit inclusion of TS information and reaction-path context affects model behaviour.

Input geometries for 3D-based representations were further grouped by fidelity to evaluate sensitivity to structural accuracy. *High-fidelity* geometries used R, P, and TS structures optimised in the gas phase at the ωB97X-D3/def2-TZVP level of theory from the Grambow dataset [23]. *Moderate-fidelity* geometries used R and P structures from the Grambow dataset, while TS geometries were predicted using React-OT from the corresponding high-fidelity R and P structures. *Low-fidelity* geometries used R and P structures generated using OpenBabel [24] and locally optimised with the generalized Amber force field (GAFF), and TS geometries predicted by React-OT from these lower-cost R and P structures. Note that React-OT is a flow-matching-based generative model that predicts TS structures from interpolated R and P geometries [19] and was selected due to its deterministic predictions and computational efficiency. For solvents, 3D structures generated using the OpenBabel force field were used consistently across all fidelitiles.

2.2 Dataset for model training and evaluation

We used the $\Delta\Delta G^{\ddagger}_{\text{solv}}$ dataset proposed by Chung and Green [9]. In that work, COSMO-RS calculations were performed using optimised gas-phase geometries of R and TS from the Grambow dataset, spanning 500, 000 reaction-solvent combinations with 13, 227 unique reactions and 295 solvents. The dataset both contains neutral closed-shell and free radical reactions, restricted to systems composed of H, C, O, and N with at most seven heavy atoms (11, 876 closed-shell; 1, 351 radical). To evaluate generalization to unseen reactions, 5% of the reaction were randomly selected and reserved as a test set. The remaining data were randomly split into a 90% training set and a 10% validation set.

2.3 Model architecture

Based on the molecular and reaction representations defined, three model architectures, namely, a 2D graph-based baseline model, a 3D geometry-based equivariant model, and a dual-modality model combining both representations. All architectures were adopted from existing frameworks and adapted to ensure a consistent comparison across input configurations.

2D-graph-based baseline model: As a baseline, we adopted a directed message passing neural network (D-MPNN) architecture [25] previously used for reaction-solvent modelling. This model independently encodes for the reaction and solvent by passing the 2D CGR and 2D solvent graph through separate D-MPNNs. A four-layer D-MPNN was used for the CGR to capture changes in atom and bond features associated with reaction, while a two-layer D-MPNN was used for the solvent graph owing to its comparatively simple structure. The resulting reaction and solvent embeddings were concatenated and passed to a feed-forward predictor consisting of six fully connected layers to generate the final prediction.

Geometry-based equivariant model: We adopted TorchMD-Net [26] as an SE(3)-equivariant encoder. TorchMD-Net takes atomic numbers and Cartesian coordinates as inputs and constructs representations using attention mechanisms based on interatomic distance and relative position vectors, without relying on explicit bond connectivity. An eight-layer equivariant encoder was applied to all molecular species. The representations obtained from the encoder for R, TS, and P were concatenated and passed through two fully connected layers to form reaction-level embedding. This embedding was then concatenated with the solvent representation and passed through the same six-layer feed-forward predictor used in the baseline model.

Duality-modality model: To evaluate whether combining 2D and 3D representations improves robustness [27, 28], we implemented a dual-modality architecture that integrates the D-MPNN-based 2D encoders and the TorchMD-Net-based SE(3)-equivariant 3D encoder. In this setting, representations obtained independently from the 2D and 3D streams were combined using a late-fusion strategy. Learnable scalar weights were applied to each modality stream, allowing the model to adjust their relative contributions during training. Specifically, two scalar parameters, i.e., w_{2D} and w_{3D}, initialized to zero, were used. These weights are passed through a softmax function to normalise them and multiplied by the embeddings from each modality stream to obtain a fused representation. The fused representation was then passed to the same feed-forward predictor used in the single-modality models (Figure 4**a**).

All models were trained for 80 epochs using a batch size of 30. Optimisation was performed with the Adam optimiser and Noam learning-rate scheduler, with an initial learning rate of 1.8×10^{4}. Mean squared error (MSE) was used as the loss function. Model implementation was performed in PyTorch 2.4.0, and training was conducted using PyTorch Lightning 1.8.6.

3. RESULTS

3.1 Effect of including TS geometries

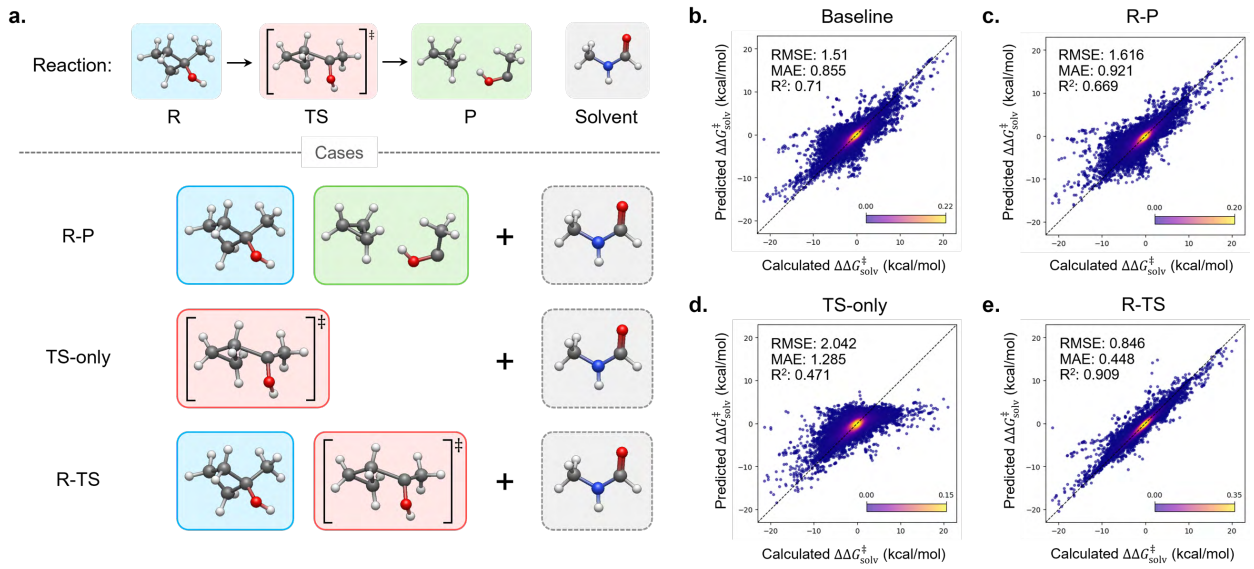

Figure 1. Overview of input configurations and prediction performance for different representations. (a) Schematic of the three 3D input configurations considered in this study: R-P, which uses the geometries of R and P; TS-only, which uses only the TS geometry; and R–TS, which uses the geometries of both R and TS. In all cases, the solvent is represented separately and combined with the reaction representation. (b-e) Parity plots comparing predicted and calculated $\Delta\Delta G^{\ddagger}_{\text{solv}}$ values for (b) the 2D graph-based baseline model, (c) the R-P model, (d) the TS-only model, and (e) the R–TS model. Root mean square error (RMSE), mean absolute error (MAE), and coefficient of determination (R^2) are reported for each model.

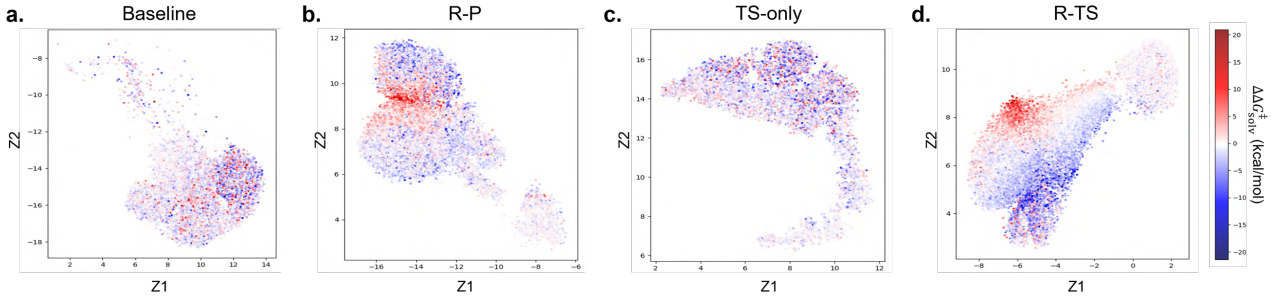

Figure 2. UMAP visualisations of the learned reaction-solvent representations for different input configurations. Latent embeddings are shown for (a) the 2D graph-based baseline model and the 3D models using (b) R-P, (c) TS-only, and (d) R-TS. Points are coloured by the corresponding $\Delta\Delta G^{\ddagger}_{\text{solv}}$ values.

To evaluate the effect of including TS geometries on $\Delta\Delta G^{\ddagger}_{\text{solv}}$ prediction, we compared the model performance using three distinct combinations of R, TS, and P geometries defined in Section 2.1—R-P, TS-only, and R-TS, with the 2D-based baseline model. The input configurations are summarised in Figure 1**a**, and all models in this comparison used high-fidelity geometries from the Grambow dataset.

Notably, the model using 3D geometries of R and P (R-P) exhibited lower prediction accuracy than the 2D-based baseline (Figure 1**b, c**). This result indicates that providing endpoint geometries alone does not improve prediction of solvent-induced activation free energy shifts. Although 3D molecular structures contain detailed geometric information, the absence of explicit TS information limits their relevance for describing kinetically meaningful structural changes. As a result, replacing 2D CGR inputs with 3D R–P geometries increase input complexities without improving representation of features directly associated with the target quantity.

The model using only TS geometry (TS-only) showed significantly lower prediction performance compared to the baseline, even though TS is a critical determinant of reaction kinetics (Figure 1**d**). This result likely stems from the fact that understanding reaction process requires not only TS geometry but also R structure as a reference state of change. The omission of the information of the reference state has led to overfitting, confirming that TS structure alone is insufficient to capture the relative structural evolution that is necessary to

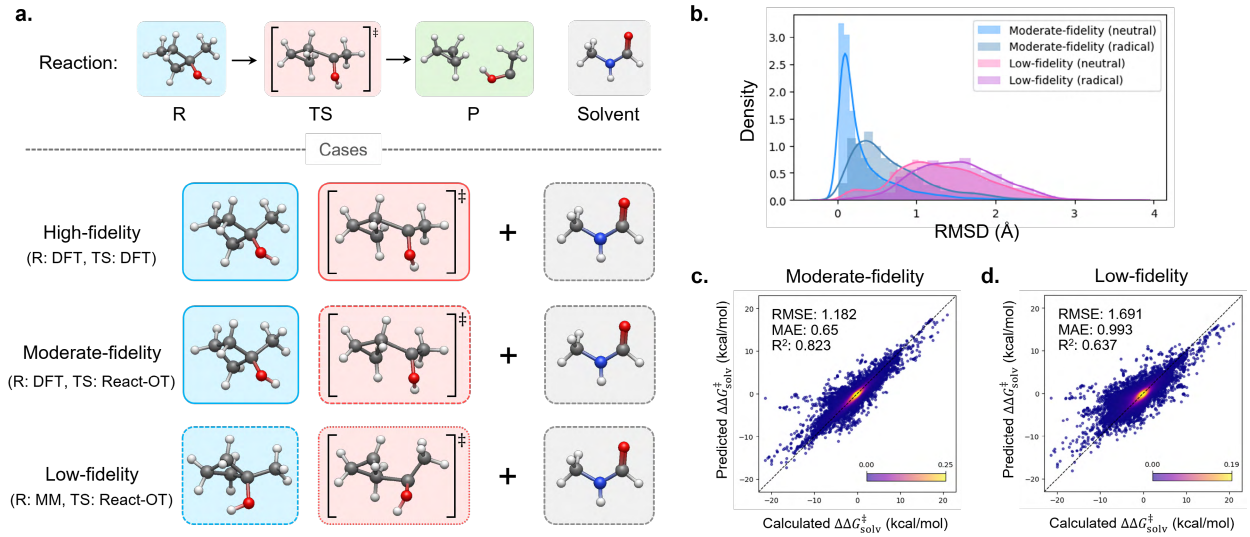

Figure 3. Effect of input structure fidelity on TS geometry accuracy and prediction performance. (a) Schematic illustration of the three input-geometry fidelity levels considered for the R–TS configuration. (b) Distributions of TS root-mean-square deviations (RMSDs) relative to the DFT reference structures for moderate- and low-fidelity cases, shown separately for neutral and radical reactions. (c,d) Parity plots comparing predicted and calculated

define a reaction path.

In contrast, incorporating both R and TS 3D geometries (R-TS) achieved the highest prediction accuracy (Figure 1**e**). This indicates that the model implicitly captured path dependent features, such as geometric distortion and charge redistribution, from the structural data of R and TS. Thus, explicitly incorporating R and TS 3D structures enables the model to encode structural differences between R and TS, which are directly relevant to solvent-induced changes.

To further examine how different input configurations influence the learned representation, the final reaction-solvent embeddings were visualised using Uniform Manifold Approximation and Projection (UMAP) [29]. For the baseline model and the TS-only model, the latent representations were broadly scattered and showed no clear patterns with respect to the target values (Figure2**a, c**). The R-P model showed a locally clustered region for reaction-solvent combinations with high $\Delta\Delta G^{\ddagger}_{\mathrm{solv}}$, but remained disordered for moderate to small values (Figure 2**b**), consistent with its larger prediction errors observed in this region. In contrast, the latent space learned by the R-TS model exhibited the clearest global ordering with respect to $\Delta\Delta G^{\ddagger}_{\mathrm{solv}}$, characterised by a continuous gradient across the embedding (Figure 2**d**). This suggest that the combined use of R and TS geometries yields representations more strongly associated with solvent-induced activation free energy shifts.

3.2 Effect of input structure fidelity

In previous section, we showed that incorporating both R and TS geometries improves prediction performance when high-fidelity structures are available. To assess the robustness of this approach under more practical conditions, we next evaluated how prediction performance depends on the fidelity of the input geometries, using the three structure-fidelity levels defined in Section 2.1.

As summarised in Figure 3**a**, model performance was compared across high-, moderate-, and low-fidelity input geometries. In the high-fidelity case, both R and TS geometries were taken from DFT-optimised structures in the Grambow dataset. In the moderate- and low-fidelity cases, TS geometries were generated using React-OT from progressively lower-fidelity R and P structures, as described in Section 2.1. Consistent with this construction, the TS RMSD relative to the DFT reference structures was smaller in the moderate-fidelity case and substantially larger in the low-fidelity case (Figure 3**b**). The broader RMSD distributions observed for radical reactions further indicate increased difficulty in accurately predicting TS geometries for open-shell systems.

Prediction accuracy decreased systematically as input structure fidelity was reduced (Figure 3**c, d**). In the low-fidelity case, the model performed worse than the 2D-based baseline, indicating that TS geometries generated from inaccurate R and P structures do not provide reliable information for learning solvent-induced changes in the activation free energy. In contrast, the moderate-fidelity case maintained superior performance relative to the baseline, demonstrating that predicted TS geometries can still be informative when the underlying R and P structures preserve essential geometric features.

Overall, these results show that models explicitly

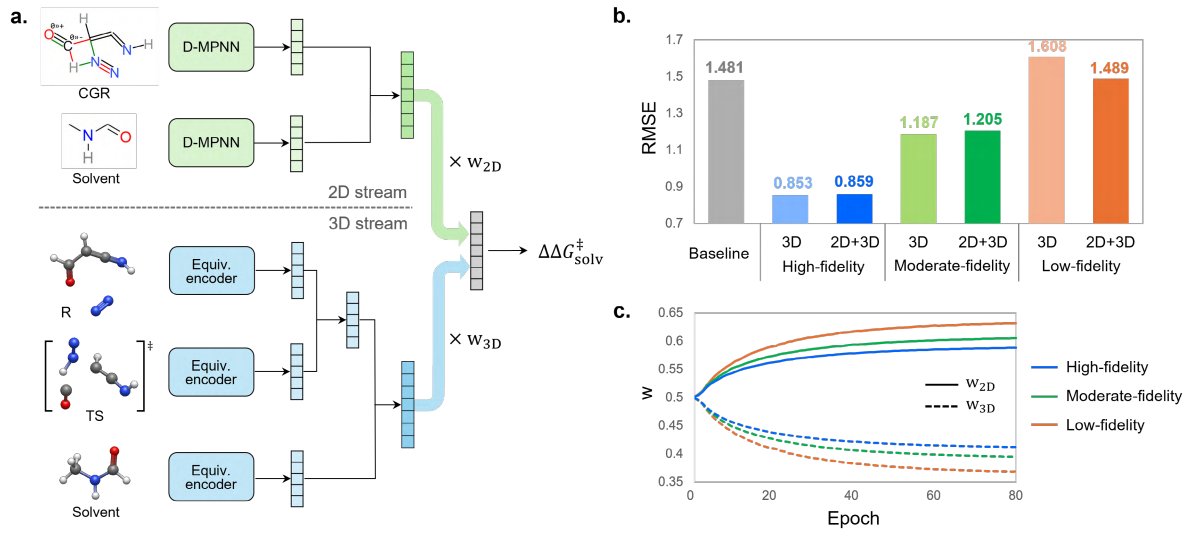

Figure 4. Dual-modality model architecture and performance across input-geometry fidelity levels. (a) Schematic illustration of the dual-modality architecture integrating 2D graph-based representations and 3D geometry-based representations. (b) Prediction performance, measured by RMSE, for the 2D-based baseline model, the 3D-only model, and the dual-modality (2D + 3D) model across high-, moderate-, and low-fidelity input geometries. (c) Evolution of the learned modality weights for the 2D and 3D streams during training for each fidelity level.

incorporating R and TS geometries are sensitive to the fidelity of the input structures. Nevertheless, the performance observed in the moderate-fidelity case indicates that practical modelling strategies based on low-cost TS generation are feasible, provided that the predicted structures retain key geometric characteristics of the activated complex.

3.3 Effect of input modality

Because prediction performance of 3D-based models depends strongly on input-geometry fidelity, we next evaluated whether combining 2D and 3D representations can improve robustness to reduced structural accuracy. Specifically, we examined a dual-modality approach that integrates 2D graph-based representations with 3D geometry-based representations, as defined in Section 2.3.

The dual-modality model integrates 2D graph-based and 3D geometry-based representations using the late-fusion architecture described in Section 2.3 (Figure 4**a**). Model performance was evaluated for high-, moderate-, and low-fidelity input geometries and compared against the corresponding 3D-only and 2D-based baseline models.

For the high- and moderate-fidelity cases, the dual modality model showed only marginal difference relative to the 3D-only model as can be seen in Figure 4**b**. This suggests that when reliable 3D structures are available, the 3D representation alone captures sufficient information for $\Delta\Delta G^{\ddagger}_{\mathrm{solv}}$, limiting the additional contribution of 2D graph-based inputs. In contrast, in the low-fidelity case, the dual modality model recovered prediction performance to a level comparable to the 2D-based baseline and outperformed the 3D-based model (Figure 4**b**). Analysis of the learned modality weights (w) provides further insight into this behaviour (Figure 4**c**). In the low-fidelity case, the weight assigned to the 2D modality increased progressively during training while the 3D weight decreased, indicating a shift toward graph-based representations as the reliability of 3D geometries decreased. Conversely, in the high-fidelity case, the modality weights stabilised early and the gap between w_{2D} and w_{3D} remained smallest, indicating a more balanced contribution of the two modalities.

Overall, these results indicate that dual-modality models can improve robustness to reduced input-geometry fidelity by adaptively reweighting information from different representations. The observation that performance in the low-fidelity case does not exceed that of the 2D-based baseline highlights a limitation of the current late-fusion strategy, suggesting that more effective integration schemes may be required to achieve synergistic gains from combined 2D and 3D inputs.

4. CONCLUSIONS

We presented systematic modelling strategies for solvent-induced change in reaction activation energy, specifically for $\Delta\Delta G^{\ddagger}_{\mathrm{solv}}$ prediction. Comprehensive analysis of prediction based on thee controlled aspects of model input design was performed: (i) the reaction-state(s) used to represent the reaction, (ii) the fidelity of the associated input geometries, and (iii) the

representation modality of 2D, 3D, and combined. The significant improvement in predictive accuracy was observed through explicit incorporation of the geometries of both R and TS. Furthermore, the visualised latent space learned by the R-TS model exhibited the clearest trend with respect to $\Delta\Delta G^{\ddagger}_{\text{solv}}$. This proves that the structural differences encoded from the R and TS geometries yield representations strongly relevant to solvent effect on activation free energy. The performance of the 3D geometry-based model was strongly influenced by input geometry fidelity. When low-fidelity geometries were used, the model showed degraded performance compared to the 2D graph-based baseline model. However, the performance under the moderate-fidelity case maintained superior performance relative to the baseline. This indicates the feasibility of the practical modelling strategies based on predicted TS geometries and concurrently highlights the importance of preserving key geometric characteristics of the activated complex. Finally, we implemented a dual modality approach that integrates 2D graph-based and 3D geometry-based representations to investigate the robustness of the model to reduced fidelity of input geometries. Under the low-fidelity cases, the model exhibited a notable recovery of prediction performance to a level comparable to the 2D-based baseline and outperformed the 3D-based model. Further analysis provided valuable insight that the improved robustness resulted from the adaptive reweighting of the information from different modalities based on their reliability. Nevertheless, the current simple late-fusion strategy failed to facilitate synergistic interaction between modalities, limiting the performance in the low-fidelity case to a similar level to the 2D-based baseline. This highlights that future works should focus on developing multi-modal frameworks that enable complementary integration of information using sophisticated techniques, such as contrastive learning and cross-attention mechanisms. We believe that such advancement can provide a practical strategy of modelling solvent-induced changes in reaction barriers by harnessing the benefits of 3D geometries while enhancing the robustness from noise-resistant inputs.

ACKNOWLEDGEMENTS

This research was supported by Korea Institute for Advancement of Technology (KIAT) grant funded by the Ministry of Trade, Industry & Energy (MOTIE), Korea Government (RS-2024-00436106, Human Resource Development Program for Industrial Innovation).

AUTHOR IDENTIFIERS

Author ORCIDs:
Shin D: 0009-0000-9044-8113
Gui L: 0000-0003-1957-1957
Na J: 0000-0002-1106-9500
Lee WB: 0000-0001-7801-083X
Lee YS: 0000-0003-1847-7838

REFERENCES

1. Spiekermann KA, Pattanaik L, Green WH. Fast predictions of reaction barrier heights: toward coupled-cluster accuracy. J. Phys. Chem. A 126:3976-3986 (2022). https://doi.org/10.1021/acs.jpca.2c02614
2. Zhao Q, Hsu HH, Savoie BM. Conformational sampling for transition state searches on a computational budget. J. Chem. Theory Comput. 18:3006-3016 (2022). https://doi.org/10.1021/acs.jctc.2c00081
3. Jackson R, Zhang W, Pearson J. Tsnet: predicting transition state structures with tensor field networks and transfer learning. Chem. Sci. 12:10022-10040 (2021). https://doi.org/10.1039/d1sc01206a
4. Ferraz-Caetano J, Teixeira F, Cordeiro MNDS. Explainable supervised machine learning model to predict solvation gibbs energy. J. Chem. Inf. Model. 64:2250-2262 (2023). https://doi.org/10.1021/acs.jcim.3c00544
5. Chung, Y., et al., Group contribution and machine learning approaches to predict Abraham solute parameters, solvation free energy, and solvation enthalpy. Journal of Chemical Information and Modeling, 2022. **62**(3): p. 433-446.
6. Low K, Coote ML, Izgorodina EI. Explainable solvation free energy prediction combining graph neural networks with chemical intuition. J. Chem. Inf. Model. 62:5457-5470 (2022). https://doi.org/10.1021/acs.jcim.2c01013
7. Gui L, Yu Y, Oliyide TO, Siougkrou E, Armstrong A, Galindo A, Sayyed FB, Kolis SP, Adjiman CS. Integrating model-based design of experiments and computer-aided solvent design. Computers & Chemical Engineering 177:108345 (2023). https://doi.org/10.1016/j.compchemeng.2023.108345
8. Struebing H, Ganase Z, Karamertzanis PG, Siougkrou E, Haycock P, Piccione PM, Armstrong A, Galindo A, Adjiman CS. Computer-aided molecular design of solvents for accelerated reaction kinetics. Nature Chem 5:952-957 (2013). https://doi.org/10.1038/nchem.1755
9. Chung Y, Green WH. Machine learning from quantum chemistry to predict experimental solvent effects on reaction rates. Chem. Sci. 15:2410-2424 (2024). https://doi.org/10.1039/d3sc05353a
10. Jorner K, Brinck T, Norrby PO, Buttar D. Machine

learning meets mechanistic modelling for accurate prediction of experimental activation energies. Chem. Sci. 12:1163-1175 (2021). https://doi.org/10.1039/d0sc04896h
11. Schwaller P, Vaucher AC, Laino T, Reymond JL. Prediction of chemical reaction yields using deep learning. Mach. Learn.: Sci. Technol. 2:015016 (2021). https://doi.org/10.1088/2632-2153/abc81d
12. Probst D, Schwaller P, Reymond JL. Reaction classification and yield prediction using the differential reaction fingerprint DRFP. Digital Discovery 1:91-97 (2022). https://doi.org/10.1039/d1dd00006c
13. Heid E, Green WH. Machine learning of reaction properties via learned representations of the condensed graph of reaction. J. Chem. Inf. Model. 62:2101-2110 (2021). https://doi.org/10.1021/acs.jcim.1c00975
14. Dobbelaere MR, Lengyel I, Stevens CV, Van Geem KM. Rxn-insight: fast chemical reaction analysis using bond-electron matrices. J Cheminform 16: (2024). https://doi.org/10.1186/s13321-024-00834-z
15. van Gerwen P, Briling KR, Calvino Alonso Y, Franke M, Corminboeuf C. Benchmarking machine-readable vectors of chemical reactions on computed activation barriers. Digital Discovery 3:932-943 (2024). https://doi.org/10.1039/d3dd00175j
16. Schwaller P, Probst D, Vaucher AC, Nair VH, Kreutter D, Laino T, Reymond JL. Mapping the space of chemical reactions using attention-based neural networks. Nat Mach Intell 3:144-152 (2021). https://doi.org/10.1038/s42256-020-00284-w
17. Mswahili ME, Jeong YS. Transformer-based models for chemical SMILES representation: a comprehensive literature review. Heliyon 10:e39038 (2024). https://doi.org/10.1016/j.heliyon.2024.e39038
18. Kim S, Woo J, Kim WY. Diffusion-based generative AI for exploring transition states from 2D molecular graphs. Nat Commun 15: (2024). https://doi.org/10.1038/s41467-023-44629-6
19. Duan C, Liu GH, Du Y, Chen T, Zhao Q, Jia H, Gomes CP, Theodorou EA, Kulik HJ. Optimal transport for generating transition states in chemical reactions. Nat Mach Intell 7:615-626 (2025). https://doi.org/10.1038/s42256-025-01010-0
20. Duan C, Du Y, Jia H, Kulik HJ. Accurate transition state generation with an object-aware equivariant elementary reaction diffusion model. Nat Comput Sci 3:1045-1055 (2023). https://doi.org/10.1038/s43588-023-00563-7
21. Choi S. Prediction of transition state structures of gas-phase chemical reactions via machine learning. Nat Commun 14: (2023). https://doi.org/10.1038/s41467-023-36823-3
22. Heid, E., et al., *Chemprop: a machine learning package for chemical property prediction.* Journal of Chemical Information and Modeling, 2023. **64**(1): p. 9-17.
23. Grambow, C.A., L. Pattanaik, and W.H. Green, Reactants, products, and transition states of elementary chemical reactions based on quantum chemistry. Scientific data, 2020. **7**(1): p. 137.
24. O'Boyle NM, Banck M, James CA, Morley C, Vandermeersch T, Hutchison GR. Open babel: an open chemical toolbox. J Cheminform 3: (2011). https://doi.org/10.1186/1758-2946-3-33
25. Yang K, Swanson K, Jin W, Coley C, Eiden P, Gao H, Guzman-Perez A, Hopper T, Kelley B, Mathea M, Palmer A, Settels V, Jaakkola T, Jensen K, Barzilay R. Analyzing learned molecular representations for property prediction. J. Chem. Inf. Model. 59:3370-3388 (2019). https://doi.org/10.1021/acs.jcim.9b00237
26. Thölke, P. and G. De Fabritiis, Torchmd-net: equivariant transformers for neural network based molecular potentials. arXiv preprint arXiv:2202.02541, 2022.
27. Luo, Y., et al., *Molfm: A multimodal molecular foundation model.* arXiv preprint arXiv:2307.09484, 2023.
28. Yu, Q., et al., *Multimodal molecular pretraining via modality blending.* arXiv preprint arXiv:2307.06235, 2023.
29. McInnes, L., J. Healy, and J. Melville, *UMAP: uniform manifold approximation and projection for dimension reduction. arXiv.* arXiv preprint arXiv:1802.03426, 2018. **10**.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

GlycoPy: An Equation-Oriented and Object-Oriented Python Framework for Process Modeling, Optimization and Optimal Control

Yingjie Ma[ab*], Jing Guo[bc], and Richard D. Braatz[a]

[a] Nanjing University, School of Robotics and Automation, Suzhou, Jiangsu Province, China
[b] Massachusetts Institute of Technology, Department of Chemical Engineering, Cambridge, MA, USA
[c] Polytechnique Montréal, Department of Chemical Engineering, Montréal, QC, Canada
* Corresponding Author: yingjie.ma@nju.edu.cn.

ABSTRACT

Nonlinear model predictive control (NMPC) can substantially improve performance and constraint handling for (bio)chemical processes, but its adoption is still limited by the effort required to build maintainable first-principles models and to implement efficient dynamic optimization-based controllers. This paper presents GlycoPy, an open-source, equation-oriented and object-oriented Python framework that supports hierarchical model construction and integrated workflows for simulation, parameter estimation, dynamic optimization, and NMPC. The case study of the monoclonal antibody glycosylation process based on a multiscale model demonstrates the capability of GlycoPy.

Keywords: Nonlinear Model Predictive Control, Dynamic Modelling, Optimization, Simulation, Multiscale Modelling, software

1. INTRODUCTION

Equation-oriented (EO) environments are well suited to modeling and optimization of (bio)chemical processes: they let researchers develop new models without implementing solvers, and they provide highly accurate derivatives through automatic differentiation [1]. Yet, given the complexity of real-world systems, existing EO software often struggles to balance user-friendliness, developer productivity, flexibility, and extensibility [2]. Commercial EO tools (e.g., gPROMS [3]) provide rich modeling capabilities but are not designed for tight integration with advanced nonlinear model predictive control (NMPC) workflows or for rapid prototyping of bespoke algorithms. Open-source packages such as GEKKO [4], do-mpc [5] and SolACE [6] offer convenient optimization and NMPC, yet lack native support for hierarchical model composition, which is important for multiscale and multi-unit systems. Pyomo supports hierarchical modeling and many advanced process systems engineering (PSE) tasks [7]. However, it is difficult for its users to develop and integrate novel simulation algorithms, which are often essential for optimization workflows driven by computationally expensive simulations.

We therefore introduce GlycoPy, an equation-oriented Python framework built on CasADi to address these gaps. CasADi provides the symbolic infrastructure for efficient numerical optimal control, although it is not itself a domain-specific language for large process models and control [8]. GlycoPy emphasizes (1) object-oriented, hierarchical modeling to ease development and improve user-friendliness and maintainability; (2) advanced PSE functionalities, including parameter estimation, dynamic optimization, and NMPC with and without parameter adaptation; (3) control-vector parameterization (CVP) approach for dynamic optimization, which is typically more robust for large and strongly nonlinear model-constrained problems; and (4) direct access to CasADi symbolic variables and equations, enabling rapid prototyping of customized algorithms. GlycoPy is released as open source.

In the remainder of this paper, we first present the software and then validate its NMPC capability by controlling monoclonal antibody (mAb) glycosylation using a multiscale model.

2. SOFTWARE DESCRIPTION

GlycoPy is built on CasADi together with NumPy/Pandas/SciPy/Matplotlib. CasADi provides symbolic variables and equations and enables efficient differential algebraic equation (DAE) and ordinary differential equation (ODE) simulations and automatic differentiation, which are essential for large-scale dynamic optimization. Compared to native CasADi, GlycoPy provides a structured framework that streamlines model development and algorithm implementation. The following sections provide more details on how GlycoPy is used for modeling, simulation, parameter estimation, dynamic optimization, and NMPC.

2.1 Nonlinear modeling

In GlycoPy, a model is defined by implementing the `build()` method of the Model class using two basic objects: Variable and Equation. GlycoPy supports three variable types: differential states x, algebraic variables z, and parameters p (time-invariant or time-varying but fixed during integration). Variables are created via add_var, which returns a GlycoPy symbol (a CasADi symbol or a dictionary of CasADi symbols). Users then write equations directly using standard CasADi-style operators and register them with `add_eq` as either ODE (`ode`) or algebraic (`alg`) equations.

To illustrate, we consider a small nonlinear DAE benchmark adapted from Andersson et al. (2019).

$$\begin{cases} \frac{dx_1(t)}{dt} = z(t)\, x_1(t) - x_2(t) + u(t) \\ \frac{dx_2(t)}{dt} = x_1(t) \\ \frac{d\mathrm{obj}}{dt} = x_1(t)^2 + x_2(t)^2 + u^2 \\ z(t) = p - x_2(t)^2 \end{cases} \quad t \in [0, T] \tag{1}$$

$$x_1(0) = 0, x_2(0) = 1, \mathrm{obj}(0) = 0, p = 1, \tag{2}$$

where $x \in \mathbb{R}^2$ denotes the differential variables, $z \in \mathbb{R}$ is the algebraic variable, and $\mathrm{obj} \in \mathbb{R}$ is the accumulated term to be used as the objective function in optimization. $u \in \mathbb{R}$ is the control action, and $T = 10$ is the considered time horizon. The GlycoPy implementation for the problem is shown below.

```
# Define model
Class SimpleDAE(Model):
 def build(self):
  xg = self.add_var('x', 'xg', m=2, val=[0.,
1.])
  obj = self.add_var('x', 'obj', val=0.)
  zg = self.add_var('z', 'zg')
  pg = self.add_var('p', 'pg', val=1.0)
  ug = self.add_var('u', 'ug', val=0.)
  self.add_eq('ode', 'xg', [zg * xg[0] -
xg[1] + ug, xg[0]])
  self.add_eq('ode', 'obj', xg[0]**2 +
xg[1]**2 + ug**2)
  self.add_eq('alg', 'zg', zg - (1 -
xg[1]**2) * pg)
```

Because each GlycoPy model is a Python class, users can naturally apply object-oriented design (inheritance and composition) and construct hierarchical models by assembling submodels. This modular structure improves development, testing, and maintenance for large chemical/bioprocess systems.

2.2 Simulation

For practical process models, time-varying inputs/parameters (e.g., piecewise-constant feed profiles) make a native CasADi workflow using built-in integrator cumbersome, especially when some input profiles are fixed while others are to be optimized. GlycoPy addresses this with the `SimulationDyn` class. With that, the code for the simulation of above DAE example is:

```
# Define simulator
initializer = Initializer('init', model)
integrator = Integrator('integrator',
'idas', dae_prob, initializer=initializer)
sim = SimulationDyn(model, integrator,
p_var=[('simple_dae.pg', 0), ])
p_var = np.random.randn(n_points-1, 1)
# Simulate
sim.simulation_with_events_numeri-
cal(t_event=time_points, p_var=p_var)
```

For numerical simulation, users specify which parameters are time-varying and provide their values at event times. The method `simulation_with_events_numerical(t_event, p_var)` runs the simulation and stores all results back into the model for convenient plotting and post-processing.

For sensitivity analysis and optimization, numerical trajectories are insufficient because automatic differentiation and symbolic Jacobians require symbolic results. Therefore, GlycoPy also provides `simulation_with_events`, which accepts symbolic input profiles and stores the resulting trajectories as CasADi expressions. Users can then directly build CasADi functions and compute Jacobians (or other derivatives) using standard CasADi operations. To further simplify robust derivative computation, GlycoPy includes `DerivativeEvaluatorSimple`, which wraps sensitivity evaluation with explicit method selection (`forward` or `adjoint`) and a safeguard: if sensitivity equations fail (a common issue for stiff, large-scale DAEs during optimization), it automatically falls back to finite differences. This design improves usability and stability for optimization workflows. The code for derivative computation is:

```
# Define derivative evaluator
evaluator = DerivativeEvaluatorSimple(y,
        p_var_sym,
  gradient_method=GradientMethod.forward)
# Evaluator derivatives
y_val = evaluator.val_fun_py(p_var)
jac_val = evaluator.jac_fun_py(p_var)
```

Finally, symbolic simulation enables developing differentiable simulation algorithms that are difficult to implement in many equation-oriented tools, such as the quasi-steady-state algorithm as will be shown later on in the case study section.

2.3 Dynamic optimization

A standard dynamic optimization problem (optimal control problem, OCP) shown below minimizes an integral (stage cost) plus a terminal cost subject to DAE dynamics, initial/terminal conditions, and bounds on states and inputs: differential states $x(t)$, algebraic states $z(t)$, control inputs $u(t)$, and parameters p.

$$\min_{x(\bullet),z(\bullet),u(\bullet)} \int_0^T L(x(t),z(t),u(t),p)dt + E(x(T),z(T),p)$$

$$\text{s.t.} \quad \left.\begin{aligned} \dot{x}(t) &= \text{ode}(x(t),z(t),u(t),p), \\ 0 &= \text{alg}(x(t),z(t),u(t),p), \\ u(t) &\in \mathcal{U}, x(t) \in \mathcal{X}, z(t) \in \mathcal{Z}, \end{aligned}\right\}, t \in [0,T]$$

$$x(0) \in \mathcal{X}_0, x(T) \in \mathcal{X}_T, z(0) \in \mathcal{Z}_0, z(T) \in \mathcal{Z}_T$$

Here, $\mathcal{U}, \mathcal{X}, \mathcal{Z}$ are given sets that the corresponding variables belong to. To solve the OCP problem numerically, the continuous-time problem is transcribed into a finite-dimensional nonlinear programming (NLP) problem:

$$\min_w f(w)$$

$$\text{s.t.}\ g_{lb} \le g(w) \le g_{ub},$$

$$w_{lb} \le w \le w_{ub},$$

where w collects decision variables (e.g., discretized controls and possibly some initial states), and $g(w)$ includes the discretized dynamics and path/terminal constraints.

To solve the OCP problem numerically, the continuous-time problem is transcribed into a finite-dimensional NLP. GlycoPy currently uses the CVP method because of its robust convergence for highly nonlinear dynamics-constrained problems. CVP-based optimization repeatedly performs DAE simulations and (for derivative-based solvers) sensitivity/derivative evaluations across iterations.

Building on GlycoPy's differentiable simulation framework, the `MPCTool` class automates the transcription from the DAE-based optimal control problem to the NLP form. Users specify: the control grid (time points), the time-varying variables to be optimized (e.g., inputs), the objective (stage/terminal terms), and constraints (e.g., path bounds on states). `MPCTool.setup()` then generates all NLP ingredients (w, f, g) and their bounds, which are passed to the `Optimization` module through a CasADi-like interface. For usability, `MPCTool` also provides a snapshot/rewind mechanism to store and restore the full model states before optimization. After solving, GlycoPy returns decision variables and key trajectories (controls and states) in convenient DataFrame structures for plotting and further analysis.

2.4 Parameter estimation

Parameter estimation (PE) identifies model parameters (and optionally uncertain initial conditions) by matching model predictions to measured data. For a DAE model, PE can be written as an optimization problem over

$$p = \begin{bmatrix} p_M \\ x(0) \end{bmatrix},$$

subject to the DAE dynamics and measurement equations. After transcription, PE becomes an NLP with decision variables p.

GlycoPy provides a modular PE workflow as shown in Fig. 1.

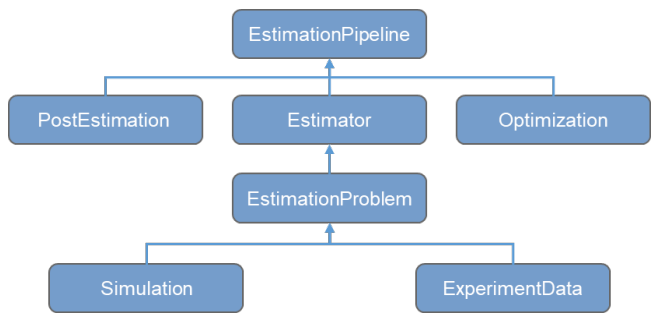

Figure 1. The structure of the parameter estiamtion module in GlycoPy.

The interface class `EstimationPipeline` defines and solves estimation tasks. Experimental data are stored in `ExperimentData`, including time-invariant operating parameters, time-varying operating profiles, measurements y, and their standard deviations σ_y. Multiple experiments are supported by creating multiple `ExperimentData` objects. For each experiment, GlycoPy runs a symbolic simulation to generate predicted responses $\hat{y}$, which enables efficient derivative-based optimization. An `EstimationProblem` object constructs residuals $r = y - \hat{y}$, stores measurement covariance Σ_y, parameter initialization and bounds (p_0, p_{lb}, p_{ub}), optional constraints, and priors $(\hat{p}, \Sigma_{\hat{p}})$ (when Bayesian estimation is used). An `Estimator` then builds the objective function according to the selected criterion:

$$\text{OLS:}\ \frac{1}{2} r^\top r, \tag{3}$$

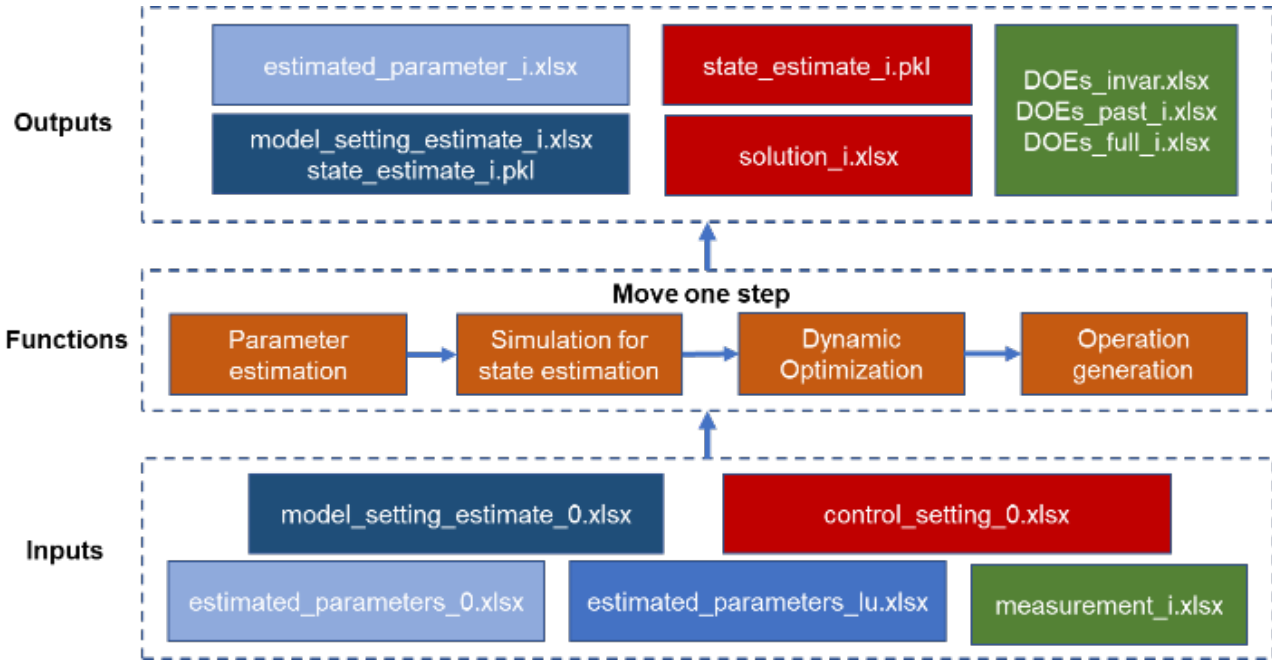

Figure 2. ANMPC implementation in GlycoPy.

$$\text{ML: } \frac{1}{2} r^{\top} \Sigma_y^{-1} r, \tag{4}$$

$$\text{MAP: } \frac{1}{2} r^{\top} \Sigma_y^{-1} r + \frac{1}{2} (p - \hat{p})^{\top} \Sigma_{\hat{p}}^{-1} (p - \hat{p}), \tag{5}$$

where OLS, ML and MAP stand for original least squares, maximum likelihood, and maximum a posteriori, respectively.

The resulting NLP is solved using GlycoPy's optimization interface. For post-analysis, `PostEstimation` computes the Fisher information matrix, parameter covariance, and confidence intervals using the output sensitivity $X = \partial \hat{y} / \partial p$.

For practical usability, GlycoPy supports: (i) estimating initial states together with parameters, (ii) optional log-parameter transforms to improve conditioning when parameters span orders of magnitude, (iii) diagnostic mode to debug simulations/experiments, and (iv) Excel-based inputs for experimental operating conditions, measurement data, initial guesses, bounds, and priors.

2.5 NMPC

The simulation, parameter estimation, and dynamic optimization modules in GlycoPy can be combined to implement state NMPC and adaptive NMPC (ANMPC) for fed-batch processes. In each control interval k, ANMPC follows a standard loop in a shrinking horizon: (i) when new measurements are available, update the parameter estimate $\hat{p}(k)$ using all collected data; (ii) simulate the model to obtain the current state estimate $\hat{x}(k)$; (iii) solve an optimal control problem over the remaining horizon $[t_k, T]$ using $\hat{p}(k)$and $\hat{x}(k)$; and (iv) apply the first control move and sample again. The algorithm terminates at the end of the batch.

A practical challenge in applying NMPC for large-scale processes is debugging and diagnosis: users often need to stop, inspect intermediate results (states, parameters, predictions, control actions), and restart from any interval without rerunning the full batch. To support this workflow, GlycoPy adopts the implementation structure in Fig. 2, which separates the NMPC loop into reproducible blocks, **inputs**, **functions**, and **outputs**.

In the **inputs** block**,** all settings and data can be provided via spreadsheets, including model/initial-condition settings, control scheduling (manipulated variables and time grids), parameter priors (values, covariance, confidence intervals), parameter bounds, and measurement files with noise levels.

In the **functions** block**,** users implement four core steps based on existing pipelines in the previous section: parameter estimation, simulation for state estimation, dynamic optimization, and operation generation. Here, the operation generation function generates files recording not only the optimized future control actions but also the actual past actions. Here, recording the actual past actions is important because low-level actuators may deviate from ideal setpoints; these recorded actions are then reused for subsequent estimation and prediction.

In the **outputs** block**,** at each interval, the framework stores updated parameter estimates and covariance, updated model settings (including refreshed initial states),

optimized decision variables (past and full horizon), and a solution file containing predicted trajectories and applied controls.

State NMPC uses the same structure but does not update parameters; it relies on simulations with fixed parameters combined with available measurements for state updates.

3. CASE STUDY

We demonstrate GlycoPy on NMPC of the mAb N-glycosylation process using the multiscale glycosylation process model, which has been studied in [9]. The model provides the mapping from extracellular conditions and operating actions to intracellular Golgi processing and the resulting extracellular glycoform distribution, which can be illustrated by Fig. 3.

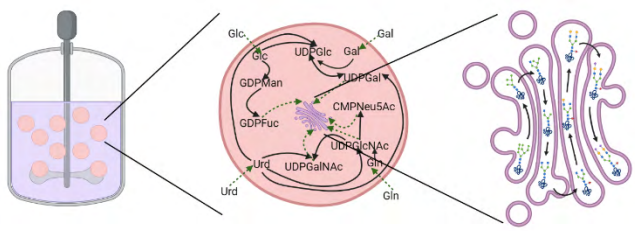

Figure 3. Multiscale mAb glycosylation process. The submodels, from left to right, are: the cell culture and metabolism model, the NSD synthesis model, and the Golgi glycosylation model. Reprinted from (Ma et al., 2025b) with permission.

The multiscale model consists of three interconnected layers [10]. At the bioreactor level, a cell culture submodel describes viable cell growth, extracellular metabolite dynamics, specific productivity, and secreted monoclonal antibody concentrations. These extracellular conditions are then passed to an intracellular nucleotide sugar donor synthesis submodel, which maps culture-level metabolite information to intracellular donor pools through a mechanistic representation of the relevant biosynthetic pathways and kinetics. Finally, a Golgi-level glycosylation submodel uses these donor species to predict the distribution of antibody glycoforms emerging from the secretory pathway by combining enzyme-mediated reaction networks with glycoprotein transport dynamics.

In the multiscale model, the Golgi subsystem is particularly complex, involving a large reaction network and spatially distributed dynamics. Implementing the full system as a single monolithic model would be difficult to maintain and reuse. Instead, GlycoPy enables a structured decomposition as shown in Fig. 4, where the overall model is built from the three main submodels, and the Golgi module is further organized into reusable components: transport proteins, enzyme, reaction, and network blocks. This hierarchy supports scalable simulation, estimation, optimization, and control for realistic biomanufacturing applications.

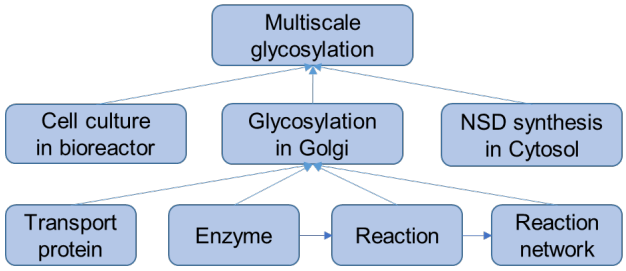

Figure 4. Implementation of the multiscale glycosylation model in GlycoPy.

The multiscale model takes the form of a large-scale partial differential–algebraic equation (PDAE) system, comprising 30 ordinary differential equations, 34 partial differential equations, and many strongly nonlinear algebraic relationships. The simulation using the finite difference method will be too slow to be used for NMPC. The symbolic simulation in GlycoPy enables the implementation of the quasi-steady-state (QSS) simulation algorithm and ensures differentiable simulations, which will be difficult in other modeling and control platforms. The steps of the QSS algorithm are as follows [11]: (i) the cell culture and NSD submodels with slow dynamics are simulated at selected time points and generate state variables env to be used for the following Golgi model simulation; (ii) the Golgi submodel is solved at steady state at each time point (parallelizable) with env as inputs, and generates intracellular glycosylation profiles $Y_{\mathrm{glycan}}^{\mathrm{intra}}$; (iii) the resulting $Y_{\mathrm{glycan}}^{\mathrm{intra}}$ are treated as time-varying parameters for the cell culture submodel to generate the extracellular glycosylation profiles $Y_{\mathrm{glycan}}^{\mathrm{extra}}$. This algorithm is packaged in `SimulationQSS` (derived from `Simulation`) and can return numerical or symbolic results depending on input types. Reported speedups exceed 300 folds for the multiscale glycosylation case. Because the QSS simulation is used, the total computation time for a single ANMPC step, including parameter estimation, state estimation, and dynamic optimization, remains below 2 h [9].

In the control problem, we maximize the galactosylation index (GI) at harvest time ($T = 288$ h),

$$\mathrm{GI}(T) = [\mathrm{FA2G1}](T) + 2\,[\mathrm{FA2G2}](T), \tag{6}$$

where $[\mathrm{FA2G1}]$ and $[\mathrm{FA2G2}]$ are the extracelluar concentrations of the corresponding glycoforms at harvest. The NMPC problem includes operational feasibility constraints on the bioreactor and product quality: (i) cell viability must remain above 60%; (ii) the media in the bioreactor volume must stay between 0.75 L and 2.25 L ; and (iii) the fraction of Man5-attached mAb is constrained by an upper bound 5% to avoid unacceptable high-mannose

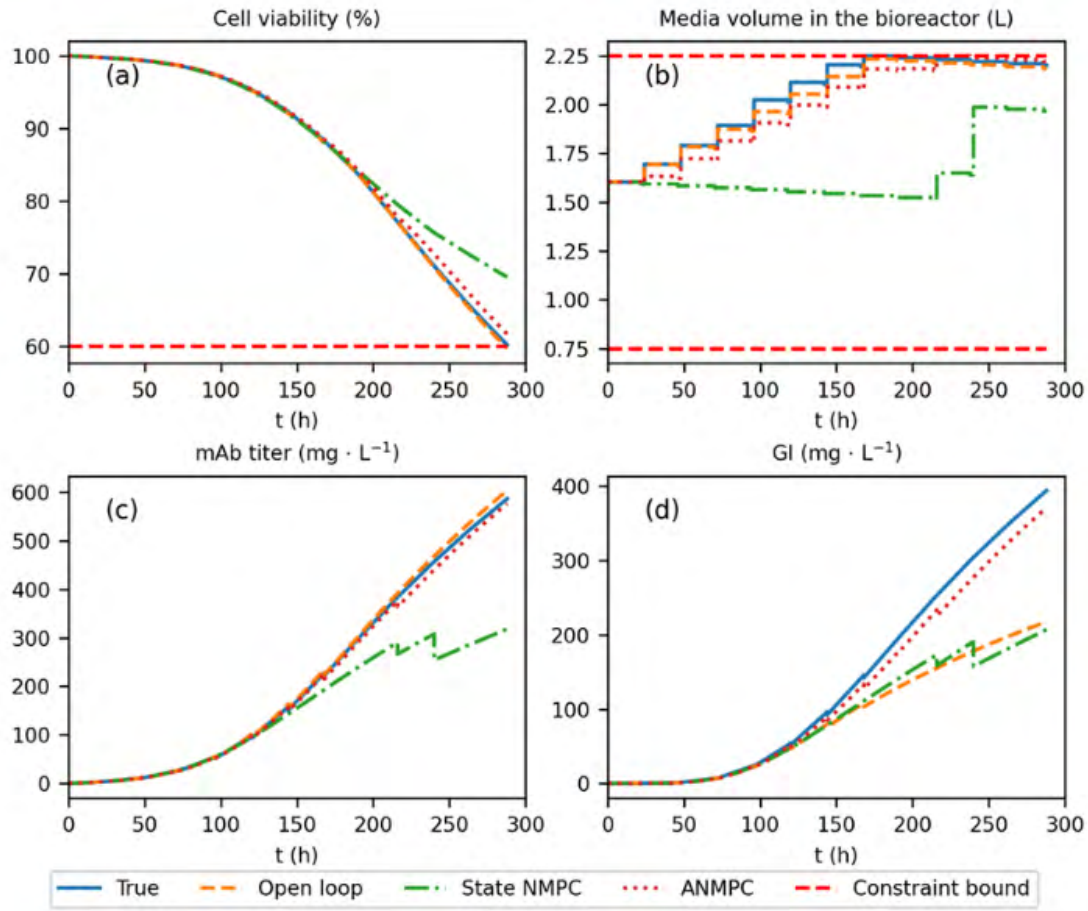

Figure 5. State variable trajectories under ground-truth optimization results and different control strategies.

content, where Man5 refers to High-mannose glycan containing five mannose residues. The controller adjusts two feeds commonly used in practice for the control purpose: a media supplement and a galactose (Gal) supplement. These inputs influence cell growth/metabolism and NSD availability, thereby shaping Golgi reaction fluxes and final glycoform composition.

The optimization and NMPC frameworks in the last section are used for the case study. We compare three control strategies: open-loop optimization, state NMPC, and ANMPC. Because of the use of the efficient quasi-steady-state simulation. The results are shown in Fig. 5.

All three strategies produce feasible trajectories satisfying the viability, volume, and Man5 constraints. However, ANMPC achieves a substantially higher terminal GI, which approximately doubles the GI achieved by open-loop optimization and state NMPC, and approaches the performance of ground-truth optimization. This highlights that, for glycosylation process control, closed-loop optimal control and adaptation can deliver large gains in product quality while maintaining process feasibility.

4. CONCLUSION

To promote the practical deployment of NMPC in real-world (bio)chemical processes, this work introduces GlycoPy, an equation-oriented (EO), object-oriented Python framework for modeling, optimization, and optimal control. GlycoPy allows users to focus on writing equations while providing hierarchical modeling capabilities that simplify the development, use, and maintenance of large-scale and multiscale process models. GlycoPy retains CasADi's powerful symbolic framework, enabling the use of symbolic simulation to develop efficient and differentiable simulation algorithms. The framework also includes parameter estimation and dynamic optimization modules, which form core building blocks for NMPC implementations. In addition, GlycoPy can export key intermediate results (e.g., parameter estimates and state trajectories), supporting transparent interpretation and efficient diagnosis of NMPC runs. Finally, a multiscale mAb glycosylation case study illustrates the effectiveness of the proposed framework.

DIGITAL SUPPLEMENTARY MATERIAL

GlycoPy is available at https://github.com/eaglema/glycopy, and the libraries for the glycosylation process modeling, optimization and control based on GlycoPy is available at https://github.com/eaglema/glycosylation.

ACKNOWLEDGEMENTS

This work was supported by a Project Award Agreement from the National Institute for Innovation in Manufacturing Biopharmaceuticals (NIIMBL), with financial assistance from the U.S. Department of Commerce, National Institute of Standards and Technology (NIST), awards70NANB17H002 and 70NANB20H037.

AUTHOR IDENTIFIERS

Author ORCIDs:
Yingjie Ma: 0009-0006-4458-9859
Jing Guo: None
Richard D. Braatz: 0000-0003-4304-3484

REFERENCES

1. P. I. Barton. The modelling and simulation of combined discrete/continuous processes. Imperial College London (University of London), London, 1992.
2. Gunnell L, Nicholson B, Hedengren JD. Equation-based and data-driven modeling: open-source software current state and future directions. Computers & Chemical Engineering 181:108521 (2024). https://doi.org/10.1016/j.compchemeng.2023.108521
3. Siemens Process Systems Engineering Ltd. *gPROMS ModelBuilder*. (2025).
4. Beal LDR, Hill DC, Martin RA, Hedengren JD. GEKKO optimization suite. Processes 6:106 (2018). https://doi.org/10.3390/pr6080106
5. Fiedler F, Karg B, Lüken L, Brandner D, Heinlein M, Brabender F, Lucia S. Do-mpc: towards FAIR nonlinear and robust model predictive control. Control Engineering Practice 140:105676 (2023). https://doi.org/10.1016/j.conengprac.2023.105676
6. S. S. Bhonsale, M. Vallerio, D. Telen, D. Vercammen, F. Logist, J. van Impe. SolACE: An open source package for nonlinear model predictive control and state estimation for (bio)chemical processes." P. 1971–1976 in Computer Aided Chemical Engineering. Vol. 38, edited by Z. Kravanja and M. Bogataj. Elsevier. (2016)
7. Hart WE, Watson JP, Woodruff DL. Pyomo: modeling and solving mathematical programs in python. Math. Prog. Comp. 3:219-260 (2011). https://doi.org/10.1007/s12532-011-0026-8
8. Andersson JAE, Gillis J, Horn G, Rawlings JB, Diehl M. Casadi: a software framework for nonlinear optimization and optimal control. Math. Prog. Comp. 11:1-36 (2018). https://doi.org/10.1007/s12532-018-0139-4
9. Ma Y, Guo J, Dubs AB, Ganko K, Braatz RD. Adaptive nonlinear model predictive control of monoclonal antibody glycosylation in CHO cell culture. Control Engineering Practice 169:106731 (2026). https://doi.org/10.1016/j.conengprac.2025.106731
10. Kotidis P, Jedrzejewski P, Sou SN, Sellick C, Polizzi K, del Val IJ, Kontoravdi C. Model?based optimization of antibody galactosylation in CHO cell culture. Biotech & Bioengineering 116:1612-1626 (2019). https://doi.org/10.1002/bit.26960
11. Ma Y, Guo J, Maloney AJ, Braatz RD. Quasi-steady-state approach for efficient multiscale simulation and optimization of mab glycosylation in CHO cell culture. Chemical Engineering Science 318:122162 (2025). https://doi.org/10.1016/j.ces.2025.122162

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Automated workflow for the configuration of modular plants and HAZOP analysis by utilizing DEXPI P&ID

Mathias Schmitz[a*], Janis Weber[a], and Norbert Kockmann[a]

[a] TU Dortmund University, Department of Biochemical and Chemical Engineering, Dortmund, North Rhine-Westphalia, Germany
* Corresponding Author: mathias.schmitz@tu-dortmund.de.

ABSTRACT

The increasing adoption of modular plant concepts in process engineering requires new strategies for design and safety evaluation. Conventional hazard analysis methods, such as HAZOP are time-consuming and must be repeated whenever a configuration changes, which contradicts the flexibility that modularization aims to achieve. This work combines concepts from the modular HAZOP method and the preHAZOP method to enable an automated, early-stage safety assessment for modular plants based on P&IDs. A workflow is developed that generates a combined P&ID for a modular plant from individual module (PEA) P&IDs provided in DEXPI format and performs an adapted preHAZOP analysis on the resulting plant representation.A key motivation is that P&IDs are always created during plant or module design and already contain relevant piping information, which can be reused for automated module interconnection and plausibility checks. In the proposed workflow, standardized interfaces enable automatic connection generation and validation. The safety assessment is performed by rule-based deviation propagation on a DEXPI-based P&ID graph and can operate with, without, or alongside process data. The workflow outputs structured scenario tables and selected design consistency checks and is implemented as a Streamlit web application, validated on a modular example system.

Keywords: Modular plants, HAZOP, DEXPI, P&ID, Safety analysis

INTRODUCTION

Globalized markets and increasingly dynamic product demands require the process industry to react faster and design plants more flexibly. Modular plants address this need by enabling reuse of standardized units and faster (re-)configuration of production systems [1]. However, the benefits of modularization are constrained by the effort required for safety studies and approvals. Conventional HAZOP is systematic but resource-intensive; repeating a full HAZOP after each configuration change undermines modularity's speed advantage. This work therefore targets an automated workflow that supports early safety assessments and helps accelerate approval processes for modular plants.

In industrial practice, each module plant or plant section is documented in a P&ID diagram during its design phase. The P&ID is a core component in the design phase during detail engineering. Manufacturers commonly annotate piping with information such as identifiers and nominal diameters, and sometimes even material information [2, 3]. In the workflow of this contribution, such data are explicitly attached to module interfaces and exported with the P&ID, enabling automatic connection generation and interface checks when assembling modules into a plant.

To achieve the goals, this contribution combines two ideas: a modular safety approach (splitting safety knowledge into reusable module/equipment-related parts) and automated, data-driven scenario detection from machine-readable P&IDs (preHAZOP). A software tool is developed that merges multiple P&IDs of modules into one plant-level DEXPI P&ID and then runs an adapted preHAZOP analysis on the combined P&ID of the modular plant.

BACKGROUND

Modular plants

Modular plants decompose a process into process

equipment assemblies (PEAs), each delivering a defined process function and equipped with own automation. This automation is done by the Module Type Package (MTP) which is comparable to a driver of peripheral devices for a computer. It encapsulates the automation interface of a PEA and supports orchestration by the Process Orchestration Layer (POL). MTP ensures the manufacturer independent compatibility between each PEA as well as to the POL. Each PEA is able to execute certain services to perform dedicated unit operations. Rearrangement of the PEAs allows to perform different processes without building a completely new plant for one new product.

While the modular concept and MTP are defined by the VDI/VDE/NAMUR 2658 [4] and 2776 [5] guidelines and different vendors are offering POLs or support to create MTP-files, there is no example of a safety assessment for the approval of a modular plant.

DEXPI P&ID

P&IDs are central documents throughout planning, construction, operation, and modification of process plants. They represent equipment, piping, valves, and instrumentation and control functions in a standardized way (e.g., via ISO/IEC-related standards), allowing consistent communication across disciplines. DEXPI e.V. is an initiative comprising various representatives from industry and academia who are working on common data exchange formats. The first implementations were carried out with the DEXPI Process Specification [6] and the DEXPI P&ID Specification [7]. In October 2025 were the DEXPI 2.0 Specification released which combines the Process Specification and P&ID Specification [8]. DEXPI is a standardized exchange format for P&IDs grounded in ISO 15926 concepts, typically realized as an XML-based schema that represents plant topology and attributes in a machine-readable form. While the DEXPI P&ID Specification relied on a Proteus-XML schema, DEXPI 2.0 uses a common XML format. The DEXPI format aligns with different international standards such as ISO 15926 [9]. The P&ID files in this work were created by using the software PlantEngineer from X-Visual Technologies GmbH.

Modular HAZOP

A HAZOP (Hazard and Operability) analysis identifies hazards and operability problems by systematically generating deviations from design intent using guidewords and parameters, then analyzing causes and consequences in an interdisciplinary team. The considered process is divided into functional units and their intended operation is described as a target function. While effective, the method is labor-intensive. Modular HAZOP extends HAZOP thinking to modular plants by structuring safety knowledge into module-oriented pieces and defining interfaces so results can be recombined for different plant configurations. For a modular plant the PEAs can be chosen as the functional units.

The ISSA (International Social Security Association) suggests seven guidewords in combination with parameters to form a deviation [10]. These guidewords are: no, more, less, as well as, partially, inversion, differently than. Three of them are considered in this work which are ‚no‘, ‚more‘ and ‚less‘. The considered parameters in this work are temperature, flow and pressure. The combination of the other guidewords with the parameters is in certain cases not reasonable or brings up an open interpretation of the deviation case and should be discussed by a HAZOP team.

Klose [11] suggested in his doctoral thesis an approach to execute a HAZOP analysis for modular plants. He considered equipment-specific safety (mHAZOP) and process safety (pHAZOP) separately in the first place. This is because each PEA is built in a more generalized way and not tailored with a certain process in mind. The operator of the modular plant knows about the planned processes and can execute the p-HAZOP. Afterwards the m- and p-HAZOP are combined to check if the selected configuration of the PEAs is suitable for the process. Klose defined a PEA as a so-called HAZOP-ode which features one or more HAZOP-Interfaces to connect to other HAZOP-Nodes (PEAs) or the plants infrastructure. The different cases of deviations are applied to the HAZOP-Nodes. The HAZOP-Interfaces are relevant to transfer deviations from one HAZOP-Node to another one. [11-13]

preHAZOP

preHAZOP is a tool approach developed by Oeing et al. [14] that reads a DEXPI-P&ID and simulation results to build a graph model of the plant, enrich it with equipment, streams, and safeguards, and then detect typical HAZOP scenarios based on a machine-readable scenario database.

DEXPI P&IDs provide structured information about plant topology, equipment and connections. This topology is transformed into a directed graph (stored in GraphML) that becomes the basis for algorithmic searches of hazardous constellations. In addition, operating conditions are incorporated by mapping data from a process simulation onto the same graph.

Based on this combined model, preHAZOP detects potential scenarios using a scenario database that formalizes common HAZOP elements (guideword, parameter, cause and consequence) together with machine-checkable requirements expressed as graph patterns and attribute constraints. The tool is able to recognize functional and instrumentation-related structures by classifying valves and pipes, identifying measurement locations and control loops, and detecting the presence and assignment of protective devices such as relief

paths, while also accounting for disabling configurations like isolating valves. The output is a structured list of scenario instances linked to specific equipment or sections of the plant, complemented by fields that indicate where automated assessment is insufficient and manual verification is needed (e.g., confirming design limits).

This contribution adopts the graph-centric idea but adapts the method to better fit modular plants, extend scenario generation, and support use cases without simulation data. The general workflow of the preHAZOP is presented in Fig. 1, where the added and adopted steps are outlined in blue.

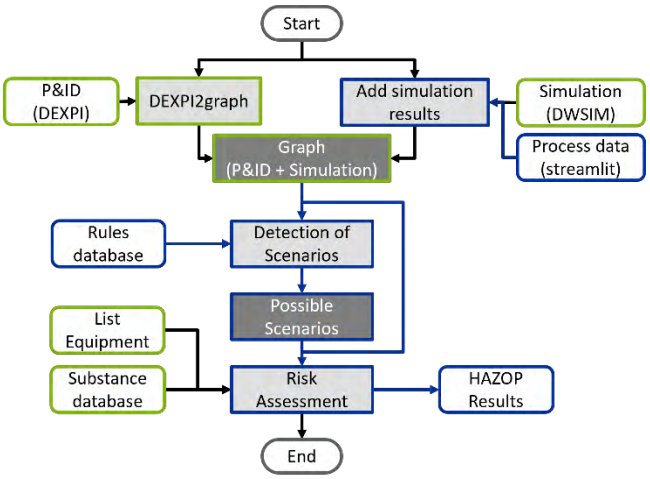

Figure 1. Suggested preHAZOP workflow. Blue marked steps are added or adapted from the original workflow.

METHODOLOGY

Laboratory set-up

The suggested workflow uses a modular reactor setup from the laboratory as an example for the P&ID flowsheets. The laboratory setup consists of four PEAs, which are shown in Fig. 2. These comprise a reactor PEA, two identical dosing units and a temperature PEA. The dosing units feed the two reactants via the temperature PEA into the reactor, where the reaction takes place before the product is discharged into another tank. The temperature PEA provides the required flow temperature for the reactants and also enables the temperature of the reactor to be controlled.

The Reactor-PEA is a continuous stirred tank reactor (CSTR) consisting of a vessel with a heating jacket, a stirrer with motor, three sensors to measure the level, the temperature and the conductivity of the substance filled in. The reactor has two inlets to feed educts inside of it and one outlet for the product with a controllable solenoid valve as well as an in- and outlet at the heating jacket.

The Dosing-PEA consists of a pump, a mass flowmeter and a storage tank. The Temperature-PEA is realized by a thermostat and is depicted here as a combination of a pump and a heater.

P&ID merge

In the workflow of Oeing [14] it was not necessary to consider multiple P&IDs. Since each PEA provides an own P&ID it is necessary to combine them to form the P&ID of the complete modular plant. To provide a user-friendly tool, the here presented workflow can executed by using a Streamlit application. Calling the python-file opens up a new tab in the user default browser, running the HMI. In the first step, the user is able to upload a DEXPI P&ID file or can select from already uploaded files.

For P&IDs of PEAs, correct and consistent interface definition is crucial. The PEA interfaces are defined as in- or outlets and are preceded by flanges. The flange objects store interface-related engineering constraints such as nominal diameter, design pressure, and temperature range.

Interface IDs follow a structured naming convention and are stored in attributes depending on whether an interface is an inlet, outlet, or bidirectional port. This convention allows the merging tool to automatically identify connection candidates, directions, and required adjustments.

Information about the PEA interfaces is extracted and listed to select how the different PEAs are connected with each other. When combining multiple DEXPI files, ID collisions must be avoided. The tool therefore prefixes each element ID with a PEA-specific shorthand and updates all references consistently across equipment, piping networks, and instrumentation functions. Simultaneously, it extracts geometric information (coordinates and extents) for PEA layout computation and stores interface coordinate data for later routing of inter-module connections. The individual P&IDs of the PEAs from the laboratory setup are shown in Fig. 2 as well as the combination of them to the complete modular setup.

To create a readable combined P&ID, PEAs must be arranged without overlaps and ideally with short connection lines. By using constraint programming in form of the CP-SAT-solver from the OR-Tools library and modeling each PEA as rectangle on the drawing sheet, enables to enforce certain rules. The first rule is that PEAs does not overlap and must keep a minimum distance. Secondly PEAs must remain within the sheet boundary and avoid the title block area. Lastly, an objective function minimizes the sum of Manhattan distances between connected interface points, discouraging small misalignments that would create unnecessary bends. This transforms a reasonable positioning into a reproducible optimization problem. The solver output is translated into shifts applied to all elements of each PEA P&ID so that the PEA internal structure remains intact while the plant-level arrangement changes. Fig. 3 shows the steps of the

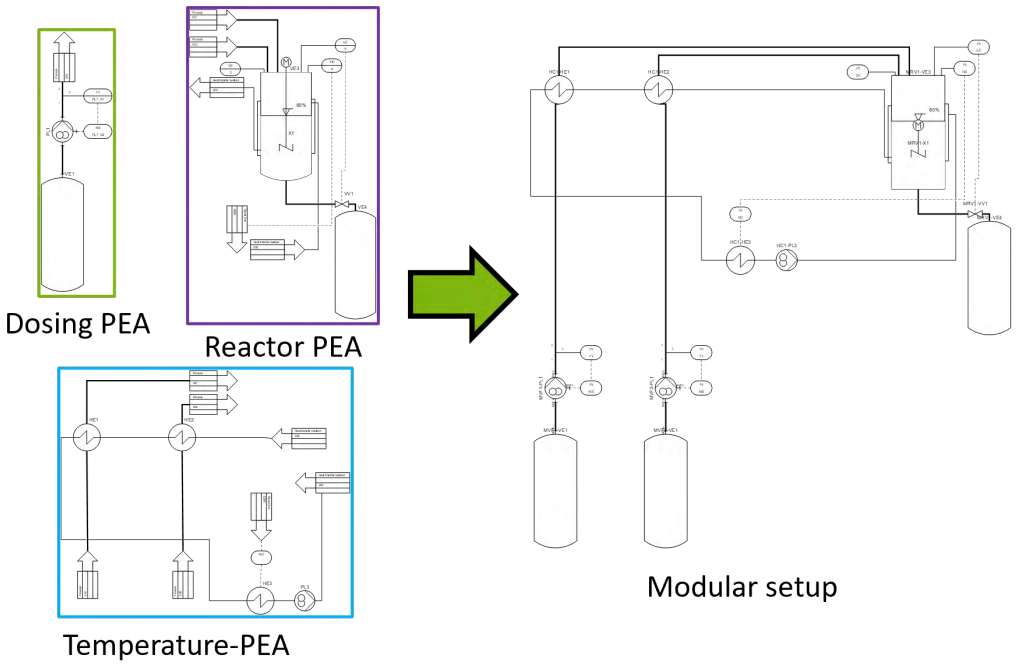

Figure 2. P&IDs of individual PEAs and combination to complete P&ID of the modular setup.

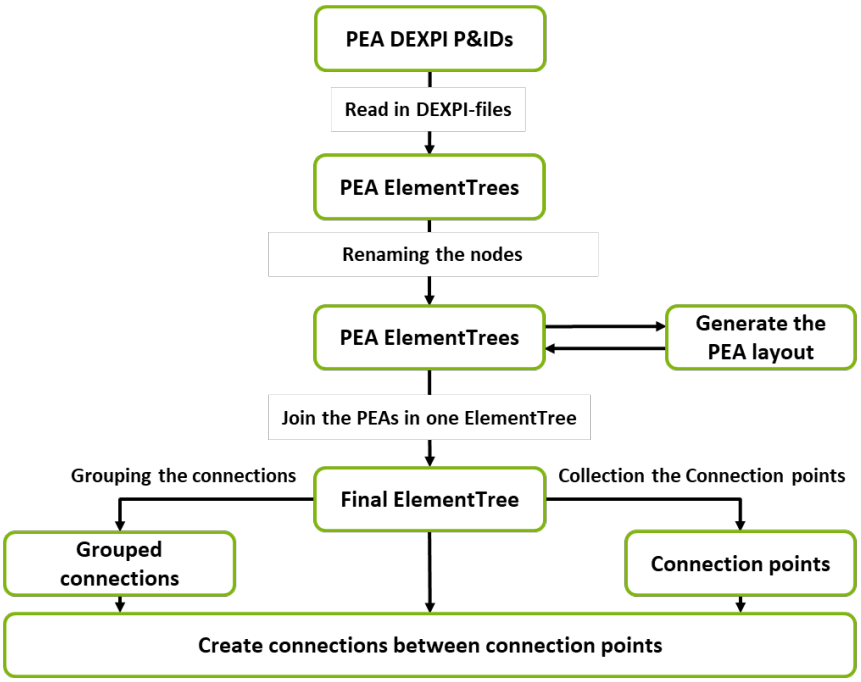

Figure 3. Procedure of reading out multiple PEA P&IDs and creating connections between their interfaces to create a combined P&ID.

programmed workflow to connect multiple P&IDs at their distinct interfaces.

After placement, the tool creates inter-module piping segments between flange coordinates. This is done by collecting all interfaces and grouping them accordingly to their function as inlet or outlet. Before the connections are created, they are grouped together. Connections with identical start and end points are combined. This can result in a 1:1 connection between two points (single start to single end) as well as 1:n or n:1 links (multiple starts or ends). For the later one T-pieces are inserted and positioned between PEAs to avoid placing

junctions inside PEA drawings. Additionally, a direction alignment is performed since individual interfaces can be defined as bidirectional. In Fig. 4 the procedure to create connections and inserting T-pieces if necessary is shown.

Before creating a connection, the tool compares flange attributes and records inconsistencies—for example mismatches in nominal diameter or minimum and maximum design pressure and temperature. This provides immediate feedback on interface compatibility and supports safer PEA coupling decisions early in engineering.

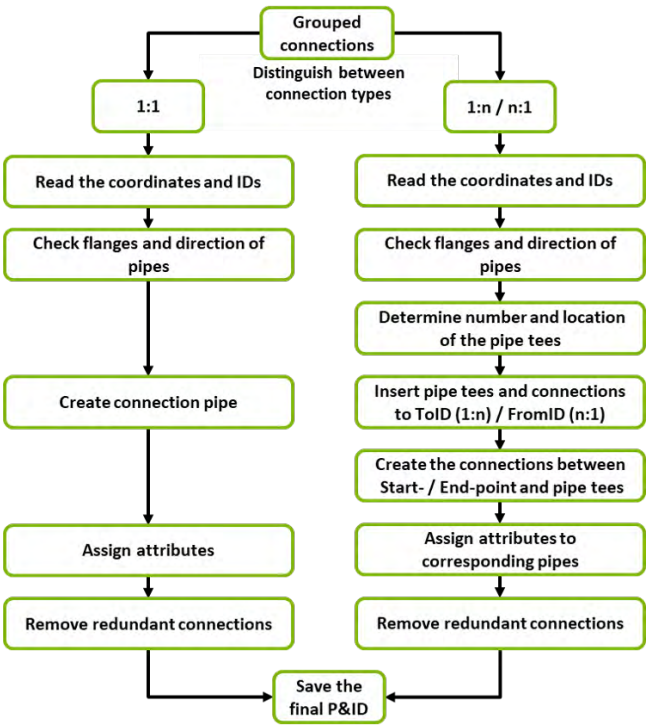

Figure 4. Procedure to check if its necessary to insert T-pieces between PEAs and how to connect them.

Adapted preHAZOP for modular plants

The original preHAZOP concept depends strongly on simulation data and a scenario database. For modular plants, simulation data may not exist when assessing a single PEA, or the final plant process may not yet be fixed. This work therefore adapts the preHAZOP so that scenario recognition remains possible without a simulation, while still benefiting from simulation or manually entered process data if available.

Instead of a purely pattern-based scenario search in a spreadsheet, this workflow implements a rule database organized by individual equipment classes in separated YAML-files. Within this rule base, different rule types are distinguished: rules that define local deviations as events originating at specific equipment, including frequency estimations and associated safeguards; rules that map these deviations to local consequences such as damage or leakage based on guideword–parameter combinations; rules that specify how deviations propagate forward or backward through the system depending on the type of piping; and rules that describe how equipment responds to deviations arriving from upstream or downstream, optionally depending on specifications such as the presence of a heating jacket, inertization, an agitator, or a control loop. This organization enables richer and more explicit logic than a flat table and aligns well with modular HAZOP thinking, because deviations can be generated locally within a PEA and then combined across PEA boundaries through propagation via the defined interfaces.

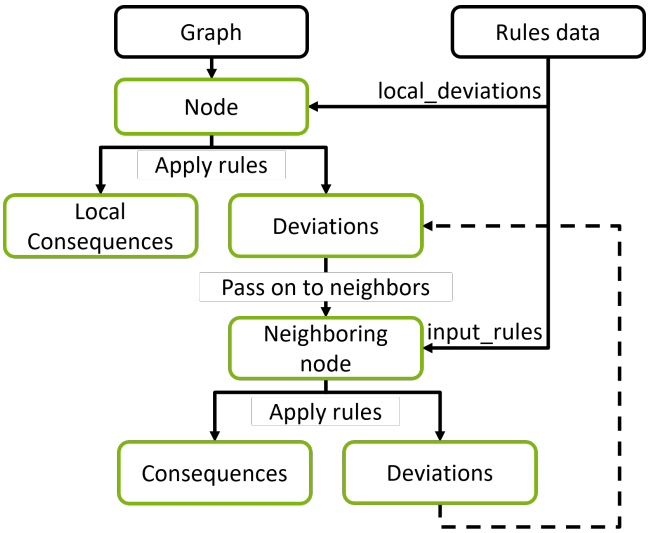

Figure 5. Workflow of applying the safety rules on the extracted graph.

The detection procedure follows two consecutive steps. First, the algorithm generates initial deviations at each node of the extracted graph based on the local deviation rules and records the corresponding local consequences immediately as scenario candidates. Second, these deviations are propagated through the graph using a queue-based mechanism. Whenever a deviation reaches another node, the engine checks the applicable input rules for that equipment, records any resulting consequences, and continues propagation if the defined propagation rules permit it. To keep the output manageable and to prevent an uncontrolled growth of near-identical scenarios, repeated events are filtered, for example when similar causes originate from the same equipment class at a node, and the propagation depth can be restricted, with a default limit of up to six nodes. The resulting scenarios are then stored in a structured, Excel-like table format consistent with modular HAZOP documentation, including deviation, cause, consequences with severity and frequency considerations, and safeguards that are either already present or recommended. Fig. 5

illustrates at which points the rules are applied in the procedure and how deviations are propagated to its neighbor nodes.

To support different stages of engineering maturity, the tool can handle process data in several ways. If simulation results are available, they can be imported in the same manner as in the original preHAZOP workflow. If no simulation file exists and operating conditions are determined manually, users can enter minimum and maximum operating pressure and temperature as well as relevant media, either as global default values or as equipment-specific overrides. If no process data are available at all, the scenario detection still runs, although parameter-based checks and media–material evaluation may be more limited. In addition to producing scenario lists, the adapted workflow performs two engineering checks that add practical value. Operating pressures and temperatures, whether derived from simulation or manual input, are compared against the design limits stored in the P&ID, so that exceedances and critical mismatches between neighboring equipment design limits can be documented as potential weak points during upset conditions. Furthermore, a media–material compatibility check is carried out using a dedicated database that flags problematic combinations with an A–D rating and can identify mechanisms such as catalytic decomposition, for example the incompatibility risk of hydrogen peroxide in contact with stainless steel. Any problematic combinations identified in this way are exported to support targeted follow-up in the subsequent engineering review.

Web Application Implementation

To make the workflow usable for engineers, the toolchain is implemented as a Streamlit web app with two separate tabs. In the first tab the user uploads PEA P&IDs, selects the PEAs, defines their interface connections via dropdowns populated from detected ports, specify attributes for connecting pipes, and export the configuration for reuse. The merged DEXPI P&ID and an interface mismatch report are returned.

In the second tab, the user runs the preHAZOP with the previously generated P&ID or is able to upload an older P&ID. Process data can be optionally entered either manually or in form of a simulation file. After running the analysis, the results can be downloaded as a table.

Validation

The workflow involving the simulation was tested by using a DWSIM simulation representing the lab set-up with a saponification reaction. In this reaction sodium hydroxide reacts with ethyl acetate to form sodium acetate and ethanol.

As a case for manual data input the oxidation of sodium thiosulfate with hydrogen peroxide is considered which serves as an example for a highly exothermic reaction [15, 16]. For 1 M sodium thiosulfate and 2.2 M hydrogen peroxide at 40 °C results in a temperature increase of 70 °C to 110 °C inside of the reactor. As a result, the HAZOP table shows that the process temperature exceeds the maximum operating temperature of the reactor. Additionally, the media and material compatibility check gives a message that hydrogen peroxide should not get contact with stainless steel.

The merging tool successfully combined multiple PEAs into a single P&ID without overlap and with short connection paths. The adapted preHAZOP detected meaningful scenarios both on the individual PEA level and on the combined plant level; when simulation data were available, the design-vs-operation checks identified a simulated design pressure exceedance (not necessarily representing the real lab setup, but validating the check mechanism). Media–material checks did not reveal issues for the lab media and materials considered.

The workflow demonstrates manual derivation and entry of operating parameters, including estimating the maximum adiabatic temperature rise for a strongly exothermic reaction. The merging workflow handled a larger number of PEAs, bidirectional interfaces, and multi-branch connections, including correct direction adjustments and T-piece insertion.

The adapted preHAZOP produced both intra-module and inter-module scenarios. Importantly, the additional checks detected a design-vs-operation temperature violation for a downstream pump (operating temperature exceeding the pump's max design temperature) and flagged a problematic media–material combination involving hydrogen peroxide and stainless steel (rated poor due to catalytic decomposition). These findings illustrate how the tool can surface concrete engineering issues early, even without detailed dynamic simulation.

SUMMARY AND CONCLUSION

This contribution builds up an the preHAZOP tool developed by Oeing et al. and adapts it for modular plants by incorporating modular safety concepts of Klose. The preHAZOP workflow is complemented with the option to combine multiple P&IDs into a complete P&ID of a plant and gives the opportunity to enter process parameters manually instead of relying on simulation files. It demonstrates that DEXPI P&IDs can serve as a practical, machine-readable backbone for two high-impact automation steps in modular plant engineering. Firstly, showing early safety assessment via graph-based, rule-driven scenario generation. Secondly, it provides a tool for rapid planning of different kinds of possible plant configurations by merging PEA P&IDs.

One of the main focuses in this work lies on the interfaces of the PEAs to connect the P&IDs with each other as well as compatibility checks between the PEAs.

Standardized connection points combined with manufacturer-provided piping attributes allow plausibility checks that are directly relevant for modular approval discussions (compatibility of design limits, nominal diameter, etc.). This helps in early phases of setting up a new plant configuration, to check if PEAs are allowed to be combined.

At the same time, limitations remain typical of knowledge-based safety automation: scenario quality depends on the completeness and correctness of rule databases, equipment classification, and available process data. The work therefore positions the tool as an accelerator for early-stage studies and as structured input for expert HAZOP teams—not as a full replacement for interdisciplinary review.

Several extensions for the preHAZOP can be proposed to improve the workflow further. The layout can be optimized by reducing line crossings, improving T-piece placement strategies and removing the current drawing-size limitation. Additionally, the rule database should be expanded to more equipment classes, guidewords, and better equipment sub-classification (e.g., distinguishing heaters vs. coolers). The media–material database could be enriched with concentration and temperature dependence. Towards an improved compatibility check, an implementation of ex-zone detection or check whether certain components are explosion-proof or not would be helpful.

Beneficial for the further automation of this workflow is the implementation of a connection to a chemical database to fetch safety relevant information of the used substances. Borgelt et al. has shown the implementation of such databases like the ChEBI, KEGG and PubChem [17].

Because DEXPI was updated during this work, the workflow should be adapted to the newer schema and—critically—support "clean" DEXPI files where properties are mapped into standard DEXPI attributes rather than vendor-specific attribute sets. This would improve interoperability with other authoring tools and strengthen the value of standardized P&ID data for safety automation.

A practical next step for usability is to leverage the HMI information contained in MTP files as a "preview" during plant configuration. Since PEAs in modular automation are already packaged with MTP-based interfaces, an automatically generated (or extracted) HMI view could guide users to select the correct connection points between PEAs more intuitively—reducing errors in manual interface selection and aligning the P&ID-based configuration step with the operational perspective of the PEA automation.

This work presents a coherent, implementable approach for accelerating modular plant safety assessments by reusing data already present in P&IDs. By merging PEA P&IDs (DEXPI) into a P&ID of a complete modular plant and applying an adapted, rule-based preHAZOP on a graph representation, the workflow enables systematic detection of deviation–cause–consequence chains across PEA boundaries and adds practical checks for design vs. operating limits and media–material compatibility. Implemented as a Streamlit web app and validated on modular examples, the workflow demonstrates how standardized engineering data can support faster, more reproducible early-stage safety assessments—an important step toward accelerating the approval of modular plants.

DIGITAL SUPPLEMENTARY MATERIAL

A collection of the created and used files can be found here: https://repository.nfdi4cat.org/dataset.xhtml?persistentId=hdl:21.11165/4cat/86xp-e338 This includes: Python-files, YAML-files with rules, Excel-files containing results of the HAZOP analysis and DWSIM simulation file.

ACKNOWLEDGEMENTS

The authors thank the German Federal Ministry for Economic Affairs and Energy (BMWE) for funding this research (funding no. 03EN2100B).

Open access funding enabled and organized by Projekt DEAL.

AUTHOR IDENTIFIERS

Author ORCIDs:
Schmitz: 0009-0006-2883-1617
Kockmann: 0000-0002-8852-3812

REFERENCES

1. Klose A, Lorenz J, Bittorf L, Stark K, Hoernicke M, Stutz A, Weinhold H, Krink N, Welscher W, Eckert M, Unland S, Menschner A, Da Silva Santos P, Kockmann N, Urbas L. Orchestration of modular plants. atp 63:68-77 (2022). https://doi.org/10.17560/atp.v63i9.2599
2. Bindel T, Hofmann D. R&i-fließschema. Springer Fachmedien Wiesbaden (2016). https://doi.org/10.1007/978-3-658-15559-9
3. Baerns M, Behr A, Brehm A, Gmehling J, Hinrichsen K, Hofmann H, Kleiber M, N. Kockmann, Onken U, Palkovits R, Renken A, Vogt D, Technische Chemie. Wiley-VCH (2023)
4. Maurmaier ASM. Dezentralisierte kommunikationsarchitektur für direkte modul-modul-interaktion auf basis von standardisierten konstrukten der vdi/vde/namur 2658. Automation

2019 :343-360 (2019). https://doi.org/10.51202/9783181023518-343
5. Verein Deutscher Ingenieure, VDI 2776 Blatt 1. Process Engineering Plants - Modular Plants - Fundamentals and Planning Modular Plants, Düsseldorf (2020).
6. DEXPI e.V., DEXPI Process Specification 1.0 (2023) https://dexpi.org/static/process_model_1.0/index.html
7. DEXPI e.V., DEXPI P&ID Specification 1.4 (2024) https://dexpi.org/static/pid_specification_1.4/
8. DEXPI e.V., DEXPI Specification 2.0 (2025). https://gitlab.com/dexpi/Specification
9. ISO 15926-2, Industrial Automation Systems and Integration – Integration of Life-Cycle Data for Process Plants Including Oil and Gas Production Facilities – Part 2: Data Model, Beuth Verlag, (2013)
10. International Social Security Association, Das PAAG-/HAZOP-Verfahren und weitere praxisbewährte Methoden, Jedermann-Verlag GmbH (2020)
11. Klose A. Modulare HAZOP für flexible anlagen der prozessindustrie. VDI Verlag (2023). https://doi.org/10.51202/9783186479204
12. Pelzer F, Klose A, Miesner J, Schmauder M, Urbas L. Safety in modular process plants: demonstration of safety concepts. Elektrotech. Inftech. 138:462-468 (2021). https://doi.org/10.1007/s00502-021-00928-8
13. Klose A, Kessler F, Pelzer F, Rothhaupt M, Kostiuk D, Kabashi A, Forkel V, Urbas L. Representing causal structures in HAZOP studies. 2021 26th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA) :1-4 (2021). https://doi.org/10.1109/etfa45728.2021.9613357
14. Oeing J, Holtermann T, Welscher W, Severins C, Vogel M, Kockmann N. Prehazop: graph?based safety analysis for early integration into automated engineering workflows. Chemie Ingenieur Technik 95:1083-1095 (2023). https://doi.org/10.1002/cite.202200222
15. Frede TA, Greive M, Kockmann N. Measuring kinetics in flow using isoperibolic flow calorimetry. Reactions 3:525-536 (2022). https://doi.org/10.3390/reactions3040035
16. Cohen W C, Spencer J L, Determination of Chemical Kinetics by Calorimetry. Chem. Eng. Prog., 58, 40–41 (1962)
17. Borgelt H, Kitel FG, Kockmann N. Ontology-based data pipeline for semantic reaction classification and research data management. Computers 14:311 (2025). https://doi.org/10.3390/computers14080311

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Accelerating Efficient Dimethyl Ether Synthesis through Machine Learning-Based Process Optimization

Mitra Jafari[a], Jefferson Santos da Silva[b], Wilson Sousa Mercês Neto[b], Lucas Fonseca Couto[b], Bogdan Dorneanu[a], Karen Valverde Pontes[b], Harvey Arellano-Garcia[a]*

[a] 1 Department of Process and Plant Technology, Brandenburg University of Technology Cottbus-Senftenberg, Cottbus
[b] Programa de Pós-Graduação em Energia e Ambiente (PGENAM), Escola Politécnica, Universidade Federal da Bahia, Salvador, Brazil
* Corresponding Author: arellano@b-tu.de

ABSTRACT

Dimethyl ether (DME) is a promising clean fuel and chemical intermediate, yet its synthesis from synthesis gas remains highly sensitive to both catalyst formulation and operating conditions. In this work, a data-driven framework is developed that combines machine learning surrogate modeling with multi-objective optimization to support systematic decision-making in DME synthesis. The novelty lies in the systematic comparison of different optimization approaches applied to an identical machine learning surrogate model for DME synthesis, thereby highlighting their respective strengths and limitations as decision-support tools under limited-data conditions. A dataset compiled from published literature includes catalyst composition, preparation methods, physicochemical descriptors, and operating conditions, with CO conversion and DME selectivity as performance indicators. After data preprocessing, feature analysis using correlation analysis and principal component analysis (PCA) is applied to explore dominant trends and interactions. Several supervised machine learning models are trained and benchmarked, and the best-performing model is selected as a surrogate for optimization. Two complementary multi-objective optimization strategies are then applied to the same surrogate model: a fuzzy-enhanced NSGA-II algorithm and a constrained multi-objective Bayesian optimization approach. Gradient boosting is found to provide the most reliable predictive performance among the tested models. Both optimization strategies identify similar catalyst formulations as optimal, while differences emerge in the recommended operating conditions. The comparative analysis highlights how different optimization paradigms influence compromise solutions when balancing CO conversion and DME selectivity.

Keywords: Dimethyl Ether, Machine Learning, Optimization, Modelling, Syngas

INTRODUCTION

The transition toward sustainable and flexible fuel production pathways has intensified interest in synthesis gas (CO/H_2) conversion processes [1]. Among the available options, dimethyl ether (DME) has attracted considerable attention due to its high cetane number, clean combustion behaviour, and compatibility with existing fuel infrastructure. DME can be produced either via a two-step route involving methanol synthesis followed by dehydration or through a direct, one-step process using bifunctional catalysts. Despite extensive experimental efforts, achieving high CO conversion and DME selectivity simultaneously remains challenging because of the strong coupling between catalyst preparation method, and operating conditions [2, 3].

From a process engineering perspective, DME synthesis constitutes a high-dimensional and strongly nonlinear optimization problem. Reaction temperature, pressure, feed composition, space velocity, and catalyst characteristics interact in complex ways that are difficult to capture using mechanistic models, particularly when kinetic data are scarce or inconsistent across studies. As a result, experimental trial-and-error approaches remain common, but are time-consuming and resource-intensive.

Machine learning (ML) techniques provide an alternative route by enabling the construction of surrogate

models that directly learn input-output relationships from experimental data [4]. When applied to catalytic systems, ML models can capture nonlinear trends and interactions that are not easily accessible through simplified kinetic expressions [5, 6].

Multi-objective optimization methods are particularly relevant for processes such as DME synthesis, where improvement in CO conversion may come at the expense of DME selectivity and vice versa. Evolutionary algorithms such as NSGA-II are widely used to generate Pareto fronts representing trade-offs between competing objectives [7], while fuzzy decision-making methods enable the selection of a single compromise solution from these fronts. However, these approaches typically treat surrogate model predictions as deterministic and do not explicitly account for uncertainty [8]. Bayesian optimization [9] offers a complementary paradigm by combining probabilistic surrogate models with acquisition functions that guide the search toward promising regions of the design space while quantifying uncertainty [10, 11]. Although Bayesian optimization has been applied successfully in catalyst and reaction optimization, its performance relative to evolutionary multi-objective methods in the context of DME synthesis has not been systematically assessed.

Motivated by this gap, the present work compares two surrogate-based multi-objective optimization strategies: a fuzzy-enhanced NSGA-II algorithm and a constrained multi-objective Bayesian optimization approach, applied to the same machine learning model trained on literature data for DME synthesis. By focusing on methodological comparison rather than introducing new experimental results, this study aims to clarify how different optimization frameworks influence the identification of optimal solution and compromise operating conditions under limited data scenarios.

METHODOLOGY

Data Collection and Preprocessing

The dataset used in this study is compiled from published experimental literature on DME synthesis from synthesis gas (CO and H_2 mixture). It includes operating conditions and catalyst descriptors commonly reported across studies, such as reaction temperature, pressure, H_2/CO feed ratio, space velocity, catalyst composition, and preparation-related parameters. CO conversion and DME selectivity are selected as the output variables, as they directly reflect process performance and selectivity toward the desired product.

Prior to model development, the dataset is screened for outliers to reduce the influence of inconsistent or erroneous entries. Categorical variables describing catalyst type and preparation method are encoded numerically. Prior to analysis, all variables are normalized using min-max scaling, which linearly transforms each feature to the range [0, 1] based on its minimum and maximum values. This step mitigates scale differences and prevents features with larger magnitudes from dominating the analysis. Missing values are imputed using automated imputation procedures in JMP. Model training and validation are performed using k-fold cross-validation to ensure robust performance assessment given the limited dataset size.

To explore relationships between descriptors and performance indicators, Spearman correlation analysis is applied to identify monotonic trends. In parallel, principal component analysis (PCA) is used to examine the structure of the dataset and to assess how catalyst-related and operating condition-related variables contribute to overall variance.

Machine Learning Surrogate Model

Several supervised ML algorithms, such as XGBoost, Gradient Boosting, CatBoost, Random Forest, and Support Vector Regression (SVR), are trained to predict CO conversion and DME selectivity from the input variables. Model performance is evaluated using standard regression metrics, including the coefficient of determination (R^2) and the root mean square error (RMSE). The model demonstrating the best balance between predictive accuracy and generalization capability is selected as the surrogate model for subsequent optimization.

Fuzzy-Enhanced NSGA-II Multi-Objective Optimization

Multi-objective optimization is formulated with the goal of simultaneously maximizing CO conversion and DME selectivity. The NSGA-II evolutionary algorithm [12] is employed to generate a set of non-dominated solutions representing trade-offs between these objectives. Since NSGA-II is inherently a minimization algorithm, the negative values of the predicted objectives are minimized. This algorithm is implemented using the Pymoo library in Python [13].

To select a single compromise solution from the Pareto front, a fuzzy logic-based decision-making approach is applied [14]. This method integrates relative preferences between objectives into a unified metric, enabling the identification of a solution that best balances CO conversion and DME selectivity. The fuzzy-optimal solution is subsequently compared with experimentally reported data by identifying the closest real datapoint in terms of Euclidean distance in feature space.

Multi-Objective Bayesian Optimization

A constrained multi-objective Bayesian optimization strategy based on the ParEGO framework [15, 16] is also applied. In this approach, the two objectives are combined into a scalarized objective function using predefined weights, and a Gaussian process model is trained to approximate this scalarized response. An Expected

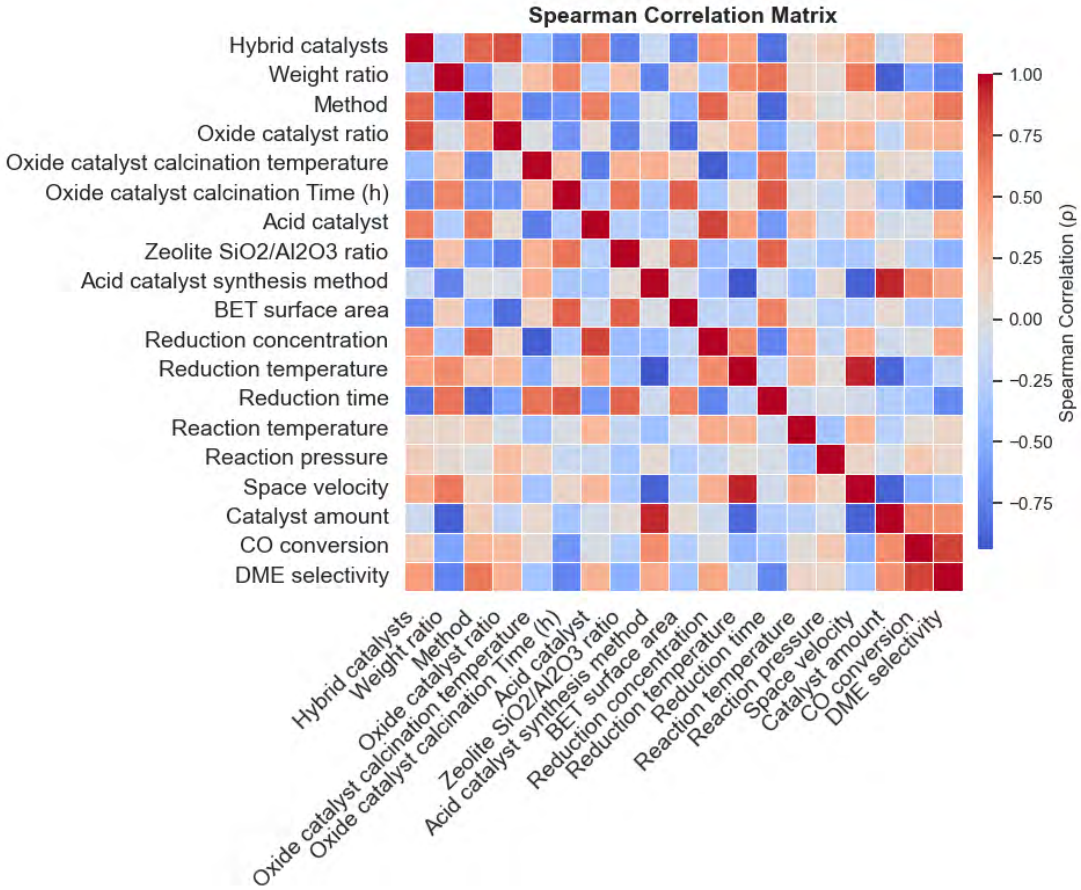

Figure 1. Spearman Correlation Analysis.

improvement acquisition function guides the selection of new candidate points while respecting process constraints [11, 17]. After completion of the optimization loop, the best compromise solution is identified based on the scalarized objective values. Presenting the optimized descriptors in their original physical units enables direct comparison with both experimental data and the fuzzy-enhanced NSGA-II results.

RESULTS

Data Preparation and Feature Analysis

The final dataset consists of 52 datapoints with 17 input features describing catalyst formulation, preparation, and operating conditions. Initial preprocessing ensures consistency across studies and prepares the data for ML analysis. After preprocessing, the data span CO conversions between 30 and 55%, and DME selectivity in the range of roughly 45-75%, reflecting the diversity of catalytic systems reported in the literature.

Spearman correlation analysis (Figure 1) provides an initial, quantitative indication of descriptor relevance. Catalyst-related, particularly synthesis method and zeolite SiO_2/Al_2O_3 ratio, show stronger monotonic correlations with DME selectivity, whereas operating variables such as space velocity exhibit higher correlation with CO conversion. However, several expected effects, such as the influence of temperature and pressure, are not fully captured, highlighting the importance of nonlinear modelling approaches.

PCA reveals further clarifies the structure of the dataset. As shown in Figure 2(a), the first four principal components account for approximately 90% of the total variance. The loading plot in Figure 2(b) indicates that the first two principal components are dominated by catalyst formulation and preparation descriptors, while PC3 and PC4 are more strongly associated with operating conditions such as reaction temperature, pressure, and space velocity. This separation quantitatively supports the need for modelling approaches capable of handling interactions between catalyst-related and operational variables.

Surrogate Model Performance

Table 1 summarizes the predictive performance of the evaluated ML models. The reduction in R^2 and increase in RMSE from training to validation are attributed

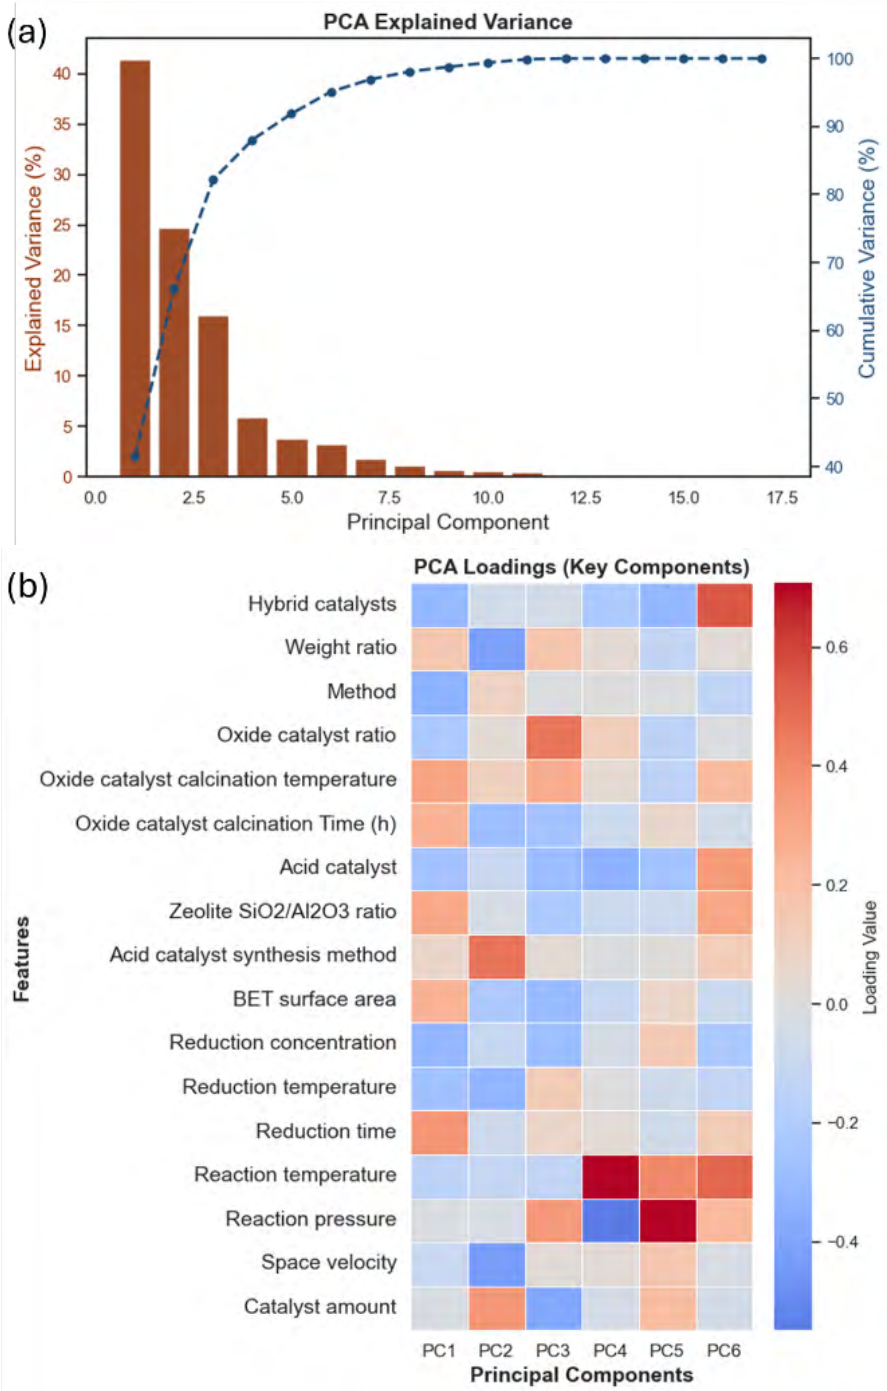

Figure 2. (a) Cumulative explained variance by PCA component, (b) PCA component loading.

to the generalization gap, indicating that models fit the training data more closely than unseen data, with some degree of overfitting. This effect can be reduced through improved hyperparameter tuning, regularization, and increasing dataset size.

While XGBoost and Gradient Boosting achieve similarly high training accuracy (R^2 of 0.99 for both CO conversion and DME selectivity), differences become apparent during validation. Gradient Boosting achieves R^2 of 0.83 for CO conversion and 0.89 for DME selectivity, with RMSE values of 0.01038 and 0.01258, respectively. In comparison, XGBoost shows lower generalization performance, particularly for CO conversion (validation R^2=0.74).

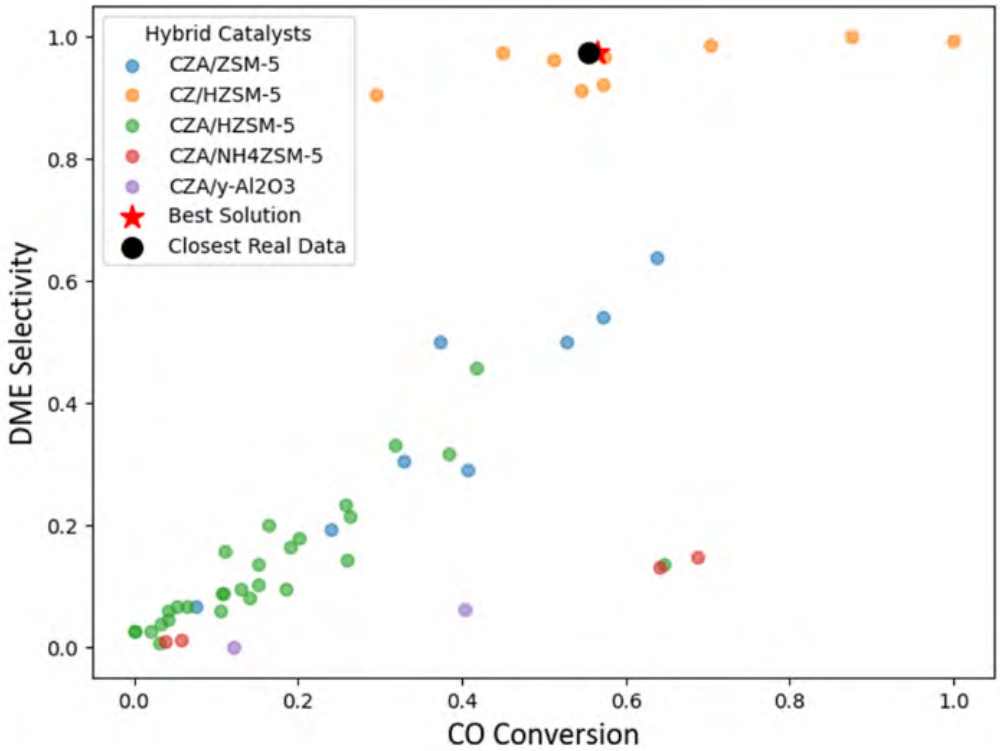

Figure 3. Comparison of the best solution and real datapoints.

Table 1: Predictive performance of the predictive models.

Training			
Method	**Parameter**	**R^2**	**RMSE**
XGBoost	CO Conversion	0.99	0.00002
	DME Selectivity	0.99	0.00001
Gradient Boosting	CO Conversion	0.99	0.00024
	DME Selectivity	0.99	0.00029
CatBoost	CO Conversion	0.97	0.00157
	DME Selectivity	0.98	0.00138
Random Forest	CO Conversion	0.96	0.00236
	DME Selectivity	0.85	0.01000
SVR	CO Conversion	0.83	0.01000
	DME Selectivity	0.91	0.01071
Validation			
Method	**Parameter**	**R^2**	**RMSE**
XGBoost	CO Conversion	0.74	0.01668
	DME Selectivity	0.82	0.01400
Gradient Boosting	CO Conversion	0.83	0.01038
	DME Selectivity	0.89	0.01258
CatBoost	CO Conversion	0.78	0.01332
	DME Selectivity	0.83	0.11604
Random Forest	CO Conversion	0.75	0.01547
	DME Selectivity	0.87	0.01360
SVR	CO Conversion	0.66	0.02115
	DME Selectivity	0.74	0.02619

These results indicate that Gradient Boosting provides the best balance between accuracy and generalization for the available dataset. Consequently, it is selected as the surrogate model for subsequent multi-objective optimization. The performance gap between training and validation metrics also reflects the limited size and heterogeneity of the literature-derived dataset, reinforcing importance of cautious interpretation of surrogate predictions.

Fuzzy-Enhanced NSGA-II Optimization Results

Using the Gradient Boosting surrogate, the fuzzy-enhanced NSGA-II algorithm generates a Pareto front representing the trade-offs between CO conversion and DME selectivity. The fuzzy decision-making step identifies a single compromise solution, which is highlighted in Figure 3 alongside the experimentally reported datapoints. The selected fuzzy-optimal solutions correspond to a CZ/HZSM-5 catalyst prepared via a core-shell method, operated at a reaction temperature of approximately 261°C and a pressure of 37 bar. The surrogate predicts a CO conversion of about 53% and a DME selectivity close to 70%. As shown in Table 2, the closest experimental datapoint exhibits very similar performance values (CO conversion close to 51% and DME selectivity of approximately 70%), lending credibility to the optimized solution.

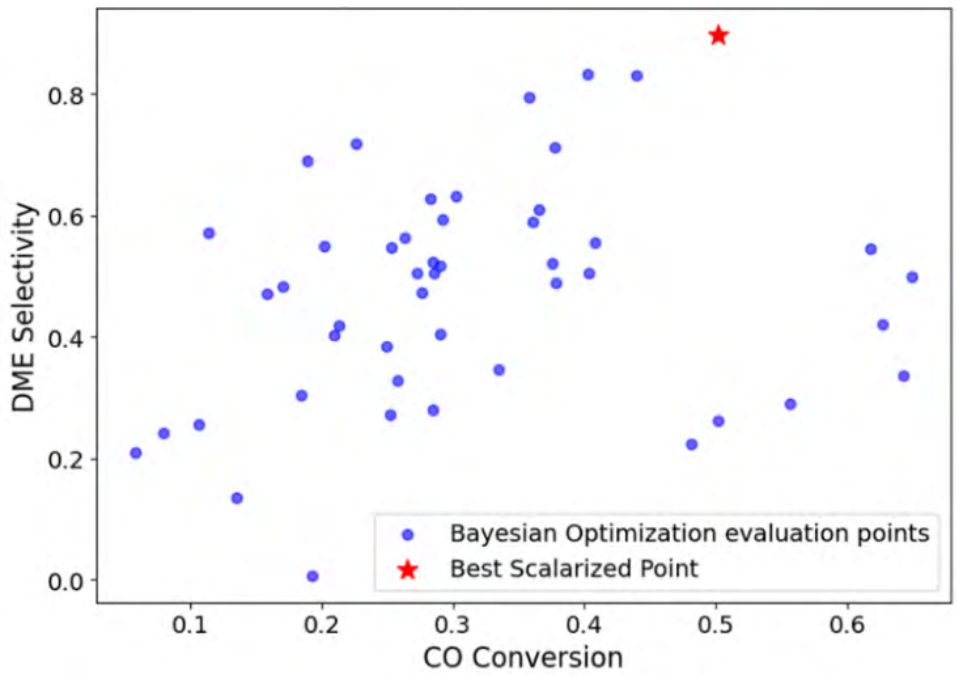

Figure 4. Constrained multi-objective Bayesian optimization.

Table 2: Feature comparison of the best solution and closest real data point.

Feature	Best Fuzzy	Closest real point
Hybrid catalysts	CZ/HZSM-5	
synthesis method	Core-shell	
Oxide catalyst ratio	69:31:0	
Oxide catalyst calcination temperature (°C)	350.8	354.9
Oxide catalyst calcination Time (h)	4.0	67.7
Zeolite SiO2/Al2O3ratio	22.0	
surface area(m2/g)	60.4	67.0
Reduction gas concentration (H2 %)	90.0	100
Reduction temperature (°C)	270	270
Reduction time (h)	1.7	1.0
Reaction temperature (°C)	261	260
Reaction pressure (bar)	37	40
Space velocity (h-1)	14181	6000
Catalyst amount (g)	0.7	0.5
CO conversion (%)	53	51
DME selectivity (%)	70	70

A detailed comparison in Table 2 shows that catalyst formulation and preparation parameters are closely aligned between the fuzzy-optimal solution and the closest real datapoint. In contrast, differences are more pronounced in operating conditions, particularly space velocity and reaction pressure. These deviations suggest that the optimization framework primarily exploits operational flexibility to improve performance while remaining consistent with experimentally validated catalyst formulations.

Bayesian Optimization Results

The constrained multi-objective Bayesian optimization results are illustrated in Figure 4, where each evaluated point represents a surrogate-predicted combination of CO conversion and DME selectivity. The best scalarized solution, indicated by the red star, reflects the optimal compromise according to the equal-weighted objective function.

As summarized in Table 3, Bayesian optimization also identifies CZ/HZSM-5 catalyst prepared by a core-shell method as optimal, confirming consistency with the evolutionary optimization results. However, the recommended operating conditions differ, with a lower reaction temperature (240°C) and a higher space velocity. The predicted performance at this point is close to 47% CO conversion and 65% DME selectivity, which is lower than the fuzzy-enhanced NSGA-II solution. The difference can be attributed to the scalarization and uncertainty-aware nature of Bayesian optimization, which tends to favour conservative solutions when model uncertainty is present. While this approach efficiently explores the design space, it may sacrifice peak performance in favour of robustness.

Table 3: Feature comparison of the best solution from multi-objective Bayesian optimization.

Feature	Best Bayesian optimized point
Hybrid catalysts	CZ/HZSM-5
synthesis method	Core-shell
Oxide catalyst ratio	69:31:0
Oxide catalyst calcination temperature (°C)	360.00
Oxide catalyst calcination Time (h)	4.0
Zeolite SiO2/Al2O3ratio	27.0
surface area(m2/g)	151
Reduction gas concentration (H2 %)	24
Reduction temperature (°C)	261
Reduction time (h)	2.71
Reaction temperature (°C)	240
Reaction pressure (bar)	38
Space velocity (h-1)	11100
Catalyst amount (g)	0.9
CO conversion (%)	46.6
DME selectivity (%)	64.5

Comparative Assessment

A direct comparison of results in Tables 2 and 3 show that both optimization strategies converge toward similar catalyst formulations, indicating that catalyst-related decisions are robust across optimization paradigms. In contrast, operating conditions vary more significantly between methods, leading to differences of up to 6-7 percentage points in predicted CO conversion and approximately 5-6 percentages points in DME selectivity.

This quantitative comparison highlights the central contribution of the present study: different surrogate-based multi-objective optimization strategies, when applied to the same dataset and predictive model, can lead to distinct yet plausible optimal operating conditions. Such differences are particularly relevant for decision-making under limited data availability, where the choice of optimization framework directly influences recommended process conditions.

CONCLUSION

This study presents a surrogate-based multi-objective optimization framework for DME synthesis from syngas that integrates ML with two distinct optimization strategies. Gradient Boosting is identified as a reliable surrogate model capable of capturing nonlinear relationships within a limited experimental dataset. Both fuzzy-enhanced NSGA-II and multi-objective Bayesian optimization successfully identify feasible and competitive operating conditions. While the two approaches agree closely on catalyst formulations, they differ in their recommended operating conditions and predicted performance levels. These differences highlight the importance of selecting an optimization strategy that aligns with the desired balance between exploration, exploitation, and decision-making under uncertainty. These results demonstrate that comparative analysis of optimization paradigms applied to the same surrogate model and dataset provides valuable insight into the robustness and limitations of data-driven decision-support tools for complex catalytic processes such as DME synthesis. Future studies will include targeted experimental validation of the optimized conditions identified by both optimization strategies, with particular emphasis on assessing robustness under realistic process variability. Extending the framework to include additional objectives, such as catalyst stability, by-product formation, or energy efficiency, would enable a more comprehensive assessment of industrial relevance.

ACKNOWLEDGEMENTS

This research was funded by: Project LOCATE - Engineering Low Carbon Technologies for Green Dimethyl Ether, Funding: German Academic Exchange Service (Deutscher Akademischer Austauschdienst - DAAD) under the Programm des Projektbezogenen Personenaustauschs Brasilien CAPES ab 2025 (Project number 57754630)

AUTHOR IDENTIFIERS

Jafari M: 0009-0004-2714-3500
Dorneanu B: 0000-0003-3553-6625
Pontes K: 0000-0002-8208-0854
Arellano-Garcia H: 0000-0002-8297-0232

REFERENCES

1. Amoo CC, Xing C, Tsubaki N, Sun J. Tandem reactions over zeolite-based catalysts in syngas conversion. ACS Cent. Sci. 8:1047-1062 (2022). https://doi.org/10.1021/acscentsci.2c00434
2. Livescu A, Navar R, Mangalindan JR, Mahnaz F, Ge Y, Shetty M, Yang X. Catalysts for clean energy: a review on current progress for the catalyzed recycling of CO2 into dimethyl ether. Top Catal 69:1633-1654 (2024). https://doi.org/10.1007/s11244-024-01913-z
3. Peinado C, Liuzzi D, Sluijter SN, Skorikova G, Boon J, Guffanti S, Groppi G, Rojas S. Review and perspective: next generation DME synthesis technologies for the energy transition. Chemical Engineering Journal 479:147494 (2024). https://doi.org/10.1016/j.cej.2023.147494

4. Das M, Ghosh A, Sunoj RB. Advances in machine learning with chemical language models in molecular property and reaction outcome predictions. J Comput Chem 45:1160-1176 (2024). https://doi.org/10.1002/jcc.27315
5. Singh S, Sunoj RB. Molecular machine learning for chemical catalysis: prospects and challenges. Acc. Chem. Res. 56:402-412 (2023). https://doi.org/10.1021/acs.accounts.2c00801
6. Shim E, Tewari A, Cernak T, Zimmerman PM. Machine learning strategies for reaction development: toward the low-data limit. J. Chem. Inf. Model. 63:3659-3668 (2023). https://doi.org/10.1021/acs.jcim.3c00577
7. Bian K, Priyadarshi R. Machine learning optimization techniques: a survey, classification, challenges, and future research issues. Arch Computat Methods Eng : (2024). https://doi.org/10.1007/s11831-024-10110-w
8. Yazdinejad A, Dehghantanha A, Parizi RM, Epiphaniou G. An optimized fuzzy deep learning model for data classification based on NSGA-II. Neurocomputing 522:116-128 (2023). https://doi.org/10.1016/j.neucom.2022.12.027
9. Garnett R. Bayesian optimization. Cambridge University Press (2023). https://doi.org/10.1017/9781108348973
10. Ramos MC, Michtavy SS, Porosoff MD, White AD. Bayesian optimization of catalysts with in-context learning. arXiv preprint arXiv:230405341. 2023;
11. Kumar A, Pant KK, Upadhyayula S, Kodamana H. Multiobjective bayesian optimization framework for the synthesis of methanol from syngas using interpretable gaussian process models. ACS Omega 8:410-421 (2022). https://doi.org/10.1021/acsomega.2c04919
12. Ma H, Zhang Y, Sun S, Liu T, Shan Y. A comprehensive survey on NSGA-II for multi-objective optimization and applications. Artif Intell Rev 56:15217-15270 (2023). https://doi.org/10.1007/s10462-023-10526-z
13. Blank J, Deb K. Pymoo: multi-objective optimization in python. IEEE Access 8:89497-89509 (2020). https://doi.org/10.1109/access.2020.2990567
14. Santiago A, Dorronsoro B, Nebro AJ, Durillo JJ, Castillo O, Fraire HJ. A novel multi-objective evolutionary algorithm with fuzzy logic based adaptive selection of operators: FAME. Information Sciences 471:233-251 (2019). https://doi.org/10.1016/j.ins.2018.09.005
15. Nakamura K, Otani N, Koike T. Multi-objective bayesian optimization of optical glass compositions. Ceramics International 47:15819-15824 (2021). https://doi.org/10.1016/j.ceramint.2021.02.155
16. Nakamura K, Otani N, Koike T. Multi-objective bayesian optimization of optical glass compositions. Ceramics International 47:15819-15824 (2021). https://doi.org/10.1016/j.ceramint.2021.02.155
17. Shields BJ, Stevens J, Li J, Parasram M, Damani F, Alvarado JIM, Janey JM, Adams RP, Doyle AG. Bayesian reaction optimization as a tool for chemical synthesis. Nature 590:89-96 (2021). https://doi.org/10.1038/s41586-021-03213-y

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Decomposition of MINLP Formulations in Process Family Design using Progressive Hedging

Ali Asger[a], Bernard Knueven[b], and Carl Laird[a]*

[a] Carnegie Mellon University, Pittsburgh, PA, USA
[b] National Laboratory of the Rockies, Golden, CO, USA
* Corresponding Author: claird@andrew.cmu.edu.

ABSTRACT

Distributed deployment of process systems can benefit from modularity and shared components across multiple variants, reducing both manufacturing costs and engineering effort. Process family design formalizes this idea by simultaneously optimizing a family of process variants while determining a shared platform of common components. This results in a large-scale mixed-integer nonlinear program (MINLP) that couples nonlinear process models with discrete platform-allocation decisions. In this work, we solve the process family design MINLP using a progressive hedging (PH)-based decomposition strategy that exploits its block-angular structure. To improve convergence for this nonconvex problem, we introduce dynamic gradient-based penalty updates, a decoupled primal-dual strategy via separate PH runs, and parallel optimization-based bounds tightening of first-stage variables. Computational results on a water desalination case study demonstrate that the proposed approach improves solution quality and reduces computational time compared to baseline PH, while achieving a lower total cost of the process family than the previous discretization-based MILP formulation [1]. These results highlight the effectiveness of tailored decomposition strategies for large-scale process family design problems.

Keywords: Process Family Design, Mixed Integer Nonlinear Programming, Progressive Hedging

INTRODUCTION

Increasing demand for energy and escalating resource depletion is going to require rapid and large-scale deployment of emerging process technologies. Existing process design methodologies are largely driven by economies of scale, with each process installation optimized independently to minimize capital and operating costs. While this paradigm is well suited for larger capacity single-site deployment, it becomes less effective for distributed deployment applications, where manufacturing standardization can play a significant role in reducing overall cost and engineering effort. In such contexts, modular design strategies that exploit standardization offer a promising alternative.

Process Family Design [1] has recently been proposed as a systematic optimization-based framework to support such modular design strategies while maintaining flexibility. Rather than optimizing a single instance of a process system, we optimize a family of process variants of a process system while promoting the use of shared components drawn from a common platform. A key feature of the optimization formulation is that the platform itself, i.e. the selection and sizing of shared components, is co-optimized with the individual process variants. This coupling enables trade-offs between individual process variant performance and the overall platform commonality to be explicitly captured within the optimization problem.

The resulting optimization problem is inherently a challenging large-scale mixed-integer nonlinear program (MINLP). Continuous nonlinear constraints arise from the process system model, e.g. process physics and cost correlations, while discrete decisions govern the platform assignment across the family of variants. The complexity arises not only from the nonlinearities within the individual variants but also from the coupling constraints that enforce platform commonality across the family. As the number of variants and potential platform components increases, the platform size grows rapidly, often rendering monolithic global optimization approaches intractable.

In prior work, Stinchfield et al [1] addressed the computational challenge of solving the complex MINLP by discretizing the design space of common modules on the platform. By pre-computing the feasibility of candidate designs and converting the problem into a mixed-integer linear program (MILP), the global optimum of the discretized space could be found efficiently. However, the discretization approach presents several limitations. First, it requires expensive pre-computation to evaluate the Cartesian product of all candidate platform designs across all variants, i.e. simulate each combinatoric instance of discretized candidate designs for each common module on the platform across each variant of the family. Second, the solution quality is limited by the resolution of the discretization of the design space; improving accuracy requires a finer grid, which reintroduces combinatorial explosion.

More recently, a decomposition strategy has been applied to the discretization-based process family design formulation. Stinchfield et al [2] leverage the block-angular structure arising from shared platform variables to decompose the MILP formulation and improve computational tractability. While this approach demonstrates the value of decomposition for large process family design problems, it remains dependent on discretization and the associated pre-computation. In this work, we address these limitations by formulating and solving the process family design problem directly as an MINLP. Building on prior decomposition ideas, we propose a decomposition strategy based on Progressive Hedging (PH) [3], a method used for stochastic programming but well-suited here due to the block-angular structure of the formulation. In contrast to applying PH to the discretization-based MILP, we apply PH directly to the original MINLP by relaxing the complicating platform consensus constraints and penalizing deviations. This enables full decomposition of the family of process variants into smaller subproblems that can be solved independently.

While PH is heuristic when applied to nonconvex problems [3], we demonstrate that a tailored algorithmic setup significantly improves convergence and solution quality. Specifically, we implement a dynamic penalty parameter ρ update scheme that utilizes gradient information at each iteration to drive the variants towards a consensus platform design. This framework is implemented using the open-source package *mpi-sppy* [4], enabling parallel solution of the subproblems. We validate this approach on a desalination process family, comparing the computational performance and solution optimality against both off-the-shelf solvers and the previous MILP discretization method.

BACKGROUND

Various product family design approaches have been extensively investigated across a wide range of industrial applications [4, 5]. In general, optimization approaches for nonlinear product family design have been focused on stochastic methods, e.g. genetic algorithms [7]. Although effective for yielding solutions to challenging problems in practice, there are no guarantees on the quality of solutions obtained.

Relatively few works have utilized mathematical programming techniques to the product family design problem. Khajavirad and Michalek [8] proposed an MINLP-based optimization approach for solving the nonlinear product family design problem, however they considered a fixed platform design. We are not aware of any contribution that directly addresses the MINLP formulation for simultaneous platform and family design.

To address the scalability limits of the process family design MINLP formulation, decomposition methods are required. The process family design problem exhibits a *block-angular* structure analogous to two-stage stochastic programming, where the variant-specific variables are independent ("second-stage"), and the platform-specific variables are coupled ("first-stage") and must be identical across the family ("non-anticipativity").

Progressive hedging (PH) [3] is a horizontal decomposition algorithm that is well-suited for this structure. By relaxing the coupling constraints and penalizing deviations from the group average, PH decomposes the problem into independent subproblems that can be solved in parallel. For a two-stage optimization, the PH algorithm can be stated as described in Algorithm 1.

Algorithm 1: Progressive hedging for two-stage stochastic programming

1. **Initialization**: Let $\nu \leftarrow 0$ and $w^{\nu} \leftarrow 0, \forall s \in \mathrm{S}$. Compute for each $s \in \mathrm{S}$

$$x^{(\nu+1)}(s) \in \operatorname{argmin}_x f_1(x) + f_2(y; \vec{x}(s))$$

   where, $f_1(x)$ is the first-stage cost and $f_2(y; \vec{x}(s))$ is the second stage cost.

2. **Iteration Update**: $\nu \leftarrow \nu + 1$
3. **Aggregation**: $\bar{x}^{(\nu)} \leftarrow \sum_{s \in S} \pi_s x^{(\nu)}(s) / \sum_{s \in S} \pi_s$
4. **Price Update**: Compute for each $s \in \mathrm{S}$

$$w^{(\nu)}(s) \leftarrow w^{(\nu-1)} + \rho[x^{\nu}(s) - \bar{x}^{(\nu)}]$$

5. **Decomposition**: Compute for each $s \in \mathrm{S}$

$$x^{(\nu+1)}(s) \in \operatorname{argmin}_x f_1(x) + f_2\big(y; \vec{x}(s)\big) + w^{(\upsilon)}(s)^{\mathrm{T}} x + \frac{\rho}{2}\left\|x - \bar{x}^{(\nu)}\right\|^2$$

6. **Termination**: If a criterion is met, Stop. Otherwise go to step 2.

The PH algorithm is initialized by independently

solving all the subproblems. Each iteration of the algorithm begins with an aggregation step, in which the probability-based weighted primal mean $\bar{x}^{(\nu)}$ is computed, which provides an estimate of the optimal decision vector that satisfies the non-anticipativity constraints. The associated dual variables $w^{(\nu)}(s)$ act as Lagrange multipliers and are updated using the primal mean $\bar{x}^{(\nu)}$, the current values of the non-anticipative variables $x^{\nu}(s)$, and quadratic penalty parameter ρ. The algorithm subsequently decomposes by subproblem: each subproblem is resolved with its objective function augmented by the updated dual terms and a quadratic penalty term. Termination of the PH algorithm may be based on internal PH criteria, such as relative optimality gap or a specified tolerance on violations of non-anticipativity constraints.

However, while PH is guaranteed to converge for convex problems, it acts as a heuristic for non-convex MINLPs [3]. The efficiency of the PH algorithm convergence is strongly impacted by the choice of the penalty parameter ρ [9]. In particular, the value of the penalty parameter ρ influences both the primal decision variables and the updates of dual weights [10]. In the update of dual price $w^{(\nu)}(s)$ (Step 4), the step size based on deviation from the mean is scaled by ρ. As a result, ρ directly controls the rate at which dual variables progress. In the following decomposition step, for each $s \in \mathrm{S}$, the algorithm minimizes an augmented objective, where the two terms of the augmented part $(w^{(\upsilon)}(s)^{\mathrm{T}}x + \frac{\rho}{2}\left\|x - \bar{x}^{(\nu)}\right\|^2)$ exert competing influences during optimization: the dual term $(w^{(\upsilon)}(s)^{\mathrm{T}}x)$ drives decision variables towards values favored by dual multipliers, and the quadratic term $(\frac{\rho}{2}\left\|x - \bar{x}^{(\nu)}\right\|^2)$ promotes proximity to the mean decision vector. Consequently, varying the penalty parameter ρ alters the trajectory of the primal variables and overall convergence of the algorithm. Larger values of ρ amplify the contribution of the quadratic penalty, leading to aggressive updates to primal convergence, whereas smaller values of ρ yield a smoother and gradual evolution of dual variables, which may be preferable in contexts such as lower bound computation.

Setting a standard static penalty parameter across first-stage variables often leads to slow convergence or cycling in non-convex applications. Watson and Woodruff [9] proposed a variable-specific heuristic for calculation of the penalty parameters for each variable based on their unit-cost, i.e. coefficient in the objective term. They further noted that per-iteration ρ values may be more appropriate for some problems. Building on this work, Naepels and Woodruff [10] combined the heuristic with a gradient-based approach to calculate ρ values and introduced a method to iteratively compute variable-dependent ρ values based on a first-order condition, allowing updates as the algorithm proceeds.

In this work, we build upon these penalty parameter ρ selection strategies within the PH framework. While prior approaches compute gradients only at the initial iteration under the assumption that the relevant objective terms are linear, and assume the availability of a feasible solution, either from an incumbent or starting point that must be specified by the user a priori in existing *mpi-sppy* implementations, we instead address nonlinear objectives by computing gradients at every iteration using the current primal iterates and then switching to incumbents as they are found. Moreover, motivated by challenges posed by the nonlinear nonconvex nature of the process family design MINLP, we extend these ideas by developing decoupled primal-dual penalty parameter ρ strategies that address the interactions between primal and dual convergence, leading to improved overall PH convergence.

PROBLEM SETUP

We consider the process family design problem as formalized by [1]. In this section, we explain how the process family design problem can be formulated as a nonlinear generalized disjunctive program (GDP) and subsequently transformed into a mixed integer nonlinear program (MINLP).

Problem Description

The goal of process family design is to simultaneously design a large number of variants of a given process system that can be deployed across a wide range of installations. For example, deploying a carbon capture system at multiple industrial sites with varying flue gas sources. While the general flowsheet configuration of the carbon capture process system will remain the same at each site, the design requirements for each site (e.g. capture rate, flue gas concentration) can be different. Thus, every variant $\upsilon \in V$ will be designed to meet specific requirements; parameterized in $\boldsymbol{r}_{\upsilon}$.

A process system is defined by all the *unit module types* necessary to build an instance of the process; defined by $m \in M$. The unit module types could include absorber, heat exchanger, regenerator, etc. Each variant $\upsilon \in V$ requires exactly one of each unit module type $m \in M$. The *unit module type design* for each unit module type m and each variant is defined by variable vector $\boldsymbol{d}_{\upsilon,m}$.

To exploit opportunities for shared components (and engineering design effort) we consider a platform of common unit modules that can be shared across the family of process variants. The unit module types M are separated into *commonly design unit modules* (set C included on the platform) and *uniquely designed unit modules* (set U), where sets C and U are disjoint ($C \cap U = \emptyset$) and their union gives the set M ($C \cup U = M$). The common designs making up the platform $\mathcal{P}$ are indexed by label

$l \in L_c$ to capture different candidate designs for each common module type $c \in C$. The design specifications for these common unit modules on the platform are captured by $\hat{\boldsymbol{d}}_{c,l}$. All unique unit module types $u \in U$ are designed variant specific.

Problem Formulation

The process system design of every single variant $v \in V$ with different set of design requirements $\boldsymbol{r}_v$ is governed by a set of nonlinear (in most applications) equations that describe the cost, physics, and performance of the process system. Our objective is to minimize the overall annualized cost of the process family while simultaneously determining (a) a platform with optimal designs of common unit module types for variants to choose from, (b) the distribution of these common unit module designs among the family of process variants, and (c) the variant-specific design variables for modules not in the platform and the optimal operating conditions for each of the variants.

Thus, the process family design problem can be formulated as a nonlinear generalized disjunctive program (GDP) that can be transformed into a mixed integer nonlinear program (MINLP):

$$\min_{p,i,y,d,\hat{d},o} \sum_{v \in V} w_v p_v \tag{1}$$

$$\text{s.t.} \quad p_v = f_v^p(\boldsymbol{r}_v, \boldsymbol{d}_v, \boldsymbol{o}_v) \;\; \forall\, v \in V \tag{2}$$

$$\boldsymbol{i}_v = f_v^i(\boldsymbol{r}_v, \boldsymbol{d}_v, \boldsymbol{o}_v) \;\; \forall\, v \in V \tag{3}$$

$$0 = h_v(\boldsymbol{r}_v, \boldsymbol{d}_v, \boldsymbol{o}_v) \;\; \forall\, v \in V \tag{4}$$

$$\bigvee_{l \in L_c} \begin{bmatrix} Y_{v,c,l} \\ \boldsymbol{d}_{v,c} = \hat{\boldsymbol{d}}_{c,l} \end{bmatrix} \;\; \forall\, v \in V, c \in C \tag{5}$$

$$\hat{\boldsymbol{d}}_{c,l-1} \le \hat{\boldsymbol{d}}_{c,l} \quad \forall\, c \in C, l \in L_c \backslash \{1\} \tag{6}$$

$$\boldsymbol{d}_v^L \le \boldsymbol{d}_v \le \boldsymbol{d}_v^U \quad \forall\, v \in V \tag{7}$$

$$\hat{\boldsymbol{d}}_{c,l}^L \le \hat{\boldsymbol{d}}_{c,l} \le \hat{\boldsymbol{d}}_{c,l}^U \quad \forall\, c \in C, l \in L_c \tag{8}$$

$$\boldsymbol{o}_v^L \le \boldsymbol{o}_v \le \boldsymbol{o}_v^U \quad \forall\, v \in V \tag{9}$$

$$\boldsymbol{i}_v^L \le \boldsymbol{i}_v \le \boldsymbol{i}_v^U \quad \forall\, v \in V \tag{10}$$

$$Y_{v,c,l} \in \{True, False\} \;\; \forall\, v \in V, c \in C, l \in L_c \tag{11}$$

The objective in (1) minimizes the weighted annualized cost aggregated over all variants $v \in V$. Constraints (2-4) embed, for each variant, the process model equations, including the cost (p_v), performance indicators ($\boldsymbol{i}_v$), and process system equations as functions of the design requirements $\boldsymbol{r}_v$, design variables $\boldsymbol{d}_v$ and operating variables $\boldsymbol{o}_v$. The formulation introduces commonality across variants through explicit logical constraints governing the selection of shared unit module designs. Constraint (5) employs disjunctions and associated Boolean variables $Y_{v,c,l}$ to assign common unit modules to variants by enforcing equality among the relevant design variables. Constraint (6) enforces a size-based ordering of the platform-level module designs as a symmetry-breaking constraint. Constraints (7 - 10) set bounds on the design variables $\boldsymbol{d}_v$, common designs $\hat{\boldsymbol{d}}_{c,l}$, operating variables $\boldsymbol{o}_v$., and performance indicators $\boldsymbol{i}_v$. Constraint (11) specifies the binary nature of the assignment variables.

PROGRESSIVE HEDGING SETUP

Problem Decomposition

Decomposition algorithms involve splitting a large problem into smaller problems and solving them iteratively until a convergence criteria is met. Stochastic programs inherently possess a block-angular structure that allows for the decomposition of a problem into smaller sub-problems. In this structure, first-stage variables link the subproblems through non-anticipativity constraints, while second-stage variables are specific to each subproblem.

Our problem can be represented with a similar block-angular structure, as shown in Figure 1. We decompose our problem into subproblems, each representing a smaller process family design problem, with the constraint that all subproblems share the same platform design. Therefore, our first-stage variables are continuous common unit module design variables in the platform ($\hat{\boldsymbol{d}}_{c,l}$) and the second-stage variables are all the local variables associated with each process family design subproblem ($p_v, \mathbf{i}_v, \mathbf{r}_v, \mathbf{d}_{v,m}, \mathbf{o}_v, \mathrm{Y}_{v,c,l}$).

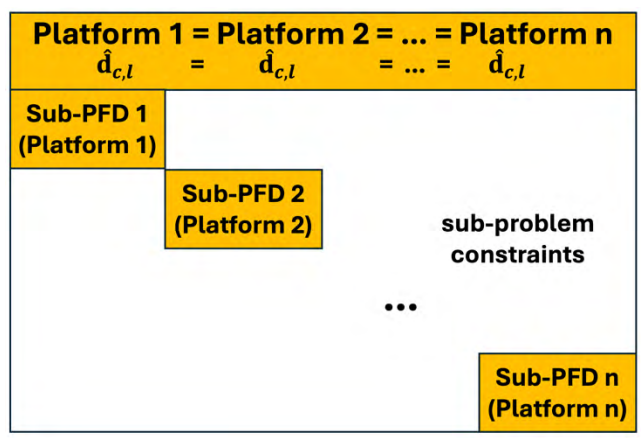

Figure 1. Block-angular structure of the process family design problem

Dynamic Gradient-based ρ Updates

The convergence of PH for non-convex MINLPs is highly sensitive to the penalty parameter ρ. A static global ρ is often insufficient: values that are too small result in slower convergence and failure to enforce consensus, while values that are too large cause the algorithm to cycle or lock into suboptimal local minima. To address this, we implement a dynamic, variable-specific update

method based on the first-order optimality conditions, as proposed by Naepels and Woodruff [10]:

$$\rho^{\nu}(s) = \left[c_D - w^{(\nu)}(s)\right] / \left[x^{(\nu)}(s) - \bar{x}^{(\nu)}\right] \quad (12)$$

with $c_D = -\nabla f(\hat{x})$

However, the effectiveness of this update rule depends critically on the availability of accurate gradient information. Naepels and Woodruff [10] compute gradients only at the first iteration under the assumption of a linear objective, and thus the gradients remain constant throughout the PH iterations. In contrast, the process family design MINLP involves nonlinear terms in the objective, for which gradient information varies with iterations. Consequently, restricting gradient evaluation to the initial iteration can lead to inaccurate updates and degraded convergence behaviour. To address this limitation, we compute gradients dynamically at every PH iteration using the current primal iterates. Specifically, the gradient term c_D in (12) is evaluated as $c_D = -\nabla f(x^{(\nu)})$, until a feasible incumbent solution becomes available. Once an incumbent is found, gradient evaluations are switched to the incumbent solution to stabilize updates and reduce oscillations in later iterations. This dynamic gradient evaluation enables penalty parameter ρ to adapt to local curvature in the nonlinear objective, improving convergence and robustness of PH for non-convex MINLPs.

While dynamic gradient-based updates provide a well-motivated mechanism for adapting the penalty parameter ρ, it also introduces increased sensitivity to primal convergence. In practice, the non-anticipative variables $x^{\nu}(s)$ may rapidly approach their consensus values during early PH iterations, after which convergence slows. In such cases, the difference $|x^{\nu}(s) - \bar{x}^{(\nu)}|$ becomes very small, causing the update in (12) to produce disproportionately large values of ρ. Overly large values of ρ are ineffective and can stall convergence in both primal and dual variables. To prevent this behaviour, Naepels and Woodruff [10] also provide explicit criteria to determine when ρ should be updated based on measures of primal and dual convergence.

Decoupled Primal-Dual Strategy

The dynamic gradient-based updates described above yields subproblem-dependent penalty parameters $\rho^{\nu}(s)$, which must be aggregated into a single subproblem-independent value to apply progressive hedging. Consequently, a rule is needed to reconcile the subproblem-specific values of $\rho^{\nu}(s)$ computed from (12).

Naepels and Woodruff [10] examine several aggregation strategies for subproblem-dependent penalty parameters, including minimum, average, and maximum values across subproblems. Consistent with preceding discussion, different aggregation rules place different emphasis on enforcing primal consensus versus stabilizing dual updates. Building on this observation, we adopt a decoupled primal-dual strategy by solving two progressive hedging procedures with distinct aggregation rules for the penalty parameter $\rho^{\nu}(s)$.

In the first PH procedure, the aggregation of the subproblem-dependent penalty parameters is chosen to promote primal consensus. Specifically, an average aggregation rule is applied,

$$\rho^{\nu}_{primal} = \sum_{s \in S} \rho^{\nu}(s) / |S| \quad (13)$$

which smooths subproblem-to-subproblem variability while maintaining sufficient penalty magnitude to enforce agreement among the nonanticipative variables. This configuration prioritizes convergence of primal decision variables.

In the second PH procedure, the aggregation rule is selected to favor stable dual behaviour. In this case the penalty parameter is defined using the minimum positive values across subproblems,

$$\rho^{\nu}_{dual} = \min\{\rho^{\nu}(s) : \rho^{\nu}(s) > 0, s \in S\} \quad (14)$$

This choice limits the magnitude of the multiplier updates and mitigates oscillations in the dual variables, particularly in later iterations when subproblem solutions are near consensus.

The two PH procedures are executed independently, each using a single, subproblem-independent penalty parameter per iteration, as required by the progressive hedging framework. Together, these complementary strategies provide a practical mechanism for improving convergence behaviour in the presence of nonconvexities and discrete decisions. Although this approach does not alter the fundamental structure of the PH algorithm, it allows different aggregation strategies to be exploited. In the results presented later, this decoupled strategy is shown to improve convergence relative to a single PH configuration.

Optimization-Based Bounds Tightening for First-Stage Variables

Bounds tightening is a critical refinement technique to strengthen relaxations and improve convergence used for solution of MINLPs [11]. Among the available approaches, optimization-based bounds tightening (OBBT) is particularly effective, albeit computationally expensive. In this work, we exploit the two-stage structure of the problem and the sensitivity of progressive hedging convergence to the bounds of the first-stage (nonanticipative) variables by applying OBBT exclusively to these variables as a pre-solve routine.

The OBBT procedure is applied to the first-stage variables independently for each subproblem. For a given first-stage variable, the lower and upper bounds are tightened by solving the auxiliary optimization problems

that minimize and maximize the variable subject to the subproblem constraints. These bound-tightening problems are solved in parallel across all subproblems, exploiting the same decomposition structure as the PH framework.

Following the solution of the OBBT subproblems, the resulting bounds are aggregated by selecting the tightest valid bounds across the subproblems. Specifically, for each first-stage variable, the updated lower bound is taken as the maximum of the subproblem-specific lower bounds, and the updated upper bound is taken as the minimum of the subproblem-specific upper bounds. This aggregation ensures that the tightened bounds remain feasible for all subproblems while reducing the feasible region of the shared variables.

The bound-tightening process is applied iteratively. After each round of OBBT, the aggregated bounds are updated, and the procedure is repeated until successive bound improvements fall below a specified tolerance. Once the updates become insignificant, the tightened bounds are fixed and used to start the PH algorithm.

By performing OBBT in parallel and restricting it to the first-stage variables, the additional computational overhead remains minimal, while accelerating convergence of PH.

CASE STUDY

We now discuss the water desalination case study and the specific process family design problem we used to demonstrate our optimization approach. The water desalination process system model was developed under the PARETO framework [12].

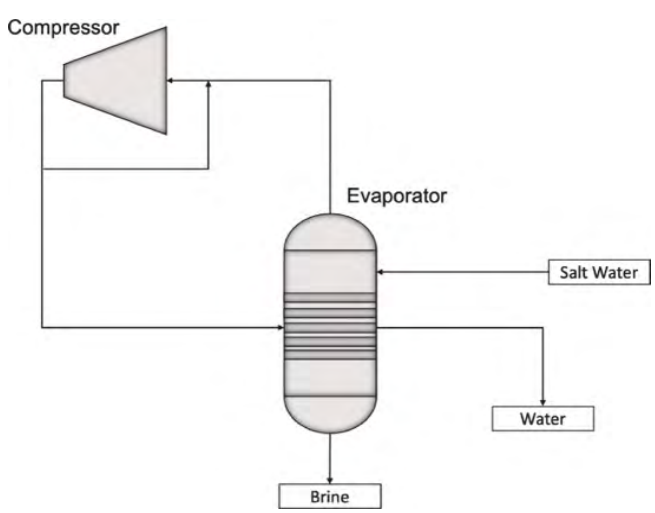

Figure 2. Water Desalination System

Around 60 million barrels of produced water per day is expected from oil and gas fracking across the United States by 2030, which will require treatment of produced water (i.e. desalination) at scale [12]. The salinity and flowrate of the produced water have high variability across wells in different geographical regions. One potential desalination technology is to first evaporate, and then mechanically recompress the vapor [13]. In our case study, we consider a desalination system with a single-effect evaporator and a single-stage adiabatic compressor as shown in Figure 2.

Building on the case study from previous work [1], we consider the process family design problem that accommodates a range of varying produced water feed conditions across the Midland, Texas area. Therefore, the set of process variants $v \in V$ is defined by different combinations of feed flowrates and feed salt concentrations. In this case study, we consider desired specifications of feed salt concentration between 20g/kg and 150g/kg and feed flowrate between 0.1 kg/s and 1 kg/s. Overall, a set of 76 process variants is considered for this process family design problem.

The compressor and the evaporator are the common unit module types to be designed in the platform, i.e. M = C = [compressor, evaporator], with design variables of compressor design flow (between 0.05 kg/s and 1 kg/s) and evaporator area (between 20 m^2 and 400 m^2).

RESULTS & DISCUSSIONS

This section presents the results of the case study introduced in the preceding section by employing the decomposition approach and the algorithmic methodology described in the Progressive Hedging Setup section. The key findings derived from these results are subsequently discussed.

The process family was designed using the optimization formulation described in (1-11). The full process family design problem of 76 variants was decomposed into 19 smaller process family subproblems with 4 variants each. We consider a platform specification that allows 3 designs, i.e. $|L_c|$ = 3, for both the compressor and evaporator. The variants were ordered by size ($\boldsymbol{r}_v$) and bundled into subproblems.

Figure 3 describes which combinations of compressor and evaporator designs offered in the platform $\mathcal{P}$ are assigned to each variant $v \in V$. The x-axis corresponds to the feed flow rate that describes a particular variant v; the y-axis corresponds to the concentration of salt in the feed at variant v.

All the computational studies were performed on the TRACE HPC cluster at Carnegie Mellon University using a single node. The node has two 64-core AMD EPYC 7713 processors (2 GHz) and 2 TB RAM. To ensure consistency in our results, we limit Gurobi© to a maximum of 8 threads.

Table 1 compares the performance of three solution strategies for the water desalination case study: (1) progressive hedging with $\hat{x}$-based gradient penalty update

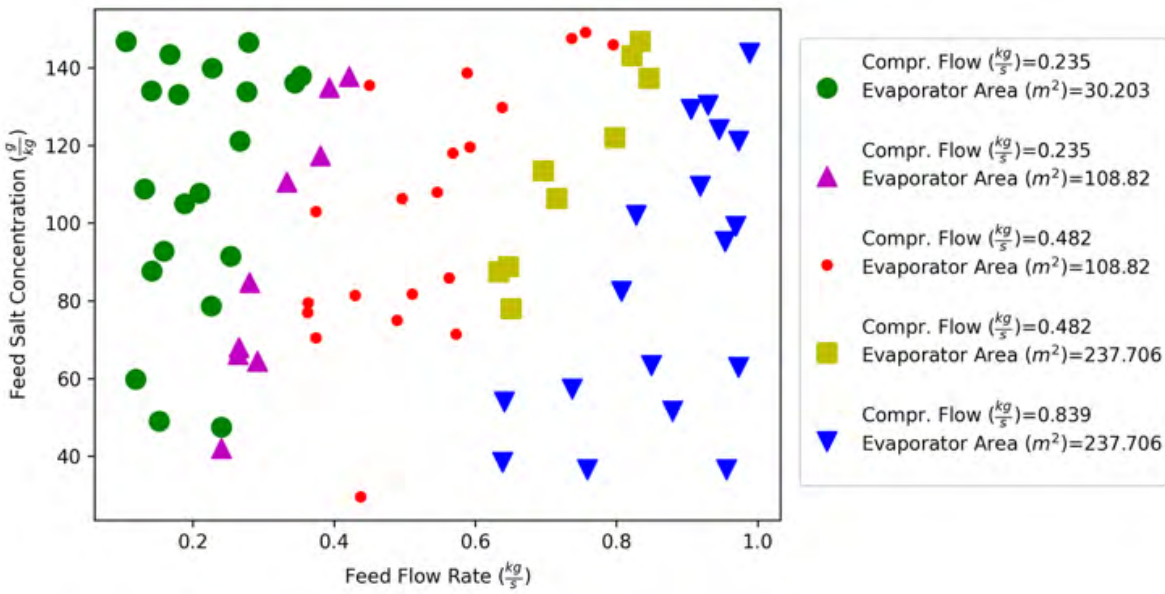

Figure 3. Water Desalination Process Family

following Naepels and Woodruff [10], (2) the proposed progressive hedging approach with dynamic $x^{(\nu)}$-based gradient penalty updates, optimization-based bounds tightening, and a decoupled primal-dual strategy, and (3) a direct MINLP solution using Gurobi© 13. For the progressive hedging-based approaches, the algorithm was run for 50 PH iterations.

Table 1: Comparison of different solution strategies on the water desalination case study.

Solution Strategy	Incumbent Objective ($M)	Dual Bound ($M)	Solution Time (secs)
PH (Grad $\rho, \hat{x}$-based) [10]	10, 131.93 8850.26	6754.36 7117.96	5372 4435
PH (Grad ρ, OBBT, decoupled $x^{(\nu)}$-based) [This work]			
Gurobi© 13 MINLP	9062.82	7744.95	7200

The baseline PH approach yields the highest incumbent objective value and the weakest dual bound, reflecting the limitations of $\hat{x}$-based gradient penalty updates in nonconvex settings. In contrast, the proposed PH configuration achieves a substantially lower incumbent objective value with reduced solution time, while also producing a stronger dual bound. The improvement in the dual bound is consistent with the intended effect of the decoupled primal-dual strategy, which stabilizes dual multiplier updates while maintaining effective progress in the primal iterates.

The direct MINLP approach using Gurobi© 13 produces a competitive incumbent solution with the imposed 7200 s time limit. However, the proposed PH method attains a better incumbent solution in a shorter computation time. The dual bound reported by Gurobi© 13 is included for context. Overall, the results indicate that the proposed PH-based approach provides an effective balance between primal solution quality, dual bound strength, and computational effort for large-scale process family design problems.

CONCLUSION

This paper investigated the direct solution of process family design problems formulated as large-scale mixed-integer nonlinear programs using a progressive-hedging-based decomposition strategy. To improve the practical performance of progressive hedging for nonconvex MINLPs, dynamic gradient-based ρ updates, a decoupled primal-dual strategy implemented through separate PH configurations, and optimization-based bounds tightening for first-stage variables were introduced. Computational results on a water desalination case study demonstrate that the proposed approach achieves improved primal solution quality, stronger dual bounds, and reduced computational time relative to a baseline progressive hedging implementation, while

achieving an approximately 2.5% reduction in total cost of process family compared to the discretized MILP formulation of Stinchfield et al [1].These results highlight the effectiveness of tailored decomposition strategies for large-scale structured nonconvex MINLPs that are challenging for off-the-shelf solvers. Future work will explore hybrid decomposition strategies that incorporate branching in the space of first-stage variables.

ACKNOWLEDGEMENTS

This work was authored in part by the National Laboratory of the Rockies for the U.S. Department of Energy (DOE), operated under Contract No. DE-AC36-08GO28308. This work was supported by the Laboratory Directed Research and Development (LDRD) Program at the National Laboratory of the Rockies. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

AUTHOR IDENTIFIERS

Author ORCIDs:
Ali Asger: 0000-0003-4434-1141
Bernard Knueven: 0000-0002-3694-6274
Carl Laird: 0000-0001-8430-1561

REFERENCES

1. Stinchfield G, Morgan JC, Naik S, Biegler LT, Eslick JC, Jacobson C, Miller DC, Siirola JD, Zamarripa M, Zhang C, Zhang Q, Laird CD. A mixed integer linear programming approach for the design of chemical process families. Computers & Chemical Engineering 183:108620 (2024). https://doi.org/10.1016/j.compchemeng.2024.108620
2. Stinchfield G, Watson, JP, & Laird CD. Progressive hedging decomposition for solutions of large-scale process family dedign problems. Comput Aided Chem Eng 53:1285-120 (2024) https://doi.org/10.1016/B978-0-443-28824-1.50215-5
3. Rockafellar RT, Wets RJB. Scenarios and policy aggregation in optimization under uncertainty. Mathematics of OR 16:119-147 (1991). https://doi.org/10.1287/moor.16.1.119
4. Knueven B, Mildebrath D, Muir C, Siirola JD, Watson JP, Woodruff DL. A parallel hub-and-spoke system for large-scale scenario-based optimization under uncertainty. Math. Prog. Comp. 15:591-619 (2023). https://doi.org/10.1007/s12532-023-00247-3
5. Simpson TW, Siddique Z, Jiao JR. Product platform and product family design. Springer US (2006). https://doi.org/10.1007/0-387-29197-0
6. Gonzalez-Zugasti JP. Models for platform-based product family design. PhD thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology (2000)
7. D'Souza B, Simpson? TW. A genetic algorithm based method for product family design optimization. Engineering Optimization 35:1-18 (2003). https://doi.org/10.1080/0305215031000069663
8. Khajavirad A, Michalek JJ. A decomposed gradient-based approach for generalized platform selection and variant design in product family optimization. Journal of Mechanical Design 130: (2008). https://doi.org/10.1115/1.2918906
9. Watson JP, Woodruff DL. Progressive hedging innovations for a class of stochastic mixed-integer resource allocation problems. Comput Manag Sci 8:355-370 (2010). https://doi.org/10.1007/s10287-010-0125-4
10. Naepels U, Woodruff DL. Gradient-based rho parameter for progressive hedging. Optimization Online, Aug. 2023.
11. http://optimization-online.org/2023/08/gradient-based-rho-parameter-for-progressive-hedging/
12. Laird C, Bynum M, Castillo A, Knueven B, Siirola J. Decomposing optimization-based bounds tightening problems via graph partitioning.. Proposed for presentation at the 2020 INFORMS Annual Meeting. : (2020). https://doi.org/10.2172/1826429
13. http://optimization-online.org/?p=16954
14. Drouven MG, Caldéron AJ, Zamarripa MA, Beattie K. PARETO: an open-source produced water optimization framework. Optim Eng 24:2229-2249 (2022). https://doi.org/10.1007/s11081-022-09773-w
15. Onishi VC, Carrero-Parreño A, Reyes-Labarta JA, Fraga ES, Caballero JA. Desalination of shale gas produced water: a rigorous design approach for zero-liquid discharge evaporation systems. Journal of Cleaner Production 140:1399-1414 (2017). https://doi.org/10.1016/j.jclepro.2016.10.012

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

CC BY SA

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Physics Constrained Machine Learning for Modeling and Optimization of Chemical Process Systems

Rahul Golder[a], Bimol Nath Roy[a], and M. M. Faruque Hasan[ab]*

[a] Texas A&M University, Artie McFerrin Department of Chemical Engineering, College Station, TX, 77843, USA
[b] Texas A&M Energy Institute, Texas A&M University, College Station, TX, 77843, USA
* Corresponding Author: hasan@tamu.edu

ABSTRACT,

Machine learning (ML) reduces reliance on computationally expensive first-principles simulation while capturing complex nonlinear behaviors. However, poor extrapolation, overfitting, limited interpretability, and lack of strict consistency with governing laws limit the use of ML models in process applications. Current methods for learning optimization policies also struggle with constraint satisfaction and optimality guarantees. Approaches such as physics-informed neural networks (PINNs) incorporate constraints "softly" and do not ensure strict constraint enforcement—an issue that can be particularly detrimental in safety-critical applications, where even minor violations may lead to unsafe or infeasible decisions. To resolve these issues, we develop an ML framework with a differential projection layer that allows computationally efficient process modeling, parameter estimation, and nonlinear optimization with feasibility and optimality guarantees. The framework is general in a sense that, depending on the loss function and the projection layer, it provides different functionalities. For instance, an ANN followed by a distance minimization-based projection leads to a hybrid mechanistic/data-driven process simulator, KKT-HardNet, that ensures inference with strict satisfaction of linear/nonlinear equality/inequality constraints. An ANN with a gradient-based loss function and a Taylor approximation-based pointwise projection leads to DAE-HardNet, which solves and/or estimates parameters of ODE/PDE-based governing equations. Lastly, an ANN with a distance-minimization based projection followed by objective-based loss function leads to KKT-OptNet, which learns to solve nonlinear optimization problems for different parameter instances. We illustrate the use of these tools for the modeling, optimization, and parameter estimation of complex processes systems.

Keywords: AI/ML, Process Modeling, Process Optimization

INTRODUCTION

Artificial intelligence and machine learning (AI/ML) have achieved remarkable success over the past decade, driven by advances in deep learning, large-scale optimization, and data-rich model training. Modern AI/ML systems now surpass human-level performance in perception, language understanding, and pattern recognition, and are increasingly deployed as computational engines for decision-making in scientific, societal, and technological settings. However, contemporary ML models are mostly designed as flexible function approximators, $y = f_\theta(x)$, without mechanisms to "strictly" enforce governing physical laws, conservation principles, or domain constraints. Physics-informed ML [1-2] incorporates these physical priors indirectly—through penalties and soft regularization—which often fail to guarantee feasibility, robustness, or resilience under distribution shift. As a result, even small violations can undermine reliability or lead to catastrophic failures in high-stakes applications such as autonomous nuclear reactors, chemical plants, molecular discovery, and next-generation autonomous cyber-physical infrastructures. These emerging domains require frameworks that are not only fast, but also provide guarantees of feasibility and optimality. Despite significant advances, a foundational gap remains between the dominant paradigm of unconstrained statistical learning and the realities of high-stakes, safety-critical, and resilience-demanding environments.

This disconnect presents a major barrier to the future of trustworthy and resilient computation. A central challenge in scientific machine learning (SciML) is to develop predictive models that are also scientifically reliable, interpretable, and physically valid under known domain knowledge. These requirements can be formalized through constraint sets $S =): h(x, y) = a, g(x, y) \leq b, x \in \Re^n, y \in \Re^m\}$, where $g(\bullet)$ represents inequality constraints, $h(\bullet)$ encodes equality constraints, and *a* and *b* encode multiparametric changing environments. A particular frontier of SciML, where feasibility and rigorous guarantees are essential, is learning to optimize—developing methods that enable the rapid, repeated solution of constrained, multiparametric nonlinear programs (mpNLP) arising in real-time decision environments. They appear in many applications in process systems engineering (PSE) including model predictive control, power grid operations, process safety surveillance, fault detection, and online energy management, all of which demand solutions under stringent constraints. In this work, we address the following questions:

- How can constraints—linear and nonlinear, differential and algebraic—be hard embedded into ANN architectures? What theoretical guarantees are possible for feasibility-aware generalization?
- How can we learn to optimize nonlinear problems with both feasibility and optimality guarantees?
- How can we design integrated learning and optimization frameworks in which constraints, feasibility, and optimality are strictly embedded rather than secondary penalties?

Addressing these requires new foundations that unify learning and optimization theory with algorithm design for the next-generation of trustworthy, reliable AI/ML applications in PSE, which is the core motivation of this work. The core of this is to establish constraint-embedded ML as a foundational paradigm for trustworthy scientific AI, in which feasibility, physical structure, and mathematical optimality are treated as first-class computational principles rather than soft regularization objectives. Enabling constrained learning that infers directly within the boundaries of admissible and physically meaningful solutions would enable "inverse modeling" through nonlinear parameter estimation and provide a compelling alternative to expensive first-principles by supporting amortized optimization, in which solution structure is learned across problem instances, enabling near-instantaneous inference.

ML-AUGMENTED GENERAL FRAMEWORK FOR PROCESS MODELING, SIMULATION AND OPTIMIZATION

Our proposed constraint-embedded ML framework (see **Figure 1**) involves augmenting a standard neural network with a projection layer acting as a corrector. Given an input vector x, the network first computes an unconstrained output vector $\hat{y} = NN(\Theta, x)$, which may not satisfy the physical constraints. The projection layer then adjusts this output to produce a corrected prediction $\tilde{y}$ that lies exactly on the constraint manifold. Here, $\hat{y}$ denotes the unconstrained network prediction and the projection operator maps it onto the feasible set defined by the equality and inequality constraints, thereby ensuring feasibility by construction while enabling downstream guarantees on optimality. Unconstrained outputs are calculated using a standard neural net, while the constrained outputs correspond to the Karush–Kuhn–Tucker (KKT) conditions of a distance minimization problem. The projection layer enforces raw neural network output to lie on the constraint manifold.

Depending on the choice of the loss function and the choice of the projection layer, the framework shown in **Figure 1** reduces to perform different specific tasks related to process modeling, simulation, optimization and parameter estimation. This is summarized in **Table 1**, and discussed below in details.

HYBRID MODELING

The framework presented in Figure 1 reduces to a foundational projection-based neural architecture that embeds and enforces hard nonlinear equality and inequality constraints within artifical neural networks (ANNs), ensuring predictions satisfy feasibility by construction. This establishes scalable algorithms for rigorous constraint satisfaction in hybrid data-driven/mechanistic process models for fast simulation purposes. In modeling for process simulation, the relationship between input variables x and output variables y is governed by mass and energy conservation laws, inter-variable dependencies, and physical limitations on the operation. These can often be expressed as linear/nonlinear equality or inequality constraints. Consider a dataset $\{(x_i, \bar{y}_i)\}_{i=1}^{N}$ representing the input–output behavior of a process system. The objective is to train a neural network that approximates the underlying functional mapping between the inputs x and the outputs y. We consider that the mass and energy conservation equations and thermodynamic and kinetic models are also available in the form of algebraic constraints that relate the inputs x and the outputs y. Formally, the goal is to learn $y = NN(\Theta, x)$ such that the predicted output y always belongs to the feasible set *S*, i.e., $y \in S$, where *S* is defined as $S = \{y | h(x, y) = 0, g(x, y) \leq 0\}$. Specifically, $h(x, y) = 0$ represents a set of algebraic equality constraints, and $g(x, y) \leq 0$ represents another

Table 1. Setup for different tasks performed by the general framework presented in Figure 1.

NN Loss function	Projection Layer	Specific Tool	Functionality
$\frac{1}{2N}\|\tilde{\boldsymbol{y}}_i - \bar{\boldsymbol{y}}_i\|^2$	"Strictly" enforces nonlinear algebraic equality and inequality constraints:	**KKT-HardNet**	Hard-constrained ML, Hybrid Data-driven/Mechanistic Process Modeling, Process Simulation
$\frac{1}{2N}\|\tilde{\boldsymbol{y}}_i - \bar{\boldsymbol{y}}_i\|^2 + \frac{\omega}{M}\|\tilde{\boldsymbol{d}} - AD(\boldsymbol{\partial})\|^2$	"Strictly" enforces pointwise gradient-based constraint satisfaction:	**DAE-HardNet**	Discovery of Governing Equation, Inverse Modeling, Parameter Estimation
min $f(\boldsymbol{x})$, or min $L(x)$, or min KKT residuals	Ensures constraint satisfaction through enforcing KKT system of the distance minimization	**KKT-OptNet**	Process Optimization, Nonlinear Multi-parametric Optimization

set of algebraic inequality constraints. The functions h and g encode domain-specific knowledge, and may be nonlinear in both x and y. Here, the tunable parameters (weights and biases) of an ANN are defined as Θ.

The strict enforcement of constraints is achieved by solving a square system of nonlinear equations derived from the KKT conditions associated with the following constrained minimization problem (Eq. 1):

$$\begin{aligned} \tilde{y}^* = \arg\min_{y} \ & \tfrac{1}{2}\|y - \hat{y}\|^2 \\ \text{s.t. } & h_k(x,y) = 0, \quad \forall k \in \mathcal{N}_E, \\ & g_k(x,y) \le 0, \quad \forall k \in \mathcal{N}_I. \end{aligned} \tag{1}$$

This searches for a *minimum distance* solution from $\hat{y}$ that lies on the feasible region defined by the equality and inequality constraints (as illustrated in **Figure 1**).

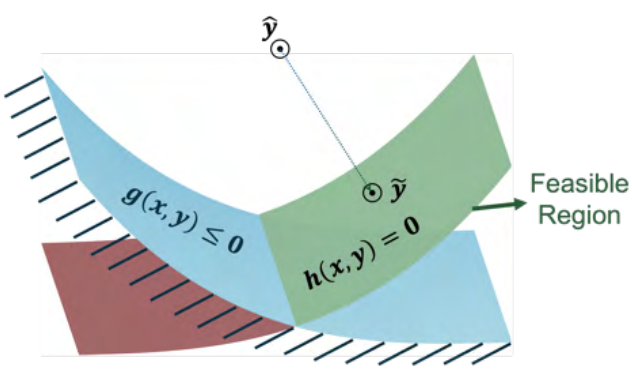

Figure 1. Projection of ANN output on the constraints manifold.

The optimal solution $\tilde{y}^*$ corresponds to a minimum distance projection of $\hat{y}$ onto the feasible region. We consider solving the system of nonlinear equations derived from the KKT conditions of (1) that yield a feasible projection $\tilde{y}$. For the special case of convexity, the KKT system is both necessary and sufficient to obtain the minimum distance solution $\tilde{y}^*$. This feasible projection can be directly embedded in the forward pass of a neural network architecture through a Newton-type iterative scheme, ensuring that only physically consistent predictions are used for loss computation and backpropagation. As a result, the loss function for the ANN is $\frac{1}{2N}\|\tilde{\boldsymbol{y}}_i - \bar{\boldsymbol{y}}_i\|^2$, in contrast to the conventional loss of $\frac{1}{2N}\|\hat{y}_i - \bar{\boldsymbol{y}}_i\|^2$. The overall implementation is refered to as KKT-HardNet [3].

The use of the projected outputs $\tilde{y}$ alters the learning directions during training, which is informed by the feasible manifold. This ultimately guides the model toward parameters that fit the data while respecting the known constraints. This approach is similar to that of Chen et el. [4] applied to linear equalities. However, we extend it to accommodate both nonlinear equality and inequality constraints. KKT-Hardnet is architecture-agnostic and can be integrated with any neural network backbone, including convolutional neural networks (CNNs) or recurrent neural networks (RNNs).

Illustrative Example of Hybrid Modeling using KKT-HardNet:

We consider the example of an extractive distillation column [3]. The goal is to learn the input to output mapping while strictly satisfying the equations described by the system of nonlinear equation given in Eq (2).

$$
\begin{cases}
x_1 + x_2 - (y_1 + y_2) = 0, \\
x_1 a_0 - (y_1 y_3 + y_2 y_6) = 0, \\
x_1 a_1 - (y_1 y_4 + y_2 y_7) = 0, \\
y_3 + y_4 + y_5 - 1 = 0, \\
y_6 + y_7 + y_8 - 1 = 0, \\
x_3 y_1 - y_9 = 0.
\end{cases} \tag{2}
$$

Along with the input vector x, we consider two fixed value a_0 and a_1 that represents the feed composition to the column. For the forward modelling example we consider $a_0 = 0.6976$ and $a_1 = 0.3024$ while simulating the system.

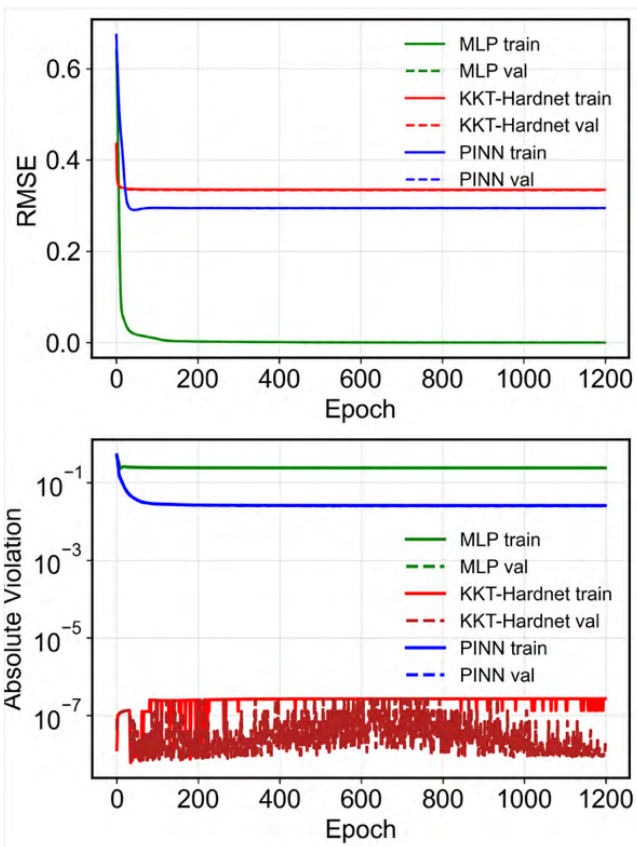

Figure 2. Learning curve for extractive distillation column. Top: mse loss between data and prediction, Bottom: absolute violation of the constraints, bottom: parameter evaluations during the training.

We compare the performance of KKT-HardNet with MLP and PINN, using identical backbone configuration and optimization settings (learning rate 10^{-4}, 1200 epochs, Adam optimizer). PINN model considers the constraint violation as physics loss while training the system with a regularization fator λ (10.0). Figure 2 shows the learning curves of three models. We observe no violation above 10^{-6} for KKT-HardNet iinferring that during training and validation the model outputs satisfies the physical constraints imposed on the model. The small oscillations below 10^{-6} arise because, in some forward passes, the Gauss–Newton iteration terminates with the residual norm dropping well below the stopping criterion of 10^{-6}. In those cases, the final residual may lie anywhere between the tolerance (10^{-6}) and the numerical precision limit. Thus, the fluctuations reflect solver termination under finite precision.

LEARNING PROCESS OPTIMIZATION

The framwork shown in Figure 1 reduces to an ML-based solver for multi-parametric constrained nonlinear programs with fast inference, stability and rigorous convergence with both feasibility and optimality guarantees. In this case, we consider convex multiparametric nonllinear programs (mpNLPs) of the form: $\min f(x)$ s.t. $x \in \mathcal{C}$ where $\mathcal{C} = \{ x \mid h(x) = a, g(x) \leq b, x \in \mathbb{R}^n \}$. We are interested in convex mpNLP, when $f: \mathbb{R}^n \to \mathbb{R}$ is convex C² and continuously differentiable, $g: \mathbb{R}^n \to \mathbb{R}^p$ is convex C¹, and $h: \mathbb{R}^n \to \mathbb{R}^q$ is affine. Convexity ensures the optimal solution during the projection from the first order neccesary conditions only.

Similar to previous, the framework consists of a neural network backbone with a projection to the feasible space trained on a new loss function. The loss function we used is the actual objective function of the optimization problem. However, the model can be trained on different loss functions such as the lagrangian or the KKT residuals. Unlike KKT-HardNet, learning optimization problem is trained in an unsupervised framework.

Illustrative Example of Nonlinear Optimization using KKT-OptNet:

For illustrative purpose, we take the classic transportation problem with parametric uncertainty. The problem is defined as follows:

$$
\min_y \ \frac{1}{2} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix}^T \begin{bmatrix} 306 & 0 & 0 & 0 \\ 0 & 324 & 0 & 0 \\ 0 & 0 & 324 & 0 \\ 0 & 0 & 0 & 252 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix} + \begin{bmatrix} 25 \\ 25 \\ 25 \\ 25 \end{bmatrix}^T \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix}
$$

$$
\text{s.t.} \ \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ -1 & 0 & -1 & 0 \\ 0 & -1 & 0 & -1 \\ -1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix} \leq \begin{bmatrix} 350 \\ 600 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ -1 & 0 \\ 0 & -1 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}. \tag{3}
$$

Here, we train the model for uncertainty bounds $0 \leq x_1, x_2 \leq 200$. For modeling efficiency we have normalized the optimization variables and the problem constraints. Additionally we solve the QP problem for 1000 parameter combinations sampled form the parameter bounds to

general optimal variable set data to validate our results.

Figure 3. shows the KKT-OptNet model performance for minimizing the objective loss function as neural network loss function parameterized by the weights and biases. Top figure shows that the model finds the optimal objective loss function with very few iterations of training. Additionally the bottom plot shows the optimality gap between the optimal solution obtained from the solver. We observe mean optimality gap of 0.0051% which shows that through neural network training the model is learning a optimal mapping from the parameter space to the variable space close to the optimal solution.

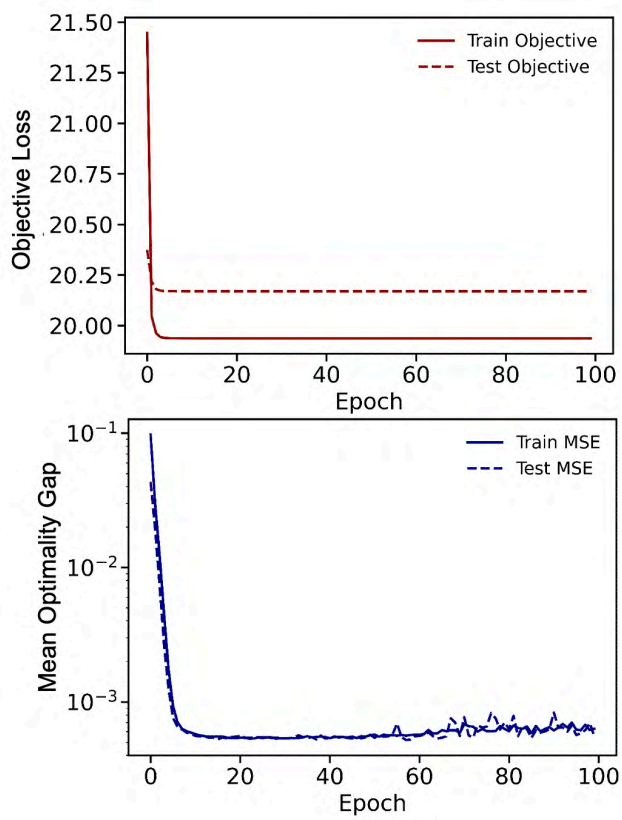

Figure 3. Learning curve for optimization learning of transportation optimization problem. Top: Objective Loss propagation through epochs during neural network training, Bottom: Mean optimality gap between predicted variables and solution obtained from solvers during training.

PARAMETER ESTIMATION

The projection based constraint satisfaction method can also be used as a tool for parameter estimation for nonlinear system. While we have discussed the scope of surrogate modeling with constraint satisfaction and learning optimization problems, some constraints might have some fixed value parameters which are unknown. Traditional parameter estimation technique involves solving a rigorous optimization problem by minimizing the prediction error. However, KKT-HardNet or DAE-HardNet is capable of doing the same with minimal changes. The unknown parameters can be introduced in the framework as learning parameters of the neural network. These parameters are optimized as the training progresses. Along with model fitting, the desired parameters are obtained that satisfied the constraints.

Parameter estimation can be challenging for system of differential equations. We illustrate how a Taylor approximation based pointwise projection can be implied to model dynamic systems and estimate unknown parameters through DAE-HardNet [6].

Illustrative Example of Parameter Estimation:

We consider a discrete-time linear time-invariant system defined by the state equation: $x_{\{k+1\}} = A_d x_k + B_d u_k$, where $x_k \in R^2$ is the state vector and $u_k \in \mathbb{R}$ is the control input.

The system matrices are parameterized as $A_d = [[1\ 1], [0\ a_1]]$ and $B_d = [0, a_0]^T$. The parameter a_1 represents the persistence or decay factor of the second state, governing how strongly $x_{k,2}$ propagates forward in time. The parameter a_0 represents the input gain and determines how strongly the control input influences the second state. Thus, a_1 controls intrinsic system dynamics while a_0 controls actuator effectiveness. For this case study we have simulated the data considering $a_0 = -1$ and $a_1 = -1$.

Given a prediction horizon N, the finite-horizon MPC problem is formulated as: minimize over $\{x_k, u_k\}$ the sum from $k = 0$ to $N - 1$ of $(x_k^T Q\, x_k + u_k^T R\, u_k)$, subject to $x_{\{k+1\}} = A_d x_k + B_d u_k$, with state constraints $-\bar{x} \leq x_k \leq \bar{x}$ and input constraints $-\bar{u} \leq u_k \leq \bar{u}$.

Stacking all predicted states and inputs into a decision vector $y = [x_1, \dots, x_N, u_0, \dots, u_{\{N-1\}}]^T$, the MPC problem can be written in quadratic program form as: minimize $\left(\frac{1}{2}\right) y^T Q y + c^T y$ subject to $Ay = b + Bx_0$ and $Cy \leq d$. The matrix A encodes stacked system dynamics, B propagates the initial condition, C represents box constraints, d defines bounds, and Q and c arise from quadratic stage costs.

In the inverse parameter estimation setting, the physical parameters $\theta = (a_0, a_1)$ are treated as unknown. Given observed trajectories $\{x_k^{data}, u_k^{data}\}$, we estimate parameters by minimizing the discrepancy between predicted and observed optimal trajectories: $\min_{a_0, a_1} \Sigma_i \left\| y\left(x_0^{(i)}; a_o, a_1\right) - y_{(i)}^{data} \right\|^2$, where $y^*(\cdot)$ denotes the solution of the MPC quadratic program.

The parameters influence feasibility through the dynamics constraints, the optimal trajectory shape, the control effort, and state propagation. Gradients $\partial y^* / \partial a_0$ and $\partial y^* / \partial a_1$ can be computed by embedding the QP

solver as a differentiable layer, enabling gradient-based learning. Figure 4 shows the parameter evaluations during the training. Table 2 includes the training results.

Table 2. Validation performance on the parameter estimation problem.

	PINN	KKT-HN	KKT-HN + PINN
MSE	1.49×10^{-1}	8.92×10^{-6}	2.43×10^{-5}
Equality violation	4.7×10^{-3}	3.1×10^{-3}	3.76×10^{-5}
Inequality violation	1.3×10^{-3}	1.5×10^{-4}	1.8×10^{-6}

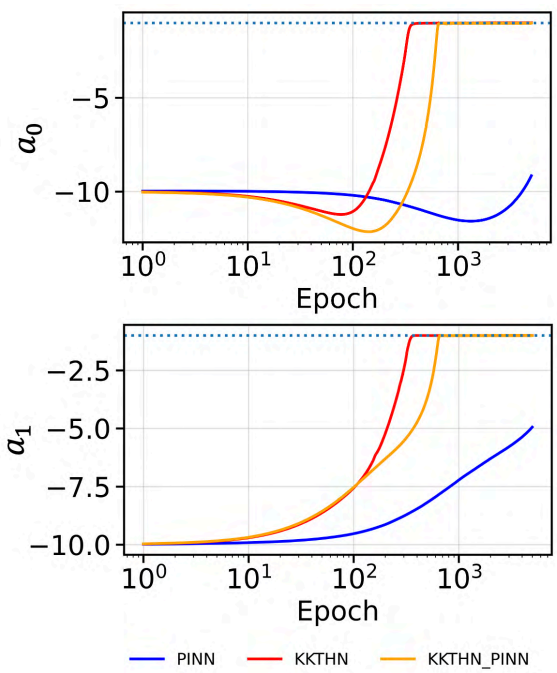

Figure 4. Learning curve for extractive distillation column. Top: mse loss between data and prediction, Bottom: absolute violation of the constraints, bottom: parameter evaluations during the training.

Figure 4 shows the parameter values through the training epochs for PINN, KKT-HardNet and KKT-HardNet augmented with PINN soft loss. We train all three models for 5000 epochs with standard hyperparameter settings. From the figure we observe that PINN even with 5000 epochs of training does not converge to the original value of the parameters. Both KKT-HardNet and KKT-HardNet augmented with soft loss converges to the original parameter values after training of only 500 epochs. Additionally Table 2 shows the loss values and constraint violation of all three models. PINN even though it achieves a constraint violation on the similar order as KKT-HardNet, the regression performance of PINN is five orders of magnitude higher than KKT-HardNet. Additionally, KKT-HardNet augmented with PINN loss shows more robust modeling performance with lower constraint violation compared to KKT-HardNet for this case study. Note that we have initialized the parameter values from -10. We observe that the KKT-HardNet converges to the actual parameter values for wide range of negative value.

CONCLUSIONS

We have presented a projection-based neural network architechture which guarantees strict feasibility of nonlinear equality and inequality constraints. The projection is based on the first order optimality conditions of a distance minimization problem. The differentiable projection layer is iteratively solved through Gauss-Newton method. This projection layer works as a corrector for the unconstrained prediction and keeps the prediction in the feasible domain. The framework can be used as a tool for robust surrogate modeling, parameter estimation and process optimization. The projection layer can act in different forms, for example, a taylor series approximation can be implemented to satisfy dynamic systems. This allows us to estimate parameters on the dynamic domain also. We provide several illustrative examples to show the applicability of our framework in chemical processes and optimization. Additionaly we observe that combining soft PINN loss with the projection layer helps achieving lower constraint violation and higher prediction performance. Moreover, it helps estimating unknown parameters. A future direction for extending this framework is to improve the computational complexity of the model with minimal iterative scheme.

AUTHOR IDENTIFIERS

Author ORCIDs:
Golder R: 0000-0003-2402-0455
Nath Roy B: 0000-0001-5387-5036
Hasan MMF: 0000-0001-9338-6069

REFERENCES

1. Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics 378:686-707 (2019). https://doi.org/10.1016/j.jcp.2018.10.045
2. Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. Nat Rev Phys 3:422-440 (2021).

https://doi.org/10.1038/s42254-021-00314-5
3. Iftakher A, Golder R, Roy BN, Faruque Hasan MM. Physics-informed neural networks with hard nonlinear equality and inequality constraints. Computers & Chemical Engineering 204:109418 (2026). https://doi.org/10.1016/j.compchemeng.2025.109418
4. Chen H, Flores GEC, Li C. Physics-informed neural networks with hard linear equality constraints. Computers & Chemical Engineering 189:108764 (2024). https://doi.org/10.1016/j.compchemeng.2024.108764
5. Stellato B, Banjac G, Goulart P, Bemporad A, Boyd S. OSQP: an operator splitting solver for quadratic programs. Math. Prog. Comp. 12:637-672 (2020). https://doi.org/10.1007/s12532-020-00179-2
6. Golder, Rahul, Bimol Nath Roy, and M. M. Hasan. "DAE-HardNet: A Physics Constrained Neural Network Enforcing Differential-Algebraic Hard Constraints." *arXiv preprint* arXiv:2512.05881 (2025).

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Section 7

Process Design, Scheduling and Optimisation

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Effect of the feed composition on the performance of a double-dividing wall distillation column

Carlos E. Guzmán-Martínez[a], Araceli G. Romero-Izquierdo[b], Claudia Gutiérrez-Antonio[a], Salvador Hernández[c], Massimiliano Errico[de], and Fernando I. Gómez-Castro[c]*

[a] Universidad Autónoma de Querétaro, Campus Juriquilla, Facultad de Ingeniería, Juriquilla, Querétaro, México
[b] Universidad Autónoma de Querétaro, Campus Amealco, Facultad de Ingeniería, Amealco de Bonfil, Querétaro, México
[c] Universidad de Guanajuato, Campus Guanajuato, División de Ciencias Naturales y Exactas, Departamento de Ingeniería Química, Guanajuato, Guanajuato, México
[d] University of Southern Denmark, Faculty of Engineering, Department of Green Technology, Odense, Denmark
[e] Università di Cagliari, Department of Mechanical, Chemical and Materials Engineering, Cagliari, Italy
* Corresponding Author: fgomez@ugto.mx

ABSTRACT

In this work, the synthesis, design and optimization of a quaternary double dividing wall distillation column (QDDWC) is presented. The effect of the feed composition over the performance of this intensified configuration is studied. The synthesis and design of the QDDWC takes place using as basis a conventional direct sequence for the separation of a n-butane/n-pentane/n-hexane/n-heptane mixture. The column is tested for three molar feed compositions: 40/10/10/40, 25/25/25/25, and 10/40/40/10. The configurations are optimized through a multiobjective genetic algorithm to simultaneously minimize the total heat duty and the number of stages. According to the results, the proposed structure allows savings in heat duty up to 59% but requiring up to 28% more stages than the conventional sequences.

Keywords: dividing wall columns, quaternary mixture, metaheuristic optimization

INTRODUCTION

Worldwide, over 90% of fluid mixtures in industrial processes are separated into different fractions or into pure components through distillation. This operation takes advantage of the differences on boiling points among the components of a given mixture using heat as separating agent. In industrial practice, this unit operation commonly takes place in distillation columns, which are devices with internals as trays or packing, whose task is to promote the contact among the phases to allow the mass and heat transfer. External heating is provided through a reboiler. Unfortunately, distillation has low thermodynamic efficiency, due to the input of high-quality energy in the reboiler which is later removed at low temperature in the condenser [1]. This implies high energy requirements, utilities costs and environmental impact.

To overcome the drawbacks of distillation in terms of its energy requirements, diverse alternatives have been developed. Among these alternatives, the thermally coupled configurations, the internally heat-integrated systems, and the cyclic distillation can be mentioned. One of the most well-known intensified distillation schemes, that has been implemented in industrial practice, is the dividing wall column (DWC). For ternary mixtures, the use of DWCs may lead to energy savings up to 40% in comparison to conventional distillation sequences [2]. In the case of quaternary mixtures, the potential in energy savings and controllability of dividing-wall distillation systems has been reported [3, 4]. In a recent study, a double dividing wall configuration has been proposed, exhibiting energy savings around 14% in comparison with a conventional direct sequence [5]. However, such study analyses the performance of the system only for a single feed composition. Indeed, it is well known that the potential for reducing energy requirements of such intensified schemes depends on the feed composition. Thus, it is mandatory to evaluate the performance of the proposed configuration for a variety of feed compositions, which is the objective of the present research.

CASE STUDY

The purification of a mixture composed by n-butane (A), n-pentane (B), n-hexane (C) and n-heptane (D) is studied in this work. A recovery of 99% is defined as target for all the components, for all the studied configurations. The separation of the mixture will take place in a conventional direct sequence (CDS) and a double-dividing wall configuration (QDDWC). A feed flowrate of 100 kmol/h is assumed, entering the distillation schemes as saturated liquid. Three feed compositions are analyzed: 40/10/10/40, 25/25/25/25, and 10/40/40/10. The first one represents a case where the middle components have a high presence in the mixture, the second case is an equimolar mixture; finally, the third case implies that the middle components have the highest contribution to the total composition.

METHODOLOGY

The synthesis of the QDDWC takes as basis a conventional direct sequence (CDS), as shown in Figure 1. The CDS has three columns, where A and B are obtained as distillate in the first and second column, respectively, while the third column produced pure C and D. Each column is divided in sections, as presented in Figure 1.a, where sections I, III and V correspond to the rectifying zone of each column, while sections II, IV and VI represent the stripping zones. By rearrangement of the sections, the intensified configuration shown in Figure 1.b is obtained, where a single reboiler is required [6].

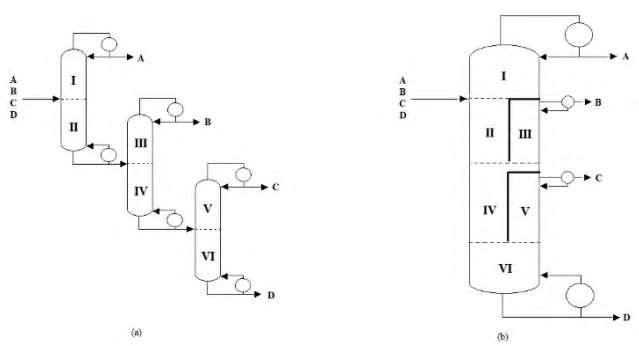

Figure 1: Stage rearrangement for the synthesis of the intensified distillation system (a) conventional direct sequence, (b) quaternary double dividing wall column

Conventional and intensified configurations are both simulated in the commercial process simulator Aspen Plus. To represent the phase equilibrium, the Chao-Seader model is employed, since it is recommended for hydrocarbon systems [7]. To simulate the CDS, the operating pressure is first estimated through bubble point calculations, assuming that the temperature on each condenser is approximately 49 °C; this ensure that cooling water can be used as utility on that equipment. Once the pressure has been determined, the number of stages and the feed stage are estimated through the Fenske-Winn-Underwood approach given by the DSTWU module, assuming that $R = 1.33R_{min}$ and aiming to obtain recoveries of 99% for each component. A pressure drop of 0.689 bar is assumed for all the conventional columns. All these calculations are performed for each of the studied feed compositions. Once the preliminary configuration of each column is obtained, the conventional sequences are rigorously simulated through the RadFrac module, which solves the MESH equations to obtain detailed information about the profiles on each column. The purities for each product are estimated from the mass balances, using the desired recovery. These purities are fixed by using design specifications, employing the reflux ratio as manipulated variable for the products obtained as distillate, and the reboiler duty for those products obtained as residues.

Once the design for the conventional sequences has been obtained, the structure of the intensified columns is generated by rearranging the sections of the CDSs. Since there is not a module to represent dividing wall systems in the simulator, the equivalent structure shown in Figure 2 is used. The split fraction for both splitters is adjusted until convergence is obtained, and the purities of the components are fixed by using design specifications.

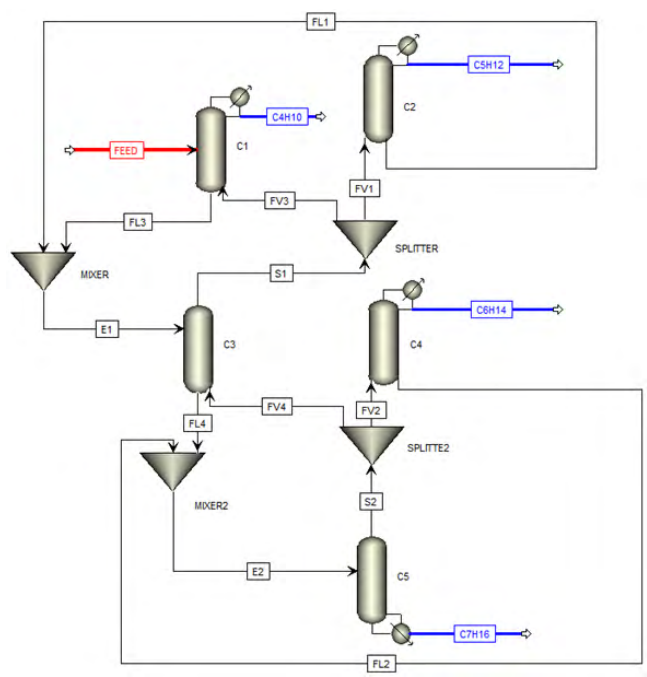

Figure 2. Representation of the QDDWC equivalence in the simulation environment.

Using the preliminary designs previously described as initial guesses, the optimization of the intensified schemes is performed through the NSGA-II metaheuristic. First, a dataset of decision-variable values (manipulated variables) is generated in a main MATLAB routine, using NSGA-II. Subsequently, these decision-variable

vectors are transferred to a MATLAB subroutine, which serves as a link between the software and where the data are received and forwarded to Aspen Plus. For each individual, Aspen Plus solves the rigorous mass and energy balances. When convergence is not achieved, a penalty value is assigned to the corresponding solution; otherwise, the simulation results are returned to the MATLAB subroutine interface. At this stage, process constraints are evaluated, and the objective-function values are computed and sent back to the MATLAB main routine. Once the objective functions are received, the genetic algorithm proceeds to the next generation. After completing the prescribed number of iterations, the Pareto front is obtained and plotted.

Figure 3 presents the implemented optimization workflow, including the information exchange between the decision variables and the corresponding response variables.

The NSGA-II parameters are set as follows: number of generations = 50, population size = 300, crossover fraction = 0.8, elite count = 15, and mutation rate = 0.01.

The optimization model simultaneously minimizes (i) the total number of stages in the column (Eq. 1) and (ii) the reboiler duty (Eq. 2). Process feasibility is ensured by imposing recovery constraints of nearly 99% for each component (Eqs. 3 and 4). As expressed in Eq. 1, the total number of stages is defined as a function of the stage allocation in Columns 1, 3, and 5 of the intensified configuration. Furthermore, Eq. 2 indicates that the reboiler duty is influenced by the total stage number, distillate flowrates, boil-up ratio, internal molar flowrates, and feed stage locations.

Minimize:

$$OF_1: TS = S_{C1} + S_{C3} + S_{C5} \tag{1}$$

$$OF_2: Q = f(D_{C1}, D_{C2}, D_{C4}, B_{C5}, FV_1, FV_2, S_{C1}, S_{C3}, S_{C5}, FS) \tag{2}$$

s.t.

$$0.988 \leq R_i \leq 0.9903 \tag{3}$$

$$R_i = \frac{\dot{n}_{i(o)}}{\dot{n}_{i(f)}} \tag{4}$$

Here, OF_1 = objective function 1, OF_2 = objective function 2, TS = total number of stages in the dividing wall column, SC_1= stages on column 1, SC_3= stages on column 3, SC_5= stages on column 5, Q = reboiler duty, DC_1= distillate rate in column 1, DC_2= distillate rate in column 2, DC_4= distillate rate in column 4, B_{C5}= boil-up ratio in column 5, FV_1= vapor molar flow from column 3 to column 2, FV_2= vapor molar flow from column 5 to column 4, FS= feed stage, R= molar recovery, $\dot{n}$= molar flow. Subscripts: i= molecular specie (butane, pentane, hexane, heptane), (o)= output stream for specie i, and (f)= feed stream for specie i.

Once the optimal points for each concentration level are selected, comparisons are carried out among the optimized solutions, against their corresponding base-case designs (i.e., the same configurations prior to the optimization algorithm), and with the conventional separation schemes.

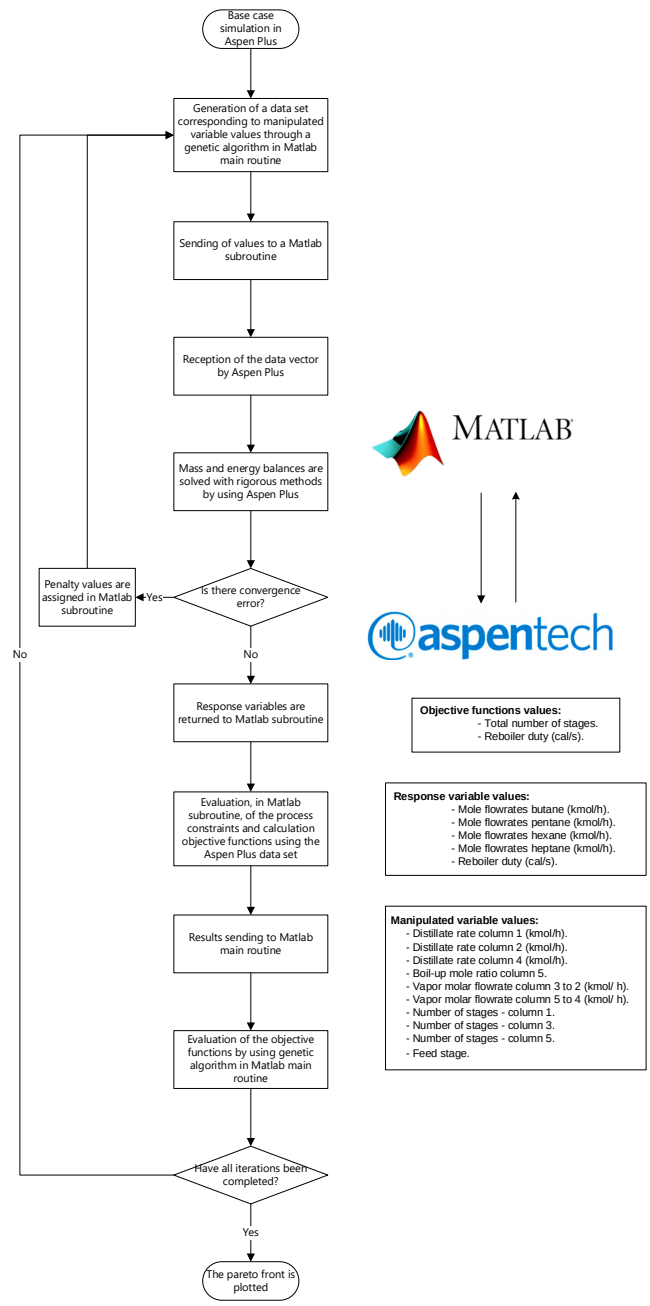

Figure 3: Optimization algorithm

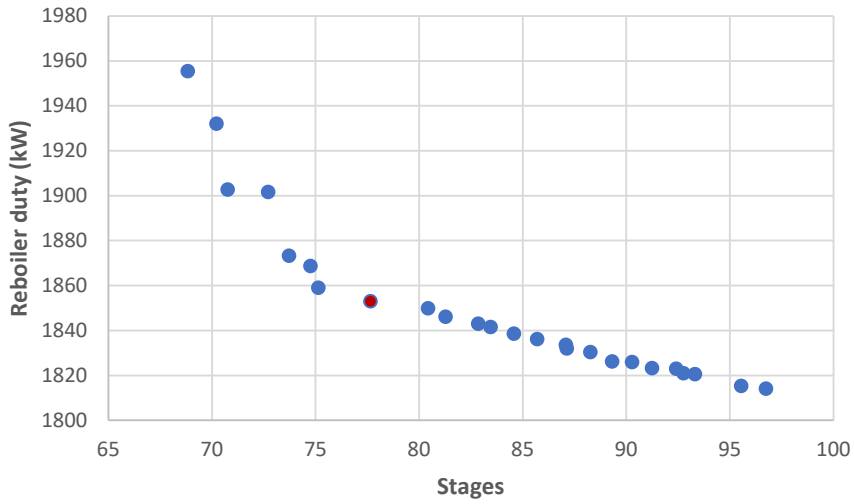

Figure 4: Pareto front of the intensified scheme for the 40/10/10/40 composition.

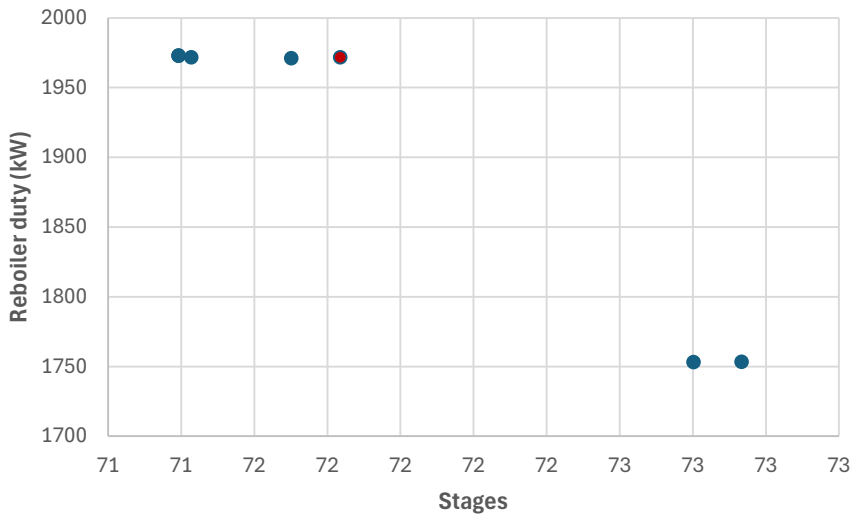

Figure 5: Pareto front of the intensified scheme for the 25/25/25/25 composition.

RESULTS

Figures 4 to 6 show Pareto fronts for each case under analysis. The red point in each graph is considered the solution with the best compromise.

Table 1 reports the values of the manipulated variables for the QDDWC base-case designs, where DDW40, DDW25, and DDW10 correspond to feed compositions of 40/10/10/40, 25/25/25/25, and 10/40/40/10, respectively. Similarly, Table 2 presents the corresponding values for the optimized designs. In this case, DDW40-o, DDW25-o, and DDW10-o refer to the optimized configurations associated with feed compositions of 40/10/10/40, 25/25/25/25, and 10/40/40/10, respectively.

Table 3 reports the response-variable values for each intensified and conventional scheme at each concentration levels. The conventional ones are CDS40, CDS25, and CDS10 for 40/10/10/40, 25/25/25/25, and 10/40/40/10 concentration, respectively.

When comparing the DDWC base-case designs with the optimized schemes, a noticeable reduction in both reboiler duty and total number of stages is observed. For instance, for the 40/10/10/40 composition, the optimized configuration achieves a 5.05% reduction in reboiler duty and a 30.03% decrease in the number of stages. Similarly, for the 10/40/40/10 case, reductions of 7.72% and 31.43% are obtained in reboiler duty and total stages, respectively. In contrast, the 25/25/25/25 composition exhibits only a 2.86% decrease in reboiler duty, while a slight increase of 2.62% in the total number of stages is

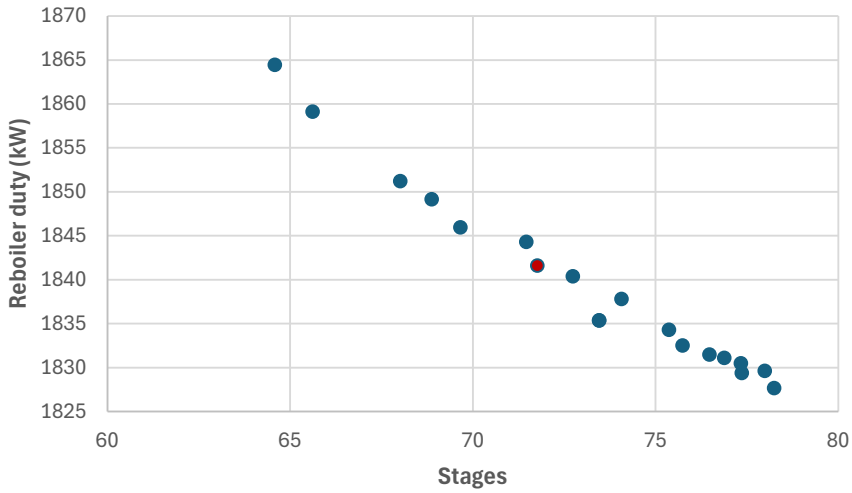

Figure 6: Pareto front of the intensified scheme for the 10/40/40/10 composition.

observed. This increase is mainly attributed to the shape of the Pareto set rather than to a failure of the optimization algorithm.

The comparison between the conventional distillation sequence (CDS) and the optimal intensified schemes (DDW-o) reveals a clear trade-off: a significant reduction in reboiler duty is achieved at the expense of an increase in the total number of stages. In terms of reboiler duty, the DDW-o configurations for 40/10/10/40, 25/25/25/25, and 10/40/40/10 exhibit reductions of 58.27%, 35.81%, and 59.69%, respectively, relative to the CDS. Conversely, stage requirements are increased by 27.32%, 28.28%, and 20.00% for the same compositions, respectively.

The energy savings observed in intensified schemes compared to conventional sequences can be explained by the inherent structure of each separation strategy. In the CDS, the separation is carried out in three columns: butane is recovered as the distillate of the first column, pentane as the distillate of the second column, and hexane as the distillate of the third column, while heptane is withdrawn from the bottom of the third column. Therefore, intermediate mixtures (e.g., pentane, hexane, and heptane) leaving the bottom of the first column are withdrawn at elevated temperatures (typically above 115 °C). When entering the second column, these streams are fed and cooled at significantly lower temperatures (approximately 70–80 °C). A similar situation occurs between the second and third columns: the bottom stream from the second column exits again at temperatures above 120 °C and is subsequently cooled to around 90 °C when introduced into the final column. In the third column, the heptane-rich bottom stream may reach temperatures up to 150 °C. Therefore, the conventional sequence is characterized by repetitive heating and cooling steps across columns, which results in avoidable energy losses.

In contrast, intensified configurations achieve the separation within a single integrated unit, which enables internal heat integration and significantly reduces the need for successive cooling and reheating of intermediate streams. This structural advantage explains the superior thermal performance of intensified schemes relative to the conventional sequence.

The increase in the total number of stages observed in the intensified schemes compared to the CDS can be mainly explained by differences in relative volatility (Figs. 5(a) to 5(f)). These results indicate that the hexane–heptane separation becomes more challenging in the intensified configurations than in the conventional sequence. In general, as relative volatility approaches unity, the separation difficulty increases substantially because the volatility difference between the light and heavy key components becomes less pronounced, thereby requiring a higher number of equilibrium stages to achieve the specified recoveries.

Figure 7 illustrates the effect of feed-composition changes on the performance of the evaluated sequences. Both plots exhibit a parabolic trend and show an inverse behavior with respect to each other. For instance, in the intensified schemes, DDW25 presents the highest reboiler duty, whereas in the conventional sequences the corresponding CDS25 case exhibits the lowest energy requirement. Conversely, DDW10 shows the lowest reboiler duty among intensified configurations, while CDS10 displays the highest energy consumption within the conventional set.

To interpret these trends, the relative-volatility profiles shown in Figs. 5(a)–5(f) and Figs. 8–10, where the heat capacity of each mixture are shown, must be considered.

Table 1: Values of manipulated variables for base-case designs

	DDW40	DDW25	DDW10
Distillate rate C1 (kmol/h)	39.713	24.757	9.94
Distillate rate C2 (kmol/h)	10.504	25.321	39.333
Distillate rate C4 (kmol/h)	10.251	25.205	40.064
Boil-up ratio	11.065	11.927	27.5
Split fraction FV1	0.3113	0.4439	0.5615
Split fraction FV2	0.4306	0.523	0.295
Stages C1	35	32	39
Stages C3 and C4	37	34	35
Stages C5	39	23	47
Feed stage	8	13	8
Stage C2	27	25	31

Table 2: Values of manipulated variables for optimized schemes

	DDW40-o	DDW25-o	DDW10-o
Distillate rate C1 (kmol/h)	39.802	25.273	9.931
Distillate rate C2 (kmol/h)	10.408	25.259	39.473
Distillate rate C4 (kmol/h)	10.235	24.702	40.074
Boil-up ratio	6.7843	11.688	25.429
Split fraction FV1	0.3918	0.516	0.549
Split fraction FV2	0.3613	0.445	0.32
Stages C1	22	27	28
Stages C3 and C4	29	36	35
Stages C5	27	9	9
Feed stage	8	20	18
Stage C2	14	7	10

Table 3: Values of response variables for all developed schemes.

Scheme	Condenser duty (kW)	Reboiler duty (kW)	Total stages
CDS40	-4464.07	4440.27	61
DDW40	-1520.14	1951.5	111
DDW40-o	-1605.45	1852.9	78
CDS25	-3177.94	3071.41	56
DDW25	-1886.3	2029.8	70
DDW25-o	-1787.64	1971.66	72
CDS10	-4783.86	4568.15	60
DDW10	-1939.2	1995.59	105
DDW10-o	-1735.92	1841.58	72

The heat capacity profiles of the butane/pentane/hexane/heptane mixtures (Figs. 8–10) show a clear compositional dependence. In particular, the mixture heat capacity increases as the fractions of pentane and hexane rise, whereas it decreases when butane and heptane become less dominant.

The behavior of the heat capacity profiles is directly associated with the phase transition of the mixture from liquid to vapor. This relationship can be demonstrated by comparing Figures 8 and 11, where Figure 11 presents the molar vapor fraction. It is important to note that these figures correspond to the 40/10/10/40 composition. The temperatures at which the liquid phase begins to evaporate coincide with the temperatures at which the heat capacity starts to increase. Although only one molar vapor fraction profile is shown, the same behavior is observed for all the different compositions analyzed in this study.

From a thermodynamic standpoint, this trend is consstent with the fact that the heat capacity of hydrocarbon mixtures is strongly influenced by the contribution of each component to the overall sensible enthalpy. Pentane and hexane, which constitute the intermediate fractions, typically exhibit higher molar heat capacities than butane due to their larger molecular size. Therefore, enriching the mixture in pentane and hexane results in a higher mixture heat capacity, meaning that more thermal energy must be supplied to increase the system temperature by 1 °C.

Additionally, the observed behavior may reflect non-ideal mixture contributions to heat capacity, which arise from compositional changes in intermolecular interactions within the multicomponent hydrocarbon system. As a result, variations in pentane and hexane content can lead to measurable changes in the mixture heat-capacity response.

Among the intensified schemes, DDW10 exhibits the lowest energy requirements, mainly due to the internal vapor-flow conditions (boil-up) established inside the column. The vapor rate is generated in the reboiler and is

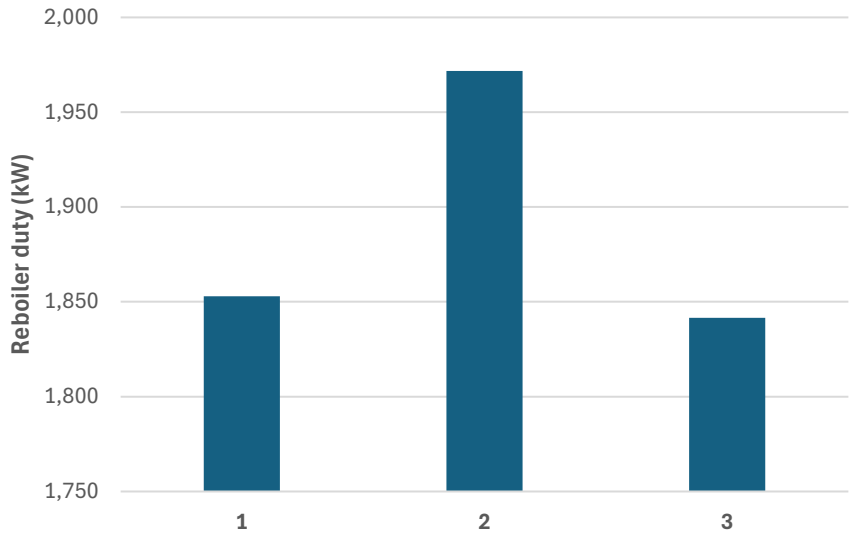

Figure 7 (a) Reboiler duty on intensified schemes. 1) DDW40, 2) DDW25, 3)DDW10.

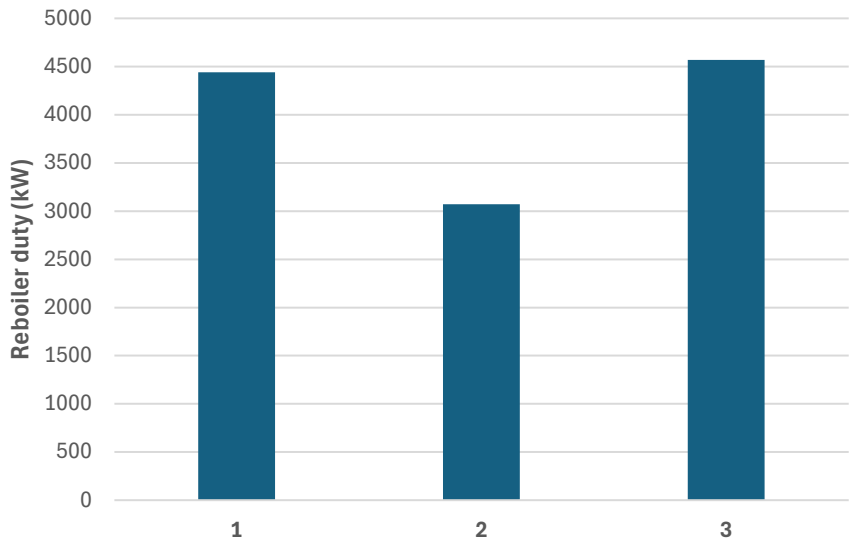

Figure 7 (b): Reboiler duty on conventional schemes. 1) CDS40, 2) CDS25, 3) CDS10.

predominantly composed of heptane. Although the heptane fraction in the vapor stream is lower in DDW10 than in DDW25 and DDW40, the generated vapor flow still provides sufficient thermal driving force to vaporize the required amount of hexane; thereby, maintaining the target vapor traffic throughout the DDW10 configuration. As a result, the overall reboiler duty is lower than that observed for DDW25 and DDW40.

In contrast, DDW40 demands higher energy than DDW10 because the hexane content in the system is lower. Consequently, a larger fraction of the boil-up must be carried by heptane to sustain the vapor flow within the column. From a thermodynamic perspective, maintaining heptane in the vapor phase is more energy-intensive than vaporizing hexane, since heptane is less volatile (higher normal boiling point), and typically requires a higher enthalpy input per unit of generated vapor. Therefore, schemes relying more heavily on heptane as the main vapor carrier led to higher reboiler duties.

In addition, the fact that DDW25 exhibits a higher energy requirement than DDW10 and DDW40 (Fig. 6) can be rationalized by analyzing the relative volatility behavior (Figs. 5(a), 5(c), and 5(e)). Although the heat-capacity trends suggest the following order: 10/40/40/10 > 25/25/25/25 > 40/10/10/40, the separation difficulty in these systems is more strongly governed by phase-equilibrium limitations than by sensible-heat effects. In particular, the relative-volatility profiles indicate a more unfavorable separation for the hexane-related split in the

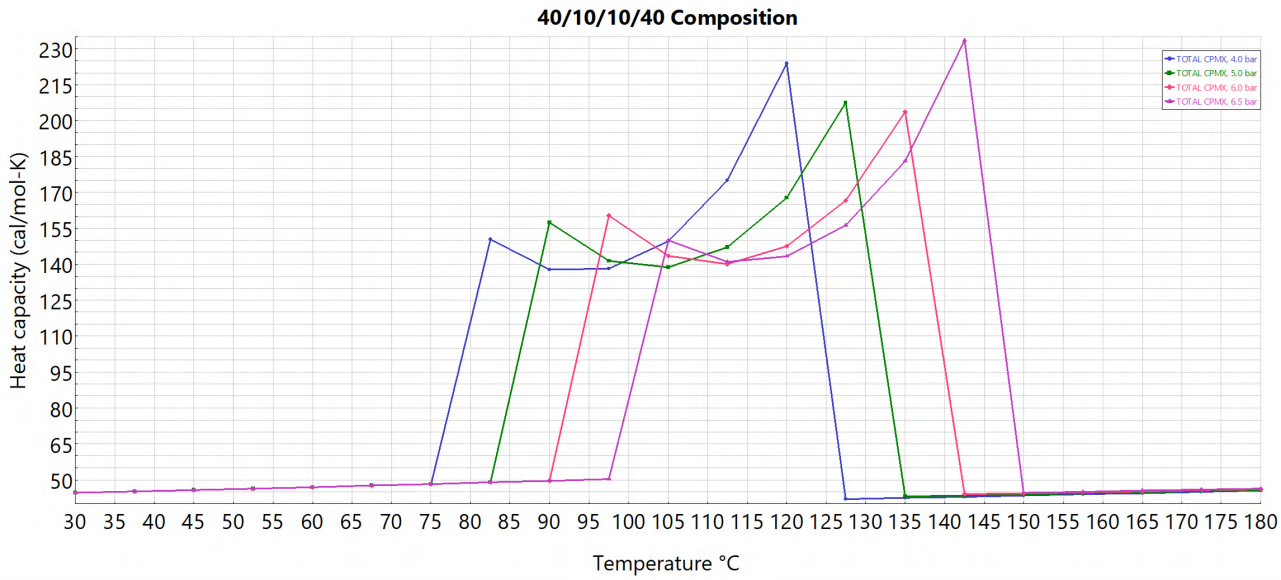

Figure 8: Heat capacity of the 40/10/10/40 mixture over the pressure range of 4–6.5 bar.

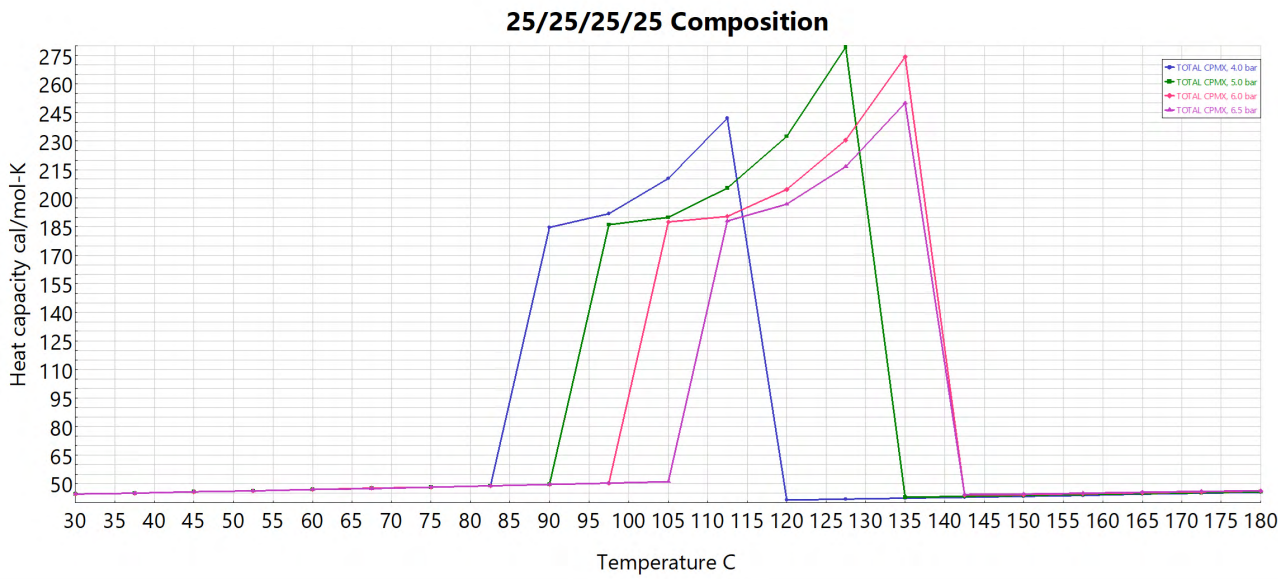

Figure 9: Heat capacity of the 25/25/25/25 mixture over the pressure range of 4–6.5 bar.

equimolar case (25/25/25/25), which increases the required vapor traffic and, consequently, the reboiler duty.

Therefore, the relative-volatility behavior supports the energy-demand trend observed in Fig. 6, namely: DDW25 > DDW40 > DDW10. This effect may be attributed to mixture non-ideality and composition-dependent intermolecular interactions, which are maximized when all hydrocarbons are present in comparable proportions; thereby, reducing the effective volatility differences between key components.

Finally, CDS10 exhibits the highest energy consumption among the conventional schemes, followed by CDS40 and CDS25. Similar to the intensified configurations, this trend is mainly governed by the relative-volatility behavior (Figs. 5(b), 5(d), and 5(f)), which determines the separation difficulty and the amount of vapor traffic required to achieve the specified recoveries. This behavior may also reflect composition-dependent intermolecular interactions that affect phase equilibrium and, consequently, the effective relative volatility between key components.

Since the conventional sequence performs the separation across different equipment units, sensible-heat effects become more relevant. In particular, for CDS10, the heptane-rich stream is cooled and reheated as it

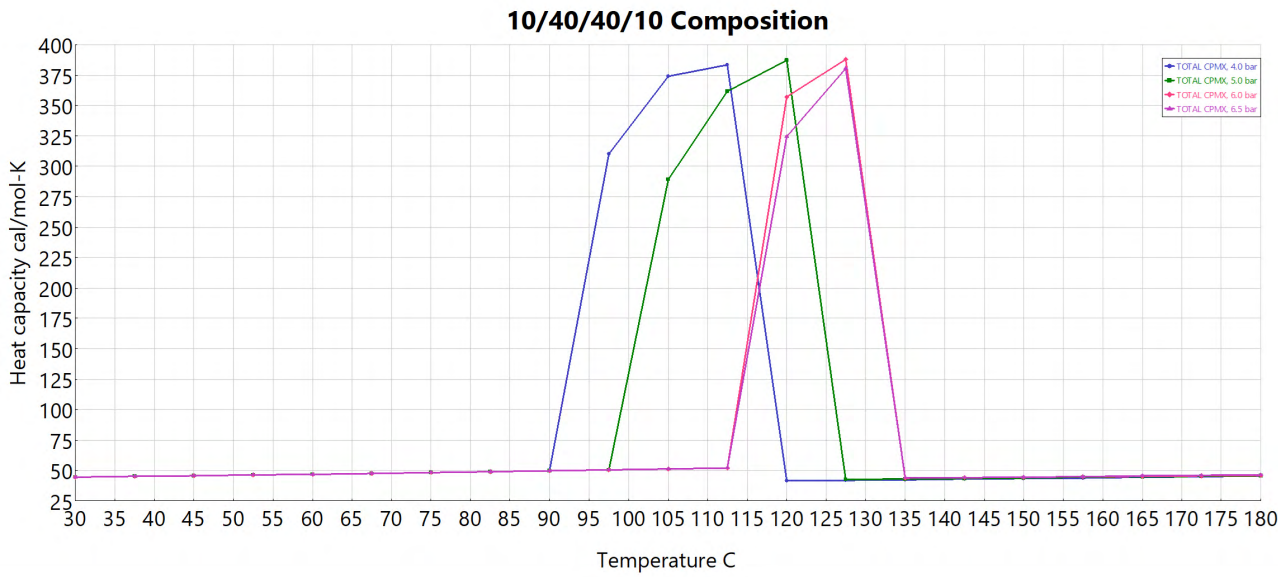

Figure 10: Heat capacity of the 10/40/40/10 mixture over the pressure range of 4–6.5 bar.

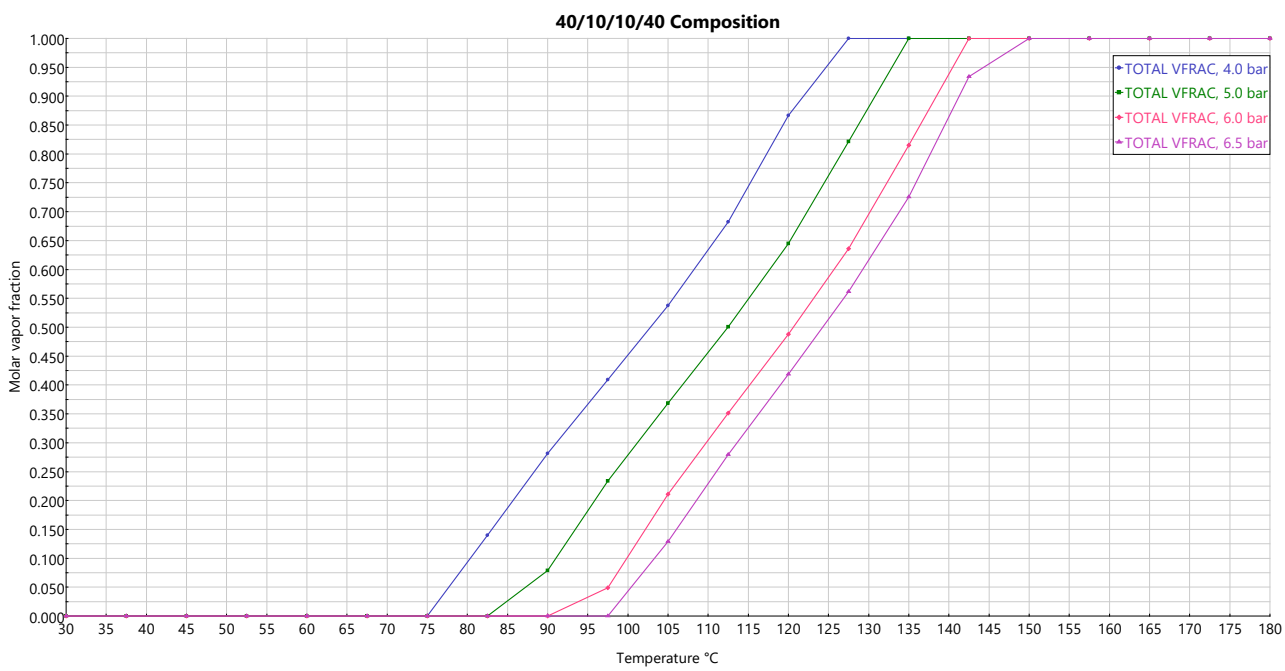

Figure 11: Molar vapor fraction of the 40/10/10/40 mixure over the pressure range of 4–6.5 bar.

passes through Columns 1 and 2 before entering Column 3, introducing additional utility consumption. Furthermore, because heptane is present in a lower proportion in CDS10 than in CDS40, a higher reboiler duty is required in the final column to supply sufficient boil-up, vaporize hexane, and maintain the internal vapor flow. These combined effects lead to a higher overall energy demand.

In contrast, CDS25 presents the lowest energy requirement among the conventional schemes. This result can be explained by the structure of the three-column sequence and the more balanced distribution of component loads, which reduces the severity of the limiting separation. Additionally, separating the mixture across different columns progressively changes the local composition, which tends to mitigate non-ideal mixture effects compared with intensified schemes where all hydrocarbons coexist within the same equipment.

CONCLUSIONS

In conclusion, the intensified schemes significantly

reduce the energy requirements required to separate the butane/pentane/hexane/heptane system. In addition, feed composition plays a critical role in the performance of both intensified and conventional configurations, strongly influencing separation difficulty and utility consumption. Therefore, a broader assessment considering additional feed compositions is recommended to provide a more comprehensive understanding of the behavior and robustness of the evaluated schemes.

Although intensified configurations tend to require a higher total number of stages, they offer an important structural advantage by reducing the number of separation units, particularly the number of columns and reboilers, compared to conventional distillation sequences. This feature highlights the potential of intensified schemes as energy-efficient alternatives for multicomponent hydrocarbon separation.

ACKNOWLEDGEMENTS

The authors acknowledge the support provided by Universidad de Guanajuato, Universidad Autónoma de Querétaro, University of Southern Denmark and Università di Cagliari. C.E. Guzmán-Martínez acknowledges to Secretaría de Ciencia, Humanidades, Tecnología e Innovación (SECIHTI, México) for the financial support granted to develop his post-doctoral research (grant no. 2397403)

AUTHOR IDENTIFIERS

Author ORCIDs:
Guzmán-Martínez CE: 0000-0002-3404-2650
Romero-Izquierdo AG: 0000-0003-1546-6033
Gutiérrez Antonio C: 0000-0002-7557-2471
Hernández S: 0000-0002-4598-3392
Errico M: 0000-0002-2172-2921
Gómez-Castro FI: 0000-0003-4906-063X

REFERENCES

1. Kiss AA, Smith R. Rethinking energy use in distillation processes for a more sustainable chemical industry. Energy 203:117788 (2020). https://doi.org/10.1016/j.energy.2020.117788
2. Czarnecki NJ, Giannetti L, Owens SA, Barnicki S, Eldridge RB. Energy and economic evaluation of wall placement for divided wall distillation columns. Ind. Eng. Chem. Res. 63:18513-18524 (2024). https://doi.org/10.1021/acs.iecr.4c01691
3. Sánchez-Ramírez E, Ramírez-Márquez C, Quiroz-Ramírez JJ, Angelina-Martínez AY, Vicente Cortazar V, Segovia-Hernández JG. Design of dividing wall columns involving sustainable indexes for a class of quaternary mixtures. Chemical Engineering and Processing - Process Intensification 148:107833 (2020). https://doi.org/10.1016/j.cep.2020.107833
4. Kong B, Zhang Q, Cui C, Sun J. Optimal design and effective control of kaibel column with liquid-only transfer streams for quaternary distillation. Separation and Purification Technology 250:117261 (2020). https://doi.org/10.1016/j.seppur.2020.117261
5. Flores-Flores M, Gómez-Castro FI, Gutiérrez-Antonio C, Romero-Izquierdo AG, Guzmán-Martínez CE, Hernández S, Errico M. Synthesis and optimization of dividing-wall distillation columns for the separation of a quaternary mixture. Separation and Purification Technology 362:131817 (2025). https://doi.org/10.1016/j.seppur.2025.131817
6. Errico M, Rong BG, Torres-Ortega CE, Segovia-Hernandez JG. The importance of the sequential synthesis methodology in the optimal distillation sequences design. Computers & Chemical Engineering 62:1-9 (2014). https://doi.org/10.1016/j.compchemeng.2013.11.010
7. Jamali MA, Bissenbay A, Nuraje N. Thermodynamic modeling and process simulation of kumkol crude oil refining. Eurasian Chem.-Technol. J. 25:183-192 (2023). https://doi.org/10.18321/ectj1521

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

System-Level CO_2 Allocation under Supply Constraints in Industrial Clusters

Razan Sawaly[a], Ahmad Abushaikha[a], and Tareq Al-ansari[ab]*

[a] College of Science and Engineering, Hamad Bin Khalifa University, Qatar Foundation, Doha, Qatar
[b] Qatar Environment and Energy Research Institute (QEERI), Hamad Bin Khalifa University, Qatar Foundation, Doha, Qatar
* Corresponding Author: talansari@hbku.edu.qa

ABSTRACT

Efficient deployment of carbon capture, utilisation, and storage (CCUS) within industrial clusters requires coordinated CO_2 allocation under economic, technical, and environmental constraints, particularly when CO_2 availability is limited. This paper presents a centralised optimisation framework for allocating captured CO_2 from nine industrial sources to six utilisation and storage sinks within an industrial park in Qatar. A multi-objective mixed-integer linear programming (MILP) model is developed to minimise total system cost while accounting for capture, purification, transport, and utilisation processes, and enforcing an environmental feasibility constraint to ensure net CO_2 reduction. The model is evaluated under four scenarios: a baseline case with sufficient CO_2 to satisfy all sink demands, and three scarcity scenarios in which 15%, 25%, and 35% of total source emissions are available. Results show that under scarcity, allocations prioritise large EOR sinks supplied by high-volume, low-purity sources, while utilisation pathways are progressively reintroduced as availability increases. The findings highlight the critical influence of CO_2 purity, sink requirements, and supply constraints on allocation outcomes and underscore the importance of centralised planning for robust CCUS deployment.

Keywords: CCUS, CO_2 allocation, Life-cycle emissions, Optimisation, CO_2 purity

INTRODUCTION

Carbon capture, utilisation, and storage (CCUS) has been widely recognised as a critical decarbonisation pathway for energy-intensive industrial systems, particularly in sectors where direct electrification and fuel switching remain technically or economically constrained. Previous system-level assessments emphasised that the effectiveness of CCUS deployment depended not only on capture technologies, but also on the coordinated planning of CO_2 transport, allocation, and end-use across interconnected industrial networks. In this context, Leonzio et al. (2022) highlighted the importance of integrated supply-chain optimisation for achieving cost-effective emissions reduction, showing that sub-optimal source-sink matching and transport design could significantly undermine the economic viability of CCUS systems.

Several studies extended CCUS optimisation frameworks to explicitly include utilisation pathways alongside storage. Zhang et al. (2024) and Iskandar & Kurihara (2025) developed co-optimisation models capturing trade-offs between CO_2 utilisation, storage, and transport modalities, while Susantoro et al. (2023) focused on CCUS-EOR configurations aimed at meeting net-zero or near-zero emissions targets. These studies demonstrated the potential economic benefits of integrating utilisation sinks within CCUS networks, but typically evaluated system performance under assumed availability conditions and aggregate emissions targets. Zhang et al. (2017) investigated the design of integrated CCS infrastructure under carbon trading schemes and highlighted the importance of spatial proximity and network configuration in determining optimal capture and transport strategies.

Recent literature placed increasing emphasis on the technical characteristics of CO_2 streams and their implications for system performance. (Alnouri & Al-Mohannadi, 2023) examined the influence of treatment design parameters on carbon integration and showed that purification requirements significantly affect capture efficiency and downstream feasibility. (Sawaly et al., 2026)

quantified the role of CO_2 purity and treatment costs in shaping optimal allocation within energy and petroleum systems, while (Kumcu et al., 2026) further highlighted that capture and purification processes often represent dominant cost components in CCUS supply chains.

Optimisation-based modelling approaches have been extensively applied to CCUS system design. Mixed-integer linear programming (MILP) formulations have been used to minimise total system cost while determining optimal capture, transport, and storage configurations at regional and national scales. Caselli et al. (2022) formulated MILP frameworks to support strategic CCUS infrastructure planning, while Wu et al. (2023) extended optimisation approaches to address transport and storage coordination under capacity constraints. More recent contributions incorporated broader system interactions, including energy-industry coupling and emissions reduction targets, to assess the role of CCUS within low-carbon transition pathways (Al-Mohannadi & Al-Mohannadi, 2023; Sawaly et al., 2024)

Despite these advances, several gaps remained in literature. First, most optimisation studies implicitly assumed that sufficient captured CO_2 was available to satisfy utilisation and storage demand, providing limited insight into allocation behaviour under supply-constrained conditions. Second, life-cycle CO_2 emissions were commonly assessed ex post, rather than enforced directly within optimisation frameworks through environmental feasibility constraints. Third, the combined influence of source CO_2 purity, sink-specific requirements, and sink-level emissions on allocation prioritisation was rarely examined in an integrated manner, particularly within industrial clusters where multiple utilisation and storage pathways compete for limited CO_2 supply.

This paper addressed these gaps by developing a centralised optimisation framework for CO_2 allocation within an industrial park in Qatar. The proposed model integrated detailed capture and purification cost functions, explicit life-cycle emissions accounting, and an environmental feasibility constraint to ensure net CO_2 reduction. The framework was evaluated under a baseline scenario with unconstrained CO_2 availability and under multiple scarcity scenarios to analyse how allocation patterns, demand satisfaction, and emissions performance evolved as CO_2 availability decreased. In doing so, the study provided new insights into cost-emissions trade-offs and allocation dynamics relevant to the planning and deployment of resilient, environmentally robust CCUS systems in industrial clusters.

2. MODELS AND OPTIMISATION

A multi-objective optimisation framework based on mixed-integer linear programming (MILP) is developed to model CO_2 allocation within an industrial park in Qatar. The framework optimises the allocation of captured CO_2 from multiple industrial sources to a set of utilisation and storage sinks while balancing economic and environmental objectives. The system consists of nine industrial CO_2 sources and six application sinks. The optimisation determines source-sink selection and associated annual CO_2 flowrates, where binary decision variables represent source-sink assignments and continuous variables describe CO_2 flows. Each sink is supplied by a single source, while each source may supply multiple sinks subject to its CO_2 availability. Economic performance is evaluated through minimisation of total system cost, and environmental performance is incorporated through explicit accounting of life-cycle CO_2 emissions and the enforcement of an environmental feasibility constraint ensuring net CO_2 reduction.

To examine system behaviour under different operating conditions, the model is solved under four distinct scenarios. The first scenario represents a baseline case in which sufficient captured CO_2 is available to fully satisfy the demand of all utilisation and storage sinks. The remaining three scenarios consider constrained CO_2 availability, where only 15%, 25%, and 35% of total source emissions are available for allocation. These scenarios enable assessment of trade-offs between cost minimisation, CO_2 utilisation levels, and emissions performance under increasingly stringent system conditions.

2.1 Available Data and Assumptions

cCapture and purification costs are calculated for each candidate source-sink match using the cost function developed in Sawaly et al. (2026). The model uses source-specific CO_2 availability and CO_2 purity, together with sink-specific CO_2 demand and minimum purity requirements. The CO_2 emissions available at each source, together with their corresponding CO_2 purity, and the CO_2 demand and purity requirements at each sink are summarised in Table 1. The spatial distribution of all sources and sinks within the industrial park is presented in Figure 1.

Table 1. Technical and economic characteristics of CO_2 sinks.

Sinks	**CO_2 demand (tCO_2/year)**	**Purity required (%)**
QAFAC H2	3.86×10^6	78%
QAFAC NG	3.90×10^5	80%
Pearl GTL	1.28×10^6	47%
Oryx GTL	1.28×10^6	38%
Dukhan Field (EOR)	9.97×10^6	95%
AlShaheen Field (EOR)	4.30×10^6	95%
Sources	**Emissions (tCO_2/year)**	**Purity available (%)**

Ras Girtas Power Company (RGPC)	4.77×10^6	6%
Mesaieed Power Company Limited (MPCL)	1.90×10^6	5%
QAFCO	5.85×10^6	52%
QATALUM	4.65×10^6	3%
Qatarenergy-LNG	3.50×10^7	17%
Qatar Steel	1.72×10^6	25%
Q-CHEM	2.16×10^6	20%
QNCC	1.50×10^6	23%
QP REFINERY	1.20×10^6	11%

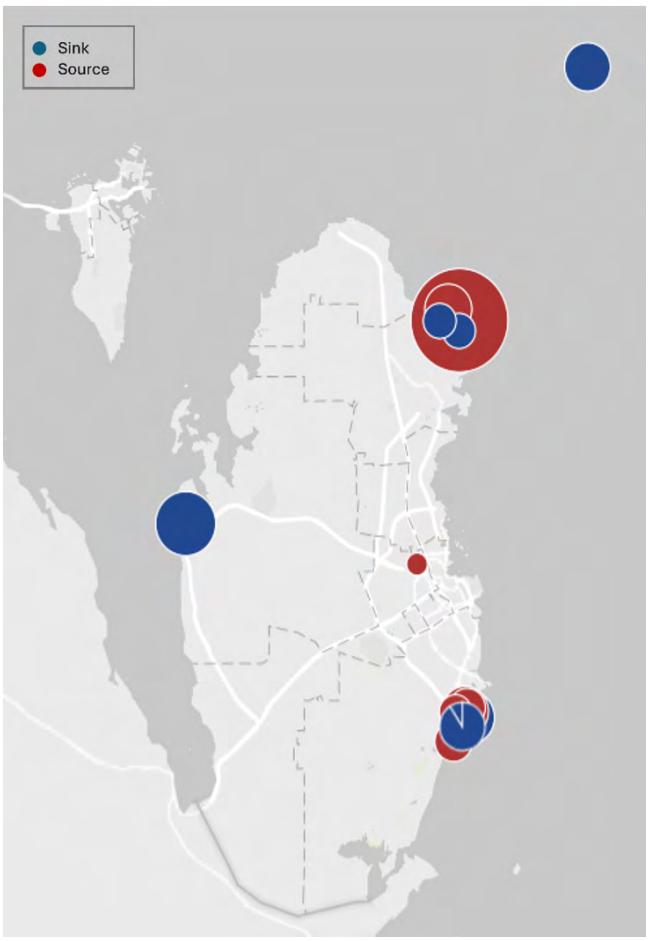

Figure 1. Spatial Distribution of CO_2 Sources and Sinks with Emissions and Demand.

The assumptions made in the model are as follows:

- CO_2 is transported via pipeline between sources and sinks.
- All source-sink connections are assumed technically feasible due to the availability of capture and purification technologies; feasibility is reflected through cost and emissions penalties rather than excluding pairs a priori.
- Differences in CO_2 purity are captured through source-sink-specific capture/purification cost and capture/purification emission factors, and through purity-adjusted transportation impacts.
- CO_2 is assumed to be transferred free of charge (no CO_2 commodity price); only operating costs are considered.
- Capital expenditures and financing effects are excluded; the optimisation accounts for annual operating costs only.
- Life-cycle CO_2 emissions include capture and purification, transport, and sink utilisation/storage emissions, consistent with the model's emissions accounting.

2.3 Model Formulation

The model includes the following decision variables:

q_{ij}: Mass flowrate of CO_2 transported from source i to sink j (tCO_2/year).

y_{ij}: Binary variable indicating whether sink j is supplied by source i.

To represent CO_2 scarcity, a non-negative slack variable u_j is introduced to capture unmet CO_2 demand at sink j.

Objective function

The optimisation minimises the total annual operating cost of the CO_2 allocation network:

$$\min Z^{\text{cost}} = \sum_{i \in S} \sum_{j \in J} [C_{ij}^{\text{cap}} + C_{ij}^{\text{tr}}] q_{ij} + \lambda \sum_{j} u_j \qquad (1)$$

where C_{ij}^{cap} and C_{ij}^{tr} denote capture-purification and transportation costs, respectively, and λ is a penalty parameter associated with unmet CO_2 demand.

The optimisation is subject to the following constraints:

(i) Source capacity constraint, limiting total CO_2 supplied by each source to its annual availability:

$$\sum_{j \in J} q_{ij} \leq E_i \; \forall i \in S \qquad (2)$$

where E_i is the annual CO_2 availability of source i.

(ii) Sink demand constraints, ensuring that CO_2 demand is satisfied or, in scarcity scenarios, captured as unmet demand:

$$\sum_{i \in S} q_{ij} + u_j = D_j \; \forall j \in J \qquad (3)$$

This constraint enforces full satisfaction of each sink's CO_2 demand D_j. It ensures that every sink receives exactly the amount of CO_2 required for its operation, preventing both under-supply and over-supply.

(iii) Single-source supply constraints, restricting each sink to be supplied by at most one source:

$$\sum_{i \in S} y_{ij} \leq 1 \; \forall j \in J \tag{4}$$

(iv) Flow-selection linking constraints, enforcing consistency between binary source-sink selection and continuous CO_2 flows:

$$q_{ij} \leq D_j \, y_{ij} \; \forall i \in S, \forall j \in J \tag{5}$$

(v) Capture and purification emissions, calculating emissions associated with CO_2 capture and purification:

$$Z^{\text{cap, emis}} = \sum_{i \in S} \sum_{j \in J} ER_{ij}^{\text{cap}} \, q_{ij} \tag{6}$$

where ER_{ij}^{cap} denote the capture and purification emission rate.

(vi) Transportation emissions, accounting for pipeline transport of CO_2:

$$Z^{\text{tr, emis}} = \sum_{i \in S} \sum_{j \in J} EF^{\text{tr}} \, d_{ij} \, \frac{1}{x_i} \, q_{ij} \tag{7}$$

where $EF^{\text{tr}} = 1.986 \times 10^{-6} \, \text{tCO}_2 \, / \, (\text{tCO}_2/\text{km})$, x_i is the CO_2 mass fraction (purity) of source i, and d_{ij} is the distance between each source-sink linkage.

(vii) Sink utilisation or storage emissions, capturing process-specific emissions at the sinks:

$$Z^{\text{sink, emis}} = \sum_{i \in S} \sum_{j \in J} ER_{j}^{\text{sink}} \, q_{ij} \tag{8}$$

where ER_{ij}^{sink} represents emissions generated during CO_2 utilisation or residual emissions during storage.

(viii) Total system emissions, aggregating all life-cycle emission components:

$$Z^{\text{emis}} = \sum_{i \in S} \sum_{j \in J} \left[ER_{ij}^{\text{cap}} + EF^{\text{tr}} \, d_{ij} \, \frac{1}{x_i} + ER_{j}^{\text{sink}} \right] q_{ij} \tag{9}$$

(ix) Environmental feasibility constraint (ε-constraint), ensuring net CO_2 reduction:

$$Z^{\text{emis}} \leq \sum_{i \in S} \sum_{j \in J} q_{ij} \tag{10}$$

This ε-constraint ensures that the total life-cycle emissions generated by the CCUS system do not exceed the amount of CO_2 captured and diverted from atmospheric release. It guarantees that all feasible solutions deliver a net reduction in CO_2 emissions relative to the baseline system without CCUS.

(x) Variable domain constraints, enforcing physical feasibility:

$$q_{ij} \geq 0 \; \forall i \in S, \forall j \in J \tag{11}$$

$$u_j \geq 0 \; \forall j \in J \tag{12}$$

$$y_{ij} \in \{0, 1\} \; \forall i \in S, \forall j \in J \tag{13}$$

These constraints ensure physical feasibility of the solution. CO_2 flows cannot be negative, and the binary variable enforces discrete selection of source-sink relationships.

RESULTS AND DISCUSSION

The optimisation results reveal that, across all scenarios, system-level performance is governed by the interaction between CO_2 availability, sink demand, and technical compatibility between sources and sinks. Figure 2 presents an aggregate comparison of total CO_2 demand, supplied CO_2, unmet demand, and overall demand satisfaction for the baseline and scarcity scenarios, highlighting the progressive decline in system-wide demand fulfilment as CO_2 availability becomes constrained. To understand the allocation mechanisms underlying these aggregate outcomes, Figure 3 illustrates the optimal source-sink CO_2 allocation pathways for the baseline and each scarcity scenario, showing how flow magnitudes and supplier-consumer relationships evolve under constrained availability.

In the baseline case, high-volume but low-purity sources such as QatarEnergy LNG (≈17% CO_2) are systematically allocated to EOR sinks, where purity requirements are less stringent and large flowrates dominate system performance. This explains why Dukhan Field and Alshaheen Field EOR jointly receive more than 14 $MtCO_2$ yr^{-1}, despite their relatively long transport distances. In contrast, chemical and GTL sinks, which impose higher purity constraints and incur higher penalties for purification, are preferentially supplied by QAFCO (≈52% CO_2) and Qatar Steel (≈25% CO_2), even when their emission capacities are smaller. This reflects the model's internal trade-off: allocating higher-purity streams reduces both purification costs and capture-related emissions, improving overall cost-emissions efficiency for utilisation pathways. As a result, sink satisfaction in the baseline scenario is not uniform but structured around technical compatibility rather than proximity alone.

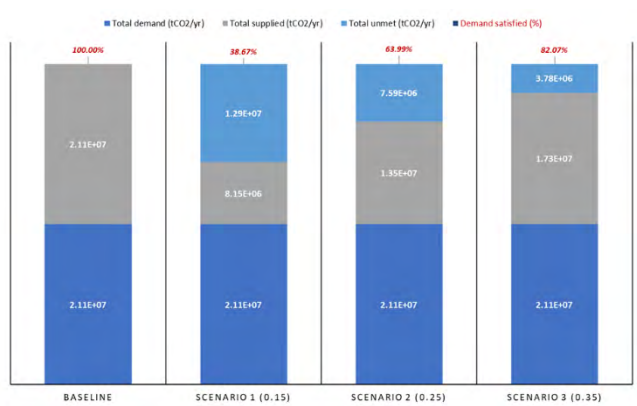

Figure 2: Scenario-Based Comparison of CO_2 Supply and Demand Satisfaction

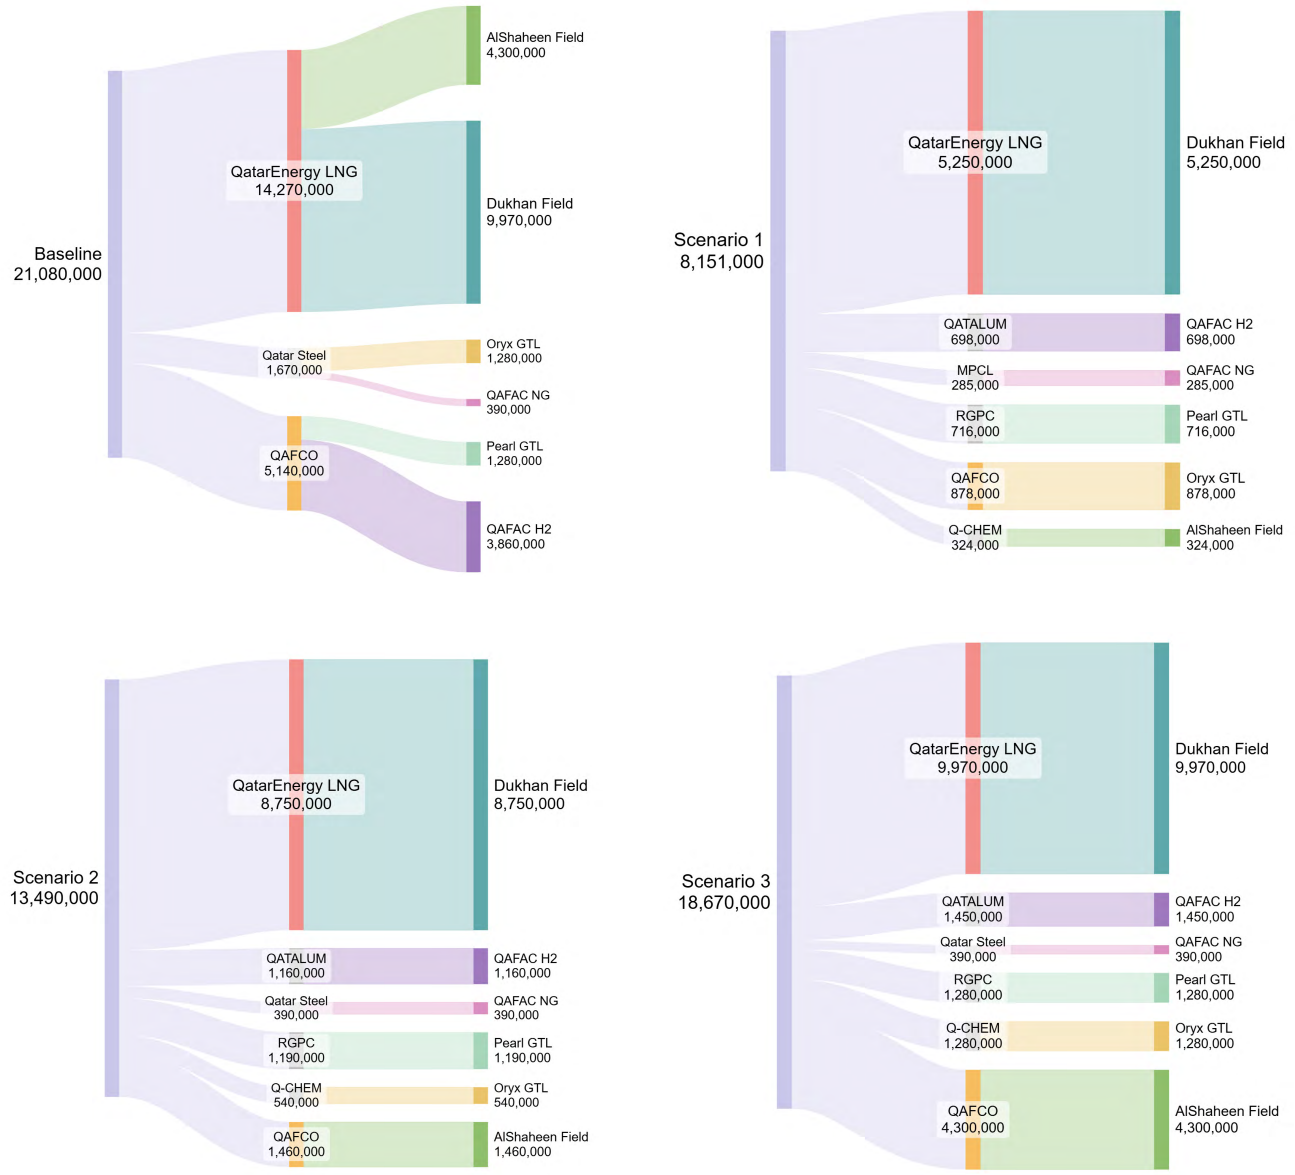

Figure 3: Source-Sink CO_2 allocation across Baseline and Scarcity Scenarios in tCO_2/year.

Under CO_2 scarcity, this compatibility-driven logic becomes more pronounced. When availability is limited to 15%, the optimisation increasingly favours sinks with lower marginal emissions per tonne utilised and greater tolerance to purity variability, particularly EOR operations. Chemical sinks such as QAFAC H_2, despite their strategic importance, are deprioritised because allocating scarce, high-purity CO_2 to these sinks displaces larger volumes that could otherwise satisfy EOR demand at lower system cost and emissions intensity. As availability increases to 25% and 35%, the reintroduction of chemical and GTL sinks reflects the point at which the marginal benefit of serving additional utilisation pathways outweighs the combined penalties of purification, transport, and sink-level emissions. Notably, sources with moderate purity and favourable locations, such as QAFCO, QATALUM, and RGPC, emerge as critical "bridging" suppliers in these scenarios, supporting diversification without sharply increasing system emissions. This behaviour highlights a key structural insight: CO_2 purity acts as an implicit scarcity multiplier, shaping allocation decisions even more strongly than absolute emission capacity. This progression, observed in Figure 3, explains the non-linear improvement in demand satisfaction reported in Figure 2. Consequently, the results demonstrate that CCUS system performance is governed not only by how much

CO_2 is available, but by how well source characteristics align with sink-specific technical and environmental constraints.

CONCLUSION

This study develops a centralised optimisation framework to examine CO_2 allocation within an industrial cluster under varying availability conditions, explicitly integrating economic costs, technical compatibility, and life-cycle emissions. By formulating the problem as a multi-objective mixed-integer linear program, the framework captures the trade-offs between cost minimisation, utilisation coverage, and environmental feasibility that arise when multiple sources and sinks compete for limited CO_2 supply. The results demonstrate that allocation outcomes are governed not only by emission volumes, but by the alignment of source CO_2 purity, sink requirements, transport distances, and sink-specific emissions profiles.

Scenario analysis reveals that CO_2 scarcity fundamentally reshapes allocation behaviour, with large EOR sinks consistently prioritised under constrained supply, while chemical and GTL utilisation pathways become viable only as availability increases. Importantly, the findings show that modest increases in captured CO_2 can yield disproportionate improvements in demand satisfaction and system balance without compromising net emissions performance. These insights highlight the limitations of assuming unconstrained CO_2 supply in CCUS planning and underscore the value of centralised allocation mechanisms, such as CO_2 hubs or banks, for managing competition among sinks.

The proposed framework provides a flexible basis for supporting CCUS infrastructure planning under uncertainty and can be extended to incorporate dynamic policy instruments, investment decisions, or decentralised market behaviour. As such, it offers a practical tool for designing resilient, cost-effective, and environmentally robust CCUS systems in industrial clusters.

ACKNOWLEDGEMENTS

The authors acknowledge the support of Qatar National Research Fund (a member of Qatar Foundation) provided by GSRA grants no. GSRA9-L-1-0531-22069.

REFERENCES

1. Al-Mohannadi SA, Al-Mohannadi DM. Qatar in the energy transition: low carbon economy challenges and opportunities. Gulf Studies :109-126 (2022). https://doi.org/10.1007/978-981-19-7398-7_7
2. Alnouri SY, Al-Mohannadi DM. The investigation of treatment design parameters on carbon integration networks. Clean Techn Environ Policy 25:2545-2559 (2023). https://doi.org/10.1007/s10098-023-02585-1
3. Caselli, G., Columbu, G., Iori, M., & Magni, C. A. (2022). A Mixed Integer Linear Programming model for CO2 emissions minimization in a waste transfer Facility Location Problem. *Advances in Database Technology - EDBT*, *2022-June*, 51-56. https://doi.org/10.48786/inoc.2022.10
4. Iskandar DEUP, Kurihara M. Co-optimization of carbon capture, utilization, and storage (CCUS) project using iterative latin hypercube sampling (ILHS). SCOG 48:253-262 (2025). https://doi.org/10.29017/scog.v48i2.1818
5. Kumcu S, Özyörük B, Bezzo F, d'Amore F. Optimization of carbon capture, utilization, and storage supply chains for the hard-to-abate industry in türkiye. International Journal of Greenhouse Gas Control 150:104568 (2026). https://doi.org/10.1016/j.ijggc.2026.104568
6. Leonzio, G., Bogle, D., Foscolo, P. U., & Zondervan, E. (2022). Optimization of CCUS supply chains in the UK: a strategic role for emissions reduction. *Chemical Engineering Research and Design*. https://www.sciencedirect.com/science/article/abs/pii/S0263876220300034
7. Sawaly R, Alherbawi M, Abd AS, AlNouss A, Abushaikha AS, Al-Ansari T. Strategic CO2 management in the energy and petroleum sector for the production of low-carbon LNG: a qatar case study. Journal of CO2 Utilization 83:102798 (2024). https://doi.org/10.1016/j.jcou.2024.102798
8. Sawaly R, AlNouss A, Abushaikha AS, Al-Ansari T. Techno-economic optimisation of carbon dioxide purification and allocation for industrial sinks: a two-stage simulation and optimisation framework. Energy Conversion and Management 347:120559 (2026). https://doi.org/10.1016/j.encoman.2025.120559
9. Tri Muji Susantoro , Sugihardjo , Wikantika K, Sunarjanto D, Pasarai U, Widarsono B, Rahmadi A, Romli M, Wahyudi P, Kepies S. CCUS-EOR optimization to achieve zero emission program targets in northwest java basin. Evergreen 10:1809-1818 (2023). https://doi.org/10.5109/7151730
10. Wu Y, Zhang H, Wang S, Zhen L. Mathematical optimization of carbon storage and transport problem for carbon capture, use, and storage chain. Mathematics 11:2765 (2023). https://doi.org/10.3390/math11122765
11. Zhang , Alhorr Y, Elsarrag E, Marafia AH, Lettieri P, Papageorgiou LG. Fair design of CCS infrastructure for power plants in qatar under carbon trading scheme. International Journal of Greenhouse Gas Control 56:43-54 (2017).

https://doi.org/10.1016/j.ijggc.2016.11.014
12. Zhang, F., Catalanotti, E., Martynov, S., Porter, R. T. J., & Mahgerefteh, H. (2024). *A Mixed-Integer Linear Programming Model for Multi-Modal CO 2 Transport*. https://ssrn.com/abstract=5059008

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article

Peer Reviewed Conference Proceeding

ESCAPE 36 - European Symposium on Computer Aided Process Engineering

Sheffield, UK. 21-24 June 2026

Virtual Plant–Model Pair as a Step Towards Real-Time Optimization of a Simulated Moving Bed System

Guilherme C. Amaral[a], Alexandre F. P. Ferreira[a], Ana M. Ribeiro[a], Idelfonso B. R. Nogueira[b], and Diogo Rodrigues[a]*

[a] LSRE-LCM, ALiCE, Faculty of Engineering, University of Porto, Rua Dr. Roberto Frias, 4200-465 Porto, Portugal
[b] Chemical Engineering Department, Norwegian University of Science and Technology, Sem Sælandsvei 4, Kjemiblokk 5, Trondheim, 793101, Norway
* Corresponding Author: dfrodrigues@fe.up.pt

ABSTRACT

Simulated Moving Bed (SMB) chromatography is widely used for a variety of separations, yet, when applicable, these systems are typically operated using offline optimization strategies. Over time, process degradation and unforeseen disturbances may cause SMB units to deviate from the calculated optimal conditions, reducing overall performance. Real-Time Optimization (RTO) offers a promising solution by continuously monitoring and adjusting operating conditions to maintain optimal performance, despite such perturbations. However, experimental implementation of RTO in industrial SMB processes is costly and requires significant interdisciplinary coordination.To address this challenge, a virtual framework is proposed for the preliminary development of a model-based RTO system. The methodology employs a virtual plant–model pair, in which a representative plant model generates *in silico* experimental data, while a structurally distinct predictive model reproduces these results. Structural mismatch was intentionally introduced to mimic real-world differences between a plant and its mathematical model, and measurement noise was added to enhance realism. Within this methodology, the *in silico* experiments were successfully generated and the parameters of the predictive model were then estimated using a Particle Swarm Optimization algorithm that sought to minimize the residuals between the *in silico* experimental data and the predictive model outputs. The parameters were successfully estimated, allowing the predictive model to closely reproduce the behavior of a structurally distinct plant model without introducing additional complications, which is expected to be analogous to a real-scenario RTO system. Hence, this work establishes a critical step towards the foundation of a virtual RTO framework.

Keywords: Optimization, Adsorption, Particle Swarm Optimization, Simulation, MATLAB

INTRODUCTION

Enantiomers have a critical role within the pharmaceutical industry, with current estimates indicating that more than half of all marketed drugs are administered as single enantiomer formulations [1]. It is also important to note that the use of racemic mixtures has been repeatedly associated with adverse, often severe, physiological effects [2]. Consequently, the isolation and production of enantiomerically pure compounds have become essential to ensure both the safety and therapeutic efficacy of chiral pharmaceuticals. Among the strategies available to obtain pure enantiomers, the separation of racemic mixtures has emerged as a practical and extensively investigated approach [3]. Of the various techniques employed for this purpose, chromatographic methods, in particular Simulated Moving Bed (SMB) chromatography, are of significant interest in this study due to their operational flexibility and cost-effectiveness.

Simulated Moving Bed

SMB chromatography relies on the interaction between two continuous countercurrent phases: a

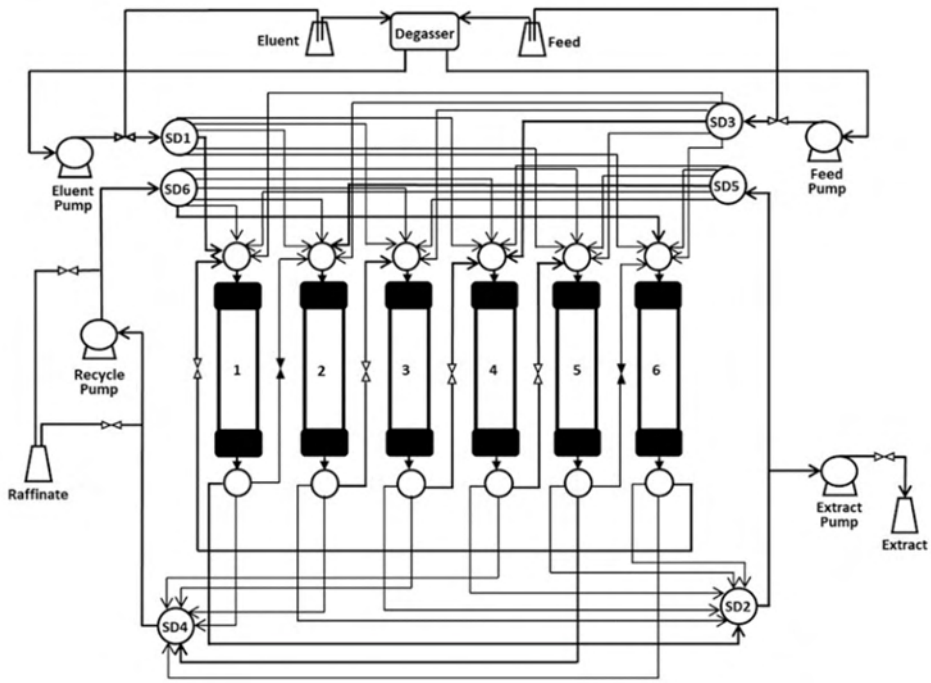

Figure 1. Diagram of the FlexSMB-LSRE®. Adapted from [4]

stationary solid phase, packed within a series of chromatographic columns and a mobile phase, typically a liquid or gas. Figure 1 displays the SMB unit studied in this work, the FlexSMB-LSRE® [4]. The interaction between phases enables the separation of adsorbates within the mobile phase. This separation is governed by the distinct affinities of individual components for the solid phase, which dictates their respective retention times. Therefore, in a binary system, preferential adsorption of one component onto the stationary phase subsequently increases its retention time, leading to the formation of three characteristic regions: two zones enriched in the individual components and an intermediate zone in which both species coexist, typically corresponding to the location of feed introduction. The extract stream, enriched in the more retained component, is withdrawn from one region, while the raffinate stream, containing the less retained component, is collected from the other. An eluent is introduced via a second inlet stream to sustain continuous separation [5, 6].

As can be seen in Figure 1, in an SMB unit, multiple fixed bed columns are connected in series and operated to emulate countercurrent contact between the solid and fluid phases. Although the solid phase remains physically stationary within the fixed beds, its motion is simulated by periodically shifting the positions of the inlet and outlet streams in the direction of the fluid flow at a fixed interval known as the switching time. This cyclic redistribution of streams creates an effective countercurrent movement of the solid relative to the mobile phase, enabling continuous separation [7, 8].

In SMB systems, the primary objectives are typically to maximize productivity and/or minimize eluent consumption while satisfying purity and/or recovery constraints. Consequently, optimization plays a vital role in enhancing process performance [9, 10]. However, conventional optimization approaches are typically conducted in an offline setting, which limits their ability to accommodate unanticipated disturbances, plant-model mismatch, or progressive system degradation. As a result, SMB units may operate under suboptimal conditions for prolonged periods, leading to reduced efficiency.

Real-Time Optimization

Real-Time Optimization (RTO) systems have emerged over the past few decades as a promising approach for mitigating process disturbances and enhancing operational efficiency by combining mathematical models with real-time process measurements [11]. As such, when integrated into a separation equipment, RTO frameworks may offer means to maintain near-optimal operating conditions under changing process environments. Despite this potential, their practical implementation remains challenging. These systems require representative process models for better performance, which can be particularly difficult to obtain due to the nonlinear behavior and complexity of the underlying dynamics, and a reliable model-updating mechanism capable of

predicting and updating process parameters to compensate plant-model mismatch. Additionally, for cyclic processes such as an SMB, RTO strategies often rely on the system reaching a cyclic steady state, further complicating and limiting the applicability of these methods [12, 13]. Overall, an RTO system consists of a group of optimization techniques designed to continuously evaluate and adjust process operating conditions, consequently enabling the autonomous implementation of business-oriented decisions aimed at enhancing process efficiency and performance, while simultaneously ensuring compliance with process and operational restrictions [14, 15]. Within the plant hierarchy, RTO operates as an intermediate layer between the business planning, scheduling, and management systems, and the lower-level process control systems. As such, RTO considers the information gathered and decisions made by the company managers, such as planning, scheduling, transportation, and stock availability of the raw materials, to assess opportunities for operational improvement in real time. The resulting optimal setpoints are then communicated to the control layer, where they are implemented within the process [16]. To address the challenges associated with RTO deployment, mentioned before, various RTO strategies have been proposed and developed over time. Depending on the specific characteristics of the problem, these approaches are typically classified into three main categories: Model-Parameter Adaptation, Modifier Adaptation, and Direct Input Adaptation [12].

This work aims to build a virtual foundation to tackle these challenges by developing a virtual model pair, with different mathematical models, to mimic a real model-based RTO framework where plant-model mismatch and structural differences are present. This virtual pair, comprised of a simpler predictive model and a more complex representative plant model as a stand-in for the SMB unit, will enable the study of these limitations and how they can affect the operation of a real-scenario RTO system implemented into an SMB. This case study was specifically designed for the optimal separation of a racemic mixture of ketoprofen enantiomers, using SMB chromatography. To help build the structure of the model pair, particularly for the predictive model, a Particle Swarm Optimization (PSO) technique was also employed.

Particle Swarm Optimization

The Particle Swarm Optimization (PSO) algorithm, introduced by Kennedy and Eberhart in 1995, was originally developed for the optimization of continuous nonlinear functions. Drawing inspiration from swarm intelligence, evolutionary computation, and genetic algorithms, PSO produces a population of candidate solutions, usually referred to as particles, that explore the search space by iteratively updating their positions based on both individual experiences and the collective knowledge of the swarm [17]. Since its inception, numerous modifications and enhancements have been proposed to improve the convergence speed, solution quality, and robustness of the original PSO algorithm. Despite the variety of PSO variants, all algorithms share a common underlying procedure. Initially, particles are distributed throughout the design space, either randomly or pseudo-randomly, to enable a controlled and comprehensive exploration of the problem's topology. Particle movement is governed by velocity vectors, which are updated based on each particle's own best-known position and the best-known position of the whole swarm. As particles traverse the search space and discover improved solutions, these newly identified best positions influence subsequent updates, dynamically guiding the swarm's exploration. Through this iterative process, the swarm progressively converges toward promising regions of the search space, facilitating the identification of near-optimal solutions for the optimization problem at hand [18].

By employing this methodology in the present study, the predictive model of the virtual pair is able to more accurately replicate the behavior of the representative plant model. This improvement is achieved by using Particle Swarm Optimization (PSO) to systematically minimize the residuals between the model's predicted outputs and the observed data, resulting in a closer and more reliable match between the two.

To summarize, the novelty of this work lies in the creation of this virtual plant–model pair, which provides a controlled environment to systematically study and understand the effects of plant-model discrepancies, providing a stepping stone toward robust RTO design. Obtaining favorable results within this study provides valuable insight into the possibility of employing similar mathematical models, with plant-model mismatch relative to the plant, for an analogous real-scenario RTO.

METHODOLOGY

Firstly, the mathematical model for an SMB unit must be built. This was accomplished by adapting an existing, experimentally validated gPROMS model into a MATLAB framework. Although straightforward, the implementation required the use of different approaches, due to fundamental differences between the gPROMS and MATLAB programming environments, each of which offers distinct advantages and limitations. Upon completing the MATLAB implementation, the model was further validated against two literature-reported case studies to ensure robust and reliable performance [19, 20].

As a subsequent step, the original MATLAB model was transformed into a virtual model pair, consisting of a simplified predictive model and a more detailed representative model that serves as a stand-in for the plant in

a typical RTO scenario, as mentioned before. While both models share a similar structural foundation, having been derived from the same original MATLAB model, structural differences were purposely inserted to reflect the discrepancies commonly observed between model and plant in practical RTO implementations.

Simulated Moving Bed Models

Both models are dictated by a set of simplifying assumptions. In this section, the differences between the developed mathematical models will be displayed. For the fixed bed columns within the SMB unit, the system was considered isothermal, with a constant mobile phase velocity, typical for diluted mixtures, and no interactions between solutes. The gel particles were considered spherical and uniform, and the fluid phase was modeled as exhibiting plug flow with axial dispersion. With these assumptions, the mass balance for the fluid phase is as follows:

For the predictive model:

$$\frac{\partial C_i}{\partial t} = D_{ax}\frac{\partial^2 C_i}{\partial z^2} - u_i\frac{\partial C_i}{\partial z} - \frac{(1-\varepsilon_b)}{\varepsilon_b}\frac{\partial \bar{q}_i}{\partial t} \tag{1}$$

where the parameter D_{ax} refers to the axial dispersion coefficient, u_i is the interstitial fluid velocity, ε_b is the fixed bed porosity, C_i is the concentration of component i in the fluid phase and $\bar{q}_i$ is the average adsorbed concentration of enantiomer i. For the representative model, the mass balance for the fluid phase is instead:

$$\frac{\partial C_i}{\partial t} = D_{ax}\frac{\partial^2 C_i}{\partial z^2} - u_i\frac{\partial C_i}{\partial z} - \frac{(1-\varepsilon_b)}{\varepsilon_b}\frac{3}{R_p}k_L(C_i - \overline{C_{p,i}}) \tag{2}$$

where R_p is the radius of the adsorbent particle, $\overline{C_{p,i}}$ is the average concentration of component i in the particle pores, and k_L is a lumped parameter, defined by:

$$\frac{1}{k_L} = \frac{1}{k_F} + \frac{1}{R_p\,\varepsilon_p\,k_p} \tag{3}$$

in which k_F is the mass transfer coefficient in the film, while k_p is the mass transfer coefficient in the pores and ε_p is the particle porosity.

For the mass balance in the particles, slight differences also appear within the models. In the predictive model, this mass balance is described by the linear driving force (LDF) equation:

$$\frac{\partial \bar{q}_i}{\partial t} = k_{LDF,i}(q_i^* - \bar{q}_i) \tag{4}$$

in which q_i^* is the adsorbed concentration in equilibrium with the fluid phase concentration of component i, while $k_{LDF,i}$ is the mass transfer coefficient. In the representative model, a more complex mass balance that utilizes the aforementioned lumped parameter k_L is used:

$$\varepsilon_p\frac{\partial \overline{C_{p,i}}}{\partial t} + (1-\varepsilon_p)\frac{\partial \bar{q}}{\partial t} = \frac{3}{Rp}\,k_L(C_i - \overline{C_{p,i}}) \tag{5}$$

Lastly, the equilibrium relation between the liquid and solid phases can be represented by any chosen isotherm equation.

$$q_i = f(C_i) \tag{6}$$

Or, for the representative model:

$$\bar{q}_i = f(\overline{C_{p,i}}) \tag{7}$$

Five different equations were considered and developed to be implemented into the SMB mathematical models. These were Linear, Langmuir, Bi-Langmuir, Langmuir + Linear and Bi-Langmuir + Linear. To account for material transport between columns and across inlet/outlet ports, where no separation occurs, transfer lines were introduced into the SMB model, in which plug-flow behavior is assumed. Therefore, this approach also physically accounts for the system's dead volumes.

The model equations were solved using MATLAB® intrinsic functions, including *ode45f* and *pdepe,* with specific adaptations performed as needed. A spatial discretization of 40 intervals was employed for the adsorption columns, while 10 intervals were used to discretize the dead volumes elements. Temporal discretization was performed using 100 time intervals.

Virtual Plant-Model Structure – Representative Model

To generate the *in silico* experiments, the representative model is utilized. In this particular study, the Linear + Langmuir isotherm was chosen. The model is executed from a primary MATLAB script that defines the operating conditions for each experiment using a Latin Hypercube Sampling (LHS) strategy, ensuring a space-filling exploration of the selected input domain. In this study, the desirable operating variables to be sampled include the feed flowrate, Q_F, the eluent flowrate, Q_D, the extract flowrate, Q_X, the raffinate flowrate, Q_R, the recycling flowrate, Q_{IV}, and the switching time, t_s. However, correct operation of the SMB unit requires that several restrictions on the operating conditions have to be strictly satisfied. The first among these is the global mass balance, as described in Equation (20). Compliance with this constraint is essential for physically meaningful operation. Therefore, only three operating conditions should be sampled, while the remaining one is computed implicitly to satisfy the overall mass balance. Additionally, for a successful SMB operation, the flowrate hierarchy shown in the following inequality must be satisfied:

$$Q_D > Q_X > Q_R > Q_F \quad (8)$$

Finally, physical constraints must also be considered. In particular, the pumps used in the FlexSMB-LSRE® system have a maximum capacity of 50 cm^3/min [4], so, none of the inlet flow rates should exceed this limit. In practice, additional margin is typically maintained for safety, and prior studies and experiments indicate that the operating flow rates remain well below this maximum. Therefore, the following upper bounds were imposed:

$$Q_{F,max} = Q_{D,max} = Q_{IV,max} = 20\ cm^3 min^{-1} \quad (9)$$

Imposing these constraints requires additional modifications to the sampled operating conditions set. Although LHS is not inherently designed to generate feasible values for dependent variables under strict inequality and mass-balance restrictions, it still offers important advantages as a sampling strategy, and its limitations can be mitigated. In particular, a rejection-based brute-force method can be introduced into the LHS procedure, such that candidate operating conditions are accepted only if all constraints are satisfied, otherwise, they are discarded and resampled. However, because the probability of pseudo-randomly generated samples meeting all constraints may be low, this approach can lead to excessive and unnecessary computational effort. For this reason, the sampling procedure was modified to increase the probability of generating feasible operating conditions as much as possible.

To improve efficiency and ensure compliance with the constraints, the desired operating conditions were generated using a sequential approach rather than directly from LHS. First, the lowest flowrate, the feed flowrate, Q_F, was sampled, as well as two increments. The sum of Q_F with the first increment I_1 would dictate the extract flowrate, Q_X, while the sum of this flowrate with the second increment I_2 would yield the eluent flowrate Q_D. The raffinate flowrate would be obtained through the global mass balance compliance. Both the flow rates and the increments were constrained by the maximum flow rate imposed by (28). In addition, a new restriction was introduced, the sum of the flowrate and the increments could not exceed the maximum allowed flowrate. This procedure ensures that all operating constraints are satisfied. The recycling flowrate and the switching time operating conditions were sampled directly through LHS. Therefore:

$$\theta_{OC,LHS} = [Q_F, I_1, I_2, Q_{IV}, t_s] \quad (10)$$

with all the remaining operating conditions being generated afterwards.

Under these constraints and boundaries, preliminary tests reveal that the LHS technique has an approximate probability of 17% of providing adequate operating conditions that ensure a correct SMB operation. Given the high efficiency of LHS in producing samples, this brute-force approach does not pose a significant computational burden within the overall methodology, and this likelihood value was deemed satisfactory.

Once the operating conditions were defined, these were employed in the representative mathematical model to generate *in silico* experiments. From these simulations, the history of the concentrations in the extract and raffinate streams were recorded and archived for subsequent use by the predictive model. To better mimic realistic plant operation and sensor limitations, measurement noise was introduced into the simulated outputs. Finally, the operating conditions and relevant metadata for each experiment were stored to facilitate the operation of the virtual model pair.

Virtual Plant-Model Structure – Predictive model

The predictive model, which has the goal of replicating the representative model behavior, is embedded within a PSO framework. The PSO algorithm minimizes the residuals between the outputs of the predictive model and the data generated by the representative model, by modifying the parameter values present in the predictive mathematical model. This is achieved by developing an appropriate objective function and then employing an adequate number of PSO particles and iterations. Each particle represents a complete set of *in silico* experiments, including their associated metadata, with the objective of identifying a set of parameter values that accurately predicts all of the representative model's outputs. The chosen parameters to be included in the estimated set, considering the option of using a Linear + Langmuir isotherm, are:

$$\theta_p = [k_{LDF,i}, D_{ax,c}, K_i, Q, b_i] \quad (11)$$

with i denoting the species, hence, $i = A, B$ in a binary mixture, and c pertaining to the column number, so, $c = 1, 2, \ldots, 6$. This means the PSO technique will estimate 13 parameters.

To create a sufficiently autonomous environment, it is necessary to implement fail-safes to address the possibility of simulation error or other computational issues. Given the structural differences between the virtual and representative models, the same operating conditions can produce different results, making two limitations particularly critical: the possibility of non-convergence, and excessively long computation times due to convergence. To address the first obstacle, the solution was to introduce a variable within the mathematical model that indicates convergence status. If the variable changes to a

state of non-convergence, a predefined objective function value is assigned, and the simulation is terminated prematurely. For the second challenge, a parallel timer was implemented using MATLAB's Parallel Computing Toolbox, allowing the simulation to be interrupted once a predefined maximum runtime is reached. If the simulation still does not finish after this time limit, the same predefined objective function value is assigned, ensuring robust and fully automated operation.

When simulations are successful, the objective function values are obtained by using a modified Mean Square Error (MSE) function:

$$J(\theta_p) = \sum_{j}^{N_{exp}} MSE_j \tag{12}$$

$$MSE = \frac{1}{N}\sum_{t}^{N}(y_{predictive,t} - y_{representative,t})^2 \tag{13}$$

in which j is the experiment number from a maximum of N_{exp} experiments and N is the amount of time instances where the residuals were determined. With this process, it may be possible to demonstrate that structural differences between the predictive and representative models, analogous to those encountered in RTO systems with a real unit and its mathematical model, do not prevent the successful operation of the framework.

RESULTS

First, *in silico* experiments were generated through the methodology described in the last section. The simulations were executed on a computer equipped with 32 GB of RAM and an Intel® Core™ i9-9900K CPU @ 3.60 GHz with eight processing cores. By taking advantage of MATLAB's Parallel Computing Toolbox, which allows one simulation to be executed per CPU core, eight experiments were generated in parallel, one per available core.

The operating conditions sampled by LHS for each experiment are presented in the Supplementary Material. The data yielded by the *in silico* experiments were properly archived and manipulated for an appropriate use by the virtual model and PSO framework. The total computational time required to generate all the *in silico* experiments was 190 seconds. Notably, since parallel computing was used, some experiments took more computational effort than others.

Before executing the PSO algorithm, a benchmark value was obtained as a means to measure the evolution of the objective function. As this separation process pertains to the separation of a racemic mixture of ketoprofen, parameters from literature were used to generate this benchmark value [20]. These parameters were introduced in the virtual model, and the alignment of the data was evaluated for each experiment, yielding a value of 0.3041 for the objective function value. It is important to note that these values were obtained comparing output average cyclic concentrations.

The PSO was then executed with 10 particles throughout 25 iterations. The set of parameters that best adjusted against the eight *in silico* experiments' datapoints are presented in Table 1.

Table 1: Set of parameter values generated by the PSO for the predictive model that best describe the plant model behavior.

Parameters	Value
$\boldsymbol{k_{LDF,A}}$ / min^{-1}	0.690
$\boldsymbol{k_{LDF,B}}$ / min^{-1}	0.303
$\boldsymbol{D_{ax,1}}$ / $cm^2 min^{-1}$	0.032
$\boldsymbol{D_{ax,2}}$ / $cm^2 min^{-1}$	0.069
$\boldsymbol{D_{ax,3}}$ / $cm^2 min^{-1}$	0.100
$\boldsymbol{D_{ax,4}}$ / $cm^2 min^{-1}$	0.021
$\boldsymbol{D_{ax,5}}$ / $cm^2 min^{-1}$	0.073
$\boldsymbol{D_{ax,6}}$ / $cm^2 min^{-1}$	0.050
$\boldsymbol{K_A}$ / $-$	0.885
$\boldsymbol{K_B}$ / $-$	0.916
$\boldsymbol{Q}$ / $mg\ cm^{-3}$	8.037
$\boldsymbol{b_A}$ / $cm^3 mg^{-1}$	0.021
$\boldsymbol{b_B}$ / $cm^3 mg^{-1}$	0.042

By employing these parameters in the virtual model, the alignment between the virtual model pair yielded a value of 0.0336, lower than the benchmark set, showcasing improvement originated by the developed methodology. Another aspect that reveals extreme importance within this work is computational time, since an RTO framework needs to operate quickly to accompany the unit's changes in behavior in real time. To operate this PSO algorithm, as described before, all 250 simulations took approximately 37000 seconds, which is equivalent to 10.5 hours. Although this value is not appropriate for an in-line optimization framework, it is expected that during the operation of the RTO, less simulations will be employed, reducing the computational effort needed and, consequently, the operation time.

A revised simulation study comprising 16 distinct experiments, 2 per CPU core, was conducted using a PSO algorithm configured with 15 particles and 25

iterations, with the aim of evaluating the reproducibility and robustness of the proposed methodology. Additionally, the feasibility of reducing the number of discretization intervals as a strategy to decrease computational time without compromising result accuracy was investigated. To this end, the adsorption columns were discretized into 20 spatial intervals, while the dead volume elements were represented using 5 intervals. The corresponding operating conditions are provided in the Supporting Information. The estimated parameters obtained from this analysis are presented in Table 2.

Table 2: Set of parameter values generated by the PSO for the predictive model that best describe the plant model behavior, for 16 in silico experiments.

Parameters	Value
$k_{LDF,A}$ / min^{-1}	0.300
$k_{LDF,B}$ / min^{-1}	0.300
$D_{ax,1}$ / $cm^2 min^{-1}$	0.096
$D_{ax,2}$ / $cm^2 min^{-1}$	0.100
$D_{ax,3}$ / $cm^2 min^{-1}$	0.017
$D_{ax,4}$ / $cm^2 min^{-1}$	0.080
$D_{ax,5}$ / $cm^2 min^{-1}$	0.089
$D_{ax,6}$ / $cm^2 min^{-1}$	0.021
K_A / $-$	1.000
K_B / $-$	1.328
Q / $mg\ cm^{-3}$	9.632
b_A / $cm^3 mg^{-1}$	0.060
b_B / $cm^3 mg^{-1}$	0.040

In this case, the objective function values were computed using all available output stream data points, instead of only calculating residuals with average cyclic compositions, enabling a direct comparison between methodologies during the development of the virtual RTO framework. The best objective function value obtained using the PSO algorithm was 11.84, representing a significant improvement over the reference value of 24.27 calculated for the reference parameters. The total computational time required for this estimation was 15870 seconds, equivalent to 4.4 hours. These results indicate that the proposed methodology is readily reproducible and capable of delivering satisfactory optimization performance, even when reduced spatial discretization intervals are employed. This conclusion, together with the previously stated assumption regarding the reduced scale of the PSO inserted into the virtual RTO, is another guarantee that computational effort and operation time will not present a detrimental issue in later studies.

With these results, it is possible to observe that the virtual model pair test provides favorable information regarding the structure distinctions between the plant and the mathematical model. Since in this case study these differences were not detrimental, the same can still be expected from an analogous real-scenario RTO + SMB system.

CONCLUSIONS

This work presented a virtual-model pair that virtually represents the intricacies of the predictive elements of a real RTO framework. The goal of this work pertained to explore the possibility of a virtual predictive model predicting the behavior of representative model simulations, denoted as *in silico* experiments, much like what happens in the predictive structure of a common RTO system. With the proposed methodology, eight different *in silico* experiments were correctly executed, each one with its own operating conditions. Under a PSO algorithm, parameters were generated that enable the virtual model to accurately predict the behavior of the representative model, even though structural differences were present between mathematical models. Similar conclusions were drawn from the study involving 16 experiments, with the additional finding that reducing spatial discretization intervals significantly decreases computational effort without compromising optimization performance.

For future work, an uncertainty evaluation is deemed necessary as the first next step. Moreover, the creation of a more robust estimation methodology is proposed, with more diverse operating conditions and with the ability of exploring other adsorption isotherm structures, while maintaining the fast simulation times, essential for an adequate operation of an RTO framework.

DIGITAL SUPPLEMENTARY MATERIAL

The tables for the Operating Conditions of each set of experiments can be seen in the document: https://PSEcommunity.org/LAPSE:2026.0031.

ACKNOWLEDGEMENTS

This work was financially supported by Fundação para a Ciência e a Tecnologia, I.P. /MCTES through national funds: LSRE-LCM, UID/50020/2025 (DOI: 10.54499/UID/50020/2025); and ALiCE, LA/P/0045/2020 (DOI: 10.54499/LA/P/0045/2020).

AUTHOR IDENTIFIERS

Author ORCIDs:
Amaral GC: 0000-0002-7736-5759
Ferreira AFP: 0000-0002-6746-8973
Ribeiro AM: 0000-0003-4269-1420
Nogueira IBR: 0000-0002-0963-6449
Rodrigues D: 0000-0001-7823-2993

REFERENCES

1. Sanz-Medel A, Blanco-González E. Chiral trace-element speciation in biological samples: present importance and application to speciation for seleno-amino acids. TrAC Trends in Analytical Chemistry 21:709-716 (2002). https://doi.org/10.1016/s0165-9936(02)01114-7
2. Ceramella J, Iacopetta D, Franchini A, De Luca M, Saturnino C, Andreu I, Sinicropi MS, Catalano A. A look at the importance of chirality in drug activity: some significative examples. Applied Sciences 12:10909 (2022). https://doi.org/10.3390/app122110909
3. Lorenz H, Seidel?Morgenstern A. Processes to separate enantiomers. Angew Chem Int Ed 53:1218-1250 (2014). https://doi.org/10.1002/anie.201302823
4. Gomes PS, Zabkova M, Zabka M, Minceva M, Rodrigues AE. Separation of chiral mixtures in real SMB units: the flexsmb?Isre®. AIChE Journal 56:125-142 (2009). https://doi.org/10.1002/aic.11962
5. Rodrigues AE, Pereira C, Minceva M, Pais LS, Ribeiro AM, Ribeiro A, Silva M, Graça N, Santos JC. Principles of simulated moving bed. Simulated Moving Bed Technology :1-30 (2015). https://doi.org/10.1016/b978-0-12-802024-1.00001-x
6. Charton F, Nicoud RM. Complete design of a simulated moving bed. Journal of Chromatography A 702:97-112 (1995). https://doi.org/10.1016/0021-9673(94)01026-b
7. Ruthven DM, Ching CB. Counter-current and simulated counter-current adsorption separation processes. Chemical Engineering Science 44:1011-1038 (1989). https://doi.org/10.1016/0009-2509(89)87002-2
8. Sá Gomes P, Rodrigues AE. Simulated moving bed chromatography: from concept to proof?of?concept. Chem Eng & Technol 35:17-34 (2011). https://doi.org/10.1002/ceat.201100281
9. Kawajiri Y, Biegler LT. Optimization strategies for simulated moving bed and powerfeed processes. AIChE Journal 52:1343-1350 (2005). https://doi.org/10.1002/aic.10736
10. Toumi A, Engell S, Diehl M, Bock HG, Schlöder J. Efficient optimization of simulated moving bed processes. Chemical Engineering and Processing: Process Intensification 46:1067-1084 (2007). https://doi.org/10.1016/j.cep.2006.06.026
11. Bonvin D. Special issue “real-time optimization” of processes. Processes 5:27 (2017). https://doi.org/10.3390/pr5020027
12. Chachuat B, Srinivasan B, Bonvin D. Adaptation strategies for real-time optimization. Computers & Chemical Engineering 33:1557-1567 (2009). https://doi.org/10.1016/j.compchemeng.2009.04.014
13. Khajuria H, Pistikopoulos EN. Optimization and control of pressure swing adsorption processes under uncertainty. AIChE Journal 59:120-131 (2012). https://doi.org/10.1002/aic.13783
14. Câmara M, Quelhas A, Pinto J. Performance evaluation of real industrial RTO systems. Processes 4:44 (2016). https://doi.org/10.3390/pr4040044
15. Marchetti A, François G, Faulwasser T, Bonvin D. Modifier adaptation for real-time optimization—methods and applications. Processes 4:55 (2016). https://doi.org/10.3390/pr4040055
16. Seborg, D.E., et al., *Process dynamics and control*. 2016: John Wiley & Sons.
17. Kennedy J, Eberhart R. Particle swarm optimization. Proceedings of ICNN'95 - International Conference on Neural Networks 4:1942-1948 (None). https://doi.org/10.1109/icnn.1995.488968
18. Venter G, Sobieszczanski-Sobieski J. Particle swarm optimization. AIAA Journal 41:1583-1589 (2003). https://doi.org/10.2514/2.2111
19. Rios AG, Ribeiro AM, Rodrigues AE, Ferreira AFP. Bovine serum albumin and myoglobin separation by size exclusion SMB. Journal of Chromatography A 1628:461431 (2020). https://doi.org/10.1016/j.chroma.2020.461431
20. Ribeiro AE, Gomes PS, Pais LS, Rodrigues AE. Chiral separation of ketoprofen enantiomers by preparative and simulated moving bed chromatography. Separation Science and Technology 46:1726-1739 (2011). https://doi.org/10.1080/01496395.2011.582070

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Rolling-Horizon Scheduling for Dynamic Market-Driven Operation of an Air Separation Plant

Kieran McKenzie[a], Christopher L. E. Swartz[a]*

[a] McMaster University, Department of Chemical Engineering, Hamilton, ON, Canada
* Corresponding Author: swartzc@mcmaster.ca.

ABSTRACT

Cryogenic air separation units (ASUs) are the primary industrial technology for producing high purity oxygen, nitrogen, and argon gases at commercial scale. Cryogenic ASUs are large consumers of electricity, making them ideal candidates for market-driven operation research in today's volatile and uncertain manufacturing environments. To maximize profitability, ASU operation must dynamically adapt to changing market conditions as they evolve. This work explores the implementation of a rolling-horizon scheduling (RHS) strategy for the real-time market-driven operation of a high-dimensional ASU model with inventory, responding to uncertainty in future plant demand and electricity price forecasts by periodically rescheduling in response to updated market information. A dynamic latent variable-based surrogate model (LV-SM) is used within the scheduling framework as a computationally efficient substitute for an existing first-principles-based ASU model. Results show that RHS and plant inventory are effective strategies for handling uncertainties in the future market forecasts, while the LV-SM shows computational performance suitable for real-time implementation. With this scheduling framework in place, extensions involving expanded case studies and uncertainty-aware optimization are planned for future work.

Keywords: Dynamic Optimization, Surrogate Modeling, Air Separation, Rolling-horizon, Scheduling, Principal Component Analysis, Neural Network

INTRODUCTION

Energy-intensive cryogenic air separation units (ASUs) are ideal candidates for real-time market-driven scheduling and optimization studies in today's uncertain and volatile manufacturing environments. To maximize profitability, ASU operation must dynamically adapt to changing market conditions as they evolve. Rolling-horizon scheduling (RHS) strategies provide a promising scheduling approach for ASU systems when responding to uncertainty in market or plant parameters. In these approaches, an initial plant schedule may be iteratively refined according to new market data or plant conditions. For instance, work by Pattison et al. [1] and Kelley et al. [2] studied the closed-loop RHS of an ASU using reduced-order scale bridging models, where rescheduling was performed in response to forecast updates or disturbances in parameters such as electricity prices, ambient temperature, and plant demand. Economic model predictive control (eMPC) schemes are also implemented in a rolling-horizon fashion and have been applied to ASU systems in work such as Schäfer et al. [3] and Caspari et al. [4]. Alternative scheduling approaches which explicitly model parameter uncertainty have also been explored, including robust optimization [5], chance-constrained programming [6], and stochastic programming techniques [7].

This work explores a rolling-horizon scheduling (RHS) strategy for the market-driven operation of a high-dimensional ASU model. To achieve computation times suitable for real-time implementation, the full-order model (FOM) originally developed by Cao et al. [8] is substituted by a dynamic surrogate model. This work implements a latent variable-based surrogate model (LV-SM), previously developed and applied in McKenzie et al. [9]. The LV-SM is trained from FOM simulation data and uses a ReLU-activated neural network to perform dynamic step-ahead predictions within a latent-variable subspace derived from principal component analysis (PCA). This modeling technique enables substantial computation time reduction, trust region enforcement to promote scheduling solution validity, and the ability to

approximate all original states of the FOM. The viability of RHS approaches using the LV-SM for real-time market-driven ASU operation under specific market conditions is assessed in this work. With the rolling-horizon framework in place, future work may consider expanded case studies which identify areas where explicit consideration of uncertainty in the problem formulation using a rolling-horizon stochastic scheduling strategy yields improved economic performance and operational robustness over deterministic alternatives.

PROCESS MODEL DESCRIPTION

A process flow diagram of the ASU model used as a basis for this study is provided in **Figure 1**.

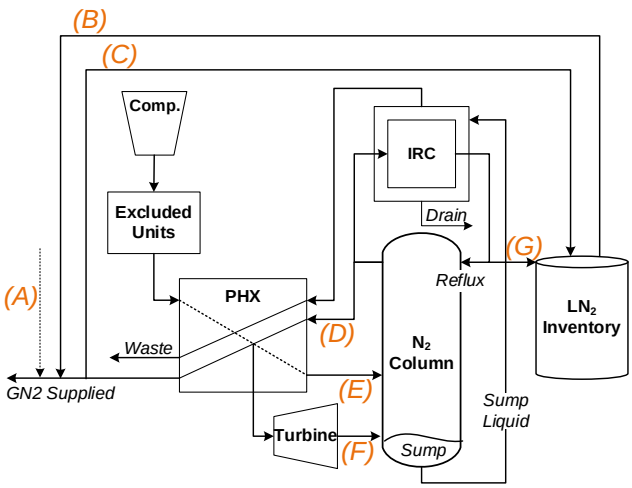

Figure 1. Cryogenic ASU diagram. Optimization-relevant variables are labelled in orange text as follows: (A) purchased GN_2, (B) vaporized LN_2, (C) condensed LN_2, (D) GN_2 produced, (E) liquid air feed, (F) vapor air feed, (G) LN_2 produced. Abbreviations: integrated reboiler-condenser (IRC), primary heat exchanger (PHX), compressor (Comp), gaseous nitrogen (GN_2), liquid nitrogen (LN_2).

The process begins with the compression and purification of ambient air which is cooled against cryogenic product and waste streams from the column in a multi-stream primary heat exchanger (PHX). A fraction of the air feed is drawn partway through the PHX and is expanded via a turbine prior to entering the column as vapor air. The remaining air enters the column as liquid air after passing entirely through the PHX. The column operates at cryogenic conditions and is configured to separate air into high-purity gaseous nitrogen (GN_2) and an oxygen-rich sump liquid. The GN_2 either exits the system as product via the PHX or continues to the integrated reboiler-condenser (IRC) where heat exchange occurs with the sump liquid. The resulting liquid nitrogen (LN_2) may return to the column as reflux or be stored as inventory. During market-driven operation, inventory may be vaporized or condensed as needed. The ASU was developed using the gPROMS software as a first-principles model containing 3819 state variables, adapted from Cao et al. [8]. For further modeling details, the reader is referred to Cao et al. [8].

SURROGATE MODEL DESCRIPTION

This work implements a latent variable-based surrogate model (LV-SM), previously developed and applied by McKenzie et al. [9]. To address the high dimensionality of the FOM, the LV-SM performs dynamic step-ahead predictions within a latent variable subspace derived from principal component analysis (PCA). PCA can be used to transform a high-dimensional matrix $X \in \mathbb{R}^{n \times j}$ into a low dimensional latent variable subspace $T \in \mathbb{R}^{n \times a}$ by:

$$T = XP \tag{1}$$

where $P \in \mathbb{R}^{j \times a}$ represents the PCA model loadings with columns representing principal components. The optimal low-rank approximation of the original matrix, $\hat{X}$, is calculated as:

$$\hat{X} = TP^T \tag{2}$$

In the LV-SM, a 10-component PCA model transforms a vector of initial FOM states at steady-state operation $\boldsymbol{x}_0^{SS} \in \mathbb{R}^{3819}$ into a vector of initial latent states $\boldsymbol{\tau}_0 \in \mathbb{R}^{10}$, drastically reducing the problem dimension. Dynamic step-ahead predictions of latent states, in addition to an estimation of the squared prediction error (SPE) statistic derived from the PCA model, are predicted as a function of latent states and surrogate model inputs at discrete time steps using a ReLU-activated multilayer perceptron (MLP) neural network:

$$\left[\mathrm{SPE}_{k+1}, \boldsymbol{\tau}_{k+1}^T\right]^T = \boldsymbol{f}_{\mathrm{N}}(\boldsymbol{\tau}_k, \boldsymbol{u}_k)\ \forall k \in K^- \tag{3}$$

where $\boldsymbol{\tau}_k$ and $\boldsymbol{u}_k$ represent the vector of latent states and surrogate model inputs at a given time step k and $\boldsymbol{f}_{\mathrm{N}}$ represents the neural network. In this work, the set $K = \{0, \dots, N_{\mathrm{K}}\}$ denotes all discrete time points, $K^- = \{0, \dots, N_{\mathrm{K}} - 1\} \subset K$ denotes the subset of discrete time points excluding the last, and $K^+ = \{1, \dots, N_{\mathrm{K}}\} \subset K$ denotes the subset of discrete time points excluding the initial point. Constrained FOM states are calculated using the PCA reverse-transform while all other FOM states are excluded from the optimization model for efficiency. The SPE metric is constrained within confidence limits during optimization to enforce a trust-region. Since the SPE metric describes the prediction error of a particular observation, an SPE value exceeding confidence limits suggests poor PCA model representation. Confidence limits were computed based on the SPE distribution of the surrogate

model's training data using critical values from a chi-squared distribution [10]. The 95% and 99% SPE confidence limits are commonly used thresholds. The 99% confidence limit is used here and was shown in previous work to promote surrogate model prediction accuracy while being less restrictive than lower limits [9].

The LV-SM was trained from FOM simulation data collected using a bounded random walk sampling method. A step input from the range $[-h_i, +h_i]$, where h_i is the allowable step size per input variable u_i, was applied to each input at an interval of $\Delta t = 0.25$h to match the control intervals used in the optimization formulation. Criteria were enforced during data collection to ensure that data remained physically realizable and within an operation region relevant to the optimization problem. The storage vessel was excluded from the FOM during data collection due to its simplicity and was modeled after LV-SM development without reduction. **Figure 2** shows the process flow diagram after substitution of the LV-SM.

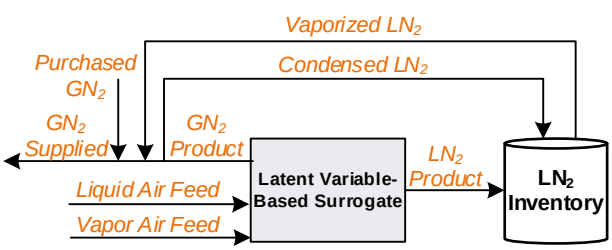

Figure 2. Substitution of the FOM by the LV-SM. Optimization-relevant variables are shown in orange text.

By operating within a reduced dimensional space, the LV-SM showed nearly a 1000-fold reduction in optimization computation time in previous work [9], while retaining high accuracy compared to the FOM in multistep-ahead simulation and dynamic optimization. The reader is referred to McKenzie et al. [9] for more detailed information on data collection, training, and validation of the LV-SM. Note that McKenzie et al. [9] embedded neural network equations in the optimization formulation using complementarity conditions. Direct inclusion of complementarity constraints often causes numerical difficulties due to unsatisfied nonlinear programming constraint qualifications [11]. Complementarity constraint reformulation using the exact penalty reformulation presented in Ralph and Wright [12] is used in this work and showed improved solver performance. This technique introduces a complementarity constraint violation term to the objective function with a tunable penalty parameter used to enforce constraint satisfaction.

MARKET-DRIVEN OPTIMIZATION UNDER UNCERTAINTY

A rolling-horizon scheduling (RHS) strategy is proposed to respond to market uncertainty by dynamically adjusting plant operation and inventory use according to evolving market forecasts. A schematic of the framework is shown in **Figure 3**.

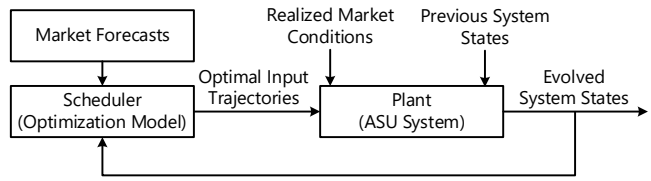

Figure 3. RHS strategy framework.

Upon each execution of the framework, an optimization model determines an optimal schedule over a specified horizon in anticipation of forecasted market conditions. An initial portion of the schedule, referred to as the rescheduling interval, is applied to the ASU plant. As time progresses forward in a rolling-horizon fashion, the remaining schedule is periodically refined upon in response to updated market forecasts and evolved system states. The LV-SM serves as the ASU model in the framework to enable real-time implementation while capturing the ASU's dynamic behavior at fine time scales. Future work may consider extension of the framework to explicitly model forecast uncertainty using stochastic optimization, such as done in Mathur et al. [13] for a hydropower system. Details of the optimization model, market forecasts, and plant model are provided in the following sections.

Optimization Model Formulation

An optimal schedule consists of economically optimal operational decisions over the scheduling horizon, determined in response to market forecasts while satisfying operational and process model constraints. This is accomplished through the minimization of an economic objective function. Additional terms accounting for input move suppression and complementarity constraint satisfaction are also minimized. The deterministic objective function z is defined as:

$$\min_{u_{i,k}} z = \Delta t \sum_{k\in K^-} \psi_k^{\text{Econ}} + m \sum_{i\in I} \sum_{k\in K} \eta_{i,k} + \frac{1}{\mu} \sum_{k\in K} \sum_{n_{\text{HL}}\in L} \zeta_{k,n_{\text{HL}}} \quad (4)$$

where ψ_k^{Econ} represents the economic objective function (negative of profit), $\eta_{i,k}$ is an input move suppression term, $\zeta_{k,n_{\text{HL}}}$ is the complementarity constraint penalty term, the parameters m and μ represent relative weights of the move suppression and complementarity constraint satisfaction terms, respectively, where $m > 0$ and $0 < \mu \leq 1$. Tuning was performed to select penalty weights $m = 0.1$ and $\mu = 0.01$. A value of $m = 0.1$ was observed to reduce nonprofitable fluctuations in the decision variables (as was observed in McKenzie et al. [9]) without

economic trade-off, while a value of $\mu = 0.01$ was observed to ensure complementarity constraint satisfaction. Further reductions in μ place higher weight on complementarity constraint satisfaction, however, were noted to increase computation time without providing additional benefit. The set $L = \{1, \dots, N_{HL}\}$ is the set of all hidden-layer nodes in the MLP network and $I = \{1, \dots, 7\}$ denotes the set of decision variables. The decision variables are defined as follows for all $k \in K^-$:

- $u_{1,k}$: Liquid air feed flowrate to the column, F_k^{LF}
- $u_{2,k}$: Vapor air feed flowrate to the column, F_k^{VF}
- $u_{3,k}$: LN_2 production from IRC to inventory, $F_k^{\mathrm{prod,\,LN_2}}$
- $u_{4,k}$: GN_2 flowrate from column to product stream via the PHX, $F_k^{\mathrm{GN_2}}$
- $u_{5,k}$: Vaporized LN_2 flowrate, $F_k^{\mathrm{vap,\,LN_2}}$
- $u_{6,k}$: GN_2 liquefaction flowrate, $F_k^{\mathrm{cond,\,GN_2}}$
- $u_{7,k}$: Amount of GN_2 purchased externally, F_k^{OS}

These variables are applied to the plant in a piecewise-constant manner in regularly spaced control intervals of $\Delta t = 0.25$h, which match the prediction interval of the LV-SM. The economic objective function ψ_k^{Econ} represents the negative of the hourly operating profit and is calculated as:

$$\psi_k^{\mathrm{Econ}} = C_k^{\mathrm{elec}} W_k^{\mathrm{comp}} + C_k^{\mathrm{cond}} F_k^{\mathrm{cond,\,GN_2}} + C_k^{\mathrm{vap}} F_k^{\mathrm{vap,\,LN_2}} + C_k^{\mathrm{OS}} F_k^{\mathrm{OS}} - C_k^{\mathrm{GN_2}} F_k^{\mathrm{sale}} \forall k \in K^- \quad (5)$$

where $C_k^{\mathrm{elec}}, C_k^{\mathrm{cond}}, C_k^{\mathrm{vap}}, C_k^{\mathrm{OS}}$, and $C_k^{\mathrm{GN_2}}$ represent prices of electricity, GN_2 liquefaction, LN_2 vaporization, GN_2 purchased from outside sources, and sold GN_2 product, respectively. Compressor shaft work is represented by W_k^{comp}; $F_k^{\mathrm{cond,\,GN_2}}$ and $F_k^{\mathrm{vap,\,LN_2}}$ represent the condensation rate and vaporization rate of GN_2; F_k^{OS} represents the amount of purchased GN_2; and F_k^{sale} represents the rate of GN_2 product sold. Note that the summation in the objective function excludes the final discrete point since the quantity $\Delta t \psi_{N_K-1}^{\mathrm{Econ}}$ accounts for costs accumulated over the final control interval. Sold GN_2 is calculated as:

$$F_k^{\mathrm{sale}} = F_k^{\mathrm{GN_2}} - F_k^{\mathrm{GN_2,\,cond}} + F_k^{\mathrm{LN_2,\,vap}} + F_k^{\mathrm{OS}} \ \forall\ k \in K^- \quad (6)$$

The move suppression term $\eta_{i,k}$ is calculated as the absolute difference between successive inputs:

$$\eta_{i,0} = |u_{i,0} - u_i^{\mathrm{prev}}| \quad \forall\ i \in I \quad (7)$$

$$\eta_{i,k} = |u_{i,k} - u_{i,k-1}| \quad \forall i \in I,\ k \in \{1, 2, \dots, N_K - 1\} \quad (8)$$

At the initial time point of each rolling-horizon execution, $\eta_{i,0}$ is calculated relative to input values u_i^{prev} which occur immediately prior to the beginning of the scheduling horizon. The values of u_i^{prev} are defined based on values from the previous rolling-horizon execution, except for the first execution of the algorithm where they are initialized using the plant's steady-state operating conditions. Since the derivatives of the absolute value function are discontinuous, these equations were reformulated using the following linear constraints:

$$\eta_{i,k} \geq 0 \quad \forall\ i \in I,\ k \in K^- \quad (9)$$

$$-\eta_{i,0} \leq u_{i,0} - u_i^{\mathrm{prev}} \leq \eta_{i,0} \quad \forall i \in I \quad (10)$$

$$-\eta_{i,k} \leq u_{i,k} - u_{i,k-1} \leq \eta_{i,k} \forall i \in I,\ k \in \{1, \dots, N_K - 1\} \quad (11)$$

Solutions of the optimization problem must satisfy the embedded LV-SM model (**Equation 3**). SPE trust region constraints constrain the predicted SPE statistic within its 99% confidence limit. Finally, several operational constraints summarized in **Table 1** are considered.

Table 1: Formulation of operational constraints.

Description	Formulation	
Purity	$\mathrm{PPMO}_k \leq \mathrm{PPMO}^{UB}$	$\forall k \in K^+$
Sump level	$l^{LB,\,\mathrm{sump}} \leq l_k^{\mathrm{sump}} \leq l^{UB,\,\mathrm{sump}}$	$\forall k \in K^+$
Reboiler level	$l^{LB,\,\mathrm{reb}} \leq l_k^{\mathrm{reb}} \leq l^{UB,\,\mathrm{reb}}$	$\forall k \in K^+$
Dew point at turbine outlet	$T_k^{\mathrm{dew,\,TBout}} \leq T_k^{\mathrm{TBout}}$	$\forall k \in K^+$
Storage level	$0 \leq V_k^{\mathrm{storage}} \leq V^{UB,\,\mathrm{storage}}$	$\forall k \in K^+$
Final storage capacity	$V_{N_k}^{\mathrm{storage}} \geq V^{\mathrm{final,\,storage}}$	$\forall k \in K^+$
Compressor surge	$P_k^{\mathrm{surge}} \leq 100\%$	$\forall k \in K^+$
Column flood	$P_k^{\mathrm{flood}} \leq 100\%$	$\forall k \in K^+$
Bubble point at liq. air feed	$P_k^{\mathrm{BP}} \leq 100\%$	$\forall k \in K^+$
Dew point at turbine inlet	$P_k^{\mathrm{DP}} \leq 100\%$	$\forall k \in K^+$
Demand met	$F_k^{\mathrm{sale}} = F_k^{\mathrm{GN_2,\,demand}}$	$\forall k \in K^-$
Input bounds	$u_i^{LB} \leq u_{i,k} \leq u_i^{UB}$	$\forall i \in I$, $k \in K^-$
Max step size in inputs	$\eta_{i,k} \leq h_i$	$\forall i \in I_{SM}$, $k \in K^-$

In **Table 1**, PPMO represents parts per million oxygen in the GN_2 product; l^{sump} and l^{reb} are sump and reboiler liquid levels; $T^{\mathrm{dew,\,TBout}}$ and T^{TBout} are dew point and discharge temperatures at the turbine outlet; V^{storage} represents storage vessel liquid volume; $P^{\mathrm{surge}}, P^{\mathrm{flood}}, P^{\mathrm{BP}}$, and P^{DP} represent values relative to constraint limits; $F^{\mathrm{GN_2,\,Demand}}$ and F^{sale} represent GN_2 demand and amount sold; and h_i are step sizes used for input variable move restriction which match those used during data collection. Lastly, $I_{\mathrm{SM}} = \{1, \dots, 4\} \subset I$ is the subset of input variables corresponding to the surrogate model inputs. The operational constraints of **Table 1** ensure that product is supplied at acceptable purity, equipment and safety

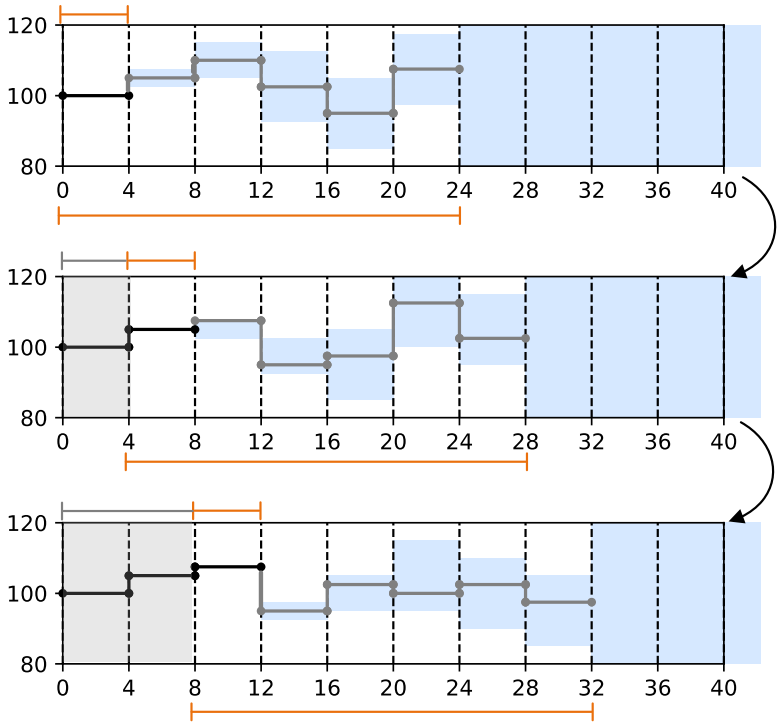

Figure 4. Example of three consecutive demand forecasts with simulated market updates. Realized demand levels are shown using black solid lines, forecasted demand levels using gray solid lines, and estimated demand regions using blue shading.

specifications are met, and modeling assumptions are upheld. Note that the storage vessel volume percentage at the end of the scheduling horizon is maintained above a threshold, $V^{\text{final, storage}}$, which is specified at 25%. An additional constraint enforces equality between the amount of product sold and the product demand. By **Equation 6**, any oversupplied product will be condensed to storage while any unmet demand will be supplied either from inventory or by purchase from outside sources. Move restrictions are enforced on surrogate model input variables to ensure agreement with the step sizes used in the surrogate model's training process.

Finally, the optimization problem is subject to the following set of initial conditions:

$$\boldsymbol{\tau_0} = \boldsymbol{\tau}_{\text{init}} \tag{12}$$

$$V_0^{\text{storage}} = V_{\text{init}}^{\text{storage}} \tag{13}$$

where $\boldsymbol{\tau}_{\text{init}}$ and $V_{\text{init}}^{\text{storage}}$ are the initial conditions for the PCA model scores and storage vessel volume percentage. For the first execution of the rolling-horizon algorithm, the initial scores are calculated as the product of a vector of FOM states at a known steady-state operation and the PCA model loadings, $\boldsymbol{x}_0^{\text{SS}}P$, whereas the initial storage vessel volume percentage is specified to be 25%. For all rolling-horizon executions after the first, these initial conditions are updated based on evolved system states from the plant.

Plant Model Description

Scheduling decisions corresponding to the rescheduling interval are applied to the ASU plant upon each execution of the rolling-horizon strategy, as shown in **Figure 3**. Within the simulation model:

- Market parameters are at their realized conditions.
- Control trajectories for decision variables $F^{LF}, F^{\text{VF}}, F^{\text{LN}_2}$, and F^{GN_2} are fixed at their values determined by the scheduler.
- Decision variables, $F^{\text{LN}_2,\text{vap}}, F^{\text{LN}_2,\text{cond}}$, and F^{OS}, can be recalculated in response to realized demand. Unmet demand is filled using inventory, if possible, or by purchasing from an external source at a high cost. Alternatively, oversupplied product is condensed to inventory to avoid wasted product.

The final item from the list above enables the plant to respond to market events that were not forecasted. As in the optimization model, the LV-SM model is used as a substitute for the FOM during simulation based on the

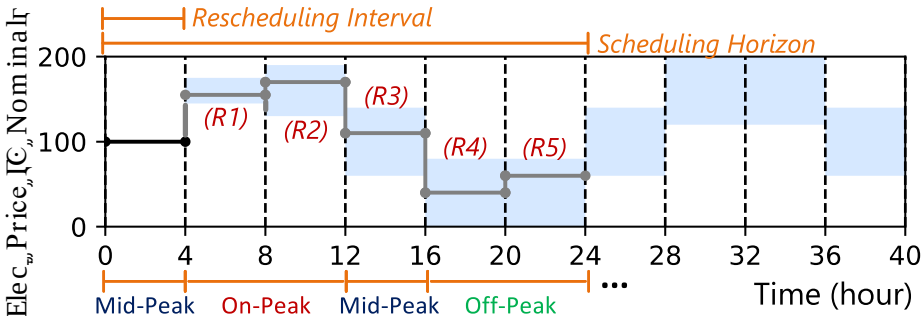

Figure 5: Example electricity price forecast. Realized prices are shown using black solid lines, forecasted prices levels using gray solid lines, and estimated price regions using blue shading.

assumption that discrepancies between simulations of the FOM and the LV-SM are sufficiently small. This assumption is supported by the high approximation accuracy of the LV-SM [9].

Market Parameter Forecasts

In the absence of historical operating data, forecast models are assumed for the product demand and electricity price. These models generate forecasts by sampling a forecast scenario from within estimated regions where market parameter values are expected to fall. As time progresses, these estimated regions are updated in response to simulated evolution in market conditions with each successive rolling-horizon execution. Lastly, forecasted values near the beginning of each execution's scheduling horizon are assumed more accurate than latter values to reflect increased market uncertainty with time. An example of three consecutive demand forecasts is shown in **Figure 4**.

The product demand is allowed to vary every four hours and is bounded within ±20% of the initial demand value. At each execution, demand is assumed known only over the rescheduling interval. Estimated demand regions are labelled R1-R5, where the size of each region reflects the amount of forecast uncertainty present and relates to how far into the future that region is. Market updates are simulated with each new execution — demand within the rescheduling interval is realized by drawing uniformly from the previous forecast's R1 region, and new R1-R5 regions are vertically centered at randomly sampled levels within ranges corresponding to the previous execution. A demand forecast is generated by drawing uniformly from the discrete set of possible realized demand levels D = {80%, 82.5%, ..., 120%} which intersect with the estimated demand regions. Two demand forecast models with differing levels of forecast uncertainty are considered. The first demand forecast model (DF1) assumes that regions R1, R2, and R3-R5 have uncertainty band widths of 5, 10, and 20 units, respectively. The second demand forecast model (DF2) assumes that regions R1, R2, and R3-R5 have widths of 20, 30, and 40 units, respectively.

Two electricity price forecast models are assumed in a similar manner; however, the electricity price is assumed to follow a structured and repeating daily cycle to reflect the seasonality present in some real-world electricity prices. Each 24-hour period is divided into four repeating segments — mid-peak (four hours), on-peak (eight hours), mid-peak (four hours), and off-peak (eight hours). An example forecast is shown in **Figure 5**.

Two electricity price forecast models with differing levels of forecast uncertainty are considered. The first model (EF1) assumes that regions R1, R2, and R3-R5 have uncertainty band widths of 30, 60, and 80 units, respectively. The second model (EF2) assumes that regions R1-R2 and R3-R5 have widths of 80 and 160 units, respectively. Uncertainty regions for off-peak, mid-peak, and on-peak prices are clipped if they exceed specified ranges to ensure that seasonality in the price profiles is maintained. Electricity prices are represented as percentages relative to an assumed nominal price. Scaled off-peak prices are assumed to always fall within the range of [0%, 80%], mid-peak prices always within the range of [80%, 120%], and on-peak prices always within the range of [120%, 200%]. Electricity price forecasts are generated by drawing uniformly from the discrete set of possible scaled price levels P = {0%, 5%, ..., 200%} which intersect with the clipped price regions. The demand and electricity price forecast models are used in the case studies of the following section.

Market-Driven Scheduling Case Studies

The RHS strategy is evaluated in terms of its economic performance and computation time requirements using several market-driven scheduling case studies. Economically optimal plant schedules are calculated over a cumulative 72-hour duration, where each rolling-horizon execution considers a 24-hour scheduling horizon, a 4-hour periodic rescheduling interval, and control interval spacing of $\Delta t = 0.25$h. The economic performance of the

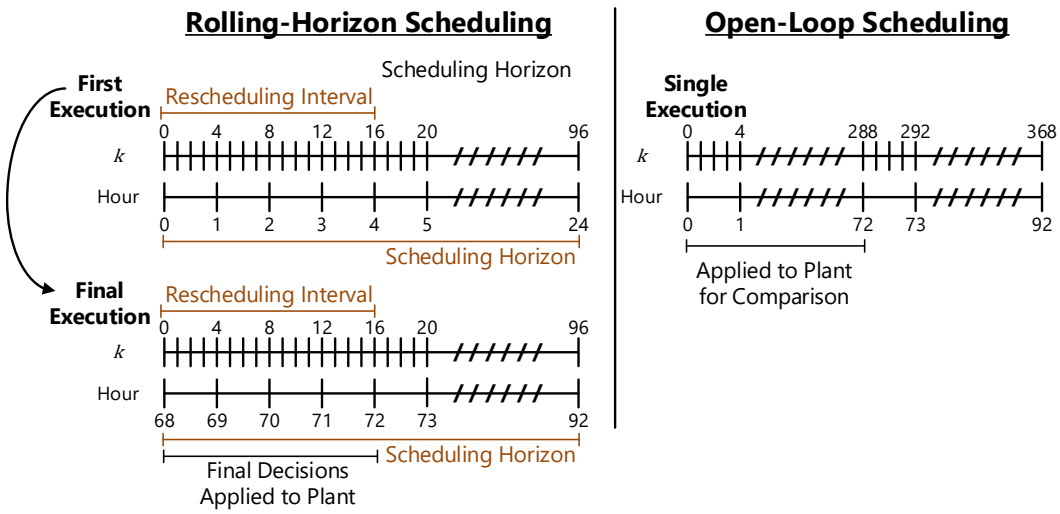

Figure 6. Discretization of RHS and OLS strategies

RHS strategy is compared to two open-loop scheduling (OLS) solutions for benchmarking. The first benchmark solution assumes perfect forecasts (PF) of market conditions across the 72-hour scheduling duration to provide an idealized upper bound on economic performance. The second OLS benchmark assumes limited knowledge of only the 24-hour initial forecast (IF), where mid-range values of demand and electricity prices are assumed for the remainder of the scheduling duration. To ensure consistency between the RHS and OLS solutions, the OLS solutions extend past the three-day scheduling window by the same duration as the final rolling-horizon execution; however, only the schedule from the first 72 hours is implemented to the plant for comparison. **Figure 6** provides a visual summary of the timing and discretization of the problem. The individual impacts of uncertainty in product demand and electricity price are evaluated using the four case studies outlined in **Table 2**. Due to the sampling method used when simulating market evolution and selecting forecasts, several unique subcases can be generated from the same initial forecast conditions (i.e. initial realized demand value, initial realized electricity price value, and initial R1-R5 region placement) for each of the DF1, DF2, EF1, and EF2 models. Therefore, each case study is repeated five times and average profits corresponding to the scheduling strategies, represented as percentages relative to profits obtained under the PF scenario, as well as the computation time required for RHS are reported in **Table 3**. Cases were executed on a Windows 10 Pro system (Intel Core i7-9700 CPU, 8 cores, 512GB SSD) using GAMSPy (v1.5.1) and the CONOPT (v4.36) solver within Python (v3.11.3).

Table 2: Summary of case studies investigated.

Case	Demand	Electricity Price
C1	Forecast (DF1)	Known
C2	Forecast (DF2)	Known
C3	Known	Forecast (EF1)
C4	Known	Forecast (EF2)

Table 3: Case results for average profit and solution time.

Case	Strategy	Avg Profit (% of PF Profits)	Avg Execution Time(s)
C1	IF	89.88	—
	RHS	99.75	35.86
	PF	100.00	—
C2	IF	87.37	—
	RHS	99.74	35.97
	PF	100.00	—
C3	IF	99.91	—
	RHS	99.94	37.63
	PF	100.00	—
C4	IF	99.91	—
	RHS	99.99	41.25
	PF	100.00	—

All scheduling strategies followed a store-then-utilize policy in their solutions, where storage was accumulated during times of low electricity prices and used during times of high electricity prices. The RHS strategy consistently demonstrated average economic performance on par with the PF solution. For example, the storage volume profile from the third realized profile in C2 is shown in **Figure 7**, showing storage usage trends which closely

align with the PF with minor discrepancies associated with overestimated or underestimated demands.

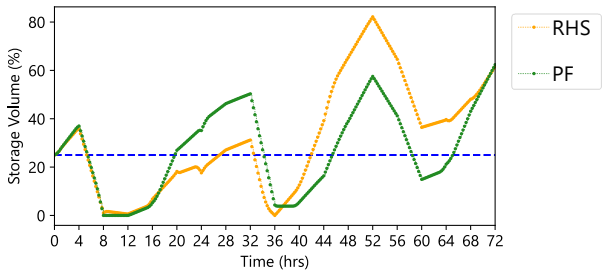

Figure 7. Storage volume profiles for the third realized profile in C2. Initial storage volume is indicated with a horizontal dashed blue line.

Interestingly, the average economic performance of the RHS strategy relative to the PF solution did not change significantly when forecast models with a greater degree of uncertainty were considered. This suggests that adequate response to changing future conditions could be achieved prior to suffering economic consequences. Furthermore, the RHS strategy significantly outperforms the IF solution in cases where demand forecast uncertainty is considered. Under demand uncertainty, the IF solution anticipated mid-range demands after the initial 24-hour forecast which resulted in periods of overestimation, underestimation, and suboptimal inventory usage, requiring the costly purchase of external GN_2 to fulfill demand on some occasions. By contrast, the RHS strategy offers only marginal improvements when electricity prices are uncertain and demand is known a priori. In these cases, no large profit losses related to unmet demand were observed and while the initial electricity price forecast was imperfect, it still followed the same periodic mid-peak, on-peak, and off-peak transitions as the realized electricity prices. This finding suggests that when plant demand and price seasonality are consistent and predictable, good scheduling decisions can be made even when specific price values are imperfectly forecasted.

Commonly activated constraints across the RHS solutions included:

- The GN_2 product purity constraint and reboiler low-level constraints, activated by increased LN_2 production and reduced column reflux rates.
- Maximum liquid feed rate and minimum vapor feed rate, activated to raise liquid content in the column for maintaining product purity.
- The compressor surge constraint, activated in response to reduced compressor air feed during periods of underproduction.
- Move restriction constraints placed on surrogate model inputs.
- Product demand satisfaction constraints.

Overall, results show that RHS combined with plant inventory are effective strategies for handling uncertainties in market forecasts under the case studies considered in this work. Notably, use of the LV-SM (implemented using a 15-minute prediction interval) offers solution times between 35-45 seconds per RHS execution, demonstrating high computational performance well within a timeframe suitable for real-time implementation considering the four-hour rescheduling interval used.

CONCLUSIONS

This work implemented a rolling-horizon scheduling strategy for the real-time market-driven operation of an ASU. Several case studies involving assumed forecast models of plant demand and electricity prices compared the economic performance of the RHS strategy to two benchmark solutions — one based on a perfect forecast of future market conditions, and one based only on information available at the time of the initial forecast. On average, the RHS strategy demonstrated economic performance on par with the perfect forecast solution and showed significant advantages over the initial forecast solution in cases considering demand uncertainty. When electricity price seasonality and plant demand were consistent and predictable over the scheduling window, RHS offered no significant advantages. Overall, these findings suggest that periodic rescheduling combined with plant inventory provides an effective means of responding to the assumed market fluctuations considered in this work. Lastly, use of the LV-SM (implemented using a 15-minute prediction interval) within the RHS strategy offered solution times between 35-45 seconds per RHS execution, demonstrating high computational performance suitable for real-time implementation.

With the rolling-horizon framework in place, future work may consider expanded case studies which identify areas where explicit consideration of uncertainty in the problem formulation (e.g., using a stochastic rolling-horizon scheduling strategy) provides improved economic performance and operational robustness over deterministic solutions. These case studies may also include forecast models based in historical data with additional sources of uncertainty, particularly those that may lead to infeasible operation if not accounted for. To mitigate the additional computational complexity introduced by this extension, techniques such as decomposition, scenario reduction, and parallelization can be explored.

ACKNOWLEDGEMENTS

Support from the Ontario Research Fund- Research Excellence (ORF-RE) Award RE09-058, and the McMaster Advanced Control Consortium, is gratefully acknowledged.

AUTHOR IDENTIFIERS

Author ORCIDs:
Swartz CLE: 0000-0003-3153-178X

REFERENCES

1. Pattison RC, Touretzky CR, Harjunkoski I, Baldea M. Moving horizon closed?loop production scheduling using dynamic process models. AIChE Journal 63:639-651 (2016). https://doi.org/10.1002/aic.15408
2. Kelley MT, Baldick R, Baldea M. An empirical study of moving horizon closed-loop demand response scheduling. Journal of Process Control 92:137-148 (2020). https://doi.org/10.1016/j.jprocont.2020.05.011
3. Schäfer P, Caspari A, Mhamdi A, Mitsos A. Economic nonlinear model predictive control using hybrid mechanistic data-driven models for optimal operation in real-time electricity markets: in-silico application to air separation processes. Journal of Process Control 84:171-181 (2019). https://doi.org/10.1016/j.jprocont.2019.10.008
4. Caspari A, Offermanns C, Schäfer P, Mhamdi A, Mitsos A. A flexible air separation process: 2. optimal operation using economic model predictive control. AIChE Journal 65: (2019). https://doi.org/10.1002/aic.16721
5. Zhang Q, Morari MF, Grossmann IE, Sundaramoorthy A, Pinto JM. An adjustable robust optimization approach to scheduling of continuous industrial processes providing interruptible load. Computers & Chemical Engineering 86:106-119 (2016). https://doi.org/10.1016/j.compchemeng.2015.12.018
6. Kelley MT, Baldick R, Baldea M. Demand response scheduling under uncertainty: chance?constrained framework and application to an air separation unit. AIChE Journal 66: (2020). https://doi.org/10.1002/aic.16273
7. Zhang Q, Cremer JL, Grossmann IE, Sundaramoorthy A, Pinto JM. Risk-based integrated production scheduling and electricity procurement for continuous power-intensive processes. Computers & Chemical Engineering 86:90-105 (2016). https://doi.org/10.1016/j.compchemeng.2015.12.015
8. Cao Y, Swartz CLE, Baldea M, Blouin S. Optimization-based assessment of design limitations to air separation plant agility in demand response scenarios. Journal of Process Control 33:37-48 (2015). https://doi.org/10.1016/j.jprocont.2015.05.002
9. McKenzie K, Swartz CLE, Corbett B. Dynamic surrogate modeling using latent variable methods and neural networks for market-driven operation of an air separation unit. Ind. Eng. Chem. Res. 65:584-599 (2025). https://doi.org/10.1021/acs.iecr.5c02735
10. Nomikos P, MacGregor JF. Monitoring batch processes using multiway principal component analysis. AIChE Journal 40:1361-1375 (2004). https://doi.org/10.1002/aic.690400809
11. Fletcher R, Leyffer S, Ralph D, Scholtes S. Local convergence of SQP methods for mathematical programs with equilibrium constraints. SIAM J. Optim. 17:259-286 (2006). https://doi.org/10.1137/s1052623402407382
12. Ralph D, Wright SJ. Some properties of regularization and penalization schemes for mpecs. Optimization Methods and Software 19:527-556 (2004). https://doi.org/10.1080/10556780410001709439
13. Mathur P, Swartz CLE, Zyngier D, Welt F. Robust online scheduling for optimal short-term operation of cascaded hydropower systems under uncertainty. Journal of Process Control 98:52-65 (2021). https://doi.org/10.1016/j.jprocont.2020.12.005

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Deep Kernel Learning with Kolmogorov–Arnold Networks for Bayesian Optimization

Zhongtao Shang[a], Zhihong Yuan[b], Lifeng Zhang[c]*, Yiyang Dai[a]*

[a] School of Chemical Engineering, Sichuan University, Chengdu, 610065, P. R. China.
[b] Department of Chemical Engineering, Tsinghua University, Beijing, 100084, China.
[c] Department of Chemical Engineering, The Sargent Centre for Process Systems Engineering, Imperial College London, London, SW7 2AZ, United Kingdom
* Corresponding Author: lifeng.zhang@imperial.ac.uk, daiyy@scu.edu.cn

ABSTRACT

Deep Kernel Learning (DKL) has emerged as a powerful framework for Bayesian Optimization (BO), via combining expressive representation learning models with typical Gaussian Processes (GPs) surrogate models. However, conventional DKL typically relies on weight-based feature extractors (e.g., multilayer perceptrons (MLPs)), which often lack interpretability and may suffer from overfitting under data scarcity or training instability, potentially leading to a degraded uncertainty quantification in GP models. Grounded in the Kolmogorov-Arnold representation theorem, this paper proposes a novel DKL-KAN framework that employs Kolmogorov-Arnold Networks (KANs) as adaptive feature extractors, formulating a DKL-KAN surrogate model. Unlike MLPs, the KANs learn data-driven univariate functions, yielding more sample-efficient and stable representations for regression under limited data regimes. Followed by the GP, the DKL-KAN facilitates end-to-end learning of expressive latent representations while maintaining robust uncertainty quantification. The proposed DKL-KAN framework is further validated with the Williams-Otto (W-O) process benchmarks. Compared to classical GPs and DKL-MLPs, the DKL-KANs consistently exhibit superior predictive accuracy and uncertainty calibration even under high-dimensional and scarce dataset. The results show that the KANs generate more structured and informative latent embeddings, thus enhancing the ability to capture complex physical nonlinearities in sparsely sampled edge regions of the design space where traditional models often falter. Moreover, the DKL-KAN is embedded into the BO and yields the efficient optimization performance on the W–O process benchmark, achieving faster early-stage improvement and converging to higher objective values. These results demonstrate the potential of DKL-KAN surrogates for accelerating BO in chemical engineering processes.

Keywords: Deep Kernel Learning, Kolmogorov-Arnold Network, Bayesian Optimization, Process Optimization

1. INTRODUCTION

Black-box optimization (BBO) has been a fundamental problem in many scientific and industrial applications, including process optimization[1], experimental design[2], drug discovery[3] and model hyperparameter tuning[4], where the objective function is expensive to evaluate and lacks accessible analytical expressions or gradient information. In the past decades, Bayesian Optimization (BO) has emerged as a premier framework in BBO, renowned for its ability to search global optima by balancing exploration and exploitation[5].

In BO, the crucial step is to develop a surrogate model to offer not only the mean predictions but also the uncertainty quantification for evaluating the acquisition function. Therefore, the Gaussian Process (GP) has been widely used in BO owing to its principled probabilistic formulation and kernel-based covariance structure. However, when applied to high-dimensional datasets, the GPs often suffer from significant computational costs for constructing the kernel matrix and kernel degeneracy, which may degrade the model performance. To resolve these limitations, Deep Kernel Learning (DKL) has been proposed to automatically learn the low-dimensional

latent embeddings and then feeding them to the GP models[6–8]. In practice, the MLPs are commonly used in most existing DKL approaches for their ability to learn flexible nonlinear feature representations from high-dimensional inputs. Singh et al.[9] proposed MLP-based DKL surrogate to convert multiple molecular representations to the latent space embedding and achieve accurate reaction yield predictions. Andrea et al.[10] developed a closed-loop DKL-based BO framework for porous lithium-ion cathode microstructure design, where a Generative Adversarial Network(GAN) generator is used to define the latent design space. Nevertheless, the limited interpretability, susceptibility to overfitting under scarce data, and poor training stability of MLPs often restrict their applicability in complex process optimization problems. Furthermore, according to prior studies, an inappropriate and over-parameterized feature extractor in DKL can lead to unreliable uncertainty estimation in BO[7].

Inspired by the Kolmogorov–Arnold representation theorem, Kolmogorov–Arnold Networks (KANs)[11], have recently emerged as a promising alternative approach to traditional MLPs. Unlike MLPs that rely on predefined activation functions to capture the nonlinearity, KANs employ learnable univariate functions as adaptive activations, enabling them to capture complex nonlinear relationships with enhanced interpretability and data efficiency. Numerous studies have demonstrated the great potential of KANs in modeling fluid dynamics, as well as in executing time-series classification and differential equation solvers[12, 13]. By providing more interpretable functional representations, KANs have shown more attraction for scientific discovery and mechanism-oriented analysis compared to MLPs[14].

Motivated by these observations, this work proposes a DKL-KAN model, in which a KAN serves as the feature extractor and is followed by a GP model within the DKL framework. The DKL-KAN model simultaneously learns the expressive latent representations and quantify the uncertainty, revealing good capture of data distribution and accurate predictions. The proposed DKL-KAN model is further integrated into the BO framework and applied to the benchmark Williams–Otto (W-O) process optimization problem. Compared to standard GPs (std-GPs) and DKL-MLP, the DKL-KAN model demonstrates a competitive performance, maintaining superior predictive accuracy, uncertainty calibration, and optimization efficiency even under high-dimensional and data-scarce regimes. Moreover, the DKL-KAN framework highlights improved interpretability and structural insight into the underlying process. The proposed DKL-KAN model offers new insights for addressing complex chemical process optimization problems.

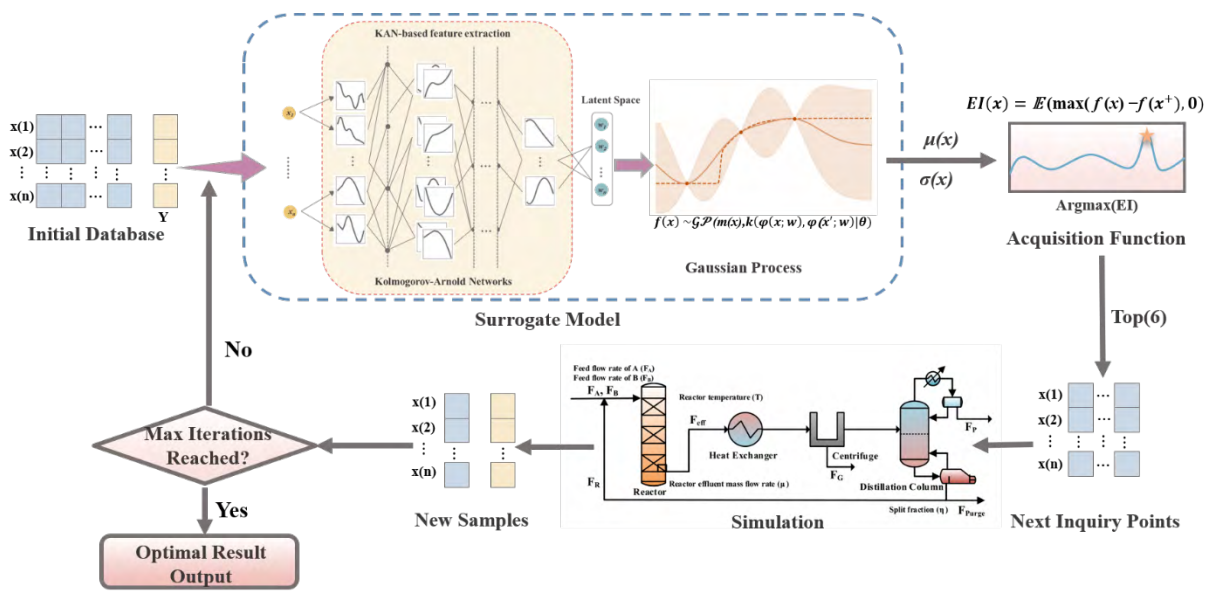

Figure 1. Schematic illustration of DKL-KAN framework for BO.

2. METHODLOGY

2.1 Optimization framework

In this work, a complex process optimization problem is considered to obtain the optimal operations under limited evaluation budget, which can be formulated as follows:

$$\max\ y = f(x) \tag{1}$$

$$s.t.\ x^{lb} \le x \le x^{ub}$$

$$x \subseteq R^n$$

Where y denotes a scalar metric to be optimized, such as

annual profit or process yields. x represents continuous decision variables, such as temperature, flow rate, and pressure. To solve the problems, the overall optimization framework which combines DKL-KAN and BO can be illustrated in Figure 1. The framework is structured as follows:

1. A small-batch dataset is generated with Sobol' sequence[15] from the design space and evaluated to get the scaler metric via process simulations.

2.With the existing dataset, the DKL-KAN surrogate model is developed by combining a KAN model as a feature extractor and a GP model which maps the normalized latent embeddings to predictive means and covariances in an end-to-end training manner.

3.Using the DKL–KAN surrogate, the expected improvement (EI) values, also known as the acquisition function, are evaluated over the candidate set. A next batch of sampling points are determined by maximizing the EI.

4.The suggested candidates from EI are subsequently evaluated in the simulation environment or retrieved from the dataset to obtain new observations, which are appended to the training set for updating the surrogate model in the next iteration.

5.The procedure terminates when the maximum number of validation iterations is reached, and the optimal solution identified so far is reported; otherwise, the workflow returns to Step 2 to update the surrogate model and continues iteratively.

2.2 Kolmogorov-Arnold Networks

KANs are theoretically grounded in the Kolmogorov-Arnold representation theorem[11] which asserts that any multivariate continuous function $f:[0,1]^n \to \mathbb{R}$ can be decomposed as a combination of univariate functions:

$$f(\mathbf{x}) = f(x_1, \dots, x_n) = \sum_{q=1}^{2n+1} \phi_q \left(\sum_{p=1}^{n} \varphi_{q,p}(x_p) \right) \tag{2}$$

where $\varphi_{q,p}$ and ϕ_q denote univariate continuous functions. While the classical theorem restricts the architecture to a fixed depth, modern KANs generalize this into a deep stacked structure. For a network with $|L|$ layers, the forward pass is represented as a composition of functional layers:

$$\mathrm{KAN}(\mathbf{x}) = \left(\phi_{|L|-1} \circ \phi_{|L|-2} \circ \dots \circ \phi_1\right)\mathbf{x} \tag{3}$$

where L refers the set of layers and each layer ϕ_l consists of a set of trainable 1-D activation functions $\{\varphi_{l,j}\}$. In practice, each univariate activation $\varphi_{l,j}(x)$ is parameterized as a combination of basis function $b_{l,j}(x)$ and a B-spline function:

$$\varphi_{l,j}(x) = \omega_{l,j} \left[b_{l,j}(x) + \sum_m c_{l,j,m} B_{l,j,m}(x) \right] \ \forall l \in L, j \in J^i \tag{4}$$

Where $m = 1, \dots, M$ indexs B-spline basis functions, and $B_{l,j,m}(x)$ denotes the mth B-spline basis defined on the corresponding grid. The spline coefficients $c_{l,j,m}$ and grid configurations in KANs are fully learnable, allowing the network to adaptively refine its function approximation capabilities.

2.3 DKL-KAN Surrogate Model

The DKL-KAN surrogate model for BO is formulated with KAN and GP where the GP converts the latent embeddings to predictive means and uncertainty quantification. Generally, the GP regression is defined as follows to obtain the black-box function $f(\mathbf{x})$[16]:

$$f(\mathbf{x}) \sim \mathcal{GP}\left(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')\right) \tag{5}$$

where $m(\mathbf{x})$ and $k(\mathbf{x}, \mathbf{x}')$ denote the mean function and kernel function. Given observations $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, GP posterior yields a predictive distribution at any candidate $\mathbf{x}_i$, whose posterior mean $\eta(\mathbf{x}|\mathcal{D})$ serves as the prediction and posterior variance $\sigma^2(\mathbf{x}|\mathcal{D})$ quantifies epistemic uncertainty. This probabilistic output is directly leveraged by acquisition functions to guide sampling in BO.

Besides, conventional GP surrogates can be limited in practical applications. Exact GP inference scales cubically with the number of evaluations due to kernel matrix computation. Moreover, the representational capacity of GP is strongly determined by the choice of kernel, standard stationary kernels may struggle to capture complex nonstationarity, sharp local variations, or high-dimensional structure without substantial kernel engineering. DKL mitigates these drawbacks by learning a task-adaptive representation prior to GP modeling. Given a KAN feature map $g(\mathbf{x}|\beta)$arameterized by β, and $k(\mathbf{x}, \mathbf{x}'|\alpha)$ representing the standard kernel parameterized by α, the kernel function can be formulated as follows[17]:

$$k_{DKL-KAN}(\mathbf{x}, \mathbf{x}') = k_{base}(g(\mathbf{x}_i;\beta), g(\mathbf{x}_i';\beta)|\alpha, \beta) \tag{6}$$

This construction preserves GP-style uncertainty quantification while improving expressiveness through data-driven feature learning, enhancing the model capacity for real world applications.

Following the DKL-KAN formulation introduced above, the surrogate is trained by jointly learning β and α via maximizing the log marginal likelihood $\mathcal{L}$ [18]. Let $\mathbf{y} \in \mathbb{R}^n$ denote the observations at inputs $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$, and let $\mathbf{K} \in \mathbb{R}^{n \times n}$ be the covariance matrix with entries $[\mathbf{K}]_{ij} = k_{\mathrm{DKL-KAN}}(\mathbf{x}_i, \mathbf{x}_j)$. Assuming Gaussian noise with variance σ_n^2, the $\mathcal{L}$ is given by:

$$\mathcal{L}(\alpha, \beta) = -\frac{1}{2}\mathbf{y}^\top(\mathbf{K} + \sigma_n^2\mathbf{I})^{-1}\mathbf{y} - \frac{1}{2}\log|\mathbf{K} + \sigma_n^2\mathbf{I}| - \frac{n}{2}\log(2\pi) \tag{7}$$

Then the surrogate model is obtained via minimizing the negative log marginal likelihood (NLL):

$$\alpha^{\star}, \beta^{\star} = \arg\min_{\alpha,\beta} -\mathcal{L} \tag{8}$$

Compared to the commonly used DKL-MLP surrogate which is usually an opaque, black-box feature extractor, the proposed DKL-KAN provides a more structured and inspectable embedding with improved data efficiency and uncertainty calibration, particularly for high-dimensional, data-scarce chemical process optimization.

2.4 Acquisition Function

A variety of acquisition functions have been proposed for BO, including probability of improvement (PI), EI, and upper confidence bound (UCB), which differ in how to trade off exploration and exploitation[19]. Among them, EI is particularly attractive because it directly quantifies the expected gain over the current best observation[20]. Given the surrogate predictive distribution $y \sim \mathcal{N}(\mu(\mathbf{x}), \sigma^2(\mathbf{x}))$, EI at a new point $\mathbf{x}$ can be written in closed-form as:

$$\mathrm{EI}(\mathbf{x}) = \left(\mu(\mathbf{x}) - f(\mathbf{x}^*)\right)\Phi(Z) + \sigma(\mathbf{x})\phi(Z) \tag{9}$$

where $f(\mathbf{x}^*)$ denotes the best objective value observed so far, Φ and ϕ denote the cumulative distribution function (CDF) and the probability density function (PDF) of the standard normal distribution, respectively. Z is a random variable defined as:

$$Z(\mathbf{x}) = \frac{\mu(\mathbf{x}) - f(\mathbf{x}^*)}{\sigma(\mathbf{x})} \tag{10}$$

In the proposed framework, $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ are provided from the DKL-KAN. At each iteration, EI is evaluated over the possible candidate set and a batch of the six points with the largest EI values is selected for querying. The resulting objective evaluations are then added to the training set to update the surrogate, and the procedure is repeated until the maximum iterations is reached.

3.CASE STUDY

3.1 Problem description

The proposed method is then applied to the W-O process optimization problem as case study. For a brief introduction, the W-O process is consist of a continuously stirred tank reactor (CSTR), a decanter, a distillation column (separator) and a splitter[21], as shown in Figure 2. The decision variables and respective bounds for W-O process can be defined in Table 1. F_A and F_B denote the feed flow rates of reactants A and B, T is the operating temperature of reactor, μ represents the total mass flow rate of the reactor effluent, and η specifies the split fraction.

Under a series of reactions, the objective for W-O process is to maximize the Return On Investment (ROI) as follows:

$$ROI = \frac{Annual\ Net\ Profit}{Total\ Capital\ Investment} \tag{11}$$

For the details of W-O process, including reaction system, mass balance constraints and reaction kinetic equations, the readers are referred to the relevant literature[21].

All computations are conducted on a workstation equipped with an Intel(R) Core(TM) i9-14900KF CPU (3.20 GHz) and 64 GB RAM. The W-O is simulated via *Python* and the KAN is constructed using the *pykan*, while GP is constructed with *GPyTorch*.

Table 1. Decision variable bounds of the optimization process.

Variables	Unit	Bounds
F_A	klb/h	[2.0, 40]
F_B	klb/h	[2.0, 40]
T	K	[520, 700]
μ	klb/h	[10, 200]
η	-	[0, 0.9]

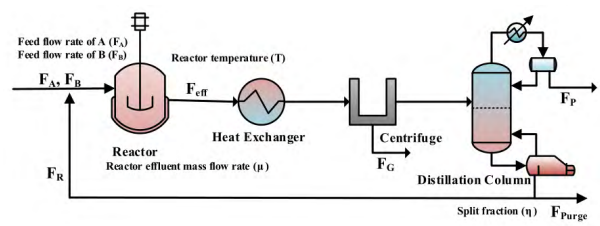

Figure 2. Schematic illustration of W-O process.

3.2 Performance of DKL-KAN

To assess the predictive capability of the proposed DKL–KAN surrogate, multiple datasets with different sample sizes are generated. For each dataset, samples were randomly split into training and test sets with an 8:2 ratio. Five models are developed for comparison, namely, std-GP, DKL–MLP, DKL–KAN, std-MLP, and std-KAN, where *std* denotes the standard model. Hyperparameter tuning is conducted by Optuna[22] to ensure a fair comparison under consistent hyperparameter optimization. The network architectures are eventually fixed as KAN: [5×4×2] and MLP: [5×64×32×2].

The R-squared (R^2) values with the dataset of 500 points are summarized in Figure 3 (a). Obviously, DKL-based surrogates improve generalization over std-GP, and DKL–KAN delivers superior performance. Worse still, std-GP is highly kernel-sensitive with R^2 ranging from 0.609±0.083(linear) to 0.794±0.068 (Radial Basis Function, RBF), indicating that the kernels cannot adequately capture the complex underlying feature structure. DKL–KAN maintains high R^2 across different kernels and data sizes, showing consistent performance from 0.919± 0.045, 0.935±0.036 and 0.916±0.056 for linear, RBF and

Matérn-2.5 kernels, respectively. This indicates that introducing KAN for DKL can significantly improve the model capacity regardless the kernels. In contrast, the baselines (std-GP, std-MLP and std-KAN) exhibit larger performance gaps and variability, whereas std-KAN attains relatively high R^2 with 0.926±0.105, but lacks the predictive uncertainty required. This superiority is further reflected in the root mean squared error (RMSE) metrics trends across dataset sizes (Figure 3 (b)), DKL–KAN yields the lowest test RMSE, with its advantage being most evident in the low-data regime and persisting as the training set grows. This enhanced performance is attributed to the integration of the KANs within the deep kernel architecture, which enables more expressive nonlinear feature mapping and superior uncertainty quantification compared to fixed-basis surrogates.

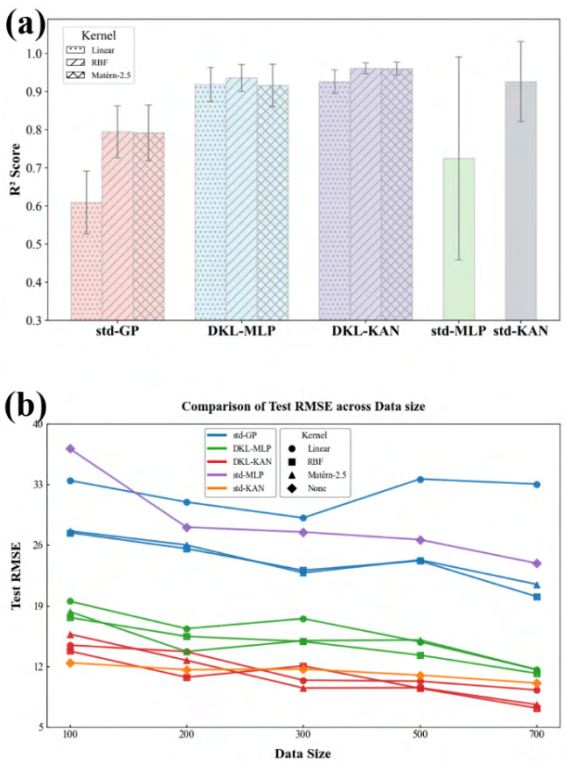

Figure 3. (a) R^2values at dataset of 500 points. (b) Test set RMSEs on different sample sizes.

3.3 Latent space embedding analysis

To gain deeper insights into the behavior of the models, the latent space embeddings of DKL–KAN model is analyzed and compared with the DKL–MLP and std-GP models. For the DKL-driven models, the two-dimensional embeddings from the final feature mapping layer are directly extracted, while for the std-GP, UMAP dimensionality reduction is applied[23]. Subsequently, k-means clustering[24] is performed to obtain four distinct clusters. As shown in Figure 4, the introduction of DKL significantly enhances the distribution of the latent space compared to the std-GP baseline, with clusters becoming more compact and well-separated[9]. The ROI color map reveals a clearer spatial distribution, indicating improvement by introducing DKL-driven surrogates. A closer examination of the cluster characteristics, the highest-ROI cluster is associated with higher reactor temperatures, larger outlet flow rates, and smaller split fractions, while lower-ROI clusters correspond to different combinations of these decision variables.

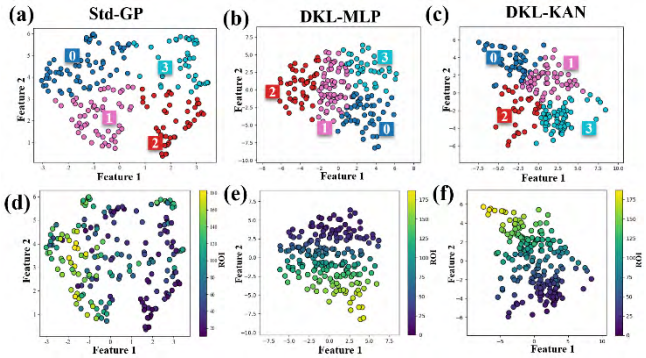

Figure 4. Analysis of the latent space of the (a) std-GP, (b) DKL-MLP and (c) DKL-KAN. The colors represent different clusters identified using k-means clustering. The distribution of ROI in each cluster formed from the (d) std-GP, (e) DKL-MLP and (f) DKL-KAN.

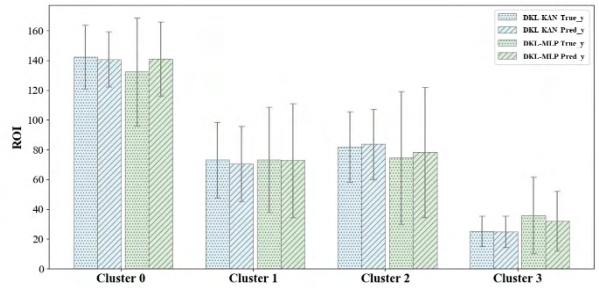

Figure 5. Comparison of ROI Prediction Consistency Across Clusters for DKL–MLP, DKL–KAN.

Figure 5 further illustrates the model predictions within each cluster. DKL–KAN demonstrates consistent ROI predictions across all clusters, performing well even in both high- and low-ROI regions, thus highlighting its ability to handle complex nonlinearities effectively. In contrast, DKL-MLP shows more significant fluctuations in ROI predictions across clusters, alternating between overestimations and underestimations, especially in the data-sparse, high-nonlinearity extreme regions, reflecting a fundamental limitation. While the integration of GP priors offers uncertainty quantification, the MLP feature extractor struggles to capture the highly nonlinear, multimodal nature of complex objective functions. The inherent smoothness of the MLP architecture, combined with its inability to adapt to sharp variations or multiple peaks,

leads to suboptimal performance, especially in data-scarce regions with high-dimensional nonlinearity. These results underscore the ability of DKL–KAN to effectively capture the underlying structure of the W–O process, enhancing both the separability of the learned feature embeddings and the accuracy of the predictions.

3.4 Optimisation performance

The DKL-KAN is further integrated with BO as the surrogate model for W-O process to validate the optimization efficiency[25]. The total optimization iteration is set to 20, with performance reported as the mean and standard deviation across five independent trials with randomized initializations, illustrated in Figure 6. Computational results from std-GP and DKL-MLP are also provided for comparison. Both DKL-MLP and DKL-KAN surrogates outperform the std-GP in terms of optimization efficiency, consistent with prior work showing that deep kernel learning can improve surrogate accuracy and thus accelerate BO convergence by capturing informative feature representations than shallow kernels alone. Notably, the proposed DKL–KAN model exhibits the most efficient performance, achieving faster improvement in the early iterations and converging to higher objective values. This superior performance is primarily attributable to the DKL–KAN surrogate, whose KAN feature extractor learns more structrued latent embeddings than MLPs and thus the GP kernel can better distinguish the datapoints and captures the highly nonlinearity. Overall, these results demonstrate that the proposed DKL–KAN model provides both excellent predictive accuracy and reliable uncertainty quantification, making it a promising alternative method for process optimization tasks where high-dimensional, nonlinear relationships are prevalent.

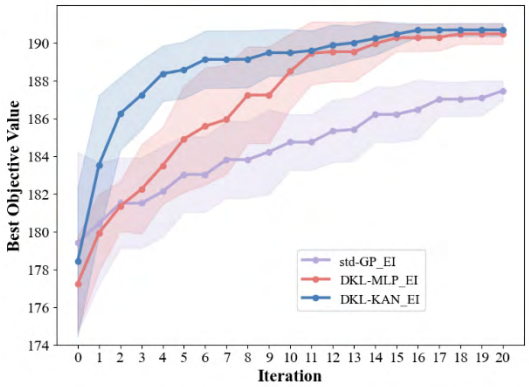

Figure 6. Comparison of optimization performance under EI acquisition function with std-GP, DKL-MLP and DKL-KAN surrogate models.

4. CONCLUSIONS AND FUTURE WORK

This work develops a novel DKL surrogate model leveraging the KAN as an adaptive feature extractor. Systematically validated on the W–O process benchmark, the proposed DKL-KAN model demonstrates superior predictive performance and more reliable uncertainty quantification compared to standard GPs and conventional DKL-MLP. Latent space analysis further elucidates that the KAN facilitates a more structured embedding distribution, maintaining high generalization and accurate uncertainty characterization even in sparse boundary regions of the sample space. When integrated with BO, the approach identifies superior optimum with significantly accelerated convergence in early iterations. This can be attributed to the more informative embeddings from KAN which better represent the highly nonlinear and strongly coupled W–O response surface. These findings underscore the significant potential of DKL-KAN surrogates as a promising method for complex chemical engineering process optimization. Future work can be done to incorporate advanced algorithms within the latent space to enhance the robustness and efficiency for global optimization exploration.

ACKNOWLEDGEMENTS

The authors are grateful for the support of the National Natural Science Foundation of China (22538010).

REFERENCES

1. Zhou J, Shi T, Ren J, He C. Accelerating operation optimization of complex chemical processes: a novel framework integrating artificial neural network and mixed-integer linear programming. Chemical Engineering Journal 481:148421 (2024). https://doi.org/10.1016/j.cej.2023.148421
2. Shields BJ, Stevens J, Li J, Parasram M, Damani F, Alvarado JIM, Janey JM, Adams RP, Doyle AG. Bayesian reaction optimization as a tool for chemical synthesis. Nature 590:89-96 (2021). https://doi.org/10.1038/s41586-021-03213-y
3. Li H, Hao J, Qiao S. Ai?driven electrolyte additive selection to boost aqueous zn?ion batteries stability. Advanced Materials 36: (2024). https://doi.org/10.1002/adma.202411991
4. S. Falkner, A. Klein, F. Hutter, BOHB: Robust and Efficient Hyperparameter Optimization at Scale, (2018). https://doi.org/10.48550/arXiv.1807.01774
5. Winz J, Fromme F, Engell S. Bayesian optimization of gray-box process models using a modified upper confidence bound acquisition function. Computers & Chemical Engineering 194:108976 (2025). https://doi.org/10.1016/j.compchemeng.2024.108976
6. A.G. Wilson, Z. Hu, R. Salakhutdinov, E.P. Xing, Deep Kernel Learning, in: Proceedings of the 19th

International Conference on Artificial Intelligence and Statistics, PMLR, 2016: pp. 370–378. https://proceedings.mlr.press/v51/wilson16.html (accessed October 10, 2025).
7. I. Achituve, G. Chechik, E. Fetaya, Guided Deep Kernel Learning, (2023). https://doi.org/10.48550/arXiv.2302.09574
8. Botteghi N, Guo M, Brune C. Deep kernel learning of dynamical models from high-dimensional noisy data. Sci Rep 12: (2022). https://doi.org/10.1038/s41598-022-25362-4
9. Singh S, Hernández-Lobato JM. Deep kernel learning for reaction outcome prediction and optimization. Commun Chem 7: (2024). https://doi.org/10.1038/s42004-024-01219-x
10. Gayon-Lombardo A, del Rio-Chanona EA, Pino-Muñoz CA, Brandon NP. Deep kernel bayesian optimisation for closed-loop electrode microstructure design with user-defined properties. Energy and AI 22:100608 (2025). https://doi.org/10.1016/j.egyai.2025.100608
11. Z. Liu, Y. Wang, S. Vaidya, F. Ruehle, J. Halverson, M. Soljačić, T.Y. Hou, M. Tegmark, KAN: Kolmogorov-Arnold Networks, (2025). https://doi.org/10.48550/arXiv.2404.19756
12. S.A. Faroughi, F. Mostajeran, A.H. Mashhadzadeh, S. Faroughi, Scientific Machine Learning with Kolmogorov-Arnold Networks, (2025). https://doi.org/10.48550/arXiv.2507.22959
13. Wang Y, Sun J, Bai J, Anitescu C, Eshaghi MS, Zhuang X, Rabczuk T, Liu Y. Kolmogorov–arnold-informed neural network: a physics-informed deep learning framework for solving forward and inverse problems based on kolmogorov–arnold networks. Computer Methods in Applied Mechanics and Engineering 433:117518 (2025). https://doi.org/10.1016/j.cma.2024.117518
14. Zeng Y, Cao W, Zuo Y, Peng T, Hou Y, Miao L, Wang Z, Shi J. Accelerating the discovery of materials with expected thermal conductivity via a synergistic strategy of DFT and interpretable deep learning. Mater. Futures 4:045602 (2025). https://doi.org/10.1088/2752-5724/ae08d0
15. E. Atanassov, S. Ivanovska, On the Use of Sobol' Sequence for High Dimensional Simulation, in: D. Groen, C. de Mulatier, M. Paszynski, V.V. Krzhizhanovskaya, J.J. Dongarra, P.M.A. Sloot (Eds.), Computational Science – ICCS 2022, Springer International Publishing, Cham, 2022: pp. 646–652. https://doi.org/10.1007/978-3-031-08760-8_53
16. Richardson RR, Osborne MA, Howey DA. Gaussian process regression for forecasting battery state of health. Journal of Power Sources 357:209-219 (2017). https://doi.org/10.1016/j.jpowsour.2017.05.004
17. S.W. Ober, C.E. Rasmussen, M. van der Wilk, The Promises and Pitfalls of Deep Kernel Learning, (2021). https://doi.org/10.48550/arXiv.2102.12108
18. C.E. Rasmussen, C.K.I. Williams, Gaussian processes for machine learning, 3. print, MIT Press, Cambridge, Mass., 2008.
19. E. Brochu, V.M. Cora, N. de Freitas, A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning, (2010). https://doi.org/10.48550/arXiv.1012.2599
20. S. Ament, S. Daulton, D. Eriksson, M. Balandat, E. Bakshy, Unexpected Improvements to Expected Improvement for Bayesian Optimization, (2025). https://doi.org/10.48550/arXiv.2310.20708
21. X. Jin, K.A. High, Decision Making for a Sustainable Chemical Process, (n.d.).
22. T. Akiba, S. Sano, T. Yanase, T. Ohta, M. Koyama, Optuna: A Next-generation Hyperparameter Optimization Framework, (n.d.).
23. L. McInnes, J. Healy, J. Melville, UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction, (n.d.).
24. Ikotun AM, Ezugwu AE, Abualigah L, Abuhaija B, Heming J. K-means clustering algorithms: a comprehensive review, variants analysis, and advances in the era of big data. Information Sciences 622:178-210 (2023). https://doi.org/10.1016/j.ins.2022.11.139
25. L.N. Sridhar, Multiobjective Nonlinear Model Predictive Control of the Williams Otto Process, Chemical Engineering 1 (2024).

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Machine Learning and Adaptive Sampling Powered Feasible Path Algorithm for Black-box Optimization

Zixuan Zhang[ab], Xiaowei Song[ab], Jiaming Li[ab], Yujiao Zeng[a], Yaling Nie[a], Min Zhu[a], Dongyun Lu[a], Yibo Zhang[a], Xin Xiao[ad]*, Jie Li[cd§]

[a] Chinese Academy of Science, Institute of Process Engineering, Beijing, China
[b] University of Chinese Academy of Science, School of Chemical Engineering, Beijing, China
[c] The University of Manchester, Department of Chemical Engineering, Manchester, United Kingdom
[d] China Institute for Innovation and Development Strategy, Work Committee on Carbon Neutral Intelligent Industry Innovation, Beijing, China
* Corresponding Author: xxiao@ipe.ac.cn.
§ Corresponding Author: jie.li-2@manchester.ac.uk.

ABSTRACT

Black-box optimization (BBO) deals with problems involving functions that are either unknown, imprecise, or costly to evaluate. Current BBO methods encounter multiple challenges, such as high computational demands from excessive function evaluations, difficulties in handling complex constraints, lack of theoretical convergence guarantees, and unstable performance due to significant variations in solution quality. This work presents a machine learning-powered feasible path (MLFP) framework for general BBO problems involving complex constraints. An adaptive sampling strategy is first proposed to explore optimal regions and pre-filter potentially infeasible points, thereby reducing the number of evaluations. Machine learning algorithms are utilized to build surrogates for black-box functions. The feasible path algorithm is integrated to accelerate theoretical convergence by updating only independent variables instead of all variables. Computational experiments demonstrate that MLFP can rapidly and robustly converge near the KKT point, even when training surrogates with small datasets. Compared to state-of-the-art BBO algorithms, MLFP stably delivers equivalent or superior solutions with fewer evaluations across benchmark examples.

Keywords: Optimization, Machine Learning, Black-box, Adaptive Sampling, Feasible Path Algorithm, Surrogate Model

INTRODUCTION

Black-box optimization (BBO) refers to a class of optimization problems where the algebraic forms of the objective function and/or constraints are unknown, imprecise, or prohibitively expensive to compute. This characteristic is ubiquitous in engineering and scientific fields, particularly in chemical engineering, where many complex systems cannot be fully captured by mechanistic models. For instance, multi-step catalytic reaction networks, molecular structure–activity relationships (SARs) linking polymer architecture to material performance, and high-fidelity simulations (e.g., density functional theory calculations, computational fluid dynamics for reactor design, and rigorous distillation modeling) all exhibit black-box properties [1-5]. These systems are foundational to chemical design and manufacturing, from catalyst development and material synthesis to plant-wide operation tuning, and optimizing them relies heavily on specialized BBO approaches. Advancing BBO methods is therefore critical to reducing development cycles, cutting costs, and unlocking innovative engineering solutions.

Despite its importance, BBO faces inherent challenges rooted in the absence of gradient information for black-box functions, which prevents the use of classical gradient-based optimization. Existing methods can be broadly categorized into three classes, each with distinct limitations. Stochastic or heuristic algorithms (e.g., PSO, GA, SA) offer global search capabilities but lack theoretical convergence guarantees and require extensive function evaluations, making them impractical for expensive

systems. Deterministic partition methods (e.g., DIRECT, pattern search) avoid randomness but suffer from the curse of dimensionality and high computational costs. For example, a modified DIRECT algorithm for LNG process optimization required 20, 000 evaluations [6]. Surrogate model-based methods, the most widely adopted class, reduce black-box calls by using cheap approximations (e.g., neural networks, Gaussian processes) but introduce new issues: implicit filtering is limited to bound-constrained problems, ReLU NN-based frameworks become computationally intractable for complex architectures [7], ALAMO performs poorly for high-dimensional inputs [8], trust-region filter (TRF) methods suffer from initial value sensitivity [9], and Bayesian optimization (BO) struggles with complex equality constraints and exhibits significant result stochasticity (variations up to 300% across runs) [10]. Additionally, many surrogate-based methods overlook the computational burden of sampling and training.

To address these limitations, this study proposes a novel machine learning-powered feasible path (MLFP) framework for general BBO problems involving complex equality and inequality constraints. The framework integrates three core components to balance sample efficiency, approximation accuracy, and convergence speed. First, an adaptive sampling strategy uses support vector machines (SVM) to filter infeasible points and focus on potential optimal regions, ensuring informative samples even with limited data. Second, machine learning models (e.g., multi-layer perceptrons with Swish activation) are employed as surrogates, leveraging their strong nonlinear approximation capabilities and twice-differentiable properties to enable gradient-based optimization. Third, the feasible path algorithm is integrated to update only independent variables instead of the entire variable set, avoiding complex nonlinear terms from full-space formulations and accelerating convergence by retaining computation graphs for efficient derivative calculation.

This work validates MLFP on 2 industrial processes (Williams-Otto process and extractive distillation process). The results demonstrate that MLFP outperforms state-of-the-art algorithms by stably achieving equivalent or better solutions with fewer evaluations.

PROBLEM STATEMENT

The general form of a BBO problem to be solved in this work is presented below, which is a nonlinear programming (NLP) problem.

$$\begin{aligned} \min_{\mathbf{x}\in\mathbb{R}^n} \quad & f(\mathbf{x}) \\ \text{s.t.} \quad & \mathbf{h}(\mathbf{x}) = 0 \\ & \mathbf{g}(\mathbf{x}) \le 0 \\ & \mathbf{x}^{lb} \le \mathbf{x} \le \mathbf{x}^{ub} \end{aligned} \tag{1}$$

where the function $f: \mathbb{R}^n \to \mathbb{R}$ is the objective function, $\mathbf{h}: \mathbb{R}^n \to \mathbb{R}^{p_E}$ is the equality constraints and $\mathbf{g}: \mathbb{R}^n \to \mathbb{R}^{p_I}$ is the inequality constraints excluding the bounding constraints. p_E , and p_I represent the dimensions of equations and inequalities respectively; $\mathbf{x}$ represent a vector of variables; the superscript ub represents the upper bound, and lb represents the lower bound. In this setting, the black-box function may pertain to the objective function f as well as to the constraint functions $\mathbf{h}$ (equalities) and $\mathbf{g}$ (inequalities), depending on the problem structure.

Without loss of generality, we divide the equality constraints $\mathbf{h}(\mathbf{x}) = 0$ into $\mathbf{h_1}(\mathbf{x}) = 0$ and $\mathbf{h_2}(\mathbf{x}) = 0$ where $\mathbf{h_1}(\mathbf{x}) = 0$ containing all black-box constraints and $\mathbf{h_2}(\mathbf{x}) = 0$ without containing any black-box functions. If any inequality constraints $\mathbf{g_j}(\mathbf{x}) \le 0$ containing black-box functions, we can introduce auxiliary variables $\mathbf{v}$ and let $\mathbf{v} = \mathbf{g_j}(\mathbf{x})$. We then add $\mathbf{v} = \mathbf{g_j}(\mathbf{x})$ into $\mathbf{h_1}(\mathbf{x}) = 0$ and remove $\mathbf{g_j}(\mathbf{x}) \le 0$ from the inequality constraint sets $\mathbf{g}(\mathbf{x}) \le 0$. The auxiliary variable $\mathbf{v}$ is then added into variables $\mathbf{x}$. The original BBO problem (1) can be reformatted into problem (2) below,

$$\begin{aligned} \min_{\mathbf{x}\in\mathbb{R}^n} \quad & f(\mathbf{x}) \\ \text{s.t.} \quad & \mathbf{h_1}(\mathbf{x}) = 0 \\ & \mathbf{h_2}(\mathbf{x}) = 0 \\ & \mathbf{g}(\mathbf{x}) \le 0 \\ & \mathbf{x}^{lb} \le \mathbf{x} \le \mathbf{x}^{ub} \end{aligned} \tag{2}$$

where $\mathbf{h_1}: \mathbb{R}^n \to \mathbb{R}^{p_{E1}}$ and $\mathbf{h_2}: \mathbb{R}^n \to \mathbb{R}^{p_{E2}}$ are the equality constraints and $\mathbf{g}: \mathbb{R}^n \to \mathbb{R}^{p_I}$ is the inequality constraints excluding the bounding constraints. $p_{E1}, p_{E2}$and p_I represent the dimensions of equations and inequalities respectively. In this setting, the black-box function may pertain to the objective function f as well as to the constraint functions $\mathbf{h_1}$ (equalities). We assume that all functions (e.g., $f, \mathbf{h_2}$ and $\mathbf{g}$) that are not the black-box functions, are twice continuously differentiable. The black-box functions $\mathbf{h_1}$ or f are unknown and their derivatives are not available; however, we assume that they can be evaluated at a given point and approximated by a twice continuously differentiable function.

METHOD

As $f(\mathbf{x})$ and/or $\mathbf{h_1}(\mathbf{x}) = 0$ in problem (2) contain black-box functions that are unknown or expensive for evaluation, we need to develop surrogate models for these black-box functions. To do this, we can divide all variables in (2) into independent variables (denoted as $\mathbf{x}_I$) and dependent variables (denoted as $\mathbf{x}_D$) where $\mathbf{x}_D$ has the same dimension as $\mathbf{h_1}$. Then, problem (2) can be transformed into problem (3).

Input: Surrogate model $\boldsymbol{s}(\cdot)$
Input: $\boldsymbol{\lambda}_{qp}^0 \leftarrow 0$, $\boldsymbol{\mu}_{qp}^0 \leftarrow 0$, $\boldsymbol{\rho}^0 \leftarrow 0$, $\boldsymbol{\nu}^0 \leftarrow 0$
Input: $\eta \leftarrow 0.1$, $tol \leftarrow 10^{-6}$, line search multiple $\beta \leftarrow 0.618$, line search maximum iteration $n_{max} \leftarrow 10$
Input: Feasible initial point $\boldsymbol{x}^0$
Output: Solution $\boldsymbol{x}^*$

$k \leftarrow 0$, $acc_{inf} \leftarrow \infty$, $acc_{step} \leftarrow \infty$, $acc_{opt} \leftarrow \infty$
while $acc_{inf} > tol$ or ($acc_{step} > tol$ and $acc_{opt} > tol$) **do**
  $k \leftarrow k + 1$
  $\boldsymbol{y}^k \leftarrow \boldsymbol{s}(\boldsymbol{x}^k)$
  Calculate $\nabla_x \boldsymbol{y}^k$ and $\nabla_x^2 \boldsymbol{y}^k$ based on Section 3
  Calculate $\nabla_x f$, $\nabla_x \boldsymbol{g}^k$, $\nabla_x \boldsymbol{h}^k$ and $\boldsymbol{H}^k$
  if $\boldsymbol{H}^k$ is positive definite **then**
    $\boldsymbol{B}^k \leftarrow \boldsymbol{H}^k$
  else
    $\boldsymbol{B}^k \leftarrow \boldsymbol{H}^k + \boldsymbol{E}^k$
  end if
  Solve (QP) subproblem, then obtain $\boldsymbol{d}^k$, $\boldsymbol{\lambda}_{qp}^k$ and $\boldsymbol{\mu}_{qp}^k$
  $\boldsymbol{\rho}^k \leftarrow \max(|\boldsymbol{\lambda}_{qp}^k|, \frac{|\boldsymbol{\lambda}_{qp}^k| + \boldsymbol{\rho}^{k-1}}{2})$
  $\boldsymbol{\nu}^k \leftarrow \max(|\boldsymbol{\mu}_{qp}^k|, \frac{|\boldsymbol{\mu}_{qp}^k| + \boldsymbol{\nu}^{k-1}}{2})$
  Calculate $\phi_1(\boldsymbol{x}^k, \boldsymbol{\lambda}_{qp}^k, \boldsymbol{\mu}_{qp}^k)$ and $D(\phi_1(\boldsymbol{x}^k, \boldsymbol{\lambda}_{qp}^k, \boldsymbol{\mu}_{qp}^k); \boldsymbol{d}^k)$
  $n \leftarrow 0$, $\alpha \leftarrow 1$
  while $n \leq n_{max}$ **do**
    $n \leftarrow n + 1$
    $\boldsymbol{x}^{k+1} \leftarrow \boldsymbol{x}^k + \alpha \boldsymbol{d}^k$
    if $\phi_1(\boldsymbol{x}^k + \alpha \boldsymbol{d}^k, \boldsymbol{\lambda}_{qp}^k, \boldsymbol{\mu}_{qp}^k) - \phi_1(\boldsymbol{x}^k, \boldsymbol{\lambda}_{qp}^k, \boldsymbol{\mu}_{qp}^k) < \alpha\eta D(\phi_1(\boldsymbol{x}^k, \boldsymbol{\lambda}_{qp}^k, \boldsymbol{\mu}_{qp}^k); \boldsymbol{d}^k)$ **then**
      break
    else
      $\alpha \leftarrow \beta\alpha$
    end if
  end while
  Update acc_{inf}, acc_{step}, acc_{opt} based on appropriate criteria
end while
return $\boldsymbol{x}^k$

Algorithm 1: Pseudocode for feasible path algorithm.

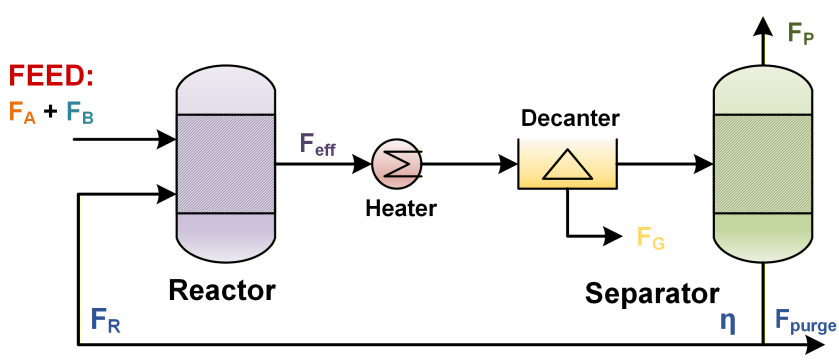

Figure 1. Flow diagram for WO process.

$$
\begin{aligned}
\min_{\boldsymbol{x} \in \mathbb{R}^n} \quad & f(\boldsymbol{x}_I, \boldsymbol{x}_D) \\
\text{s.t.} \quad & \boldsymbol{h}_1(\boldsymbol{x}_I, \boldsymbol{x}_D) = 0 \\
& \boldsymbol{h}_2(\boldsymbol{x}_I, \boldsymbol{x}_D) = 0 \\
& \boldsymbol{g}(\boldsymbol{x}_I, \boldsymbol{x}_D) \leq 0 \\
& \boldsymbol{x}_I^{lb} \leq \boldsymbol{x}_I \leq \boldsymbol{x}_I^{ub} \\
& \boldsymbol{x}_D^{lb} \leq \boldsymbol{x}_D \leq \boldsymbol{x}_D^{ub}
\end{aligned}
\tag{3}
$$

Adaptive sampling strategy

Adaptive sampling is pivotal for surrogate-assisted BBO, as it generates high-quality data, builds accurate models, and reduces function evaluations. Most existing strategies are iteratively integrated with optimization, selecting the next sampling point based on current results, which is effective for coupled sampling-optimization algorithms but unsuitable for ones that construct offline surrogate before optimization. To address this, we propose a new adaptive sampling strategy to build a global dataset for pre-optimization surrogate training, focusing on capturing details near optimal regions and variable bounds.

In this work we proposed a adaptive sampling strategy that focus on potential optimal regions and onstraint

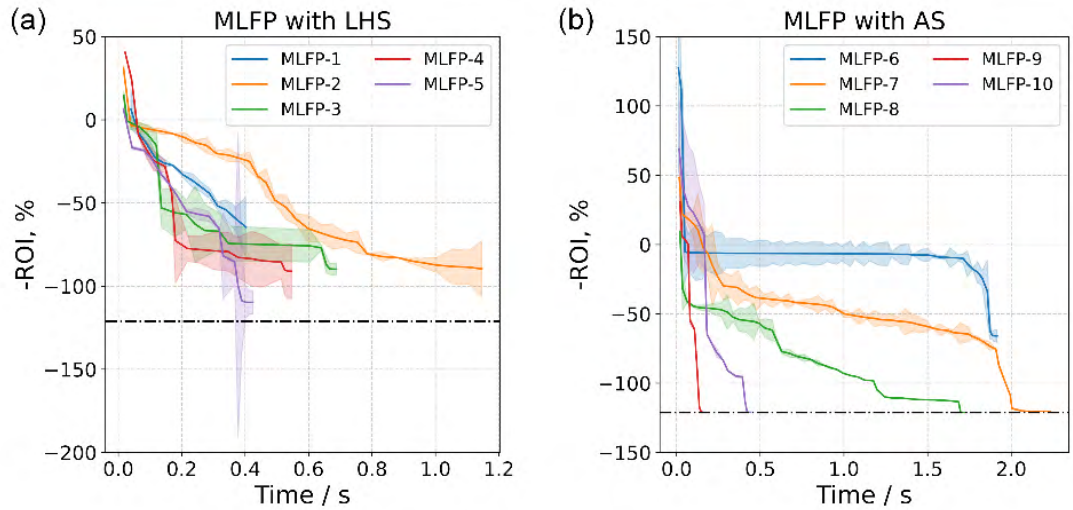

Figure 2. Convergence curves for WO process with error band of MLFP 1-10.

boundaries. The strategy proceeds in three steps:

1. Slightly expand the bounds of each independent variable (e.g., ±5% of original) and perform initial Latin hypercube sampling (LHS) to collect a small sample set (e.g., 10% of total required). Evaluate black-box functions at these samples to obtain dependent variable values, understand function behaviors, and identify promising regions.

2. Label samples as feasible (satisfying constraints like $\mathrm{g}(\mathrm{x_I}, \mathrm{x_D}), \mathrm{h}_2(\mathrm{x_I}, \mathrm{x_D})$ and bound constraints for variables) or infeasible (1 or 0), then train a support vector machine (SVM) classifier. To mitigate inaccuracy from limited initial samples, extract support vectors from the trained SVM and retrain the classifier by treating all support vectors as feasible, which includes near-boundary infeasible points to ensure precise boundary characterization.

3. Establish a sub-region around the best feasible point, conduct LHS within this area (e.g., 10% of total samples per iteration), and use the trained SVM to filter out infeasible points (avoiding wasted black-box evaluations). Update the best feasible point and repeat until the required sample size is reached.

Feasible path optimization algorithm

After the surrogate models for $\mathbf{h}_1(\mathbf{x}) = 0$ and/or the objective function are developed, then $\mathbf{h}_1(\mathbf{x}) = 0$ and/or the objective function in problem (3) are replaced by the surrogates denoted as $\boldsymbol{x}_D = \mathbf{s}(\mathbf{x}_I)$ and/or $fs(\mathbf{x}_I)$ and problem (4) below is obtained. We also assume that $\mathbf{s}(\mathbf{x}_I)$ and/or $fs(\mathbf{x}_I), \mathbf{h}_2$ and $\mathbf{g}$ are all twice continuously differentiable. The remaining constraints $\mathbf{h}_2$ and $\mathbf{g}$, and the objective function f can be calculated without calling the black-box.

$$\begin{aligned} \min_{\mathbf{x}\in\mathbb{R}^n} \quad & f(\mathbf{x}_I, \mathbf{x}_D) \text{ or } fs(\mathbf{x}_I) \\ \text{s.t.} \quad & \mathbf{x}_D = \mathbf{s}(\mathbf{x}_I) \\ & \mathbf{h}_2(\mathbf{x}_I, \mathbf{x}_D) = 0 \\ & \mathbf{g}(\mathbf{x}_I, \mathbf{x}_D) \le 0 \\ & \mathbf{x}_I^{lb} \le \mathbf{x}_I \le \mathbf{x}_I^{ub} \\ & \mathbf{x}_D^{lb} \le \mathbf{x}_D \le \mathbf{x}_D^{ub} \end{aligned} \tag{4}$$

To solve problem (4) effectively, we employ the feasible path optimization algorithms developed by Ma et al. [11]. Different from full-space formulation, which directly employs NLP algorithms to solve problem (4), feasible path optimization algorithms divide the entire problem (4) to two subproblems including a small optimization problem in the outer level and a process simulation problem in the inner level. While the small optimization problem in the outer level can be solved using the existing gradient-based optimization algorithms such as improved SQP, the process simulation problem in the inner level could be solved using solution algorithms for nonlinear equations such as a Newton-based algorithm or PTC approach [11]. In this work, the inner level only requires the use of a surrogate model to predict $\mathbf{x}_D$, which is considerably faster than rigorous simulation or black-box evaluation. Moreover, when predicting $\mathbf{x}_D$, the computation graph will be retained to compute the derivative of $\mathbf{x}_D$ with respect to $\mathbf{x}_I$, instead of calculating a large Jacobi matrix.

Detailed algorithm parameters and pseudocode are presented in the **Algorithm 1**.

In this work, the ℓ_1 merit function, defined as Eq.(5), is used to evaluate convergence criteria.

$$\phi_1(\mathbf{x}_I^k; \boldsymbol{\lambda}_{qp}^k, \boldsymbol{\mu}_{qp}^k) = f(\mathbf{x}^k) + (\boldsymbol{\rho}^k)^{\mathrm{T}} |\mathbf{h}(\mathbf{x}_I^k)| + (\boldsymbol{\nu}^k)^{\mathrm{T}} \mathbf{g}(\mathbf{x}_I^k)^{-} \tag{5}$$

where, $\mathbf{g}(\mathbf{x}^k)^{-} := \max\left(0, -\mathbf{g}(\mathbf{x}^k)\right)$, and the penalty

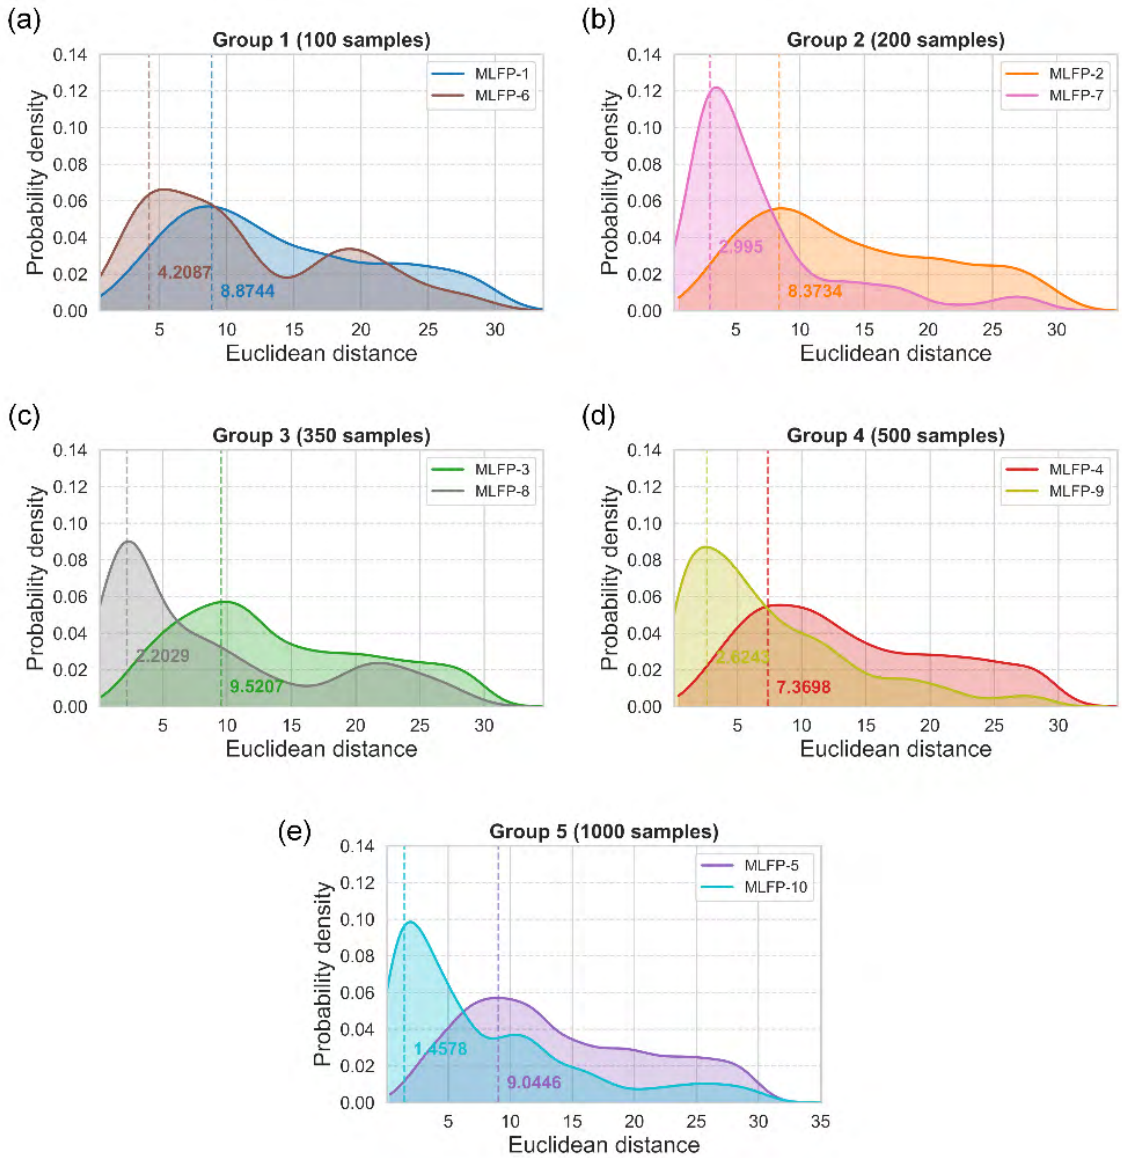

Figure 3. Probability density function of Euclidean distance from sampling points to rigorous model-based solution of WO process.

parameters are defined as follows [12]:

$$\boldsymbol{\rho}^k = max\left(|\boldsymbol{\lambda}_{qp}^k|, \frac{\boldsymbol{\rho}^{k-1}+|\boldsymbol{\lambda}_{qp}^k|}{2}\right) \tag{6}$$

$$\boldsymbol{\nu}^k = max\left(|\boldsymbol{\mu}_{qp}^k|, \frac{\boldsymbol{\nu}^{k-1}+|\boldsymbol{\mu}_{qp}^k|}{2}\right) \tag{7}$$

Given the ℓ_1 merit function being not differentiable everywhere, the Eq.(8) describes the directional derivative of $\phi_1(\mathbf{x}^k;\boldsymbol{\lambda}_{qp}^k,\boldsymbol{\mu}_{qp}^k)$ along the direction $\mathbf{d}^k$ generated by the SQP subproblem.

$$D(\phi_1(\mathbf{x}_I^k;\boldsymbol{\lambda}_{qp}^k,\boldsymbol{\mu}_{qp}^k);\mathbf{d}^k) = \nabla f(\mathbf{x}_I^k)^{\mathsf{T}}\mathbf{d}^k - (\boldsymbol{\rho}^k)^{\mathsf{T}}|\mathbf{h}(\mathbf{x}_I^k)| - (\boldsymbol{\nu}^k)^{\mathsf{T}}\mathbf{g}(\mathbf{x}_I^k)^{-} \tag{8}$$

The Armijo condition guarantees a sufficient decrease of the merit function:

$$\phi_1(\mathbf{x}_I^k + \alpha\mathbf{d}^k;\boldsymbol{\lambda}_{qp}^k,\boldsymbol{\mu}_{qp}^k) - \phi_1(\mathbf{x}_I^k;\boldsymbol{\lambda}_{qp}^k,\boldsymbol{\mu}_{qp}^k) < \alpha\eta D(\phi_1(\mathbf{x}_I^k;\boldsymbol{\lambda}_{qp}^k,\boldsymbol{\mu}_{qp}^k);\mathbf{d}^k) \tag{9}$$

Where $\eta \in (0, 0.5)$ is a hyper-parameter. In this work, we choose $\eta = 0.1$ according to Ma et al. [12] and set line search multiple $\beta = 0.618$, maximum number of line search iterations $n_{max} = 10$. However, it is possible that no step length α can satisfy the Armijo condition, regardless of how small it is. Under such circumstances, we accept the last α when the maximum number of line searches has been reached. According to the backtracking method,

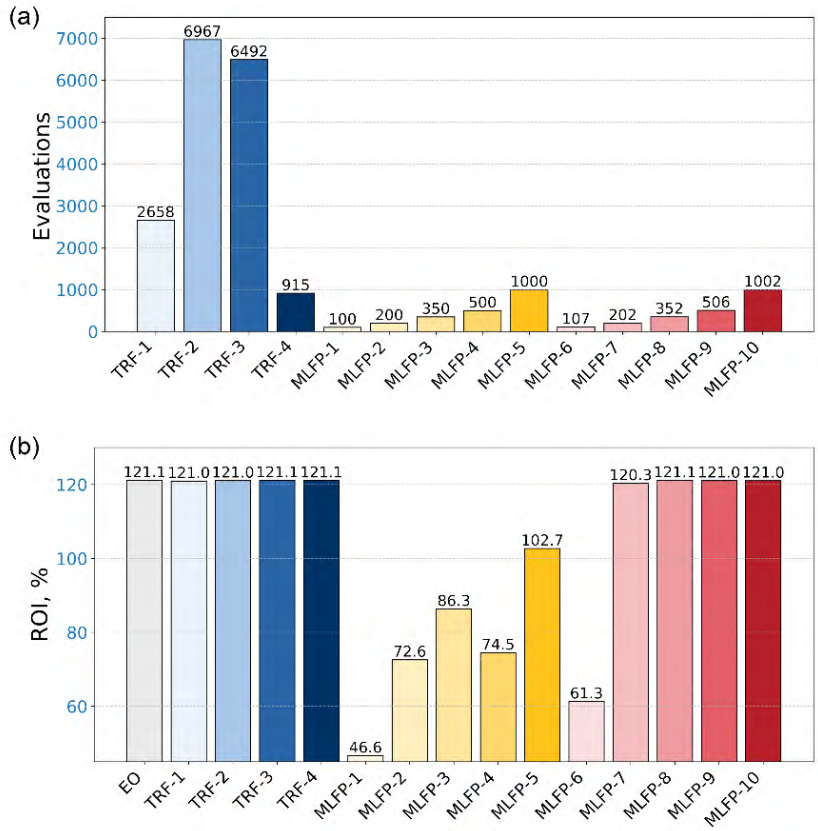

Figure 4. Comparasion for WO process from MLFPs and TRF1-4 on (a) the number of function evaluations of the black-box function and (b) objective function value.

after 10 reduction steps in the line search, the step size becomes approximately 8×10^{-3}. At this point, the step size is sufficiently small compared to the initial step size, yet it exerts little impact on convergence. Meanwhile, it remains larger than the tolerance, which ensures the iteration does not terminate prematurely before approaching the optimal value.

For practical engineering problems, due to numerical stability and other factors, the KKT conditions as criteria may be too strict, making the algorithm difficult to achieve convergence. Therefore, based on previous research, we check whether the following criteria are met after every line search:

$$acc_{inf} = \left\|\mathbf{h}\left(\mathbf{x}_I^k + \alpha \mathbf{d}^k\right)\right\|_1 + \left\|\mathbf{g}\left(\mathbf{x}_I^k + \alpha \mathbf{d}^k\right)^-\right\|_1 < tol \quad (10)$$

$$acc_{step} = \left|f\left(\mathbf{x}_I^k + \alpha \mathbf{d}^k\right) - f\left(\mathbf{x}_I^k\right)\right| < tol \quad (11)$$

$$acc_{opt} = \left\|\mathbf{d}^k\right\|_2 < tol \quad (12)$$

where, tol represents the convergence tolerance. The solution is considered optimal if Eqs.(10), (11) or Eqs.(10), (12)are satisfied. The satisfaction of Eq.(10) indicates that the optimal solution is feasible. The satisfaction of either Eq.(11) or Eq.(12) suggests that there is little potential for further decrease in the objective function.

CASE STUDY

In This work, MLPs are used to construct surrogate models with the Swish activation function. All examples are solved on a desktop with Intel(R) Core(TM) i9-14900K (3.20 GHz) processor and 128 GB of RAM running Windows 11 64-bit operating system.

Williams-Otto process

The Williams Otto (WO) process is a classic benchmark that has been widely used to evaluate optimization algorithms. The schematic diagram of the WO process is shown in **Figure 1**. The goal is to maximize the return on investment (ROI) by adjusting the reactor volume V, reaction temperature T, as well as the feed flow rates ($F_A$and F_B) and material flow cycle ratio η of components A and B. The inputs of the surrogate model are V, T, as well as F_A, F_B and η, which are also independent variables in the optimization problem. The output includes the objective function value (ROI) and the product material flow rate F_P, which are dependent variables subject to boundary constraints. In this optimization problem, there are 27

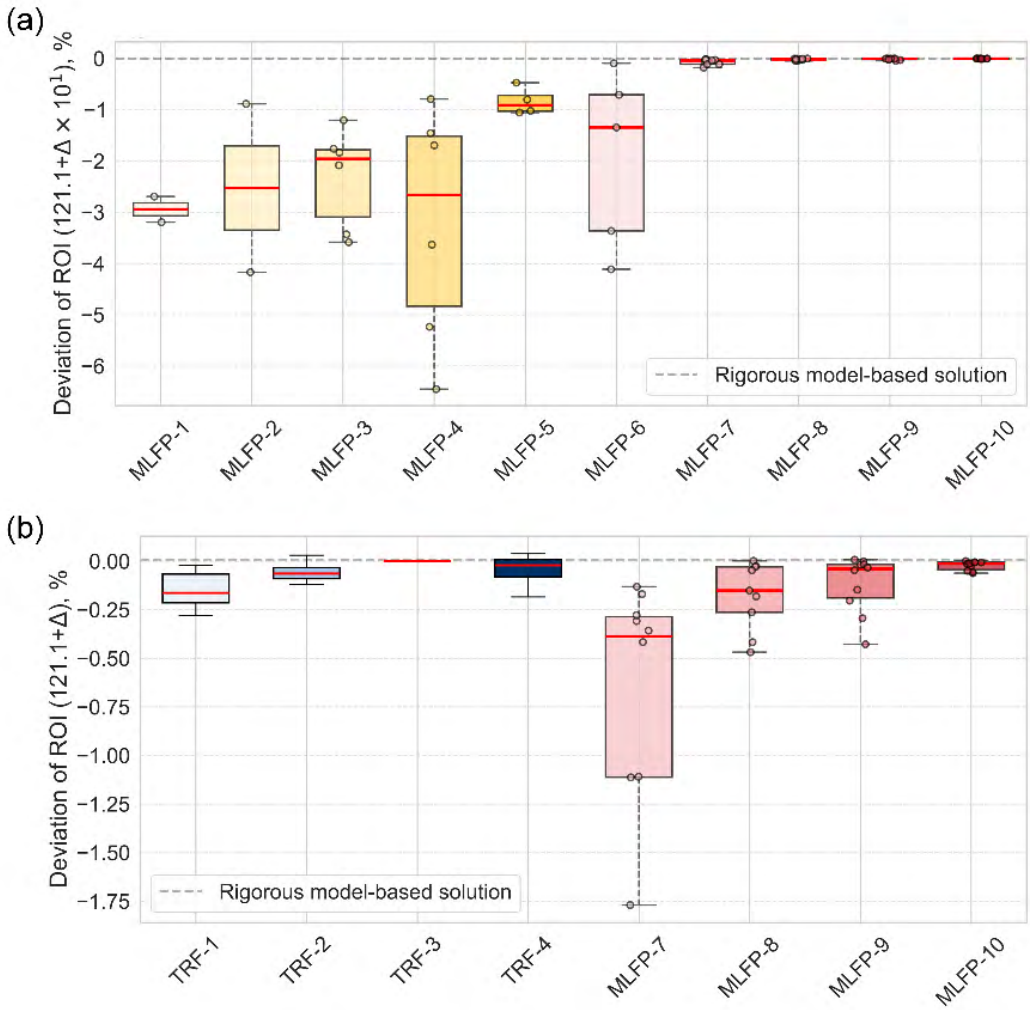

Figure 5. Boxplot of best ROI for WO process obtained from (a) MLFP1-10, and (b) TRF 1-4 and MLFP 7-10 with 10 runs.

variables and 54 constraints.

In this example, we compared the performance of TRF 1-4 and MLFP frameworks under different sampling sizes, and also compared the performance of the AS strategy and LHS strategy proposed in the MLFP framework. In the MLFP framework using the LHS strategy, we generate 100, 200, 350, 500, and 1000 valid sampling points respectively (simulation convergence successful), corresponding to MLFP 1-5; In the MLFP framework using the AS strategy, we also generate 100, 200, 350, 500, and 1000 valid sampling points, respectively, labeled as MLFP 6-10. The benchmark experimental results of the TRF method were taken from the study of Liang et al. [9]. It should be noted that the initial points used by Liang et al. [9] may differ from those of the MLFP framework.

The convergence curve of the MLFP series methods is shown in **Figure 2**. According to **Figure 2** (b), it can be seen that MLFP-7 to MLFP-10 can converge to the vicinity of the KKT point (as indicated by the black dashed line in the figure), while the results in **Figure 2** (a) show that MLFP-1 to MLFP-5 did not converge to the vicinity of the KKT point. The shaded area next to the convergence curve in **Figure 2** represents the difference between the predicted and actual values of the model. It can be seen that as the number of effective sampling points increases, the shaded area gradually narrows, indicating that the difference between the predicted and actual values continues to decrease. In other words, increasing the number of effective sampling points helps improve the prediction accuracy of the model. It can also be observed that as the iteration point approaches the true optimal value, the shaded area gradually shrinks, indicating that the samples generated by the adaptive sampling strategy are mainly concentrated near the optimal solution based on the strict model. To verify this hypothesis, this article combines samples obtained from 30 running experiments, divides the sampling results of the MLFP method into 5 groups, and assigns the same number of sampling points

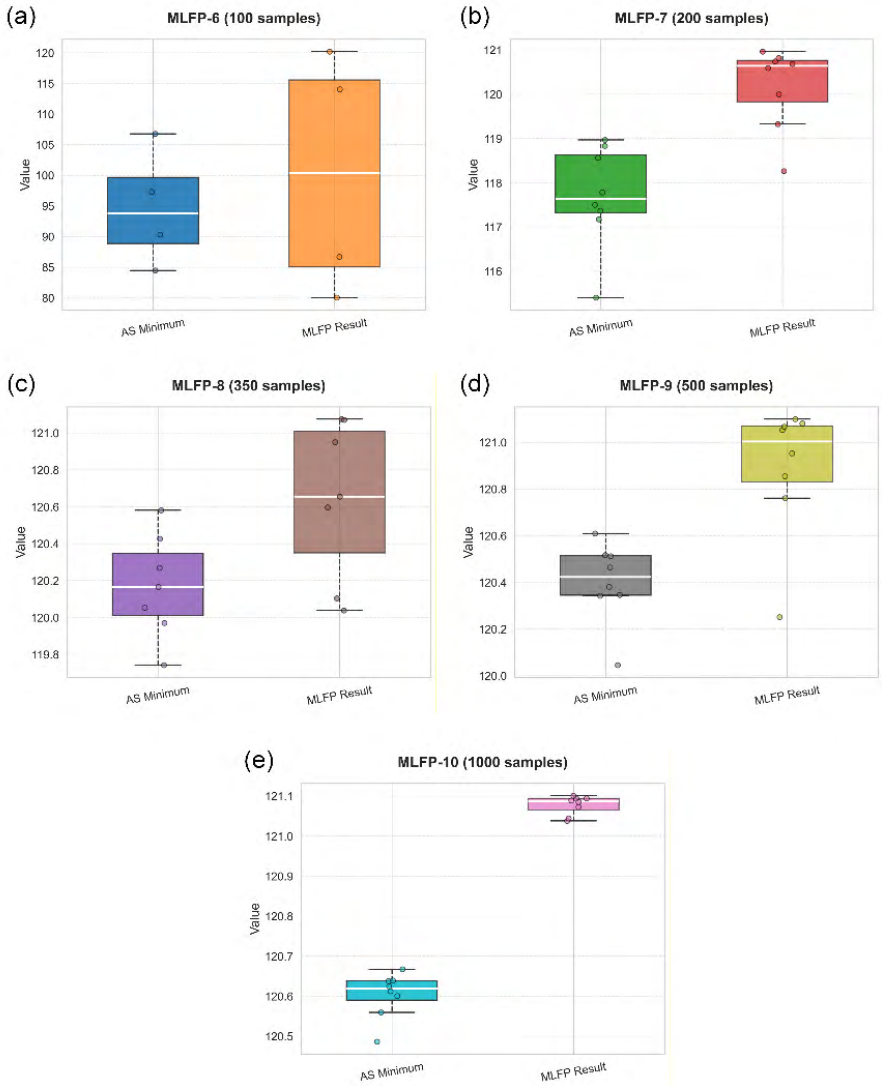

Figure 6. Comparison of the minimum objective value obtained with adaptive sampling and MLFP optimization results in WO process.

to each group. Then, LHS and AS strategies are used to conduct sampling experiments on each group, and the probability density function of the obtained samples is shown in **Figure 3**. **Figure 3** shows that under the same sampling size, the peak value of AS strategy sampling results is closer to 0 compared to LHS, indicating that its sample distribution is closer to the optimal solution based on the strict model; And as the number of sampling points increases, the peak of AS strategy sampling results will further approach 0.

The performance comparison results between TRF 1-4 and MLFP 1-10 are shown in **Figure 4**, where the dashed line represents the optimal solution obtained by using IPOPT solver to solve the strict model. Compared with TRF 1-4, MLFP 7-8 can achieve an optimal return on investment (ROI) of 121.1 and require significantly fewer function evaluations. **Figure 5** shows the optimization results of TRF 1-4 and MLFP 1-10 in 10 independent running experiments. As shown in **Figure 5** (a), under the same number of function evaluations, the stability of MLFP 6-10 is significantly better than that of MLFP 1-5, indicating that the MLFP method using the proposed adaptive sampling strategy performs better than the MLFP method using the LHS strategy; And as the number of sampling points increases, the stability of MLFP 6-10 further improves, as evidenced by the decrease in variance of the optimized ROI obtained from multiple experiments. Combining **Figure 4** (a) and **Figure 5**, it can be observed that TRF-4 and MLFP-10 require similar number of function evaluations (915 and 1002, respectively), and both can converge to the optimal solution based on the strict model; But the stability of MLFP-10 is better because its box line height, i.e. interquartile range (IQR) and whisker plot range, are much narrower than TRF-4.

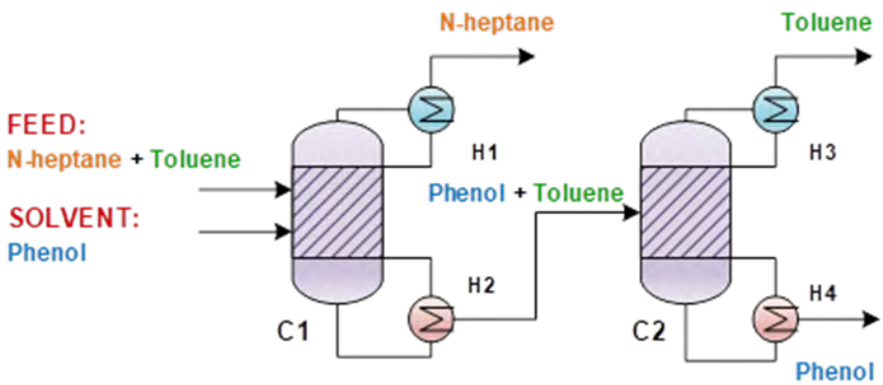

Figure 7. Process flow diagram for separation of toluene form n-heptane using solvent phenol in Ma et al.'s work [13].

In addition, we also compared the minimum target values obtained by AS and MLFP 6-10, as shown in **Figure 6**. In 10 runs, compared with the minimum value obtained by adaptive sampling, the median optimization results of MLFP 6-10 increased by 6.87%, 1.97%, 0.50%, 0.48%, and 0.39%, respectively.

Extractive distillation (ED) for toluene and n-heptane separation

Process simulation-based optimization is a typical BBO problem involving expensive function evaluations or unknown functions. In this subsection, we test an example of extractive distillation to separate toluene form n-heptane using solvent phenol provided by Ma et al. [13]. The process flowsheet is provided in **Figure 7**. The process mainly consists of two columns. The first column (C1) is used to separate n-heptane, and the second column (C2) is used for separation of toluene form solvent phenol. The purity requirement for both products is not less than 0.98. All specific constraints and utility prices are referenced from Fan et al. [10]. The optimization problem using the surrogate models mainly achieves the minimum operating cost by adjusting the amount of extractant, reflux ratio, and distillate flow rate, as shown in Eq.(13).

$$\begin{aligned} \min_{F,r_1,D_1,r_2,D_2\in\mathbb{R}} \quad & c_{hu}(Q_2+Q_4)-c_{cu}(Q_1+Q_3) \\ s.t. \quad & (x_n,x_t,Q_1,Q_2,Q_3,Q_4)=s(F,r_1,D_1,r_2,D_2) \\ & 0.98\le x_n\le 1 \\ & 0.98\le x_t\le 1 \end{aligned} \tag{13}$$

where, c_{hu} and c_{cu} is the cost coefficients for cold and hot utilities; Q_1 and Q_3 are duties of condensers C1 and C2, Q_2 and Q_4 are duties of reboilers C1 and C2, kW; F is the mass flow of solvent phenol, kg·hr^{-1}; r_1 and r_2 are the reflux ratio of C1 and C2; D_1 and D_2 are the distillate flow of C1 and C2; x_n and x_t are the mole fraction of n-heptane in C1 distillate and toluene in C2 distillate; $s(\cdot)$ is the surrogate model for extractive distillation process. The variable range and parameter values are obtained from Fan et al. [10]. In this optimization problem, there are 1, 290 variables and 1289 constraints.

According to Ma et al. [13], the surrogate models constructed using either ReLU neural networks or ALAMO exhibit great relative errors ranging from 0.2% to 25%. Despite employing complex basis functions, these models still fail to fully capture behaviors of the process system due to inherent limitations in their nonlinear representation capabilities. We use the proposed MLFP framework to generate a surrogate model for the entire process where 1014 valid sampling points are generated.

We then use MLFP to solve the problem (13). The convergence curves of the objective value, and n-heptane and toluene molar fractions are depicted in **Figure 8**. As illustrated in **Figure 8**, the predicted objective value, and n-heptane and tolumen molar fractions from the surrogated models are in good agreement with those validated from Aspen rigorous simulation, indidating the high model prediction accuracy. The best energy cost is 134.60 \$ hr^{-1}. The achieved purities of n-heptane and toluene products are 0.98 and 0.98 respectively, which satisfy the purity requirements.

To further illustrate the capability of the proposed MLFP framework, we also use genetic algorithm (GA), particle swarm optimization (PSO) to solve problem (13) where the surrogate model is obtained from the proposed MLFP framework. The reported network structure with the highest accuracy in Ma et al. [13] is also used to develop ReLU NN by using the sampling points generated from LHS, which is then reformulated into an MILP problem and optimized by Gurobi. All optimization results using the surrogate models are validated in Aspen Plus. The

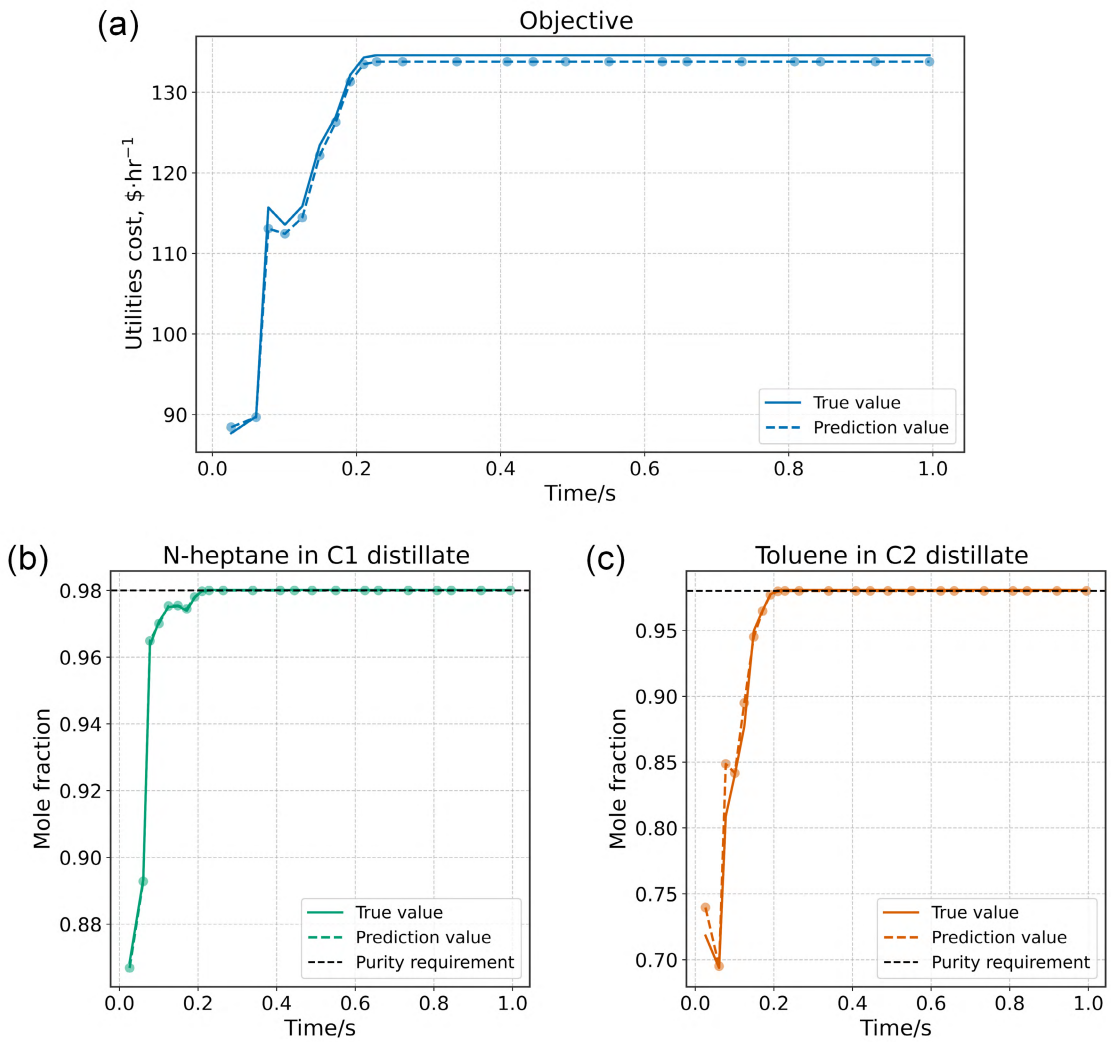

Figure 8. Convergence curves for ED process where (a) objective values, (b) n-heptane molar fraction in the top distilate of C1 and (c) toluene molar fraction in the top distillate of C2.

Table 1: Validated optimization results for ED process from different optimization strategies.

Algorithms	Objective ($·hr−1)	Purity		No. of samples/Initial points	Optimization time (sec)	Total time (min)
		x_n	x_t			
Surrogate-GA	154.77	0.99	0.98	2400	3.59	62.14
Rigorous-GA	177.07	0.98	1.00	-	2898.55	48.31
Surrogate-PSO	134.73	0.98	0.98	2400	3.94	62.15
Rigorous-PSO	151.60	0.98	0.98	-	2919.47	48.66
SQP (Aspen)	231.18	0.99	0.99	1	Failed	-
ASBO-I (best)	248.60	valid	valid	160	-	-
ReLU NN	154.13	0.96	0.96	2400	2.57	79.11
MLFP	134.60	0.98	0.98	2400	0.99	62.08

validated results are provided in **Table 1**. The best results from the adaptive sampling Bayesian optimization (ASBO-I) reported by Fan et al. [10] also provided in **Table 1** for comparison.

From Table 1, it is shown that SQP (Aspen) fails to converge and halts at the initial point provided. ASBO-I requires the fewest samples, but its best objective is 248.60 $\$\cdot hr^{-1}$, which is 84.7% higher than that of 134.60 $\$\cdot hr^{-1}$ obtained from MLFP. Although the reformulated ReLU NN using the network structure reported in Ma et al. [13] obtains the best objective of 154.13 $\$\cdot hr^{-1}$, which is lower than ASBO-I, but still higher than that of 134.60 $\$\cdot hr^{-1}$ from MLFP. Moreover, the purity of n-heptane and toluene is 0.96, which does not meet the purity requirements. In contrast, MLFP exhibits the best performance, generating the lowest objective value of 134.60 $\$\cdot hr^{-1}$ and all purity constraints are met. It should be noted that the total computational time for MLFP reported in **Table 1** includes the optimization time, the effort for evaluating the black-box function and the time for training and pre-training the SVM.

CONCLUSION

In this work, we proposed a novel machine learning powered feasible path optimization (MLFP) framework for solving general BBO problems where complex equality and inequality constraints are involved. The framework includes three main porcedures. First, an efficient adaptive sampling strategy is proposed, which filters out potentially infeasible samples with SVM and focuses on regions likely to contain optimal solutions. Second, machine learning models are leveraged to develop surrogates of black-boxes. Third, the feasible path algorithm is employed, which divides the entire optimization problem into two subproblems including a small optimization problem in the outer level and a simulation problem in the inner level (refering to the use of surrogate models for prediction in this work). Therefore, there are only independent variables that need to be updated rather than all, thus accelerating the theoretical convergence.

Computational studies show that the proposed adaptive sampling approach needs far fewer samples than the commonly used LHS strategy. The feasible path method greatly boosts computational speed, finishing optimization in just a few seconds. For benchmark problems with known global optima, the MLFP converges to solutions around the KKT point. For Williams-Otto optimization problem, MLFP shows similar performance to TRFs, but requires fewer evaluations. For an extractive distillation case, the optimal objective value obtained with MLFP is 84.7% lower than the best result of ASBO-I from its 10 runs.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge financial support from the National Key Research and Development Program of China (2024YFC3909800). Jie Li would like to appreciate the financial support from Engineering and Physical Sciences Research Council (EP/T03145X/1, EP/V051008/1).

REFERENCES

1. Shields BJ, Stevens J, Li J, Parasram M, Damani F, Alvarado JIM, Janey JM, Adams RP, Doyle AG. Bayesian reaction optimization as a tool for chemical synthesis. Nature 590:89-96 (2021). https://doi.org/10.1038/s41586-021-03213-y
2. Rangarajan S, Bhan A, Daoutidis P. Identification and analysis of synthesis routes in complex catalytic reaction networks for biomass upgrading. Applied Catalysis B: Environmental 145:149-160 (2014). https://doi.org/10.1016/j.apcatb.2013.01.030
3. Meijer HEH, Govaert LE. Mechanical performance of polymer systems: the relation between structure and properties. Progress in Polymer Science 30:915-938 (2005). https://doi.org/10.1016/j.progpolymsci.2005.06.009
4. Dedeyne JN, Geerts M, Reyniers PA, Wéry F, Van Geem KM, Marin GB. Computational fluid dynamics?based optimization of dimpled steam cracking reactors for reduced co₂ emissions. AIChE Journal 66: (2020). https://doi.org/10.1002/aic.16255
5. Lang YD, Malacina A, Biegler LT, Munteanu S, Madsen JI, Zitney SE. Reduced order model based on principal component analysis for process simulation and optimization. Energy Fuels 23:1695-1706 (2009). https://doi.org/10.1021/ef800984v
6. Na J, Lim Y, Han C. A modified DIRECT algorithm for hidden constraints in an LNG process optimization. Energy 126:488-500 (2017). https://doi.org/10.1016/j.energy.2017.03.047
7. Bertsimas D, Margaritis G. Global Optimization: A Machine Learning Approach. *arXiv*. Preprint posted online November 3, 2023. Accessed January 12, 2024. http://arxiv.org/abs/2311.01742
8. Cozad A, Sahinidis NV, Miller DC. Learning surrogate models for simulation?based optimization. AIChE Journal 60:2211-2227 (2014). https://doi.org/10.1002/aic.14418
9. Liang R, Han Y, Hu H, Chen B, Yuan Z, Biegler LT. Efficient trust region filter modeling strategies for computationally expensive black-box optimization. Computers & Chemical Engineering 189:108816 (2024). https://doi.org/10.1016/j.compchemeng.2024.108816
10. Fan S, Hong X, Liao Z, Ren C, Yang Y, Wang J, Yang Y. Adaptive sampling bayesian algorithm for constrained black?box optimization problems.

AIChE Journal 71: (2025). https://doi.org/10.1002/aic.18715
11. Ma Y, McLaughlan M, Zhang N, Li J. Novel feasible path optimisation algorithms using steady-state and/or pseudo-transient simulations. Computers & Chemical Engineering 143:107058 (2020). https://doi.org/10.1016/j.compchemeng.2020.107058
12. Ma Y, Gao X, Liu C, Li J. Improved SQP and SLSQP algorithms for feasible path-based process optimisation. Computers & Chemical Engineering 188:108751 (2024). https://doi.org/10.1016/j.compchemeng.2024.108751
13. Ma K, Sahinidis NV, Bindlish R, Bury SJ, Haghpanah R, Rajagopalan S. Data-driven strategies for extractive distillation unit optimization. Computers & Chemical Engineering 167:107970 (2022). https://doi.org/10.1016/j.compchemeng.2022.107970

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Work and Heat Exchanger Networks as a General Energy-Integration Strategy for Chemical Processes

José A. Caballero[a]*, Zinet Mekidiche-Martínez[a] and Juan A. Labarta[a]

[a] University of Alicante, Institute of Chemical Process Engineering, Alicante, Spain
* Corresponding Author: caballer@ua.es.

ABSTRACT

The integrated recovery of heat and mechanical work has gained increasing importance in process integration due to the strong thermodynamic coupling between temperature and pressure changes in many industrial systems. This work presents a rigorous framework for the simultaneous synthesis of Work and Heat Exchanger Networks (WHEN), in which heating, cooling, compression, expansion, throttling, and pumping are optimized in a coordinated manner. The problem is formulated using Generalized Disjunctive Programming (GDP), allowing the explicit representation of alternative thermodynamic paths, phase-dependent behavior, and logical equipment choices. Process streams are defined by supply and target states, while only bounds are imposed on intermediate pressures, temperatures, and flow rates. Streams may change classification between hot and cold multiple times and may undergo several phase transitions.Rigorous thermodynamic correlations obtained from Aspen HYSYS are embedded in the optimization model, enabling a physically consistent treatment of condensation, vaporization, and isenthalpic valve expansion without restrictive assumptions. The heat exchanger network is represented implicitly through the Pinch Location Method, avoiding the combinatorial complexity of explicit exchanger matching while accommodating multiple utilities and unclassified streams.The approach is demonstrated through a self-refrigerated alkylation process case study. The optimal solution reveals a non-intuitive thermodynamic path involving alternating vapor and liquid compression stages and multiple phase changes, which enhances internal heat recovery. Compared to a conventional design strategy, the proposed framework reduces total annual energy integration costs by 12.5%. The resulting nonconvex MINLP is solved to global optimality within seconds, demonstrating both robustness and computational tractability.

Keywords: Work exchanger networks, Heat exchanger networks, WHEN, Energy efficiency, Process Integration

INTRODUCTION

In recent years, the integrated treatment of heat and work recovery has emerged as a key topic in process integration, motivated by the need to further reduce energy consumption and carbon emissions beyond what can be achieved through conventional heat integration alone. While classical Heat Exchanger Network (HEN) synthesis focuses on the recovery of thermal energy and Work Exchange Network (WEN) synthesis addresses the reuse of mechanical energy, many industrial processes exhibit strong thermodynamic coupling between heat and work. This coupling has led to the development of Work and Heat Exchanger Networks (WHEN), in which heat and shaft work are considered in a coordinated and sequential manner rather than as independent subsystems [1].

The relevance of WHEN becomes particularly evident in processes involving significant pressure changes, such as gas processing, refrigeration and liquefaction systems, air separation units, and high-pressure chemical reactors. In these systems, compression and expansion steps not only consume or produce mechanical work but also generate substantial temperature changes. As a result, decisions related to pressure manipulation directly affect the thermal structure of the process and, consequently, the opportunities for heat recovery.

According to Haoshui et al. [1] The systematic

integration of power and heat has evolved significantly over several decades, moving from basic component design to complex network synthesis. In 1983, the principles of Pinch Analysis were expanded through the "appropriate placement" concept, allowing for the strategic integration of heat engines and heat pumps [2]. This conceptual framework was further formalized in 1987 through the development of a superstructure-based optimization model for the integration of these thermal machines [3].

The specific concept of Work Exchanger Networks (WENs) was formally proposed in 1996 [4], marking a shift toward treating work as a primary utility. In 2007, the ExPAnD procedure advanced the field by combining heuristics, pinch, and exergy analyses for subambient design [5]. By 2014, the field matured into the simultaneous design of Work and Heat Exchanger Networks (WHENs) through superstructure-based MINLP optimization [6].

Recent advancements have focused on refining these synthesis methods to handle increasing complexity. In 2015, new insights regarding the appropriate placement of compressors and expanders led to the development of manual design procedures [7–10]. This was followed by the introduction of an extensive superstructure for WHENs. Perhaps the most comprehensive superstructure is attributed to Nair et al. [11]. In their model, streams are not pre-classified as hot or cold, or as high-pressure or low-pressure. Instead, the model determines their behavior, which may even change across the various stages a stream can undergo between its supply and target conditions. Furthermore, unlike previous models, it accounts for liquid-vapor phase changes. However, it only considers the use of turbines, compressors, and valves, omitting pumps to increase liquid pressure (for instance, in systems that favor condensation prior to pressure elevation to avoid the use of costly compressors).

WHEN FRAMEWORK

The WHEN problem we are dealing with can be stated as follows:

Given a set of process streams with specified supply and target states (pressure, temperature, and phase), the objective is to determine the optimal thermodynamic path for each stream—incorporating heating, cooling, compression, and expansion stages—to synthesize a network of pressure-change equipment, heat transfer units, mixers, and splitters. This configuration is optimized based on a relevant performance metric, such as minimum energy cost, minimum total annualized cost, or minimum total exergy consumption.

Several relevant features address and overcome common limitations in existing models:

- Variable Bounds: The values for state variables (P, T) or flow rates do not need to be fixed beforehand; only the lower and upper bounds for these parameters are required.
- Stream Classification: Whenever possible, streams are pre-classified as hot (requiring cooling) or cold (requiring heating) to improve the numerical performance of the model. If the behavior is unknown a priori, streams can be defined as "unclassified, " allowing the optimization to determine their role.
- Variable Thermodynamic Paths: Streams undergoing both temperature and pressure changes may follow complex routes involving compressions, expansions, heating, or cooling. Consequently, their behavior may switch between "hot" and "cold" multiple times along the path. Therefore, these streams are generally treated as "unclassified."
- Phase-Independent Pressure Changes: Pressure changes are not restricted to gaseous streams (a common simplification in many models); the model includes also liquid compression and expansion through valves, which may result in liquid-vapor mixtures.
- Rigorous Thermodynamic Correlations: The thermodynamic behavior of each stream between its supply and target states is rigorously correlated using Aspen HYSYS. Integrating these correlations into the optimization model provides a realistic, physically consistent description of process behavior, which is particularly important for streams undergoing condensation, vaporization, or significant density variations.
- Isenthalpic Valve Expansion: Expansion through valves is assumed to be isenthalpic. By utilizing rigorous correlations, it is unnecessary to assume constant Joule-Thomson coefficients or to restrict valve expansion to temperatures below the inversion point. Besides, expansions that end in a multiple phase state can be also incorporated (i.e. typical in refrigeration cycles)
- Multiple Utilities: The model allows for the simultaneous use of multiple hot and cold utilities.
- Multistage Compression: If multistage compression is required, it emerges naturally as a result of the optimization process.

Superstructure and model formulation

The overall WHEN model is formulated using Generalized Disjunctive Programming (GDP), which provides a rigorous mathematical structure to represent logical conditions, operational modes, and phase-dependent behavior.

Process streams are allowed to undergo

compression, expansion, throttling, or pumping across multiple pressure-change stages, potentially accompanied by condensation and/or evaporation. Indeed, a single stream may experience several phase changes as it evolves from its initial to final conditions. This behavior is represented through a stage wise model of pressure-change units, schematically shown in Figure 1. Each stage may include compressors, turbines, pumps, and valves. Compressors and turbines operate with vapor-phase streams, while pumps are restricted to liquid-phase segments. Valves may handle vapor or liquid feeds; in the latter case, partial or total condensation may occur due to isenthalpic expansion. A bypass option represents cases in which a stage does not require pressure modification. The outlet of each pressure-change stage is subsequently sent to the Heat Exchanger Network for potential heat recovery or rejection, as indicated in Figure 1

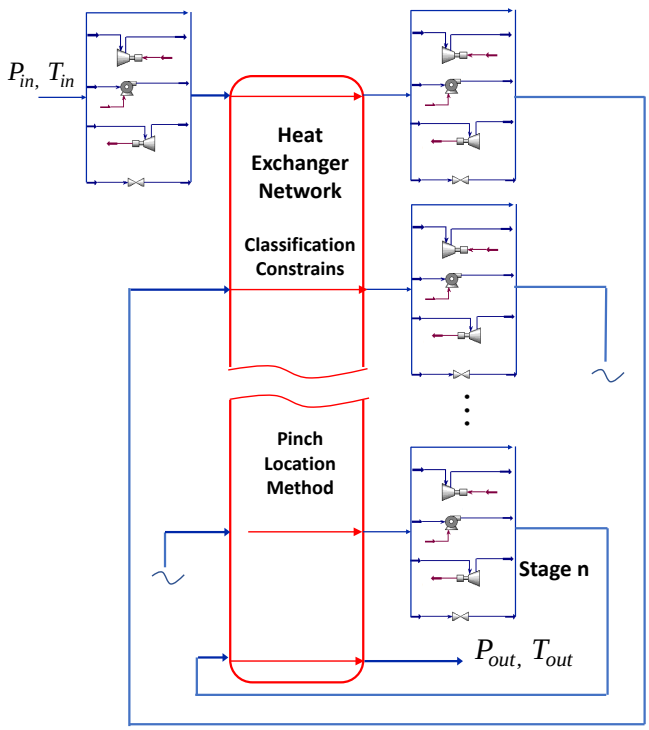

Figure 1. WHEN Superstructure involving streams with change in temperature and presure.

Each stage of a process stream is described as a five-term disjunction, in which each term includes a rigorous model for each one of the possible equipment involved: compression, assumed isentropic with fixed efficiency; expansion in a turbine, also isentropic with fixed efficiency; throttling through a valve assumed isenthalpic; liquid pumping assuming constant liquid density calculated at the inlet conditions; or bypass that simple equals inlet to outlet state variables.

The heat exchanger network itself is not modeled explicitly as a set of predefined exchanger matches, as this would lead to an intractable number of discrete combinations, especially when stream classification and phase behavior are uncertain. Instead, the network is represented implicitly through the Pinch Location [12, 13] Method, which efficiently handles variable heat loads, flexible temperatures, unclassified streams, and multiple hot and cold utilities without enumerating individual exchangers.

Case Study

To illustrate the model, a case study was conducted using data extracted from the process flowsheet presented by Luyben [14] for the self-refrigerated alkylation of C4 olefins with isobutane to produce high-octane C8 components.

In this process, the alkylation reactions occur in a series of stages, which are typically simulated using boiling continuous stirred-tank reactors (CSTRs) such that the reaction product is in the vapor phase. The pressure of this stream must be increased, in this case from approximately 300 kPa to 1480 kPa, for introduction into the separation system. Given that the energy required to compress vapor is significantly higher than that required for liquid compression, and that compressors are among the most expensive equipment in the chemical industry, standard practice involves either condensing the vapor and pumping the resulting liquid or—if the condensation temperature is too low—utilizing a compressor until the condensation temperature is high enough to allow for condensation and subsequent liquid pumping. The original flowsheet uses this second approach.

While this strategy is conceptually sound, it does not account for the potential benefits of integration with the remainder of the process.

Table 1 presents the data for process streams with temperature changes only. Table 2 shows the available utilities while Table 3 details the streams involving simultaneous pressure and temperature changes.

Costs for compressors, turbines, pumps, and other equipment in the optimization model are correlated using the procedure by Turton et al. [15]. Efficiencies for compressors, turbines, and pumps are set at 75%, consistent with the Aspen.Hysys default values.

Table 1. Data for streams without pressure change.

Name	Tin (°C)	Tout (°C)	FCp (kW/°C)
H1	47.01	45.77	130.887
H2	83.26	47.22	7.342
H3	74.19	47.63	245.313
H4	47.22	37.22	63.25
C1	80.69	83.26	75.714
C2	62.56	82.26	367.980
C3	21.00	36.67	28.469

Table 2. Data for available utilities.

Name	Tin(°C)	Tout (°C)	Cost (k$/kW·y)
LP-Steam	120	119	666.4
Water	20	30	20.0
Refrig.	1	5	100
Electricity			1500
Minimum Approach Temperature (ΔT_{min}) = 10 ºC			

Table 3. Data for streams involving simultaneous pressure and temperature changes.

Stream name	St1
Supply Temperature (ºC)	20
Target Temperature (ºC)	75
Intermediate WEN Stages temperature bounds (ºC)	15 – 100
Supply Pressure (kPa)	299.2
Target Pressure (kPa)	1480
Intermediate WEN Stages Pressure bounds (kPa)	250-1500
Supply Phase	Vapor
Target Phase	Liquid
Mean molar liquid density (kmol/m^3)	9.871
Number of WEN Stages	3
Flow (kmol/h)	225.9
Composition (molar fractions)	
Propane	0.1230
i-Butane	0.8415
n-Butane	0.0351
2, 2, 4-Mpentane	0.0004

The optimal solution includes the selection of hot and cold utilities, energy costs (thermal and electrical), and the annualized costs of pressure-change equipment (pumps, turbines, compressors, and drivers); however, it excludes the heat exchanger network (HEN) investment. Once all stream temperatures and flows are determined, designing a HEN that meets all requirements is relatively straightforward.

A simplified flowsheet identifying the process streams involved in the heat and work exchange network is shown in Figure 2.

The optimal solution includes the selection of hot and cold utilities, energy costs (thermal and electrical), and the annualized costs of pressure-change equipment (pumps, turbines, compressors, and drivers); however, it excludes the heat exchanger network (HEN) investment. Once all stream temperatures and flows are determined, designing a HEN that meets all requirements is relatively straightforward. Figure 3 presents the optimal configuration, including the pressure-change equipment for the relevant streams.

The optimal solution shows counterintuitive behavior. Initially, the vapor stream is compressed to the minimum pressure required for condensation with available cold utilities (~309 kPa at approximately 15 ºC). Compression then continues using a liquid-phase pump. However, instead of reaching the final target pressure, the stream is only pumped to ~937 kPa, where it is evaporated and superheated by process streams. A second compressor then raises the pressure to the final specification. This final compression also increases the temperature to approximately 15 ºC above the target temperature.

The second condensation together with the final superheating are used by the system to increase the degree of heat integration, reducing both hot and cold utilities.

The total energy integration cost using the strategy described above is 4, 395 k$/y. In contrast, following the strategy presented by Luyben [15] results in a cost of 4, 945 k$/y, which is 12.5% higher. Furthermore, decoupling the final heat exchanger network design from the energy and major equipment cost targets (such as compressors and turbines) ensures model reliability and robustness. This case study is formulated as a non-convex Mixed-Integer Non-Linear Programming (MINLP) problem, consisting of 2, 752 equations, 2, 065 variables, and 45 binary variables. It was solved to global optimality using the GUROBI solver [16] in 6.5 seconds using a personal computer (Intel(R) Core(TM) i7-10700 2.90GHz with 32 GB RAM)

ACKNOWLEDGEMENTS

The authors gratefully acknowledge financial support from the «Ministerio de Ciencia, Innovación y Universidades» from Spain under the project PID2024-155232NB-I00 -Net-0-Cheml-.

AUTHOR IDENTIFIERS

Author ORCIDs:
Caballero JA: 0000-0001-6470-2907
Mekidiche-Martínez Z: 0009-0000-0373-5588
Labarta JA: 0000-0002-4870-2031

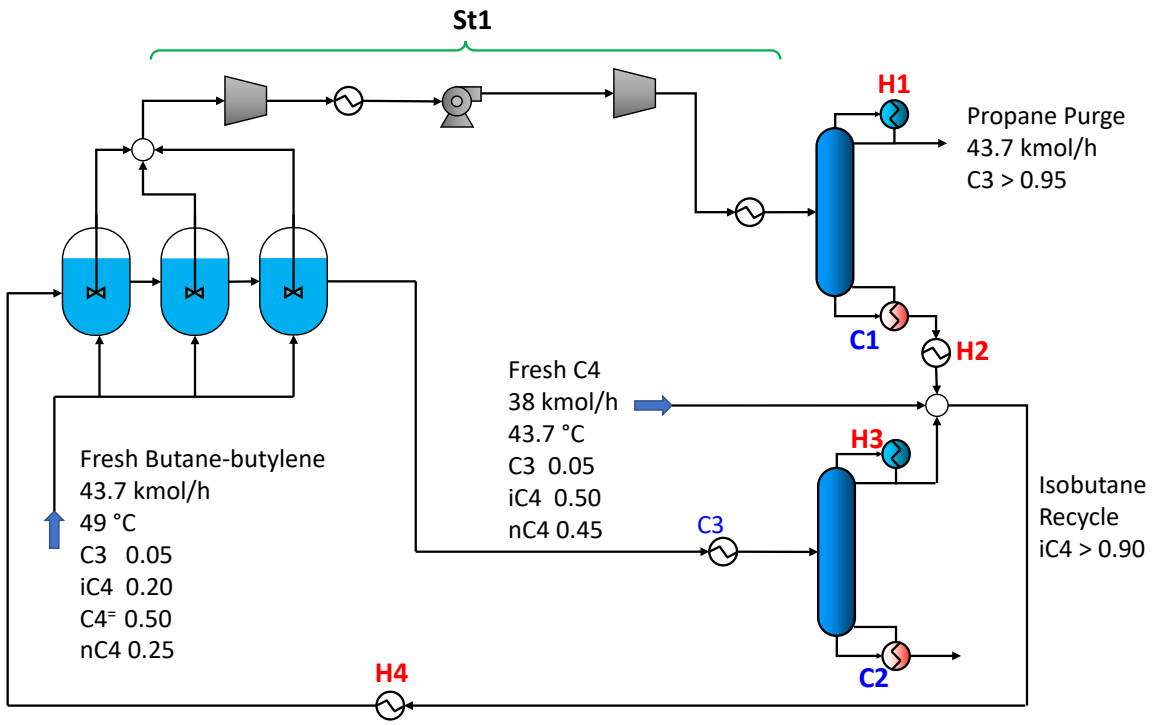

Figure 2. Flowsheet scheme indicating the streams that form part of the work and heat exchanger network

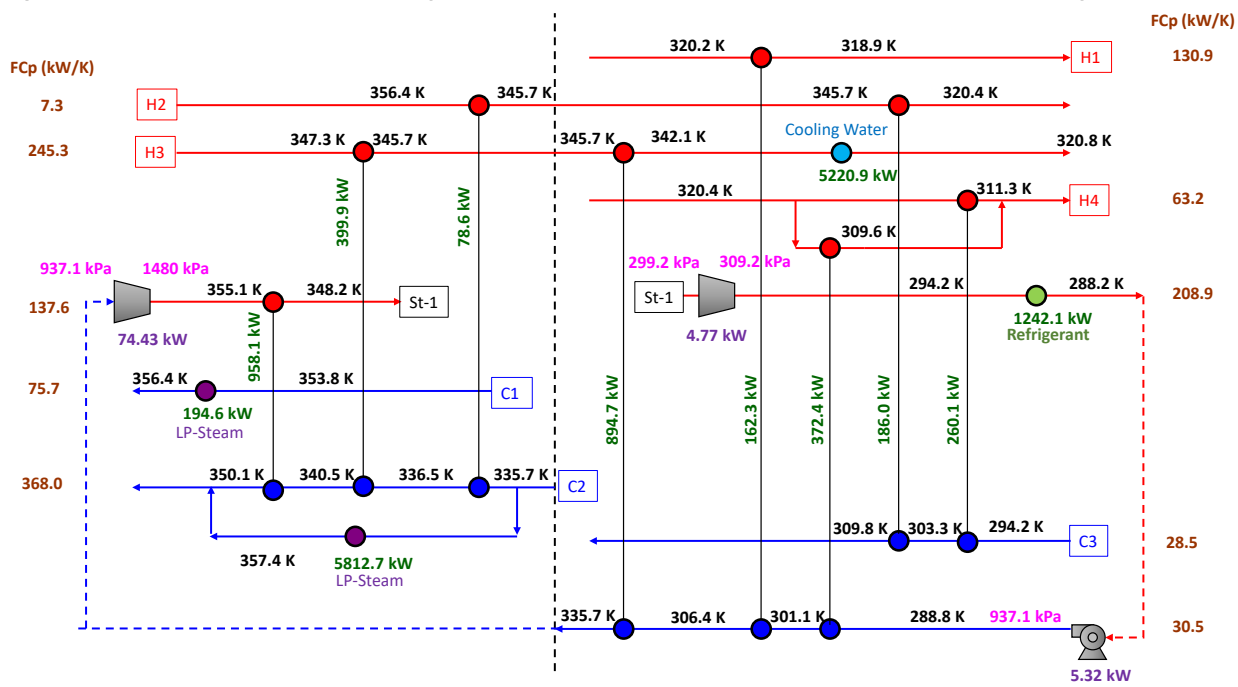

Figure 3: Heat exchanger network for the Case Study 1 optimal solution. Streams H1–H4 are hot process streams,

REFERENCES

1. Yu H, Fu C, Gundersen T. Work exchange networks (wens) and work and heat exchange networks (whens): a review of the current state of the art. Ind. Eng. Chem. Res. 59:507-525 (2019). https://doi.org/10.1021/acs.iecr.9b04932
2. Townsend DW, Linnhoff B. Heat and Power Networks in Process Design Part I: Criteria for Placement of Heat Engines and Heat Pumps in Process Networks. AIChE J 1983;29:742–7.

3. Colmenares TR, Seider WD. Heat and power integration of chemical processes. AIChE Journal 33:898-915 (2004). https://doi.org/10.1002/aic.690330604
4. Huang YL, Fan LT. Analysis of a work exchanger network. Ind. Eng. Chem. Res. 35:3528-3538 (1996). https://doi.org/10.1021/ie9507383
5. Aspelund A, Berstad DO, Gundersen T. An extended pinch analysis and design procedure utilizing pressure based exergy for subambient cooling. Applied Thermal Engineering 27:2633-2649 (2007). https://doi.org/10.1016/j.applthermaleng.2007.04.017
6. Onishi VC, Ravagnani MASS, Caballero JA. Simultaneous synthesis of heat exchanger networks with pressure recovery: optimal integration between heat and work. AIChE Journal 60:893-908 (2013). https://doi.org/10.1002/aic.14314
7. Fu C, Gundersen T. Integrating compressors into heat exchanger networks above ambient temperature. AIChE Journal 61:3770-3785 (2015). https://doi.org/10.1002/aic.15045
8. Fu C, Gundersen T. Integrating expanders into heat exchanger networks above ambient temperature. AIChE Journal 61:3404-3422 (2015). https://doi.org/10.1002/aic.14968
9. Fu C, Gundersen T. Sub-ambient heat exchanger network design including compressors. Chemical Engineering Science 137:631-645 (2015). https://doi.org/10.1016/j.ces.2015.07.022
10. Fu C, Gundersen T. Sub-ambient heat exchanger network design including expanders. Chemical Engineering Science 138:712-729 (2015). https://doi.org/10.1016/j.ces.2015.09.010
11. Nair SK, Nagesh Rao H, Karimi IA. Framework for work?heat exchange network synthesis (WHENS). AIChE Journal 64:2472-2485 (2018). https://doi.org/10.1002/aic.16129
12. Quirante N, Caballero JA, Grossmann IE. A novel disjunctive model for the simultaneous optimization and heat integration. Computers & Chemical Engineering 96:149-168 (2017). https://doi.org/10.1016/j.compchemeng.2016.10.002
13. Onishi VC, Quirante N, Ravagnani MASS, Caballero JA. Optimal synthesis of work and heat exchangers networks considering unclassified process streams at sub and above-ambient conditions. Applied Energy 224:567-581 (2018). https://doi.org/10.1016/j.apenergy.2018.05.006
14. Luyben WL. Principles and case studies of simultaneous design. Wiley (2011). https://doi.org/10.1002/9781118001653
15. Turton R, Bailei RC, Whiting WB, Shaeiwitz, J. A., & Bhattacharyya D. Analysis, Synthesis and Design of Chemical Processes. 5th editio. Upper Saddle River, NJ. USA.: Pearson Education, Inc.; 2018.
16. Gurobi Optimization L. Gurobi Optimizer Reference Manual 2025.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Designing a Load-Flexible Renewable Ammonia Plant for Variable Green Hydrogen Supply

Niklas Groll[a]* and Gürkan Sin[a]

[a] Process and System Engineering Center (PROSYS), Department of Chemical and Biochemical Engineering, Technical University of Denmark (DTU), 2800, Kgs. Lyngby/Denmark
* Corresponding Author: nigro@kt.dtu.dk.

ABSTRACT

Decarbonizing ammonia by replacing grey with green hydrogen directly affects the operation of the Haber-Bosch (HB) process. When directly coupled to green hydrogen production from renewable energy, the HB process should be able operate flexibly to match variable hydrogen supply. This study presents a structured approach for designing a load-flexible HB plant, supported by a rigorous process model. First, we screen 2, 000 designs at high (100%) and low (10%) hydrogen loads to assess operability. Only 1, 100 designs are feasible for both loads, underscoring the need to account for multivariable interactions during design. Next, we assess the economic feasibility of a base design, comparing HB operation under constant and flexible loads. Flexible operation reduces the levelized cost of ammonia (LCOA) by about 5.8%, primarily by lowering green hydrogen production costs. This cost reduction results from downregulating hydrogen production during periods of high electricity prices. By contrast, HB design improvements yield only small LCOA reductions (about 0.4%), though lower reactor pressure and a larger reactor volume remain the best HB design options to further reduce renewable ammonia costs.

Keywords: Process Design, Process Operations, Renewable and Sustainable Energy, Green Ammonia

INTRODUCTION

Ammonia is one of the most-produced chemicals worldwide. It is primarily used as a fertilizer, but could also serve as an energy carrier or fuel in the future [1, 2]. Most ammonia is produced from hydrogen via steam methane reforming, also called grey hydrogen. Since grey hydrogen is produced from natural gas, ammonia production inherently accounts for approximately 2% of global greenhouse gas emissions [2]. To reduce the environmental impact of ammonia production, studies suggest replacing grey hydrogen with green hydrogen as the primary feedstock [1, 2]. Pure grey and green hydrogen are chemically identical; however, their production behavior can significantly affect ammonia plant operations. Today's commercial ammonia plants all rely on the Haber-Bosch (HB) process, in which hydrogen and nitrogen are converted to ammonia over a metal catalyst at high pressure and temperature. The HB process is expected to remain the preferred option, at least in the mid-term, even when coupled to a green hydrogen supply.

The main difference between producing ammonia from green or grey hydrogen lies in the availability and flexibility of hydrogen supply. Typically, ammonia production is coupled directly with a hydrogen production process to minimize the cost of hydrogen storage and transport [3, 4]. Grey hydrogen, produced continuously from methane, enables the HB process to operate continuously at a constant hydrogen load. In contrast, green hydrogen production via electrolysis powered by renewable energy can exhibit significant load variations, mainly due to the intermittent nature of renewable electricity supply. Although modern electrolyzers can rapidly adapt to changing conditions [5], incorporating this flexibility into load-flexible HB plant design and operation remains an open challenge, increasing the design complexity.

Many previous studies have focused on the core of the HB process, the reactor system [6]. The reactor system usually consists of a combination of adiabatic beds and cooling systems to drive the strongly exothermic reaction $3\,H_2 + N_2 \rightarrow 2\,NH_3$. Most conventional HB plants today still rely on the reactor system, which couples multiple adiabatic reactor beds with quench cooling (AQCR). However, Morud and Skogestad first noted that this

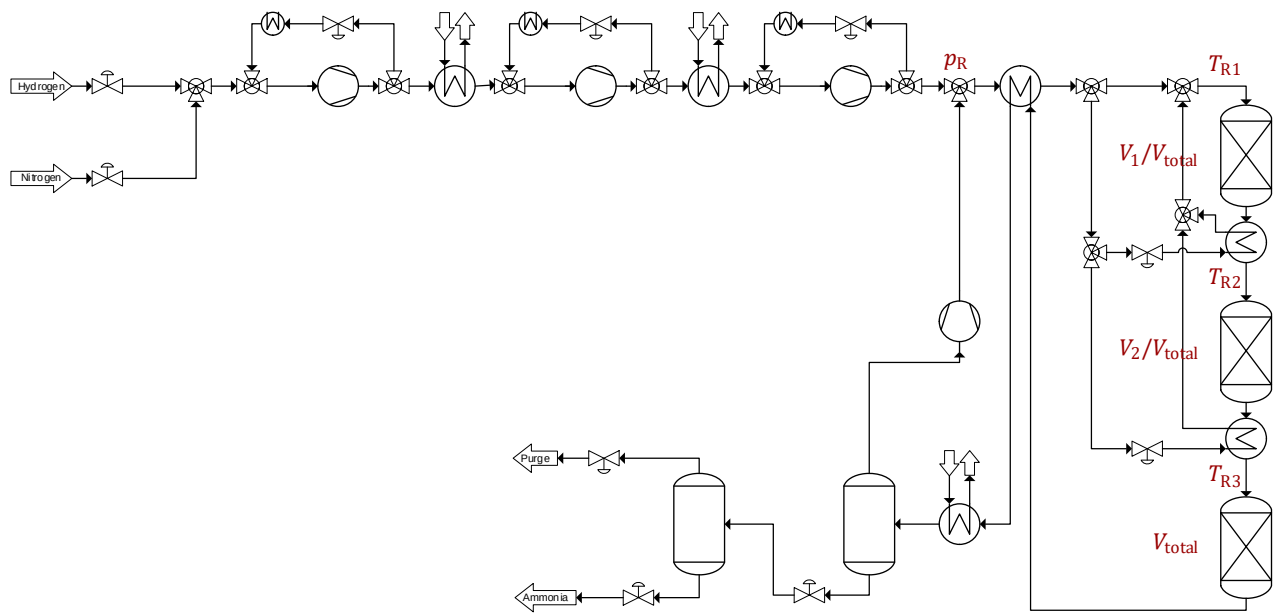

Figure 1: Haber-Bosch process flowsheet with indication of the seven key design variables.

AQCR system could experience stability issues during load changes [7]. In contrast, Rosbo *et al.* found that adiabatic indirect-cooled reactor (AICR) systems can handle load changes well and achieve higher conversion rates than AQCR systems [6]. Since then, several studies have focused on operating HB plants for green ammonia production using AICR systems, assuming variable hydrogen supply rates ranging from 100% to 10% of the nominal capacity [4, 8]. Furthermore, given the described advantages, ammonia plant licensors are increasingly focusing on AICR systems rather than AQCR systems, particularly for green hydrogen applications [8, 9]. Recently, the world's first large-scale demonstration plant for dynamic ammonia production opened in Denmark, operating solely on renewable energy and using an AICR system [9]. AICR systems use three reactor beds with interbed cooling in series. The coolers are usually heat-integrated, resulting in a complex HB design. Key design choices include compressor dimensions, reactor bed volumes, and heat exchanger sizes. However, dimensions are rarely publicly available.

Given the degrees of freedom, traditional process design focuses on a single steady-state operating point. In the HB process, this point corresponds to a nominal hydrogen supply rate. But, coupling the HB process to a green hydrogen source often results in a variable feed supply, making single-point designs potentially infeasible or uneconomical. Studies show that certain HB process designs can operate at reduced hydrogen feed rates through various low-load operating procedures [4, 10]. Opinions diverge on whether the reactor pressure should be reduced from the nominal setpoint during low-load operation [4] or remain constant [10]. While some studies showed better performance under varying pressure [4], we believe that, given concerns about mechanical stress from frequent pressure changes, operating at constant reactor conditions is preferable.

In this study, we provide guidance for designing a flexible HB plant to produce green ammonia using green hydrogen. We explore the range of feasible AICR system designs using a detailed, rigorous simulation model and demonstrate the potential economic advantages of flexible operation. We also examine how HB process design can be adapted and improved to further reduce costs when coupled to a green hydrogen supply.

METHODOLOGY

In this chapter, we first detail the HB process model and its key design variables, including reactor pressure, bed inlet temperatures, and reactor bed volumes. We then introduce a sampling algorithm to efficiently explore the feasible design space defined by these variables. Further, we present how to calculate and quantify the costs of green ammonia production. Finally, we explain how dynamic electricity prices can serve as decision variables for operating the green hydrogen production and HB plant at variable loads.

Green Ammonia Process Design

The green ammonia process design follows the same basic flowsheet as the well-known HB process, developed by Fritz Haber and Carl Bosch in the early 20th century, and largely unchanged since then. Our model, shown in Figure 1, comprises three main parts: (i) the feed compression of reactants (hydrogen and nitrogen), (ii)

the reactor system for converting the reactants to ammonia, and (iii) the separation system for recovering unreacted gases and product purification. The plant is designed to produce 500, 000 tonnes of ammonia per year using the commercial AVEVA Process Simulation (APS) software. For modelling details, we refer to our previous work [11], as the process design here investigated follows a very similar setup with similar assumptions.

In the process, three compressors with intermediate cooling in series compress the reactant gas from 10 bar to the specified reactor pressure. Each compressor is equipped with an anti-surge recycle that opens at feed flow rates below 60% of the nominal load to prevent surging. The reactor system is based on an AICR design with three adiabatic fixed beds and interbed cooling. Due to the exothermic reaction, the effluent from each reactor bed is hot and therefore cooled by the cold reactor feed stream. The remaining heat from the last bed is integrated to heat the reactor feed stream via the feed-effluent heat exchanger. Since single-pass conversion in the AICR system is limited by chemical equilibrium to about 30% (depending on operating conditions) [6], the reactor effluent still contains high concentrations of nitrogen and hydrogen. These gases must be recovered and recycled for efficient operation. To maximize recovery, the reactor effluent stream is further cooled and separated into two phases in a flash drum. The reactant gases in the distillate stream are recompressed to reactor pressure and recycled, while the ammonia-rich bottom stream is further purified. A second flash at lower pressure vaporizes the remaining impurities from the product, achieving 99.8 mol-% ammonia purity.

The most impactful and complex design decisions concern the selection of the AICR system's size and operating conditions. Key variables include reactor pressure, individual reactor beds' inlet temperatures, and the individual volumes of the three reactor beds. These key variables not only affect reactor conversion but also influence equipment specifications, such as compressor dimensions, reactor bed volumes, heat exchanger sizes, and, consequently, equipment costs and the required amount of utilities.

Sampling of Design Alternatives

The values of the seven previously defined key design variables must be carefully selected by the responsible design engineer, as not all possible combinations of design variable values yield feasible design solutions; certain regions of the design space may be inherently infeasible due to process constraints. This problem becomes even more complex if the design must be operable over a broad range, rather than at a single operating point, as required for green ammonia production.

To efficiently explore the design space, we define broad bounds for all key design variables (see Table 1). We assume uniform distributions within these bounds and use Latin Hypercube Sampling (LHS), a quasi-random, space-filling technique, to generate numerous design alternatives. Each alternative is a unique combination of the key design variables and is evaluated using the rigorous process model in APS. We classify the results as feasible or infeasible based on process system constraints. To account for different operating loads, we test the feasibility at both the high hydrogen load (same as nominal) and the low hydrogen load (10% of nominal).

Table 1: Base design and bounds for the HB process key design variables.

Design Variable	Base	Bounds
Reactor Pressure (p_R) in bar	218	[200, 250]
Total Reactor Vol. (V_{total}) in m^3	65	[50, 80]
Ratio Volume Bed 1 (V_{R1}/V_{total})	0.3	[0.1, 0.4]
Ratio Volume Bed 2 (V_{R2}/V_{total})	0.3	[0.2, 0.4]
Inlet Temp. Bed 1 (T_{R1}) in K	670	[620, 700]
Inlet Temp. Bed 2 (T_{R2}) in K	620	[620, 700]
Inlet Temp. Bed 3 (T_{R2}) in K	620	[620, 700]

Economic Evaluation

Several economic metrics can evaluate green ammonia production, such as net present value (NPV), payback time (PBT), minimum selling price (MSP), and levelized cost of ammonia (LCOA). NPV and PBT depend on financing, discount rate, and market prices, making them suitable for evaluating individual investments. In contrast, MSP and LCOA measure economic performance per unit, allowing comparison across projects. We use LCOA as our primary metric because it provides the average cost of ammonia production over the plant's lifetime, independent of market prices. Table 2 summarizes all prices and economic parameters used.

The LCOA is calculated using Equation 1: The annualized total capital investment (TCI) is split over the plant's lifetime using the capital recovery factor (CRF) (see Equation 2), then operational expenditures (OPEX) are added, and the result is divided by the annual ammonia production size.

$$LCOA = \frac{TCI * CRF + OPEX}{M_{NH_3,annual}} \tag{1}$$

$$CRF = \frac{r(1+r)^n}{(1+r)^n - 1} \tag{2}$$

The TCI is calculated by summing the fixed capital investment (FCI) and additional working capital (WC) costs (see Equation 3). For the FCI calculation, the costs for all process equipments are summed up and scaled by the respective 2024 Chemical Engineering Plant Cost Index (CEPCI) [12]. In the final step, we multiply this sum by the Lang factor for fluid-processing plants (see Table 2) to estimate total installation costs [13].

$$TCI = \sum(C_{\text{equip}} * \frac{\text{CEPCI}_{2024}}{567}) * \text{f}_\text{l} * (1 + \text{f}_\text{WC}) \quad (3)$$

We calculate the cost for each piece of equipment using the equations and assumptions implemented in APS. Due to the harsh conditions (high pressure and temperature) in the HB plant, we assume stainless steel as the construction material for all equipment. The compressors are modeled as centrifugal units driven by an electric motor, with an overall compression efficiency of 95%. For the heat exchangers, we assume a counter-current design and, if necessary, cooling with cooling water. To calculate the heat exchanger area, we use an overall heat transfer coefficient (value depends on the fluids, see Table 2), along with the logarithmic mean temperature difference between the tube and shell sides, and assume a minimum temperature difference of 10 K. The flash units are assumed to be vertical drums with hemispherical heads, each with a residence time of 50 min and a length-to-diameter ratio of three. The AICR system is assumed to be placed in a vertical shell with inner and outer bed diameters of 2 and 3.5 m, and a total length determined by the selected reactor bed volumes. The required thicknesses of the hemispherical heads of the AICR and flash vessels are determined in accordance with the ASME Boiler and Pressure Vessel Code, Section VIII, Division 1, UG-32(f) [14].

Table 2: Parameters for the techno-economic assessment.

Parameter	Value
CEPCI 2024 (CEPCI_{2024})	800 [12]
Lang factor (f_l)	5.93 [13]
Working capital factor (f_WC)	15% [15]
Plant lifetime (n)	25 yr
Discount rate (r)	6%
Overall heat transfer coefficient pressured gas – pressured gas	0.5 kW/m²/K
Overall heat transfer coefficient pressured gas – liquid	0.4 kW/m²/K
Price cooling water	0.35 €/GJ [15]
Price catalyst	15.36 €/kg [16]
Energy demand nitrogen	0.37 kWh/kg [17]
Catalyst lifetime	5 yr
Electricity fees and taxes	26.26 €/MWh [18]

We calculate OPEX (Equation 4) from reactants costs (hydrogen and nitrogen) and utilities (compression energy, cooling water, and catalyst replacement as a continuous cost). We use fixed prices for cooling water and catalyst per unit, and variable electricity prices to calculate hydrogen, nitrogen, and compression costs. The levelized cost of operation of hydrogen (LCOH) are calculated using the European Hydrogen Observatory calculator with default values for a PEM electrolyzer [19], splitting the cost into CAPEX, electricity, and OPEX (see Equation 5). We source electricity prices from Danish databases [20] and adjust operating times to match the plant operation as described in the next subsection.

$$OPEX = (C_{\text{H}_2} + C_{\text{N}_2}) + (C_{\text{comp}} + C_{\text{cw}} + C_{\text{cat}}) \quad (4)$$

$$C_{\text{H}_2} = \dot{m}_{\text{H}_2} * (c_{\text{H}_2,\text{CAPEX}} + c_{\text{H}_2,\text{Elec}} + c_{\text{H}_2,\text{OtherOPEX}}) \quad (5)$$

For load-flexible operation, we average OPEX across the two main operating points (high and low load), assuming instantaneous changes between. In reality, the HB plant, compared to the PEM system, is the limiting factor for transition time between operating points, and transitions are not rapid [4–6]. However, recent studies show that fast transitions of 3% load change per minute are possible [4, 8]. Further, we use representative electricity prices at the two operating loads to account for realistic costs when using load-flexible operation.

Load-Flexible Operation

Green hydrogen and ammonia production depend heavily on renewable energy supply. In practice, a green ammonia plant would most likely be built alongside electrolyzer units close to renewable electricity sources, such as solar or wind farms. However, modeling renewable electricity generation is heavily dependent on location and technology and is beyond the scope of this study. To still account for variability in renewable electricity generation, we consider dynamic electricity prices driven by regional demand and supply. In particular, we use historic hourly distributed prices from the day-ahead market in West Denmark from 2024 [20]. Figure 2 displays the mean price and the histogram of all prices, which range from slightly negative to above 800 €/MWh. With these prices we calculate LCOH and LCOA to approximate the dynamic price effect on renewable production plants.

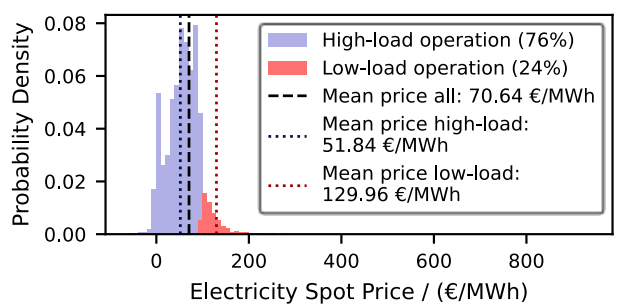

Figure 2. Electricity spot price probability distribution from hourly day-ahead market data from West Denmark in 2024 [20], including the indication of mean price and representative prices for load-flexible operation.

For constant operation, each electricity price would have the same impact on average production costs. However, with a load-flexible system, production at high prices can be reduced to a low-load operating point,

while at low prices, the production can be shifted back to the high operating point. In this study, we compare a constant operation at high load with the annual average electricity price to a load-flexible approach using two representative prices. The lower of these prices represents the lower 76% of the historic prices, while the high price represents the upper 24%.

RESULTS AND DISCUSSION

In this section, we systematically explore the design space of the HB plant to identify feasible process designs that account for different loads. Then, we evaluate the economics of constant ammonia production relative to load-flexible operation and, lastly, show how an improved HB process design can influence the total LCOA.

Finding a Feasible Process Design

The most important task when designing a load-flexible process is to identify a design that is operable over the full operating range, rather than at a single operating point. Thus, different designs must be evaluated under different loads. Since we assume that the HB process is primarily operated at two extreme load levels (high and low), we evaluate the designs at these two operating points. In total, we sample 2000 different designs and thus perform 4000 evaluations with our rigorous process model. Of the sampled designs, 1103 yield feasible solutions for both loads. Figure 3 visualizes the feasible design space by projecting it onto planes of two variables at a time and differentiating between all sampled (grey) and feasible designs (blue).

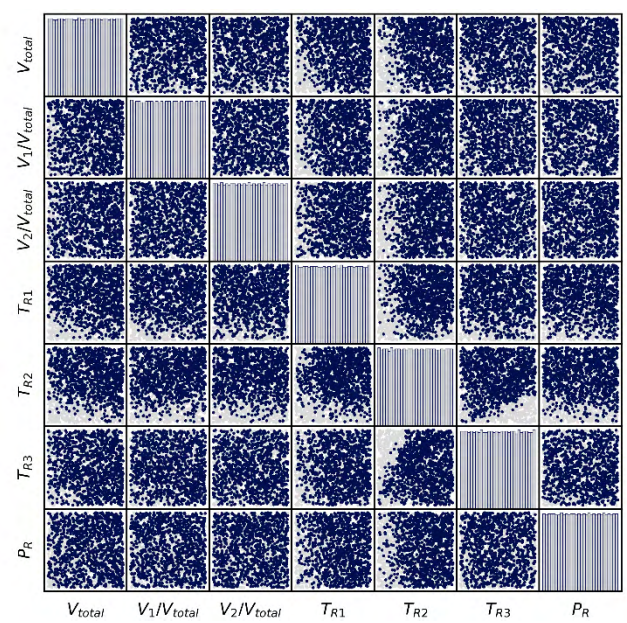

Figure 3. Input (grey) and feasible (blue) design-space projections in two-dimensional space.

The success of 55% of all sampled designs implies that a feasible design solution satisfying all process constraints at both loads is indeed possible, but not straightforward. Feasible design solutions exist for most combinations of one or two fixed design variables since only some regions of the two-dimensional projections of the full design space appear completely or largely infeasible, like very low T_{R2} values or combinations of low V_{total} and low T_{R1}. However, the contrast between the largely feasible region in plane projections and the large number of infeasible designs indicates that interactions among multiple variables are significant for feasibility. These multivariable interactions should be considered during the process design procedure. We also note that nearly all infeasibilities arise from high-load simulations, suggesting that the design at this operating point should be assessed first.

Based on the insights from the design space evaluation, we select a feasible base design near the center of the total design space for the following assessments. The values of the design variables for the base design are listed in Table 1.

Comparison of Operational Approaches

Moving forward, when designing a load-flexible plant, the most important objective is that flexible operation reduces costs relative to constant operation at high load. In the following, we compare two operational approaches for integrating green hydrogen and the HB plant based on the LCOA, using the feasible base design identified earlier. Before, in the next subsection, we replace the base design with an improved alternative.

As explained earlier, for flexible operation, the plant operates at high load (76% of the time) with lower electricity prices and at low load (24% of the time) with higher prices. Electricity prices affect green hydrogen costs and other OPEX, such as nitrogen and compression costs. The total equipment costs do not scale with the load since the nominal plant size remains unchanged, but operating costs vary. However, compressor power does not scale proportionally with the load because, for low loads, anti-surge recycling is required. Table 3 highlights key differences between the operation of the base design at high and low loads.

With above assumptions for the two operational modes, the first two bars in Figure 4 directly compare the LCOA for constant versus load-flexible HB process operation using the base design. The figure details the separate contributions of HB TCI and all OPEX components. Overall, the load-flexible approach yields an LCOA about 5.8% lower than constant operation. Comparing this LCOA to literature is challenging because key parameters, such as electrolyzer technology and electricity price, vary widely. For example, our LCOA matches the values reported by Sánchez and Martín [21] (1.35 €/kg),

which are well above current grey ammonia prices. In contrast, Hartvigsen *et al.* [22] report an LCOA of 792 €/kg for green ammonia in Denmark, using alkaline electrolyzers with far lower electricity prices for 2030.

Table 3: Comparison between high and low load operation of HB process.

	High	Low
Faction of time in operation	76%	24%
Hydrogen feed rate in t/d	250	25
Ammonia production rate in t/d	1368.6	136.81
Average electricity price + fees and taxes in €/MWh	78.10	156.22
Compression Demand in MW	26.18	15.51
Cost for Compression in €/h	2045	2423
TCI in M€	285.85	285.85

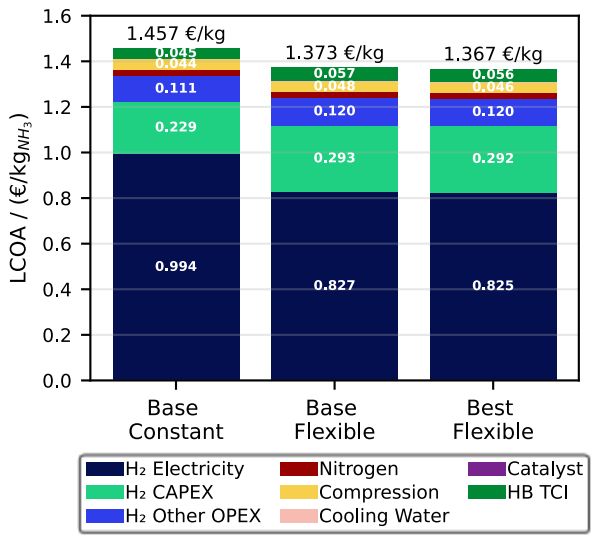

Figure 4. LCOA and its components for the three different case studies: Constant operated base design, load-flexible operated base design, and load-flexible operated best (LCOA) design.

As shown by the individual contributions, the green hydrogen price accounts for more than 91.5% of the total LCOA in the constant and 90.3% in the load-flexible case. The LCOH itself depends heavily on the price of green electricity; thus, using load-flexible operation for green hydrogen production reduces the share of electricity costs in the LCOH from 74.5% to 66.7%. This reduction is then inherently responsible for the lower LCOA when the HB process also operates flexibly. In the future, adding hydrogen storage between the PEM system and the HB process could be a case worth comparing. Introducing a buffer offers the advantage that PEM and HB process operations can be flexibilized separately, which could yield further LCOA reductions, despite high investment in the hydrogen storage system [5].

Another result highlighted by Figure 4 is that under flexible operation, the shares of CAPEX for the hydrogen generation process and the HB process increase. The higher share of equipment costs stems from the lower total output from the same plant during periods of low load. Thus, in particular, lower electrolyzer costs could significantly affect the LCOH and LCOA, especially when operated flexibly. In contrast, the effect of reducing the TCI or utility costs of the HB process itself on the LCOA is limited, accounting for only about 10% of the total LCOA. However, the ratio of hydrogen consumed to ammonia produced in the HB process could have a significant impact, and the HB design should be closely assessed in this regard.

Benefit of HB Process Design Enhancements

Having examined differences in LCOA between constant and load-flexible operation, we investigate how improvements of the HB process design could further reduce the green ammonia costs. In particular, we focus on the three most influential factors related to the HB process identified in the analysis of the LCOA components: the ratio of hydrogen feed to ammonia produced (at high load), the TCI of the HB plant, and the cost for electric compression (at both loads). Figure 5 shows the LCOA for each of the three influencing factors across all feasible designs, identified during the LHS-based design evaluation. The scatter plots show that all three factors indeed influence the total LCOA, and that, starting from the base design, minimizing each factor individually also reduces the LCOA.

The LCOA of the selected base design equals the average LCOA of all feasible designs and the standard deviation of LCOA across designs is very low (0.003 €/kg), supporting the claim that the HB process design has little influence on the total LCOA. Selecting the design with the lowest LCOA yields a 0.4% reduction compared to the base design (see Figure 4, second and third bar).

The best design alternative with the lowest LCOA (Best LCOA) also yields the lowest TCI, indicating that a focus on reducing equipment costs can have a positive overall effect. However, the design with the lowest hydrogen-to-ammonia ratio (Best HtA) and the lowest compression power has only a negligibly higher LCOA. Table 4 compares the two best designs and shows similarities in the values of the key design variables.

Both designs operate almost at the lower bound of the pressure range to minimize compression and compressor costs. This suggests that, in the future, even lower pressures could be assessed if compatible with the chosen catalyst. Recently, research has focused on so-called low-pressure HB processes using special catalysts that enable pressures below 150 bar, potentially reducing overall compression costs [2]. Further, our selected two best design candidates also almost maximize the total

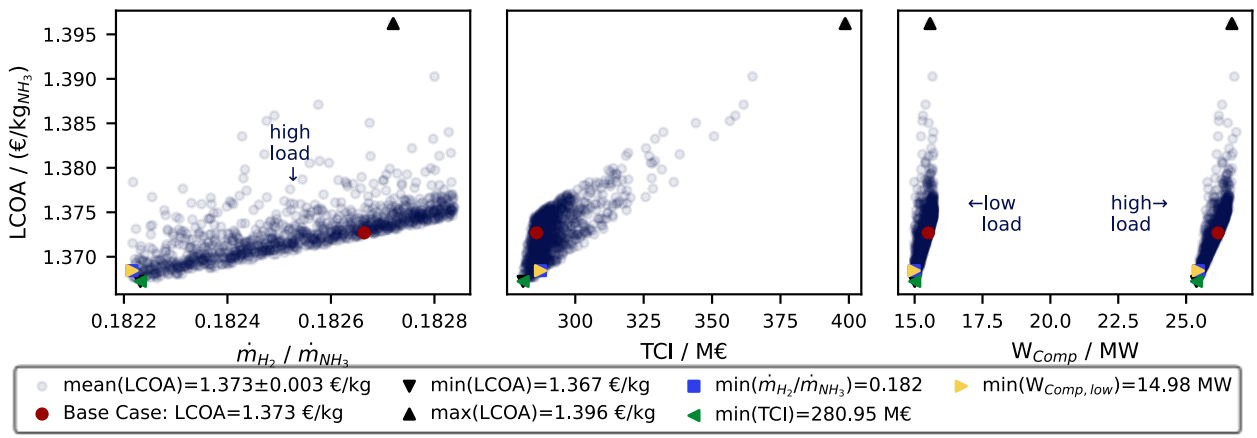

Figure 5: LCOA related to three HB process metrics: hydrogen-to-ammonia ratio, TCI, and compression demand.

reactor volume to its upper limit. As seen in Figure 4, catalyst costs affect the LCOA only minimally; thus, a higher reactor volume that improves single-pass conversion and reduces the demand for recycle stream compression could potentially lead to a reduction of the total LCOA. Lastly, comparing the volume shares among the three reactor beds and the reactor bed inlet temperatures for both design candidates does not reveal close similarities. However, the choice of reactor bed volume together with the reactor bed inlet temperature is essential for the feasibility of the design, as discussed earlier.

Table 4: Best feasible HB process designs related to the lowest LCOA and the lowest hydrogen-to-ammonia (HtA) ratio.

Design Variable	Best LCOA	Best HtA
Reactor Pressure (p_R) in bar	200.7	200.0
Total Reactor Vol. (V_{total}) in m^3	80.0	79.4
Ratio Volume Bed 1 (V_{R1}/V_{total})	0.263	0.333
Ratio Volume Bed 2 (V_{R2}/V_{total})	0.373	0.388
Inlet Temp. Bed 1 (T_{R1}) in K	654.0	638.9
Inlet Temp. Bed 2 (T_{R2}) in K	675.2	679.3
Inlet Temp. Bed 3 (T_{R2}) in K	667.3	628.5
LCOA	**1.367**	**1.368**

In summary, while the total impact of all key design variables on the LCOA could be evaluated more quantitatively by calculating sensitivity indices, the overall impact of the HB process design on the total LCOA remains small relative to the cost of green hydrogen. Thus, we find this effort is not worthwhile. The same applies to the use of optimization techniques to minimize the LCOA, treating the key design variables as decision variables.

CONCLUSION

This work addresses the central challenge of coupling ammonia synthesis to variable green hydrogen: high fluctuations in renewable electricity and, hence, stochastic hydrogen availability and price. We proposed a structured design-space exploration to identify Haber–Bosch (HB) designs that remain operable at extreme load variations. Further, we compared constant production with load-flexible operation based on the levelized cost of ammonia (LCOA). The key finding is that load-flexible operation can lower the LCOA relative to constant operation because electricity-driven hydrogen costs dominate overall economics, and shifting production toward periods with low electricity prices reduces the average green hydrogen price. By contrast, HB process design variations have only a secondary influence on the total LCOA. We indicate that a lower reactor pressure (within catalyst constraints) and a larger total reactor volume can reduce the overall LCOA by lowering compression duty and compressor sizes, yet the aggregate cost benefits of improved HB designs remain modest relative to the price of green hydrogen.

These conclusions should be interpreted in light of the scope and modeling choices. The analysis compared two steady-state load points and did not include dynamic transients. Control constraints were enforced implicitly by feasibility, not by explicit plant-wide control design. Moreover, exogenous electricity prices and renewable availability were also not modeled at hourly resolution. To address these limitations, future work should focus on rigorous modeling of renewable generation, electrolyzer dynamics, and plant controls, while accounting for time-series market and weather data. Stochastic, multi-objective assessments (cost, emissions, reliability) can then quantify when added design complexity is justified. Overall, the dominant levers for renewable ammonia lie upstream, in electricity price trajectories, electrolyzer cost, and integration strategies, while HB design refinements provide only incremental gains.

DIGITAL SUPPLEMENTARY MATERIAL

All code and simulations are available on GitHub: https://github.com/gsi-lab/FlexiblePtA_Economics

ACKNOWLEDGEMENTS

This research was funded by the European Union Horizon Europe 2022 Research and Innovation Programme under the Marie Skłodowska-Curie Grant Agreement No. 101119358 (PROSAFE).

AUTHOR IDENTIFIERS

Author ORCIDs:
Groll N: 0000-0003-0493-2270
Sin G: 0000-0003-0513-4502

REFERENCES

1. Olabi AG, Abdelkareem MA, Al-Murisi M, Shehata N, Alami AH, Radwan A, Wilberforce T, Chae KJ, Sayed ET. Recent progress in green ammonia: production, applications, assessment; barriers, and its role in achieving the sustainable development goals. Energy Conversion and Management 277:116594 (2023). https://doi.org/10.1016/j.enconman.2022.116594
2. Smith C, Hill AK, Torrente-Murciano L. Current and future role of haber–bosch ammonia in a carbon-free energy landscape. Energy Environ. Sci. 13:331-344 (2020). https://doi.org/10.1039/c9ee02873k
3. Cheema II, Krewer U. Operating envelope of haber–bosch process design for power-to-ammonia. RSC Adv. 8:34926-34936 (2018). https://doi.org/10.1039/c8ra06821f
4. Fahr S, Kender R, Bohn JP, Rehfeldt S, Peschel A, Klein H. Dynamic simulation of a highly load-flexible haber–bosch plant. International Journal of Hydrogen Energy 102:1231-1242 (2025). https://doi.org/10.1016/j.ijhydene.2025.01.039
5. Armijo J, Philibert C. Flexible production of green hydrogen and ammonia from variable solar and wind energy: case study of chile and argentina. International Journal of Hydrogen Energy 45:1541-1558 (2020). https://doi.org/10.1016/j.ijhydene.2019.11.028
6. Rosbo JW, Jensen AD, Jørgensen JB, Huusom JK. Comparison, operation and cooling design of three general reactor types for power-to-ammonia processes. Chemical Engineering Journal 496:153660 (2024). https://doi.org/10.1016/j.cej.2024.153660
7. Morud JC, Skogestad S. Analysis of instability in an industrial ammonia reactor. AIChE Journal 44:888-895 (2004). https://doi.org/10.1002/aic.690440414
8. Casale SA. FlexAMMONIA. v04.10.24 (2024)
9. Topsoe A/S. World's first dynamic green ammonia plant starts operations in Denmark. www.topsoe.com/press-releases/worlds-first-dynamic-green-ammonia-plant-starts-operations-in-denmark. Accessed 2 Feb 2026
10. Speth CH, Hultquist M, Pat AH. Method for the control of pressure in a loop for the preparation of ammonia or methanol. US Patent 2023/0137755 A1.
11. Groll N, Sin G. A benchmark simulation model of ammonia production: enabling safe innovation in the emerging renewable hydrogen economy. Systems and Control Transactions 4:1354-1359 (2025). https://doi.org/10.69997/sct.164574
12. Jenkins S. CEPCI updates. *Chemical Engineering magazine*. (2024) www.chemengonline.com/2024-cepci-updates-june-prelim-and-may-final/ Accessed 2 Feb 2026
13. Peters M, Timmerhaus K, West R. Plant Design and Economics for Chemical Engineers. 5th ed. McGraw Hill (2002)
14. Pastor T. Section viii—division 1: rules for construction of pressure vessels. Companion Guide to the ASME Boiler & Pressure Vessel Code :21-1-21-110 (None). https://doi.org/10.1115/1.859872.ch21
15. Turton R, Shaeiwitz JA, Bhattacharyya D, Whiting WB. Analysis, synthesis, and design of chemical processes. Prentice Hall (2018)
16. Bañares-Alcántara R, Dericks G, Fiaschetti M, et al. Analysis of Islanded Ammonia-based Energy Storage Systems. University of Oxford (2015)
17. Aneke M, Wang M. Potential for improving the energy efficiency of cryogenic air separation unit (ASU) using binary heat recovery cycles. Applied Thermal Engineering 81:223-231 (2015). https://doi.org/10.1016/j.applthermaleng.2015.02.034
18. Statistics Denmark. Prices of electricity for non-households by annual. www.statbank.dk. Accessed 2 Feb 2026
19. Rezaei M, Akimov A, Gray EMA. Levelised cost of dynamic green hydrogen production: a case study for australia's hydrogen hubs. Applied Energy 370:123645 (2024). https://doi.org/10.1016/j.apenergy.2024.123645
20. Stepanov S. Transformation of x-ray server from a set of www-accessed programs into www-based library for remote calls from x-ray data analysis software. Thin Solid Films 515:5700-5703 (2007). https://doi.org/10.1016/j.tsf.2006.12.011
21. Sánchez A, Martín M. Optimal renewable production of ammonia from water and air. Journal of Cleaner Production 178:325-342 (2018).

https://doi.org/10.1016/j.jclepro.2017.12.279
22. Hartvigsen C, Wenzel H, Nami H. Techno-economic comparison of green ammonia imports from chile and australia versus domestic production in denmark. Energy Conversion and Management: X 29:101541 (2026). https://doi.org/10.1016/j.ecmx.2026.101541

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A framework for dynamic rescheduling under disruptions and resource constraints

David Robins[a], Farshid Babaei[a], Joan Cordiner[a], and Solomon F. Brown[a]*

[a] The University of Sheffield, School of Chemical, Material, and Biological Engineering, Sheffield, United Kingdom
* Corresponding Author: S.F.Brown@sheffield.ac.uk

ABSTRACT

Manufacturing disruptions can be a major driving factor in the wastage of resources and delays which result in spiralling costs and cancelled orders. Operational decision making should therefore consider the potential for disruptions from as many sources as possible, encouraging improvements to operational resilience and agility. Our work presents a scheduling and rescheduling framework formulated as a rolling horizon problem for the emulation of real time decision making within a dynamically changing scenario. The framework is applied to a complex multistage problem with parallel lines susceptible to disruptions as a result of process or equipment failures, or ineffective inventory management that results in material shortages. The framework is demonstrated for a simple example case which highlights the impact of disruptions on the time taken to complete orders and the associated costs. It is observed that the inclusion of disruptions can alter equipment congestion, shifting focus for future process improvements. A scenario with intermittent raw material availability is explored, with greater mean and range for a large number of simulations performed, compared to the case with constant availability. A compounding effect is observed, whereby disruptions lead to a greater likelihood of further disruptions as machine runtimes increase and tasks are repeated. The presented framework presents a strong basis from which future works could be performed in a range of scenarios and with different operating policies.

Keywords: Scheduling, Modelling, Rolling Horizon Optimization, Supply chain, Pyomo

INTRODUCTION

Process scheduling is a key component of the field of operations management, influencing decisions surrounding the timing of tasks in order to meet objectives using finite resources. Significant research efforts have been made globally in all almost all chemical and industrial engineering sectors to achieve improvements in order to minimise costs and increase efficiency. However, the practical reality of manufacturing often differs to the ideal case, and the ability to understand, anticipate, and respond to disruptions can prove to be influential in the ability to meet objectives. In this work, we investigate the role of disruptions in the scheduling of process tasks.

One popular approach for solving larger mixed integer scheduling problems is the fix-and-relax metaheuristic, which involves the fixing of variables from previous iterations and relaxing those for future periods [1]. This allows for the partitioning of large problems while keeping the current window optimisable. Prior works utilising this approach include Brahimi [2] who use the approach to handle a large problem with optional order acceptance and flexible due dates and Bouzid et al. [3] who use it to assess order acceptance under time-of-use costs and carbon taxation.

Works investigating disruptions or changing scheduling scenarios include that of Moghaddam and Saitou [4] who utilised a dynamic pegging approach to handle the predictive and reactive rescheduling of jobs upon changes to order arrivals, due dates, cancellations, and equipment failure. Hidri and Tlija [5] investigated the role of lag times (delays between tasks on a given piece of equipment) in determining the overall makespan (time taken to complete a set list of orders) of a facility. Zhao [6] investigated the handling of material in a multistage facility with uncertain job additions. The use of material handling equipment was dynamically rescheduled in response to orders, with the aim of minimising makespan.

Zheng et al. [7] investigated a dynamic problem with processing delays and demand changes. A Bayesian Network was used to understand the probability distributions for disturbances, informing ongoing decision making.

Our work presents a novel rolling horizon scheduling framework for complex multistage processes with parallel lines. The framework incorporates uncertain disruptions to process outcomes and equipment downtime, with alternative rescheduling responses based on the type of task impacted (setup, main process, or cleaning). Resource constraints and an inventory management system (with associated reorder policy) are incorporated to assess their potential for causing disruptions. An approach inspired by the fix-and-relax metaheuristic is incorporated to minimise computational requirements whilst maintaining dynamic responses that emulate real time decision making. The outcomes of the problems considered are quantified in terms of key metrics including makespan and cost.

METHODOLOGY

Problem statement

The showcased framework is targeted at the scheduling and rescheduling of a dynamically changing process. Unanticipated orders are received at set times, with disruptions in the form of process and equipment failures. Resource constraints are also considered, with an inventory management system and operating policy. The following assumptions and assertions are made to define the problem being considered.

Scheduling

- Processes are organised into stages which can perform tasks on dedicated equipment. Tasks are performed following defined recipes to complete batches
- The overall framework is split into defined time increments for the performance of additional jobs (batch addition, disruption testing, material ordering) but tasks are scheduled continuously
- Setup tasks immediately precede main process tasks which are followed by cleaning tasks
- In the presence of storage capabilities, tasks on a given stage have a defined minimum duration. Otherwise, fixed durations are defined
- Variables for tasks completed or undergoing completion are fixed in the absence of disruptions
- Potential infeasibility following disruptions is handled through the systematic unfixing of batches until feasibility is regained

Disruptions

- Disruptions occur in the form of process failures (where quality or yield deviations occur) and equipment failures (which require maintenance to be performed)
- A defined time (relative to the task start point) during material holding in storage is used to perform testing for process failures
- Failure likelihoods are known and defined as fixed probabilities
- Known distributions are used to determine the duration of equipment downtime following failure
- A defined likelihood of recovery is used to determine if repeatable tasks (e.g. purification tasks) are repeated following process failure
- Non-recovery results in material discarding and the entire batch being rescheduled. The same procedure is used should equipment failure occur during a main processing task
- Equipment failures result in equipment being blocked until maintenance is completed
- Setup, main process, and cleaning tasks are rescheduled in the event of equipment failure during a setup task
- The main process tasks are unimpacted in the event of equipment failure during a cleaning task

Costing

- The disposal of rejected material has no associated cost
- Labour costs are incurred upon the start of a main process task
- Payments for raw materials are made at the time of reorder
- Equipment repairs have no associated costs

Orders

- Orders are received at defined points and contain the order quantity, the order due date, and the desired product
- Deliveries are always completed to their fullest possible extent (based on available inventory) at the due date. Shortfalls are marked for completion at a defined later time

Inventory management

- Sufficient unallocated material must be present in order to schedule a batch

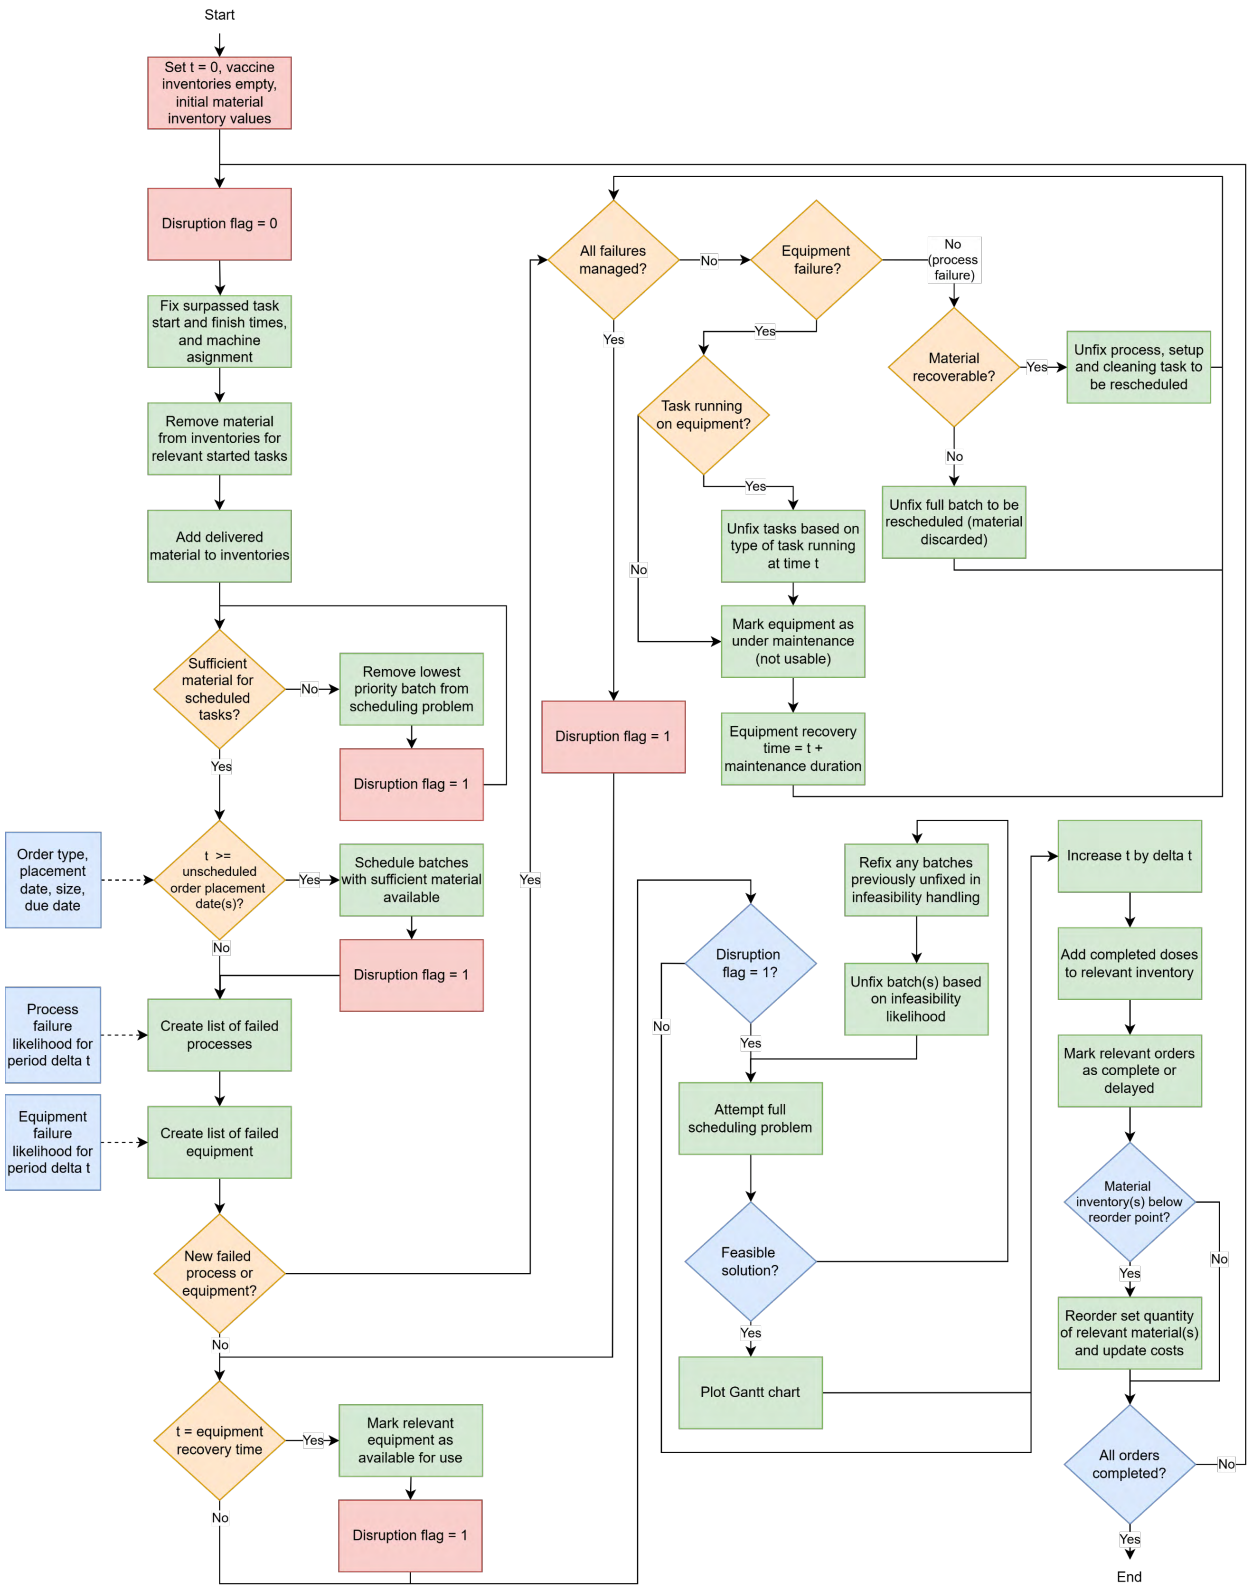

Figure 1. An overview of the presented scheduling framework used in this study

- In the case of disruption resulting in insufficient material for currently scheduled batches, the batch with the furthest due date is temporarily unscheduled

- A reorder policy is implemented based on a reorder point. Should the stored quantity of a given material reach this point, a defined quantity is reordered
- Intermittent supply refers to a case where a material is only available to reorder at defined times, unknown to the scheduler

Scheduling formulation

For brevity, the full formulation for the core scheduling problem will not be presented. A summarised version will be showcased. The key novelty of the formulation is its capacity to facilitate inputs (such as the status of a disruption) from the larger framework into the scheduling problem and allow for effective variable fixing. The objective function is the minimisation of overall makespan as shown in Equations 1. Equation 2 shows the constraint used to enforce the definition of makespan.

$$min(Makespan) \tag{1}$$

$$Makespan \geq Task\ finish\ time\ \forall\ Task\ finish\ times \tag{2}$$

The remaining constraints for the model are used to achieve the following outcomes:

1. Ensure each task is only scheduled upon a single piece of equipment
2. Ensure that the finish time of a task is equal to (for fixed duration tasks) or greater than (for variable duration tasks) the start time plus the minimum duration
3. Ensure that the start time of a successive task in a recipe for a batch is equal to the finish time of the prior task minus the duration of overlap between tasks
4. In the case of a process recovery, the prior constraint is overruled to be greater than or equal to the stated condition
5. Two tasks cannot be performed on the same piece of equipment at the same time
6. The finish times of setup tasks are equal to the start time of the associated main processing tasks, and the start time of cleaning tasks is equal to the process task finish time
7. Cleaning and setup tasks are performed on the same equipment as the main process task

Dynamic framework

An overview of the framework used to handle inventory management, order addition and completion, equipment and process disruptions, costing, and the rolling horizon model structure can be seen in Figure 1.

The framework begins with a definition of initial inventory levels followed by the main loop. The main loop first handles the fixing of variables associated with tasks completed or under completion, before managing the addition and removal of material from their associated inventories. Next, a check is performed to ensure sufficient allocated material for the currently scheduled batches, with the furthest due batch removed should current inventories be insufficient. Following this, batches from received orders that are yet to be scheduled are added to the scheduling problem, should sufficient unallocated material be available.

Process and equipment failures are determined based on their defined probabilities, creating lists of failed tasks and equipment. Should either of these occur during the current time increment, the failure handling procedure is initiated. In the event of equipment failure, currently running tasks are handled following the rules defined in the problem statement, and in all instances the equipment is marked as not usable until a later time. Should a process failure have occurred, a secondary probability is used to determine if recovery is possible, with task or full batch rescheduling as a result. Any recovered equipment is then marked as available for scheduling.

Should batch addition or removal, a process or equipment failure, or equipment recovery have occurred within the current time increment, the scheduling problem is run. Should an infeasibility occur, batches are temporarily unfixed starting with those that have been disrupted and moving onto those that are running. One batch at a time is initially unfixed, increasing once all combinations have been investigated. The first feasible solution is taken forward.

Finally, the time increment is increased, doses added to inventories, and orders completed or delayed. If inventories fall below their reorder points, those materials are reordered. This loop is repeated until all orders have been completed.

RESULTS

Demonstration of disruption impacts

In order to demonstrate the role of disruptions within the framework, a small example problem is showcased featuring the same case with and without disruptions. The example case is an eight stage process, with two pieces of parallel equipment assigned to each stage. The equipment is assigned in order, with stage 1 containing the equipment k_1 and k_2 and stage 2 containing k_3 and k_4 and so on. The stages contain a mixture of task types, with stages 1 and 1 being fixed in length, and all others containing storage and so are variable in length (above the minimum duration). Stages 3, 4, 5, and 8 represent

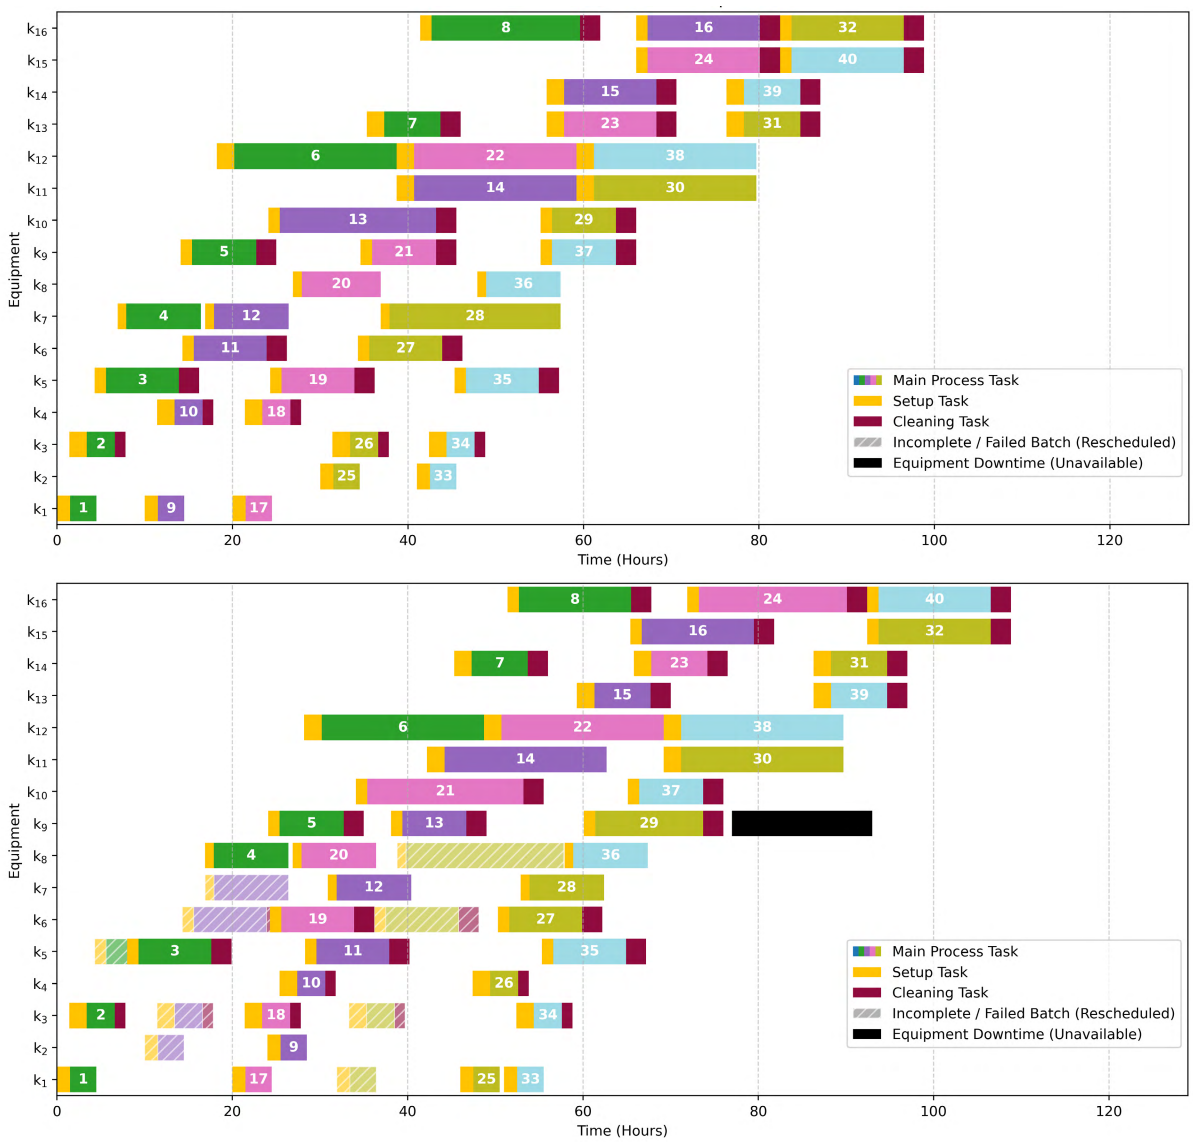

Figure 2. Gantt charts showing the scheduling of tasks for the example scenario without (above) and with (below) disruptions

purification tasks, and so have a probability of recovery following process failure. Stages 1, 4, and 6 represent processing using disposable equipment and so no dedicated cleaning is performed. The scenario shown is for the scheduling of 5 identical batches, the orders for which are received at 10 hour increments. A reorder point and reorder size of 3 batches worth of each material is used for the inventory management policy. A lead time of 10 hrs is used for the reordering of material.

Figure 2 shows the impact of disruptions on the scheduling of tasks in a singular instance of the problem. It can be seen that in the case without disruptions, the primary bottleneck is at stage 6 (equipment k_{11} and k_{12}) which sees tasks performed back-to-back. This is to be expected given that these tasks have the longest minimum duration of those within the problem. For stages 1 and 2, the use of parallel equipment is not necessary as all tasks could be performed one a singular piece of equipment without any other task being impacted. The case with disruptions demonstrates several instances of process and equipment failure. The first failure is that of task 3 at the time t=7 hrs. This task is recovered, with the remainder of the batch delayed as a result. At t=24 hrs, task 12 on equipment k_7 fails and is not recovered. This results in a full rescheduling of the batch, including the repeat of tasks 9-11. This task repetition results in a delay in the scheduling of the batch starting with task 33 at the time t=40 hrs as there is insufficient unallocated material available. At the time t=46 hrs, task 28 on equipment k_8 fails and is not recovered, resulting in a full batch reschedule in which some of the same tasks are repeated on different equipment. Upon the delivery of further

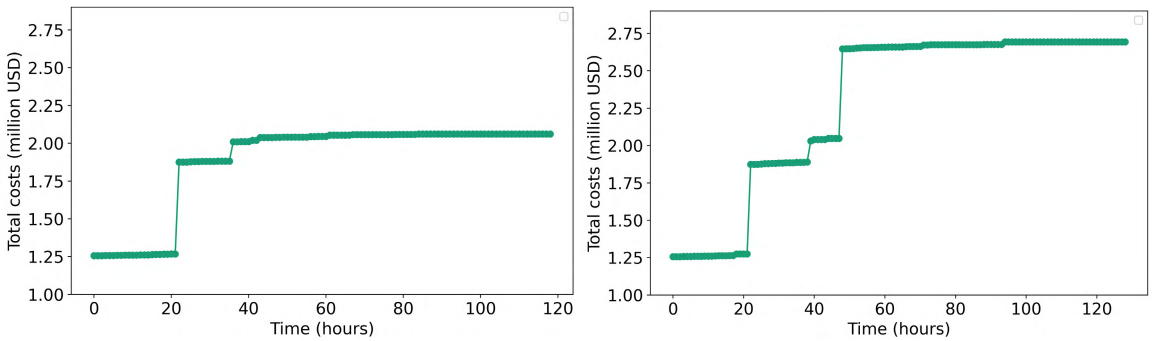

Figure 3. Cost increases over time for the example scenario without (left) and with (right) disruptions

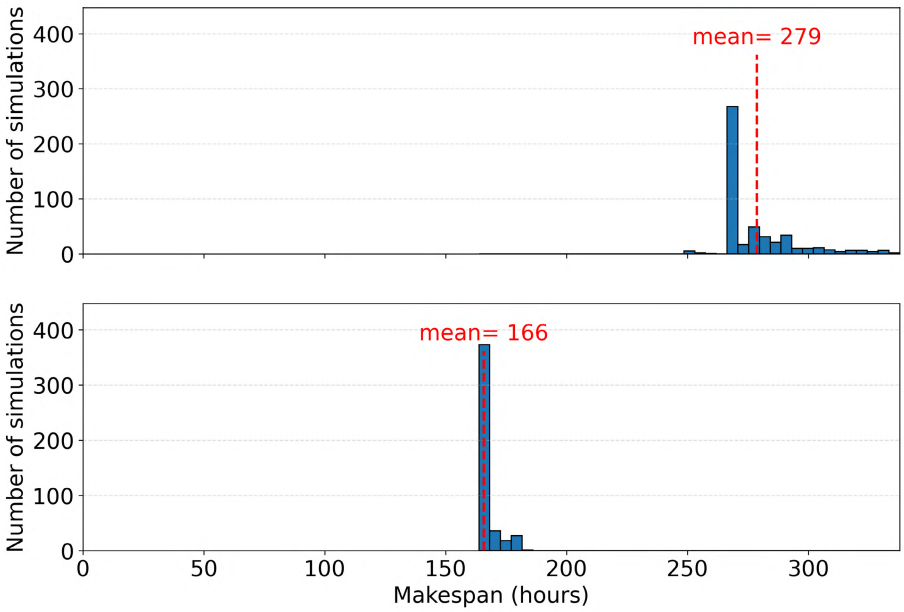

Figure 4. Makespan distributions for the case with (above) and without (below) intermittent material supply. 500 simulations were used in each case.

materials, the batch beginning with task 33 is able to be scheduled. At the time t=77 hrs, equipment k_9 fails. There is no process running and no further tasks were to be scheduled on that equipment, meaning it has no impact on overall makespan. This would not be the case in all instances of the problem.

The most noticeable impact of the incorporated disruptions is the increase in overall makespan from 98.8 hrs to 108.8 hrs. It can also be seen that the grouping of tasks changes, shifting the limiting factor from purely equipment congestion on a single stage to resource availability, changing equipment congestion, and delays themselves.

Figure 3 shows the cost in each of the same scenarios, incorporating the cost of initial raw material purchase at t=0 hrs. The case without disruptions shows that the a large increase occurs shortly after t=20 hrs. This corresponds to the purchase of materials for stage 1, which are overwhelming the highest single cost in this scenario. It can be seen that with disruptions, this reorder is repeated at t=48 hrs. This is a result of tasks being repeated, with this impact driving an overall cost increase of 30.7% as a result of disruptions.

The impact of operational policy on makespan distributions

Figure 4 shows the same case as presented previously (with disruptions) but incorporates an intermittent supply of one of the materials used in stage 6. The material is periodically available and unavailable in 8 hour increments (starting as available). To maximise the visible impact, a reorder point and quantity of 1 batches worth of material is used in both instances.

The most readily available observation is that intermittent material supply severely increases the mean makespan, representing a 68% increase in this case.

It can be seen that in both cases the distribution is right-skewed. This is partially a result of Poisson distributions used to determine equipment downtimes, but also demonstrates the compounding effect of disruptions

and breaks in production due to material unavailability. The higher frequency with which these events occur, the longer the makespan and the greater the total number of tasks completed (due to repetition following rescheduling), which creates further instances for these disruptions to occur. This is also reflected in the spread of the distributions shown, with the case with intermittent supply showing a much larger range of makespans as disruptions interact with material unavailability, causing significant delays in some simulations.

CONCLUSIONS

A dynamic framework was showcased for multi-product scheduling across parallel lines with a rolling horizon approach. The scenarios explored using the framework feature several continuously changing aspects including order receival and completion, uncertain disruptions in the form of process and equipment failures, and variations in raw material availability.

The framework was applied to a small example scenario, demonstrating its capability to emulate real time decision making in response to these aspects. The inclusion of disruptions was shown to increase the makespan in the instance shown, but also results in cost increases. Equipment congestion also changed, indicating that the disruptions may alter which parts of the process would be most desirable to receive further investment or improvement. Intermittent supply of materials was found to increase both the mean and range of makespans when applied to a large number of simulations, also demonstrating a compounding impact as a result of disruptions causing further disruptions.

Further work utilising the framework could investigate different operational policies and problem structures and case studies.

ACKNOWLEDGEMENTS

This research is funded by the Department of Health and Social Care using UK International Development funding and is managed by EPSRC. The views expressed in this publication are those of the author(s) and not necessarily those of the Department of Health and Social Care.

AUTHOR IDENTIFIERS

Author ORCIDs:
Robins D: 0009-0008-5914-3977
Babaei F: 0009-0002-2179-9723
Cordiner J: 0000-0002-9282-4175
Brown SF: 0000-0001-8229-8004

REFERENCES

1. Amiri-Aref M, Doostmohammadi M. Relax-and-fix and fix-and-optimise algorithms to solve an integrated network design problem for closing a supply chain with hybrid retailers/collection centres. Computers & Operations Research 177:106981 (2025). https://doi.org/10.1016/j.cor.2025.106981
2. Brahimi N, Aouam T, Aghezzaf EH. Integrating order acceptance decisions with flexible due dates in a production planning model with load-dependent lead times. International Journal of Production Research 53:3810-3822 (2014). https://doi.org/10.1080/00207543.2014.993045
3. Bouzid M, Masmoudi O, Yalaoui A. Exact methods and heuristics for order acceptance scheduling problem under time-of-use costs and carbon emissions. Applied Sciences 11:8919 (2021). https://doi.org/10.3390/app11198919
4. Moghaddam SK, Saitou K. A novel predictive-reactive rescheduling method for products assembly lines with optimal dynamic pegging. Computers & Industrial Engineering 171:108496 (2022). https://doi.org/10.1016/j.cie.2022.108496
5. Hidri L, Tlija M. Multi-stage hybrid flow shop scheduling problem with lag, unloading, and transportation times. PeerJ Computer Science 10:e2168 (2024). https://doi.org/10.7717/peerj-cs.2168
6. Zhao C, Fu J, Xu Q. Real?time dynamic hoist scheduling for multistage material handling process under uncertainties. AIChE Journal 59:465-482 (2012). https://doi.org/10.1002/aic.13852
7. Zheng, T., Li, D., & Li, J. (2025). Bayesian dynamic scheduling of multipurpose batch processes under incomplete look-ahead information. https://doi.org/10.48550/arXiv.2512.01093

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimization of Large-Scale Lycopene Production from Tomato Waste: A Comparative Study of Different Processing Technologies

Nereyda V. Hernández-Camacho[a], Fernando I. Gómez-Castro[a]*, Mariano Martín[b], Ehecatl A. del Rio-Chanona[c], Oscar D. Lara-Montaño[d]

[a] Universidad de Guanajuato, Departamento de Ingeniería Química, Guanajuato, Guanajuato, México.
[b] Universidad de Salamanca, Departamento de Ingeniería Química, Salamanca., España.
[c] Imperial College London, Centre for Process Systems Engineering (CPSE), Department of Chemical Engineering, United Kingdom.
[d] Universidad Autónoma de Querétaro, Facultad de Ingeniería, Juriquilla, Querétaro, México.
* Corresponding Author: fgomez@ugto.mx

ABSTRACT

For process simulation, Python and Aspen Plus can be combined to leverage the advantages of both. This work utilizes the integration of Python and Aspen Plus for the design and optimization of a lycopene production process from tomato waste. Three production pathways are studied: acetone and hexane as solvents, enzymes with ethyl acetate, and supercritical CO_2 with ethanol. This allows for the scaling of laboratory-scale studies into industrial-scale analyses. Genetic algorithms are used for optimization, enabling the determination of the optimal process design, costs, and operating conditions, while minimizing the total annual cost. The process with acetone and hexane yields a final production of 0.21 kg/h, the process with enzymes and ethyl acetate, 5.13 kg/h, and the process with supercritical CO_2 and ethanol, 0.13 kg/h. It is shown that the process with ethyl acetate has a higher production and the process with supercritical CO_2 and ethanol results in lower production and higher costs.

Keywords: Tomato waste, lycopene, solvent extraction, unconventional solvents

INTRODUCTION

Aspen Plus is a powerful tool that helps to design processing units or even complete chemical processes [1, 2, 3]. It enables rigorous designs, especially in separation systems where thermodynamics plays a key role. However, the simulator has difficulties when handling solids that are not included in its database and when there is no rigorous information about the studied compounds. Python is, on the other hand, a language with a wide variety of characteristics that made it a useful tool for solving and analyzing mathematical models of chemical processes. It allows greater freedom in coding processes for unknown compounds, but one of its drawbacks is that to have a high level of rigor in its thermodynamics and separation systems it would involve coding huge amounts of databases and some of them are not freely accessible, leading to a less rigorous modelling. There are studies that combine the use of Aspen Plus and Python for process design and optimization, but the design is typically done entirely in Aspen Plus and the optimization in Python, so they do not use a combination of both tools for design [4, 5].

This work proposes a combination of Aspen Plus and Python, seeking to leverage the advantages offered by each tool to design and optimize various lycopene production processes from tomato waste. The decision to work with tomato waste was made due to the high production in Mexico and Spain, with reports indicating that around 30% of crops end up as waste [6]. Tomatoes have a high content of lycopene, a biopigment used primarily in the food, pharmaceutical, and cosmetics industries [7]. Additionally, an increase in demand for naturally sourced lycopene has been reported, with the food and beverage sector expected to account for 45.6% of the total by 2023 [8]. Some works that have addressed the extraction of lycopene from tomato have used various techniques from traditional solvents to techniques with unconventional

solvents [9, 10, 11]. This article compares three lycopene extraction processes based on experimental studies from literature. These processes are scaled to an industrial level to determine the economic viability of these processes.

METHODOLOGY

Case studies

The first lycopene extraction process uses a mixture of acetone and hexane as solvents. This case is based on the work of Poojary et al. [11], which presents a laboratory-scale study on which the industrial-scale process design is based. This process consists of passing the tomato waste through a dryer, in which the moisture is reduced to 90%, then it passes to a mill in which a particle size of 78 mm is reached, which allows the biomass to meet with its most exposed surfaces. The process then moves to the extraction stage, where a liquid-solid extraction takes place between the biomass and an acetone-hexane mixture. The solvents extract the lycopene, separating it from the solid biomass. Following this stage, there is one outlet containing the remaining solid residue and another outlet containing a solvent mixture with the lycopene. This liquid mixture is then transferred to a flash tank, which allows for the simple separation of most of the solvents. Subsequently, the mixture is transferred to a distillation column where a lycopene purity of 95% is achieved.

The second extraction process uses enzymes and ethyl acetate and is considered a more environmentally friendly extraction method for end consumers. This process was based on experimental work by Catalkaya et al. [9], and with the help of this work, the process could be defined on an industrial scale. This process involves first passing the waste through a dryer and then through a mill with the same characteristics as the previous process. Subsequently, the waste is transferred to an enzymatic reactor, where it is mixed with a solution of water and a mixture of enzymes, specifically Celluclast (cellulolytic enzyme) and Pectinex (pectinolytic enzyme). This allows an enzymatic reaction to take place, breaking down the fibers of the biomass present and thus better preparing the raw material for extraction. The mixture then passes through a filtration centrifuge, from which the liquid fraction, including water and enzymes, is removed. The solid residue goes to the extraction stage, where it is mixed with ethyl acetate as a solvent, allowing the lycopene to be extracted. The extraction process has two outlets: the first carries the solid residue, and the second carries the lycopene-ethyl acetate mixture. Subsequently, the liquid mixture is transferred to two flash evaporation tanks. The first tank removes most of the solvent, while the second achieves a lycopene purity of 95%.

The third extraction process opts for a novel alternative that includes the use of supercritical CO_2 with ethanol. This process is based on the experimental work of Huang et al. [10]. They did help us an industrial-scale process that could be defined. In this process, tomato waste undergoes a pretreatment of drying and grinding, followed by a saponification reactor where it is mixed with 0.5 M NaOH. Subsequently, the solution passes to a second reactor where it is mixed with a 1.0 M H_2SO_4 solution, resulting in a neutralization reaction. These two reaction stages allow the biomass to break down its lignocellulosic structure in a manner like that used with enzymes in the previous section. Once neutralization is achieved, the liquid phase is removed using a filtration centrifuge, and the solid phase is transferred to the extraction stage. In this stage, a mixture of supercritical CO_2 at 25 °C and 65 bar with 16% ethanol, initially at 1 bar and 25 °C, is added. These two solvents are preconditioned using a pressure and temperature system before being mixed with the biomass, ensuring they enter the extractor under the specified conditions. The extraction stage has two outlets: one for the solids and the other for the mixture of supercritical CO_2, ethanol, and lycopene. A complete system is used to separate the three components and purify the lycopene to 95%. This system requires four main expansion tanks with intermediate turbines to perform pressure changes. The domes of the expansion tanks allow for the recovery of supercritical CO_2 and ethanol in another complete separation system.

Mathematical modeling

In all three cases, a combination of Python and Aspen Plus was used. Figures 1, 2, and 3 show the process diagrams for each case described, as well as the identification of the parts developed in Python and Aspen Plus.

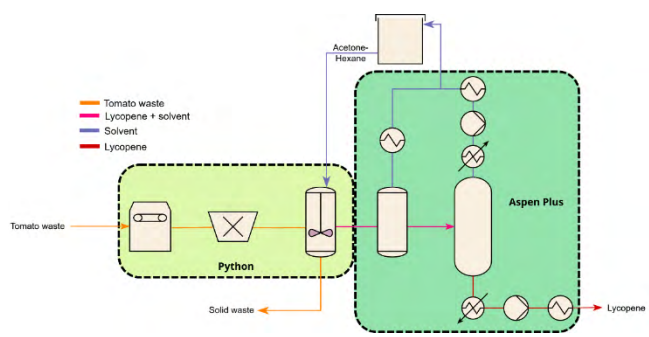

Figure 1. Distribution of process modeling with acetone-hexane in Python and Aspen Plus.

The parts developed in Python (light green) were chosen due to the possibility of mathematically modeling the equipment, mass and energy balances using design equations, due to the presence of solids in the process, and the greater flexibility offered by working with mathematical equations. The parts developed in Aspen Plus (darker green) correspond to separation and purification

processes due to their rigor in this type of equipment and ease of working with liquid mixtures.

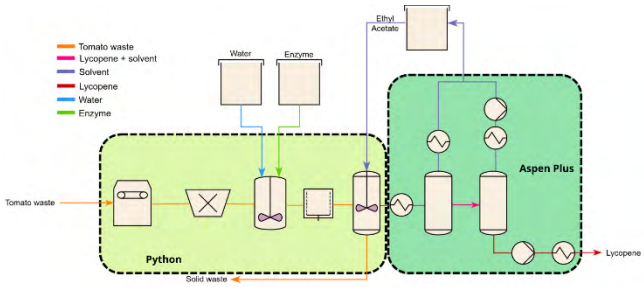

Figure 2. Distribution of process modeling with enzymes and ethyl acetate in Python and Aspen Plus.

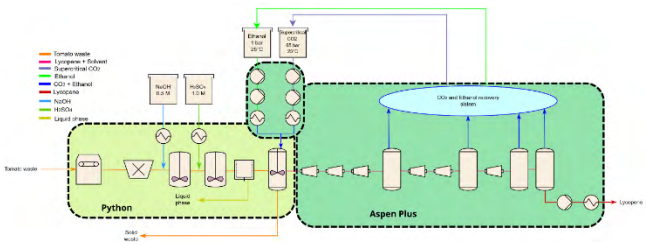

Figure 3. Distribution of process modeling with supercritical CO_2 and ethanol in Python and Aspen Plus.

Operating conditions

To define the operating and production conditions of lycopene, it was necessary to use an equation that describes the extraction stage according to its experimental functioning; equations 1, 2 and 3 represent these behaviors for the three processes mentioned [11, 9, 10].

$$\begin{aligned} Y = {} & 0.009621 - 0.0028125A + 0.0010875B - 0.0014875C \\ & + 0.00525D + 0.00045AB + 0.00265AC \\ & - 0.0022875AD - 0.000775BC \\ & + 0.0000375BD - 0.0009625CD \end{aligned} \tag{1}$$

$$\begin{aligned} Y = {} & 3.115 - 0.275X_1 + 0.125X_2 - 0.645X_3 - 0.918X_4 - 0.364X_5 \\ & - 0.112X_6 + 0.004{X_1}^2 + 0.667{X_2}^2 \\ & - 0.107{X_3}^2 + 0.279{X_4}^2 + 0.234{X_5}^2 \\ & + 0.889{X_6}^2 - 0.237X_1X_2 + 0.220X_1X_3 \\ & + 0.337X_1X_4 - 0.637X_1X_5 + 0.087X_1X_6 \\ & + 0.121X_2X_3 + 0.038X_2X_4 - 0.531X_2X_5 \\ & - 0.065X_2X_6 - 0.109X_3X_4 + 0.368X_3X_5 \\ & - 0.045X_3X_6 + 0.896X_4X_5 + 0.230X_4X_6 \\ & + 0.020X_5X_6 \end{aligned} \tag{2}$$

$$Y = 24.05 - 1.50P - 1.33P^2 - 2.64T^2 + 2.67t - 2.46PT + 1.62Tt \tag{3}$$

These equations depend on different variables that affect the amount of lycopene extracted, so they are selected as optimization variables along with other variables depending on the process. Tables 1, 2, and 3 show the optimization variables for each case study.

Table 1 includes variables for the extraction stage, the distillation column, and the dryer. Table 2 includes variables in the enzyme reactor, extractor, and dryer. Finally, Table 3 includes optimization variables only in the extraction stage and in the dryer.

Table 1: Optimization variables for the lycopene production process with acetone and hexane.

Variable	Symbol		Units	Limits
Extraction temperature	A	T_{ext}	°C	[30, 50]
Residence time	B	t	min	[1, 60]
Acetone-hexane ratio	C	A:H	%	[25, 75]
Volume of solvent mixture	D	V_{solv}	L	[10, 30]
Number of stages	-	N_t	-	[5, 20]
Feeding stage	-	N_F	-	[1, N_t]
Reflux ratio	-	RR	-	[0.1, 2]
Dryer temperature	-	T_{dry}	°C	[30, 50]

Table 2: Optimization variables for the lycopene production process with enzymes and ethyl acetate.

Variable	Symbol		Units	Limits
Enzymatic reaction temperature	X_1	T_{react}	°C	[40, 60]
Enzymatic reaction time	X_2	tR	h	[1, 5]
Enzymes/substrate ratio	X_3	E:S	v/w	[0.2, 2]
Solvent/substrate ratio	X_4	S:S	v/w	[5, 30]
Extraction time	X_5	t:E	h	[1, 4]
Enzyme/enzyme ratio	X_6	E:E	v/v	[1, 3]
Dryer temperature	T_{dry}	TS	°C	[30, 50]

Table 3: Optimization variables for the lycopene production process with supercritical CO_2 and ethanol.

Variable	Symbol		Units	Limits
Pressure	P	P_{extr}	MPa	[40, 60]
Temperature	T	T_{extr}	°C	[1, 5]
Extraction time	t	t_{extr}	h	[0.2, 2]
Dryer temperature	-	T_{dry}	°C	[5, 30]

Optimization procedure

The optimization for the three processes is executed from Python, always considering the link with Aspen Plus for the calculation of the separation section in each case. The objective function for this case is the Total Annual Cost (TAC) (Equation 4), which involves equipment costs (Equation 5) and operating costs depending on the process: Equation 6 for acetone-hexane, equation 7 for enzymes and ethyl acetate, and equation 8 for supercritical CO_2 and ethanol. such as refrigeration, heating, solvents, among others, in addition to the payback

period, which in this case was defined as 5 years.

$$TAC = CT_{oper} + \left(\frac{CT_{equip}\left(\frac{CEPCI\ 2023}{CEPCI\ 2001}\right)}{payback\ time}\right) \quad (4)$$

Where:

$$CT_{equip} = \sum_{i=1}^{7} C_{i,equip} \quad (5)$$

$$CT_{oper} = \sum_{i=1}^{n} C_{i,cooling} + \sum_{i=1}^{n} C_{i,steam} + \sum_{i=1}^{n} C_{i,elect} + \sum_{i=1}^{n} C_{i,solvent} \quad (6)$$

$$CT_{oper} = \sum_{i=1}^{n} C_{i,cooling} + \sum_{i=1}^{n} C_{i,steam} + \sum_{i=1}^{n} C_{i,elect} + \sum_{i=1}^{n} C_{i,solvent} + \sum_{i=1}^{n} C_{i,water} + \sum_{i=1}^{n} C_{i,enzym} \quad (7)$$

$$CT_{oper} = \sum_{i=1}^{n} C_{i,cooling} + \sum_{i=1}^{n} C_{i,steam} + \sum_{i=1}^{n} C_{i,elect} + \sum_{i=1}^{n} C_{i,solvent} + \sum_{i=1}^{n} C_{i,react} \quad (8)$$

The optimization was performed in Python using the PYMOO library, which works well for both single-objective, as performed in this study, and multi-objective optimizations. Genetic algorithms (GA) were employed to find the solution with the minimum total annual cost (TAC). GA was chosen for their ease of use in Python programs that link to Aspen Plus, since the equations in Aspen Plus are not explicitly defined. Figure 4 shows a flowchart that represents the general operation of the process optimization process.

RESULTS

For the first process using acetone and hexane, the genetic algorithm was employed with a population of 100 and 21 generations, resulting in a total of 2100 evaluations. This optimization required a computation time of 117.64 minutes. Figure 5 shows a comparison between the iterations and the TAC, illustrating how the optimization process seeks the minimum TAC, marked with the red line. This yielded a TAC of 305, 425.5 USD/year and a lycopene production of 0.21 kg/h.

The second process, using enzymes and ethyl acetate, required 4, 000 evaluations. This optimization required a computation time of 125.09 minutes, yielding a TAC of 6, 038, 983.34 USD/year and a lycopene production of 5.13 kg/h. Figure 7 shows the cost distribution for each piece of equipment involved in this process, allowing us to identify the stages that contribute most to the TAC. The extraction stage accounts for the largest percentage, at 51% of the TAC, followed by the enzyme reactor at 26.1%, the first flash tank at 6.6%, and the filtration centrifuge at 6.3%.

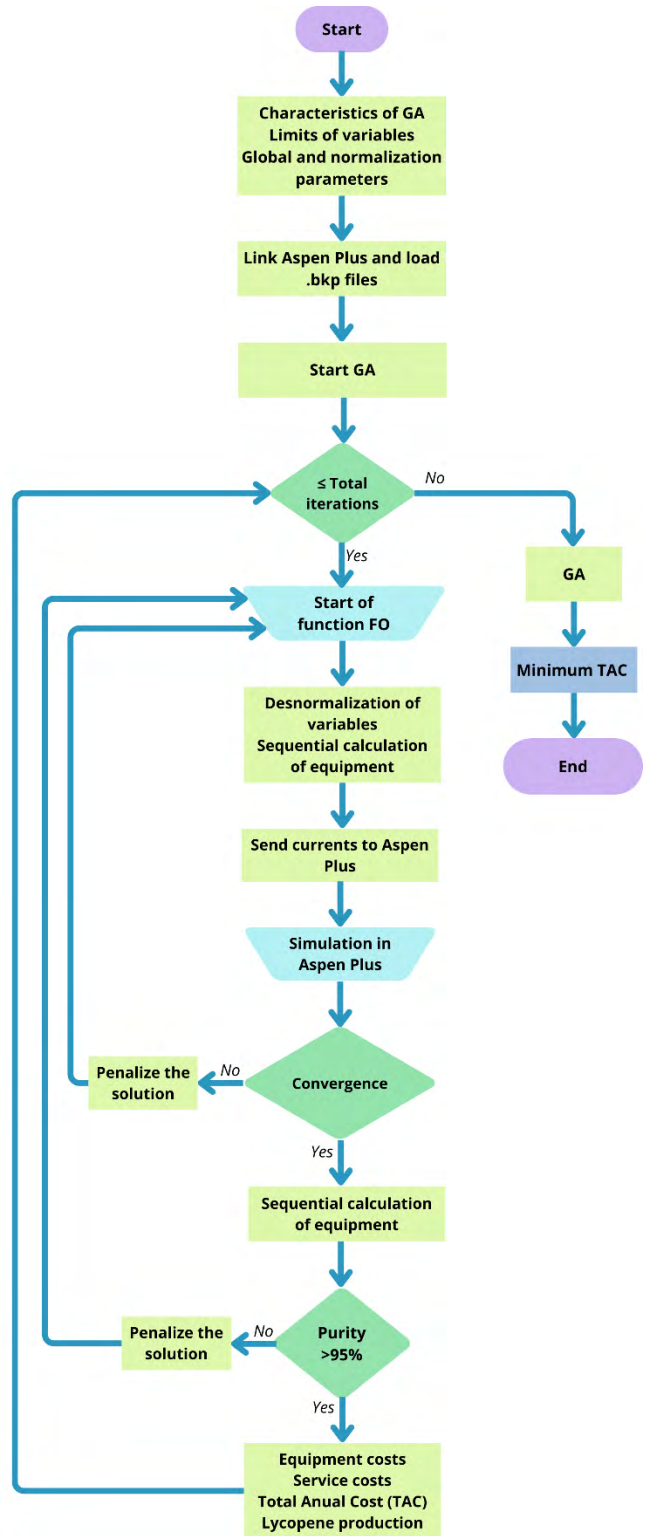

Figure 4. Optimization process flowchart.

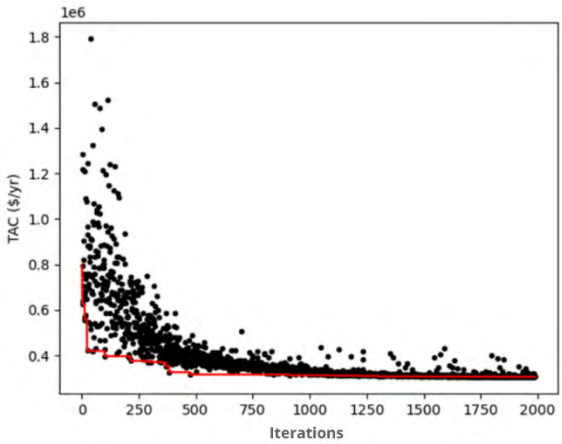

Figure 5. Comparison of iterations vs TAC for process with acetone-hexane.

Figure 6. shows a breakdown of the costs for each piece of equipment involved in this process, allowing us to identify the stages that contribute most to the total annual costs. As can be seen, the extraction stage accounts for the largest percentage at 20% of the TAC, followed by the mill and the flash tank at 19.4% and 18.1%, respectively. Additionally, heat exchangers 1 and 2 contribute approximately 10.8% of the TAC.

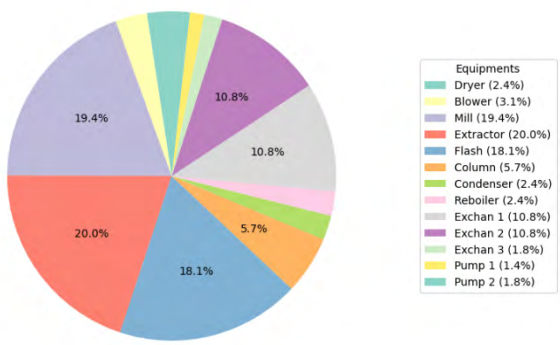

Figure 6. Percentage distribution of what each equipment contributes to the TAC.

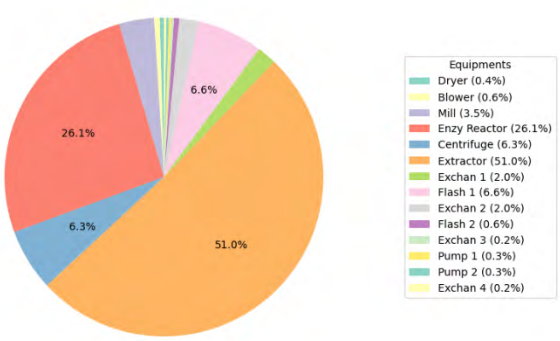

Figure 7. Percentage distribution of what each equipment contributes to the TAC.

For the third process with supercritical CO_2 and ethanol the genetic algorithm was used with a total of 2100 evaluations. This optimization required a computation time of 369.85 minutes. The TAC was 70.30 MUSD/year and a production of 0.13 kg/h of lycopene.

Figure 8 shows the cost distribution for each piece of equipment involved in this process up to the extraction stage, which corresponds to the first part of the process. This includes the cost of the entire CO_2 and ethanol separation system, allowing us to identify the stages that contribute most to the total annual costs. It can be observed that the extraction stage represents the largest percentage, with 81% of the TAC, followed by the neutralization reactor at 5.5% and the saponification reactor at 4.3%. This demonstrates the clear predominance of the extraction stage in this first part of the process due to the necessary operating conditions. Figure 9 now shows a comparison of the costs of this first stage of the process with the costs of the set of flash units, heat exchangers, pumps, and compressors that make up the ethanol and CO_2 separation and recovery section, as well as the lycopene purification. The first stage of the process accounts for most operating costs, at 62.6%, while the remainder corresponds to the second stage of the separation process, where the expansion tank section represents the largest share of these costs, at 30.8%. These graphs suggest the possibility of implementing improvements in both the extraction stage and the expansion tank separation stage.

Tables 4, 5 and 6 show the results obtained for each of the optimization variables and the objective function for the three processes studied.

Of the three processes presented, it can be observed that the one that allows the highest yield of lycopene is the process with enzymes and ethyl acetate, noting that the least viable solution is to work with supercritical CO_2 and ethanol due to the high operating costs of the system and the low production of lycopene.

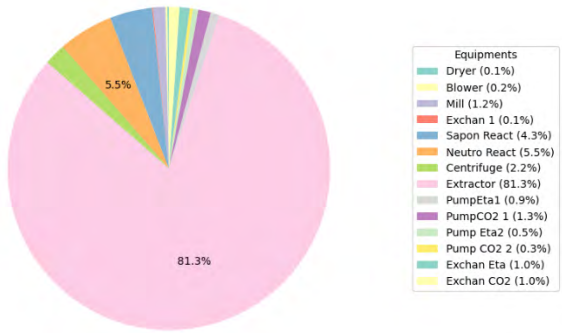

Figure 8. Percentage distribution of what each equipment contributes to the TAC, from the first part of the process.

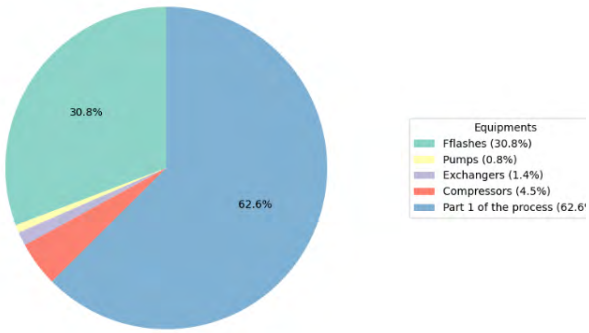

Figure 9. Percentage distribution of what each equipment contributes to the TAC, from the second part of the process.

Table 4: Results of the process optimization with acetone and hexane.

Variable	Units	Values
Operating costs	USD/yr	40, 366.58
Equipment costs	USD/yr	265, 058.86
TAC	USD/yr	305, 425.45
Lycopene production	Kg/h	0.21

Table 5: Results of the process optimization with enzymes and ethyl acetate.

Variable	Units	Values
Operating costs	USD/yr	4, 582, 107.58
Equipment costs	USD/yr	1, 456, 875.75
TAC	USD/yr	6, 038, 983.34
Lycopene production	kg/h	5.13

Table 6: Results of the process optimization with supercritical CO_2 and ethanol.

Variable	Units	Values
Operating costs	USD/yr	61, 032, 519.41
Equipment costs	USD/yr	9, 275, 715.16
TAC	USD/yr	70, 308, 234.57
Lycopene production	kg/h	0.13

CONCLUSIONS

A mathematical algorithm combining Python and Aspen Plus was developed for the design and optimization of lycopene production from tomato waste using three different extraction solvents. The process design was obtained, including the optimal operating conditions to minimize the total annual cost of the processes, allowing for the evaluation of the amount of lycopene obtainable under these conditions. The lycopene production process using enzymes and ethyl acetate yielded the best results in terms of lycopene production, at 5.13 kg/h, for a total annual cost of 6, 04 MUSD/yr. On the other hand, the least recommended process is the one that requires supercritical CO_2 and ethanol due to the high operating costs and the minimal amount of lycopene produced with 0.13 kg/h.

ACKNOWLEDGEMENTS

The authors of this work gratefully acknowledge the support of the University of Guanajuato, the University of Salamanca, Imperial College London, and the Autonomous University of Querétaro for the development of this project. N.V. Hernández Camacho gratefully acknowledges the financial support of SECIHTI through a scholarship for his doctoral studies, as well as the support of the Junta de Castilla y León research project SA085P2A, and the Asociación Universitaria Iberoamericana de Posgrado (AUIP) for his research stay at the University of Salamanca, Spain.

AUTHOR IDENTIFIERS

Author ORCIDs:
Hernández-Camacho NV: 0009-0000-0259-3696
Gómez-Castro FI: 0000-0003-4906-063X
Martín M: 0000-0001-8554-4813
Del Río-Chanona EA: 0000-0003-0274-2852
Lara-Montaño OA: 0000-0003-2055-1216

REFERENCES

1. Flores-Flores M, Gómez-Castro FI, Gutiérrez-Antonio C, Romero-Izquierdo AG, Guzmán-Martínez CE, Hernández S, Errico M. Synthesis and optimization of dividing-wall distillation columns for the separation of a quaternary mixture. Separation and Purification Technology 362:131817 (2025). https://doi.org/10.1016/j.seppur.2025.131817
2. Fonseca-Pérez RM, Almena A, Ramírez-Márquez C, Bonilla-Petriciolet A, Martín M. Techno-economic and environmental comparison of processes for the production of grape oil. Journal of Cleaner Production 441:141041 (2024). https://doi.org/10.1016/j.jclepro.2024.141041
3. Oliva-Gonzalez I, Romero-Izquierdo AG, Gutierrez-Antonio C, Gómez-Castro FI, Hernández S. Computer-Aided Desing of Intensified Separation Sequences for a Complex Mixture of Renewable Hydrocarbons. *Computer Aided Chemical Engineering* 53:1483-1488 (2024)

https://doi.org/10.1016/B978-0-443-28824-1.50248-9
4. Espinosa-Vázquez YM, Gómez-Castro FI, Sánchez-Ramirez E, Romero-Izquierdo AG. Economic optimization of intensified processes to produce bioethanol from lignocellulosic biomass. *Computer Aided Chemical Engineering* 53:2635-2640 (2024) https://doi.org/10.1016/B978-0-443-28824-1.50440-3
5. Herrera-Velázquez JJ, Cabrera-Ruiz J, Alcántara-Avila JR, Hernández-Hernández EJ, Perez-Gudiño MA, Luna-Perez R, Hernández S. Optimization of the bioethanol dehydration process using mgcl2 as mass separating agent. Chemical Engineering and Processing - Process Intensification 201:109809 (2024). https://doi.org/10.1016/j.cep.2024.109809
6. Méndez-Carmona JY, Ascacio-Valdes JA, Alvarez-Perez OB, Hernández-Almanza AY, Ramírez-Guzman N, Sepúlveda L, Aguilar-González MA, Ventura-Sobrevilla JM, Aguilar CN. Tomato waste as a bioresource for lycopene extraction using emerging technologies. Food Bioscience 49:101966 (2022). https://doi.org/10.1016/j.fbio.2022.101966
7. Kuvendziev S, Lisichkov K, Marinkovski M, Stojchevski M, Dimitrovski D, Andonovikj V. Valorization of tomato processing by-products: predictive modeling and optimization for ultrasound-assisted lycopene extraction. Ultrasonics Sonochemistry 110:107055 (2024). https://doi.org/10.1016/j.ultsonch.2024.107055
8. Kings Research. Mercado del licopeno. (2025) https://www.kingsresearch.com/es/lycopene-market-1036
9. Catalkaya G, Kahveci D. Optimization of enzyme assisted extraction of lycopene from industrial tomato waste. Separation and Purification Technology 219:55-63 (2019). https://doi.org/10.1016/j.seppur.2019.03.006
10. Huang W, Li Z, Niu H, Li D, Zhang J. Optimization of operating parameters for supercritical carbon dioxide extraction of lycopene by response surface methodology. Journal of Food Engineering 89:298-302 (2008). https://doi.org/10.1016/j.jfoodeng.2008.05.006
11. Poojary MM, Passamonti P. Optimization of extraction of high purity all-trans-lycopene from tomato pulp waste. Food Chemistry 188:84-91 (2015). https://doi.org/10.1016/j.foodchem.2015.04.133

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Auxiliary flexibility in an integrated green steel plant participating in Day-ahead and Intra-day electricity markets

Santeri Vaara[a], Iiro Harjunkoski[ab]

[a] Aalto University, Department of Chemical and Metallurgical Engineering, Espoo, Finland
[b] Hitachi Energy Research, Mannheim, Germany
* Corresponding Author: iiro.harjunkoski@aalto.fi.

ABSTRACT

In the pursuit of decarbonisation, process industries are turning to electrification as a solution to avoid fossil fuels for heating and processing raw material. Transitioning to renewable electricity couples the processes to varying electricity availability and requires more consideration for production timing and scheduling to support grid stability and avoid high electricity prices. However, practical challenges limit the capability for unforeseen rescheduling for large processes. This paper explores the idea of auxiliary flexibility in an electrified steel production process, where only the auxiliary systems can react to changing conditions. We model an H2-DRI-EAF inspired process with controllable Air-Separation unit, water electrolysis, pressurized hydrogen storage, gas liquefaction units, and a battery energy storage system to react to a production related demand delay. First, we compare hourly and 15-minute DA pricing and observe that without fast flexibility the cost difference is marginal, while fast-reacting flexibility (electrolyser ramping and batteries) enables a small additional benefit from 15-minute pricing, around 0.3% on average and up to 1% at maximum. Second, without a demand delay, allowing DA+ID decisions provides significant additional cost reduction via intraday arbitrage, yielding up to about 3% lower total costs compared to DA-only operation under perfect price knowledge. Third, when introducing 15–45-minute demand delays close to delivery, the model shows little to no cost increase on average compared to no delay. This highlights that price-taker models yield optimistic results in small markets, where open positions over 100 MW would realistically influence the market price significantly.

Keywords: Process Operations, Energy Management, Optimization, Scheduling

INTRODUCTION

The global transition from fossil fuels places process industries under pressure to decarbonize their operations. Electrification presents an key pathway for decarbonization for many processes but the increasing share of supply side volatility complicates operations [1]. At the same time, electrification intensifies the coupling between processes and grid conditions. This growing interdependence makes large industrial loads candidates for contributing to grid stability through demand-side management (DSM).

As of September 2025, the European day-ahead energy markets moved to a 15-minute market time unit (MTU), from a one-hour MTU, to improve supply-demand balance [2]. The shorter MTU requires more accurate control over consumption timing and -volume. Research on industrial DSM has focused on established power-intensive processes, such as air-separation units (ASU) [3], steel industries [4], and chlor-alkali electrolysis [5], using hourly electricity pricing. The 15-minute markets present potential for increased value of short-duration storage for electricity, raw materials, intermediate and final products and operational flexibility, compared to hourly markets.

Green steel represents one of the harder challenges for industrial decarbonization, as currently steel production accounts for around 7% of global emissions. Alternate production routes to replace the blast furnace and basic oxygen furnace (BF-BOF) process have been

presented, with different configurations of a hydrogen-based shaft furnace and electric-arc furnace (H2-DRI-EAF) are recognized as most promising future pathways [6]. However, reacting to 15-minute price change frequency is likely unrealistic for such large plants with quality-sensitive products.

Both iron- and steel-making processes consume large amounts of gases and are typically equipped with ASUs to meet this demand. Scheduling ASUs for providing DSM has been studied from multiple angles [7]. Similarly, hydrogen production via water electrolysis has received significant attention as a source of power system flexibility [8], where certain types of electrolysers can change load levels without excessive wear. For green steel production both hydrogen and ASUs present continuous processes that can dynamically adjust their load without disrupting the steel plant operations itself.

A practical challenge in electricity market participation is that day-ahead bids are submitted before the market clearing price is known, while scheduling decisions can be adjusted once prices are revealed. Simultaneous optimization of market bidding and process schedules has therefore been proposed [9], [10] to minimize imbalances and electricity costs. In 15-minute markets this becomes increasingly challenging, as predicting consumption and prices with higher accuracy is more difficult. Additionally, the ability of large processes to react to price changes is also likely limited. This means that changes in production schedules that occur after the day-ahead market position is locked in make auxiliary load adjustments potentially a practical tool for managing market position imbalances.

In this paper, we evaluate the demand-side flexibility of a green steel inspired process through its auxiliary units in 15-minute electricity markets. We model controllable electrolysers, an ASU, gas liquefaction units, and a battery energy storage system (BESS) within an integrated process. We take the electricity, hydrogen, argon, and oxygen consumption of the steel production process as exogenous demands and examine the flexibility generated by optimizing the auxiliary units under 15-minute day-ahead (DA) and intraday (ID) electricity markets. Specifically, we address three questions: (1) How does the transition from hourly to 15-minute DA pricing affect operational costs and the value of flexibility? (2) What is the potential cost reduction from participating in the ID market with a fixed day-ahead position? (3) Can auxiliary flexibility mitigate the settlement imbalance cost caused by short-term production delays? We use a deterministic scheduling approach based on the state-task-network (STN) framework, with a two-stage decision process where the first stage performs day-ahead scheduling and the second stage determines the intraday recourse to minimize the total electricity costs.

THE GREEN STEEL PROCESS

We model the gas and electricity demands based on an integrated H2-DRI-EAF process configuration, which is regarded as one of the viable routes for green steel production. In the process, raw iron ore is reduced in a shaft furnace with hydrogen forming direct-reduced iron (DRI). The DRI is then fed to a steel production process with electric arc furnaces (EAF), argon-oxygen decarburization (AOD) unit and continuous casting (CC). We consider the integrated steel production to work as an uncontrollable electricity and gas demand for controllable air separation units providing the process with nitrogen, argon and oxygen. Similarly, the water electrolysis unit that provides hydrogen for the DRI process is accounted as a controllable unit. Figure 1 shows the basic process configuration.

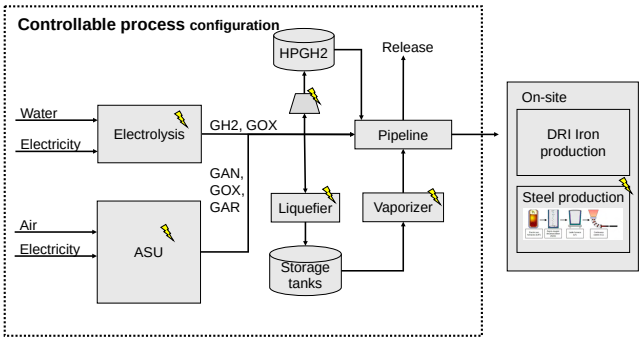

Figure 1. The modelled process configuration.

The process model is based on the STN framework presented in [11]. Each state represents a material in some specific state, and they can be converted to other states through tasks. The processes (tasks) are modeled with convex region surrogate models, allowing us to capture a wide range of process operating points and subsequent consumption of materials and electricity, while retaining convexity.

To account for short-term operational flexibility, we assume that gaseous product pipelines provide storage capacity (line-packing), corresponding to 15-minutes of average consumption [12]. Additionally, gaseous products can be converted to liquid products with a liquefaction step to be stored and again later vaporized for consumption. The gaseous hydrogen from electrolysis can be stored in a high-pressure storage unit, which requires compression. Gaseous products can be vented out of the system at any time.

For the iron and steel production processes, we consider a medium-sized integrated plant with an annual production rate of 1 Mton of steel, and base the demands for oxygen, and hydrogen from [6] and the argon and nitrogen consumption as described in [13]. We utilize the 24-heat steel plant schedule from [14], determining the demands for electricity, argon, oxygen, and nitrogen. We

assume EAF consumption is raised to full power in two steps: first at 70%, second at 90%.

The flexible ASU, liquefiers and vaporizers are parametrized based on [12] and extended such that gaseous nitrogen to oxygen ratio is 1.5:1 and argon about 5% of oxygen production. The electrolyser is modeled based on the model presented by [15], and hydrogen compression energy consumption calculated assuming adiabatic compression. We assume that of all hydrogen demand, we can produce 10% with variable load electrolysers giving us a controllable electrolyser capacity of 48.7 MW. We assume the remaining hydrogen demand is met through steady-state external supply or inflexible baseload production, isolating the impact of flexible capacity on the local energy system. The BESS is modeled separately on top of the STN model. We use a simplified model with a constant charge and discharge efficiency of 95%. For simplicity, no degradation effects are considered.

We model the BESS according to equations (1-4).

$$SOC_t = SOC_{t=1} + SOC_{t-1} + \eta^{BESS} \bullet p_t^{CH} - \frac{p_t^{DC}}{\eta^{BESS}} \quad (1)$$

$$0 \leq SOC_t \leq SOC^{max} \quad (2)$$

$$p_t^{CH} \leq P^{max}\alpha_t \quad (3)$$

$$p_t^{DC} \leq P^{max}(1-\alpha_t) \quad (4)$$

Here SOC_t depicts the state of charge, p_t^{CH} and p_t^{DC} the charge and discharge energies, respectively, for each timestep t, and η^{BESS} the charge and discharge efficiency. To limit the rate of discharge from the battery, P^{max} is used to set the maximum energy draw per time-step and α_t as an on/off binary variable to limit simultaneous charging and discharging.

MARKET PARTICIPATION MODEL

We consider two market participations for electricity procurement from the day-ahead (DA) and continuous intra-day (ID) markets. First, we procure the required energy from DA market based on demand forecast for the next day, without considering possible ID prices. This fixes a physical position that is carried over to the next stage where a possible demand forecast error is balanced by further procurement from the ID market, and where the final operational decisions are made. We utilize the real DA and ID power prices from Finland in October 2025 from Entso-e [16] and Nord-Pool continuous intra-day market [17], assuming a perfect knowledge of both prices in advance. For the hourly prices we take the average of each 15-minute price within an hour, forming a static price for each hour. It is assumed that the process can act as a price taker. We acknowledge that this assumption is a very generous one, considering the large demand of a steel plant in Finland, especially in the continuous ID market, where the average 15-minute transaction volume is in the range 150 MW for the studied period.

We use this approach to compare the effects of the transition from hourly to quarter-hourly (15-minute) DA markets to assess ability of a large process to react to faster changing prices. Additionally, we analyze the balancing capability of a plant in the case of demand forecast error. It is assumed that the steel plant acts as a balance responsible party towards the grid operator, therefore it must ensure an energy balance at each 15-minute time step, reflecting the practice in Finland. A balance responsible is expected to make power procurement from the energy markets or other contracts such that the consumed or produced energy matches the market position for each balance period. However, energy procured from DA market based on forecasts can differ drastically in the balance settlement, if there are significant errors in the consumption or production forecast if no action is taken. An imbalance in the balance settlement is penalized by fees from the TSO. We model the forecast error as a shop floor delay of 15, 30 and 45 minutes by delaying demand accordingly and forcing a power balance in the ID stage. The demand delay is shown in Figure 2.

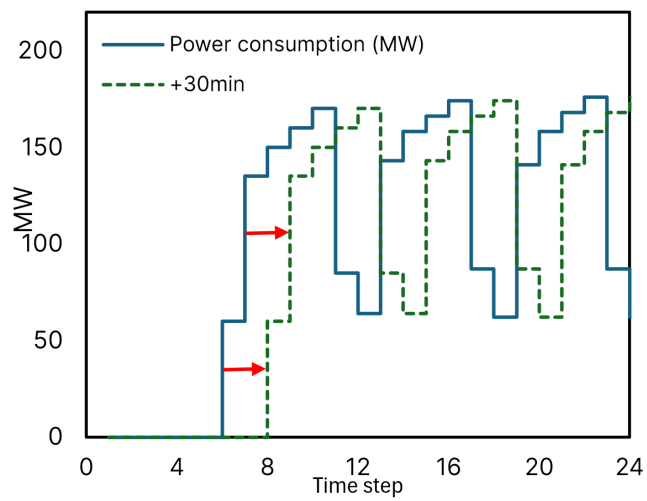

Figure 2. Demand delay for 30 minutes

$$\min \sum_{t \in T} C_t^{DA}\left(e_t^{buy} - e_t^{sell}\right) \quad (5)$$

$$s.t. e_t^{buy} - e_t^{sell} = E_t + p_t^{CH} - p_t^{DC} \quad (6)$$

The DA stage is governed with equations (5-6). Where C_t^{DA} is the DA electricity cost for each time step, e_t^{buy}, e_t^{sell} are the electricity purchase and selling variables respectively. E_t is the sum of energy consumption over all units. After solving the DA stage, we have a net electricity position, which we pass over to a second optimization stage where ID operation is solved, reacting to a possible demand delay. The ID stage is represented by equations (7-8).

$$\min \sum_{t \in T} C_t^{ID}\left(e_t^{buy} - e_t^{sell}\right) \tag{7}$$

$$s.t. e_t^{buy} - e_t^{sell} + \Phi^{DA} = E_t + p_t^{CH} - p_t^{DC} \tag{8}$$

Here ID prices are represented by C_t^{ID}, and the DA position Φ^{DA} in equation 8. Additionally, simultaneous buying and selling is limited by a binary variable, similar to eq. 3 in both stages. The intraday stage yields a net adjustment (positive or negative) to the day-ahead outcome, which is summed to form the total electricity cost over the horizon.

STUDIED CASES

We start by forming a base scenario where we set storage capacities to 0 and allow only steady state operation for the electrolyser and ASU. Similarly, we limit liquefaction and BESS capacity to 0. This gives us a baseline cost comparison between the quarterly and hourly DA prices and provides a sense of the costs a production delay might cause.

For further case studies we introduce flexibility into our system; for case 1 we allow flexible ASU operation. The control of an ASU is an intricate process, and changes in production setpoints are slow, so the introduction of ASU flexibility provides multi hour electricity consumption flexibility. The ASU is over-dimensioned by 10%, compared to the maximum oxygen demand of 20 000 Nm^3/h.

In Case 2, we introduce liquefaction and vaporization units, which allow for liquefying oxygen and nitrogen for storage and later use through vaporization.

For Case 3, we allow electrolysis units to operate flexibly with storage by giving them the ability to reduce their production rate by 100% an hour, meaning a reduction of power consumption of 2.5 MW per time-step.

Finally for Case 4, we consider BESS sizing of 5, 10, 25 and 50 MWh, with a fixed charge/discharge rate of 5 MW for 5- and 10-MWh units and 10 MW for 25- and 50-MWh units. This configuration deliberately decouples the power and energy capacities (C-rates), allowing us to analyze different storage durations and how power and energy limit capacities limit the operation.

It is important to note that this study only focuses on the operational electricity cost savings enabled by these auxiliary units. The capital investment, maintenance, and other costs associated are considered out of the scope of this work.

All cases are run for 14 separate days, with identical demands and with four demand delays of 15, 30, and 45 minutes. The cases are listed in Table 1.

Table 1: Studied cases. SS: steady state, F: flexible

Case	ASU	Lique-faction	H2 stor-age	Elec-tro-lyser	BESS (MWh/ MW)
Base	SS	no	no	SS	no
1	F	no	no	SS	no
2	F	yes	no	SS	no
3	F	yes	yes	F	no
4.1	F	yes	yes	F	5 /5
4.2	F	yes	yes	F	10 /5
4.3	F	yes	yes	F	25 /10
4.4	F	yes	yes	F	50 /10

RESULTS

15-minute pricing

We first analyze the DA costs to assess how the more frequent pricing affects our results and whether our modeled system can benefit from it. Figure 3 shows the cost distribution across the days studied in the base, steady state system and shows a very large variability between days when consumption pattern is in fact equal each day. The cost variation is caused by the variation in DA prices. Though this kind of price risk is in reality reduced by utilizing longer term contracts to control price risk, here it illustrates well how drastically electricity prices can change day-to-day costs. Since the cost variance is high, it makes sense to compare relative cost reduction. Figure 4 illustrates the relative cost difference between the mean pricings for each case and day with error bars showing the deviation between days.

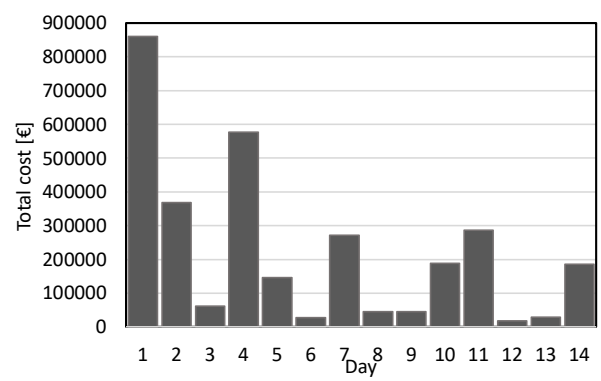

Figure 3. Absolute total cost for studied days in baseline case for 15-min pricing

We use DA to refer to only DA cost, q and h to refer to DA pricing granularity quarterly (15-min) and hourly respectively, and ID refers to DA+ID cost with a specified DA granularity and delay in minutes.

A comparison of the differences of the hourly and 15-minutes pricing scenarios for DA stage is shown in

Table 2. The cases are compared to base case, so we observe additional cost differences with the added components. With steady state operation, 15-minute pricing causes 0.09% more costs on average. This is likely caused by the chance of how pricing peaks hit consumption peaks with minimum and maximum values showing differences between extreme days. This can be viewed as an insignificant error associated with the method generating the hourly dataset. As price peaks and valleys in 15-minute data are averaged into hourly data, it is by chance how they line up with the consumption peaks, since with a continuous consumption average total cost would be equal.

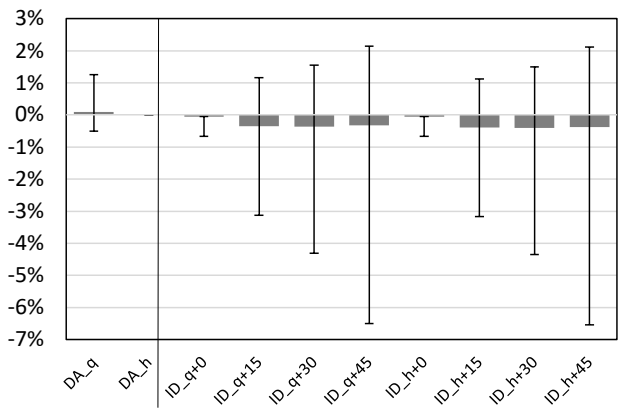

Figure 4. Relative cost difference in the baseline case. DA_q is relative to DA_h.

Table 2: Mean, minimum, and maximum cost difference of 15-min DA pricing to hourly pricing

Case	mean diff	min diff	max diff
Base	0.09%	-0.59%	1.17%
1	0.00%	0.00%	0.00%
2	0.00%	0.00%	0.00%
3	-0.20%	-0.72%	-0.02%
4.1	-0.30%	-1.00%	-0.04%
4.2	-0.29%	-0.99%	-0.04%
4.3	-0.38%	-1.28%	-0.05%
4.4	-0.37%	-1.23%	-0.06%

In Case 3, which has the capability to ramp the electrolysers, the shorter pricing interval allows a 0.20% cost reduction on average. Adding a BESS increases the gap from Case 3 but doesn't grow with added capacity until Case 4.3. Case 4.4 benefits less from the shorter pricing interval, compared to Case 4.3, which is explained by the increased ability of the case with hourly pricing to utilize the BESS capacity. We can also note that slow reacting flexibility by the ASU and liquid storage don't allow the utilization of the shorter pricing interval. This suggests that with increasing fast reacting flexibility a cost decrease can be achieved with the shorter pricing time interval. However, this cost decrease is likely insignificant in many industrial applications but might increase the viability for very fast reacting components such as BESS units.

Table 3: DA costs per studied case with difference to baseline and subsequent ID cost reduction.

Case	DA_q	ID_q+0	ID_h+0	Diff to base
Base	222397	-0.05%	-0.05%	0.00%
1	221977	-0.19%	-0.19%	-0.19%
2	221977	-0.49%	-0.49%	-0.19%
3	219923	-1.66%	-1.68%	-1.11%
4.1	219179	-2.08%	-2.20%	-1.45%
4.2	218770	-2.17%	-2.30%	-1.63%
4.3	217282	-2.75%	-2.90%	-2.30%
4.4	216589	-3.00%	-3.13%	-2.61%

Auxiliary Flexibility - Intraday Market participation

Next, we assess the capability of auxiliary systems to perform DSM activities to reduce electricity costs. First, without consideration of the demand delay. Table 3 shows that on average added flexibility can reduce costs by up to 2.61%. Furthermore, adding the ID stage where changing prices can be exploited against the DA position further reduces costs, up to about 3% of the DA cost.

We demonstrate this in figure 5 that shows how the electricity consumption and subsequent DA and ID prices affect the operating decisions of the production units for day 4 in Case 4.1. Here, DA operation reduces costs by 1.2%, compared to the baseline and then a further 1.49% with the ID operation. During the first price peak of the DA stage, the electrolyser reduces production significantly but is considerably limited due to insufficient time to fill the hydrogen storage. The ID prices on this day did not follow the trend of the DA prices and allowed continuous hydrogen production to deviate from the DA schedule and later utilize stored hydrogen to sell the position during the peak of the ID price on that day. The figure also visualizes the question of auxiliary flexibility for such a production process. It shows how massive the electricity consumption of the electric arc furnaces is compared to the ASU consumption, of about 10%, and other units even smaller. The majority of the shiftable demand is provided by the electrolyser units, which also shows the enormous electricity consumption of hydrogen production through electrolysis, as this only represents 10% of total hydrogen production required for DRI.

Managing Settlement Imbalances

Figure 6 illustrates the relative cost difference between the ID costs with no demand delay and with 15-, 30-, and 45-minute demand delay. The cases 1-4 are shown as additional difference compared to the baseline difference. Against expectation, the delay causes in

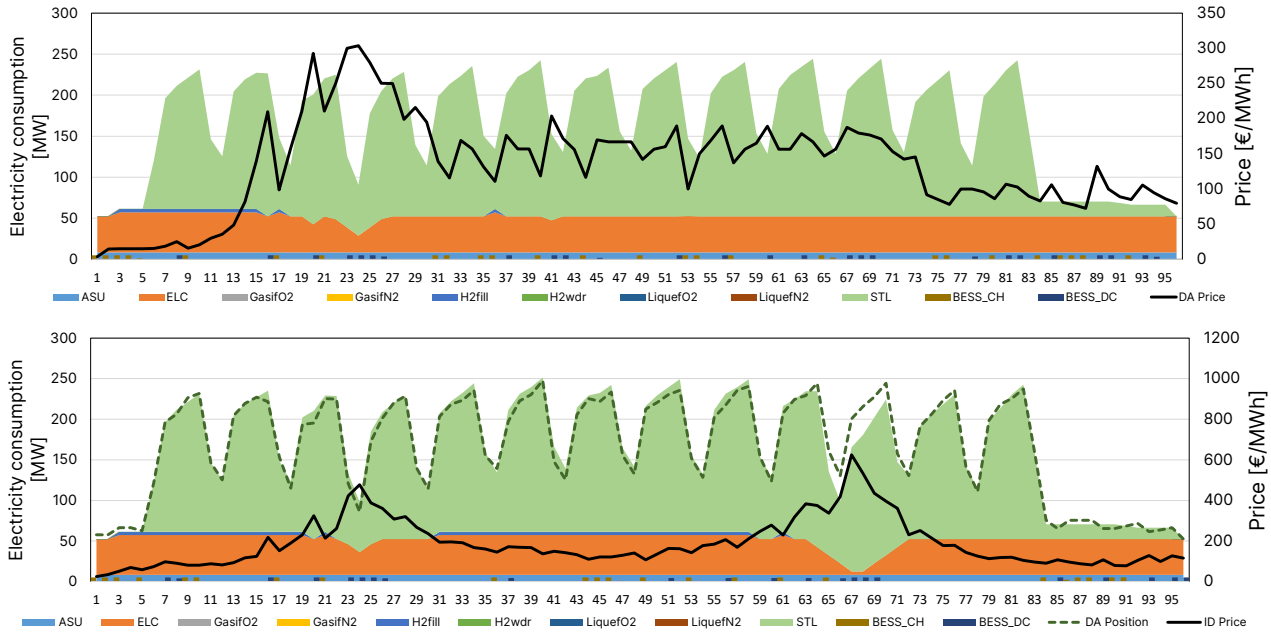

Figure 5. Day 4 Case 4.1 electricity consumption per unit. DA = upper, ID = lower.

these cases reduced costs on average, compared to the base case. On individual days, Figure 4 shows large variance with growing delay, as in some cases we see up to a 2% cost increase, for some a 7% cost decrease in the baseline case. In the baseline scenario this is caused just by how the shifted consumption peaks happen to hit the ID price each day as the shift causes a symmetric open position. For cases with flexibility, we see even more variation between days caused by the frequent negative prices in the ID market allowing the system to ramp consumption on those time points, even with no practical use. For the studied time-period, this creates on average a positive impact on the total costs. With a constant price, the additional cost caused by such a delay would amount to zero.

This behavior is unrealistic and highlights the limitation that comes with the price-taker assumption. The shifted demand frequently causes over 100 MW buy or sell bids in the ID market, which is close to the average 15-minute trading volume. This means that such a transaction volume would likely affect the price significantly against our favor causing large costs from a demand shift, incentivizing the auxiliary flexibility even more. We also assume that the imbalance position is corrected exclusively from the continuous market, whereas there are other ID markets available to handle such cascading delays.

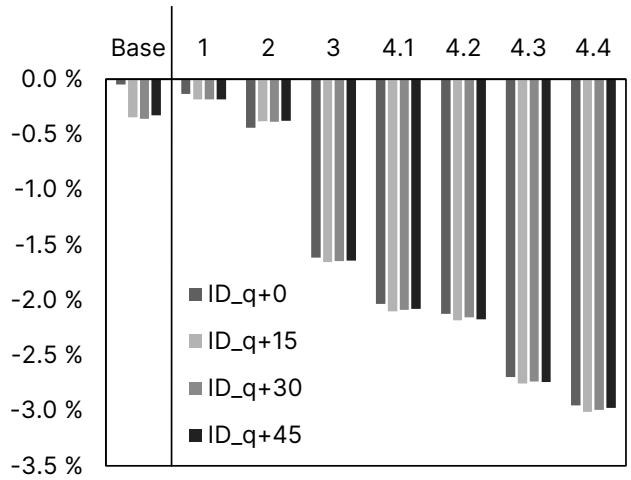

Figure 6. Average relative cost difference to DA_q costs**.** Cases 1-4 are additive to Base.

Looking at individual days, we observe that the ASU and the liquefiers are not utilized to fill the shifted position and are only turned on to increase the consumption during negative prices, which are frequent in the continuous ID market, or down during price peaks. We see in Cases 1 and 2 only little additional cost decrease compared to the base case. Similarly, while Cases 3-4 can further lower total costs, they seem to react to changes in prices rather than to the shifted demand, indicating that such a demand shift doesn't impose additional costs

that incentivize the system to correct the imbalance. This is illustrated in Figure 7, showing the ID operation of day 14 in Case 3, where the system only avoids the price spikes. Figure 7 also illustrates how a small demand delay introduces large open positions for almost 200 MW in the beginning and ends of the steel process operation period, which could occur similarly even if we perform optimal scheduling for the process for the DA markets.

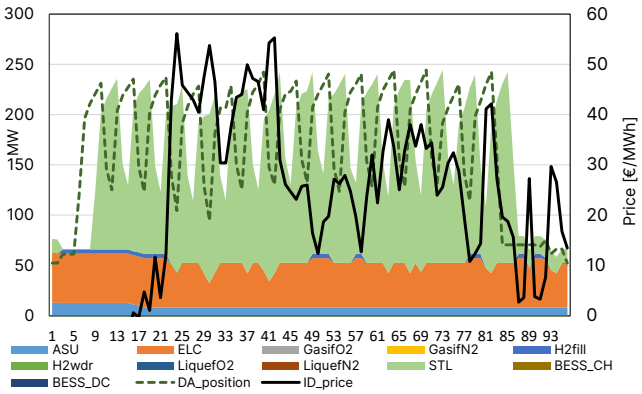

Figure 7. Day 14 Case 3 production profile

Our approach is very optimistic in the sense that we only consider the price-taker role in both markets. We also have perfect, risk-free, knowledge of upcoming prices. While the model shows a 3% saving (approx. 10 k€ on average cost), this benefit is dwarfed by the market risks illustrated in Figure 3. With a price volatility reaching 500%, optimizing an average day yields negligible benefits if the system remains exposed to daily costs of up to 800 k€. Therefore, future development should consider focusing on finding multi-day flexibility, as the ability to avoid consumption on high-priced days is likely more valuable than improving efficiency during them.

CONCLUSIONS

In this work, a green steel inspired process configuration was analyzed in the view of auxiliary flexibility for demand-side management. We considered a system with controllable electricity consumption for an ASU, nitrogen and oxygen liquefiers, electrolyser and gaseous hydrogen storage, and various BESS sizes, which we studied under different combinations.

We first quantified the relative differences that hourly and 15-minute DA pricing cause on the system total cost, and whether the more frequent pricing interval could be exploited by a flexible system to reduce operational costs. It is noted that slowly reacting flexibility cannot utilize the more frequent pricing. For components that are able to react to changing prices, we see only a small decrease in total costs (around 0.3% on average and 1% at maximum). This suggests that the 15-minute pricing might increase the viability of BESS units, with more arbitrage opportunities.

Second, we looked at how arbitrage in the ID markets can be used to decrease costs with a naive DA+ID market strategy and showed that the additional arbitrage opportunity is able to provide us with up to 3% of additional cost decrease on average.

Finally, we analyzed how a delay in the steel process between the acquired DA position and delivery can be mitigated with flexibility provided by the auxiliary units. We see that with the current setup; the demand delay causes neither additional costs nor profits on average compared with a case with no delay. This exposes the limitation in assuming a price-taker scenario as we see no cost increase on average with the delay even though it creates frequent open positions in the range of 100 MW on a market that likely has limited liquidity.

AUTHOR IDENTIFIERS

Author ORCIDs:
Vaara S: 0009-0002-0752-3915
Harjunkoski I: 0000-0003-0428-0751

REFERENCES

1. Baldea M, Endler EE, Hale E, Maravelias CT, Barolo M, Harjunkoski I, Mercangoz M, Shah SL, Soroush M, Young BR, Zhang Q. Transforming the process industries through electrification: challenges and opportunities. Ind. Eng. Chem. Res. 64:16466-16478 (2025). https://doi.org/10.1021/acs.iecr.5c01256
2. "Single Day-ahead Coupling (SDAC)." Accessed: Oct. 15, 2025. [Online]. Available: https://www.entsoe.eu/network_codes/cacm/implementation/sdac/
3. Zhang Q, Grossmann IE, Heuberger CF, Sundaramoorthy A, Pinto JM. Air separation with cryogenic energy storage: optimal scheduling considering electric energy and reserve markets. AIChE Journal 61:1547-1558 (2015). https://doi.org/10.1002/aic.14730
4. Hadera H, Harjunkoski I, Sand G, Grossmann IE, Engell S. Optimization of steel production scheduling with complex time-sensitive electricity cost. Computers & Chemical Engineering 76:117-136 (2015). https://doi.org/10.1016/j.compchemeng.2015.02.004
5. Roh K, Brée LC, Perrey K, Bulan A, Mitsos A. Flexible operation of switchable chlor-alkali electrolysis for demand side management. Applied Energy 255:113880 (2019). https://doi.org/10.1016/j.apenergy.2019.113880
6. von Schéele J. Pathways towards full use of

hydrogen as reductant and fuel. Matériaux & Techniques 111:405 (2023). https://doi.org/10.1051/mattech/2023030
7. Kelley MT, Tsay C, Cao Y, Wang Y, Flores-Cerrillo J, Baldea M. A data-driven linear formulation of the optimal demand response scheduling problem for an industrial air separation unit. Chemical Engineering Science 252:117468 (2022). https://doi.org/10.1016/j.ces.2022.117468
8. Huang C, Torres JLR, Zong Y, You S, Jin X. A review of alkaline electrolyzer technology modeling and applications for decision-making optimization in energy systems. Renewable and Sustainable Energy Reviews 224:116005 (2025). https://doi.org/10.1016/j.rser.2025.116005
9. Leo E, Dalle Ave G, Harjunkoski I, Engell S. Stochastic short-term integrated electricity procurement and production scheduling for a large consumer. Computers & Chemical Engineering 145:107191 (2021). https://doi.org/10.1016/j.compchemeng.2020.107191
10. Zhang Q, Cremer JL, Grossmann IE, Sundaramoorthy A, Pinto JM. Risk-based integrated production scheduling and electricity procurement for continuous power-intensive processes. Computers & Chemical Engineering 86:90-105 (2016). https://doi.org/10.1016/j.compchemeng.2015.12.015
11. Zhang Q, Sundaramoorthy A, Grossmann IE, Pinto JM. A discrete-time scheduling model for continuous power-intensive process networks with various power contracts. Computers & Chemical Engineering 84:382-393 (2016). https://doi.org/10.1016/j.compchemeng.2015.09.019
12. Zhou D, Zhou K, Zhu L, Zhao J, Xu Z, Shao Z, Chen X. Optimal scheduling of multiple sets of air separation units with frequent load-change operation. Separation and Purification Technology 172:178-191 (2017). https://doi.org/10.1016/j.seppur.2016.08.009
13. "Blastr Green Steel Oy, Vihreä terästehdas, Inkoo." Accessed: Dec. 11, 2025. [Online]. Available: https://www.ymparisto.fi/fi/osallistu-ja-vaikuta/ymparistovaikutusten-arviointi/blastr-green-steel-oy-vihrea-terastehdas-inkoo
14. Zhang X, Hug G, Harjunkoski I. Cost-effective scheduling of steel plants with flexible eafs. IEEE Trans. Smart Grid 8:239-249 (2017). https://doi.org/10.1109/tsg.2016.2575000
15. Varela C, Mostafa M, Zondervan E. Modeling alkaline water electrolysis for power-to-x applications: a scheduling approach. International Journal of Hydrogen Energy 46:9303-9313 (2021). https://doi.org/10.1016/j.ijhydene.2020.12.111
16. "Transparency Platform." Accessed: Jan. 24, 2026. [Online]. Available: https://transparency.entsoe.eu
17. "Nord Pool | Intraday constract statistics." Accessed: Jan. 24, 2026. [Online]. Available: https://data.nordpoolgroup.com/intraday/intraday-market-statistics?deliveryDate=2026-01-17&deliveryArea=FI

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Development of a process modeling library for the design and optimization of beverage production plants

Valentin Becher[a*], Christian Prommesberger[a], Ulrike Paap[a], Anna Afanasev[a], Anna Bechtold[a], Jörg Zacharias[a]

[a] KRONES AG, Neutraubling, Germany
* Corresponding Author: valentin.becher@krones.com.

ABSTRACT

Today, beverage production plants are planned and designed from the material-handling context as a packaged-goods production facility, not as a process plant. Therefore, a lot of potential for optimization exists. This paper presents a new approach to the design of beverage production plants according to the design of process plants. A component library for the simple creation of beverage production plant process models is developed. All steps in the plant design process can be accelerated and automated to be used for the high number of existing and new installations around the world. As first use case an energy optimization upgrade for existing Carbonated-Soft-Drink production lines is described to save cooling and heating energy in warm climates.

Keywords: Modelling and Simulations, Dynamic Modelling, Energy Efficiency, Matlab, Process Design, Simulation, Technoeconomic Analysis

INTRODUCTION

In the past, the focus for optimizing beverage production plants was on costs, production output and availability, while flexibility, energy and media consumption played a lower role. The entire basic design process was centered on material-flow considerations. A production plant is built as a combination of connected machines, but the connections and dependencies were only considered in a very simple way. Different modelling approaches are available for the process part of beverage preparation before filling and for the material-flow part of the packaged-goods streams after filling.

With the increasing sustainability focus of plant operators and other stakeholders, the pressure to optimize all aspects during the life cycle of a plant is growing. The main driver of course is still the Total Cost of Ownership (TCO) for the produced goods, but the optimization focus shifts from site to site depending on local material, energy and labor costs (**Figure 1**).

KRONES AG sells hundreds of beverage production lines worldwide each year. The number of quotations runs into thousands. The workflow for creating these quotations is very fragmented today, with multiple people involved and a multitude of separate, mostly Excel-based design tools. As a result, only very few iterations for optimization purposes can be carried out.

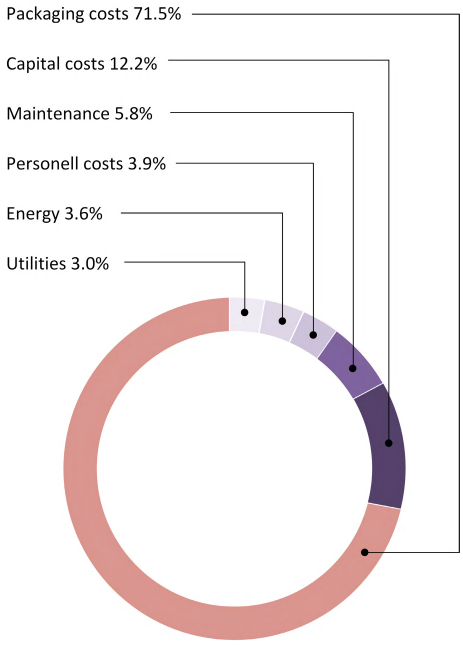

Figure 1. Typical Total-Cost-of-Ownership (TCO) distribution for a PET CSD or water line.

The reason for the splintered approach is the background of the main industry players in the manufacturing industry and the selling of single production machines. The design of lines and plants is often done by independent, specialized engineers. The result is that the know-how for the design of the machines and machine overarching dependencies is distributed between different areas of responsibility. KRONES AG itself emerged as an integrated enterprise with separate business units organized around the machine types in beverage production lines (blow moulder, filler, labeler, packing machines, washing machines, pasteurization machines and process equipment). This leads to today's situation, that design and simulation models are distributed over different organisations and exist in different levels of validity.

One promising option to address these challenges is the use of plant wide modeling and simulation [1]. At the same time, the current state of research shows few approaches for analyzing and simulating the energy and media consumption in the beverage industry, but these typically focus only on individual machines, specific subprocesses, or selected types of energy [2-6].

As the aim of KRONES AG is to sell turnkey plants and lines, an important value proposition is the possibility of deep line optimizations due to the available machine and process knowledge.

To this end, the authors started to develop a central cross-sectoral modelling library for basic design purposes of machines, lines, and plants. The library should enable the building of holistic design and optimization models for many use cases in the lifecycle of a plant by a straightforward drag-and-drop functionality.

METHODOLOGY

Definition of basic plant design process

The main design parameters for a new beverage plant are the beverage type, the nominal container volumes and the annual production schedule. The container material is known due to the existing or planned logistics, selling and recycling infrastructure. For aluminum cans and glass bottles, the containers are produced by specialized suppliers, and therefore the brand owner mostly chooses from their standard designs. In the case of PET bottles, the exact design of the containers is more flexible, as the final bottle is produced directly before filling in the production line.

The first step in the plant design process (**Figure 2**) is to design the beverage container according to the requirements during its life cycle:

1. recognition value by the final customer,
2. conservation of beverage properties (carbon dioxide concentration, no oxygen uptake to avoid oxidation reactions),
3. container stability against internal pressure and top load during storage and transport and
4. microbiologic stability

This first design step is done by the supplier of the container production machines together with the plant operator or brand owner.

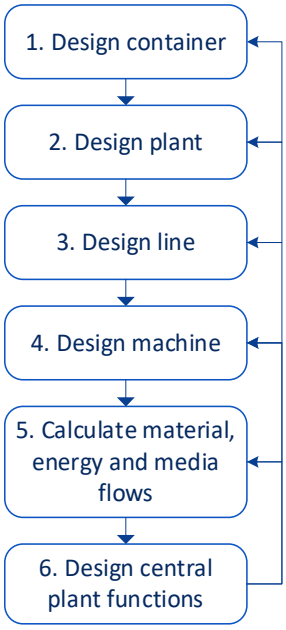

Figure 2. Basic design workflow for beverage production plants

The second step is to design the plant, i.e. what kind and number of production lines are required, what line production speeds are necessary, the sizing of the storage facilities for the produced beverages, and the logistics of the material supply and final product distribution.

In the third design step, the line is defined, i.e. type and nominal production speed of the used machines, together with the necessary buffer capacities.

The second and third design steps are done by plant design engineers and end with the specification documents for the machines and equipment.

In the fourth step, the machines and equipment itself are designed and sized. This step needs the most knowledge and can therefore best be carried out by the manufacturer.

When all equipment is known, the overall material, energy and media flows of the plant can be calculated in the fifth step again by the plant design engineers. Optimizations according to the customers wishes towards costs, energy and media consumption, flexibility or availability are possible by iteration of the preceding design workflow steps.

After the input and output flows are optimized, the central plant functions (water, energy, compressed air and other material supply, central cleaning and disposal

facilities) are designed in the next step again by the plant design engineers.

The optimization possibilities in the design process are greatly increased if the design steps and data interfaces are automatized as much as possible. The knowledge and methods of the different involved parties must be incorporated. Then a high number of iterations is possible to find the best configuration. With the help of AI-methods, the resulting data masses can be handled.

Selection of design software

The authors did a market search over the available process simulation software for building the modeling library. The software should be able to incorporate new components on a well readable engineering-equations basis without requiring deep programming knowledge. It should construct large equation systems of networks to be solved, created by drag-and-drop of components from a library. The models should be nestable in different hierarchical levels to master complexity while avoiding redundancy. New stream types with variable material properties should be definable. And it should be possible to simulate stationary as well as dynamic scenarios. The most important requirement was the possibility to create a simple and well-maintainable code base, preferably manageable by a distributed version-control system such as git or subversion.

The existence of predefined components would be beneficial but not required, as none of the software tools found were used for the intended purpose in the beverage industry. The resulting models should be able to be incorporated into PLM and other software landscapes and easy to be deployed within the company.

Open-source tools were not considered due to support and service over lifetime concerns.

In the end, MATLAB/Simscape® was selected due to cost reasons, its good availability within the company, and the fact that most engineers encounter MATLAB® during their education. The drawback of fewer predefined components and material properties was put up with. One of the useful extensions was the MATLAB App Designer combined with MATLAB® Web app server, which allows fast creation of graphical user interfaces (GUIs) and easy deployment throughout the company without requiring software licences for the end user. The result was an agile and fast development, validation and testing process for the resulting engineering design tools.

Library structure

The library consisted of physical models for simulation of physical processes such as heat transfer, gas solubility, gas permeation, mass transfer and many others. The models were implemented as custom components in the Simscape® language. Using the physical models, equipment models were created as composite components for the machines and process equipment in the beverage plant. Both physical and equipment models were centrally stored and managed in the new library (**Figure 3**).

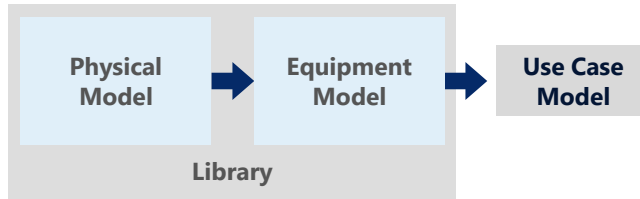

Figure 3. Modeling workflow and library structure

For each use case, a new Simulink/Simscape® model was created. The easy combination of the new components with all possibilities provided by Simulink/Simscape® toolboxes allowed the creation of digital twins for the design process in different levels of fidelity from basic design models to detailed Hardware-In-the-Loop (HIL) models. With the App Designer and standard MATLAB® coding the design workflows, the management of the design parameters and the new models were incorporated into use-case-specific GUIs.

USE CASE

The first commercial use case for the new modeling library was an energy-savings upgrade for Carbonated-Soft-Drink production lines in tropical climates [77].

Motivation

The main issue while filling carbonated beverages is foaming. The beverages are oversaturated with carbon dioxide at atmospheric pressure to achieve the fizzle mouth feeling while drinking.

The solution of the carbon dioxide into the beverage is performed at elevated pressures well above the total gas saturation pressure. The beverage containers are filled as well at pressures above the total gas saturation pressure to dissolve formed gas bubbles during the filling process and keep gas losses low. The beverage temperature is the parameter with the strongest effect on the saturation pressure, i.e. cold temperatures (5-16°C) lead to low process pressures. The container is pressurized before the filling process in the filler and vented afterwards. The foaming issues occur directly after the venting step. Small bubbles grow due to the pressure release. The bubbles grow further due to the oversaturated state of the gases in the liquid, if they are above a critical diameter [8]. Due to buoyancy effect the grown bubbles rise, and if they reach the liquid surface, they form foam. Colder filling temperatures lead to less pressure release and therefore less bubbles are above the critical diameter for further growth and foam formation.

Cold filling of carbonated beverages is state-of-the-art in the industry to handle foaming problems. The beverage is cooled down after mixing and before carbonation and filling. If the production line is in a hot and humid climate with a dewpoint above the filling temperature, the filled and capped containers are heated up again in a tunnel heater to avoid hygienic problems due to condensation (**Figure 4**).

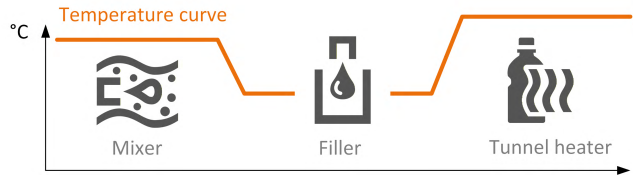

Figure 4. State-of-the-art temperature profile for Carbonated-Soft-Drink (CSD) filling in tropical climates

Historically, the filling temperature was a parameter that the plant operator required at a level based on experience, or that the filler supplier set as high as possible, also based on experience. However, there was no logical deduction for determining the best possible value.

In recent years, plant operators started to focus on energy efficiency, especially at sites with high energy prices, even if the energy costs are a small part of the TCO (**Figure 1**). The cooling and heating energy for CSD filling has come into focus, and ways to reduce it have been explored.

Process model

The simplest solution to save cooling and heating energy was to raise the filling temperature as much as possible. To avoid excessive gas losses, the filling pressure had to be raised as well and led to the limiting parameter: the maximum pressure the container can withstand. There is a design pressure for cans (6.2 bar) and glass containers (10 bar) for carbonated beverages, but not for PET containers.

PET containers are widely customized, and the design pressure is not a common parameter used during the design process. Instead, each customer has its own specifications regarding maximum fill-level drop during storage or other criteria. The so-called “thermo test” is also popular, i.e. the container is stored for 24 hours at an elevated temperature of 38°C and then inspected. If the PET container is deformed beyond the given specifications, the container must be redesigned.

The reason for the “thermo test” is that a filled and capped carbonized beverage container will reach its highest internal pressure during uncooled storage in the distribution chain from the production plant to the end customer (**Figure 5**). The gases in the liquid and the gaseous head space reach equilibrium, and the pressure in the container rises according to the dissolved gas content in the beverage (carbon dioxide, oxygen and nitrogen).

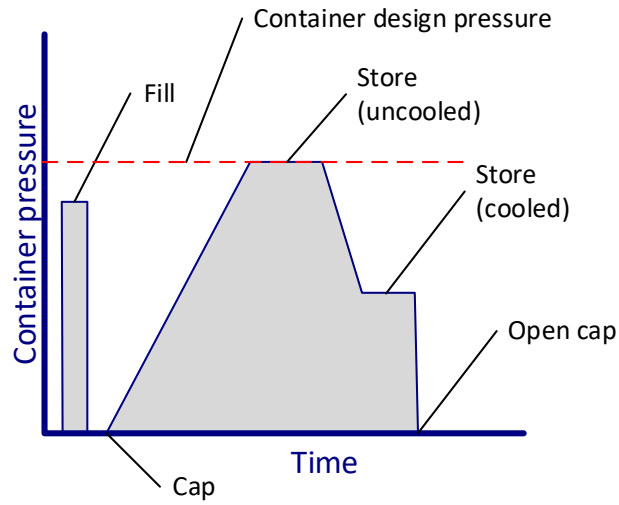

Figure 5. Schematic inner pressure vs. time diagram of a carbonized beverage container. The inner pressure in the container is increased before filling and decreased to atomospheric pressure after filling. The inner pressure raises slowly after capping due to reaching gas equilibration between head space and liquid up to saturation pressure at storage temperature.

The basis for the developed process model is the calculation of the highest pressure in the container during its lifetime. If the container cannot withstand this pressure, it will fail its main purpose. The calculated storage pressure provides the upper limit of the filling pressure, which allows the calculation of a conservative maximum filling temperature.

The process model accounts for the factors with most influence on the maximum filling pressure (**Figure 6**).

Site specific:

- Dew point for calculation of the heating energy
- Altitude for correction of the internal pressure at uncooled storage
- Maximum temperature during uncooled storage
- Product water temperature as a limit of the filing temperature
- Heating costs
- Electricity costs

Customer specific:

- CO_2 measurement method. While independent bottlers use standard devices and correlations [99], some larger companies have their own methods and correlations.

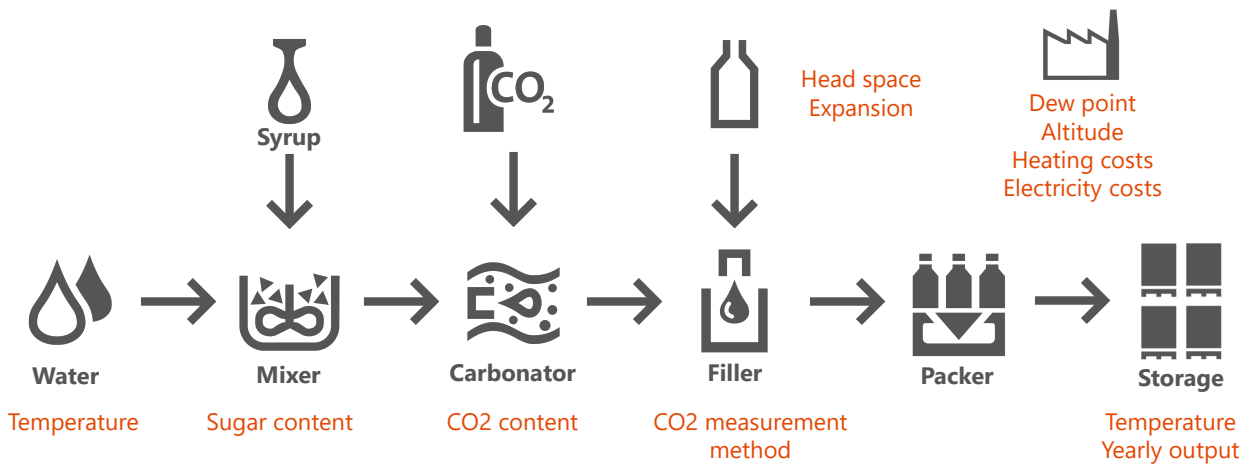

Figure 6. Process model for warm fill upgrades. Red are the site, customer and Stock-Keeping-Unit (SKU) specific model parameters.

Stock-Keeping-Unit (SKU) specific:

- Sugar content
- CO2 content
- Headspace expansion due to PET elasticity
- Yearly output

Results

With the process model the maximum feasible filing pressures could be found for each SKU at a specific production site with the limit of the calculated inner container pressure in uncooled storage and compared to the current process value (**Figure 7**).

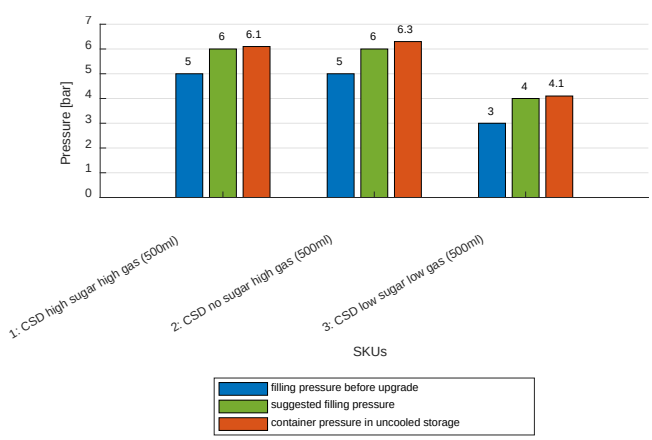

Figure 7. Pressures in PET container for different CSD beverages

The maximum filling pressure lead then to the calculation of the maximum filling temperature range depending on the unknown foaming potential of the beverage and compared again to the current process value (**Figure 8**).

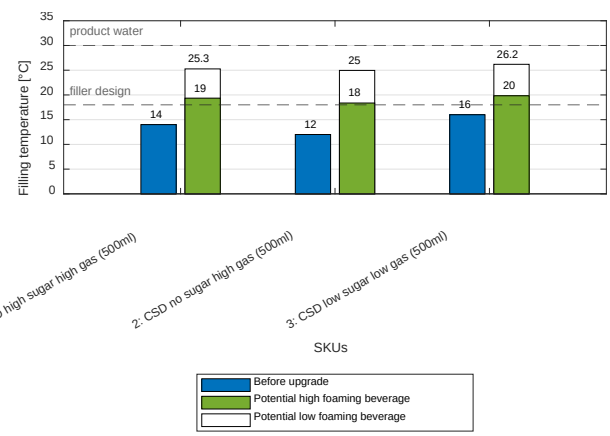

Figure 8. Filling temperature analysis for different CSD beverages.

Through the necessary cooling and heating energy, their associated costs and the yearly production out could then the annual energy costs for each SKU calculated (**Figure 9**) and the final Return-on-Invest (ROI) with the upgrade costs for each analysed production line.

The upgrade has been available for PET filling lines since the end of 2025. The first on-site implementations showed even higher possible filling temperatures and lower ROI numbers as calculated due to the fact, that especially small PET containers have higher stability against internal pressure than required, and the filling pressure could therefore be increased further than predicted by the process model. There will be a future update for the process model to incorporate the real container pressure limit based on simplified expansion.

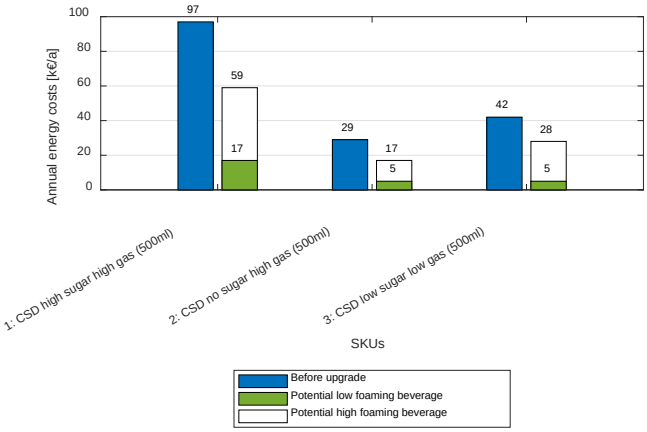

Figure 9. Annual energy costs for different CSD beverages.

CONCLUSIONS

The authors started to build up a process modelling library for the beverage industry to speed up the design process and unlock new line- and plant-wide optimization possibilities. The software MATLAB/Simscape® was selected as the tool, and the machines and equipment implemented in a structured order. The first realized use case was the creation of a process model and a tool for calculating the ROI of an energy-saving upgrade for CSD filling lines in warm climates.

AUTHOR IDENTIFIERS

Author ORCIDs:
Becher V: 0009-0005-5112-1907
Prommesberger C: 0009-0002-3664-2147
Paap U: 0000-0002-1442-0239
Afanasev A: 0009-0000-7391-6595
Bechtold A: 0009-0001-6808-5149
Zacharias J.: 0009-0006-3564-6380

REFERENCES

1. Bär RM, Voigt T. Analysis and prediction methods for energy efficiency and media demand in the beverage industry. Food Eng Rev 11:200-217 (2019). https://doi.org/10.1007/s12393-019-09195-y
2. Bär RM, Zeilmann M, Nophut C, Kleinert J, Beyer K, Voigt T. Simulation of energy and media demand of beverage bottling plants by automatic model generation. Sustainability 13:10089 (2021). https://doi.org/10.3390/su131810089
3. Hubert S.R., Unternehmensweite Modellierung und Simulation von Produktionssystemen der Prozessindustrie am Beispiel von Brauereien. PhD-Thesis, FAU University Press, Erlangen (2018)
4. Osterroth I, Klein S, Nophut C, Voigt T. Operational state related modelling and simulation of the electrical power demand of beverage bottling plants. Journal of Cleaner Production 162:587-600 (2017). https://doi.org/10.1016/j.jclepro.2017.06.006
5. Osterroth I.A., Electrical Energy Consumption of Beverage Bottling Plants: Analysis, Modeling, and Forecast. PhD-Thesis, TUM School of Life Sciences, München (2023)
6. Rädler T, Weisser H. Stochastische modellierungvon abfüllinien. Chemie Ingenieur Technik 72:627-630 (2000). https://doi.org/10.1002/1522-2640(200006)72:6<627::aid-cite627>3.0.co;2-#
7. Leserer J, Filling warmer and saving energy. KRONES magazine https://www.krones.com/en/company/press/magazine/innovation/filling-warmer-and-saving-energy.php (2025)
8. Fischer S, Blasenbildung von in Flüssigkeiten gelösten Gasen. PhD-Thesis, Technical University Munich. (2001)
9. Rammert M., Pahl M.H., Die Löslichkeit von Kohlendioxid in Getränken. Brauwelt 12:489-499 (1991)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Discrete multi-criteria optimisation of a modular heterogeneous electrolysis system

Hannes Lange[a*], Lukas Furtner[a], Michael Große[a], Isabell Viedt[a], and Leon Urbas[a]

a Technische Universität Dresden, Process-to-Order Group, Helmholtzstraße 10, 01069 Dresden, Saxony, Germany
* Corresponding Author: hannes.lange1@tu-dresden.de.

ABSTRACT

To effectively distribute power to a system of multiple electrolyzer stack units, control strategies have been developed that now need to be applied to heterogeneous electrolysis systems. These are the 'segment principle', the 'slow start principle' and the 'start-stop principle'. As there are many possible combinations to the system composition of a modular heterogeneous electrolysis system together with the most suitable control strategy, a discrete multi-criteria optimisation problem can be formulated. To solve this discrete multi-criteria optimisation problem, two discrete decision variables are introduced. One is the electrolysis system composition, represented by the power ratio/configuration (C). A total of 17 different configurations were used for this, consisting of different proportions of alkaline electrolysis (AEL) and proton exchange membrane electrolysis (PEMEL). The other one are the control strategies (R). For the control strategies, the conventional strategies, mentioned above, have been adjusted to deal with the heterogeneity of the electrolysis system. Variations are made in the priorization of certain technologies for power distribution and the availability of a base load operation. The objective function was designed to be the weighted sum of the three highest ranked decision criteria in the categories 'flexibility', 'cost', and 'process engineering' from an empirical expert survey, which led to the respective decision criteria LGC, LCOH and the reciprocal of η_{sys}. The weights were derived from the ranking, a global grid search identified the combination of (R) and (C) which will be discussed in the paper together with a detailed analysis of the optimization problem.

Keywords: Energy Systems, Hydrogen, Optimization, Multi-Criteria, Discrete, Modular Heterogeneous Systems

INTRODUCTION

The use of renewable energy leads to a volatile power supply and fluctuating energy prices, which affect energy-intensive industries (e.g., steel production). Large-scale electrolysis systems can contribute to energy flexibility by consuming electricity when demand and prices are low. Today, industrial concepts typically focus on a single electrolysis technology per electrolysis system. Alkaline electrolysis (AEL) is generally regarded as cost-effective and well-suited for base-load operation, whereas proton exchange membrane electrolysis (PEMEL) is regarded as more suitable for flexible operation. A heterogeneous plant layout could therefore combine AEL for base load and PEMEL for peak load, exploiting the strengths of each technology [1]. However, this design space, together with operational control choices, naturally results in a multi-criteria optimisation problem.

This paper aims to close a methodological gap in the design of modular heterogeneous electrolysis systems by formulating the selection of technology shares and operating concepts as a discrete multi-criteria optimization problem. While heterogeneous layouts have been discussed in prior work, the choice of technology fractions is often predefined heuristically (e.g., "reasonable" splits) or implicitly constrained by the authors' available system sizes, rather than being derived from a transparent, reproducible optimization procedure [2–4]. In contrast, we define the electrolysis configuration and the control strategy as discrete decision variables and identify trade-offs between key engineering objectives such as operational flexibility (LGC), hydrogen production cost (LCOH), and system efficiency, aggregated through expert-informed weighting and evaluated across the full

combinatorial design space.

METHOD OF DISCRETE MULTI-CRITERIA OPTIMIZATION

Definition of a use case

The application should meet several requirements. First, it should be equipped with a fluctuating power input, as continuously operated electrolysis systems have been state of the art for many years [5]. Furthermore, more than one electrolysis stack unit is required for the application in order to scale an electrolysis system consisting of several electrolysis stack units in a modular, heterogeneous way.

One application that meets the above requirements is a green steel model presented in [3]. As a boundary condition, the electrolysis system must be capable of supplying hydrogen to produce approximately 125, 000 t/h of hot briquetted iron. This corresponds to a total electrolysis capacity of approx. 400 MW. Since the largest electrolysis stack unit available on the market has a size of 5 MW, 80 such electrolysis stack units are required for a steel mill [6]. Since the LGC depends by definition on the incoming input profile, a synthetic power input profile was developed to ensure comparable power consumption and to separate the LGC from the scenario. In this way, the LGC can be applied to all electrolysis systems of different power ratings. [7].

Definition of the optimization Problem

The approach follows a decision-analytical structure: definition of (i) discrete decision variables (or criteria), (ii) objectives an the objective function that aggregates objectives into a single scalar score for ranking the proposed solutions, (iii) constraints, and (iv) the complete mathematical formulation of the optimization problem [8]. First, the design and operation of a modular heterogeneous electrolysis system are defined as discrete decision variables as follows:

- Configuration (C): The system composition, expressed by the power ratio of AEL to PEMEL. Seventeen configurations (Conf. 0–16) cover edge cases from fully PEMEL to fully AEL and include variants with finer granularity. The power ratio is introduced here to enable a comparison between systems of different sizes, for example between the large-scale electrolysis system shown here and a small laboratory system. Another application for this would be the comparability of configurations with the same system power but with individual electrolysis stack units of different sizes.
- Control strategy (R): Control strategies adapted from modular electrolysis operation, including 'start-stop', 'slow-start' and 'segment' strategies, each refined by variants to prioritise base load and switching between AEL and PEMEL [9].

The decision variables are regarded as time-invariant. The configuration (C) is not changed during operation, no electrolysis systems are connected or disconnected. The control strategy (R) is not changed during operation, as changing the control strategy requires a different system layout (in terms of sensors, actuators, etc.) [10].

Secondly, objectives are defined by which the modular heterogeneous electrolysis system is to be optimized. Multiple objectives are selected and weighted base on an empirical expert survey [11]. In this paper, these shall be the following:

- process engineering: performance, represented by the reciprocal of system efficiency ($1/\eta_{sys}$).
- Flexibility: Flexibility for dynamic operation: represented by the Lange-Große coefficient (LGC) [2, 12]
- Costs: Product costs, represented by the Levelized Cost of Hydrogen (LCOH)

Based on these objectives the objective function can be formulated as in formular (1), where (a, b, c) are weighting factors derived from the survey. In the analysed case, a = 0.50, b = 0.34 and c = 0.16.

$$Z = a * \frac{1}{\eta_{sys}} + b * LGC + c * LCOH \tag{1}$$

Thirdly, the constraints are defined. These are to be defined for the LCOH for green hydrogen at 7.7€/kg_{H2}, as this corresponds to the average value of currently estimated costs by project developers in the field of green hydrogen [13, 14]. The LCG is limited to a maximum of 2%, as already described in the literature. System efficiency is not to be further restricted [15].

In the fourth step, the optimization problem M, which is represented in formula (2), is defined. The proposed solutions are evaluated using a weighted, normalized sum with ideal point normalization.

$$M = \begin{cases} \min\limits_{C\&R} (a * \frac{1}{\eta_{sys}} + b * LGC + c * LCOH) \\ \quad s.t.\ LGC \leq 2\% \\ \quad s.t.\ LCOH \leq 7.7€/kg_{H_2} \end{cases} \tag{2}$$

This optimisation problem M is solved by a comprehensive grid search across the discrete space, combining all configurations C with all control strategies R. After applying the constraints, the remaining solutions are evaluated using the objective function and visualised as a heat map to determine the weighted score for each (C, R) pair.

The illustration Figure 1 shows the basic structure of the optimisation framework. This shows the two discrete decision variables (C) and (R) (different shapes), which in

this case consist of only two different discrete points (white and black). Each discrete point of a decision variable is then combined with all other decision variables, resulting in four different discrete points in this case. Z can now be calculated for each combination of the discrete decision variables. The values η_{sys}, LGC, and LOHC can then be entered in a three-dimensional diagram. The two constraints can now be entered on the axes for LGC and LOHC (blue lines). The green point at the origin of the coordinates represents the ideal point if all objectives were perfectly reached at the same time.

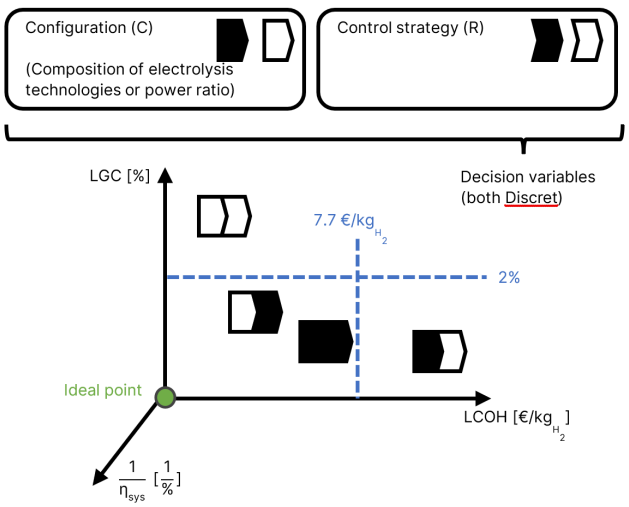

Figure 1. Schematic representation of the optimisation framework

Description of the simulation models

The behavioural models used were developed to determine the operability of connected heterogeneous electrolysis systems [12, 16]. They simulate the lumped, dynamic behaviour of a connected electrolysis system, including stack-units, electrolyte circulation, gas separations and water supply. Internal control strategies enable the evaluation of cold start-up, warm start-up, stationary operation, load changes and shut-down to accurately predict the dynamic behaviour and power consumption of the systems. Operational setpoints in the models, such as temperature, flowrate and pressure were held constant across all simulation runs. The models were validated using literature data from [17].

Table 1 provides an overview of the stack unit parameters used, including their source in the literature. All of these parameters are necessary included in behavioural models, that form the basis for calculating the selected objective of system efficiency, system flexibility and the levelized cost of hydrogen for each combination of (C) and (R).

Table 1: List of stack unit parameters used for the simulation runs

			AEL	PEMEL	Source
Con. 0-16	Efficiency@100% Load	[%]	0.6345	0.5837	[6, 9]
	Start-Up (Cold)	[%/s]	~0, 175	~0, 25	[6, 9, 17]
	Flexibiliy (Ramp-Up)	[%/s]	~0, 25	~5	[6, 9, 17]
	Flexibiliy (Ramp-Down)	[%/s]	~0, 25	~5	[6, 9, 17]
	CAPEX	[€/kW]	850	1050	[18]
	Stack durability	[h]	80000	60000	[19]
	Stack degradation	[%/1000h]	0.12	0.15	[19]
	Stack replacement costs	[%/CAPEX]	15	15	[19]
	Other OPEX	[%/CAPEX]	2	2	[19]
	Average electricity costs +Grid costs	[€/MWh]	105.9	105.9	[20]
	Electricity tax	[€/MWh]	11	11	[20]
	Life span	[a]	20	15	[19]
	Operating time	[h/a]	8000	8000	[-]
	Total scenario duration	[h]	13.2347	13.2347	[-]
Con. 0-10	EL/Block	[-]	8	2	[6, 21]
	EL-Power	[MW]	5	2	[6]
Con. 11-16	EL/Block	[-]	2	2	[6, 21]
	EL-Power	[MW]	5	0.7	[6]

Definition of the configurations and the control strategies

Once the use case, optimization problem, and models have been defined, the decision evaluation points of the defined decision variables will be specified. Table 2 provides an overview of the composition of the individual configurations (C). The number of electrolysis system blocks and their block output are specified. The last column provides an overview of the power ratio between the electrolysis technologies AEL and PEMEL.

The basic control strategies from [9] are used as the control strategy (R). These are further refined by

additions. A base load share has been defined, as the minimum operational load of one electrolysis system, for each strategy in order to always generate a minimum amount of hydrogen. The opposite of this is a strategy without base load (woBL), where the entire electrolysis system does not produce any hydrogen as soon as the power input falls below the minimum operational load. In the segment principle, a further subdivision was made based on the control strategy. AEL1 stands for AEL having priority. First, all AEL blocks are started up to 80%, then to 100%. This is followed by the PEMEL blocks, which are started up to 80% and then to 100%. N stands for a slightly different variant of this. Here, all AEL blocks are first started up to 80%, then all PEMEL blocks to 80%. Then the AEL blocks are first run at 100%, followed by the PEMEL blocks at 100%. These strategies serve to make the partial load efficiencies of the stack units usable for longer [22].

Table 2: Comparison of the different configurations. No. AEL/No. PEMEL represents the number of system blocks, the specified power refers to the system power of the individual block, and the power ratio indicates the ratio of AEL:PEMEL.

Configuration (C)	No. AEL	AEL [MW]	No. PEMEL	PEMEL [MW]	Power ratio
Con. 0	0	40	109	4	0.000
Con. 1	1	40	98	4	0.102
Con. 2	2	40	87	4	0.230
Con. 3	3	40	76	4	0.395
Con. 4	4	40	66	4	0.606
Con. 5	5	40	55	4	0.909
Con. 6	6	40	44	4	1.364
Con. 7	7	40	33	4	2.121
Con. 8	8	40	22	4	3.636
Con. 9	9	40	11	4	8.182
Con. 10	10	40	0	4	∞
Con. 11	34	10	47	1.4	5.167
Con. 12	35	10	40	1.4	6.250
Con. 13	36	10	32	1.4	8.036
Con. 14	37	10	24	1.4	11.012
Con. 15	38	10	16	1.4	16.964
Con. 16	39	10	9	1.4	30.952

The NewS strategy was introduced specifically for the start-stop strategy. This aims to distribute the load by switching from PEMEL blocks to AEL blocks when all PEMEL blocks are at 100% capacity and AEL blocks are not yet fully utilized. This strategy is intended to help make better use of the more efficient AEL blocks.

Evaluation of simulation results

First, the objectives of operational efficiency (including increased efficiency during partial load operation and energy loss, due to minimum load constraints), LGC (taking into account the energy loss due to the control strategy and the minimum load constraints) and LCOH are calculated for each configuration (C) in each control strategy (R), using the simulation setup, described in [12], and the standardized input profile, described in [7]. The objective values are normalized using their global maxima across the control strategy (R) and configuration (C) area, to determine a weighted score comparable across all configurations (C) and control strategies (R).

RESULTS

This section summarises the investigated configuration space and the resulting weighted optima across all combinations of configuration C and control strategy R. Across all feasible configuration strategy pairs considered in the grid search, the global minimum of the weighted score is achieved by Con. 16, which operates under the segment strategy and achieves a weighted score of 0.801. In general, the optimal combinations of (C) and (R) are located in the right area of the heat map, which means a ratio of AEL to PEMEL of 11.012-30.952. This shows that a minimally added amount of PEMEL technology in a large-scale AEL system offers advantages over a pure AEL system, as the configurations become worse with increasing AEL share. Furthermore, the results show that exploiting the partial load efficiencies of the individual stack units, such as in the 'segment' or 'slow-start' strategy, offers an advantage over pure 'start-stop' operation.

Finally, finer configurations across all evaluated (C) and (R) have a more consistent weighted score, suggesting that the sensitivity of system performance to the operating concept decreases with increasing modular resolution. This still needs to be confirmed in further research.

DISCUSSION

The optimisation results highlight the central coupling between design (configuration/technology mix) and operation (control strategy) in modular electrolysis systems. Even within a relatively small discrete design space—defined here by 17 configurations and 10 control strategies—the ranking of configurations is not invariant to the chosen operating concept. This observation is consistent with process-systems intuition: control strategies influence transient loading, start–stop events, and part-load operation, which in turn affect flexibility metrics, efficiency, and ultimately the levelized hydrogen cost.

A key outcome is that PEMEL-only concepts do not emerge as optimal under the selected objective and constraints. This indicates that, for the examined use case and weighting, economic and efficiency advantages associated with AEL can dominate the aggregated

	Configuration (C)	Con. 0	Con. 1	Con. 2	Con. 3	Con. 4	Con. 5	Con. 6	Con. 7	Con. 8	Con. 11	Con. 12	Con. 13	Con. 9	Con. 14	Con. 15	Con. 16	Con. 10
	Power Ratio	0.000	0.102	0.230	0.395	0.606	0.909	1.364	2.121	3.636	5.167	6.250	8.036	8.182	11.012	16.964	30.952	100.000
Control Strategy (R)	Slow-Start_woBL					0.907	0.895	0.882	0.870	0.857	0.823	0.820	0.816	0.844	0.812	0.809	0.806	0.839
	Slow-Start_BL				0.902	0.895	0.891	0.880	0.869	0.856	0.823	0.819	0.816	0.844	0.812	0.809	0.806	0.839
	Segment_woBL_N				0.894	0.885	0.876	0.867	0.858	0.849	0.816	0.813	0.810	0.840	0.806	0.804	0.801	0.842
	Segment_BL_N			0.896	0.891	0.883	0.874	0.866	0.858	0.849	0.815	0.812	0.809	0.841	0.806	0.803	0.801	0.834
	Segment_woBL_AEL1			0.901	0.891	0.881	0.872	0.862	0.853	0.845	0.814	0.811	0.809	0.838	0.805	0.803	0.801	0.842
	Segment_BL_AEL1			0.901	0.891	0.881	0.872	0.862	0.853	0.845	0.814	0.811	0.809	0.838	0.805	0.803	0.801	0.842
	Start-Stop_woBL										0.843	0.840	0.837		0.833			
	Start-Stop_BL										0.843	0.840	0.837		0.833			
	Start-Stopp_woBL_NewS				0.922	0.907	0.880	0.878	0.868	0.862	0.825	0.825	0.822	0.855	0.819	0.818	0.816	
	Start-Stopp_BL_NewS				0.922	0.907	0.880	0.878	0.868	0.862	0.825	0.825	0.822	0.855	0.819	0.818	0.816	

Figure 2. Heat Map: The color spectrum reflects the distribution from the global minimum (dark green) through a mid-range (yellow) to the global maximum (dark red). The empty cells represent all combinations of (C) and (R) that failed to meet a constraint (LGC or LCOH) and are therefore not displayed, even if the weighted score would be lower.

objective. Importantly, this is not a universal statement about technology superiority; rather, it reflects the normalization, the constraint set (especially the LCOH).

The reported benefit of finer granularity—more uniform weighted score across strategies—has a clear engineering interpretation. Smaller module sizes increase combinatorial resolution in the feasible operating region, enabling closer matching of the demand profile while respecting gradient constraints. This can reduce the penalty associated with a suboptimal control choice. However, finer granularity may increase balance-of-plant complexity, footprint, and maintenance burden. Therefore, granularity should be assessed jointly with efficiency, flexibility and LCOH considerations.

Several limitations should be considered when interpreting the results. First, the analysis excludes energy storage and focuses on direct load-following. Storage can decouple electrolysis operation from renewable variability and may change the trade-off between flexibility and cost. Second, cost and efficiency parameters are taken from literature and technology assessments; uncertainty ranges and learning-curve effects are not propagated through the optimisation. Third, the electrolysis models used were validated using literature data rather than actual operating data. Since the model from [17] was validated using real operating data, the significance of the model used is considered to be high. Despite these caveats, the discrete grid-search framework is valuable for decision making because it is transparent, reproducible, and directly compatible with modular design spaces. It also provides a consistent basis for comparing large-scale and laboratory-scale electrolysis systems, provided that decision criteria and constraints are defined in a scale-aware manner.

CONCLUSION

A discrete multi-criteria optimisation framework was applied to identify favourable combinations of technology composition (AEL/PEMEL ratio) and control strategy for a modular heterogeneous electrolysis system in a green-steel use case. Using survey-derived weights and constraints on flexibility (LGC) and hydrogen cost (LCOH), the results show that (i) the optimal configuration depends on the operating strategy, (ii) PEMEL-only systems are not optimal under the applied objective, (iii) homogeneous AEL systems can be more optimal than homogeneous PEMEL systems when used in conjunction with suitable control concepts., and (iv) finer-granularity configurations tend to reduce sensitivity to the chosen control strategy.

OUTLOOK

Future work should extend the configuration space to include additional technologies (e.g., high-temperature electrolysis) and heterogeneous layouts with three or more technology types. Integrating hydrogen and electrical storage, as well as grid services, would enable assessing system flexibility under broader operational paradigms beyond direct load-following. Finally, experimental validation on small-scale heterogeneous systems can support model calibration and help translate the optimisation framework to practical control implementation in modular electrolysis plants.

ACKNOWLEDGEMENTS

We would like to thank the German Federal Ministry of Research, Technology and Space and the project management organization Jülich for their financial support within the framework of the eModule research project (FKZ 03HY116A) of the H2Giga lead platform. Furthermore, the P2O-Lab Learning Factory of the TUD Dresden University of Technology is thanked for its support.

REFERENCES

1. Mock M, Viedt I, Lange H, Urbas L. Heterogenous electrolysis plants as enabler of efficient and flexible Power-to-X value chains. Comput. Aided Chem. Eng., vol. 53, Elsevier; 2024, p. 1885–90. https://doi.org/10.1016/B978-0-443-28824-1.50315-X
2. Lange H, Gro?e M, Viedt I, Urbas L. Modular and heterogeneous electrolysis systems: a system flexibility comparison. Systems and Control Transactions 4:759-764 (2025). https://doi.org/10.69997/sct.177861
3. Lange H, Meltzner FM, Große M, Viedt I, Klose A, Urbas L. Design of a modular and heterogeneous electrolysis system to increase flexibility in the use case of direct iron reduction. Energy Reports 15:108801 (2026). https://doi.org/10.1016/j.egyr.2025.11.060
4. Chernikova OP, Zlatitskaya YA. Resource efficiency of metallurgical production. Steel Transl. 52:561-567 (2022). https://doi.org/10.3103/s0967091222060031
5. Chatenet M, Pollet BG, Dekel DR, Dionigi F, Deseure J, Millet P, Braatz RD, Bazant MZ, Eikerling M, Staffell I, Balcombe P, Shao-Horn Y, Schäfer H. Water electrolysis: from textbook knowledge to the latest scientific strategies and industrial developments. Chem. Soc. Rev. 51:4583-4762 (2022). https://doi.org/10.1039/d0cs01079k
6. Lange H, Klose A, Beisswenger L, Erdmann D, Urbas L. Modularization approach for large-scale electrolysis systems: a review. Sustainable Energy Fuels 8:1208-1224 (2024). https://doi.org/10.1039/d3se01588b
7. Lange H, Meltzner FM, Große M, Viedt I, Klose A, Urbas L. Design of a modular and heterogeneous electrolysis system to increase flexibility in the use case of direct iron reduction. Energy Reports 15:108801 (2026). https://doi.org/10.1016/j.egyr.2025.11.060
8. Sá Rêgo PM. Mcintyre, L. (2018). post-truth. cambridge, MA: MIT press. MJ 19:333-335 (2019). https://doi.org/10.14195/2183-5462_34_25
9. Lange H, Klose A, Lippmann W, Urbas L. Technical evaluation of the flexibility of water electrolysis systems to increase energy flexibility: a review. International Journal of Hydrogen Energy 48:15771-15783 (2023). https://doi.org/10.1016/j.ijhydene.2023.01.044
10. Lange H, Schirmer C, Mock M, Viedt I, Klose A, Urbas L. Operation strategies for modular large-scale alkaline electrolysis systems. 2024 IEEE 22nd International Conference on Industrial Informatics (INDIN) :1-7 (2024). https://doi.org/10.1109/indin58382.2024.10774361
11. Lange H, Gro?e M, Viedt I, Urbas L. Modular and heterogeneous electrolysis systems: a system flexibility comparison. Systems and Control Transactions 4:759-764 (2025). https://doi.org/10.69997/sct.177861
12. Gro?e M, Viedt I, Lange H, Urbas L. Dynamic operability analysis of modular heterogeneous electrolyzer plants using system co-simulation. Systems and Control Transactions 4:53-59 (2025). https://doi.org/10.69997/sct.102556
13. Mohr S, Schmidt-Achert T, Ziemsky V, Pichlmaier S. Von der Theorie zur Praxis: Warum grüner Wasserstoff teurer ist als gedacht. Forschungsstelle für Energiewirtschaft (FfE); 2025.
14. Hoffmann B, Ziemsky V, Mohr S. Was kostet das „Grün“ im Wasserstoff? Forschungsstelle für Energiewirtschaft (FfE); 2025.
15. Liu J, Bao X, Zhou Y, Xu Z, Guo W, Wu J, Guan X. Optimizing the design of combined alkaline and proton exchange membrane electrolyzers to fully utilize fluctuating renewable energy. 2024 IEEE 20th International Conference on Automation Science and Engineering (CASE) :335-340 (2024). https://doi.org/10.1109/case59546.2024.10711405
16. Große M. H2Giga-eModule-Simulationsmodelle, https://github.com/p2o-lab/H2Giga-eModule-standard-integration-profile 2025.
17. Sakas G, Ibáñez-Rioja A, Ruuskanen V, Kosonen A, Ahola J, Bergmann O. Dynamic energy and mass balance model for an industrial alkaline water electrolyzer plant process. International Journal of Hydrogen Energy 47:4328-4345 (2022). https://doi.org/10.1016/j.ijhydene.2021.11.126
18. sakthimurugan V, Lakshmikanth G, Balaji N, Roopashree R, Kumar D, Devarajan Y. Green hydrogen revolution: advancing electrolysis, market integration, and sustainable energy transitions towards a net-zero future. Results in Engineering 26:104849 (2025). https://doi.org/10.1016/j.rineng.2025.104849
19. Rezaei M, Akimov A, Gray EMA. Levelised cost of dynamic green hydrogen production: a case study for australia's hydrogen hubs. Applied Energy 370:123645 (2024). https://doi.org/10.1016/j.apenergy.2024.123645
20. SMARD B für E Gas, Telekommunikation, Post und Eisenbahnen. Energiemarkt aktuell - Entwicklung der Industriestrompreise. 2025.
21. Lange H, Schirmer C, Mock M, Viedt I, Klose A, Urbas L. Operation strategies for modular large-scale alkaline electrolysis systems. 2024 IEEE 22nd International Conference on Industrial Informatics (INDIN) :1-7 (2024). https://doi.org/10.1109/indin58382.2024.10774361
22. Buttler A, Spliethoff H. Current status of water electrolysis for energy storage, grid balancing and

sector coupling via power-to-gas and power-to-liquids: a review. Renewable and Sustainable Energy Reviews 82:2440-2454 (2018). https://doi.org/10.1016/j.rser.2017.09.003

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimizing Flexible Operation of Grid-Connected Electrolyzers: Storage Capacity as the Key to Economic Viability

Julian Pamperin[a], Hannes Lange[a], Michael Große[a], and Leon Urbas[a]

[a] Dresden University of Technology, Process-to-Order Group, Dresden, Saxony, Germany
* Corresponding Author: julian.pamperin@mailbox.tu-dresden.de.

ABSTRACT

Grid-connected electrolyzers with intermediate hydrogen storage offer significant potential for reducing electricity costs through flexible operation under dynamic pricing. A threshold-based scheduling optimization approach is developed that derives interpretable on/off production rules from electricity price signals. The method identifies local price thresholds separating high-price from low-price periods, yielding binary production schedules. Adaptive horizon partitioning—subdividing the scheduling horizon when constant thresholds become infeasible—is combined with a receding horizon strategy that implements only a portion of each optimized schedule before re-optimization. This procedure enables systematic investigation of how characteristics of Integrated Electrolyzer–Storage Systems (IESS) influence cost reduction potential while maintaining computational tractability for both offline analysis and online implementation. A case study applying the approach to historical German electricity prices across 2, 688 scenarios reveals that storage capacity is the primary determinant of cost-saving potential, while start-up and shut-down times within the tested range (up to 120 minutes) show negligible impact. Systems with storage capacities allowing for hydrogen production stops of up to 1.4 days achieved electricity cost reductions of up to 80% compared to steady-state production. Production cycles in optimal schedules predominantly follow 24-hour rhythms, with approximately 30 cycles per month observed for high-performing configurations. These findings indicate that conventional flexibility metrics—focused on ramp rates and start-up times—are inadequate for assessing the cost-saving potential of flexibly operated, grid-connected electrolyzers under dynamic pricing. Instead, equipment wear from the frequent production cycling required by optimal schedules emerges as a key consideration for economic operation.

Keywords: Scheduling, Hydrogen, Process Design, Rolling Horizon Optimization, Modelling and Simulations

BACKGROUND

The large-scale deployment of green hydrogen requires flexible electrolysis systems that can adapt to fluctuating renewable energy supply and dynamic electricity markets. Germany targets an electrolyzer capacity of 10 GW by 2030 for applications in green steel production and the chemical industry [1], yet specific flexibility requirements remain insufficiently understood despite extensive technical characterizations [2-4]. Current research explores control strategies for adaptive hydrogen production [4, 5] but rarely quantifies how start-up and shut-down dynamics impact techno-economic performance across different operational contexts [6]. This knowledge gap hinders optimal system design—especially for grid-connected electrolyzers with intermediate hydrogen storage, limiting the assessment of cost reduction potentials for electricity markets [7].

Existing scheduling approaches predominantly focus on behind-the-meter configurations, where electrolyzers adjust consumption based on local renewable availability rather than market signals [8]. Recent work demonstrates cost reductions through adaptive scheduling under dynamic prices [9–11], but these approaches

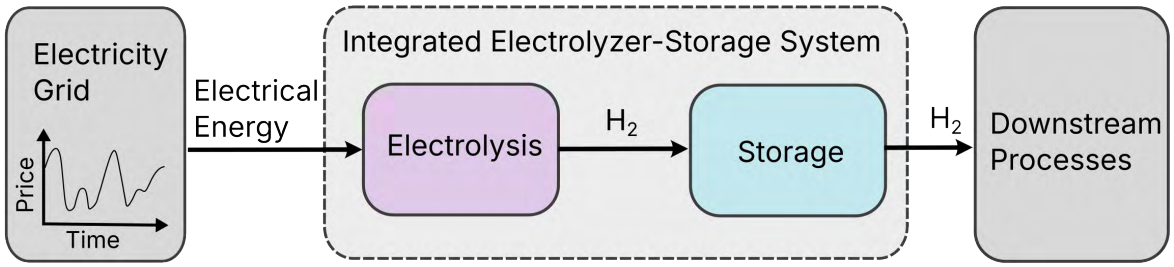

Figure 1. Integrated Electrolyzer–Storage System (IESS) and periphery

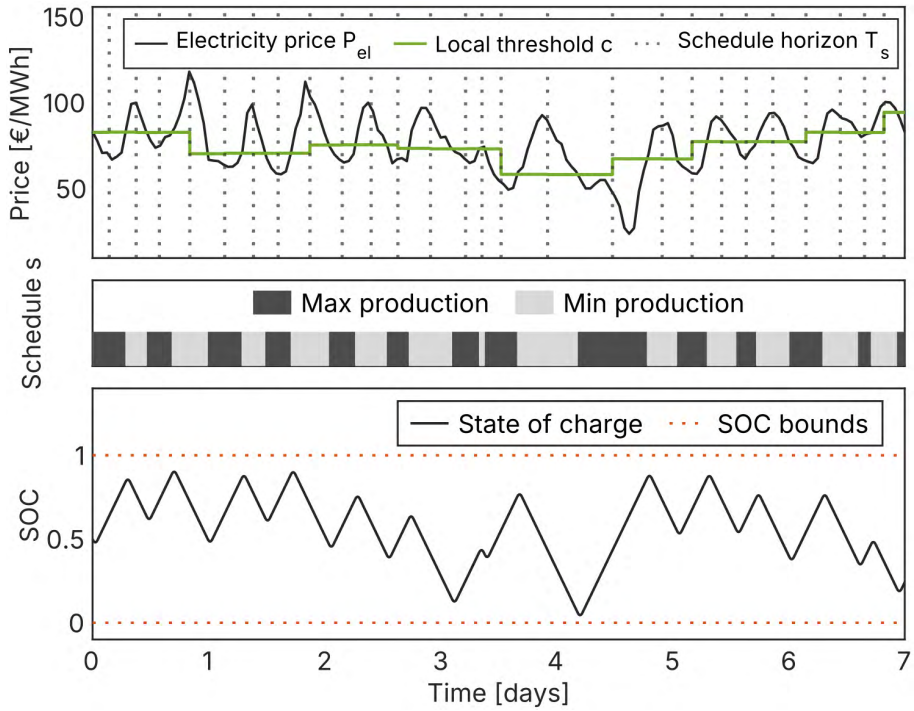

Figure 2. Production schedule generated by the optimization procedure. Top: electricity price P_{el} with local threshold c and schedule horizon boundaries. Middle: resulting binary production schedule with maximum and minimum production phases. Bottom: hydrogen storage state of charge (SOC) evolution within operational bounds.

yield schedules without revealing underlying decision logic—limiting online applicability. Methods for electrolyzer-storage sizing have been proposed [12], yet critical gaps remain: existing approaches lack systematic feasibility analysis, neglect electrolyzer dynamics, and operate over fixed horizons without mechanisms to prevent dead-end states. Consequently, the relationship between system characteristics and achievable cost savings remains poorly understood.

SCHEDULING OPTIMIZATION

Fluctuating electricity prices create a complex decision space balancing immediate costs against storage constraints and continuous downstream demand. The interdependence between production capacity, storage size, and price characteristics creates a nonlinear optimization landscape.

This work investigates optimal scheduling of grid-connected Integrated Electrolyzer–Storage Systems (IESS) under dynamic electricity prices (Figure 1). The system comprises a grid-connected electrolyzer, intermediate hydrogen storage, and downstream processes with constant demand. Variable hydrogen demand is excluded to isolate the scheduling problem.

Two main questions are addressed: (1) how to optimally schedule hydrogen production for a given IESS configuration, (2) which IESS characteristics influence the potential to minimize electricity costs.

The focus is exclusively on electricity procurement costs from the grid; capital expenditure, operational expenditure, and wear from load cycling are not considered in the optimization criterion. Additionally, behind-the-meter consumption of renewable energy sources is

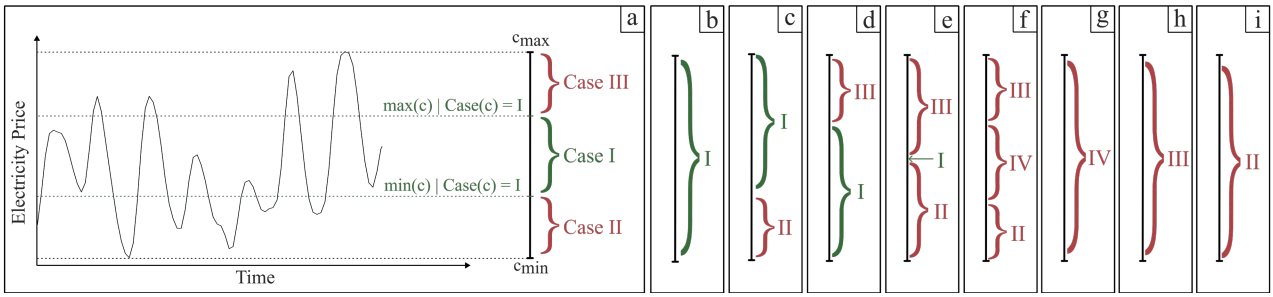

Figure 3. Visualisation of all circumstances relevant for Algorithm 1 early termination and reasoning for error preventing strategies.

neglected in this work. Such consumption does not eliminate the scheduling problem but rather modifies the price input: grid prices act as a competing alternative for local generation, creating a mixed effective electricity price with analogous dynamic patterns. Since only the price input changes—not the problem structure—the proposed method remains directly applicable.

The Optimal Production Schedule

The optimal production schedule minimizes electricity costs while maintaining continuous hydrogen supply for downstream processes. For cost functions based solely on electricity prices—deliberately excluding degradation, start-up costs, and cycling penalties to isolate dynamic pricing effects—clear patterns emerge. Cost is assumed linear with consumption; demand-dependent grid fees introduce nonlinearities but are negligible when consumption concentrates in low-price periods of generation excess, where flexible loads stabilize the grid. Intermediate load levels offer no advantage: any partial load hour could be improved by shifting to a cheaper hour at full load, leading to an on/off control concept. However, identifying the threshold separating low from high prices is non-trivial: storage capacity limits shutdown duration, production capacity determines deficit recovery speed, and start-up dynamics affect schedule precision. Thresholds must be chosen locally based on IESS state, forecast prices, and demand—creating a multi-level optimization problem. Scenarios where on/off schedules cannot achieve optimality are discussed in the Fatal Error Analysis.

Scheduling Optimization Concept

Figure 2 illustrates a production schedule generated by the proposed optimization procedure. The interdependencies between electricity price fluctuations, production schedule, and system state are central to the approach. The ***production schedule*** defines periods of maximum and minimum production. The ***production rate set point*** translates this schedule into control commands, accounting for electrolysis system inertia to align power consumption with intended production periods. The ***hydrogen production rate*** follows the load evolution but deviates due to efficiency dependencies and transient offsets. The ***State of Charge*** (SOC) of hydrogen storage evolves according to the balance between production and downstream demand, closing the optimization loop.

Iterative refinement of the production schedule requires simulation of the IESS to evaluate SOC trajectories for candidate schedules. SOC limits impose hard boundaries on feasible schedules, while electricity price evolution defines the cost landscape. Storage size bounds the horizon over which local decisions influence future optimality, motivating two application goals: extensive offline investigation of IESS configuration impacts, and online applicability with comparable schedule quality.

The optimization procedure follows a receding horizon framework conceptually related to Model Predictive Control [13], but operates at scheduling timescales (hours to days) rather than control timescales (seconds to minutes). An outer iteration scheme adaptively determines prediction horizon lengths based on feasibility conditions rather than using fixed horizons. Three challenges must be addressed: (1) identifying local electricity price thresholds defining high and low price periods, (2) selecting appropriate prediction horizons for which constant thresholds yield feasible schedules, and (3) preventing dead-end states arising from system inertia.

Estimation of Local Electricity Price Threshold

For a given time period, the electricity price threshold c that defines high and low price periods is bounded by the electricity price P_{el} extrema, i.e., $c \in [c_{\min}, c_{\max}]$ where $c_{\min} = \min(P_{\mathrm{el}})$ and $c_{\max} = \max(P_{\mathrm{el}})$. The production schedule $s(t)$ is then derived from the threshold according to:

$$s(t) = \begin{cases} 1 & \text{if } P_{\mathrm{el}}(t) \leq c \\ 0 & \text{if } P_{\mathrm{el}}(t) > c \end{cases} \tag{1}$$

where $s = 1$ corresponds to maximum production and $s = 0$ to minimum production.

Within the range of $[c_{\min}, c_{\max}]$, four cases are distinguished based on the resulting SOC trajectory (see Table

1): Case I corresponds to fully feasible operation within storage limits; Case II occurs when storage runs empty at least once but never exceeds capacity; Case III occurs when storage exceeds capacity but never runs empty; and Case IV represents infeasible operation where storage both runs empty and overflows. Multiple circumstances determine which cases can occur simultaneously for a given IESS configuration. These depend on its state, price trajectory and H_2 demand; the relevant circumstances are shown in Figure 3. Circumstances a–e admit feasible solutions where at least one c yields Case I; circumstances f–i admit no feasible threshold. Based on this classification, the search algorithm identifies the feasible interval $[c^I_{\min}, c^I_{\max}]$.

Table 1: Classification of Storage Operation Cases Based on State of Charge (SOC).

Case	Description
I	If $\forall t \in [t_{start}, t_{end}]$ it holds that $0 < \mathrm{SOC}(t) < 1$, the H_2 storage operates entirely within its admissible limits and neither runs empty nor exceeds its capacity. Consequently, c is considered sufficient.
II	If $\exists t \in [t_{start}, t_{end}]$ such that $\mathrm{SOC}(t) < 0$, while $\mathrm{SOC}(t) \leq 1$ for all $t \in [t_{start}, t_{end}]$, the H_2 storage runs empty at least once but never exceeds its capacity.
III	If $\mathrm{SOC}(t) \geq 0$ for all $t \in [t_{start}, t_{end}]$ and $\exists t \in [t_{start}, t_{end}]$ such that $\mathrm{SOC}(t) > 1$, the H_2 storage capacity is exceeded at least once, while the storage never runs empty.
IV	If $\exists t_1, t_2 \in [t_{start}, t_{end}]$ such that $\mathrm{SOC}(t_1) < 0$ and $\mathrm{SOC}(t_2) > 1$, the H_2 storage both runs empty and exceeds its capacity during the considered interval. In this case, no value of c can satisfy the H_2 demand constraint while remaining within the storage relaxation limits.

The minimum feasible threshold $c^I_{\min}$ is identified via interval bisection (Algorithm 1). Each candidate is evaluated by simulating the IESS: Case I indicates feasibility, Case II indicates threshold too low, Case III too high, and Case IV infeasibility. The interval narrows until convergence $|c_n - c_{n-1}| < \delta$; $c^I_{\max}$ is obtained analogously. For nonlinear IESS models, non-monotonic relationships may cause bisection failure, requiring algorithm extension.

Identification of the feasible threshold interval $[c^I_{\min}, c^I_{\max}]$ enables generation of multiple valid production schedules with different SOC characteristics. A threshold near $c^I_{\min}$ produces schedules that minimize hydrogen storage utilization, while $c^I_{\max}$ yields schedules that tend to fill the storage. To allow control over SOC behavior within the feasible range, a scheduling goal parameter $g \in [0, 1]$ is introduced:

$$c_{set} = c^I_{\min} + \left(c^I_{\max} - c^I_{\min}\right) \bullet g \tag{2}$$

where c_{set} is the threshold used to generate the production schedule according to Equation 1. Setting $g = 0$ selects $c^I_{\min}$, minimizing mean SOC; setting $g = 1$ selects $c^I_{\max}$, maximizing mean SOC.

The influence of g depends strongly on initial conditions and the width of the feasible interval. For circumstance e (Figure 3), where the feasible interval is narrow, the scheduling goal has limited effect on SOC trajectory. This basic threshold estimation performs well when electricity prices follow uniform oscillations with period lengths compatible with storage capacity. More complex scenarios result in fatal errors discussed in the following.

Fatal Error Analysis

Electricity price evolutions exhibit global trends, local trends, and randomness, inherently limiting periods for which a feasible constant threshold exists. A fatal error occurs when storage underflows ($\mathrm{SOC} < 0$) or overflows ($\mathrm{SOC} > 1$). Five failure modes are identified:

- **Long Periods:** When shutdown phases exceed storage hold-up duration, underflow becomes inevitable regardless of threshold selection—the system must produce during high-price periods, rendering binary switching unsuitable. Consequently, long-term price variation exploitation to reduce electricity costs is limited.
- **Short Periods:** Threshold-based scheduling triggers switching at every price-threshold crossing. When resulting cycles occur at timescales below start-up and shut-down times, the system cannot follow the schedule and these price variations cannot be exploited.
- **Changing Characteristics:** Periods with uniform price fluctuations admit relatively constant thresholds; when global or local trends shift, previously valid thresholds become infeasible.
- **Unfavorable Phase:** Initial SOC misaligned with price phase can cause SOC oscillations violating capacity limits despite otherwise suitable price characteristics.
- **Dead-Ends:** These correspond to circumstances f–i (Figure 3), where SOC approaches limits while production state is too inert to prevent violations. Dead-ends are critical when feasible schedules yield end states near circumstance e, where subsequent operation potentially fails. These failure modes motivate the error preventing strategies presented in the following section.

Algorithm 1: Minimum electricity price threshold for demand-feasible scheduling

Input : $P_{\mathrm{el}}(t)$: electricity price function, δ: convergence tolerance
Output: $c^I_{\min}$: optimal minimal threshold, or $\emptyset$ if no solution exists

```
/* Initialization */
c_min ← min(P_el),   c_max ← max(P_el)
c_1 ← c_min,   c_2 ← c_max,   c_3 ← (c_min + c_max)/2
Evaluate Case(c_1), Case(c_2), Case(c_3)
if circumstances indicate early termination then
    return corresponding result
n ← 3
/* Main iteration loop */
while true do
    if Case(c_n) = I then
        if |c_n − c_{n−1}| < δ and Case(c_{n−1}) = II then
            return c* ← c_n
        else
            c^+ ← max{c_i : Case(c_i) = II, i ≤ n}
            c_{n+1} ← (c^+ + c_n)/2
    else if Case(c_n) = II then
        c^+ ← min{c_i : Case(c_i) ∈ {I, III}, i ≤ n}
        c_{n+1} ← (c^+ + c_n)/2
    else if Case(c_n) = III then
        c^+ ← max{c_i : Case(c_i) = II, i ≤ n}
        c_{n+1} ← (c^+ + c_n)/2
    else
        return ∅
    n ← n + 1
```

Error Preventing Strategies

Adaptive horizon partitioning and receding horizon strategies ensure robustness of the global optimization procedure (Algorithm 2). By adaptive partitioning, the scheduling horizon is subdivided into segments, each admitting a feasible constant threshold c^I. Starting from initial segment length T_0 the algorithm iteratively reduces segment size by factor $\alpha \in (0,1)$ until a feasible threshold exists:

$$T_{\mathrm{k+1}} = \alpha \bullet T_{\mathrm{k}} \tag{3}$$

This continues until the segment admits Case I or minimum segment length is reached. For non-dead-end circumstances (a–e in Figure 3), this guarantees a feasible schedule. Segment boundaries may be suboptimally placed, but near-optimal results in the case study indicate limited room for improvement.

Adaptive partitioning alone cannot prevent dead-end states at segment boundaries nor account for future price evolution. A locally optimized schedule may drive SOC toward extreme values, leaving insufficient flexibility for subsequent segments (e in Figure 3). To introduce predictiveness, a receding horizon framework distinguishes the *prediction horizon* T_p from the *schedule horizon* T_s [13]. At each iteration, the algorithm optimizes over T_p, applies only T_s, then shifts forward and re-optimizes—preventing myopic decisions that lead to dead-ends.

Algorithm 2: Optimized production schedule generation with adaptive horizon partitioning and receding horizon

Input : $P_{\mathrm{el}}(t)$: electricity price trajectory, T_0: initial segment length, α: decreasing factor, o_f: overlap factor, g: scheduling goal, $T_{\min}$: minimum segment length, $\mathbf{x}_0$: initial IESS state
Output: $s(t)$: production schedule over scheduling horizon

```
t_current ← 0
x ← x_0
s(t) ← ∅
while t_current < t_end do
    T_p ← T_0
    c^I_min ← ∅
    while c^I_min = ∅ and T_p ≥ T_min do
        c^I_min ← EstimateCmin(P_el, [t_current, t_current + T_p], x)
        if c^I_min = ∅ then
            T_p ← α · T_p
    if c^I_min = ∅ then
        return infeasible
    c^I_max ← EstimateCmax(P_el, [t_current, t_current + T_p], x)
    c_set ← c^I_min + (c^I_max − c^I_min) · g
    s_p(t) ← s(t; c_set) for t ∈ [t_current, t_current + T_p]
    T_s ← T_p (1 − o_f/(1+o_f))
    Append s_p(t) for t ∈ [t_current, t_current + T_s] to s(t)
    x ← SimulateIESS(x, s_p, [t_current, t_current + T_s])
    t_current ← t_current + T_s
return s(t)
```

The ratio between prediction and schedule horizons is defined through an overlap factor o_f, which specifies the fractional overlap between consecutive prediction horizons of equal size:

$$T_s = T_p\left(1 - \frac{o_f}{1+o_f}\right) \tag{4}$$

For $o_f = 0.5$, half of consecutive prediction horizons overlap. This parameterization provides intuitive tuning: higher overlap increases predictiveness at the cost of computational effort.

The combination of adaptive horizon partitioning and receding horizon implementation robustly generates feasible production schedules. These schedules are, due to their heuristic nature, often but not necessarily close to global optimality. Therefore, schedule quality must be assessed independently, as addressed in the following section.

OPTIMALITY SCORE

To evaluate how effectively a schedule exploits price variability, an optimality score $\phi \in [0,1]$ is introduced comparing achieved performance to best-case (optimal) and worst-case (pessimal) consumption allocations, given identical total energy demand, temporal consumption structure, and system limitations. The score is independent of absolute price levels, instead quantifying how well local fluctuations are utilized through the dispersion of the consumption-weighted price distribution relative to global mean price $\bar{P}$.

Let $\tilde{P}_{\mathrm{el}}(t)$ denote the electricity price signal after removing non-exploitable long-term trends imposed by storage limitations, while preserving the global mean $\bar{P}$. For a given production schedule with consumption rate $\dot{E}(t)$ and total energy consumption

$$E_{\mathrm{total}} = \int \dot{E}(t)\,\mathrm{d}t, \tag{5}$$

the induced consumption-weighted price distribution is defined through the measure

$$\mathrm{d}E = \dot{E}(t)\,\mathrm{d}t \quad \text{at price} \quad \tilde{P}_{\mathrm{el}}(t) \quad . \tag{6}$$

For the same total energy consumption and the same temporal consumption profile $\dot{E}(t)$, two limiting reference distributions are defined:

- The *optimal distribution* $\mathcal{T}_{\mathrm{opt}}$, obtained by reassigning the largest consumption magnitudes to the lowest available values of $\tilde{P}_{\mathrm{el}}(t)$, with consumption magnitudes decreasing monotonically as prices approach $\bar{P}$.
- The *pessimal distribution* $\mathcal{T}_{\mathrm{pes}}$, obtained analogously by reassigning the largest consumption magnitudes to the highest available values of $\tilde{P}_{\mathrm{el}}(t)$.

In both cases, the distribution of consumption magnitudes over time is preserved; only its association with electricity prices is reordered. As a result, ramping behavior due to start-up and shut-down dynamics is inherently respected, while unnecessary production cycling does not artificially inflate the achievable optimality. The quality of each distribution is quantified using its central moment relative to the global mean price:

$$\mu_{\mathrm{opt}} = \frac{1}{E_{\mathrm{total}}} \int_{\mathcal{T}_{\mathrm{opt}}} \left(\tilde{P}_{\mathrm{el}} - \bar{P}\right)^2 \mathrm{d}E, \tag{7}$$

$$\mu_{\mathrm{pes}} = \frac{1}{E_{\mathrm{total}}} \int_{\mathcal{T}_{\mathrm{pes}}} \left(\tilde{P}_{\mathrm{el}} - \bar{P}\right)^2 \mathrm{d}E, \tag{8}$$

$$\mu_{\mathrm{schedule}} = \frac{1}{E_{\mathrm{total}}} \int \left(\tilde{P}_{\mathrm{el}}(t) - \bar{P}\right)^2 \dot{E}(t)\,\mathrm{d}t. \tag{9}$$

The optimality score is defined by normalizing the achieved second moment between the pessimal and optimal bounds:

$$\phi = \frac{|\mu_{\mathrm{schedule}} - \mu_{\mathrm{pes}}|}{|\mu_{\mathrm{opt}} - \mu_{\mathrm{pes}}|}. \tag{10}$$

By construction, $\phi = 1$ corresponds to optimal reassignment of consumption to low-price periods, $\phi = 0$ to pessimal reassignment, and $\phi \approx 0.5$ to steady-state production. The score captures theoretical scheduling effectiveness while accounting for storage capacity and production inertia limitations, deliberately excluding start-up costs, efficiency variations, cycling degradation, and capital expenditure.

CASE STUDY

The proposed scheduling optimization was applied to a range of IESS configurations under historical electricity price data to evaluate performance across diverse scenarios. The study investigates three aspects: (1) performance of the scheduling optimization procedure, (2) IESS characteristics influencing electricity cost-saving potentials, and (3) operational strain indicated by production cycling. The section first describes the simulation model and experimental setup, then presents metrics for evaluation.

IESS Model

For computational tractability, a simplified dynamic IESS model was employed. The production schedule is directly translated to production rate setpoints, resulting in anticipated phase lag between scheduled and actual production. Electricity consumption and hydrogen production are coupled through a linear relationship. Start-up and shut-down dynamics are represented by linear ramps following the production rate setpoint.

Table 2: IESS configuration parameters. Relative production and storage capacities are normalized to downstream hydrogen demand.

Parameter	Tested Values
Relative production capacity [–]	1.1, 1.5, 2.0, 2.6, 3.4, 4.4, 5.5, 6.0
Start-up/shut-down time [min]	0, 40, 80, 120
Relative storage capacity [days]	0.05, 0.25, 0.5, 0.7, 1.4, 2.8, 4.9

Experimental Setup

IESS configurations were varied in a full factorial design according to Table 2. The ***relative production capacity*** is defined as the ratio of maximum hydrogen production rate to the constant hydrogen demand of downstream processes; a value of 2.0 indicates the electrolyzer can produce twice the demanded flow rate. The ***relative storage capacity*** is defined as the ratio of storage volume to one day of downstream hydrogen demand; a value of 1.0 indicates storage sufficient to supply downstream processes for 24 hours without production. Each

configuration was tested using historical electricity price data for each month of 2024, sourced from the German Federal Network Agency [14], yielding 2, 688 optimization runs. Parameters of the scheduling optimization were set to an overlap factor of o_f = 1, scheduling goal g = 0.5, and segment size reduction factor α = 0.7.

Metrics

Three metrics quantify scheduling performance and electricity cost-saving potential of IESS configurations. The previously defined ***optimality score*** ϕ measures schedule quality relative to theoretical bounds. The ***relative electricity cost*** C_{rel} captures the ratio of electricity costs under the optimized schedule to costs under steady-state production:

$$C_{\text{rel}} = \frac{\sum_t P_{\text{el}}(t) \cdot \dot{E}(t)}{\bar{P}_{\text{el}} \cdot E_{\text{total}}} \tag{11}$$

where $\dot{E}(t)$ is the instantaneous electricity consumption rate. This metric should be interpreted as an indicator of cost-saving potential rather than a prediction of actual savings, as realized electricity costs are subject to considerable uncertainty including capacity charges, grid fees for flexible loads, and evolving market dynamics. The ***number of production cycles*** counts complete start-up/shut-down sequences per month, indicating operational strain relevant to equipment wear.

Results and Discussion

The optimality score ϕ exhibits strong dependence on relative storage capacity (Figure 4). Configurations with relative storage capacities below 0.5 days consistently yielded $\phi < 0.8$, while capacities exceeding 1.4 days achieved $\phi > 0.9$ across most scenarios. Low storage capacity leads to the long periods problem. This indicates general unsuitability of low-storage configurations for exploiting electricity market price fluctuations, independent of the scheduling method employed.

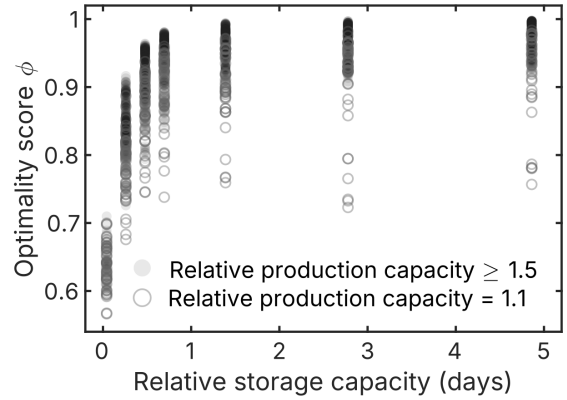

Figure 4. Optimality score ϕ as a function of relative storage capacity. The dependency to relative production capacity indicated by filled and unfilled markers.

Electricity cost-saving potential mirrors optimality score behavior (see Figure 5). Systems with storage capacities below 0.5 days achieved only marginal cost reductions, while capacities above 1.4 days overcame the storage bottleneck. Production capacities near unity (1.1) severely constrain flexibility due to insufficient surplus for shutdown periods. Cost-saving potential increases asymptotically with production capacity: values of 2.0 resulted in 70% relative electricity costs compared to steady-state production, while 4.0 achieved up to 20% relative costs in favorable months.

Start-up and shut-down times up to 120 minutes showed negligible impact on cost savings. Whether marginal observed dependencies result from the phase mismatch inherent in the simplified IESS control or reflect genuine limitations of process inertia on cost-saving potential cannot be conclusively determined from the present study.

For configurations with storage capacities above 1.4 days, a mean of 35 production cycles per month was observed (Figure 6), suggesting price fluctuations relevant

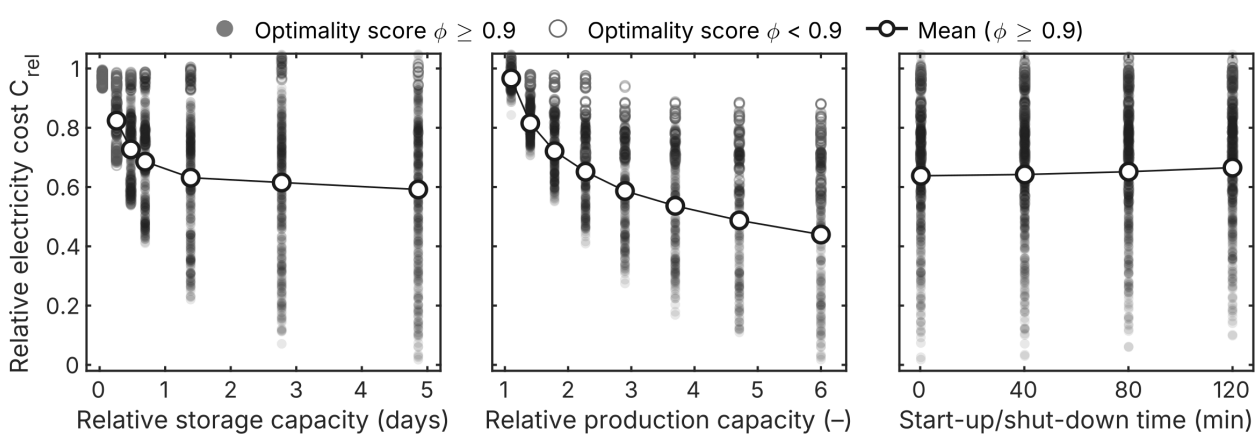

Figure 5. Relative electricity cost C_{rel} as a function of relative storage capacity, relative production capacity, and start-up/shut-down time.

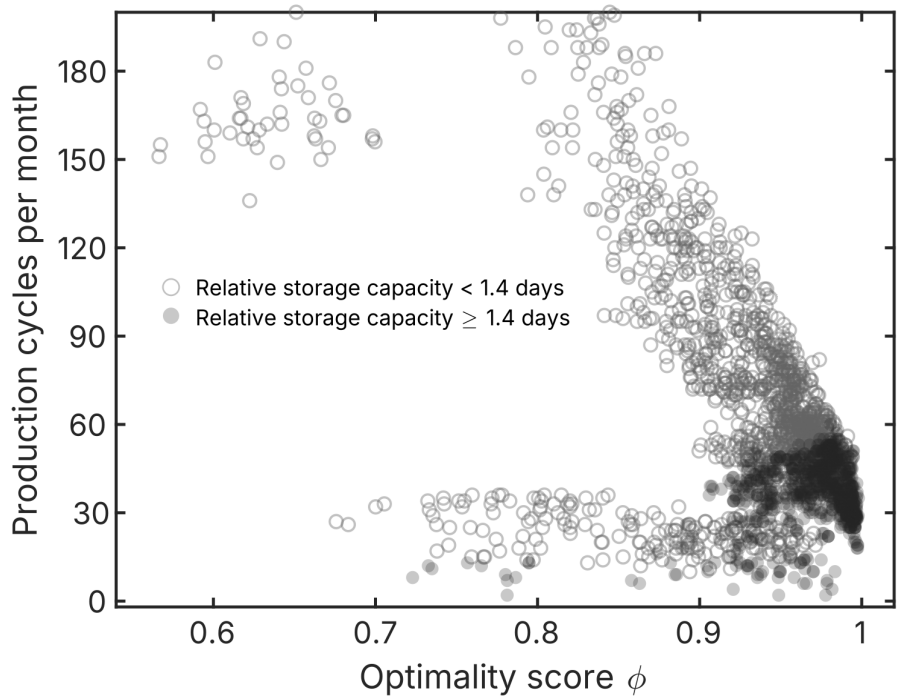

Figure 6. Production cycles per month as a function of optimality score ϕ.

for cost optimization occur predominantly at daily timescales—consistent with diurnal patterns in markets with significant renewable penetration. Production cycle count remained largely independent of start-up and shut-down times.

The scheduling optimization procedure demonstrated robust performance across the parameter space. Except for configurations combining small storage capacities (< 0.25 days) with high production capacities (> 4.0), where scheduled cycle durations approached or fell below start-up and shut-down times, the algorithm reliably generated feasible production schedules. For all other configurations, the adaptive horizon partitioning and receding horizon strategies successfully prevented dead-end states and maintained SOC within admissible bounds. Computational cost was dominated by IESS simulation, with cumulative simulated time consistently between 100 and 160 times the generated schedule length for well-performing configurations ($\phi > 0.9$). While sequential dependencies between receding horizon iterations limit full parallelization, independent IESS simulations within each iteration are parallelizable, offering speedup potential by orders of magnitude.

CONCLUSION

This work presents a threshold-based scheduling optimization approach for grid-connected electrolyzer-storage systems that derives interpretable production rules from electricity price characteristics, enabling both offline analysis and online application. The theoretical framework establishes necessary conditions for feasible schedule existence, and the optimality score ϕ provides a system-independent quality metric for integration into higher-level optimization.

Results challenge conventional flexibility metrics for electrolysis systems: start-up and shut-down times up to 120 minutes showed negligible impact on cost-saving potentials, while storage capacity emerged as the primary determinant. For grid-connected electrolyzers, the number of production cycles—rather than process dynamics—emerges as the critical constraint requiring assessment for economic operation.

OUTLOOK

Future work should quantify how cycling frequency limits affect cost-saving potential by imposing minimum on-period durations to constrain cycle counts, enabling economic assessments. Extending the optimization criterion to include capital expenditure and cycling-induced degradation would support comprehensive techno-economic evaluation. Application to nonlinear IESS models, variable demand, and diverse market conditions with increasing price volatility [15] would broaden applicability.

ACKNOWLEDGEMENTS

This work was supported by the German Federal Ministry of Research, Technology and Space and the

project management organization Jülich within the framework of the eModule research project (FKZ 03HY116A) of the H2Giga lead platform. We also gratefully acknowledge the support of the P2O-Lab Learning Factory at TUD Dresden University of Technology.

REFERENCES

1. BMWK (2023). National Hydrogen Strategy Update. Federal Ministry for Economic Affairs and Climate Action, Germany.
2. Brauns J, Turek T. Alkaline water electrolysis powered by renewable energy: a review. Processes 8:248 (2020). https://doi.org/10.3390/pr8020248
3. Buttler A, Spliethoff H. Current status of water electrolysis for energy storage, grid balancing and sector coupling via power-to-gas and power-to-liquids: a review. Renewable and Sustainable Energy Reviews 82:2440-2454 (2018). https://doi.org/10.1016/j.rser.2017.09.003
4. Lange H, Klose A, Lippmann W, Urbas L. Technical evaluation of the flexibility of water electrolysis systems to increase energy flexibility: a review. International Journal of Hydrogen Energy 48:15771-15783 (2023). https://doi.org/10.1016/j.ijhydene.2023.01.044
5. Götz M, Lefebvre J, Mörs F, McDaniel Koch A, Graf F, Bajohr S, Reimert R, Kolb T. Renewable power-to-gas: a technological and economic review. Renewable Energy 85:1371-1390 (2016). https://doi.org/10.1016/j.renene.2015.07.066
6. Lorenz J, Klose A, Lange H, Kock T, Urbas L. Process control principles for scalable electrolysis systems. 2023 IEEE 28th International Conference on Emerging Technologies and Factory Automation (ETFA) :1-4 (2023). https://doi.org/10.1109/etfa54631.2023.10275330
7. Yuan T, Tan J, Teng Y. A review of dynamic modeling and control of grid-connected hydrogen production units using water electrolysis. Renewable and Sustainable Energy Reviews 226:116296 (2026). https://doi.org/10.1016/j.rser.2025.116296
8. Y. Zhang et al., "Water electrolyzer operation scheduling for green hydrogen production: A review, " Renewable and Sustainable Energy Reviews, vol. 205, Art. no. 114857, 2024, doi: 10.1016/j.rser.2024.114857.
9. Madsen FW, Jørgensen BN, Ma ZG. Integrating electricity market granularity and sector coupling: adaptive power-to-x scheduling optimization under dynamic electricity markets. Energies 18:6182 (2025). https://doi.org/10.3390/en18236182
10. A. Lesniak, A. G. Johnsen, N. Rhodes, and L. Roald, "Advanced Scheduling of Electrolyzer Modules for Grid Flexibility, " arXiv:2412.19345, 2024.
11. J. W. Law, B. K. Mignone, and D. S. Mallapragada, "Design Space for PEM Electrolysis for Cost-Effective H2 Production Using Grid Electricity, " Industrial & Engineering Chemistry Research, vol. 63, no. 16, pp. 7258–7270, 2024, doi: 10.1021/acs.iecr.4c00123.
12. Saars L, Madsen M, Meyer J. Optimizing the operation of an electrolyzer with hydrogen storage using two different methods: a trade-off between simplicity and precision in minimizing hydrogen production costs using day-ahead market prices. Energies 17:5546 (2024). https://doi.org/10.3390/en17225546
13. J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, Model Predictive Control: Theory, Computation, and Design, 2nd ed. Madison, WI: Nob Hill Publishing, 2017.
14. SMARD – Strom- und Gasmarktdaten, "Marktdaten visualisieren, " Bundesnetzagentur, Germany. [Online]. Available: https://www.smard.de/home/marktdaten. Accessed: Mar. 12, 2025.
15. Wang G, Sbai E, Wen , Selena Sheng M. The impact of renewable energy on extreme volatility in wholesale electricity prices: evidence from organisation for economic co-operation and development countries. Journal of Cleaner Production 484:144343 (2024). https://doi.org/10.1016/j.jclepro.2024.144343

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Comparative Techno-Economic Assessment of Hybrid-Green Ammonia Layouts for Available-to-Date Decarbonization of the Fertilizer Industry

Andrea Isella and Davide Manca*

PSE-Lab, Process Systems Engineering Laboratory, Dipartimento di Chimica, Materiali e Ingegneria Chimica "Giulio Natta", Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy
* Corresponding Author: davide.manca@polimi.it.

ABSTRACT

Industrial ammonia synthesis remains a high-impact decarbonization target due to the combined effect of large production volumes and reliance on CO_2-intensive fossil-based hydrogen. Replacing it partially with renewable-based (green) electrolytic hydrogen can minimize direct emissions; however, the intermittency of solar and wind complicates retrofit strategies for existing continuous plants. This paper addresses a techno-economic study centered on a retrofit-oriented hybrid concept: a conventional natural-gas ammonia plant is operated within a pre-defined "hybridization envelope", where part of the process hydrogen is supplied by a renewable Power-to-Hydrogen subsystem. A steady-state process model (gray/blue baselines) is coupled with an hourly energy dispatch and a sizing optimization of solar/wind park, electrolyzer, battery, and hydrogen storage. A case study based on 2024 California data for renewables and carbon pricing illustrates how hybridization can reduce carbon intensity while ensuring all process constraints are met, and how the optimal renewable mix trades off between overbuild and storage. Results highlight the role of solar-wind complementarity in lowering the levelized cost of hydrogen and, consequently, the levelized cost of ammonia, while leaving a clear pathway for combining hybrid-green operation with carbon capture (hybrid-blue-green) as a transitional strategy.

Keywords: Green ammonia, Electrolysis, Hydrogen storage, Energy storage, Retrofit, Power-to-Hydrogen

INTRODUCTION

Ammonia is central to the global fertilizer value chain and is increasingly considered as an energy carrier for hard-to-abate sectors. However, the scientific community consistently identifies ammonia synthesis as a priority petrochemical process for decarbonization, alongside steam cracking and methanol [1, 2]. Conventional ("gray") ammonia production relies on natural gas reforming to supply hydrogen and on high-pressure synthesis loops that operate most efficiently in continuous mode. As a result, deep decarbonization of existing ammonia plants faces a dual challenge: (i) abating CO_2 emissions from hydrogen production and process energy, and (ii) integrating low- or zero-carbon energy vectors without destabilizing plant operation.

Green ammonia pathways typically assume renewable power fully supplies hydrogen via water electrolysis. While conceptually straightforward, the intermittency of solar and wind imposes operational and economic penalties unless extensive flexibility (grid exchange, mass and energy storage, electricity curtailment) is provided [3, 4]. Conversely, blue ammonia combines natural-gas-based hydrogen with carbon capture and storage (CCS). Blue routes can achieve meaningful emission cuts but remain vulnerable to concerns about upstream methane leakage and the associated costs of capture/transport/storage [5, 6]. Hybrid-green pathways, where a fraction of the process hydrogen is supplied by renewable electricity and the rest of the plant remains fossil-fed, can provide a pragmatic bridge, especially for retrofits [7, 8]. The hybrid approach aims to decouple "how much" renewable energy is used from "how the synthesis loop operates", by enforcing an operating envelope compatible with safe and efficient plant behavior.

Therefore, this work details the hybrid-green retrofit

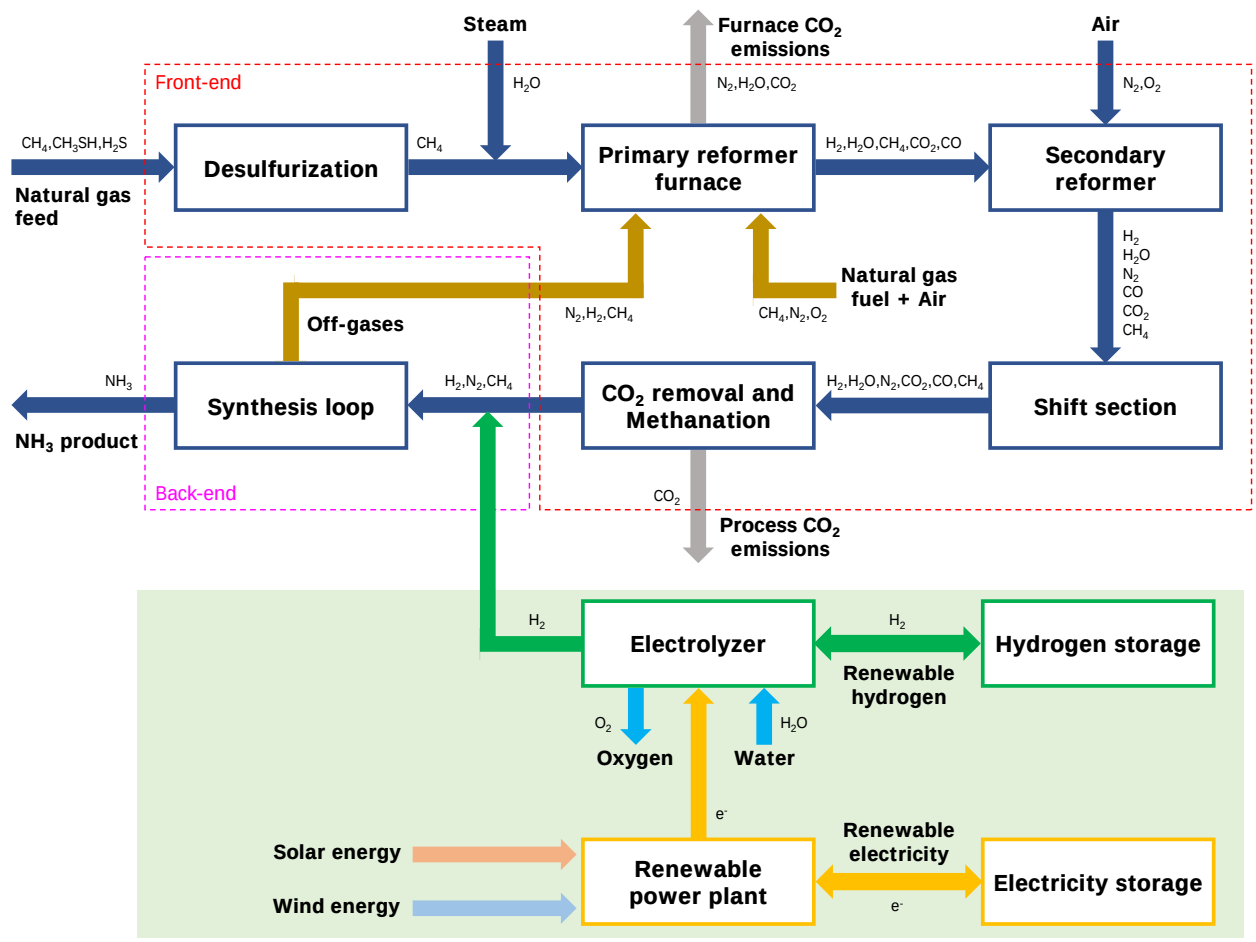

Figure 1. Block flow diagram of the hybrid-green retrofit for an ammonia plant. The green rectangle encloses the additional units introduced by the hybrid concept.

concept for existing ammonia plants and also provides a quantitative assessment of its techno-economic and environmental performance in the 2024 California context. Specifically, the analysis refers to published best-available-technology baselines for ammonia synthesis and recent cost data for renewables, electrolysis, and mass/energy storage. Furthermore, it expands an optimization-based sizing methodology for Power-to-Hydrogen (PtH) assets previously developed by the Authors [8].

HYBRID-GREEN RETROFIT CONCEPT AND OPERATING ENVELOPE

The hybrid concept retrofits a conventional ammonia plant by adding a PtH subsystem (renewable electricity, electrolyzer, battery, H_2 storage) that can displace a fraction of the fossil-based hydrogen produced by the reformer (see **Figure 1**). To keep downstream units and the synthesis loop close to design conditions, we investigate the hybridization factor as the fraction of the plant's hydrogen demand supplied by the PtH system on average.

Rather than allowing arbitrary turndown, the retrofit is constrained to an "operating envelope" expressed by (i) minimum acceptable "throughput" (*i.e.*, the mass flow passing through the primary reformer), and (ii) maximum acceptable changes of key state variables (*i.e.*, the operating temperature in the secondary reformer and the fuel gas mass flow to the primary reformer furnace). This envelope is designed to mimic the practical constraints that arise from catalyst stability, heat-integration requirements, and process units' operating ranges [8].

More specifically, we have the following KPIs:

A. The throughput must always be above 80% of the primary reformer's nominal capacity.
B. The operating temperature in the secondary reformer must not exceed design values by more than 100°C.
C. The fuel gas mass flow to the primary reformer furnace must never drop to zero.

Accordingly, a 0–20% range of allowable hybridization factors was identified (as shown in **Table 1**).

Table 1: Evolution of the 3 KPIs (see the previous list) with increasing hybridization factor [8]. Values highlighted in red exceed operational constraints.

Case	Hybrid. factor	KPI A	KPI B	KPI C
Gray (baseline)	0%	100%	±0°C	100%
Hybrid-green (mild)	10%	89.3%	+47°C	92.3%
Hybrid-green (heavy)	20%	78.8%	+109°C	82.8%

MODELING AND OPTIMIZATION FRAMEWORK

Our assessment links a steady-state process model of the ammonia plant with an hourly electricity-to-hydrogen dispatch model and a capacity-sizing optimization. The goal is to quantify the cost and emissions impact of hybridization at the plant level, while explicitly representing the variability of renewable generation and market electricity prices.

Baseline process models (gray and blue)

A conventional natural-gas ammonia plant is simulated in UniSim® Design R500, following a typical flowsheet: desulfurization, primary and secondary reforming, shift conversion, CO_2 removal, methanation, compression, and ammonia synthesis with recycle and purge (*i.e.*, the Haber-Bosch process). Property methods rely on the Soave-Redlich-Kwong equation of state, consistent with high-pressure systems. Baseline performance and operating ranges are cross-checked against industrial manuals and engineering references.

For the blue baseline, building on the gray ammonia model, CO_2 capture is implemented on both the high-purity process stream (*i.e.*, the emissions from CO_2 removal) and flue gas streams (*i.e.*, the emissions generated by the combustion of fuel gas in the primary reformer furnace), using cost and performance assumptions aligned with public techno-economic studies for industrial capture [9, 10] (see **Table 2**). As our assessment is restricted to the production site battery limits following a gate-to-gate approach, the present analysis excludes both CO_2 transportation and storage.

Power-to-Hydrogen subsystem

The PtH subsystem includes: (i) a renewable generation (solar and/or wind) power plant, (ii) an alkaline electrolyzer, (iii) a lithium-ion battery for short-term smoothing, and (iv) a pressurized hydrogen storage system for multi-hour to multi-day balancing. Component costs follow recent public reports and vendor data for low-temperature electrolysis as well as electricity and hydrogen storage technologies (see **Table 2**). Moreover, the design and flexible operating schedule of batteries and hydrogen tanks are supported by frameworks developed by the Authors for Power-to-X applications [11, 12].

Table 2: The complete set of input values and cost assumptions for the present techno-economic assessment.

RENEWABLE ENERGY DATASET		
Location	California	[-]
Year	2024	[-]
Solar capacity factor	29.67%	[-]
Wind capacity factor	30.68%	[-]
Time discretization	1	[h]
TECHNO-ECONOMIC DATA		
Interest rate	5%	[-]
Solar (Photovoltaics)		
Lifetime	25	[y]
CAPEX	1058	[USD/kW]
OPEX	1.7%	[CAPEX/y]
Wind (Onshore)		
Lifetime	25	[y]
CAPEX	1568	[USD/kW]
OPEX	2%	[CAPEX/y]
Electrolyzer (Alkaline)		
Lifetime	10	[y]
CAPEX	700	[USD/kW]
OPEX	2%	[CAPEX/y]
Electrical efficiency	51.7	[kWh/kg$_{H_2}$]
Maximum load flexibility	100%	[-]
Minimum load flexibility	5%	[-]
Battery (Lithium-ion)		
Lifetime	10	[y]
CAPEX	234	[USD/kW]
OPEX	2%	[CAPEX/y]
Hydrogen storage (Tanks @ 200 bar)		
Lifetime	20	[y]
CAPEX	1898	[USD/kg$_{H_2}$]
OPEX	1%	[CAPEX/y]
Ammonia synthesis loop (Haber-Bosch process)		
Target green H_2 feed	1.74	[t$_{H_2}$/h]
Minimum green H_2 feed	0	[t$_{H_2}$/h]
Maximum green H_2 feed	3.55	[t$_{H_2}$/h]
Maximum green H_2 ramp	100%	[load/h]
Carbon capture and compression system		
Process CO_2 emissions	1.17	[t$_{CO_2}$/t$_{NH_3}$]
Furnace CO_2 emissions	0.467	[t$_{CO_2}$/t$_{NH_3}$]
Furnace CO_2 capture costs	60	[USD/t$_{CO_2}$]
Total CO_2 compression costs	15	[USD/t$_{CO_2}$]
California market scenario		
Natural gas price (2024 avg)	37.48	[USD/MWh]
Carbon tax (2024 avg)	38.59	[USD/t$_{CO_2}$]

Electricity market and policy context

Hourly renewable profiles for 2024 are sourced from the California Independent System Operator (CAISO),

representing a high-renewables grid with pronounced diurnal and seasonal patterns [13]. Policy signals are modeled through a carbon price consistent with California's cap-and-trade mechanism [14]. Natural gas prices are based on industrial price indicators referring to industrial delivery rates for large-scale chemical facilities in California [15]. These cost assumptions are detailed in **Table 2**.

Optimization problem

Decision variables are the installed capacities of the solar and wind farm, and the objective is to minimize the levelized cost of hydrogen (LCOH) produced by the PtH plant while meeting the hybridization target and respecting the retrofit envelope constraints.

The LCOH is computed by **Eq. (1)** from annualized CAPEX and OPEX contributions of the PtH subsystem (*i.e.*, solar plant, wind plant, electrolyzer, battery, and hydrogen storage; hereafter denoted by S, W, E, B, and H).

$$LCOH = \frac{\sum_{j=S,W,E,B,H}\left(\frac{CAPEX_j}{\sum_{y=1}^{NY_j}(1+ir)^{-y}}+OPEX_j\right)}{8760\ [\mathrm{h/y}] \bullet \dot{m}_{H_2}^{target}} \quad (1)$$

Optimization proceeds through a systematic evaluation of 1681 solar-wind installed capacity combinations (0–400 MW range, at 10 MW resolution) using MATLAB R2024a, followed by an estimation of the required electricity and hydrogen storage capacities according to the frameworks established by Isella and Manca [11, 12].

RESULTS AND DISCUSSION

Optimal renewable and PtH sizing

The sizing optimization reveals a clear benefit of combining solar and wind. Solar-only systems require larger batteries and hydrogen storage to bridge diurnal variability, while wind-only systems reduce storage needs but still show considerable electrolysis capacities and power generation costs. Lastly, the mixed solar-and-wind case leverages complementarity, reducing both overbuild and storage requirements, which translates into a lower LCOH (see **Table 3**).

Table 3: Optimal installed capacities and resulting LCOH for solar, wind, and solar-wind PtH configurations.

	Solar	Wind	Solar-Wind
Solar [MW]	330	0	110
Wind [MW]	0	310	210
Electrolyzer [MW]	243.20	147.55	129.88
Battery [MW]	216.45	16.82	8.47
H_2 storage [t_{H_2}]	19.52	2.16	2.95
LCOH [USD/kg_{H_2}]	3.75	3.25	2.94

Ammonia production costs

Figure 2 shows the ensuing levelized cost of ammonia (LCOA) for gray, blue, and hybrid pathways, including hybrid-green and "hybrid-blue-green" (*i.e.*, hybrid-green plus carbon capture). The gray baseline reflects fuel and feedstock costs plus the carbon price burden. The blue pathway benefits from captured process CO_2 but includes carbon capture costs. Hybrid-green pathways show higher LCOA in this scenario, primarily because renewable hydrogen remains costlier than reforming-based hydrogen under 2024 assumptions, despite rapidly declining renewable generation costs.

Numerically, the LCOA values (in USD/t_{NH_3}) resulting from this case study are: gray 373.44, blue 364.51, hybrid-green (solar-only) 415.04, hybrid-green (wind-only) 404.50, hybrid-green (solar-wind) 398.03, and hybrid-blue-green (solar-wind) 395.27. Such results emphasize that, in a high-gas-price moderate-carbon-price context, blue ammonia can be competitive with gray while still reducing direct emissions. On the other hand, hybrid-green becomes more competitive as renewable electricity and electrolysis costs fall further or carbon prices rise.

The cost breakdown provided helps interpret which components dominate the incremental cost of hybridization. In the hybrid-green cases, noteworthy contributors are renewable CAPEX and electrolyzer CAPEX, followed by storage (electricity and hydrogen) to buffer variability. In the hybrid-blue-green case, the addition of CCS introduces capture costs but reduces carbon liabilities and can improve competitiveness depending on policy and CO_2 storage or utilization availability.

Emissions performance and CCS leverage

The operating envelope enables decarbonization progress without forcing deep turndown of the synthesis loop. In the study, increasing the hybridization factor from 0 to 10% and 20% reduces the plant-level CO_2 intensity from 1.64 to 1.48 and 1.31 t_{CO_2}/t_{NH_3}, respectively. The reduction is driven by the partial displacement of reformer-derived hydrogen by renewable electricity-derived hydrogen.

A key insight is that hybridization and CCS are not mutually exclusive. Hybrid-blue-green combines renewable hydrogen displacement with CO_2 capture on remaining fossil-derived streams. This stacking can reduce residual emissions beyond what either measure achieves alone while maintaining high utilization of existing assets. Such combinations align with CCS roadmaps for industrial decarbonization, where early deployment targets high-concentration CO_2 sources and leverages existing infrastructure.

At the same time, the environmental narrative of blue pathways depends on upstream methane emissions and on the overall capture rate achieved at the plant level. Recent debate emphasizes that "how green" blue

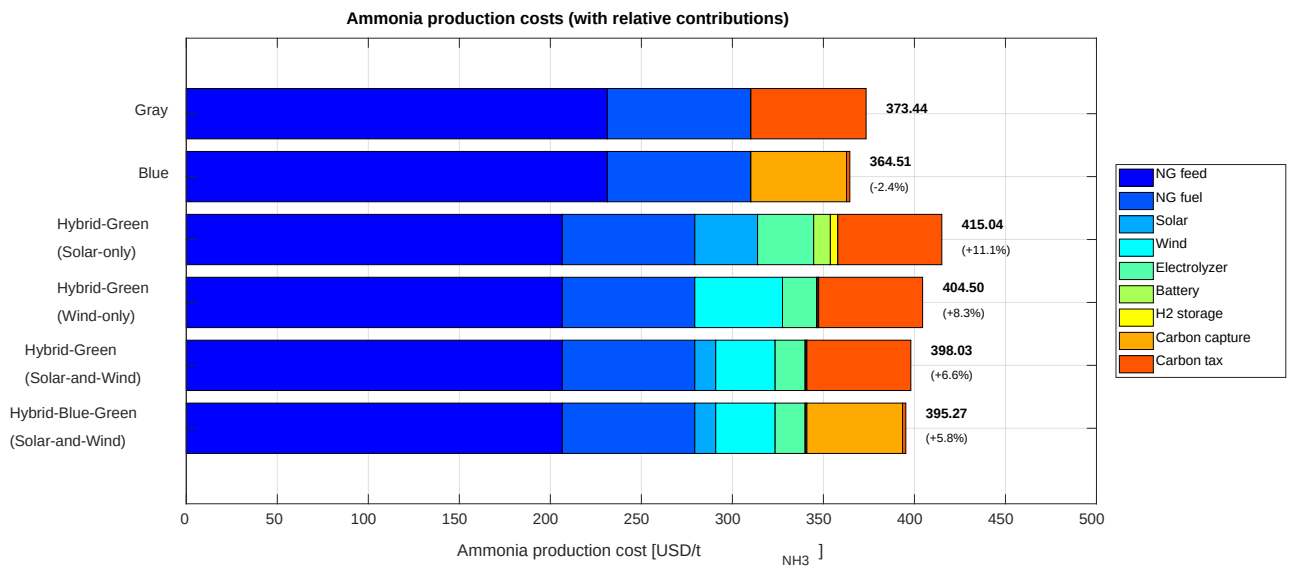

Figure 2. Ammonia production costs and relative contributions for the six assessed configurations. Percentages in parentheses next to the bars indicate the relative difference from the gray ammonia baseline. (NG=Natural Gas)

hydrogen and blue ammonia are depends critically on leakage rates and on capture scope [5]. For this reason, hybridization can be viewed as a hedge: increasing renewable substitution reduces exposure to upstream gas supply emissions even if CCS deployment faces constraints.

Operational insights for retrofit deployment

From an operational perspective, the hybrid retrofit can be implemented with minimal disruption if the PtH subsystem is treated as an auxiliary unit that delivers hydrogen at a suitable mass flow. The process integration challenge is to ensure that hydrogen displacement does not violate reformer heat-balance constraints or alter the synthesis loop's inert levels and purge management. Both scientific literature and industrial practice suggest that maintaining stable inlet compositions and temperature profiles is crucial to reliable operation.

Storage plays a key role: it buffers variability so that the plant receives a smoother electricity and hydrogen supply that is also compliant with the operating requirements of the electrolyzer and the Haber-Bosch process. Battery storage is most valuable for short-term smoothing, whereas hydrogen storage provides longer balancing. In this respect, recent design criteria for flexible electricity and hydrogen storage in PtX systems were applied to select storage capacity without overinvesting in rarely used inventory [11, 12].

The optimal sizing results show that solar-only operation requires significantly larger battery and hydrogen storage capacities than wind-only or mixed operation. This indicates that, for retrofit cases with limited space or where minimizing additional equipment is essential, the availability of a diversified renewable portfolio (either on-site or via PPAs) can be as important as absolute generation cost.

CONCLUSIONS

A hybrid-green retrofit concept can meaningfully reduce ammonia plant CO_2 intensity (about 20% decrease) while preserving continuous operation through an explicit operating envelope. When PtH assets are optimally sized, solar-wind complementarity cuts storage capacities and delivers the lowest LCOH among the tested cases. In the California 2024 scenario, hybrid-green remains more expensive than blue ammonia; however, the gap narrows with lower electrolyzer costs and higher carbon prices. Hybrid-blue-green offers an attractive transitional option by stacking renewable displacement and CCS.

From a deployment viewpoint, our framework supports retrofit screening by quantifying (i) the additional equipment required for a given hybridization target, (ii) the cost drivers that dominate incremental LCOA, and (iii) the emission reductions achievable without major process redesign. The same methodology can be extended to other regions by updating market prices, carbon policies, and technology costs.

AUTHOR IDENTIFIERS

Author ORCIDs:
Isella A: 0000-0002-3699-4698
Manca D: 0000-0003-2055-9752

REFERENCES

1. The Royal Society. Ammonia: Zero-Carbon Fertiliser, Fuel and Energy Store. The Royal Society (2020).
2. Isella A, Manca D. GHG emissions by (petro)chemical processes and decarbonization priorities—a review. Energies 15:7560 (2022). https://doi.org/10.3390/en15207560
3. Armijo J, Philibert C. Flexible production of green hydrogen and ammonia from variable solar and wind energy: case study of chile and argentina. International Journal of Hydrogen Energy 45:1541-1558 (2020). https://doi.org/10.1016/j.ijhydene.2019.11.028
4. Isella A, Manca D. Integrating economic, environmental, and social sustainability in power-to-ammonia plants: a multi-objective optimization methodology. Computers & Chemical Engineering 198:109082 (2025). https://doi.org/10.1016/j.compchemeng.2025.109082
5. Howarth RW, Jacobson MZ. How green is blue hydrogen?. Energy Science & Engineering 9:1676-1687 (2021). https://doi.org/10.1002/ese3.956
6. Smith E, Morris J, Kheshgi H, Teletzke G, Herzog H, Paltsev S. The cost of CO2 transport and storage in global integrated assessment modeling. International Journal of Greenhouse Gas Control 109:103367 (2021). https://doi.org/10.1016/j.ijggc.2021.103367
7. Isella A, Lista A, Colombo G, Ostuni R, Manca D. Gray and hybrid green ammonia price sensitivity to market fluctuations: the Russia-Ukraine war case. In: Computer Aided Chemical Engineering. Ed: Kokossis AC, Georgiadis MC, Pistikopoulos E. Elsevier (2023) https://doi.org/10.1016/B978-0-443-15274-0.50364-4
8. Isella A, Ostuni R, Manca D. Towards the decarbonization of ammonia synthesis – a techno-economic assessment of hybrid-green process alternatives. Chemical Engineering Journal 486:150132 (2024). https://doi.org/10.1016/j.cej.2024.150132
9. Cornils B. Haldor topsøe low?energy ammonia process. Catalysis from A to Z : (2020). https://doi.org/10.1002/9783527809080.cataz07654
10. Global CCS Institute. Advancements in CCS Technologies and Costs. Global CCS Institute (2025).
11. Isella A, Manca D. A framework for the design of battery energy storage systems in power-to-x processes. Journal of Energy Storage 123:116744 (2025). https://doi.org/10.1016/j.est.2025.116744
12. Isella A, Manca D. A general criterion for the design and operation of flexible hydrogen storage in power-to-x processes. International Journal of Hydrogen Energy 99:649-660 (2025). https://doi.org/10.1016/j.ijhydene.2024.12.228
13. California Independent System Operator (CAISO). http://www.caiso.com
14. World Bank Group. https://carbonpricingdashboard.worldbank.org/compliance/price
15. U.S. Energy Information Administration (EIA). https://www.eia.gov/dnav/ng/ng_pri_sum_a_EPG0_PIN_DMcf_a.htm

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimizing Steam flux for Energy efficiency in Ammonia Recovery during Sodium carbonate production

Ediane S. Alves[a], Mohamad A. Chahine[a], Denis Guillaume[a], and Julien Gornay[a*]

[a] IFP Energies nouvelles, Rond-point de l'échangeur de Solaize, BP 3, 69360 Solaize, France
* Corresponding Author: julien.gornay@ifpen.fr.

ABSTRACT

Industrial decarbonization is crucial to reducing global emissions. Efficient processes lower energy use and reduce the environmental impacts, such as material use and waste, decreasing the overall industrial footprint. In this context, the present study explores the impact of reducing steam consumption (thermal energy) during the ammonia regeneration process in the production of sodium carbonate. A key feature of the Solvay process is ammonia recycling, which significantly reduces raw material consumption and ensures both economic and environmental sustainability. However, this stage is highly energy-intensive. To enhance energy efficiency in soda ash production, a study was conducted to analyze variation in temperature, pressure, and steam flow introduced into the ammonia regeneration system. The objective is to understand its impact on both ammonia recovery and the process's energy consumption. Variations in steam pressure do not impact on energy consumption of the process. By reducing the temperature and steam flow introduced into the system, in order to reduce the energy consumption, it was found that there is no decrease in ammonia recovery. The reduction in both the total injected steam flow and operating temperature led to lower heat exchange requirements and a decreased the cooling energy demand, highlighting the potential for optimizing steam flow as an effective strategy to improve process efficiency. Instead, it underscores a potential for energy optimization, promoting a more sustainable and cost-effective operation.

Keywords: Energy Efficiency, Modelling and Simulations, Aspen Plus, Energy, Optimization

INTRODUCTION

Sodium carbonate is a key industrial chemical with widespread applications across several sectors, including glass manufacturing, paper production, detergents, and the chemical industry. Its strategic importance is reflected in the steady growth of the global sodium carbonate market, primarily driven by the expansion of the glass and detergent industries [1]). This trend is further reinforced by increasing urbanization and industrialization, particularly in emerging economies. According to market projections, the global sodium carbonate market is expected to reach USD 14.42 billion in 2025 and expand to approximately USD 25.74 billion by 2035 [1].

The Ammonia Process or Solvay Process, a widely used and efficient method to produce sodium carbonate also known as soda ash. The production of sodium carbonate through the Solvay process is divided into four stages. The first stage is limestone calcination, in which calcium carbonate ($CaCO_3$) is thermally decomposed to produce carbon dioxide (CO_2) and calcium oxide (CaO). The carbon dioxide generated in this step is subsequently utilized in the following stage. In the second stage, carbonation, the carbon dioxide is bubbled through a brine solution containing sodium chloride (NaCl), ammonia (NH_3), and water. Within this medium, carbon dioxide reacts with ammonia and sodium ions, leading to the formation of sodium bicarbonate ($NaHCO_3$), which precipitates due to its low solubility. Sodium bicarbonate constitutes the desired intermediate product, whereas ammonium chloride (NH_4Cl) is simultaneously formed as a by-product. The third stage involves the thermal decomposition of sodium bicarbonate. The solid obtained from carbonation is separated by filtration and subsequently subjected to calcination. Upon heating, sodium bicarbonate decomposes into sodium carbonate

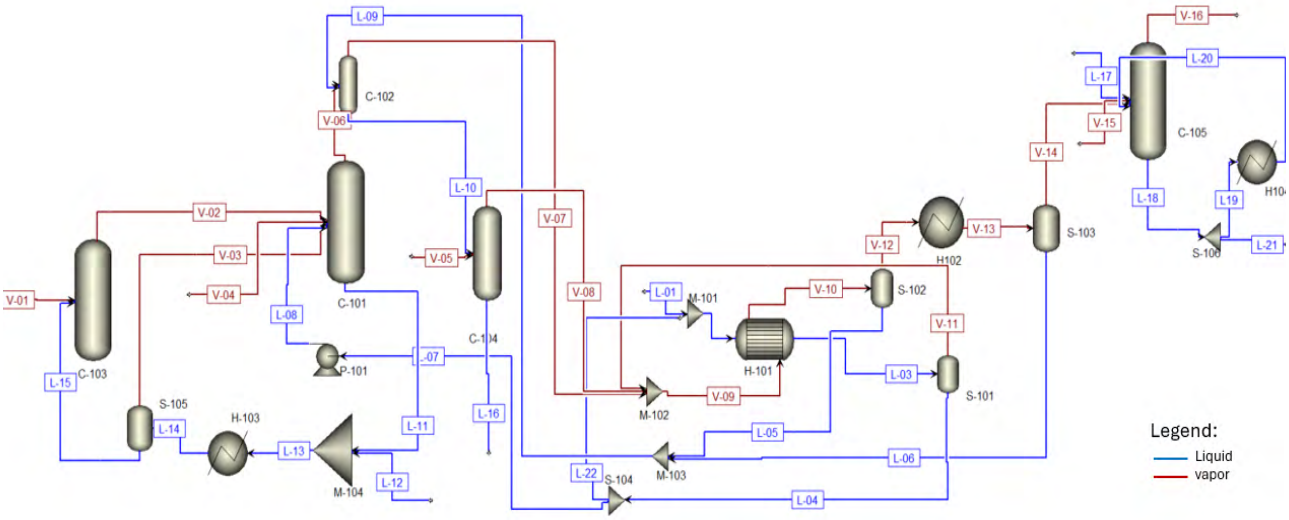

Figure 1: Ammonia regeneration process flow diagram in sodium carbonate production.

(Na_2CO_3), water, and carbon dioxide. The regenerated carbon dioxide is recycled within the process, contributing to its overall sustainability. Finally, the fourth stage is ammonia regeneration. In this step, ammonium chloride reacts with hydrated calcium oxide ($Ca(OH)_2$), yielding free ammonia and calcium chloride ($CaCl_2$). The recovered ammonia is recycled into the process through absorption in brine, thus closing the operational cycle and ensuring the economic and environmental viability of the Solvay process [2].

The production of sodium carbonate is both energy-intensive and CO2-intensive [3], in part because the ammonia recovery step requires a substantial amount of steam. Despite this significant energy requirement, ammonia recovery is the operation that ensures the economic viability of the Solvay process when compared with alternative sodium carbonate production routes (Book: Manufacture of Soda). Nevertheless, even though it plays an important role in process efficiency and overall performance, the ammonia recovery step remains insufficiently addressed in the scientific literature. Most existing studies focus on eliminating the ammonia recovery stage [4, 5], replacing ammonia as the process intermediate [6, 7] or integrating carbon capture strategies into the Solvay process [8]. However, although such strategies can reduce environmental impacts, they require major modifications to existing industrial plants, which in turn demand substantial capital investment. For this reason, these modified methods have generally not been favored by manufacturers due to economic considerations[9].

The ammonia regeneration process investigated in this study presents significant advantages with respect to process efficiency, operational costs, and environmental performance. The effective recovery of ammonia and carbon dioxide reduces waste streams and decreases the requirement for fresh ammonia make-up, thereby lowering operating expenses. To further improve the energy performance of soda ash production, this study focuses on assessing the impact of steam flow rate variations in the ammonia regeneration stage. The analysis evaluates the relationship between steam input, ammonia recovery efficiency, and overall energy consumption, providing a technical basis for the optimization of industrial-scale sodium carbonate production systems.

METHODOLOGY

Thermodynamic model

In this study, process modeling is performed in Aspen Plus® V14.0. The Electrolyte Wizard is used to identify possible reactions, components, and intermediates, and to calculate equilibrium constants and properties. For systems with inorganic electrolytes, such as sodium bicarbonate, ionic species and reactions are also considered. The thermodynamic model ELECNRTL-RK is applied, combining the NRTL electrolyte model for the aqueous phase with the Redlich-Kwong equation of state for the vapor phase. This thermodynamic model has been validated in previous studies, such as the methodology used by Yoshi et al. (2022) [10], specifically in the simulation of the Solvay process.

Process simulation

The simulated system (Figure 1) is based on ammonia regeneration, focusing on the processing of cold

mother liquor (L-01), which is predominantly composed of water and ammonium chloride (NH_4Cl), also known as fixed ammonia, as well as dissolved carbon dioxide (CO_2) and free ammonia (NH_3). Initially, L-01 is directed to the heat exchanger (H-101), where it absorbs heat for preheating, causing part of the dissolved gasses to gas phase. After this stage, the heated fluid proceeds to a gas separator (S-101), where the liquid and vapor phases are separated. The liquid phase is then sent to the column C-101.

Column C-101, designed as a stripping unit, functions to remove carbon dioxide (CO_2) and free ammonia present in the mother liquor. The liquid effluent from this stage is sent to the mixer (M-104), where calcium hydroxide ($Ca(OH)_2$) is added. The chemical reaction occurring here can be detailed as follows: $Ca(OH)_2 + 2NH_4Cl \rightarrow 2NH_3 + CaCl_2 + 2H_2O$. The chemical reaction between calcium hydroxide and ammonium chloride results in the formation of free ammonia and calcium chloride ($CaCl_2$). Subsequently, the mixture is transferred to a foam separator, whose primary function is to remove the gases formed, mainly composed of ammonia. The vapor extracted in this separator is redirected to column C-101, while the liquid phase proceeds to column C-103.

Column C-103, also configured as a stripping unit, is responsible for removing the free ammonia generated in the previous stage (M-104). The liquid effluent from this column is now characterized as de-ammoniated mother liquor. The vapors generated in C-103 are sent back to column C-101, along with the vapor from the foam separator and an additional flow of steam introduced into the system.

The vapor from the top of column C-101 is directed to column C-102, another stripping unit. C-102 receives condensed water from the heat exchangers in the ammonia regeneration system (H-101 and H-102). At this stage, ammonia and carbon dioxide present in the condensates are recovered and reincorporated into the vapor flow. The liquid collected at the bottom of C-102 is sent to a second stripping column (C-104), where residual NH_3 and CO_2 in the condensate are transferred to the vapor phase.

The ammonia recovered in systems C-102, C-104, and S-101 is routed to the heat exchangers (H-101 and H-102) for cooling before being introduced into the absorber. In this equipment, the gas is absorbed by the NaCl brine, allowing its reintegration into the sodium carbonate production process.

The formation of ammonia brine is accompanied by the release of a significant amount of heat, characterizing it as an exothermic process. To control the temperature and optimize NH_3 absorption, a pre-cooled ammonia brine (H-104) is recycled in the system, ensuring ideal conditions for the continuity of the process.

Test optimization

The process has three steam inputs: the V-01 stream, which is added to the stripping column C-103 and represents the main steam flow of the system, with a flow rate of 38, 700 kg/h; the V-04 stream, a supplemental steam flow with a flow rate of 3, 000 kg/h, directed to column C-101; and the V-05 stream, with a flow rate of 5, 573 kg/h, introduced into column C-104.This study analyzed the variation in temperature, pressure, and mass flow rate of steam introduced into the fixed ammonia regeneration process, with the aim of understanding the impact of these variables on ammonia recovery efficiency. The temperature, pressure, and mass flow conditions of these three steam streams were modified as presented in Table 1, one variable at a time. The proposed variations are based on values practiced by the soda ash production industry.

Table 1: Variation in the temperature, pressure, and mass flow rate of the process's steam inputs.

Stream	Temperature	Pressure	Mass Flow
V-01	116-145 °C	1.68-1.80 bar	10% reduction
V-04	116-145 °C	1.30-1.50 bar	10% reduction
V-05	116-145 °C	1.30-1.50 bar	10% reduction

RESULTS AND DISCUSSION

The effects of temperature, pressure, and mass flow rate variations on the vapor flow rate the main equipment of the ammonia regeneration process are analyzed, and the effect of mass flow variation is presented in Figure 2. The mass flow rate results of the system's vapor streams were normalized with respect to the mass flow rate obtained when the analyzed variable was at its maximum value. In addition, the results highlight the effect of reducing the analyzed variable on system behavior; therefore, in Figures 2, the *x*-axis is presented from the highest to the lowest value of the variable.

The variation in temperature and pressure has a minimum impact on ammonia recovery Similarly to the influence on the process steam flow, some variations are so small that they can be attributed to changes caused by model convergence. The most significant change observed was a 2% reduction in the steam flow exiting the C-105 column, which does not impact ammonia recovery or the stream entering the absorption column.

The low influence of temperature and pressure variations on ammonia recovery and steam mass flow rate in the system indicates that, despite the reduction in thermal energy transported by the steam stream fed into the system columns, the energy supplied is sufficient to achieve ammonia recovery. Under these conditions, the steam acts predominantly as an entraining agent, and

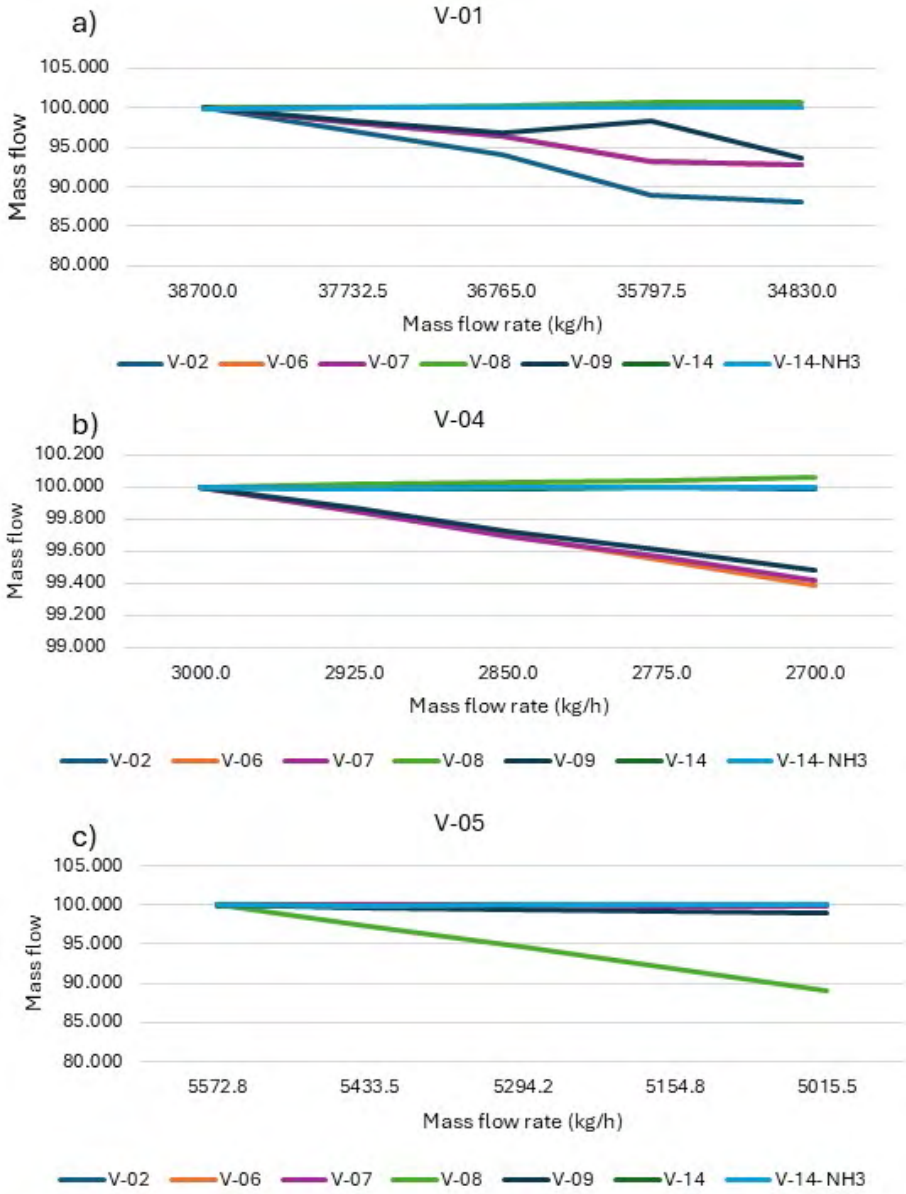

Figure 2. Influence of vapor mass flow rate input on the vapor mass flow rate for different vapor outputs of equipment: a) vapor input V-01, b) vapor input V-04, and c) vapor input V-05.

small variations in its temperature or pressure do not result in significant changes in the thermodynamic gradient required for ammonia vaporization. Despite the reduction in thermal energy introduced into the system, it continues to operate without major alterations. More details about the influence of temperature and pressure variations on ammonia recovery and steam mass flow rate in the system in the supplementary material.

The reduction in steam flow introduced into the system (Figure 2) implies changes in the dynamics of the vapor streams, although with variable impacts depending on the stream analyzed. In the case of stream V-04, the reduction in steam flow is minimal, less than 1% in streams V-06, V-07, and V-09, demonstrating little sensitivity to this change. On the other hand, stream V-05 shows a significant reduction of approximately 11% in the V-08 flow, which originates from column C-104, a column where the V-05 vapor stream is added, directly impacting the amount of vapor exiting this column. For the other system streams, including the absorption column's inlet stream, no significant impact is observed, with the NH_3 percentage remaining constant. The low influence of the reduction in the flow of V-04 and V-05 is due to the small change in the total energy supplied to the system. The

flow of V-04 represents only 6% of the steam introduced into the system, and a 10% reduction in this value results in a decrease of less than 1% in the total heat supplied to the system.

In the case of stream V-01, which is the main source of steam for the system, its reduction has a considerable impact on several vapor streams. This occurs because the steam introduced into column C-103 flows through multiple units within the system and accounts for approximately 82% of the total steam supplied to the process. This reduction, however, does not compromise ammonia recovery, which remains constant, as does the performance of the vapor stream feeding the absorption column (line v-14 and V-14NH3, which corresponds to the mass flow rate of ammonia entering the absorption column).

The observed behavior can be attributed to the thermodynamic properties of the involved compounds. Due to the high volatility of ammonia, even with the reduction in steam flow in the stripping columns, ammonia continues to efficiently migrate to the vapor phase. On the other hand, a larger amount of water remains in the liquid phase, as water requires more energy to transition to the vapor phase.

Table 2: Energy balance exchanged by the equipment streams in different equipment for various pressure of steam stream introduced into the system.

Pressure	V-01	C101	C102	C103	C104	H101	H102
	bar	Gcal/h	Gcal/h	Gcal/h	Gcal/h	Gcal/h	Gcal/h
V-01	1.80	38.86	12.44	23.33	3.73	-11.27	-8.77
	1.68	38.86	12.45	23.33	3.73	-11.27	-8.77
V-04	1.50	38.86	12.45	23.33	3.73	-11.28	-8.78
	1.30	38.85	12.45	23.33	3.73	-11.28	-8.77
V-05	1.50	38.85	12.46	23.33	3.74	-11.28	-8.78
	1.30	38.85	12.45	23.33	3.73	-11.28	-8.78

In the system's heat exchangers, the reduction in steam mass flow rate impacts the performance of heat transfer. In heat exchanger H-101, it is observed that the effluent steam exhibits a lower temperature compared to the scenario where the system operates with a higher steam flow rate, and a smaller amount of steam is condensed during the heat exchange process. More details regarding the temperature variation in the steam streams are available in the supplementary material. In the subsequent heat exchanger, H-102, designed to ensure that the steam reaches a temperature of 64 °C, required for entry into the absorption column, the reduction in steam mass flow rate does not alter the final steam flow rate coming from the ammonia regeneration system or the ammonia concentration present. The only modification occurs in the fraction of steam that is condensed at this stage of the process.

To better understand the heat exchange occurring within the process equipment, Tables 2, 3, and 4 present the energy transfer in the liquid phase between process streams at the maximum and minimum values of the following variables: pressure, temperature, and mass flow rate, respectively. The analysis of energy variation in a stripping column is highly complex, since the enthalpy of the liquid stream is a function of temperature, pressure, and chemical composition. The latter changes along the column as a result of mass transfer phenomena between the liquid and vapor phases. Consequently, an increase in the energy associated with the outlet stream may be attributed, among other factors, to the transfer of vapor to the liquid phase, which supplies latent heat to the system. In order to simplify the overall energy analysis, the parameter adopted to quantify the reduction of the system's thermal energy is the heat duty removed in the heat exchangers responsible for vapor cooling. These units are located at the end of the process and are designed to remove heat from the vapor stream, thus constituting the primary mechanism for energy removal from the system.

Table 3: Energy balance exchanged by the equipment streams in different equipment for various temperature of steam stream introduced into the system.

Temperature	V-01	C101	C102	C103	C104	H101	H102
	°C	Gcal/h	Gcal/h	Gcal/h	Gcal/h	Gcal/h	Gcal/h
V-01	145	38.94	13.31	19.75	3.75	-11.27	-9.25
	116	38.86	12.43	23.33	3.73	-11.27	-8.77
V-04	145	38.86	12.43	23.33	3.73	-11.27	-8.77
	116	39.13	12.42	23.34	3.73	-11.27	-8.74
V-05	145	38.85	12.46	23.33	3.73	-11.28	-8.78
	116	38.86	12.40	23.33	4.16	-11.27	-8.72

In the heat exchangers, the thermal behavior shows specific variations. Heat exchanger H-101 maintains its thermal exchange capacity almost unchanged in the three tables, as the equipment operates within its design range, allowing energy exchange between the hot and cold fluid streams to occur small variations. On the other hand, in heat exchanger H-102, designed to cool the vapor before its entry into the absorption column, the reduction in mass flow rate and vapor temperature results in a lower amount of heat exchanged. This reduction represents the most significant thermal alteration observed in the system, indicating a direct adaptation to the decrease in vapor flow introduced.

In Table 2, it is observed that the heat exchanged in the H-102 exchanger does not present significant variation, except for a reduction of 0.01 Gcal/h in stream V-04, indicating that pressure changes do not result in relevant

alterations in the system's thermal exchange. Changes in the temperature and mass flow rate of the vapor stream show, as presented in Tables 3 and 4, a reduction in energy consumption in the H-102 exchanger, with the most significant impact observed when modifying stream V-01, as it is the most important vapor stream in the process. Although both strategies promote a reduction in energy consumption, the reduction in vapor flow led to more pronounced decreases compared to temperature variation, considering the range of variation proposed in this study.

Table 4: Energy balance exchanged by the equipment streams in Different Equipment for Various Mass Flow Rates of V-01 Introduced into the System.

Mass flow	V-01	C101	C102	C103	C104	H101	H102
	Kg/h	Gcal/h	Gcal/h	Gcal/h	Gcal/h	Gcal/h	Gcal/h
V-01	38700	38.86	12.43	23.33	3.73	-11.27	-8.77
	34830	38.48	10.40	23.39	3.66	-11.27	-6.88
V-04	3000	38.85	12.90	23.33	3.73	-11.28	-8.79
	2700	38.88	12.67	23.33	3.73	-11.28	-8.66
V-05	5573	38.86	12.48	23.33	3.73	-11.27	-8.78
	5015	38.86	12.06	23.33	3.73	-11.27	-8.54

Additionally, the decrease in energy required to cool the vapor in the absorption column reflects the impact of the reduced amount of vapor generated in the system. This change not only reduces the overall energy consumption of the process but also improves the thermal efficiency of the system by operating with lower energy demands without compromising operational stability.

This scenario reinforces the robustness of the system in response to adjustments in vapor flow, demonstrating that the reduction in the mass flow rate of vapor V-01 does not compromise critical operations, such as ammonia removal or the efficiency of the heat exchangers. On the contrary, it highlights an opportunity for energy optimization, enabling sustainable and economically advantageous operation.

However, it is important to emphasize that the present work is strictly theoretical in nature and is therefore subject to the inherent limitations of mathematical models and the simplifying assumptions adopted. While the findings indicate potential operational benefits, the analyzed system exhibits high complexity, as it involves an electrolytic system where the formation of solid species (salts) may occur depending on the operational conditions of composition, temperature, and pressure. The precipitation of these species can lead to the formation of deposits on the system's surfaces, resulting in a loss of operational efficiency, changes in mass and energy transfer coefficients, as well as potential hydrodynamic impacts. Although the model calculates the formation of solids, such phenomena were not incorporated into the model used, which may limit the applicability of the results to real systems.

CONCLUSION

Process optimization plays a fundamental role in industrial decarbonization, establishing itself as a viable strategy for reducing energy consumption and inputs in industrial processes. The results of this study demonstrate that variations in temperature and pressure do not significantly impact the system's operational performance (flow rate, temperature, composition), allowing adjustments in these variables without compromising overall operation. However, only the reduction in temperature resulted in a 5% decrease in energy consumption. Among the evaluated strategies, reducing the flow of steam injected into the process proved to be the most effective, promoting considerable reductions in steam flows, temperatures, and heat exchange, without compromising the efficiency of ammonia recovery, a crucial factor for process sustainability. These results suggest that there is significant potential for operational adjustments in the steam used in the system, even within the limits already practiced in the industry. However, to achieve even more significant advances in the process's energy efficiency, more detailed investigations involving other operational parameters are necessary. Future studies may focus on the interaction between different process variables, the impact of extreme conditions, and the application of emerging technologies, aiming to identify additional opportunities for optimization and emission reduction, thereby strengthening the commitment to industrial sustainability.

DIGITAL SUPPLEMENTARY MATERIAL

Further details on the effects of variations in steam pressure and temperature on the system temperature and mass flow rate are provided in the digital supplementary material (LAPS.2026.0023)

ACKNOWLEDGEMENTS

This work was made possible thanks to the partnership with Novacarb, which provided technical support for the industrial representativeness of this study. This work was supported by a grant from the French government, managed by the National Research Agency under the France 2030 initiative: DCARBO project within the PEPR SPLEEN program, reference number 22-PESP-0002.

AUTHOR IDENTIFIERS

Author ORCIDs:

Alves ES: 0000-0002-2398-6226
Chahine MA: 0009-0000-9509-5476
Guillaume D: 0000-0003-0531-5071
Gornay J: 0009-0000-7772-5446

REFERENCES

1. Precedence Research. https://www.precedenceresearch.com/sodium-carbonate-market
2. Te-Pang H. Manufacture of soda, with special reference to the ammonia process: A practical treatised. 2ª ed. Reinhold Publishing Corporation, (1942).
3. Liu Z. National carbon emissions from the industry process: production of glass, soda ash, ammonia, calcium carbide and alumina. Applied Energy 166:239-244 (2016). https://doi.org/10.1016/j.apenergy.2015.11.005
4. Petrauskas JV, Faria SHB, Moreira WM, Bonfim-Rocha L. Simulation of sodium bicarbonate purification via the sodium sulfate carbonation route. Processes 12:1687 (2024). https://doi.org/10.3390/pr12081687
5. Abd El?Samad HM, Shaddad MAK. Comparative effect of sodium carbonate, sodium sulphate, and sodium chloride on the growth and related metabolic activities of pea plants. Journal of Plant Nutrition 19:717-728 (2008). https://doi.org/10.1080/01904169609365155
6. Mourad AAI, Mohammad AF, Altarawneh M, Al?Marzouqi AH, El?Naas MH, Al?Marzouqi MH. Effects of potassium hydroxide and aluminum oxide on the performance of a modified solvay process for CO ₂ capture: a comparative study. Int J Energy Res 45:13952-13964 (2021). https://doi.org/10.1002/er.6737
7. Mourad AAHI, Mohammad AF, Al-Marzouqi AH, El-Naas MH, Al-Marzouqi MH, Altarawneh M. Koh-based modified solvay process for removing na ions from high salinity reject brine at high temperatures. Sustainability 13:10200 (2021). https://doi.org/10.3390/su131810200
8. Mohammad AF, Mourad AAH, Al-Marzouqi AH, El-Naas MH, Van der Bruggen B, Al-Marzouqi MH. Multistage modified solvay process based on calcium oxide for carbon dioxide capture and reject brine desalination. Separation and Purification Technology 328:125000 (2024). https://doi.org/10.1016/j.seppur.2023.125000
9. Rahimpour H, Fahmi A, Zinatloo-Ajabshir S. Toward sustainable soda ash production: a critical review on eco-impacts, modifications, and innovative approaches. Results in Engineering 23:102399 (2024). https://doi.org/10.1016/j.rineng.2024.102399
10. Hungaro Yoshi HCM, Silva AB, de Souza M, Bernardo de Faria SH, Bernardo A, Bonfim-Rocha L. Multi-criteria assessment of sodium bicarbonate optimized production through CO2 utilization strategies. Journal of Cleaner Production 349:131419 (2022). https://doi.org/10.1016/j.jclepro.2022.131419

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Structural Constraints with the P-graph Framework: Application to an Ammonia Synthesis Process

Darrick Hillaby[a], Andrés Piña Martinez[a]*, Jean-François Portha[a]*, and Laurent Falk[a]

[a] Université de Lorraine, CNRS, LRGP, Nancy, F-54000, France
* Corresponding Authors: andres-david.pina-martinez@univ-lorraine.fr, jean-francois.portha@univ-lorraine.fr.

ABSTRACT

An optimized flowsheet can be generated by numerous approaches. Process optimization via superstructure is one of the methods used to provide solutions that consider the interactions between different decision layers. A process simulator-based optimization is considered in this work, as it offers a reliable and rigorous modeling environment. It is then coupled with a P-graph-based framework to reduce the tedious mathematical writing of the logical constraints to guarantee the structural coherence of a sequence of unit operations.The developed framework consists of three algorithms. The first algorithm transforms the superstructure flowsheet into a P-graph. The second algorithm gets process sub-structures from the superstructure by searching for active units corresponding to a set of decisions made, for example, by an optimizer. The third one checks structural feasibility by verifying that the resulting structure satisfies the five axioms of the original P-graph framework and two additional connectivity tests proposed in this work.The proposed methodology is then applied to an ammonia synthesis process. The results obtained in this work enable a reduction in structural logical constraints while achieving outcomes equivalent to those of a previous study, proving the methodology's efficacy.

Keywords: Process simulator-based optimization, superstructure, P-graph framework, process synthesis, MINLP problem

INTRODUCTION

Climate ambition announced at COP28 requires drastic actions to be taken, such as tripling global renewable energy capacity and doubling the global average annual rate of energy efficiency improvements [1]. One solution to achieve the ambition is the use of modular and flexible plants [2]. Robust processes capable of handling variable inputs, flexible capacity, product variations and decentralized plants are highly valuable to the modern industry. However, modular processes require robust design methods and different optimization strategies as they must adapt to variable inputs or process parameters while maintaining increasing process efficiency. It is therefore essential to employ process design techniques, specifically process synthesis, to design modular processes.

Process synthesis is the selection and interconnection of units involving physical and chemical phenomena to obtain products from raw materials to satisfy economic, environmental and/or social objectives [3]. Main approaches include hierarchical decomposition, mathematical programming and even hybrid approaches combining both heuristics and mathematical programming. Mathematical programming-based strategies rigorously search through a proposed design space for the optimal configuration, unlike hierarchical decomposition. Depending on the definition of the search space, mathematical approaches can be categorized as superstructure-based or superstructure-free optimization [4, 5]. The superstructure-based approach consists of (a) postulating a set of process alternatives in superstructure form, (b) formulating it as a mathematical programming model, and (c) determining the optimal configuration with an optimization algorithm. Conversely, superstructure-free approaches create new flowsheets either from existing ones or from scratch.

In superstructure-based process synthesis, the superstructure can be modeled with different levels of detail. Models can be characterized into three main categories: high-level aggregate, shortcut, and detailed rigorous models. When detailed modelling is required, process

simulators offer a reliable and rigorous modeling environment [6]. In this context, superstructure optimization may be performed in two ways: in an external optimization environment or as a flowsheet in the process simulator itself.

This work focuses on the latter. To avoid extensive mathematical formulations for ensuring structural coherence, a P-graph-based framework is proposed [7]. The process superstructure is implemented in a simulator with the defined optimization problem (objective function, variables, and constraints). As shown in **Figure 1**, a P-graph-based feasibility block interacts with the simulator and an external optimizer to obtain the optimal solution. This feasibility block can be further decomposed into three algorithms.

The paper is structured as follows. First, the methodology is presented, including definitions and algorithm descriptions. Next, its application to an ammonia synthesis process is outlined. Finally, the approach is evaluated against a non–P-graph method, and the optimization results are discussed.

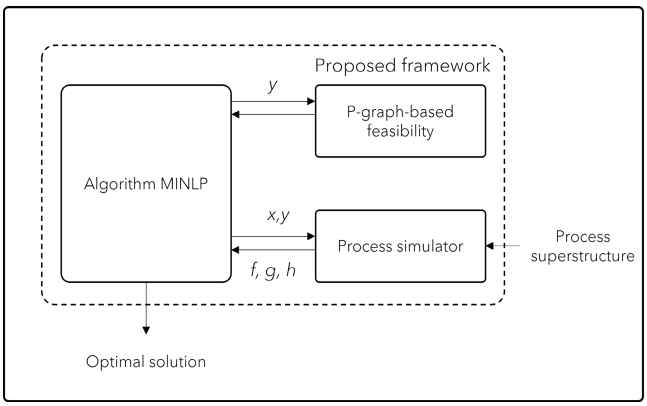

Figure 1. P-graph integrated into a process simulator-based optimization approach. Figure taken from Hillaby et al. [8].

METHODOLOGY

The P-graph-based framework proposed intends to improve the generic optimization procedure by: (1) implicitly postulating the structural constraints through the superstructure topology, (2) incorporating three procedures or algorithms, to evaluate the structural feasibility of a given process throughout the optimization. The improved methodology enables general users to leverage P-graph without requiring a direct knowledge of it. There are different technical terms that are defined in this work: permanent and conditional units; opening and closing switches; conditional stream; bypass units; structural feasibility.

Permanent and conditional units

Permanent units are defined as essential components of the structure, meaning they are necessary for the operation of the process and cannot be removed without compromising its functionality. Conditional units, on the other hand, are units that can be added or modified to optimize the structure. These units are introduced to explore different configurations or improvements in the process. If conditional units are removed, the result is the simplest or base case of the process.

Conditional, opening switch, and closing switch

The only intricacy that exists between a typical process simulator flowsheet and a P-graph process representation is the addition of a fictive operating unit or a bypass unit, herein labeled UF. Since a P-graph is a bipartite graph, a material cannot be connected to a material, hence for a structural switch that has a bypass (i.e., the material input and the material output are equal), a bypass unit has to be defined. Technically, no material transformation is done, it is only to respect the P-graph bipartite representation and its eventual algorithm that will be used. When an opening switch is identified, the directly connected streams are treated as candidate or conditional streams. This implies that, among the alternative pathways, only one stream can be activated at a time. An opening switch is a fictive unit that indicates that a decision has to be made among two candidate pathways. A close switch is also a fictive unit that requires a decision to be made from two candidate streams, coming from the same or different opening switches.

Bypass units

Bypass units are fictive units used to represent material flows that bypass any transformation processes. In other words, these units are used to model situations where a material goes directly from one point to another without undergoing any changes. However, as explained before, in a P-graph representation, the connection of one material node to another is not possible. Therefore, a bypass unit is introduced to facilitate this representation.

Structural feasibility

A feasible structure indicates that the process network adheres to the principles and logic of process synthesis, meaning that the arrangement of process units and material flows forms a valid and functional sequence. In the context of P-graph framework, a feasible structure implies that the five axioms are satisfied. For example, materials must flow into operating units, and these units must either produce outputs that feed into other units or generate final products. The maximal structure generated by the Maximal Structure Generation (MSG) method ensures that the search space is defined, from which the

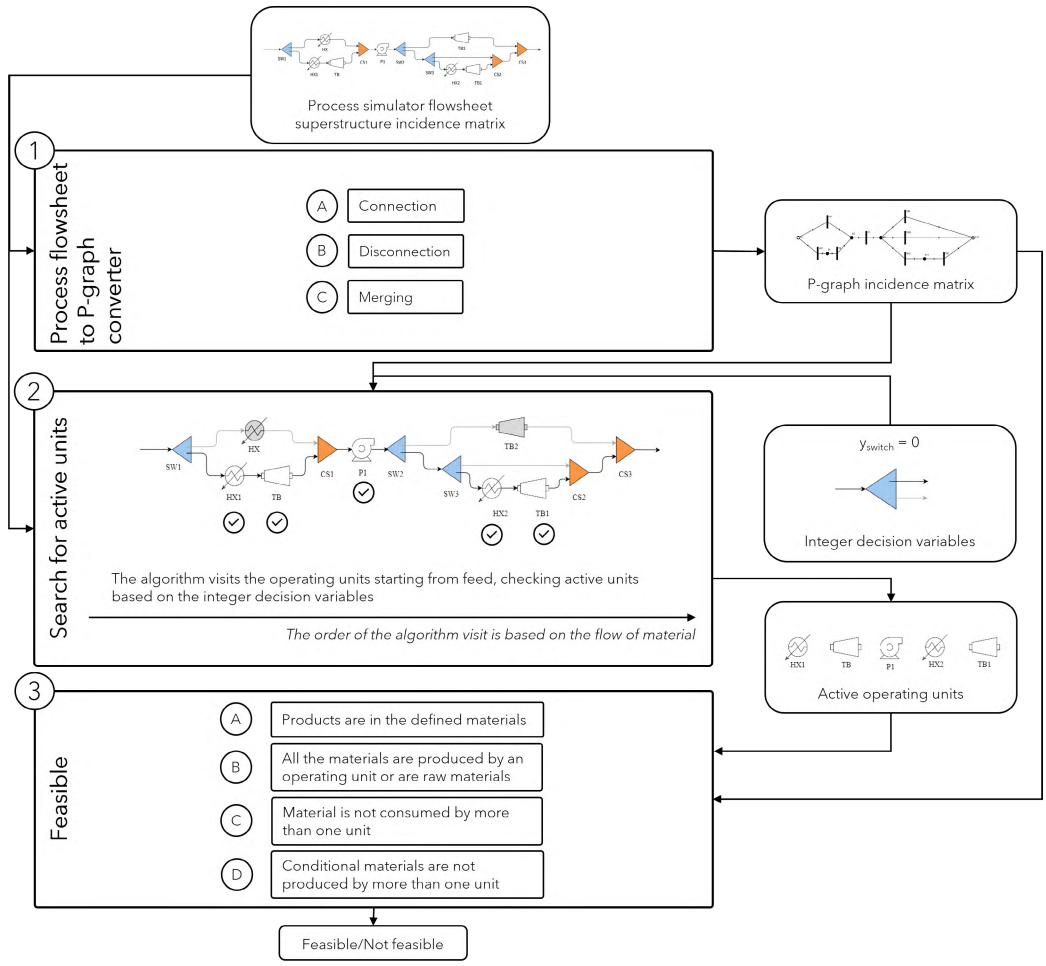

Figure 2. P-graph-based framework proposed, with three algorithms in three different blocks, labeled 1, 2, and 3. Figure adapted from Hillaby et al. [8].

optimal solutions can be identified.

MSG consists of two phases: the reduction phase and the composition phase. The algorithm used in the reduction phase is central to the methodology, with additional modifications to enhance compatibility with structural switches (i.e., opening and closing switches). The algorithm begins by extracting the information of the raw materials, products, and the superstructure (the potential arrangement of material flows and operating units). The algorithm initiates by performing the first feasibility test: verifying whether the defined products exist within the superstructure. The second test ensures that there are no intermediate products that are neither produced by any operating unit nor specified as raw materials, thus preventing the creation of unaccounted-for materials. These two tests enforce the five axioms of the P-graph framework. The third and fourth tests are added modifications, which ensure that when an opening structural switch is encountered, a decision must be made to select only one process path. Similarly, when a closing structural switch is encountered, a choice is required, meaning only one path can be accepted to continue the process sequence.

Algorithms

In general, a process optimization via superstructure starts with the user generating the superstructure itself. The problem is then defined in the simulator's optimizer interface, i.e., objective functions, constraints, and action variables. Depending on the commercial simulator used, an external optimizer could be coupled with it because the optimization tools provided by default may not be as effective or as flexible as necessary. The action variables are then assigned by the optimizer, presenting a new configuration of the flowsheet. This new structure then gets checked for feasibility. If the structure is infeasible, a penalty is applied. On the other hand, if it is feasible, the objective function is evaluated. The algorithm reassigns new action variables or exits the loop, depending on the stopping criteria such as the maximum number of evaluations, optimization time limit, or tolerance on the objective function. As the loop exits, the optimal values

are assigned and the final optimization report is obtained.

Block 1 corresponds to the Process Flowsheet to P-graph algorithm. This procedure converts a process adjacency matrix into an incidence P-graph matrix. There are three main operations that are executed: (1) connection between two units, i.e., disconnection between a unit and an open switch or a close switch (3) merging close switches' inputs.

Once the P-graph superstructure is obtained, it is possible to determine the process structure defined by the binary variables assigned by the optimizer (block 2). This is achieved by employing a graph traversal algorithm such as Depth-First Search (DFS) algorithm. The active operating units are determined and sent to the final algorithm.

Finally, the current process structure is then checked for structural feasibility (block 3). Four tests are done to ensure the structure the new process structure is feasible:

1. The products of the new structure are found in the initial defined materials.
2. All the streams are produced by an operating unit or they are defined as raw materials.
3. A stream is not consumed by more than one unit.
4. Conditional streams are not produced by more than one unit.

If the four conditions are respected, the process structure is considered structurally feasible.

By integrating the framework, i.e., the three main blocks shown in **Figure 2**, into the complete process optimization via superstructure procedure, the need to manually define structural constraints is reduced, decreasing human intervention at the start of the process simulator-based optimization. This procedure is iterative and continues until the optimization criteria are met.

A list of feasible sub-structures can be made by implementing the 2nd and 3rd algorithm blocks prior to the optimization. This precomputation reduces redundancy during the optimization stage, as repeatedly proposed structures need not be revalidated. For example, if a known feasible sub-structure is suggested multiple times by the optimizer, feasibility can be confirmed by referring to the precomputed list, instead of running the full evaluation.

AMMONIA SYNTHESIS SUPERSTRUCTURE

A case study based on Quintero-Masselski's work [9] on optimizing the Haber-Bosch (HB) process through reactors configurations is conducted to apply the developed methodology. The original ammonia process superstructure is restructured and the MINLP formulation is modified to adapt to the methodology developed. The adapted structure is shown in **Figure 3**.

Hydrogen and nitrogen from renewable sources are used to produce ammonia, which is then separated and stored. The feeds enter at 10 bar and 293.15 K and are compressed to match reaction conditions, with pressures ranging from 50 to 200 bar.

The reaction stage has different structural alternatives of adiabatic reactors, including Adiabatic Indirect Cooling Reactor (AICR) and Adiabatic Quench Cooling Reactor (AQCR). In the AICR configuration, a heat exchanger is used between the reactors to cool down the outlet stream of each reactor to the desired temperature of the following unit. This is achieved by using an external utility. In the AQCR configuration, the heat exchange is carried out by mixing the outlet stream of each reactor with a fraction of fresh process stream that is not used for reaction and bypassed from the previous reactors. There are five configurations that are implemented in the current work: (1) 1 reactor without cooling, (2) 2 adiabatic reactors in AICR configuration, (3) 3 reactors in AICR configuration, (4) 2 reactors in AQCR configuration, and (5) 3 reactors in AQCR configuration.

Iron-based (Fe) and ruthenium-based (Ru) catalysts are considered due to their industrial relevance and extensive literature availability. Process configurations (binary switches) and catalyst selection (integer variable) are modeled as discrete decision variables, as summarized in **Table 1**.

Separation by condensation is achieved by adjusting the reactor outlet mixture to its dew-point temperature, thereby separating a liquid phase composed of 99.95 % vol. ammonia with trace amounts of hydrogen and nitrogen.

These configurations are evaluated using two key performance indicators: Global Levelized Cost of Ammonia production (LCOA) and energetic efficiency. Two optimizations were performed: without the P-graph and with the P-graph methodology. The LCOA is the price of ammonia production considering every cost involved in the process, i.e., the sum of investment and operational expenditures of raw materials, electricity, and equipment.

$$\eta_E = \frac{(\dot{m}_{NH_3} LHV_{NH_3}) + \Delta\dot{W}^- + \Delta\dot{Q}^-}{(\dot{m}_{H_2} LHV_{H_2}) + \Delta\dot{W}^+ + \Delta\dot{Q}^+} \quad (1)$$

The Pré-Estime method is used to calculate the equipment cost [10]. The process energy efficiency indicator, Eq. (1), is the ratio between the useful produced energy and the supplied energy to the process: where η_E is the overall energy efficiency, $\dot{m}$ is the mass flowrate in kg/s, LHV is the Low Heating Value in kJ/kg, $\Delta\dot{W}$ is the overall balance of mechanical power in kW, and $\Delta\dot{Q}$ is the balance of thermal energy in kW. The energy terms use superscripts to indicate production (+) and consumption (-) [11].

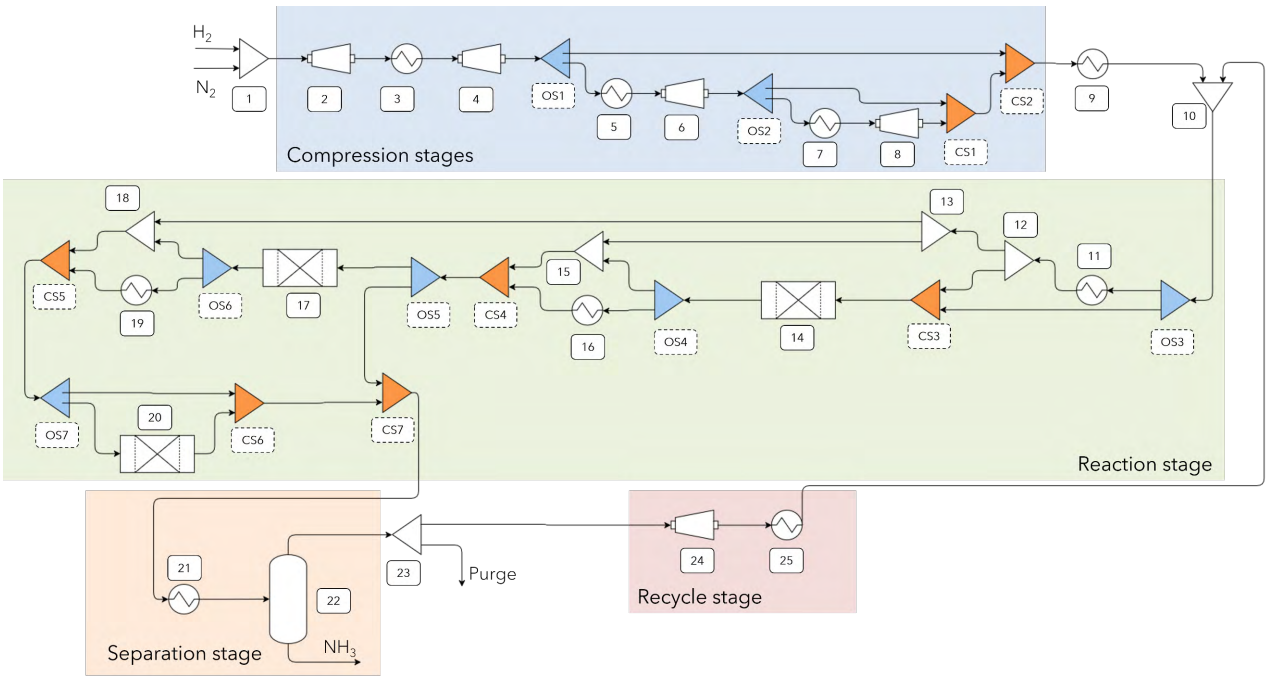

Figure 3. Superstructure representation for the evaluation of the adiabatic reactor configuration. The compression stage includes the selection of the number of compressors. The reaction stage has five structural variables to determine the number of reactors and the type of colling. Module nomenclature corresponds to: Mixers (1, 10, 15, 18), Splitters (12, 13, 23), Compressors (2, 4, 6, 8, 24), Heat exchangers (3, 5, 7, 9, 11, 16, 19, 21, 25), Adiabatic reactors (14, 17, 20), Separation drum (22), Opening switches (OS1, OS2, OS3, OS4, OS5, OS6, OS7), and Closing switches (CS1, CS2, CS3, CS4, CS5, CS6, CS7).

There are seven inequality process constraints that are defined in this problem, see **Table 2**. The maximum LCOA is fixed to discard solutions with high costs of production whereas the minimum H2 conversion is defined to avoid low H1 conversions per pass and important hydrogen losses on the purge stream. A maximal reactor temperature is defined to increase the catalyst lifetime.

Table 1. Discrete action variables defined in this study.

Variable	Details
SW1	Enables the 3rd compressor
SW2	Enables the 4th compressor
SW3	Enables the quench setting
SW4	Enables either the 1st quench or the 1st external cooling heat exchanger
SW5	Enables the 2nd reactor
SW6	Enables the 2nd quench or the 2nd external cooling heat exchanger
SW7	Enables the 3rd reactor
Catalyst	0: Fe-based; 1: Ru-based

Table 2. Process constraints and their values, with the catalysts' corresponding pressure intervals being the disjunctive constraints.

Constraint	Value	Units
Maximum LCOA	850	€/t
Minimum H_2 global conversion	90	%
Maximal temperature in 1st reactor	800	K
Maximal temperature in 2nd reactor	800	K
Maximal temperature in 3rd reactor	800	K
Fe catalyst pressure range	150 – 250	bar
Ru catalyst pressure range	50-100	bar

The optimizer handles the constraint violations in two layers: internally and externally. The first layer is an external layer, in which the structural feasibility is assessed. If the structure is not feasible, i.e., the P-graph axioms are violated, a death penalty (objective functions, $f(x) = 10^{300}$; constraints, $g(x) = 0$) is applied.

The second layer is the internal layer, in which the Oracle Penalty Method [12] is applied. If the candidate solutions are structurally feasible, the method ranks the solutions using both objective value and constraint residual.

Table 3. Continuous action variables defined in the optimization, together with the stage at which they are defined, the associated process units, and their bounds.

Variable	Stage	Process unit	Bounds	Units
Pressure	Compression	2, 3, 6, 8	50 – 200	bar
Inlet temperature at 1st reactor	Reaction	9, 25	636 – 736	K
Inlet temperature at 2nd reactor	Reaction	16	636 – 736	K
Inlet temperature at 3rd reactor	Reaction	19	636 – 736	K
Length of 1st reactor	Reaction	14	0.5 – 7.0	m
Length of 1st reactor	Reaction	17	0.5 – 7.0	m
Length of 1st reactor	Reaction	20	0.5 – 7.0	m
Separation temperature	Separation	21	253 – 275	K
Split ratio at first quench splitter	Reaction	12	0.4 – 0.8	-
Split ratio at second quench splitter	Reaction	13	0.4 – 0.8	-

Implementation

The optimizations were performed on ProSimPlus 3.7.9 coupled with an external optimizer (MIDACO [13]) on an Intel(R) Xeon(R) w5-3433, 1.99 GHz, 256 GB RAM, Windows 10 64-bit operating system. The optimizations stopped when 15000 consecutive evaluations no longer improved the current solution. The optimization is done with two seeds, improving the exploration of the search space and reduce sensitivity to stochastic convergence.

RESULTS

Methodology comparison

Table 4. A comparison of the optimal point intervals between Quintero-Masselski's study [9] and this work.

Catalyst	Objective functions	Quintero-Masselski	This work
Fe	LCOA (€/t)	752 – 816	751 – 753
	Efficiency (%)	54.4 – 56.4	53.9 – 54.8
Ru	LCOA (€/t)	754 – 804	752 – 804
	Efficiency (%)	55.3 – 58.1	54.8 – 58.2

Based on the results of the optimizations performed, as shown in **Table 4**, it can be said that the modified superstructure is equivalent to the one used in Quintero-Masselski's work [9]. Hence, only the comparison between optimizations done with and without the modified P-graph methodology is done hereinafter.

The runs without the P-graph procedures and with the proposed methodology (P-graph based) took 47.7 and 36.6 hours respectively, resulting in a 22 % reduction of optimization time. The results are further detailed in **Table 5**.

90 % of the solutions proposed by the optimizer with the P-graph-based methodology implemented are feasible structurally, representing most of the evaluations performed. This is because a death penalty is applied when the solution proposed is deemed infeasible, steering the optimizer toward a feasible region. The percentage of simulator-feasible solutions among the structurally feasible points, represented by $\frac{n_{simF}}{n_{structF}}$, is similar for both cases. This suggests that the structural feasibility had more impact on the simulator convergence than the combination of continuous action variables for a given candidate design. Finally, the methodology shows a clear advantage when comparing the fraction of design-feasible over simulator-feasible solutions, at 0.94 over 0.68.

Furthermore, based on **Figure 4**, there are 45 global Pareto points, of which 35 are with P-graph and 10 are without P-graph. This shows that more non-dominated Pareto points are attained with the modified methodology.

Table 5. Comparison of the methodology with different fractions to analyze the efficacy of the methodology, including structural-feasible over total data points (3rd column), simulation-feasible over structural-feasible (4th column), and design-feasible over simulation-feasible (5th column).

Case	n_D	$\frac{n_{structF}}{n_D}$	$\frac{n_{simF}}{n_{structF}}$	$\frac{n_{designF}}{n_{simF}}$
Without P-graph	27397	0.27	0.96	0.68
With P-graph	18053	0.90	0.94	0.94

Process optimization discussions

There are several interpretations that can be considered as an extended version to that of Quintero-Masselski and colleagues' work. Only the results from the P-graph methodology are studied in this section.

Figure 5 shows that the two-compressor configuration dominates the lower efficiency and lower LCOA zone

whereas the three-compressor and a four-compressor configurations zone are mostly overlapped. This indicated that the main performance gain in energy efficiency comes from increasing the compressor number to at least three stages, while a fourth stage provides limited improvement. But increasing the number of compressors will increase the LCOA.

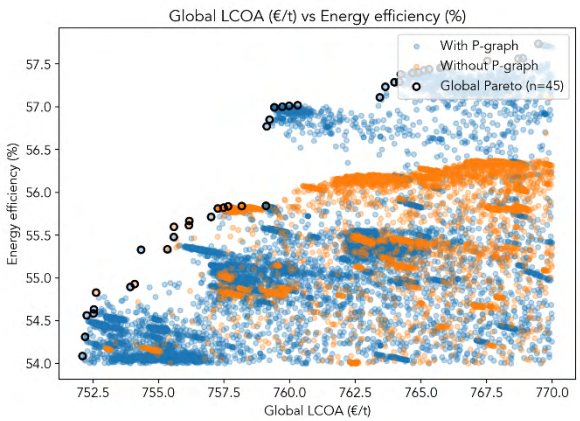

Figure 4. Energy efficiency vs. process LCOA. Blue points correspond to solutions obtained with the P-graph methodology, orange points to solutions obtained without it. Pareto-optimal points are shown as hollow black markers.

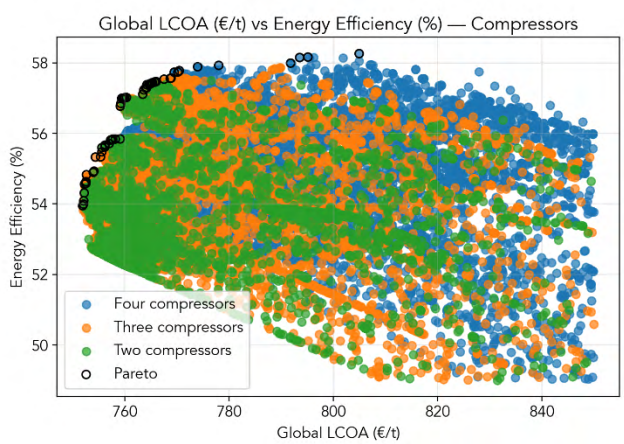

Figure 5. LCOA versus process energy efficiency, colored by the number of compressors (2, 3, or 4). Global Pareto-optimal solutions are indicated by hollow black markers.

Figure 6 exhibits that the configurations with 1, 2 or 3 AICR dominates the solution space as 2 and 3 AQCR are not as visible in the figure. When comparing the AICR configurations, the systems with 1 reactor have clearly the lower efficiency and the lower LCOA values. When there is only one stage of reaction, this means lower conversion per pass, requiring massive recycle load and increasing recycle compressor duty.

Figure 7 compares the solutions based on the type of catalyst used. There are 36643 and 44363 solution points for the Fe-based and the Ru-based catalysts respectively. There are 8 solutions on the Pareto front with Fe-catalyst ranging from 752 to 753 €/t in LCOA and from 53 % to 55 % in energy efficiency while 47 Ru-based solutions on the Pareto front ranging from 752 to 805 €/t in LCOA and from 55 % to 58 %. The increased value of energy efficiency using Ru can be explained by a lower pressure range which lead to a decrease in the compression power.

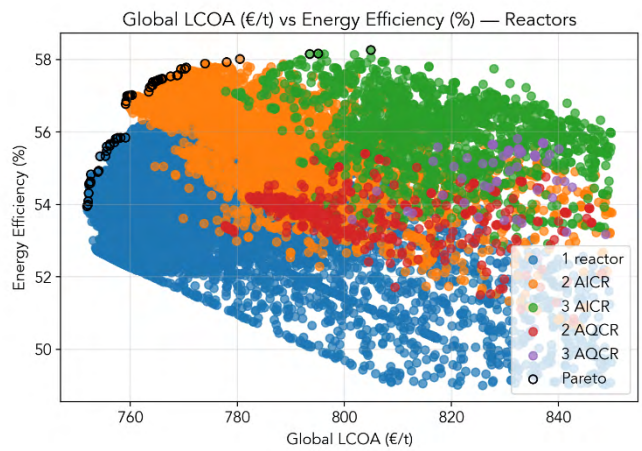

Figure 6. LCOA versus process energy efficiency, colored by the number (1, 2, or 3) and configuration (with or without cold shot) of reactors. Global Pareto-optimal solutions are indicated by hollow black markers.

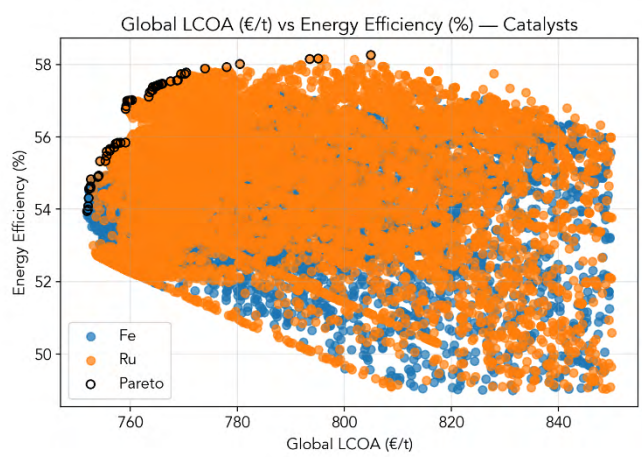

Figure 7. LCOA versus process energy efficiency, colored by the type of catalysts (Fe-based or Ru-based). Global Pareto-optimal solutions are indicated by hollow black markers.

CONCLUSION

This work demonstrates a novel methodology of process simulator-based optimization via superstructure with an application for ammonia synthesis through the Haber Bosch process.

Binary opening and closing switches as well as three algorithms are introduced in this study. The binary switches are used to propose different structural alternatives in a process while the three algorithms are used to determine the feasibility of sub-structures in a superstructure.

The approach proposed not only reduced the time required to optimize a process superstructure, especially when comparing the total evaluated solutions that are design-wise feasible but also reduced the need for users to define the structural constraints.

Future work is going to be more focused on benchmarking additional case studies using the proposed methodology, as well as conducting a more holistic optimization in which a Key Performance Indicator (KPI) of the Life Cycle Assessment (LCA) is included in the objective functions.

ACKNOWLEDGEMENTS

The authors address their acknowledgements to Jean Pimentel Losada of the Budapest University of Technology and Economics, for his valuable insights concerning the P-graph framework.

AUTHOR IDENTIFIERS

Author ORCIDs:
Hillaby D.: 0009-0007-0721-3972
Piña Martinez A.: 0000-0003-0691-0075
Portha J.-F.: 0000-0001-8698-8635

REFERENCES

1. UNFCCC, Summary of Global Climate Action at COP 28, 2023. https://unfccc.int/documents/636485
2. Montastruc L, Belletante S, Pagot A, Negny S, Raynal L. From conceptual design to process design optimization: a review on flowsheet synthesis. Oil Gas Sci. Technol. – Rev. IFP Energies nouvelles 74:80 (2019). https://doi.org/10.2516/ogst/2019048
3. Chen Q, Grossmann IE. Recent developments and challenges in optimization-based process synthesis. Annu. Rev. Chem. Biomol. Eng. 8:249-283 (2017). https://doi.org/10.1146/annurev-chembioeng-080615-033546
4. Mencarelli L, Chen Q, Pagot A, Grossmann IE. A review on superstructure optimization approaches in process system engineering. Computers & Chemical Engineering 136:106808 (2020). https://doi.org/10.1016/j.compchemeng.2020.106808
5. Boonstra S, van der Blom K, Hofmeyer H, Emmerich MTM, van Schijndel J, de Wilde P. Toolbox for super-structured and super-structure free multi-disciplinary building spatial design optimisation. Advanced Engineering Informatics 36:86-100 (2018). https://doi.org/10.1016/j.aei.2018.01.003
6. Corbetta M, Grossmann IE, Manenti F. Process simulator-based optimization of biorefinery downstream processes under the generalized disjunctive programming framework. Computers & Chemical Engineering 88:73-85 (2016). https://doi.org/10.1016/j.compchemeng.2016.02.009
7. Friedler F, Orosz Á, Pimentel Losada J. P-graphs for process systems engineering. Springer International Publishing (2022). https://doi.org/10.1007/978-3-030-92216-0
8. Hillaby D, Piña-Martinez A, Falk L, Portha JF. Structural constraint reduction in process simulator-based optimisation: leveraging the p-graph framework. Computers & Chemical Engineering 201:109254 (2025). https://doi.org/10.1016/j.compchemeng.2025.109254
9. C.S. Quintero-Masselki, Process design by superstructure optimization for the industry of the future, Université de Lorraine, 2022.
10. A. Chauvel, G. Fournier, C. Raimbault, A. Pigeyre, Manual of Process Economic Evaluation, Éditions Technip, Paris (2001).
11. Peduzzi E, Boissonnet G, Haarlemmer G, Maréchal F. Thermo-economic analysis and multi-objective optimisation of lignocellulosic biomass conversion to fischer–tropsch fuels. Sustainable Energy Fuels 2:1069-1084 (2018). https://doi.org/10.1039/c7se00468k
12. Schlüter M, Gerdts M. The oracle penalty method. J Glob Optim 47:293-325 (2009). https://doi.org/10.1007/s10898-009-9477-0
13. Schlüter M, Egea JA, Banga JR. Extended ant colony optimization for non-convex mixed integer nonlinear programming. Computers & Operations Research 36:2217-2229 (2009). https://doi.org/10.1016/j.cor.2008.08.015

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article: Revised Version
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

High Performance Heat Pumps Using Tailored Refrigerants

Finlay M. Sandham, Andrew Muumbo, Kenneth Mathew, Sarthak Sinha and Smitha Gopinath*

University of Sheffield, School of Chemical, Materials, & Biological Engineering, Sheffield, United Kingdom
* Corresponding Author: s.gopinath@sheffield.ac.uk

ABSTRACT

Heat Pumps (HPs) can play a vital role in the decarbonization of heating in industry. The performance of a HP strongly depends on the refrigerant, the working fluid within the HP. In order to maximize HP performance, systematic selection of the refrigerant is key. Refrigerant choice affects the very feasibility of employing a HP to deliver heating to a process. A flexible and robust method is required to select refrigerants that are the best fit for a given heating application. A computer-aided molecular & process design (CAMPD) method is developed to design the optimal refrigerant that is tailored to process needs. The method is applied to three case studies across which the HP performance objectives and constraints, and heat source and heat sink temperatures are varied. In addition, the design of refrigerants with low (<150) global warming potentials and zero ozone depletion potentials is investigated. For all applications across all case studies, the CAMPD approach successfully identifies high-performance refrigerants including those that are known refrigerants, other known fluids and novel molecules.

Keywords: optimization, molecular design, decarbonization, process design

INTRODUCTION

Heat consumption in industries worldwide is around 110 EJ per year [1]. Over 90% of this heat is generated by the combustion of fossil fuels, resulting in 21% of global carbon dioxide emissions [2]. Heat pumps (HPs) that are powered by electrical energy and offer high energy efficiencies are a promising alternative heating solution. HPs harness heat from a low-temperature heat source and upgrade it to an increased, useful temperature by utilizing (electro-)mechanical work. HPs are able to utilize both ambient heat and waste heat due to which the work input into the HP is only a fraction of the heat they supply. HPs can be driven by low-emissions electricity, yielding low-carbon heating.

The working fluid within the HP is known as the refrigerant. HPs based on the subcritical vapour compression cycle are the focus of this work. In the evaporator of the HP, the low-pressure refrigerant takes up thermal energy from a (low temperature) heat source, at temperature T_{source}, and vaporizes. The gaseous refrigerant is compressed to a higher pressure and temperature in the compressor. In the condenser, the high-temperature refrigerant releases thermal energy to the heat sink, which is at temperature T_{sink}, and condenses. To complete the cycle, the refrigerant passes through the expansion valve (sometimes called a throttling valve) that causes a drop in pressure and temperature. This process is illustrated in figure 1.

The application of HPs in the process industries necessitates innovation in the design of HPs. HPs may be used to deliver heating to a variety of unit operations such as distillation columns, dryers, and reactors whose temperatures are below 200 °C [3]. At the same time, several (waste) heat sources whose temperature vary between 30 °C and 200 °C are available in the process industries [3]. Given the large number of heat sources and sinks and the diversity of their temperatures and heat duties, HPs and their refrigerants need to be tailored to each application to maximize performance. The refrigerant in the HP must satisfy manifold requirements: (i) the refrigerant must be able to undergo phase transitions in both the evaporator and condenser; ii) the boiling point of the refrigerant in the evaporator must be below T_{source} at pressures above 1 atm (to prevent vacuum operations); iii) the critical temperature of the

refrigerant must be higher than T_{sink} to allow the refrigerant to deliver heat by condensation; and iv) the refrigerant must be compatible with compressor seals and lubricant.

Apart from the performance of the refrigerant in the HP, there are several environmental and safety-related implications to the choice of the refrigerant. Refrigerants containing chlorine can have high ozone depletion potentials (ODPs) and have been phased out by the Montreal protocol [4]. They have been replaced by fluorinated refrigerants, many but not all of which have very high global warming potentials (GWPs), exceeding the regulatory limit of 150 [5]. Several fluorinated molecules are either expected to be phased out or are already banned due to their high GWPs, toxic degradation products, and pervasiveness in the environment [5]. Natural refrigerants, which include molecules such as alkanes, alkenes, and ammonia, are a promising alternative class of refrigerants. However, several of these are highly flammable, restricting their use across settings. Water and carbon dioxide are non-flammable natural refrigerants that offer high safety and good performance for a limited range of temperatures.

Given the shortcomings of known refrigerants, the discovery of new refrigerants is essential. Further, it is crucial to systematically select the best-fit refrigerant considering both process-specific needs (heat duties, temperatures) as well as the availability of waste heat streams (heat capacities and temperatures). In addition, the refrigerant must comply with environmental and safety-related regulations, e.g., the GWP of the refrigerant must be lower than the legal limit.

Most previous computational techniques for refrigerant selection use a screening process, where all molecules within a chemical database are examined to find those with desirable properties or performance. Examples of such studies include [6], [7], [8].

Computer-aided molecular & process design (CAMPD) is an alternative that can be used to identify best-fit refrigerants. In this framework, an optimization problem is formulated to find the optimal refrigerant and process conditions of the HP. [9] Applied CAMPD to find refrigerant candidates for a domestic HP using detailed process models including a sophisticated compressor model.

In this work we simultaneously find the refrigerant and the process variables that satisfy performance constraints and maximize the thermodynamic performance of a subcritical vapor-compression industrial HP, using a CAMPD approach. The best possible pure-component refrigerant, be it a known refrigerant, other known fluid, or a novel molecule, is identified from within a large molecular design space. To the best of our knowledge, this is the first CAMPD study on the basis of rigorous thermodynamic models that is applied to high-temperature industrial HPs and that also considers the environmental properties of the refrigerant. In this work, we systematically investigate the use of CAMPD to design HPs by i) varying heat source and heat sink temperatures; ii) varying HP performance objectives; and iii) restricting refrigerants to be low-GWP and zero-ODP. We also benchmark the performance of designed refrigerants against known refrigerants to study the effectiveness of this CAMPD approach.

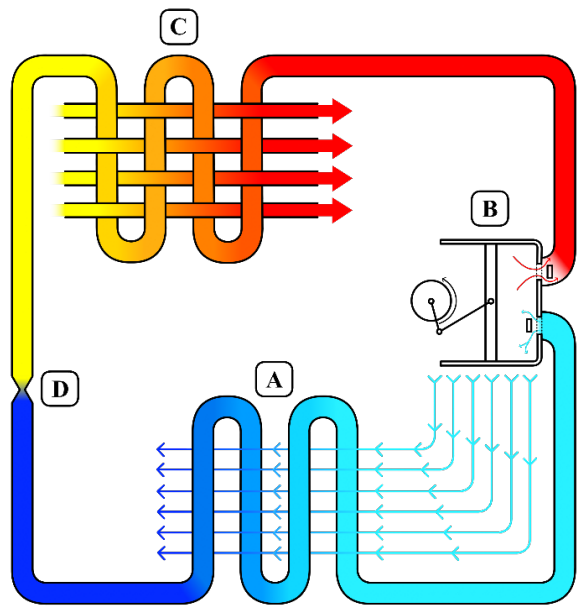

Figure 1. A diagram of a vapor-compression HP cycle. A: Evaporator; B: Compressor; C: Condenser; D: Expansion valve

CAMPD OF HP

The HP cycle is modelled in gPROMS Process 2024.1.0 using MESH (Mass, Equilibrium, Summation, and Enthalpy) models. These models do not consider sizing or costing of HP equipment. Several assumptions are made in modelling the HP. The heat source and heat sink are isothermal reservoirs at temperatures T_{source} and T_{sink}, respectively. The evaporator and condenser are isobaric at pressures P_{evap} and P_{cond}, respectively. The compressor has an isentropic efficiency of 70%, and the motor and drivetrain have a combined efficiency of 95%. The expansion valve is isenthalpic. There are no pressure drops or heat losses in the connections between unit operations.

The performance of the HP cycle is quantified using two metrics: coefficient of performance (COP) and volumetric heating capacity (VHC). The COP, a measure of energy efficiency, is the ratio of heat delivered to the heat sink (Q_{out}) to compressor power consumption (W_{in}). The VHC is the energy delivered to the heat sink per unit volume of refrigerant at the compressor inlet ($V_{comp.inlet}$),

i.e., $VHC = \frac{Q_{out}}{V_{comp.inlet}}$. VHC is indicative of the sizing of the HP. The temperature lift ΔT_{lift} of the HP is given by $\Delta T_{lift} = T_{sink} - T_{source}$.

Process-Level Degrees of Freedom.

In the optimization of this HP model, the following process-level degrees of freedom are considered as continuous decision variables:

The condenser pressure P_{cond}

The condenser subcool ΔT_{sub}

The evaporator pressure P_{evap}

The evaporator superheat ΔT_{super}

In addition, the chemical composition of the refrigerant is a decision variable. The refrigerant is represented using a vector of integers $\boldsymbol{n}$, where each element of the vector n_i represents the number of occurrences of functional/atomic group i in the refrigerant. For example, to represent butanol, the vector elements corresponding to -OH, -CH_2-, and -CH_3 would be set to 1, 3, and 1, respectively, and all other elements of $\boldsymbol{n}$ would be set to 0. 20 chemical groups are allowed in the refrigerant, as outlined in Table 1. Up to 10 occurrences of each group are allowed in the molecule, other than >N- which can take a maximum value of 1 to avoid complex branching structures. The group -cO- refers to an oxygen atom that is attached to two -CH_2- groups, and the -eO- group is an (end) ether group in which the oxygen atom is connected to a -CH_3 group.

Table 1: The groups considered in the CAMPD problem.

Chemical family	Groups
Hydrocarbons	-CH_2-, -CH_3,
Olefins	=CH_2, =CH-
Alcohols, carboxylic acids ethers, esters, and ketones	-OH, -eO-, -cO-, >C=O, -COO, -COOH
Chlorocarbons	-$CHCl_2$, -CH_2Cl
Fluorocarbons	-CF_3 -CHF_2, -CH_2F, -CF_2-, -CHF-
Amines	-NH_2, -NH-, >N-

Group contribution methods (GCMs)

Group contribution methods (GCMs) are used to estimate physicochemical properties of the refrigerant, which is described by its chemical composition vector $\boldsymbol{n}$. The group contribution equation of state (EoS) SAFT-γ Mie [10] is used to predict the thermodynamic properties of the refrigerant molecule both as a function of $\boldsymbol{n}$ and its state (temperature, pressure) in the process. SAFT-γ Mie was chosen for its accuracy in predicting both calorific properties and vapour-liquid equilibria, its ability to describe a wide range of refrigerants and fluids, and its proven use in similar CAMPD problems [11]. The SAFT-γ Mie EoS is included in the inbuilt gProperties tool in gPROMS Process, which is used to compute thermodynamic properties. The parameters for the fluorinated groups are taken from Martínez [12] and all others are taken from the gProperties databank.

While the SAFT-γ Mie EoS can predict all thermodynamic properties of fluids, the accuracy of critical point predictions is typically low. Thus, a first-order semi-empirical group contribution model was chosen to model the critical point. [13] was chosen for its higher accuracy in comparison to two other methods [14], [15] for a small test of molecules.

The work of [16] is used for GWP prediction. Again, only first-order group contributions are considered. For groups with no computed GWP contribution, such as -CH_2-, the contribution is assumed to be negligible and is set to 0. However, this approach did not account for the significantly reduced GWP of molecules with double bonds - the atmospheric lifespan of these molecules is greatly reduced which causes them to have a lower GWP. Therefore, we parameterized first-order contributions for =CH- and =CH_2 groups by regressing to GWP data from [17]. The contribution calculated for =CH- is -1.618305, and for =CH_2 is -0.502935.

Objectives & Constraints

The CAMPD problem is formulated to maximize one of two objective functions: COP, and VHC. The design problem is subject to process-related constraints to ensure feasible operation of the subcritical HP cycle. Molecular feasibility constraints are added to ensure that the octet rule is satisfied for all atoms in the molecule. Several new molecular feasibility constraints are introduced to prevent molecules that are known to be unstable stable such as geminal diols, geminal amines, α-haloalcohols, peroxides and peracids. Further, for molecules that exhibit keto-enol and amine-imine tautomerism, the enol and amine forms that are typically unstable in aliphatic molecules are excluded. The process constraints and molecular feasibility constraints are listed in the supplementary information (SI).

Experiments

The full MINLP is implemented in gPROMS Process 2024.1.0 using gProperties v2.8 and solved using the MINLPOA solver which is based on the outer approximation algorithm [25] with equality relaxation and augmented penalty. This algorithm works by sequentially solving a primal continuous nonlinear program (NLP) and a master mixed-integer linear program (MILP). The NLP is solved using the NLPSQP algorithm.

Only local optima of this non-convex problem are obtained by solving the MINLP. To obtain high-quality

solutions, integer cuts are used to re-solve each design problem while excluding previous solutions. The integer cut is as is shown in Equation 1 ([18], [19]), where $\boldsymbol{p}$ is a vector of the first 20 prime numbers, M is a sufficiently large number, 10,000 for this problem, $\boldsymbol{n}^{cut}$ is the solution to be excluded, y is a binary decision variable, and ϵ is a tolerance, which is set to 2 in this case.

$$My + \epsilon < \boldsymbol{p}^T \boldsymbol{n}^{cut} - \boldsymbol{p}^T \boldsymbol{n} < M(1-y) - \epsilon \quad (1)$$

Computational experiments are performed to observe how the solution changes across different source-sink temperatures, design objectives, and constraints. The experiments are run from source temperatures of 340 K to 420 K and sink temperatures of 360 K to 440 K at temperature intervals of 20 K. For each source-sink temperature pair, the MINLP optimization is solved 4 times using 3 successive integer cuts to generate 4 different candidate solutions. The best-found solution found over the 4 runs is then reported for each temperature pair. The initial guess for all MINLP solves is fixed and given in the SI. In order to facilitate these experiments a computer program was developed in C++ to call gPROMS optimizations through the gO:RUN feature at the specified source-sink temperatures and to automatically apply the integer cuts.

Table 2: Summary of experiments

Case	Objective	Additional Restrictions
1	VHC	None
2	COP	None
3	COP	Constraint on GWP (<150 as per EU F-gas regulations), zero ODP enforced by removal of chlorine containing groups. Restricted design space through removal of nitrogenated and fluorinated groups.

The results are benchmarked against four existing refrigerants that are already being used in HPs [20]. The benchmark refrigerants are ammonia (R-717), pentane (R-601), 1,1,1,3,3-pentafluoropropane (R-245fa), and water (R-718). These were chosen as benchmarks in this study because they collectively cover a wide range of feasible operating temperatures and chemical families. For each source-sink temperature pair, the refrigerant was fixed to each benchmark fluid and the COP and the VHC were maximized.

RESULTS AND DISCUSSION

Benchmarks

In Table 3, the best of the four benchmark fluids with the highest value of COP is shown for each source-sink temperature pair. Similarly, the best benchmark fluid with VHC as the design objective has also been identified (results shown in SI). The name of the highest-COP benchmark for each source-sink pair is given. Below the name is the optimized value of the objective. In Table 3, pentane dominates at lower source temperatures, being the top benchmark for 6 out of 15 source-sink temperature pairs. At higher temperatures, water emerges as the best benchmark fluid.

Table 3: Best of four benchmark refrigerants for COP at varying source and sink temperatures.

COP		T_{source} (K)				
		340	360	380	400	420
T_{sink} (K)	360	Ammonia **6.99**	–	–	–	–
	380	Pentane **3.35**	R-245fa **5.88**	–	–	–
	400	Pentane **2.19**	Pentane **3.15**	Water **8.56**	–	–
	420	Pentane **1.53**	Pentane **2.02**	Water **5.28**	Water **9.03**	–
	440	No solutions **N/A**	Pentane **1.34**	Water **3.87**	Water **5.55**	Water **9.47**

Case 1: Maximization of VHC

In Table 4, for each source-sink temperature pair, the best-found refrigerant molecule for maximizing the VHC, after applying 3 integer cuts, is shown. The value of the VHC is given in bold in MJ m^{-3}. In brackets next to this, is the percentage change in VHC compared to the benchmark fluid that achieves the highest VHC for a given source-sink temperature. At T_{source} = 360 K, T_{sink} = 380 K, the best-found refrigerant was butadiene with a VHC of 8.03; however, butadiene polymerizes easily, therefore the second-best solution is displayed instead and highlighted with an asterisk.

Typically, for any fixed source temperature, the VHC reduces as the ΔT_{lift} increases. We note that this trend is not followed in Table 4 at T_{source} = 360 K as the second-best solution is reported for ΔT_{lift} = 20 K. The maximum VHC observed was 15.14 at T_{source} = 420 K, T_{sink} = 440 K , the smallest ΔT_{lift} value of 20 K. The minimum was 1.88 at T_{source} = 340 K, T_{sink} = 440 K which is at the largest ΔT_{lift} value of 100 K.

The results indicate that the refrigerant that maximizes the VHC varies with the source-sink temperatures. 7 distinct optimal refrigerants were identified. In 1 out of the 15 tested instances, the VHC value obtained by CAMPD was 57.8% lower than that of ammonia, the best benchmark fluid for that instance. It is worth noting that ammonia was not included in the design

Table 4: Maximization of VHC for case 1 at varying source and sink temperatures.

VHC (MJ m^{-3})		T_{source} (K)				
		340	360	380	400	420
T_{sink} (K)	360	1,1-Difluoroethane **9.78** (-57.8%)	–	–	–	–
	380	1,1,2-Trifluoroethane **5.57** (+104.8%)	1,2-Difluoropropane* **3.70** (-24.9%)	–	–	–
	400	3,3-Difluoropropene **3.98** (+323.4%)	3,3-Difluoropropene **6.26** (+247.8%)	3,3-Difluoropropene **9.59** (+185.4%)	–	–
	420	1,1-Difluoroethanamine **2.84** (+323.9%)	1-Chloro-2,2-difluoroethane **3.72** (+188.4%)	Methanol **6.39** (+163.0%)	Methanol **10.55** (+133.9%)	–
	440	1-Chloro-2,2-difluoroethane **1.88** (N/A)	Methanol **3.40** (+286.4%)	Methanol **5.89** (+145.4%)	Methanol **9.66** (+137.9%)	Methanol **15.14** (+132.6%)

space of the CAMPD problem and, thus, could not have been found as a solution. For 12 out of 15 instances, the use of CAMPD resulted in large improvements of VHC over the best of the benchmark fluids. 1,1-Difluoroethanamine yielded a 323.9% increase in VHC when compared with pentane (R-601), the best benchmark fluid at $T_{source} = 340$ K and $\Delta T_{lift} = 80$ K.

All of the identified refrigerants were novel. One olefin, 3,3-difluoropropene, was found in this case study. 3,3-Difluoropropene belongs to the hydrofluoroolefin family of refrigerants, several of which have been examined in HPs, due to their high COPs, low GWPs and low flammability. One novel oxygenated refrigerant, methanol, was identified; methanol is found to be an excellent candidate in other case studies as well. The novel hydrochlorofluorocarbon (HCFC) 1-chloro-2,2-difluoroethane was also identified; HCFCs have been widely used as refrigerants.

The average pressure ratio across all 15 instances was 6.45 with a maximum pressure ratio of 17.72. Most compressor types are able to achieve pressure ratios up to 10, and screw compressors can achieve pressure ratios up to 20 [20]. ΔT_{super} was negligible for all instances; however, ΔT_{sub} was consistently high, reaching a maximum of 40.37 K at $T_{source} = 380$ K, $T_{sink} = 420$ K. These results are all expected because condenser pressures are increased in order to maximize the energy that is delivered by unit volume of refrigerant. We planned to add 3 integer cuts to generate designs for each of the 15 source-sink temperature pairs shown in Table 4, resulting in a total of 60 optimizations. However, for 2 of the solutions in Table 4, optimizations failed to converge after the addition of the first and second integer cuts respectively, and thus further cuts were not added. The 57 optimizations took an average of 702.2 s, with a maximum of 7138.1 s and minimum of 13.5 s.

Case 2: Maximization of COP

In Table 5, the refrigerant that maximizes the COP for each source-sink temperature pair is shown. This refrigerant corresponds to the best-found solution after 3 integer cuts. The value of the COP achieved by the refrigerant is shown in bold. In brackets, the percentage change of this COP value with respect to the COP of the best benchmark fluid for that source-sink temperature pair is shown.

The COP values in Table 5 illustrate the dependence of the performance of the HP on source-sink temperatures and temperature lifts. For each T_{source}, the COP decreases as T_{sink} increases, i.e. as ΔT_{lift} increases. For any fixed value of ΔT_{lift}, higher COPs are obtained when T_{sink} is higher. The maximum obtained COP is 9.46 by the refrigerant 2-fluoroethanol at $T_{source} = 420$ K, $T_{sink} = 440$ K. Fluorinated alcohols are not common refrigerants.

The results demonstrate that the choice of objective function strongly impacts designs, including the choice of optimal refrigerant. For 14 out of the 15 source-sink pairs that are tested, the refrigerant that maximizes the COP (Table 5) is distinct from the refrigerant that maximizes the VHC (Table 4).

The results indicate the refrigerant that maximizes the COP varies with source-sink temperatures. For the 15 total source-sink pairs that are tested, 5 distinct optimal

Table 5: Maximization of COP for case 2 at varying source and sink temperatures.

COP		T_{source} (K)				
		340	360	380	400	420
T_{sink} (K)	360	3-Fluoropropene **7.81** (+11.7%)	–	–	–	–
	380	Methanol **4.81** (+43.6%)	Methanol **8.16** (+38.8%)	–	–	–
	400	Methanol **3.52** (+60.7%)	3-Fluoropropene **4.98** (+58.1%)	Fluoromethanamine **8.73** (+2.0%)	–	–
	420	3-Fluoropropene **2.72** (+77.8%)	Ethanol **3.64** (+80.2%)	Fluoromethanamine **5.39** (+2.1%)	Fluoromethanamine **9.11** (+0.9%)	–
	440	3-Fluoropropene **2.22** (N/A)	Methanol **2.84** (+111.9%)	Ethanol **3.70** (-4.4%)	Fluoromethanamine **5.59** (+0.7%)	2-Fluoroethanol **9.46** (-0.1%)

refrigerants were identified. 3-Fluoropropene (also known as allyl fluoride) was a common solution in case 2, appearing as the best-found solution for 4 out of 15 temperature pairs. This is a known compound but has no documented use as a HP refrigerant. A novel fluorinated amine molecule, fluoromethanamine, was identified in this case study.

In case 2, improvements to the COP have been observed for 12 out of 15 source-sink pairs when compared with the best benchmark, illustrating the benefits of selecting a refrigerant tailored to the application. At $T_{source} = 360$ K, $T_{sink} = 440$ K, a COP improvement of 111.9% was obtained by choosing methanol as the refrigerant. Screening studies [6] and [8] which examined a small set of molecules have also found methanol to be a promising refrigerant. It is highly promising that the CAMPD framework which is using optimization to implicitly search through a large design space with $2 * 11^{19}$ possible combinations of groups has also been able to identify this high-performance fluid.

In case 2, the pressure ratio increased as ΔT_{lift} increased. The maximum pressure ratio observed was 14.5 at the maximum ΔT_{lift} of 100 K, for 3-fluoropropene. The minimum pressure ratio was 2.15 observed at $\Delta T_{lift} =$ 20 K and $T_{source} = 400$ K. The pressure ratios were all below 20 and were below 10 for 12 out of the 15 source-sink pairs. Pressure ratios above ten were only observed at high values of ΔT_{lift} of 80 K and 100 K. This highlights that the design of high temperature lift HP cycles is indeed more challenging than cycles with low ΔT_{lift} and necessitates the addition of equipment-related constraints. ΔT_{super} and ΔT_{sub} remained negligible for all results in case 2. Case 2 had an average total solve time of 36.3 s over a total of 60 optimizations, with a maximum solve time of 165.0 s and a minimum of 12.0 s.

Case 3: Maximizing COP with GWP, ODP, & Design Space Constraints

In Table 6, the best low-GWP and zero-ODP refrigerant that is found for each source-sink temperature pair is shown. GWP values are constrained to be below 150 in accordance with environmental regulations [5]. The COP value is shown in bold. In brackets, the percentage change of this COP value with respect to the best benchmark fluid for that source-sink temperature pair is shown. The refrigerant design space is much more restricted in case 3. Chlorinated molecules are removed from the search space to obtain zero ODP refrigerants; fluorinated molecules are disallowed as several of these may be phased out; and we also exclude nitrogenated molecules.

For case 3 all of the obtained solutions are oxygenated molecules, as shown in Table 6. Ethanol was found to be the best-fit low-GWP refrigerant in 9 out of 15 instances. Methanol was found to be the top solution for 5 out of 15 points. It is not currently used in vapour-compression HPs but several computational studies ([6], [8]) have indicated methanol to be a high-performance refrigerant for HPs. [6] Found methanol to be a high performing refrigerant for high temperature HPs, for COP at condensing temperatures of around 200 °C. [8] Found methanol to be a top performing subcritical HP refrigerant alongside acetone and ethanol, the other

optimal refrigerants identified in this case study. While methanol is not yet used in industrial vapour-compression HPs, it is likely to be widely used given its superior heat transfer properties and low viscosity. It is also used in adsorption HPs, and organic Rankine cycles.

Table 6: Maximization of COP for case 3 at varying source and sink temperatures.

COP		T_{source} (K)				
		340	360	380	400	420
T_{sink} (K)	360	Acetone **7.79** (+11.4%)	–	–	–	–
	380	Methanol **4.81** (+43.6%)	Ethanol **8.22** (+39.8%)	–	–	–
	400	Methanol **3.52** (+60.7%)	Ethanol **5.04** (+60.0%)	Ethanol **8.55** (-0.1%)	–	–
	420	Methanol **2.79** (+82.4%)	Ethanol **3.64** (+80.2%)	Ethanol **5.18** (-1.9%)	Ethanol **9.08** (+0.6%)	–
	440	Methanol* **2.31** (N/A)	Methanol **2.84** (+111.9%)	Ethanol **3.70** (-4.4%)	Ethanol **5.54** (-0.2%)	Ethanol **9.41** (-0.6%)

While acetone, ethanol, and methanol are all stable compounds with several industrial uses, their chief drawback is their flammability. These flammable refrigerants may be used in industry and even homes with appropriate controls and safety features.

Interestingly, imposing constraints on GWP, ODP and restricting the molecular design space led to no reductions in COP on average; instead, a small improvement in COP of 0.04% on average was observed. COP in case 3 differed from case 2 by -3.9% to +4.1%.. For the third integer cut at $T_{source} = 340$ K and $T_{sink} = 440$ K, the optimization failed to converge to a solution after 12 hours, so the reported result (marked with an asterisk) is the best of 2 integer cuts. Across all 59 converged optimizations in case study 3, the average solve time was 18.6 s, with a maximum of 90.3 s and a minimum of 7.2 s.

CONCLUSIONS

This paper takes a multi-scale process systems engineering approach to the design of industrial HPs leveraging process modelling, equations of state, and optimization. Industrial HPs with high COPs and high VHCs were designed in this work for a range of source-sink temperature pairs, demonstrating their potential to decarbonise heat in the chemical industry. New stability constraints were introduced into the CAMPD formulation for the first time to obtain meaningful solutions (see SI). In addition, low-GWP and zero-ODP refrigerants that maximize the COP were designed. The results of case studies 1 and 2 illustrate that there is no single universal optimal refrigerant for all use cases: the optimal refrigerant and process conditions vary with source and sink temperatures, objectives, and design constraints. The results highlight the intrinsic link between refrigerant and heat-pump performance. COP was maximized in cases 2 and 3 in an average of 27.5 s. The optimization implicitly searched through a vast molecular design space with $2 * 11^{19}$ combinations of groups to identify high-performing refrigerants for each source-sink pair. Refrigerants identified in this work include known fluids that other authors have found to be promising and a few novel molecules. First-order GCMs were used to predict physical properties in this work, which may lead to lower-accuracy solutions. Validation steps that consider structural information of the molecule are necessary. For case study 3, we used literature data where available to certify that the identified known molecules indeed had low GWP and zero-ODP. Compressor ratios and discharge temperatures were not limited in this study, and this will be the subject of future work.

DIGITAL SUPPLEMENTARY MATERIAL

The digital supplementary material includes the process constraints, VHC benchmarks, and initial values for decision variables. This can be found at the following link: https://psecommunity.org/LAPSE:2026.0039

REFERENCES

1. International Energy Agency, *Renewables 2023*. IEA Publications, 2023.
2. Thiel GP, Stark AK. To decarbonize industry, we must decarbonize heat. Joule 5:531-550 (2021). https://doi.org/10.1016/j.joule.2020.12.007
3. Marina A, Spoelstra S, Zondag HA, Wemmers AK. An estimation of the european industrial heat pump market potential. Renewable and Sustainable Energy Reviews 139:110545 (2021). https://doi.org/10.1016/j.rser.2020.110545
4. United Nations, The Montreal Protocol on Substances that Deplete the Ozone Layer, Kigali update 2016. Kigali, 2016. doi: 10.4324/9781003597681.
5. European Parliament and Council, 'Regulation (EU) 2024/573 on fluorinated greenhouse gases', *Official Journal of the European Union*. 2024. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/573/oj
6. Zini A, Socci L, Vaccaro G, Rocchetti A, Talluri L.

Working fluid selection for high-temperature heat pumps: a comprehensive evaluation. Energies 17:1556 (2024). https://doi.org/10.3390/en17071556
7. Abedini H, Vieren E, Demeester T, Beyne W, Lecompte S, Quoilin S, Arteconi A. A comprehensive analysis of binary mixtures as working fluid in high temperature heat pumps. Energy Conversion and Management 277:116652 (2023). https://doi.org/10.1016/j.enconman.2022.116652
8. Vieren E, Demeester T, Beyne W, Arteconi A, De Paepe M, Lecompte S. The thermodynamic potential of high-temperature transcritical heat pump cycles for industrial processes with large temperature glides. Applied Thermal Engineering 234:121197 (2023). https://doi.org/10.1016/j.applthermaleng.2023.121197
9. Neumaier L, Roskosch D, Schilling J, Bauer G, Gross J, Bardow A. Refrigerant selection for heat pumps: the compressor makes the difference. Energy Tech 11: (2023). https://doi.org/10.1002/ente.202201403
10. Papaioannou V, Lafitte T, Avendaño C, Adjiman CS, Jackson G, Müller EA, Galindo A. Group contribution methodology based on the statistical associating fluid theory for heteronuclear molecules formed from mie segments. The Journal of Chemical Physics 140: (2014). https://doi.org/10.1063/1.4851455
11. Bowskill DH, Tropp UE, Gopinath S, Jackson G, Galindo A, Adjiman CS. Beyond a heuristic analysis: integration of process and working-fluid design for organic rankine cycles. Mol. Syst. Des. Eng. 5:493-510 (2020). https://doi.org/10.1039/c9me00089e
12. J. Santiuste Martínez, 'Development of a SAFT-γ Mie group contribution approach to modelling working fluids and its application in the computer aided molecular and process design of ORCs', Jul. 2020, [Online]. Available: https://uvadoc.uva.es/bitstream/handle/10324/42747/TFM-I-1655.pdf
13. Hukkerikar AS, Sarup B, Ten Kate A, Abildskov J, Sin G, Gani R. Group-contribution+ (GC+) based estimation of properties of pure components: improved property estimation and uncertainty analysis. Fluid Phase Equilibria 321:25-43 (2012). https://doi.org/10.1016/j.fluid.2012.02.010
14. JOBACK KG, REID RC. ESTIMATION OF PURE-COMPONENT PROPERTIES FROM GROUP-CONTRIBUTIONS. Chemical Engineering Communications 57:233-243 (2007). https://doi.org/10.1080/00986448708960487
15. Y. Nannoolal, J. Rarey, and D. Ramjugernath, 'Estimation of pure component properties Part 2. Estimation of critical property data by group contribution'.
16. Hukkerikar AS, Kalakul S, Sarup B, Young DM, Sin G, Gani R. Estimation of environment-related properties of chemicals for design of sustainable processes: development of group-contribution⁺ (gc⁺) property models and uncertainty analysis. J. Chem. Inf. Model. 52:2823-2839 (2012). https://doi.org/10.1021/ci300350r
17. Hodnebrog Ø, Aamaas B, Fuglestvedt JS, Marston G, Myhre G, Nielsen CJ, Sandstad M, Shine KP, Wallington TJ. Updated global warming potentials and radiative efficiencies of halocarbons and other weak atmospheric absorbers. Reviews of Geophysics 58: (2020). https://doi.org/10.1029/2019rg000691
18. Samudra AP, Sahinidis NV. Optimization-based framework for computer-aided molecular design. AIChE Journal 59:3686-3701 (2013). https://doi.org/10.1002/aic.14112
19. S. Gopinath, 'Molecular design, process design and process synthesis of separation systems', Imperial College London, 2018.
20. B. Zühlsdorf *et al.*, 'High Temperature Heat Pumps, Task 1 - Technologies', Danish Technological Institute, HPT-AN58-2, Aug. 2023. Accessed: Jan. 30, 2026. [Online]. Available: https://heatpumpingtechnologies.org/content/uploads/sites/70/2023/09/annex-58-task-1-technologies-task-report.pdf
21. Wasik DO, Lasala S, Herbinet O, Samukov K, Calero S, Vlugt TJH. Multiscale modeling of dimerization thermodynamics of formic acid. Fluid Phase Equilibria 594:114356 (2025). https://doi.org/10.1016/j.fluid.2025.114356

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Generative AI for the optimal design of seawater desalination processes

Valentin ZARLENGA[abc]*, Antonio ROCHA AZEVEDO[ab], Alvaro MARTINEZ-TRIANA[b] and Thibaut NEVEUX[ab]

[a] Université de Lorraine, CNRS, LRGP, F-54000, Nancy, France
[b] EDF R&D Chatou, 6 quai Watier, 78400 Chatou, France
[c] Université de Strasbourg, ECPM, 25 rue Becquerel, 67200 Strasbourg, France
* Corresponding Author: valentin.zarlenga@univ-lorraine.fr

ABSTRACT

In recent years, research for systematic process design approaches has gained traction, especially with the rise in popularity of generative machine learning models and reinforcement learning. However, works from the literature will often focus on proof-of-concept studies, limited to a specific process synthesis problem. Despite showing promising results, it is not clear how easily these methodologies could be transposed to new applications, and whether they would be successful. In this context, this work evaluates the possibility of using a Natural Language Processing model, which has already proven itself for thermodynamic cycle generation, for another different case: seawater desalination. The processes generated by this model will initially be those using reverse osmosis processes aimed at desalinating a seawater solution containing 25000 ppm of NaCl. Results show that the model has been successful in designing structural reverse osmosis desalination processes without defining assembly rules or a superstructure. Given the large number of these generated processes, a clustering method based on their similarity has been developed. This made it possible to identify different known heuristics (like multi-stage) in process engineering. This adaptation was made possible by modifying aspects external to the original model: a dedicated vocabulary, design rules and objective function.

Keywords: Process synthesis, Seawater desalination, Artificial Intelligence, SFILES, Space visualization

1. INTRODUCTION

In the work by Nabil et al. [1] and Rocha Azevedo et al. [2], a Generative Artificial Intelligence (AI) model for process synthesis was successfully applied to the synthesis of supercritical CO_2 power cycles. It works by treating the process synthesis problem as a Natural Language Processing one: flowsheets are represented by sequences ("sentences") of tokens ("words"). Then, data augmentation and fine-tuning methods are used to achieve targeted generation – generation of relevant flowsheets for a given objective function (defined by the user for each case study). The model identified known design heuristics despite having no prior domain knowledge of the application, while also discovering a novel, counter-intuitive heuristic not previously reported in the literature. This capability proves how useful these methods may be for innovative applications.

This work proposes to extend the method used in [1] and [2] for the application of water desalination. Indeed, water resources have been under increasing stress leading to their depletion. The scarcity of this resource is due to the combined effects of global warming and overexploitation in irrigation, industrial and urban sectors. Faced with this challenge, it is becoming necessary to adopt more sustainable water resource management strategies. The recovery of industrial effluents therefore appears to be an interesting solution.

However, its deployment is often hindered by the high energy and equipment costs associated with conventional water treatment processes. The specific case of sodium chloride / water solution treatment for drinking water production, generally named desalination with seawater, is one of the most investigated problems. Designing optimal process architectures that balance energy, environmental, and economic criteria becomes a major

challenge considering the emerging water management difficulties. Furthermore, the development of new technologies for water treatment, such as membrane distillation, requires further studies to determine their relevance. The use of tools for autonomously evaluating process alternatives could be useful for quickly identifying new and useful solutions.

The primary objective is to assess how this method behaves to this new case study, with a focus on a single-technology seawater desalination system using Reverse Osmosis (RO). The second objective will focus on the generation of the results and how can they be visualized and interpreted. This well-known case study will allow for a critical analysis of its results and comparison with traditional industrial practices. Then, a section will be dedicated to future perspectives.

This generative synthesis approach is expected to explore a wide design space without prior assumptions, potentially identifying novel and highly performant process configurations. This design space will vary depending on the objective function and the constraints defined by each scenario.

2. METHODS

2.1. Data representation

The work by Rocha Azevedo et al. [2] is based on the use of a symbolic language dedicated to the representation of processes: SFILES 2.0. This language converts a process graph, comprising unit operations and their connections, into a sequence of tokens ("words") representing equipment, branches, recycles and tags in the process [3]. It only allows the process flowsheet to be represented without internal parameters.

This notation had previously been modified to facilitate the model's learning process for cyclic processes. The main difference between this version and the original is that a unit operation hierarchy is defined for the SFILES generation (canonical order), improving the consistency of the representation and the model's performance [4]. This version was retained, defining the seawater solution inlet as the starting point and prioritizing the equipment defined for transcribing the process into a sequence of tokens.

In the case of this study, the processes generated will include a raw seawater entry and two global exits represented by the permeate (solution with the lowest salt concentration after treatment) and retentate (solution with the highest salt concentration after treatment). The units available for the synthesis of these processes are: mixer, splitter, pump and reverse osmosis.

2.2. Dataset generation

Currently, there is no large-scale SFILES dataset dedicated to desalination processes. Developing this database is a crucial task. It must contain unique, random (no bias in token assembly) and valid sequences.

The algorithm proposed by Nabil et al. [1] for the generation of supercritical CO_2 cycles is adapted to the case of seawater desalination. The objective is to generate directed graphs (representing the processes), where the nodes correspond to unit operations and the edges to connection flows. The algorithm begins by creating the most basic desalination process: a raw seawater solution, a reverse osmosis or a membrane distillation (equal occurrence) and two exits.

An edge is then randomly selected before applying a mutation. Each mutation is characterized by the random insertion of one of the units available between two nodes linked by an edge. This step is repeated according to a random number between 4 and 20. When the maximum number of mutations is reached, the additional exits added during the mutation step are connected with mixers to ensure that there are only two exits named Retentate and Permeate (randomly). Finally, for each separation operation, the exit edges are randomly labelled "{P}" and "{R}" to distinguish the permeate and retentate flows. This directed graph is then converted into a SFILES sequence and a pre-established set of rules (see Supplementary Material) is used to certify the validity of the processes. 300 000 unique, random and valid sequences were thus generated for both cases.

The definition of the base dataset used to train the model is of utmost importance, as it strongly conditions the research space on which the model will operate. For this reason, it is desired that the dataset is as varied as possible. While in this work we try to achieve this by using purely random mutations (avoiding bias in this step), there isn't a clear metric for studying the variety of such a dataset. As the number of available and total units increases, this task becomes all the more complicated, but the use of exploration methods (such as thermal sampling, which will be mentioned in the next section) can also help the model explore the research space and even go beyond its original scopes.

2.3. Model architecture and implementation detail

In order to explore a large design space, the original article is based on the sequential model including a long short-term memory (LSTM) proposed by Nabil et al. [1]. The model is trained iteratively to produce valid SFILES sequences. There are two distinct phases in this model.

The first phase consists in training the model on the syntactic rules (what makes a SFILES string represent a valid process). In it, the model is trained with random but valid SFILES, stemming from the dataset described on section 2.2, plus other valid sequences generated by the model itself, in a data augmentation loop.

The second phase will finetune the model to focus

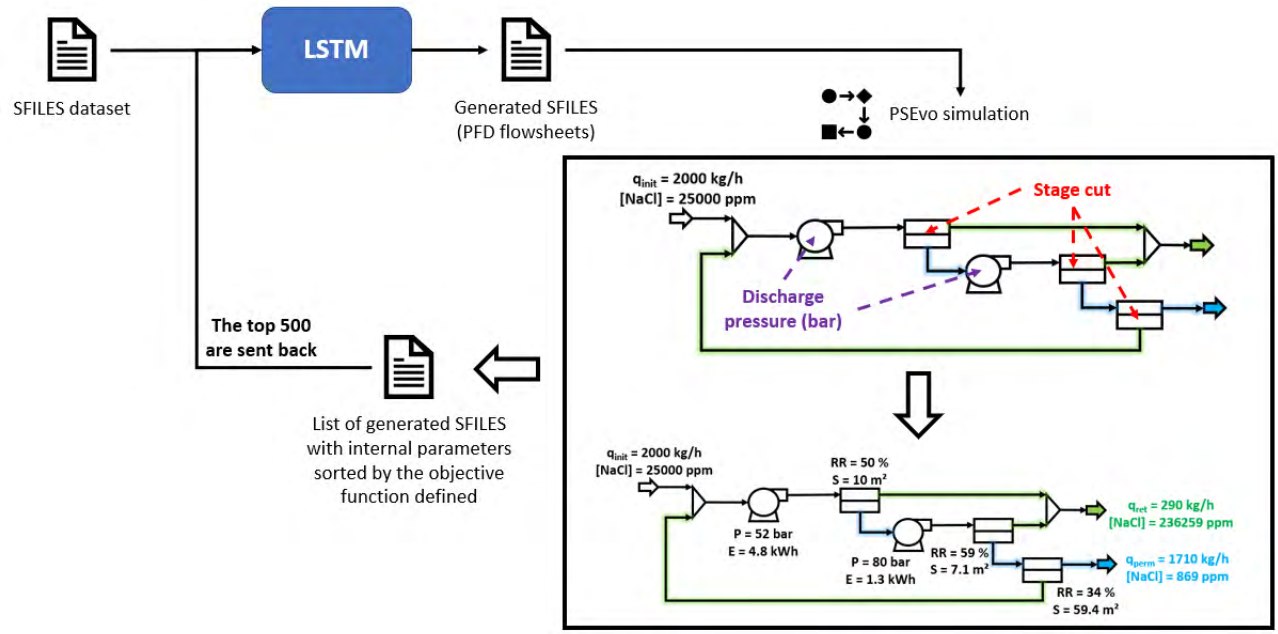

Figure 1. Relationship between the generative AI model and process optimization/evaluation

on the generation on relevant solutions according to a given objective function (which will be discussed in the next section). In this phase, the model is used for generating flowsheets in the form of SFILES. These SFILES processes are sent to the internal EDF simulator (PSEvo) where a simulation is launched. The aim is to maximize the objective function that has been defined by interacting with the design data (such as discharge pressure for the pumps or recovery rate for the reserve osmosis). After the optimization is complete, the process indicators (such as membrane surface, permeate NaCl concentration, permeate flow, ...) and value of the objective function will be extracted in order to evaluate the process efficiency. This internal simulator was crafted for efficiency and distributed calculations, which allowed the method to optimize over 20 000 processes in around 8 000 CPU hours (just a few real-time days thanks to the parallel computations performed by 10 Intel Xeon Platinum 8260 processors). The details of how this software works have been detailed in the article of Rocha Azevedo et al. [2].

Then, following the strategy proposed in [2], the best 500 flowsheets generated so far are used to fine-tune the model (see Figure 1). This process is repeated 10 times, iteratively refining the model and improving the quality of the flowsheets it generates. To improve the flowsheet exploration, a strategy of thermal sampling has been developed. This method allows the distribution of the tokens probability predictions to be influenced. A high temperature flattens the probability distribution, allowing more randomness and novelty in the generated SFILES sequences, but can also lead to more mistakes. Preliminary tests have shown that a temperature of 1.5 generates new design areas (see Supplementary Material), for that reason, it is used throughout the results presented in this work.

2.4. Objectives functions and weights

In the case of seawater desalination by reverse osmosis, we used various criteria to assess the suitability of the generated processes. These criteria are:

- The recovery rate of pure water produced (noted RR)
- The total membrane surface area used by the desalination units (noted S)
- The electrical power consumption required to generate a pressure greater than the necessary osmotic pressure (noted E)
- The NaCl concentration of the treated solution (noted x). This solution must be less than 1000 ppm (noted c_0) total dissolved salts for the process to be considered interesting.

Mathematically, all of these criteria were represented using the following formula:

$$f_{obj} = k_{Tr} \cdot RR - \left(\frac{1}{q_{perm}} \cdot (k_w \cdot E + k_S \cdot S) \right) - \begin{cases} \dfrac{k_c \cdot e^{m \cdot x}}{m \cdot e^{m \cdot c_0}} & \text{if } x < 1000 \\ k_c \cdot (x - c_0) + \dfrac{k_c}{m} & \text{if } x \geq 1000 \end{cases} \quad (1)$$

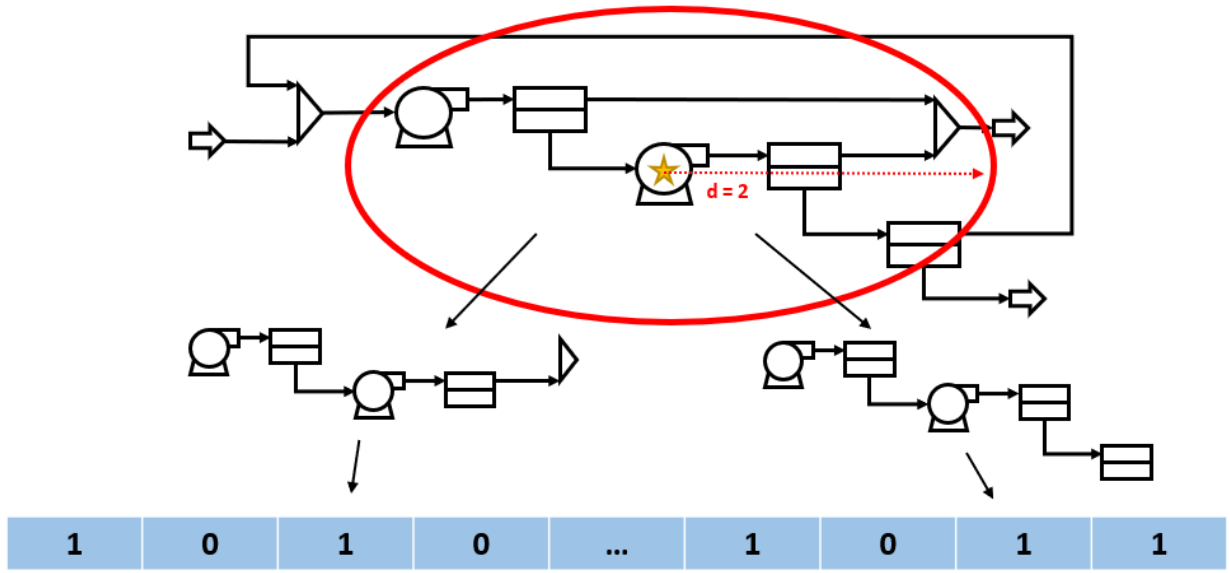

Figure 2. Example of fragmentation for a node in the graph (symbolized by a star) and its vector representation.

Here, weights (k) have been assigned to each of these parameters, in order to standardize their physical magnitude and/or prioritize certain parameters over others (see Supplementary Material). The NaCl concentration penalty will be strong if the permeate concentration is greater than 1000 ppm, and quickly (exponentially) otherwise. The term "m" is another weight that influences the slope of the exponential decrease.

2.5. Visualization of the generated design space

The AI method presented generates a large number of processes, making it impossible to analyze each process individually. It is therefore necessary to develop a global analysis strategy to extract a set of relevant information. The essential information is the common substructures of the processes and the search space explored.

Many articles in computational chemistry using AI models transform molecules into binary vectors (vectors whose elements are 0 or 1 called fingerprints) of fixed dimension. This representation is particularly suitable for visualization or similarity analysis. There are various vectorization processes (the most well known being Extended Connectivity Fingerprints [5]) based on rules specific to chemistry. The aim is not to replicate these algorithms but to redefine a digital representation strategy adapted to the processes field.

To do this, the SFILES sequences generated by the model are transformed into a directed graph. For each node in the graph, a subgraph containing all nodes at a distance less than or equal to d from the first node is isolated. This fragmentation is applied to all the processes to be studied. When all the unique subgraphs have been extracted, they are associated with a component of a vector, the dimension of which is fixed by the number of unique subgraphs. For each process, this vector is initialized with zeros. If the process includes a particular subgraph, its component takes the value 1, otherwise it remains at 0. This mechanism allows the sequences generated by the model to be represented as a binary vector indicating the presence of the local patterns that constitute it. It is important to note that the distance d is a parameter set by the operator (see Figure 2). However, it should be not too low (low discriminating patterns) or too high (many long and unique patters creating sparse vectors with little information). The value set in this work is 2.

To visualize the search space, the UMAP dimensionality reduction method is applied. UMAP (Uniform Manifold Approximation and Projection) allows complex nonlinear data to be projected into a two-dimensional space while preserving the local and global topologies of the processes (suitable for revealing significant clustering patterns) [6]. The more structurally similar two processes are, the most similar the points representing them in the reduce space will be. Once the projection had been made on all the processes, a density-based clustering technique (HDBSCAN [7]) was applied to group similar processes together. Thanks to the binary vectors, we can quickly identify the most common or dominant substructures in these process families.

3. RESULTS AND DISCUSSION

This section focuses on the reverse osmosis desalination processes generated. The first subsection will

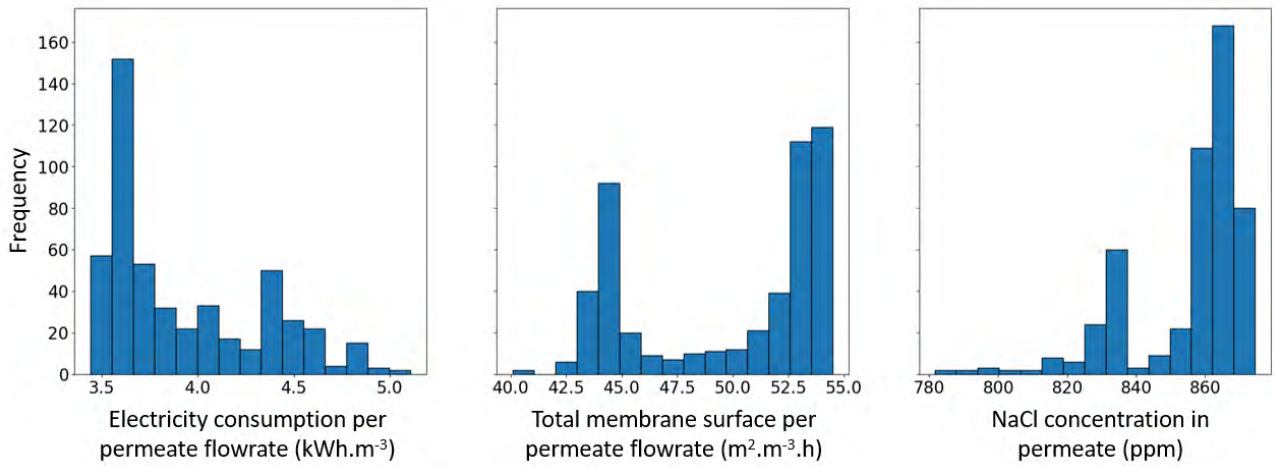

Figure 3. Histogram of internal parameter for the 500 best processes studied in section 3.2

investigate the steps necessary to adapt the initial model. Then, the next subsection focuses on the model exploration to determine if it converges towards a single zone (meaning that a single global architecture is found) or if several zones are explored. The 500 best solutions proposed by the model will be then investigated. Because this case study is well documented, it will be the basis for evaluating the model, as we can compare the flowsheets the model generates to known relevant processes. Finally, the perspectives of this work will be presented.

3.1. Model adaptation

This section lists the steps required to modify the Rocha Azevedo et al. model [2] in order to apply it to our case study.

The first step was to create the tokens used in the application. Each token has its own name, as well as an inlet number, an outlet number and a tag name for these currents (see Supplementary Material). A set of rules was created to verify the physical validity of the generated processes (see Supplementary Material). Otherwise, non-physical processes could have been simulated. When the valid token sequence needs to be optimized and then evaluated, it must be transformed into an incidence matrix where each column corresponds to a specific equipment. The algorithm used to do this had therefore to be adapted to assign the correct token to the correct simulator equipment. The same applies to the permeate/retentates flowrates. Another aspect that been modified is the post processing of the processes. Equipment that is not used for the process or unused flows has already been removed, but this post processing has been extended. The objective function has been also modified.

The LSTM model has not been modified except for the integration of thermal sampling strategy.

3.2. Evaluation of model exploration

In order to visualize how the model explores the research space, the search space has been modeled (see section 2.5) for the original random dataset (iteration “0”) plus the flowsheets generated by the method at each fine-tuning iteration (iterations 1 – 10). Only the flowsheets with a feasible value of objective function (equation 1) are considered. The objective is to highlight the structural evolution of the processes and the detection of areas that improve the objective function.

The results (see Figure 4) show that the generated processes differ significantly from those sent to the model in terms of structure and efficiency. Analysis of the objective function of the processes has highlighted regions that are not particularly relevant to the application. Future processes generated will tend to move away from these regions in order to exploit new designs in known regions or explore other regions.

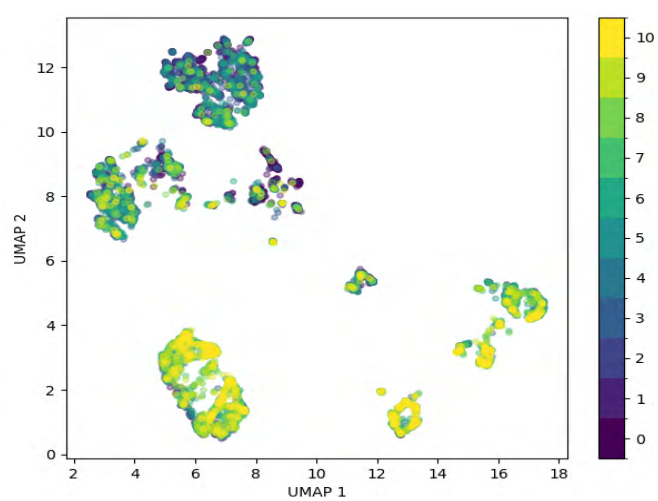

Figure 4. : Visualization of the processes generated by the model based on iteration (colorbar).

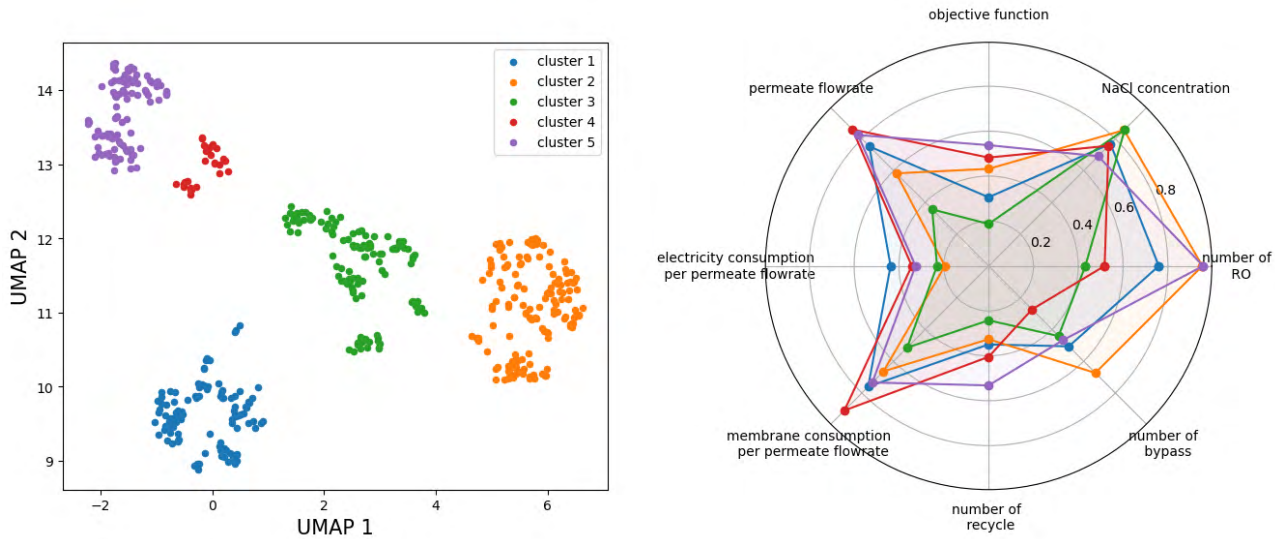

Figure 5. Process clustering (left) and Kiviat diagram of normalized data for each cluster (right) for the most relevant processes.

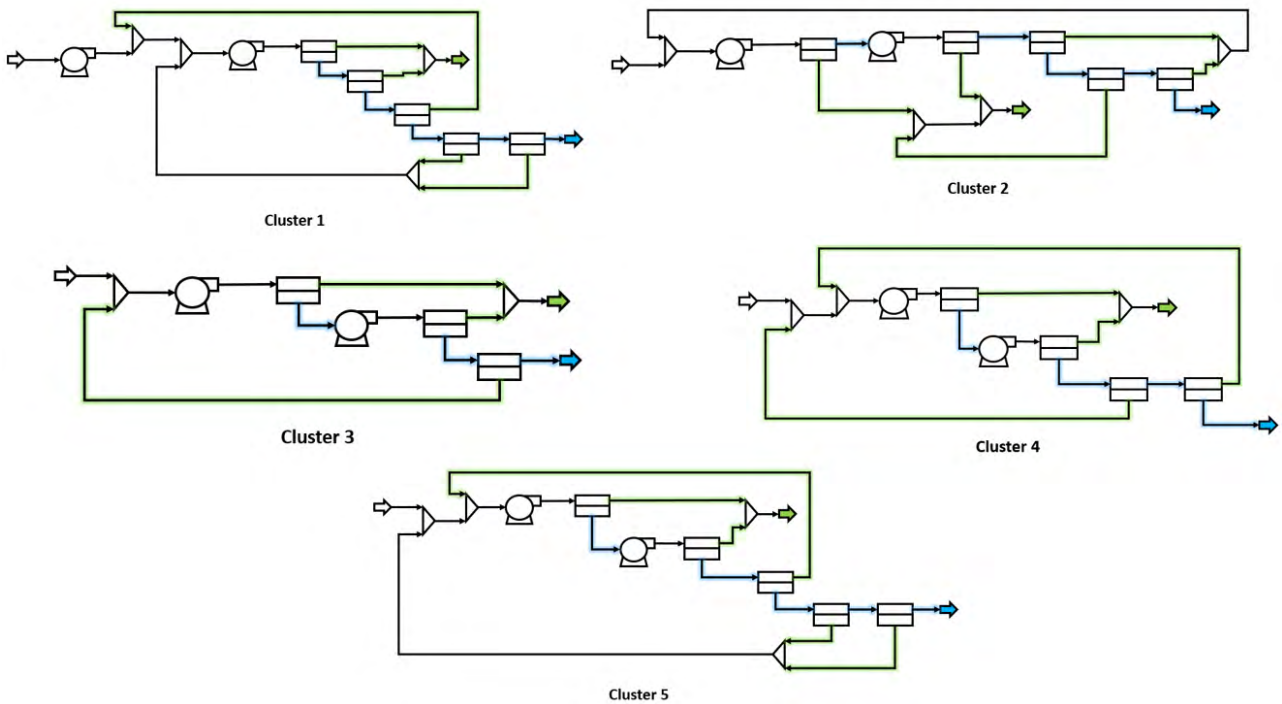

Figure 6. Randomly selected process of the each cluster. The green element correspond to retentate flows and the blue element correspond to permeate flows.

3.3. Exploitation of the most relevant processes

The 500 best processes have been selected for this study. Analysis of the results show that electricity consumption ranges from 3 and 5 kWh/m^3 and total membrane surface used is between 40 and 55 m^2/m^3. All these processes produce water with less than 1000 ppm in salt (see Figure 3).

Using the method described in section 2.5, five clusters were detected, each with different design architectures. According to Kiviat diagram (see Figure 6), the elements of cluster 5 have the highest objective function. A random process was therefore selected for each cluster in order to analyze its architecture in detail (overview of these processes are given in Supplementary Material).

As illustrated in the Figure 6, a number of well-

known and widely documented heuristics were found in these processes, such as:

- Use of multi-stage reverse osmosis (between 3 and 5 stages) allowing separation requirements to be distributed across several levels.
- Excessively concentrated rejects are not reprocessed to avoid the use of very high hydraulic pressure or excessive membrane use by sending them directly to the retentate.
- Low concentration effluents are reprocessed to maximize water recovery.
- Use of partial recirculation

The processes of clusters 2, 3, 4 and 5 present a heuristic that those in cluster 1 do not share: a second pressurization stage after the first reverse osmosis unit. This allows us to reduce the level of electricity needed to obtain the same permeate flowrate (see Supplementary Material). The cluster 1 processes start directly at the highest possible pressure, which requires much more electricity to operate.

Unlike the other processes, those generated in cluster 3 have less reserve osmosis. Looking at the appearance of the processes, they are generated at the beginning (iteration 1) and then their number are falling in favor of processes incorporating more complex designs (see Figure 7). This proves that the model tends to move away from simple solutions to explore possibilities for designs that are more relevant to a process engineering expert.

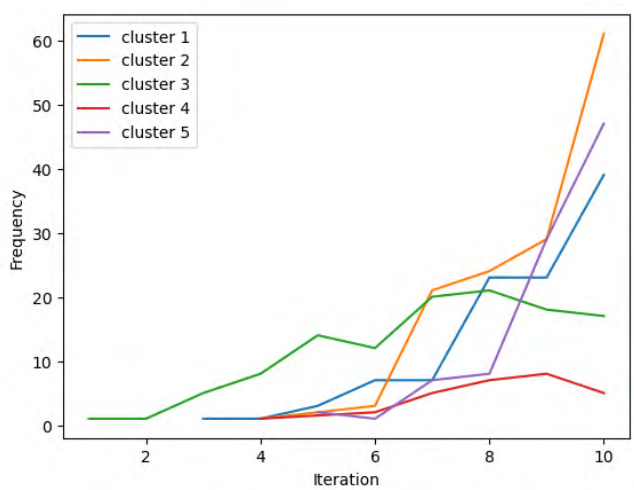

Figure 7. Appearance of processes in each cluster based on iterations.

3.4. Future perspectives

This work has demonstrated the ability of AI to generate processes incorporating reverse osmosis and certain known heuristics with strong potential for industrial application. However, several perspectives can be implemented in order to refine the results. First, the objective function can be redefined to include more elements related to CAPEX or OPEX. This economic dimension could potentially lead to identifying simpler processes (with less separation stages) for achieving the same performance indicators. The penalty for NaCl concentration in the permeate could also be increased to obtain a weaker NaCl solutions in the permeate for the most interesting processes. It would also be interesting to have a multi-objective function and study the Pareto fronts of the generated solutions.

Nevertheless, scientific literature mentions numerous studies based exclusively on this separation technology. The integration of multi-technology processes appears to be a more relevant study with the use of membrane distillation. This type of hybrid process is generating interest despite of a lack of knowledge. A working perspective of readjusting the model and then retraining it to design this type of hybrid process (reverse osmosis and membrane distillation) should also be explored.

Finally, one could imagine coupling this synthesis approach with others methods. This could either be done: upstream, mixing known heuristics to the dataset during training, for nudging the method in the right direction; or downstream, automatically extracting useful patterns from the data generated for the development of superstructures which can better refine the research space.

4. CONCLUSION

The objective of this work was to evaluate the integration of AI to propose relevant processes for water recovery. Although the original model was proposed for thermodynamic cycle design, processes using reverse osmosis for seawater desalination were developed. These processes offer satisfactory energy and material efficiency and incorporate design architectures commonly used in industry (multi-stage reserve osmosis, ...). Generative AI has therefore made it possible to find known heuristics without prior knowledge. These heuristics were easily highlighted thanks to the development of a visualization method using UMAP and then clustering.

5. NOMENCLATURE

- RR: Recovery Rate
- S: membrane surface used
- E: electricity power consumption
- q_{perm}: permeate flowrate
- [NaCl]: NaCl concentration
- P: Pressure discharge
- q_{ret}: retentate flowrate

- q_{init}: initial flowrate

6. DIGITAL SUPPLEMENTARY MATERIAL

The supplementary material contains additional data and incorporates additional references [8]. See https://psecommunity.org/LAPSE:2026.0030.

7. ACKNOWLEDGEMENTS

The authors would like to express gratitude to the EDF R&D for supporting this research.

8. AUTHOR IDENTIFIERS

Author ORCIDs:
Valentin ZARLENGA: 0009-0006-6937-6884
Antonio ROCHA AZEVEDO: 0000-0002-0633-1252
Alvaro MARTINEZ-TRIANA: 0000-0002-0259-9833
Thibaut NEVEUX: 0000-0002-7460-702X

9. REFERENCES

1. Tahar Nabil, Mohamed Noaman et Tatiana Morosuk. "Data-driven structural synthesis of supercritical CO2 power cycles". In: Frontiers in Chemical Engineering 5 (16 oct. 2023). Publisher: Frontiers. issn: 2673-2718. doi: http://dx.doi.org/10.3389/fceng.2023.1144115. url: https://www.frontiersin.org/journals/chemical-engineering/articles/10.3389/fceng.2023.1144115/full.
2. Rocha Azevedo A, Nabil T, Loubière V, Privat R, Neveux T, Commenge JM. Comparing generative process synthesis approaches with superstructure optimization for the conception of supercritical CO2 brayton cycles. Computers & Chemical Engineering 201:109255 (2025). https://doi.org/10.1016/j.compchemeng.2025.109255
3. Vogel, G., Hirtreiter, E., Schulze Balhorn, L., Schweidtmann, A.M., 2023. SFILES 2.0: An extended text-based flowsheet representation. Optim. Eng. 24 (4), 2911–2933. http://dx.doi.org/10.1007/s11081-023-09798-9. URL: https://link.springer.com/10.1007/s11081-023-09798-9.
4. Antonio Rocha Azevedo et al. Representation Learning for flowsheets: generating structures for Process Synthesis. Publication Title: 34th European Symposium on Computer Aided Process Engineering / 15th International Symposium on Process Systems Engineering (ESCAPE34- PSE24). Mai 2024. 372-377. doi: http://dx.doi.org/10.3303/BOA2401. Url: https://hal.science/hal-04733044.
5. David Rogers et Mathew Hahn. "Extended-Connectivity Fingerprints". In: Journal of Chemical Information and Modeling 50.5 (24 mai 2010), p. 742-754. issn: 1549-9596, 1549 960X. doi: http://dx.doi.org/10.1021/ci100050t. Url: https://pubs.acs.org/doi/10.1021/ci100050t.
6. Leland McInnes, John Healy et James Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. 18 sept. 2020. doi: https://doi.org/10.1038/s43586-024-00369-5 Url: http://arxiv.org/abs/1802.03426.
7. McInnes L, Healy J. Accelerated hierarchical density based clustering. 2017 IEEE International Conference on Data Mining Workshops (ICDMW) :33-42 (2017). https://doi.org/10.1109/icdmw.2017.12
8. Hinton, Geoffrey & Dean, Jeff & Vinyals, Oriol & Rachmad, Yoesoep. (2014). Distilling the Knowledge in a Neural Network. 1-9. url: https://arxiv.org/pdf/1503.02531.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Assessing the Impact of Solvent Recycling in Cooling Crystallization using Computer-Aided Molecular and Process Design

Gaurav Seth[a], Saman Naseri Boroujeni[a], Shubhani Paliwal[a], Amparo Galindo[a], George Jackson[a], and Claire S. Adjiman[a*]

[a] Department of Chemical Engineering, The Sargent Centre for Process Systems Engineering, Imperial College London, London SW7 2AZ, United Kingdom

* Corresponding Author: c.adjiman@imperial.ac.uk.

ABSTRACT

Although solvent-based crystallization is widely adopted for separation and purification of crystalline pharmaceutical products, solvent choice and utilisation critically influence product quality, manufacturing cost, and the environmental performance of the pharmaceutical process. Escalating demands to reduce process mass intensity (PMI), together with increasing vulnerabilities in the supply chains, necessitate the development of more efficient and resilient process designs, incorporating solvent and active pharmaceutical ingredient (API) recycling. The conceptual design of crystallization processes offers a viable route to identify flowsheets with substantially reduced solvent consumption. In this paper we present a computer-aided molecular and process design (CAMPD) formulation to explore the benefits of solvent/API recycle for two processes/APIs: (i) a continuous cooling crystallization process for mefenamic acid (MA) employing a binary solvent mixture and (ii) a batch cooling crystallization for paracetamol using a single solvent. At the conceptual design stage, the process is assumed to operate at thermodynamic equilibrium and the SAFT-γ Mie group contribution approach is used as a rigorous thermodynamic model. The optimal design for each process is determined based on the trade-off between two key performance indicators (KPIs), crystallization yield and solvent consumption. The resulting bi-objective optimization (BOO) problem is implemented in gPROMS. In addition to the solvent identities the proposed methodology also enables the simultaneous optimization of process temperatures, and the feed composition. The results demonstrate that solvents and API recycling can deliver significant benefits, leading to reduced process mass intensity (PMI) while maximizing the crystallization yield.

Keywords: Crystallization, Solvent selection, SAFT, Process design, Optimization

INTRODUCTION

Over 80% of small-molecule pharmaceuticals are delivered in crystalline form, including tablets and capsules [1]. Consequently, crystallization and purification are crucial components in pharmaceutical manufacturing. The design of isolation and purification steps impacts the physical properties of active pharmaceutical ingredients (APIs), which influence the therapeutic performance of the drug through its efficacy, absorption, and bioavailability. Additionally, crystalline form properties affect the performance of downstream operations such as filtration, drying, milling, and tableting [2].

Most industrial crystallization processes are solvent based, primarily due to the operational simplicity of such systems [3]. However, the choice of solvent(s) has a profound impact on the overall efficiency and sustainability of the crystallization process. From a thermodynamic perspective, solvent selection influences the solubility of the API, thereby affecting both the attainable crystal yield and the solvent volume required for crystallization. Morover, solvent usage accounts for over 70% of the total energy demand and nearly 50% of the greenhouse gas emissions associated with API production [4].

Consequently, systematic approaches to solvent choice are fundamental to developing crystallization processes that are both efficient and environmentally sustainable.

Considering these challenges, increasing attention has been directed towards integration of solvent and API recycling strategies within crystallization processes [5-7]. Among these, Wong et al. [5] and Ferguson et al. [6], have carried out experimental studies on single stage, mixed-suspension, mixed-product removal (MSMPR) crystallizers with recycle and have shown that recycling mother liquor back to the crystallizer and manipulating the recycle ratio can boost the crystallization yield.

From a computer-aided design perspective, several systematic methods for integrated solvent (or solvent mixture) and process design have been reported for standalone crystallization [2, 8-11]. Various strategies have been adopted to solve such computer-aided molecular and process design (CAMPD) problems. These include generate-and-test, in which property prediction methods or databases are used to generate a list of candidate solvents to be evaluated [11], decomposition-based approach [8, 10], where optimization subproblems are used to identify molecules for which the process is further optimized; and simultaneous optimization of solvent/solvent blend and process [2, 9], where a mixed-integer nonlinear program (MINLP) is solved with respect to the process variables and the molecular structure. As expected, optimal binary solvent mixtures often results in greater solubility than pure solvents [12].

Comparatively limited attention has been given to the design of crystallization processes with solvent and API recycling. Wang and Lakerveld [7] have recently presented a continuous molecular targeting (CoMT-CAMPD) approach for solvent selection in a continuous antisolvent crystallization process. Their process configuration integrated a crystallizer with a flash drum to recycle and reuse the antisolvent. Their approach was based on the perturbed-chain statistical associating fluid theory (PC-SAFT) equation of state [13] as the underlying thermodynamic model. The resulting MINLP problem was relaxed into a sequence of nonlinear programming (NLP) problems via continuous mapping of solvent parameters.

Building on these advances, the aim of this work is to employ a simultaneous CAMPD formulation to investigate the potential benefits of solvent/API recycling for two process configurations/APIs. The first process configuration considered, Case 1, is adapted from the continuous cooling crystallization design proposed by Wang and Lakerveld [7] using a binary solvent mixture. Two solvent and API recycling strategies are examined: direct recycling of the mother liquor and recovery via a flash separation step. The performance of these recycling strategies is compared with that of standalone crystallization, with solvent identity and process operating conditions optimised in each case. The second process configuration examined, Case 2, is adapted from a laboratory-scale batch cooling crystallization setup for paracetamol using a single solvent [19]. Here pure solvent recovery is achieved via a rotary evaporator, after which the recovered solvent is reused for washing the filtered crystal cake and subsequently recycled to the crystallizer together with the fresh feed. For both case studies, the thermodynamic properties of the pure solvent and the mixtures are evaluated using SAFT-γ Mie [14].

The design approach, including the process models and design constraints for both configurations, is discussed in the next section. The results of BOO are presented in subsequently the Case studies section and the final conclusions are drawn in the last section.

DESIGN APPROACH

Problem definition

The CAMPD problem addressed in this work can be stated as: Given a crystallization process flowsheet for a pharmaceutical compound, a specified product purity, and a list of potential solvents, determine the optimal combination of solvent/solvent mixture, mixture composition and process conditions. At this stage of conceptual design, the process model is formulated based solely on thermodynamic equilibrium considerations and material balance equations.

To facilitate the formulation of the process models, a number of sets are defined as follows.

- Q denotes the set of solvents used in the process. In this work, case studies with one or two solvents are used so that $Q = \{s1, s2\}$ or $Q = \{s1\}$.
- I denotes the set of impurities present in the feed stream.
- C denotes the set of all components that may appear in any liquid stream within the process, i.e., $C = Q \cup I \cup \{API\}$, where API represents the desired product.
- $C^k \subseteq C$ represents the set of components that may appear in the solid phase.
- S represents the set of candidate solvents from which the solvents used in the process (Q) are selected.
- $\mathcal{T} = \{1, 2, \dots, N\}$ is the set of process streams, where $N = 10$ for Case 1 and $N = 22$ for Case 2.
- G is the set of functional groups used to represent the solvents in set S.

Model formulation

Continuous Cooling Crystallization – Case 1

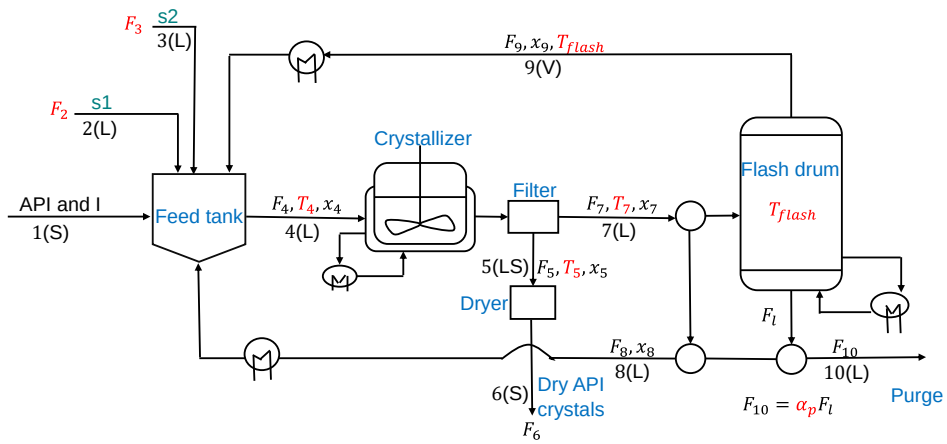

Figure 1: Flowsheet for continuous cooling crystallization of mefenamic acid. The letters L-liquid, S-solid, LS-liquid and solid, and V-vapour shown in parentheses after each stream number indicate the phases present in the corresponding process streams.

The continuous cooling crystallization flowsheet for a binary solvent mixture with solvent and API recycle via a flash drum is shown in Figure 1. The mass of component $i \in C$ in stream $t \in \mathcal{T}, M_{t,i}$, is given by

$$M_{t,i} = F_t x_{t,i} MW_i, i \in C; t \in T \tag{1}$$

where F_t is the molar flow rate of stream t, and $x_{t,i}$ represents the mole fraction of the ith component in stream t, such that $\sum_{i \in C} x_{t,i} = 1, \forall t \in \mathcal{T}$.

Stream 1 is the feed entering following upstream synthesis and contains only API and impurities; hence, $x_{1,i} = 0, i \in C \backslash \{API\} \cup I$. Streams 2 and 3 correspond to pure solvent $s1$ and $s2$, respectively, $x_{2,i} = 0, i \in C \backslash \{s1\}$ and $x_{3,i} = 0, i \in C \backslash \{s2\}$. The material balance around the feed tank is given as:

$$F_1 x_{1,i} + F_2 x_{2,i} + F_3 x_{3,i} + F_8 x_{8,i} + F_9 x_{9,i} = F_4 x_{4,i}, i \in C \tag{2}$$

while that around the crystallizer and filter is given as:

$$F_4 x_{4,i} = F_5 x_{5,i} + F_7 x_{7,i}, i \in C \tag{3}$$

Stream 5 consists of the wet API cake and is modelled as a two-phase mixture of solid and liquid phases:

$$F_5 x_{5,i} = F_5^{ml} x_{5,i}^{ml} + F_5^{s} x_{5,i}^{s}, i \in C \tag{4}$$

where superscripts ml and s denote the mother liquor and solid phases, respectively. The solvent holdup in the wet cake is characterized by a mass fraction α_h:

$$\Sigma_{i \in C} M_{5,i}^{ml} = \alpha_h \left[\Sigma_{i \in C} M_{5,i}^{ml} + \Sigma_{i \in C} M_{5,i}^{s} \right] \tag{5}$$

It is further assumed that no compositional change occurs in the liquid phase during filtering; therefore, the mole fraction of mother liquor retained within the wet API cake is equal to that of the liquid stream exiting the crystallizer, stream 7, i.e., $x_{5,i}^{ml} = x_{7,i}$. Additionally, 100% impurity rejection is assumed in the crystallizer, such that only the API reaches solid–liquid equilibrium (SLE), while all impurities remain undersaturated in the mother liquor (stream 7), i.e., $x_{i,5}^{s} = 0, i \in C \backslash \{API\}$. Liquid stream 7 is at solid liquid equilibrium (SLE) with the solid APIs in stream 5, and the API solubility is calculated as:

$$\ln x_{7,API} \gamma_{7,API} = \frac{\Delta H_{API}^{m}}{R} \left(\frac{1}{T_{API}^{m}} - \frac{1}{T_7} \right) \tag{6}$$

where $\gamma_{7,API}$ denotes the activity coefficient of the API at equilibrium obtained using SAFT-γ Mie, ΔH_{API}^{m} is the molar heat of melting of the API, T_{API}^{m} is the melting temperature of the API, R is the ideal gas constant, and T_7 is the crystallizer operating temperature.

Stream 5 is sent to a dryer to produce dry API crystals (stream 6), with residual solute impurities originating from the solvent holdup within the wet cake.

Material balance around the flash drum is given as:

$$F_5 x_{5,i} = F_9 x_{9,i} + F_{10} x_{10,i} + F_8 x_{8,i}, \forall i \in C \tag{7}$$

where F_9 is vapour outlet flow rate from the flash drum. The liquid stream leaving the flash drum is split into a recycle stream (F_8) and a purge stream (F_{10}), with purge fraction defined as α_p, such that $\alpha_p = (1 - \alpha_p) F_{10} / F_8$. Phase equilibrium at the flash temperature, T_{flash} and atmospheric pressure is determined via a temperature-pressure flash (TP-flash). Furthermore, API and the impurities are assumed to be non-volatile and therefore, $x_{9,i} = 0, i \in C \backslash \{s1, s2\}$.

In some streams, at least one component in C^k should remain in the liquid phase. For a stream t, the set $\hat{C}_t^k \subseteq C^k$ of such component is defined and the solubility limit of a component $i \in \hat{C}_t^k$ is computed by assuming that only pure crystals can form (no co-crystals) [15] and by defining a hypothetical mixture with mole fraction vector $\hat{\boldsymbol{x}}_t^{(i)}$ in which the relative proportions of all components are identical to those in stream t:

$$\hat{x}_{t,i'}^{(i)}/\hat{x}_{t,s1}^{(i)} = x_{t,i'}/x_{t,s1}, i' \in C\backslash\{s1, i\};\ i \in \hat{C}_t^k\ ; t \in \{4,7,8\} \quad (8)$$

with the normalization condition $\Sigma_{i' \in C}\hat{x}_{t,i'}^{(i)} = 1, i \in \hat{C}_t^k$ and component i is assumed to be at SLE:

$$\ln \hat{x}_{t,i}^{(i)}\hat{\gamma}_{t,i}^{(i)} = \frac{\Delta H_i^m}{R}\left(\frac{1}{T_i^m} - \frac{1}{T_7}\right), i \in \hat{C}_t^k; t \in \{4,7,8\} \quad (9)$$

where $\hat{\gamma}_{t,i}^{(i)}$ is the activity coefficient of component i in the hypothetical stream. The undersaturation is then imposed as:

$$x_{t,i} \leq \hat{x}_{t,i}^{(i)} - \epsilon_c, i \in \hat{C}_t^k \quad (10)$$

where ϵ_c is a small positive tolerance. These equations are applied to stream 7, where only the API crystallizes, ($\hat{C}_7^k = C^k\backslash\{API\}$), and streams 4 and 8, where no component crystallizes ($\hat{C}_4^k = \hat{C}_8^k = C^k$).

Process design constraints: Liquid range constraints are enforced so that the solvent mixture is in the liquid state in specific units:

$$T_t \leq T_{mix,t}^{bub} - T_0, \forall t \in \mathcal{T}^L \quad (11)$$

Here, T_t and $T_{mix,t}^{bub}$ are the operating temperature and the bubble-point temperature of stream t, respectively, and $\mathcal{T}^L$ is a relevant subset of $\mathcal{T}$. For continuous crystallization $\mathcal{T}^L = \{4,7\}$. Bubble-point temperatures are calculated using SAFT-γ Mie. Constraints on the melting points of the solvent(s) are not considered as the solvents have been chosen to have melting points well below the minimum process temperature. Additionally, the quaternary mixture (s1, s2, API, and I) is required to form a homogeneous liquid phase in streams $t \in \{4,7,8\}$. This is ensured by constraining a logical stability function $\phi(T_t, P_t, \boldsymbol{x}_t)$ to unity. $\phi = 1$ denotes that a TP-flash calculation at the specified conditions identifies neither a second liquid phase nor a vapour phase, while $\phi = 0$ indicates the presence of an additional fluid phase.

Switching configurations: In the equations presented so far, it is assumed that recycling using flash drum takes place. To model the direct-recycle configuration, the following constraints are introduced:

$$T_{flash} \leq Mz + T_7 + \epsilon_z \quad (12)$$

$$T_{flash} \geq T_7 + \epsilon_z(1 - z) \quad (13)$$

$$T_{flash} \geq zT_{mix,7}^{bub} + \epsilon_z \quad (14)$$

ϵ_z is a small positive constant, M is a large positive constant, z parameter indicating the status of the flash drum: (active for $z = 1$, in-active for $z = 0$). denotes a binary parameter indicating the operating status of the flash drum.

Batch Cooling Crystallization – Case 2

Figure 2 depicts the batch cooling crystallization process. The crystallization is performed in a pure solvent and is followed by filtration, washing and drying. Due its batch nature, the process requires a few iterations to reach a steady state. While this is considered via a batch index b and a total number of batches N_b, the detailed dynamics of the process units are neglected at the conceptual design stage. Each unit is modelled by considering the initial time (inlet) and final time (outlet). As a result, the mass balances and phase-equilibrium formulations for the crystallizer, filter, and dryer are analogous to those described previously (with the addition of the batch index b) and are not repeated here.

As in the continuous crystallization model, 100% impurity rejection is assumed during batch crystallization. Specifically, only the API is allowed to reach SLE, while impurities are constrained to remain undersaturated in the relevant mother liquor streams (Streams 1, 2, 5, 17, 18, and 22). Furthermore, temperature constraints are also imposed for streams $t \in \mathcal{T}^L = \{1,2\}$ (cf. Eq. (11)). After filter 1, the cake is washed using a fraction of pure solvent from the solvent recovery tank. The mass of wash solvent in batch b, M_{14}^b is modelled using wash mass fraction, α_w:

$$M_{14}^b = \alpha_w M_{3,API}^{s,b} \quad (15)$$

where $M_{3,API}^{s,b}$ denotes the mass of solid API in stream 3.

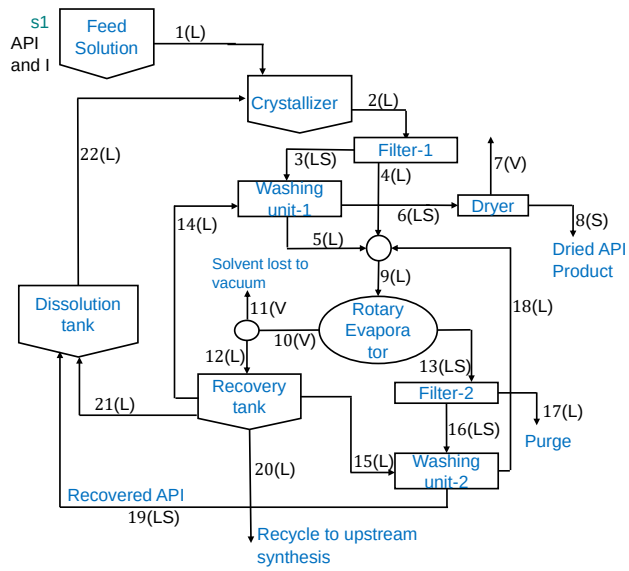

Figure 2: Flowsheet for Batch cooling crystallization

The washing operation is assumed to occur at room temperature (i.e., $T_5^b = T_5 = 293\ \mathrm{K}$) with the resulting mother liquor is at solid-liquid equilibrium (SLE) with respect to the API. Accordingly, API solubility in stream 5 is given by:

$$\ln x_{5,API}^b \gamma_{5,API}^b = \frac{\Delta H_{API}^m}{R}\left(\frac{1}{T_{API}^m} - \frac{1}{T_5}\right) \quad (16)$$

The resulting liquid streams from filter 1 (stream 4) and washing units 1 and 2 are combined and fed to a

rotary evaporator, where a mass fraction, α_{evap}, of the total solvent amount in the combined stream (stream 9) is evaporated and collected in the solvent recovery tank. The corresponding mass balances are given by

$$M_{9,i}^b = M_{5,i}^b + M_{4,i}^b + M_{18,i}^b, i \in C \quad (17)$$

$$M_{10}^b = \alpha_{evap} M_{9,s1}^b = M_{11}^b + M_{12}^b \quad (18)$$

$$M_{11}^b = \alpha_c M_{10}^b \quad (19)$$

where α_c represents the fraction of evaporated solvent lost in the vacuum system. For clarity, the component index i is omitted for the streams containing pure solvent. The concentrated mother liquor leaving the rotary evaporator is cooled to room temperature ($T_{13}^b = T_{13} = 293$ K) and sent to filter unit 2 to recover additional API. The filter cake is washed using solvent from the solvent recovery tank at room temperature. Next, the washed cake is redissolved and recycled to the crystallizer. The filtrate is discarded (stream 17) whereas the wash liquor (stream 18) is recycled back to the evaporator.

The composition of the recycle stream (stream 22) is constrained such that the mass ratio of API to the solvent is identical to that of the fresh feed stream. i.e.,

$$M_{API,22}^b / M_{s1,22}^b = M_{API,1}^b / M_{s1,1}^b \quad (20)$$

This ensures consistent feed composition upon recycle.

Design Constraints and MINLP formulation

Additional constraints and model equations are required to describe the selection of a solvent from the set of candidates S and to quantify process performance.

Solvent assignment constraints: Logical constraints are imposed on the binary variables representing the choice of solvent. Exactly one solvent from set candidate S is assigned to the solvent component(s).

$$\Sigma_{s \in S} y_{q,s} = 1, \forall q \in Q \quad (21)$$

For continuous cooling crystallization configuration, $Q = \{s1, s2\}$, whereas for batch cooling crystallization, $Q = \{s1\}$. To prevent the same candidate solvent from being assigned to both solvent components in a binary solvent mixture, the following constraint is imposed:

$$\Sigma_{q \in Q} y_{q,s} \leq 1, \forall s \in S \quad (22)$$

Since thermodynamic properties are computed using a group contribution approach, each solvent candidate s is represented by a set of functional groups $n_{s,k}$, where k is a functional group in set G. The solvents $q \in Q$ are then defined as:

$$\tilde{n}_{q,k} = \Sigma_{s \in S} y_{q,s} n_{s,k}, \forall q \in Q, \forall k \in G \quad (23)$$

where $\tilde{n}_{q,k}$ denotes the number of functional group of type k present in solvent q.

KPIs: The KPIs considered in this work are process mass intensity (PMI), product purity, and crystal yield.

- PMI is sustainability metric that quantifies material efficiency. For continuous crystallization process, it is calculated as:

$$PMI = \Sigma_{t \in \{1,2,3\}} \Sigma_{i \in C} M_{i,t} / M_{API,6}\,, \quad (24)$$

For batch crystallization, the overall PMI is evaluated over N_b batches by accounting explicitly for the solvent recovered and recycled to upstream synthesis:

$$PMI = \left(\Sigma_{b=1}^{N_b} \left(\Sigma_{i \in C} M_{i,1}^b - M_{s1,20}^b\right)\right) / \Sigma_{b=1}^{N_b} M_{API,8}^b \quad (25)$$

- Product purity (PP) is defined as the fraction of desired product in the dry API stream. For continuous crystallization process this is defined as

$$PP = M_{API,6} / (M_{API,6} + M_{I,6}) \times 100 \quad (26)$$

An analogous expression is used for batch crystallization process. In this work, product purity is imposed as a constraint, with minimum target value set to 99.5%.

- Crystal yield (Y_c, %) is the mass fraction of API recovered in the product stream relative to that in the feed. For the continuous crystallization process, the yield:

$$Y_c = (M_{API,6} / M_{API,1}) \times 100 \quad (27)$$

For batch crystallization, the yield is defined over N_b batches and calculated as:

$$Y_c = \Sigma_{b=1}^{N_b} M_{API,8}^b / (N_b M_{API,1}^b) \quad (28)$$

Crystal yield is also constrained, with a minimum value of 90% imposed for both the systems.

CASE STUDIES

To understand the trade-offs between crystal yield and material consumption, a BOO-CAMPD formulation is constructed to simultaneously minimize the PMI and maximize the yield. It is solved using the ϵ-constraint method [16], reformulating a series of single level optimization problems, each maximizing crystal yield subject to a specified upper bound (ϵ_{PMI}) on PMI. The formulation is implemented in gPROMS Process, version 2023.1.0.

The optimization problem is solved using OAREAP, the implementation of the outer approximation algorithm [17] in gPROMS. To increase the likelihood of finding the global solution, the NLP primal subproblems are solved using a multi-start strategy, with 16 initial guesses, and the NLPMSO local solver.

Continuous Cooling Crystallization – Case 1

The continuous cooling crystallization of mefenamic acid (MA), with toluic acid (TA) as an impurity is

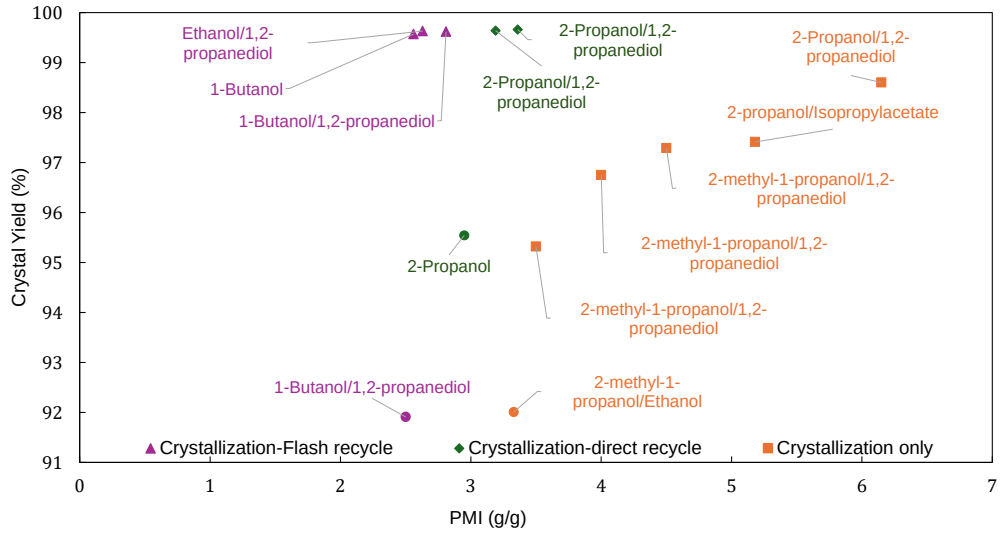

Figure 3: Pareto-optimal solutions for continuous crystallization of mefenamic acid, showing the optimal solvent/solvent pairs obtained for different cases.

investigated. Since all analyses are performed on relative basis, stream 1 is assigned a total molar flow of 1 mol/s, with an impurity to API weight ratio of $1:30$. The identity and composition of the binary solvent mixture, and the process operating conditions, are optimized in an integrated manner. The decision variables, their bounds, and the set of 15 candidate solvents are shown in Table 1. The candidates are selected based on the availability of the required SAFT-γ Mie interaction parameters [18]. The optimization problem is solved for three configurations: crystallization with direct recycle, crystallization with flash recycle, and stand-alone crystallization. For all cases, the minimum allowable PMI is fixed at 2.5 g/g, to avoid the formation of highly viscous slurries. Upper bound on PMI is varied to generate multiple instances of the ϵ-constraint formulation for maximizing crystal yield.

Table 1: Case 1: decision variables and bounds

Variables	Range
$y_{q,s}$	{0, 1}
F_2, F_3(moles/s)	[0, 1000]
T_4, T_7(K)	[290, 360]
T_{flash}	[290, 420]
α_p	[0.2, 0.9]
S	{water, acetic acid, cyclohexanone, acetone, ethanol, methanol, butanol, pentanol, 2-propanol, 1, 2-propanediol, 2-methyl-1-propanol, ethyl acetate, butyl acetate, isobutyl acetate, isopropyl acetate}

The ϵ_{PMI} values are varied between the PMI at the solution of the minimization of PMI solution and the PMI at maximum yield. For stand-alone crystallization (crystallization only in Figure 3) case, ϵ_{PMI} values are set at intervals of 0.5 between these bounds. For both recycling cases, ϵ_{PMI} values are set with an increment of 0.2. The Pareto-optimal solutions for the three configurations are compared in Figure 3, where the optimal solvent(s) corresponding to each solution are also reported along with the associated PMI and yield values. For certain solutions, only one solvent is reported as the optimal flow rate of the second solvent is zero, implying that a pure-solvent is optimal at those conditions. The circled points denote the solutions corresponding to the minimum PMI.

For most solutions, the temperature range attains its maximum value ($T_4 = 360$ K, $T_7 = 290$ K), reflecting the increased solubility range needed to enhance API recovery. Further, the purge ratio (α_p) is driven to lower bound of 0.2 corresponding to the maximum allowable recycle to the crystallizer.

As expected, increasing the PMI results in higher solvent usage and process yield. Incorporating a recycle stream retains API within the process and reduces the fresh solvent demand, thereby enhancing crystal yield without proportionately increasing PMI. This is seen as a narrower range of PMI values on the Pareto front and a decrease in the number of Pareto-optimal solutions. The range of PMI values is narrowest with the use of flash drum as this reduces the fresh solvent requirement.

Batch Cooling Crystallization – Case 2

Batch cooling crystallization of paracetamol with metacetamol and acetanilide as impurities is studied. The feed contains 16.5 g of paracetamol, corresponding to the laboratory scale to which this process is based, with an impurity to API weight ratio of $1:30$ for both impurities. The design objective is to determine the optimal solvent,

solvent mass, and process conditions that maximize the crystal yield while minimizing PMI. The decision variables and the variable bounds are detailed in Table 2. The details of the SAFT-γ Mie model used for paracetamol is given in the supplementary material.

The crystallization process is examined over $N_b = 20$ batches, with overall PMI constrained to a lower bound of 2.2 g/g. Multiple ϵ-constraint optimization instances are generated by varying the upper bound on PMI to maximize crystal yield. The resulting BOO-CAMPD optimal solutions are presented in Figure 4. Three Pareto-optimal solutions are shown as points with $b = 1, \ldots, 20$. Ethanol yields the minimum overall PMI, while isopropyl acetate yields the maximum crystal yield. The range of operating temperatures is maximized in all cases ($T_1 = 343$ K, $T_2 = 263$ K). Similar to the continuous crystallization case, a narrow Pareto front is observed for the batch crystallization process.

Table 2: Case 2: decision variables and bounds

Variables	Range
$y_{q,s}$	{0, 1}
$M_{s1,1}$ (g)	[0, 1000]
T_1, T_2(K)	[263, 343]
α_{evap}	[0, 0.8]
$\alpha_h, \alpha_w, \alpha_c$	0.25, 0.5, 0.1
S	{water, acetic acid, acetone, ethanol, methanol, butanol, pentanol, 2-propanol, 1, 2-propanediol, 2-methyl-1-propanol, ethyl acetate, butyl acetate, isobutyl acetate, isopropyl acetate}

The variation of PMI and crystal yield with the number of batches is also examined in Figure 4. As the Pareto-optimal solutions transition from minimum PMI (ethanol) to maximum yield (isopropyl acetate), the number of batches required to reach steady state decreases. In the case of isopropyl acetate, the changes in both PMI and yield with batch number are negligible. This behaviour arises because a large amount of solvent is used in the first batch, leading to maximum product recovery during crystallization. Consequently, very little material is available for recycle, thereby leading to minimal variation across batches. In other cases, the introduction of recycling leads to simultaneous improvement in crystal yield and PMI. Overall, the Pareto front moves closer to the utopia point as b increases.

CONCLUSIONS

We have presented a systematic methodology for the identification of optimal solvent/solvent mixtures, and process conditions for integrated crystallization and solvent/API recovery processes. A general CAMPD framework was formulated to design the crystallization processes and was applied to a continuous crystallization process for mefenamic acid and a batch crystallization process for paracetamol. Both processes incorporate solvent and API recovery and recycle. The process models were based on rigorous thermodynamics and the imposition of a number of thermodynamic constraints. The results have demonstrated the potential benefits of solvent recycling. This is reflected by a narrowing of the Pareto front upon the introduction of a recycle stream in the continuous case, indicating that improvements in crystal yield can be achieved with only marginal increases in PMI. In the batch case, this is seen as the Pareto front moves closer to the utopia after a small number of batches.

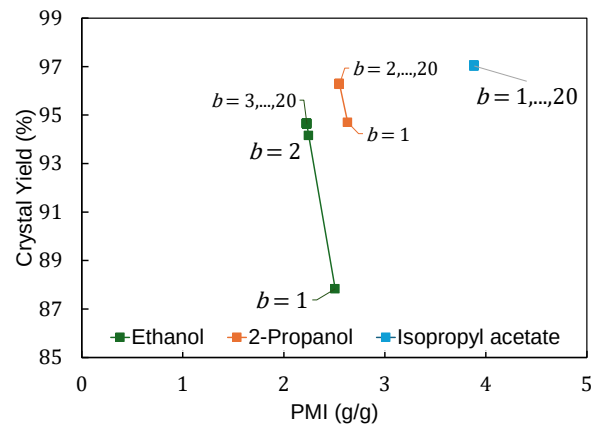

Figure 4: Pareto-optimal solutions obtained for paracetamol batch crystallization, showing the variation in PMI and crystal yield as a function of batch cycle b.

These results confirm that low-PMI operation can be realized while maintaining high levels of product recovery. Despite the sustainability advantages of recycling, the case studies also indicate that solvent selection plays a critical role and recycle may offer limited benefits for certain solvents. Integrating energy consumption into the multi-objective optimization framework will enable a more comprehensive sustainability assessment and facilitate the identification of solvent–process combinations that balance material efficiency with energy efficiency.

DIGITAL SUPPLEMENTARY MATERIAL

The supplementary material is available at LAPSE Id: **LAPSE:2026.0014**.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge funding from the UK Engineering and Physical Sciences Research Council (EPSRC) research grant (EP/W01923X/1 and EP/Z532964/1).

AUTHOR IDENTIFIERS

Author ORCIDs:
Seth G: 0009-0006-3466-9345
Naseri S: 0000-0003-2830-8624
Paliwal S: 0000-0002-1993-591X
Jackson G: 0000-0002-8029-8868
Galindo A: 0000-0002-4902-4156
Adjiman CS: 0000-0002-4573-7722

REFERENCES

1. Variankaval N, Cote AS, Doherty MF. From form to function: crystallization of active pharmaceutical ingredients. AIChE Journal 54:1682-1688 (2008). https://doi.org/10.1002/aic.11555
2. Watson OL, Jonuzaj S, McGinty J, Sefcik J, Galindo A, Jackson G, Adjiman CS. Computer aided design of solvent blends for hybrid cooling and antisolvent crystallization of active pharmaceutical ingredients. Org. Process Res. Dev. 25:1123-1142 (2021). https://doi.org/10.1021/acs.oprd.0c00516
3. Tung H, Paul EL, Midler M, McCauley JA. Crystallization of organic compounds. Wiley (2023). https://doi.org/10.1002/9781119879497
4. Gani, R., Jimenez-Gonzalez, C., ten Kate, A., Crafts, P.A., Jones, M., Powell, L., Atherton, J.H. and Cordiner, J.L., 2006. A modern approach to solvent selection: although chemists' and engineers' intuition is still important, powerful tools are becoming available to reduce the effort needed to select the right solvent. Chemical Engineering, 113(3), pp.30-44.
5. Wong SY, Tatusko AP, Trout BL, Myerson AS. Development of continuous crystallization processes using a single-stage mixed-suspension, mixed-product removal crystallizer with recycle. Crystal Growth & Design 12:5701-5707 (2012). https://doi.org/10.1021/cg301221q
6. Ferguson S, Ortner F, Quon J, Peeva L, Livingston A, Trout BL, Myerson AS. Use of continuous MSMPR crystallization with integrated nanofiltration membrane recycle for enhanced yield and purity in API crystallization. Crystal Growth & Design 14:617-627 (2013). https://doi.org/10.1021/cg401491y
7. Wang J, Lakerveld R. Integrated solvent and process design for continuous crystallization and solvent recycling using PC?SAFT. AIChE Journal 64:1205-1216 (2017). https://doi.org/10.1002/aic.15998
8. Karunanithi AT, Achenie LEK, Gani R. A computer-aided molecular design framework for crystallization solvent design. Chemical Engineering Science 61:1247-1260 (2006). https://doi.org/10.1016/j.ces.2005.08.031
9. Watson, O.L., Galindo, A., Jackson, G. and Adjiman, C.S., 2019, November. Computer-Aided Design of Solvent Blends for the Cooling and Anti-Solvent Crystallization of Active Pharmaceutical Ingredients. In 2019 AIChE Annual Meeting. AIChE.
10. Karunanithi AT, Acquah C, Achenie LEK, Sithambaram S, Suib SL. Solvent design for crystallization of carboxylic acids. Computers & Chemical Engineering 33:1014-1021 (2009). https://doi.org/10.1016/j.compchemeng.2008.11.003
11. Foli? M, Gani R, Jiménez-González C, Constable DJC. Systematic selection of green solvents for organic reacting systems. Chinese Journal of Chemical Engineering 16:376-383 (2008). https://doi.org/10.1016/s1004-9541(08)60092-0
12. Jonuzaj S, Akula PT, Kleniati P, Adjiman CS. The formulation of optimal mixtures with generalized disjunctive programming: a solvent design case study. AIChE Journal 62:1616-1633 (2016). https://doi.org/10.1002/aic.15122
13. Gross J, Sadowski G. Application of perturbation theory to a hard-chain reference fluid: an equation of state for square-well chains. Fluid Phase Equilibria 168:183-199 (2000). https://doi.org/10.1016/s0378-3812(00)00302-2
14. Haslam AJ, González-Pérez A, Di Lecce S, Khalit SH, Perdomo FA, Kournopoulos S, Kohns M, Lindeboom T, Wehbe M, Febra S, Jackson G, Adjiman CS, Galindo A. Expanding the applications of the saft-? mie group-contribution equation of state: prediction of thermodynamic properties and phase behavior of mixtures. J. Chem. Eng. Data 65:5862-5890 (2020). https://doi.org/10.1021/acs.jced.0c00746
15. Muhieddine MH, Viswanath SK, Armstrong A, Galindo A, Adjiman CS. Model-based solvent selection for the synthesis and crystallisation of pharmaceutical compounds. Chemical Engineering Science 264:118125 (2022). https://doi.org/10.1016/j.ces.2022.118125
16. Laumanns M, Thiele L, Zitzler E. An efficient, adaptive parameter variation scheme for metaheuristics based on the epsilon-constraint method. European Journal of Operational Research 169:932-942 (2006). https://doi.org/10.1016/j.ejor.2004.08.029
17. Viswanathan J, Grossmann IE. A combined penalty function and outer-approximation method for MINLP optimization. Computers & Chemical Engineering 14:769-782 (1990). https://doi.org/10.1016/0098-1354(90)87085-4
18. Febra SA, Bernet T, Mack C, McGinty J, Onyemelukwe II, Urwin SJ, Sefcik J, ter Horst JH,

Adjiman CS, Jackson G, Galindo A. Extending the saft-? mie approach to model benzoic acid, diphenylamine, and mefenamic acid: solubility prediction and experimental measurement. Fluid Phase Equilibria 540:113002 (2021). https://doi.org/10.1016/j.fluid.2021.113002

19. Scott A. Brown, Yusuf Khan, Nicola Voiculescu, Anna Jawor-Baczynska, Kirstie Milne, Jorge Calderon De Anda, Helen P. Wheatcroft, Christopher Burcham, Stanley Kolis, Jun Li, Jan Sefcik and Chris J. Price. Enhancing Product Yield, Purity & Sustainability by API & Solvent Recycling in Cooling Crystallisation. Under review.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Reinforcement Learning-driven Process Intensification Synthesis – Design and Optimization of Reaction/Separation Systems

Dylan Nice[a], Daniel Wenck Ribeiro[bc], Kristina Savitskaya[bc], Rahul Bindlish[d], Efstratios N. Pistikopoulos[bc], Yuhe Tian[a*]

[a] West Virginia University, Department of Chemical and Biomedical Engineering, Morgantown, West Virginia, United States
[b] Texas A&M University, Artie McFerrin Department of Chemical Engineering, College Station, Texas, United States
[c] Texas A&M University, Texas A&M Energy Institute, College Station, Texas, United States
[d] The Dow Chemical Company, Technical Expertise and Support Technology Center, Houston, Texas, United States
* Corresponding Author: yuhe.tian@mail.wvu.edu.

ABSTRACT

This work aims to systematically generate intensified process designs by integrating reinforcement learning (RL)-driven process synthesis and phenomena-based modeling via Generalized Modular Framework (GMF). Rather than considering flowsheet synthesis with conventional unit-operations, GMF utilizes fundamental building blocks, also known as mass and heat exchange modules, to describe the physiochemical phenomena and to enhance novel process discovery. At its core are driving forces which characterize the mass transfer feasibility based on the total change in Gibbs free energy of the system. RL is integrated with this phenomena-based modeling strategy to drive flowsheet generation by exploring much of the total action space and minimizing pre-postulation of stream connections. All possible inlets, outlets, and interconnections between modules are contained in a stream matrix. Deep Q-Network is used as the RL agent, which contains a multi-layer convolution neural network followed by a multi-layer feedforward neural network. This network computes q values for each possible stream connection contained in the stream matrix and determines an optimal action from the maximum q-value (e.g., varying the stream interconnections). New designs can thus be generated and iteratively improved. This approach is demonstrated using two case studies: (1) binary separation of benzene and toluene, and (2) membrane-assisted reaction for methanol production.

Keywords: Process Intensification, Machine Learning, Optimization, Process Synthesis, Process Design

1. INTRODUCTION

With the demand for sustainable chemical production rising, novel processes are more pertinent now than ever before. These technologies should simultaneously improve cost efficiency, energy consumption, and environmental outlook. An emerging process design technique known as process intensification is a promising innovation in the chemical sector [1-2]. It is capable of significant improvements versus conventional designs by intensifying physiochemical fundamentals to overcome equilibrium limitations, often leading to equipment featuring multifunctional tasks. Representative process intensification examples are reactive distillation, membrane reactors, dividing wall columns, etc. Despite these advantages, there are limited examples of intensified systems deployed in industry. This is due to the high complexity for the conceptual process design of these units and their integration with existing processes/plants [3].

Recent research has been investigating reinforcement learning (RL) as a tool to drive process synthesis to address computational complexity associated with traditional approaches using mixed-integer nonlinear programming (MINLP) [4-7]. The RL-based methods typically start with a pool of unit operations in which connections between each can be defined using an inlet-outlet stream matrix, graphical representation, etc. without explicit postulation of a superstructure. The RL agent will

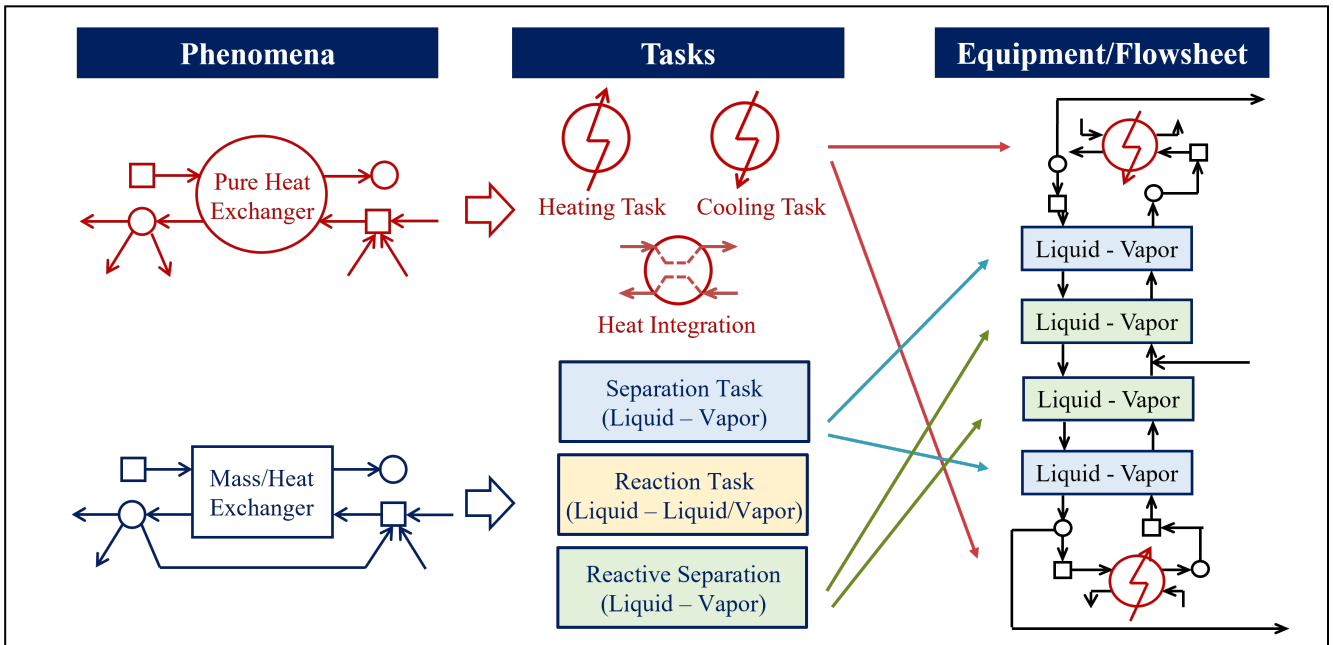

Figure 1. Phenomena-based modeling via GMF.

explore the design space and make decisions on e.g., changing unit operation selections or stream connections, thus generating new process designs.

This work aims to extend RL-driven process synthesis towards systematic process intensification. An integrated framework is presented which augments phenomena-based process modelling though the Generalized Modular Framework (GMF) and automated flowsheet generation through RL's exploration capabilities. The remainder of this paper is organized as follows: Section 2 introduces the RL-GMF methodology for process intensification synthesis. Section 3 showcases the proposed approach on two representative case studies including benzene/toluene binary separation and CO_2 to methanol production via membrane-assisted reaction. Section 4 concludes this paper with key findings and ongoing work.

2. RL-DRIVEN PROCESS INTENSIFICATION SYNTHESIS

2.1. Overview of Generalized Modular Framework

The Generalized Modular Framework is a phenomena-based modeling strategy which, rather than considering conventional chemical processing equipment, considers two fundamental building blocks describing general physiochemical phenomena [8]. As shown in Fig. 1, these phenomena building blocks include: (1) a pure heat exchange module (HE) and (2) a mass/heat exchange module (MH). The pure heat exchange module acts as a heat exchanger where the process stream can be heated, cooled, or utilized in heat integration. The mass/heat exchange module can represent any of the following tasks: separation, reaction, or reactive separation. These tasks can then be synergistically combined to model conventional/intensified units or flowsheets.

The core of GMF is its Gibbs-based driving forces which enforce the total Gibbs free energy of the system to always be decreasing or at zero for feasible mass/heat transfer to occur. While equilibrium is considered, it is not strictly enforced. Eqs. 1-3 present the driving force constraints for reaction/separation systems driven by vapor-liquid equilibrium (VLE), such as distillation-based systems [9]. For these systems, $G1$ defines the net mass transfer direction and $G2$ defines the mass transfer feasibility which is the distance from physical and chemical equilibrium. The product of $G1$ and $G2$ approximates the total change in Gibbs free energy of the system.

$$d(nG)^{tot}_{T,P} \approx G1_i \times G2_i \leq 0 \tag{1}$$

$$G1_i = dn^L_{ei} \approx f^{LO}_e x^{LO}_{ei} - f^{LI}_e x^{LI}_{ei} \tag{2}$$

$$G2_i = \left[\frac{\partial(nG)^{tot}}{d(n^L_{ei})}\right]_{T,P} \approx \ln\left[\frac{\gamma^L_{ei} x^L_{ei} P^{sat,L}_{ei}}{\phi^V_{ei} x^V_{ei} P_{tot}}\right] + \sum_i \sum_k \left[\frac{\upsilon_{ik}\Delta G^f_i}{RT^L_e} + \upsilon_{ik}\ln\left(\phi^V_{ei} x^V_{ei} P_e\right)\right]\frac{\partial \epsilon_k}{\partial n^L_{ei}} \tag{3}$$

Eqs. 4-5 give the driving force constraints for membrane-assisted reaction/separation systems [10]. μ^V_i in Eq. 4 is the chemical potential in the reaction zone and is determined from the mass transfer across the membrane calculated by Eq. 5. The product of the chemical potential and the change in material approximates the total change in Gibbs free energy for the reaction zone.

$$d(nG)^{tot}_{T,P} = \sum_i \mu^V_i dn^V_i \leq 0 \tag{4}$$

$$J_i = \lambda_i A(\mu_i^M - \mu_i^V) \quad (5)$$

The GMF process synthesis model is built with a physical model, structural model, and objective function. The physical model specifies a set of physical and/or chemical constraints of the system. These are broken down into: (1) mass balances, (2) energy balances, (3) thermodynamic properties, (4) phase defining, (5) driving forces, etc. Thermodynamic properties can be defined through any given thermodynamic model and phase defining to ensure that liquid and vapor streams are accurately represented. The structural model represents the existence (or not) of modules and streams along with the interconnectivity of the system, defined through binary variables. Mixers are placed before the module inlet stream, and splitters are placed after the module outlet stream. These allow for interconnecting streams between modules. The objective function can be defined to maximize or minimize any desired performance metric.

2.2. RL-GMF Methodology

The RL-driven process synthesis algorithm applied in this work is built on that developed by Wang et al. [5] for unit operation-based process design. Tian et al. [11] presented a follow-up work to accelerate the RL design search. This work introduces the RL-GMF process intensification synthesis strategy, as shown stepwise in Fig. 2.

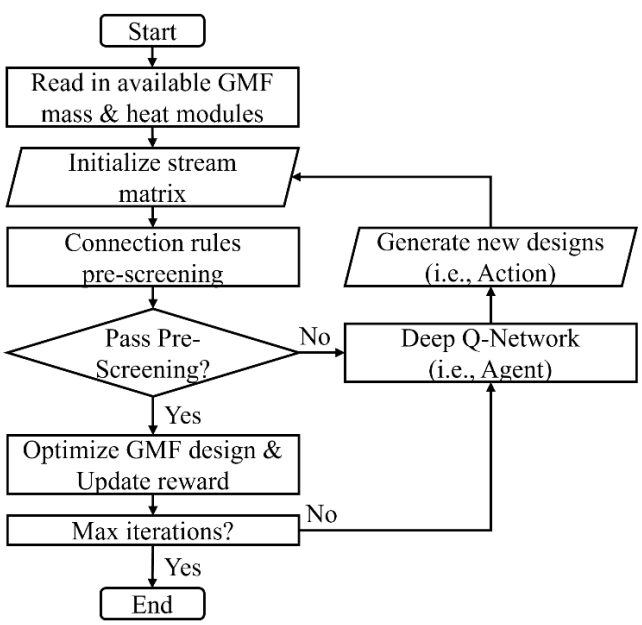

Figure 2. RL-GMF algorithm for process intensification synthesis.

2.2.1 Stream Matrix Initialization

The methodology starts off by reading in the available GMF MH and HE modules where the maximum number of available modules are user defined. Compared to our prior work for GMF process intensification synthesis using mixed-integer nonlinear optimization (MINLP) [8-9], the binary variables are redefined as parameters in RL and used to construct an inlet-outlet stream matrix. For example, the matrix for a two MH module system is shown below in Table 1. The rows correlate to inlets and the columns correlate to outlets. 18 binary parameters are present, generating a 6×3 matrix.

Table 1. Stream matrix representation: An example of 2 MH modules.

	Liquid Outlets		
Liquid Inlets	$y^{IP,L}$	y_1^{LP}	y_2^{LP}
	y_1^{IL}	$y_{1,1}^{LL}$	$y_{2,1}^{LL}$
	y_2^{IL}	$y_{1,2}^{LL}$	$y_{2,2}^{LL}$
	Vapor Outlets		
Vapor Inlets	$y^{IP,V}$	y_1^{VP}	y_2^{VP}
	y_1^{IV}	$y_{1,1}^{VV}$	$y_{2,1}^{VV}$
	y_2^{IV}	$y_{1,2}^{VV}$	$y_{2,2}^{VV}$

2.2.2 Connection Rules Pre-Screening

As per Fig. 2, after the stream matrix gets initialized, it is sent to a series of pre-screening checks to ensure that the constructed flowsheet satisfies the most basic requirements on connectivity. This helps to reduce computational time by completely skipping the GMF modelling of any nonsensical solution generated by RL. Examples of these checks include:

- There exists at least one feed stream connection to the system.
- There exists at least one product stream connection from the system.
- There is at least one liquid inlet and one vapor outlet for heater modules to exist.
- There is at least one vapor inlet and one liquid outlet for cooler modules to exist.
- Inlet and outlet streams from both sides of each module must exist to enforce module existence.

The last connection rule requires that there should be vapor and liquid inlets/outlets or permeate and retentate inlets/outlets on both sides for VLE-driven systems or membrane assisted reaction-driven systems respectively. Failure of any of these checks will assign either the minimum pre-score or incur a penalty to the pre-score, resulting in a failed pre-screening and negative reward.

2.2.3. GMF Optimization & Reward Update

Upon successfully passing pre-screening, the associated stream matrix will be translated to a GMF phenomena-based flowsheet for optimization. If there is no inlet or no outlet streams to/from a module, the module is treated as non-existing. A reward is generated based on the status of the optimizer and the value of the objective function. If the optimizer fails to converge (e.g., infeasible

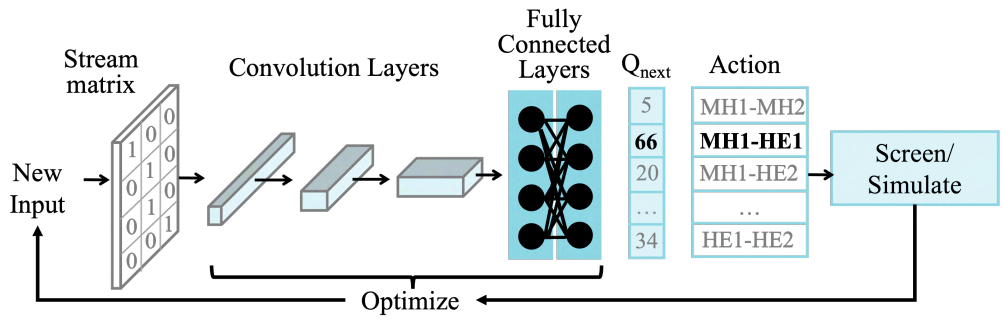

Figure 3. RL agent with Deep Q-Network.

solution), the minimum reward is assigned. If it successfully converges to a solution, a reward value is calculated based on e.g. objective value and process variables.

2.2.4. Deep Q-Network and Generation of New Design

Upon failing pre-screening or passing through GMF optimization, the current state and reward is sent to the DQN [5]. This step uses the agent in RL to learn from given conditions and intelligently make changes to the stream matrix to drive the algorithm towards improved designs. As illustrated in Fig. 3, the DQN is comprised of a three-layer one-dimensional convolution neural (CNN) network followed by a two-layer feed forward neural network (FNN). The state is a two-dimensional matrix of the batch number and flattened stream matrix (e.g., Table 1). This is used as the input to the CNN where key features are extracted, reshaped, and subsequently processed through the fully connected layers which determine the q-values for each row. An optimal decision is made where the index of the maximum q-value in each row takes a value of one. To open the possibility for every index in a single row to take a value of zero, an additional column is generated in this step to act as a placeholder for the optimal choice. These decisions translate to the switching of module inter-connections, elimination of modules, and adjustment of inlet/product streams. In this way, a new design can be generated with a unique inlet-outlet stream matrix and is then sent back to the workflow as shown in Fig. 2.

This integration of RL and GMF allows for automated flowsheet synthesis, scanning much of the total action space. Through the reduction of flowsheet pre-postulation, enhanced discovery of innovative process designs and non-intuitive flowsheets can be observed. Additionally, GMF is inherently a MINLP problem where solution convergence is time consuming and may be infeasible without good initial guesses. In this proposed approach, RL takes care of parameterizing the GMF binary variables, and the problem can then be solved as a nonlinear problem (NLP) with enhanced solution convergence. This framework is built in Python version 3.8.20 using TensorFlow for RL development, IPOPT as the NLP solver for Case Study 1 (VLE-driven GMF), and BARON as the NLP solver for Case Study 2 (membrane reaction-driven GMF).

3. CASE STUDIES

This section discusses two case studies performed to test the RL-GMF framework. These are representative of the two modelling trajectories of GMF with VLE-driven systems tested through the binary separation of benzene and toluene and membrane assisted reaction-driven systems tested with modular methanol production.

3.1. Binary Separation of Benzene/Toluene

3.1.1. Problem Statement

As a first motivating example, this case study describes a conventional separation scheme where a high purity product is to be obtained for at least 99% benzene from a liquid benzene/toluene mixture feed. Ideal vapor-liquid equilibrium conditions are used in the phenomena-based modelling with purely separation considered, neglecting the need for a reactive term in the driving force (Eq. 3). Table 2 shows the feed conditions of this scheme adopted from Proios and Pistikopoulos [12]. The GMF objective function is formulated to maximize the mole fraction of benzene from the column distillate (Eq. 6). The reward for RL is defined as the difference between the distillate rate and scaled reflux ratio (Eq. 7), which was the objective function used in [12].

$$\max\left(x_2^{VO,benzene}\right) \tag{6}$$

$$\text{reward} = f_{2,p_{liq}}^{CP} - 50\frac{f_{2,2}^{CL}}{f_{2,p_{liq}}^{CP}} \tag{7}$$

Table 2. Feed Stream Conditions for VLE-driven System

Liquid Feed	
Flowrate (kmol/hr)	360.00
Temperature (K)	359.89
Pressure (bar)	1.01
Benzene	0.70
Toluene	0.30

Proios and Pistikopoulos [12] solved this GMF process synthesis problem via MINLP. The resulting optimal design used two MH modules and two HE modules shown in Fig. 4. It featured a compact representation of a trayed distillation column. In what follows, we number the MH modules in ascending order. Namely, the MH module at the bottom is numbered as MH 1, the existence or not defined by binary parameter y_1. The MH module at the top is numbered as MH 2, the existence or not defined by binary parameter y_2. Each MH module has two associated HE modules. For the optimal MINLP solution shown in Fig. 4, the HE module for MH 1 acts as the reboiler. The HE module for MH 2 acts as the condenser. MH 1 defines the stripping section capturing separation below the feed stage. MH 2 defines the rectifying section capturing separation above the feed stage.

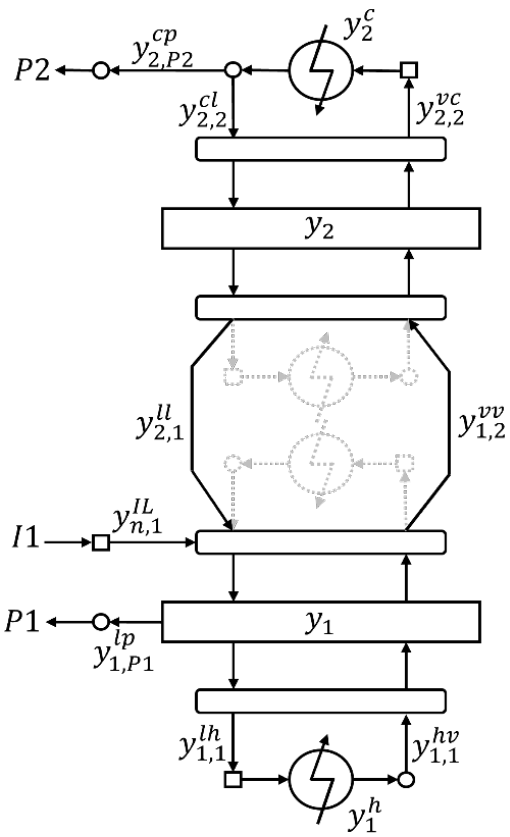

Figure 4. GMF representation of trayed distillation column (reproduced from [12]).

3.1.2. RL-GMF Implementation

The goal of this case study is to test if the RL-GMF methodology can successfully generate good design configuration(s) for this binary separation problem. A maximum of two mass heat exchange modules and four pure heat exchange modules are provided to the RL-GMF process synthesis. There are forty-two total binary variables associated with the design. All inlets, outlets, and interconnections are available for RL to explore resulting in a stream matrix of size 16×3. Note that for the stream matrix shown in Table 1, the current RL algorithm only allows one index in each row of the stream matrix to take a value of one. If it is possible that multiple indices in a single row would take the value of one, the stream matrix can be restructured.

Over 100, 000 episodes, eight unique feasible solutions were found as shown in Fig. 5. The total run time for 100, 000 episodes is 146 minutes. Out of this total run time, 139 minutes are spent on GMF optimization to evaluate the more than 2, 700 configurations which have passed the pre-screening step.

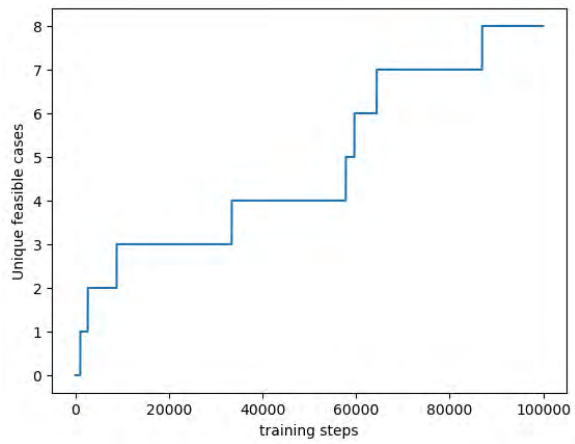

Figure 5. Benzene/toluene binary separation: Unique feasible solutions during RL-GMF training.

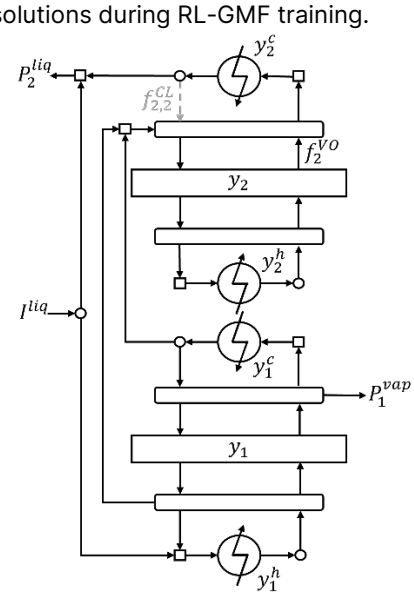

Figure 6. Benzene/toluene binary separation: RL-GMF design with the highest reward value.

The highest reward value achieved is 3120, the design of which is shown in Fig. 6. Using the trayed distillation column structure as base case (Fig. 4), its reward value corresponds to 2094. From optimization point of view, the RL algorithm successfully identifies a better

solution with non-intuitive design structure. However, from physical point of view, we notice that part of the initial feed stream (I^{liq}) is directly mixed with the product stream (P_2^{liq}). This happens as the problem statement only poses constraint for the purity of benzene vapor outlet from module 2 (f_2^{VO}). Additionally, the choice of the reward function (Eq. 3) strives to maximize the flowrate of the distillate product (P_2^{liq}) and minimize that of the reflux ($f_{2,2}^{CL}$), enforcing nonexistence of the reflux stream. This altogether drives the RL algorithm to identify Fig. 6 as the best solution.

The 2, 700 GMF configurations generated by RL do not include the base case configuration presented in Fig. 4. We hypothesize that this may be due to the large combinatorial design space with the forty-two total binary variables. To test this, a simplified case study is run with the HE module connections predefined and the MH connections open for RL to explore. This leads to a total of eighteen binary parameters. This simplified scenario is able to identify twelve unique feasible design structures over only 10, 000 episodes, and the base case configuration is generated on the 2, 308th episode. The increase from eighteen to forty-two parameters significantly adds to the design space (especially the infeasible design space), requiring longer run times to identify good design solutions.

3.2. Membrane-Assisted Methanol Production

3.2.1. Problem Statement

This case study investigates the design of a water selective membrane-assisted reaction scheme to produce methanol from CO_2 direct hydrogenation [13]. Kinetic models are adopted from Van den Bussche and Froment [14] using commercial catalyst $Cu/ZnO/Al_2O_3$. This kinetic model considers stoichiometric methanol production (Eq. 8) with reverse water gas shift (rWGS) as a side reaction (Eq. 9). The reactions take place in gas phase. An adiabatic reactor is considered with pressure drop negligible per the low flowrate. Feed conditions are listed in Table 3 at pilot scale. A total of ten membrane tubes are used with water permeance as 10^{-4} mol/s/m^2/Pa. Note that this permeance value is adopted higher than the reported values to more clearly showcase the impact of the membrane. A parametric study has been performed to quantify the impact of water permeance from 10^{-6} – 10^{-4} mol/s/m2/Pa but is not included here due to the page limit. The objective function for GMF optimization is to maximize CO_2 conversion. The reward function for RL is adopted as 10^4 times of the CO_2 conversion, to drive the design toward higher conversion.

$$CO_2 + 3H_2 \rightleftharpoons CH_3OH + H_2O \quad (8)$$

$$CO_2 + H_2 \rightleftharpoons CO + H_2O \quad (9)$$

Table 3. Feed conditions for methanol production.

	Feed	Sweep Gas
Flowrate (kmol/hr)	44.5	2.5
Temperature (K)	483.15	483.15
Pressure (bar)	77.9	1.00
CO_2	0.1850	0
H_2	0.7561	0
CH_3OH	0.0017	0
CO	0.0567	0
H_2O	0.0004	0
N_2	0	1

3.2.2. RL-GMF Implementation

A maximum of three MH modules are allowed for this study, which generates a stream matrix with thirty-two parameters of shape 8×4. The length of each MH module is set to 0.5 m. This model is run through the RL algorithm over 30, 000 episodes finding three unique feasible designs as shown in Fig. 7. The total computational time is 20 minutes.

Fig. 8(1-3) depict RL-GMF design solutions. Fig. 8(1) features the use of a single MH module with a CO_2 conversion of 9.55%. Fig. 8(2-3) use all three MH modules with CO_2 conversions as 26.51% and 26.34% respectively. The sweep gas recycle streams formed in Fig. 8(3) are not physically correct, but often created by the RL algorithm to obtain higher rewards numerically. This is consistent with the observation in our prior work for RL-driven unit operation-based process synthesis [11]. As such, Fig. 8(2) presents the best design solution identified by RL-GMF.

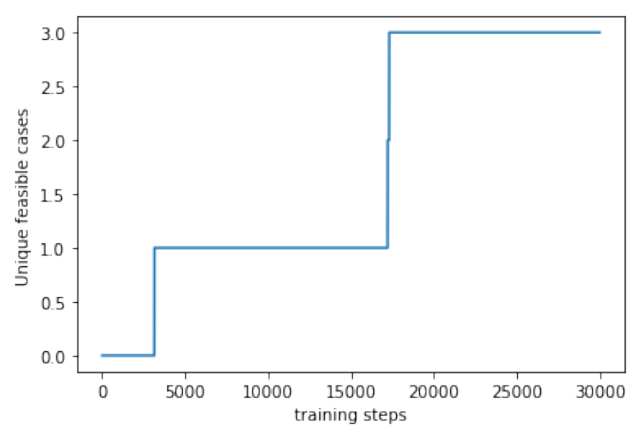

Figure 7. Unique feasible solutions during RL-GMF training – membrane assisted reaction-driven model.

The base case design is adopted as a counter-current membrane reactor as that shown in Fig. 8(4). The corresponding CO_2 conversion is 17.39%, which is lower than the 26.51% conversion provided by the design in Fig. 8(2). Table 4 compares the reaction rates and membrane permeance flux for the two designs. In the counter-flow

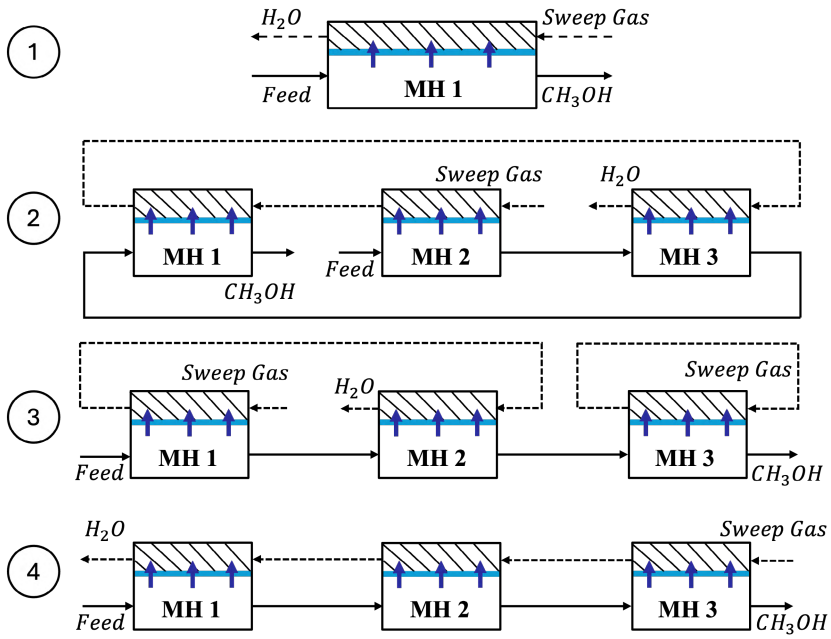

Figure 8. Membrane-assisted methanol production: (1-3) Feasible designs from RL-GMF, (4) base case design.

Table 4. Comparison of reaction rates and membrane permeance flux.

	Reaction Rate (Eq. 8) Unit: kmol/h/kg catalyst		**Membrane Permeance Flux (Water)** Unit: kmol/h/m^2	
	RL-GMF best design (Fig. 8(2))	Base case counter-current (Fig. 8(4))	RL-GMF best design (Fig. 8(2))	Base case counter-current (Fig. 8(4))
MH 1	0.80	1.89	0.55	0.84
MH 2	1.89	0.52	1.00	0.35
MH 3	2.43	0.94	1.09	0.63

structure in Fig. 8(4), MH 2 presents a bottleneck with small driving forces for both reaction and membrane separation. However, the more non-intuitive design structure in Fig. 8(2) successfully debottlenecks this by altering the sequences of feed and sweep gas flowing through the three MH modules.

4. CONCLUDING REMARKS

This paper introduced the use of reinforcement learning and phenomena-based modeling to enhance the discovery of novel process designs. Problem complexity was reduced from MINLP to NLP with the parameterization of GMF binary variables through the RL stream matrices. Two case studies were presented to demonstrate the RL-GMF algorithm, including benzene/toluene binary separation and membrane-assisted methanol production. For both case studies, the RL-GMF algorithm was able to identify non-intuitive process design solutions which provided higher rewards than the base case designs. The large infeasible design space still presented challenges for large-scale process synthesis. Future work should address this to reduce the run time and drive more efficient solution generation. Ongoing efforts are incorporating masking methods to RL and investigating classification neural networks to rule out the infeasible designs.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the financial support from NSF CBET-2401564, NSF EPSCoR RII Track-4 OIA-2327303, ACS PRF 67606-DNI9, West Virginia University, and Texas A&M Energy Institute.

REFERENCES

1. Tian Y, Demirel SE, Hasan MMF, Pistikopoulos EN. An overview of process systems engineering approaches for process intensification: state of the art. Chemical Engineering and Processing - Process Intensification 133:160-210 (2018). https://doi.org/10.1016/j.cep.2018.07.014
2. Ramírez-Márquez C, Al-Thubaiti MM, Martín M, El-

Halwagi MM, Ponce-Ortega JM. Processes intensification for sustainability: prospects and opportunities. Ind. Eng. Chem. Res. 62:2428-2443 (2023). https://doi.org/10.1021/acs.iecr.2c04305
3. Patel BA, Pereira CS. Process intensification at scale: an industrial perspective. Chemical Engineering and Processing - Process Intensification 181:109098 (2022). https://doi.org/10.1016/j.cep.2022.109098
4. Khan AA, Lapkin AA. Designing the process designer: hierarchical reinforcement learning for optimisation-based process design. Chemical Engineering and Processing - Process Intensification 180:108885 (2022). https://doi.org/10.1016/j.cep.2022.108885
5. Wang D, Bao J, Zamarripa-Perez MA, Paul B, Chen Y, Gao P, Ma T, Noring AA, Iyengar AKS, Schwartz DT, Eggleton EE, He Q, Liu A, Marina OA, Koeppel B, Xu Z. A coupled reinforcement learning and IDAES process modeling framework for automated conceptual design of energy and chemical systems. Energy Adv. 2:1735-1751 (2023). https://doi.org/10.1039/d3ya00310h
6. Gao Q, Schweidtmann AM. Deep reinforcement learning for process design: review and perspective. Current Opinion in Chemical Engineering 44:101012 (2024). https://doi.org/10.1016/j.coche.2024.101012
7. Reynoso?Donzelli S, Ricardez?Sandoval LA. A reinforcement learning approach with masked agents for chemical process flowsheet design. AIChE Journal 71: (2024). https://doi.org/10.1002/aic.18584
8. Pistikopoulos EN, Tian Y. Computer-aided modular process intensification: design, synthesis, and operability. Synthesis and Operability Strategies for Computer-Aided Modular Process Intensification :19-41 (2022). https://doi.org/10.1016/b978-0-32-385587-7.00011-7
9. Tian Y, Pistikopoulos EN. Synthesis of operable process intensification systems—steady-state design with safety and operability considerations. Ind. Eng. Chem. Res. 58:6049-6068 (2018). https://doi.org/10.1021/acs.iecr.8b04389
10. Tian, Y., Akintola, A., & Akoh, B. (2023). A Process Design, Intensification, and Modularization Approach for Membrane-Assisted Reaction Systems. *In Comput. Aided Chem. Eng.* (Vol. 52, pp. 3159-3164). Elsevier.
11. Tian Y, Akintola A, Jiang Y, Wang D, Bao J, Zamarripa MA, Paul B, Chen Y, Gao P, Noring A, Iyengar A, Liu A, Marina O, Koeppel B, Xu Z. Reinforcement learning-driven process design: a hydrodealkylation example. Systems and Control Transactions 3:387-393 (2024). https://doi.org/10.69997/sct.119603
12. Proios P, Pistikopoulos EN. Generalized modular framework for the representation and synthesis of complex distillation column sequences. Ind. Eng. Chem. Res. 44:4656-4675 (2005). https://doi.org/10.1021/ie040163m
13. Hamedi H, Brinkmann T, Shishatskiy S. Membrane-assisted methanol synthesis processes and the required permselectivity. Membranes 11:596 (2021). https://doi.org/10.3390/membranes11080596
14. Bussche KMV, Froment GF. A steady-state kinetic model for methanol synthesis and the water gas shift reaction on a commercial cu/zno/al2o3catalyst. Journal of Catalysis 161:1-10 (1996). https://doi.org/10.1006/jcat.1996.0156

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Graph Reinforcement Learning Framework for Batch Process Scheduling in State-Task Networks

Syu-Ning Johnn[a], Victor-Alexandru Darvariu[b], and Vassilis M. Charitopoulos[a]*

[a] Department of Chemical Engineering & The Sargent Centre for Process Systems Engineering, University College London, London, United Kingdom
[b] Oxford Robotics Institute, Department of Engineering Science, University of Oxford, Oxford, United Kingdom
* Corresponding Author: v.charitopoulos@ucl.ac.uk.

ABSTRACT

Batch production scheduling of resources to meet fluctuating product demand is a critical topic in the process industry. Existing optimisation approaches, based on heuristic and exact methods, trade off solution optimality and scalability to large problems. In this work, we investigate deep reinforcement learning as a powerful alternative in order to *learn* heuristics for batch scheduling. We formulate the batch scheduling problem as a Markov decision process operating on a state-task network representation encoded using graph neural networks, capturing relevant structural inductive biases. We propose a centralised training with decentralised execution architecture, in which agents placed on machines individually choose which tasks to complete using a global view of the network, cooperating towards task schedules that optimise the final production quantity. Preliminary results demonstrate that the proposed end-to-end framework learns to construct task schedules comparable to the optimal solution on small instances unseen during training, exhibiting strong potential for extension to more general graph structures and better scalability.

Keywords: Batch Process Scheduling, Reinforcement Learning, Markov Decision Process, Deep-Q Networks, Graph Neural Networks

INTRODUCTION

Batch scheduling is a critical functionality in the supply chain of process industries that deals with the allocation of limited resources over time in order to meet market demand for products that follow a batch recipe (Harjunkoski et al., 2010). As demand for diverse products continues to rise in the chemical industries, developing adaptive batch process operations constitutes a fundamental problem in these sectors (Yoo et al., 2021). Particularly, the scheduling of batch processes is a critical factor in process operations and plays a vital role in optimising production system performance (Wu and Maravelias, 2021).

While batch scheduling problems can be effectively tackled by existing exact and heuristic methods, learning-based AI methods, and in particular Reinforcement Learning (RL), provide a powerful alternative. RL is capable of solving complex, large-scale problems in dynamic environments and discovering complex scheduling policies beyond human intuition (Dogru et al., 2024; Johnn and Charitopoulos, 2025). RL algorithms such as Deep Q-Networks (DQN) and Proximal Policy Optimization (PPO) can be used to train an agent to learn the optimal decision-making policy through trial-and-error interactions with an environment that provides feedback. Given that a natural framing of batch scheduling is through State-Task Networks (STNs), neural architectures explicitly designed to operate on graphs are highly suitable. The combination of Graph Neural Network (GNN) architectures and RL algorithms can yield a powerful approach capable of solving complex graph optimisation problems that lack satisfactory traditional algorithms (Darvariu et al., 2024).

RL has demonstrated strong adaptability and substantial potential for improving industrial-level applications (del Real Torres et al., 2022). The integration of RL into process scheduling has recently garnered emerging attention in the literature. Hameed et al. (2023) developed a PPO framework with a GNN-based scheduler for production planning problems to minimise the production makespan. Zhang et al. (2023) introduced a DQN framework with multi-agent graphs for a job-shop scheduling

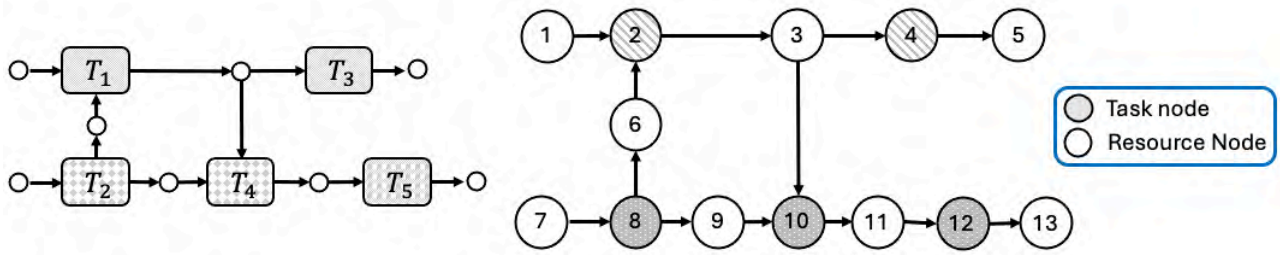

Figure 1. The classic STN model (left) can be converted into the directed graph structure (right) with nodes characterised by different features. Tasks TA1 and TA3, converted into nodes 2 and 4, are assigned to machine 1. Tasks TA2, TA4, TA5, converted to nodes 8, 10, 12 are controlled by machine 2. During the scheduling proces, the initial resources stored at node 1 and 7 are converted into intermediate nodes (3, 6, 9, 11) via different production processes, and eventually converted into final products as nodes 5 and 13.

process, where the job and machine-associated agents collaboratively work on job allocation, sequencing and routing. More recently, Rangel-Martinez and Ricardez-Sandoval (2025) proposed a PPO framework for multi-product batch scheduling processes, in which a hybrid agent handles multiple discrete and continuous decisions to build online schedules for STNs under uncertainty.

In this work, we propose an end-to-end DQN framework with a GNN encoder to address the batch process scheduling problem. We adopt the STN formalism by representing each production task as a node and each material flow as a link. The key differentiating characteristic of our work is the use of a multi-agent decomposition to manage the extensive decision space that arises when the problem is considered from a global perspective, enhancing scalability and computational tractability. Our approach obtains solutions close to optimality on problem instances unseen during training, significantly outperforms a single-agent baseline using joint actions, and offers strong potential for scalability and generalisability.

PROBLEM FORMULATION

Batch Process Scheduling Problem

STN graphs are a common way to visualise and model the batch scheduling problem (Méndez et al, . 2006). In this work, we reformulate classic STNs as directed graph structures, as shown in Figure 1. We preserve the graph topology and transform all task and resource components into indexed nodes distinguished by their node features (categories).

In this problem setting, we are given a set of machines that collaboratively schedule tasks to produce final products within a fixed time horizon. Each machine can be assigned more than one task over the time horizon. However, machines can process at most one task at any given time, which transforms the input resource(s) into output resource(s). When a task node is activated at a specific non-final time step, we assume it consumes exactly one unit of material from each precedent resource node at the start of this step, and produces exactly one unit of intermediate (or final) resource to each sequential resource node at the end of step. We assume only one type of material is stored at each resource node, and only one type of action is performed at a task node. A task node is available to process when all required input resource nodes have sufficient inventory levels and the time horizon has not yet been reached. Each task takes exactly one fixed time step to finish. Multiple machines collaboratively schedule task execution, with each machine processing at most one task at any given time to achieve maximum final production output.

Mixed-Integer Linear Programming Model

We adapt the STN formulation introduced by Kondili et al. (1993) and present a mixed-integer linear programming (MILP) model for the basic batch scheduling problem compatible with the directed graph representation given in Figure 1.

Let T denote the discrete set of time steps $\{0, 1, 2, \dots, T_{max}\}$. K is the set of tasks and L is the set of resource nodes, which can be partitioned into initial (L_{init}), intermediate (L_{inter}), and final product nodes (L_{final}). The set of machines is denoted M, where each machine $m \in M$ has a predefined set of tasks $K_m \subseteq K$ to execute. The binary variable $x_{k,t}$ takes the value 1 if task k is being executed at the start of time step t. The continuous variable $J_{l,t} \geq 0$ represents the inventory level of resource $l \in L$ at the start of time t. The objective is to maximise the total production of final products at the end of the time horizon. The MILP formulation is presented below:

$$\max \sum_{l \in L_{final}} J_{l,T_{max}} \tag{1}$$

$$\sum_{k \in K_m} x_{k,t} \leq 1, \forall m \in M, \forall t \in T \tag{2}$$

$$V_{l,t+1} = V_{l,t} + \sum_{k \in K_{BN}} P_{k,l} x_{k,t} - \sum_{k \in K_{FN}} C_{k,l} x_{k,t}, \forall l \in L \tag{3}$$

$$V_{l,0} = I_{l,0}, \forall k \in L_{init} \tag{4}$$

$$V_{l,t} \geq \sum_{k \in K_{FN}} C_{k,l} x_{k,t}, \forall l \in L \backslash L_{final} \quad (5)$$

$$V_{min} \leq V_{l,t} \leq V_{max}, \forall l \in L, \forall t \in T \quad (6)$$

Constraint (2) restricts each machine to choose no more than one task per step. Constraint (3) ensures flow conservation between task and resource nodes, enforcing and updating the inventory balance. Each resource node l is updated based on its precedent input task nodes producing resources and the subsequent output task nodes that consume it. Constraint (4) specifies the case in which all initial resource nodes start with fixed inventory levels prior to any task execution. Constraints (5) and (6) ensure that any task can only be selected if the required input resources are available and that storage limits are respected.

RL METHODOLOGY

Multi-agent Markov Decision Process

RL is a methodology in which an agent iteratively improves a decision-making policy through trial-and-error interactions with a dynamic environment (Sutton and Barto, 2018). The agent's objective is to map states to actions such as to maximise expected long-term rewards. Markov Decision Processes (MDP) provide the underlying model of the sequential decision-making task and can be solved using a variety of RL techniques.

An MDP is defined as a tuple (S, A, P, R), where a state $s_t \in S$ represents a complete environment snapshot at a specific time step. At each non-terminal state s_t, the agent chooses an action $a_t \in A$ from the set of available actions $A(s_t)$ and receives an immediate reward r_t specified by a reward function $R(s_t, a_t)$ for the selected state-action pair at step t. The subsequent state s_{t+1} is governed by a transition function P that captures movement from one state to another given an action.

An agent interacts with the environment through episodes, each of which represents a sequential chain of states, actions, and rewards over time. Reward feedback from the environment is incorporated iteratively to refine strategic decisions and train the agent to learn an optimal policy $\pi(a_t|s_t)$ that maximises cumulative future rewards. The state-action value function $Q(s_t, a_t)$ captures the expected reward associated with a given state-action pair under the optimal policy π. This quantity is also referred to as Q-value.

Multi-agent RL (MARL) extends classic RL by modelling how multiple agents interact within a shared environment and learn strategies to maximise their expected returns. In this setting, each agent's selected actions and the associated expected rewards can be influenced by the actions of all agents. The dynamic of the shared global state evolves by the joint actions of all agents following a Markovian transition function. Through this process, agents are trained to optimise individual or collective rewards, depending on the decision-making context being competitive, collaborative, or hybrid. Concretely, we consider a fully collaborative framing of batch scheduling, which is appropriate as all machines are under the control of a single organisation.

Deep Q-Learning and Graph Neural Networks

Q-learning is a class of RL algorithms that estimate the state-action value function (Q) using samples of agent experience. Estimates are iteratively refined using the update rule below, where the discount factor γ mediates the trade-off between immediate and future rewards:

$$Q^{new}(s,a) \leftarrow Q(s,a) + \alpha(r + \gamma \bullet \max_{a' \in A(s')} Q(s',a') - Q(s,a)) \quad (7)$$

Upon convergence, a greedy policy with respect to the learned Q-function is optimal.

Deep Q-learning (DQN) is a model-free RL approach that uses deep neural networks for scaling to large state and action spaces (Mnih et al. 2015). It incorporates a replay buffer and periodically updated target network. The state-action Q-value $Q(s,a)$ is estimated by the main network, while the maximum Q-value $Q(s', a')$ at the subsequent step is estimated by the target network. The target network parameters are periodically copied from the main network.

Graph Neural Networks (GNNs) are a learning representation designed to operate on graph-structured data (Velickovic, et al 2017). A GNN computes an embedding that captures node and edge features as well as underlying structural regularities. It aggregates information from neighbouring nodes through multiple message-passing layers to derive a higher-dimensional representation for the graph that exhibits desirable properties such as permutation invariance. GNNs can learn representations that achieve strong extrapolation capabilities to feature values and topologies unseen in the training set.

End-to-end DQN Framework

Our formulation of the batch scheduling problem operates on a directed graph $G = (V, E)$ with N_n nodes and N_e edges, where $V = K \cup L$. In each state, every machine $m \in M$ is responsible for a disjoint subset of task nodes K_m. The edges E connect resource nodes to task nodes, representing direct resource supply and production relationships.

We propose a multi-agent MDP framework as an alternative to using a single agent to control all machine nodes simultaneously in order to reduce the dimensionality of the problem. Instead, agents are placed at each machine and asked to individually act in a loop. The full

production schedule is established by the agents taking turns to select tasks given those chosen by the previous agents in the loop. This inner loop is executed once per time step t until the final time step T_{max} is reached.

We obtain a fully observable, finite-horizon, turn-based Markov game (Littman, 1994) with shared rewards. Therefore, we formulate the problem as the tuple $(S, A^{(i)}, P, R)$ containing an *agent-indexed* set of valid actions. The components are defined as follows:

- States: each state at time step t of an MDP episode is denoted $s_t^{(i)} = (G, X_t^{node}, X_t^{edge}, X_t^{global}, J_t, \tau_t)$, with i corresponding to index of the agent whose turn it is to act. $X_t^{node}, X_t^{edge}, X_t^{global}$ are node, edge, and global graph features respectively. Some features are updated as the episode progresses while others remain constant (see Table 2 for details). J_t denotes the set of inventory levels for each resource node at step t, and τ_t denotes whether an agent is the final decision-maker in the agent loop (i.e., if $i = |M|$).
- Actions: at time step t, agent i selects a task $a_t^{(i)}$ from the action space $A^{(i)}(s_t^{(i)})$ containing all production tasks available to it, along with an additional no-op action that allows the corresponding machine to stay idle for this step. Action masks are applied that prevent an agent from selecting tasks outside its allocated set K_m, and to dynamically filter out infeasible actions due to insufficient input resources.
- Transitions: once an agent has chosen an action $a_t^{(i)}$ that is not no-op, the selected task node is activated. The resource inputs from preceding resource nodes are converted into output resources stored at the sequential resource nodes, yielding J_{t+1}. The node, edge, and global features are also updated accordingly. If no-op is chosen, the machine status is set to "idle" for step t. Two types of transitions occur depending on the value of τ_t. If $\tau_t = 0$, meaning that this is not the last agent to act at global time step t, the turn passes to the next agent. If $\tau_t = 1$, the global time step advances to $t + 1$, and the turn transfers to the first agent. The terminal state $s_{T_{max}}$ is reached and the episode completes when either all time steps are exhausted or the remaining time is insufficient for the minimum batch to process. The terminal state captures the complete production schedule.
- Rewards: at the terminal state, the environment generates a reward based on the production schedule $R(s_t, a_t) = F(s_{T_{max}})$. As specified in (1), this corresponds to the total number of final products at the final time step. All intermediate rewards are 0.

Figure 1 provides a visualisation of a sample process schedule involving 2 machine agents, 5 task nodes, and 7 resource nodes. We treat each machine as an individual agent that controls the execution of task k at each time step t, representing the same decision as the binary variables $x_{k,t}$ in the MILP formulation.

A visualisation of the end-to-end framework is presented in Figure 2. To begin with, the state is represented as a graph and encoded by the GNN. The resulting feature vectors are passed to multilayer perceptrons (MLPs) representing the agents' Q-networks. These networks are iteratively updated based on the observed rewards to maximise the total amount of final products produced at the end of the time horizon.

For this work, we apply a centralised training, decentralised execution (CTDE) architecture (Amato, 2024). As mentioned, this is appropriate because all machine agents collaborate on the same graph topology to determine the schedule for a single problem instance. Given each agent will correspond to a different node representation based on the assigned task nodes and global state, parameterising each agent's policy separately is unnecessary. Instead of training N MLPs based on the same graph to yield N different agent policies $\pi_{i \in N}$, we therefore only train one global MLP that estimates the Q-values for the allowed actions. This is achieved by a mask layer after the MLP that is applied such that Q-values are only generated for the allocated task nodes, whereas the infeasible actions are filtered out by the mask. The CTDE framework is advantageous as only a single deep neural network is trained, yielding good computational efficiency and training stability, while still allowing for placing the neural network on every machine in an operational deployment setting.

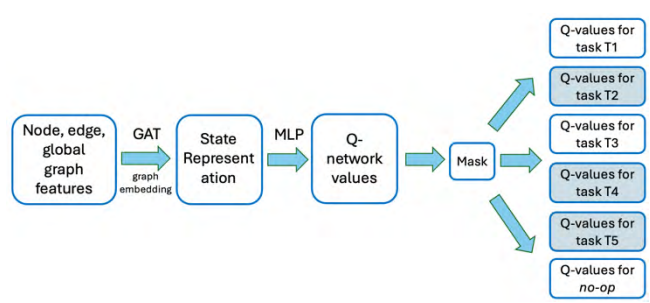

Figure 2. Visualisation of our end-to-end training framework with the CTDE architecture. All agents use the same global information and share the GNN encoder and MLP layers. A mask layer filters out disallowed tasks that are not allocated to this agent. In this example, corresponding to Figure 1, agent 1 is acting and therefore tasks T2, T4, and T5 are masked out.

Table 2: Node, edge and global features used as input for the graph representation of our end-to-end framework.

Name	Type	Re-scaled	Static over training	Feature description
agent_pinpoint	node	F	F	Indicator of the active agent at this node. Uses one-hot encoding: type 1 - agent is currently acting; type 2 - not acting; type 3 - resource node.
decision_status	node	F	F	Indicates the decision status of the agent associated with this node. Uses one-hot encoding: type 1 - currently deciding; type 2 - task nodes already decided; type 3 - task nodes not yet decided; type 4 - resource nodes.
node_type	node	F	T	Node category. Uses one-hot encoding: type 1 - task node, type 2 - resource node.
agent_allocation	node	T	T	Index of the machine agent ID assigned to a task node (1, 2, ...); 0 for resource nodes.
task_node_status	node	F	F	Task occupancy status. Uses one-hot encoding: type 1 - task occupied, type 2 - task idle, type 3 - resource nodes.
product_rate	node	F	T	Production rate associated with the node.
consume_rate	node	F	T	Consumption rate associated with the node.
inventory_level	node	T	F	Normalised inventory level for resource nodes; −1 for task nodes.
final_sales_price	node	T	T	Unit profit for each final resource node, 0 otherwise.
job_start_time	node	T	F	Job starting time for task node at this step; -1 for resource nodes.
job_end_time	node	T	F	Job completion time for task node at step; -1 for resource nodes.
current_time	global	T	F	Tracks the current state's time step for an episode.
remaining_time	global	T	F	Tracks the remaining time in this episode: $(T_{max} - T_{current}) / T_{max}$.
current_agent_id	global	T	F	Tracks the current agent id in one-hot encoding within the agent loop.
remaining_agent	global	T	F	Tracks the agent decision status sequence to count how many more agents should act within the agent loop.
task_size	global	T	T	Tracks the number of task nodes for this instance.
resource_size	global	T	T	Tracks the number of resource nodes for this instance.
physical_edges	edge	F	T	Stores node linkage showing consumption (resource to task) or production (task to resource).

To highlight the benefits of the proposed multi-agent decomposition, a single-agent baseline is also considered in which a DQN agent handles the scheduling of all production tasks. Instead of applying the proposed decomposition, the baseline uses a centralised architecture with a single agent selecting from a combinatorial joint action space as opposed to iterating M times over individual agents. Consequently, the replay buffer records the global state s_t only, together with the joint action tuple $\left(a_t^{(1)}, a_t^{(2)}, \ldots, a_t^{(m)}\right)$ generated by the single DQN agent. For the MLP output layer, rather than assign outputs to individual tasks (layer size $|K|+1$), the baseline model uses an output layer over every task combination of size $\prod_{i \in M}\left(|A^{(i)}|+1\right)$. In the illustrative example in Figure 1, this corresponds to an MLP output dimension of 6 for the multi-agent decomposition framework, and $3 \times 4 = 12$ for the single-agent baseline. A comparative performance analysis is presented in the next section.

CASE STUDY

Experiment Setting

We evaluate the proposed end-to-end framework on a range of process scheduling problems modelled as STN graphs, thereby demonstrating its effectiveness across varying problem parameters.

We adapt the batch scheduling problem instance given in Figure 1, with a fixed time horizon consisting of

10 steps for each episode. The graph topology remains fixed while the initial inventory level at each intermediate resource node is uniformly randomly generated from the interval $[0, 5]$. The initial resource nodes are unlimited and the final product quantities are set to zero.

Three distinct sets of training, validation and evaluation instances are generated. The training set is used to update the policy parameters during learning. The validation set is periodically used to quantify the performance of the policy specified by the learned parameters at a particular learning step. Finally, the evaluation set is used for reporting the quality of the best-performing learned model. The training, validation, and evaluation sets contain distinct problem instances.

As a GNN, we opt for the graph attention network (GAT) architecture (Velickovic et al., 2017). Parameters are updated by stochastic gradient descent using the Q-learning loss and the Adam optimiser. Hyperparameter values obtained as a result of hyperparameter tuning and other settings are specified in Table 1 below. The detailed list of GNN input features is given in Table 2.

We train our end-to-end DQN framework using an epsilon-greedy policy to ensure both exploration and exploitation. At evaluation time, a greedy policy is adopted, which always selects the action with the highest estimated Q-value.

We consider the following baselines. To derive an upper bound on performance, we use the MILP formulation in (1)-(6) to find the optimal solution for each problem configuration. The MILP model is implemented in Pyomo and solved with the Gurobi optimizer. As a lower bound, we use a uniform random policy that allows a machine to select any task with equal chance. This uses the same time horizon and also benefits from the action masking to filter out disallowed tasks. Lastly, we compare against the centralised single-agent method described in the previous section. This agent uses identical hyperparameters to the proposed method.

Table 1: end-to-end framework hyperparameter values.

Hyperparameter	**Value**
Learning rate	0.0005
Discount factor γ	1.0
Initial exploration epsilon	1.0
Final exploration epsilon	0.1
GAT attention heads	4
GAT number of layers	5
GAT hidden dimension	64
MLP 1st hidden layer dimension	256
MLP layer number	2
Sample batch size	4
Replay buffer size	20% of training
Validation frequency (episodes)	25

Experiment Results

We train the policy for 50, 000 episodes. Performance on the validation set containing 16 instances is checked every 25 episodes and the best-performing model is stored. Figure 3 shows the performance on the validation set during training. The best-performing policy is stored and evaluated on the separate evaluation set containing 64 instances.

We conduct a Monte Carlo simulation for the trained end-to-end DQN framework and compare the results with the three baseline methods. The lines shown in Figure 4 represent the average value accumulated over the evaluation instances. Notably, the single-agent baseline yields fewer production outputs and therefore less profit given the same number of training steps. Due to the sharp increase in action space size, we expect this effect to be even more pronounced on larger problem instances.

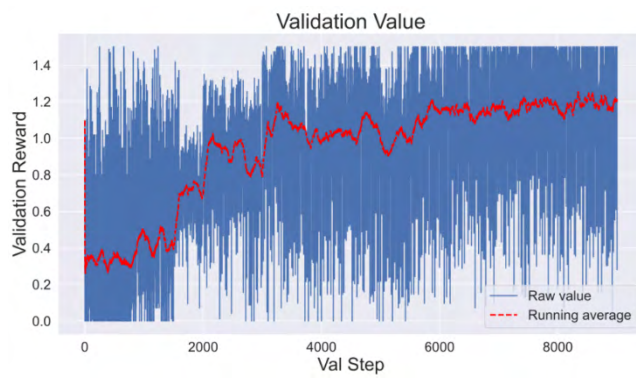

Figure 3. Validation performance during training based on the average reward computed from 32 predefined problem instances.

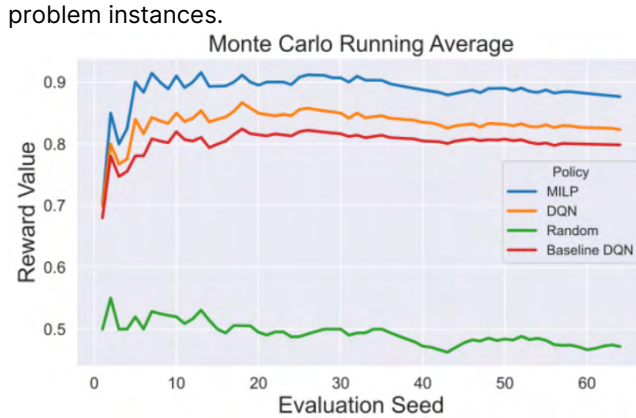

Figure 4. Average reward obtained on the evaluation set by the 3 considered methods. Averaged from 10 rounds.

From the box plot in Figure 5, we observe that our method constructs optimal solutions in 60.9% of cases and schedules with at most one task missing in 98.4% of cases. The random policy constructs solutions that, on average, only meet half of the production quantity compared to the optimal solution. Our framework leads to

solutions whose quality is clearly an improvement over the random baseline.

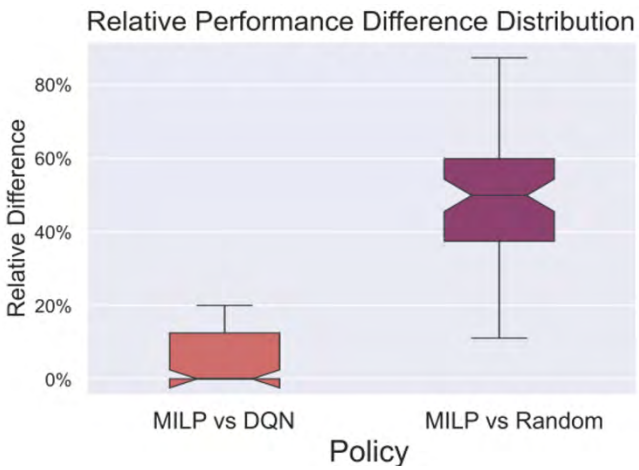

Figure 5. Evaluation performance comparing the MILP against the trained policy and a uniform random policy.

The Gantt charts in Figures 6 to 8 showcase the batch schedules constructed by the three methods using the problem instance on which the MILP scheduling method obtains the median performance among all problem instances.

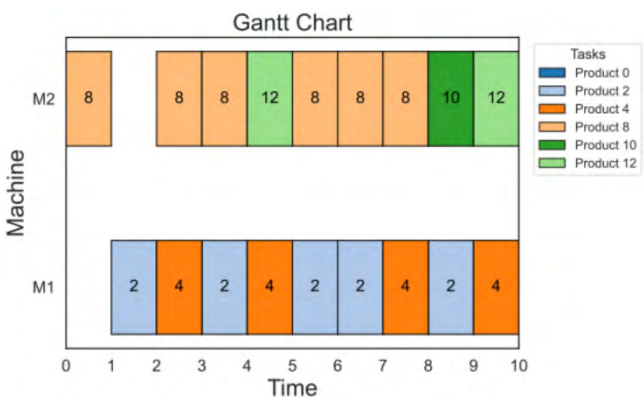

Figure 6. MILP scheduling plan on the median problem instance as a Gantt chart. Optimal solution.

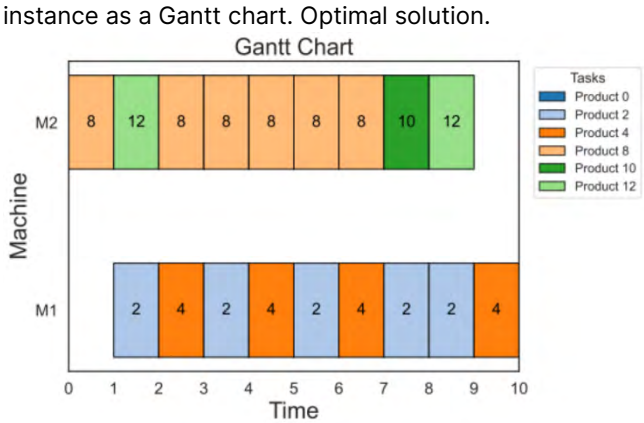

Figure 7. Scheduling plan output by our method on the median problem instance.

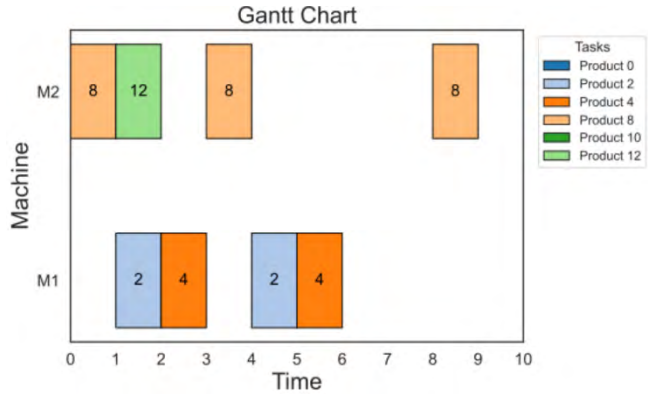

Figure 8. Scheduling plan output by uniform random action choices on the median problem instance.

CONCLUSION

In this work, we addressed the batch process scheduling problem using a deep reinforcement learning method. We formulated it as a Markov decision process and proposed an end-to-end DQN framework underpinned by a GNN encoder, in which multiple agents collaboratively control manufacturing machines to process tasks and construct work sequences to maximise final production quantity. Preliminary results on small-scale instances with up to 13 nodes demonstrate strong performance, with the trained agents achieving high-quality solutions in 98.4% of the evaluation scenarios and reaching optimality in 60.9% of scenarios.

In future research, we aim to leverage the generalisation ability of GNNs to train and evaluate our model on graph topologies with different sizes and characteristics. Our goal is to demonstrate the scalability of our framework to problem instances that are too large to be solved optimally using a MILP formulation but are amenable to a learning-based approach that can produce outputs quickly.

ACKNOWLEDGEMENTS

Financial support from the UK EPSRC grant EP/V051008/1 is gratefully acknowledged. V.-A. Darvariu acknowledges support from the Natural Environment Research Council (NERC) Twinning Capability for the Natural Environment (TWINE) Programme NE/Z503381/1 and the Innovate UK AutoInspect Grant 1004416.

AUTHOR IDENTIFIERS

Author ORCIDs:
Johnn S.-N.: 0000-0001-7958-2434
Darvariu V.-A.: 0000-0001-9250-8175
Charitopoulos V. M.: 0000-0001-9051-917X

REFERENCES

1. Amato, C., 2024. An introduction to centralized training for decentralized execution in cooperative multi-agent reinforcement learning. arXiv preprint arXiv:2409.03052.
2. Darvariu, V.-A., Hailes, S., & Musolesi, M. (2024). Graph reinforcement learning for combinatorial optimization: A survey and unifying perspective. *Transactions on Machine Learning Research (TMLR)*.
3. del Real Torres A, Andreiana DS, Ojeda Roldán Á, Hernández Bustos A, Acevedo Galicia LE. A review of deep reinforcement learning approaches for smart manufacturing in industry 4.0 and 5.0 framework. Applied Sciences 12:12377 (2022). https://doi.org/10.3390/app122312377
4. Dogru O, Xie J, Prakash O, Chiplunkar R, Soesanto J, Chen H, Velswamy K, Ibrahim F, Huang B. Reinforcement learning in process industries: review and perspective. IEEE/CAA J. Autom. Sinica 11:283-300 (2024). https://doi.org/10.1109/jas.2024.124227
5. Hameed MSA, Schwung A. Graph neural networks-based scheduler for production planning problems using reinforcement learning. Journal of Manufacturing Systems 69:91-102 (2023). https://doi.org/10.1016/j.jmsy.2023.06.005
6. Harjunkoski I, Maravelias CT, Bongers P, Castro PM, Engell S, Grossmann IE, Hooker J, Méndez C, Sand G, Wassick J. Scope for industrial applications of production scheduling models and solution methods. Computers & Chemical Engineering 62:161-193 (2014). https://doi.org/10.1016/j.compchemeng.2013.12.001
7. Johnn SN, Charitopoulos VM. A hybrid deep q-learning approach to online planning and rescheduling of single-stage multi-product continuous processes. Computers & Chemical Engineering 204:109415 (2026). https://doi.org/10.1016/j.compchemeng.2025.109415
8. Kondili E, Pantelides CC, Sargent RWH. A general algorithm for short-term scheduling of batch operations—i. MILP formulation. Computers & Chemical Engineering 17:211-227 (1993). https://doi.org/10.1016/0098-1354(93)80015-f
9. Méndez CA, Cerdá J, Grossmann IE, Harjunkoski I, Fahl M. State-of-the-art review of optimization methods for short-term scheduling of batch processes. Computers & Chemical Engineering 30:913-946 (2006). https://doi.org/10.1016/j.compchemeng.2006.02.008
10. Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, Petersen S, Beattie C, Sadik A, Antonoglou I, King H, Kumaran D, Wierstra D, Legg S, Hassabis D. Human-level control through deep reinforcement learning. Nature 518:529-533 (2015). https://doi.org/10.1038/nature14236
11. Rangel-Martinez D, Ricardez-Sandoval LA. Recurrent reinforcement learning strategy with a parameterized agent for online scheduling of a state task network under uncertainty. Ind. Eng. Chem. Res. 64:7126-7140 (2025). https://doi.org/10.1021/acs.iecr.4c04900
12. Sutton, R. and Barto, A., 2018. Reinforcement learning: An introduction. MIT Press.
13. Casanova A, Cucurull G, Drozdzal M, Romero A, Bengio Y. On the iterative refinement of densely connected representation levels for semantic segmentation. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) :1091-109109 (2018). https://doi.org/10.1109/cvprw.2018.00144
14. Littman ML. Markov games as a framework for multi-agent reinforcement learning. Machine Learning Proceedings 1994 :157-163 (1994). https://doi.org/10.1016/b978-1-55860-335-6.50027-1
15. Wu Y, Maravelias CT. A general framework and optimization models for the scheduling of continuous chemical processes. AIChE Journal 67: (2021). https://doi.org/10.1002/aic.17344
16. Yoo H, Byun HE, Han D, Lee JH. Reinforcement learning for batch process control: review and perspectives. Annual Reviews in Control 52:108-119 (2021). https://doi.org/10.1016/j.arcontrol.2021.10.006
17. Zhang JD, He Z, Chan WH, Chow CY. Deepmag: deep reinforcement learning with multi-agent graphs for flexible job shop scheduling. Knowledge-Based Systems 259:110083 (2023). https://doi.org/10.1016/j.knosys.2022.110083

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Two-stage stochastic programming optimization of natural gas pipeline network under cost and carbon emission reduction

Huiyu H, Bohong Wang*

[a] Zhejiang Ocean University, National & Local Joint Engineering Research Center of Harbor Oil & Gas Storage and Transportation Technology, Zhoushan, Zhejiang, China
* Corresponding Author: wangbh@zjou.edu.cn.

ABSTRACT

The pipeline transportation and distribution process of natural gas, from production sources to end consumers, can be divided into three stages: upstream production and supply, midstream storage and pipeline transportation, and downstream distribution and end-use. In the optimization of natural gas pipeline networks, considering the full life cycle, multiple uncertainties in planning, design, operation, and maintenance often affect the efficiency and quality of model optimization. This study addresses the uncertainty in end-user demand during the operation of natural gas pipeline net-works and investigates a scheduling method that simultaneously meets user demand while achieving coordinated optimization of operational costs and carbon emissions. Based on one year of historical demand data, a normal distribution is fitted to characterize the demand, and several representative scenarios with corresponding probabilities are extracted through clustering to capture demand uncertainty. A two-stage stochastic programming model is developed. In the first stage, the basic operational strategy of the pipeline network is determined, while in the second stage, flow, pressure, and compressor power are adjusted according to each demand scenario. The model prioritizes minimizing operational costs and employs the ε-constraint method to transform the carbon emission objective into a carbon emission cap. By systematically varying the value of ε, the Pareto frontier between cost and carbon emissions is delineated. Furthermore, nonlinear constraints are addressed using a combination of piecewise linearization and McCormick re-laxation. The proposed model is implemented in a real-world case study in China using Python and the Gurobi solver. Results demonstrate that, compared to conventional methods, the proposed approach effectively reduce carbon emissions by 34.96%, 42.35%, and 36.84% in high, medium, and low demand scenarios. The method provides clear optimization solutions and supports in-depth analysis of results, thereby enhancing data management capabilities of pipeline systems. The model systematically accounts for uncertainties in natural gas demand, offering valuable research support and practical references for scheduling decisions in the natural gas pipeline industry.

Keywords: Process Operations, Natural Gas, Stochastic Optimization, Uncertainty, ε-Constraint Method

INTRODUCTION

According to the latest statistics from the American Gas Association, natural gas has accounted for 25% of the global energy supply, particularly demonstrating significant application advantages in fields such as power generation, industrial manufacturing, and urban gas utilization [1].

As the "main artery" of the energy transmission system, the scientific nature of natural gas pipeline network planning and design, as well as the quality of operational efficiency, directly relate to national energy security, market stability, and people's livelihood protection [2]. he modern pipeline transportation system, by establishing a nationwide transmission and distribution network, has achieved efficient matching between resource supply

and market demand, providing stable and reliable gas supply assurance for various users including industrial and civil sectors [3]. Systematic pipeline network optimization design and intelligent dispatch management can reduce gas transmission losses and operational costs by up to 33.77% [4], significantly enhancing the market competitiveness of natural gas supply; it can also effectively reduce carbon emissions resulting from energy waste, thereby providing positive support for achievement of carbon peak and carbon neutrality goals.

Alinia Kashan and Molaei [5] confirmed through research that there is a direct correlation between carbon emissions and operating costs. Based on this, they incorporated carbon tax costs into the core composition of operating costs, achieving significant reductions in carbon emissions from pipeline network systems. Deng et al. [6], in their research on risk constraints in coupled power-gas networks, introduced a carbon tracking analysis algorithm that efficiently quantifies the carbon emission intensity at each node and the carbon flow in pipeline segments, achieving minimization of the total system emissions. Cheng et al. [7] proposed an intelligent method for intelligent low-carbon operation and maintenance of integrated energy systems. This method can accurately and in real-time perceive the system's operational status, optimization potential, and abnormal conditions, providing an effective technical pathway for natural gas pipeline networks to dynamically adjust operation and maintenance strategies based on supply and demand conditions. Based on this, the study further introduced intelligent algorithms such as the BP neural network, simultaneously achieving functions such as system energy output prediction, load forecasting, scheduling plan generation, and carbon emission optimization. Zhou et al. [8] proposed a hybrid-driven automatic response framework. When coupled with an optimization model, the resulting solution can reduce compressor energy consumption by 4.6% to 10.8% compared to actual operational schemes. Wen et al. [9], by formulating peak-shaving and flow distribution strategies, constructed a multi-objective optimization model to achieve collaborative scheduling optimization for multiple gas sources. The model considers both economic and environmental benefits, and its feasibility was verified through three case studies.

From the current research progress, the methods for low-carbon optimization of natural gas pipeline network operations can be mainly categorized into two types: one is to incorporate carbon emissions into the objective function for optimization, and the other is to treat carbon emissions as constraints for solving. Current related research mostly focuses on analyzing scenarios with certain user demands, while research on low-carbon optimization for large-scale uncertain demand scenarios is still insufficient. Existing multi-objective optimization studies often set carbon emissions as the objective function or simplified constraints, lacking a methodological system that can efficiently solve the Pareto optimal solutions for energy consumption and cost under different scales of carbon emission constraints, making it difficult to meet the low-carbon operation requirements in dynamic demand scenarios. In studies where carbon emissions are treated as the objective function, this method can only characterize the direct correlation between carbon emissions and operating costs, but lacks a detailed analysis of the dynamic coupling relationship between energy consumption and costs under different carbon emission constraints. Meanwhile, research on the coupled impact of compressor configuration and start-stop strategies on energy consumption and carbon emissions is relatively scarce. Additionally, there is a significant research gap in solving the problems of energy consumption optimization and carbon emission control under uncertain demand in large-scale natural gas pipeline network systems, and a unified methodology that can bridge the optimization research between certain and uncertain demand pipeline networks has not yet been formed.

This paper further constructs an optimization model for natural gas pipeline networks under uncertain demand scenarios. Given that the research object has characteristics such as a large number of stations, long pipeline segments, and complex topological structures, leading to optimization variables in the model that are diverse in type, numerous in nodes, and large in scale, this paper adopts a two-stage stochastic optimization model to quantitatively characterize the demand uncertainty in the pipeline network system. This model sets carbon emission indicators as rigid constraints and establishes multi-step gradient carbon emission constraint conditions. By solving the cost-minimizing optimization model constructed under different demand scenarios through optimization algorithms, it outputs the algorithm iteration processes and detailed optimization schemes corresponding to each scenario. Subsequently, combined with the pipeline flow distribution results under each scenario, the dynamic changes in pressure and flow along the pipeline network in the optimization scheme are verified to meet the system operation constraint requirements. By separately calculating the quantitative values of pipeline network energy consumption and carbon emissions for actual operational schemes and optimization schemes under different carbon emission constraint levels, the carbon emission-cost Pareto frontier curve corresponding to the optimization scheme is finally determined. Meanwhile, the formulation logic and core basis for compressor configuration and start-stop strategies under different demand scenarios, with energy consumption optimization as the core orientation and carbon emissions as the constraint boundary, are deeply analyzed. The optimization deduction process of the model

and implementable detailed implementation plans are fully presented.

MATHEMATICAL MODEL

Objective Function

The objective function of the two-stage stochastic programming consists of two parts: the first stage involves fixed costs related to decision variables, and the second stage involves the expected variable costs realized in response to different scenarios.

(1) First stage

Before uncertainty emerges, initial decisions (such as facility location and resource allocation) are formulated prior to the manifestation of uncertainty, requiring consideration of adaptability to multiple possible future scenarios.

$$C_{\text{first}} = \sum_{(i,j)\in \text{A}^{\text{C}}} Y_{i,j} * C_{\text{comp}} \tag{1}$$

where,

C_{first}—the fixed cost in the first stage, yuan;

$Y_{i,j}$—whether the compressor is on or off; when $Y_{i,j} = 1$, it indicates that the compressor remains on; when $Y_{i,j} = 0$, it indicates that the compressor remains off;

C_{comp}—the fixed cost per compressor, yuan.

(2) Second stage

The cost in the second stage mainly includes: the cost of multi-source gas supply calculated based on the actual demand at user nodes, and the energy consumption cost of electric-driven compressors calculated based on real-time electricity prices.

$$C_{\text{second}} = C_{\text{gas}} + C_{\text{com}} \tag{2}$$

$$C_{\text{gas}} = \sum_{(i,j)\epsilon \text{A}^{\text{P}}} Q_{i,j} * c_{\text{g}} \tag{3}$$

$$C_{\text{com}} = \sum_{(i,j)\epsilon \text{A}^{\text{C}}} N^{C}{}_{i,j} * c_{\text{c}} \tag{4}$$

where, C_{second}—the cost of the second stage, yuan.

The objective function of two-stage stochastic programming is the sum of the two-stage costs is Eq. (5).

$$\min f_2 = C_{\text{first}} + C_{\text{second}} \tag{5}$$

Constraints

The operational constraints include pipeline flow constraints, pipeline pressure constraints, compression ratio constraints, and compressor operation constraints. Due to the division of three representative scenarios during uncertainty modeling, stochastic variables are introduced in the modeling process.

Pipeline flow constraints

The flow on pipeline (*i*, *j*) must have upper and lower bounds:

$$Q_{i,j}^{\max} \le Q_{i,j,s} \le Q_{i,j}^{\max} \tag{6}$$

where, $Q_{i,j,s}$—pipeline flow from *i* to *j* under scenario *s*, m³/h.

Since supply and demand vary across scenarios, flow balance constraints must be set per scenario as follows:

$$\sum_{i|(i,j)\in \text{N}} Q_{j,i,s} + q_{i,s}^{\text{s}} = \sum_{j|(i,j)\in \text{N}} Q_{i,j,s}, i \in \text{N}_{\text{s}} \tag{7}$$

$$\sum_{i|(i,j)\in \text{N}} Q_{j,i,s} = \sum_{j|(i,j)\in \text{N}} Q_{i,j,s} + q_{i,s}^{\text{d}}, i \in \text{N}_{\text{d}} \tag{8}$$

$$\sum_{i|(i,j)\in \text{N}} Q_{j,i,s} = \sum_{j|(i,j)\in \text{N}} Q_{i,j,s}, i \in \text{N}_{\text{n}} \tag{9}$$

where, $q_{i,s}^{\text{s}}$—supply flow at node *i* under scenario *s*, m³/h;

$q_{i,s}^{\text{d}}$—demand flow at node *i* under scenario *s*, m³/h.

Node pressure constraints

$$p_i^{\min} \le p_{i,s} \le p_i^{\max} \tag{10}$$

where, $p_{i,s}$—pressure at node *i* under scenario *s*, MPa.

Compression ratio constraints

$$z_{\min} \le z_{i,s} \le z_{\max} \tag{11}$$

where, $z_{i,s}$—compressor pressure ratio at node *i* under scenario *s*.

Power constraints

$$N^{C}{}_{\min} \le N^{C}{}_{i,j,s} \le N^{C}{}_{\max} \tag{12}$$

where, $N^{C}{}_{i,j,s}$—power increment from node *i* to *j* under scenario *s*, kW.

Pipeline pressure drop constraints

$$p_{i,s}^2 - p_{j,s}^2 - \frac{\lambda \text{Z}_{\text{avg}} \rho \text{T}}{\text{C}_0^2 D_{i,j}^5} L_{i,j} Q_{i,j,s}^2 = 0 \tag{13}$$

Power calculation

$$N^{C}{}_{i,j,s} = \frac{\text{RTZ}_{\text{avg}}}{\eta} Q_{i,j,s} \frac{\upsilon}{\upsilon - 1}\left(z^{\frac{\upsilon}{\upsilon-1}} - 1\right) \tag{14}$$

$$N^{C}{}_{i,j,s} = C_1 \text{Q}_{i,j,s} w \tag{15}$$

Specific constraints and linearization processes can be found in reference [1].

Solver Configuration

The two-stage stochastic optimization model is implemented using Python programming and solved with the Gurobi solver in a Python environment. Key solution parameters are configured as follows: numerical precision focus mode is set to 2, thread count is 20, and node file initial memory is 0.5 GB.

Value for ε-Constraints

Based on the model optimization results, the ε-constraint method is further applied to impose carbon emission restrictions of varying intensities. Using the gradient

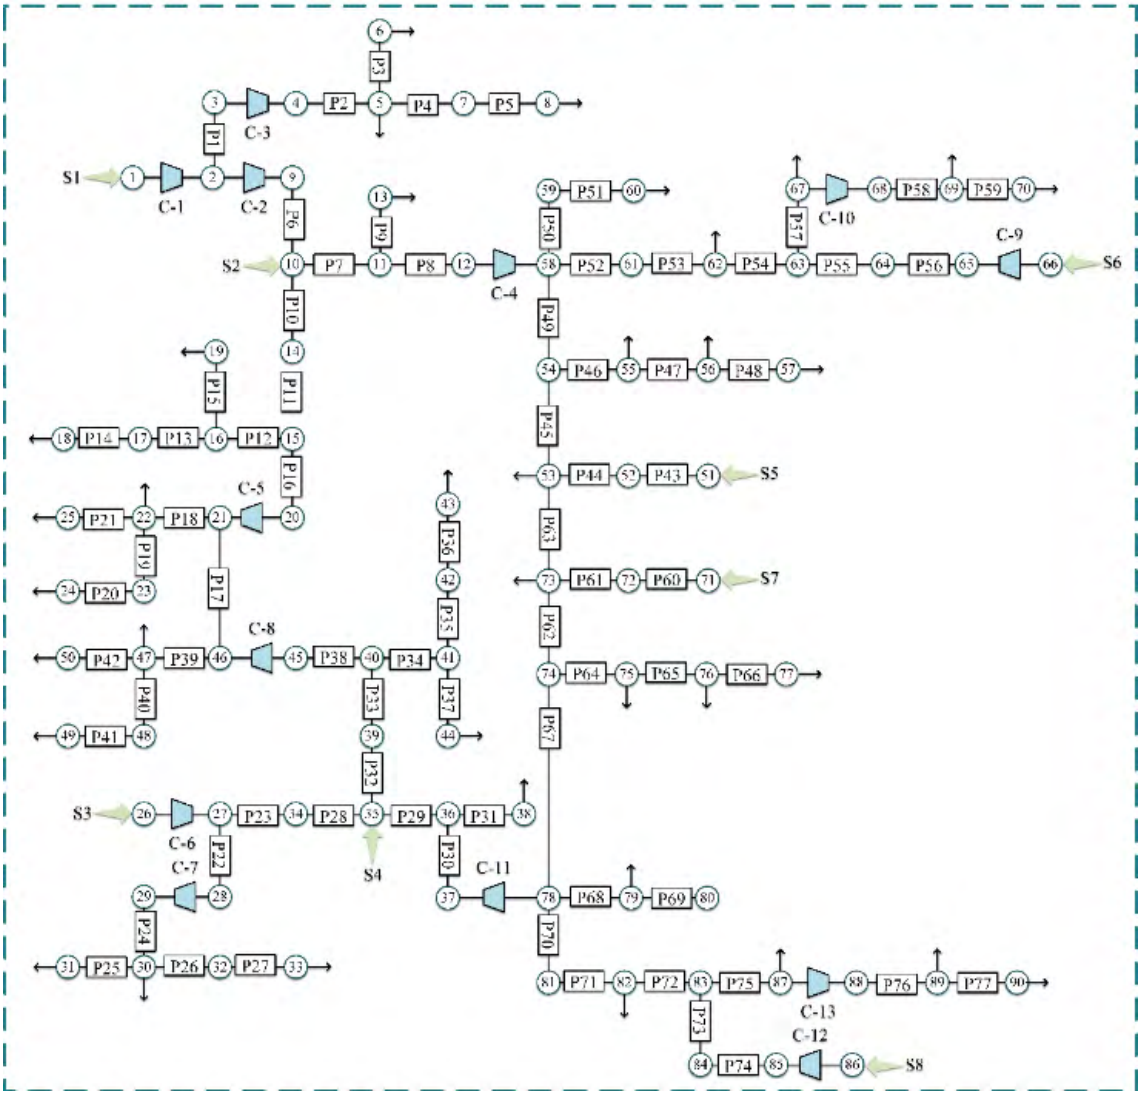

Figure 1. Natural gas pipeline network topology

method, ε values are divided into four constraint levels, covering a spectrum from loose to strict carbon emission constraint ranges. Combined with case study solution analysis, on the basis of obtaining the actual optimal energy consumption values, for different scenarios, upper and lower constraint intervals with an interval of approximately 2000 kg are set around the optimal carbon emission values to systematically investigate the impact of carbon constraints on operational strategies.

CASE STUDY

The case study involves a natural gas pipeline network comprising 8 gas sources (S1–S8), 90 nodes, 77 pipelines, and 13 compressors, as shown in Figure 1. The upper limit of node pressure is set at 2.9 MPa, and the lower limit at 0.8 MPa; the maximum transmission capacity for both pipelines and compressors is 566, 340 m^3/h. Based on natural gas production and consumption coefficients in northern China, the gas supply capacity of each source and the flow distribution in the pipeline network were calculated.

Based on the processing of one-year data, probability distribution fitting was performed to analyze its statistical characteristics. As shown in Figure 2, the annual natural gas demand distribution indicates that the actual demand distribution does not fully conform to a normal distribution but exhibits a certain degree of right-skew, meaning the frequency of higher demand days is higher than theoretical expectations. This distribution characteristic suggests the following: (1) During winter or extreme weather conditions, natural gas consumption significantly increases due to heightened heating demand; (2) Holidays, seasonal changes, or specific events may

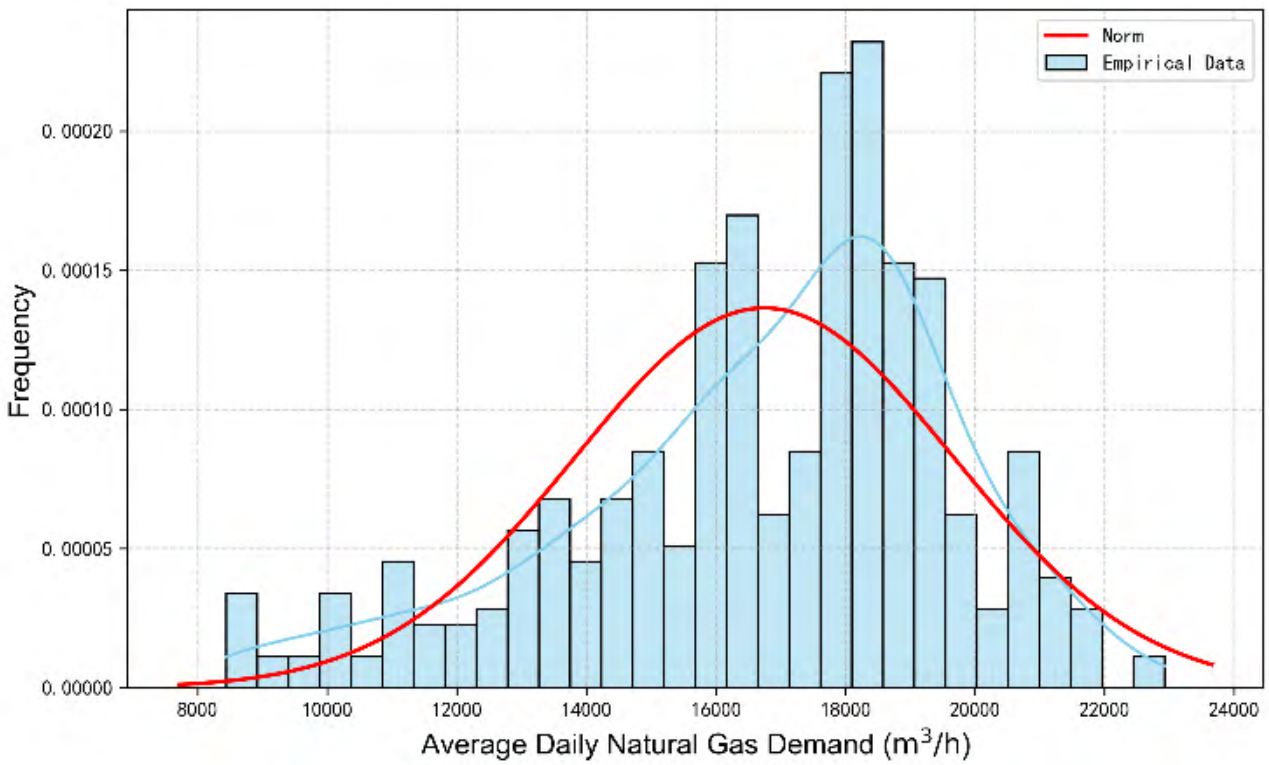

Figure 2. Natural gas demand distribution throughout the year.

lead to periodic fluctuations in demand; (3) In energy system planning, full consideration must be given to "tail risks" in such distributions to ensure supply security and system stability.

RESULTS AND DISCUSSION

All obtained results meet the requirement of an optimality gap of 0. The three scenarios were solved within 30 seconds, with the longest single-solution time being 26.74 seconds.

Cost-Carbon Emission Analysis

When determining the carbon emission factor values, this paper analyzes regional disparities based on the provincial grid emission factor data released by the Ministry of Ecology and Environment and the National Bureau of Statistics. Significant differences exist among provinces, with higher emission factors in North China and relatively lower ones in South China. To ensure consistency and comparability of the model, this paper ultimately adopts the national average electricity emission factor, i.e., 0.608 $kgCO_2/kWh$, as a unified calculation parameter.

Through analysis and testing of multiple sets of experimental data, the optimization algorithm obtained a series of optimal solutions. To evaluate the trade-off between economic and environmental objectives, this study employs the Pareto frontier analysis method and utilizes the ε-constraint method to transform the multi-objective optimization problem of carbon emissions and costs into a single-objective optimization problem with cost as the primary objective. Ultimately, Pareto frontier curves under different demand scenarios, as shown in Figure 3, were plotted.

The results indicate that the system exhibits a set of optimal Pareto solutions, including the carbon emission minimization solution. The Pareto frontier curve generally shows a downward trend from left to right, indicating a negative correlation between cost and carbon emissions. This relationship is primarily driven by compressor start-stop strategies: as carbon emission constraints tighten, the cost of adjusting compressor operating states and flow distribution to reduce emissions gradually increases, leading to a rising marginal cost of carbon reduction. The environmental optimal solution shows that to achieve emission reduction targets, the system needs to increase the utilization ratio of compressors in the natural gas pipeline network, resulting in higher operating costs. The rings in Figure 3 illustrate the relative proportion of optimal cost values under different scenarios. It is clearly observable that higher user demand corresponds to higher associated costs.

Analysis of Optimization Strategies under Different Carbon Emission Constraints

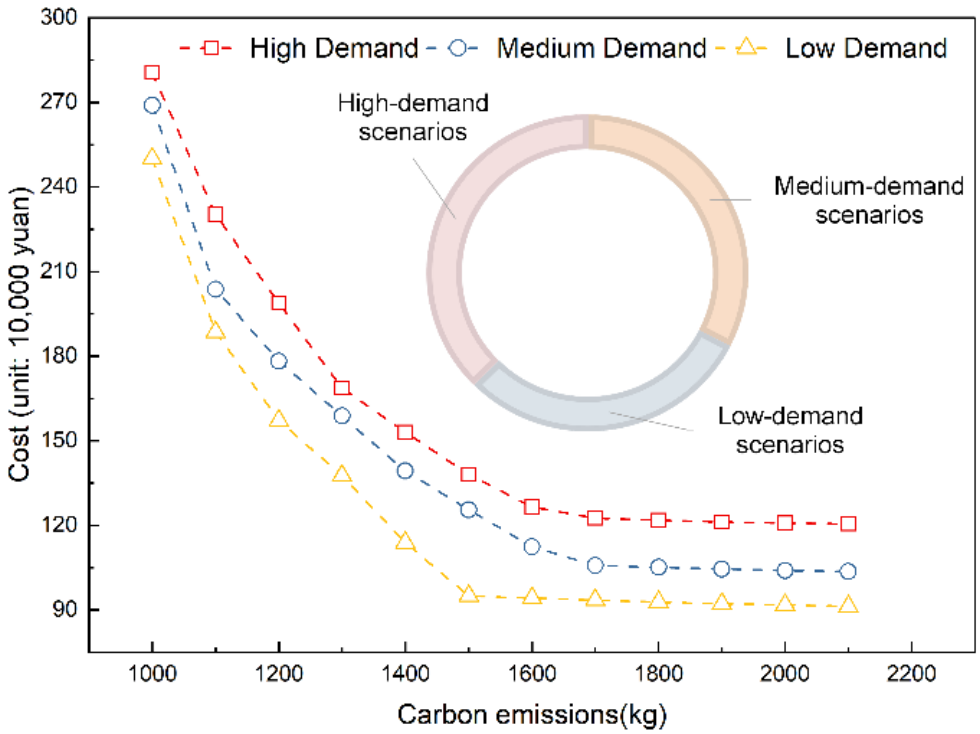

Figure 3. Pareto optimal curves for different scenarios.

Under different scenario conditions, by setting differentiated carbon emission constraint levels, the natural gas pipeline transmission system exhibits distinct compressor scheduling strategies and carbon emission constraint response characteristics during optimized operation when coordinating transmission load and multi-objective trade-offs as user demand varies from high to low. Details are as follows:

(1) Scenario 0: In high-demand scenarios, the natural gas transmission system faces high load and flow requirements. Here, the optimization model must balance large-flow transmission tasks, economic cost minimization objectives, and carbon emission constraints. It optimizes compressor start-stop strategies based on multi-unit collaborative operation. When carbon emission constraints are loose, the model prioritizes transmission capacity assurance, activating a sufficient number of compressors (including some inefficient units) to maintain system supply-demand balance, prioritizing flow supply in multi-objective conflicts while striving for economic cost minimization. Under moderate carbon emission constraints, the number of activated compressors is maintained at the minimum level required for flow balance. The system reduces excess carbon emissions and costs from inefficient operation by excluding inefficient and redundant units while ensuring transmission capacity. When constraints are strict, compressors are precisely matched to actual demand, avoiding redundant unit activation to reduce operational energy consumption. This satisfies carbon emission targets and node flow balance while lowering long-term system energy costs.

(2) Scenario 1: In medium-demand scenarios, system transmission loads are relatively moderate with smoother flow fluctuations. Compressor start-stop strategies here emphasize real-time optimization and dynamic adjustment to achieve dual objectives of economic cost minimization and carbon emission control. When constraints are loose (large ε values), compressors typically maintain basic flow balance and steady operation without activating reserve units. Only during significant load fluctuations are partial compressors fine-tuned for start-stop, ensuring stable pipeline operation at lower costs with naturally met carbon emission levels. Under moderate constraint strength, the system dynamically adjusts compressor operating states according to load fluctuations, achieving synergy between economy and carbon reduction while avoiding unnecessary energy consumption and maintenance costs. When constraints are strict (small ε values), compressors activate only high-efficiency units with strong power regulation capabilities, prioritizing minimum carbon emissions to achieve comprehensive optimization of system energy consumption and economic costs under flow balance and constraint conditions.

(3) Scenario 2: In low-demand scenarios, pipeline system transmission loads are small, and the core

challenge is avoiding excess energy consumption. Compressor start-stop decisions focus on streamlining operating units and reducing invalid start-stop to lower energy consumption and operating costs. When carbon emission constraints are loose, the system activates only efficient compressors to meet basic transmission needs under the goal of minimizing operating costs, indirectly reducing emissions through lower energy consumption. Under moderate carbon emission constraints, the system further controls emissions based on energy conservation, achieving cost-carbon emission dual-objective synergy through optimized unit start-stop strategies. When constraints are strict, the system adopts a mode combining high-efficiency compressors with intermittent operation, avoiding invalid emissions under low loads and continuously reducing energy consumption and related costs through precise compressor state regulation.

Analysis of Optimization Strategies under Different Carbon Emission Constraints

Finally, during the model solving process for each scenario, Pareto optimal solutions are screened from the solution set through non-dominated analysis. The following table summarizes the Pareto solution sets composed of system operation target costs and corresponding carbon emissions under each scenario based on the ε-constraint method and the two-stage stochastic programming approach, corresponding to high-demand, medium-demand, and low-demand scenarios, shown in Table 1. It reflects the actual carbon emissions, energy consumption, and total cost results under different constraint intensities.

Table 1: Impact of carbon emission constraints on cost in high-demand scenarios

Scenario	Total Cost (kCNY)	Carbon Emissions (kg)	Energy Consumption (MW)	Carbon Emission Reduction (%)
High Demand	1232.8	1640.91	112.4	-34.96%
Medium Demand	1066.5	1419.49	97.3	-42.35%
Low Demand	979.2	1303.29	89.3	-36.84%

Under high, medium, and low demand scenarios, compared to the unconstrained baseline scenario, the model optimization achieves carbon emission reduction rates of 34.96%, 42.35%, and 36.84% respectively. The results demonstrate that this optimization model is significantly effective in coordinating carbon emission control and economic cost optimization for large-scale natural gas pipeline networks, providing feasible scheduling strategies for low-carbon operation under varying demand intensities.

CONCLUSIONS

Addressing the uncertainty characteristics of user demand, this paper completes typical scenario division based on the probability distribution characteristics of annual demand data. Focusing on the operational optimization of natural gas pipeline systems, this study first constructs an optimization model minimizing energy consumption, then builds a two-stage stochastic programming model to solve for stochastic variables under different demand scenarios. The model selects pipeline flow and compressor station start-up quantity as core decision variables, incorporating multiple constraints including pipeline strength, pressure, flow balance, and compressor power/speed. During model solving, the ε-constraint method is employed to convert carbon emission indicators into constraints, achieving synergistic optimization of energy consumption and carbon emissions. The solution results show significant heterogeneity in the reduction of optimal system energy consumption under differentiated demand scenarios and low-carbon constraints. Pareto frontier analysis reveals that carbon emissions in high, medium, and low demand scenarios are reduced by 34.96%, 42.35%, and 36.84% respectively compared to actual levels.

ACKNOWLEDGEMENTS

This work is supported by the National Natural Science Foundation of China (52302420).

REFERENCES

1. American Gas Association: Statistical Review of World Energy, https://www.aga.org/research-policy/resource-library/statistical-review-of-world-energy/, (2025)
2. Yu W, Huang W, Wen Y, Li Y, Liu H, Wen K, Gong J, Lu Y. An integrated gas supply reliability evaluation method of the large-scale and complex natural gas pipeline network based on demand-side analysis. Reliability Engineering & System Safety 212:107651 (2021). https://doi.org/10.1016/j.ress.2021.107651
3. Research on reliability assessment methods for large gas pipeline network. Oil & Gas Storage and Transportation. 7–15 (2016)
4. Peng J, Zhou J, Liang G, Li C, Qin C. Multi-period integrated scheduling optimization of complex natural gas pipeline network system with underground gas storage to ensure economic and

environmental benefits. Energy 302:131837 (2024). https://doi.org/10.1016/j.energy.2024.131837
5. Alinia Kashani AH, Molaei R. Techno-economical and environmental optimization of natural gas network operation. Chemical Engineering Research and Design 92:2106-2122 (2014). https://doi.org/10.1016/j.cherd.2014.02.006
6. Deng Z, Zhang Y, Guo F, Pan X, Zhu Z. Towards a low-carbon and low-risk community: risk-constrained optimization of energy pricing in electricity and gas networks with flexible carbon-responsive loads. Journal of Cleaner Production 437:140534 (2024). https://doi.org/10.1016/j.jclepro.2023.140534
7. Cheng S, Nie T, Hui Q, Lu G, Teng Y, Chen Z. Integrated optimization method for HIES energy and carbon based on hybrid neural network. International Journal of Electrical Power & Energy Systems 172:111237 (2025). https://doi.org/10.1016/j.ijepes.2025.111237
8. Zhou J, Qin C, Fu T, Liu S, Liang G, Li C, Hong B. Automatic response framework for large complex natural gas pipeline operation optimization based on data-mechanism hybrid-driven. Energy 307:132610 (2024). https://doi.org/10.1016/j.energy.2024.132610
9. Wen K, Qiao D, Nie C, Lu Y, Wen F, Zhang J, Miao Q, Gong J, Li C, Hong B. Multi-period supply and demand balance of large-scale and complex natural gas pipeline network: economy and environment. Energy 264:126104 (2023). https://doi.org/10.1016/j.energy.2022.126104

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimization-based design of distillation processes with embedded pressure drop and HETP correlations

Sina Bertram[a], Jonas Schnurr[a], and Mirko Skiborowski[a]*

[a] Hamburg University of Technology, Institute of Process Systems Engineering, Hamburg, Germany
* Corresponding Author: mirko.skiborowski@tuhh.de.

ABSTRACT

To improve the energy efficiency of distillation processes, various process intensification concepts have been proposed, including direct heat integration and thermal coupling. Identifying the most suitable alternative for a given separation task requires a rigorous and consistent techno-economic optimization. Superstructure models typically rely on isobaric operation and fixed HETP values, in order to avoid treating column hydraulics when solving the already challenging mixed-integer nonlinear optimization problems. In order to overcome this limitation and evaluate the effect of the simplification, the current work extends a rigorous equilibrium-stage superstructure model to account for tray-specific pressure drop and HETP values. A polylithic solution approach is implemented to improve the convergence for the resulting optimization problems. The proposed approach is demonstrated for the optimization of heat-integrated distillation sequences operated at close to atmospheric and vacuum conditions, enabling a closer investigation of the impact of the classically applied simplifications. As the results illustrate that the overall energy demand and total annualized costs are only marginally affected for the considered wide-boiling mixtures the evaluation of competing process configurations will not be affected by the simplifications. However, considerable changes in column height and heat exchanger areas are observed, particularly under vacuum operation, such that column hydraulics should be considered in design optimization, in case column size restrictions or other considerations require a more accurate equipment sizing.

Keywords: Distillation, Energy integration, Superstructure, Optimization, Pressure drop

INTRODUCTION

Distillation remains the most widely applied separation technology in the chemical and process industries, accounting for approximately 90-95% of all fluid separations tasks [1] and a substantial share of the overall industrial energy demand. Despite its high energy consumption, direct substitution of distillation by alternative separation technologies is largely impractical due to its broad applicability, wide throughput range and capability to produce multiple high-purity products within a single unit [2]. In the context of rising energy prices and increasingly stringent environmental regulations, the development of energy-efficient distillation processes is therefore of central importance.

Numerous energy integration concepts have been proposed and demonstrated as effective means to reduce energy demand [3]. One prominent approach is direct heat integration between columns, in which the heat of condensation of the top vapor from one column is used to supply the heat for evaporating part of the bottoms stream of another column. This is achieved by adjusting the operating pressure of at least one column such that the top temperature of the heat-supplying column exceeds the bottom temperature of the heat-receiving column [4]. For two columns with comparable heat duties, direct heat integration has the potential to reduce energy consumption by up to 50% [4].

Another option to reduce the energy requirement of distillation processes is thermal coupling in which heat exchangers and liquid streams between adjacent distillation columns are replaced by bidirectional liquid and vapor transfer. Equipment-integrated dividing wall column (DWC) implement thermally coupled columns in a single shell and have been shown to reduce capital cost by approximately 30% compared to conventional, non-

integrated distillation sequences [5].

The identification of the best performing concept requires a systematic evaluation of all available process configurations. However, performing a detailed and rigorous analysis of all alternatives with rigorous models is usually intractable, such that shortcut models are the preferred choice for an initial screening [6]. Even the separation of nonideal mixtures can be addressed effectively when making use of thermodynamically sound pinch-based methods [7]. Yet this screening provides only an estimate of the performance, mandating that the most promising alternatives are further evaluated by more rigorous models and a techno-economic assessment.

The optimal design of distillation processes can be determined either through simulation-based approaches or via rigorous numerical optimization [8]. Simulation-based approaches can readily account for column hydraulics, as the column design, especially the number of stages in each section, are fixed prior to simulation. However, typically applied sequential optimization methods are rather cumbersome [5, 9], while gradient-free optimization using surrogate models with response surface methodologies [10] become computationally expensive when design spaces are large.

Mathematical programming applied to superstructure models posed as mixed-integer nonlinear programming (MINLP) problems enable a simultaneous design of the column structure and operating conditions building on gradient-based optimization [5, 8, 11]. While this approach allows for proven (local) optimality, the reported models usually build on isobaric conditions and constant height equivalent of a theoretical plate (HETP) values. Incorporating column hydraulics for pressure drop is a challenge as the actual number of stages and column dimensions are decision variables. Pressure variations resulting from friction losses along the column height affect the boiling temperature, relative volatilities, and ultimately separation performance [12]. Additionally, the pressure drop can increase the heating and cooling demands, and influence the relationship between the number of stages and energy demand, which may become nonlinear at higher stage numbers [13]. In direct heat-integrated distillation sequences, changes in pressure directly impact the reboiler temperature and, consequently, the available temperature driving force in the integrated heat exchanger, thereby significantly influencing feasibility and efficiency. For DWCs, the pressure drop in the parallel sections impacts the vapor split below the dividing wall and thus potentially impairs operation at minimum energy demand.

Although the common simplifications improve the numerical solution of the resulting MINLP problems, they may impair the identification of feasible/optimal designs and may bias the comparison of different process alternatives. While some recent studies have incorporated pressure drop correlations into rigorous distillation models [14] or superstructure optimization of retrofit problems with fixed column dimensions [15], the direct consideration of pressure drop and HETP correlations in superstructure-based design optimization with unknown column dimensions has not been addressed.

Therefore, the current work presents a proof of concept by extending an existing superstructure-based design optimization [16-19] by incorporating correlations for pressure drop and HETP as functions of the vapor load, expressed via the F-factor. By accounting for these effects at the conceptual design stage, a more realistic evaluation of energy-integrated distillation processes is enabled. Furthermore, an investigation of the impact of these correlations on the optimal design is investigated for representative case studies on the design of a heat integrated direct sequence for the separation of two wide boiling mixtures at vacuum and close to atmospheric conditions.

METHODOLOGY

In this work, a rigorous techno-economic optimization is preformed using an equilibrium-stage superstructure model based on the MESH equations and implemented in the General Algebraic Modeling System (GAMS). The resulting MINLP problem is solved as a series of successively relaxed NLP problems [16]. The following subsections describe the superstructure model for individual distillation columns, its extension to directly heat-integrated sequences, as well as the modifications for the implementation of pressure drop and HETP correlations.

Superstructure model for individual distillation columns

The superstructure model for an individual distillation column and its implementation is illustrated in **Figure 1.** It builds on the initial work of Kraemer et al. [16] and Skiborowski et al. [17], which was inspired by the work of Bartttfeld et al. [11]. The sizing of the stripping and rectifying sections is enabled by variable locations for the feed, reflux and boil-up streams. Stages above the reflux and below the boil-up tray are effectively bypassed and do not contribute to the separation. They are consequently not considered for equipment sizing and cost calculations.

Each equilibrium stage is described by the MESH equations, comprising mass balances, vapor-liquid equilibrium relations, summation constraints, and enthalpy balances. The thermodynamic properties and vapor-liquid equilibrium calculations are modeled via external equations implemented via a dynamic link library (DLL), enabling a reduced-space formulation that significantly decreases the number of variables and equations

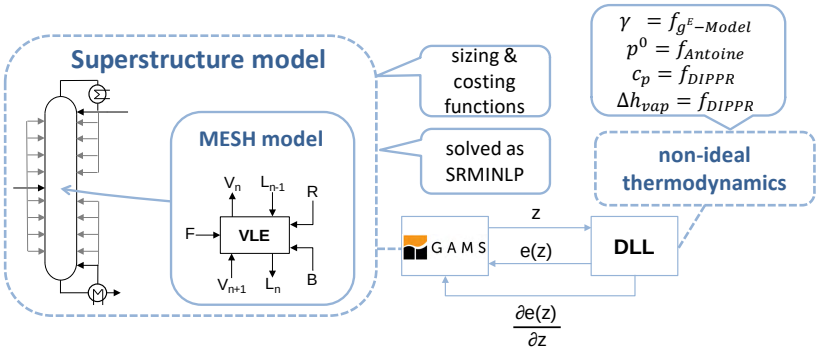

Figure 1. Superstructure model of an individual distillation column and implementation in GAMS with external equations for VLE computations.

managed by GAMS [17].

The optimization objective is to minimize the total annualized cost (TAC). Utility prices are derived from Turton et al. [20], while equipment costs are estimated using Guthrie's method, as outlined by Biegler et al. [21].

Modeling of direct heat-integrated sequences

The model of the heat-integrated direct sequence is based on the work of Waltermann and Skiborowski [18] with the adaptions described by Adami et al. [19]. The operating pressure of the second column is determined from a pre-defined approach temperature and flash calculations for the bottom stream of the first column and top vapor of the second column. Depending on the individual heat requirements either an additional condenser for the second column or an additional reboiler for the first column is implemented to compensate for a mismatch, while an intermediate heat exchanger provides the necessary heat to reach saturation conditions for the connecting stream at the increased pressure.

Pressure drop and HETP correlations

While the columns were considered isobaric in all preceding work on optimization-based design, the current work employs pressure drop and HETP correlations as functions of the vapor load, represented by the F-factor

$$FF_n = u_n^{superficial} \cdot \sqrt{\rho_n^{vapor}} = \frac{4 \cdot \dot{V}_n}{\pi \cdot D_{Col}^2} \cdot \sqrt{\frac{R \cdot T_n \cdot \sum_i y_{i,n} \cdot M_i}{p_n}}, \quad (1)$$

following the retrofit design approach described by Horsch et al. [15]. The column diameter is determined by the maximum vapor load across all active stages, excluding condenser and reboiler stages, following Douglas correlation [22]

$$D_{Col} \geq \sqrt{\frac{4 \cdot \dot{V}_n}{\pi \cdot FF_n} \cdot \sqrt{\frac{R \cdot T_n \cdot \sum_i y_{i,n} \cdot M_i}{p_n}}}, n = 2, \dots, n_{max} - 1. \quad (2)$$

The actual column diameter is considered to be the largest required stage diameter, ensuring that the maximum allowable F-factor is not exceeded

$$FF_n \leq FF_{max}, n = 2, \dots, n_{max} - 1. \quad (3)$$

The pressure drop per height and the HETP of the individual stages are estimated using polynomial correlations, which are fitted to vendor data for Sulzer BX structured packing derived at 960 mbar [23]:

$$\left(\frac{\Delta p}{h}\right)_n = \left(1 + b_{R,n} - \sum_{m=n+1}^{n_{max}} b_{R,m} + \sum_{m=1}^{n-1} b_{B,m}\right) \cdot (0.3617 \cdot FF_n^3 - 0.6838 \cdot FF_n^2 + 0.7572 \cdot FF_n) \quad (4)$$

$$HETP_n = \left(1 - \sum_{m=n+1}^{n_{max}} b_{R,m} + \sum_{m=1}^{n-1} b_{B,m}\right) \cdot (0.0120 \cdot FF_n^3 - 0.0666 \cdot FF_n^2 + 0.1657 \cdot FF_n + 0.0872) \quad (5)$$

Stages that are effectively bypassed by either the reflux or the boil-up stream do not contribute to the separation performance, as either the liquid or vapor stream is zero, such that the mass- and energy balance enforce the same for the liquid or vapor stream leaving the stage [16]. In order to maintain this property and implement the stage-specific pressure drop and HETP, the values of the bypassed stages are also set to zero, by accounting for the binary decision variables of the reflux ($b_{R,n}$) and boil-up ($b_{B,n}$) location as listed in eq. (4) and (5).

In context of the condenser stage ($b_{R,n} = 1$) and the reboiler stage (n_{max}), a doubled pressure drop is assumed to account for the effects of additional friction and phase change in the heat exchangers

$$p_{n_{max}} = p_{n_{max}-1} + 2 \cdot \left(\sum_{m=1}^{n_{max}} b_{B,m} \cdot \frac{\Delta p}{h}_m\right) \cdot \left(\sum_{m=1}^{n_{max}} b_{B,m} \cdot HETP_m\right). \quad (6)$$

The resulting pressure profile along the column is computed recursively as

$$p_n = p_{n-1} + \left(\frac{\Delta p}{h}\right)_n \bullet HETP_n, n = 2, \dots, n_{max} - 1. \quad (7)$$

In the current work the pressure at the condenser is fixed at the operating pressure, which is determined by means of a preceding shortcut calculation. However, this constraint could also be relaxed to allow for further improvements.

Initialization and solution procedure

The resulting MINLP problem is solved using a successively relaxed continuous reformulation approach, as outlined by Kraemer et al. [16]. The binary decision variables are relaxed and discrete values are successively enforced through penalty terms based on the Fischer-Burmeister functions. To enhance numerical robustness, a polylithic solution strategy (cf. **Figure 2**) is employed following the concept described by Waltermann and Skiborowski [18].

After a first initialization of the individual column models using all available stages (with fixed reflux and boil-up stages), the columns are connected to a direct sequence without the additional heat integration. This is further optimized for minimum external heat demand. In subsequent steps, the superstructure model is refined to the direct sequence with heat integration. Afterwards an initial estimation of the pressure profile is performed using fixed F-factor values to ensure the numerical robustness of the model, after which the HETP and pressure drop correlations, as well as sizing and cost correlations are introduced to the model. After solving the extended model for minimum heat demand and total annualized cost, resulting in a fully initialized model with reasonable values for all variables, the binary decision variables are relaxed, and a minimization of the TAC is performed with successively increasing penalties until a MINLP solution is obtained. The sequence of model refinement and NLP solutions is performed automatically, such that the overall MINLP problem is started from an initial problem definition up to the final design, representing a proven (local) minimum of the cost objective, which satisfies the predefined product purities.

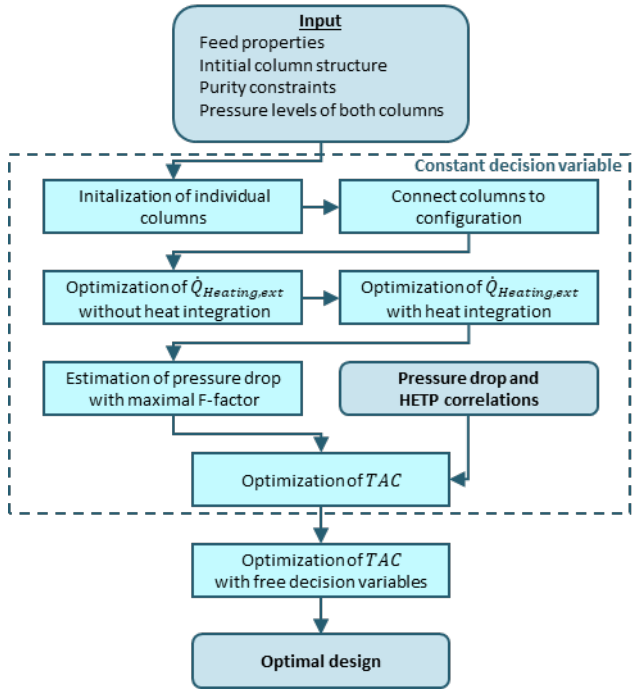

Figure 2. Initialization and solution procedure for direct heat integrated distillation sequence.

CASE STUDY AND RESULTS

To investigate the influence of stage-dependent pressure drop and HETP on the techno-economic design of heat-integrated direct distillation sequences, the separation of a saturated liquid feed into three product

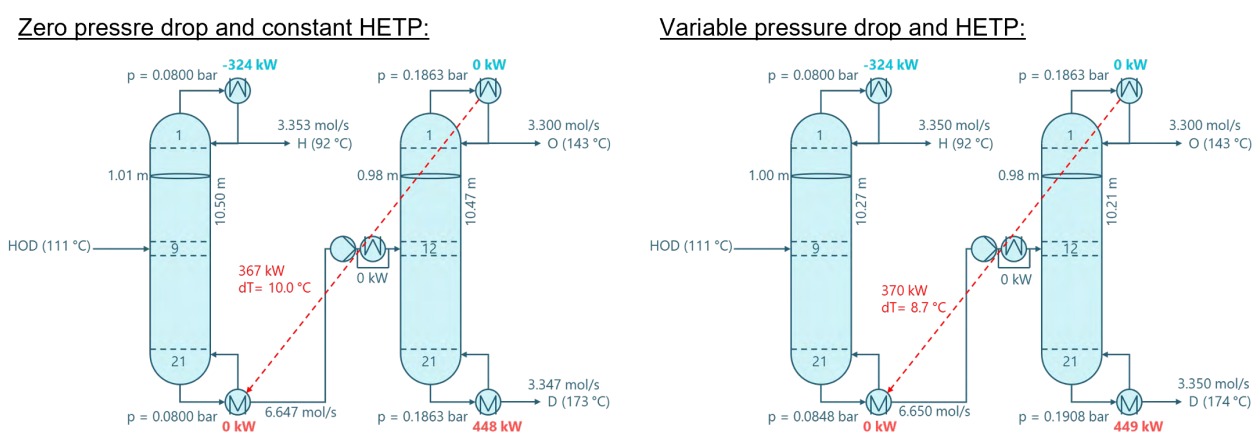

Figure 3. Optimalized heat-integrated direct distillation sequences for vacuum separation with isobaric (left) and pressure-drop-aware (right) superstructure.

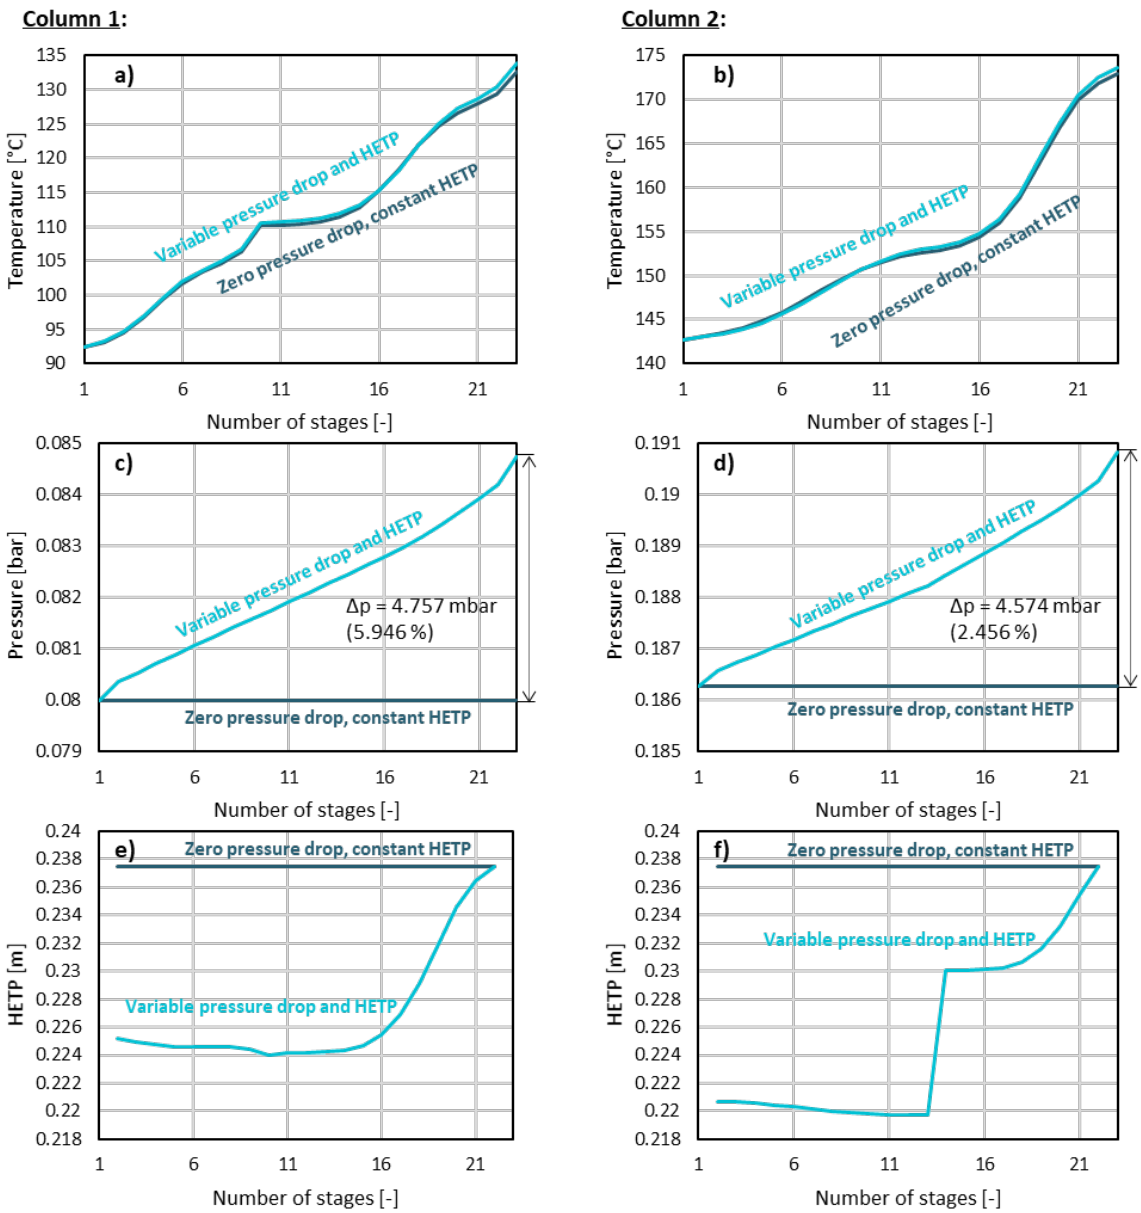

Figure 4. Temperature (a,b), pressure (c,d) and HETP (e,f) profiles inside of heat-integrated direct distillation sequence for vacuum separation with isobaric and pressure-drop aware superstructures.

streams with 99 mol% purity is considered.

The separation of hexanol, octanol, and decanol is investigated under vacuum conditions at 80 mbar [24], while the well-studied separation of benzene, toluene, and p-xylene is examined under atmospheric conditions [7, 19]. For both case studies, an equimolar feed and a feed flow rate of 10 mol/s are assumed. In the reference approach assuming isobaric columns with a constant HETP, the HETP is estimated for a maximum allowable F-factor of 1.75 $\sqrt{Pa}$ with the correlation stated in equation (5), resulting in a value of 0.238 m.

Cost data and economic assumptions are adopted from Adami et al. [19], while the VLE is modeled with the NRTL activity coefficient model and the Redlich-Kwong equation of state for the vapor phase fugacity. Pure component vapor pressure, heat capacities and heats of vaporization, are calculated with the extended Antoine equations and DIPPR correlations. All calculations are performed on an Intel i7-14700 CPU with GAMS 49.6.0 using SNOPT as NLP solver.

Influence on vacuum separation

The optimized flowsheets obtained from the

Table 1. Results of optimized flowsheet for atmospheric separation and vacuum separation with isobaric and pressure-drop-aware superstructure.

	Zero pressure drop, constant HETP		**Variable pressure drop and HETP**		**Relative difference**	
Vacuum separation:	Column 1	Column 2	Column 1	Column 2	Column 1	Column 2
Number of stages [-]	21	21	21	21	0%	0%
Column height [m]	10.50	10.47	10.27	10.21	-2.19%	-2.49%
Reboiler duty [kW]	366.65	448.01	370.21	448.88	+0.97%	+0.19%
Reboiler area [m²]	32.24	35.56	37.66	37.94	+16.81%	+6.69%
Bottom temperature [°C]	132.70	172.90	133.99	173.58	+0.97%	+0.39%
Bottom pressure [bar]	0.080	0.186264	0.0847568	0.190838	+5.95%	+2.46%
Temperature difference heat integration [K]	10.1		8.7		-13.86%	
Total annualized cost [k$/a]	285.324		286.417		+0.38%	
Atmospheric separation:	Column 1	Column 2	Column 1	Column 2	Column 1	Column 2
Number of stages [-]	28	44	28	44	0%	0%
Column height [m]	11.76	18.04	11.67	17.78	-0.78%	-1.47%
Reboiler duty [kW]	323.99	360.49	324.24	360.56	+0.08%	+0.02%
Reboiler area [m²]	28.23	14.03	29.05	14.19	+2.90%	+1.13%
Bottom temperature [°C]	122.65	161.37	122.93	161.61	+0.23%	+0.15%
Bottom pressure [bar]	1.01325	1.81977	1.02094	1.8304	+0.76%	+0.58%
Temperature difference heat integration [K]	10.1		9.8		-2.97%	
Total annualized cost [k$/a]	224.830		224.763		-0.03%	

optimization assuming isobaric operation and constant HETP and variable pressure drop and HETP are shown in **Figure 3**. Both approaches result in the same number of equilibrium stages in the optimal design. However, slight differences in the column height can be noticed. The variations are elucidated in **Figure 4**, which compares the resulting temperature profiles and highlights the variation in stage-specific pressure and HETP values for the resulting designs for both columns. While the temperature profiles are nearly identical, a slight increase in temperature toward the column bottom reduces the temperature difference for heat integration. Consequently, a considerable increase in the reboiler heat exchanger area is observed when pressure drop and HETP are included. Although the pressures increase almost linearly along the column height, which may also be reflected quite well with a constant stage specific pressure drop, the total pressure drop has no significant influence on the reboiler duty, which remains nearly unchanged between both approaches, indicating that isobaric operation and constant HETP can be appropriate approximations.

As the maximum F-factor, which determines the column diameter, occurs at the stage where the boil-up enters in both designs (cf **Figure 4**) the column diameter is basically equal for the two approaches. It is apparent that the HETP varies significantly between the rectifying and stripping sections, reflecting changes in vapor load along the column height and highlighting the importance for the resulting column sizing. This is also apparent in the key quantitative differences between the two designs, which are summarized in **Table 1**. Overall, it can be concluded that the consideration of column hydraulics does not show a significant effect on the energy demand and TAC, but has a noticeable effect on the HETP and the approach temperature in the heat exchanger. The latter might become more pronounced in case of a close boiling system, with higher stage requirements.

Influence on atmospheric separation

For the atmospheric separation, the qualitative trends observed in the vacuum case persist, but the impact of pressure drop is further reduced. As shown in **Table 1**, the relative pressure increase at the column bottom is below 1%, resulting in negligible effects on temperature profile, energy demand, and TAC.

However, analogous to the vacuum separation, the final column sizing is slightly influenced by the consideration of the stage-specific pressure drop and HETP values. In the context of the current investigation, the variations in investment cost exhibit no substantial influence on the overall TAC, as the operating costs are the prevailing factor in the objective function. In any case it is important to have a proper estimate of the pressure drop and HETP per tray to obtain a reliable estimate of the

column sizing and cost estimation.

Discussion and limitations

The current results indicate that the consideration of pressure drop and variable HETP can have a noticeable effect on the actual sizing of the column. However, the effect on the overall performance, in terms of heat duties and TAC, was limited for the wide boiling mixtures considered in the case studies. This indicates that assuming isobaric operation with constant HETP values is not expected to have a significant effect on the evaluation of competing configurations. While the overall economic impact is modest for the investigated wide-boiling mixtures, the effect may become larger for close-boiling or azeotropic mixtures and decreasing temperature differences for heat exchangers may become more decisive for thermally-coupled heat pump assisted distillation.

It also needs to be considered that the current proof-of-principle study applied simplified pressure drop and HETP correlations based on vendor data, which solely correlate with the F-factor. Consequently, variations in transport properties like liquid viscosity, surface tensions, as well as liquid load are not explicitly considered. Such effects may alter mass transfer efficiency and friction losses leading to a change in the correlation.

Nevertheless, the present implementation represents a first step towards incorporating pressure drop and variable HETP into an equation-oriented superstructure-optimization. It enables a more realistic assessment of structural design decisions at an early stage and provides valuable insights into the necessity of accounting for pressure-dependent effects in techno-economic optimization of energy-integrated distillation processes.

CONCLUSION

The presented extension of a superstructure-based optimization model for distillation processes by explicitly incorporating pressure drop and variable HETP values enables a systematic assessment of the influence of these effects on the techno-economically optimized grassroot design.

For the wide-boiling mixtures investigated, the inclusion of pressure drop and HETP correlations has only a minor impact on the overall energy demand and total annualized cost of heat-integrated direct distillation sequences. However, the results clearly indicate the effects on the specific sizing in terms of column height and heat exchanger areas. These effects become more pronounced under vacuum operation, where relative pressure changes along the column are significantly higher.

The presented case studies highlight that the relevance of pressure drop and stage-specific HETP strongly depends on the specific separation task and the chosen optimization objective. Consequently, these effects should be considered when evaluating alternative process configurations, especially in scenarios involving vacuum operation, close-boiling mixtures, or capital-intensive designs.

Future work will investigate the effect of the extension in the proposed approach for dividing wall columns, where vapor splits are inherently linked to pressure drop and column dimensions and heat pump assisted distillation. In addition, the robustness of different process alternatives with respect to multiple operating points and feed variation can be investigated. In order to account for changing properties, like viscosity and surface tension, and the liquid load over the column, generalized approaches, such as those proposed by Stichlmair et al. [25], will be incorporated and compared to the simplified, vendor-based correlations employed in the present study.

ACKNOWLEDGEMENTS

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 523327609.

AUTHOR IDENTIFIERS

Author ORCIDs:
Betram S: 0009-0007-2326-469X
Schnurr J: 0009-0005-2907-8551
Skiborowski M: 0000-0001-9694-963X

REFERENCES

1. Oak Ridge National Laboratory. Materials for separation technologies, Energy Emiss. Reduct. Opport. (2005) https://doi.org/10.2172/1218755
2. Kiss AA, Smith R. Rethinking energy use in distillation processes for a more sustainable chemical industry. Energy 203:117788 (2020) https://doi.org/10.1016/j-energy.2020.117788
3. Blahušiak M, Kiss AA, Babic K, Kersten SRA, Bargeman G, Schuur B. Insights into the selection and design of fluid separation processes. Separation and Purification Technology 194:301-318 (2018). https://doi.org/10.1016/j.seppur.2017.10.026
4. Rix A, Hecht C, Paul N, Schallenberg J. Design of heat-integrated columns: industrial practice. Chemical Engineering Research and Design 147:83-89 (2019). https://doi.org/10.1016/j.cherd.2019.05.009
5. Waltermann T, Sibbing S, Skiborowski M. Optimization-based design of dividing wall columns with extended and multiple dividing walls for three- and four-product separations. Chemical Engineering and Processing - Process

Intensification 146:107688 (2019). https://doi.org/10.1016/j.cep.2019.107688
6. Ramapriya GM, Selvarajah A, Jimenez Cucaita LE, Huff J, Tawarmalani M, Agrawal R. Short-cut methods versus rigorous methods for performance-evaluation of distillation configurations. Ind. Eng. Chem. Res. 57:7726-7731 (2018). https://doi.org/10.1021/acs.iecr.7b05214
7. Adami M, Espert D, Skiborowski M. Rapid multi-criteria screening of energy-integrated distillation processes for nonideal mixtures. Sep. Purif. Technol. 377:134463 (2025) https://doi.org/10.15480/882.15774
8. Skiborowski M, Harwardt A, Marquardt W. Conceptual design of azeotropic distillation processes. Distillation :305-355 (2014). https://doi.org/10.1016/b978-0-12-386547-2.00008-9
9. Dejanovi? I, Matijaševi? L, Jansen H, Oluji? Ž. Designing a packed dividing wall column for an aromatics processing plant. Ind. Eng. Chem. Res. 50:5680-5692 (2011). https://doi.org/10.1021/ie1020206
10. Sangal VK, Kumar V, Mishra IM. Optimization of structural and operational variables for the energy efficiency of a divided wall distillation column. Computers & Chemical Engineering 40:33-40 (2012). https://doi.org/10.1016/j.compchemeng.2012.01.015
11. Barttfeld M, Aguirre PA, Grossmann IE. Alternative representations and formulations for the economic optimization of multicomponent distillation columns. Computers & Chemical Engineering 27:363-383 (2003). https://doi.org/10.1016/s0098-1354(02)00213-2
12. Jobson M. Energy considerations in distillation. Distillation :225-270 (2014). https://doi.org/10.1016/b978-0-12-386547-2.00006-5
13. Luyben WL. Effect of tray pressure drop on the trade-off between trays and energy. Ind. Eng. Chem. Res. 51:9186-9190 (2012). https://doi.org/10.1021/ie300634a
14. Li F, Luo Y, Yuan X. Equation-oriented optimization of a distillation column considering stage hydraulics. Ind. Eng. Chem. Res. 59:13657-13668 (2020). https://doi.org/10.1021/acs.iecr.0c01972
15. Horsch AS, Acevedo AR, Skiborowski M. Optimal Retrofit of Simple Distillation Sequences to Thermally Coupled Side-Stream Configurations. In: Proceedings of the 33rd European Symposium on Computer Aided Process Engineering. Ed: Kokssis A, Georgiadis MC, Pistikopoulos EN. Elsevier (2023) https://doi.org/10.1016/b978-0-443-15274-0.50053-6
16. Kraemer K, Kossack S, Marquardt W. Efficient optimization-based design of distillation processes for homogeneous azeotropic mixtures. Ind. Eng. Chem. Res. 48:6749-6764 (2009). https://doi.org/10.1021/ie900143e
17. Skiborowski M, Harwardt A, Marquardt W. Efficient optimization-based design for the separation of heterogeneous azeotropic mixtures. Computers & Chemical Engineering 72:34-51 (2015). https://doi.org/10.1016/j.compchemeng.2014.03.012
18. Waltermann T, Skiborowski M. Efficient optimization-based design of energy-integrated distillation processes. Computers & Chemical Engineering 129:106520 (2019). https://doi.org/10.1016/j.compchemeng.2019.106520
19. Adami M, Bertram S, Espert D, Skiborowski M. Beyond dividing wall columns: improved process intensification through liquid-only transfer and heat integration. Chemical Engineering and Processing - Process Intensification 218:110559 (2025). https://doi.org/10.1016/j.cep.2025.110559
20. Turton R, Bailie RC, Whiting WB, Shaeiwitz JA. Analysis, Synthesis, and Design of Chemical Processes. Prentice Hall PTR (2009)
21. Biegler LT, Grossmann IE, Westerberg AW. Systematic Methods of Chemical Process Design. Prentice Hall PTR (1997)
22. Douglas JM. Conceptual design of chemical processes. McGraw-Hill (1988)
23. Sulzer brochure for structure packings (accessed 26.01.2026). https://www.sulzer.com/en/-/media/files/products/separation-technology/brochures/english/structured_packings_e10686_en_web.pdf
24. Schnurr J, Adami M, Skiborowski M. Refrigerant selection and cycle design for industrial heat pump applications exemplified for distillation processes. Systems and Control Transactions 4:1306-1311 (2025). https://doi.org/10.69997/sct.119494
25. Stichlmair J, Bravo JL, Fair JR. General model for prediction of pressure drop and capacity of countercurrent gas/liquid packed columns. Gas Separation & Purification 3:19-28 (1989). https://doi.org/10.1016/0950-4214(89)80016-7

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Data-Driven Multi-Objective Optimization of Energy, Environmental, and Economic Performances in Manufacturing with Physics-Consistent Deep Learning

Hyeonrok Choi[ab], Jaewook Lee[a], Won Yang[a] and Seong-il Kim[ac]*

[a] Low-Carbon Emission Control R&D Department, Research Institute of Sustainable Development Technology, Korea Institute of Industrial Technology, Cheonan 31056, Republic of Korea
[b] School of Mechanical Engineering, Sungkyunkwan University, Suwon 16419, Republic of Korea
[c] Green Process and Energy System Engineering, University of Science and Technology, Cheonan, Chungnam, Republic of Korea
* Corresponding Author: ksungil85l@kitech.re.kr.

ABSTRACT

Aluminium cold rolling is an energy-intensive process that has a substantial impact on CO_2 emissions and production cost, yet plant-level optimization remains challenging due to strong process nonlinearities and various operational constraints. This study develops a physics-consistent hybrid model that combines a Stone–Hitchcock–Ludwik analytical rolling-energy formulation with a residual deep neural network to predict the daily electricity consumption of three single-stand cold rolling mills. Using plant raw data, the hybrid model achieves lower prediction errors than conventional data driven model and yields line-specific physical parameters that agree well with the observed behaviour of each mill. On this basis, an NSGA-II-based tri-objective optimization is carried out to minimise daily energy use, CO_2 emissions, and specific production cost (SPC) by adjusting pass-wise reduction and tension schedules and line-wise production allocation. Case studies on a representative operating day and additional plant data show that the optimised operating strategy shifts production load from less efficient to more efficient lines and smooths pass-wise operating conditions, thereby consistently reducing daily energy consumption and unit cost while moderately decreasing CO_2 emissions without any hardware modifications. The proposed hybrid prediction–optimization framework thus provides a practical decision-support tool for integrated energy–environment–economic optimization in multi-line aluminium cold rolling operations.

Keywords: Aluminium cold rolling, Physics-consistent hybrid modelling, Rolling energy consumption, Multi-objective optimization, NSGA-II (Non-dominated Sorting Genetic Algorithm II)

INTRODUCTION

Global efforts to mitigate climate change have tightened greenhouse gas (GHG) reduction requirements toward carbon neutrality by 2050 across the manufacturing sector [1]. The aluminium industry, responsible for about 2% of global industrial emissions (≈1.1 $GtCO_2$-eq/year), is a representative carbon- and energy-intensive sector in which process electricity use is a major driver of emissions [2, 3]. In power systems still dominated by fossil fuels, electricity consumption translates directly into GHG emissions, making renewable penetration and fuel switching key mitigation options. In the Republic of Korea, however, a high fossil share in the generation mix and strong industrial dependence on electricity [4] limit the pace of power-sector decarbonisation, so process-level energy efficiency improvement and operational optimization emerge as particularly effective short- to mid-term options.

Cold rolling of aluminium is a power-intensive process that determines final strip thickness (tens of micrometres), flatness, and mechanical properties, and therefore requires high-speed, high-tension operation with precise control [5]. In single-stand foil mills operated under multi-pass schedules, the combination of reduction, entry/exit tension, and rolling speed at each pass jointly determines rolling load, torque, and thus electrical power

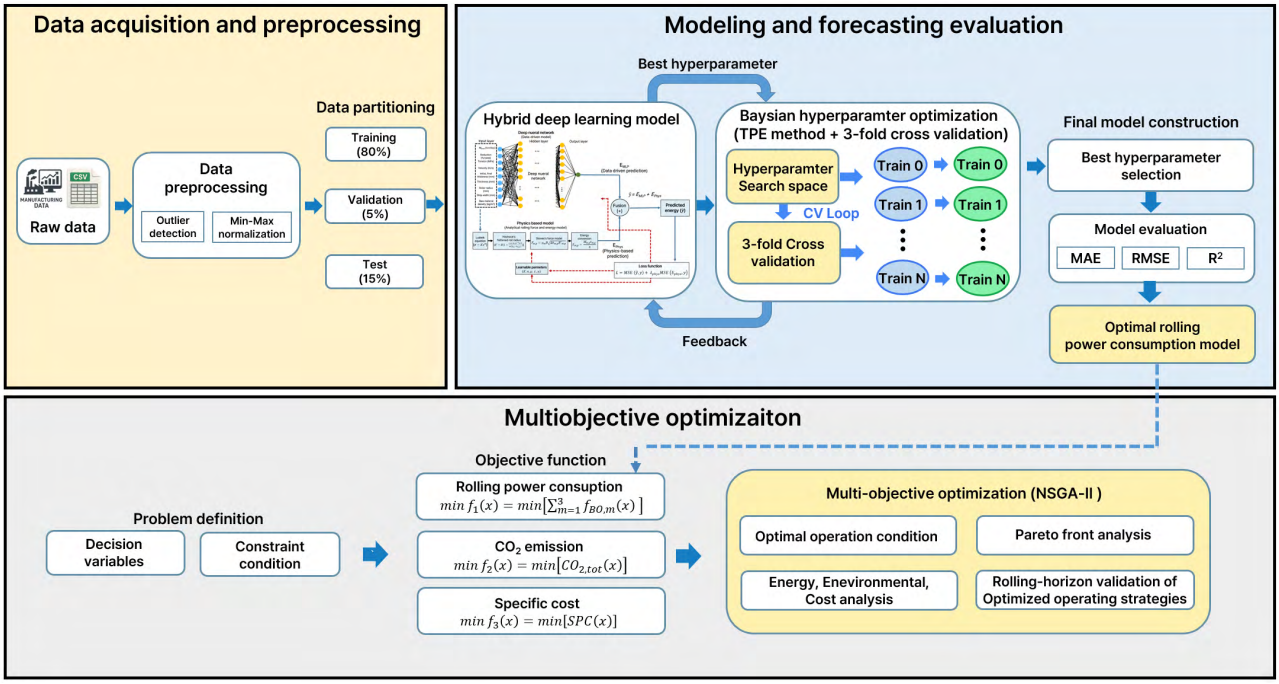

Figure 1. Overall framework of data-driven rolling power prediction and multi-objective optimization for aluminum rolling.

demand, under equipment constraints such as motor rating, drivetrain and inverter capacity, and available power supply. Designing pass schedules and allocating production among parallel lines is therefore a constrained decision problem, in which energy performance must be optimised concurrently with process stability and product requirements. Predictive models that reliably quantify the impact of process variables and production allocation on power demand and daily energy use are a key enabler for such optimization.

Modelling approaches for rolling load and power have evolved from analytical formulations and finite element method (FEM) simulations to artificial intelligence (AI)-based methods. Analytical models are computationally efficient and interpretable, but they require calibration for friction, work hardening, and drive efficiency, and their simplifying assumptions can limit accuracy and general applicability. FEM can represent roll–strip contact and deformation in detail, but it is too computationally expensive for repeated evaluation in plant-wide optimization. Data-driven AI models, including artificial neural networks (ANNs), have been applied to rolling load, power, and process monitoring tasks [6–9]; however, their black-box nature can limit interpretability and reduce reliability under extrapolative conditions. To address these limitations, hybrid modelling approaches that embed data-driven correction layers within physics-based structures have been proposed to balance computational efficiency, physical consistency, and predictive accuracy. In particular, previous studies have reported FEM–ANN hybrid models, genetic algorithm–backpropagation neural network (GA–BP) models, and stacking ensemble models combining extreme random tree (ET), AdaBoost regression tree (AdaBoost-RT), linear regression (LR), and backpropagation neural network (BPN) base learners with an LR meta-learner for load and flatness prediction [10–13].

As predictive models have improved, multi-objective optimization studies that embed such models into decision-making frameworks have become more common. In rolling applications, differential evolution (DE)- and Non-dominated Sorting Genetic Algorithm II (NSGA-II)-based methods have been used to derive Pareto-optimal solutions considering energy reduction, strip quality, slip prevention, and motor load balancing, while more recent studies have explored reinforcement learning-based pass schedule design and energy-oriented scheduling under time-of-use (TOU) tariffs and peak-demand constraints [14–16]. However, process-condition scheduling, such as pass-wise load and tension control, and TOU-based operating-time scheduling have mostly been treated separately. Consequently, integrated frameworks that link a physics-informed daily energy model with plant-level energy, environmental, and economic indicators for single-stand, multi-pass, high-speed, high-tension aluminium cold rolling mills remain very limited. To address this gap, this study develops an integrated framework that combines actual plant operating data, including pass-wise reduction, tension, speed, and thickness, with manufacturing execution system (MES)-based

production data, including daily throughput and line-wise allocation, to predict mill-wise daily electricity consumption and to use these predictions in NSGA-II-based tri-objective optimization of energy consumption, carbon dioxide (CO_2) emissions, and specific production cost (SPC) across operating conditions and inter-line production allocation. The overall workflow is summarised in Figure 1.

METHODOLOGY

Overview of the target process

The case study considers an industrial aluminium sheet cold rolling plant with a nominal capacity of approximately 500 ton/day. The plant consists of three single-stand cold rolling mills (RM4, RM5, and RM6). In each line, intermediate sheet stock supplied from upstream processes is reduced to the target thickness over multiple passes. The layout of the three mills, coil transfer routes, and pass directions are illustrated in Figure 2.

Problem definition

The plant energy management system (EMS) and process databases provide operating and production data, including alloy grade, strip geometry, pass schedule, key operating variables such as reduction, tension, and speed, and line-wise daily electricity consumption. Among the frequently used production routes, this study focuses on AA8075-series sheets operated under a six-pass cold rolling schedule.

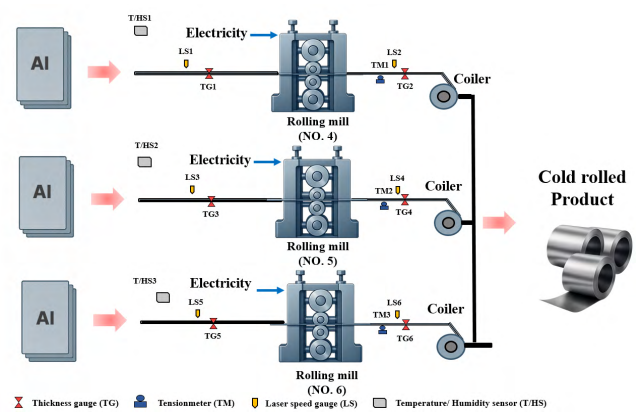

Figure 2. Schematic of three parallel aluminum rolling mills (No. 4–6) and product material flow.

The statistical characteristics of the full raw dataset are summarised in Supplementary Material S1. The process variables exhibit typical multi-pass cold rolling behaviour: the reduction ratio is highest in the initial passes and decreases with pass number; tension starts at an intermediate level in pass 1, increases up to pass 5, and slightly decreases in pass 6; and roll speed increases gradually with pass number. The exit thickness for all three lines is tightly distributed around the target value, indicating stable thickness control.

Table 1: Detailed values of the model parameters

Parameter	Description	Unit	Value
ρ	Density of strip (A8075)	kg/m^3	2710
N_m	Number of rolling mills	-	3
N_p	Number of passes	-	6
h_0	Initial thickness of foil	mm	0.3
h_{aim}	Target thickness of foil of final production	mm	0.03
R	Work roll radius	m	0.375
B	Width of strip	mm	1690
ψ	Force arm coefficient	-	0.45
η_m	Mechanical efficiency	-	0.85-0.995
μ	Friction coefficient	-	0.05-0.08
K	Strength coefficient	MPa	110-200
n	Strain hardening exponent	-	0.2-0.25

A representative operating condition is defined such that the total daily throughput of the three mills is 219.3 ton/day, with the incoming strip thickness $h_0=0.3$ *mm* and target final thickness $h_{aim}=0.03$ *mm* over six passes. The material, geometric, and operating parameters used in the analytical cold rolling model for this base case are summarised in Table 1, including strip density ρ, number of passes N_p, number of mills N_m, initial and target thickness h_0, h_{aim}, roll radius R_w, strip width B, arm coefficient ψ, mechanical efficiency η_m, friction coefficient μ, strength coefficient K, and strain-hardening exponent n. These parameters, summarised in Table 1, constitute the nominal parameter set and are applied to both the analytical cold rolling model and the hybrid energy prediction model, except where indicated otherwise.

Data preprocessing

Process raw data can contain noise, input errors, and missing values, which may deteriorate the stability and reproducibility of model training. To address this, pass-wise reduction, tension, speed, thickness, daily throughput, and electricity consumption are preprocessed by removing physically unrealistic outliers and imputing missing values with the variable-wise median, followed by min–max normalisation of all input variables.

Normalisation is defined as

$$X' = \frac{X - X_{min}}{X_{max} - X_{min}}, \quad (1)$$

X' is the normalised value, X is the original value, and $X_{\max}, X_{\min}$ are the maximum and minimum of each variable. The preprocessed data are split into 80% for training, 5% for validation, and 15% for testing.

Model development

The model architecture used to predict the daily electricity consumption of the cold rolling mills is illustrated in Figure 3. Normalised process inputs x (Production, pass-wise reduction, speed, tension, exit thickness, etc.) are fed in parallel to a data-driven deep neural network $f_{DL}(x, \theta)$ and to an analytical rolling-energy model $E_{phys}(x, \phi)$ based on the Stone–Hitchcock–Ludwik formulation. The network is a fully connected deep neural network (DNN) with rectified linear unit (ReLU) activation and a linear output layer for regression, trained using mean squared error (MSE) and the Adam optimiser with early stopping. The analytical module computes per-pass rolling force, torque, power, and energy for each mill and aggregates them to daily energy using mill-specific physical parameters ϕ=(K, n, μ, η).

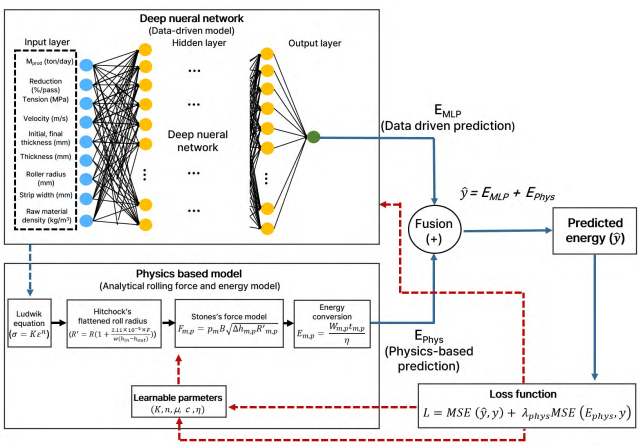

Figure 3. Physics-consistent hybrid model for rolling-energy prediction

The detailed equations and formulation of the analytical rolling-energy model are summarised in Supplementary Material S2. The hybrid predictor is formulated in residual form and trained with the composite loss.

$$\hat{y}(x; \theta, \phi) = f_{DL}(x, \theta) + E_{phys}(x, \phi) \quad (2)$$

$$L = MSE\,(\hat{y}, y) + \lambda_{phys} MSE\,(E_{phys}, y) \quad (3)$$

where $\lambda_{phys} > 0$ controls the weight on physics consistency. The physical parameters ϕ are treated as learnable and updated jointly with θ, so that the DNN mainly captures residual patterns while the analytical model provides the physical baseline and global energy scale.

Bayesian optimization

Bayesian hyperparameter optimization was employed to enhance the predictive performance of the physics-consistent hybrid model. A Tree-structured Parzen Estimator (TPE)–based algorithm, implemented via the Optuna library, was used to efficiently explore a mixed continuous–discrete search space, tuning key hyperparameters such as the number of hidden units, learning rate, weight decay, physics-loss weight (λ_{phys}), and batch size. Each trial was evaluated using the hybrid loss $L(\theta, \phi)$ on the validation set, and the hyperparameter configuration yielding the minimum validation loss over all trials was selected.

Model evaluation

The dataset was randomly shuffled and split into 80% for training, 5% for validation (Bayesian hyperparameter optimization), and 15% for testing. Test-set performance was evaluated using mean absolute error (MAE), root mean squared error (RMSE), and the coefficient of determination (R^2), defined for N test samples as

$$MAE = \frac{1}{N} \sum_{i=1}^{N} |\hat{y}_i - y_i| \quad (4)$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (\hat{y}_i - y_i)^2} \quad (5)$$

$$R^2 = 1 - \frac{\sum_{i=1}^{N} (\hat{y}_i - y_i)^2}{\sum_{i=1}^{N} (y_i - \bar{y}_i)^2}, \bar{y} = \frac{1}{N} \sum_{i=1}^{N} y_i \quad (6)$$

MULTI-OBJECTIVE OPTIMIZATION

Objective function

The multi-objective optimization problem is solved using the NSGA-II evolutionary algorithm. The previously developed physics-informed hybrid surrogate model is used as the evaluator: for each candidate operating condition (pass-wise reduction, tension, and production allocation), it predicts the daily electrical energy use of each mill. NSGA-II then searches for a set of Pareto-optimal solutions that simultaneously minimise energy consumption, CO_2 emissions, and economic cost.

$$\min f_1(x) = \min \sum_{m=1}^{3} f_{BO,m}(x) \quad (7)$$

$$\min f_2(x) = \min \alpha_{elec.} E_{total,day}(x) + \alpha_{raw} M_{raw,day}(x) + \alpha_{lub} M_{lub}(x) + \alpha_{add} M_{add}(x)\ [ton/day] \quad (8)$$

$$\min f_3(x) = \frac{C_{elec.}(x) + C_{ETS}(x) + C_{raw}(x) + C_{O\&M}}{M_{prod,day}(x)} \quad (9)$$

The objective functions in this study are defined at the plant level as follows: (i) total daily electricity consumption, evaluated as the sum of the hybrid-model-predicted daily electricity use of the three mills; (ii) total daily

CO_2 emissions arising from electricity consumption as well as raw material, lubricant, and additive use; and (iii) the specific production cost (SPC), defined as the ratio of total daily operating cost to daily production. The coefficients required for evaluating the CO_2 and SPC objective functions are compiled and derived in Supplementary Tables S3 and S4, respectively.

Table 2: Prediction accuracy of energy consumption by different machine learning models

Sheet	Model	MAE	RMSE	R^2
Rolling Mill 4	SVR	0.357	0.653	0.964
	RFR	0.349	0.660	0.964
	XGB	0.314	0.637	0.966
	DNN	0.434	0.684	0.961
	Hybrid DNN	0.202	0.400	0.987
Rolling Mill 5	SVR	0.284	0.366	0.983
	RFR	0.307	0.406	0.980
	XGB	0.285	0.376	0.982
	DNN	0.302	0.402	0.980
	Hybrid DNN	0.258	0.357	0.988
Rolling Mill 6	SVR	0.153	0.274	0.963
	RFR	0.178	0.294	0.957
	XGB	0.154	0.254	0.968
	DNN	0.136	0.216	0.967
	Hybrid DNN	0.115	0.169	0.986

Decision variables and industrial constraints

The decision vector comprises (i) daily production allocated to each mill, $M_{m,\ day}$(m=1, 2, 3) (ii) pass-wise reduction ratios R_i (i=1, ..., 6), and (iii) pass-wise tensions T_i (i=1, ..., 6), and can be written for each mill as

$$x = [\, M_{i,day},\ R_1, R_2 R_3 R_4 R_5 R_6, T_1, T_2, T_3, T_4, T_5, T_6] \qquad (10)$$

Search bounds are chosen to reflect the feasible operating window of the actual equipment: lower and upper limits for $M_{m,\ day}$ R_i, and T_i are set to the minimum–maximum values observed in plant data, thereby preventing unrealistic operating combinations. In addition, the total daily production is constrained to remain close to the reference target to avoid trivial improvements obtained by artificially changing throughput.

RESULTS AND DISCUSSION

Performance comparison of hybrid model

Table 2 compares the daily energy prediction performance of support vector regression (SVR), random forest regression(RFR), extreme gradient boosting (XGB), pure DNN, and the proposed Bayesian optimization (BO)-hybrid DNN for the three foil rolling lines (RM4–RM6) using identical inputs. The BO-hybrid DNN yields the lowest MAE/RMSE and highest R^2 across all lines, indicating that incorporating physical load–torque–power–energy relations improves both accuracy and robustness at the daily energy scale.

Figure 4 shows the BO-hybrid DNN predictions of

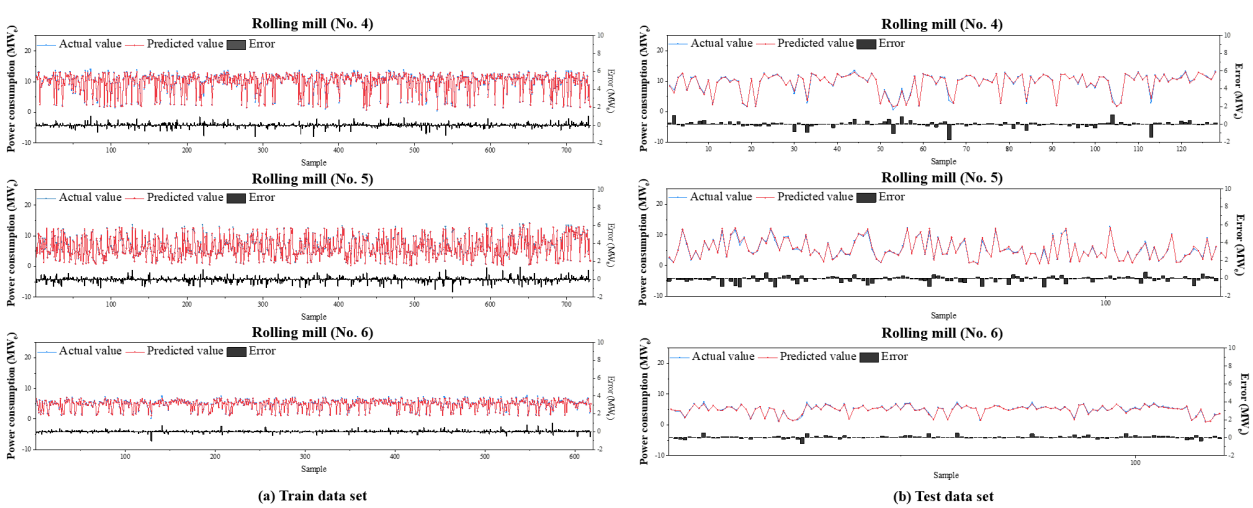

Figure 4. Actual and predictive value calculated by BO-Hybrid model.

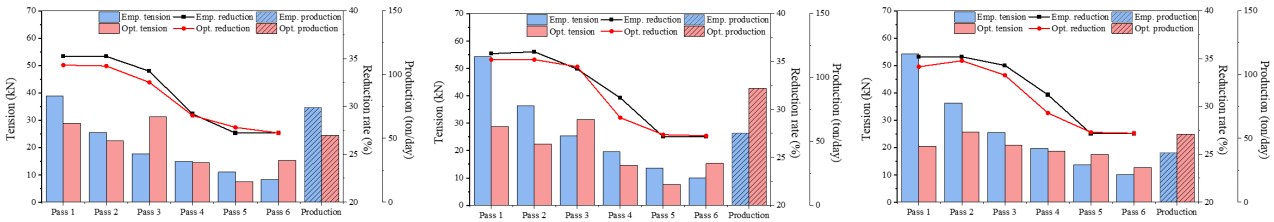

(a) Rolling mill (No. 4) (b) Rolling mill (No. 5) (c) Rolling mill (No. 6)

Figure 5. Optimization results of tension, reduction, and daily production for rolling mills No. 4–6 (empirical vs. optimized).

daily energy consumption for each line, with training results on the left and test results on the right. In each subplot, red markers denote measured values, black markers denote model predictions, and blue bars indicate the absolute error per sample. For all three lines, the predicted curves closely follow the measured trends, with most absolute errors remaining within a narrow range and only limited deviations at a few peak-load or non-standard operating points.

Table 3: Inferred physical parameters of the rolling process for the three mills

Parameter	Unit	Description	RM4	RM5	RM6
K_{MPa}	MPa	Strength coefficient	180	178	181
n	-	Strain-hardening exponent	0.24	0.25	0.24
μ		Friction coefficient	0.05	0.05	0.03
η	%	Mechanical efficiency	90.8	93.1	92.2

***RM : Rolling mill**

In addition to predicting energy, the BO-hybrid DNN simultaneously back-estimates the physical parameters (K, n, μ, η) for each mill, as summarised in Table 3. The inferred values fall within similar ranges but exhibit mill-specific variations, indicating that the model not only improves daily energy prediction but also yields interpretable, line-dependent physical parameters rather than purely statistical correlations.

Multi-objective optimization results

Case study of a representative operating day

Energy consumption and CO_2 emissions are strongly positively correlated, because CO_2 is computed directly from the predicted daily electricity use via a fixed grid emission factor; thus, reducing energy immediately reduces CO_2. In contrast, energy (or CO_2) and SPC exhibit an overall inverse relationship, since SPC is defined as total daily operating cost per unit throughput and includes fixed capital/ operations and maintenance (O&M) components as well as energy-dependent costs.

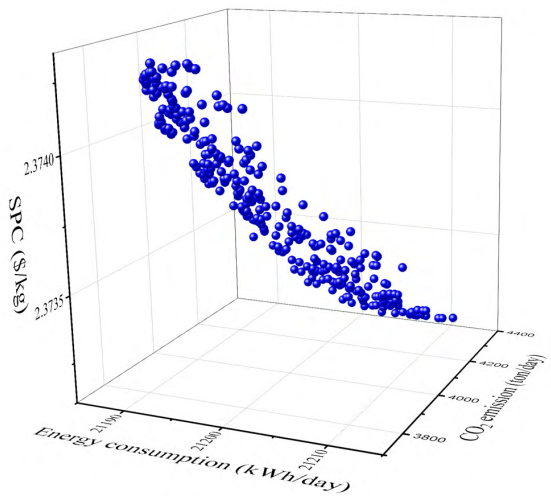

Figure 6. Pareto solution set in objective space for (a) 3D of energy consumption, CO_2 emission, and SPC

Along the Pareto front, some operating conditions slightly increase energy use but allow higher daily production and more favourable load sharing among mills, thereby diluting fixed costs per ton and yielding lower SPC despite higher energy and CO_2.

Figure 7 compares the three objective values—daily energy consumption, daily CO_2 emissions, and SPC—between the empirical operating condition and the optimized condition, shown in normalized form with absolute values annotated on each bar. Daily energy consumption decreases from 25.91 to 21.19 MWh/day (≈18.22%), and SPC from 2.51 to 2.502 USD/kg (≈0.32%), indicating a simultaneous reduction in unit production cost, including electricity and emissions trading system (ETS)-related charges. CO_2 emissions change only marginally, from 4158.5 to 4156.4 kg/day (≈0.05%), as a substantial fraction of total emissions arises from non-electricity contributions such as materials and process auxiliaries. Overall, these results summarise that appropriate adjustment of line-wise production allocation and pass-wise operating

conditions can improve energy, environmental, and economic performance simultaneously even under unchanged equipment and system configurations.

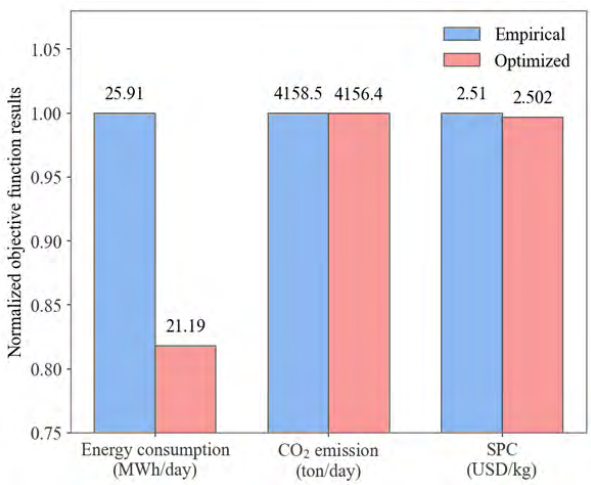

Figure 7. Relative comparison of normalized objective functions (Energy, CO_2, SPC)

Case study with additional operating data

In this section, the proposed hybrid prediction–optimization framework is validated on 40 additional operating days that were not used for training, under the actual plant constraints of each day (target throughput, mill capacities, speed/tension limits, etc.). For each day, NSGA-II is applied with pass-wise reduction, tension, speed, and line-wise production allocation as decision variables. As shown in Figure 8(a), although the total daily production varies between approximately 288 and 489 ton/day, the optimised solutions consistently redistribute load away from Mill 4 and toward the more energy-efficient Mills 5 and 6. Consequently, Figure 8(b) indicates that empirical daily electricity use, which ranges from 20.9 to 33.1 MWh, is reduced and flattened to about 20.3–21.9 MWh under the optimised operation, corresponding to an average energy reduction of roughly 21.12%. Figures 8(c) and (d) show that, despite relatively small percentage reductions in CO_2 (0.01–0.068%, but on 2.56 ton/day in absolute terms) and SPC also decreased, with the relative reduction ranging from 0.21 to 0.283% and averaging approximately 0.24%. Despite the dominance of raw material cost in the total cost structure, the optimization consistently improves energy, environmental, and economic performance across all 40 days while satisfying the production constraints.

CONCLUSION

This study developed a physics-consistent hybrid energy surrogate for daily electricity prediction in aluminium cold rolling by coupling a Stone–Hitchcock–Ludwik analytical rolling-energy model with a residual deep neural network. The hybrid model improved predictive capability while preserving physically interpretable mill-specific parameters consistent with actual plant behaviour. On this basis, an NSGA-II-based tri-objective optimization framework was formulated to minimise energy consumption, CO_2 emissions, and SPC through the adjustment of pass-wise operating conditions and line-wise production allocation. The results showed that the proposed framework can systematically improve plant operation by shifting production toward more energy-efficient mills and by smoothing pass-wise schedules under practical process constraints. Furthermore, out-of-sample validation confirmed that the framework remains effective under varying operating conditions. These findings demonstrate that the proposed approach offers a practical and scalable tool for integrated energy–environment–economic optimization of industrial aluminium cold rolling plants without additional hardware modification.

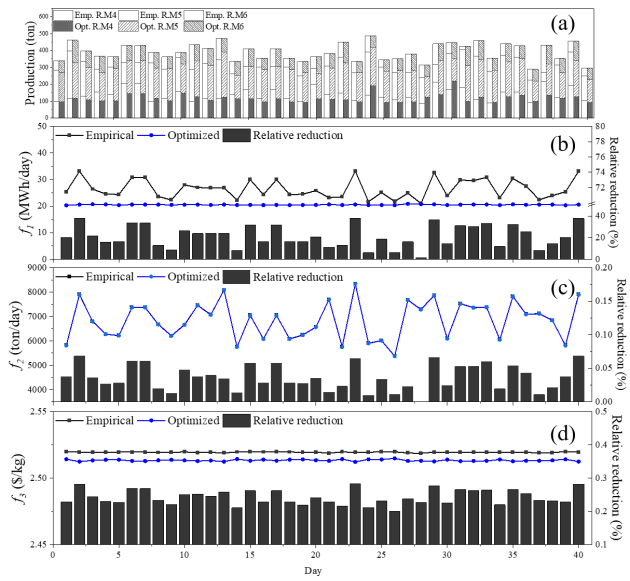

Figure 8: Effect of the optimized rolling schedules on (a) line-wise production, (b) daily energy consumption, (c) daily CO_2 emissions, and (d) daily SPC
DIGITAL supplementary material

SUPPLEMENTARY MATERIALS

Digital supplementary material is included in the Supplementary Material file which is downloadable at https://PSEcommunity.org/LAPSE:2026.0041.

AUTHOR IDENTIFIERS

Author ORCIDs:
Choi HR: 0009-0006-6282-4736

REFERENCES

1. International Energy Agency (IEA), Net Zero by 2050, IEA, Paris, 2021. Available at: https://www.iea.org/reports/net-zero-by-2050 (accessed 01.12.2026).
2. International Aluminium Institute (IAI), Greenhouse gas emissions: Aluminium sector, 2024. Available at: https://international-aluminium.org/statistics/greenhouse-gas-emissions-aluminium-sector/ (accessed 01.12.2026).
3. Kvande H, Saevarsdottir G, Welch BJ. Direct and indirect CO2 equivalent emissions from primary aluminium production. The Minerals, Metals & Materials Series :998-1003 (2022). https://doi.org/10.1007/978-3-030-92529-1_130
4. U.S. Energy Information Administration (EIA), South Korea, Country Analysis Brief, EIA, Washington, DC, 2023 (accessed 05.12.2025).
5. D. Meech, P. Lawlor, Aluminium rolling: improving production and quality with minimum energy consumption, Mater. Technol. 24 (2009) 144–145.
6. Yang L, Zhiying G, Xiaomin Z, Qingdong Z. Industrial data-driven intelligent forecast for chatter of cold rolling of thin strip with LSTM recurrent neural network. Journal of Mechanical Engineering 56:121 (2020). https://doi.org/10.3901/jme.2020.11.121
7. Chen Z, Liu Y, Valera-Medina A, Robinson F, Packianather M. Multi-faceted modelling for strip breakage in cold rolling using machine learning. International Journal of Production Research 59:6347-6360 (2020). https://doi.org/10.1080/00207543.2020.1812753
8. Chen D, Zhang R, Li Z, Li Y, Yuan G. Temperature distribution prediction in control cooling process with recurrent neural network for variable-velocity hot rolling strips. Int J Adv Manuf Technol 120:7533-7546 (2022). https://doi.org/10.1007/s00170-022-09065-8
9. Lakshmanan K, Borghini E, Beckmann A, Pleydell-Pearce C, Giannetti C. Data modelling and remaining useful life estimation of rolls in a steel making cold rolling process. Proced Comp Sci 2022;207:1057–66. https://doi.org/10.1016/j procs.2022.09.16
10. Lv C, Wang GD, Liu XH, et al. High-precision prediction of rolling load of finishing stands with neural networks. Iron Steel 1998;03:33–5. https://doi.org/10.3321/j issn:0449-749X.1998.03.009.
11. He YY, Yan X, Li LX, et al. Prediction model of rolling force of CSP line based on BP neural network. MaterMechEng 2014;38(010):79–82. CNKI:SUN:GXGC.0.2014-10- 020.
12. Zhang SH, Deng L, Che LZ. An integrated model of rolling force for extra-thick plate by combining theoretical model and neural network model. Journal of Manufacturing Processes 75:100-109 (2022). https://doi.org/10.1016/j.jmapro.2021.12.063
13. Cui CG. Intelligent prediction model based on neural network algorithm and mechanism model for rolling force in tandem cold rolling. Comput Modernization 2019;000(008):74–78, 91. https://doi.org/10.3969/j.issn.1006-2475.2019.08.014
14. Jia S, Li W, Liu X, Du B. Multi-objective load distribution optimization for hot strip mills. J Iron Steel Res Int 2013;20:27–61. https://doi.org/10.1016/S1006-706X (13)60052-7.
15. Hu ZY, Yang JM, Zhao ZW, Sun H, Che HJ. Multi-objective optimization of rolling schedules on aluminum hot tandem rolling. Int J Adv Manuf Technol 85:85-97 (2015). https://doi.org/10.1007/s00170-015-7909-1
16. L. Chen, L.L. Cao, Y.M. Wen, H. Chen, S.L. Jiang, A knowledge-based NSGA-II algorithm for multi-objective hot rolling production scheduling under flexible time-of-use electricity pricing, J. Manuf. Syst. 69 (20)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Joint Optimization of Feedstock Procurement and Production Planning in AD: A Deep Learning-Integrated Stochastic Programming Framework

Ruosi Zhang[a], Michael Short[s]*

[a] School of Chemistry and Chemical Engineering, University of Surrey, Stag Hill, Guildford GU2 7XH, UK
* Corresponding Author: m.short@surrey.ac.uk.

ABSTRACT

Anaerobic digestion (AD) across Europe and the UK faces increasing economic and operational pressure from volatile feedstock supply under climate extremes. Existing stochastic programming (SP) approaches for feedstock planning often rely on limited historical observations and/or simplify yield uncertainty in ways that miss the joint, non-linear response of crops to weather variability, thereby understating downside supply risk. We develop an integrated decision-support framework that links climate uncertainty to AD procurement planning by coupling mechanistic crop simulation, generative surrogate modelling, and stochastic optimization. First, APSIM is used offline to generate a mechanistic yield knowledge base across weather trajectories and discrete planting-density choices. Then, a conditional GAN (CGAN) is trained to produce non-parametric joint yield samples for multi crops conditioned on scenario features and management, enabling fast Monte Carlo evaluation. At last, these samples are embedded in a two-stage SP that optimizes first-stage land allocation and planting densities, with second-stage recourse represented by spot-market purchases to cover shortfall. The architecture is designed for stage-based rolling updates as forecasts are progressively replaced by observations. We demonstrate the framework on a 15-ha unit under different contracts of biogas-equivalent output. Results reveal a target-induced regime shift in optimal procurement. Under moderate production targets, monocropping solutions minimize cost with negligible loss in reliability, reflecting broad operational indifference across several land allocation patterns. As contract levels approach the biophysical limits of a 15-ha system (~110×10^3 m^3), the optimizer transitions into a risk-reducing regime where maize–rye double cropping becomes increasingly necessary. At high targets, the feasible set collapses toward near-complete double cropping. At a representative contract of 140×10^3 m^3, the least-cost rye-only plan achieves 34.3% supply confidence, while full double cropping increases confidence to 85.7% but with higher cultivation cost. Even under full intensification, an irreducible tail risk remains (14.3% shortfall frequency), implying unavoidable reliance on spot-market procurement under extreme seasons. The APSIM–CGAN–SP integration translates climate-driven biophysical uncertainty into actionable procurement strategies. It reveals threshold behavior, tail-risk exposure, and the limits of intensification under fixed land constraints. The framework supports both pre-seasons planning and rolling in-season updates, providing a quantitative basis for contract feasibility assessment, hedging design, and resilient feedstock procurement.

Keywords: Biomass, Stochastic Optimization, Energy Systems, Surrogate Model, Planning, CGAN, Anaerobic Digestion

INTRODUCTION

Anaerobic digestion (AD) is a mature bioenergy technology that converts organic matter into biogas [1]. Unlike intermittent renewable sources such as solar or wind power, AD facilities often operate as baseload units, which means they rely on a steady and predictable feedstock supply. In practice, even short-term feedstock

shortages can trigger substantial economic losses or force unplanned shutdowns [2]. This dependence turns feedstock availability into a planning problem, operators must decide how much to grow, how much to contract, and how much risk to leave to the spot market.

Energy crops such as maize and rye are widely used as high-energy feedstock in AD systems [3]. Beyond providing high-energy biomass, these crops support greenhouse gas mitigation, energy security, and rural development [3]. However, climate-driven yield variability and extreme events increasingly translate agricultural uncertainty into operational risk for AD sites [4], [5]. Yield forecasting serves not as a standalone objective but as a foundational input for the generation of supply chain scenarios. These scenarios are subsequently embedded in procurement and production-planning frameworks to enable hedging strategies that account for operational variability and upstream supply uncertainty [6].

Process-based crop models (e.g., Agricultural Production Systems sIMulator (APSIM) [7]) can provide realistic yield responses but are computationally too expensive for real-time stochastic optimization. Purely data-driven machine learning (ML) models are faster, yet often lack physical consistency and can be unreliable under extremes such as heatwaves or floods [8]. This limits their usefulness for decision support, especially when yield uncertainty and needs to be propagated into procurement and production planning. This creates a need for surrogate-based modeling strategies, balancing physical rigor with mathematical tractability in process systems engineering [9].

Existing AD supply-chain studies have largely focused on logistics or facility siting under deterministic assumptions, overlooking the stochastic nature of biophysical production [10]. While integrated crop–operations frameworks have appeared in the ethanol sector [11], such crop-to-decision integration remains limited in AD. Mechanistic models like APSIM remain accurate but infeasible for real-time optimization, and no operational framework currently connects biophysical yield uncertainty to tactical procurement decisions [12].

In mature electricity markets, stochastic programming (SP) is routinely used to hedge fuel contract risks under price volatility [13], [14]. AD procurement faces a similar make-or-buy challenge, but the dominant uncertainty arises not from market prices but from biophysical variation in crop yields. Yield variability translates into supply shortfalls and, consequently, greater pressure from volatile spot-market purchases. SP is therefore a natural framework for trading off upfront commitments (planting density and contracts) against scenario-dependent recourse actions and supply-gap penalties. If future yields can be predicted with adequate physical realism, managers can jointly optimize planting densities and purchase contracts to reduce exposure to costly spot-market feedstock.

To address the trade-off between biophysical realism in crop modeling and the computational demands of large-scale system decision-making, this paper proposes an integrated framework to support feedstock procurement decisions for ADs. This method utilizes the APSIM simulator to generate a high-fidelity dataset that captures the nonlinear response of crop yields to extreme weather events (such as heat waves). Based on this data, we train a physically-aware Conditional Generative Adversarial Network (CGAN) as a surrogate model, that simulates the complex mechanistic relationships between crops and conditions such as soil and weather. Finally, we embed this surrogate model generating probabilistic crop yield scenarios into a two-stage SP framework to evaluate land use and cultivation decisions under climate-driven uncertainty. By incorporating ensemble weather forecasts into the scenario generation process, this framework enables AD operators to optimize land allocation and planting density based on production objectives. The resulting decision-making tools reduce the risks associated with adverse yield outcomes and recourse costs related to volatile spot market feedstock procurement.

METHODOLOGY

The proposed framework (Figure 1) optimizes AD feedstock procurement by coupling biophysical process simulation, deep learning surrogates, and two-stage stochastic programming SP.

Problem Definition

A standard feedstock procurement agreement (FPA) for an AD site has been formalized, which specifies a contracted cultivation capacity A_{tot} (ha) and an annual baseload production target expressed as a feedstock requirement Q_{target} (m^3 biogas or t DM-equivalent). Internal supply is secured via fixed cultivation fees (covering seeds, management and land use). To ensure baseload operation, a minimum supply guarantee is effectively enforced by shortfall recourse.

The decision process is a two-stage stochastic problem. The first stage is here-and-now decision. Prior to the growing season, the manager determines the area allocation $\boldsymbol{A} = \{A_M, A_R, A_{DR}\}$ and planting densities $\boldsymbol{D} = \{D_m, D_r\}$ for maize (M), rye (R), and double-cropping (DR). It denotes sequential maize–rye double cropping within the same year, where rye and maize occupy less overlapping growing seasons on the same land base. DR is approximated in this prototype by combining monocrop yields with an annual reduction factor $\eta_{DR} \in (0, 1]$ that reflects a possible reduction in crop yield due to a shortened growing season or insufficient nutrients.

And the second stage is wait-and-see decision.

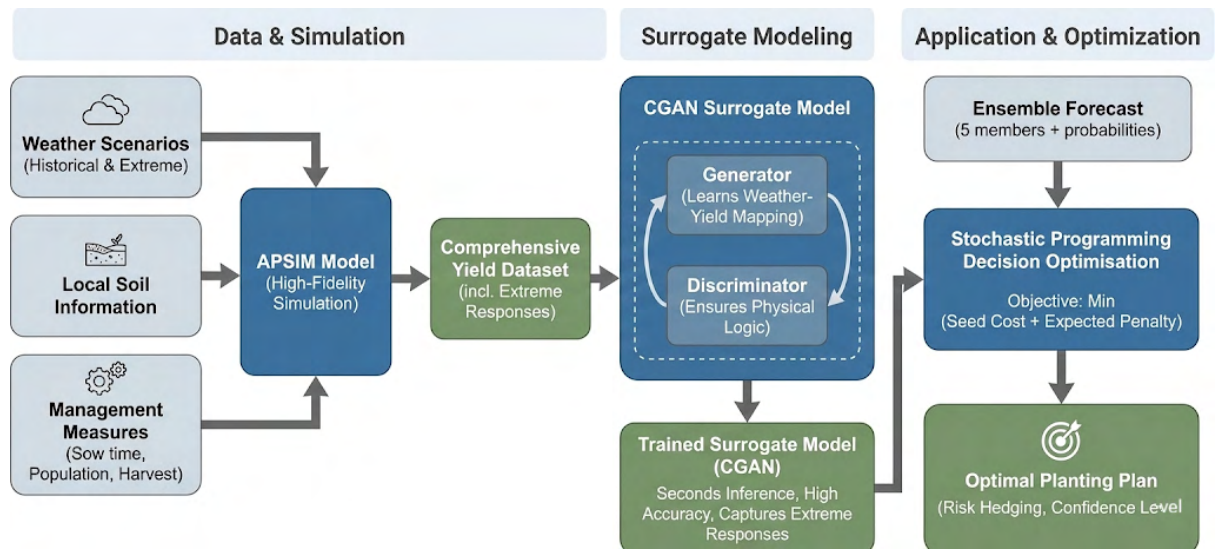

Figure 1. Integrated decision-support framework for AD feedstock procurement optimization under climate uncertainty, coupling biophysical simulation (APSIM), surrogate modelling (CGAN), and two-stage SP.

After weather scenarios s are realized, the yield Q_s is observed. If realized yields are insufficient to meet Q_{target} due to adverse weather, shortfalls are mitigated via spot-market purchases at a premium price P_{buy}, while surpluses are stored or sold at a credit price P_{sell}.

To address data scarcity under extreme climates, the APSIM is used to generate a mechanistic knowledge base. We incorporate a pre-simulation spin-up using historical weather data to initialize the earlier soil environment (e.g., plant-available water capacity). This captures the biophysical memory, where residual moisture conditions from previous years interact with current climate anomalies to produce non-linear yield responses

DR yields are represented by $\hat{Y}_{DR,s} = \eta_{DR}(\hat{Y}_{M,s} + \hat{Y}_{R,s})$, where $\eta_{DR} \in (0,1]$ is an operational efficiency factor capturing inter-crop competition and harvesting windows.

Surrogate Modelling via CGAN

To bypass the computational cost of APSIM, a CGAN is employed as a surrogate. Unlike standard regressions, the CGAN maps the conditional inputs (weather W_s and density D) to a probabilistic yield manifold. The Generator G takes a latent noise vector $z^{(j)} \sim \mathcal{N}(0, I)$ as an additional input. By drawing N independent samples of z for each weather scenario s, the model generates a yield cloud:

$$\left[\hat{Y}_{M,s}^{(j)}, \hat{Y}_{R,s}^{(j)}\right] = G\left(W_s, D, z^{(j)}\right), j = 1, \dots, N \quad (1)$$

where $W_s \in R^{12}$ represents standardized weather anomalies (temperature and precipitation deviations) across key growth stages. Here $s$indexes the five climate regimes and $j$indexes CGAN samples within each regime. Only the weather space is reduced by clustering historical meteorology into five climate regimes (empirical probabilities $\{\pi_s\}$). No post-CGAN reduction is applied. For regime s, the CGAN generates N_s joint samples $\left[\hat{Y}_{M,s}^{(j)}, \hat{Y}_{R,s}^{(j)}\right]$, all retained; hence $S = \sum_s N_s$. Scenario weights are $p_{s,j} = \pi_s / N_s$. The CGAN internalises physiological overturn points where high density exacerbates water stress, ensuring the surrogate is physics aware.

This N-fold sampling is critical as it captures the inherent stochasticity of the biological system (e.g., non-linear interactions between soil moisture and micro-climates) that weather anomalies alone cannot fully explain. It transforms the surrogate from a point-estimator into a stochastic generator, providing the density information required for risk-aware optimization.

The optimization employs a sample average approximation approach to reconcile weather uncertainty. Climate uncertainty is represented by $s \in S$ representative ensemble members (e.g., $S = 5$), capturing large-scale climatic shifts. These encompass worst-case (dry/hot), best-case (wet/cool), and median trajectories. We apply a scenario reduction logic where the probability P_s of discarded ensemble members is redistributed to the closest representative member.

To bridge the gap between high-fidelity mechanistic simulation and stochastic optimization, we apply an impact-based scenario discretization approach. The initial APSIM simulations, which represent the full population of potential biophysical outcomes, are partitioned into five representative climate regimes based on their total annual biomass.

Specifically, we utilize a quantile-based binning method to ensure that the entire variance of climate-driven risks, from extreme failures to optimal harvests, is captured. Each bin represents 20% of the probability

space ($P_s = 0.20$), effectively discretizing the continuous yield distribution into five distinct states: Worst-case, Lower-Mid, Median, Upper-Mid, and Best-case.

The mean feature vector (weather anomalies) and mean biomass of each quintile serve as the representative scenarios for the CGAN surrogate and the subsequent two-stage stochastic model. This discretization ensures that the hedging decisions are grounded in the actual statistical density of the 30-year climatological baseline, rather than arbitrary weather samples.

Two-stage stochastic programming formulation

The objective is to minimize the expected total cost (ETC), reconciling deterministic investments with probabilistic recourse impacts:

$$minC_{cultivation}(x) + \sum_s \sum_{j=1}^{N_s} p_{s,j}\ C_{\text{recourse}}(x, s, j), \text{with} \sum_s \sum_j p_{s,j} = 1 \quad (2)$$

Subject to:

$$A_M + A_R + A_{DR} \leq A_{tot}, A_k \geq 0 \quad (3)$$

The recourse cost $\mathcal{L}_s^{(j)}$ captures the make-or-buy logic:

$$\mathcal{L}_s^{(j)} = P_{buy} \cdot max\left(0, Q_{target} - \hat{Q}_s^{(j)}\right) - P_{sell} \bullet max\left(0, \hat{Q}_s^{(j)} - Q_{target}\right) \quad (4)$$

where $\hat{Q}_s^{(j)}$ is the total simulated yield for scenario s and Monte Carlo draw j, $p_{s,j} = \pi_s/N$. The first term represents the shortfall penalty, and the second term is the credit from surplus sales or recovery. All results and comparisons refer to this total expected cost. Hence, plans with very low cultivation expenditures may still be suboptimal if they induce higher expected shortage or imbalance penalties in the recourse stage.

$$\hat{Q}_s^{(j)} = \sum_{c \in \{m,r\}} A_c \hat{Y}_{c,s}^{(j)} \eta_c\,, \quad (5)$$

where $Y_{c,s}$denotes the realised dry-matter yield (t/ha) under scenario s, and $\eta_c$is the crop-specific conversion to methane (e.g., $\eta_c = VS_c \times BMP_c, Nm^3\, CH_4/t$). This maps biomass outputs to the contractual unit consistently across monocrop and double-crop cases.

To ensure resilience against tail-risks, we evaluate the Cost Value-at-Risk ($VaR_{0.95}$). Let $TC_s^{(j)}$ be the realised total cost sample:

$$TC_s^{(j)} = Cost_{seed} + \mathcal{L}_s^{(j)} \quad (6)$$

where $VaR_{0.95}(TC)$ is defined as the 95th percentile of the scenario-weighted cost distribution. This metric allows the operator to quantify the financial exposure during synchronized biophysical failures, beyond simple expected values.

The framework supports rolling updates at crop growth-stage decision epochs τ (e.g., sowing, early growth, flowering, pre-harvest). At each τ, forecast components in $W_{s,\tau}$ are updated by replacing them with realised observations and refreshed ensemble information, yielding updated scenario inputs and (if required) updated spot price $P_{buy,\tau}$ from latest market quotes. The two-stage problem is resolved with updated inputs. After sowing, A and D remain fixed, while the rolling update refines the procurement hedge under increasingly accurate yield outlooks.

RESULTS AND DISCUSSION

Biophysical Yield Response and Sensitivity Analysis

Before optimising the procurement strategy, we first quantify the biophysical response of crop yield to planting density under climatic stress. Figure 2 shows APSIM-simulated dry matter yields of maize and rye as functions of planting density for five representative scenarios selected by ranking all ensemble members by APSIM-implied total annual biomass and taking predefined yield quantiles. In the upper-quantile (favourable) case, maize averages $17.5\ t\ DM \cdot ha^{-1}$and rye $11.0\ t\ DM \cdot ha^{-1}$. In contrast, in the lower-quantile (adverse) case the combined biomass availability declines markedly, with total dry matter reduced to approximately $20.6\ t\ DM \cdot ha^{-1}$. This quantile-based spread provides the empirical basis for the subsequent hedging optimisation, which trades off deterministic cultivation costs against the expected cost of shortfall procurement under low-yield realisations.

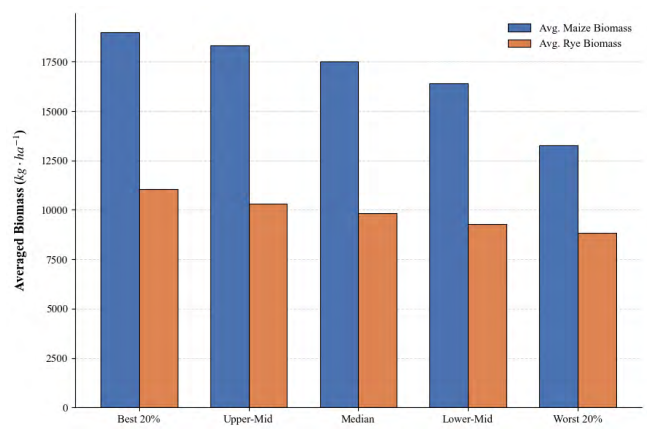

Figure 2. Mean crop biomass (dry matter) response for maize and rye under 5 climate ensemble scenarios, categorised by total yield performance quantiles ($kg \cdot ha^{-1}$).

The yield–density relationship is inherently non-linear and strongly regulated by inter-annual climate variability, where increased density does not monotonically

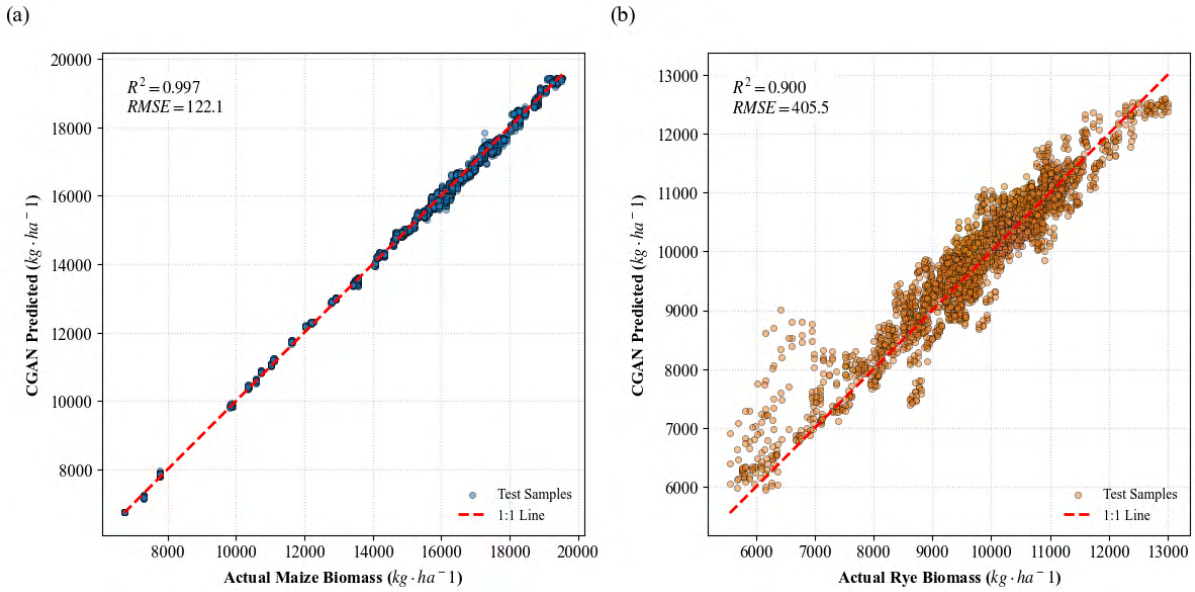

Figure 3. Validation of the CGAN surrogate model on the independent test dataset ($n = 3024$). (a) Maize biomass yield. (b) Rye biomass yield.

translate into higher productivity under water-limited conditions [15]. Increasing planting density generally accelerates canopy closure and increases leaf area index (LAI), thereby enhancing the potential for radiation interception. However, LAI is only a proxy for light capture and is not sufficient to determine final biomass or economic yield. Our mechanistic simulations confirm that while higher densities expand radiation interception capacity, they simultaneously exacerbate transpiration demand, potentially triggering yield collapse during synchronized heat-water stress. This non-linear coupling necessitates the use of a mechanistic solver like APSIM to capture the stochastic tipping points that purely statistical models often overlook. Empirically, this often results in a plateauing (or even declining) yield response beyond an intermediate density, with the optimal density shifting across weather scenarios, consistent with our motivation to treat density and area allocation as hedging decisions under uncertainty.

The non-linear yield response and the pronounced quantile gap indicate that a deterministic, average-year calibration is insufficient for procurement planning. Instead, the decision must account explicitly for the distributional tail where shortfall costs concentrate. This motivates the stochastic optimisation model, which hedges the plant's operational costs against low-probability, high-impact adverse seasons. Implementing this framework with APSIM directly is infeasible because the optimisation requires repeated evaluation of yield distributions under many scenarios and decision candidates. To enable fast probabilistic sampling while retaining biophysical realism, we train a CGAN as a surrogate of APSIM.

The predictive performance of the CGAN surrogate model was evaluated using an independent test dataset ($n = 3024$ scenarios). The parity plots in Figure 3 illustrate the model's ability to replicate the complex biophysical outputs of the APSIM platform across diverse climatic conditions and management intensities.

As shown in Figure 3a, the CGAN achieved superior predictive fidelity for maize biomass ($R^2 = 0.997$). The data points are tightly clustered along the parity line across the entire productivity spectrum. This indicates that the generative network successfully captures the primary physiological drivers of maize growth under summer conditions. In contrast, the prediction for rye biomass (Figure 3b) exhibited higher variance ($R^2 = 0.90$). While the general trend remains robust, the increased dispersion suggests that the secondary crop in the system is subject to more complex stochastic influences that are harder for the model to generalize.

Beyond parity and RMSE/R^2, we compared the marginal density of maize and rye yields (empirical vs CGAN) and observed close alignment overall; modest overestimation of rye in the lowest decile may slightly favor double-crop hedging under severe-drought scenarios, a point considered in the sensitivity discussion.

A critical observation in the rye validation is the systematic over-estimation in low-productivity regimes. This phenomenon is driven by two main factors. First, unlike maize, which grows during the relatively stable summer window, rye's phenology spans the winter-spring transition. It is highly sensitive to events such as late-spring frosts or erratic rainfall distributions. While APSIM's

mechanistic equations may trigger near-zero yields once stress thresholds are exceeded. As a function approximator, the CGAN tends to soften such hard threshold behavior, producing a less discontinuous lower tail (regression-to-the-mean). The smaller absolute magnitude of rye yields further reduces the signal-to-noise ratio in the joint-output model, potentially biasing learning toward the dominant maize gradients. Importantly, the surrogate shows no drift and preserves the directional response to planting density, supporting its use as a computationally efficient generator of yield samples for Monte Carlo-based stochastic optimization.

We conducted a sensitivity analysis by sweeping the contracted annual production target from 80 to $200 \times 10^3\ m^3$, spanning regimes from readily achievable to practically infeasible for a 15-ha farm system. For each target level, we solved the land-allocation problem over three cropping options: maize monocropping, rye monocropping, and maize–rye double cropping. Visualised the resulting optimal area shares in Figure 4. Each point on the curve corresponds to the minimum-cost solution (under the defined reliability metric) obtained by enumerating feasible density pairs and searching over land allocations with 1 ha resolution.

Figure 4 reveals a clear regime shift in the optimal allocation as the production target increases. At low targets ($\approx 80 - 110 \times 10^3\ m^3$), the optimiser predominantly allocates land to maize and rye monocropping, indicating that under moderate contractual requirements, a single-crop strategy provides sufficient biomass supply without relying on intensification. An additional feature in Figure. 4 is the intermittent appearance of a small DR share within the low-target regime, embedded in an otherwise single crop dominated allocation. This should be interpreted as a near-indifference band rather than a qualitative regime change. Under the objective of minimizing expected total cost (seed expenditure plus expected shortfall purchasing), several allocation patterns can become quasi-equivalent in this range. Introducing a marginal area of the DR system may slightly reduce tail shortfall exposure under adverse realizations, which can offset its yield-penalty factor and additional seed inputs. Because the objective differences are small and land is discretized at 1-ha resolution, the optimizer may switch between near-tied allocations across adjacent targets, producing the visually inserted sliver. Importantly, this indicates operational robustness in the low-target regime, that small changes in allocation do not materially affect expected cost or reliability.

Around $110 \times 10^3\ m^3$, the allocation begins to transition. This marks a threshold for operational sustainability, as further increases in target yields bring mono-crop strategies close to their biophysical limits, posing significant risks under unfavourable yield expectations. Beyond this threshold, the optimiser progressively reallocates area away from monocropping and toward intensification via double cropping, which becomes the dominant strategy as targets approach $135 - 140 \times 10^3\ m^3$.

For higher targets ($> 140 \times 10^3\ m^3$), the solution space collapses toward an all-out intensification regime in which almost all land is assigned to double cropping. This indicates that, for a fixed land base, meeting high contractual requirements relies primarily on stacking production within the year rather than marginal tuning of a single crop. Overall, Figure 4 summarises how increasing contractual pressure drives a qualitative change in the optimal procurement strategy, from simple monocropping to increasingly intensive multi-crop deployment. This marks a clear shift in land-allocation behaviour and reveals distinct planning regimes.

To illustrate the decision complexity at a risk-sensitive contract level, we fix the production target at $Q_{tar} = 140 \times 10^3 m^3$ for the 15-ha system and evaluate the cost–reliability trade-off over the feasible land-allocation and density options. We define supply confidence as the probability that internal production meets the contractual target:

$$Confidence = Pr(\hat{Q} \geq Q_{tar}) \tag{7}$$

It is estimated via CGAN-based Monte Carlo sampling (scenario-weighted) as the fraction of simulated realizations in which annual supply exceeds the target.

Table 1 summarises three representative strategies spanning the economic-to-reliability spectrum. Land is reported as D/M/R ha, where D denotes maize–rye double cropping, M maize monocropping, and R rye monocropping. Pop denotes planting densities for maize/rye. At the low-input extreme (0/0/15), allocating all land to rye minimises cultivation cost (£5, 536) but achieves only 34.3% confidence, implying frequent reliance on external procurement under adverse yield realizations. A balanced strategy (7/1/7) increases confidence to 59.4% at a moderate cost (£7, 009), reflecting that reallocating hectares toward intensification and diversification can materially reduce downside exposure. The high-reliability extreme (15/0/0) assigns all land to double cropping, achieving the highest confidence observed in this set (85.7%) at the highest cost (£8, 364). Notably, even the most intensive plan does not reach 100% confidence, indicating residual exposure to extreme climate realizations under a fixed 15-ha land constraint. In those cases, recourse actions (spot-market purchases) remain necessary. Practically, these strategies correspond to different managerial priorities. If the operator is cost-constrained and can tolerate frequent recourse purchases, the low-input rye plan is preferred. If a balanced hedge is required, the mixed allocation offers a mid-cost improvement in reliability. If the minimum supply guarantee compliance is the dominant objective, the solution shifts toward full intensification via double cropping.

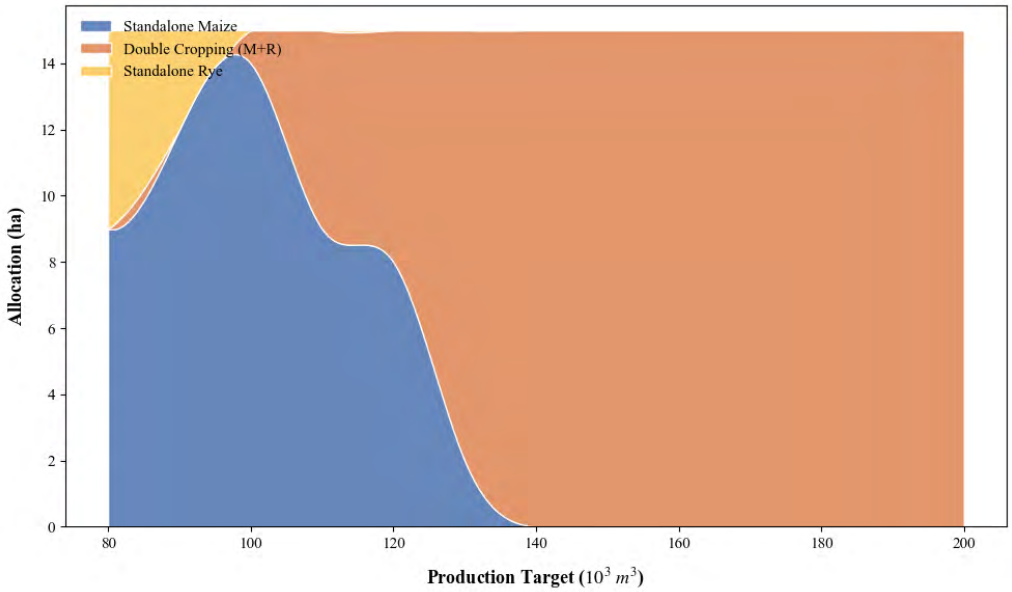

Figure 4. Optimal land allocation across production targets for a 15-ha system. For each annual production target ($80 - 200 \times 10^3$ m^3), the optimization selects the area (ha) allocated to maize monocropping (blue), rye monocropping (yellow), and maize–rye double cropping (orange). The stacked areas visualise the evolution of the optimal strategy as contractual requirements increase, revealing a threshold-like transition from monocropping-dominated solutions to intensification via double cropping.

Overall, the $140 \times 10^3 m^3$ target represents a regime where the system transitions from cost-first planning to explicit risk mitigation. Hectare-level allocation choices produce large changes in confidence, while marginal cost increases eventually yield decreasing reliability gains. This provides a quantitative basis for contract feasibility discussions, either accepting a target with residual procurement exposure, or complementing agronomic optimisation with target flexibility and/or external buffering arrangements.

Table 1: Strategic decision matrix for optimal land allocation under a high production target (140×10^3 m^3). The table compares the performance of three Pareto-optimal regimes: Economic, Balanced, and Reliable, extracted from the climate-resilience simulations. Allocations are presented as hectares of double-cropping (D), standalone maize (M), and standalone rye (R).

Land (D/M/R ha)	Pop (M/R)	Total Cost (£)	Confidence (%)
0.0/0.0/15.0	8/250	5, 535.71	34.30
7.0/1.0/7.0	8/250	7, 008.57	59.40
15.0/0.0/0.0	8/250	8, 364.29	85.70

Although the numerical experiments in this paper focus on pre-season planning, the proposed architecture is explicitly designed for stage-based rolling updates. As the season progresses, forecast components in the scenario feature vector $W_{s,\tau}$ can be progressively replaced by realized observations, and procurement price inputs ($P_{spot,\tau}$) can be refreshed from market quotes. This update decreases predictive uncertainty in the CGAN-generated yield distribution and enables rapid re-evaluation of shortfall exposure under the same committed planting plan. Each solve is a two-stage SP with first-stage cropping/procurement decisions and second-stage recourse. Designing explicit mid-season controls is left for the journal version. The surrogate's low evaluation cost is precisely what makes such sequential re-optimisation operationally feasible.

CONCLUSION

This study addresses the robustness of AD feedstock procurement under extreme climate uncertainty by developing a closed-loop decision support system encompassing weather scenario synthesis, mechanistic yield generation and stochastic planting decision optimization with rolling updates. Through typical case study, the study concludes that:

- The CGAN surrogate trained on APSIM simulations reproduces non-linear yield responses to climate variability and management intensity with strong out-of-sample agreement. This provides realistic, distributional yield inputs for downstream

stochastic optimization, enabling planners to hedge against downside under-yield tail events rather than relying on mean-year assumptions.

- The optimal planting structure (land allocation across monocropping vs double cropping, and discrete planting densities) is highly target dependent. As the contracted gas target increases, the decision space compresses and the optimal policy exhibits regime shifts. Low targets admit multiple low-cost feasible allocations, whereas high targets push the system toward intensification and diversification, which reveals threshold-like behavior consistent with biophysical and land constraints.
- A cost–reliability trade-off evaluation at the intensified target level ($140 \times 10^3 m^3$) shows residual exposure even under maximal intensification. In our case study, the most reliable strategy still requires recourse procurement in approximately 14.3% of realizations. It quantifies the remaining vulnerability to extreme climate outcomes under a fixed land base and highlights a practical carrying-capacity ceiling for contract design.
- In practice, the framework yields three actionable strategy classes: economical, balanced, and reliability oriented. Multiple strategies allow decision-makers to select policies aligned with budget and risk preferences. The rolling (stage-based) update capability allows the procurement hedge to be revised as forecast uncertainty collapses through the season, improving decision timeliness without changing already-committed sowing decisions.

AUTHOR IDENTIFIERS

Author ORCIDs:
Ruosi Zhang: 0000-0002-4931-9771

REFERENCES

1. Alengebawy A, Ran Y, Osman AI, Jin K, Samer M, Ai P. Anaerobic digestion of agricultural waste for biogas production and sustainable bioenergy recovery: a review. Environ Chem Lett 22:2641-2668 (2024). https://doi.org/10.1007/s10311-024-01789-1
2. Espinoza Pérez AT, Camargo M, Narváez Rincón PC, Alfaro Marchant M. Key challenges and requirements for sustainable and industrialized biorefinery supply chain design and management: a bibliographic analysis. Renewable and Sustainable Energy Reviews 69:350-359 (2017). https://doi.org/10.1016/j.rser.2016.11.084
3. Souza GM, Ballester MVR, de Brito Cruz CH, Chum H, Dale B, Dale VH, Fernandes ECM, Foust T, Karp A, Lynd L, Maciel Filho R, Milanez A, Nigro F, Osseweijer P, Verdade LM, Victoria RL, Van der Wielen L. The role of bioenergy in a climate-changing world. Environmental Development 23:57-64 (2017). https://doi.org/10.1016/j.envdev.2017.02.008
4. Nyéki A, Neményi M. Crop yield prediction in precision agriculture. Agronomy 12:2460 (2022). https://doi.org/10.3390/agronomy12102460
5. Mall RK, Gupta A, Sonkar G. Effect of climate change on agricultural crops. Current Developments in Biotechnology and Bioengineering :23-46 (2017). https://doi.org/10.1016/b978-0-444-63661-4.00002-5
6. Lepnaan Dayil J, Akande O, Mahmoud AED, Kimera R, Omole O. Challenges and opportunities in machine learning for bioenergy crop yield prediction: a review. Sustainable Energy Technologies and Assessments 73:104057 (2025). https://doi.org/10.1016/j.seta.2024.104057
7. D. P. Holzworth *et al.*, "APSIM – Evolution towards a new generation of agricultural systems simulation, " *Environ. Model. Softw.*, vol. 62, pp. 327–350, Dec. 2014, doi: 10.1016/j.envsoft.2014.07.009.
8. Tamayo-Vera D, Wang X, Mesbah M. A review of machine learning techniques in agroclimatic studies. Agriculture 14:481 (2024). https://doi.org/10.3390/agriculture14030481
9. Bhosekar A, Ierapetritou M. Advances in surrogate based modeling, feasibility analysis, and optimization: a review. Computers & Chemical Engineering 108:250-267 (2018). https://doi.org/10.1016/j.compchemeng.2017.09.017
10. Mayerle SF, Neiva de Figueiredo J. Designing optimal supply chains for anaerobic bio-digestion/energy generation complexes with distributed small farm feedstock sourcing. Renewable Energy 90:46-54 (2016). https://doi.org/10.1016/j.renene.2015.12.022
11. Kim S, Kim S. Hybrid simulation framework for the production management of an ethanol biorefinery. Renewable and Sustainable Energy Reviews 155:111911 (2022). https://doi.org/10.1016/j.rser.2021.111911
12. Y. Shi *et al.*, "Deep Learning Meets Process-Based Models: A Hybrid Approach to Agricultural Challenges, " *ArXiv Prepr. ArXiv250416141*, 2025.
13. Gal N, Milstein I, Tishler A, Woo CK. Fuel cost uncertainty, capacity investment and price in a competitive electricity market. Energy Economics 61:233-240 (2017). https://doi.org/10.1016/j.eneco.2016.11.014

14. Chabar R, Pereira M, Granville S, Barroso LA, Iliadis N. Optimization of fuel contracts management and maintenance scheduling for thermal plants under price uncertainty. 2006 IEEE PES Power Systems Conference and Exposition :923-930 (2006). https://doi.org/10.1109/psce.2006.296437
15. Lobell DB, Roberts MJ, Schlenker W, Braun N, Little BB, Rejesus RM, Hammer GL. Greater sensitivity to drought accompanies maize yield increase in the U.S. midwest. Science 344:516-519 (2014). https://doi.org/10.1126/science.1251423

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Method for Uniquely Determining Robust Operating Conditions in Simulated Moving Bed Chromatography

Kensuke Suzuki[ab]*, Tomoyuki Yajima[a], and Yoshiaki Kawajiri[ac]

[a] Nagoya University, Department of Materials Process Engineering, Nagoya, Japan
[b] The University of Tokyo, Department of Chemical System Engineering, Tokyo, Japan
[c] LUT University, School of Engineering Science, Mukkulankatu, Finland
* Corresponding Author: k-suzuki@pse.t.u-tokyo.ac.jp.

ABSTRACT

In this study, we propose a method to uniquely determine robust operating conditions for simulated moving bed (SMB) chromatography, an essential continuous liquid-phase separation technique in the pharmaceutical industry, in the form of explicit algebraic equations. The proposed method incorporates process robustness—defined as the probability of meeting the target purities under flow-rate uncertainty due to pump errors—without requiring computationally expensive dynamic simulations. In a computational demonstration, the method achieved a joint probability of 0.960 for simultaneously attaining 99.9% purity in both extract and raffinate products.

Keywords: Adsorption, Process Design, Process Operations, Optimization, Design Under Uncertainty

INTRODUCTION

Simulated moving bed (SMB) chromatography is a continuous liquid-phase separation technique widely used in the pharmaceutical and petrochemical industries. SMB is based on the concept of true moving bed (TMB) chromatography (Fig. 1a), in which continuous separation is achieved by countercurrent contact between a liquid phase and a solid adsorbent. In SMB, the countercurrent separation in TMB is mimicked by connecting multiple chromatographic columns in a semi-closed loop and periodically switching the inlet/outlet ports (Fig. 1b). These ports serve to supply the feed and desorbent and to withdraw the two product fractions, extract and raffinate. The port roles are switched simultaneously in the direction of the internal fluid flow at fixed time intervals, referred to as the step time.

This dynamic port switching adds substantial operational complexity to SMB; consequently, it is highly challenging to determine robust operating conditions—four zone flow rates and the step time—that simultaneously enhance productivity and guarantee high purity as required in pharmaceutical applications, under unavoidable uncertainties such as pump errors. Mazzotti et al. proposed the well-known triangle theory [1], which enables

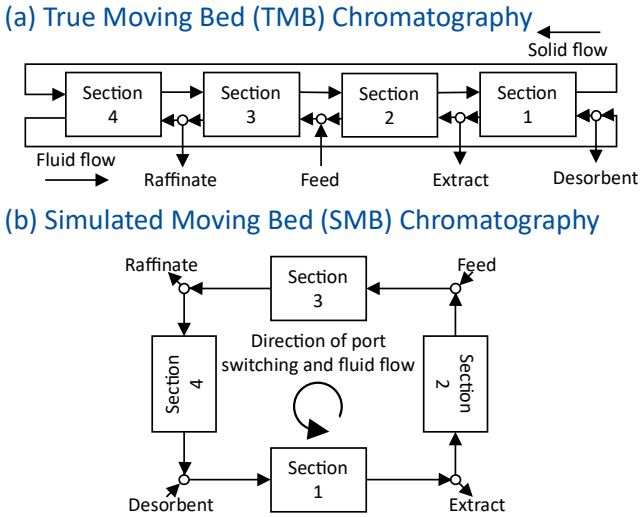

Figure 1. Schematic diagram of (a) true moving bed chromatography process and (b) simulated moving bed chromatography process.

SMB operating conditions to be determined solely from algebraic equations based on a TMB process model. Although Triangle Theory remains the most widely used method for selecting SMB operating conditions, it does not provide a systematic way to identify robust

conditions in the presence of flow-rate uncertainties caused by pump errors. Note that they pointed out that the triangle theory can account for robust operation in a trial-and-error manner. Recently, sampling-based uncertainty quantification for SMB has been explored [2]; however, uncertainty quantification alone may require a large number of dynamic simulations, making the computational cost prohibitive for robust operating-condition determination.

In this study, we propose a method to uniquely determine robust SMB operating conditions under flow-rate uncertainty, without requiring any dynamic simulations. Specifically, the robust operating conditions under flow-rate uncertainty are derived analytically as explicit algebraic equations. This formulation enables a substantial reduction in computational cost for operating-condition selection. The robustness of the derived conditions is quantified by probability of achieving the target purities for the extract and raffinate products [3] and is validated computationally using dynamic simulations. In our case study, the probability that a purity of 99.9% is simultaneously achieved in both products was 0.960.

METHODS

Derivation of Unique Operating Condition

In this section, we present a method for uniquely determining robust SMB operating conditions in the presence of flow-rate uncertainty. Building on the triangle theory [1] as the theoretical foundation, we derive explicit algebraic equations for the robust operating conditions, expressed as dimensionless flow rates of the four zones (Fig. 1a, b), which are normalized by step time. According to the triangle theory, the dimensionless flow rate in jth zone, m_j, is defined as:

$$m_j = \left(Q_j t^{\mathrm{step}} - V_{\mathrm{zone}}\epsilon\right)/V_{\mathrm{zone}}(1-\epsilon)\,, j = 1,2,3,4, \quad (1)$$

where $Q_j, t^{\mathrm{step}}, V_{\mathrm{zone}}$, and ϵ are the volumetric flow rate in jth zone, step time, total zone volume, and the total porosity, respectively. In this study, the adsorption equilibrium of the two components is described by the following competitive Langmuir isotherm:

$$q_i^j = H_i C_i^j / \left(1 + b_1 C_1^j + b_2 C_2^j\right), i = 1,2, \quad (2)$$

where q_i^j and C_i^j represent the solid-phase and the liquid-phase concentrations for ith component in jth zone, respectively; H_i and b_i refer to the Henry's constant and the affinity coefficient for ith component. Under Eq. (2), the triangle theory provides the conditions on m_j required to achieve complete separation—i.e., 100% purity for both extract and raffinate—through the following three inequalities:

$$H_2 < m_1 < \infty, \quad (3)$$

$$m_{2,\min}(m_2, m_3) < m_2 < m_3 < m_{3,\max}(m_2, m_3), \quad (4)$$

$$-\epsilon/(1-\epsilon) < m_4 < m_{4,\max}(m_2, m_3), \quad (5)$$

where $m_{2,\min}(\bullet), m_{3,\max}(\bullet)$, and $m_{4,\max}(\bullet)$ stand for algebraic functions defined by the triangle theory [1].

In this study, the uncertainty in the dimensionless flow rates of the four zones arising from pump errors is modeled by a four-dimensional multivariate normal distribution:

$$\tilde{\boldsymbol{m}} \sim \mathrm{p}(\boldsymbol{m}) = N(\boldsymbol{m}, \boldsymbol{\Sigma}), \boldsymbol{m} = [m_1, m_2, m_3, m_4]^T \quad (6)$$

where $\tilde{\boldsymbol{m}}$ denotes the uncertain (random) vector of dimensionless flow rates, and $\boldsymbol{m}$ is its mean vector. The matrix $\boldsymbol{\Sigma}$ is the covariance matrix; its entries $\Sigma_{k,l}$ are not necessarily zero. Derivation of $\boldsymbol{\Sigma}$ can be found in the literature [4].

The robust operating condition $\boldsymbol{m}^*$—i.e., the setpoint flow-rate vector to be specified for the pumps while accounting for pump errors—is determined by the following multiobjective optimization problem:

$$m_1^* = \arg\min_{\boldsymbol{m}} m_1 - H_2 \quad (7)$$

$$(m_2^*, m_3^*) = \arg\max_{\boldsymbol{m}} m_3 - m_2 \quad (8)$$

$$m_4^* = \arg\min_{\boldsymbol{m}} m_{4,\max}(m_2^*, m_3^*) - m_4 \quad (9)$$

s.t. all set of $\boldsymbol{m}$ in the following Eq. (10) satisfy Eqs. (3)-(5) simultaneously:

$$(\tilde{\boldsymbol{m}} - \boldsymbol{m})^T \boldsymbol{\Sigma}^{-1} (\tilde{\boldsymbol{m}} - \boldsymbol{m}) \le \chi^2_{n=4}(\mu). \quad (10)$$

Here, the three objective functions in Eqs. (7)–(9) represent desorbent consumption, productivity, and enrichment, respectively. In Eq. (10), $\chi^2_{n=4}(\mu)$ is the μ-quantile of the chi-square distribution and n refers to the degrees of freedom. Because $\mu \in [0,1)$, the conservativeness can be arbitrarily adjusted by controlling how much flow-rate uncertainty is considered. Note that this optimization problem cannot be solved analytically and therefore requires a numerical solution. In this study, we derive an analytical solution by reformulating it as a two-stage optimization [4]: Eq. (8) is treated as the primary objective function, while Eqs. (7) and (9) are considered as subsequent objectives.

Consequently, we derive a closed-form expression for the robust operating condition $\boldsymbol{m}^*$, given by the following explicit algebraic equations:

$$\begin{bmatrix} m_2^* \\ m_3^* \end{bmatrix} = -\boldsymbol{A}^{-1}\left(\boldsymbol{\Gamma}_{23} + \boldsymbol{n}_{23}\sqrt{\chi^2_{n=2}(\mu)}\right), \quad (11)$$

$$\begin{bmatrix} m_1^* \\ m_4^* \end{bmatrix} = \boldsymbol{\Gamma}_{14} + \boldsymbol{n}_{14}\sqrt{\chi^2_{n=2}(\mu)}, \quad (12)$$

$$\boldsymbol{A} = \begin{bmatrix} H_2 - H_1\left(1 + b_2 C_2^F\right) & b_2 C_2^F H_1 \\ H_2 - \omega_G\left(1 + b_2 C_2^F\right) & b_2 C_2^F \omega_G \end{bmatrix}, \quad (13)$$

$$\boldsymbol{\Gamma}_{23} = \begin{bmatrix} -H_1(H_2 - H_1) \\ -\omega_G(H_2 - \omega_G) \end{bmatrix}, \tag{14}$$

$$\boldsymbol{n}_{23} = \begin{bmatrix} \sqrt{[A_{11}, A_{12}]\begin{bmatrix} \Sigma_{22} & \Sigma_{23} \\ \Sigma_{32} & \Sigma_{33} \end{bmatrix}[A_{11}, A_{12}]^T} \\ \sqrt{[A_{21}, A_{22}]\begin{bmatrix} \Sigma_{22} & \Sigma_{23} \\ \Sigma_{32} & \Sigma_{33} \end{bmatrix}[A_{21}, A_{22}]^T} \end{bmatrix}, \tag{15}$$

$$\boldsymbol{\Gamma}_{14} = \begin{bmatrix} H_2 \\ m_{4,\max}(m_2^*, m_3^*) \end{bmatrix}, \tag{16}$$

$$\boldsymbol{n}_{14} = \begin{bmatrix} \sqrt{\Sigma_{11}} \\ \sqrt{\Sigma_{44}} \end{bmatrix}, \tag{17}$$

where C_i^F represents feed concentration for ith component and ω_G is a solution of quadratic equation given by the triangle theory [1]. Thus, without performing computationally expensive SMB dynamic simulations, the robust operating condition $\boldsymbol{m}^*$ can be determined simply by substituting the a priori identifiable parameters C_i^F, H_i, b_i, Σ, and μ into Eqs. (11) and (12).

Process Robustness Quantification

In this study, SMB process robustness is defined as the probability of achieving purities at or above the specified target purity, and this metric is used to evaluate the robust operating conditions $\boldsymbol{m}^*$ determined by Eqs. (11) and (12). This quantified process robustness is computed from the joint predictive distribution of the extract and raffinate purities, $\mathrm{p}(P_E, P_R)$, obtained under the flow-rate uncertainty $\mathrm{p}(\boldsymbol{m})$, as follows:

$$\mathrm{p}(P_E, P_R) = \int_{\mathcal{M}} \mathrm{p}(P_E, P_R|\boldsymbol{m})\mathrm{p}(\boldsymbol{m})d\boldsymbol{m} \tag{18}$$

where P_E and P_R refer to extract and raffinate purities, respectively; $\mathcal{M}$ denotes the uncertainty region defined by $\mathrm{p}(\boldsymbol{m})$; $\mathrm{p}(P_E, P_R|\boldsymbol{m})$ is a probabilistic model based on the SMB process model, which is formulated as a system of coupled partial differential–algebraic equations [3]. Additionally, the predictive distribution of the purity for a single product can be obtained by marginalizing the joint distribution $p(P_E, P_R)$, i.e., by integrating out the other purity variable, as follows:

$$\mathrm{p}(P_s) = \int_0^{100} \mathrm{p}(P_E, P_R)dp_{\hat{s}}, (s, \hat{s}) = (E, R), (R, E). \tag{19}$$

The process robustness is computed by integrating the joint predictive distribution $\mathrm{p}(P_E, P_R)$ from the specified target purities r_E and r_R up to 100 as follows:

$$\mathrm{P}(P_E \geq r_E, P_R \geq r_R) = \int_{r_E}^{100} \int_{r_R}^{100} \mathrm{p}(P_E, P_R)dP_E dP_R, \tag{20}$$

where $\mathrm{P}(P_E \geq r_E,\ P_R \geq r_R)$ is the joint probability that the purities of extract and raffinate simultaneously meet or exceed the target values r_E and r_R, respectively. Similarly to Eq. (19), the probability for a single product purity (i.e., the marginal probability) is obtained as follows:

$$\mathrm{P}(P_s \geq r_s) = \int_{r_s}^{100} \mathrm{p}(P_s)dP_s, s = R, E. \tag{21}$$

RESULTS

We determined the robust operating conditions $\boldsymbol{m}^*$ using Eqs. (11) and (12), and quantified the process robustness of $\boldsymbol{m}^*$ by the marginal probability (Eq. (21)) or the joint probability (Eq. (20)). The SMB scale and model parameters were taken from the study by Grosfils *et al.* [5], as shown in Table 1. Here, σ_{pump} denotes the standard deviation of the pump error used to construct the covariance matrix $\boldsymbol{\Sigma}$ in Eq. (6). Details of the derivation of $\boldsymbol{\Sigma}$ from σ_{pump} are available in the literature [4].

Table 1: Parameters used in this study [5].

parameter	value	parameter	value
C_1^F [g/L]	1.20	C_2^F [g/L]	1.20
H_1 [-]	3.23	H_2 [-]	6.50
b_1 [L/g]	0.210	b_2 [L/g]	0.422
ϵ [-]	0.620	V_{zone} [mL]	141
t^{step} [s]	300	σ_{pump} [mL/s]	0.124

When the conservativeness parameter μ is varied over $[0, 1)$, Eqs. (11) and (12) determine the four dimensionless flow rates $\boldsymbol{m}^*$ by reconciling the three objective functions in Eqs. (7)–(9), in particular at the expense of the primary objective, productivity (Eq. (8)). In this study, μ was set to six values: 0, 0.05, 0.25, 0.5, 0.75, and 0.99. Note that at $\mu = 0$, Eqs. (11) and (12) reduce exactly to the theoretical optimum in the triangle theory. As can be seen from Table 2, increasing μ leads to a simultaneous increase in m_1^*, m_2^*, and m_4^*, and a decrease in m_3^*; i.e. increasing μ results in higher desorbent consumption, lower productivity, and lower enrichment.

Table 2: Determined robust operating conditions $\boldsymbol{m}^*$ by the proposed method (Eqs. (11) and (12)). Note that $\boldsymbol{m}^*$ are shown to three decimal places for comparison across different μ values.

μ	m_1^*	m_2^*	m_3^*	m_4^*
0	6.500	2.491	3.968	2.556
0.05	6.547	2.609	3.970	2.555
0.25	6.611	2.771	3.972	2.556
0.50	6.672	2.925	3.974	2.560
0.75	6.744	3.105	3.977	2.570
0.99	6.944	3.609	3.984	2.644

For the target purities $r_E = r_R = 99.9$, the proposed method (Eqs. (11) and (12)) enabled the determination of robust operating conditions $\boldsymbol{m}^*$ with an arbitrary level of conservativeness by varying μ. Using dynamic simulations of the SMB process model [3], we estimated Eqs. (18) and (19) and quantified the process robustness of each of the six $\boldsymbol{m}^*$ in Table 2 using Eqs. (20) and (21). As shown in Table 3, $\mathrm{P}(P_E \geq r_E), \mathrm{P}(P_R \geq r_R)$, and $\mathrm{P}(P_E \geq r_E,\ P_R \geq$

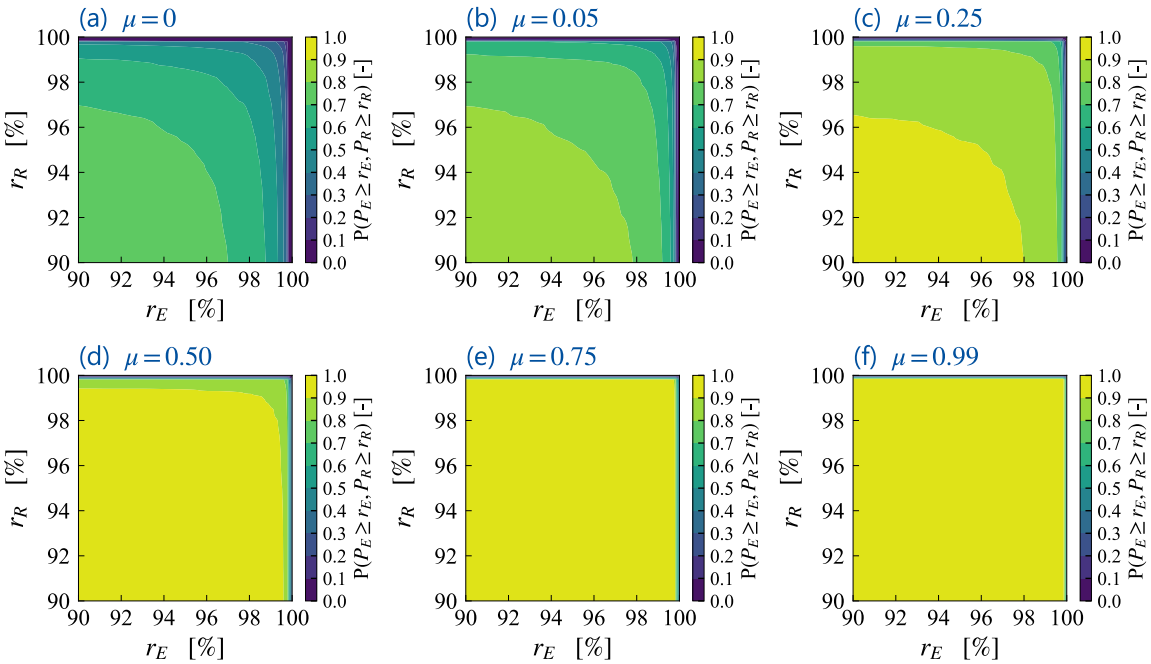

Figure 3: Comparison of six robust operating conditions by a contour plot of $P(P_E \geq r_E, P_R \geq r_R)$: (a) $\mu = 0$. (b) $\mu = 0.05$. (c) $\mu = 0.25$. (d) $\mu = 0.50$. (e) $\mu = 0.75$. (f) $\mu = 0.99$. The color bar on the right side of each contour plot refers to the value of $P(P_E \geq r_E, P_R \geq r_R)$.

r_R) all increase with μ. Particularly, at $\mu = 0.99$, all values exceed 0.950. Thus, $\mu = 0.99$ yields a robust operating condition $\boldsymbol{m}^*$ for $r_E = r_R = 99.9$.

Table 3: Values of process robustness at μ quantified by Eqs. (20) and (21) at the target purity as $r_E = r_R = 99.9$.

μ	$P(P_E \geq 99.9)$	$P(P_R \geq 99.9)$	$P\begin{pmatrix} P_R \geq 99.9, \\ P_E \geq 99.9 \end{pmatrix}$
0	0.113	0.133	0.0104
0.05	0.268	0.239	0.0646
0.25	0.485	0.436	0.245
0.50	0.651	0.617	0.462
0.75	0.802	0.789	0.689
0.99	0.973	0.978	0.960

When examining the response of the process robustness with the target purity treated as a variable, the trend of robustness of extract purity, $P(P_E \geq r_E)$, with respect to the target value r_E changes markedly depending on μ. As observed in Fig. 2a, increasing μ transforms $P(P_E \geq r_E)$ from a gently varying curve into a more sharply curved profile that remains close to $P(P_E \geq r_E) = 1.0$ over a wide range of r_E. This change in curvature indicates that the robustness of the extract purity is strongly influenced by μ.

In contrast, the robustness of the raffinate purity, $P(P_R \geq r_R)$, with respect to its target value r_R is less sensitive to μ than $P(P_E \geq r_E)$. As shown in Fig. 2b, even as μ is varied from 0 to 0.99, the overall profile remains similar: $P(P_R \geq r_R)$ decreases linearly for $r_R < 98.0$ and then exhibits a sharp drop at $r_R \approx 99.0$. The difference between the shapes of the $P(P_E \geq r_E)$and $P(P_R \geq r_R)$ curves in Figs. 2a and 2b arises from the shape of the shock fronts in the internal concentration profiles of the SMB process, as discussed in detail in the literature [3].

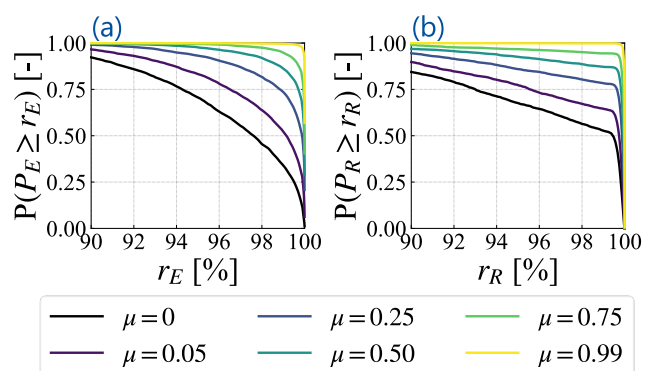

Figure 2. Comparison of process robustness with μ. (a) Change of robustness $P(P_E \geq r_E)$ through the target extract purity r_E. (b) Change of robustness $P(P_R \geq r_R)$ through the target raffinate purity r_R.

By comparing the joint probability $P(P_E \geq r_E, P_R \geq r_R)$ across the six μ values, we found that an appropriate choice of μ for efficient SMB operation under the flow-

rate uncertainty depends on the target purities r_E and r_R. As shown in Fig. 3, an area where $\mathrm{P}(P_E \geq r_E,\, P_R \geq r_R) \geq 0.90$ first appear at $\mu = 0.25$ (Fig. 3c). With further increases in μ, the area expands; however, at $\mu = 0.75$ (Fig. 3e) most of the domain already attains $\mathrm{P}(P_E \geq r_E,\, P_R \geq r_R) \geq 0.90$, and little additional expansion is observed when increasing μ to 0.99 (Fig. 3f). Although, as shown in Table 2, the difference between $\mu = 0.75$ and $\mu = 0.99$ is pronounced for stringent targets of 99.9%, this difference becomes negligible for moderate targets such as 95.0 or 99.0. These results indicate that μ should be selected by simultaneously considering the desired process robustness at the specified target purities and the other performance metrics in Eqs. (7)–(9).

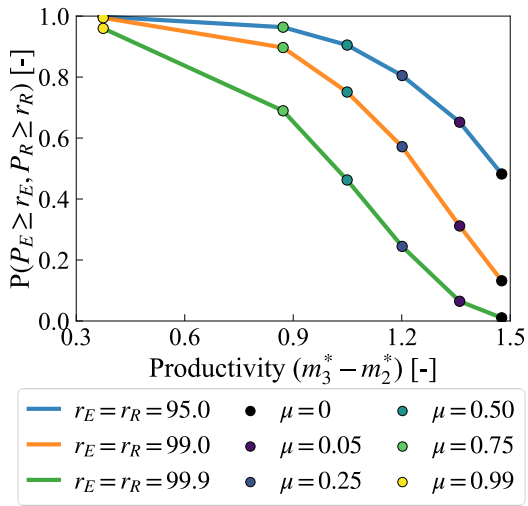

Figure 4. Trade-off relationship between productivity and robustness $\mathrm{P}(P_E \geq r_E, P_R \geq r_R)$ at three target purities.

The choice of an appropriate μ should be guided by the trade-off between process robustness (Eq. (20)) and productivity, which is treated as the primary objective function (Eq. (8)). As shown in Fig. 4, the trade-off curve for $r_E = r_R = 99.9$ is nearly linear, indicating that robustness can be increased at the expense of productivity in an approximately proportional manner. In contrast, the trade-off curves for targets of 95.0 and 99.0 are convex toward the upper right, implying diminishing returns: increasing μ beyond 0.75 yields little improvement in robustness. In particular, for the 95.0 target, the joint probability already satisfies $\mathrm{P}(P_E \geq 95.0,\, P_R \geq 95.0) \geq 0.90$ at $\mu = 0.50$. These results highlight that, within the proposed framework, selecting μ appropriately under the robustness–productivity trade-off is essential and should be tailored to the desired target purities.

CONCLUSION

We proposed a methodology to determine robust SMB operating conditions under flow-rate uncertainty in the form of explicit algebraic equations. We further evaluated the robustness and discussed an appropriate conservativeness parameter. Experimental validation of the method constitutes an important direction for future work.

ACKNOWLEDGEMENTS

This work was supported by the Japan Society for the Promotion of Science (JSPS) Grants-in-Aid for Scientific Research, Grant Number JP25K23521.

AUTHOR IDENTIFIERS

Author ORCIDs:
Kensuke Suzuki: 0000-0002-3628-5314
Tomoyuki Yajima: 0000-0002-7240-1693
Yoshiaki Kawajiri: 0000-0002-7124-1704

REFERENCES

1. Mazzotti M, Storti G, Morbidelli M. Optimal operation of simulated moving bed units for nonlinear chromatographic separations. Journal of Chromatography A 769:3-24 (1997). https://doi.org/10.1016/s0021-9673(97)00048-4
2. He Q, Zhao L. Bayesian inference based process design and uncertainty analysis of simulated moving bed chromatographic systems. Separation and Purification Technology 246:116856 (2020). https://doi.org/10.1016/j.seppur.2020.116856
3. Suzuki K, Yajima T, Kawajiri Y. Process robustness evaluation for various operating configurations of multi-column chromatography processes with nonlinear isotherm. Chemical Engineering Science 309:121395 (2025). https://doi.org/10.1016/j.ces.2025.121395
4. Suzuki K, Yajima T, Kawajiri Y. Quick robust design for simulated moving bed chromatography under comprehensive uncertainty via robust triangle theory. Separation and Purification Technology 391:136979 (2026). https://doi.org/10.1016/j.seppur.2026.136979
5. Grosfils V, Levrie C, Kinnaert M, Vande Wouwer A. A systematic approach to SMB processes model identification from batch experiments. Chemical Engineering Science 62:3894-3908 (2007). https://doi.org/10.1016/j.ces.2007.04.015

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Design of Fluid Distribution Devices Using Topology Optimization

Osamu Tonomura*, Shunya Doi, Naohiro Akashi, and Ken-ichiro Sotowa

Kyoto University, Department of Chemical Engineering, Kyoto, Japan
* Corresponding Author: tonomura@cheme.kyoto-u.ac.jp.

ABSTRACT

A topology optimization approach is applied to the design of a compact fluid distribution device that achieves uniform flow distribution among parallel channels while minimizing the mean residence time under a pressure loss constraint. Such devices are widely used in applications including parallel microreactors, thermal management systems, and energy devices, where both compactness and precise flow control are required. The design problem is formulated using a density-based method in which the design variable represents the local void fraction of a fictitious porous medium. The governing equations are based on the incompressible Navier–Stokes equations with an additional Darcy-type resistance term. The objective function is defined as the spatially averaged design variable, corresponding to the minimization of the effective flow passage volume and thus the mean residence time. Constraints are imposed on both flow rate uniformity among parallel channels and the total pressure loss. A two-stage topology optimization strategy is proposed to obtain physically realizable geometries. The results show that the proposed strategy effectively removes non-functional regions while preserving flow uniformity, leading to compact designs. Furthermore, numerical investigations reveal that designs with fewer early-stage bifurcations tend to achieve smaller flow passage areas under identical pressure loss conditions. These findings demonstrate that topology optimization provides a rational and flexible framework for the systematic design of compact fluid distribution devices.

Keywords: Topology optimization, Fluid distribution device, Mean residence time, Flow uniformity, Pressure loss constraint, Parallel channels

INTRODUCTION

The performance of many engineering systems strongly depends on their geometrical configuration, yet the optimal shape is rarely obvious. In such cases, systematic design methods based on mathematical optimization provide a powerful alternative to conventional trial-and-error approaches. Among various structural optimization techniques, topology optimization has attracted increasing attention because it allows the material distribution within a prescribed design domain to be optimized, enabling the generation of novel and non-intuitive structures.

In recent years, topology optimization has been successfully extended from structural mechanics to fluid mechanics. Several studies have demonstrated the potential of topology optimization for fluid flow problems, including flow channel design, heat transfer enhancement, and microfluidic devices. Comprehensive reviews of topology optimization methods for fluid-based systems can be found in recent literature [1, 2]. Since the pioneering work on topology optimization for Stokes flow [3], the method has been successfully applied to a wide range of fluid-related problems, including channel flow optimization [4] and catalytic microreactors [5]. These studies have shown that topology optimization can significantly enhance performance compared with conventional heuristic designs.

Fluid distribution devices, which distribute an inlet flow into multiple parallel channels, are widely used in various applications such as parallel microreactors, fuel cell stacks, heat exchangers, and thermal management systems [6, 7]. In such systems, uniform flow distribution among channels is crucial for achieving high performance and operational safety. Uneven flow distribution can lead to issues such as local overheating, reduced efficiency,

and degradation of device lifetime [6, 7].

Although various design strategies for fluid distribution devices have been proposed, including geometrical optimization based on CFD simulations [8], biomimetic or fractal-inspired designs [9], and empirically optimized manifold structures [10], most conventional approaches rely on predefined geometrical configurations and parametric tuning. As a result, the design freedom is inherently limited, and the obtained solutions may not be globally optimal.

To overcome these limitations, topology optimization offers a promising framework for the rational design of fluid distribution devices. A previous study demonstrated the feasibility of achieving target flow distributions using topology optimization [11]. However, this study primarily focused on objectives such as energy dissipation minimization, while the compactness of the device, which is closely related to the mean residence time, has received less attention. In practical applications, the mean residence time is a key performance indicator because it directly affects reactor productivity, system responsiveness, and overall device size. Therefore, designing a compact fluid distribution device while maintaining uniform flow distribution is of great practical importance.

In the present study, topology optimization is applied to the design of a fluid distribution device that achieves uniform flow distribution among parallel channels while minimizing the mean residence time under a pressure loss constraint. The optimization problem is formulated using a density-based method, and a two-stage optimization strategy is introduced to obtain physically realizable geometries. Furthermore, the effects of pressure loss constraints, initial design variables, and the number of parallel channels on the optimized designs are systematically investigated.

PROBLEM FORMULATION AND METHODOLOGY

Design problem

In this study, the topology optimization of a fluid distribution device is conducted using the same two-dimensional geometry as that employed in [11], in order to enable a direct comparison with the previous study.

A schematic of the computational domain is shown in Fig. 1. The design domain is defined between a single inlet and multiple outlet channels, and topology optimization is performed only within this region. The inlet velocity is specified as a uniform profile, and pressure boundary conditions are imposed at the outlets. No-slip boundary conditions are applied on all solid walls.

The optimization problem is formulated such that the objective function corresponds to the minimization of the spatially averaged design variable over the design domain. Since the design variable represents the local void fraction of the porous medium, minimizing its average value leads to a reduction in the effective flow passage volume, which is equivalent to minimizing the mean residence time of the fluid. Two constraints are imposed in the optimization problem. The first constraint ensures uniformity of flow rates among all outlet channels, which is evaluated by the normalized deviation of individual outlet flow rates from the target value. The second constraint is imposed on the total pressure loss between the inlet and the outlets to ensure practical operating conditions. Accordingly, the design problem is formulated as a constrained optimization problem with the objective of minimizing the mean residence time while satisfying the constraints on flow rate uniformity and pressure loss.

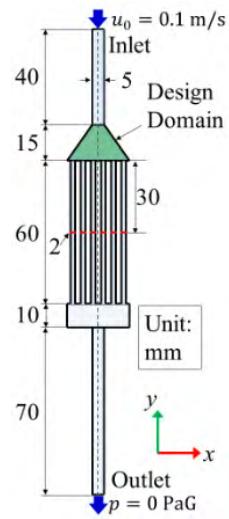

Figure 1. Computational domain and design domain [11].

Topology optimization method

In this study, a density-based topology optimization method is employed to determine the optimal material distribution within the design domain. The design variable is defined as a continuous scalar field that represents the local void fraction of a fictitious porous medium, taking values between 0 (solid region) and 1 (fluid region).

To model the presence of the porous medium, a Darcy-type resistance term is introduced into the momentum equations. The inverse permeability is interpolated as a function of the design variable, such that high resistance is imposed in solid-like regions and negligible resistance is applied in fluid-like regions. This approach enables a smooth transition between solid and fluid regions during the optimization process.

To avoid numerical instabilities such as checkerboard patterns and mesh dependency, a PDE-based filtering technique is applied to the design variable. Furthermore, a smoothed Heaviside projection is employed to promote clear separation between solid and fluid regions and to reduce gray-scale regions in the optimized designs.

The optimization problem is solved using a gradient-based algorithm. Sensitivities of the objective function and constraints with respect to the design variables are evaluated using the adjoint method. The design variables are then updated iteratively until convergence is achieved.

Governing equations

The fluid flow in the design domain is modeled as a steady, incompressible laminar flow. The governing equations are the incompressible Navier–Stokes equations with an additional Darcy-type resistance term to represent the effect of the fictitious porous medium.

The continuity equation is given by

$$\nabla \cdot \boldsymbol{u} = 0 \tag{1}$$

where $\boldsymbol{u}$ denotes the velocity vector. The momentum equation is expressed as

$$\rho(\boldsymbol{u} \cdot \nabla)\boldsymbol{u} = -\nabla p + \mu \nabla^2 \boldsymbol{u} - \alpha(\gamma)\boldsymbol{u} \tag{2}$$

where p is the pressure, ρ is the fluid density, and μ is the dynamic viscosity. The last term represents the Darcy-type resistance introduced to model the porous medium. The inverse permeability $\alpha(\gamma)$ is defined as a function of the design variable γ, such that high resistance is imposed in solid-like regions and low resistance is imposed in fluid-like regions.

Specifically, the inverse permeability is interpolated using a convex function of the design variable, which allows a smooth transition between solid and fluid phases during the optimization process. This formulation enables the flow field to be continuously modified as the design variable evolves.

These governing equations are solved numerically using the finite element method at each optimization iteration to evaluate the objective function and constraints.

Objective function and constraints

The optimization problem is formulated to minimize the mean residence time of the fluid while satisfying constraints on flow rate uniformity and pressure loss.

As described in the previous section, the design variable γ represents the local void fraction of the fictitious porous medium. Therefore, the mean residence time is proportional to the effective flow passage volume in the design domain. In this study, the objective function is defined as the spatially averaged design variable over the design domain Ω_d, given by

$$J(\gamma) = \frac{1}{\Omega_d} \int_{\Omega_d} \gamma d\Omega \tag{3}$$

where $|\Omega_d|$ denotes the area of the design domain. Minimizing this objective function corresponds to minimizing the effective flow volume and thus the mean residence time.

To ensure uniform flow distribution among the outlet channels, a constraint on the outlet flow rates is imposed. The flow rate through the i-th outlet, Q_i, is evaluated by integrating the normal velocity over the outlet boundary. The flow uniformity constraint is defined based on the normalized deviation of each outlet flow rate from the target value Q_{target}, expressed as

$$g_1 = \frac{1}{N} \sum_{i=1}^{N} \left[\frac{Q_i - Q_{\text{target}}}{Q_{\text{target}}} \right] \leq \varepsilon \tag{4}$$

where N is the number of outlet channels and ε is a prescribed tolerance.

In addition, a constraint is imposed on the total pressure loss between the inlet and the outlets to ensure practical operating conditions. The pressure loss Δp is defined as the difference between the spatially averaged inlet pressure and the averaged outlet pressure, and the constraint is expressed as

$$g_2 = \Delta p \leq \Delta p_{\max} = \beta \Delta p_0 \tag{5}$$

where $\Delta p_{\max}$ denotes the allowable maximum pressure loss, Δp_0 is the pressure loss that would be obtained if the entire design domain is completely open to flow, and β is the design parameter.

Accordingly, the optimization problem is formulated as a constrained minimization problem with the objective of minimizing the mean residence time subject to the constraints on flow rate uniformity and pressure loss.

Optimization procedure

The optimization problem is solved using a gradient-based optimization algorithm. At each iteration, the governing equations described in the previous section are solved to obtain the velocity and pressure fields, from which the objective function and constraint values are evaluated. All simulations were performed using COMSOL Multiphysics® 6.0, which provides a finite element framework for solving the governing equations and implementing topology optimization.

The sensitivities of the objective function and constraints with respect to the design variables are computed using the adjoint method. Gradient-based optimization methods are widely used in topology optimization because they efficiently handle large numbers of design variables through adjoint sensitivity analysis, making them computationally more efficient than population-based optimization methods for high-dimensional design problems. This approach enables efficient sensitivity evaluation with a computational cost that is independent of the number of design variables, making it suitable for large-scale topology optimization problems.

The design variables are updated iteratively using the Method of Moving Asymptotes (MMA), which is widely used in topology optimization due to its robustness and numerical stability. During the optimization process, the filtering and projection techniques described in

the previous section are applied at each iteration to ensure mesh-independent solutions and to promote clear separation between fluid and solid regions.

The optimization is terminated when the relative change in the objective function over successive iterations becomes sufficiently small, indicating convergence of the solution. The resulting optimized distribution of the design variable is then interpreted to obtain the final geometry of the fluid distribution device.

RESULTS AND DISCUSSION

This section presents the results of the topology optimization of the fluid distribution device. First, the effectiveness of the proposed two-stage optimization strategy is demonstrated. Subsequently, the effects of pressure loss constraints, initial design variables, and the number of parallel channels on the optimized designs are systematically investigated.

Two-stage topology optimization

In this study, water (298K) was used as the working fluid. The values N, ε, and β were set to 6, 0.001, and 1.25, respectively. The initial value of the design variable, γ_0, was set to 0.5. The topology optimization problem was solved using COMSOL Multiphysics® 6.0. The resulting solution is referred to as the first-stage optimization result. The obtained distribution of the design variable is shown in Fig. 2 (left), where red and blue regions correspond to fluid and solid phases, respectively. The black lines indicate streamlines.

In the region from the inlet to just upstream of the first bifurcation point, the streamlines are nearly parallel to the y-axis and the velocity component in the x-direction is negligible. This indicates that the flow in this region can be regarded as fully developed and that the region does not contribute to the flow distribution process. From the viewpoint of reducing the mean residence time, such a region should therefore be eliminated.

Based on the first-stage optimization result, a clear channel geometry without gray-scale regions was extracted following the procedure proposed in Ref. [9]. In addition, the region identified as unnecessary from the viewpoint of mean residence time reduction was removed, and the geometry shown in Fig. 2 (right) was generated. However, subsequent CFD analysis confirmed that this geometry no longer satisfied the flow uniformity constraint that was achieved in the first-stage result.

Therefore, the geometry shown in Fig. 2 (right) was adopted as a new design domain, and the second-stage topology optimization was performed. The optimized result obtained after the second stage is shown in Fig. 3. CFD analysis confirmed that the final design satisfied the flow uniformity constraint. As a result, a fluid distribution device was designed with a reduction of 76.3% in the design domain area compared with the initial configuration.

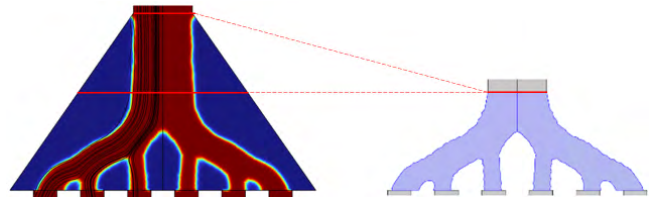

Figure 2. First-stage optimization result (left) and the design space generated for second-stage optimization (right).

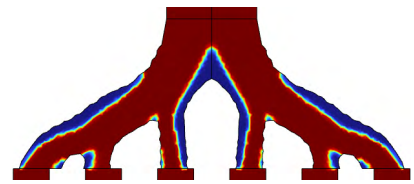

Figure 3. Second-stage optimization result.

These results demonstrate that the proposed two-stage topology optimization strategy is effective for obtaining physically realizable and compact fluid distribution device designs.

A comparison with the results reported in Ref. [11] provides further insight into the characteristics of the optimized designs. In Ref. [11], the optimization was performed using the dissipation function as the objective function, whereas the present study employs the mean residence time as the objective. Although both studies yield designs that initially bifurcate into two branches, noticeable differences can be observed in the resulting geometries. The design reported in Ref. [11] exhibits wider channel widths and a shorter distance from the inlet to the first bifurcation. Moreover, the flow is gradually redirected, resulting in a smooth branching into six channels. These features can be considered advantageous from the viewpoint of minimizing the dissipation function, as they contribute to reducing flow resistance and energy loss. In contrast, the present study aims to reduce the effective flow passage volume by minimizing the mean residence time. This objective promotes the formation of narrower channels and longer straight sections prior to branching, leading to a more compact distribution region. These differences highlight that the resulting topology is strongly influenced by the choice of the objective function, and demonstrate the flexibility of topology optimization in tailoring designs to specific performance requirements.

Effects of initial design variables and pressure loss constraint

In topology optimization problems, the obtained solution often depends on the initial values of the design variables, and the optimization process may converge to

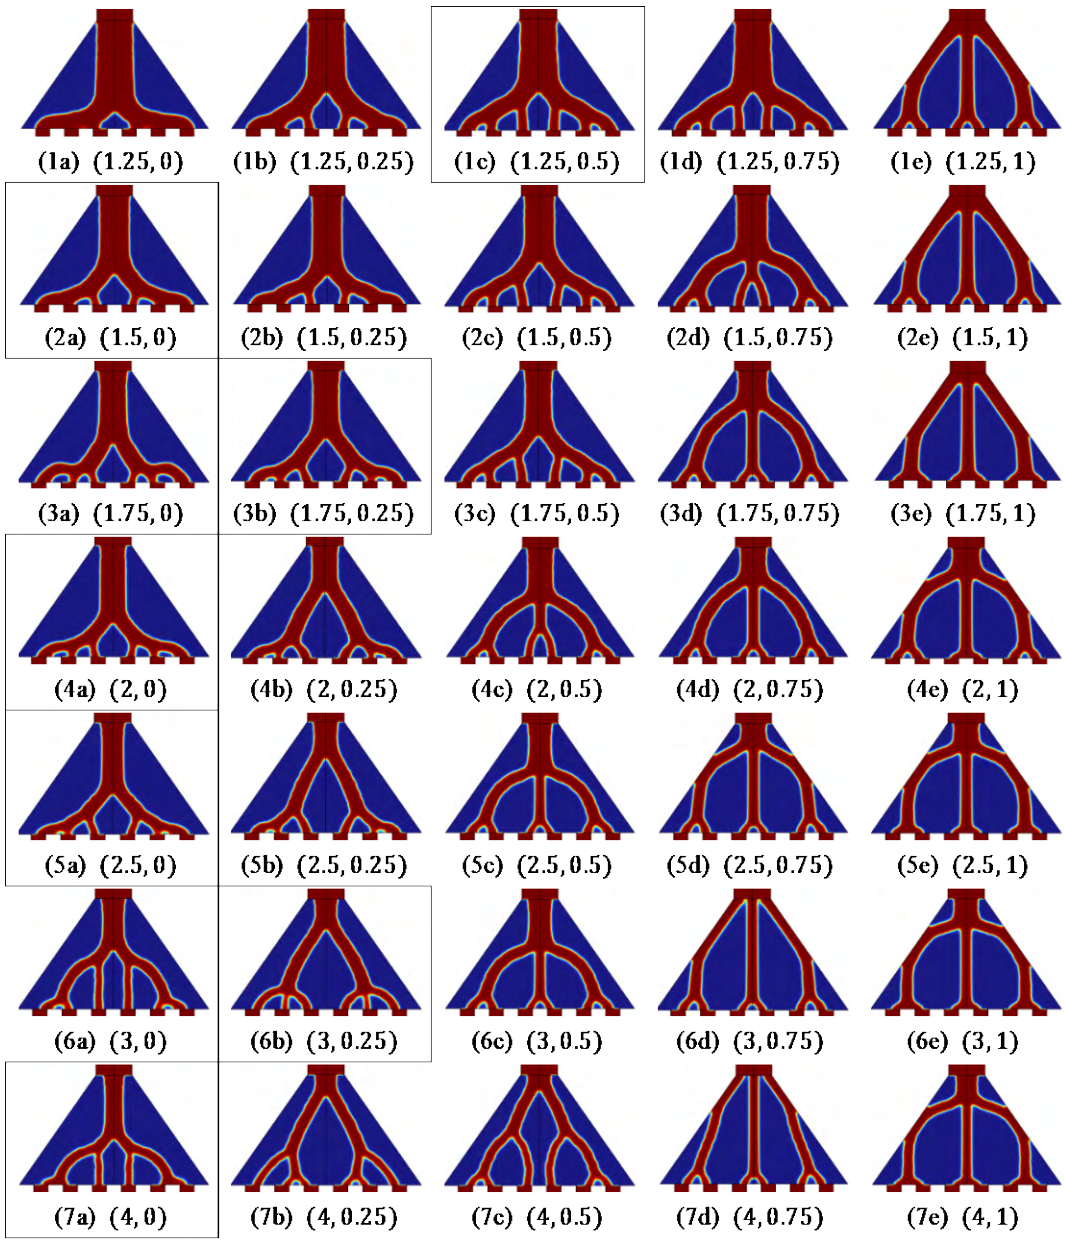

Figure 4. Optimization results for all combinations of seven values of β (= 1.25, 1.5, 1.75, 2, 2.5, 3, 4) and five values of γ_0 (= 0, 0.25, 0.5, 0.75, 1). The red regions represent the fluid, and the blue regions represent the solid. The designs highlighted with black frames correspond to the cases that achieved the minimum objective value for each β. The legend indicates (β, γ_0).

local optima. In addition, the flow channel structures obtained under different pressure loss constraints may exhibit characteristics different from those of the design derived in the previous section.

In this section, the first-stage optimization described in the previous section was performed for a total of 35 cases, consisting of five different initial values of the design variable $\gamma_0$and seven different values of β.

The resulting optimized channel structures are summarized in Fig. 4. As shown in the figure, the obtained flow channel structures can be classified into two groups based on their topological features: those that first bifurcate into two branches (e.g., Fig. 4(1d), hereafter referred to as Group A) and those that first bifurcate into three branches (e.g., Fig. 4(1e), hereafter referred to as Group B). It is found that the designs belonging to Group A tend to exhibit smaller objective function values than those belonging to Group B. Moreover, for each value of β, the design with the minimum objective function value always belongs to Group A. These results indicate that, under

identical pressure loss conditions, the channel geometries in Group A tend to have smaller flow passage areas in the branching region compared with those in Group B.

Subsequently, for all 35 first-stage optimization results, channel geometries were generated by eliminating gray-scale regions and removing flow-irrelevant regions using the same procedure as described in the previous section. Among these cases, the geometry generated from the result shown in Fig. 4(4a) exhibited the smallest flow passage area. A second-stage topology optimization was then performed using this geometry as the new design domain. As a result, a fluid distribution device with a reduction of 79.6% in the initial design domain area was obtained.

Effect of the number of parallel channels

In this section, the effect of the number of parallel channels on the optimized flow channel topology was investigated. Topology optimization calculations were performed for cases with $N = 4$ and $N = 8$ in order to examine how the flow channel configurations in the fluid distribution design domain change with the number of outlets. The governing equations, objective function, and constraints were identical to those used in the design problem with $N = 6$. However, the inlet flow rate was adjusted such that the flow rate per outlet channel was the same as that in the $N = 6$ case.

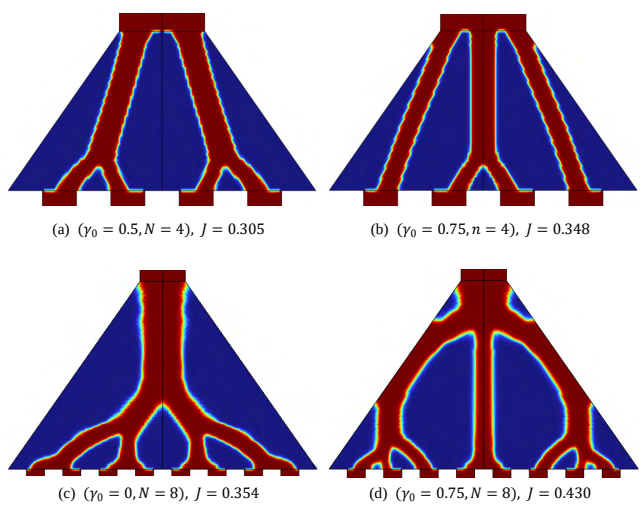

Figure 5. First-stage topology optimization results for different numbers of parallel channels.

For each value of N, first-stage topology optimization was performed for five different initial values of the design variable, $\gamma_0 = 0, 0.25, 0.5, 0.75$, and 1. For the case of $N = 4$, the distributions of the design variable corresponding to the minimum and maximum objective function values are shown in Fig. 5(a) and Fig. 5(b), respectively. The corresponding results for $N = 8$ are shown in Fig. 5(c) and Fig. 5(d).

It is observed that the designs with the minimum objective function value initially bifurcate into two branches, whereas those with the maximum objective function value initially bifurcate into three branches. This tendency is consistent with the results obtained for the $N = 6$ case, indicating that the relationship between the initial branching structure and the objective function value is preserved even when the number of parallel channels is changed.

CONCLUSIONS

In this study, a topology optimization approach was applied to the design of a compact fluid distribution device that achieves uniform flow distribution among parallel channels while minimizing the mean residence time under a pressure loss constraint. The optimization problem was formulated using a density-based method combined with incompressible Navier–Stokes equations including a Darcy-type resistance term.

A two-stage topology optimization strategy was proposed to obtain physically realizable geometries. In the first stage, an optimized material distribution was obtained, while gray-scale regions and flow-irrelevant areas were identified. In the second stage, these regions were removed and the optimization was repeated using the reduced design domain. The results demonstrated that the proposed strategy effectively eliminates non-functional regions while maintaining flow uniformity, leading to highly compact fluid distribution devices.

The effects of pressure loss constraints and initial design variables were systematically investigated through multiple optimization runs. The optimized flow channel topologies were classified into two characteristic groups based on their initial branching structures. Designs that initially bifurcate into two branches were found to yield smaller objective function values than those that initially bifurcate into three branches. This tendency was consistently observed under identical pressure loss conditions.

Furthermore, the effect of the number of parallel channels was examined by performing topology optimization for cases with four and eight outlets. The same relationship between the initial branching structure and the objective function value was observed, indicating that the identified design tendency is robust with respect to the number of parallel channels.

A comparison with a previous study using dissipation minimization as the objective function highlighted that the resulting topology strongly depends on the choice of the objective function. While dissipation minimization favors wider and smoother channels, minimizing the mean residence time promotes narrower flow passages and a more compact distribution region. These results demonstrate the flexibility of topology optimization

in tailoring fluid distribution devices to specific performance requirements.

Overall, the present study provides a systematic framework for designing compact and high-performance fluid distribution devices and offers insights into the relationship between objective functions and resulting flow channel topologies.

ACKNOWLEDGEMENTS

This work was partially supported by the Grant-in-Aid for Scientific Research (B) (No. 23H01754) and a project, Development of Continuous Production and Process Technologies of Fine Chemicals, commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

AUTHOR IDENTIFIERS

Author ORCIDs:
Tonomura O: 0000-0002-0094-519X
Sotowa K: 0000-0002-9636-8306

REFERENCES

1. Sigmund O, Maute K. Topology optimization approaches. Struct Multidisc Optim 48:1031-1055 (2013). https://doi.org/10.1007/s00158-013-0978-6
2. Alexandersen J, Andreasen CS. A review of topology optimisation for fluid-based problems. Fluids 5:29 (2020). https://doi.org/10.3390/fluids5010029
3. Borrvall T, Petersson J. Topology optimization of fluids in stokes flow. Numerical Methods in Fluids 41:77-107 (2002). https://doi.org/10.1002/fld.426
4. Gersborg-Hansen A, Sigmund O, Haber RB. Topology optimization of channel flow problems. Struct Multidisc Optim 30:181-192 (2005). https://doi.org/10.1007/s00158-004-0508-7
5. Okkels F, Bruus H. Scaling behavior of optimally structured catalytic microfluidic reactors. Phys. Rev. E 75: (2007). https://doi.org/10.1103/physreve.75.016301
6. Facão J. Optimization of flow distribution in flat plate solar thermal collectors with riser and header arrangements. Solar Energy 120:104-112 (2015). https://doi.org/10.1016/j.solener.2015.07.034
7. Kim B, Kim K, Kim S. Numerical study on novel design for compact parallel-flow heat exchanger with manifolds to improve flow characteristics. Energies 13:6330 (2020). https://doi.org/10.3390/en13236330
8. Wei M, Fan Y, Luo L, Flamant G. Cfd-based evolutionary algorithm for the realization of target fluid flow distribution among parallel channels. Chemical Engineering Research and Design 100:341-352 (2015). https://doi.org/10.1016/j.cherd.2015.05.031
9. Mazur M, Bhatelia T, Kuan B, Patel J, Webley PA, Brandt M, Pareek V, Utikar R. Additively manufactured, highly-uniform flow distributor for process intensification. Chemical Engineering and Processing - Process Intensification 143:107595 (2019). https://doi.org/10.1016/j.cep.2019.107595
10. Barbosa ISO, Santos RJ, Silva CG, Brito MSCA. Consecutive flow distributor device for mesostructured reactors and networks of reactors. Chemical Engineering and Processing - Process Intensification 167:108541 (2021). https://doi.org/10.1016/j.cep.2021.108541
11. Zeng S, Lee PS. A header design method for target flow distribution among parallel channels based on topology optimization. 2018 17th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm) :156-163 (2018). https://doi.org/10.1109/itherm.2018.8419642

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

An Extended Superstructure Formulation for Non-Isobaric Flowsheet Synthesis

Harrison A. Fraser[a], Smitha Gopinath[b], Jan Sefcik[c], George Jackson[a], Amparo Galindo[a], and Claire S. Adjiman[a*]

[a] Department of Chemical Engineering, Sargent Centre for Process Systems Engineering, Imperial College London, SW7 2AZ, London, United Kingdom
[b] School of Chemical, Materials and Biological Engineering, University of Sheffield, S1 3JD, Sheffield, United Kingdom
[c] Department of Chemical and Process Engineering, University of Strathclyde, G1 1XJ, Glasgow, United Kingdom
* Corresponding Author: c.adjiman@imperial.ac.uk

ABSTRACT

Flowsheet synthesis is an integral step in process design, entailing the selection of a set of unit operations and their connectivity to convert raw materials to products. Superstructure optimisation represents a promising class of synthesis approaches, allowing for the systematic exploration of the flowsheet design space. Despite this, many superstructure formulations suffer from numerical instabilities, combinatorial explosion, and/or rely on restrictive assumptions on the types of flowsheet alternatives that can be considered. The modified state-operator network (MSON) formalism has recently been proposed to address some of these issues for isobaric flowsheets. The constant-pressure assumption restricts the applicability of the MSON to real process applications as pressure is a key process variable in many unit operations, such as distillation, reaction, and extrusion, and is necessary to elicit flow. In this work, we present the extended MSON (E-MSON) which inherits the numerical stability of MSON, whilst removing the isobaric assumption. This is achieved through the introduction of new constraints that further improve the numerical behaviour of the MSON. The E-MSON is then applied to a simple non-isobaric superstructure optimisation problem, the results of which demonstrate that the E-MSON can serve as a framework for non-isobaric flowsheet synthesis, enabling a broader range of flowsheet alternatives to be considered.

Keywords: Superstructure Optimisation, Process Synthesis, Optimisation, Process Design, gPROMS, MINLP

INTRODUCTION

Flowsheet synthesis is a key step in the broader problem of chemical process design, involving the selection of a feasible or even optimal configuration of unit operations to convert raw materials into valuable products. Traditional approaches to flowsheet synthesis and design are hierarchical and heuristic-based [1]. In these methods, the design procedure is decomposed into increasingly detailed stages that progress from high-level structural decisions to detailed unit design. These stages rely on ad hoc rules-of-thumb to eliminate potential configurations without performing detailed design calculations [2]. Although it is relatively straightforward to apply such an approach, decomposing the design problem makes accounting for the interdependence of decisions made at different stages largely intractable [1].

With the advancement of nonlinear and mixed-integer programming, superstructure optimisation has become an increasingly promising alternative to hierarchical process synthesis techniques [3]. Superstructure optimisation first involves the representation of all candidate flowsheet topologies in a superstructure. This superstructure is then translated into a mathematical programming problem which can then be solved to yield the optimal flowsheet configuration [4], allowing for a systematic exploration of the process design space.

In general, due to the binary decision variables representing unit selection or deselection in flowsheet alternatives, and nonlinearities in the models that describe process phenomena, superstructure optimisation problems can be represented as mixed-integer nonlinear programming (MINLP) problems. How to represent the superstructure is, however, non-trivial, having a significant

influence on the solvability of the subsequent problem, and on the quality of the result obtained [1]. For instance, one commonly encountered issue which depends on the formulation of choice is combinatorial complexity.

Among superstructure representations, there exist several graph-theoretic approaches, including the State Task Network [5], State Equipment Network [6], and the State Operator Network [7]. More recent approaches instead involve abstraction of the flowsheet into building blocks, such as the phenomena building block approach of Lutze et al. [8] or the building-block superstructure presented by Demirel et al. [9]. These abstract representations facilitate intensified or novel unit operations to be yielded as a solution of the optimisation problem, but this generality comes at the expense of the solvability of the subsequent mathematical programming problem [1].

In the State Operator Network (SON) [7], which is the subject of this work, nodes of the graph represent process units, this may include unit operations, e.g., a distillation column, discrete subsets of a unit operation, e.g., a stage of a distillation column, or an intensified unit operation, e.g., a vapour compression distillation which includes the distillation column and heat pump. The edges of the graph are streams, representing material flows from unit operation to the other. To allow for material to flow between units, conceptual mixers and splitters are introduced. Each inlet of each process unit is associated with a conceptual mixer that combines the properties of connected streams (edges) into the input to the units' model. Similarly, a splitter that splits material flows across streams is defined at each outlet.

To ensure that meaningful results are obtained upon the solution of superstructure optimisation problems, rigorous and high-fidelity models must be utilised [7]. In the SON, each process unit is described by a detailed process model that not only includes mass, energy, summation of composition, and heat balance equations, detailed thermodynamic models, and kinetic models when applicable, but also costing and sizing correlations. The model equations that represent each unit are therefore coupled, nonlinear, and high-dimensional, making the solution of superstructure optimisation problems challenging.

Superstructure optimisation problems may be solved in equation-oriented optimisation paradigms, e.g., Pyomo, or by using simulation-based optimisation. Simulation-based optimisation offers the advantages of in-built-process models and thermodynamic routines, as well as initialisation strategies which can be crucial to obtaining high-quality solutions. On the other hand, equation-oriented optimisation paradigms offer great flexibility in formulating the problem and a greater choice of numerical solvers. They also provide complete access to functional forms, allowing the construction of convex relations for global optimisation. Despite this, due to the sheer number of variables and constraints in these problems, global optimisation is typically intractable, meaning superstructure optimisation problems are generally only solved to local optimality.

Another complication associated with superstructure optimisation is numerical failure. When sections of the superstructure are deselected, whether that be building blocks or entire units, mass flows into these units must be set to zero. Thus, even the seemingly simple calculation of a mole fraction can lead to a numerical singularity and hence the failure of the optimisation solver. Several examples of numerically undefined behaviour or singularity arise when process units are deselected [10].

Several approximate methods to avoid numerical issues have been presented, including the relaxation of problematic equations to continuous and differentiable approximations [11], and the introduction of small non-zero flows into deselected units [12]. Exact, systematic formulations that are guaranteed to avoid numerical issues due to unit deselection have also been proposed. Yeomans and Grossmann [13] suggested to formulate the optimisation problem using generalised disjunctive programming (GDP), representing the binary decisions of unit selection and deselection as logical disjunctions. Using the logic-based outer approximation (LBOA) algorithm, the primal problem constructed at each iteration only includes the model equations pertaining to selected units with non-zero inlet flowrates. This avoids numerical failure entirely, however, the automatic generation of the primal problem in process simulation software can be challenging, as noted by Navarro-Amorós et al. [14]. Further issues pertaining to initialisation of the master problem of the LBOA algorithm may also occur [15].

Another promising solution to the problem of numerical failure arising from unit deselection is to modify the superstructure representation itself, as performed in the modified SON (MSON), presented by Gopinath and Adjiman [15, 16]. In the MSON, mixers are modified in the to have a fictitious stream. Fictitious stream flowrates are non-zero when unit operation is deselected and zero otherwise; splitters are similarly modified, ensuring that the solution (optimal point, objective function value) is unaffected by these fictitious flows.

The MSON formulation leads to improved numerical stability, eliminating the possibility of singularities in deselected units. Applied to the design of a counter-current carbon capture column, the optimisation of an MSON superstructure yielded a high-performance solution, indicating its applicability to complex synthesis problems. The MSON may be used in both equation-oriented and simulation-based optimization paradigms. Being based on the SON [7], the MSON also inherits an implicit reduction of combinatorial complexity, as compared to task-based representations such as the State Task Network.

In the definition of both the SON [7] and MSON [15], the pressure is assumed to be constant throughout the

flowsheet. Although an extension to treat pressure as an optimisation variable has been presented for the specific case of the synthesis of counter-current columns only [16], the MSON cannot directly be applied generally to non-isobaric flowsheet synthesis without modification.

Pressure is a key operating variable in many unit operations, due to its influence on chemical and phase equilibria [17], and being essential to elicit flow. Many unit operations, such as pressure-swing distillation, membrane-based processes, and high-pressure synthesis also cannot be modelled without allowing pressure to vary across the flowsheet [18-20]. Furthermore, compression and expansion can represent significant proportions of process costs, e.g., up to 75% of the capital expenditure for refrigeration systems [21]. Rigorous flowsheet modelling also requires accurately accounting for pressure drops.

Therefore, a superstructure formulation that inherits the mathematical tractability of the MSON, whilst relaxing its isobaric assumption, is highly desirable. It would broaden the scope of flowsheet alternatives and objectives that can be considered in a superstructure optimisation problem. Thus, in this work, we present the extended MSON (E-MSON), establishing a framework which facilitates the systematic and rigorous synthesis of non-isobaric flowsheets. We demonstrate the application of E-MSON on a test problem, successfully allowing the process units to operate at different pressures. In the next section, we present the superstructure representation of the E-MSON, highlighting the variables and sets of relevance, before giving its full derivation.

E-MSON: SUPERSTRUCTURE REPRESENTATION

The E-MSON formulation avoids numerical singularities through the introduction of strictly positive fictitious streams into deselected units [15]. Each stream has a mass flowrate, $f \in \mathbb{R}^+$, composition, $\boldsymbol{q} \in [0,1]^K$, where K is the number of components considered in the problem, and is at temperature, $T \in \mathbb{R}^+$, and pressure, $P \in \mathbb{R}^+$.

Let the set of all process units in the superstructure be represented by $\mathcal{L}$, with each unit $u \in \mathcal{L}$ being described by rigorous model equations (mass and energy balances, sizing, costing, etc.). We represent the set of conditional units, referring to units that can be deselected in a particular alternative, as $\mathcal{T} \subseteq \mathcal{L}$, such that the set of permanent units that must be included in all alternatives is given by $\mathcal{L} \backslash \mathcal{T}$. The set of units include all unit operations, including pressure-change equipment such as pumps, valves, and compressors, which we represent as $\mathcal{PC} \subset \mathcal{L}$. The set of all nodes of the graph can then be defined as $\mathcal{U} = \mathcal{L} \cup \mathcal{F} \cup \mathcal{P}$, where $\mathcal{F}$ and $\mathcal{P}$ refer to the sets of feeds and products containing the material sources and sinks to and from the superstructure respectively.

Each unit $u \in \mathcal{U}$ is also associated with inlet and outlet indices corresponding to the sets $\mathcal{IN}_u$ and $\mathcal{OUT}_u$, the union of which give, for all $u \in \mathcal{U}$, the sets of all mass inlet and outlet indices $\mathcal{I}$ and $\mathcal{O}$. At each inlet and outlet, we additionally specify an isenthalpic mixer and splitter, respectively, to facilitate connectivity between units. Permanent units $u \in \mathcal{L} \backslash \mathcal{T}$ are associated with standard mixers and splitters, denoted by m_i and s_o at inlet $i \in \mathcal{I}$ and outlet $o \in \mathcal{O}$ respectively. The inlet flowrate, composition, temperature, and pressure to a unit from a mixer at inlet i are labelled as $f_i^{in}, \boldsymbol{q_i^{in}}, T_i^{in}$, and P_i^{in}, and the outlet flowrate, composition, temperature and pressure from a unit to a splitter at outlet o as $f_o^{out}, \boldsymbol{q_o^{out}}, T_o^{out}, P_o^{out}$.

As in the MSON [15, 16], conditional units are connected to modified isenthalpic mixers and splitters, labelled as $\tilde{m}_i$ and $\tilde{s}_o$. At modified mixer $\tilde{m}_i$, the flowrate of the fictitious stream, f_i^M, takes the auxiliary value, f_i^A, when the unit is deselected. Additionally, when $f_i^M > 0$, the corresponding unit's inlet composition, temperature, and pressure are constrained to take the auxiliary values, $\boldsymbol{q_i^A}, T_i^A$, and P_i^A, defined *a priori* for the fictitious stream. Analogously, at the modified splitter, $\tilde{s}_o$, the variables $f_o^S, \boldsymbol{q_o^S}, T_o^S$, and P_o^S define the fictitious stream. An example of a completely labelled conditional unit following the E-MSON formalism can be seen in Figure 1.

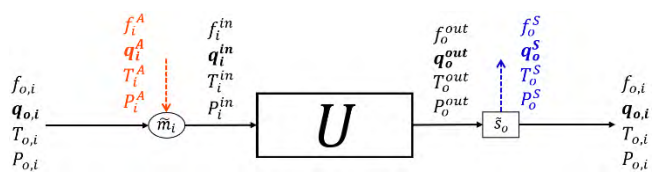

Figure 1. Example representation of a conditional unit $\underline{U}$ in the E-MSON formulation, including all variables of relevance. The fictitious streams of the modified mixer and splitter are coloured and marked by dashed lines.

The connectivity between mixers and splitters is captured in the sets $\mathcal{M}_i$ and $\mathcal{S}_o$, where $\mathcal{M}_i \subseteq \mathcal{O}$ represents the set of outlet indices connected to mixer m_i or $\tilde{m}_i$ and $\mathcal{S}_o \subseteq \mathcal{I}$ represents the set of inlet indices connected to splitter s_o or $\tilde{s}_o$. As with all other streams, the connections between mixers and splitters are parameterised in terms of flowrate, composition, temperature, and pressure, denoted as $f_{o,i}, \boldsymbol{q_{o,i}}, T_{o,i}$, and $P_{o,i}$, respectively.

To enable the selection and deselection of conditional units in the superstructure, each conditional unit $u \in \mathcal{T}$ is also associated with a binary variable z_u that describes whether it is selected in ($z_u = 1$) or deselected from ($z_u = 0$) a particular flowsheet configuration. It is important to also define $\mathcal{C}_u$, representing the set of "output" variables, that is, those variables corresponding to conditional unit $u \in \mathcal{T}$ that are included in flowsheet-level constraints or the objective function.

In the case where continuous or discrete *design*

variables are present in the model of a conditional unit (such as a reboiler duty or the number of stages in a column), additional auxiliary variables must be defined to ensure model convergence. We represent these continuous and discrete design variables in the vectors $\boldsymbol{x}_{\mathcal{A}_u}$ and $\boldsymbol{y}_{\mathcal{A}_u}, \forall\, u \in \mathcal{T}$, respectively, and store the auxiliary values similarly in the vectors $\boldsymbol{x}^{\boldsymbol{A}}_{\mathcal{A}_u}$ and $\boldsymbol{y}^{\boldsymbol{A}}_{\mathcal{A}_u}, \forall\, u \in \mathcal{T}$. As with other auxiliary variables in the E-MSON approach, these should be assigned values in advance such that, in combination with $f_i^A, \boldsymbol{q}_i^A, T_i^A$, and P_i^A for all inlets $i \in \mathcal{IN}_u$, a solution is obtained to the model equations of process unit $u \in \mathcal{T}$.

We also define a binary variable, $z_{o,i}$ for all $o \in \mathcal{M}_i, i \in \mathcal{IN}_u, u \in \mathcal{L} \cup \mathcal{P}$. When outlet o and inlet i are connected, meaning $f_{o,i} > 0$, the binary variable $z_{o,i}$ takes a value of 1 and otherwise the binary variable is zero. Further, when $z_{o,i} = 1$, the pressure at the inlet i must be equal to the pressure $P_{o,i}$ of the stream that exits outlet o (cf. Figure 2b), this ensures the mixing of streams does not result physically in unintended flows, i.e.: from a high-pressure unit into a low-pressure one. When outlet o and inlet i are disconnected, when $z_{o,i} = 0$ and $f_{o,i} = 0$, the downstream pressure, P_i^{in}, is allowed to vary independently of $P_{o,i}$ (c.f Figure 2c).

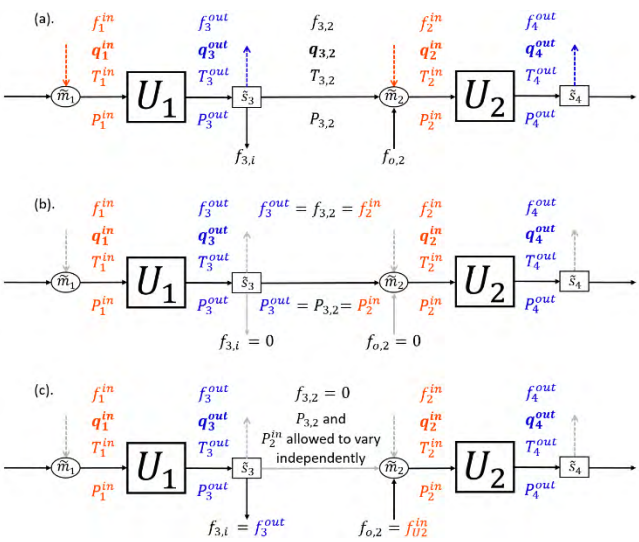

Figure 2. Representation of how pressures are connected within the E-MSON approach. (a) denotes the relevant variables of two conditional units, U_1, which has a single inlet and two outlets, and U_2, which has two inlets and a single outlet. (b) shows the pressure relationship when the two units are connected (when $z_{3,2} = 1$), and (c) illustrates the case when they are not connected (when $z_{3,2} = 0$).

Note that, due to Equations 9b to 9d and 10b to 10g, all streams leaving a splitter are at a composition, temperature, and pressure equal to the splitter's inlet.

Note that meaningful relations between the pressures of two units cannot be derived solely based on their selection in a flowsheet (as denoted by z_u). Inlet and outlet pressures should only be equal when two units are directly connected via flow, as indicated by $z_{o,i}$. In addition to enforcing correct pressure relationships, the introduction $z_{o,i}$ allows for the consideration of flowsheet alternatives where the same units are selected but different streams are active. For instance, the comparative performance of two flowsheet alternatives including the same set of conditional units, but with or without a recycle stream, can now directly be compared using the E-MSON formalism. Lastly, we mention the following assumptions made in the derivation of the MSON, which carry across to the E-MSON: all flows are assumed isenthalpic and frictionless, implying no pressure losses are incurred between units; all units are at steady state; no reaction occurs in mixers and splitters; and all mixers and splitters are isenthalpic and the work of mixing is zero.

E-MSON: FORMULATION

What follows is a complete definition of the E-MSON model equations; although many of these constraints are in common with the MSON approach [15, 16], they are included here for completeness. The MINLP reformulation of the E-MSON is given in Big-M form in this work, but alternative reformulations, such as those using convex hulls [22], could also be utilised. To constrain that the flowrate $f_{o,i}$ between outlet o and inlet i must be non-zero when $z_{o,i} = 1$, and zero when $z_{o,i} = 0$, the following constraints are defined, $\forall\, o \in \mathcal{M}_i, \forall\, i \in \mathcal{IN}_u, \forall\, u \in \mathcal{L} \cup \mathcal{P}$:

$$f_{o,i} \leq \bar{f} z_{o,i} \tag{1a}$$

$$\varepsilon_f z_{o,i} \leq f_{o,i} \tag{1b}$$

where $\bar{f}$ and ε_f are upper and lower bounds on flowrate respectively. Additionally, as in the MSON derivation, we constrain that all inter-unit flowrates directed to a deselected unit are zero, as follows:

$$f_{o,i} \leq \bar{f} z_u, \forall\, o \in \mathcal{M}_i, \forall\, i \in \mathcal{IN}_u, \forall\, u \in \mathcal{T} \tag{2}$$

and that all flowrates into process units are positive:

$$\varepsilon_f \leq f_i^{in}, \forall\, i \in \mathcal{IN}_u, \forall\, u \in \mathcal{L} \tag{3}$$

Further to this, we can impose the following linear relations between unit and stream binary variables, such that the number of feasible combinations of binary variables is reduced significantly. If a unit is deselected, the connected flows should accordingly be inactive:

$$z_{o,i} \leq z_u, \forall\, o \in \mathcal{M}_i, \forall\, i \in \mathcal{IN}_u, \forall\, u \in \mathcal{T} \tag{4a}$$

$$z_{o,i} \leq z_u, \forall\, i \in \mathcal{S}_o, \forall\, o \in \mathcal{OUT}_u, \forall\, u \in \mathcal{T} \tag{4b}$$

and, if a unit is selected, it should have at least one active inlet and one active outlet stream:

$$z_u \leq \sum_{i \in \mathcal{IN}_u} \sum_{o \in \mathcal{M}_i} z_{o,i}, \forall\, u \in \mathcal{T} \quad (5a)$$

$$z_u \leq \sum_{o \in \mathcal{OUT}_u} \sum_{i \in \mathcal{S}_o} z_{o,i}, \forall\, u \in \mathcal{T} \quad (5b)$$

The mass, component, and enthalpy balances around standard mixers are defined as follows, $\forall\, i \in \mathcal{IN}_u, u \in (\mathcal{L} \backslash \mathcal{T}) \cup \mathcal{P}$:

$$f_i^{in} = \sum_{o \in \mathcal{M}_i} f_{o,i} \quad (6a)$$

$$f_i^{in} q_{i,c}^{in} = \sum_{o \in \mathcal{M}_i} f_{o,i}\, q_{o,i,c}, \forall\, c \in \{1, \dots, K\} \quad (6b)$$

$$f_i^{in} h(T_i^{in}, \boldsymbol{q_i^{in}}, P_i^{in}) = \sum_{o \in \mathcal{M}_i} f_{o,i}\, h(T_{o,i}, \boldsymbol{q_{o,i}}, P_{o,i}) \quad (6c)$$

where $h(T, \boldsymbol{q}, P)$ is the enthalpy per unit mass of a stream at temperature T, pressure P, and composition $\boldsymbol{q}$. Note the sum of all component fractions need not be constrained to one at mixers and splitters, given that it is satisfied by the input specifications at sources.

To ensure the desired behaviour of fictitious streams at modified mixers is enforced, at all inlets $i \in \mathcal{IN}_u, \forall\, u \in \mathcal{T}$, modified mixer constraints are defined. The mass, component, and enthalpy balances are given by:

$$f_i^{in} = \sum_{o \in \mathcal{M}_i} f_{o,i} + f_i^M \quad (7a)$$

$$f_i^{in} q_{i,c}^{in} = \sum_{o \in \mathcal{M}_i} f_{o,i}\, q_{o,i,c} + f_i^M q_{i,c}^A, \forall\, c \in \{1, \dots, K\} \quad (7b)$$

$$f_i^{in} h(T_i^{in}, \boldsymbol{q_i^{in}}, P_i^{in}) = \sum_{o \in \mathcal{M}_i} f_{o,i}\, h(T_{o,i}, \boldsymbol{q_{o,i}}, P_{o,i}) + f_i^M h(T_i^A, \boldsymbol{q_i^A}, P_i^A) \quad (7c)$$

To ensure that the fictitious flowrate and inlet pressure are only equal to the auxiliary value when the unit is deselected, and that when the unit is selected, the fictitious flowrate is zero, the following constraints are also necessary, where M is a suitably large positive constant:

$$f_i^M = f_i^A (1 - z_u) \quad (7d)$$

$$-M z_u \leq P_i^A - P_i^{in} \leq M z_u \quad (7e)$$

Additionally, across *all* mixers, the inlet and outlet pressures should only be equal when the corresponding flowrate is non-zero (i.e.: $z_{o,i} = 1$), therefore, $\forall\, o \in \mathcal{M}_i, \forall\, i \in \mathcal{IN}_u, u \in \mathcal{L} \cup \mathcal{P}$:

$$-M(1 - z_{o,i}) \leq P_i^{in} - P_{o,i} \leq M(1 - z_{o,i}) \quad (8)$$

Equation 8 implies that if there are multiple non-zero flows from different outlets connected to an inlet, each of those outlets are at the same pressure. At each standard splitter, we require mass balance constraints and enforce the equality of compositions, temperatures, and pressures across the splitter, s_o. Hence, $\forall\, o \in \mathcal{OUT}_u, \forall\, u \in (\mathcal{L} \backslash \mathcal{T}) \cup \mathcal{F}$, we define:

$$f_o^{out} = \sum_{i \in \mathcal{S}_o} f_{o,i} \quad (9a)$$

$$q_{o,c}^{out} = q_{o,i,c}, \forall\, c \in \{1, \dots, K\}, \forall\, i \in \mathcal{S}_o \quad (9b)$$

$$T_o^{out} = T_{o,i}, \forall\, i \in \mathcal{S}_o \quad (9c)$$

$$P_o^{out} = P_{o,i}, \forall\, i \in \mathcal{S}_o \quad (9d)$$

At modified splitters, we define the following, including the fictitious stream, $\forall\, o \in \mathcal{OUT}_u, \forall\, u \in \mathcal{T}$:

$$f_o^{out} = \sum_{i \in \mathcal{S}_o} f_{o,i} + f_o^S \quad (10a)$$

$$q_{o,c}^{out} = q_{o,i,c}, \forall\, c \in \{1, \dots, K\}, \forall\, i \in \mathcal{S}_o \quad (10b)$$

$$T_o^{out} = T_{o,i}, \forall\, i \in \mathcal{S}_o \quad (10c)$$

$$P_o^{out} = P_{o,i}, \forall\, i \in \mathcal{S}_o \quad (10d)$$

$$q_{o,c}^{out} = q_{o,c}^S, \forall\, c \in \{1, \dots, K\} \quad (10e)$$

$$T_o^{out} = T_o^S \quad (10f)$$

$$P_o^{out} = P_o^S \quad (10g)$$

As fictitious mass flow is added to the unit upon its deselection, to ensure that exact results are obtained from the remainder of the flowsheet, we must also enforce that when $z_u = 0, f_o^S = f_o^{out}$, by:

$$f_o^{out} - f_o^S \leq \bar{f} z_u \quad (11)$$

Furthermore, we must also ensure that no mass leaves in the fictitious stream when the unit is selected, restricting f_o^S to zero as follows:

$$0 \leq f_o^S \leq \bar{f}(1 - z_u) \quad (12)$$

To ensure the design variables of conditional units ($\boldsymbol{x}_{\mathcal{A}_u}$ and $\boldsymbol{y}_{\mathcal{A}_u}$) take the required auxiliary values upon deselection ($\boldsymbol{x}_{\mathcal{A}_u}^{\boldsymbol{A}}$ and $\boldsymbol{y}_{\mathcal{A}_U}^{\boldsymbol{A}}$), we introduce, $\forall\, u \in \mathcal{T}$:

$$-M z_u \leq x_{\mathcal{A}_u} - x_{\mathcal{A}_u}^A \leq M z_u \quad (13a)$$

$$-M z_u \leq y_{\mathcal{A}_u} - y_{\mathcal{A}_u}^A \leq M z_u \quad (13b)$$

As well as this, the "output" variables of a conditional unit, $x_i, i \in \mathcal{C}_u$, should only impact flowsheet-wide constraints if the unit is active. This is achieved by introducing an auxiliary vector of continuous variables, $\boldsymbol{x_i^S}$, and the following constraints, $\forall\, i \in \mathcal{C}_u, \forall\, u \in \mathcal{T}$:

$$-M(1 - z_u) \leq x_i - x_i^S \leq M(1 - z_u) \quad (14a)$$

$$-M z_u \leq x_i^S - c_{u,i} \leq M z_u \quad (14a)$$

where, $c_{u,i}$ is the relevant member of a vector of values $\boldsymbol{c_u}$, often of zeroes, which "output" variables should take upon unit deselection. Further, the flowsheet wide constraints and objective are expressed in term of auxiliary variables x_i^S instead of $x_i\, \forall\, i \in \mathcal{C}_u, \forall\, u \in \mathcal{T}$.

In addition to enforcing that physically meaningful values of pressure are obtained from the E-MSON, it is also of interest to include constraints related to the pressure-change units, $u \in \mathcal{PC}$. If $\mathcal{PC}$ is partitioned into a set $\mathcal{COM}$ of compressing units that lead to a positive pressure difference across them, and analogously a set $\mathcal{EXP}$

of expanding units that lower pressure, such that $\mathcal{PC} = \mathcal{EXP} \cup \mathcal{COM}$, it is possible to enforce that only the correct pressure-change occurs across them. Although the detailed models of such units should ensure that only the correct pressure difference can occur in compressors and expanders at feasible solutions (i.e. an increase across compressors, and a decrease across expanders), numerical difficulties and thus solver failure can still be encountered. For instance, if poor initialisation leads to a negative ΔP across a compressor or an NLP solver progresses in this direction, this can result in a negative value for power, which in many cost correlations would then be raised to a non-integer exponent [21] and thus cause a numerical singularity. To avoid this, we define the following for compressing units, $\forall\, i \in \mathcal{IN}_u, \forall\, o \in \mathcal{OUT}_u, \forall\, u \in \mathcal{COM}$:

$$\varepsilon_P \leq P_o^{out} - P_i^{in} \tag{15a}$$

where ε_P is a small positive constant to ensure a positive ΔP is obtained. Analogously for expanding units, $\forall\, i \in \mathcal{IN}_u, \forall\, o \in \mathcal{OUT}_u, \forall\, u \in \mathcal{EXP}$:

$$\varepsilon_P \leq P_i^{in} - P_o^{out} \tag{15b}$$

AN APPLICATION OF E-MSON

We now illustrate the use of the E-MSON formalism for variable pressure flowsheet synthesis on a simple test problem. The E-MSON representation of the problem is given in Figure 3.

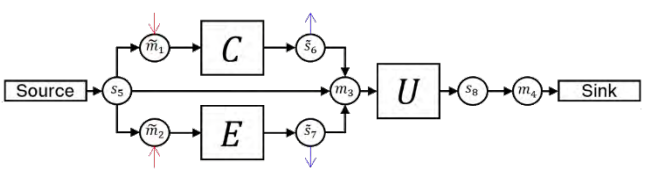

Figure 3. E-MSON representation of the test problem, where units C and E are a conditional compressor and expander, respectively, and U is an arbitrary permanent unit. Splitter s_8 and mixer m_4 are shown for completeness and may be eliminated in practice.

Three feasible flowsheet alternatives can be obtained from the superstructure, where exactly one of the compressor, C, the expander, E, or a bypass stream is active. Within the problem, the key decision variables are the two unit selection binary variables, z_C and z_E, the three flow activity binaries to indicate whether compression, expansion, or neither occurs ($z_{6,3}, z_{7,3}, z_{5,3}$), as well as the inlet pressure to the permanent unit, P_U^{in}.

For simplicity, streams in this problem are described only in terms of mass flowrate and pressure, implying any constraints involving composition and temperature from the full E-MSON derivation are not included. Simplified models are also used for each unit operation. The purpose of this problem is to provide initial evidence of the applicability of the E-MSON to non-isobaric superstructure optimisation problems. Indeed, as the pressure across the compressor and expander must vary from the source pressure, the MSON and SON formulations are inapplicable to even this rudimentary case.

The compressor and expander are assigned a fixed cost, and no mass loss is assumed across them. Notably, as the pressure difference across pressure-change units is constrained to be in the correct direction for each unit, neither the inlet pressure nor the outlet pressure is involved in the toy compressor and expander models. These models are given by, $\forall\, u \in \{C, E\} = \mathcal{PC} = \mathcal{T}$:

$$f_i^{in} = f_o^{out}, \forall\, i \in \mathcal{IN}_u, \forall o \in \mathcal{OUT}_u \tag{20a}$$

$$c_u = \mu_u \tag{20b}$$

where μ_u is a unit-specific constant.

The permanent unit is modelled as following, based on modification of the toy unit models in the case studies on which the MSON was first tested [15], $\forall\, u = U = \mathcal{L} \backslash \mathcal{T}$:

$$x_u f_i^{in} = f_o^{out}, \forall\, i \in \mathcal{IN}_u, \forall\, o \in \mathcal{OUT}_u \tag{21a}$$

$$x_u = \frac{P_i^{in}}{P_{UB}}, \forall\, i \in \mathcal{IN}_u \tag{21b}$$

$$c_u = \mu_u \left(f_i^{in} P_i^{in}\right)^{\upsilon_u}, \forall\, i \in \mathcal{IN}_u \tag{21c}$$

$$P_i^{in} = P_o^{out}, \forall\, i \in \mathcal{IN}_u, \forall\, o \in \mathcal{OUT}_u \tag{21d}$$

where υ_u is a unit-specific constant, x_u is a scaled pressure, and P_{UB} is an upper bound on pressure.

The objective is to maximise "profit", modelled as the revenue (the value of the stream being processed, given by the inlet flowrate into the sink and corresponding price c_f) minus the cost of each unit included in the selected configuration. It is given as:

$$\max\left(c_f f_{sink}^{in} - c_U - c_E^S - c_C^S\right) \tag{22}$$

where c_f is set to 1 without loss of generality, and c_E^S and c_C^S are auxiliary variables representing the cost of the expander and compressor, respectively, taking a value of zero when the unit is deselected. To complete the definition of the model, the unit-specific parameters are given in Table 1, and the additional E-MSON specific parameters are given in Table 2.

Table 1. Unit-specific parameters of the E-MSON test problem, where f_i^A, P_i^A, μ_u, and υ_u are the fictitious flowrate, fictitious pressure, and unit-specific parameters of the compressor, C, expander, E, and unit, U.

Unit	f_i^A	P_i^A	μ_u	υ_u
C	25.0	1.1	20.0	-
E	25.0	9.9	5.0	-
U	-	-	0.3	0.7

Table 2. Additional constants necessary to define the E-MSON test problem. Where f_{source} and P_{source} are the source flowrate and pressure, respectively, $\bar{f}$ and P_{UB} are the flowsheet-wide upper bounds on flow rate and pressure, P_{LB} is the lower bound, and ε_F and ε_P are small constants to ensure positive flowrates and pressure differences.

Parameter	Value
f_{source}	50.0
P_{source}	5.0
$\bar{f}$	50.0
P_{UB}	10.0
P_{LB}	1.0
ε_F	$1/\bar{f}$
ε_P	0.1

Table 3. Results from the E-MSON test problem, where P_U^{in} is the inlet pressure to the unit U in the superstructure.

Primal Iteration #	Selected Configuration	Objective Value	P_U^{in}
1	Bypass active	10.689	5
2	Expander active	5.390	1
3	Compressor active	6.751	10

The test problem was implemented in gPROMS Process 2025.2.0, utilising MINLPOA, a proprietary implementation of OAERAP [23], as the MINLP solver, NLPSQP as the primal NLP subproblem solver, and LPSOLVE as the master MILP problem solver. All the solvers utilised were run with standard settings, except for the convexity-based termination criterion of MINLPOA, which was disabled. The problem involved 72 variables, 8 of which were binary. The solver ran for 4 master problem iterations, obtaining solutions for each feasible combination of unit selection and flow activity binaries (3 combinations out of a possible 32), also indicating that the master problem at the final (fourth) iteration becomes infeasible. Convergence was achieved in less than 1 second. Further detail on the three primal problems that were solved is provided in Table 3, with the alternative in which no pressure-change unit was selected being found as optimal.

The results reported in Table 3 were obtained independently of which of the three configurations was selected as the initial guess to the problem. This was also found to be the case upon initialisation from an infeasible point, where all binary variables were set to zero.

CONCLUSION

In this work, we have presented E-MSON, a framework for the superstructure optimisation of non-isobaric flowsheets. It is a generalisation of the MSON approach in which modified isenthalpic splitters and mixers are introduced to avoid numerical difficulties, with the addition of pressure-change units and new constraints that ensure physically meaningful solutions. The extended MSON has then been tested on a toy superstructure containing an expander and a compressor, yielding promising results. Although this illustrative problem is small, it represents a non-isobaric flowsheet synthesis problem and is thus, to the authors' knowledge, beyond the scope of other exact and rigorous graph-theoretic superstructure formulations such as the SON, or MSON.

In future work the E-MSON will be applied to a broader range of problems, including more realistic process applications in which pressure represents a key operating variable. Beyond this, integrating such rigorous and exact superstructure formulations into molecular design problems is of interest, facilitating simultaneous molecular and process design. Additionally, constraints to further improve numerical convergence for units which require a particular inlet phase will be considered.

Furthermore, given the time-consuming nature of implementing such optimisation problems, automating the writing and execution of E-MSON superstructure optimisation problems is also of value. A Python library to do so within gPROMS Process is currently in development, permitting the user to visually represent the superstructure and generate the corresponding E-MSON problem formulation quickly.

DIGITAL SUPPLEMENTARY MATERIAL

Digital supplementary material, including the gPROMS Process 2025.2.0 file used in the solution of the test problem, and relevant case files, is available here: https://psecommunity.org/LAPSE:2026.0034

ACKNOWLEDGEMENTS

The authors would like to thank the CMAC EPSRC Centre for Doctoral Training in Cyber-Physical Systems for Medicines Development and Manufacturing (CEDAR) (Grant Ref. EP/Y035593/1) for funding this work.

AUTHOR IDENTIFIERS

Author ORCIDs:
Fraser HA: 0000-0002-3894-2154
Gopinath S: 0009-0004-2872-2440
Sefcik J: 0000-0002-7181-5122
Jackson G: 0000-0002-8029-8868
Galindo A: 0000-0002-4902-4156
Adjiman CS: 0000-0002-4573-7722

REFERENCES

1. Mencarelli L, Chen Q, Pagot A, Grossmann IE. A review on superstructure optimization approaches in process system engineering. Computers & Chemical Engineering 136:106808 (2020). https://doi.org/10.1016/j.compchemeng.2020.106808
2. Douglas JM. A hierarchical decision procedure for process synthesis. AIChE Journal 31:353-362 (2004). https://doi.org/10.1002/aic.690310302
3. Alcántara-Avila JR. 50 years of optimization in japan using superstructures. Journal of Chemical Engineering of Japan 56: (2023). https://doi.org/10.1080/00219592.2023.2188112
4. Cremaschi S. A perspective on process synthesis: challenges and prospects. Computers & Chemical Engineering 81:130-137 (2015). https://doi.org/10.1016/j.compchemeng.2015.05.007
5. Kondili E, Pantelides CC, Sargent RWH. A general algorithm for short-term scheduling of batch operations—i. MILP formulation. Computers & Chemical Engineering 17:211-227 (1993). https://doi.org/10.1016/0098-1354(93)80015-f
6. Yeomans H, Grossmann IE. A systematic modeling framework of superstructure optimization in process synthesis. Computers & Chemical Engineering 23:709-731 (1999). https://doi.org/10.1016/s0098-1354(99)00003-4
7. Smith EMB, Pantelides CC. Design of reaction/separation networks using detailed models. Computers & Chemical Engineering 19:83-88 (1995). https://doi.org/10.1016/0098-1354(95)87019-9
8. Lutze P, Babi DK, Woodley JM, Gani R. Phenomena based methodology for process synthesis incorporating process intensification. Ind. Eng. Chem. Res. 52:7127-7144 (2013). https://doi.org/10.1021/ie302513y
9. Demirel SE, Li J, Hasan MMF. A general framework for process synthesis, integration, and intensification. Ind. Eng. Chem. Res. 58:5950-5967 (2019). https://doi.org/10.1021/acs.iecr.8b05961
10. Dowling AW, Biegler LT. A framework for efficient large scale equation-oriented flowsheet optimization. Computers & Chemical Engineering 72:3-20 (2015). https://doi.org/10.1016/j.compchemeng.2014.05.013
11. Schilling J, Horend C, Bardow A. Integrating superstructure?based design of molecules, processes, and flowsheets. AIChE Journal 66: (2020). https://doi.org/10.1002/aic.16903
12. Lee S, Grossmann IE. Global optimization of nonlinear generalized disjunctive programming with bilinear equality constraints: applications to process networks. Computers & Chemical Engineering 27:1557-1575 (2003). https://doi.org/10.1016/s0098-1354(03)00098-x
13. Yeomans H, Grossmann IE. Disjunctive programming models for the optimal design of distillation columns and separation sequences. Ind. Eng. Chem. Res. 39:1637-1648 (2000). https://doi.org/10.1021/ie9906520
14. Navarro-Amorós MA, Ruiz-Femenia R, Caballero JA. Integration of modular process simulators under the generalized disjunctive programming framework for the structural flowsheet optimization. Computers & Chemical Engineering 67:13-25 (2014). https://doi.org/10.1016/j.compchemeng.2014.03.014
15. Gopinath S, Adjiman CS. Superstructure optimization with rigorous models via an exact reformulation. Computers & Chemical Engineering 194:108972 (2025). https://doi.org/10.1016/j.compchemeng.2024.108972
16. Gopinath S, Adjiman CS. Advances in process synthesis: new robust formulations. Systems and Control Transactions 3:145-152 (2024). https://doi.org/10.69997/sct.169290
17. Poling BE, Prausnitz JM, O'Connell JP. Properties of Gases and Liquids, Fifth Edition. McGraw-Hill Education (2020).
18. Liang S, Cao Y, Liu X, Li X, Zhao Y, Wang Y, Wang Y. Insight into pressure-swing distillation from azeotropic phenomenon to dynamic control. Chemical Engineering Research and Design 117:318-335 (2017). https://doi.org/10.1016/j.cherd.2016.10.040
19. Wang R, Lin S. Pore model for nanofiltration: history, theoretical framework, key predictions, limitations, and prospects. Journal of Membrane Science 620:118809 (2021). https://doi.org/10.1016/j.memsci.2020.118809
20. Wang X, Liu X. High pressure: a feasible tool for the synthesis of unprecedented inorganic compounds. https://doi.org/10.1039/d0qi00477d
21. Pundir SS, Kumar S. Improving cost estimation for compressors: evaluating existing and proposing new comprehensive correlations. Chemical Engineering Research and Design 224:79-90 (2025). https://doi.org/10.1016/j.cherd.2025.10.050
22. Lee S, Grossmann IE. New algorithms for nonlinear generalized disjunctive programming. Computers & Chemical Engineering 24:2125-2141 (2000). https://doi.org/10.1016/s0098-1354(00)00581-0
23. Viswanathan J, Grossmann IE. A combined penalty

function and outer-approximation method for MINLP optimization. Computers & Chemical Engineering 14:769-782 (1990). https://doi.org/10.1016/0098-1354(90)87085-4

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Superstructure Optimization of CCUS Value Chain: Case Study on Sohar Freezone in Oman

Shaima Al-Salmani[a], and Muhammad Abdul Qyyum[a]*

[a] Sultan Qaboos University, Department of Petroleum and Chemical Engineering, P.O. Box 33, Al-Khoud, Muscat, Oman
* Corresponding Author: m.qyyum@squ.edu.om.

ABSTRACT

Sohar Port and Freezone is well-positioned to drive Oman`s economic future through continued expansion, increased reliance on sustainable energy, and the integration of advanced digital tools. Sohar contributes 25% of the current GHG emissions in Oman. Decarbonizing Sohar industrial clusters is a strategic challenge for emerging economies striving to attain net-zero objectives. This research develops a superstructure optimization framework for designing an integrated Carbon Capture, Utilization and Storage (CCUS) value chain, specifically applied to the Sohar Port and Freezone in Oman. A mixed-integer linear programming (MILP) formulation that simultaneously selects capture technologies, identifies optimal transport modes (pipeline, trucking, and shipping), and allocates CO_2 flows among utilization and geological storage alternatives. The novelty of this work lies in the multi-scale integration of techno-economic, spatial, and policy dimensions, coupled with GIS-based route generation into the optimization model. Currently, the total CO_2 emission from Sohar Freezone is around 20.3 $MtCO_2$/yr, after applying the CCUS value chain optimization, 7 – 14 $MtCO_2$ annual net reduction could be achieved, corresponding to 30-70% cluster-level decarbonization at 50 – 85 USD/tCO_2 levelized cost of CO_2 capture, transport, and storage. Sensitivity analysis reveals that port-based CO_2 hubs, combined with short-haul shipping to offshore storage, reduce the system cost by 23–31% compared to single-mode configurations. The framework offers decision–support tools for policymakers and industrial stakeholders, integrating techno-economic optimization with regional carbon control policies. The Sohar case study demonstrates the capabilities of computational superstructure design to expedite CCUS implementation in the industrial clusters in the Gulf region.

Keywords: CCUS, Superstructure Optimization, MILP, Sohar Freezone

INTRODUCTION

The reduction of carbon dioxide emissions became a major international concern, considering that in 2024 emissions were 53.2 Gt CO_2eq, a 1.3% rise from 2023 levels. Globally, fossil fuel CO_2 emissions, the primary contributors to the global green house gases (GHG) emissions, continue to rise despite international climate change mitigation agreements [1]. Achieving the Paris Agreement's goal of limiting the rise in temperature of earth to 1.5 °C, GHG emissions must decline to net zero by mid of this century. To achieve targeted GHG reduction, International Energy Agency (IEA) has outlined a range of technological alternatives, such as, carbon capture, utilization and storage (CCUS), switching to low-carbon fuel, process efficiency improvements, and renewable energy with CCUS accounting for 14% of the overall GHG reduction potential [2]. As per the Global CCS Institute, as of mid-2025, CCUS projects captured 64 million tons of CO_2 globally in 2025, 25.5% increase in capturing capacity compared to 2024 (51 Mtpa) [3].

Industrial clusters with emissions that are hard-to-abate are progressively prioritized for carbon capture, utilization, and storage (CCUS) deployment. This is due to their substantial CO_2 concentrations and continuous CO_2 point sources, which facilitates the economic viability of shared transport and storage infrastructure. According to IEA, the pipeline for CCUS projects have significantly increased in recent years, and policy and business model support is essential for achieving bankability

Nomenclature

i	CO_2 source	$F^{SH}_{i,j,tt}$	Continuous, CO_2 flow from source i to hub j using transport technology tt (tCO_2/yr)
j	Hub location	$F^{HS}_{j,k,tt}$	CO_2 flow from hub j to storage k using transport technology tt (tCO_2/yr)
k	Storage site	$F^{HU}_{j,m,tt}$	CO_2 flow from hub j to utilization m using transport technology tt (tCO_2/yr)
m	Utilization option	F^{H}_{j}	total CO_2 flow at hub j (tCO_2/yr)
ct	Capture Technology	$\mathrm{Cap}^{SH}_{i,j,tt}$	Installed pipeline capacity from source i to hub j using transport tt (tCO_2/yr)
tt	Transport technology	$\mathrm{Cap}^{HS}_{j,k,tt}$	Installed pipeline capacity from hub j to storage k using transport tt (tCO_2/yr)
$x_{i,ct}$	Binary,1 if capture technology ct is installed at source i, otherwise, 0	$\mathrm{Cap}^{HU}_{j,m,tt}$	Installed pipeline capacity from hub j to utilization m using transport tt (tCO_2/yr)
x^{hub}_{j}	Binary,1 if hub j is developed, otherwise, 0	E_i	Annual CO_2 emissions at source i (tCO_2/yr)
x^{stor}_{k}	Binary, 1 if storage k is developed, otherwise, 0	ε_{ct}	Capture efficiency of technology ct
x^{util}_{m}	Binary,1 if utilization m is developed, otherwise, 0	$\mathrm{tech_applicable}_{i,ct}$	Binary, returns to 1 if capture technology ct is applicable to source I based on CO_2 concentration constraints, otherwise 0.
$x^{SH}_{i,j,tt}$	Binary, 1 if a pipeline of type tt is built from source i to hub j, otherwise 0	$\underline{C}_{tt}$	Minimum feasible CO_2 capacity of pipeline type tt per year
$x^{HS}_{j,k,tt}$	Binary, 1 if a pipeline of type tt is built from hub j to storage k, otherwise 0	$\bar{C}_{tt}$	Maximum annual CO_2 injection rate
$x^{HU}_{j,m,tt}$	Binary, 1 if a pipeline of type tt is built from hub j to utilization m, otherwise 0	I^{max}_{k}	Maximum annual injection rate of storage k (tCO_2/yr)
$x^{S\to H}_{i,j}$	Binary, 1 if source i is connected to hub j, otherwise 0	U^{cap}_{m}	Maximum annual CO_2 utilization capacity by utilization m (tCO_2/yr)
$f_{i,ct}$	Continuous, CO_2 captured at source i using technology ct (tCO_2/yr)	M	Big -M constant used to activate or deactivate the flow

at commercial scale [4]. Concurrently, hub and cluster strategies are being actively evaluated, wherein shared pipelines, compression, and injection infrastructure minimize redundancy and decrease unit transportation and storage expenses.

Oman has established a national strategy for achieving net-zero emissions by 2050, which is backed by policy frameworks and other supporting initiatives [5]. Sohar Freezone and industrial area are ideal location for a "clustered" decarbonization approach due to the concentration of emissions-producing industries there, including refinery, metals, cement, and petrochemicals. Initiatives for industrial decarbonization in Sohar are receiving more attention as it emits more than 20.3 $MtCO_2$/yr which accounts for 20% of total Oman's CO_2 emissions in 2024 (98.51 $MtCO_2$/yr) [6], [1]. CCUS system-level optimization is required to determine (i) what technologies should be put in place at each source to capture CO_2, (ii) which hubs should be turned on, and (iii) how CO_2 should be routed to utilization and/or stored while keeping capacity constraints into consideration and lowering the total annualized cost.

The optimization in literature on optimization offers a robust basis for CCUS network planning, employing mixed-integer linear programming (MILP) and superstructure formulations. These applications encompass CCUS supply chain optimization under uncertainty, rollout planning, and source-sink matching with capacity constraints [7]. Recent research highlights the need for strong decision-support systems for actual industrial clusters by reviewing the state of source-sink matching techniques and problems and further developing spatial

and cluster-based optimization [8].

The main purpose of this study is to develop a superstructure optimization model that aims for the design and planning of a CCUS supply chain network featuring the minimum annualized net cost for Scope 1 emissions in the Sohar Freezone industrial area in Oman. The model is implemented in GAMS using Mixed Linear Integer Programming (MILP) for different scenarios.

METHODOLOGY

In this study, to achieve CO_2 reduction target aligning with Oman's Pathway to Net-Zero 2050 report and make use of captured CO_2 to produce value-added products, and optimization based CCUS supply chain model for major CO_2 emitters in Oman's industrial sector is developed using GAMS program. The methodology consists of five major steps.

Step 1: Identify the major industries (sources) that emits CO_2 in a large quantity, possible capture technologies, CO_2 hubs, storage and utilization nodes.

Step 2: Create a superstructure of CCUS network by connecting the identified CO_2 emission sources, hubs, storage sites and utilization routes.

Step 3: Construct CCUS supply chain optimization model in GAMS using mixed-integer linear programming (MILP) formulation for three scenarios to achieve 7% reduction level in CO_2 emission by 2030, 54% reduction by 2040 and 94% reduction in CO_2 emissions by 2050 aligning with Oman's net zero pathways' targets.

Step 4: Optimize and identify the best CCUS network configuration for relevant scenarios with various targeted criteria and conduct techno-economic analysis.

Data Collection

CO_2 sources data related to emission, concentration of CO_2 and factories' geological location, storage and utilization nodes in the Sohar Freezone are collected in Tables 1, 2, and 3 and analyzed. The study focuses on SFZ industrial area due to it considers as the largest CO_2 emitters in Oman, as described in Table 1. Geographic locations of these factories, hubs, utilization and storage sites are also collected and used in the optimization model. The potential CO_2 sinks for Sohar Freezone, including those long-term CO_2 storage and utilization routes are identified and provided in Tables 2 and 3. An offshore deep saline aquifer is being considered as a potential long-term sink due to its high injection capacity of up to 20 $MtCO_2$/yr, while an onshore geological storage site is being considered as secondary alternative with 10 $MtCO_2$/yr capacity.

Moreover, three utilization paths are evaluated and studies: Enhanced oil recovery (EOR), chemical production, and concrete manufacturing, with capacities ranging from 0.8 to 5 $MtCO_2$/yr as shown in Table 3.

Table 2: Possible CO_2 Storage Sites at Sohar Freezone

Storage options	Type	Max. Injection rate $MtCO_2$/yr	Longitude	Latitude
Storage 1	Offshore deep saline aquifer	20	56.70	24.52
Storage 2	Geological storage	10	56.45	24.30

Table 3: Possible CO_2 Utilization Route at SFZ

Utilization Options	Type	Capacity $MtCO_2$	Longitude	Latitude
Utilization 1	EOR Field	5.0	56.969	23.619
Utilization 2	Chemical Plant	2.0	56.631	24.475
Utilization 3	Concrete Facility	0.8	56.579	24.429

OPTIMIZATION MODEL AND ASSUMPTIONS

Model Assumptions

Figure 1 represents the full supply chain of CCUS for this case study. The superstructure CCUS supply chain model is developed in GAMS under the following assumptions and basis:

- Capture plants are located near the emissions source to prevent accounting transportation cost of fuel gas.
- Each industrial source is connected to at most one capture technology, which is installed directly at source location the capture technology is selected based on the CO_2 concentration at feed stream, energy consumption, CAPEX and OPEX and technology efficiency.
- All capture technologies alternatives have the same availability.
- The targeted emissions reduction is 7% as an early phase decarbonization by 2030, and 50% as a realistic carbon reduction through CCUS that shall be attained by 2040, and 94% as an ambitious target to decarbonize emissions from scope 1 industries by 2050. These reduction level % is aligned with Oman Pathway towards net-zero objectives [5].
- Onshore and offshore pipelines are the transport mode that is considered for Sohar Freezone and for each transport link (source to hub, hub to utilization

Table 1: CO_2 Emissions Data for 17 Major Emitters from Sohar Freezone

Factory No.	Industry Type	CO_2 Emissions tCO_2/yr	Plant Capacity tone/yr	CO_2 Conc. (mol%) in Flue Gas	Capture Priority Score	Emissions Intensity tCO_2/tone of Product	Longitude	Latitude
F1	Refining	3945768	9400000	13	81%	0.42	56.608	24.483
F2	Chemicals	2798090	2400000	12	62%	1.17	56.616	24.486
F3	Iron and steel	2520000	2400000	22	74%	1.05	56.604	24.503
F4	Chemicals	2409000	1050000	25	77%	2.29	56.631	24.475
F5	Rubber and Plastic	2114019	880000	16	58%	2.40	56.623	24.462
F6	Refining	1869000	600000	13	49%	3.12	56.552	24.455
F7	Chemicals	1833387	1016000	14	50%	1.80	56.608	24.483
F8	Iron and steel	1228500	650000	22	54%	1.89	56.601	24.502
F9	Iron and steel	1056304	10000000	15	40%	0.11	56.590	24.506
F10	Refining	212931	600000	14	26%	0.35	56.608	24.483
F11	Iron and steel	113400	60000	18	31%	1.89	56.599	24.519
F12	Iron and steel	94500	50000	20	33%	1.89	56.556	24.468
F13	Iron and steel	67682	70000	12	20%	0.97	56.606	24.419
F14	Iron and steel	20625	75000	22	36%	0.28	56.556	24.460
F15	Iron and steel	15120	8000	18	29%	1.89	56.563	24.461
F16	Iron and steel	13750	50000	22	35%	0.28	56.553	24.462
F17	Iron and steel	12150	9000	10	16%	1.35	56.563	24.461

or storage), one transport type shall be selected.

- One emission source node can only be connected to one hub. However, a hub node can get CO_2 from more than one emission source node.
- Energy consumption is modelled using technology-specific energy intensities and a uniform energy price, with transport and storage implicitly included in OPEX terms to ensure parsimonious system-level optimization.
- The basis of energy price is 70% energy come from natural gas as heat and 30% from electricity, the cost of electricity in Oman for non-residential tariff is around 18.1 USD/GJ, while the cost of heat from natural gas 3.5 USD/GJ, hence, the energy price is considered as 7.88 USD/GJ [9].
- Based on various techno-economic studies in CCUS value chain, the discount rate for standard commercial projects is between 7% – 8% and for this study, 8% is considered [10], [11].
- The carbon capture cost is 50 USD/tCO_2 as a baseline as an exogenous policy parameter in the optimization to represent the economic value of permanent CO_2 storage; this reflects mid-range prices used in global net-zero and CCUS techno-economic studies [12].
- The hub cost is defined as including compression/ liquefaction, intermediate buffer storage, on/off – loading facilities, mode switching (rail – ship – pipeline) and export terminals. Capture plant CAPEX / OPEX, injection wells and subsurface monitoring, and utilization plant costs are excluded from the hub cost.
- The overall CCUS network is economically optimized for 20 years' life span to provide the geographic distribution and scale of capture, transportation, utilization and sequestration sites.

Overall Optimization Model and Model Constraints

The purpose of optimization model is to minimize the total cost of CCUS value chain that is associated with cost of CO_2 capture, transportation, CO_2 hub, storage, utilization and energy consumption., as described in equation 1.

$$\min. obj = Total\ Annualized\ net\ cost = CCUS\ cost - CCUS\ revenue = Capture\ cost + Tansport\ cost + Hub\ Cost + Storage\ Cost + Utlization\ Cost + Energy\ Cost - Revenue\ Carbon\ Credit - Revenue\ Utilization \quad (1)$$

In addition, the MILP superstructure model includes four main constraints to guarantee the precision and viability of the results, i.e., capture allocation, mass balance, network connectivity / capacity and sink capacity. CO_2 capture at each source is bounded by emission capacity and limited to at most one technology, as described in equation 2.

Figure 1. Superstructure for Sohar Freezone CCUS Supply Chain Network

$$\sum_{ct} f_{i,ct} \leq E_i, \sum_{ct} x_{i,ct} \leq 1, f_{i,ct} \leq E_i \times \varepsilon_{ct} \times x_{i,ct} \times tech_applicable_{i,ct} \quad (2)$$

Second constraint is the mass conservation, where the captured CO_2 is routed through hubs and then either to utilization or storage,

$$\sum_{j,tt} F^{SH}_{i,j,tt} = \sum_{ct} f_{i,ct},\ F^{H}_{j} = \sum_{i,tt} F^{SH}_{i,j,tt}, \sum_{k,tt} F^{HS}_{j,k,tt} + \sum_{m,tt} F^{HU}_{j,m,tt} = F^{H}_{j} \quad (3)$$

Third, connectivity and activation link flows to discrete pipeline choices (no free transport is considered) using big-M logic, which is a constant used to activate / deactivate the flow when transport type from source to hub, hub to storage, and hub to utilization is selected. The hub will open if any CO_2 flow from any source is connected ($x^{S\to H}_{i,j} \leq x^{hub}_{j}$).

$$\sum_{tt} F^{SH}_{i,j,tt} \leq M \sum_{tt} x^{SH}_{i,j,tt},\ F^{SH}_{i,j,tt} \leq M\, x^{SH}_{i,j,tt} \quad (4)$$

The last constraint is pipeline sizing and sink limits enforce capacity feasibility, and the utilization and storage locations are bounded by maximum injection rate and utilization capacities refer to equations (5, 6).

Table 4: CCUS value chain optimal network and CO_2 flow in pipelines for different CO_2 reduction level scenarios

Scenarios	Source	Capture Technology	Connection	Amount of CO_2 Captured tCO_2/yr	Transport Type	End User
Scenario A: 7% Emissions Reduction by 2030	F6	Advanced Solvent-based Absorption Technology	F6 to Hub 2	1422696	Onshore pipeline	Hub 2 to Utilization 2 and 3
Scenario B: 54% Emissions Reduction by 2040	F1	Advanced Solvent-based Absorption Technology	F1 to Hub 1	3748480	Offshore pipeline	Hub 1 to Storage 2, Utilization 2 and 3
	F2		F2 to Hub 1	2658186	Offshore pipeline	
	F3		F3 to Hub 1	1457340	Onshore pipeline	
	F7		F7 to Hub 1	1741718	Onshore pipeline	
	F8		F8 to Hub 1	1167075	Onshore pipeline	
	F10		F10 to Hub 1	202284	Onshore pipeline	
Scenario C: 94% Emissions Reduction by 2050	F1	Advanced Solvent-based Absorption Technology	F1 to Hub 1	3748480	Onshore Pipeline	Hub 1 to Storage 1, Storage 2, and Utilization 3
	F2		F2 to Hub 1	2658186	Onshore Pipeline	
	F3		F3 to Hub 1	2394000	Onshore Pipeline	
	F4		F4 to Hub 1	2288550	Onshore Pipeline	
	F5		F5 to Hub 1	2008318	Onshore Pipeline	
	F6		F6 to Hub 1	1775550	Onshore Pipeline	
	F7		F7 to Hub 1	1741718	Onshore Pipeline	
	F8		F8 to Hub 1	1167075	Onshore Pipeline	
	F9		F9 to Hub 1	1003489	Onshore Pipeline	
	F10		F10 to Hub 1	202284	Offshore Pipeline	
	F11		F11 to Hub 1	107730	Onshore Pipeline	
	F13		F13 to Hub 1	9393	Onshore Pipeline	

$$\underline{C}_{tt} x_{i,j,tt}^{SH} \leq Cap_{i,j,tt}^{SH} \leq \overline{C}_{tt} x_{i,j,tt}^{SH}, F_{i,j,tt}^{SH} \leq Cap_{i,j,tt}^{SH} \quad (5)$$

$$\sum_{j,tt} F_{j,k,tt}^{HS} \leq I_k^{max} x_k^{store}, \sum_{j,tt} F_{j,m,tt}^{HU} \leq U_m^{cap} x_m^{util} \quad (6)$$

These constraints collectively ensure that the chosen capture technologies, hub locations, pipeline types, and sink options must all be physically consistent in order for CCUS network to achieve the desired capture target.

RESULTS AND DISCUSSION

The superstructure optimization model findings indicate a distinct cost-effective sequence for CCUS deployment in one major industrial area in Oman, Sohar Freezone, driven mainly by scale, infrastructure availability, and distances. As described in Table 4, in short-term scenario A (7% reduction target), CO_2 reduction is predominantly accomplished by localized capture and utilization, hence, reducing the need for transport and storage sites as the utilization is preferred. Table 5 and Figure 2 represent the cost components of CCUS supply chain in the three scenarios. For scenario A, the total cost is 282 million USD/ yr, with capture component account for 70% of total cost. Transport cost is less than 2 MUSD/yr and storage are not preferred as the utilization path is nearby the decarbonized factory.

As the goal for mitigation rises to 54% reduction in CO_2 emission, the CCUS system experiences a structural shift towards centralized hubs and permanent geological storage. As depicted from Table 5, in scenario B the cost of capture escalates significantly to 1.56 billion USD annually, although CO_2 transport and storage (18 and 125 MUSD/yr, respectively) are identified as economically viable system components. Despite a significant increase in total cost, the CO_2 avoidance cost decreases to 134 USD/tCO_2 compared to scenario A (224.8 USD/tCO_2), this indicates for the importance of economies of scale and efficiency improvements from shared CCUS infrastructure. Moreover, the implementation of carbon credit revenues (500 MUSD/yr) is essential in mitigating rising CAPEX and energy costs, demonstrating that integrated hub-based carbon capture, utilization and storage systems are more effective than disjoined, source-specific solutions after mitigation approaches moderate levels

Table 5: Total cost for the optimal CCUS networks for three studied scenarios

Cost item	Scenario A	Scenario B	Scenario C
Total Cost, MUSD/yr	281.59	2007.42	3553.53
Capture, MUSD/yr	202.01	1558.36	2712.71
Hub cost, MUSD/yr	14.85	47.18	71.57
Transportation, MUSD/yr	1.97	17.95	28.53
Utilization, MUSD/yr	33.61	33.61	18.33
Storage, MUSD/yr	0	125.46	330.97
Energy, MUSD/yr	29.15	224.86	391.42
Revenue income, MUSD/yr	82.68	69.25	64
Carbon Credit Revenue, MUSD/yr	0	500	915.24
CO_2 Avoidance cost, USD/tCO_2	224.87	133.69	135.03

[13].

The model approaches a fully integrated CCUS network with numerous storage sinks and a dominant hub under conditions of deep decarbonization (scenario C, 94% level reduction), with practically all major emissions sources linked to this network, as result indicated in Table 4. Table 5 shows that in scenario C, the capture cost (2.71 billion USD/yr) and storage cost (331 MUSD/yr) constitute the predominant system expenses, although energy consumption emerges as an essential factor (391 MUSD/yr), indicative of the thermodynamic penalty associated with near-complete abatement. In comparison to scenarios A and B, the utilization cost in scenario C (18.33 million USD/yr) is reduced, indicating that the storage option is the most practical choice for large scale CCUS network deployments in the Sohar Freezone industrial area. The CO_2 avoidance cost is stabilized at around 135 USD/tCO_2, indicating that after infrastructure is constructed, marginal abatement costs remain relatively constant even at high capture levels.

These findings align with Oman's strategies for achieving net zero pathway and global climate emissions reduction objectives. The utilization-driven configuration in the 2030 scenario aligns with Oman's Nationally Determined Contribution (NDC), highlighting cost-efficient industrial mitigation while constraining infrastructure lock-in. The 2040 scenario B depicts a transitional phase consistent with Oman vision 2040, wherein shared CO_2 transport and storage infrastructure become economically and strategically viable. The captured CO_2 required to be transported to the end-users, either storage sites or utilization plant via pipelines, ships or trucks. The cost of pipeline is highly influenced by the economies of scale, and it varies from country to country. However, shipping can be the best options for very long transport distances, and mid-range CO_2 capacities compared to pipelines. In this case study, shipping is not feasible option since the CO_2 will be stored and utilized within the same region, hence, onshore and offshore pipelines are considered in this study. As well as, to deploy the CCUS network for large scale, CO_2 storage sites must be selected and established. The onshore storage sites are usually less expensive compared to offshore storage sites, due to reuse of depleted oil & gas reservoirs and less data availability for offshore storage sites [14]. The same results are obtained in this case study, geographical storage sites are preferable, compared to onshore sites, Table 4. By 2050, network pipelines and geological storage will be prevailed, aligning with the global net-zero goals established by IEA and IPCC [15]. These pathways emphasise the importance of CCUS for decarbonization hard – to-abate industrial sector.

The results collectively establish the Sohar Freezone as a credible foundation for Oman's long-term decarbonization policy, capable of progressing from initial mitigation to complete net-zero compliance in accordance with international climate commitments.

CONCLUSION

This study developed a mixed-integer linear programming to evaluate the cost effective of CCUS network for major industrial region in Oman, Sohar Freezone. The study's main objective was to determine the extent to which the CCUS infrastructure, technology selection, and system configuration should evolve under different emissions-reduction targets that are aligned with Oman transition towards net-zero by 2030, 2040 and 2050.

The model results indicate that the implementation of CCUS in Sohar is most economically viable when executed in a phased and system-integrated approach. At low mitigation levels, scenario A, localized capture and utilization, which requires short distance transportation and no geological storage, is the main method of reducing CO_2 emissions. Centralized nodes, shared pipeline networks, and permanent storage become economically

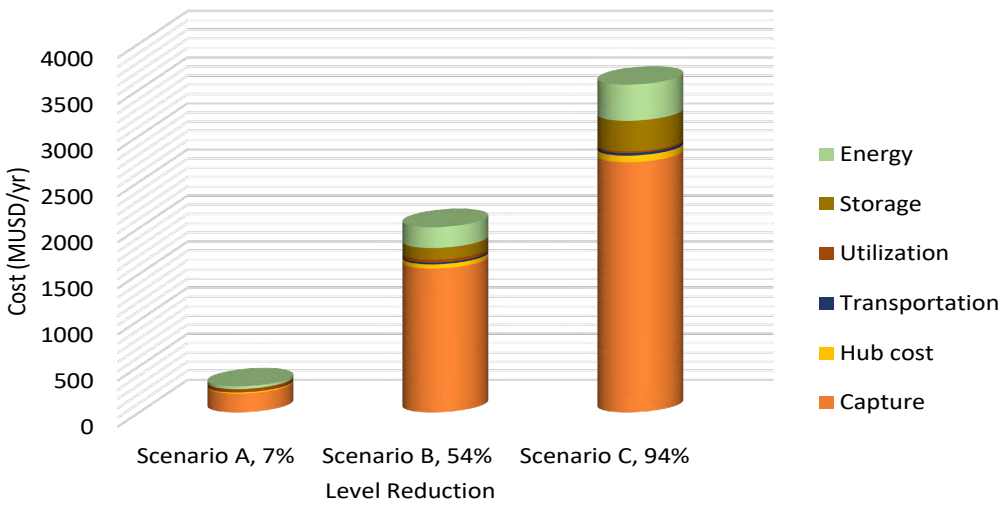

Figure 2. Cost Components Results for the Studies Scenarios

viable as ambition increases, scenario B, resulting in substantial reduction in CO_2 avoidance costs because of economic of scale. In scenarios of deep decarbonization, scenario C, the system evolves into a sophisticated CCUS network primarily characterized by extensive capture and geological storage, with utilization serving a supplementary component. Moreover, it is important to highlight that, marginal abatement costs remain relatively stable at very high capture rates once the infrastructure is established.

The results emphasise the potential of Sohar Freezone to function as a strategic CCUS anchor for the decarbonization of Oman's industrial sector. The initial deployment of CCUS value chain is driven by industries and availability of utilization paths, however, sustained success relies on the government and supported policies for planning shared transportation, building storage infrastructures and availability of carbon market mechanisms. These findings show how CCUS may move from being a short-term mitigation alternative to being a foundation of Oman's net-zero industrial strategy, in line with global climate goals, and provide practical guidance for policymakers and business leaders [5].

Although the model offers significant insights regarding to the deployment of CCUS, future research should enhance the framework by integrating the uncertainty in the carbon pricing and technology learning rates, dynamic multi-period investment strategies, and interactions with emerging low carbon alternatives like, green hydrogen and process electrification, to future strengthen long-term planning applicability.

ACKNOWLEDGEMENTS

This research was supported by Hanbal Masstech Co., Ltd., Korea in collaboration with the Petroleum and Chemical Engineering Department at the College of Engineering, Sultan Qaboos University. The project was funded under the grant number [CR/ENG/PCED/23/04].

REFERENCES

1. M. Crippa *et al.*, Eds., *GHG emissions of all world countries: 2025*. Luxembourg: Publications Office, 2025. doi: 10.2760/5917997.
2. "Climate action support 2025, " International Renewable Energy Agency (IRENA), Abu Dhabi, UAE, 2025. [Online]. Available: https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2025/Nov/IRENA_COU_Climate_Action_Report_2025.pdf
3. "Global-Status-of-CCS-2025-report-9-October.pdf." Accessed: Feb. 02, 2026. [Online]. Available: https://www.globalccsinstitute.com/wp-content/uploads/2025/10/Global-Status-of-CCS-2025-report-9-October.pdf
4. "Carbon Capture Utilisation and Storage - Energy System, " IEA. Accessed: Jan. 19, 2026. [Online]. Available: https://www.iea.org/energy-system/carbon-capture-utilisation-and-storage
5. "The Sultanate of Oman's National Strategy for an Orderly Transition to Net Zero." Accessed: Feb. 02, 2026. [Online]. Available: https://www.ea.gov.om/media/aaslyc3l/oman-net-zero-report-2022_screen.pdf
6. Foreign Ministry of Oman, "Sohar Port and Freezone joins alliance for eliminating carbon emissions, " FM.gov.om. Accessed: Jan. 19, 2026. [Online]. Available: https://www.fm.gov.om/sohar-port-and-freezone-joins-alliance-for-eliminating-

carbon-emissions/
7. Leonzio G, Foscolo PU, Zondervan E. Optimization of CCUS supply chains for some european countries under the uncertainty. Processes 8:960 (2020). https://doi.org/10.3390/pr8080960
8. Zhang J, Zhao L, Jin L, Zhu C, Wang H, Wang L. Optimizing regional CCUS clusterization deployment for multi-industrial sectors: a carbon neutrality pathway for emission-intensive region. Carbon Capture Science & Technology 17:100495 (2025). https://doi.org/10.1016/j.ccst.2025.100495
9. "CCUS deployment challenges and opportunities for the GCC." Accessed: Feb. 02, 2026. [Online]. Available: https://www.ogci.com/wp-content/uploads/2023/04/Gulf-CCUS-challenges-and-opportunities-Mar-2022_compressed.pdf
10. Zhang S, Liu L, Zhang L, Zhuang Y, Du J. An optimization model for carbon capture utilization and storage supply chain: a case study in northeastern china. Applied Energy 231:194-206 (2018). https://doi.org/10.1016/j.apenergy.2018.09.129
11. "2024-03 Techno-Economic Assessment of Small-Scale Carbon Capture for Industrial and Power Systems.pdf." Accessed: Feb. 02, 2026. [Online]. Available: https://publications.ieaghg.org/technicalreports/2024-03%20Techno-Economic%20Assessment%20of%20Small-Scale%20Carbon%20Capture%20for%20Industrial%20and%20Power%20Systems.pdf
12. Bazzana D, Comincioli N, Gusperti C, Legrenzi D, Rizzati M, Vergalli S. Gotta catch 'em all: CCUS with endogenous technical change. SSRN Journal : (2023). https://doi.org/10.2139/ssrn.4647837
13. ZHAO X, XIAO J, HOU J, WU J, LYU X, ZHANG J, LIU Y. Economic and scale prediction of CO2 capture, utilization and storage technologies in china. Petroleum Exploration and Development 50:751-764 (2023). https://doi.org/10.1016/s1876-3804(23)60425-2
14. D. H. Liu, "TECHNOLOGY READINESS AND COSTS OF CCS, " Global CCS Institute, Mar. 2021. Accessed: Jan. 26, 2026. [Online]. Available: https://scienceforsustainability.org/w/images/b/bc/Technology-Readiness-and-Costs-for-CCS-2021-1.pdf
15. "Net Zero by 2050 - A Roadmap for the Global Energy Sector, " International Energy Agency, Oct. 2021. Accessed: Jan. 15, 2026. [Online]. Available: https://iea.blob.core.windows.net/assets/deebef5d-0c34-4539-9d0c-10b13d840027/NetZeroby2050-ARoadmapfortheGlobalEnergySector_CORR.pdf

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Design Optimization of Shell-and-Tube Heat Exchangers Under Operational Uncertainty: A Comparative Study Across Three Paradigms

Fernando Israel Gómez-Castro[a], Sergio Iván Martínez-Guido[b], Claudia Gutiérrez-Antonio[b], Oscar Daniel Lara-Montaño[c]*

[a] Universidad de Guanajuato, Departamento de Ingeniería Química, Guanajuato 36050, Mexico
[b] Universidad Autónoma de Querétaro, Facultad de Ingeniería, Querétaro 76230, Mexico
[c] Universidad Autónoma de Querétaro, Campus Amealco, Querétaro 76850, Mexico
* Corresponding Author: oscar.lara@uaq.mx.

ABSTRACT

Shell-and-tube heat exchangers are critical assets in the process industries, yet their design optimization typically relies on deterministic formulations that ignore operational variability. This study presents a systematic comparison of four optimization approaches, spanning three paradigms for decision-making under uncertainty, applied to shell-and-tube heat exchanger design. The objective minimizes total annualized cost (comprising capital and pumping costs) subject to thermal duty, pressure drop, velocity, and geometric constraints. Six uncertain parameters are modeled across three categories: mass flowrates (coefficient of variation = 10%), inlet temperatures (coefficient of variation = 2%), and fouling resistances (uniform distribution). Shell-side heat transfer is computed via the Bell–Delaware method, while tube-side correlations follow Sieder–Tate. The four approaches benchmarked are: (1) deterministic optimization as a baseline, (2) sample-average approximation, a classical stochastic programming method, (3) box-constrained robust optimization, which guards against worst-case realizations without probabilistic assumptions, and (4) Wasserstein distributionally robust optimization, which hedges against distributional ambiguity. Results reveal a clear cost–robustness hierarchy: deterministic designs yield the lowest expected cost but lack formal guarantees, sample-average approximation balances cost and reliability with modest computational overhead, the distributionally robust approach achieves the tightest cost dispersion at a moderate premium, and robust optimization delivers the highest feasibility guarantees at the greatest cost. These findings provide practitioners with quantitative guidance for selecting an appropriate optimization paradigm based on reliability requirements, data availability, and computational budget.

Keywords: optimization under uncertainty, heat exchanger design, stochastic programming, distributionally robust optimization, robust optimization

INTRODUCTION

Shell-and-tube heat exchangers (STHEs) are the most widely deployed heat transfer equipment in the chemical and petrochemical industries, representing a significant share of capital investment. Operating costs (dominated by pumping power) contribute substantially to facility-wide energy consumption [1]. Traditional deterministic design assumes fixed operating conditions, yet industrial reality involves persistent parametric uncertainty: upstream disturbances introduce flowrate fluctuations, control systems exhibit finite bandwidth, and fouling evolves stochastically. Designs optimized under nominal conditions often violate pressure-drop or thermal-duty constraints when conditions deviate, whereas conservative designs with arbitrary safety factors incur unnecessary capital costs [2].

Recent advances in STHE optimization have ex-

plored metaheuristic algorithms[3], data-driven surrogates [4], and multi-objective frameworks [5]. However, these efforts largely employ deterministic formulations or ad-hoc robustness strategies without rigorous uncertainty quantification. Optimization under uncertainty encompasses three principal paradigms [6]: stochastic programming, which optimizes over a known probability distribution, robust optimization, which hedges against worst-case realizations within a deterministic uncertainty set, and distributionally robust optimization, which guards against the worst-case distribution within an ambiguity set [7]. Although each paradigm has been applied individually to process design problems, systematic comparative studies evaluating multiple paradigms on identical case studies remain scarce [8].

Despite theoretical progress, systematic comparative studies evaluating multiple methodologies on identical problems under consistent uncertainty specifications remain limited. Existing literature presents individual methods in isolation, with incomparable assumptions about uncertainty, constraint formulations, and evaluation protocols, hindering practical adoption. Engineers lack actionable guidance on method selection that accounts for application requirements, data availability, and computational resources.

This work addresses these gaps through a systematic benchmarking of four optimization strategies spanning the three main paradigms for decision-making under uncertainty, applied to STHE design under realistic operational variability. The objectives are: (i) to provide the first consistent comparison of deterministic, sample-average approximation (SAA), box-constrained robust optimization (RO), and Wasserstein distributionally robust optimization (WDRO) approaches for STHE design under identical uncertainty specifications; (ii) to quantify cost–robustness trade-offs across paradigms; (iii) to assess computational scalability of each approach; (iv) to identify failure modes and practical limitations; and (v) to develop actionable method selection guidelines for industrial practitioners.

PROBLEM FORMULATION AND METHODS

Deterministic Design Model

The STHE design problem minimizes total annualized cost (TAC) comprising annualized capital expenditure (via a capital recovery factor accounting for interest rate and plant lifetime) plus annual operating expenditure. The capital recovery factor is $CRF = \frac{i(1+i)^n}{(1+i)^n - 1}$, with interest rate i = 5% and lifetime n = 20 years. TAC is calculated with equation (1).

$$\mathrm{TAC} = \frac{\mathrm{CAPEX}}{\tau} + 8760 C_{\mathrm{elec}} \frac{P_s + P_t}{\eta} \tag{1}$$

where C_{elec} is electricity cost ($/kWh), P_s and P_t are shell/tube-side pumping power (kW), η is pump-motor efficiency, and 8, 760 represents annual operating hours.

The design vector x comprises 11 variables: shell diameter D_s, tube length L, baffle spacing L_{bc}, tube-side passes N_p, plus some discrete selections. Thermal-hydraulic performance employs the Bell-Delaware method for shell-side, accounting for flow bypassing, baffle leakage, and adverse gradients through correction factors J_c, J_l, J_b, J_s, J_r typically ranging 0.6-0.9. This can be observed in equation (2).

$$h_s = h_{\mathrm{ideal}} \cdot J_c \cdot J_l \cdot J_b \cdot J_s \cdot J_r \tag{2}$$

Tube-side forced convection employs Sieder-Tate for turbulent flow ($\boldsymbol{Re} > \mathbf{10,000}$), shown in equation (3).

$$\mathrm{Nu}_t = 0.027 \mathrm{Re}^{0.8} \mathrm{Pr}^{1/3} \left(\frac{\mu}{\mu_w}\right)^{0.14} \tag{3}$$

Operational constraints enforce: (i) pressure drop $\Delta P_s, \Delta P_t \leq 75$ kPa (typical available head); (ii) velocity $0.8 \leq v_t \leq 3.0$ m/s (fouling mitigation and erosion prevention); (iii) length-to-diameter $L/D_s \leq 15$; (iv) heat duty $Q \geq Q_{\mathrm{req}}$.

Uncertainty Quantification

We model six uncertain parameters across three physical categories using physically-motivated probability distributions. Shell-side and tube-side mass flowrates exhibit the largest variability (10% CV) due to upstream disturbances and control limitations. Shell-side and tube-side inlet temperatures are regulated more tightly (2% CV). Shell-side and tube-side fouling resistances follow uniform distributions reflecting epistemic uncertainty about the fouling state within cleaning cycles. Training data comprises 30 scenarios sampled via Latin hypercube sampling (LHS) to ensure coverage of the uncertainty space. Testing employs 100 independent scenarios for out-of-sample validation. Stress scenarios (N=10) represent extreme combinations to assess robustness under tail events, shown in equation (4).

$$F_{m,s} \sim \mathrm{TN}\left(27.8\ \mathrm{kg/s}, (2.78\ \mathrm{kg/s})^2, 19.46, 36.14\right) \tag{4}$$

The $\pm 30\%$ truncation bounds capture process transients while preventing physically unrealistic scenarios. Flowrate affects both thermal (heat transfer $\propto Re^{0.6}$) and hydraulic ($\Delta P \propto Re^{1.8}$) performance simultaneously.

Shell-side inlet temperature is modeled with the same distribution structure (2% CV), as shown in equation (5). Tube-side inlet temperature follows equation (6).

$$F_{m,s} \sim \mathrm{TN}\left(27.8\ \mathrm{kg/s}, (2.78\ \mathrm{kg/s})^2, 19.46, 36.14\right) \tag{5}$$

$$T_{t,\,\mathrm{in}} \sim \mathrm{TN}\left(298.15\ \mathrm{K}, (5.96\ \mathrm{K})^2, 286.19, 310.11\right) \tag{6}$$

Fouling resistance accumulates monotonically from clean to heavily fouled state. We employ a uniform distribution reflecting epistemic uncertainty about the actual fouling state within cleaning cycles, as shown in equation (7).

$$R_s \sim \mathrm{Uniform}\left(0.00033\ \mathrm{m}^{2\cdot\mathrm{K/W}}, 0.00050\ \mathrm{m}^{2\cdot\mathrm{K/W}}\right) \tag{7}$$

These bounds correspond to moderate fouling service (treated cooling water, light oils) per TEMA standards.

LHS generates scenarios with improved convergence relative to crude Monte Carlo. A three-tier evaluation protocol is employed: (i) training ($N = 30$) drives optimization following guidelines for 4-6 uncertain parameters; (ii) test ($N = 100$) assesses out-of-sample performance; (iii) stress ($N = 10$) probes extreme events (deterministic worst-cases plus tail samples beyond 95th percentile).

OPTIMIZATION APPROACHES UNDER UNCERTAINTY

Deterministic Optimization

As a reference baseline, deterministic optimization solves the design problem at nominal (expected-value) parameters ξ_0, as in equation (8).

$$\min_x \mathrm{TAC}\,(x, \xi_0) \quad \text{s.t.} \quad g(x, \xi_0) \le 0,\ x \in \mathcal{X} \tag{8}$$

Where g represents constraints. Differential evolution with adaptive penalty functions, shown in equation (9).

$$\Phi(x, \xi_0) = \mathrm{TAC}(x, \xi_0) + \rho_k \sum_i \max\left(0, g_i(\boldsymbol{x}, \xi_0)\right)^2 \tag{9}$$

where $\rho_k = \rho_0 \cdot (1.2)^k$ (initialized $\rho_0 = 1000$, capped at 10, 000) encourages exploration then ensures constraint satisfaction. The adaptive penalty mechanism steers the optimizer away from tight constraint boundaries, which may result in incidental margins under perturbations. However, since no uncertainty information enters the formulation, any out-of-sample feasibility is problem-specific rather than guaranteed, and this approach serves strictly as a cost baseline.

Sample-Average Approximation

Sample-average approximation is a classical stochastic programming method that replaces the true expected value with an empirical average over a finite scenario set drawn from the known distribution. The penalized formulation is given in equation (10).

$$\min_{\mathrm{x}} \mathrm{E}_\xi\,[\mathrm{TAC}(\mathrm{x}, \xi)] \quad \text{s.t.} \quad \mathrm{Prob}(\mathrm{g}(\mathrm{x}, \xi) \le 0) \ge \alpha \tag{10}$$

using sample average over N=30 scenarios:

$$\min_x \frac{1}{N} \sum_{n=1}^{N} \mathrm{TAC}(\boldsymbol{x}, \xi^{\boldsymbol{n}})$$

Penalized objective aggregates across scenarios:

$$\Phi_{\mathrm{SAA}}(x) = \frac{1}{N} \sum_{n=1}^{N} \Bigg[\mathrm{TAC}(x, \xi^n) + \rho_k \sum_i \max\left(0, g_i(\boldsymbol{x}, \xi^{\boldsymbol{n}})\right)^2 \Bigg]$$

By the law of large numbers, the sample-average objective converges to the true expectation as $N \to \infty$. Computational cost increases 30 times, but parallelizes trivially over scenarios.

Box-Constrained Robust Optimization

Robust optimization hedges against the worst-case realization within a deterministic uncertainty set, without invoking probability distributions. The formulation seeks a design that minimizes cost under the most adverse conditions within the hyperrectangle defined by physical parameter bounds, as expressed in equation (11).

$$\min_x P_{90}\,[\mathrm{TAC}(x, \xi)] \quad \text{s.t.} \quad P_{90}[g_i(x, \xi)] \le 0 \tag{11}$$

The uncertainty set is defined as the hyperrectangle bounded by physical parameter limits, as expressed in equation (12).

$$\widehat{P_{90}}[\mathrm{TAC}(x, \xi)] = \mathrm{TAC}_{(\lceil 0.9N \rceil)}(x) \tag{12}$$

The inner maximization is solved via vertex enumeration of the 2^6 = 64 box corners, followed by Nelder–Mead refinement to capture interior worst cases arising from model nonlinearities. By construction, this formulation provides deterministic feasibility guarantees for all realizations within the uncertainty set, without requiring probabilistic assumptions.

Wasserstein Distributionally Robust Optimization

WDRO hedges against distributional ambiguity [7], this is shown in equation (13).

$$\min_x \sup_{P \in \mathcal{D}} E_P\,[\mathrm{TAC}(x, \xi)] \tag{13}$$

The ambiguity set is defined as a Wasserstein ball centered on the empirical distribution, as given in equation (14).

$$\mathcal{D} = \left\{P : d_W\left(P, \widehat{P_N}\right) \le \varepsilon\right\} \tag{14}$$

The radius is calibrated via a confidence parameter to ensure that the true distribution lies within the ball with high probability:

$$\varepsilon = \sqrt{\frac{2}{N}\ln\left(\frac{1}{\alpha}\right)}\left(1 + \frac{1}{\sqrt{N}}\right)$$

With $\alpha = 0.05$ ensures that the true distribution lies within the Wasserstein ball with at least 95% probability. A tractable mean–variance approximation yields the surrogate objective in equation (15).

$$\Phi_{\mathrm{WDRO}}(x) = \frac{1}{N}\sum_{n=1}^{N}\mathrm{TAC}(x,\xi^n) + \lambda\sqrt{\frac{1}{N}\sum_{n=1}^{N}\left[\mathrm{TAC}(x,\xi^n) - \mu\right]^2} \tag{15}$$

Where $\lambda = \varepsilon/\sqrt{N}$ acts as a variance penalty, promoting cost consistency across distributional perturbations.

RESULTS AND DISCUSSION

Comparative Performance Analysis

Table 1 summarizes the out-of-sample performance of each method, evaluated on 2, 000 independent test scenarios. Cost statistics are reported without outlier filtering, preserving the tail behavior that uncertainty-aware methods are designed to address.

Table 1: Comparative Performance Metrics

Method	Avg Cost ($)	CV (%)	Test Feas (%)	Time (s)
Deterministic	10, 518	0.84	76.5	0.1
SAA	10, 553	0.91	76.7	3.8
WDRO	10, 833	0.49	77.4	3, 737
RO	28, 359	0.57	78.8	16.0

Deterministic optimization yields the lowest expected cost, as anticipated from a formulation that ignores uncertainty. When the resulting design is evaluated a posteriori under stochastic scenarios, it achieves 76.5% feasibility, a consequence of the penalty-based solver steering designs away from constraint boundaries, rather than an inherent robustness property. The cost variability is moderate (CV = 0.84%), and sub-second runtime enables rapid design-space exploration. However, this approach provides no formal guarantees on out-of-sample performance.

SAA delivers balanced performance at a negligible cost premium (+0.3%) over the deterministic baseline. Test feasibility (76.7%) is comparable to the deterministic result, suggesting that for this case study, SAA drives the optimizer toward a similar region of the design space while incorporating scenario-based penalty information. The modest computational overhead (3.8 s) remains practical for iterative design workflows.

WDRO achieves the tightest cost dispersion (CV = 0.49%) at a 3.0% cost premium, validating the variance-penalizing effect of the distributional robustness formulation. Feasibility improves to 77.4%, and the 95th-percentile cost is notably compressed relative to the mean. The distributional guarantee ensures reliable performance even when the true distribution deviates from the empirical estimate, a property particularly valuable for processes subject to seasonal or operational drift. The computational cost (~62 min) is substantially higher than SAA, reflecting the inner Nelder–Mead optimization required at each candidate evaluation.

Box-constrained RO yields the highest-cost design (US$28, 359, a 169% premium), reflecting the well-known conservatism of worst-case optimization. The optimizer selects a substantially larger shell diameter (D_s ≈ 0.96 m vs. 0.49–0.52 m for other methods) to guarantee feasibility across the entire uncertainty set. This conservatism produces the highest out-of-sample feasibility (78.8%) and the best performance under distributional drift (51.1% feasibility at κ = 1.0 vs. 45–46% for other methods). While the cost premium is prohibitive for routine applications, the RO paradigm is appropriate for safety-critical equipment where constraint violations carry severe consequences.

Table 2 compares the key geometric and thermal–hydraulic parameters of the optimal designs obtained by each method. Although the deterministic, SAA, and WDRO formulations all converge to compact geometries with similar heat transfer areas (131–145 m²), they differ in the balance between shell diameter and tube length. The deterministic design favors a short, wide configuration (D_s = 0.49 m, L = 2.5 m), while SAA opts for a narrower but longer exchanger (D_s = 0.45 m, L = 6.3 m) approaching the L/D limit. WDRO selects an intermediate geometry (D_s = 0.52 m, L = 5.4 m). In contrast, the RO design is fundamentally different: a shell diameter nearly twice as large (D_s = 0.96 m) with 596 m² of transfer area—roughly four times the other designs. This oversizing is the mechanism by which RO guarantees feasibility under worst-case conditions but at a substantial cost penalty. All designs operate within the velocity and pressure-drop constraints, with RO exhibiting the lowest tube-side velocity (0.86 m/s), close to the minimum bound.

Cost-Robustness Trade-offs

Figure 1 illustrates the cost–robustness trade-off. Deterministic and SAA cluster near US$10, 500 with CV around 0.84–0.91%, while WDRO achieves the lowest dispersion (CV = 0.49%) at a modest 3% cost premium. RO, at US$28, 359, lies far outside the cluster (indicated by arrow), reflecting the substantial cost of worst-case pro-

tection. The plot confirms that WDRO offers the most favorable cost–dispersion balance among uncertainty-aware methods.

Table 2: Optimal design parameters by optimization method

Parameter	Units	DET	SAA	RO	WDRO
Shell diameter	m	0.491	0.450	0.964	0.518
Tube length	m	2.54	6.33	10.41	5.44
Number of tubes	–	2583	518	1436	668
Heat transfer area	m^2	130.7	130.8	596.5	145.0
Tube velocity	m/s	1.12	1.20	0.86	0.93
Overall U	W/m^2K	1185	1185	260	1069
L/D ratio	–	5.2	14.1	10.8	10.5
TAC	US$/yr	10,506	10,548	28,367	10,830

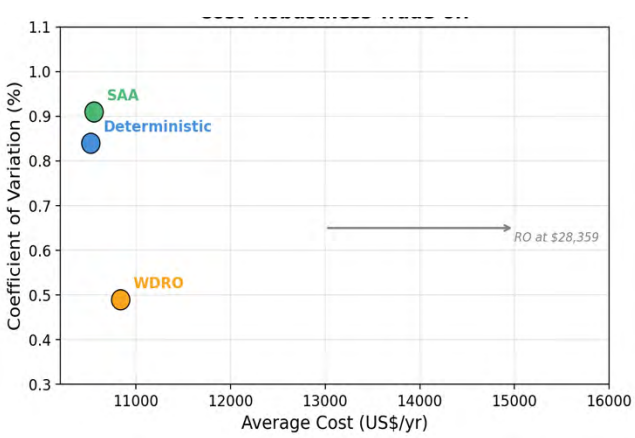

Figure 1. Cost–robustness trade-off across the four optimization approaches.

The observed CV range (0.49–0.91% for SP and DRO methods) indicates that cost variability is well controlled under the specified uncertainty levels. The 3.0% WDRO premium is modest relative to typical 10–20% industrial contingency factors, suggesting that formal distributional guarantees offer superior risk mitigation compared to ad-hoc safety margins.

Figure 2 compares feasibility across methods. Under nominal conditions (κ = 0), all methods achieve 76–79% feasibility, with RO performing marginally better due to its worst-case design margins. As distributional drift increases (κ → 1), feasibility degrades for all methods, but RO retains a consistent advantage (51% vs. 45–46%), confirming that worst-case guarantees translate into superior drift resilience.

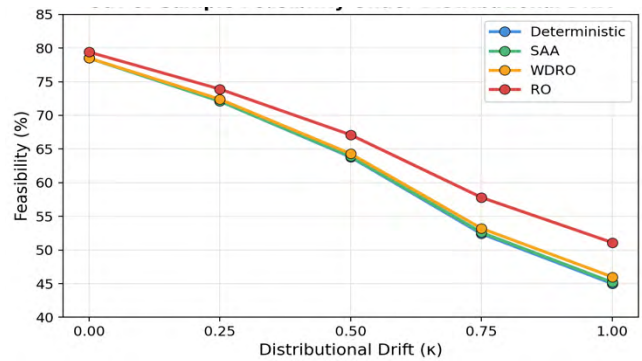

Figure 2. Out-of-sample feasibility as a function of distributional drift.

Method Selection Guidelines

Based on the comparative analysis, the following guidelines are proposed for practitioners. Deterministic optimization is suited to preliminary design and rapid prototyping where sub-second runtime enables extensive sensitivity analysis. It is acceptable when operational variability is low (CV < 5%) and equipment operates well within constraint boundaries, but should not be relied upon for safety-critical applications.

SAA is recommended as the default method for general industrial practice. It provides a principled framework for incorporating uncertainty with modest computational overhead (~4 s) and negligible cost premium. It is particularly effective when reliable historical data is available to characterize the uncertainty distribution and the process exhibits approximately stationary behavior.

WDRO is recommended for applications that demand cost predictability and robustness to distributional uncertainty. It is especially justified when training data is limited, the true distribution is imperfectly known or non-stationary, or regulatory requirements mandate high reliability. The 3% premium purchases formal distributional guarantees and the tightest cost dispersion among all methods tested.

RO is reserved for safety-critical applications where constraint satisfaction must be guaranteed across all foreseeable operating conditions. The worst-case focus provides deterministic feasibility guarantees but at a substantial cost premium (169% in this case study), making it impractical for routine design. RO is most appropriate when probability distributions are unavailable or unreliable and when the consequences of constraint violation are severe.

CONCLUSION

This study benchmarks four optimization approaches (spanning three paradigms for decision-making under uncertainty) on the same STHE design problem. The principal findings are: (1) deterministic optimization provides the lowest-cost baseline but offers no formal robustness guarantees; (2) SAA delivers comparable cost at negligible premium while incorporating scenario-based uncertainty information; (3) WDRO achieves the tightest cost dispersion (CV = 0.49%) at a modest 3% premium, confirming the practical value of distributional robustness when uncertainty characterization is imperfect; and (4) box-constrained RO guarantees feasibility across the entire uncertainty set but at prohibitive cost (169% premium), illustrating the conservatism inherent to worst-case formulations. Across all methods, the cost premium for uncertainty awareness is modest relative to typical industrial contingency factors, while providing formal guarantees that ad-hoc safety margins cannot.

ACKNOWLEDGMENTS

The authors gratefully acknowledge financial support from the Secretariat of Education, Science, Humanities, Technology, and Innovation (SESHITI, Mexico), computational resources provided by Universidad Autónoma de Querétaro, and institutional support from Universidad de Guanajuato.

REFERENCES

1. Sinnott and G. Towler, Chemical Engineering Design, Sixth. in Coulson and Richardson's Chemical Engineering Series. Elsevier Science & Technology, 2019.
2. Müller-Steinhagen H, Zhao Q. Investigation of low fouling surface alloys made by ion implantation technology. Chemical Engineering Science 52:3321-3332 (1997). https://doi.org/10.1016/s0009-2509(97)00162-0
3. Venkatesh B, Kiran A, Khan M, Rahmani MKI, Upadhyay L, Babu JC, Narayana TL. Performance optimization for an optimal operating condition for a shell and heat exchanger using a multi-objective genetic algorithm approach. PLoS ONE 19:e0304097 (2024). https://doi.org/10.1371/journal.pone.0304097
4. Han L, Yang K, Yang J, Li R, Li Y, Deng L, Che D. A thermal calculation model for tubular condensing heat exchanger. Applied Thermal Engineering 244:122701 (2024). https://doi.org/10.1016/j.applthermaleng.2024.122701
5. Zhao K, Zhao L, Lu W, Tang QQ, He C, Chen QL, Zhang BJ. Multi-objective optimization design of heat exchanger networks with simultaneous evaluation of economic performance and complexity via a novel two-step optimization approach. Computers & Chemical Engineering 200:109212 (2025). https://doi.org/10.1016/j.compchemeng.2025.109212
6. Shapiro A, Dentcheva D, Ruszczy?ski A. Lectures on stochastic programming. Society for Industrial and Applied Mathematics (2012). https://doi.org/10.1137/1.9780898718751
7. Mohajerin Esfahani P, Kuhn D. Data-driven distributionally robust optimization using the wasserstein metric: performance guarantees and tractable reformulations. Math. Program. 171:115-166 (2017). https://doi.org/10.1007/s10107-017-1172-1
8. Sudhanshu Rahul, Kumawat Piyush Kumar, and Chaturvedi Nitin Dutt, "Robust Optimization of Heat Exchanger Network with Uncertainty in Inlet Temperatures of Streams, " Chemical Engineering Transactions, vol. 88, pp. 307–312, Nov. 2021, doi: 10.3303/CET2188051.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Reactor network synthesis of enzymatic cascades using superstructure optimization

Swastik Chandra[+], Leandros Paschalidis[+], Siv Kinau, and Mirko Skiborowski*

Institute of Process Systems Engineering, Hamburg University of Technology, Am Schwarzenberg-Campus 4, Hamburg, 21073, Germany
* Corresponding Author: mirko.skiborowski@tuhh.de.
[+] These authors contributed equally

ABSTRACT

While classical heuristics can be applied to decide on the preferred reactor concept for simple reaction schemes, more complex reaction networks require more sophisticated methods, such as the multilevel reactor design approach or superstructure optimization. Based on an analysis of the existing methods a nonlinear programming framework for a superstructure-based reactor network synthesis is presented, emphasizing numerical robustness and flexible network representation without relying on integer decisions. The approach, which is implemented in GAMS, allows for the combination of continuous stirred-tank and cross-flow reactor models. An exemplary application for the classical Van de Vusse reaction is first shown for validation, prior to the application to an enzymatic cascade based on the Weimberg pathway. Assuming fast co-factor regeneration, the performance of the resulting PFR cascade, which can also be interpreted as a sequence of batch reactions, is compared with a commonly applied single batch reactor. The results show that the two-reactor configuration consistently achieves higher product formation by mitigating inhibition effects, demonstrating the potential of reactor network synthesis for complex reaction systems like enzymatic cascades.

Keywords: reactor network, NLP, GAMS, enzymatic cascades, reaction engineering

INTRODUCTION

The systematic design of reactor networks is a central and enduring challenge in chemical reaction engineering because the selection and interconnection of reactors fundamentally determines achievable yields, selectivities, and overall performance for complex reaction systems. Even when the problem is restricted to kinetics-based descriptions, the number of feasible reactor configurations grows rapidly with increasing system complexity, rendering the use of heuristics or intuitive or case-specific design approaches insufficient and highlighting the need for systematic, theoretically grounded methodologies for reactor network synthesis.

A unifying framework for assessing the limits of reactor performance is provided by the concept of the attainable region (AR), first introduced by Horn [1] as the "totality of physically possible reactors." The AR represents the complete set of achievable outlet states for a given reaction system under admissible operating conditions, enabling the identification of optimal reactor configurations through geometric arguments. Building on this concept, Feinberg and Hildebrandt [2, 3] demonstrated that the complete attainable region can be constructed from a finite set of elementary reactor building blocks—namely the continuously stirred tank reactor (CSTR), the plug-flow reactor (PFR), and the differential side stream reactor (DSR)—in combination with ideal mixers and splitters. A key implication of this result is that any optimal reactor performance can, in principle, be realized by a reactor network composed exclusively of these elementary units [4].

Despite its conceptual power, the explicit construction of the full attainable region is computationally demanding and becomes impractical for systems of realistic dimensionality. Consequently, over the past four decades, numerous optimization-based targeting methods have been developed to identify optimal or near-optimal reactor configurations without explicitly constructing the entire AR. These approaches follow two different lines,

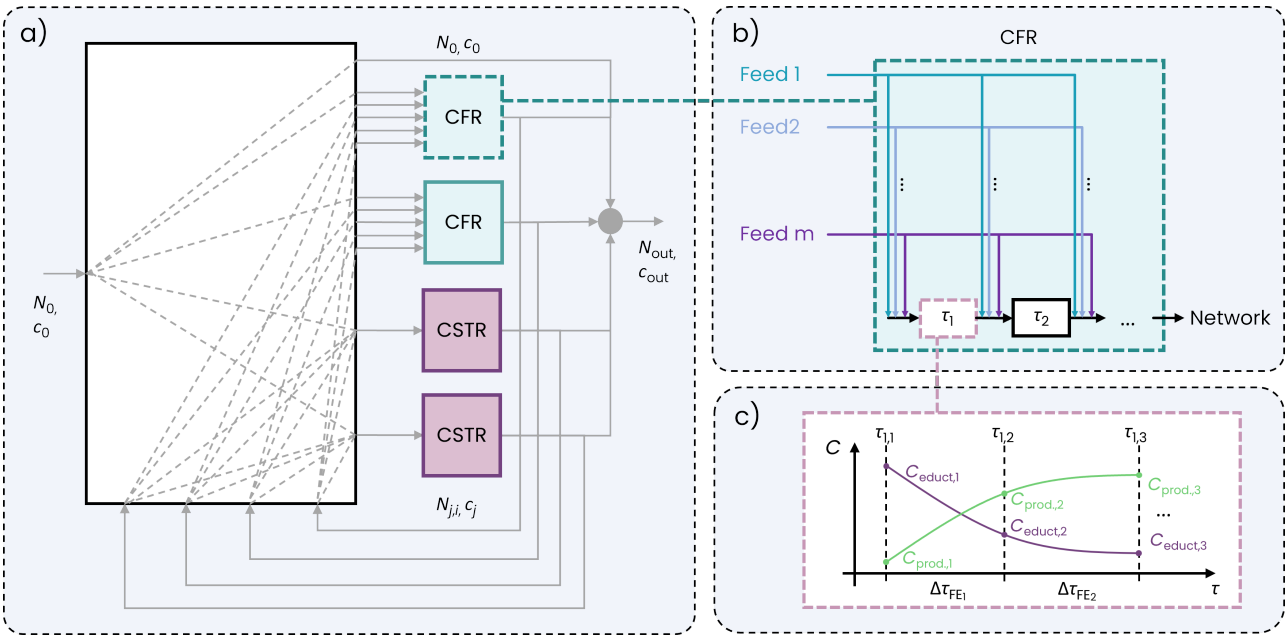

Figure 1. a) An example of the superstructure network for a network with 2 CFRs and 2 CSTRs. b) The connections of the multiple inlet streams inside the CFRs. c) The collocation method for each segment of the CFR (FE stands for finite element).

either pursuing a dynamic optimization of e.g. elementary process functions (EPF) [5], subsequently analyzing the result in a multilevel reactor design approach [6], or pursuing a superstructure-based optimization to determine the optimal reactor network. The latter typically formulate the reactor network synthesis problem as a nonlinear (NLP) or mixed-integer nonlinear programming (MINLP) problem. Examples include discretized recycle reactor modules that interpolate between PFR and CSTR behavior [7], combinations of CSTRs and PFRs [8, 9], networks of just CSTRs [10] and networks composed of CSTRs and DSRs with optional side streams [4, 11]. To render these models tractable, governing differential equations are transformed into algebraic equation systems, either by discretizing PFRs and DSRs as series of CSTRs [8, 12] or by employing direct transcription techniques such as orthogonal collocation on finite elements [4]. The resulting nonlinear and mixed-integer optimization problems are typically solved using local deterministic methods or stochastic metaheuristics, such as genetic algorithms, each with their own limitations in reliability and scalability.

The current contribution revisits the reactor network synthesis problem from both a modeling and numerical robustness perspective and proposes a refined superstructure model that:

- avoids unnecessary integer decisions while retaining the generality needed to capture arbitrary network connectivity,
- relies on continuous stirred tank and cross-flow reactors, and
- is solved by means of a polylithic solution approach in GAMS.

A key novelty of our work is that we demonstrate how many numerical difficulties traditionally associated with reactor network synthesis are not inherent to the problem itself but are instead induced by specific modeling choices.

To validate the proposed methodology and benchmark its performance against established approaches, we first illustrate the application to the classical Van de Vusse reaction network [14] that has been extensively investigated for reactor network synthesis and attainable region analyses.

As a second innovative application, we consider the design of an optimal reactor network for an enzymatic cascade. Such cascades involve a sequence of biocatalytic transformations catalyzed by multiple enzymes. They are of growing interest in biocatalysis and process intensification due to their potential for high selectivity and environmentally friendly operation. One example that is considered in the current study, such as the oxidative Weimberg pathway for pentose conversion to α-ketoglutarate [14], which builds on recent advances in in-vitro enzyme cascade design for metabolic pathways. Up to today, model-based optimization methods for enzymatic cascade reaction systems focus on optimizing the single pot approach [15-18]. Even though, these multi-step reactions can often be carried out in a single pot, this is not

necessarily optimal under all conditions. Inhibition or allosteric effects are only two possible reasons a more complex reactor network can be a beneficial choice. These systems provide an excellent application for reactor network synthesis methods because they combine strongly nonlinear kinetics with interlinked reaction steps.

METHODOLOGY

Model formulation

Figure 1 provides an illustration of the superstructure model used in this work. A generic reactor network superstructure following the state-space approach is formulated to allow arbitrary interconnections between feed streams, reactors, bypasses, and outlet streams via a distribution network. The individual reactors in the network are assumed to operate at steady state. An overall molar balance is imposed over the complete reactor network, ensuring consistency between all feed, reactor, bypass, and outlet flows. Unlike the work of Recker et al. [10], which considers solely ideal CSTRs, also tubular cross flow reactors (CFRs), which can implement a PFR or DSR [11] are considered as building blocks. The system of differential equations for the CFRs is modeled via a direct transcription using orthogonal collocation on finite elements, resulting in a NLP formulation. Intermediate feeding is feasible in between the finite elements, while a Radau IIA collocation captures axial variations in composition along the length of each finite element. For CSTR-only models' large cascades of CSTRs are required to accurately approximate the PFR performance. For kinetically sensitive systems, minor changes in residence time distribution and mixing assumptions can induce disproportionate deviations in the state trajectory, while networks with a large number of CSTRs pose combinatorial challenges for optimization. With the combination of CSTRs and CFRs a superstructure with two of these reactors is rich enough for most applications.

Unlike the EPF-approach that requires an interpretation of internal fluxes to identify CSTRs [19], the superstructure implements CSTRs as perfectly mixed reactors, with outlet compositions determined by the steady-state reaction rates and an explicit residence time. The superstructure formulation allows flexible stream splitting, mixing, and recycling, enabling the optimization of network topology, reactor selection, and sizing, while ensuring that material balances and component mole fractions are satisfied throughout the network.

Model implementation

The superstructure model is implemented in GAMS. A polylithic modeling and solution strategy is employed, which starts with the definition of the reaction kinetics that are included for each reactor type, with the flexibility to incorporate multiple reaction mechanisms. On this basis each individual reactor is initialized based on the overall feed and a small residence time. After the definition of potential additional constraints, e.g. for the definition of a minimum conversion or an upper bound for some intermediate composition, an initial solution for the full superstructure model is determined for a distribution network with a fixed structure that distributes the feed to all available reactors and does the same for the reactor products. After this initialization, the fixed split factors in the distribution network are relaxed and the reactor network optimization is performed.

Optimization of the superstructure is performed using nonlinear programming solvers (e.g., CONOPT), maximizing a specified performance objective, while satisfying all reactor and network constraints. Note that solver options, scaling factors, and finite-element discretization have to be selected to ensure numerical stability depending on the specific application. The respective model implementations for the subsequent case studies are provided via Gitlab (link in the appendix).

CASE STUDIES

Van de Vusse reaction

A commonly studied example for reactor network optimization is the so-called Van de Vusse reaction shown in Figure 2:

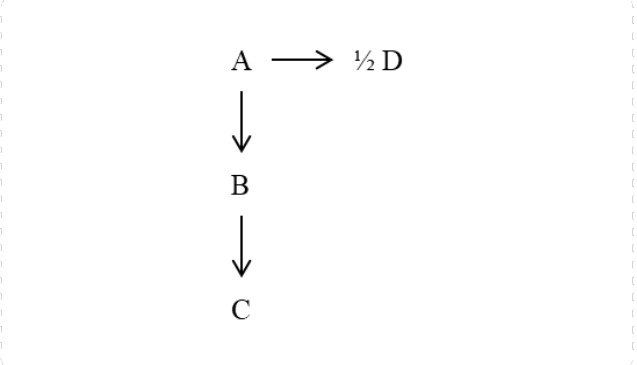

Figure 2. The chemical reaction network of the widely studied Van de Vusse reaction.

All reactions are catalyzed by the same catalyst. The kinetic and feed-related parameters are reflecting the case study 1 considered in the work of [19]. The Yield of B is selected as an objective to be maximized by selecting the specific reactor network structure and the residence time and temperature of the isothermal reactors.

Weimberg pathway

Reactor network synthesis is particularly interesting for applications to enzymatic cascades, as they involve multiple enzymes whose concentrations must be carefully balanced to achieve optimal performance. The

reaction network of the Weimberg pathway is shown in Figure 3. This cascade is well characterized kinetically, with each step catalyzed by a distinct enzyme. Two reactions in the pathway are cofactor dependent.

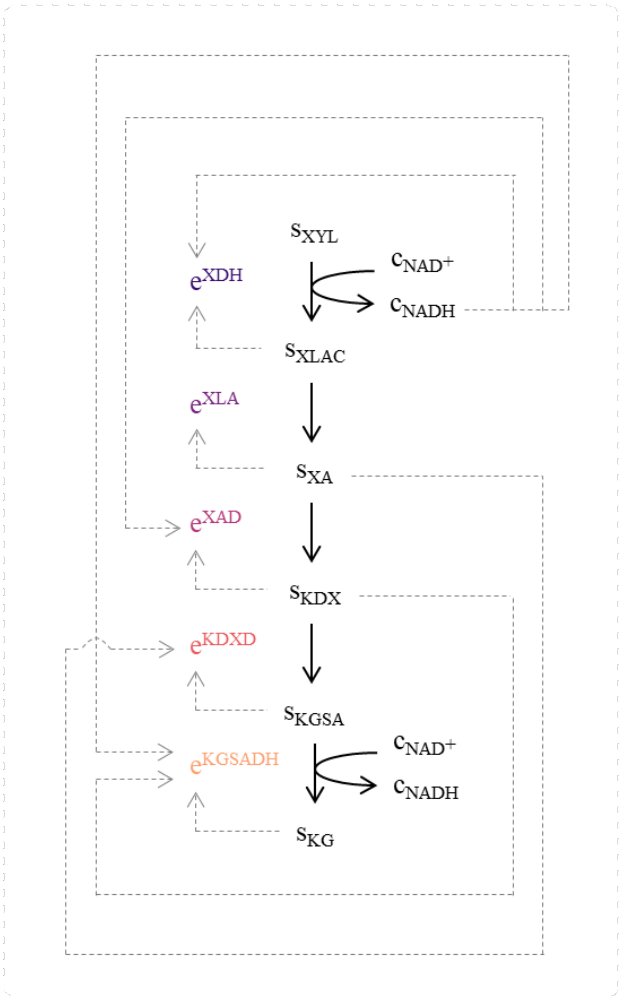

Figure 3: The chemical reaction network of the enzymatic cascade reaction of the Weimberg pathway. The black arrors show the reaction steps, while the grey, dashed arrows show inhibition effects.

Figure 3. also presents the substrate (s) and enzyme (e) abbreviations, together with the chemical reaction network and known inhibitory interactions between substrates and enzymes (dashed lines). A detailed kinetic model of the Weimberg pathway was developed by Shen et al. [14]. In their study, cofactor regeneration was implemented via a secondary reaction converting pyruvate to L-lactate catalyzed by L-lactate dehydrogenase. Under the assumption that this regeneration step proceeds rapidly—corresponding to the use of sufficiently high L-lactate dehydrogenase levels—we apply the proposed methodology to evaluate the performance gains achievable by employing multiple batch reactors instead of a single batch reactor for the Weimberg pathway. The well-known analogy between batch reactors and plug-flow reactors was used to apply our steady-state centric methodology to this dynamic problem.

The objective of the optimization problem is to maximize product formation (s_{KG}). CONOPT3 is employed as NLP solver in GAMS analyzing two scenarios: (i) a single-reactor configuration, in which the optimizer determines the optimal enzyme titers, and (ii) a two-reactor configuration in series, in which the optimizer determines the optimal allocation of a fixed enzyme budget between the reactors. The GAMS models are available in a GitLab repository (link provided in the appendix).

RESULTS AND DISCUSSION

Van de Vusse reaction

To validate the reactor network synthesis approach, the case study with parameters considered by Xie et al. [19] was solved. The resulting optimal reactor network matched the previously reported configuration, consisting of a CSTR followed by a PFR in series (cf. Figure 4), which also reflects the objective function and design variables reported by Xie et al [19] (objective function value: $Y_{A\rightarrow B} = 63.5\ \%$).

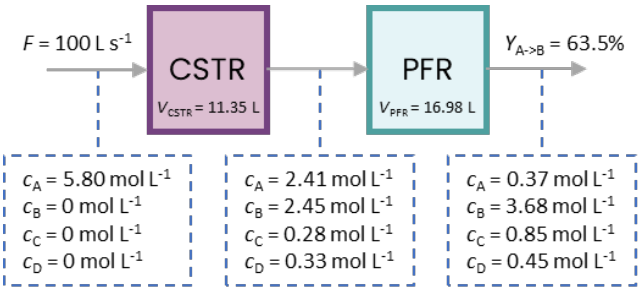

Figure 4. The results of the Van de Vusse case study.

Optimization of the Van de Vusse case study also allows for the identification of some hidden pitfalls for the flexible CFR model. If side-feeding and side-draws locations are feasible and the residence time in each finite element is treated as decision variable the solver exploits this by activating feed injections and effluent withdrawal at intermediate locations, effectively bypassing finite elements and reducing the numerical accuracy of the integration of the differential equation system. This concentrates residence time only in the most sensitive segments, resulting in a degenerate optimum: comparable objectives obtained with different apparent reactor sizes and partial volume utilization, which is not a faithful realization of the intended CFR behavior. Consequently, the different degrees of freedom in the CFR model should be used with care, to avoid such compromising bypass streams.

Weimberg pathway

The results of the comparison between the single batch reactor and the superstructure with two PFR (batch reactors) are summarized in Figure 5. The x-axis represents the total enzyme loading, while the y-axis denotes the total substrate concentration, for which each 3 distinct values have been evaluated. The percentage improvement in product formation achieved by the resulting two-reactor configuration is shown for each combination as a third dimension in the heat-map representation.

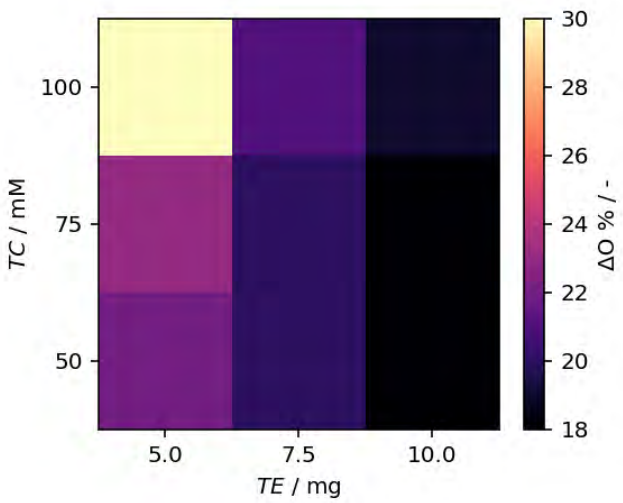

Figure 5. Comparison of the produced product with two batch reactors against the produced product with one reactor for different substrate and enzyme amounts in a raster-type contour plot.

For all conditions, the resulting two-reactor network consistently outperforms the single-reactor configuration, with up to 30% more product with the same enzyme budget. The performance advantage becomes increasingly pronounced at higher substrate loadings and lower total enzyme amounts.

Further insight is provided by examining the enzyme allocation strategy identified by the optimizer. Figure 6 illustrates the optimal enzyme distributions for the case of 75 mM substrate and 7.5 mg total enzyme. In addition to determining the optimal relative enzyme ratios, the optimizer also identifies the optimal partitioning of enzymes between the reactors when the two-reactor network is considered.

The optimizer preferentially distributes enzymes between the two reactors to mitigate inhibitory effects, particularly those affecting the final two enzymes in the cascade. This result is noteworthy, as it challenges the prevailing assumption in enzymatic cascade design that co-localizing all enzymes within a single reactor is invariably optimal.

CONCLUSIONS

A continuous NLP-based framework for reactor network synthesis was presented and demonstrated on both classical and biocatalytic reaction systems. By avoiding integer decisions and relying on carefully formulated CSTR and CFR models, the approach enables the systematic identification of optimal reactor configurations while maintaining numerical robustness.

Application to the Weimberg enzymatic cascade showed that splitting the reaction across two batch reactors can provide substantial kinetic advantages over conventional single-pot operation. The optimized solutions consistently favored enzyme distributions that alleviate inhibition effects, particularly in the downstream cascade steps, resulting in higher product formation under constrained enzyme budgets. This behavior challenges the commonly held assumption that enzymatic cascades should always be operated in a single reactor and highlights the importance of considering reactor network structure as an additional design degree of freedom.

Overall, the results demonstrate that reactor network synthesis can provide valuable insights into both achievable performance limits and optimal process configurations for complex reaction systems. When combined with careful modeling choices, the proposed framework offers a reliable and flexible tool for the design of advanced reactor networks in both chemical and biochemical applications.

DIGITAL SUPPLEMENTARY MATERIAL

The GAMS code used in this article is given in the following repository.

https://collaborating.tuhh.de/clp3832/escape-36-reactor-network-synthesis-of-enzymatic-cascades-using-superstructure-optimization/-/tree/b48a99412abe854ae54642805688e7451c7a8eff/

ACKNOWLEDGEMENTS

Financial support from the Collaborative Research Center (CRC) “SMART Reactors for Future Process Engineering, ” funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—project number 503850735—at the Hamburg University of Technology (TUHH), is gratefully acknowledged.

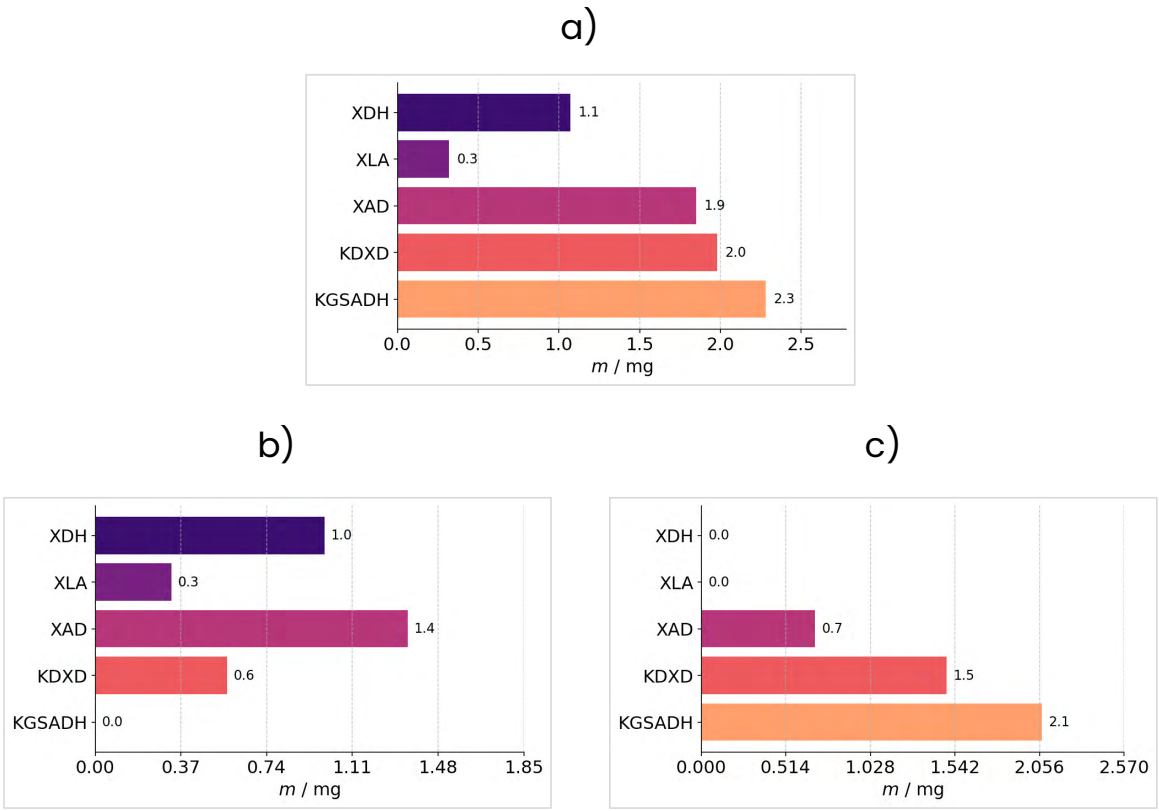

Figure 6. a) The suggested enzyme ratios in the single reactor network. b) The suggested enzyme ratios in the first reactor of the two reactor network. c) The suggested enzyme ratios in the second reactor of the two reactor network.

AUTHOR IDENTIFIERS

Author ORCIDs:
Leandros Paschalidis: 0000-0003-1578-4154
Mirko Skiborowski: 0000-0001-9694-963X

REFERENCES

1. G. Horn, “Attainable and Non-Attainable Regions in Chemical Reaction Technique, ” Proc. Eur. Symp. on Chemical Reaction Engineering, Pergamon Press, London, 1964.
2. Feinberg M. Toward a theory of process synthesis. Ind. Eng. Chem. Res. 41:3751-3761 (2002). https://doi.org/10.1021/ie010807f
3. Hildebrandt D, Glasser D. The attainable region and optimal reactor structures. Chemical Engineering Science 45:2161-2168 (1990). https://doi.org/10.1016/0009-2509(90)80091-r
4. Lakshmanan A, Biegler LT. Synthesis of optimal chemical reactor networks. Ind. Eng. Chem. Res. 35:1344-1353 (1996). https://doi.org/10.1021/ie950344b
5. Freund H, Sundmacher K. Towards a methodology for the systematic analysis and design of efficient chemical processes. Chemical Engineering and Processing: Process Intensification 47:2051-2060 (2008). https://doi.org/10.1016/j.cep.2008.07.011
6. Freund H, Maußner J, Kaiser M, Xie M. Process intensification by model-based design of tailor-made reactors. Current Opinion in Chemical Engineering 26:46-57 (2019). https://doi.org/10.1016/j.coche.2019.08.003
7. L. E. K. Achenie and L. T. Biegler, “Algorithmic synthesis of chemical reactor networks using mathematical programming, ” Ind. Eng. Chem. Fundam., vol. 25, no. 4, pp. 621–627, 1986.
8. Gholizadeh Z, Soltani H, Javid M, Azar MS. A new robust approach for reactor network synthesis by combination of mathematical method and NSGAII. International Journal of Chemical Reactor Engineering 18: (2020). https://doi.org/10.1515/ijcre-2019-0090
9. Wen Y, Biegler LT, Ochoa MP, Matthews L, Ferrio J, Weston J, Nikbin N. Continuous reactor network design for rigid polyol production. Chemical Engineering Science 230:116189 (2021).

https://doi.org/10.1016/j.ces.2020.116189
10. Recker S, Skiborowski M, Redepenning C, Marquardt W. A unifying framework for optimization-based design of integrated reaction–separation processes. Computers & Chemical Engineering 81:260-271 (2015). https://doi.org/10.1016/j.compchemeng.2015.03.014
11. Schweiger CA, Floudas CA. Synthesis of optimal chemical reactor networks. Computers & Chemical Engineering 23:S47-S50 (1999). https://doi.org/10.1016/s0098-1354(99)80013-1
12. Marcoulaki E, Linke P, Kokossis A. Design of separation trains and reaction-separation networks using stochastic optimization methods. Chemical Engineering Research and Design 79:25-32 (2001). https://doi.org/10.1205/026387601528499
13. Biki? D, Butinar B, Glavi? P. Optimal reactor systems for van de vusse reaction scheme with multicomponent feed. Computers & Chemical Engineering 26:1335-1343 (2002). https://doi.org/10.1016/s0098-1354(01)00759-1
14. Shen L, Kohlhaas M, Enoki J, Meier R, Schönenberger B, Wohlgemuth R, Kourist R, Niemeyer F, van Niekerk D, Bräsen C, Niemeyer J, Snoep J, Siebers B. A combined experimental and modelling approach for the weimberg pathway optimisation. Nat Commun 11: (2020). https://doi.org/10.1038/s41467-020-14830-y
15. Paschalidis L, Beer B, Sutiono S, Sieber V, Burger J. Design of enzymatic cascade reactors through multi-objective dynamic optimization. Biochemical Engineering Journal 181:108384 (2022). https://doi.org/10.1016/j.bej.2022.108384
16. L. Paschalidis, D. Fröschl, M. Ibañez, S. Sutiono, V. Sieber, and J. Burger, “Boosting of enzymatic cascades by intermediates: theoretical analysis and model-based optimization, ” Biochem. Eng. J., p. 109440, 2024.
17. Paschalidis L, Arana?Peña S, Sieber V, Burger J. Modeling enzymatic cascade reactions immobilized in plug?flow reactors for flow biocatalysis. Chemie Ingenieur Technik 96:741-748 (2024). https://doi.org/10.1002/cite.202300159
18. Milker S, Fink MJ, Oberleitner N, Ressmann AK, Bornscheuer UT, Mihovilovic MD, Rudroff F. Kinetic modeling of an enzymatic redox cascade in vivo reveals bottlenecks caused by cofactors. ChemCatChem 9:3420-3427 (2017). https://doi.org/10.1002/cctc.201700573
19. Xie M, Freund H. Fast synthesis of optimal chemical reactor networks based on a universal system representation. Chemical Engineering and Processing - Process Intensification 123:280-290 (2018). https://doi.org/10.1016/j.cep.2017.11.011

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Designing Multi-Objective Optimization Models for Vaccine Supply Chains: Economic, Environmental, and Social Trade-offs in the COVID-19 Context

Jonathan Jair Cuevas Lopez[a], Sofía De-León Almaraz[b], Alberto A. Aguilar Lasserre[c], Catherine Azzaro-Pantel[a]

[a] Laboratoire de Génie Chimique, Université de Toulouse, CNRS, INPT, UPS, Toulouse, France
[b] Corvinus University of Budapest, Institute of Operations and Decision Sciences, 8 Fővám tér, 1093 Budapest, Hungary
[c] Tecnológico Nacional de México, Instituto Tecnológico de Orizaba, Orizaba, Veracruz, México
[*] Corresponding Author: catherine.azzaropantel@toulouse-inp.fr

ABSTRACT

Pharmaceutical supply chains face increasing pressure to deliver high service levels while meeting environmental and social expectations. Vaccine supply chains amplify these challenges due to strict cold-chain requirements, demand uncertainty driven by acceptance and preferences, and the urgency of public-health objectives. This paper develops a multi-objective mixed-integer linear programming (MILP) framework for national-scale vaccine distribution that explicitly integrates economic cost, service level, greenhouse-gas emissions, and population-level vaccine effectiveness. Behavioral realism is incorporated by modeling vaccine acceptance and brand preferences as operational constraints rather than ex-post indicators. Trade-offs are explored using an ε-constraint method that preserves the MILP structure and enables systematic recovery of Pareto-optimal solutions. The framework is applied to a 52-week national case study for metropolitan France during the 2021 COVID-19 vaccination campaign, focusing on Pfizer-BioNTech and Moderna mRNA vaccines. Results show that population-level effectiveness saturates early once acceptance-driven coverage is achieved, making cost and emissions the primary decision levers under service guarantees. Balanced solutions maintain near-maximal protection while achieving substantial emission reductions at moderate additional cost, whereas deeper decarbonization requires explicit service concessions. The proposed framework provides transparent, policy-relevant insights and supports informed decision-making in large-scale vaccine logistics.

Keywords: Vaccine supply chain, Multi-objective optimization, MILP, Sustainability, ε-constraint

INTRODUCTION

Pharmaceutical supply chains (PSCs) are a cornerstone of global healthcare systems, ensuring the timely delivery of essential medicines, vaccines, and other healthcare products. Their performance directly conditions public health outcomes, particularly under conditions of uncertainty and time pressure. Over recent decades, PSCs have become increasingly complex due to stringent quality requirements, strict regulatory oversight, and the perishable nature of many pharmaceutical products [1–3]. Additional challenges arise from demand variability, emergency preparedness, and growing expectations regarding sustainability and equity [4].

Within this broad pharmaceutical context, vaccine supply chains (VSCs) constitute a particularly critical and challenging subset. Vaccines are highly perishable and often require strict cold-chain management, sometimes at ultra-low temperatures (e.g., mRNA vaccines). Distribution performance is tightly linked to public acceptance and preferences, which directly affect uptake and, ultimately, immunization. At the same time, energy-intensive storage and transport generate non-negligible environmental impacts that are rarely considered explicitly in distribution planning. Unlike other pharmaceuticals, vaccines are often deployed at national or global scale during public health emergencies, where both operational efficiency and social equity are paramount.

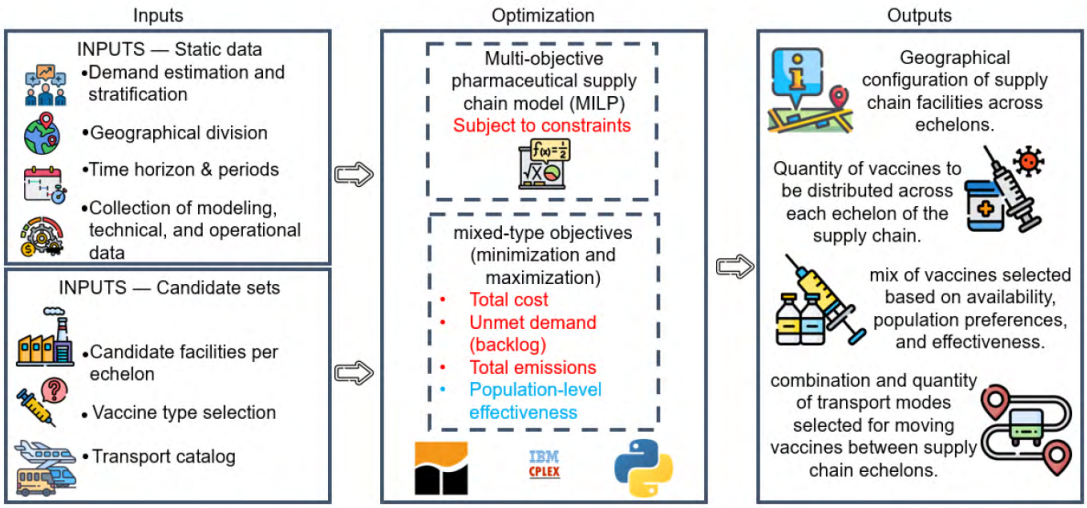

Figure 1. General framework of the proposed optimization model: inputs (static data and candidate sets) feed the MILP, which produces outputs.

The COVID-19 pandemic highlighted these challenges at an unprecedented scale. Beyond the health crisis, it exposed structural weaknesses in vaccine logistics when systems are stressed by uncertainty, compressed timelines, and capacity constraints. The emergency deployment of mRNA vaccines at the end of 2020 introduced ultra-low temperature requirements that amplified bottlenecks in cold-chain infrastructure, staffing, testing and release, and transportation. As a result, vaccine distribution emerged not only as a medical issue, but as a large-scale operation and decision-support problem, where trade-offs between speed, coverage, cost, and environmental performance became explicit.

Despite significant advances in pharmaceutical and vaccine supply chain modeling, several limitations remain:

- Limited integration of social and behavioral factors, such as public acceptance and preference heterogeneity, in optimization-based approaches;
- Partial treatment of sustainability, with greenhouse gas emissions rarely assessed alongside economic and operational objectives;
- Scarcity of empirically grounded, national-scale models combining demographic data, logistics constraints, and cold-chain requirements.

This work addresses these gaps by proposing a national-scale, multi-echelon, multi-objective optimization framework for vaccine distribution. The model explicitly integrates cold-chain and capacity constraints, behavioral features related to vaccine acceptance and preferences, and multiple performance indicators, including cost, backlog, emissions, and vaccination effectiveness. The framework is designed as a decision-support tool, enabling the exploration of trade-offs rather than the identification of a single optimal solution.

A key insight emerging from the analysis is that vaccination effectiveness saturates early, making cost and emissions the primary decision levers for large-scale deployment strategies. This result highlights the importance of combining epidemiological effectiveness with economic and environmental criteria when designing resilient and sustainable vaccine distribution policies.

The remainder of this paper is organized as follows. Section 2 presents the proposed multi-echelon MILP framework, including the network structure, operational constraints, objective functions, and the ε-constraint solution strategy used to explore trade-offs among cost, service, emissions, and effectiveness. Section 3 describes the national-scale case study for metropolitan France during the 2021 COVID-19 vaccination campaign, detailing data sources, assumptions, and calibration choices. Section 4 reports and discusses the main results, with a focus on Pareto-optimal trade-offs and their implications for large-scale vaccine deployment strategies. Finally, Section 5 concludes with key insights for decision-makers and outlines perspectives for extending the framework.

METHODOLOGICAL FRAMEWORK

The proposed optimization framework models a pharmaceutical supply chain (PSC) in a structured, multi-echelon, multi-period setting. The chain is represented as a sequence of operational stages, capturing flows from origin to end users under capacity, inventory, and routing

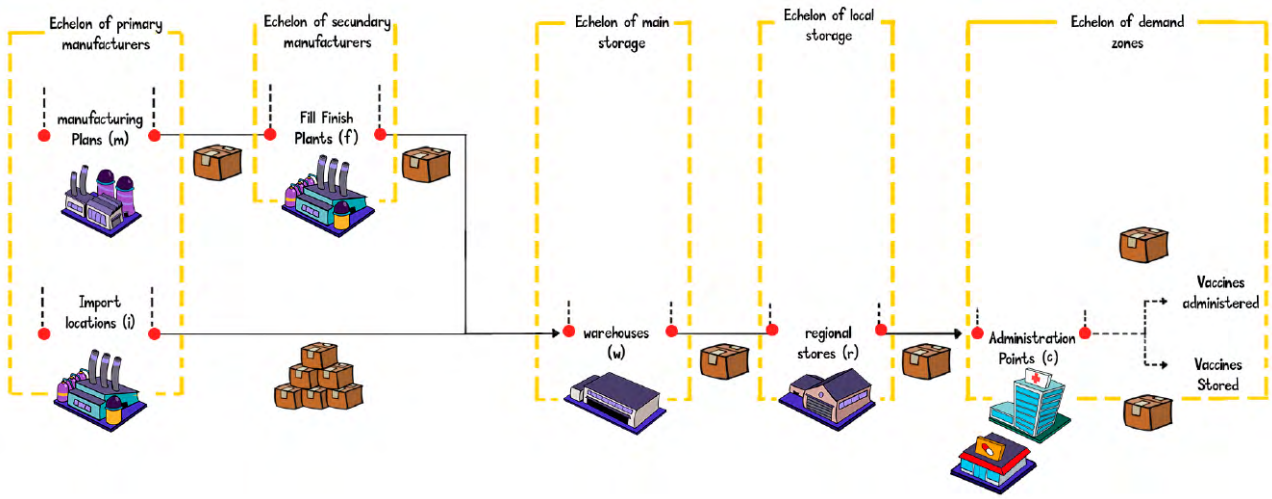

Figure 2: Schematic representation of the five-echelon pharmaceutical supply chain network (nodes and arcs) used in the MILP formulation.

constraints. Figure 1 illustrates the modeling pipeline, linking input data, the MILP optimization core, and resulting outputs. Inputs include static data (e.g., demand, geographical information, time horizon, costs, emission factors) and candidate sets (e.g., facility options, vaccine types, transport modes) from which the model selects. The optimization core is formulated as a Mixed-Integer Linear Program (MILP) with binary, integer, and continuous decision variables, solved using GAMS/CPLEX. Outputs comprise facility locations, vaccine allocation over time, selected vaccine mix, and transport configuration.

Network structure and flows

The network comprises five echelons: primary manufacturers, secondary manufacturers (fill–finish), main storage, regional storage, and demand zones (administration points). Vaccines flow along the sequences *(manufacturer → fill–finish → warehouse)*, *(import → warehouse)*, and *(warehouse → regional → administration)*, ensuring consistent representation of sourcing, processing, storage, and last-mile distribution over multiple periods (Figure 2).

MILP formulation and operational constraints

Key decisions include facility activation, vaccine flows, transport logistics, and inventory management. Binary variables indicate active facilities and routes, integer variables capture shipment counts and number of facilities. Continuous variables describe vaccine flows, inventory, and administered doses. Slack variables allow minor deviations to preserve feasibility.

Core assumptions include deterministic data (demand, travel times, costs, capacities), discrete planning periods, strict cold-chain compliance, and stable inventory without spoilage beyond standard holding costs. Operational constraints ensure feasibility:

- **Routing and transport:** shipments occur only if facilities and routes are active; shipment sizes respect predefined bounds.
- **Inventory:** standard stock–flow identity with min/max bounds enforced at operational facilities.
- **Processing:** production, fill–finish, and imports respect throughput limits; transport trips are derived from shipment quantities relative to maximum capacity.

Objective functions

The model considers four main objectives to capture vaccination campaign priorities:

1. **Economic efficiency (cost minimization):** Includes production/procurement, facility activation, transport, cold-chain, inventory, administration, and unmet demand penalties. Fixed and variable costs across all echelons are aggregated into a single additive objective [1].
2. **Service-level compliance (backlog minimization):** Defined as the total number of doses not administered despite demand. Backlog is computed across all vaccine types, locations, and periods. Minimizing backlog serves as a proxy for service reliability and ensures high-risk populations are protected [5].
3. **Environmental sustainability (GHG minimization):** Accounts for emissions from transport (van, trailer, airplane), refrigeration, dry ice, and packaging materials. Fuel consumption and emissions depend on distance, vehicle load, and speed, justifying a dedicated environmental objective [6].
4. **Population-level vaccine effectiveness** (VE) **(maximization):** Consistent with overall (population)

effectiveness measures used in vaccine study design, weighting administered doses by product-specific efficacy provides a tractable proxy for population protection—even when indirect (herd) effects are not explicitly modeled [7].

$$VE(\%) = 100 \times \frac{\sum_{v \in V, c \in C, t \in T} q_{v,c,t} . ef_v}{tpd} \quad (1)$$

where:

V = set of vaccine types

C = set of vaccination centers (demand zones)

T = set of planning periods

$q_{v,c,t}$ = number of doses of vaccine v administered at center c during period t

ef_v = clinical efficacy of vaccine v

tpd = total planned doses for the campaign

This metric integrates coverage and per-dose effectiveness into a single, linear indicator that can be incorporated alongside cost, backlog, and emissions in a multi-objective optimization framework.

Multi-objective solution approach: the ε-constraint method

The vaccine supply chain model involves multiple, often conflicting objectives: Cost (min), Backlog (min), Emissions (min), and Effectiveness (max). Using a single weighted sum may miss relevant trade-offs, especially in mixed-integer problems. We employ the ε-constraint method, which: (i) optimizes one objective (primary) while converting the others into constraints; (ii) preserves the MILP structure and ensures extreme Pareto points can be recovered; (iii) avoids heuristic weight tuning. The implementation follows best-practice recommendations from [8].

The solution procedure follows four main steps:

1. **Payoff table (ideal and nadir anchors)**
   - Optimize each objective individually and record all objective values.
   - Compute lower/upper anchors for each objective to define ε bounds.
2. **Grid of ε values**
   - Generate a uniform grid of ε points between the anchors.
   - Provides a transparent schedule of target levels for the constrained objectives.
3. **ε-constrained solves**
   - At each grid point, optimize the primary objective while enforcing ε constraints for the others.
   - Retain solutions only if the MILP reaches optimality (integer-optimal).
4. **Rotate the primary objective**
   - Cycle the primary objective across runs to explore the Pareto frontier evenly.
   - Reduces bias and avoids gaps in the recovered solutions.

In the national-scale application, 50 grid points (g_o = 1, ..., 50) were generated using the ε-constraint method. Only solutions that achieved full optimality—i.e., feasible and meeting the solver's convergence criteria—were retained for analysis, ensuring that all reported trade-off points are valid and reliable.

Trade-offs were quantified by identifying non-dominated (Pareto) solutions. Dominance was assessed in the full four-dimensional space of Cost (min), Backlog (min), Emissions (min), and Effectiveness (max) [13]. A solution A dominates B if it is no worse in every objective and strictly better in at least one. Conversely, a solution is non-dominated if no other solution dominates it. When applying an ε-constraint grid, weakly efficient points were filtered out to ensure that the retained set represents true Pareto candidates [8].

This approach highlights configurations that are not strictly inferior to any others across all objectives and focuses the analysis on a compact set of meaningful alternatives.

CASE STUDY: FRANCE (2021)

France was selected for this study due to its extensive availability of detailed public vaccination data, mature logistics infrastructure, and clearly defined administrative regions. During the COVID-19 pandemic, collaboration among public health institutions, statistical offices, and other agencies ensured high-quality, timely, and comprehensive open datasets, enabling accurate data-driven analysis [10].

Within the European context, France faced complex logistical and public health requirements, including cross-border vaccine imports, coordination with EU procurement initiatives, and compliance with stringent cold-chain conditions for large-scale distribution [11]. The combination of detailed open data, multi-tiered administrative structures, and a robust logistics system makes France a suitable case for modeling vaccine distribution, capturing interactions across import flows, warehouses, regional hubs, and local administration points.

Geographic Scope and Spatial Representation

This study considers metropolitan France only: 13 regions, 96 departments, and approximately 34, 945 communes. Administrative divisions allow efficient mapping and resource allocation, enabling accurate distance calculations and detailed logistical planning [12]. Spatial

fragmentation was applied across echelons:

- National/port nodes for imports (i),
- Plant locations (f),
- Regional hubs (w),
- Department-level stores (r),
- Commune-level administration points (c).

Clusters of communes follow official geographical breakdowns to consolidate nearby localities while maintaining reporting consistency. The analysis focuses on the population aged 15 years and older.

Vaccine selection

This study focuses on two widely deployed mRNA vaccines, Pfizer-BioNTech (BNT162b2) and Moderna (mRNA-1273), selected for their central role in the French vaccination campaign, the availability of detailed logistical data, and contrasting cold-chain requirements. Pfizer-BioNTech requires ultra-cold storage and specialized transport, while Moderna can be stored at less stringent temperatures, allowing for the evaluation of different logistical strategies. National delivery records corroborate this choice. Excluding Pfizer's paediatric presentation, Pfizer and Moderna together account for >96% of cumulative doses delivered in metropolitan France up to 2024 and approximately 98% during 2021, the core year modelled. The predominance of mRNA products thus makes the focus on BNT162b2 and mRNA-1273 representative for modelling the French campaign [14].

Supply chain nodes

Nodes represent specific points in the vaccine supply chain, organized into five echelons in the generic model:

1. Primary manufacturers,
2. Secondary manufacturers (fill–finish),
3. Main storage (warehouses),
4. Local storage (regional stores),
5. Demand zones (administration points).

For France, manufacturing and fill–finish are inactive (no domestic facilities), so all inbound supply enters via import locations: Puurs (Belgium, Pfizer) and Visp (Switzerland, Moderna). The operational chain reduces to i → w → r → c. Warehouse candidates were selected among departments exceeding a 3 M-dose cumulative delivery threshold. Regional stores correspond to the 13 administrative regions, and administration centers to the 22 former regional capitals. Figure 3 illustrates the spatial organization of the national vaccine supply chain, highlighting the geographical alignment between echelons and the main logistic flows across the French territory.

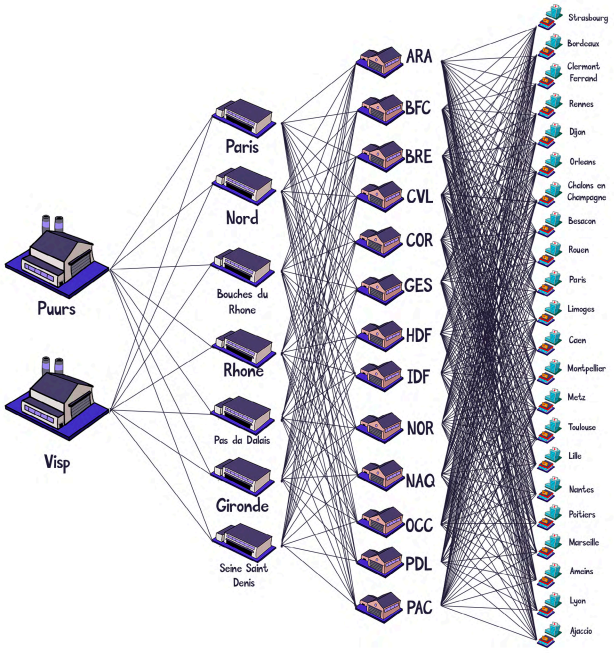

Figure 3. National supply chain structure and geographical alignment across echelons.

Demand Parameters

The demand parameter for the vaccine supply chain model was based on demographic data for individuals aged 15 years and older. The year 2021 was chosen to align with the vaccine shipment timeline, with the first shipment arriving on 27 December 2020 and the initial vaccination recorded on 5 January 2021. The model simulates a 52-week period corresponding to this timeframe, ensuring temporal consistency.

Demographic data were extracted from the INSEE database and aggregated according to the department codes comprising the 22 former administrative regions. Several departments were grouped within a former region and labeled using the name of the main city in the cluster. The selected labels were: Strasbourg, Bordeaux, Clermont-Ferrand, Rennes, Dijon, Orléans, Châlons-en-Champagne, Besançon, Rouen, Paris, Limoges, Caen, Montpellier, Metz, Toulouse, Lille, Nantes, Poitiers, Marseille, Amiens, Lyon, and Ajaccio.

Population Acceptance and Hesitancy

Vaccine acceptance and hesitancy rates were derived from representative international surveys [9], with global averages of approximately 79% acceptance and 21% hesitancy. Booster acceptance among already vaccinated individuals averaged 88%, with a hesitancy rate of 12%.

For the French case study, these values were adapted to the two main mRNA vaccines, Pfizer-

BioNTech and Moderna, based on their relative usage. (Pfizer = 85.3%, Moderna = 14.7%). To avoid over-constraining the model, these normalized values are then implemented as upper bounds (preference caps) with a small modeling slack: 80% for Pfizer and 10% for Moderna. This preserves empirical guidance while allowing flexibility in the optimization.

RESULTS AND DISCUSSION

ε-constraint exploration and anchor solution

The multi-objective problem was explored using an ε-constraint grid composed of 50 feasible configurations. For each grid point, one objective was optimized while the remaining ones were enforced through ε-bounds. Only solutions reaching proven integer optimality were retained, ensuring that all reported configurations correspond to valid and comparable trade-offs.

For the national-scale French instance (5.088M equations and 1.941M variables, solved with 2 threads), the four single-objective baseline runs required 0.01–5.82 h and closed with MIP gaps between 0.00% and 0.153%; the representative balanced Pareto point $g_0 = 50$ required 7.43 h and closed with a 0.14% MIP gap.

As a preliminary step, each objective was optimized independently to compute the corresponding single-objective optima. These solutions, referred to as *anchor solutions*, delineate the extreme boundaries of the attainable performance space—namely the minimum cost, minimum backlog, minimum emissions, and maximum effectiveness configurations (Table 1).

Table 1: Single-objective optima (SO) used as anchor points for ε-grid construction

Objective function	Cost (USD)	Backlog (doses)	Emissions (kgCO2e)	Effectiveness (%)
Cost	2.339B	120.149K	7.540M	71.39
Backlog	2.427B	118.847K	8.113M	71.38
Emissions	2.707B	8.349M	7.095M	63.92
Effectiveness	3.040B	118.847K	8.457M	71.3

Beyond defining the ε-grid bounds, they provide natural reference policies against which multi-objective compromises can be interpreted. Recovering these extremes within the constrained runs also validates the correctness of the ε-constraint implementation.

These anchor solutions clearly illustrate the inherent trade-offs: minimizing cost, for instance, does not guarantee low emissions or maximal effectiveness, highlighting the need for a multi-objective perspective.

Structure of the trade-off space

The global structure of the trade-offs is illustrated in Figure 4, which shows all feasible grid points across the four objectives. A parallel coordinates plot is used to visualize the multi-dimensional performance, highlighting patterns across cost, backlog, emissions, and effectiveness. Two consistent trends emerge.

First, within the retained feasible ε-grid, population-level effectiveness exhibits very limited variability, ranging from approximately 70.5% to 71.4%. This indicates that, once acceptance-driven coverage and the induced vaccine mix are satisfied, the system enters a high-protection plateau in which additional logistical effort yields negligible incremental protection. Effectiveness is therefore retained not because it strongly shapes the local Pareto geometry, but because it acts as a public-health safeguard: it certifies that candidate policies remain in the high-protection regime while cost and emissions are traded off. This distinction matters because the single-objective anchors used to define the ε-bounds span a wider effectiveness interval, and the emissions-optimal anchor drops to 63.92%, showing that stronger environmental pursuit can indeed decrease protection when service is relaxed.

Second, most variation occurs along the economic and environmental dimensions. Cost and emissions become the dominant levers shaping the frontier, revealing that post-coverage decisions are primarily logistical rather than epidemiological. This shift indicates that improving sustainability or reducing expenditure involves reallocating storage, routing, and transport choices rather than increasing administered doses.

Non dominated solutions

Table 2 reports the non-dominated configurations extracted from the feasible ε-grid. Among the 50 evaluated points, only three solutions remain Pareto-efficient, indicating that most intermediate configurations are dominated. This small set substantially simplifies interpretation and provides decision-makers with a concise portfolio of representative policies, ranging from cost-oriented to environmentally efficient strategies.

Table 2: Four-dimensional non-dominated solutions extracted from the feasible ε-grid.

g_0	Cost (USD)	Backlog (doses)	Emissions (kg CO_2e)	Effectiveness (%)
38	2.367B	118.847K	7.943M	71.31
44	2.365B	162.653K	7.872M	71.35
50	2.369B	118.847K	7.858M	71.39

Among these, $g_0 = 50$ clearly stands out as the most policy-relevant "clean-and-served" solution. It achieves the lowest carbon footprint (7.858M kg CO_2e), meets the

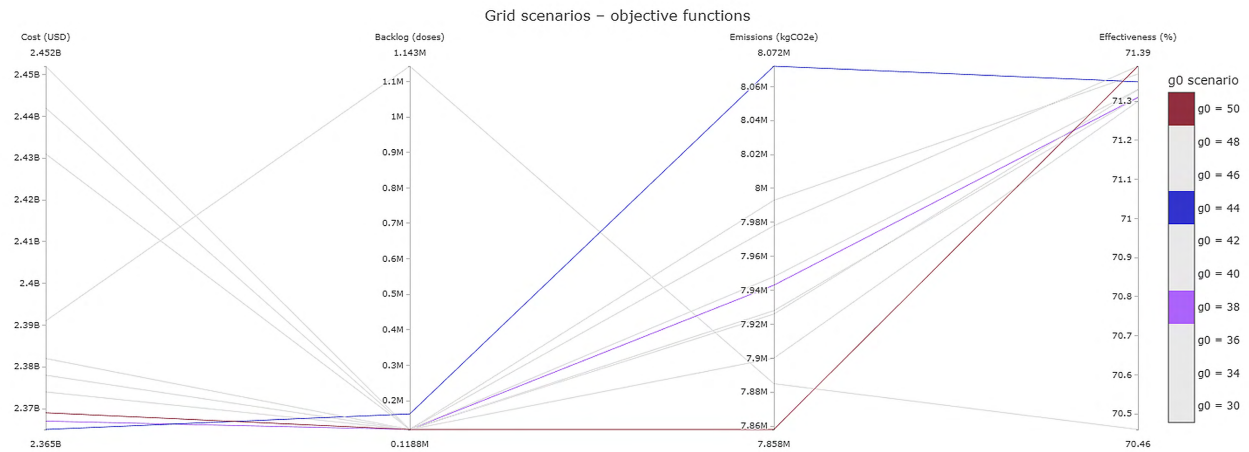

Figure 4. Trade-offs across the four objectives (cost, backlog, emissions, effectiveness).

backlog floor (118.847K doses), and attains the highest effectiveness (71.39%), all while maintaining total cost near the lower bound (2.369B USD). This configuration illustrates how tight ε-constraints on backlog and cost, combined with emissions as the primary objective, can produce a solution that simultaneously balances service quality, cost efficiency, effectiveness, and environmental performance.

The other non-dominated solutions highlight the trade-offs involved in prioritizing a single objective. g_0 = 38 also respects the backlog floor but exhibits slightly higher emissions (7.943M kg CO_2e) and marginally lower effectiveness (71.31%). By contrast, g_0 = 44 achieves the absolute minimum cost (2.365B USD) but at the expense of a higher backlog (162.653K doses) and increased emissions (7.872M kg CO_2e). Compared to these alternatives, g_0 = 50 occupies the most balanced region of the multi-objective frontier, providing the best combination of cost, service, effectiveness, and environmental impact.

CONCLUSIONS AND PERSPECTIVES

This study developed a multi-objective optimization framework for national-scale vaccine supply chains, explicitly integrating economic, environmental, and public-health objectives while capturing real operational constraints. By combining a multi-period, multi-echelon MILP with cold-chain realism and behavior-aware elements reflecting vaccine acceptance and preferences, the framework provides a practical tool to explore trade-offs in large-scale vaccination campaigns. The ε-constraint method enabled systematic recovery of Pareto-optimal solutions without altering the linear structure of the problem, ensuring methodological rigor and operational relevance.

Application to the French national case study yields several key insights. Population-level vaccine effectiveness saturates early once acceptance-driven coverage is achieved, making cost and emissions the main levers for decision-making under service guarantees. Among Pareto-efficient solutions, certain configurations achieve minimal carbon footprint, high coverage, and controlled costs simultaneously, illustrating that well-designed strategies can balance economic, environmental, and public-health objectives. Solutions focusing solely on cost or another single objective either increase emissions or backlog, highlighting the trade-offs inherent in multi-objective planning.

Strategies with lighter cold-chain requirements can substantially reduce emissions without compromising protection, helping identify where genuine trade-offs begin. Procurement remains the primary cost driver, while behavior-aware constraints ensure realistic vaccine mixes and stable flows across the supply network.

Beyond these case-specific insights, the contribution is methodological and transferable. The framework combines a solver-neutral formulation, a reproducible data pipeline, and a rigorous ε-constraint protocol, enabling transparent, policy-relevant decisions. Decision-makers can compare cost, carbon, and coverage on the same scale, making trade-offs explicit and defensible. Future research could enhance realism and impact by :

- Integrating bilevel optimization to explicitly model policy interventions.
- Modeling vaccine loss and perishability along the supply chain.

- Incorporating finer-grained emission accounting across transport, storage, and packaging.
- Applying geospatial optimization for facility siting to improve distribution efficiency.

Overall, this framework provides a robust, adaptable decision-support tool, capable of guiding vaccine supply strategies today and evolving with new data and public-health priorities. Importantly, it demonstrates that multi-objective planning can identify strategies that achieve high coverage, reduce environmental impact, and control costs simultaneously, offering actionable insights for large-scale vaccine logistics.

REFERENCES

1. Shah N. Pharmaceutical supply chains: key issues and strategies for optimisation. Computers & Chemical Engineering 28:929-941 (2004). https://doi.org/10.1016/j.compchemeng.2003.09.022
2. Masoumi AH, Yu M, Nagurney A. A supply chain generalized network oligopoly model for pharmaceuticals under brand differentiation and perishability. Transportation Research Part E: Logistics and Transportation Review 48:762-780 (2012). https://doi.org/10.1016/j.tre.2012.01.001
3. Papageorgiou LG. Supply chain optimisation for the process industries: advances and opportunities. Computers & Chemical Engineering 33:1931-1938 (2009). https://doi.org/10.1016/j.compchemeng.2009.06.014
4. Zavvar Sabegh MH, Mohammadi M, Naderi B. Multi-objective optimization considering quality concepts in a green healthcare supply chain for natural disaster response: neural network approaches. Int J Syst Assur Eng Manag 8:1689-1703 (2017). https://doi.org/10.1007/s13198-017-0645-1
5. Hirsch P. Book reviews : competitive communciation: a rhetoric for modern business barry eckhouse. new york: oxford university press, 1999. 288 pages. corporate communication, second edition. paul A. argenti. boston: irwin mcgraw-hill, 1998.258 pages. Business Communication Quarterly 63:111-113 (2000). https://doi.org/10.1177/108056990006300112
6. Bekta? T, Laporte G. The pollution-routing problem. Transportation Research Part B: Methodological 45:1232-1250 (2011). https://doi.org/10.1016/j.trb.2011.02.004
7. Martell L. A review of: “design and analysis of vaccine studies, by M. E. halloran, I. M. longini, jr., and C. J. struchiner,”. Journal of Biopharmaceutical Statistics 20:1077-1079 (2010). https://doi.org/10.1080/10543406.2010.495703
8. Mavrotas G. Effective implementation of the ?-constraint method in multi-objective mathematical programming problems. Applied Mathematics and Computation 213:455-465 (2009). https://doi.org/10.1016/j.amc.2009.03.037
9. Lazarus JV, Ratzan SC, Palayew A, Gostin LO, Larson HJ, Rabin K, Kimball S, El-Mohandes A. A global survey of potential acceptance of a COVID-19 vaccine. Nat Med 27:225-228 (2020). https://doi.org/10.1038/s41591-020-1124-9
10. Etalab. Comment les administrations ont collaboré à l’ouverture des données du coronavirus : le cas français, 2020. URL https://www.etalab.gouv.fr/comment-les-administrations-ont-collabore-a-louverture-des-donnees-du-coronavirus-le-cas-francais/. Consultation: April 6, 2025.
11. Khodaee V, Kayvanfar V, Haji A. A humanitarian cold supply chain distribution model with equity consideration: the case of COVID-19 vaccine distribution in the european union. Decision Analytics Journal 4:100126 (2022). https://doi.org/10.1016/j.dajour.2022.100126
12. INSEE. Dossier complet — france métropolitaine. https://www.insee.fr/fr/statistiques/2011101?geo=METRO-1, 2025. Consultation: May 22, 2025.
13. Marler RT, Arora JS. Survey of multi-objective optimization methods for engineering. Structural and Multidisciplinary Optimization 26:369-395 (2004). https://doi.org/10.1007/s00158-003-0368-6
14. Plateforme ouverte des données publiques françaises. data.gouv.fr - plateforme ouverte des données publiques françaises. https://www.data.gouv.fr/fr/, 2024. Consulté le 30 mai 2025.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Multiperiod optimisation of a European CCS supply chain under capture-cost uncertainty.

José A. Álvarez-Menchero*, Rubén Ruiz-Femenia, Raquel Salcedo-Díaz, José A. Caballero

Department of Chemical Engineering, University of Alicante. Ap. Correos 99, E-03080, Alicante. Spain.
* Corresponding Author: joseantonio.alvarez@ua.es

ABSTRACT

This paper presents a Europe-wide optimisation framework for designing and operating a multi-period Carbon Capture and Storage (CCS) supply chain across Europe. A MATLAB preprocessing pipeline constructs an auditable techno-economic dataset (emission nodes, ports, aquifers, candidate pipeline/shipping arcs and costs) and exports it to a GAMS optimisation model. The planning problem is formulated as a two-stage stochastic MILP, where scenario-independent first-stage decisions select discrete pipeline and shipping capacity bands and port operating modes, while scenario-dependent second-stage decisions allocate capture, transport and sequestration flows. Uncertainty is represented through correlated scenarios of capture unit costs for four capture technologies ($CV = 0.35, \rho = 0.8, N_s = 20$). To address the computational burden induced by inter-temporal binary investments and scenario replication, we apply a two-phase arc-screening heuristic: an LP relaxation on the full network identifies promising corridors, and a reduced MILP is solved on the resulting candidate arc sets. The resulting stochastic model contains 1, 278, 661 continuous variables and 48, 860 binaries and is solved in 1, 347.9 s with a 4.85% optimality gap.

Keywords: Optimization, Supply Chain, Carbon Dioxide Capture, GAMS, Design Under Uncertainty

1. INTRODUCTION

As a consequence of the Paris Agreement reached at the 2015 United Nations Climate Change Conference, 196 Parties pledged to reduce their emissions and to collaborate to limit global temperature rise well below 2 °C above pre-industrial levels, aiming to stay below 1.5 °C [1]. Nonetheless, the climate-change mitigation scenarios presented by the Intergovernmental Panel on Climate Change suggest that even rapid decarbonisation may not be enough to achieve this stringent target [2]. A promising strategy to address both climate and economic challeng-es is the implementation of Carbon Capture and Storage (CCS) [3–5]. CCS is widely recognised as an indispensa-ble technology for limiting global warming below 1.5 °C [6, 7]. Consequently, a significant body of research has focused on designing CCS supply chains in recent years. For instance, [8] proposes a temporal decomposition scheme for CCS supply-chain design.

In this work, a multi-period mixed-integer linear programming (MILP) model for a CCS supply chain across Europe is formulated and solved in full space (i.e., without decomposition) to minimise overall cost under uncertainty. The model is grid-based and spatially explicit, simultaneously optimising capture, multi-mode transport, and geological sequestration nodes over a planning horizon. Transport technologies include onshore pipelines, offshore pipelines, and maritime shipping. Emissions from power, iron and steel and cement industries are projected from the EDGAR database [9]; potential storage comes from JRC/CO2Stop [10]. In the figures 1-4 are shown the regional discretisation used for Europe (Figure 1), the location of the European ports (Figure 2), the place of the aquifers (Figure 3) and the large stationary emission sources considered (Figure 4).

Uncertainty is introduced explicitly in unit capture costs for four technologies—post-combustion (coal), post-combustion (gas), oxy-combustion (coal), and pre-combustion. These four uncertain parameters are correlated, reflecting shared cost drivers; for example, a decrease in coal post-combustion cost is expected to co-move with a decrease in gas post-combustion cost.

Compared with Álvarez-Menchero et al. [8], this study (i) expands the spatial scope to continental Europe,

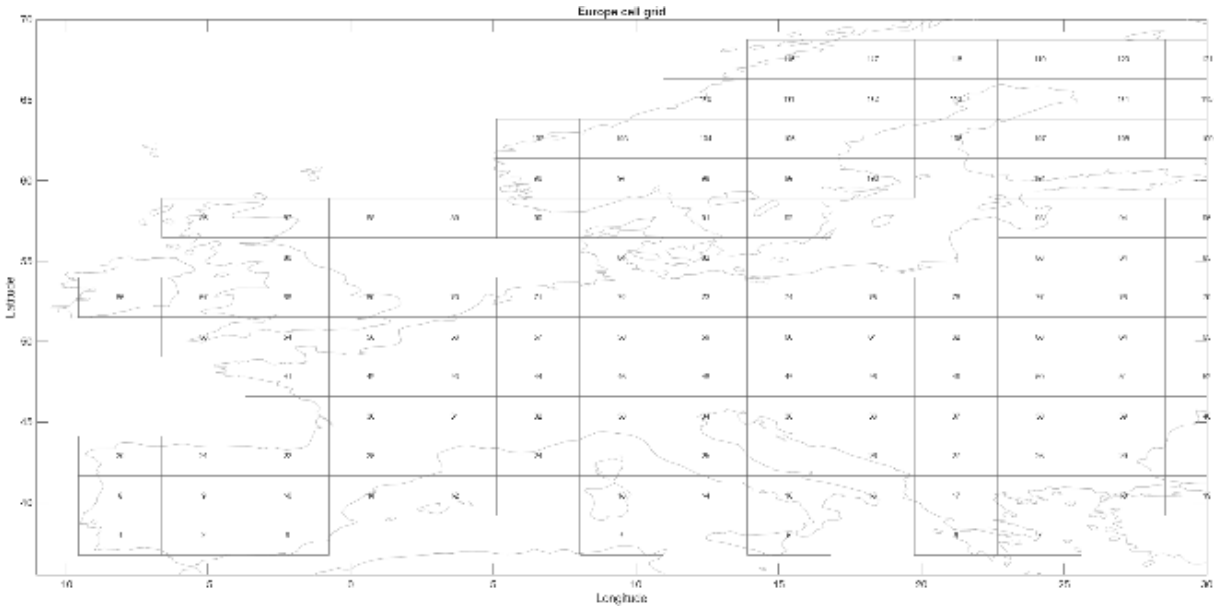

Figure 1. Regional discretization

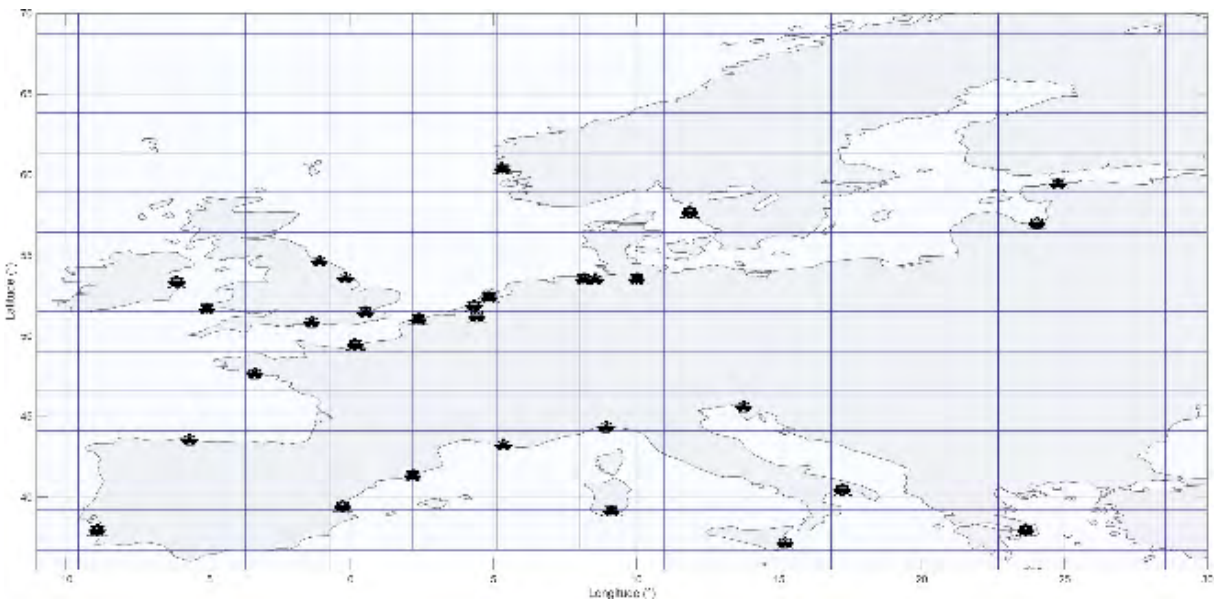

Figure 2. Port locations.

(ii) adds ship transport alongside onshore/offshore pipelines, and (iii) considers correlated uncertainty in the costs of the capture technologies.

2. PROBLEM STATEMENT

2.1 Model formulation

The proposed optimisation model can be interpreted as the interaction of three tightly coupled subproblems:

- a **capture subproblem**, which allocates emissions to capture technologies and enforces a time-dependent capture target;
- a **transport subproblem**, which designs a discrete infrastructure (pipelines and shipping arcs) and routes captured CO_2;
- a **sequestration subproblem**, which injects CO_2 into aquifers under rate and capacity limits.

The overall model is formulated as a two-stage stochastic MILP. The first-stage variables represent *strategic* infrastructure decisions—pipeline and shipping capacity band activations and port operating modes—and are therefore shared across all scenarios. The second-

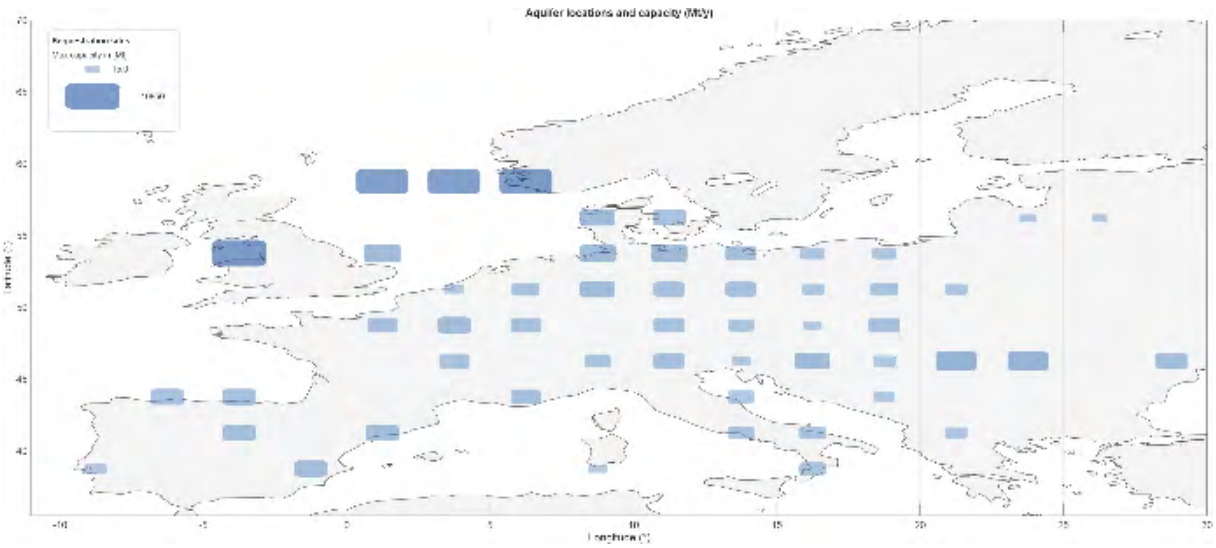

Figure 3. Aquifer locations.

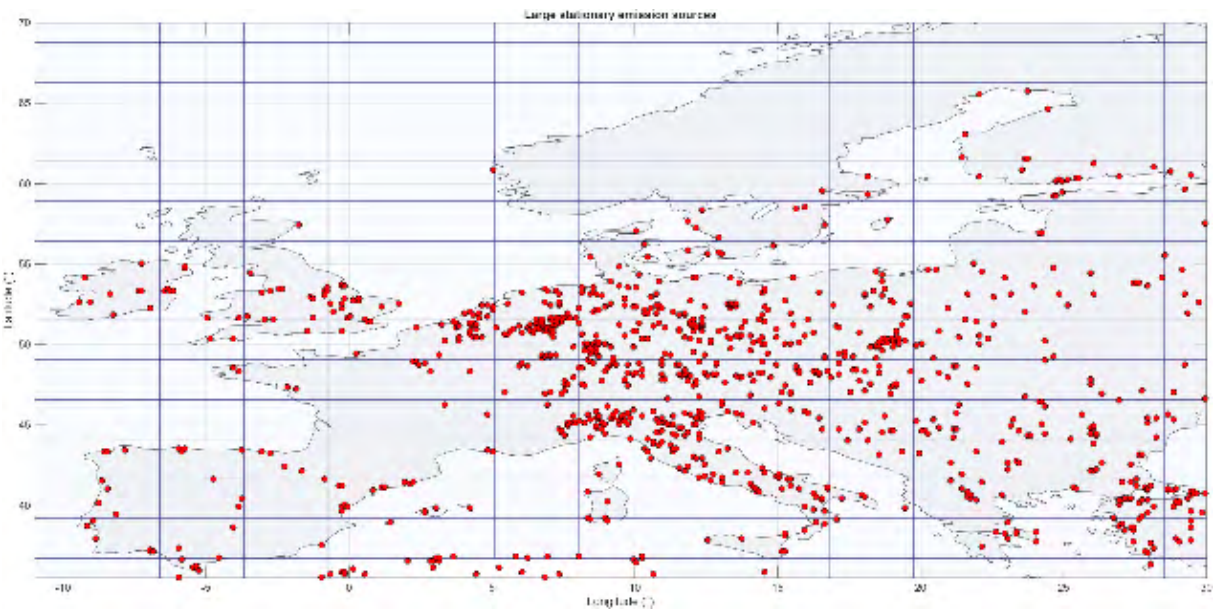

Figure 4. Large stationary emission sources.

stage variables are *operational* and scenario-dependent—capture levels, transported flows, and sequestration—so that the system can adapt to each realisation of uncertain capture costs. The three subproblems are coupled through a node-wise mass balance and an expected-cost objective.

Because the formulation combines (i) large-scale discrete network design and (ii) scenario replication of operational decisions, the full stochastic MILP can become computationally demanding. This motivates the two-phase solution approach presented later (LP screening followed by a reduced MILP).

2.2.1 Capture subproblem

At each emitting node n and year t, the model decides how much CO_2 is processed by each capture technology c, denoted $pCO2_{c,n,t,s}$, and the corresponding captured amount $cCO2_{c,n,t,s}$. Processing is limited by the local emission availability, and by technology-specific feasibility rules (e.g., fuel-share-based envelopes for the power sector).

First, processed CO_2 cannot exceed total emissions:

$$\sum_{c} pCO2_{c,n,t,s} \leq E_{n,t}^{total} \quad \forall n, t, s \tag{1}$$

Second, technology eligibility is enforced through an admissible envelope $E_{c,n,t}^{allow}$:

$$pCO2_{c,n,t,s} \leq E^{allow}_{c,n,t} \ \forall c, n, t, s \tag{2}$$

Captured CO_2 is obtained from processed CO_2 via the technology efficiency ce_c:

$$cCO2_{c,n,t,s} = ce_c pCO2_{c,n,t,s} \ \forall c, n, t, s \tag{3}$$

Finally, a year-dependent capture target $\alpha(t)$ imposes a minimum captured fraction of total emissions:

$$\sum_n \sum_c cCO2_{c,n,t,s} \geq \alpha_t \sum_n E^{total}_{n,t} \ \forall t, s \tag{4}$$

In addition to these core constraints, pre-combustion capture is modelled as a niche option through deployment caps. A node-level cap limits the local pre-combustion contribution relative to power emissions:

$$cCO2_{pre^{comb},n,t,s} \leq preCapFrac^{node} E^{power}_{n,t} : \forall n, t, s \tag{5}$$

and an EU-wide cap limits aggregate pre-combustion deployment:

$$\sum_n cCO2_{pre^{comb},n,t,s} \leq PreCapTotal_t \ \forall t, s \tag{6}$$

2.2.2 Transport subproblem

Captured CO_2 can be transported through (i) pipelines on candidate arcs $A_{n,nn}$ and (ii) maritime shipping on candidate port-to-port arcs $SArc_{n,nn}$. Transport capacity is discretised into throughput bands q with bounds $\left[Q^{min}_q, Q^{max}_q\right]$. The activation of a band is controlled by binary variables $y^{pipe}_{q,n,nn,t}$ and $y^{ship}_{q,n,nn,t}$, which are first-stage (scenario-independent) decisions. Operational flows—$qCO2^{pipe}_{q,n,nn,t,s}$ and $qCO2^{ship}_{q,n,nn,t,s}$—are second-stage (scenario-dependent).

Pipeline capacity selection and linking

If a pipeline band is activated, the corresponding flow must remain within the band bounds:

$$qCO2^{pipe}_{q,n,nn,t,s} \geq Q^{min}_q y^{pipe}_{q,n,nn,t} \ \forall q, (n, nn) \in Aout, \forall t, s \tag{7}$$

$$qCO2^{pipe}_{q,n,nn,t,s} \leq Q^{max}_q y^{pipe}_{q,n,nn,t} \ \forall q, (n, nn) \in Aout, \forall t, s \tag{8}$$

where $Aout \subseteq A$ excludes arcs whose origin is an aquifer node (aquifers are sinks in the network).

At most one capacity band can be chosen per arc-year:

$$\sum_q y^{pipe}_{q,n,nn,t} \leq 1 \ \forall (n, nn) \in Aout, t \tag{9}$$

To avoid redundant designs, anti-parallel construction is forbidden:

$$\sum_q y^{pipe}_{q,n,nn,t} + \sum_q y^{pipe}_{q,nn,n,t} \leq 1 \ \forall (n, nn) \in A, \ n < nn, \forall t \tag{10}$$

Finally, infrastructure is persistent over time: once activated, it cannot be deactivated in later years:

$$\sum_q y^{pipe}_{q,n,nn,t} \geq \sum_q y^{pipe}_{q,n,nn,t-1} \ \forall (n, nn) \in Aout, \forall t > 1 \tag{11}$$

Shipping capacity selection and linking

Shipping arcs follow the same band-activation logic:

$$qCO2^{ship}_{q,n,nn,t,s} \geq Q^{min}_q y^{ship}_{q,n,nn,t} \ \forall q, (n, nn) \in SArc, \forall t, s \tag{12}$$

$$qCO2^{ship}_{q,n,nn,t,s} \leq Q^{max}_q y^{ship}_{q,n,nn,t} \ \forall q, (n, nn) \in SArc, \forall t, s \tag{13}$$

with one band per arc-year,

$$\sum_q y^{ship}_{q,n,nn,t} \leq 1 \ \forall (n, nn) \in SArc, \forall t \tag{14}$$

anti-parallel exclusion,

$$\sum_q y^{ship}_{q,n,nn,t} + \sum_q y^{ship}_{q,nn,n,t} \leq 1 \ \forall (n, nn) \in SArc, \ n < nn, \forall t \tag{15}$$

and persistence,

$$\sum_q y^{ship}_{q,n,nn,t} \geq \sum_q y^{ship}_{q,n,nn,t-1} \ \forall (n, nn) \in SArc, \forall t > 1 \tag{16}$$

2.2.3 Sequestration subproblem, coupling constraints, and objective function

Sequestration takes place at aquifer nodes. The key coupling between capture, transport and sequestration is the node-wise mass balance, enforced for every node, year, and scenario.

$$\begin{aligned} &\sum_{q,nn:(nn,n)\in A} qCO2^{pipe}_{q,nn,n,t,s} + \sum_{q,nn:(nn,n)\in SArc} qCO2^{ship}_{q,nn,n,t,s} \\ &\qquad + \sum_c cCO2_{c,n,t,s} \\ &= \sum_{q,nn:(n,nn)\in A} qCO2^{pipe}_{q,n,nn,t,s} \\ &\qquad + \sum_{q,nn:(n,nn)\in SArc} qCO2^{ship}_{q,n,nn,t,s} \\ &\qquad + Seq_{n,t,s} \ \forall n, t, s \end{aligned} \tag{17}$$

Sequestration is further limited by an annual injection cap:

$$Seq_{n,t,s} \leq SeqRate \ \forall n, t, s \tag{18}$$

and by finite site capacity over the full horizon:

$$\sum_t Seq_{n,t,s} \leq S_n \ \forall n, \forall s \tag{19}$$

Total cost in each year and scenario is the sum of capture, transport, and storage costs. Capture costs are scenario-dependent via $UCC_{c,s}$:

$$TCC(t, s) = \sum_{n,c} UCC_{c,s} \ cCO2_{c,n,t,s} \ \forall t, s \tag{20}$$

Pipeline transport costs are proportional to distance and include an offshore multiplier $aco_{n,nn}$:

$$TTC_{t,s}^{pipe} = \sum_{q} \sum_{(n,nn)\in Aout} qCO2_{q,n,nn,t,s}^{pipe} \; d_{n,nn} UTC_q aco_{n,nn} \; \forall t, s \tag{21}$$

Shipping costs are calculated in the following way:

$$TTC_{t,s}^{ship} = \sum_{q} \sum_{(n,nn)\in SArc} qCO2_{q,n,nn,t,s}^{ship} \; (UTC^{ship,slope} d_{n,nn} + UTC_q^{ship,int}) \; \forall t, s \tag{22}$$

Storage costs are linear in injected CO_2, scaled by an offshore storage multiplier:

$$TSC_{t,s} = \sum_{n} Seq_{n,t,s} USC \, USC_n^{off} \; \forall t, s \tag{23}$$

Let $TC_{t,s} = TCC_{t,s} + TTC_{t,s}^{pipe} + TTC_{t,s}^{ship} + TSC_{t,s}$. The objective minimises the expected total cost across the planning horizon using scenario probabilities π_s:

$$\min \; Z = \sum_{s} \pi_s \sum_{t} TC_{t,s} \tag{24}$$

NOMENCLATURE

Sets

t: years {1, 2, ..., 10}

n: nodes {1, 2, ..., 203}

c: capture technology $\{post_{coal}^{comb}, post_{coal}^{comb}, pre^{comb}, oxy^{fuel}\}$

$\mathcal{Q} = \{q1, q2, q3, q4\}$:set of discrete throughput bands, $|\mathcal{Q}| = 4$
q: throughput bands $\{q_1, q_2, q_3, q_4\}$

$\mathcal{S}$:set of scenarios, $|\mathcal{S}| = 20$

s: scenarios $\{s_1, s_2, \ldots, s_{20}\}$

A: set of feasible pipeline arcs

$SArc$: set of feasible shipping arcs (port-to-port)

$pred$: predecessor relation ($(t, tt) \in pred \Rightarrow tt$ precedes t)

PARAMETERS

$coord_{n,p}$:coordinates of node n (degrees)

$d_{n,nn}$:distance between nodes n and nn (km)
$aco_{n,nn}$:offshore pipeline cost multiplier on arc (n, nn) (dimensionless)

$E_{n,t}^{total}$:total emissions [Mt/y]

$E_{n,t}^{power}$:power-sector emissions [Mt/y]

$E_{n,t}^{cement}$:cement-sector emissions [Mt/y]

$E_{n,t}^{steel}$:iron&steel emissions [Mt/y]

α_t:required capture fraction in year t (target on captured CO2

ce_c:capture efficiency of technology c

S_n:total storage capacity at aquifer node n over the horizon [Mt

$SeqRate$:maximum annual injection rate (scalar) [Mt/y]
USC:unit storage (sequestration) cost [€/Mt]

USC_n^{off}:offshore storage cost multiplier at aquifer n (dimensionless
Q_q^{min}:minimum throughput in band q [Mt/y]

Q_q^{max}:maximum throughput in band q [Mt/y]

UTC_q:pipeline unit transport cost in band q [€/Mt·km]
$UTC^{ship,slope}$:shipping distance slope [€/Mt·km]

$UTC_q^{ship,int}$:shipping intercept in band q [€/Mt]

$UCC_{c,s}$:scenario capture unit cost [€/Mt]

π_s: scenario probability, $\sum_{s\in\mathcal{S}} \pi_s = 1$

sh_n^{coal}:coal share of power emissions at node n (dimensionless sh_n^{gas}:gas share of power emissions at node n (dimensionless
$E_{c,n,t}^{allow}$:eligible-emissions envelope for tech c [Mt/y]

$preCapFrac^{node}$:node-level cap fraction for pre-combustion (dimensionless
$preCapFrac^{EU}$:EU-wide cap fraction for pre-combustion (dimensionless)

$PreCapTotal_t$:EU-wide pre-combustion cap in year t [Mt/y]

Continuous decision variables (scenario - dependent)

$pCO2_{c,n,t,s}$: processed CO2 [Mt/y]

$cCO2_{c,n,t,s}$: captured CO2 [Mt/y]

$qCO2_{q,n,nn,t,s}^{pipe}$: pipeline flow on arc $(n \to nn)$ [Mt/y]

$qCO2_{q,n,nn,t,s}^{ship}$: shipping flow on arc $(n \to nn)$ [Mt/y]

$Seq_{n,t,s}$: sequestered CO2 at aquifer node n [Mt/y]

$TCC_{t,s}$: total capture cost [€/y]

$TTC_{t,s}^{pipe}$: pipeline transport cost [€/y]

$TTC_{t,s}^{ship}$: shipping transport cost [€/y]

$TTC_{t,s}$: total transport cost (=pipeline+ship) [€/y]

$TSC_{t,s}$: total storage cost [€/y]

$TC_{t,s}$: total system cost [€/y]

Z: expected objective value [€], $Z = \sum_{s\in\mathcal{S}} \pi_s \sum_{t\in\mathcal{T}} TC_{t,s}$

Binary decision variables (scenario - independent)

$y^{pipe}_{q,n,nn,t} \in \{0,1\}$:1 if pipeline band q is activated on arc $(n \rightarrow nn)$ in year t

$y^{ship}_{q,n,nn,t} \in \{0,1\}$:1 if shipping band q is activated on arc $(n \rightarrow nn)$ in year t

$z^{out}_{n,t} \in \{0,1\}$:port mode (1=ship-export, 0=ship-import)

3. UNCERTAINTY MODELLING.

Uncertainty is introduced in the capture unit costs through a scenario-based representation. Let UCC_c^{mean} denote the nominal capture cost for technology c. We generate a set of scenarios $scen$in MATLAB by sampling correlated cost multipliers with coefficient of variation $CV = 0.35$ and common correlation $\rho = 0.8$ across the four capture technologies (post-combustion coal, post-combustion gas, oxy-combustion coal, and pre-combustion), reflecting shared cost drivers (e.g., energy penalty, solvent/oxygen supply costs, and scale effects). Scenario-specific capture costs are then computed as

$$UCC_{c,s} = UCC_c^{mean} \xi_{c,s} \tag{25}$$

where $\xi_{c,s}$ is the sampled multiplier, and scenario probabilities are taken as uniform. The stochastic MILP minimises the expected total cost, with first-stage infrastructure decisions shared across all scenarios and second-stage operational decisions adapting to each scenario realisation.

4. SOLUTION APPROACH.

As has been mentioned previously, the stochastic CCS network design problem introduced in Section 2 results in a large-scale mixed-integer linear programme (MILP). To keep the model tractable, we adopt a two-phase solution heuristic.

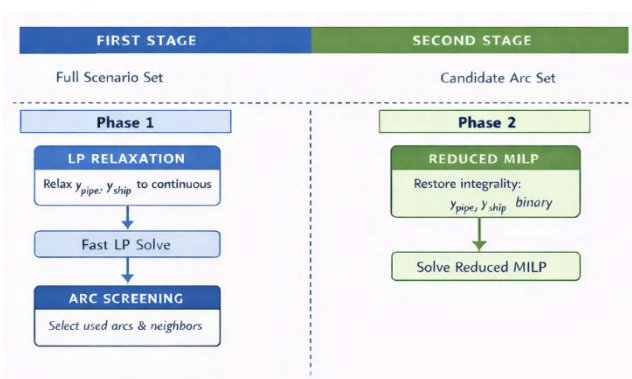

Figure 5. Two-phase heuristic diagram

4.1 Phase 1: LP relaxation on the full scenario set + arc screening

In Phase 1, we solve a *relaxation* of the full stochastic MILP in binaries $y^{pipe}_{q,n,nn,t}$ and $y^{ship}_{q,n,nn,t}$ are relaxed. This converts the original MILP into a large but comparatively fast linear programme (LP) over the full scenario set and the complete set of candidate arcs.

The objective of Phase 1 is to get structural information about where the optimal solution tends to route CO_2. The LP solution provides two key outputs:

- Arcs that carry non-negligible CO_2 flow in at least one time period and scenario are flagged as *used*.
- A conservative candidate set construction. To avoid overly aggressive pruning, we expand the used-arc set by adding a controlled neighbourhood around the endpoints of used arcs (i.e., we keep used arcs and arcs departing from nodes that appear as endpoints of used arcs). This step is critical in practice: it preserves alternative routings that the LP may not select strongly, but that can become essential once integrality is enforced.

The output of Phase 1 is therefore a Candidate Arc Set (for pipelines and for shipping), together with the Phase-1 operational solution (e.g., captured CO_2 levels) that can later be used to stabilise Phase 2.

4.2 Phase 2: reduced stochastic MILP on the candidate arc set

In Phase 2, integrality is restored and the model is solved as a reduced stochastic MILP defined on the Candidate Arc Set identified in Phase 1 (Figure 5).

Phase 2 returns a fully implementable solution: discrete infrastructure investments (pipelines, shipping capacity bands, and port mode), and consistent scenario-dependent operational decisions (capture, transport flows, and sequestration) that meet the decarbonisation targets.

4.3 Validation of the solution method

The proposed solution method was validated on reduced-size instances for which the integrated full-space stochastic MILP could still be solved directly. These simplified instances retained the main structural features of the original problem, including multi-period decisions, uncertainty in capture costs, and the coupling between capture, transport, and sequestration, while reducing the number of nodes, candidate arcs, time periods, and/or scenarios to keep the integrated formulation computationally manageable.

The comparison focused on the consistency of the objective value and on the main first-stage investment decisions, namely the activation of pipelines, shipping

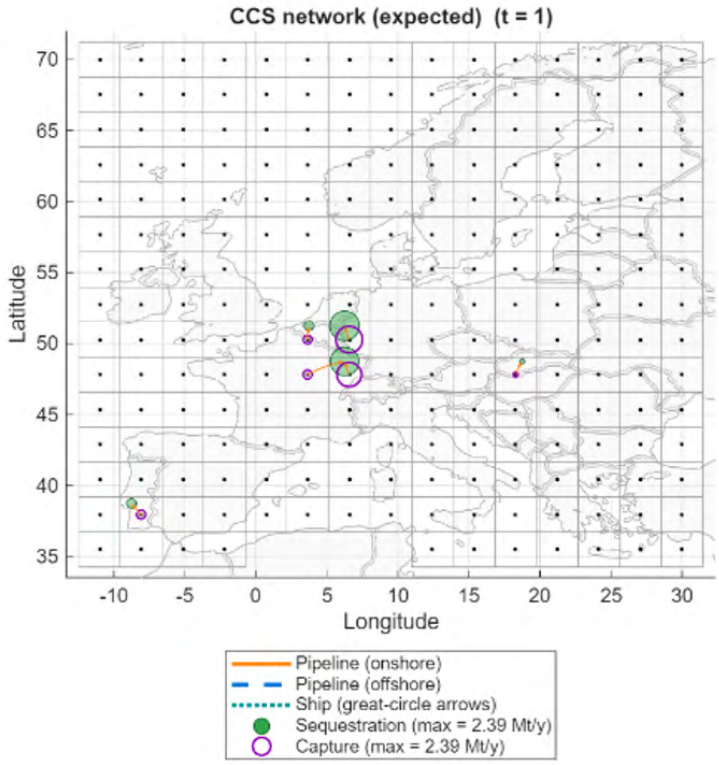

Figure 6. CCS network in year 1

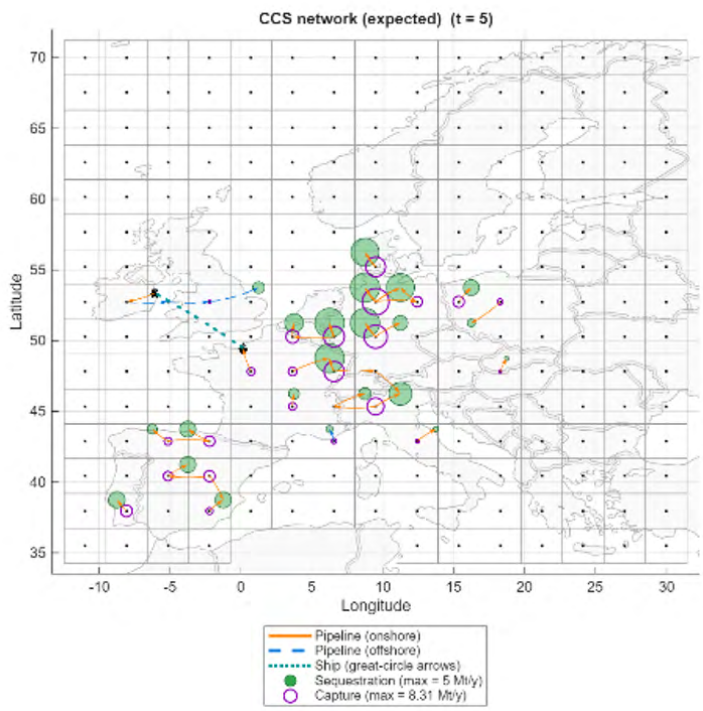

Figure 7. CCS network in year 5

arcs, and capacity levels. The reduced-instance tests showed that the proposed two-phase approach reproduced the same qualitative infrastructure patterns as the integrated model and provided solutions of comparable quality. This confirms that the candidate-arc screening step preserves the key design information required by

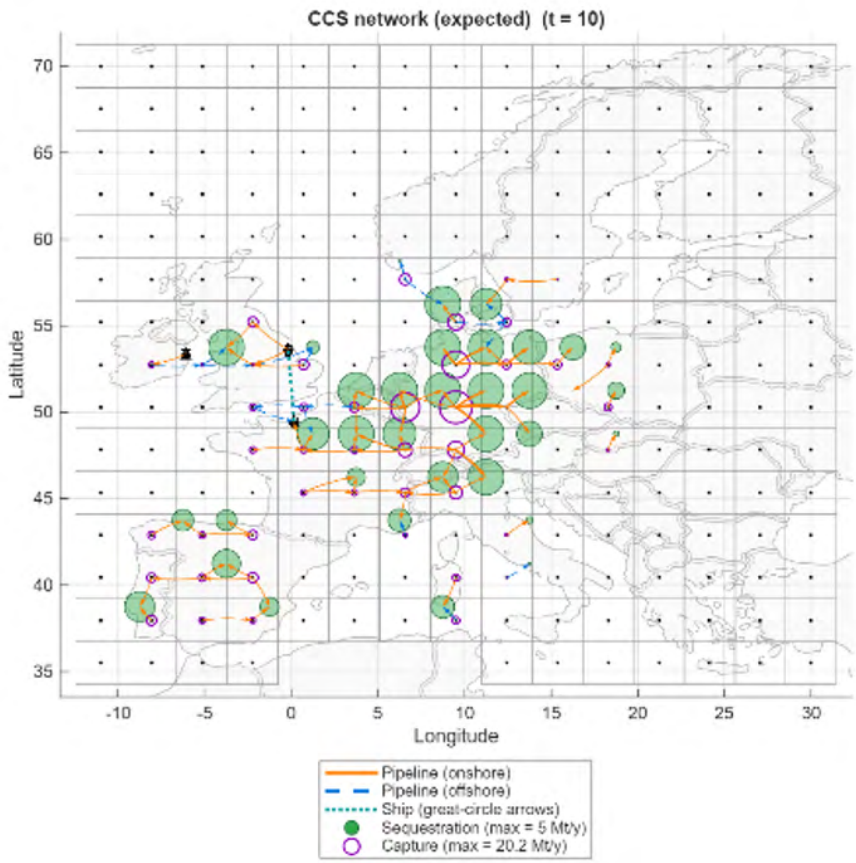

Figure 8. CCS network in year 10

the full-space formulation for the instances examined.

These validation runs were used during model development to verify the soundness of the proposed solution strategy before applying it to the large European-scale cases considered in this work. For such large stochastic instances, the direct solution of the integrated model becomes computationally very demanding. Therefore, the proposed method should be understood as a tractable solution strategy that preserves the essential structure and strategic decisions of the integrated formulation while enabling the treatment of realistically sized case studies.

Table 1. Model size and computational performance

No. of continuous variables	No. of binary variables	Solution time (s)	Optimality gap (%)
1, 278, 661	48, 860	1, 347.9	4.85

5. RESULTS

This section summarises the key outputs of the optimisation. To visualise the spatial evolution of the infrastructure, Figures 6–8 depict the designed CCS network in years 1, 5, and 10, respectively, highlighting how the selected pipeline/shipping corridors and utilised storage sites develop over time. Figure 9 presents the annual captured CO_2 by technology over the 10-year horizon. Figure 10 reports the total annual system cost together with its decomposition into capture, transport, and sequestration contributions. Finally, Table 1 reports the model size and computational performance.

6. CONCLUSION

We developed an integrated MATLAB–GAMS framework for Europe-wide CCS network design that combines multi-technology capture, multimodal transport (pipelines and shipping), and geological storage in a multi-period setting. Capture-cost uncertainty is modelled via correlated scenarios and embedded in a two-stage stochastic MILP, enabling here-and-now infrastructure decisions. To make the large stochastic design problem tractable, a two-phase LP-based arc-screening strategy reduces the candidate transport network before solving the reduced MILP. The approach delivers a feasible time-

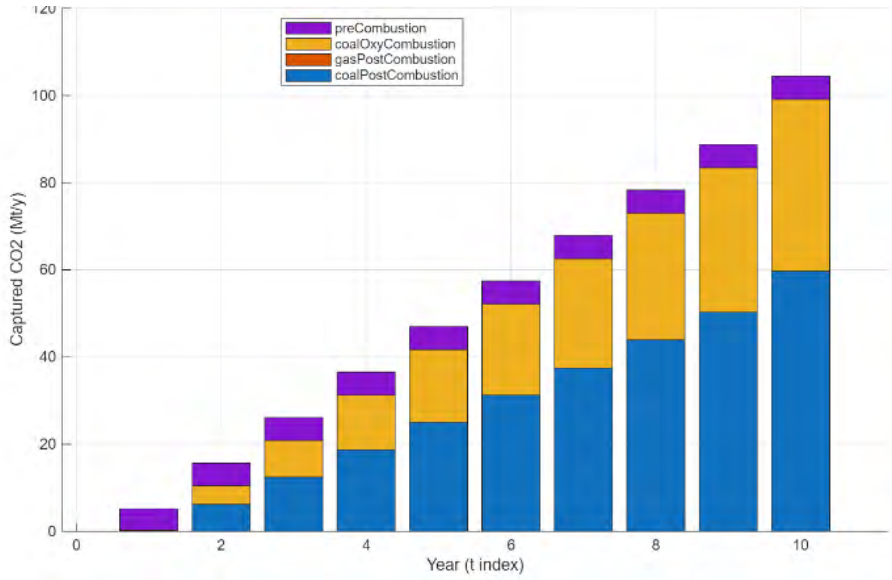

Figure 9. Annual captured CO_2

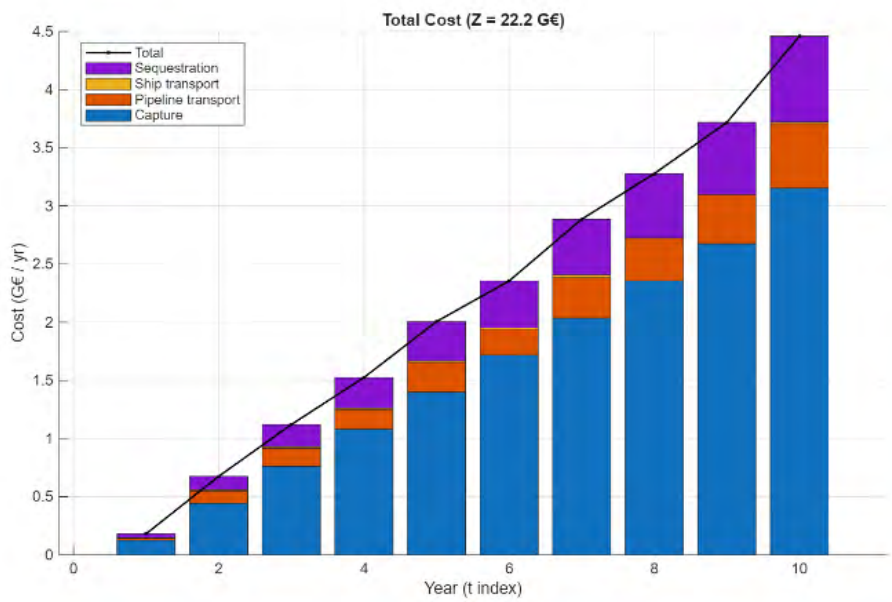

Figure 10. Total annual cost.

evolving infrastructure plan together with technology and cost insights, while achieving practical runtimes for a problem of this scale. Future work will focus on tightening optimality guarantees and extending uncertainty to additional parameters (e.g., transport and storage costs, storage capacity).

ACKNOWLEDGEMENTS

The authors gratefully acknowledge financial support to the Spanish “Ministerio de Ciencia, Innovación y Univer sidades” under project PID2024-155232NB-I00. The authors would also like to thank “Generalitat Valenciana: Conselleria d’Educació, Universitats y Educació” and the European Social Fund Plus (ESF+) for the PhD grant (CIACIF/2023/117).

REFERENCES

1. COP-21. In: Paris Agreement, United Nations Framework Convention on Climate, Conference of the Parties 21. Paris, France 2015. https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement.

2. IPCC. Special Report on the Impacts of Global Warming of 1.5 °C Above Pre-Industrial Levels and Related Global Greenhouse Gas Emission Pathways. https://www.ipcc.ch/sr15/ (Intergovernmental Panel on Climate Change, 2018).
3. IPCC Special Report on Carbon Dioxide Capture and Storage (eds Metz, B., Davidson, O., de Coninck, H. C., Loos, M. & Meyer, L. A.) Prepared by Working Group III of the Intergovernmental Panel on Climate Change (IPCC) (Cambridge University Press, Cambridge, UK, 2005).
4. Boot-Handford ME, Abanades JC, Anthony EJ, Blunt MJ, Brandani S, Mac Dowell N, Fernández JR, Ferrari MC, Gross R, Hallett JP, Haszeldine RS, Heptonstall P, Lyngfelt A, Makuch Z, Mangano E, Porter RTJ, Pourkashanian M, Rochelle GT, Shah N, Yao JG, Fennell PS. Carbon capture and storage update. Energy Environ. Sci. 7:130-189 (2014). https://doi.org/10.1039/c3ee42350f
5. Bui M, Adjiman CS, Bardow A, Anthony EJ, Boston A, Brown S, Fennell PS, Fuss S, Galindo A, Hackett LA, Hallett JP, Herzog HJ, Jackson G, Kemper J, Krevor S, Maitland GC, Matuszewski M, Metcalfe IS, Petit C, Puxty G, Reimer J, Reiner DM, Rubin ES, Scott SA, Shah N, Smit B, Trusler JPM, Webley P, Wilcox J, Mac Dowell N. Carbon capture and storage (CCS): the way forward. Energy Environ. Sci. 11:1062-1176 (2018). https://doi.org/10.1039/c7ee02342a
6. Luderer G, Vrontisi Z, Bertram C, Edelenbosch OY, Pietzcker RC, Rogelj J, De Boer HS, Drouet L, Emmerling J, Fricko O, Fujimori S, Havlík P, Iyer G, Keramidas K, Kitous A, Pehl M, Krey V, Riahi K, Saveyn B, Tavoni M, Van Vuuren DP, Kriegler E. Residual fossil CO2 emissions in 1.5–2 °C pathways. Nature Clim Change 8:626-633 (2018). https://doi.org/10.1038/s41558-018-0198-6
7. Boitier B, Nikas A, Gambhir A, Koasidis K, Elia A, Al-Dabbas K, Aliba? ?, Campagnolo L, Chiodi A, Delpiazzo E, Doukas H, Fougeyrollas A, Gargiulo M, Le Mouël P, Neuner F, Perdana S, van de Ven DJ, Vielle M, Zagamé P, Mittal S. A multi-model analysis of the eu's path to net zero. Joule 7:2760-2782 (2023). https://doi.org/10.1016/j.joule.2023.11.002
8. ?lvarez-Menchero JA, Ruiz-Femenia R, Salcedo-Diaz R, Moreno-Palancas IF, Caballero JA. Temporal decomposition scheme for designing large-scale CO2 supply chains using a neural network-based model for forecasting CO2 emissions. Systems and Control Transactions 4:918-923 (2025). https://doi.org/10.69997/sct.193502
9. Joint Research Centre. Emission Database for Global Atmospheric Research (EDGAR). https://edgar.jrc.ec.europa.eu/index.php

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimal Operation of an Alkaline Electrolyzer in an Industrial Setting Using Effective Linearization Techniques

Jonaed Bin Mustafa Kamal[a]*, Loukas Kyriakidis[a], Saskia Bublitz[a], Bogdan Dorneanu[b], and Harvey Arellano-Garcia[b]

[a] German Aerospace Center, Institute of Low-Carbon Industrial Processes, Simulation and Virtual Design, Cottbus, Germany
[b] Brandenburg University of Technology, Chair of Process and Plant Technology, Cottbus, Germany
* Corresponding Author: jonaed.mustafakamal@dlr.de.

ABSTRACT

Renewable powered water electrolysis offers a promising strategy to decarbonize industrial sectors with high demand for hydrogen. Operational optimization of industrial electrolyzer systems is often formulated as mixed-integer linear programming (MILP) problems, where a constant hydrogen production to electrical power consumption ratio is assumed instead of the nonlinear relationship. Incorporating a nonlinear electrolyzer model into a linear optimal hydrogen dispatch framework remains a significant challenge. This study addresses this challenge by formulating the optimization problem in two ways. First, the model is solved as a nonlinear programming (NLP) problem by incorporating a nonlinear model of an alkaline electrolyzer (AEL) into the optimization framework. The binaries and integers are relaxed to continuous variables and associated penalty terms are added to the objective function to enforce integrality. The NLP is solved using the local nonlinear solver Interior Point OPTimizer (IPOPT) using a multi-start approach, where the solver is executed from random initial points and the minimum objective value is taken as the best guess for the global minimum. Second, a linearized model of the AEL is used in an MILP formulation. The univariate nonlinear terms within the model are approximated using three piecewise linearization techniques, i.e., convex combination, convex combination with SOS2 variables and a slope based big-M formulation. The bilinear terms are relaxed using the McCormick envelope. Results show that the MILP formulation achieves faster solution times, with varying accuracy depending on the linearization method, whereas the NLP approach more accurately captures the hydrogen production curve, albeit having longer convergence times.

Keywords: Alkaline electrolyzer, Renewable hydrogen system, Piecewise linearization, McCormick relaxation, Operational optimization

INTRODUCTION

Optimal dispatch problems for electrolysis require robust and quickly solvable process models to enable real-time control decisions. For this reason, mixed-integer linear programming (MILP) models are often employed [1-3]. Among the simplest models are those assuming a constant electrolyzer efficiency or a fixed hydrogen production to electrical power consumption ratio. Baumhof et al. [4] showed that such models exhibit significant deviations in the partial load range of the electrolyzer resulting in sub-optimal solutions. In fact, there exists a nonlinear relationship between power consumption and hydrogen production that can be modeled using detailed physical models or semi-empirical models of the electrolyzer [5]. Baumhof et al. [4] incorporated the widely adopted semi-empirical model of an alkaline electrolyzer (AEL) first proposed by Ulleberg [6] and later extended by Sánchez et al. [7] into their MILP framework by piecewise linearizing the power consumption vs hydrogen production curve obtained from a separate nonlinear AEL model. Furthermore, the study builds the MILP by introducing different operational states, i.e., on, off and stand-by, of the AEL.

Due to the inherent complex nature, incorporating a nonlinear component model into a linear optimization problem remains a challenge. On one hand nonlinear solvers like Interior Point OPTimizer (IPOPT) only works

with continuous variables [8] and thus cannot handle integers and binary variables. On the other hand linear solvers only work with linear objective function and linear constraints [9]. This study incorporates the nonlinear AEL model into a linear operational optimization problem by first solving the problem as a nonlinear program (NLP), where the integer and the binary variables are relaxed to continuous variables and associated penalty terms are added to the objective function to enforce integrality. Second, the problem is reformulated as an MILP by linearizing the univariate nonlinear terms within the AEL model using three different piecewise linearization techniques, and relaxing the multilinear terms using McCormick envelope. Furthermore, this paper extends the work of Baumhof et al. [4] by developing an AEL system model tailored to a multi-stack configuration. This work is part of the development of a digital twin for a renewable powered 40 MW AEL system designed to supply hydrogen to a refinery in Corinth, Greece within the framework of the EU project EPHYRA [10].

The paper has two primary objectives. First, it tries to identify the best linearization method when linearizing the nonlinear semi-empirical AEL model when applied to a hydrogen dispatch problem, evaluating performance in terms of accuracy and computational speed. Second, it compares the results of the MILP with the NLP to assess the merits of each approach and evaluates if the linearized model can accurately reproduce the variable trajectories when compared to the NLP formulation.

DESCRIPTION OF THE SYSTEM

The industrial system, illustrated in **Figure 1**, has a 40 MW pressurized AEL at its core that produces hydrogen for a refinery. The AEL is primarily powered by a photovoltaic (PV) plant, with supplementary power produced through waste heat recovery from the refinery via an Organic Rankine Cycle (ORC). When required, power can be drawn from the grid for meeting additional hydrogen demand. The AEL system is divided into 8 stacks with each operating at a pressure of 20-25 bar and consuming 5 MW power when operating at full capacity. The produced hydrogen is compressed and routed through a hydrogen storage system operating at 350 bar before being delivered to the refinery. The system operates in such a way that it meets the hydrogen demand of the refinery at each time step.

MODEL FORMULATION

The optimal hydrogen dispatch problem is first formulated as a mixed-integer nonlinear problem (MINLP) and solved as an NLP by relaxing the binary and integer variables to continuous variables and then introducing penalty terms within the objective function to encourage integrality. Then it is converted to an MILP by linearizing the nonlinear terms within the AEL model. The optimization framework is implemented in the Python based optimization framework Pyomo. The NLP is solved using the local nonlinear solver Interior Point OPTimizer (IPOPT), whereas the MILP formulation is solved using the linear solver Gurobi. The objective function and the constrains within the optimization problem are described in detail in the following paragraphs.

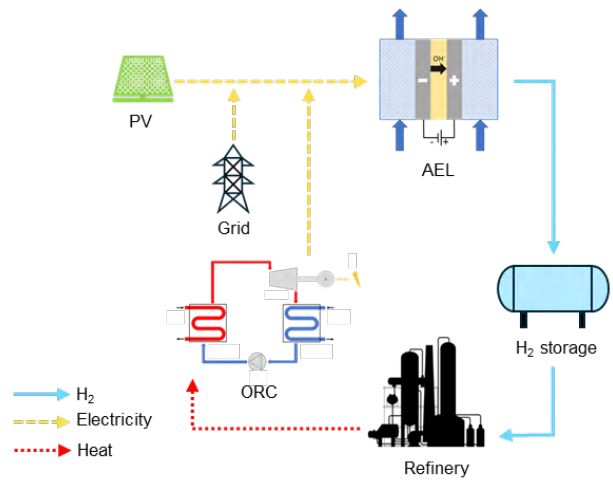

Figure 1. Process flow diagram of the industrial setup.

Objective function

The objective of the optimization problem is to minimize the total operating costs associated with the system over a period of n hours. The objective function is thus formulated as:

$$\min \sum_{t=1}^{n} \left[P_{\text{grid}}^{t} c_{\text{grid}}^{t} + \sum_{s=1}^{N_{\text{stack}}} z_{\text{su}}^{t,s} c_{\text{su}} \right] \tag{1}$$

where, P_{grid} is the power drawn from the grid, c_{grid} the cost to purchase one unit of grid electricity, z_{su} a binary variable denoting the start-up of one AEL stack, discussed in detail in the following section, c_{su} the cost associated with the start-up of a stack, n the total number of time steps, and N_{stack} the total number of stacks. The superscripts t and s are the time step and the AEL stack indices and are used throughout the model. The objective function in the NLP adds further penalty terms for enforcing integrality of the binary and the integer variables.

Operational States of the AEL Stacks

The operational states of the AEL stacks are modeled in a similar way to Baumhof et al. [4]. However, in this work, the model is further extended to be compatible with a multi-stack AEL system. Each stack has three operational states, i.e., on, off and standby. These operational states are described by the three binary variables $z_{\text{on}}, z_{\text{off}}$, and z_{sb}. At any time step, only one of the states can be active, which can be written as:

$$z_{\text{on}}^{t,s} + z_{\text{off}}^{t,s} + z_{\text{sb}}^{t,s} = 1 \quad (2)$$

The cold start-up of a stack is indicated by the binary variable z_{su}. The variable z_{su} is 0 at the start and becomes 1 whenever transitioning from off to on state between two consecutive time steps. The constraints are thus written as:

$$z_{\text{su}}^{t=1,s} = 0 \quad (3)$$

$$z_{\text{su}}^{t,s} \geq z_{\text{on}}^{t,s} - z_{\text{on}}^{t-1,s} - z_{\text{sb}}^{t-1,s} \text{ for } t = 2, \dots, n \quad (4)$$

An additional constraint is added to avoid a stack from transitioning from off to standby state:

$$z_{\text{off}}^{t-1,s} + z_{\text{sb}}^{t,s} \leq 1 \text{ for } t = 2, \dots, n \quad (5)$$

AEL Model

The AEL is modeled using the widely adopted semi-empirical model proposed first by Ulleberg [6] and later extended by Sánchez et al. [7]. For brevity, only the equations that were linearized are mentioned in this paper.

The reversible cell voltage of an AEL cell can be written as:

$$U_{\text{rev}}^{t,s} = U_{\text{rev0}}^{t,s} + \frac{RT_{\text{stack}}^{t,s}}{2F} \ln\left(\frac{(p_{\text{cell}}^{t,s} - p_{\text{H}_2\text{O}})^{1.5}}{\alpha_{\text{H}_2\text{O}}}\right) \quad (6)$$

where, U_{rev} is the reversible cell voltage, U_{rev0} the standard reversible cell voltage, T_{stack} the stack temperature, F the Faraday number, p_{cell} the cell pressure, $p_{\text{H}_2\text{O}}$ the partial vapor pressure, and $\alpha_{\text{H}_2\text{O}}$ the water activity. The cell voltage U_{cell} is expressed as:

$$U_{\text{cell}}^{t,s} = U_{\text{rev}}^{t,s} + K_1 \frac{I_{\text{cell}}^{t,s}}{A_{\text{cell}}} + s\log\left(K_2 \frac{I_{\text{cell}}^{t,s}}{A_{\text{cell}}} + 1\right) \quad (7)$$

where, I_{cell} is the cell current, A_{cell} the cell area, and s an empirical coefficient. K_1 and K_2 are two nonlinear terms that are dependent on T_{stack}. The Faraday efficiency η_{Faraday} is calculated as:

$$\eta_{\text{F}}^{t,s} = B_1 + B_2 \exp\left(K_3 \frac{A_{cell}}{I_{cell}^{t,s}}\right) \quad (8)$$

where, η_{F} is the Faraday efficiency, B_1 and B_2 two empirical coefficients, and K_3 a nonlinear term that depends on T_{stack}. At each time step, the total amount of hydrogen produced by an AEL stack $\dot{m}_{\text{prod, st}}$ is described by the Faraday law as:

$$\dot{m}_{\text{prod, st}}^{t,s} = \frac{\eta_{\text{F}}^{t,s} I_{\text{cell}}^{t,s}}{2F} M_{H_2} N_{\text{cell}} z_{\text{on}}^{t,s} \quad (9)$$

where, N_{cell} is the number of cells in the stack, and M_{H2} the molar mass of hydrogen. The total amount of hydrogen produced $\dot{m}_{\text{prod}}$ is the summation of the hydrogen produced by each stack:

$$\dot{m}_{\text{prod}}^{t} = \sum_{s=1}^{N_{\text{stack}}} \dot{m}_{\text{prod, st}}^{t,s} \quad (10)$$

The power consumed by an AEL stack P_{stack} can be written as:

$$P_{\text{stack}}^{t,s} = N_{\text{cell}} U_{\text{cell}}^{t,s} I_{\text{cell}}^{t,s} z_{\text{on}}^{t,s} + P_{\text{sb}} z_{\text{sb}}^{t,s} \quad (11)$$

where, U_{cell} is the cell voltage that is assumed to be same for all the cells in a stack, and P_{sb} the power consumed by a stack in its standby state. P_{stack} has to be between the maximum capacity P_{cap} and the minimum load P_{min} of the stack:

$$P_{\text{stack}}^{t,s} \leq P_{\text{cap}} z_{\text{on}}^{t,s} + P_{\text{sb}} z_{\text{sb}}^{t,s} \quad (12)$$

$$P_{\text{stack}}^{t,s} \geq P_{\text{min}} z_{\text{on}}^{t,s} + P_{\text{sb}} z_{\text{sb}}^{t,s} \quad (13)$$

Constraints (6) to (8) were approximated using a composite linearization approach, where piecewise approximation [11] and the McCormick envelope method [12] were used sequentially. Furthermore, the multilinear terms in equations (9) and (11) were relaxed using the same McCormick relaxation. The linearization methods used in this study are discussed in detail in the next section.

Power Balance

At every time step, the power balance of the whole system can be expressed by:

$$P_{\text{PV}}^{t} + P_{\text{ORC}}^{t} + P_{\text{grid}}^{t} = P_{\text{e}}^{t} + P_{\text{comp}}^{t} \quad (14)$$

where, P_{PV} is the power produced by the PV plant, P_{ORC} is the power produced by the ORC, P_{e} is the total power consumed by the AEL system, and P_{comp} is the power consumed by the hydrogen storage compressor. P_{e} is the summation of the power consumption of each stack, expressed as:

$$P_{\text{e}}^{t} = \sum_{s=1}^{N_{\text{stack}}} P_{\text{stack}}^{t,s} \quad (15)$$

P_{comp} is a function of the specific power consumption of the compressor w_{comp} and the storage inlet mass flow rate $\dot{m}_{\text{sto, in}}$ [13]. To keep the model simple, w_{comp} is taken as a constant parameter. Therefore, the constraint is written as:

$$P_{\text{comp}}^{t} = w_{\text{comp}} \dot{m}_{\text{sto, in}}^{t} \quad (16)$$

P_{PV} is modeled as a function of solar irradiance and module efficiency [14]. The polynomial equation to calculate P_{ORC}, expressed as a function of the temperature of the heat source and the cooling air (ambient) temperature, was obtained from the manufacturer of the ORC unit while the heat source temperature data were provided by the refinery.

Mass Balance

The mass balance of the system at every time step is written as:

$$\dot{m}^t_{\text{prod}} = \dot{m}^t_{\text{d}} + \dot{m}^t_{\text{sto, in}} - \dot{m}^t_{\text{sto, out}} \quad (17)$$

where, $\dot{m}_{\text{prod}}$ is the total hydrogen produced, $\dot{m}_{\text{d}}$ the hydrogen demand, and $\dot{m}_{\text{sto, out}}$ the mass flow rate of the hydrogen exiting the storage.

Hydrogen Storage

The operation of the hydrogen storage at each time step can be expressed by the following two constraints:

$$m^t_{\text{sto}} = m^{t-1}_{\text{sto}} + (\dot{m}^t_{\text{sto, in}} - \dot{m}^t_{\text{sto, out}})\Delta t \quad (18)$$

$$m^{t=0}_{\text{sto}} = m_{\text{sto, ini}} \quad (19)$$

where, m_{sto} is the amount of stored hydrogen at any time step, and $m_{\text{sto, ini}}$ is the amount of hydrogen stored at the start of the operation. Both $\dot{m}_{\text{sto, in}}$ and $\dot{m}_{\text{sto, out}}$ are limited by the maximum inlet and outlet flowrates of the storage, $\dot{m}_{\text{sto, in, max}}$ and $\dot{m}_{\text{sto, out, max}}$, respectively:

$$\dot{m}^t_{\text{sto, in}} \leq \dot{m}_{\text{sto, in, max}} \quad (20)$$

$$\dot{m}^t_{\text{sto, out}} \leq \dot{m}_{\text{sto, out, max}} \quad (21)$$

The maximum amount of stored hydrogen is limited by the volumetric storage capacity $V_{\text{sto, max}}$ and is given by:

$$m^t_{\text{sto}} \leq V_{\text{sto, max}}\rho_{H_2} \quad (22)$$

where, the density of hydrogen, ρ_{H_2}, is calculated using the standardized equation for hydrogen gas density, given in [15]. Two further constraints are added to prevent the storage from being charged and discharged at the same time:

$$\dot{m}^t_{\text{sto, in}} \leq z^t_{\text{ch}}\dot{m}_{\text{sto, in, max}} \quad (23)$$

$$\dot{m}^t_{\text{sto, out}} \leq (1 - z^t_{\text{ch}})\dot{m}_{\text{sto, out, max}} \quad (24)$$

where, z_{ch} is a binary variable indicating the charging of the storage.

LINEARIZATION METHODS

The nonlinearities in the optimization problem, described by equation (1) to (24), can be divided into two categories. The first category contains the univariate nonlinear terms that are linearized using piecewise linear approximation methods. Three different piecewise linearization techniques [11] are implemented and are discussed below. The second category consists of the multilinear terms, i.e. the bilinear and trilinear terms, that are approximated by the McCormick relaxation [12].

Piecewise Linearization Methods

Convex combination (CC)

A nonlinear function $f(x)$ is approximated by discretizing the domain of x into $k-1$ segments with k being the number of breakpoints. The linear approximations of x and $f(x)$ are weighted sums (convex combination) of their values at the breakpoints, where the weights λ_i sum to 1. In (29) to (32), the binary variables z_i ensure that at most two adjacent λ_i can be non-zero. The convex combination method can be written as:

$$f = \sum_{i=1}^{k} f(x_i)\,\lambda_i \quad (25)$$

$$x = \sum_{i=1}^{k} x_i\,\lambda_i \quad (26)$$

$$\lambda_i \geq 0 \quad (27)$$

$$\sum_{i=1}^{k} \lambda_i = 1 \quad (28)$$

$$\sum_{i=1}^{k-1} z_i = 1 \quad (29)$$

$$\lambda_1 \leq z_1 \quad (30)$$

$$\lambda_i \leq z_{i-1} + z_i \text{ for } i = 2, \dots, k-1 \quad (31)$$

$$\lambda_k \leq z_{k-1} \quad (32)$$

Convex combination with SOS2 variables (CC-SOS2)

To avoid the use of binary variables z_i, as seen in (29) to (32), the set of variables $\{\lambda_1, \lambda_2, \dots, \lambda_n\}$ can be formulated as a Special Ordered Set of type 2 (SOS2), which ensures the adjacency rule of the λ_i variables. With $\lambda_i \in$ SOS2, the formulation is then fully described by (24) to (27). It has to be noted that a suitable solver capable of handling SOS2 variables needs to be chosen.

Slope based Big-M formulation

To avoid the use of continuous variables λ_i in (25) to (32), the nonlinear function $f(x)$ can be approximated using a big-M formulation. Constraints (33) and (34) restrict the variable x to lie within the active segment, while (35) and (36) employ a Big-M formulation to define the function value f by the slope and the intercept of the selected segment.

$$x \geq x_i - (1 - z_i)(x_{\max} - x_{\min}) \quad (33)$$

$$x \leq x_{i+1} + (1 - z_i)(x_{\max} - x_{\min}) \quad (34)$$

$$f \geq a_i x + b_i - (1 - z_i)M \quad (35)$$

$$f \leq a_i x + b_i + (1 - z_i)M \quad (36)$$

$$\sum_{i=1}^{k-1} z_i = 1 \quad (37)$$

where, $a_i = \frac{f(x_{i+1}) - f(x_i)}{x_{i+1} - x_i}$, and $b_i = f(x_i) - a_i x_i$. M is a sufficiently large positive constant used to activate or deactivate constraints (35) and (36).

McCormick Relaxation

The McCormick envelope relaxes the product of two continuous variables, $w = xy$, where $x \in [x_{\text{lb}}, x_{\text{ub}}]$ and $y \in [y_{\text{lb}}, y_{\text{ub}}]$, by defining the convex hull of the feasible set through the following linear inequalities:

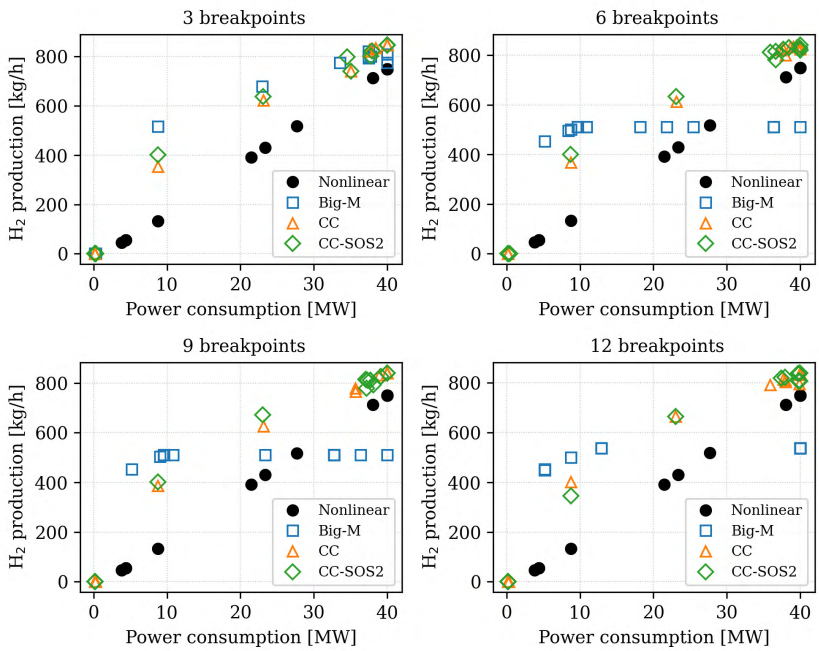

Figure 2. Hydrogen production-power consumption curve: NLP vs. MILP formulations using different linear approximation methods with 3 breakpoints (**top left**), 6 breakpoints (**top right**), 9 breakpoints (**bottom left**) and 12 breakpoints (**bottom right**).

$$w \geq x_{\mathrm{lb}} y + y_{\mathrm{lb}} x - x_{\mathrm{lb}} y_{\mathrm{lb}} \tag{38}$$

$$w \geq x_{\mathrm{ub}} y + y_{\mathrm{ub}} x - x_{\mathrm{ub}} y_{\mathrm{ub}} \tag{39}$$

$$w \leq x_{\mathrm{ub}} y + y_{\mathrm{lb}} x - x_{\mathrm{ub}} y_{\mathrm{lb}} \tag{40}$$

$$w \leq x_{\mathrm{lb}} y + y_{\mathrm{ub}} x - x_{\mathrm{lb}} y_{\mathrm{ub}} \tag{41}$$

If the bilinear term consists of a continuous variable $x \in [x_{\mathrm{lb}}, x_{\mathrm{ub}}]$ and a binary variable $z \in \{0,1\}$, the McCormick relaxation reduces to the following inequalities:

$$x_{\mathrm{lb}} z \leq w \leq x_{\mathrm{ub}} z \tag{42}$$

$$x - x_{\mathrm{ub}}(1-z) \leq w \leq x - x_{\mathrm{lb}}(1-z) \tag{43}$$

RESULTS

The NLP is solved using a multi-start approach using the local nonlinear solver IPOPT, in which the optimization problem is executed 50 times from random initial points and the solution with the minimum objective value is taken as the best guess for the global minimum. The MILP is solved with the linear solver Gurobi using the three piecewise linearization methods, described in the section above, by gradually increasing the number of breakpoints to potentially enhance the accuracy of the decision variables. The optimization was done over a 15-hour period. The dimensions of the different formulations are listed in **Table 1**. All the simulations were run using a computer with a 13th Gen Intel(R) Core(TM) i7-1365U processor.

Table 1: Dimensions of the MILP formulations and the NLP

Method	Variable count	Constraint count
CC (3 breakpoints)	8040	11257
CC (6 breakpoints)	12360	13417
CC (9 breakpoints)	16680	15577
CC (12 breakpoints)	21000	17737
CC-SOS2 (3 breakpoints)	6600	8377
CC-SOS2 (6 breakpoints)	8760	8377
CC-SOS2 (9 breakpoints)	10920	8377
CC-SOS2 (12 breakpoints)	13080	8377
Big-M (3 breakpoints)	5880	12457
Big-M (6 breakpoints)	8040	21097
Big-M (9 breakpoints)	10200	29737
Big-M (12 breakpoints)	12360	38377
NLP	2655	2287

We first assess the accuracy of the linearization methods by comparing the hydrogen production curve (hydrogen production versus power consumption) obtained from the MILP formulations using the three linearization methods with that of the NLP. Given that the NLP formulation retains the original nonlinear model, we take the hydrogen production curve from the NLP as the reference. We see from **Figure 2** that the curves from all of

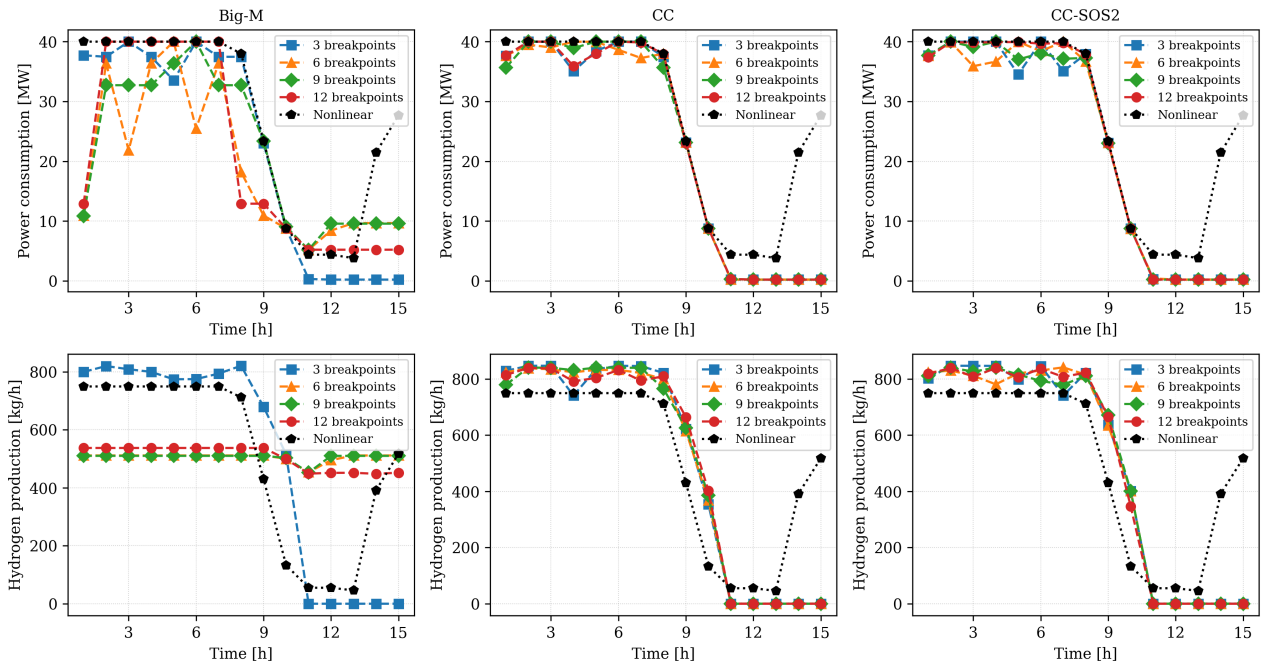

Figure 3. Power consumption (**top**) and hydrogen production (**bottom**) profiles: MILP formulations with different linearization methods vs. NLP.

the MILP formulations deviate from that of the NLP, but to varying degree. As expected, the CC and the CC-SOS2 produce very similar results since their underlying principles are the same. The performance of the big-M formulation heavily depends on the choice of the big-M parameter, which is difficult to appropriately determine. As a rule of thumb, the big-M was initially set equal to the difference between the maximum and minimum values of the nonlinear function being linearized. While this was sufficiently large for three breakpoints, the big-M value had to be increased every time the number of breakpoints was increased to prevent infeasibility, causing the feasible region to increase resulting in loose relaxation. Consequently, the big-M linearization method exhibits larger deviation in the hydrogen production curve as the number of breakpoints increases. This is also evident in **Figure 3**, where we see that beyond three breakpoints, the big-M formulation produces unusable results. Due to the difficulty in choosing an adequately large value of big-M, we come to the conclusion that the CC and the CC-SOS2 methods are the better choices for linearizing the nonlinear AEL model.

A further interesting observation is the impact of the number of breakpoints on the accuracy of the hydrogen production curve. As can be seen in **Figure 2**, increasing the number of breakpoints does not noticeably improve the accuracy of the curve, which is somehow unexpected. This behavior can be explained by the fact that the univariate nonlinear terms might exhibit only mild nonlinearities within the specified variable bounds. Future investigations will be carried out to understand this behavior in more detail.

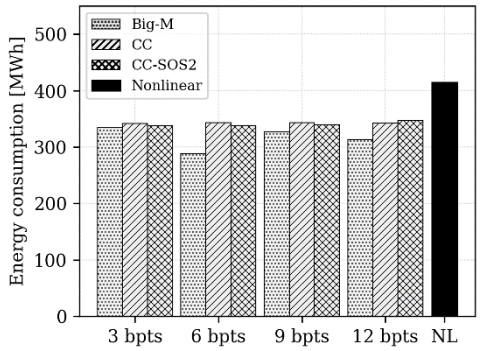

Figure 4. Total energy cosumption over a 15-hour period: MILP formulations vs. NLP.

We also see from both **Figure 2** and **Figure 3** that regardless of the linearization method used, compared to the NLP, the MILP formulations always underestimate the power consumption per kilogram of hydrogen produced. The constraints for the hydrogen production (equation 9) and the power consumption (equation 11) contain bilinear and trilinear terms, which were relaxed in the MILP using McCormick envelope. Since the McCormick relaxation provides an outer approximation of a multilinear term, the relaxation can lie below the actual function value causing the underestimation. This is further evident in the total energy consumption of the AEL over the 15-hour

optimization period, shown in **Figure 4**. However, a partitioned McCormick envelope, that will be implemented in the future, could provide a more accurate approximation.

Finally, we examine the total computational time of the different solution approaches. From **Table 2** we see that CC-SOS2 performs the fastest when having three breakpoints. However, as the number of breakpoints increases, the computational time of CC-SOS2 grows more rapidly than that of the standard CC formulation, making CC the more flexible option. As discussed previously, the Big-M method with more than three breakpoints does not produce satisfactory results. Therefore, its computation times are not considered meaningful for comparison. The reported MILP solution times were consistent across multiple runs. In contrast, the time taken for each run of the NLP varies widely with the multi-start strategy taking a total of 257 seconds for 50 runs. This wide spread of the objective values, ranging from 7909.76 to 28620.28, indicates the presence of local optima and makes it difficult to estimate in advance how many runs are required to reliably identify the best solution. Advanced optimization algorithm, such as BO-IPOPT [16, 17], which is a hybrid algorithm based on Bayesian optimization and IPOPT, will be employed in future work to address this issue.

Table 2: Solution times of the MILP formulations and the NLP.

Method	Solution time (seconds)
CC (3 breakpoints)	1.12
CC (6 breakpoints)	10.08
CC (9 breakpoints)	54.45
CC (12 breakpoints)	100.19
CC-SOS2 (3 breakpoints)	0.92
CC-SOS2 (6 breakpoints)	33.95
CC-SOS2 (9 breakpoints)	101.37
CC-SOS2 (12 breakpoints)	148.47
Big-M (3 breakpoints)	1.72
Big-M (6 breakpoints)	7.45
Big-M (9 breakpoints)	9.96
Big-M (12 breakpoints)	40.39
NLP	0.82-61.14 (single run) 257.03 (50 runs)

CONCLUSION

This study addresses a major challenge of integrating a nonlinear AEL model into a linear optimal dispatch problem of a renewable powered industrial AEL system. Two approaches were taken and compared. The first approach includes solving the operational optimization problem as an NLP by relaxing the binaries and integers within the model to continuous variables and then adding associated penalty terms to promote integrality. The NLP is then run using a multi-start strategy from random initial points. The minimum objective value is taken as the best guess for the global minimum. The second approach formulates the optimal dispatch problem as an MILP by linearizing the univariate nonlinear terms within the AEL model using piecewise approximation and relaxing the multilinear terms using McCormick envelope. The hydrogen production curve of the NLP is taken as reference since the NLP uses the original nonlinear AEL model.

We see that out of the three piecewise linearization methods, namely, CC, CC-SOS2 and big-M formulation, CC and CC-SOS2 perform very similarly in reproducing the hydrogen production curve irrespective of the number of breakpoints. The big-M method performs well only with three breakpoints. Beyond three breakpoints, choosing the big-M value becomes significantly difficult, requiring large values, and producing unusable results.

Another important, albeit surprising, observation is that increasing the number of breakpoints does not have any meaningful impact on the accuracy of the solution, possibly due to the univariate nonlinear terms within the AEL model exhibiting mild nonlinearities within the specified variable bounds.

Furthermore, we see that the MILP always underestimates the power consumption of the AEL when compared to the NLP. The McCormick relaxations of the hydrogen production and power consumption constraints can be attributed to this underestimation.

Finally, when taking a look at the computational times, we observe that the CC method is generally the fastest when working with a varied number of breakpoints. While the CC-SOS2 is marginally faster when having three breakpoints, the computational time increases much faster as the number of breakpoints increases making CC the more flexible option. We consider the big-M method not suitable for this optimization problem as it cannot produce usable results when more than three breakpoints are used and exhibits the longest computation time even in the three-breakpoint case. Comparing the computation time of the MILP with the NLP is difficult as the NLP is modeled with a multi-start approach with the time taking for each run varying widely. The presence of local optima in the NLP means it is hard to estimate how many runs are necessary to reliably obtain the best solution.

Future work includes potentially increasing the accuracy of the linearized AEL model by using the partitioned McCormick envelope that yields a tighter relaxation of multilinear terms by subdividing the variable bounds into smaller intervals. Furthermore, the MILP approach will be compared with global hybrid modeling approaches, e.g., BO-IPOPT [16, 17], and investigated as a potential initialization scheme for these hybrid methods. Additionally, these models will be validated with real data from a real plant before being evaluated for real-time digital twin implementation

ACKNOWLEDGEMENTS

The authors acknowledge the funding provided through the Establishing **E**uropean **P**roduction of **HY**drogen **R**enew**A**ble energy and integration into an industrial environment (EPHYRA) project, supported by the Clean Hydrogen Partnership in collaboration with Hydrogen Europe and Hydrogen Europe Research, under Grand Agreement no. 101112220.

AUTHOR IDENTIFIERS

Jonaed Bin Mustafa Kamal: 0009-0001-8155-0471
Loukas Kyriakidis: 0009-0003-6634-8579
Saskia Bublitz: 0000-0002-9050-7862
Bogdan Dorneanu: 0000-0003-3553-6625
Harvey Arellano-Garcia: 0000-0002-8297-0232

REFERENCES

1. Matute G, Yusta JM, Beyza J, Correas LC. Multi-state techno-economic model for optimal dispatch of grid connected hydrogen electrolysis systems operating under dynamic conditions. International Journal of Hydrogen Energy 46:1449-1460 (2021). https://doi.org/10.1016/j.ijhydene.2020.10.019
2. Bolívar Jaramillo L, Weidlich A. Optimal microgrid scheduling with peak load reduction involving an electrolyzer and flexible loads. Applied Energy 169:857-865 (2016). https://doi.org/10.1016/j.apenergy.2016.02.096
3. Korpas M, Holen AT. Operation planning of hydrogen storage connected to wind power operating in a power market. IEEE Trans. On Energy Conversion 21:742-749 (2006). https://doi.org/10.1109/tec.2006.878245
4. Baumhof MT, Raheli E, Johnsen AG et al. Optimization of hybrid power plants: when is a detailed electrolyzer model necessary? https://doi.org/10.48550/arXiv.2301.05310 (2023)
5. Hu S, Guo B, Ding S, Yang F, Dang J, Liu B, Gu J, Ma J, Ouyang M. A comprehensive review of alkaline water electrolysis mathematical modeling. Applied Energy 327:120099 (2022). https://doi.org/10.1016/j.apenergy.2022.120099
6. Ulleberg O. Modeling of advanced alkaline electrolyzers: a system simulation approach. International Journal of Hydrogen Energy 28:21-33 (2003). https://doi.org/10.1016/s0360-3199(02)00033-2
7. Sánchez M, Amores E, Rodríguez L, Clemente-Jul C. Semi-empirical model and experimental validation for the performance evaluation of a 15 kw alkaline water electrolyzer. International Journal of Hydrogen Energy 43:20332-20345 (2018). https://doi.org/10.1016/j.ijhydene.2018.09.029
8. Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math. Program. 106:25-57 (2005). https://doi.org/10.1007/s10107-004-0559-y
9. Dantzig G. Linear programming and extensions. RAND Corporation (1963). https://doi.org/10.7249/r366
10. EPHYRA. http://www.ephyraproject.eu
11. Lin MH, Carlsson JG, Ge D, Shi J, Tsai JF. A review of piecewise linearization methods. Mathematical Problems in Engineering 2013:1-8 (2013). https://doi.org/10.1155/2013/101376
12. McCormick GP. Computability of global solutions to factorable nonconvex programs: part i — convex underestimating problems. Mathematical Programming 10:147-175 (1976). https://doi.org/10.1007/bf01580665
13. Zheng Y, You S, Bindner HW, Münster M. Optimal day-ahead dispatch of an alkaline electrolyser system concerning thermal–electric properties and state-transitional dynamics. Applied Energy 307:118091 (2022). https://doi.org/10.1016/j.apenergy.2021.118091
14. Kyriakidis L, Kansara R, Roldán Serrano MI. Energy management of industrial energy systems via rolling horizon and hybrid optimization: a real-plant application in germany. Energies 18:3977 (2025). https://doi.org/10.3390/en18153977
15. Lemmon EW, Huber ML, Leachman JW. Revised standardized equation for hydrogen gas densities for fuel consumption applications. J. Res. Natl. Inst. Stand. Technol. 113:341 (2008). https://doi.org/10.6028/jres.113.028
16. Kyriakidis L, Mendez MA, Bähr M. A hybrid algorithm based on bayesian optimization and interior point optimizer for optimal operation of energy conversion systems. Energy 312:133416 (2024). https://doi.org/10.1016/j.energy.2024.133416
17. Kyriakidis L, Bähr M, Mendez MA. Enhanced hybrid algorithm based on bayesian optimization and interior point optimizer for constrained optimization. Optim Eng : (2025). https://doi.org/10.1007/s11081-025-09975-y

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Set-based Formulations for the State Task Network Scheduling Problem

David A. Liñán[a], Georgia Stinchfield[b], Carl D. Laird[b], and Jan Kronqvist[a]*

[a] KTH Royal Institute of Technology, Department of Mathematics, Stockholm, SE-10044, Sweden
[b] Carnegie Mellon University, Department of Chemical Engineering, Pittsburgh, PA-15213, USA
* Corresponding Author: jankr@kth.se.

ABSTRACT

The state task network (STN) representation is a widely used modeling approach for optimal multipurpose batch production scheduling. In practice, STNs have been traditionally formulated as mixed-integer programming (MIP) problems and solved using general-purpose MIP solvers relying on branch-and-bound and branch-and-cut. In the meantime, alternative modeling and solution paradigms for optimization have been developed, enabling the incorporation of alternative variable types and optimization algorithms. Specifically, this work relies on the Hexaly software, which introduced set-based models and their solution through general-purpose hybrid algorithms, i.e., methods that combine traditional MIP with constraint programming, local search, large neighborhood search, among other tools. So far, Hexaly has shown promising results when tackling optimal scheduling problems, however, set-based models and solution approaches for STN optimization have not been studied in the literature. Aiming to fill this gap, this work introduces the first set-based models for the STN problem and performs preliminary benchmarking tests with Hexaly's general-purpose hybrid algorithms. Results from two case studies suggest that set-based models may perform better than the traditional MIP STN formulation when dealing with a simple STN with sequential connectivity and a long-term scheduling instance of a STN with fixed integer batching variables. Overall, this work establishes the foundations for advancing research on set-based approaches for network scheduling, opening new directions beyond traditional MIP optimization.

Keywords: Scheduling, Optimization, Modelling, Process Operations, Batch Systems

INTRODUCTION

Scheduling problems represent some of the most prominent applications of mathematical optimization [1]. In particular, the state task network (STN) [2] scheduling model is a seminal mixed-integer programming (MIP) formulation that led to the development of a variety of network scheduling MIP formulations such as the resource task network (RTN) [3] model, among other extensions [4, 5]. The main practical feature of network scheduling models is their flexibility in supporting complex tasks with multiple inputs and outputs, enabling material mixing and splitting, i.e., the STN and RTN models do not require materials to be handled sequentially [6]. This modeling flexibility has led to a variety of real-world applications, e.g., optimal scheduling in a vegetal oil refinery [7], multipurpose batch plant scheduling in the chemical-pharmaceutical industry [8], electricity scheduling for cement plants [9], among others [10, 11]. Despite several extensions and variants are available in the literature, this work concentrates on the standard STN formulation.

Traditional mathematical programming has powered the generation of optimal STN scheduling solutions in chemical engineering. Most works in the literature rely on MIP modeling and algorithms, e.g., branch-and-bound and branch-and-cut [10]. This is because STN scheduling applications are usually linear and only require binary or integer variables combined with continuous variables. Nonetheless, the generalized disjunctive programming (GDP) modeling formalism has also been proposed to derive STN formulations, through the incorporation of Boolean variables, disjuncts and logic relationships [12, 13]. In addition to MIP and GDP, a higher-level modeling formalism available in the literature is set-based modeling, which not only enables all the modeling features from MIP and GDP, but also introduces the idea of using sets as

variables. Set-based optimization has resulted in compact formulations yielding computational advantages compared to MIP in a variety of scheduling applications, e.g., job-shop, open-shop, resource-constrained, among others [14]. However, the field of optimal set-based STN scheduling has not yet been explored.

This work brings novelty by deriving and testing the first set-based formulations for the STN scheduling problem. Set-based models introduce three key variable types, namely set variables, list variables and interval variables. In this work, we aim to contribute novel mixed-interval disjunctive program (MIDP) and mixed-interval-list disjunctive program (MILDP) STN models. Our secondary objective is to evaluate the computational performance of commercial software Hexaly given these set-based formulations. How we choose to model the STN directly contributes to the efficiency with which the optimization problem is solved. However, deriving the best set-based model (e.g., intervals, intervals + sets, intervals + lists) is not straightforward. While Hexaly exploits this type of modeling, the progress made by their proprietary algorithms and heuristics is not directly visible to the practitioner. Thus, observing how different formulations of the STN interact with Hexaly is vital for determining which combinations of set-based modeling paradigms are most efficient.

This study is organized as follows: first we provide modeling background. Then, we describe the proposed MIDP and MILDP STN models. Two case studies are then used to illustrate the performance of the proposed set-based formulations. Concluding remarks and future work are provided at the end.

BACKGROUND

Background on STN MIP modeling is provided first, followed by background on set-based modeling.

Mixed-integer program (MIP)

Optimal STN scheduling relies on four key fixed sets that are used to define a problem: tasks $i \in I$ (e.g., $i \in$ {manufacture product A, separate A from B, ...}); units $j \in J$ available to perform these tasks (e.g., $j \in$ {reactor 1, reactor 2, distillation column, ...}); states $k \in K$ that act as initial, intermediate or final storage defining the feasible interconnection between tasks, and a scheduling horizon that is discretized into $n_T + 1$ time points $t \in T = \{0, 1, 2, \dots, n_T\}$. Based on these four sets, additional fixed sets are used to specify the interconnections between tasks, units and states, namely

$$IJ^{\times} = \{(i,j) \in I \times J |\ i \text{ can be performed in } j\}, \quad (1)$$

$$IK^{+} = \{(i,k) \in I \times K |\ i \text{ feeds } k\}, \quad (2)$$

$$IK^{-} = \{(i,k) \in I \times K |\ k \text{ feeds } i\}. \quad (3)$$

Generally, two variable types need to be optimized in a STN problem: discrete timing variables and continuous batch size variables. In a traditional MIP model, timing variables $x_{i,j,t} \in \{0,1\}$ decide if task i is executed in unit j starting at time t. Non-overlapping constraints are needed to ensure that a given unit can only perform one task at a time, i.e.,

$$\sum_{(i,j)\in IJ^{\times}} \sum_{\substack{i\in I \\ t'\in T \\ t-\tau_{i,j}+1\le t'\le t}} x_{i,j,t'} \le 1, \forall j \in J, \forall t \in T, \quad (4)$$

where $\tau_{i,j}$ is the number of discrete time points needed to execute task i in unit j, i.e., a discrete processing time.

Batch size variables $b_{i,j,t} \in [0, \beta_j^{max}] \subset \mathbb{R}$ define the amount of material processed in task i by unit j at time t, where $\beta_{i,j}^{max}$ and $\beta_{i,j}^{min}$ are the fixed maximum and minimum capacities of unit j when executing task i. Variable $b_{i,j,t}$ is bounded by $\beta_{i,j}^{max}$ and $\beta_{i,j}^{min}$ when a task is executed, as defined by the following capacity constraints:

$$\beta_{i,j}^{min} x_{i,j,t} \le b_{i,j,t} \le \beta_{i,j}^{max} x_{i,j,t}, \forall (i,j) \in IJ^{\times}, \forall t \in T. \quad (5)$$

STN problems stand out due to the consideration of states k, representing storage facilities in which incoming materials from previous task executions can be stored until they are released to subsequent tasks. This means that there is a mass balance inherently embedded within the scheduling problem. These mass balances are needed to ensure that, at a given time point, enough materials are available for the execution of subsequent tasks and storage facilities do not overflow. Mathematically, storage modeling requires intermediate storage states $s_{k,t} \in \mathbb{R}$ to define the mass balance constraints in (6) and (7), as well as the state tracking constraints in (8).

$$s_{k,0} = \gamma_k^{init} - \sum_{\substack{(i,j)\in IJ^{\times} \\ (i,k)\in IK^{-}}} \rho_{i,k}^{-} b_{i,j,0}, \forall k \in K, \quad (6)$$

$$s_{k,t} = s_{k,t-1} + \sum_{\substack{(i,j)\in IJ^{\times} \\ (i,k)\in IK^{+}, t-\tau_{i,j}\ge 0}} \rho_{i,k}^{+} b_{i,j,t-\tau_{i,j}} - \sum_{\substack{(i,j)\in IJ^{\times} \\ (i,k)\in IK^{-}}} \rho_{i,k}^{-} b_{i,j,t}, \forall k \in K, \forall t \in T\backslash\{0\}, \quad (7)$$

$$\gamma_k^{min} \le s_{k,t} \le \gamma_k^{max}, \forall k \in K, \forall t \in T. \quad (8)$$

These constraints introduce parameter γ_k^{init} denoting the initial storage in state k; the maximum (γ_k^{max}) and minimum (γ_k^{min}) storage capacity of state k, and the fraction of material in state k consumed ($\rho_{i,k}^{-}$) or produced ($\rho_{i,k}^{+}$) by task i.

A common completion constraint is that all executed tasks must have finished by the end of the scheduling horizon, which can be expressed as

$$x_{i,j,t} = 0, \forall (i,j) \in IJ^{\times}, \forall t \in \{n_T - \tau_{i,j} + 1, \dots, n_T\}. \quad (9)$$

The constraints in (4)-(9) are standardly used to model STN problems, however, additional requirements are often imposed depending on specific application

requirements. Other modeling features considered in the case study section include constraint (10) to ensure that every task is executed at most once, or constraint (11) to ensure that the number of times task i is executed in unit j matches integer batching parameter $\psi_{i,j}$.

$$\sum_{\substack{j\in J\\(i,j)\in IJ^{\times}}} \sum_{t\in T} x_{i,j,t} \leq 1, \forall i \in I, \forall t \in T, \tag{10}$$

$$\sum_{t\in T} x_{i,j,t} = \psi_{i,j}, \forall (i,j) \in IJ^{\times}. \tag{11}$$

Set-based modeling

Set-based modeling relies on three variable types that are not used in MINLP or GDP modeling defined as:

Interval variables

Let $n, m \in \mathbb{Z}$ be fixed numbers with $m \geq n$ and $(\!(n,m)\!) = \{(n',m') \in \mathbb{Z}^2 |\ n' \geq n, m' \leq m, m' \geq n'\}$. We define an *interval variable* $(x,y) \in (\!(n,m)\!)$, i.e., (x,y) can take tuples in $(\!(n,m)\!)$ as values.

For instance, take $n = 0$ and $m = 2$. Then, it is possible to define interval variable $(x,y) \in (\!(n,m)\!)$, where $(\!(n,m)\!)$ is equal to $\{(0,1),(0,2),(1,2),(0,0),(1,1),(2,2)\}$. Note that interval variables are herein defined as tuples, but they resemble intervals in the sense that the second entry y (the "end") is always greater than or equal to the first entry x (the "start"). By convention, an interval variable is said to be empty when its start and end points match.

Set variables

Consider some fixed set A and take $\mathbb{P}(A)$ as the power set of A, i.e., the set of subsets of A. We define a *set variable* $X \in \mathbb{P}(A)$, i.e., X can take sets in $\mathbb{P}(A)$ as values.

For example, take $A = \{0,1\}$. Then, it is possible to define set variable $X \in \mathbb{P}(A)$ with $\mathbb{P}(A) = \{\{0\},\{0,1\},\{1\},\{\}\}$. Note that sets $\{0,1\}$ and $\{1,0\}$ are the same, i.e., set variables do not distinguish between different orderings of elements. List variables are similar to set variables, except that these account for permutations.

List variables

Consider some fixed set A and its power set $\mathbb{P}(A)$. Also, take $\mathbb{S}(A) = \bigcup_{X\in\mathbb{P}(A)} \mathcal{S}(X)$, where $\mathcal{S}(X)$ is the set of lists obtained by permuting the elements of X. We define a *list variable* $y \in \mathbb{S}(A)$, i.e., y can take lists in $\mathbb{S}(A)$ as values.

Again, take $A = \{0,1\}$ as an example. Then, it is possible to define list variable $y \in \mathbb{S}(A)$ with $\mathbb{S}(A) = \{\langle 0\rangle, \langle 0,1\rangle, \langle 1,0\rangle, \langle 1\rangle, \langle\rangle\}$. Note that set and list variables may take the empty set $\{\}$ or empty list $\langle\rangle$ as valid values.

In the previous definitions, A is called a "fixed set" in the sense that it is not a set variable. In fact, Hexaly only supports A being a finite range set of the form $\{0,1,2,\dots,n-1\}$. However, Hexaly supports establishing one-to-one relationships between A and other sets of the same size, even if they contain variables, e.g., a one-to-one function between A and a set of interval variables that needs to be shuffled through a list variable. For notation simplicity, this work avoids one-to-one functions and allows A being a finite set of variables explicitly.

MATHEMATICAL MODELS

In this section, we introduce the constraints needed to formulate MIDP and MILDP set-based STN models.

Mixed-interval disjunctive program (MIDP)

We propose a STN MIDP model that replaces timing variables $x_{i,j,t} \in \{0,1\}$ by timing interval variables. To do that, we define execution subindex $q \in Q_{i,j}$, where $Q_{i,j}$ is the fixed set of all potential executions of task i in unit j defined as

$$Q_{i,j} = \{1,2,\dots,\lfloor n_T/\tau_{i,j}\rfloor\}, \forall (i,j) \in IJ^{\times}, \tag{12}$$

with $\lfloor\bullet\rfloor$ being the floor operator. This allows to define interval variables $(\mathfrak{x}^{start}_{i,j,q}, \mathfrak{x}^{end}_{i,j,q}) \in (\!(0,n_T)\!), \forall (i,j) \in IJ^{\times}, \forall q \in Q_{i,j}$. Interval $(\mathfrak{x}^{start}_{i,j,q}, \mathfrak{x}^{end}_{i,j,q})$ represents the potential execution q of task i in unit j, starting at time point $\mathfrak{x}^{start}_{i,j,q}$ (inclusive) and finishing at time point $\mathfrak{x}^{end}_{i,j,q}$ (exclusive).

For some $(i,j) \in IJ^{\times}$, not every interval $q \in Q_{i,j}$ is necessarily going to be executed. Therefore, we use the convention that the length of $(\mathfrak{x}^{start}_{i,j,q}, \mathfrak{x}^{end}_{i,j,q})$ is 0 (i.e., $\mathfrak{x}^{end}_{i,j,q} - \mathfrak{x}^{start}_{i,j,q} = 0$) when interval q is not executed. Otherwise, the interval is executed and its length is equal to the processing time of the task-unit pair ($\mathfrak{x}^{end}_{i,j,q} - \mathfrak{x}^{start}_{i,j,q} = \tau_{i,j}$), i.e., we impose the following interval presence constraints:

$$\left[\mathfrak{x}^{end}_{i,j,q} - \mathfrak{x}^{start}_{i,j,q} = \tau_{i,j}\right] \vee \left[\mathfrak{x}^{end}_{i,j,q} - \mathfrak{x}^{start}_{i,j,q} = 0\right],$$
$$\forall (i,j) \in IJ^{\times}, \forall q \in Q_{i,j}, \tag{13}$$

where $\vee$ is the or operator. Now, we introduce a new operator defined as $\mathbf{if}(\Omega) = \{1 \text{ if } \Omega = \text{True}, 0 \text{ if } \Omega = \text{False}\}$, where Ω is a logical condition. Hexaly supports this operator for a variety of logical conditions. This operator allows us to define non-overlapping constraints as

$$\sum_{\substack{i\in I\\(i,j)\in IJ^{\times}}} \sum_{q\in Q_{i,j}} \mathbf{if}\left(t \geq \mathfrak{x}^{start}_{i,j,q}\right) - \mathbf{if}\left(t \geq \mathfrak{x}^{end}_{i,j,q}\right) \leq 1,$$
$$\forall j \in J, \forall t \in T. \tag{14}$$

Next, we introduce MIDP capacity constraints as

$$\mathbf{if}\left(\bigvee_{q\in Q_{i,j}} \begin{bmatrix} \mathfrak{x}^{start}_{i,j,q} = t \\ \mathfrak{x}^{end}_{i,j,q} - \mathfrak{x}^{start}_{i,j,q} \neq 0 \end{bmatrix}\right) \beta^{min}_{i,j} \leq b_{i,j,t},$$

$$b_{i,j,t} \leq \mathbf{if}\left(\bigvee_{q\in Q_{i,j}} \begin{bmatrix} \mathfrak{x}_{i,j,q}^{start} = t \\ \mathfrak{x}_{i,j,q}^{end} - \mathfrak{x}_{i,j,q}^{start} \neq 0 \end{bmatrix}\right)\beta_{i,j}^{max},$$

$$\forall (i,j) \in IJ^{\times}, \forall t \in T, \tag{15}$$

which link timing interval variables with continuous batch size variables that remain unchanged with respect to the traditional MIP STN model. Note that $x_{i,j,t}$ is equivalent to the **if** operator multiplying $\beta_{i,j}^{min}$ and $\beta_{i,j}^{max}$ in (15), i.e., if $x_{i,j,t} = 1$, then there is some interval $q \in Q_{i,j}$ of non-zero length starting at time t.

Since batch size variables $b_{i,j,t}$ remain unchanged in this formulation, mass balance and state tracking constraints (6)-(8) remain unchanged with respect to the MIP formulation. Furthermore, MIP completion constraints in (9) are readily enforced by $(\mathfrak{x}_{i,j,q}^{start}, \mathfrak{x}_{i,j,q}^{end}) \in (\!(0, n_T)\!)$, i.e., the end of an interval $\mathfrak{x}_{i,j,q}^{end}$ is bounded above by n_T. In summary, MIP constraints (4), (5) and (9) are the counterpart of MIDP constraints (13)-(15), while both MIP and MIDP share constraints (6)-(8).

After defining the main components of the MIDP STN model, we now examine how the additional MIP modeling features in (10) and (11) can be incorporated into the MIDP formulation. An equivalent version of constraint (10) in the MIDP model to ensure that each task is executed at most once is of the form

$$\sum_{\substack{J\in I \\ (i,j)\in IJ^{\times}}} \sum_{q\in Q_{i,j}} \mathbf{if}\left(\mathfrak{x}_{i,j,q}^{end} - \mathfrak{x}_{i,j,q}^{start} > 0\right) \leq 1, \forall i \in I. \tag{16}$$

Regarding constraint (11), major modifications are needed if the number of times task i is executed in unit j is known and equal to $\psi_{i,j}$. Firstly, set $Q_{i,j}$ in (12) must be redefined as

$$Q_{i,j} = \{1, 2, \dots, \psi_{i,j}\}, \forall (i,j) \in IJ^{\times}, \tag{17}$$

So that only $\psi_{i,j}$ intervals are considered in the formulation. Secondly, these intervals must be scheduled, meaning that interval presence constraints in (13) simplify to the following interval length constraints:

$$\mathfrak{x}_{i,j,q}^{end} - \mathfrak{x}_{i,j,q}^{start} = \tau_{i,j}, \forall (i,j) \in IJ^{\times}, \forall q \in Q_{i,j}. \tag{18}$$

Lastly, all intervals have length $\tau_{i,j}$, thus the condition within the **if** operator in (15) simplifies and capacity constraints reduce to

$$\mathbf{if}\left(\bigvee_{q\in Q_{i,j}} \left[\mathfrak{x}_{i,j,q}^{start} = t\right]\right)\beta_{i,j}^{min} \leq b_{i,j,t},$$

$$b_{i,j,t} \leq \mathbf{if}\left(\bigvee_{q\in Q_{i,j}} \left[\mathfrak{x}_{i,j,q}^{start} = t\right]\right)\beta_{i,j}^{max},$$

$$\forall (i,j) \in IJ^{\times}, \forall t \in T. \tag{19}$$

Mixed-interval-list disjunctive program (MILDP)

The interval variables from the MIDP formulation can be included in sets that define list variables, so that some constraints of the MIDP formulation can be reformulated in terms of list variables in a MILDP model. To do that, we create the following sets with $Q_{i,j}$ defined in (12):

$$A_j = \{(\mathfrak{x}_{i,j,q}^{start}, \mathfrak{x}_{i,j,q}^{end}) \in (\!(0, n_T)\!) \mid (i,j) \in IJ^{\times}, q \in Q_{i,j}\},$$

$$\forall j \in J, \tag{20}$$

which for every unit j contain all the intervals (i.e., tasks) that can be potentially performed in that unit. To decide which intervals are going to be scheduled and their order, we define list variables $y_j \in \mathbb{S}(A_j)$. We denote $y_j = \langle y_j(1), y_j(2), \dots, y_j(|y_j|)\rangle$, where the components $y_j(p), \forall p \in \{1, 2, \dots, |y_j|\}$ provide ordered access to the individual intervals in y_j. With this new variable, interval presence constraints take the following form:

$$\mathfrak{x}_{i,j,q}^{end} - \mathfrak{x}_{i,j,q}^{start} = \mathbf{if}\left((\mathfrak{x}_{i,j,q}^{start}, \mathfrak{x}_{i,j,q}^{end}) \in y_j\right)\tau_{i,j},$$

$$\forall (i,j) \in IJ^{\times}, \forall q \in Q_{i,j}. \tag{21}$$

Also, non-overlapping constraints can now be compactly formulated as

$$y_j(p) < y_j(p-1), \forall p \in \{2, 3, \dots, |y_j|\}, \forall j \in J, \tag{22}$$

where the strict $<$ constraint between two intervals means that the second entry of the right-hand-side interval tuple is less than or equal to first entry of the left-hand-side interval tuple. Note that having $|y_j|$ as variable in (22) is an accepted cardinality expression in Hexaly's set-based modeling platform.

Capacity constraints remain unchanged with respect to those defined for the MIDP model in (15). Also, mass balance and state tracking constraints (6)-(8) remain unchanged in the MILDP formulation. Similarly to the MIDP model, completion constraints of tasks within the scheduling horizon are not needed because these are readily satisfied by interval variables. In summary, MIP constraints (4), (5) and (9) are the counterpart of MILDP constraints (15), (21) and (22), while both MIP and MILDP share constraints (6)-(8).

Now that the main MILDP constraints have been defined, we explain how constraints (10) and (11) are translated from the MIP formulation to the MILDP model. On the one hand, constraint (10) takes the form of constraint (16) from the MIDP model. On the other hand, imposing a constraint equivalent to (11) in the MILDP model requires defining set $Q_{i,j}$ as in (17) instead of (12), imposing interval length constraint (18) instead of (21), imposing capacity constraints (19) instead of (15), and enforcing the following constraint on the size of lists:

$$|y_j| = \sum_{\substack{i\in I \\ (i,j)\in IJ^{\times}}} \psi_{i,j}, \forall j \in J. \tag{23}$$

CASE STUDY

We assess the performance of the proposed MIDP and MILDP formulations in comparison with the traditional MIP formulation through two case studies. The objective functions and parameters defining these case studies are introduced below.

Case Study 1

This case study was proposed by Hirn and Grapendal [15], who proposed a STN model for the optimal scheduling of electromagnetic compatibility tests with limited scheduling time and equipment. That is, 15 devices of different sizes are scheduled to be tested in 3 testing chambers with different capacities. For this case study, each task represents a device that needs to be tested once; therefore, special constraint (10) is enforced. Additionally, a profit maximization objective is considered. For the MIP formulation the objective function is defined as

$$\max \sum_{(i,j)\in IJ\times} \sum_{t\in T} \upsilon_i x_{i,j,t}, \tag{24}$$

where υ_i is the fixed profit generated by one execution of task i, which readily considers the difference between revenue and task execution costs. Objective function (24) translates into

$$\max \sum_{(i,j)\in IJ\times} \sum_{q\in Q_{i,j}} \upsilon_i \mathbf{if}\left(\mathfrak{x}_{i,j,q}^{end} - \mathfrak{x}_{i,j,q}^{start} > 0\right), \tag{25}$$

for set-based formulations. The resulting optimization problem is shown in Table 1, while all the fixed parameters are available as supplementary material (Appendix A).

Table 1: MIP, MIDP and MILDP problem formulations.

MIP	MIDP	MILDP
Eq. (24), s.t. Eqs. (4)-(10).	Eq. (25), s.t. Eqs. (6)-(8), (12)-(15), (16).	Eq. (25), s.t. Eqs. (6)-(8), (12), (15), (16), (21), (22).

Case Study 2

This case study was adapted from the work by Velez and Maravelias [16], who studied MIP formulations for long-term STN scheduling. It considers the well-studied Kondili et al. [2] STN, which includes one mixing task performed in a separation unit, three reaction tasks scheduled in two reactors and one separation task with its corresponding separation unit. This STN has three inputs for which purchasing costs are considered and two products that generate revenue. For this case study, it is assumed that the number of times each task is executed in each unit is known, thus, special constraint (11) is enforced. In this case, the profit maximization objective takes the following form:

$$\max \sum_{k\in K} \omega_k s_{k,n_T} - \sum_{(i,j)\in IJ\times} \alpha_{i,j} \psi_{i,j}, \tag{26}$$

where ω_k is the revenue from selling the material in state k and $\alpha_{i,j}$ is the cost of running task i in unit j. Altogether, the optimization problem considered is summarized in Table 2, while all the fixed parameters are available in the supplementary material (Appendix B).

Table 2: MIP, MIDP and MILDP problem formulations.

MIP	MIDP	MILDP
Eq. (26), s.t. Eqs. (4)-(9), (11).	Eq. (26), s.t. Eqs. (6)-(8), (14), (17)-(19).	Eq. (26), s.t. Eqs. (6)-(8), (17)-(19), (22), (23).

RESULTS AND DISCUSSION

This section evaluates the performance of the MIP, MIDP, and MILDP models by solving multiple instances of case studies 1 and 2 across varying levels of complexity. All optimization runs were performed in the Python API of Hexaly 14.0 with a time limit of 3600 seconds on a 1.7 GHz, 64 GB, Intel Core® Ultra™ 7 165U, Dell Inc. Latitude 7350 laptop.

Case Study 1

Different instances of this case study were generated by testing different time horizons n_T from 5 to 17, while keeping the scheduling time step fixed. Since tasks can be executed only once, increasing n_T results in increased flexibility to accommodate more tasks within the scheduling horizon. The performance plots obtained with solution time and objective gap as performance metrics are shown in Figures 1 and 2, respectively.

In the y-axis, Figure 1 shows the fraction of problems solved to global optimality within the solution time shown in the x-axis. Similarly, Figure 2 shows the fraction of problems solved to either feasibility or global optimality within the percentual objective gap in the x-axis. From Figure 1, the MIDP formulation clearly outperforms the MIP and MILDP formulations for this case study. Within the time limit, the MIDP formulation can solve 11 out of the 13 tested instances to global optimality. On the contrary, the MIP and MILDP formulations are unable to close the objective gap to satisfactory levels (see Figure 2).

Figure 2 shows that the MILDP model achieves better effectiveness at closing the optimality gap reported by the solver compared to the standard MIP, despite not performing as well as the MIDP formulation. Furthermore, all formulations reached the same primal objective within the imposed time limit. This outcome suggests that the set-based formulations used enable Hexaly's optimization algorithm to compute tighter objective lower bounds for this test problem. Based on publicly available documentation, the observed increased performance of set-

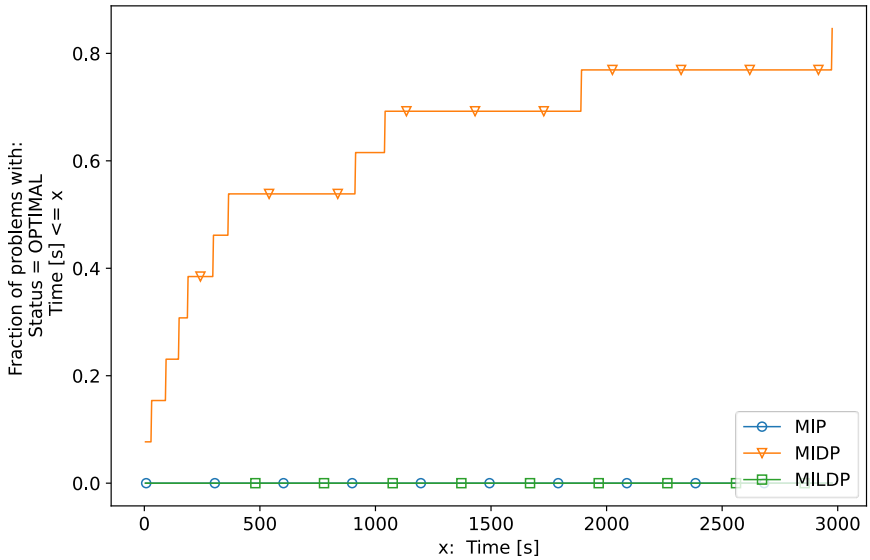

Figure 1. Solution time performance profile for case 1.

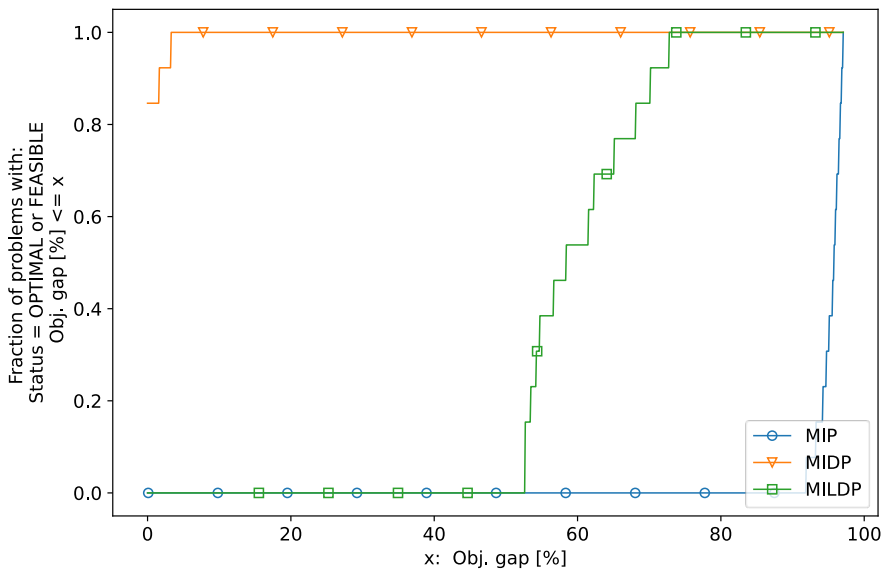

Figure 2. Objective gap performance profile for case 1.

based formulations at solving the dual problem and finding better objective bounds is attributed to a combination of column generation, column elimination and/or branch-cut-price strategies tailored to list and interval variables. Particularly, the energetic relaxation implemented in Hexaly has been shown to be effective for deriving lower bounds in scheduling problems [17].

The results obtained in this work align with the previous findings by Hirn and Grapendal [15], who highlighted how increasing the time horizon affects the number of nodes that a standard MIP algorithm (GLPK) must explore to close the optimality gap. As the time horizon n_T increases, more devices can be scheduled, which adds combinatorial complexity and forces GLPK to move from exploring roughly 1 node to about 1×10^7 nodes before finding the optimal solution, until all devices can finally be scheduled at n_T=16. This is in line with the best-performing set-based model (MIDP), which only struggles to close the optimality gap for the instances with n_T=14 and n_T=15, reporting gaps of 3.21% and 1.61%, respectively. In contrast, the MIP formulation solved with Hexaly consistently reports optimality gaps larger than 90%, further highlighting the computational benefits of the proposed MIDP model.

Case Study 2

A total of 20 different instances of this case study were generated by adjusting the fineness (or accuracy)

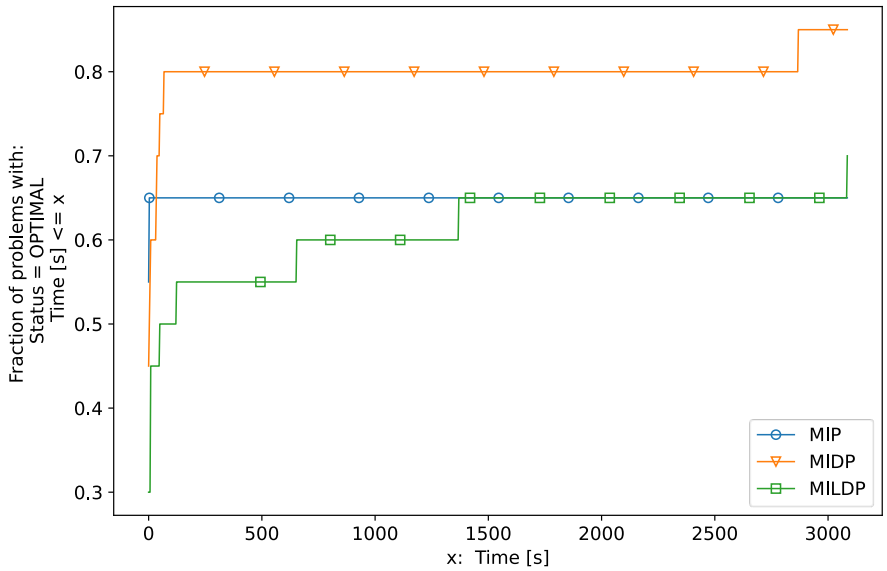

Figure 3. Solution time performance profile for case 2.

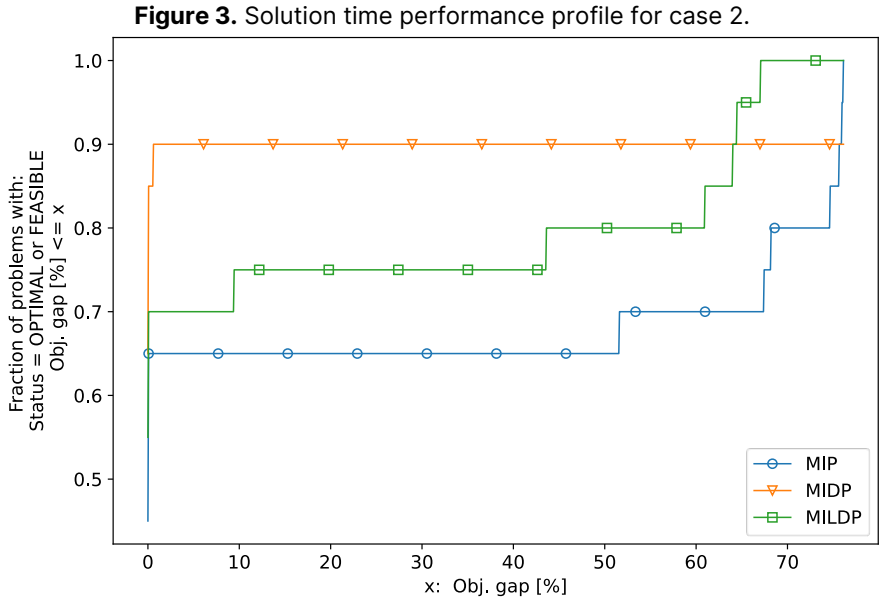

Figure 4. Objective gap performance profile for case 2.

for time discretization, ranging from 2 (i.e., $n_T = 1$), to 121 (i.e., $n_T = 120$) discretization points. The remaining parameters of the formulation were adjusted to maintain a physical scheduling horizon of 120 hours. The performance plots obtained with solution time and objective gap as performance metrics are shown in Figures 3 and 4, respectively.

Figure 3 illustrates improvements in the performance of the MIP formulation relative to the previous case study. Specifically, the MIP formulation outperforms the MILDP formulation for solution times below 1500 seconds. However, at higher time limits the MILDP model exhibits a slight increase in performance, solving one additional instance to global optimality within the 3600 seconds limit. In contrast, the MIDP formulation outperforms both the MIP and MILDP models, solving more than 80% of the test problems to global optimality.

Similarly to the previous case study, both the MIDP and MILDP models outperform the original MIP model in closing the objective gap (see Figure 4). However, not all formulations reach the same primal solution in this case study. The MILDP model consistently identified the best-known primal solution across all instances. This indicates that the advantage of set-based formulations arises not only from dual strategies that yield tighter objective lower bounds, but also from primal heuristics that quickly generate high-quality feasible solutions. Specifically, Hexaly leverages local search and large neighborhood search

strategies that operate directly on interval and list variable domains, eliminating the need for binary reformulations. For example, consider a task-unit pair (i, j). The MIDP and MILDP formulations avoid introducing one binary variable $x_{i,j,t}$ for each time point $t \in \{0, 1, 2, \dots, n_T\}$, and instead replace them with $\psi_{i,j}$ interval variables, where $\psi_{i,j}$ is a fixed integer. In terms of variable reduction, the best-case scenario in the set-based formulations compared to the MIP formulation would occur when $\psi_{i,j} = 1$ and $n_T = 120$. In that case, 121 binary variables in the traditional MIP formulation are replaced by a single interval variable in the MIDP and MILDP formulations. At solving time, Hexaly does not reformulate interval and list variables as binary variables, but applies powerful primal heuristics combined with exact algorithms to the set-based representation directly. Thus, this study shows that exploiting set-based formulations rather than relying on standard binary variable models enables the primal solver within Hexaly to scale better with large scheduling instances and get good quality solutions faster.

The results discussed through this and the previous section highlight instances where at least one of the proposed MIDP and/or MILDP formulations outperform a traditional MIP formulation when using Hexaly. Nevertheless, set-based formulations are not always expected to bring computational improvements. To illustrate this, we have performed a computational experiment that is a modified version of the current case study, with the only difference being that integer batching variables are now left unfixed, i.e., constraint (11) is removed from the MIP. In the MIDP formulation, this translates to using definition (12) instead of (17), constraint (13) instead of (18), and constraint (15) instead of (19). Similarly, the MILDP formulation needs to be modified by using definition (12) instead of (17), constraint (21) instead of (18), constraint (15) instead of (19), and removing constraint (23).

Assuming a 300-second time limit, the solving-time performance profiles of this modified case study are shown in Figure 5. Recall that, within 300 seconds, the MIDP formulation solves 80% of the test problems to optimality, compared to 65% solved by the MIP model in the original version of this case study (see Figure 3). In this modified version, the MIP performance-profile line lies above both the MIDP and MILDP lines in Figure 5, indicating that the MIP formulation is now outperforming the set-based formulations, e.g., the MIP formulation solves 55% of the tested instances to optimality within 40 seconds, compared to 25% solved to optimality by the MIDP formulation. These results highlight that STN optimal scheduling benefits from the proposed MIDP and MILDP formulations when dealing with problems with fixed integer batching variables; however, this outcome does not generalize to problems with unfixed integer batching variables for the case study under consideration.

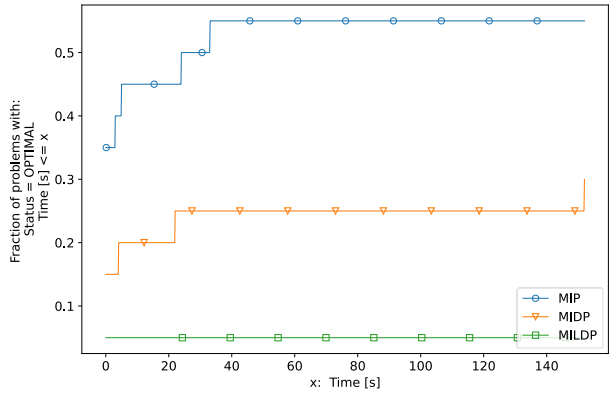

Figure 5. Solution time performance profile for the modified version of case 2.

CONCLUSIONS

The optimal STN scheduling problem has traditionally been addressed through standard MIP formulations and classical solution algorithms. In contrast, this work introduced the first set-based STN models, opening a new avenue for research in optimal network scheduling. Documentation suggests Hexaly performs well for specific classes of industrial-scale problems, such as scheduling and vehicle routing, precisely because it exploits this set-based modeling paradigm effectively. We have shown promising results for determining a strong set-based model for representing and solving STN formulations using Hexaly, namely a mixed-interval disjunctive programming (MIDP) formulation.

As future work, we aim to continue refining and studying the proposed set-based formulations. Also, a systematic comparison of set-based, constraint programming, and MIP optimization beyond Hexaly and across multiple modeling and solution platforms was outside the scope of this work. The solver embedded in Hexaly is known to be competitive with commercial and open-source solvers such as Gurobi, CPLEX, HiGHS, and OR Tools, e.g., in job shop scheduling [18]. Consequently, a broader benchmarking study with a larger set of test instances is necessary to further clarify how STN scheduling benefits from the proposed set-based formulations.

DIGITAL SUPPLEMENTARY MATERIAL

A file containing the parameters needed to define the case studies and to replicate the sensitivity analyses shown in the result section are available at https://psecommunity.org/LAPSE:2026.0003. All results can be replicated by running the code available at https://psecommunity.org/LAPSE:2026.0002.

ACKNOWLEDGEMENTS

The authors would like to thank the Hexaly team for

their valuable feedback on the manuscript.

David A. Liñán and Jan Kronqvist acknowledge the financial support provided by LKAB and Digital Futures.

AUTHOR IDENTIFIERS

Author ORCIDs:
Liñán DA: 0000-0002-3190-7612
Stinchfield G: 0009-0004-5130-5661
Kronqvist J: 0000-0003-0299-5745

REFERENCES

1. Agnetis A, Billaut JC, Pinedo M, Shabtay D. Fifty years of research in scheduling — theory and applications. European Journal of Operational Research 327:367-393 (2025). https://doi.org/10.1016/j.ejor.2025.01.034
2. Kondili E, Pantelides CC, Sargent RWH. A general algorithm for short-term scheduling of batch operations—i. MILP formulation. Computers & Chemical Engineering 17:211-227 (1993). https://doi.org/10.1016/0098-1354(93)80015-f
3. C.C. Pantelides. Unified frameworks for optimal process planning and scheduling. In: Proceedings of the Second Conference on Foundations of Computer Aided Operations. CACHE (1994)
4. Maravelias CT. Chemical production scheduling. Cambridge University Press (2021). https://doi.org/10.1017/9781316650998
5. Méndez CA, Cerdá J, Grossmann IE, Harjunkoski I, Fahl M. State-of-the-art review of optimization methods for short-term scheduling of batch processes. Computers & Chemical Engineering 30:913-946 (2006). https://doi.org/10.1016/j.compchemeng.2006.02.008
6. Samadi A, Maravelias CT. A comprehensive chemical production scheduling representation. Computers & Chemical Engineering 181:108552 (2024). https://doi.org/10.1016/j.compchemeng.2023.108552
7. Lekkerkerk K. Scheduling optimization in a refinery for vegetable oils. TU Delft (2020)
8. Moniz S, Barbosa-Póvoa AP, de Sousa JP. Simultaneous regular and non-regular production scheduling of multipurpose batch plants: a real chemical–pharmaceutical case study. Computers & Chemical Engineering 67:83-102 (2014). https://doi.org/10.1016/j.compchemeng.2014.03.017
9. Huang Y, Su R, Yang J, Li J, Du Z, Huang Z. An intelligent electricity scheduling model for cement plants based on STN. 2024 IEEE 2nd International Conference on Power Science and Technology (ICPST) :618-624 (2024). https://doi.org/10.1109/icpst61417.2024.10602163
10. Georgiadis GP, Elekidis AP, Georgiadis MC. Optimization-based scheduling for the process industries: from theory to real-life industrial applications. Processes 7:438 (2019). https://doi.org/10.3390/pr7070438
11. Perez HD, Amaran S, Iyer SS, Wassick JM, Grossmann IE. Applications of the RTN scheduling model in the chemical industry. Simulation and Optimization in Process Engineering :365-400 (2022). https://doi.org/10.1016/b978-0-323-85043-8.00006-4
12. Mostafaei H, Harjunkoski I. Continuous?time scheduling formulation for multipurpose batch plants. AIChE Journal 66: (2019). https://doi.org/10.1002/aic.16804
13. Castro PM, Grossmann IE. Generalized disjunctive programming as a systematic modeling framework to derive scheduling formulations. Ind. Eng. Chem. Res. 51:5781-5792 (2012). https://doi.org/10.1021/ie2030486
14. Blaise L, Artigues C, Benoist T. Solution repair by inequality network propagation in localsolver. Lecture Notes in Computer Science :332-345 (2020). https://doi.org/10.1007/978-3-030-58112-1_23
15. Grapendal L, Hirn C. Mathematical Optimization of Scheduling Problems: A State-Task Network and Mixed Integer Linear Programming Approach. KTH (2025)
16. Velez S, Maravelias CT. Reformulations and branching methods for mixed-integer programming chemical production scheduling models. Ind. Eng. Chem. Res. 52:3832-3841 (2013). https://doi.org/10.1021/ie303421h
17. Laborie P. Bornes rapides pour l'ordonnancement dans LocalSolver. ROADEF (2023)
18. Abreu LR. Mixed-integer versus constraint programming solvers. Proceeding Series of the Brazilian Society of Computational and Applied Mathematics : (2026). https://doi.org/10.5540/03.2026.012.01.0308

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Ammonia as Fuel for Gas Turbines – The Impact of Heat Integrated Partial Decomposition

Julian Straus[a]*, John C. Morud[b], and Elettra Vantaggiato[a]

[a] SINTEF Energi, Trondheim Norway
[b] SINTEF Industri, Trondheim, Norway
* Corresponding Author: julian.straus@sintef.no.

ABSTRACT

Ammonia has received in recent years significant attention as potential carbon free fuel. However, its combustion properties limit its direct application for both providing heat and in power generation through gas turbines. Ammonia cracking is one potential solution to circumvent the problem by producing hydrogen. When using the ammonia in gas turbines, it is possible to heat integrate the endothermic decomposition reaction with the exhaust gas from the gas turbine. Thermodynamic and kinetic limitations have however a major impact on the achievable ammonia conversion. Based on the consideration of these limitations, this paper presents a detailed investigation of key design parameters affecting the overall process efficiency utilizing both an equilibrium reactor model and a reactor model based on detailed kinetics and heat transfer. Ammonia decomposition should occur at sufficiently high pressure to avoid a) the com-pression energy demand for achieving the pressure of the combustion chamber and b) to reduce equipment size although the increased pressure results in a reduced conversion when considering an equilibrium reactor. It is crucial for the energy efficient integration of ammonia decomposition with gas turbines to avoid a partial combustion of the decomposed ammonia even as an increased temperature of the reactor results in an increased ammonia conversion. Furthermore, it is beneficial to operate the compression stages without intercooling to reduce the required fuel flow to the gas turbine.

Keywords: Alternative Fuels, Process Design, Modelling and Simulations, Surrogate Model

INTRODUCTION

The application of ammonia in the energy sector has received significant attention in recent years [1]. Ammonia can be considered as both a hydrogen carrier and a fuel in itself. The utilization of Ammonia as fuel results in no direct CO_2 emissions and has the potential to drastically reduce the emissions associated with energy use when produced from decarbonized hydrogen, *i.e.*, hydrogen produced from natural gas reforming with CO_2 capture or electrolysis [2]. In addition, it has in its liquid state the advantage of a relatively high volumetric energy density and established logistics chains. Hence, it can solve the problem of transporting hydrogen from regions with high renewable potential to regions with high energy demand [3]. Ammonia is also seen as direct fuel and not only as hydrogen carrier. It is considered as one alternative for dispatchable power generation [4] with a benefit at low capacity utilizations compared to other fuels when utilized in gas turbines [5]. Its potential application in gas turbines for electricity generation is however hindered by its poor combustion properties [6]. Consequently, decomposition of ammonia into hydrogen and nitrogen is required to maintain a stable flame [5]. Ammonia decomposition, that is the reverse of its synthesis, is an endothermic reaction given by

$$2NH_3 \rightleftharpoons N_2 + 3H_2 \tag{1}$$

High conversions can either be achieved at high temperatures or at a low pressure. Furthermore, the activity of ammonia decomposition catalysts at the low temperatures of gas turbine flue gas is low [7]. As a consequence, analyses normally focus on providing the heat for the decomposition reaction by burning a part of the (partially) decomposed) ammonia as fuel [5]. The decomposition of the ammonia leads however to an efficiency penalty, resulting in calls for the industry to develop novel

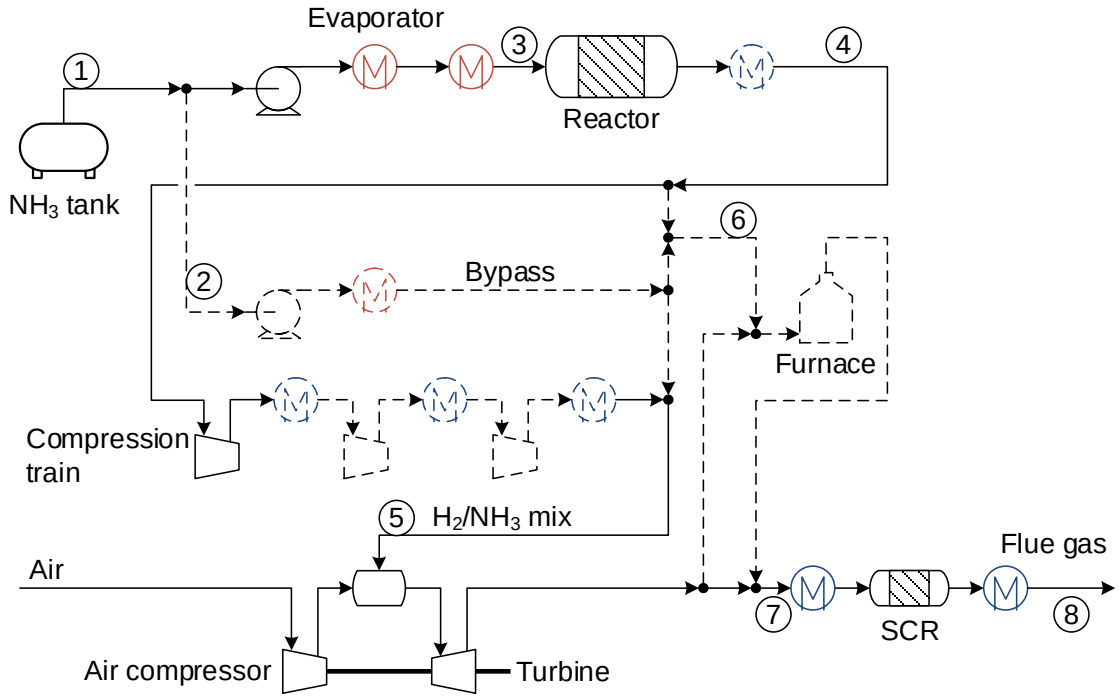

Figure 1. Process flow diagram of the proposed process. Blue heat exchangers correspond to coolers while red heat exchangers to heaters or boilers.

gas turbine burner types to be able to utilize pure ammonia. An alternative route for increasing the efficiency is to focus on partial decomposition of ammonia [8]. The partial decomposition has the advantage of a) having a reduced energy demand for the reaction, b) smaller reactor sizes as full conversion is not required, and c) the potential for heat integration with the flue gas of the gas turbine as less energy is required at lower temperatures.

The application of ammonia in gas turbines has however still open questions. While Cesaro et al. investigated as well the efficiency of partially cracked ammonia [5], their investigation did not consider the thermal integration of the gas turbine flue gas with the decomposition reactor. Instead, the decomposed ammonia (in this case a mixture of hydrogen and nitrogen) is used for providing the energy to the endothermic reaction while the flue gas is only utilized for preheating. Similarly, Paschchenko focused extensively on the operation of the gas turbine, but assumed an equilibrium reactor for the conversion of ammonia [8]. The investigation of the partial decomposition of ammonia has, however, still open questions with respect to the optimal process configuration due to the large number of design parameters.

Hence, we present in this study a thorough analysis of a novel integrated process for partial ammonia decomposition coupled with a gas turbine for exhaust heat recuperation. A focus of the work is to understand how the different design choices may affect the process energy efficiency and potential complexity. A detailed model of the decomposition reactor at low temperatures is included in the process flowsheet as surrogate model to account for the catalytic performance and potential limitations in the heat transfer. This allows the analyses of heat transfer limitations, and hence, both direct integration for heat recuperation and indirect through a heating fluid with a high heat capacity are investigated.

PROCESS DESCRIPTION

Modelling approach

The overall process is modelled in Aspen HYSYS and illustrated in Figure 1. All dashed components are optional depending on the chosen process configuration. Ammonia is stored (*NH_3 tank*) at refrigerated conditions (-33 °C and 1.1 bar pressure) to avoid explosive spillage in the case of a leak [9]. The liquid ammonia (*stream 1*) is first pumped before being evaporated (*Evaporator*) to avoid excessive compression requirements. The gaseous ammonia is then heated to the inlet temperature of the *Reactor*. In this study, two different reactor models are considered; initial analyses focusing on investigating the reactor pressure utilize an equilibrium reactor with an approach to equilibrium of 250 K while in-depth analyses are conducted utilizing a conversion reactor in which the conversion and outlet temperature of the different streams are calculated using a table lookup method given the inlet temperature to the reactor of ammonia and the heating medium (exhaust gas or a heating liquid). The detailed reactor model for creating the table is implemented

in MATLAB and models a shell and tube reactor with experimentally calculated kinetics for low temperature ammonia decomposition using a Co-Ce catalyst. These kinetic values are seen as optimistic as the used weight-hourly space velocity is below the experimental range, but potential trends in the analyses are considered to be independent of the exact kinetic values. After the decomposition reactor, the partially cracked ammonia is cooled (to *stream 4*) and compressed to the combustion chamber pressure at 30 bar (*Compression train*). Both the cooling after the reactor and interstage cooling between the compression stages can be deactivated to analyze their impact. The number of compressor stages is automatically calculated based on a maximum allowed compression ratio of 2.6. The process allows furthermore the inclusion of a bypass of the decomposition reactor (*stream 2*). The bypass corresponds to pumping the liquid ammonia directly to the combustion chamber pressure before evaporating it. The bypass can be seen as an additional degree of freedom for controlling the composition to the gas turbine.

Subsequently, the *Bypass* and the stream from the *Compression train* are mixed and fed to the combustion chamber of the gas turbine (*stream 5*). The mass flow of the *Air* to the combustion chamber is used for adjusting the temperature in the combustion chamber to 1160 °C. Electricity is subsequently generated using an expander model corresponding to a turbine. The exhaust gas contains significant amounts of NO_x which is removed by first adding ammonia (not shown in the Figure but modelled) to the exhaust gas and subsequent selective catalytic reduction (*SCR*) modelled as an equilibrium reactor. Heat recuperation from the exhaust gas is achieved both before and after the *SCR* reactor to allow for the required temperature in the *SCR* reactor.

As the temperature of the exhaust gas of the turbine may be too low to provide heat to the decomposition reaction, a potential *Furnace* is included. The fuel to the furnace (*stream 6*) is partially cracked ammonia at the same composition as the fuel to the combustion chamber while the air is a fraction of the exhaust gas. However, as the furnace is operating at a lower pressure, it utilizes the partially cracked ammonia before the compression to the pressure of the combustion chamber to avoid unnecessary energy demands. The flow to the furnace is controlled by pinch analysis and a specified minimum temperature approach ΔT_{min} of 25 °C. In the case of the detailed reactor model, we specify the required temperature of *stream 7* as input parameter while the furnace flow is calculated through a piecewise linear surrogate model. If the outlet temperature of the gas turbine exceeds the specified value, we use instead the outlet temperature directly. The hot liquid requires furthermore a ΔT_{min} of 12.5 °C for the heat transfer between the exhaust gas and the hot liquid. This assures the potential for designing an optimal heat exchanger network. Different approaches are chosen as the minimum temperature approach can occur at the reactor inlet in the case of the equilibrium reactor, while this is already included in the detailed reactor model.

Design parameters

The complete process has several important design parameters. These design parameters are explained in detail in the following sections.

Reactor pressure and temperatures

The reactor pressure has an impact on both the potential parasitic electricity requirements for the compression of the partially cracked ammonia and the equilibrium conversion of ammonia. Higher conversions can be achieved at low reactor pressures, but additional electricity for compression is required in this case. In addition, lower pressures result in larger equipment sizes. Similarly, higher reaction temperatures result in higher conversions, while they may require more combustion of partially cracked ammonia to achieve the temperatures as the temperature of the exhaust gas may be too low for providing the heat to the decomposition reaction.

The equilibrium reactor approach utilizes a constant temperature change in the reactor to approximate a counter current reactor while the table look-up method includes the temperature and relative flow of the heating medium as input parameters. The latter is implemented through the temperature of *stream 7* in Figure 1.

Cooling of the partially cracked ammonia and interstage cooling

The outlet stream from the ammonia decomposition reactor can be cooled to a desired temperature in *stream 4*. In addition, inter-stage cooling in the *compression train* can be applied or not. If the outlet stream is not cooled and no interstage cooling is applied, the efficiency of the compressors is reduced due to the increased temperature. The benefit of this approach is a simplified compressor design (albeit with potentially more expensive materials due to elevated temperatures) and a reduction in required fuel flow to the combustion chamber due to the increased temperature of the fuel.

Bypass of the decomposition reactor

The bypass of the decomposition reactor (*stream 2*) is an important factor for both the overall cracking ratio and the required compression energy demand. Introducing a bypass reduces the compression demand after the decomposition reactor and increases the conversion in the decomposition reactor due to a longer residence time for a given reactor size. In addition, it reduces the equipment size for compressors and heat exchangers However, care must be taken to avoid a significant reduction in overall cracking ratio to maintain a stable flame in the

combustion chamber. The bypass does not affect the heat required for the cracking ratio for a constant overall cracking ratio.

Key performance indicators

The overall cracking ratio X_P provides information regarding the conversion of ammonia. It has an impact on the flame in the combustion chamber as well as the potential NO_x formation of the process. It is calculated through the conversion in the decomposition reactor X_R and the bypass ratio, $BR = \dot{m}_2/\dot{m}_1$. $\dot{m}_1$ corresponds here to the complete ammonia feed, including the bypass and the required ammonia for the SCR reaction, while $\dot{m}_2$ represents the flow through the *Bypass*. The stream numbers are also included in Figure 1. It is given by

$$X_P = BR \times X_R \tag{2}$$

The overall lower heating value efficiency η_{LHV} can be calculated as

$$\eta_{LHV} = \frac{P_{turb} - P_{para}}{\dot{m}_1 \times LHV_1} \tag{3}$$

where P_{turb} corresponds to the electricity generation in the *Turbine* and P_{para} to parasitic electricity demand for pumping and compression. It is closely related to the energy efficiency of the gas turbine η_{GT}, calculated as

$$\eta_{GT} = \frac{P_{turb} - P_{comp}}{\dot{m}_5 \times LHV_5} \tag{4}$$

with P_{comp} representing the electricity demand for the *Air compressor* while the denominator is based on the inlet flow to the combustion chamber (*stream 5*). The gas turbine efficiency is higher than the overall efficiency of the process. It can provide an overview of the losses incurred by the partial ammonia decomposition.

The ratio of feed stream to the furnace stream, $FR = \dot{m}_6/\dot{m}_1$, is another important key performance indicator. While it is indirectly included in the overall efficiency η_{LHV}, it is still important to understand what fraction of ammonia must be combusted to achieve the desired decomposition reactor temperature.

ANALYSIS

The analysis is split into two distinctive sections, an analysis of the impact of the reactor pressure on the performance of the process through the application of the equilibrium reactor and a subsequent in-depth analysis utilizing the detailed reactor model and a table lookup method for additional parameters. In both cases, one-dimensional sensitivity analyses are conducted. The reference parameters are provided in Table 1. Note that the heating fluid temperature is only relevant for the table lookup method.

Table 1: Reference parameters for the analysis.

Parameter	Value	Unit
Reactor inlet temperature T_3	400	°C
Reactor inlet pressure p_3	10	bar
Cooling temperature T_4	23	°C
Bypass ratio BP	0	
Heating fluid temperature T_7	455	°C

Equilibrium reactor

The equilibrium reactor model is first utilized for analyzing the impact of the reactor pressure as the impact of the reactor pressure is significantly affected by the chosen reactor size in the detailed reactor model. Interstage cooling is utilized to avoid high temperatures after the compression train affecting the results significantly. Instead, it is assumed that the compressed gas after the compression section is cooled to 100 °C.

Figure 2 shows the impact of the reactor inlet pressure on the performance of the integrated process. Due to the chemical equilibrium of the decomposition reaction, low pressures are beneficial for achieving a high ammonia conversion; the conversion is reduced from 79.5 % at a reactor pressure of 2 bar to 32.9 % at a reactor pressure of 15 bar. However, due to the increased conversion at lower pressures, it is necessary to have a higher furnace bypass ratio to provide sufficient high temperature heat for the decomposition reaction resulting in a similar behavior of both KPIs as shown in Figure 2 a). In the case of a high conversion, the minimum temperature approach is at the reactor inlet corresponding to an increased ΔT at the reactor outlet. While the turbine efficiency is largely unaffected by the reactor inlet pressure, we see large increases in the process efficiency from 36.7 % to 43.0 % when increasing the reactor inlet pressure from 2 bar to 15 bar.

The increase is especially pronounced in the range of 2 bar to 10 bar where the latter already results in 42.6 %. The efficiency increase of 17 % can be explained by 1) the reduction in the required furnace bypass ratio due to the reduced head demand in the decomposition reaction and 2) the reduced compression energy demand. The latter is especially important as the compression energy requirement at a reactor pressure of 2 bar corresponds to 7 % of the lower heating value of the partially cracked ammonia. This electricity demand is reduced to 1.9 % at a pressure of 10 bar and 1.2 % at a pressure of 15 bar.

Given these initial analyses utilizing the simplified equilibrium reactor model, detailed analyses are conducted utilizing a reactor pressure of 10 bar. This pressure can be seen as a beneficial compromise between the required conversion of ammonia for the gas turbine and the process efficiency.

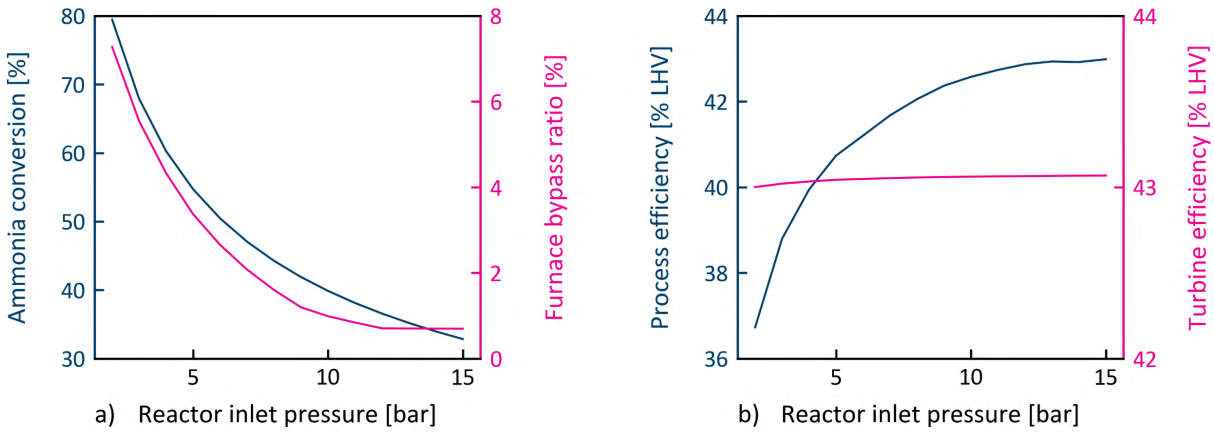

Figure 2. Results for the eqilibrium reactor model when varying the reactor inlet pressure. The results show a) (blue) the ammonia conversion, a) (pink) the furnace bypass ratio, b) (blue) the process efficiency, and b) (pink) the gas turbine efficiency.

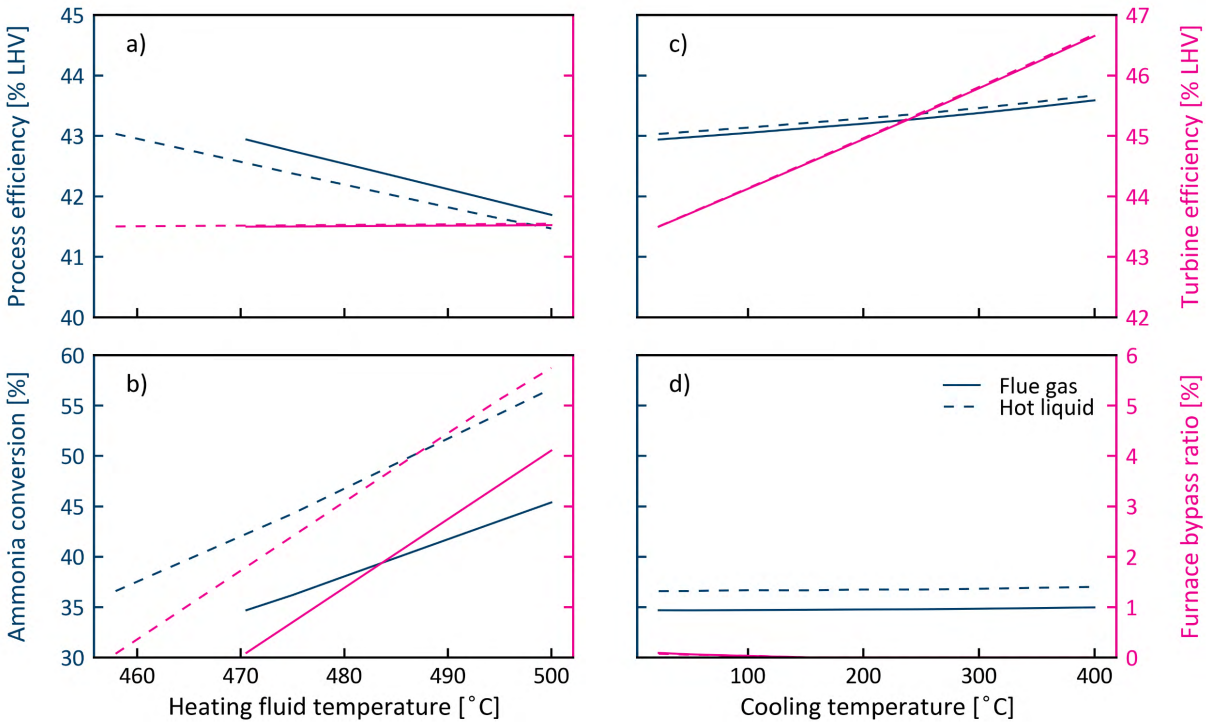

Figure 3. Results for the detailed reactor model when varying in a) and b) the heating fluid temperature or in c) and d) the cooling temperature when utilizing the exhaust gas (solid) or a hot liquid (dashed) as heat transfer medium. The results show a) and c) (blue) the process efficiency, a) and c) (pink) the gas turbine efficiency, b) and d) (blue) the ammonia conversion, and b) and d) (pink) the furnace bypass ratio.

Detailed reactor model

Figure 3 illustrates the results for varying the heating fluid temperature, a) and b), and the cooling temperature, c) and d) for both using the flue gas and the hot liquid for providing heat to the decomposition reactor. When investigating the impact of the heating fluid temperature, shown in Figure 3 a) and b), we see directly the difference of 12.5 °C for the minimum value used for the heating fluid temperature, representing the chosen ΔT_{min} when using the hot liquid. This difference leads to a shift of the complete analysis. Increasing the heating fluid temperature results in a linear decrease in the process efficiency while the turbine efficiency increases slightly. The linear decrease of the process efficiency is nearly independent of the chose heating fluid when considering

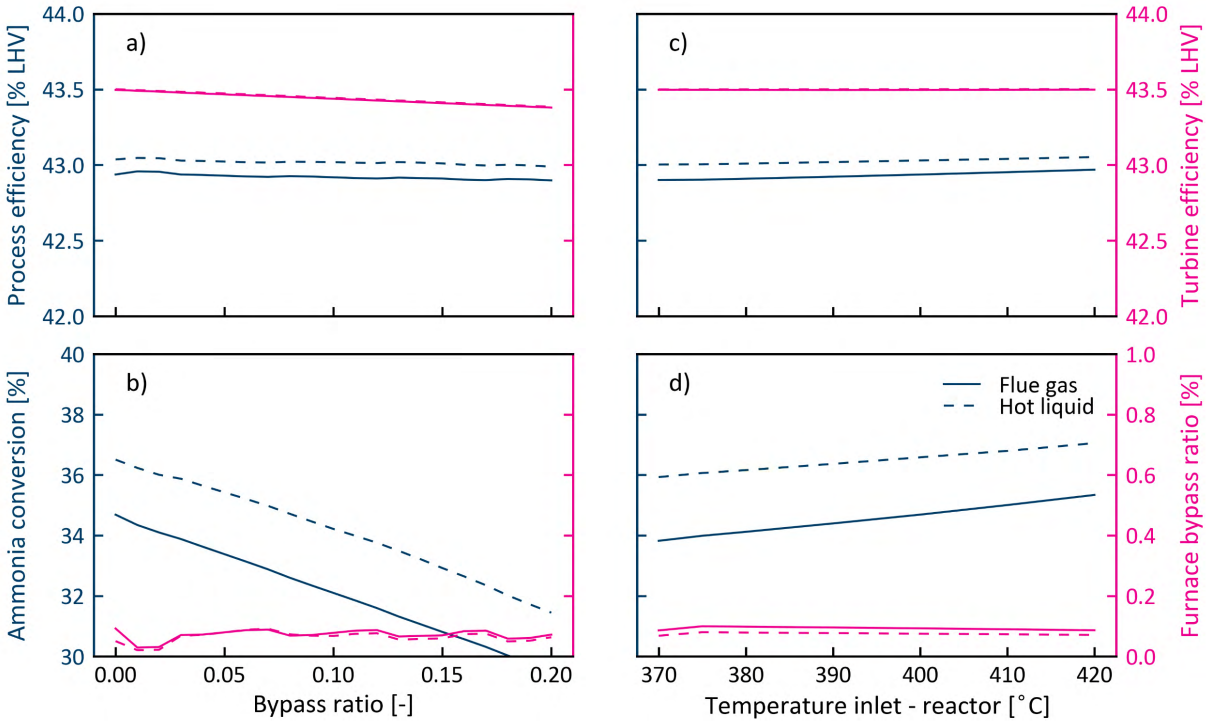

Figure 4. Results for the detailed reactor model when varying in a) and b) the heating bypass ratio or in c) and d) the reactor inlet temperature of the NH_3 feed when utilizing the exhaust gas (solid) or a hot liquid (dashed) as heat transfer medium. The results show a) and c) (blue) the process efficiency, a) and c) (pink) the gas turbine efficiency, b) and d) (blue) the ammonia conversion, and b) and d) (pink) the furnace bypass ratio.
Note the change in y-axis ticks compared to Figure 3.

the 12.5 °C temperature difference. When increasing the heating fluid temperature, the temperature of the exhaust gas from the gas turbine is not sufficient for providing the heat to the decomposition reactor. It is hence necessary to combust part of the partially cracked ammonia to reach the required temperature level. The observed increase in the cracking ratio is furthermore different between utilizing the exhaust gas directly or utilizing a hot liquid due to the differing heat transfer properties with a steeper slope when using the hot liquid.

The impact of the cooling temperature on the efficiencies is more pronounced as shown in Figure 3 c). Increasing the cooling temperature, *i.e.*, cooling less after the reactor, has a profound effect on both the turbine efficiency and the process efficiency, although the effect on the latter is smaller. The increase in the gas turbine efficiency can be explained by the chosen boundary conditions. As the temperature of the feed increases with increasing cooling temperature, less fuel must be provided to the combustion chamber for achieving the desired temperature level after the combustion chamber. The benefit for the process efficiency is, however, reduced as the increased temperature in the compressor results in an increased electricity demand of the compressor by 120 % when increasing the temperature before the compression from 23 °C to 400 °C. There is however still a significant net benefit of not cooling the partially cracked ammonia. On the other hand, the cooling temperature has next to no effect on both the ammonia conversion and the furnace bypass ratio as the reactor conditions remain unchanged, and hence, the conversion and required additional heat through the furnace bypass.

Figure 4 shows the results when varying the bypass ratio, a) and b) and the inlet temperature to the reactor, c) and d). The reactor bypass ratio has only a minor direct impact on both the gas turbine and the process efficiency as shown in Figure 4 a). Increasing the bypass ratio reduces both efficiencies, but the reduction is only 0.04 percentage points for the process efficiency and 0.1 percentage for the gas turbine efficiency. Both numbers are considered to be within the margin of error in the table lookup method. The slight change in the efficiency is on the one hand explainable by the reduced ammonia conversion as shown in Figure 4 b). Due to the bypass, we reduce the overall ammonia conversion. However, the reduction from 34.7% to 29.5 % in the case of using the exhaust gas directly and from 36.5 % to 31.4 % in the case of using an intermediate hot liquid does not reflect the actual loss as the conversion within the reactor is increased. If the conversion in the reactor remains the same, one would expect a conversion of 27.7 and 29.2 %, respectively. Hence, the conversion in the reactor is

increased if the mass flow of the heating medium is increased compared to the mass flow of ammonia. As a consequence, using a bypass would allow for a smaller reactor size at the cost of a slightly reduced process efficiency. The furnace bypass ratio is in practice independent of the chosen bypass of the decomposition reactor. The observed non-linearities can be attributed to steps within the table lookup method resulting in less accurate simulations. The observed value of 0.1 % is also rather small.

Similar to the reactor bypass ratio, we do not observe a significant impact of the reactor inlet temperature on the efficiencies as shown in Figure 4 c). While the process efficiency experiences a small increase with a higher reactor inlet temperature, the turbine efficiency remains unaffected by this change despite the increased ammonia conversion as shown in Figure 4 d). The furnace bypass ratio is again small and unaffected by the changes, similar to the results shown in Figure 4 b) for the cooling temperature. The result can be explained by the chosen size for the reactor model. The reactor outlet temperature of the partially cracked ammonia is for the chosen size nearly independent of the reactor inlet temperature within the given range

DISCUSSIONS AND LIMITATIONS

Impact of the reactor model approach

On the one hand, the chosen reactor model and its representation has a significant impact on the identification of key operating parameters in the analysis. Specifically, utilizing an equilibrium reactor requires assumptions regarding the approach to equilibrium and the heat transfer within the reactor to approximate the temperature change ΔT in the reactor. Neither parameters can be directly deduced from kinetic studies. Especially the assumed ΔT in combination with the inlet temperature has a major effect as the equilibrium is calculated at the reactor outlet conditions. In addition, when analyzing the reactor bypass, we cannot observe the increased conversion in the reactor when using an equilibrium reactor.

On the other hand, it is not possible to utilize a kinetic reactor when designing a process due to the requirement of specifying a size for the reactor. As can be seen in the results for varying the inlet temperature to the reactor, the assumption of a constant size provides a suboptimal configuration in some of the values for the inlet temperature.

One improvement would hence be to operate iteratively between the detailed reactor model and the process model to adjust the size of the reactor model based on the analyses from the process model or alternatively use a complete surrogate model of the detail reactor model.

Cooling before compression

As shown in the analyses, with a required combustion chamber pressure of 30 bar, cooling before the compression chamber is disadvantageous. Although compression at high temperatures results in an increased electricity demand, it is compensated by the reduced feed flow rate for a constant electricity generation in the turbine. However, the feed gas to the combustion chamber has without cooling a temperature of 675 °C resulting in potential problems when compressing a mixture consisting of ammonia and hydrogen. When assuming cooling of the partial cracked ammonia post the reactor to 23 °C, the temperature is reduced to 175 °C without interstage cooling. While temperatures of 600 °C are observed for air compression, the corrosive nature of ammonia and potential hydrogen embrittlement can lead to problems. Consequently, it is important to consider the practical limits for the compression section.

Inclusion of a steam cycle for combined cycle operation

The current model does not include a steam cycle utilizing the still available exhaust heat for increased electricity generation. Due to the utilization of parts of the heat for the decomposition reaction, less heat is available for the steam cycle. Specifically, we observe a reduction in the temperature of the exhaust gas stream available for steam generation. Consequently, electricity generation in a potential steam cycle will be reduced compared to the standard process without ammonia decomposition. When not cooling the reactor before the compression steps, ammonia is heated using the exhaust flue gas from the reactor. This is especially a problem when choosing high reactor inlet temperatures for ammonia as the temperature exceeds in this case the exhaust of the heating medium from the reactor. This implies that the optimal process parameters may differ when including a steam cycle for an increased efficiency.

Dynamic flexibility of the process and operation at part-load

The overall process requires a constant heat supply to the decomposition reactor for maintaining a constant cracking ratio to the combustion chamber. As this is not necessarily feasible, introducing the bypass will allow a tight control of the composition to the combustion chamber. In this situation, operational flexibility can be achieved through sacrificing efficiency.

Part-load operation of gas turbines is in general achieved through a reduced pressure in the combustion chamber. As a consequence, the temperature of the exhaust gas is reduced. Hence, more partially cracked ammonia must be combusted to compensate for the reduced temperature of the exhaust gas. This implies that the efficiency penalty of the gas turbine at reduced load

is furthermore exacerbated by the need for additional combustion of ammonia for achieving the required temperature in the decomposition reactor.

CONCLUSION AND FUTURE WORK

In this work, we provide an in-depth analysis of a coupled process for electricity generation from ammonia. The process, in which the ammonia decomposition reactor is heat-integrated with the gas turbine exhaust gas, allows for a high efficiency in the operation of the gas turbine. Lowering the required ammonia conversion results in improved overall performance due to a reduction in the required high temperature heat, and hence, the removal of combusting a fraction of the partially cracked ammonia for providing the heat to the decomposition reactor. Similarly, higher operating pressures for the reaction are beneficial to avoid the compression of the partially decomposed ammonia.

Future work will focus on part-load operation and analyzing the dynamic behavior of the complete process to propose potential control strategies given the increasing share of variable renewable power generation and the demand for adjustable dispatchable generators.

ACKNOWLEDGEMENTS

This publication is based on results from the research project **DECAMMP– Decomposed ammonia for carbon-free power generation**, performed under the PETROMAKS2 programme. The authors acknowledge the following parties for financial support: Research Council of Norway (336549), Johnson Matthey and the following partners of the LowEmission Research Centre: TotalEnergies, Equinor, Vår Energi, Aker BP, Repsol Norge, Wintershall DEA Norge, DNO Norge, OMV Norge, Sval Energi and Siemens Energy

AUTHOR IDENTIFIERS

Author ORCIDs:
Straus J: 0000-0001-8622-1936
Morud JC: 0000-0003-1098-0404.
Vantaggiato E: 0009-0005-7586-1168

REFERENCES

1. Valera-Medina A, Amer-Hatem F, Azad AK, Dedoussi IC, de Joannon M, Fernandes RX, Glarborg P, Hashemi H, He X, Mashruk S, McGowan J, Mounaim-Rouselle C, Ortiz-Prado A, Ortiz-Valera A, Rossetti I, Shu B, Yehia M, Xiao H, Costa M. Review on ammonia as a potential fuel: from synthesis to economics. Energy Fuels 35:6964-7029 (2021). https://doi.org/10.1021/acs.energyfuels.0c03685
2. Arnaiz del Pozo C, Cloete S. Techno-economic assessment of blue and green ammonia as energy carriers in a low-carbon future. Energy Conversion and Management 255:115312 (2022). https://doi.org/10.1016/j.enconman.2022.115312
3. Ishimoto Y, Voldsund M, Nekså P, Roussanaly S, Berstad D, Gardarsdottir SO. Large-scale production and transport of hydrogen from norway to europe and japan: value chain analysis and comparison of liquid hydrogen and ammonia as energy carriers. International Journal of Hydrogen Energy 45:32865-32883 (2020). https://doi.org/10.1016/j.ijhydene.2020.09.017
4. Valera-Medina A, Xiao H, Owen-Jones M, David WIF, Bowen PJ. Ammonia for power. Progress in Energy and Combustion Science 69:63-102 (2018). https://doi.org/10.1016/j.pecs.2018.07.001
5. Cesaro Z, Ives M, Nayak-Luke R, Mason M, Bañares-Alcántara R. Ammonia to power: forecasting the levelized cost of electricity from green ammonia in large-scale power plants. Applied Energy 282:116009 (2021). https://doi.org/10.1016/j.apenergy.2020.116009
6. Herbinet O, Bartocci P, Grinberg Dana A. On the use of ammonia as a fuel – a perspective. Fuel Communications 11:100064 (2022). https://doi.org/10.1016/j.jfueco.2022.100064
7. Lamb KE, Dolan MD, Kennedy DF. Ammonia for hydrogen storage; a review of catalytic ammonia decomposition and hydrogen separation and purification. International Journal of Hydrogen Energy 44:3580-3593 (2019). https://doi.org/10.1016/j.ijhydene.2018.12.024
8. Pashchenko DI. Ammonia as a fuel for gas-turbine units with thermochemical recuperation of exhaust gas heat. Therm. Eng. 72:1-7 (2025). https://doi.org/10.1134/s0040601524700691

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Research on Dynamic Scheduling of Multi-line Polyolefin Production Based on Deep Reinforcement Learning

Zhineng Tao[a], Tong Qiu[a]*, Zhenzhi Gong[ab], Fenglian Dong[c], Zhiwei Wei[c], Yunlong Guan[b]

[a] Department of Chemical Engineering, Tsinghua University, Beijing 100084, China;
[b] PetroChina Dushanzi Petrochemical Company, Xinjiang 833699, China;
[c] Petrochina Planning And Engineering Institute, Beijing 100083, China;
* Corresponding Author: qiutong@tsinghua.edu.cn

ABSTRACT

The scheduling of multi-line polyolefin production is a complex decision-making process characterized by sequence-dependent changeovers, strict physicochemical constraints, and dynamic market environments. Traditional optimization methods often suffer from high computational costs and a lack of flexibility in online adjustments. To address these challenges, this paper proposes a Deep Reinforcement Learning (DRL) framework for dynamic scheduling tasks. We first construct a high-fidelity simulation environment that meticulously models realistic industrial constraints, including transition materials, shutdowns, and inventory limits. A Soft Actor-Critic (SAC) agent with a tuple-based action space is employed to mitigate the combinatorial explosion associated with multi-line decisions. Furthermore, a dynamic action masking mechanism embedded with domain knowledge is introduced to strictly enforce hard constraints and significantly improve sample efficiency. Case studies based on real-world industrial data demonstrate that the proposed method can autonomously generate valid schedules that satisfy complex production requirements. Comparative experiments further reveal that the action masking mechanism accelerates training convergence, and the DRL agent exhibits superior adaptability to dynamic price fluctuations.

Keywords: Scheduling, Modelling and Simulations, Optimization, Reinforcement learning, Polyolefin production

1 INTRODUCTION

Polyolefins are the most widely produced chemical materials globally. Consequently, the efficiency of polyolefin production scheduling directly impacts economic performance and market responsiveness [1]. This process, characterized by continuous operations and frequent grade transitions, faces several challenges, including complex transition processes, sequence-dependent changeover costs, and strict process constraints. Although traditional methods, such as mathematical programming and heuristic algorithms, have achieved some success, they encounter significant limitations in practical applications. The primary issues are high computational costs and a lack of flexibility [2]. Fundamentally, these methods rely on static scheduling, which involves offline optimization based on fixed initial states and demands over a specific time horizon. When production disturbances occur or boundary conditions change, the entire schedule must be recalculated from scratch. This process is computationally expensive and time-consuming, resulting in response delays that fail to meet the requirements of modern chemical manufacturing for efficient and rapid decision-making [3].

Reinforcement Learning (RL) presents a promising approach for solving such complex sequential decision-making problems. Previous studies have demonstrated the potential of RL in industrial scheduling [4, 5]. In this framework, the scheduler is modeled as an agent that learns by observing states from the environment, executing actions, and receiving reward signals that reflect economic benefits or operational efficiency. Unlike traditional static methods that solve specific scenarios independently, RL trains a generalized end-to-end policy by interacting with a simulation environment. This policy maps high-dimensional production states directly to near-optimal scheduling actions. Leveraging the "offline training, online execution" mechanism, the trained agent can generate schedules via simple forward inference when facing different horizons or initial states. This data-

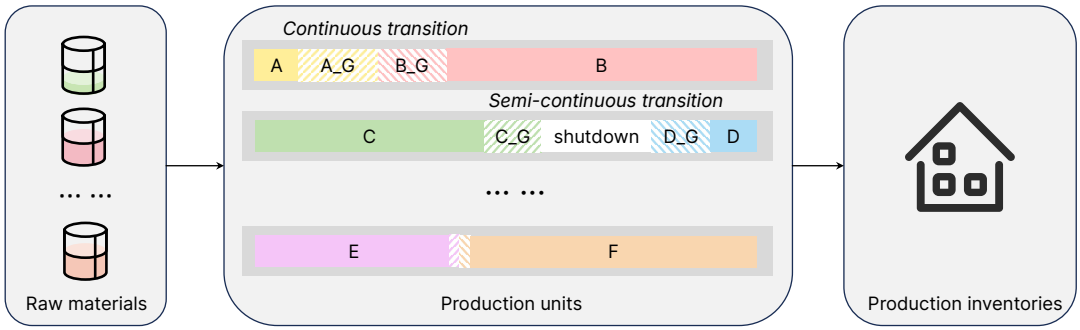

Figure 1: Schematic of the polyolefin production process.

driven approach avoids computationally expensive re-optimization for every new scenario, thereby demonstrating a significant advantage in computational efficiency over traditional methods.

Motivated by these advantages, this study constructs a high-fidelity simulation environment and proposes a Deep Reinforcement Learning (DRL) framework for multi-line polyolefin scheduling in real-world industrial scenarios. The main contributions are summarized as follows:

1. Design of a tuple-based action space with a multi-head architecture: To tackle the combinatorial explosion inherent in multi-line scheduling, we design a tuple-based action structure implemented via a multi-head neural network. This approach effectively decomposes high-dimensional decision variables, enabling precise and collaborative control over parallel production units.

2. Implementation of a domain-knowledge-embedded dynamic action masking mechanism: We introduce a real-time masking mechanism that dynamically adjusts the feasible action space based on process states. By rigorously filtering out actions that violate physicochemical constraints, this mechanism ensures operational validity and significantly accelerates training convergence.

3. Introduction of a dual-level stochastic mechanism for robust generalization: We incorporate episode-level domain randomization and step-level market price volatility into the environment. This dual uncertainty prevents overfitting to deterministic conditions and compels the agent to develop "market acumen, " thereby maximizing economic benefits in non-stationary environments.

The remainder of this paper is organized as follows: Section 2 details the methodology, including the mathematical problem description and the construction of the simulation environment. Section 3 presents the case study and analyzes the experimental results. Finally, Section 4 concludes the paper.

2 METHOD

2.1 Problem Description

The polyolefin production system consists of multiple parallel production units, as illustrated in Figure 1. These units share upstream raw material inventories and produce various downstream polyolefin product grades. This study explicitly models detailed grade transition processes, accounting for both transition material generation and unit shutdowns. As depicted in the first unit of Figure 1, the transition from prime grade A to prime grade B involves producing transition materials for both grades (denoted as A_G and B_G) without interrupting production. This mode is defined as a continuous transition. In contrast, the second unit illustrates a transition from grade C to grade D, which requires a shutdown period in addition to generating transition materials. This mode, termed a semi-continuous transition, typically occurs when significant physicochemical differences between products necessitate equipment cleaning to remove catalyst residues.

Consequently, the polyolefin production scheduling problem is characterized as a complex decision-making process involving multi-line, multi-product, and mixed continuous-discrete dynamics. The core objective is to maximize the total economic benefit over a finite horizon by collaboratively optimizing product sequencing, production rates, and transition operations across parallel units, subject to market demands, raw material availability, and strict process constraints.

2.2 Mathematical Formulation

To strictly define the scheduling problem, the aforementioned polyolefin production system is mathematically formulated as follows:

Given Parameters:

1. Basic Parameters: Consider a scheduling horizon of T days ($t \in \{1, 2, \dots, T\}$). There are U production units ($u \in \{1, 2, \dots, U\}$) capable of consuming MR types of raw materials ($mr \in \{1, 2, \dots, MR\}$) to produce MP different product grades ($mp \in \{1, 2, \dots, MP\}$). Each prime product grade is associated with a specific transition material ($mp_G \in \{1_G, 2_G, \dots, MP_G\}$). The set of products compatible with unit u is denoted as $linkum_{u,mp}$, and the set of feasible transitions between products is denoted as $linkmm_{mp,mp'}$.

2. let $Recipe_{mp,mr}$ denote the raw material consumption ratio, $Pcost_{mp}$ the processing cost, and $Price_{mp}$ the selling price. The transition material shares the same recipe ($Recipe_{mp,mr}$) and processing cost ($Pcost_{mp_G}$) as its corresponding prime grade, but commands a lower selling price ($Price_{mp_G}$). The unit cost for each raw material is denoted as $Rcost_{mr}$.

3. Production Process Parameters: During a grade changeover, the quantity of transition material and the necessity of a shutdown are determined by both the preceding and succeeding products. For a transition from product mp to mp', the transition attributes $Changeover_{mp,mp'}$ are defined by a 6-tuple $(T_{BG}, Q_{BG}, T_{AG}, Q_{AG}, T_{stop}, Cost_{stop})$. These elements represent the duration and quantity of the preceding transition material, the duration and quantity of the succeeding transition material, the shutdown duration, and the shutdown cost, respectively.

Operating Rules and Constraints:

1. Product Demand Limits: By the end of the scheduling horizon, the accumulated production yield of each product must fall within the specified demand range $[Demand_{mp}^{min}, Demand_{mp}^{max}]$.

2. Production Rate Limits: The production rate for each product must operate within the allowable range $[Rate_{mp}^{min}, Rate_{mp}^{max}]$.

3. Raw Material Inventory Limits: The inventory levels of specific raw materials must remain within the safety limits $[Inv_{mr}^{min}, Inv_{mr}^{max}]$ throughout the entire horizon. The daily external supply of raw material is denoted as $S_{mr,t}$, and the daily external consumption is $EC_{mr,t}$.

4. Workforce and Timing Constraints: Due to maintenance crew limitations, simultaneous changeovers on multiple production units are prohibited. Furthermore, changeovers are restricted to specific time windows on designated working days Fd_d.

Decision Variables:

$ch_{mp,mp',u,t}$: Binary variable indicating whether a changeover from product mp to mp' starts on unit u on day t.

$t_{mp,u,t}$: Production duration of product mp on unit u on day t.

$r_{mp,u,t}$: Production rate of product mp on unit u on day t.

Objective Function:

The objective function is formulated to maximize the total profit, which comprises seven distinct components: sales revenue from prime products, sales revenue from transition materials, raw material costs for prime products, processing costs for prime products, raw material costs for transition materials, processing costs for transition materials, and shutdown costs associated with semi-continuous transitions.

The objective function is expressed as:

$$\max Profit = Revenue_{mp} + Revenue_{mp_G} - Cost_{mp,raw} - Cost_{mp,prod} - Cost_{mp_G,raw} - Cost_{mp_G,prod} - Cost_{trans} \quad (1)$$

These components are calculated using Eqs. (2)–(8) as follows:

$$Revenue_{mp} = \sum_t \sum_u \sum_{mp} r_{mp,u,t} \cdot t_{mp,u,t} \cdot Price_{mp} \quad (2)$$

$$Revenue_{mp_G} = \sum_t \sum_u ch_{mp,mp',u,t}(Q_{BG} \cdot Price_{mp_G} + Q_{AG} \cdot Price_{mp'_G}) \quad (3)$$

$$Cost_{mp,raw} = \sum_t \sum_u \sum_{mp} \sum_{mr} r_{mp,u,t} \cdot t_{mp,u,t} \cdot Recipe_{mp,mr} \cdot Rcost_{mr} \quad (4)$$

$$Cost_{mp,prod} = \sum_t \sum_u \sum_{mp} r_{mp,u,t} \cdot t_{mp,u,t} \cdot Pcost_{mp} \quad (5)$$

$$Cost_{mp_G,raw} = \sum_t \sum_u \sum_{mr} ch_{mp,mp',u,t}(Q_{BG} \cdot Recipe_{mp_G,mr} + Q_{AG} \cdot Recipe_{mp'_G,mr}) \cdot Rcost_{mr} \quad (6)$$

$$Cost_{mp_G,prod} = \sum_t \sum_u ch_{mp,mp',u,t}(Q_{BG} \cdot Pcost_{mp_G} + Q_{AG} \cdot Pcost_{mp'_G}) \quad (7)$$

$$Cost_{trans} = \sum_t \sum_u ch_{mp,mp',u,t} \cdot Cost_{stop} \quad (8)$$

2.3 Proposed DRL Framework for Dynamic Scheduling

To address the complexity of multi-line polyolefin scheduling, we formulate the problem as a Markov Decision Process (MDP) [6]. The overall framework of the proposed method is illustrated in Figure 2. This section first details the mathematical fundamentals of the MDP and then introduces three key mechanisms designed to handle combinatorial explosion, strict constraints, and non-stationary processes: a multi-head tuple action space strategy, a domain-knowledge-embedded dynamic masking mechanism, and a dual-level stochastic mechanism for generalization.

2.3.1 State Space and Reward Function

The state space $s_t \in \mathcal{S}$ serves as the foundational input for the agent's decision-making, providing a comprehensive snapshot of the production system. To accommodate the multi-line characteristics of polyolefin manufacturing, the state vector is structured hierarchically, comprising global states shared across all units and local states specific to each individual unit, as defined in Eqs. (9)–(11):

$$s_t^{global} = [\tau_t, inv_{mr,t}, \rho_t] \quad (9)$$

$$s_{u,t}^{unit} = [\delta_{u,t}, \nu_{u,t}, \Pi_{u,t}] \quad (10)$$

$$s_t = \left[s_t^{global}, s_{1,t}^{unit}, s_{2,t}^{unit}, \dots, s_{U,t}^{unit}\right] \quad (11)$$

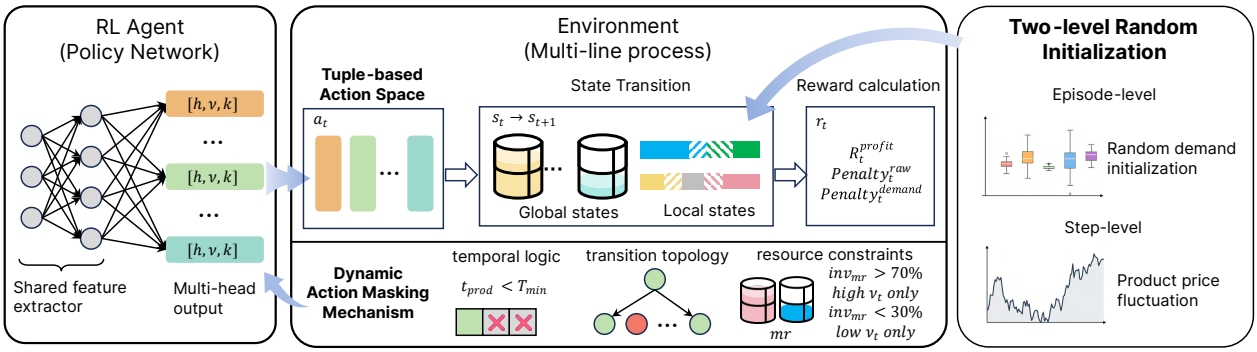

Figure 2. Overview of the proposed Deep Reinforcement Learning framework for multi-line polyolefin scheduling.

Specifically, The global state s_t^{global} encompasses the normalized time progress τ_t, the raw material inventory levels $inv_{mr,t}$, and the product market price ratios ρ_t. The local state for unit u, denoted as $s_{u,t}^{unit}$, captures the detailed production status of the u-th unit. This comprises the production duration ratio relative to the minimum requirement T_{min} (denoted as $\delta_{u,t}$), the normalized current production rate $\nu_{u,t}$, and a binary indicator $\Pi_{u,t}$ signaling whether the unit is currently in a transition state.

To guide the agent towards optimal policies, the reward function $R(s_t, a_t)$ aims to achieve a multi-objective trade-off between economic benefits and the risk of constraint violations. It is defined as follows:

$$R_t = \omega_1 R_t^{profit} - \omega_2 Penalty_t^{raw} - \omega_3 Penalty_t^{demand} \quad (12)$$

where $R_t^{profit}, Penalty_t^{raw}, Penalty_t^{demand}$ represent the daily net profit from production, the penalty for raw material inventory overflow or shortage, and the penalty for violating product demand limits, respectively. ω_1, ω_2, and ω_3 denote the weight hyperparameters. The calculation of R_t^{profit} follows Eq. (1), while the other two terms are calculated as follows:

$$Penalty_t^{raw} = \sum_{mr}\left(\max(0, inv_{mr,t} - Inv_{mr}^{max})^{1.5} + \max(0, Inv_{mr}^{min} - inv_{mr,t})^{1.5}\right) \quad (13)$$

$$Penalty_t^{demand} = \sum_{mp}\left(\max(0, inv_{mp,t} - Demand_{mp,t}^{max})^{0.5} + \max(0, Demand_{mp,t}^{min} - inv_{mp,t})^{2}\right) \quad (14)$$

Different degrees of non-linear terms are introduced in the above calculations to reflect the varying levels of tolerance for different boundary violations.

2.3.2 Tuple-based Action Space

Direct application of DRL to multi-line scheduling faces a critical challenge: the curse of dimensionality. In a system with U parallel units, a naive flattening of decision variables results in an action space size that grows exponentially, rendering exploration inefficient and convergence nearly impossible. To overcome this combinatorial explosion, we propose a decomposed tuple-based action formulation implemented via a multi-head neural network architecture.

Instead of treating the global action as a monolithic high-dimensional vector, we structure the action at time t as a joint tuple containing distinct decision components for each unit, as defined in Eq. (15):

$$a_t = \{[h_{1,t}, \nu_{1,t}, k_{1,t}], [h_{2,t}, \nu_{2,t}, k_{2,t}], \dots, [h_{u,t}, \nu_{u,t}, k_{u,t}]\} \quad (15)$$

This formulation decomposes the decision process for each unit u into three strategically discretized components (illustrated in Figure 3):

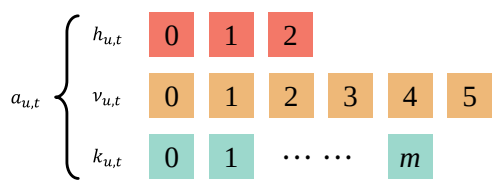

Figure 3. Architecture of the tuple-based action space for a single production unit.

Specifically, $h_{u,t}$ determines the transition timing on unit u. It takes values from the set $\{0, 1, 2\}$, corresponding to no transition, a transition at 11:00 AM, and a transition at 4:00 PM, respectively. $\nu_{u,t}$ represents the discretized production rate level for the subsequent product. It comprises six integer values (0-5), mapping to the 0th, 20th, 40th, 60th, 80th, and 100th percentiles of the allowable production rate range. Finally, $k_{u,t}$ identifies the target product grade to be produced following a transition. Its domain consists of feasible options corresponding to the specific products compatible with unit u, as defined by the set $linkum_{u,mp}$.

To realize this tuple-based control, we design a specific multi-head neural network. The network utilizes a shared feature extractor to capture global states, which then branches into U independent output heads. Each head autonomously generates the probability

distributions for the sub-actions $\langle h, v, k \rangle$ of a specific unit. This architecture effectively linearizes the decision complexity from exponential to linear relative to the number of units, significantly enhancing the scalability of the algorithm for large-scale industrial scheduling .

2.3.3 Dynamic Action Masking Mechanism

To address the hard constraints inherent in industrial production and enhance the sample efficiency of reinforcement learning, this study introduces a dynamic action masking mechanism. Prior to the generation of the action probability distribution, this mechanism evaluates the validity of actions in real-time based on the current state. By zeroing out the probabilities of invalid actions, it dynamically constructs a subspace of feasible actions.

Specifically, the action mask addresses three dimensions: temporal logic, transition topology, and resource constraints.

First is the temporal logic mask, which strictly enforces minimum production duration constraints. If the continuous production time of unit u for the current grade has not reached the minimum requirement, the mask restricts the transition timing action to 0. Simultaneously, the product selection mask permits only the current grade, effectively precluding premature interruptions or transitions, thus ensuring process continuity and stability.

Second is the transition mask based on topology and demand. For product selection, the mask activates only those target grades with direct physical connections to the current product, as defined by the transition set $linkmm_{mp,mp'}$, thereby preventing illegal cross-series transitions. Furthermore, an inventory-driven intervention logic is introduced: when the global inventory of the current product approaches its maximum limit (e.g., reaching 85%), the mask automatically disables the "maintain current product" option. This compels the agent to execute a transition, fundamentally mitigating the risk of inventory overflow.

Finally, the resource constraint mask governs material balance by mapping raw material inventory states to the feasible space of production rates. By monitoring inventory levels in real-time, this mechanism dynamically truncates the action space when critical thresholds are breached. Specifically, to prevent stockouts, high-production rate actions are masked when raw material inventory is critically low; conversely, to mitigate overflow risks, low-rate options are disabled when inventory approaches its upper limit, thereby compelling the unit to adjust its consumption rate to maintain stability.

Through this multi-dimensional masking mechanism, complex physical and business constraints are internalized as inherent rules of the environment. This eliminates futile exploration of invalid actions during the early stages of training, significantly accelerating policy convergence.

2.3.4 Stochastic Dynamics and Generalization

A core challenge in deploying DRL to industrial scheduling is the "reality gap" and the propensity for agents to overfit to fixed training scenarios. To cultivate an agent capable of mastering generalized scheduling logic rather than merely memorizing specific case data, we introduce a dual-level stochastic mechanism comprising episode-level domain randomization and step-level market dynamics.

At the episode level, to prevent reliance on deterministic initial conditions, the environment state is rigorously randomized upon initialization. For each unit, the initial product grade is sampled uniformly from the feasible set, while the initial production rate and raw material inventories are randomly instantiated within the specific operational bounds defined in Section 2.2 . Concurrently, to simulate diverse market pressures, the upper limit of product demand undergoes a random drift around its base value, as formalized in Eqs. (16) and (17) :

$$Demand'_{max} = \max(Demand_{min}, Demand_{max} + \Delta_{demand}) \tag{16}$$

$$\Delta_{demand} \sim U\left(-D_{factor}, +D_{factor} \times (Demand_{max} - Demand_{min})\right) \tag{17}$$

where D_{factor} represents the demand fluctuation factor. This high-variance initial distribution forces the agent to rely on the current state s_t for decision-making rather than memorizing historical trajectories, thereby enhancing the model's generalization capability.

Complementing this initialization strategy, we introduce step-level stochasticity to evaluate and improve the agent's adaptability to non-stationary environments. Unlike traditional scheduling assumptions where prices remain static throughout the horizon, we model the product selling price $Price_{mp,t}$ as a time-variant stochastic process following a Bounded Random Walk, as described in Eqs. (18) and (19):

$$Price_{mp,t+1} = \max\left(0, Price_{mp,t} + \Delta Price_{mp,t}\right) \tag{18}$$

$$\Delta Price_{mp,t} = \mathrm{clip}(\epsilon_t, -3\%, 3\%) \times Price_{mp,t} \tag{19}$$

where $\epsilon_t \sim \mathcal{N}(0, \delta^2)$ represents the daily random perturbation term. This dynamic structure introduces strict non-stationarity into the environment, implying that the economic return of a decision is coupled with fluctuating market signals. Consequently, this compels the agent to develop "market acumen"—learning to predictively build up inventory of high-value products during upward price trends or shift to lower-cost grades during downturns. By training under these dual stochastic conditions, the agent learns a robust policy that maximizes total economic benefits even in highly dynamic and uncertain market environments .

3 CASE STUDY

Table 1: Operational parameters and product specifications for Polyethylene (PE) production units.

unit code	Num of product	initial product	initial rate	key raw material
FDP1-1	6	DG0	42.5	ethylene
FDP1-2	8	D8H	43.0	limits:
LPU1-1	1	TZG	10.5	[2000,8000]
LPU1-2	1	POE	4.2	initial:5250
LPU2-1	2	TUN	39.3	

Table 2: Operational parameters and product specifications for Polypropylene (PP) production units.

unit code	Num of product	initial product	initial rate	key raw material
PPU1-1	3	PP3	8.5	propylene
PPU1-2	4	K33	8.5	limits:
PPU2-1	3	T44	28.5	[1200,4000]
PPU2-2	7	TF7	27.5	initial:3550

3.1 Base Case

To validate the effectiveness of the proposed framework, we conducted a case study utilizing real-world data from a polyolefin production plant. The scheduling horizon was set to $T = 30$ days, covering two distinct production scenarios: Polyethylene (PE) and Polypropylene (PP). The PE scenario comprises 5 production units and 12 product grades, whereas the PP scenario consists of 4 units and 13 grades. These scenarios exhibit distinct operational characteristics: the PE process involves semi-continuous transitions (requiring shutdowns) for specific grade changes, while the PP process is characterized exclusively by continuous transitions. Detailed parameters for these production units are provided in Table 1 and Table 2.

The model was trained using the discrete Soft Actor-Critic (SAC-discrete) algorithm implemented on the Tianshou reinforcement learning platform. For reproducibility, the model was trained for 500 epochs with a batch size of 512, a fixed random seed of 42. Both actor and critic networks utilized a [512, 512] hidden layer architecture with learning rates of 10^{-4} and 3×10^{-4}, respectively, alongside a discount factor $\gamma = 0.99$. The reward function weights were configured as $\omega_1 = 10^{-6}, \omega_2 = 2.0$, and $\omega_3 = 0.5$.

Figure 4 presents the testing results for both scenarios, depicting scheduling Gantt charts, comparisons between realized production yields and market demands, and inventory trajectories for critical raw materials. As evidenced by the visualization, the schedules generated by the trained agent for both PE and PP lines strictly adhere to demand requirements and raw material inventory constraints. These results demonstrate the method's capability to successfully accomplish complex scheduling tasks within industrial specifications.

3.2 Comparative Experiments

We designed two sets of comparative experiments to systematically validate the effectiveness of the action masking mechanism and to evaluate the adaptability of the proposed reinforcement learning framework in handling dynamic environments.

3.2.1 Effectiveness of Action Masking

To assess the contribution of the action masking mechanism, we conducted an ablation study by training a variant of the model with the masking mechanism disabled on the Polyethylene scenario, while maintaining identical hyperparameters. The comparison of training curves between the unmasked baseline and the proposed masked model is illustrated in Figure 5.

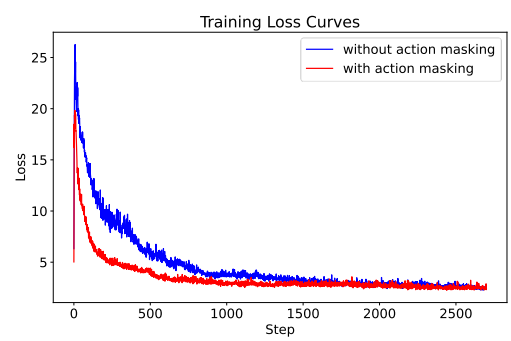

Figure 5. Comparison of training loss convergence with and without the action masking mechanism.

As illustrated in the figure, the incorporation of the action masking mechanism yields a significantly more stable loss profile and an accelerated convergence rate during training. This indicates that action masking effectively constrains the decision space, preventing the agent from engaging in futile exploration of invalid states during the early training stages, thereby substantially improving sample efficiency.

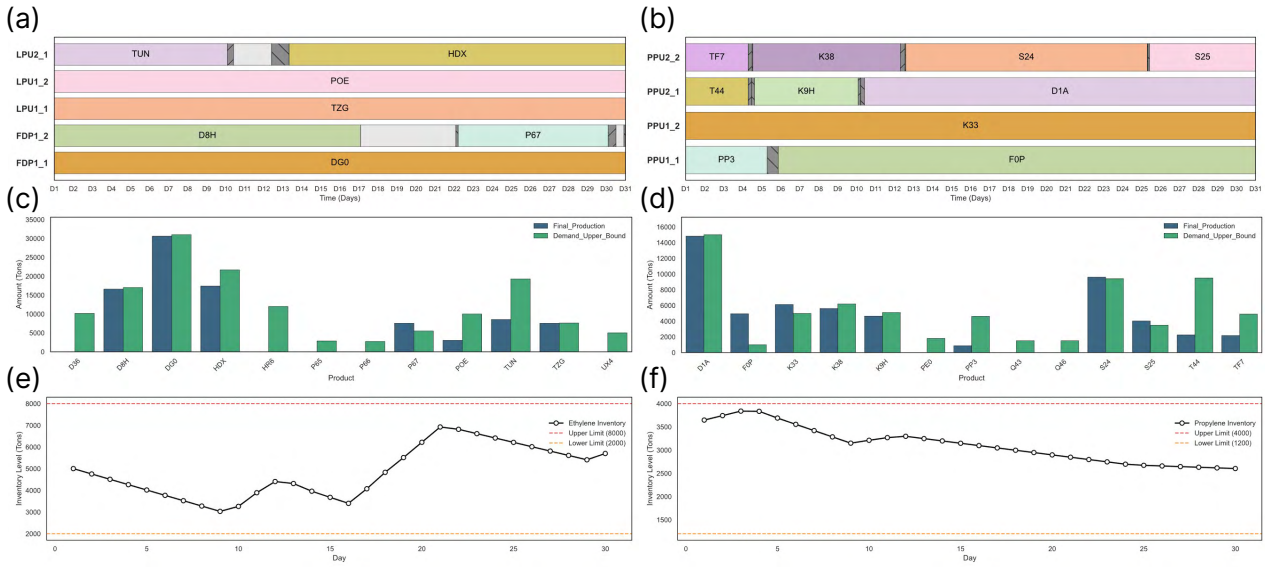

Figure 4: Scheduling optimization results for Polyethylene (PE) and Polypropylene (PP) scenarios: (a) Scheduling Gantt chart for PE units; (b) Scheduling Gantt chart for PP units; (c) Comparison of realized production yields vs. market demand for PE units; (d) Comparison of realized production yields vs. market demand for PP units; (e) Inventory trajectory of the key raw material (Ethylene) for PE units; (f) Inventory trajectory of the key raw material (Propylene) for PP units.

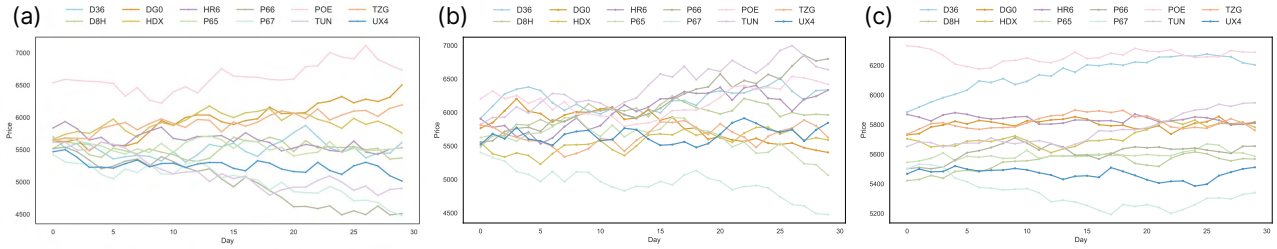

Figure 6: Three distinct price fluctuation scenarios generated by the Random Walk model for the Polyethylene line: (a) Scenario 1(($\delta = 0.15$); (b) Scenario 2(($\delta = 0.1$); (c) Scenario 3(($\delta = 0.05$).

3.2.2 Adaptability to Dynamic Price Fluctuations

To evaluate the agent's responsiveness to market dynamics, we incorporated the daily random walk for product prices (as defined in Eq. (18)-(19)) into the training environment and trained a specific "Dynamic Agent" . For the testing phase, we generated three distinct test instances characterized by different price fluctuation profiles, as shown in Figure 6.

We compared the performance of the "Dynamic Agent" against a "Static Agent" (trained without price volatility) on these instances. The resulting Gantt charts are compared in Figure 7.

The results reveal a distinct contrast in behavioral patterns. The Static Agent generated rigid schedules that remained nearly invariant across different price scenarios. In contrast, the Dynamic Agent proactively adjusted its production sequence and inventory strategy in response to varying price signals, thereby achieving superior economic returns (higher objective function values). This confirms the model's capability to effectively handle dynamic manufacturing environments.

4 CONCLUSION

This study addresses the complex scheduling problem in multi-line polyolefin production by proposing a Deep Reinforcement Learning (DRL) solution grounded in a high-fidelity simulation environment. By integrating a tuple-based action space with a domain-knowledge-embedded dynamic action masking mechanism, the proposed framework effectively overcomes the challenges of combinatorial explosion and strict hard constraint satisfaction inherent in industrial processes. The case study utilizing real-world data verifies that the trained agent can autonomously generate optimal schedules that

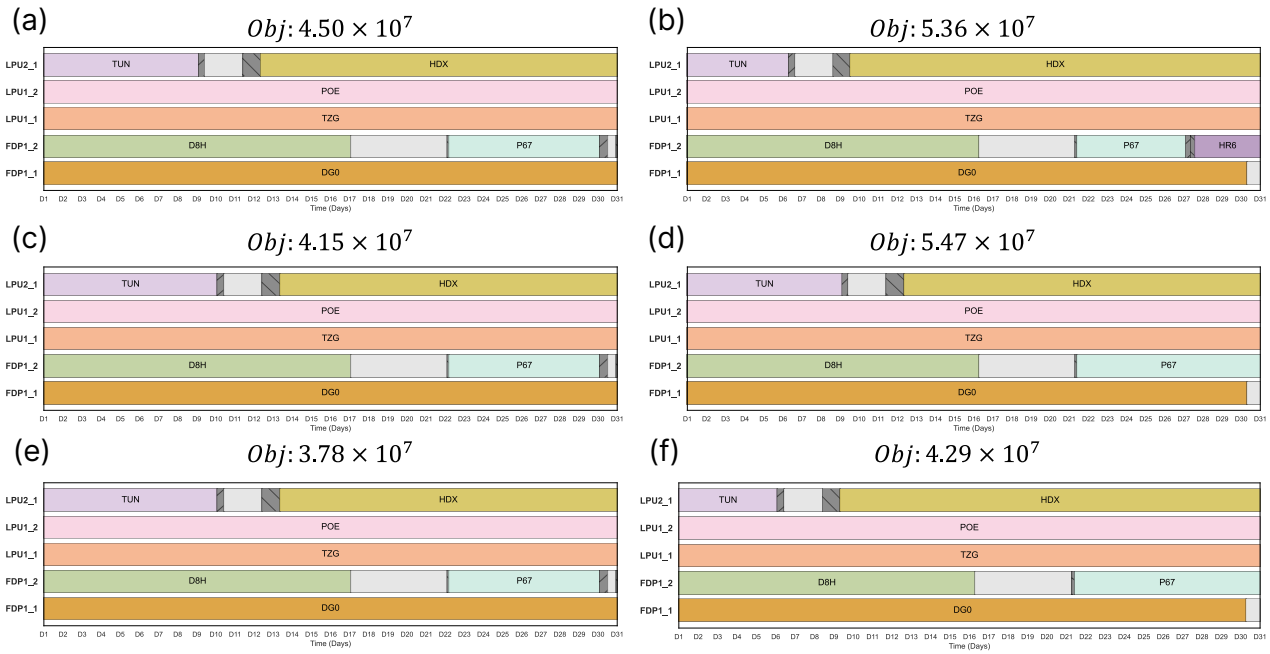

Figure 7: Comparison of scheduling strategies under different price fluctuation scenarios: (a, c, e) Schedules generated by the static model (trained without price volatility) under Scenarios 1, 2, and 3, respectively; (b, d, f) Schedules generated by the adaptive model (trained with price volatility) under Scenarios 1, 2, and 3, respectively.

strictly adhere to complex operational rules, such as transition constraints and inventory limits. Furthermore, comparative experiments demonstrate that the action masking mechanism significantly accelerates training convergence by eliminating futile exploration. Crucially, the agent trained under dynamic conditions exhibits robust adaptability to non-stationary market price fluctuations, proactively adjusting production strategies to maximize economic benefits. These findings confirm that the "offline training, online execution" paradigm of DRL offers a highly efficient and flexible alternative to traditional static optimization methods for intelligent industrial decision-making.

AUTHOR IDENTIFIERS

Author ORCIDs:
Zhineng Tao: 0009-0009-3929-6455
Tong Qiu: 0000-0003-4294-9671

REFERENCES

1. Modrak V, Sudhakarapandian R, Balamurugan A, Soltysova Z. A review on reinforcement learning in production scheduling: an inferential perspective. Algorithms 17:343 (2024). https://doi.org/10.3390/a17080343
2. Mowbray, M., Zhang, D., Chanona, E.A.D.R., 2022. Distributional Reinforcement Learning for Scheduling of Chemical Production Processes. https://doi.org/10.48550/arXiv.2203.00636
3. Yan J, Zhao T, Zhang T, Chu H, Yang C, Zhang Y. A dynamic scheduling method combining iterative optimization and deep reinforcement learning to solve sudden disturbance events in a flexible manufacturing process. Mathematics 13:4 (2024). https://doi.org/10.3390/math13010004
4. Hoss, J., Schelling, F., Klarmann, N., 2025. A Production Scheduling Framework for Reinforcement Learning Under Real-World Constraints. https://doi.org/10.48550/arXiv.2506.13566
5. Li H, Qiu T. Supply responsive scheduling for ethylene cracking furnace systems based on deep reinforcement learning. AIChE Journal 70: (2024). https://doi.org/10.1002/aic.18563
6. Altman, E., 2021. Constrained Markov Decision Processes: Stochastic Modeling, 1st ed. Routledge, Boca Raton. https://doi.org/10.1201/9781315140223

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Multi-Objective Optimisation of Pressure Swing Adsorption Systems via Symbolic Regression

Carine Menezes REBELLO[a*], Amilton Barbosa BOTELHO JUNIOR[a], Anderson Rapello DOS SANTOS[bc], and Idelfonso B. R. NOGUEIRA[a]

[a] Department of Chemical Engineering, Norwegian University of Science and Technology, Trondheim, Norway
[b] Petroleo Brasileiro S.A. (PETROBRAS), Wells/Well Engineering, Brazil.
[c] Pontifical Catholic University of Rio de Janeiro (PUC-Rio), Department of Mechanical Engineering (DEM), Rio de Janeiro, RJ, Brazil.
* Corresponding Author: carine.m.rebello@ntnu.no.

ABSTRACT

This work explores symbolic regression (SR) as an interpretable surrogate modelling approach for the multi-objective optimisation of pressure swing adsorption (PSA) systems for CO_2 capture. A first-principle model was used as a virtual plant to generate synthetic datasets covering the operating space defined by cycle step durations. Two surrogate frameworks were developed and compared: SR models derived through evolutionary search and deep neural networks (DNNs) trained via Hyperband-based tuning. Both surrogates were used as simulation models within an optimisation procedure based on a particle swarm optimisation (PSO) algorithm to maximise CO_2 purity and recovery. While DNNs achieved the lowest prediction errors (MSE ≈ $10^{-6)}$, the SR surrogates provided compact analytical representations and significantly faster optimisation. The SR framework yielded a denser and more diverse Pareto front (4345 vs 508 points). It was about 34 times faster (38.6 s vs 1331 s), confirming its efficiency for surrogate-based optimisation of cyclic adsorption processes.

Keywords: PSA, multi-objective optimisation, surrogate models, optimality, symbolic regression

INTRODUCTION

Pressure swing adsorption (PSA) systems play a key role in gas separation technologies due to their flexibility, high efficiency, and low maintenance costs (Capra et al., 2018; Habib et al., 2016). They have been extensively employed in processes ranging from air separation (Kamin et al., 2022; Tian et al., 2022) and hydrogen purification (Luberti and Ahn, 2022; Kalman et al., 2022) to carbon dioxide (CO_2) capture from flue gases (Nikolaidis et al., 2016, 2015). Despite their industrial maturity, optimisation of PSA units remains challenging due to their inherently cyclic operation, strong multiscale dynamics, and the tight coupling between adsorption, desorption, and switching phases, which lead to highly nonlinear and constrained process behaviour (Khajuria & Pistikopoulos, 2013; Leperi et al., 2019). The coupled mass, heat, and momentum balances that govern adsorption and desorption transitions must be solved repeatedly over multiple cycles until the system reaches a cyclic steady state (CSS), resulting in a computationally expensive optimisation problem governed by large sets of stiff partial differential equations (PDEs) (Ward and Pini, 2022; Rebello & Nogueira, 2024).

Surrogate modelling techniques have emerged as an alternative to accelerate process optimisation while retaining sufficient physical accuracy. Surrogate-based optimisation replaces the high-fidelity first-principles model with an approximate data-driven representation that captures the input-output relationships between decision variables (e.g., step durations, pressures, and reflux ratios) and process performance metrics (e.g., CO_2 purity and recovery).

These surrogate models have been successfully applied to PSA, generating Pareto fronts that describe the trade-offs between competing objectives, such as purity, recovery, and energy consumption (De Witte et al., 2021; Panerati et al., 2019; Zhang et al., 2022).

Among the various surrogate modelling strategies, deep neural networks (DNNs) have received particular attention (Costa et al., 2024; Nogueira et al., 2022; Subraveti et al., 2019). As demonstrated by Rebello and

Nogueira (2024), DNN-based surrogate frameworks can reliably approximate the nonlinear behaviour of cyclic adsorption systems, drastically reducing simulation time and enabling systematic multi-objective optimisation. The resulting operating maps help identify optimal process configurations and inform decision-making. Despite their predictive accuracy, these data-driven models remain largely opaque, offering limited interpretability and little insight into the physical or phenomenological mechanisms underlying adsorption dynamics (Rebello et al., 2024).

To overcome these limitations, recent research has begun to explore more transparent machine learning paradigms, in which the emphasis shifts from pure prediction to computational efficiency and analytical tractability. In this context, SR offers a promising approach (Makke & Chawla, 2024; Quade et al., 2016; Angelis et al., 2023). Unlike neural networks, which encode relationships implicitly through parameters distributed in hidden layers, SR explicitly searches for analytical expressions that best fit the available data (Rebello et al., 2024; Hu & Zhang, 2023). By combining evolutionary search and sparse regression principles, symbolic regression identifies compact, readable equations that capture system behaviour while revealing underlying physical trends and nonlinear dependencies.

Symbolic regression has been increasingly adopted across a wide range of engineering applications. In process monitoring and fault detection, SR has been used to derive interpretable diagnostic expressions directly from process data (Hale et al., 2022). In soft sensing, SR has been applied to construct explicit input-output mappings for inferential estimation of difficult-to-measure variables (Kay et al., 2024). In the context of control, SR has been employed to identify compact symbolic controllers and model predictive control surrogates that are both computationally lightweight and physically transparent (Scheffold et al., 2021). For optimisation, SR-derived surrogates have been embedded into mathematical programming frameworks to enable gradient-based search and reduce evaluation cost (Sofronova et al., 2021). More broadly, SR has demonstrated strong potential for scientific discovery, recovering governing equations from data across physics, chemistry, and materials science (Angelis et al., 2023).

Beyond its direct application as a surrogate modelling tool, SR offers several promising extensions, particularly relevant to computationally intensive multi-objective optimisation problems. The analytical nature of SR expressions enables direct embedding into algebraic optimisation solvers, avoiding the matrix operations and numerical differentiation overhead inherent to black-box models such as DNNs, which becomes especially advantageous when objective function evaluations must be performed repeatedly across large swarms or populations (Cozad et al., 2014). Furthermore, SR-derived surrogates naturally produce smooth, differentiable expressions that facilitate gradient-based search and yield denser, more diverse Pareto fronts at a lower computational cost, while also producing a significantly richer set of optimal solutions. These properties suggest that SR could be extended to higher-dimensional multi-objective problems, to online re-optimisation under changing operating conditions, and to real-time decision support in CO_2 capture systems, where computational efficiency is a critical practical constraint.

However, symbolic regression models have not yet been investigated for PSA systems. This represents a notable gap in the emerging literature on machine learning applications in chemical process systems, where the potential of faster surrogate models remains largely unexplored. Moreover, given the highly nonlinear nature of PSA dynamics, the behaviour of SR-based surrogates requires careful assessment, as these models may introduce spurious or artificial local minima in the optimisation landscape if not properly regularised and validated.

In this context, the objective of this work is to compare the performance of DNN and SR surrogate models for the optimisation of PSA systems, and to evaluate whether SR-based models can effectively serve as reliable and computationally efficient surrogates, where interpretability is understood as a consequence of the analytical nature of SR expressions rather than a claim of physical correspondence to adsorption dynamics.

METHODOLOGY

The overall methodology adopted in this study is illustrated in Figure 1. The methodology is structured into four main stages: data curation (Section 2.1), SR model (Section 2.2), DNN model (Section 2.3), and multi-objective optimisation (Section 2.4).

The case study investigated in this work concerns the optimisation of a single-column PSA unit for capturing CO_2 from a binary CO_2/N_2 mixture. The PSA process operates cyclically, relying on the preferential adsorption of CO_2 over N_2 onto a solid adsorbent, zeolite 13X, under alternating pressure conditions. Each cycle comprises a sequence of steps (Figure 2): pressurisation (pres), adsorption (ads), heavy reflux (HR), depressurisation (depress) and light reflux (LR). During the pressurisation stage, the column pressure is increased to the adsorption level; in the adsorption step, CO_2 is selectively retained by the zeolite, while N_2 exits as the light product. The subsequent depressurisation and purge steps promote CO_2 desorption and restore the adsorbent's capacity for the next cycle. The process is governed by key decision variables, including the durations of each step (t_{pres}, t_{ads}, t_{depres}, and t_{purge}). The leading performance indicators considered are CO_2 purity and CO_2 recovery, which depend nonlinearly on these operational conditions and on the

thermodynamic and kinetic properties of the zeolite 13X adsorbent.

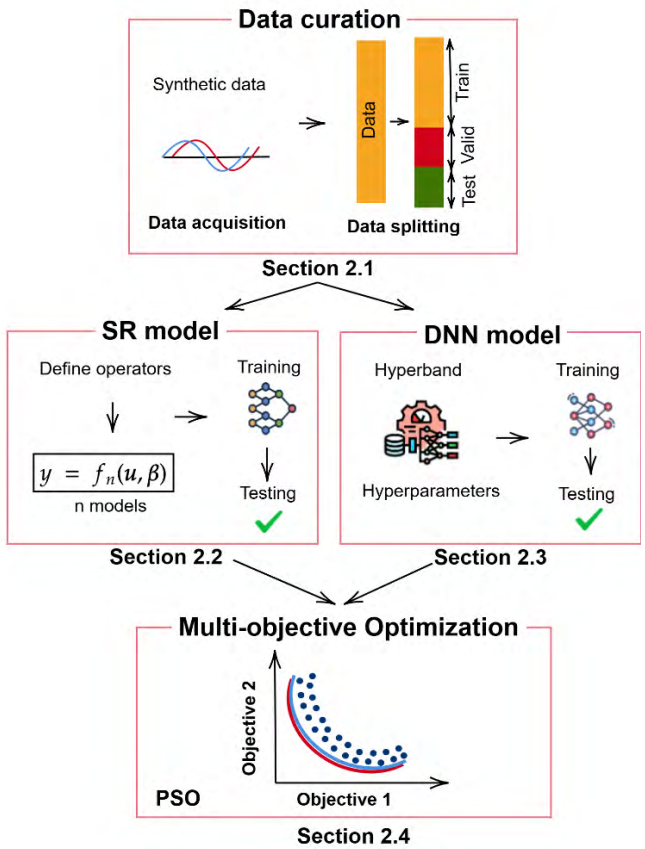

Figure 1. Diagram describing the methodology adopted in this work.

The detailed formulation of the mechanistic PSA model used in this work, including the PDE system, modelling assumptions, boundary and initial conditions, and the adopted kinetic and process parameters, is presented in Rebello and Nogueira (2024).

Data curation

Synthetic data were generated using a mechanistic PSA model implemented as a virtual plant. The sampling of operating conditions was performed using Latin Hypercube Sampling (LHS) to ensure a uniform and statistically independent exploration of the design space. The selected input variables correspond to the step durations of the PSA cycle, pressurisation (t_{pres}), adsorption (tads), heavy reflux (t_{HR}), depressurisation (t_{depres}), and light reflux (t_{LR}), which were systematically varied within physically feasible limits. For each LHS sample, the virtual plant was simulated until a cyclic steady state (CSS) was reached, and the corresponding process performance indicators were recorded. The model outputs, or target variables, are CO_2 purity and recovery, which quantify the PSA process's separation efficiency and product quality.

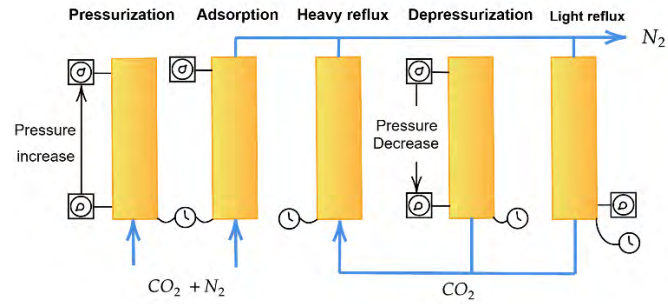

Figure 2. Schematic representation of the PSA cycle for CO_2 capture from a CO_2/N_2 gas mixture, illustrating the successive steps.

The complete dataset was then curated to remove incomplete or non-convergent simulations and subsequently normalised to improve numerical stability during model training. Finally, the curated data were randomly split into training (70 %), validation (15 %), and testing (15 %) subsets. This partition ensures that the surrogate models, both SR and DNN, are trained on representative data while preserving independent sets for unbiased performance evaluation and generalisation assessment.

SR model

The SR framework was employed to identify analytical surrogate models that capture the nonlinear relationship between the PSA operating variables and the process performance indicators, namely CO_2 purity and CO_2 recovery. Symbolic regression formulates this task as a nonlinear regression problem of the form:

$$\boldsymbol{y} = \boldsymbol{f}(\boldsymbol{u_n}, \boldsymbol{\beta}) \tag{1}$$

where u denotes the vector of decision variables (step durations), and β represents the fitted parameters. Candidate expressions are iteratively constructed by combining elementary mathematical operators: {+, −, ×, ÷, sin, cos, exp, log} and constants through an evolutionary symbolic search. This process explores a large function space while progressively selecting the most representative equations based on their predictive accuracy and structural simplicity.

Model selection was guided by a multi-objective fitness criterion that balances prediction error, quantified by the mean squared error (MSE) on the validation dataset, with a complexity penalty proportional to expression length.

The evolutionary search was terminated once convergence was achieved, i.e., when no significant improvement in model accuracy or parsimony was observed.

The best-performing symbolic expressions for CO_2 purity and recovery were then independently validated using the testing subset. These analytical surrogates

replace the high-fidelity PSA model in the multi-objective optimisation framework.

The SR models were developed in Python using the PySRRegressor package (Tonda, 2025) (see: https://astroautomata.com/PySR/), an open-source library for SR that combines genetic programming with simulated annealing for evolutionary search. The algorithm was run with random initialisation, and multiple repetitions were performed to ensure convergence to consistent solutions.

DNN model

The same dataset generated by the mechanistic PSA model was also used to identify the DNN models. The network architecture was defined using the HYPERBAND algorithm (Li et al., 2018) (see: https://keras.io/keras_tuner/api/tuners/hyperband/), which efficiently explores the hyperparameter space to identify near-optimal configurations. The DNNs adopted a multiple-input single-output (MISO) structure, where the input layer receives the PSA decision variables t_{pres}, t_{ads}, t_{HR}, t_{depres}, and t_{LR}, and the output layer provides a single process indicator, either CO_2 purity or CO_2 recovery.

The search space considered the number of hidden layers (ranging from 1 to 5), the number of neurons per layer (from 50 to 300 in steps of 50), activation functions (ReLU and tanh), and learning rates (0.0001, 0.001, and 0.01). Each candidate model was trained using the MSE as the loss function, with the Adam optimiser configured according to the selected learning rate.

Model training aimed to minimise the MSE between predicted and target values, while the mean absolute error (MAE) was also monitored as an auxiliary metric. Each model was trained for up to 1000 epochs, with early stopping and dropout regularisation applied to prevent overfitting and improve generalisation.

After training, the resulting DNN models were validated on the independent test set to assess predictive accuracy. The best-performing networks were subsequently integrated into the multi-objective optimisation framework.

Multi-objective optimisation

The optimisation stage aims to identify the operating conditions that maximise CO_2 purity and recovery, as predicted by the surrogate models. The problem is formulated as a multi-objective optimisation (MOO) task, in which both objectives are maximised simultaneously subject to operational constraints. Mathematically, it can be expressed as:

$$\begin{cases} \max Pur_{CO_2}(\boldsymbol{u}, \boldsymbol{\beta}) \\ \max Rec_{CO_2}(\boldsymbol{u}, \boldsymbol{\beta}) \end{cases} \tag{2}$$

subject to:

$$\boldsymbol{u} \in [\boldsymbol{u}_{min}, \boldsymbol{u}_{max}] \tag{3}$$

where **u** represents the decision variable vector containing the step durations of the PSA cycle, and β denotes the surrogate model parameters. The objective functions are defined as follows:

$$Pur_{CO_2} = \frac{moles_{CO_2,out}\big|_{depress} + moles_{CO_2,out}\big|_{LR}}{moles_{CO_2,total}\big|_{depress} + moles_{CO_2,total}\big|_{LR}} \tag{4}$$

$$Rec_{CO_2} = 1 - \frac{moles_{CO_2,out}\big|_{ads}}{moles_{CO_2,in}\big|_{press} + moles_{CO_2,in}\big|_{ads}} \tag{5}$$

The multi-objective problem was solved using the particle swarm optimisation (PSO) algorithm, chosen for its efficiency in handling nonlinear and nonconvex search spaces. Each surrogate model, SR, and deep neural network, DNN, was embedded within the optimisation routine to compute the objective functions at each iteration. The algorithm simultaneously explores the trade-off between CO_2 purity and recovery, producing a Pareto front that reflects the optimal balance between these competing objectives. This formulation enables a direct comparison of the optimisation performance obtained with SR- and DNN-based surrogates in terms of accuracy, smoothness, and convergence behaviour.

RESULTS AND DISCUSSIONS

The statistical characteristics of the dataset used for surrogate model identification were first analysed to verify the independence and coverage of the sampled operating variables. The correlation heatmap (Figure 3) indicates weak pairwise correlations among inputs, minimising multicollinearity during training.

Figure 4 displays the LHS generated step durations for t_{pres}, t_{depres}, t_{ads}, t_{LR}, and t_{HR}. The plots show patterns consistent with stratified, quasi-independent draws. Note that the horizontal axis represents the sample order of LHS points, evidencing broad coverage of the operating domain for each step.

The analytical surrogate expressions obtained using the SR approach for predicting CO_2 purity and recovery are presented in Equations 6 and 7. The resulting symbolic forms represent compact yet accurate approximations of the nonlinear mappings between the process decision variables.

$$Pur_{CO_2} = \sin\left(\frac{t_{ads} + 0.058}{sin^3\left(\frac{1}{cos(t_{depres}^3)}\right) + sin\left(t_{depres} e^{cos(0.609)}\right)}\right) \tag{6}$$

$$Rec_{CO_2} = 0.91\cos^4\left(e^{-4(\sin(t_{depres}/0.94)\sin(t_{ads}))^2} t_{ads}\right) \quad (7)$$

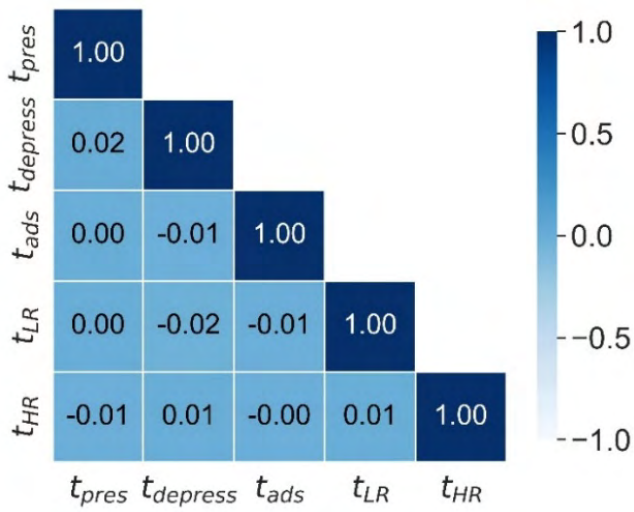

Figure 3. Correlation heatmap for the input variables of the surrogate models.

Both SR-derived equations capture the nonlinear dependence of CO_2 purity and recovery on the durations of the adsorption (t_{ads}) and depressurisation (t_{depres}) steps. Their explicit analytical structure facilitates low-complexity surrogate representation suitable for direct embedding into optimisation routines.

The DNN developed for predicting CO_2 purity used a fully connected feedforward architecture with five input features and four hidden layers, employing the ReLU activation function. The hidden layers comprised 300, 200, 50, and 200 neurons, respectively, followed by a single linear output neuron. This configuration resulted in 82451 trainable parameters. The model was trained using the Adam optimiser with a learning rate of 0.001 and MSE as the loss function.

Similarly, the DNN model designed for CO_2 recovery employed the same input structure and ReLU-activated hidden layers, but with a more compact topology. The network consisted of two hidden layers containing 300 and 100 neurons, respectively, also followed by a single linear output neuron. The total number of trainable parameters in this configuration was 32001. This model was also trained with the Adam optimiser at a learning rate of 0.001, ensuring consistent training dynamics and convergence across both networks.

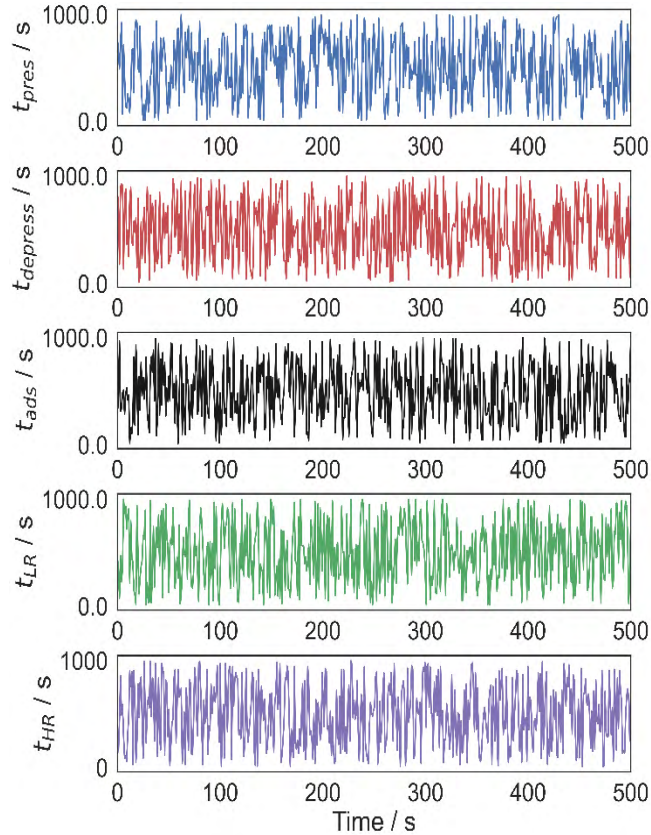

Figure 4. Sampled PSA Step Durations Generated by LHS.

Figure 5 compares the parity plots for the two surrogate modelling approaches used to predict CO_2 purity and recovery. Subfigures 5a and 5b correspond to the SR-based models, while subfigures 5c and 5d refer to the DNN surrogates. In all cases, the predicted values exhibit a good agreement with the reference data. The mechanistic PSA model has the points tightly distributed along the identity line. A slightly better predictive performance is observed for the DNN models, which exhibit lower dispersion and greater alignment along the diagonal, reflecting their enhanced capacity to capture complex nonlinear relationships.

Table 1 summarises the prediction errors obtained for both surrogate modelling strategies. The metrics considered were the MSE and MAE, both computed in the denormalised domain for direct physical interpretation. The DNN models outperform the SR surrogates in predictive accuracy, with MSE values one to two orders of magnitude lower. Nevertheless, the symbolic regression models also demonstrate satisfactory performance, achieving low error rates and accurately reproducing the PSA system's main nonlinear trends.

The multi-objective optimisation procedure was carried out using the PSO algorithm, configured with a swarm of 100 particles and 100 iterations. Both surrogate-based frameworks (SR and DNN models) were

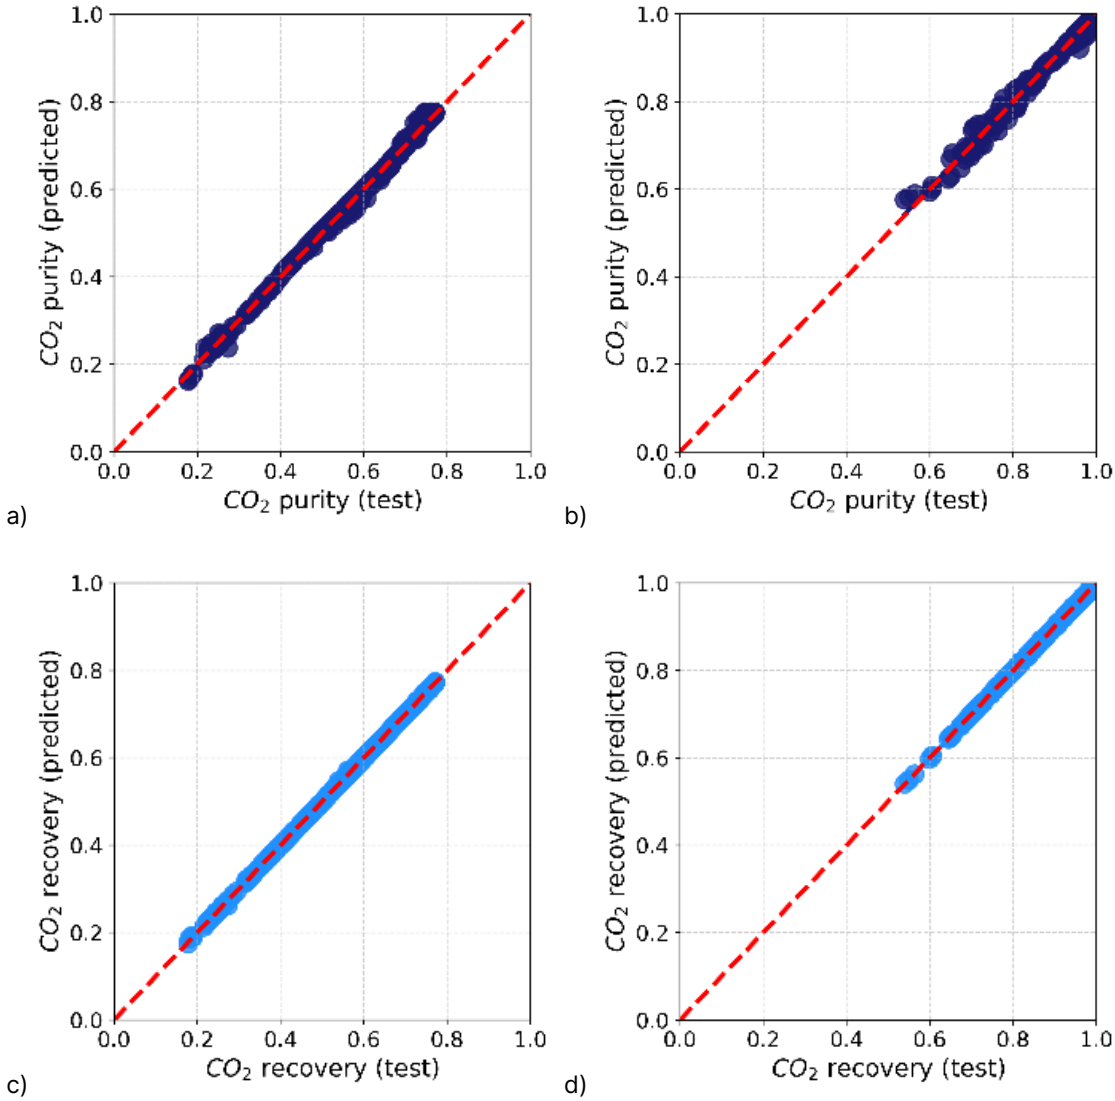

Figure 5. Parity plots comparing predicted and reference values for CO_2 purity and recovery obtained with SR (a, b) and DNN (c, d) surrogate models.

employed in the optimisation problem defined in Section 2. Additionally, the optimisation results were compared with those obtained using the mechanistic PSA model, which was optimised under the same conditions but with 50 particles and 50 iterations. The decision variable values corresponding to the Pareto-optimal points from each surrogate model were further evaluated in the mechanistic simulator to verify the consistency and physical plausibility of the predicted solutions. The resulting Pareto fronts, illustrated in Figure 6, display the trade-off between CO_2 purity and CO_2 recovery for each modelling strategy.

Table 1: Final MSE and MAE values for the test dataset obtained with the SR and DNN surrogate models.

Model	Variable	MSE	MAE
SR	CO_2 purity	5.0×10^{-5}	4.33×10^{-3}
SR	CO_2 recovery	5.9×10^{-5}	4.03×10^{-3}
DNN	CO_2 purity	1.24×10^{-6}	5.70×10^{-4}
DNN	CO_2 recovery	5.22×10^{-7}	5.11×10^{-4}

The optimisation using the SR surrogate produced a Pareto front containing 4345 optimal solutions, while the DNN surrogate yielded 508 points. Despite both optimisations using the same number of particles and iterations, the SR-based approach produced a denser, more diverse Pareto front, indicating superior exploration of the solution space compared to the DNN model.

This is attributed to the continuous and differentiable nature of the analytical expressions derived via SR,

which induces a smoother optimisation landscape, facilitating uniform exploration of the trade-off region by the PSO algorithm. In contrast, DNN surrogates may introduce local irregularities and discontinuities due to nonlinear activation functions such as ReLU, causing the swarm to concentrate on fewer regions of the Pareto front and resulting in a sparser front.

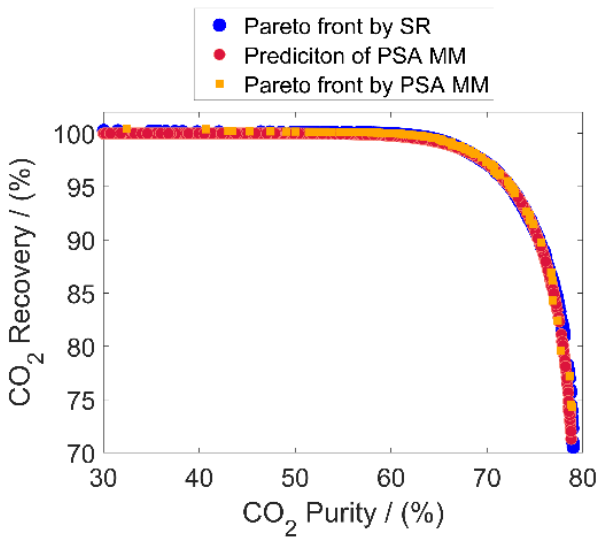

a) Pareto front by SR surrogate

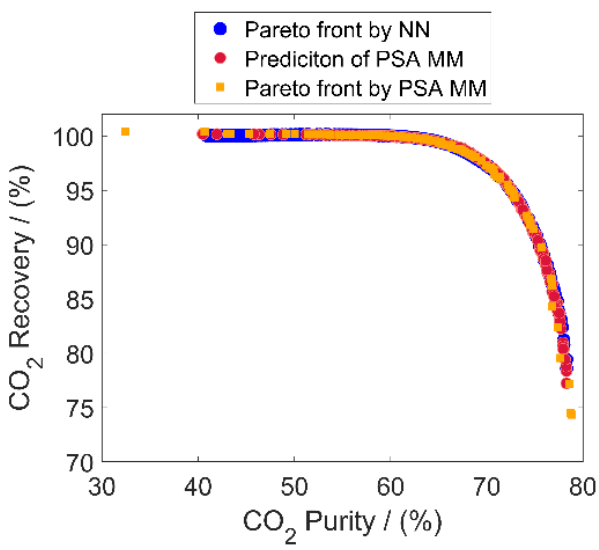

b) Pareto front by DNN surrogate.

Figure 6. Comparison of Pareto fronts obtained with SR and DNN surrogates against the PSA MM.

Figure 7 presents the cumulative optimisation time per iteration obtained for both surrogate-based optimisation frameworks. Each curve represents the total computation time accumulated across all particles up to a given PSO iteration. The SR-based optimisation required a total of 38.6 s to complete the 100 iterations, whereas the DNN-based framework required 1331.5 s under the same conditions.

These results highlight a substantial difference in computational efficiency: the SR approach was approximately 34.5 times faster than the DNN model. This performance gap arises mainly from the analytical nature of the SR surrogates, which allow direct function evaluations without the matrix operations and backpropagation overhead inherent to neural networks. Consequently, SR models enable a much faster exploration of the search space while maintaining sufficient predictive accuracy for the multiobjective optimisation task.

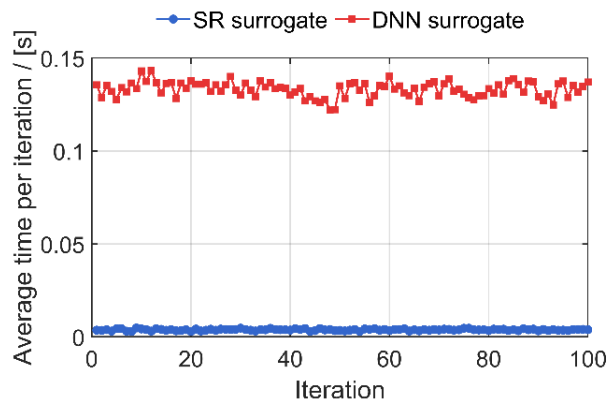

Figure 7. Average computational cost per iteration of optimisation using surrogate SR and DNN models.

CONCLUSION

This study demonstrated that SR can serve as an efficient surrogate modelling strategy for optimising PSA processes. Compared to DNN, the SR-based framework achieved faster convergence and produced a more diverse Pareto front while maintaining good predictive accuracy. These findings highlight the potential of SR as a computationally lightweight alternative for data-driven optimisation of cyclic adsorption systems, offering a balance between robustness, and efficiency.

ACKNOWLEDGEMENTS

The present work contributes to completion of a sub project at SUBPRO, a research-based innovation center within Subsea Production and Processing at the Norwe gian University of Science and Technology. The authors would like to express their gratitude for the financial support received from SUBPRO, funded by the Research Council of Norway through grant number 237893, major industry partners, and NTNU.

REFERENCES

1. Angelis D, Sofos F, Karakasidis TE. Artificial

intelligence in physical sciences: symbolic regression trends and perspectives. Arch Computat Methods Eng 30:3845-3865 (2023). https://doi.org/10.1007/s11831-023-09922-z
2. Capra, F., Gazzani, M., Joss, L., Mazzotti, M., and Martelli, E. (2018). MO-MCS, a derivative-free algorithm for the multiobjective optimisation of adsorption processes. Technical report.
3. Costa EA, Rebello CM, Schnitman L, Loureiro JM, Ribeiro AM, Nogueira IBR. Adaptive digital twin for pressure swing adsorption systems: integrating a novel feedback tracking system, online learning and uncertainty assessment for enhanced performance. Engineering Applications of Artificial Intelligence 127:107364 (2024). https://doi.org/10.1016/j.engappai.2023.107364
4. Cozad A, Sahinidis NV, Miller DC. Learning surrogate models for simulation?based optimization. AIChE Journal 60:2211-2227 (2014). https://doi.org/10.1002/aic.14418
5. De Witte N, Denayer JFM, Van Assche TRC. Effect of adsorption duration and purge flowrate on pressure swing adsorption performance. Ind. Eng. Chem. Res. 60:13684-13691 (2021). https://doi.org/10.1021/acs.iecr.1c02291
6. Ben-Mansour R, Habib MA, Bamidele OE, Basha M, Qasem NAA, Peedikakkal A, Laoui T, Ali M. Carbon capture by physical adsorption: materials, experimental investigations and numerical modeling and simulations – a review. Applied Energy 161:225-255 (2016). https://doi.org/10.1016/j.apenergy.2015.10.011
7. Hale WT, Safikou E, Bollas GM. Inference of faults through symbolic regression of system data. Computers & Chemical Engineering 157:107619 (2022). https://doi.org/10.1016/j.compchemeng.2021.107619
8. Hu W, Zhang L. First-principles, machine learning and symbolic regression modelling for organic molecule adsorption on two-dimensional cao surface. Journal of Molecular Graphics and Modelling 124:108530 (2023). https://doi.org/10.1016/j.jmgm.2023.108530
9. Kalman V, Voigt J, Jordan C, Harasek M. Hydrogen purification by pressure swing adsorption: high-pressure PSA performance in recovery from seasonal storage. Sustainability 14:14037 (2022). https://doi.org/10.3390/su142114037
10. Kamin Z, Bahrun MHV, Bono A. A short review on pressure swing adsorption (PSA) technology for nitrogen generation from air. AIP Conference Proceedings 2656:050006 (2022). https://doi.org/10.1063/5.0099702
11. Kay H, Kay S, Mowbray M, Lane A, Mendoza C, Martin P, Zhang D. Constructing a symbolic regression-based interpretable soft sensor for industrial data analytics and product quality control. Ind. Eng. Chem. Res. 63:4083-4092 (2024). https://doi.org/10.1021/acs.iecr.3c04021
12. Khajuria H, Pistikopoulos EN. Optimization and control of pressure swing adsorption processes under uncertainty. AIChE Journal 59:120-131 (2012). https://doi.org/10.1002/aic.13783
13. Leperi KT, Yancy-Caballero D, Snurr RQ, You F. 110th anniversary: surrogate models based on artificial neural networks to simulate and optimize pressure swing adsorption cycles for co₂ capture. Ind. Eng. Chem. Res. 58:18241-18252 (2019). https://doi.org/10.1021/acs.iecr.9b02383
14. Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., & Talwalkar, A. (2018). Hyperband: A novel bandit-based approach to hyperparameter optimization. *Journal of machine learning research*, *18*(185), 1-52.
15. Luberti M, Ahn H. Review of polybed pressure swing adsorption for hydrogen purification. International Journal of Hydrogen Energy 47:10911-10933 (2022). https://doi.org/10.1016/j.ijhydene.2022.01.147
16. Makke N, Chawla S. Interpretable scientific discovery with symbolic regression: a review. Artif Intell Rev 57: (2024). https://doi.org/10.1007/s10462-023-10622-0
17. Nikolaidis, G.N., Kikkinides, E.S., and Georgiadis, M.C. (2015). Modelling and simulation of pressure swing adsorption (psa) processes for post-combustion carbon dioxide (co2) capture from flue gas. In K.V. Gernaey, J.K. Huusom, and R. Gani (eds.), 12th.
18. International Symposium on Process Systems Engineering and 25th European Symposium on Computer Aided Process Engineering, volume 37 of Computer Aided Chemical Engineering, 287–292. Elsevier. doi: https://doi.org/10.1016/B978-0-444-63578-5.50043-8
19. Nikolaidis GN, Kikkinides ES, Georgiadis MC. Model-based approach for the evaluation of materials and processes for post-combustion carbon dioxide capture from flue gas by PSA/VSA processes. Ind. Eng. Chem. Res. 55:635-646 (2016). https://doi.org/10.1021/acs.iecr.5b02845
20. Nogueira IBR, Dias ROM, Rebello CM, Costa EA, Santana VV, Rodrigues AE, Ferreira A, Ribeiro AM. A novel nested loop optimization problem based on deep neural networks and feasible operation regions definition for simultaneous material screening and process optimization. Chemical Engineering Research and Design 180:243-253

(2022). https://doi.org/10.1016/j.cherd.2022.02.013
21. Panerati J, Schnellmann MA, Patience C, Beltrame G, Patience GS. Experimental methods in chemical engineering: artificial neural networks–anns. Can J Chem Eng 97:2372-2382 (2019). https://doi.org/10.1002/cjce.23507
22. Quade M, Abel M, Shafi K, Niven RK, Noack BR. Prediction of dynamical systems by symbolic regression. Phys. Rev. E 94: (2016). https://doi.org/10.1103/physreve.94.012214
23. Rebello CM, Costa EA, Fontana M, Schnitman L, Nogueira IBR. Interpretable scientific machine learning approach for correcting phenomenological models: methodology validation on an ESP prototype. Ind. Eng. Chem. Res. 63:19030-19050 (2024). https://doi.org/10.1021/acs.iecr.4c02104
24. Rebello CM, Nogueira IBR. Optimizing 2 capture in pressure swing adsorption units: a deep neural network approach with optimality evaluation and operating maps for decision-making. Separation and Purification Technology 340:126811 (2024). https://doi.org/10.1016/j.seppur.2024.126811
25. Scheffold L, Finkler T, Piechottka U. Gray-box system modeling using symbolic regression and nonlinear model predictive control of a semibatch polymerization. Computers & Chemical Engineering 146:107204 (2021). https://doi.org/10.1016/j.compchemeng.2020.107204
26. Subraveti SG, Li Z, Prasad V, Rajendran A. Machine learning-based multiobjective optimization of pressure swing adsorption. Ind. Eng. Chem. Res. 58:20412-20422 (2019). https://doi.org/10.1021/acs.iecr.9b04173
27. Sofronova E, Diveev A. Universal approach to solution of optimization problems by symbolic regression. Applied Sciences 11:5081 (2021). https://doi.org/10.3390/app11115081
28. Tian T, Wang Y, Liu B, Ding Z, Xu X, Shi M, Ma J, Zhang Y, Zhang D. Simulation and experiment of six-bed PSA process for air separation with rotating distribution valve. Chinese Journal of Chemical Engineering 42:329-337 (2022). https://doi.org/10.1016/j.cjche.2021.03.027
29. Tonda A. Review of pysr: high-performance symbolic regression in python and julia. Genet Program Evolvable Mach 26: (2024). https://doi.org/10.1007/s10710-024-09503-4
30. Ward A, Pini R. Efficient bayesian optimization of industrial-scale pressure-vacuum swing adsorption processes for co₂ capture. Ind. Eng. Chem. Res. 61:13650-13668 (2022). https://doi.org/10.1021/acs.iecr.2c02313
31. Zhang R, Shen Y, Tang Z, Li W, Zhang D. A review of numerical research on the pressure swing adsorption process. Processes 10:812 (2022). https://doi.org/10.3390/pr10050812

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

GPU-Accelerated Nonlinear Multi-Period AC Optimal Power Flow for Large-Scale Power–Hydrogen Systems

Geunseo Song[ab], Dirk Lauinger[bc], Sungho Shin[b]*, and Jonggeol Na[a]*

[a] Ewha Womans University, Department of Chemical Engineering and Materials Science, Seoul, Republic of Korea
[b] Massachusetts Institute of Technology, Department of Chemical Engineering, Cambridge, Massachusetts, USA
[c] Massachusetts Institute of Technology, MIT Energy Initiative, Cambridge, Massachusetts, USA
* Corresponding Authors: sushin@mit.edu (S. Shin), jgna@ewha.ac.kr (J. Na)

ABSTRACT

The growing penetration of renewable energy sources and power-to-hydrogen (P2H) systems demands high-fidelity, large-scale optimization frameworks that capture the nonlinear physics of both AC power flow and hydrogen thermodynamics. However, existing approaches rely on DC approximations and simplified electrolyzer models, neglecting critical operational constraints. As a result, accurately modeling such systems leads to large-scale nonlinear programs that are computationally intractable for conventional CPU-based solvers. This motivates the need for scalable optimization frameworks capable of handling both physical fidelity and computational complexity. This paper proposes a fully GPU-native framework for solving large-scale multi-period AC optimal power flow (AC-OPF) problems with integrated power-to-hydrogen systems. High-fidelity thermodynamic models of hydrogen production, compression, cooling, and storage are coupled with AC power flow constraints, resulting in large-scale nonlinear programs (NLPs) with up to 14.4 million variables and 22.4 million constraints in the largest benchmark case (9, 591-bus network with a 168-hour horizon). To enable scalable solutions, condensed-space KKT reformulations and GPU-accelerated sparse Cholesky factorization are employed within an end-to-end GPU optimization pipeline. Numerical results on benchmark networks up to 9, 591 buses demonstrate 10–600× speedups over CPU solvers, whereas CPU-based solvers fail to converge on the largest instances. Operational studies further highlight the importance of thermodynamic constraints in realistic hydrogen system scheduling.

Keywords: AC optimal power flow, power-to-hydrogen, integrated energy systems, GPU-accelerated optimization, nonlinear programming

1. INTRODUCTION

Global electricity demand is growing rapidly due to the electrification of industrial processes, the expansion of data centers, and the increasing penetration of electric transportation systems [1]. Achieving carbon neutrality under these trends requires a fundamental transformation of power systems toward renewable energy sources (RES). However, the inherent intermittency and uncertainty of RES introduce significant temporal mismatches between electricity supply and demand, posing critical challenges to secure and efficient system operation.

Addressing these challenges necessitates flexible energy storage solutions that enable sector coupling across multiple energy domains. Among the available options, Power-to-Hydrogen (P2H) technology offers a promising solution by converting surplus electricity into hydrogen, enabling long-duration storage and cross-sectoral decarbonization [2].

Despite the potential of integrated systems, existing optimization frameworks often rely on significant simplifications to manage computational complexity. On the grid side, many studies utilize Direct Current (DC) power flow approximations, which ignore voltage limits and reactive power requirements [3]. On the hydrogen side, electrolyzers are frequently represented as idealized, controllable electrical loads [4]. Such formulations neglect critical physical characteristics of hydrogen systems, including the nonlinear auxiliary power

consumption of balance-of-plant components (e.g., compressors and cooling systems) and the thermodynamic dynamics of hydrogen storage vessels, such as pressure and temperature evolution. Ignoring these effects can lead to inaccurate assessments of system flexibility and suboptimal operational decisions.

To address these limitations, this paper proposes a high-fidelity, integrated optimization framework that couples a multi-period Alternating Current Optimal Power Flow (AC-OPF) model with detailed hydrogen production and storage dynamics. The AC-OPF formulation explicitly captures voltage magnitudes, active and reactive power flows, and transmission losses, while the hydrogen subsystem is modeled with rigorous operational and thermodynamic constraints. This integrated formulation enables a physically consistent representation of the coupled power–hydrogen system across multiple time periods.

The resulting optimization problem is a large-scale, nonlinear program (NLP) characterized by strong temporal coupling and complex network constraints. Solving such problems using conventional CPU-based nonlinear solvers becomes computationally prohibitive for realistically sized systems, particularly at the national scale.

To overcome this computational challenge, the proposed framework is implemented using ExaModels.jl, a GPU-based algebraic modeling system, and solved with MadNLP.jl, a GPU-accelerated nonlinear solver [5]. By leveraging the massive parallelism of modern graphics processing units (GPUs), this computational strategy enables the efficient solution of large-scale, high-fidelity integrated energy system models that were previously intractable.

The proposed approach is validated through case studies ranging from regional to national-scale power networks. The results demonstrate both the operational importance of detailed hydrogen process modeling and the substantial computational advantages of GPU-accelerated nonlinear optimization.

The remainder of this paper is organized as follows. Section 2 presents the mathematical formulation of the multi-period AC-OPF and hydrogen system models. Section 3 describes the GPU-accelerated computational framework. Section 4 reports numerical results and performance analyses, and Section 5 concludes the paper.

2. MATHEMATICAL MODELING

This section presents the mathematical formulation of the proposed multi-period AC-OPF problem with integrated hydrogen production and storage. Figure 1 illustrates the overall structure of the coupled power-hydrogen system, including the electrical network, electrolyzer, compressor, cooling system, and hydrogen storage tank, as well as their physical and operational interactions.

The proposed formulation extends the conventional AC power flow representation [6] by explicitly incorporating the inter-temporal dynamics and operational constraints of hydrogen subsystems. The model is implemented using ExaModelsPower.jl, an open-source modeling library for creating GPU-compatible nonlinear AC optimal power flow models [7].

The integrated system is formulated as a large-scale NLP problem over a discrete time horizon $t \in \mathcal{T}$. Building upon [8], the proposed formulation augments the basic power-to-hydrogen model with rigorous thermodynamic constraints for compression and cooling. At selected buses $i \in \mathcal{H} \subseteq \mathcal{N}$, hydrogen production and storage facilities are installed and electrically coupled to the power grid. Each hydrogen facility consists of a proton exchange membrane (PEM) electrolyzer, a compressor for pressurized hydrogen storage, and a cooling system to maintain safe operating temperatures. The coupling between the power grid and the hydrogen subsystem is realized through the electrical power consumption of these components, which directly enters the nodal power balance equations of the grid.

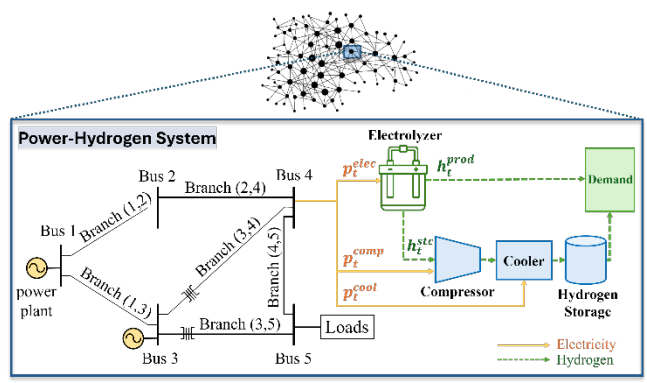

Figure 1. Schematic of the integrated power–hydrogen system. Selected buses in the transmission network are equipped with hydrogen production facilities, each comprising a PEM electrolyzer, compressor, and cooler. The electrical power consumption of these components the nodal power balance equations of the grid.

2.1 Multi-period AC-OPF with Integrated Hydrogen Loads

The core of the integrated energy system is the nonlinear AC power flow model, which ensures the physical feasibility of power system operation. For each bus $i \in \mathcal{N}$ and time period $t \in \mathcal{T}$, the nodal active and reactive power balance equations are formulated as

$$\sum_{g \in \mathcal{G}_i} p_{i,t}^g - p_{i,t}^d - p_{i,t}^{H2} = \sum_{j \in \mathcal{N}_i} p_{ij,t}, \forall i \in \mathcal{N}, t \in \mathcal{T} \tag{1}$$

$$\sum_{g \in \mathcal{G}_i} q_{i,t}^g - q_{i,t}^d = \sum_{j \in \mathcal{N}_i} q_{ij,t}, \forall i \in \mathcal{N}, t \in \mathcal{T} \tag{2}$$

where $p_{i,t}^g$ and $q_{i,t}^g$ denote the active and reactive power

generation, and $p_{i,t}^{d}$ and $q_{i,t}^{d}$ represent the fixed active and reactive power demands, respectively. The term $p_{i,t}^{H2}$ denotes the total active power consumption of the hydrogen facility. This term is non-zero only for the subset of buses equipped with hydrogen systems, $i \in \mathcal{H} \subseteq \mathcal{N}$, and is decomposed as

$$p_{i,t}^{H2} = p_{i,t}^{elec} + p_{i,t}^{comp} + p_{i,t}^{cool}, \tag{3}$$

where $p_{i,t}^{elec}, p_{i,t}^{comp}$, and $p_{i,t}^{cool}$ denote the electrical power consumption of the PEM electrolyzer, hydrogen compressor, and cooling system, respectively.

The branch active and reactive power flows between buses i and j are defined using the standard polar form of the AC power flow equations as

$$p_{ij,t} = G_{ij}V_{i,t}^2 - V_{i,t}V_{j,t}(G_{ij}cos\theta_{ij,t} + B_{ij}sin\theta_{ij,t}), \tag{4}$$

$$q_{ij,t} = -B_{ij}V_{i,t}^2 - V_{i,t}V_{j,t}(B_{ij}cos\theta_{ij,t} + G_{ij}sin\theta_{ij,t}), \tag{5}$$

where $V_{i,t}$ and $\theta_{i,t}$ are the voltage magnitude and angle at bus i, and G_{ij} and B_{ij} are the elements of the bus admittance matrix. Operational security is ensured by enforcing voltage magnitude limits, generator output limits, and branch flow limits.

2.2 Hydrogen Production Model

For notational simplicity, the bus index i is omitted in the hydrogen subsystem equations. All hydrogen-related variables are defined for each hydrogen-equipped bus $i \in \mathcal{H}$ and time period $t \in \mathcal{T}$, unless otherwise stated. The hydrogen production system is based on a PEM electrolyzer, modeled based on electrochemical principles described in [9]. The hydrogen generation rate is modeled as a function of the electrolyzer current. The hydrogen production rate h_t^{prod} is given by

$$h_t^{prod} = \frac{3600 \cdot N_c \cdot A \cdot M_{H2} \cdot i_t \cdot \eta_F(i_t)}{2F}, \tag{6}$$

where N_c is the number of cells, A is the electrode active area, M_{H_2} is the molar mass of hydrogen, i_t is the current density, F is Faraday's constant, and $\eta_F(i_t)$ is the Faradaic efficiency.

The Faradaic efficiency is modeled as a logarithmic function of the current density following,

$$\eta_F(i_t) = 0.98 - 0.02ln(i_t), \tag{7}$$

where the current density i_t is restricted to the practical operating range $0.5 \le i_t \le 2.5\,\mathrm{A/cm^2}$.

The electrical power consumption of the electrolyzer is modeled as

$$p_t^{elec} = V_c(i_t) \cdot i_t \cdot A \cdot N_c, \tag{8}$$

where $V_c(i_t)$ is the voltage of a single cell, which depends nonlinearly on the operating current density. The electrochemical parameters for cell voltage are adopted from [10].

2.3 Compressor and Cooling Power Models

Hydrogen produced by the electrolyzer is compressed prior to injection into the storage tank. The compressor power consumption is modeled based on isentropic compression:

$$p_t^{comp} = \frac{\gamma R T_{elec}}{(\gamma - 1)\eta_{comp}} \bullet n_t^{stc} \bullet \left[\left(\frac{P_t^{comp,out}}{P_{H_2}}\right)^{\frac{\gamma-1}{\gamma}} - 1\right], \tag{9}$$

where γ is the specific heat ratio, R is the gas constant, T_{elec} is the electrolyzer outlet temperature, η_{comp} is the compressor efficiency, n_t^{stc} is the molar flow rate entering the storage tank, and $P_t^{comp,\,out}$ and P_{H_2} denote the compressor outlet and inlet pressures, respectively.

The outlet temperature of the compressor is given by

$$T_t^{comp,out} = T_{elec} + \frac{T_{elec}}{\eta_{comp}}\left[\left(\frac{P_t^{comp,out}}{P_{H_2}}\right)^{\frac{\gamma-1}{\gamma}} - 1\right], \tag{10}$$

The cooling system removes excess thermal energy to maintain safe tank temperatures. The cooling power is derived from a simple energy balance:

$$p_t^{cool} = \frac{Q}{COP} = \frac{h_t^{stc} \bullet c_p \bullet \Delta T}{COP}, \tag{11}$$

where c_p is the specific heat capacity, ΔT is the required temperature reduction, and COP denotes the coefficient of performance of the cooling system.

2.4 Hydrogen Storage Dynamics and Thermodynamic Constraints

The hydrogen storage tank is modeled using dynamic mass and energy balance equations. The hydrogen inventory evolves as

$$h_t^{st} = h_{t-1}^{st} + h_t^{stc} - h_t^{std}, \tag{12}$$

where h_t^{stc} and h_t^{std} denote hydrogen charged into and discharged from the storage tank, respectively. The hydrogen inventory in the storage tank is bounded by its minimum and maximum storage capacities:

$$S^{min} \le h_t^{st} \le S^{max}, \tag{13}$$

The tank temperature dynamics are governed by the internal energy balance:

$$\begin{aligned} c_v \bullet n_t^{st} \bullet T_t^{tank} &= c_v \bullet n_{t-1}^{st} \bullet T_{t-1}^{tank} \\ &+ c_p \bullet n_t^{stc} \bullet T_{tank,in} - c_p \bullet n_t^{std} \bullet T_t^{tank}, \end{aligned} \tag{14}$$

where c_v is the constant-volume specific heat, n_t^{st} is the molar quantity stored in the tank, and $T_{tank,in}$ denotes the inlet temperature after cooling.

The tank pressure is determined by the ideal gas law:

$$P_t^{tank} = \frac{n_t^{st} R T_t^{tank}}{V_{tank}}, \tag{15}$$

where V_{tank} denotes the storage volume.

To ensure operational safety and compatibility with downstream systems, the compressor outlet pressure satisfies

$$P_t^{comp,out} = P_t^{tank} + \Delta P_{margin}, \tag{16}$$

where ΔP_{margin} represents the pressure margin required for reliable injection into the storage vessel.

3. GPU-BASED COMPUTATIONAL FRAMEWORK

This section presents the computational framework developed to solve the proposed large-scale, multi-period nonlinear AC-OPF problem with integrated hydrogen production and storage. The framework adopts a fully GPU-native optimization pipeline based on ExaModels.jl for algebraic modeling, MadNLP.jl for nonlinear programming, and GPU-accelerated sparse linear solvers for Newton step computation. This end-to-end GPU execution enables scalable solution of tightly coupled, high-dimensional energy system optimization problems.

3.1 Computational Challenges in Large-Scale Nonlinear Optimal Power–Hydrogen Flow

The integrated power–hydrogen optimization problem formulated in Section 2 can be expressed as a large-scale NLP problem:

$$\begin{aligned} &\min_x f(x) \\ &\text{s.t.}\, g(x) = 0, \\ &\quad h(x) \le 0 \end{aligned}$$

where x comprises generator outputs, bus voltages, branch flows, and hydrogen system states (e.g., current density, tank pressure, temperature, and stored mass) across all buses and time periods.

Dimensionality and inter-temporal coupling: The problem dimension scales with both network size and time horizon length. The incorporation of hydrogen dynamics introduces additional state variables and constraints with strong inter-temporal coupling, substantially increasing memory requirements and computational complexity. For large-scale instances, this leads to NLPs with tens of millions of variables and constraints.

Nonlinearity and nonconvexity: The AC power flow equations introduce strong nonlinearities and nonconvexity, necessitating the use of Newton-type iterative methods to achieve reliable convergence. Each iteration requires solving a Karush–Kuhn–Tucker (KKT) system:

$$\begin{bmatrix} \nabla^2 L(x,\lambda) & J(x)^T \\ J(x) & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta \lambda \end{bmatrix} = -\begin{bmatrix} \nabla L(x,\lambda) \\ g(x) \end{bmatrix}$$

where $L(x,\lambda)$ is the Lagrangian and $J(x)$ is the Jacobian of constraints.

These systems are typically ill-conditioned due to wide variable scaling, tight operational constraints, and strong coupling between electrical and thermodynamic subsystems.

Large-scale sparse linear algebra: On conventional CPU architectures, sparse factorization of these systems dominates the computational cost, severely limiting scalability. These challenges motivate the development of a GPU-based computational framework capable of efficiently handling large-scale nonlinear optimization problems.

3.2 Algebraic Modeling Using ExaModels.jl

The proposed formulation is implemented using ExaModels.jl, which provides a GPU-oriented abstraction for algebraic modeling and parallel automatic differentiation. All nonlinear objective and constraint functions are formulated within this algebraic framework, enabling the automatic generation of first- and second-order derivative information required by Newton-type optimization methods.

ExaModels employs a single-instruction multiple-data (SIMD) parallel abstraction that maps residual evaluation, Jacobian assembly, and Hessian computation directly onto GPU threads. This design enables highly parallel evaluation of derivative information and significantly accelerates the linearization steps performed at each Newton iteration [5].

The hydrogen production and storage models introduced in Section 2 are incorporated into this algebraic framework, enabling a unified representation of the electrical and thermodynamic subsystems within a single optimization model. This unified formulation preserves consistency across coupled physical domains and facilitates efficient execution on GPU architectures.

3.3 GPU-Accelerated Nonlinear Optimization with MadNLP.jl

The resulting nonlinear programming problem is solved using MadNLP.jl, a high-performance primal–dual interior-point solver designed for large-scale and GPU-accelerated nonlinear optimization. At each iteration, MadNLP solves a structured Newton system derived from the KKT optimality conditions.

A central challenge in implementing interior-point methods on GPUs lies in solving the associated KKT linear systems. In standard formulations, the KKT matrix is indefinite and requires sparse symmetric factorization with numerical pivoting to ensure numerical stability. However, pivoting introduces data-dependent branching

and irregular memory access patterns, which are poorly suited for massively parallel GPU architectures. Moreover, as the interior-point iterations approach the solution, barrier terms associated with active constraints lead to increasingly ill-conditioned KKT systems, making purely iterative solvers difficult to deploy robustly.

To address these challenges, MadNLP adopts condensed-space KKT reformulation strategies that transform the original indefinite KKT system into symmetric positive definite (SPD) forms, thereby enabling the use of pivoting-free Cholesky factorizations that are well suited for GPU execution. Two complementary strategies are employed.

MadNLP employs condensed-space KKT reformulations — including Lifted KKT and HyKKT strategies — that transform the indefinite KKT system into symmetric positive definite form, enabling pivoting-free Cholesky factorization suitable for GPU execution [11, 12].

These reformulations allow MadNLP to achieve robust and scalable convergence for highly ill-conditioned multi-period optimization problems with substantial performance improvements over conventional CPU-based implementations.

3.4 GPU-Based Newton Step Computation and Sparse Linear Solvers

The condensed-space KKT reformulations described in Section 3.3 produce SPD systems, which can be efficiently factorized using Cholesky-based sparse direct solvers such as NVIDIA's cuDSS library [13].

Since the condensed KKT matrices are positive definite, cuDSS performs Cholesky factorization without numerical pivoting, avoiding the data-dependent branching that hinders parallel execution. The factorization consists of a symbolic analysis phase — performed only once since the sparsity structure remains fixed across interior-point iterations — and a numerical factorization phase that is reused at each iteration.

The computational workflow proceeds as follows: sparse Jacobian and Hessian matrices are evaluated in parallel using ExaModels and assembled directly into GPU-resident memory, the KKT matrix is factorized using cuDSS, and the Newton direction is recovered via GPU-accelerated triangular solves. Throughout this process, all data remains GPU-resident, minimizing host–device data transfers.

4. CASE STUDY

This section presents numerical experiments to evaluate the computational performance of the GPU-based framework and to analyze the operational characteristics of the integrated power–hydrogen system. The results demonstrate the scalability of the proposed approach and reveal the operational trade-offs arising from rigorous thermodynamic modeling of hydrogen production and storage.

4.1 Experimental Setup

This subsection outlines the optimization problem formulation, benchmark networks, hydrogen system modeling assumptions, and computational setup adopted in the case study. These components collectively define the experimental environment for evaluating both computational performance and operational behavior.

4.1.1 Optimization Problem

The objective of the case study is to minimize total generation cost over a multi-period planning horizon, subject to electrical and hydrogen operational constraints. The objective function is defined as:

$$\min \sum_{t\in\mathcal{T}} \sum_{i\in\mathcal{N}} \left(a_i (p^g_{i,t})^2 + b_i p^g_{i,t} + c_i \right) \tag{17}$$

where $p^g_{i,t}$ denotes the active power output of generator i, and a_i, b_i, c_i are the corresponding cost coefficients. This objective is minimized subject to the AC power flow constraints (Eqs. 1-5) and the hydrogen system constraints (Eqs. 6-16) detailed in Section 2.

Additionally, a fixed hydrogen demand profile is imposed at each hydrogen bus to represent downstream industrial consumption requirements:

$$h^{prod}_t - h^{stc}_t + h^{std}_t \geq D_t \tag{18}$$

where D_t is the hydrogen demand at time t.

4.1.2 Hydrogen System Specifications

The parameters for the hydrogen subsystems are listed in Table 1. A nominal electrolyzer configuration with a maximum hydrogen production capacity of approximately 1000 kg/h is assumed at each installation site. The storage tank is sized to buffer fluctuations over the planning horizon.

Table 1. Hydrogen production and storage system specifications.

Parameter		Value
Number of cells	N_c	5000
Active cell area	A	2000 cm^2
Hydrogen production capacity	h^{prod}_{max}	1000 kg/h
Operating temperature	T_{elec}	323 K
Hydrogen outlet pressure	P_{H_2}	30 bar
Tank volume	V_{tank}	25 m^3
Operating pressure range	P^{tank}	0-700 bar
Operating temperature range	T^{tank}	270-375 K

4.1.3 Modeling Configurations

Two modeling configurations are considered to assess the impact of thermodynamic detail on both solution accuracy and computational requirements.

The isothermal model assumes constant temperature and pressure throughout the hydrogen process. Adiabatic compression work is approximated using a simplified linear correlation [14], and storage thermal dynamics are neglected. This simplification reduces problem dimensionality and significantly improves computational efficiency.

The thermodynamic model incorporates the full rigorous formulation presented in Section 2. This includes isentropic compression (Eq. 9-10), cooling power requirements (Eq. 11), storage mass balance (Eq. 12), and temperature dynamics via the energy balance (Eq. 14). Tank pressure is determined by the ideal gas law (Eq. 15). This model captures the operational flexibility limits and auxiliary power consumption of the integrated system.

4.1.4 Test Networks and Problem Scale

We evaluate the framework on seven benchmark networks from the PGLib library [15], ranging from the 30-bus IEEE system to a 9591-bus synthetic grid. To simulate a wide deployment of P2H technology, hydrogen production facilities (electrolyzer, compressor, storage) are assumed to be colocated with all active power generators. This placement strategy is motivated by the practical consideration that hydrogen production facilities require access to substantial electrical infrastructure, which is typically available at generation sites.

Table 2. Problem dimensions for benchmark cases with integrated hydrogen systems (60-min resolution). k=10^3, m=10^6

Case	Periods	Isothermal		Thermodynamic	
		nvars	ncons	nvars	ncons
case30	24	6.0k	8.9k	6.3k	9.1k
case118	24	29.8k	41.8k	32.0k	44.1k
case300	24	68.1k	96.0k	75.0k	102.8k
case793	24	164.0k	227.3k	185.0k	248.3k
case2312	24	490.6k	701.9k	540.2k	751.6k
case4020	24	897.8k	1.4m	908.5k	1.4m
case9591	24	2.0m	3.2m	2.1m	3.2m
case4020	168	6.3m	9.7m	6.4m	9.8m
case9591	168	14.3m	22.2m	14.4m	22.4m

Table 2 summarizes the problem dimensions for each test case under both modeling configurations. We consider two planning horizons: a standard 24-hour day (24 periods) and a week-long 168-hour horizon (168 periods), both with hourly resolution. The inclusion of thermodynamic variables and constraints significantly increases the problem size. For the largest instance (*case9591*, 168h), the thermodynamic formulation involves over 14.4 million variables and 22.4 million constraints, a scale that poses significant challenges for conventional CPU-based solvers.

4.1.5 Computational Platform

All GPU experiments are conducted on an NVIDIA H200 GPU using MadNLP v0.8.6 as the nonlinear programming solver and cuDSS v0.4.0 for sparse Cholesky factorization. For comparison, CPU-based solvers including MadNLP with UMFPACK and Ipopt with MUMPS are evaluated on a dual Intel Xeon 8480C system (112 cores total, 3.8 GHz max turbo).

4.2 Computational Performance

We evaluate the computational efficiency and scalability of the proposed GPU-native optimization framework (MadNLP-GPU) against two CPU-based alternatives. Figure 2 reports the wall-clock solution times for all benchmark networks under both the isothermal and thermodynamic formulations.

For the isothermal formulation, both CPU-based solvers exhibit reliable convergence on small- to medium-scale instances (up to *case793*). However, their performance degrades rapidly as network size and planning horizon increase. Ipopt successfully solves the 24-hour *case9591* instance in approximately 4, 200 seconds, while MadNLP-CPU fails on instances larger than *case793*. In contrast, MadNLP-GPU demonstrates consistently strong performance across all test cases, with solution times ranging from less than 1 second for *case30* to approximately 470 seconds for the largest instance (*case9591*, 168-hour), highlighting the excellent scalability of the GPU-based approach.

The thermodynamic formulation introduces tight bounds on temperature and pressure states, as well as strong nonlinear coupling between electrical and thermodynamic variables. These factors, combined with the wide range of variable scales, significantly degrade the numerical conditioning of the KKT system, consistent with the conditioning challenges discussed in Section 3, and pose severe challenges for CPU-based solvers.

Under this formulation, Ipopt fails to converge for all benchmark instances, and MadNLP-CPU solves only up to *case793*, requiring approximately 3 hours. By contrast, MadNLP-GPU maintains robust convergence across the full test set, solving the largest instance (14.4 million variables, 22.4 million constraints) in approximately 6, 650 seconds.

On instances where both GPU and CPU solvers converge, the GPU achieves speedups of 10-600× depending on problem size and model complexity. More importantly, the GPU-based framework extends the tractable problem size: from *case2312* onward, CPU-based solvers fail, whereas MadNLP-GPU scales efficiently to problems with millions of variables. These results confirm that the proposed framework enables high-fidelity modeling of integrated power-hydrogen systems at scales previously inaccessible to conventional approaches.

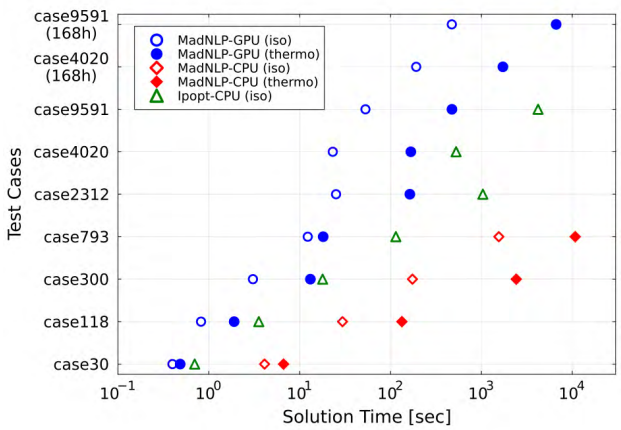

Figure 2. Wall-clock solution times for GPU and CPU solver configurations. Hollow and filled markers denote isothermal and thermodynamic formulations, respectively. GPU refers to MadNLP-GPU. Missing data points indicate solver failure.

4.3 Hydrogen System Operational Analysis

To illustrate the operational behavior captured by the proposed framework, Figure 3 presents the optimal 24-hour dispatch for a hydrogen production facility at bus 69 of the *case118* network. This bus is connected to a generator and is well-connected to neighboring buses, providing sufficient electrical capacity for hydrogen production. The optimization minimizes total generation cost while satisfying a fixed hydrogen demand of 42 kg/h.

Figure 3(a) shows the hydrogen production rate alongside the baseline electrical load profile. Production varies between 39 and 45 kg/h, exhibiting an inverse correlation with system load. During off-peak hours (hours 1–8), the electrolyzer operates above the required demand at 43–45 kg/h, with the surplus directed to storage. As system load increases (hours 9–10 and 16–23), production decreases to 39–41 kg/h, falling below the fixed demand. This pattern is consistent with both generation cost minimization and network capacity constraints.

Figure 3(b) illustrates the storage dynamics. The tank inventory follows a diurnal cycle, increasing from 50 kg to 60 kg during low-load periods, then decreasing as stored hydrogen compensates for reduced production during high-load periods. The storage thus decouples instantaneous production from the fixed demand.

Figure 3(c) presents the power consumption breakdown. The electrolyzer dominates at 1.86–2.04 MW, exhibiting approximately 9% variation over the horizon. Auxiliary equipment (compressor and intercooler) consumes less than 1 kW under the moderate pressure ratio (30 to 60 bar) considered in this study, representing a

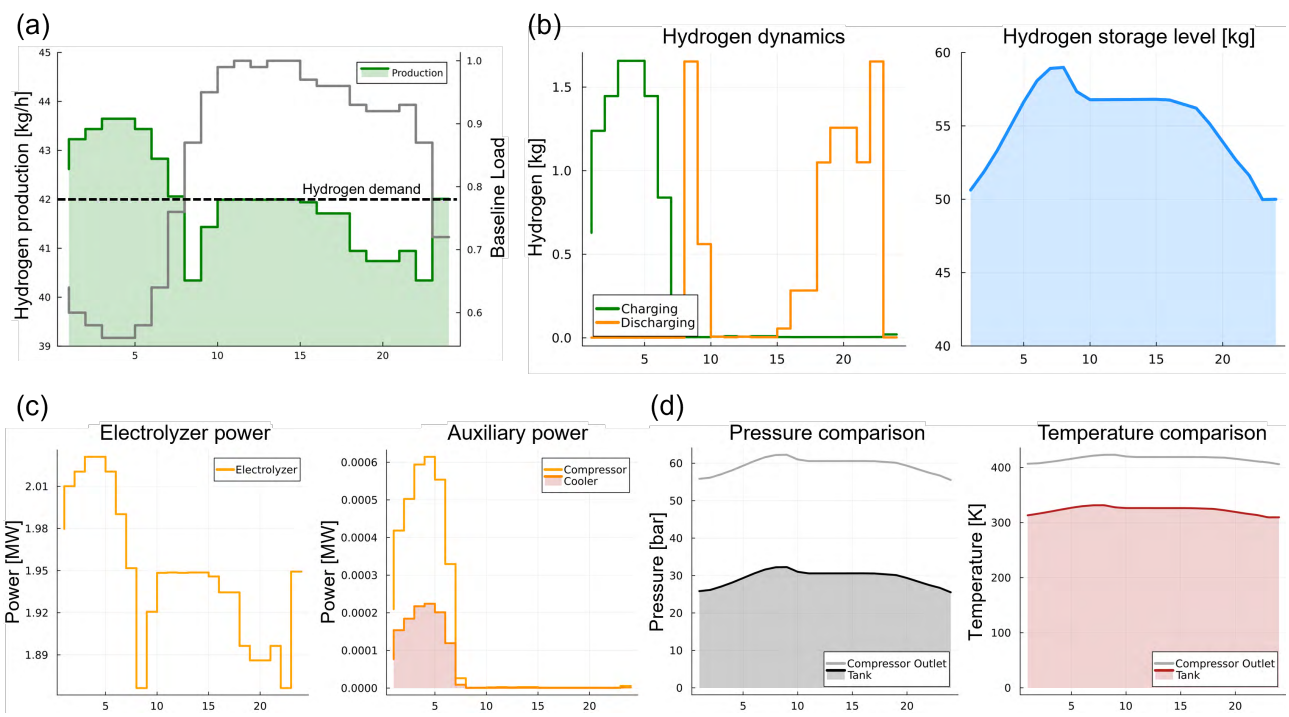

Figure 3. 24-hour operational profiles at bus 69 for the IEEE 118-bus system: (a) hydrogen production rate and electrical baseline load with fixed hydrogen demand, (b) storage charging/ discharging dynamics and inventory level, (c) electrolyzer and auxiliary (compressor and cooler) power consumption, and (d) compressor outlet and storage tank pressure/temperature profiles.

negligible fraction of total consumption. For higher storage pressures (350–700 bar), compression power would become significant, and the detailed thermodynamic model would be essential.

Figure 3(d) shows pressure and temperature profiles. The compressor outlet reaches pressures above 60 bar and temperatures exceeding 400 K during active charging. Following intercooling, the storage tank operates at approximately 30 bar and 300–320 K, within the specified limits (0–700 bar, 270–375 K).

The compressor outlet temperature, exceeding 400 K, must be managed through intercooling to maintain safe storage conditions. Isothermal models would not capture these thermal requirements, potentially yielding schedules that violate equipment limits. This justifies the thermodynamically consistent formulation proposed in this work, demonstrating the practical value of high-fidelity modeling for integrated power–hydrogen system operation.

5. CONCLUSION

This paper proposed and demonstrated a GPU-accelerated framework for multi-period AC-OPF with integrated hydrogen systems, incorporating high-fidelity thermodynamic constraints for electrolyzer, compressor, cooling, and storage operations. The GPU-based solver achieved 10–600× speedups over CPU alternatives, enabling the solution of problems with over 14 million variables, while CPU-based solvers failed to scale beyond medium-sized instances. Operational analysis on the IEEE 118-bus system demonstrated that hydrogen production is inversely correlated with electrical load, with storage buffering production and demand. The thermodynamic model captured critical constraints—such as compressor outlet temperatures exceeding 400 K—that simplified isothermal formulations fail to capture.

Future work will extend this framework to optimal placement of hydrogen facilities, large-scale integration of renewable energy sources, and CO_2 emission mapping for decarbonization-oriented power–hydrogen system planning.

ACKNOWLEDGEMENTS

This research was supported by Korea Institute for Advancement of Technology(KIAT) grant funded by the Ministry of Trade, Industry & Energy(MOTIE), Korea Government (RS-2024-00436106, Human Resource Development Program for Industrial Innovation

REFERENCES

1. Li X, Lepour D, Heymann F, Maréchal F. Electrification and digitalization effects on sectoral energy demand and consumption: a prospective study towards 2050. Energy 279:127992 (2023). https://doi.org/10.1016/j.energy.2023.127992
2. Shiraishi K, Park WY, Kammen DM. The role of hydrogen as long-duration energy storage and as an international energy carrier for electricity sector decarbonization. Environ. Res. Lett. 19:084011 (2024). https://doi.org/10.1088/1748-9326/ad5856
3. Singh A, Kumar A, Chinmaya KA, Maulik A. Optimal operation of an electricity-hydrogen DC microgrid with integrated demand response. Sustainable Energy, Grids and Networks 39:101451 (2024). https://doi.org/10.1016/j.segan.2024.101451
4. Martínez L, Fernández D, Mantz R. Two layer control strategy of an island DC microgrid with hydrogen storage system. International Journal of Hydrogen Energy 50:365-378 (2024). https://doi.org/10.1016/j.ijhydene.2023.09.009
5. Shin S, Anitescu M, Pacaud F. Accelerating optimal power flow with gpus: SIMD abstraction of nonlinear programs and condensed-space interior-point methods. Electric Power Systems Research 236:110651 (2024). https://doi.org/10.1016/j.epsr.2024.110651
6. G-L F, Desforges R. Limoges. essai d'analyse demographique, economique et sociale (etude sommaire). Population (French Edition) 17:806 (1962). https://doi.org/10.2307/1526309
7. Johnson, Sanjay, et al. "ExaModelsPower. jl: A GPU-Compatible Modeling Library for Nonlinear Power System Optimization." *arXiv preprint arXiv:2510.12897* (2025). https://doi.org/10.48550/arXiv.2510.12897
8. Baumhof MT, Raheli E, Johnsen AG, Kazempour J. Optimization of hybrid power plants: when is a detailed electrolyzer model necessary?. 2023 IEEE Belgrade PowerTech :1-10 (2023). https://doi.org/10.1109/powertech55446.2023.10202860
9. Falcão DS, Pinto AMFR. A review on PEM electrolyzer modelling: guidelines for beginners. Journal of Cleaner Production 261:121184 (2020). https://doi.org/10.1016/j.jclepro.2020.121184
10. Yigit T, Selamet OF. Mathematical modeling and dynamic simulink simulation of high-pressure PEM electrolyzer system. International Journal of Hydrogen Energy 41:13901-13914 (2016). https://doi.org/10.1016/j.ijhydene.2016.06.022
11. Pacaud, François, et al. "Condensed-space methods for nonlinear programming on GPUs." *arXiv preprint arXiv:2405.14236* (2024). https://doi.org/10.48550/arXiv.2405.14236
12. Regev, Shaked, et al. "A hybrid direct-iterative method for solving kkt linear systems." *arXiv*

preprint arXiv:2110.03636 (2021). https://doi.org/10.48550/arXiv.2110.03636
13. NVIDIA, cuDSS: CUDA Direct Sparse Solver Library, https://developer.nvidia.com/cudss
14. Zheng Y, You S, Bindner HW, Münster M. Optimal day-ahead dispatch of an alkaline electrolyser system concerning thermal–electric properties and state-transitional dynamics. Applied Energy 307:118091 (2022). https://doi.org/10.1016/j.apenergy.2021.118091
15. Babaeinejadsarookolaee, Sogol, et al. "The power grid library for benchmarking ac optimal power flow algorithms." *arXiv preprint arXiv:1908.02788* (2019). https://doi.org/10.48550/arXiv.1908.02788

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Foundation Model-Guided Optimization of Chemical Reaction Spaces for Autonomous Experimentation

Youhyun Kim[abc] and Jonggeol Na[abc]*

[a] Ewha Womans University, Department of Chemical Engineering and Materials Science, Seoul 03760, Republic of Korea
[b] Ewha Womans University, Graduate Program in System Health Science and Engineering, Seoul 03760, Republic of Korea
[c] Ewha Womans University, Institute for Multiscale Matter and Systems (IMMS), Seoul 03760, Republic of Korea
* Corresponding Author: jgna@ewha.ac.kr.

ABSTRACT

The optimization of chemical reactions requires navigating a high-dimensional design space composed of both discrete and continuous variables. Although one-hot encoding has been widely adopted, it lacks chemically meaningful information and suffers from sparsity and poor generalization. To address these limitations, we explored the use of pretrained molecular foundation models to generate latent representations as input variables for optimization. However, rigorously comparing different combinations of reaction representations and optimization algorithms remains a time- and resource-intensive challenge. In this work, we developed an end-to-end benchmarking platform that systematically evaluates diverse encoding schemes and optimization strategies under identical conditions. The platform automates the entire workflow from data preprocessing to result analysis, supporting fair comparison across multiple representation–optimizer combinations. Furthermore, we designed a custom reaction representation that integrates a 3D equivariant encoder with a bidirectional cross-attention module to explicitly capture interactions between reaction components. The proposed platform provides a scalable foundation for reaction optimization and advances the feasibility of autonomous experimental systems.

Keywords: Reaction optimization, Molecular representation, Black-box optimization, Autonomous experimentation, Benchmarking platform

INTRODUCTION

Organic chemical reactions underpin key processes across fine chemicals, pharmaceutical manufacturing, and the synthesis of functional materials. Achieving economic and environmentally responsible production of target molecules under constrained time and resource budgets therefore requires precise optimization of reaction conditions. In strongly condition-dependent, multivariable systems—such as metal-catalyzed cross-coupling and asymmetric catalysis—reaction yield and selectivity can vary dramatically with the choice and combination of substrates, ligands, catalysts, solvents, temperature, and concentration [1]. However, the reaction-condition space comprises mixed discrete and continuous variables, resulting in a highly combinatorial search space. This makes trial-and-error experimentation or limited design-of-experiments protocols insufficient for efficiently identifying globally optimal conditions [2].

Enabled by AI-accelerated experimental planning and high-throughput automation, the paradigm of self-driving laboratories (SDLs) has rapidly emerged [3]. Within SDL frameworks, the mapping from experimental conditions to observed reaction outcomes is typically evaluated in a closed loop of experimentation, and condition selection is guided by optimization strategies. Despite these advances, the effectiveness of these approaches depends critically on how reaction conditions are represented. Widely used one-hot encodings (OHE) induce sparse search spaces that fail to reflect chemical similarity [4], while molecular fingerprints, being hash-based and discrete, can miss three-dimensional geometry and subtle variations of the local chemical environment, and distances in fingerprint space often correlate poorly with chemically meaningful similarity [5]. These limitations motivate the reformulation of the search space through chemically informed representations that enable the extraction of meaningful features.

Recent advances in deep learning have substantially improved data-driven prediction of reaction yield, particularly with graph neural networks and attention-based models [6, 7]. However, comparatively limited attention has been paid to the role of learned latent representations in shaping the optimization landscape and their geometric influence on search efficiency in practical optimization loops . This gap suggests that reaction representation learning and optimization strategy design should not be treated as independent problems, but rather as a coupled, end-to-end design challenge.

Here, we propose an integrated framework that unifies AI-driven reaction representations with optimization for organic synthesis. Using pretrained foundation models, reaction components—substrates, ligands, catalysts, bases, and solvents—are embedded as continuous vectors. A chemically informed cross-attention mechanism is introduced to explicitly model interactions between the reaction center and surrounding reagents, establishing a meaningful continuous alternative to OHE. We couple this latent representation with a range of derivative-free optimizers and benchmark encoder–optimizer combinations on representative cross-coupling reactions. Finally, we investigate the relationship between latent-space structure and optimization performance, deriving practical design principles for representation-centric search spaces.

METHODOLOGY

Molecular representation models and embedding extraction

In this study, we benchmarked multiple pretrained molecular foundation models to transform discrete reaction-condition components into continuous vector representations. The following pretrained encoders were considered: MoLFORMER [8], KPGT [9], UMA [10], Frad [11], KERMT [12], and SMI-TED$_{289M}$ [13]. To enable a consistent comparison across encoders, all input molecules were standardized to canonical SMILES or reaction SMILES through a single preprocessing pipeline, and the same set of reaction components was provided to each encoder. All encoders were used in inference mode with frozen weights, and embeddings were extracted without fine-tuning. Each encoder outputs a fixed-length vector representation at the molecular level, and the encoders used along with their output dimensions are summarized in Table 1. Most encoders directly output molecule-level embeddings, which were used without modification. For UMA, which outputs atom/node-level representations, node embeddings were aggregated to yield a single molecular vector.

Table 1. Summary of molecular representation models used in this study.

Model	Core Architecture	Output Dimension
MoLFormer	Transformer encoder with linear attention	768
KPGT	Graph Transformer	2304
UMA	3D equivariant message passing (eSCN-based)	256
FRAD	Equivariant Transformer (TorchMD-Net)	256
KERMT	Graph Transformer	3200
SMI-TED$_{289M}$	Encoder-decoder Transformer	768

Reaction-Level Representation

The reaction vector was constructed by concatenating the component embeddings $\mathbf{e}_{c_i} \in \mathbb{R}^{d_i}$. Continuous process variables were scaled to [0, 1] via min–max normalization and included as additional inputs.

Since each encoder has different output embedding dimensions, we applied a VAE to project all embeddings to a common latent dimension of $d = 256$, preventing performance comparisons from being confounded by differences in dimensionality. The VAE loss function consists of KL divergence loss and component-wise reconstruction loss:

$$L_{\text{total}} = \lambda_{\text{recon}} \cdot L_{\text{recon}} + \lambda_{\text{KL}} \cdot L_{\text{KL}} \tag{1}$$

The reconstruction loss employs cross-entropy loss over the candidate embedding set for each component. For absent components, a learnable "NONE" embedding $\mathbf{n}_i$ was included in the candidate set to represent missing inputs as a meaningful class. The logits are then computed by scaling the Euclidean distance between the reconstructed component vector $\hat{\mathbf{e}}_i$ and the candidate set $\mathcal{C}_i = \left\{\mathbf{c}_i^{(1)}, \dots, \mathbf{c}_i^{(N_i)}, \mathbf{n}_i\right\}$ by temperature τ:

$$\text{logit}_{\text{k}} = -\frac{\left\|\hat{\boldsymbol{e}}_i - \boldsymbol{c}_i^{(k)}\right\|_2}{\tau} \tag{2}$$

Here, $\hat{\mathbf{e}}_i$ is the embedding vector of the i-th component reconstructed by the decoder, $\mathbf{c}_i^{(k)}$ is the embedding of the k-th candidate molecule, and $\mathbf{n}_i$is the learnable NONE embedding. The ground truth label y_i is the index of the candidate closest to the original embedding, with the NONE class assigned for missing components:

$$\boldsymbol{L_{recon}} = \frac{\mathbf{1}}{\boldsymbol{K}} \sum_{\boldsymbol{i=1}}^{\boldsymbol{K}} \boldsymbol{CrossEntropy}(\boldsymbol{softmax}(\boldsymbol{logit_i}), \boldsymbol{y_i}) \tag{3}$$

Designed Reaction Representation

The benchmarked molecular representation models

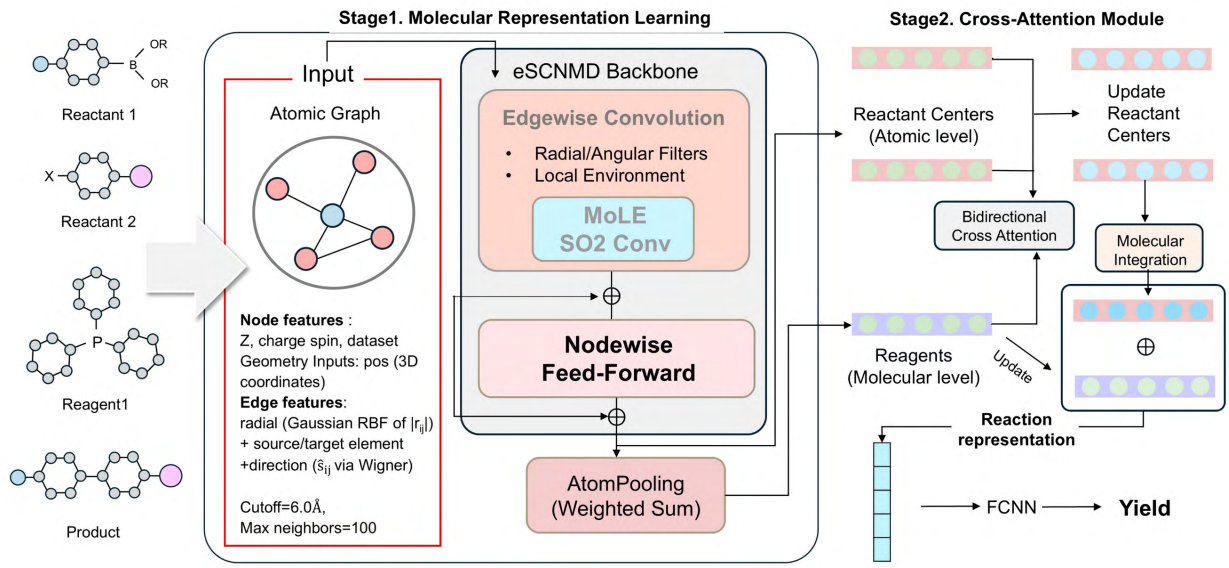

Figure 1. Architecture of the custom reaction representation model. (Left) Stage 1 extracts atomic and molecular embeddings using an eSEN backbone with eSCN convolution over 3D geometric features. (Right) Stage 2 employs a bidirectional cross-attention module that enables mutual information exchange between reaction centers and reagents.

encode each molecule independently and represent reactions through simple concatenation, followed by VAE-based dimensionality reduction. However, this approach does not explicitly capture chemical interactions between reaction components. Based on the benchmark results, UMA was selected as the backbone encoder due to its consistently competitive optimization performance. Moreover, its 3D equivariant architecture, which leverages both interatomic distances and directional information, is well suited to capturing the geometric features critical for stereoselectivity and coordination environments in chemical reactions. Building upon UMA, we integrated a bidirectional cross-attention module to explicitly capture inter-component interactions (Figure 1).

Cross-Attention Module

We employ a parallel bidirectional cross-attention mechanism between reaction centers and reagents. The two attention pathways are computed simultaneously with swapped query and key/value roles:

$$\boldsymbol{Q} = \boldsymbol{W_Q h_q};\ \boldsymbol{K} = \boldsymbol{W_K h_{kv}};\ \boldsymbol{V} = \boldsymbol{W_V h_{kv}} \tag{4}$$

$$\boldsymbol{Attention}(\boldsymbol{Q},\boldsymbol{K},\boldsymbol{V}) = \boldsymbol{softmax}\left(\frac{\mathbf{Q} \cdot \mathbf{K}^{\mathbf{T}}}{\sqrt{\mathbf{d_k}} \cdot \tau}\right) \cdot \boldsymbol{V} \tag{5}$$

$$\begin{aligned} \boldsymbol{h'_{center}} &= \boldsymbol{Attn}(\boldsymbol{h_{center}}, \boldsymbol{h_{reagent}}, \boldsymbol{h_{reagent}}) \\ \parallel\ \boldsymbol{h'_{reagent}} &= \boldsymbol{Attn}(\boldsymbol{h_{reagent}}, \boldsymbol{h_{center}}, \boldsymbol{h_{center}}) \end{aligned} \tag{6}$$

where $\mathbf{h}_{\text{center}}$ and $\mathbf{h}_{\text{reagent}}$ denote the embeddings of reaction center atoms and reagent molecules, respectively, and $\parallel$ denotes parallel computation. This bidirectional design enables mutual information exchange between reaction centers and reagents.

Problem Formulation

The reaction optimization problem is defined over a latent design space constructed from discrete reaction components. Each reaction condition is represented as a continuous vector $\mathbf{z} \in \mathbb{R}^d$ obtained by encoding the discrete components through a pretrained molecular encoder. The optimization problem is formulated as:

$$\boldsymbol{z}^* = \boldsymbol{arg} \max_{\boldsymbol{z} \in \mathbb{R}^d} \boldsymbol{f}(\boldsymbol{Decode}(\boldsymbol{z})) \tag{7}$$

At each iteration, the optimizer proposes a latent vector $z_t \in \mathbb{R}^d$, which is mapped to the nearest candidate condition via:

$$\mathbf{Decode}(\boldsymbol{z_t}) = \mathbf{arg}\ \min_{\mathbf{z_i} \in \mathcal{Z}} \parallel \mathbf{z_t} - \mathbf{z_i} \parallel_2 \tag{8}$$

where $\mathcal{Z} = \{z_i = \text{Encode}(c_i) \mid c_i \in \mathcal{C}\}$ is the set of encoded candidate reaction conditions. The corresponding yield $f(\text{Decode}(z_t))$ is then retrieved from the pre-existing experimental dataset and returned to the optimizer as the observed objective value.

Black-box Optimization Algorithms

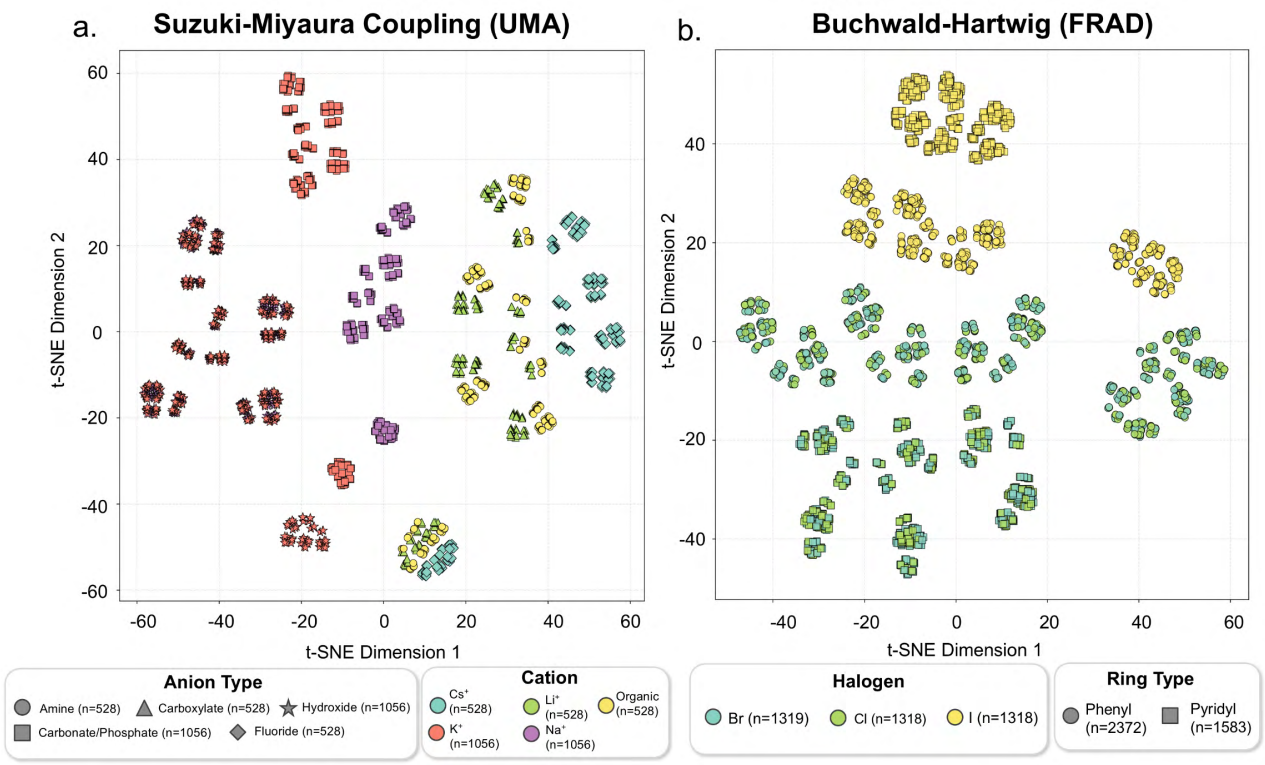

Figure 2. t-SNE visualization of reaction vectors for (a) Suzuki–Miyaura coupling encoded by UMA, where colors and marker shapes represent cation and anion types of the base, respectively, and (b) Buchwald–Hartwig amination encoded by FRAD, where colors and marker shapes represent halogen and aromatic ring types of the aryl halide, respectively.

To efficiently explore high-yield conditions within a limited experimental budget, we systematically benchmarked a variety of black-box optimization algorithms. In previous reaction optimization studies, Bayesian optimization (BO), which learns a surrogate model from observed data and proposes the next experimental conditions through an acquisition function, has been widely adopted. However, chemical reaction spaces are characterized by a high proportion of categorical/discrete variables, and the objective function is often non-smooth or multimodal, accompanied by experimental noise and constraints. These characteristics can conflict with the smoothness and stationarity assumptions underlying standard GP-based BO, potentially degrading its performance. Moreover, BO has frequently been reported to encounter challenges with high-dimensional bottlenecks, where exploration becomes more challenging as the number of dimensions increases [14]. Additionally, its computational demands grow with the accumulation of data, which restricts its scalability in closed-loop settings. Therefore, rather than relying on a specific algorithm, this study broadly compared and evaluated algorithms with different search mechanisms from the perspective that optimal strategies may vary depending on the structure of the reaction space and experimental scenarios. The optimization algorithms used in this study, along with their classification, are summarized in Table 2.

Table 2. Categories of black-box optimization algorithms benchmarked in this study.

Category	Algorithms
Bayesian Optimization	EDBO [2], Gryffin [15], BoTorch [16], ProcessOptimizer [17], ATLAS [18]
Evolutionary Algorithms	DEAP [19], PSO [20], CMA-ES [21], PADDY [22]
Direct Search Methods	NOMAD [23], DIRECT [24]
Surrogate-assisted DFO	BOBYQA [25], SNOBFIT [26], BAM [27], PWAS [28]

Benchmarking Platform

To enable quantitative comparison of reaction representation and optimization algorithm combinations under identical conditions, we developed an end-to-end benchmarking platform that streamlines the entire workflow from reaction data input to result analysis and visualization. Users simply input a reaction dataset and select encoding methods and optimization algorithms; the platform then automatically performs search space

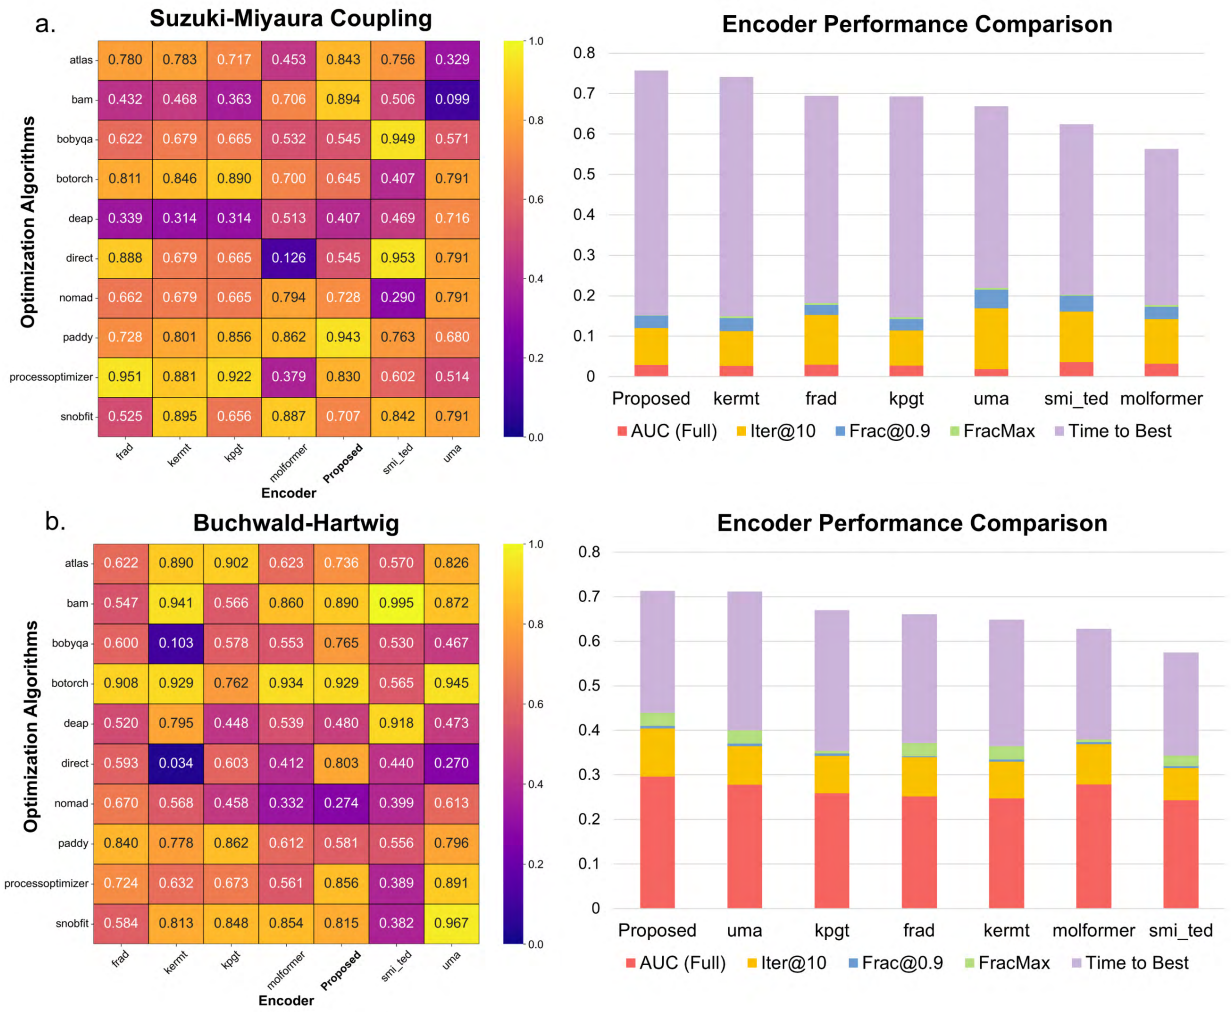

Figure 3. Optimization performance benchmarking for (a) Suzuki–Miyaura coupling and (b) Buchwald–Hartwig amination. Left: heatmap of final scores (weighted sum of five evaluation metrics) for each encoder–algorithm combination. Right: aggregated encoder performance across five evaluation metrics.

construction, optimization iterations, performance evaluation, and result reporting. The platform analyzes the input dataset to identify the objective function and decision variables, classifies each variable as categorical/discrete or continuous, and applies appropriate encoding and scaling rules to construct the search space. For reaction representation, the platform employs the pipeline described in the previous sections. To resolve package version compatibility issues between pretrained models, independent Docker-based virtual environments were established for each model.

All algorithms are evaluated under identical budgets and a unified closed-loop protocol. At each iteration, the optimizer proposes the next experimental conditions based on accumulated observations, and the resulting objective values are appended to the history for use in subsequent iterations. To ensure fair comparison and reproducibility, the platform automatically logs all key information generated during optimization, including proposed conditions, observed objective values, and best-so-far trajectories. In addition, visualizations based on t-distributed stochastic neighbor embedding (t-SNE) and Uniform Manifold Approximation and Projection (UMAP) provide qualitative insight into exploration behavior over iterations.

RESULTS

Reaction Latent Space Analysis

Suzuki-Miyaura Cross-Coupling

Suzuki–Miyaura cross-coupling forms C–C bonds between organoboron compounds and organic halides under palladium catalysis [29]. The dataset used in this study was constructed through high-throughput nanomole-scale reaction screening [30]. Reaction vectors

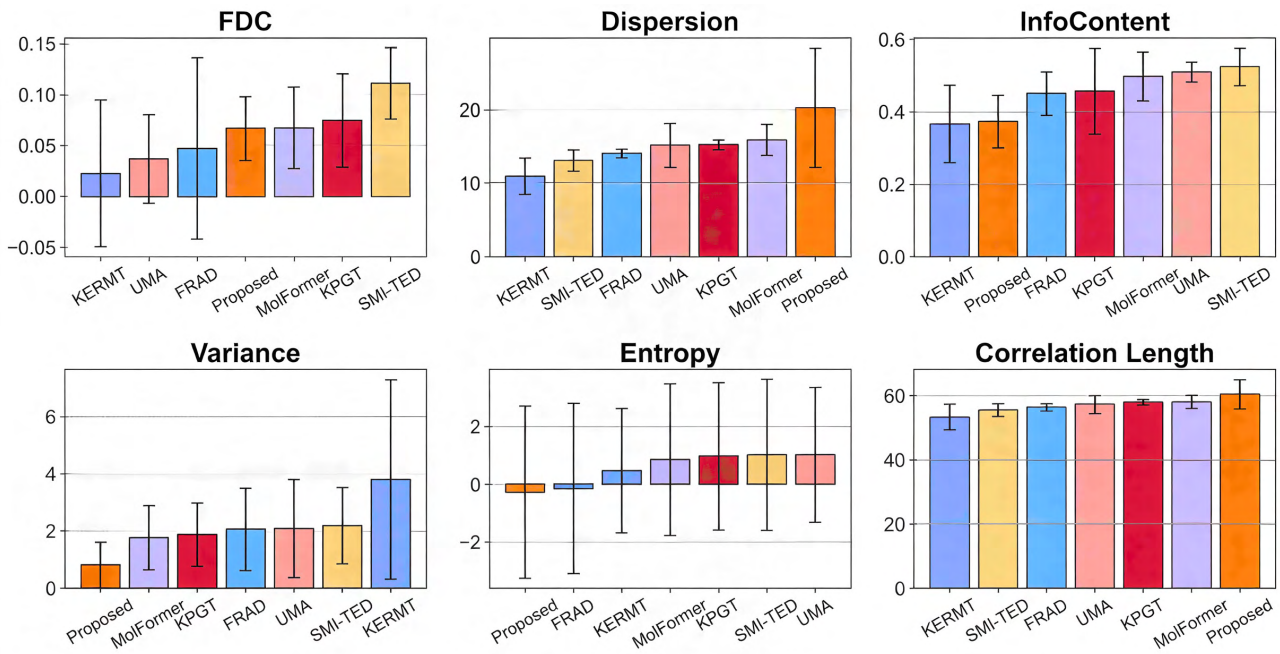

Figure 4: Fitness landscape metrics for each encoder, averaged across all reaction datasets. Six metrics are reported: fitness-distance correlation (FDC), dispersion, information content, variance, entropy, and correlation length. Bars indicate the mean value and error bars represent ±1 standard deviation.

were generated using UMA and visualized via t-SNE. The results revealed a hierarchical clustering structure dominated by the cation type of the base (Cs^+, K^+, Li^+, Na^+, and organic cations), with secondary sub-clustering by anion type (amine, carbonate/phosphate, carboxylate, fluoride, and hydroxide) within each cation region (Figure 2a).

This hierarchical structure reflects the underlying reaction mechanism. In transmetalation, the cation modulates the reaction rate through complexation with Pd-bound hydroxide [31, 32], while the anion determines the activation pathway [33]. The emergence of this chemically meaningful hierarchy without explicit supervision demonstrates that pretrained molecular encoders can capture mechanistically relevant features of reaction components.

Buchwald-Hartwig Amination

Buchwald–Hartwig amination is a cross-coupling reaction that forms C–N bonds between aryl halides and amines under palladium catalysis [34]. The dataset was constructed through high-throughput, nanomole-scale reaction screening [35]. Reaction vectors generated by FRAD and visualized via t-SNE revealed that the structural characteristics of aryl halides governed the latent space distribution, with phenyl and pyridyl derivatives occupying distinctly separated regions (Figure 2b).

This separation reflects the underlying reaction mechanism. The nitrogen atom in pyridine induces stereoelectronic stabilization during oxidative addition, leading to a fundamentally different pathway from phenyl halides [36]. Within each ring type, further separation by halogen type was observed—iodine derivatives formed independent clusters, while bromine and chlorine derivatives intermixed, consistent with the reactivity order ArI > ArBr > ArCl. Meanwhile, ligands, bases, and additives exhibited uniform distributions, confirming that aryl halide structure is the dominant factor. These results demonstrate that pretrained encoders can identify the most mechanistically influential components without explicit reaction knowledge.

Comparative Analysis of Encoder–Optimizer Combinations

For each reaction dataset, combinations of 7 encoders and 10 optimization algorithms were systematically compared under identical evaluation budgets and initial conditions. Performance was assessed using five metrics: AUC Full (area under the best-so-far curve over all iterations), Iter@10 (best-so-far yield at iteration 10), Frac@0.9 (number of iterations required to reach 90% of the maximum yield), FracMax (whether the global maximum yield was achieved), and Time to Best (the iteration at which the best value was first discovered). The final score for each encoder–algorithm combination was computed as a weighted sum of min-max normalized metrics. To verify that results are not sensitive to the weighting scheme, 100 random weight combinations were sampled from a Dirichlet distribution, confirming that the proposed

representation consistently ranked first regardless of weight choice. Figure 3 presents the heatmap of final scores and the aggregated encoder performance for each reaction.

In the Suzuki–Miyaura reaction, the proposed representation achieved the highest aggregated performance across all encoders, while maintaining a relatively stable performance on average across the 10 optimization algorithms. In contrast, other encoders exhibited larger performance fluctuations depending on the algorithm. For instance, MoLFormer scored 0.887 with SNOBFIT but only 0.126 with DIRECT, and UMA achieved 0.680 with PADDY yet dropped to 0.099 with BAM. Notably, the performance range of the proposed representation across algorithms (0.407–0.943) was narrower than those of MoLFormer (0.126–0.887) and UMA (0.099–0.791), indicating that the cross-attention-based reaction representation provides a latent space less sensitive to the choice of optimization algorithm.

In Buchwald–Hartwig reaction, the proposed representation also achieved the highest aggregated performance. It recorded high final scores with BoTorch (0.929), BAM (0.890), ProcessOptimizer (0.856), and SNOBFIT (0.815), indicating that the representation, which explicitly models interactions between reaction components, improves the fitting quality of surrogate models. Meanwhile individual molecular encoders exhibited high sensitivity to algorithm selection. For example, KERMT achieved 0.941 with BAM but dropped to 0.034 with DIRECT, and UMA recorded 0.967 with SNOBFIT but declined to 0.270 with DIRECT.

Overall, the proposed reaction representation demonstrated advantages in both robustness across algorithms and overall performance compared to simple molecular concatenation approaches. This is attributed to the fact that explicitly reflecting inter-component interactions in the latent space provides a structure that enables optimization algorithms to explore high-yield regions more efficiently. However, the proposed representation did not rank first in every reaction scenario, and cases where specific encoder–algorithm combinations were more advantageous depending on reaction characteristics were also observed. This supports the conclusion that no single optimal strategy exists, and the selection of an appropriate representation–algorithm combination tailored to the reaction scenario is important.

Landscape Characterization of Latent Representations

To understand the structural basis of these performance differences, we analysed the fitness landscape induced by each encoder using six metrics derived from the length scale framework [37]: fitness-distance correlation (FDC), dispersion, information content, variance, entropy, and correlation length. FDC indicates how strongly better fitness values align with proximity to the global optimum. Dispersion describes whether good solutions are clustered or widely scattered in the space. Information content measures how rugged or irregular the landscape is to small perturbations. Variance represents the diversity of local gradients across flat and steep regions. Entropy captures the overall complexity and unpredictability of the landscape structure. Correlation length reflects how smoothly neighboring regions are connected. Figure 4 reports the average and standard deviation of each metric across all reaction datasets.

CONCLUSION

The proposed representation exhibited the lowest variance, the lowest entropy, and the highest correlation length among all encoders. This combination indicates that the latent space formed yields a smooth, structurally simple landscape where the objective function changes predictably over long distances. Such properties are directly favorable for surrogate-based optimization, as they allow surrogate models to fit the underlying function more accurately with limited observations. The proposed representation also recorded the highest dispersion, suggesting that high-quality solutions are distributed broadly rather than confined to a narrow region, which benefits global search strategies. These landscape characteristics offer a structural explanation for the robust optimization performance observed in Figure 3.

This study proposed an integrated framework that combines foundation molecular representations with black-box optimization for organic reaction optimization problems and systematically analyzed its effectiveness. We demonstrated that the continuous latent representations from foundation models reflect molecular similarities learned from large-scale chemical data, thereby aligning the combinatorial space of reaction conditions into an optimization-friendly continuous representation space and enabling chemically meaningful exploration during the optimization process. In particular, by adopting UMA as the backbone and designing an interaction-aware reaction representation that explicitly models interactions between reaction components through bidirectional cross-attention, we confirmed that this design can address the dimensional increase and lack of interaction modeling in simple concatenation approaches. Furthermore, by utilizing the proposed representation as input design variables for optimization algorithms, we quantitatively demonstrated that various derivative-free optimization methods can effectively explore high-dimensional, mixed-integer reaction conditions under limited budgets, and showed that the advantageous representation–algorithm combinations may vary depending on the structure of the reaction scenario. Moreover, by analyzing the latent space generated by foundation

models, we demonstrated that the characteristics of the representation space can be closely related to optimization efficiency, providing empirical evidence for representation evaluation criteria from an optimization perspective that complements prediction-focused assessments.

In future work, we plan to extend this framework by integrating it with automated experimental equipment and SDLs, developing a closed-loop autonomous optimization system where robots execute the proposed experimental conditions and the results are fed back into learning and exploration. Through this, we aim to enhance the feasibility of autonomous reaction optimization that can identify high-performance conditions with minimal iterations even under limited experimental budgets.

REFERENCES

1. Rinehart NI, Saunthwal RK, Wellauer J, Zahrt AF, Schlemper L, Shved AS, Bigler R, Fantasia S, Denmark SE. A machine-learning tool to predict substrate-adaptive conditions for pd-catalyzed C–N couplings. Science 381:965-972 (2023). https://doi.org/10.1126/science.adg2114
2. Shields BJ, Stevens J, Li J, Parasram M, Damani F, Alvarado JIM, Janey JM, Adams RP, Doyle AG. Bayesian reaction optimization as a tool for chemical synthesis. Nature 590:89-96 (2021). https://doi.org/10.1038/s41586-021-03213-y
3. Kim Y, Doo H, Shin D, Lee SY, Roh Y, Park S, Song H, Jung Y, Yoo HJ, Han SS, Kim JW, Besenhard MO, Lee YS, Na J. Self-driving laboratories with artificial intelligence: an overview of process systems engineering perspective. Computers & Chemical Engineering 203:109266 (2025). https://doi.org/10.1016/j.compchemeng.2025.109266
4. Rankovi? B, Griffiths RR, Moss HB, Schwaller P. Bayesian optimisation for additive screening and yield improvements – beyond one-hot encoding. Digital Discovery 3:654-666 (2024). https://doi.org/10.1039/d3dd00096f
5. Wigh DS, Goodman JM, Lapkin AA. A review of molecular representation in the age of machine learning. WIREs Comput Mol Sci 12: (2022). https://doi.org/10.1002/wcms.1603
6. Kwon Y, Jung Y, Choi YS, Kang S. Interpretation of chemical reaction yields with graph neural additive network. Mach. Learn.: Sci. Technol. 6:025054 (2025). https://doi.org/10.1088/2632-2153/addfaa
7. Shi R, Yu G, Chen L, Yang Y. Yieldfcp: enhancing reaction yield prediction via fine-grained cross-modal pre-training. Artificial Intelligence Chemistry 3:100085 (2025). https://doi.org/10.1016/j.aichem.2025.100085
8. Ross J, Belgodere B, Chenthamarakshan V, Padhi I, Mroueh Y, Das P. Large-scale chemical language representations capture molecular structure and properties. Nat Mach Intell 4:1256-1264 (2022). https://doi.org/10.1038/s42256-022-00580-7
9. Li H, Zhang R, Min Y, Ma D, Zhao D, Zeng J. A knowledge-guided pre-training framework for improving molecular representation learning. Nat Commun 14: (2023). https://doi.org/10.1038/s41467-023-43214-1
10. Wood, B.M., et al. Uma: A family of universal models for atoms. arXiv preprint arXiv:2506.23971: (2025) https://doi.org/10.48550/arXiv.2506.23971
11. Ni Y, Feng S, Hong X, Sun Y, Ma WY, Ma ZM, Ye Q, Lan Y. Pre-training with fractional denoising to enhance molecular property prediction. Nat Mach Intell 6:1169-1178 (2024). https://doi.org/10.1038/s42256-024-00900-z
12. Adrian, M., et al. Multitask finetuning and acceleration of chemical pretrained models for small molecule drug property prediction. arXiv preprint arXiv:2510.12719: (2025) https://doi.org/10.48550/arXiv.2510.12719
13. Soares E, Vital Brazil E, Shirasuna V, Zubarev D, Cerqueira R, Schmidt K. An open-source family of large encoder-decoder foundation models for chemistry. Commun Chem 8: (2025). https://doi.org/10.1038/s42004-025-01585-0
14. Binois M, Wycoff N. A survey on high-dimensional gaussian process modeling with application to bayesian optimization. ACM Trans. Evol. Learn. Optim. 2:1-26 (2022). https://doi.org/10.1145/3545611
15. Häse F, Aldeghi M, Hickman RJ, Roch LM, Aspuru-Guzik A. Gryffin: an algorithm for bayesian optimization of categorical variables informed by expert knowledge. Applied Physics Reviews 8: (2021). https://doi.org/10.1063/5.0048164
16. Balandat, M., et al. Botorch: A framework for efficient monte-carlo bayesian optimization. Advances in neural information processing systems 33:21524-21538 (2020) https://doi.org/10.48550/arXiv.1910.06403
17. Bertelsen, S., et al. Processoptimizer, an open-source python package for easy optimization of real-world processes using bayesian optimization: Showcase of features and example of use. Journal of Chemical Information and Modeling 65:1702-1707 (2025) https://doi.org/10.1021/acs.jcim.4c02240
18. Hickman RJ, Sim M, Pablo-García S, Tom G, Woolhouse I, Hao H, Bao Z, Bannigan P, Allen C, Aldeghi M, Aspuru-Guzik A. Atlas: a brain for self-driving laboratories. Digital Discovery 4:1006-1029 (2025). https://doi.org/10.1039/d4dd00115j

19. De Rainville FM, Fortin FA, Gardner MA, Parizeau M, Gagné C. DEAP. Proceedings of the 14th annual conference companion on Genetic and evolutionary computation :85-92 (2012). https://doi.org/10.1145/2330784.2330799
20. Kennedy J, Eberhart R. Particle swarm optimization. Proceedings of ICNN'95 - International Conference on Neural Networks 4:1942-1948 (None). https://doi.org/10.1109/icnn.1995.488968
21. Hansen, N. The cma evolution strategy: A tutorial. arXiv preprint arXiv:1604.00772: (2016) https://doi.org/10.48550/arXiv.1604.00772
22. Beck AG, Iyer S, Fine J, Chopra G. Paddy: an evolutionary optimization algorithm for chemical systems and spaces. Digital Discovery 4:1352-1371 (2025). https://doi.org/10.1039/d4dd00226a
23. Audet C, Dennis JE Jr. Mesh adaptive direct search algorithms for constrained optimization. SIAM J. Optim. 17:188-217 (2006). https://doi.org/10.1137/040603371
24. Jones DR, Perttunen CD, Stuckman BE. Lipschitzian optimization without the lipschitz constant. J Optim Theory Appl 79:157-181 (1993). https://doi.org/10.1007/bf00941892
25. Powell, M.J. The bobyqa algorithm for bound constrained optimization without derivatives. Cambridge NA Report NA2009/06, University of Cambridge, Cambridge 26:26-46 (2009) https://doi.org/
26. Huyer W, Neumaier A. SNOBFIT -- stable noisy optimization by branch and fit. ACM Trans. Math. Softw. 35:1-25 (2008). https://doi.org/10.1145/1377612.1377613
27. Ma K, Rios LM, Bhosekar A, Sahinidis NV, Rajagopalan S. Branch-and-model: a derivative-free global optimization algorithm. Comput Optim Appl 85:337-367 (2023). https://doi.org/10.1007/s10589-023-00466-3
28. Zhu M, Mroz A, Gui L, Jelfs KE, Bemporad A, del Río Chanona EA, Lee YS. Discrete and mixed-variable experimental design with surrogate-based approach. Digital Discovery 3:2589-2606 (2024). https://doi.org/10.1039/d4dd00113c
29. Miyaura N, Suzuki A. Palladium-catalyzed cross-coupling reactions of organoboron compounds. Chem. Rev. 95:2457-2483 (2002). https://doi.org/10.1021/cr00039a007
30. Perera D, Tucker JW, Brahmbhatt S, Helal CJ, Chong A, Farrell W, Richardson P, Sach NW. A platform for automated nanomole-scale reaction screening and micromole-scale synthesis in flow. Science 359:429-434 (2018). https://doi.org/10.1126/science.aap9112
31. Lima CFRAC, Rodrigues ASMC, Silva VLM, Silva AMS, Santos LMNBF. Role of the base and control of selectivity in the suzuki–miyaura cross?coupling reaction. ChemCatChem 6:1291-1302 (2014). https://doi.org/10.1002/cctc.201301080
32. Zhang H, Kwong FY, Tian Y, Chan KS. Base and cation effects on the suzuki cross-coupling of bulky arylboronic acid with halopyridines: synthesis of pyridylphenols. J. Org. Chem. 63:6886-6890 (1998). https://doi.org/10.1021/jo980646y
33. Thomas AA, Denmark SE. Pre-transmetalation intermediates in the suzuki-miyaura reaction revealed: the missing link. Science 352:329-332 (2016). https://doi.org/10.1126/science.aad6981
34. Ruiz-Castillo P, Buchwald SL. Applications of palladium-catalyzed C–N cross-coupling reactions. Chem. Rev. 116:12564-12649 (2016). https://doi.org/10.1021/acs.chemrev.6b00512
35. Ahneman DT, Estrada JG, Lin S, Dreher SD, Doyle AG. Predicting reaction performance in C–N cross-coupling using machine learning. Science 360:186-190 (2018). https://doi.org/10.1126/science.aar5169
36. Kania MJ, Reyes A, Neufeldt SR. Oxidative addition of (hetero)aryl (pseudo)halides at palladium(0): origin and significance of divergent mechanisms. J. Am. Chem. Soc. 146:19249-19260 (2024). https://doi.org/10.1021/jacs.4c04496
37. Morgan R, Gallagher M. Analysing and characterising optimization problems using length scale. Soft Comput 21:1735-1752 (2015). https://doi.org/10.1007/s00500-015-1878-z

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Techno-Economic Optimization of Electrified Airports as Collaborative Energy Hubs

Mohammadreza Babaei[a]*, Stavros Vouros[a], Konstantinos Kyprianidis[a], and John D. Hedengren[b]

[a] Department of Engineering Sciences, Mälardalen University, Västerås, Sweden
[b] Department of Chemical Engineering, Brigham Young University, Provo, UT, USA
* Corresponding Author: mohammad.reza.babaei@mdu.se.

ABSTRACT

The electrification of regional aviation requires coordinated planning of airport energy systems that integrate renewable generation, energy storage, and hydrogen technologies in a cost-efficient and resilient manner. This paper presents a scalable techno-economic optimization framework that models multiple airports as collaborative energy hubs. An object-oriented mixed-integer linear programming (MILP) formulation is combined with a genetic algorithm (GA) to optimize infrastructure sizing and energy dispatch. The framework is applied to three Swedish regional airports—Västerås, Jönköping, and Visby. A set of scenarios, including parties operating under shared wind-energy contracts using power purchase agreements (PPAs) and dynamic pricing (DP), was studied. Detailed representations of battery energy storage, hydrogen production and storage, and market interactions are included. Results show that coordinated operation and airport collaboration under a smart energy management system can reduce total annual cost by up to 5% while enhancing operational flexibility and resilience to demand peaks and grid disturbances.

Keywords: Optimization, Genetic Algorithm, Energy Systems, Renewable and Sustainable Energy, Hydrogen

INTRODUCTION

Over the past decade, there has been an increasing effort to reduce greenhouse gas emissions. Under the Paris Agreement and net-zero goals, countries should reduce their carbon emissions by at least 55% relative to 1990 levels by 2030. Transport remains a major source of anthropogenic CO_2, and aviation presently accounts for about 2.5 % of total emissions; without effective mitigation, this share could increase to around 22 % by 2025 [1]. Although aviation's footprint is comparatively modest today, expanding demand, further stimulated by new air mobility services, could drive substantial emissions growth.

In response, researchers are developing low-carbon propulsion technologies for aviation. Hydrogen-fueled aircraft, which emit only water vapor, rely on advances in fuel cell technology and cryogenic storage systems, while electric propulsion for short-haul routes requires significant improvements in battery energy density and power-to-weight ratios [2].

Electrified aviation starts with short-haul flights, which show how regional airports can serve as energy hubs and provide an affordable transportation option. However, transitioning to electrified airports requires significant infrastructure investments, including airport energy systems, grid connections, power generation, and storage. This transition will be especially tough for medium-sized and smaller regional airports.

The resilience of the energy system is another challenge in electrification. The term resilience was first introduced by Holling et al. [3] to describe the capacity of ecosystems to absorb change. Molyneaux et al. [4] adopt the concept of resilience in energy systems and review different scenarios to measure the system's resilience. Resilience is the ability to adapt to change and uncertainty while maintaining a consistent performance during the transformation [5].

The integration of hydrogen into the energy system poses challenges, including balancing cost efficiency with infrastructure scalability. These challenges are compounded by the need for flexibility to adapt to dynamic market conditions and by the integration of renewable energy sources. In a study, Almansoori et al. [6-7] propose a mixed-integer linear programming (MILP) framework to optimize hydrogen supply chains by minimizing

sizing and dispatching costs across different demand profiles.

Multiple objectives are often involved in optimizing hybrid energy systems. In a study, Zheng et al. [8] designed a multi-objective framework to assess the potential of photovoltaic (PV) systems and battery energy storage for grid integration and deployment in the charging infrastructure. Yan et al. [9] investigated the combined use of wind and solar energy with hydrogen systems for off-grid applications like airports. Despite the potential benefits of integrating PV panels and wind turbines, some studies [10] highlight conflicts with existing safety regulations.

Xiang et al. [11] employed MILP to model an integrated hydrogen-solar storage microgrid and evaluate scenarios involving combinations of hydrogen production, PV energy, and battery storage systems. Their study highlighted the need for further investigation into effective grid integration strategies.

The resilience of the hybrid energy system was evaluated by Zhao et al. [12], focusing on the airport with the hydrogen-integrated technologies. The study introduces a MILP framework to optimize hydrogen storage and operation under various levels of disruptions in grid capacity and hydrogen permeability. As a result, integrating hydrogen increases the system's resilience but also raises the total annual cost.

A more detailed MILP model for the airport energy network was investigated by Vouros et al. [13] for a regional airport, where the total annual cost was evaluated across different resilience levels.

This paper advances the techno-economic assessment of airport energy systems through a scalable, object-oriented framework combined with a hybrid MILP–genetic algorithm approach. Unlike conventional single-site studies, it captures inter-airport collaboration, renewable energy flexibility markets, and modular system expansion, enabling a more realistic and extensible analysis of coordinated multi-airport energy hubs.

Electricity Pricing Schemes and Flexibility-Oriented Operation

Modern airport microgrids require the capability to adjust electricity supply and demand in near real time. Europe's electricity-market reform recognizes the role of flexibility markets in ensuring fair compensation for flexible resources such as energy storage, demand response, and dispatchable generation, while providing clear price signals for investment [14].

Long-term power-purchase agreements (PPAs) represent a common mechanism for valuing renewable electricity. Under a PPA, a buyer commits to purchasing a predefined quantity of electricity from a generator at a fixed price for 10–15 years. This agreement hedges against market volatility and allows risk sharing between contracting parties [15]. The contract price is often benchmarked against average wholesale market prices or a long-term forecast, with margins reflecting financing costs and risk preferences. Airports such as JFK [16] and Bristol [17] have adopted long-term renewable energy contracts.

The alternative dynamic pricing (DP) approach links electricity prices to real-time or day-ahead market outcomes. Dynamic tariffs reflect market-clearing prices and convey short-term supply conditions to consumers, incentivizing them to shift loads or operate storage systems [18]. In practice, suppliers often add a margin above the wholesale price to cover project financing and risk, which directly affects the economic attractiveness of storage and demand-response technologies that exploit price variability under dynamic pricing schemes.

While flexibility markets typically involve explicit participation of distributed energy resources through aggregators bidding into local or regional platforms operated by distribution system operators (DSOs), this study does not model such market structures explicitly. Instead, flexibility is represented implicitly through dynamic electricity pricing and coordinated dispatch of storage and conversion technologies at airport energy hubs. Long-term power purchase agreements (PPAs) are considered as bilateral contracts that provide price stability but do not constitute flexibility markets. On the other hand, dynamic pricing reflects changing system conditions over time and can encourage flexible operation of airport assets. This allows for system-level flexibility without the need for direct market involvement.

METHODS

This section describes the assumptions, algorithms, and model structure. First, the details of the algorithm are presented, followed by the model structure and problem assumptions.

Algorithm

The algorithm proposed in this work determines the optimal capacity and dispatch of the energy network for a multi-objective problem by combining MILP with the genetic algorithm (GA). The primary objective is to enable scalable modeling of energy networks for multiple airports. To achieve this, an object-oriented framework is implemented that leverages inheritance and encapsulation principles. Each component in the system inherits common base features, such as mass and energy balance constraints, while also retaining its own specific attributes. The overall problem formulation and optimization process are summarized as follows:

1. Airport object:

   An airport object is created and equipped with the required components, including energy storage

units, production technologies, market interfaces, and transportation systems. The capacity of each component can be defined as either a decision variable or a fixed parameter. A single problem instance may include multiple airport objects.

2. System connections:

   System components are interconnected either by user-defined instructions or by automatic configuration. These connections are subsequently used to formulate the mass and energy balance equations.

3. Energy network formulation:

   The energy network model is formulated as an MILP framework. Unit capacities can be optimized either directly within the MILP solver or through an outer-loop GA that explores the design space.

4. Optimization and results:

   The optimizer determines the optimal values of the decision variables according to the defined objective functions. The resulting solution is stored as a problem object for post-processing and analysis. The same object can subsequently be used for dispatch optimization under fixed unit capacities.

The final MILP model is generated based on the added components and their corresponding connections. Therefore, the implemented framework can be used for various applications, such as smart charging stations.

Problem Definition

We study three potential airport infrastructures to support battery-electric and hydrogen-electric aircraft.

Västerås Airport (VST) is located in an industrial city near major cities such as Stockholm and Uppsala. The airport could use excess energy generated by local power plants and contribute to the network. In addition to handling short-haul training flights, this airport transports approximately 100, 000 passengers per year.

Visby Airport (VBY), located on the island of Gotland, experiences significant seasonal passenger inflow and supports occasional international traffic, particularly during the tourist season. While not a primary international hub, its island location and wind energy potential make it strategically relevant for decentralized electrification planning. According to Swedavia statistics, the airport handled around 300 flights per month in 2025. While recent studies focus on the current status and technical potential, several planning and expansion studies identify significant potential for increasing wind capacity on the island, both offshore and onshore. An offshore floating wind farm is planned for 2032 with a capacity of around 1.75 GW.

The third airport is Jönköping (JKG), which enables short-range flights to southern cities. Recent studies focus on the airport's future role by investigating various charging and battery-swapping scenarios.

Figure 1. Three airports in contract with a wind energy supplier collaborated to improve resilience under a smart energy management control system

In addition, all the candidates are based on the SE3 grid region and operate under the same grid company. Therefore, agreements between energy providers and airports that serve as hubs are easier to establish.
This paper examines different scenarios in which a smart energy management system dispatches energy for three airports with a PPA or DP contract with a wind energy supplier. The wind farm can provide up to 2 MW of capacity during peak production hours, as the energy must be shared among all three airports. This limit reflects near-term availability and grid access constraints, rather than the island's full wind potential.

Table 1: Number of weekly flights for Västerås, Jönköping, and Visby airports

Airport	2025	2035*
Västerås (VST)	70	140 [13]
Jönköping (JKG)	18	58 [20]
Visby (VBY)	76	224 [21]

Flight Schedules

The flight schedule information was obtained from the OpenSkyNetwork [19] API. Figure 2 demonstrates estimated hourly flight schedules for VST, JKG, and VBY airports for a week. The number of flights and the fleet mix in 2035 are estimated based on the Swedavia Net Zero Roadmap and previous studies on Swedish airports, as shown in Table 1.

Hybrid Airport Energy Network Architecture

The airport's hybrid energy system can be divided into multiple sections. The storage and generation units, hydrogen conversion technologies, demand, and supply. Figure 3 shows the detailed structure of the hybrid airport energy network.

PV cells and wind turbines are the main electricity generation units. The weather data, including solar

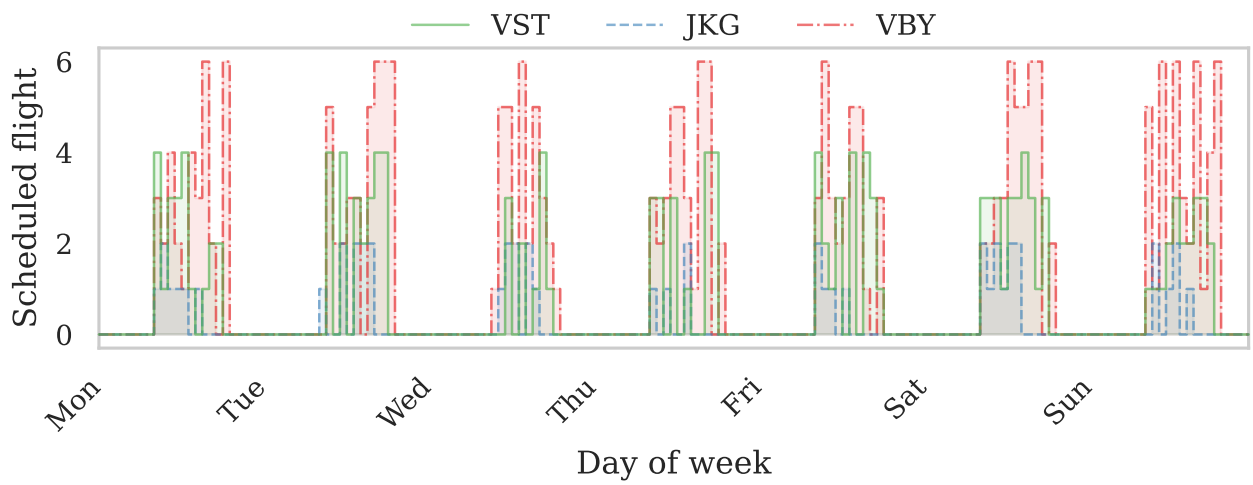

Figure 2. A representation of hourly flight schedules during a week

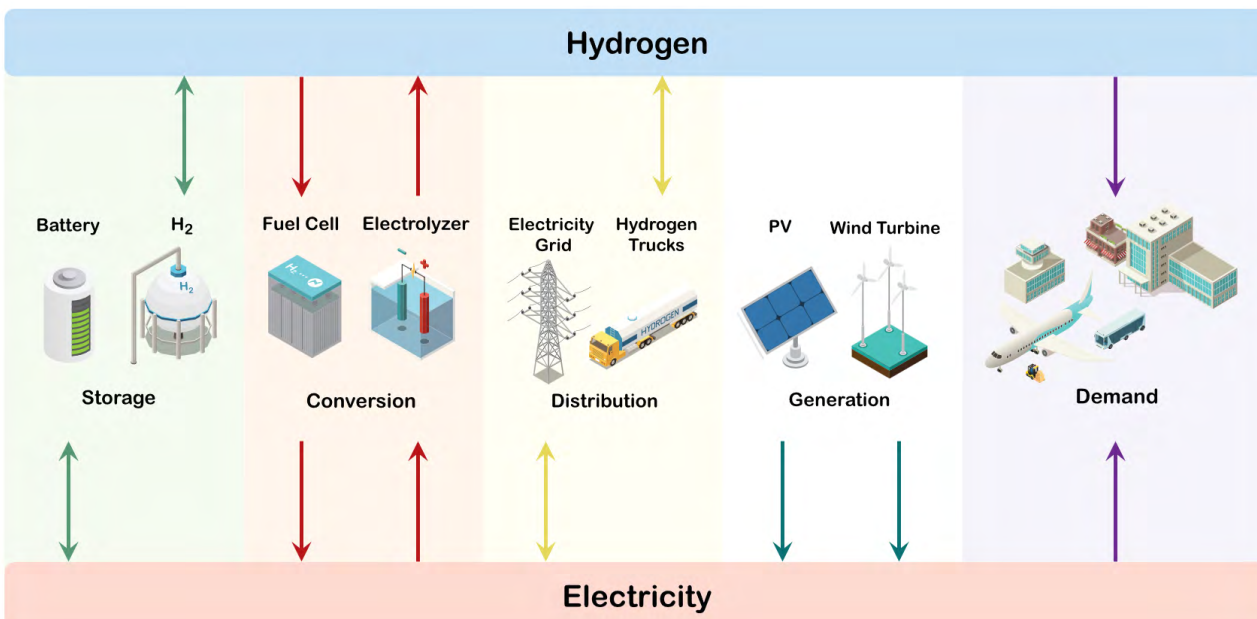

Figure 3. A schematic of mass and energy flow within an electrified airport energy network

irradiation and wind speed, are obtained from public databases such as SMHI [22] and PVLib [23]. Cut-in and cut-off wind speeds, as well as power curves, are considered during the data processing. Solar energy generation varies throughout the day based on the irradiation level and PV surface temperature. Equations (1-3) simplified, assuming panel and ambient temperatures are identical. Based on [24, 25], the equations for wind and solar energy are defined as:

$$P_k^{PV} = P_{STP}\frac{G_k}{G_{STP}}\left(1+\alpha_{pv}(T_{PV,k}-T_{STP})\right) \quad (1)$$

$$P_k^{WT} = P_k^{WT}(v_k) \quad (2)$$

The annual average accumulated global solar radiation is aligned with the Swedish records [26, 27].

BESS plays a critical role in peak shaving and providing short-term flexibility and buffering for the energy system. The battery model formulation, derived from [28], is outlined in Equation (3).

$$S_k^{bat} = S_{k-1}^{bat}(1-\beta^{bat}) + \left(\zeta_{ch}^{bat}\eta_{ch}^{bat}P_{ch,k}^{bat} - \zeta_{dch}^{bat}\frac{P_{dch,k}^{bat}}{\eta_{dch}^{bat}}\right)\Delta t \quad (3)$$

Where, the charging (η_{ch}^{bat}) and discharging efficiencies (η_{dch}^{bat}) in addition to a self-discharge term (β^{bat}) considered in the model to represent losses during the battery operation.

In order to isolate charging and discharging states, a pair of binary variables $\zeta_{ch,dch}$ is considered in (3), where:

$$\zeta_{ch} + \zeta_{dch} \leq 1 \quad (4)$$

Hydrogen tanks in both liquid (H_2^l) and vapor (H_2^v) phases are formulated in Equations (5-6). The term $\beta_{evap}^{H_2^l}$ shows the rate of evaporation for liquid hydrogen in the mass balance equations.

$$S_k^{H_2^l} = S_{k-1}^{H_2^l}\left(1 - \beta_{evap}^{H_2^l}\right) + \left(\zeta_{in}^{H_2^l} F_{in}^{H_2^l} - \zeta_{out}^{H_2^l} F_{out}^{H_2^l}\right)\Delta t \quad (5)$$

$$S_k^{H_2^v} = S_{k-1}^{H_2^v} + \beta_{evap}^{H_2^l} S_{k-1}^{H_2^l} + \left(\zeta_{in}^{H_2^v} F_{in}^{H_2^v} - \zeta_{out}^{H_2^v} F_{out}^{H_2^v}\right)\Delta t \quad (6)$$

The stored energy and hydrogen can be utilized in an electrolyzer and a fuel cell, respectively, to produce hydrogen and generate electricity during peak hours and provide flexibility for the system. The formulation for the fuel cell and electrolyzer is provided in (7-8).

$$F_k^{Elyz} = \frac{\eta^{Elyz}}{\alpha_{LVH}^{H_2}} P_k^{Elyz} \quad (7)$$

$$P_k^{FC} = \eta^{FC} \alpha_{LHV}^{H_2} F_k^{FC} \quad (8)$$

The airport can interact with the electricity and hydrogen market using the electricity grid connection and hydrogen transportation units, such as tube trucks and tank trucks. The hydrogen transportation units are limited and carry discrete amounts of hydrogen, with lower and upper limits. The electricity spot prices were retrieved from the Nord Pool platform. A synthetic DSO price is calculated taking into account the local energy tax and the DSO fee. While power purchases follow DSO prices, the selling value to the market is assumed to equal spot prices. In general form, each energy or mass stream in the system follows the balance equations. Mathematical formulation depends on the structure of hybrid energy networks. Equations (9-10) show the mass and energy balance.

$$P_k^i = \sum_j P_k^{ji} \quad (9)$$

$$F_k^i = \sum_j F_k^{ji} \quad (10)$$

Where ji represents the stream from item j to i. A non-negative constraint applied to the streams.

$$D_k^E \leq \sum_i P_k^i \quad \forall\, i \in \{bat, grid, FC, PV, WT\} \quad (11)$$

$$D_k^{H_2^l} \leq \sum_i F_k^i \quad \forall\, i \in \{liquef, H_2^l\} \quad (12)$$

$$D_k^{H_2^v} \leq \sum_i F_k^i \quad \forall\, i \in \{comp, H_2^v\} \quad (13)$$

The energy hub prioritizes airport demand while offering excess energy to the local grid to balance the load. The demand equations are shown in (11, 13).
The optimization includes design variables defined as:

$$N_{min}^i \leq N^i < N_{max}^i \quad \forall\, i \in \{bat, FC, Elyz, H_2^l, H_2^v, comp, liquef\} \quad (14)$$

The objective of the design and optimization problem is to minimize total annual cost (TAC) while avoiding the power outage in (19).

$$\text{CAPEX} = \sum_i \left(\frac{L_{sys}}{L_{cyc}^i}\right) N^i C_{cap}^i \quad \forall\, i \in \{bat, FC, Elyz, H_2^l, H_2^v, comp, liquef\} \quad (15)$$

$$\text{OPEX} = \sum_i \left(Q_{buy,k}^i C_{buy,k}^i - Q_{sell,k}^i C_{sell.k}^i\right) \quad \forall\, Q \in \{P^{grid}, F^{tat}, F^{tut}\}, \forall\, i \in \{grid, tat, tut, wt\} \quad (16)$$

$$CRF = \frac{r(1+r)^{L_{sys}}}{(1+r)^{L_{sys}} - 1} \quad (17)$$

$$\text{TAC} = \text{CRF} \times \text{CAPEX} + \left(\frac{\alpha}{\tau_{\text{sim}}}\right) \times OPEX \quad (18)$$

$$\min \quad TAC + \sum_i \sum_{k=1}^{\tau_{sim}} w^i \delta_k^i \quad \forall\, i \in \left\{D^E, D^{H_2^l}, D^{H_2^v}\right\} \quad (19)$$

System Parameters

This section covers both technical and economic parameters. Design variable constraints set upper and lower limits on decision variables. Table 2 presents the limits for the system's decision variables.

Table 2. Design variable constraints

Symbol	Description	Value
C_{max}^{bat}	Maximum battery capacity	200 MWh
$C_{max}^{H_2^v}$	Maximum vapor H_2 tank capacity	4000 kg
$C_{max}^{H_2^l}$	Maximum liquid H_2 tank capacity	6000 kg
C_{max}^{FC}	Maximum fuel cell capacity	7 MW
C_{max}^{Elyz}	Maximum electrolyzer capacity	10 MW
C_{max}^{comp}	Maximum compressor capacity	200 kg/h
C_{max}^{liquef}	Maximum liquefaction unit capacity	60 kg/h

To model airport energy system infrastructures incorporating hydrogen and energy storage technologies, it is essential to adopt up-to-date cost projections for both capital expenditures (CAPEX) and operation and maintenance (O&M) costs. Accordingly, this study synthesizes current and future cost estimates for key components, including green hydrogen production, battery energy storage systems, fuel cells, electrolyzers, hydrogen storage tanks in both compressed gaseous and liquid forms, and hydrogen delivery infrastructure. At present, green hydrogen production remains relatively costly due to its strong dependence on renewable electricity inputs. However, continued technological learning, economies of scale, and declining renewable energy costs are expected to substantially reduce production costs, with several projections indicating that green hydrogen could reach prices below 2 €/kg$_{H_2}$ by 2035.

Table 3. Cost estimation per unit for key equipment in the hybrid energy system

Equipment	Unit	2025	2035	2050	Ref.
Electrolyzer (Alkaline/PEM)	€/kW	270 – 900 [600]	200 – 700 [450]	120 – 270 [200]	[29,30]
Fuel Cell (PEM, stationary)	€/kW	600 – 1200 [900]	250 – 600 [500]	80 – 300 [180]	[12,29]
H_2^v tanks (350-700 bar)	€/kg_{H_2}	500 – 900 [700]	400 – 600 [500]	250 – 450 [350]	[11,13]
H_2^l tanks	€/kg_{H_2}	1150 – 1450 [1350]	1200 – 1300 [1250]	900 – 1200 [1000]	[11,13]
Battery Storage (Li-ion)	€/kWh	260 – 350 [300]	180 – 270 [220]	95 – 275 [160]	[21,29]

Scenarios

This section describes the scenarios studied in this research. Table 4 lists the scenarios analyzed in the following section. The first scenario, which includes design and dispatch optimization, is considered the basis for comparison. Scenarios 2-5 ($S_{2\text{-}5}$) represent base-case sizing as we evaluate various combinations of contract type and involvement.

The term "Collab" indicates the involvement of all three airports in contact with the wind energy supplier. While no collaboration considers only the VBY airport for accessing wind energy. S_{10} and S_{11} use the same design as S_7, and S_6 to evaluate the effect of switching to PPA while the system is optimized for the DP contract.

Table 4. Description of different scenarios

S_i	Sizing	Wind	Collab	Contract	Margin [%]
1	✓	-	-	-	-
2	-	✓	✓	PPA	0, ±5, ±10
3	-	✓	-	PPA	0, ±5, ±10
4	-	✓	✓	DP	0, 5, 10
5	-	✓	-	DP	0, 5, 10
6	✓	✓	-	DP	0, 5, 10
7	✓	✓	✓	DP	0, 5, 10
8	✓	✓	✓	PPA	0, ±5, ±10
9	✓	✓	-	PPA	0, ±5, ±10
10	-	✓	✓	PPA	0, ±5, ±10
11	-	✓	-	PPA	0, ±5, ±10

Computational Setup

The framework is implemented in Python 3.9 [29] using GEKKO Optimization Suite [30] for MILP formulation within an object-oriented structure, with AMPL [31] and CPLEX [32] as external solvers. Its modular design supports extension to larger airport networks and other multi-energy hub systems. The full GA optimization (40 generations, population 200) requires around 20 hours on an Intel i7-13700HX processor and 32 GB of RAM.

RESULTS AND DISCUSSION

The selection of the optimization time window can impact the final result. While a larger window yields a more general solution, it also increases computational cost. Therefore, the effect of window size is studied in Figure 4. A one-week window should reflect the system's demand and supply, given that flight schedules follow a weekly pattern at most airports.

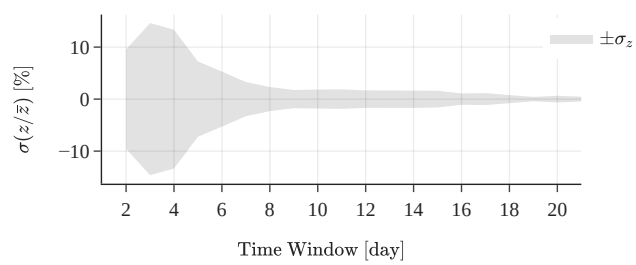

Figure 4. Variance of the difference between objective functions against different optimization window sizes for multiple starting dates in a year

Table 5 presents a selection of scenarios observed in this study, assuming a 5% margin for under-contract cases to ensure mutual benefits. The complete results report is available in the supplementary materials section. The MILP models for all scenarios were solved using CPLEX, with an optimality gap of 0.5%.

Table 5. CAPEX and OPEX values for selected scenarios

S_i	CAPEX [MEuro]	OPEX [MEuro]	$\Delta OPEX$ [%]	TAC [MEuro]	ΔTAC [%]
1	41.25	8.45	0.0	11.76	0
2	41.25	8.1	-4.1	11.42	-2.9
4	41.25	8.2	-2.9	11.51	-2.1
6	42.17	8	-5.4	11.38	-3.2
7	40.9	8	-5.3	11.29	-4
8	39.77	8.22	-2.7	11.41	-3
9	40.73	8.26	-2.3	11.53	-2
10	40.83	8.15	-3.5	11.43	-2.8

The base scenario has the highest capital and operating costs and suggests higher storage capacity to meet

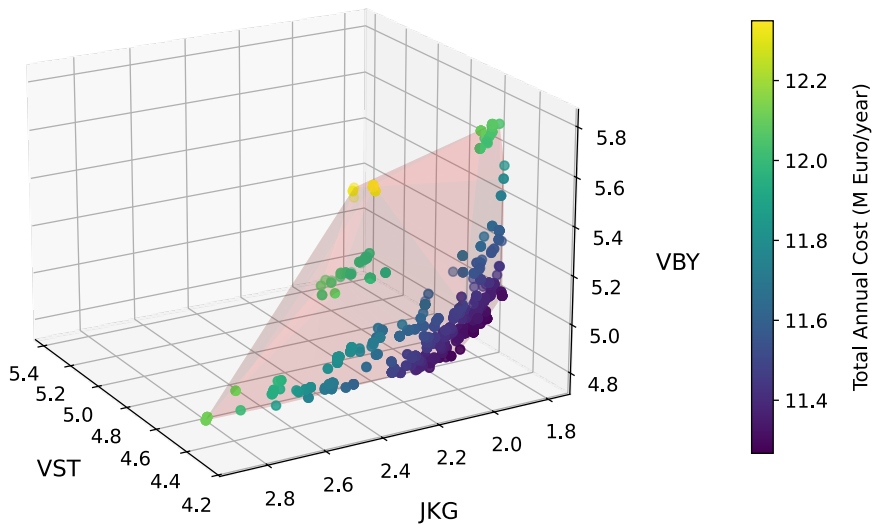

Figure 5. The optimal design and dispatch candidates from the last population of GA based on the total annual cost for VST, JKG, and VBY airports

demand during peak hours.

S_2 and S_3 show the effect of the contract with the wind farm, considering a fixed optimal design from the base scenario. In addition, S_2 represents the advantage of airport collaboration, reducing operational costs under normal conditions and providing room for resilience during unforeseen surges.

The scenarios with equipment sizing, while keeping up with demand, eliminate design redundancy and identify an optimal dispatch profile to reduce both OPEX and CAPEX. While the DP contract offers lower OPEX, a market study must still be conducted to reduce planning risk. On the other hand, PPA offers a fixed price for energy over a long period to mitigate spot price volatility. Considering the cost gap between S_7 and S_8, and the long-term energy price, one of these two scenarios may be the optimal choice. Despite a reduction of up to 5% in TAC, airports can meet demand and remain more flexible with the same base investment in the face of potential system disruptions.

All the case studies so far assume TAC is the sum of TAC across all individual airports. However, when it comes to collaboration and resilience, airports have the right to invest independently to achieve their own goals. To address conflicts among objectives, a three-objective optimization problem was solved using a GA to minimize TAC for the studied airports, under the assumptions in S_7. Figure 5 shows the optimal GA population after 50 generations with a population size of 200.

Three samples from the corners of the optimal surface were selected to represent cases in which each airport takes a major part in the investment. Another sample was selected to demonstrate minimal investment across all airports, bringing the total to four. The selected candidates are shown in Figure 6. By moving away from the optimal area, we offset the economic decision by increasing the system's flexibility. However, the pool of optimal candidates provides a guideline for moving toward a more resilient design.

While the TAC is almost identical across these cases, the individual's share of the total investment ranges from 0% to 90%, compared with the minimum-investment case. The TAC for these candidates is shown in Figure 7. An airport network represents a selected candidate and the objective value associated with each airport. The amount of TAC for all three airports aligns with their weekly load.

Figure 8 shows the decision variables for all airports among the selected candidates. One option to increase the VST airport's resilience is to invest in a fuel cell to handle sudden grid interruptions or utilize hydrogen to generate electricity on demand. This results in an 18% increase in TAC, representing the worst optimal point for VST. The same information applies to the VBY airport in describing candidate three.

The second network has the highest investment share for JKG, suggesting that investing in Battery is optimal. Based on the observations, BESS capital cost dominates the problem's objective value. In other words, the cost of battery storage technologies has a high impact on the future of electrification. In this case, increasing battery storage capacity improves disturbance rejection and peak shaving at JKG airport. As a result, BESS is a short-term solution to meet demand and handle grid outages for a couple of hours.

Hydrogen storage, on the other hand, despite its

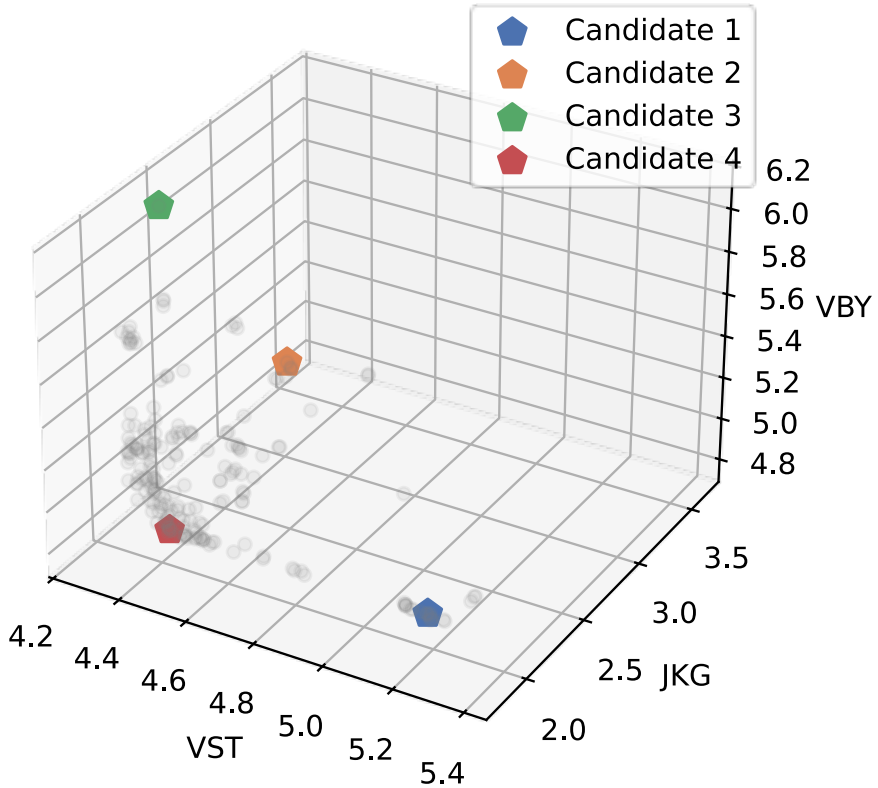

Figure 6. Total annual cost of VST, JKG, and VBY airports with respect to optimal candidates from the GA final population

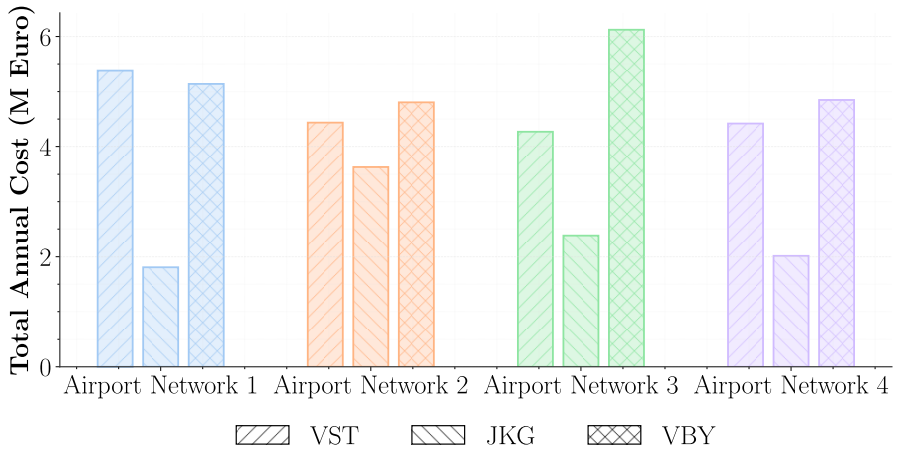

Figure 7. Total annual cost of VST, JKG, and VBY airports for the selected optimal candidates

lower efficiency at converting to electricity, is considered better at withstanding higher load spikes and sustaining longer grid outages. Therefore, investing in hydrogen technology is recommended for larger airports such as VST and VBY to increase the system's resilience.

The fourth airport network suggests greater investment in hydrogen storage, a liquefaction unit, and an electrolyzer. The result is promising, as this candidate has the lowest TAC across all airports. In conclusion, airports can take advantage of hydrogen production units rather than relying on the market. A detailed market study that accounts for safety and operating costs is required to reach a final conclusion.

From a system perspective, coordinated airport operation provides implicit benefits to the electricity grid by smoothing demand peaks and absorbing excess renewable generation. Optimized dispatch of battery and hydrogen storage enables load shifting away from high-

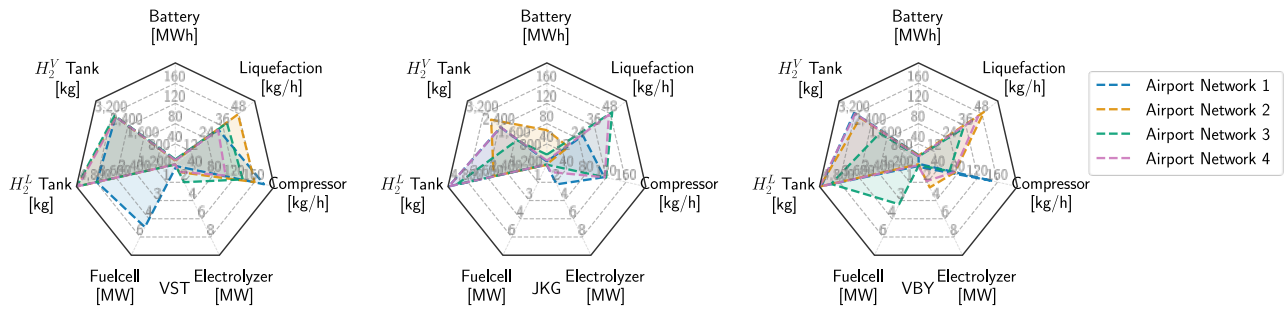

Figure 8. A comparison of decision variable values for the selected optimal population

priced periods and reduces reliance on grid imports during peak hours. In addition, the ability to export surplus electricity during periods of high renewable generation helps balance the local grid and mitigate congestion. Although these services are not monetized through an explicit flexibility market in this study, the results demonstrate that airport collaboration can provide flexibility-relevant services that are valuable for distribution system operation.

Beyond airport applications, the proposed framework is transferable to other multi-energy systems with distributed assets and flexible market interactions, such as smart charging networks, industrial energy hubs, and islanded microgrids. Its object-oriented structure enables scalable representation of interconnected nodes and shared resources, making it suitable for analyzing coordinated operation in energy communities and regional energy clusters. This highlights the framework's potential as a scalable decision-support tool for designing and operating complex, multi-vector energy systems under evolving market and regulatory conditions.

CONCLUSION

This work formulates a mixed-integer linear programming (MILP) model, embedded in an object-oriented framework and combined with a genetic algorithm, to optimize the sizing and operation of electrified airport energy hubs. Applied to a network of three regional airports, the approach demonstrated that coordinated operations and shared renewable contracts can reduce total annual costs by up to 5% while increasing flexibility during demand surges. Battery energy storage emerged as the primary contributor to capital cost and is effective for short-term resilience, whereas on-site hydrogen production and storage provide longer-duration security and can further lower costs for larger facilities. The study also illustrates that contract selection (dynamic pricing versus fixed price) influences the balance between operating costs and market exposure and that coordinated dispatch yields ancillary benefits for the local grid by smoothing peaks and absorbing excess renewable generation. These results suggest that future airport energy systems should prioritize coordinated operation, hybrid storage integration, and market-aware design to achieve cost-effective and resilient electrification.

DIGITAL SUPPLEMENTARY MATERIAL

Digital supplementary material for this study, including a spreadsheet of all scenarios, is available here.

ACKNOWLEDGEMENTS

This work is supported by the Green4iFED project under the RESILIENT competence center, financed by the Swedish Energy Agency (project number 2021-90273).

AUTHOR IDENTIFIERS

Author ORCIDs:
Babaei M.R.: 0000-0003-0709-3599
Vouros S.: 0000-0002-2737-3769
Kyprianidis K.: 0000-0002-8466-356X
Hedengren J.D.: 0000-0002-5535-5277

REFERENCES

1. Dolšak N, Prakash A. Different approaches to reducing aviation emissions: reviewing the structure-agency debate in climate policy. Clim Action 1: (2022). https://doi.org/10.1007/s44168-022-00001-w
2. Xue D, Chen XM, Yu S. Sustainable aviation for a greener future. Commun Earth Environ 6: (2025). https://doi.org/10.1038/s43247-025-02222-3
3. Holling CS. Resilience and stability of ecological systems. Annu. Rev. Ecol. Syst. 4:1-23 (1973). https://doi.org/10.1146/annurev.es.04.110173.000245
4. Molyneaux L, Brown C, Wagner L, Foster J. Measuring resilience in energy systems: insights from a range of disciplines. Renewable and Sustainable Energy Reviews 59:1068-1079 (2016). https://doi.org/10.1016/j.rser.2016.01.063
5. Folke C. Resilience (republished). E&S 21: (2016).

https://doi.org/10.5751/es-09088-210444
6. Almansoori A, Shah N. Design and operation of a future hydrogen supply chain. Chemical Engineering Research and Design 84:423-438 (2006). https://doi.org/10.1205/cherd.05193
7. Almansoori A, Shah N. Design and operation of a future hydrogen supply chain: multi-period model. International Journal of Hydrogen Energy 34:7883-7897 (2009). https://doi.org/10.1016/j.ijhydene.2009.07.109
8. Zheng P, Fang Z, Li H, Pan Y, Luo D, Zhang Z. Multi-objective optimization of hybrid energy management system for expressway chargers. Journal of Energy Storage 54:105233 (2022). https://doi.org/10.1016/j.est.2022.105233
9. Yan C, Zou Y, Wu Z, Maleki A. Effect of various design configurations and operating conditions for optimization of a wind/solar/hydrogen/fuel cell hybrid microgrid system by a bio-inspired algorithm. International Journal of Hydrogen Energy 60:378-391 (2024). https://doi.org/10.1016/j.ijhydene.2024.02.004
10. Morini M, Saletti C, Gambarotta A. A REVIEW OF SMART ENERGY PRACTICES AT AIRPORTS: CHALLENGES AND OPPORTUNITIES FOR SUSTAINABLE AVIATION. 37th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems (ECOS 2024) :144-155 (2024). https://doi.org/10.52202/077185-0013
11. Xiang Y, Cai H, Liu J, Zhang X. Techno-economic design of energy systems for airport electrification: a hydrogen-solar-storage integrated microgrid solution. Applied Energy 283:116374 (2021). https://doi.org/10.1016/j.apenergy.2020.116374
12. Zhao H, Xiang Y, Shen Y, Guo Y, Xue P, Sun W, Cai H, Gu C, Liu J. Resilience assessment of hydrogen-integrated energy system for airport electrification. IEEE Trans. on Ind. Applicat. 58:2812-2824 (2022). https://doi.org/10.1109/tia.2021.3127481
13. Vouros S, Margaritari K, Kyprianidis K. CO-OPTIMIZATION OF RESILIENT AIRPORT ENERGY HUBS FOR FLIGHT ELECTRIFICATION AND HYDROGEN PROPULSION. 37th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems (ECOS 2024) :120-131 (2024). https://doi.org/10.52202/077185-0011
14. Granado PC, Rajasekharan J, Pandiyan SV, Tomasgard A, Kara G, Farahmand H, Jaehnert S. Flexibility characterization, aggregation, and market design trends with a high share of renewables: a review. Curr Sustainable Renewable Energy Rep 10:12-21 (2023). https://doi.org/10.1007/s40518-022-00205-y
15. Kandpal B, Backe S, Crespo del Granado P. Power purchase agreements for plus energy neighbourhoods: financial risk mitigation through predictive modelling and bargaining theory. Applied Energy 358:122589 (2024). https://doi.org/10.1016/j.apenergy.2023.122589
16. Port Authority of New York and New Jersey, “Port Authority and TotalEnergies Launch New Solar Carport and Storage System at JFK Airport.” Apr. 2024. [Online]. Available: https://pv-magazine-usa.com/2024/04/24/a-second-solar-project-takes-off-at-jfk-airport/
17. Bristol Airport, “Bristol Airport enters milestone solar energy agreement with Luminous Energy.” Mar. 2024. [Online]. Available: https://www.bristolairport.co.uk/corporate/news-and-media/news-and-media-centre/2024/3/bristol-airport-enters-milestone-solar-energy-agreement-with-luminous-energy/
18. D. Holmberg, F. Omar, and T. Roth, “Impact of dynamic prices on distribution grid power quality, ” National Institute of Standards and Technology (U.S.), Gaithersburg, MD, NIST TN 2261, (2023). https://doi.org/10.6028/NIST.TN.2261
19. Schafer M, Strohmeier M, Lenders V, Martinovic I, Wilhelm M. Bringing up opensky: a large-scale ADS-B sensor network for research. IPSN-14 Proceedings of the 13th International Symposium on Information Processing in Sensor Networks :83-94 (2014). https://doi.org/10.1109/ipsn.2014.6846743
20. A. Holmberg, “Comparative Cost Analysis of PlugIn Charging and Battery Swapping Systems for Electric Aircraft: A Case Study of Jönköping Airport.” 2024. [Online]. Available: https://urn.kb.se/resolve?urn=urn%3Anbn%3Ase%3Akth%3Adiva-353827
21. Ollas P, Sigarchian SG, Alfredsson H, Leijon J, Döhler JS, Aalhuizen C, Thiringer T, Thomas K. Evaluating the role of solar photovoltaic and battery storage in supporting electric aviation and vehicle infrastructure at visby airport. Applied Energy 352:121946 (2023). https://doi.org/10.1016/j.apenergy.2023.121946
22. Swedish Meteorological and Hydrological Institute (SMHI), “Observed Meteorological Data from Swedish Weather Stations.” 2026. [Online]. Available: https://opendata.smhi.se/apidocs/metobs
23. Anderson KS, Hansen CW, Holmgren WF, Jensen AR, Mikofski MA, Driesse A. Pvlib python: 2023 project update. JOSS 8:5994 (2023). https://doi.org/10.21105/joss.05994
24. Waite M, Modi V. Modeling wind power curtailment with increased capacity in a regional electricity grid

supplying a dense urban demand. Applied Energy 183:299-317 (2016). https://doi.org/10.1016/j.apenergy.2016.08.078
25. Skoplaki E, Palyvos JA. On the temperature dependence of photovoltaic module electrical performance: a review of efficiency/power correlations. Solar Energy 83:614-624 (2009). https://doi.org/10.1016/j.solener.2008.10.008
26. J. Lindahl, "National Survey Report of PV Power Applications in Sweden".
27. Good C, Shepero M, Munkhammar J, Boström T. Scenario-based modelling of the potential for solar energy charging of electric vehicles in two scandinavian cities. Energy 168:111-125 (2019). https://doi.org/10.1016/j.energy.2018.11.050
28. Zou B, Peng J, Yin R, Li H, Li S, Yan J, Yang H. Capacity configuration of distributed photovoltaic and battery system for office buildings considering uncertainties. Applied Energy 319:119243 (2022). https://doi.org/10.1016/j.apenergy.2022.119243
29. Python Software Foundation, *Python 3.9*. Wilmington, DE, USA: Python Software Foundation, (2020). [Online]. Available: https://www.python.org
30. Beal LDR, Hill DC, Martin RA, Hedengren JD. GEKKO optimization suite. Processes 6:106 (2018). https://doi.org/10.3390/pr6080106
31. Fourer R, Gay DM, Kernighan BW. A modeling language for mathematical programming. Management Science 36:519-554 (1990). https://doi.org/10.1287/mnsc.36.5.519
32. IBM Corporation, *IBM ILOG CPLEX Interactive Optimizer 22.1.1.0*. Armonk, NY, USA: IBM Corporation, (2022). [Online]. Available: https://www.ibm.com/products/ilog-cplex-optimization-studio

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Particle Swarm Optimization for simultaneous design and optimization of heat pumps considering Mixed Integer problems

Beatriz C. da Silva[a]*, Ana M. Ribeiro[a], Alírio E. Rodrigues[a], Alexandre F.P. Ferreira[a], Diogo Rodrigues[a], Idelfonso B.R. Nogueira[b]

[a] LSRE-LCM, ALiCE, Faculty of Engineering, University of Porto, Rua Dr. Roberto Frias, 4200-465 Porto, Portugal
[b] Chemical Engineering Department, Norwegian University of Science and Technology, Sem Sælandsvei 4, Kjemiblokk 5, Trondheim 7491, Norway
* Corresponding Author: up201707213@fe.up.pt.

ABSTRACT

This study presents different approaches for introducing mixed integer problems into a meta-heuristic algorithm. The algorithms are developed to address the simultaneous design and optimization of a heat pump unit. A distinction is made between integer variables such as nominal tube diameters and the adsorbent employed in the process. The choice of adsorbent is named as a "key variable" due to its high impact on the process. To optimize the selection of these "key variables", a branched version of Particle Swarm Optimization (PSO) is presented and compared with the non-Branched version and a deterministic solver (IPOPT). Advanced Convergence Criterion is also implemented to mitigate the computational effort of these approaches. In the studied cases, Branch_PSO presents a higher degree of consistency and can even outperform the traditional PSO in simultaneous process optimization and material screening. However, its computational effort in cases with a large number of branches might be an obstacle.

Keywords: Optimization, Particle Swarm Optimization, Adsorption, Energy Systems, Energy Efficiency, Mixed Integer nonlinear problems, Key Variables, Heat Pumps, Material Screening

INTRODUCTION

Mixed Integer Nonlinear Programs (MINLP) are complex optimization problems that combine nonlinear objectives or constraints with integer variables. They frequently arise in fields such as engineering, economy, and logistic planning. Classic examples of such problems are the development of superstructures to determine which equipment shall be used in a certain process to maximize profit or operation scheduling to make the most of the resources available according to the requested shipments [1].

A possible approach to MINLP is the adaptation to a Nonlinear Program (NLP) via the relaxation of the integer variables into continuous ones. Another approach relies on the model adaptation by converting the nonlinear equations into linear ones, diminishing the problem's complexity. When no simplification is used, MINLP optimization can be achieved through deterministic methods, such as Branch-and-Bound (BB). A comprehensive review of those methods and commercial solvers applied to MINLPs can be found in Kronqvist et al. [2].

Nonetheless, there are some limitations to the deterministic methods. For example, they are usually unreliable when it comes to nonconvex problems, forcing a conversion into a convex problem, which is not always possible, or they fail to reach the global optimum by getting stuck in local optima or due to the premature cut of the branch that contained it (in the case of branch-based methods) [3].

Therefore, Evolutionary Algorithms have emerged for MINLP to avoid the previously mentioned problems. In the case of genetic algorithms (GA) [4] and Artificial Bee Colony (ABC) [5], the coding and decoding processes can be modified to encompass the existence of discrete variables. For Particle Swarm Optimization (PSO), the methodologies vary. The most common one is considering the discrete variables as continuous in terms of

evolution and rounding them off before the evaluation [6]. The rounding operation can be performed by simply using the round operator in the code language or by vertex approximation, as described further in the work of Chowdhury et al. [7]. An alternative approach uses operators such as sigmoid functions that affect the velocity of binary or discrete variables, ensuring their integer nature [8].

It is also common to encounter cases in the literature where multiple optimization strategies are combined to take advantage of their individual strengths. An example of this can be seen in Nema et al. [9], where a Hybrid PSO-BB is implemented. In that case, branch-and-bound is used to enhance convergence and PSO ensures the exploration of the problem domain to avoid being trapped in local minima.

However, the integer variables are rarely used as a key to altering the properties of the problem. From a process engineering perspective, the integer side of the problem is more commonly associated with variables such as the nominal pipe diameter or number of trays within a distillation column. This means that the variables considered have a sense of continuity despite their discontinuous nature, since the insufficiency of a variable value leads to the evaluation of the next higher value available, much similar to climbing a ladder.

When other types of integer variables such as the type of catalyst, solvent, or adsorbent materials are considered, the evaluation may not have that continuity. The optimal operating conditions for each material are rarely the same due to their different properties. By changing the evaluated material, the optimum for the process operation is likely to change, hampering the process optimization. As such, the sense of continuity is not naturally present in their evaluation. It is easier to think of such variables as keys on a key chain. When looking for the right key, trial-and-error is usually the chosen approach. It is not uncommon for the first try (key B) not leading to an open door, but neighboring keys (A and C) have no correlation with key B and choosing the next key to try is not evident.

Some techniques can be employed to mitigate this problem. One can use logic criteria to order the variables (ex. sort the adsorbents by their maximum capacity of adsorption). In the specific case of evolutionary algorithms that rely on random attributes for determining the values to be used in the next iteration, such as PSO, it is also possible to increase the number of evaluations to ensure proper exploration. That increase comes at the cost of higher computational effort, but can be an advantage for statistical analysis of the results obtained. At last, one can isolate the "key variables" and solve the optimization problem individually for each of the possibilities and then analyze the overall results. This last approach might not be viable if the problem considers multiple "key variables" (e.g. evaluating multiple catalysts for a reactor superstructure)

The present work aims to explore the impact of the "key variables" in MINLP optimization when using PSO algorithms. For that purpose, a base PSO algorithm considering rounding of operations was compared to a Branch_PSO. Additionally, advanced convergence criterion was also considered. A deterministic solver (IPOPT) was also considered in the comparison. The different algorithms were compared based on the performance obtained during the optimization of a material screening with simultaneous process optimization for adsorption heat pumps.

METHODS

In this work, four different PSO algorithms were compared:

- Base_PSO
- Branch_PSO
- Base_PSO_ACC
- Branch_PSO_ACC

These four algorithms were also compared with the deterministic method known as IPOPT (Interior Point OPTimizer). The case study considered in this study was the simultaneous design and optimization of adsorption heat pumps, as presented by Silva et al. [6].

The algorithms and the case studies are detailed in the following subsections.

Base_PSO

Amongst the different variations of the PSO algorithms presented in the literature, the one chosen as a comparison base was Constraints-based Sliding Particle Swarm Optimizer (CSPSO). CSPSO was introduced by Nogueira et al. [10] and adapted for later works in the optimization of chemical processes [6, 11-14]. CSPSO will be referred to as Base_PSO for the remainder of this paper.

The Base_PSO evaluates the MINLP by considering the integer variables as continuous ones and then rounding up their value upon the objective function evaluation.

For better understanding of Base_PSO, the corresponding equations and algorithm can be consulted in the referred papers or in the Supporting Information.

Branch_PSO

To evaluate the impact of the rounding operation on the previous algorithm, the integer variables were updated externally through the logic statements presented in Algorithm 1.

Algorithm 1. Branch_PSO operation.

```
Define N_d.
Create C.
Define N_C.
for i=1:N_C:
   Run PSO for C(i), updating only the con-
   tinuous variables.
   Save Y_globest_branch and X_globest_branch
end
Find Y_globest = min or max(Y_globest_branch) and re-
spective X_globest
```

In this algorithm, N_d is the number of discrete variables, C is the list of all the possible combinations (branches) of values for the discrete variables, N_C is the number of possible combinations (branches) obtained, $C(i)$ is the branch that is being evaluated in the ith run of the PSO algorithm, $Y_{globest_branch}$ is the vector that stores the global optimum found for each branch correspondent with the particle position $X_{globest_branch}$, $Y_{globest}$ is the best optimum amongst all branches in the $X_{globest}$ position. The position of the N_p particles defines the values for the decision variables.

Advanced Convergence Criterion (ACC)

During a PSO run, the number of iterations required to reach the optimal solution is often overestimated. Due to that, the PSO algorithm keeps on running even after effective convergence has been achieved, merely concentrating the particles around the optimal region without significant improvement in the global best solution. To overcome this limitation, an advanced convergence criterion is proposed.

Instead of relying exclusively on a predefined maximum number of iterations ($N_{it_{max}}$), the proposed criterion monitors the variation of the global best value between successive iterations. Convergence is assumed when the difference between the current and previous best values remains below a predefined threshold ($error_{crit}$) for a given number of consecutive iterations ($c_{it_{min}}$). This approach ensures that the optimization process is terminated only after a consistent solution has been identified, avoiding redundant and unnecessary iterations. This procedure is presented in Algorithm 2.

Algorithm 2. Advanced Convergence Criterion.

```
Set Y_globest_before = ∞
Set k = 0
Set c_it = 0
while k < N_it_max
    Run one PSO iteration.
    error_k = Y_globest − Y_globest_before
    if error_k < error_crit
    then Set c_it = c_it + 1
    else Set c_it = 0
    end
    if c_it = c_it_min
    then break
    end
    Set Y_globest_before = Y_globest
    Set k = k + 1
end
```

In this algorithm, c_{it} is the number of consecutive iterations that were assumed converged based on the cut-of criterion.

The application of this criterion to the Base_PSO and Branch_PSO resulted in Base_PSO_ACC and Branch_PSO_ACC algorithm, respectively.

IPOPT solver

To compare the performance of the developed algorithms, the IPOPT solver was also implemented to the same case study as the PSO.

The IPOPT (Interior Point OPTimizer) is used for large-scale Nonlinear Problems and is available as an open-source software package. The solver was implemented through the `cyipopt` library in Python [15].

The solver is based on barrier methods and therefore requires information about the derivatives of the constraints and objective function, unlike the PSO algorithm [16].

To approximate the Branch_PSO behavior, the solver was used in a similar way as described in Algorithm 1, with the integer ("key") variables being handled externally.

The calculation of the Objective Function derivative was made based on the forward finite-differences method with a step of 10^{-6}.

Case study: Design and optimization of adsorption heat pumps

Since the major distinction made between "ladder" and "key" variables lies with the physical and processual meaning the integer variables acquire within a process or operation, the different algorithms were compared within a design and optimization problem of a given processual unit. The comparison of the algorithm's performance under typical benchmark problems could not demonstrate the impact of the "key" variables evaluation methods and is not presented in the present paper.

The processual unit chosen for the case study was

an adsorption heat pump (AHP). Adsorption-based heat pumps offer a sustainable alternative to conventional heating technologies due to their low electricity consumption and ability to operate with renewable or waste heat sources [17, 18]. Their performance strongly depends on the choice of the adsorbent–adsorbate pair, making material screening and optimization essential to identify high-efficiency systems [19-21]. Usually, the operation cycle of AHPs is generically represented as pictured in Figure 1.

Isosteric heating and isobaric desorption are the steps where heat is supplied to the adsobent bed in the amounts Q_{heat} and Q_{des}, respectively, and Q_{cond} is produced as the working fluid is condensed. During the isosteric cooling and isobaric adsorption, heat is released to the environment in the amounts Q_{cool} and Q_{ads} and the working fluid is once again evaporated with the heat provided Q_{evap}.

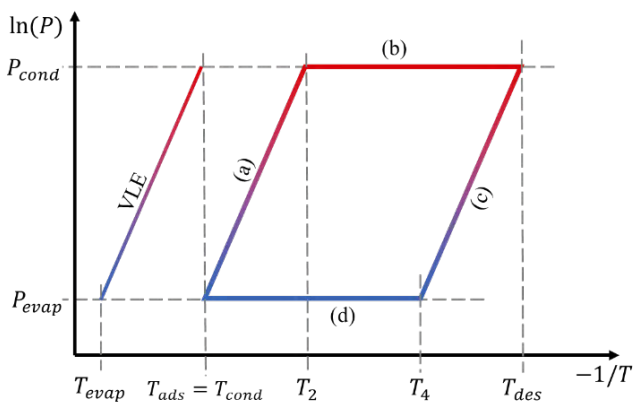

Figure 1. Diagram of ln(*P*) vs −1/*T* of an adsorption heat pump's cycle operation [vapor liquid equilibrium, isosteric heating (a), isobaric desorption (b), isosteric cooling (c), and isobaric adsorption (d)].

The performance of the equipment is measured by the coefficient of performance (COP_{heat}), calculated by Equation 1. This coefficient expresses the amount of heat produced by unit of heat required to operate the unit. Q_{evap} is considered ambient provided heat and therefore not considered within the required heat.

$$\text{COP}_{\text{heat}} = \frac{Q_{\text{cond}}+Q_{\text{ads}}+Q_{\text{cool}}}{Q_{\text{des}}+Q_{\text{heat}}} \quad (1)$$

In the study of Silva et al. [6], it is possible to encounter more details about the model equations, the adsorbent materials and the parameters considered in the model simulation. These are not included here for the sake of brevity.

The optimization problem is summarized below.

$$\max_{T_{\text{evap}},T_{\text{cond}},T_{\text{des}},\text{Material}} \text{COP}_{\text{heat}} \quad (2)$$

s.t.

$$g_1 = q_{\min} - q_{\max} \leq 0 \quad (3)$$

$$g_2 = T_{\text{evap}} - T_{\text{cond}} + 5 \leq 0 \quad (4)$$

$$g_3 = T_{\text{cond}} - T_{\text{des}} \leq 0 \quad (5)$$

$$g_4 = T_4 - T_{des} \leq 0 \quad (6)$$

$$g_5 = T_{cond} - T_4 \leq 0 \quad (7)$$

$$g_6 = T_2 - T_{des} \leq 0 \quad (8)$$

$$g_7 = T_{cond} - T_2 \leq 0 \quad (9)$$

$$T_{evap} \in [280.0, 303.0] \text{ K} \quad (10)$$

$$T_{cond} \in [309.0, 365.0] \text{ K} \quad (11)$$

$$T_{des} \in [345.0, 473.0] \text{ K} \quad (12)$$

$$\text{Material} \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\} \quad (13)$$

In this formulation, q_{min} and q_{max} are the minimum and maximum adsorbed amount (dependent on all the decision variables), $T_2 = T_{cond}^2/T_{evap}$ and $T_4 = T_{evap}T_{des}/T_{cond}$. In the PSO algorithms, the constraints were enforced by assigning the value of 100 (penalization) to the objective function whenever the constraints were not verified. In the IPOPT implementation, the constraints g_2-g_7 were applied within the solver algorithm. The same did not occur with the constraint g_1 because the IPOPT algorithm requires information about the derivatives of the constraints and that handling was not perceived as trivial, given the intricate equations at stake. Therefore, g_1 was handled with the penalization system in both PSO and IPOPT.

Simulation

The material screening with simultaneous process optimization for adsorption heat pumps was implemented in a computer with 32 GB of RAM and an Intel(R) Core(TM) i7-14650HX (2.20 GHz), using Spyder 6.1.0, associated with Python 3.10.13 64-bit.

The parameters used in the four PSO algorithms considered can be found in Table 1.

Table 1. Parameters used in the PSO algorithms.

	N_p	N_{it_max}	$error_{crit}$	$c_{it_{min}}$
Base_PSO	3000	200		
Base_PSO_ACC	3000	200	10^{-14}	20
Branch_PSO	300	200		
Branch_PSO_ACC	300	200	10^{-14}	20

When it comes to the number of particles and iterations in the branched and base operation, the common rule is to maintain the number of iterations constant and then calculate the number of particles by ensuring that $N_{p_branched_PSO}\, N_c = N_{p_Base_PSO}$ for better comparison.

In IPOPT, the convergence criterion was user-

defined to better compare with the PSO algorithms. The absolute tolerance was defined as 10^{-14}, and the bounds relaxation factor was considered as 0.0. The initial guess for the continuous decision variables was considered as the mid-interval value.

RESULTS

Material Screening of Adsorption Heat Pumps

The results for the maximization of the performance coefficient for heating operations (COP_{heat}) are presented in Table 2. The optimal values obtained by each optimizer are reported with a high number of significant figures to clearly highlight the differences between them, with the final digits shown in bold for added emphasis.

Table 2. Computation time and $\max(COP_{heat})$ obtained for the different algorithms applied.

	Base_PSO
t/s	3870
$\max(COP_{heat})$	1.863846575137443118
	Base_PSO_ACC
t/s	198
$\max(COP_{heat})$	1.863846575137442896
	Branch_PSO
t/s	1650
$\max(COP_{heat})$	1.863846575137444228
	Branch_PSO_ACC
t/s	524
$\max(COP_{heat})$	1.863846575137442896
	IPOPT
t/s	51
$\max(COP_{heat})$	1.863846575137385

As can be perceived, Branch_PSO surpasses Base_PSO used in the previous study. The computation time is reduced by 57%. In addition to this, the most probable value (MPV) for the optimum is slightly better in Branch_PSO for the maximization of COP_{heat}, although all the differences in results displayed in Table 2 are not significant when one considers a practical application.

In the case of the advanced convergence algorithms, the time reduction in computational time is quite significant when compared to the algorithms where the number of iterations was set from the beginning (95% for Base_PSO and 68% for Branch_PSO). However, the solution in cases that do not consider accelerated convergence is slightly more accurate than the ACC version. Nevertheless, since only one run was performed for each case, the different randomness of the initial particle colocation amongst algorithms can also be a factor in this analysis. It can also be seen that Base_PSO_ACC is 62% faster than Branch_PSO_ACC. This is due to the fact that Branch_PSO_ACC needs to meet the convergence criterion in all the branches and Base_PSO has 10x more particles in the unconstrained search for the global optimum.

Comparing the IPOPT results with the PSO ones, IPOPT is the fastest one (74% faster than Base_PSO_ACC) but is also the one where the optimal value of the objective function found is the lowest.

To better analyze the different methods, the heatmaps of the evaluations made by each algorithm were plotted and are presented in Figure 2.

Several notes can be taken from the data in the heatmaps. The first one is that Base_PSO algorithms did not execute any evaluation in the quadrants where the optimum for materials 1 and 7 were located (signaled with the pink dotted boxes). This highlights the problem of "key" variables because if the algorithm cannot locate the optimum within a given discrete option of adsorbent or configuration of the unit, then the comparison is not being performed fairly. It is important to note that these materials were not the worst two of the ten that were evaluated, as materials 9 and 2 display a lower or equal $\max(COP_{heat})$ value.

The second note worth mentioning is the inability of IPOPT to locate the optimum for material 9. In fact, IPOPT failed to make a single evaluation that is not penalized. The penalization on its own is not an exclusive limitation to the success of the optimization, as material 8 also had some penalized evaluations and the algorithm still got within the optimal quadrant. The fact that the evaluation of material 9 could not produce any non-penalized evaluation is cause for raising doubts on the effectiveness of IPOPT for the evaluation of more complex problems where the equations are not easily differentiated with respect to the decision variables. One can argue that the initial guess of the optimum has an effect, but making a more educated guess also requires more information on the problem, or a trial-and-error approach. With Branch_PSO, all materials were evaluated fairly, with a very similar number of evaluations in the optimal regions of all materials, even with the ACC. It is also interesting that material 9 was the most penalized material for the PSO algorithms as well, but they still converged to the optimal region (including the Base_PSO ones that fail in the other two materials, 1 and 7).

Hence, despite the lower computational time, IPOPT has a worse performance when considering optimum evaluations and the flexibility provided in the problem evaluation stage. Within the PSO methods, Branch_PSO outperformed Base_PSO in terms of balance across the evaluation of the different materials, Base_PSO_ACC can be considered a viable alternative as well for fast evaluation, but the convergence problems due to the "key" variables misevaluation are still a possibility.

CONCLUSIONS

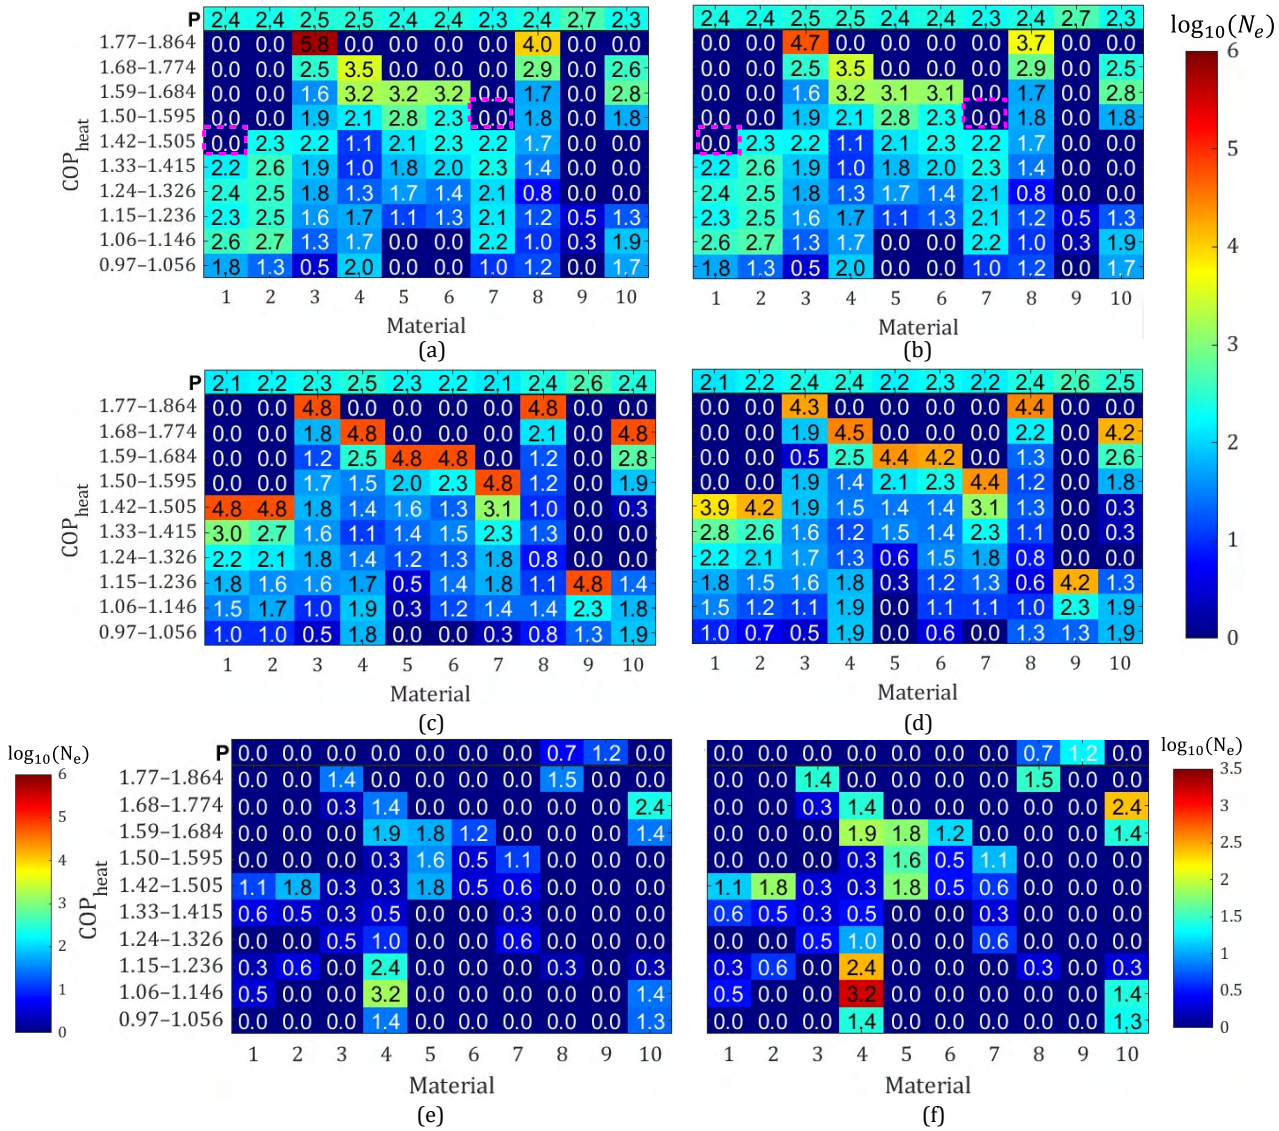

Figure 2. Heatmaps for the evaluations performed in Base_PSO (a), Base_PSO_ACC (b), Branch_PSO (c) Branch_PSO_ACC (d), IPOPT with the same scale as the PSO maps (e) and with autoscale (f). The values in the Heatmap are expressed in $\log_{10}$(Number of experiments (N_e)). The pink dotted squares represent quadrants that were not evaluated in the Base_PSO algorithms (correspondent with the location of that material optima). P stands for the evaluations that resulted in penalizations due to the non-observation of the constraints (only g_1 in the IPOPT cases).

Four different PSO algorithms were presented to solve MINLP. The algorithms were designed to assess the impact of "key variables" on optimization and process design problems. Their performance was compared to a deterministic solver (IPOPT).

For the design and optimization of adsorption heat pumps, Branch_PSO outperformed Base_PSO and highlighted the discontinuous problem of "key variables". Advanced convergence (ACC) had a large advantage in shorter computation time and can be employed in real-life applications if the error values are not prohibitive upon comparison with the remaining approaches.

Although IPOPT was the fastest algorithm, it delivered the lowest optimal value and showed limited robustness in constrained regions. IPOPT failed to produce any non-penalized evaluations for material 9, indicating difficulties in handling complex, non-smooth problems with stringent constraints.

Heatmap analyses showed that Base_PSO failed to explore the optimal regions of materials 1 and 7, which highlights the issue of key-variable misidentification, i.e., the inability of the algorithm to correctly explore variables

such as material selection. In contrast, Branch_PSO ensured a balanced evaluation of all materials, even with ACC enabled.

Overall, Branch_PSO_ACC provides the best compromise between robustness and computational efficiency, while Base_PSO_ACC remains a viable faster alternative with some risk of missing optima due to incomplete exploration of key variables.

As a final note, and despite the better performance of Branch_PSO, the algorithm presents a limitation in the incremental computational effort when a larger number of branches is considered. The breaking point when the computational effort nullifies the other advantages is a parameter that can only be evaluated case by case. However, with the increase in computational power and core-based parallelization, this problem can be surpassed with multiple particles being evaluated simultaneously, reducing computational time.

DIGITAL SUPPLEMENTARY MATERIAL

Base_PSO algorithm can be found in Supporting Information (LAPSE:2026.0035).

ACKNOWLEDGEMENTS

This work was financially supported by Fundação para a Ciência e a Tecnologia, I.P./MECI through national funds: LSRE-LCM, UID/50020/2025 (https://doi.org/10.54499/UID/50020/2025); and ALiCE, LA/P/0045/2020 (https://doi.org/10.54499/LA/P/0045/2020). Beatriz Silva acknowledges her Ph.D. scholarship 2023.00674.BD (DOI: 10.54499/2023.00674.BD) funded by FEDER funds through NORTE 2020 and by national funds through FCT/MCTES.

AUTHOR IDENTIFIERS

Author ORCIDs:
Beatriz C. da Silva: 0000-0001-6740-8762
Ana M. Ribeiro: 0000-0003-4269-1420
Alírio E. Rodrigues: 0000-0002-0715-4761
Alexandre F.P. Ferreira: 0000-0002-6746-8973
Diogo Rodrigues: 0000-0001-7823-2993
Idelfonso B.R. Nogueira: 0000-0002-0963-6449

REFERENCES

1. Grossmann IE, Quesada I, Raman R, Voudouris VT. Mixed-integer optimization techniques for the design and scheduling of batch processes. Batch Processing Systems Engineering :451-494 (1996). https://doi.org/10.1007/978-3-642-60972-5_21
2. Kronqvist J, Bernal DE, Lundell A, Grossmann IE. A review and comparison of solvers for convex MINLP. Optim Eng 20:397-455 (2018). https://doi.org/10.1007/s11081-018-9411-8
3. Molina-Pérez D, Mezura-Montes E, Portilla-Flores EA, Vega-Alvarado E, Calva-Yañez B. A differential evolution algorithm for solving mixed-integer nonlinear programming problems. Swarm and Evolutionary Computation 84:101427 (2024). https://doi.org/10.1016/j.swevo.2023.101427
4. Shreem SS, Turabieh H, Al Azwari S, Baothman F. Enhanced binary genetic algorithm as a feature selection to predict student performance. Soft Comput 26:1811-1823 (2022). https://doi.org/10.1007/s00500-021-06424-7
5. Akay B, Karaboga D, Gorkemli B, Kaya E. A survey on the artificial bee colony algorithm variants for binary, integer and mixed integer programming problems. Applied Soft Computing 106:107351 (2021). https://doi.org/10.1016/j.asoc.2021.107351
6. Silva BC, Rebello CM, Rodrigues AE, Ribeiro AM, Ferreira AFP, Nogueira IBR. Metaheuristic framework for material screening and operating optimization of adsorption-based heat pumps. ACS Omega 8:19874-19891 (2023). https://doi.org/10.1021/acsomega.3c01797
7. Chowdhury S, Tong W, Messac A, Zhang J. A mixed-discrete particle swarm optimization algorithm with explicit diversity-preservation. Struct Multidisc Optim 47:367-388 (2012). https://doi.org/10.1007/s00158-012-0851-z
8. Yiqing L, Xigang Y, Yongjian L. An improved PSO algorithm for solving non-convex NLP/MINLP problems with equality constraints. Computers & Chemical Engineering 31:153-162 (2007). https://doi.org/10.1016/j.compchemeng.2006.05.016
9. Nema S, Goulermas J, Sparrow G, Cook P. A hybrid particle swarm branch-and-bound (HPB) optimizer for mixed discrete nonlinear programming. IEEE Trans. Syst., Man, Cybern. A 38:1411-1424 (2008). https://doi.org/10.1109/tsmca.2008.2003536
10. Nogueira IBR, Martins MAF, Requião R, Oliveira AR, Viena V, Koivisto H, Rodrigues AE, Loureiro JM, Ribeiro AM. Optimization of a true moving bed unit and determination of its feasible operating region using a novel sliding particle swarm optimization. Computers & Industrial Engineering 135:368-381 (2019). https://doi.org/10.1016/j.cie.2019.06.020
11. Rebello CM, Martins MAF, Loureiro JM, Rodrigues AE, Ribeiro AM, Nogueira IBR. From an optimal point to an optimal region: a novel methodology for optimization of multimodal constrained problems and a novel constrained sliding particle swarm optimization strategy. Mathematics 9:1808 (2021). https://doi.org/10.3390/math9151808

12. Rebello CM, Martins MAF, Santana DD, Rodrigues AE, Loureiro JM, Ribeiro AM, Nogueira IBR. From a pareto front to pareto regions: a novel standpoint for multiobjective optimization. Mathematics 9:3152 (2021). https://doi.org/10.3390/math9243152
13. Nogueira IBR, Dias ROM, Rebello CM, Costa EA, Santana VV, Rodrigues AE, Ferreira A, Ribeiro AM. A novel nested loop optimization problem based on deep neural networks and feasible operation regions definition for simultaneous material screening and process optimization. Chemical Engineering Research and Design 180:243-253 (2022). https://doi.org/10.1016/j.cherd.2022.02.013
14. Rebello CM, Martins MAF, Rodrigues AE, Loureiro JM, Ribeiro AM, Nogueira IBR. A novel standpoint of pressure swing adsorption processes multi-objective optimization: an approach based on feasible operation region mapping. Chemical Engineering Research and Design 178:590-601 (2022). https://doi.org/10.1016/j.cherd.2021.12.047
15. ipopt 1.0.3. 2021 [Access; Available from: https://pypi.org/project/ipopt/.
16. Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math. Program. 106:25-57 (2005). https://doi.org/10.1007/s10107-004-0559-y
17. Canivet J, Fateeva A, Guo Y, Coasne B, Farrusseng D. Water adsorption in mofs: fundamentals and applications. Chem. Soc. Rev. 43:5594-5617 (2014). https://doi.org/10.1039/c4cs00078a
18. Dias JMS, Costa VAF. Adsorption heat pumps for heating applications: a review of current state, literature gaps and development challenges. Renewable and Sustainable Energy Reviews 98:317-327 (2018). https://doi.org/10.1016/j.rser.2018.09.026
19. Liu Z, Li W, Cai S, Tu Z, Luo X, Li S. Screening versatile water/adsorbent working pairs for wide operating conditions of adsorption heat pumps. Sustainable Energy Fuels 6:309-319 (2022). https://doi.org/10.1039/d1se01819a
20. de Lange MF, Verouden KJFM, Vlugt TJH, Gascon J, Kapteijn F. Adsorption-driven heat pumps: the potential of metal–organic frameworks. Chem. Rev. 115:12205-12250 (2015). https://doi.org/10.1021/acs.chemrev.5b00059
21. Demir H, Mobedi M, Ülkü S. A review on adsorption heat pump: problems and solutions. Renewable and Sustainable Energy Reviews 12:2381-2403 (2008). https://doi.org/10.1016/j.rser.2007.06.005

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Techno-economic analysis of hydrogen refueling station with on-site production from a novel blue H_2 and N_2 production system

Adrian R. Irhamna[ab] and George M. Bollas[a]*

[a] Department of Chemical and Biomolecular Engineering, University of Connecticut, Storrs, CT 06269, USA.
[b] Faculty of Mechanical and Aerospace Engineering, Institut Teknologi Bandung, Jawa Barat, 40132, Indonesia.
* Corresponding Author: george.bollas@uconn.edu.

ABSTRACT

This study presents a techno-economic modeling framework integrating a modular blue H_2N_2 production unit with a hydrogen refueling station (HRS) across capacities ranging from 0.1 to 4.0 tpd. A model-based approach is used to size key process and refueling components and to estimate the resulting hydrogen retail cost. The analysis indicates that hydrogen retail costs range from 4.6 to 10.8 USD kgH_2^{-1} over the considered capacity range. Relative to alternative on-site hydrogen production pathways, the proposed system demonstrates better cost-effectiveness while meeting clean hydrogen production standards. The approach is particularly suitable for regions with established natural gas infrastructure, as it leverages existing supply chains. Overall, the results provide actionable insights for policymakers and industry stakeholders in planning future hydrogen refueling infrastructure.

Keywords: Hydrogen, Hydrogen Refueling Stations, Distributed Production, Modelling, Technoeconomic Analysis.

INTRODUCTION

The decarbonization of heavy-duty transportation presents a distinct challenge within the broader energy transition. Unlike passenger vehicles, long-haul freight operations require high energy density, fast refueling, and minimal downtime [1]. These constraints limit the applicability of battery-electric solutions at scale. As a result, hydrogen has emerged as a candidate for deep decarbonization of freight corridors. Even though vehicle technology allows for hydrogen adoption, the heavy transportation sector is increasingly shaped by the feasibility and cost of hydrogen at refueling locations [2].

A central barrier to hydrogen development in freight transportation lies in the reliance on centralized hydrogen production coupled with long-distance delivery [3]. While large-scale hydrogen facilities benefit from economies of scale, they introduce significant logistical complexity, including high-pressure storage, liquefaction or compression, and dedicated transport infrastructure [3-5]. These downstream requirements contribute substantially to delivered hydrogen costs and pose scalability challenges for geographically distributed hydrogen refueling station (HRS), particularly along highway freight corridors where demand is spatially dispersed.

In response to these challenges, onsite and distributed hydrogen production has gained increasing attention as a complementary infrastructure strategy [6]. By producing hydrogen directly at refueling stations, onsite systems can eliminate transportation-related costs and losses, improve supply reliability, and enable incremental deployment aligned with localized demand growth. Such approach is attractive for freight applications, where hydrogen demand is concentrated at specific highway nodes rather than centralized hubs.

Onsite hydrogen production can be realized through both green and blue pathways, depending on regional resource availability and infrastructure constraints. Green hydrogen offers clear long-term decarbonization potential but remains cost-intensive at small scales due to the combined capital requirements of electrolyzers, and renewable electricity. In contrast, blue hydrogen pathways can leverage existing natural gas infrastructure to achieve improved cost competitiveness at distributed

scales, provided that high CO_2 capture efficiencies are maintained and methane emissions are effectively controlled [7]. This infrastructure compatibility positions blue hydrogen as a potential transitional solution for accelerating early-stage hydrogen deployment in freight applications [7].

Building on this premise, our previous work developed a modular blue hydrogen production based on a fixed-bed chemical looping process, designed explicitly for small-scale and distributed operation [8]. The system produces high-purity hydrogen while separating CO_2 and N_2 into concentrated streams, enabling efficient carbon capture without additional separation units [8, 9]. Prior techno-economic analysis demonstrated hydrogen production costs in the range of 2.8 – 3.9 USD kgH_2^{-1} at the plant gate for modular deployments [9, 10]. While these costs remain higher than those of large, centralized facilities, the modular system is well suited for integration with downstream users where hydrogen delivery costs dominate total supply-chain economics.

Despite growing interest in onsite hydrogen production, limited studies have systematically evaluated the integration of modular blue hydrogen systems with hydrogen refueling stations. In particular, the combined effects of production scale, refueling station equipment sizing, operating costs, and carbon intensity on retail hydrogen prices remain insufficiently explored.

We address this gap by developing an integrated techno-economic framework that couples a modular onsite blue H_2 and N_2 production system with a hydrogen refueling station across a range of distributed production scales. By explicitly modeling equipment sizing, operating requirements, hydrogen retail cost, and lifecycle CO_2 emissions, this work provides comprehensive assessment of the role of onsite blue hydrogen production in enabling cost-effective and low-carbon hydrogen refueling infrastructure for long-haul freight transportation.

SYSTEM DESCRIPTION

The Blue H_2-N_2 Production System

Figure 1 presents an overview of the blue H_2-N_2 production platform investigated in this study. The system is composed of three primary units: (i) a multi-reactor reforming section, (ii) an integrated heat recovery network, and (iii) a downstream gas separation and purification unit. The reforming section uses an intensified fixed-bed reactor previously reported in [8], which operates according to the chemical looping reforming (CLR) concept and cyclically undergoes oxidation (OX), reduction (RED), and reforming (REF) steps.

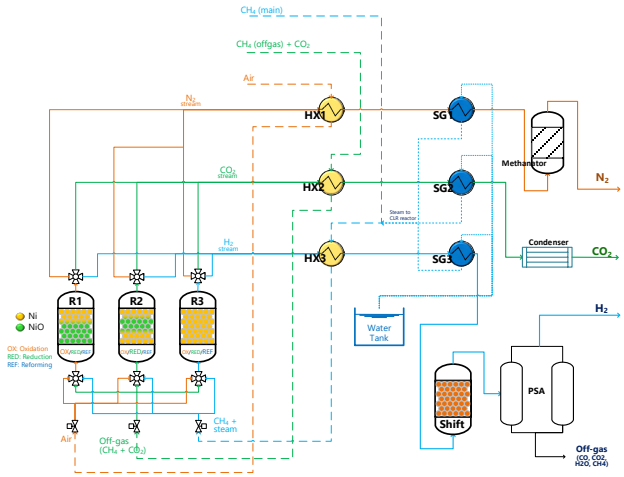

Figure 1. Simplified flow diagram of the blue H_2-N_2 production system considered in this study [9].

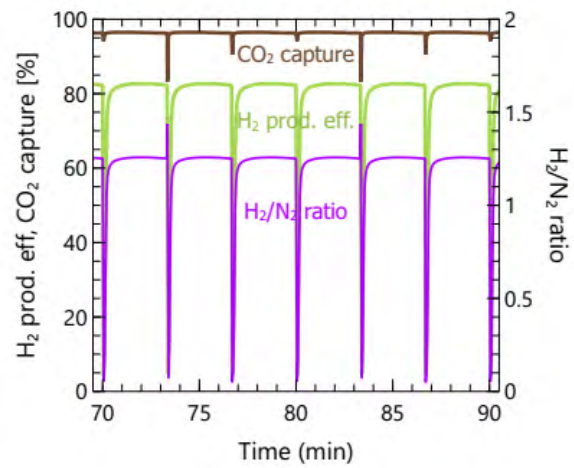

Figure 2. Performance metrics of the blue H_2N_2 production system discussed in this study [9].

Due to the inherently cyclic nature of CLR operation in fixed-bed configurations, continuous product generation cannot be achieved with a single reactor. To overcome this limitation, at least three reformers are operated in parallel, enabling semi-continuous system operation. This multi-reactor arrangement produces three separated outlet streams: a N_2-rich stream, a CO_2-rich stream, and a H_2-rich stream. Figure 2 presents the performance metrics of the blue H_2N_2 production system, where the system can produce H_2 with 82% efficiency and potential 98% of CO_2 capture [9]. A detailed description of the reactor sequencing, system integration, operating strategy, and system performance is provided in [9].

Hydrogen Refueling Station (HRS)

The HRS configuration comprises four main components: (i) multi-stage hydrogen compressors; (ii) a three-stage cascade hydrogen storage system; (iii) a hydrogen

dispenser; and (iv) a cooling system, following the configuration reported in [11]. The HRS is designed to accommodate onsite hydrogen production from both blue and green hydrogen pathways.

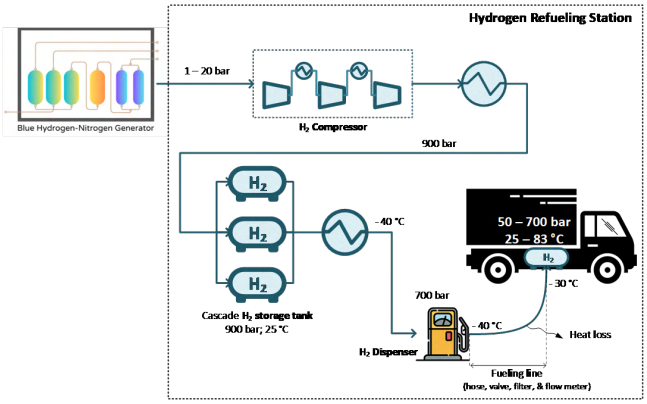

Figure 3. Simplified schematic of the integrated blue H_2-N_2 generator co-located with HRS.

METHODOLOGY

The techno-economic analysis in this study spans the system boundary, from natural gas supply to downstream treatment of the captured CO_2 by-product. Both capital expenditures (CAPEX) and operating expenditures (OPEX) are considered for the integrated blue H_2-N_2 generator and the hydrogen refueling station (HRS).

The techno-economic analysis of the blue H_2N_2 generator has been reported earlier [9, 10], and is summarized in Table 1. The reported hydrogen production cost includes the downstream treatment of the captured CO_2.

In the present study, the economic analysis focuses primarily on the hydrogen refueling station, including costs associated with hydrogen compression, storage, and dispensing to fuel electric vehicles (FCEVs). The HRS analysis considers the major system components: hydrogen compressors, hydrogen storage tanks, hydrogen dispensers, and refrigeration systems.

Table 1: H_2 production costs via the modular blue H_2N_2 generator estimated in our previous studies [9, 10].

Capacities	H_2 retail cost [USD kgH_2^{-1}]	Daily production [kgH_2 day^{-1}]
0.5 tpd	3.806	570
2.8 tpd	3.042	2, 882
7 tpd	2.882	7, 100
17 tpd	2.736	17, 107

Compressors

Piston and diaphragm compressors are commonly used for high-pressure hydrogen compression. In this work, a piston compressor is considered with an inlet pressure of 0.1 bar and a discharge pressure of 90 bar. This configuration enables hydrogen dispensing at 87.5 MPa, which corresponds to the maximum allowable pressure specified by the SAE J2601 standard for 70 MPa-class HRS [12, 13]. The high overall compression ratio is achieved through a multi-stage compression process, in which hydrogen is compressed sequentially across multiple stages with lower individual pressure ratios.

The compressor power requirement is calculated assuming adiabatic compression, expressed as:

$$\dot{W}_{comp} = c_{H2}\frac{T_{inlet}}{\eta_{comp}}\left(\left(\frac{P_{outlet}}{P_{inlet}}\right)^{\frac{r-1}{r}} - 1\right) \bullet \dot{m}_{H2} \quad (1)$$

where c_{H2} is the specific heat of hydrogen at constant pressure (14.304 kJ kg^{-1} K^{-1}), η_{comp} is the compressor efficiency (assumed to be 70%), T_{inlet} is the compressor inlet temperature in Kelvin, P_{inlet} and P_{outlet} are the inlet and outlet pressures, respectively, r is the specific heat ratio of hydrogen (1.4), and $\dot{m}_{H2}$ is the hydrogen mass flow rate through the compressor in kg s^{-1}.

Storage tanks

The hydrogen storage system is designed as a cascade configuration with multiple pressure levels. Storage tanks are sequentially filled by the compressor, starting from the low-pressure level and progressing to higher pressure levels, following a conventional with a conventional cascade refueling strategy [11, 12, 13]. The required storage capacity is calculated as:

$$C_{tank} = \dot{m}_{H2} t_{dispensing} \quad (2)$$

where C_{tank} is the total hydrogen storage capacity in kg or metric ton, and $t_{dispensing}$ is the daily refueling duration, assumed to be 18 h. In this study, a two-level pressure storage system is considered, with the total storage capacity evenly distributed between the two pressure levels.

Refrigeration and dispensing unit

The hydrogen dispenser system is designed to comply with the SAE J2601 fueling protocol, which specifies a hydrogen precooling temperature range of -30 to -40°C. Precooling is required to prevent the onboard hydrogen tank temperature in FCEVs from exceeding 85°C during refueling, a threshold above which tank safety may be compromised. The temperature increase during refueling is primarily attributed to the negative Joule-Thompson effect of hydrogen, which causes gas heating during throttling process. Accordingly, the hydrogen delivery temperature must remain between -40°C and -17.5°C [12, 13].

The main components of the dispenser system include a refrigeration unit for hydrogen precooling and a dispensing nozzle that regulates the hydrogen mass flow rate into the FCEVs storage tank. The cooling load for a single dispenser during vehicle fueling is calculated as:

$$\dot{Q}_{cool} = \dot{m}_{H2,D} \bullet (h_{storage} - h_{dispenser}) \quad (3)$$

where $\dot{m}_{H2,D}$ is the hydrogen mass flow rate at the dispenser (equal to 16.7 g s^{-1}) [11], $h_{storage}$ and $h_{dispenser}$ are the specific enthalpies at the storage (900 bar and 25°C) and at the dispenser exit (900 bar and -40°C), respectively. This calculation is based on following assumptions: (i) a dispensing pressure of 700 bar; (ii) two dispensers are required for station capacities of 100-200 kg day^{-1}; (iii) a fueling duration of 5 min per vehicle, and (iv) on board storage capacity of 5 kg.

The electrical power requirement for the refrigeration system is calculated as:

$$\dot{W}_{refrigeration} = \frac{Q_{cool}}{COP} \quad (3)$$

where the coefficient of performance (COP) is assumed to be 1.0 based on reference in [12].

Table 2: Economic formulations and cost data used in the techno-economic analysis of HRS.

Parameter
Capital cost
• Compressor [USD] [12]
$Cost_{compressor} = 45,626 \bullet \dot{W}_{comp}^{0.5861}$
• Storage tank [USD]
• High pressure [13]
$Cost_{storageHP} = 1,001 \bullet \dot{m}_{H2}$
• Low pressure [12]
$Cost_{storageLP} = 1,709 \bullet \dot{m}_{H2}$
• Refrigerator [USD] [12]
$Cost_{refrigerator} = 5,588 \bullet \dot{Q}_{cool}$
• Dispenser [USD] [12] *)
$Cost_{dispenser} = 65,000 \bullet n_{dispenser}$
O&M cost
• Compressor [USD][12]
$O\&M_{compressor} = 8\% \bullet Cost_{compressor}$
• Refrigerator [USD][12]
$O\&M_{refrigerator} = 3\% \bullet Cost_{refrigerator}$
• Dispenser [USD][12]
$O\&M_{dispenser} = 3\% \bullet Cost_{dispenser}$
*) $n_{dispenser} = 2$ for each additional 100–200 kgH_2 day^{-1} of station capacity [12].

Table 3: Framework used to estimate H_2 production costs, adapted from the U.S. DOE/NETL [14].

Components	**Formula**
Capital Cost	
- Bare Erected Cost of j-components ($BECj$)	$\sum BEC_j$
Total BEC	
Engineering Procurement and Construction (EPC)	$10\%\ (BEC)$
Contingency (CON)	$20\%(BEC\ +\ EPC)$
Total Plant Cost (TPC)	$BEC + EPC + CON$
Owner Cost & Other Financial Cost ($O\&OFC$)	$20\%\ (BEC)$
Total Overnight Cost (TOC)	$TPC + O\&OFC$
Total as spent Cost ($TASC$)	$TOC + 18\%TOC$
Annualized Capital Cost ($TASCpa$)	$TASC\ \times\ f$
Annual Operation and Maintenance Cost ($O\&M$)	$O\&M$
Total Annual Cost (TAC)	$TASCpa + O\&M$
H_2 retail cost (USD/kg)	$TAC/_{H2prod.}$

Economic and emissions analysis

Capital and operating costs for each major HRS component are estimated using cost correlations and assumptions compiled from the literature [12, 13], as summarized in Table 2. These cost correlations are applied consistently across all HRS capacities considered. The cost of individual HRS components are combined with those of the blue H_2 production unit to estimate Bare Erected Cost (BEC). Additional costs, including contingency, procurement, overnight cost, owner cost and other financial costs, are subsequently estimated.

The total capital cost is annualized using a capital recovery factor f of 0.159, corresponding to a project lifetime of 20 years and a discount rate of 15%, reflecting the perceived investment risk associated with hydrogen refueling infrastructure. The annualized capital cost is then combined with the operation and maintenance (O&M) costs of both HRS and blue H_2 production unit. This overall cost estimation framework is summarized in Table 3.

A sensitivity analysis is conducted on the estimated hydrogen retail cost with respect to key variable operating costs, including natural gas price, CO_2 transportation cost, and electricity price. The variability ranges for these operating costs are based on unit cost values reported in the literature [15-17].

Finally, the carbon dioxide equivalent (CO_2e) emissions associated with the hydrogen refueling station and onsite blue H_2 production are evaluated. The primary sources of CO_2 emissions arise from electricity consumption for hydrogen and CO_2 compression and refrigeration

Table 4: Key parameters used for equipment price estimation, obtained from dynamic simulations of the proposed system at hydrogen production rates of 0.1-, 0.5-, 1.0-, 2.0-, and 4.0-ton H_2 per day.

Equipment	Scaling variable	H_2 production rate				
		0.1 tpd	0.5 tpd	1.0 tpd	2.0 tpd	4.0 tpd
CL Reformer	Volume [m^3]	0.08	0.42	1.06	1.16	3.39
Shift Reactor	Shift gas flow [kg s^{-1}]	0.01	0.05	0.1	0.15	0.31
Heat recovery system	Heat duty per temp. [kW K^{-1}]	0.053	0.275	0.671	1.074	2.146
PSA	H_2 flow [liter s^{-1}]	3.2	22.2	39.7	63.4	127.0
CO_2 compressors	Power [kW]	5.4	27.4	66.9	107.1	214.2
Heat rejection unit	Amount of heat [kW]	6.3	33.0	78.2	125.2	250.4
HRS compressors	Power [kW]	47.8	236.4	594.6	955.1	1903.1
HRS storage tanks	Capacity [kgH_2 day^{-1}]	86.4	427.6	1075.6	1727.6	3442.3
HRS refrigeration	Cooling load [kW]	1.3	6.45	16.2	26.0	51.9
HRS dispenser	# of dispenser(s)	2	6	10	20	40

system. Electricity is assumed to be supplied by the grid, and emissions are calculated using a grid emission factor of 287 $kgCO_2e$ MWh^{-1} [18].

RESULTS AND DISCUSSION

The HRS and onsite blue H_2 production units were sized to match the onsite hydrogen production capacities of 0.1, 0.5, 1.0, 2.0, and 4.0 tH_2pd. Table 4 summarizes the sizing results for the main equipment in both the on-site blue hydrogen production system and the HRS, which are subsequently used to estimate hydrogen retail costs. The resulting retail hydrogen cost estimates are presented in Table 5.

Table 5: Estimated hydrogen retail costs for HRS with on-site hydrogen production.

Capacities	H_2 retail cost [USD kgH_2^{-1}]	Daily production [kgH_2 day^{-1}]
0.1 tpd	10.871	115
0.5 tpd	6.122	570
1.0 tpd	4.923	1, 434
2.0 tpd	4.821	2, 303
4.0 tpd	4.607	4, 589

The retail costs obtained from the proposed onsite blue hydrogen pathway are compared with alternative onsite hydrogen production technologies reported in the literature, as shown in Figure 4. Overall, the results clearly demonstrate the benefits of economies of scale, with higher production capacities yielding more competitive hydrogen retail costs. Most onsite hydrogen production technologies operate at relatively small scales (0.1 – 0.5 tpd), with corresponding retail hydrogen costs ranging between 8 and 10 USD kgH_2^{-1}. In particular, HRS with on-site electrolysis powered by renewable electricity operate at capacities of 0.1 tpd or lower and report retail hydrogen costs exceeding 12 USD kgH_2^{-1}.

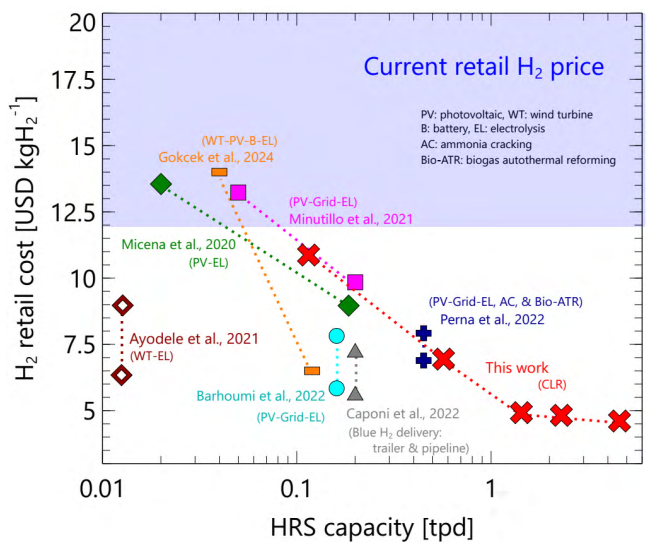

Figure 4. Comparison of hydrogen retail costs across different on-site hydrogen production technologies reported in the literature [6, 12, 13, 19-22].

Figure 4 indicates that HRS supplied by onsite blue hydrogen production becomes cost-effective at production rates above 1.0 tpd, achieving hydrogen retail costs of approximately 4.8 USD kgH_2^{-1}. To date, limited studies have reported onsite blue or green hydrogen pathways at comparable production scales. At smaller production capacities, the cost advantage of the proposed system reduces. Specifically, at a production scale of 0.5 tpd, the retail hydrogen cost increases to approximately 6.1 USD kgH_2^{-1}. As the production scale is further reduced to 0.1 tpd, the retail cost rises to approximately 10 USD kgH_2^{-1}, placing the proposed system within similar cost range to onsite green hydrogen production technologies.

Figure 5 shows the sensitivity analysis of the estimated retail price for a 0.5 tpd HRS with respect to key variable operating costs, including natural gas price, CO_2 transportation cost, and electricity prices. Overall, the retail hydrogen cost exhibits a linear response to variations in all operating cost components. Among these factors, natural gas is the most influential, followed by electricity price, and CO_2 transportation cost. Variations in natural gas price result in changes to the retail hydrogen cost ranging from -9% to 21%, while CO_2 transportation cost leads to variations of -1% to 10%. Electricity price fluctuations result in retail cost changes of approximately -13% to 13%.

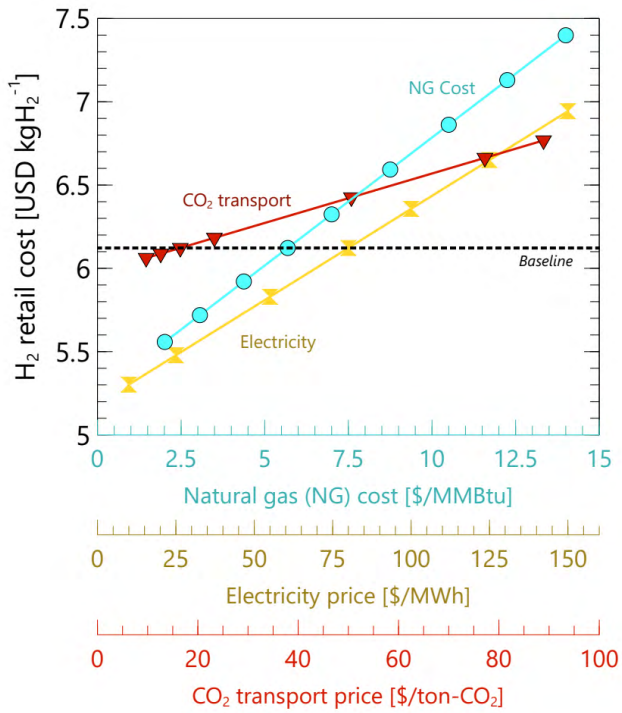

Figure 5. Sensitivity analysis of blue H_2 retail costs with respect to variations in natural gas price, electricity price, and CO_2 transportation cost for a HRS with 0.5 tpd onsite blue hydrogen production.

Figure 6 illustrates the CO_2e intensity associated with the HRS supplied by onsite blue hydrogen production considered in this study. The estimated emissions intensity remains relatively constant across all production scales, at approximately 3.05 $kgCO_2$ kgH_2^{-1}. At this level, more than 75% of total emissions originate from the HRS unit, primarily due to the electricity consumption for hydrogen compression and refrigeration. Overall, the resulting CO_2e intensity meets the U.S. Department of Energy (DoE) criteria for clean hydrogen production [23]. However, it should be noted that this value is achieved under the assumptions of a 95% CO_2 capture and negligible methane leakage during natural gas transmission.

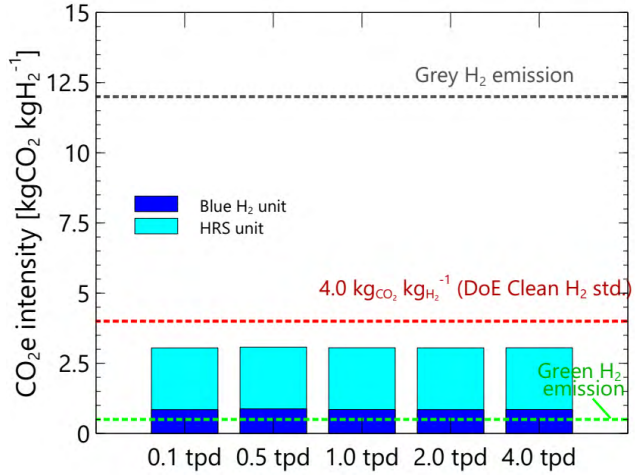

Figure 6. CO_2e intensity associated with HRS supplied by onsite blue H_2 production discussed in this study.

CONCLUSIONS

This paper presented a techno-economic analysis model that integrates a modular blue H_2N_2 generator with a HRS model across production scales. Dynamic simulations were performed to size the main equipment of the blue H_{2+}-N_2 generator and its integration with HRS components, enabling estimation of the resulting hydrogen retail cost. The retail cost was estimated to range from 4.6 – 10.8 USD kgH_2^{-1} across capacities of 0.1 – 4.0 tpd. Compared to other on-site production technologies, this process option was shown to be the most cost-effective while meeting clean H_2 production standards. The proposed system could be particularly attractive to regions with extensive natural gas infrastructure, like the United States, as it leverages existing supply chains. However, reliance on natural gas raises concerns over fugitive methane emissions, which remain subject to regulatory scrutiny. Despite these considerations, the findings provide valuable insights for policymakers and industry stakeholders in shaping future hydrogen infrastructure.

AUTHOR IDENTIFIERS

Author ORCIDs:
Irhamna, Adrian: 0000-0001-9752-7017
Bollas, George: 0000-0002-1960-431X

REFERENCES

1. Fuels Institute, “The Easiest and Hardest Commercial Vehicles to Decarbonize APRIL 2 0 22, ” 2022. Accessed: Jan. 15, 2026. [Online]. Available: https://transportationenergy.org/wp-content/uploads/2022/10/FI_Report_Med-Heavy-Duty_FINAL.pdf

2. World Economic Forum, “The hidden 85%: How to make hydrogen cost-effective.” Accessed: Jan. 14, 2026. [Online]. Available: https://www.weforum.org/stories/2024/05/hydrogen-hidden-costs-energy-transition/
3. Di Lullo G, Giwa T, Okunlola A, Davis M, Mehedi T, Oni AO, Kumar A. Large-scale long-distance land-based hydrogen transportation systems: a comparative techno-economic and greenhouse gas emission assessment. International Journal of Hydrogen Energy 47:35293-35319 (2022). https://doi.org/10.1016/j.ijhydene.2022.08.131
4. Andersson J, Grönkvist S. Large-scale storage of hydrogen. International Journal of Hydrogen Energy 44:11901-11919 (2019). https://doi.org/10.1016/j.ijhydene.2019.03.063
5. Usman MR. Hydrogen storage methods: review and current status. Renewable and Sustainable Energy Reviews 167:112743 (2022). https://doi.org/10.1016/j.rser.2022.112743
6. Caponi R, Bocci E, Del Zotto L. Techno-economic model for scaling up of hydrogen refueling stations. Energies 15:7518 (2022). https://doi.org/10.3390/en15207518
7. Mayer P, Ramirez A, Pezzella G, Winter B, Sarathy SM, Gascon J, Bardow A. Blue and green ammonia production: a techno-economic and life cycle assessment perspective. iScience 26:107389 (2023). https://doi.org/10.1016/j.isci.2023.107389
8. Irhamna AR, Bollas GM. Intensified reforming reactor for blue hydrogen and nitrogen production. International Journal of Hydrogen Energy 80:1477-1492 (2024). https://doi.org/10.1016/j.ijhydene.2024.05.428
9. A. R. Irhamna and G. M. Bollas. “Enhancing Hydrogen Infrastructure for Domestic Use: A Small-scale Blue Hydrogen and Nitrogen Production System”, submitted (under-review).
10. Irhamna AR, Bollas GM. Techno-economic analysis of a novel small-scale blue H2 and N2 production system. Systems and Control Transactions 4:159-165 (2025). https://doi.org/10.69997/sct.134013
11. Reddi K, Elgowainy A, Rustagi N, Gupta E. Impact of hydrogen refueling configurations and market parameters on the refueling cost of hydrogen. International Journal of Hydrogen Energy 42:21855-21865 (2017). https://doi.org/10.1016/j.ijhydene.2017.05.122
12. Minutillo M, Perna A, Forcina A, Di Micco S, Jannelli E. Analyzing the levelized cost of hydrogen in refueling stations with on-site hydrogen production via water electrolysis in the italian scenario. International Journal of Hydrogen Energy 46:13667-13677 (2021). https://doi.org/10.1016/j.ijhydene.2020.11.110
13. Micena RP, Llerena-Pizarro OR, de Souza TM, Silveira JL. Solar-powered hydrogen refueling stations: a techno-economic analysis. International Journal of Hydrogen Energy 45:2308-2318 (2020). https://doi.org/10.1016/j.ijhydene.2019.11.092
14. US-DOE & NETL, “Assessment of hydrogen production with CO2 capture volume 1: baseline state-of- the-art plants, ” *US-Department of Energy/National Energy Technical Laboratory-2010/1434*, vol. 1, 2011, [Online]. Available: https://www.osti.gov/servlets/purl/1767148
15. Trading Economics, “Natural Gas Price.” Accessed: Jan. 14, 2026. [Online]. Available: https://tradingeconomics.com/commodity/natural-gas,
16. Smith E, Morris J, Kheshgi H, Teletzke G, Herzog H, Paltsev S. The cost of CO2 transport and storage in global integrated assessment modeling. International Journal of Greenhouse Gas Control 109:103367 (2021). https://doi.org/10.1016/j.ijggc.2021.103367
17. EIA, “Electric Power Monthly.” Accessed: Jan. 14, 2026. [Online]. Available: https://www.eia.gov/electricity/monthly/epm_table_grapher.php?t=epmt_5_6_a.
18. US EPA, “GHG Emission Factors Hub.” Accessed: Jan. 14, 2026. [Online]. Available: https://www.epa.gov/climateleadership/ghg-emission-factors-hub.
19. Ayodele TR, Mosetlhe TC, Yusuff AA, Ntombela M. Optimal design of wind-powered hydrogen refuelling station for some selected cities of south africa. International Journal of Hydrogen Energy 46:24919-24930 (2021). https://doi.org/10.1016/j.ijhydene.2021.05.059
20. Gökçek M, Paltrinieri N, Liu Y, Badia E, Dokuz A?, Erdo?mu? A, Urhan BB, Yolda? Ö. Optimum sizing of hybrid renewable power systems for on-site hydrogen refuelling stations: case studies from türkiye and spain. International Journal of Hydrogen Energy 59:715-729 (2024). https://doi.org/10.1016/j.ijhydene.2024.02.068
21. Barhoumi EM, Okonkwo PC, Ben Belgacem I, Zghaibeh M, Tlili I. Optimal sizing of photovoltaic systems based green hydrogen refueling stations case study oman. International Journal of Hydrogen Energy 47:31964-31973 (2022). https://doi.org/10.1016/j.ijhydene.2022.07.140
22. Perna A, Minutillo M, Di Micco S, Jannelli E. Design and costs analysis of hydrogen refuelling stations based on different hydrogen sources and plant configurations. Energies 15:541 (2022). https://doi.org/10.3390/en15020541
23. US DOE, “Clean Hydrogen Production Standard Guidance.” Accessed: Jan. 14, 2026. [Online].

Available: https://www.hydrogen.energy.gov/library/policies-acts/clean-hydrogen-production-standard.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Distributed low-carbon hydrogen for freight corridors: siting hydrogen refueling station with onsite production on New England highways

Adrian R. Irhamna[ab], Burcu Beykal[a], and George M. Bollas[a]*

[a] Department of Chemical and Biomolecular Engineering, University of Connecticut, Storrs, CT 06269, USA.
[b] Faculty of Mechanical and Aerospace Engineering, Institut Teknologi Bandung, Jawa Barat, 40132, Indonesia.
* Corresponding Author: george.bollas@uconn.edu.

ABSTRACT

This work presents an integrated geospatial-technoeconomic optimization framework for siting modular blue and green hydrogen production units co-located with hydrogen refueling stations (HRS) along U.S highways, with a case study focused on New England. The workflow identifies geospatial highway networks and natural gas infrastructure intersections, estimates hydrogen demand based on heavy-duty truck flows from U.S. Freight Analysis Framework, and formulates a mixed-integer linear program (MILP) that selects technology candidates and their capacities to minimize total cost, subject to corridor coverage and supply-demand constraints. Two onsite hydrogen production scenarios are evaluated: a green hydrogen-only production case and a mixed configuration combining modular green and blue hydrogen. Results indicate that, under a 5% hydrogen adoption scenario in truck traffic, 29 HRS with onsite hydrogen production are needed in the New England region. These findings highlight the benefits of integrating local hydrogen production with HRS planning to reduce the reliance on centralized hydrogen delivery and storage infrastructure.

Keywords: Low-carbon hydrogen, hydrogen refueling stations, onsite production, highways, geospatial-technoeconomic optimization, MILP.

INTRODUCTION

Hydrogen is projected to play a significant role in meeting future energy demand, driven by emerging applications across multiple sectors, with approximately 27% of total hydrogen demand attributed to transportation [1]. Among transportation modes, heavy-duty applications, particularly long-haul trucking, are widely regarded as one the most adopters of hydrogen [2]. This sector currently contributes nearly 30% of global transportation-related CO_2 emissions [3], making deep decarbonization particularly challenging with conventional electrification approaches alone.

The large-scale deployment of low-carbon hydrogen for transportation must be supported by the development of reliable hydrogen infrastructure capable of delivering fuel at points of use. However, centralized hydrogen production requires hydrogen storage and long-distance transportation, both of which remain technically challenging [4, 5], creating a major bottleneck for hydrogen market penetration. As an alternative, distributed hydrogen production tailored to local demand has been proposed as a viable alternative. This approach is especially well aligned with transportation applications, which require networks of hydrogen refueling stations (HRS) along heavy-duty vehicles corridors.

Onsite hydrogen production can be achieved through either green hydrogen or blue hydrogen, depending on local resource availability and infrastructure characteristics. While green hydrogen offers superior environmental performance, its production cost remains less favorable than that of blue hydrogen [6]. The economic advantage of blue hydrogen, however, depends critically on minimizing methane leakage and achieving high CO_2 capture efficiency throughout the production process [7]. Leveraging the existing infrastructure of natural gas pipeline networks can further increase the com-

petitiveness of onsite blue hydrogen production by potentially eliminating the need for hydrogen transportation.

In our previous work, we developed and optimized a fixed-bed chemical looping reactor for hydrogen production that generates high-purity hydrogen while simultaneously producing concentrated CO_2 and N_2 streams in three separate outlets [8, 9]. The system is designed for small-scale operation to support distributed production scenarios that utilize existing natural gas infrastructure. A techno-economic assessment of this technology estimated hydrogen production costs in the range of 2.8 – 4.0 USD kgH_2^{-1} [9, 10]. Although these costs are higher than those reported for large-scale centralized blue hydrogen plants, it is important to note that onsite production eliminates downstream transportation and distribution costs, resulting in a delivered hydrogen cost that is closer to the price paid by end users.

In this work, we develop an integrated geospatial-technoeconomic optimization framework to site modular blue and green hydrogen units co-located with HRS along U.S. highway corridors. Modular blue hydrogen uses the blue H_2-N_2 technology developed in prior work [8, 9], while modular green hydrogen is supposed to be supplied via renewable-powered electrolysis. Heavy-duty truck flow and highway networks are extracted from the Freight Analysis Framework (FAF) [11], and the US Energy Information Administration (EIA) natural gas pipeline network [12] is used to assess proximity between highways and existing pipeline infrastructure. A mixed-integer linear programming (MILP) model is formulated to optimally select facility locations and technology capacities that minimize total system cost while satisfying corridor coverage and supply-demand constraints.

This study addresses a key gap in the literature by providing a system-level framework that explicitly integrates existing natural gas infrastructure with onsite hydrogen production to support heavy-duty hydrogen demand along US highway networks.

SYSTEM DESCRIPTION

The Blue H_2-N_2 Production System

Figure 1 illustrates the novel blue H_2-N_2 production system discussed in this work. The system consists of three major subsystems: (i) multi-reformer system; (ii) a heat recovery system; and (iii) a gas purification system. The multi-reformer system uses the intensified reforming reactor discussed in [8], which operates based on the chemical-looping reforming (CLR) principle and consists of successive oxidation (OX), reduction (RED), and reforming (REF) stages.

To address the intrinsic operational discontinuity of a CLR fixed-bed reactor, a minimum of three reactors is used to achieve a semi-continuous product stream. This multiple-reactor configuration generates three distinct outlet streams: a N_2-rich stream, a CO_2-rich stream, and a H_2-rich stream. Additional details of the system design and operation are provided in [9].

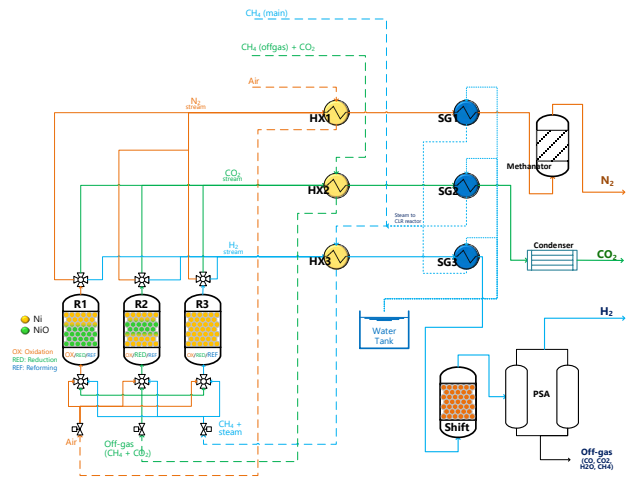

Figure 1. Simplified flow diagram of the blue H_2-N_2 production system considered in this study.

Hydrogen Refueling Station (HRS)

The HRS configuration comprises four main components: hydrogen compressors, a three-stage cascade hydrogen storage system, a hydrogen dispenser and a cooling system, following configuration reported in [13]. The HRS is designed to accommodate onsite hydrogen production from both blue and green hydrogen pathways.

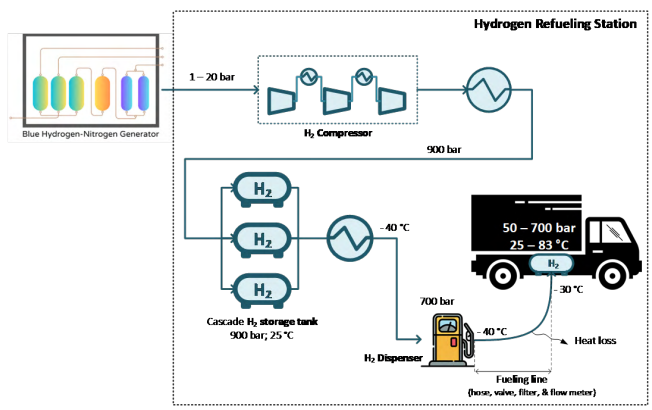

Figure 2. Simplified schematic of the integrated blue H_2-N_2 generator co-located with HRS.

Supply Chain Schematic

Figure 3 illustrates the overall hydrogen supply chain system considered in this study. The system boundary extends from energy primary resources, including renewable electricity as well as natural gas pipelines, through onsite hydrogen production units, and finally to hydrogen demand centers and CO_2 storage sites.

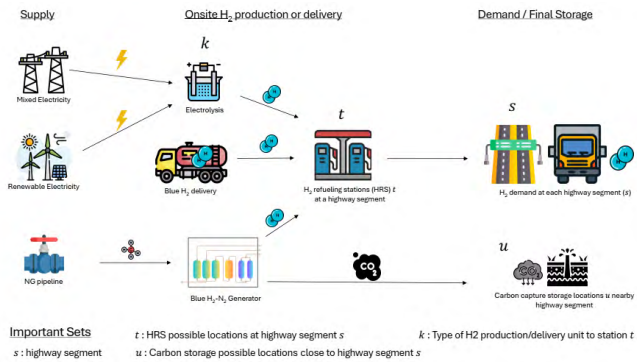

Figure 3. Overall hydrogen supply chain schematic discussed in this work.

METHODOLOGY

Figure 4. FAF highway network and source of truck flow data overlaid with natural gas pipeline infrastructure.

This study uses the Freight Analysis Framework (FAF) dataset provided by the U.S. Department of Transportation [11], which contains detailed information on interstate highway networks and heavy-duty truck flow statistics. Natural gas pipeline data are obtained from the U.S. Energy Information Administration (EIA) [12]. All datasets are spatially aligned and integrated within a unified geographic information system (GIS) environment, QGIS [14], as illustrated in Figure 4.

This study focuses on the New England region, comprising Connecticut (CT), Massachusetts (MA), Rhode Islands (RI), New Hampshire (NH), Vermont (VT), and Maine (ME). This region is selected as a representative case without access to large-scale carbon capture, utilization, and storage (CCUS) infrastructure and because it encompasses the northern extent of the major Interstate 95 (I-95) freight corridor along the U.S. East Coast. Figure 5 shows the major interstates highways across New England considered in this study, while Figure 6 depicts annual heavy-duty truck traffic volumes across the region.

The FAF dataset is available at the level of individual interstate segments, including segment identifiers, lengths, route names, truck traffic volumes, etc. Based on these data, a structured data preparation workflow is developed to transform the raw network datasets into a form suitable for optimization. The major steps are summarized below:

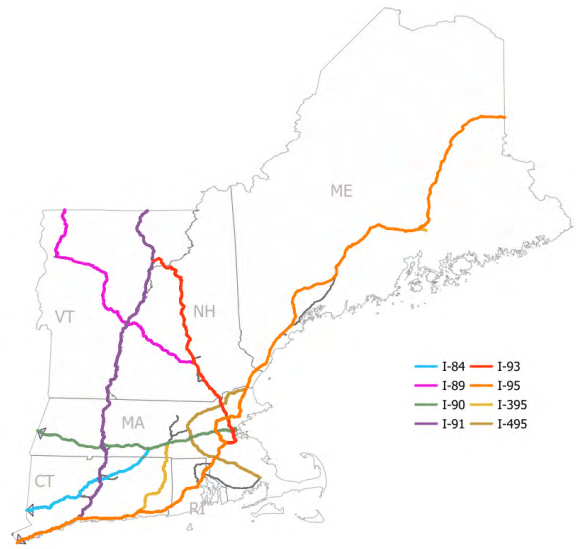

Figure 5. Major interstate highway across New England discussed in this study. For reference, segment labels are ordered along each interstate from south to north, and west to east within the map.

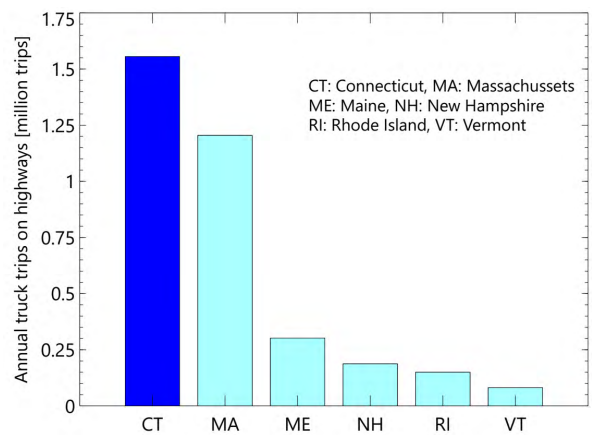

Figure 6. Annual heavy-duty truck trips across New England based on FAF data [11].

1. Segment labeling and ordering
   Each interstate is sorted from upstream (starting node) to downstream (terminal node), and every segment is assigned a unique identifier. For example, segments along Interstate 84 (I-84) are labeled as: `I84_s1, I84_s2, …, I84_sN`.
2. Identification of interstate branches
   For interstate that intersect, segment on each route is assigned connectivity information to represent branching and cross-interstate connections.
3. Hydrogen demand estimation

Hydrogen demand for each highway segment is estimated from segment-level truck flow data by accounting for segment length and fuel consumption:

$$demand_s = \alpha \times TruckFlows \times length_s \times FuelConsumption \quad (1)$$

where α represents the hydrogen adoption rate. In this study, a 5% adoption rate is assumed, with a fuel consumption of 15 kg H_2 per mile for hydrogen-powered trucks.

4. Selection of HRS candidate locations
   Candidate HRS locations are selected based on proximity to existing natural gas pipelines. Using GIS tools, highway segments whose centroids lie within 0.04° longitude-latitude radius of a natural gas pipeline are identified as candidate sites.
5. Construction of the influence matrix
   An influence matrix is constructed to quantify the ability of a candidate station t to serve hydrogen demand at highway segment s. Influence matrix, infl_{st}, ranges from 0 to 1 and is defined based on the distance between segment s and candidate station t:
   - Self-influence (same segment):

$$\mathrm{infl}_{st} = 1.0 \quad (2)$$

   - Different segments:

$$\mathrm{infl}_{st} = \left(1 - {}^{d}/_{radius_{max}}\right) \bullet f_{\mathrm{connect}} \quad (3)$$

     where $f_{\mathrm{connect}} = 1.0$ if segment s and candidate station t are on the same interstate, and $f_{\mathrm{connect}} = 0.8$ if they are on different but directly connected interstates.
   - Beyond maximum service radius:
     $\mathrm{infl}_{st} = 0$ for distance greater than 125 miles.

   To reduce computational complexity, the influence matrix is pruned using the following rules:
   - Remove influence pairs with $\mathrm{infl}_{st} \leq 0.1$;
   - Ensure each segment s retains at least one candidate station t;
   - Limit the number of influence pairs to at most 40 per segment s;
   - Always retain the top 5% strongest influence values;
   - Enforce a minimum spacing of at least 25 miles between candidate stations.

After dataset processing, the MILP optimization problem is formulated. Tables 1 and 2 summarize the sets, indices, decision variables, and parameters used in the model.

The objective of the optimization model is to determine the optimal configuration of HRS, including their capacity levels, the selection of onsite hydrogen production, and the associated CO_2 transport to nearby CCUS locations, while satisfying regional hydrogen demand and infrastructure constraints. The total system cost, TC, is defined as:

$$\begin{aligned} TC = \Big(& \textstyle\sum_{t\in T}\sum_{c\in C} y_{tc}^{\mathrm{HRS}} \mathrm{cap}_c \mathrm{Cost}_c^{\mathrm{station}} \\ & + \textstyle\sum_{t\in T}\sum_{c\in C}\sum_{k\in K} y_{tck}^{\mathrm{H2unit}} \mathrm{cap}_c \mathrm{Cost}_{ck}^{\mathrm{H2unit}} \\ & + \textstyle\sum_{t\in T} x_t^{\mathrm{add}} \mathrm{Cost}^{\mathrm{add}} \\ & + \textstyle\sum_{(t,u)\in TU} f_{tu}^{\mathrm{CO2}} \mathrm{Cost}^{\mathrm{CO2del}} \Big) \\ & \bullet 365\ {}^{\mathrm{days}}/_{\mathrm{year}} \end{aligned} \quad (4)$$

TC represents the annualized total cost of the HRS with onsite production, including both capital and operating expenditures, assuming 20-year project lifetime. The total cost accounts for hydrogen production, hydrogen compression, and CO_2 transportation to storage.

Table 1: Sets and indices used in the MILP formulation for HRS placement with onsite production on the U.S. highways.

Sets and indices	Description
$s \in S$	Set of all highway segments; index s
$t \in T \subseteq S$	Set of candidate HRS locations; index t
$u \in U \subseteq S$	Set of CCUS locations mapped to highway segments; index u
$ST = \{(s,t)\epsilon S \times T\}$	Set of feasible influence pairs between demand segment s and candidate locations t
$TU = \{(s,t)\epsilon S \times T\}$	Set of feasible transport pairs between candidate locations t and CCUS site u
$c \in C = \{0.1, 0.5, 1.0, 2.0, 4.0\}$	Set of HRS capacity levels [tpd]; index c
$k \in K = \{\mathrm{CLR}, \mathrm{ELRE}\}$	Set of H_2 production unit types, CLR-based (blue H_2) or ELRE (electrolysis-renewable energy); index k
$K_{NG} = \{\mathrm{CLR}\} \subseteq K$	Subset of natural gas based H_2 production technologies

The optimization minimizes the total system cost by jointly selecting the capacity of HRS installation the station, y_{tc}^{HRS}, and the type of hydrogen production unit, $y_{tck}^{\mathrm{H2unit}}$. Station capacities are discretized at 0.1, 0.5, 1.0, 2.0, and 4.0 tH_2pd. A continuous variable, x_t^{add}, is introduced to relax the discrete capacity selection and ensure demand satisfaction for segments with hydrogen demand levels that fall between the predefined capacity options.

The optimization is subject to a set of supply-demand, capacity, infrastructure, and emissions constraints. Hydrogen demand at each highway segment $s \in$

S must be satisfied by nearby candidate locations t, as defined by the influence set ST:

$$\text{demand}_s = \sum_{t:(s,t)\in ST} \text{supply}_{st}\ ;\ \forall\, s \in S \tag{5}$$

Table 2: Variables and parameters used in the MILP formulation for HRS placement.

Variables/ Parameters	Description
Decision Variables	
$y_{tc}^{\text{HRS}} \in \{0,1\}$	1 if a HRS of capacity c is installed at candidate station t
$y_{tck}^{\text{H2unit}} \in \{0,1\}$	1 if a H_2 unit type k of capacity c is installed at candidate station t
$x_t^{\text{add}} \geq 0$	Additional low-carbon H_2 delivered to candidate station t
State Variables	
avail_t	Total H_2 available at candidate station t
supply_{st}	H2 supplied from candidate station t to segment s
f_{tu}^{CO2}	CO_2 transported from candidate station t to segment s
Parameters	
demand_s	H_2 demand at segment s
$\text{infl}_{st} \in [0,1]$	Influence factor of station t on meeting demand at segment s
cap_c	Nominal HRS capacity level c (tpd)
$\text{Cost}_c^{\text{station}}$	Cost associated with a HRS of capacity c
$\text{Cost}_{ck}^{\text{H2unit}}$	Cost of H_2 production unit type k for capacity c
Cost^{add}	Cost of additional delivered low-carbon H_2
$\text{Cost}^{\text{CO2del}}$	Cost per unit of CO_2 transported from HRS t to CCUS site u
$\text{loc}_t^{\text{HRS}} \in \{0,1\}$	1 if segment t is eligible to host a new HRS
M	Big-M parameter (sufficiently large constant)
ρ^{NGH2}	Natural gas consumption per unit of H_2 produced (CLR technology)
$\text{NG}^{\text{extract}}$	Fractional limit of extractable natural gas from the nearby pipeline
$\text{NG}_t^{\text{avail}}$	Natural gas availability at candidate station t
$\text{EF}_{ck}^{\text{H2unit}}$	Emission factor of H_2 production/delivery unit k at capacity c
EF^{add}	Emission factor associated with additional delivered H_2
$\text{Em}^{\max}$	Maximum allowable emission intensity: 4 $kgCO_2\ kgH_2^{-1}$
ccusdist_{tu}	Distance between HRS candidate t and CCUS site u

Hydrogen supply from station t to segment s is permitted only if the station lies within the influence of the segment:

$$\text{supply}_{st} \leq \text{infl}_{st}\text{avail}_t\ ;\ \forall\, s \in S, t \in T \tag{6}$$

The total hydrogen supplied by a station cannot exceed its available hydrogen:

$$\sum_{s:(s,t)\in ST} \text{supply}_{st} \leq \text{avail}_t;\ \forall\, t \in T \tag{7}$$

Hydrogen availability at a station t is defined as the sum of onsite production and externally delivered hydrogen:

$$\text{avail}_t \leq M \sum_{c\in C} \sum_{k\in K} y_{tck}^{\text{H2unit}}\ ;\ \forall\, t \in T \tag{8}$$

Available hydrogen at station t is defined as the sum of onsite production:

$$\text{avail}_t = \sum_{c\in C} \sum_{k\in K} y_{tck}^{\text{H2unit}}\, \text{cap}_c + x_t^{\text{add}};\ \forall\, t \in T \tag{9}$$

Additional hydrogen delivery is permitted only when a production unit exists and is limited by the installed station capacity:

$$x_t^{\text{add}} \leq M \sum_{c\in C} \sum_{k\in K} y_{tck}^{\text{H2unit}}\ ;\ \forall\, t \in T \tag{10}$$

$$x_t^{\text{add}} \leq M \sum_{c\in C} \sum_{k\in K} y_{tc}^{\text{HRS}} \text{cap}_c\ ;\ \forall\, t \in T \tag{11}$$

A station may only be constructed at a valid candidate location:

$$\sum_{k\in K} y_{tck}^{\text{H2unit}} \leq \text{loc}_t^{\text{HRS}};\ \forall\, t \in T \tag{12}$$

Furthermore, a hydrogen production unit of capacity c may be installed only if corresponding station capacity is selected, and at most one production unit can be installed at each station:

$$\sum_{c\in C} \sum_{k\in K} y_{tck}^{\text{H2unit}} \leq y_{tc}^{\text{HRS}};\ \forall\, t \in T \tag{13}$$

$$\sum_{c\in C} \sum_{k\in K} y_{tck}^{\text{H2unit}} \leq 1;\ \forall\, t \in T \tag{14}$$

For blue hydrogen production route K_{NG}, natural gas consumption must not exceed the available extraction limit:

$$\sum_{c\in C} \sum_{k\in K_{NG}} y_{tck}^{\text{H2unit}} \text{cap}_c \rho^{\text{NGH2}} \leq \text{NG}^{\text{extract}} \text{NG}_t^{\text{avail}};\quad \forall\, t \in T \tag{15}$$

System-level lifecycle emissions from all hydrogen production routes are constrained by a maximum emissions intensity threshold:

$$\sum_{t\in T}\left(\sum_{c\in C} \sum_{k\in K} y_{tck}^{\text{H2unit}} \text{cap}_c \text{EF}_{tck}^{\text{H2unit}} + x_t^{\text{add}} \text{EF}^{\text{add}}\right) \leq \sum_{t\in T}\left(\sum_{c\in C} \sum_{k\in K} y_{tck}^{\text{H2unit}} \text{cap}_c + x_t^{\text{add}}\right) \text{Em}^{\max} \tag{16}$$

For stations using blue hydrogen production, all generated CO_2 must be transported to feasible CCUS locations:

$$\sum_{u:(t,u)\in TU} f_{tu}^{CO2} = \sum_{c\in C} \sum_{k\in K_{NG}} y_{tck}^{\text{H2unit}} \text{cap}_c \text{EF}_{tck}^{\text{H2unit}}\ ;\quad \forall\, t \in T \tag{17}$$

Since no potential CCUS exist within New England, captured CO_2 is transported outside the regional boundary

via the first segments of Interstate 84 (I84_s1), Interstate 90 (I90_s1), or I-95 (I95_s1). Tables 3 and 4 summarize the key parameters used in the MILP formulation. Cost and emission factors for the blue H_2N_2 generators follows our prior work in [9, 10], while the electrolysis from renewable are adopted from the literature [6, 15].

Table 3: Estimated costs and emission factors (EF) for HRS and the onsite hydrogen production unit at different capacities. Reported costs include both capital and operating expenses. Detailed regarding this is provided in [20].

Capacities	HRS-only		H_2 prod. unit	
	CLR	ELRE	CLR	ELRE
Cost [USD kgH_2^{-1}]				
0.1 tpd	2.683		8.188	8.079
0.5 tpd	2.016		4.264	8.079
1.0 tpd	1.948		3.066	8.371
2.0 tpd	1.948		2.873	8.079
4.0 tpd	1.948		2.735	8.074
Emission Factor [USD kgH_2^{-1}]				
0.1 tpd	2.2	--	0.84	--
0.5 tpd	2.2	--	0.84	--
1.0 tpd	2.2	--	0.84	--
2.0 tpd	2.2	--	0.84	--
4.0 tpd	2.2	--	0.84	--

Table 4: Other important parameters used in the formulation.

Parameters	Values	Units	Refs
NG_t^{avail}	4.16×10^3	tNGpd	[16]
$\rho^{NG\text{-}H2}$	1.23	kgNG kgH_2^{-1}	[9, 10]
$NG^{extract}$	0.1	-	assumption
Em^{max}	4.0	$kgCO_2$ kgH_2^{-1}	[17]
$Cost^{CO2del}$			
- dist. < 100 mi.	9.7	USD tCO_2^{-1}	[18]
- dist. 100-200 mi.	12.6	USD tCO_2^{-1}	[18]
- dist. 200 – 500 mi	24.1	USD tCO_2^{-1}	[18]

To evaluate the impact of different hydrogen pathways, two HRS onsite production scenarios are considered:

1. IDEAL

   This scenario represents a fully renewable, zero-carbon pathway in which only onsite electrolyzers powered by renewable electricity are permitted at all HRS locations.

2. MIXED

   This scenario allows HRS to be supplied with low-carbon hydrogen from either electrolyzers-renewable electricity or blue H_2-N_2 generators.

The MILP optimization problem is implemented in GAMS (version 52.3.0) [19] and solved using the CPLEX. The optimization results identify the selected highway segments for HRS deployment, along with the corresponding station capacities and hydrogen production unit. These results are subsequently visualized using GIS-based mapping tools.

RESULTS AND DISCUSSION

Figures 7 and 8 illustrate the optimal placement of HRS including their associated onsite production technologies, for the two evaluated scenarios: IDEAL and MIXED. Despite the different technology assumptions, both scenarios identify the same 29 HRS locations across New England. Of these, 12 stations are located in MA, 6 in CT, 5 in ME, 3 in NH, 2 in VT, and 1 in RI. While the total number and spatial distribution of HRS remain identical, the allocation of station capacities differs between the two scenarios.

HRS deployment is more densely concentrated in CT and MA, driven by high truck traffic volumes along corridors. In contrast, VT and NH host only five stations in total, despite their long interstate corridors (I-89 and I-93), primarily due to limited access to natural gas pipeline infrastructure. Notably, both scenarios consistently identify the largest HRS (4.0 tH_2pd) at segment I89_s92 in VT, which serves hydrogen demand associated with truck traffic through the mountainous regions of VT and NH.

Figure 8 further indicates that CLR-based onsite production is favored in high-demand regions of CT and MA in the MIXED scenario, owing to both demand density and access to CO_2 transport routes exiting the region via I95_s1 and I84_s1. Conversely, electrolyzer-based hydrogen production is selected in more remote locations, such as along I-89 in northern VT, where access to CO_2 transport options is limited.

From a cost perspective, the IDEAL scenario results in an annualized total system cost of 81.19 million USD $year^{-1}$, whereas the MIXED scenario reduces this cost to 53.62 million USD $year^{-1}$. The corresponding levelized cost of hydrogen (LCOH) at the station level is 9.24 USD kgH_2^{-1} for the IDEAL scenario and 5.96 USD kgH_2^{-1} for the MIXED scenario. As illustrated in Figure 9, this cost reduction is accompanied by increased CO_2 emissions. The IDEAL scenario achieves lower emissions at higher cost, while the MIXED scenario offers improved economic performance at the expense of higher CO_2 intensity, highlighting a clear trade-off between cost and emissions across hydrogen production pathways.

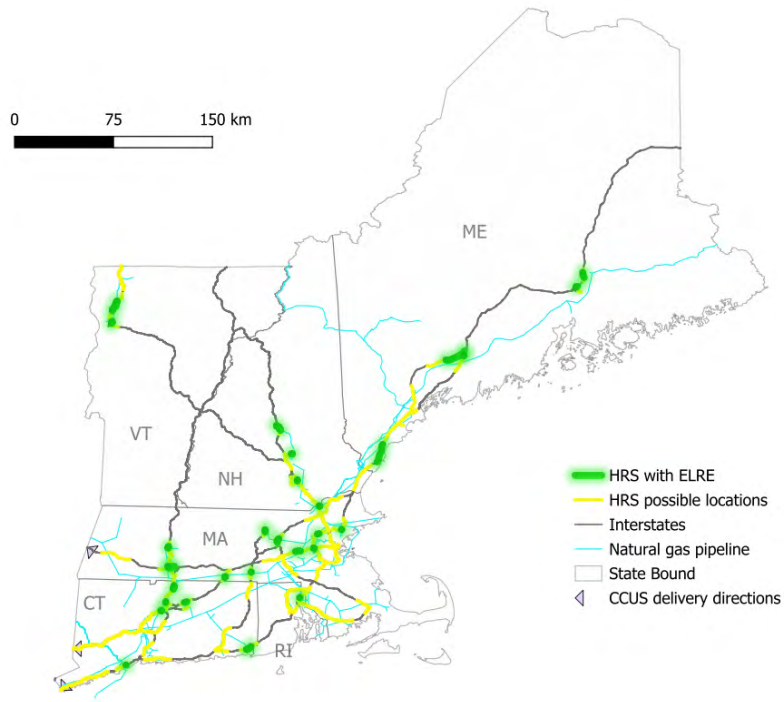

Figure 7. Optimal HRS siting with onsite hydrogen production across New England for the IDEAL scenario, which permits only low-carbon H_2 via renewable-powered electrolysis.

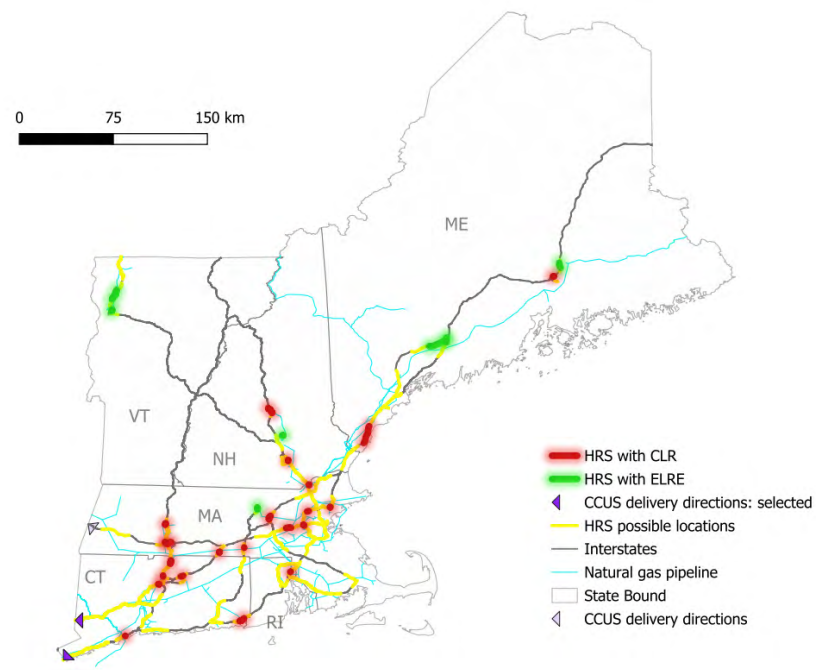

Figure 8. Optimal HRS siting with onsite hydrogen production across New England for the MIXED scenario, which allows low-carbon H_2 from both renewable-powered electrolysis (green) and blue H_2-N_2 generators (red).

CONCLUSIONS

The paper presents the workflow integrating geospatial of highway networks and natural gas intersections, freight activities, and a mixed-integer linear programming (MILP) model that selects sites and sizes HRS

and H_2 technologies. By minimizing the total cost, 29 HRS with onsite production were identified in New England to cover H_2 demand from 5% of total truck flowing through the region. Most HRS are concentrated in CT and MA due to their high truck flow. The final H_2 cost that must be paid by the users varies from 5.96 to 9.24 USD kgH_2^{-1} depending on the onsite H_2 production. The IDEAL and MIXED scenarios highlight the trade-offs that must be considered between achieving cost-effectiveness and reducing CO_2 emissions.

This study showcases the benefit of onsite production leveraging the existing infrastructure to provide hydrogen at the point of use. Further work should examine additional case studies and regional contexts to evaluate how implementation pathways may vary under different infrastructure and market conditions. Overall, this study provides valuable insights for policymakers in shaping strategies for the decarbonization of long-haul freight transportation.

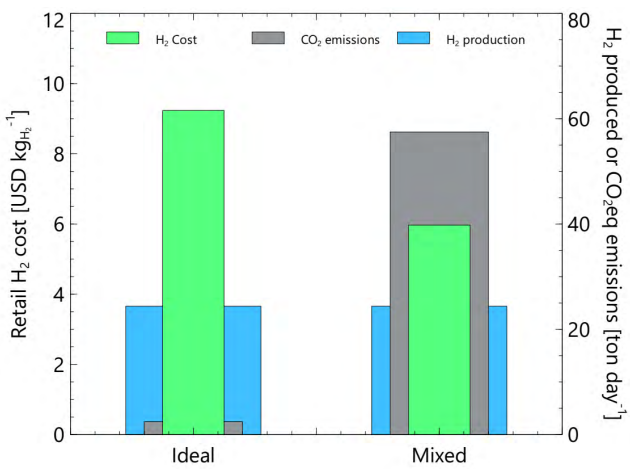

Figure 9. H_2 retail cost at HRS for the IDEAL and MIXED onsite production scenarios [USD kgH_2^{-1}].

AUTHOR IDENTIFIERS

Author ORCIDs:
Irhamna, Adrian: 0000-0001-9752-7017
Beykal, Burcu: 0000-0002-6967-6661
Bollas, George: 0000-0002-1960-431X

REFERENCES

1. International Energy Agency (IEA), "Global hydrogen demand by sector in the Sustainable Development Scenario, 2019-2070, " Paris, 2020.
2. S. Satyapal, "US DOE Hydrogen Program Annual Merit Review (AMR) Plenary Remarks, " 2023. doi: 10.1353/leg.2003.0006.
3. H. Ritchie, "Cars, planes, trains: where do CO_2 emissions from transport come from?, " 2020. [Online]. Available: https://ourworldindata.org/co2-emissions-from-transport#article-citation
4. Andersson J, Grönkvist S. Large-scale storage of hydrogen. International Journal of Hydrogen Energy 44:11901-11919 (2019). https://doi.org/10.1016/j.ijhydene.2019.03.063
5. Usman MR. Hydrogen storage methods: review and current status. Renewable and Sustainable Energy Reviews 167:112743 (2022). https://doi.org/10.1016/j.rser.2022.112743
6. Caponi R, Bocci E, Del Zotto L. Techno-economic model for scaling up of hydrogen refueling stations. Energies 15:7518 (2022). https://doi.org/10.3390/en15207518
7. Mayer P, Ramirez A, Pezzella G, Winter B, Sarathy SM, Gascon J, Bardow A. Blue and green ammonia production: a techno-economic and life cycle assessment perspective. iScience 26:107389 (2023). https://doi.org/10.1016/j.isci.2023.107389
8. Irhamna AR, Bollas GM. Intensified reforming reactor for blue hydrogen and nitrogen production. International Journal of Hydrogen Energy 80:1477-1492 (2024). https://doi.org/10.1016/j.ijhydene.2024.05.428
9. Irhamna, A. R., & Bollas, G. M. Enhancing Hydrogen Infrastructure for Domestic Use: A Small-scale Blue Hydrogen and Nitrogen Production System (under-review)
10. Irhamna AR, Bollas GM. Techno-economic analysis of a novel small-scale blue H2 and N2 production system. Systems and Control Transactions 4:159-165 (2025). https://doi.org/10.69997/sct.134013
11. US Department of Transportation: Federal Highway Administration, "Major Truck Routes on the National Highway System: 2040 Map." [Online]. Available: https://ops.fhwa.dot.gov/freight/freight_analysis/nat_freight_stats/nhsmajortrkrts2040.htm
12. EIA, "Natural Gas Interstate and Intrastate, " 2023. [Online]. Available: https://atlas.eia.gov/datasets/eia::natural-gas-interstate-and-intrastate-pipelines/explore
13. Reddi K, Elgowainy A, Rustagi N, Gupta E. Impact of hydrogen refueling configurations and market parameters on the refueling cost of hydrogen. International Journal of Hydrogen Energy 42:21855-21865 (2017). https://doi.org/10.1016/j.ijhydene.2017.05.122
14. QGIS Association, "QGIS Geographic Information System, " 2026, http://www.qgis.org/, QGIS.org.
15. Moreno-Benito M, Agnolucci P, Papageorgiou LG. Towards a sustainable hydrogen economy: optimisation-based framework for hydrogen infrastructure development. Computers & Chemical Engineering 102:110-127 (2017).

https://doi.org/10.1016/j.compchemeng.2016.08.005
16. US Department of Transportation, “Gas Distribution, Gas Gathering, Gas Transmission, Hazardous Liquids, Liquefied Natural Gas (LNG), and Underground Natural Gas Storage (UNGS) Annual Report Data, ” https://www.phmsa.dot.gov/data-and-statistics/pipeline/gas-distribution-gas-gathering-gas-transmission-hazardous-liquids.
17. U.S. Department of Energy, “U.S. Department of Energy Clean Hydrogen Production Standard (CHPS) Guidance, ” 2023. [Online]. Available: https://greet.es.anl.gov/greet.models.
18. Smith E, Morris J, Kheshgi H, Teletzke G, Herzog H, Paltsev S. The cost of CO2 transport and storage in global integrated assessment modeling. International Journal of Greenhouse Gas Control 109:103367 (2021). https://doi.org/10.1016/j.ijggc.2021.103367
19. GAMS Development Corp., “GAMS, ” 2025, 52.3.0. Accessed: Jan. 21, 2026. [Online]. Available: https://www.gams.com/52/docs/RN_52.html.
20. Irhamna AR, Bollas GM. Techno-economic analysis of a novel small-scale blue H2 and N2 production system. Systems and Control Transactions 4:159-165 (2025). https://doi.org/10.69997/sct.134013

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Superstructure Modelling of Membrane Systems for the Optimization and Flexible Design of Post-combustion Carbon Capture Processes

Stefania Bempeli[a], Marina Micari[a]*

[a] Laboratory of Advanced Separations (LAS), École Polytechnique Fédérale de Lausanne (EPFL), Sion, Switzerland, 1950
* Corresponding Author: marina.micari@epfl.ch.

ABSTRACT

Membranes provide an efficient method for treating flue gases to capture CO_2 from various point sources, achieving high recovery and purity rates. However, the lack of systematic process-level design tools has limited the translation of advanced membrane materials into large-scale technical and economic metrics. Thus, in this study, we present a superstructure model for the design of membrane-based carbon capture, both from highly energy-intensive industries and from power plants. The superstructure model enables the flexible design and global optimization of multi-stage membrane systems. Multiple membranes are compared under technical performance indicators (specific energy and specific area), while the already commercialized polymeric membranes Polaris and PolyActive are taken into consideration for estimating their economic performance. The presented framework establishes a robust link between material innovation and optimal process design, providing a key tool for the large-scale deployment of membrane-based carbon capture.

Keywords: carbon capture, membrane systems, superstructure, optimization

INTRODUCTION

As 2050 approaches, the efforts to achieve the net-zero transition should be intensified. Under these terms, Carbon Capture, Utilization, and Storage (CCUS) technologies should be strengthened across technical, economic, and environmental dimensions and implemented at an industrial scale. Post-combustion carbon capture can assist energy-intensive industries and existing fossil-fuel power plants in their decarbonization. Various technologies have been proposed and evaluated for this purpose, with membrane process systems constituting a promising option [1].

More specifically, membrane systems are characterized by high energy efficiency, compact and modular design, operational simplicity, and significant environmental benefits relative to other state-of-the-art technologies, without the use of heating utilities, which can be both costly and environmentally impactful. Most recent advances focus on material development, while significant gaps remain in optimizing membrane systems that utilize these materials at process level, and measure their performance under industrial scales [2].

In particular, a comprehensive framework that finally links advanced material properties and large-scale process design is still missing. This work addresses this critical gap by introducing a superstructure modelling framework that systematically optimizes membrane systems for post-combustion carbon capture.

METHODOLOGY

Single-stage membrane model

Any membrane network is based on the combination of multiple single stages. To model such a single-stage system, the following assumptions have been made:

- Isothermal conditions along the length of the membrane.
- Negligible pressure drops or concentration polarization phenomena occur on the feed side of the membrane module.

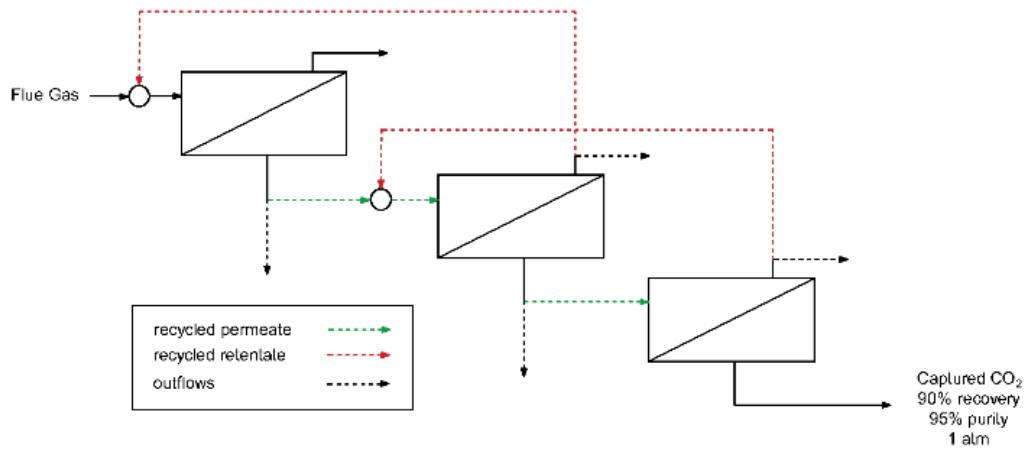

Figure 1. Developed membrane module network.

- A cross-current flow pattern is implemented.

In a single-stage membrane process, the membrane sheet is discretized along its length (z-axis) into a specific number of elements of length dz and area dA. In each one of them, the transmembrane flux of a component in a flue gas can be estimated using the solution-diffusion model.

$$J_i(z) = P_i(P_{feed}x_{feed,i}(z) - P_{perm}x_{perm,i}(z)) \quad (1)$$

where P_i the permeance of the component i, and is a characteristic of the membrane material used, P_{feed} and P_{perm} the pressure applied in the feed and the permeate channel, respectively, $x_{feed,i}$ and $x_{perm,i}$ the molar fractions of the component i in the feed stream and right after crossing the membrane. A multicomponent mixture is considered, so i corresponds to CO_2, N_2, O_2, and H_2O.

The total flux can be estimated by the summation of the fluxes of the individual components (Equation (2)). By implementing the total and partial mass balances among the elements, the rest of the flows and compositions can be determined, through Equations (3-6):

$$J_{tot}(z) = \sum_{i=1}^{components} J_i(z) \quad (2)$$

$$Q_{ret}(z) = Q_{feed}(z) - J_{tot}(z)\, dA \quad (3)$$

$$x_{i,ret}(z)\, Q_{ret}(z) = x_{i,feed}(z)\, Q_{feed}(z) - J_i(z)\, dA \quad (4)$$

$$Q_{perm} = \sum_{z=0}^{L} J_{tot}(z)dA \quad (5)$$

$$X_{i,perm} = \frac{\sum_{z=0}^{L} J_i(z)dA}{Q_{perm}} \quad (6)$$

where Q_{feed} and x_{ifeed}, Q_{ret} and $x_{i,ret}$ and Q_{perm} and $x_{i,perm}$ the flow rates and the molar fractions comprising the feed the retentate and the permeated side streams [3, 4].

Membrane module network

The implemented superstructure model is comprised of 3 membrane stages. To model this network, Equations (1)-(6) are developed once for each stage, and additional mass balances are introduced to specify the connections between them. The retentate stream of each stage can either exit the system or be recycled to the feed of the previous stage. The retentate from the 1st-stage has the lowest CO_2 content, thus it is assumed to exit the system directly. For the permeate streams concentrated in CO_2, the option of feeding them to the next stage or bypassing this route is available for all 3 stages. This feature enables the determination of the presence or the absence of a membrane stage without introducing integer variables, despite the discrete nature of the decision [3, 5].

In Figure 1, a schematic of the developed network is shown, with all allowable connections depicted. Dashed lines correspond to plausible flows, whose presence or absence gives the final configuration of the system, and is a subject of the optimization performed.

To increase the driving force at each stage, compressors and vacuum pumps are installed on the feed and permeate sides, respectively. To determine their energy consumption (kW), Equation (7) is applied to either the compressors (c) or the vacuum pumps (vp), using the corresponding efficiency and flow rate.

$$W_{c/vp} = \frac{10^{-3}}{\eta_{c/vp}} \cdot \frac{\gamma}{\gamma - 1} \cdot RTQ \cdot \left[\left(\frac{P_{high}}{P_{low}}\right)^{\frac{\gamma-1}{\gamma}} - 1\right] \quad (7)$$

Turbines are also included in the network at the side

of the retentate streams that either exit the system or being recycled, to recover part of the electricity used for compression. The recovered energy (kW) is calculated through Equation 8.

$$W_{turb} = \frac{\eta_{turb}}{10^3} \bullet \frac{\gamma}{\gamma - 1} \bullet RTQ \bullet \left[1 - \left(\frac{P_{low}}{P_{high}}\right)^{\frac{\gamma-1}{\gamma}}\right] \quad (8)$$

The heat capacity ratio γ (-) is determined based on the composition of each gaseous stream, while the efficiency for the compressors, the vacuum pumps and the turbines are equal to 0.85, 0.7 and 0.9 [6].

Key performance indicators

In this analysis, the developed superstructure membrane module network is used to estimate both the technical and economic performance of different membrane materials. Thus, appropriate key performance indicators are defined. Based on them, the evaluation of the designed systems is determined.

The primary key process performance metrics are specific energy (MJ/kg_{CO_2}) and specific area (m^2/(kg_{CO_2}/s)). These two metrics are strongly influenced by the properties of the membrane material and the composition of the flue gas being treated, and they exhibit a trade-off.

$$sE = \frac{W_{total}}{X_{CO_2,perm}\, Q_{perm} MW_{CO_2}} \quad (9)$$

$$sA = \frac{A_{total} 10^3}{X_{CO_2,perm}\, Q_{perm}\, MW_{CO_2}} \quad (10)$$

Moreover, during the design of a carbon capture system regardless of the technology implemented, specific purity and recovery rates should be achieved. For this reason, the corresponding metrics can be calculated through Equations (11-12), respectively. These two metrics are introduced as constraints during the design of the membrane systems studied here.

$$Pu_{CO2} = \frac{X_{CO_2,perm}}{1 - X_{H_2O,perm}} \quad (11)$$

$$Rec_{CO2} = \frac{X_{CO_2,perm} \bullet Q_{perm}}{X_{CO_2,feed} \bullet Q_{feed}} \quad (12)$$

From literature, the purity of the final stream of captured CO_2 should be at least 95% to enable storage or utilization as raw material, while recovery should reach values of 90%.

For cost calculations, the methodology presented in previous works [4, 7] is followed. In the present work, we adopted fixed values for the electricity price and for the membrane material, equal to 0.1$/kWh and 50$/m^2 [4, 7].

Membrane materials and case studies

In Table 1, the material properties of CO_2 permeance (GPU) and the selectivity values for CO_2/N_2, CO_2/O_2, and CO_2/H_2O (-) required by the process model are reported for 5 materials. The selected membranes belong to different material classes, resulting in markedly different combinations of their properties. Two are polymeric and already at commercial scale (Polaris and PolyActive); two are facilitated-transport-based (FSC and FT-PVAm); and one is an inorganic, graphene-based membrane (NSLG, nanoporous single-layer graphene).

Table 1: Material properties of membranes tested during the analysis/optimization [4, 8].

Material	P_{CO_2} (GPU)	a_{CO_2/N_2}	a_{CO_2/O_2}	a_{CO_2/H_2O}
Polaris	2200	50	10	0.1
PolyActive	1450	46	17	0.1
FSC	838	95	21	0.1
FT-PVAm	1456	165	10	0.1
NSLG	10000	30	13	1.0

These materials and their performances are tested for the treatment of a number of flue gases, ranging from highly concentrated point-sources, with CO_2 content of 13-24% (coal power plant, cement plant, oil refinery) to more diluted ones with <8% CO_2 content (natural gas power plant, aluminium industry).

Optimization framework

The aforementioned membrane module network is formulated as a nonlinear programming (NLP) problem in the GAMS-Python environment, GamsPy. No integer variables are introduced into the system of equations, and discrete decisions, such as the presence or absence of a membrane module in the final system configuration, are determined entirely by the model's continuous variables. Despite the absence of integer variables, the final problem remains non-convex due to the nature of the equations that comprise the model of the developed superstructure. Such optimization problems can be difficult to solve and to guarantee global optimality. To prevent such issues, the global solver BARON [9] is used.

Two different optimization problems are solved for each material and case study:

- the bi-objective optimization of specific energy and specific area, by implementing the weighting method, and
- the single objective minimization of the capture cost.

In both cases, optimization is performed subject to two constraints: purity (95%) and recovery (90%), which are implemented as inequality constraints.

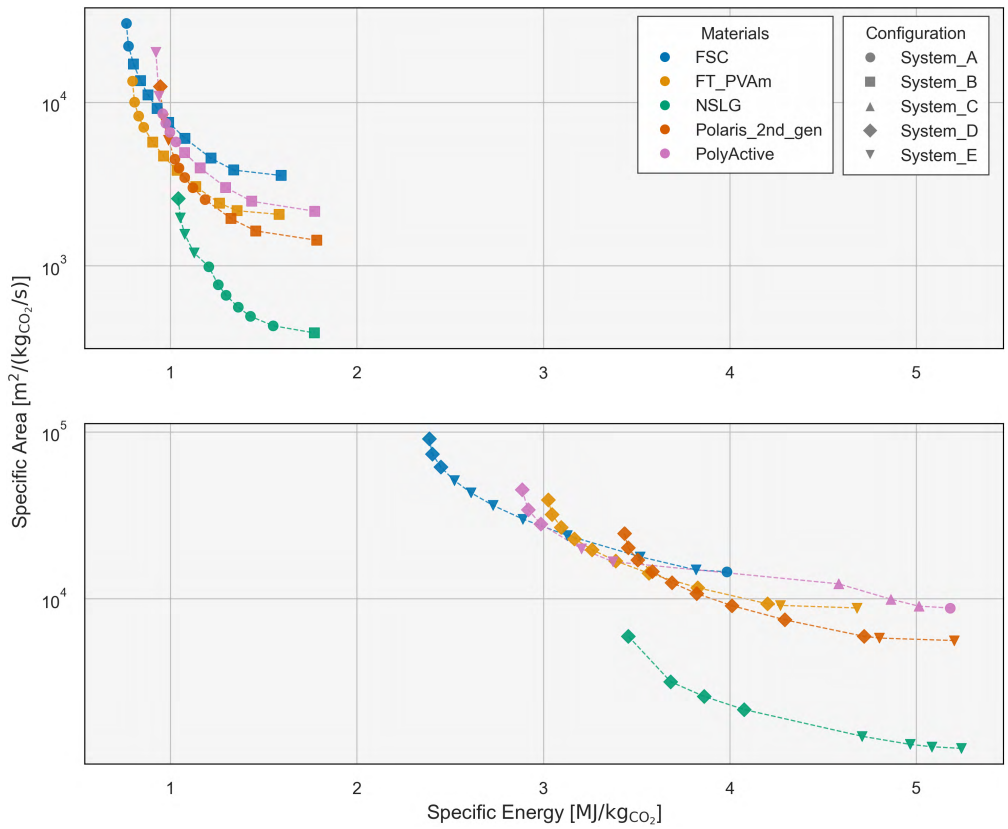

Figure 2: Technical Pareto fronts when treating a flue gas from a coal power plant with composition x_{CO_2}=13.7%, x_{N_2}=72.8%, x_{O_2}=3.6% (up) and from a natural gas power plant with composition x_{CO_2}=3.9%, x_{N_2}=75.2%, x_{O_2}=12.3% (below), under purity and recovery constraints of 905 and 90% respectively. Each curve is comprised of points with a variety of system configurations.

RESULTS AND DISCUSSION

Technical bi-objective optimization

In this series of optimization problems, the Pareto front curves for the specific energy and specific area of the studied membrane materials are generated for all membrane materials listed in Table 1. We report the results for two case studies with high and low CO_2 concentrations (Figure 2). The two objectives are closely correlated with the composition of the treated flue gas and the material properties. By using the superstructure model rather than a fixed-configuration model for membrane process, we produce Pareto fronts along which the points represent optimal configurations with different numbers of membrane stages and connections among them. The exact configurations suggested by the optimization are described in Table 2 and are shown in the Pareto fronts of Figure 2 with different symbols.

This approach finally allows a detailed comparison between materials, and the identification of limiting properties that determine each material's performance depending on the case study.

From this analysis, we observe that the CO_2/O_2 selectivity of the membrane material affects the allowable and feasible configurations on the Pareto front, particularly when the O_2 content is high. For example, in the Pareto fronts of specific energy-specific area of the FSC and FT-PVAm membranes, for the coal power plant case study, all the points correspond to systems with 2 stages thanks to the high selectivity. At the same time, the lower CO_2 permeance of these materials leads to higher specific area values.

For diluted flue gases with CO_2 concentrations < 8%, as in a natural gas power plant, the estimated specific energy and specific area values increase significantly. More specifically, the curves of most materials consist of points corresponding to systems with 3 stages, and only the curves for FSC and PolyActive (which have higher CO_2/O_2 selectivity) contain points corresponding to systems with 2 stages.

Moreover, we notice that some of the tested membranes achieve similar performances in terms of specific energy and area, by adopting different operating conditions or system configurations depending on the membrane properties. However, with diluted flue gases, the

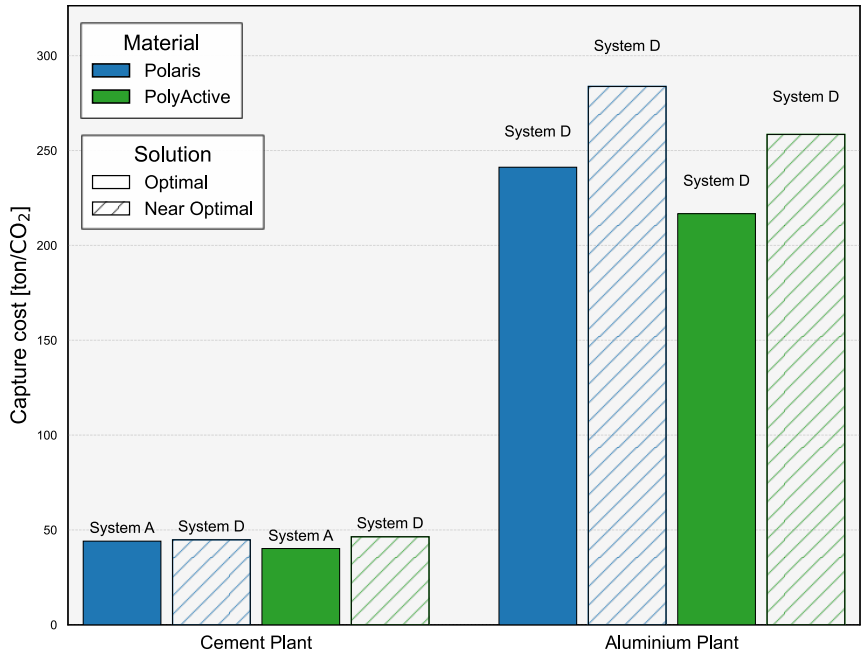

Figure 3. Best two minimum capture cost values for Polaris and PolyActive for capturing CO_2 from a cement plant, and an aluminium plant.

separation is more challenging, and all membranes require more complex configurations to meet the constraints of recovery and purity.

Table 2: Description of the membrane networks suggested after the technical and economic optimization.

Double Stage Systems	
System A	Recycle of the 2nd stage retentate to the 1st stage feed
System B	Recycle of the 2nd stage retentate to the 1st stage feed and bypass of the 1st permeate to the final outlet
System C	Recycle of the 2nd stage retentate to the 1st stage feed and recycle of the 2nd stage permeate to the 2nd stage feed
Triple Stage Systems	
System D	Recycle of the 2nd and 3rd stages retentates to the feed of the previous stage
System E	Recycle of the 2nd and 3rd stages retentates to the feed of the previous stage and bypass of the 2nd stage permeate to the final outlet

Cost minimization

The evaluation of economic performance is limited to the commercialised membrane materials, Polaris and PolyActive, for which price data are available. In this analysis, the capture cost is used as the objective function. In Figure 3, we report the results for two case studies: one with high CO_2 content (cement plant: x_{CO_2}=18%, x_{N_2}=63%, x_{O_2}=10%) and one with lower CO_2 content (aluminium plant: x_{CO_2}=4%, x_{N_2}=74.6%, x_{O_2}=20.5%). In particular, we show the minimum capture cost (optimal), along with the next-best optimal value (near-optimal).

In the case of the Cement plant, with high CO_2 content, the minimum capture cost can be secured with a double-stage system. However, for the membrane material Polaris, a slightly higher capture cost (~2%) can be achieved with a triple-stage system. The analysis of near-optimal solutions can be very helpful during decision-making. Indeed, a triple-stage system has typically lower specific energy consumption, thus it can be preferred in cases of an increase in the electricity price, which is always prone to a plethora of uncertainties, to enhance robustness in the economic performance of the system.

For the more dilute flue gas from aluminium plants, the next-best points for both materials have higher specific energy and specific area values, under the same system configuration. This means that, for challenging separations with low CO_2 and high O_2 concentrations, the points with minimum cost lie at the very end of the Pareto front, in the region of lower specific energy. In this case, achieving a similar capture cost with a different configuration would require adding a 4th stage, which would lower specific energy and increase specific area.

Robustness of the computational framework

The presented framework constitutes a key tool for

evaluating the performance of both established and emerging membrane materials at the process level. To improve the robustness and the computational efficiency of the model simultaneously, we have built the model as a non-linear program, avoiding discrete variables and defining a maximum number of possible recycle and bypass streams based on preliminary analysis. Also, we use the global solver BARON to ensure global optimality, given the non-convex behaviour of the superstructure that introduces many local minima in the solution domain.

Finally, the framework allows to perform both technical and economic optimization of different material properties under multiple case studies. By determining the optimal membrane network and operating conditions for each available material, we can determine the best option for each case study.

These results may also contribute to identify material properties that are desirable to improve the performance of the membrane process and to drive material enhancements in this direction. Finally, the ability to identify alternative designs with similar performance metrics provides the flexibility to consider both the operability of each design in real-life applications and its robustness under uncertainty.

CONCLUSION

We present an inclusive and detailed superstructure model capable of addressing multiple case studies across membrane materials and flue gas compositions. The framework enables us to identify optimal process configurations through fast and robust multi-objective technical and economic optimization. A range of membrane materials with varying material properties is tested, and the optimal configuration at both the structural and operational levels is determined. Through technical optimization, material performance targets are established to improve process feasibility. Such a framework constitutes an important advance in bridging the gap between small-scale material development and large-scale implementation of membrane process systems.

ACKNOWLEDGEMENTS

The authors thank the Swiss National Science Foundation (SNSF) for funding this work via the SNSF Ambizione grant n. 223517.

REFERENCES

1. Chatziasteriou CC, Kikkinides ES, Georgiadis MC. Recent advances on the modeling and optimization of CO2 capture processes. Computers & Chemical Engineering 165:107938 (2022). https://doi.org/10.1016/j.compchemeng.2022.107938
2. Chiwaye N, Majozi T, Daramola MO. Superstructure-based optimization of membrane gas separation processes: a review. Ind. Eng. Chem. Res. 64:13905-13919 (2025). https://doi.org/10.1021/acs.iecr.4c03693
3. Arias AM, Scenna NJ, Mores PL. Membrane superstructure optimization for carbon capture from cement plants. water content influence on optimal solution. International Journal of Greenhouse Gas Control 129:103964 (2023). https://doi.org/10.1016/j.ijggc.2023.103964
4. Micari M, Dakhchoune M, Agrawal KV. Techno-economic assessment of postcombustion carbon capture using high-performance nanoporous single-layer graphene membranes. Journal of Membrane Science 624:119103 (2021). https://doi.org/10.1016/j.memsci.2021.119103
5. Hasan MMF, Baliban RC, Elia JA, Floudas CA. Modeling, simulation, and optimization of postcombustion co₂ capture for variable feed concentration and flow rate. 1. chemical absorption and membrane processes. Ind. Eng. Chem. Res. 51:15642-15664 (2012). https://doi.org/10.1021/ie301571d
6. Micari M, Hsu KJ, Bempeli S, Agrawal KV. Energy- and cost-efficient CO2 capture from dilute emissions by pyridinic-graphene membranes. Nat Sustain 9:164-175 (2025). https://doi.org/10.1038/s41893-025-01696-5
7. Arias AM, Mussati MC, Mores PL, Scenna NJ, Caballero JA, Mussati SF. Optimization of multi-stage membrane systems for CO 2 capture from flue gas. International Journal of Greenhouse Gas Control 53:371-390 (2016). https://doi.org/10.1016/j.ijggc.2016.08.005
8. He X, Chen D, Liang Z, Yang F. Insight and comparison of energy-efficient membrane processes for CO2 capture from flue gases in power plant and energy-intensive industry. Carbon Capture Science & Technology 2:100020 (2022). https://doi.org/10.1016/j.ccst.2021.100020
9. Zhang Y, Sahinidis NV. Solving continuous and discrete nonlinear programs with BARON. Comput Optim Appl 92:1123-1161 (2024). https://doi.org/10.1007/s10589-024-00633-0

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Multi-scenario Optimization of Groundwater-Sourced Water Production Networks With Daily Well Shutdown Requirements

Pedro H. CALLIL-SOARES[a], René P. SCHNEIDER[a], and Galo A. CARRILLO LE ROUX[a]*

[a] University of São Paulo, Department of Chemical Engineering, São Paulo, São Paulo, Brazil
* Corresponding Author: galoroux@usp.br.

ABSTRACT

Water supply in the countryside of São Paulo state, Brazil, is based on groundwater resources that can be contaminated with substances such as heavy metals or fluoride, requiring the usage of water treatment technologies such as Reverse Osmosis (RO); however, RO systems create a stream of high-salinity brine, with negative environmental consequences. Besides, regulatory constraints demand that well operations must be interrupted for a daily contiguous period. In this work, a Mixed-Integer Linear program (MILP) was implemented to define water network topologies and well exploitation schedules, under these downtime constraints, aiming the minimization of RO plant capacity (and, therefore, of brine discharges). This model was then applied to the water supply of a small city in the São Paulo countryside, with around 8000 inhabitants, where high fluoride concentrations warranted the implementation of an RO system. Demand variations between weekdays and weekends (with demands 52.7% higher) were modeled using a multi-scenario formulation. This model was optimized under different conditions of tank capacity and membrane technologies. Different membrane systems (with salt retention rates varying from 85% to 99.5%) associated with different storage tanks (with volumes between 50 and 1000 m^3) led to different system topologies, and a trade-off was observed between plant capacity, tank capacity, and operating pressure: while a low pressure system with a small tank capacity required larger plant capacities (up to 44.02 m^3/h), high pressure systems associated with large storage tanks required smaller plants for treatment (with a capacity as small as 11.15 m^3/h).

Keywords: Water Networks, Membranes, Planning & Scheduling, Optimization, Water, Reverse Osmosis, Multi-scenario Optimization, MILP

INTRODUCTION

Water Production Networks and Reverse Osmosis Plants

In an urban center, the water supply system is comprised of networks of water abstraction, treatment and consumption units. The abstraction units send raw water to the treatment units, where it is converted to drinking water for distribution to the consumers.

The production of drinking water under adequate quality constraints and in sufficient volume to match the demands of the consumers is a key challenge for several urban centers: climate change-fueled water scarcity, growing population, and improvements in living standards increase design and operation requirements of water production and distribution systems. Hence, the enhancement of currently existing systems, the exploration of alternative water resources and the adoption of different water treatment technologies are often needed.

Among these new technologies, the usage of Reverse Osmosis (RO) systems for the treatment of groundwater sources is an attractive possibility. In these membrane-based separation plants, a brackish water stream is split in two: a high-salinity stream (the “reject” or “concentrate”) and a low-salinity stream (the “permeate”). This separation technology is one of the possible

alternatives for the purification of water resources with high concentrations of monovalent ions, such as fluoride, but some disadvantages (such as its high energy needs and the discharge of a considerable volume of brine) warrant careful considerations before implementation [1].

Optimization of Water Production Networks

Due to the widespread presence of water networks, in the most varied engineering fields, several optimization strategies have been applied to their design and operation. These methods can be categorized in two groups: "pinch" approaches, which require less computational resources but don't necessarily lead to an optimal solution, and systematic approaches, in which all mass balances of each equipment in the system become constraints in a mathematical program. This approach can yield demonstrably optimal solutions, and makes the selection of generic objective functions straightforward; however, more computational resources are required [2]. For instance, when the concentrations of contaminants are important to the problem (as is the case in urban water production networks), the resulting mathematical program can be classified as a pooling problem, which must either be solved as a non-convex, nonlinear optimization program, or linearized through the usage of integer variables [3].

Classical analyses of water networks usually focused on systems that operate in a steady-state regime, relevant to industrial contexts. However, in an urban water supply context the dynamic operation of the system needs to be taken into account: not only the water demand is variable, the water resources might only be available in selected periods of the day. Hence, a more sophisticated mathematical tooling is required, and these systems must be modeled through an Optimal Control formulation. In this technique, the temporal domain is discretized, and the optimization variables become the state of each equipment in the system in a selected set of time instants [4].

São Paulo State Legal Requirements

One additional operational requirement, particular to São Paulo state, is the legally-mandated contiguous daily downtime period: every day, pumping from each well operation has to be interrupted for a single period of four hours, to allow the water level in the well to recover. This constraint is absent from previously published models, but is paramount to the water management company, allowing it to replace *ad-hoc* heuristics with results from a deterministic optimization model.

MATERIALS AND METHODS

Mathematical Programming

The optimal control problem was initially converted to a Disjunctive Programming formulation: the optimization horizon was divided in H intervals and the water production network was modeled as a set of equipments (nodes) connected between themselves. For every equipment E, for every instant $h \in \{1, 2, \ldots, H\}$, were defined flowrates Q_E^h (in m³/h) and contaminant loads F_E^h (in g/h) for each of its inputs and outputs. Each equipment of the system was modeled linearly, as shown below, and the inputs and outputs of connected equipments were subjected to equality constraints. The logical constraints were converted to a Mixed-Integer Linear Programming (MILP) formulation, and the resulting program was solved.

Equipments and connections

All equipments were modeled in a straightforward fashion; however, a more sophisticated formulation was required for wells (contiguous downtime requirement), splitters (linearization of bilinearities) and tanks (multi-scenario optimization with the usage of design days).

Wells

For a well W, its downtime fraction D_W was known, as well as its contaminant concentration c_W and its maximum water production $Q_{\max,W}$. Initially, were defined output flowrates Q_W^h and contaminant loads F_W^h for each instant h, as wells as their limits (1 and 2, respectively).

$$0 \leq Q_W^h \leq Q_{\max,W} \forall h \in \{1, 2, \ldots, H\} \tag{1}$$

$$F_W^h = c_W Q_W^h \tag{2}$$

For the downtime modeling, two boolean variables were defined for each instant h: X_W^h was True if and only if the well was inactivated in instant h (3), and Y_W^h was True if and only if X_W^h was True and X_W^{h-1} was False (4).

$$\begin{bmatrix} X_W^h \\ Q_W^h = 0 \end{bmatrix} \vee \begin{bmatrix} \neg X_W^h \\ 0 \leq Q_W^h \leq Q_{\max,W} \end{bmatrix} \tag{3}$$

$$\begin{gathered} X_W^h \wedge \neg X_W^{h-1} \Rightarrow Y_W^h \forall h \in \{2, 3, \ldots, H\} \\ X_W^1 \wedge \neg X_W^H \Rightarrow Y_W^1 \end{gathered} \tag{4}$$

The downtime (5) and its contiguity (6) were constrained through the number of True X_W^h and Y_W^h, respectively.

$$\text{card}\{h \in \{1, 2, \ldots, H\} : X_W^h\} \geq H D_W \tag{5}$$

$$\text{card}\{h \in \{1, 2, \ldots, H\} : Y_W^h\} \leq 1 \tag{6}$$

RO plants and brine discharges

A RO plant P is characterized by its minimum and maximum feed flowrates $Q_{\min,P}$ and $Q_{\max,P}$, its yield η_P, defined as the ratio between permeate and feed flowrates, and its salt retention ρ_P, defined as the rate between brine and feed contaminant loads. For each

instant h, were defined feed, permeate and concentrate flowrates ($Q^h_{\text{in},P}, Q^h_{\text{perm},P}$ and $Q^h_{\text{brine},P}$, respectively) and contaminant loads ($F^h_{\text{in},P}, F^h_{\text{perm},P}$ and $F^h_{\text{brine},P}$, respectively). The yield (7), salt retention (8), mass balances (9) and operation limits (10) were modeled.

$$Q^h_{\text{perm},P} = \eta_P Q^h_{\text{in},P} \tag{7}$$

$$F^h_{\text{brine},P} = \rho_P F^h_{\text{in},P} \tag{8}$$

$$\begin{aligned} Q^h_{\text{in},P} &= Q^h_{\text{perm},P} + Q^h_{\text{brine},P} \\ F^h_{\text{in},P} &= F^h_{\text{perm},P} + F^h_{\text{brine},P} \end{aligned} \tag{9}$$

$$Q_{\min,P} \leq Q^h_{\text{in},P} \leq Q_{\max,P} \tag{10}$$

Each plant requires a brine discharge unit R with input flowrates and contaminant loads (respectively, $Q^h_{\text{in},R}$ and $F^h_{\text{in},R}$) with a maximum contaminant load c_R, which limits brine concentration (11).

$$F^h_{\text{in},R} \leq c_R Q^h_{\text{in},R} \tag{11}$$

Mixers, splitters and selectors

A mixer M is associated with mass balances constraints (12), that ensure that its output flowrates and contaminant loads ($Q^h_{\text{out},M}$ and $F^h_{\text{out},M}$, respectively) match the sum of its n_M input flowrates $i_1, i_2, \ldots, i_{n_M}$ and contaminant loads ($Q^h_{\text{in},i_j,M}$ and $F^h_{\text{in},i_j,M}$, respectively, for $j \in \{1, 2, \ldots, n_M\}$).

$$\begin{aligned} \textstyle\sum_{j=1}^{n_M} Q^h_{\text{in},i_j,M} &= Q^h_{\text{out},M} \\ \textstyle\sum_{j=1}^{n_M} F^h_{\text{in},i_j,M} &= F^h_{\text{out},M} \end{aligned} \tag{12}$$

Splitters, on the other hand, required a more sophisticated model. For a splitter, S, with two output streams (o_1 and o_2), mass balances were introduced through constraint (13), but the same-concentration constraint (14) had to be replaced by a linear formulation.

$$\begin{aligned} Q^h_{\text{in},S} &= Q^h_{\text{out},o_1,S} + Q^h_{\text{out},o_2,S} \\ F^h_{\text{in},S} &= F^h_{\text{out},o_1,S} + F^h_{\text{out},o_2,S} \end{aligned} \tag{13}$$

$$Q^h_{\text{in},S} F^h_{\text{out},o_1,S} = Q^h_{\text{out},o_1,S} F^h_{\text{in},S} \tag{14}$$

Constraint (14) was linearized with a logarithmic piecewise discretization [5]. Initially, the number $G_S = 2^{N_S}$ of possible split fractions was chosen, with $N_S \in \mathbb{N}$. Then, for each $n \in \{1, 2, \ldots, N_S\}$ two continuous variables were defined ($Q^h_{\text{internal},n,S}$ and $F^h_{\text{internal},n,S}$), as well as a boolean variable $W^h_{n,S}$. These variables were constrained with a Disjunctive Programming formulation (15) and the continuous variables were combined to yield output flowrates and contaminant loads (16).

$$\begin{bmatrix} W^h_{n,S} \\ Q^h_{\text{internal},n,S} = \frac{2^{-n} Q^h_{\text{in},S}}{1-2^{-N_S}} \\ F^h_{\text{internal},n,S} = \frac{2^{-n} F^h_{\text{in},S}}{1-2^{-N_S}} \end{bmatrix} \vee \begin{bmatrix} \neg W^h_{n,S} \\ Q^h_{\text{internal},n,S} = 0 \\ F^h_{\text{internal},n,S} = 0 \end{bmatrix} \tag{15}$$

$$\begin{aligned} Q^h_{\text{out},o_1,S} &= \textstyle\sum_{n=1}^{N_S} Q^h_{\text{internal},n,S} \\ F^h_{\text{out},o_1,S} &= \textstyle\sum_{n=1}^{N_S} F^h_{\text{internal},n,S} \end{aligned} \tag{16}$$

Finally, for topology optimization, a simplified splitter, with integer and constant split fraction, is introduced. This unit defines to which of its two possible outputs all input flowrate and contaminant loads are directed. Constraint (17) makes use of a boolean variable (V_S) to model a given selector, S.

$$\begin{bmatrix} V_S \\ Q^h_{\text{out},o_1,S} = Q^h_{\text{in},S} \\ Q^h_{\text{out},o_2,S} = 0 \\ F^h_{\text{out},o_1,S} = F^h_{\text{in},S} \\ F^h_{\text{out},o_2,S} = 0 \end{bmatrix} \vee \begin{bmatrix} \neg V_S \\ Q^h_{\text{out},o_1,S} = 0 \\ Q^h_{\text{out},o_2,S} = Q^h_{\text{in},S} \\ F^h_{\text{out},o_1,S} = 0 \\ F^h_{\text{out},o_2,S} = F^h_{\text{in},S} \end{bmatrix} \tag{17}$$

Consumers and associated tanks

The water supplied to a population center must obey the legal requirements for contaminant load concentration, and the supplied volume must be enough to cover the local water demand. In this model, consumers and storage tanks were modeled as a single unit: while offering less flexibility, by constraining contaminant load requirements in the tank input, and demand in its output, no bilinearity is introduced, reducing the complexity of the mathematical program.

For a consumer C, are known the capacity of the associated tank ($V_{\max,C}$) the demand in each instant $h \in \{1, 2, \ldots, H\}$ ($Q^h_{\text{demand},C}$) and the minimum and maximum contaminant concentrations ($c_{\min,C}$ and $c_{\max,C}$, respectively). For each instant h were defined input and output flowrates (respectively, $Q^h_{\text{in},C}$ and $Q^h_{\text{out},C}$), input contaminant load ($F^h_{\text{in},C}$), and volume of water stored (V^h_C). Constraints (18) and (19) model the behavior of the tank, constraint (20) ensures that water demand is met, and constraint (21) ensures that input contaminant concentration lies between the legally allowed limits. Constraint (22) ensures that the daily operation of the tank is cyclical, that is, the difference ΔV_C between the volumes stored in the tank at the start (instant 0) and end (instant H) of the day is equal to 0.

$$0 \leq V^h_C \leq V_{\max,C} \tag{18}$$

$$\begin{aligned} &Q^{h-1}_{\text{in},C} - Q^{h-1}_{\text{out},C} + V^{h-1}_C = V^h_C \\ &\qquad \forall h \in \{2, 3, \ldots, H\} \\ &Q^H_{\text{in},C} - Q^H_{\text{out},C} + V^H_C = V^1_C + \Delta V_C \end{aligned} \tag{19}$$

$$Q^h_{\text{out},C} = Q^h_{\text{demand},C} \tag{20}$$

$$F^h_{\text{in},C} \le c_{\max,C} Q^h_{\text{in},C}$$
$$F^h_{\text{in},C} \ge c_{\min,C} Q^h_{\text{in},C} \quad (21)$$

$$\Delta V_C = 0 \quad (22)$$

Superstructure optimization

To assess the inclusion or exclusion of a given unit U, a binary variable z_U is defined, and constraint (23) forces the inactivation of the unit if this binary variable is zero.

$$z_U = 0 \Rightarrow Q^h_{\text{in},U} = 0 \forall h \in \{1, \ldots, H\} \quad (23)$$

Multi-scenario modeling

Assuming that a given period of time can be split in D sets of N_d days each with similar water supply and demand (for instance, week days and weekends, or days in different months of the year), a multi-scenario approach can be employed, instead of optimizing the system for each time instant [6]. Under this model, d "clones" of the system are optimized simultaneously, with the inclusion of extra constraints to ensure that the optimal design is the same for the entire period: for all pairs (d, d') of sets of days, all selectors S must activate/inactivate the same output (23), all units U which could be included or excluded from the system are included/excluded (24) and the volume of each tank C must be a continuous function of time (25) (constraint (22), defined in the previous section, must be discarded). Finally, the volume stored in the tank must be greater than zero and smaller than its capacity (25).

$$V_{S,d} = V_{S,d'} \quad (24)$$

$$z_{U,d} = z_{U,d'} \quad (25)$$

$$V^H_{C,d} + (N_d - 1)\Delta V_{C,d} + Q^H_{\text{in},C,d} - Q^H_{\text{out},C,d} = V^1_{C,d+1} \quad (26)$$

$$0 \le V^h_{C,d} + (N_d - 1)\Delta V_{C,d} \le V_{\max,C} \quad (27)$$

The contiguity of the downtime periods required a few more constraints, however. As defined in constraint (5), the number of intervals in which some well W had to be inactivated is $A_W = \lceil H D_W \rceil$. If $D = 2$, it can be shown that the contiguity of the downtime periods is assured (even if the downtime is divided between two days) if Constraint (26) applies.

$$Y^{H-A_W+n}_{W,d=1} = Y^{H-A_W+n}_{W,d=2} \forall n \in \{2, \ldots, A_W\} \quad (28)$$

Finally, to avoid global deficit or surplus, $V^1_{C,D+1}$ is defined as $V^1_{C,1}$ in Constraint (26).

Objective function

Finally, the objective function was defined as the total permeate production in the plant (27), as a proxy for capital and operational costs. However, the model accepts different objective functions, with the possibility of minimizing water usage or other system properties.

$$\begin{aligned} \min \quad & \textstyle\sum_{d=1}^{D} \left(N_d \sum_{h=1}^{H} Q^h_{\text{perm},P} \right) \\ \text{s.t.} \quad & \text{constraints (1-13)} \\ & \text{constraints (15-21)} \\ & \text{constraints (23-28)} \end{aligned} \quad (29)$$

Case Study

The model defined in the previous sections was applied to the optimization of the design and operation of the water production system in a small city in São Paulo state, Brazil (Figure 1), where the water demands of a population of *c.* 8000 people was supplied by five wells (Table 1) with high fluoride concentrations. The modeling scenarios investigated contemplated fluoride removal by RO (plant P, in the figure) and daily contiguous periods of four hours in which water production from each well was interrupted. The permeate produced by this RO plant, with a low concentration of fluoride, was then mixed with the raw well water, to adjust the fluoride concentration in the distributed water to the legally allowed concentraion interval (between 0.6 and 0.8 mg/L). This stream supplied a water storage tank T, which supplied the city. In the base case, the storage tank had a capacity of 500 m³, and the plant operated with a recovery of 80% and a salt retention of 95%. Afterwards, sensitivity analyses were performed to assess the impact of tank capacity and salt retention.

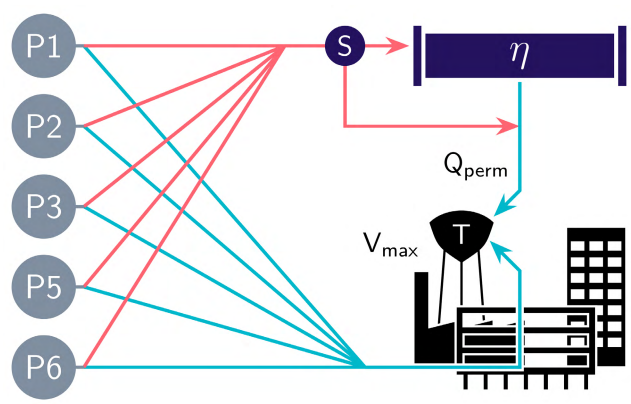

Figure 1. Schematic representation of the water production network studied, showing an RO plant with a permeate production Q_{perm} and a salt retention η, supplying a tank with capacity $V_{\max}$. Two water pipes must be implemented: one (a subset of the structure shown in pink) feeds the water treatment system and the other (in teal) supplies the storage tank T. A splitter S splits the stream that feeds the plant in an excess water stream (which is sent to the storage tank) and a plant feed stream. Mixers, selectors, and the brine discharge are not shown, for clarity.

Table 1: Properties of the wells.

Well	Maximum flowrate (m^3/h)	Fluoride concentration (mg/L)
P1	16.0	6.15
P2	28.0	6.03
P3	23.0	6.62
P5	30.0	1.09
P6	16.0	1.04

The water demand exhibited hourly and daily variations, with the weekend demand of 1, 608 m^3/day being 52.7% higher than demand in weekdays (equal to 1, 056 m^3/day). Hence, a multi-scenario approach was adopted, with two sets of days (N_1=5 week days and N_2=2 weekend days).

The goal of the optimization was to determine the optimal plant capacity and network topology, obeying legal requirements for minimum and maximum fluoride concentrations and producing drinking water, with water quality within the legal drinking water standards, to supply the demand of the city. The water storage tank accumulated excess drinking water produced during the week to cover the needs of the (higher demand) weekends. Finally, plant input and output flowrates were set to be constant, requiring the introduction of a new constraint (30), where $Q_{\text{perm},P}$ was a continuous variable also constrained by (10). As the flowrates were the same in every period, the objective-function was simply the sum of of the permeate output in all instants considered, without need for the weights N_d adopted in (29).

$$Q^H_{\text{perm},P,d} = Q_{\text{perm},P} \quad \forall d \in \{1,2\}, \forall h \in \{1,\ldots,H\} \tag{30}$$

In this model, constraint (27) was adapted, replacing the factor $(N_d - 1)$ by the factor N_d. This modification, which allowed one extra day of weekend or week day demand, was required to improve the robustness of the model, permitting the operation of the system in "long weekend" conditions.

In this model, a bi-hourly discretization was adopted (that is, H=12), and the granularity of the logarithmic approximation of the splitter was set to G_S=16=2^4. This granularity leads to small (<5%) differences between the optimal split fractions obtained from the linearized model and the results which would be obtained from the nonlinear program.

The model was implemented in the modeling language CasADi [7], in a Python script, and solved with the Mixed-Integer Linear solver CPLEX 22.1.1.0. The Mixed-Integer gap was set to 0.01%. The script was executed in a personal computer running Debian GNU/Linux 12, with an 8-core/16-thread AMD Ryzen 7 2700 CPU and 48 GB of RAM.

RESULTS AND DISCUSSION

Base case

Initially, the system was implemented adopting ρ=95% and $V_{\max,T}$=500 m^3. It was possible to optimize the resulting mathematical program to optimality in a reasonable time frame, obtaining a minimal plant capacity of $Q_{\text{perm},P}$=17.76 m^3 (Table 2).

Table 2: Properties of the mathematical program and software performance.

MILP properties (after presolve)	
Number of variables	2077
Continuous variables	1741
Binary variables	336
Constraints	1294
Software performance	
Feasibility	Feasible
Mixed-Integer Gap	0.00%
Execution time (s)	107.69
Best objective function value	1065.74

In the optimal solution, the storage tank, as expected, accumulated excess drinking water production during the week days and used this water to cover the increased weekend demand (Figure 2a). In the optimal solution, the production of wells 1, 2 and 3 was directed to the RO plant, while the production of wells 5 and 6 was sent directly to the tank (as illustrated in Figure 3, point "e"). The downtimes of each well obeyed the legal requirements (Figure 2b).

Tank capacity and plant salt retention effects

To assess the impacts on plant capacity requirements of changing tank capacity or the selection of membrane elements (properties that lead to complex financial trade-offs: larger plants require higher volumes of feedwater and lead to higher reject flowrates, membrane elements with higher salt retention rates have higher energy demands, and the CAPEX of a tank grows with its capacity) a sensitivity analysis was performed. 300 pairs $(\eta, V_{\max,T})$ of salt retention and tank capacity values were selected in a centered square lattice with η between 85 and 99.5% and $V_{\max,T}$ between 50 and 1000 m^3. Then, the resulting mathematical program was optimized, as the base case was in the previous section. Of the initial 300 points, 274 were in the feasible region.

By comparing the optimal plant capacity values and the optimal system topologies in each analyzed point, some global trends could be observed (Figure 3). First, the optimal topology of the network depended on the storage tank capacity and membrane design (both its

capacity and its rejection characteristics). In a scenario with higher plant capacity requirements (Figure 3a) all wells were connected to the plant, but in a scenario where a smaller plant was required (Figure 3f) only the wells with higher contaminant concentrations demanded RO treatment.

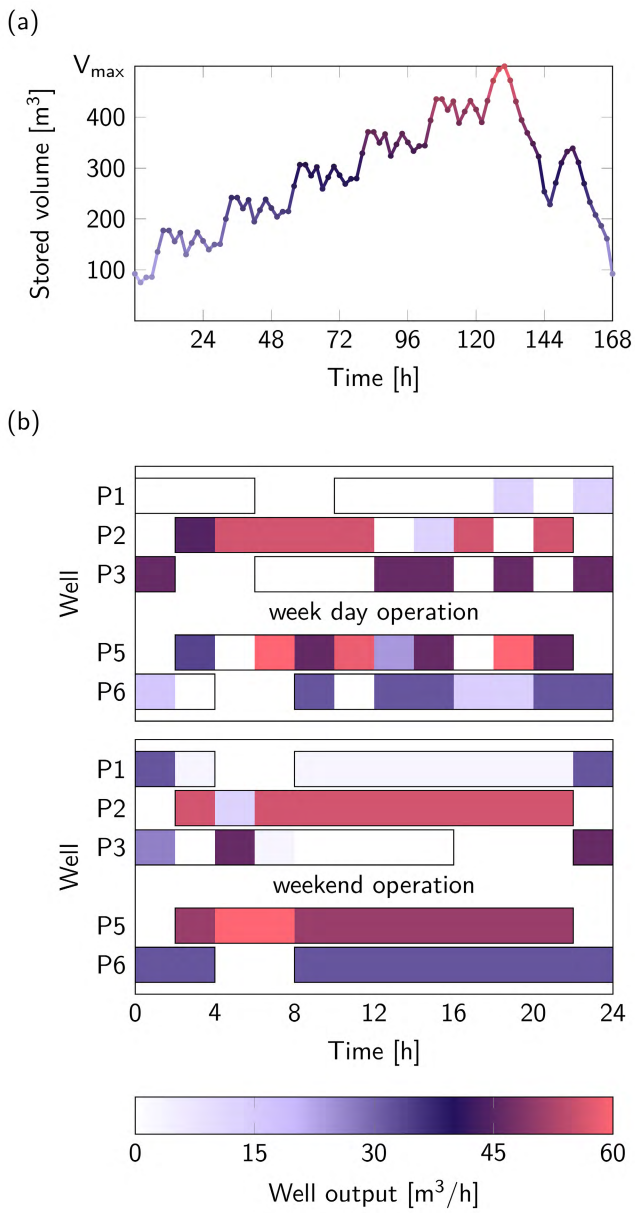

Figure 2. Operation of the tank (a) and wells (b) on week days and weekend days.

Second, while $Q_{\mathrm{perm},P}$ changed smoothly with both η and $V_{\mathrm{max},T}$ for smaller values of η, there were some discontinuous regions for η<90%; besides, as this parameter grows, eventually an unfeasible region was reached, initially for systems with smaller tanks, and then for systems with progressively larger storage capacities. However, no general relations between plant capacity discontinuities and changes in optimal topology were observed.

Finally, the complexity of the mathematical landscape, illustrated in the figure, demonstrated that no global recommendations for the improvement of the system could be made: while increasing salt retention (by selecting different membrane elements) or storage capacity might lead to a significantly smaller plant, there are diminishing returns for these modifications, specially for higher values of η and $V_{\mathrm{max},T}$.

CONCLUSION

A Mixed-Integer approach was developed and implemented for the optimization of water production networks with Reverse Osmosis water treatment units, in which there are legal requirements for daily contiguous downtime periods. This model was then applied to a case study, were large demand variations between weekends and week days required a multi-scenario optimization strategy with the usage of “design days” to model demand in each scenario without needing to optimize the operation of the system over the entire optimization horizon.

The application of the model to the case study demonstrated its reasonable performance (with global optimization achieved in less than two minutes, in the base case), and yielded important results: first, it was shown that the implementation of an RO system is a feasible alternative to lower the high contaminant concentration in the water supply of the city; second, a sensitivity analysis performed by optimizing systems with different membrane technologies and equipped with different storage capacities exhibited pathways for system improvement, defining a landscape of trade-offs between tank CAPEX, energy usage, and water abstraction and waste.

However, further research is warranted: given the random nature of water demand, which is affected by several environmental and social factors, a stochastic optimization approach would increase the robustness of the proposed schedules. This modeling approach would also allow the consideration of contaminant concentration variations, which can be significant.

Finally, while the current mathematical program minimizes plant flowrate, as a proxy for water treatment expenses, a future research avenue lies in minimizing the system costs as a whole, considering CAPEX and OPEX for each equipment (wells, plants, tanks) and minimizing total annualized costs.

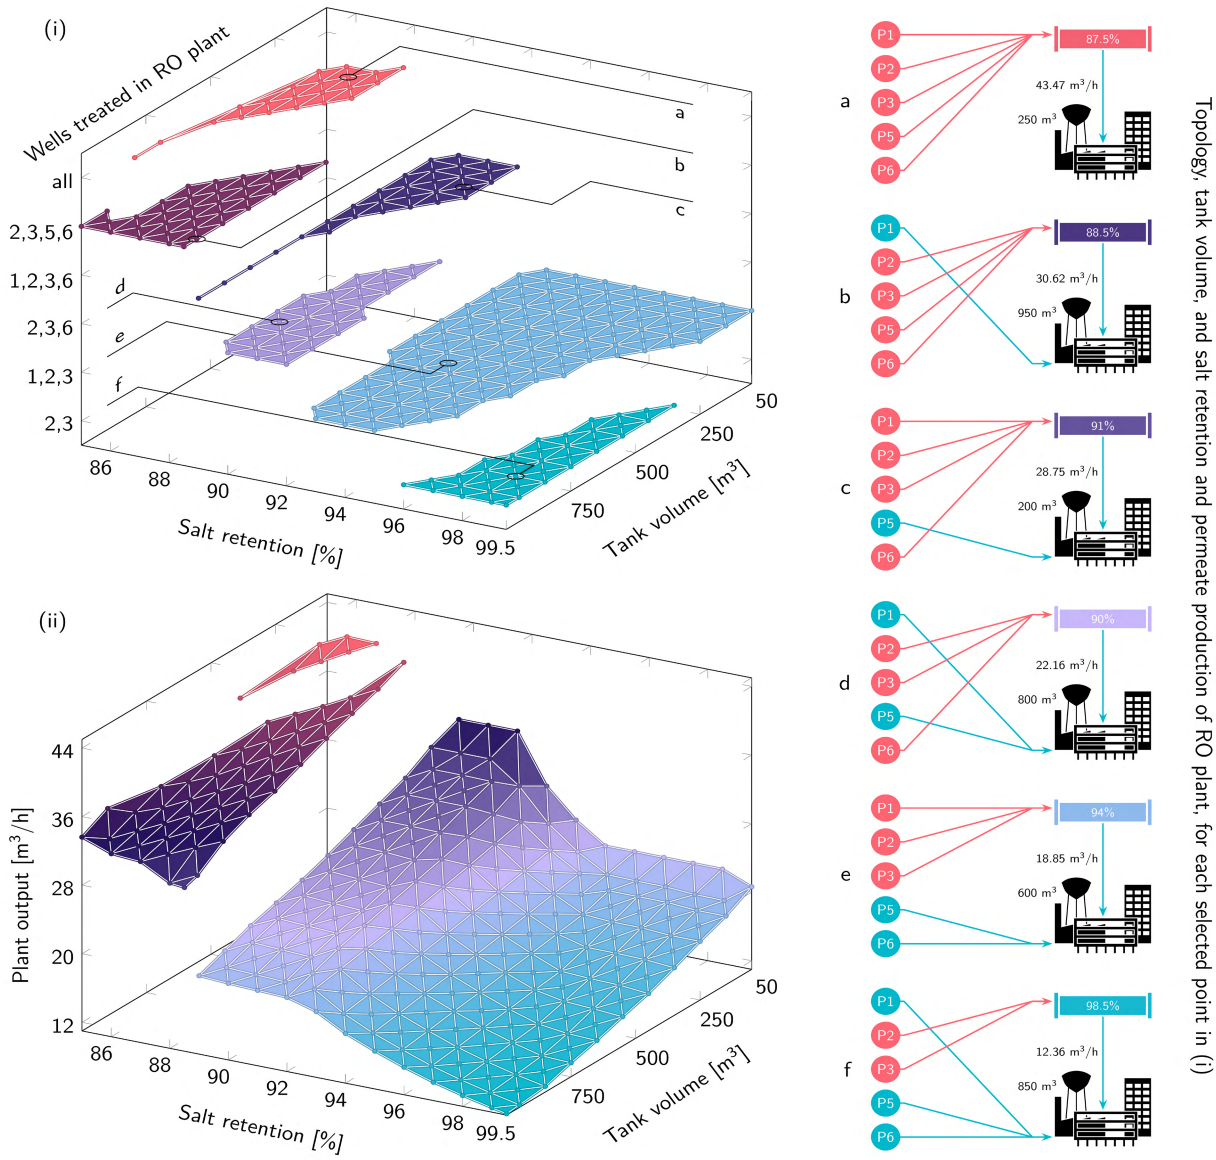

Figure 3: Relation between salt retention, tank volume, optimal network topology (i), and plant capacity requirements, with illustrative examples of the different network topology recommendations (a-f).

ACKNOWLEDGEMENTS

The study was financed by the *Fundação de Amparo à Pesquisa do Estado de São Paulo* (FAPESP), grants 2023/04683-7 (PHCS) and 2020/12275-8 (RPS), and by the *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior* (CAPES) — Finance Code 88887.227909/2025-00 (PHCS). The authors acknowledge the water management company of São Paulo state (SABESP) for supplying well and demand data for the case study.

AUTHOR IDENTIFIERS

Author ORCIDs:
Callil-Soares PH: 0000-0002-9949-3532
Schneider RP: 0000-0001-8436-0270
Carrillo Le Roux GA: 0000-0001-7227-5271

REFERENCES

1. Malaeb L, Ayoub GM. Reverse osmosis technology for water treatment: state of the art review. Desalination 267:1-8 (2011). https://doi.org/10.1016/j.desal.2010.09.001
2. Je?owski J. Review of water network design methods with literature annotations. Ind. Eng.

Chem. Res. 49:4475-4516 (2010). https://doi.org/10.1021/ie901632w
3. Misener R, Thompson JP, Floudas CA. APOGEE: global optimization of standard, generalized, and extended pooling problems via linear and logarithmic partitioning schemes. Computers & Chemical Engineering 35:876-892 (2011). https://doi.org/10.1016/j.compchemeng.2011.01.026
4. Hull DG. Optimal control theory for applications. Springer Science & Business Media (2013).
5. Kolodziej S, Castro PM, Grossmann IE. Global optimization of bilinear programs with a multiparametric disaggregation technique. J Glob Optim 57:1039-1063 (2013). https://doi.org/10.1007/s10898-012-0022-1
6. Gabrielli P, Gazzani M, Martelli E, Mazzotti M. Optimal design of multi-energy systems with seasonal storage. Applied Energy 219:408-424 (2018). https://doi.org/10.1016/j.apenergy.2017.07.142
7. Andersson JAE, Gillis J, Horn G, Rawlings JB, Diehl M. Casadi: a software framework for nonlinear optimization and optimal control. Math. Prog. Comp. 11:1-36 (2018). https://doi.org/10.1007/s12532-018-0139-4

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Towards the Resilient Design of Power-to-Ammonia Systems via Linear Optimization Tools

José M. Pires[ab], João Fortunato[a], Henrique A. Matos[a] and Diogo A. C. Narciso[a]*

[a] Centro de Recursos Naturais e Ambiente, Departamento de Engenharia Química, Instituto Superior Técnico, Universidade de Lisboa, Lisboa, Portugal
[b] Centro de Química Estrutural, Institute of Molecular Sciences, Departamento de Engenharia Química, Instituto Superior Técnico, Universidade de Lisboa, Lisboa, Portugal
* Corresponding Author: diogo.narciso@tecnico.ulisboa.pt.

ABSTRACT

The design of Power-to-Ammonia systems (P2A) is a challenging task. While the technology for all its components, including renewable energy harnessing, electrolysis, Haber-Bosch synthesis, and auxiliary buffers, is mature, assembling such a system to meet the challenges of varying power profiles is not trivial. To ensure resilient, cost-effective designs, careful selection of unit capacities and coordination of all system operations are required. Specifically, this requires modeling system behaviour and enforcing operational constraints to capture system flexibility over a representative time frame. The first steps towards a novel P2A design framework are presented, with a focus on enhanced process operations and exploring new options for process flexibility. A general methodology is proposed, where the full system can be customized, namely by enabling or disabling: (i) multiple renewable energy harnessing sources, (ii) grid operations, and (iii) buffers, including battery and/or an H_2 tank. The model is defined as a linear function of a set of design and operation variables, expressed through a set of matrices and vectors, which are then used to formulate and solve a Linear Programming problem. The design methodology is presented and illustrated via 3 distinct P2A configurations, with key results and analyses of system flexibility and performance. A discussion of future steps towards developing this design framework is also presented.

Keywords: Power-to-Ammonia, Process Design, Optimization, Sustainability, Resilience

INTRODUCTION

Ammonia (NH_3) is a key building block in the modern economy, with a global production of 183 million metric tons in 2020. Its relevance stems from its role as a precursor in fertilizer production, supporting an estimated 50% of food chains [1]. Given its potential as an energy vector, demand is forecast to increase and to capture new markets, such as maritime transportation [1].

The Haber-Bosch (HB) process is the industry's standard for ammonia production, accounting for more than 90% of the total output [2]. It reports roughly 63% energy efficiency and 1.5-1.6 ton CO_2 / ton NH_3 for the Best Available Technology, via the Steam Methane Reforming (SMR) process [3]. While the HB process has long achieved full technological maturity, the industry accounts for an estimated 1-2% of CO_2 emissions, making it a primary target for decarbonization and achieving the UN's Sustainable Development Goals [4].

The primary source of carbon dioxide (CO_2) emissions in this process is related to the operation of the SMR section, where a 3:1 molar syngas mixture of hydrogen (H_2) and nitrogen (N_2) is produced. Natural gas (primarily composed of methane – CH_4) is used in this section as a feedstock to manufacture H_2, and as a fuel to achieve high temperatures in the reforming units [5].

The last section of the process is the HB synthesis, where the syngas from the SMR section is catalytically converted to NH_3 in a reaction loop. Direct CO_2 emissions are much lower and are mainly due to CH_4 purge from the loop, where it remains chemically inert. The synthesis reaction is autothermal, avoiding fuel consumption except during process start-up, but it requires a power supply to achieve the high pressures in the synthesis loop [2].

Water electrolysis is a more sustainable route to hydrogen production, as it avoids the use of fossil fuels.

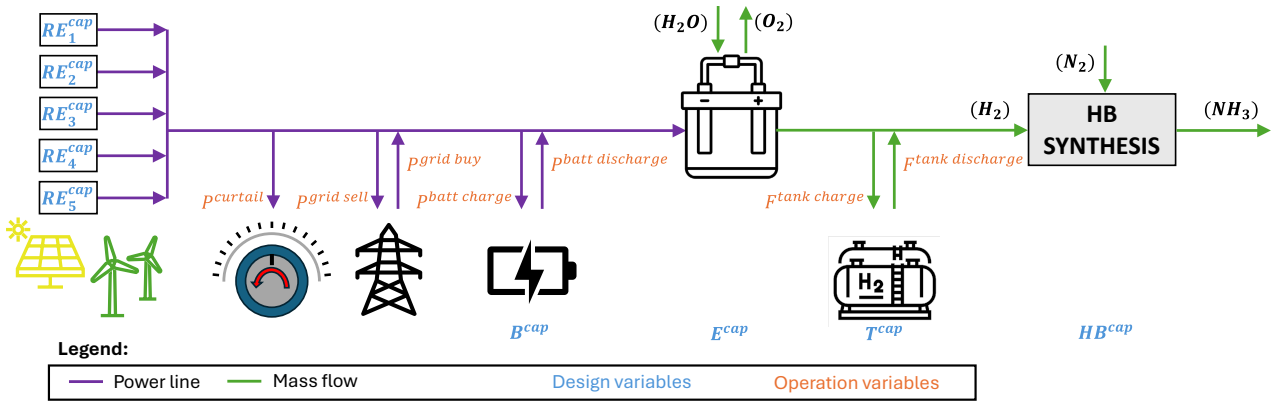

Figure 1. P2A system superstructure (adapted from [4]). Multiple renewable sources (e.g., solar and wind) may be enabled to power the system. Their installed capacities are represented as RE_i^{cap} (i = 1, ..., 5). Power curtailment ($p^{curtail}$) may be activated at periods when a large excess of energy is available from the renewable sources to protect the system from excessive power. Buying ($p^{grid\ buy}$) and selling ($p^{grid\ sell}$) energy to the grid may also be enabled, subject to the local power grid's restrictions. A battery (B^{cap}) may be used as a buffer to store energy ($p^{batt\ charge}$) during peak energy availability and supply ($p^{batt\ discharge}$) the system when the harnessed power is low. These 3 auxiliary services may be enabled or disabled in the design framework, though at least one is generally required to stabilise power in the system. The electrolyzer and HB synthesis capacities are denoted as E^{cap} and HB^{cap}, respectively. A buffer H_2 tank (T^{cap}) may also be enabled to stabilise the H_2 flowrate to the HB synthesis section, via charge ($F^{tank\ charge}$) or discharge ($F^{tank\ discharge}$) operations. A set of constraints may also be enabled to capture the availability (e.g., grid) or flexibility (e.g., HB synthesis) of all units in managing their loads, in accordance with industrial experience. System design consists of taking all requirements above and the complexity of managing this system over time to achieve the most economical NH_3 manufacturing cost.

However, a fully green route can only be achieved when all process equipment is powered by renewable energy sources, thereby minimizing direct and indirect CO_2 emissions. These processes, known as Power-to-Ammonia (P2A), are part of the broader Power-to-X class [6] and are schematically depicted in Figure 1.

Renewable sources, electrolysis, and HB synthesis are the core components of P2A systems. Collectively, they provide power for operation and enable a chemical route for producing ammonia from water (H_2O) and N_2. Multiple auxiliary services, including buffers (battery and H_2 storage tank), power curtailment, and grid transactions, may also be considered. Since renewable energy production is not constant, one or more services are generally required to stabilize energy and mass flows, and ensure that the core units operate within their safety specifications.

Moreover, the design of these systems is challenging because the fluctuating power profiles of renewables preclude the use of conventional methods tailored for full-load operation. While this approach would deliver the most cost-effective sizes for electrolysis and HB synthesis, it may require oversizing buffers and/or renewable-harvesting infrastructure (highlighted in blue in Figure 1), thereby increasing the costs of green NH_3 manufacturing.

Operating process units at variable loads can be advantageous in this context, provided time-varying operations in auxiliary units (highlighted in orange in Figure 1) are considered to achieve resilient, competitive project designs. Several contributions have been proposed in the last decade on this subject [3, 7, 8]. A review of design strategies for P2A systems is presented in [4].

This work presents the first steps towards a novel design methodology for P2A systems, informed by gaps identified in the scientific literature, as follows:

- Full system customization, including enabling/disabling process units and setting parameters to test and solve a wide range of process configurations, which is the main focus of this paper.
- Robust design by factoring uncertainty on renewable power, process costs, and key model parameters.
- Enhanced chemical process operations (e.g., electrolysis and/or HB synthesis startup/shutdown) for increased accuracy of system behaviour.
- Explore the potential of H_2 process loading/unloading for improved process flexibility and project economics.

A strong focus is placed on capturing the system's performance over a representative time interval, while high-

level models are used to capture the behaviour of all the system's units. A design methodology based on Linear Programming (LP) tools is proposed. Its building blocks are presented in the next section, and its application is illustrated in the results section.

METHODOLOGY

Design and operation variables

The P2A superstructure used in this work is presented in Figure 1. It includes all the core and auxiliary units presented above, which may be enabled/disabled individually. For each of these units, a design variable representing their capacities is defined in the design problem (highlighted in blue in Figure 1, e.g., E^{cap} for the nominal electrolyzer capacity).

To capture the impact of short-term fluctuations and seasonal patterns in renewable energy harvesting on the operation of these systems, 1 year of operation, discretised in 1h intervals, is considered [4]. This amounts to a total of 8760h in the defined time-horizon, and as many variables in the design problem for each activated system operation (shown in orange in Figure 1, e.g., $P^{curtail}$ for power curtailment); the focus of the design problem is not on delivering the optimal values for this extended set of operation variables, but rather incorporating them on the design problem to more accurately capture the challenges of time-varying operation. This includes the variability of renewables and constraints on unit operations, to achieve a robust and economic system design.

P2A modelling

Renewable sources

Up to 5 independent renewable sources may be selectively enabled in the design framework. For all sources, a technical specification is selected, and the corresponding power profiles are obtained from historical data at the selected project location(s) consistently with the time structure presented above (e.g., from [9]. These are obtained and specified in the design problem as fractions of the renewable source capacities, such that, at all times in the time horizon, renewable power is given by the product of the corresponding source capacity and these fractions. For all sources, a constant energy price (€/MWh) is assumed.

Power curtailment is a mechanism by which the power generated from a renewable source may be reduced. Formally, this corresponds to adjusting the power from all sources individually; from a modelling perspective, a single variable per time interval is considered, bounded between 0 and the total renewable power.

Electrolyzer

Several technologies are available for water electrolysis [9]. Proton Exchange Membrane (PEM) electrolysis is considered, given the very low O_2 concentrations achieved in the H_2-rich stream, which acts as a catalyst poison in the HB synthesis. A PEM cell includes an anode and a cathode compartment, connected via a proton-exchange membrane and an external electron circuit as depicted in Figure 2.

An oxidation reaction (Equation 1) and a reduction reaction (Equation 2) occur at the electrolyzer's anode and cathode, respectively. This electrochemical reaction is driven by an externally applied voltage, in which H_2O is globally converted to H_2 and O_2, as shown in Equation 3.

$$H_2O \rightarrow 2H^+ + \frac{1}{2}O_2 + 2e^- \quad (1)$$

$$2H^+ + 2e^- \rightarrow H_2 \quad (2)$$

$$H_2O \rightarrow H_2 + \frac{1}{2}O_2, \Delta H^\circ = 241.93\ kJ/mol \quad (3)$$

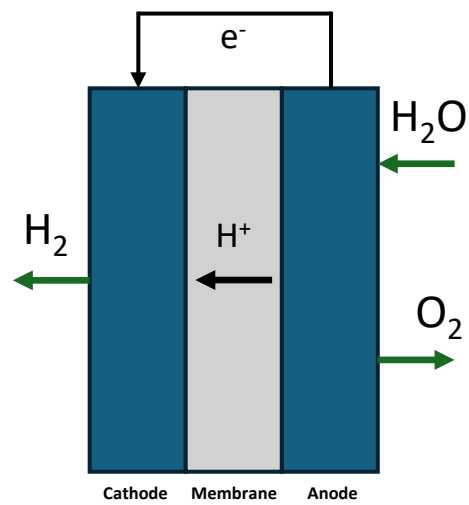

Figure 2. Schematic representation of a PEM cell.

A high-level model is used to capture system behaviour rather than to rigorously model all phenomena. The use of renewable energy to drive this reaction is the key focus, via the definition of a specific energy consumption – expressed in MWh / ton H_2 – where the reaction enthalpy in Equation 3 defines its lower bound. The output H_2 flowrate ($F^{H2\ elect}$) is then calculated as the product of this high-level parameter and the power input to the electrolyzer (P^{elect}).

H_2O is assumed to be available on demand, and O_2 is released from the P2A system (not sold). None is considered in the total cost of green NH_3 production, on the assumption that their prices would represent a small fraction of total manufacturing costs, and particularly the electrolyzer's operating costs.

HB synthesis

The process enables the conversion of a 3:1 molar mixture of H_2 and N_2 (syngas) to NH_3:

$$3H_2 + N_2 \rightleftarrows NH_3, \Delta H^{\circ} = -92.28\ kJ/mol \quad (4)$$

A simplified process layout is schematically depicted in Figure 3. The syngas is first compressed and fed to a converter, which typically includes 2 or 3 iron-based catalytic beds. Several designs for this converter have been proposed, in which the feed gas is conveniently distributed and preheated to sustain autothermal operation [2]. Conversions per pass of around 25% are commonly reported, which requires gas recycle to achieve high overall conversions. The mixture is then cooled to promote the condensation and removal of NH_3 from the process stream. A purge is commonly employed to remove inert components, such as Argon (Ar); in the SMR-based process, it also includes purging CH_4. The unreacted gas is compressed and mixed with the make-up gas, thus completing the synthesis loop. Operating pressures in the range 150-300 bar and temperatures 623-773 K have been reported for this process [10].

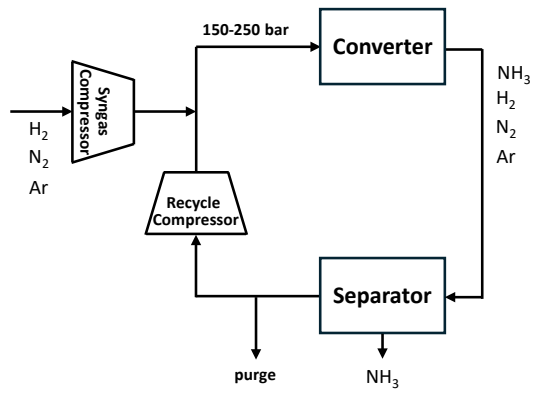

Figure 3. Schematic representation of the HB synthesis.

N_2 for the HB synthesis is supplied via an Air Separation Unit (ASU) or a Pressure Swing Adsorption (PSA) unit, depending on the scale of NH_3 production. It is assumed that the capital costs of the HB synthesis also include the ASU or PSA unit, and that the operational costs and operational constraints of this unit are not significant, in contrast with the operation of the electrolyzer and HB synthesis. Formally, the HB synthesis's compressors also require a power supply; some studies suggest this accounts for a small fraction of the electrolyzer's power demand and is not considered in this study.

The full conversion of the feed syngas is assumed, from which the NH_3 flowrate ($F^{NN3\,HB}$) is calculated directly from the stoichiometry in Equation 4.

Grid

Two 24h price profiles are assumed for: (i) energy bought from the grid, and (ii) energy sold to the grid. These profiles are repeated for the 365 days of the time horizon. The primary modelling consideration concerning grid interaction is setting the upper bounds of $P^{grid\ buy}$ and $P^{grid\ sell}$, which limit the extent of these operations.

Buffers

Battery and H_2 tank, if enabled, are modelled as conventional storage devices according to Equations 5 and 6, respectively:

$$C_{j+1} = C_j + \left(P_j^{batt\ charge} - P_j^{batt\ discharge}\right)\Delta t$$
$$j = 1, \dots, 8760 \quad (5)$$

$$H_{j+1} = H_j + \left(F_j^{tank\ charge} - F_j^{tank\ discharge}\right)\Delta t$$
$$j = 1, \dots, 8760 \quad (6)$$

where C_j and H_j denote the battery's energy charge and H_2 tank's holdup, respectively, at time j, and Δt = 1h. An initial condition for both buffers is also required.

Energy and mass balances

Energy and mass balances are defined at all nodes of the system, based on the associated power and flowrates in the system's main line. This includes the supply of renewable energy and interaction with auxiliary units. Power and flowrates are calculated for all streams over the 8760 h of the time horizon, assuming no power transmission losses occur.

From a modelling perspective, energy utilization in the electrolyzer to produce H_2 is the primary consideration; this is quantified by the electrolyzer's specific energy consumption. This is highlighted in Figure 1, where the modelling focus is on power/flowrate upstream/downstream of the electrolyzer, respectively. In the HB synthesis, H_2 is converted to NH_3 according to the stoichiometry in Equation 4.

Design problem via Linear Programming

The design problem must deliver the capacities of all hardware units (design variables) while minimizing ammonia production cost by fully articulating them with all operation variables. This problem is particularly well-suited to optimization, as it involves tens of thousands of process variables and operational constraints. This choice enables a systematic route to finding the best design via formal optimization principles. Since all model equations and inequations are expressed as linear functions of the design and operation variables, Linear Programming is a particularly suitable choice for the design problem.

Objective function

The objective function for the design problem, assuming all optional units are enabled, is shown below:

$$f = \left(E^{cap}k^E + HB^{cap}k^{HB} + B^{cap}k^B + T^{cap}k^T\right) + \sum_{i=1}^{5} RE_i^{cap}k_i^{RE} + \sum_{j=1}^{8760} P_j^{grid\ buy}k_j^{grid\ buy} - \sum_{j=1}^{8760} P_j^{grid\ buy}k_j^{grid\ buy}$$

(7)

The objective function comprises 3 types of costs. Firstly, the CAPEX and OPEX costs for all units enabled in the P2A design. All are expressed in the objective function as the product of a constant annualised cost factor (e.g., k^E in M€ / per unit of installed capacity – Table 1) with the corresponding design variable (e.g., E^{cap}). OPEX costs include maintenance and, in the case of the electrolyzer, exclude the power supply, which is accounted for through renewable power sources.

The second cost type relates to the total energy harnessed in 1 year of operation. For all enabled sources, this is calculated as the product of the selected source capacity (RE_i^{cap}, $i = 1, \ldots, 5$) with the sum of all load fractions (full year) and the corresponding energy price (k^{RE}).

The last cost component considers grid interaction. If buy from and/or sell to the grid service(s) are enabled, the costs are expressed in the objective for all hours in the time horizon as the product of $P^{grid\ buy}$ or $P^{grid\ sell}$, respectively, and the corresponding price vectors ($k^{grid\ buy}$ and $k^{grid\ sell}$). Note that for consistency, selling energy to the grid represents a monetary income, and the corresponding product in Equation 7 is multiplied by "-1".

Constraints

The design problem also includes several constraints. These relate mainly to the allowed flexibility in the operation of the core and buffer units as presented earlier. Two types of constraints may be imposed on electrolyzer operation, as follows:

- Power must be within a lower and an upper bound, defined as fractions of the nominal capacity.

$$E^{cap}\phi^E_{min} \leq P^{elect}_j \leq E^{cap}\phi^E_{max}, j = 1, \ldots, 8760 \quad (8)$$

- Power change between consecutive time intervals must be within a lower and an upper bound, defined as fractions of its nominal capacity.

$$E^{cap}\mu^E_{min} \leq P^{elect}_{j+1} - P^{elect}_j \leq E^{cap}\mu^E_{max},$$
$$j = 2, \ldots, 8760 \quad (9)$$

where ϕ_{min}^E and ϕ_{max}^E denote the minimum and maximum electrolyzer loads, expressed as a fraction of the electrolyzer capacity to enforce a valid power operating range. Equation 9 limits the extent of ramping power in this unit, by imposing a similar type of constraint via the difference of power between consecutive times, and where μ_{min}^E and μ_{max}^E enforce the corresponding limits.

A similar approach is adopted to HB synthesis operation via parameters ϕ_{min}^{HB}, ϕ_{max}^{HB}, μ_{min}^{HB} and μ_{max}^{HB}:

$$HB^{cap}\phi^{HB}_{min} \leq F^{NH3\ HB}_j \leq HB^{cap}\phi^{HB}_{max},$$
$$j = 1, \ldots, 8760 \quad (10)$$
$$HB^{cap}\mu^{HB}_{min} \leq F^{NH3\ HB}_{j+1} - F^{NH3\ HB}_j \leq HB^{cap}\mu^{HB}_{max},$$
$$j = 2, \ldots, 8760 \quad (11)$$

System design also includes a key equality constraint that enforces that the sum of NH_3 manufactured across all time intervals matches a predefined target (M^{NH3}, in kt NH_3):

$$\sum_{j=1}^{8760} F^{NH3\ HB}_j \frac{\Delta t}{1000} = M^{NH3} \quad (12)$$

Global limits on the total energy bought from and sold to the grid annually may also be enabled to further constrain operations in line with local regulations, energy availability, and limits on green H_2 production. These are expressed as fractions of the total energy consumed in electrolyzer operation.

Additional constraints are also defined to ensure that the energy charge and H_2 mass holdup are within the minimum and maximum limits of the corresponding storage devices' capacities. The maximum discharge and charge rates are specified via the upper bounds on $P^{batt\ charge}$, $P^{batt\ discharge}$, $F^{tank\ charge}$, and $F^{tank\ discharge}$.

Software/workflow

The design framework is implemented in Python. This includes calculating all matrices and vectors that define the design problem's objective function and constraints and using them to solve the corresponding LP problem. The numpy and scipy libraries [11, 12] deliver this functionality. All P2A plant specifications are imported to Python, and all LP matrices are automatically generated. At the end, the optimizer is returned, from which the design variables can be extracted, and some analyses of the system's performance may be performed.

Key unit specifications

A summary of the key parameters required for configuring P2A systems is presented in Table 1.

RESULTS AND DISCUSSION

To illustrate the application of the P2A design framework, 3 examples are configured based on the general superstructure presented in Figure 1 and solved via LP. A small-scale 10 kt NH_3/year production target is considered. The electrolyzer has a specific energy consumption of 50 MWh/tH_2 and is constrained to operate between 0% and 100% of the nominal load. Wind and solar sources are assumed to be available in Sines (Portugal) as the location for energy generation. Renewable harnessing data was collected for both sources from [9]; the average load factors for the solar and wind series are 0.189 and 0.447, respectively. Grid prices for a representative day of operation on an hourly basis are shown in Figure 4, and the remaining cost factors are estimated as listed in Table 2.

Table 1: Key user-defined parameters in the proposed

design framework for P2A systems.

Description	Units
Renewable energy prices (per source).	€/MWh
Renewables load fractions (per source, for all 8760 h in a year).	-
Grid buy and sell prices (24h series).	€/MWh
Electrolyzer specific energy consumption.	MWh/tH_2
Electrolyzer min and max loads[1].	-
Electrolyzer min and max load change[2].	-
HB synthesis min and max loads[1].	-
HB synthesis min and max load change[2]	-
CAPEX and OPEX cost factors (all units).	[3]
NH_3 target production.	kt NH_3

[1] – expressed as a fraction of the unit's capacity (ϕ_{min}^E and ϕ_{max}^E; ϕ_{min}^{HB}, ϕ_{max}^{HB}); [2] – same as [1] and referring to the load increments between consecutive hours (μ_{min}^E and μ_{max}^E; μ_{min}^{HB} and μ_{max}^{HB}); [3] – expressed in M€ per unit of installed capacity (annualised for a 30-year project duration).

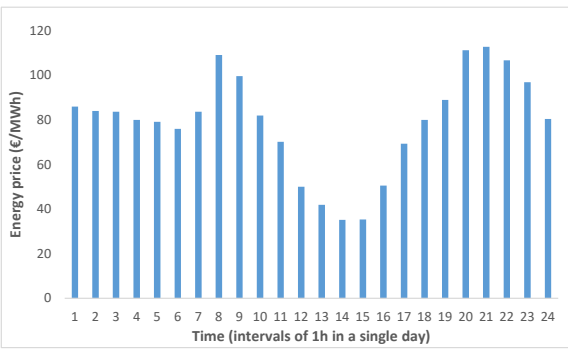

Figure 4. Grid prices for a representative day of operation, on an hourly basis.

Table 2: Infrastructure cost factors and renewable energy prices used in all scenarios.

Equipment	Cost factor	Units
Electrolyzer	$1.33*10^5$	€/MW
HB synthesis	$1.80*10^5$	€/(kt NH_3/y)
H_2 tank	$2.80*10^4$	€/t H_2
Solar source	40	€/MWh
Wind source	45	€/MWh

In the 3 illustrative examples presented in sequence, distinct auxiliary units and constraints are enabled. The key results for all scenarios are shown in Table 3.

Scenario 1: Full load P2A with grid

This design aims to reproduce the full-load operation of the classic HB process via SMR. This is achieved by enforcing the HB synthesis to operate at 100% of its nominal capacity for all hours of the time horizon. This scenario establishes a base case for comparison with the remaining P2A examples. Buying energy from the grid is enabled and subject to a 10% limit on the annual contribution of grid energy for green H_2 manufacture.

The design capacities and KPIs for all scenarios are listed in Table 3. The HB synthesis operates, as expected, at an average 100% load (always at full load), and the minimum installed capacity of 10 kt NH_3/y is achieved to deliver the target production. Since no buffer is active between the electrolyzer and the HB synthesis, the electrolyzer must also operate at 100%, thereby achieving the minimum costs for the core units. The 10% limit on grid energy utilization is active. The grid acts as an indispensable buffer to ensure full load requirements are met, particularly when renewable generation is low. The burden to deliver an optimal design under these conditions falls on the contracted power from renewables, which largely exceeds the actual system needs, as shown in the energy and cost shares in the second column of Table 3. Grid operation and curtail power for this scenario are illustrated in Figure 5.

A manufacturing cost of 1080 €/tNH_3 is achieved. If the 10% limit on grid usage is removed, the price would fall from 1080 to 841 €/tNH3, with 33.7% grid utilization.

Scenario 2: Flexible P2A with grid

In this and the next scenarios, the full-load requirement is relaxed, and instead, the HB synthesis is required to operate at 70%-100%, with load increments between consecutive hours limited to 10% of its nominal capacity. The same grid utilization conditions as defined in the first scenario are also considered.

From Table 3, the 10% grid utilization limit remains active. However, by adding flexibility to the HB synthesis, the design of the electrolyzer and HB synthesis are now slightly oversized, with an average load of 88.6%. These units contribute more significantly to overall costs but enable a significant reduction in the solar and wind contracted capacities from 17.9 MW to 13.6 MW and from 32.3 MW to 21.0 MW, respectively. Power curtailment is significantly reduced, and a much lower manufacturing cost of 894 €/tNH_3 is obtained. By further extending flexibility to 50-100% loads and a maximum of 20% loads' increment, the cost can be further reduced to 861 €/tNH_3.

Scenario 3: Flexible P2A with H_2 tank

An islanded system is considered, precluding the utilization of the grid to assist P2A system operation. An H_2 tank is considered to provide the necessary flexibility and meet the load requirements of HB synthesis operation as presented in scenario 2.

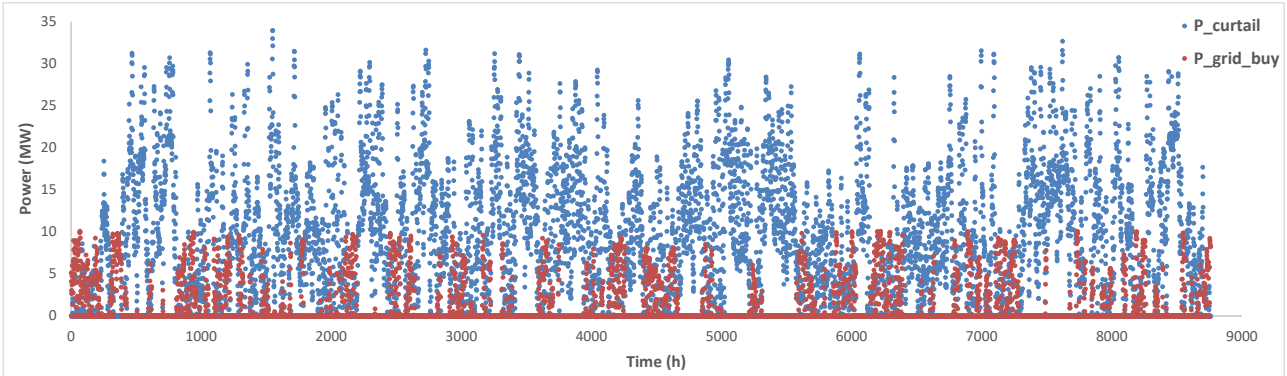

Figure 5. Curtailment and grid power bought from the grid over the full time-horizon in scenario 1.

The H_2 tank decouples the electrolyzer and HB synthesis operations: in the previous scenarios, there are no buffers between the two units, and their operating loads are always synchronized. Now, they are to some extent independent of one another, as confirmed by their average operating loads in Table 3. This change effectively enables full flexibility for the electrolyzer unit (0-100% load). A new energy mix is obtained, promoting the slightly cheaper solar source as listed in Table 2. Note that despite optimal capacities of 18.7 MW and 17.8 MW for solar and wind, respectively, wind power still dominates energy shares due to the higher load fractions in the corresponding renewables' series.

A 14.0 t H_2 tank capacity is delivered from the optimal design. This capacity would be sufficient to continuously supply the HB synthesis for 65 h at the nominal rate, with no H_2 supplied from the electrolyzer. The tank holdup is depicted in Figure 6, where several cycles between full charge and no charge are identified.

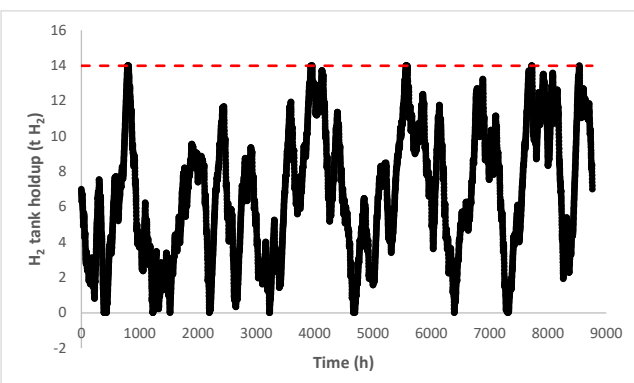

Figure 6. H_2 holdup over the time horizon for scenario 3.

A cost of 861 €/tNH_3 is achieved. While this is a small improvement over scenario 2, a fully green P2A system has been developed.

Design results

The design results and summary of the Key Performance Indicators (KPIs) in the 3 scenarios presented above are listed in Table 3.

Table 3: Summary of design results and KPIs for all scenarios presented earlier.

Design variables & KPIs		Scenario 1	Scenario 2	Scenario 3	Units
Design capacities	RE_{solar}^{cap}	17.9	13.6	18.7	MW
	RE_{wind}^{cap}	32.3	21.0	17.8	MW
	E^{cap}	10.1	11.4	15.0	MW
	HB^{cap}	10.0	11.3	10.6	kt NH_3/y
	T^{cap}	-	-	14.0	t H_2
KPIs	NH_3 cost	1080	894	868	€/tNH_3
	Elect load[1]	100	88.6	67.5	%
	HB load[1]	100	88.6	94.8	%
Energy shares	Solar	33.5	25.5	34.9	%
	Wind	142.4	92.7	78.7	%
	Grid	10	10	0	%
	Curtail	85.9	28.2	13.5	%
Cost shares	Renewables	63.6	51.5	50.5	%
	Electrolyzer	12.5	17.1	23.1	%
	H_2 tank	0	0	4.5	%
	Grid	7.2	8.7	0	%
	HB synthesis	16.7	22.7	21.9	%

[1] – average loads for 1 year of operation.

CONCLUSION AND FUTURE WORK

This work presents the key principles of a novel and fully customizable methodology for P2A design. It consists of a superstructure of core and auxiliary units, the latter of which may be enabled or disabled to test a multitude of P2A configurations. High-level models are defined for all units, and model complexity arises mainly from considering a full year of operation with hourly resolution. The design problem is formulated as an LP, containing both design and operation variables; while calculating the former is the main goal of the design problem, the full consideration of operation variables, and

particularly their operation in the full time-horizon, is critical to achieve design results considering the full complexity of renewables' variability.

The initial results provide useful insights into how unit flexibility, renewable variability, buffers, and objective function definition affect unit capacities, thereby offering a systematic approach to optimal and robust P2A designs. While the initial results are promising, multiple key developments are desirable to enable a more detailed modelling and optimization of these systems. Several phenomena, including: (i) power transmission, (ii) reaction conversion in electrolyzer and HB synthesis, (iii) H_2O and N_2 supply, (iv) powering the HB synthesis' compressors from renewables, (v) neutrality conditions on buffers' operation, and (vi) setting multiple production objectives, will be considered in future work.

Furthermore, this framework enables more advanced utilization modes. For instance, assuming all design variables are set, the framework may be used in P2A systems to facilitate process management and articulation of all units to achieve the most economical manufacturing costs. One of the most critical elements in this framework is setting a priori load fractions for renewable sources based on historical data, as is common practice in the field. Understanding how the design methodology can be improved to deliver robust designs that operate efficiently across a broad range of meteorological conditions remains largely unexplored and may be explored in the future through this design framework.

Given the focus on design modularity presented in this design framework, it readily enables more advanced features than those captured in Figure 1. Work is ongoing to capture the impact of process shutdowns on P2A operation and economics. This is an important consideration because the HB synthesis may require maintenance for 10% of the total operation time [8]. To achieve better process economics and flexibility, buying/selling H_2 is also being developed as an optional feature to make the best use of the investments in renewable power and electrolyzer infrastructure.

ACKNOWLEDGEMENTS

Project co-financed by the PRR – Recovery and Resilience Plan by the European Union, in particular, the Mobilizing Agenda "Moving2Neutrality" (Project no. 32, with reference no. C644927397-00000038). The authors gratefully acknowledge the support of CQE (UID/00100/2025 and UID/PRR/100/2025) and IMS the Associated Laboratory funded by Project LA/P/0056/2020 (DOI 10.54499/LA/P/0056/2020), and of the CERENA through FCT Project UID/04028/2025 (https://doi.org/10.54499/UID/04028/2025).

AUTHOR IDENTIFIERS

Author ORCIDs:
Pires J: 0009-0008-4598-1574
Fortunato J: 0009-0002-6251-6427
Matos H: 0000-0001-8128-7346
Narciso D: 0000-0002-2776-7477

REFERENCES

1. IRENA & AEA. https://www.irena.org/publications/2022/May/Innovation-Outlook-Renewable-Ammonia.
2. Appl M. Ammonia, 2. production processes. Ullmann's Encyclopedia of Industrial Chemistry : (2011). https://doi.org/10.1002/14356007.o02_o11
3. Smith C, Hill AK, Torrente-Murciano L. Current and future role of haber–bosch ammonia in a carbon-free energy landscape. Energy Environ. Sci. 13:331-344 (2020). https://doi.org/10.1039/c9ee02873k
4. Narciso DAC, Pires JM, Fortunato J, Teixeira P, Castro PM, Pinheiro CIC, Matos HA. Design and operation of power-to-ammonia systems: a review. Energy Conversion and Management 327:119494 (2025). https://doi.org/10.1016/j.enconman.2025.119494
5. Fortunato J, Narciso DAC, Matos HA. Towards the Decarbonization of Ammonia Production through the Gradual Incorporation of Green Hydrogen (in Press). *Ind Eng Chem Res*.
6. Palys MJ, Daoutidis P. Power-to-x: a review and perspective. Computers & Chemical Engineering 165:107948 (2022). https://doi.org/10.1016/j.compchemeng.2022.107948
7. Wang H, Daoutidis P, Zhang Q. Harnessing the wind power of the ocean with green offshore ammonia. ACS Sustainable Chem. Eng. 9:14605-14617 (2021). https://doi.org/10.1021/acssuschemeng.1c06030
8. Campion N, Nami H, Swisher PR, Vang Hendriksen P, Münster M. Techno-economic assessment of green ammonia production with different wind and solar potentials. Renewable and Sustainable Energy Reviews 173:113057 (2023). https://doi.org/10.1016/j.rser.2022.113057
9. Renewables.ninja. https://www.renewables.ninja/.
10. Cheema II, Krewer U. Operating envelope of haber–bosch process design for power-to-ammonia. RSC Adv. 8:34926-34936 (2018). https://doi.org/10.1039/c8ra06821f
11. Harris CR et al. Array programming with NumPy. *Nature* 585, 357–362 (2020). https://www.nature.com/articles/s41586-020-2649-2.
12. Virtanen P et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*

17(3), 261-272 (2020). https://www.nature.com/articles/s41592-019-0686-2.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Transfer Learning–Enhanced Deep Probabilistic Surrogates for Scalable Multi-Fidelity Bayesian Optimisation in Process Design

Jaewook Lee, Ethan Errington, and Miao Guo*

Department of Engineering, King's College London, London, WC2R 2LS, United Kingdom
* Corresponding Author: miao.guo@kcl.ac.uk.

ABSTRACT

Self-driving laboratories (SDLs) increasingly use Bayesian optimisation (BO) to navigate expensive design spaces, yet high-fidelity simulations and experiments remain too costly to query at scale. Multi-fidelity Bayesian optimisation (MFBO) alleviates this by combining abundant low-fidelity evaluations with scarce high-fidelity observations. However, Gaussian process (GP) surrogates can become computational bottlenecks as data volume and dimensionality increase, motivating scalable alternatives. Here, we assess transfer learning based deep neural network (DNN) surrogates that pretrain on low-fidelity data and fine-tune on high-fidelity observations. We construct a chemical process benchmark for glacial acetic acid separation and purification with paired low- and high-fidelity flowsheets. The optimisation considers eight decision variables and minimises the minimum selling price (MSP), while enforcing a product purity threshold via a quadratic penalty. To reflect realistic resource constraints, we define seven training scenarios with approximately constant total simulation time, given the large cost gap between fidelities. Across all scenarios, the transfer learning DNN outperforms a GP regression baseline trained only on high-fidelity data. Test-set R^2 increases from 0.41 to 0.71 for the DNN, compared with 0.32 to 0.51 for GP, with gains already visible in the most data-scarce setting. These results suggest that transfer learning–enhanced DNN surrogates can offer accurate, scalable multi-fidelity surrogate models for process optimisation, and could be a practical component in BO-based automated discovery settings where low- and high-fidelity data coexist.

Keywords: Multi-fidelity Bayesian optimisation, Transfer learning, Deep surrogate models, Process optimisation

INTRODUCTION

Self-driving laboratories (SDLs) and simulation-driven design workflows increasingly rely on sequential decision-making algorithms to navigate expensive design spaces [1-3]. Bayesian optimisation (BO) has been widely adopted in this context for optimising expensive black-box objectives in chemical process design, as it balances exploration and exploitation through a probabilistic surrogate model and an acquisition function, achieving competitive solutions with relatively few evaluations [4-7]. However, high-fidelity process simulations often involve rigorous thermodynamic models, equilibrium-stage calculations, recycle loops, and strong nonlinearities, making each evaluation computationally expensive and limiting the number of feasible evaluations within a realistic compute budget. Multi-fidelity Bayesian optimisation (MFBO) alleviates this by integrating abundant, inexpensive low-fidelity evaluations with scarce high-fidelity observations within a single optimisation loop [8-10]. The key technical challenge is to learn a surrogate that can exploit large amounts of biased low-fidelity data while achieving high predictive accuracy in the high-fidelity regime.

Gaussian processes (GPs) have been a dominant surrogate choice in BO and MFBO due to their principled uncertainty quantification and strong small-data performance [11-13]. However, standard GP training scales cubically with the number of training points, which becomes increasingly problematic as low-fidelity datasets grow or

as optimisation campaigns accumulate data over time [14]. Deep neural networks offer a scalable alternative: they can leverage large datasets, learn flexible representations, and train efficiently on modern hardware. In multi-fidelity settings, transfer learning is particularly appealing—a model pretrained on abundant low-fidelity data can learn general input–output structure, while fine-tuning on limited high-fidelity observations corrects systematic bias and calibrates the model toward the target objective [15]. However, a careful evaluation is needed to determine whether this strategy provides reliable improvements over GP baselines under realistic multi-fidelity constraints.

In this work, we evaluate transfer learning–enhanced deep surrogate models as scalable alternatives to GP regression for multi-fidelity optimisation in chemical process design. We construct a benchmark for glacial acetic acid separation and purification with paired low- and high-fidelity BioSTEAM flowsheets, considering eight decision variables and minimising a penalised minimum selling price. To enable fair comparison, we define seven budget-matched scenarios in which total simulation time is kept approximately constant while the ratio of high- to low-fidelity data varies. The contributions are threefold: (i) a reproducible multi-fidelity process benchmark with strong cost asymmetry between fidelity levels; (ii) a transfer learning surrogate strategy that leverages low-fidelity pretraining and high-fidelity fine-tuning; and (iii) a controlled comparison against a GP regression baseline across multiple budget-matched scenarios, demonstrating how predictive accuracy scales with increasing data availability.

METHODS

Overview of the study design

This work evaluates surrogate modelling strategies for multi-fidelity optimisation under realistic computational constraints. We construct a paired low-fidelity (LoFi) and high-fidelity (HiFi) process simulation benchmark for glacial acetic acid separation and purification, and compare two surrogate training paradigms: (i) Single-fidelity baseline (HiFi-only): a GP surrogate trained only on HiFi samples. (ii) Multi-fidelity transfer learning (LoFi pretraining + HiFi fine-tuning): a DNN surrogate pretrained on abundant LoFi data and subsequently fine-tuned using a smaller set of HiFi samples. To ensure a fair comparison aligned with practical usage, we define seven scenarios in which the total simulation time budget is approximately constant, exploiting the large cost disparity between LoFi and HiFi evaluations (approximately 0.5 s vs 167 s per evaluation).

Process benchmark: acetic acid separation and purification

We use a BioSTEAM-based flowsheet representing a hybrid liquid-liquid extraction (LLE) and distillation process to produce glacial acetic acid [16, 17]. A dilute aqueous feed (10, 000 kg h^{-1}, 10 wt% acetic acid) is contacted with ethyl acetate (solvent-to-feed mass ratio 1.5) to yield an acetic-acid-rich extract and a water-rich raffinate. The extract is processed in an extract-distillation column with phase separation to enable solvent recycle, and subsequently purified in an acetic acid purification column to meet the product specification. In parallel, the raffinate and water-rich streams are distilled to recover ethyl acetate for recycle and discharge a wastewater stream. Unless otherwise specified, pressure is set by BioSTEAM defaults (1 atm), and thermodynamic phase equilibria determine unit temperatures except for a heat-exchanger temperature decision variable. The overall process flowsheet is shown in Figure 1.

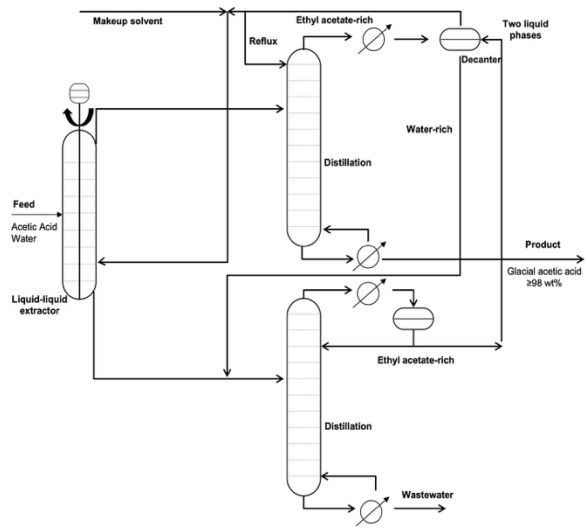

Figure 1. Process flow diagram for glacial acetic acid separation and purification. Solid lines denote process streams; dashed lines denote recycle streams.

Fidelity definitions: LoFi and HiFi models

Two fidelity levels are implemented to reflect computational cost and model fidelity. In the LoFi model, the extractor is represented by a single-stage equilibrium unit (bst.MixerSettler), and key columns are modelled using shortcut distillation (bst.ShortcutColumn, Fenske-Underwood-Gilliland formulation) [18-20]. In the HiFi model, the extractor is modelled as a multi-stage system (bst.MultiStageMixerSettlers, $n_{\text{stages}} \in [10, 50]$) and the two main columns are simulated using equilibrium-stage distillation (bst.MESHDistillation) solved via MESH equations, with tightened convergence tolerances (sys.set_tolerance(rmol=1e-3, mol=1e-3, subsystems=True)) [21]. The raffinate distillation column is implemented as a shortcut column in both fidelities. A key implementation detail is that, because the LoFi extractor

is single-stage, the decision variable n_{stages}does not affect LoFi outputs and is only active in the HiFi model.

Decision variables and optimisation objective

Let $\mathrm{x} \in \mathbb{R}^8$denote the vector of decision variables: the number of extraction stages (n_{stages}, integer 10–50), six recovery ratios for light and heavy keys in three distillation columns (each in $[0, 0.9999]$), and the heat-exchanger temperature $T_{\text{hex}} \in [273, 350]$K. The optimisation objective is to minimise a penalised MSP proxy:

$$f(\mathrm{x}) = \mathrm{MSP}(\mathrm{x}) + \lambda \max\left(0, p_{\min} - p(\mathrm{x})\right)^2,$$

where $p(\mathrm{x})$is the acetic acid mass fraction in the final product stream and $p_{\min} = 0.98$. We use a quadratic penalty with $\lambda = 10$. The MSP is computed from BioSTEAM cost outputs as:

$$\mathrm{MSP}(\mathrm{x}) = \frac{\mathrm{CAPEX}(\mathrm{x}) + \mathrm{OPEX}(\mathrm{x})}{F_{\text{prod}}(\mathrm{x})},$$

with annual production F_{prod}computed using 7, 920 operating hours per year. Because a full techno-economic analysis model is not invoked, this MSP serves as a consistent cost proxy for comparing surrogate performance across fidelities. Numerical safeguards are applied to handle unstable simulations (e.g., failed simulations return a large objective value, and extreme MSP values are clipped).

Data generation and budget-matched scenarios

Table 1. Scenario-wise training set sizes for surrogate models.

	GP	DNN	
Case No.	HiFi train set size	LoFi train set size	HiFi train set size
1	100	5000	50
2	200	10000	100
3	400	20000	200
4	800	40000	400
5	1600	80000	800
6	3200	160000	1600
7	6400	320000	3200

Training datasets are generated by sampling $\mathbf{x}$and running the simulator at the selected fidelity to obtain targets $f(\mathbf{x})$. Initial designs are generated using Latin hypercube sampling (LHS; `scipy.stats.qmc.LatinHypercube`, seed=42) over $[0, 1]^8$, which is then scaled to the variable bounds; the integer variable $n_{\text{stages}}$is obtained by rounding [22]. We define seven scenarios with increasing HiFi data availability. In each scenario, we compare a HiFi-only GP model trained with N_{HiFi}samples to a transfer learning DNN trained with N_{LoFi}LoFi samples for pretraining and N'_{HiFi}HiFi samples for fine-tuning, where $(N_{\text{HiFi}}, N_{\text{LoFi}}, N'_{\text{HiFi}})$ are chosen to keep the total simulation budget approximately constant. The specific dataset sizes are reported in Table 1.

Surrogate models

We compare two surrogate modelling approaches. As a single-fidelity baseline, we train a GP surrogate using only HiFi data, assuming $\hat{f}_{\text{GP}}(x) \sim \mathcal{GP}(m(x), k(x, x'))$, with kernel hyperparameters learned by maximising the marginal likelihood on the HiFi training set [13]. The GP was implemented using scikit-learn's GaussianProcessRegressor with a radial basis function (RBF) kernel. For the proposed approach, we train a DNN surrogate using two-stage transfer learning: model parameters are first obtained by minimising a regression loss on the LoFi dataset,

$$\theta^{(0)} = \arg\min_{\theta} \sum_{i \in \mathcal{D}_{\text{LoFi}}} \mathcal{L}(g(x_i; \theta), f_{\text{LoFi}}(x_i)),$$

and are then fine-tuned on the HiFi dataset starting from $\theta^{(0)}$,

$$\theta^* = \arg\min_{\theta} \sum_{j \in \mathcal{D}_{\text{HiFi}}} \mathcal{L}(g(x_j; \theta), f_{\text{HiFi}}(x_j)).$$

where $g(\cdot; \theta)$is a fully connected network. In our implementation, the DNN was implemented in PyTorch with a 2-layer feature network (64-64 hidden units with ReLU) and is trained using Adam with mean-squared error loss. Input features were normalised to [0, 1] using the physical bounds of each decision variable. Pretraining is performed for 200 epochs (learning rate 10^{-3}), followed by fine-tuning for 100 epochs (learning rate 10^{-4}), updating all layers [15].

All surrogate training and optimisation experiments were performed on the King's College London CREATE HPC cluster (SLURM scheduler), with each job allocated a single NVIDIA A40 GPU (48 GB VRAM, CUDA 12.2), 4 CPU cores, and 16 GB of system memory.

Evaluation protocol and metric

Model performance is evaluated against HiFi ground truth using the coefficient of determination [23]:

$$R^2 = 1 - \frac{\sum_n (y_n - \hat{y}_n)^2}{\sum_n (y_n - \bar{y})^2}.$$

Reporting R^2across the seven budget-matched scenarios quantifies how efficiently each approach leverages limited HiFi data and how predictive accuracy scales under fixed total compute.

RESULTS AND DISCUSSION

Predictive performance across budget-matched scenarios

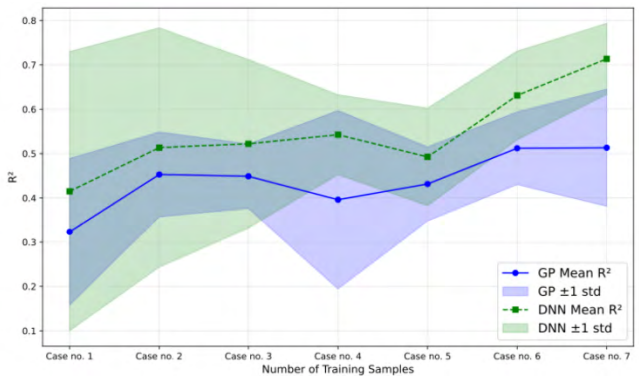

Figure 2. HiFi test-set R^2 across seven budget-matched scenarios comparing a HiFi-only GP baseline and a transfer-learning DNN (LoFi pretrain + HiFi fine-tune). Error bars denote mean ±one standard deviation over 10 random seeds.

Figure 2 summarises the predictive accuracy of the two surrogate models across seven budget-matched scenarios. Each case corresponds to an increasing training budget (Table 1), where the GP baseline is trained using HiFi evaluations only, while the transfer learning DNN is pretrained on a substantially larger LoFi dataset and then fine-tuned using a smaller HiFi set. Generalisation performance is reported as the coefficient of determination R^2 on held-out HiFi data. To quantify robustness, each model was trained with 10 different random seeds, and the figure reports the mean R^2 with a ± 1 standard deviation band.

Across all cases, the transfer learning DNN achieves higher mean R^2 than the GP baseline. In the most data scarce setting (Case 1), the DNN already shows a clear advantage (mean $R^2 \approx 0.41$ vs 0.32), indicating that LoFi pretraining provides useful inductive bias that improves HiFi generalisation even when only a small number of HiFi samples are available for fine tuning. As the training budget increases, the performance gap persists and becomes more pronounced. For example, in the largest budget setting (Case 7), the DNN reaches a mean $R^2 \approx 0.71$, whereas the GP baseline saturates around $R^2 \approx 0.51$. Overall, the GP shows only modest gains with additional data, while the DNN continues to improve, consistent with better scalability in data rich regimes.

Importantly, this improvement is achieved under matched simulation budgets, where the DNN uses fewer HiFi evaluations than the GP in each case (for example, 50 vs 100 in Case 1, and 3, 200 vs 6, 400 in Case 7) by exploiting inexpensive LoFi data. These results suggest that transfer learning enables more effective utilisation of multi fidelity information, yielding higher HiFi predictive accuracy than a HiFi only GP surrogate across a wide range of training budgets.

Learning Dynamics in the Low-Data Regime

To explain why the transfer learning DNN outperforms the GP surrogate in the most data-scarce scenario, we analyse the learning dynamics in Case 1 using the training and held-out HiFi validation loss curves shown in Figure 3. In this setting, the proposed model is pretrained on 5, 000 simulations and then fine-tuned using only 50 HiFi simulations, whereas both the GP baseline and the from-scratch DNN are trained using 100 HiFi simulations only. Validation performance is measured on an independent held-out HiFi set, making it a direct indicator of surrogate quality for the rigorous model.

The left panel (training loss) primarily reflects optimisation behaviour rather than generalisation. As expected, the GP baseline exhibits near-zero training loss

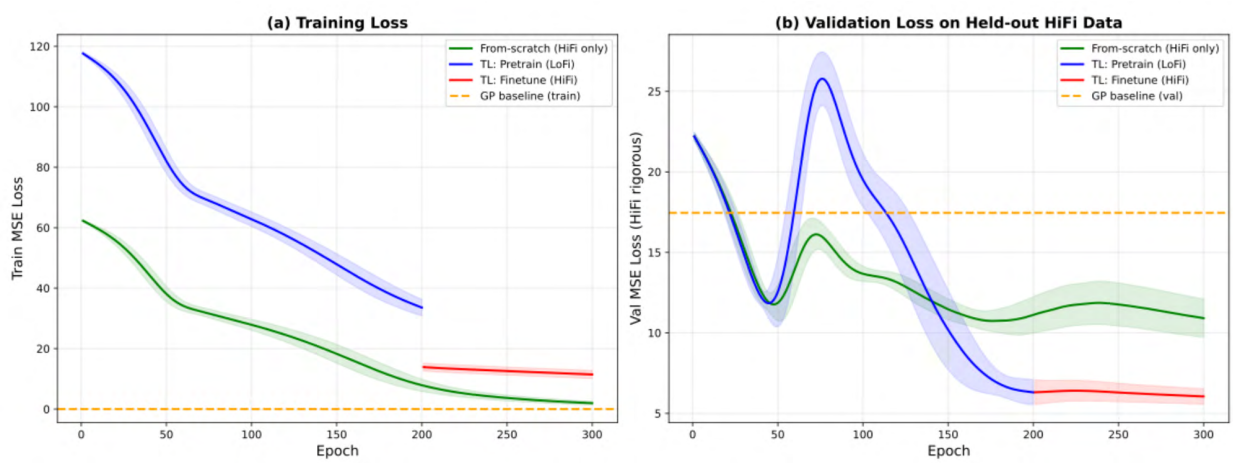

Figure 3. Loss curves in Case 1 showing learning dynamics and positive transfer: (a) training MSE and (b) held-out HiFi validation MSE. The TL-DNN is pretrained on 5,000 LoFi samples and fine-tuned on 50 HiFi samples, whereas the GP and from-scratch DNN are trained on 100 HiFi samples.

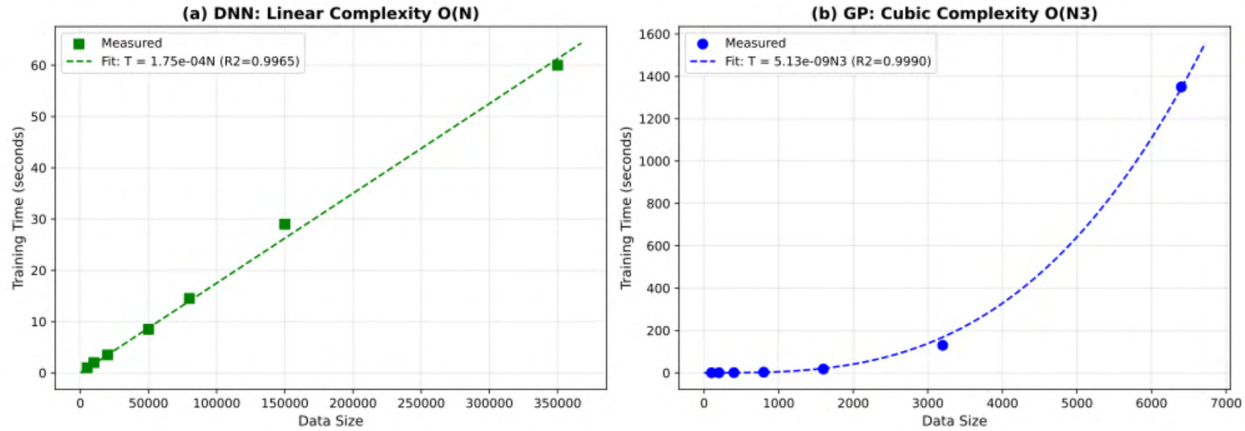

Figure 4. Empirical surrogate training time versus dataset size: (a) DNN training shows approximately linear scaling $O(N)$with a linear fit, and (b) GP training exhibits cubic scaling $O(N^3)$with a cubic fit.

because kernel regression can interpolate the training data, which does not imply strong predictive performance away from observed points. The from-scratch DNN steadily drives the training loss down to very small values over epochs, indicating that the network can fit the limited HiFi training set extremely well. However, this apparent advantage is not mirrored in generalisation.

The right panel of Figure 3 provides the key evidence. The from-scratch DNN reduces validation loss initially but then plateaus and remains substantially above the transfer learning model, consistent with limited generalisation under scarce HiFi supervision. In contrast, LoFi pretraining yields a rapid decrease in HiFi validation loss even before any HiFi fine-tuning is applied, suggesting that the low-fidelity data already encode transferable structure of the process response surface. After switching to HiFi fine-tuning, validation loss decreases further and stabilises at the lowest level among all methods. Importantly, this improvement is achieved despite using half as many HiFi samples as the baselines (50 vs 100), indicating genuine positive transfer rather than an advantage from greater access to high-fidelity supervision.

These dynamics support a mechanistic interpretation of the Case 1 results in Figure 2. LoFi pretraining provides a useful representation of the input–output relationship that acts as an implicit regulariser, constraining the model to physically plausible response patterns and reducing overfitting when only a small HiFi dataset is available. Fine-tuning with sparse HiFi samples then serves to calibrate systematic LoFi bias toward the rigorous objective. Together, this two-stage training strategy improves sample efficiency in the HiFi regime, explaining why transfer learning yields superior predictive accuracy compared to both a HiFi-only GP and a from-scratch DNN in low-data multi-fidelity settings.

Computational scalability with increasing data

Figure 4 compares the empirically measured surrogate training time as the dataset size increases. The DNN surrogate shows approximately **linear** scaling, well described by the fitted relationship $T = 1.75 \times 10^{-4} N$ seconds $(R^2 = 0.9965)$, where N is the number of training samples. This trend indicates that enlarging the LoFi pretraining set remains computationally practical. For example, the fit corresponds to roughly 56 s of training time at $N = 320,000$, which matches the largest LoFi pretraining set used in our scenarios.

In contrast, the GP baseline follows a much steeper **cubic** scaling with dataset size, consistent with the computational cost of exact GP training. The measured times are accurately captured by $T = 5.13 \times 10^{-9} N^3$ seconds $(R^2 = 0.9990)$. Under this scaling, training time increases rapidly even within the sample sizes considered here, reaching approximately 168s at $N = 3,200$ and 1, 345 s at $N = 6,400$.

Together with the accuracy trends in Figure 2, these timing results highlight a key practical advantage of the transfer-learning DNN: it can absorb large LoFi datasets at modest training cost while avoiding the rapidly growing computational burden that limits GP surrogates as datasets expand over iterative optimisation campaigns.

Scope, limitations, and implications for end-to-end MFBO

The results above demonstrate that transfer learning with a DNN surrogate can deliver consistently higher HiFi predictive accuracy than a HiFi-only GP baseline under approximately matched simulation budgets, while also offering substantially better computational scalability as training datasets grow. However, the conclusions of this study should be interpreted within the scope of

the present evaluation, which focuses on surrogate model quality rather than full closed-loop optimisation performance.

First, our comparisons are based on offline predictive metrics (primarily R^2on held-out HiFi data). While improved predictive accuracy is a necessary condition for improved optimisation performance, it is not sufficient by itself. In end-to-end MFBO, overall efficiency depends not only on the surrogate's mean predictions but also on how uncertainty is estimated and utilised by the acquisition function to select the next evaluation. Therefore, a direct next step is to integrate the proposed surrogate into an MFBO loop and quantify optimisation outcomes such as best-so-far objective value, sample efficiency, and HiFi evaluation savings under realistic stopping criteria.

Second, the multi-fidelity pairing in this benchmark is intentionally cost-asymmetric and includes an important modelling detail: the LoFi extractor is represented as a single-stage equilibrium unit, meaning the extraction stage decision variable does not influence LoFi outputs and becomes active only in the HiFi model. This setting reflects a realistic scenario in which low-fidelity approximations omit certain degrees of freedom or physics. Nevertheless, it implies that transfer learning primarily leverages LoFi data to learn shared response structure across the remaining variables, while HiFi fine-tuning is responsible for learning the effect of extraction staging. This reinforces the interpretation of transfer learning as an approach that provides a useful inductive bias and regularisation rather than a complete substitute for HiFi supervision.

Third, the objective function used here is a consistent process-economic proxy with penalty-based feasibility enforcement and numerical safeguards for unstable simulations. Although this formulation is appropriate for controlled benchmarking and surrogate comparison, it does not represent a full techno-economic analysis with annualised CAPEX and explicit financial assumptions. Consequently, the reported values should be interpreted comparatively rather than as absolute economic predictions.

It is worth noting that eight decision variables are well within the capacity of traditional flowsheet optimisation tools such as sequential quadratic programming or evolutionary algorithms. However, these methods typically operate at a single fidelity level, most commonly using shortcut or simplified models, and require numerous function evaluations to converge. When high-fidelity simulations are employed, involving rigorous thermodynamic models and equilibrium-stage calculations, the cost per evaluation becomes a dominant bottleneck. The focus of this work is therefore not on handling dimensionality, but on improving the trade-off between predictive accuracy and computational cost by leveraging abundant low-fidelity data to reduce the number of expensive high-fidelity evaluations required.

Overall, these findings support the use of transfer learning–based deep surrogates as MFBO-ready components for process optimisation workflows where low-fidelity data are abundant and high-fidelity evaluations are costly. Future work should evaluate end-to-end MFBO performance with fidelity selection policies, extend the benchmark to additional flowsheets and fidelity definitions, and further examine uncertainty calibration for acquisition-driven decision making.

CONCLUSION

This study evaluated transfer learning–enhanced deep surrogate models as scalable alternatives to GP surrogates for multi-fidelity optimisation in a BioSTEAM-based acetic acid separation and purification benchmark. Under budget-matched scenarios that reflect the large cost disparity between shortcut (LoFi) and rigorous (HiFi) process simulations, the proposed approach pretrained a DNN on abundant LoFi data and fine-tuned it using a much smaller HiFi dataset.

Across seven scenarios, the transfer learning DNN consistently achieved higher predictive accuracy on held-out HiFi data than a HiFi-only GP baseline, with the advantage already evident in the most data-scarce setting. Analysis of learning dynamics in the low-data regime provided evidence of positive transfer: LoFi pretraining yielded a representation that improved HiFi generalisation and reduced sensitivity to limited HiFi supervision, while subsequent fine-tuning corrected residual bias toward the rigorous objective. In addition to accuracy gains, empirical timing results showed that DNN training scaled approximately linearly with dataset size, whereas GP training exhibited steep cubic scaling, indicating a practical computational advantage as datasets grow.

Taken together, these findings suggest that transfer learning–based deep surrogates can improve both sample efficiency and scalability in multi-fidelity settings where low-fidelity evaluations are plentiful and high-fidelity evaluations are expensive. While this work focused on surrogate modelling performance, integrating the proposed surrogate into end-to-end MFBO loops with acquisition and fidelity selection policies is a natural next step to quantify downstream optimisation gains in realistic closed-loop workflows.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the financial support from the King's College London Net Zero Centre PhD Scholarship scheme. Computational experiments were carried out using the King's College London CREATE HPC cluster [24].

REFERENCES

1. Häse F, Roch LM, Aspuru-Guzik A. Next-generation experimentation with self-driving laboratories. Trends in Chemistry 1:282-291 (2019). https://doi.org/10.1016/j.trechm.2019.02.007
2. MacLeod BP, Parlane FGL, Morrissey TD, Häse F, Roch LM, Dettelbach KE, Moreira R, Yunker LPE, Rooney MB, Deeth JR, Lai V, Ng GJ, Situ H, Zhang RH, Elliott MS, Haley TH, Dvorak DJ, Aspuru-Guzik A, Hein JE, Berlinguette CP. Self-driving laboratory for accelerated discovery of thin-film materials. Sci. Adv. 6: (2020). https://doi.org/10.1126/sciadv.aaz8867
3. Nikolaev P, Hooper D, Webber F, Rao R, Decker K, Krein M, Poleski J, Barto R, Maruyama B. Autonomy in materials research: a case study in carbon nanotube growth. npj Comput Mater 2: (2016). https://doi.org/10.1038/npjcompumats.2016.31
4. E. Brochu, V.M. Cora, N. De Freitas, arXiv preprint arXiv:1012.2599, (2010).
5. Shahriari B, Swersky K, Wang Z, Adams RP, de Freitas N. Taking the human out of the loop: a review of bayesian optimization. Proc. IEEE 104:148-175 (2016). https://doi.org/10.1109/jproc.2015.2494218
6. J. Snoek, H. Larochelle, R.P. Adams, Advances in neural information processing systems, 25 (2012).
7. Shin Y, Lee M, Kim D, Lee J, Lee JW. A hybrid evolutionary and bayesian framework for automated process synthesis of energy-efficient multicomponent azeotropic separation networks. Separation and Purification Technology 384:136280 (2026). https://doi.org/10.1016/j.seppur.2025.136280
8. K. Kandasamy, G. Dasarathy, J. Schneider, B. Póczos, in: International conference on machine learning, PMLR, 2017, pp. 1799-1808.
9. Kennedy M. Predicting the output from a complex computer code when fast approximations are available. Biometrika 87:1-13 (2000). https://doi.org/10.1093/biomet/87.1.1
10. J. Wu, M. Poloczek, A.G. Wilson, P. Frazier, Advances in neural information processing systems, 30 (2017).
11. R.B. Gramacy, H.K. Lee, Statistics and Computing, 22 (2012) 713-722.
12. J. Snoek, K. Swersky, R. Zemel, R. Adams, in: International conference on machine learning, PMLR, 2014, pp. 1674-1682.
13. Rasmussen CE, Williams CKI. Gaussian processes for machine learning. The MIT Press (2005). https://doi.org/10.7551/mitpress/3206.001.0001
14. Patwary MMA, Byna S, Satish NR, Sundaram N, Luki? Z, Roytershteyn V, Anderson MJ, Yao Y, Prabhat , Dubey P. BD-CATS. Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis :1-12 (2015). https://doi.org/10.1145/2807591.2807616
15. Pan SJ, Yang Q. A survey on transfer learning. IEEE Trans. Knowl. Data Eng. 22:1345-1359 (2010). https://doi.org/10.1109/tkde.2009.191
16. Cortes-Peña Y, Kumar D, Singh V, Guest JS. Biosteam: a fast and flexible platform for the design, simulation, and techno-economic analysis of biorefineries under uncertainty. ACS Sustainable Chem. Eng. 8:3302-3310 (2020). https://doi.org/10.1021/acssuschemeng.9b07040
17. J.D. Seader, E.J. Henley, D.K. Roper, Separation process principles, wiley New York, 1998.
18. M. Fenske, Industrial & Engineering Chemistry, 24 (1932) 482-485.
19. E. Gilliland, Industrial & Engineering Chemistry, 32 (1940) 1101-1106.
20. Underwood AJV. Fractional distillation of multi-component mixtures. Chem. Eng. Prog. 44:603-314 (1948).
21. L.T. Biegler, I.E. Grossmann, A.W. Westerberg, (1997).
22. Mckay MD, Beckman RJ, Conover WJ. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. Technometrics 42:55-61 (2000). https://doi.org/10.1080/00401706.2000.10485979
23. Kuhn M, Johnson K. Applied predictive modeling. Springer New York (2013). https://doi.org/10.1007/978-1-4614-6849-3
24. in, King's College London. (2022). King's Computational Research, Engineering and Technology Environment (CREATE). Retrieved March 2, 2022, from https://doi.org/10.18742/rnvf-m076

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Optimal Biogas Utilization Planning in a Pig Farm Under Sustainability Indicators

Jaime David Ponce-Rocha[ab], Martín Picón-Núñez[a], César Ramírez-Márquez[c], José María Ponce-Ortega[c], Ricardo Morales-Rodriguez[a*]**

[a] Universidad de Guanajuato, Department of Chemical Engineering, Guanajuato, Guanajuato, Mexico.
[b] CIATEQ A.C.-Centro de Tecnologia Avanzada, Department of Engineering and Plant Construction, Queretaro, Queretaro, Mexico.
[c] Universidad Michoacana de San Nicolás de Hidalgo, Department of Chemical Engineering, Morelia, Michoacan, Mexico.
* Corresponding Author: ricardo.morales@ugto.mx*, jaime.ponce@ciateq.mx**

ABSTRACT

This work proposes a two-stage optimization framework for the optimal utilization of biogas from pig manure, integrating process-level design with short-term operational planning under dynamic electricity tariff schemes in Mexico. In the first stage, a multi-objective optimization based on 3E (Exergy, Environment, and Energy) analysis was performed. The results demonstrate that increasing the biogas split fraction for upgrading significantly reduces the environmental and exergy indices, enhancing thermodynamic and environmental performance without compromising the energy index. High upgrading flows (split > 0.7) emerged as the most favorable compromise across the evaluated metrics. In the second stage, support vector regression (SVR) surrogate models were developed to approximate nonlinear relationships between the operational split and process outputs. These surrogates were embedded in a Mixed-Integer Linear Programming (MILP) formulation to optimize weekly scheduling under the Mexican Net Billing scheme, incorporating an economic objective (+1E index). The SVR-MILP framework successfully captured market-driven decisions, prioritizing biomethane upgrading during base and intermediate tariff periods and switching to electricity generation during peak hours to maximize economic profit. Biomethane upgrading accounts for approximately 90-100% of the total positive economic benefits, while electricity generation contributes the remaining 10%, depending on the exposure to peak electricity tariffs. Thus, the proposed methodology offers a robust, computationally efficient decision-support tool for flexible biogas systems, bridging the gap between steady-state design and dynamic market responsiveness based on a 4E approach.

Keywords: Biogas upgrading , Power generation, Simulation-optimization , 4E analysis, Planning & scheduling.

INTRODUCTION

Population growth drives a substantial increase in energy demand, directing development plans toward energy security, environmental impact management, and economic growth. Accordingly, international treaties like the 2030 Agenda for Sustainable Development and the Kyoto Protocol emphasize the transition to renewable energy and the reduction of greenhouse gas emissions [1].

Livestock farming represents a key opportunity for waste valorization, particularly manure, which is rich in organic matter and nutrients. However, improper management significantly contributes to the global carbon footprint, water eutrophication, soil acidification, and air pollution [2].

Among livestock residues, pig manure is particularly suitable for on-site energy systems due to its high methane yields, stable biogas production profiles, and organic fertilizer [3]. While currently regarded as a carbon-neutral energy source, its valorization routes vary. Conventional pathways, such as electricity generation using internal combustion engines or combined heat and power (CHP) systems, are mature but often exhibit limited efficiencies. Conversely, upgrading biogas for transportation fuels, grid injection, as well as hybrid configurations, represents a more effective approach to harnessing its potential [4-5].

Although various biogas upgrading techniques exist (e.g., adsorption, membrane separation), absorption remains the most prevalent in the chemical industry due to its high efficiency and low operating pressures. However, absorbent selection is a critical factor directly influencing process performance and operating costs [6]. In addition, process selection depends on methane concentration and plant requirements. For instance, when upgrading is not viable, integrating power generation systems, requiring only the removal of corrosive contaminants such as hydrogen sulfide and siloxanes, is an alternative [7]. However, these valorization processes involve complex trade-offs between technical, economic, and environmental factors. Therefore, a comprehensive evaluation using sustainability metrics, such as 4E analysis (Energy, Exergy, Economy, and Environment) within a dynamic system-level framework, is essential.

Thus, the objective of this work is to design a biogas valorization system using agro-industrial residues, employing a multi-objective optimization framework based on 4E analysis. Additionally, this study integrates optimal planning and scheduling strategies for feedstock utilization to maximize economic performance under dynamic electricity pricing schemes.

CASE STUDY

The biogas feed stream from anaerobic digestion has a mass composition of methane (0.599), carbon dioxide (0.399), siloxanes (0.001), and others (<0.0001), with a volumetric flow rate of 150 Nm^3/h [8]. The proposed utilization scheme integrates two alternative pathways, as illustrated in Figure 1.

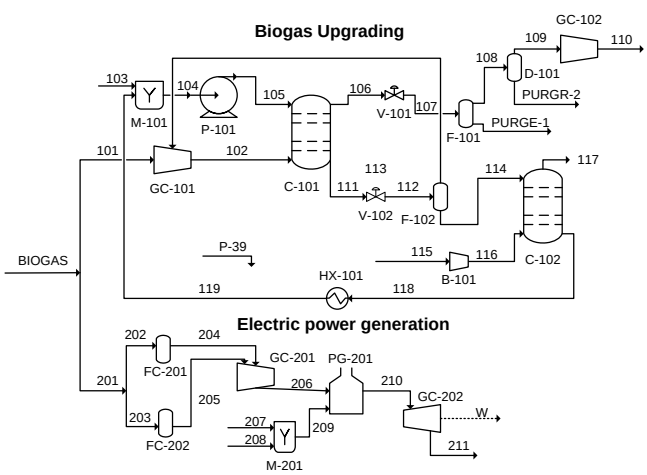

Figure 1. Process Flow Diagram for plant biogas utilization.

The first (upper) pathway corresponds to a biogas upgrading process based on high-pressure water scrubbing (HPWS). This process relies on physical absorption and desorption principles to produce biomethane with high purity (P ≥ 97 %) and recovery (R ≥ 95 %). The biogas is compressed using a gas compressor (GC-101), while water is pressurized by a pump (P-101); both streams are fed into an absorption packed column (C-101) at 5 barg and 5 °C. The top stream consists of methane-rich gas, which is first flashed (F-101), subsequently conditioned in a dehydration column (D-101), and further pressurized in a second compressor (GC-102) at 250 barg for final utilization and transport. The bottom stream of C-101, containing water and absorbed gas impurities, is directed to a flash vessel (F-102) at 2 barg to recover residual methane, which is recycled back to the absorption column. The liquid phase from the flash vessel is sent to a desorption packed column (C-102) at atmospheric pressure, where air supplied by a blower (B-101) strips the absorbed gases. The desorbed gases are released to the atmosphere through the top stream, while the regenerated water is recirculated to the absorption system (M-101 to P-101). The second pathway involves power generation through the combustion of cleaned biogas in a gas turbine to produce electricity for on-site consumption and grid injection. In this configuration, the biogas is first purified in activated carbon columns (FC-201/202) and subsequently pressurized by a low-pressure compressor (GC-201). The fuel is then fed into the power generation unit, which comprises an air supply system (M-201), a combustion chamber (PG-201), and a gas turbine (GC-202). The generated electricity is supplied to the national grid, while the flue gases are discharged into the atmosphere.

Both processes were modeled and simulated in Aspen Plus™, employing the NRTL thermodynamic model and the equation of state Redlich-Kwong to describe liquid-vapor equilibria and incorporating Henry's law to accurately predict gas solubility in the absorption and desorption units.

METHODOLOGY

Two optimization stages were considered in the proposed methodology, which are described as follows:

First stage: Process optimization

The first stage of this methodology consists of optimizing the process operating conditions, and the multi-objective optimization problem is formulated as follows (Equation 1):

$$\underset{x}{Min}\ \underline{Z} = \left[\, Z_i(\underline{x}) \right] i = 1, \dots, j.$$
$$s.t.\, h(\underline{x}) = 0; g(\underline{x}) \leq 0 \qquad (1)$$

where $\underline{Z}$ represents a vector containing j objective functions, and $(\underline{x})$ denotes the vector of decision variables. Furthermore, $\boldsymbol{h}(\underline{x})$ and $\boldsymbol{g}(\underline{x})$ represent the vectors for

equality and inequality constraints associated with the decision variables.

The selected objective functions were energetic, exergetic, and environmental indicators, which are used to identify optimal scenarios for on-site biogas utilization. The decision variables included water feed flow rate (F_{H2O}), air feed flow rate (F_{AIR}), and combustion air feed flow rate (F_{AIRC}). Additionally, inequality constraints were defined for the purity (P) and recovery (R) of biomethane and process water ($P_{CH4} \geq 0.97, P_{H2O} \geq 0.99, R_{CH4} \geq 0.95, R_{H2O} \geq 0.99$), with all values expressed as mass fractions.

Energy

In this work, the energy intensity (EI) index was used, defined as the ratio of electrical energy consumption to the product mass flow rate, expressed in kWh/kg (Equation 2). This indicator enables the assessment of the energy penalty associated with biogas upgrading and power generation.

$$EI\left[\frac{\text{kWh}}{kg}\right] = \frac{Energy\ consumption}{Product\ mass\ flow} \quad (2)$$

Exergy

Exergy represents a metric of the quantity and quality of energy. The exergy losses (σT_0) were quantified as an indicator of system irreversibility; lower values indicate improved thermodynamic performance and a more efficient plant.

To determine the exergy loss, the difference between the physical exergy of the inlet streams [$\Delta \dot{E}x_{In,j}$] and outlet streams [$\Delta \dot{E}x_{Out,j}$] was considered, in addition to the exergy destroyed by heat generation [$\dot{E}x_{q,j}$] and work [$\dot{W}_j$], where j represents each process equipment within the plant (Equation 3).

$$\sigma T_0[\text{kW}] = \Delta \dot{E}x_{In,j} - \Delta \dot{E}x_{Out,j} + \dot{E}x_{q,j} - \dot{W}_j \quad (3)$$

Environment

The objective function associated with the environment integrates Life Cycle Assessment (LCA), allowing the calculation of environmental impacts associated with the construction and operation of the biogas utilization plant. LCA is a methodology established in the ISO 14040 standard and is typically applied in four stages:

1. **Goal and scope definition:** Minimization of the environmental impact for biomethane and energy production, where the Eco-indicator 99 (EI-99) methodology was integrated as the appropriate tool for the LCA.
2. **Inventory analysis:** This step considers the inputs and outputs of matter and energy, construction materials, and the consumption of auxiliary services associated with the process, where all loads must be expressed per reference unit (Life Cycle Inventory, LCI).
3. **Environmental impact assessment:** In this step, process data are translated into environmental impact information; Equation 4 represents the Eco-indicator 99 calculation. Where β represents the total amount of chemical substance b released per unit of reference flow via direct emissions, α represents the damage caused in each category k per unit of chemical substance b released into the environment, ω is the weighting factor for damage in category d, and δ is the normalization factor for damage categories.
4. **Interpretation:** The results are analyzed, and conclusions and recommendations for the system are generated.

$$EI\text{-}99\left[\frac{\text{mPt}}{year}\right] = \sum_b \sum_d \sum_k \alpha_{b,k} \beta_b \delta_d \omega_d \quad (4)$$

The Eco-indicator 99 (EI-99) method was selected due to its suitability for comparing environmental trade-offs in energy and process systems at an early design stage [10].

To perform the optimization, a computational interface COM was implemented using Aspen Plus™ and Matlab™, allowing automated simulation-optimization using a multi-objective genetic algorithm based on NSGA-II (gamultiobj). The optimization hyperparameters included 100 generations, 100 individuals, and a crossover fraction of 0.8. Eleven optimization scenarios were evaluated, varying the split fraction (s) of biogas flow rate destined to each process from 0.00 to 1.00 with increments of 0.1.

To identify the best solution within the optimal set, a global criterion decision-making method was employed. This approach identifies the point on the Pareto front that minimizes the sum of the normalized deviations from the ideal point (Equation 5) [11].

$$\min\ [\underline{Y_j}]$$

$$Y_j = \sum_{i=1}^{n} \left\| \frac{\Phi_{i,j} - \Phi_i^{min}}{\Phi_i^{max} - \Phi_i^{min}} \right\| \quad (5)$$

where Y_j represents the normalized L_1 distance of solution j to the ideal point, $\Phi_{i,j}$ notes the value of the i-th objective function for the j-th optimal solution, Φ_i^{max} and Φ_i^{min} correspond to the maximum and minimum values achieved by each objective function within the Pareto set, and n represents the number of objective functions (3).

Second stage: Optimal biogas planning utilization

The second stage focused on planning and scheduling optimization under economic performance metrics. Using the dataset generated in the first stage, a surrogate model was implemented using support vector regression (SVR), enabling the approximation of the process behavior over the evaluated operating range.

Particularly, SVR was selected due to its robustness under limited training data, its ability to capture nonlinear trends, and its smooth extrapolation behavior within the training domain, making it suitable for embedding within a mixed-integer linear programming (MILP) framework solved using Gurobi™ Optimizer, both implemented in Python™.

The MILP model was formulated to determine the optimal utilization of the biogas flow rate between upgrading and power generation pathways. The objective function aimed to maximize economic profit under time-of-use electricity pricing of Mexico's Net Billing scheme (independent generation and consumption rates). Electricity prices were defined as follows: electricity consumption (base: $ 0.066, intermediate: $ 0.109, and peak: $ 0.130 USD/kWh), electricity generation (base: $ 0.029, intermediate: $ 0.053, and peak: $ 0.101 USD/kWh) [9]. Additionally, a biomethane sales price of $ 1.73 USD/Nm^3 was considered.

Economy

The economic performance of the system was assessed through an economic index defined as the net economic profit [USD], accounting for revenues from biomethane upgrading and electricity export, as well as costs associated with electricity consumption. Equation (6) presents the mathematical formulation of the economic objective function as a function of electricity prices and operational decisions:

$$\underset{x}{Max}\ Z = \sum_{t\in\tau}(P_{CH4}bm_t + p_t^{gen}E_t^{exp} - p_t^{cons}E_t^{imp})\Delta t \quad (6)$$

where bm_t denotes the biomethane production rate, E_t^{exp} and E_t^{imp} represent the exported and imported from the grid, at time interval (Δt), respectively, and P_{CH4}, p_t^{gen}, and p_t^{cons} correspond to the biomethane selling price, electricity generation, and electricity consumption tariffs, respectively.

Table 1 summarizes the applicable electricity tariffs by day of the week and time of operation, which were explicitly integrated into the optimization framework.

RESULTS

Figure 2 shows the three-dimensional trade-off between energy intensity (EI), Eco-indicator 99 (EI-99), and exergy losses (σT_0) for the evaluated scenarios, while Table 2 summarizes the corresponding optimal values for each scenario. Distinct marker shapes were used to differentiate the results of each optimization scenario. The dashed line represents the system trajectory.

As the operational split increases from 0.0 to 1.0, a pronounced and non-linear reduction in EI-99 and σT_0 is observed. A higher biogas flow directed to upgrading results in a reduction in overall environmental impacts and a more thermodynamically efficient system without a significant penalty in the energy intensity index, in contrast to the electricity generation process. While the transition from low to intermediate splits (0.0-0.7) shows substantial reductions in EI-99 and σT_0, the marginal benefits diminish for higher split values (above 0.7). Consequently, the region corresponding to high operational splits can be identified as the most attractive compromise between the indicators reported for the system studied.

Table 1: Electricity rates by day and hour.

	Day-tariff category		
	Base	**Intermediate**	**Peak**
Consumption	$ 0.066 USD/kWh	$ 0.109 USD/kWh	$ 0.130 USD/kWh
Generation	$ 0.029 USD/kWh	$ 0.053 USD/kWh	$ 0.101 USD/kWh
	Day-tariff periods		
	Base	**Intermediate**	**Peak**
Monday to Friday	00:00-6:00	6:00-19:30, 23:30-00:00	19:30-22:30
Saturday	00:00-7:00	7:00-00:00	
Sunday and Holidays	00:00-19:00, 23:00-00:00	19:00-23:00	

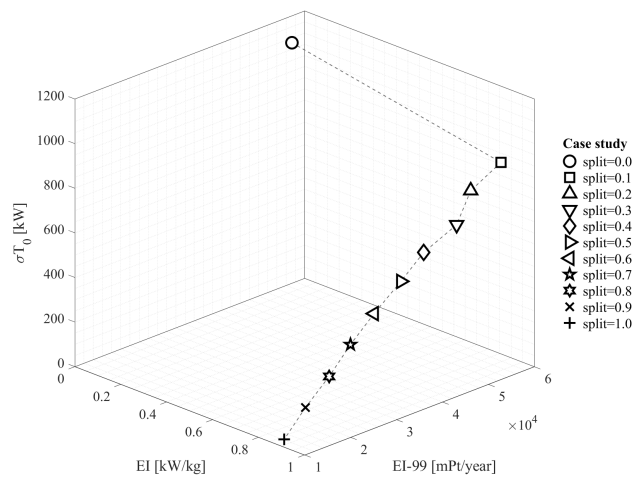

Figure 2. First stage optimization results.

σT_0 values exhibit a decrease of more than 98% when shifting from full electricity generation to upgrading, while EI-99 shows a similarly strong downward trend, decreasing by approximately 81%. In addition, the energy intensity index remains relatively stable for values above 0.3 within a narrow range between 0.89 and 0.94 kWh/kg.

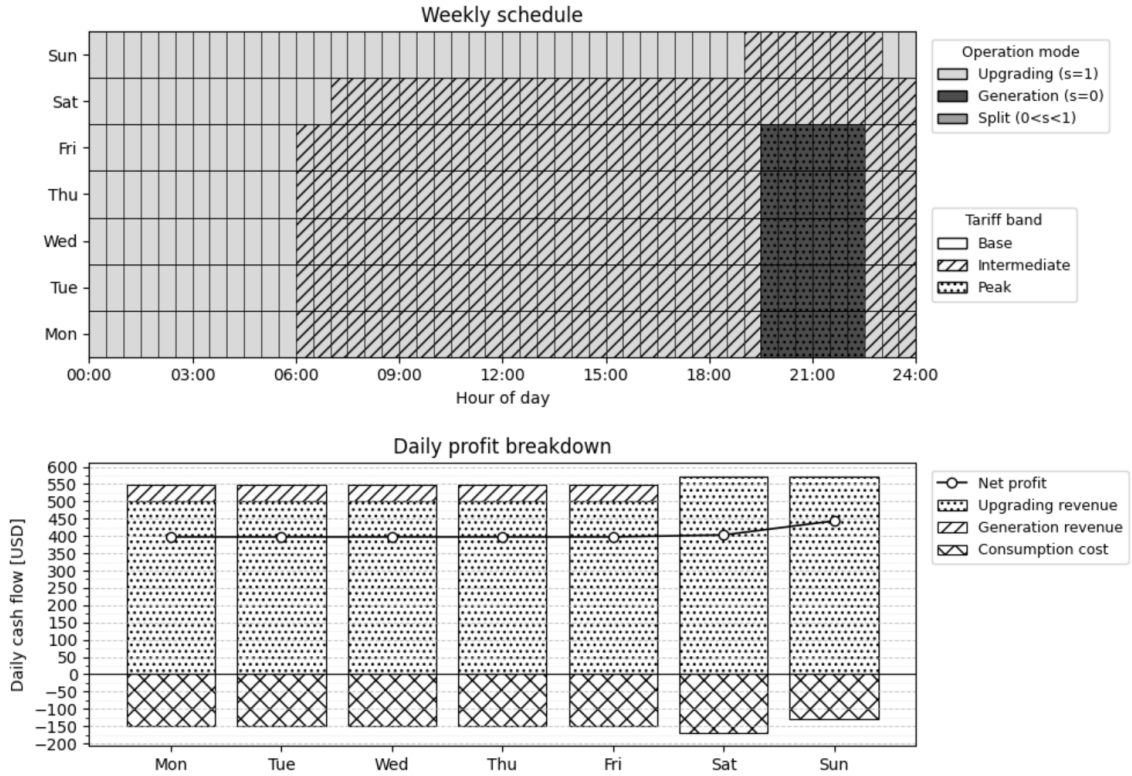

Figure 3. Weekly scheduling and daily profitability under time-of-use electricity rates.

Table 2: Optimal objective function values.

Split	$EI\left[\frac{\mathrm{kWh}}{kg}\right]$	$EI\text{-}99\left[\frac{\mathrm{mP}}{year}\right]$	$\sigma T_0[\mathrm{kW}]$
0.00	0.05	57, 269.00	1, 078.20
0.10	1.00	52, 783.00	972.98
0.20	0.96	48, 188.00	867.11
0.30	0.99	43, 620.00	761.73
0.40	0.94	38, 948.00	655.13
0.50	0.92	34, 990.00	548.76
0.60	0.91	29, 650.00	442.51
0.70	0.90	25, 010.00	336.24
0.80	0.90	20, 370.00	230.11
0.90	0.89	15, 710.00	123.67
1.00	0.89	11, 077.00	17.60

In the second stage of the proposed framework, the SVR allowed capturing nonlinear relationships between the operational decision variables and key process outputs, specifically, electricity generation and electricity consumption. This approach significantly reduces computational effort compared to direct process simulation and optimization. The surrogate models exhibited high predictive performance for operational variables, objective function values, and operating conditions, with coefficients of determination (R^2) above 0.99 over the evaluated operating range (0≤Split≤ 1). Thus, the SVR-based surrogate demonstrated stable behavior across the entire operational range, producing smooth and physically consistent values that enabled its direct embedding into the MILP formulation.

Figure 3 illustrates the optimal weekly operational schedule and the corresponding daily economic metrics obtained from the SVR-assisted MILP optimization under time-of-use electricity tariffs in Mexico. The upper panel represents the selected operational mode for each 30-minute interval throughout the week, revealing a consistent operational pattern during low-price periods (base and intermediate tariff), the system operates in upgrading mode (s=1), prioritizing biomethane production. During peak tariff hours, the system consistently switches to power generation mode (s=0). In this time-based context, electricity generation represents revenues that exceed the benefits of upgrading and methane production. This confirms that the optimization framework correctly captures the economic trade-off imposed by the time-varying electricity process and improves operational decisions in the short-term market.

The lower panel of Figure 3 represents the daily cash flow breakdown. Across all days, upgrading revenue represents the dominant contribution to total profit; therefore, biomethane upgrading remains the primary

economic driver of the system. Electricity generation provides an additional revenue income stream, mainly concentrated on days with extended peak tariff exposure; even then, the electricity consumption cost is higher than base and intermediate tariffs. Biomethane upgrading represents the dominant source of profit, accounting for 91 % of the total positive cash flow from Monday to Friday and 100 % during the weekend.

CONCLUSIONS

This work proposed and validated a two-stage optimization framework for the sustainable valorization of biogas under 4E metrics, integrating process-level optimization with operational planning and scheduling under dynamic electricity pricing schemes.

In the first stage, a multi-objective optimization based on 3E analysis was implemented to evaluate alternative biogas utilization scenarios, revealing a strong trade-off between electricity generation and upgrading pathways. Increasing the split fraction toward upgrading led to a pronounced reduction in EI-99 and σT_0 indexes, representing a more thermodynamically efficient and environmentally friendly process. Consequently, a high biogas flow rate to upgrading (split upper to 0.7) was identified as the most attractive compromise region.

In the second stage, a surrogate-based optimization strategy using SVR and a MILP framework demonstrated that the optimal operation presents a clear operational pattern: biomethane upgrading is consistently favored during base and intermediate tariff periods (s=1), where stable biomethane revenues outweigh electricity-rated benefits, whereas power generation becomes economically attractive during peak tariff intervals (s=0), when electricity generation revenues compensate for the lower overall efficiency of this pathway. The resulting weekly schedule confirms that the proposed framework effectively aligns short-term operational decisions with market-driven price signals.

From an economic perspective, biomethane upgrading emerged as the dominant contributor to overall profitability throughout the week, while electricity generation served as a complementary revenue stream during high-price periods. Based on the optimal weekly schedule obtained from the SVR-MILP framework, biomethane upgrading accounts for approximately 90-100 % of the total positive economic benefits, while electricity generation contributes the remaining 10 %, depending on the exposure to peak electricity tariffs. This highlights the importance of hybrid utilization strategies, where flexibility in biogas allocation enables enhanced economic resilience under variable energy markets.

ACKNOWLEDGEMENTS

Jaime D. Ponce-Rocha acknowledges the Secretariat of Science, Humanities, Technology, and Innovation (SECIHTI) for the scholarship support received to develop this project (Grant No. 4029745).

AUTHOR IDENTIFIERS

Author ORCIDs:
Ponce-Rocha Jaime D.: 0000-0001-7822-941X
Picón-Núñez Martín: 0000-0002-0793-192X
Ramírez-Márquez César: 0000-0002-2321-0436
Ponce-Ortega José María: 0000-0002-3375-0284
Morales-Rodriguez Ricardo: 0000-0001-5909-2155

REFERENCES

1. Kabeyi MJB, Olanrewaju OA. Biogas production and applications in the sustainable energy transition. Journal of Energy 2022:1-43 (2022). https://doi.org/10.1155/2022/8750221
2. Pexas G, Mackenzie SG, Wallace M, Kyriazakis I. Environmental impacts of housing conditions and manure management in european pig production systems through a life cycle perspective: a case study in denmark. Journal of Cleaner Production 253:120005 (2020). https://doi.org/10.1016/j.jclepro.2020.120005
3. Grossi G, Goglio P, Vitali A, Williams AG. Livestock and climate change: impact of livestock on climate and mitigation strategies. Animal Frontiers 9:69-76 (2018). https://doi.org/10.1093/af/vfy034
4. Bai D, Jain V, Tripathi M, Ali SA, Shabbir MS, Mohamed MAA, Ramos-Meza CS. Performance of biogas plant analysis and policy implications: evidence from the commercial sources. Energy Policy 169:113173 (2022). https://doi.org/10.1016/j.enpol.2022.113173
5. Golmakani A, Ali Nabavi S, Wadi B, Manovic V. Advances, challenges, and perspectives of biogas cleaning, upgrading, and utilisation. Fuel 317:123085 (2022). https://doi.org/10.1016/j.fuel.2021.123085
6. Maile OI, Tesfagiorgis H, Muzenda E. The potency of monoethanolamine in biogas purification and upgrading. South African Journal of Chemical Engineering 24:122-127 (2017). https://doi.org/10.1016/j.sajce.2017.06.004
7. Schweigkofler M, Niessner R. Removal of siloxanes in biogases. Journal of Hazardous Materials 83:183-196 (2001). https://doi.org/10.1016/s0304-3894(00)00318-6
8. Calbry-Muzyka A, Madi H, Rüsch-Pfund F, Gandiglio M, Biollaz S. Biogas composition from agricultural sources and organic fraction of municipal solid waste. Renewable Energy 181:1000-

1007 (2022).
https://doi.org/10.1016/j.renene.2021.09.100
9. CFE, “Esquema tarifario vigente - Industria.” Accessed: Jan. 27, 2026. [Online]. Available: https://app.cfe.mx/Aplicaciones/CCFE/Tarifas/TarifasCREIndustria/Industria.aspx
10. Sánchez-Ramírez E, Quiroz-Ramírez JJ, Segovia-Hernández JG, Hernández S, Ponce-Ortega JM. Economic and environmental optimization of the biobutanol purification process. Clean Techn Environ Policy 18:395-411 (2015). https://doi.org/10.1007/s10098-015-1024-8
11. Belegundu AD, Chandrupatla TR. Optimization concepts and applications in engineering. Cambridge University Press (2019). https://doi.org/10.1017/9781108347976

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Process design for the recovery of valuable organic compounds from pyrolysis oil aqueous phase

Matteo Gilardi[a]*, Filippo Bisotti[a], Stefan Schmidt[a], Rune Myrstad[a], Camilla Otterlei[a], Trung Trinh[a], Bernd Wittgens[a]

[a] SINTEF Industry, Department of Process Technology, Trondheim, Trøndelag, Norway
* Corresponding Author: matteo.gilardi@sintef.no.

ABSTRACT

Pyrolysis, a key waste-to-X technology, enables converting a wide portfolio of biomass waste into valuable chemicals and fuels. However, raw pyrolysis oils are chemically and physically unstable. A multi-step stabilisation is necessary to reduce acidity and the content of reactive components, mainly carbonyls and carboxylic acids. During stabilisation, which involves deoxygenation and hydrogenation as the main steps, an aqueous phase is generated as a by-product. This stream contains mainly water, but relevant amounts of methanol and ethanol (2-8 wt%) are also present, together with minor concentrations of higher alcohols, C_1-C_4 carboxylic acids, and light esters. The aim of this work is to design and optimise a process to isolate the methanol and ethanol embodied in the aqueous phase and exploit them as intermediates to generate biofuels, biochemicals, and pharmaceutical products. The process consists of a train of four distillation columns to maximise the recovery rate and purity of the two targeted products. COCO-COFE, a CAPE-OPEN process simulator, is used for flowsheet development. The optimal number of theoretical stages and the optimal positioning of the feed to each unit have been optimised, considering a trade-off CAPEX-OPEX as the criterion. Results show that up to 95% and 98% yield can be achieved for methanol and ethanol, respectively. Methanol purity is limited to a maximum of 92 wt% due to the presence of volatile impurities, while ethanol contains residual propanol (6 wt%). The overall heat demand is in the order of 1 MJ/kg feed.

Keywords: waste valorisation, process design and modelling, process optimisation, organics recovery from wastewater, multistep distillation

INTRODUCTION

Fast pyrolysis is considered a high-potential route to produce sustainable advanced biofuels. However, Fast Pyrolysis Bio Oil (FPBO) is not suitable for direct use as a transportation fuel, and further chemical upgrading is required to enhance stability. Different approaches can be adopted to reduce the concentration of acidic compounds, including co-processing in Fluid Catalytic Cracking (FCC), and hydrotreating. The Biomass Technology Group (BTG), a private Dutch company specialised in converting biomass into sustainable biofuels, has developed a proprietary stabilisation process using a Ni-Cu-based catalyst called Picula™ [1]. This process, which is shown in Figure 1, includes the following main steps [2]:

1. Selective hydrogenation at mild temperatures (200-250°C). The obtained product is called Stabilised Pyrolysis Oil (SPO);
2. Deoxygenation at high temperature (350°C). The obtained product is called Stabilised Deoxygenated Pyrolysis Oil (SDPO);
3. Final hydrogenation (refinement step at 425°C) to produce Hydrogenated Pyrolysis Oil (HPO).

This FPBO stabilisation process generates a liquid product that comprises both an organic and an aqueous phase. The presence of a water phase is largely determined by the substantial water content present in the biomass feedstock. Aqueous phase from pyrolysis is a complex mixture of water (> 60 wt%) and various dissolved organic compounds, including alcohols (methanol, ethanol, propanol, and butanol), sugars (levoglucosan and others), carboxylic acids (mainly acetic, propionic,

and butyric acid), esters, furfural, guaiacol, and minor quantities of aldehydes and ketones. The aqueous phase is generally disposed of as waste or potentially utilised as co-feed in biogas plants. However, given its composition, it could be considered a valuable resource for producing various bio-based chemicals and fuels, though the chemicals of interest are extremely diluted in water. Therefore, up-concentration and purification are needed, which makes the process economically challenging. The aim of this work, which is carried out within the framework of the FUEL-UP project [3] funded by the European Commission, is to design and optimise a process to isolate the chemicals embodied in the aqueous phase.

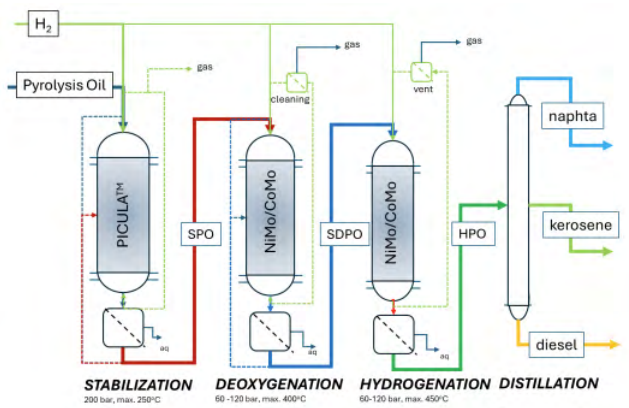

Figure 1. The FPBO upgrading process developed by BTG [2].

The idea is to separate aqueous/organic phase and subsequently treat the two phases individually [1]. The water phase can be further processed to separate and recover valuable chemicals. However, water-alcohols-carboxylic acid mixtures have an azeotropic nature, which makes compounds separation challenging.

One viable option is extractive distillation. This approach can effectively isolate polar components, such as alcohol and acetic acid, enabling their direct use as valuable products or in subsequent processes like biofuel upgrading [4]. Alternatively, azeotropic distillation can be used to remove water from the aqueous phase of pyrolysis oil by utilising the formation of azeotropic mixtures [5]. However, these technologies suffer from the introduction of a third component (the entrainer) during the separation process as a major drawback [4]. This can lead to relevant final product purity issues. Besides, the plant layout complexity considerably increases, and there is a need for heat integration to minimise the energy consumption of product recovery and solvent regeneration. Recently, other advanced separation options have also been investigated, including freeze concentration, reverse osmosis, or bioconversion. However, these technologies are all characterised by a much lower Technology Readiness Level (TRL).

For these reasons, in this work, a train of distillation columns in series is identified as the simplest and most robust separation scheme:

- The first distillation column allows a bulk removal of C_1-C_3 alcohols and eventual other species whose boiling point is lower than water;
- A variable number of downstream columns can be used to refine the separation.

FEEDSTOCK CHARACTERISATION

The composition of the water phase obtained during Stabilised Pyrolysis Oil (SPO) and Stabilised Deoxygenated Pyrolysis Oil (SDPO) production at BTG has been measured using GC-FID (Gas Chromatography - Flame Ionization Detector) technique in SINTEF, and it is shown in Table 1. 14 different species were identified, and their concentration was quantified through calibration curves. The remaining peaks of unidentified species were combined with the peaks of the nearest species observed in the chromatogram (closest retention time).

Table 1: Composition of the aqueous phase samples from SPO (SPO_{Aq}) and SDPO ($SDPO_{Aq}$) measured with GC-FID. Acronyms: water (H2O), methanol (MeOH), ethanol (EtOH), 1-propanol (1-PrOH), 1-butanol (BuOH), acetic acid (AcAc), tetrahydrofuran (THF), Methyl acetate (MeA), ethyl acetate 8EtA), acetone (Acet), formic acid (FoAc), isopropanol (2-PrOH), propionic acid (PrAc), butyric acid (BuAc), 1-pentanol (PeOH).

	GC-FID results		For simulation	
Component	**SPO_{Aq}**	**$SDPO_{Aq}$**	**SPO_{Aq}**	**$SDPO_{Aq}$**
H2O	72.50	89.90	85.60	94.22
MeOH	4.95	0.33	5.84	0.35
Ethanol	2.70	1.85	3.19	1.94
1-Propanol	0.32	0.55	0.38	0.58
Butanol	0.13	0.03	0.15	0.03
Acetic Acid	0.22	0.00	0.26	0.00
THF	0.23	0.19	0.27	0.20
Methyl Acetate	0.25	0.00	0.30	0.00
Ethyl Acetate	0.59	0.49	0.70	0.51
Acetone	0.05	0.08	0.06	0.08
Formic Acid	0.05	0.05	0.06	0.05
Isopropanol	0.20	0.10	0.24	0.10
Propionic Acid	0.23	0.47	0.27	0.49
Butyric Acid	0.37	0.50	0.44	0.52
Pentanol	1.91	0.87	2.26	0.91
SUM	84.70	95.41	100.00	100.00

For simplicity, alcohols heavier than C_5

(concentration below 0.1 wt%) were lumped into pentanol, and carboxylic acids higher than C_4 were lumped into butyric acid. The concentrations of the compounds detected with the GC cover 85% and 95% of the mass balance for SPO_{Aq} and $SDPO_{Aq}$, respectively. For modelling purposes, the composition has been normalised to 100%, as shown in Table 1.

PROCESS DESIGN

Thermodynamic model

In water-alcohols-carboxylic acids systems, strong non-idealities are present due to the azeotropic nature of most of the involved binary interactions. Activity models like Non-Random Two-Liquids (NRTL) and UNIQUAC are the most suitable to describe VLE of this system. However, it has been verified that the existing default NRTL models available in Aspen Plus and COCO-COFE V3.7, a license-free software released by AmsterChem, do not accurately reproduce the equilibrium measurements of most of the binary and ternary mixtures involving the 14 compounds of interest for this application. Therefore, the NRTL model has been re-parametrised to increase predictability. The NIST ThermoData Engine V10.3 was used as the experimental data source. The re-fitted model shows a significantly improved capability to reproduce experimental data, as the Average Relative Deviation (ARD) in the prediction of vapor phase composition at fixed liquid composition and temperature (for isothermal data) or pressure (for isobaric data) drops from 16% and 12.1% of the current COFE and Aspen models, respectively, down to 7%.

Details on the procedure followed for the mode repparametrisation and the outcomes of its performance evaluation are described in a separate contribution.

Process Flow Diagram

The Block Flow Diagram of the designed process for efficient recovery of methanol and ethanol at the highest possible purity is drawn in Figure 2. It consists of a train of four distillation columns in series:

- The first column (Col. 1) is a bulk water removal;
- The second column (Col. 2) separates methanol from ethanol;
- The third column (Col. 3) is a methanol purification step. The target is to minimise the residual concentration of acetone, acetates, and THF. Given the azeotropic nature of most mixtures including the mentioned compounds, there is a limit in the maximum amount of impurities that can be removed by simple distillation. Two scenarios have been investigated: scenario A has maximum methanol purity as the target, while in scenario B methanol recovery is maximised.
- The last column (Col. 4) is an ethanol purification step. The target is to minimise the residual concentration of water and heavier alcohols, with a special focus on propanol.

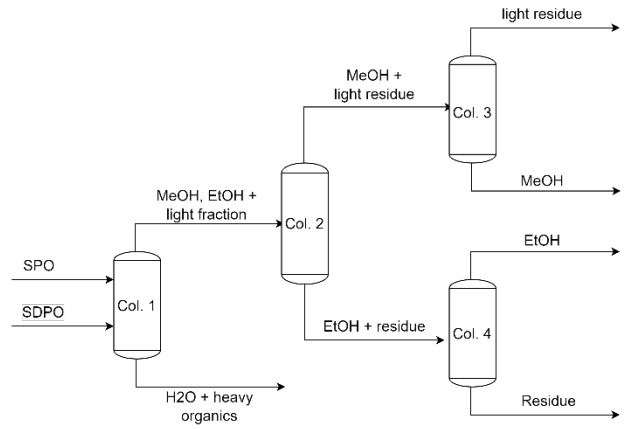

Figure 2. BFD of the distillation process designed for methanol and ethanol recovery from pyrolysis oil stabilisation wastewater.

Three case studies are analysed considering three different feedstocks:

- Only SPO_{Aq} stream (CS1);
- Only $SDPO_{Aq}$ stream (CS2);
- Combined SPO_{Aq} + $SDPO_{Aq}$ streams (CS3).

The same process layout and the same specifications are used to run each case study. A small variation is introduced when $SDPO_{Aq}$ only is treated: given its very low methanol content, column 3 is removed in this configuration.

The specifications set to design each distillation column are gathered in Table 2.

Table 2: List of the specifications chosen to model each distillation column.

Column	Top / bottom	Specification	Value
Column 1	Top	EtOH recovery	99%
	Bottom	H2O recovery	99%
Column 2	Top	MeOH recovery	99.99%
	Bottom	EtOH recovery	99.99%
Column 3, scenario A	Top	THF recovery	95%
	Bottom	EtA recovery	20%
Column 3, scenario B	Top	THF recovery	55%
	Bottom	MeOH recovery	95%
Column 4	Top	EtOH recovery	99%
	Bottom	1-PrOH recovery	95%

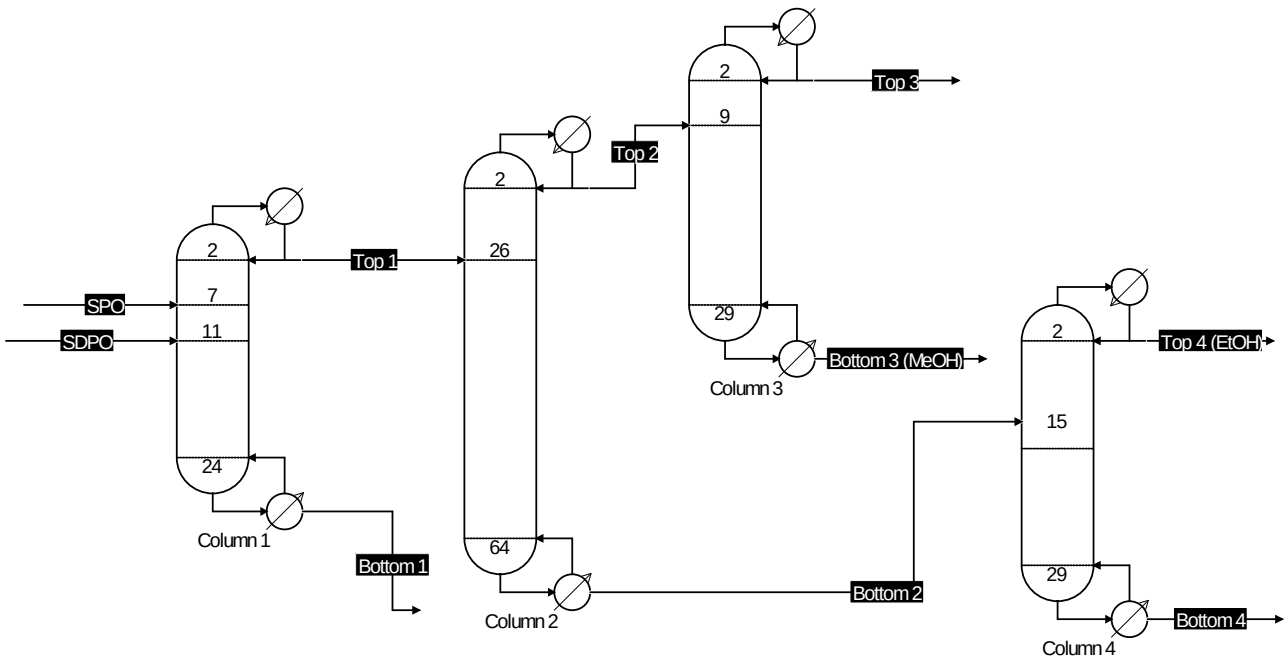

Figure 2. Process flowsheet implemented in COCO-COFE process simulator for methanol and ethanol recovery from combined SPO_{Aq} and $SDPO_{Aq}$ streams.

Process optimisation

The optimal number of theoretical stages and the optimal positioning of the feed to each column are optimised, considering a trade-off CAPEX-OPEX as the criterion. Indeed, increasing the number of theoretical stages results in a decrease in the energy demand for the separation, but it also generates an increase in the investment costs to build the column. More specifically, the minimum number of stages for which an increment lower than 5% is observed in the reboiler duty compared to a configuration with 5 additional stages is selected as the optimum. Noteworthily, only the reboiler duty is considered for the optimisation, while the cooling energy demand for the condenser is neglected. This choice can be justified based on the utilities needed in the reboiler and condenser. Indeed, the operating temperatures in each of the four columns are such that low-pressure steam is required as the utility in the reboiler, while standard cooling water can be used for distillate condensation.

As for the feed positioning, feeding from the top to the bottom of the column is investigated (with a step of two stages), and the feed position associated with the lowest reboiler duty is chosen as the optimum.

Flowsheet implementation

The Block Flow Diagram (BFD) shown in Figure 2 is used as the basis to develop a process flowsheet. The simulation is implemented in COCO-COFE, a license-free and CAPE-OPEN flowsheet environment released by Amsterchem. A size of 3 ton/h of SPO_{Aq} and 3 ton/h of $SDPO_{Aq}$ feed is assumed as representative of a commercial-scale application of the proposed technology.

A screenshot from the simulation developed for the combined SPO_{Aq} and $SDPO_{Aq}$ feed is drawn in Figure 2.

RESULTS

Optimal column sizing

The optimal theoretical stage number and the optimal feed positioning identified for each column are gathered in Table 3. These are the same for all three case studies.

Table 3: Results of distillation column layout optimisation.

Column	Optimal N stages	Optimal feed stage
Col. 1	25	7, 11
Col. 2	65	26
Col. 3, scenario A	20	2
Col. 3 scenario B	35	25
Col. 4	35	22

Product recovery rate and purity

The yield and purity of methanol and ethanol obtained from the optimised process model are gathered in Table 4.

Table 4: Methanol and ethanol yield (%) and purity (wt%) obtained in the different modelled process configurations. Components' acronyms are retrieved from Table 1.

Comp.	Scen.	KPI	CS1	CS2	CS3
MeOH	A	Yield	56.0	-	51.0
MeOH	A	Purity	95.5	-	92.2
MeOH	B	Yield	94.6	-	94.9
MeOH	B	Purity	88.1	-	81.1
EtOH	-	Yield	98.0	93.1	98.0
EtOH	-	Purity	85.5	85.7	86.0

Results indicate that there is a trade-off between methanol yield and purity. Indeed, a high recovery rate can be achieved (95%, scenario B), but the corresponding product quality is relatively poor (81-88 wt%). On the other hand, increasing the methanol purity up to 95 wt% is possible (scenario A), but in this case, only 56% of the total methanol content present in the feed is recovered in the final product stream. The reason behind this observation is that the presence of volatile impurities negatively affects the product quality, as these compounds have boiling points close to methanol, and some of them form non-ideal interactions, which makes their separation via distillation more challenging. The detailed composition of the methanol product stream obtained when using SPO_{Aq} only, or combined SPO_{Aq} and $SDPO_{Aq}$ feed, is reported in Table 5.

Table 5: Methanol composition (wt%) in the different modelled process configurations. Components' acronyms are retrieved from Table 1.

Component	CS1	CS1	CS3	CS3
Scenario	A	B	A	B
MeOH	95.50	88.14	92.20	81.06
EtA	4.08	8.94	7.11	15.88
THF	0.40	2.36	0.67	2.85
Acet	-	0.32	-	0.17

The most abundant impurity is ethyl acetate, which is predicted to reach concentrations above 7 wt% in the combined feed case study. Indeed, the re-parametrised NRTL model used in this work indicates that methanol (MeOH) and ethyl acetate (EtA) form a minimum-boiling azeotrope at a methanol concentration of 70 vol% (see Figure 3), and the bubble and dew curves become flat and close for higher methanol concentrations, which indicates that the mixture is difficult to separate by conventional distillation methods. Notably, MeOH-EtA VLE represents the most constraining azeotropic behaviour on the methanol product purity, and this observation holds for all the analysed case studies.

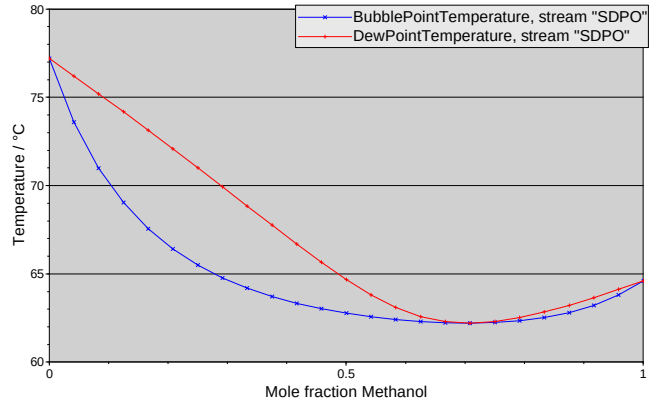

Figure 3. Temperature-composition VLE diagram for the binary mixture MeOH-EtA obtained with the re-fitted NRTL model.

Ethanol is obtained at a purity ranging between 85 wt% and 86 wt%. It contains residual water and C_3 alcohols, mainly isopropanol. Indeed, isopropanol forms an azeotrope with a comparable composition and boiling point of the water-ethanol one, which hinders their separation. The detailed ethanol product stream composition is gathered in Table 6.

Table 6: Ethanol composition (wt%) in the different modelled process configurations. Components' acronyms are retrieved from Table 1.

Component	CS1	CS2	CS3
EtOH	85.49	85.74	85.98
H2O	7.93	7.88	7.94
2-PrOH	6.28	4.90	5.68
1-PrOH	0.28	0.66	0.39
MeOH	-	0.81	-

The highest methanol purity is achieved in CS1. Given that a substantial amount of residual volatile organics is present in the final product, maximising yield (scenario B) looks the most promising option. On the other hand, ethanol yield and purity prove to be not affected by the type of feedstock, with comparable results achieved in all the investigated case studies.

Energy requirements

The total thermal energy demand of the process is calculated as the sum of the reboiler duties of all four columns (three for CS2), and the results are shown in Table 7. As the bottom temperature ranges from 63°C to 99°C in all columns, low-pressure steam is the required utility.

CS2 is the option associated with the lowest energy demand, but in this case only ethanol is recovered as a product. Recovering methanol and ethanol from both SPO and SDPO aqueous streams in the same plant (CS3) is more energy efficient than treating the SPO aqueous stream only. The total duty is around 1 MJ/kg of treated

aqueous stream, resulting in a 16% reduction in the heat demand.

Table 7: Specific energy requirement (in MJ/kg feed) of the distillation process for methanol and ethanol recovery.

Case study	Scenario A	Scenario B
CS1	1.16	1.15
CS2	0.73	0.73
CS3	0.95	0.97

CONCLUSIONS

A distillation-based process for the recovery of valuable organic compounds, namely alcohols, from pyrolysis oil wastewater stream has been designed, modelled, and optimised. This study shows that there is potential for ethanol and methanol recovery for applications in fuel blending. These alcohols can be effectively separated from water and other heavier compounds: high recovery rates (> 95%) can be achieved by adopting a simple process layout made up of four distillation columns in series, which minimises investment costs. Moreover, no third compound is introduced to avoid contamination effects.

However, the presence of detectable concentrations of esters (i.e., methyl acetate and ethyl acetate), and furans (like tetrahydrofuran) together with alcohols in the SPO and SDPO-derived aqueous mixtures represents a major challenge since all these species have a similar boiling point range as methanol and ethanol. As a result, the obtained alcohol purity is limited.

To target applications requiring a higher degree of purity, such as in the chemical industry, the separation scheme becomes more complex, and an unspecified number of additional purification steps/columns need to be added in a future process layout.

ACKNOWLEDGEMENTS

The FUEL-UP project has received co-funding from the European Union's Horizon Europe Research and Innovation Programme under Grant Agreement N° 101136123.

The authors express their gratitude to BTG (Netherlands) for their precious contribution concerning feedstock supply and to Avecom (Belgium) for their support and cooperation in sample analysis.

AUTHOR IDENTIFIERS

Author ORCIDs:
Gilardi Matteo: 0000-0003-2456-7549
Bisotti Filippo: 0000-0001-9106-1797
Rune Myrstad: 0000-0002-8694-496
Bernd Wittgens: 0000-0001-5705-9255

REFERENCES

1. https://www.btgworld.com/media/xpqpaucv/advanced-biofuels-from-fpbo.pdf
2. BTG Biomass Technology Group BV. Production of 200 L HPO. Deliverable 3.4 of REFOLUTION project. file:///C:/Users/matteog/Downloads/Attachment_0%20(3).pdf
3. FUEL-UP project. Powering a greener future for aviation and marine transport. https://www.fuelup-project.eu/
4. Mora M, Fàbregas E, Céspedes F, Bartrolí J, Puy N. Production and separation of value-added compounds from pine wood using pyrolysis and biorefinery techniques. Fuel Processing Technology 238:107509 (2022). https://doi.org/10.1016/j.fuproc.2022.107509
5. Smets K, Adriaensens P, Vandewijngaarden J, Stals M, Cornelissen T, Schreurs S, Carleer R, Yperman J. Water content of pyrolysis oil: comparison between karl fischer titration, gc/ms-corrected azeotropic distillation and 1H NMR spectroscopy. Journal of Analytical and Applied Pyrolysis 90:100-105 (2011). https://doi.org/10.1016/j.jaap.2010.10.010

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Intensified liquid-liquid process design for critical metals extraction from e-waste

Konstantinos Katsoulas[a], Arun Pankajakshan[a], Malik Olasinde[b], Cong Chao[b], Federico Galvanin[a], Panagiota Angeli[b], and Eric S. Fraga[a]*

[a] Sargent Centre for Process Systems Engineering, Department of Chemical Engineering, University College London, Torrington Place, WC1E 7JE, London, UK
[b] ThAMeS Multiphase, Department of Chemical Engineering, University College London, Torrington Place, WC1E 7JE, London, UK
* Corresponding Author: e.fraga@ucl.ac.uk.

ABSTRACT

Critical metals are essential for clean energy technologies but, due to being mainly sourced through mining, the critical metal supply chain is susceptible to geopolitical risks. Electronic waste (e-waste), however, can serve as an alternative "urban mine", but the recovery at high purities requires complex and resource-intensive processing. This work explores the modeling and optimization-based design for the intensification of liquid-liquid extraction in small channels as a means to recover critical metals from e-waste. Small channels can achieve high mass transfer rates while mitigating the environmental impact. A superstructure-based approach is employed to represent the alternative system configurations, while a plant propagation algorithm is used to optimize the multi-objective problem to recover Neodymium (Nd) and Samarium (Sm). The multi-objective problem aimed to tackle product quality, process economics, and environmental impact. The results demonstrated that optimally designed extraction can support efficient recovery of critical metals; however, decision making support may be essential for the selection of one process design from those identified as optimal in a multi-objective design approach.

Keywords: Superstructure optimization, multi-objective optimization, extraction, critical metals, process intensification

INTRODUCTION

Critical metals are integral to renewable low carbon and hydrogen technologies and essential for the global energy transition since they are necessary to produce solar panels, batteries, wind turbines, electronics and catalysts (Department for Business & Trade, 2025). These metals are mainly sourced through mining; however, only few countries have access to available ores, with China controlling more than 90% of the global rare earth market and South Africa controlling most of the platinum group metals market (*Recycling Rare Earth Metals - Mathys & Squire LLP*, n.d.). With the global demand in critical metals being higher than ever, their availability has become susceptible to geopolitics, thereby threatening the supply chains. Alternatives to mining include the recovery of critical metals from the increasing volumes of electronic waste (e-waste), which can serve as "urban" mines with critical metal concentrations often higher than in primary ores (Hagelüken, 2012). However, the recovery of metals from e-waste at required purities necessitates multiple processing steps, including solvent extraction and recovery steps.

Metals recovery from waste is traditionally carried out via pyrometallurgical or hydrometallurgical routes. Hydrometallurgical routes are preferred since these do not involve hazardous emissions and offer high recovery and selectivity for a wide range of metals. Hydrometallurgical routes involve the dissolution of the solid waste in acidic solutions followed by a solvent extraction step responsible for selectively separating the metals into solvent phases. Traditionally, dissolution and extraction are carried out in batch systems, and thus are usually characterized by low efficiency, while they are difficult to monitor and digitize due to non-uniformities associated with the batch processing. In lieu of batch dissolution and extraction, intensified flow processing in small channels presents an alternative that can enable the transition to

fully monitored continuous processes, owing to enhanced mass transfer rates that can eliminate flow non-uniformities, increase efficiency, and decrease solvent usage. Throughput in intensified small channel processing can be increased in a modular way by increasing the number of channels and not their volume, thus also facilitating an easier transfer from lab to pilot to industrial settings and ensuring process safety.

In this work, we explore the multi-objective optimal design for the separation of Neodymium (Nd) and Samarium (Sm), primarily used for magnets found in electronics and electric vehicles, using continuous flow liquid-liquid extraction (LLE) in small channels (Garciadiego-Ortega et al., 2020). Mass transfer in the channels is modeled as a set of differential equations. We follow a superstructure-based approach to determine the optimal flowsheet by accounting for multiple extraction steps (Mencarelli et al., 2020). Each step may consist of several stages. The multiple objectives reflect final products quality, process economics and environmental footprint (e.g. solvent usage). The optimization problem is solved employing Fresa, a nature inspired plant propagation optimization algorithm (Fraga, 2021).

This manuscript is structured as follows: first, an overview of the extraction system under consideration is presented. Then, the modeling components, simulation and optimization methods are provided. The proposed system is then used to implement a case study and the results are discussed. Finally, some concluding remarks and future directions are provided.

METHODOLOGY

System under consideration

For this work, we considered the continuous flow LLE system in small channels adapted from Pineda et al., (2021). In the channels, high mass transfer rates can be achieved by inducing segmented flow, that is a repeating pattern of two immiscible liquids moving as elongated drops (plugs) separated by continuous phase slugs. The reader is referred to the work of Garciadiego-Ortega et al. (2020) for a detailed presentation of segmented flow in small channels. This extraction system employs two extraction steps and one scrubbing step. A feed of known concentration and a stream of fresh organic solvent enter the first extraction step, E1. For the organic solvent, we assumed a mixture of 2-ethylhexyl diamide amine in n-dodecane. The organic outlet of E1 is introduced into a scrubbing step, SCR, to further purify the organic stream, whilst the aqueous outlet of E1 is introduced into subsequent extraction step, E2. The mixers in the flowsheet of Fig. 1 are used to combine the two distinct organic and aqueous streams, respectively. Note that the mixers used in this work do not appear in the work of Pineda et al. (2021) since the considered system was used to optimally detect trivalent Samarium and Neodymium in nuclear waste, while also considering uncertainty. In the work of Pineda et al. (2021), each extraction step was equivalent to a single stage, and thus a single channel. The current work aims to identify optimal designs for the recovery of the two metals employing multi-stage extraction steps. The number of stages is a design variable as explained in a later section of this paper. Fig. 2 depicts a multi-stage configuration where the flow is introduced co-currently in each stage (channel) but counter-currently overall in a given processing step: the feed is introduced in the first stage of a step whilst the solvent is introduced in the last stage. This pseudo-countercurrent configuration aids in more efficient separations (Bascone et al., 2019) and thus is preferred to co-current configurations.

Modeling

Modeling solvent extraction in small channels is a problem well addressed in the open literature. Although computational fluid dynamics (CFD) models offer rigor and accuracy (Zhang et al., 2019), they can hinder the solution of the optimization problem due to the computational burden that arises from such models. Instead, we can use a simplified model that consists of two differential mass balances, for the aqueous and organic phases (Pineda et al., 2021). Thus, the differential mass balance for the aqueous phase concentration $C_{a,i,j}$ of the i-th component of the j-th processing step reads:

$$v_{a,j}\frac{\partial}{\partial x}C_{a,i,j}(x) = -K_j\left[C_{a,i,j}(x) - C^*_{a,i,j}(x)\right] \text{ for } x = 0, \dots, L_j \quad (1)$$

where $C^*_{a,i,j}$ is the equilibrium concentration given by $C^*_{a,i,j} = C_{o,i,j}/a_i$ with a_i being the distribution coefficient of the i-th component. We used values of $a_{Nd} = 3.7$ and $a_{Sm} = 0.3$, for the extraction and scrubbing steps described above, that come from the work of Suzuki et al. (2017). The mass balance of the organic phase concentration $C_{a,i,j}$ of the i-th component of the j-th step reads:

$$\frac{\partial}{\partial x}C_{o,i,j}(x) = -q_j\frac{\partial}{\partial x}C_{a,i,j}(x) \text{ for } x = 0, \dots, L_j \quad (2)$$

where K_j is the mean overall volumetric mass transfer coefficient, L_j is the channel length, $q_j = \frac{\dot{Q}_{a,j}}{\dot{Q}_{o,j}}$ is the ratio between the volumetric flow rate of the aqueous phase, $\dot{Q}_{a,j}$, and the organic phase, $\dot{Q}_{o,j}$, of each step, $v_{a,j} = \frac{\dot{Q}_{a,j}}{A_j}$ is the superficial velocity of the aqueous phase, and $A_j = \pi D_j^2/4$ is the cross sectional area of each channel in a processing step. To estimate the mean overall volumetric mass transfer coefficient, K_j, a hydrodynamics correlation is used (Kashid et al., 2010):

$$K_j = \left(\frac{\dot{Q}_{a,j} + \dot{Q}_{o,j}}{A_j}\right)\left(\frac{Ca^{-0.09}Re^{-0.09}}{L_j}\right)\left(\frac{D_j}{L_j}\right)^{-0.1} \quad (3)$$

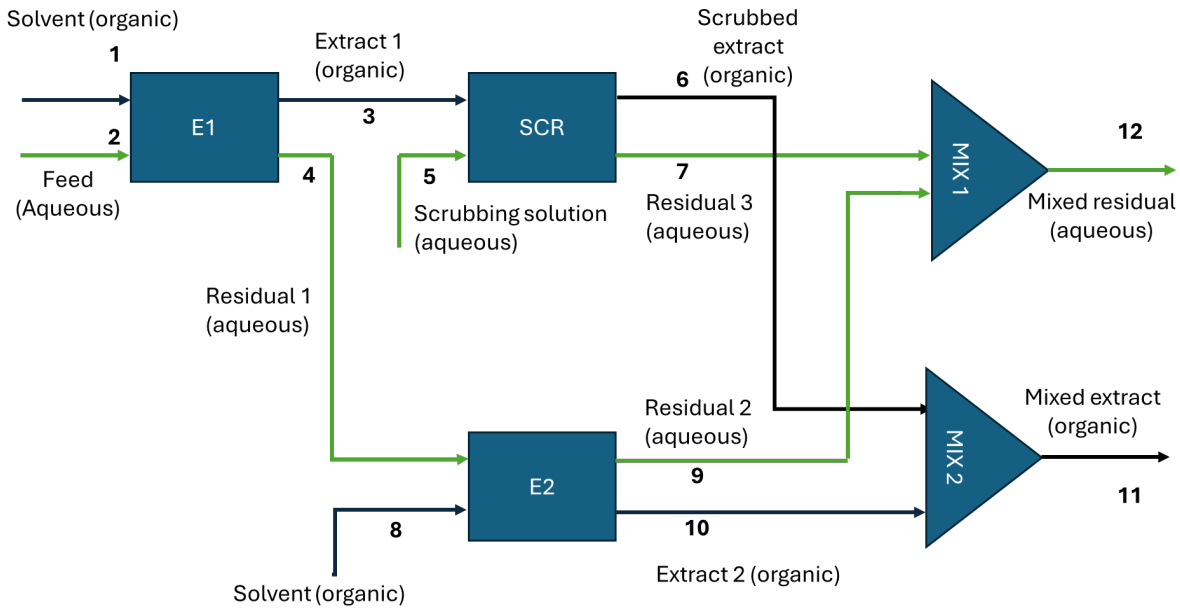

Figure 1. Schematic of the continuous flow LLE system considered in this work.

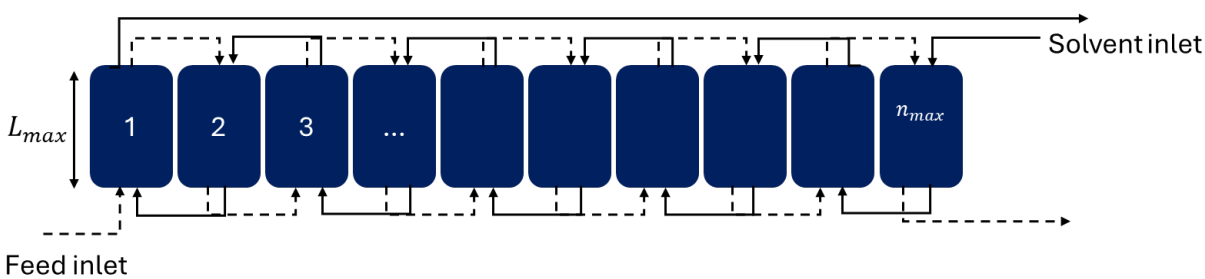

Figure 2. Schematic of the pseudo-countercurrent multi-stage configuration within each step of the system considered in this work.

where Ca and Re are the capillary and Reynolds numbers, respectively, and D_j is the diameter of the channels of each step.

Simulation and optimization

We employ a superstructure-based approach to determine the optimal flowsheet for the separation process (Mencarelli et al., 2020). The process flowsheet involves two distinct extraction steps and a scrubbing step as depicted in Fig. 1. Each extraction step may comprise of several stages (in this work: single-channel stages). The Julia programming language (Bezanson et al., 2017) is used to model the flowsheet as a superstructure of extraction steps (scrubbing included), and in turn, each extraction step as a superstructure of stages. This approach offers good scalability and modularity. The pseudo-countercurrent configuration (see Fig. 2), implemented in Julia, results in multiple recycle streams. The streams have to be torn and converged in order to simulate the superstructure. Therefore, to converge the streams we have implemented Broyden's method in Julia (Broyden, 1965).

Optimal designs are generated by considering a search space defined by the decision variables vector $\mathbf{x} = [q_j, \xi_j, D]$, where ξ_j is a decision variable that defines both the number of stages n_j as well as the length L_j of the j-th step (i.e., channel). Note that we assume the length stays uniform within a given step, that is all channels share the length, although it can vary across the different processing steps. This assumption aims to facilitate structural simplicity of the flowsheet in order to maintain low capital cost and promote scalability of the plant. In parallel, a distinct channel length variable for each stage would drastically increase the search space and potential solutions, thus leading to an exponential increase in computational time.

Therefore, given $\xi_j \in [0,1], L_{\max,j}, L_{\min,j}, n_{\max,j}$:

$$n_j \leftarrow \lceil \xi_j n_{\max,j} \rceil \tag{4}$$

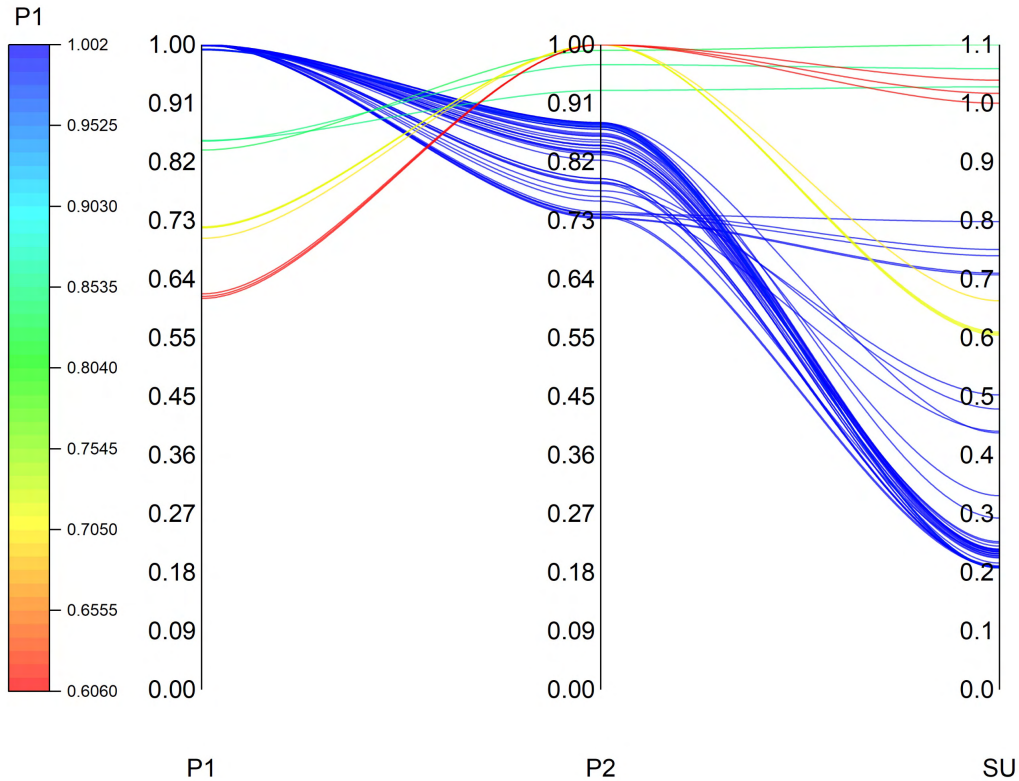

Figure 3. Parallel coordinate plot of the non-dominated solutions obtained by Fresa for the settings in Table 4 to solve the three-objective problem defined by Eqs. (6)-(8). P1 and P2 correspond to the purities of Nd and Sm, respectively, while SU is the organic solvent usage in m^3/yr.

$$L_j \leftarrow (\xi_j n_{\max,j} - n_j + 1)(L_{\max,j} - L_{\min,j}) + L_{\min,j} \qquad (5)$$

The formulation described in Eqs. 4 and 5 is useful because a single value for ξ_j (not physically meaningful) is passed down to define two physically meaningful design variables, n_j and L_j. This decreases the dimensionality of the problem, thus reducing the potential computational cost. Moreover, the above formulation deals with discrete decisions (e.g., the number of stages per step) by relying solely on continuous variables. This mitigates the combinatorial explosion that would be caused if we employed discrete or binary variables instead. This representation of the design space is effective since it includes all the designs of interest, yet it mitigates any infeasible designs. However, there will be potential designs, that albeit feasible, might be impractical.

We employed Fresa to solve the optimization problem. Fresa is an implementation of a plant propagation algorithm in the Julia programming language (Fraga, 2021). Fresa is a nature-inspired evolutionary algorithm that mimics the propagation of strawberry plants. In fertile soil, strawberry plants send runners to exploit (i.e. refine) the surroundings, while in less fertile soil, fewer but longer runners are sent to explore the *design space* to find more fertile soil. We refer the reader to original paper Salhi & Fraga (2011) for a detailed description of the plant propagation algorithm.

CASE STUDY AND RESULTS

Neodymium (Nd) and Samarium (Sm) are metals that can be considered as rare earth elements. Applications of these metals include, but are not limited to, magnets for electronic devices, electric vehicles (EVs), and wind turbines. Hence, these metals are pivotal in energy transition. Pineda et al. (2021) have explored optimal design of continuous flow LLE in small channels with the aim to optimally *detect* Nd and Sm radionuclides during the decommissioning of nuclear power plants. In this work, we drew data from the work of Pineda et al. (2021) to develop designs to achieve the best *separation* of Nd and Sm.

As stated above, the main goal of this work is to intensify the separation of critical metals in order to achieve high metals purities and recoveries, while minimizing the environmental impact and operational cost by minimizing the organic solvent usage. The above goals are addressed through the definition of a three-objective

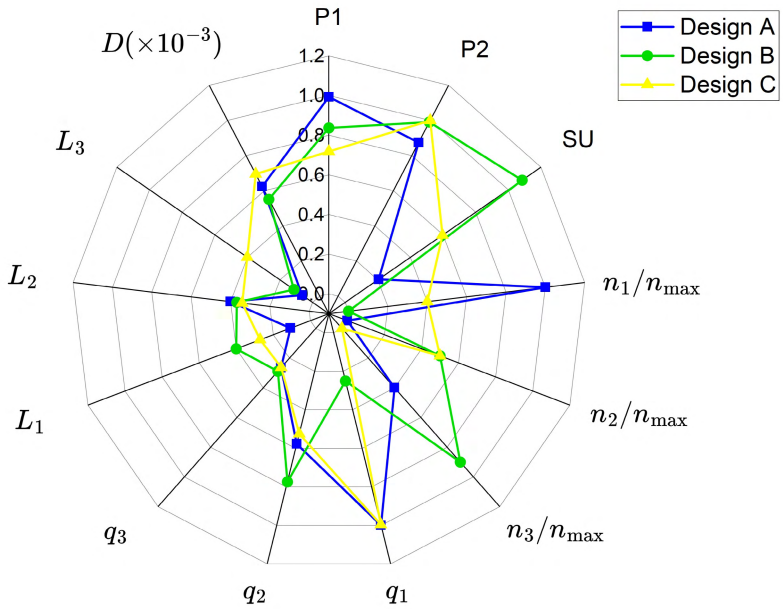

Figure 4. Radar chart of representative designs drawn from the set of non-dominated solutions (objectives and design variables) of the three-objective problem defined by Eqs. (6)-(8).

optimization problem:

$$\max_{\mathbf{x}} f_1(\mathbf{x}) = \frac{C_{\text{Nd, 11}}}{C_{\text{Nd, 11}} + C_{\text{Sm, 11}}} \equiv \text{Purity}_{\text{Nd}} \tag{6}$$

$$\max_{\mathbf{x}} f_2(\mathbf{x}) = \frac{C_{\text{Sm, 12}}}{C_{\text{Nd, 12}} + C_{\text{Sm, 12}}} \equiv \text{Purity}_{\text{Sm}} \tag{7}$$

$$\min_{\mathbf{x}} f_3(\mathbf{x}) = (\dot{Q}_1 + \dot{Q}_8)\frac{\text{m}^3}{\text{s}} \times 31,536,000\ \frac{\text{s}}{\text{yr}} \tag{8}$$

Note that while only the purities of the two components are explicitly optimized, the component recoveries are also implicitly optimized, due to the Purity of Nd being associated to the Recovery of Sm and vice versa. To gain better insight into why this happens, the recoveries of the components read:

$$\text{Recovery}_{\text{Nd}} \equiv \frac{C_{\text{Nd,11}}\dot{Q}_{11}}{C_{\text{feed}}\dot{Q}_{\text{feed}}} \tag{9}$$

$$\text{Recovery}_{\text{Sm}} \equiv \frac{C_{\text{Sm,12}}\dot{Q}_{12}}{C_{\text{feed}}\dot{Q}_{\text{feed}}} \tag{10}$$

From Eqs. (6), (7), (9) and (10), we can deduce that to increase the purity of (for instance) Nd the concentration of Sm in Stream 11 must be reduced. Since there are only two final outlet streams as depicted in Fig. 1 (Streams 11 and 12), reducing the concentration of Sm in Stream 11 translates into concentrating Sm in Stream 12, and thus increasing its recovery. The same stands for the Purity of Sm and Recovery of Nd. Note also that Eq. 8 targets only the organic solvent streams (Streams 1 and 8) and accounts for solvent consumption in a yearly basis operation.

The extraction system consists of four streams of known concentration. The feed enters E1 at a given metal concentration and known flow rate. A stream of fresh organic solvent also enters E1 whose flow rate is determined by the design variable, q_1., Stream 8 that contains fresh organic solvent enters a subsequent extraction step, E2, at a flowrate determined by q_2, whilst the fresh aqueous solvent stream that enters the scrubbing step, Stream 5, enters at a flowrate determined by q_3.Table 1 summarizes the streams of known concentrations.

Table 1: Streams of known concentrations in the liquid-liquid extraction flowsheet. Flow rate values that depend on the optimal designs are denoted by To Be Defined (TBD).

Stream	Description	Concentration (M)		Flow rate (m^3/s)
		Nd	Sm	
Feed (2)	Aqueous	0.001	0.001	1.3×10^{-8}
1	Organic	0.0	0.0	TBD
8	Organic	0.0	0.0	TBD
5	Aqueous	0.0	0.0	TBD

The mass balance of Eq. 1 features the mean overall volumetric mass transfer coefficient which is calculated based on Eq. 3. The dimensionless numbers featured in Eq. 3 require values for the physical properties of the mixture that flows through the channel. Thus, we used values for the interfacial tension, the viscosity, and density of the mixture that we drew from the work of Pineda et al. (2021). The values for both the organic and the

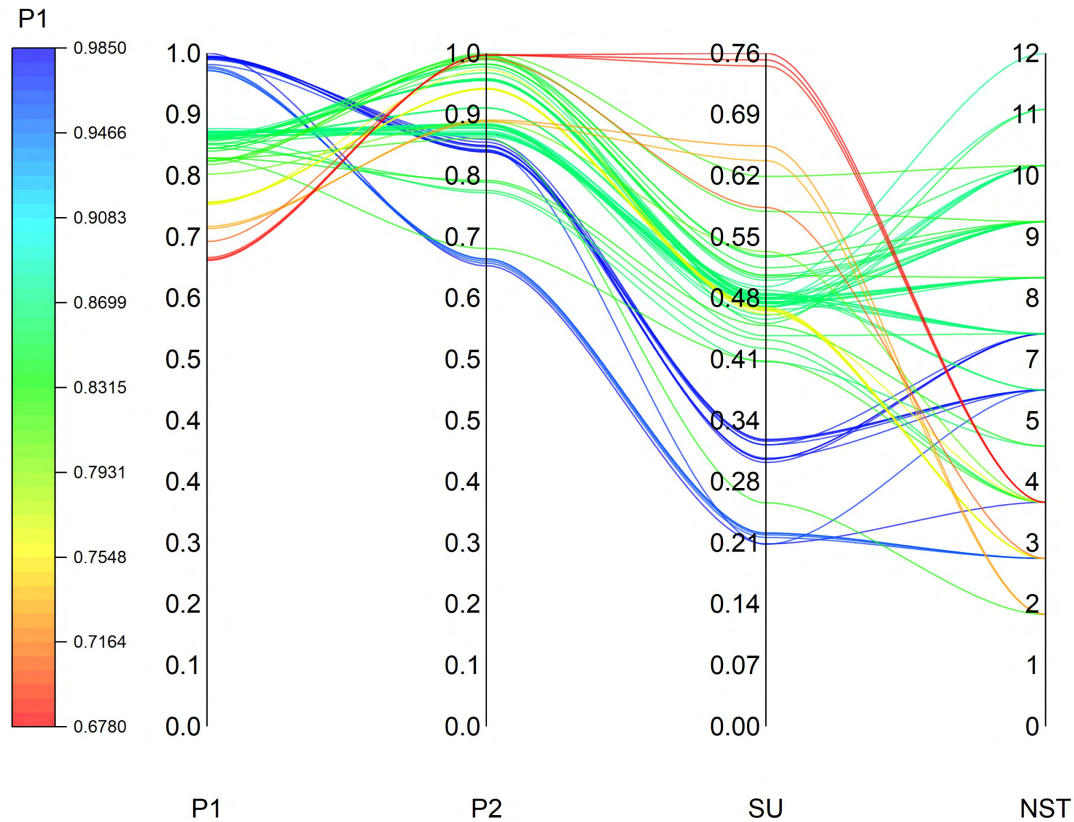

Figure 5. Parallel coordinate plot of the non-dominated solutions obtained by Fresa to solve the four-objective problem defined by Eqs. (6)-(9). P1 and P2 correspond to the purities of Nd and Sm, respectively, SU is the organic solvent usage in m^3/yr, and NST is the number of stages total.

aqueous mixtures are summarized in Table 2.

As discussed previously, a vector of design variables, $\mathbf{x} = [q_j, \xi_j, D]$, is employed to optimize the metal recovery process. These variables are bounded due to operational (e.g., flow rate ratio to maintain segmented flow) or structural (e.g., channel length) constraints. Although, channel length, L_j, and number of stages, n_j, are not directly leveraged, they are implicitly defined through the ξ_j variable; hence, upper and lower bounds for those design pseudo-variables are also required. Table 3 summarizes the lower and upper bounds for the physical and pseudo- design variables.

Table 2: Thermophysical properties of the liquid mixtures in the process flowsheet.

Property	Value	
	Organic	Aqueous
Interfacial tension, γ $(\mathrm{N\,m^{-1}})$	0.01	0.01
Viscosity, μ $(\mathrm{kg\,m^{-1}s^{-1}})$	0.0023	0.001
Density, ρ $(\mathrm{kg\,m^{-3}})$	844	997

Table 3: Lower and upper bounds for the design (pseudo-) variables. Regular variables appear in bold while pseudo-variables in plain font.

Variable	Value	
	Lower	Upper
$\boldsymbol{q_j}$ $(-)$	**0.5**	**2.0**
$\boldsymbol{\xi_j}$ $(-)$	**0.0**	**1.0**
L_j (m)	0.05	0.4
n_j $(-)$	0	10
$\boldsymbol{D}$ $(\mathbf{m})$	**0.001**	**0.005**

Solving the three-objective problem resulted in 57 non-dominated solutions in 10101s (≈ 168 min or ≈ 2.8 hrs) on a 48-core Intel® Xeon® w7-2495X @4800 MHz and 250 GB of memory. The non-dominated solutions are summarized in a parallel coordinate plot of Fig. 3. The solutions in Fig. 3 appear color-coded in a scale from low to high purity of Nd. Fig. 3 shows a clear trade-off between the purity of the two components, demonstrated by the intersection of lines between the purity of Nd, P1, and the

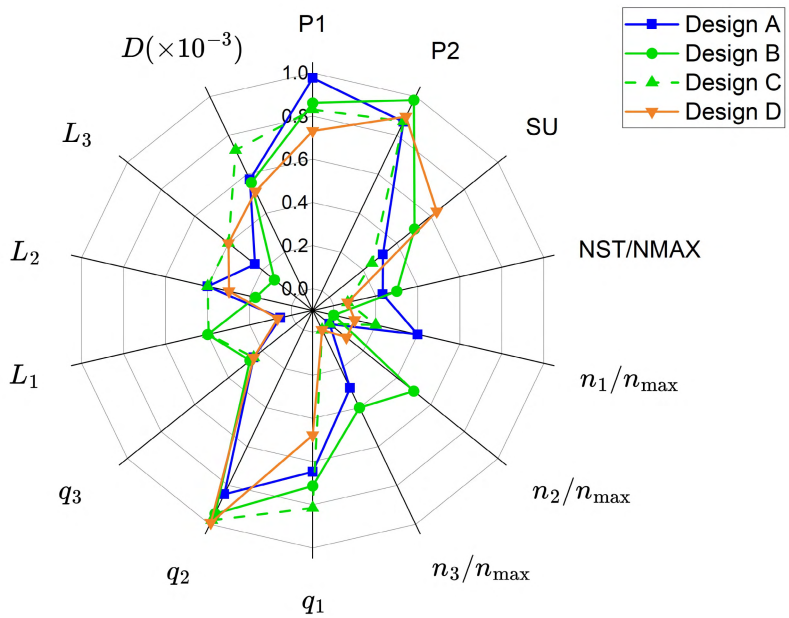

Figure 6. Radar chart of representative designs drawn from the set of non-dominated solutions (objectives and design variables) of the four-objective problem defined by Eqs. (6)-(9).

purity of Sm, P2. Therefore, the third objective, solvent usage, SU, is an important deciding factor, assuming that the first two objectives are of equivalent significance.

To gain better insight into the optimal designs, we have selected representative solutions to plot on the radar chart of Fig. 4, which shows the values of the objectives together with the values of the design variables. Design A results in a P1 ≈ 0.99 and P2 ≈ 0.87, whilst keeping SU ≈ 0.20 requiring ten stages for E1 (the reader is referred to Fig. 1), no stages for E2, and four stages for E3. Employing E3, the scrubbing step, as well as not employing the second extraction step, E2, results in a very high purity of Nd at the expense of lowering the purity of Sm. Designs that had non-zero stages for E2 led to potentially lower P1, mainly due to back extraction and higher SU, since an additional stream of fresh organic solvent is employed, and thus were dominated by other solutions. Design B, on the other hand, is characterized by P1 ≈ 0.84, P2 ≈ 0.99, and SU ≈ 1.09. While the trend in purities has reversed, vis-à-vis Design A, SU is significantly higher. Specifically, Design B disabled E1 (no stages), while it required five and nine stages for E2 and E3, respectively. However, enabling step E2 without enabling step E1 is essentially incoherent. If there are no stages for E1, no separation takes place, and the streams entering E3 are empty; hence, E3 stages in this design have no impact whatsoever in the objective values. The reason this design appeared in the set of non-dominated solutions is potentially due to the stochasticity of the optimization method that did not potentially converge to the global Pareto frontier, and although one would not normally consider it, it still presents as a valid design. Another representative solution is that of Design C, which comprises of P1 ≈ 0.72, P2 ≈ 0.99, and SU ≈ 0.6. Although this design achieves lower P1 compared to Design B, it achieves savings in SU by employing four stages for E1 and five stages for E2. To select an optimal design amongst the three representatives, one must weigh in process profitability (e.g., component purity *vs* price) and operational cost/environmental footprint (e.g., solvent cost). However, capital cost may be considered, as well. Although the total number of allowable stages is relatively small (i.e., 30) and might not impact considerably the economics of a pilot plant, the scaling out of the process, that is installing multiple channels in parallel for every stage to increase throughput, can have a big impact on capital cost. Therefore, we must investigate the minimization of total stages in addition to the three objectives we have already considered.

Note that although the superstructure does not consider the scaled-out process, capital cost is reflected through the number of stages per step. Thus, the fourth objective, in addition to the original three of Eqs. (6)-(8) reads:

$$\min_{\mathbf{x}} f_4(\mathbf{x}) \equiv \sum_{j=1}^{3} n_j \tag{11}$$

The four-objective problem was solved using the same feeds, concentrations, and thermophysical properties as before, yielding 75 non-dominated solutions in 7136 s (≈ 118 min or ≈ 2 hrs). The resulting Pareto set, shown in Fig. 5, exhibits the same trade-off between Nd and Sm

purities observed in the three-objective case. Four representative designs are presented in Fig. 6. Design A achieves high purities (P1 ≈ 0.98, P2 ≈ 0.87) with moderate solvent usage (SU≈ 0.32) using seven total stages allocated to E1 and E3. Design B prioritizes Sm purity (P2 ≈ 0.98) but incurs higher solvent usage (SU ≈ 0.5) and exhibits counterintuitive stage allocation, again disabling E1. Designs C and D highlight low-complexity alternatives: Design C achieves balanced purities with minimal solvent usage using only two stages in E1, while Design D trades Nd purity for a small gain in Sm purity at substantially higher solvent consumption.

The results show that no single optimal configuration exists; instead, families of non-dominated optimal designs emerge, reflecting trade-offs between separation performance, solvent usage, and total stages. While near-perfect Nd and Sm purities typically require multiple stages, comparable performance can often be achieved with fewer stages, albeit sometimes at higher solvent consumption. Overall, the multi-objective framework provides clear decision support for balancing performance, operational sustainability, and capital cost.

CONCLUSIONS

This work focused on modeling and multi-objective optimization to intensify the design of continuous flow LLE in small channels for critical metals recovery from e-waste. A superstructure-based approach was used to represent alternative flowsheet configurations with variable numbers of extraction and scrubbing stages, while mass transfer in segmented flow was described using simple differential balances to balance accuracy and computational efficiency. Discrete design decisions, such as stage allocation, were handled through a continuous pseudo-variable formulation to limit combinatorial complexity. The resulting nonlinear optimization problem was solved using the plant propagation algorithm implemented in Fresa.

The framework reveals key trade-offs between separation performance, solvent usage, and process complexity. Several non-dominated solutions show that high metal purities can be achieved with relatively simple configurations, highlighting opportunities for more compact and scalable designs. Overall, the approach supports informed selection of designs that balance performance with practical implementation constraints. Future work will consider tackling the intensification of metals recovery under feed variability conditions and parametric uncertainty.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge funding received from UK Research & Innovation through the grant EP/Z533798/1 which supported the research presented in this paper.

AUTHOR IDENTIFIERS

Author ORCIDs:
Katsoulas K: 0009-0004-9833-0127
Pankajakshan A: 0000-0002-3365-2668
Olasinde M: 0000-0001-6001-1404
Galvanin F: 0000-0002-3296-1581
Angeli G: 0000-0001-9575-4259
Fraga ES: 0000-0002-5732-6082

REFERENCES

1. Bascone D, Angeli P, Fraga ES. Process intensification applied to spent nuclear fuel reprocessing: an alternative flowsheet using small channels. Chemical Engineering and Processing - Process Intensification 143:107618 (2019). https://doi.org/10.1016/j.cep.2019.107618
2. Bezanson J, Edelman A, Karpinski S, Shah VB. Julia: a fresh approach to numerical computing. SIAM Rev. 59:65-98 (2017). https://doi.org/10.1137/141000671
3. Broyden CG. A class of methods for solving nonlinear simultaneous equations. Math. Comp. 19:577-593 (1965). https://doi.org/10.1090/s0025-5718-1965-0198670-6
4. Department for Business & Trade. (2025). *Vision 2035: Critical Minerals Strategy - GOV.UK*. https://www.gov.uk/government/publications/uk-critical-minerals-strategy/vision-2035-critical-minerals-strategy
5. Fraga, E. S. (2021). International Journal on Engineering Technologies and Informatics, 2(4).
6. Garciadiego-Ortega E, Tsaoulidis D, Pineda M, Fraga ES, Angeli P. Hydrodynamics and mass transfer in segmented flow small channel contactors for uranium extraction. Chemical Engineering and Processing - Process Intensification 153:107921 (2020). https://doi.org/10.1016/j.cep.2020.107921
7. Hagelüken BC. Recycling the platinum group metals: a european perspective. Platinum Metals Review 56:29-35 (2012). https://doi.org/10.1595/147106712x611733
8. Kashid MN, Gupta A, Renken A, Kiwi-Minsker L. Numbering-up and mass transfer studies of liquid–liquid two-phase microstructured reactors. Chemical Engineering Journal 158:233-240 (2010). https://doi.org/10.1016/j.cej.2010.01.020
9. Mencarelli L, Chen Q, Pagot A, Grossmann IE. A review on superstructure optimization approaches in process system engineering. Computers &

Chemical Engineering 136:106808 (2020). https://doi.org/10.1016/j.compchemeng.2020.106808

10. Pineda M, Tsaoulidis D, Filho PIO, Tsukahara T, Angeli P, Fraga ES. Design optimization of microfluidic-based solvent extraction systems for radionuclides detection. Nuclear Engineering and Design 383:111432 (2021). https://doi.org/10.1016/j.nucengdes.2021.111432
11. *Recycling rare earth metals - Mathys & Squire LLP*. (n.d.). Retrieved January 15, 2026, from https://www.mathys-squire.com/insights-and-events/news/recycling-rare-earth-metals/
12. Salhi, A., & Fraga, E. S. (2011). Nature-Inspired Optimisation Approaches and the New Plant Propagation Algorithm.
13. Zhang H, Wang H, Luo X, Leung DYC, Pang Q, Xu H, Zhang L, Xuan J. Toward a mechanistic understanding of microfluidic droplet-based extraction and separation of lanthanides. Chemical Engineering Journal 356:673-679 (2019). https://doi.org/10.1016/j.cej.2018.09.043

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Sensitivity of MPC Performance to Component Scaling in a Battery–Hydrogen Storage System

Moritz Zebenholzer[a]*, Paul Schwarzmayr[a], Alexander Schirrer[b], Stefan Jakubek[b] and René Hofmann[a]

[a] TU Wien, Institute of Energy Systems and Thermodynamics, Vienna, Austria
[b] TU Wien, Institute of Mechanics and Mechatronics, Vienna, Austria
* Corresponding Author: moritz.zebenholzer@tuwien.ac.at.

ABSTRACT

With the increasing share of volatile renewable energies, there is a growing need for flexible storage systems to balance fluctuations between generation and demand. Multi-energy systems, featuring battery and hydrogen storage systems, provide an efficient and scalable solution for this purpose. While rule-based control is primarily used in industry, model predictive control (MPC) is considered the most promising strategy for cost-optimized and safe operation. The performance of such controllers depends heavily on the capacities, power rates, and degradation behavior of the storage systems; for hydrogen systems, it can also depend on minimum switch-on/off times and ramping rates. This work aims to use sensitivity analysis to quantify the influence of variations in these component parameters on control quality. To this end, a lab-scale battery-hydrogen storage system (TU Wien) is modeled and operated using MPC to minimize grid exchange and thus increase self-sufficiency. Experimental data from electrolysers and fuel cells were used to calibrate the model. Actual energy production and consumption data from one year are condensed into five representative weeks, serving as test cases. In a systematic study, the capacities and power rates of the components are varied. For the given lab-scale system, the maximum power of the electrolyzer has the greatest influence on control performance, followed by the battery and hydrogen storage capacity. The power of the fuel cell and battery has no limiting effect. It can thus be concluded that the electrolyzer is undersized, and the battery should have greater storage capacity to improve performance.

Keywords: Model-predictive-control (MPC), Mixed-integer linear programming (MILP), Sensitivity Analysis, Optimization

INTRODUCTION

The energy transition is increasing the share of renewable sources in electricity generation, particularly from photovoltaics and wind [1, 2]. Their volatile, intermittent feed-in requires storage to balance discrepancies between generation and demand and reduce strain on the grid [2]. Hybrid battery-hydrogen storage systems offer a solution here: batteries with high round-trip efficiency can buffer short-term fluctuations [2]. At the same time, hydrogen enables longer-term shifts thanks to its greater storage capacity [1]. Surplus electricity is converted into chemical energy (hydrogen) in the electrolyzer and later converted back into electricity in the fuel cell [3]. This creates a sector-coupled multi-energy system (MES) with multiple energy sources and control variables [1, 3].

An energy management system is required for optimal and safe operation, which is often implemented in industry through rule-based control; state-of-the-art research, however, focuses on model predictive control (MPC). This control method uses a plant model to calculate optimal control inputs based on, e.g., forecasts of generation and demand [4]. The control quality of MPC depends mainly on the system specification (e.g., storage capacities, component power rates) and the operating behavior (e.g., system dynamics, such as minimum on/off times, ramping constraints, and selected degradation-preventing strategies) [2].

In literature, coupled design and operational optimization for given system requirements, such as a specified demand, is often proposed, leading to cost- and energy-

optimal dimensioning [2]. This raises the following question: how does a given operating control strategy for a specified plant respond to variations in system size around the design point?

For varying generation and demand, Holtwerth et al. [5] demonstrated the dependence of the MPC performance of a hybrid battery-hydrogen storage system. However, they did not investigate the sensitivity of individual design parameters to optimization; instead, they considered the overall system under varying disturbance inputs.

The investigation of design parameter variations is conducted at TU Wien's new hybrid battery–hydrogen storage laboratory, which operates in the low-kW power range. An MPC based on mixed-integer linear programming (MILP) formulations is being implemented for this system. The studied operational optimization goal is to minimize grid exchange, i.e., to shift generation and demand in time through storage management, thereby minimizing grid strain. This requires accurate models of the relevant components, in particular the electrolyzer and fuel cell. Electricity from photovoltaics and wind power is fed into the grid via rooftop systems, and test consumption trajectories can be emulated.

The main contributions of this publication are:

- Modeling of lab-scale system components of electrolyzer and fuel cell
- Numerical experiments with actual generation and demand data for five representative weeks of a year
- Analysis of control quality depending on variations in storage capacities and maximum power rates of components

The remainder of the paper is structured as follows: first, the use case is presented, *TU Wien's* hybrid battery-hydrogen energy storage lab. The electrolyzer and fuel cell are then characterized, followed by the generation and demand curves. The energy management system, a MILP-based MPC, is then introduced. Then, the sensitivity of the control quality to variations in storage capacities and maximum power rates is investigated. The paper concludes with a summary and outlook.

MATERIAL & METHODS

The investigation of variations in design parameters is conducted by modeling and controlling the use case using a MILP-based MPC approach, adapted from [6].

Use Case

The system comprises a battery and a hydrogen energy storage system, consisting of an electrolyzer, a hydrogen storage tank, and a fuel cell, as shown in Figure 1.

The difference between generation (P_{gen}) and demand (P_{de}) in the energy system is compensated for by storage,

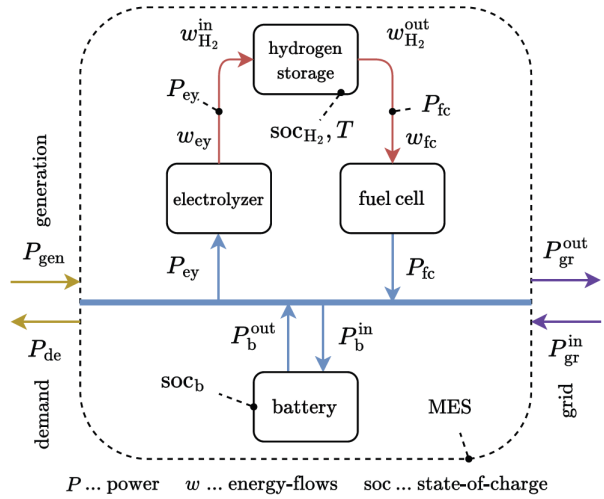

Figure 1. Hybrid battery-hydrogen energy storage system flow diagram with battery (b), electrolyzer (ey), hydrogen storage (H_2) and fuel cell (fc) adapted from [6].

and the grid balances the remaining energy discrepancy ($P_{\mathrm{gr}}^{\mathrm{in}}, P_{\mathrm{gr}}^{\mathrm{out}}$). The battery is defined by its state-of-charge ($\mathrm{soc}_{\mathrm{b}}$), which changes according to

$$\mathrm{soc}_{\mathrm{b},t+1} = \mathrm{soc}_{\mathrm{b},t} + \frac{\Delta t}{C_{\mathrm{b}}} \begin{cases} P_{\mathrm{b},t}^{\mathrm{in}} \bullet \eta_{\mathrm{b}}, \text{ charge} \\ -P_{\mathrm{b},t}^{\mathrm{out}} / \eta_{\mathrm{b}}, \text{ discharge} \end{cases} \tag{1}$$

with charging and discharging power rates ($P_{\mathrm{b}}^{\mathrm{in}}, P_{\mathrm{b}}^{\mathrm{out}}$), respectively, as well as capacity (C_{b}) and efficiency (η_{b}). The time span between consecutive time steps (t) and ($t+1$) is denoted as (Δt). The hydrogen storage state-of-charge ($\mathrm{soc}_{\mathrm{H_2}}$) is modeled analogously,

$$\mathrm{soc}_{\mathrm{H_2},t+1} = \mathrm{soc}_{\mathrm{H_2},t} + \frac{\Delta t \bullet \eta_{\mathrm{H_2}}}{C_{\mathrm{H_2}}} \left(w_{\mathrm{H_2},t}^{\mathrm{in}} - w_{\mathrm{H_2},t}^{\mathrm{out}} \right), \tag{2}$$

with hydrogen flow entering and exiting ($w_{\mathrm{H_2}}^{\mathrm{in}}, w_{\mathrm{H_2}}^{\mathrm{out}}$), respectively, as well as capacity ($C_{\mathrm{H_2}}$) and efficiency ($\eta_{\mathrm{H_2}}$). The power balance links all power values according to

$$P_{\mathrm{gen}} - P_{\mathrm{de}} + P_{\mathrm{gr}}^{\mathrm{in}} - P_{\mathrm{gr}}^{\mathrm{out}} - P_{\mathrm{b}}^{\mathrm{in}} + P_{\mathrm{b}}^{\mathrm{out}} + P_{\mathrm{fc}} - P_{\mathrm{ey}} = 0. \tag{3}$$

Figure 2 shows the relationship between the input power of the electrolyzer (P_{ey}) and the hydrogen produced (w_{ey}). Figure 3 shows the function between the hydrogen entering the fuel cell (w_{fc}) and the electrical power generated (P_{fc}). The maximum efficiencies for the electrolyzer and the fuel cell are 56% and 53%, respectively. These efficiencies were measured to account for all losses.

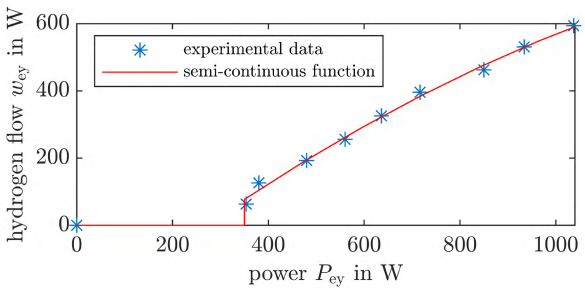

Figure 2. Fitted experimental data for semi-continuous non-linear electrolyzer function linking power (P_{ey}) in W and energy equivalent hydrogen flow (w_{ey}) in W.

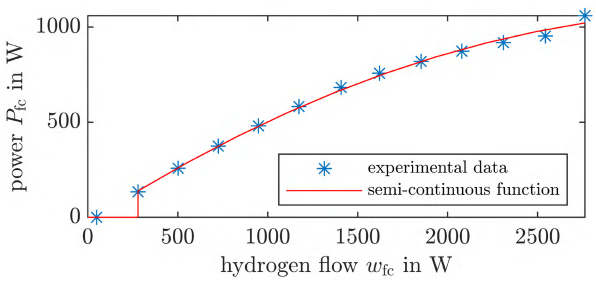

Figure 3. Fitted experimental data for semi-continuous non-linear fuel cell function linking power (P_{fc}) in W and energy equivalent hydrogen flow (w_{fc}) in W.

Parameters & Data

The energy system is specified based on the component sizes, as shown in Table 1. These are hereinafter referred to as the design point.

Table 1. TU Wien's Hybrid Energy System - Properties

Battery System		Hydrogen System		
Property	Value	Property	Value	Unit
C_b	2.6	C_{H_2}	14.3	kWh
P_b^{max}	500	P_{ey}^{max}	1040	W
		P_{fc}^{max}	1020	W
soc_b^{min}	0.2	$soc_{H_2}^{min}$	0.1	-
soc_b^{max}	0.8	$soc_{H_2}^{max}$	0.9	-
η_b	0.9	η_{H_2}	1	-

The optimization goal of providing storage flexibility is tested using five representative weeks. These are obtained as the best approximation of the entire year and are listed in Table 2. Actual photovoltaic and wind production data, as well as energy consumption data, from Austria for 2023 are used [7].

It was rescaled that on average over the year, this results in a generation of $\overline{P_{gen}}$ = 1000 W and a demand of $\overline{P_{de}}$ = 750 W, i.e., a surplus of $\overline{P_{gen}} - \overline{P_{de}}$ = 250 W of energy that is fed into the system.

Table 2. Representative Periods [7]

Name	Time period (2023)	Residual $\overline{P_{gen}} - \overline{P_{de}}$	Factor
Week 1	13.03-19.03	344 W	9.03
Week 2	19.06-25.06	679 W	8.16
Week 3	02.10-08.10	425 W	14.50
Week 4	23.10-29.10	37 W	10.83
Week 5	27.11-03.12	-290 W	9.48
average	01.01-31.12	250 W	52

Control Engineering

This chapter is based on [4]. Model predictive control (MPC) utilizes a design model (MILP formulations) to predict future system behavior based on the current state and a sequence of future control inputs and disturbance predictions over a prediction and optimization horizon N_{pred}, as shown in Figure 4.

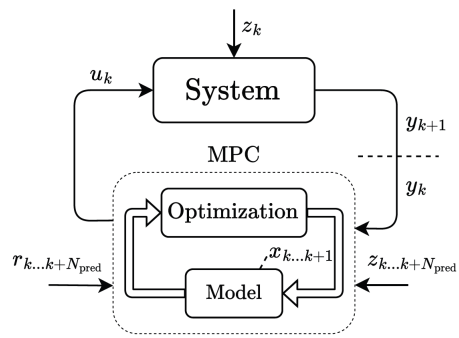

Figure 4. Model predictive control loop with control variable u, system state x, system output y, disturbance z and reference value r for discrete time step k over prediction horizon N_{pred} adapted from [4, 6].

At each discrete time step k, an optimization problem minimizes a cost function J, in this case, the energy exchange with the grid, according to

$$\min. J = \min. \sum_{k}^{k+N_{pred}} \frac{\left(P_{gr,\,k}^{in} + P_{gr,\,k}^{out}\right)}{\left(\overline{P_{gen}} - \overline{P_{de}}\right)}. \tag{4}$$

This optimization is subject to system dynamics and input and state constraints, described as equality constraints $h(x,u) = 0$ and inequality constraints $g(x,u) \leq 0$. Here, the system output y_k includes the system state x_k, comprising the states-of-charge of the battery and hydrogen storage, as well as grid exchange, as seen in equations (5) and (6).

$$y_k = [x_k \quad P_{gr}^{in} \quad P_{gr}^{out}]_k^T \tag{5}$$

$$x_k = [soc_b \quad soc_{H_2}]_k^T \tag{6}$$

Only the first control action u_k is applied, and the process

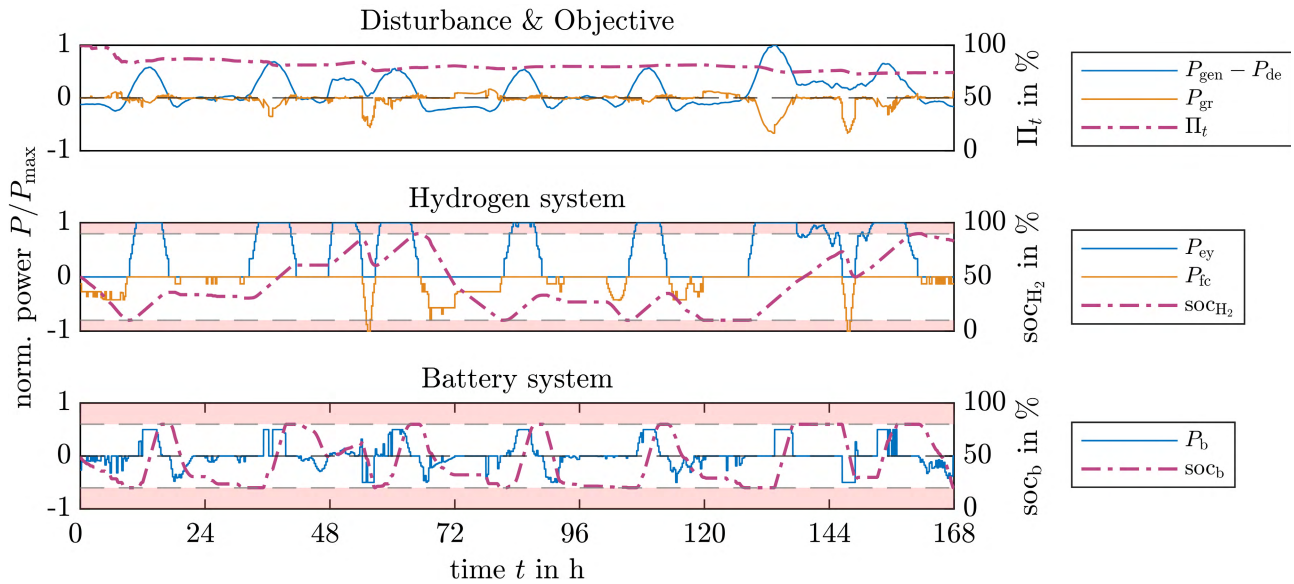

Figure 5. Exemplary process variables for week 3, 2-8.10.2023, with imposed disturbance ($P_{\mathrm{gen}} - P_{\mathrm{de}}$), setpoints of electrolyzer (P_{ey}), fuel cell (P_{fc}) and battery ($P_{\mathrm{b}}^{\mathrm{in}}, P_{\mathrm{b}}^{\mathrm{out}}$) power rates, states-of-charge ($soc_{\mathrm{H_2}}, soc_{\mathrm{b}}$) and resulting grid exchange ($P_{\mathrm{gr}}^{\mathrm{in}}, P_{\mathrm{gr}}^{\mathrm{out}}$) and performance factor (Π).

is repeated at the next step $k+1$, using updated measurements, known as receding horizon control.

$$u_k = [P_{\mathrm{b}}^{\mathrm{in}} \quad P_{\mathrm{b}}^{\mathrm{out}} \quad P_{\mathrm{ey}} \quad P_{\mathrm{fc}}]_k^T. \tag{7}$$

Energy generation and demand act as a disturbance on the system; the predictions can be written as

$$z_k = [P_{\mathrm{gen}} \quad P_{\mathrm{de}}]_{k \ldots k+N_{\mathrm{pred}}}^T. \tag{8}$$

MPC is well-suited for constrained, multivariable systems with predictable dynamics and disturbances. To evaluate the control quality, a dimensionless metric, the performance factor, is built, according to

$$\Pi_t = 1 - \frac{\sum_{\tau=0}^{\tau=t} P_{\mathrm{gr},\tau}^{\mathrm{in}} + P_{\mathrm{gr},\tau}^{\mathrm{out}}}{\sum_{\tau=0}^{\tau=t} |P_{\mathrm{gen},\tau} - P_{\mathrm{de},\tau}|}. \tag{9}$$

Since there is no explicit degradation modeling as a regularization, the ramping of electrolyzer and fuel cell power is limited, as shown in Equations (10) and (11). The maximum rate of change was chosen to be $\delta P_{\mathrm{ey,\,lim}} = P_{\mathrm{ey}}^{\max}/3600$ in W/s for the electrolyzer and $\delta P_{\mathrm{fc,\,lim}} = P_{\mathrm{fc}}^{\max}/3600$ in W/s for the fuel cell. Therefore, extensive power changes are restricted, thereby minimizing one part of degradation [8, 9].

$$\frac{|P_{\mathrm{ey},k+1} - P_{\mathrm{ey},k}|}{t_{k+1} - t_k} \leq \delta P_{\mathrm{ey,\,lim}} \tag{10}$$

$$\frac{|P_{\mathrm{fc},k+1} - P_{\mathrm{fc},k}|}{t_{k+1} - t_k} \leq \delta P_{\mathrm{fc,\,lim}} \tag{11}$$

RESULTS & DISCUSSION

To evaluate the dependence of individual component size variation on system performance, the capacities of the battery and hydrogen storage tank, as well as the maximum outputs of the electrolyzer and fuel cell, are varied relative to the design point (see Table 1) over a range of -40% to +40%. For this purpose, the terminal value (168 h) of the performance factor (see Equation (9)) is evaluated for the five representative weeks. The weighted sum is then calculated using the previously determined factors for the respective representative week (see Table 2) and used for subsequent evaluations.

Figure 5 illustrates, exemplarily, the simulated process variable resulting from the optimized control input and imposed disturbance for week 3 (see Table 2) at the design point. In the upper tile, the energy residual (blue), grid exchange (orange), and performance factor (magenta, dashed) are shown. The middle tile displays the hydrogen subsystem, with normalized electrolyzer power (blue), fuel cell power (orange), and storage state-of-charge (magenta, dashed).

Analogously, the lower tile displays the battery subsystem, with normalized charging and discharging power (blue) and state-of-charge (magenta, dashed).

The variation in battery and hydrogen storage capacities, C_{b} and $C_{\mathrm{H_2}}$, respectively, influences performance, resulting in a three-dimensional surface, shown in Figure 6. The performance factor varies between 56.8% and 68.5% over the considered parameter range, with the battery having a greater relative influence on the

result ($\bar{\Delta}$ = 7.5%) than the hydrogen storage system ($\bar{\Delta}$ = 4.4%). To answer the question of which component has the most limiting effect on performance, the variations around the design point for the maximum power rates, as well as the capacities at the design point, are plotted in Figure 7. The upper tile shows the performance, and the lower tile shows its partial derivative with respect to each component sizing, over the sizing ratio relative to the design point.

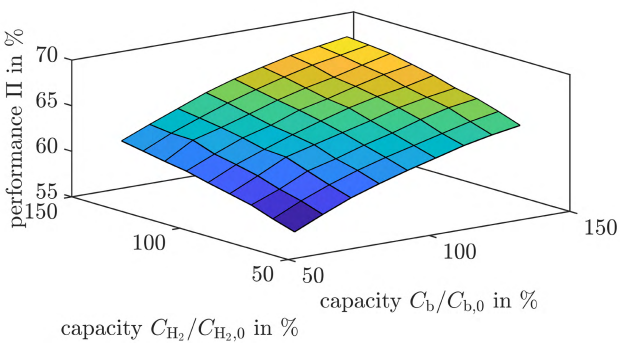

Figure 6. Performance value over battery and hydrogen storage capacity ratio variation in %.

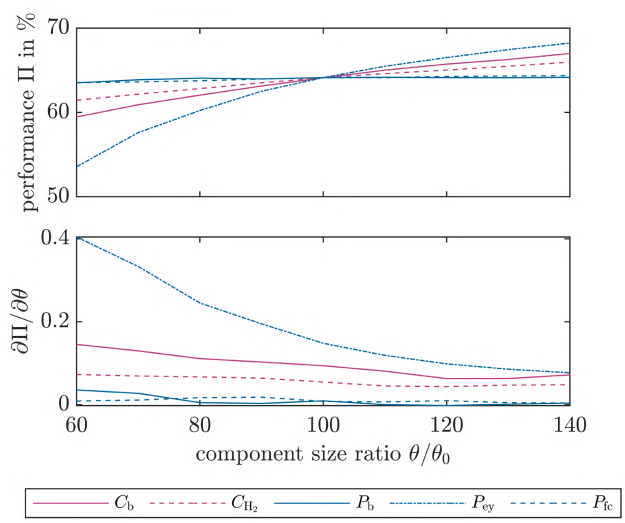

Figure 7. Upper tile: performance value; lower tile: partial derivative of performance value over component size ratio variation in %.

It can be seen that the electrolyzer has the greatest influence; this has a limiting effect, suggesting that it was undersized compared to the other storage components. This can also be explained by the fact that converting electrical energy into hydrogen and reconverting it to electricity are subject to losses due to low efficiency. Based on the determined efficiencies, an electrolyzer power input approximately four times the fuel cell power output would be appropriately dimensioned.

At maximum electrolyzer power, the capacities of the battery and hydrogen storage tank have the greatest influence on control quality. The maximum power rates of the battery and fuel cell have no influence on the performance factor.

CONCLUSION & OUTLOOK

In this study, the performance sensitivity with respect to component sizing of a model predictive control strategy designed to minimize grid exchange and thereby increase self-sufficiency was evaluated. To this end, a simulation and optimization model based on mixed-integer linear programming formulations for a lab-scale hybrid battery-hydrogen energy storage system was developed. Here, experimental data were used to fit non-linear semi-continuous functions for the electrolyser and fuel cell. The variation in battery and hydrogen storage capacity, as well as the maximum power rates of the electrolyser, fuel cell, and battery, was evaluated with respect to control performance.

The results indicate that for the given lab-scale system, the electrolyzer's maximum power has the greatest influence, followed by the battery and hydrogen storage capacity. The power of the fuel cell and battery does not affect the controller's performance. It can be concluded that the electrolyzer is undersized, and the battery should have greater storage capacity to improve performance. In general, the electrolyzer power input should be approximately 4 times the fuel cell's power output, accounting for conversion losses.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the support provided by the doctoral school Smart Industrial Concept (https://www.tuwien.at/doc/sic) and the TU Wien Library for financial support through its Open Access Funding Programme.

AUTHOR IDENTIFIERS

Author ORCIDs:
Moritz Zebenholzer: 0009-0000-1753-2167
Paul Schwarzmayr: 0000-0003-4735-3447
René Hofmann: 0000-0001-6580-4913

REFERENCES

1. Van LP, Chi KD, Duc TN. Review of hydrogen technologies based microgrid: energy management systems, challenges and future recommendations. International Journal of Hydrogen Energy 48:14127-14148 (2023). https://doi.org/10.1016/j.ijhydene.2022.12.345
2. Le TS, Nguyen TN, Bui DK, Ngo TD. Optimal sizing

of renewable energy storage: a techno-economic analysis of hydrogen, battery and hybrid systems considering degradation and seasonal storage. Applied Energy 336:120817 (2023). https://doi.org/10.1016/j.apenergy.2023.120817

3. E. Eriksson, E. Gray, Optimization and integration of hybrid renewable energy hydrogen fuel cell energy systems – A critical review, Applied Energy 202 (2017) 348–364. doi:10.1016/j.apenergy.2017.03.132. URL https://linkinghub.elsevier.com/retrieve/pii/S0306261917306256
4. C. Bordons, F. Garcia-Torres, M. A. Ridao, Model Predictive Control of Microgrids, Advances in Industrial Control, Springer International Publishing, Cham, 2020. doi:10.1007/978-3-030-24570-2. URL http://link.springer.com/10.1007/978-3-030-4570-2
5. A. Holtwerth, A. Xhonneux, D. Müller, Closed loop model predictive control of a hybrid battery-hydrogen energy storage system using mixed-integer linear programming, Energy Conversion and Management: X 22 (2024) 100561. doi:10.1016/j.ecmx.2024.100561. URL https://linkinghub.elsevier.com/retrieve/pii/S2590174524000394
6. Zebenholzer M, Kasper L, Schirrer A, Hofmann R. Optimal energy scheduling for battery and hydrogen storage systems using reinforcement learning. Systems and Control Transactions 4:1201-1207 (2025). https://doi.org/10.69997/sct.134052
7. Entso-e transparency platform, https://transparency.entsoe.eu/, accessed: 2025-10-14 (2025).
8. S. M. Alia, S. Stariha, R. L. Borup, Electrolyzer Durability at Low Catalyst Loading and with Dynamic Operation, Journal of The Electrochemical Society 166 (15) (2019) F1164–F1172. doi:10.1149/2.0231915jes. URL https://iopscience.iop.org/article/10.1149/2.0231915jesN.
9. P. Ren, P. Pei, Y. Li, Z. Wu, D. Chen, S. Huang, Degradation mechanisms of proton exchange membrane fuel cell under typical automotive operating conditions, Progress in Energy and Combustion Science 80 (2020) 100859. doi:10.1016/j.pecs.2020.100859. URL https://linkinghub.elsevier.com/retrieve/pii/S0360128520300691

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A novel decomposition-based approach to solve heterogeneous capacitated vehicle routing problems

Vakil Vamsi Krishna[a], Mangesh Kapadi[b], Pankaj Verma[b], Shamik Misra[a]*

[a] Department of Chemical engineering, Indian Institute of Technology Tirupati, Andhra Pradesh, India – 517619
[b] Linde South Asia Services Limited, Pune, Maharashtra, India - 411045
* Corresponding Author: misra@iittp.ac.in.

ABSTRACT

The Heterogeneous Capacitated Vehicle Routing Problem (HCVRP) is a fundamental extension of the classical Vehicle Routing Problem in which customer demands must be satisfied using a fleet of vehicles with varying capacities and costs. In this paper, a novel and intuitive decomposition-based formulation for HCVRP is presented that decomposes the problem into two tractable sub-problems: (i) a route generation and an optimal customer sequencing problem and (ii) a vehicle route assignment problem. In the first stage, all feasible customer combinations are constructed as routes, and for each route an optimization problem is solved to identify the optimal customer sequence that results in the minimum distance travelled. In the second stage, the optimal routes are selected, and vehicles are assigned using a mixed integer linear programming (MILP) formulation that minimizes the fixed cost of vehicle utilisation and total transportation costs, ensuring demand satisfaction for all customers while meeting vehicle capacity limitations. Computational results demonstrate that the decomposition strategy preserves optimality while significantly reducing solution time.

Keywords: Vehicle Routing, Mixed integer linear programming, Decomposition, Optimization

INTRODUCTION

Vehicle routing problem (VRP), introduced by Dantzig and Ramser (1959), determines optimal routes for delivering products from a depot to multiple locations with minimum cost [1]. In classical VRP, a set of homogeneous vehicles is supposed to start from the depot, visit the customer nodes once to fulfil their demand and come back to the depot [2]. The total products to be delivered are constrained by the vehicle capacity, generating the capacitated vehicle routing problem (CVRP) [3]. A more practical extension of this problem is where customers are served by a heterogeneous set of vehicles with different capacities and costs. In these heterogeneous capacitated vehicle routing problems (HCVRP), the fleet composition also becomes a decision variable [4]. The schematic for the HCVRP is presented in Figure 1. Since its introduction in Golden et al. (1984) [5], HCVRP is one of the most studied problems in literature, with multiple variants involving different complexities[6], [7], [8], [9], [10].

The contributions on VRP and its variants in the last two decades are surveyed in Golden et al. (2023)[11] and Archetti et al. (2026)[12]. The solution method used to solve VRP can mostly be segregated into three types: 1) Exact, 2) heuristics-assisted, and 3) metaheuristic methods[13]. With the advancements in AI and ML, recently, large-scale VRPs have been solved by incorporating AI and ML techniques into the solution methods[14], [15], [16].

Despite HCVRP being a well-known problem, due to its NP-hard nature, finding an optimal solution remains a challenge, especially with exact methods. Yaman (2006) [17] proposed four formulations for the HCVRP based on the subtour elimination constraints of Miller, Tucker, and Zemlin (1960) [18]. Later, Baldacci et al. (2009) proposed some valid inequalities that strengthen the linear relaxation of the MILP formulation[19]. Pessoa et al. (2008)[20] proposed an improved branch-cut-and-price model (BCP) based on the formulation presented by Fukasawa et al. (2006)[21]. Recently, several improvements on the BCP framework have been proposed that can solve

large-scale HCVRP instances[22], [23], [24], [25]. However, the design and implementation of such a complex and sophisticated algorithm require substantial effort, limiting the application to real-world problems, which sometimes deviate from classical HCVRPs discussed in the literature[22], [26].

To address these problems, this paper presents a novel decomposition-based formulation for the HCVRP. The proposed formulation is inspired by the set-partitioning approach of Balinski and Quandt (1964)[27]. The key contribution of our work is to decompose the HCVRP into two problems: 1) A route generation, invalidation and optimal customer sequencing problem, and 2) Vehicle and route assignment problem. At first, based on the customer locations, all possible route combinations are generated. Each route starts and ends at the same depot. For each customer combination, a mathematical optimization problem is solved to find the minimum distance option and the corresponding customer sequence.

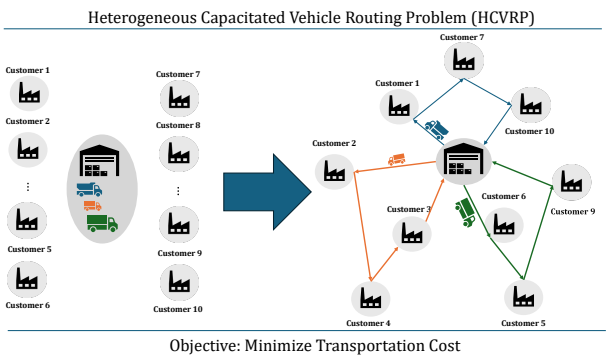

Figure 1. General Depiction of HCVRP

Once the optimal customer sequence for each customer combination is found, these routes are sent to the heterogeneous vehicle route assignment MILP problem.

In the vehicle assignment problem, each vehicle is assigned to go only on one route in such a way that all the customer demands on that route can be met. Constrained by the capacity of the vehicle, multiple vehicles might be required to cater to the demands of all customers. The objective function of the assignment MILP problem minimizes the fixed cost of using the vehicles along with the overall transportation cost.

The rest of the paper is organised as follows: in Section 2, the HCVRP problem is described, and a traditional arc-based formulation with MTZ subtour elimination constraints, taken from the literature, is presented. The novel decomposition-based HCVRP model is described in detail in Section 3, followed by the computational results and a discussion about the efficacy of the proposed model in Section 4. Section 5 summarizes the results and key takeaways.

PROBLEM DESCRIPTION

An HCVRP problem consists of $l \in L$ locations in which location L^0 represents the depot and other locations $l \in L^C$ denotes the customer locations. Each customer $l \in L^C$ has a known demand of De_l. The demand of depot is considered to be 0. The distance between each pair of locations ($d_{l,l'}$) or the transportation cost ($c_{l,l'}$) are given. Multiple vehicles ($v \in V$) with heterogeneous capacity (VQ_v^{max}) are used to travel from the depot to customer locations to cater to the demands. The objective is to find the optimal path to serve the customers with minimal cost.

The traditional mathematical model to solve the abovementioned HCVRP problem is taken from Yaman (2006), which uses Miller-Tucker-Zemlin (MTZ) based sub-tour elimination constraints. We define $yAl_{l,l',v}$ be the binary variable that assumes a value of 1 if the vehicle v has travelled from the location l to l' or 0 otherwise. All vehicles need not be used, and whether a vehicle is used for delivery or not is captured through a binary variable yS_v. We define zP_l be the variable for the position or order of the stops in that location.

The objective function minimizes the cost of transportation, along with the fixed cost associated with vehicle use and is presented in Equation 1.

$$\min xTC = \sum_{l \in L} \sum_{l' \in L} \sum_{v \in V} c_{l,l'} * yAl_{l,l',v} + \sum_{v \in V} FC_v * yS_v \quad (1)$$

$$\sum_{l \in L, l \neq l'} yAl_{l,l',v} = \sum_{l \in L, l \neq l'} yAl_{l',l,v} \ \forall l' \in L^C, \forall v \in V \quad (2)$$

$$\sum_{l \in L, v \in V, l \neq l'} yAl_{l,l',v} = 1 \ \forall l' \in L^C \quad (3)$$

$$\sum_{l' \in L^C} yAl_{l,l',v} \leq 1 \ \forall v \in V, l \in L^0 \quad (4)$$

$$\sum_{l' \in L^C} yAl_{l',l,v} \leq 1, \forall v \in V, \forall l' \in L^C \quad (5)$$

$$\sum_{l \in L^C} De_l * \sum_{l \in L^C, l \neq l'} yAl_{l,l',v} \leq VQ_v^{max}, \forall v \in V \quad (6)$$

$$\sum_{l \in L^0} \sum_{l' \in L^C, l \neq l'} yAl_{l',l,v} \leq yS_v, \forall v \in V \quad (7)$$

$$zP_l - zP_{l'} + VQ_v^{max} yAl_{l,l',v} \leq VQ_v^{max} - De_l$$
$$\forall l \in L^C, l' \in L^C l, l \neq l', v \in V \quad (8)$$

The equation (2) is the flow conservation equation, stating if a vehicle enters a node, it must leave, equation (3) guarantees that every customer is visited by only once and by only one vehicle, equation (4) ensures that each vehicle starts at most one route, equation (5) ensures that each vehicle ends at depot, equation (6) links the vehicle to the route stating that vehicle is only used if and only if it is selected, equation (7) defines the vehicle capacity constraint stating total demand served by the

vehicle does not exceeds its capacity and equation (8) eliminates subtours by the use of MTZ subtour elimination constraints.

```
Input: Locations, Distance, Demand
Output: R {Route List with optimal Sequence}
1.  Route Generation Stage:
Begin R ← [] {empty route set}
Initialize NCTC = 0
CTC_{NCTC,p} ← [] {empty customer trip
combination set}
for l in 1 to |L^c|
    for p in 1 to l
        Initialize Customer Position (p) = p
    end for
    repeat
        NCTC = NCTC + 1
        for p in 1 to l
          CTC_{NCTC,p} ← l_p
        end for
        p ← l
        while (p > 0 and pos(l_p) = p + |L^c| - l)
              p = p − 1
        end-do
        if p = 0
              break  (All combinations generated)
        endif
        pos(l_p) = pos(l_p) + 1
        for (p' in (p+1) to l)
            pos(l_p') = pos(l_{p'-1}) + 1
        end for
    until false
end for
for r in 1 to NCTC
    R ← r
    n_r = 1
    L_{r,n_r} = 'depot'
    for l in 1 to |L^c|
        n_r = n_r + 1
        L_{r,n_r} ← CTC_{r,l}
end for
    n_r = n_r + 1
    L_{r,n_r} = 'depot'
End for

2. Route Invalidation Logic:
for r ∈ R : n_r > 3
    if total demand exceeds max vehicle capacity,
discard that route
        R ← R\r
    end if
end for

3. Route Sequencing Stage:
for r ∈ R : n_r > 3
        call route sequencing Optimization
Algorithm
        find optimal sequence for the customer
combinations
        Calculate minimum distance for the
customer combinations
        Assign it as Route Distance (D_r)
end for
```

Find the rightmost index that can be incremented

Update the Customer Position ($pos(l_p)$) to the right of p

Create all routes and make first and last position of Route R as depot

Figure 2. Route Generation, Invalidation and Sequencing Algorithm

PROPOSED FRAMEWORK

The proposed framework consists of three stages: 1) Route generation and invalidation, 2) optimal customer sequencing stage and 3) optimal route selection and vehicle allocation.

In the first stage, based on the location details of depot(L^0) and the customers ($l \in L^C$), the pairwise distances ($d_{l,l'}$), all possible combinations of customers are generated. The algorithm for the same is described in Figure 2.

The algorithm creates a list of feasible routes by exploring all possible combinations of the customers with the depot. These combinations are created using nested loops in which the outer loop determines the size of the combination and the inner loop iterates from right to left, checking for an increment in the right-most position. If it is not possible, it moves to the left, resetting the right position in a lexicographical manner. It needs to be noted that at this stage, all possible combinations regarding the sequence of customer locations are explored and stored as routes, with the default starting and ending point for each route being the depot.

It should be noted that till this point, routes are essentially a set of locations (combining the depot and customer locations) without the sequence in which these locations can be travelled. Once the routes are generated, using the customer demands and maximum vehicle capacity, a route invalidation logic can be implemented, which limits the maximum number of customer locations that can be travelled, essentially eliminating the routes in which the sum of customer demands exceeds the maximum vehicle capacity available to the depot.

It should be noted that the route invalidation logic depends on customer demand and vehicle capacity. Consequently, although this logic significantly reduces the computational time and memory requirements of the proposed algorithm, it necessitates solving the algorithm separately for each demand scenario. In contrast, excluding the invalidation logic may increase memory usage or execution time; however, in this case, the algorithm's outcome becomes independent of demand scenarios. As long as the customer set remains unchanged (i.e., no customer is added or removed), the route generation and sequencing algorithm needs to be executed only once.

In the route sequencing stage, an optimisation problem is solved to find the optimum sequence in which the distances travelled to cover all customer locations on that route are minimum. We define $y_{l,l'}$, a binary variable, that assumes the value 1 if the vehicle travels directly from l to l', and to be 0 otherwise, and zP_l be an integer variable that indicates the position of the location l on that route. It should be noted that the following distance minimization optimization problem is solved for each route r and the corresponding optimum distance is stored in the variable $xDist_r$.

$$\min\ xDist_r = \sum_{l=1}^{n_r} \sum_{l'=1:(l\neq l')}^{n_r} D_{ll'} * y_{l,l'} \tag{9}$$

$$\sum_{l'=1:(l\neq l')}^{n_r} y_{l,l'} = 1, \forall l \in \{1, \dots, n_r - 1\} \tag{10}$$

$$\sum_{l'=1}^{n_r-1} y_{n_r,l'} = 0 \tag{11}$$

$$\sum_{l=2}^{n_r} y_{l,1} = 0 \tag{12}$$

$$\sum_{l=1:(l\neq l')}^{n_r} y_{l,l'} = 1, \forall l' \in \{2, \dots, n_r\} \tag{13}$$

$$zP_{n_r} = n_r \tag{14}$$

$$zP_1 = 1 \tag{15}$$

$$zP_{l'} - zP_l - y_{l,l'} - (n_r - 2)(1 - y_{l,l'}) \leq 0 \quad \forall l, l' \in 1, \dots, n_r \text{ and } l \neq l' \tag{16}$$

Where equation (9) is the objective function which minimizes the distance, equation (10) and equation (11) enforces exactly one leaving arc from each node except the terminal node, equation (12) ensures no incoming arc to the depot, whereas equation (13) ensure only one outgoing arc from depot. Equation (14) and (15) calculates the position of the depot as starting node and terminal node in the route. Equation (16) essentially determines the positions of other nodes and eliminates subtours using big-M constraint. The proposed mathematical model provides the optimal sequence of customers for each route and the corresponding minimum distance of the entire route. Customer positions are stored and sent to the third stage, in which an optimal route will be chosen for each assigned vehicle that minimises the overall distribution cost while meeting the demands of all customers.

For finding the optimal route and the corresponding vehicle allocation an efficient mixed integer linear programming (MILP) approach is proposed. Here, we define $yVR_{v,r}$ be a binary variable, and it takes the value of 1 if vehicle v is assigned to route r and assumes 0 otherwise. Another binary variable $yVSl_v$ is defined to denote whether the vehicle v is utilized or not. $xVDL_{v,l}$ captures the amount delivered to the customer(l) by vehicle v and $xVTS_v$ denotes the total amount supplied by the vehicle v, the total route distance as calculated in the route sequencing stage($xDist_r$) are captured as an input parameter D_r, and the fixed cost of using a vehicle is denoted as FC_v and the total cost be xTC. The objective function to be minimised is as follows:

$$\min xTC = \sum_{r\in R, v\in V} D_r * yVR_{v,r} + \sum_{v\in V} FC_v * yVSl_v \tag{17}$$

Subject to

$$\sum_{r\in R} yVR_{v,r} \leq yVSl_v\ \forall v \in V \tag{18}$$

$$xVTS_v \leq VQ_v^{max} * yVSl_v\ \forall v \in V \tag{19}$$

$$xVTS_v \leq VQ_v^{max} * \sum_{r\in R} y_{v,r}\ \forall v \in V \tag{20}$$

$$xVTS_v = \sum_{l\in L^C} xVDL_{v,l}\ \forall v \in V \tag{21}$$

$$xVDL_{v,l} \leq De_l * \sum_{r\in R_l} y_{v,r}\ \forall v \in V, \forall l \in L^c \tag{22}$$

$$De_l = \sum_{v\in V} xVDL_{v,l}\ \forall l \in L^c \tag{23}$$

Where equation(18) ensures that the vehicle is only assigned to a maximum of one route if it is selected, equation(19) ensures that the vehicle carries a load only if it is selected, and equation(20) defines that the total load carried by the vehicle equals the total demand in that route. Equation(21) equates the total load of a vehicle to the combined demand of customers it caters to. Equation(22) ensures that the vehicle delivers only to the customers belonging to the assigned route, and Equation(23) ensures full demand satisfaction of the customer locations. It should be noted that Equation (23) facilitates that the demand of one customer can be quenched by multiple vehicles.

RESULT AND DISCUSSIONS

To demonstrate the efficacy of the abovementioned approach, we designed 9 instances of HCVRP problems spanning between 5 and 15 customer nodes with 2-3 heterogeneous vehicles. We also include 3 instances from CVRP library[28] and modify them as HCVRP instances. The computational efficiency of the proposed model is demonstrated by comparing it with the HCVRP model proposed by Yaman (2006), which uses MTZ subtour elimination routines. All instances are solved in Fico Xpress Optimization Suite[29] on an Intel®Core™ i9-14900K (3.7 GHz) machine with 64 GB RAM. The termination criteria were set to either 0% Optimality Gap or 3600 s, whichever is attained early. The problem characteristics and the result comparison are shown in Table 1. In the case of the smaller instances 1-5, each model attains the same solution (Figure 3) with 0% optimality gap, proving that the proposed decomposition strategy does not affect the optimal solution. In the large-scale instances 6-9, a significant reduction in computational time is observed in the proposed model, substantiating the computational efficiency.

Table 1: Model Performance and Characteristics

			Yaman (2006)			Proposed Model				
Instance	Customer Nodes	Vehicles	Best Known Solution	Optimality Gap (%)	Time (s)	Best Known Solution	Optimality Gap (%)	Time (s)		
								RE	VR	Total
1	5	2	1172.46	0	0.12	1172.46	0	1.50	0.01	1.51
2	6	2	1177.52	0	0.39	1177.52	0	3.2	0.01	3.21
3	7	2	1177.55	0	7.3	1177.55	0	7.9	0.02	7.92
4	8	2	1736.72	0	341.4	1736.72	0	16.4	0.02	16.42
5	9	3	1629.92	76.5	3600	1590.75	0	28.5	0.4	28.9
6	10	3	1814.05	78.9	3600	1795.77	0	52.1	0.5	52.6
7	12	3	2043.47	82.5	3600	2014.19	0	269	2.2	271.2
8	14	3	2599.17	86.3	3600	2586.5	0	1891	21	1912
9	15	3	2682.81	86.5	3600	2606.31	0	3402	38	3440
10	15	8	440.16	40.32	3600	411.83	0	12	6	18
11	20	4	372	21.10	3600	358	0	1124	5.2	1129.2
12	21	4	390	21.6516	3600	373	0	1424	10	1434

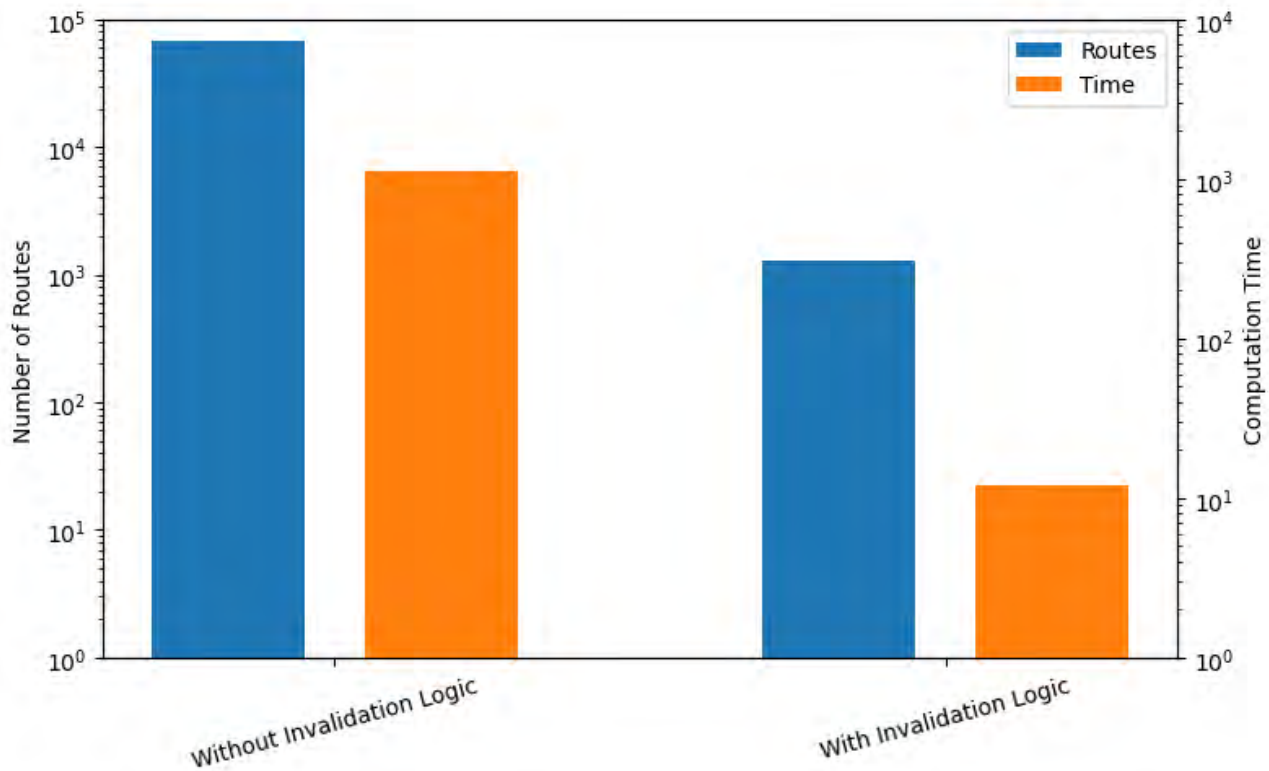

Figure 2. Impact of Invalidation Logic on Resultant Routes and Computational Time

It should also be noted that Instances 1-9 are solved without using the route invalidation logic mentioned in Figure 2. The majority of the computational time is devoted to the route enumeration algorithm, as can be seen from Table 1. In a real-world situation, this route enumeration algorithm only needs to be solved if a new customer enters the supply chain or leaves the supply chain. Compared to this, the vehicle route assignment problems take minimal time and can be solved daily based on the customer demand.

However, with increasing customer locations (>20), the route generation and sequencing step may require a significant amount of memory and time, which can be reduced by implementing the route invalidation logic. To test the efficacy of the proposed algorithm combined with route invalidation logic, 3 instances from the CVRP

library are taken and suitably modified to create HCVRP instances. Figure 3 demonstrates the efficacy of route invalidation logic on the reduction of route generation and sequencing efforts when applied to Instance 10. The other two instances taken from CVRP Library are also solved using the invalidation logic and are solved to optimality in a reasonable time.

CONCLUSION

This paper presented a novel decomposition-based formulation for the heterogeneous capacitated vehicle routing problem aimed at improving computational efficiency while preserving optimality. By separating the original NP-hard problem into two intuitive and modular subproblems: 1) route generation and optimal customer sequencing, and 2) optimal route selection and vehicle assignment problem, the proposed approach alleviates the combinatorial burden typically encountered in monolithic MILP formulations. Unlike sophisticated exact methods such as branch-cut-and-price algorithms, the proposed framework is conceptually simple and requires substantially less modelling and implementation effort, making it more suitable for practical decision-support applications.

Furthermore, due to the simplicity of the proposed decomposition-based formulation, it can be easily extended to incorporate the complexities of HCVRP, such as time windows, split deliveries, multiple depots, etc. One major limitation of the proposed algorithm, however, is the memory- and time-intensive route-generation stage, which renders the problem intractable as the number of locations grows (>25), thereby motivating future research to address this challenge.

AUTHOR IDENTIFIERS

Author ORCIDs:
Krishna VV: 0009-0008-7500-4723
Misra S: 0000-0002-1684-4174

REFERENCES

1. Dantzig GB, Ramser JH. The truck dispatching problem. Management Science 6:80-91 (1959). https://doi.org/10.1287/mnsc.6.1.80
2. Baldacci R, Toth P, Vigo D. Exact algorithms for routing problems under vehicle capacity constraints. Ann Oper Res 175:213-245 (2009). https://doi.org/10.1007/s10479-009-0650-0
3. Fitzpatrick J, Ajwani D, Carroll P. A scalable learning approach for the capacitated vehicle routing problem. Computers & Operations Research 171:106787 (2024). https://doi.org/10.1016/j.cor.2024.106787
4. Koç Ç, Bektaş T, Jabali O, Laporte G. Thirty years of heterogeneous vehicle routing. European Journal of Operational Research 249:1-21 (2016). https://doi.org/10.1016/j.ejor.2015.07.020
5. Golden B, Assad A, Levy L, Gheysens F. The fleet size and mix vehicle routing problem. Computers & Operations Research 11:49-66 (1984). https://doi.org/10.1016/0305-0548(84)90007-8
6. Baldacci R, Toth P, Vigo D. Exact algorithms for routing problems under vehicle capacity constraints. Ann Oper Res 175:213-245 (2009). https://doi.org/10.1007/s10479-009-0650-0
7. Irnich S, Schneider M, Vigo D. Chapter 9: four variants of the vehicle routing problem. Vehicle Routing :241-271 (2014). https://doi.org/10.1137/1.9781611973594.ch9
8. Hoff A, Andersson H, Christiansen M, Hasle G, Løkketangen A. Industrial aspects and literature survey: fleet composition and routing. Computers & Operations Research 37:2041-2061 (2010). https://doi.org/10.1016/j.cor.2010.03.015
9. Golden B, Assad A, Levy L, Gheysens F. The fleet size and mix vehicle routing problem. Computers & Operations Research 11:49-66 (1984). https://doi.org/10.1016/0305-0548(84)90007-8
10. Laporte G. Fifty years of vehicle routing. Transportation Science 43:408-416 (2009). https://doi.org/10.1287/trsc.1090.0301
11. Golden B, Wang X, Wasil E. The evolution of the vehicle routing problem—a survey of VRP research and practice from 2005 to 2022. Synthesis Lectures on Operations Research and Applications :1-64 (2023). https://doi.org/10.1007/978-3-031-18716-2_1
12. Archetti C, Coelho LC, Speranza MG, Vansteenwegen P. Beyond fifty years of vehicle routing: insights into the history and the future. European Journal of Operational Research 330:355-372 (2026). https://doi.org/10.1016/j.ejor.2025.06.014
13. El-Sherbeny NA. Vehicle routing with time windows: an overview of exact, heuristic and metaheuristic methods. Journal of King Saud University - Science 22:123-131 (2010). https://doi.org/10.1016/j.jksus.2010.03.002
14. Hildebrandt FD, Thomas BW, Ulmer MW. Opportunities for reinforcement learning in stochastic dynamic vehicle routing. Computers & Operations Research 150:106071 (2023). https://doi.org/10.1016/j.cor.2022.106071
15. Morabit M, Desaulniers G, Lodi A. Learning to repeatedly solve routing problems. Networks 83:503-526 (2023). https://doi.org/10.1002/net.22200
16. Guan Q, Cao H, Tan J, Ou B, Wang Y, Yan X, Chen

B. Dynamic-aware deep reinforcement learning for heterogeneous fleet-based capacitated electric vehicle routing problems. Applied Energy 413:127757 (2026). https://doi.org/10.1016/j.apenergy.2026.127757
17. Yaman H. Formulations and valid inequalities for the heterogeneous vehicle routing problem. Math. Program. 106:365-390 (2005). https://doi.org/10.1007/s10107-005-0611-6
18. Miller CE, Tucker AW, Zemlin RA. Integer programming formulation of traveling salesman problems. J. ACM 7:326-329 (1960). https://doi.org/10.1145/321043.321046
19. Baldacci R, Battarra M, Vigo D. Valid inequalities for the fleet size and mix vehicle routing problem with fixed costs. Networks 54:178-189 (2009). https://doi.org/10.1002/net.20331
20. Pessoa A, Uchoa E, Poggi de Aragão M. A robust branch-cut-and-price algorithm for the heterogeneous fleet vehicle routing problem. Networks 54:167-177 (2009). https://doi.org/10.1002/net.20330
21. Fukasawa R, Longo H, Lysgaard J, Aragão MP, Reis M, Uchoa E, Werneck RF. Robust branch-and-cut-and-price for the capacitated vehicle routing problem. Math. Program. 106:491-511 (2005). https://doi.org/10.1007/s10107-005-0644-x
22. Pessoa A, Sadykov R, Uchoa E, Vanderbeck F. A generic exact solver for vehicle routing and related problems. Math. Program. 183:483-523 (2020). https://doi.org/10.1007/s10107-020-01523-z
23. Pessoa A, Sadykov R, Uchoa E. Enhanced branch-cut-and-price algorithm for heterogeneous fleet vehicle routing problems. European Journal of Operational Research 270:530-543 (2018). https://doi.org/10.1016/j.ejor.2018.04.009
24. Pecin D, Pessoa A, Poggi M, Uchoa E. Improved branch-cut-and-price for capacitated vehicle routing. Math. Prog. Comp. 9:61-100 (2016). https://doi.org/10.1007/s12532-016-0108-8
25. N. Errami, E. Queiroga, R. Sadykov, and E. Uchoa, “VRPSolverEasy: A Python Library for the Exact Solution of a Rich Vehicle Routing Problem, ” *INFORMS J. Comput.*, vol. 36, no. 4, 2024, doi: 10.1287/ijoc.2023.0103.
26. Fitzpatrick J, Ajwani D, Carroll P. A scalable learning approach for the capacitated vehicle routing problem. Computers & Operations Research 171:106787 (2024). https://doi.org/10.1016/j.cor.2024.106787
27. Balinski ML, Quandt RE. On an integer program for a delivery problem. Operations Research 12:300-304 (1964). https://doi.org/10.1287/opre.12.2.300
28. Uchoa E, Pecin D, Pessoa A, Poggi M, Vidal T, Subramanian A. New benchmark instances for the capacitated vehicle routing problem. European Journal of Operational Research 257:845-858 (2017). https://doi.org/10.1016/j.ejor.2016.08.012
29. C. Guéret, C. Prins, M. Sevaux, and S. Heipcke, *Applications of Optimization with Xpress-MP*. Dash Optimization Limited, 2002. [Online]. Available: https://books.google.co.in/books?id=2hDrPQAACAAJ

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Network Design Optimization for Biomethane Feed-in from Decentralized Production Sites

Lena Lessig[a]*, Martin Fischer[a], Paul Schwarzmayr[a] and René Hofmann[a]

[a] TU Wien, Institute of Energy Systems and Thermodynamics, Vienna, Austria
* Corresponding Author: lena.lessig@tuwien.ac.at.

ABSTRACT

Biomethane can serve as an alternative to fossil fuels. Pipeline and compressor infrastructure must be built to enable the feed-in of biomethane from decentralized production sites into the existing gas network. Costs can be reduced by sharing this transport infrastructure among individual biogas plant operators. To find the cost-optimal layout, a mixed integer linear programming (MILP) problem is formulated. The model is expanded to show the effect of considering not only pipe cost but also compressor cost in the objective. It is applied to an Austrian region to find the optimal network for 14 biogas plants. After optimization, two allocation methods for sharing investment costs among biogas plants are explored: an equal cost approach and a volumetric share approach. Optimal network topology changes due to the modified objective, resulting in significant cost savings of 9 % for the total region and 19 % for a subregion. Investment costs for individual biogas plant operators vary widely, depending on network configuration and chosen allocation method.

Keywords: Energy Systems, Biofuels, Optimization, MILP

INTRODUCTION

To mitigate climate change, anthropogenic greenhouse gas emissions must be drastically reduced. A great share of these emissions is caused by fossil fuels used in industries. While many sectors can be decarbonized through electrification, eliminating fuels is difficult in others. For example, many industrial processes require high temperature heat [1].

Biofuels can serve as an alternative to fossil fuels, particularly biomethane. Biomethane is produced by processing biogas that results from the fermentation of biogenic residues [2, 3].

In Austria, 74.37 TWh of gas were consumed in 2024. Around 124.52 GWh or 0.16 % of this consumption was covered through biomethane [4]. A law proposal to increase the quota of renewable gases failed in 2024 [5]. However, the current Austrian government has announced a new attempt to pass such a law in its government programme for 2025-2029. The target of 6.5 TWh of renewable gases per year should be reached by 2030 [6]. Currently, most of the biogas produced in Austria is used for on-site electricity and heat generation in combined heat and power (CHP) plants instead of being upgraded to biomethane. In 2024, almost 300 plants generated electricity from biogas [7]. These plants are located decentrally and are usually small. More than two-thirds of the mentioned biogas CHP plants have a maximum capacity lower than 500 kW_{el}, with the average maximum capacity of these being 164 kW_{el} [8]. In comparison, only 14 plants delivered biomethane to the gas grid [9]. To increase the quota of biomethane, investment in upgrading and transport infrastructure is necessary. Individual biogas plant operators can benefit from economies of scale when cooperating with each other along the biomethane supply chain [10].

A primary decision is how to find a cost-effective pipeline network topology to connect biogas plants and enable feed-in. Hengeveld et al. [11] found that biogas transport costs can be reduced when connecting plants using a fishbone layout. Recent publications presented mixed integer linear programming (MILP) models to determine network topology via mathematical optimization. These models are based on mathematical formulations for finding minimal spanning trees. Schröer et al. [12] provided a model to cluster decentralized biogas plants and connect them to central upgrading facilities and feed-in points while maximising profit. Again, this publication

found great potential for cost savings due to shared infrastructure compared to individual solutions. A recent publication of Blanco et al. [13] presented a MILP model for designing waste-to-biomethane logistics, covering large parts of the biomethane supply chain from waste collection to feed-in and biomethane liquefaction. However, none of these models include feed-in conditions in detail. When clustering projects are implemented in practice, feed-in capacities as well as minimal pressure levels of feed-in points must be met.

In this paper, a minimum spanning tree MILP formulation for biogas networks proposed by Huber et al. [14] was enhanced to include individual feed-in capacity limits and pressure requirements into the model. In this way, the cost of compression is considered in detail. The model is applied to a real use case in Austria. To demonstrate the effects of model extension, the optimization results of both the previous and enhanced models are compared.

MATERIAL AND METHODS

Use Case

The study examines clustering and connecting 14 biogas plants to the existing gas grid. 15 possible feed-in points are considered in the optimization. Locational data, production and capacity volumes, as well as pressure levels, correspond to existing plants and connection points in Austria. The 14 biogas plants are defined by their location and their biogas production volume, which lies between 83 and 402 Nm^3/h. The 15 feed-in points are characterized by their location, maximum feed-in capacity, and the required pressure level for feed-in. Feed-in capacities range from 291 to 2910 Nm^3/h. Pressure levels are defined by the operational pressure of the existing gas network and are 3.6, 8, 12.5, or 67 barg.

MILP Model

The MILP model is based on the network optimization formulation proposed by Huber et al. [14]. Constraints regarding feed-in point capacity limits and compressors were added.

The problem is formulated using incidence matrices introduced in Equation (1). The adjacency matrix C_{BP} indicates connections of biogas plants to each other. The matrix element $C_{\mathrm{BP},ii'}$ takes the value 1 if biogas plant i is directly connected to plant i' and biogas is transported from plant i to i'. The adjacency matrix C_{FP} indicates connections of biogas plants to feed-in points identified with index j. Equation (2) limits the self-connection of biogas plants.

d_{BP} and d_{FP} are matrices with spherical distances between locations of biogas plants or biogas plants and feed-in points. As the problem is highly combinatoric, Equations (3) and (4) improve computing time by excluding connections longer than the median distance between biogas plants $\tilde{d}_{\mathrm{BP}}$ – respectively the median distance between biogas plants and feed-in points $\tilde{d}_{\mathrm{FP}}$ – from the feasible solution space.

$$C_{\mathrm{BP},ii'}, C_{\mathrm{FP},ij} \in \{0,1\} : \forall\, i, i', j \tag{1}$$

$$C_{\mathrm{BP},ii} = 0 : \ \forall\, i \tag{2}$$

$$C_{\mathrm{BP},ii'} = 0 : \ \forall\, (i,i') \in \left\{(i,i') \mid d_{\mathrm{BP},ii'} > \tilde{d}_{\mathrm{BP}}\right\} \tag{3}$$

$$C_{\mathrm{FP},ij} = 0 : \ \forall\, (i,j) \in \left\{(i,j) \mid d_{\mathrm{FP},ij} > \tilde{d}_{\mathrm{FP}}\right\} \tag{4}$$

Binary edge matrix C_{E} describes how many edges are used to connect individual plants to feed-in points. For each plant i, the index e of the column containing a non-zero entry is the number of edges between the plant and the feed-in point. The construction of the edge matrix is given in Equation (6). The operator $\otimes$ is defined as a Boolean product of a matrix and a vector. Further description of this operation can be found in [14]. Equation (7) ensures that each biogas plant is connected to the grid.

$$C_{\mathrm{E},ie} \in \{0,1\} : \forall\, i, e \tag{5}$$

$$C_{\mathrm{E},ie} = \begin{cases} \sum_j C_{\mathrm{FP},ij} & \text{for } e = 1 \\ C_{\mathrm{BP}} \otimes C_{\mathrm{E},ie-1} & \text{for } e > 1 \end{cases} : \forall\, i \tag{6}$$

$$\sum_e C_{\mathrm{E},ie} = 1 : \forall\, i \tag{7}$$

Matrix $\dot{V}$ is a helper variable to calculate accumulated flows in the network from biogas plant production flows $\dot{V}_{\mathrm{BP},i}$. Equation (8) describes the construction of the helper variable. Vector $\dot{V}_{\mathrm{FP}}$ contains accumulated flows at feed-in points and is calculated by Equation (9). Compliance with capacity limits of feed-in points is ensured by Equation (10).

$$\dot{V}_{ie} = \begin{cases} \dot{V}_{\mathrm{BP},i} & \text{for } e = 1 \\ \dot{V}_{\mathrm{BP}}(C_{\mathrm{BP}} \otimes C_{\mathrm{E},ie-1}) & \text{for } e > 1 \end{cases} : \forall\, i \tag{8}$$

$$\dot{V}_{\mathrm{FP},j} = \sum_i (C_{\mathrm{FP},ij} \sum_e \dot{V}_{ie}) : \ \forall\, j \tag{9}$$

$$\dot{V}_{\mathrm{FP},j} \le \dot{V}_{\mathrm{max},j} : \ \forall\, j \tag{10}$$

Vector p_{out} is a binary variable that defines the chosen pressure category for each feed-in point that is used. According to Equation (12), only one pressure category k can be chosen for each feed-in point j. Equation (13) constrains the chosen compressor outlet pressure to be greater than the necessary pressure at feed-in points $p_{\mathrm{FP},j}$. The upper limit of each pressure category is denoted in parameter vector $\mathrm{p}_{\mathrm{out}}^{\mathrm{lim}}$ and given in the next subsection.

$$p_{\mathrm{out},jk} \in \{0,1\} : \ \forall\, j, k \tag{11}$$

$$\sum_k p_{\mathrm{out},jk} \le 1 : \ \forall\, j \tag{12}$$

$$\sum_k p_{\mathrm{out},jk}\, p_{\mathrm{out},k}^{\mathrm{lim}} \ge p_{\mathrm{FP},j} : \ \forall\, j \tag{13}$$

Annual investment and operational costs for pipes $C_{\mathrm{CAPEX}}^{\mathrm{pipe}}$ and $C_{\mathrm{OPEX}}^{\mathrm{pipe}}$ were calculated after Equations (14) and (15) with an annualization factor AF_{inv}. Parameters $c_{\mathrm{CAPEX}}^{\mathrm{pipe}}$ and $c_{\mathrm{OPEX}}^{\mathrm{pipe}}$ are given in the next subsection.

$$C_{\mathrm{CAPEX}}^{\mathrm{pipe}} = AF_{\mathrm{inv}}\, c_{\mathrm{CAPEX}}^{\mathrm{pipe}} \sum_i \left(\sum_{i'} C_{\mathrm{BP},ii'} d_{\mathrm{BP},ii'} + \sum_j C_{\mathrm{FP},ij} d_{\mathrm{FP},ij} \right) \tag{14}$$

$$C_{\mathrm{OPEX}}^{\mathrm{pipe}} = c_{\mathrm{OPEX}}^{\mathrm{pipe}} \sum_i \left(\sum_{i'} C_{\mathrm{BP},ii'} d_{\mathrm{BP},ii'} + \sum_j C_{\mathrm{FP},ij} d_{\mathrm{FP},ij} \right) \tag{15}$$

Compressor capital investment $C_{\mathrm{CAPEX}}^{\mathrm{comp}}$ is described by the exponential expression in Equation (16). Constant $c_{\mathrm{CAPEX}}^{\mathrm{comp}}$ and exponent $k_{\mathrm{CAPEX}}^{\mathrm{comp}}$ are given in the following subsection. This nonlinear equation was approximated through piecewise-linearization and implemented with SOS2. Operational costs $C_{\mathrm{OPEX}}^{\mathrm{comp}}$ of compressor units at feed-in points are calculated through Equation (17). The energy demand of pressure category k is $e_{\mathrm{comp},k}$. This and other parameters, namely the annual operating hours t_{op}, the electricity price c_{el}, and the share of investment cost for maintenance c_{m} are given in the next subsection.

$$C_{\mathrm{CAPEX},j}^{\mathrm{comp}} = AF_{\mathrm{inv}}\, c_{\mathrm{CAPEX}}^{\mathrm{comp}}\, \dot{V}_{\mathrm{FP},j}^{\;1-k_{\mathrm{CAPEX}}^{\mathrm{comp}}} \;:\; \forall\, j \tag{16}$$

$$C_{\mathrm{OPEX},j}^{\mathrm{comp}} = \sum_k \left(e_{\mathrm{comp},k}\, p_{\mathrm{out},jk} \right) \dot{V}_{\mathrm{FP},j}\, t_{\mathrm{op}}\, c_{\mathrm{el}} + c_{\mathrm{m}} C_{\mathrm{CAPEX},j}^{\mathrm{comp}} \;:\; \forall\, j \tag{17}$$

The optimization model is run with two different objectives, both minimizing total annual cost. The first objective in Equation (O1) only considers the cost for pipes and is therefore used to calculate the benchmark case where the pressure levels of feed-in points are not considered during optimization. The second objective in Equation (O2) includes the costs of compressors, exploiting the full potential of the enhanced model. By comparing the results of both optimization runs, the effect of the model enhancement can be evaluated.

$$\min\; C_{\mathrm{CAPEX}}^{\mathrm{pipe}} + C_{\mathrm{OPEX}}^{\mathrm{pipe}} \tag{O1}$$

$$\min\, C_{\mathrm{CAPEX}}^{\mathrm{pipe}} + C_{\mathrm{OPEX}}^{\mathrm{pipe}} + \sum_j C_{\mathrm{CAPEX},j}^{\mathrm{comp}} + \sum_j C_{\mathrm{OPEX},j}^{\mathrm{comp}} \tag{O2}$$

Parametrization

CAPEX and OPEX of pipes and compressors are annualized. Useful life is assumed to be 20 years, so the annualization factor AF_{inv} is 0.05 1/a.

Pipe costs in Table 1 are selected according to expert recommendation from the industrial project partner ÖVGW [15]. The investment and operational costs of compressors were modelled based on an internal report from the project partner ÖVGW [16]. According to this study, capital expenditures for compressors can be described by the nonlinear relation in Equation (16). The coefficient $c_{\mathrm{CAPEX}}^{\mathrm{comp}}$ and exponent $k_{\mathrm{CAPEX}}^{\mathrm{comp}}$ of the compressor CAPEX equation are dependent on compressor inlet pressure, which is assumed to be 0.1 barg for all compressors in the described use case. Therefore, coefficients are set to the values in Table 2.

Table 1: Pipe parameters [15].

Parameter	Value	Unit
$c_{\mathrm{CAPEX}}^{\mathrm{pipe}}$	275	EUR/m
$c_{\mathrm{OPEX}}^{\mathrm{pipe}}$	1	EUR/m/a

Table 2: Compressor parameters [16]. Electricity cost for non-households between 20 and 500 MWh/a in the first half of 2025 [17].

Parameter	Value	Unit
$c_{\mathrm{CAPEX}}^{\mathrm{comp}}$	83 688	EUR*h/Nm³
$k_{\mathrm{CAPEX}}^{\mathrm{comp}}$	0.648	-
t_{op}	8500	h/a
c_{el}	0.28	EUR/kWh
c_{m}	4	%

Equation (17) describes operational expenditures of compressors, including energy cost and maintenance costs. Energy cost depends on the electricity price and the amount of energy necessary for compression. Table 3 contains compressor pressure categories and the respective energy demand for compression. Electricity price and maintenance cost can be found in Table 2.

Table 3: Compressor outlet pressure categories and respective energy demand [16].

$p_{\mathrm{out}}^{\mathrm{lim}}$ in barg	e_{comp} in kWh/Nm³
3.6	0.1094
8	0.1450
12.5	0.1807
28	0.2369
70	0.2888

Cost Allocation

When clustering biogas plants, it is essential to find a fair method of allocating investment and operating costs of the shared infrastructure. Different methods have been proposed, e.g., equal cost for all cluster participants, sharing costs depending on the volumetric flow rate of participants, or on the distance to a feed-in point [18].

In this paper, the equal cost (EC) and volumetric share (VS) allocation method is applied after optimization to split CAPEX for pipes and compressors among participants in the clusters. Equation (18) describes how CAPEX are allocated to individual biogas plants using the EC approach. $C_{\mathrm{CAPEX},i}^{\mathrm{EC}}$ is the CAPEX allocated to biogas

plant i, while $C^{\text{pipe}}_{\text{CAPEX},Cl}$ and $C^{\text{comp}}_{\text{CAPEX},Cl}$ are total CAPEX of pipes and compressors of the cluster Cl that the respective plant is part of. $N_{BP,Cl}$ is the number of biogas plants belonging to the cluster. Analogously, Equation (19) presents CAPEX allocation using the VS method where $C^{\text{VS}}_{\text{CAPEX},i}$ is the CAPEX allocated to biogas plant i and $\dot{V}_{\text{FP},Cl}$ is the total flow at the feed-in point of the cluster.

$$C^{\text{EC}}_{\text{CAPEX},i} = (C^{\text{pipe}}_{\text{CAPEX},Cl} + C^{\text{comp}}_{\text{CAPEX},Cl})/N_{BP,Cl} \quad (18)$$

$$C^{\text{VS}}_{\text{CAPEX},i} = (C^{\text{pipe}}_{\text{CAPEX},Cl} + C^{\text{comp}}_{\text{CAPEX},Cl}) * \dot{V}_{\text{BP},i}/\dot{V}_{\text{FP},Cl} \quad (19)$$

RESULTS AND DISCUSSION

The MILP formulation is applied to the previously introduced Austrian use case. The optimization problem is run twice, using the different objectives in Equations (O1) and (O2). Figure 1 depicts the resulting networks, and the corresponding system costs are shown in Figure 2.

When only pipe costs are minimized with the objective (O1), 7 clusters are formed as depicted in Figure 1a. In contrast, minimizing pipe and compressor costs with objective (O2) results in 5 clusters as shown in Figure 1b. A cluster of three biogas plants in the northwest, and two single biogas plants without connections to other plants in the middle and in the southeast remain unchanged by the modified objective. The boxes in Figure 1a and Figure 1b are used to enlarge areas where the cluster configuration changes due to the change of the objective function: In the northern area (Region 1), two clusters are combined into one. The feed-in point of the eastern cluster continues to be used. In the southern area (Region 2), two clusters are united into one. It is especially notable that none of the feed-in points used when minimizing objective (O1) in the east are used. Instead, biogas plants are connected to a previously unused feed-in point in the west.

Figure 2 discloses the result of the modified objective on system cost. For the total area, specific system costs from optimization with objective (O1) are 14.53 EUR/MWh of biomethane. With objective (O2), they reduce to 13.28 EUR/MWh, which is a reduction by 9 %.

When comparing costs obtained by using objective (O1) to those resulting from using (O2) on component level, costs for pipes increase by 0.48 EUR/MWh. At the same time, the decrease is 1.37 EUR/MWh in compressor OPEX and 0.36 EUR/MWh in compressor CAPEX. As cost reduction is caused by changes to the network in the two enlarged areas, costs for these areas are analysed separately in the lower four bars of the plot. The changed configuration in the north reduces total system cost from 13.74 to 12.94 EUR/MWh or 6 %. There, pipe costs increase minimally, but investment costs for compressors decrease by 0.48 EUR/MWh. In the southern area, total system costs can decrease from 14.55 to 11.87 EUR/MWh or 19 % due to the modified layout. While pipe costs grow by 1.29 EUR/MWh, investment and operating costs of compressors can be reduced by 0.52 EUR/MWh and 3.45 EUR/MWh, respectively.

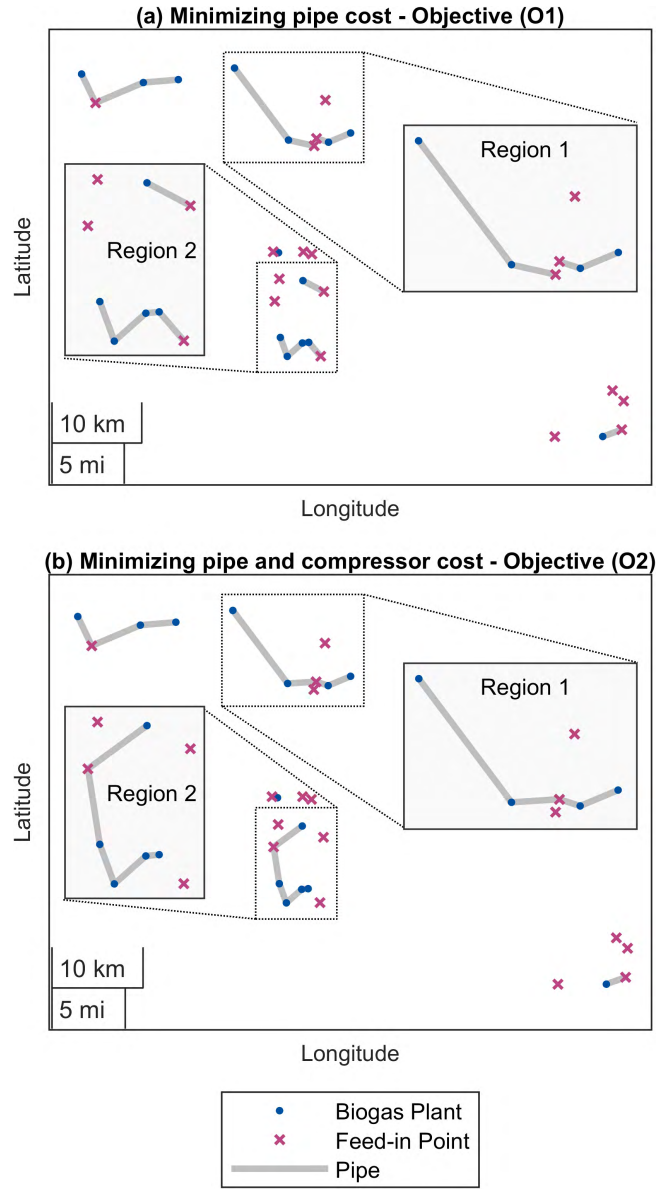

Figure 1. Network topology obtained by (a) optimization of pipe cost using objective (O1) and (b) optimization of pipe and compressor cost using objective (O2). Regions where clusters change due to the modified objective are magnified.

For Region 1, the modification of the objective leads to a change in network configuration. This is because compressor CAPEX can be reduced when only one compressor is built. However, operating cost changes only slightly, as the compressed amount of biomethane and, therefore, energy demand does not decrease. For Region

2, the cost structure changes more noticeably. Again, compressor CAPEX can be reduced by installing only one instead of two compressors. However, optimal feed-in locations change due to the modified objective. The reason for this is that eastern feed-in points operate at much higher pressure levels than the chosen feed-in point in the west. The great reduction of operating costs leads to a changed network combination with a slightly longer pipeline network.

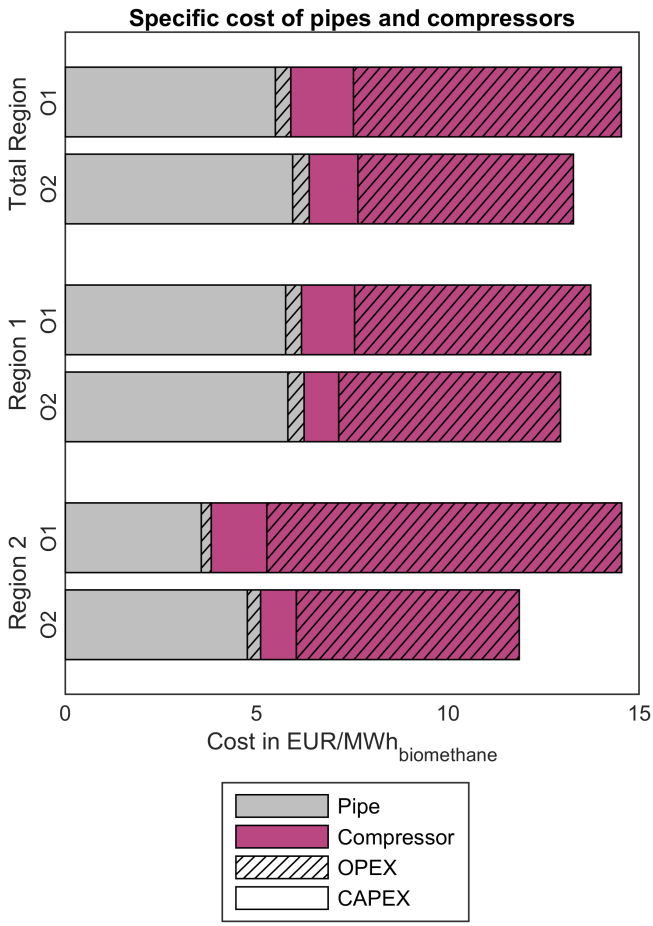

Figure 2. Specific cost of pipe and compressor infrastructure per MWh biomethane for the total region and magnified regions from Figure 1. The upper bar of each pair shows the cost of the network when minimizing only the pipe cost using the objective (O1). The lower bar shows the cost of the network when minimizing pipe and compressor cost using the objective (O2).

Figure 3 shows the results of applying the different allocation methods given in Equations (18) and (19) to the CAPEX of clusters in Region 1. Figure 3a and Figure 3b depict clusters from optimization with objectives (O1) and (O2). Biogas plants are numbered for identification. In Figure 3c, the allocated CAPEX cost per MWh of biomethane is plotted for each biogas plant in the region in a radar chart. Orange lines represent allocated cost after optimization with objective (O1), corresponding to clusters in Figure 3a, where the same color is used for depicted pipes. Blue lines show allocated cost after optimization with objective (O2), corresponding to the cluster in Figure 3b that is plotted in blue as well. Dashed lines indicate values obtained by equal cost allocation, and solid lines represent values calculated by volumetric share allocation.

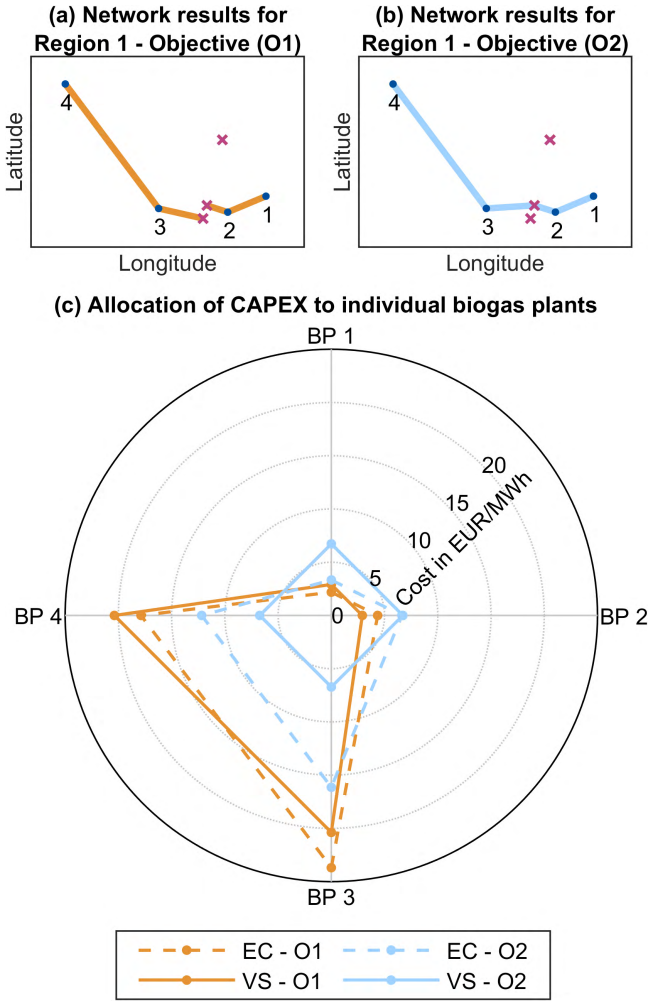

Figure 3. Network topology for Region 1 resulting from (a) optimization with objective (O1) and (b) optimization with objective (O2). Biogas plants (BP) are numbered for identification. Pipe colors are chosen corresponding to CAPEX allocation results in (c), where the equal cost (EC) and volumetric share (VS) methods are applied to both optimization results for all four biogas plants in the region.

Optimization with objective (O1) results in two clusters. CAPEX for the cluster containing biogas plants 1 and 2 are low, independent of the allocation method. This is caused by short pipeline distances and high biogas production volume in this cluster. Biogas plant 2 profits from VS allocation because its production flow is lower than the production flow of biogas plant 1. CAPEX for the

cluster containing biogas plants 3 and 4 are high, because pipe distances are long and the total biogas production volume of this cluster is low. Again, plant 3 profits from VS allocation because it has a lower production volume than plant 4. Optimization with objective (O2) results in only one cluster, so the total CAPEX now have to be allocated among all four biogas plants. Biogas plants 3 and 4 now reach much lower CAPEX costs than in the previous case, while the situation worsens for plants 1 and 2. This can again be reasoned with the lower production volumes of biogas plants 3 and 4 that additionally have much longer necessary pipe connections to the feed-in point. Independent of the allocation method, plants 3 and 4 profit from combining two clusters into one.

CONCLUSION AND OUTLOOK

In this paper, a mixed-integer linear programming-based network optimization problem for biomethane feed-in is extended to include modelling of compressor infrastructure. After optimization, two cost allocation methods are applied to the resulting networks to share the investment cost among participating biogas plants.

By including compressors in the optimization model, network layouts change, and system cost is reduced significantly. For the total region under investigation, a 9 % cost reduction is achieved; for one of the subregions, savings of 19 % are accomplished. The results undermine the importance of carefully selecting which components to model in energy system design optimization.

As a next step, additional system components, such as processing plants, should be included in the model. This is especially important because different upgrading technologies operate at different output pressures. Hence, further cost reductions can be expected when selecting the best technological fit for each feed-in point.

Moreover, this publication shows that investment cost for individual biogas plants can vary significantly based on the resulting network configuration and the chosen cost allocation method. It highlights that lowering total system cost and thereby achieving a better overall result does not always correspond to improving the situation for every participant when using established allocation methods. With equal cost and volumetric share allocation, some biogas plant operators can expect higher profits if they exclude other plants from the cluster.

To significantly increase the amount of biomethane fed into the grid and exploit the full potential in Austria, many biogas plant operators must participate. Two factors are detrimental to this development: Firstly, smaller biogas plants that are far from the grid face high investment costs. Secondly, large biogas plants close to the grid can take the opportunity and increase profits by excluding disadvantaged biogas plants from the network. To aid the design of effective governmental incentives and regulatory frameworks, further research should investigate cost allocation and cooperation among biogas plant operators.

ACKNOWLEDGEMENTS

The authors would like to express their sincere thanks to ÖVGW and Netz OÖ for providing data used in this study and for their trust and cooperation throughout the project. Their support was essential for the successful completion of this study.

AUTHOR IDENTIFIERS

Author ORCIDs:
Lena Lessig: 0009-0005-5976-8393
Martin Fischer: 0000-0002-9038-6773
Paul Schwarzmayr: 0000-0003-4735-3447
René Hofmann: 0000-0001-6580-4913

NOMENCLATURE

Acronyms

BP	biogas plant
CHP	combined heat and power
EC	equal cost
FP	feed-in point
MILP	mixed integer linear programming
SOS2	special ordered sets of type 2
VS	volumetric share

Indices/Sets

$i \in \{1, \ldots, N_{BP}\}$	biogas plant
$i' \in \{1, \ldots, N_{BP}\}$	biogas plant
$j \in \{1, \ldots, N_{FP}\}$	feed-in point
$e \in \{1, \ldots, N_E\}$	edges
$k \in \{1, \ldots, N_C\}$	pressure category

Optimization Variables

C_{BP}	binary matrix for connections between BPs
C_{E}	binary matrix indicating the number of edges to the FP
C_{FP}	binary matrix for connections between BPs and FPs
$C_{\mathrm{CAPEX}}^{\mathrm{comp}}$	vector of compressor CAPEX at FPs
$C_{\mathrm{OPEX}}^{\mathrm{comp}}$	vector of compressor OPEX at FPs
$C_{\mathrm{CAPEX}}^{\mathrm{pipe}}$	total pipe CAPEX
$C_{\mathrm{OPEX}}^{\mathrm{pipe}}$	total pipe OPEX
p_{FP}	vector of pressure at FPs
p_{out}	binary matrix of pressure category at FPs
$\dot{V}$	matrix of flows at edges
$\dot{V}_{\mathrm{FP}}$	vector of flows at FPs
$\dot{V}_{\mathrm{max}}$	vector of maximal capacities of FPs

Parameters

AF_{inv} annualization factor
$c_{\mathrm{CAPEX}}^{\mathrm{pipe}}$ pipe investment cost parameter
$c_{\mathrm{OPEX}}^{\mathrm{pipe}}$ pipe operating cost parameter
$c_{\mathrm{CAPEX}}^{\mathrm{comp}}$ compressor investment cost parameter
c_{el} electricity price
c_{m} maintenance cost parameter
d_{BP} matrix of distances between BPs
$\tilde{d}_{\mathrm{BP}}$ median distance between BPs
d_{FP} matrix of distances between BPs
$\tilde{d}_{\mathrm{FP}}$ median distance between BPs and FPs
e_{comp} compressor energy parameter
$k_{\mathrm{CAPEX}}^{\mathrm{comp}}$ compressor investment cost parameter
N_{BP} number of BPs
N_{FP} number of FPs
N_{E} maximum number of edges between BP and FP
N_{C} number of compressor pressure categories
$p_{\mathrm{out}}^{\mathrm{lim}}$ compressor outlet category parameter
t_{op} operating hours
$\dot{V}_{\mathrm{BP}}$ vector of BP production flows

Cost Allocation

$C_{\mathrm{CAPEX},i}^{\mathrm{EC}}$ cost allocated to BP i through EC method
$C_{\mathrm{CAPEX},i}^{\mathrm{VS}}$ cost allocated to BP i through VS method
$C_{\mathrm{CAPEX},Cl}^{\mathrm{pipe}}$ pipe cost of total cluster
$C_{\mathrm{CAPEX},Cl}^{\mathrm{comp}}$ compressor cost of total cluster
$N_{BP,Cl}$ number of BPs in cluster
$\dot{V}_{\mathrm{FP},Cl}$ flow at FP of cluster

REFERENCES

1. Intergovernmental Panel On Climate Change (IPCC). Energy Systems. In: Climate Change 2022 - Mitigation of Climate Change: Working Group III Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change, 1st ed. Cambridge University Press, Cambridge, pp 613–746 (2023). https://doi.org/10.1017/9781009157926.008
2. IEA. Outlook for Biogas and Biomethane. IEA, Paris (2025). https://www.iea.org/reports/outlook-for-biogas-and-biomethane
3. Da Costa Gomez C. Biogas as an energy option: an overview. The Biogas Handbook :1-16 (2013). https://doi.org/10.1533/9780857097415.1
4. E-Control. BERICHT 2025 gem. § 28 Abs. 3 E-ControlG Zur Versorgungssicherheit im Gasbereich (2025). https://www.e-control.at/documents/1785851/1811582/Versorgungssicherheitsbericht-Gas-2025.pdf
5. Parlament Österreich. Grün-Gas-Quote: Bundesrat schickt lückenhaftes Gesetz zurück an den Nationalrat. Parlamentskorrespondenz Nr. 817 vom 10.07.2024. [Accessed: 12 Jan 2026] https://www.parlament.gv.at/aktuelles/pk/jahr_2024/pk0817#XXVII_I_02455.
6. Bundeskanzleramt. Jetzt das Richtige tun. Für Österreich. Regierungsprogramm 2025-2029 (2025). https://www.bundeskanzleramt.gv.at/bundeskanzleramt/die-bundesregierung/regierungsdokumente.html
7. E-Control. Anlagenregister. [Accessed: 12 Jan 2026] https://www.e-control.at/anlagenregister.
8. Bundesministerium für Klimaschutz, Umwelt, Energie, Mobilität, Innovation und Technologie. Marktbericht 2024 Servicestelle für Erneuerbare Gase (2024). https://www.erneuerbaresgas.at/jart/prj3/erneuerbare_gase/data/uploads/SEG-Marktbericht%202024%20final%202025%2003%2026.pdf
9. AGCS Gas Clearing and Settlement AG. Statistik Biomethan - AGCS Bioregisterstelle Österreich. [Accessed: 12 Jan 2026] https://www.biomethanregister.at/de/statistik.
10. Bekkering J, Broekhuis TA, van Gemert WJT. Operational modeling of a sustainable gas supply chain. Engineering in Life Sciences 10:585-594 (2010). https://doi.org/10.1002/elsc.201000066
11. Hengeveld EJ, Bekkering J, van Gemert WJT, Broekhuis AA. Biogas infrastructures from farm to regional scale, prospects of biogas transport grids. Biomass and Bioenergy 86:43-52 (2016). https://doi.org/10.1016/j.biombioe.2016.01.005
12. Schröer D, Herlicka L, Heinold A, Latacz-Lohmann U, Meisel F. A network design problem for upgrading decentrally produced biogas into biomethane. Journal of Cleaner Production 452:142049 (2024). https://doi.org/10.1016/j.jclepro.2024.142049
13. Blanco V, Hinojosa Y, Zavala VM. The waste-to-biomethane logistic problem: a mathematical optimization approach. ACS Sustainable Chem. Eng. 12:8453-8466 (2024). https://doi.org/10.1021/acssuschemeng.4c01429
14. Huber D, Birkelbach F, Hofmann R. Network optimization for sustainable integration of decentralized biogas production. ASME 2024 18th International Conference on Energy Sustainability : (2024). https://doi.org/10.1115/es2024-121756
15. Huber D, Hofmann R. Optimierung der Integration von Biogasanlagen an das bestehende Gasverteilnetz (2024)
16. Zerbes I. Waffengleichheit im strafverfahren: traditionelles konzept vor aktuellen herausforderungen. Festschrift 50 Jahre ÖRAK :265-288 (2024).

https://doi.org/10.5771/9783214260125-265
17. E-Control. Strompreisentwicklungen. [Accessed: 18 Jan 2026] https://www.e-control.at/statistik/e-statistik/archiv/marktstatistik/preisentwicklungen.
18. Fachagentur Nachwachsende Rohstoffe e. V. (FNR). Clusterung von Biogasanlagen - Leitfaden für die Bereitstellung von Biomethan (2025). https://www.fnr.de/fileadmin/Projekte/2025/Mediathek/SR41_BGA-Cluster_250512_web_bf.pdf

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Simulation-Optimization vs. MILP Approaches for Real-Time Scheduling of Multiproduct Batch Plants

Engelbert Pasieka[a*], Sebastian Engell[bc]

[a] INOSIM Software GmbH, Joseph-von-Fraunhofer-Str. 20, 44227 Dortmund, Germany
[b] TU Dortmund, Department of Biochemical and Chemical Engineering, Emil-Figge-Str. 70, 44221 Dortmund, Germany
[c] ZEDO e.V., Joseph-von-Fraunhofer-Str. 20, 44227 Dortmund, Germany
* Corresponding Author: engelbert.pasieka@inosim.com.

ABSTRACT

Production scheduling in the process industry is often treated as a static optimization problem, although real plants require frequent rescheduling due to disturbances such as rush orders, equipment breakdowns, and changes in processing times. This paper compares a simulation-optimization approach that couples a discrete-event simulator with an evolutionary algorithm (EA) with a sequence-based mixed-integer linear programming (MILP) formulation for real-time scheduling of multistage batch systems. Both methods are embedded in an event-driven rolling-horizon framework under strict computation time limits.In static experiments for a 3-stage, 2-machine flow-shop setting (10 products, 20 orders, random processing times), the EA achieved lower makespans across all tested time budgets, improving results by about 7–13% on average compared to the MILP approach. In real-time experiments (40 initial orders, maintenance, three rush orders, 10 s and 60 s periodic updates), the solution quality of the MILP approach was lower after disturbances under restricted computation times. Each experiment was repeated 25 times with identical randomly generated processing times across methods per run; in 4% of all runs no solution was obtained within the available response time. Including explicit allocation decisions for the EA increases flexibility but can reduce short-term responsiveness within the rolling-horizon setting.Overall, the results indicate that MILP can be competitive in stable scenarios when sufficient solver time is available, whereas simulation-optimization is better suited to reactive scheduling in real-time settings where rapid schedule adaptation is critical.

Keywords: Simulation, Scheduling, Rolling Horizon Optimization

INTRODUCTION

Production scheduling is used in the process industry to improve resource utilization and to increase operational efficiency in dynamic situations. In the scientific literature, scheduling problems are often addressed as static optimization problems, where a fixed set of orders, resources, and processing times is assumed, and the problem is solved in advance for a given state of the plant, or for set of predefined scenarios to evaluate the performance of the plant. In these cases MILP is effective in computing high-quality or even optimal schedules if enough time is given or if the problem is decomposed.

In industrial practice, production environments encounter frequent and unforeseen disturbances, such as arrivals of rush orders, machine breakdowns, main–tenance activities, and variations in processing times. These events render precomputed schedules suboptimal or even infeasible, and they require adaptations of the schedule of operations, often under strict time constraints, to regain feasibility and to reduce losses in plant efficiency. In reality, production scheduling systems must frequently react to changes in the state of the production system rather than solve a one-time static optimization task [1]. This setting is commonly referred to as real-time or online scheduling.

Some research effort has been devoted to extend MILP formulations to online or rolling-horizon scheduling frameworks [2, 3]. Nevertheless, the performance of MILP in real-time environments is hampered by the need to frequently re-start the optimization with limited computation time in reaction to disturbances and structural

changes, such as new order arrivals or unavailability of resources, modify the objective or the feasible region and preclude a continuation of the branch-and-bound search, ruling out a reuse of the previously explored search tree [4].

Simulation-optimization provides an alternative paradigm for real-time scheduling. By coupling a detailed discrete-event simulation model with a metaheuristic optimization algorithm, it can incorporate complex plant constraints and adapt to changing system states [5]. EAs coupled to schedule builders generate feasible and executable schedules quickly and retain the state of the search across updates, which is advantageous in reactive environments [6] and can be used for high frequency scheduling tasks [7].

In this paper, we compare a standard MILP approach with a simulation-optimization approach based on an EA [6] across problem instances of increasing complexity. The available computation time is varied to reflect the requirement of real-time scheduling that feasible and high-quality solutions must be generated within strict time limits, even for complex problem instances. Static scheduling experiments are used to analyze the performance of both methods in the non-reactive case, while real-time scheduling experiments are used to assess their responsiveness and robustness in the presence of disturbances.

The MILP-based approach uses a sequence-based formulation, because sequence-based models are among the most compact and widely used formulations for scheduling problems. This formulation avoids the use of a discrete time grid and eliminates time discretization errors which facilitates initialization from an existing system state in online or real-time settings. Compared to time-indexed formulations, sequence-based models require fewer decision variables and constraints, which can lead to reduced computational effort and improved scalability.

SCHEDULING APPROACHES

We consider a multistage batch scheduling problem with a set of orders that have to be processed sequentially through multiple production stages. Each stage consists of one or more parallel units, and each order is assigned to one unit at every stage. Units process at most one order at a time, and preemption is not allowed. Processing times differ between units, and all orders are subject to release-time constraints.

Sequence-Based MILP Formulation

The MILP scheduling approach is a sequence-based formulation for multistage batch plants as in [8]. We set M to an upper bound on the scheduling horizon, e.g., $M = \sum_i \max_j \tau_{ij}$.

Sets and Indices

- $i, i' \in I$: orders
- $k \in K$: processing stages
- $j \in \mathrm{J}$: units

Decision Variables

- $x_{ij} \in \{0,1\}$: equals 1 if i is assigned to j
- $y_{ii'j} \in \{0,1\}$: equals 1 if i precedes i' on j
- $s_{ij} \geq 0$: start time of order i on j
- $s_{ik} \geq 0$: start time of i on k
- $ms \geq 0$: makespan

Parameters

- τ_{ij}: Processing time of order i on unit j

Objective Function

The objective is to minimize the makespan which is the completion time of the last processed operation across all orders (completion at final stage):

$$MS \geq s_{ik_{last}} + \sum_{j \in J_{ik_{last}}} \tau_{ij} x_{ij} \ \forall i \tag{1}$$

Constraints

Assignment constraints

Each order must be assigned to one unit at each stage:

$$\sum_{j \in J_{ik}} x_{ij} = 1 \ \forall \mathrm{i}, \mathrm{k} \tag{2}$$

Sequencing consistency constraints

One order must precede the other on a unit:

$$y_{ii'j} + y_{i'ij} \geq x_{ij} + x_{i'j} - 1 \ \forall \ i, i', j, k \ with \ i < i' \tag{3}$$

Start Time per Stage

$$s_{ik} = \sum_{j \in J_{ik}} s_{ij} \ \forall \ i, k \tag{4}$$

Activation of start time

$$s_{ij} \leq M x_{ij} \ \forall \ i, j \tag{5}$$

Disjunctive (non-overlap) constraints

Orders assigned to the same unit cannot overlap in time.

$$s_{ij} + \tau_{ij} \leq s_{i'j} + M(1 - y_{ii'j}) \ \forall \ i, i', j, k \tag{6}$$

Stage precedence constraints

Each order must complete processing at stage k:

$$s_{ik} + \sum_{j \in J_{ik}} \tau_{ij} x_{ij} \leq s_{ik+1} \ \forall \ i, k \tag{7}$$

Release-time constraints

Orders cannot start before their release times:

$$s_{ik} \geq \rho_i \ \forall i, k \tag{8}$$

Real-Time Scheduling Loop for the MILP Approach

The MILP formulation is embedded in an event-driven rolling-horizon framework. Scheduling updates are triggered periodically as well as by disruptive events such as rush orders.

At each update event, the start time of the scheduling horizon is aligned with the current state of the production system. Orders that have completed or are currently being processed are fixed by constraining their assignment X_{ij} and start-times S_{ij}. Newly released orders are included by updating the order set and the corresponding constraints. Lower bounds on start times are imposed to ensure consistency with executed operations and release dates.

When a maintenance event occurs on a machine, an additional disjunctive time constraint is introduced dynamically to ensure that no operation is scheduled within the maintenance window. For each operation o that can be assigned to machine j, a binary decision variable $z_{oj}(o,j)$ is created together with two bounding equations. The lower and upper constraints are formulated as

$$s_{oj} + \tau_{oj} \leq \mathrm{t}_j^{\mathrm{maint,start}} + M \cdot (z_{oj} + (1 - \mathrm{x_{oj}}))$$
$$s_{oj} \geq \mathrm{t}_j^{\mathrm{maint,end}} - M((1 - z_{oj}) + (1 - x_{oj}))$$

The constraint is conditioned on the assignment variable $x_{oj}(o,j)$ and is therefore enforced only if operation o is scheduled on machine j where $s_{oj}(o,j)$ denotes the start time of operation o on machine $j, \tau_{oj}(o,j)$ represents its processing time, and M is a sufficiently large constant. The interval $[\mathrm{t}_j^{\mathrm{maint,start}}, \mathrm{t}_j^{\mathrm{maint,end}}]$ defines the maintenance period on the respective unit j. By incorporating the processing time in the first constraint, the model ensures that an operation assigned before maintenance must also be completed before the maintenance window begins. If the operation is assigned to the affected machine, the binary variable activates a big-M disjunction that forces it to be scheduled either entirely before or entirely after the maintenance period. If $z_{oj}(o,j) = 0$, the completion time is restricted to occur before maintenance starts; if $z_{oj}(o,j) = 1$, the first inequality is relaxed, and the operation is forced to start after the maintenance period.

The MILP model is rebuilt at each update and solved within the available computation time. To improve performance, warm-start information from the previously computed solution is provided to the solver in the form of initial values for assignment variables and, optionally, for start times. However, due to structural changes in the model and the limited ability of MILP solvers to retain internal search state across re-optimizations, the solver effectively restarts after each disturbance.

The real-time MILP scheduling procedure can be summarized as follows:

1. Initialize the scheduling horizon based on the current plant state
2. Generate an event list consisting of periodic and event-driven updates
3. At each update event, fix the completed and the ongoing operations
4. Update orders, resource availability, and constraints
5. Initialize the solver using information from the previous solution
6. Solve the MILP within the available computation time and transfer the best available schedule to the production system
7. Repeat until the end of the simulation horizon.

Simulation-Optimization Approach

The simulation-optimization–based real-time scheduling approach used in this work is based on a previously published framework [6] that couples a detailed discrete-event simulator with an EA. The method operates as a continuously running optimization process that preserves the search state across updates and incrementally adapts schedules in response to changes in the production system.

This approach is used as a reference method for comparison with the MILP-based approach. The EA design, population management strategies, and simulation integration are described in detail in the original publication and are only briefly explained here.

The EA is embedded in a simulation–optimization real-time scheduling loop: schedules are encoded as genotypes (primarily a global order sequence π_{seq} and, possibly resource-allocation decisions π_{alloc}), and each genome is evaluated by executing it in a discrete-event simulator with a detailed model that represents the current state of the production system and acts as a schedule builder, using priority rules and enforcing logic or plant constraints (including repair mechanisms to keep the schedules executable). The EA runs continuously, iterating through selection, recombination, and mutation to gradually improve schedule quality on the current simulation model. At periodic update times and when disruptive events occur (e.g., rush orders, maintenance, delays), the production system feeds new state information to the scheduler, which updates the simulation model and reinitializes the population using strategies such as carrying over the full population, keeping the best individual, or random restart. The best available schedule can already be exported after one iteration and is then further improved as optimization continues.

CASE STUDY

The experimental study is designed to compare the performance of the sequence-based MILP approach and simulation-optimization based on an EA under strict computation time limits. Two types of experiments are conducted. First, static scheduling experiments are used to analyze the behavior of both approaches independent of the effect of real-time updates due to disturbances. Second, real-time scheduling experiments are performed to evaluate the performance under dynamic operating conditions with disruptive events. In all experiments, the fitness criterion is the minimization of the makespan.

To ensure a fair comparison, both approaches are implemented on a single CPU core, and the schedules generated by the EA are evaluated sequentially. All experiments were performed on a Ryzen 9800X3D CPU with 64GB of RAM. The MILP model was solved using the IBM ILOG CPLEX Optimizer 20.1.0.0 via the GAMS command-line interface.

Static Scheduling Experiments

The static experiments are performed for a multistage batch scheduling problems as described above. The production system consists of three stages, each equipped with two parallel units of different efficiency. A total of ten product types and twenty orders are considered in each experiment.

The processing times were generated randomly for each order–machine combination, with values uniformly distributed between 0 and 4 time units.

Objective and Constraints

Both the MILP and EA approaches solve identical scheduling problem All precedence constraints, machine capacity constraints, and order release constraints are enforced in both approaches.

Computation Time Limits

To reflect the requirements of real-time and near-real-time scheduling, the available computation time is strictly limited. Both methods are executed with computation time limits of 10, 15, 20, and 25 seconds. Each configuration is repeated ten times using different randomly chosen processing times.

The MILP solver runs until the specified time limit is reached or an optimal solution has been found. The EA runs continuously during the available time and improves the candidate schedules incrementally.

Evolutionary Algorithm Configuration

The EA uses a custom programmed compiled discrete-event simulator for schedule building and evaluation. In the static scheduling experiments, the decision variables of the EA comprise either only the sequencing of orders by a global priority parameter or in addition also the allocation of orders to machines. If the allocation is not controlled by the EA, the schedule builder assigns the machine with the earliest completion time. The EA parameter settings are taken from the prior study reported in [6] and remain fixed across all experiments.

Real-Time Scheduling Experiments

The real-time scheduling experiments are performed for the same multistage batch plant with three stages and two units per stage. Forty orders are released at the beginning of the simulation. All experiments follow an event-driven rolling-horizon scheme with a periodic update interval of either 10 seconds or 60 seconds. At each update event, rescheduling is performed based on the current state of the production system. The update interval also defines the maximum computation time available for solving the scheduling problem at each event. In addition, rush orders and maintenance occur within the production horizon, these are not known a priori.

For each update interval, three scheduling approaches are evaluated: the MILP formulation, the EA with allocation genome, and the EA without allocation genome. Each experiment is repeated 25 times with randomly generated processing times. The processing times for each order–unit combination are sampled randomly with a uniform distribution in the interval [1, 12] minutes.

The EA is implemented with the same settings as in the static experiments.

Disruptive Events and Update Policy

During execution in the real-time setting, three rush orders are introduced, representing unplanned disturbances that require immediate rescheduling. Rush orders become known at the 1st, 3rd, and 5th periodic update events.

In addition, a maintenance request is introduced temporarily disabling individual machines. A maintenance downtime of 500 seconds starting at simulation time 2000 seconds is requested for machine M01 at the second update event. At the fourth update event, a maintenance downtime of 500 seconds starting at simulation time 3000 seconds is requested for machine M04. Operations that are still running at the requested maintenance times are completed before the maintenance started but during the full maintenance interval no new operation can be started.

Experimental Variants

The EA-based simulation-optimization approach is evaluated in two configurations:

1. Using a genome encoding only of the sequence of orders, and
2. Using a genome encoding of both the sequence of orders and the allocation of orders to machines.

In the first configuration, machine allocation is

determined by a First-In-First-Out (FIFO) (earliest completion time) dispatching rule.

The MILP-based approach employs the sequence-based formulation described in Section 2 and is embedded in the same real-time scheduling framework. At each update event, the MILP model is rebuilt so that it reflects the current system state and then is solved within the available computation time.

RESULTS

In this section we present and discuss the results of the static and real-time scheduling experiments described in Section 3.

Table 1: Static scheduling results for a fixed problem structure under varying computation time limits. For each time limit, 10 instances with randomized processing times were solved using the MILP-based approach and the EA. The table reports the average relative difference in solution quality over all instances. The objective was the minimization of the makespan (see eq. 1).

Time limit (s)	EA better (count)	MILP better (count)	Mean relative difference (%)	Std. dev. of rel. diff. (%)
10	9 / 10	1 / 10	12.49	11.30
15	9 / 10	1 / 10	13.12	8.97
20	8 / 10	2 / 10	7.37	12.86
25	10 / 10	0 / 10	9.79	5.18

Table 2: Real-time scheduling results for update intervals of 10 seconds and 60 seconds. A Each experiment was repeated 25 times. In every run, the processing times were randomly generated but kept identical across the compared methods for the same run.

Update interval	Method	Mean makespan [h]	Std. dev.	Improvement vs. MILP %
10 s	MILP	4.72	0.57	-
	EA w. alloc.	2.24	0.22	52.4
	EA w.o. alloc.	2.06	0.21	56.2
60 s	MILP	2.94	0.50	-
	EA w. alloc.	2.17	0.24	26.1
	EA w.o. alloc.	1.96	0.26	33.5

Static Scheduling Results

Table 1 summarizes the results of the static scheduling experiments for a fixed problem structure under varying computation time limits for the minimization of the makespan. For each time limit, multiple problem instances were generated, and the table reports the resulting average solution quality obtained by each approach.

For all computation time limits, the EA achieves a better average solution quality than the MILP-based approach. As the available computation time increases, the performance gap between the two approaches does not decrease monotonically.

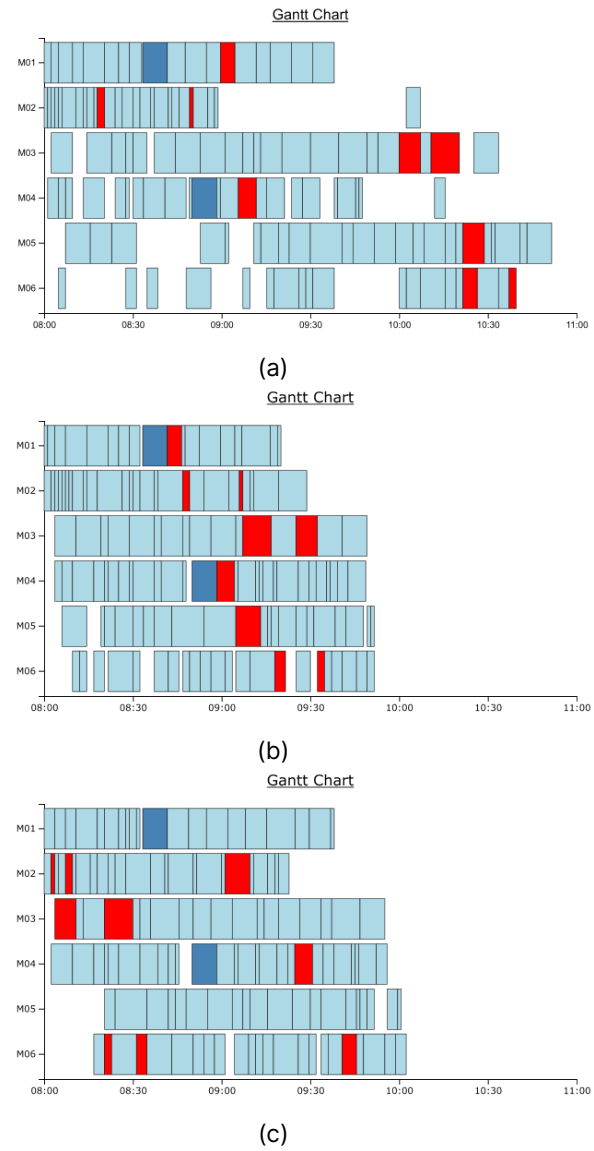

Figure 1: Gantt charts of the schedules generated by the MILP approach (a), EA with allocation genome (b), and EA without allocation genome (c) after 10 periodic updates with an update interval of 60 seconds. The rush orders are highlighted in red, and the maintenance downtimes are highlighted in dark blue.

Real-Time Scheduling Results

Table 2 presents the results of the real-time

scheduling experiments for the two different update intervals. The number of update cycles is fixed at 10 for all experiments. After the tenth update, no further rescheduling is performed, and the last computed schedule is simulated until the end of the production horizon. The table reports the average makespan values for the full simulation horizon. Runs in which the MILP solver did not return a feasible solution within the specified response time were excluded from the averaging procedure, as no valid schedule could be extracted in these cases. Across all experiments, this occurred in 4% of all runs.

The EA-based approaches achieved a better performance under real-time conditions with limited computation time per update.

When the update interval is increased to 60 seconds, the gap between the performance of the EA and of the MILP formulation becomes smaller. In 4% of all runs, the MILP approach did not provide a feasible solution within the available computation time. This occurred equally across both response-time settings.

For both response-times, the inclusion of allocation decisions in the EA genome led to slightly higher makespan values compared to the sequence-only encoding, so the higher flexibility of this formulation is outweighed by the increased size of the search-space.

The results are illustrated by the Gantt charts shown in Figures 1 a-c. All figures show results for the same simulation run, i.e., they are based on identical randomly generated processing times and the same disruptive events, including rush-order arrivals and maintenance downtimes.

The schedule generated by the MILP-based approach (Figure 1a) exhibits an uneven distribution of the workload across parallel units, showing idle periods on some machines. In contrast, the EA-based schedules (Figures 1b and 1c) show a more balanced allocation of processing tasks across the available machines.

CONCLUSIONS

This paper compared a sequence-based MILP approach with a simulation-optimization approach based on an EA for scheduling of multiproduct batch plants under strict computation time limits. The comparison was performed using static experiments as well as real-time scheduling scenarios with periodic updates and disruptive events.

The static results show that the EA achieves competitive or superior makespan values under tight time budgets.

In real-time scheduling, the MILP-based approach also exhibits a lower solution quality under restricted computation times. Increasing the available solver time improves the performance of the MILP formulation, however, the EA variants show a consistent advantage across both response-time settings. In a small number of cases (4% of all runs across all experiments), no feasible solutions were obtained by the MILP formulation within the available computation time. This is an additional weakness of the MILP approach which is not reflected in the quantitative comparisons where these runs were ignored. These infeasibilities can be related to the limited ability of the MILP solver to retain internal search state information across successive re-optimizations.

In contrast, the EA preserves information on the search state between updates and adapts the schedules incrementally. The EA only handles few degrees of freedom, the remaining ones are handled by a rule-based schedule builder. While theoretically including allocation decisions should provide better solutions than handling the allocations by the rule-based schedule builder, the increase of the size of the search space outweighs this advantage for short computation times. The use of a schedule builder also increases the robustness of the optimization, at the price of a loss of optimality. Overall, the results indicate that simulation-optimization is a competitive approach for real-time scheduling when fast schedule adaptations are required.

Future work will extend the analysis to different problem settings and will investigate hybrid strategies combining MILP-derived baseline schedules with simulation-optimization for real-time adaptation.

REFERENCES

1. Espinaco F, Henning GP. Industrial rescheduling approaches: where are we and what is missing?. Proceedings of the 11th International Conference on Production Research – Americas :461-467 (2023). https://doi.org/10.1007/978-3-031-36121-0_58
2. Subramanian K, Maravelias CT, Rawlings JB. A state-space model for chemical production scheduling. Computers & Chemical Engineering 47:97-110 (2012). https://doi.org/10.1016/j.compchemeng.2012.06.025
3. Chu Y, You F. Rolling horizon optimization for process scheduling under uncertainty. *AIChE J* 59:2110–2127 (2013) https://doi.org/10.1002/aic.13963
4. Gamrath G, Hiller B, Koch T, Pfetsch ME. Reoptimization techniques for mixed integer programming. *Math Program Comput* 8:1–29 (2016) https://doi.org/10.1007/s12532-015-0094-1
5. Klanke C, Engell S. Scheduling and batching with evolutionary algorithms in simulation–optimization of an industrial formulation plant. Computers & Industrial Engineering 174:108760 (2022). https://doi.org/10.1016/j.cie.2022.108760

6. Pasieka E, Engell S. Evolutionary algorithm based real-time scheduling via simulation-optimization for multiproduct batch plants. Systems and Control Transactions 4:894-899 (2025). https://doi.org/10.69997/sct.183349
7. Gaina RD, Devlin S, Lucas SM, Perez-Liebana D. Rolling horizon evolutionary algorithms for general video game playing. IEEE Trans. Games 14:232-242 (2022). https://doi.org/10.1109/tg.2021.3060282
8. Maravelias CT. Chemical production scheduling. Cambridge University Press (2021). https://doi.org/10.1017/9781316650998

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Design and Optimization of Supply Chain for Citrus Biorefineries: A Regional Approach for Waste Valorization in Brazil

Marilia G. L. Cavenaghi[a]*, Larissa T. Bruschi[a], and Moises T. dos Santos[a]

[a] University of São Paulo, Department of Chemical Engineering São Paulo, Brazil
* Corresponding Author: marilia.cavenaghi@usp.br.

ABSTRACT

Brazil is the world's largest producer of orange juice, generating significant peel residues that are currently underutilized. This study proposes a mixed-integer linear programming (MILP) framework for optimal supply chain design, utilizing Special Ordered Sets of type 2 (SOS2) to accurately represent non-linear investment costs. The model maximizes Net Present Value (NPV) by integrating production costs with multi-echelon logistics, including inland transport, port handling, and international maritime freight. Applied to a case study in São Paulo, the framework evaluates pathways for the co-production of D-limonene, pectin, and bioenergy. Results indicate a positive NPV of BRL 1.27 billion, with pectin contributing over 65% of total revenue. The optimization favors centralized configurations in Araraquara or Matão to exploit economies of scale while minimizing the transport of high-volume, wet biomass. Notably, total transportation costs represent only 1.13% of expenditures, as the high value-density of the bioproducts effectively absorbs global distribution costs. This work demonstrates that in regions with high biomass abundance, strategic success is primarily governed by process efficiency and industrial scale rather than logistical constraints, providing a robust tool for decision-making in the citrus bioeconomy.

Keywords: Supply chain optimization, biorefinery, orange peel waste, mathematical programming

INTRODUCTION

The industrial processing of oranges in São Paulo generates massive volumes of peel residues, which are typically diverted to low-value uses despite their high concentration of pectin and D-limonene. These residues support the evaluation of biorefinery concepts to improve economic performance through high-value product diversification. At a regional scale, the deployment of such systems is governed by complex trade-offs between biomass availability, conversion technologies, and market demand, requiring mathematical programming for integrated assessment.

Mixed-integer linear programming allows the explicit representation of discrete decisions, such as facility siting and technology selection, alongside continuous material flows and economic performance metrics. Such models enable the systematic evaluation of alternative network configurations and support the identification of coordinated and economically consistent solutions for the regional deployment of citrus biorefineries [1, 2].

Recent literature has extensively explored the optimization of biorefinery supply chains within the Brazilian context, primarily focusing on the sugarcane and energy sectors. Bruschi *et al*. (2026) [3] provided a comprehensive review of MILP architectures for biomass supply chains, establishing robust modeling foundations for regional networks. Aba *et al.* (2022) [4] proposed a multi-period MILP model to manage the integrated supply chain of fossil fuels and biofuels, addressing seasonality and demand fluctuations for corn and sugarcane ethanol. Similarly, [5] developed a bi-objective MILP framework to evaluate the trade-offs between financial profit and environmental impact, identifying optimal technology configurations for the transition from sugar-only to mixed-product plants. More recently, [6] utilized robust MILP optimization to assess low-carbon ammonia and hydrogen production from ethanol and biomethane,

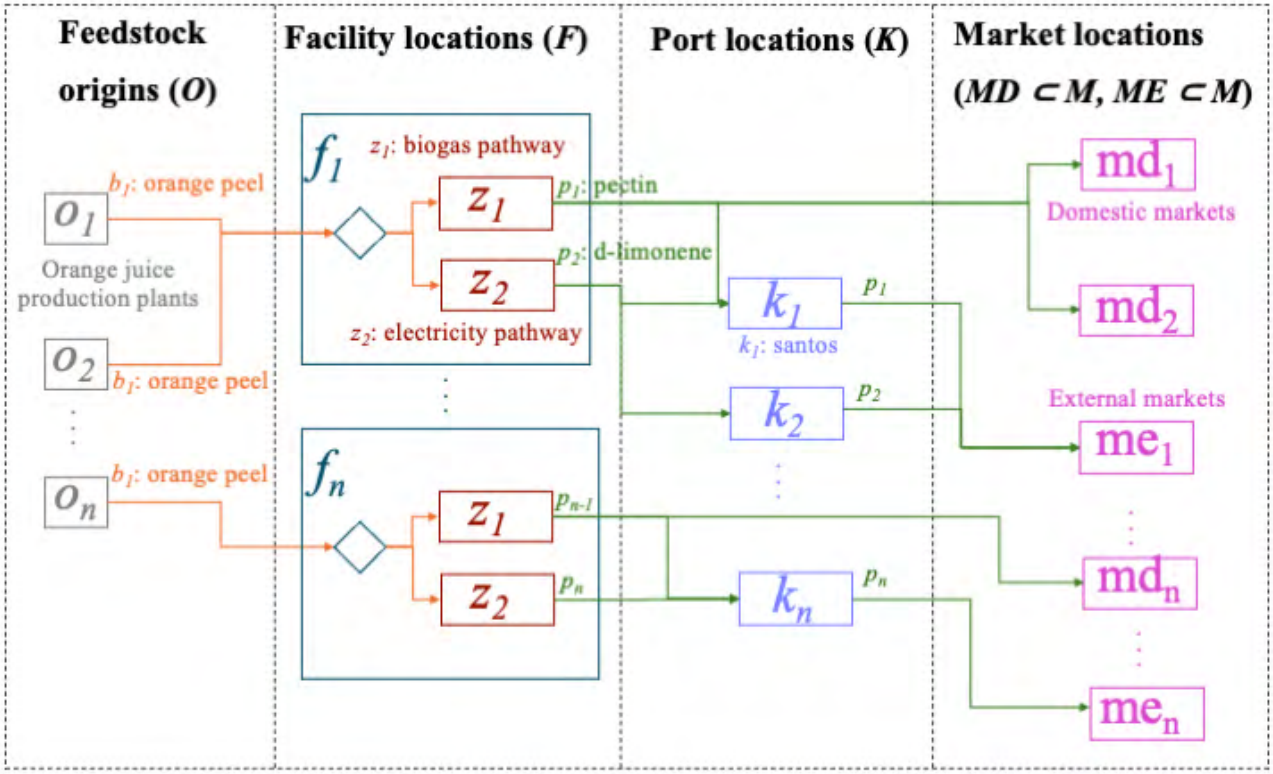

Figure 1. Schematic representation of the citrus biorefinery supply chain and associated modeling sets. Orange juice processing plants $o \in O$ generate orange peel residues $b \in B$, which are supplied to candidate biorefinery facility locations $f \in F$ for conversion through technologies $z \in Z$. Marketable products $p \in P$ are distributed to domestic market ($MD \subset M$) and external ($MD \subset M$)markets, with the latter routed via port nodes $k \in K$. Optimization variables determine facility siting, technology selection, and material flows across the network layers.

highlighting the potential for advanced transition products within existing industrial infrastructures. Although these models provide sophisticated insights into regional energy and fuel distribution, they typically prioritize domestic markets and focus on energy products or traditional sugarcane derivatives.

The proposed optimization framework integrates high-value biochemical production with residual energy valorization under a global logistics perspective. By expanding the scope beyond regional energy self-sufficiency, this model incorporates maritime freight indices and port-specific costs to evaluate citrus biorefineries as export-oriented hubs. The use of Special Ordered Sets of type 2 (SOS2) for non-linear investment costs allows the framework to assess how regional biomass density in São Paulo facilitates a strategic shift from market-oriented to feedstock-oriented configurations. This approach seeks to provide a more detailed understanding of the industrial valorization of citrus processing residues by capturing the economic trade-offs between process scale and global distribution requirements.

METHODOLOGY AND MODEL FORMULATION

A generic supply-chain-level optimization framework is adopted to analyze alternative configurations for the valorization of residues from the orange juice industry. The problem is formulated as a mixed-integer linear programming (MILP) model, designed to support strategic and operational decisions within integrated biorefinery supply chains. The formulation adopts a modular, set-based structure comprising biomass types (set B), biomass origins (set O), candidate facility locations (set F), conversion technologies (set Z), products (set P), potential domestic and external markets (set M), logistic port nodes (set K), wastes generated in conversion (set W), process inputs (set I), consumed utilities (set U), and time periods (set T). **Figure 1.** presents a schematic overview of the main components and modeling sets considered in the proposed framework.

The supply chain is represented as a multi-echelon network composed of biomass supply origins O, candidate

processing locations F, domestic markets $MD \subset M$, external markets $ME \subset M$, and port nodes K, which act as export gateways. Processing facilities and conversion technologies may be installed in selected municipalities through binary decision variables, while material flows are represented by continuous variables linking the different network layers.

Transportation costs are differentiated by the nature of the origin–destination pair and the specific logistics echelon. For biomass sourcing and domestic market ($MD \subset M$) product distribution, costs are calculated as a linear function of the distance between nodes and the transported mass, using a standardized cost-per-ton-kilometer parameter. Products destined for external markets ($ME \subset M$) are first routed from processing facilities F to port nodes K via road transport. This inland segment follows the same distance-based cost logic as domestic flows, but is supplemented by an export-specific cost layer that incorporates port handling fees and maritime freight to final international destinations. This multi-echelon formulation explicitly distinguishes between domestic and export-oriented logistics, ensuring that the model captures the economic trade-offs associated with different market allocation strategies.

Market allocation is governed by demand limits defined according to the geographic nature of the consumer nodes. For domestic markets ($MD \subset M$), demand is specified at the municipal level, reflecting local consumption. For the export segment ($ME \subset M$), demand is aggregated at representative port nodes for each destination continent. It is assumed that all products are sold within the same period, with no storage of biomass or products. This structure allows the model to capture the trade-offs between fragmented domestic supply and high-volume, port-centric export logistics

Processing capacity is modeled as a continuous decision variable and is directly associated with capital investment. Investment costs are represented using a non-linear cost function with a scaling exponent, following standard capital cost scaling principles [7]. To ensuring the representation of economies of scale within a tractable MILP framework, this non-linear investment function is approximated using a piecewise linear formulation. The approximation is implemented using Special Ordered Sets of type 2 (SOS2) [8], relying on a convex combination of discretized points. The capacity domain is discretized using log-spaced breakpoints, allowing a more accurate representation of the non-linear cost function over a wide range of capacities. Fixed operational costs are assumed to be proportional to the investment cost and are incurred over the planning horizon.

The objective function of the model is defined as the maximization of the net present value (NPV) over the planning horizon and is expressed as follows:

$$max\,\mathrm{NPV} = -\sum_{t\in T}\frac{C_t^{\mathrm{inv}}}{(1+r)^t} + \sum_{t\in T}\frac{\left(R_t - C_t^{\mathrm{op}} - C_t^{\mathrm{tr}}\right)}{(1+r)^t} \tag{1}$$

Where $t \in T$ denotes the time periods of the planning horizon and r is the discount rate. C_t^{inv} represents the investment cost incurred in period t, associated with the installation and capacity of processing facilities. R_t denotes the revenue obtained from product sales in period t. C_t^{op} corresponds to operating costs, including fixed and variable components, while C_t^{tr} represents transportation costs associated with biomass supply and product distribution. All cash flows are discounted to their present value using the discount rate r.

Case study: regional citrus waste valorization in São Paulo

The optimization model was applied to a regional case study focused on the valorization of orange peel generated by the orange juice industry in the state of São Paulo, Brazil. This region concentrates the largest installed orange juice processing capacity in the country and exhibits a well-defined spatial structure for biomass generation, logistics, and market access, which made it suitable for a supply-chain-level assessment of alternative biorefinery configurations.

Wet orange peel residue was assumed to be acquired directly from orange processing plants and supplied to conversion facilities without intermediate storage. The alternative valorization routes evaluated in the case study were based on the process configurations developed by [9], describe an integrated biorefinery scheme for citrus processing residues. In this scheme, orange peel is initially processed for the recovery of D-limonene and pectin. After the hydrolysis stage, a solid–liquid separation step is applied to remove residual solids from the hydrolysate prior to pectin precipitation. The resulting solid fraction, hereafter referred to as filter cake, constitutes the residual biomass directed to energy valorization.

In [9], authors propose two alternative energy valorization pathways for the filter cake generated after pectin recovery. In the first pathway, the solid residue is directed to anaerobic digestion for biogas production. In the second pathway, the same residue is routed to a gasification process followed by electricity generation. Accordingly, biogas route integrates the extraction of D-limonene and pectin with biogas production, while electricity route integrates the extraction of D-limonene and pectin with gasification-based electricity generation.

Technological parameters required to represent these routes within the optimization model, including biomass characterization, material conversion factors,

product yields, residue generation rates, and electricity consumption associated with individual process units, were obtained from the study of [9]. Data related to utility consumption, specifically cooling water and steam requirements, were obtained from [10].

Spatially explicit data on biomass availability were obtained from sectoral information reported by CitrusBR [11], including the location of juice processing plants, installed processing capacity, and processed orange volumes.

Market nodes and demand limits were defined to reflect both regional and global consumption patterns. For domestic markets (MD), the analysis focused on municipalities in the Southeast region and Goiás, where food, pharmaceutical, and cosmetic industries are concentrated. The domestic pectin and D-limonene markets were estimated using trade data and industrial profiles, resulting in municipal-level demand limits.

To rigorously assess export-oriented configurations (ME) the port of Santos was selected as the primary gateway for international flows. External demand was disaggregated into representative port nodes for each destination continent: Houston (North America), Rotterdam (Europe), and Shanghai (Asia). For these destinations, external demand was estimated based on the average historical export volumes of pectin and orange essential oils to each respective continent.

Based on these demand estimations, a base case scenario was defined by satisfying 30% of the modeled domestic demand (Southeast region and Goiás) and 30% of the external demand for each continent (North America, Europe, and Asia).

Logistics costs for the export segment were explicitly modeled to include both inland and maritime stages. The maritime transportation cost factor (BRL/mt.nm) was derived from Drewry's World Container Index [12] and associated maritime route distances. Additionally, port handling costs per product tonne were incorporated based on official reports from the National Agency for Waterway Transportation [13]. This approach ensures a more realistic economic evaluation compared to simplified models, capturing the impact of port-centric logistics on the overall Net Present Value (NPV).

For energy products, it was assumed that all generated output was fully commercialized. Candidate locations for processing facilities were restricted to municipalities with access to natural gas infrastructure or power grid connection, allowing for the direct injection of biogas or the full dispatch of electricity. Under these assumptions, transportation costs for both energy carriers were neglected, as demand was effectively localized at the point of generation. Based on these infrastructure requirements, candidate locations were selected among municipalities either close to biomass sources or major consumption centers, enabling the assessment of trade-offs between feedstock proximity and market access.

Prices for pectin and D-limonene were primarily derived from international trade data. The price of biogas was assumed equivalent to the natural gas price with an additional premium for its renewable origin. Electricity prices were based on representative average values reported by the [14].

To avoid trivial solutions with no installed capacity while preserving the ability of the model to return economically unfavorable configurations, a sufficiently small minimum demand constraint was imposed, allowing scenarios with negative net present values (NPV) to be obtained when economic conditions were not favorable.

Model implementation

The optimization model was implemented in Python using the Pyomo modeling framework, with input data provided through structured Excel files. The model was solved using Gurobi (version 12.0.3) under an academic license, with a relative optimality gap tolerance of 5.0×10^{-3}. Computations were performed on a machine equipped with an Intel® Core™ i7-8559U CPU (4 cores, 8 threads) and 16 GB of RAM. The resulting MILP comprises approximately 15, 946 constraints and 60, 800 variables, including 30 binary variables and SOS2 constraints for piecewise linear cost approximation. The base case was solved within the specified tolerance in approximately 274 s of computation time.

RESULTS AND DISCUSSION

The base case scenario considers the supply of 30% of domestic demand and 30% of external demand and serves as a reference for evaluating the economic performance and structural configuration of the citrus biorefinery. The optimal processing capacity is 87, 676.06 mt/year of wet orange peel, resulting in the production of around 3, 621.93 mt/year of pectin, 2, 576.43 mt/year of limonene, and 11, 485.56 MW of biogas. Under these conditions, model results indicate a positive net present value (NPV) of approximately BRL 1.27 billion, demonstrating high economic potential of orange residue valorization within the current market framework.

An analysis of the economic structure presented in **Figure 2** reveals that total revenues reach approximately BRL 4.6 billion, with pectin accounting for the largest share (65.23%), followed by limonene (34.70%), while energy products contribute negligibly to total revenue (0.08%). This distribution highlights the strong dependence of system profitability on high-value chemical products, particularly pectin.

Regarding operating expenditures costs at present value, represent the dominant component, totaling approximately BRL 1.88 billion. Within this category, production costs are the primary driver, representing

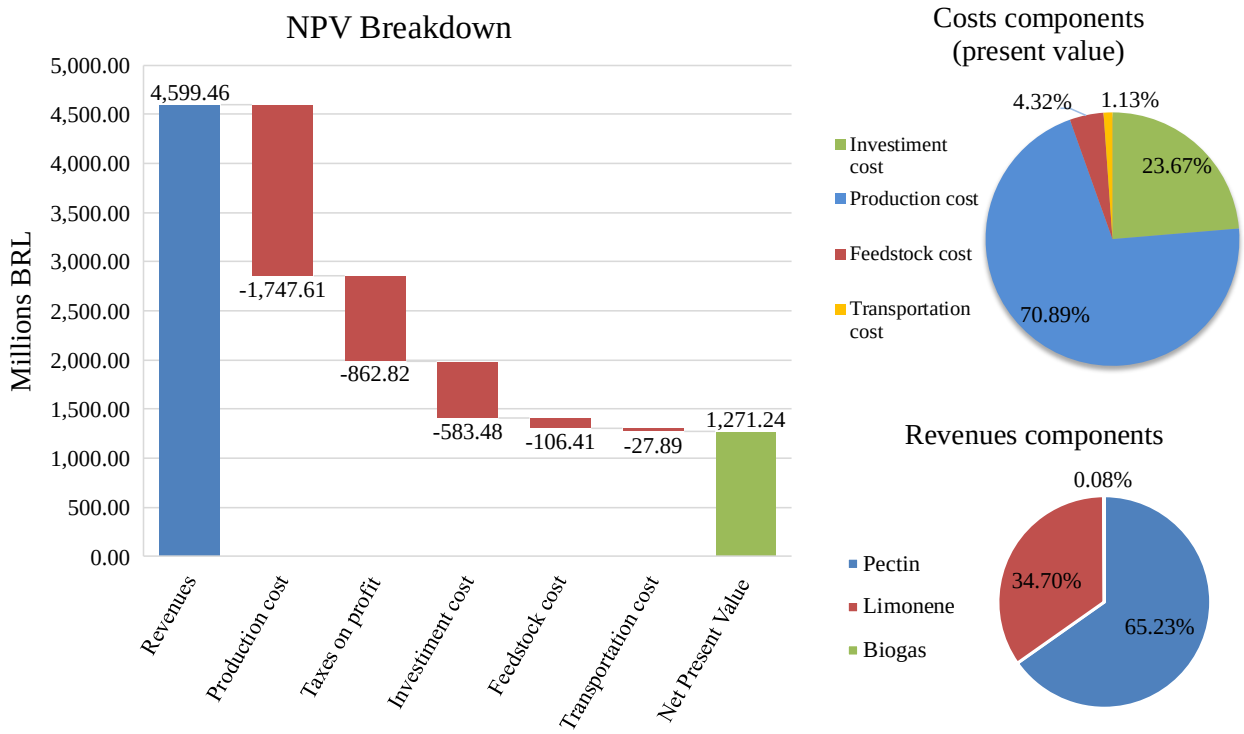

Figure 2. Base case NPV breakdown, costs composition at present value, and revenues composition.

70.89% of the total, followed by capital-related expenditures (23.67%) and feedstock acquisition (4.32%). This cost composition indicates that the economic performance of the system is dictated by the high utility requirements and energy intensity of the extraction and solvent recovery stages, rather than by biomass supply costs or logistics. Capital investment (CAPEX) for base case is BRL 5.83 billion and represents a relatively small share of total costs compared to operating expenditures, suggesting that economies of scale are effectively captured in the optimal solution.

The optimal network configuration selects a single processing facility located in Araraquara, utilizing the biogas-based conversion route. This centralized configuration reflects the combined effect of economies of scale and the spatial distribution of biomass availability, which is sufficient to sustain a large-scale facility without requiring decentralization.

The selection of the biogas route is not driven by direct revenues from energy products, which are marginal, but rather by its lower capital investment and utility consumption compared to the gasification-based electricity pathway.

Overall, the base case results highlight that the chain profitability is strongly dependent on high-value product markets, particularly pectin, and production costs, which represent the most significant financial burden. While logistics and feedstock costs play a secondary role in the overall economic outcome, these findings provide a clear baseline for interpreting the sensitivity of the system to key economic parameters explored in the following sections.

Sensitivity analysis

The sensitivity analysis was performed to evaluate how variations in key economic and operational parameters affect both the profitability and structural configuration of the system, using the base case as reference. **Figure 3** summarizes the impact of each parameter on the net present value (NPV), relative to base case.

The results indicate an asymmetry in the economic drivers of the system, with market-related parameters exerting a stronger influence on performance than cost-related variables. Among these, pectin price has the largest impact on profitability. A 30% reduction in pectin price decreases the NPV from BRL 1.27 billion to BRL 5.14 million, while a 30% increase raises the NPV to about BRL 2.28 billion. These variations occur without changes in system configuration, indicating that price fluctuations primarily affect revenue levels rather than structural decisions. This behavior is consistent with the base case results, in which pectin accounts for the largest share of total revenues.

Market demand affects both economic performance and network configuration. A 75% reduction in demand leads to a decrease in NPV to approximately BRL 59.13

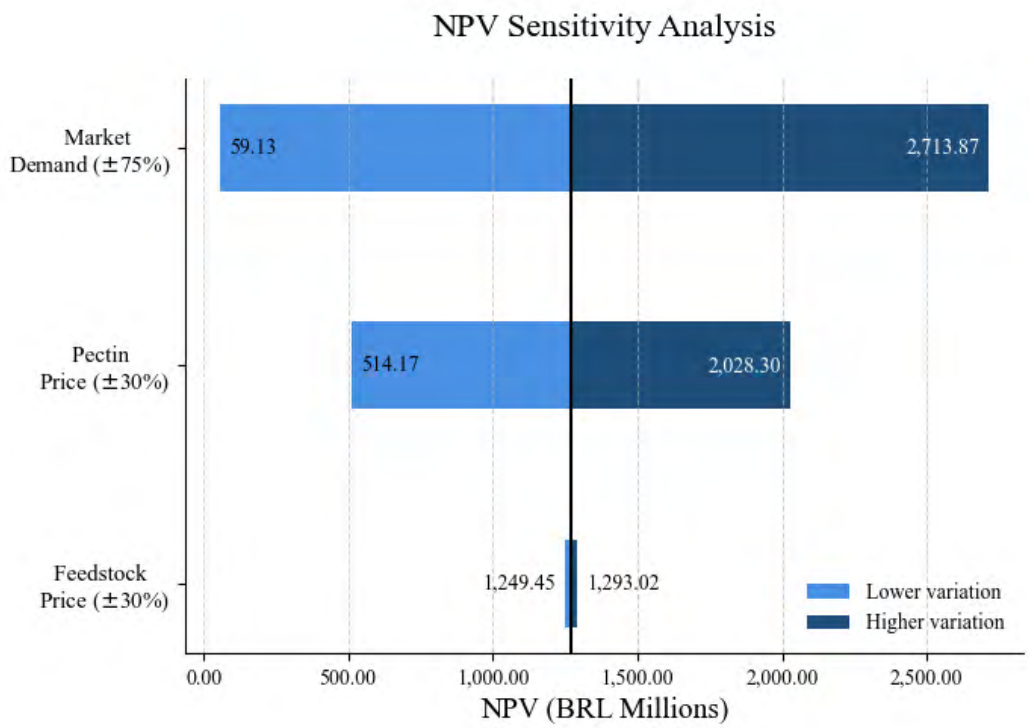

Figure 3. Tornado chart illustrating the sensitivity of the Net Present Value (NPV) to variations in market demand (±75%), pectin price (±30%), and raw material price (±30%).

million and a reduction in processing capacity from about 87, 676.06 mt/year to 21, 919.01 mt/year, with the optimal location shifting to Rio Claro. Conversely, a 75% increase in demand raises the NPV to approximately BRL 2.71 billion and increases processing capacity to about 153, 433.10 mt/year, with the model selecting Matão as the optimal location. In the base case, the selected location is Araraquara, with an intermediate capacity level. These findings show that demand influences both production scale and spatial configuration.

This behavior can be explained by the interaction between economies of scale and the spatial distribution of biomass resources. As demand increases, the model favors larger processing capacities, reducing specific investment costs and improving the dilution of fixed operating costs. Although transportation costs represent a small share of total operating costs, they influence spatial decisions due to differences in unit transportation costs. Due to its density, the cost of transporting orange peel residue (BRL 0.83 per mt.km) is higher than that of processed products such as pectin (BRL 0.22 per mt.km) and limonene (BRL 0.21 per mt.km). As a result, the model prioritizes locations with sufficient local biomass availability to supply the plant, reducing the need for long-distance feedstock transport.

At higher demand levels, the impact of economies of scale significantly favors centralized network configurations. The model tends to select a single, large-capacity facility positioned as close as possible to the port of Santos, provided that local biomass availability is sufficient to sustain the required throughput. This siting strategy reflects a fundamental trade-off of the system, where the framework prioritizes the exploitation of scale-related cost reductions while minimizing the transportation of low-density feedstock. In the context of São Paulo, the high regional concentration of residues allows the facility to achieve these economies of scale without triggering the need for a decentralized or market-oriented shift.

In contrast, feedstock-related parameters have a smaller impact on system performance. A ±30% variation in biomass purchase cost leads to changes in NPV between approximately BRL 1.25 billion and BRL 1.29 billion, without affecting plant location, technology selection, or processing capacity. This is consistent with the cost structure observed in the base case, where feedstock costs represent a relatively small share compared to production costs.

Alternative scenarios

In addition to the sensitivity analysis, a set of alternative scenarios was evaluated to further investigate the structural behavior of the system under specific operational and modeling constraints. These scenarios aim to

assess (i) the robustness of the network under reductions in biomass availability, and (ii) the conditions under which decentralized configurations may emerge due to capacity limitations.

To evaluate the robustness of the system under feedstock scarcity, scenarios with reductions of 30% and 50% in biomass availability were considered under base case demand conditions. The results show negligible changes in economic performance, with NPV remaining approximately constant at around BRL 1.27 billion in both cases. However, a shift in plant location is observed, from Araraquara in the base case to Matão under reduced availability conditions.

These results indicate that biomass availability, while not limiting overall system feasibility, can still influence spatial decisions when local supply becomes relatively more constrained. Under such conditions, the model exhibits patterns similar to those observed in high-demand scenarios, shifting facility locations toward regions with a higher concentration of biomass resources to ensure sufficient feedstock for the required processing capacity. Accordingly, biomass availability does not act as a binding constraint, but can still drive localized adjustments in the optimal network configuration.

To investigate scenarios in which decentralized configurations may emerge, a constraint was imposed on the maximum processing capacity of individual plants, limiting it to 75, 000.00 mt/year of biomass. Under base case demand conditions, the model continues to select a centralized configuration with a single facility in Araraquara, operating at the imposed capacity limit. In this case, total system capacity is reduced from approximately 87, 676.06 mt/year to 75, 000.00 mt/year, leading to a decrease in NPV from BRL 1.27 billion to approximately BRL 1.01 billion. Although the capacity constraint becomes binding, the model still favors a single-facility configuration, even at the expense of reduced throughput and economic performance. Under these demand conditions, the loss of economies of scale associated with splitting capacity across multiple plants outweighs the potential gains from increased product supply.

To further assess decentralization effects, a second scenario was evaluated under higher demand conditions (50% of domestic and 50% of external demand), maintaining the same capacity constraint. In this case, the required processing capacity exceeds the maximum allowable capacity of a single plant, leading to the installation of two facilities, located in Araraquara and Matão. The total installed capacity reaches approximately 146, 126.76 mt/year, with an NPV of approximately BRL 1.94 billion. Despite enabling full demand satisfaction, this decentralized configuration presents lower economic performance compared to the corresponding centralized configuration without capacity constraints. This reduction is primarily associated with the loss of economies of scale and the increase in total capital investment resulting from capacity splitting across multiple plants.

Comprehensively, decentralized configurations emerge only when capacity constraints become binding at higher demand levels.

General considerations on network configuration

The results obtained across the evaluated scenarios identify consistent structural patterns governing the optimal configuration of the supply chain. Rather than being driven by isolated parameters, facility siting and capacity decisions emerge from a complex interaction between market scale, regional biomass availability, and economies of scale.

Market demand plays a central role in determining both the scale and spatial orientation of the system. At lower demand levels, the model favors locations that balance proximity to major consumption centers with sufficient biomass supply. In this context, Rio Claro emerges as the preferred location, as it represents the most suitable area with significant biomass availability near key demand centers such as Araras, Campinas, Jundiaí, São Paulo, and Santos (external market), which together account for approximately 87% of total demand across all evaluated scenarios.

As demand increases, the system transitions toward plant locations capable of supporting larger processing capacities while remaining reasonably close to the primary consumption hubs. Under these conditions, Araraquara becomes the preferred location, as it represents the nearest region to the main demand centers with sufficient biomass availability to sustain higher throughput levels. This shift reflects the model's behavior of securing feedstock at scale as demand rises. Importantly, this tendency to source the full feedstock requirement from a single location is an emergent behavior of the system under the specific regional context rather than a pre-defined model constraint. In this study, such transitions are possible because biomass is highly abundant and spatially concentrated, allowing several candidate locations to independently supply the full feedstock requirement over a broad range of demand levels. Consequently, minimizing wet biomass transportation remains an important driver of location decisions, as it avoids the significantly higher transport efforts associated with feedstock compared to final products.

Under more restrictive or higher-demand conditions, alternative locations may also emerge. Matão, for example, is selected in specific scenarios due to its substantial biomass supply, stemming from its position as a major national hub for orange processing. This vast local availability allows the facility to sustain the large-scale operations required by high-demand levels without the need for decentralization. Although Matão is less

favorably located in relation to the primary consumption centers compared to Araraquara or Rio Claro, it becomes economically attractive when demand increases significantly, such as in scenarios exceeding the base case by 50% (representing approximately 60% of total combined demand). In these instances, the benefit of concentrated biomass availability at a single node offsets the logistical disadvantage of being further from the market hubs.

Furthermore, the optimization results exhibit a consistent structural preference for centralized configurations as long as individual plant capacity constraints remain non-binding. This behavior is a direct manifestation of the economies of scale embedded in both capital investment and process operations, which prioritize higher throughput levels to maximize the Net Present Value (NPV). Decentralized networks emerge only as a necessary alternative when capacity saturation prevents a single facility from meeting the aggregate demand. In such instances, the forced splitting of capacity across multiple sites leads to a measurable degradation of economic performance, as the system loses the significant cost-dilution benefits inherent to large-scale, centralized hubs.

Ultimately, the supply chain architecture is governed by a clear hierarchy of strategic drivers: market demand dictates the required industrial scale, regional biomass density defines the feasible spatial anchors, and economies of scale drive the system toward centralization. The dynamic interplay between these factors provides a robust explanation for the spatial transitions observed between Rio Claro, Araraquara, and Matão across the evaluated scenarios. These findings underscore that for high-value bioproducts like pectin, the citrus-producing region of São Paulo offers a resilient landscape where process-scale advantages outweigh logistical complexities, provided the network design remains aligned with these primary economic drivers.

CONCLUSIONS

A MILP-based optimization framework was developed and applied to assess the supply chain configuration of citrus residue biorefineries in the state of São Paulo, incorporating a multi-echelon logistics structure that captures both domestic and global market dynamics. The results demonstrate that the proposed biorefinery is highly profitable under the evaluated conditions, achieving a Net Present Value (NPV) of BRL 1.27 billion in the base case. This economic performance is fundamentally anchored by the high-value biochemical market, specifically pectin, which accounts for over 65% of total revenue.

Sensitivity and scenario analyses reveal that market scale and product pricing serve as the fundamental pillars of the supply chain architecture. As demand rises, the network exhibits a distinct preference for facilities sited at locations capable of satisfying the entire feedstock requirement locally, rather than prioritizing proximity to high-demand consumer centers. This structural pattern is intrinsically tied to the regional context of São Paulo, where the exceptional spatial concentration of residues allows the system to minimize the significant logistical effort of high-volume, wet biomass while maximizing economies of scale through centralized configurations. Furthermore, the comprehensive inclusion of multi-echelon logistics including inland transport, port handling, and international maritime freight reveals that total transportation costs represent a mere 1.13% of total expenditures. This marginal share is particularly striking when contrasted with production costs, which constitute the primary economic bottleneck of the system. Given the high utility and energy intensity of the extraction and recovery stages, the relative impact of product distribution is significantly diluted within the overall cost structure. Ultimately, the high value-density of products like pectin effectively absorbs global logistics costs, reinforcing that in regions with high biomass density, strategic success is governed by process efficiency and industrial scale rather than logistical constraints.

Building on these economic interdependencies, this research provides a robust baseline for strategic planning in the citrus bioeconomy. The findings indicate that the most resilient configurations are those that maximize throughput in centralized hubs to exploit non-linear investment scaling. While the system is logistically robust to biomass scarcity due to regional abundance, its long-term viability is strictly contingent upon process efficiency, especially regarding the management of production costs and utility consumption, as well as favorable market positioning for high-value bioproducts. This framework offers a comprehensive tool for sustainable waste valorization in major agro-industrial clusters, supporting the transition toward a circular bioeconomy.

ACKNOWLEDGEMENTS

This study was financed, in part, by the São Paulo Research Foundation (FAPESP), Brazil. Process Number #2024/17847-0.

REFERENCES

1. Floudas CA, Niziolek AM, Onel O, Matthews LR. Multi?scale systems engineering for energy and the environment: challenges and opportunities. AIChE Journal 62:602-623 (2016). https://doi.org/10.1002/aic.15151
2. Garcia DJ, You F. Supply chain design and optimization: challenges and opportunities. Computers & Chemical Engineering 81:153-170

(2015). https://doi.org/10.1016/j.compchemeng.2015.03.015
3. Bruschi LT, Kulay L, Santos MT. An optimization framework for renewable liquefied gas supply chain: a case study on glycerol-based production. Computers & Chemical Engineering 205:109481 (2026). https://doi.org/10.1016/j.compchemeng.2025.109481
4. Aba MM, Le Roux GAC, Menezes BC. Integrated ethanol and gasoline supply chain planning under environmental constraints: a case study of brazil. Computers & Chemical Engineering 164:107875 (2022). https://doi.org/10.1016/j.compchemeng.2022.107875
5. Macowski DH, Bonfim-Rocha L, Orgeda R, Camilo R, Ravagnani MASS. Multi-objective optimization of the brazilian industrial sugarcane scenario: a profitable and ecological approach. Clean Techn Environ Policy 22:591-611 (2020). https://doi.org/10.1007/s10098-019-01802-0
6. Garcia VF, Palacios R, Ensinas A. Optimisation of ammonia production and supply chain from sugarcane ethanol and biomethane: a robust mixed-integer linear programming approach. Processes 12:2204 (2024). https://doi.org/10.3390/pr12102204
7. Turton R, Shaeiwitz JA, Bhattacharyya D, Whiting WB. *Analysis, Synthesis, and Design of Chemical Processes*. 5th ed. Pearson Education (2018)
8. Vielma JP, Ahmed S, Nemhauser G. Mixed-integer models for nonseparable piecewise-linear optimization: unifying framework and extensions. Operations Research 58:303-315 (2010). https://doi.org/10.1287/opre.1090.0721
9. Suzuki LL, Cunha ILC, Teixeira CE, Kulay L, dos Santos MT. The modeling and simulation of waste recovery processes applied to the orange juice industry. Sustainability 15:15225 (2023). https://doi.org/10.3390/su152115225
10. Marchette IR, Cavallari RBC, Rothier FDC, Bojorge NIBR, Alhadeff EM, Young AF, dos Santos BF. Economic assessment and process design to valorize citrus juice industry waste. Biofuels Bioprod Bioref 19:616-633 (2025). https://doi.org/10.1002/bbb.2722
11. CitrusBR – Associação Nacional dos Exportadores de Sucos Cítricos. CitrusBR. (2026). https://citrusbr.com/
12. Vielma JP, Ahmed S, Nemhauser G. Mixed-integer models for nonseparable piecewise-linear optimization: unifying framework and extensions. Operations Research 58:303-315 (2010). https://doi.org/10.1287/opre.1090.0721
13. Agência Nacional de Transportes Aquaviários (ANTAQ). Final report – comparative study of THC (terminal handling charge) prices in container terminals in Brazil and worldwide. ANTAQ (2019). https://sophia.antaq.gov.br/terminal/acervo/detalhe/27610?guid=163891118832
14. Câmara de Comercialização de Energia Elétrica (CCEE). CCEE. (2026). https://www.ccee.org.br/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Logistics Management of Agri-Industrial Waste for Energy Valorization in Uruguay

Milena Lagarmilla[a], Ivan Guchin[a]*, Mauro Gambetta[a], Darío Huelmo[b], Adrián Ferrari[a], Soledad Gutierrez[ab]

[a] Universidad de la República, Chemical and Process Systems Engineering Group – Chemical Engineering Institute - Engineering School, Montevideo, Uruguay
[b] Universidad de la República, Industrial Project Department – Chemical Engineering Institute - Engineering School, Montevideo, Uruguay
* Corresponding Author: ivanguchin1@gmail.com.

ABSTRACT

The energy recovery of agro-industrial residual biomass offers a pathway to reduce fossil fuel emissions in thermal processes while valorizing waste. In practice, however, the primary bottleneck is logistical: feedstocks are geographically dispersed, with low bulk density and high moisture content, driving up collection, pretreatment, and transport costs. This work combines geospatial processing with mathematical optimization to design a multi-stage logistics network. The model incorporates intermediate densification options and technology selection (chipping, pelletizing, or briquetting) to supply one or more final waste-to-energy plants. The case study focuses on Northeastern Uruguay, considering forestry residues, meat-processing waste, and rice husks. We formulate a multi-period Mixed-Integer Linear Programming (MILP) model aimed at minimizing the total annualized cost, encompassing transportation, logistical operations, capital investment, and plant O&M, subject to supply constraints, capacities, and maximum transport distances. Results show that the optimal configuration corresponds to the briquetting scenario, which includes a fixed densification facility and processes 200, 160 t/y of biomass from 18 sawmills. Under this configuration, the system generates 276, 954 MWh, with a net margin of USD 5.77/MWh and a net profit of USD 1, 596, 662 over the planning horizon. Sensitivity analysis indicates that the electricity selling price is a critical parameter, with a tolerance margin of only 8%, while mobile densification is not economically viable under current market conditions. Beyond the specific case study, the main contribution of this work is the optimization model itself, whose general and configurable formulation allows its application to other regions, residue types, and market conditions without reformulating its mathematical structure.

Keywords: Biomass, Supply Chain, Energy Systems, Modelling and Simulations, Optimization, Technoeconomic Analysis

INTRODUCTION

The energy valorization of agro-industrial residues is a key strategy for industrial decarbonization, yet its implementation is often limited by logistical costs. Because of dispersed biomass availability, low bulk density, and high moisture content, the expenses associated with collection and transport can be prohibitive. While densification addresses these issues by increasing volumetric energy density, it introduces significant design challenges: determining where densification should occur, at what scale, and how to optimize biomass flows to the end-user. Optimizing these interrelated decisions requires a spatially explicit framework that integrates road-network data with operational and geographic constraints.

This work proposes a multi-stage logistics network design for residual biomass with intermediate densification. The model is applied to the Tacuarembó–Rivera corridor in Uruguay, evaluating pre-treatment alternatives to identify cost-efficient configurations for facility locations, capacities, and flow allocations. Regarding the state of the art, O'Neill, Geissler, and Maravelias (2024)

Residue	Generation rate (t/year)	Recoverable fraction (%)	Heating value (kcal/kg w.b.)	Bulk density (kg/m³)	Moisture content (% w.b.)	Harvest period / Seasonality
Rice husk	261,200	80%	3,520.9	125	8.50%	February – June
Sawmill residues	412,027	54.40%	4,269.0	275	15.5%	Annual
Rumen content	81,195	100%	2,153.0	770	76–82%	Annual

Table 1. Characteristics for the studied residues

integrated logistics networks with biorefinery location and design, demonstrating high sensitivity to spatial variables and incentives. O'Neill et al. (2025) compared bioenergy landscape configurations and logistics networks under transport constraints and mitigation targets, while Zervopoulou et al. (2025) formulated a spatially explicit MILP model to locate plants and allocate flows in biomass-to-fuel supply chains. Similar to our work, these studies address network design and highlight the dependence on resource geography and logistical constraints. Our contribution differs by including intermediate densification as a core design decision, incorporating mobile alternatives with period-by-period relocation, pre-treatment technology selection, and a multi-period planning horizon.

CONTEXT

Residual biomass can be understood as any organic matter of animal or plant origin that can be valorised for energy or material recovery. This classification facilitates the analysis of agro-industrial residues with different properties and seasonal patterns, all sharing a common objective: their transformation into energy carriers or chemical feedstock.

Nevertheless, the existence of a residue does not guarantee its logistical feasibility. Therefore, this study narrows its scope to three specific streams characterized by high availability, concentration and low current market value: sawmill residues, slaughterhouse rumen content, and rice husks

Sawmill residues

Sawmill operations produce byproducts such as bark, trimmings, and wood shavings, which are characterized by low moisture content and relatively concentrated generation points. National surveys estimate annual production of approximately 412, 000 tonnes on a dry basis (Gutiérrez et al., 2017).

Slaughterhouse rumen content

Bovine rumen content stands out among slaughterhouse waste due to its high volume, estimated at 36 kg per head. Its primary logistical constraint is a high moisture content (80%), which significantly inflates transportation costs and limits conditioning alternatives. In 2024, Uruguay processed 2, 256, 123 cattle (INAC, 2024), consolidating beef as the country's primary slaughter category. A local benchmark is the ONTILCOR meat-packing plant, where rumen content is co-fired in boilers alongside firewood without additional drying. This facility reportedly valorizes its entire waste generation (440 t/month), achieving a substitution ratio of 2 kg of residue for every 1 kg of firewood while maintaining steady steam production.

Rice husk

Generated during the dehulling process, rice husk is concentrates at milling facilities, which provides favorable collection logistics. This residue represents represents roughly 20% of the grain weight and has a reported HHV of 3, 300 kcal/kg (Faroppa, 2010). Its main drawback is a low bulk density (120–130 kg/m³), making the evaluation of densification alternatives essential (Díaz Tovar, 2019). In Uruguay, irrigated rice cultivation maintains average yields of 9, 000 kg/ha, with a harvested area of nearly 150, 000 hectares during the 2023/2024 season (MGAP-OPYPA, 2024).

Table 1 summarizes the key characteristics of the three residues, which are used to parameterize availability and the conditioning and densification alternatives within the model.

Densification

The model evaluates three densification alternatives. None of these densification alternatives are compatible with rumen content due to its high moisture content.

Shredding/Grinding

This is included as a lower-complexity option, though it yields lower improvements in density, heating value, and moisture reduction. It is a particularly relevant technology for processing sawmill residues.

Pelletizing

A compaction process that transforms biomass into

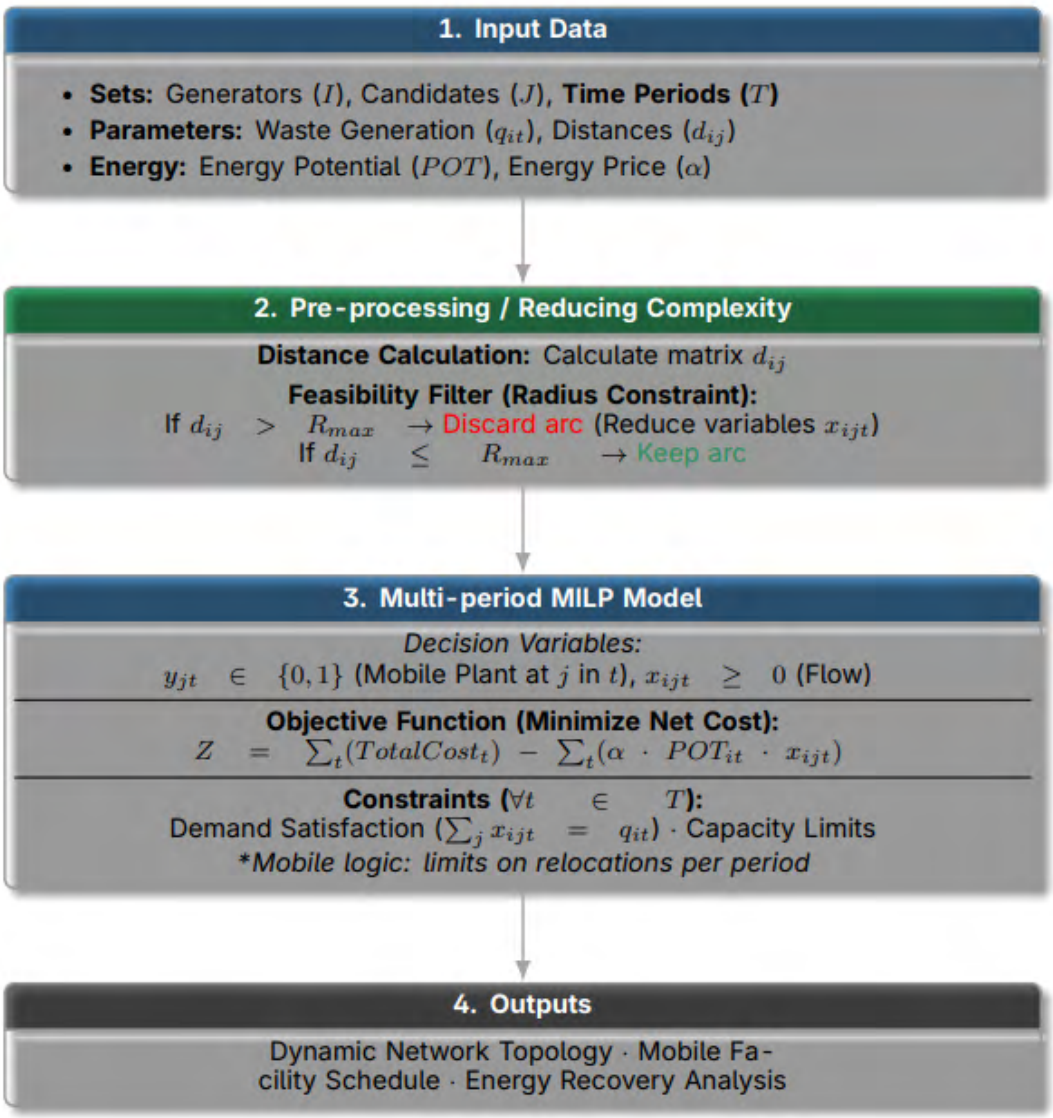

Figure 1. Structure of the multi-period optimization framework for mobile plant location.

pellets, significantly increasing both bulk density and volumetric energy density. This alternative requires higher capital investment and more rigorous operational control than the others, often necessitating strict feedstock moisture ranges to achieve acceptable quality standards (Hernández et al., 2017).

Briquetting

A compaction technology similar to pelletizing but resulting in a larger final product. The mechanical quality of the densified material depends heavily on compaction pressure and the specific properties or blending of the biomass. These factors are critical for technology selection and the definition of operating parameters (Forero-Núñez et al., 2014).

STUDY CASE

Based on the initial residual biomass survey in Uruguay, a nationwide georeferenced database was developed. This database includes all surveyed meat-packing plants, sawmills, and rice mills, recording specific coordinates and monthly residue generation for each facility. The next step involved identifying regions characterized by both high concentrations of generation points and the logistical viability required to supply a centralized plant. To achieve this, a computational procedure was implemented to analyze spatial clusters within a predefined

radius, calculating the aggregate energy potential by residue type and prioritizing areas with the highest supply density.

The clustering radius was defined based on energy viability criteria for transport: for each residue, the energy content of a fully loaded truck was compared to the energy consumed during transit, yielding a maximum feasible transport distance specific to each material. Among the identified regions, the Tacuarembó-Rivera corridor was selected as the primary case study, with its logistical axis centered on the national highway traversing the area.

Once the region was defined, potential land parcels for an energy valorisation plant were evaluated. Key watercourses within the study area were mapped, and candidate parcels were pre-selected using a weighted multi-attribute scoring scheme. In general terms, the attributes considered included road accessibility for heavy-duty vehicle traffic, proximity to water sources, available water flow, and topographical characteristics. Based on this assessment, the five highest-scoring parcels were selected as candidate sites for the model.

Network Structure

The study case is defined by the following components, which constitute the input data for the mathematical model developed in the following section. A total of 22 georeferenced generation points are considered: 18 sawmills, 2 rice mills, and 2 slaughterhouses, located in the northeastern corridor of Uruguay between the departments of Tacuarembó and Rivera, with National Route 5 as the main logistical axis. The logistics network includes 6 candidate locations for the final energy recovery plant, 16 candidate nodes for fixed densification plants, and up to 4 mobile densification units. The three densification alternatives evaluated are shredding, pelleting, and briquetting, while the selected valorization pathway is electricity generation through biomass combustion in a steam boiler coupled to a power generation turbine. Biomass transport is carried out using standard trucks.

METHODOLOGY

The logic of the proposed model is illustrated in Figure 1. The framework is structured as a multi-period Mixed-Integer Linear Programming (MILP) problem, designed to capture the temporal variability of seasonal feedstock supply and to allow for the relocation of mobile plants between periods according to the spatial availability of resources. The model considers three decision levels: generation points ($p \in P$), fixed ($n \in N$) or mobile ($k \in K$) densification plants, and final energy plants ($f \in F$), across time periods ($t \in T$). It incorporates 35 parameters describing costs, capacities, distances, process yields, energy consumption, operational times, and transport factors, assuming trucking as the sole transport mode available in the study region.

The primary decision variables are:

- tp, t- tons of biomass utilized at generation point *p* in period *t* [ton]
- zp, n, f, t binary variable; 1 if residues from generation point *p* are densified at location *n* and sent to final plant *f* in period *t* ; 0 otherwise.
- mk, n, t binary variable; 1 if mobile densification unit kkk is operating at location nnn in period ttt, and 0 otherwise.
- yp, f, t binary variable; 1 if residues are sent directly from generation point *p* to final plant *f* in period *t*; 0 otherwise.
- un, t binary variable; 1 if a fixed densification plant is activated at location nnn, and 0 otherwise.
- dd binary variable; 1 if densification method ddd is selected, and 0 otherwise.
- ff binary variable; 1 if the final energy plant is activated at location fff, and 0 otherwise.

The objective function (see Figure 1) seeks to minimize the total net system cost while accounting for the energetic valorization of the available biomass. The formulation includes the revenues associated with energy production together with the annualized costs of the supply chain. These costs comprise transportation and logistical operations, truck fleet requirements, mobile plant relocation, and the CAPEX and O&M of both densification and final waste-to-energy facilities. As a result, negative objective-function values indicate economically viable configurations in which energy revenues exceed total system costs over the planning horizon.

The model incorporates a comprehensive set of balance, capacity, and technological compatibility constraints to ensure the physical and operational consistency of the network. This includes mass balances at each node and period—accounting for yields and losses associated with densification—alongside capacity limits for fixed plants, mobile units, and the transport fleet. Spatial constraints define maximum allowable distances based on the type and state of the residue. Furthermore, the formulation enforces exclusivity and unique flow allocation conditions, as well as technological compatibility rules; specifically, only sawmill and rice mill residues are eligible for densification, while rumen content is restricted to drying. These constraints guarantee system feasibility and accurately capture the interplay between transport, densification, and energy recovery.

The model was implemented in Python 3.13 using the Pyomo algebraic modeling framework and the HiGHS

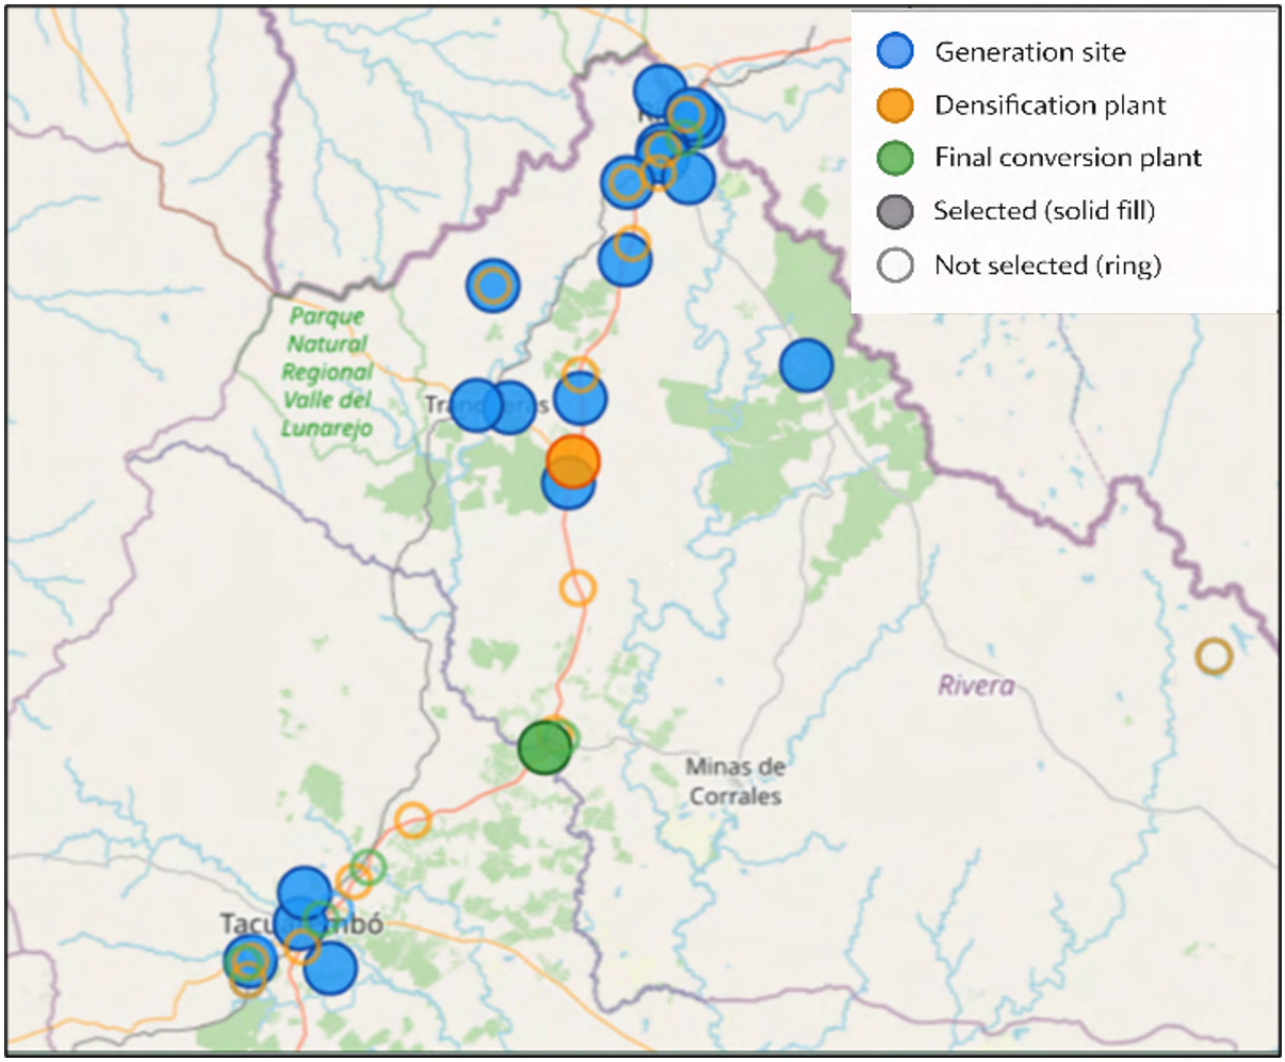

Figure 2. Location of the selected plants and active generation points.

MILP solver. Input data are provided through Excel files containing biomass generation points, candidate plant locations, and model parameters. Based on these data, the optimization model is executed to provide the following outputs: the selection of collection points, optimal flow allocation, the location and activation of both densification and final plants, and an estimation of the total energy potential.

Results

The optimal solution corresponds to the briquetting scenario, in which the model selects the opening of one fixed densification plant and one final energy recovery plant. The resulting logistics network combines direct flows to the final facility with flows requiring an intermediate densification stage, depending on distance and overall network configuration. Under this solution, the system processes 200, 160 tons of biomass per year from 18 active sawmills and generates 276, 954 MWh over the annual planning horizon. The objective function reaches −USD 1, 596, 662, indicating that revenues from energy sales exceed total system costs. Mobile densification is not economically viable under current market conditions.

CONCLUSION

This work developed and applied a mathematical optimization model to evaluate logistics strategies for the energy valorization of residual biomass in Uruguay. A key finding is that densification is not justified primarily by its direct energetic benefit, but by its logistical role within the supply chain: it expands the effective reach of the network and enables the incorporation of biomass that would otherwise remain outside the feasible system.

The analysis also shows that transport, although not dominant in the base-case cost structure, becomes much more influential as biomass sources become more

geographically dispersed. In addition, the economic viability of the system is highly sensitive to the electricity selling price and to the conversion efficiency of the final plant, with a tolerance margin of only 8%. Under the evaluated conditions, mobile densification does not constitute a viable alternative due to its higher operating cost and limited processing capacity. Beyond the specific results of the case study, the main contribution of this work lies in the optimization model itself, whose general and configurable formulation makes it applicable to other regions, residue types, and market conditions without requiring changes to its mathematical structure.

REFERENCES

1. Böthig S, Arrejuria S, Bonfiglio F, Cagno M, Delgado Y, Martínez S, Rey F, Cabot P, Martínez D. By-product valorization project for the mechanical wood processing industry: final report. Latitud–LATU (2021). https://www.gub.uy/ministerio-industria-energia-mineria/publicaciones/informe-general-final
2. Carroll M, Borges M. Towards a green economy in Uruguay: stimulating sustainable production practices and low-emission technologies in prioritized sectors. UNIDO (2021).
3. Pena vergara G, Durante A, Curto-Risso P, TANCREDI N, Amaya A, Pina Riveiro AC, Franco E. Characterization of residual biomass from agricultural and agroindustrial activities. Procceedings of the 24th ABCM International Congress of Mechanicl Engineering : (2017). https://doi.org/10.26678/abcm.cobem2017.cob17-0668
4. Díaz Tovar D. Potential uses of rice husk in the department of Casanare. UNAD Repository (2019).
5. Faroppa C. Evaluation of the availability of biomass residues or by-products at a national level. Ministry of Industry, Energy, and Mining, National Energy Directorate (2010). https://www.gub.uy/ministerio-industria-energia-mineria/publicaciones/evaluacion-de-la-disponibilidad-de-residuos-o-subproductos-de-biomasa
6. Forero-Núñez CA, Camargo-Vargas G, Sierra-Vargas FE. Modelos de compresión aplicados al proceso de densificación de combustibles sólidos binarios carbón-madera. ITECKNE 11: (2014). https://doi.org/10.15332/iteckne.v11i2.733
7. Gutiérrez S, Philippi C, Kreimerman R, Ures P, Torres AI. Survey of available biomass in Uruguay usable as feedstock for chemical production. Paper presented at the VI Regional Meeting of Chemical Engineering (2017).
8. Hernandez Solorzano LC, Forero Nuñez CA, Sierra FE. Biomass densification: a review of the market and recent r&d trends. TECCIENCIA 12:81-92 (2017). https://doi.org/10.18180/tecciencia.2017.23.10
9. INAC – National Meat Institute. 2024 statistical results: slaughter and exports of the meat sector (2025). https://www.inac.uy/innovaportal/v/27775/37/innova.bs/exportacion
10. MGAP-OPYPA. Rice (Chapter 11). In: 2024 OPYPA Yearbook. Ministry of Livestock, Agriculture, and Fisheries, MGAP (2024).
11. MIEM. Biovalor demonstration projects: Ontilcor. Ministry of Industry, Energy, and Mining (2025). https://www.gub.uy/ministerio-industria-energia-mineria/comunicacion/publicaciones/proyectos-demostrativos-biovalor-ontilcor
12. MIEM. Plants in operation: Galofer. Ministry of Industry, Energy, and Mining, National Energy Directorate (2022). https://www.gub.uy/ministerio-industria-energia-mineria/publicaciones/plantas-operacion-galofer
13. Nunes LJR, Matias JCO, Loureiro LMEF, Sá LCR, Silva HFC, Rodrigues AM, Causer TP, DeVallance DB, Ciolkosz DE. Evaluation of the potential of agricultural waste recovery: energy densification as a factor for residual biomass logistics optimization. Applied Sciences 11:20 (2020). https://doi.org/10.3390/app11010020
14. O'Neill EG, Geissler CH, Maravelias CT. Large-scale spatially explicit analysis of carbon capture at cellulosic biorefineries. Nat Energy 9:828-838 (2024). https://doi.org/10.1038/s41560-024-01532-8
15. O'Neill EG, Lark TJ, Xie Y, Kim S, Dale BE, Robertson GP, Basso B, Maravelias CT. Available land for cellulosic biofuel production: a supply chain centered comparison. Environ. Res. Lett. 20:014067 (2025). https://doi.org/10.1088/1748-9326/ada301
16. Zervopoulou S, Papadokonstantakis S, J?rvinen M, Inayat M. Lignocellulosic waste supply chain network design for sustainable aviation fuels production through solar pyrolysis. Systems and Control Transactions 4:2019-2025 (2025). https://doi.org/10.69997/sct.152432

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Section 8

Process Control and Operation

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Design and Control of Heat Pump Assisted Distillation Processes for Flexible E-methanol Production

Lucas A.T. Poker[a], Marija Saric[b], Jan Wilco Dijkstra[b], Vladimir Dikic[b], Anton A. Kiss[a*]

[a] Department of Chemical Engineering, Delft University of Technology, Van der Maasweg 9, 2629 HZ, Delft, The Netherlands
[b] Sustainable Technologies for Industrial Processes, Energy Transition, TNO, Petten, The Netherlands
* Corresponding author: L.A.T.Poker@tudelft.nl.

ABSTRACT

This study investigates control strategies for the flexible operation of heat pump-assisted distillation processes, focusing on the heat integrated distillation column configuration. The methanol/water separation system was selected as a case study and modelled to achieve 99.9 wt% AA-grade methanol purity. A limiting piece of equipment for flexible operation of heat pump assisted distillation is the compressor. To assess its impact on flexible operation, dynamic simulations in Aspen Dynamics were conducted for two heat integrated distillation column control strategies: one using fixed compressor duty and one using variable compressor duty. The control performance for a 20% throughput disturbance, as well as for a 50% turndown ratio scenario was investigated. Results show that fixed-duty operation maintains robust stability and rapid disturbance recovery even at 50% turndown, while variable-duty operation delivers higher efficiency for moderate load changes but cannot sustain low-load stability. This work supports the electrification of distillation by enhancing the operational flexibility of heat pump-assisted distillation, enabling better integration with intermittent renewable electricity grids.

Keywords: Distillation, Energy Efficiency, Dynamic Modelling, Aspen Dynamics, Methanol

INTRODUCTION

E-methanol produced from green hydrogen and captured CO_2 can be used as a versatile chemical building block, transportation fuel, and steel carburizing agent, making it central to the transition toward circular economy [1]. Heat pump assisted distillation (HPAD) enables fully electric purification using compressor shaft work, while significantly lowering energy consumption compared to conventional distillation. A common used HPAD configuration is mechanical vapor recompression (MVR) [2]. Here, the pressure of the overhead vapor of the column is increased such that its latent heat can be used to drive the reboiler. However, the application of MVR is limited to separations of close boiling point mixtures (ΔT_b = 20 °C), as large compression ratios are needed to overcome the entire temperature lift over the column [2]. Due to the large boiling point difference for methanol-water (ΔT_b = 34.7 °C) the heat-integrated distillation column (HIDiC) becomes an attractive HPAD technology. HIDiC reduces the required temperature lift and compression ratio by coupling the whole rectifying and stripping sections via internal heat exchange [3]. Nevertheless, very few HIDiC configurations exist at plant scale due to its complex design and structure [3]. Wakabayashi et al. [4] proposed a discretely heat-integrated distillation column (D-HIDiC), removing the need for impractical continuous heat exchange between stripping and rectifying section of HIDiC. This resulted in the first commercial HIDiC. Commercial operation has shown that D-HIDiC can increase control robustness of HIDiC, while maintaining high energy efficiency [4]. Cui et al. have investigated the design of a fully electrified D-HIDiC for the separation of methanol-water [5]. Previous studies on HIDiC process control proposed control structures based on expensive online composition analysers [6]. Cui et al. proposed a single end inferential composition control strategy based on temperature control for D-HIDiC [7]. This strategy manipulates compressor duty for control and showed robust performance under ±20% disturbances in throughput and feed composition. Recently, flexible operation of electrified processes to accommodate variations in electricity prices has gained increased attention, but this has not been investigated for HIDiC.

PROBLEM STATEMENT

E-methanol plants could operate flexibly to reduce costs by adapting to fluctuating electricity prices, as green hydrogen production is the main cost driver. To prevent the need for large hydrogen holdups, downstream processes could be operated flexibly. In this work, D-HIDiC is chosen as a technology of interest for full electrification of e-methanol purification due to its significant energy savings of up to 80%. Flexible operation of D-HIDiC is challenging, as D-HIDiC's strong thermal integration introduces complex process dynamics: disturbances in feed throughput or composition cause rapid, system-wide shifts in temperature, pressure, and composition profiles [8]. These tightly coupled responses challenge conventional control strategies, especially under flexible operation with frequent turndowns of up to 50%. While Cui et al. demonstrated control with variable compressor duty [7], their strategy requires the compressor to operate over a wide range when operating flexibly (e.g. 50-120% of design capacity), which is often infeasible due to compressor surge/stall limits. They mention that the modular nature of compressor systems could be used to enhance flexibility, but using this approach would increase CAPEX as well as control complexity. Therefore, this study aims to explore an alternative control strategy for D-HIDiC using fixed compressor operation, which is capable of operation at high turndown ratios of up to 50% to enable flexible operation.

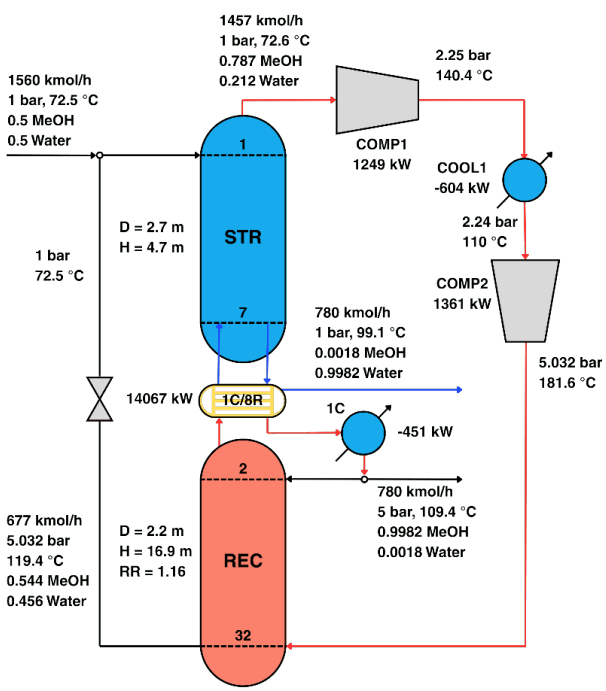

Figure 1. D-HIDiC configuration studied

METHODOLOGY

Process description

The HIDiC studied in this paper is shown in Figure 1 and is a modified version of the design proposed by Cui et al [5]. The following thermodynamic and hydrodynamic models are used for the process simulations in Aspen Dynamics, based on the work of Cui et al.:

- The thermodynamic model used is non-random two-liquid (NRTL), which is suitable and validated for the methanol water system [9].
- Column pressure drop was based on rigorous calculations using Aspen Plus column internals. To ensure high energy efficiency, low pressure drop Mellapak 350Y packing is used for the column internals.
- The hydrodynamic model is based on equilibrium stage, which is used in various process dynamics studies [10]. It is therefore assumed that each stage reaches vapor liquid equilibrium.
- Liquid holdup is estimated based on design heuristics by Luyben, ensuring a 5 minute holdup time at 50% full [10].

This design is a D-HIDiC, which features heat integration between the reboiler of the stripping section and the condenser in the rectifying section to realize full electrification of the column. Instead of utilizing the compressor interstage cooling duty for heat integration with the column, an interstage cooler is used for improved dynamic performance at a slight decrease in energy efficiency. Still, this design reduces the primary energy consumption compared to a conventional distillation column (CDiC) by 81.6% and has a heat pump coefficient of performance (COP) of 5.39. The D-HIDiC is designed to separate the effluent of an e-methanol reactor consisting of equimolar amount of methanol and water up to an AA grade methanol purity of 99.85 %wt [11]. Due to the maximum allowed mass fraction of water in AA grade methanol, the final purity was set at 99.9 %wt. A 200 kt/y methanol capacity is chosen based on the capacities of planned hydrogen electrolysers in Europe [12].

The D-HIDiC consists of 40 stages, split into an 8-stage low pressure stripping column (STR) and a 32-stage high pressure rectifying column (REC). A compression ratio of 5 was chosen to raise the boiling point of methanol in the high-pressure rectifying column, ensuring a minimum ΔT of 10 °C in the reboiler/condenser heat exchanger. A packed height of 0.5 meter per stage is used. A 20% column cross-section increase is applied to account for factors such as liquid holdup for surge capacity, based on recommendations by Luyben [10]. Furthermore, isentropic compressors with an isotropic

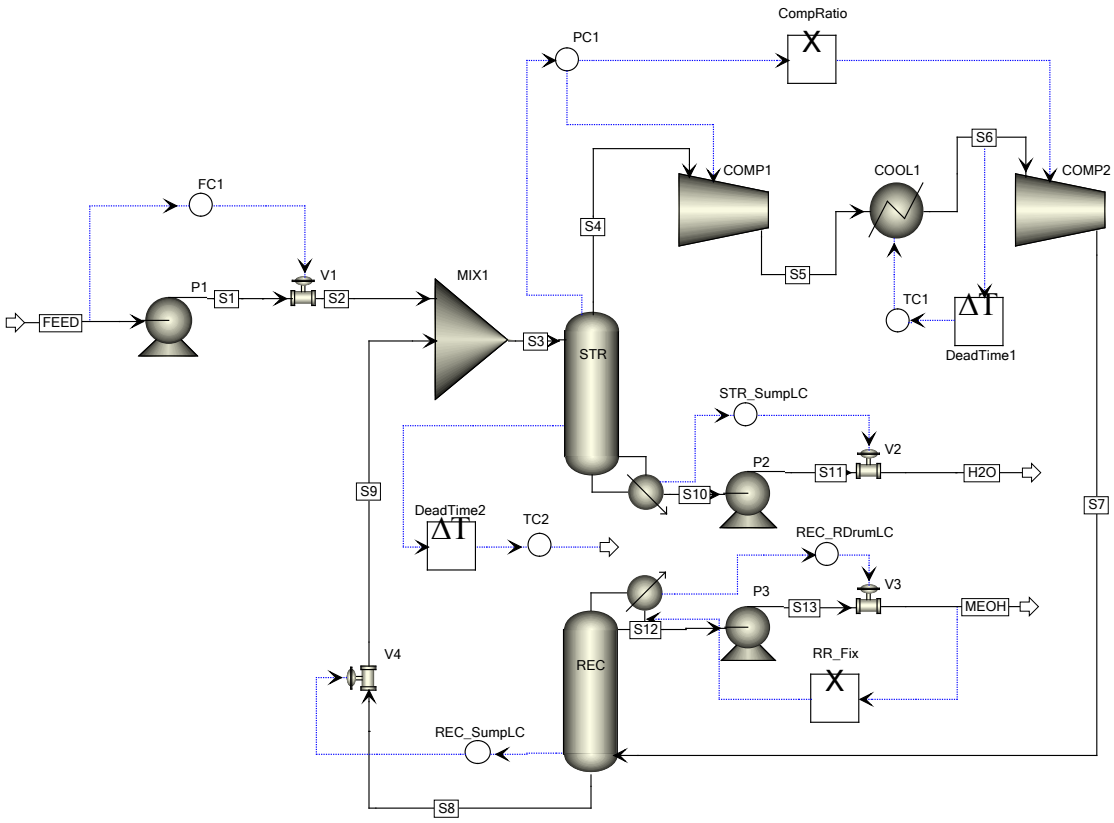

Figure 2. Process control structure 1 (PCS 1) utilizing variable compressor duty

efficiency of 80% were used.

Dynamics and control

Two control structures were evaluated in this work: one employing variable compressor duty based on the work by Cui et al. [7], and one novel control structure utilizing fixed compressor duty. The dynamic performance of both control structures was evaluated for ±20 % throughput disturbances, typical for industrial operation [10]. Composition disturbances were not investigated, as the effluent of a CO_2 hydrogenation methanol reactor typically has minor variations in composition. A sensitivity study was done on compressor mass throughput, since compressors have a limited range of operation. The sensitivity study was done for turndown ratios up to 50%, to account for a scenario with high energy prices and/or reduced availability of hydrogen to produce methanol.

Control basis

The process control structures were modelled using pressure driven simulations in Aspen Dynamics V14. Control valve pressure drops are 3 bar with the valve half open at design flow rates. PI controllers are used, with flow controllers having a K_c of 0.5 and τ_i of 0.3 min, the level controllers a K_c of 2 and τ_i of 9999 min based on recommendations by Luyben [10]. The other controllers were tuned using Tyreus-Luyben tuning rules. Temperature controllers are set at a 1 minute deadtime. Stage 5 of the stripper is identified as the most temperature sensitive stage based on the slope criterium [10]. Condenser and reboiler heat integration is done using flowsheet equations (Figure 4 & 5). Reboiler duty is calculated using (1), and then total condenser duty is calculated by adding auxiliary condenser duty to the reboiler duty (2). UA is kept constant. Condenser and reboiler duties then have to be changed from 'fixed' to 'free' in Aspen Dynamics.

$$Q_{reb} = UA(T_{cond} - T_{reb}) \quad (1)$$

$$Q_{cond} = -Q_{reb} + Q_{aux\ cond} \quad (2)$$

Process control structure 1 (PCS 1) is shown in Figure 2 and uses the following control loops:

- Fresh feed is flow controlled (FC1).
- L_{sump} is controlled using bottoms flowrate for both columns (STR_SumpLC & REC_SumpLC).
- $L_{refluxdrum,\ REC}$ is controlled using distillate flowrate (REC_RDrumLC)
- $T_{interstage}$ is controlled by manipulating duty of interstage cooler 'COOL1' (TC1).

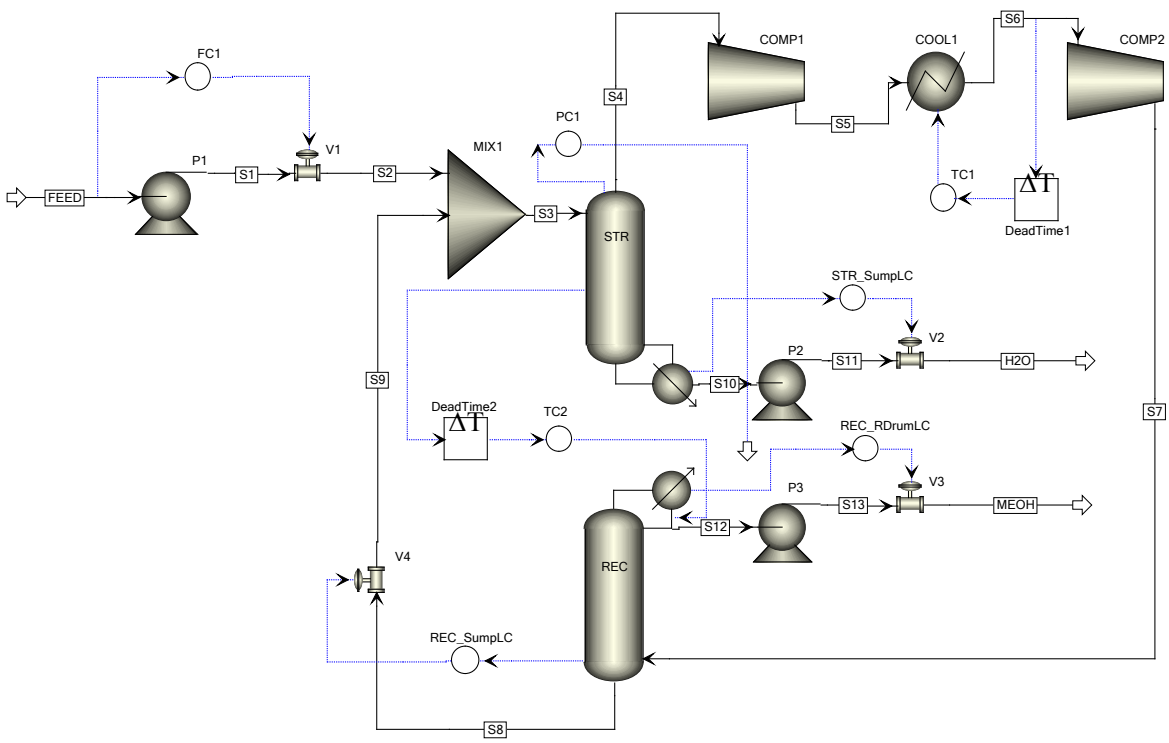

Figure 3: Process control structure 2 (PCS 2) utilizing fixed compressor duty

Constraints - Flowsheet

```
CONSTRAINTS
  // Flowsheet variables and equations...
  blocks("STR").Qreb = 4.944989695 * (blocks("REC").Stage(1).T-blocks("STR").Stage(8).T);

  blocks("REC").Condenser(1).Q = -blocks("STR").QReb + blocks("TC2").op;
```

Figure 4. Flowsheet equations PCS 1

Constraints - Flowsheet

```
CONSTRAINTS
  // Flowsheet variables and equations...
  blocks("STR").Qreb = 4.944989695 * (blocks("REC").Stage(1).T-blocks("STR").Stage(8).T);

  blocks("REC").Condenser(1).Q = -blocks("STR").QReb + blocks("PC1").op;
```

Figure 5: Flowsheet equations PCS 2

- $T_{stage\ 5,\ STR}$ is controlled by manipulating auxiliary condenser duty, see Figure 4 (TC2).
- $P_{stage\ 1,\ STR}$ is controlled by manipulating compressor duty, using a ratio controller to evenly change both compressor duties (PC1).
- P & T in the rectifier floats to allow for effective heat integration between rectifier and stripper
- Reflux ratio is kept constant using a ratio controller setting reflux flow as a ratio of distillate flow.

Process control structure 2 (PCS 2) in Figure 3 uses the same flow, interstage temperature, and level control loops, but has different control loops for stripper T and P:

- $T_{stage\ 5,\ STR}$ is controlled by manipulating reflux flow (TC2).
- $P_{stage\ 1,\ STR}$ is controlled by manipulating auxiliary condenser duty, see Figure 5 (PC1).

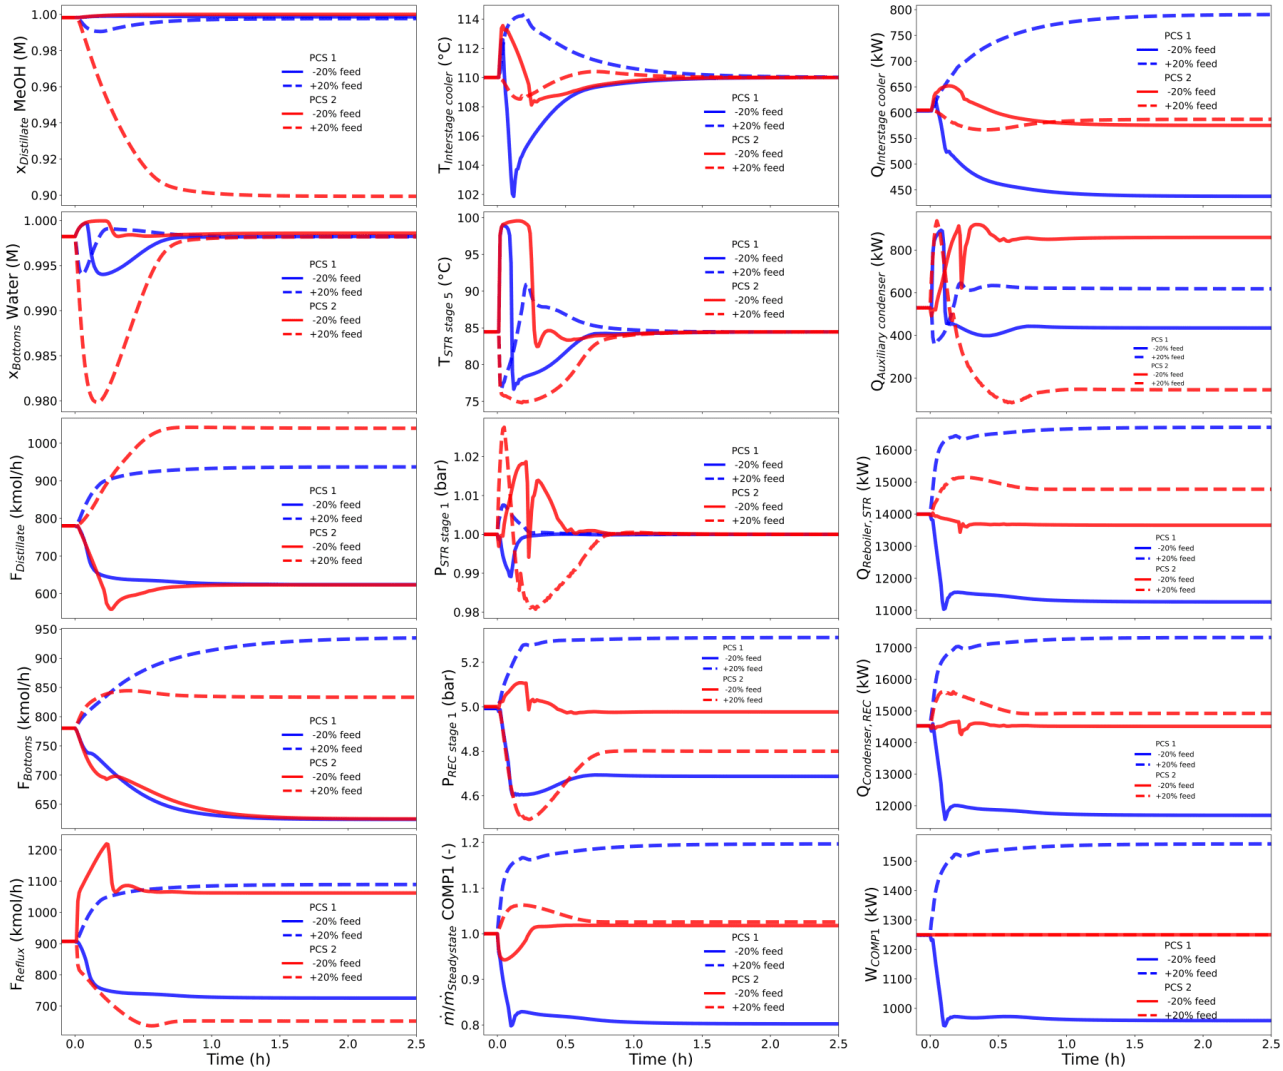

Figure 6. Control responses of PCS 1 and PCS 2 for ±20 % feed throughput disturbances

- Reflux ratio is no longer fixed, but instead compressor duty is kept constant.

The closed-loop system is continuously disturbed after 0.1 hr, for a total of 2.5 hr. The results are shown in Figure 6.

Compressor sensitivity analysis

A compressor feed disturbance sensitivity analysis was done by finding the maximum change of corrected mass throughput for the compressor for step change feed disturbances ranging from 0.5 feed fraction to 1.2 feed fraction. For each step change feed disturbance, a separate simulation run in Aspen Dynamics was done using a continuous disturbance for 2.5 h, after which the maximum change of corrected compressor mass throughput was found. The compressor mass throughput is corrected based on (3) to account for the effect of inlet temperature and pressure changes:

$$\dot{m}_{corr} = \dot{m}\sqrt{\frac{T_{in}}{T_{ref}}} * \frac{P_{ref}}{P_{in}} \tag{3}$$

The results can be seen in Figure 7. An open source high efficiency centrifugal compressor (HECC) map from NASA was included (Figure 8) as an example for surge margins [13]. The data for a vaneless diffuser with 0.018 inch tip clearance was used, as this was the largest dataset with larger surge margin.

RESULTS AND DISCUSSION

System control responses

Figure 6 shows that both control structures achieve stable control for feed throughput disturbances, with fast responses and settling times of about 1 hr, due to short liquid holdup times of the packed columns. PCS 1 shows faster pressure control than PCS 2, while temperature

control settling times are similar. For both strategies, a large auxiliary condenser operating range is recommended for fast and robust control. The studied D-HIDiC has an auxiliary condenser duty to total condenser duty ratio of only 4.3%. A feed preheater or higher compressor interstage temperature could be used to increase this ratio, improving control robustness at an efficiency cost.

PCS 2 is a material balance-based control strategy and therefore cannot maintain product purity at increased feed throughput, as the reflux ratio must decrease to keep compressor flow constant. For a 20% feed throughput increase, the reflux ratio changes from 1.16 to 0.63, decreasing distillate purity by 10 mol% and dropping methanol below fuel grade (95 wt%, max 1 wt% water). Because the steady-state reflux ratio is close to 1, the absolute changes in reflux and distillate flow produce a large relative change in reflux ratio; for high reflux ratio columns, this effect would be smaller. Conversely, feed decreases cause over-purification at the same energy expenditure as steady-state production, meaning a 50% turndown ratio results in twice the energy cost per unit of purified methanol compared to steady state.

Despite these disadvantages, PCS 2 has one significant merit: compressor mass throughput is stabilized and step-change feed disturbances are substantially damped at the compressor. PCS 2 is therefore suitable for processes where the economic benefits of flexible operation of electricity-intensive upstream processes outweigh the increased purification cost per unit during turndown. For e-methanol production, about 90–95% of all electric power is consumed by water electrolysis [14], making the cost savings from electrolyzer turndown during electricity price spikes larger than the increased purification costs per unit of methanol incurred by PCS 2.

In general, PCS 2 should only be applied for turndown and the subsequent return to steady-state production, as product purity is maintained during these transient periods. Ramping above design capacity with PCS 2 should be avoided; a different strategy such as PCS 1 should be used instead to prevent distillate purity loss.

Compressor sensitivity analysis

PCS 1 shows a proportional change of the corrected compressor mass throughput for feed rate changes (Figure 7). This is expected, as the compressor duty is altered to maintain stripping column pressure, matching the change of feed inflow. For PCS 2, the change in compressor mass throughput is damped (Figure 7). For example, a -20% feed step change causes a -5.7% change in compressor mass throughput. Figure 8 shows that for any given operating point, if the pressure ratio is maintained constant and the corrected mass flow rate is halved, the new operating point is beyond the surge limit. Based on this compressor map, PCS 1 is incapable of a 50% turndown ratio. This shows the potential of PCS 2 for flexible operation, especially when using feed ramp changes. Ramp time could be chosen such that the compressor experiences minimal changes in mass throughput. Nevertheless, it should be mentioned that a compressor map is unique for every compressor. Choices can be made in compressor design to increase surge margin at a decrease in efficiency and pressure ratio.

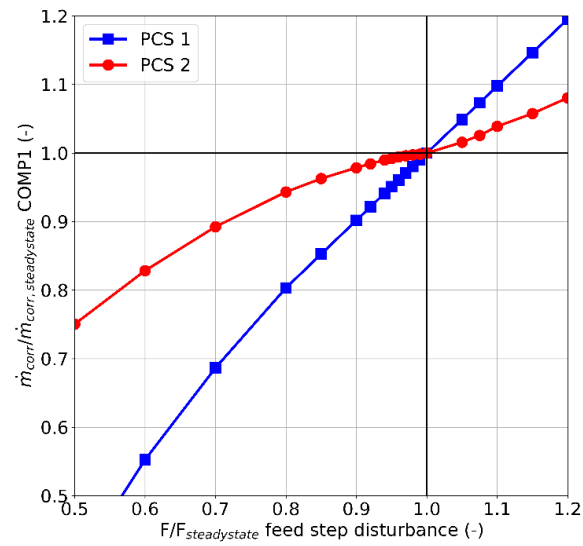

Figure 7. Compressor sensitivity analysis

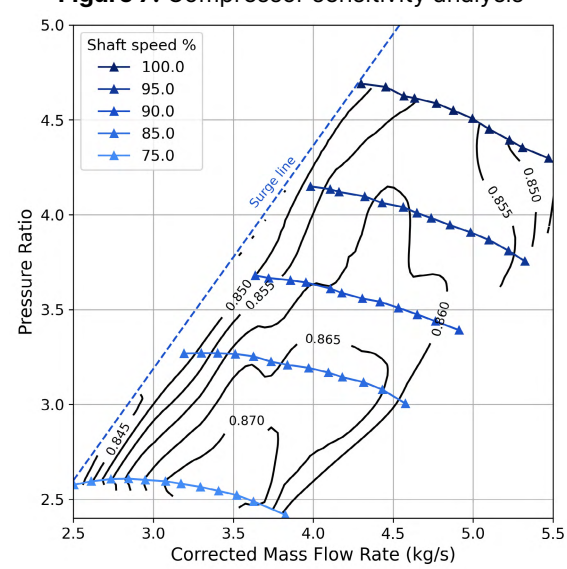

Figure 8: HECC compressor map

CONCLUSIONS

This study has provided a thorough analysis of a control structure based on variable compressor duty (PCS 1) and for a control structure based on fixed compressor duty (PCS 2) for D-HIDiC. Both control structures are based on inferential temperature control and show robust performance for ±20 % feed throughput disturbances. Compressor sensitivity analysis has shown that

PCS 2 damps throughput disturbances for the compressor. PCS 2 can maintain stable compressor operation at 50% turndown ratio, whereas PCS 1 requires a compressor capable of operating at 50% capacity. PCS 2 enables high turndown, but at the cost of losing purity control during production increases and incurring an energy efficiency penalty during turndown. Its application is therefore best suited for scenarios where the economic value of flexibility outweighs these drawbacks.

These findings highlight the potential application of D-HIDiC technology in flexible chemical production systems. This work contributes to further electrification of distillation processes by making D-HIDiC compatible with variations in power availability.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge financial support through the EMPHATICaL project. the EMPHATICaL project has received funding from the European Union's horizon Europe research and innovation programme under grant agreement no. 101177725.

REFERENCES

1. Olah GA. Beyond oil and gas: the methanol economy. Angew Chem Int Ed 44:2636-2639 (2005). https://doi.org/10.1002/anie.200462121
2. Kiss AA, Flores Landaeta SJ, Infante Ferreira CA. Towards energy efficient distillation technologies – making the right choice. Energy 47:531-542 (2012). https://doi.org/10.1016/j.energy.2012.09.038
3. Jana AK. Advances in heat pump assisted distillation column: a review. Energy Conversion and Management 77:287-297 (2014). https://doi.org/10.1016/j.enconman.2013.09.055
4. Wakabayashi T, Yoshitani K, Takahashi H, Hasebe S. Verification of energy conservation for discretely heat integrated distillation column through commercial operation. Chemical Engineering Research and Design 142:1-12 (2019). https://doi.org/10.1016/j.cherd.2018.11.031
5. Cui C, van Reisen J, Tyraskis I, Kiss AA. Efficient heat integration within discretely heat integrated distillation columns using liquid injection. AIChE Journal 71: (2025). https://doi.org/10.1002/aic.18861
6. Ho TJ, Huang CT, Lin JM, Lee LS. Dynamic simulation for internally heat-integrated distillation columns (hidic) for propylene–propane system. Computers & Chemical Engineering 33:1187-1201 (2009). https://doi.org/10.1016/j.compchemeng.2009.01.004
7. Cui C, Li Q, Luyben WL, Kiss AA. Dynamics and control of discretely heat integrated distillation columns. Computers & Chemical Engineering 199:109144 (2025). https://doi.org/10.1016/j.compchemeng.2025.109144
8. Huang K, Nakaiwa M, Akiya T, Aso K, Takamatsu T. A numerical consideration on dynamic modeling and control of ideal heat integrated distillation columns.. J. Chem. Eng. Japan / JCEJ 29:344-351 (1996). https://doi.org/10.1252/jcej.29.344
9. Kiss AA, Pragt JJ, Vos HJ, Bargeman G, de Groot MT. Novel efficient process for methanol synthesis by CO2 hydrogenation. Chemical Engineering Journal 284:260-269 (2016). https://doi.org/10.1016/j.cej.2015.08.101
10. Luyben WL. Distillation design and control using aspentm simulation. Wiley (2006). https://doi.org/10.1002/0471785253
11. Methanol institute, 'Methanol-Technical-Data-Sheet'. 2025. [Online]. Available: https://methanolorg.wpenginepowered.com/wp-content/uploads/2016/06/Methanol-Technical-Data-Sheet.pdf
12. 'Hydrogen Production and Consumption Projects | European Hydrogen Observatory'. Accessed: Dec. 09, 2025. [Online]. Available: https://observatory.clean-hydrogen.europa.eu/hydrogen-landscape/projects-and-valleys/hydrogen-production-and-consumption-projects
13. Robles Vega G, Bosshart AJ, Ni M, Ni RH, Harrison HM, Nguyen-Huynh T. NASA HECC geometry and performance review part 1: validation of a computational model for the vaneless diffuser configuration with as-manufactured impeller geometry. Volume 12D: Turbomachinery — Multidisciplinary Design Approaches, Optimization, and Uncertainty Quantification; Radial Turbomachinery Aerodynamics; Unsteady Flows in Turbomachinery : (2024). https://doi.org/10.1115/gt2024-125128
14. Battaglia P, Buffo G, Ferrero D, Santarelli M, Lanzini A. Methanol synthesis through CO2 capture and hydrogenation: thermal integration, energy performance and techno-economic assessment. Journal of CO2 Utilization 44:101407 (2021). https://doi.org/10.1016/j.jcou.2020.101407

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Long-Cycle Operation for Residue Hydrotreating Processes with Bayesian Optimization

Pengcheng Zhu[a], Han Wang[b], Gang Chen[b], Bo Chen[b], Fei Zhao[a]*, Xi Chen[a]

[a] Zhejiang University, College of Control Science and Engineering, Hangzhou, Zhejiang, China
[b] Sinopec (Dalian) Research Institute of Petroleum and Petrochemicals Limited Company, Dalian, Liaoning, China
* Corresponding Author: zhaofeizju@zju.edu.cn.

ABSTRACT

For the long-cycle process industry, operational cycles can be severely affected by equipment aging, catalyst deactivation, and safety limitations. As illustrated by the residue hydrotreating process, metal impurities gradually deposit on the catalyst during residue purification, leading to catalyst poisoning and eventual process shutdown. Such long-cycle processes require dynamic adjustments of operating conditions to balance immediate economics with long-term sustainability. While current practice relies on empirical tuning based on historical data, this work focuses on studying how to obtain an optimal operating trajectory to guide the monthly adjustments of operating variables. The long-cycle simulation of the residue hydrotreating process can be performed using the commercial software, PetroSIM. After adjusting the feed conditions, its embedded mechanistic model can calculate the deviation of average bed temperature from the set point and output the remaining operating time. Since Bayesian optimization (BO) is well-suited to address complex processes with unknown mechanisms and high computational costs, and can effectively seek optimal solutions, a BO framework incorporating constraints evaluated by PetroSIM is proposed in this work to optimize the monthly feed composition and maximize the total profit over the entire operational cycle. The results indicate that the total profit achieved through BO-optimized operation is 17.8% higher than that from empirical operation. The optimized strategy demonstrates practical rationality: in the early stage of high catalyst activity, more residue can be processed; in the later stage, increasing the light oil ratio helps extend the processing cycle. This strategy provides theoretical support for advancing the research toward industrial applications.

Keywords Process Operations, Derivative Free Optimization, Petroleum, Hydrotreating processes

INTRODUCTION

Residue hydrotreating (RHDT) is a process that removes sulfur and metal impurities from heavy oil via catalytic reactions [1]. As the process proceeds, the catalyst gradually deactivates until its activity drops below a specified threshold, requiring a unit shutdown and catalyst replacement [2]. RHDT therefore operates over finite operational cycles. Given the high cost of the catalyst and the time-consuming replacement procedure, it is economically imperative to purify as much residuum as possible within each cycle. The key operational variable is the residuum content in the feedstock. Increasing this fraction lowers feed costs, thereby raising short-term processing profits. However, feeding a higher-impurity feedstock accelerates catalyst deactivation, which shortens the operable duration in the future. Conversely, decreasing the fraction yields the opposite effect: it extends catalyst life and increases process sustainability at the expense of short-term operating profit. This poses a significant trade-off problem: determining how to optimize the operations each month to maximize the total profit across the entire RHDT running cycle.

In this optimization, a critical constraint is the relationship between feed composition and the remaining useful life (RUL) of the catalyst. PetroSIM, a professional software in this field, contains an RHDS module capable of simulating the long-cycle operation of residue hydrotreating [3]. This module incorporates state-space correlations to model catalyst deactivation: it predicts catalyst activity for the subsequent month based on the current state (e.g., catalyst activity) and process inputs

(e.g., feed properties). Using this framework, PetroSIM calculates the minimum temperature rise required to satisfy the product composition constraint, from which the weighted average bed temperature (WABT) is determined. The distance between this value and the predefined threshold, combined with the catalyst deactivation rate, is then used to estimate the current RUL. Hence, distinct operating strategies yield different RUL trajectories. These RUL estimates are generated by the deterministic model embedded within PetroSIM. However, due to the closed architecture of this commercial software, neither the model nor its underlying database is accessible. As a result, analytical derivatives cannot be obtained. This study does not seek to construct a surrogate model for this internal relationship. Instead, we focus on solving the long-cycle operational optimization problem using this black-box simulator, and it is a scenario commonly encountered in industrial practice.

Bayesian optimization (BO) is a widely used black-box optimization algorithm grounded in Bayesian theory [4]. As an iterative framework, it integrating modeling and optimization. Specifically, the algorithm first fits a surrogate model (e.g., Gaussian process regression) for the objective function using an initial set of observed points. Based on this model, the subsequent point is selected by maximizing an acquisition function (e.g., expected improvement) that balances exploration and exploitation. Each newly evaluated sample is incorporated to update the surrogate, thereby progressively reducing uncertainty over the objective function. By repeatedly executing this 'decision-modeling' iteration, Bayesian optimization converges to the optimum with as few function evaluations as possible. Consequently, it is particularly suitable for black-box issues where evaluations are computationally expensive.

For the RHDT process, the operational optimization is conducted via the PetroSIM simulator, which is characterized by its opaque mechanisms and significant computational cost (a full-cycle simulation takes over one minute). If an evolutionary strategy (e.g., genetic algorithm) requiring numerous iterations is adopted, the optimization process becomes time-consuming—at least 16.7 hours for 1000 iterations. Thus, to minimize the number of calls to PetroSIM, this study utilizes BO method to address this operating scenario.

This section outlines the theoretical foundation and overall framework of this study. The remainder of the paper is organized as follows: Section 2 details the methodology for simulating the long-cycle RHDT process using PetroSIM; Section 3 formulates the objective of this study by defining the optimization problem, along with its associated constraints and boundary conditions; Section 4 presents the results and discusses the operational benefits achieved; Section 5 concludes this study and suggests potential directions for future research.

RHDT PROCESS SIMULATION

Process simulation of residue hydrotreating can be accomplished using the RHDS module in PetroSIM. Figure 1 shows the RHDT process flowsheet developed on this platform. Residuum and hydrogen are fed into the reactor for hydrotreating. The effluent then passes through an ideal separator (modeled as a component splitter) to simulate a high-pressure high-temperature flash drum, assuming complete separation of non-condensable gases (H_2S, NH_3) from the liquid hydrocarbons based on a defined split fraction. The Peng-Robinson (PR) equation of state with built-in binary interaction parameters is applied throughout the process, as it is widely accepted for hydrocarbon systems containing sour components (H_2S, NH_3). The liquid stream is subsequently fractionated into products of varying boiling points (gas, naphtha, diesel, and bottoms). In addition, H_2O serves as a wash water stream to rinse metal deposits from the catalyst. HPS functions as a high-pressure vent line to withdraw excess hydrogen feedstock.

Table 1: Properties of the feed oil.

Property	Value
Sulfur content	3.0 ~ 6.0 wt%
Conradson carbon content	11.20 wt%
Nickel content	21.80 ppmwt%
Vanadium content	64.40 ppmwt%
Nitrogen content	0.24 wt%
Distillation 1% (TBP)	273.4 °C
Distillation 5% (TBP)	349.4 °C
Distillation 10% (TBP)	384.6 °C
Distillation 30% (TBP)	459.0 °C
Distillation 50% (TBP)	523.2 °C
Distillation 70% (TBP)	597.6 °C
Distillation 90% (TBP)	694.4 °C
Distillation 99% (TBP)	722.2 °C
Specific gravity (Std. Cond.)	0.977 -
Temperature	260 °C
Pressure	101 kPa
Volume flow	535.7 m^3/h

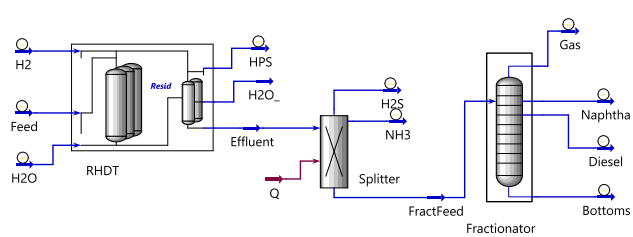

Figure 1. Flowsheet of the RHDT process in PetroSIM.

Table 1 presents the properties of the feedstock. The sulfur content is identified as the key factor responsible for catalyst deactivation. Hence, adjusting the

proportion of heavy oil in the feed can be regarded as effectively modifying the sulfur content. According to the review [5], this study specifies the sulfur content of the heaviest feed as 6.0 wt% and that of the lightest feed constituents as 3.0 wt%. As the only manipulated variable, feed sulfur content dominates catalyst RUL and thus drives profitability. Fixed true boiling point (TBP) cut points minimize product yield variation from the separation system, so the ideal splitter simplification has negligible impact on economic conclusions. Due to the requirements of the downstream processes in RHDT, the sulfur content in the bottoms product needs to be maintained below 0.5 wt%, regardless of the feedstock. However, there are price variations: heavy feedstock costs about 350 USD/m³ (P_H), light feedstock costs about 600 USD/m³ (P_L), and the purified effluent is valued at approximately 800 USD/m³ (P_B).

Table 2: Parameters of the reactor.

Parameter	Value
WABT (Start of Run)	360 °C
WABT (End of Run)	420 °C
Number of beds	5 -
Catalyst density	493.7 kg/m³
Catalyst weight	1656.37 ton

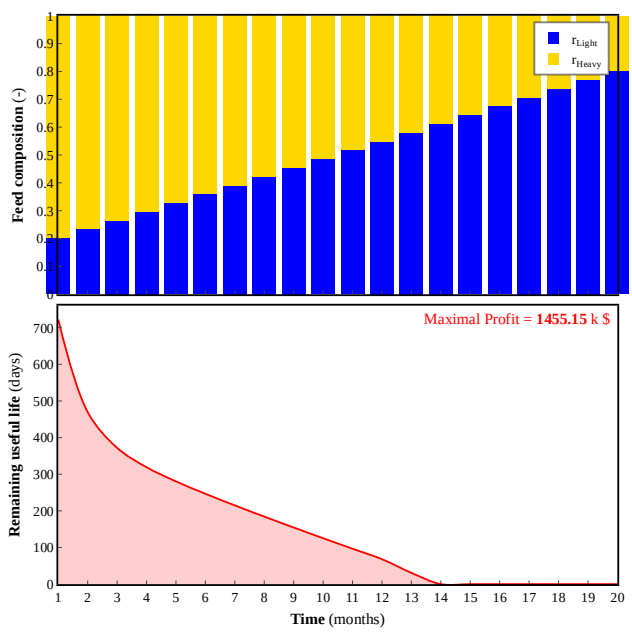

Figure 2. Feed composition and RUL trajectory under empirical operation.

In addition, the parameters of the reactor remain constant, as listed in Table 2. It can be seen that the catalyst inventory is held constant. Thus, its activity progressively decreases during residue hydrotreating. In PetroSIM, this effect is reflected by a rise in WABT, as the process must maintain the product sulfur specification. Starting from the first month (WABT=360°C), operating decisions are updated at monthly intervals. The process stops once the WABT exceeds the set value (WABT=420°C). Figure 2 illustrates an empirical operational approach: linearly increasing the light feedstock fraction from 0.2 to 0.8 to compensate for the declining catalyst activity in later periods. However, as can be seen from PetroSIM's estimation, the RUL reaches zero by the 14th month, implying that no profit can be generated from subsequent operations. Therefore, the total profit, calculated as the cumulative monthly product revenue minus feedstock costs, is 1455.15 thousand USD.

OPERATIONAL OPTIMIZATION

In this study, given that hydrogen is typically supplied in excess during industrial residue hydrotreating and its precise consumption is difficult to quantify, hydrogen-related costs are treated as fixed overheads and excluded from the objective function. Operating conditions influence utility consumption through their effect on WABT; however, the integrated value of total utility cost remains largely unchanged across different operating strategies over the entire processing cycle, allowing minor fluctuations to be neglected in the model. Moreover, catalyst are expensive and consumed in large quantities; frequent replacement not only increases operating expenses but also disrupts the refinery-wide material balance. Industrial practice therefore prioritizes the full utilization of catalyst within each operating cycle. Accordingly, this study optimizes profit over a full single catalyst cycle; as the one-time catalyst replacement cost is independent of operating decisions, it is neglected in the analysis. Based on these assumptions, the objective is to maximize the total profit by optimizing the heavy component fraction in the monthly feed. The optimization problem is expressed in Equation 1. This study sets the total duration to no more than 20 months. The heavy oil fraction in the feed for each month is represented by the set $\boldsymbol{r}_H = \{r_H^{(i)} \mid i = 1, \ldots, 20\}$. The light oil fraction is complementary, such that $r_H^{(i)} + r_L^{(i)} = 1$. A binary decision variable $z_i \in \{0, 1\}$ governs whether the monthly processing profit contributes to the total profit; its value is assigned based on the RUL estimated via PetroSIM: $z_i = 1$ when $RUL > 0$, and $z_i = 0$ when $RUL = 0$. The total feed flow rate F_{in} is set to 535.7 m³/h. The product flow rate of the bottoms $F_B^{(i)}$ is simulated by PetroSIM. P_B, P_H, and P_L correspond to the prices of the bottoms, heavy feedstock, and light feedstock. Additionally, c is defined as 24*30 (hours*days).

$$\max_{\boldsymbol{r}_H} \sum_{i=1}^{20} z_i * \left(F_B^{(i)} P_B - r_H^{(i)} F_{in} P_H - r_L^{(i)} F_{in} P_L\right) * c \qquad (1)$$

$$z_i, F_B^{(i)} = \mathrm{PetroSIM}\left(r_H^{(i)}, r_L^{(i)}\right), i = 1, \ldots, 20$$

$$r_H^{(i)} + r_L^{(i)} = 1, i = 1, \ldots, 20$$

$$0 \le r_{\mathrm{H}}^{(i)} \le 1$$

$$0 \le r_{\mathrm{L}}^{(i)} \le 1$$

The Bayesian optimization algorithm is applied to optimize this problem. The procedure is illustrated in Figure 3. First, 30 sets of $\boldsymbol{r}_{\mathrm{H}}$ are sampled within the domain using Latin hypercube sampling, and the light feedstock fraction $\boldsymbol{r}_{\mathrm{L}}$ is derived accordingly. For each set, PetroSIM computes the corresponding $F_{\mathrm{B}}^{(i)}$ and $RUL^{(i)}$ (i.e., z_i). Subsequently, the values of the objective function are calculated, and the sample with the highest value is denoted as $\boldsymbol{r}_{\mathrm{H}}^*$. Gaussian process regression is utilized to fit the mapping from the decision variables to the objective function. The next point to be evaluated is determined via the expected improvement. This process iterates until a predetermined iteration count is met, at which point the optimal solution $\boldsymbol{r}_{\mathrm{H}}^*$ is returned.

This framework is implemented via the BoTorch library in Python, and the parameters are presented in Table 3. The RBF kernel, known for its universal approximation capability, is adopted for the Gaussian process, while expected improvement (EI) serves as the acquisition function. A total of 550 iterations are performed, with the initial samples used to initialize the Gaussian process regression constituting about 5.5% of the total iterations.

Table 3: Parameters of BO.

Parameter	Value
Gaussian kernel	RBF kernel
Acquisition function	EI
Number of initial points	30
Number of iterations	520

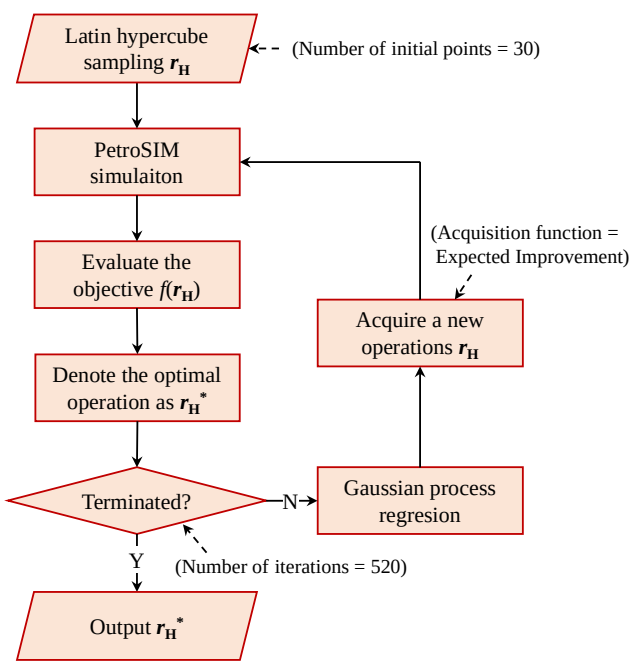

Figure 3. Workflow of the BO strategy.

RESULTS AND DISCUSSION

Figure 4 illustrates the iterative process of the algorithm. The horizontal axis represents the number of objective function evaluations, comprising 30 initial evaluations and 520 iterative evaluations. Total revenue and total cost are shown by cyan and pink bars, respectively. The difference between them, which is the total profit (i.e., objective function value), is plotted as a blue curve. In addition, the green line represents the magnitude of the acquisition function, i.e., the EI value. It is mapped to the right vertical axis, and non-zero values of this line suggest that the subsequent sampling point can improve upon the current $\boldsymbol{r}_{\mathrm{H}}^*$. It can be observed from the figure that the green line displays multiple peaks in the initial iterations. The BO focuses on exploring the decision space during this period. This drives the algorithm to quickly locate the promising region for the optimum. In the middle and later stages, BO emphasizes exploitation by sampling around the optimal solution, which serves to improve the accuracy of the surrogate model. The algorithm identified the optimal value at iteration 432, requiring merely 7.2 hours of runtime.

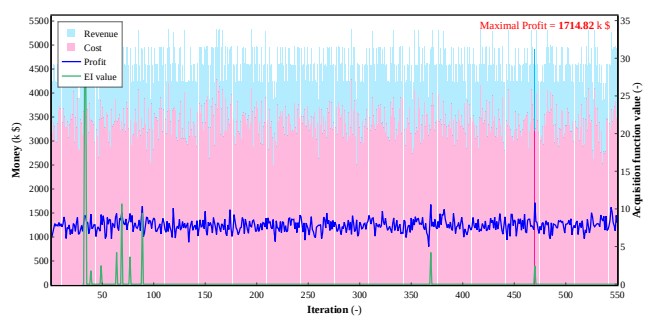

Figure 4. Iteration process of BO for Equation 1.

The bars representing revenue and cost associated with the optimal operation are highlighted in Figure 4. The total profit obtained is 1714.82 thousand USD, reflecting a 17.8% improvement over that achieved by the empirical operation. Figure 5 depicts the optimal operation and the estimated RUL. As shown, a higher fraction of heavy oil is fed during the early processing stage, thereby maintaining a lower cost level. Contrary to the empirical operation that continuously increases the fraction of light oil feed, optimized strategies recommend utilizing periods of high catalyst activity to process more residue. Subsequently, the fraction of light oil feed is progressively increased in the mid-to-late periods to prolong the processing time. Compared with the results in Figure 2, the RUL here is sustained for one additional month, resulting in additional profit.

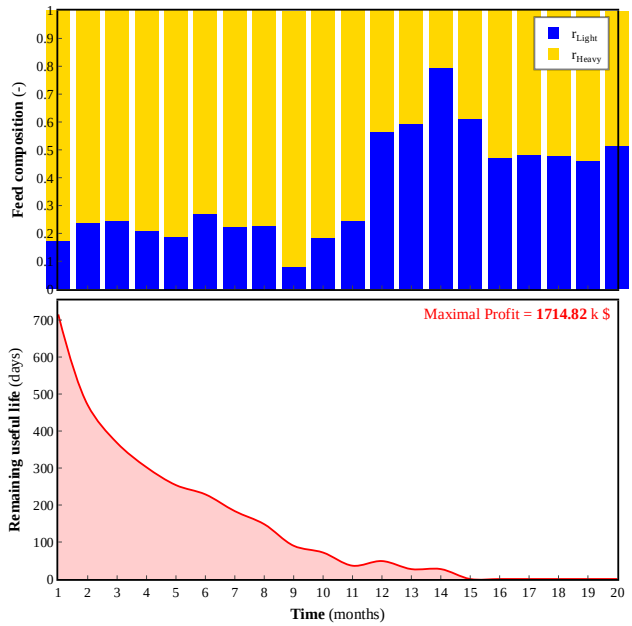

Figure 5. Feed composition (with varying r_H) and RUL trajectory after BO optimization.

Figure 6 illustrates the iterative process, showing that BO identified the optimum at a relatively early stage. This is attributed to the reduced dimensionality and simpler surrogate modeling in the single-variable problem. As shown in Figure 7, the results suggest that residue should constitute the majority of the feed. Although only 12 months of operation are valid, the profit reaches 1667.83 thousand USD, which is 212.68 thousand USD higher than that of empirical operation. However, this result is inferior to that of the free-feed optimization. The rationale lies in the fact that the optimal solution under this problem is a lower bound solution of Equation 1. Consequently, plants are advised to optimize the supply of light and heavy oils at the supply chain level.

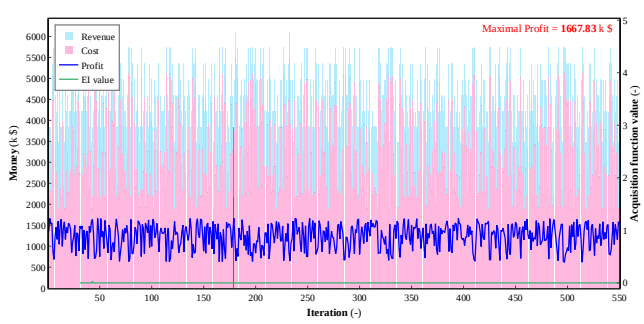

Figure 6. Iteration process of BO for Equation 1 under constant r_H constraint.

Additionally, in practical applications, the supply of light oil often lacks flexibility. This implies that the feed composition must be predetermined prior to startup. For comparative purposes, another optimization scenario is investigated: maximizing total profit by optimizing a constant heavy oil fraction in the feed. Mathematically, this is equivalent to adding a constraint to Equation 1: $r_H^{(1)} = r_H^{(2)} = \ldots = r_H^{(20)}$.

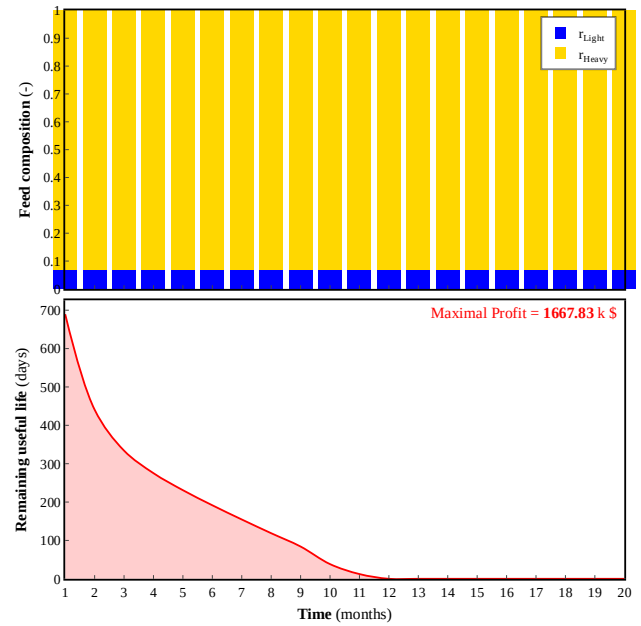

Figure 7. Feed composition (with constant r_H) and RUL trajectory after BO optimization.

CONCLUSION

The main contribution of this work lies in developing an optimization framework for long-cycle operating trajectories of residue hydrotreating based on the PetroSIM, a problem essentially characterized by high-dimensional black-box optimization. BO, leveraging a probabilistic surrogate model to actively sample and iteratively update the posterior, enables efficient optimization with minimal objective function evaluations. This study employs BO to address the problem. Following 432 iterations (7.2 hours), BO identifies the optimal solution, yielding a 17.8% profit improvement compared to empirical operation. The optimized results exhibit an operating pattern with a higher heavy oil ratio in the early stage and a higher light oil ratio in the later stage. This pattern is consistent with the decreasing catalyst activity over time and the stronger impact of feed when activity is low. This provides practical guidance: operators can initially set a high throughput and then gradually reduce it based on WABT feedback. Moreover, this study investigates a lower-bound case of the problem by optimizing with fixed operations each month. This is designed for refineries with limited flexibility in light feed supply. The results, though not as good as those from unrestricted operational optimization, still show a 14.6% increase over the unoptimized practice.

Similar optimization problems, balancing immediate

economic returns and future sustainability, are prevalent in various industrial contexts. Additionally, commercial software such as PetroSIM, PRO/II, and Aspen HYSYS do not provide analytical derivatives. The proposed BO framework is general and inherently incorporates uncertainty considerations, laying a foundation for methodological innovations and adaptations to more complex industrial problems.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the financial support from the Zhejiang Provincial Natural Science Foundation of China (No. LZYQ25F030001).

REFERENCES

1. Ferreira C, Tayakout-Fayolle M, Guibard I, Lemos F. Hydrodesulfurization and hydrodemetallization of different origin vacuum residues: new modeling approach. Fuel 129:267-277 (2014). https://doi.org/10.1016/j.fuel.2014.03.056
2. Kohli K, Prajapati R, Maity SK, Sau M, Garg MO. Deactivation of hydrotreating catalyst by metals in resin and asphaltene parts of heavy oil and residues. Fuel 175:264-273 (2016). https://doi.org/10.1016/j.fuel.2016.02.036
3. Fan M, Zhao Y, Wang Y, He S. Simulation and application of fixed-bed residuum hydrotreating unit process model based on RHDS-SIM. *Chem. Ind. Eng. Prog.* 40:5514– 5522 (2021) https://link.cnki.net/doi/10.16085/j.issn.1000-6613.2020-2152
4. Frazier PI. A tutorial on Bayesian optimization. *arXiv preprint* arXiv:1807.02811 (2018) https://doi.org/10.48550/arXiv.1807.02811
5. Marafi A, Albazzaz H, Rana MS. Hydroprocessing of heavy residual oil: opportunities and challenges. Catalysis Today 329:125-134 (2019). https://doi.org/10.1016/j.cattod.2018.10.067

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Decentralized Causal Monitoring in High-Dimensional Systems: Revealing the Topological Drivers behind Fault Detection Performance

Rodrigo Paredes[a] and Marco S. Reis[a*]

[a] University of Coimbra, CERES, Department of Chemical Engineering, 3030-790 Coimbra, Portugal
* Corresponding Author: marco@eq.uc.pt.

ABSTRACT

Centralized monitoring methods experience reduced fault detection sensitivity in large-scale industrial systems due to the masking effect arising from the aggregation of many interconnected variables. Decentralized monitoring, where variables are grouped into subsystems, has been shown to effectively address these limitations. However, the performance of this class of methods critically depends on how the network is partitioned, and the role of its structural factors on fault detection remains poorly understood. This work studies how network topology and causal structure affect decentralized monitoring in high-dimensional systems. Using *SimCaNet*, a DAG-based data simulator, where large-scale systems with 100-1000 variables were generated, we rigorously compared the performance of centralized and decentralized causal log-likelihood monitoring methods under process perturbations and sensor bias faults. Network partitioning is performed using the Leiden community detection algorithm and characterized at the network, community, and node levels. Fault Detection Sensitivity (AUC_{NORM}) is quantified per variable, using the normalized area under the curve, which is computed from the true positive rate profiles across fault magnitudes. Results show that AUC_{NORM} is primarily driven by the number and the size of communities across fault types. Smaller communities consistently achieve higher AUC_{NORM} than larger communities, while excessive fragmentation or high modularity degrade performance. Community density and inter-community coupling further enhance fault detection sensitivity. The strength of node-level causal relationships explains most of the variability observed at root nodes for sensor bias faults and their robustness to information loss. Furthermore, ablation studies confirm the robustness of distributed systems to major disruptions in the system: more than 32% of community information can be removed without significantly impacting the fault detection sensitivity. These topology-driven partitioning findings constitute novel and valuable contributions to the theoretical foundations for designing better distributed monitoring systems in complex industrial plants.

Keywords: Fault Detection, Modelling and Simulations, Industry 4.0, Big Data, Decentralized Monitoring, Network Topology, Structural Causal Models, Community Detection

INTRODUCTION

Fault detection in large-scale industrial systems remains a significant challenge with direct implications for safety, efficiency, and operational reliability. As modern processes scale in size and complexity, hundreds or thousands of interdependent variables must be monitored simultaneously. In such high-dimensional settings, small and localized faults are often statistically diluted within global monitoring indices, leading to delayed detection or complete loss of sensitivity [1, 2].

Decentralized monitoring has emerged as a promising strategy to mitigate this limitation, by decomposing large systems into smaller groups monitored locally. By reducing dimensionality and focusing on localized dynamics, decentralized schemes can detect faults before they propagate across the network [3, 4]. However, the effectiveness of decentralization critically depends on system partitioning, where poorly designed partitions can replicate the limitations of centralized monitoring or

introduce excessive false alarms.

Fault propagation and its proper diagnostics are inherently governed by causal relationships between process variables. This motivates the use of causal network representations for monitoring, where directed interactions encode how disturbances propagate through the system [5-7]. While causal monitoring frameworks provide a principled basis for understanding fault propagation, their performance in high-dimensional systems remains limited.

Recent advances in network science have shown that the topological structure of complex systems shapes information flow [8, 9]. Properties such as community structure, internal connectivity, modularity, and node centrality have been shown to influence system dynamics. Fault detection sensitivity in large-scale systems is not only a statistical problem, but also a structural one. Despite this, the role of network topology in decentralized fault detection has not been systematically addressed or characterized. It is still unclear which topological properties govern the sensitivity gains and under what conditions decentralization may fail. Consequently, current monitoring strategies lack topology-aware design principles, limiting their application, scalability, and robustness.

To address this gap, in this work we provide a quantitative analysis of how network topology and causal structure jointly affect decentralized monitoring in high-dimensional systems, with the building blocks described next. A graph-based causal data simulator generates networks reflecting structural characteristics commonly observed in industrial processes. Two representative fault classes are considered: process perturbations and sensor bias faults. Network partitions are obtained using the Leiden algorithm, a modularity-based community detection method, allowing controlled variation of partition size and structure [10]. Monitoring performance is evaluated using a Fault Detection Sensitivity metric that captures node-level fault detection across fault magnitudes, enabling fair comparison between centralized and decentralized schemes.

Our results reveal a consistent relationship between fault detection sensitivity and topological descriptors, including community size, internal density, modularity, and node centrality. Decentralized causal monitoring systematically outperforms centralized approaches in large-scale networks. However, suboptimal partitions may reduce these benefits, highlighting the need to consider structural trade-offs. These findings provide fundamental insights to the systematic design of topology-aware monitoring and demonstrate that fault detection in high-dimensional systems depends fundamentally on network topology.

METHODOLOGY

This section presents the framework used to assess fault detection sensitivity in large-scale systems. The methodology combines a high-dimensional data simulator, centralized and decentralized monitoring schemes, network partitioning, and a unified fault detection sensitivity metric. The goal is to isolate the impact of network topology, partition structure, and fault type on fault detection performance under fully controlled conditions.

Network Data Simulator and Fault Types

To assess fault detection sensitivity in large-scale systems, a high-dimensional causal data simulator, called SimCaNet, is used. Existing benchmark simulators for process monitoring are typically limited to low-dimensional settings, which hinders systematic analysis of how network topology affects fault detection sensitivity. Although latent-variable models can handle high-dimensional data, they ignore explicit causal relationships, and are therefore less sensitive and have more limited diagnostic capabilities [7, 11].

SimCaNet generates synthetic datasets using Directed Acyclic Graphs (DAGs) implemented as Structural Causal Models (SCMs) [12]. Each variable is defined as a function of its causal parents, ensuring understandable propagation mechanisms and realistic fault patterns.

Network structures are generated using a preferential-attachment mechanism based on the Barabási-Albert model, producing heterogeneous, hub-dominated topologies closely related to large-scale industrial systems. This choice reflects the non-uniform connectivity and consequently the asymmetric propagation paths, commonly observed in practice. All network structures preserve the DAG property to ensure causal consistency.

Data are generated under Normal Operating Conditions (NOC) following a linear and static SCM formulation, with additive measurement noise to emulate realistic observational conditions. This simplification enables focusing on the identification of critical topological features driving fault detection sensitivity, without the interference of other factors, such as non-linearity and multiscale dynamics. Structural parameters are generated once for each network and kept fixed across experiments, so that observed variations in monitoring performance can be attributed to network topology, partitioning, and fault characteristics.

Two representative fault classes are considered, as summarized in Table 1. Process perturbations are modeled as step changes in root variables (data drift), inducing process disturbances that propagate through the causal network according to its topology. Sensor bias faults are modeled as additive biases affecting only local measurements, without altering the underlying process. Fault severity is controlled through a fault magnitude

parameter proportional to the variability of the faulty variable.

Table 1: Types of faults implemented in SimCaNet.

Type of fault	Meaning	Modified factor
Process Perturbation	A step change in root nodes, representing e.g., abnormal (but real) variation in raw materials	$\mu_i = \mu_{NOC,i} + k\sigma_i$
Sensor bias	Biased readings caused by a fault in the sensors	$X_i = X_i + k\sigma_{e_i}$

Large-scale Causal Monitoring Methods

This section presents the two causal-based Statistical Process Monitoring methods used in this work to evaluate how network dimensionality and structure affect fault detection sensitivity. Both approaches follow the standard SPM pipeline, including phases of process stability analysis, calculation of monitoring statistics, and fault detection of a new observation via control limits.

Centralized Causal Log-likelihood monitoring

The Centralized Causal Log-likelihood monitoring (CNc-MSPC) method is based on log-likelihood indices and was adopted as the benchmark for the fault detection performance study. In ref. [13], this methodology was compared with the classical PCA-MSPC approach using both a simulated dataset and an industrial semiconductor manufacturing dataset. The results showed that the causal approach achieved superior fault diagnosis, and better or at least similar detection performance.

In practice, the causal network structure is inferred through causal discovery methods [14]. Here, the causal graph is assumed to be known, as the purpose is to study the impact of network topology (and not discovery accuracy).

Each variable X_i is modelled through a SCM as a function of its causal parents, $\mathrm{Pa}(X_i)$. These yields, $X_i = f(\mathrm{Pa}(X_i)) + e_i$, where the conditional distribution of $X_i|\mathrm{Pa}(X_i)$ corresponds to the error term distribution, e_i.

A local log-likelihood index is computed for each observation *j*. Assuming a normal distribution, the log-likelihood for an individual observation $L_i^L(j)$ is defined as:

$$L_i^L(j) = -\frac{\ln(2\pi)}{2} - \frac{\ln(\hat{\sigma}_{e_i}^2)}{2} - \frac{(e_i(j) - \mu_{e_i})^2}{2\hat{\sigma}_{e_i}^2} \quad (1)$$

where $e_i(j)$ is the error term of variable X_i conditioned onto its causal parents, and $\hat{\mu}_{e_i}$ and $\hat{\sigma}_{e_i}$ are the sample mean and standard deviation of e_i calculated from a reference dataset corresponding to normal operating conditions (NOC).

A global log-likelihood index $L^G(j)$ is then obtained by aggregating the local indices across all nodes. Leveraging on the chain rule, the joint (or global) likelihood factorizes into the product of the conditional likelihoods. Therefore, applying logarithms, the global log-likelihood, $L^G(j)$, is obtained in terms of the individual log-likelihoods, as follows:

$$L^G(j) = \sum_{i=1}^{n_{variables}} L_i^L(j) \quad (2)$$

Control limits for fault detection are established using the values of L^G obtained from a validation dataset. A kernel density estimation (KDE) is fitted to approximate the probability density function of this global log-likelihood statistic [14]. The cumulative distribution function (CDF) of the KDE is then used to determine the Lower Control Limit (LCL) by solving the following equation for a predefined Type I error rate, α:

$$\int_{-\infty}^{\mathrm{LCL}} \hat{f}_{KDE}(L^G)\, \mathrm{d}L^G = \alpha \quad (3)$$

where, $\hat{f}_{KDE}(L^G)$ is the estimated probability density function. Since high log-likelihood values indicate proximity to the reference distribution, only the Lower Control Limit (LCL) is needed for fault detection. This threshold is used for monitoring future observations.

Decentralized Causal Log-likelihood monitoring

The Decentralized Causal Log-likelihood monitoring (CNd-MSPC) method also starts by inferring the causal network. Then, the variables are divided into disjoint communities (groups or clusters) using a community detection algorithm that considers network topology and modularity. Such grouping into strongly connected communities effectively enhances the fault detection sensitivity. However, the selection of the partition and consequently its quality directly impacts the monitoring performance.

Given a network partition, a global log-likelihood metric (Equation 2) and the control limits are computed for each community (sub-level monitoring), denoted as L_c^G for $c = 1, 2, \ldots, n_{communities}$, where $n_{communities}$ is the total number of communities.

To aggregate information from multiple communities, an OR-gate approach is adopted, meaning that a global warning is issued as soon as at least one community-level out-of-control signal is produced.

Each community's control limits are set independently, with the significance level adjusted to maintain the overall false alarm rate across multiple monitoring statistics. To this end, the Šidák correction is applied, setting the significance level for each community as $\alpha_c = 1 - (1-\alpha)^{1/n_{communities}}$ [15]. This ensures that the global alarm reflects true anomalies while controlling the overall Type I error rate (also known as the Family-Wise Error Rate, FWER). CNd-MSPC is particularly well-suited for

Table 2: Network-, community-, and node-level descriptors used to characterize the structural, statistical, and causal properties of each potentially faulty variable for a fixed network and partition.

Feature	Description	Level
Network Size	Total number of nodes in the network	Network
Modularity	Quantifies the quality of the network's division into communities	Network
Number of communities	Total number of communities in a partition	Community
Number of nodes in the community	Total number of nodes in the community of the faulty variable *i*	Community
Density	Quantifies the intra-connectivity in the community of the faulty variable *i*	Community
Conductance	Measures the fraction of edges leaving the community of the faulty variable *i* relative to all edges within and leaving the community	Community
Cohesiveness	Quantifies the maximum conductance among all connected subsets (size >2) of the community of the faulty variable *i*, capturing the weakest internal connectivity within the community	Community
Separability	Measures the ratio of external edges to the internal edges in the community of the faulty variable *i*	Community
Clustering Coefficient	The average clustering of nodes within the community containing the faulty variable *i*, quantifying the tendency to form tightly connected triplets	Community
Normalized Causal Strength	Quantifies the relative causal influence of the faulty node *i* on its children, normalized by the total parents' influence on each affected child, based on SCM coefficients	Node
In-degree	Number of incoming edges to faulty node *i*, indicating how many parent nodes directly influence it	Node
Out-degree	Number of outgoing edges from node *i*, indicating how many child nodes it directly influences	Node
Is root-node?	Binary indicator for the faulty node *i*	Node
Degree centrality	Fraction of nodes that are connected to the faulty node *i*	Node
Closeness centrality	Measures how easily other vertices can be reached from the faulty node *i*	Node
Betweenness centrality	A measure of how much the faulty node *i* acts as an intermediary in a shortest path between the other node pairs	Node
Eigenvector centrality	Measures the influence of faulty node *i* based on the influence of its neighbors	Node
Page Rank	Measures the influence of faulty node *i* based on the influence of incoming neighbors, adjusted by the damping factor α	Node
Shortest Path to Centroid	Measures how close faulty node *i* is to the three most degree central nodes of the network	Node

high-dimensional networks, where centralized monitoring often becomes statistically insensitive to localized faults. Although the offline phase (causal discovery, causal inference, process stability analysis) scales with network dimensionality, the online monitoring phase is computationally efficient.

Network Partitioning and Topological Descriptors

Partitioning the inferred causal graph is a crucial step in CNd-MSPC, as the resulting communities define the structural units used for decentralized monitoring. To generate these partitions, community detection algorithms are employed, leveraging the graph representation of the process. A broad range of algorithms is

Table 4: Summary of network partitions used in the Fault Detection Sensitivity study for Process Perturbation (PP) and Sensor Bias (SB) faults.

Network	Partition	#Communities	Modularity	#Faulty variables (PP/SB)
Network 1	Centr.	1	-	48/100
	#1	4	0.365	48/100
	#2	7	0.417	48/100
	#3	15	0.397	48/100
	#4	21	0.358	48/100
	#5	29	0.310	48/100
	#6	41	0.258	48/100
Network 2	Centr.	1	-	152/300
	#1	11	0.374	152/300
	#2	31	0.325	152/300
	#3	46	0.292	152/300
	#4	84	0.244	152/300
	#5	144	0.175	152/300
	#6	191	0.122	152/300
Network 3	Centr.	1	-	277/500
	#1	5	0.347	277/500
	#2	11	0.364	277/500
	#3	28	0.343	277/500
	#4	46	0.312	277/500
	#5	99	0.257	277/500
	#6	179	0.201	277/500
Network 4	Centr.	1	-	333/1000
	#1	7	0.341	333/1000
	#2	12	0.348	333/1000
	#3	37	0.322	333/1000
	#4	60	0.304	333/1000
	#5	183	0.230	333/1000
	#6	391	0.157	333/1000
			Total of tests:	56

available, including the Girvan-Newman and Louvain methods, which differ in the principles used to uncover community structure. In this work, the communities were generated through the Leiden algorithm.

The Leiden method improves on the Louvain method by addressing its main limitation: after the aggregation step, the Louvain method can yield communities that are internally disconnected or weakly connected. As for the Louvain method, the Leiden method also optimizes modularity, Q, but introduces a refinement step that guarantees connectivity within each community.

$$Q = \frac{1}{2m}\sum_{i,j}\left(w_{ij} - \gamma\frac{k_i k_j}{2m}\right)\delta(c_i, c_j) \qquad (4)$$

In Equation 4, modularity function is represented where *m* is the number of graph edges, γ is the resolution parameter controlling the granularity of identified communities, w_{ij} is the edge weight between nodes *i* and *j*, k_i and k_j are their respective degrees, and $\delta(c_i, c_j)$ is a community indicator function (1 if *i* and *j* belongs to the same community, *i.e.*, $c_i = c_j$, and 0 otherwise).

The Leiden method tends to produce communities that are densely connected internally while minimizing connections between communities.

Topological Descriptors

For a given network and a fixed partition, fault detection sensitivity may vary across fault locations due to the differences in the structural and causal embedding of variables within the network. To systematically analyze these effects, each faulty variable is characterized using a set of network-, community-, and node-level descriptors, summarized in Table 2. These descriptors are subsequently used for exploratory data analysis and pairwise correlation studies, together with the Fault Detection Sensitivity metric, to identify structural properties that drive fault detection under decentralized monitoring.

Fault Detection Sensitivity Assessment Framework

We developed a unified evaluation framework used to compare fault detection performance across all network sizes, partitioning strategies, fault types, and monitoring schemes. The framework ensures that all results are directly comparable, enabling the analysis of the effects of network topology and partitioning on fault

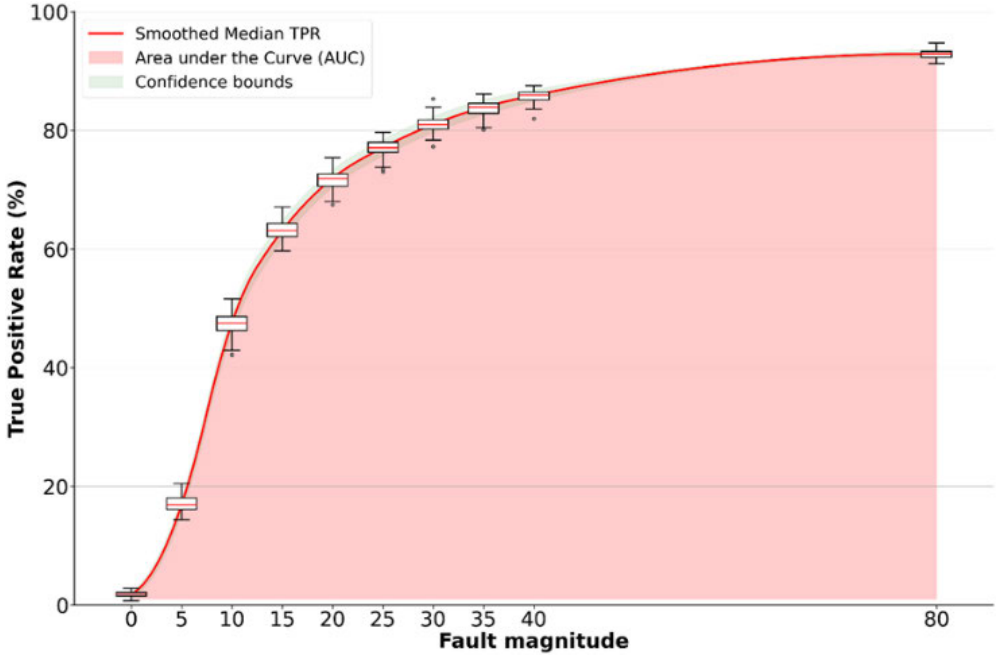

Figure 1. Fault detection sensitivity assessment for a sensor bias fault injected in a single network variable. Box plots show the distribution of True Positive Rate values for each fault magnitude. The red curve represents the median TPR versus magnitude, and the shaded red area corresponds to the AUC.

detection sensitivity.

For each network, both centralized and decentralized monitoring schemes are evaluated. Different partition granularities are obtained by varying the resolution parameter, γ, while centralized monitoring (i.e., no partitioning) is used as a benchmark configuration. Some descriptors of the tested networks are summarized in Table 3.

Table 3: Characterization of the four networks under study.

Descriptor	Network 1	Network 2	Network 3	Network 4
# of nodes	100	300	500	1000
# of edges	212	875	1564	3880
# of root nodes	48	152	277	333
Graph density	0.0214	0.0098	0.0063	0.0039

For process faults, perturbations are injected exclusively at root nodes. Ten fault magnitudes are considered, with a fixed magnitude range across all networks, but that is adapted to each fault type. To account for stochastic variability, each magnitude is simulated 50 times (independent replicates, each consisting of 1000 observations). The signal-to-noise ratio for the additive noise is fixed at 10 dB.For sensor bias faults, the same evaluation pipeline implemented for process perturbation faults is applied, with faults injected independently at each network variable. All monitoring schemes are calibrated to operate at a fixed false alarm rate of α = 0.05 using independent fault-free validation datasets.

Fault detection sensitivity performance is quantified using a Normalized Area Under the Curve (AUC_{NORM}) metric. For a given faulty variable (see the Figure 1 as an example), the True Positive Rate (TPR) is computed for each fault magnitude and replicate. The median TPR across replicates is then calculated for each magnitude, yielding a single TPR vs. Fault magnitude curve. The area under this curve defines the AUC, which summarizes the sensitivity of the monitoring scheme across fault magnitudes. To ensure comparability across configurations, AUC values were normalized using a max-min approach, where the maximum AUC corresponds to perfect detection (TPR = 100% for all magnitudes) and the minimum was equivalent to the false positive rate (FPR) when no faults were introduced.

The evaluation pipeline applies this framework to each network and partition, as illustrated in Figure 2. The workflow starts with network selection and data generation under normal operating conditions, followed by partitioning for decentralized monitoring and control limit estimation. The TPR vs Fault magnitude curve is then computed for each variable under the selected partition, and the corresponding AUC_{NORM} is obtained. This procedure yields a single fault detection sensitivity score per variable, per network, and per partition.

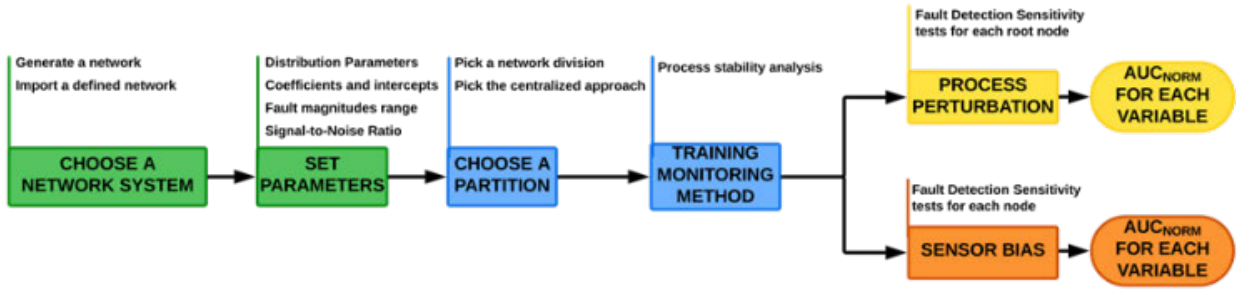

Figure 2. Workflow for Fault Detection Sensitivity evaluation.

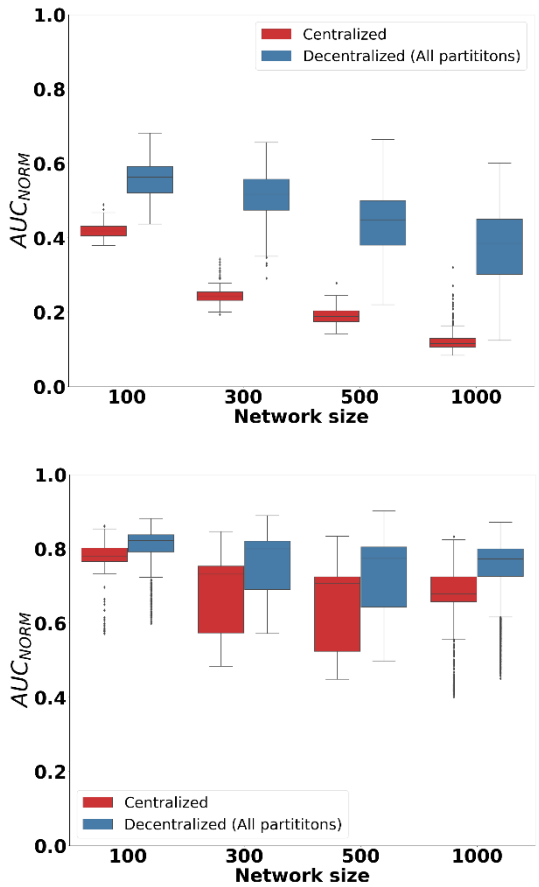

Figure 3. Comparison of AUC_{NORM} metric distributions obtained with centralized and decentralized monitoring schemes across different network sizes for process perturbation (top) and sensor bias (bottom) faults. Box plots aggregate AUC_{NORM} values over all faulty nodes and all partitions.

For each network size, six decentralized partitions are considered in addition to the centralized configuration. These partitions are selected to span a broad range of community sizes and modularity values, allowing a controlled analysis of the effect of the resolution parameter γ on fault detection sensitivity. Table 4 summarizes the resulting network partitions, reporting the number of communities, modularity values, and the number of faulty variables considered for both Process Perturbation and Sensor Bias faults.

RESULTS AND DISCUSSION

This section analyzes how network partitioning and topological descriptors impact fault detection sensitivity (AUC_{NORM}). Rather than focusing on absolute performance gains, the results aim to identify the structural factors that condition and drive fault detectability and that justify poor performances. Specifically, we show that fault detection is primarily governed by the size and number of communities, followed by internal and external connectivity, and finally by node-level causal factors.

Figure 3 compares the AUC_{NORM} distributions obtained with centralized and decentralized monitoring approaches across all network sizes. For both fault types, centralized monitoring exhibits a clear degradation in fault detection sensitivity as network size increases. This behavior reflects the well-known loss of sensitivity to localized faults in high-dimensional scenarios. In contrast, decentralized monitoring consistently attenuates the AUC_{NORM} degradation across all networks and partitions. The gain is more pronounced for process perturbation faults. Sensor bias faults also benefit from decentralization, although with a smaller relative improvement.

Figure 4 illustrates the relationship between AUC_{NORM} and the *number of nodes in the community* containing the faulty variable for both fault types. A strong and consistent trend is observed, with fault detection sensitivity decreasing monotonically as community size increases. This relationship holds across fault types and network sizes, indicating that community size acts as a first-order control parameter for fault detectability.

Figure 5 explores the relationship between partition *modularity* and fault detection sensitivity for process perturbation faults. Contrary to common assumptions, higher modularity does not systematically lead to improved fault detection performance. In several

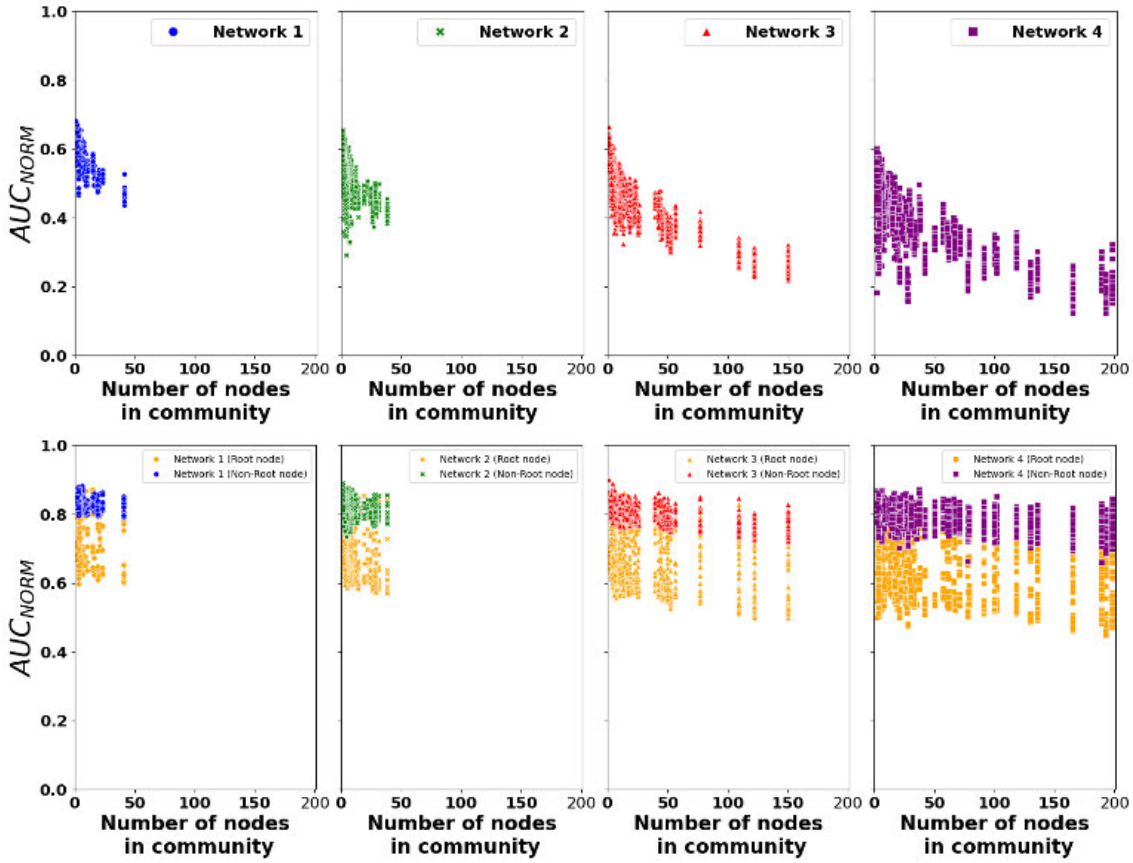

Figure 4. AUC_{NORM} vs. Number of nodes in community containing the faulty node for process perturbation (top) and sensor bias (bottom) faults. Each point represents the AUC_{NORM} for each faulty node.

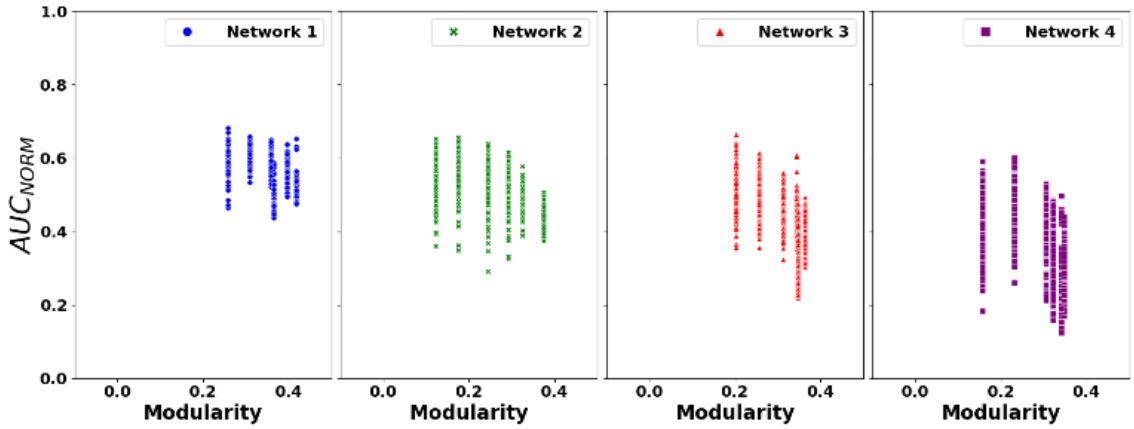

Figure 5. AUC_{NORM} vs. Modularity of the tested partition for process perturbation faults.

regimes, partitions with higher modularity exhibit lower AUC_{NORM} scores. This highlights the existence of an optimal range of modularity, that must be considered when establishing design principles for decentralized monitoring. Maximizing modularity alone may therefore yield structurally isolated communities that are suboptimal for monitoring. Sensor bias follows the same trend but is not present here due to space constraints.

The impact of internal community connectivity is analyzed in Figure 6, which relates AUC_{NORM} to the *density* of the community containing the faulty node for process perturbation faults. Higher internal density is associated with increased fault detection sensitivity, indicating that dense communities better preserve the statistical dependencies. A similar qualitative trend was observed for sensor bias faults. These findings reinforce that effective

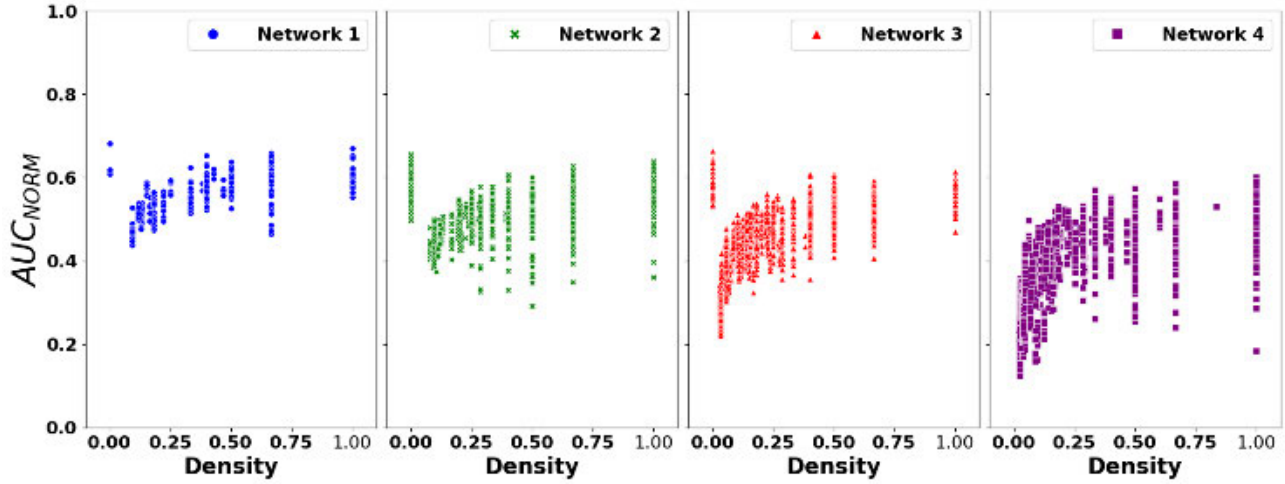

Figure 6. AUC_{NORM} vs. Density in the community containing the faulty node for process perturbation faults.

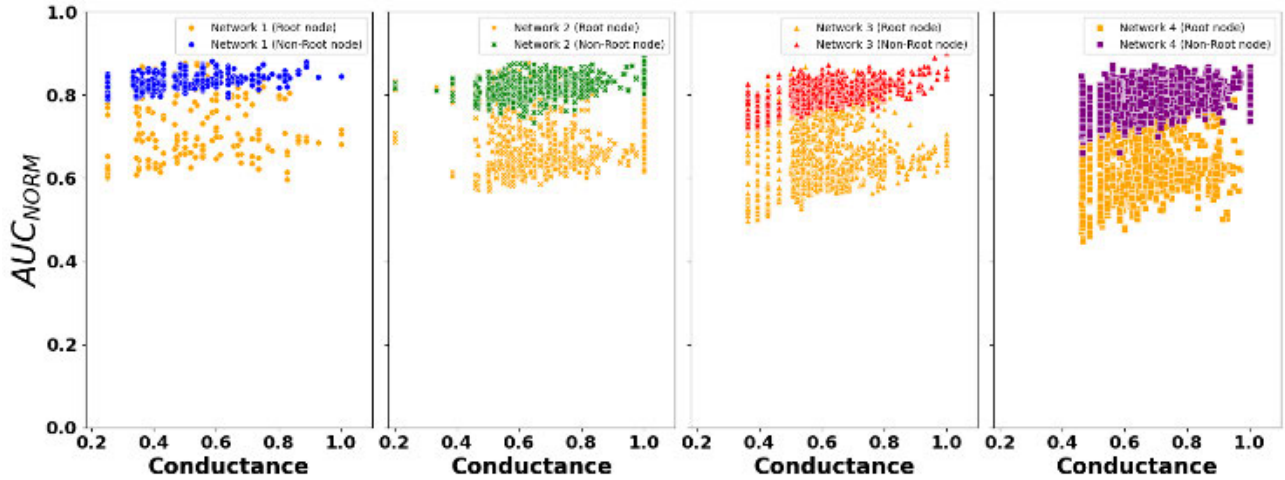

Figure 7. AUC_{NORM} vs. Conductance in the community containing the faulty node for sensor bias faults.

decentralized monitoring requires a balance between dimensionality reduction and sufficient internal connectivity.

Figure 7 focuses on the role of *conductance* for sensor bias faults. Communities with low conductance, corresponding to communities well-isolated from the rest of the network, tend to exhibit lower AUC_{NORM} values. Excessively isolated communities hinder the fault propagation, limiting the monitoring performance. It emphasizes that fault detection is not a purely local problem, but that an information exchange between communities is essential to preserve sensitivity. The same trend is observed in process perturbation faults.

Together, density and conductance reveal a fundamental trade-off: effective decentralized monitoring requires communities that are internally dense while remaining sufficiently coupled to the rest of the network to allow fault information to propagate.

In sensor bias faults, root-nodes lead to a broad range of AUC_{NORM} values, indicating that topological descriptors alone are insufficient to fully explain fault detection performance. Figure 8 relates AUC_{NORM} to the normalized causal strength ($\beta_{NORMALIZED}$) of the faulty node for sensor bias faults. Nodes with stronger causal influence on the network exhibit significantly higher detection sensitivity, even when community-level properties are comparable. This explains why purely topological descriptors fail to fully predict fault detectability and motivates the inclusion of causal metrics in topology-aware monitoring design.

Figure 9 presents an ablation study in which community-level information is selectively removed for a sensor bias fault affecting a representative node. Fault detection performance deteriorates sharply only when the community containing the faulty node is removed, while the exclusion of unrelated communities has a limited impact on sensitivity. For instance, in process perturbation faults, excluding up to 15 non-faulty communities

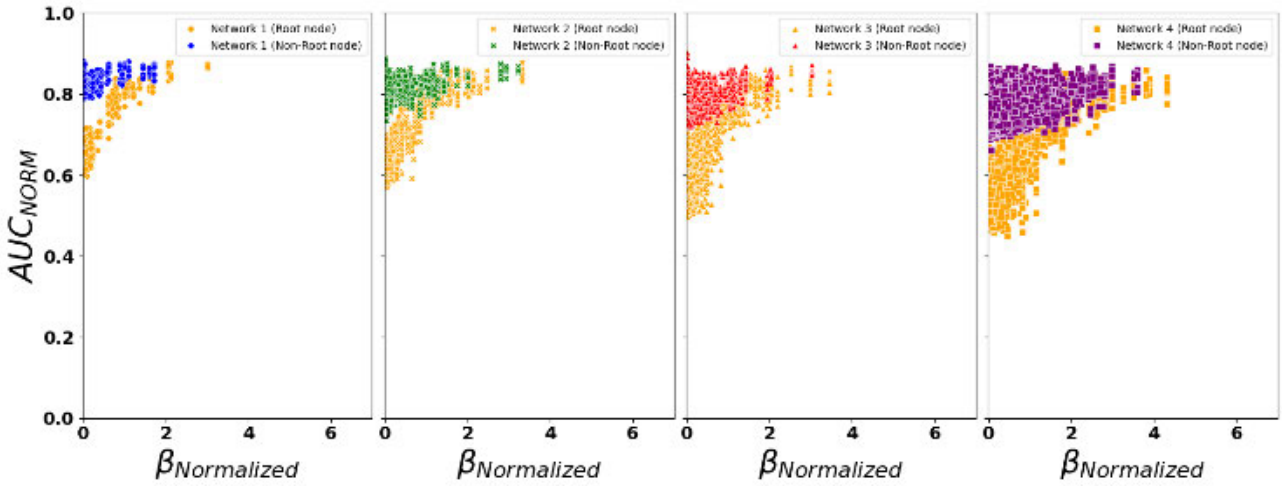

Figure 8. AUC_{NORM} vs. Normalized Causal Strength ($\beta_{NORMALIZED}$)of the faulty node for sensor bias faults.

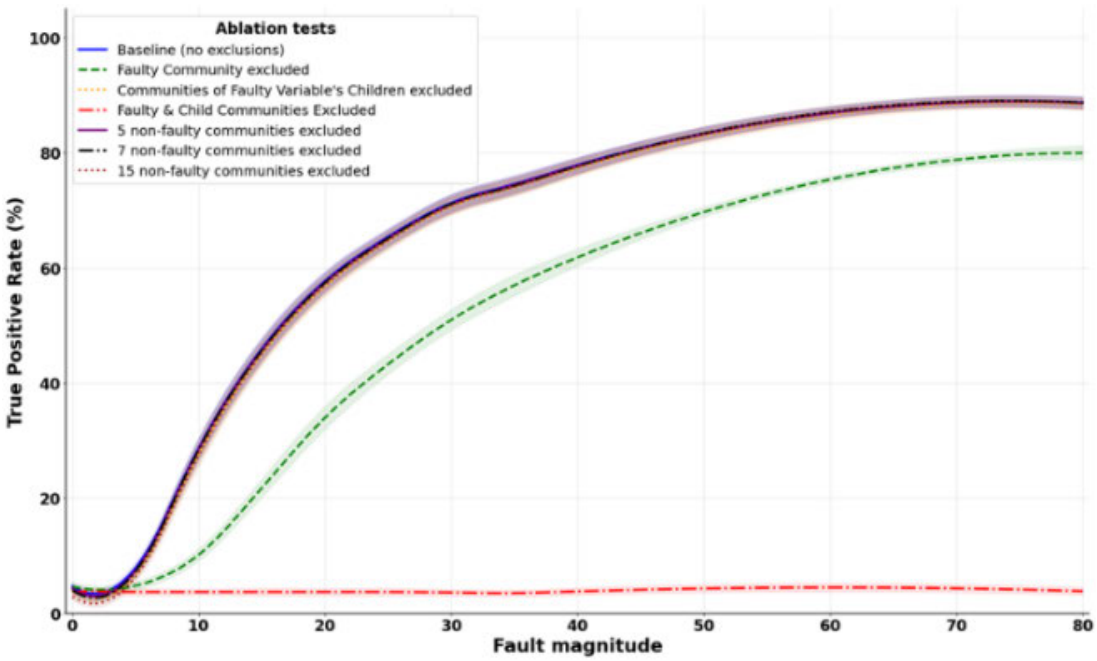

Figure 9. Median True Positive Rate vs Fault magnitude for sensor bias faults affecting node *X226* in Network 3, Partition #4, under different community-level information loss scenarios.

(representing over 32% of network communities) results in only marginal reductions in TPR at low fault magnitudes. This confirms that decentralized monitoring relies primarily on local causal information and is inherently robust to partial observability, a critical requirement for scalable industrial monitoring.

CONCLUSIONS

This study solidly establishes that decentralized causal monitoring substantially enhances fault detection sensitivity in high-dimensional networks compared to centralized approaches. The improvement is statistically significant ($p < 0.00001$) for both process perturbation and sensor bias faults, confirming that network partitioning itself provides a robust mechanism for mitigating the well-known loss of sensitivity in large-scale systems.

Community size and the number of communities emerge as the primary determinants of fault detection: smaller communities preserve sensitivity by limiting aggregation dilution, whereas excessively large communities or overly modular partitions reduce performance. Other topological drivers include internal connectivity and inter-community metrics. Higher community density,

conductance, and separability consistently improve fault detection, while excessive isolation impairs fault detection sensitivity.

Node-level causal influence, quantified by the normalized causal strength ($\beta_{NORMALIZED}$), explains residual variability in the root nodes. For sensor bias faults, root nodes with $\beta_{NORMALIZED}$ < 1.6 exhibit markedly reduced sensitivity, whereas higher $\beta_{NORMALIZED}$ secure high sensitivity through residual propagation to descendants.

Ablation studies further confirm the robustness of decentralized monitoring. Fault detection relies primarily on the local community containing the faulty node. Minimal degradation is observed when unrelated communities, even up to 32% of the network, are removed. For sensor bias, only the simultaneous loss of the faulty community and all descendant communities eliminates fault detection capability, emphasizing the interplay between causal influence and community structure.

Altogether, these results show a range of quantitative design principles for decentralized industrial monitoring: *(i)* maintain community sizes of few nodes (10–25 nodes), *(ii)* balance internal density with moderate inter-community connectivity, *(iii)* avoid excessive modularity (>0.35) and cohesiveness (>0.4), and *(iv)* identify and specially monitor root nodes with weak causal influence.

These guidelines not only support practical implementation of decentralized monitoring but also provide mechanistic foundations for predictive models of fault detection sensitivity, enabling topology-aware, adaptive, and resilient monitoring architectures.

Future work will extend this framework to nonlinear and nonstationary systems, exploring and testing whether the same community-based and causal principles hold in more complex, real-world industrial networks.

ACKNOWLEDGEMENTS

Rodrigo Paredes acknowledges the support from the doctoral Grant (PRT/BD/154267/2022) financed by the Portuguese Foundation for Science and Technology (FCT), under MIT Portugal Program. The authors acknowledge support from CERES funded by FCT (UID/00102/2025) and PRR – Recovery and Resilience Program of the Portuguese Republic (UID/PRR/00102/2025).

AUTHOR IDENTIFIERS

Author ORCIDs:
Paredes, R: 0000-0002-0447-3898
Reis, Marco S: 0000-0002-4997-8865

REFERENCES

1. Isermann R. Process fault detection based on modeling and estimation methods—a survey. Automatica 20:387-404 (1984). https://doi.org/10.1016/0005-1098(84)90098-0
2. Reis M, Gins G. Industrial process monitoring in the big data/industry 4.0 era: from detection, to diagnosis, to prognosis. Processes 5:35 (2017). https://doi.org/10.3390/pr5030035
3. Ge Z, Song Z. Distributed PCA model for plant-wide process monitoring. Ind. Eng. Chem. Res. 52:1947-1957 (2013). https://doi.org/10.1021/ie301945s
4. Peng X, Ding SX, Du W, Zhong W, Qian F. Distributed process monitoring based on canonical correlation analysis with partly-connected topology. Control Engineering Practice 101:104500 (2020). https://doi.org/10.1016/j.conengprac.2020.104500
5. Pearl J. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann (1988).
6. Chiang LH, Jiang B, Zhu X, Huang D, Braatz RD. Diagnosis of multiple and unknown faults using the causal map and multivariate statistics. Journal of Process Control 28:27-39 (2015). https://doi.org/10.1016/j.jprocont.2015.02.004
7. Paredes R, Rato TJ, Reis MS. Causal network inference and functional decomposition for decentralized statistical process monitoring: detection and diagnosis. Chemical Engineering Science 267:118338 (2023). https://doi.org/10.1016/j.ces.2022.118338
8. Li Y, He J, Chen C, Guan X. Topology inference for network systems: causality perspective and nonasymptotic performance. IEEE Trans. Automat. Contr. 69:3483-3498 (2024). https://doi.org/10.1109/tac.2023.3303816
9. Newman MEJ. Networks: An Introduction. Oxford Univ. Press (2010) https://doi.org/10.1093/acprof:oso/9780199206650.001.0001
10. Traag VA, Waltman L, van Eck NJ. From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports* 9:5233 (2019) https://doi.org/10.1038/s41598-019-41695-z
11. Rato TJ, Reis MS. Markovian and non-markovian sensitivity enhancing transformations for process monitoring. Chemical Engineering Science 163:223-233 (2017). https://doi.org/10.1016/j.ces.2017.01.047
12. Paredes R, Yang WT, Reis MS. Decentralized causal-based monitoring for large-scale systems: sensitivity and robustness assessment. IFAC-PapersOnLine 59:127-132 (2025). https://doi.org/10.1016/j.ifacol.2025.07.133
13. Yang WT, Reis MS, Borodin V, Juge M, Roussy A. An interpretable unsupervised bayesian network model for fault detection and diagnosis. Control Engineering Practice 127:105304 (2022). https://doi.org/10.1016/j.conengprac.2022.105304
14. Paredes R, Reis MS. Causality in process systems engineering: fundamentals, applications, and

emerging trends. Computers & Chemical Engineering 203:109345 (2025). https://doi.org/10.1016/j.compchemeng.2025.109345

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Control Structure Design of Novel Microwave-Catalyzed Process for Simultaneous Production of Ammonia and Ethylene

Md Mizanur Rahman[a], Omar Almaraz[a], Snehitha Baddam[a], Jianli Hu[a], and Srinivas Palanki[a]*

[a]Department of Chemical and Biomedical Engineering, West Virginia University,
Morgantown, WV 26506, USA
* Corresponding Author: srinivas.palanki@mail.wvu.edu

ABSTRACT

This work demonstrates the application of a pulsed microwave system for single-step co-production of ethylene and ammonia from methane. To mitigate inherent production fluctuations from pulsed microwave reactors, a staggered manifold configuration was utilized to stabilize effluent flow for industrial-scale compatibility. Dynamic validation of the ammonia and ethylene purification columns confirmed that a rigorously tuned control strategy effectively rejects ±10% feed disturbances while maintaining process stability and product purity. Ultimately, this systematic approach establishes a robust foundation for the sustainable, electrified production of foundational chemicals by bridging the gap between laboratory-scale pulsing phenomena and industrial-scale operational reliability.

Keywords: Process Control, Ethylene, Aspen Dynamics

INTRODUCTION

Ethylene (C_2H_4) is a simple, flammable, colorless, and highly reactive hydrocarbon containing a carbon–carbon double bond, which makes it the smallest and most important alkene. As one of the most widely produced organic chemicals, it is a foundational building block of the petrochemical industry and a crucial feedstock for manufacturing polyethylene plastics, ethanol, ethylene oxide, ethylene glycol, and vinyl chloride [1]. Ammonia (NH_3) is a colorless gas with a pungent odor, composed of one nitrogen atom bonded to three hydrogen atoms, with a boiling point of approximately −33°C at ambient pressure. It is highly soluble in water, forming ammonium hydroxide (NH_4OH), a weak base, and is a critically important industrial chemical with diverse applications across agriculture, energy, and industrial sectors [2]. Globally, more than 80% of ammonia production is used for nitrogen-based fertilizers such as urea, ammonium nitrate, and ammonium sulfate, which are essential for sustainable agricultural productivity, while additional applications include refrigeration systems, cleaning agents, and emerging roles as a hydrogen carrier for carbon-free energy storage, transportation, and fuel use [2]. Ethylene production currently tops 200 million metric tons annually, fueling a market exceeding $180 billion. This growth is sustained by a 4.1% rise in demand and a 6.2% increase in production capacity, largely propelled by the global requirement for plastics and consumer products [3]. In a similar vein, global ammonia output has surpassed 180 million metric tons per year, anchoring a fertilizer market valued at more than $100 billion while simultaneously representing nearly 2% of total global energy consumption [4]. The production of ethylene and ammonia represents the backbone of the modern chemical industry, though traditional methods remain significantly energy-intensive and carbon-heavy due to their reliance on high-temperature, endothermic reactions.

Microwave (MW) reactor technology represents a sophisticated process-intensification strategy that aligns with the global shift toward chemical process electrification as a replacement for traditional thermal heating [5]. By leveraging microwave catalysis, it is possible to simultaneously activate the highly stable CH_4 and N_2 molecules, enabling the direct co-production of ammonia and ethylene within a single reactor step [6]. This technology

not only enhances energy efficiency and reduces CO_2 emissions by selectively heating the catalyst surface but also introduces operational flexibility through pulsed microwave energy. By deploying microwave reactor technology, it is possible to operate the reactor at two different temperature modes by pulsing the microwave power in periodic intervals. Under microwave power-on mode (heating), an endothermic reaction to produce ethylene can take place whereas under the microwave power-off mode (quenching), an exothermic reaction to produce ammonia can occur.

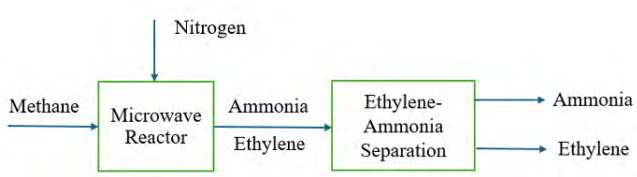

Figure 1. Block flow diagram of the simultaneous production of ethylene and ammonia.

The reactor outlet streams are combined and sent to the distillation column train for separation. During plant operation, due to switching between power on and power off modes in the two microwave reactors, the reactor outlets will not be steady. This fluctuation in the flowrate to the distillation column train leads to an inherently dynamic operation and hence it is necessary to design a control system for the separation section that keeps the quality of the products ethylene and ammonia at the desired levels. Based on the concept in Figure 1, a plant-wide steady-state model was constructed and subsequently converted into a dynamic model using Aspen Plus Dynamics. This model integrates a base control strategy consisting of tuned PID loops and parameters to provide a robust foundation for industrial-scale process simulations.

It was recently shown by our group that producing ethylene via novel microwave (MW) technology is economically competitive, with a levelized cost of USD 0.51/kg compared to USD 0.56/kg for conventional methods [7]. While a formal technoeconomic analysis (TEA) for the simultaneous production of ethylene and ammonia is not part of this study, this integrated process has the potential to achieve higher net present value (NPV) than standalone ethylene production since the hydrogen required for ammonia production is produced in the ethylene reaction cycle.

MICROWAVE REACTOR DESIGN ISSUES

Conventional processes to produce ethylene mostly operate at equilibrium conditions with a continuous heating process to supply the heat requirement for the reaction conditions. But a pulsed microwave heating process is a non-equilibrium process that has the potential to switch between exothermic and endothermic reactions together for simultaneous production of two commodity chemicals. In the pulsing process, the power-on mode (heating) favors endothermic reactions, and the power-off mode (quenching) favors exothermic reactions. The main reactions involved are as follows:

$$2CH_4 \rightarrow C_2H_4 + 2H_2 \text{ (endothermic, 500-700}^{o}\text{C)}$$

$$3H_2 + N_2 \rightarrow 2NH_3 \text{ (exothermic, 250-400}^{o}\text{C)}$$

As shown experimentally in Rahman et al. [7], small quantities of additional compounds such as acetylene, ethane, carbon dioxide, benzene, and toluene are also produced via side reactions, which need to separated from the main products via a train of distillation columns. The methane coupling reaction produces ethylene and hydrogen comparatively at higher temperature range (500-700°C) in power-on mode. The produced hydrogen reacts with nitrogen at lower temperature (250-400°C) to produce ammonia in the power-off mode. The transition from laboratory-scale pulsed microwave (MW) reactors to industrial-scale production introduces significant complexities in reactor design and process control, primarily due to the inherent fluctuations in product flow that alternately produce ethylene and ammonia. In a continuous industrial facility, equipment such as compressors, separation columns, and heat exchangers is designed for steady-state flows. Rapid fluctuations in mass flow can lead to severe pressure surges, thermal cycling fatigue, and significant challenges in maintaining product purity within distillation sumps. To rectify these fluctuations and stabilize the production rate, a process intensification strategy involving a series of MW reactors operating in a staggered, out-of-phase configuration must be implemented. By coordinating the pulsing intervals of multiple reactors, the peak production of one unit can compensate for the quenching phase of another, effectively dampening the oscillation and delivering a more consistent feed to the downstream processes. Despite this staggered operation, the outlet of the microwave reactor train can fluctuate up to ±10% and it is necessary to develop a control system for the separation train to work effectively despite this fluctuation.

METHODOLOGICAL FRAMEWORK

A systematic approach was adopted integrating steady-state modeling, dynamic simulation, and advanced process control. Initially, steady-state benchmarks for the intensified microwave-assisted routes were developed in Aspen Plus v14.0 using the NRTL property model [7]. This thermodynamic selection is critical to accurately account for the highly non-ideal liquid-

phase activity coefficients of ammonia within the complex hydrocarbon and light-gas matrix. The outlet of the microwave reactor train has the composition of methane, ethylene, hydrogen, nitrogen, ammonia, acetylene, ethane, carbon dioxide, benzene, and toluene that were determined experimentally and fluctuates with microwave pulsing [5, 6, 7]. Then, a dynamic model in Aspen Plus Dynamics (APDs) v14, was constructed for the separation train based on the steady-state process model [8]. During this stage, all equipment in the process were appropriately sized. Building upon the steady-state architecture, a high-fidelity dynamic model was developed in Aspen Plus Dynamics (APD) v14.0, where all unit operations were sized based on established literature [7, 8]. This phase included the implementation of a foundational control layer and the rigorous fine-tuning of control parameters to achieve optimal response characteristics across diverse operational transients. To validate the model's reliability, it was subjected to extensive disturbance testing, specifically focusing on variations in feed flow rates and compositions [8]. The effectiveness of this control strategy was verified through robustness evaluations under dynamic scenarios, demonstrating a consistent ability to preserve optimal operating conditions.

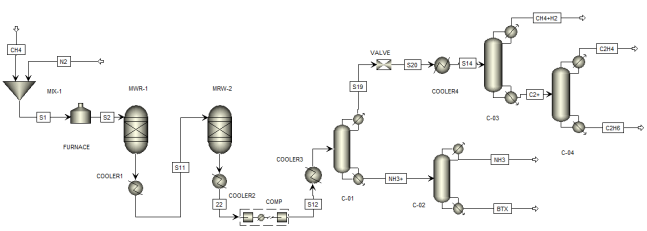

Figure 2. Process Flow Diagram of Microwave Assisted Process to Convert Methane to Ethylene and Ammonia.

Figure 2 illustrates the process flow diagram (PFD) that produces ethylene and ammonia simultaneously. In this industrial-scale simulation, a feed stream of methane and nitrogen is initially processed in a yield reactor at 600°C to facilitate co-production. To prevent ammonia decomposition, the effluent is immediately quenched to 280°C before entering a secondary hydrogenation reactor, which operates at the same temperature to convert byproduct acetylene into ethylene, thereby enhancing overall yield. The resulting stream is pressurized to 15 bar via a multistage compressor and fed into the first distillation column (C_01), operating at -15°C and 15 bar. The ammonia-rich bottom stream from C_01 is directed to a purification column (C_02) operating at 39°C and 15 bar to recover high-purity ammonia. On the other hand, the ethylene and methane-rich overhead stream is processed through two subsequent distillation columns (C_03 and C_04) for ethylene purification, following the same separation technique by Rahman et al [7]. A buffer tank to smooth out the reactor outlet fluctuations created by pulsing was not considered primarily because gas-phase streams are highly compressible and occupy large volumes, making physical storage for stabilization both economically and spatially impractical. To evaluate the direct conversion efficiency of the microwave-assisted process, no recycle streams were utilized in this study. Under these conditions, a methane feed of 2500 kg/h successfully yielded 2122 kg/h of ethylene and 955 kg/h of ammonia, both achieving a 99.9% purity level. The rigorous design specifications for the entire four-column separation train are summarized in Table 1.

Table 1: Feed and Column Design Specifications.

	C_01	C_02	C_03	C_04
Purpose	Distillation	Distillation	Distillation	Distillation
No. of stages	18	24	17	51
Feed stage	12	15	10	20
Pressure (bar)	15	15	8	8
Reflux ratio	0.317	0.358	0.220	2.690
Condenser duty (MW)	-1.850	-0.416	-0.670	-0.650
Condenser temperature (ºC)	-97.77	38.53	-131.62	-58.40
Distillate rate (kmol/h)	1610.71	55.98	1511.16	75.56
Reboiler duty (MW)	0.201	0.418	0.490	0.900
Reboiler temperature (ºC)	39.13	177.34	-48.45	-29.38
Bottom rate (kmol/h)	57.30	1.31	99.53	23.93
Tray spacing (m)	0.61	0.61	0.61	0.61
Tray type	Sieve	Sieve	Sieve	Sieve
Column height (m)	9.75	13.41	9.14	23.77
Diameter (m)	1.28	0.28	1.33	0.68

A dynamic model was developed to account for the fluctuations in the reactor outlet stream. In this simulation, the physical dimensions of the unit operations were established using industry-scale design principles. Distillation column heights were calculated through established heuristics and relevant correlations, while diameters were determined based on estimated flooding velocities to ensure operational stability [8]. Following the completion of equipment sizing, the validated steady-state simulation model was transitioned into a dynamic framework using Aspen Plus Dynamics V14 [8]. To ensure operational stability and robust disturbance rejection, a plant-wide control structure was implemented. This strategy integrates conventional regulatory loops with advanced regulatory schemes, where all controller parameters were optimized using the Tyreus–Luyben autotune variation (ATV) method adjusting the controller's gain (K_C) integral time constant (τ_I) and derivative time constant (τ_D) [9]. As illustrated in Figure 3, each distillation column implements a standard control configuration, including top pressure control, reflux drum level control,

and sump level control. For column C_01, condenser temperature is regulated through utility flow adjustment, the reflux drum level is controlled via reflux flow manipulation, and the bottom-stage temperature is maintained by varying the reboiler duty. Additionally, a sump level is controlled by applying a sump level controller at the bottom. An identical control philosophy is implemented for the remaining columns (C_02, C_03, and C_04).

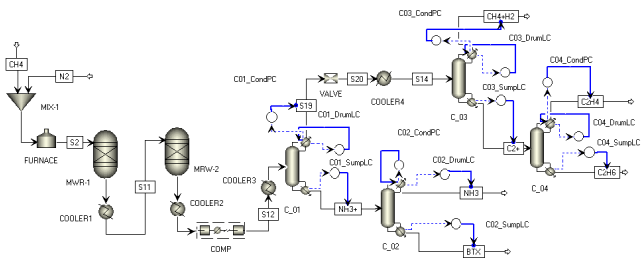

Figure 3. Plant wide control strategy for simultaneous ammonia and ethylene production.

RESULTS AND DISCUSSION

This study evaluates the dynamic performance and controller structure of the ammonia (C_02) and ethylene (C_04) purification columns to ensure process stability during simultaneous production. Given the fluctuating nature of the upstream pulsed microwave reactor effluent, achieving a robust control response in the separation train is critical for maintaining product purity and operational safety.

The initial controller response of column C_02 & C_04 exhibited significant oscillations and sluggish stabilization in both the reflux drum (C_02DrumLC) and the column sump (C_02SumpLC) levels. As shown in Figure 4, the mass flow rates and temperatures showed substantial deviations before reaching a settled state. To mitigate process instability, a rigorous tuning phase was conducted. The controller outputs were optimized to aggressively dampen oscillations, primarily utilizing Proportional-Integral (PI) control structures. The controller gains K_C and τ_I were evaluated using the Tyreus–Luyben ATV method. Table 2 summarizes the optimized parameters for the reflux drum and sump level controllers for both purification columns.

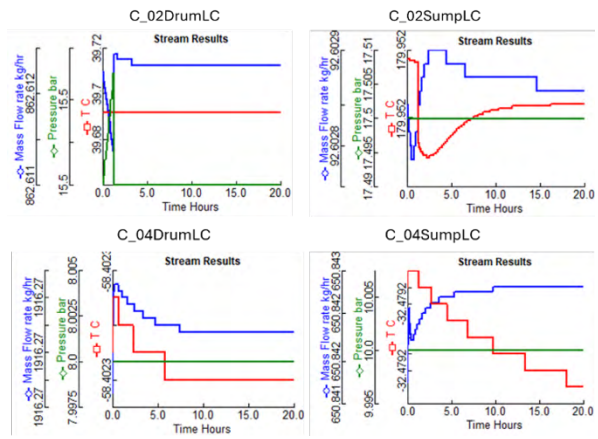

Figure 4. Controller response of ammonia and ethylene purification column (C_02 & C_04) before tuning.

Table 2. Optimized Controller's Gain and Integral Time Constant

Column	Controller Location	Gain (K_C)	Integral Time (τ_I)
C_02 (Ammonia)	Reflux Drum	32.55	10.56
	Sump	24.73	3.96
C_04 (Ethylene)	Reflux Drum	46.35	3.96
	Sump	25.67	2.64

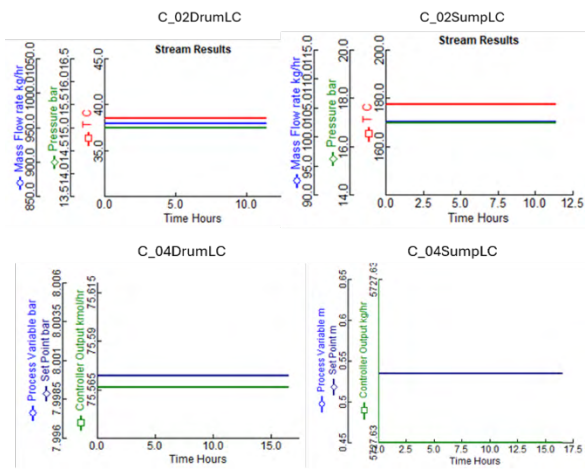

Figure 5. Controller response of ammonia and ethylene purification column (C_02 & C_04) after tuning.

Figure 5 shows that both the drum and sump levels reached their respective setpoints with minimal overshoot and zero steady-state error. This stable baseline is essential for the intensified process, as it ensures that the energy-intensive separation of ammonia and ethylene from the hydrocarbon matrix remains consistent despite the inherent transients of the microwave-assisted route

Controller Performance Analysis

To validate the reliability of the established control strategy, the column was subjected to step changes of ±10% in the feed flow rate. As illustrated in Figure 6, the C_02 control system effectively rejected these disturbances with minimal deviation. Upon a +10% increase in feed, the condenser and reboiler duties adjusted dynamically to stabilize inventory levels, reaching a new steady state shortly after the onset of the disturbance. This transition resulted in a proportional impact on the top and bottom product flow rates, directly reflecting the change in feed throughput. Similarly, the -10% feed variation demonstrated a smooth transition in both the top and bottom product flow rates. This robustness confirms that the integrated control layer can effectively manage process fluctuations while maintaining product quality and system stability.

To verify the reliability of the control strategy for the C_04 column, the system was tested with ±10% step changes in the feed flow rate. As shown in Figure 7, the control loops successfully rejected these disturbances, keeping all process variables stable. When the feed increased by 10%, the reboiler and condenser duties adjusted automatically to maintain the liquid inventory in the sump and reflux drum. This allowed the system to reach a new steady state shortly after the change, with the top and bottom product flow rates increasing to match the higher throughput. A similarly smooth transition occurred during the -10% feed reduction, where the system adjusted to a lower-load operating point without any instability. These results confirm that the integrated control layer effectively preserves column stability and product purity, even under the fluctuating conditions typical of electrified processes.

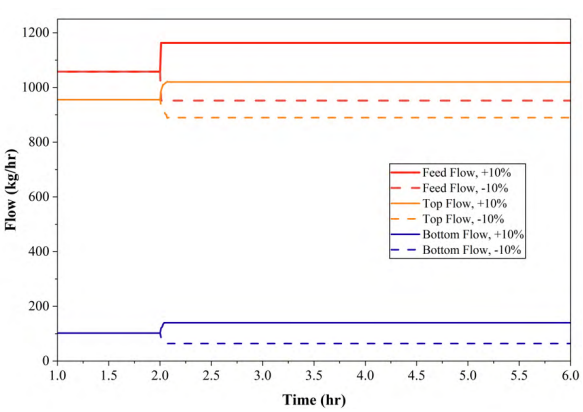

Figure 6. Controller response of the ammonia purification column (C_02) to a ±10% step change in feed flow rate

Controller Sensitivity Matrix for Ammonia Purification Column

Table 3 and corresponding Figure 8 illustrate a parametric sensitivity analysis of the ammonia (C_02) purification column, where Case-1 (Optimized TL-ATV) serves as the performance benchmark. In the dynamic simulation graphs, Case-1 is presented as the most efficient response, that returns the reflux and sump levels to their setpoints with the shortest settling time and minimal overshoot. In contrast, the low-gain scenarios (Case-2 and Case-5) appear on the graph as sluggish, long-duration curves; because a K_C of 5 lacks the proportional response to move control valves decisively, these cases take significantly more time to stabilize and allow for large deviations from the setpoint. On the other hand, the high-gain scenarios (Case-3 and Case-4) demonstrate better response compared to case-2 and case-5.

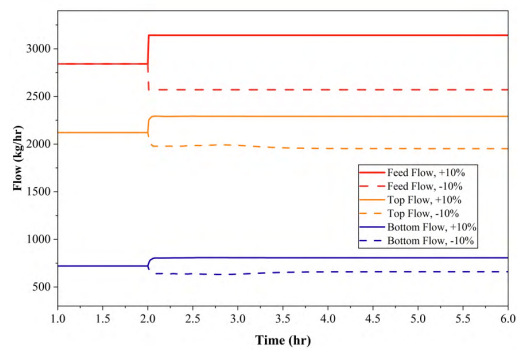

Figure 7. Controller response of the ethylene purification column (C_02) to a ±10% step change in feed flow rate

Table 3. Controller Sensitivity Matrix for Ammonia Column (C_02)

test Case	reflux drum Level (K_C, τ_I)	sump level (K_C, τ_I)
case-1: optimized (TL-ATV)	K_C = 32.55, τ_I =10.56	K_C = 24.73, τ_I = 3.96
case-2: low gain / fast reset	K_C = 5, τ_I = 10.56	K_C = 5, τ_I = 3.96
case-3: high gain / fast reset	K_C = 60, τ_I = 10.56	K_C = 60, τ_I = 3.96
case-4: high gain / Slow Reset	K_C = 60, τ_I = 60	K_C = 60, τ_I = 60
case-5: low gain / slow reset	K_C = 5, τ_I = 60	K_C = 5, τ_I = 60

case-1

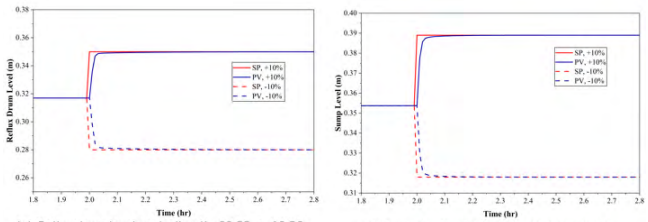

(a) Reflux drum level controller: K_c=32.55, τ_I=10.56 (b) Sump level controller: K_c=24.73, τ_I=3.96

case-2

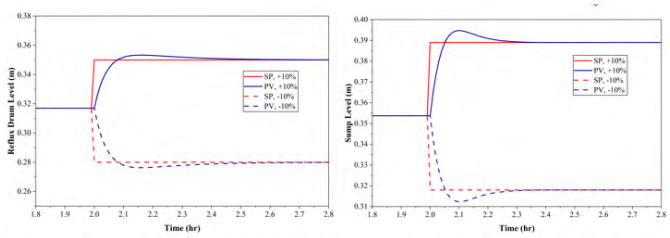

(a) Reflux drum level controller: K_c=5, τ_I=10.56 (b) Sump level controller: K_c=5, τ_I=3.96

case-3

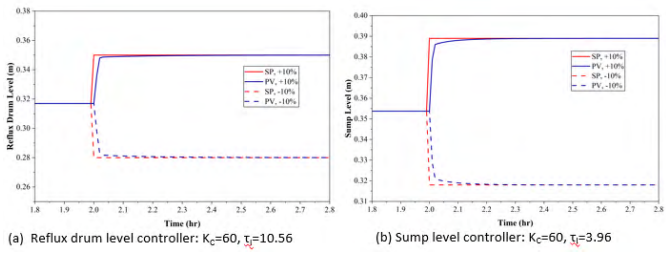

(a) Reflux drum level controller: K_c=60, τ_I=10.56 (b) Sump level controller: K_c=60, τ_I=3.96

case-4

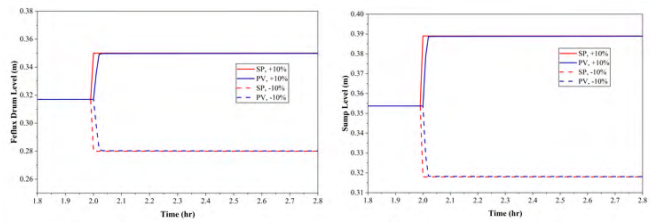

(a) Reflux drum level controller: K_c=60, τ_I=60 (b) Sump level controller: K_c=60, τ_I=60

case-5

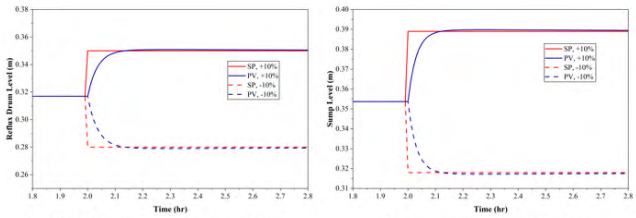

(a) Reflux drum level controller: K_c=5, τ_I=60 (b) Sump level controller: K_c=5, τ_I=60

Figure 8. Controller response of ammonia column (C_02)

Controller Sensitivity Matrix for Ethylene Purification Column

Table 4 and the corresponding Figure 9 illustrate the parametric sensitivity analysis for the ethylene purification column (C_04), where Case-1 (Optimized TL-ATV) consistently serves as the most efficient performance benchmark.

In the dynamic simulation graphs, Case-1 demonstrates the superior control response, characterized by high proportional gain and tight integral times (τ_I = 3.96 and 2.64) that restore the reflux and sump levels to their setpoints with the shortest settling time and minimal overshoot. In contrast, the low-gain scenarios (Case-2 and Case-5) appear as sluggish, long-duration curves; because a K_C of 5 lacks the proportional magnitude to drive the control valves effectively, these cases result in significantly prolonged stabilization periods and allow for large deviations from the setpoint. While the high-gain scenarios (Case-3 and Case-4) demonstrate a more robust initial response compared to the low-gain cases.

CONCLUSIONS

This study presents a comprehensive framework for the electrification of commodity chemical manufacturing by integrating steady-state and high-fidelity dynamic modeling to identify and rectify the associated challenges. The work successfully demonstrates a single-step co-production pathway that leverages microwave-assisted technology to simultaneously activate stable methane and nitrogen molecules for ethylene and ammonia synthesis. To address the inherent production fluctuations of pulsed microwave reactors, the study highlights a staggered manifold configuration to stabilize effluent flow, ensuring compatibility with industrial-scale downstream operations. Dynamic validation exemplified by the ammonia and ethylene purification column (C_02 & C_04) confirmed that a rigorously tuned control strategy can effectively reject ±10% feed disturbances while maintaining process stability and product purity. Ultimately, this systematic approach bridges the gap between laboratory-scale phenomena and industrial-scale operational reliability, establishing a robust foundation for the sustainable, electrified production of foundational chemicals

Table 4. Controller Sensitivity Matrix for Ethylene Column (C_04)

test Case	reflux drum Level (K_C, τ_I)	sump level (K_C, τ_I)
case-1: optimized (TL-ATV)	K_C = 46.35, τ_I = 3.96	K_C = 25.67, τ_I = 2.64
case-2: high gain / fast reset	K_C = 5, τ_I = 3.96	K_C = 5, τ_I = 2.64
case-3: high gain / slow reset	K_C = 80, τ_I = 3.96	K_C = 60, τ_I = 2.64
case-4: low gain / fast reset	K_C = 46.35, τ_I = 80	K_C = 25.64, τ_I = 60
case-5: low gain / slow reset	K_C = 5, τ_I = 80	K_C = 5, τ_I = 60

Case-1

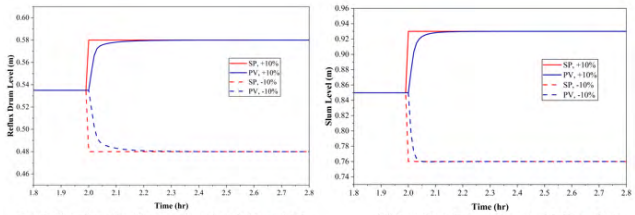

(a) Reflux drum level controller: K_c=46.35, τ_I=3.96 (b) Sump level controller: K_c=25.67, τ_I=2.64

case-2

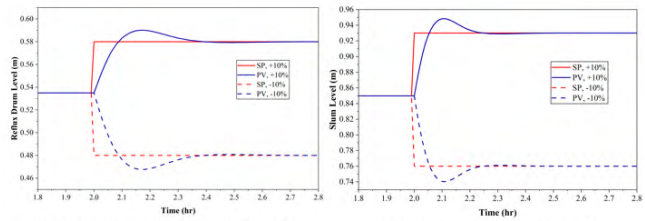

(a) Reflux drum level controller: K_c=5, τ_I=3.96 (b) Sump level controller: K_c=5, τ_I=2.64

case-3

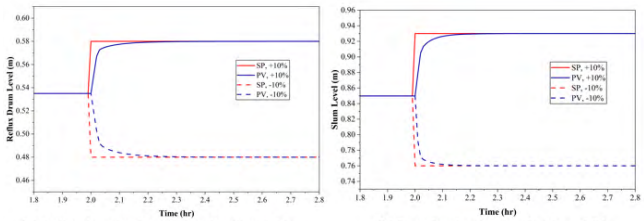

(a) Reflux drum level controller: K_c=80, τ_I=3.96 (b) Sump level controller: K_c=60, τ_I=2.64

case-4

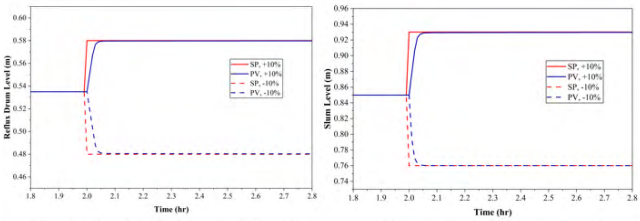

(a) Reflux drum level controller: K_c=46.35, τ_I=80 (b) Sump level controller: K_c=25.64, τ_I=60

case-5

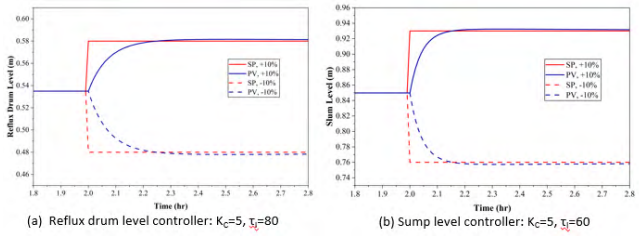

(a) Reflux drum level controller: K_c=5, τ_I=80 (b) Sump level controller: K_c=5, τ_I=60

Figure 9. Controller response of ethylene purification column (C_04) for different gains

ACKNOWLEDGEMENTS

This study is supported by the United States Department of Energy (DOE) under grant no. DE-EE0011195.

REFERENCES

1. Chen Y, Kuo MJ, Lobo R, Ierapetritou M. Ethylene production: process design, techno-economic and life-cycle assessments. Green Chem. 26:2903-2911 (2024). https://doi.org/10.1039/d3gc03858k
2. Retablo T. Ammonia, production, applications, and the effect of its phase-out [Internet]. Available from: https://epocbelgium.be/sites/epoc/files/Ammonia%20phase%20out%20consequence%20_%202.pdf
3. Chauhan R, Sartape R, Minocha N, Goyal I, Singh MR. Advancements in environmentally sustainable technologies for ethylene production. Energy Fuels 37:12589-12622 (2023). https://doi.org/10.1021/acs.energyfuels.3c01777
4. Elgowainy A, Mintz M, Lee U, Stephens T, Sun P, Reddi K, Zhou Y, Zang G, Ruth M, Jadun P, Connelly E, Boardman R. Assessment of Potential Future Demands for Hydrogen in the United States [Internet]. 2020 Oct [cited 2025 Nov 15] p. ANL--20/35, 1710201, 163944. Report No.: ANL--20/35, 1710201, 163944.
5. Available from:https://www.osti.gov/servlets/purl/1710201/
6. Baddam SR, Jiang C, Poreddy MR, Tewari K, Robinson B, Wang Y, Palanki S, Hu J. Microwave-driven nonoxidative and selective conversion of methane to ethylene over mn-based catalysts. Ind. Eng. Chem. Res. 64:22102-22114 (2025). https://doi.org/10.1021/acs.iecr.5c02894
7. Tiwari S, Khan TS, Tavadze P, Hu J. Activation of two highly stable molecules – nitrogen and methane to co-produce ammonia and ethylene. Chemical Engineering Journal 413:127501 (2021). https://doi.org/10.1016/j.cej.2020.127501
8. Rahman MM, Haque ME, Baddam SR, Hu J, Palanki S. Technoeconomic analysis of a novel microwave process to produce ethylene from methane. ACS Omega 11:8376-8390 (2026). https://doi.org/10.1021/acsomega.5c11191
9. Haque ME, Palanki S. Advanced process control strategies for efficient methanol production from natural gas. Processes 13:424 (2025). https://doi.org/10.3390/pr13020424
10. Luyben WL. Distillation design and control using Aspen simulation. 2nd ed. Hoboken, N.J:

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Towards Safety-Intelligent Cyber-Physical Systems: A Real-time Monitoring and Control Framework

Zhane Ann Tizon[abc], Yuanxing Liu[abc], Sahithi Srijana Akundi[abc], Austin Braniff[d], Beatriz Dantas[d], Yuhe Tian[d], Faisal I. Khan[bc], Efstratios N. Pistikopoulos[ab]*

[a] Texas A&M Energy Institute, Texas A&M University, College Station, TX, USA
[b] Artie McFerrin Department of Chemical Engineering, Texas A&M University, College Station, TX, USA
[c] Mary Kay O'Connor Process Safety Center (MKOPSC), Texas A&M University, College Station, TX, USA
[d] Department of Chemical and Biomedical Engineering, West Virginia University, Morgantown, WV, USA
* Corresponding Author: stratos@tamu.edu.

ABSTRACT

A safety-intelligent framework is presented for developing a multiple-input multiple-output (MIMO) risk-based explicit model predictive control (R-eMPC) for metal hydride storage systems (MHSS). These systems are susceptible to thermal runaway during the charging process as a result of the exothermic adsorption reaction within the metal alloy. To address this issue, deterministic and stochastic safety-intelligent control algorithms are designed and implemented by explicitly embedding a dynamic risk index (RI) or risk tolerance (α) into the control law and decision-making. In closed-loop analysis, the deterministic R-eMPC regulates both core temperature and hydrogen storage capacity by forecasting fault occurrence, triggering alarms, and reducing the risk index by adjusting the optimal control actions, supply pressure and water flowrate. Meanwhile, the stochastic R-eMPC accounts for uncertainties in core temperature variation by incorporating risk tolerance through chance-constraints. When the temperature exceeds the strict operational limit, stochastic R-eMPC enables continued operation beyond conservative safety boundaries via dynamic updated probabilistic limits. Closed-loop responses demonstrate that the safety-intelligent controllers effectively minimize evolving thermal risks while maximizing storage capacity during charging operation of the MHSS. Overall, the risk-based controllers improve both safety and charging performance, demonstrating the potential of the proposed framework for real-time, safety-intelligent operation of metal hydride systems.

Keywords: hydrogen, metal hydrides, model predictive control, process safety, multi-parametric programming, cyber-physical system

INTRODUCTION

Chemical and energy processes undergo different types of reactions, such as exothermic and endothermic reactions, during operation. When these reactions become uncontrollable, they can cause runaway incidents, leading to catastrophic damages such as fires or explosions, resulting in serious process safety issues [1]. Therefore, it is critical that reactive processes are properly monitored and controlled. However, most existing process control systems do not explicitly incorporate holistic safety metrics, instead treating safety through separate supervisory layers.

Hydrogen energy systems, encompassing production, storage, and utilization are inherently susceptible to hazardous events. Hydrogen storage represents a major safety bottleneck due to high energy density and reactive nature of hydrogen. Solid-state hydrogen storage using metal hydrides has attracted significant attention for fuel cell electric vehicles (FCEVs) because it can store hydrogen at relatively low pressures [2]. However, charging or adsorption process of metal hydride storage system (MHSS) is exothermic, where significant amount of heat is released, leading to rapid temperature rise. The reaction can be expressed in Eq. (1).

$$M + \frac{x}{2}H_2 \leftrightarrow MH_x + \Delta H \tag{1}$$

where M is metal alloy, H_2 is hydrogen gas, MH_x is metal hydride formation, and ΔH is heat of enthalpy.

In transportation application, the US Department of Energy (US DOE) specifies that the operating temperature of hydrogen storage should remain within between 233 K to 358 K [3]. Exceeding this temperature window may degrade system performance or compromise safety. Consequently, effective thermal management and control strategies are essential for MHSS.

Ali et al. [4] introduced a deterministic risk-based explicit MPC (R-eMPC) in which dynamic risk modeling as a function of a safety-critical variable was embedded in the control optimization. The R-eMPC was able to trigger the alarm system ahead of time to ensure enough time for the operators to shut down the system before any safety events occur. Akundi et al. [5] extended the safety-control theory by explicitly accounting for uncertainties by using Bayesian inference techniques and chance-constraints programming. In this study, the controller is allowed to be able to operate by exceeding the safety limit, thereby providing more flexibility in the presence of uncertainties. More recently, Liu et al. [6, 7] demonstrated the applicability of both online risk-based MPC and offline R-eMPC in a cyber-physical system of proton exchange membrane water electrolysis (PEMWE). The dynamic risk index was shown to effectively quantify evolving thermal risk and support real-time safety-aware operation.

Despite these advances, risk-based explicit MPCs have not yet been developed or validated for metal hydride hydrogen system (MHSS), where strongly exothermic reactions and strict thermal limits pose significant runaway risks. Hence, this study aims to develop a safety-intelligent control framework using deterministic and stochastic risk-based R-eMPC for MHSS, with particular emphasis on charging operation, while maintaining applicability to discharging operation and addressing the trade-offs between safety and performance.

METAL HYDRIDE STORAGE SYSTEM

This study focuses on metal hydride storage system (MHSS) to store hydrogen gas. Figure 1 demonstrates the schematic diagram of the charging process in MHSS with the metal hydride reactor, water tank, and process instrumentations. This analysis follows the schematic for modeling and simulation of the charging (adsorption) process, in which hydrogen gas is supplied from an external source, e.g. water electrolysis or refueling station. Since the reaction is exothermic and releases significant amount of heat, thermal management is incorporated through water supply is embedded in this system. Further, accurate modeling and real-time control implementation are crucial to ensure both safety and performance of the MHSS.

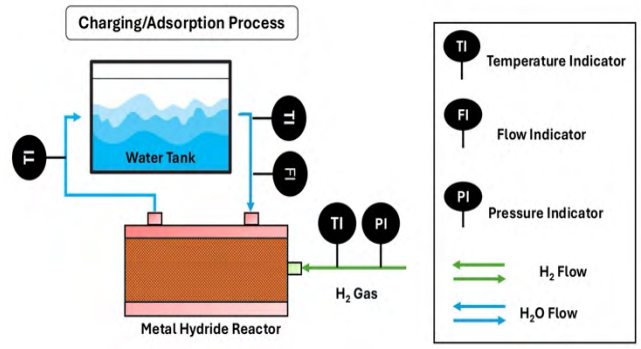

Figure 1. Schematic diagram of charging operation of metal hydride storage system.

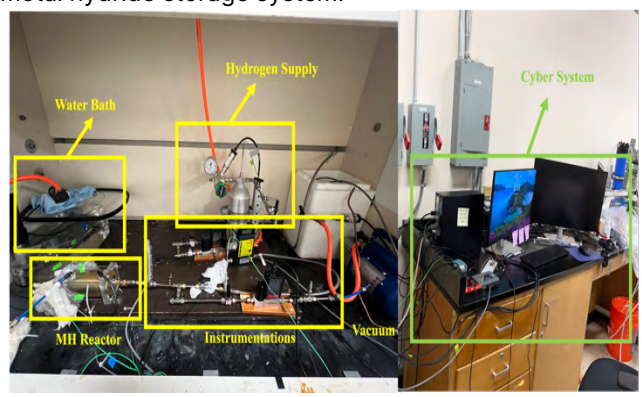

Figure 2. Experimental setup of metal hydride storage system.

High-Fidelity Model

The metal hydride storage model is developed using first-principles based on thermodynamics, reaction kinetics, and heat transfer mechanisms [8, 9, 10]. The governing equations are presented in Eqs. (2)-(10). These equations capture the coupled behavior of hydrogen uptake, reaction heat generation, and heat exchange between the metal hydride bed, reactor wall, and cooling medium.

$$\frac{dm_{H2,g}}{dt} = f_{in,H2} - rm_s \frac{MW_{H2}SC}{MW_{MH}} \tag{2}$$

$$m_{MH} = rm_{solid} \tag{3}$$

$$\left(m_{H2,g}C_{p,H2} + m_s C_{p,s}\right)\frac{dT_{MH}}{dt} = Q - \Delta H_a r_{ms}\frac{SC}{MW_{MH}} \tag{4}$$

$$r = C_a e^{-\frac{Ea}{RT}} \ln\left(\frac{Pa}{Peq}\right)\left(1 - \frac{m_{MH}}{m_s}\right) \tag{5}$$

$$P_{eq} = e^{\left(\frac{\Delta Ha}{RT} - \frac{\Delta Sa}{R} + 0.5\left(\pi\left(\frac{m_{MH}}{m_s} - 0.5\right)\right) + 0.1\right)} P_0 \tag{6}$$

$$m_{wall}C_{p,wall}\frac{dT_{wall}}{dt} = -h_{wall}A_{wall}(T_{wall} - T_{MH}) + Q \tag{7}$$

$$Q = f_{H2O}C_{p,H2O}\left(T_{H_2O,in} - T_{wall}\right) \tag{8}$$

$$v_j \rho_{H2O} C_{p,H2O} \frac{dT_{H2O,out}}{dt} = f_{H2O} C_{p,H2O} \left(T_{H_2O,in} - T_{H2O,out}\right) - Q \tag{9}$$

$$m_{H2,g} = \frac{PV_g}{RT} MW_{H2} \tag{10}$$

Cyber-Physical System

The experimental setup of the metal hydride storage system (MHSS) is illustrated in Figure 2. It consists of several integrated units, including a hydrogen supply (canister tank), metal hydride reactor, water bath, and associated process instrumentation. The system also incorporates a computer, data acquisition hardware, and a microcontroller to enable real-time data exchange and signal processing between the physical and cyber layers.

The cyber component primarily comprises the software and control algorithms used for monitoring, decision-making, and supervisory control. In this work, LabVIEW is employed for real-time data acquisition, signal conditioning, visualization, system identification, control implementation, and alarm monitoring. Measured process variables are continuously communicated to the control layer, where safety-intelligent control actions are computed and transmitted back to the physical system for real-time operation.

SAFETY-INTELLIGENT CYBER-PHYSICAL CONTROL ARCHITECTURE

For designing the safety-intelligent control strategies, this work adapts the PAROC framework [11] for both modeling and simulation and experimention. Figure 3 presents the framework on implementation of control strategies in cyber-physical system.

First, the system is modeled to capture its dynamic behavior under open-loop simulations. These simulations enable the a *priori* identification of potential safety events, such as thermal runaway, overpressurization, or fuel starvation, before conducting any physical experiments. The high-fidelity model is developed in Siemens Process Academics (Research). Subsequently, a cyber-physical system is developed by integrating both the experimental and cyber setups. The experimental component is designed based on a process and instrumentation diagram (P&ID) or a process flow diagram (PFD). The cyber component consists of a human–machine interface implemented using a customized LabVIEW platform, data acquisition for communication and data exchange, and a microcontroller for real-time control and operation.

Using the outputs and inputs of the open-loop data, a linear or surrogate model is derived using techniques such as system identification, model reduction, machine learning, or Jacobian-based linearization [12, 13, 14]. This linearized model then serves as the basis for designing safety-intelligent explicit model predictive control (MPC) strategies.

The formulations of deterministic and stochastic risk-based explicit MPC (R-eMPC) are presented in Eqs. (11)-(18) and Eqs. (19)-(26), respectively. The objective function aims to minimize the deviations in system outputs (*y*) from their corresponding setpoints, while simultaneously penalizing control efforts of the inputs (*u*). Eqs. (11) and (19) defines the weights for outputs and inputs matrices, *QR* and *R*, and output horizon (*OH*) and control horizon (*CH*), where *k* denotes the discrete timestep in seconds. The state-spaced models are given in Eqs. (12)-(13) and (20)-(21), where *x* represents state vector and *A, B, C,* and *D* are the system matrices. For the charging operation of the metal hydride, the outputs (*y*) correspond to the core temperature and hydrogen adsorbed capacity while inputs (*u*) are supply pressure and water flowrate.

Within the deterministic R-eMPC, dynamic risk assessment is conducted to compute a risk index, *RI*, which is defined as the product of event probability and severity as shown in Eq. (15). The resulting *RI* can be expressed as a function of safety-critical variable, enabling the controller to evaluate the safety level of the system during operation. Based on this assessment, R-eMPC mitigates risk propagation speed and severity, or raise the alarm when runaway reactions cannot be prevented during the operation.

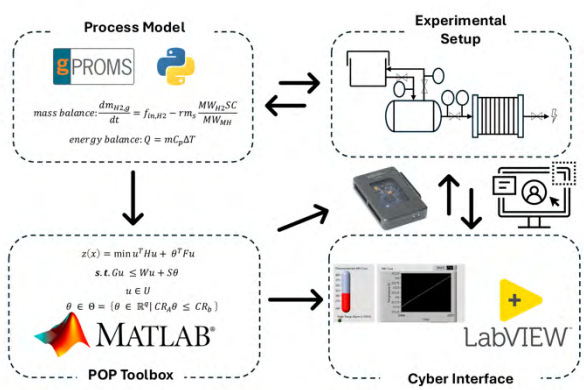

Figure 3. Safety-intelligent framework for cyber-physical systems.

$$min. J = \sum_{k=1}^{OH-1} QR\left(y_k - y_k^R\right)^2 + \sum_{k=1}^{CH-1} R\left(u_k - u_k^R\right)^2 \tag{11}$$

$$s.t. x_{k+1} = A_{xk} + B_{uk} \tag{12}$$

$$y_k = C_{xk} + D_{uk} \tag{13}$$

$$y_{min} \leq y \leq y_{max} \tag{14}$$

$$RI = P(y_{0,k}) \times S(y_{0,k}) \tag{15}$$

$$u_{min} \leq u \leq u_{max} \tag{16}$$

$$\Delta u_{min} \leq \Delta u \leq \Delta u_{max} \quad (17)$$

$$x_{min} \leq x \leq x_{max} \quad (18)$$

$$RI_{min} \leq RI(k) \leq RI_{max} \quad (19)$$

To account for the uncertainties during the operation of a system, stochastic R-eMPC (sR-eMPC) extends the deterministic R-eMPC through incorporation of risk tolerance (α) and chance-constraints within the optimization and decision-making framework. The design of sR-eMPC is based on Eqs. (20)-(27). The risk tolerance (α) can be calculated based on Bayesian inference in Eq. (22) to capture uncertainties of core temperature variations in metal hydride. As incorporated in Eq. (24), the chance-constraint update dynamic probabilistic limits to allow the system to operate beyond strict limits, thereby preventing unnecessary interventions or shutdown of the system. Further details on designing control optimization of the risk-based controllers can be found in Ali et al. and Akundi et al. [4, 5].

$$min. J = \sum_{k=1}^{OH-1} [QR(y_k - y_k^R)^2 + W(\alpha_k - \alpha_k^R)^2] + \sum_{k=1}^{CH-1} R(u_k - u_k^R)^2$$

$$s.t. x_{k+1} = A_{xk} + B_{uk} \quad (20)$$

$$y_k = C_{xk} + D_{uk} \quad (21)$$

$$P[Event|[y_k]] \leq \alpha_k \quad (22)$$

$$y_{min} \leq y_k \leq y_{max} \quad (23)$$

$$y_{0,k} - y_{0,threshold} \leq F^{-1}(\alpha_k) \quad (24)$$

$$u_{min} \leq u_k \leq u_{max} \quad (25)$$

$$\Delta u_{min} \leq \Delta u_k \leq \Delta u_{max} \quad (26)$$

$$x_{min} \leq x_k \leq x_{max} \quad (27)$$

The final step in this framework is to solve the risk-based control optimization via multi-parametric programming using POP Toolbox in MATLAB® [15]. This is an emerging technique to solve MPC explicitly, which means the optimization problem is solved once instead of iterating it every sampling time [16, 17]. Finally, during closed-loop implementation, the control actions is evaluated by computing the parameter vector, which includes hidden states, inputs, outputs, the risk index or risk tolerance, and setpoints. This parameter vector is then used to identify the feasible critical region, from which the appropriate control action is directly applied to the model or system. With the explicit R-eMPC solutions in place, the closed-loop performance of the system can be analyzed to demonstrate the effectiveness of both deterministic and stochastic risk-based control strategies.

RESULTS AND DISCUSSION

Open-loop Analysis

The dynamic behavior of the metal hydride storage system (MHSS) during charging is analyzed under open-loop conditions, without any control feedback. The charging or adsorption process involves supplying hydrogen gas to the metal hydride reactor, where H_2 is adsorbed within the lattices of the metal alloy, resulting in an exothermic reaction. A comparison between simulated and dynamically optimized responses of the core temperature and hydrogen adsorption capacity is presented.

For the open-loop simulation, the input variables, supply pressure and water flowrate, are set to prescribed values. In contrast, dynamic optimization aims to maximize the hydrogen adsorbed capacity while satisfying the operational constraints specified by the US DOE [3] within 300 seconds charging window.

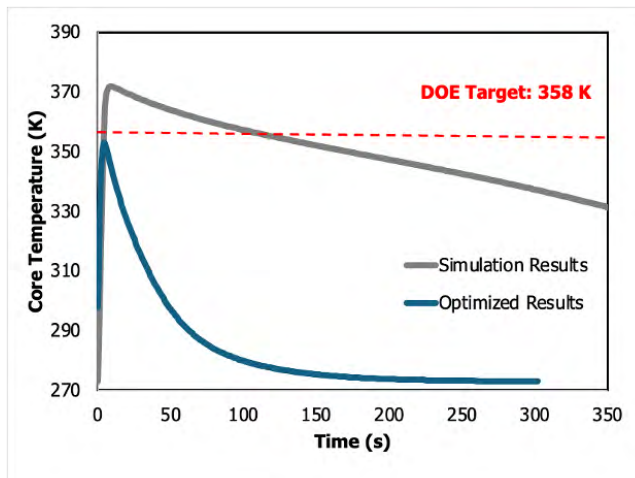

Figure 4. Core temperature trajectory during charging of metal hydride storage system during simulation and dynamic optimization.

Figure 4 compares the open-loop responses of core temperature obtained from simulation and optimization. In the simulated case, the supply pressure and water flowrate are set to 14 bar and 5 g/s, respectively. Under these conditions, the core temperature increases and reaches a peak of 378 K within 20 s, exceeding the recommended operating temperature limit of 358 K [3]. By comparison, the optimized response of core temperature peaks at 354 K, below the safety threshold, and achieves rapid cool down of the metal hydride. The optimal input profile determined by the optimization problem corresponds to a supply pressure of 14 bar and 12.46 g/s of water flowrate.

The corresponding hydrogen adsorbed capacity during open-loop simulation and dynamic optimization is shown in Figure 5. At a supply pressure of 14 bar and a water flow rate of 5 g/s, the simulated hydrogen storage

capacity reaches only 0.9 wt.% within 300 s, indicating that the metal hydride does not reach saturation. In this case, the core temperature exceeds the thermal limit while the maximum storage capacity remains unattained. In contrast, the open-loop optimized response achieves a higher hydrogen storage capacity of 1.34 wt.% within the same time horizon. Therefore, these behaviors highlight a critical trade-off between faster charging rate and thermal safety, emphasizing the need for effective thermal management.

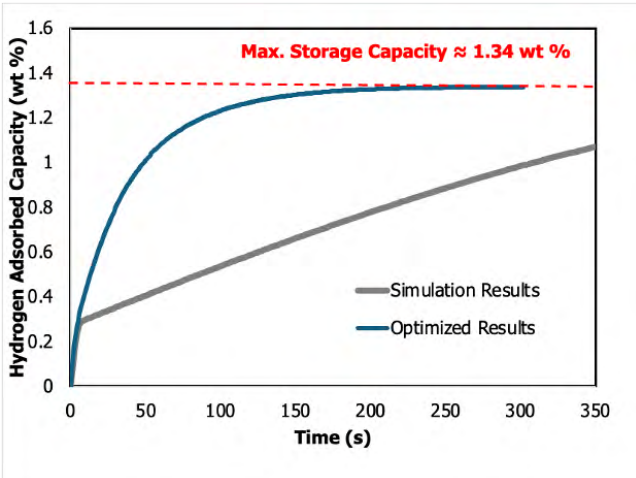

Figure 5. Hydrogen stored capacity trajectory during charging of metal hydride storage system during simulation and dynamic optimization.

The open-loop results demonstrate that while supply pressure improves charging efficiency and storage capacity, it also exacerbates thermal risks. These observations motivate the development of risk-based explicit model predictive control strategies to balance performance objectives with safety constraints under dynamic operating conditions.

Design of Risk-based Control Optimization

Deterministic and stochastic risk-based explicit control strategies were formulated based on the safety-intelligent control and optimization framework. The deterministic risk-based explicit MPC (R-eMPC) is designed according to Eqs. (11)-(18), where the core temperature is identified as the safety-critical variable due to the risk of thermal runaway during the charging process of the metal hydride storage system (MHSS). Figure 6 illustrates the dynamic risk analysis assuming a normal distribution of the core temperature, with a calculated mean (μ) of 290 K, a standard deviation (σ) of 11.35 K, and an upper control limit (μ + 3σ) of 324.05 K. Based on the analysis, the risk index limit is *RI* = 7.36, corresponding to the temperature limit = 358 K. Since the dynamic risk index is a temperature-dependent function that grows exponentially, a piecewise linear approximation is employed to obtain linear expressions suitable for embedding within the decision-making optimization problem.

Meanwhile, the stochastic R-eMPC (sR-eMPC), formulated in Eqs. (19)-(27), extends the deterministic framework by explicitly accounting for uncertainty through risk tolerance (α) based on Bayesian inference. In MHSS, the variation of core temperature is considered to analyze the risk tolerance during dynamic operation. In this formulation, α = 1 represents a fully safe operating condition, while α = 0 indicates a high-risk state. Compared to deterministic R-eMPC, the stochastic formulation provides additional operational flexibility by allowing controlled and temporary exceedance of temperature limits through chance-constraints. The chance-constraints incorporate the dynamically updated risk tolerance and are computed using inverse cumulative distribution function, as shown in Eq. (24), with a threshold temperature limit ($y_{0,\ threshold}$) of 358 K. Therefore, this allows the controller to satisfy the safety limits probabilistically, ensuring that constraints are met with a specified confidence level rather than deterministically.

Overall, the objective function of both safety-intelligent control strategies is to minimize excessive temperature peaks while maximizing hydrogen storage capacity within a finite time horizon during MHSS operation. Once the risk-based explicit MPC problem is solved using multi-parametric programming, explicit solutions to the optimization problem are obtained [15, 17]. These solutions are then implemented for closed-loop analysis of the metal hydride storage system.

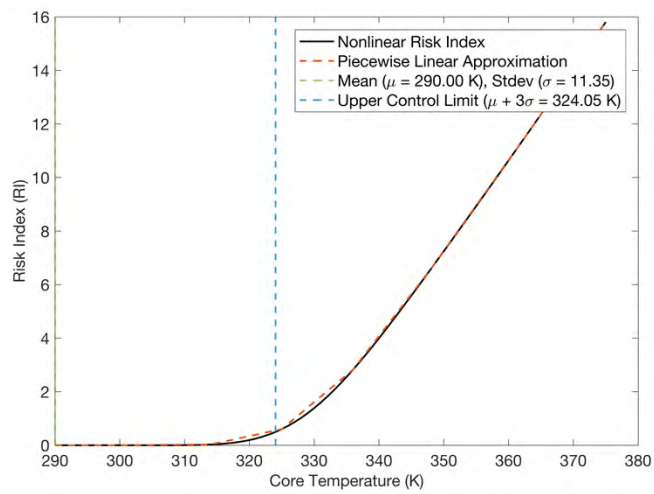

Figure 6. Dynamic risk analysis of core temperature during charging of metal hydride storage system.

Closed-loop Analysis

The resulting closed-loop responses are validated through simulation (*in-silico*) first using gO:MATLAB as the communication protocol between the system model in gPROMS and R-eMPC in MATLAB. The following subsection presents the analysis of both deterministic and

stochastic R-eMPC strategies during charging operation of the metal hydride storage system.

Figure 7 presents the closed-loop responses of the metal hydride system under deterministic risk-based explicit MPC (R-eMPC) in *in-silico* simulations. The core temperature increases gradually and reaches a peak value of 372 K, exceeding the setpoint limit of 358 K. Simultaneously, the hydrogen storage capacity attains the maximum of 1.3 wt.% within 300 s. The risk index exhibits a trend similar to that of the core temperature and exceeds the specified risk limit when the temperature peaks. Although thermal runaway is not completely prevented, the controller effectively attenuates both the risk propagation rate and consequence severity.

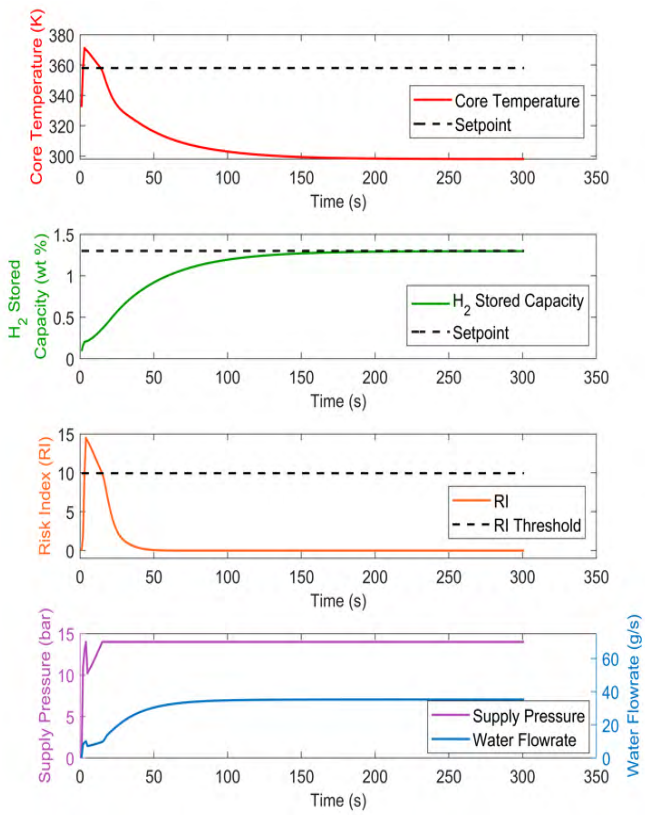

Figure 7. *In-silico* application of deterministic risk-based explicit MPC for closed-loop analysis of core temperature, hydrogen stored capacity, risk index and control actions.

At *t = 2 s,* the R-eMPC forecasts the risk index will exceed the limit in the next 10 seconds, which enable to trigger the alarm system ahead of fault occurrence time. While the risk index reaches a high value of *RI* = 14.5 at *t = 5 s,* the controller continues to act and successfully reduces both the risk and fault severity within 15 seconds. Here, the control actions are adjusted accordingly to manage the trade-offs between safety and performance. As a result, the R-eMPC effectively mitigates thermal risk while maximizing hydrogen storage capacity.

Figure 8 shows the in-silico performance of stochastic risk-based explicit MPC (sR-eMPC) during charging of metal hydride. The core temperature reaches a peak value of 367 K. Although this value exceeds the nominal temperature limit of 358 K, it remains below the maximum probabilistic limit of 371 K, demonstrating the increased operational flexibility of the stochastic controller and avoiding the need for a system shutdown.

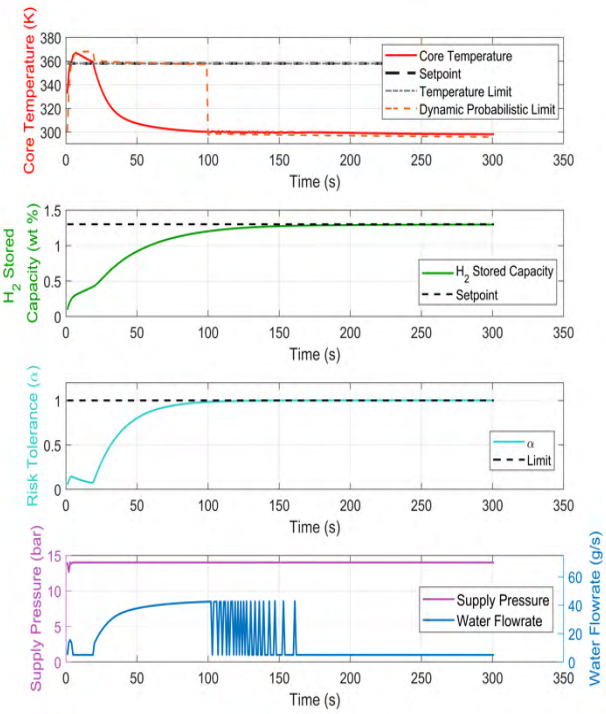

Figure 8. *In-silico* application of stochastic risk-based explicit MPC for closed-loop analysis of core temperature, hydrogen stored capacity, risk tolerance and control actions.

Unlike the risk index, the risk tolerance (α) evolves differently from the temperature profile. When it peaked at maximum, α = 0.1 indicating system is at state of risk. As adsorption process completes, the temperature gradually cools down and risk tolerance stabilizes to α =1. When the system operates in safe mode, the controller switches to predefined minimum control actions, corresponding to a supply pressure of 14 bar and a water flow rate of 5 g/s. At *t* = 170 s, the control system engages these predefined actions once the system is classified as safe according to the updated dynamic risk tolerance.

Moreover, the optimal control actions during closed-loop operation are provided by the controller to manage safety and performance. To achieve maximum storage capacity, the supply pressure maintains to be 14 bar. Meanwhile, as the temperature rises, the water flowrate increases to dissipate heat and maintain safe

operation. As a result, sR-eMPC is able to allow the system to exceed the temperature limit and achieving full saturation within 300 s while satisfying the dynamic probabilistic limits.

CONCLUSION AND FUTURE WORK

This work presented dynamic modeling, experimental setup, control and operations of hydrogen storage using metal hydride in real-time. The safety-intelligent controllers, deterministic and stochastic R-eMPC, utilize risk indicators such as risk index and tolerance to assess safety level of the system and account for uncertainties, respectively. These were implemented to minimize the evolving thermal risks while maximizing the hydrogen capacity of metal hydride tank. Future work will include implementation of the stochastic risk-based explicit MPC into the cyber-physical setup of metal hydride storage for closed-loop analysis and validation.

ACKNOWLEDGEMENTS

The authors acknowledge financial support from NSF RETRO Project CBET- 2312457, Texas A&M Energy Institute, Mary Kay O'Connor Process Safety Centre (MKOPSC), and Department of Chemical and Biomedical Engineering at West Virginia University.

AUTHOR IDENTIFIERS

Author ORCIDs:
Tizon ZA: 0009-0006-9728-8224
Liu Y: 0000-0002-5895-8428
Akundi SS: 0000-0001-7711-7264
Braniff A: 0009-0006-6835-320X
Dantas B: 0009-0001-0446-7906
Tian Y: 0000-0001-7937-5785
Khan F: 0000-0002-5638-4299
Pistikopoulos EN: 0000-0001-6220-818X

REFERENCES

1. Saada R, Patel D, Saha B. Causes and consequences of thermal runaway incidents—will they ever be avoided?. Process Safety and Environmental Protection 97:109-115 (2015). https://doi.org/10.1016/j.psep.2015.02.005
2. Bellosta von Colbe J, Ares JR, Barale J, Baricco M, Buckley C, Capurso G, Gallandat N, Grant DM, Guzik MN, Jacob I, Jensen EH, Jensen T, Jepsen J, Klassen T, Lototskyy MV, Manickam K, Montone A, Puszkiel J, Sartori S, Sheppard DA, Stuart A, Walker G, Webb CJ, Yang H, Yartys V, Züttel A, Dornheim M. Application of hydrides in hydrogen storage and compression: achievements, outlook and perspectives. International Journal of Hydrogen Energy 44:7780-7808 (2019). https://doi.org/10.1016/j.ijhydene.2019.01.104
3. U.S. Department of Energy. Hydrogen Storage | Department of Energy n.d. https://www.energy.gov/eere/fuelcells/hydrogen-storage (accessed July 19, 2025).
4. Ali M, Cai X, Khan FI, Pistikopoulos EN, Tian Y. Dynamic risk-based process design and operational optimization via multi-parametric programming. Digital Chemical Engineering 7:100096 (2023). https://doi.org/10.1016/j.dche.2023.100096
5. Akundi SS, Liu Y, Braniff A, Dantas B, Shayan Niknezhad S, Tian Y, Khan F, Pistikopoulos EN. Towards intelligent safety-aware control systems – a stochastic risk-based explicit model predictive control strategy. Chemical Engineering Science 322:123052 (2026). https://doi.org/10.1016/j.ces.2025.123052
6. Liu Y, Akundi SS, Braniff A, Dantas B, Niknezhad SS, Tian Y, Khan F, Pistikopoulos EN. Process monitoring and safety?informed control of a proton exchange membrane water electrolysis system. AIChE Journal 71: (2025). https://doi.org/10.1002/aic.18909
7. Liu Y, Akundi SS, Tizon ZA, Braniff A, Dantas B, Tian Y, Niknezhad SS, Khan F, Pistikopoulos EN. Next?generation safety?aware cyber?physical system integration: a proton exchange membrane water electrolysis demonstration. AIChE Journal 72: (2026). https://doi.org/10.1002/aic.70211
8. Krokos CA, Nikolic D, Kikkinides ES, Georgiadis MC, Stubos AK. Modeling and optimization of multi-tubular metal hydride beds for efficient hydrogen storage. International Journal of Hydrogen Energy 34:9128-9140 (2009). https://doi.org/10.1016/j.ijhydene.2009.09.021
9. Talagañis BA, Meyer GO, Aguirre PA. Modeling and simulation of absorption–desorption cyclic processes for hydrogen storage-compression using metal hydrides. International Journal of Hydrogen Energy 36:13621-13631 (2011). https://doi.org/10.1016/j.ijhydene.2011.07.139
10. Aruna R, Jaya Christa ST. Modeling, system identification and design of fuzzy PID controller for discharge dynamics of metal hydride hydrogen storage bed. International Journal of Hydrogen Energy 45:4703-4719 (2020). https://doi.org/10.1016/j.ijhydene.2019.11.238
11. Pistikopoulos EN, Diangelakis NA, Oberdieck R, Papathanasiou MM, Nascu I, Sun M. Paroc—an integrated framework and software platform for the optimisation and advanced model-based control of process systems. Chemical Engineering

Science 136:115-138 (2015). https://doi.org/10.1016/j.ces.2015.02.030
12. Kravaris C, Kookos IK. Understanding process dynamics and control. Cambridge University Press (2021). https://doi.org/10.1017/9781139565080
13. Wang W, Wang Y, Tian Y, Wu Z. Explicit machine learning-based model predictive control of nonlinear processes via multi-parametric programming. Computers & Chemical Engineering 186:108689 (2024). https://doi.org/10.1016/j.compchemeng.2024.108689
14. Dantas B, Braniff A, Liu Y, Akundi SS, Tizon ZA, Niknezhad SS, Khan F, Pistikopoulos EN, Lima FV, Tian Y. Hybrid mechanistic/data-driven modeling for enhancing dynamic operations in electrolysis-based hydrogen production. International Journal of Hydrogen Energy 207:153544 (2026). https://doi.org/10.1016/j.ijhydene.2026.153544
15. Oberdieck R, Diangelakis NA, Papathanasiou MM, Nascu I, Pistikopoulos EN. POP - Parametric Optimization Toolbox. Ind Eng Chem Res 2016;55:8979–91. https://doi.org/10.1021/ACS.IECR.6B01913/ASSET/IMAGES/LARGE/IE-2016-01913A_0011.JPEG
16. Tøndel P, Johansen TA, Bemporad A. An algorithm for multi-parametric quadratic programming and explicit MPC solutions. Automatica 39:489-497 (2003). https://doi.org/10.1016/s0005-1098(02)00250-9
17. Pistikopoulos EN. Multi?parametric optimization and control. Wiley (2020). https://doi.org/10.1002/9781119265245

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Extremum seeking control by perturb and observe applied to dividing wall column pilot

Ivar J. Halvorsen*[ab], Bart M. A. Bergers[b], Giovanni Merlo[b], Leontine I.M. Aarnoudse[b], Mark A.M. Haring[a], Sigurd Skogestad[b]

[a] SINTEF, Trondheim, Norway
[b] Norwegian University of Science and Technology (NTNU), Trondheim, Norway
* Corresponding Author: Ivar.J.Halvorsen@sintef.no

ABSTRACT

The Dividing Wall Column (DWC) offers significant potential in saving both energy- and capital cost compared to conventional distillation sequences. However, there are some issues regarding flexibility and control that require attention in reducing the risks or uncertainties in achieving the potential benefits in practical operation. This calls for control and optimization methods that rely on the available measurement data and less on simulation models. The “Perturb and Observe” method is a simple algorithm that seems suitable for this on-line optimisation task. A series of experiments have been carried out at the Kaibel-column pilot at NTNU and some key results are presented. The method is combined with a conventional control structure at the regulatory layer.

Keywords: Distillation, Optimization, Process Control, Dividing Wall Column, On-line, Perturb & Observe

INTRODUCTION

The Dividing Wall Column (DWC) is an attractive multicomponent distillation technology that enables energy savings in the ballpark of 20-30% and even similar numbers for capital cost due to more compact size and construction [1, 2]. Still, robust practical control is often listed as a challenge in practical operation. It is therefore needed to show how industrial DWCs can be controlled in realistic environments.

The plant manager needs to be assured that the operation and control can be realized within the organization without extraordinary effort such that not only the plant owner, but also operating personnel such as process engineers and operators can benefit from this kind of energy-saving technology. The operating condition that enables high purity recovery of all products and at the same time minimizes the overall energy consumption is well described theoretically, but in practice it is not so easy. This is because the inevitable real-world’s uncertainties make it difficult to infer the actual operating condition and bring the operation back to and maintain the optimal energy-saving track with high purity products. In a lab-pilot, as well as in industry, there are significant uncertainties in separation performance, heat loss, internal hydraulics and mixing, and in the accuracy of setting the manipulative variables and in the measurements. This implies that it is impossible in practice to pre-calculate the settings that provide the desired operating condition. Some kind of feedback control strategy is required for both the basic control levels and the optimising control levels.

This uncertainty and variation in conditions calls for control and optimization methods that rely on the available measurement data and less on simulation models. The “Perturb and Observe” method is an available algorithm which is simple to realise on-line, and seems suitable for this on-line optimisation task [3]. During a previous simulation study [4], the reboiler vapor rate, that is closely related to column energy consumption, was used as the objective function to be minimized, while simple control loops maintained compositions. In industrial settings, the boilup rate is usually set constant at the capacity limit in the column and a different approach for defining the cost function is needed. In this paper, we will present how the on-line perturb and observe method can be combined with a basic control structure that handles the typical industrial operation condition, and we will show some samples from pilot-column experiments.

DWC OPERATION CHARACTERISTICS

Consider a typical ternary DWC as illustrated in Figure 1. The feed (F) enters the prefractionator sections (1, 2) and is described by its molar fractions (z) of pure components (A, B, C) and its liquid fraction (q). There are three product specifications, given by the main component purities in each product stream (D, S, B). There are five manipulated variables (MVs): 1) boilup vapor rate (V), 2) reflux rate (L), 3) side product draw rate (S), 4) the liquid split ratio (R_l) above the wall and 5) the vapor split ratio (R_v) below the dividing wall. The prefractionator shall separate the light key (A) from the heavy key (C) while the intermediate key (B) is distributed to both ends of the "dividing" wall. This leaves separation of A from B in the upper part of the main column (sections 3, 4) and B from C in lower part (sections 5, 6).

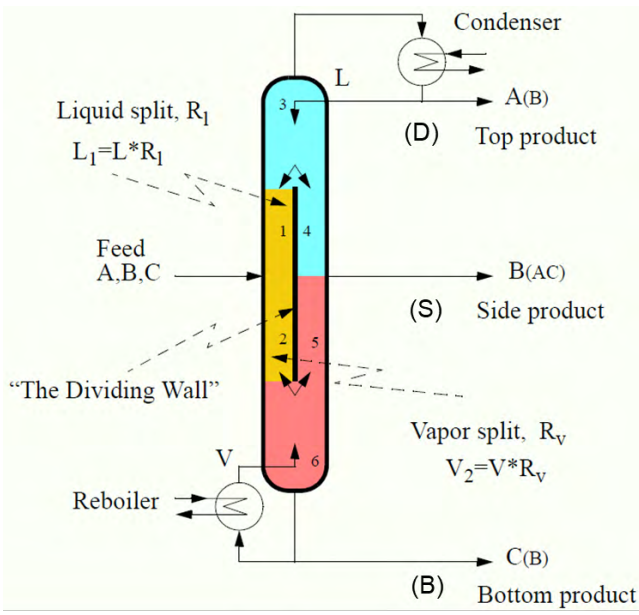

Figure 1. Ternay Dividing Wall Column (DWC)

This sounds straightforward, but to obtain the potential minimum vapor operation in practice, it is vital that the prefractionator is operated precisely within its "optimality region". The optimal operating characteristics are described in previous publications [5–7]. Here, just a summary is given as preparation for the optimising control experiments. The Vmin-diagram in Figure 2 is based on the feed properties only and shows how the feed components are distributed to the product ends in an infinite stage binary column as a function of the operational degrees of freedom, selected as vapor rate (V/F) and net top product (D/V=(V-L)/F). For a ternary feed, the left peak represents minimum vapor for a sharp A/BC split, and the right correspond to the AB/C split. In the "valley" between these peaks, we fine the so-called "Preferred split" which is minimum vapor for a sharp A/C split while B is distributed to both ends. The solid lines represent distribution boundaries, that is when a component is about to appear or disappear in one column end as V and/or D is changed such that the operating point crosses the actual line as illustrated with the small columns in **Figure 2**.

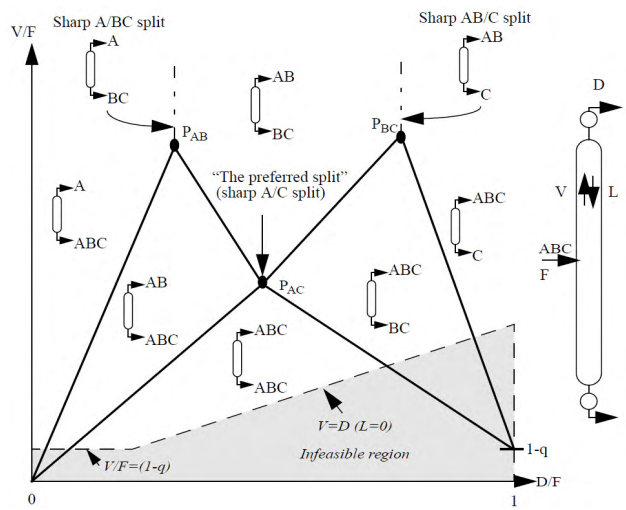

Figure 2. Vmin diagram for a ternary feed ABC [6]

With the assumption of ideal mixtures, the information in this diagram can be calculated analytically by the Underwood equations. It is important to recognise that these distribution boundaries are also there for real zeotropic mixtures, which requires some more care when doing the calculations, and which does not have these nice and linear boundaries. Thus, the concept of the distribution regions will still be valid. This holds also when assuming finite number of stages.

Minimum vapor rate in a DWC

In a DWC, the minimum overhead vapor rate is given by the highest of these two peaks. However, there is a very important requirement to the prefractionator operation: it must be operated in the vicinity of the preferred split when we apply the Vmin-diagram for the prefractionator itself. This is what we denote as the "optimality region". If this requirement is not fulfilled, it will not be feasible to produce all three products at specified purities with the given minimum vapor rate, despite otherwise perfect column sizing for the given feed conditions.

In the previous P&O study [4], the boilup itself has been used as the cost function value to be minimised in search for the optimality region. Minimum boilup is still very important in design calculations, because this is what gives the operating conditions that are the basis for column design. It is also straightforward in simulations without any limits on the available internal vapor and liquid rates, since sub-optimal settings of the split rates will force increased overall reboiler vapor rate to enable pure products.

In [4], the control structure applied the reflux (L), side-draw (S) and boilup (V_b) for product purity control, leaving the liquid and vapor split for optimisation. Since most industrial columns are operated with constant vapor split, this was also applied such that we only consider the one-dimensional search for the minimum cost (V_b) by perturbing the liquid split (R_l). Feed rate and conditions were regarded as disturbances.

Proposed control structure for the lab column

In practical operation, the vapor capacity (V_b)will be limited, and to ensure as high throughput of feed rate as possible, the reboiler duty is usually set close to maximum capacity. Then, the consequence of sub-optimal splits will *not* be increased vapor rates, but instead it will become infeasible to produce pure products. Thus, the degree of sub-optimality should be measured by the actual product impurities. This could also be "converted" to energy or vapor rate by calculating the amount of extra vapor rate needed to take the actual impure products and purify these in a theoretical DWC distillation column. But minimising directly on the impurity will give the same result.

A basic control loop that uses the top reflux as MV for maintaining a certain temperature location in the top section (3) ensures that the top product becomes reasonably pure even for sub-optimal prefractionator operation. This is equivalent to typical top product purity control in conventional columns. In the bottom, however, the boilup rate is not an available MV, so we may choose to use the side draw rate (S) instead to affect the bottom product purity via the liquid rate towards the bottom.

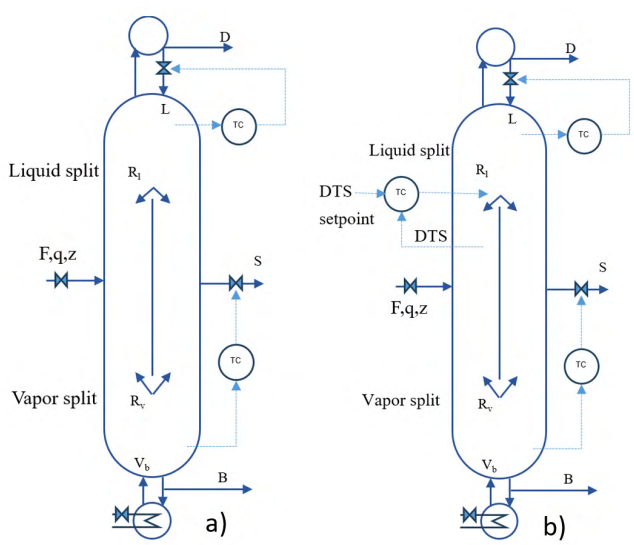

Figure 3. a) Basic control structure ensures pure D&B, while side product (S) depends on prefractionator operation. b) Added stabilising feedback for prefractionator by a combination for temperatures (DTS).

This control structure has been analysed by Qian et. al. [8] and show satisfying dynamic performance. The two control loops in the top and bottom are illustrated in Figure 3a. The side impurity is not handled by this basic structure at all. Luckily, this will be taken care of in the prefractionator control since, when it reaches its optimality region such that it performs the desired AB/BC split (no heavy C above the wall and no light A below the wall), the basic 2-loop structure in Figure 3a) will automatically recover all the light (A) in the top product, and all the heavy (C) in the bottom product, and the intermediate (B) has no other exit than through the side stream (S).

The resulting optimisation task becomes slightly different form the previous study [4] when the boilup is fixed. The cost function may be the side-product impurity or some other measure that is closely related to this.

Since we fix the boilup (V_b) it is important to set a realistic feed. The minimum vapor to feed ratio is related to the highest peak in the Vmin-diagram, here denoted $(V/F)_{min}$. The real ratio is denoted by just (V/F). This gives the following simple relation for the maximum feed capacity that still enables three pure products:

$$F_{max} = V_{b,max} \Big/ \left(V/F\right) \leq V_{b,max} \Big/ \left(V/F\right)_{min}$$

This might look obvious, but it is very practical. If the feed is too high, the side stream cannot become pure since the (V/F) in the Vmin-diagram will be below the highest peak and full recovery of all products becomes infeasible. If the feed is lower, this implies that the real $\left(V/F\right) > \left(V/F\right)_{min}$ and we have more reboiler duty than required for pure products (assuming that optimal prefractionator operation is reached). Thus, it is an option to either increase the feed, or in fact reduce the reboiler duty. If not, some of the potential savings are not realised. But there is also some trade-off by running with somewhat higher V/F ratio than required since the optimality region is also widened to a certain extent, allowing for simpler operation at the cost of abandoning some of the potential savings [9].

We should also consider how the real vapor split is set according to the requirement given from the Vmin-diagram. If too low, it can be partly compensated for by reducing the feed (implicitly increasing (V/F))

About cost function selection

It has been shown that an alternative cost function that characterises close to optimal operation is the temperature difference between the lower and upper end of the dividing wall [5], where a maximal difference corresponds to optimal operation. This cost function is usually easily available in practice and is also an intensive measurement.

Stabilising control of the prefractionator

It is possible to use the liquid split (R_l) directly as the MV for optimisation. However, the dynamic response on direct manipulation of internal flow rates in distillation is known to be slow. Then, a better alternative is to use a feedback controller to stabilise the profile and use the setpoint of this controller as the optimiser MV. This is illustrated in Figure 3b). The feedback may be a single point on the prefractionator temperature profile or, as in Figure 3b), the "Difference Temperature Symmetry" measure (DTS) [4, 5], which is a measure of the temperature differences on both side of the wall. This measurement also has so-called "self-optimising" properties, since keeping this value constant maintains operation close to an optimum just by feedback control. The setpoint may need some adjustments from the optimising level since the optimal value depends on the feed properties.

PERTURB AND OBSERVE (P&O)

The model-free Perturb and Observe (P&O) optimising search method is described in more detail in the previous work [4] and in the three student internship reports [10–12] for the methods used in the lab. The simplest explanation is that P&O perturbs the MV (u) and then measures the change in the measured cost function y=f(u) and from that does a directed search towards a local minimum while obtaining additional information for each step. The method can be refined in how uncertainty can be modelled [13]. Use of model information for feedforward actions is also an option. P&O can also be multidimensional, but only a single dimension is tried out in the lab experiments. In the presented DWC experiments, the perturbing MV (u) is either the manual set value for the liquid split, or when the stabilising DTS controller is used, the P&O-MV is the setpoint to this controller, that uses the liquid split as the controller MV. (See Figure 3)

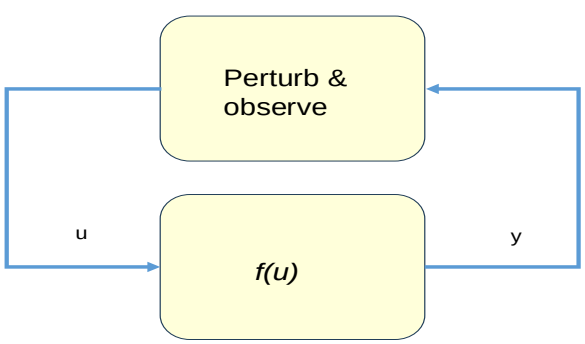

Figure 4. Conceptual P&O loop

LAB COLUMN PILOT

The lab pilot at NTNU (see Figure 5) was first built in 2005 by Strandberg. It has been used for various experimental studies by several PhD and master's students. It has the structure of a 4-product Kaibel-column. In this project, it is used as a ternary DWC by simply closing one side-stream outlet (S2). The feed is Ethanol, 1-propanol and 1-Butanol with boiling points at 78 °C, 100 °C and 117 °C, respectively. The feed can be preheated to a controlled temperature. Changes in feed rate, composition and temperature can be regarded as disturbances.

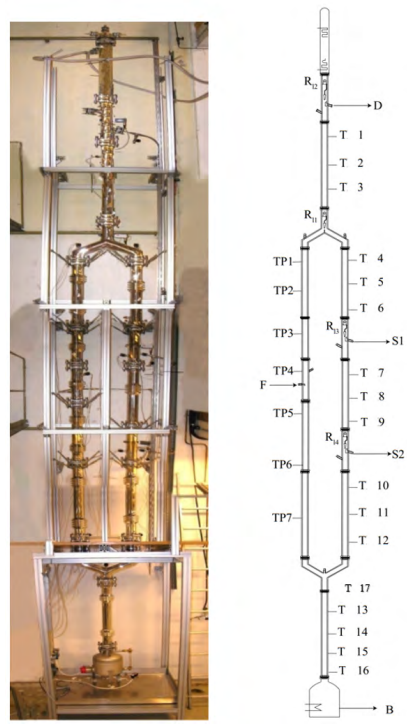

-Height ~12m
-Internal Diameter:
Bottom: 7 cm
Branches: 5 cm
Top: 5cm
-Atmospheric pressure at condenser
-Vacuum insulated glass sections produced by Normag Labortechnik.
-Electric heater 2.5 kW. Max about 2
-Feed rate 0-10 l/hour, normal 2-3.
-24 temp locations
-6mm Raschig glass ring plus some structured elements.
-Labview control
-Switching funnels for split and products products.
-Active vapor split

Figure 5. The Kaibel column Pilot at NTNU. The left branch is the prefractionator, and the rest makes up the main column. Only one side-draw is used for ternary feed.

A theoretical Vmin diagram is shown in Figure 6. This is just indicative as it will become distorted with real (very few) stages and uncertain packing performances with varying rates. The figure illustrates three possible operational paths (dashed) for three different vapor split values.

However, there is still a minimum side-stream impurity, even for sub-optimal settings for vapor split values so the one-dimensional algorithm will find a local optimum. For a compete DWC optimisation both vapor-split and boilup-to-feed ratio (V/F) need to be used as MVs in addition to the liquid split. This is not considered herein.

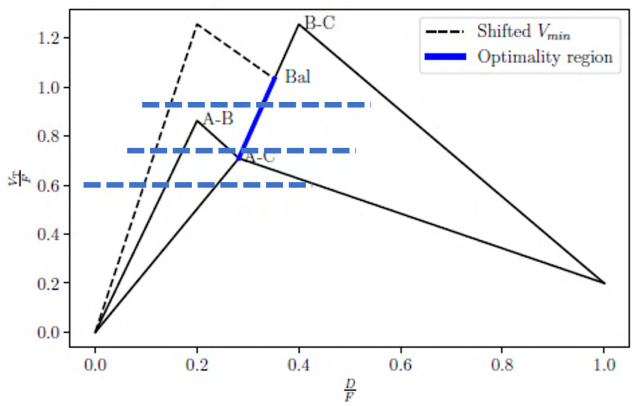

Figure 6. Ideal Vmin diagram for the actual feed with molar fractions [0.2, 0.2, 0.6] at room temperature. Changing liquid split moves the operating point horizontally (dashed blue lines). We do not really now the position of this line in practice, and we do not know the real diagram. Consider the real diagram is hidden in "fog".

EXPERIMENTAL RESULTS

The experiments are described in more detail in the internship reports by Kooy, Bergers and Merlo [10–12]. These reports are available on request. Just a few samples are included here.

The first example from Kooy [10] is an experiment based on the structure in Figure 3a) by using liquid split as MV (u in Figure 4) and an approximated temperature difference over the "wall" (or branches on the wall) as a cost function. See Figure 7, where the cost is given by $\text{Cost}=(TP1+T4)/2-T17$. (The desired cost is temperature difference between liquid split and vapor split locations. Unfortunately, there is no sensor at the liquid split location so the average (TP1+T4)/2 is used instead.) Note that the column settling time will be quite slow when manipulating directly on the liquid split. This gives a rather slow underlying mode during the experiment, resulting in a slow-moving cost function optimum value. Even then the P&O algorithm manage to track the optimum as shown in the sequence in Figure 7.

The second experiment, see Figure 8, is with the stabilising DTS loop as in Figure 3b). The DTS-loop results in a considerably faster settling time for each perturbation of the DTS controller setpoint value.

Bergers and Merlo worked together at several experiments in the autumn of 2025. Figure 9 shows an example of cost function repeatability experiment for the case when DTS control is applied. The setpoint is changed upwards and downwards to explore the cost function and confirm that the function has a minimum. The match is acceptable. However, when running a sequence of experiments where the products are recirculated to the feed there may be some variations in feed properties e.g. due to some material losses (leakages) etc. Thus, it is not so straightforward to ensure the same starting / operating conditions each time.

Bergers also looked into uncertainty based "uP&O" [13]. He proved that the perturb and observe with a closed loop manipulation of R_l tracks the optimum of the same $\text{Cost}=(TP1+T4)/2-T17$.

It is quite important that the underlying regulatory control loops for top and bottom products purity control is properly tuned. An experiment illustrating improved response after tuning is given in Figure 10.

They also performed some basic characterisations for model fitting and measurement calibration and Merlo also worked on fitting an Aspen Plus model in a rather detailed manner.

Merlo verified through Gas Chromatography a mismatch between side stream purity and the optimum of the cost and verified this mismatch via the Aspen Plus model.

The desired cost function, being the side impurity, could not be measured on-line. In some experiments an off-line Gas Chromatography analysis could be performed.

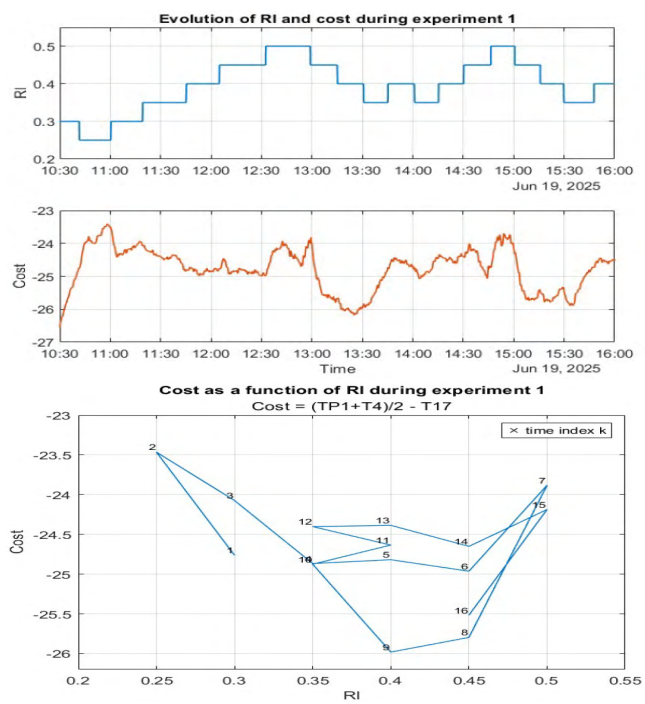

Figure 7. P&O experiment by manipulating liquid split (RI) [10]

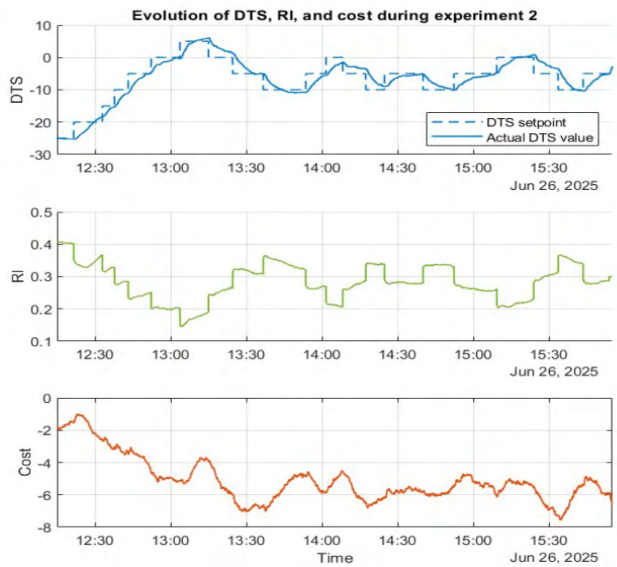

Figure 8. P&O experiments with simple P&O with DTS [10]

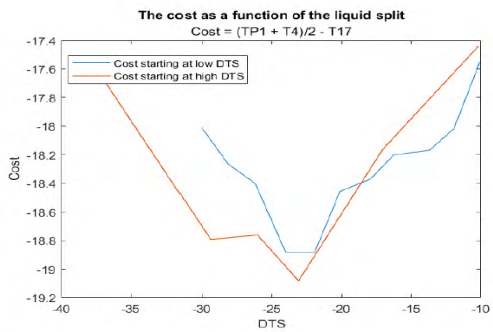

e 7.1: The relation between DTS and cost based on two experiments.

Figure 9. Experiments showing cost function vs DTS [11].

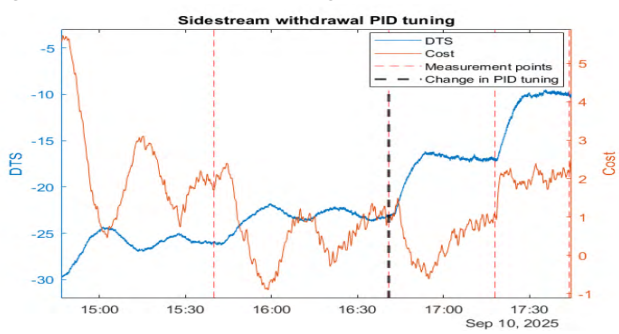

Figure 10. Improving basic regulatory PI-tuning [11]

The gas chromatography measurements provided a reference to evaluate the behaviour of the cost function compared to the real objective, as shown in Figure 11.

The minimum of the cost indicator is shifted to higher R_l values compared to the minimum of the impurity. The same phenomenon could be verified in the fitted Aspen Plus model. Figure 12 shows the result of simulations of the fitted model: for analogous operating conditions to those implemented experimentally, taking the average of the stage temperatures around the wall to approximate the measurement at the upper junction along the wall yields a slightly different optimal liquid split than the desired one.

As the available temperature inferential did not perfectly match the side impurity's optimality and online concentration measurements are not available, Merlo suggested to use the total side draw flow rate as a cost function. Since the top and bottom purities are regulated via temperature control, suboptimal steady-state operation will necessarily result in part of the top component A and the bottom component C being recovered together with the middle distillate B in the side stream. Optimal operation, on the other hand, signifies that only B is recovered via the side stream, thus making the side stream flow rate minimal. This is easier to measure than the composition. It is also possible to infer this amount directly from the S-controller outputs and the liquid split.

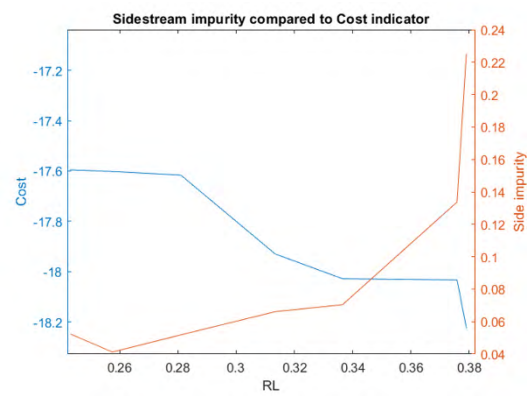

Figure 11. Side-stream impurity compared to Cost indicator [12]. Off-line analysis

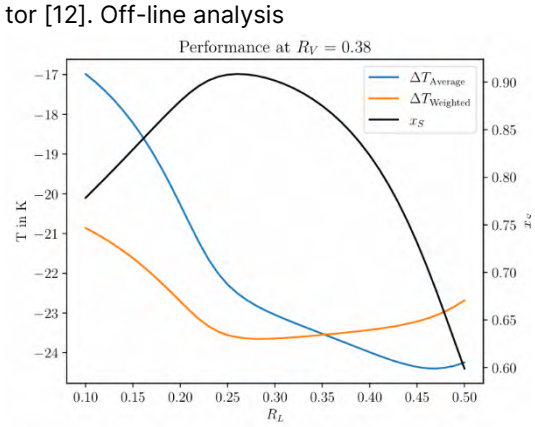

Figure 12. Comparing cost function alternatives by simulations with the fitted Aspen Plus model. [12]

DISCUSSION

The task of searching for the optimal operating condition for a DWC can be compared to navigating in mountain terrain in dense fog. We know that the mountains (Vmin-diagram) are there, but we do not know exactly how they look like (due to model uncertainty) and where our current position is just now (which prefractionator

operating point). However, we have some communication to the outside that can tell the impact of moving our position (the cost function measurement and other external measurements).

If we have an accurate GPS and map, we could easily move to the desired position. This could be for example equivalent of measuring the light and heavy impurities in the bottom and top of the prefractionator. This can be used tell how far the operating point is from the optimality region at the Vmin-diagram boundary line.

The experience from the lab indicates that it is difficult to develop a precise model and even a constant model due to significant uncertainties. First there are some heat losses, partly due to that some section may have lost the vacuum. This has been indicated by IR thermography. There may also be rate dependent variations in HETP and variations due to hydraulics maldistributions. All these uncertainties are likely to be more dominating on this pilot, than in an industrial setting.

The simplest P&O method is able to track the defined cost function despite quite uncertain and non-steady process conditions. The step size and the period between each perturbation step are parameters that need to be tuned.

Experiments were also conducted to investigate possible feed-forward actions. In the simplest form, e.g. if the feed temperature is changed, we might know from the model how the optimum is expected to change. Then, this information about an optimality-change direction can be helpful to move faster to the new optimum. More detailed model knowledge may allow also the expected magnitude of change of the optimisation MV.

In the presented experiments we have used a single perturbation variable, and thereby only considered one-dimensional optimisation search. In further work it is of interest to include impact from the vapor split and to maximise the feed throughput.

CONCLUSION

A series of experiments have been carried out on the NTNU Kaibel-column pilot during 2025. The goal was to try out the on-line model-free “Perturb and Observe” algorithm in an environment that is similar to an industrial plant. Unlike in simulations where we have quite good control of the external “upsets” when we study the algorithm performance, the lab also includes some of the real-world uncertainties, and sometimes unforeseen and unmodeled behaviour. Despite this, the results are promising.

In addition to collect experiences on how to improve the optimisation algorithms, the results also have resulted in new knowledge on DWC control structure and performance.

The use of P&O is a low-hanging fruit that provides a systematic optimum tracking 24/7 compared to if operators should follow up this manually, enabling realising the benefits of optimisation not only for design but in practical plant operation. It is easily combined with conventional basic control.

At the basic level, some attention should be given to obtaining reasonably fast settling time upon disturbances. This enables faster optimisation too.

And lastly, since the optimisation algorithm requires a measurable cost function, it is required to ensure that the available measured properties used for this function are well representative for the actual process operating cost that we really want to minimise.

Compared to full model-based real time optimisation (RTO), the P&O does not require any on-line model, and functionality for model updating from measurements. Thus, P&O and should thereby be much simpler to implement on an industrial plant. Some risk assessment on the impact of applying expected perturbations must of course be included as for all control upgrade projects.

DWC columns are built to save energy and using on-line optimisation methods like Perturb and Observe give the managers and operators better tools to realise the potential savings in practice.

ACKNOWLEDGEMENTS

Funded by The Research Council of Norway, Project No.: 345272: Data-driven optimization of industrial processes in time-varying environments. Thanks to visiting internship students and their institutions for carrying out work in the lab: Benedikt Kloss, TUM, (2023), Lars Kooy, TUE, Bart Bergers TUE, Giovanni Merlo, TUM, (2025).

AUTHOR IDENTIFIERS

Author ORCIDs:
Halvorsen IJ: 0000-0002-3215-1810
Haring MAM: 0000-0002-3340-6731
Aarnoudse LIM: 0000-0003-1143-6195
Merlo G: 0009-0002-9349-0055
Bergers BMA: 0009-0001-7770-4327
Skogestad S: 0000-0001-6187-8261

REFERENCES

1. Dejanovi? I, Matijaševi? L, Oluji? Ž. Dividing wall column—a breakthrough towards sustainable distilling. Chemical Engineering and Processing: Process Intensification 49:559-580 (2010). https://doi.org/10.1016/j.cep.2010.04.001
2. Asprion N, Kaibel G. Dividing wall columns: fundamentals and recent advances. Chemical Engineering and Processing: Process Intensification 49:139-146 (2010).

https://doi.org/10.1016/j.cep.2010.01.013
3. Femia N, Petrone G, Spagnuolo G, Vitelli M. Optimization of perturb and observe maximum power point tracking method. IEEE Trans. Power Electron. 20:963-973 (2005). https://doi.org/10.1109/tpel.2005.850975
4. Halvorsen IJ, Aarnoudse LIM, Haring MAM, Skogestad S. Extremum seeking control applied to operation of dividing wall column ? DWC. Systems and Control Transactions 4:1157-1162 (2025). https://doi.org/10.69997/sct.105824
5. Halvorsen IJ, Skogestad S. Optimal operation of petlyuk distillation: steady-state behavior. Journal of Process Control 9:407-424 (1999). https://doi.org/10.1016/s0959-1524(99)00009-8
6. Halvorsen IJ, Skogestad S. Minimum energy consumption in multicomponent distillation. 1. v_{min} diagram for a two-product column. Ind. Eng. Chem. Res. 42:596-604 (2003). https://doi.org/10.1021/ie010863g
7. Halvorsen IJ, Skogestad S. Minimum energy consumption in multicomponent distillation. 2. three-product petlyuk arrangements. Ind. Eng. Chem. Res. 42:605-615 (2003). https://doi.org/10.1021/ie0108649
8. Qian X, Jia S, Skogestad S, Yuan X. Comparison of stabilizing control structures for dividing wall columns. IFAC-PapersOnLine 49:729-734 (2016). https://doi.org/10.1016/j.ifacol.2016.07.271
9. Ränger LM, Halvorsen IJ, Grützner T, Skogestad S. What can go wrong in a dividing wall column and how to detect it. Separation and Purification Technology 354:129151 (2025). https://doi.org/10.1016/j.seppur.2024.129151
10. Kooy L (2025) Optimisation of a Solar Array and Dividing Wall Column Using Perturb-and-Observe methods. NTNU, TU/e, Trondheim, Norway
11. Bergers B (2025) Optimisation and Control of a Dividing Wall Distillation Column. NTNU, TU/e, Trondheim, Norway
12. Merlo G (2026) Modelling and Control of an experimental Dividing Wall Column. NTNU/TUM
13. Aarnoudse L, Haring M, van de Wouw N, Pavlov A. Uncertainty-based perturb and observe for fast optimization of unknown, time-varying processes [*]. 2025 IEEE 64th Conference on Decision and Control (CDC) :2268-2273 (2025). https://doi.org/10.1109/cdc57313.2025.11313008

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Managing Renewable Energy Uncertainty in Green Hydrogen Production Systems

Matteo Lea Casagrande, Andrea Isella, and Davide Manca*

PSE-Lab, Process Systems Engineering Laboratory, Dipartimento di Chimica, Materiali e Ingegneria Chimica "Giulio Natta", Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy
* Corresponding Author: davide.manca@polimi.it.

ABSTRACT

The extensive use of renewable energy to supply hydrogen production for chemical processes is hindered by the uncertainty in power generation and by strict operational limits.These challenges are addressed through a real-time dynamic optimization approach based on a receding-horizon strategy that provides optimal decision variables. The framework explicitly relies on imperfect weather forecasts and dynamically adapts the hydrogen reference production to guarantee the final productivity target. The optimization methodology focuses on minimizing grid electricity imports, limiting excessive equipment stress, and preventing constraint violations, while ensuring that the hydrogen production target is satisfied.The proposed approach yields competitive economic and environmental performance, with a levelized cost of hydrogen of 3.31 USD/kg_{H_2}, well within literature values (1.50–7.50 USD/kg_{H_2}) and below typical industrial costs (4–12 USD/kg_{H_2}). At the same time, carbon dioxide emissions are reduced by more than 90% compared to steam methane reforming, resulting in specific emissions below 1 kg_{CO_2}/kg_{H_2}. These results are supported by a year-long simulation based on CAISO renewable energy data. The optimized system achieves the annual hydrogen production target while respecting all operational constraints, whereas the absence of operational strategy and buffering systems leads to systematic constraint violations. The strategy can be implemented in existing industrial facilities with limited modifications.

Keywords: Green hydrogen, Real-time dynamic optimization, Forecast uncertainty, Receding horizon, Power-to-Hydrogen

INTRODUCTION

Global scenario

The urgency surrounding industrial decarbonization has intensified as CO_2 emissions from human activities continue to increase. Recent data from 2024 indicate that emissions from energy generation, fuel combustion, and manufacturing operations have reached 37.6 Gt_{CO_2} annually, an 8.4% increase relative to levels observed ten years earlier [1]. Direct emissions from industrial sources (Scope 1) represent approximately 23% of this global total, not accounting for the substantial indirect emissions associated with energy consumption across these operations.

Among industrial sectors, chemical manufacturing presents particularly high carbon intensity. Traditional synthesis pathways for several foundational chemicals result in considerable emission factors: ammonia production generates up to 4.6 kg_{CO_2}/kg_{NH_3}, methanol up to 2.8 kg_{CO_2}/kg_{CH_3OH}, nitric acid up to 1.6 kg_{CO_2}/kg_{HNO_3}, and hydrogen up to 26 kg_{CO_2}/kg_{H_2}. When accounting for both on-site emissions and those associated with purchased energy, the chemical industry is responsible for approximately 6.4% of global greenhouse gas releases [2].

Concurrently, the renewable energy sector has undergone substantial expansion. European markets illustrate this transition: the share of renewables in total energy consumption increased from 9.6% in 2004 to 24.5% by 2023 [3]. While this growth trajectory presents opportunities for reducing carbon emissions in chemical production, it simultaneously introduces operational challenges due to the intermittent nature of wind and solar power generation.

Renewable energy sources

Solar photovoltaic and onshore wind technologies exhibit favorable characteristics for industrial deployment due to their modularity and minimal site-specific requirements. Nevertheless, the temporal variability inherent to these sources necessitates effective energy conversion and storage infrastructure to ensure process continuity.

For renewable-powered hydrogen generation, proton exchange membrane (PEM) electrolysis systems demonstrate superior performance under dynamic operating conditions compared to alkaline technologies, despite the latter's higher technology readiness level and lower capital expenditure. Regarding energy storage architectures, battery energy storage systems (BESS) deliver high round-trip efficiencies and rapid power response characteristics, while compressed hydrogen storage represents the most established and cost-effective approach for hydrogen inventory management.

Plant design and operation in literature

Existing literature on renewable-hydrogen system design predominantly emphasizes techno-economic optimization under steady-state or quasi-static operating assumptions [4, 5]. Reported levelized costs for hydrogen production typically range from 3 to 4 USD/kg_{H_2}, with regional variations driven by resource availability: lower values observed in high-resource regions and elevated costs documented across European installations [6]. The operational performance of these systems under transient renewable supply conditions remains insufficiently characterized.

Current research on operational control strategies frequently relies on deterministic renewable energy forecasts or omits critical process-side constraints [7, 8]. As a result, methodologies capable of addressing renewable variability, equipment operational limits, downstream process requirements, and economic objectives within a unified framework remain underdeveloped.

This study presents a real-time, multi-objective optimization framework for integrated renewable-hydrogen production systems operating under uncertainty. The framework incorporates forecast error propagation, dynamic control variable adjustment, degradation-aware operation, and hydrogen demand satisfaction constraints. The methodology enables techno-economically optimized operation under realistic industrial conditions and is directly applicable to commercial-scale green hydrogen facilities with minimal adaptation requirements.

METHODOLOGY

System description

The integrated system for hydrogen production comprises renewable power generation units (wind and solar photovoltaic) coupled with a PEM water electrolysis unit. To attenuate high-frequency power fluctuations, lithium-ion BESS and compressed hydrogen storage vessels are implemented as buffering components. The system interfaces with both the electrical grid and the downstream hydrogen-consuming process.

As illustrated in **Figure 1**, power from the renewable generation and grid connection is distributed between the BESS and the electrolyzer. Hydrogen produced by the electrolyzer is partitioned between the storage vessels and the downstream facility. During periods of insufficient renewable generation, electrical energy and hydrogen are supplied from BESS discharge and tank depletion, respectively.

The system dynamics can be expressed mathematically as follows:

$$E_{ren} + E_{grid} = E_{bess} + E_{el} \tag{1}$$

$$F_{el} = c_{el} \bullet E_{el} \tag{2}$$

$$F_{el} = F_{tank} + F_{plant} \tag{3}$$

$$\frac{dQ_{bess}}{dt} = \eta_{bess} \bullet E_{bess} - \kappa_{bess} \bullet Q_{bess} \tag{4}$$

$$\frac{dm_{tank}}{dt} = \eta_{tank} \bullet F_{tank} - \kappa_{tank} \bullet m_{tank} \tag{5}$$

Each variable corresponds to the mass and energy streams depicted in **Figure 1**. The parameter c_{el} represents the electrolyzer conversion coefficient for hydrogen production, while η and κ denote the charging efficiency and self-discharge rate for the BESS and hydrogen storage.

Since E_{ren} is determined by exogenous renewable resource availability, the system contains 8 variables. Given 5 governing equations, the formulation exhibits 3 degrees of freedom.

Approach to the problem

Given weather forecast data over the next $t_{regulation}$ time steps, the system has $3 \bullet t_{regulation}$ degrees of freedom. Following a receding horizon strategy, the optimal decision variable trajectory is computed via optimization, with only the initial time step implemented and subsequent values discarded upon re-optimization [9].

This formulation shares conceptual similarities with model predictive control (MPC) frameworks [10]. However, whereas MPC constitutes an advanced real-time control architecture, the proposed methodology generates optimal reference production trajectories for implementation within existing control systems. In other words, the proposed methodology constitutes a supervisory optimization layer that provides optimal setpoints corresponding to the optimal configuration of the decision variables, while real-time tracking and disturbance rejection are handled by the underlying control system. The design and implementation of such control systems

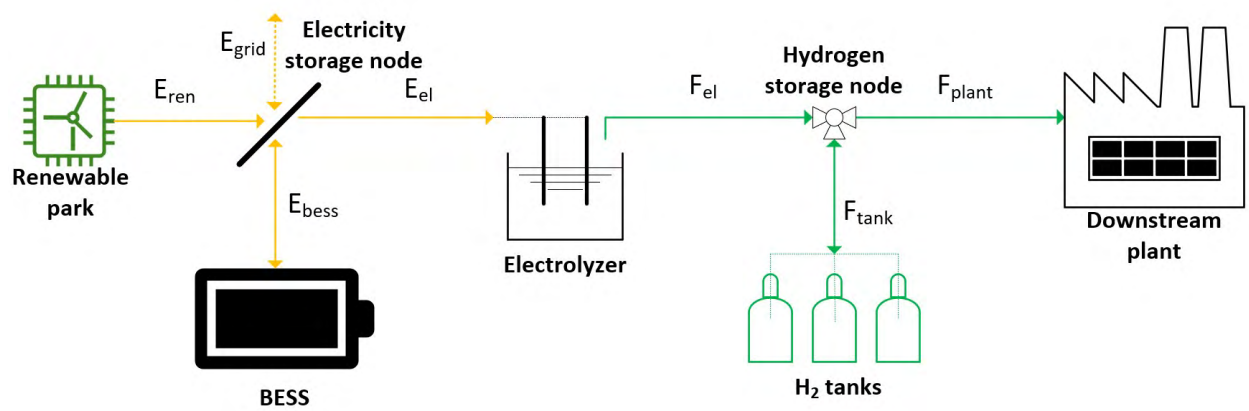

Figure 1. Hydrogen production plant scheme.

are beyond the scope of this work, as the focus is on the optimization-based determination of the setpoint.

The optimization requires specification of the three manipulated variables per time step, selected as E_{bess}, F_{tank}, and E_{grid} based on their direct controllability. The primary objective ensures that F_{plant} tracks the specified hydrogen demand profile for each time step.

$$f_{rp} = \sum_{t=1}^{t_{regulation}} k_1 \bullet \left| \frac{F_{plant}^{t} - F_{plant}^{rp}}{F_{plant}^{rp}} \right| \tag{6}$$

Additionally, to minimize equipment degradation, solutions exhibiting reduced high-frequency variability are favored. This applies specifically to hydrogen flow to the downstream process, electrical power to the electrolyzer, and charging/discharging rates of the buffer systems.

$$\begin{aligned} f_{var} = \sum_{t=1}^{t_{regulation}} \Bigg(& k_2 \bullet \left| \frac{F_{plant}^{t} - F_{plant}^{t-1}}{F_{plant}^{rp}} \right| + k_3 \\ & \bullet \left| \frac{E_{bess}^{t} - E_{bess}^{t-1}}{Q_{bess}^{max}} \right| + k_4 \\ & \bullet \left| \frac{F_{tank}^{t} - F_{tank}^{t-1}}{m_{tank}^{max}} \right| + k_5 \\ & \bullet \left| \frac{E_{el}^{t} - E_{el}^{t-1}}{E_{el}^{nominal}} \right| \Bigg) \end{aligned} \tag{7}$$

The economic impact of grid interaction must also be incorporated, accounting for electricity import costs and export revenues (which may be positive, negative, or zero depending on market conditions).

$$\begin{aligned} f_{grid} = \sum_{t=1}^{t_{regulation}} \big(& k_6 \bullet \max(E_{grid}^{t}, 0) + k_7 \\ & \bullet \min(E_{grid}, 0) \big) \end{aligned} \tag{8}$$

Operational constraints are imposed to ensure solution feasibility. Two constraint categories are defined: bounds on state and control variables to enforce operational limits and physical constraints, and rate-of-change limitations to restrict ramping behavior. Each penalty term equals the constraint violation magnitude multiplied by an appropriate penalty coefficient, or zero when the constraint is satisfied.

$$P_{value}: \begin{cases} Q_{bess}^{min} \le Q_{bess}^{t} \le Q_{bess}^{max} \\ m_{tank}^{min} \le m_{tank}^{t} \le m_{tank}^{max} \\ E_{el}^{min} \le E_{el}^{t} \le E_{el}^{max} \\ F_{plant}^{min} \le F_{plant}^{t} \le F_{plant}^{max} \end{cases} \tag{9}$$

$$P_{var}: \begin{cases} |E_{bess}^{t} - E_{bess}^{t-1}| \le \delta_{bess}^{max} \\ |F_{tank}^{t} - F_{tank}^{t-1}| \le \delta_{tank}^{max} \\ |E_{el}^{t} - E_{el}^{t-1}| \le \delta_{el}^{max} \\ |F_{plant}^{t} - F_{plant}^{t-1}| \le \delta_{plant}^{max} \end{cases} \tag{10}$$

The complete objective function is formulated as:

$$\Phi = f_{rp} + f_{var} + f_{grid} + P_{value} + P_{var} \tag{11}$$

Minimization of the objective function

Minimization of the objective function presents computational challenges due to its inherent non-convexity and potentially large dimensionality. Furthermore, variations in the weighting coefficients (k_1 to k_7) can result in objective function values spanning several orders of magnitude.

The scaling issue is addressed through a double-logarithmic transformation of the objective function, with an offset term to prevent numerical singularities:

$$\hat{\Phi} = \log_{10}(\log_{10}(\Phi + 10)) \tag{12}$$

The non-convexity is handled by exploiting the cusp-like topology of the objective function, as identified through preliminary grid-search analysis. Under these conditions, Powell's conjugate direction method demonstrates robust convergence behavior. Alternatively, genetic algorithms can be employed to guarantee global optimality, albeit with higher computational expense.

The high-dimensional search space is handled through a warm-start initialization strategy [11], wherein the discarded portion of the previous optimization solution provides the initial guess for subsequent iterations:

$$X_{first\ guess}^{\theta+1} = [X_2^{\theta}, X_3^{\theta}, X_4^{\theta}, \dots, X_{t_{regulation}-1}^{\theta}, X_{t_{regulation}}^{\theta}, X_{t_{regulation}}^{\theta}] \quad (13)$$

where X represents an arbitrary decision variable stream. The final time step is replicated since $t_{regulation}$ initial values are required while only $t_{regulation} - 1$ values are available from the previous solution. For the initial simulation step, where no prior solution exists, genetic algorithms are employed due to their lack of initialization requirements.

Renewable energy forecast modeling

Given the reliance on imperfect renewable energy forecasts and the difficulty in obtaining extended-horizon predictions at high temporal resolution, synthetically generated forecast data are employed in this analysis.

Artificial forecasts are derived from historical renewable generation profiles through application of two sequential perturbations: Savitzky-Golay filtering and time-dependent stochastic deviation.

The Savitzky-Golay filter attenuates high-frequency components while preserving underlying trends, thereby replicating the smoothing effect inherent to forecast models with limited temporal resolution.

Time-dependent stochastic perturbations are subsequently applied to represent forecast degradation with increasing prediction horizon. The filtered profile is modulated by a scaling factor incorporating both temporal distance from the present and normally distributed random variability:

$$\varepsilon_{eff} = (u_1 + u_2 \bullet t) \bullet rand_{normal} \quad (14)$$

$$E_{ren}^{forecast} = E_{ren}^{SG-filtered} \bullet (1 + \varepsilon_{eff}) \quad (15)$$

where the parameters u_1 and u_2 characterize forecast accuracy, and $E_{ren}^{SG-filtered}$ denotes the Savitzky-Golay-filtered generation profile.

This methodology enables a realistic representation of operational forecast behavior, yielding results consistent with actual forecasting system performance.

Constraint-correction procedure

Due to discrepancies between optimal solution computed using forecast data and actual system behavior under real renewable generation profiles, constraint violations may occur during implementation despite satisfying constraints during optimization.

To address this issue, a post-processing correction procedure is applied at each time step to enforce constraint satisfaction based on realized operating conditions.

The correction algorithm proceeds hierarchically from downstream process requirements toward upstream renewable generation. For each system variable (flows and inventories), feasibility is evaluated against operational bounds. When violations are detected, the manipulated variables are adjusted to restore constraint satisfaction.

For a generic variable X, the corrected value X^{final} and the constraint function ψ_j to be satisfied are defined as:

$$X^{final} = \begin{cases} X^{max} & if\ X > X^{max} \\ X & if\ X^{min} \le X \le X^{max} \\ X^{min} & if\ X < X^{min} \end{cases} \quad (16)$$

$$\psi_j = model^X(E_{bess}, F_{tank}, E_{grid}) - X^{final} \quad (17)$$

Where $model^X$ represents the governing equation for variable X.

The hierarchical structure processes constraints sequentially. Initially, the constraint function ψ for F_{plant} is enforced through adjustment of F_{tank}. Subsequently, two independent subsystems address constraints on m_{tank} and E_{el}, coupled respectively with F_{plant} through F_{tank} and E_{bess} adjustments.

Finally, to satisfy the BESS charge constraint Q_{bess}, the corresponding function ψ is incorporated into a system coupling F_{plant} and E_{el} constraints, with all three manipulated variables available for adjustment.

During resolution of the E_{el} system, the previously corrected m_{tank} value may not be preserved; however, this does not compromise system feasibility to all operational bounds. Physically, when energy surplus causes tank saturation, excess power is redirected to the BESS rather than the electrolyzer. In the limiting case where electrolyzer power remains above maximum capacity, power reduction yields a tank inventory slightly below maximum but always exceeding the uncorrected optimization result. Analogous reasoning applies to renewable energy deficit conditions, ensuring no operational conflicts arise.

Productivity target enforcement

The methodology enforces production of a target hydrogen quantity over the simulation period:

$$M_{H_2}^{total} = F_{plant}^{nominal} \bullet \theta_{tot} \quad (18)$$

where $F_{plant}^{nominal}$ denotes the downstream facility nominal flow rate and θ_{tot} represents the total simulation duration.

To achieve this cumulative production target, the F_{plant} reference production requires dynamic adjustment at each time step to account for deviations from nominal production. The objective is precise attainment of $M_{H_2}^{total}$

rather than production maximization.

The cumulative hydrogen production up to time θ can be expressed in discrete form as:

$$M_{H_2}(\theta) = M_{H_2}^{\theta} = \sum_{i=0}^{\theta} F_{plant}^{i} \bullet \Delta\theta \tag{19}$$

To satisfy the productivity constraint, the updated reference production $F_{plant}^{rp\ new}$ must compensate for the remaining production requirement:

$$M_{H_2}^{\theta} + F_{plant}^{rp\ new} \bullet (\theta_{tot} - \theta) = F_{plant}^{nominal} \bullet \theta_{tot} \tag{20}$$

Yielding:

$$F_{plant}^{rp\ new} = \frac{F_{plant}^{nominal} \bullet \theta_{tot} - M_{H_2}^{\theta}}{\theta_{tot} - \theta} \tag{21}$$

The updated reference production is subsequently bounded within operational limits F_{plant}^{min} and F_{plant}^{max}.

This formulation exhibits fundamental limitations: when the computed reference production saturates at either bound, the productivity target becomes unattainable, as the physically required reference production exceeds the admissible operating range. Additionally, stochastic fluctuations cause the average production to systematically deviate from the reference production when operating at the constraints, progressively amplifying the deviation between ideal and constrained reference production values.

An alternative formulation addresses this issue by distributing the production deficit over a fraction of the remaining horizon rather than the entire duration. The original formulation is designated the "ideal reference production, " representing the constant flow rate that would achieve the target productivity. The modified formulation, with α denoting the fraction of remaining time considered, becomes:

$$F_{plant}^{rp\ new} = F_{plant}^{nominal} + \frac{F_{plant}^{nominal} \bullet \theta - M_{H_2}^{\theta}}{(\theta_{tot} - \theta) \bullet \alpha} \tag{22}$$

This expression is also subject to feasibility constraints within operational bounds.

The revised formulation allows the ideal reference production to converge toward nominal conditions, thereby reducing cumulative productivity deviations.

Both formulations exhibit pronounced oscillatory behavior near simulation termination, as illustrated in **Figure 2**, attributable to the diminishing remaining time in the denominator, which amplifies the sensitivity to production deviations.

Figure 2 also demonstrates that reference productions persistently exceed nominal production rates throughout most of the simulation period. This phenomenon arises because renewable energy deficits necessitate elevated reference productions, while surplus renewable generation is preferentially allocated to buffer charging rather than production intensification.

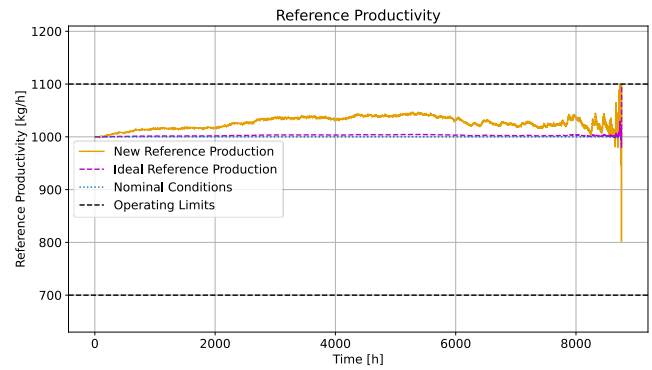

Figure 2: Comparison of the reference productions: “new” (solid orange), “ideal” (dashed magenta) and nominal conditions (dotted blue).

RESULTS AND DISCUSSION

The proposed methodology is validated through a case study focused on hydrogen production for a downstream ammonia plant. Renewable energy availability is derived from data provided by the California Independent System Operator (CAISO) [12] from 2020.

The optimal sizing of the units, the definition of nominal operating conditions for the downstream plant, and the identification of operational constraints are determined using the methodologies developed by Isella and Manca, as described in [13, 14] and subsequently integrated in [15]. Constant parameters are obtained from technical reports and peer-reviewed scientific literature. Weighting factors are fine-tuned based on a combination of parametric analysis and engineering judgment, reflecting the relative importance of each objective term. Parameters such as $t_{regulation}$ and α are selected as compromises between, respectively, model accuracy and computational efficiency, and between error correction speed and reference production oscillations.

The two buffer units and the electrolyzer are assumed to be highly flexible, with no limitations related to rapid variations in ramping rates. Additionally, the self-discharge of both buffer units is considered negligible, and the charging and discharging efficiencies of the hydrogen tank are assumed to be unity, implying no hydrogen losses.

The one-year dynamic simulation required 2106 s of CPU time on a conventional Windows computer (Intel Core i7-1195G7 @ 2.9 GHz, 11th generation, 16 GB RAM), running Python code pre-compiled using Numba JIT. The average computation time per simulation step was 0.24 s, demonstrating that the proposed methodology is compatible with real-time execution.

The results of the case study show that the average renewable and grid power supplied to the battery energy

storage system (BESS) and the electrolyzer is approximately 54.16 MW. This value is consistent with the electrolyzer energy demand under nominal operating conditions, as expressed by:

$$E_{el}^{nominal} = F_{plant}^{nominal}/c_{el} \approx 54\ MW \tag{23}$$

where $F_{plant}^{nominal} = 1000\frac{kg_{H_2}}{h}$ and $c_{el} = 18.5\ kg_{H_2}/MWh$.

Both the BESS and the hydrogen tank exhibit similar seasonal trends, with higher state-of-charge levels during the summer and nearly depleted conditions during the winter. The BESS, however, experiences larger fluctuations due to its smaller capacity. The sizing methodology [15] results in a BESS capacity of 31.81 MWh and a tank capacity of 14.58 t_{H_2}. Consequently, the BESS frequently reaches full charge or discharge to mitigate renewable power variability.

As expected, the electrolyzer operates predominantly near its nominal capacity to ensure the target hydrogen production rate of 1000 kg_{H_2}/h. Owing to its high flexibility, the electrolyzer exhibits frequent load variations, which contribute to a reduced required BESS capacity relative to the hydrogen tank. The larger tank size is also justified by the stringent constraints imposed by the downstream plant, which require the hydrogen flow rate to remain within the range of 700 to 1100 kg_{H_2}/h, as can be seen in **Figure 3**. In this context, the tank plays a crucial role in ensuring operational stability and preventing sudden and large fluctuations in hydrogen supply.

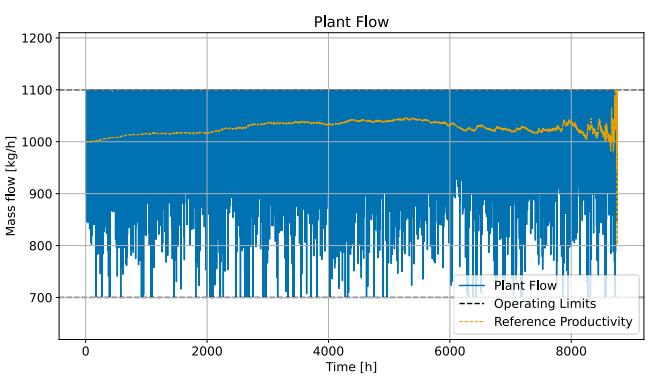

Figure 3: Hydrogen flow rate to the downstream plant.

A comparison between the hydrogen flow rate obtained in this scenario and that of a configuration without scheduling and buffer systems highlights the importance of the proposed approach. In the absence of buffers and scheduling, operational constraints are systematically violated, and the plant fails to meet its required operating limits.

From an economic standpoint, the minimum levelized cost of hydrogen (LCOH), accounting only for capital expenditure, is given by:

$$LCOH_{min} = \frac{CAPEX}{M_{H2}^{total}} = \frac{\sum_j size_j \bullet \frac{cost_j}{life_j}}{M_{H2}^{total}} = 3.00\frac{USD}{kg_{H_2}} \tag{24}$$

where the capital cost is defined as the sum of the purchase costs of each unit divided by their respective expected lifetimes.

The actual LCOH also includes the cost of electricity consumption and the effective hydrogen production. In the present analysis, revenues from selling excess electricity to the grid are not considered, and the excess renewable generation is curtailed; therefore, the resulting LCOH represents a conservative estimate. Assuming a constant electricity price of 100 USD/MWh, the final LCOH is:

$$LCOH_{eff} = \frac{CAPEX + cost_{grid} \bullet E_{grid}^{buy}}{M_{H2}^{\theta_{tot}}} = 3.31\frac{USD}{kg_{H_2}} \tag{25}$$

The total hydrogen production over the simulation period amounts to 8759.949 t_{H_2}, compared to a target value of 8760 t_{H_2}. The discrepancy of 51 kg_{H_2} is considered negligible and can be attributed to forecast–reality mismatches. First, industrial flowmeters typically exhibit an accuracy of approximately 0.25%, corresponding to a minimum detectable deviation of about 22 t_{H_2}, which is several orders of magnitude larger than the observed difference. Second, contractual agreements with the downstream plant are assumed to allow for limited operational flexibility, such that a deviation of 0.0006% does not incur penalties. Finally, at the end of the simulation, approximately 670 kg_{H_2} of hydrogen remains available in the tank, allowing the missing production to be compensated without violating any operational constraints and fully compensating for the observed overproduction.

From an environmental perspective, the only emissions associated with hydrogen production in the proposed system originate from grid electricity generation. The simulation indicates a grid electricity consumption of 26.94 GWh, which corresponds to estimated emissions of 0.92 kg_{CO_2}/kg_{H_2}, assuming a grid emission factor of 0.3 kg_{CO_2}/kWh.

This emission intensity can be compared with conventional hydrogen production pathways, such as steam methane reforming and water-gas shift processes, whose emissions are typically estimated to range from 10 kg_{CO_2}/kg_{H_2} for natural gas-based systems to 26 kg_{CO_2}/kg_{H_2} for coal-based systems. The proposed approach therefore achieves a reduction of more than one order of magnitude relative to conventional technologies.

The current system configuration does not explicitly include the compression stage between the electrolyzer, the hydrogen tank, and the downstream plant. However, a post-processing analysis indicates that incorporating

the compression system would increase the final LCOH by less than 3%.

The specific compression work is evaluated based on the inlet and outlet pressures (corresponding to the electrolyzer outlet pressure and the tank inlet pressure, representing the case with the largest pressure differential), the electrolyzer operating temperature, the compressor efficiency, and the thermodynamic properties of hydrogen:

$$l_{spec}^{comp} = \frac{1}{\eta_{comp}} \cdot \frac{\gamma}{\gamma - 1} \cdot \frac{R}{MW_{H_2}} \cdot T_1 \cdot \left[\left(\frac{P_2}{P_1} \right)^{\frac{\gamma-1}{\gamma}} - 1 \right] \quad (26)$$

Assuming a pressure ratio of 30, an inlet temperature of 60 °C, and a compressor efficiency of 0.6, the resulting specific compression work is 0.004 MWh/kg_{H_2}.

If the required compression energy is supplied exclusively from the grid and only in the absence of renewable energy surplus, the updated LCOH becomes 3.42 USD/kg_{H_2}. This corresponds to an increase of less than 3%, or 0.11 USD/kg_{H_2}, relative to the configuration without the compression unit.

CONCLUSIONS

This work presents a novel approach for the scheduling of hydrogen production plants. The proposed framework introduces a robust methodology that solves a multidimensional optimization problem over a receding horizon, enabling the delivery of a target hydrogen flow rate to the downstream plant while simultaneously reducing unit stress and grid electricity consumption, and ensuring strict compliance with all operational constraints.

The use of imperfect renewable energy forecasts significantly enhances the realism of the analysis and demonstrates the robustness of the methodology in managing discrepancies between forecasted and actual energy availability. Despite these mismatches, the system consistently satisfies all operational constraints. A further key feature of the proposed approach is its ability to guarantee the target hydrogen productivity through dynamic adjustment of the operating reference production.

From an economic standpoint, the resulting levelized cost of hydrogen (LCOH) is consistent with values reported in the literature and remains below current industrial hydrogen production costs. From an environmental perspective, the methodology achieves a reduction in carbon dioxide emissions of approximately one order of magnitude compared to conventional hydrogen production pathways, such as natural gas- or coal-based processes.

The proposed approach is fully automated and does not require human intervention during operation, thereby reducing the risk of human-related errors. The operator's role is limited to system supervision and fault detection. Furthermore, the methodology is general in nature and can be readily adapted to existing industrial facilities, offering a practical pathway to improve operational efficiency and overall system performance.

Despite these advantages, several limitations should be acknowledged. First, the renewable energy forecast model lacks full accuracy, as it is derived as a post-processing step from real energy profiles. Additionally, the stochastic component is defined as a fixed percentage of the filtered profile, resulting in increasing absolute deviations at higher renewable generation levels (e.g., during summer periods), even though the relative error remains constant. Second, operational and maintenance costs are not included in the LCOH calculation, which would lead to a higher overall cost estimate. Third, the electricity price is assumed to be constant, whereas in practice it is time-dependent and subject to market variability. Finally, although its contribution is relatively minor, the compression system should be explicitly integrated into the main system model.

Future developments may include the extension of the methodology to islanded systems, where access to grid electricity is not available. In addition, a more detailed and explicit modeling of downstream processes, such as ammonia or methanol synthesis plants, would further enhance the industrial relevance and applicability of the proposed approach.

AUTHOR IDENTIFIERS

Author ORCIDs:
Lea Casagrande M: 0009-0009-5672-3041
Isella A: 0000-0002-3699-4698
Manca D: 0000-0003-2055-9752

REFERENCES

1. IEA. Global Energy Review 2025 (2025). https://www.iea.org/
2. WRI. Where Do Emissions Come From? 4 Charts Explain Greenhouse Gas Emissions by Sector (2024). https://www.wri.org/insights/4-charts-explain-greenhouse-gas-emissions-countries-and-sectors
3. Eurostat. Renewable Energy Statistics (2024). https://ec.europa.eu/eurostat/statisticsexplained/
4. Hasan MM, Genç G. Techno-economic analysis of solar/wind power based hydrogen production. Fuel 324:124564 (2022). https://doi.org/10.1016/j.fuel.2022.124564
5. Isella A, Ostuni R, Manca D. Towards the decarbonization of ammonia synthesis – a techno-economic assessment of hybrid-green process alternatives. Chemical Engineering Journal 486:150132 (2024).

https://doi.org/10.1016/j.cej.2024.150132
6. Roy TK. Techno-economic and environmental assessment of green hydrogen production in multiple australian regions using different electrolyzer technologies. Renewable Energy 256:123879 (2026). https://doi.org/10.1016/j.renene.2025.123879
7. Salmon N, Bañares-Alcántara R. Impact of process flexibility and imperfect forecasting on the operation and design of haber–bosch green ammonia. RSC Sustainability 1:923-937 (2023). https://doi.org/10.1039/d3su00067b
8. Schmitt T, Engel J, Kopp M, Rodemann T. Implicit incorporation of heuristics in mpc-based control of a hydrogen plant. 2023 IEEE International Conference on Power Electronics, Smart Grid, and Renewable Energy (PESGRE) :1-8 (2023). https://doi.org/10.1109/pesgre58662.2023.10404171
9. García CE, Prett DM, Morari M. Model predictive control: theory and practice—a survey. Automatica 25:335-348 (1989). https://doi.org/10.1016/0005-1098(89)90002-2
10. Morari M, H Lee J. Model predictive control: past, present and future. Computers & Chemical Engineering 23:667-682 (1999). https://doi.org/10.1016/s0098-1354(98)00301-9
11. Marcucci T, Tedrake R. Warm start of mixed-integer programs for model predictive control of hybrid systems. IEEE Trans. Automat. Contr. 66:2433-2448 (2021). https://doi.org/10.1109/tac.2020.3007688
12. CAISO. Production and curtailments data (2025). https://www.caiso.com/library/production-curtailments-data
13. Isella A, Manca D. A framework for the design of battery energy storage systems in power-to-x processes. Journal of Energy Storage 123:116744 (2025). https://doi.org/10.1016/j.est.2025.116744
14. Isella A, Manca D. A general criterion for the design and operation of flexible hydrogen storage in power-to-x processes. International Journal of Hydrogen Energy 99:649-660 (2025). https://doi.org/10.1016/j.ijhydene.2024.12.228
15. Isella A, Manca D. Integrating economic, environmental, and social sustainability in power-to-ammonia plants: a multi-objective optimization methodology. Computers & Chemical Engineering 198:109082 (2025). https://doi.org/10.1016/j.compchemeng.2025.109082

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Forecasting Time-to-Cyclic Steady State in Periodic Bioprocesses via a Multi-Feature k-Nearest Neighbours Framework

Yasser Algoufily[ab*], Foteini Michalopoulou[ab], Maria M. Papathanasiou[ab], Mehmet Mercangöz[ab]

[a] Imperial College London, Department of Chemical Engineering, London, United Kingdom
[b] Sargent Center for Process Systems Engineering, London, United Kingdom
* Corresponding Author: y.algoufily22@imperial.ac.uk.

ABSTRACT

Early and reliable prediction of convergence to cyclic steady state (CSS) is increasingly important in periodic downstream bioprocessing, where switching and cut decisions are tuned for a repeatable cyclic regime. This work addresses time-to-CSS (TCSS) forecasting and CSS-existence classification for multicolumn countercurrent solvent gradient purification (MCSGP) systems under run-to-run feed variability. We propose a Multi-Feature k-Nearest Neighbours (MF-kNN) framework that performs long-horizon one-shot trajectory forecasting from an early run segment. CSS outcomes are inferred by reapplying a peak-based convergence rule to the predicted trajectory, while CSS existence is predicted via neighbour-label voting. The approach uses multivariate, standardised features, run-level splits, and a windowed neighbour search to reduce computation. Hyperparameters are tuned with a CSS-oriented objective function that balances trajectory fidelity, TCSS error, and misclassification penalties. On an in-silico MCSGP dataset (98 runs; Δt = 0.2 s; 6800 steps/run) with varying initial modifier concentration, MF-kNN produces accurate full-run forecasts from early data and enables operationally useful early go/no-go decisions. Across outlets, results support accurate CSS timing inference and high CSS-existence classification accuracy (up to 100% on selected outlets), indicating MF-kNN as a transparent and deployment-ready complement to cycle-to-cycle CSS monitoring and control.

Keywords: cyclic steady state, periodic bioprocessing, MCSGP, time-to-CSS forecasting, k-nearest neighbours, one-shot forecasting

INTRODUCTION

Biomanufacturing is progressively shifting from conventional batch workflows toward increasingly integrated operations that deliver higher productivity with reduced footprint (e.g., via higher volumetric productivity, reduced hold steps, and tighter integration of unit operations) [1, 2]. This evolution is visible in industrial biopharmaceutical manufacturing, particularly in down-stream purification where chromatographic separations have moved from single-column, non-periodic operation to cyclic multicolumn switching schemes, such as capture purification transitioning from batch Protein A chromatography toward periodic counter-current chromatography (PCC) [3], challenging polishing tasks extending from single-column batch gradient operation to multicolumn countercurrent solvent gradient purification (MCSGP) [4], and, more broadly, simulated moving bed (SMB) chromatography employing cyclic switching to sustain periodic operation with improved utilisation relative to single-column batch processing [10]. Because these schemes rely on repeated switching cycles, their analysis and performance assessment hinge on whether the process has reached a repeatable periodic regime, commonly termed cyclic steady state (CSS).

CSS generally refers to the regime in periodic operation where cycle-to-cycle behaviour becomes repeatable within defined tolerances, so that outlet profiles and cycle-level outcomes no longer exhibit systematic drift [5]. In MCSGP, this notion is operationalised by requiring that consecutive-cycle chromatograms and cut outcomes meet repeatability criteria, typically assessed by

comparing feature summaries such as peak maxima, peak areas, retention times, and pool compositions against acceptance thresholds tied to process objectives and analytical variability [5]. This repeatability-based workflow also aligns with process-development and performance-qualification practice, where multiple consecutive cycles are evaluated at a candidate operating point to confirm stable cyclic performance before re-porting yields and purities [6].

With CSS as the target operating regime in such periodic processes, an operationally important question is whether the process will attain a repeatable cyclic pattern within the available operating horizon and, if so, when [5]. The answer affects column switching schedules, fraction collection, solvent usage, and product quality in cyclic downstream operations [5]. Overly conservative actions can waste resources and extend runtime, while premature decisions can degrade purity or yield and disrupt intended continuous operation [5].

These considerations motivate reliable early CSS forecasting and classification to support (i) early run-level go/no-go decisions on continued operation, (ii) anticipatory adjustment of fractionation and recycle policies during the transient approach to CSS, and (iii) improved planning of pooling and resource utilisation (e.g., buffer and hold capacities) in integrated downstream trains.

MCSGP provides a representative case study for periodic CSS decision support: it is a semi-continuous chromatography platform for biomolecule purification that exhibits strongly periodic dynamics under coordinated column switching and solvent gradients [4]. These dynamics are shaped by nonlinear adsorption and mass transfer, while feed variability and initial conditions in-fluence convergence toward CSS [5]. In practice, operating poli-cies such as gradient programs, cut criteria, and recycle routing are typically tuned for the cyclic regime; therefore, uncertainty about CSS convergence and timing propagates into decision un-certainty and disturbances [5]. Accordingly, this work examines whether early-cycle information can enable prospective inference of CSS convergence and timing in MCSGP, shifting beyond retrospective confirmation toward prospective decision support that reduces the latency inherent to cycle-to-cycle monitoring and UV-based control [5].

1.1 Literature review: CSS detection and management in periodic bioprocesses

A substantial body of MCSGP work has therefore focused on CSS management via closed-loop operation rather than purely offline certification. Optimisation and control studies treat MC-SGP as a cyclic, hybrid process in which switching and gradient decisions must remain robust under nonlinear adsorption, mass transfer, and delayed measurements, motivating cycle-to-cycle formulations and multi-rate measurement use [5]. Building on this, UV-based closed-loop strategies have been proposed to regulate key cycle outcomes via online peak-tracking and cut adjustment, demonstrating practical controllers that mitigate drift in chromatographic signatures once operation is near the cyclic regime [5]. More recent implementations extend these concepts to improve robustness under disturbances such as feed variability and retention-time shifts, targeting earlier convergence and stable cyclic operation using limited online information [6]. Overall, the MCSGP literature converges on a consistent operational picture: CSS is detected and maintained primarily through cycle-to-cycle repeatability of chromatographic signatures and control actions that keep cycle-level KPIs in-spec [5, 6].

When CSS is addressed predictively, it is most often through model-based simulation and surrogate acceleration rather than prospective inference from early online data. Mechanistic models have been used to simulate MCSGP start-up, CSS attainment, and shutdown behaviour and to predict performance parameters under specified operating conditions, at the cost of repeated-cycle simulation [7]. To reduce this computational burden, hybrid or data-driven surrogates have been developed that target rapid prediction of CSS behaviour for optimisation and decision support (for example, hybrid formulations that achieve CSS predictions substantially faster than high-fidelity models) [8]. Closely related efforts in cyclic adsorption more broadly also emphasise accelerating cyclic-process prediction and synthesis, for instance by physics-informed neural network surrogates designed for cyclic adsorption calculations [9]. These strands primarily support faster CSS computation given a model or training set, rather than early-cycle, forecast-driven go/no-go decisions using only initial-cycle measurements.

Beyond MCSGP, cyclic steady state is a foundational concept in other periodic separations and bioprocess settings. In simulated moving bed (SMB) and related multicolumn chromatography, CSS is often defined formally via periodicity constraints (end-of-cycle states equal start-of-cycle states) and assessed via convergence of cycle-to-cycle outlet profiles and KPIs [10]. The corresponding literature includes dedicated numerical methods for computing CSS efficiently, reflecting the importance of cycle-level periodicity in SMB design and optimisation [11, 10]. Similar convergence logic appears in other cyclic bioprocesses, for example sequencing batch reactors, where steady cyclic operation is assessed when cycle-aggregated performance indicators become repeatable over successive cycles, enabling consistent operation and control [12, 13]. Across these domains, the dominant CSS handling paradigm remains cycle-to-cycle convergence of either (i) observable signatures (peaks, areas, outlet profiles) or (ii) cycle-aggregated KPIs, with control-oriented studies prioritising robust repeatability

under disturbances [5, 10].

1.2 Research gaps and contributions

In MCSGP, and more broadly in periodic downstream and cyclic bioprocess operations, CSS decision support is still dominated by repeatability-based confirmation rather than prospective pre-diction from early-cycle information. Despite the mature body of CSS monitoring and control work surveyed in the previous section, a central limitation for early decision support is that CSS is commonly confirmed retrospectively via consecutive-cycle repeatability checks. This retrospective paradigm is not ideal because it:

- Provides limited quantitative information about proximity to CSS and remaining stabilisation time, constraining proactive planning during start-up.
- Introduces an inherent confirmation latency, since a candidate CSS convergence point can only be validated after observing subsequent cycles (or peaks) that satisfy the tolerance.
- Ties switching, cutting, and recycle actions to confirmed repeatability, delaying high-value decisions (e.g., collection start, recycle routing).

While mechanistic models and advanced nonlinear surrogates can, in principle, predict CSS behaviour, such approaches are often computationally heavy and difficult to deploy online under run-to-run variability [15]. As a result, operators may delay high-value actions until CSS is verified or take conservative actions that reduce productivity. These limitations motivate forecast-driven methods that deliver early, computationally light decision support in periodic operation.

This work addresses these gaps with a multi-feature k-Nearest Neighbours (MF-kNN) framework [14], developed and evaluated on MCSGP, for TCSS forecasting and CSS-existence classification. By leveraging multivariate, physically meaningful signals, MF-kNN provides long-horizon one-shot predictions from early-cycle data, enabling earlier decisions than retrospective cycle-to-cycle confirmation and UV-triggered actions. Performance is benchmarked against classical and learning-based forecasting models, alongside a univariate kNN model, to contextualise accuracy and runtime. Specifically, MF-kNN is tailored to CSS decision support through:

- CSS-aligned inference: CSS is inferred from forecasts using a peak-based convergence rule yielding both TCSS estimates and a CSS-existence classification.
- Multivariate similarity over measured chromatographic signals: a feature-weighted distance over multiple online variables (e.g., outlet concentration traces and modifier profiles), with standardisation to keep signals comparable and avoid dominance by any single measurement.
- Localised neighbour search: a ±W windowed neighbour search around the forecast start time that improves relevance and reduces runtime, supporting online use.
- CSS-oriented validation and tuning: a validation objective that jointly scores trajectory accuracy and TCSS error, with explicit penalties for incorrect CSS-existence predictions to reduce optimistic false positives.
- Early-cycle feasibility: early, one-shot prediction using only the initial segment of the run.

METHODOLOGY

2.1 Process overview and dataset

We consider periodic MCSGP runs, where coordinated column switching and a solvent-modifier gradient induce repeated elution cycles. The data was generated in silico using a validated gPROMS model of the MCSGP process under run-to-run variability in the initial modifier setting. The dataset comprises N_{run} = 98 runs totalling 666, 400 timesteps, sampled every $\Delta t = 0.2\,s$, with fixed run length T = 6800 timesteps (approximately 22.67 min/run).

For each run $r \in \{1, \ldots, N_{run}\}$, the recorded signals include inlet and outlet concentration traces for the four tracked components, $\{C_{\text{in},c,r}(t), C_{\text{out},c,r}(t)\}_{c=1}^{4}$, where $c = 1$ denotes the modifier (buffer), $c = 2$ weak impurities, $c = 3$ product, and $c = 4$ strong impurities. The initial modifier value is set at the start of each run and varies across runs; we denote this run-level preset by $C_{\text{in},1,r}(0)$. The corresponding modifier trajectory $C_{\text{in},1,r}(t)$ evolves over the run according to the imposed gradient program. This exogenous run-to-run variability influences both whether CSS is reached within the finite run horizon and the resulting TCSS, thereby complicating reliable early detection and forecasting. Figures 1 and 2 illustrate representative periodic concentration profiles and the distribution of $C_{\text{in},1,r}(0)$ across runs, respectively.

2.2 Problem definition

For run r, let $\boldsymbol{x}_r(t)$ denote a designated target concentration profile and let $\{\boldsymbol{u}_{r,j}(t)\}_{j=2}^{p}$ denote auxiliary measured trajectories over $t = 0, \ldots, T$. We define the TCSS for run r as $TCSS_r$. At an early initialisation time $t_0 < T - h$, where h denotes the forecast horizon (timesteps), the objectives are to: (i) generate a one-shot forecast of the remainder of the run, i.e., predict $\boldsymbol{x}_r(t)$ for $t = t0 + 1, \ldots, T$; (ii) classify whether CSS will be achieved within the run horizon; and, if CSS is declared, then (iii) estimate $TCSS_r$.

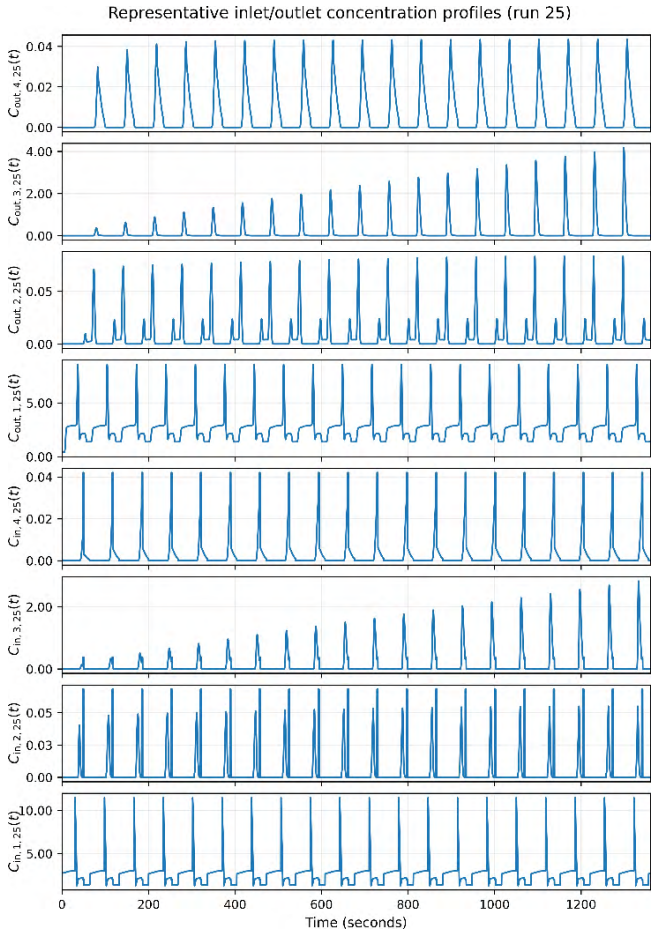

Figure 1. Representative concentration profiles for run 25 showing periodic elution cycles and convergence variability.

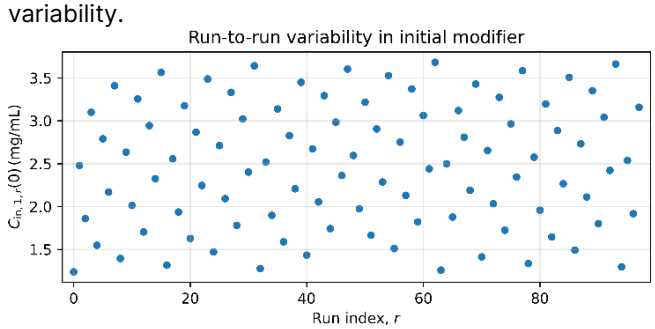

Figure 2. Run-to-run variability in the initial modifier concentration

2.3 Peak-based CSS detection

Peaks are detected on the target series $\boldsymbol{x}_r(t)$ and filtered to retain only the large peaks. Smaller peaks with amplitudes below a fixed cutoff α are discarded, and only peaks with $\boldsymbol{x}_r(t_{peak}) \geq \alpha$ are retained. Let $\{t_{\text{peak},i}, A_i\}_{i=1}^{N_{\text{peak}}}$ denote the retained peak times and amplitudes, where A_i is the amplitude of the i-th large peak. CSS is declared when consecutive large-peak amplitudes remain within tolerance for q successive peak pairs:

$$|A_{i+\ell} - A_{i+\ell+1}| < \varepsilon, \qquad \ell = 0, \dots, q-1 \tag{1}$$

where ε is the peak-amplitude tolerance and q is the required number of consecutive satisfactions. If Eq. (1) is first satisfied at peak index $i^\star$, then the TCSS is defined as:

$$T_{\text{CSS},r} := t_{\text{peak},i^\star} \tag{2}$$

A binary CSS-existence label $z_r \in \{0,1\}$ is assigned based on whether CSS is reached within the finite run horizon. As illustrated in Figure 3, the TCSS is taken as the time index of the first peak that satisfies the convergence criterion.

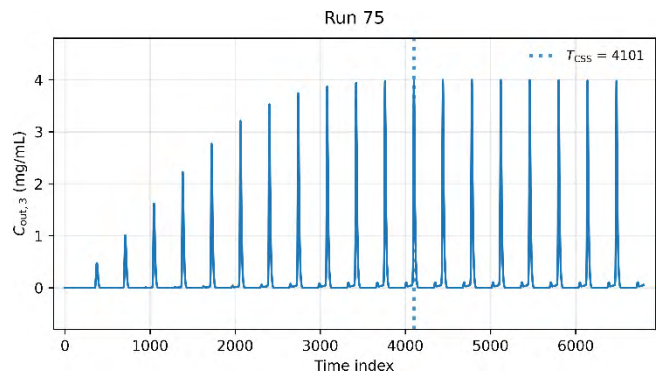

Figure 3. Peak-based CSS definition

2.4 MF-kNN one-shot forecasting

For run r, we construct a $p \times (T+1)$ data matrix by stacking the target and auxiliary signals:

$$Y_r(T) = \begin{bmatrix} x_r(0) & x_r(1) & \cdots & x_r(T) \\ u_{r,2}(0) & u_{r,2}(1) & \cdots & u_{r,2}(T) \\ \vdots & \vdots & \ddots & \vdots \\ u_{r,p}(0) & u_{r,p}(1) & \cdots & u_{r,p}(T) \end{bmatrix} \tag{3}$$

Here, feature index $j = 1$ corresponds to the target $\boldsymbol{x}_r(t)$, and $j = 2, \dots, p$ correspond to auxiliary trajectories $\boldsymbol{u}_{r,j}(t)$. At an initialisation time $t_0 < T - h$, we represent the recent evolution of each feature using a unified embedding length m (i.e., $m_j \equiv m$ for $j = 1, \dots, p$). For the target ($j = 1$), we define the embedded evolution segment:

$$\boldsymbol{x}_r(t_0) := [x_r(t_0 - m + 1), \dots, x_r(t_0)]^\top \in R^m \tag{4}$$

Similarly, for each auxiliary feature $j = 2, \dots, p$, we define:

$$\boldsymbol{u}_{r,j}(t_0) := [u_{r,j}(t_0 - m + 1), \dots, u_{r,j}(t_0)]^\top \in R^m \tag{5}$$

The multi-feature evolution segment is then formed by stacking the embedded segments as rows:

$$Y_r(t_0) = \begin{bmatrix} \boldsymbol{x}_r(t_0)^\top \\ \boldsymbol{u}_{r,2}(t_0)^\top \\ \vdots \\ \boldsymbol{u}_{r,p}(t_0)^\top \end{bmatrix} \in R^{p \times m} \tag{6}$$

Multi-feature distance

Given a query run r at initialisation time t_0, candidate neighbours are drawn from memory runs $r' \in R_{mem}$. For each memory run r' and admissible endpoint τ, the candidate segment $Y_{r'}(\tau) \in R^{p\times m}$ is constructed identically to $Y_r(t_0)$ by taking the most recent m samples of each feature up to τ, with τ constrained to $m-1 \le \tau \le T-h$. Similarity between the query segment $\boldsymbol{Y_r}(t_0)$ and a candidate segment $\boldsymbol{Y_{r'}}(\tau)$ is computed via a feature-weighted Euclidean distance. Defining $\Delta\mathbf{Y} := \boldsymbol{Y_r}(t_0) - \boldsymbol{Y_{r'}}(\tau)$, we compute:

$$D\big(\boldsymbol{Y_r}(t_0), \boldsymbol{Y_{r'}}(\tau)\big) = \sqrt{\sum_{j=1}^{p} w_{f,j} \left\|\Delta\boldsymbol{Y_{j,:}}\right\|_2^2} \tag{7}$$

where $\boldsymbol{Y_{j,:}}$ denotes the j-th row and $w_{f,j} \ge 0$ with $\sum_{j=1}^{p} w_{f,j} = 1$.

Neighbour selection

The k candidates with smallest distance D are selected as neighbours, indexed by pairs $\{r_n, \tau_n\}_{n=1}^{k}$. These neighbours are then weighted to favour candidates that are (i) close in the embedded signal space and (ii) close in the run-level initial modifier setting. First, we define the distance proximity score:

$$s_n^{(D)} := \frac{1}{2}\left(1 - \frac{D_n}{\sum_{\ell=1}^{k} D_\ell}\right) \tag{8}$$

Here $D_n := D\big(Y_r(t_0), Y_{r_n}(\tau_n)\big)$. Next, define the modifier proximity score:

$$s_n^{(C)} := \frac{1}{2}\left(1 - \frac{\Delta C_n}{\sum_{\ell=1}^{k} \Delta C_\ell}\right) \tag{9}$$

Here $\Delta C_n := |C_{\text{in},1,r}(0) - C_{\text{in},1,r_n}(0)|$. If a normalising denominator is zero, the corresponding score is set uniformly over the k neighbours. The final aggregation weights are obtained by summing the two scores and normalising to sum to one:

$$\omega_n := \frac{s_n^{(D)} + s_n^{(C)}}{\sum_{l=1}^{k}\left(s_\ell^{(D)} + s_\ell^{(C)}\right)}, n = 1, \ldots, k \tag{10}$$

One-shot forecast

Finally, the one-shot forecast of the target trajectory is obtained by weighted aggregation of the neighbours' future continuations:

$$\hat{\boldsymbol{x}}_{\boldsymbol{r}}(t_0+1{:}\,t_0+h) = \sum_{n=1}^{k} \omega_n\, x_{r_n}(\tau_n+1{:}\,\tau_n+h) \tag{11}$$

CSS inference

Because historical runs are labelled with CSS existence, the neighbour set supports classification without an additional model. Let $z^{(n)}$ be the CSS-existence label for neighbour n. The CSS-existence decision is then:

$$\hat{z} = I\left(\sum_{n=1}^{k} \omega_n\, z^{(n)} \ge \eta\right) \tag{12}$$

with voting threshold η (default $\eta = 0.5$). When $\hat{z} = 1$, the predicted $\hat{T}_{\text{CSS}}$ is obtained by applying the same peak rule (Eq. (1)) to the full reconstructed trajectory; i.e. using measured values for $t < t_0$ and the one-shot forecast for $t \ge t_0$.

2.5 Hyperparameter tuning objective

We tune k, feature weights w_f, and any embedding or window parameters, to balance trajectory fidelity, TCSS accuracy, and CSS-existence correctness. Let $\mathcal{L}_{traj}$ denote a trajectory loss over the forecast horizon. We define:

$$\text{Obj}_1 = \frac{1}{N}\sum_{r=1}^{N} \mathcal{L}_{traj}(\boldsymbol{x_r}, \hat{\boldsymbol{x}}_{\boldsymbol{r}}) \tag{13}$$

$$\text{Obj}_2 = \frac{1}{N}\sum_{r=1}^{N} \begin{cases} |T_{\text{CSS},r} - \hat{T}_{\text{CSS},\,r}| & if\ z_r = \hat{z}_r = 1 \\ 0 & if\ z_r = \hat{z}_r = 0 \\ penalty & if\ z_r \ne \hat{z}_r \end{cases} \tag{14}$$

Let N denote the number of validation runs. We combine the trajectory term Obj_1 and the CSS-aligned term Obj_2 into a single scalar objective:

$$\text{Obj}_{\text{overall}} = (1-\beta)\text{Obj}_1 + \beta\text{Obj}_2 \tag{15}$$

where $\beta \in [0, 1]$ controls the trade-off.

2.6 Experimental protocol and implementation notes

We use a run-wise split of $80/10/10\%$ for memory/validation/test. For N_{run} = 98, this yields N_{mem} = 78, N_{val} = 10, and N_{test} = 10. The features are standardised to prevent feature scale dominance. Forecasts are computed one-shot using only the first 100 timesteps of the unfolding run.

Hyperparameters are tuned on validation runs using grid search. Tuned parameters include the neighbour count $k \in \{1, 3, 5, \ldots, k_{max}\}$, feature weights $w_{f,j} \in [0, 1]$, the neighbour aggregation scheme (uniform vs. distance-weighted), and window/embedding settings. Candidate combinations are scored using the composite objective in Eq. (15). For computational efficiency, the neighbour search is restricted to a local temporal window around the forecast initialisation time such that when forecasting a run at t_0, candidate historical segments are only drawn from memory runs with segment endpoints $\tau \in [max(0, t_0 - W), min(T-1, t_0 + W)]$, where W is a window half-width in timesteps. Distances are computed

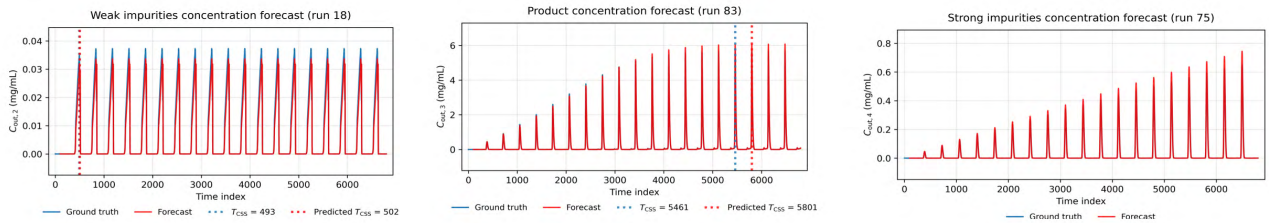

(a) $C_{\text{out},2,18}$ (b) $C_{\text{out},3,83}$ (c) $C_{\text{out},4,75}$

Figure 4. Representative one-shot forecasts for the outlet concentrations initialised at t_0 = 100.

between the query segment ending at t_0 and each candidate segment ending at τ, and the k nearest neighbours are selected by which a one-shot forecast is then generated as defined in Eq. (11). We report trajectory error in MSE, TCSS error in RMSE, and CSS-existence classification accuracy in percentage.

2.7 Benchmark models and comparative evaluation criteria

To contextualise MF-kNN performance, a comparative one-shot forecasting study was conducted against representative classical and learning-based baselines, alongside a univariate kNN model. All models were evaluated under a common run-wise protocol, where forecasts were initialised at the same t_0 timestep, and the prediction horizon was set to the remainder of the run. The forecasting target is the outlet signal $C_{\text{out},3,r}(t)$. Hyperparameters were tuned on the validation runs using model-specific searches.

Benchmarked models. The benchmark suite comprises: (i) MF-kNN; (ii) Univariate kNN; (iii) ARMA as a classical stochastic time-series baseline for periodic signals; (iv) Theta as a lightweight statistical forecasting baseline; and (v) LSTM as a representative recurrent neural forecasting model.

Tuning and scoring. Hyperparameters were tuned via a fixed number of trials on the validation runs. MF-kNN tunes the neighbour count and feature weights; univariate kNN tunes the neighbour count; ARMA tunes (p, q); Theta tunes θ; and the LSTM tunes hidden-unit configuration. Evaluation considers both (a) trajectory prediction error for $C_{\text{out},3,r}(t)$ over the forecast horizon and (b) error on the derived CSS timing metric from the corresponding predictions, so that models are compared on the same decision-support objective.

RESULTS AND DISCUSSION

Qualitative performance. Figure 4 illustrates representative one-shot forecasts for outlet concentration profiles using only the first 100 timesteps. Across these examples, MF-kNN reproduces the dominant periodic dynamics over the full horizon, correctly distinguishes CSS from no-CSS runs, and yields TCSS estimates that are consistent with the peak-based rule.

Quantitative performance. We summarise quantitative results on the held-out test set using three primary metrics. Table 1 reports trajectory accuracy as MSE, TCSS error as RMSE, and CSS-existence classification accuracy (%). Table 2 reports the CSS-existence confusion matrix for $C_{\text{out},3,r}(t)$, from which precision and recall can be computed. As shown in Table 1, MF-kNN achieves consistently low trajectory errors across the outlet signals, indicating that the one-shot forecasts reproduce the dominant periodic structure and amplitude of the concentration profiles. The CSS-existence classification performance is strong overall (average accuracy 96.7%), with perfect classification reported for $C_{\text{out},3,r}(t)$ and $C_{\text{out},4,r}(t)$ and a small degradation for $C_{\text{out},2,r}(t)$ at 90%, suggesting that early segments contain sufficient information to discriminate convergent from non-convergent behaviour in most cases. In contrast, the TCSS RMSE exhibits a clearer dependence on the outlet signal: $C_{\text{out},2,r}(t)$ and $C_{\text{out},3,r}(t)$ show substantially lower timing errors than $C_{\text{out},4,r}(t)$, indicating that CSS timing is harder to infer for $C_{\text{out},4,r}(t)$, likely due to weaker observability of peak stabilisation or greater sensitivity of its peak features to transient dynamics and noise. Overall, this trend highlights that accurate CSS existence prediction can be achieved even when the precise timing of convergence is more uncertain for specific outlets.

Benchmark comparison. Under the common one-shot protocol described in Section 2.7, Table 3 shows that MF-kNN provides the most reliable decision-support performance for $C_{\text{out},3,r}(t)$ among the benchmarked methods. Beyond achieving the lowest CSS timing error, MF-kNN is the only method in the benchmark set that simultaneously maintains low trajectory error and high CSS-existence accuracy, which is the key requirement for early operational decisions. The univariate kNN baseline degrades both timing and classification performance relative to MF-kNN, indicating that leveraging additional measured signals in the distance metric improves neighbour relevance and reduces ambiguity in early-cycle matching. The LSTM and ARMA baselines perform poorly in this setting, which is consistent with the difficulty of

learning stable long-horizon dynamics from limited early segments under run-to-run variability without extensive model capacity and careful regularisation. The Theta model performs worst on both timing and classification, suggesting that lightweight univariate statistical structure is insufficient to capture the nonlinear, multivariate transient CSS behaviour of MCSGP. Overall, the benchmark results support the central premise of this work: multi-feature similarity enables robust early matching to historical trajectories, yielding more reliable long-horizon forecasts and CSS timing inference than both univariate similarity and representative statistical or learning-based baselines in the one-shot regime.

Table 1: Quantitative performance on the test set.

Signal	TCSS RMSE (steps)	Time-series MSE ($\times 10^{-3}$)	CSS classification accuracy (%)
$C_{\text{out},2}(t)$	157.3	0.02	90.0
$C_{\text{out},3}(t)$	340.0	1.15	100.0
$C_{\text{out},4}(t)$	1356.0	0.04	100.0
Avg.	**617.8**	**0.40**	**96.7**

Table 2: CSS-existence confusion matrix for $C_{\text{out},3}(t)$.

	Predicted CSS	Predicted No-CSS
Actual CSS	20.0%	0.0%
Actual No-CSS	0.0%	80.0%

Table 3: Benchmark performance for $C_{\text{out},3}(t)$.

Model	TCSS RMSE (steps)	Time-series MSE ($\times 10^{-3}$)	CSS classification accuracy (%)
MF-kNN	340.0	1.15	100.0
kNN	537.90	14.95	40.0
LSTM	2917.74	57.86	30.0
ARMA	2937.48	60.62	10.0
Theta	4455.30	98.34	10.0

CONCLUSION

This paper introduced a MF-kNN framework for early, one-shot forecasting CSS behaviour in periodic downstream bioprocesses, demonstrated on MCSGP system runs subject to run-level variability in the initial modifier concentration. Using only an early segment of each run, MF-kNN forecasts the full-run outlet trajectories, infers TCSS, and predicts CSS existence via neighbourhood voting. The methodology is designed for practical deployment through run-wise splitting, feature standardisation, and a windowed neighbour search that reduces runtime while preserving local relevance.

Across the evaluated outlet signals, MF-kNN yields consistent qualitative agreement with the true periodic profiles and supports accurate early decision support for CSS convergence. In the benchmark study, MF-kNN outperformed representative statistical and learning-based baselines on the CSS timing objective under a common one-shot protocol, reinforcing the value of similarity-based forecasting when long-horizon predictions must be derived from limited early-cycle information.

From an operational perspective, the proposed framework enables earlier switching and cut decisions, earlier initiation of collection, and early termination or alternative actions when convergence is unlikely, with the potential to reduce solvent usage and wasted runtime during non-convergent start-up.

Future work

Future work will (i) validate the proposed MF-kNN framework on experimental/industrial MCSGP datasets, beyond the current in-silico study; (ii) quantify operational benefits under realistic decision policies, e.g., solvent usage, time savings, and yield/purity impacts from earlier go/no-go and collection-start decisions; (iii) integrate MF-kNN predictions into closed-loop operation, for example as a supervisory layer that adapts cut/collection logic or informs receding-horizon MPC for cyclic operation; and (iv) extend the feature set used for neighbour retrieval to improve robustness and to provide calibrated confidence estimates for TCSS and CSS-existence predictions.

ACKNOWLEDGEMENTS

The authors acknowledge Imperial College London for providing the research environment and computational resources that supported this work. We also thank colleagues in the Department of Chemical Engineering.

AUTHOR IDENTIFIERS

Author ORCIDs:
Algoufily Y: 0009-0000-8945-6028
Michalopoulou F: 0009-0003-0335-3178
Papathanasiou M: 0000-0002-8886-0624
Mercangöz M: 0000-0002-4449-0414

REFERENCES

1. Xu J, Xu X, Huang C, Angelo J, Oliveira CL, Xu M, Xu X, Temel D, Ding J, Ghose S, Borys MC, Li ZJ. Biomanufacturing evolution from conventional to intensified processes for productivity improvement: a case study. mAbs 12: (2020). https://doi.org/10.1080/19420862.2020.1770669
2. Shukla AA, Wolfe LS, Mostafa SS, Norman C. Evolving trends in mab production processes.

Bioengineering & Transla Med 2:58-69 (2017). https://doi.org/10.1002/btm2.10061
3. Godawat R, Brower K, Jain S, Konstantinov K, Riske F, Warikoo V. Periodic counter?current chromatography – design and operational considerations for integrated and continuous purification of proteins. Biotechnology Journal 7:1496-1508 (2012). https://doi.org/10.1002/biot.201200068
4. Müller?Späth T, Aumann L, Melter L, Ströhlein G, Morbidelli M. Chromatographic separation of three monoclonal antibody variants using multicolumn countercurrent solvent gradient purification (MCSGP). Biotech & Bioengineering 100:1166-1177 (2008). https://doi.org/10.1002/bit.21843
5. Papathanasiou MM, Steinebach F, Morbidelli M, Mantalaris A, Pistikopoulos EN. Intelligent, model-based control towards the intensification of downstream processes. Computers & Chemical Engineering 105:173-184 (2017). https://doi.org/10.1016/j.compchemeng.2017.01.005
6. Eisenhuth R, Müller-Späth T. Process characterization and performance qualification of MCSGP. Processes 13:3950 (2025). https://doi.org/10.3390/pr13123950
7. Müller-Späth T, Ströhlein G, Aumann L, Kornmann H, Valax P, Delegrange L, Charbaut E, Baer G, Lamproye A, Jöhnck M, Schulte M, Morbidelli M. Model simulation and experimental verification of a cation-exchange igg capture step in batch and continuous chromatography. Journal of Chromatography A 1218:5195-5204 (2011). https://doi.org/10.1016/j.chroma.2011.05.103
8. Michalopoulou F, Papathanasiou MM. Accelerated process optimization of chromatographic separations using a hybrid modeling approach. IFAC-PapersOnLine 59:427-432 (2025). https://doi.org/10.1016/j.ifacol.2025.07.183
9. Subraveti SG, Li Z, Prasad V, Rajendran A. Physics-based neural networks for simulation and synthesis of cyclic adsorption processes. Ind. Eng. Chem. Res. 61:4095-4113 (2022). https://doi.org/10.1021/acs.iecr.1c04731
10. Rajendran A, Paredes G, Mazzotti M. Simulated moving bed chromatography for the separation of enantiomers. Journal of Chromatography A 1216:709-738 (2009). https://doi.org/10.1016/j.chroma.2008.10.075
11. Minceva M, Pais LS, Rodrigues AE. Cyclic steady state of simulated moving bed processes for enantiomers separation. Chemical Engineering and Processing: Process Intensification 42:93-104 (2003). https://doi.org/10.1016/s0255-2701(02)00038-7
12. Ibrahim G, Abasaeed AE. Modelling of sequencing batch reactors. Water Research 29:1761-1766 (1995). https://doi.org/10.1016/0043-1354(94)00317-z
13. Dionisi D, Rasheed AA, Majumder A. A new method to calculate the periodic steady state of sequencing batch reactors for biological wastewater treatment: model development and applications. Journal of Environmental Chemical Engineering 4:3665-3680 (2016). https://doi.org/10.1016/j.jece.2016.07.032
14. Algoufily Y, Borghesan F, Mercangöz M. Predictive modelling of desiccant drying processes using multi-feature k-nearest neighbours algorithm. IFAC-PapersOnLine 59:205-210 (2025). https://doi.org/10.1016/j.ifacol.2025.07.146
15. Daoutidis P, Lee JH, Rangarajan S, Chiang L, Gopaluni B, Schweidtmann AM, Harjunkoski I, Mercangöz M, Mesbah A, Boukouvala F, Lima FV, del Rio Chanona A, Georgakis C. Machine learning in process systems engineering: challenges and opportunities. Computers & Chemical Engineering 181:108523 (2024). https://doi.org/10.1016/j.compchemeng.2023.108523

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Enhancing Control in Chemical Processes using Reinforcement from Human Feedback

Hilde Gerold[a]*, Dean Brandner[a], and Sergio Lucia[a]

[a] Technische Universität Dortmund, Laboratory of Process Automation Systems, Emil-Figge-Straße 70, Dortmund 44227, Germany
* Corresponding Author: hilde.gerold@tu-dortmund.de.

ABSTRACT

Reinforcement learning (RL) presents a promising alternative to model-based advanced control schemes, such as model predictive control (MPC), whose application can be limited by highly complex system models. However, incorporating constraints in RL remains challenging and formulating a suitable optimization objective is not straightforward. Reinforcement learning from human feedback (RLHF) offers an approach to derive the RL reward function from human expert preferences, enabling the incorporation of process knowledge. In this work, we present the application of RLHF to fine-tune an approximate MPC controller with suboptimal performance. We demonstrate that combining conventional reward formulations with RLHF, along with varying trajectory segment lengths for collecting human feedback, improves the control methodology for a batch bioreactor by enhancing safety and accounting for long-term effects. Furthermore, direct-preference based policy optimization (DPPO) represents a promising alternative for directly fine-tuning learning-based controllers while circumventing explicit reward model design.

Keywords: Model predictive control, reinforcement learning, human feedback

INTRODUCTION

Advanced process control methods, such as model predictive control (MPC), represent a powerful approach as they enable constraint handling and are applicable to nonlinear multivariable processes. However, the high computational demand resulting from repeatedly solving the associated optimization problem makes real-time application challenging for high-dimensional and complex systems. Deep reinforcement Learning (RL) [1] presents a promising data-based approach, as the resulting control policy enables fast evaluation. This policy is learned through interactions with the system environment and guided by a scalar reward. However, formulating a suitable reward encoding the underlying control objective is not always straightforward. Additionally, constraint incorporation remains challenging, posing a problem for safety-critical process industries. Reinforcement learning from human feedback (RLHF) offers a promising method to derive the RL reward from human expert preferences [2] and is commonly applied for fine-tuning large language models [3], as well as in games [4] and robotics [5]. A detailed survey of RLHF is provided by Kaufmann et al. [6].

This work builds on the recent success of RLHF approaches and investigates different fine-tuning strategies for learning-based controllers using human feedback. First, we provide a short introduction to RL, by an overview of methodologies for incorporating human feedback into the fine-tuning process. Subsequently, we illustrate the application of these methods for improving an approximate MPC controller to control a batch bioreactor. We investigate various design parameters of the RLHF approach to improve process security and incorporation of process knowledge. Finally, we summarize our findings and discuss the observed effects.

BACKGROUND ON REINFORCEMENT LEARNING

In reinforcement learning (RL), an agent learns a mapping from states to actions by interacting with an environment and receiving feedback through a scalar reward $r \in \mathbb{R}$ [1]. Although RL employs distinct notation conventions, we adopt standard control terminology throughout this paper. Accordingly, the terms state and

observation are used interchangeably, as are input and action. The goal in RL is to learn a policy $\pi: \mathcal{X} \rightarrow \mathcal{U}$, mapping the states $x \in \mathbb{R}^{n_x}$ to optimal actions $u \in \mathbb{R}^{n_u}$. The optimization objective is defined to not only maximize the immediate rewards but the accumulated rewards for the following transitions. The RL objective therefore is formalized as:

$$\pi^* = arg \max_{\pi} \mathbb{E}_{\pi}[G] = \arg\max_{\pi} \mathbb{E}_{\pi}\left[\sum_{t=0}^{T} \gamma^t r_{k+t+1}\right]. \quad (1)$$

The optimal policy π^* maximizes the expected return G, which is defined as the accumulated reward r over an episode of length T, where the time instance is denoted by k. The discount parameter $\gamma \in (0, 1]$ reduces the influence of future rewards on the cumulative return.

BACKGROUND ON FINE-TUNING VIA HUMAN FEEDBACK

Reinforcement learning from human feedback

Designing a suitable reward function in RL that aligns with the underlying goal of the process can be challenging. To circumvent excessive reward shaping, reinforcement learning from human feedback (RLHF) incorporates human preferences to obtain a reward function estimate $\hat{r}: \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ [2]. Specifically, human feedback is obtained through comparison of two trajectory segments $\{\sigma^1, \sigma^2\}$, sampled from a set of trajectories $\{\tau^1, \ldots, \tau^i\}$. Each snippet σ consists of a sequence of state-action pairs of length L_s:

$$\sigma = ((x_0, u_0), (x_1, u_1), \ldots, (x_{L_s-1}, u_{L_s-1})) \in (\mathcal{X} \times \mathcal{U})^{L_s}. \quad (2)$$

The human expert compares both trajectories and selects the preferred one. Both segments, together with the indicated label μ, present the triple $(\sigma^1, \sigma^2, \mu) \in \mathcal{D}$, which is stored in a dataset $\mathcal{D}$. In case that the first segment is preferred over the second one, the corresponding judgmental label is defined as with $\mu(1) = \mu(\sigma^1 \succ \sigma^2) = \{1, 0\}$. If the second segment is preferred, it is marked by $\mu(2) = \mu(\sigma^2 \succ \sigma^1) = \{0, 1\}$. When no distinct preference exists, the clips are rated equally with $\mu(\sigma^1 \sim \sigma^2) = \{0.5, 0.5\}$. Following the methodology of Christiano et al. [2], it is assumed that the human preference for one segment is reasoned in a higher sum of underlying reward of the presented state-action pairs. The preference reward estimate $\hat{r}^{\Phi}$, which is typically presented by a neural network (NN) with its parameters Φ, is then trained to maximize the likelihood (MLE) of reward predictions and true estimates, formalized as:

$$\min_{\phi} - \sum_{(\sigma^1, \sigma^2, \mu) \in \mathcal{D}} \mu(1) \log \hat{P}\left[\sigma^1 \succ \sigma^2, \phi\right] + \mu(2) \log \hat{P}\left[\sigma^2 \succ \sigma^1, \phi\right], \quad (3)$$

where

$$\hat{P}[\sigma^1 \succ \sigma^2, \Phi] = \frac{\exp \sum_t \hat{r}^{\Phi}(x_t^1, u_t^1)}{\exp \sum_t \hat{r}^{\Phi}(x_t^1, u_t^1) + \exp \sum_t \hat{r}^{\Phi}(x_t^2, u_t^2)}. \quad (4)$$

The trained preference reward NN can be implemented into any arbitrary RL framework. The described procedure can be iteratively performed until the desired policy performance is achieved.

Direct preference-based policy optimization

Several challenges arise when using RLHF [7, 8]. The reward model derived from human preferences can fail to capture the true underlying reward of the task [9]. Furthermore, the quality of the learned reward $\hat{r}^{\Phi}$ can only be assessed through evaluation of the resulting fine-tuned policy. Therefore, the entire optimization process not only relies on the supervised reward training, but also on the maximization problem during policy learning [7]. As a result, direct preference-based policy optimization (DPPO) [9] has emerged as an alternative to RLHF, directly optimizing the policy parameters using the preference dataset $\mathcal{D}$ without requiring an explicit reward model. In DDPO the applied score metric quantifies the distance between policy predictions and actions of the labeled trajectory segments by any aggregation of distance d. Applying MLE to adjust the policy parameters, the training loss is defined similar to Equation (3):

$$\min_{\theta} - \sum_{(\sigma^1, \sigma^2, \mu) \in \mathcal{D}} \mu(1) \log \tilde{s}\left(\pi^{\theta}, \sigma^1 \succ \sigma^2, \lambda_{\tilde{s}}\right) + \mu(2) \log \tilde{s}\left(\pi^{\theta}, \sigma^2 \succ \sigma^1, \lambda_{\tilde{s}}\right), \quad (5)$$

where

$$\tilde{s}\left(\pi^{\theta}, \sigma^1 \succ \sigma^2\right) = \frac{\exp\left(-d(\pi^{\theta}, \sigma^1)\right)}{\exp\left(-d(\pi^{\theta}, \sigma^1)\right) + \exp\left(-\lambda_{\tilde{s}} d(\pi^{\theta}, \sigma^2)\right)}. \quad (6)$$

The regularization factor $\lambda_{\tilde{s}} \in (0, 1)$ is introduced to penalize the increase of both distances. Without this regularization, the ratio $\tilde{s}$ could stay in a similar range even when the overall distance to both segments increases [9]. A comparison of RLHF with such a direct policy optimization (DPO) [10] method has found that DPO has the potential to outperform RLHF, especially with a sufficient large preference dataset [11]. However, DPO can be sensitive to distribution shifts between the provided preference data and can lead to insufficient generalization [12]. As in RLHF the reward model might provide incorrect high rewards to unseen data, this misspecification can also happen for DPO but directly impacts the policy [12].

IMPROVEMENT OF APPROXIMATE MPC FOR BATCH BIOREACTOR

Batch bioreactor model

To investigate the fine-tuning capabilities using human feedback, we consider the batch bioreactor system

from Srinivasan et al. [13], where penicillin is produced by microorganisms. Assuming an isothermal and ideally mixed reactor, the system is described by the following set of ordinary differential equations:

$$\dot{X}_s = \mu_s(S_s)X_s - \frac{u_{\text{inp}}}{\text{V}_\text{s}}\text{X}_\text{s}, \tag{7a}$$

$$\dot{S}_s = -\frac{\mu_s(S_s)X_s}{Y_x} - \frac{\upsilon X_s}{Y_p} + \frac{u_{\text{inp}}}{V_s}(S_{\text{in}} - S_s), \tag{7b}$$

$$\dot{P}_s = \upsilon X_s - \frac{u_{\text{inp}}}{V_s}P_s, \tag{7c}$$

$$\dot{V}_s = u_{\text{inp}}, \tag{7d}$$

where

$$\mu_s(S_s) = \frac{\mu_m S_s}{K_m + S_s + \left(\frac{S_s^2}{K_i}\right)}. \tag{8}$$

The concentration of biomass producing the pharmaceutical product is denoted by X_s and the concentration of penicillin by P_s. Substrate at concentration S_s is supplied to the microorganisms via the input feed u_{inp}. The reactor volume is described by V_s. Overall, the system is defined by the four physical states $x_{\text{phys}} = (X_s, S_s, P_s, V_s)^T$ and one control input $u_{\text{phys}} = \left(u_{\text{inp}}\right)^T$. The kinetics of the bioprocess are defined by μ_m, K_m, K_i and υ. The microorganisms exhibit substrate inhibition as described by Equation (8), meaning that a higher substrate concentration leads to a decrease of the kinetic term $\mu_s(S_s)$ for biomass growth. The yield coefficients of biomass and penicillin are Y_x and Y_p, respectively. The concentration of the inlet feed is denoted by S_{in}.

The process aim is to produce as much penicillin as possible in a fixed batch time, where the concentration of biomass is restricted to $X_{s,\,\max} = 3.7$ mol L^{-1} to ensure sufficient supply of oxygen. The system can be controlled using MPC, where the corresponding stage cost is formalized by:

$$l(x_t, u_t) = -\left(P_{s,t} - \Delta u_{\text{inp},t}^2\right). \tag{9}$$

For smooth control trajectories changes of the control input are penalized, where $\Delta u_{\text{inp},t} = u_{\text{inp},t} - u_{\text{inp},t-1}$.

Design of approximate MPC

Approximate model predictive control (AMPC) circumvents the computational limitations inherent to conventional MPC by employing NN-based function approximators to learn the mapping from system states to optimal control actions [14]. This enables the deployment of predictive control strategies in time-critical applications where online optimization would be computationally prohibitive. However, the achieved accuracy highly depends on the available training data. We therefore investigate the fine-tuning capabilities using human feedback for an AMPC controller with suboptimal performance.

A MPC controller is designed for the batch bioreactor system according to the stage cost described in Equation (9) and the described constraint with a prediction horizon of 20 min**.** We apply trajectory-based sampling to obtain a dataset for approximating the optimal control law of the MPC. The input data for the trained NN $\text{x}_{\text{AMPC},t} = \left(x_{\text{phys},t}^T, u_{\text{phys},t-1}^T\right)^T \in X_{\text{AMPC}}$ is defined by the states and preceding control inputs to enable learning the control objective as defined in Equation (9). The output data $u_{\text{AMPC},t} = \left(u_{\text{phys},t}\right)^T \in Y_{\text{AMPC}}$ is defined according to the control actions. To obtain an AMPC with sufficient potential for improvement, $25,000$ data points are used to train a NN with parameters θ to minimize the following mean square error (MSE) objective:

$$\min_{\theta} \frac{1}{N}\sum_{i=1}^{N} \| u_{\text{AMPC},\,0}^{*,[i]} - \pi\left(x_{\text{AMPC},\,0}^{[i]}; \theta\right)\|_2^2. \tag{10}$$

Reward model from human preferences

We now design a reward model from human preference according to Christiano et al. [2] to reflect the intentions of a human expert. Following this approach, a set of trajectories $\{\tau^1, \dots, \tau^i\}$ is sampled to cut out trajectory segments σ out of it. We use MPC trajectories, which are available from training the AMPC, as well as closed-loop evaluations of the AMPC. Overall, multiple iterations are required for reward design until the preference dataset sufficiently covers the state space and the reward model adequately reflects the intended behavior. If the trajectory segments labeled for the preference dataset do not cover sufficiently the state space, incorrect high rewards are might be given to undesired states and control actions. Fine-tuning the policy with this inadequate reward leads to undesired control behavior and is additionally used to generate trajectories. Thus, a broad range of different control trajectories are included in the preference dataset. Different initial conditions from the feasible state space are used to generate the trajectories and the start of each segment is selected uniformly and randomly out of those trajectories. The segment length L_s is set to ten, reflecting 10 min, sufficient for evaluating the system's current state and estimating its development in this initial investigation. For labeling, 350 snippets are generated. Two segments are randomly selected from the trajectory set $\{\sigma_1, \sigma_2, \dots, \sigma_{350}\}$, and presented to a process expert, who indicates which trajectory is preferable for system control, or labels them as equally. The segment data and corresponding label are saved automatically and indistinguishable pairs are skipped.

In total $2,133$ preferences are collected for the reward dataset according to the following labeling criteria: Segments with constraint violations (CVs), where biomass concentration exceeds 3.7 mol L^{-1}, are disfavored, though if both segments exhibit CVs, they are rated equally. Similarly, substrate concentrations $S_s \geq 1$ mol L^{-1} are penalized due to substrate inhibition. Further, if one

trajectory's initial penicillin concentration exceeds the other's final concentration, that segment is preferred, as maximizing penicillin production is the primary objective. These decision rules were derived from observed human expert labeling behavior and subsequently implemented as automatic labeling criteria, thereby augmenting the manual expert annotations. Comparisons not meeting these criteria rely on human judgment. The human's expertise is especially interesting in this case study when deciding between the trade-off of higher penicillin concentration and a greater amount of biomass, which promises faster penicillin production in the future.

The collected preferences build the training data for the NN-based reward model. Based on the current observations and actions taken, a reward is obtained from the reward network. Hence, the NN is set up with six inputs $\mathrm{x}_{\hat{r},t} = \left(x_{\mathrm{phys},t}^T, u_{\mathrm{phys},t-1}^T, u_{\mathrm{phys},t}^T\right)^T$, consisting of the four states, the previous and the current control input. The output layer of the network consists of one neuron for the scalar reward $\hat{r}$. The parameters of the preference reward network Φ are trained by minimizing the RLHF objective as defined in Equation (3).

An ensemble of three separately trained NNs is used, a common approach in RLHF [2, 15] that helps counteract reward hacking [16] and overoptimization [17]. Predictions are obtained by averaging across all models.

Investigation of different fine-tuning approaches

Comparison of reward model designs

In a first investigation, the preference reward derived in the previous section is compared to a rigorous encoded reward as typically defined for RL training. Further, a hybrid model combining both approaches is evaluated. The three rewards are formalized as:

$$\hat{r}_{\mathrm{pref}}(x,u) = NN(x,u;\Phi), \tag{13a}$$

$$\hat{r}_{\mathrm{rig}}(x,u) = -l(x,u) - c(x,u), \tag{13b}$$

$$\hat{r}_{\mathrm{hybrid}}(x,u) = \hat{r}_{\mathrm{rig}}(x,u) + NN_{\mathrm{hybrid}}(x,u;\Phi). \tag{13c}$$

The preference reward $\hat{r}_{\mathrm{pref}}(x,u)$ is obtained from human feedback as described in the previous section. The rigorous reward $\hat{r}_{\mathrm{rig}}(x,u)$ is defined as the negative inverse stage cost from Equation (6.4), transforming the MPC minimization objective into an RL maximization objective. A penalization term $c(x,u)$ for constraint violations is added, weighting violations by factor 20. Both strategies are merged via a hybrid reward $\hat{r}_{\mathrm{hybrid}}$. For this hybrid reward, an NN is trained similarly to how it was described for the preference reward, using the same labeled dataset and loss function for training. However, instead of just passing the NN predictions as reward estimates for the preference calculation according to Equation (3), it is added first to the rigorous reward. The reward for the trajectory segment therefore consists of the rigorous reward penalizing CVs and the NN as an expansion. This hybrid reward is designed to extent the rigorous reward, encoding the main goal and limitations, by the expert knowledge included through the preference dataset.

Evaluating reward models derived from human preference to determine whether they suitably capture the task presents a major challenge in RLHF [7, 8]. As the dimensionality of the present control task is manageable, the reward behavior of reward models can be investigated using heatmaps. Predictions of the three applied reward models are displayed in Figure 1 and 2. Each plot displays the reward prediction dependent on two of the six inputs for reward estimation. The remaining inputs are kept constant.

Figure 1 displays the reward prediction as a function of biomass and penicillin concentration. High concentrations of penicillin are expected to be rewarded, and violations of constraints penalized with small rewards. The rigorous model meets this expectation. As long as the amount of biomass is not violating the constraint, it does not influence the rigorous reward. However, CVs are clearly reflected by the model, as the reward significantly decreases. The penalty for CVs is comparatively much higher than the reward for increasing penicillin concentration. The hybrid model is almost identical to the rigorous reward, but a small influence of the biomass concentration is noticeable. Smaller amounts of penicillin receive higher rewards if a high biomass concentration is present, showing the influence of the preference dataset. The reward predictions from the preference model display a different mapping. The distinction between higher and lower penicillin concentrations becomes more noticeable as CVs are relatively less penalized compared to other points, especially when compared to the rigorous and hybrid reward. Similar to the hybrid model, higher amounts of biomass are favored. From analyzing the dependency of biomass and penicillin concentration, the hybrid model seems to mainly align with the rigorous reward and adopt the rewarding of high biomass concentrations from the human preferences.

Figure 2 plots rewards against penicillin and substrate concentration, depicting more clearly the difference between hybrid and rigorous rewards. The rigorous reward is not influenced by the amount of substrate. A high concentration of substrate hindrances biomass growth and thereby penicillin production. It has been observed that some RL agents get stuck in local maxima, feeding a lot of substrate. This behavior is specifically included in the preference dataset by sampling trajectories from these agents. Therefore, both, preference and hybrid reward yield smaller rewards for large amounts of substrate.

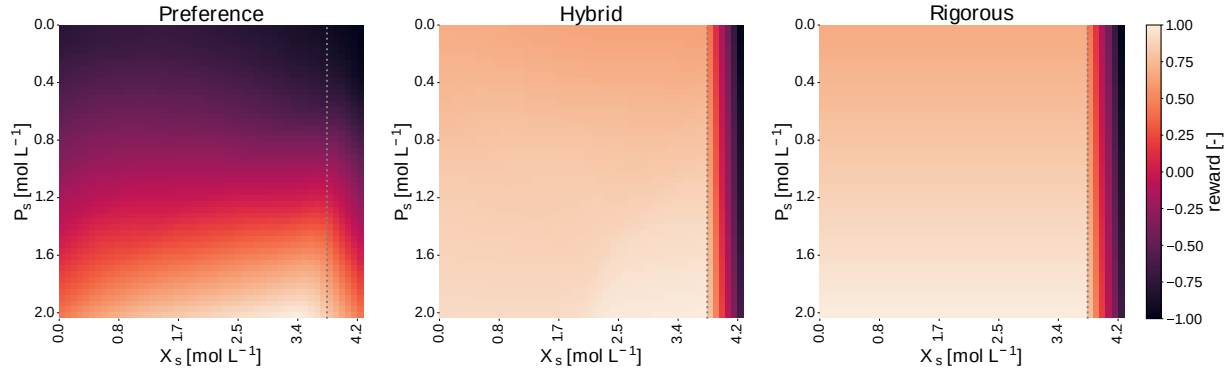

Figure 1: Evaluation of reward predictions for the preference reward model from human feedback, hybrid reward and rigorous reward. The reward is predicted in dependence of penicillin concentration P_s and biomass concentration X_s, where all other reward inputs are kept constant. Constraint indication by grey dotted line.

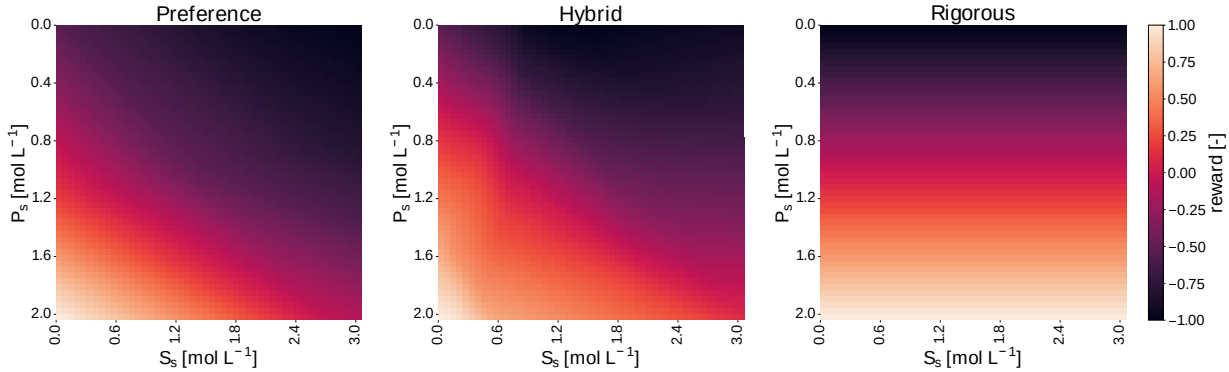

Figure 2: Evaluation of reward predictions for the preference reward model from human feedback, hybrid reward and rigorous reward. The reward is predicted in dependence of penicillin concentration P_s and substrate concentration S_s, where all other reward inputs are kept constant.

To further investigate the rewards, their fine-tuning capabilities with respect to the previously derived AMPC are investigated. The pretrained AMPC is implemented as the policy in the DDPG algorithm [19] using Stable Baselines3 [20]. For each reward model, 25 agents are learned in parallel with different seeds. Each agent is learned for $5 \cdot 10^5$ transitions, frequently evaluated and the best fine-tuned agents are saved. The closed-loop evaluations of the best three models obtained from each reward are listed in Table 1. The MPC control law π^{MPC} and the AMPC π^{AMPC} serve as benchmarks. Indices refer to the reward used for fine-tuning. To evaluate all obtained policies from the fine-tuned agents, closed-loop trajectories with a length of 150 steps for 100 initial conditions are analyzed. The average return $\bar{G}$ is defined by the undiscounted, accumulated rigorous reward and is used as the evaluation metric, along with its standard deviation. Penalization of CVs is not considered in the return. Furthermore, the number of CVs for all initial conditions is summed up as n_{CV} and the average deviation $\bar{d}_{\mathrm{CV}}$ from the constraint is computed and listed. None of the best fine-tuned agents violate the constraint in any of the closed-loop trajectories and all agents improve the average return of the AMPC. The hybrid reward leads to the highest average reward. The best agent obtained with the rigorous reward yield a slightly higher average return than the agent learned with the preference reward.

Table 1: Evaluation of fine-tuning capabilities of AMPC with different reward designs.

Reward	Policy	$\bar{G}$	CV	
			n_{CV}	$\bar{d}_{\mathrm{CV}}$
	π^{MPC}	116.16 ± 21.66	0	-
	π^{AMPC}	96.36 ± 15.97	3	0.053
Preference	$\pi^{\mathrm{AMPC}}_{\mathrm{pref}}$	106.83 ± 22.32	0	-
		106.48 ± 18.83	0	-
		103.46 ± 18.14	0	-
Rigorous	$\pi^{\mathrm{AMPC}}_{\mathrm{rig}}$	108.85 ± 19.57	0	-
		105.76 ± 21.80	0	-
		104.19 ± 17.95	0	-
Hybrid	$\pi^{\mathrm{AMPC}}_{\mathrm{hybrid}}$	110.05 ± 21.89	0	-
		104.32 ± 18.27	0	-
		100.93 ± 16.75	0	-

Influence of length of trajectory segments

The previous investigation of different models

revealed that the preference model does not clearly localize the constraint. A snippet is labeled unfavored if a CV takes place at any point. The applied labeling strategy does not differentiate between the severity of constraint violations or when they occur within the segment. Consequently, the preference model may inadequately reward states with favorable biomass concentrations for penicillin production if they are associated with segments containing any constraint violation. It is presumed that with smaller snippets, the localization of the constraint can be more clearly encoded. To investigate the influence of the length of trajectory segments, shorter trajectory snippets of $L_s = 5$ are labeled with the aim to clearly localize the constraint. Longer snippets of $L_s = 20$ are labeled to investigate the representation of long-term effects in the reward model. Furthermore, a combined dataset is compiled by sampling equally from each of the three segment lengths, with each contributing one third of the total data. Both new datasets are generated following the same procedure and triplet count as the initial dataset. Figure 3 shows heatmaps of reward dependency on biomass and penicillin concentration, illustrating the reward mapping and constraint adaptation of all models.

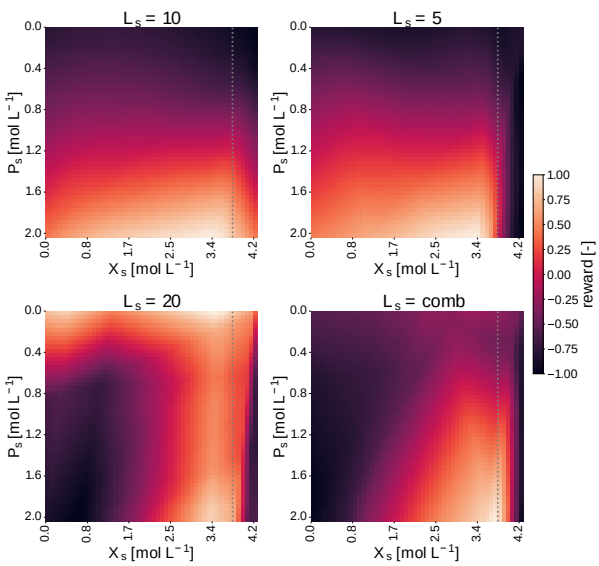

Figure 3: Evaluation of reward predictions for the preference reward models derived from different segment length. The reward is predicted in dependence of penicillin concentration P_s and biomass concentration X_s, where all other reward inputs are kept constant. Constraint indication by grey dotted line.

The model obtained from shorter trajectory length $L_s = 5$ displays the constraint restriction more clearly than the previous model from a segment length of $L_s = 10$. Slightly before reaching the maximum permitted biomass concentration, the reward decreases. Additionally, the penalization is much higher. Therefore, it is concluded that the shorter segment lengths help to identify and penalize constraints more clearly. The reward model, which is trained based on longer preference segments, also does not identify the constraints sufficiently. However, the penalization for larger violations is stronger than for the preference reward from $L_s = 10$. Furthermore, the reward from $L_s = 20$ differs from the previous models as it mainly rewards high biomass concentrations instead of product concentration. Longer segments better illustrate the development of the system and the slope of penicillin increase can be observed in more detail. Therefore, the human expert often prefers trajectories with a higher biomass concentration, as the slope of the penicillin production indicates a faster production, which is more advantageous in the long term. The incorrect high reward for very low amounts of penicillin is probably reasoned in a sparse representation of this data. The reward model obtained from the mixed dataset combines the different reward estimations. High penicillin concentrations in combination with a high amount of biomass are encouraged the most with a high reward, as a high amount of product is obtained quickly. Similar to the reward model trained with long segments, higher amounts of biomass are favored for constant amounts of product. The constraint for the biomass is clearly encoded and the reward decreases as soon as the constraint is violated.

Table 2: Evaluation of fine-tuning capabilities of AMPC with preference rewards trained with different segment lengths L_s.

Reward	Policy	$\bar{G}$	CV	
			n_{CV}	$\bar{d}_{\mathrm{CV}}$
	π^{MPC}	116.16 ± 21.66	0	-
	π^{AMPC}	96.36 ± 15.97	3	0.053
$L_s = 10$	$\pi^{\mathrm{AMPC}}_{\mathrm{pref}}$	106.83 ± 22.32	0	-
		106.48 ± 18.83	0	-
		103.46 ± 18.14	0	-
$L_s = 5$	$\pi^{\mathrm{AMPC}}_{L_s=5}$	101.68 ± 31.02	0	-
		99.20 ± 23.52	0	-
		94.11 ± 31.22	0	-
$L_s = 20$	$\pi^{\mathrm{AMPC}}_{L_s=20}$	101.44 ± 17.31	0	-
		109.21 ± 21.23	171	0.026
		111.53 ± 20.91	283	0.041
$L_s =$ comb	$\pi^{\mathrm{AMPC}}_{L_s=\mathrm{comb}}$	108.37 ± 19.45	0	-
		108.34 ± 20.86	0	-
		108.18 ± 21.94	0	-

The three best policies obtained from the different reward models are listed in Table 2, with the previous references. Agents learned with the reward from a segment length of $L_s = 5$ achieve the lowest average return, with only some exceeding AMPC performance. However, no CV occur for these agents. This supports the

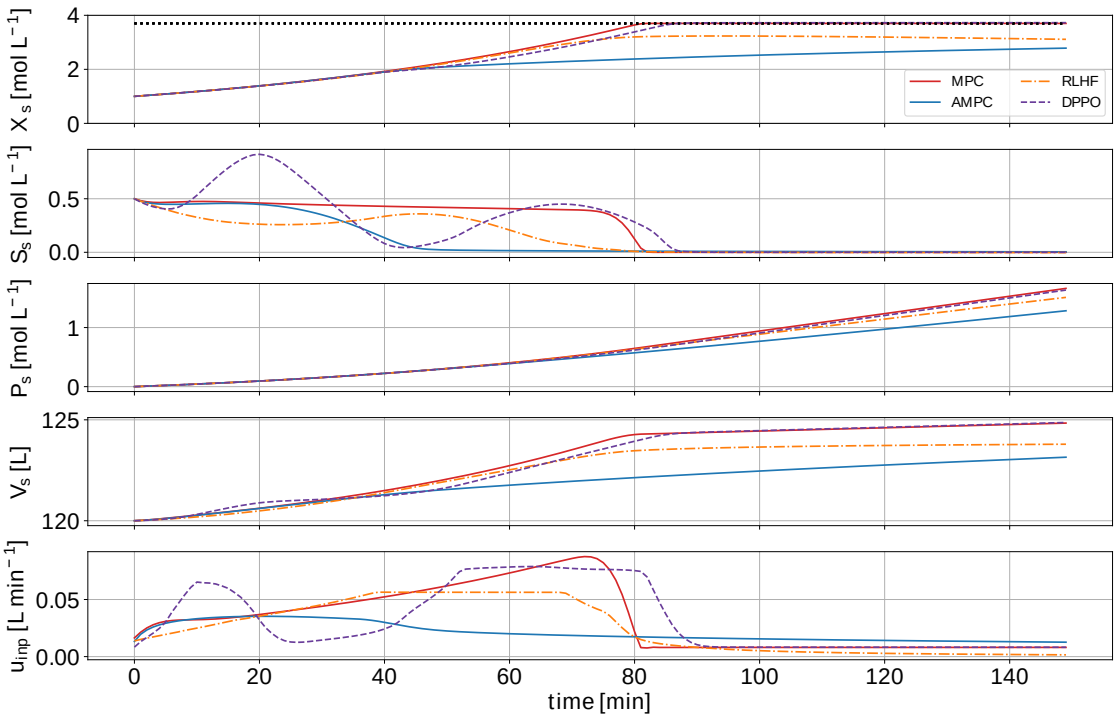

Figure 4: Closed-loop trajectories of MPC (red, solid), AMPC (blue, solid), reference policy from fine-tuned policy using RLHF (orange, dash-dot), and optimized policy using DPPO (purple, dashed). The constraint is marked by the dotted black line.

observations from the heatmap, that the constraints are recognized, but high amounts of biomass are not rewarded sufficiently. Correspondingly, the agents, learned with the reward of length $L_s = 20$, produce higher concentrations of penicillin in the same period of time under more CVs. The agents learned with the reward from the mixed dataset achieve the best results in terms of a high return without CVs, combining the strengths of all datasets.

Comparison of RLHF and DPPO

As the presented RLHF approach is prone to errors arising from the interdependence of reward model design and RL training, we investigate DPPO as an alternative approach to enable human preferences for fine-tuning. To ensure a valid comparison, the same dataset with $L_s = 10$ used for deriving the first preference reward is applied. It has been found that DPPO is sensitive to hyperparameters, especially the regularization factor $\lambda_{\bar{s}}$, which is used in the probability distribution to prevent both distances from increasing [9]. Therefore, a small grid search is performed and the best three fine-tuned agents obtained are listed together with the benchmark policies and reference policies from RLHF in Table 3.

The average return of the best policy optimized with DPPO outperforms the fine-tuned agents from RLHF in terms of the achieved return. However, the total number of CVs is very high for almost all models. The cause for that many CVs becomes evident when considering a closed-loop trajectory as shown in Figure 4. The development of biomass concentration, controlled by the policy optimized with DPPO, shows very similar behavior to the optimal MPC trajectory. It also converges to a high biomass concentration, which leads to a high product concentration. However, the DPPO policy converges to a biomass concentration which is slightly higher than the constraint, resulting in the high return in combination with many, but small CVs. As the optimized policy still presents an approximation, small approximation errors are to be expected.

Table 3: Evaluation of fine-tuning capabilities of AMPC with RLHF and DPPO.

Reward	Policy	$\bar{G}$	CV	
			n_{CV}	$\bar{d}_{CV}$
	π^{MPC}	116.16 ± 21.66	0	-
	π^{AMPC}	96.36 ± 15.97	3	0.053
RLHF	π^{AMPC}_{pref}	106.83 ± 22.32	0	-
		106.48 ± 18.83	0	-
		103.46 ± 18.14	0	-
DPPO	π^{AMPC}_{DPPO}	112.19 ± 20.79	1, 857	0.016
		111.29 ± 20.86	4, 376	0.022
		109.73 ± 19.60	1, 575	0.013

CONCLUSION

In this work, we demonstrated how to implement and enhance reinforcement learning from human feedback for controlling a biochemical batch reactor. Deriving the reinforcement learning reward from a human process expert's feedback enables encoding domain knowledge regarding both safety-critical and long-term behavior. Our case study visualized how different reward model designs and trajectory segment lengths influence the optimization objective. A hybrid reward function effectively combines formalizable constraints with intuitive expert knowledge, while varying segment lengths allows the incorporation of both short- and long-term control strategies. Finally, direct preference-based policy optimization offers a promising fine-tuning alternative that circumvents explicit reward model design.

In future research, we will investigate the potential of multi-trajectory ranking over pairwise comparisons for improving learning-based controllers.

ACKNOWLEDGEMENTS

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 466380688 – within the Priority Program “SPP 2331: Machine Learning in Chemical Engineering”.

AUTHOR IDENTIFIERS

Author ORCIDs:
Gerold H: 0009-0002-9287-0128
Brandner D: 0000-0003-1500-7064
Lucia S: 0000-0002-3347-5593

REFERENCES

1. Sutton RS, Barto AG. Reinforcement Learning: An Introduction, Second edition. The MIT Press (2018)
2. Christiano PF, et al. Deep reinforcement learning from human preferences. *Adv Neural Inf Process Syst* 30 (2017)
3. Agarwal S, Almeida D, Askell A, Christiano P, Hilton J, Jiang X, Kelton F, Leike J, Lowe R, Miller L, Mishkin P, Ouyang L, Ray A, Schulman J, Simens M, Slama K, Wainwright C, Welinder P, Wu J, Zhang C. Training language models to follow instructions with human feedback. Advances in Neural Information Processing Systems 35 :27730-27744 (2022). https://doi.org/10.52202/068431-2011
4. Palan M, Shevchuk G, Charles Landolfi N, Sadigh D. Learning reward functions by integrating human demonstrations and preferences. Robotics: Science and Systems XV : (2019). https://doi.org/10.15607/rss.2019.xv.023
5. Hejna J, Sadigh D. Inverse preference learning: preference-based RL without a reward function. Advances in Neural Information Processing Systems 36 :18806-18827 (2023). https://doi.org/10.52202/075280-0825
6. Kaufmann T, et al. A survey of reinforcement learning from human feedback. *arXiv* 2312.14925 (2023) https://doi.org/10.48550/arXiv.2312.14925
7. Casper S, et al. Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv* 2307.15217 (2023) https://doi.org/10.48550/arXiv.2307.15217
8. Lambert N, Calandra R. The alignment ceiling: Objective mismatch in reinforcement learning from human feedback. *arXiv* 2311.00168 (2023) https://doi.org/10.48550/arXiv.2311.00168
9. An G, Kim KM, Kosaka N, Lee J, Song HO, Zuo X. Direct preference-based policy optimization without reward modeling. Advances in Neural Information Processing Systems 36 :70247-70266 (2023). https://doi.org/10.52202/075280-3078
10. Ermon S, Finn C, Manning CD, Mitchell E, Rafailov R, Sharma A. Direct preference optimization: your language model is secretly a reward model. Advances in Neural Information Processing Systems 36 :53728-53741 (2023). https://doi.org/10.52202/075280-2338
11. Nika A, et al. Reward model learning vs. direct policy optimization: A comparative analysis of learning from human preferences. *arXiv* 2403.01857 (2024) https://doi.org/10.48550/arXiv.2403.01857
12. Xu S, et al. Is DPO superior to PPO for LLM alignment? A comprehensive study. *arXiv* 2404.10719 (2024) https://doi.org/10.48550/arXiv.2404.10719
13. Srinivasan B, et al. Dynamic optimization of batch processes: II. Role of measurements in handling uncertainty. *Comput Chem Eng* 27:27-44 (2003) https://doi.org/10.1016/S0098-1354(02)00117-5
14. Karg B, Lucia S. Efficient representation and approximation of model predictive control laws via deep learning. IEEE Trans. Cybern. 50:3866-3878 (2020). https://doi.org/10.1109/tcyb.2020.2999556
15. Chittepu Y, Finn C, Hejna J, Knox W, Niekum S, Park R, Rafailov R, Sikchi H. Scaling laws for reward model overoptimization in direct alignment algorithms. Advances in Neural Information Processing Systems 37 :126207-126242 (2024). https://doi.org/10.52202/079017-4009
16. Zhang S, et al. Improving reinforcement learning from human feedback with efficient reward model ensemble. *arXiv* 2401.16635 (2024) https://doi.org/10.48550/arXiv.2401.16635
17. Coste T, et al. Reward model ensembles help

mitigate overoptimization. *arXiv* 2310.02743 (2024) https://doi.org/10.48550/arXiv.2310.02743
18. Henderson P, Islam R, Bachman P, Pineau J, Precup D, Meger D. Deep reinforcement learning that matters. AAAI 32: (2018). https://doi.org/10.1609/aaai.v32i1.11694
19. Lillicrap TP, et al. Continuous control with deep reinforcement learning. *arXiv* 1509.02971 (2019) https://doi.org/10.48550/arXiv.1509.02971
20. Raffin A, et al. Stable-Baselines3: Reliable reinforcement learning implementations. *J Mach Learn Res* 22:1-8 (2021)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Open-Source Optimization Algorithm for the Simulated Moving Bed Process using CasADi

João Nunes*, Ana M. Ribeiro, Alexandre Ferreira, Diogo Rodrigues, Alírio E. Rodrigues

LSRE-LCM, ALiCE, Faculty of Engineering, University of Porto, Rua Dr. Roberto Frias, 4200-465, Porto, Portugal
* Corresponding Author: up201900588@edu.fe.up.pt.

ABSTRACT

In modern industrial systems, increasing performance requirements and sustainability constraints have intensified the need for advanced optimization methodologies capable of efficiently handling complex process models. The Simulated Moving Bed (SMB) process is a well-established technology for continuous chromatographic separations, offering high productivity and reduced solvent consumption compared to batch operations. However, its optimization is challenging due to the underlying distributed-parameter nature of the process.This work presents the development of a dynamic simulation and parameter optimization framework for the SMB process, implemented in Python using the open-source CasADi framework. The SMB model accounts for axial dispersion and mass transfer using a linear driving force formulation and is discretized in space using the method of lines, resulting in a state-space representation compatible with CasADi's numerical tools. Model accuracy was validated by reproducing a benchmark case from literature and comparing concentration profiles and product purities against results obtained using MATLAB's PDEPE solver.The proposed framework was further applied to a parameter optimization problem involving the separation of fructose and glucose. The objective was to maximize the feed velocity subject to purity and recovery constraints for fructose, with the superficial velocities in each SMB section and the switching time treated as decision variables. The resulting nonlinear programming problem was solved using a direct single-shooting approach and the interior-point optimizer IPOPT. The optimized operating conditions and performance metrics closely match published results, while achieving the solution in less than 90% of the computational time reported for a gPROMS-based implementation.

Keywords: Simulated Moving Bed, Optimization, CasADi, Partial Differential Equations, Dynamical Systems

INTRODUCTION

In modern industrial systems, the increasing complexity and performance demands of engineering processes have made advanced optimization approaches essential. Optimizing process performance is required not only to improve profitability, but also to enhance sustainability through more efficient use of energy and raw materials, as well as reduced waste and resource consumption.

Simulated Moving Bed

The simulated moving bed (SMB) process is widely used for continuous chromatographic separations, as it simulates the countercurrent flow of the fluid and solid phases to separate components based on their affinity with the stationary phase [1]. A schematic representation of the process for a four-column configuration is shown in Figure 1. The SMB process can produce high-purity products even under low selectivity conditions, while achieving higher productivity and lower eluent consumption compared to batch chromatography [2]. Representative applications include the separation of chiral molecules [1], such as in the production of formoterol, an antiasthmatic drug [3], functional isomers, such as fructose and glucose [1, 4], and ionic molecules, such as in the production of lysine, an essential amino acid [5].

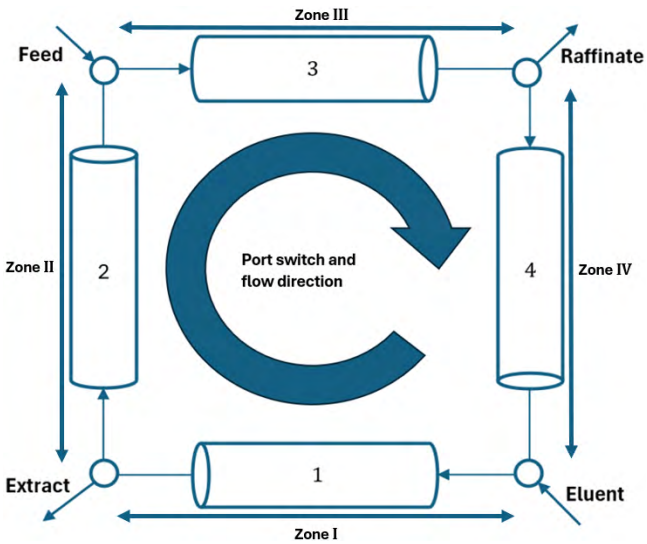

Figure 1. Four-column SMB configuration.

CasADi

CasADi is an open-source software toolbox available for Python that provides automatic differentiation and numerical optimization capabilities for complex dynamical systems described by ordinary differential equations (ODEs) and subject to constraints. The package is user-friendly and has been successfully applied to a wide range of industrial and academic problems, including optimal control of batch chromatographic separation processes [6], modeling of highly complex dynamical systems [7], and optimal control of non-isothermal continuously stirred tank reactors [8].

This work presents the development of a dynamic simulation and optimization framework for the SMB process, implemented in Python using the CasADi framework. First, a detailed simulation model is developed and validated against reference results from the literature. Subsequently, the proposed framework is applied to reproduce and solve a parameter optimization case study reported in the scientific literature.

SMB COMPUTATIONAL MODELING

The case study considered in this work is the chromatographic resolution of bi-naphthol enantiomers, as reported in the book *Simulated Moving Bed Technology: Principles, Design and Process Applications* [1]. The adsorption equilibrium isotherms, measured at 25 °C, follow a bi-Langmuir model and are presented below. For notational simplicity, the explicit dependence on time is omitted in the subsequent equations, although all state variables remain time-dependent.

Isotherm of component A

$$q^*_{A,k}(x) = \frac{2.69\, C_{A,k}(x)}{1+0.0336\, C_{A,k}(x)+0.0446\, C_{B,k}(x)} + \frac{0.1\, C_{A,k}(x)}{1+C_{A,k}(x)+3\, C_{B,k}(x)} \tag{1}$$

Isotherm of component B

$$q^*_{B,k}(x) = \frac{3.73\, C_{B,k}(x)}{1+0.0336\, C_{A,k}(x)+0.0446\, C_{B,k}(x)} + \frac{0.3\, C_{B,k}(x)}{1+C_{A,k}(x)+3\, C_{B,k}(x)} \tag{2}$$

The mathematical model employed considers axial dispersion and mass transfer between fluid and solid described by the linear driving force (LDF) model. Radial dispersion and pressure effects are neglected, and the system is assumed to be isothermal at 25 °C, ensuring the validity of the adsorption isotherms. Under these assumptions, the model equations are given by:

Fluid-phase mass balance

$$\frac{\partial C_{i,k}(x)}{\partial \theta} = \gamma_k\left(\frac{1}{Pe_k}\frac{\partial^2 C_{i,k}(x)}{\partial x^2} - \frac{\partial C_{i,k}(x)}{\partial x}\right) - \frac{(1-\varepsilon)}{\varepsilon}\frac{\partial q_{i,k}(x)}{\partial \theta} \tag{3}$$

Solid-phase mass balance

$$\frac{\partial q_{i,k}(x)}{\partial \theta} = \alpha_i\left(q^*_{i,k}(x) - q_{i,k}(x)\right) \tag{4}$$

Boundary conditions

$$C_{in,i,k} = C_{i,k}(0) - \frac{1}{Pe_k}\frac{\partial C_{i,k}(0)}{\partial x} \tag{5}$$

$$\frac{\partial C_{i,k}(1)}{\partial x} = 0 \tag{6}$$

Initial conditions

$$C_{i,k}(x) = 0;\ \theta = 0;\ \forall x \in [0,1] \tag{7}$$

$$q_{i,k}(x) = 0;\ \theta = 0;\ \forall x \in [0,1] \tag{8}$$

Parameters

$$\gamma_m = \frac{u_m t^*}{L} \tag{9}$$

$$\alpha_i = k_{h,i} t^* \tag{10}$$

$$Pe_m = \frac{u_m L}{D_{ax,m}} \tag{11}$$

$$A = \frac{\pi D^2}{4} \tag{12}$$

$$u_m = \frac{Q_m}{\varepsilon A} \tag{13}$$

Feed inlet port

$$C_{in,i,N_F} = \frac{u_F}{u_{III}} C_{F,i} + \frac{u_{II}}{u_{III}} C_{i,N_F-1}(1) \tag{14}$$

$$u_F = u_{III} - u_{II} \tag{15}$$

Eluent inlet port

$$C_{in,i,N_E} = \frac{u_{IV}}{u_I} C_{i,N_E-1}(1) \tag{16}$$

$$u_E = u_I - u_{IV} \tag{17}$$

Extract outlet port

$$C_{X,i} = C_{i,N_X-1}(1) \tag{18}$$

$$u_X = u_I - u_{II} \tag{19}$$

Raffinate outlet port

$$C_{R,i} = C_{i,N_R-1}(1) \tag{20}$$

$$u_R = u_{III} - u_{IV} \tag{21}$$

Internal nodes

$$C_{in,i,k} = C_{i,k-1}(1);\ \forall k = 1, \dots, N_T;\ k \neq N_F, N_E \tag{22}$$

Performance parameters

$$Purity_{X,B} = \frac{\int_\theta^{\theta+1} u_X C_{X,B}(1) d\theta}{\sum_{i=A,B} \int_\theta^{\theta+1} u_X C_{X,i}(1) d\theta} \tag{23}$$

$$Purity_{R,A} = \frac{\int_\theta^{\theta+1} u_R C_{R,A}(1) d\theta}{\sum_{i=A,B} \int_\theta^{\theta+1} u_R C_{R,i}(1) d\theta} \tag{24}$$

Here, $C_{i,k}(x)$ and $q_{i,k}(x)$ denote the system state variables, namely the bulk-phase and adsorbed-phase concentrations, respectively. The indices i and k refer to the component and column number. The variables x and θ represent the dimensionless axial coordinate and dimensionless time, defined by normalization with the column length L and the switching time t^*, respectively. The equilibrium adsorbed concentration is denoted by $q^*_{i,k}(x)$ and is obtained from the isotherm relations.

The inlet concentration of component i to column k is denoted by $C_{in,i,k}$, while $C_{F,i}$ represents the feed concentration of component i. The interstitial fluid velocity in section m (I–IV) of the SMB is denoted by u_m, and u_E, u_X, u_R, and u_F correspond to the eluent, extract, raffinate, and feed velocities, respectively.

The parameter α_i denotes the number of mass transfer units, where $k_{h,i}$ is the mass transfer coefficient of the LDF model. The column cross-sectional area is denoted by A, calculated from the column diameter D. The Peclet number in section m, Pe_m, is defined using the axial dispersion coefficient $D_{ax,m}$, while γ_m represents the ratio between fluid and equivalent solid interstitial velocities in section m. The bed porosity is denoted by ε, and Q_m is the volumetric flow rate in section m of the SMB.

Finally, N_T denotes the total number of columns, while N_E, N_X, N_R, N_F denote the column indices corresponding to the eluent, extract, raffinate, and feed ports, respectively. Depending on the location of the ports and sections that results from port switch, the Peclet number Pe_k and the interstitial velocity ratio γ_k in column k are computed accordingly from $Pe_I, \dots, Pe_{IV}$ and $\gamma_I, \dots, \gamma_{IV}$, respectively. The purity is defined as $Purity_{X,B}$, representing the purity of component B in the extract stream with outlet concentration $C_{X,B}$, and $Purity_{R,A}$, representing the purity of component A in the raffinate stream with outlet concentration $C_{R,A}$.

A key challenge in this work is that the SMB process is described by partial differential equations (PDEs), whereas CasADi can only handle ODEs. To address this limitation, the structure of MATLAB's PDEPE solver was adapted to the CasADi framework. PDEPE is a numerical solver for one-dimensional, time-dependent parabolic PDEs. If additional spatial dimensions were required, an alternative modeling and numerical approach would be necessary.

The solver applies the method of lines, discretizing the spatial domain and transforming the PDEs into a system of time-dependent ODEs. The resulting equations, functions and isotherms are presented below.

$$\bar{C}_{i,j,k} = 0.5\left(C_{i,j+1,k} + C_{i,j,k}\right);\ 1 \le j \le n-1 \tag{25}$$

$$\bar{q}_{i,j,k} = 0.5\left(q_{i,j+1,k} + q_{i,j,k}\right);\ 1 \le j \le n-1 \tag{26}$$

Isotherms

$$\bar{q}^*_{A,j,k} = \frac{2.69\,\bar{C}_{A,j,k}}{1+0.0336\,\bar{C}_{A,j,k}+0.0446\,\bar{C}_{B,j,k}} + \frac{0.1\,\bar{C}_{A,j,k}}{1+\bar{C}_{A,j,k}+3\,\bar{C}_{B,j,k}} \tag{27}$$

$$\bar{q}^*_{B,j,k} = \frac{3.73\,\bar{C}_{B,j,k}}{1+0.0336\,\bar{C}_{A,j,k}+0.0446\,\bar{C}_{B,j,k}} + \frac{0.3\,\bar{C}_{B,j,k}}{1+\bar{C}_{A,j,k}+3\,\bar{C}_{B,j,k}} \tag{28}$$

Mass balance equations

$$\frac{\partial C_{i,j,k}}{\partial \theta} = \frac{f^{L,C}_{i,j,k} - f^{R,C}_{i,j,k}}{\Delta x} + 0.5\left(s^{R,C}_{i,j,k} + s^{L,C}_{i,j,k}\right);\ 2 \le j \le n-1 \tag{29}$$

$$\frac{\partial q_{i,j,k}}{\partial \theta} = 0.5\left(s^{R,q}_{i,j,k} + s^{L,q}_{i,j,k}\right);\ 2 \le j \le n-1 \tag{30}$$

$$f^{R,C}_{i,j,k} = -\gamma_k\left(\frac{1}{Pe_k}\frac{C_{i,j+1,k} - C_{i,j,k}}{\Delta x} - \bar{C}_{i,j,k}\right);\ 1 \le j \le n-1 \tag{31}$$

$$s^{R,C}_{i,j,k} = -\frac{(1-\varepsilon)}{\varepsilon}\alpha_i\left(\bar{q}^*_{i,j,k} - \bar{q}_{i,j,k}\right);\ 1 \le j \le n-1 \tag{32}$$

$$s^{R,q}_{i,j,k} = \alpha_i\left(\bar{q}^*_{i,j,k} - \bar{q}_{i,j,k}\right);\ 1 \le j \le n-1 \tag{33}$$

$$f^{L,C}_{i,j,k} = f^{R,C}_{i,j-1,k};\ 2 \le j \le n \tag{34}$$

$$s^{L,C}_{i,j,k} = s^{R,C}_{i,j-1,k};\ 2 \le j \le n \tag{35}$$

$$s^{L,q}_{i,j,k} = s^{R,q}_{i,j-1,k};\ 2 \le j \le n \tag{36}$$

Boundary conditions

$$\frac{\partial C_{i,1,k}}{\partial \theta} = \frac{f^{L,C}_{i,1,k} - f^{R,C}_{i,1,k}}{0.5\Delta x} + s^{R,C}_{i,1,k} \tag{37}$$

$$f^{L,C}_{i,1,k} = \gamma_k C_{in,i,k} \tag{38}$$

$$\frac{\partial q_{i,1,k}}{\partial \theta} = s^{R,q}_{i,1,k} \tag{39}$$

$$\frac{\partial C_{i,n,k}}{\partial \theta} = \frac{f^{L,C}_{i,n,k} - f^{R,C}_{i,n,k}}{0.5\Delta x} + s^{L,C}_{i,n,k} \tag{40}$$

$$f^{R,C}_{i,n,k} = \gamma_k C_{i,n,k} \tag{41}$$

$$\frac{\partial q_{i,n,k}}{\partial \theta} = s^{L,q}_{i,n,k} \tag{42}$$

Boolean switching functions

$$\gamma_k = \sum_{m=I}^{IV} H_{m,k}\gamma_m \tag{43}$$

$$Pe_k = \sum_{m=I}^{IV} H_{m,k} Pe_m \tag{44}$$

$$C_{i,k}^{F} = H_k^F\left(\frac{u_F}{u_{III}} C_{F,i} + \frac{u_{II}}{u_{III}} C_{i,n,k-1}\right) \quad (45)$$

$$C_{i,k}^{E} = H_k^E\left(\frac{u_{IV}}{u_I} C_{i,n,k-1}\right) \quad (46)$$

$$C_{in,i,k} = C_{i,k}^{F} + C_{i,k}^{E} + \left(1 - H_k^F - H_k^E\right) C_{i,n,k-1} \quad (47)$$

$$C_{X,i} = \sum_{k=1}^{N_T} H_k^X C_{i,n,k-1} \quad (48)$$

$$C_{R,i} = \sum_{k=1}^{N_T} H_k^R C_{i,n,k-1} \quad (49)$$

The variables $\bar{q}_{i,j,k}, \bar{C}_{i,j,k}$, and $\bar{q}^*_{i,j,k}$ represent the average values of the adsorbed concentration, bulk-phase concentration, and equilibrium adsorbed concentration, respectively, as defined by the employed discretization scheme. The index j denotes the spatial discretization point along the column axis and ranges from 1 to n, where n is the total number of spatial nodes per column. The terms $f_{i,j,k}^{R,C}$ correspond to convective–dispersive fluxes in the bulk-phase mass balance, while $s_{i,j,k}^{R,C}$ and $s_{i,j,k}^{R,q}$ denote source terms associated with mass transfer between the fluid and solid phases. The spatial discretization step is given by Δx.

The Boolean variables $H_{m,k}, H_k^F, H_k^E, H_k^X$, and H_k^R are used to represent the port-switching logic of the SMB. These variables take binary values and activate the appropriate terms depending on the column position and section configuration. The parameter γ_k denotes the column-dependent velocity ratio, obtained as a combination of the section-dependent γ_m values through the Boolean mapping $H_{m,k}$. The inlet concentration $C_{in,i,k}$ accounts for all possible inlet contributions to column k, including feed, eluent, and internal node streams. The auxiliary variables $C_{i,k}^F$ and $C_{i,k}^E$ represent the contributions of the feed and eluent streams, respectively, with the corresponding Boolean variables ensuring that only the active one is considered at each switching configuration. Finally, the extract and raffinate outlet concentrations are computed as a sum of the columns' outlet concentrations, where the Boolean variables ensure that only the column connected to the corresponding port contributes to each stream.

Once the model formulation was completed, simulation results obtained with CasADi were validated against those generated using MATLAB's PDEPE solver, which has been previously validated for the simulation of SMB systems [1]. The model parameters and operating conditions used in the simulations are summarized in Table 1, while the integrator settings adopted in this work are reported in Table 2. All remaining integrator settings were kept at their default values, and a configuration of three columns per section was considered.

The concentration profiles of species A (less-retained) and B (more-retained), evaluated at cyclic steady state at the end of a switching period, together with the corresponding product purities, are compared with the reference results in Figure 2 and Table 3.

Table 1: Model parameters and operating conditions for the simulation

Parameter	Value
L / cm	7
Pe_m	666.67
D / cm	2.6
$(1-\varepsilon)/\varepsilon$	1.5
$k_{h,A}$ / min^{-1}	6
$k_{h,B}$ / min^{-1}	6
Q_I / cm^3 • min^{-1}	56.83
Q_{II} / cm^3 • min^{-1}	38.85
Q_{III} / cm^3 • min^{-1}	42.49
Q_{IV} / cm^3 • min^{-1}	35.38
t^* / min	2
$C_{F,A}$ / mg • cm^{-3}	2.9
$C_{F,B}$ / mg • cm^{-3}	2.9

Table 2: Integrator settings for the simulation.

Setting	
Time integrator	CVodes
Time step size	10^{-2}
Spatial step size	$2 \cdot 10^{-2}$
Relative tolerance (CVodes)	10^{-4}
Absolute tolerance (CVodes)	10^{-6}

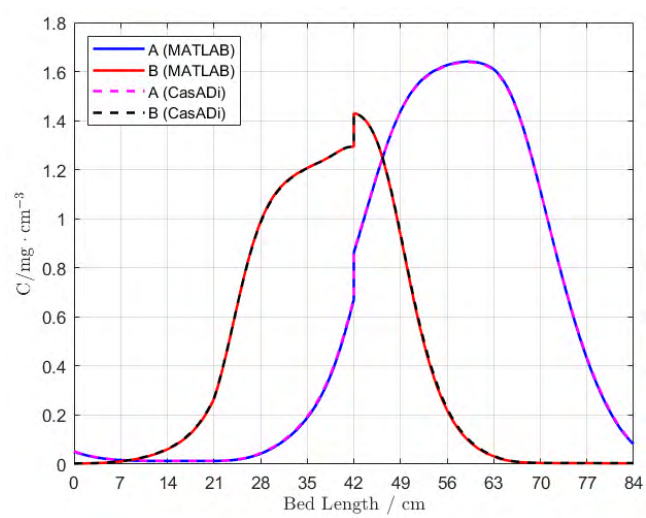

Figure 2. MATLAB and CasADi's concentration profiles at the end of a cycle.

Table 3: Product purity.

Performance parameter	MATLAB	CasADi
A purity in the raffinate / %	99.1	99.1
B purity in the extract / %	96.8	96.7

The profiles obtained with CasADi closely match those from MATLAB's PDEPE solver. Similarly, the products purities reported in Table 3 are in excellent agreement, with a maximum deviation of 0.1%, which is consistent with differences expected from the use of distinct numerical integration methods. These results confirm that the CasADi-based extension of the PDEPE approach is functioning correctly.

SMB OPTIMIZATION PROBLEM

The parameter optimization problem considered in this work focuses on the separation of fructose and glucose. The main objective was to maximize the feed velocity while satisfying terminal cycle purity and recovery constraints for the primary product (fructose) in the extract. The optimization variables included the internal superficial fluid velocities in zones II, III and IV of the SMB and the switching time between steps [9].

It should be noted that, while the overall model structure remains unchanged, the adsorbent mass balance in this problem follows a liquid-phase-based LDF formulation, consistent with the isotherm parameters provided in the reference case. In this approach, the mass transfer coefficient is expressed in terms of fluid-phase concentrations, and the adsorbed concentration is obtained through a liquid-based isotherm relation. Additionally, axial dispersion effects are implicitly accounted for within the mass transfer formulation, and therefore no explicit axial dispersion term or Peclet number appears in the model equations.

Based on the validated model and the formulation described above, the dynamic optimization problem was then defined. The objective function, decision variables, purity and recovery constraints are presented below.

$$\max_{u_{II},u_{III},u_{IV},t^*} u_F \tag{47}$$

$$Purity_{X,f} = \frac{\int_{\theta}^{\theta+1} u_X C_{X,f} d\theta}{\sum_{i=f,g} \int_{\theta}^{\theta+1} u_X C_{X,i} d\theta} \geq 90\% \tag{48}$$

$$Recovery_{X,f} = \frac{\int_{\theta}^{\theta+1} u_X C_{X,f} d\theta}{\int_{\theta}^{\theta+1} u_F C_{F,f} d\theta} \geq 80\% \tag{49}$$

In addition, non-negativity constraints were imposed on all inlet and outlet velocities to ensure physically meaningful operation and to prevent flow reversal during the optimization.

The model parameters, decision variable bounds, and initial guesses used in the optimization are summarized in Tables 4, 5 and 6. The numerical integrator settings employed in the optimization, including time and spatial discretization and solver tolerances, are reported in Table 7.

The adsorption equilibrium for each component is described by linear isotherms. The parameters K_g and K_f denote the equilibrium adsorption coefficients for glucose and fructose, respectively. These coefficients are assumed to be constant and were taken directly from the reference case.

Table 4: Model parameters and operating conditions for the optimization problem

Parameter	Value
L / m	2
ε	0.389
$k_{h,f}$ / s^{-1}	$6.84 \cdot 10^{-3}$
$k_{h,g}$ / s^{-1}	$6.84 \cdot 10^{-3}$
K_g	0.518
K_f	0.743
$C_{F,f}$ / %	50
$C_{F,g}$ / %	50
u_I / m • h^{-1}	8
Cycles	60
Columns per section	1

Table 5: Decision variables' initial guesses

Decision variable	Initial guess
u_{II} / m • h^{-1}	5
u_{III} / m • h^{-1}	5.1
u_{IV} / m • h^{-1}	3.5
t^* / s	1140

Table 6: Decision variables' upper and lower bounds

Decision variable	Lower bound	Upper bound
u_{II} / m • h^{-1}	0.1	8
u_{III} / m • h^{-1}	0.1	8
u_{IV} / m • h^{-1}	0.1	8
t^* / s	0	-

Table 7: Integrator settings for the optimization problem

Setting	
Time integrator	CVodes
Time step size	1
Spatial step size	$2 \cdot 10^{-2}$
Relative tolerance (CVodes)	10^{-4}
Absolute tolerance (CVodes)	10^{-4}

The parameter optimization problem was transcribed using a direct single-shooting approach and solved using the interior-point optimizer IPOPT. The optimal operating conditions and corresponding performance metrics obtained are summarized in Table 8, together with the associated constraint values and the reference results reported in the original article [9]. The species concentration profiles at cyclic steady state, evaluated at the end of a switching period, are shown in Figure 3. All computations were performed on a Windows 11

workstation equipped with an Intel® Core™ i7 processor and 16 GB of RAM, using Python 3.12.3 and CasADi version 3.7.2.

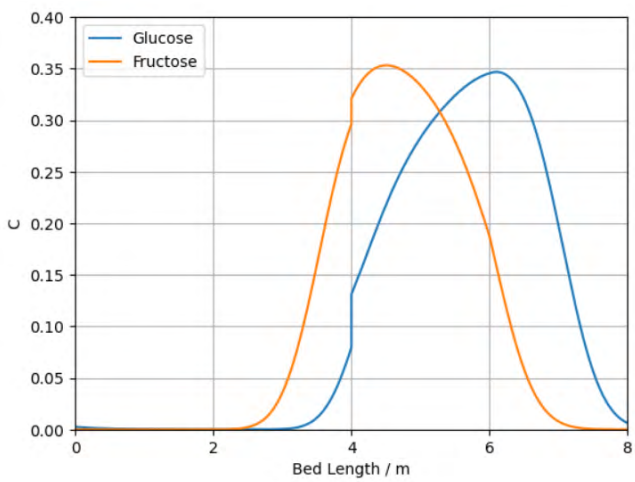

Figure 3. Fructose and glucose concentration profiles at the end of a cycle.

Table 8: Optimization results and reference values

	CasADi IPOPT	gPROMS SRQPD
u_F / $\mathrm{m \bullet h^{-1}}$	0.5216	0.5216
u_{II} / $\mathrm{m \bullet h^{-1}}$	3.7345	3.7361
u_{III} / $\mathrm{m \bullet h^{-1}}$	4.2561	4.2577
u_{IV} / $\mathrm{m \bullet h^{-1}}$	3.1298	3.1259
t^* / s	1416	1415
Fructose purity / %	89.98	90
Fructose recovery / %	79.99	80
Optimization time / min	10.7	111.8
Number of iterations	21	39

The parameter optimization results obtained with the proposed CasADi-based framework are in very good agreement with those reported in the reference study. Both approaches converged to the same objective function value, while small differences were observed in the individual decision variables, which can be attributed to the use of different numerical integrators, solvers and tolerance settings. The purity and recovery constraints for fructose in the extract were satisfied to within 0.02% of the specified limits. These minor deviations result from the looser integration tolerances and reduced number of integration steps adopted in the formulation of this optimization problem, reflecting a deliberate trade-off between computational speed and numerical precision. Tighter tolerances and finer temporal resolution could be employed if stricter constraint satisfaction were required.

Despite these minor differences, the overall performance of the optimized solution remains consistent with the reference case. Moreover, the total optimization time was significantly reduced, requiring 10.7 minutes compared to 111.8 minutes reported for the gPROMS-based implementation. This corresponds to a reduction of approximately 90% in computational time, highlighting the efficiency of the CasADi-based approach for parameter optimization of the SMB system.

It should be noted that the trade-off between computational speed and numerical precision was explored only for the CasADi implementation, while the gPROMS results correspond to those reported in the reference study.

A potential limitation of this proposed framework is related to usability. Commercial modelling environments such as gPROMS or MATLAB provide highly integrated modelling interfaces, built-in generic solvers for ODEs and PDEs, and graphical tools that simplify model implementation and debugging. In contrast, Python-based implementations using CasADi typically require a greater level of programming expertise and manual configuration of the simulation and optimization workflow. Also, in the case of models described by systems of PDEs such as the SMB process, the model must be transformed into a system of ODEs via spatial discretization to enable the use of the CasADi framework.

CONCLUSION AND FUTURE WORK

The proposed CasADi-based framework provides a robust alternative to traditional iterative optimization techniques, such as the triangle theory, for the analysis and optimization of Simulated Moving Bed processes. The developed model is highly flexible and can be readily extended to other tubular flow systems governed by partial differential equations, including plug flow reactors and pressure swing adsorption columns. Moreover, the implementation relies exclusively on open-source tools, enabling the simulation and optimization of complex process systems without the need for commercial software such as gPROMS or MATLAB, thus improving accessibility for process systems engineering applications.

Future work will focus on extending the proposed state-space formulation toward dynamic and real-time optimization of PDE-governed systems, allowing operation under optimal conditions even in the presence of plant-model mismatch.

ACKNOWLEDGEMENTS

This paper is a result of the project “WAVES - Seawater & Wastewater: Turning Challenges into Resources”, with operation code at the Funds Platform (Balcão dos Fundos), NORTE2030-FEDER-02716600, co-financed by the European Union through the NORTE 2030 Regional Program, of Portugal 2030. This research was also supported by Fundação para a Ciência e a

Tecnologia, I.P./MECI through national funds: LSRE-LCM, UID/50020/2025 (https://doi.org/10.54499/UID/50020/2025); and ALiCE, LA/P/0045/2020 (https://doi.org/10.54499/LA/P/0045/2020).

AUTHOR IDENTIFIERS

Author ORCIDs:
Nunes J: 0009-0008-6504-9469
Ribeiro AM: 0000-0003-4269-1420
Ferreira A: 0000-0002-6746-8973
Rodrigues D: 0000-0001-7823-2993
Rodrigues AE: 0000-0002-0715-4761

REFERENCES

1. Rodrigues AE, Pereira C, Minceva M, Pais LS, Ribeiro AM, Ribeiro A, Silva M, Graça N, Santos JC. Principles of simulated moving bed. Simulated Moving Bed Technology :1-30 (2015). https://doi.org/10.1016/b978-0-12-802024-1.00001-x
2. Kim KM, Lee JW, Kim S, Santos da Silva FV, Seidel-Morgenstern A, Lee CH. Advanced operating strategies to extend the applications of simulated moving bed chromatography. Chem Eng & Technol 40:2163-2178 (2017). https://doi.org/10.1002/ceat.201700206
3. Juza M, Mazzotti M, Morbidelli M. Simulated moving-bed chromatography and its application to chirotechnology. Trends in Biotechnology 18:108-118 (2000). https://doi.org/10.1016/s0167-7799(99)01419-5
4. Ruthven DM, Ching CB. Counter-current and simulated counter-current adsorption separation processes. Chemical Engineering Science 44:1011-1038 (1989). https://doi.org/10.1016/0009-2509(89)87002-2
5. Blehaut J, Nicoud RM. Recent aspects in simulated moving bed. Analusis 26:60-70 (1998). https://doi.org/10.1051/analusis:199826070060
6. Holmqvist A, Magnusson F. Open-loop optimal control of batch chromatographic separation processes using direct collocation. Journal of Process Control 46:55-74 (2016). https://doi.org/10.1016/j.jprocont.2016.08.002
7. Andersson J, Åkesson J, Diehl M. Dynamic optimization with CasADi. IEEE CDC (2012) https://doi.org/10.3384/ecp11063218
8. Andersson J, Åkessona J, Casellad F, Diehl M. Integration of casadi and jmodelica.org. Linköping Electronic Conference Proceedings 63:218-231 (2011). https://doi.org/10.3384/ecp11063218
9. Kawajiri Y, Biegler LT. Optimization strategies for simulated moving bed and powerfeed processes. AIChE Journal 52:1343-1350 (2005). https://doi.org/10.1002/aic.10736

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Connecting the Dots: A Graph-based Approach for Unsupervised Learning and Adaptive Process Monitoring with LLM-assisted Fault Diagnosis

Kyle Territo[a], Jose Romagnoli[a*]

[a] Louisiana State University, Department of Chemical Engineering, Baton Rouge, LA, United States
* Corresponding Author: jose@lsu.edu.

ABSTRACT

The convergence of artificial intelligence (AI) and chemical process systems engineering is creating unprecedented opportunities to transform current refineries from conventionally operated plants into intelligent, automated, and resilient systems. However, the practical deployment of AI in these complex industrial environments faces several critical challenges. First, most existing process datasets contain minimal labeled data, making it difficult to apply supervised learning techniques that require extensive annotations to generate meaningful insights. Furthermore, refinery data typically consists of high-dimensional, multivariate time series, which pose additional complexities in capturing temporal dynamics and system interactions. Traditional Fault Detection and Diagnosis (FDD) frameworks often struggle to address these complexities, lacking adaptability to evolving process conditions, scalability to large plant networks, and explainability in their diagnostic reasoning. To address these gaps, we propose an integrated framework for unsupervised knowledge discovery, adaptive process monitoring, and AI-assisted fault diagnosis. The approach combines Dynamic Time Warping (DTW) with graph-based layout algorithms to extract and visualize temporal structures and similarities within complex process datasets. The resulting graphs function as an unsupervised model, enabling the identification of operational regimes through process clustering. These models can be deployed online to detect anomalies using adaptive fault thresholding, ensuring robustness to changing process dynamics. When faults are detected, the framework leverages large language models (LLMs) as AI assistants to support fault diagnosis. By bridging domain knowledge—such as piping and instrumentation diagrams (P&IDs) and signal flow diagrams—with data-driven insights, including machine learning metrics and feature contribution scores, the system delivers explainable and accurate diagnostic outputs. This hybrid methodology enhances interpretability and decision-making, offering a scalable and adaptive solution for next-generation AI-enabled refinery operations.

Keywords: Unsupervised Learning, Fault Detection, LLMs, Machine Learning, Graph Networks

INTRODUCTION

Background and Motivation

Chemical refineries are among the most complex and highly automated industrial environments in modern manufacturing. Over several decades, these facilities have adopted advanced process control (APC) technologies—such as proportional–integral–derivative (PID) controllers, multivariable control systems, and model predictive control (MPC)—to maintain operational efficiency, stability, and safety under tightly constrained process conditions. As a result, today's refineries operate with extensive digitization, high sensor density, and continuous streams of multivariate data originating from interconnected unit operations [1]. Despite this progress, current refinery optimization and monitoring practices still rely heavily on deterministic models, periodic calibrations, and human expertise to interpret process behavior and respond to abnormal conditions. These systems, while robust and well-established, show diminishing

returns as they approach their performance limits, motivating the need for more adaptive and intelligent data-driven solutions [1].

The emergence of Industry 4.0 has accelerated interest in leveraging artificial intelligence (AI), machine learning (ML), and advanced analytics to complement or extend the capabilities of conventional refinery control architectures. AI-driven approaches promise improved performance for nonlinear, high-dimensional, and dynamic systems by uncovering patterns not easily expressed in first-principles or optimization-based models. However, the practical implementation of AI in refinery environments remains challenging. Most downstream oil and gas facilities generate large quantities of operational data but very limited labeled data, restricting the use of supervised learning techniques that dominate the ML landscape [1]. Moreover, the datasets are typically highly-dimensional, multivariate time series, often with irregular sampling, complex cross-correlations, and contain evolving process dynamics such as catalyst degradation, feedstock variability, and equipment aging. These characteristics complicate the training of predictive models and reduce their robustness when deployed in real-world operations.

Prior work in industrial digitalization has shown that while Industry 4.0 technologies—such as IoT sensing, big data infrastructures, and cyber-physical systems—can provide transformative value, they must be carefully adapted for continuous process industries such as refining [1]. Traditional machine learning techniques, often developed for domains with abundant annotations and static data distributions often struggle with poor scalability to plant-wide networks, inadequate handling of non-stationary behavior, difficulty capturing temporal relationships, and a lack of explainability in diagnostic outcomes. These shortcomings contribute to slow adoption rates of AI-based FDD systems in safety-critical industrial environments.

At the same time, broader advancements in artificial intelligence highlight both the urgency and the opportunity for integrating more sophisticated AI techniques into process systems engineering. The 2025 Artificial Intelligence Index Report indicates that AI models are rapidly improving in efficiency, accessibility, and reasoning capabilities, with significant adoption growth across industrial sectors [3]. Industries worldwide are increasingly leveraging AI for predictive maintenance, anomaly detection, and decision support, yet concerns about data quality, transparency, and reliability remain prevalent. These concerns are particularly acute in refinery operations, where process deviations can propagate quickly and incur substantial economic or safety consequences.

Finally, sustainability pressures—including energy efficiency demands, emission reduction targets, and global competitiveness—further amplify the need for intelligent, adaptive process monitoring solutions. As highlighted by Ghobakhloo, Industry 4.0 technologies can unlock substantial economic and environmental benefits, but only when deployed through architectures that support scalability, robustness, and meaningful knowledge generation [2]. These requirements align closely with the challenges facing modern refineries, where operators must navigate massive data volumes, interconnected process units, and evolving operational regimes.

Together, these factors illustrate a clear need for frameworks that can:

(1) operate effectively in high-dimensional, minimally labeled data environments

(2) capture complex temporal structures

(3) adapt to shifting process conditions, and

(4) provide explainable and actionable diagnostic insight for the refinery engineers

The framework proposed in this work directly addresses these gaps by integrating unsupervised knowledge discovery, adaptive anomaly monitoring, and AI-assisted diagnostic reasoning within a unified and interpretable architecture. From a process systems engineering perspective, the methodological contribution lies in structuring these components into a cohesive workflow that links data-driven temporal representations with adaptive monitoring and context-aware diagnostic reasoning to support real-time, plant-level decision making.

PROPOSED FRAMEWORK

Unsupervised Graph Construction

The first step towards building the unsupervised models is to shift the focus from individual time steps to windows of data. Before this step, the data were preprocessed through standard cleaning procedures, including handling of missing values and normalization to ensure consistent scaling across variables. Then, by analyzing these windows of data—as opposed to just pointwise data—the model can learn temporal dependencies from the sequences of observations. The parameters for this step are straightforward: a window size, W, and a stride, S. Window size is just the fixed length of data in each window and stride is the step size, or percentage of overlap, between consecutive windows. The choice of W and S determine both the temporal resolution of the data and ultimately the computational cost. The next step is to perform dynamic time warping (DTW) calculations in a pairwise fashion for each window of data. DTW is a widely used approach for assessing similarity between time-series sequences that do not share the same duration or pacing. First proposed by Itakura (1975) and later expanded by Sakoe and Chiba (1978), the technique

aligns two signals by flexibly stretching or compressing the time axis to achieve the minimum alignment cost, also known as the DTW distance [4-5]. The resulting alignment costs are then stored in a matrix as shown in **Figure 1**. These will be the basis for the graph model shown later. For these results, the *tslearn* [6] python package was utilized.

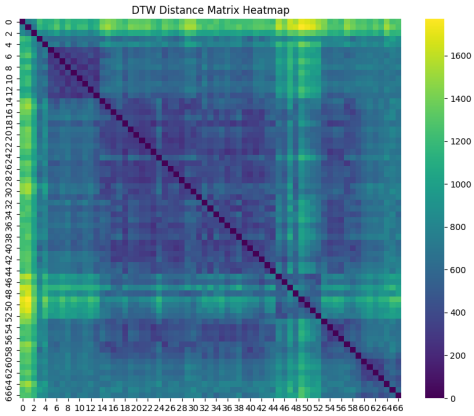

Figure 1. Heatmap of pairwise DTW distances between all windows of an example dataset

Once the DTW matrix is obtained, the next step is to represent the similarity structure from the distances as a graph. Utilizing the well-known k-NN algorithm, each window is treated as a node and edges are created by connecting each node to its k most similar neighbors—determined by the DTW distances. This creates a mapping that connects the temporal relationships within the dataset. These maps can be expressed as graphs by utilizing the *NetworkX* package in python [7]. An arbitrary example is shown in **Figure 2**.

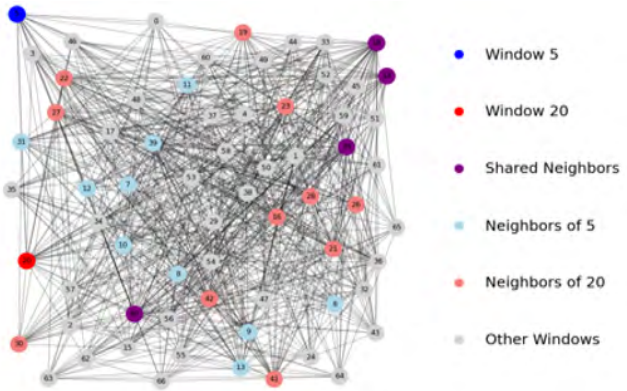

Figure 2. Initial graph highlighting nodes 5 and 20 connected by their k=15 neighbors

The issue now is that we have the mapping—however, it is just a tangled mess. It makes sense mathematically, but it serves no use yet as a visual graph. Since the goal is to ultimately cluster the windows based on different operational regimes, the nodes need to be organized in such a way that this can be achieved. *NetworkX* has several graph layout algorithms that help with this arrangement. The spring layout, also referred to as the force-directed layout, was chosen as it is a physics inspired algorithm in which nodes are modeled as mutually repelling particles and edges act as attractive springs. The implementation in *NetworkX* utilizes the Fruchterman-Reingold [8] algorithm in which the node positions are iteratively updated until a low-energy, near-equilibrium state is reached or if the set number of iterations is achieved. The resulting graph clusters similar nodes closer together and sparsely connected nodes and outliers are placed towards the periphery. This can be seen in **Figure 3** when 50 iterations of the spring layout algorithm were applied. The most similar nodes to window 5 are mostly grouped near window 5 (blue cluster), the most similar nodes to window 20 are grouped near window 20 (red cluster), and some of the intermediate windows are placed at the boundary between the clusters (purple). At this point a graph-based embedding of the dataset is created, based upon the warped temporal distances. From here, traditional clustering techniques, KMEANS, DBSCAN, etc., could be applied to formally label the data from the unsupervised graphs generated.

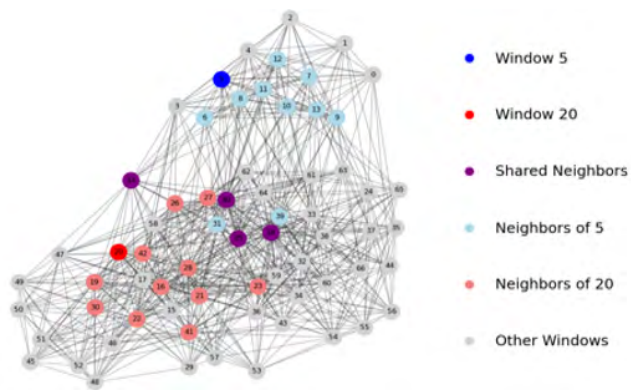

Figure 3. Resulting graph after 50 iterations of the force-directed algorithm was applied

Adaptive Online Monitoring

Since the goal of this framework is to expand beyond just offline modeling to online fault detection, the next step was to see how these graphs could be utilized online. First, the offline models need to be studied and the clusters need to be verified, perhaps by speaking with someone at the plant that is knowledgeable about changes in the process. When a normal operational regime is established, this cluster will be utilized as the baseline for the initial, online graph. Since this is considered normal operations, the natural variation among the distances between these windows will serve great purpose for calculating the initial fault detection threshold. By drawing from the normal cluster, the corresponding

DTW distances for each window can be collected and the distribution can be shown such as in ***Figure 4***. From the distribution, statistical calculations can be performed to highlight the mean, standard deviation, percentiles, etc.

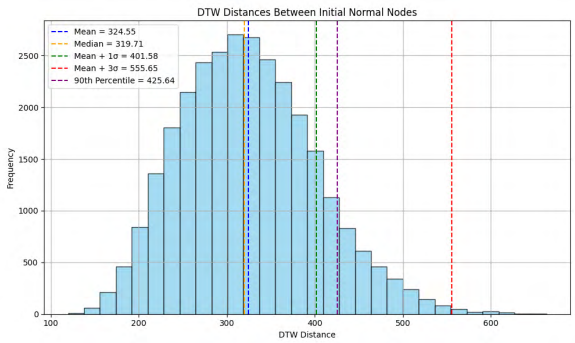

Figure 4: Distribution of pairwise DTW distances between initial normal nodes

Ultimately, the threshold was chosen to be adjustable based on the mean and standard deviations as shown in **Equation 1**. Here μ is the average of the distances, σ is the standard deviation, and X is a multiplier of the standard deviation. These parameters can be tuned based on the dataset at hand and the required fault detection sensitivity.

$$Threshold = \mu + X\sigma \tag{1}$$

To tackle the issue with evolving process conditions, the fault threshold is updated dynamically using an exponential moving average (EMA) of recent DTW values. This allows for the threshold to gradually adjust to changes in process conditions while still remaining sensitive to sudden anomalies. A parameter α is chosen as a smoothing factor the controls the responsiveness of the threshold to recent changes. Let μ_t and σ_t denote the running mean and standard deviation of current DTW distance d_t:

$$\mu_t = \alpha \bullet d_t + (1-\alpha) \bullet \mu_{t-1} \tag{2}$$

$$\sigma_t = \alpha \bullet |d_t - \mu_t| + (1-\alpha) \bullet \sigma_{t-1} \tag{3}$$

Then the adaptive threshold at time t becomes:

$$Threshold_t = \mu_t + X \bullet \sigma_t \tag{4}$$

If the current distance d_t exceeds the current threshold, the window is flagged as an anomaly. If a window is flagged as normal or if the current threshold is less than the initial threshold, then the threshold is updated. This allows the system to only learn during normal operations while also preventing the threshold from being stuck in a minimum. A third updating condition is applied if a fault persists across a specified number of consecutive windows. This mechanism will serve as an update trigger that will ensure new operating conditions are not persistently flagged as faults. By incorporated them into the threshold learning, the model is prevented from becoming obsolete and can evolve as plant operations change—tackling one of the main objectives of this framework.

LLM Diagnostic Agents

Up to this point, two stages in the framework are achieved. First, the graph-based embeddings are created and utilized as an unsupervised model. Second, the graph model is prepared and subjected to live data for fault detection. The third stage is to extend to the fault diagnosis portion.

Fault diagnosis is tricky due to complex causal relationships, feedback loops, and misleading indicators. Traditional approaches—physics-based, expert systems, and data-driven ML—each fall short. Models may be rigid, they may be too complicated to create, or they may lack causal grounding. Data-driven methods such as Principal Component Analysis (PCA) contributions, Shapley Additive exPlanations (SHAP) analysis or Subspace Greedy Search (SGS) are often utilized due to their ease of use. By comparing a set of normal conditions to a set of faulty conditions, features with the most contribution can be highlighted. If these always worked, that would be great, however, these methods are purely black-box and often run into the "false contribution" problem. For example, a SHAP analysis may point to a feature that is highly important, but it is just a reactive symptom and not the true root cause of the fault. To overcome this, large language models (LLMs) are suggested as diagnostic agents.

The latest generation of LLMs—such as OpenAI's ChatGPT models, Anthropic's Claude models, or Google's Gemini models, present a significant opportunity for advancing fault diagnosis in process monitoring systems. LLMs are trained on very large text corpora that make them excellent at predicting human like text sequences. In addition, the current stage of LLMs can handle multimodal inputs such as tables, process diagrams, etc. This capability opens the door for LLMs that act as agents that can perform contextual reasoning that mirrors the diagnostic thought process of a human engineer. Venkatasubramanian and Chakraborty deem these context rich LLMs as "Large Knowledge Models (LKMs)" as opposed to the traditional "LLM" nomenclature [9]. These "LKM" agents can bridge the gap between the domain specific knowledge and data-driven ML outputs to support root cause identification, hypothesis generation, and human-in-the-loop decision making in industrial environments.

Overall Framework

The overall pipeline for this work is shown in **Figure 5**. This summarizes the three stages of the process monitoring workflow including the offline portion where historical data is utilized to build a DTW-based graph, an

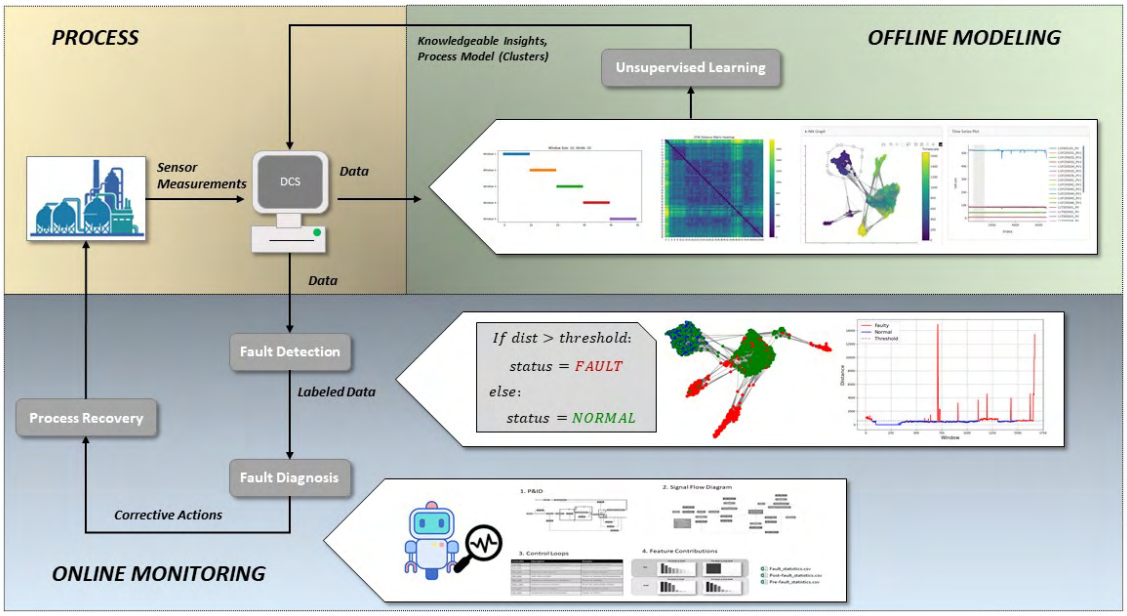

Figure 5. Overall Framework for Process Monitoring icluding the three stages: (1) Unsupervised Graph-based Learning, (2)Adaptive Fault Detection, and (3) LLM-assisted Fault Diagnosis

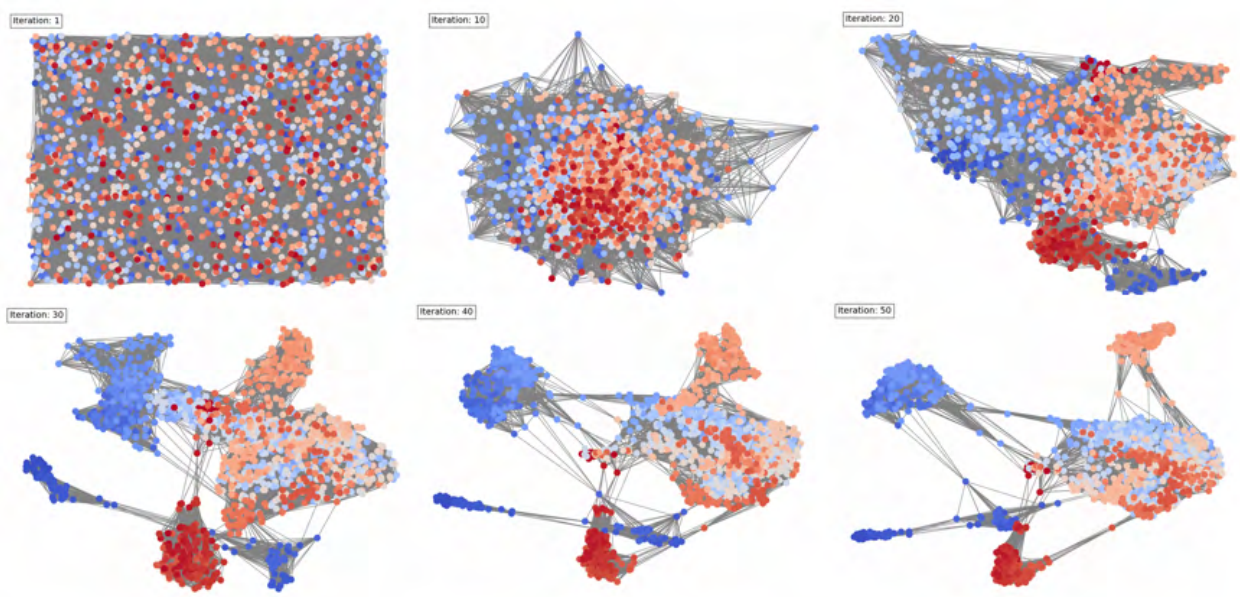

Figure 6. Graph evolution over successive iterations of the physics-based layout (colored by timescale).

online portion where live data can be subjected to the graph to generate fault labels while dynamically updating the fault threshold, and a diagnostic portion where LLM's serve as AI-assistants in real time diagnosis.

RESULTS

The framework was first tested for unsupervised process clustering and process monitoring using data from an industrial-scale pyrolysis furnace that is involved in the production of ethylene from naphtha [10]. The reactor system is housed within a fired furnace where energy is supplied by burning fuel gas. The system consists of six reactor coils through which a mixture of naphtha and steam flows. The mixture is exposed to temperatures approaching 1000 °C, allowing for thermal cracking to produce ethylene. A total of 27 key process variables were monitored at 20-min intervals over a continuous

90-day operating period. Within the batch, there are known start-up and shutdown regimes, noise from a faulty sensor, and gradual drift due to accumulation of coke on the reactor tubes.

Building Graph Models

The dataset was split into windows of length 5 and a stride of 4, corresponding to window spans of 100 minutes with 20 minutes of overlap. This resulted in 1667 individual windows that were used to create the pairwise DTW matrix. A sensitivity analysis of the k parameter was performed and ultimately a value of k=20 was selected as it provided the most well-structured representation of the underlying graph. The evolution of the graph with W=5, S=4, and k=20 is shown in **Figure 6**. At iteration 1, the nodes are all randomly positioned. As iterations of the spring layout algorithm were applied, the similar nodes were slowly pulled toward each other and dissimilar nodes were pushed away from one another—resulting in an embedding with distinct regimes. Furthermore, the two-dimensional coordinates from the produced by the force directed layout were used as inputs to the HDBSCAN clustering algorithm in order to generate real cluster labels. The clustering hyperparameters—min_cluster_size and min_samples—were again tuned via a sensitivity analysis of the resulting graphs. The values of min_cluster_size and min_samples were eventually set to 5 and 10 respectively for the best result. Under this configuration, six distinct clusters were found, along with a small subset of nodes labeled as noise as shown in ***Figure 7***. This cluster identification remained consistent after repeated trials as well.

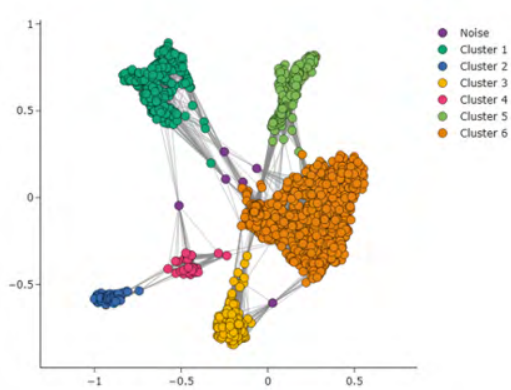

Figure 7. Process Model colored by Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) labels resulting in 6 clusters + noise

Evaluating Graphs for Online Monitoring

To extend the model for online analysis, Cluster 1 was deemed as the normal operational regime. The pairwise DTW distances for this cluster were collected and shown in **Figure 4**. Based on the statistical analysis, the initial fault threshold was set at the mean plus three standard deviations. The graph was then initialized with only the nodes corresponding to Cluster 1. To monitor for faults, the full dataset was then subjected to this initial model and the results are shown in **Figures 8-9**. **Figure 8** shows the graph evolution and **Figure 9** shows the live run chart with the DTW distances for fault detection.

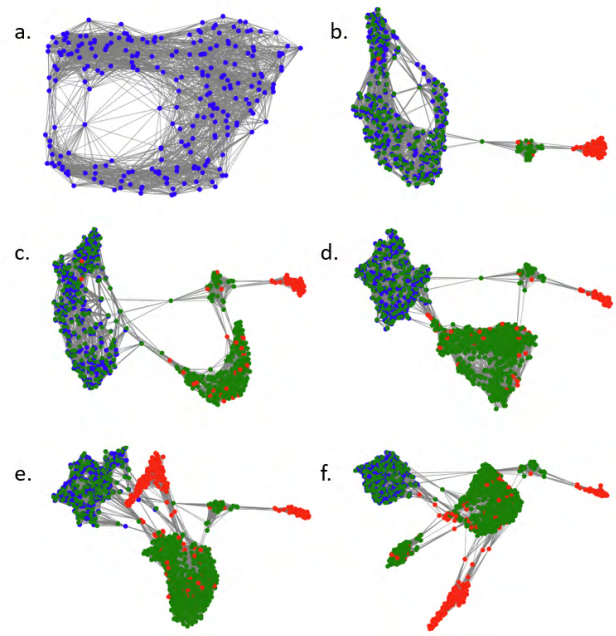

Figure 8. Evolution of the online graph for monitoring (BLUE is initial training nodes, GREEN is live normal nodes, and RED is live faulty nodes)

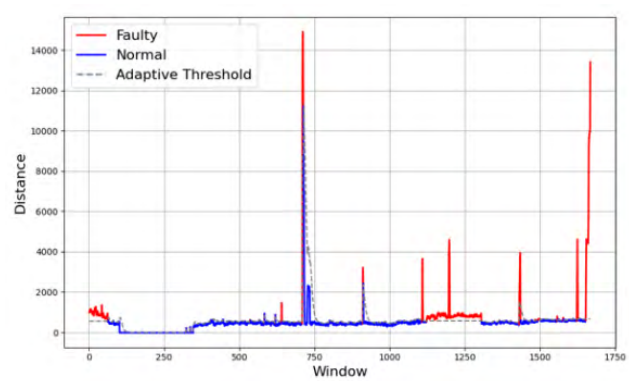

Figure 9. Online Fault Chart with Adaptive Threshold ($\alpha = 0.9, X = 3$)

From these results, several things were accomplished. First, the final resulting graph is similar to the initial graph that was built. This shows that the results are reproducible. Second, the faults that were found were accurate to what we know about the process. It flagged the start-up cluster, the spikes or noise cluster from the sensor noise, a separate operational period cluster, and the final shutdown cluster. Importantly, it did not classify

the process drift as noted in Cluster 6 as faults. This is due to the third point—that the online adaptions were successful. This can be better seen by zooming in on the run chart to see how the dynamic threshold was adapting. This can be observed in **Figure 10**.

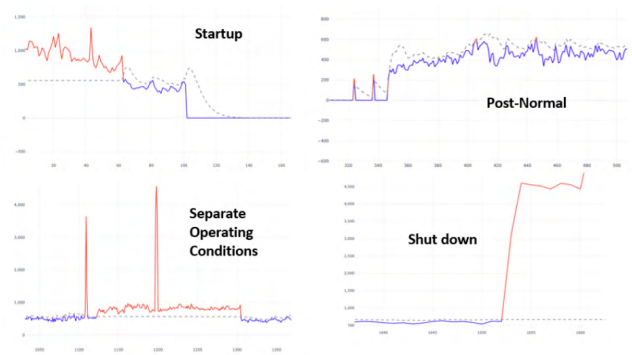

Figure 10. Zooming in on run chart to see the online threshold adaptions

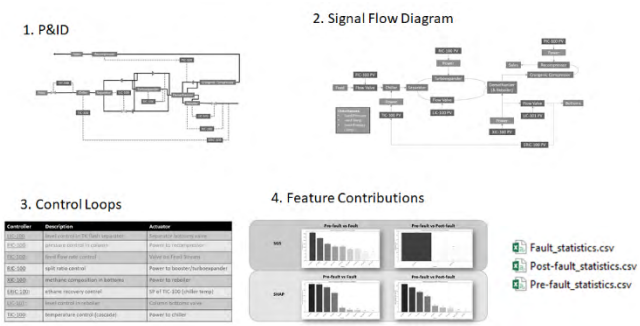

Figure 11. Inputs for Agent Prompting

Utilizing LLMs as Diagnostic Agents

In the final stage, we decided to utilize an ASPEN HYSYS® simulation for an LNG plant for the AI-assisted fault diagnosis. This was chosen since we had the ability to introduce our own faults into the simulation as opposed to relying on only historical datasets. This simulation monitors the operation of a demethanizer column operating under CRR separation conditions. Notably, the simulation has a feed stream defined as natural gas at 4980 kmol/hr supplied at a pressure of 5818 kPa and a temperature of 35 °C. The central unit is a 30-stage distillation tower with an integrated reboiler. A dataset was first created by introducing several faults to the system after running at normal conditions. These faults included a valve malfunction, a sensor bias, a feed disturbance, and a set-point (SP) change. Just like before, the graph model was constructed and the fault threshold was set. Now, when the model was online, the AI-assistant was incorporated. When a set number of faults were labeled, a snapshot of the data was taken and several inputs were collected. These inputs correspond to contextual information such as the P&ID diagram, signal flow diagram, and control loop table as well as the data-driven metrics—the feature contributions from the fault and statistics from the data as shown in **Figure 11**.

These inputs were then sent to the AI-assistant, via OpenAI's API system and using the GPT-4o model. The agent was instructed to use the domain knowledge and the data-driven knowledge to find the root cause of the fault. Specifically, it was prompted to produce a one-page diagnostic report that identifies the root cause, highlights insights, grounds reasoning in domain knowledge, and recommends operator actions. Using this setup, the agent was tested for several fault scenarios. Here, the results are shown for the valve malfunction in **Figures 12-13**.

GPT Response (VALVE MALFUNCTION):

📌 **Fault Type**: Control Error

🔍 **Evidence from inputs**:

- The OP monitoring metrics indicate that FIC-100 is a significant candidate for control loop errors detected by OP trends.
- In the SGS analysis, XIC-100_OP has the highest contribution (18.95%), followed by PIC-100_OP (12.98%) and TIC-100_OP (12.19%).
- However, FIC-100_OP has a very low contribution (0.26%), indicating it is not the root cause.
- The SHAP analysis also shows that FIC-100_OP has a contribution of 10.07%, which is not significantly larger than the others.

🧠 **Root Cause Explanation**:

- The control error is likely due to a malfunction in the FIC-100 control loop, which is responsible for maintaining a stable feed flow rate into the process.
- Given that no significant OP metrics were detected, it suggests that the feed valve may not be responding correctly to the control signals, leading to instability in the feed flow.
- This could be due to a valve malfunction, such as sticking or improper calibration, or a sensor bias affecting the flow measurement.

🛠 **Recommended Next Steps**:

- Operators should inspect the FIC-100 control loop, focusing on the feed valve and its associated flow sensor.
- Calibration checks should be performed, and the valve should be tested for proper operation.
- If necessary, maintenance should be conducted to ensure the valve is functioning correctly and that the sensor readings are accurate.
- Additionally, monitoring should continue to assess the impact of any adjustments made.

Figure 12. Diagnostic report generated by LLM for Valve Malfunction

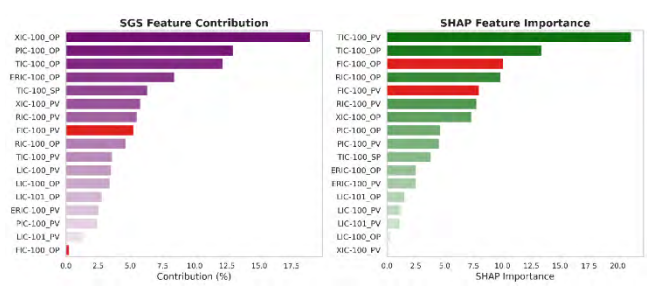

Figure 13. Feature Contributions for SGS and SHAP (true fault features in RED)

In this fault, the feed control valve (FIC-100) was

plugged by decreasing the valve coefficient by 10%. The agent analyzed the inputs that were collected and, in the report, accurately described the fault. It began by first discussing evidence from the inputs. For example, it said that "FIC-100 is a significant candidate for control loop errors detected by OP trends". It then went on to say that FIC-100 was not highly ranked in the feature contributions for both SHAP and SGS, causing a discrepancy. It then went on to the root cause hypothesis. Here it ultimately decided that there was most likely a "malfunction in the FIC-100 control loop." Following this, it suggested corrective actions such as the valve should be "tested for proper operation" and specifically "focus on the feed valve and its associated flow sensor."

Overall, this was a pretty good report as the agent only had knowledge based on its initial knowledge and the knowledge learned from the inputs. The prompt included detailed instructions, but no actual information of what faults occurred. In addition, it is important to note that just the standard SGS and SHAP plots would not have captured the true root cause of the fault as it highlighted the FIC-100 control loop first only as the 8th most important in the SGS plot and 3rd most important in the SHAP plot.

CONCLUSION

In summary, the proposed framework "connects the dots" across key challenges in process monitoring by addressing limited labeled data, capturing temporal dynamics in multivariate time series, adapting to changing process conditions through a dynamic fault threshold, and autonomously generating root-cause hypotheses using LLMs. Together, these elements form a unified, interpretable, and scalable approach for modern refinery operations. The results show that DTW-based graph embeddings effectively reveal operational regimes, the adaptive thresholding reliably separates true faults from natural drift, and the LLM-assisted diagnosis produces context-aware explanations grounded in both domain knowledge and data-driven evidence. Future work will aim at addressing scalability concerns, prompting strategies, and collecting more quantitative results for both the detection and diagnosis portions of the framework. Overall, this work demonstrates the value of integrating unsupervised learning, adaptive monitoring, and AI-assisted diagnosis into a cohesive workflow. Future work will refine LLM prompting strategies, expand multimodal inputs, and evaluate the framework's scalability for broader industrial deployments.

AUTHOR IDENTIFIERS

Author ORCIDs:
Territo K: 0009-0002-1133-915X
Romagnoli JA: 0000-0003-3682-1305

REFERENCES

1. Olaizola I, Quartulli M, Unzueta E, Goicolea J, Flórez J. Refinery 4.0, a review of the main challenges of the industry 4.0 paradigm in oil & gas downstream. Sensors 22:9164 (2022). https://doi.org/10.3390/s22239164
2. Ghobakhloo M. Industry 4.0, digitization, and opportunities for sustainability. Journal of Cleaner Production 252:119869 (2020). https://doi.org/10.1016/j.jclepro.2019.119869
3. Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J. and Ligett, K., 2025. Artificial intelligence index report 2025. arXiv preprint arXiv:2504.07139.
4. Itakura F. Minimum prediction residual principle applied to speech recognition. IEEE Trans. Acoust., Speech, Signal Process. 23:67-72 (1975). https://doi.org/10.1109/tassp.1975.1162641
5. Sakoe H, Chiba S. Dynamic programming algorithm optimization for spoken word recognition. IEEE Trans. Acoust., Speech, Signal Process. 26:43-49 (1978). https://doi.org/10.1109/tassp.1978.1163055
6. Tavenard, R., Faouzi, J., Vandewiele, G., Divo, F., Androz, G., Holtz, C., Payne, M., Yurchak, R., Rußwurm, M., Kolar, K. and Woods, E., 2020. Tslearn, a machine learning toolkit for time series data. Journal of machine learning research, 21(118).
7. Hagberg AA, Schult DA, Swart PJ. Exploring network structure, dynamics, and function using networkx. Proceedings of the Python in Science Conference :11-15 (2008). https://doi.org/10.25080/tcwv9851
8. Fruchterman T. M. J., Reingold E. M. (1991). Graph drawing by force-directed placement. Software: Practice and Experience, 21(11), 1129–1164.
9. Venkatasubramanian V, Chakraborty A. Quo vadis chatgpt? from large language models to large knowledge models. Computers & Chemical Engineering 192:108895 (2025). https://doi.org/10.1016/j.compchemeng.2024.108895
10. Romagnoli JA, Briceño-Mena L, Manee V. AI in chemical engineering. CRC Press (2024). https://doi.org/10.1201/9781003455905

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Hybrid Data-Driven Approach for the Optimization of an Industrial Alkylation Unit

Rastislav Fáber[a]*, Karol Ľubušký[b], and Radoslav Paulen[a]

[a] Faculty of Chemical and Food Technology, Slovak University of Technology in Bratislava, 812 37 Bratislava, Slovakia
[b] Slovnaft, a.s., 824 12 Bratislava, Slovakia
* Corresponding Author: rastislav.faber@stuba.sk

ABSTRACT

We develop a multi-fidelity soft-sensing framework to reconcile online (low-fidelity) industrial measurements with sparse (high-fidelity) laboratory samples from an alkylation unit in a refinery. A first-principles model is used to generate an additional low-fidelity dataset and train a surrogate that predicts the output variable. We investigate whether incorporating sparse high-fidelity laboratory data with the low-fidelity data improves prediction accuracy. A multi-fidelity predictor forms a corrected output by learning the residual between the high-fidelity observations and the low-fidelity surrogate using Gaussian process regression. The simplest model structure performs best, reducing test prediction error (computed against laboratory samples) by 31.1% relative to the currently deployed industrial analyzer and outperforming a standard high-fidelity-only model trained on laboratory data. Overall, the simplified surrogate model captures the main industrial trends well enough to serve as a reliable low-fidelity input for the multi-fidelity soft sensor.

Keywords: Process Simulation, Deep Learning, Alkylation, Process Optimization, Energy Efficiency, Data-Driven Modeling

INTRODUCTION

Industrial processes operate under tight economic and safety constraints, and their units are expected to run close to optimal conditions for long periods [1]. Deviations from the desired operating point reduce yield, increase utility consumption, and can accelerate equipment wear [2]. At the same time, plants face feed variability, seasonal effects, catalyst aging, and operational interventions that change the process behavior over time [3]. This combination drives operators to rely on monitoring and decision support tools [4] that are reliable, and easy to maintain.

Modern plants are heavily instrumented. Flow, pressure, and temperature measurements are usually available at high sampling rates and are directly usable for monitoring and control. Other variables readings can be harder to obtain. Product quality indicators and stream compositions are often measured by online sensors or by laboratory analyses [5]. Online analyzers provide frequent updates but may drift, show bias, or suffer from maintenance related faults, whereas lab measurements are less frequent but are typically treated as a reference signal. This results in a practical mismatch between what is needed for operation, and what is available as reliable ground truth [6]. Since quality measurements are not collected continuously, plants often operate with wider safety margins, which typically reduces profitability [7].

Data-driven models are a common way to bridge the gap between high-rate process measurements and sparse, reliable quality data, for example through soft sensors that can infer difficult-to-measure outputs from easy-to-measure inputs. In industrial operation, the implemented soft sensor must handle missing values, sensor faults, and operating mode changes. They also need to be interpretable to diagnose failures. Moreover, many industrial soft sensors are deliberately simple and use a small set of engineered regressors that aim to approximate the dominant material and energy balances [8, 9]. Even the best-performing offline model is not useful if it depends on variables that are unreliable, unavailable in real time, or unstable after maintenance [10]. Some recent deep learning soft sensors aim to improve prediction by extracting relevant hidden features while eliminating redundancy and maintaining structure simplicity [11, 12].

First-principles process models can also address

some of these issues because they encode mass and energy conservation laws, phase equilibria, and known reaction pathways. They can provide consistent data under systematic perturbations. In large-scale industrial settings, however, fully accurate mechanistic models are difficult to obtain. Parameter uncertainty, unknown disturbances, unmeasured states, and changing plant conditions often limit predictive accuracy. This makes a direct mechanistic model unrealistic when the goal is a reliable, daily decision-making support [13]. Broader reviews report that implementations of many first-principles models remain constrained by uncertainty and integration or maintenance burden [14].

These constraints motivate hybrid workflows that combine mechanistic structure with data-driven learning. In such workflows, mechanistic models are used to provide physically consistent trends and to expand coverage of the operating space, while plant data compensate unmeasured disturbances [15]. Surrogate model frameworks (e.g., deep neural networks trained on simulation data) are commonly used to retain process trends while reducing computational cost for downstream tasks [16]. In this work, a deliberately simplified first-principles simulator is used to generate an additional fidelity level that supplies trend information beyond routine operation and supports surrogate training. Each of these data sources can be treated as a separate information stream within an MF framework. High-frequency measurements act as a low-fidelity (LF) data that capture dynamics and disturbances. Sparse laboratory samples provide high-fidelity (HF) reference to correct bias. The objective is to explore whether a simple surrogate trained on simulated inputs can explain the measured industrial output.

In this work, we build a workflow that (*i*) uses a mechanistic process model to generate a large synthetic dataset under controlled variations of key control parameters, (*ii*) trains a surrogate model that reproduces the simulator trends at low computational cost, and (*iii*) integrates the predictions with industrial measurements in an MF framework. The main focus is to improve monitoring and a deployable structure with limited number of input variables, and a predictor that can be retrained when an operation point changes.

PROBLEM STATEMENT

Many industrial regression problems involve multiple information sources with different fidelity. Let $x \in \mathbb{R}^p$ denote the regressor vector (measured variables) and $y \in \mathbb{R}$ the target output variable. An industrial LF dataset provides T online samples $\{(x_t^{\mathrm{LF}}, y_t^{\mathrm{LF}})\}_{t=1}^{T}$, while a HF reference provides only K trusted laboratory samples $\{(x_k^{\mathrm{HF}}, y_k^{\mathrm{HF}})\}_{k=1}^{K}$, with $K \ll T$. In addition, an auxiliary LF' dataset may be available from an alternative source, or model simulation $\{(x_n^{\mathrm{LF'}}, y_n^{\mathrm{LF'}})\}_{n=1}^{N}$, with $K \ll T \ll \mathrm{N}$. The objective is to use the sparse HF reference to correct predictions from the frequent LF data within an MF framework.

We consider three information sources. Simulation dataset (LF' hereafter referenced as X^{SIM}) generated from variations of key control variables within defined limits. Here, p denotes the number of monitored variables:

$$X^{\mathrm{SIM}} = \{(x_n^{\mathrm{SIM}}, y_n^{\mathrm{SIM}})\}_{n=1}^{N}, x_n^{\mathrm{SIM}} \in \mathbb{R}^p, y_n^{\mathrm{SIM}} \in \mathbb{R}. \tag{1}$$

Industrial LF variables $X^{\mathrm{IND}} = \{(x_t^{\mathrm{LF}}, y_t^{\mathrm{LF}})\}_{t=1}^{T}$ consist of a time series of online measurements with regressors x_t^{IND}, and the output approximation y^{LF} (e.g., online analyzer) that may be biased or faulty available in real time. Industrial HF samples $X^{\mathrm{LAB}} = \{(x_k^{\mathrm{HF}}, y_k^{\mathrm{HF}})\}_{k=1}^{K}$ are available at sparse time instants, where y^{HF} is obtained from laboratory analyses.

Let θ^{LF} denote the parameters of the surrogate $f^{\mathrm{SIM}}(\cdot;\theta^{LF})$ learned from the simulation dataset X^{SIM} and then evaluated on the industrial data X^{IND}:

$$\hat{y}_t^{\mathrm{LF}} = f^{\mathrm{SIM}}(x_t^{\mathrm{LF}}; \theta^{\mathrm{LF}}). \tag{4}$$

Two model classes, linear and nonlinear, are considered. The role of this step is to provide a dense, physics-consistent prediction over all industrial time points $t = 1, \dots, T$.

We model the HF-to-LF discrepancy at HF time points k as:

$$\Delta_k = y_k^{\mathrm{HF}} - \hat{y}_k^{\mathrm{LF}}. \tag{5}$$

We train a correction model $g^{\mathrm{IND}}(\cdot)$ on HF residuals, and predict all differences as:

$$\hat{\Delta}_t = g^{\mathrm{IND}}(z_t^{\mathrm{LF}}; \phi_t^{\mathrm{LF}}), \tag{6}$$

where z_t^{LF} is defined as a representation of inputs with added information from a best performing surrogate model $f^{\mathrm{SIM}}(\cdot)$:

$$z_t^{\mathrm{LF}} = [x_t^{\mathrm{LF}^{\mathrm{T}}}, \hat{y}_t^{\mathrm{LF}}]^{\mathrm{T}}. \tag{7}$$

The final MF predictor is then defined as:

$$\hat{y}_t^{\mathrm{MF}} = \hat{y}_t^{\mathrm{LF}} + \hat{\Delta}_t. \tag{8}$$

If we assume a perfect prediction that satisfies $\hat{\Delta}_{\mathrm{t}} \approx \Delta_t$, then, based on (5) and (8), we can state:

$$\hat{y}_t^{\mathrm{MF}} = \hat{y}_t^{\mathrm{LF}} + \hat{\Delta}_t \approx \hat{y}_t^{\mathrm{LF}} + (y_t^{\mathrm{HF}} - \hat{y}_t^{\mathrm{LF}}) = y_t^{\mathrm{HF}}, \tag{9}$$

where y_t^{HF} denotes the theoretical HF data. Therefore, the MF prediction can be interpreted as a real-time estimate of the HF output $\hat{y}_t^{\mathrm{MF}} \approx y_t^{\mathrm{HF}}$.

METHODOLOGY

Following our previous work [17], all models use the

same feature-engineered regressor vector $x \in \mathbb{R}^p$, consisting only of variables that are both statistically relevant and operationally trusted.

Rows containing unusable (i.e. NaN, or out-of-range outliers) values in any regressor or in the corresponding output are removed prior to model training and evaluation. A standardization (zero mean, unit variance) is applied to scale and anonymise the data.

Principal Component Analysis (PCA) is a linear dimensionality-reduction method applied to z_t^{LF} as a preprocessing step to obtain a smaller set of principal components (PCs). By keeping only the most relevant q components ($q \ll p$). PCA reduces noise and collinearity and yields a compact representation that is easier to model.

We train the surrogate model $f^{\mathrm{SIM}}(\cdot)$ model using two separate approaches:

1. Linear Ordinary Least Squares (OLS) regression is trained and validated on the SIM inputs x_n^{SIM} and output y_n^{SIM}.

2. Nonlinear Neural network (NN) uses a feed-forward network trained on the same SIM pairs $(x_n^{\mathrm{SIM}}, y_n^{\mathrm{SIM}})$ to provide an alternative $\hat{y}_t^{\mathrm{LF}}$ prediction that, unlike linear models, can capture complex, non-linear relationships and feature interactions within the dataset.

Gaussian Process regression (GPR) is a nonlinear regression method that models an unknown function as a distribution over functions. The model is fully specified by mean $m(x)$, and a covariance function $k(x, x')$ (kernel).

$$g(x) = \mathcal{GP}\big(m(x), k(x, x')\big) \tag{10}$$

The kernel defines how strongly two inputs are correlated based on their distance, and tuned hyperparameters (e.g., length scale ℓ and signal variance δ^2). A commonly used RBF kernel can be defined as:

$$k(x, x') = \delta^2 exp\big(-||x - x'||^2/2\ell^2\big). \tag{11}$$

In this work, the GP is used to learn the correction term Δ_k.

Performance is evaluated only at time indices $t(k)$, where HF measurements are available, using Root Mean Squared Error (RMSE) regression metric:

$$\mathrm{RMSE} = \sqrt{\frac{1}{K}\sum_{\forall k}\big(y_k^{\mathrm{HF}} - \hat{y}_{t(k)}\big)^2}. \tag{12}$$

CASE STUDY

We study an industrial alkylation unit of our industrial partner, Slovnaft, a.s., (Bratislava, Slovakia), that produces high-octane alkylate by reacting C_3-C_4 olefins with i-C_4 isobutane in an H_2SO_4 acid-catalyzed process. Fig. 1 shows the simplified layout, where the reactant mixture is cooled and fed to the reactor section, and the effluent is separated to recover and recycle isobutane while the alkylate product is withdrawn.

We focus on the isobutane-to-olefin ratio control as it strongly influences product quality and plant economy. Excess of isobutane is needed to suppress undesired side reactions, however, a key online analyzer (y_t^{LF}) on a recycle stream provides unreliable measurements of i-C_4 affecting the amount of recycled material. Simply increasing the recycle flows would put unnecessary load on the downstream separation, deteriorating plant economics. A better monitoring strategy is needed.

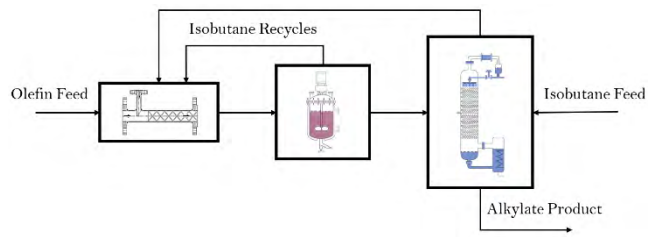

Figure 1. Overview of the simplified industrial process.

IMPLEMENTATION

A simplified plant flow-sheet model is built in the gPROMS ProcessBuilder [18]. A simulation dataset X^{SIM} is then generated by systematically varying operating parameters in the range of plant operating conditions to obtain the simulation dataset $x_n^{\mathrm{SIM}} \in \mathbb{R}^p$ ($p = 6$) and y_n^{SIM}. Although more steady states could be generated, we limited the simulation set to $N = 5{,}000$ because the simplified first-principles model is used only to supply

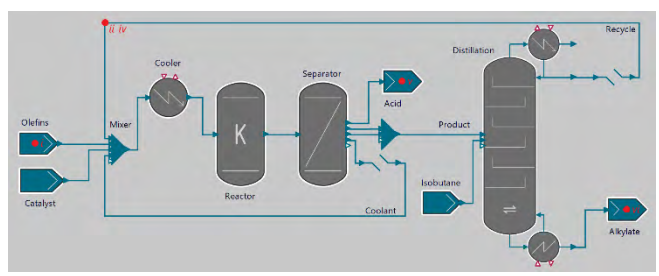

Figure 2. Overview of the simplified alkylation unit model developed in the gPROMS environment.

mass-balance-consistent trends for training, while the industrial dataset already provides dense coverage in time ($T > N$). The inputs (red dots in Fig. 2) were iteratively selected in previous work as most relevant to explain the output variable y_t^{IND}. They include four concentrations measured by online analyzers (olefin feed propylene (i), deisobutanizer recycle propane (ii), isobutane (iii), and n-butane (iv)) and two flow rates (recycled fresh acid (v), and alkylate to storage (vi)).

We compare the simulated and industrial data distributions. The simplified gPROMS model is able to

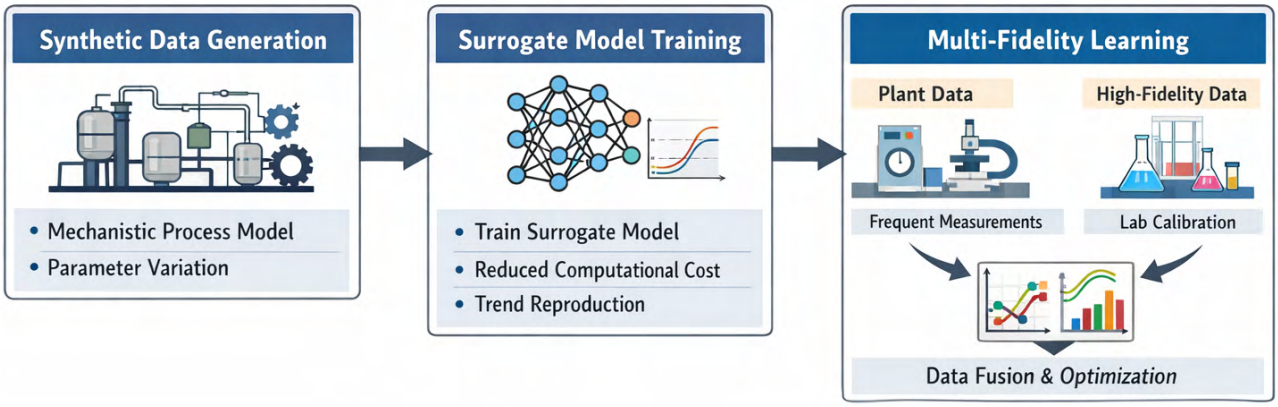

Figure 3. Overview of the proposed hybrid multi-fidelity workflow combining multiple fidelity sources.

reproduce the dominant trends of the plant, although some differences in variability and distribution remain.

We partition the simulation data X^{SIM} $(80/20\%)$ into training (TR) and validation subsets to learn the initial surrogate model. Similarly, the industrial dataset X^{IND} provides $T = 13,000$ online samples, while the laboratory dataset X^{LAB} provides $K = 18$ reference samples. Data is split $(60/40\%)$ based on balanced HF indices into training and testing (TS) sets by alternating the HF measurements. Following the flowchart shown in Fig. 3, we train and validate the surrogate models $f^{\mathrm{SIM}}(\cdot)$ in Python using a nonlinear feed-forward NN implemented by `Keras` [19] with four fully connected layers of sizes 6-3-2-1 and activation functions SELU, tanh, SELU, and linear, respectively. The linear OLS regression, i.e., minimizing the sum of squares, is implemented by `scikit-learn` [20].

Trained prediction models are evaluated on the LF industrial dataset X^{IND} to predict the output $\hat{y}_t^{\mathrm{LF}}$; see (4). The most accurate prediction is subsequently used as an additional input in the MF correction model (8).

To reduce dimensionality, we use PCA on this extended input vector $z_t^{\mathrm{LF}} \in \mathbb{R}^7$. The resulting 4 PCs are used to train a GPR on the HF training points to predict the HF-to-LF differences; as defined in (6).

The MF estimate (8), is obtained by correcting the $\hat{y}_t^{\mathrm{LF}}$ prediction, and performance is reported using RMSE and parity plots. A simple HF model (GPR trained only on HF data) is also trained directly on the HF data X^{LAB} to replicate standard practice and serve as a reference. The final soft-sensor predictions are evaluated against laboratory measurements y_k^{HF}.

RESULTS

The presented workflow encompasses two evaluation stages. First, we verify surrogate models on the SIM dataset, where OLS model achieves RMSE (TR) = 0.1772 and RMSE (TS) = 0.1879, while the NN model achieves much lower RMSE (TR) = 0.0136 and RMSE (TS) = 0.0127.

Secondly, we evaluate the trained predictors on the IND dataset against laboratory measurements y_k^{HF} using RMSE on the TR and TS splits. This is visualized in Figure 4 and quantitatively summarized in Table 1. The OLS-based MF model (black points) and NN-based MF model (red points) are presented in the left and right columns, respectively. The top row is TR and the bottom row is TS. The blue diagonal represents a perfect fit.

As a reference, the online analyser y_t^{LF} evaluated at the HF indices k yields RMSE (TR) = 0.6759 and RMSE (TS) = 0.5557. This confirms a noticeable mismatch between the LF and the HF data.

Table 1: Predictors compared to laboratory data y_k^{HF}.

	RMSE (TR)	RMSE (TS)	reduction % (TS)
$y_{t(k)}^{\mathrm{LF}}$	0.6759	0.5557	-
$\hat{y}_{t(k)}^{\mathrm{HF}}$	0.4757	0.4735	↓ 14.8%
	OLS		
$\hat{y}_{t(k)}^{\mathrm{LF}}$	0.5373	0.4741	↓ 14.7%
$\hat{y}_{t(k)}^{\mathrm{MF}}$	0.4605	0.3831	**↓ 31.1%**
	NN		
$\hat{y}_{t(k)}^{\mathrm{LF}}$	1.1921	1.3297	↑139.3%
$\hat{y}_{t(k)}^{\mathrm{MF}}$	0.5566	0.6127	↑ 10.2%

The simple OLS predictor $\hat{y}_t^{\mathrm{LF}}$ improves the output monitoring, achieving RMSE (TR) = 0.5373 and RMSE (TS) = 0.4741, i.e., a 14.7% reduction on the TS split relative to the analyser. This suggests that, even with a simple linear mapping, the simulator trends and engineered regressors carry useful information that partially explains the HF output.

Adding the OLS $\hat{y}_t^{\mathrm{LF}}$ to the MF correction provides the best overall performance. The MF estimate reaches RMSE (TR) = 0.4605 and RMSE (TS) = 0.3831, corresponding to a 31.1% TS improvement. The improvement over the HF-only model (0.3831 vs 0.4735 on TS)

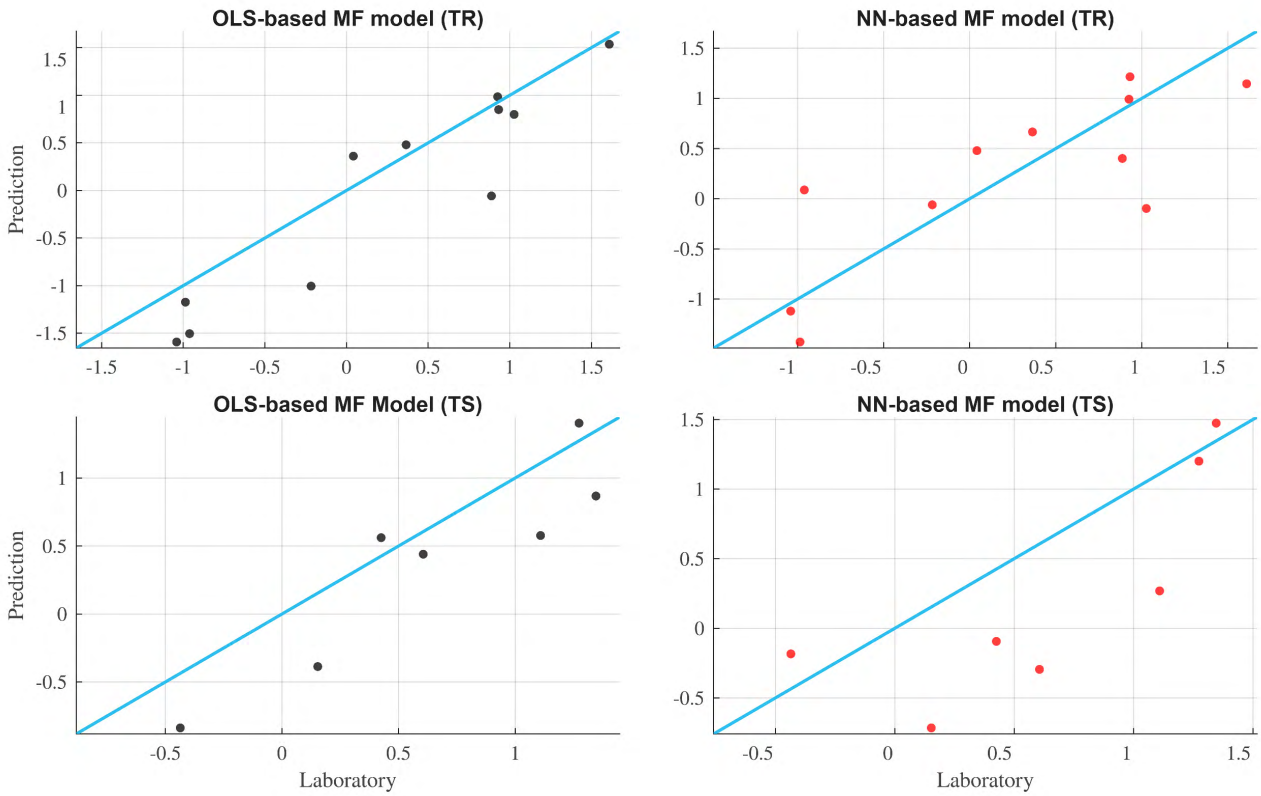

Figure 4. Parity plots of MF predictions $\hat{y}^{\mathrm{MF}}_{t(k)}$ compared to the laboratory data y^{HF}_{k}.

indicates that the correction benefits from the dense LF data, rather than relying on sparse HF data alone.

The simple HF model yields RMSE (TR) = 0.4757 and RMSE (TS) = 0.4735 (14.8% improvement on TS) is comparable with the OLS predictor, but it is worse than MF on the TS split.

Although the NN fits the simulation data accurately (near-zero simulation RMSE), its prediction on the industrial data is poor; RMSE (TR) = 1.1921 and RMSE (TS) = 1.3297 (worse by 139.3%). This may be caused by the domain-shift, where the NN learns simulator-specific nonlinearities that do not map clearly to the real plant (different disturbances, unmodelled constraints).

The MF correction with the NN $\hat{y}^{\mathrm{LF}}_{t}$ partially improves this mismatch (RMSE = 0.6127 on TS), but it still does not surpass the $\hat{y}^{\mathrm{HF}}_{t(k)}$, nor the $y^{\mathrm{LF}}_{t(k)}$.

Overall, the results show that MF correction is most effective when the predictor provides a reasonable trend-consistent estimate.

CONCLUSION

A multi-fidelity correction framework was applied to reconcile dense industrial measurements with sparse laboratory references using a simulator-trained surrogate and a Gaussian process corrector. The ordinary least squares-based multi-fidelity model achieves the best performance, reducing test prediction error against laboratory measurements by 31% relative to the industrial analyzer at high-fidelity time points, and outperforming a standard high-fidelity model (soft sensor). These results indicate that, even with a deliberately simplified simulator, the surrogate is able to capture the underlying trend, enabling the multi-fidelity correction to further improve predictions toward laboratory accuracy. The neural network surrogate, despite excellent fit to simulation data, did not transfer robustly to the industrial dataset; limiting the multi-fidelity correction.

Future work should focus on increasing physical realism and coverage of the simulation model to provide more information in the training phase.

ACKNOWLEDGEMENTS

This work is funded by the Slovak Research and Development Agency under the project APVV-21-0019, by the Scientific Grant Agency of the Slovak Republic under the grant 1/0263/25. We acknowledge the financial support provided by the European Union through the NextGenerationEU programme under the Recovery and Resilience Plan for Slovakia, projects no. 09I01-03-V05-00002, no. 09I01-03-V04-00024, and the Early Stage

Grant no. 09I03-03-V05 (23-04-06-A).

AUTHOR IDENTIFIERS

Author ORCIDs:
Fáber R: 0000-0003-0088-3467
Paulen R: 0000-0002-1599-2634

REFERENCES

1. Engell S. Feedback control for optimal process operation. Journal of Process Control 17:203-219 (2007). https://doi.org/10.1016/j.jprocont.2006.10.011
2. Qin SJ. Survey on data-driven industrial process monitoring and diagnosis. Annual Reviews in Control 36:220-234 (2012). https://doi.org/10.1016/j.arcontrol.2012.09.004
3. Adams JN, van Zelst SJ, Rose T, van der Aalst WMP. Explainable concept drift in process mining. Information Systems 114:102177 (2023). https://doi.org/10.1016/j.is.2023.102177
4. Thalmann S, Mangler J, Schreck T, Huemer C, Streit M, Pauker F, Weichhart G, Schulte S, Kittl C, Pollak C, Vukovic M, Gashi M, Rinderle-Ma S, Suschnigg J, Jekic N, Lindstaedt S. Data Analytics for Industrial Process Improvement: A Vision Paper. *IEEE 20th Conference on Business Informatics (CBI)* 92–96
5. Kadlec P, Grbi? R, Gabrys B. Review of adaptation mechanisms for data-driven soft sensors. Computers & Chemical Engineering 35:1-24 (2011). https://doi.org/10.1016/j.compchemeng.2010.07.034
6. Farsang B, Gomori Z, Horvath G, Nagy G, Németh S, Abonyi J. Simultaneous Validation of Online Analyzers and Process Simulators by Process Data Reconciliation. *Chemical Engineering Transactions* 32, 1303–1308 (2013)
7. Colling E, Farenzena M, Trierweiler JO. Data-driven prediction of intermediate product properties in petroleum refining. Ind. Eng. Chem. Res. 64:23510-23521 (2025). https://doi.org/10.1021/acs.iecr.5c02838
8. Syauqi A, Nagulapati VM, Chaniago YD, Ningtyas JA, Andika R, Lim H. Advancement in power-to-methanol integration with steel industry waste gas utilization through solid oxide electrolyzer cells: surrogate model-based approach for optimization. Sustainable Energy Technologies and Assessments 73:104160 (2025). https://doi.org/10.1016/j.seta.2024.104160
9. Fáber R, Vaccari M, Bacci di Capaci R, ?ubuský K, Pannocchia G, Paulen R. Data-based multi-fidelity modeling for online sensors correction. Computers & Chemical Engineering 211:109665 (2026). https://doi.org/10.1016/j.compchemeng.2026.109665
10. Kadlec P, Gabrys B, Strandt S. Data-driven soft sensors in the process industry. Computers & Chemical Engineering 33:795-814 (2009). https://doi.org/10.1016/j.compchemeng.2008.12.012
11. Jin H, Dong X, Qian B, Wang B, Yang B, Chen X. Soft sensor modeling using deep learning with maximum relevance and minimum redundancy for quality prediction of industrial processes. ISA Transactions 159:293-311 (2025). https://doi.org/10.1016/j.isatra.2025.02.010
12. Zhang L, Ren G, Li S, Du J, Xu D, Li Y. A novel soft sensor approach for industrial quality prediction based TCN with spatial and temporal attention. Chemometrics and Intelligent Laboratory Systems 257:105272 (2025). https://doi.org/10.1016/j.chemolab.2024.105272
13. Peterson L, Gosea IV, Benner P, Sundmacher K. Digital twins in process engineering: an overview on computational and numerical methods. Computers & Chemical Engineering 193:108917 (2025). https://doi.org/10.1016/j.compchemeng.2024.108917
14. Iranshahi K, Brun J, Arnold T, Sergi T, Müller UC. Digital twins: recent advances and future directions in engineering fields. Intelligent Systems with Applications 26:200516 (2025). https://doi.org/10.1016/j.iswa.2025.200516
15. Sun B, Yang C, Wang Y, Gui W, Craig I, Olivier L. A comprehensive hybrid first principles/machine learning modeling framework for complex industrial processes. Journal of Process Control 86:30-43 (2020). https://doi.org/10.1016/j.jprocont.2019.11.012
16. Nargesian F, Samulowitz H, Khurana U, Khalil EB, Turaga D. Learning feature engineering for classification. Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence :2529-2535 (2017). https://doi.org/10.24963/ijcai.2017/352
17. F?ber R, Vaccari M, Capaci RB, Lubu?k? K, Pannocchia G, Paulen R. Soft-sensor-enhanced monitoring of an alkylation unit via multi-fidelity model correction. Systems and Control Transactions 4:1389-1395 (2025). https://doi.org/10.69997/sct.125527
18. Siemens, gPROMS Digital Process Design and Operations.
19. Abadi M. Tensorflow: learning functions at scale. SIGPLAN Not. 51:1-1 (2016). https://doi.org/10.1145/3022670.2976746

20. Pedregosa F, *et al.*, Scikit-learn: Machine Learning in Python, *JMLR 12* pp. 2825-2830 (2011)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Energy Management of a Renewable-Powered Alkaline Electrolyzer System: A Comparative Study of Nonlinear Optimization Methods

Loukas Kyriakidis[a]*, Jonaed Bin Mustafa Kamal[a], Saskia Bublitz[a], Bogdan Dorneanu[b], and Harvey Arellano-Garcia[b]

[a] German Aerospace Center, Institute of Low-Carbon Industrial Processes, Cottbus, Germany
[b] Brandenburg University of Technology, Chair of Process and Plant Technology, Cottbus, Germany
* Corresponding Author: loukas.kyriakidis@dlr.de.

ABSTRACT

Energy management plays a crucial role in achieving efficient and sustainable operation of industrial energy systems. With the increasing integration of renewable electricity and the growing complexity of hydrogen production networks, effective control strategies are required to minimize operational costs and carbon footprint. However, the uncertain nature of renewable energy sources, such as photovoltaic (PV) power, complicates their accurate forecasting and challenges the optimal energy management of system components. To deal with uncertainties, the rolling horizon approach (RHA) provides a practical framework for adaptive decision-making by repeatedly solving optimization problems over moving time windows while updating system data in real time. In RHA-based energy management, linear or linearized system models are often employed and optimized by linear methods to reduce computational complexity; however, these simplifications can compromise physical realism and lead to suboptimal decisions. Although RHA can also incorporate local, or global deterministic and stochastic algorithms for nonlinear problems, such approaches frequently suffer from high computational effort, slow convergence, local optima, and difficulty in ensuring constraint satisfaction in large-scale nonlinear systems. To overcome these limitations, this work employs the novel hybrid optimization method "BO-IPOPT"—a combination of Bayesian Optimization (BO) for global exploration and the Interior Point OPTimizer (IPOPT) for rapid local refinement. Applied to an industrial hydrogen production system, BO-IPOPT outperforms state-of-the-art approaches in accuracy and robustness by achieving lower operational costs at the same CPU time while satisfying all constraints. Finally, the influence of the uncertainties in PV generation on the performance of the energy management system is analyzed.

Keywords: Energy management, Rolling horizon approach, BO-IPOPT, Renewable hydrogen system

INTRODUCTION

Energy utility systems play a central role in industrial operations, and their efficient management is increasingly crucial as sustainable energy technologies are integrated. The growing penetration of renewable energy sources and energy storage systems has increased the system complexity, making optimal energy management essential [1]. Key objectives include minimizing operating costs, reducing CO_2 emissions, and improving the overall energy efficiency [2]. However, these goals are challenged by the variability of renewable generation, fluctuating process demands, grid stability requirements, power quality constraints, and uncertainty in energy forecasts [3].

Conventional control strategies, such as PID controllers and rule-based scheduling [4, 5], remain widely used but are inherently limited. These approaches react only to current errors, lack predictive capabilities, struggle with multi-objective optimization and system-wide constraints, and are poorly suited to dynamic and uncertain operating conditions. As a result, more advanced

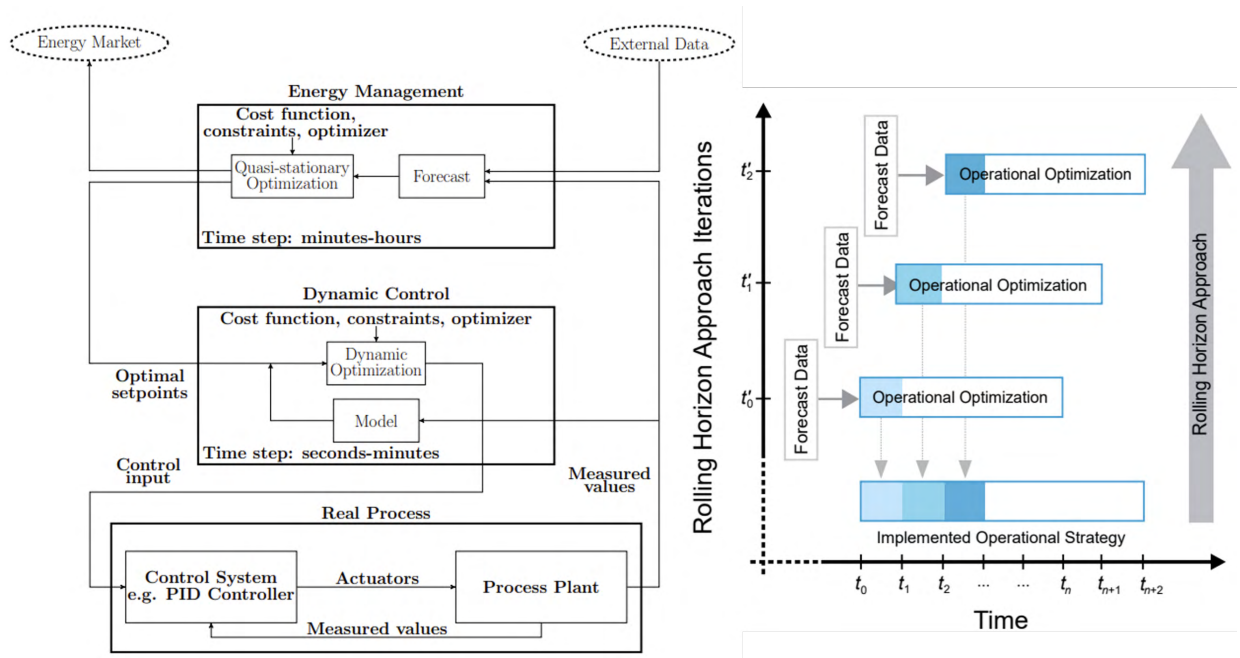

Figure 1. Detailed block structure of the different control (time) levels in real-time optimization of industrial processes (**left**) and visualization of the RHA (**right**) [6].

energy management strategies that can anticipate future conditions, coordinate multiple system components, and explicitly account for uncertainty are increasingly necessary.

The role of energy management within the broader framework of real-time optimization is displayed in **Figure 1** on the left [6]. At the highest level, predictive optimization computes optimal setpoints for energy flows, temperatures, and power exchange over time horizons ranging from minutes to hours, with the objective of minimizing operating costs and emissions. These setpoints are passed to the second level, where component-level controllers ensure stable and reliable system behavior by tracking the prescribed targets. The third level describes the physical process itself, including the plant, instrumentation, and actuators that execute the control actions in real time. While all three levels are essential for system operation, this work focuses exclusively on the system (highest) level, i.e., the energy management layer.

To address the uncertainties in energy systems, various methods have been proposed, including stochastic programming, robust optimization, fuzzy logic, and Monte Carlo simulation [6]. Although effective, these methods often entail high computational effort, limiting their applicability in time-sensitive operational contexts. To address this limitation, the rolling horizon approach (RHA) has emerged as a practical and widely adopted solution [6]. By repeatedly solving a finite-horizon optimization problem using updated forecasts and measurements (see **Figure 1** on the right), rolling horizon optimization enables adaptive, real-time decision-making under uncertainty. This makes it particularly well suited for industrial energy systems with high renewable penetration and rapidly changing operating conditions.

The performance of RHA is closely linked to the choice of the optimization method, which governs how decisions are updated as forecasts and system conditions evolve. Most RHA applications rely on linear or mixed-integer linear formulations, as these enable efficient handling of continuous and discrete variables and offer favorable computational scalability [7, 8]. However, the simplifications inherent in linear and linearized models often reduce physical accuracy, particularly in systems exhibiting strong nonlinear behavior, wide operating ranges, or complex component interactions. In such cases, linear approximations may lead to suboptimal operational decisions. In addition, linearization requires careful balance between accuracy and computational efficiency for the grid fineness. Introducing additional variables—particularly in systems exhibiting pronounced nonlinear dynamics or complex component interactions—rapidly transforms linearization into a computationally demanding process.

Nonlinear optimization methods are less frequently adopted in RHA-based energy management due to their computational burden. Deterministic global solvers are generally impractical under real-time constraints because of exhaustive region enumeration and slow convergence, especially for large-scale nonconvex problems [6]. Stochastic global optimization techniques, such as genetic algorithms, have been explored to better capture nonlinearities and avoid local optima, particularly in

renewable and hybrid energy systems. Nevertheless, their high computational cost, extensive function evaluations, and slow convergence limit their suitability for high-dimensional, constrained, real-time applications [6].

To overcome the computational and convergence limitations of conventional stochastic optimization methods, this work incorporates the recently proposed BO–IPOPT algorithm into the energy management framework for industrial systems. BO–IPOPT is a hybrid approach that combines the global exploration capability of Bayesian optimization with the efficient local convergence and constraint handling of IPOPT [9, 10]. Previous studies have shown that this method can reliably solve large-scale, nonlinear, and constrained optimization problems, delivering significantly improved solution quality and robustness compared to existing approaches under comparable computational effort [6, 11].

In this work, BO–IPOPT is applied to the energy management of an industrial hydrogen production system that includes a 40 MW alkaline electrolyzer (AEL), photovoltaic (PV) generation, grid electricity, an organic Rankine cycle (ORC) unit for waste heat recovery, a high-pressure hydrogen (H_2) storage, a configuration that is inherently nonlinear. The study benchmarks BO–IPOPT against state-of-the-art optimization methods to evaluate its performance in terms of solution quality, robustness, and computational efficiency under realistic operating conditions. Additionally, the investigation addresses the impact of solar data uncertainties on energy management decisions, examining how prediction errors in PV generation influence operational strategies and the overall system performance.

This paper is organized as follows: Section 2 introduces the considered industrial system and Section 3 the optimization methods whose performance is investigated in this work. In Section 4, the prediction model and strategy for multi-step-ahead forecasting are presented. Section 5 shows the influence of different settings on the energy management performance of the industrial system. Finally, Section 6 summarizes the main conclusions and provides a perspective for future work.

USE CASE

In this section, the underlying industrial hydrogen production system is briefly introduced as well as the corresponding optimization problem.

System Description

The industrial system illustrated in **Figure 2** is built around a 40 MW pressurized AEL that supplies hydrogen to a refinery. Its main power source is a PV plant, complemented by electricity generated from refinery waste heat via an ORC unit. When required, additional power can be drawn from the grid as an auxiliary energy source. The AEL comprises eight stacks, each operating at a pressure of 20-25 bar and consuming 5 MW when running at full capacity. The produced hydrogen is compressed and stored in a high-pressure system operating at 350 bar before being delivered to the refinery. The overall operation ensures that the refinery's hydrogen demand is continuously met at each time step.

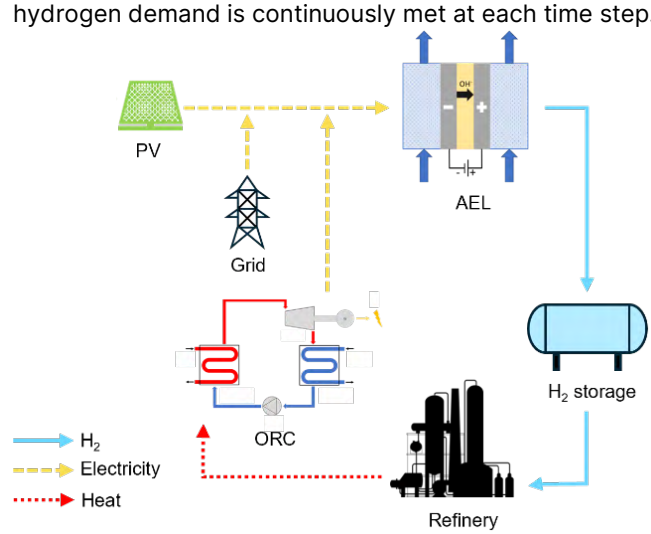

Figure 2. Visualization of the industrial hydrogen production system.

Optimization Problem

For the cost-optimal operation of the industrial hydrogen production system, a constrained, nonlinear, and nonconvex optimization problem is formulated within a multi-period framework. The optimization horizon is discretized into n uniform time steps with the duration Δt, such that the decision variables are evaluated at discrete time points t_k, with $k = 1, \dots, n$. The objective is to minimize the total operational costs consisting of grid electricity costs (first term) and AEL startup costs (second term):

$$\min \sum_{k=1}^{n} P_{\text{grid}}^{k} g_{\text{grid}}^{k} \Delta t + \sum_{k=1}^{n} \sum_{s=1}^{N_{\text{stack}}} g_{\text{su}}^{k,s} z_{\text{su}}^{k,s}, \tag{1}$$

where P_{grid}^{k} represents the grid electricity consumption, g_{grid}^{k} the grid electricity price, $g_{\text{su}}^{k,s}$ the cost per startup of an AEL stack, and $z_{\text{su}}^{k,s}$ a binary variable denoting the startup of a stack, with s describing the AEL stack indices. At each time step, the power balances among all generation and consumption units are maintained as:

$$P_{\text{PV}}^{k} + P_{\text{ORC}}^{k} + P_{\text{grid}}^{k} = P_{\text{e}}^{k} + P_{\text{comp}}^{k}, \tag{2}$$

where P_{PV}^{k} and P_{ORC}^{k} denote the PV and ORC power generation, P_{e}^{k} the AEL consumption, and P_{comp}^{k} the power consumption for hydrogen compression modeled based on [12]. The hydrogen production rate of each AEL stack is given by Faraday's law as:

$$\dot{m}^{k,s}_{\text{prod, stack}} = \frac{\eta^{k,s}_{\text{Faraday}} I^{k,s}_{\text{cell}}}{2F} M_{H_2} N_{\text{cell}} z^{k,s}_{\text{on}}, \quad (3)$$

where η_{Faraday} is the Faraday efficiency, F the Faraday constant, I_{cell} the cell current, N_{cell} describes the number of cells in the stack, and M_{H2} denotes the molar mass of hydrogen [13]. The total hydrogen production is the summation over all stacks:

$$\dot{m}^{k}_{\text{prod}} = \sum_{s=1}^{N_{\text{stack}}} \dot{m}^{k,s}_{\text{prod, stack}}. \quad (4)$$

The AEL power consumption P^k_{e} is modeled according to [13]. In addition, the AEL operates in on/off/standby modes, which are enforced through AEL operational constraints. The PV power generation is modeled as a function of solar irradiance and ambient temperature:

$$P^k_{\text{PV}} = A_{\text{PV}} G^k_{\text{PV}} \eta_{\text{PV, inv}} \eta^k_{\text{PV, mod}} (1 - x^k_{\text{PV, cur}}), \quad (5)$$

where A_{PV} is the module area, G^k_{PV} denotes the solar irradiance, $\eta_{\text{PV, inv}}$ the inverter efficiency, and $\eta^k_{\text{PV, mod}}$ the module efficiency. The latter depends on its part-load behavior, nominal module efficiency, nominal solar irradiance, ambient temperature, and nominal temperature [6]. The variable $x^k_{\text{PV, cur}}$ accounts for the restriction or deactivation of the PV operation. All PV model parameters are taken from [6].

The ORC unit is represented by a surrogate model as a quadratic function of its hot-source temperature $T^k_{\text{ORC,source}}$ and ambient temperature T^k_{amb}:

$$P^k_{\text{ORC}} = f(T^k_{\text{ORC,source}}, T^k_{\text{amb}}). \quad (6)$$

The dynamics of the hydrogen storage are modeled by the differential mass balance:

$$m^k_{\text{sto}} = m^{k-1}_{\text{sto}} + (\dot{m}^k_{\text{sto, in}} - \dot{m}^k_{\text{sto, out}})\Delta t, \quad (7)$$

where m^k_{sto} is the amount of hydrogen stored, $\dot{m}^k_{\text{sto, in}}$ denotes the hydrogen mass flow entering the storage, and $\dot{m}^k_{\text{sto, out}}$ the hydrogen mass flow withdrawn from the storage. Additional constraints prevent simultaneous charging and discharging, while $\dot{m}^k_{\text{sto, in}}$ and $\dot{m}^k_{\text{sto, out}}$ are limited by the maximum inlet and outlet flowrates of the storage. The resulting formulation describes a nonlinear, nonconvex optimization problem with continuously differentiable constraints, avoiding binary or integer variables through smooth approximations in the hydrogen storage and the AEL, added as penalty terms in the objective function.

This optimization problem employs two different time resolutions to accommodate the available input data characteristics. Grid electricity consumption, PV, storage dynamics, and AEL operation are controlled at 15-minute intervals to capture short-term fluctuations, while the ORC power output follows an hourly resolution by maintaining the heat source temperature $T^k_{\text{ORC,source}}$ constant over each hour (via additional equality constraints). The 15-minute timestep aligns with common intraday electricity market trading practices [14]. The optimization requires input data including electricity prices, solar irradiance, ORC heat source temperature, hydrogen demand profile, and ambient temperature, with only the solar data forecasted in this work while all other inputs are treated as known.

OPTIMIZATION METHODS

To solve the aforementioned constrained, nonlinear, and nonconvex optimization problem at each iteration of the RHA, several optimization algorithms are examined. These include one widely known stochastic approach, namely GA; the local deterministic optimizer IPOPT applied in single-start form (with and without warm-starting); and the recently proposed hybrid algorithm BO-IPOPT. Although local deterministic solvers are not commonly employed within RHA-based energy management frameworks, their performance is explored here due to their broad applicability and strong efficiency in solving large-scale problems. Owing to the size and complexity of the optimization problem, global deterministic solvers such as BARON are not considered, as they become computationally impractical for problems of this scale (over 100 decision variables) [10]. The following subsections briefly describe the optimization methods evaluated in this study.

Stochastic Method

GA is a well-established stochastic optimization technique inspired by natural processes. GA [15] operates by evolving a population of candidate solutions through selection, crossover, and mutation operators. Constraints are typically addressed by incorporating equality and inequality violations into the objective function via penalty terms [16]. While GA is well suited for exploring large and complex search spaces and capable of escaping local optima, it generally requires a large number of function evaluations, leading to high computational effort and long convergence times, particularly in complex optimization settings.

Deterministic Methods

IPOPT [17] is a deterministic optimization algorithm designed for large-scale nonlinear programming problems. It employs an interior-point framework to enforce constraints and efficiently converges to locally optimal solutions. IPOPT is widely used in scientific and engineering applications due to its robustness and computational efficiency; however, as a local solver, it may converge to suboptimal local minima depending on the initial conditions. Warm-start (WS) strategies are often applied to accelerate convergence by exploiting information from previous solutions, especially when a sequence of related optimization problems is solved, as in RHA-based

approaches [18]. Nevertheless, inappropriate or poorly chosen warm-start information can adversely affect the solution quality.

Hybrid Method

BO-IPOPT is a hybrid optimization method originally proposed in [9] and further developed in [10], which integrates Bayesian optimization (BO) with IPOPT. The method combines BO's capability for global exploration with IPOPT's strength in local refinement and constraint handling. In this optimizer, BO constructs a surrogate model—based on Gaussian processes—to approximate both the objective function and constraints, guiding the search toward promising regions of the decision space without requiring a highly accurate global surrogate. Even if the surrogate model does not fully capture the objective landscape in high-dimensional spaces, its role in BO-IPOPT is primarily to guide the search, while its impact on the performance of the hybrid method will be further investigated in future work. Constraint handling within BO is achieved using an augmented Lagrangian formulation, where inequality constraints are converted into equalities through slack variables, resulting in a sequence of simpler optimization subproblems. IPOPT is then employed to perform local optimization within these regions while efficiently enforcing the constraints. By alternating between BO-driven global exploration and IPOPT-based local refinement, the BO-IPOPT framework is able to effectively address high-dimensional, constrained optimization problems, overcoming many of the limitations associated with purely stochastic or purely deterministic approaches.

UNCERTANTIES IN SOLAR DATA

Solar irradiance is inherently uncertain due to its strong dependence on atmospheric conditions and cloud dynamics, directly affecting photovoltaic power generation and operational decision-making. Reliable short-term forecasting is therefore essential for the RHA considered in this work.

As discussed in [6], physical forecasting methods rely on atmospheric modeling, numerical weather prediction, and satellite data, resulting in high computational complexity and significant data requirements. Moreover, physical methods are generally more suitable for long-term prediction horizons (> 6 hours ahead), while data-based approaches have been shown to perform better for short-term forecasts of up to six hours ahead. These characteristics make physical models unsuitable for real-time and iterative optimization frameworks, and they are therefore not considered in this study. Instead, data-based forecasting methods are considered, as they leverage historical time-series data and require substantially lower computational effort.

In [6], several data-driven models were evaluated, including linear and nonlinear regression, statistical methods, and machine and deep learning approaches. The results demonstrated that linear regression provides the most favorable trade-off between forecasting accuracy and computational efficiency, especially for both one-step and multi-step ahead predictions. Consequently, linear regression is selected as the solar irradiance forecasting model in this work. Based on the feature selection presented in [6], the resulting multiple linear regression model (MLR) used to predict solar irradiance is expressed as:

$$G_t = f(G_{t-1}, G_{t-2}, G_{t-3}, G_{t-4}, G_{t-\text{day}}, G_{t-\text{year}}, SD_{t-\text{day}}, T_{t-\text{day}}, p_{t-\text{day}}), \quad (8)$$

where G denotes the solar irradiance, SD the sunshine duration, T the air temperature at two meters, and p the air pressure station level.

To propagate uncertainty over the prediction horizon, a recursive multi-step forecasting strategy is employed. Following the approach described in [6], a single one-step-ahead model is trained and applied iteratively, with each predicted value fed back as an input for the subsequent time step. This process continues until the full forecasting horizon is covered. The one-step-ahead forecast at time t is given by:

$$\hat{G}_t = f(G_{t-1}, G_{t-2}, G_{t-3}, G_{t-4}, G_{t-\text{day}}, G_{t-\text{year}}, SD_{t-\text{day}}, T_{t-\text{day}}, p_{t-\text{day}}), \quad (9)$$

while the two-step-ahead forecast is obtained recursively as:

$$\hat{G}_{t+1} = f(\hat{G}_t, G_{t-1}, G_{t-2}, G_{t-3}, G_{t+1-\text{day}}, G_{t+1-\text{year}}, SD_{t+1-\text{day}}, T_{t+1-\text{day}}, p_{t+1-\text{day}}). \quad (10)$$

ENERGY MANAGEMENT

This section introduces the scenarios used for the energy management of the industrial hydrogen production system and the corresponding results.

Scenarios

In this work, two one-week solar data scenarios are considered, representing winter and summer conditions with lower and higher solar irradiation respectively, with a 15-minute resolution. Moreover, the electricity price g_{grid} is fixed at 0.12 €/kWh and the price g_{su} per startup event in each AEL stack at 435 €. The chosen scenarios and values are representative for the considered refinery.

Results

The performance of the energy management strongly depends on the selected optimization algorithm. In this section, we assess the stochastic algorithm "GA",

IPOPT in single-start mode (with and without warm-starting), and the hybrid BO-IPOPT within the RHA-based energy management framework. All methods are implemented in Python 3.12; IPOPT is accessed via Pyomo with default settings, GA uses parameter values recommended in [19], and BO-IPOPT follows the configuration proposed in [10], with the four best candidates selected per outer iteration to match the four available CPU cores. To enhance computational efficiency, the IPOPT evaluations within BO-IPOPT are executed in parallel on four cores using the multiprocessing library, whereas the stochastic search component of BO-IPOPT is kept sequential to ensure a fair comparison with GA. All simulations are performed on a machine equipped with an Intel(R) Core(TM) i7-8665U processor.

In this study, the optimizer CPU running time denotes the fixed time budget assigned to GA and BO-IPOPT at each RHA iteration to compute the best available minimum, i.e., the optimal trajectories for the system setpoints. By contrast, IPOPT and WS-IPOPT are started only once and perform a local search, typically terminating before the end of the fixed time window (average termination time about 40 seconds). Since all considered algorithms contain inherent randomness (e.g., random initializations, sampling), each numerical experiment is repeated 10 times, and the reported results are based on these repeated runs to average out stochastic effects. As a reference configuration, we employ a 24-step optimization horizon, corresponding to 6 hours with 15-minute intervals, resulting in an optimization problem with 4, 248 decision variables. This horizon strikes a balance between incorporating enough future steps for robust operational planning and the accuracy of data-based solar forecasts (limited to ≤ 6 hours ahead). Since the true global optimum is unknown, the best solution obtained among BO-IPOPT and GA with a CPU time of 5 minutes per RHA iteration is used as a benchmark for evaluating all algorithms.

We first investigate the effect of the different optimizers when included in the RHA, neglecting any uncertainty in the input data to ensure a fair comparison. In agreement with the refinery, the CPU time for each optimizer is fixed to 50 seconds per RHA iteration, which was chosen so that the resulting control actions and system responses can be updated sufficiently fast for practical implementation. The optimization outcomes for the winter and summer weeks are summarized in **Figure 3**, which presents box plots of the minimum accumulated objective function values for each method over the repeated trials. The results indicate that BO-IPOPT achieves the lowest operating costs and the tightest spread of values among all algorithms. On average (relative to the median), BO-IPOPT attains a relative error of about 0.01 % with respect to the benchmark in the winter week, whereas GA, IPOPT, and WS-IPOPT exhibit errors of approximately 314 %, 3 %, 0.2 %, respectively. For the summer week, the average error of BO-IPOPT is about 0.2 %, while GA, IPOPT, and WS-IPOPT show relative errors of around 505 %, 4 %, 2 %, respectively.

BO-IPOPT is the only method that consistently reaches the benchmark solution in both weeks. As expected, IPOPT and WS-IPOPT show worse performance than BO-IPOPT due to their single-start limitations and suboptimal warm-starts, which lead to wide variability in objective values, higher operating costs, and the need for multiple restarts.

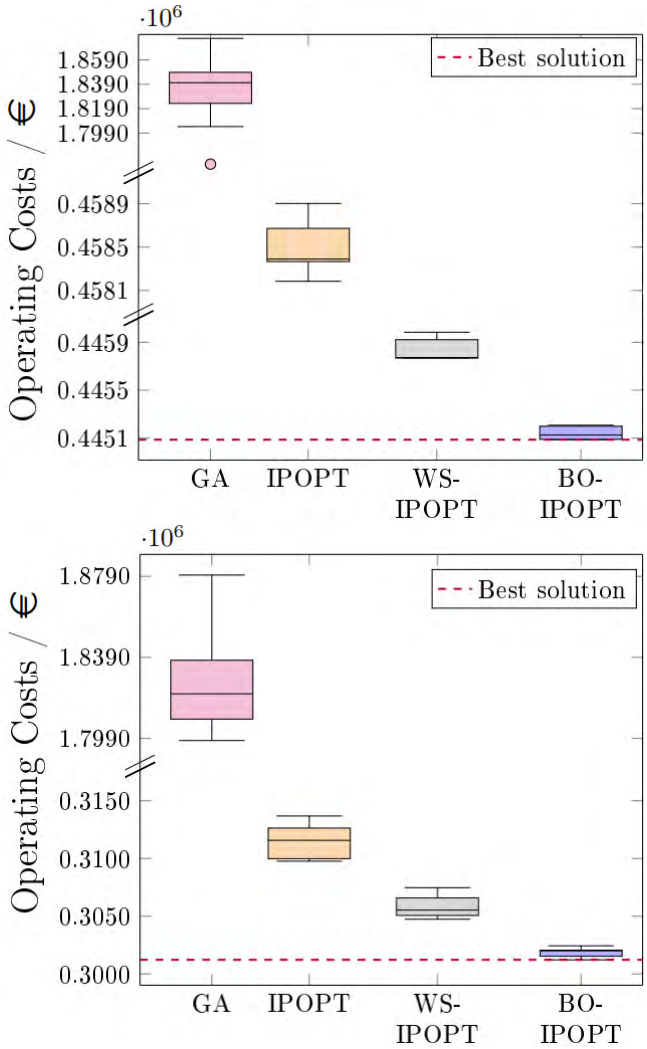

Figure 3. Comparison of operating costs over 10 trials for different optimizers in winter (**top**) and summer (**bottom**). The circles in the box plots represent outliers.

Among single-start methods, WS-IPOPT delivers lower objective function values and a tighter range of results than IPOPT, benefiting from more informed starting guesses that reduce sensitivity to starting points and enhance consistency across repeated runs. The stochastic optimizer "GA" encounters significant challenges: its population-based mechanism and random perturbations

demand numerous iterations to find feasible, low objective function solutions. Within the tight per-RHA-iteration time limits, it fails to adequately search the space or address constraint violations—yielding infeasible results with violations on the order of 10^{12} and much higher operating costs than the other solvers.

It is worth noticing that higher solar irradiation in summer makes the optimization problem more challenging, adding system variability and operational flexibility, thus widening performance gaps between the solvers.

Although IPOPT and WS-IPOPT are single-start local solvers, they frequently fail to converge on the first try, requiring restarts. In this system, IPOPT needed them in 25 % of winter RHA iterations and 43 % in summer. WS-IPOPT improved to 3 % (winter) and 12 % (summer) restart rates. This further supports the advantage of BO-IPOPT, which improves both the solution quality and the convergence to a feasible solution at each RHA iteration.

Although this analysis assumes perfect input data (no uncertainties), **Figure 4** demonstrates that solar uncertainties handled via MLR remain close to the perfect case, with 0.2 % cost deviations in the winter week and 1.6 % cost deviations in the summer week—making this forecasting model suitable for real-time applications (CPU training time < 1 second).

CONCLUSION

This study evaluates the effectiveness of the novel hybrid BO-IPOPT method in the energy management of an industrial hydrogen production system introduced in this work. We first compared BO-IPOPT against state-of-the-art optimizers—IPOPT, WS-IPOPT, and GA—assessing their accuracy and robustness. BO-IPOPT consistently outperformed the alternatives, achieving higher accuracy and robustness, without facing any convergence issues or constraint violations. IPOPT and WS-IPOPT underperformed due to reliance on single-start strategies and suboptimal WS initialization. In addition, both solvers often struggled with convergence failures, frequently needing multiple restarts at individual RHA iterations. Beyond these difficulties, both deterministic methods showed substantial variability in the objective values across repeated runs compared to BO-IPOPT. Although WS-IPOPT offered better accuracy and robustness than IPOPT by reducing the restart frequency and outcome variability, these issues were not fully resolved. Moreover, GA failed to satisfy the system's constraints within the defined computational limits, compromising solution quality.

Another important observation from this energy management analysis is that scenarios with higher solar power availability made the optimization problem more challenging due to greater system variability and operational flexibility. Thus, employing a reliable and efficient optimizer like BO-IPOPT proves especially critical in such cases. In addition, uncertainties in solar data have also been incorporated into this work. The simple MLR delivered consistently robust results with very low computational demands. Across both solar scenarios, the linear model limited the operating cost deviations up to 1.6 % relative to perfect foresight cases, rendering it highly suitable for real-time applications requiring rapid retraining.

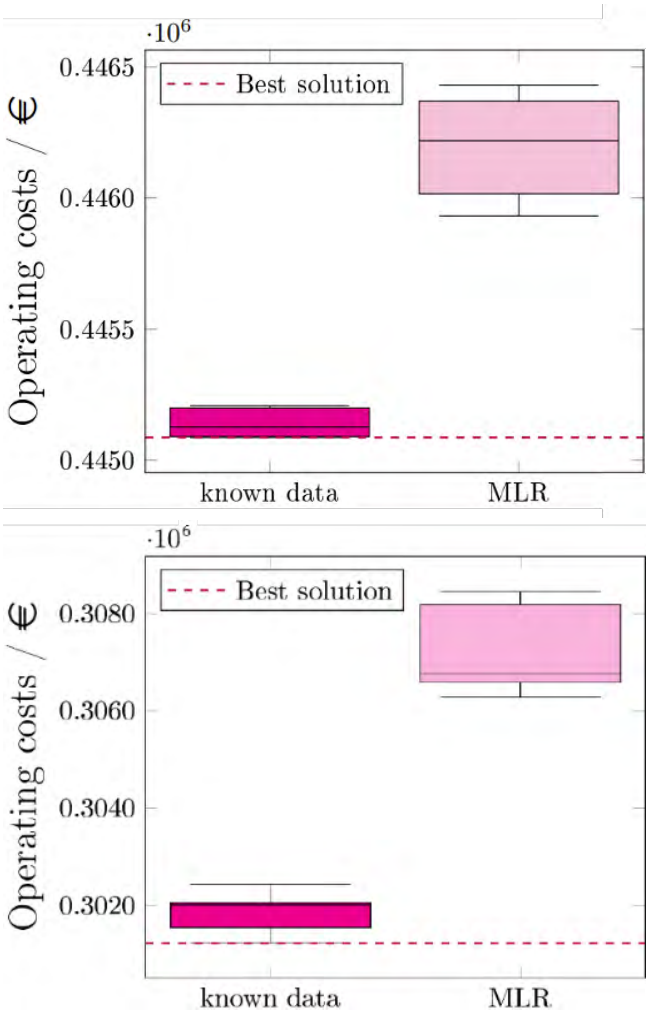

Figure 4. Comparison of operating costs over 10 trials using BO–IPOPT with and without uncertainties modeled via MLR for the solar data in winter (**top**) and summer (**bottom**).

Future work will first benchmark BO-IPOPT against Gurobi on a linearized system using experimental data to determine the superior approach, and investigate hybrid methods such as using Gurobi solutions to initialize BO-IPOPT. Next, different optimization horizons will be investigated to balance forecasting accuracy of input data (which decreases with longer horizons) with computational feasibility (which may worsen as the problem size grows with the horizon length) and predictive planning

(larger horizons help incorporate future conditions into the optimal operating strategy). In addition, we aim to incorporate variable electricity prices alongside uncertainties in hydrogen demand and ORC source temperature. Moreover, implementation in a Hardware-in-the-Loop or real plant setup will then evaluate closed-loop performance with real-time sensor feedback. Finally, the performance of BO-IPOPT will be assessed on even larger systems in future research, including the impact of its surrogate accuracy on the method's performance.

ACKNOWLEDGEMENTS

The authors acknowledge the funding provided through the Establishing **E**uropean **P**roduction of **HY**drogen **R**enew**A**ble energy and integration into an industrial environment (EPHYRA) project, supported by the Clean Hydrogen Partnership in collaboration with Hydrogen Europe and Hydrogen Europe Research, under Grand Agreement no. 101112220.

AUTHOR IDENTIFIERS

Loukas Kyriakidis: 0009-0003-6634-8579
Jonaed Bin Mustafa Kamal: 0009-0001-8155-0471
Saskia Bublitz: 0000-0002-9050-7862
Bogdan Dorneanu: 0000-0003-3553-6625
Harvey Arellano-Garcia: 0000-0002-8297-0232

REFERENCES

1. Lund H, Werner S, Wiltshire R, Svendsen S, Thorsen JE, Hvelplund F, Mathiesen BV. 4th generation district heating (4GDH). Energy 68:1-11 (2014). https://doi.org/10.1016/j.energy.2014.02.089
2. Olatomiwa L, Mekhilef S, Ismail MS, Moghavvemi M. Energy management strategies in hybrid renewable energy systems: a review. Renewable and Sustainable Energy Reviews 62:821-835 (2016). https://doi.org/10.1016/j.rser.2016.05.040
3. Khalid M. Smart grids and renewable energy systems: perspectives and grid integration challenges. Energy Strategy Reviews 51:101299 (2024). https://doi.org/10.1016/j.esr.2024.101299
4. Shanmugarjah V, Lidula NWA, Rajapakse AD. A comparative analysis of rule-based and optimization-based energy management for a campus microgrid in sri lanka. 2024 International Conference on Sustainable Energy: Energy Transition and Net-Zero Climate Future (ICUE) :1-6 (2024). https://doi.org/10.1109/icue63019.2024.10795646
5. Hu N. The limitations of traditional PID controllers and modern optimization methods. ACE 147:238-244 (2025). https://doi.org/10.54254/2755-2721/2025.22912
6. Kyriakidis L, Kansara R, Roldán Serrano MI. Energy management of industrial energy systems via rolling horizon and hybrid optimization: a real-plant application in germany. Energies 18:3977 (2025). https://doi.org/10.3390/en18153977
7. Bischi A, Taccari L, Martelli E, Amaldi E, Manzolini G, Silva P, Campanari S, Macchi E. A rolling-horizon optimization algorithm for the long term operational scheduling of cogeneration systems. Energy 184:73-90 (2019). https://doi.org/10.1016/j.energy.2017.12.022
8. Corinaldesi C, Schwabeneder D, Lettner G, Auer H. A rolling horizon approach for real-time trading and portfolio optimization of end-user flexibilities. Sustainable Energy, Grids and Networks 24:100392 (2020). https://doi.org/10.1016/j.segan.2020.100392
9. Kyriakidis L, Mendez MA, Bähr M. A hybrid algorithm based on bayesian optimization and interior point optimizer for optimal operation of energy conversion systems. Energy 312:133416 (2024). https://doi.org/10.1016/j.energy.2024.133416
10. Kyriakidis L, Bähr M, Mendez MA. Enhanced hybrid algorithm based on bayesian optimization and interior point optimizer for constrained optimization. Optim Eng : (2025). https://doi.org/10.1007/s11081-025-09975-y
11. Kyriakidis, L., & Bähr, M. ENERGY MANAGEMENT OF AN INDUSTRIAL POWER-TO-HEAT SYSTEM USING A HYBRID OPTIMIZATION ALGORITHM. *Proc. ECOS 2025*:1-12 (2025, accepted)
12. Zheng Y, You S, Bindner HW, Münster M. Optimal day-ahead dispatch of an alkaline electrolyser system concerning thermal–electric properties and state-transitional dynamics. Applied Energy 307:118091 (2022). https://doi.org/10.1016/j.apenergy.2021.118091
13. Sánchez M, Amores E, Rodríguez L, Clemente-Jul C. Semi-empirical model and experimental validation for the performance evaluation of a 15 kw alkaline water electrolyzer. International Journal of Hydrogen Energy 43:20332-20345 (2018). https://doi.org/10.1016/j.ijhydene.2018.09.029
14. Corinaldesi C, Schwabeneder D, Lettner G, Auer H. A rolling horizon approach for real-time trading and portfolio optimization of end-user flexibilities. Sustainable Energy, Grids and Networks 24:100392 (2020). https://doi.org/10.1016/j.segan.2020.100392
15. Hayes-Roth F. Review of "adaptation in natural and artificial systems by john H. holland", the U. of

michigan press, 1975. SIGART Bull. :15-15 (1975). https://doi.org/10.1145/1216504.1216510
16. Long Q, Wu C. A hybrid method combining genetic algorithm andhooke-jeeves method for constrained global optimization. JIMO 10:1279-1296 (2014). https://doi.org/10.3934/jimo.2014.10.1279
17. Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math. Program. 106:25-57 (2005). https://doi.org/10.1007/s10107-004-0559-y
18. Olivares DE, Lara JD, Canizares CA, Kazerani M. Stochastic-predictive energy management system for isolated microgrids. IEEE Trans. Smart Grid 6:2681-2693 (2015). https://doi.org/10.1109/tsg.2015.2469631
19. Solgi RM. https://pypi.org/project/geneticalgorithm/

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Real-Time Estimation and Optimal Control of Supersaturation in Sugar Crystallization using Model-based Soft Sensor

Ananya Lohani[a]*, Adam Fedor[a], Július Kurucz[b], and Radoslav Paulen[a]

[a] Faculty of Chemical and Food Technology, Slovak University of Technology in Bratislava, 812 37 Bratislava, Slovakia
[b] FUZZY s.r.o., 925 81 Diakovce, Slovakia
* Corresponding Author: ananya.lohani@stuba.sk.

ABSTRACT

Maintaining mother liquor supersaturation at a setpoint within the metastable range is vital for achieving the best production yield in industrial sugar production. However, precise online measurement and control is challenging. In this work, we develop a model-based soft sensor for supersaturation monitoring, and we propose a new feedforward-feedback control structure for batch sugar crystallization. Supersaturation is estimated using standard process measurements, enabling a soft sensor that can be readily adapted to different production units. The soft sensor continuously estimates supersaturation from standard process signals, and the control strategy ensures it remains within the desired operating range, enabling simple and straightforward application to other sugar production units.

Keywords: Mass Balance, Energy Balance, Soft Sensor, Feedback Control, Feedforward Control, Supersaturation

INTRODUCTION

Sugar crystallization is a crucial unit operation in the sugar manufacturing industry, directly influencing the yield and quality of product, and overall process efficiency [1]. Supersaturation is one of the key variables determining crystallization performance, as it regulates crystal growth, nucleation behaviour, and dissolution phenomena. It must be maintained within the metastable zone, typically in the range of $1.00 < \sigma < 1.35$ for the best yield and quality of sugar [2]. Operating outside this window may result in the excessive nucleation, crystal agglomeration, or dissolution, leading to reduced sugar recovery and off-spec crystal size distribution.

Despite its significance, supersaturation is challenging to measure online in industrial crystallizers due to its reliance on several major factors, like temperature, concentration (Brix), purity of the mother liquor, and pressure [3]. As direct measurement is not viable, industrial practice largely relies on indirect indicators, operator experience, or exclusive instruments. At present, the SeedMaster series transmitter [4] is the only commercially available instrument that provides an online estimation of supersaturation. Its applicability is, however, limited by cost, proprietary algorithms, and lack of integration with advanced feedback control strategies.

Several research efforts have addressed the estimation and regulation of other hard-to-measure process properties, such as crystal size distribution and solution conductivity, which are closely linked to supersaturation and can be viewed as different expressions of the same crystallization behaviour, to improve the sugar yield quality using data-driven, model-based, and hybrid soft sensors. Data-driven approaches offer major advantages while capturing complex, non-linear process behaviour using historical or online plant data, thereby improving prediction accuracy without requiring detailed physical descriptions. Leveraging these benefits, Devogelaere et al. [5] demonstrated the use of feedforward neural networks as soft sensors for predicting massecuite electrical conductivity that serves as an indicator of supersaturation, supporting real-time control setpoint adjustments. Likewise, Lu et al. [6] presented an online modeling technique based on multi-input multi-output partial least-

squares support vector machines, achieving simultaneous estimation of massecuite brix and supersaturation.

Model-based approaches contribute to improved interpretability by incorporating fundamental physical relationships into the process description. Bakir et al. [7] proposed a continuous-discrete observer for estimating crystal size distribution in batch crystallization using a reduced population balance equation model. This formulation enables real-time monitoring while supporting control of filterability and flowability. Similarly, Damour et al. [2] introduced a model-based soft sensor for C-stage crystallization, providing online estimates of crystal mass, sucrose concentration, and purity aimed at maximizing sugar exhaustion.

Despite the advantages of data-driven and model-based approaches, their performance is often limited in crystallization processes involving high impurity levels and complex kinetic behaviour. Model-based approaches often suffer from parameter uncertainty, whereas data-driven models typically lack robustness, physical interpretability, and extrapolation capability beyond their training domain. Hybrid modeling frameworks address these limitations by merging mechanistic equations with data-driven components, enabling accurate representation of complex phenomena with limited data requirements. Georgieva et al. [8] proposed a knowledge-based hybrid model for fed-batch crystallization, integrating mass, energy, and population balance equations with artificial neural networks to accurately describe growth, nucleation, and agglomeration processes. Similarly, Meng et al. [9] demonstrated the applicability of data-driven and hybrid models for online estimation of supersaturation, mother liquor purity, crystal content, and crystal size distribution.

Despite the significant progress in data-driven, model-based, and hybrid soft-sensing techniques, their utilization in closed-loop control structures remains limited in industrial environments. The complexity of crystallization kinetics, coupled with uncertainty in process parameters and industrial implementation constraints, has limited the transition from monitoring to active control. Consequently, industrial sugar crystallization processes continue to rely mainly on conventional PI controllers for regulation of level and temperature, while the true control objective—maintaining supersaturation within the metastable zone- is only indirectly addressed through these variables. To the best of the authors' knowledge, no recorded studies have implemented a feedback control framework that explicitly manages supersaturation within the metastable range using model-based estimation. This represents a significant gap, particularly given the increasing demand for improved energy efficiency, higher extraction rates, and consistent product quality in sugar factories [10].

In this work, a model-based soft sensor is utilized to explicitly estimate supersaturation using mass and energy balance equations, incorporating key process variables such as purity, crystal mass, temperature, and concentration. Furthermore, supersaturation is directly regulated through a combined feedback-feedforward control strategy based on a standard PID controller, with the feed flow rate and heating steam supply rate considered as the manipulated inputs. The strength of this study lies in the practical use of simple, well-established control techniques to explicitly estimate and control supersaturation, thereby bridging the gap between theoretical approaches and industrial implementation. Overall, the proposed framework enables online supersaturation estimation and regulation without reliance on specialized measurement systems, offering a practical and industry-ready solution.

THEORY

This section is focused on the development of a model-based soft sensor of supersaturation using mass and energy balance equations.

Mass and Energy Balance Equations

The C-stage crystallization process is modelled using the set of five ordinary differential equations, with only the operating conditions and parameter values adjusted to reflect the C-stage operation [2], [12].

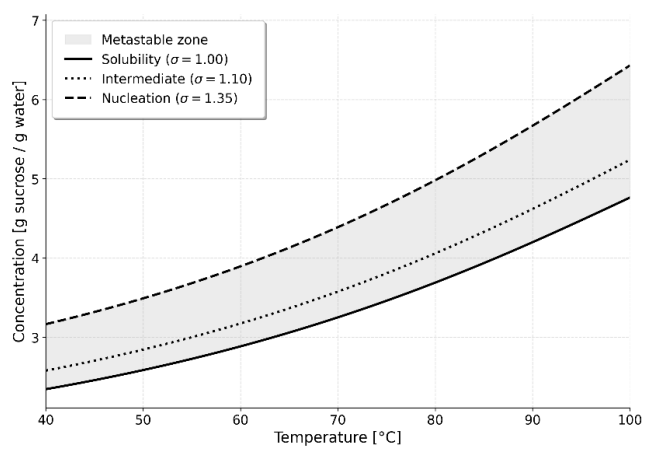

Figure 1. Temperature-dependent solubility curve of sucrose [13]

$$\frac{dm_w}{dt} = \rho_{feed} F_{feed}(1 - Bx_{feed}) - \dot{m}_{vap} \quad (1)$$

$$\frac{dm_i}{dt} = \rho_{feed} F_{feed} Bx_{feed}(1 - Pur_{feed}) \quad (2)$$

$$\frac{dm_c}{dt} = cc\left(\rho_{feed} F_{feed} - \dot{m}_{vap}\right) + \alpha_c \quad (3)$$

$$\frac{dm_s}{dt} = (\rho_{feed} F_{feed} Bx_{feed} Pur_{feed}) - \frac{dm_c}{dt} \quad (4)$$

$$\frac{dT_{mag}}{dt} = \frac{W + Q_{hs} + \rho_{feed} F_{feed}(h_{feed} - h_{ml}) - \dot{m}_{vap} L_{vap} + \frac{dm_c}{dt} L_c}{m_{total} Cp_{mag}} \quad (5)$$

Table 1. summarizes the typical operating conditions of an industrial sugar crystallization. Throughout the work, the subscript 'feed' refers to properties associated with the inlet feed stream. Here m_w, m_i, m_c and m_s denote the masses of water, impurities, crystal and dissolved sucrose, respectively. The feed density is depicted by ρ_{feed}, while F_{feed} represents the feed flow rate. The stirring power W is fixed at 15 kW. In addition, h denotes the specific enthalpy, L the specific latent heat, and Cp_{mag} the specific heat capacity of magma. $\dot{m}_{vap}$ represents mass flow rate of emitted vapour which is defined as $\dot{m}_{vap} = \rho_{cw} F_{vap}$, where ρ_{cw} is the density of condensed water and F_{vap} the emitted vapour mass flow rate, m_{total} the total mass in the magma and is defined as $m_{total} = m_c + m_w + m_i + m_s$, and Q_{hs} the heating power which is brought by the condensation of heating steam and defined as $Q_{hs} = \rho_{cw} F_{cs} L_{cw}$. All the thermal and physical correlations are mentioned below in the appendix. The RHS of eq (3) is the crystallization rate, which is calculated using the moment equations computed from the population balance equations, where cc is the crystal content which can be defined as $cc = m_c/m_{total}$ and α_c is an adjusted parameter according to crystal growth kinetic [14]. All the typical operating conditions are mentioned in Table 1.

Table 1: Typical operation conditions of a sugar plant [12]

Parameter	Description	Values
p_v	Vacuum Pressure	$0.12 - 0.15$ bar
Bx_{feed}	Feed Brix	$72 - 76°$ Bx
Bx_{mag}	Magma Brix	$88 - 95°$ Bx
Pur_{feed}	Feed Purity	$70 - 75\%$
Pur_{mag}	Magma Purity	$60 - 65\%$
Pur_{ml}	Mother Liquor Purity	$0 - 65\%$
T_{mag}	Magma Temperature	$60 - 75°\text{C}$
T_w	Water Temperature	$60 - 65°\text{C}$
T_{feed}	Feed Temperature	$60 - 65°\text{C}$

Direct measurement of gas flow rates in industrial equipment is typically infeasible. Consequently, both the heating steam and emitted vapour flow rates are commonly computed through indirect measurements using simplifying assumptions. The authors [2] suggest simplifying involved heat and mass transfer phenomena by introducing simplifications in a form of correlation factors α

1. First, the dynamics between heating steam supply flow rate F_{hs} and condensed steam F_{cs} is considered negligible with respect to the kinetics of the process. Thus:

$$F_{hs} = \alpha_Q F_{cs} \tag{6}$$

2. Second, the emitted vapor flow rates F_{vap} is proportional to the condensed steam flow rate F_{cs} via:

$$F_{vap} = \alpha_{vap} F_{cs} \tag{7}$$

The model framework is strictly governed by mass and energy conservation laws, utilizing standardized correlations (in Appendix) for the enthalpy and density of technical sugar solutions.

Supersaturation in crystallization processes

Supersaturation is described as the ratio between the mass of sucrose to the mass of water contained in a supersaturated solution ($s = m_s/m_w$) and the mass of the sucrose to the mass of the water which would be present in a saturated solution (s_{sat}) at the same temperature and purity [9], [15].

$$SS = \frac{s}{s_{sat} C_{sol}} \tag{8}$$

For a pure saturated solution, it is known that s_{sat} depends on temperature via the brix at saturation Bx_{sat}: [13]

$$s_{sat} = \frac{Bx_{sat}}{100 - Bx_{sat}} \tag{9}$$

However, considering the presence of impurities in industrial sucrose solutions, the s_{sat} is adjusted through the introduction of a solubility coefficient C_{sol}. The solubility coefficient is a parameter representing the impact of impurities on sucrose solubility. Several researchers have proposed correlations to calculate it [10], [12]. In this work, C_{sol} is determined by the well-known empirical correlation,

$$C_{sol} = a\frac{m_i}{m_w} + b + (1-b)e^{-c\frac{m_i}{m_w}} \tag{10}$$

where a, b and c are quality parameters of the mother liquor that depend on the amount and composition of local impurities.

Since supersaturation cannot be measured directly and is driven by the strongly interacting dynamics of concentration and temperature [11], the adopted formulation offers a practical and transparent estimation strategy suitable for industrial crystallization processes.

METHODOLOGY

The proposed model is validated through numerical simulations using synthetic datasets embedded with Gaussian noise to mimic industrial conditions. The flow rate of feed F_{feed} and heating steam supply F_{hs} are implemented as time-varying, piecewise-constant inputs. This strategy ensures controlled supersaturation while maintaining temperature and level % for model validation.

Given the complexity of the multivariable system, we determined the kinetic parameters using well-established thermodynamic definitions and empirical

correlations reported in the literature (Appendix)[2], [10]. To evaluate the sensitivity of the estimated α parameters to experimental noise, a Monte Carlo Simulation was performed [16]. For each iteration, the parameter set was re-estimated using a Levenberg-Marquardt least-squares minimization (8). This process allows the uncertainty and correlation between parameters to be analyzed.

Here for $\mathrm{Bx}_{ml,exp}$, $\mathrm{Bx}_{mag,exp}$, and $T_{mag,exp}$, we used simulated measurements with added gaussian noise to mimic experimental variability. The parameters are determined by minimizing the difference between these noisy "experimental" values and the corresponding model predictions.

$$\min_{\alpha_c,\alpha_{vap},\alpha_Q}\left(\frac{1}{2}\sum_{j=1}^{N}\left(\mathrm{Bx}_{ml,exp}^{j}-\mathrm{Bx}_{ml,sim}^{j}\right)^2+\frac{1}{2}\sum_{j=1}^{N}\left(\mathrm{Bx}_{mag,exp}^{j}-\mathrm{Bx}_{mag,sim}^{j}\right)^2+\frac{1}{2}\sum_{j=1}^{N}\left(T_{mag,exp}^{j}-T_{mag,sim}^{j}\right)^2\right) \tag{11}$$

To determine the optimal vector of coefficients, present in supersaturation equation $\theta=[a,b,c]$, a minimization problem is solved using Nelder-Mead optimization algorithm, based on values reported in previous literature. The goal is to align the model's calculated supersaturation at the final batch state with a target supersaturation value ($S_{target}=1.10$). The targeted supersaturation value is taken from industrial perspective. The objective function, $J(\theta)$, is defined as the sum of the squared error between the calculated and target supersaturation, plus a regularization term, which is added so to keep the estimated value closer to the values available in previous literature [12]:

$$J(\theta)=\left(S_{calc}(\theta)-S_{target}\right)^2+\lambda\sum_{i\in a,b,c}\left(\theta_i-\theta_{prior}\right)^2 \tag{12}$$

Where $S_{calc}(\theta)$ is the supersaturation computed using the trial parameters, λ is a penalty weight and θ_{prior} represents the original estimates for the coefficients which are taken from the previous literature [10].

Note that these parameters vary from plant to plant, depending on plant design and operating conditions.

Table 3: Model parameters and their values

Parameter	Values
α_{vap}	0.98
α_Q	1.65
α_c	0.29
a	-0.00573
b	1.00939
c	2.09860

Control strategy

Soft sensor will act as a feedforward part of the controller adjusting the supersaturation measurement with respect to the disturbances such as impurities, variations in temperature, etc. For feedback control procedure, a PID controller is implemented to regulate the supersaturation level inside the batch crystallizer. The control objective is to track a target supersaturation setpoint while compensating for disturbances induced by evaporation and non-linear solubility dynamics.

Here SS is the controlled variable, whereas F_{feed} is the manipulated variable. The supersaturation tracking error is defined as:

$$e(t)=SS_{target}-SS(t) \tag{13}$$

Where SS_{target} is the desired supersaturation setpoint.

The feed flow rate $F_{\mathrm{feed}}(t)$ is calculated using a standard PID algorithm. The control effort relies entirely on the error dynamics and the integral accumulation to sustain the flow.

$$F_{\mathrm{feed}}(t)=K_p e(t)+K_i\int_0^t e(\tau)\,d\tau+K_d\frac{de(t)}{dt} \tag{14}$$

Where K_p, K_i, and K_d are the proportional, integral, and derivative gains, respectively, and are tuned based on simulation studies [17], [18].

In practice, the controller is implemented in discrete time, with the integral and derivative terms in equation (15) approximated numerically at each sampling instant. The heating steam supply rate $F_{hs,forw}$ is applied as a feedforward control signal, based on typical operating settings. To ensure the physical realism, the actual flow rate of feed $F_{\mathrm{feed}}(t)$ applied to the crystallizer is subject to three specific constraints implemented in the simulation:

1. Integral Anti-Windup: To prevent excessive accumulation of error during saturation, the integral term is clamped to a fixed range chosen based on simulation studies and K_i value [19]:

$$-20\le\int_0^t e(\tau)\,d\tau\le 20 \tag{15}$$

2. Actuator Rate Limiting: The model limits how quickly the feed valve can open or close. The change in output (Δu) per time step is constrained by a maximum rate ($R_{max}=0.0001$) for smooth control of the manipulated input. [20]

$$|F_{\mathrm{feed}}(t)-F_{\mathrm{feed}}(t-\Delta t)|\le R_{max} \tag{16}$$

3. Output Saturation: The final flow rate is bounded by the physical limitations of the valve (minimum and maximum flow):

$$F_{\mathrm{feed},min}\le F_{\mathrm{feed}}(t)\le F_{\mathrm{feed},max} \tag{17}$$

4. Shutdown logic: To be in safety thermal limit and level%, the system calculates a dimensionless Shutdown factor, denoted as $\alpha(L)$, which is the function of the crystallizer level percentage $L(t)$. This factor is used to ramp down the steam injection and trigger a

hard stop for the feed.

$$\alpha(L) = \begin{cases} 1.0 & if\ L(t) < 98.0\% \\ \frac{99.5-L(t)}{99.5-98.0} & if\ 98.0\% \leq L(t) < 99.5\% \\ 0.0 & if\ L(t) \geq 99.5\% \end{cases} \quad (18)$$

The steam flow rate, which normally follows a step-change Feed forward control $F_{hs,forw}$, is directly multiplied by this factor. This gradually reduces evaporation energy as the tank fills, preventing overheating at the end of the batch:

$$F_{hs}(t) = F_{hs,forw}(t) \times \alpha(L) \quad (19)$$

The feed flow is controlled by the PID loop described above, but it is subject to a separate "Hard Safety Cutoff" strictly when the tank reaches maximum capacity:

$$F_{feed}(t) = \begin{cases} 0.0 & if L(t) \geq 100 \\ u_{PID}(t) & otherwise \end{cases} \quad (20)$$

Without such constraints, the controller becomes aggressive, producing control actions that violate physical realism. This can lead to excessive dilution or overheating, ultimately degrading crystak quality. Variations in F_{feed} affect the water content and dissolved sucrose concentration, which in turn alter the supersaturation level.

The control strategy was implemented during the batch crystallization, and its performance was evaluated by tracking the evolution of supersaturation, magma temperature, and mass components.

RESULTS

The simulation was conducted over a batch duration of 8 hours (28800 seconds) with a sampling integration step (Δt) of 10 seconds.

For α parameter sensitivity and robustness, a Monte Carlo simulation with 200 iterations was performed, introducing 2% gaussian noise to the nominal simulation data. The results (Figure 2) present the pairwise correlation plots, which illustrate the robustness of the estimated α parameters. Most of the 200 parameters estimates are tightly clustered around the true value, indicated by the red marker, showing that not only the Levenberg-Marquardt algorithm reliably avoids convergence to local minima but also that the parameters are well-identifiable.

A correlation analysis was conducted to examine the relationship among α_Q, α_{vap}, and α_c. The results indicate a moderate negative correlation between α_Q and α_{vap} $(r = -0.65)$, and a strong negative correlation between α_Q and α_c $(r = -0.93)$. In contrast, α_{vap} and α_c exhibit a strong positive correlation $(r = 0.88)$. The correlation results can be justified by looking at Eq. (3). The parameters mutually align to accurately track crystal mass accumulation. Table 2. suggests that the estimation procedure is numerically stable and that the identified parameters are robust with respect to uncertainty and noise in the data.

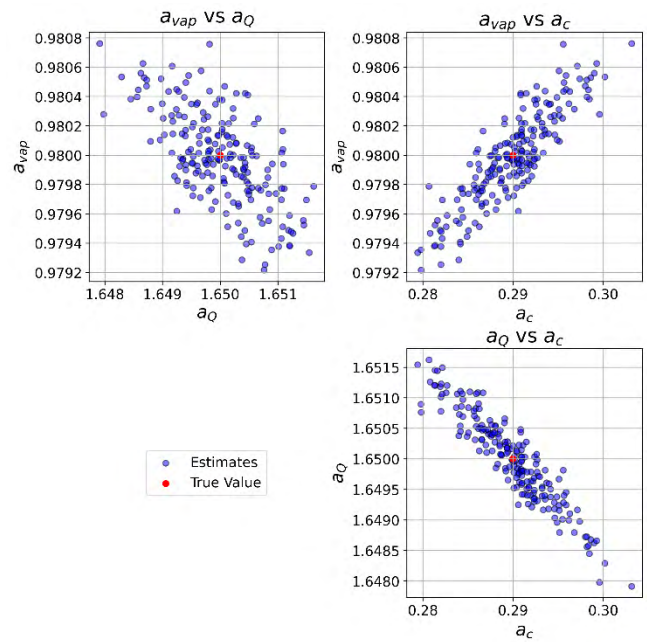

Figure 2. Pairwise parameter correlations ($N = 200$ and noise= 2%)

Table 2. Parameter estimation statistics

Parameter	Mean	Standard deviation
α_c	0.2899	0.0045
α_{vap}	0.9800	0.0003
α_Q	1.6500	0.0007

The impurity-dependent coefficients a, b, and c in the supersaturation were estimated via Nelder-Mead optimization algorithm (12). Estimated values are mentioned in Table 3. The resulting coefficients were subsequently validated by confirming that the simulated supersaturation dynamics and temperature-dependent solubility profiles aligned with consistent physical behaviour.

The controller performance was analyzed in terms of its ability to regulate supersaturation close to the desired operating region despite varying thermal conditions and changing liquor composition. Since the feedforward heating steam supply flow rate is not violating safety limits, hence the $\alpha(L) = 1$, which concludes that $F_{hs} = F_{hs,forw}$ here (Figure 3.)

The results (figure 3.) indicate that the control strategy provides smooth and stable trajectories, avoiding abrupt oscillations in supersaturation and temperature that could otherwise compromise crystal growth with the supersaturation profile remaining predominantly within the metastable region, indicating that the interaction

between feed addition and evaporation is well balanced. Brief disturbances are observed primarily during adjustment of steam flows; however, these deviations are short-lived and do not lead to uncontrolled nucleation or crystal dissolution. The results are further validated in the solubility curve Figure 3.

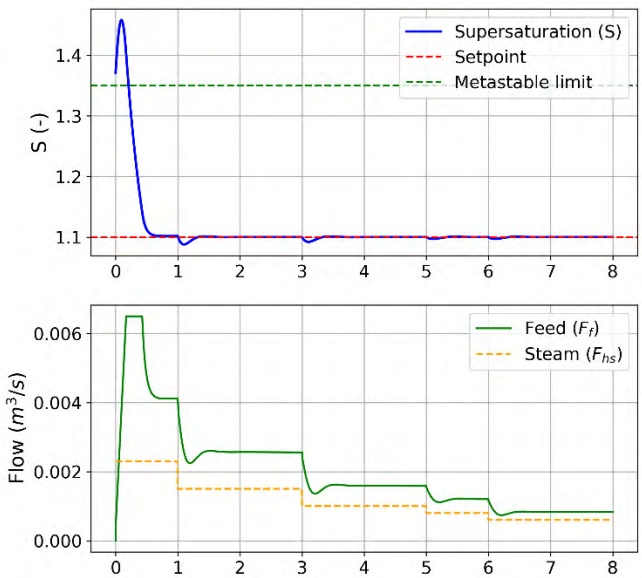

Figure 3. Control performance: supersaturation over time (top), manipulated variables (bottom). In addition, the controlled feed policy results in a gradual and consistent increase in crystal mass, reflecting favorable growth-dominated conditions. The absence of sharp oscillations in supersaturation and temperature suggests that the controller effectively mitigates process non-linearities and time-varying behavior inherent to batch crystallization.

Overall, the dynamic profiles indicate a stable operating regime that supports uniform crystal growth and improved product quality. The magma temperature varies within safe thermal limits across a wide operating range, indicating that the proposed soft sensor and control strategy operate effectively under different thermal conditions and maintain proper regulation of the energy balance through coordinated feed-flow control.

CONCLUSION AND FUTURE WORK

This study shows that supersaturation in batch sugar crystallization can be effectively regulated using constrained feedback control of the feed flow rate combined with feedforward steam supply. The proposed dynamic model captures the coupled mass and energy balances and maintains supersaturation within the metastable range under stable operation. Thermodynamic inconsistencies in commonly used parameter correlations were resolved, resulting in a physically consistent and numerically robust framework. Simulation results reproduce key industrial trends, indicating strong potential for practical application. Future work will focus on parameter estimation using real plant data to further improve predictive accuracy and assess robustness under realistic operating conditions.

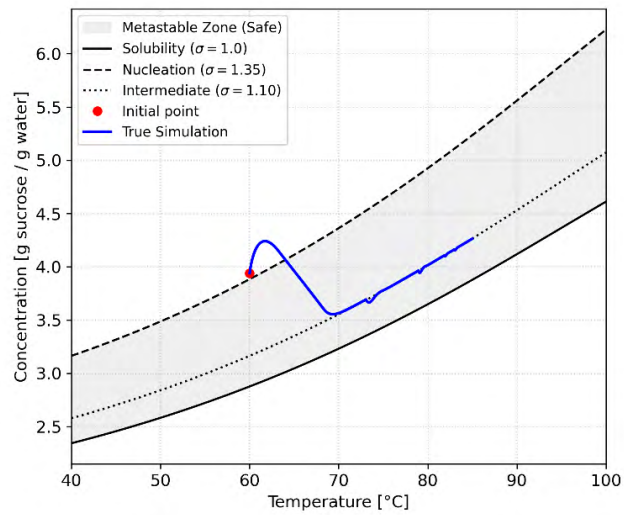

Figure 4. Feedback control effect on the solubility curve

APPENDIX

Brix and Purity for mother liquor and magma

$$\mathrm{Bx}_{ml} = \frac{m_i + m_s}{m_w + m_i + m_s}$$

$$\mathrm{Bx}_{mag} = \frac{m_i + m_{c_{m_s}}}{m_w + m_i + m_c + m_s}$$

$$\mathrm{Pur}_{ml} = \frac{m_s}{m_i + m_s}$$

$$\mathrm{Pur}_{mag} = \frac{m_c + m_s}{m_i + m_c + m_s}$$

Density (in $\mathbf{kg/m^3}$)

$$\rho_{\mathrm{feed}} = \left(1000 + \frac{(100Bx_{\mathrm{feed}})(200 + (100Bx_{\mathrm{feed}}))}{54}\right)\left(1 - 0.036\frac{T_{\mathrm{feed}} - 20}{160 - T_{\mathrm{feed}}}\right) - \left(1 - e^{(1-Pur_{\mathrm{feed}})\left(1.164\times10^{-4}(100Bx_{\mathrm{feed}}) + 6.927\times10^{-6}(100Bx_{\mathrm{feed}})^2\right)}\right)$$

$$\rho_{cw} = 1016.7 - 0.57T_w$$

Specific heat capacity (in $\boldsymbol{J/(kgK)}$)

$$Cp_{mag} = 4187 - 29.309Bx_{mag}$$

$$Cp_c = 1163.2 - 3.488T_{mag}$$

$$Cp_{ml} = 4186.8 - 29.7Bx_{ml} + 4.61Bx_{ml}Pur_{ml} + 0.075Bx_{ml}T_{mag}$$

$$Cp_{\text{feed}} = 4186.8 - 29.7Bx_{\text{feed}} + 4.61Bx_{\text{feed}}Pur_{\text{feed}} + 0.075Bx_{\text{feed}}T_{\text{feed}}$$

Specific enthalpy (in J/kg)

$$h_{\text{feed}} = Cp_{\text{feed}}T_{\text{feed}}$$

$$h_{ml} = Cp_{ml}T_{mag}$$

$$h_c = Cp_cT_{mag}$$

$$h_w = 2323.3 + 4106.7T_w$$

$$h_{vap} = 2499980 - 24186p_v + (1891.1 + 106.1p_v)T_v$$

Specific latent heat (in J/kg)

$$L_c = h_{ml} - h_c$$

$$L_{vap} = h_{vap} - h_w$$

Brix at Saturation

$$Bx_{sat} = 64.447 + 0.08222T_{mag} + 1.6169 \times 10^{-3}T_{mag}^2 - 1.558 \times 10^{-6}T_{mag}^3 - 4.63 \times 10^{-8}T_{mag}^4$$

ACKNOWLEDGEMENT

This work is funded by the Slovak Research and Development Agency under the project APVV-24-0007 and by the Scientific Grant Agency of the Slovak Republic under the grant 1/0263/25, and by the European Union under the grant scheme NextGenerationEU project no. 09I01-03-V05-00002.

REFERENCES

1. RICHARDSON JF, HARKER JH, BACKHURST JR. Crystallisation. Chemical Engineering :827-900 (2002). https://doi.org/10.1016/b978-0-08-049064-9.50026-6
2. Damour C, Benne M, Grondin-Perez B, Chabriat JP. Soft-sensor for industrial sugar crystallization: on-line mass of crystals, concentration and purity measurement. Control Engineering Practice 18:839-844 (2010). https://doi.org/10.1016/j.conengprac.2010.03.005
3. Damour C, Benne M, Boillereaux L, Grondin-Perez B, Chabriat JP. Multivariable linearizing control of an industrial sugar crystallization process. Journal of Process Control 21:46-54 (2011). https://doi.org/10.1016/j.jprocont.2010.10.002
4. L. Rózsa, J. Rózsa, Z. Ltd, S. Kilpinen, E. Mielonen, and K.-P. Oy, "Selection of the Operating Parameters in Sugar Crystallization Control".
5. Devogelaere D, Rijckaert M, Leon OG, Lemus GC. Application of feedforward neural networks for soft sensors in the sugar industry. VII Brazilian Symposium on Neural Networks, 2002. SBRN 2002. Proceedings. :2-6 (None). https://doi.org/10.1109/sbrn.2002.1181426
6. Meng Y, Yu X, He H, Tang Z, Wang X, Chen J. Knowledge-based modeling for predicting cane sugar crystallization state. International Journal on Smart Sensing and Intelligent Systems 7:942-965 (2014). https://doi.org/10.21307/ijssis-2017-689
7. Bakir T, Othman S, Puel F, Hammouri H. Continuous-discrete observer for crystal size distribution of batch crystallization process. Proceedings of the 44th IEEE Conference on Decision and Control :6240-6244 (None). https://doi.org/10.1109/cdc.2005.1583161
8. Georgieva P, Meireles MJ, Feyo de Azevedo S. Knowledge-based hybrid modelling of a batch crystallisation when accounting for nucleation, growth and agglomeration phenomena. Chemical Engineering Science 58:3699-3713 (2003). https://doi.org/10.1016/s0009-2509(03)00260-4
9. Meng Y, Yu S, Zhang J, Qin J, Dong Z, Lu G, Pang H. Hybrid modeling based on mechanistic and data-driven approaches for cane sugar crystallization. Journal of Food Engineering 257:44-55 (2019). https://doi.org/10.1016/j.jfoodeng.2019.03.026
10. Morales H, di Sciascio F, Aguirre-Zapata E, Amicarelli A. Crystallization process in the sugar industry: a discussion on fundamentals, industrial practices, modeling, estimation and control. Food Eng Rev 16:441-469 (2024). https://doi.org/10.1007/s12393-024-09377-3
11. Simoglou A, Georgieva P, Martin EB, Morris AJ, Feyo de Azevedo S. On-line monitoring of a sugar crystallization process. Computers & Chemical Engineering 29:1411-1422 (2005). https://doi.org/10.1016/j.compchemeng.2005.02.013
12. Morales H, Sciascio F, Aguirre-Zapata E, Amicarelli AN. A model-based supersaturation estimator (inferential or soft-sensor) for industrial sugar crystallization process. Journal of Process Control 129:103065 (2023). https://doi.org/10.1016/j.jprocont.2023.103065
13. Crestani CE, Bernardo A, Costa CBB, Giulietti M. Experimental data and estimation of sucrose solubility in impure solutions. Journal of Food Engineering 218:14-23 (2018). https://doi.org/10.1016/j.jfoodeng.2017.08.023
14. Hulburt HM, Katz S. Some problems in particle technology. Chemical Engineering Science 19:555-574 (1964). https://doi.org/10.1016/0009-2509(64)85047-8
15. Frew JA. Optimal control of batch raw sugar crystallization. Ind. Eng. Chem. Proc. Des. Dev. 12:460-467 (2002). https://doi.org/10.1021/i260048a013

16. Harrison RL, Granja C, Leroy C. Introduction to monte carlo simulation. AIP Conference Proceedings :17-21 (2010). https://doi.org/10.1063/1.3295638
17. "Development of the PID controller, " *IEEE Control Syst.*, vol. 13, no. 6, pp. 58–62, Dec. 1993, doi: 10.1109/37.248006.
18. Borase RP, Maghade DK, Sondkar SY, Pawar SN. A review of PID control, tuning methods and applications. Int. J. Dynam. Control 9:818-827 (2020). https://doi.org/10.1007/s40435-020-00665-4
19. "An analysis package comparing PID anti-windup strategies, " *IEEE Control Syst.*, vol. 15, no. 2, pp. 34–40, Apr. 1995, doi: 10.1109/37.375281.
20. Klyde D, McRuer D, Myers T. PIO analysis with actuator rate limiting. 21st Atmospheric Flight Mechanics Conference : (1996). https://doi.org/10.2514/6.1996-3432

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Data Reconciliation for Inventory Monitoring in a Petrol Refinery

Jakub Gaborčík[a*], Karol Ľubušký[b], and Radoslav Paulen[a]

[a] Faculty of Chemical and Food Technology, Slovak University of Technology in Bratislava, Bratislava, Slovakia
[b] Slovnaft, a.s., Bratislava, Slovakia
* Corresponding Authors: xgaborcik@stuba.sk

ABSTRACT

We study a data reconciliation problem in a petrol refinery. The problem is to reconcile inventory and flow measurements to estimate true values of measured and unmeasured flows respecting the mass conservation. The problem is formulated as a mixed-integer quadratic program (MIQP). Upon successful problem resolution, a neural network (NN) is trained to mimic the MIQP solver to study potential improvements in CPU time without compromising the solution quality. The results show a significant improvement in refinery monitoring and feasibility of NN-based reconciliation.

Keywords: data reconciliation, optimization, neural networks, oil refinery

INTRODUCTION

A typical inventory infrastructure in a refinery consists of multitude interconnected technological blocks with an extensive network of liquid storage tanks. Although measurements of input and output flows to some of these blocks and tank levels might be available, the precise flows to, from, and between individual storage tanks are not necessarily directly metered. The goal of this work is to design a data reconciliation system coupled to mass balance to estimate the incoming, interconnecting and outflowing streams for product storage tanks in the Slovnaft refinery, an industrial partner.

Effective data reconciliation [1] is a vital tool in process industry to support operational and economic decisions [2]. Over the past years, data reconciliation has been successfully applied in petrol industry [3, 4, 5]. The problem is usually resolved by formulating an optimization problem in a static or dynamic fashion [6, 7]. The optimization problem rectifies the values of measured variables by enforcing the mass and energy balance conservation laws as problem constraints.

In this work, we leverage historical operational data, technical data sheets, and fundamental physical laws to derive the necessary relationships and introduce an optimization problem to design a data reconciliation framework. The data reconciliation needs to assign network topology (pipeline connections) and thus the problem boils down to a mixed-integer quadratic program (MIQP). To study whether a simplified workaround can be used for data reconciliation, we develop a neural network model that mimics the rigorous MIQP solution.

PROBLEM DEFINITION

Figure 1 shows a simplified scheme of the product storage block that we study in this work. The measured quantities (flows, inventory levels) are depicted in green and unknowns are marked with red question marks. The technological block consists of $m = 7$ tanks with measured hold-ups and there are $n = 17$ inflows into the block. The inflows into tanks are assigned based on operational requirements. The assignment is done manually and is unknown to the plant database. It follows the rule that each input stream leads into a single tank while multiple input streams can be assigned into the same tank. There are $o = 3$ outflowing streams from the block and, analogically, the outflow from a tank leads to a single output flow. Moreover, there are connections between the tanks, whose assignment is also unknown but a connection can be established only between two tanks at a time.

Data Reconciliation

Data reconciliation is the process of comparing, matching, and balancing data from two or more independent sources to ensure accuracy and consistency across systems [7]. The use of reconciliation in this work can be represented by a flowchart in Figure 2, showing

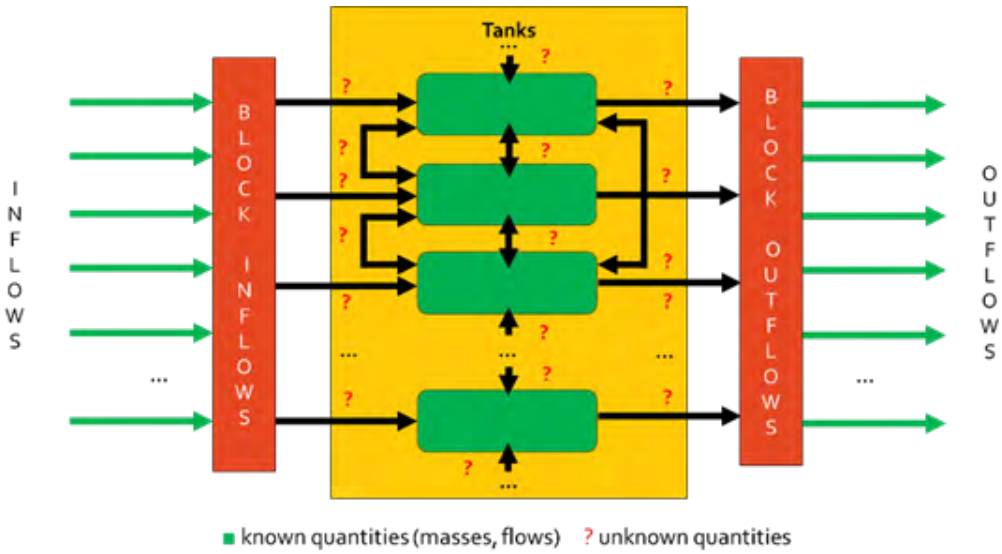

Figure 1. Schematic representation of a studied inventory problem, showing measured and unknown quantities.

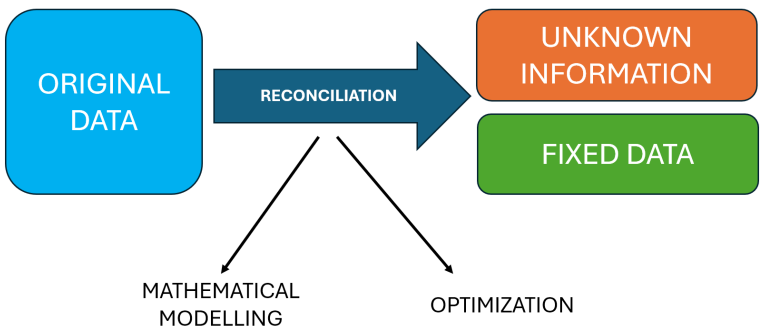

Figure 2. Schematic representation of data reconciliation methodology.

that we attempt to estimate unmeasured quantities from the measured quantities. At the same time, we also aim to rectify the value of measured variables so they would follow physical laws such as mass balance. All this is achieved using mathematical modelling and optimization.

The data reconciliation problem is formulated upon the following conditions and assumptions:

- The measured variables are denoted as y^{meas}
- The reconciled variables are denoted as y
- The dynamic mass balance model assumes the knowledge of actual state (hold-up) of the tanks and the state in the next sampling point
- According to the problem setup, we complement inflow and outflow of each tank with a binary variable in relation to block streams – a stream can either enter the tank or not and the outflow from the tank can enter only one outflow block stream
- The representation of interactions between tanks is realized with and m x m anti-symmetric matrix, representing all the possible inflows/outflows between each tank. The matrix diagonals are zero since they represent flow from/to the tank itself.

The described rules are transformed into the following mathematical relations. The vector

$$y^{meas} = \begin{Bmatrix} \dot{m}_{in}^{meas} \in \mathbb{R}^n \\ m_{tank,k}^{meas} \in \mathbb{R}^m \\ m_{tank,k+1}^{meas} \in \mathbb{R}^m \\ \dot{m}_{out}^{meas} \in \mathbb{R}^o \end{Bmatrix} \tag{1}$$

involves measurements, i.e., masses m and mass flows $\dot{m}$. Subscripts $in, out, tank, k$, and $tank, k+1$ represent inflows, outflows, and current and future tank hold-ups.

Similarly, the vector

$$y = \begin{Bmatrix} \dot{m}_{in} \in \mathbb{R}^n \\ m_{tank,k} \in \mathbb{R}^m \\ m_{tank,k+1} \in \mathbb{R}^m \\ \dot{m}_{out} \in \mathbb{R}^o \\ \dot{m}_{tank} \in \mathbb{R}^{m \times m} \end{Bmatrix} \tag{2}$$

describes the reconciled quantities together with $\dot{m}_{tank}$ signifying interaction flows between the tanks. Further,

$$z = \begin{Bmatrix} Z_{in} \in \{0,1\}^{n\times m} \\ Z_{tank} \in \{0,1\}^{m\times m} \\ Z_{out} \in \{0,1\}^{o\times m} \end{Bmatrix} \tag{3}$$

involves binary variables that represent the network connections. This specifically means the optimization problem for each time step involves a total of 272 variables: 189 binary (119 for inflows, 21 for outflows, and 49 for inter-tank connections) and 83 continuous (representing all mass flows and inventory levels at time k and $k+1$).

The objective of data reconciliation can be formulated as:

$$\min_{y,z\in\{0,1\}} (y^{meas} - y)^{\top} Q(y^{meas} - y) \tag{4}$$

where

$$Q_{ii} := \sigma_y^{-2},\ Q_{ij} := 0 \tag{5}$$

with σ_y representing the vector of standard deviations of measurement errors.

The objective function is subject to the constraints:

$$m_{tank,k+1} = m_{tank,k} + \Delta m \tag{6}$$

$$\Delta m = Z_{in}^{\top}\dot{m}_{in} + Z_{tank}\dot{m}_{tank}\mathbf{1}_{\boldsymbol{m}\times\mathbf{1}} - Z_{out}^{\top}\dot{m}_{out} \tag{7}$$

$$m_{tank,k+1},\ m_{tank,k},\ \dot{m}_{in},\ \dot{m}_{out} \geq 0 \tag{8}$$

$$\forall i \in \{1,\dots,n\}: \quad \sum_{j=1}^{m} Z_{in}(i,j) = 1 \tag{9}$$

$$\forall j \in \{1,\dots,m\}: \quad \sum_{l=1}^{o} Z_{out}(l,j) = 1 \tag{10}$$

$$\forall j \in \{1,\dots,m\}: \quad Z_{tank}(j,j) = 0 \tag{11}$$

$$\forall j \in \{1,\dots,m\}: \quad \sum_{k=1}^{m} Z_{tank}(j,k) \leq 1 \tag{12}$$

$$\forall j \in \{1,\dots,m\}: \quad \sum_{i=1}^{m} Z_{tank}(i,j) \leq 1 \tag{13}$$

$$Z_{\text{tank}}(i,j) + Z_{\text{tank}}(j,i) \leq 1 \tag{14}$$

that represent mass balance (Eqs. (6) and (7)), non-negativity constraints (Eq. (8)), assignment constraints in Eqs. (9) and (10), tank to tank constraints in Eqs. (11) – (13) and binary anti-symmetry in Eq. (14). Note that in Eqs. (6) and (7), we assume aligning the temporal difference Δm unit with the flow units.

Reconciled data can be potentially used for neural network training with original data as inputs and reconciliated data as outputs, giving us another tool beside optimization solver. Neural networks could improve real-time reconciliation should the MIQP solver turn out slow.

IMPLEMENTATION

We use MATLAB for framework implementation. The workflow contains the following steps:

1. Processing time series data with half-hour sampling time, covering several months of refinery operation (13, 726 half-hourly data points in total). Outliers are removed. Missing or removed values are linearly interpolated.
2. Unit and quantity conversion, where volumes and volumetric flows are transformed into masses and mass flows using measured densities.
3. Setting up the problem (4) – (14) using YALMIP [9] and solving it using GUROBI [10].
4. Training a neural network using *feedforwardnet* function in MATLAB (data split training / validation / testing = 70% / 15% / 15%). A feedforward network is created with three hidden layers sized 250, 150, and 80. The training algorithm is set to resilient backpropagation (*trainrp*). Each hidden layer's transfer function is set to *poslin*, which is ReLU. We exclude binary variables and use only continuous ones, because while binary variables are essential in mixed-integer problem definition, it is the set of continuous variables that carries the physical importance of the system. This step makes the problem simpler for neural network training, too. However, this approach makes the NN model topology-specific. If the physical network connections change (e.g., a new pipeline is installed or a tank is rerouted), the NN cannot dynamically adapt its binary logic and the surrogate model must be retrained using a new dataset generated by the updated MIQP formulation. Another specific settings, mostly for early stopping, are specified in Table 1.

Table 1: Feedforward neural network settings.

Parameter	Value
Maximum no. of epochs	2, 000
Sufficient performance	1E-05
Sufficient gradient	1E-06
Maximum validations check	15

RESULTS

Data Reconciliation using Optimization Solver

We demonstrate results of reconciliation via optimization in Figure 3. Figure 3.a) shows that the MIQP solver was able to find optimal solution in around 98% of the instances, with remaining 2% representing failed attempts, mostly due to numerical problems (98%) and few because of reached time/iteration limit of 60 s (2%).

Figure 3.b) represents the empirical cumulative distribution function of mass balance residuals between reconciled and measured variables. It should be noted that in the MIQP formulation, the reconciled flow is defined by the product of a binary existence variable and a continuous flow rate. While the optimization ensures that

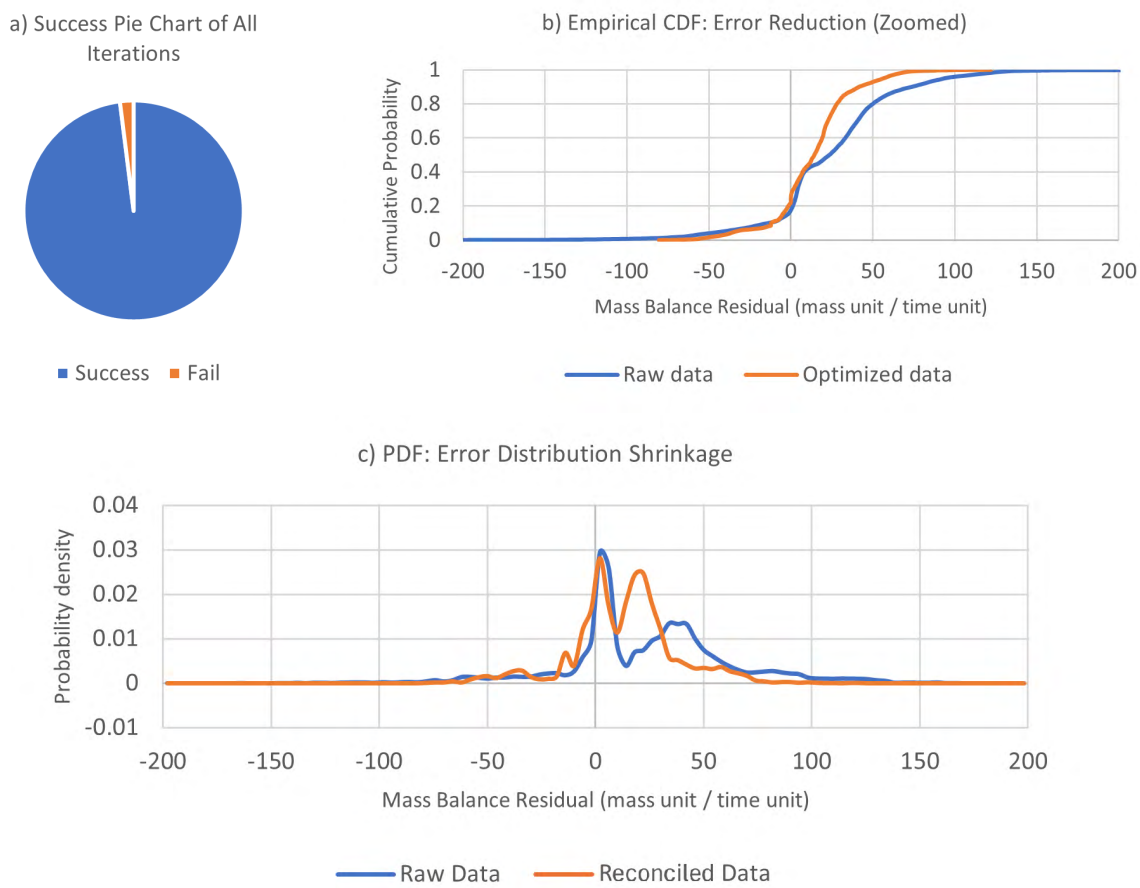

Figure 3. Performance plots of data reconciliation.

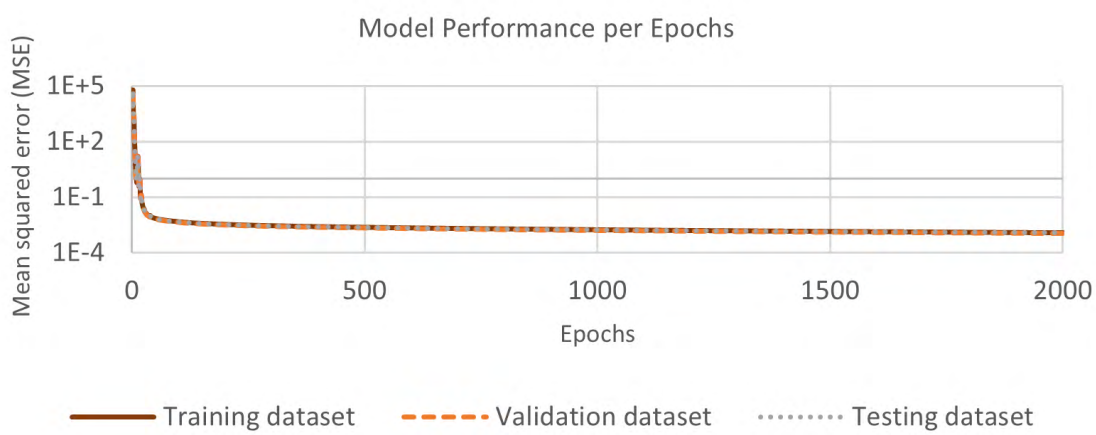

Figure 4. NN model performance over the training epochs.

mass conservation is strictly satisfied for these reconciled values, the residuals plotted here represent the deviation of raw plant measurements from this physically consistent logic. The steep slope around zero confirms high reconciliation precision for most observations. Approximately 70% of residuals are positive, indicating that the reconciliation process systematically compensates for under-measured inflows or unmeasured gains to satisfy mass conservation constraints. This indicates systematic (gross) errors and provides plant personnel with valuable monitoring information.

Complementarily, Figure 3.c) illustrates the probability density function (PDF), which showcases a significant reduction of long tails within the reconciliation

process. While the raw measurements exhibit a broad and dispersed noise profile, the optimized residuals are concentrated into two dominant peaks near zero, confirming the effective suppression of random measurement noise. The presence of the secondary peak in the positive region of the horizontal axis further validates the model's handling of the binary-continuous products; it demonstrates the aforementioned systematic inflow errors, indicating that the model successfully identifies and compensates for mass balance inconsistencies within the monitored system by reconciling the raw data toward the enforced physical structure.

Data Reconciliation using Neural Network

We demonstrate the neural network training and validation performance in Figure 4. The plot illustrates the evolution of the Mean Squared Error (MSE) over the maximum 2, 000 epochs, since requirements for early stopping were not satisfied. The logarithmic scale of the graph highlights the rapid initial convergence and subsequent fine-tuning of the model. The training process successfully reduced the performance metric from an initial value of 8.24×10^3 to a final termination value of 1.15×10^{-3} as indicated by the best validation performance achieved at the final epoch.

The high fidelity of the resulting model is further quantified by the regression analysis (not reproduced here), where the network reached a correlation coefficient of $R = 0.9982$ on the independent test dataset. This corresponds to a coefficient of determination $R^2 \approx 0.9964$, signifying that the neural network accounts for approximately 99.64% of the variance in the optimization data. The close tracking of the training, validation, and test curves throughout the 2, 000 epochs confirms a robust learning process with no signs of overfitting, demonstrating that the network has effectively internalized the underlying physical constraints and mass balance logic of the MIQP solver.

Table 2 shows comparison of CPU time as a comparison between optimization solver execution time and neural network inference time, representing potential of neural network for real-time optimization since it is, based on mean values, around 300-times faster. Both minimum and maximum CPU times are exponentially smaller for neural network than for optimization solver.

Table 2: CPU time comparison for optimization solver vs neural network.

CPU time (s)	Optimization solver	Neural network
Minimum	0.059903	0.007845
Mean	2.787424	0.009257
Maximum	60.480676	0.025399

CONCLUSION

We studied inventory data reconciliation in an oil refinery. The presented approach led to a significant increase in the accuracy of tank inventory monitoring and allowed for the quantification of previously unknown flows within the entire refinery system. Consequently, the resulting model provides a valuable tool for the early detection of anomalies, improvement of mass balance calculations, and real-time optimization of liquid inventory management. The future goal would encompass a use of an SCIP solver [11] that can be licensed freely not only for academic but also for commercial purposes.

ACKNOWLEDGEMENT

This work is funded by the Slovak Research and Development Agency under the project APVV-24-0007 and by the Scientific Grant Agency of the Slovak Republic under the grant 1/0263/25, and by the European Union under the grant scheme NextGenerationEU projects no. 09I01-03-V05-00002 and no. 09I01-03-V04-00024.

AUTHOR IDENTIFIERS

Author ORCIDs:
Paulen R: 0000-0002-1599-2634

REFERENCES

1. Leibman MJ, Edgar TF, Lasdon LS. Efficient data reconciliation and estimation for dynamic processes using nonlinear programming techniques. Computers & Chemical Engineering 16:963-986 (1992). https://doi.org/10.1016/0098-1354(92)80030-d
2. Galan A, De Prada C, Gutierrez G, Sarabia D, Gonzalez R. Real-time reconciled simulation as decision support tool for process operation. Journal of Process Control 100:41-64 (2021). https://doi.org/10.1016/j.jprocont.2021.02.003
3. Plácido J, Campos AA, Monteiro DF. Data reconciliation practice at a petroleum refinery company in Brazil. Comput Aided Chem Eng 27:777-782 (2009) https://doi.org/10.1016/S1570-7946(09)70350-5
4. de Oliveira EC, Lourenço FR. Data reconciliation applied to the conformity assessment of fuel products. Fuel 300:120936 (2021). https://doi.org/10.1016/j.fuel.2021.120936
5. Lid T, Skogestad S. Data reconciliation and optimal operation of a catalytic naphtha reformer. Journal of Process Control 18:320-331 (2008). https://doi.org/10.1016/j.jprocont.2007.09.002
6. Taylor JH, del Pilar Moreno R. Nonlinear dynamic

data reconciliation: in-depth case study. 2013 IEEE International Conference on Control Applications (CCA) :746-753 (2013). https://doi.org/10.1109/cca.2013.6662839
7. Bai S, Thibault J, McLean DD. Dynamic data reconciliation: alternative to kalman filter. Journal of Process Control 16:485-498 (2006). https://doi.org/10.1016/j.jprocont.2005.08.002
8. DAMA International. DAMA-DMBOK: Guide to the Data Management Body of Knowledge. Technics Publications (2017).
9. Löfberg J. YALMIP: A Toolbox for Modeling and Optimization in MATLAB. In: Proceedings of the CACSD Conference. Taipei, Taiwan (2004). https://yalmip.github.io/
10. Gurobi Optimization LLC. Gurobi Optimizer Reference Manual. Gurobi Optimization LLC (2025). https://www.gurobi.com
11. Bestuzheva Z. et al. The SCIP Optimization Suite 9.0. Technical Report, Optimization Online (2023). https://www.scipopt.org

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Advanced Process Control Structures for Energy-Efficient Downstream Processing in HMF Biorefineries

Norbert B. Mihály[a]*, Miruna Prodan[a], Vasile M. Cristea[a] and Anton A. Kiss[b]

[a] Babes-Bolyai University of Cluj-Napoca, 1 Mihail Kogalniceanu Street, 400028 Cluj-Napoca, Romania
[b] Delft University of Technology, van der Maasweg 9, 2629 HZ Delft, Netherlands
* Corresponding Author: norbert.mihaly@ubbcluj.ro

ABSTRACT

This research presents a novel framework for the surrogate-based dynamic optimization of control schemes within chemical separation and purification processes such as the biorefinery downstream processing. The current study investigated the downstream of an enzymatic bioreactor responsible for the synthesis of 5-hydroxymethylfurfural value-added derivatives, focusing on the critical balance between operational costs and productivity. Two high-fidelity long short-term memory neural network-based surrogate models were developed to predict energy consumption and economic gain, both achieving a coefficient of determination (R^2) exceeding 0.97. These models were subsequently integrated into a multi-objective optimization architecture to address an operating efficiency testing scenario characterized by stepwise inflow parameter changes. By exploring the resulting Pareto front, an optimal set of operational (control) settings was identified and validated. The results demonstrate that while energy consumption remained nearly constant, the total economic benefit increased by approximately 20% over the whole studied timeframe. Critically, during the transient period between steady-states, the economic benefit surged by nearly 60%, highlighting the potential of surrogate-based dynamic optimization. The high computational efficiency of the developed models (optimal solution obtained in less than three minutes) suggests significant real-time process control application possibilities. Unlike computationally expensive first-principles models, the developed neural architectures allow for virtually instantaneous setpoint adjustments. This paves the way for their integration into model predictive control or real-time process control layers, enabling industrial plants to respond adaptively to feed fluctuations, minimize off-spec product during transitions, and maximize profitability.

Keywords: Aspen Plus, Aspen Dynamics, Artificial Intelligence, Optimization, Process Control

INTRODUCTION

Classic control strategies often fail to maintain their performance under variable conditions as nonlinear processes become more interlinked and data-heavy [1]. This holds particularly true in case of novel complex processes, such as the separation and purification of 5-hydroxymethylfurfural (HMF) derivatives. The design of economical and process-efficient separation sequences with minimal environmental impacts is crucial for the future of such promising green biocatalytic pathways [2]. A shift toward dynamic operation optimization is required due to the intrinsic non-linearity of these separation units and the variability of bio-reactor effluents. Artificial Neural Networks (ANNs), that describe intricate nonlinear input-output interactions without the need for the deep understanding of the underlying process dynamics [3], represent the forefront of this shift.

Recent research on neural network-based surrogate models showed their high-fidelity at orders of magnitude smaller computational requirements, emphasizing their potential for real-time applications [4]. These networks can forecast ideal setpoints of the control loops by training on synthetic and historical information, allowing for proactive rather than reactive control. Furthermore, these surrogate models can be developed to directly forecast the evaluation criteria or objectives for the whole process [5]. The incorporation of temporal dependencies using Long Short-Term Memory (LSTM) networks has been a major development in this field. Due to

their design, LSTMs are most appropriate for handling sequences and time-series data [6]. LSTMs were shown to outperform conventional ANNs in handling the stochastic nature of changing inputs in industrial power research [7]. An excellent surrogate model was developed utilizing LSTM networks for multi-objective optimization in multigeneration systems, increasing operational efficiency by up to 12% in some aspects [8].

In the context of Computer-Aided Process Engineering (CAPE), the smooth integration of advanced surrogates into optimization frameworks remains a challenge, particularly for processes outside the energy sector. To address this, the current study proposes an improved LSTM surrogate model-based dynamic optimization architecture designed for the continuous refinement of operational parameters. The proposed framework aims to enhance the energy efficiency and economic viability of a previously proposed and investigated biorefinery downstream process for bio-based HMF derivatives [2], combining recent advances in machine learning and practical process systems engineering.

METHODOLOGY

Case study description

The development of a downstream process for the enzymatic conversion of HMF and 1-phenyle-thylamine (PEAM) was previously proposed and studied. The separation and purification of the resulting mixture of 2, 5-bis(hydroxymethyl)furan (BHMF), 1-phenylethylalcohol (PEA), 1-phenyle-thylamine (PEAM) and 5-(aminomethyl)-2-furanmethanol (HMFA) was discussed [2]. The mentioned process was since enhanced in terms of heat integration and product quality as shown in Figure 1. The enhanced version of the Aspen Plus simulation was implemented in Aspen Plus Dynamics in order to study the control of the process. The dynamic implementation was performed separately for the two non-interacting process lines responsible for the alcohol and amine products, respectively, as can be observed in Figure 1.

Dynamic model development

The proper functioning of each column was ensured by four control loops responsible for the material and energy balances. Columns C1 and C3 use three mass-balance loops plus a temperature control loop, while columns C2 and C4 employ cooling agent pressure manipulation specific to total-condenser configurations. Further control loops were employed to maintain the optimal conditions of the other unit operations, e.g., flow ratio controller for the liquid-liquid separation. The utilized controllers were conventional PID which were tuned based on various methods, e.g., closed loop (Ziegler-Nichols), open loop (internal model control). The current study focuses solely on the alcohol line of the proposed biorefinery downstream process. The regulatory control layer is driven by the upper control layer, where dynamic optimization of energy and economic indices is performed to provide setpoints for the control loops of the bottom layer.

Data generation

Two types of simulation were performed in order to obtain the database required for the development of data-driven models. In case of the first type, the parameters of the incoming stream were subjected to a step change, while the operational parameters were kept constant. The inflow parameters considered were the mass flowrate, temperature and mass fractions of the two alcohol products. The boundaries of change for the mass flowrate and mass fractions were defined as -10% and +10%, whereas the temperature was changed by -5 °C to +5 °C compared to the steady-state value. In order to encompass a variety of steady-state conditions, the simulation was not reset before each step changes, but only after three changes were implemented.

In case of the second type of simulation, the operational parameters were subjected to step changes in a similar fashion. The setpoints of the most important control loops were considered, i.e., temperature setpoints of columns C1 and C2 and the pressure setpoint of the COMP1 compressor. The limits for the step change were chosen as -10 °C to +10 °C for column C1, -12.5 °C to +12.5 °C for column C2 and -0.05 bar to +0.05 bar in case of the pressure setpoint. Similarly to the previously described data generation procedure, the simulation was reset to the original steady-state values only after a step change for each of the operational parameters was performed. The proposed tuned control structure resulted in a dynamic transition period of around one and a half hours after a step change (at the upper or lower boundary) was introduced in all cases before reaching steady-state.

In order to best cover the region defined by the step change boundaries in the least number of scenarios, a combined design of experiments (DoE) approach was utilized. The D-optimal design with three levels and the Halton quasi-random sequence were employed to fill out the input space. In this manner, 25 scenarios of the first type and another 20 scenarios of the second type were defined. The sequences of individual step changes were extracted at a five-minute sampling time from the mentioned scenarios. Each sequence contained half an hour of steady-state followed by the step change and a three-hour period which contained both the dynamic evolution and resulting steady-state.

LSTM model development

A total of 118 data sequences were obtained through the previously described methodology which

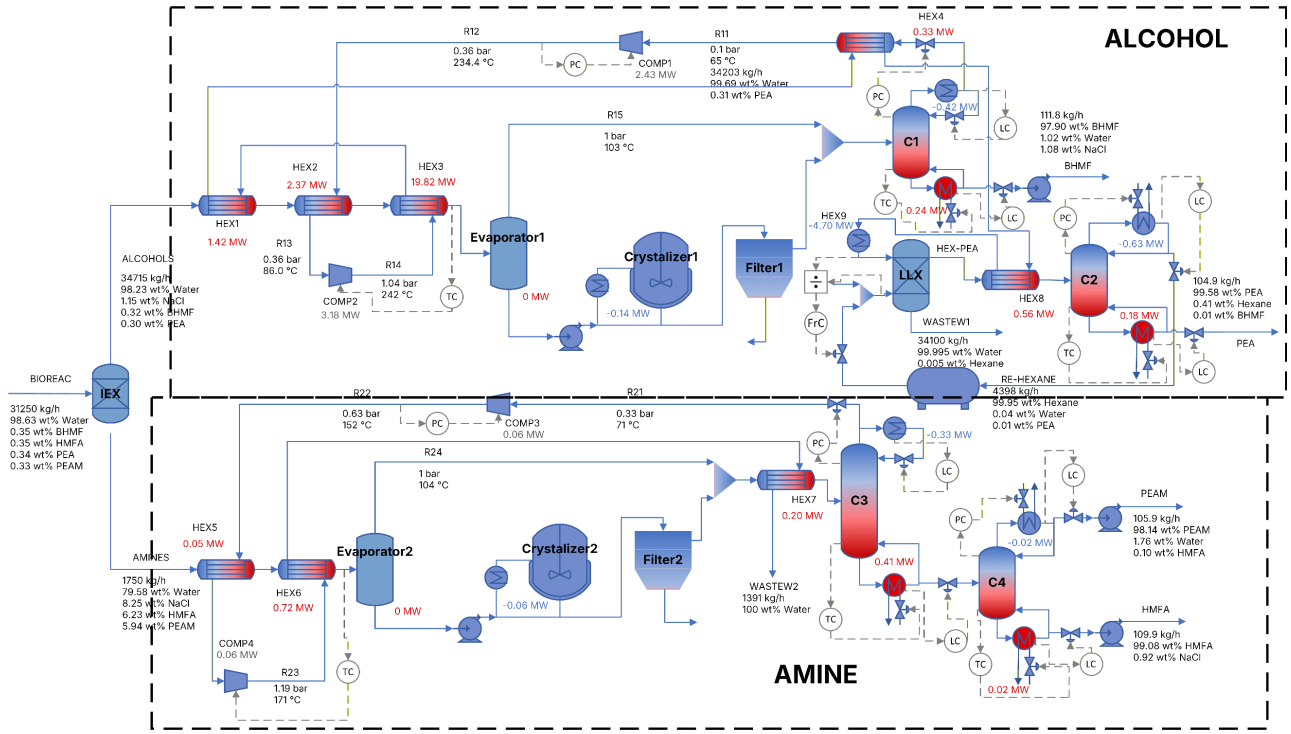

Figure 1. Schematic of product separation and purification processes and their control structures implemented in the Aspen Plus and Aspen Plus Dynamics simulation software.

were utilized in the ANN development.

The input features of these models are the seven parameters of interest i.e., the four inflow parameters and the three operational parameters which were varied in the data generation scenarios. The output features of these models were the objective function results for the studied sequences.

These objective functions are presented in the section below. Based on the collected data, a sequence-to-sequence regression model was intended to be developed. Consequently, long short-term memory (LSTM) type networks were chosen as these were developed to map an input sequence to a continuous-valued output sequence of possibly varying length [9]. The prediction capability of this kind of models was proven for similar modelling problems such as forecasting exact numerical values via signal processing [10] or for multi-step time series predictions [11]. The LSTM models were developed as single-output models in the sense that the sequences of only one of the two output features are predicted based on the input feature sequences. Thus, an LSTM model was developed for the prediction of each output feature as single output models were shown to outperform multiple output models [4]. The LSTM models were developed by utilizing 85% (i.e., 100 sequences) of the generated dataset for training, while the remaining 15% was equally divided between validation and testing.

Dynamic optimization

The optimization of the dynamic downstream process focuses on two aspects of great importance: the energy efficiency and the economic viability. In terms of energy efficiency, both the electric and thermal energy necessities were taken into consideration and expressed in equivalents of thermal energy. Equation 1 shows how this objective was calculated in $kWh_{thermal}$ unit.

$$obj_1 = reb_{C1} + reb_{C2} + 2.5 \cdot (bp_{COMP1} + bp_{COMP2}) \quad (1)$$

where, reb_{C1} and reb_{C2} represent the reboiler duties of columns C1 and C2, while bp_{COMP1} and bp_{COMP2} refer to the brake power of compressors COMP1 and COMP2, and 2.5 was the primary energy factor utilized to convert electric energy to thermal energy equivalents.

In case of the second objective, the economic benefit of the studied downstream process (alcohol line) was evaluated. This gain was expressed as the difference between the value of produced outputs and operational expenses strictly regarding the operation of the process line calculated in € and shown in Equation 2. However, the products have to meet the commercial standards, i.e., BHMF > 97 wt%, PEA > 99 wt% and $BHMF_{Temp}$ < 120 °C. In case these constrains are violated, the calculation of the objective was penalized by considering that the value of the product is zero in that timeframe.

$$obj_2 = 10 \cdot fm_{BHMF} + 3 \cdot fm_{PEA} - 0.1902 \cdot (bp_{COMP1} + bp_{COMP2}) - 25 \cdot (st_{C1} + st_{C2}) \quad (2)$$

Table 1. LSTM models evaluation criteria results calculated both on the training and testing datasets.

Model	Training			Testing		
	R^2	RMSE	MAPE (%)	R^2	RMSE	MAPE (%)
LSTM-obj1	0.9969	41.233	0.1285	0.9983	33.152	0.1261
LSTM-obj2	0.9732	20.248	1.3866	0.9798	19.515	3.0366

where, fm_{BHMF} is the mass flowrate of BHMF with a value of 10 €/kg, fm_{PEA} is the mass flowrate of PEA with a value of 3 €/kg, 0.1902 is the EU average price of electricity in the first half of 2025 [12], st_{C1} and st_{C2} are the steam necessities for the reboilers of columns C1 and C2 expressed in GJ/h with a true cost of 25 €/GJ.

RESULTS AND DISCUSSION

LSTM model results

Several deep neural network architectures were studied; however, the best results were obtained in case of both objectives with the LSTM model containing a lstm layer, followed by fully connected and dropout layers two times. The networks were trained with normalized input and output features and the adaptive moment estimation training algorithm was employed. The best performing models were selected based on their adjusted coefficient of determination (R^2), root mean square error (RMSE) and mean absolute percentage error (MAPE). The values obtained in case of the best models for each output feature are summarized in Table 1, demonstrating the satisfactory generalization capabilities of the surrogate-models.

It can be observed that both developed LSTM models maintain the evaluation criteria values obtained at training even in case of a previously unseen dataset, i.e., at testing. In terms of precision, the first objective can be forecast at a higher precision based on the presented results. This can be attributed to a smaller variation of the objective throughout the simulated scenarios. In order to better illustrate the capabilities of the developed models, the linear regression between the modeled and targeted variables are presented in Figure 2 for both output features, in case of both training and testing datasets.

Optimization results

The two LSTM models were utilized in optimization concerning a new scenario of step changes for the inflow parameters. Conversely to the previous scenarios, all four parameters were subjected to a step change in this case. The final values for these parameters were chosen to be different from those that appear in the generated dataset. The parameter changes applied were of -7.5% for mass flowrate, -4.5 °C for temperature, -5% and -7% for the PEA and BHMF mass fractions, respectively. The operational parameters were chosen as the decision variables, as such the optimal value of the three setpoints was searched for. As the two objectives are in trade-off, due to the fact that at lower energy input the product purity constraints are violated, a multi-objective optimization was performed.

The multi-objective optimization was performed in MATLAB with a genetic algorithm approach, a variant of the non-dominated sorting genetic algorithm II. In this manner, a set of noninferior solutions are found, called the Pareto front, considering the objective function described in Equation 3.

$$min(f_{obj}) = \begin{cases} obj_1 \\ -obj_2 \end{cases} \tag{3}$$

The Pareto front presented in Figure 3 shows the best set of solutions found. The first objective is but slightly affected by the decision variables. These results suggest that the energy consumption of the studied downstream process line is defined almost entirely by the inflow parameters. At the same time, further reduction of the energy consumption results in the decrease of the process' economic viability. This can be attributed to the reduction of the columns reboiler duties, which in turn impact the product purity. At a certain point their reduction results in the violation of purity constraints, i.e., penalties for the second objective function.

Assuming a decision-making process focused on the economic aspect of the downstream process of the alcohol products, the right-most solution (solution 11) was selected for validation on the dynamic model. Solution 11 was chosen as the economic gain is forecast to be best in this case, with a marginal increase in energy consumption compared to the other solutions of the Pareto front.

Firstly, the values of the investigated objectives were averaged over the studied period of time after the step change occurs. The utility of the developed LSTM models and optimization based on these models is evidenced, as the proposed solution 11 increases the economic gain from 209.8 €/h to 250.2 €/h meaning a percentual improvement of more than 19%. As previously mentioned, the two objectives are in a trade-off, consequently, the increase in economic viability comes at the cost of an increased energy consumption. However, the change in the first objective is barely above 1%, thus the chosen solution 11 stands firm considering the benefits

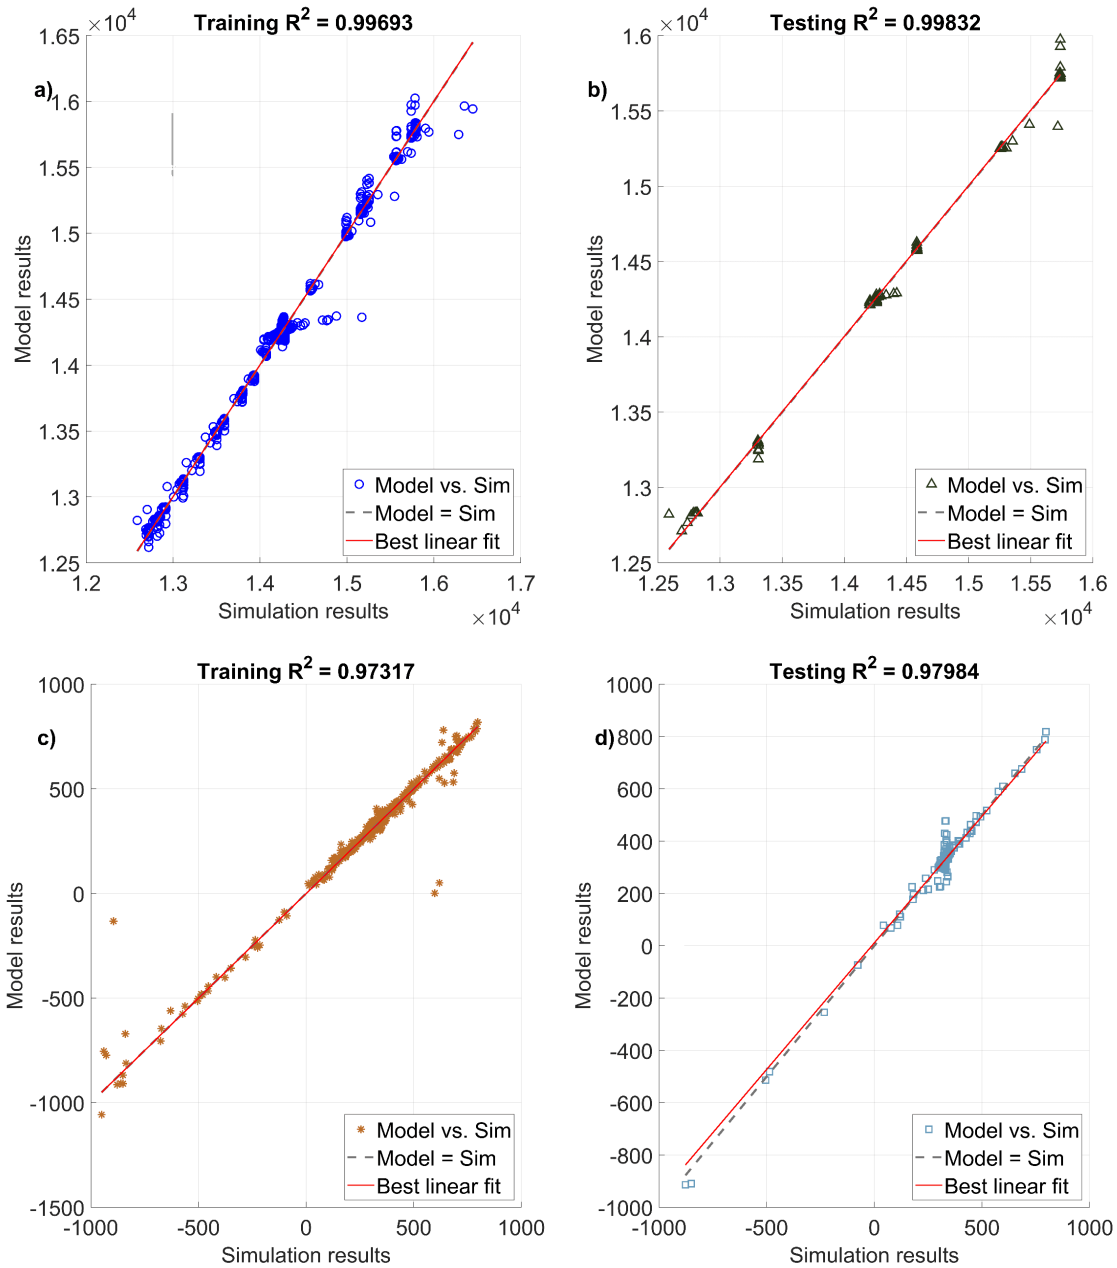

Figure 2. Regression of model and simulation results in case of: a) objective 1, training data; b) objective 1, testing data; c) objective 2, training data; d) objective 2, testing data

obtained in the economic aspect.

Secondly, the dynamic simulation was studied in two segments for a better understanding of when and how the proposed solution outperforms the reference scenario. The two segments considered were: i) the first half of the simulation time after the step change occurs (i.e., the first 1.5 hours) when the majority of the dynamic behavior occurs; ii) the second half of the dynamic simulation time when slight changes can occur, however, most of this timeframe can be considered steady-state. Based on the results presented in Table 2, the tested solution heavily outperforms the reference case in the first half of the studied scenario. The optimized control loop setpoints increase the economic gain by 58% at the cost of a 1% greater energy consumption. It is also important to note that the in case of solution 11, the overall and first half values of the objectives are very similar. As such, it is clear that the optimized setpoint values result in a smoother dynamic behavior. The large improvement in economic gain can be attributed to the objective penalties as the reference case settings result in sub-standard product quality, whereas the settings of solution 11

Table 2. Dynamic simulation results of the reference case and the tested solution.

Case	Overall = 3h		First half = 1.5h	
	obj_1 ($kWh_{thermal}$)	**obj_2 (€/h)**	**obj_1 ($kWh_{thermal}$)**	**obj_2 (€/h)**
Reference	15258	209.8	15279	157.7
Solution 11	15443	250.2	15432	249.3

avoid any constraint violation. These results also imply that the reference case outperforms solution 11 in the second half of the dynamic simulation. In which case, the initial operational parameter settings outperform those of solution 11 after steady-state is achieved.

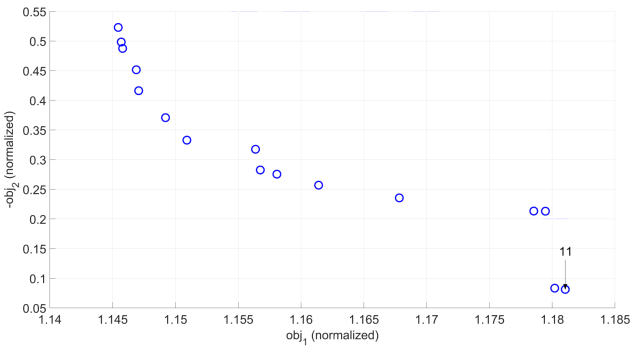

Figure 3. Multi-objective optimization results (Pareto front) as the second objective (negative economic gain) in function of the first objective (energy consumption)

The currently developed tools can address the above-mentioned limitation by performing the same optimization process for more sets of decision variables, e.g., one set for each half of the studied simulation timeframe. In this manner, the implementation would resemble that of a real-time optimization (RTO) layer with an optimization horizon of 1.5 hours. On the other hand, this integrated tool offers valuable information on the trade-off between process-wide performance criteria. Such an RTO layer would be able to take into account the dynamic behavior of the system, unlike classic RTO which rely on steady-state models. Another important aspect for the real-time implementations of the developed models is that the solutions presented in Figure 3 were obtained in less than 3 minutes. In these conditions, a selected solution can be adopted at the same time as the changes occur if information is available from upstream and the hydraulic retention time is of at least 3 minutes. However, even without such information an optimal solution can be adopted at the next point in time (as the frequency of the dataset is of 5 minutes), in which case the optimal solutions should be searched for starting from t+1 timestep if the step changes occur at step t. Due to the promptness of the underlying models, the integrated tool shows possible applications in dynamic real-time optimization of process-wide criteria, a combination of RTO and model predictive control (MPC).

CONCLUSIONS

This study demonstrates the possible advantages and applications of the integrated surrogate model-based multi-objective optimization of a biorefinery downstream process line. The studied tool can be utilized at a RTO hierarchy level with possible economic gain improvements of almost 20% at the cost of around 1% increase in energy consumption, due to its dynamic modelling capability. It shows even greater promise as a complementary tool to existing RTOs, being responsible only for the optimization of dynamic transitions in case of inflow variations. In this type of application in the case study, the economic gain was increased by 58%, with a marginal growth of 1% of the energy consumption. The developed integrated tool can be considered as a first step in combining RTO and MPC for the real-time dynamic optimization of complex industrial systems. Ultimately, the current study proposes a roadmap for actively taking advantage of inflow variations to optimize economic performance (with close to no compromise) in terms of energy efficiency, rather than simply surviving them.

ACKNOWLEDGEMENTS

This study was supported from the project “Advanced (multi)-enzymatic synthesis and purification processes for bio-based furan derivatives – ASPIRE” (CF 25/14.11.2022). The financial support from the International Advanced Fellowship CNFIS-FDI-2021-0061 is also gratefully acknowledged.

AUTHOR IDENTIFIERS

Author ORCIDs:
Mihaly NB: 0000-0002-0126-2657
Cristea VM: 0000-0001-9557-1970
Kiss AA: 0000-0001-5099-606X

REFERENCES

1. Javaid M, Haleem A, Singh RP, Suman R. Artificial intelligence applications for industry 4.0: a literature-based study. J. Ind. Intg. Mgmt. 07:83-

111 (2021).
https://doi.org/10.1142/s2424862221300040
2. Mih?ly NB, Prodan M, Cristea VM, Kiss AA. Sustainable downstream process design for HMF conversion to value-added chemicals. Systems and Control Transactions 4:491-496 (2025). https://doi.org/10.69997/sct.185398
3. Comas P, Morell A, Vilanova R, Vicario JL. Predicting transfer learning suitability in ann-based control of FOPDT industrial processes. Journal of Process Control 152:103476 (2025). https://doi.org/10.1016/j.jprocont.2025.103476
4. Mihály NB, Simon-Várhelyi M, Cristea VM. Data-driven modelling based on artificial neural networks for predicting energy and effluent quality indices and wastewater treatment plant optimization. Optim Eng 23:2235-2259 (2022). https://doi.org/10.1007/s11081-022-09724-5
5. Mihály NB, Luca AV, Simon-Várhelyi M, Cristea VM. Improvement of air flowrate distribution in the nitrification reactor of the waste water treatment plant by effluent quality, energy and greenhouse gas emissions optimization via artificial neural networks models. Journal of Water Process Engineering 54:103935 (2023). https://doi.org/10.1016/j.jwpe.2023.103935
6. Xu Y, Gao W, Qian F, Li Y. Potential analysis of the attention-based LSTM model in ultra-short-term forecasting of building HVAC energy consumption. Front. Energy Res. 9: (2021). https://doi.org/10.3389/fenrg.2021.730640
7. Hwang J, Suh D, Otto MO. Forecasting electricity consumption in commercial buildings using a machine learning approach. Energies 13:5885 (2020). https://doi.org/10.3390/en13225885
8. Ghafariasl P, Mahmoudan A, Mohammadi M, Nazarparvar A, Hoseinzadeh S, Fathali M, Chang S, Zeinalnezhad M, Garcia DA. Neural network-based surrogate modeling and optimization of a multigeneration system. Applied Energy 364:123130 (2024). https://doi.org/10.1016/j.apenergy.2024.123130
9. Graves A. Supervised sequence labelling. Studies in Computational Intelligence :5-13 (2012). https://doi.org/10.1007/978-3-642-24797-2_2
10. Zeng S, Ma C, Liu J, Li M, Gui H. Sequence-to-sequence based LSTM network modeling and its application in thermal error control framework. Applied Soft Computing 138:110221 (2023). https://doi.org/10.1016/j.asoc.2023.110221
11. Mu Y, Wang M, Zheng X, Gao H. An improved lstm-seq2seq-based forecasting method for electricity load. Front. Energy Res. 10: (2023). https://doi.org/10.3389/fenrg.2022.1093667
12. Eurostat. Electricity price statistics. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Electricity_price_statistics

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

CMLM: A Cascade of Machine Learning Models to detect and diagnose the performance of model predictive controllers

Elizabeth V. Melo[a]*, Argimiro R. Secchi[ab], and Maurício B. de Souza Jr.[ab]

[a] Chemical Engineering Program, PEQ/COPPE, Universidade Federal do Rio de Janeiro, Rio de Janeiro, Brazil
[b] EPQB, School of Chemistry, Universidade Federal do Rio de Janeiro, Rio de Janeiro, Brazil
* Corresponding Author: evmelo@peq.coppe.ufrj.br.

ABSTRACT

In this work, we propose a methodology for monitoring the performance of model predictive controllers (MPCs). A sequence of binary classification machine learning models, organized in cascade, called Cascade Machine Learning Models (CMLM), is evaluated to give a diagnosis of the control conditions. The proposed methodology was assessed using two case studies: a benchmark problem (the van de Vusse reactor under nonlinear MPC, NMPC) and a simulated industrial debutanizer column under commercial MPC. The ML models evaluated were the Random Forest and the Multilayer Perceptron. The results show that the proposed approach outperforms both a single multiclass model and traditional MPC performance monitoring methodologies, while remaining adaptable and scalable to larger applications.

Keywords: Machine Learning, Artificial Intelligence, Nonlinear Model Predictive Control, Process Monitoring, Fault Detection

INTRODUCTION

Model predictive controllers (MPCs) have become popular in industrial applications. However, just like other control strategies, this kind of controller depends on a healthy closed-loop configuration to perform properly. The satisfactory implementation of an MPC depends on its tuning parameters and mainly on the model representation, and a plant-model mismatch is still a challenge to be identified [1, 2].

Many metrics have been developed to detect and diagnose performance issues in controllers [3, 4]. Indeed, commercial software already has some tools to help monitor the closed-loop operation. However, obtaining a quick and accurate diagnosis remains challenging in industrial processes [5]. Improving data visualization is one of the approaches that has been explored for easing control loop diagnosis [6]. Machine learning (ML) is one of the classes of techniques employed to detect and diagnose control performance issues.

Control degradation is considered part of the fault detection and diagnosis (FDD) field [4], which is already widely exploring ML applications in abnormalities classification [6-8]. Some studies have been conducted to evaluate the applicability of ML tools in monitoring control loops, most of them focusing on SISO loops [10-13], and some on MPC [14-17]. However, there is still a lack of studies that evaluate the use of ML tools applied to nonlinear MPCs (NMPCs) and commercial MPCs.

In this study, we explore the use of ML algorithms applied to detect and diagnose the performance of MPCs. We propose a sequential procedure to evaluate the usage of independent ML models, named Cascade Machine Learning Models (CMLM), to detect and diagnose the performance of model predictive controllers using a data-based approach. To evaluate this approach, a comparison between Random Forest (RF) and Multilayer Perceptron (MLP) Neural Network as candidate ML models is carried out. As case studies, the van de Vusse reactor controlled by an NMPC, a nonlinear benchmark, and a simulated industrial debutanizer column controlled by a commercial MPC, with more variables involved, were used to generate data to train, validate, and test the learning capability of the ML models.

METHODOLOGY

System Architecture

Figure 1 shows a schematic representation of the process controlled by an (N)MPC and integrated with the AI-based monitoring system. This proposal aims to evaluate the capability of an IA agency to apply online to an industrial process, giving alerts and suggesting corrective actions to improve control performance.

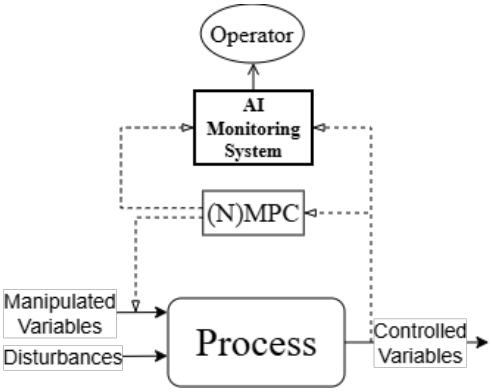

Figure 1. Controlled process integrated with the AI-based monitoring system.

Models in Cascade

Models in cascade have already been used for monitoring MPC performance [15]. However, this methodology has not been widely explored, especially for NMPC's performance monitoring. Figure 2 illustrates how these ML models are organized for the van de Vusse reactor system, while Figure 3 shows how it was designed for the debutanizer column, the case studies of this work. The specific adaptation of the CMLM is based on the physical knowledge of the process and on the main typical source of degradation. Process particularities must be taken into consideration to design the cascade arrangement.

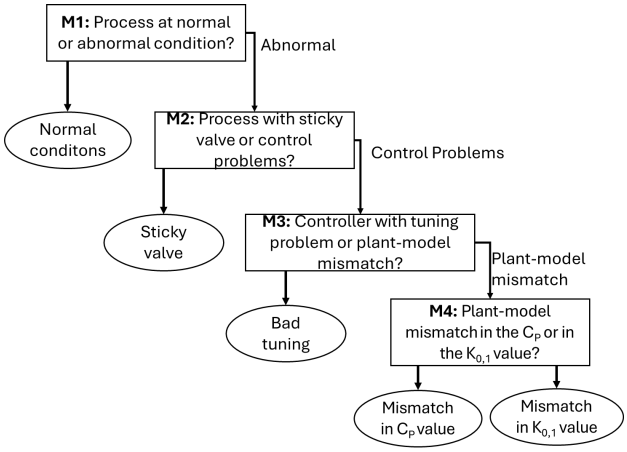

Figure 2. Models (M) in cascade for the classification problem for the van de Vusse reactor.

In Figure 2, the CMLM was designed to discriminate between the instrumentation and process problems and the control configuration problem. The first ML model is responsible for detecting whether the closed-loop performance is abnormal. Once abnormalities are detected, the process data goes to the second ML model, which is responsible for diagnosing whether the issue is due to the process actuator, such as valve stiction, or internal control configuration. Once an internal control

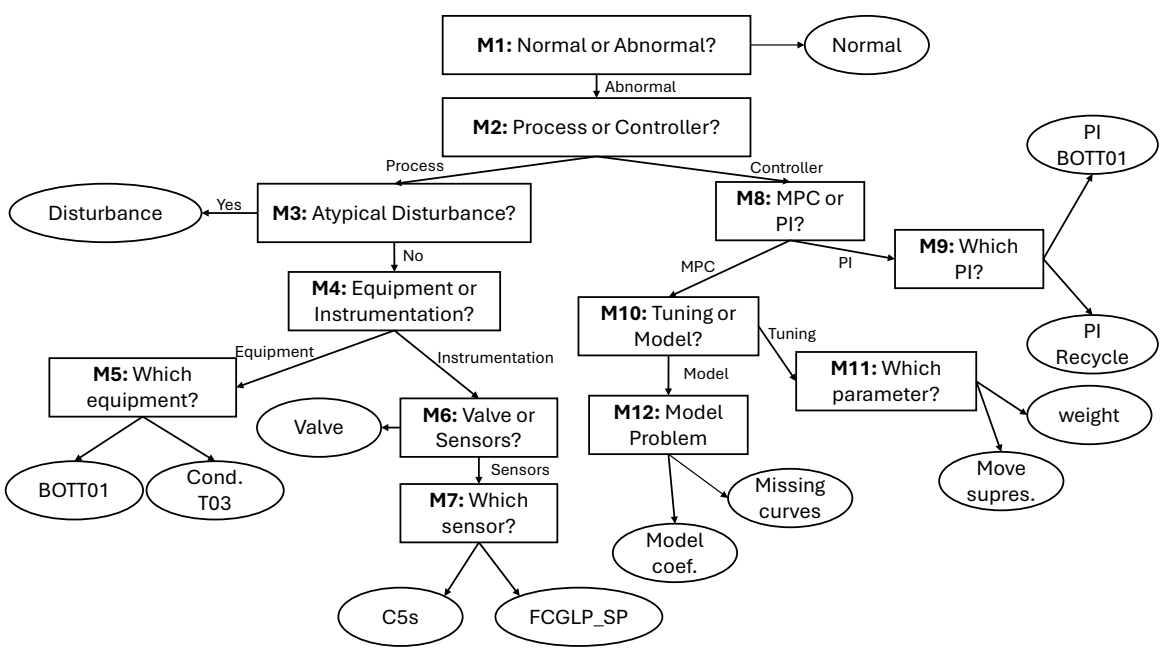

Figure 3. Models (M) in cascade for the classification problem for the debutanizer column.

configuration issue is detected, the third ML model is used for diagnosing poor control tuning or plant-model mismatch. Finally, if a plant-model mismatch is identified, the final model is responsible for diagnosing which of the two process model parameters should be evaluated.

The CMLM for the distillation column presented in Figure 3 follows, in general terms, its counterpart developed for the benchmark reactor. However, as the model is now more complex, additional models were included, such as M8 and M9, which deal with PI regulatory control. Additionally, to make the FDD problem more complete, models designed to detect atypical disturbances (M3) and sensor problems (M7) were also included.

This methodology is compared with a single machine learning model (SMLM) for a multiclass classification, in which detection and diagnosis occur at the same model, according to Figure 4. The multiclass approach is used in other studies for MPC status classification.

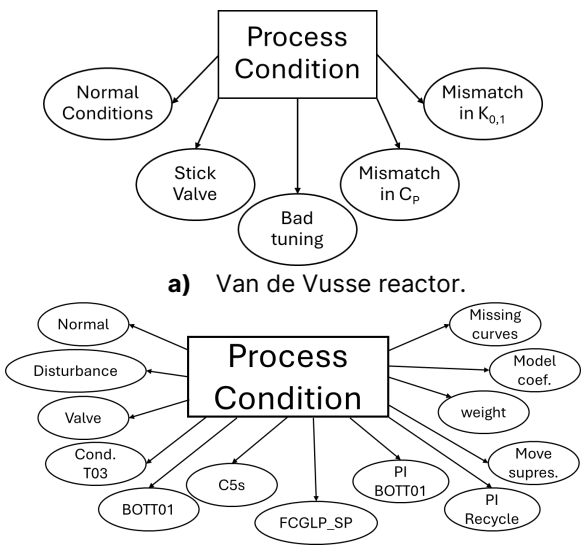

a) Van de Vusse reactor.

b) Debutanizer column.

Figure 4. ML representation for multiclass classification (diagnosis and detection) problem for each case study.

The main advantage of using CMLM is its facility for maintenance and adaptations. Chemical processes with many variables and controllers in cascade may suffer from a huge number of possible issues in the closed loop. In this case, there is a necessity to update the SMLM structure, and the whole model should be retrained, while the CMLM allows the possibility to update just one of the structures used or include a new model for a new diagnosis.

Besides, CMLM allows the use of different ML algorithms combined, as one technique can be better for one specific diagnosis than others. Furthermore, the technical team can also contribute to the precise diagnosis, removing models from the cascade. For example, if it is certain that the process is not under normal conditions, the operators and engineers can use only the models regarding the diagnosis, skipping the detection model.

Moreover, for each diagnosis, different features are expected to affect each abnormally. For the SMLM, all features that might contribute must be used, resulting in a complex ML model. For the CMLM, in case there is a necessity of reducing model complexity, different features could be used for each model in the cascade, allowing the construction of simpler models.

Case Study 1: van de Vusse Reactor

To evaluate the CMLM methodology, the van de Vusse reactor was used as a case study. The reaction system produces a substance B from a reagent A, having C and D as byproducts, according to Equation (1).

$$A \xrightarrow{1} B \xrightarrow{2} C \; 2A \xrightarrow{3} D \qquad (1)$$

The inlet flow enters a CSTR reactor with a temperature T_{in}, a concentration $C_{A,in}$ of reagent A, and a flowrate F, according to Figure 5. The outlet B concentration (C_B) and the reactor temperature (T) are the controlled variables (CVs). The heat exchange rate $\dot{Q}$ and the flowrate and reactor volume ratio F/V are used as manipulated variables (MVs). All system parameters were described by Engell and Klatt (1993) [18].

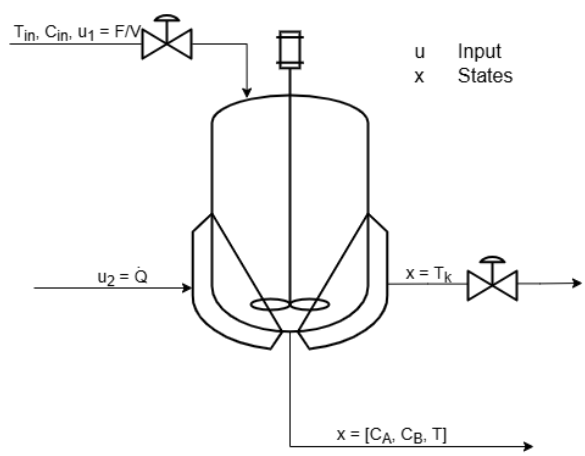

Figure 5. van de Vusse reactor.

The reactor is controlled by an NMPC. All simulations were run considering the same steady state as the initial process condition. The simulations were carried out using the same structure used by Lima *et al.* [19]. The concentration $C_{A,in}$ is considered an unmeasured disturbance. A sampling time of 2.5×10^{-3} h was adopted.

Data Generation

The simulation for generating the data was carried out considering four different cases:

1. **Normal conditions:** absence of faults.
2. **Valve stiction:** a dead band was considered.

3. **Poor tuning:** intentional poor tuning was made at the controller.
4. **Plant-model mismatch in C_P or $K_{0,1}$:** difference between the NMPC internal model and the simulated plant for the specific heat of the mixture C_P or for the pre-exponential constant of the B formation reaction rate ($K_{0,1}$) were considered.

Simulations were performed to produce 10, 000 samples for each class. The presence of Gaussian noise in the measurement of the control variables was also applied.

Case Study 2: Debutanizer Column

The system is based on a Natural Gas Processing Unit (NGPU) that produces Liquefied Petroleum Gas (LPG) from Natural Gas. The system is composed of three distillation columns: a demethanizer, a deethanizer, and a debutanizer column. The main column is the debutanizer column, which processes the liquid from the bottom of the deethanizer to produce LPG and light naphtha.

Figure 6 shows the simplified scheme of the process. The simulation was carried out using Aspen Dynamics, controlled by DMC3, the AspenTech's controller.

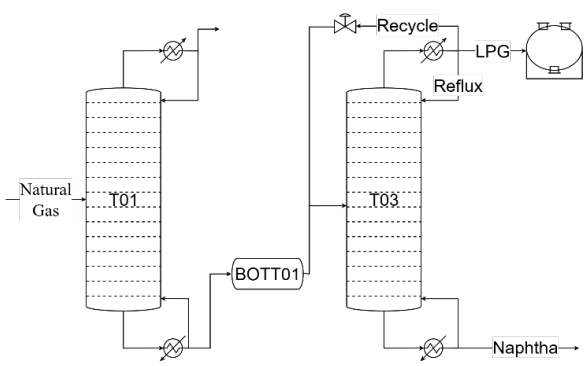

Figure 6. Simplified scheme of the NGPU used.

The control system was designed to have an optimization layer, given by the DMC3, and a regulatory layer, represented by different PI controllers present in the process. The main CVs of the process are:

- The ethane content of LPG
- The pentane content of LPG
- The Reid Vapor Pressure (RPV) of Naphtha
- The valve opening of the recycled flow rate of LPG back to the feed flow of the debutanizer column
- The condenser duty at the top of the debutanizer column
- The temperature at the top of the debutanizer column

The MVs for MPC are the setpoints of the variables:

- The reflux flow rate
- The temperature at the bottom of the demethanizer (T01)
- The temperature of the sensitive tray in the debutanizer (T03)
- The recycle flow rate
- The pressure at the top of the debutanizer
- The pressure at the bottom of the demethanizer

The simulation did not consider the presence of measurement noise. All variables were measured every minute without dead time. However, an artificial random noise was added to each sample of each variable to improve generalization.

Data Generation

The simulation for generating the data was carried out considering 13 different scenarios with 10, 000 samples each. The scenarios include:

1. **Normal conditions:** absence of faults.
2. **Unmeasured disturbance:** a different disturbance applied to the system.
3. **Condenser loss of efficiency:** loss of efficiency in the condenser of the debutanizer column.
4. **Reboiler loss of efficiency:** loss of efficiency in the reboiler at the bottom of the demethanizer column.
5. **Valve stiction:** a dead band was considered in the recycled flow of LPG.
6. **Pentane content measurement error:** an offset in the measurement of pentane content was implemented.
7. **Recycle flow rate measurement error:** an offset in the measurement of recycled flow was implemented.
8. **Change in the PI tuning at the demethanizer column:** changes in the control tuning of the pressure controller.
9. **Change in the PI tuning at the debutanizer column:** changes in the control tuning of the recycle controller.
10. **Change in the MPC weights:** the weights related to the pentane content were modified to decrease the concern regarding this variable.
11. **Change in the MPC move suppression:** more free and aggressive movement of both the recycled flow rate and the temperature of the

sensitive tray in the debutanizer was considered.b

12. **MPC internal curves removed:** the curves describing the relationship between the setpoint of the recycled flow rate and the pentane content, the RPV of Naphtha, and the temperature at the top of the debutanizer were removed to simulate a plant-model mismatch.

13. **Change in model coefficients:** the coefficients of the curve that describe the relationship between the pentane content and the temperature at the sensitive tray of the debutanizer were changed to simulate an outdated model.

Machine learning algorithms

The original dataset for both case studies was split randomly into data for training, validation, and testing, in the proportions of 64%, 16%, and 20%, respectively.

The input features for ML models, selected based on the physical knowledge of the system, were: the MVs included with up to two time lags; the integral of the time-weighted absolute errors (ITAE), considering the last 5 points of the controlled variables (CVs) with respect to their setpoints, and the prediction errors of these controlled variables. The ITAE of each CV was chosen to capture the error persistence through time, while the lagged MVs were chosen to capture dynamic behavior and the control actions. The prediction error, on the other hand, was chosen to capture model changes and the bias between the prediction and the measurement. This approach leads to a total of 10 features for the van de Vusse case study and 30 features for the debutanizer column.

RF and MLP were used to evaluate their performance in monitoring the NMPC and MPC closed loops. The implementations were carried out in Python. The architecture of each one was obtained by Optuna [20] after 100 trials, with the search conducted as follows:

- RF: The Python function RandomForestClassifier, from Scikit-learn, was used. Cross-validation was performed during optimizations, considering stratified K-folds with 5 folds. The objective function was the average accuracy of the test dataset. The optimized hyperparameters were:
    - Number of estimators between 10 and 100;
    - Maximum depth between 5 and 15;
    - Maximum features among "none", "sqrt", and "log2";
    - Minimum samples at each leaf between 5 and 10;
    - Criterion among "gini", "entropy", and "log_loss".
- MLP: The Python Tensor-Flow library was used. Early Stopping with a patience of 100 and a restore of the best weights was also used, considering the validation accuracy as a parameter. The reduction of the learning rate when accuracy achieved a plateau was also considered. The optimizer considered was Adam, and the batch size was 64. In each layer, a L_2 regularizer of 0.0001 was also used. The hyperparameters optimized were:
    - Number of layers between 2 and 4;
    - Number of neurons at each layer between 8 and 64;
    - Activation function among ReLU, LeakyReLU, and ELU;
    - Learning rate between 10^{-4} and 10^{-2};
    - Batch normalization between each layer: yes or no;
    - Dropout rate between 0.1 and 0.3.

F1-score and Accuracy, Equations (3) and (4), were used to compare RF and MLP results. TP denotes true positives, FP represents false positives, TN indicates true negatives, and FN refers to false negatives. An ideal scenario is one in which FP and FN are rare, while TP and TN are the majority. In such cases, these metrics tend to approach 1.

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \quad (3)$$

$$F1\text{-}score = \frac{2\times TP}{2\times TP+FP+FN} \quad (4)$$

The F1-score average used in this work, for both CMLM and SMLM, is the macro average, which calculates the F1-score for each label and then takes their unweighted mean. This average was used because, in this work, all classes are equally important, so they all have the same weight in the final value. Usually, on imbalanced datasets, the high metric score might be due to the correct classification of the majority class, while the minority class might have a poor result. The use of the unweighted average avoids this misinterpretation.

A total of 10 independent training sessions were performed for each ML model.

RESULTS

Table 1 presents the metrics obtained for the multiclass classification and for each ML model of the CMLM for the Vande Vusse case study. The best results in Table 1, when comparing RF and MLP, are in bold font. It is possible to observe that models specializing in each fault have generally better performance than just a multiclass ML model. Only the first model, responsible for detecting abnormalities, when using MLP, has a slightly lower F1-score. The low standard deviation observed in each model training indicates that the metric values are consistent over the 10 independently trained models.

Table 1: Metrics of each ML model for the van de Vusse case study.

ML	**RF**		**MLP**	
Metric	Accuracy	F1 score	Accuracy	F1 Score
Multiclass	0.8087 ± 0.0037	0.8128 ± 0.0038	**0.8491 ± 0.0367**	**0.8502 ± 0.0397**
M1: Detection	0.8379 ± 0.0023	**0.8351 ± 0.0024**	**0.8789 ± 0.0047**	0.8073 ± 0.0108
M2: Valve Stiction	0.9104 ± 0.0029	0.8885 ± 0.0038	**0.9442 ± 0.0059**	**0.9244 ± 0.0078**
M3: Tuning	0.9772 ± 0.0020	0.9746 ± 0.0022	**0.9918 ± 0.0026**	**0.9908 ± 0.0029**
M4: Mismatch	**0.9679 ± 0.0019**	**0.9679 ± 0.0019**	0.9580 ± 0.0025	0.9579 ± 0.0025

Before analyzing the results, it is important to understand how Accuracy and F1-score values should be compared. Accuracy is less sensitive to imbalanced datasets, and its high value might be due to the excellent performance of the majority class, independently of the minority class results. However, since the F1-score is taking the macro average, this metric is fairer to take into consideration the minority class, as all labels are treated as equally important. In M4 (diagnosing mismatch), where data is perfectly balanced, both the F1-score and Accuracy have similar values. M1 (detection), on the other hand, has a discrepancy between the values of those two metrics, which might be due to the artificial imbalance created in the cascade. Once the problem is turned into a binary classification, normal data becomes the minority class when compared to the other labels for this model. For this reason, F1-score is used as the priority metric to compare MLP and RF, while Accuracy is used to evaluate how well the majority class is being classified.

By comparing the performance of RF and MLP for each diagnostic model, it is possible to observe that RF models were better in diagnosing the parameter of mismatch (M4), while MLP models had slightly better performance in diagnosing valve stiction and poor tuning (M2 and M3). In the first model, responsible for detecting abnormalities in the control loop, the performance of MLP was better in terms of accuracy. However, the F1-score shows that this first model has better results with RF.

The reason for RF having better results in some situations might be due to the nature of this algorithm. RF based its predictions on hard boundaries of each feature, while MLPs are better at capturing nonlinear relationships between variables. For detection (M1) and diagnosis of mismatch (M4), the diagnosis is more based on whether the variables are higher than a given value. For the diagnosis of valve stiction (M2) and poor tuning (M3), which both cause oscillation in the closed-loop, MLP might be better suited to capture this nonlinear interaction.

The main advantage of using CMLM is that it gives a prompt diagnosis without exhaustive data investigation. One of the traditional methodologies applied in industry is using the standard deviation of variables to capture variability changes. This analysis is based on the minimum variance control theory and on the Harris Index [21] to monitor control performance. Aspen Watch and Schneider Electric's EcoStruxure Control Expert also use the variables' standard deviation as one of the metrics to evaluate the closed loop. Figure 7 shows the standard deviation time history of each MV (F/V and $\dot{Q}$) and CV (C_B and T) using a moving window of 100 samples (15 min of simulation). Visually, both the normal (blue area) and the valve stiction (green area) have the same range of standard deviation values, except for some samples equal to 0 for the F/V. Changes in the NMPC tuning and in its internal model parameters cause the standard deviation range for some variables to be different than those observed in normal conditions, especially regarding the tuning changes. However, the standard deviation gives information regarding a specific variable and does not give a clue to the root cause of this performance degradation.

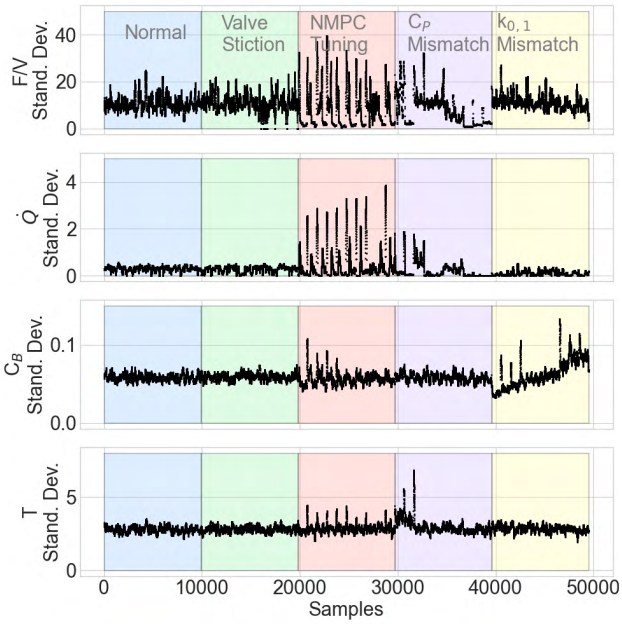

Figure 7. Standard deviation of the van de Vusse reactor.

To evaluate the method's scalability for industrial MPC implementations, a similar approach was tested in a debutanizer column. The results obtained are in Table 2, considering the cascade presented in Figure 3. The low standard deviation indicates the stability of the results. The metric values show that most of the models in the

Table 2: Metrics of each ML model for the debutanizer column case study.

ML	**RF**		**MLP**	
Metric	Accuracy	F1-score	Accuracy	F1-score
Multiclass	0.9350 ± 0.0013	0.8972 ± 0.0016	0.9479 ± 0.0006	0.9151 ± 0.0031
M1: Detection	0.9239 ± 0.0041	0.7554 ± 0.0079	**0.9977 ± 0.0003**	**0.9871 ± 0.0017**
M2: Process or controller	0.9407 ± 0.0011	0.9393 ± 0.0012	**0.9487 ± 0.0010**	**0.9466 ± 0.0015**
M3: Disturbance	0.9983 ± 0.0003	0.9945 ± 0.0009	**0.9998 ± 0.0001**	**0.9995 ± 0.0002**
M4: Equipment or Instrumentation	0.9996 ± 0.0001	0.9996 ± 0.0001	**0.9999 ± 0.0001**	**0.9999 ± 0.0001**
M5: Which equipment	**0.9999 ± 0.0001**	**0.9999 ± 0.0001**	0.9999 ± 0.0001	0.9999 ± 0.0001
M6: Valve or sensors	0.9987 ± 0.0003	0.9985 ± 0.0003	**0.9995 ± 0.0002**	**0.9995 ± 0.0002**
M7: Which sensor	**0.9999 ± 0.0001**	**0.9999 ± 0.0001**	0.9999 ± 0.0001	0.9999 ± 0.0001
M8: MPC or PI	0.9995 ± 0.0002	0.9993 ± 0.0003	**0.9998 ± 0.0002**	**0.9997 ± 0.0002**
M9: Which PI	**0.9950 ± 0.0010**	**0.9950 ± 0.0010**	0.9950 ± 0.0003	0.9721 ± 0.0018
M10: MPC tuning or model	0.9956 ± 0.0007	0.9956 ± 0.0007	**0.9985 ± 0.0005**	**0.9985 ± 0.0005**
M11: Which tuning parameter	0.9954 ± 0.0006	0.9954 ± 0.0006	**0.9989 ± 0.0004**	**0.9989 ± 0.0004**
M12: Which model issue	0.9997 ± 0.0002	0.9996 ± 0.0002	**0.9998 ± 0.0002**	**0.9998 ± 0.0003**

CMLM approach have a better performance than the SMLM. Only the RF M1, responsible for the detection, has significantly lower performance. This may have occurred because the RF algorithm cannot capture the more complex differences between normal and abnormal samples using thresholds.

By comparing the performance between RF and MLP models, the MLP models outperformed RF in most models of the CMLM, even though some of them have just a slight difference. Curiously, the RF models in the CMLM that have a similar or better performance than the MLP models are those responsible for diagnosing between two alternatives of the same type of equipment (M5: which process equipment, M7: which sensor, M9: which PI). This might have happened due to the characteristic of the RF algorithm, which is better suited to classify problems related to a given threshold. Since those alternatives affect different variables in the process, RF can handle the classification properly.

Regarding the computational cost, for the evaluation of 1, 000 random samples for the debutanizer column, the CMLM spends an average of 0.4 seconds per sample for the MLP, with a maximum time of 1.4 seconds. This was evaluated on a computer with Intel(R) Core (TM) i7-12700, CPU 2.10 GHz, and 16.0 GB of RAM. This evidence reinforces the capability of CMLM to be applied to online monitoring.

CONCLUSIONS

One of the goals of this study is to highlight the potential of using ML models for monitoring and diagnosing MPC performance. The method is more accurate than those currently used in industrial applications. The results show that ML models have the potential to diagnose control performance issues precisely, with CMLM having superior performance to a single multiclass model approach by isolating fault characteristics. By combining independent ML models, it is possible to use different model structures than a single type of ML model, taking advantage of the strengths of each ML tool. The two case studies show that the methodology can be implemented in both small and large industrial MPC applications. Even though the plant-model mismatch classification is too narrow and the source is usually not known, the case studies assume that the source of MPC degradation was limited to a few possibilities to explore how accurate the diagnosis of the ML models is. The diagnosis might be adapted to require the identification of only part of the process, instead of requiring the reidentification of the whole plant.

ACKNOWLEDGEMENTS

This study was financed in part by CAPES - Finance Code 001 and by Petrobras S.A. (Cooperation term no. 0050.0125244.23.9). Professors Maurício B. de Souza Jr. and Argimiro R. Secchi are grateful to financial support from CNPq (Grants No. 304190/2025-0 and 300744/2025-0), and Prof. Maurício B. de Souza Jr. is also thankful to FAPERJ (Grant No. E-26/200.532/2026).

AUTHOR IDENTIFIERS

Author ORCIDs:
Melo EV: 0009-0006-5566-8413
Secchi AR: 0000-0001-7297-3571
Souza MB Jr: 0000-0002-1090-8958

REFERENCES

1. Forbes MG, Patwardhan RS, Hamadah H, Gopaluni RB. Model predictive control in industry: challenges and opportunities. IFAC-PapersOnLine 48:531-538 (2015). https://doi.org/10.1016/j.ifacol.2015.09.022
2. Rawlings JB. Tutorial Overview of Model Predictive Control, *IEEE Control Systems Magazine*, v. 20, pp. 38–52, 2000. doi: 10.1109/37.845037.
3. Bauer M, Horch A, Xie L, Jelali M, Thornhill N. The current state of control loop performance monitoring – a survey of application in industry. Journal of Process Control 38:1-10 (2016). https://doi.org/10.1016/j.jprocont.2015.11.002
4. Qin SJ. Survey on data-driven industrial process monitoring and diagnosis. Annual Reviews in Control 36:220-234 (2012). https://doi.org/10.1016/j.arcontrol.2012.09.004
5. De Campos MCMM, Gomes MVC, Perez JMGT. Controle avançado e otimização na indústria do petróleo. Rio de Janeiro, Editora Interciência, 2013.
6. Elnawawi S, Siang LC, O'Connor DL, Gopaluni RB. Interactive visualization for diagnosis of industrial model predictive controllers with steady-state optimizers. Control Engineering Practice 121:105056 (2022). https://doi.org/10.1016/j.conengprac.2021.105056
7. Venkatasubramanian V, Rengaswamy R, Kavuri SN, Yin K. A review of process fault detection and diagnosis. Computers & Chemical Engineering 27:327-346 (2003). https://doi.org/10.1016/s0098-1354(02)00162-x
8. Özdemir P, Y?ld?r?m R. Ml@cheme: past, present, and future of machine learning in chemical engineering. ChemBioEng Reviews : (2025). https://doi.org/10.1002/cben.70012
9. Schweidtmann AM, Esche E, Fischer A, Kloft M, Repke J, Sager S, Mitsos A. Machine learning in chemical engineering: a perspective. Chemie Ingenieur Technik 93:2029-2039 (2021). https://doi.org/10.1002/cite.202100083
10. Ya?c? M, Forsman K, Böling JM. A machine learning classifier for detection of performance issues in industrial closed-loop PID controllers. 2024 10th International Conference on Control, Decision and Information Technologies (CoDIT) 19:1231-1236 (2024). https://doi.org/10.1109/codit62066.2024.10708598
11. Grelewicz P, Khuat TT, Czeczot J, Nowak P, Klopot T, Gabrys B. Application of machine learning to performance assessment for a class of pid-based control systems. IEEE Trans. Syst. Man Cybern, Syst. 53:4226-4238 (2023). https://doi.org/10.1109/tsmc.2023.3244714
12. Grelewicz P, Nowak P, Khuat TT, Czeczot J, Klopot T, Gabrys B. Practical implementation of computationally-efficient machine learning-based control performance assessment system for a class of closed loop systems. Applied Soft Computing 146:110690 (2023). https://doi.org/10.1016/j.asoc.2023.110690
13. Reynoso-Meza G, Carrillo-Ahumada J, Marques T. Multiobjective Machine Learning for control performance assessment in PID control loops, *IFAC-PapersOnLine*, v. 58, n. 7, pp. 168–173, 2024. doi: 10.1016/j.ifacol.2024. 08.029. 4th IFAC Conference on Advances in Proportional-Integral-Derivate Control PID 2024.
14. Lu Wang NL, Li S, Li K. Neural network based model predictive control performance monitoring-data-driven approach. 2013 9th Asian Control Conference (ASCC) :1-6 (2013). https://doi.org/10.1109/ascc.2013.6606358
15. Loquasto F, Seborg DE. Monitoring model predictive control systems using pattern classification and neural networks. Ind. Eng. Chem. Res. 42:4689-4701 (2003). https://doi.org/10.1021/ie020870k
16. Lu Q, Gopaluni RB, Forbes MG, Loewen PD, Backström JU, Dumont GA. Model-plant mismatch detection with support vector machines. IFAC-PapersOnLine 50:7993-7998 (2017). https://doi.org/10.1016/j.ifacol.2017.08.999
17. Dambros JWV, Trierweiler JO, Farenzena M, Kloft M. Oscillation detection in process industries by a machine learning-based approach. Ind. Eng. Chem. Res. 58:14180-14192 (2019). https://doi.org/10.1021/acs.iecr.9b01456
18. Engell S, Klatt KU. Nonlinear control of a nonminimum-phase CSTR. In: *Proceedings of the 1993 American Control Conference*, pp. 2041 – 2045, 1993.
19. Lima FARD, Faria RR, Curvelo R, Cadorini MCF, Echeverry CAG, de Souza MB Jr, Secchi AR. Influence of estimators and numerical approaches on the implementation of nmpcs. Processes 11:1102 (2023). https://doi.org/10.3390/pr11041102
20. Akiba T, Sano S, Yanase T, Ohta T, Koyama M. Optuna. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining :2623-2631 (2019). https://doi.org/10.1145/3292500.3330701
21. Harris TJ, Boudreau F, Macgregor JF. Performance

assessment of multivariable feedback controllers. Automatica 32:1505-1518 (1996). https://doi.org/10.1016/s0005-1098(96)00108-2

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Utilization of Additional Equipment Information for Drift Diagnosis in Chemical Plants

Linda Eydam[a]*, Julius Lorenz[b], and Leon Urbas[ab]

[a] TU Dresden, Process-to-Order Group, Process Systems Engineering Group, Dresden, 01069 Germany
[b] TU Dresden, Process-to-Order Group, Chair of Process Control Systems, Dresden, 01069 Germany
* Corresponding Author: linda.eydam@tu-dresden.de

ABSTRACT

Predictive maintenance is a promising approach to increase safety and productivity in chemical plants. One notoriously difficult problem in predictive maintenance are hard to predetermine, non-deterministic changes such as drifts. The term "drift" can be found with different definitions in this context. Therefore, it is defined here as changes in variables and parameters that occur orders of magnitude slower than the nominal process dynamics and are not directly measurable. Previous research resulted in a hybrid method that detects and diagnoses drifts from two sources: process and equipment. This method combines model-based and statistical approaches and additional information from the equipment, such as measurement gain or power consumption, is envisioned to reduce uncertainty about the drift cause [1]. First case studies revealed significant problems regarding economically viable integration of additional information. These problems arise due to the amount of information in scenarios with multiple devices, making analysis costly and time-consuming. For solving these problems, an automated evaluation of the information is introduced. It analyzes which additional equipment information is most relevant to distinguish between different drifts, based on sensitivity analyses. Additionally, a drift index - a unique mathematical label for drifts - is defined. The extended method is then successfully applied to a scenario closer to the reality in chemical plants with multiple adjacent devices. It is shown that the additional information is evaluated automatically and drifts are diagnosed more efficiently than in the previous method.

Keywords: Fault Detection, Process Monitoring, Drift Diagnosis, Additional Information, Predictive Maintenance

1 MOTIVATION

Reliability plays an important role in ensuring the safety and productivity of chemical processes. Increasing reliability by predictive maintenance is especially promising because it helps identifying equipment degradation and health status throughout its life cycle [2], [3]. A significant and challenging issue in predictive maintenance is the occurrence of drifts. The term “drift” has not yet become common in the literature on predictive maintenance. Therefore, it is defined here as not directly measurable changes in variables that occur orders of magnitude slower than the nominal process dynamics. This makes drift-related faults difficult to detect and diagnose. Examples of drifts include fouling of heat exchanger surfaces or the de-calibration of sensors.

To help solving the problem of drift detection and diagnosis, a hybrid fault detection and diagnosis (FDD)-method combining model-based and statistical approaches was developed [1]. This method selects and utilizes additional (freely) available data from the equipment, such as its power consumption, to reduce the uncertainty of assigning drifts as cause of observed faults. But how exactly does this information help distinguish between different drifts? Which additional equipment data is most useful for differentiating between different drifts? How can this information be evaluated?

These questions will be answered during this paper. First, the background of drifts is explained, including their types and definitions. Additionally, the fundamentals of the aforementioned method are presented. Next, the availability and utility of additional equipment information

are examined in more detail. Then the method is extended with an automated evaluation of additional data using sensitivity analysis. Our case study uses real data to investigate whether the automated evaluation of additional information and its implementation are reasonable and satisfactory, particularly in cases involving multiple devices and large amounts of additional information. Finally, the implications of extending the method regarding the overall method and potential future developments, is discussed.

2 BACKGROUND

The FDD-methods applied in predictive maintenance detect and diagnose faults for instance by data-driven approaches using principal component analysis and autoencoders [4]. Faults are defined as deviations of a set of observable process variables from an acceptable range [5] that are caused by failures or malfunctions of technical equipment. Faults caused by drifts, that is slow changes in not directly measurable variables are notoriously hard to detect. Firstly, drifts are, per our definition, typically much slower than the nominal dynamics of the process. Secondly, drifts may be attributed both to slow changes in the process (e.g. fouling) or the equipment (e.g. sensor de-calibration). Thirdly, different drifts may occur at the same time and influence or superposition each other.

A hybrid method combining statistical and model-based approaches was developed in [1] for the detection and diagnosis of drifts, even if they occur simultaneously. The model-based part of this approach utilizes a process simulation to generate empirical observations for known drift phenomena. The data provides the basis for a drift-related cluster analysis. The resulting clusters then-after are used to classify the data collected from the real plant. The procedure of the method is shown in Figure 1.

This hybrid method uses additional equipment information to identify characteristics that help distinguish between different drifts. Because drifts manifest themselves in the form of indirect consequences, such as increased energy consumption, less accurate positioning or changes in dynamic behavior. Therefore, with the help of additional information the space for classifying data and diagnosing drifts is defined. But this space is spanned by vectors that specify the different possible drifts. These drift vectors are defined in all available dimensions, that means with help of all measured variables m, consisting of primary and additional information. Thus, for n possible drifts, n drift vectors in m dimensions are spanned. So, an m-dimensional space is spanned with all available data. But for an unambiguously diagnosis of drifts, the n drift vectors must span an n-dimensional space. Therefore, they must be linearly independent of each other. Usually more measured variables are accessible than the required dimensions, so m > n. But when m > n, which information should be used?

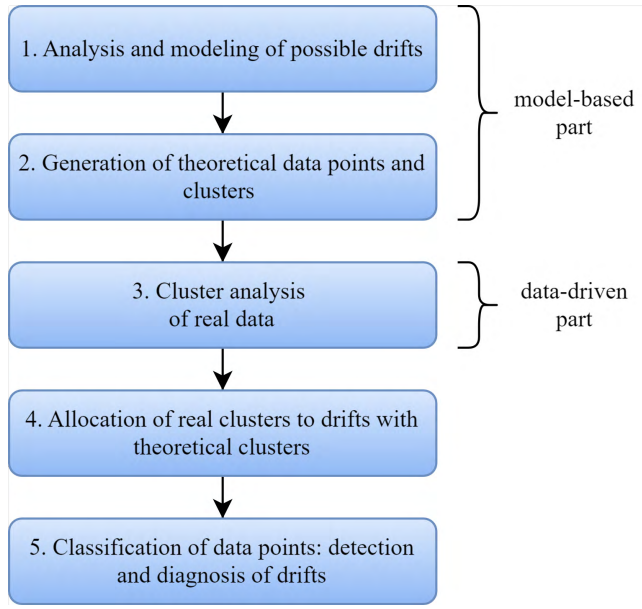

Figure 1: Workflow of the hybrid method from [1].

So, depending on how the space is defined in terms of its dimensions and axes, it becomes easier or more difficult to differentiate between drifts. Until now, this has been analyzed manually based on the independence of the drift vectors in different spanned n-dimensional spaces. But when n variables must be filtered out of m variables, the number of possible combinations increases significantly, especially as the number of variables increases. For example, when three drifts occur and five variables are measured, there are 10 possible combinations. If there are six variables, then there are 20 possible combinations. With seven variables, there are 35 combinations. The drift vectors must be defined for each combination of variables, and their independence must be checked. If, in addition, several combinations of variables define independent drift vectors, which combination makes the most sense? A process expert must decide this.

However, this cannot be used in practice because it requires too much computational effort and is too complex, especially in industrial plants with large amounts of data. That's why our goal is to expand the hybrid method with an automated evaluation of all available information to identify the most beneficial information for defining an n-dimensional space and efficiently diagnosing drifts.

3 ADDITIONAL EQUIPMENT INFORMATION AND THEIR EVALUATION

Modern devices typically record a lot of additional information to complement the primary measurement

values. Depending on the sensor or actuator being examined, various data is available, such as

- Power consumption,
- Measurement gain,
- Signal-to-Noise Ratio (SNR),
- Feedback of the valve position,
- Pressure to open a valve,
- Internal temperature of the device.

3.1 Availability of additional information

There are several ways to make this additional equipment data accessible to the FDD-methods. The traditional way would be to sample the data points via the process control system (PCS) and making it available for FDD on the PCS' process data historian. This is sometimes not possible in chemical plants because effort and demand are too high related to the amount of data. Therefore, there are two options: data reduction to acceptable amount or using a (standardized) side channel. But to reduce the amount of data, it must already be known at this point which data is needed, and which is not. A side channel via the NAMUR open architecture, for example, better addresses the additionality and expected changeability [6], [7].

3.2 Evaluation of additional information with the help of sensitivity analysis

In order to identify the most beneficial information for defining an n-dimensional space and efficiently diagnosing drifts, it is necessary to know which measured variables are most sensitive to drifts. A sensitivity analysis (SA) is suggested to do so. Such an analysis is used to investigate how the variation in the outputs of a model come from the variation in the input factors. Sensitivity analysis can be local or global, quantitative or qualitative and one-at-a-time (OAT) or all-at-a-time (AAT) [8]. Different purposes according to [8] are:

- Ranking / Factor Prioritization: the input factors are ranked based on their relative contribution to the variability of the output,
- Screening / Factor Fixing: input factors with negligible influence on the variability of the output are identified,
- Mapping: detection of the area of the input variability space where significant output values are produced.

To find the most beneficial additional equipment information, a SA is suggested that is global, AAT and prioritizes input factors. With a global SA, variations in the entire variability space of the input factors are considered. This is important, because otherwise a nominal value for the input factors would have to be specified. In an AAT method the input factors are varied simultaneously to get variations in the output. So, direct influences of each factor as well as joint influences through interactions are considered [8]. Factor prioritization ranks the contributions of input factors, or in this case, the available information. Thus, it is possible to identify which information is most sensitive to the output, or in this case, the different drifts.

Possible approaches are correlation & regression analysis or variance-based & density-based methods [8]. Here correlation analysis methods are used. In these methods the correlation coefficient between input factor x_i and output y is calculated with the help of statistical analysis of input/output datasets. These correlation coefficients are used as the sensitivity measures [8]:

$$S_i = correlation\,(x_i, y) \tag{1}$$

One type of correlation is the Pearson correlation coefficient (CC). It is for linear relations between the input factors $\boldsymbol{x}$ and the output y. Another type is the Spearman rank correlation coefficient (SRCC) for nonlinear relations that are monotonic [8]. Since we initially assume linear relationships, the correlations between the measured variables and the drifts are calculated with the Pearson correlation coefficient.

In this case, the input factors of the SA are the measured variables. This means all primary and additional information. The outputs are the various drifts, along with their corresponding theoretical data points, including the area without a drift. To determine the most sensitive information regarding the SA outputs, they require a unique mathematical labeling. As such a labeling the drift index $Drift_{ind}$ is defined. Each theoretical data point that belongs to a modeled clusters have such an index that is calculated with the following formula:

$$Drift_{ind} = Drift_{num} + l_{Vec,scal} \tag{2}$$

The number of the respective drift to which the data point belongs is denoted as $Drift_{num}$. It is an integer between 0 and n. The data points without drift have a value of 0, and the drift data points have values from 1 to n. Furthermore, the length of the respective data point vector l_{Vec} is calculated. This data point vector is spanned from the zero point to the respective data point. Its scaled length is denoted by $l_{Vec,scal}$. An example of the drift index of theoretical data points in two-dimensional space is shown in Figure 2.

Then the Pearson correlation coefficients are computed and compared. The n measured variables with the highest coefficients are chosen. These are the n out of m variables to which the drifts are most sensitive.

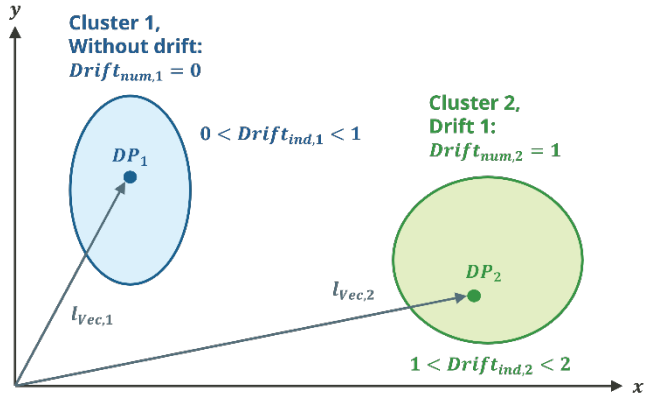

Figure 2. Example of the drift index for two data points DP in two dimensions, which are in two different clusters. The indexes are composed of the respective drift number $Drift_{num}$ and the scaled length of the respective data point vector $l_{Vec,scal}$. DP_1 has no drift and therefore a drift index between zero and one. DP_2 is part of the drift cluster and has a drift index between one and two.

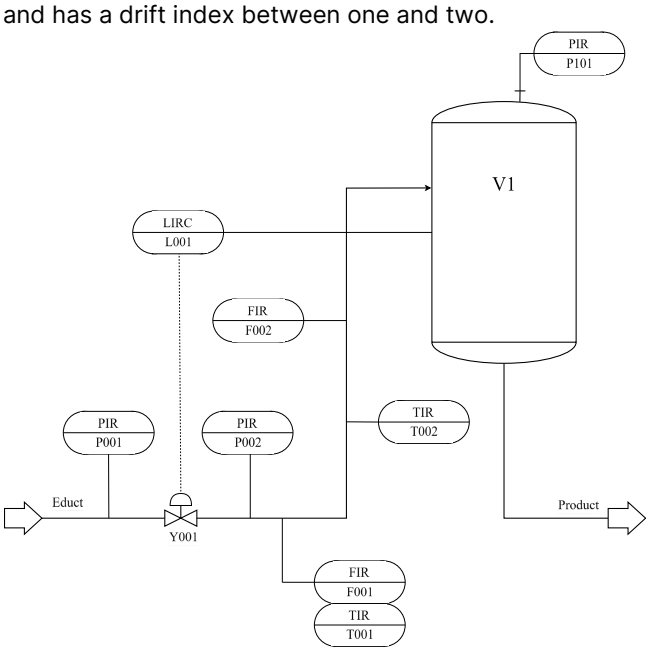

Figure 3: P&ID of the investigated case.

The global sensitivity analysis is implemented in an automated manner behind the cluster analysis of real data in Figure 1. In this step it is analyzed which measured variables are most sensitive regarding modeled drifts. Then the dimensions of the real and modeled clusters are adapted. So, all available information is evaluated automatically to identify the most beneficial information. That means only n of the m dimensions are used for assigning the clusters to drifts and the classification of new data points for detecting and diagnosing drifts. Since this selection is made automatically, the computational power and time to get detection and diagnosis results, is reduced.

4 CASE STUDY

A case study is conducted to demonstrate that the method with automatic evaluation can be applied to increase the efficiency compared to [3].

Table 1: Measured variables of the scenario considered.

Measured value	TAG	Symbol	Type of information
Flow rate ultrasound	F001	F_1	Primary
Flow rate inductive	F002	F_2	Primary
Temperature 1 of medium	T001	T_1	Primary
Temperature 2 of medium	T002	T_2	Primary
Position of position controller	Y001	D	Additional
Pressure of position controller	Y001	P_{Pos}	Additional
Pressure control air	Y001	P_{Air}	Additional
Calculated actuating value	Y001	Y_{Calc}	Additional
Temperature position controller	Y001	T_{Pos}	Additional
Measurement gain ultrasound	F001	G	Additional
Signal to noise ratio (SNR)	F001	SNR	Additional
Sound velocity	F001	V	Additional
Pressure before valve	P001	P_1	Primary
Pressure behind valve	P002	P_2	Primary
Actuating value of valve	Y001	Y	Additional
Level of tank	L001	L_{tank}	Primary
Pressure in tank	P101	P_{tank}	Primary

The case study refers to a real-world chemical plant using the data from several adjacent devices as shown by the P&ID of this configuration in Figure 3. A three-phase liquid with solids due to crystallization and gas bubbles flows through the pipes. The actuator is the control valve Y001 in front of a tank V1. The flow rate into V1 is measured using the ultrasound sensor F001 as well as with the inductive flow sensor F002. The temperature of the medium is measured redundantly with the sensors T001 and T002. The pressure before and after the control valve is measured with the sensors P001 and P002.

Furthermore, there are sensors that measure the level L001 and pressure P101 in the tank. The system runs continuously and automated. The level in the tank L001 is controlled with the valve Y001. Table 1 lists all the variables of the measuring devices considered. It also indicates whether the information is a primary or additional measurement value.

Various drifts can occur in this scenario under the assumption that the temperature and pressure sensors are free from deposits and abrasions. The possible drifts and their effects are described below:

- Equipmentdrift 1 (Eq1): Due to deposits that have slowly accumulated, the control valve is stuck and does not open properly.
- Equipmentdrift 2 (Eq2): Due to aging processes in the valve, it closes less effectively over time until it finally fails to close completely.
- Processdrift 1 (Pd1): The quality of the educts changes over time and deteriorates. Reasons could be increased proportions of gas bubbles or solids as well as other medium changes.
- Processdrift 2 (Pd2): Due to changes in the process, the flow rate has fallen below the desired specification.
- Processdrift 3 (Pd3): Due to changes in the process, the flow rate has increased above the desired specification.
- Processdrift 4 (Pd4): An unforeseen slow change in the process has reduced the temperature of the medium.

The area without any drifts as well as the drift areas are modelled using a mechanistic approach. But parameters, nominal process values and relationships between variables are calculated with the help of available real data. Then, theoretical data points and clusters are generated. Following the available data from the real plant is normalized and clustered. The data points are divided into three clusters and outliers. After that the evaluation of the measured variables with the method described in Section 3 resulted in the following ranking:

1. Signal to noise ratio SNR,
2. Sound velocity V,
3. Measurement gain ultrasound G,
4. Temperature 2 of medium T_2,
5. Temperature 1 of medium T_1,
6. Flow rate inductive F_2.

These six variables are used for all subsequent steps in the method. Half of the variables are additional equipment information, and the other half are primary measurement values. The measurement gain, which was identified as important through sensitivity analysis, can detect the deterioration of the reactant quality over time. This occurs because the presence of additional gas bubbles or solids interferes with the ultrasound measurement and slowly increases its gain over time. Figure 4 shows this phenomenon. The gain increases slowly over two and a half hours, rising from 33 to over 40. This example shows that useful information can indeed be identified through the automatic evaluation.

Then the real clusters are assigned to the drifts. In this case, a cluster is assigned to Processdrift 1. Another cluster is identified as drift-free. The final cluster cannot be assigned to a drift. It is close to Processdrift 1 and far from all the other modeled drifts. The data points of the unassigned cluster, as well as the outliers, do not belong to any drift. They are outside the assigned clusters and therefore possible combinations of drifts. These data points are categorized and their causes are diagnosed using the further procedure from [1]. For this purpose, a combination of real and modeled clusters is used. Only one real drift cluster exists. For the others', modeled clusters are used. The process of diagnosing combinations of drifts will not be explained in detail here. For more information, see [1]. As a result, the following combinations of drifts are diagnosed:

- Eq2 and Pd1,
- Pd1 and Pd3,
- Pd1 and Pd4.

Based on these diagnostic results, it can be concluded that the most common drift is Processdrift 1: The quality of the educts changes over time and deteriorates. One real cluster is assigned to this drift, and it occurs in every combination of drifts. In contrast, equipment drifts rarely occur alone.

But how efficient is the method with automatic evaluation compared to without? If the method were applied to this case without automated evaluation, the following steps would have had to be taken:

- Combination of all possibilities to span a six-dimensional space from 17 measured variables (12, 376 possible combinations)
- Definition of drift vectors for all combinations,
- Checking if the drift vectors are independent,
- Collection of all combinations with independent drift vectors and manual decision which combination is the best by a process expert.

These steps are time-consuming and expensive. In contrast, automated evaluation only requires assigning a drift index to all theoretical data points and performing a

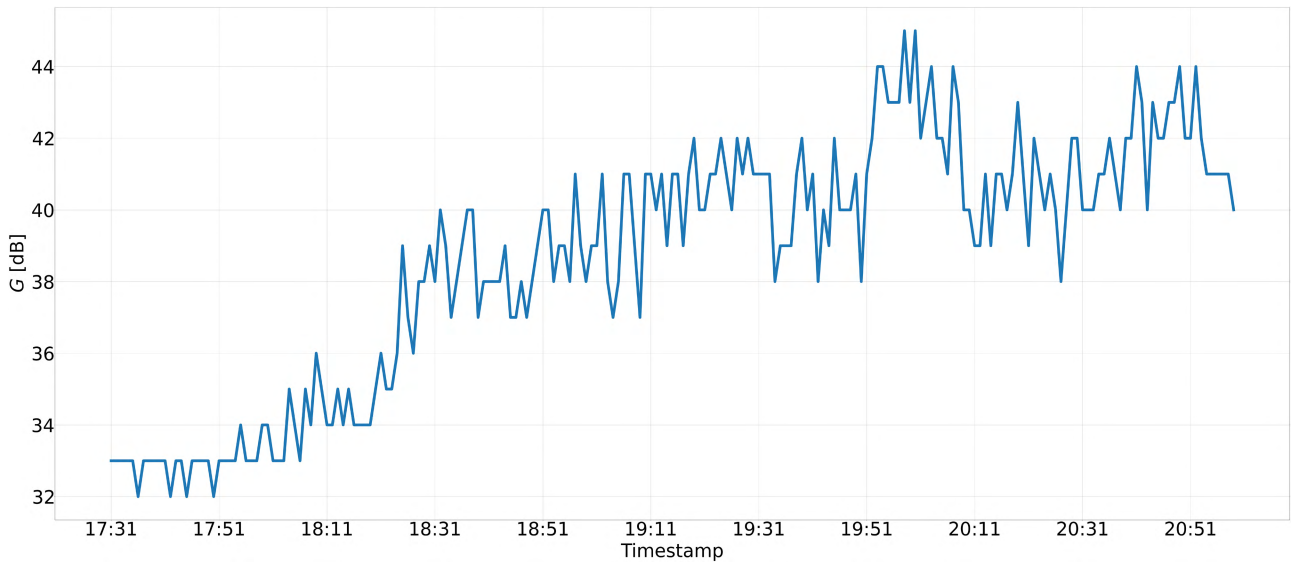

Figure 4. The measurement gain G increases slowly over two and a half hours, rising from 33 to over 40. Through the increase in this additional information a change in the quality of the educts can be diagnosed (Processdrift 1).

sensitivity analysis. The variables with the highest sensitivities are then selected. This procedure must only be applied once, making it significantly less time-consuming and costly.

5 DISCUSSION AND OUTLOOK

This paper demonstrates how additional equipment information can be used in the most beneficial way for an efficient detection and diagnosis of drifts. A hybrid method for detecting and diagnosing drifts was developed in [1]. This method uses additional information to distinguish between drifts. But to do so, it requires manually selecting the most useful information. This process is time-consuming and costly. This renders the method impractical for use in real-world applications. Thus, an automated evaluation of the additional information was introduced to expand the hybrid method. This evaluation is based on sensitivity analysis; in particular, correlation coefficients are used. Besides that, the drift index is introduced during this process. It represents a unique mathematical designation for data points that have been assigned to specific drifts.

The extended method was successfully applied to a scenario involving multiple adjacent devices. A control valve in front of a tank, along with various flow, pressure, and temperature measurements, were examined. Thus, the method was successfully applied to a more realistic case. Furthermore, it was demonstrated that the automated evaluation of the information increases the method's efficiency compared to [1].

However, the method continues to have problems with processes that have varying process conditions or different dynamics. Changes in drifts or new additions are not considered. Future research should focus on developing strategies for dealing with changes in process conditions and higher process dynamics.

ACKNOWLEDGEMENTS

The authors thank the partners of the industrial project RUBENS: LANXESS Deutschland GmbH, SAMSON AG and KROHNE Messtechnik GmbH for the collaboration and discussion.

AUTHOR IDENTIFIERS

Author ORCIDs:
Linda Eydam: 0009-0000-5609-4606
Julius Lorenz: 0000-0001-7124-4844
Leon Urbas: 0000-0001-5165-4459

REFERENCES

7. Eydam L, Furtner L, Lorenz J, Urbas L. Differentiation between process and equipment drifts in chemical plants. Systems and Control Transactions 4:1396-1402 (2025). https://doi.org/10.69997/sct.124377
8. Tan JS, Kramer MA. A general framework for preventive maintenance optimization in chemical process operations. Computers & Chemical Engineering 21:1451-1469 (1997). https://doi.org/10.1016/s0098-1354(97)88493-1
9. Hashemian HM. State-of-the-art predictive maintenance techniques. IEEE Trans. Instrum. Meas. 60:226-236 (2011). https://doi.org/10.1109/tim.2010.2047662

10. Selim MM, Althobaiti S, Ismail HMA, Mohamed SMY. Design of an optimized sensor fault identification within the limited computing, memory, and energy capabilities. AIP Advances 14: (2024). https://doi.org/10.1063/5.0216212
11. V. Venkatasubramanian, R. Rengaswamy, S. N. Kavuri, and K. Yin, "A review of process fault detection and diagnosis: Part I: Quantitative model-based methods, " *Computers & Chemical Engineering*, vol. 27, no. 3, Art. no. 3, Mar. 2003, doi: 10.1016/S0098-1354(02)00160-6.
12. J. de Caigny, T. Himmelsbach, and R. Huck, "NOA-Konzept, NE 175, " in *NAMUR Open Architecture (NOA) Das Konzept zur Öffnung der Prozessautomatisierung*, T. Tauchnitz, Ed., Essen: Vulkan Verlag, 2021, pp. 5–9.
13. Rahm L, Lorenz J, Khaydarov V, Maywald T, Glas T, Urbas L. Predictive maintenance with NOA: application and insights for rotating equipment. 2020 25th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA) :761-767 (2020). https://doi.org/10.1109/etfa46521.2020.9212014
14. Pianosi F, Beven K, Freer J, Hall JW, Rougier J, Stephenson DB, Wagener T. Sensitivity analysis of environmental models: a systematic review with practical workflow. Environmental Modelling & Software 79:214-232 (2016). https://doi.org/10.1016/j.envsoft.2016.02.008

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Reinforcement Learning Supervisory Control with Fuzzy-Logic Reward for Multistage Gas Compression

José R. Torraca Neto[a]*, Sergio A. C. Giraldo[a], Mario C. M. M. Campos[c], Gustavo L. R. Caldas[b], Bruno D. O. Capron[a], and Argimiro R. Secchi[b]

[a] EPQB, School of Chemistry, LADES—Software Development Laboratory, Universidade Federal do Rio de Janeiro, Rio de Janeiro, RJ, Brazil.
[b] Chemical Engineering Program, COPPE, Universidade Federal do Rio de Janeiro, Rio de Janeiro, RJ, Brazil.
[c] SmartAutomation, Rio de Janeiro, RJ, Brazil.
* Corresponding Author: joseneto@eq.ufrj.br.

ABSTRACT

Offshore natural gas compression systems are characterized by strong hydraulic coupling, non-linear behaviour, and strict safety constraints, particularly in high-CO_2 production environments. Conventional decentralized PID control with anti-surge protection ensures reliable local regulation but often leads to poor plant-wide coordination and persistent offsets when multiple compression trains, recycle loops, and separation units interact dynamically. Although multivariable control strategies such as model predictive control can address these issues, their industrial application remains limited by modeling effort, computational demand, and robustness concerns. This work presents a hybrid supervisory control framework in which reinforcement learning (RL) augments an existing PI-based architecture for an offshore gas compression system with membrane-based CO_2 separation. A Proximal Policy Optimization (PPO) agent is trained on a dynamic digital-twin model of export, CO_2, injection, and bypass compression trains with shared headers, recycle flows, and thermal constraints. The RL agent acts exclusively at the supervisory level, providing bounded incremental adjustments to pressure and bypass setpoints, while all regulatory PI loops, anti-surge protections, and safety logic remain unchanged. Closed-loop simulations under representative disturbance scenarios show that the RL-supervised architecture significantly improves plant-wide coordination. Compared to fixed-setpoint PI operation, total tracking error (IAE) and time-weighted error (ITAE) are reduced by approximately 54% and 56%, respectively, using smooth, low-magnitude supervisory actions. In addition, RL achieves performance close to that of a non-linear MPC benchmark based on a random-shooting, MPPI (Model-Predictive Path Integral) formulation, while requiring substantially lower online computational effort. The results demonstrate that RL can be effectively integrated as a supervisory layer in offshore gas compression systems without replacing established industrial control infrastructure, although additional safety mechanisms are still required to ensure stricter constraint satisfaction.

Keywords: Offshore gas compression, reinforcement learning, supervisory control, PI control

INTRODUCTION

Natural gas compression systems play a central role in offshore oil and gas production, particularly in fields with high CO_2 content such as Brazil's pre-salt reservoirs. In these environments, compression is essential for enabling CO_2 separation, gas export, reinjection, and downstream processing, while simultaneously representing one of the largest contributors to onboard energy consumption and operational cost [1, 2]. Modern offshore facilities typically rely on multistage centrifugal compressors operating under tight constraints on pressure, flow rate, rotational speed, and surge margin, which must be satisfied under highly variable operating conditions.

In industrial practice, gas compression trains are commonly controlled using layered architectures based on proportional–integral–derivative (PID) controllers, complemented by anti-surge protection systems and

hard safety interlocks [3]. While such architectures provide reliable local control and are well established in industry, their performance deteriorates in highly coupled systems where multiple compressors, recycle valves, membranes, and headers interact dynamically. In offshore facilities, interactions between parallel compression trains create strong loop interactions, often resulting in oscillatory behaviour, conservative operation, and limited plant-wide coordination [4].

To overcome these limitations, advanced control strategies have been extensively investigated in offshore applications over the last two decades. Model Predictive Control (MPC), in particular, has been successfully applied to gas processing and compression systems, improving constraint handling, multivariable coordination, and operational efficiency compared to decentralized PID control [5, 6]. Nevertheless, the deployment of MPC in large-scale offshore compression systems remains challenging due to the high modeling effort required, the nonlinear nature of compressor maps, frequent changes in operating regimes, the need for robust performance under significant disturbances and uncertainty, and the required computational infrastructure.

At the same time, increasing attention has been given to the environmental impact of offshore gas production. Compression systems are responsible for a substantial fraction of platform-wide energy consumption, typically supplied by gas turbines, which account for the majority of CO_2 emissions in offshore installations [7]. In addition, inefficient operation may lead to increased flaring, recycle flows, and pressure violations, further increasing emissions and reducing overall efficiency. As a result, improving the operational efficiency and stability of gas compression systems is directly aligned with broader efforts to reduce greenhouse gas emissions and comply with international climate targets [8].

In this context, data-driven control methods based on Reinforcement Learning (RL) have emerged as a promising alternative for the control of complex, nonlinear, and highly coupled process systems. By learning control policies through interaction with a dynamic environment, RL techniques can address multivariable coordination problems without relying on explicit linearized models, making them attractive for systems characterized by nonlinear dynamics, strong interactions, and changing operating conditions [9, 10]. Recent studies have reported encouraging results for RL-based control and optimization in energy systems and chemical processes, including continuous control problems with constraints, disturbances, and nonlinear dynamics [11]. These works show the potential of modern policy-gradient and actor–critic methods to outperform conventional controllers in terms of tracking performance, efficiency, and adaptability.

Nevertheless, the direct application of RL to safety-critical industrial systems remains limited. Key challenges include guaranteeing constraint satisfaction, ensuring robustness under unmodelled disturbances, maintaining interpretability for operators, and safely integrating learning-based components with existing control infrastructures [12]. These concerns have motivated increasing interest in hybrid and hierarchical architectures that combine RL with conventional control layers.

A pragmatic approach to mitigate these limitations consists in embedding RL within hierarchical control architectures, in which learning-based agents operate at a supervisory level, while conventional controllers remain responsible for direct actuator manipulation. In such hybrid schemes, the RL agent adjusts setpoints or high-level targets, whereas low-level PID controllers, anti-surge loops, and safety interlocks enforce fast dynamics and operational constraints [13]. This structure preserves the reliability and transparency of industrial control systems while enabling adaptive, plant-wide coordination and optimization. Related concepts have been explored in supervisory, economic, and learning-based control layers; however, their application to multistage offshore gas compression systems with CO_2 separation and strong hydraulic coupling between trains remains relatively underexplored.

This work proposes a hybrid PI–RL supervisory control strategy for an offshore gas compression system. It comprises multiple interacting compression trains and a membrane-based CO_2 separation unit. A dynamic plant model was developed as a simplified digital twin, incorporating centrifugal compressor models, composition-dependent thermodynamics, valve flow dynamics, recycle loops, and operational constraints representative of industrial practice [1, 4]. The RL agent, trained using Proximal Policy Optimization (PPO) [14], acts exclusively on supervisory setpoints related to pressure and capacity, while conventional PI controllers and anti-surge protections remain unchanged.

To ensure learning stability and reflect operational priorities, a structured reward-shaping strategy was adopted. In addition to tracking export flow and key suction pressures, the reward incorporates soft penalties for violations of compressor discharge temperature limits, excessive anti-surge recycle activity, undesirable valve usage, and abrupt setpoint movements. Temperature constraints are enforced through continuous penalties activated beyond predefined soft limits, discouraging sustained operation at elevated discharge temperatures that could compromise equipment integrity and efficiency. This multi-term reward formulation, inspired by fuzzy-logic concepts and robust loss functions, introduces tolerance regions around nominal operating targets while penalizing asymmetric deviations, thereby preventing overly aggressive control actions and improving training robustness under disturbances [15].

The proposed approach was evaluated using closed-loop simulations based on real offshore gas compression system subject to disturbances in inlet flows and operating conditions. Performance was assessed relative to conventional fixed-setpoint PI operation using standard control performance metrics.

METHODOLOGY

1. Case study and control objective

This work considers the offshore gas-compression and high-CO_2 management system of an offshore platform. The unit has strong interactions between separation, compression, export, and reinjection. Figure 1 presents the process flow diagram (PFD) used as the reference for the modeling and control design. The system comprises: (i) a membrane-based CO_2 separation unit generating retentate (treated gas) and permeate (CO_2-rich gas), (ii) export compression, (iii) CO_2 compression for reinjection/EOR, (iv) injection compression, and (v) bypass routing that allows partial or full diversion around the membrane unit depending on the operating mode (export vs. injection priority).

The routing and coupling structure that drives the control problem can be seen in Figure 1: after the membrane split step, the retentate flow is routed for compression and exportation, while the permeate flows for CO_2 compression. The bypass stream modifies both membrane feed and downstream load. Shared headers link export and injection trains; thus, changes in exportation demand or bypass fraction propagate through suction pressures, discharge pressures, and recycle activity.

The operational goal is to maintain safe and efficient operation under disturbances (production flow swings, routing changes, header pressure shifts) while ensuring:

- Header and train stability: suction/discharge pressures remain within allowable envelopes to prevent compressor trips and surge;
- Export delivery tracking: export flow (or export capacity) follows a target consistent with pipeline constraints and gas quality;

CO_2 management: membrane operation and bypass routing are coordinated so that the CO_2 content in the treated gas (when exporting) is kept within specification, while CO_2-rich permeate is compressed for reinjection when required;

- Coupling-aware coordination: shared headers and routing decisions couple export and injection demands, and membrane permeate pressure is coupled to the suction setpoint of the CO_2 train, impacting separation driving force and downstream compression load.

These couplings motivate a hierarchical control architecture. Fast stabilization and protection remain handled by conventional regulatory control (PI loops, anti-surge and operational interlocks), while a supervisory layer performs slower coordination by adjusting a small set of economically and operationally meaningful setpoints (pressures and bypass routing), targeting global feasibility, constraint compliance, and reduced recycle/energy waste.

2. Dynamic simulation model

A dynamic simulator was implemented in Python to enable closed-loop testing and training of the supervisory controller. The model follows a simplified digital-twin philosophy: it preserves the dominant pressure–flow accumulation dynamics, train couplings (shared headers, membrane routing), and actuator nonlinearities needed for control design, while remaining computationally efficient for RL rollouts.

2.1 Compressor trains and stage calculations

Each compression section is modeled as a series of centrifugal compressor stages, each represented by a performance map embedded within a dynamic mass and energy balance. For each stage, an iterative procedure matches volumetric flow and operating point to compute polytropic head, efficiency, power, and discharge temperature.

Composition effects are captured through mixture properties (molecular weight, Mw, and compressibility factor, Z), allowing the simulator to represent performance shifts due to varying CO_2 fraction (especially relevant across the membrane split and bypass changes).

2.2 Vessel/header pressure dynamics

Dynamic pressure behavior is represented by vessels/headers that capture the dominant accumulation effects in each section, providing the state evolution used by the control loops and the RL observation vector. (These dynamics are exposed to the RL agent primarily through suction and discharge pressures per train, see Section 4.)

2.3 Valve and recycle-flow modeling

Valve flows are modeled using a compressible-flow relationship, given by Equation (1).

$$\dot{m} = \phi(u) K p_{up} C_v \sqrt{\frac{(p_{up} - p_{down})^+}{p_{up} T Z}} \quad (1)$$

where u is the valve opening, $\phi(u)$ represents the installed valve characteristic (linear, equal-percentage, or quick-opening), p_{up} and p_{down} are upstream and downstream pressures, T is absolute temperature, Z is compressibility factor, C_v is the valve coefficient, and K is a unit-conversion constant. The "$(\bullet)^+$" operator prevents non-physical negative pressure drops.

Anti-surge behavior is represented through recycle flows per train/stage, which are also included as monitored variables and (optionally) penalized in the reward

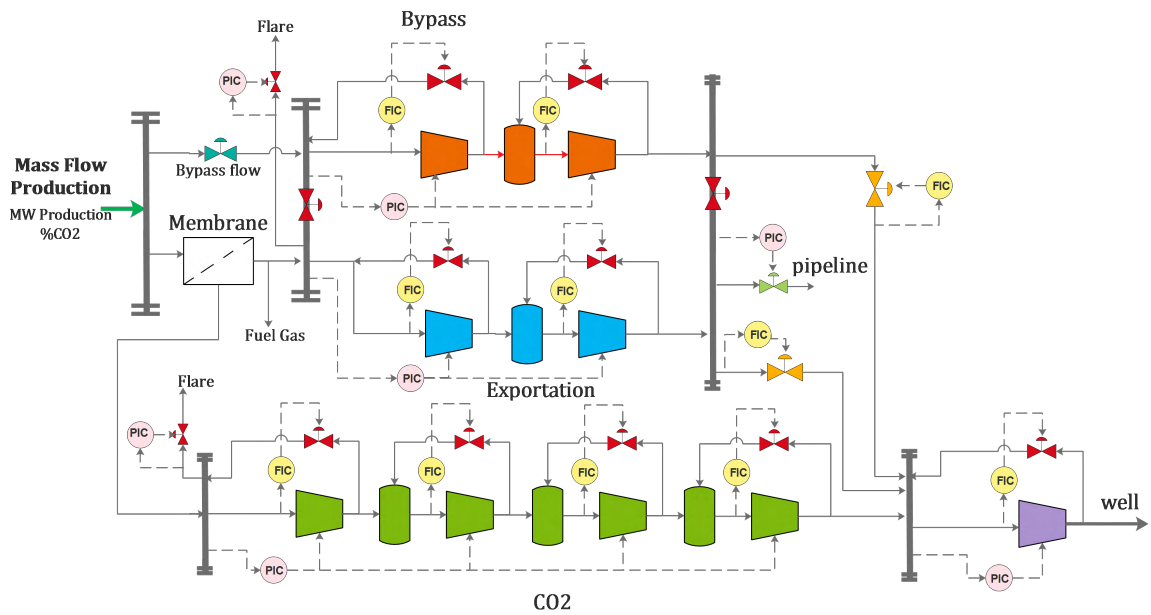

Figure 1. Flowchart of the compression system of the real offshore platform.

to discourage unnecessary recycle activity.

2.4 Membrane CO_2 separation unit

The membrane unit is modeled as a pressure-driven separation process that splits the feed into retentate (treated gas) and permeate (CO_2-rich) streams. An empirical bias term is included to represent non-ideal effects such as aging or performance degradation, while outlet molecular weights are updated based on the resulting gas composition.

The retentate is routed to the export compression section, while the permeate is routed to the CO_2 compression section. The permeate-side pressure is directly linked to the CO_2 compression suction setpoint, explicitly coupling membrane separation performance with downstream compression conditions and supervisory control actions.

3. Disturbance modelling

To assess the robustness and system-level coordination of the supervisory controller, a severe load disturbance is considered. At time t=400s, the production mass flow is subjected to a step increase of 20 kg/s. This scenario imposes an immediate thermal and mechanical load on the compression system, leading to increased discharge pressures and temperatures across all compression trains, including the export, CO_2, and injection sections.

4. Hierarchical control architecture

The proposed control structure is hierarchical and explicitly preserves the regulatory and protection logic of the compression system. The inner layer consists of conventional PI controllers, anti-surge loops at each compression stage, actuator saturation, and split-range logic.

The RL agent operates strictly at the supervisory level and does not bypass or override the inner-loop protections. Instead, it provides bounded adjustments to selected setpoints that are already used by the PI controllers.

4.1 Action space

The RL agent operates in delta (incremental) mode, outputting bounded changes to supervisory setpoints at each control interval. This choice enforces smooth behavior, avoiding abrupt operating-point shifts. The action vector and its numerical bounds are defined in Table 1.

Table 1: Action space configuration.

Action	Description	Lower Bound	Upper Bound	Units
$\Delta\dot{m}_{bypass}$	Bypass flow target	-1	+1	kg/s
$\Delta p^{sp}_{s,exp}$	Export suction pressure setpoint	-500	+500	kPa
$\Delta p^{sp}_{d,exp}$	Export discharge pressure setpoint	-500	+500	kPa
$\Delta p^{sp}_{s,co2}$	CO_2 suction pressure setpoint	-5	+5	kPa
$\Delta p^{sp}_{s,inj}$	Injection suction pressure setpoint	-50	+50	kPa

4.2 Observation vector

The observation vector consists of 28 continuous variables, selected to expose the dominant coupling mechanisms, constraint indicators, and dynamic states relevant for supervisory coordination. Specifically, the environment returns:

- Time $t(\mathrm{s})$;
- Suction, intermediate, and discharge pressures for the bypass, export, CO_2, and injection sections;
- Compressor "rotation proxies" for bypass, export, CO_2, and injection trains;
- Mass flows: bypass, membrane outputs to export and CO_2, and export inlet flow;
- Treated-to-injection valve opening;
- Recycle volumetric flows for bypass and each compression section/stage;
- Discharge temperatures for export stage 2 and CO_2 stage 4;
- Current export capacity setpoint (included as a state variable for consistency).

5. Reward design

A multi-term reward is used to align learning with operational goals: (i) tracking/coordination, (ii) anchoring around desired operating targets, and (iii) safety penalties dominated by discharge temperature constraints. The instantaneous reward is defined by Equation (2).

$$r = r_{track} + r_{temp} + r_{soft} \tag{2}$$

5.1 Tracking

Tracking terms penalize deviations in suction/discharge pressures and export flow, with weights assigned to reflect operational priorities. In particular, CO_2 pressures receive stronger emphasis because they directly affect membrane driving force. An additional target anchoring term penalizes excessive deviation of supervisory setpoints from nominal operating targets.

5.2 Temperature constraints

Thermal safety is encoded explicitly through continuous penalties activated beyond stage-specific discharge temperature limits. The following limits are used:

- $T_{lim,bypass} = 390\ K, T_{lim,export} = 410\ K,$
- CO_2 stages: $[400, 440, 420, 420]\ K$ (stages 1-4),
- weights: $w_{temp,co2} = 6.0, w_{temp,export} = 2.0, w_{temp,bypass} = 2.0, w_{temp,injection} = 1.0.$

A soft margin of 5 K is applied before the penalty increases steeply, and temperature exceedances are normalized by a 10 K scale.

5.3 Soft penalties

Additional soft penalties are included to promote operational efficiency and smooth control behavior. Recycle activity is aggregated from recycle flows across sections: bypass, export (stages 1–2), CO_2 (stages 1–4), and injection. Valve usage is represented by the treated-to-injection valve opening, and action smoothness is penalized via the L_1 norm of the action vector.

6. RL algorithm and training procedure

The supervisory policy is trained using PPO [14] with a multilayer perceptron (MLP) policy consisting of two hidden layers of 256 neurons each.

The main hyperparameters are: learning rate 3×10^{-4}, discount factor γ = 0.99, GAE λ = 0.95, clipping range 0.2, n_{steps} = 2048, n_{epochs} = 10, batch size 256, and random seed 42.

Each episode lasts 3000 simulation steps with Δt=0.1 s, corresponding to 300 s of plant operation. Training was performed on a CPU-only workstation (24 logical cores) for a total of 200 episodes ($\approx 6\times10^5$ steps in total).

During training, the best-performing policy is selected based on the rolling mean of episode rewards and saved for post-training evaluation. After training, the selected policy is evaluated in closed-loop simulations using deterministic action selection.

7. NMPC benchmark controller

As a benchmark, a stochastic nonlinear model predictive control (NMPC) strategy was implemented using a random-shooting approach inspired by Model Predictive Path Integral (MPPI) control [16].

The same supervisory action space and objective function used by the RL controller was adopted. The final configuration used a prediction horizon of 8 steps and a control horizon of 2 steps, with 16 candidate sequences evaluated every 6 simulation steps, and 4 elite samples used for action selection.

RESULTS AND DISCUSSION

8 Training convergence

The PPO-based supervisory policy converged within approximately 80–100 episodes (Figure 2). With an average simulation throughput of 129 steps/s, the complete training required approximately 77.5 minutes.

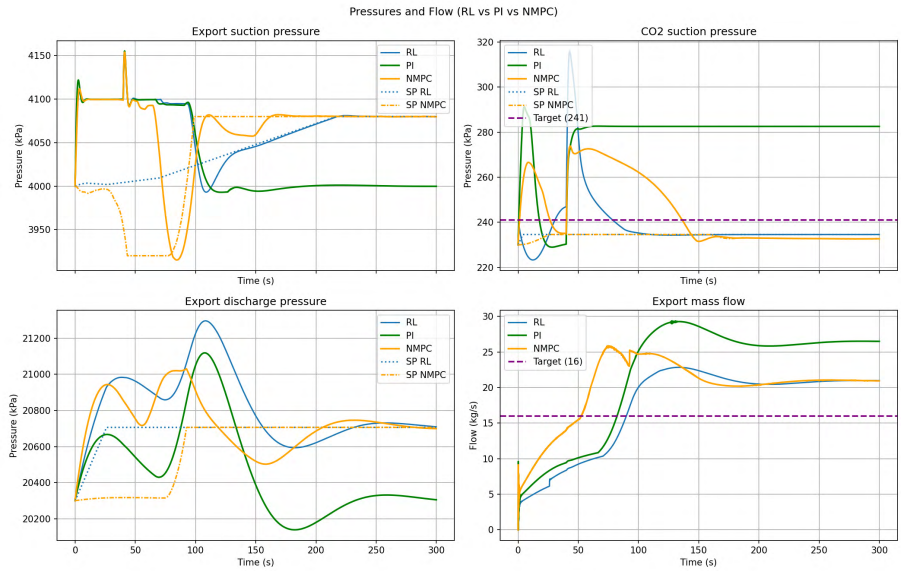

Figure 3. Closed-loop pressures and export flow under PI, RL, and NMPC control.

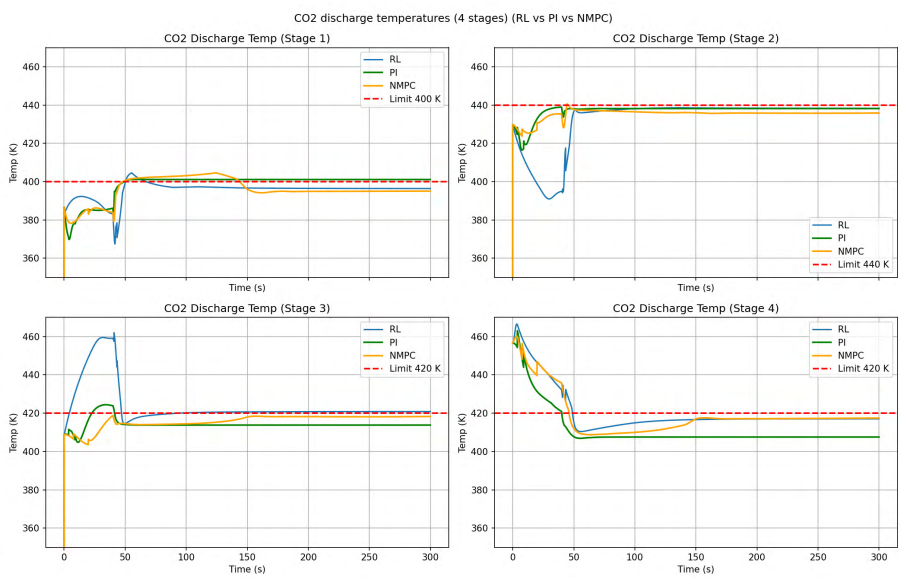

Figure 4. CO_2 compressor discharge temperatures (stages 1–4) under PI, RL, and NMPC control; dashed lines indicate temperature limits.

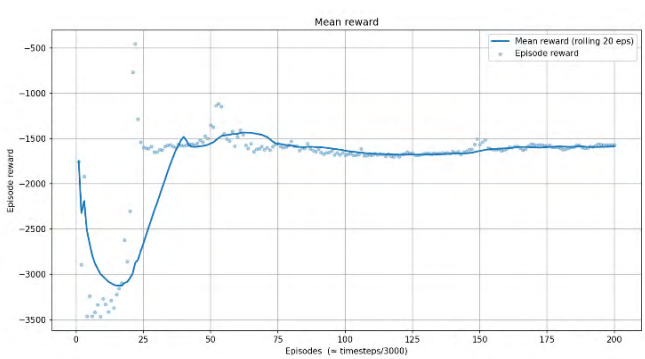

Figure 2. Training mean reward versus episodes.

9 Closed-loop control

Figure 3 shows that RL and NMPC converge to similar steady states, both outperforming PI. Their main difference is in the transient response: NMPC yields smoother but slower corrections, while RL applies more aggressive initial actions before stabilizing.

Figure 4 highlights that RL produces larger temperature excursions, particularly at CO_2 stage 3, where constraint violations are more frequent. These results indicate that RL's aggressive transient actions improve coordination but weaken constraint satisfaction.

10 Quantitative performance comparison

Table 2 compares RL, PI, and NMPC in terms of tracking, energy, safety, control effort, and computational cost.

Table 2: Controllers' performances.

Metric	RL	PI	NMPC
IAE	322, 462	699, 682	310, 400
ITAE	48.94×10^6	111.48×10^6	46.54×10^6
Mean power (HP)	25, 979	25, 090	24, 212
Total abs. power change (HP)	28, 241	30, 572	31, 115
Temp. violation total count	5, 701	3, 178	3, 053
Temp. violation total magnitude	43, 252	35, 423	30, 931
$\int u^2\,dt$	1202.20	0.00	2765.63
Evaluation time (s)	71.04	65.28	3895.77

In terms of tracking, RL and NMPC achieve comparable performance and clearly outperform PI, with similar IAE and ITAE values. From an energy perspective, NMPC achieves the lowest mean power, while RL operates at a slightly higher level. However, RL exhibits the lowest total absolute power variation, indicating less aggressive load changes despite similar steady-state performance.

Regarding safety, RL shows the highest temperature-violation count and magnitude, confirming that the current reward-based formulation allows stronger transient excursions. In terms of control effort, RL is significantly smoother than NMPC, with much lower integral quadratic control effort ($\int u^2 dt$), indicating less aggressive supervisory actions.

Finally, computational cost highlights a key advantage of RL. While RL and PI have similar evaluation times, NMPC is more than one order of magnitude slower.

11 Limitations and scope of applicability

The present work is based on a simplified digital twin and does not account for all sources of uncertainty present in real offshore systems, such as sensor noise, actuator degradation, and long-term equipment aging. The evaluation was conducted under a limited set of disturbance scenarios. Although the RL policy generalizes well within this range, its robustness to broader operating conditions remains to be assessed.

Safety is handled through soft penalties, which leads to transient constraint violations. Practical deployment would require additional mechanisms such as hard constraint enforcement or action shielding. Finally, the study considers a single-agent supervisory policy. Extensions to multi-agent formulations, economic objectives, and online adaptation remain open research directions.

CONCLUSIONS

This work showed the feasibility and benefits of integrating reinforcement learning as a supervisory layer in a multistage offshore gas compression system, improving plant-wide coordination while preserving conventional PI control and safety logic.

The RL-supervised strategy significantly outperforms PI, reducing IAE and ITAE by approximately 54% and 56%, respectively. Its performance is also close to NMPC, with differences of only about 4–5%, and both approaches converge to similar steady-state operating points.

From an operational perspective, RL provides smoother actuation, as the quadratic control effort is reduced by approximately 57% compared to NMPC. In terms of energy, NMPC achieves the lowest mean power consumption, while RL operates about 7% above NMPC and slightly above PI.

A key limitation of the current approach is safety performance. RL exhibits higher temperature violation frequency and magnitude, with approximately 80–90% more violation events and about 40% higher total violation magnitude compared to NMPC. These results reflect the use of soft penalties in the reward function, which allow stronger transient excursions (particularly during early adaptation phases) before converging to safer operation.

Finally, computational cost strongly favors RL. While NMPC requires approximately 55× more evaluation time, the RL policy can be executed with negligible online overhead. Combined with a moderate offline training time (≈77.5 minutes on CPU), this makes RL a computationally efficient alternative for real-time supervisory control.

Future work will focus on strengthening safety guarantees through hard constraint enforcement, extending evaluation across multiple disturbance scenarios, and assessing robustness and transferability toward real offshore operating conditions.

ACKNOWLEDGEMENTS

Jose Rodrigues Torraca Neto reports financial support provided by CAPES, Finance Code 001, and Petrobras (Cooperation term 0050.0124501.23.9). Argimiro R. Secchi Argimiro R. Secchi is grateful to financial support from CNPq (Grants No. 300744/2025-0). Other authors declare that they have no known competing

financial interests or personal relationships that could have appeared to influence the work reported in this paper.

AUTHOR IDENTIFIERS

Author ORCIDs:
Torraca JR Neto: 0009-0000-4127-5682
Giraldo SAC: 0000-0001-6253-564X
Campos MCMM: 0000-0002-5746-6915
Caldas GLR: 0000-0003-2554-9561
Capron BDO: 0000-0001-5497-9769
Secchi AR: 0000-0001-7297-3571

REFERENCES

1. Mokhatab S, Poe WA, Mak JY. Phase separation. Handbook of Natural Gas Transmission and Processing :137-167 (2015). https://doi.org/10.1016/b978-0-12-801499-8.00004-3
2. Jones K, Cortinovis A, Mercangoez M, Ferreau HJ. Distributed model predictive control of centrifugal compressor systems. IFAC-PapersOnLine 50:10796-10801 (2017). https://doi.org/10.1016/j.ifacol.2017.08.2343
3. Plucenio A, Normey-Rico J, Vettorazzo C, Campos MM, De Lima ML. MPC advanced control of an offshore gas compression system. OTC Brasil : (2017). https://doi.org/10.4043/28129-ms
4. Gravdahl JT, Willems F, de Jager B, Egeland O. Modeling for surge control of centrifugal compressors: comparison with experiment. Proceedings of the 39th IEEE Conference on Decision and Control (Cat. No.00CH37187) 2:1341-1346 (None). https://doi.org/10.1109/cdc.2000.912043
5. Ribeiro CHP, Miyoshi SC, Secchi AR, Bhaya A. Model predictive control with quality requirements on petroleum production platforms. Journal of Petroleum Science and Engineering 137:10-21 (2016). https://doi.org/10.1016/j.petrol.2015.11.004
6. Giraldo SAC, Supelano RC, d'Avila TC, Capron BDO, Ribeiro LD, Campos MM, Secchi AR. Model predictive control with dead-time compensation applied to a gas compression system. Journal of Petroleum Science and Engineering 203:108580 (2021). https://doi.org/10.1016/j.petrol.2021.108580
7. International Energy Agency. CO_2 emissions in 2023. https://www.iea.org/reports/co2-emissions-in-2023
8. United Nations Framework Convention on Climate Change. Paris agreement. https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement
9. Sutton RS, Barto AG. Reinforcement Learning: An Introduction. 2nd ed. MIT Press (2018).
10. Lillicrap TP, Hunt JJ, Pritzel A, Heess N, Erez T, Tassa Y, Silver D, Wierstra D. Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971 (2015).
11. Devarakonda VS, Sun W, Tang X, Tian Y. Recent advances in reinforcement learning for chemical process control. Processes 13:1791 (2025). https://doi.org/10.3390/pr13061791
12. García J, Fernández F. A comprehensive survey on safe reinforcement learning. J Mach Learn Res 16:1437–1480 (2015).
13. Qin SJ, Badgwell TA. A survey of industrial model predictive control technology. Control Engineering Practice 11:733-764 (2003). https://doi.org/10.1016/s0967-0661(02)00186-7
14. Schulman J, Wolski F, Dhariwal P, Radford A, Klimov O. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017).
15. Ng AY, Harada D, Russell SJ. Policy invariance under reward transformations: theory and application to reward shaping. In: Proceedings of the Sixteenth International Conference on Machine Learning. Morgan Kaufmann Publishers, San Francisco, CA:278–287 (1999).
16. Williams G, Aldrich A, Theodorou EA. Model predictive path integral control: from theory to parallel computation. Journal of Guidance, Control, and Dynamics 40:344-357 (2017). https://doi.org/10.2514/1.g001921

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Relating Loss Geometry to Empirical Generalization in Recurrent Neural Net Surrogates: Three Tanks Case Study

Ricardo M. Roxas II[a] and Karl Ezra Pilario[a]*

[a] Process Systems Engineering Laboratory, Department of Chemical Engineering, University of the Philippines Diliman, Quezon City, 1101, Philippines
* Corresponding Author: kspilario@up.edu.ph.

ABSTRACT

Recurrent neural nets (RNNs) are now commonly used for the surrogate modeling of process systems, leading to better control and faster real-time optimization. However, when trained with small training data sets, most experiments show that RNNs exhibit poor generalization abilities outside the range of the training data space. Nonetheless, recent advances in deep learning research have shown that certain characteristics of the loss landscape of trained models, such as the flatness around the local minimum, tend to relate to generalization ability. This paper investigates this phenomenon for the case of RNN surrogates of the well-known Three Tanks case study, which is representative of many continuous processes. We trained a total of 200 LSTMs (long short-term memory networks) differing in initialization, architecture, and training dynamics on the same data of 500 samples. The number of model parameters ranges from 238 to 11, 353. We estimated the loss curvature of each trained model using the Hessian-vector products method and then evaluated their test accuracy on 100 pre-generated data sets with 50% larger input amplitude than the training data to force extrapolation. We found that the estimate of the top eigenvalue of the Hessian matrix is 75% rank-correlated to the one-step-ahead prediction accuracy averaged on all test data sets. We report other insights on the impact of other Hessian-based metrics to generalization ability. Our study can potentially inform how to control the training dynamics of RNN surrogates to improve their extrapolation ability in the absence of first-principles knowledge.

Keywords: System Identification, Machine Learning, Artificial Intelligence, Derivative Free Optimization, Dynamic Modelling, Hessian vector products, Generalization

INTRODUCTION

Nonlinear system identification tools are now used for the surrogate modelling of continuous industrial processes. Common tools include Koopman models, kernel subspace methods, and neural nets [1]. The recurrent neural net (RNN) is a prominent type of deep learning model suitable for multivariate time series data. Previous works already used RNNs to achieve nonlinear model predictive control (MPC) [2], fault detection [3-4], and process optimization [5] in various case studies.

Despite these applications, the generalizability of RNNs is still poorly understood. In cases where industrial data is scarce or has poor quality, achieving a high validation accuracy does not guarantee good performance during deployment [2]. When the model encounters out-of-distribution data or if it is forced to extrapolate from the training data space, model errors may amplify as it makes future predictions. Yet researchers aim to train valid RNN models from start-up to shutdown [6].

One way to understand the validity of RNNs is to derive generalization bounds. For instance, based on the Rademacher complexity approach [7], RNNs were proven to generalize better with more training samples and simpler architectures. Meanwhile, vanilla RNNs with ReLU (rectified linear unit) activation were shown to generalize better with lower matrix 1-norm and Fisher-Rao norm of the weights, which is consistent with improvements via weight decay and gradient clipping. Recently, [8] obtained tighter bounds on the RNN generalization error using the Frobenius norms of the weights, still pointing to the importance of weight decay in practice.

While theoretical results make clear steps forward in analyzing the validity of RNNs, empirical results are also

helpful especially for finding RNN architectures that generalize well. A recent work validated LSTM surrogates for a boiler, a debutanizer, and two nonlinear reactors [9]. In a polymerization case study, [10] recommends the use of partially input convex RNNs for better accuracy and computational efficiency under MPC. To find good LSTM architectures, [11] advocates for the use of state-of-the-art hyper-parameter optimization methods. This idea is now widely accepted in many LSTM surrogate modeling studies [6], [12].

Aside from architectural choices, there is a growing interest in relating the geometry of the training loss landscape of neural nets to their generalizability. As early as [13], the flatness of the local minima of the training loss was already tied to better generalization (the authors are also the originators of the LSTM), i.e. models that converged on flat minima generalize better than those on sharp minima. In the past, many flatness or sharpness metrics have been defined based on measuring the loss curvature at the minimum, such as the Hessian eigenspectrum [14], Hessian spectral radius, Hessian trace, top Hessian eigenvalue [15], and ϵ-sharpness [16]. Some definitions were motivated by empirical observations while others are more theoretical, such as those based on PAC-Bayes (Probably Approximately Correct-Bayesian) generalization bounds [17]. Some sharpness metrics, however, can be arbitrarily changed by scaling the weights without changing the neural net behavior such as the Hessian eigenspectrum [15]. Hence, scale-invariant metrics are more favorable such as the normalized sharpness [17] or the effective dimensionality (effective rank) of the Hessian of the loss [18], [19]. Due to these works, sharpness has been used to control neural net training to give models that generalize better [20-22].

In terms of regularization, weight decay is known to improve generalization, but the most common L^2 regularization is only one approach. Spectral norm regularization of weight matrices [23] is another approach based on the notion that models generalize better if they are insensitive to large input perturbations. For the case of RNNs, penalizing the largest singular value of the recurrent weight matrix was found to provide smoother state transitions [24] which is another form of regularization.

In this paper, we are interested in relating the training loss curvature to generalization in RNNs for surrogate modeling in continuous industrial processes. Our belief that generalization is related to some form of smoothness in RNNs is motivated by our previous work that pitted RNNs against a kernel method for MPC, namely mixed kernel canonical variate analysis (MKCVA) [25]. MKCVA was found to generalize better in the extrapolation regime, despite being at par in terms of median accuracy with LSTMs at the interpolation regime. Hence, MKCVA gave better control than LSTMs at various setpoints. We attributed this superior performance to the use of kernels that encode smoothness properties. Knowing this, our present aim is to attempt to find some geometry-based indicators of test-time error in trained RNNs. Consistent with [18], we do not view kernel methods and neural nets to be competing, rather, we are interested in finding elements of any representation learner that can improve generalization in surrogate modeling.

To achieve this aim, we take a representative industrial case study, namely the Three Tanks system [25], and train numerous LSTMs on a single training data set from a noisy simulation of this system. These LSTMs differ in architecture, training dynamics, and weight initialization. We then compute loss curvature metrics such as the top Hessian eigenvalue, the anisotropy ratio, the effective rank of the Hessian, and including the spectral norm of the recurrent weights of each LSTM, then connect them to 1-step ahead prediction error averaged over 100 pre-generated test data sets from the same noisy simulator but with 50% larger input signals. We hope that our results can help inform the control of LSTM training towards models with better generalization.

Preliminaries

We consider the vanilla LSTM network defined by:

$$\boldsymbol{z}_k = \sigma_{io}(\boldsymbol{W}_z\boldsymbol{x}_k^* + \boldsymbol{R}_z\boldsymbol{y}_{k-1}^* + \boldsymbol{b}_z) \tag{1}$$

$$\boldsymbol{i}_k = \sigma_{rec}(\boldsymbol{W}_i\boldsymbol{x}_k^* + \boldsymbol{R}_i\boldsymbol{y}_{k-1}^* + \boldsymbol{b}_i) \tag{2}$$

$$\boldsymbol{f}_k = \sigma_{rec}(\boldsymbol{W}_f\boldsymbol{x}_k^* + \boldsymbol{R}_f\boldsymbol{y}_{k-1}^* + \boldsymbol{b}_f) \tag{3}$$

$$\boldsymbol{c}_k = \boldsymbol{z}_k \odot \boldsymbol{i}_k + \boldsymbol{c}_{k-1} \odot \boldsymbol{f}_k \tag{4}$$

$$\boldsymbol{o}_k = \sigma_{rec}(\boldsymbol{W}_o\boldsymbol{x}_k^* + \boldsymbol{R}_o\boldsymbol{y}_{k-1}^* + \boldsymbol{b}_o) \tag{5}$$

$$\boldsymbol{y}_k^* = \sigma_{io}(\boldsymbol{c}_k) \odot \boldsymbol{o}_k \tag{6}$$

where $\boldsymbol{x}_k^* = \{[\boldsymbol{u}_t^T, \boldsymbol{y}_{t-1}^T]^T, t \in [k - \text{lookback}, \dots, k]\}$ are the LSTM inputs, $\boldsymbol{y}_k^* = \{\boldsymbol{y}_t, t \in [k - \text{lookback}, \dots, k]\}$ are the LSTM outputs, $\boldsymbol{u}_t$ and $\boldsymbol{y}_t$ are the inputs and outputs of the process, $\boldsymbol{W}_z, \boldsymbol{W}_i, \boldsymbol{W}_f, \boldsymbol{W}_o$ are input weights, $\boldsymbol{R}_z, \boldsymbol{R}_i, \boldsymbol{R}_f, \boldsymbol{R}_o$ are the recurrent weights, $\boldsymbol{b}_z, \boldsymbol{b}_i, \boldsymbol{b}_f, \boldsymbol{b}_o$ are the biases, $\boldsymbol{c}_k, \boldsymbol{i}_k, \boldsymbol{f}_k, \boldsymbol{o}_k$ are the cell state and gate outputs, σ_{io} is the input-output activation, σ_{rec} is the recurrent activation, and $\odot$ is element-wise multiplication.

For the metrics derived for LSTMs, we denote $\boldsymbol{H}$ as the local Hessian matrix of the loss, $tr(\boldsymbol{H})$ is the trace, and λ is an eigenvalue of $\boldsymbol{H}$. We also measure the singular values of weights matrix in a trained LSTM, which we denote as σ with no subscript, to distinguish them from activation functions σ_{io} and σ_{rec} in the LSTM equations.

METHODOLOGY

To find the best LSTM surrogate for the process system while also encouraging diversity in the set of models from which to reason about curvature effects, we

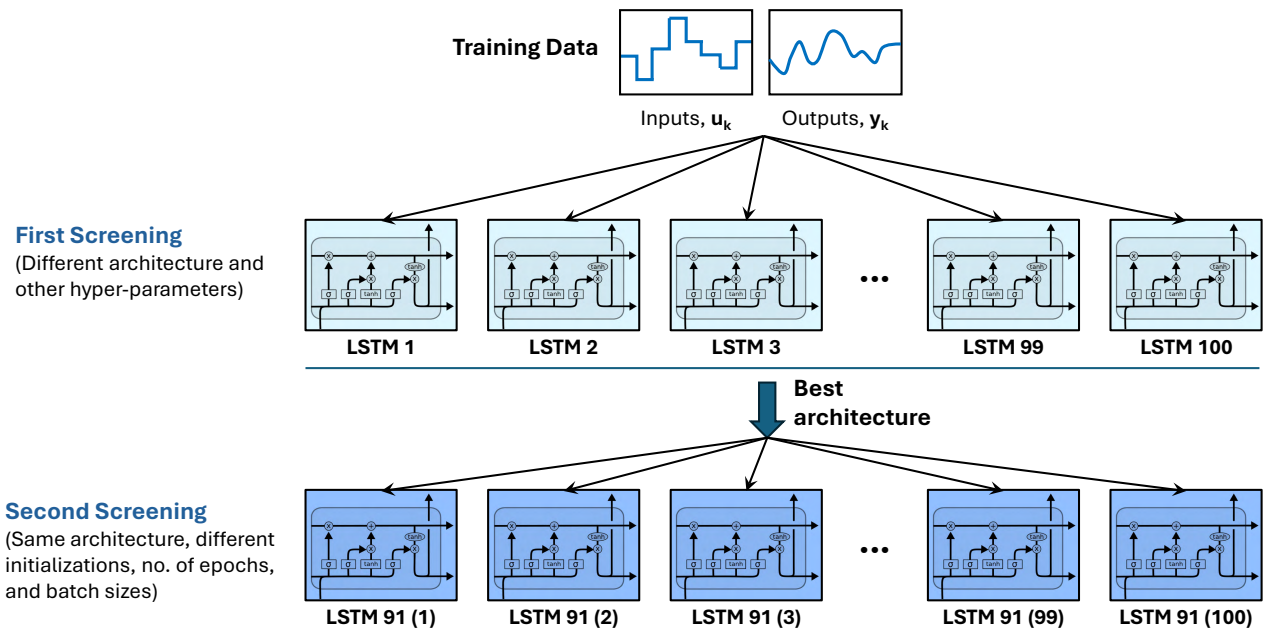

Figure 1. Methodology for investigating the relationship of LSTM architecture, training dynamics, and training loss curvature to model generalizability in this work.

designed an experiment that includes two screening stages (see Fig. 1). In the first stage, we train 100 LSTMs all with hyper-parameter tuning via Optuna [11] on the same fixed training data of 500 samples. Table 1 lists the hyper-parameters and their search ranges. To enforce more diversity, we deliberately narrowed the search into various regions of the search range for some LSTMs. For example, some LSTMs were forced to have ReLU input activations only, some were tuned up to 20 Optuna trials only rather than 50, some were stopped at 50 epochs rather than 200, some were forced to search within 30 to 40 hidden units rather than 5 to 45, and so on. Note that all the LSTMs in this study used *sigmoid* as the recurrent activation to maintain the notion that gates must output a value only from 0 to 1. All LSTMs also have some form regularization by using dropout and L^2 penalty.

For the second screening stage, we then picked the best LSTM from the first screening and kept its architecture. This stage now trains 100 additional LSTMs on the same training data of 500 samples, all of which have the same architecture (i.e. number of model parameters, activation, and regularization) but with different number of training epochs, training batch size, and initial weights.

The purpose of this stage is to remove the effects of model size and architecture choice and instead focus on sampling the training loss landscape diversely. Table 2 shows the range of number of epochs and batch sizes tried in this stage. Note that even though the batch size differs between LSTMs during training, we evaluated all curvature metrics using the entire batch of 500 samples. All generated data sets start with the same known initial state to remove any errors due to different initial conditions. The lookback is also fixed at 23, as with [25].

Table 1: Hyper-parameter search ranges in 1st screening.

Name	Search Range
No. of LSTM units	5, 6, 7, ..., 45
Input-output activation, σ_{io}	Tanh, ReLU, Linear
Dropout rate	0.0, 0.1, ..., 0.5
Recurrent dropout rate	0.0, 0.1, ..., 0.5
Optimizer	Adam, SGD, RMSProp, Adadelta
Batch Size	2^3, 2^4, ..., 2^8
Validation split	0.1 or None
L^2 regularization	0.001 to 1
No. of training epochs	50 to 200
No. of Optuna trials	20 to 50

Table 2: Ranges used for training LSTMs in 2nd screening.

Name	Search Range
Batch Size	2^4, 2^5, ..., 2^8
No. of training epochs	100, 200, 300

Curvature and Smoothness Metrics

We measure the following curvature and local state transition smoothness metrics of all trained LSTMs.

Top Hessian Eigenvalue, $\lambda_{\max}$

A measure of the local curvature of the loss in the parameter direction where it is sharpest.

Hessian Trace, $\mathrm{tr}(\boldsymbol{H})$

The sum of the eigenvalues of the Hessian.

Anisotropy Ratio, $\lambda_{\max}/\mathrm{tr}(\boldsymbol{H})$

Top Hessian eigenvalue divided by the Hessian trace, indicating the fraction of the total curvature concentrated in the sharpest direction.

Effective Rank of the Hessian, R_{eff}

Indicates the true dimensionality of the loss curvature, which we compute as the ratio of the squared Frobenius norm to the squared spectral norm:

$$R_{\mathrm{eff}} = (\Sigma\lambda_i)^2/\Sigma\lambda_i^2 \tag{7}$$

Recurrent Kernel Spectral Norm, $\sigma_{\max}(\boldsymbol{R})$

The largest singular value of the recurrent kernel matrix $[\boldsymbol{R}_z, \boldsymbol{R}_i, \boldsymbol{R}_f, \boldsymbol{R}_o]$ from (1), (2), (3), (5).

Our hypotheses are as follows. If the top Hessian eigenvalue of the loss landscape is low, then the loss is flat along the worst-case (sharpest) direction. The Hessian trace measures the mean overall curvature on all directions, so a low trace also indicates flatness. An anisotropy ratio near 0 means that the curvature is more *spread* out on all directions, and if it is near 1 then the minima is likely a narrow ridge that is flat except for a single sharp direction. Hence, if the anisotropy ratio is near 0 and λ_{max} is also low, then the minima is a flat and wide basin. Meanwhile, the R_{eff} measures the effective number of dominant directions in the curvature. Lower R_{eff} can also mean flatter minima since the *number* of dominant (sharp) directions is low while the rest of the directions are flat. A lower R_{eff} can also indicate that the model is more overparametrized, possibly admitting a simplicity bias [18]. By Occam's Razor, simpler models tend to generalize well.

We include the recurrent kernel spectral norm [23] in the analysis since we consider it to be a smoothness metric. The lower the maximum singular value of the recurrent weights matrix, the more insensitive the model is to large changes in the hidden vector $\boldsymbol{y}_k^*$, which is the data being fed back to the LSTM. This means that upon making multi-step-ahead predictions, smoother trajectories are expected.

For the case of different architectures, as in the first screening stage (see Fig. 1), the Hessian trace and the effective rank must be normalized by the number of model parameters to be comparable.

Approximating the Hessian Metrics

The time and memory cost of computing the full Hessian is quadratic with the model size, making it infeasible for large LSTMs. To address this, most literature use the Hessian-vector products (HVP) method which computes unbiased estimates of the product of the Hessian and an arbitrary vector $\boldsymbol{Hv}$ in linear time [26].

The full set of algorithms are not explained here, but we refer readers to the implementation in PyHessian [27]. In particular, the $\lambda_{\max}$ is calculated using the power iteration method, while the Hessian trace and effective rank use the Hutchinson method, all of which use HVP. We arbitrarily set 30 iterations in the power iteration method and 10 samples in the Hutchinson method.

Generalization Metrics

To measure generalization, 100 additional data sets each having 500 samples were generated from the simulation of the process that differ in the input signals and noise. Inputs in both training and testing sets are amplitude modulated pseudo-random binary signals, but the test data input amplitudes are 50% larger in range than in the training data to test for extrapolation ability. All test data sets are fixed in both screening stages in Fig. 1.

We calculate the mean squared error (MSE) and R^2 accuracies in the 1-step-ahead and 500-step-ahead (or infinite-step) prediction scenarios, as follows:

$$\mathrm{MSE} = \frac{1}{N}\Sigma\left(y_{\mathrm{true}} - y_{\mathrm{pred}}\right)^2 \tag{8}$$

$$\mathrm{R}^2 = 1 - \Sigma\left(y_{\mathrm{true}} - y_{\mathrm{pred}}\right)^2/\Sigma\left(y_{\mathrm{true}} - \mu_{\mathrm{true}}\right)^2 \tag{9}$$

which are averaged on all outputs, then averaged on all 100 test data sets for a single LSTM. An MSE near 0 and an R^2 near 1 means that a model generalizes well.

Three Tanks Case Study

The system is governed by same differential equations as those written in [25]. Similarly, this paper uses the same two inputs and three outputs. The exact training data set used in this work is shown in Fig. 2.

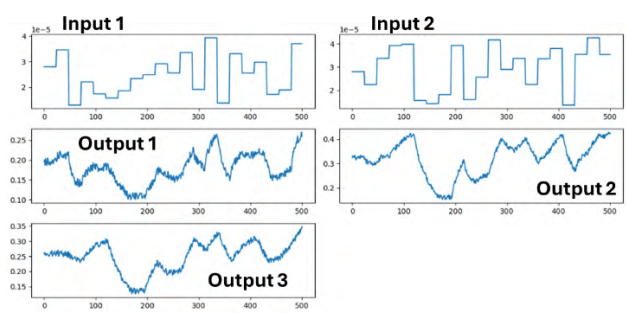

Figure 2. The training data set of 500 samples used to train all 200 LSTMs in this work.

RESULTS AND DISCUSSION

Correlations between Metrics

Quantifying the correlation between metrics already brings many insights. Fig. 3 shows the Spearman rank correlations between the choices during LSTM training,

curvature and smoothness metrics, and their generalizability for the 100 LSTMs with different architectures in the first screening stage. We chose rank-based correlation rather than Pearson because: (i) some accuracy metrics are too worse that they appear as outliers, hence, we analyzed more about the directions of the interaction rather than their magnitudes, and (ii) we wish to find possibly nonlinear interactions rather than linear ones.

First, we find that a lower top Hessian eigenvalue appears in LSTMs that generalize well, but the normalized trace, anisotropy, and normalized effective rank of the Hessian do not indicate the same. The rest of the Hessian metrics are weakly correlated with the mean MSE and mean R^2 for both 1-step-ahead and Inf-step-ahead accuracy. This means that for LSTMs with different architectures, the flatness at the direction of sharpest curvature dictates extrapolation behavior more effectively than the rest of the Hessian metrics. We further checked that the ranges of normalized $tr(\boldsymbol{H})$ (0-0.12) and normalized R_{eff} (0-0.03) are both close to 0. Hence, we attribute their weak correlation with accuracy to either: (i) possible saturation at already low values for trained LSTMs, or (ii) poor choice of normalizing constant, i.e. the number of model parameters. This needs further study.

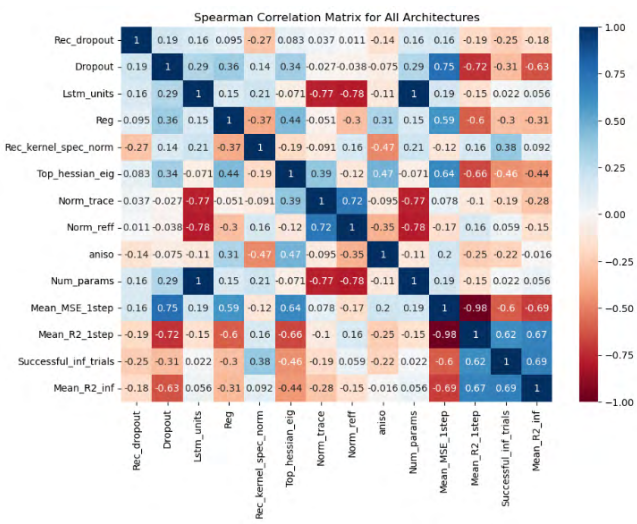

Figure 3. Spearman correlation matrix between metrics for LSTMs in the 1st screening (different architectures).

In terms of architecture, Fig. 3 shows that the number of model parameters is also weakly correlated with the test accuracy, meaning that there are extrapolative LSTMs from both small and large networks. We also find that the larger the network, the lower the normalized trace and normalized effective Hessian rank (correlation=-0.77). This means that bigger LSTMs tend to be more over-parametrized after training, i.e. with larger networks, fewer weights matter against the total number.

We chose a particular LSTM with the best accuracy for inf-step-ahead prediction, namely LSTM-91, in the first screening stage. LSTM-91 gained a mean R^2 of 82.7% and 85.67% for inf-step-ahead and 1-step-ahead prediction, respectively. Table 3 gives more details on this LSTM. Then, we trained 100 more LSTMs with the same hyper-parameters as LSTM-91, except with different initial weights, no. of epochs, and batch sizes.

Table 3: Details of the best LSTM from the 1st screening.

Hyper-parameter	Value
No. of LSTM units	31
No. of model parameters	4684
No. of training epochs	100 (Adam)
Input-output activation	Linear
Batch size	2^4
Recurrent dropout rate	0.1
Dropout rate	0.0
L^2 regularization	0.001126
No. of Optuna trials	50

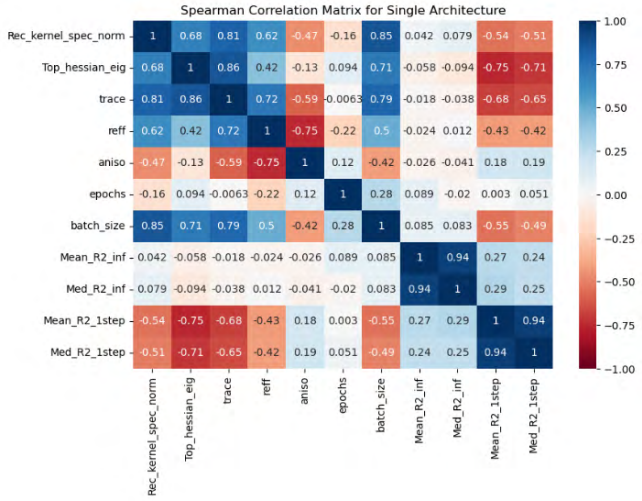

Figure 4. Spearman correlation matrix between metrics for LSTMs in the 2nd screening (same architecture).

Fig. 4 shows the Spearman rank correlation matrix for the 2nd screening stage. For the same architecture, we see stronger correlations between curvature, smoothness, and generalizability. Low values of $\sigma_{\max}(\boldsymbol{R}), \lambda_{\max}$, and $tr(\boldsymbol{H})$ all indicate better 1-step-ahead prediction accuracy (in terms of mean and median R^2). These metrics are also found to be correlated among each other: smoother model dynamics (low $\sigma_{\max}(\boldsymbol{R})$) correspond with flatter minima (low $\lambda_{\max}$). However, we find that anisotropy and effective rank still do not indicate generalizability. To understand why, we looked at the relationship between the metrics more closely in Fig. 5.

Studying 100 LSTMs with Same Architecture

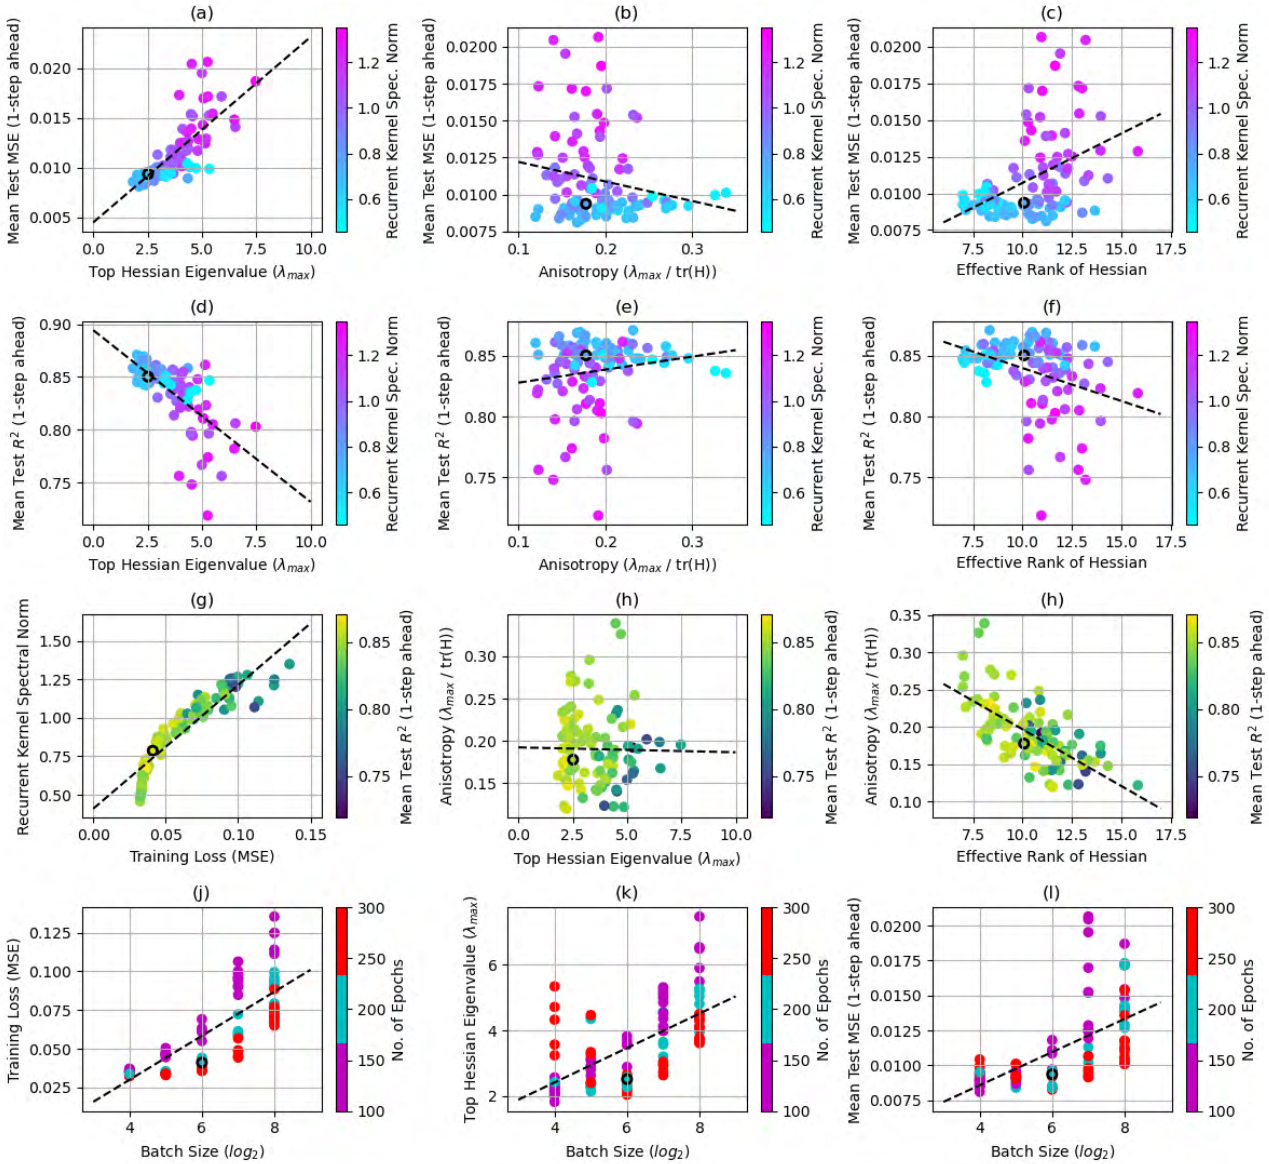

Figure 5. Second Screening Results: Scatter plots relating Hessian metrics, recurrent kernel spectral norm, training loss, batch size, and number of epochs to the 1-step-ahead accuracies averaged on 50 test data sets. Dashed line: Best-fit line; Black circle: LSTM with best inf-step-ahead accuracy.

In Fig. 5(a, d), we see that LSTMs with $\lambda_{\max} < 2.5$ and $\sigma_{\max}(\boldsymbol{R}) < 0.8$ are good extrapolators, and the linear correlation with test accuracy is clearer at lower values of $\lambda_{\max}$. Fig. 5(b, c, e, f) show that there are models that do not generalize well despite having lower anisotropy ratio and lower effective rank. Upon inspecting the trend of the recurrent kernel spectral norm in these plots, the reason for weak correlation of $tr(\boldsymbol{H})$ and R_{eff} with accuracy is the presence of low trace, low rank Hessian models with simultaneously high recurrent kernel spectral norm (pink).

Fig. 5(g) shows that as the model reaches lower minima in the training loss, the recurrent kernel spectral norm also tends to decrease, leading to more accurate models. Fig. 5(h) tells us that anisotropy is not correlated with $\lambda_{\max}$, meaning that the $\lambda_{\max}$ and Hessian trace can be independent in fully trained LSTMs. Next, Fig. 5(i) shows that anisotropy is anti-correlated with R_{eff}, which means that the trained LSTMs converge to either higher effective-dimensional but relatively uniform curvature minima or lower effective-dimensional but non-uniform curvature minima. Still, neither of these ends are indicative of generalizability for 1-step-ahead prediction.

Lastly, we tested some notions from literature about the relationship of batch size, flatness, and test accuracy.

According to [16], large batches are known to drive models to sharper minima and smaller batches lead to flatter minima. Fig. 5(j, k, l) shows that smaller batch sizes tend to converge to lower training loss, but not necessarily flat. We found 3 LSTMs, in particular, each trained with a batch size of 16 using 300 epochs that reached sharper minima ($\lambda_{\max} > 4$) but with roughly the same training loss (~0.03) as those on flat minima ($\lambda_{\max} < 2.5$). These 3 also had the lowest recurrent kernel spectral norms among all LSTMs ($\sigma_{max}(\boldsymbol{R}) < 0.5$, lightest blue dots in Fig. 5(a)). Consistent with [15], these are the sharp-minima LSTMs that still generalize well, but this time owing to low recurrent kernel spectral norms. On the other hand, models with large batch sizes indeed tend to converge to sharper minima, especially those trained with only 100 epochs.

Because we have 100 LSTMs with the same number of model parameters, we can investigate which models are similar in the weight space. We projected the 4, 684 weights of each LSTM in 2-D using t-distributed stochastic neighborhood embedding (t-SNE), a widely used data visualization method. Fig. 6 shows that there is a cluster of LSTMs with roughly equal training loss and similar weights (meaning that they may be "close" to each other in the loss landscape). The other more isolated LSTMs can each be thought of having very different weights distributions. All LSTMs that clustered were found to generalize well in terms of 1-step-ahead test accuracy, but not all of them have flat minima since the 3 sharp-minima LSTMs are also in this cluster. By finding several similar LSTMs clustered together (37 out of 100) despite starting with various initial weights, we believe that there likely exists a low-lying mostly flat basin (with a few sharp minima) in the loss landscape that is relatively easy to find via Adam when training the LSTM-91 architecture (see Table 3) given the fixed training data set.

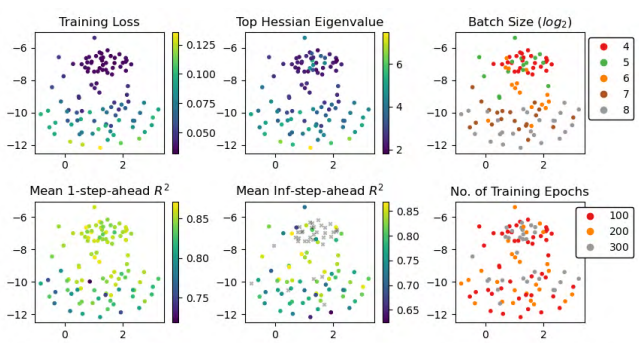

Figure 6. A 2-D projection of 4, 684 weights of the LSTMs in the 2nd screening stage via t-SNE, colored by the titles.

We also found that most of the clustered LSTMs did not perform well for multi-step-ahead prediction (see gray x's in the middle bottom box denoting models with negative mean R^2 for more than 10 out of 100 data sets). This means that steering LSTMs to be accurate 1-step-ahead predictors still does not guarantee stability.

CONCLUSION

This work aims to empirically relate the flatness of the local minima of the loss and smoothness of the model local dynamics to generalizability in LSTMs for the Three Tanks case study. We summarize our findings into the following, subject to the limitations of our experiments:

LSTMs that converged to flatter minima tend to generalize better, but mostly when flatness is measured as the top Hessian eigenvalue. This metric rank-correlates with the mean test R^2 (at -0.66) regardless of architecture. Better rank-correlation (-0.75) was found when analyzing the loss landscape of a single architecture.

Lower Hessian trace (overall loss curvature) also indicates better 1-step-ahead prediction accuracy (rank-correlation=-0.68), but only among LSTMs with the same architecture.

LSTMs with smooth local dynamics marked by a low recurrent kernel spectral norm also tend to generalize better given a fixed architecture.

Based on these results, it might then be worthwhile to steer the training of LSTM surrogates while being aware of both sharpness and local transition smoothness in terms of the metrics that we found useful in this work. Nonetheless, future work must also validate our results in other processes and with more diverse RNNs.

ACKNOWLEDGEMENTS

Authors are grateful for the support of the Philippine Department of Science and Technology, Engineering Research and Development for Technology (DOST-ERDT).

AUTHOR IDENTIFIERS

Author ORCIDs:
Roxas RM II: 0000-0002-1675-3926
Pilario KE: 0000-0001-5448-0909

REFERENCES

1. Pilario KES, Cao Y, Shafiee M. A kernel design approach to improve kernel subspace identification. IEEE Trans. Ind. Electron. 68:6171-6180 (2021). https://doi.org/10.1109/tie.2020.2996142
2. Wu Z, Christofides PD, Wu W, Wang Y, Abdullah F, Alnajdi A, Kadakia Y. A tutorial review of machine learning-based model predictive control methods. Reviews in Chemical Engineering 41:359-400 (2024). https://doi.org/10.1515/revce-2024-0055

3. Sun W, Paiva ARC, Xu P, Sundaram A, Braatz RD. Fault detection and identification using bayesian recurrent neural networks. Computers & Chemical Engineering 141:106991 (2020). https://doi.org/10.1016/j.compchemeng.2020.106991
4. Chen H, Cen J, Yang Z, Si W, Cheng H. Fault diagnosis of the dynamic chemical process based on the optimized CNN-LSTM network. ACS Omega 7:34389-34400 (2022). https://doi.org/10.1021/acsomega.2c04017
5. Zhu W, Chebeir J, Webb Z, Romagnoli J. A Deep Learning Approach on Surrogate Model Optimization of a Cryogenic NGL Recovery Unit Operation. *Computer Aided Chemical Engineering*, 48:1285–1290 (2020)
6. Esche E, Weigert J, Brand Rihm G, Göbel J, Repke JU. Architectures for neural networks as surrogates for dynamic systems in chemical engineering. Chemical Engineering Research and Design 177:184-199 (2022). https://doi.org/10.1016/j.cherd.2021.10.042
7. Alhajeri MS, Alnajdi A, Abdullah F, Christofides PD. On generalization error of neural network models and its application to predictive control of nonlinear processes. *Chem. Eng. Res. Des.* 189:664–679 (2023) https://doi.org.10.1016/j.cherd.2022.12.001
8. Cheng X, Huang K, Ma S. Generalization and risk bounds for recurrent neural networks. Neurocomputing 616:128825 (2025). https://doi.org/10.1016/j.neucom.2024.128825
9. Li J, Qin SJ. Applying and dissecting LSTM neural networks and regularized learning for dynamic inferential modeling. Computers & Chemical Engineering 175:108264 (2023). https://doi.org/10.1016/j.compchemeng.2023.108264
10. ?awry?czuk M. Input convex neural networks in nonlinear predictive control: a multi-model approach. Neurocomputing 513:273-293 (2022). https://doi.org/10.1016/j.neucom.2022.09.108
11. Pravin PS, Tan JZM, Yap KS, Wu Z. Hyperparameter optimization strategies for machine learning-based stochastic energy efficient scheduling in cyber-physical production systems. Digital Chemical Engineering 4:100047 (2022). https://doi.org/10.1016/j.dche.2022.100047
12. Shakouri B, Murthy S, De Clippeleir H, Lesnik K, Massoudieh A. Surrogate modeling of ASM1 using feedforward and LSTM neural networks. Journal of Water Process Engineering 81:109246 (2026). https://doi.org/10.1016/j.jwpe.2025.109246
13. Hochreiter S, Schmidhuber J. Flat minima. Neural Computation 9:1-42 (1997). https://doi.org/10.1162/neco.1997.9.1.1
14. Chaudhari P, Choromanska A, Soatto S, LeCun Y, Baldassi C, Borgs C, Chayes J, Sagun L, Zecchina R. Entropy-SGD: Biasing Gradient Descent Into Wide Valleys. (2017) http://arxiv.org/abs/1611.01838
15. Dinh L, Pascanu R, Bengio S, Bengio Y. Sharp Minima Can Generalize For Deep Nets. (2017) http://arxiv.org/abs/1703.04933
16. Keskar NS, Mudigere D, Nocedal J, Smelyanskiy M, Tang PTP. On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima. (2017) http://arxiv.org/abs/1609.04836
17. Tsuzuku Y, Sato I, Sugiyama M. Normalized Flat Minima: Exploring Scale Invariant Definition of Flat Minima for Neural Networks Using PAC-Bayesian Analysis. *Proceedings of the 37th International Conference on Machine Learning* 119 (2020)
18. Wilson AG. Deep Learning is Not So Mysterious or Different. (2025) http://arxiv.org/abs/2503.02113
19. Shoham N, Mor-Yosef L, Avron H. Flatness After All? (2025) http://arxiv.org/abs/2506.17809
20. Foret P, Kleiner A, Mobahi H, Neyshabur B. Sharpness-Aware Minimization for Efficiently Improving Generalization. (2021) http://arxiv.org/abs/2010.01412
21. Zhang X, Xu R, Yu H, Dong Y, Tian P, Cu P. Flatness-Aware Minimization for Domain Generalization. (2023) http://arxiv.org/abs/2307.11108
22. Lee S, He C, Avestimehr S. Achieving small-batch accuracy with large-batch scalability via hessian-aware learning rate adjustment. Neural Networks 158:1-14 (2023). https://doi.org/10.1016/j.neunet.2022.11.007
23. Yoshida Y, Miyato T. Spectral Norm Regularization for Improving the Generalizability of Deep Learning. (2017) http://arxiv.org/abs/1705.10941
24. Saab S Jr, Fu Y, Ray A, Hauser M. A dynamically stabilized recurrent neural network. Neural Process Lett 54:1195-1209 (2021). https://doi.org/10.1007/s11063-021-10676-7
25. Pilario KE, Wu Z. Fast mixed kernel canonical variate analysis for learning-based nonlinear model predictive control. Chemical Engineering Research and Design 219:19-33 (2025). https://doi.org/10.1016/j.cherd.2025.05.050
26. Pearlmutter BA. Fast exact multiplication by the hessian. Neural Computation 6:147-160 (1994). https://doi.org/10.1162/neco.1994.6.1.147
27. Yao Z, Gholami A, Keutzer K, Mahoney MW. Pyhessian: neural networks through the lens of the hessian. 2020 IEEE International Conference on Big Data (Big Data) :581-590 (2020). https://doi.org/10.1109/bigdata50022.2020.9378171

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Active-Constraint Regions and Power Distribution in Multi-Stack PEM Water Electrolysis Systems

Marius Fredriksen[a], Johannes Jäschke[a]*

[a] Norwegian University of Science and Technology, Department of Chemical Engineering, Trondheim, Norway
* Corresponding Author: johannes.jaschke@ntnu.no

ABSTRACT

Multi-stack proton exchange membrane (PEM) water electrolysis systems are increasingly deployed to improve the scalability and flexibility of green hydrogen production. However, sharing balance-of-plant equipment introduces coupling between stacks, and differences in stack performance increase the complexity of plantwide operation. In particular, non-identical efficiencies and safety constraints, such as hydrogen-to-oxygen (HTO) ratio limits, can render single-stack or equal-power-sharing control strategies suboptimal. In this work, the steady-state optimal operating regime of a two-stack PEM electrolysis system is characterized using a plantwide optimization approach and active constraint mapping over a range of system power loads. Performance differences between the stacks are represented through variations in Faraday efficiency to emulate simplified degradation. For identical stacks, the system behaves similarly to a single large electrolyzer, where equal power distribution is optimal, and the active constraint regions closely resemble those of a single-stack system. As the stack performance differences increase, the optimal power distribution becomes asymmetric, with the more efficient stack preferentially loaded. However, HTO safety constraints in the degraded stack may limit the utilization of the more efficient stack and introduce additional active constraint regions, resulting in more complex operating regimes.

Keywords: Hydrogen, Energy Management, Process Optimization, PEM Electrolysis, Active Constraint Regions

INTRODUCTION

Hydrogen is expected to play a key role in decarbonization, and the development of large-scale, safe, and efficient production technologies is therefore required to meet zero-emission targets. Proton exchange membrane (PEM) water electrolysis is particularly attractive due to its high efficiency, fast dynamics, and suitability for intermittent renewable power. To reach industrially relevant capacities, modern PEM systems employ multiple stacks operating in parallel, often connected through a shared balance of plant (BoP). Such multi-stack configurations improve scalability and cost-effectiveness at high power loads compared to single-stack systems [2]. However, parallel stacks significantly increase system complexity and make plantwide operation more challenging due to coupling effects.

Despite this, most studies on PEM electrolysis focus primarily on single-stack performance. While several works have investigated power distribution strategies for multi-fuel-cell systems, considerably less research has addressed multi-stack PEM water electrolysis [2]. Among the existing studies, Xinyu et al. [11] proposed a daisy-chain power-distribution strategy, Cheng et al. [2] developed a power-adaptive fuzzy PID approach to improve load flexibility, and Guo et al. [5] introduced a rule-based strategy for hybrid PEM–alkaline systems utilizing the hydrogen-to-oxygen (HTO) impurity ratio.

Furthermore, due to manufacturing tolerances, aging, and thermal differences, individual stacks may exhibit different efficiencies. As a result, equal power sharing can be suboptimal and, in the worst case, violate operational or safety constraints. These performance differences should therefore be considered when determining power allocation among stacks [11].

In this work, we determine the steady-state optimal operating regime of a two-stack PEM electrolyzer system by solving a plantwide optimization problem and mapping the resulting active constraint regions across a range of power inputs. This analysis reveals how power should be

distributed between two PEM stacks operating in parallel and how differences in Faraday efficiency affect the optimal power allocation. The goal is to quantify how differences in stack performance affect the system's operating conditions and to provide a foundation for future control development.

METHODOLOGY

In this section, we briefly describe the two-stack PEM electrolysis system, the underlying process model, and the optimization framework used to determine the optimal power allocation and corresponding active constraint regions.

Case study and system layout

The system consists of two 60 kW PEM water electrolyzer stacks operating in parallel. Each stack has its own heat exchanger for temperature regulation and a dedicated power converter. The stacks share common BoP equipment, including two gas–liquid separators, a deionized water buffer tank, four pumps, four valves, and a hydrogen compressor. The piping is neglected as its effect on the steady-state operating conditions is negligible. The overall system layout is shown in Figure 1.

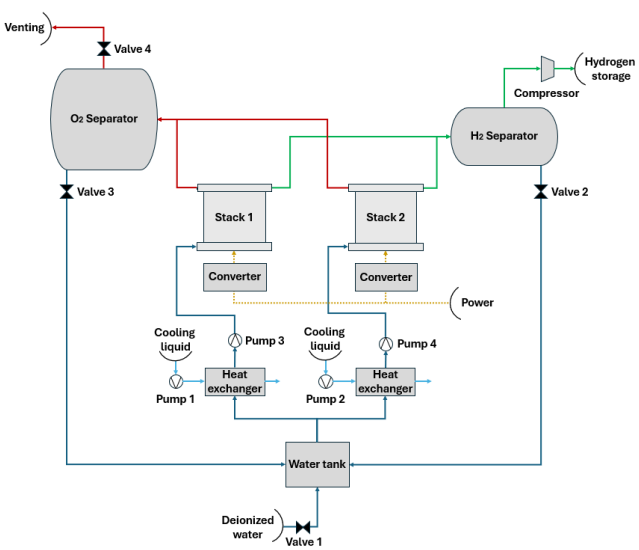

Figure 1. System layout of the two-stack PEM water electrolysis system with parallel stacks and shared balance-of-plant equipment.

Deionized water from the buffer tank is cooled in the heat exchangers and then supplied to the stacks, where it is electrolyzed to produce hydrogen and oxygen. The produced gases are subsequently separated from the liquid water in the gas–liquid separators, after which the oxygen is vented, and the hydrogen is compressed for downstream treatment. Finally, the remaining water is recycled back to the buffer tank.

Sharing separators introduces coupling between the stacks and additional operational constraints. In particular, both stacks must operate at similar pressures, limiting independent adjustment of membrane pressure gradients for gas crossover control. Furthermore, the HTO ratio measured at the separator does not directly reflect the HTO within each stack; therefore, stack-level HTO values must be estimated.

Process model description

The PEM electrolysis system model is adopted from our previous work [4], which combines electrochemical, thermal, and mass-transport stack submodels with standard mass and energy balances for the BoP equipment. For brevity, only the relations most relevant to the steady-state optimization are summarized in this work; for a more detailed model description, we refer to [4].

Electrochemical relations

The Stack performance is described using the electrochemical submodel that relates cell voltage (V_{cell}) to the current density (i_{el}) at a given power. The cell voltage is expressed as the sum of the reversible potential (V_{rev}), and the ohmic (V_{ohm}), activation (V_{act}), and concentration (V_{con}) overpotentials,

$$V_{cell} = V_{rev} + V_{ohm} + V_{act} + V_{con}, \tag{1}$$

where the reversible voltage represents the thermodynamic minimum energy required to drive the reaction, and the remaining terms account for resistive, kinetic, and mass-transport losses. The reversible potential is computed from the Gibbs free energy and corrected for pressure effects using the Nernst equation. Ohmic losses are estimated using an empirical membrane resistance correlation for Nafion 117, while activation and concentration losses are described using Tafel-type kinetics [1, 3, 8].

Faraday efficiency and degradation

Hydrogen production is calculated using Faraday's law of electrolysis. To represent performance differences between stacks, a Faraday efficiency factor (η_F) is introduced as a multiplicative term on the electrochemical gas generation rates,

$$\dot{n}_{\mathrm{gen},H_2} = \frac{i_{\mathrm{el}}A}{zF}\eta_F, \dot{n}_{\mathrm{gen},O_2} = \frac{i_{\mathrm{el}}A}{2zF}\eta_F. \tag{2}$$

A is the active membrane area, z is the number of electrons transferred, and F is the Faraday's constant [3]. Values $\eta_F < 1$ represent parasitic charge consumption (e.g., crossover-related recombination and other side currents), resulting in reduced hydrogen production per unit current. The associated electrical losses are assumed to be dissipated as heat. However, since this study focuses primarily on HTO constraints and steady-state power allocation, additional heat generation is neglected for simplicity.

HTO ratio and gas solubility

The HTO ratio is defined as the volumetric fraction of hydrogen relative to oxygen. Assuming ideal-gas behavior, this ratio can be estimated directly from the gas-phase composition,

$$HTO = \frac{V_{H_2}}{V_{O_2}} = \frac{y_{H_2}}{y_{O_2}}, \tag{3}$$

where V is the gas volume and y the gas-phase mole fraction.

The gas inventories in the separators are determined by subtracting dissolved gases from the total species inventories. The amount of dissolved gases is estimated using Henry's law with temperature-dependent solubility constants (H) obtained from Sander [9],

$$H = H_0 \cdot \exp\left(B\left(\frac{1}{T} - \frac{1}{T_0}\right)\right). \tag{4}$$

H_0 is the Henry constant at the reference temperature $T_0 = 298.15$ K, T is the stack/separator temperature, and B is an empirical temperature parameter. Gas solubility within the membrane is described using the correlations presented by Afshari et al. [1] and is assumed constant.

For simplicity, the stack-level HTO is computed using the total hydrogen and oxygen inventories. Due to oxygen's higher solubility in water, HTO values are slightly lower inside the stacks than in the separators, particularly at low pressure and temperature.

Compression and pump work

Compression and pumping contribute to the total system power consumption and therefore influence the optimal operating regime. Hydrogen leaving the cathode is compressed to 30 barg for downstream handling, while liquid water and coolant are circulated through the stacks and heat exchangers by pumps. Including these auxiliary loads allows the optimization to capture the trade-off between electrochemical stack efficiency and mechanical energy requirements for compression and circulation.

The compressor is modeled assuming multistage compression with intercooling, resulting in approximately isothermal behavior. The required compression work is estimated as,

$$W_c = \frac{\dot{n}_{H_2} R T_{\text{sep}}}{\eta_c} \ln\left(\frac{P_{\text{storage}}}{P_{\text{sep}}}\right), \tag{5}$$

where $\dot{n}_{H_2}$ is the hydrogen molar flow rate, T_{sep} is the separator temperature, P_{storage} and P_{sep} denote the storage and separator pressures, and η_c is the compressor efficiency (80%).

The pump work is approximated by,

$$W_p = \frac{\dot{m}}{\rho} \frac{(P_{\text{final}} - P_{\text{initial}})}{\eta_p}. \tag{6}$$

$\dot{m}$ is the mass flow rate, ρ is the liquid density, and η_p is the pump efficiency (75%).

Steady state optimization formulation

To determine the optimal operating policy, a steady-state optimization problem is solved over a range of system power inputs. The total electrical power is treated as an external disturbance representing the available power supply, while the internal power distribution between stacks is optimized. For each load level, the objective is to maximize the hydrogen output subject to the process model and operational constraints.

$$\begin{aligned} \max_{\text{MVs}} \quad & \dot{n}_{H_2} \\ s.t. \quad & \text{Process model equations} \\ & \text{Operational constraints} \end{aligned} \tag{7}$$

The manipulated variables (MVs) correspond to the available actuators in the system, namely valves, pumps, power converters, and the hydrogen compressor. With the liquid inventories in the separators and buffer tank regulated by local control loops and the feedwater flow rates fixed, five MVs remain for the optimizer. These are the oxygen outlet valve (z_4), the compressor work (W_c), the work of the two cooling liquid pumps (W_{p3}, W_{p4}), and the stack power distribution factor (α), which determines the fraction of total power supplied to each stack,

$$MVs = [z_4, W_c, W_{p3}, W_{p4}, \alpha]. \tag{8}$$

Operational limits are imposed on stack pressure, temperature, and the HTO ratio to ensure safe operation. Although the lower explosion limit of hydrogen in oxygen is approximately 4 vol.%, the HTO constraint is conservatively set to 2 vol.% to provide a safety margin. Upper bounds on stack pressure and temperature are adopted from Crespi et al. [3]. Additional constraints are imposed on minimum stack power (5% of nominal load) and cooling liquid flow rates to ensure physically feasible operation. The resulting bounds are summarized in Table 1.

Table 1: PEM system operating constraints.

Condition	Lower bound	Upper bound
Stack power	3 kW	60 kW
Anode pressure	0.2 barg	2.7 barg
Cathode pressure	0.2 barg	30 barg
Temperature	10 °C	60 °C
Cooling liquid	0 mol/s	15 mol/s
HTO ratio	0 vol.%	2 vol.%

The optimization problem is solved for 120 evenly spaced values of total power input, ranging from 6 to 120 kW. For each operating point, all constraints are scaled between 0 and 1 and used to identify the corresponding active constraint region. A constraint is considered active when its scaled value reaches either 0 for the lower limit or 1 for the upper limit. This procedure is repeated for multiple Faraday efficiency scenarios to evaluate the

impact of stack degradation and performance differences.

The steady-state optimization problem is implemented in Julia (v1.9.3) using JuMP (v1.29.0) [7]. The resulting problem is solved with the Ipopt (v1.11.0) interior-point solver [10], using the HSL MA57 linear solver [6].

RESULTS AND DISCUSSION

To evaluate the impact of stack performance differences on optimal operation, active constraint regions are analyzed for three efficiency scenarios. In Scenario 1, both stacks operate with identical Faraday efficiencies of 100%. In Scenarios 2–3, the efficiency of Stack 2 is reduced to 90% and 80%, respectively, while Stack 1 remains at nominal performance. Although these reductions are more severe for the Faraday efficiency than typically observed in practice, they are intentionally introduced to clearly illustrate how stack mismatch influences the optimal power distribution and operating regime. The considered scenarios and the corresponding number of active constraint regions are summarized in Table 2.

Table 2: Faraday efficiency scenarios for Stack 2 and the corresponding number of active constraint regions. The efficiency of Stack 1 is fixed at 100%.

Scenario	Stack 2 Faraday efficiency [%]	Active constraint regions [#]
1	100	2
2	90	5
3	80	6

Scenario 1 – Identical stacks

The optimal operating regime, active constraint regions, and corresponding power distribution for Scenario 1 are shown in Figure 2.

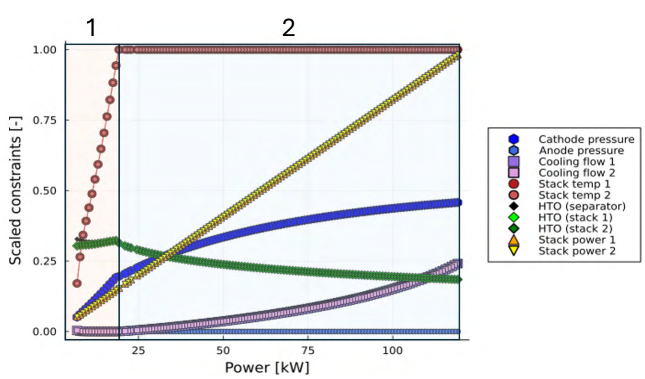

Figure 2. Optimal operating regime and active constraint regions for Scenario 1.

When both stacks have identical efficiencies, the optimizer distributes the power evenly between them. This is expected, since operating each stack at a lower individual power load improves hydrogen production per unit of power. Consequently, the two-stack system behaves similarly to a single larger stack.

At low loads, the stack temperatures increase rapidly with power until the upper temperature constraint becomes active at approximately 20 kW. Because higher temperatures improve electrochemical efficiency, the optimizer keeps the cooling flow rates at their lower bounds until the maximum temperatures are reached.

The anode pressure remains at its lower bound across the entire operating range due to the small efficiency penalty associated with elevated pressures. In contrast, the cathode pressure gradually increases with load. This behavior arises because the optimizer utilizes the stack pressure to partially compress the produced hydrogen, thereby reducing the work required by the downstream compressor. Although in-stack compression is generally more energetically favorable than mechanical compression, the compressor is modeled as a highly efficient isothermal unit. As a result, the optimizer prefers using both the compressor and the stack for compression. However, a less efficient compressor would likely result in higher optimal cathode pressures.

The HTO constraint remains inactive throughout this scenario. This is partly due to the low pressure gradients at small loads and partly due to the equal power distribution between the stacks, which results in similar crossover behavior. At very low power levels, the separator HTO is slightly higher than the stack HTO. This difference arises because the stack HTO is computed from the total gas inventories, whereas the separator measurement excludes gases dissolved in the liquid phase.

As both stacks operate under identical conditions, the active constraint regions closely resemble those of a single-stack system [4]. One of the main differences is that the single-stack system exhibits a slightly higher anode pressure at high loads. In the multi-stack configuration, however, increasing the anode pressure requires additional pump work due to the extra feedwater pump, making such operation less favorable.

Consequently, two active constraint regions are identified for the multi-stack system. In Region 1, the temperature constraints are inactive, and the cooling liquid flow rates remain at their lower bounds. As the load increases and the upper temperature constraints become active, the system transitions to Region 2, where additional cooling is required to maintain safe operation. Both regions have two degrees of freedom (DoFs), defined as the number of MVs minus the number of active constraints.

Scenario 2 – Stack 2 at 90 %

In Scenario 2, the Faraday efficiency of Stack 2 is reduced to 90%, introducing asymmetry between the two stacks. The optimal operating regime, active constraint

regions, and power distribution are shown in Figure 3.

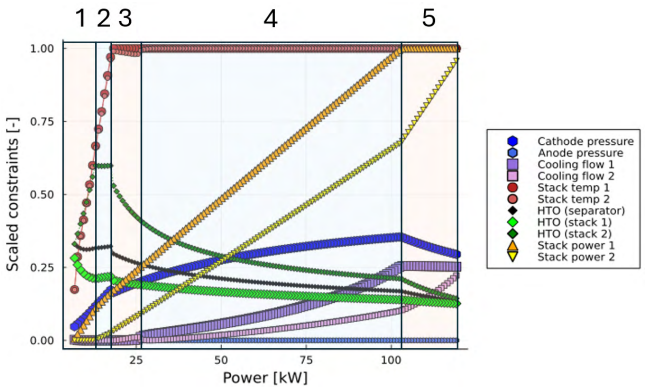

Figure 3. Optimal operating regime and active constraint regions for Scenario 2.

Stack 1 is more efficient than Stack 2, therefore the optimizer allocates a larger share of the power load to this electrolyzer. This behavior is expected and consistent with prior work recommending preferential use of the most efficient stack [11].

The HTO ratio of Stack 2 becomes significantly higher than that of Stack 1 at low loads. This occurs because reduced Faraday efficiency lowers oxygen production, thereby reducing hydrogen crossover dilution. At low loads, the HTO in Stack 2 increases rapidly until additional power is allocated to the stack, demonstrating a clear coupling between stack loading and the HTO ratio.

The temperature behavior largely follows Scenario 1. However, Stack 2 operates slightly below its upper temperature limit when Stack 1 reaches its constraint. Since Stack 2 contributes less to the overall hydrogen production in this region, operating at a lower temperature incurs a relatively minor efficiency penalty while slightly reducing downstream compression work. The optimizer also concentrates cooling on Stack 2, effectively allowing one cooling circuit to assist both stacks by overcooling the less efficient electrolyzer.

The cathode pressure remains inactive throughout the operating range and follows the same trend as in Scenario 1, but at lower absolute values. The maximum cathode pressure decreases from 14.84 bar in Scenario 1 to 11.75 bar, while the anode pressure remains fixed at its lower bound.

Scenario 2 has five active constraint regions. Regions 1–2 resemble Region 1 from Scenario 1, where the temperature constraints are inactive. However, in Region 1, Stack 2 operates at minimum power, while in Region 2, its load gradually starts increasing. In Region 3, the temperature constraint on Stack 1 becomes active, and Stack 2 is slightly overcooled. In Region 4, the temperature constraint on Stack 2 also becomes active. Finally, in Region 5, Stack 1 reaches its maximum allowable power, limiting further redistribution. Regions 2–4 have two DoFs, while Regions 1 and 5 have one DoF.

Scenario 3- Stack 2 at 80 %

In Scenario 3, the Faraday efficiency of Stack 2 is reduced to 80%, further increasing the performance mismatch between the two stacks. The resulting optimal operating regime, active constraint regions, and power distribution are shown in Figure 4.

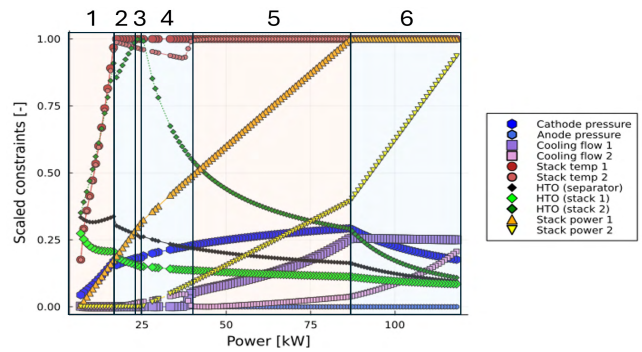

Figure 4. Optimal operating regime and active constraint regions for Scenario 3.

The overall behavior follows the same trends observed in Scenario 2, but with more pronounced asymmetry. The optimizer allocates an even larger fraction of the total load to Stack 1, causing Stack 2 to operate at its minimum power over a wider range of low system loads, while Stack 1 remains at its upper power limit for a larger portion of the high-load region.

As a result of this extended low-power operation, the HTO ratio in Stack 2 increases and reaches its upper constraint, which becomes active over a small portion of the operating range. When the HTO limit is reached, the cathode pressure stagnates slightly before the power load of Stack 2 begins to increase. This behavior indicates a clear relationship between the stack HTO, the pressure gradient, and the power load, where either the stack power must be increased or the cathode–anode pressure gradient reduced to maintain the HTO ratio of the degraded stack within safe limits.

Consistent with the reduced utilization of Stack 2, its operating temperature decreases further compared to Scenario 2. The cathode pressure continues to decrease, with the maximum dropping to 9.92 bar, while the anode pressure remains at its lower bound.

The increased stack asymmetry adds another active constraint region as the HTO limit of Stack 2 becomes active, yielding six active constraint regions in Scenario 3. In Region 1, Stack 2 operates at minimum power, with the anode pressure and both coolant flow rates at their lower bounds. In Region 2, the temperature constraint of Stack 1 becomes active, and the coolant flow of Stack 2 leaves its bound. In Region 3, the HTO constraint for Stack 2 becomes active, and it subsequently relaxes in

Region 4 as the lower power limit for Stack 2 becomes inactive. In Region 5, both temperature constraints are active, and in Region 6, Stack 1 reaches its maximum power limit. Regions 1, 2, 3, and 6 have one DoF, while Regions 4 and 5 have two.

CONCLUSION AND FUTURE WORK

This work investigates how differences in stack performance affect the optimal operating regime and active constraint regions of a two-stack PEM water electrolysis system. For identical stacks, the plant behaves similarly to a large single-stack system, and equal power allocation remains optimal across the entire operating region.

As performance differences are introduced, the optimal allocation becomes asymmetric, with the more efficient stack preferentially loaded. However, this utilization is restricted by safety constraints, particularly the HTO limit in the degraded stack. To maintain safe operation, the system must balance power allocation and pressure levels. This coupling between HTO, pressure, and power leads to additional active constraint regions, fragmenting the operating space and increasing the complexity of the optimal operating policy. These findings highlight the importance of accounting for stack asymmetry when designing control strategies for multi-stack systems.

Future work will use the identified operating regimes and active constraint regions to develop appropriate control structures for multi-stack PEM electrolysis systems using dynamic plant models. Additionally, the impact of extending the framework to systems with more stacks will be investigated.

ACKNOWLEDGEMENTS

This work was supported by the HYDROGENi Research Centre (hydrogeni.no) under the Norwegian research program FMETEKN. The authors acknowledge the industry partners in HYDROGENi for their contributions and the Research Council of Norway (project no. 333118).

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work, the authors used generative AI tools (ChatGPT and GitHub Copilot) to assist with phrasing, language improvement, and to generate code suggestions. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

REFERENCES

1. Afshari E, Khodabakhsh S, Jahantigh N, Toghyani S. Performance assessment of gas crossover phenomenon and water transport mechanism in high pressure PEM electrolyzer. International Journal of Hydrogen Energy 46:11029-11040 (2021). https://doi.org/10.1016/j.ijhydene.2020.10.180
2. Cheng K, He S, Hu B. Power adaptive control strategy for multi-stack PEM photovoltaic hydrogen systems considering electrolysis unit efficiency and hydrogen production rate. Sustainable Energy Technologies and Assessments 75:104200 (2025). https://doi.org/10.1016/j.seta.2025.104200
3. Crespi E, Guandalini G, Mastropasqua L, Campanari S, Brouwer J. Experimental and theoretical evaluation of a 60 kw PEM electrolysis system for flexible dynamic operation. Energy Conversion and Management 277:116622 (2023). https://doi.org/10.1016/j.enconman.2022.116622
4. Fredriksen M, J?schke J. Advanced regulatory control structure for proton exchange membrane water electrolysis systems. Systems and Control Transactions 4:1011-1016 (2025). https://doi.org/10.69997/sct.130330
5. Guo Y, Qi P, Zhang Q, Li M, Liu J, Sun H. Control strategy for hydrogen production system using hto-based hybrid electrolyzers. Energy Reports 13:2354-2364 (2025). https://doi.org/10.1016/j.egyr.2025.01.012
6. HSL. A collection of Fortran codes for large scale scientific computation. http://www.hsl.rl.ac.uk/
7. Lubin M, Dowson O, Garcia JD, Huchette J, Legat B, Vielma JP. Jump 1.0: recent improvements to a modeling language for mathematical optimization. Math. Prog. Comp. 15:581-589 (2023). https://doi.org/10.1007/s12532-023-00239-3
8. Ojong ET, Kwan JTH, Nouri-Khorasani A, Bonakdarpour A, Wilkinson DP, Smolinka T. Development of an experimentally validated semi-empirical fully-coupled performance model of a PEM electrolysis cell with a 3-D structured porous transport layer. International Journal of Hydrogen Energy 42:25831-25847 (2017). https://doi.org/10.1016/j.ijhydene.2017.08.183
9. Sander R. Compilation of henry's law constants (version 5.0.0) for water as solvent. Atmos. Chem. Phys. 23:10901-12440 (2023). https://doi.org/10.5194/acp-23-10901-2023
10. Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math. Program. 106:25-57 (2005). https://doi.org/10.1007/s10107-004-0559-y
11. Xinyu L, Banghua D, Bo Z, Lqilei Z, & Changjun X. Design and control of wind-hydrogen coupled

system based on chain distribution strategy [J]. Journal of Solar Energy, 43(06), 05-413. (2022)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Exploiting the line pack potential of gaseous CO_2 pipelines

Archana Kumaraswamy[a], Johannes Jäschke*[a]

[a] Norwegian University of Science and Technology (NTNU), Department of Chemical Engineering, Trondheim, Norway
* Corresponding Author: johannes.jaschke@ntnu.no.

ABSTRACT

Carbon dioxide transport is a critical component of the carbon capture and sequestration (CCS) supply chain. Given the substantial energy requirements and dispersed locations of CCS facilities, optimizing pipeline operations is critical to minimize costs. Although CO_2 in dense phase is typically favored for long-distance transport, gaseous phase transport is also a possibility for shorter distances and volumes. This study models a gaseous CO_2 pipeline system. Since CO_2 gas pipelines provide the benefit of line packing, owing to gas compressibility, this work leverages it to maximize throughput in the presence of disturbances. Pipeline pressures within each segment are perceived as an inventory (i.e. form of storage) and a model predictive control (MPC) formulation for optimal inventory management is implemented to maximize throughput. This study applies the formulation to pipelines arranged in series and parallel. It effectively maximizes throughput and optimally drains pipeline pressures in the presence of an inlet flow rate disturbance.

Keywords: Process Control, Optimization, Nonlinear Model Predictive Control, Carbon Dioxide Gas Pipelines

INTRODUCTION

CO_2 transport via pipelines is an integral component of carbon capture and sequestration (CCS) due to the spatial distribution of capture and sequestration sites. Additionally, the urgency for decarbonization necessitates minimizing the cost and energy intensity of CO_2 transport especially since it accounts for approximately 25% of total CCS costs [1]. One approach to minimize operational costs is to exploit the line pack potential of gaseous CO_2 pipelines [2] to maximize throughput in the presence of shifting bottlenecks. Bottleneck here refers to a process constraint preventing further increase in throughput. In essence, this study investigates how pipeline pressures can be used as buffers to delay the impact of disturbances on throughput. Thus, a model predictive approach for inventory (pipeline pressure) control based on Turan, et al. [3] and Kumaraswamy, et al. [4] is implemented in this study to maximize throughput in gaseous CO_2 pipelines.

Although CO_2 can be transported in various other forms (e.g. ships, trucks, and rail), pipeline transport is one of the most mature and cost-effective technologies for large-scale transport [1, 5]. CO_2 is typically transported in liquid or dense phase at high-pressure over long distances as this is cost-effective [6]. Lower pressures result in lower densities and higher pressure drops. However, gaseous CO_2 pipelines are a possibility for shorter distances with operating pressures below 48 bar [7]. For example, OGE is developing a 28 km pipeline in Germany from Lägerdorf to Brunsbüttel, operating at a design pressure of 40 bar [8]. With the development of industrial clusters for gathering captured CO_2, short distance and low volume low-pressure gaseous pipelines might be preferable for certain pipelines. Several industrial clusters for gathering captured CO_2 are in development including industrial sources in the Oslo-fjord region [9], the Coda terminal [10], and the industry in Rotterdam [11]. Such industrial clusters are expected to experience disturbances in CO_2 capture flow rates and injection flow rates, owing to the scale and interconnectedness of the system [12]. The proposed control approach in this study is therefore well suited for throughput maximization within industrial clusters, where variability in capture and injection rates must be effectively managed.

Within the field of CO_2 transport and injection research, there has been great emphasis on steady-state optimal supply-chain design. Since several CCS projects have been staged for construction and implementation in the upcoming years, optimal design of transportation networks to minimize costs is a critical challenge. Studies have attempted to address this challenge considering

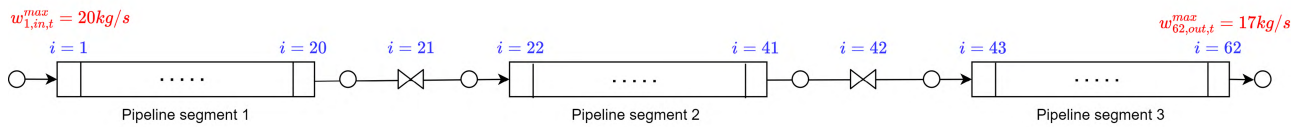

Figure 1. Sequential arrangement of pipelines (Case A)

different geographical scales, transport modes, and model complexity. Zhou, et al. [13] provide a detailed review of superstructure-based models for CCS pipeline network design. Ismail, et al. [14] review several studies centered on injection optimization and well control strategies incorporating detailed reservoir modelling.

Although there have been studies carrying out dynamic simulations to investigate the hydraulic and thermodynamic behavior of CO_2 in pipelines [15, 16], few studies have focused on methods for control and operational optimization for CO_2 transport and injection [17, 18]. This is the focus of the present paper.

Pressures within CO_2 pipelines can vary significantly depending on throughput, temperature, and transportation distance. Martynov, et al. [2] developed a correlation estimating line packing times for CO_2 pipelines under various pressure and distance conditions. The line packing time can vary between 30 mins to 2 hours for pipeline inlet pressures between 20 to 120 bar. The pressure threshold within the pipeline was set to 10 bar above the inlet pipeline pressure. This emphasizes the operational flexibility that CO_2 pipelines can offer in the presence of bottlenecks. It underscores the significance for investigating control strategies to optimally operate low-pressure gaseous CO_2 pipelines. While control strategies for optimal operation of gaseous CO_2 pipelines remains largely unexplored, it has been studied for natural gas pipeline networks [19–21].

Given the scale and interconnectedness of natural gas pipeline networks, which is also expected for CO_2 pipeline networks, studies have investigated the application of model predictive control (MPC) strategies to optimally operate such networks. Baumrucker, et al. [19] developed a non-linear dynamic model for mass flow and pressure in a natural gas pipeline network. An economic MPC was implemented to optimize compressor operation in the presence of varying electricity prices. However, to our knowledge, no study has adopted an optimal inventory control approach to the gas pipeline network process. In this study, pipeline pressures are treated as inventory and an inventory control formulation [3] is applied to maximize throughput from the gas pipeline while respecting pressure constraints.

Inventory control has been explored previously in the literature [22]. The term inventory is defined as any unit with hold-up (e.g. pressure, level). Some work has focused on inventory control strategies for bottleneck isolation. This concept can be simply explained as follows. Consider a pipeline transporting CO_2 from a capture site to a terminal for further transportation to a sequestration site. Additionally, consider that the inlet and outlet flow rates for this pipeline are determined by an optimizer. The optimizer is built on the principle that maximizing CO_2 transport to the terminal equates profit maximization. If a CO_2 capture plant were to shut down temporarily, the optimal inventory management strategy would be to maximize the pipeline outlet flow rate until pressures in the pipeline are at their lower limit. Subsequently, outlet flow rates would be lowered. This approach to isolate bottlenecks in a process has been described in the literature and applied to liquid tank levels. A PID control strategy, termed as the bidirectional control strategy, was formulated in Zotică, et al. [23] for sequential liquid tanks. An MPC approach achieving similar objectives was proposed by Turan, et al. [3] and Kumaraswamy, et al. [4].

In this work, an MPC inventory control approach is selected owing to the non-linearities, scale, and multi-input-multi-output nature of a CO_2 gas pipeline network. The inventory control formulation approach from Turan, et al. [3] is applied to two arrangements of CO_2 gas pipelines – sequential and parallel. We seek to investigate the effectiveness of this control strategy to optimally operate a CO_2 gas pipeline network.

MODELLING

The non-linear dynamic model in this paper is taken and adapted from a natural gas pipeline network model in [19]. We follow their exposition closely. Consider a pipeline network consisting of pipeline arcs, P, and choke valves, V. These elements are modelled as a system of arcs connected by nodes. Each arc in the network is indexed by $\forall i \in A$, with its endpoints identified by $\forall k \in \{in, out\}$, and nodes indexed by $\forall j \in J$. For the purposes of this paper, we define the term "pipeline segment" as a control volume formed by aggregating multiple pipeline arcs connected in series. An overview of the two pipeline arrangements investigated in this study are shown in Figures 1 and 2, where the distinction between pipeline arcs and segments are illustrated in these figures. The time horizon is partitioned into several discrete intervals, indexed by $t \in T$, each corresponding to one time step. Material and momentum conservation equations are imposed at arc endpoints. We also assume that each arc

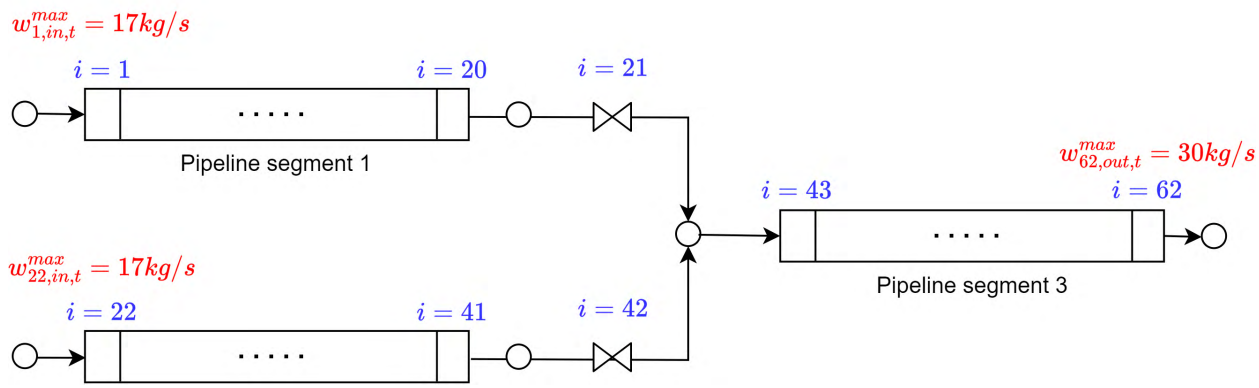

Figure 2. Parallel arrangement of pipelines (Case B)

endpoint shares the same pressure as its connected node. Additionally, no mass accumulation occurs at the nodes.

The governing equations are presented below. The material balance equation (eq. 1) is written in integrated form and approximated using the trapezoidal rule.

$$V_i\left(\overline{\rho}_{i,t+1} - \overline{\rho}_{i,t}\right) = \frac{1}{2}[(w_{i,in,t} - w_{i,out,t}) + (w_{i,in,t+1} - w_{i,out,t+1})]\Delta t,$$

$$\forall i \in \mathrm{P},\ t = 1, \dots, T-1. \quad (1)$$

The material balance equation determines the change in average pipeline arc density, $\overline{\rho}_{i,t}$ $(\frac{kg}{m^3})$, within a pipeline arc of volume, V_i (m^3). This is dependent on the flows entering, $w_{i,in,t}$ $(\frac{kg}{s})$, and leaving, $w_{i,out,t}$ $(\frac{kg}{s})$, the pipeline arc. Pressure drop due to friction is incorporated by implementing the one-dimensional plug flow equation (eq. 2) which calculates the pressure differential, $\left(\frac{dp}{dz}\right)_{i,k,t}$ $(\frac{bar}{m})$, at pipeline arc endpoints.

$$\left(\frac{dp}{dz}\right)_{i,k,t} = -\frac{f_{i,k,t}\, w^2_{i,k,t}}{2D_i A_i^2 \rho_{i,k,t}\, 10^5},$$

$$\forall \mathrm{i} \in \mathrm{P}, \mathrm{t} = 2, \dots, \mathrm{T}. \quad (2)$$

Note that only horizontal pressure variations along the pipeline (z) are considered and gravitational pressure losses are neglected under the assumption of perfectly horizontal pipelines. The frictional pressure drop is governed by flow-dependent parameters - specifically the friction factor, $f_{i,k,t}$, flow rate, $w_{i,k,t}$, and density, $\rho_{i,k,t}$- as well as pipeline characteristics, namely the cross-sectional area, A_i (m^2), and diameter, D_i (m). Each pipeline arc is assumed to have a diameter of 0.4m. A Reynolds-dependent friction factor is computed based on the Haaland correlation (eq. 3). Absolute roughness, ϵ (m), is assumed to be $2 \times 10^{-5}\ m$. Since CO_2 exhibits non-ideal gas behavior over the pressure range considered in this study, we have modified the dynamic model by Baumrucker, et al. [19] to incorporate a real gas formulation. Fluid properties – density and viscosity – for each pipe segment is computed using a real gas formulation, where the compressibility factor is represented by a second-order polynomial fitted from regression. The polynomial fit is valid over a range of 5 to 45 bar, corresponding to expected operating ranges of gaseous CO_2 pipelines.

$$\frac{1}{\sqrt{f_{i,k,t}}} = -1.8 \log_{10}\left[\left(\frac{\epsilon/D_i}{3.7}\right)^{1.11} + \frac{6.9}{Re_{i,k,t}}\right]$$

$$\forall \mathrm{i} \in \mathrm{P}, \mathrm{t} = 2, \dots, \mathrm{T}. \quad (3)$$

Pressure changes along pipeline arcs with length, L_i, of 1 km, are described using a 2-point Lobatto collocation scheme. This results in expressions for the outlet pressure and average pressure of each arc (eq. 4 and 5 respectively). A detailed derivation can be found in Baumrucker, et al. [19].

$$p_{i,\,\mathrm{out},t} = p_{i,\,\mathrm{in},t} + \frac{L_i}{2}\left[\left(\frac{dp}{dz}\right)_{i,\,\mathrm{in},t} + \left(\frac{dp}{dz}\right)_{i,\,\mathrm{out},t}\right],$$

$$\forall \mathrm{i} \in \mathrm{P}, \mathrm{t} = 2, \dots, \mathrm{T}. \quad (4)$$

$$\overline{p}_{i,t} = \frac{p_{i,\,\mathrm{in},t} + p_{i,\,\mathrm{out},t}}{2} + \frac{L_i}{12}\left[\left(\frac{dp}{dz}\right)_{i,\,\mathrm{in},t} - \left(\frac{dp}{dz}\right)_{i,\,\mathrm{out},t}\right],$$

$$\forall i \in \mathrm{P}, t = 2, \dots, T. \quad (5)$$

MPC CONTROL FORMULATION

Having presented the dynamic model, we now define the control problem. In this study, we apply an inventory control MPC formulation to maximize system outflow and weighted inventories [3]. We seek to maximize the outlet flow rate from a pipeline system and the weighted pressure within each pipeline segment,

$$\min J = \sum_{t=2}^{T}\left\{\underbrace{\sum_{i\in\mathcal{O}} \gamma^t \alpha_i \left(w_{i,out,t} - w^{max}_{i,out,t}\right)^2}_{Outflow\ max.} - \right.$$

$$\left.\underbrace{\sum_{i\in S}\gamma^t\beta_i p_{i,in,t}}_{Inventory\ max.}\right\} + \underbrace{R}_{penalty\ term} \quad (6)$$

In accordance with the weight selection guidelines in Turan, et al. [3], inventories closest to the system outflow are assigned the highest weights. This prioritization emphasizes a high pressure in the pipeline segment connected to the outlet. Consequently, bottleneck isolation can be achieved by sequentially lowering pipeline segment pressures, starting from the segment located furthest from the outlet. While it may initially appear counterintuitive – as flow is typically driven by a pressure gradient from the inlet to outlet – the pressure constraint limits are defined based on expected operating pressures. For instance, in Case A, the maximum pressure limit is highest in segment 1, followed by segments 2 and 3 respectively. By prioritizing inventory weights for segments near the outlet, these critical sections—essential for maintaining flow—are preferentially driven toward their maximum pressure limits.

For the purposes of this paper, pipeline pressures at the midpoint of each segment are chosen for maximization and control. Let us define S as the set of arc indices representing the midpoints of each pipeline segment and O as the set of arc indices representing system outflows. α_i and β_i represent the weights for system outflow and inventory (i.e. pressure) level maximisation. Throughput maximisation should be preferred over inventory maximisation to ensure inventories are drained in the presence of an upstream bottleneck. Hence, α_i is always assigned a larger value than β_i. The term, γ^t, is a time discount factor which prioritizes minimizing the objective function at earlier time steps.

A fundamental difference between liquid tank inventories, studied earlier and pipeline networks in the present work, is the requirement for boundary conditions. While liquid tank systems can be modelled using only flow rates, pipeline networks require at least one specified pressure boundary condition [24]. Hence, the outlet pressure is fixed in all simulations in this paper. Consequently, the outlet flow rate is determined by the system dynamics and MV movements. Hence, system outflow maximization in this study is formulated as a reference-tracking objective, unlike the liquid tank network where the system outflow is modelled as a manipulated variable (MV). As shown in Figures 1 and 2, the maximum outlet flow rate, $w_{i,out,t}^{max}$, is the setpoint for system outflow reference tracking.

The MVs in this system are the inlet flow rate(s) and valve openings between pipeline segments. These variables are assigned lower and upper bounds. Valves are assumed to have lower and upper bounds of 0.05 and 1 respectively. Lower bounds for the inlet flow rates are 5 $\frac{kg}{s}$ while upper bounds are shown on Figures 1 and 2.

The penalty term comprises of soft constraints on pipeline segment pressures and a damping term on network flow rates, with $\lambda_1$and λ_2 denoting the corresponding penalty factors.

$$R = \lambda_1 \sum_{t=2}^{T}\sum_{i\in S}\left(p_{i,in,t}^{slack} + p_{i,in,t}^{surplus}\right) + \lambda_2 \sum_{t=2}^{T}\sum_{i=1}^{I}\gamma^t\left(w_{i,in,t} - w_{i,in,t-1}\right)^2 \quad (7)$$

Slack and surplus variables corresponding to violations of the minimum and maximum pressure limits are denoted as $p_{i,in,t}^{slack}$ and $p_{i,in,t}^{surplus}$. At the initial steady-state starting point, all pipeline segment midpoint pressures are at their maximum limit. The pressure range for each segment is defined as 4 bar. In addition, a damping term is introduced to penalize large fluctuations in pipeline flow rates across all arcs. This term promotes smoother MV movement and is included specifically in the pipeline pressure system. Unlike the liquid tank system explored in Turan, et al. [3], gas pipeline networks are sensitive to large and rapid changes in the MVs, which can cause pressure oscillations and instability. Thus, regulated MV movement is necessary. Moreover, inventory levels in gas pipeline networks are inherently coupled. While the liquid level in a tank depends only on its inlet and outlet flow rates, pressure variations in one pipeline segment affect flow rates and pressures across the entire network. Hence, the inclusion of a term to smooth MV movement is desirable.

RESULTS

The control strategy described above was implemented for two pipeline arrangements – sequential (Case A, Figure 1) and parallel with merger (Case B, Figure 2). In this study, the MPC receives pipeline pressure measurements at intervals of 1 km along the pipeline length. Simulations were carried out with a disturbance in the maximum allowable inlet flow rate to the pipeline at $t = 1h$. In this system, maximum flow rates for the inlet are assumed to be higher than the outlet to demonstrate both pipeline depressurization and pressurization as the bottleneck shifts between the outlet and inlet of the pipeline. Simulations begin at steady-state conditions with all pipeline segment midpoint pressures and the outlet flow rate maximized. Steady-state conditions are obtained by solving the steady-state form of the dynamic model presented above using a non-linear solver, Ipopt. The total simulated time is 22 hours with a time step of 5 minutes for all cases. All computations were performed on a 64-bit system with an Intel Core Ultra 5 135H processor (3.60 GHz) and 16 GB RAM. Each run requires approximately 5 to 10 mins. The dynamic model and MPC were formulated in the Julia programming language, using the JuMP optimization framework, and Ipopt as the non-linear programming solver.

Sequential arrangement of pipelines (Case A)

As shown in Figure 3, the inlet flow rate to pipe segment 1 reduces immediately because the upper bound on the inlet flow rate was reduced from 17 kg/s to 14 kg/s. Subsequently, the optimizer opens valve 1 to depressurize pipe segment 1. As the pressure in pipe segment 1 approaches its lower limit, valve 2 is opened and valve 1 is closed to depressurize pipe segment 2. Note that although pipe segment 3 has a lower limit of approximately 18 bar, it is limited by the outlet pressure boundary condition of 22 bar. Hence, pipe segment 3 is limited in its ability to depressurize. When further depressurization of the pipe segments is no longer possible, the valve openings reach a new steady-state such that the outlet flow rate reduces to 14 kg/s. At $t = 12h$, when the maximum inlet flow rate is restored to its original value of 20 kg/s, the optimizer increases the inlet flow rate and valve openings to maximize throughput and pressurize the system.

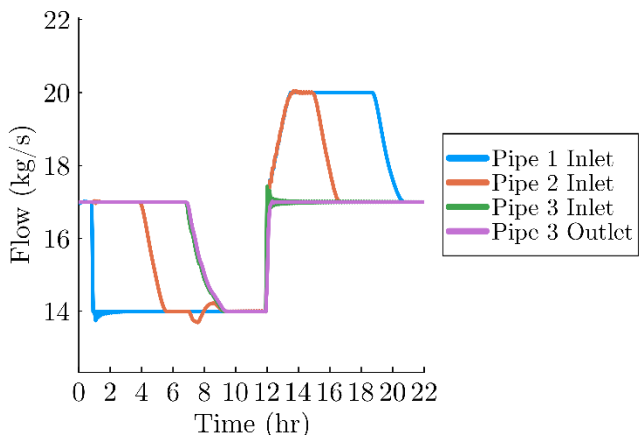

Figure 3. Case A – System flow rates

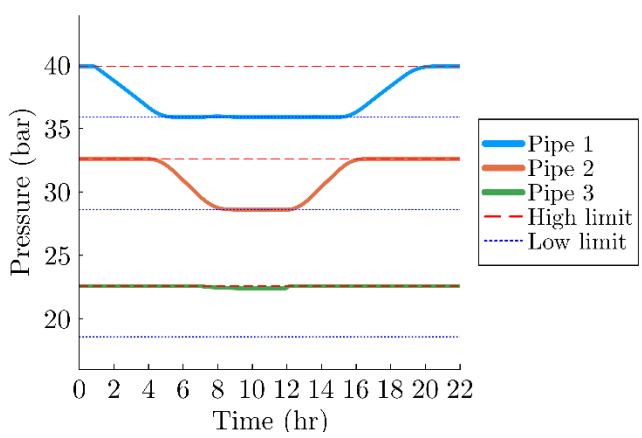

Figure 4. Case A – Pipeline segment midpoint pressures

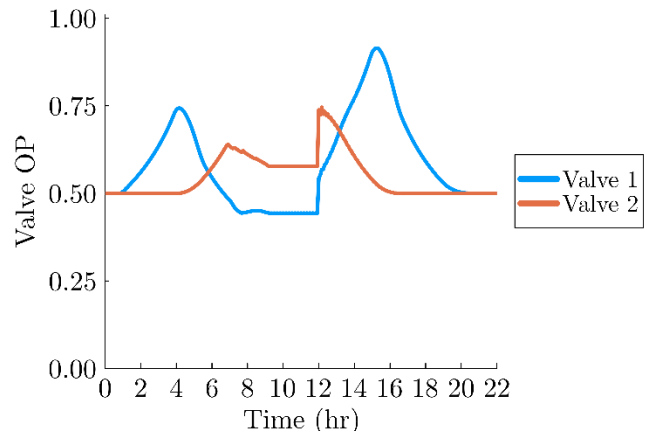

Figure 5. Case A – Valve openings

Parallel arrangement of pipelines (Case B)

For pipeline segments arranged in parallel, the MPC exhibits a similar response. At $t = 1h$, the upper bound on the inlet flow rate to pipe segment 1 is reduced to 10 kg/s. As a result, the optimizer reduces the inlet flow rate to pipe segment 1, increases the inlet flow rate to pipe segment 2, and manipulates the valve openings to depressurize pipe segment 1. Note that we assign the inventory weighting factors (β_i) such that the pressure in pipe segment 3 has the highest priority for maximisation, followed by pipe segment 2 and pipe segment 1. As shown in Figure 7, pipe segment 1 is depressurized followed by pipe segments 2 and 3. Subsequently, the outlet flow rate is also reduced when pipeline inventories are at their minimum pressures. When the upper bound on the inlet flow rate to pipe segment 1 is restored, the optimizer proceeds to maximize pressures in pipe segment 2 followed by pipe segment 1. The original steady-state throughput of 30 kg/s is restored.

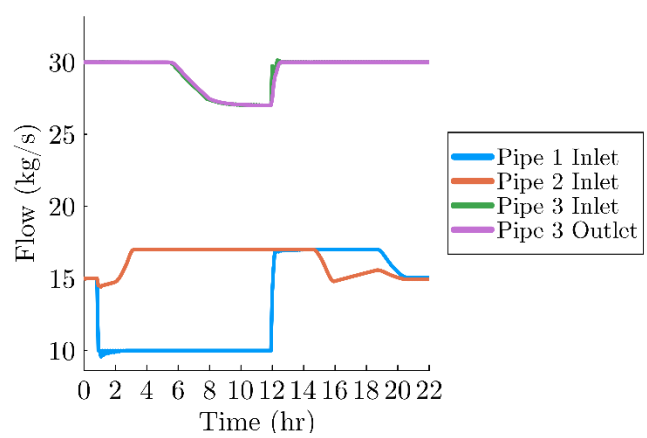

Figure 6. Case B – System flow rates

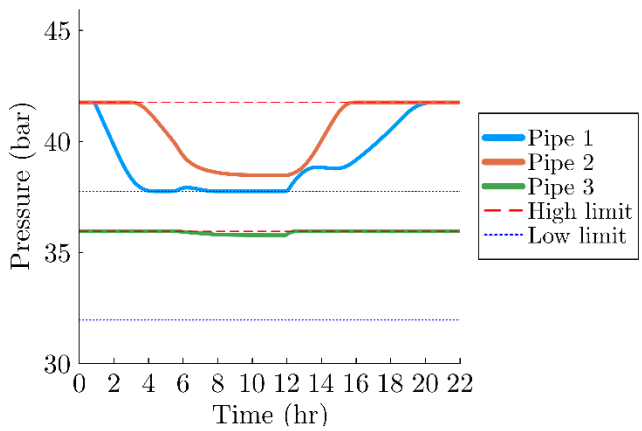

Figure 7. Case B – Pipeline segment midpoint pressures

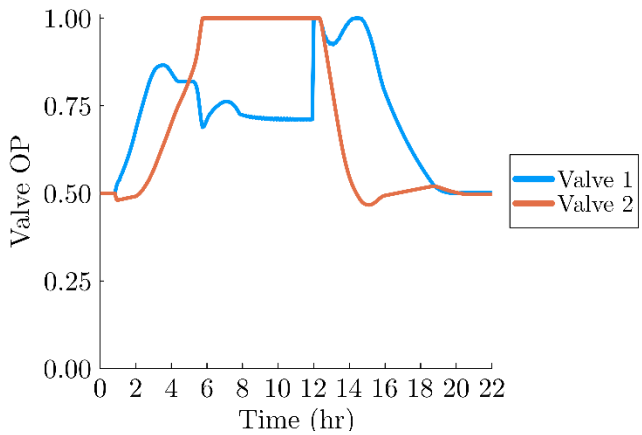

Figure 8. Case B – Valve openings

DISCUSSION

Although throughput maximisation can be economically desirable, this may come at a cost. Firstly, maximising throughput requires aggressive manipulation of the available degrees of freedom (see Fig. 5 and 8). However, this comes with operational wear and higher energy consumption. In systems where rapid MV movement is undesirable or destabilizing, some degree of optimality must be sacrificed. Hence, throughput maximisation should be pursued only as long as the marginal benefit of increased flow outweighs the marginal cost of achieving it.

From a practical perspective, the non-linear dynamic model developed in this study carries inherent assumptions that may limit its application to real-world conditions. Firstly, the pipeline network is assumed to be horizontal. However, real terrain is rarely uniform and incorporating a gravity term into the pressure drop formulation (eq. 2) can improve the model accuracy. In general, a key limitation in the proposed framework is model-plant mismatch, which is an inherent challenge in any model-based control formulation. Despite detailed pressure-flow modelling and the consideration of variable CO_2 compressibility, model-plant mismatch could still result in prediction errors and sub-optimal MV movement. In practice, however, this challenge is well-managed through state estimation where real-time measurements are used to continuously reconcile model predictions. Therefore, the inventory control formulation investigated in this study remains a viable and practical approach for throughput maximization in CO_2 gas pipeline networks.

CONCLUSION

In this paper, an optimal inventory MPC formulation for bottleneck isolation was implemented for CO_2 gas pipelines. The formulation sought to maximize throughput and the weighted sum of pipeline pressures. The formulation was applied with modifications, including the addition of a pressure boundary condition to accommodate pressure and flow dynamics in a pipeline. A flow damping term was also incorporated to ensure regulated MV movement as well as stable flow and pressure variations. Two arrangements of pipelines – sequential and parallel – were investigated and the MPC effectively maximized throughput in the presence of an inlet flow rate disturbance. Future work will extend the application of this inventory control MPC formulation for throughput maximization to larger-scale and more complex low pressure CO_2 gas pipeline networks.

ACKNOWLEDGEMENTS

This work was carried out as a part of SUBPRO Zero, a Research Centre at the Norwegian University of Science and Technology (project code 100589).

REFERENCES

1. Alcalde J. Transportation of CO2 [Internet]. Drax Group Plc; 2025. https://www.drax.com/wp-content/uploads/2025/06/DR1907-6_Transportation-of-CO2_GB_V002.pdf
2. Martynov SB, Porter RTJ, Mahgerefteh H. Estimating the line packing time for pipelines transporting carbon dioxide. Carbon Capture Science & Technology 11:100188 (2024). https://doi.org/10.1016/j.ccst.2023.100188
3. Turan EM, Skogestad S, Jäschke J. Model predictive control for bottleneck isolation with unmeasured faults. IFAC-PapersOnLine 58:767-774 (2024). https://doi.org/10.1016/j.ifacol.2024.08.430
4. Kumaraswamy A, Turan EM, Jäschke J. Optimal inventory control for bottleneck isolation in general processes. Journal of Process Control 155:103557 (2025). https://doi.org/10.1016/j.jprocont.2025.103557

5. Svensson R, Odenberger M, Johnsson F, Strömberg L. Transportation systems for co2—application to carbon capture and storage. Energy Conversion and Management 45:2343-2353 (2004). https://doi.org/10.1016/j.enconman.2003.11.022
6. Onyebuchi VE, Kolios A, Hanak DP, Biliyok C, Manovic V. A systematic review of key challenges of CO2 transport via pipelines. Renewable and Sustainable Energy Reviews 81:2563-2583 (2018). https://doi.org/10.1016/j.rser.2017.06.064
7. Metz B, Davidson O, Coninck H de, Loos M, Meyer L. IPCC special report on carbon dioxide capture and storage [Internet]. Intergovernmental Panel on Climate Change; 2005. https://www.ipcc.ch/report/carbon-dioxide-capture-and-storage/
8. OGE. CLB pipeline construction project [Internet]. https://oge.net/en/press-public-information/construction-projects/current-construction-projects-co2/clb-pipeline-construction-project
9. Equinor. Northern lights [Internet]. https://www.equinor.com/energy/northern-lights
10. Carbfix. A terminal coda [Internet]. https://www.carbfix.com/operations/codaterminal
11. Porthos. CO2 reduction through storage under the north sea [Internet]. https://www.porthosco2.nl/en/
12. Spitz T, Avagyan V, Ascui F, Bruce ARW, Chalmers H, Lucquiaud M. On the variability of CO2 feed flows into CCS transportation and storage networks. International Journal of Greenhouse Gas Control 74:296-311 (2018). https://doi.org/10.1016/j.ijggc.2018.04.008
13. Zhou C, Liu P, Pistikopoulos EN. Superstructure-based optimal design of pipeline network for CO2 transport in large-scale carbon capture and sequestration [Internet]. Computer aided chemical engineering. 2015. p. 225–52. https://doi.org/10.1016/B978-0-444-63472-6.00009-4
14. Ismail I, Gaganis V. Carbon capture, utilization, and storage in saline aquifers: subsurface policies, development plans, well control strategies and optimization approaches—a review. Clean Technol. 5:609-637 (2023). https://doi.org/10.3390/cleantechnol5020031
15. Yang Z, Fahmi A, Drescher M, Teberikler L, Merat C, Solvang S, et al. Improved understanding of flow assurance for CO_2 transport and injection. Proceedings of the 15th greenhouse gas control technologies conference. 2021.
16. Aursand P, Hammer M, Munkejord ST, Wilhelmsen Ø. Pipeline transport of CO2 mixtures: models for transient simulation. International Journal of Greenhouse Gas Control 15:174-185 (2013). https://doi.org/10.1016/j.ijggc.2013.02.012
17. Zhou L, Preindl M. Mpc-based harmonic injection techniques for reconfigurable single/three-phase inverters with grid neutral point connection. 2024 IEEE Transportation Electrification Conference and Expo (ITEC) :1-6 (2024). https://doi.org/10.1109/itec60657.2024.10598838
18. Xu Q, Zhang X, Sun B. Modeling and model predictive control of CO2 salt cavern storage process. Fuel 405:136633 (2026). https://doi.org/10.1016/j.fuel.2025.136633
19. Baumrucker BT, Biegler LT. MPEC strategies for cost optimization of pipeline operations. Computers & Chemical Engineering 34:900-913 (2010). https://doi.org/10.1016/j.compchemeng.2009.07.012
20. Gopalakrishnan A, Biegler LT. Economic nonlinear model predictive control for periodic optimal operation of gas pipeline networks. Computers & Chemical Engineering 52:90-99 (2013). https://doi.org/10.1016/j.compchemeng.2012.11.011
21. Moetamedzadeh HR, Khanmirza E, Pourfard A, Madoliat R. Intelligent nonlinear model predictive control of gas pipeline networks. Transactions of the Institute of Measurement and Control 41:4569-4589 (2019). https://doi.org/10.1177/0142331219864190
22. Aske EMB, Skogestad S. Consistent inventory control. Ind. Eng. Chem. Res. 48:10892-10902 (2009). https://doi.org/10.1021/ie801603j
23. Zotic? C, Forsman K, Skogestad S. Bidirectional inventory control with optimal use of intermediate storage. Computers & Chemical Engineering 159:107677 (2022). https://doi.org/10.1016/j.compchemeng.2022.107677
24. Çengel YA, Cimbala JM. Fluid mechanics: Fundamentals and applications [Internet]. McGraw-HillHigher Education; 2006. https://books.google.no/books?id=4RVEAQAAIAAJ

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Causal Discovery for the Spatial Autoregressive Model: Application to Defect Analysis in the Plastic Injection Molding Process

Ryosuke Tanaka[ab], Koichi Fujiwara[acd*]

[a] Nagoya University, Department of Materials Process Engineering, Nagoya, Aichi, Japan
[b] MCC Advanced Moldings Co., Ltd., Yokkaichi, Mie, Japan
[c] Hokkaido University, Research Institute for Electronic Science, Sapporo, Hokkaido, Japan
[d] Nara Institute of Science and Technology, Medilux Research Center, Ikoma, Nara, Japan
* Corresponding Author: fujiwara.koichi@hps.material.nagoya-u.ac.jp.

ABSTRACT

Plastic injection molding is a widely used polymer-processing method. As the requirements for processing accuracy have become increasingly stringent, defect analysis in plastic injection molding is necessary to improve the product yield. Causal discovery has recently gained attention for defect analysis in many processes. Because injection molding is a spatial process involving the distribution of physical quantities, spatial autocorrelation should be considered. Although the linear non-Gaussian acyclic model (LiNGAM) is a well-known causal discovery method, it cannot properly model spatial autocorrelation. In this study, a new causal discovery method for a spatially autocorrelated dependent variable, referred to as the Causal Structure Search for the Spatial Autoregressive Model (CASSPAR), is proposed. It models the causal relationships among the observed points without prior knowledge of the spatial structure. The proposed method represents the causal relationships among the observed points as a causal graph and estimates the adjacency matrix of the graph. The adjacency matrix is estimated using LiNGAM, and the model parameters are estimated using two-stage least squares (2SLS). The usefulness of the proposed CASSPAR was demonstrated using simulation data of a plastic injection molding process to identify the root cause of warpage.

Keywords: Polymers, Modelling and Simulations, Machine Learning, Algorithms

INTRODUCTION

Injection molding is a widely used polymer processing method in which a molten resin is injected into a mold and solidified to form a product. Warpage is a post-molding product deformation that hinders the manufacture of high-quality products. In Figure 1, warpage may occur in completely opposite directions depending on the process parameters, even for identical products, which makes it difficult to analyze and control the causes of defects. Warpage may be influenced by multiple factors, including the temperature distribution, pressure distribution, product geometry, and materials [1]. Because these mechanisms often vary across individual products, identifying the causes of warpage and modifying the process operations are crucial for improving product quality.

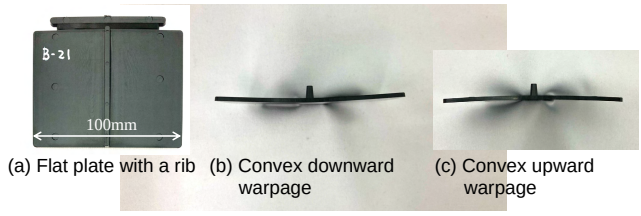

(a) Flat plate with a rib (b) Convex downward warpage (c) Convex upward warpage

Figure 1: An example of warpage. A product of the same shape exhibits warpage in opposite directions depending on process parameters

Simulation studies have been successfully applied to predict warpage and optimize the process operations. Chen *et al.* investigated process variables affecting warpage using a design of experiments (DOE) on

simulation data [2]. In addition, the response surface method (RSM) in multi-objective optimization for warpage suppression and cycle time has been reported [3].

Machine learning (ML) techniques have also been applied to defect analysis. For instance, a warpage prediction model was constructed using random forests [4], and neural networks (NN) were trained on time-series data of in-mold pressure [5].

However, ML models and ML model-based optimization often become black boxes. Although they may achieve good performance in warpage prediction, their lack of interpretability makes it challenging to identify the root causes of warpage and develop fundamental countermeasures.

Causal discovery has gained attention for analyzing industrial process data [6, 7]. The linear non-Gaussian acyclic model (LiNGAM) is a well-known causal discovery method that estimates the causal structure among variables as a graph [8]. A key feature of LiNGAM is its statistical identifiability; it can uniquely identify the causal graph from observational data alone without prior knowledge. Various LiNGAM extensions have been proposed, including time-series data [9]. However, existing variants do not take spatial information into account.

Injection molding is considered a spatial process because physical quantities, such as temperature and pressure, are distributed across the product and mold. For example, spatial nonuniformity in temperature and stress can induce warpage [10, 11].

Spatial autocorrelation needs to be considered in spatial data, which means that observations at neighboring points are mutually dependent. This characteristic violates the independence assumption underlying many statistical methods, such as linear regression [12]. LiNGAM estimates also become biased from spatial data because each variable is no longer independent due to spatial autocorrelation. Therefore, a new causal discovery method for addressing spatial autocorrelation is required.

To express the relationship among observed points, a spatial weight matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$ is used, where $\mathbf{W}$ can be defined as the adjacency or distance matrix between two points, and N is the number of points. The spatial autoregressive (SAR) model is a widely used regression method that uses $\mathbf{W}$ to incorporate the dependent variable y at neighboring locations as a predictor.

In most spatial models, $\mathbf{W}$ has to be given in advance, which means that the spatial autocorrelation is already known; however, this assumption is unrealistic for injection molding processes because it is difficult to observe such a relationship during molding. To estimate the spatial weight matrix $\mathbf{W}$ without prior constraints becomes difficult when the number of elements in $\mathbf{W}$ exceeds the number of observed points N. Thus, the structure of $\mathbf{W}$ has to be constrained with some assumptions [13]. However, it is unclear whether the injection molding process data satisfy these assumptions.

A key idea of this study is to interpret $\mathbf{W}$ as a representation of causal relationships because the spatial distributions in the injection molding process always follow the causal relationship by physicochemical mechanisms. Thus, $\mathbf{W}$ is regarded as an adjacency matrix of a causal graph. By applying LiNGAM, $\mathbf{W}$ can be uniquely identified and estimated directly from the data.

This study proposes a new causal discovery method for the spatial autoregressive model, referred to as the Causal Structure Search for the Spatial Autoregressive Model (CASSPAR). The proposed CASSPAR reveals causal relationships among locations without prior knowledge of spatial structure. This capability enables more precise defect analysis of warpage, leading to improved product yield and reduced energy consumption and costs. The estimation performance of CASSPAR is demonstrated through its application to simulation data of an injection-molding process.

METHOD

The proposed CASSPAR interprets the spatial weight matrix $\mathbf{W}$ in the SAR model as a spatial causal structure and estimates $\mathbf{W}$ using LiNGAM. The background on LiNGAM and the SAR model is provided, and the estimation procedure of CASSPAR is explained in the following sections.

Linear non-Gaussian acyclic model (LiNGAM)

LiNGAM is a causal discovery method that represents causal relationships among variables using a directed acyclic graph. Figure 2 shows an example of a causal graph used to investigate the causes of warpage in injection molding. The strength of the causal relationship is denoted by a coefficient b. In LiNGAM, each observed variable x_i is expressed as a linear sum of other observed variables x_j and an error term e_i. The causal graph in Figure 2 corresponds to the following structural equations:

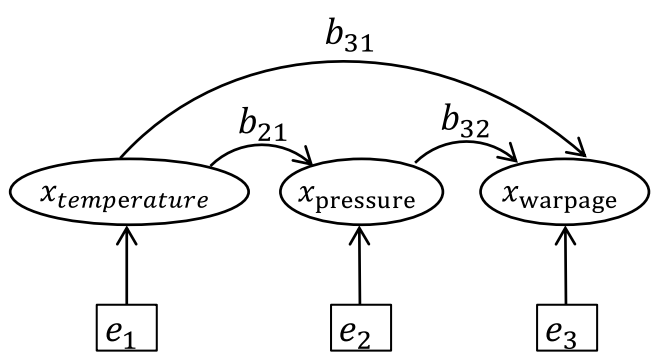

Figure 2: An example of a causal structure estimated by LiNGAM, represented as a causal graph.

$$
\begin{aligned}
x_{temperature} &= e_1 \\
x_{pressure} &= b_{21} x_{temperature} + e_2 \\
x_{warpage} &= b_{31} x_{temperature} + b_{32} x_{pressure} + e_3.
\end{aligned}
\tag{1}
$$

This figure illustrates the causal relationships, indicating that warpage is affected by both the process temperature and pressure, and that pressure is also affected by the temperature. In matrix form, Eq. (1) is written as

$$
\mathbf{x} = \mathbf{B}\mathbf{x} + \mathbf{e} \tag{2}
$$

where $\mathbf{B} \in \mathbb{R}^{K \times K}$ denotes the adjacency matrix of the causal graph to be estimated, $\mathbf{x} \in \mathbb{R}^K$ is the vector of K observed variables, and $\mathbf{e} \in \mathbb{R}^K$ represents the vector of error terms. To ensure the identifiability of the causal graph, the following assumptions are required: 1) The error variables follow non-Gaussian distributions and are mutually independent. 2) The causal graph is acyclic, which means that no directed cycles exist in the causal graph.

Spatial autoregressive (SAR) model

Spatial autocorrelation is a key characteristic of spatial data, which is the tendency for observations at neighboring locations to be more similar (or dissimilar) than those at more distant locations. Positive spatial autocorrelation corresponds to a pattern in which neighboring locations have similar values.

The spatial weight matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$ is used to model spatial autocorrelation, which is a matrix representation of the adjacency relationships among N locations. Figure 3 shows typical definitions of $\mathbf{W}$. The element w_{ij} in $\mathbf{W}$ represents the influence of location j on location i. The diagonal elements w_{ii} are set to zero to avoid influencing itself. For reasons of computational stability and interpretability, $\mathbf{W}$ is commonly row-standardized so that each row sums to one. Figure 3 (a)–(c) show an example of a typical $\mathbf{W}$ indicating which locations affect the central cell shown in red. Color intensity represents the strength of the influence.

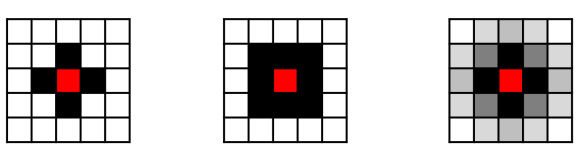

(a) Rook adjacency (b) Queen adjacency (c) Distance based W

Figure 3. Typical definitions of the spatial weight matrix $\mathbf{W}$ in 3 × 3 lattice space.

Various spatial regression models can be defined depending on which variable, dependent variable, independent variable, or error term exhibits spatial autocorrelation. This study focuses on the spatial autoregressive (SAR) model, also called the spatial lag model (SLM), which assumes that the dependent variable y is spatially autocorrelated. This assumption is natural for warpage prediction, and the dependent variable y represents warpage. Warpage exhibits positive spatial autocorrelation because deformation at one location typically causes the surrounding areas to deform in the same direction. The SAR model is given by

$$
y_i = \rho \sum_{j=1}^{N} w_{ij}\, y_j + \mathbf{X_i}\beta + e_i \tag{3}
$$

where y_i is the dependent variable at location i, ρ denotes the spatial autocorrelation coefficient, w_{ij} represents the spatial weight between locations i and $j, \mathbf{X_i} \in \mathbb{R}^K$ is a vector of K independent variables at location $i, \boldsymbol{\beta} \in \mathbb{R}^K$ is a coefficient vector, and e_i is the error term at location i.

Figure 4 presents a schematic of the SAR model used to predict warpage in injection molding. As shown, the dependent variable y_i, which represents warpage, is influenced by neighboring warpages, such as y_j, y_k, y_l. In addition, y_i is influenced by other variables $\mathbf{X_i}$ at the same location (e.g., product temperature and pressure), as well as by the error term e_i.

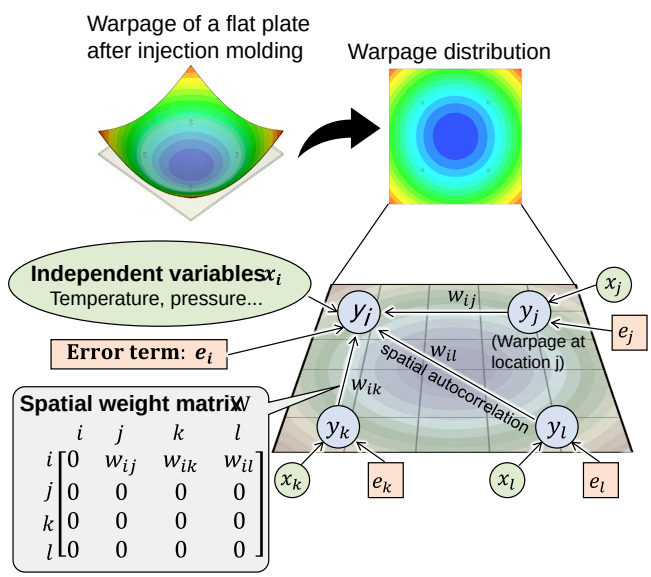

Figure 4: A schematic of SAR model for predicting warpage in injection molding.

The spatial autocorrelation coefficient ρ denotes the strength of the influence that a location receives from its neighbors. A large positive value of ρ represents a positive spatial autocorrelation. The term $\sum_{j=1}^{N} w_{ij}\, y_j$ is called the spatial lag. The variables ρ and $\boldsymbol{\beta}$ are the model parameters to estimate.

Causal Structure Search for the Spatial Autoregressive Model (CASSPAR).

CASSPAR estimates the spatial weight matrix $\mathbf{W}$ and the SAR parameters: $\rho, \boldsymbol{\beta}$. To apply LiNGAM to spatial data, the SAR model is reformulated. The matrix form of the SAR model is defined as

$$\mathbf{Y} = \rho\mathbf{W}\mathbf{Y} + \mathbf{X}\boldsymbol{\beta} + \mathbf{e} \tag{4}$$

where $\mathbf{Y} \in \mathbb{R}^{T\times N}$ and $\mathbf{X} \in \mathbb{R}^{T\times N\times K}$ denote the measurements of the dependent and independent variables, respectively. $\mathbf{W} \in \mathbb{R}^{N\times N}$ is a spatial weight matrix, ρ and $\boldsymbol{\beta} \in \mathbb{R}^{K}$ are spatial autoregressive coefficient and regression coefficient vectors, and $\mathbf{e}$ is an error term. Note that $\mathbf{W}$ is assumed to be row-standardized and spectrally normalized.

Algorithm 1 Step 1: Estimation of spatial weight matrix $\widehat{\mathbf{W}}$

1. **Input:**
   Observed dependent variable matrix $\mathbf{Y} \in \mathbb{R}^{T\times N}$
2. **Output:**
   Estimated spatial weight matrix after normalization $\widehat{\mathbf{W}} \in \mathbb{R}^{N\times N}$
3. Apply the Direct-LiNGAM algorithm to $\mathbf{Y}$ and obtain an estimated adjacency matrix $\widehat{\mathbf{A}} \in \mathbb{R}^{N\times N}$.
4. $\widehat{\mathbf{A}}$ is first row-standardized and then spectrally normalized to obtain the normalized matrix $\widehat{\mathbf{W}}$.
5. Return $\widehat{\mathbf{W}}$

By defining $\mathbf{A} = \rho\mathbf{W}$ and $\mathbf{e}' = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}$, SAR can be reformulated in the same form as the LiNGAM model:

$$\mathbf{Y} = \mathbf{A}\mathbf{Y} + \mathbf{e}' \tag{5}$$

This reformulation allows us to estimate the spatial structure using LiNGAM. However, the following assumptions are required for LiNGAM identification:

Assumption 1: Spatial weight matrix $\mathbf{W}$ (and thus $\mathbf{A}$) corresponds to an adjacency matrix of a directed acyclic graph (DAG). In other words, $\mathbf{A}$ can be permuted into a strictly lower triangular matrix. This assumption is natural in the context of injection molding, where the resin flows in one direction, that is, the influence propagates only from upstream to downstream locations.

Assumption 2: Each component of the external error vector $\boldsymbol{e}'$ follow a non-Gaussian distribution, and they are mutually independent.

This study proposes a two-step estimation procedure. In step 1, the structure of the spatial weight matrix $\mathbf{W}$ is estimated, and in step 2, the SAR parameters ρ and $\boldsymbol{\beta}$ are estimated. Figure 5 illustrates the procedure of the proposed CASSPAR.

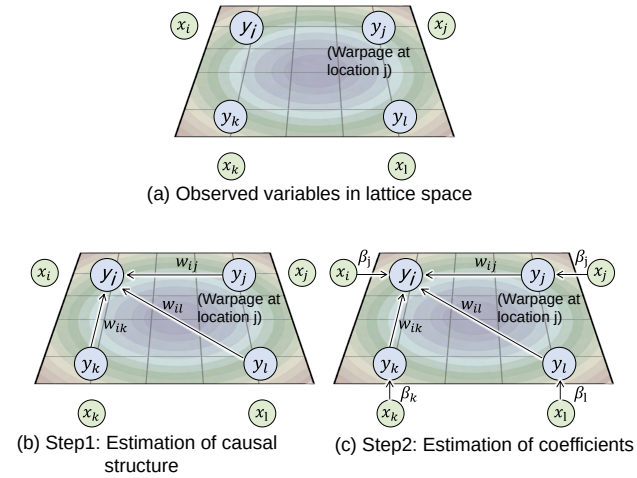

Figure 5: Conceptual illustration of the proposed two-step estimation procedure.

Step 1: Estimation of the spatial weight matrix

The Direct-LiNGAM algorithm [14] is applied to the reformulated LiNGAM model in Eq. (5), and the estimated adjacency matrix $\widehat{\mathbf{A}} \in \mathbb{R}^{N\times N}$ is obtained. Although step 1 estimates the adjacency matrix $\widehat{\mathbf{A}}$, a normalized weight matrix $\mathbf{W}$ is required for step 2. Therefore, $\widehat{\mathbf{W}}$ is obtained by row-standardizing and spectral normalizing $\widehat{\mathbf{A}}$ (Algorithm 1).

Step 2: Estimation of SAR Parameters via 2SLS

Given $\widehat{\mathbf{W}}$, the SAR parameters ρ and $\boldsymbol{\beta}$ are estimated by an instrumental variables (IV) approach based on two-stage least squares (2SLS). Using the normalized spatial weight matrix $\widehat{\mathbf{W}}$ obtained in Step 1, the original SAR equation is formulated as:

$$\mathbf{Y} = \rho\widehat{\mathbf{W}}\mathbf{Y} + \mathbf{X}\boldsymbol{\beta} + \mathbf{e}. \tag{6}$$

The endogenous problem arises in Eq. (6) because the spatial lag term $\widehat{\mathbf{W}}\mathbf{Y}$ is simultaneously determined with dependent variable $\mathbf{Y}$. Consequently, $\widehat{\mathbf{W}}\mathbf{Y}$ is correlated with the error term $\mathbf{e}$. Due to this endogeneity, OLS estimates biased ρ and $\boldsymbol{\beta}$. Thus, this study adopts the IV method to estimate these parameters, which is a standard estimation method for SAR models [15].

For the convenience of the IV method, Eq. (6) is rearranged as follows: let $\mathbf{y}_t \in \mathbb{R}^{N}$ and $\mathbf{X}_t \in \mathbb{R}^{N\times K}$ denote the t-th sample of $\mathbf{Y}$ and $\mathbf{X}$, respectively, $\mathbf{y}_{\text{stack}} \in \mathbb{R}^{TN}$ and $\mathbf{X}_{\text{stack}} \in \mathbb{R}^{TN\times K}$ are vertically stacked matrices of $\{\mathbf{y}_t\}_{t=1}^{T}$ and $\{\mathbf{X}_t\}_{t=1}^{T}$. The corresponding stacked spatial lag term is defined as

$$\mathbf{S} = [(\widehat{\mathbf{W}}\mathbf{y}_1)^\top, \dots, (\widehat{\mathbf{W}}\mathbf{y}_T)^\top]^\top \in \mathbb{R}^{TN}. \tag{7}$$

Using $\mathbf{S}$, the SAR model can be expressed as

$$\mathbf{y}_{\text{stack}} = \rho\mathbf{S} + \mathbf{X}_{\text{stack}}\boldsymbol{\beta} + \mathbf{e}_{\text{stack}} \tag{8}$$

where $\mathbf{e}_{\text{stack}} \in \mathbb{R}^{TN}$ denotes the stacked error term.

The instrument variable $\mathbf{Z}$ is constructed from the exogenous variables $\mathbf{X}$ and their spatial lags. For each lag order $\ell = 1, \dots, L$, the matrices

$$\mathbf{Z}_t^{(\ell)} = \widehat{\mathbf{W}}^{\ell}\mathbf{X}_t \in \mathbb{R}^{N\times K} \tag{9}$$

are defined, and $\mathbf{Z}^{(\ell)} \in \mathbb{R}^{TN\times K}$ is obtained by stacking $\{\mathbf{Z}_t^{(\ell)}\}_{t=1}^{T}$ vertically. The full instrument matrix is then given by

$$\mathbf{Z} = [\,\mathbf{X}_{\text{stack}}, \mathbf{Z}^{(1)}, \dots, \mathbf{Z}^{(L)}\,] \in \mathbb{R}^{TN\times K(L+1)}. \tag{10}$$

This IV procedure requires the following assumptions about the instruments:

Assumption 3: The exogenous variables are uncorrelated with the error term $e, \mathrm{Cov}[Z, e] = 0$, to ensure that the instruments are exogenous and valid for the two-stage least squares (2SLS) estimation.

Assumption 4: The instruments are correlated with the endogenous spatial lag term $S, \mathrm{Cov}[Z, S] \neq 0$, which means that the instruments have sufficient relevance for the spatial lag term.

The SAR parameters ρ and β can be estimated via

Algorithm 2 Step 2: Estimation of SAR parameters ρ and β via 2SLS

1. **Input:**
   Normalized spatial weight matrix $\widehat{\mathbf{W}} \in \mathbb{R}^{N\times N}$.
   Observed dependent variable matrix $\mathbf{Y} \in \mathbb{R}^{T\times N}$.
   Observed independent variables tensor $\mathbf{X} \in \mathbb{R}^{T\times N\times K}$.
2. **Output:**
   Spatial autoregressive coefficient $\hat{\rho} \in \mathbb{R}$.
   Regression coefficients $\hat{\boldsymbol{\beta}} \in \mathbb{R}^{K}$.
3. Stack $\mathbf{Y}$ and $\mathbf{X}$ over samples to obtain $\mathbf{y}_{\text{stack}} \in \mathbb{R}^{TN}$ and $\mathbf{X}_{\text{stack}} \in \mathbb{R}^{TN\times K}$.
4. Consider the SAR model with the estimated spatial weight, $\mathbf{y}_{\text{stack}} = \rho\mathbf{S} + \mathbf{X}_{\text{stack}}\boldsymbol{\beta} + \mathbf{e}_{\text{stack}}$ here, $\mathbf{S}$ is the stacked spatial lag term.
5. Construct the instrument matrix for the t-th sample: $\mathbf{Z}_t = \left[\mathbf{X}_t,\ \widehat{\mathbf{W}}\mathbf{X}_t,\ \widehat{\mathbf{W}}^2\mathbf{X}_t,\ \dots,\ \widehat{\mathbf{W}}^L\mathbf{X}_t\right] \in \mathbb{R}^{TN\times K(L+1)}$
6. Apply two-stage least squares (2SLS) to the stacked SAR with the spatially lagged term as the endogenous regressor and $\mathbf{Z}$ as the instrument matrix, Obtain the estimates $\hat{\rho}$ and $\hat{\boldsymbol{\beta}}$.
7. Return $\hat{\rho}, \hat{\boldsymbol{\beta}}$

the 2SLS method. In the first stage of 2SLS, the endogenous spatial lag term $\mathbf{S}$ is regressed on the instrument matrix $\mathbf{Z}$ by OLS, and the fitted values $\hat{\mathbf{S}}$ are obtained. In the second stage, $\mathbf{y}_{\text{stack}}$ is regressed on $[\hat{\mathbf{S}}, \mathbf{X}_{\text{stack}}]$ by OLS. The coefficient on $\hat{\mathbf{S}}$ and the coefficient vector on $\mathbf{X}_{\text{stack}}$ yield the 2SLS estimates $\hat{\rho}$ and $\hat{\boldsymbol{\beta}}$, respectively. Under Assumptions 1-4, these estimators are consistent with the parameters of the original SAR. Algorithm 2 describes the procedure of SAR parameter estimation.

APPLICATION TO SIMULATION OF INJECTION MOLDING

In this section, the application results of the proposed CASSPAR to simulation data of injection molding to verify its capability to identify the causes of warpage are reported.

One of the primary causes of warpage is a non-uniform temperature distribution. Figure 6 (a) shows the 3D model of the injection-molding process used in the simulation, and (b)-(d) illustrate the mechanism of product deformation: (b) a temperature difference exists between the top and bottom of the product, which (c) leads to non-uniform resin shrinkage during solidification, and (d) consequently causes product deformation.

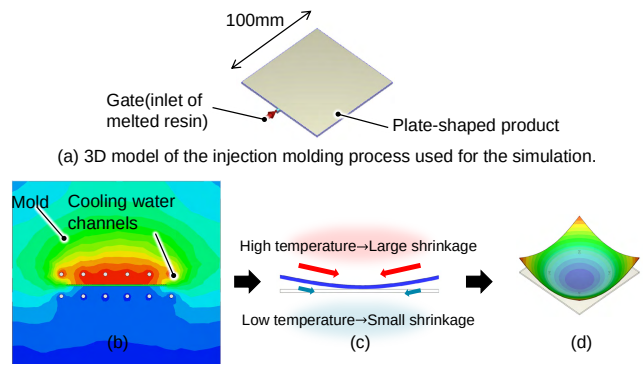

Figure 6: (a) 3D model of the injection molding process used for the simulation; (b)-(d) show the mechanisms of warpage.

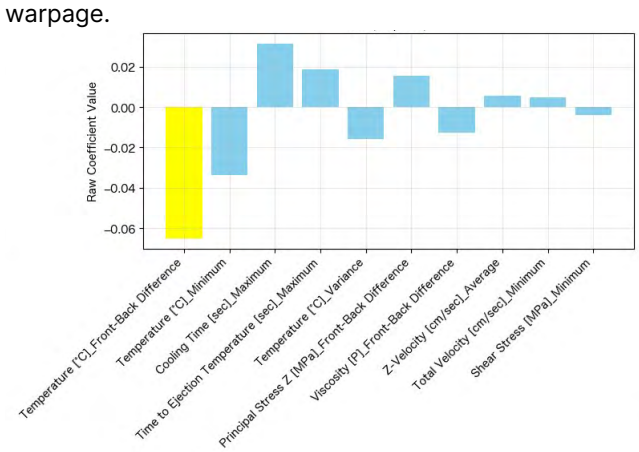

Figure 7: The top 10 largest absolute coefficient values estimated by the proposed CASSPAR.

In the simulation, a temperature difference between the top and bottom surfaces was intentionally introduced. A 9×9 grid of measurement points ($N = 81$) was placed on the plate for data collection. The data included variables such as the warpage, temperature difference between the top and bottom of the product, pressure, and resin flow velocity. The warpage at the 81 locations was treated as the endogenous spatial variable Y,

whereas the temperature difference, pressure, and flow velocity were treated as the exogenous variables X.

Figure 7 shows the top 10 largest absolute coefficient values estimated by the proposed CASSPAR. The top-ranked variable, "Temperature [℃]_Front-Back Difference, " represents the temperature difference between the top and bottom surfaces of the product. This result indicates that the proposed CASSPAR successfully identified the intentionally introduced temperature difference as the top-ranking causal factor of warpage.

Figure 8 shows the estimated spatial weight matrix W for the warpage, which is expected to exhibit positive spatial autocorrelation because a deformation at one location typically induces deformations in the same direction in the surrounding areas. The estimated $\mathbf{W}$ exhibits a local neighborhood structure similar to contiguity-based matrices, such as rook or queen adjacency.

This result suggests that the proposed CASSPAR successfully captured the physically reasonable spatial interactions underlying the warpage deformation.

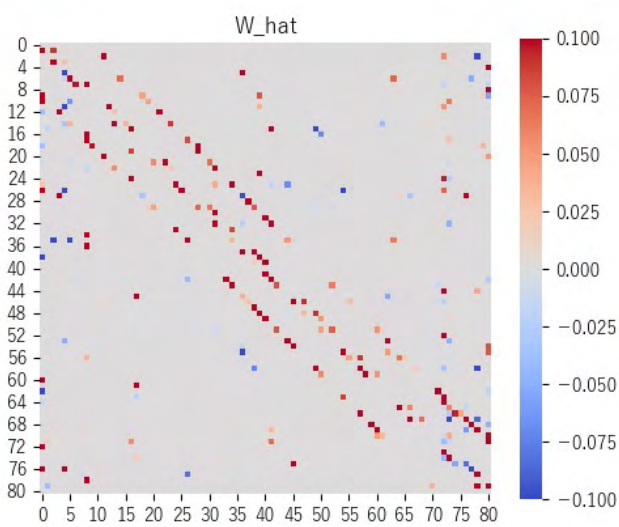

Figure 8. Heatmap of the estimated spatial weight matrix $\mathbf{W}$.

CONCLUSIONS

In this study, a new causal discovery method for spatial data, referred to as CASSPAR, was proposed. A limitation of this study is that CASSPAR considers spatial autocorrelation only in the dependent variable Y. However, in many industrial processes, the independent variables X may also exhibit spatial dependence. CASSPAR can also be applicable to other spatially dependent processes, such as fault diagnosis in chemical plants and semiconductor wafer manufacturing. Future work includes validation in real-world experiments and extension to industrial processes beyond injection molding.

ACKNOWLEDGEMENTS

This work was based on the results obtained from project JPNP20004, subsidized by the New Energy and Industrial Technology Development Organization (NEDO).

REFERENCES

1. Coca-Gonzalez M, Jimenez-Martinez M. Warpage: causes, manufacturing processes and future challenges: a review. Proceedings of the Institution of Mechanical Engineers, Part L: Journal of Materials: Design and Applications 239:1201-1217 (2024). https://doi.org/10.1177/14644207241285399
2. Chen CP, Chuang MT, Hsiao YH, Yang YK, Tsai CH. Simulation and experimental study in determining injection molding process parameters for thin-shell plastic parts via design of experiments analysis. Expert Systems with Applications 36:10752-10759 (2009). https://doi.org/10.1016/j.eswa.2009.02.017
3. Kitayama S, Yokoyama M, Takano M, Aiba S. Multi-objective optimization of variable packing pressure profile and process parameters in plastic injection molding for minimizing warpage and cycle time. Int J Adv Manuf Technol 92:3991-3999 (2017). https://doi.org/10.1007/s00170-017-0456-1
4. Ahmed, T., Sharma, P. & Karmaker, C. Warpage prediction of Injection-molded PVC part using ensemble machine learning algorithm. Mater. Today Proc. 50, (2020).
5. Gim J, Lin CY, Turng LS. In-mold condition-centered and explainable artificial intelligence-based (IMC-XAI) process optimization for injection molding. Journal of Manufacturing Systems 72:196-213 (2024). https://doi.org/10.1016/j.jmsy.2023.11.013
6. Uchida Y, Fujiwara K, Saito T, Osaka T. Causal plot: causal-based fault diagnosis method based on causal analysis. Processes 10:2269 (2022). https://doi.org/10.3390/pr10112269
7. Xue J, Zhang T, Ye H. Kpi-oriented process monitoring based on causal-weighted partial least squares. Information Sciences 689:121470 (2025). https://doi.org/10.1016/j.ins.2024.121470
8. Shimizu, S., Hoyer, P. O. & Hyvärinen, A. A Linear Non-Gaussian Acyclic Model for Causal Discovery. Journal of Machine Learning Research vol. 7 (2006).
9. Hyvärinen, A., Zhang, K., Shimizu, S. & Hoyer, P. O. Estimation of a Structural Vector Autoregression Model Using Non-Gaussianity. J. Mach. Learn. Res. 11, 1709–1731 (2010).
10. Wang J, Liu T, Zheng K, Liu H, Cui H, Li H. Local

thermal warpage deformation of polypropylene injection molded flat part and neural network prediction model. Front. Mater. 11: (2024). https://doi.org/10.3389/fmats.2024.1421546
11. Choi S, Im YT. Prediction of shrinkage and warpage in consideration of residual stress in integrated simulation of injection molding. Composite Structures 47:655-665 (1999). https://doi.org/10.1016/s0263-8223(00)00045-3
12. Pace R, LeSage J. Spatial econometrics. Chapman & Hall/CRC Handbooks of Modern Statistical Methods :245-260 (2010). https://doi.org/10.1201/9781420072884-c15
13. Ahrens A, Bhattacharjee A. Two-step lasso estimation of the spatial weights matrix. Econometrics 3:128-155 (2015). https://doi.org/10.3390/econometrics3010128
14. Shimizu, S. et al. DirectLiNGAM: A Direct Method for Learning a Linear Non-Gaussian Structural Equation Model. The Journal of Machine Learning Research 12, 1225–1248 (2011).
15. Kelejian HH, Prucha IR. A generalized spatial two-stage least squares procedure for estimating a spatial autoregressive model with autoregressive disturbances. The Journal of Real Estate Finance and Economics 17:99-121 (1998). https://doi.org/10.1023/a:1007707430416

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Industrial batch process online fault detection using machine learning

Oliver Pennington[a], Adam Wilson[b], Carolina Cruz[b], and Dongda Zhang[a]*

[a] University of Manchester, Department of Chemical Engineering, Manchester, Greater Manchester, UK
[b] KEIT Industrial Analytics, Didcot, Oxfordshire, UK
* Corresponding Author: dongda.zhang@manchester.ac.uk.

ABSTRACT

As industries pursue more sustainable and flexible manufacturing strategies, batch processes continue to play a vital role across a wide range of applications. Batch operations offer the ability to handle diverse feedstocks and accommodate varying product specifications. These processes are broadly used in sectors such as pharmaceuticals, specialty chemicals, food production, and bioprocesses, where precise control over reaction conditions and product quality is essential. However, maintaining optimal conditions in a batch process can be challenging due to the minimal opportunities for mid-batch interference. This work focuses on a real industrial batch process that frequently sees batches with poor yields resulting in large financial losses. Despite utilizing a mid-infrared spectrometer analyzing the batch medium in real-time, the reduced product accumulation observed in faulty batches is not evident until over a third of the batch time has passed, by which point the batch is not economically viable to abort. This study applies and compares various machine learning based fault detection strategies, including multiple Principal Component Analysis (PCA) variants and autoencoder variants, to identify faulty batches as soon as possible, since resetting reset the batches earlier can maximize overall productivity. The findings of this study offer faster and more robust fault detection than typical PCA for this industrial batch process, reducing time lost to faulty batches and improving overall productivity. This work supports the transition towards autonomous and digitalized batch manufacturing while providing an in-depth comparison between several online fault detection strategies on real industrial data.

Keywords: Fault Detection, Batch Process, Machine Learning, Process Control, Modelling

INTRODUCTION

Batch processes remain a cornerstone of modern manufacturing systems due to their versatility and robustness, particularly in applications requiring product customization, process flexibility, or adaptability to evolving feedstock supply chains. Despite these advantages, fault detection in batch processes remains challenging due to inherent process complexity, time-varying dynamics, and measurement noise. As a result, distinguishing genuine process faults from normal operational variability requires robust monitoring frameworks capable of reliably identifying true anomalies in real time.

PCA is a well-established multivariate statistical method that has been widely applied to fault detection and process monitoring in batch processes [1–4]. Owing to its simplicity, interpretability, and strong theoretical foundation, PCA serves as a natural benchmark against which more advanced fault detection strategies can be compared. However, conventional PCA treats observations as independent and does not explicitly account for process dynamics, which limits its effectiveness for batch operations.

To address this limitation, Dynamic PCA (DPCA) extends the traditional PCA framework by incorporating time-lagged measurements, such that consecutive datapoints are combined into a single model input. This enables DPCA to capture temporal correlations that are intrinsic to batch processes. In addition, this study considers a gradient-space PCA approach, motivated by the observation that rates of change in process variables can be more informative than raw concentration measurements for identifying abnormal process behavior.

Furthermore, a moving-window Dynamic Functional

Principal Component Analysis (DFPCA) framework is also investigated. DFPCA employs functional representations, that can be nonlinear, to model batch trajectories, allowing intricate process dynamics to be captured while simultaneously smoothing noisy measurements. These capabilities are particularly advantageous in batch environments where measurement noise and nonlinearity can obscure early fault signatures.

Finally, autoencoder-based methods are explored due to their ability to model nonlinear relationships that cannot be captured by linear PCA-based techniques. Autoencoders are Artificial Neural Networks (ANNs) that reconstruct input data after compression through a lower-dimensional latent space. In this study, a dynamic (or lagged) autoencoder variant is considered, where consecutive time points are treated as a single input, analogous to DPCA. As with DPCA and DFPCA, the number of incorporated lags can be adjusted to control the extent to which process dynamics are embedded within the latent representation.

The primary objective of this study is to detect faulty batches as early as possible – ideally before deviations become apparent in the raw process data – while maintaining a clear distinction between normal and faulty operation to ensure confidence in fault detection outcomes. The proposed methodologies are therefore evaluated based on their ability to quickly and reliably distinguish faulty batches from normal operation in real time. The remainder of this paper is organized as follows: the *Methodology* section describes the data structure, the calculation of fault detection metrics, and implementation details of each faut detection strategy; the *Results and discussion* section compares the performance of the different fault detection strategies; and overall *Conclusions* are provided in the final section.

METHODOLOGY

Data availability

The thirteen batches recorded in this study span at least 3000 minutes each with concentration measurements of eight components taken every 2.1 minutes using a mid-infrared spectrometer. Three faulty batches have previously been identified; the remaining ten non-faulty batches are partitioned into eight randomly selected training batches and two validation batches. Faulty batches visibly deviate from normal operation around the 1000-minute mark, which this study aims to improve upon using the various process monitoring strategies. Data is passed through a short-span median filter and random-walk Kalman filter, which is well suited to processes with frequent measurements and slow dynamics [5], to smooth data while still being suitable to online application.

Fault detection metrics

Fault detection can be conducted by monitoring two different metrics over a moving window. If we consider the process data $\mathbf{X}$ to be reconstructed as $\widetilde{\mathbf{X}}$ with residual error $\mathbf{E}$ as in Equation (1):

$$\mathbf{X} = \widetilde{\mathbf{X}} + \mathbf{E}, \tag{1}$$

then through a dimensionality reduction framework, such as PCA, we can calculate both latent space and reconstruction space monitoring statistics. The latent space Hotelling's T^2 statistic and reconstruction space Q-residual can be calculated as in Equations (2) and (3):

$$T^2 = \mathbf{z}^\top \mathbf{\Lambda} \mathbf{z}, \tag{2}$$

$$Q = \|\mathbf{e}\|_2^2, \tag{3}$$

where $\mathbf{z}$ is the latent score vector, $\mathbf{\Lambda}$ contains latent variances, and $\mathbf{e}$ is the residual vector.

In this study, we consider a 50-input moving window for fault detection as opposed to single points to improve the robustness of fault detection. This study primarily focuses on the Hotelling's T^2 metric due to its superior overall performance on this case study.

Fault detection strategies

In this section, all fault detection strategies are applied to the dataset which is standardised such that the training dataset has a mean of zero and a standard deviation of one, with the exception of gradient-space PCA which is standardized after variable gradients are approximated.

Principal Component Analysis

Under PCA, the reconstruction of the standardised data matrix $\mathbf{X}$ shown in Equation (1) is therefore written as in Equation (4):

$$\mathbf{X} = \mathbf{Z}\mathbf{P}^\top + \mathbf{E}, \tag{4}$$

where $\mathbf{Z}$ is the matrix of scores in the latent space that correspond to $\mathbf{X}$, and $\mathbf{P}$ is the loading matrix. In this work, Singular Value Decomposition (SVD) is utilized to compute the loading matrix due to its efficiency and computational stability [6].

To keep only essential linear relationships and avoid capturing noise, the Principal Components (PCs) explaining the least variance are removed. The reconstructed dataset using the first k PCs is shown in Equation (5):

$$\widetilde{\mathbf{X}} = \sum_{i=1}^{k} \mathbf{Z}_i \mathbf{P}_i^\top. \tag{5}$$

Although PCA is a widely applied method for process monitoring and fault detection, it is limited by its inability to capture nonlinear and temporal relationships, preventing PCA from incorporating dynamic and complex relationships that can dominate batch processes.

Dynamic Principal Component Analysis

One way in which PCA can be improved is to incorporate past information, not just current measurements. Dynamic Principal Component Analysis (DPCA) is similar to PCA; the only difference is that the input data for DPCA is not a single vector of current variable values, it is instead a series of consecutive vectors of the variables up to the current measurement. The number of consecutive measurements included in each input is referred to as the number of 'lags'. In this study, we use a 100-lag DPCA approach that allows the DPCA model to capture temporal relationships that can be more informative than raw measurements in batch processes. The PCA dataset $\mathbf{X}$ with N measurements of each process variable can be augmented into a DPCA dataset $\mathbf{W}$ with L lags as in Equations (6) and (7):

$$\mathbf{w}(n) = \left[\mathbf{x}(n)^{\top}, \mathbf{x}(n-1)^{\top}, \dots, \mathbf{x}(n-L+1)^{\top}\right]^{\top}, \tag{6}$$

$$\mathbf{W} = \begin{bmatrix} \mathbf{w}(L-1)^{\top} \\ \mathbf{w}(L)^{\top} \\ \vdots \\ \mathbf{w}(N-1)^{\top} \end{bmatrix}. \tag{7}$$

Gradient-Space Principal Component Analysis

Another way in which temporal information can be incorporated is using gradient-space PCA, which is conducted on the derivatives of process variables, as opposed to their actual values. Since batch processes often show first order dynamics, for example in chemical, fermentation, and some crystallization systems, the gradient-space can be more informative than the absolute value of process variables.

The main challenge with gradient-space PCA is the fact that insufficiently smoothed data will see noise amplification in the gradient-space [7]. Therefore, a Locally Weighted Scatterplot Smoothing (LOWESS) filter is used so that smoothness can be tuned by altering the filter span [8]. However, the LOWESS filter span causes a delay between when a measurement is taken and when the corresponding smoothed value is returned. Therefore, gradient-space PCA can only be implemented if the delay is sufficiently small. In this study, measurements are taken frequently, relative to the batch time horizon, and a relatively small span is required, thus facilitating the application of gradient-space PCA.

After data smoothing, the gradient of each process variable $\dot{x}_j^{(v)}(t)$ is approximated as in Equation (8):

$$\dot{x}_j^{(v)}(t) = \frac{x_j^{(v)}(t) - x_j^{(v)}(t-\Delta t)}{\Delta t}, \tag{8}$$

where $x_j^{(v)}(t)$ is the measurement of variable v in batch j at time t, and Δt is the time between measurements. Thereafter, the gradient-space data is standardised to allow PCA to be conducted.

Functional Principal Component Analysis

Functional Principal Component Analysis (FPCA) considers each process trajectory as a continuous function over time, rather than a vector of measurements. By modeling the smoothed temporal evolution of the profile, FPCA can better capture nonlinear dynamic patterns. For the multivariate instance where each batch trajectory consists of V functional variables $x_j^{(v)}(t), v = 1, \dots, V$, the variation profile of each process variable is described by the weighted summation of B_v basis functions $\phi_b^{(v)}(t)$ as in Equation (9):

$$\tilde{x}_j^{(v)}(t) = \mu^{(v)}(t) + \sum_{b=1}^{B_v} c_{jb}^{(v)} \phi_b^{(v)}(t). \tag{9}$$

Each observation j collects the vector of coefficients as $\mathbf{c}_j = \left(c_{j1}^{(1)}, \dots, c_{jB_1}^{(1)}, c_{j1}^{(2)}, \dots c_{jB_2}^{(2)}, \dots, c_{j1}^{(V)}, \dots, c_{jB_V}^{(V)}\right)$ which are collected in the matrix $\mathbf{C}$. PCA is applied to the coefficient matrix to yield the loading matrix $\mathbf{P}$ and the matrix of shared multivariate FPCA scores, as in Equation (10):

$$\mathbf{C} = \mathbf{Z}\mathbf{P}^{\top} + \mathbf{E}. \tag{10}$$

The eigenfunctions $\psi_i^{(v)}(t)$ of variable v are obtained by expressing the appropriate loading block in its basis as in Equation (11):

$$\psi_i^{(v)}(t) = \sum_{b=1}^{B_v} P_{bi}^{(v)} \phi_b^{(v)}(t), \tag{11}$$

where $P_{bi}^{(v)}$ represents the loading rows of $\mathbf{P}$ associated with variable v. Multivariate FPCA scores Z_{ji} therefore satisfy Equation (12):

$$Z_{ji} = \sum_{v=1}^{V} \sum_{b=1}^{B_v} c_{jb}^{(v)} P_{bi}^{(v)}, \tag{12}$$

which demonstrates that a single set of scores summarizes the joint variation across all variables. Similarly to PCA, the coefficients that explain the least variance are removed, meaning only k scores remain, as in $\mathbf{z}_j = (Z_{j1}, Z_{j2}, \dots, Z_{jk})$, giving the approximate reconstruction of the coefficients $\tilde{\mathbf{c}}_j$ as in Equation (13):

$$\tilde{\mathbf{c}}_j = \sum_{i=1}^{k} Z_{ji} P_{\bullet i}. \tag{13}$$

Ultimately, variable profiles are reconstructed using the multivariate scores and eigenfunctions as in Equation (14):

$$\tilde{x}_j^{(v)}(t) = \mu^{(v)}(t) + \sum_{i=1}^{k} Z_{ji} \psi_i^{(v)}(t), \tag{14}$$

where the scores Z_{ji} are shared across all variables, unlike the coefficients $c_{jb}^{(v)}$ in Equation (11).

Within the FPCA framework described in Equations (9-14) there exist several options, namely different basis function types as well as the period over which the basis functions are fit. This study considers B-spline basis functions due to their flexibility potentially being able to capture complex batch dynamics. However, the

additional flexibility can capture noise, and so this study also considers an orthonormal polynomial approach as such functions can provide a smoother fit. Since this study focuses on thirteen batches, fitting the basis functions to the entire profile would yield just thirteen observations per variable, which is impractical for trying to identify statistically significant faults. Therefore, this study considers a dynamic (moving-window) FPCA (DFPCA) approach where basis functions are fit to the 100-lag segments from the DPCA data $\mathbf{W}$ in Equation (7).

Autoencoders

Autoencoders are a type of ANN that learn compact representations of higher-dimensional data through data reconstruction. Autoencoders can capture nonlinear relationships, meaning they can potentially uncover the more complex relationships in batch processes that standard PCA cannot. As illustrated in Figure 1, an autoencoder is composed of an encoder that compresses the original data $\mathbf{x}$ onto the lower-dimensional latent space $\mathbf{z}$, which is passed into the decoder that then reconstructs the original input data $\tilde{\mathbf{x}}$.

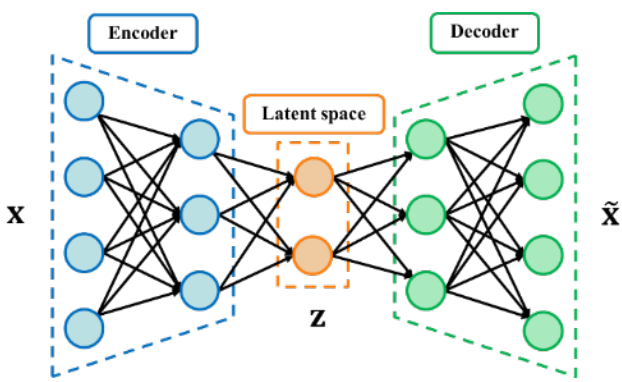

Figure 1. Example of an autoencoder architecture with 4-dimensional input, a 3-dimensional hidden layer in the encoder, and a 2-dimensional latent space. The decoder architecture always reflects that of the encoder, but the decder has its own unique weights and biases.

The autoencoder can be described by Equations (15a) and (15b):

$$\mathbf{z} = f_{\text{enc}}(\mathbf{x}), \tag{15a}$$

$$\tilde{\mathbf{x}} = f_{\text{dec}}(\mathbf{z}), \tag{15b}$$

where $f_{\text{enc}}(\cdot)$ is the function representing the encoder and $f_{\text{dec}}(\cdot)$ is the function representing the decoder. Autoencoders can also be trained on the 100-lag segments from the DPCA data $\mathbf{W}$ in Equation (7), so this study also considers a 100-lag autoencoder to capture temporal relationships, not just an autoencoder for modelling $\mathbf{X}$.

Autoencoder training focuses on minimizing the reconstruction loss described by the Mean Squared Error (MSE) loss function over J observations in Equation (16):

$$MSE = \frac{1}{J}\sum_{j=1}^{J}\left\|\mathbf{x}_j - \tilde{\mathbf{x}}_j\right\|_2^2. \tag{16}$$

The number of epochs (times passed through the training dataset during training) can be increased to improve fitting to the training dataset. However, using a higher number of epochs is both more computationally expensive and can lead to overfitting of the training data and poor fitting of the validation dataset. Another way in which the validation performance of the autoencoder can be improved is to introduce a penalization term into the loss function that penalizes the size of the autoencoder weights and biases during training, thus shrinking redundant parameters. In this study, 250 epochs are used alongside L2 regularization (sum of squared weights and biases). Finally, different autoencoder architectures are tested to identify the best performer on validation data.

It is also worth noting that since the autoencoder latent variables are not guaranteed to be orthonormal or centred like PCA scores, the Hotelling-type T^2 fault detection metric is approximated using a slightly different formula shown in Equation (17):

$$T_{AE}^2 = (\mathbf{z} - \bar{\mathbf{z}})^{\top}\boldsymbol{\Sigma}_{\mathbf{z}}^{-1}(\mathbf{z} - \bar{\mathbf{z}}), \tag{17}$$

where $\boldsymbol{\Sigma}_{\mathbf{z}}$ is the latent covariance matrix, and $\bar{\mathbf{z}}$ is the latent mean vector.

RESULTS AND DISCUSSION

Principal Component Analysis performance

As mentioned in the *Fault detection metrics* section, we consider a 50-input moving window for fault detection. In this window we monitor the percent of datapoints that exceed the 99% confidence limit of the T^2 statistic for both normal (training and validation) and faulty batches. The earliest window that demonstrates the clearest distinction between normal and faulty batches is chosen as the best fault detection window, the results of which are illustrated in Figures 2-4.

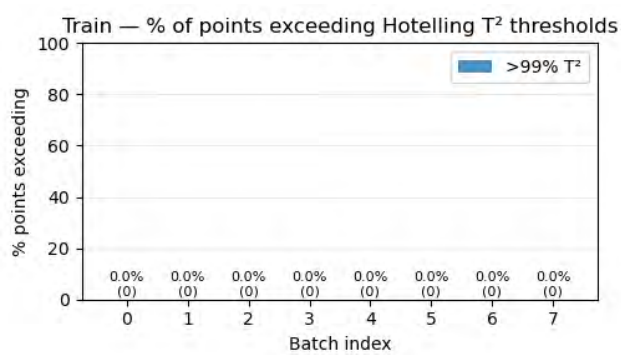

Figure 2. Bar chart showing percent of datapoints in each (normal operation) training batch where T^2 threshold is violated.

Looking at Figures 2 and 3, it is evident that the PCA does not flag any of the datapoints within this 50-datapoint window as exceeding the 99% confidence limit of the T^2 statistic, meaning PCA can correctly classify both the training and validation batches as normal operation.

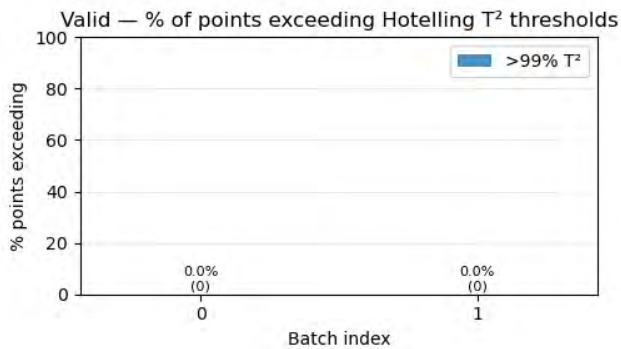

Figure 3. Bar chart showing percent of datapoints in each (normal operation) validation batch where T^2 threshold is violated.

Moreover, Figure 4 shows that PCA uncovers several instances where data within the fault detection window exceeds the 99% confidence limit of the T^2 statistic, therefore providing clear distinction between the normal and faulty batches.

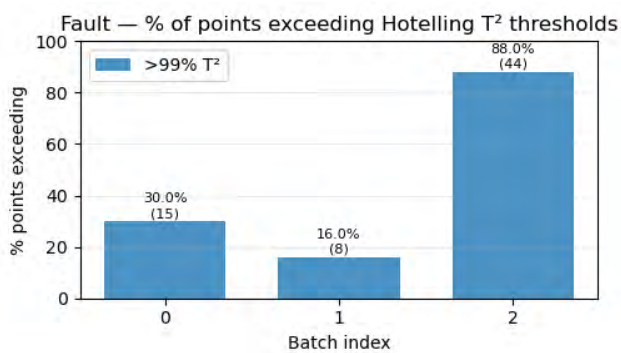

Figure 4. Bar chart showing percent of datapoints in each faulty test batch where T^2 threshold is violated.

The optimal fault detection window for the benchmark PCA requires data up to the 638th minute, a significant improvement upon observing faulty batches through raw data at around 1000th minute. If $\%FP$ is the highest percent of faults detected in any single normal batch and $\%TP$ is the lowest percent of faults detected in any single faulty batch, the fault detection margin $\%FDM$ is calculated as in Equation (18):

$$\%FDM = \%TP - \%FP. \tag{18}$$

The fault detection margin for the PCA benchmark is therefore 16% (since $\%FP = 0\%$, and $\%TP = 16\%$ for the benchmark PCA). The higher the fault detection margin, the more robustly the faulty batches can be distinguished from normal operation. Despite PCA reducing the time for faulty batch detection, there is still room for improvement in terms of both fault detection time and fault detection margin. This motivates the need to investigate more state-of-the-art process monitoring strategies.

Comparison of fault detection strategies performances with benchmark PCA

For the remaining fault detection strategies, performance is compared to the benchmark PCA in Table 1.

Table 1: Overall comparison between fault detection methods, showing online fault detection time and fault detection margin. The DFPCA abbreviations *Spln*. And *Poly*. stand for *B-spline* and *orthonormal polynomials*, respectively.

Model	Online fault detection time / min	Fault detection margin / %
PCA (*benchmark*)	638	16
100-lag DPCA	460	100
Gradient-space PCA	355	42
100-lag DFPCA (*Spln.*)	439	52
100-lag DFPCA (*Poly.*)	407	100
Autoencoder	344	0
100-lag Autoencoder	386	52

Looking at Table 1, all dynamic methods (those involving lagged data or operating in the gradient-space) outperform the benchmark PCA in terms of online fault detection time and margin. In fact, the only method that does not improve upon the benchmark PCA is the regular autoencoder. The stronger performance of gradient-space and lag-based strategies is likely because they better incorporate temporal relationships which can be more informative of the true state of a batch process.

It can also be seen that the two best performing approaches, in terms of fault detection margin, are 100-lag DPCA and 100-lag orthonormal polynomial DFPCA. Although the lag-based models use a substantially larger input space, due to the inclusion of time-lagged variables, it can achieve stronger dimensionality reduction than standard PCA, yielding a latent space that more effectively captures the system dynamics. For example, to explain 95% of the data variance, the benchmark PCA has an input dimension of 8 and latent space dimension of 6 (25% reduction), whereas the 100-lag DPCA has an input dimension of 800, and a latent dimension of just 18 (97.75% reduction). The inability for the latent space to capture system dynamics for lower-dimension inputs could be the reason that the standard autoencoder fails to distinguish faulty batches from normal operation, meanwhile the higher-dimensional 100-lag autoencoder

performed strongly with a 250 minute reduction in the fault detection time and 36-point reduction to the fault detection margin with respect to benchmark PCA.

The 100-lag DFPCA performs stronger when using orthonormal polynomials than B-splines, with a fault detection time reduction of over 30 minutes and a 48-point improvement in fault detection margin. Evidently, the use of smoother functions better describes the system which implies that the more flexible B-splines are capturing excessive noise instead of additional process information.

Although the fastest fault detection strategy is gradient-space PCA, it performs weaker than all the lag-based fault detection strategies in terms of fault detection margin, which is likely due to the aforementioned noise amplification in the gradient-space. This indicates that more smoothing is required to reduce the noise in the original data, however increased smoothing will both require a larger LOWESS filter span, thus raising the fault detection time, while also inadvertently removing some useful process dynamics.

Overall, the best performing fault detection strategy for this industrial batch process is found to be that using the 100-lag orthonormal polynomial DFPCA approach due to its excellent fault detection margin of 100% and a competitively low fault detection time of 407 minutes.

CONCLUSIONS

As batch processing continues to play a vital role in modern manufacturing and sustainable production, ensuring consistent and normal operation is becoming increasingly important. Through the use and comparison of these machine learning techniques, faulty batch detection can be achieved within the first 400 minutes; a significant improvement upon the 1000 minutes taken when observing the raw data and 638 minutes when utilizing standard PCA. In addition, a more statistically significant distinction can be found between faulty and non-faulty batches, meaning that fault detection robustness has been notably improved in comparison with the benchmark PCA. This study also demonstrates the benefits of utilizing fault detection strategies that exploit lag-based and gradient-space foundations for fault detection in batch processes. This study therefore paves the way for fault detection in industry-scale batch processes.

ACKNOWLEDGEMENTS

The authors thank BBSRC for providing funding for this project, along with KEIT Industrial Analytics who provided data and helped oversee the project.

AUTHOR IDENTIFIERS

Author ORCIDs:
Pennington O: 0009-0009-8119-871X
Zhang D: 0000-0001-5956-4618

REFERENCES

1. Gunther JC, Seborg DE, Baclaski J. Fault detection and diagnosis in industrial fed-batch fermentation. 2006 American Control Conference :6 pp. (2006). https://doi.org/10.1109/acc.2006.1657601
2. Zhang H, Lennox B. Fault detection and isolation in a fed-batch penicillin fermentation process. 2003 European Control Conference (ECC) :1670-1675 (2003). https://doi.org/10.23919/ecc.2003.7085204
3. Bicciato S, Bagno A, Soldà M, Manfredini R, Di Bello C. Fermentation diagnosis by multivariate statistical analysis. ABAB 102-103:049-062 (2002). https://doi.org/10.1385/abab:102-103:1-6:049
4. Tates AA, Louwerse DJ, Smilde AK, Koot GLM, Berndt H. Monitoring a PVC batch process with multivariate statistical process control charts. Ind. Eng. Chem. Res. 38:4769-4776 (1999). https://doi.org/10.1021/ie9901067
5. Li JY, Ambikasaran S, Darve EF, Kitanidis PK. A kalman filter powered by h2-matrices for quasi-continuous data assimilation problems. Water Resour. Res. 50:3734-3749 (2014). https://doi.org/10.1002/2013wr014607
6. Bakarji J. Singular Value Decomposition (SVD) and Its Relation to PCA. Riad El-Solh, Beirut 1107, Lebanon; 2025. Online lecture notes. (2020) Available from: https://www.ml4science.com/content/notes/05_svd/svd_demo
7. Chartrand R. Numerical differentiation of noisy, nonsmooth data. ISRN Applied Mathematics 2011:1-11 (2011). https://doi.org/10.5402/2011/164564
8. Cleveland WS. Robust locally weighted regression and smoothing scatterplots. Journal of the American Statistical Association 74:829-836 (2012). https://doi.org/10.1080/01621459.1979.10481038

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Enhancing plasma etching efficiency via physics-based modeling and machine learning

Eneri Boniakou[a], Yao Xue[b], Tzannis Vasileiadis[c], Sotiris Mouchtouris[ad], Katerina Oikonomou[e], Chloi Zormpa[d], Antonios Armaou[cf*], Vassilios Constantoudis[d], Evangelos Gogolides[d], George Kokkoris[a]

[a] National Technical University of Athens, Department of Chemical Engineering, Athens 15780, Greece
[b] The Pennsylvania State University, Department of Mechanical Engineering, University Park, PA 16802, USA
[c] University of Patras, Department of Chemical Engineering, Rio, 26504, Greece
[d] National Centre for Scientific Research Demokritos, Institute of Nanoscience and Nanotechnology, Athens 15341, Greece
[e] National Centre for Scientific Research Demokritos, Institute of Informatics and Telecommunications, Athens 15341, Greece
[f] The Pennsylvania State University, Department of Chemical Engineering, University Park, PA 16802, USA
* Corresponding Author: armaou@psu.edu; armaou@upatras.gr.

ABSTRACT

Modern semiconductor manufacturing requires extreme precision as yield margins narrow in the "More-than-Moore" era. While physics-based models (PBMs) provide high-fidelity insights into plasma etching, their computational intensity—often requiring hours per simulation—renders them impractical for direct iterative optimization. This work demonstrates a hybrid framework that utilizes data-driven surrogate models to enable rapid, cost-effective process optimization. A 2D axisymmetric fluid model of an inductively coupled O_2 plasma (ICP) reactor was developed to generate a training dataset for two neural architectures: a Multi-Layer Perceptron (MLP) and a Kolmogorov-Arnold Network (KAN). These surrogates predict radial etching rates across a wide operating window of power, pressure, gas flow, and bias voltage. By replacing the expensive PBM with these high-speed surrogates, derivative-free optimization algorithms (Nelder-Mead and Powell) successfully identified a profit-maximizing operating point (2000 W, 10 mTorr) orders of magnitude faster than direct physical simulation. The results confirm that surrogate-based optimization effectively captures dominant physical trends, such as ion-flux limited regimes, while providing a "Confidence Gap" through model disagreement to flag epistemic uncertainty. This methodology offers a scalable blueprint for reducing the computational burden of process design, transitioning from expensive trial-and-error to efficient, physics-validated autonomous discovery.

Keywords: Modelling and Simulations, Machine Learning, Optimization, Industry 4.0, Plasma process

INTRODUCTION

The semiconductor industry is currently navigating a “More-than-Moore” era. Until recently, critical dimensions kept shrinking until they reached a scale below the 5 nm node, where the economic and technical constraints on device fabrication have intensified. Yield margins are becoming increasingly narrow, requiring strict process control and optimization. Thus, the ability to precisely model and optimize fabrication steps is no longer merely advantageous but essential for maintaining economic viability in high-volume manufacturing [1].

Low pressure plasma, an electrically neutral ionized gas, is formed in a low pressure environment (typically 1-100 mTorr) by applying voltage between two electrodes or through coil excitation. Characterized by complex physical and chemical interactions, these discharges enable unique surface modifications unmatched by other commercial methods [2]. Among the various configurations, inductively coupled plasmas (ICPs) are of particular importance in semiconductor manufacturing [3].

Plasma etching is an essential dry fabrication process in modern integrated circuit production. It uses reactive species and energetic ions to selectively remove material from surfaces, enabling highly anisotropic and controllable pattern transfer from lithographic masks [2]. In an ICP reactor, an oscillating magnetic field is generated by a coil which in turn induces a strong electric field

capable of sustaining high plasma densities. Furthermore, wafer biasing controls the energy of ions reaching the wafer, enabling high etching rates and selectivity [4].

Numerical simulations of these reactors have become an important tool [5], which enables an improved understanding of the process and allows replacing in-situ measurements which are expensive or even impossible [6]. However, accurate predictions of spatially varying etching rates are challenging due to the complex interplay of the governing physics. This challenge is compounded by the significantly nonlinear nature of the problem, which elevates the complexity of optimization process requiring multiple iterations of the computationally expensive physics-based model (PBM).

To address these challenges, the industry is increasingly looking toward "virtual metrology" and digital twin technologies [1]. Currently, two distinct modeling paradigms exist. PBMs offer high fidelity and interpretability but are computationally expensive, often requiring hours to simulate a single operating point. Data-driven Machine Learning (ML) model approaches offer rapid inference speeds suitable for real-time control but typically suffer from poor generalizability outside their training data, a lack of physical interpretability and reliability which is directly correlated to the accuracy and the volume of the training data.

PHYSICS BASED MODEL

Fluid model plasma simulations have been proved to give valid approaches of plasma processes with typical ICP chambers of pressures in the order of $10\ mTorr$ [6].

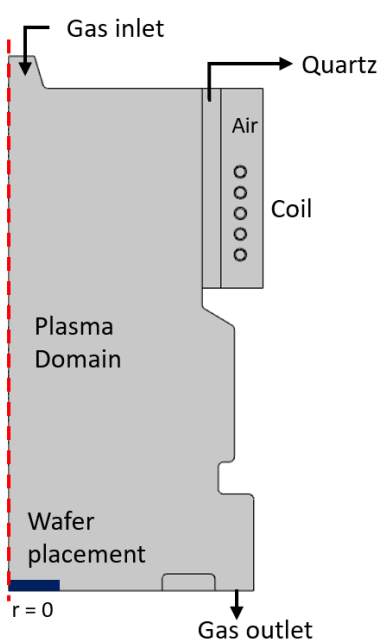

Figure 1. Computational domain of the *COBRA* reactor showing its major parts.

In this paper, a fluid model is developed, for the simulation of the plasma gas phase at low pressure conditions, focusing on etching PMMA blanket samples by O_2 gas within Oxfords Instruments' *PlasmaPro 100 Cobra* ICP plasma reactor. In this system, a 5-turn coil, placed behind a dielectric window, driven by a 13.56 MHz RF generator, couples energy electromagnetically into the plasma.

In the fluid model, species densities, electron density, and electron energy evolve in response to key operating conditions: coil power, pressure, gas feed rate, and wafer bias voltage. Given the reactor's axial symmetry, the problem is modeled in 2D axisymmetric coordinates. Figure 1 illustrates the computational domain representing the reactor geometry.

Briefly, for the electromagnetic field, Maxwell's equation set is solved. Ampere's law [eqs. (1)-(2)] is used to calculate the magnetic vector potential in the plasma region, in the dielectric window, in the coil domain and in the air contained in the domain covering the antenna

$$\nabla \times \mathbf{B} = \mu_0 \mu_r \mathbf{J_{tot}} \tag{1}$$

$$\mathbf{B} = \nabla \times \mathbf{A} \tag{2}$$

where $\mathbf{A}$ is the magnetic vector potential, μ_0 and μ_r are the permeability of free space and the domain material (quartz or air) respectively, and $\boldsymbol{B}$ is the magnetic flux density. $\mathbf{J_{tot}}$ is the vector of the total current density in each domain. For the electric potential, Poisson's equation [eq. (3)] is utilized

$$-\nabla^2 V = \frac{\rho}{\epsilon_0 \epsilon_r} \tag{3}$$

where ϵ_0 and ϵ_r is the permittivity of free space and material respectively, and ρ is the charge density in the plasma domain.

The electrical conductivity of each domain is a key model parameter. Known values are used for air and quartz, while copper's conductivity is assigned to the coil. For the plasma domain, the conductivity is updated dynamically at each time step according to eq. (4)

$$\sigma_{plasma} = \frac{N_e q^2}{m_e(\upsilon + j\omega)} \tag{4}$$

where N_e refers to the electron density, q is its charge, j is the unit imaginary number and ω stands for the angular frequency of the antenna. The total collision frequency υ is computed by summing all electron impact reaction rates and dividing by the electron density.

Regarding the electron impact reactions, their rate constants are calculated by integrating the product of the electron energy distribution function (EEDF) and the reaction-specific cross-section over the electron energy spectrum as given by eq. (5). For the O_2 plasma, the EEDF is approximated as Maxwellian, consistent with ref.

[7], thus

$$k_{e,i} = \int_0^\infty \sqrt{\frac{2q}{m_e}} \epsilon \sigma_i f(\epsilon) d\epsilon \quad (5)$$

where ϵ is the electron energy, σ_i is the i-th electron impact reaction cross section and f is the EEDF.

The mass balances for heavy species (neutral or ionic) are also solved in the model as shown, for a species k, in eq. (6)

$$\frac{\partial(\rho w_k)}{\partial t} = R_k - \nabla \cdot \mathbf{\Gamma}_k - \nabla \cdot (\rho \boldsymbol{u} w_k) \quad (6)$$

where ρ, w_k, R_k, and $\mathbf{\Gamma}_k$ are the gas mixture density, the mass fraction of species k, the production/consumption rate due to reactions in the gas phase, and the flux of k respectively. $\boldsymbol{u}$ is the fluid velocity. The flux is given from the drift-diffusion approximation [5] and the velocity is computed from the continuity and momentum conservation equations. The gas temperature is calculated by solving the gas energy balance, which accounts for heating contributions from electron–neutral collisions, thermochemical reactions, Joule heating, and the Frank–Condon mechanism, based on ref. [8]. The PBM is implemented and numerically solved using COMSOL Multiphysics.

The set of reactions employed for the O_2 plasma is based on ref. [7]. A comprehensive description of the underlying reaction mechanisms is omitted for brevity. Briefly, it includes elastic collisions of O_2 and O with electrons, electron impact ionization, O_2 dissociation, and excitation reactions. At the walls, ions are neutralized and excited states return to their corresponding ground state. Electrically, the inner walls are set to ground potential, while the inlet and outlet boundaries maintain zero net current. A no-slip condition is applied to all wall surfaces, and the inlet and outlet are modeled with a fully developed velocity profile.

At the wafer level, the PMMA etching surface model follows the framework of ref. [9]. The dominant mechanism is ion-enhanced etching, driven by the interaction of adsorbed neutral oxygen with ions. Under typical plasma etching conditions, ion enhanced etching dominates over physical sputtering and pure chemical etching. In this mechanism, the oxidation of carbon in the PMMA polymer acts as the rate-determining step, while the desorption of volatile reaction products is rapid. Consequently, the overall etching rate depends on the fluxes of both ions and neutrals, as well as the ion energy.

SURROGATE MODELING

Multi-Layer Perceptron

The trained Multi-Layer Perceptron (MLP) has as input the 4 operating conditions and predicts 10 etching rates across the wafer radius.

Table 1: Process operation parameters' range

Parameter	Lower Bound	Higher Bound
ICP Power (W)	500	2000
Pressure (mTorr)	10	100
Feed (sccm)	40	150
Voltage Bias (V)	0	700

The PBM was run iteratively 2000 times under varying operating conditions to generate an ensemble of data for model training. The simulation inputs were sampled from the operation windows presented in Table 1. The solution of the PBM for a single set of operating conditions requires ~1h (Threadripper 3960X, 264GB RAM). Following each simulation, 10 radial etching rate values along the wafer are exported. The dataset was partitioned into training (60%), validation (20%), and testing (20%) sets. To rigorously assess model generalization, the test set deliberately occupied a distinct subspace comprising edge cases of the parameter space. This ensured the model was evaluated on challenging, unseen conditions rather than interpolated data points.

To ensure stable and efficient network training, all input and output variables were min-max normalized to a $[0, 1]$ range [10]. This scaling prevents features with wider numerical ranges (e.g., etching rate) from dominating the learning process and helps the optimizer converge faster by mitigating gradient instability. The implemented model was a fully connected MLP with a rather simple architecture of an input layer with 4 neurons, two hidden layers with 20 neurons each, and an output layer with 10 neurons. The hidden layers utilized the Rectified Linear Unit (ReLU) activation function to introduce non-linearity while maintaining stable gradient flow during backpropagation. Experimental fine-tuning indicated that ReLU hidden layers with an Exponential Linear Unit (ELU) [11] activated output layer, had the best performance with a smooth, continuous regression at the output.

Weighted mean squared error (MSE) was implemented with L1 regularization which is ideal for regression tasks as a loss function [12]. The L1 penalty term, controlled by the hyperparameter λ, was applied to network weights to improve generalization. During training λ was set to 10^{-3}. The function is formalized as

$$L = \frac{1}{N} \sum_{m=1}^{N} \sum_{j=1}^{10} \frac{1}{10} \left(y_{m,j} - \hat{y}_{m,j}\right)^2 + \lambda |W_k| \quad (7)$$

where 10 is the set of output etching rate points, N the batch size, W_k refers to the model weights and $y_{m,j}, \hat{y}_{m,j}$ are the true and predicted values for the sample m and radial position j.

Training was performed using the Adam optimizer [13], with an initial learning rate of $3 \cdot 10^{-3}$ and a weight decay of $2 \cdot 10^{-5}$. A scheduler was configured to monitor the validation loss and reduce the learning rate by a factor of 10 if no improvement was observed over 5 epochs.

To prevent overfitting, early stopping was employed based on validation loss, to terminate the training if no improvement was observed, reverting the model to the state with the best validation performance. The batch size of training was set to 32. The initial learning rate, the weight decay and the batch size were optimized via grid search based on minimizing the loss. Figure 2 shows the MLP model accuracy with high R^2 (0.97). Points cluster tightly along the ideal line, confirming the model captures process physics well. Minor divergence at data edges is expected due to sparse training in those regions and greater uncertainty in the PBM itself. The training and inference times are quite low; 10 min and < 10^{-3} s respectively (Threadripper 3960X, 264GB RAM).

Table 2: Comparison of KAN and MLP architectures

Aspect	MLP	KAN
Activation Location	On nodes	On edges
Activation Type	Fixed (ReLU, etc.)	Learnable (B-splines)
Nonlinearity source	Stacked layers	Flexible edge functions
Smoothness	Not guaranteed	Inherently smooth
Parameter efficiency	Lower	Higher
Interpretability	Limited	Improved

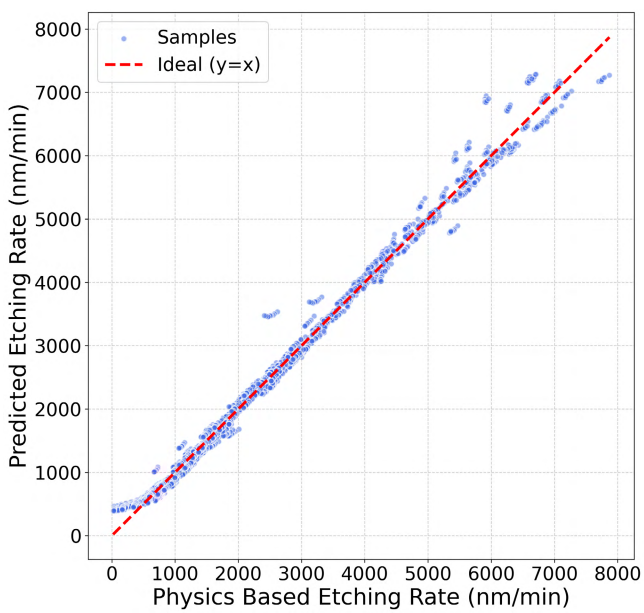

Figure 2. Scatter plot of the MLP model (4 inputs, 10 outputs). Model scores: $R^2 = 0.97$ and a MSE of 0.002.

Kolmogorov-Arnold Network

We also employed Kolmogorov-Arnold Networks (KANs) [14] as our base architecture for etch rate prediction. This model's training time was 5 minutes on a GeForce Titan RTX, requiring about 1GB of GRAM.

Unlike traditional MLPs that apply fixed activation functions (e.g., ReLU, sigmoid) on nodes, KANs place learnable activation functions on edges. Each edge activation is parameterized by B-spline basis functions

$$\phi(x) = \sum_{j=1}^{G+k} c_j B_j^k(x) \tag{8}$$

where B_j^k are B-spline basis functions of order k defined on a grid with G intervals, and c_j are learnable coefficients. This allows KANs to adaptively learn activation shapes suited to specific input-output relationships, rather than relying on predefined fixed nonlinearities.

Advantages of B-Spline Parameterization

The B-spline parameterization offers several advantages for scientific modeling tasks, such as smoothness, local support and parameter efficiency. B-splines are inherently smooth functions, making them well-suited for approximating continuous physical quantities where abrupt discontinuities are not expected. Also, each basis function has compact support, enabling efficient local adaptation to input variations without causing global interference in other regions of the input space. Lastly, KANs can achieve comparable or superior accuracy with fewer parameters than MLPs, as the flexible learnable activation functions can capture complex nonlinear relationships more efficiently than stacking many layers with fixed activations.

Plasma-Encoded Formulation

To enable continuous spatial prediction, we adopt a position-encoded formulation

$$\hat{r}(p) = KAN([x, p]) \tag{9}$$

where the radial position $p \in [-1, 1]$ is concatenated with process parameters as network input, $x \in R^4$ denotes the same operating parameters as previously, and $\hat{r}(p)$ is the predicted etch rate at normalized radial position p.

This pointwise formulation transforms the problem from predicting a fixed-length output vector to learning a continuous function over spatial coordinates. Rather than training a model to output $[r_1, r_2, \ldots, r_{10}]$ simultaneously, we train it to predict the etching rate at any given position, allowing:

- **Arbitrary resolution:** Prediction can be made at any spatial location, not just the measurement points.
- **Implicit spatial coherence:** The continuous function representation naturally enforces smooth transitions between nearby positions.
- **Better generalization:** The network learns the underlying continuous mapping rather than memorizing discrete point correspondences.

The KAN's smooth B-spline activations further

promote smooth spatial variations in the output, which is consistent with physical expectations for plasma etching processes where etching rates vary continuously across the wafer surface.

OPTIMIZATION

Objective Function Formulation

The optimization problem is formulated to maximize the net economic value generated by the etching process per unit time (€/s). The objective function, $J(x, y)$, accounts for both the value of the processed wafer and the operational expenditures. It is defined as

$$J(x, y) = P_{Efficiency} - P_{Cost} \tag{10}$$

$$x = Power, Pressure, Feed, VBias \tag{11}$$

$$y = Er_{radial} \tag{12}$$

where the decision vector x includes the controllable operation parameters (Source Power, Pressure, Feed Flow, Bias Voltage) and y represents the predicted radial etching rate profile derived from the surrogate models.

The efficiency term $(P_{Efficiency})$ maps the physical output to monetary gain. It is defined as the product of the mean etching rate (Er_{mean}) normalised to the wafer area, the effective process yield (Y), and the economic value per unit of etched thickness (G_h):

$$P_{Efficiency} = Y \cdot Er_{mean} \cdot G_h \tag{13}$$

In this context, yield represents the fraction of wafer surface area that falls within strict uniformity tolerance bounds; while high etching rates increase throughput, the function degrades rapidly if the radial profile deviates from the wafer center maximum, ensuring speed is not prioritized at the expense of critical dimension uniformity.

The cost term (P_{Cost}) consists of the power, material and environmental harm costs. The power cost (C_{Power}) includes the RF power (source and bias) and facility baseloads. We observe that the total power cost variance is minimal across the parameter space, as the fixed power consumption of auxiliary systems (e.g., pumps, cleanroom infrastructure) outweighs the variable RF load.

$$P_{Cost} = C_{Power} + C_{Mat} + C_{Env} \tag{14}$$

The material cost (C_{Mat}) accounts for the consumption of process gases. While the optimization strictly controls the flow of the reactive precursor (Feed), the total material cost is largely driven by the steady-state consumption of inert carrier gases required for the reactor's operation. Similar to the power term, the relative variance of material costs remains low across the parameter bounds, despite constituting a significant fraction of the absolute operational expenditure. The environmental term (C_{Env}) calculates costs associated with the Global Warming Potential (GWP) of the effluent. In the studied O_2 plasma, GWP costs are negligible, though the framework is designed to scale for more complex chemistries where environmental abatement is a dominant cost driver.

The objective function can be readily augmented to include upstream lithography parameters (e.g., mask selectivity interactions) or auxiliary economic terms such as equipment maintenance cycles, thereby enabling a holistic optimization of the broader patterning module without altering the fundamental control logic.

Optimization Strategy

The optimization of the reactor setpoints is implemented within Python as a derivative-free optimization (DFO) problem. Given the non-convex nature of the neural network response surfaces and the discontinuities introduced by the yield penalty function (Y), gradient-based methods are prone to instability or premature convergence to local optima. Consequently, we validate our findings by executing the optimization task using two independent direct-search algorithms: Nelder-Mead (NM) and Powell's Conjugate Direction (PCD) methods [15, 16].

Rather than relying on a single solver, conducting separate, independent optimization searches with these distinct methods serves as a critical verification. The simplex-based approach of NM and the directional line-search approach of PCD operate via different mechanisms; their convergence to a common basin of attraction, despite their differing search trajectories, increases our confidence in the optimality of the identified solution and independence from algorithmic artifacts.

The optimization is constrained within a hyper-rectangle defined by the physical limits of the training data, presented in Table 1. To facilitate unbiased exploration, both solvers are initialized using a "Warm Start" strategy originating from the centroid of the hyperparameter space. This neutral initialization point prevents the optimizer from being prejudiced toward specific high-power or high-pressure regimes, allowing the objective function gradient to naturally drive the trajectory toward the most profitable operating window.

Optimization Results

The optimization results for both surrogate architectures are summarized in Tables 3 and 4. For each respective surrogate model, both the NM and PCD algorithms converged to the same operating point. Combined with the multistart method, this confirms that the identified solutions represent the true global maximums of the respective neural network response surfaces, independent of the optimization method employed.

Table 3: MLP Optimization Results

Optimal Conditions	J [=] €/s	10.55

Power (W)	2000	$P_{Efficiency}$	13.64
Pressure (mTorr)	10	C_{Power}	$3.15\ e^{-4}$
Feed (sccm)	40	C_{Mat}	3.09
Voltage Bias (V)	700	C_{Env}	$2.18\ e^{-8}$

Table 4: KAN Optimization Results

Optimal Conditions		$J\ [=]\ €/s$	**10.47**
Power (W)	2000	$P_{Efficiency}$	13.57
Pressure (mTorr)	10	C_{Power}	$3.14\ e^{-4}$
Feed (sccm)	105	C_{Mat}	3.09
Voltage Bias (V)	686	C_{Env}	$5.72\ e^{-8}$

Figure 3 illustrates the optimization convergence, highlighting the prohibitive cost of direct optimization. Since a PBM prediction takes about an hour, the numerous iterations required would be computationally impractical to run. In contrast, the surrogate models allowed the entire process to complete on a mid-tier CPU in approximately 3-4 minutes per method employed.

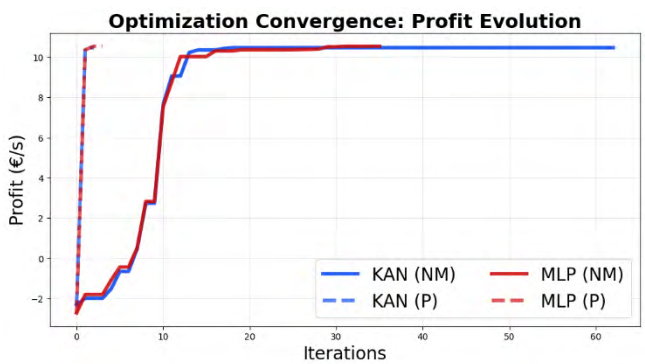

Figure 3. Optimization Convergence of each method

To validate the accuracy of these solutions, the optimal etching rate profiles predicted by both surrogates are compared against the PBM ground truth in Figure 4. The PBM reference profile was generated using the optimal conditions identified by the MLP surrogate. As illustrated, both the MLP and KAN surrogates yield radial profiles that closely align with the physical baseline, despite the divergence in their secondary Feed Flow parameters.

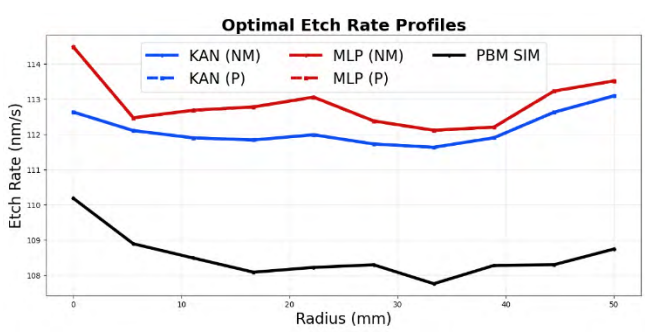

Figure 4. Etching rate profiles for both models and optimization methods, compared to PBM profile.

Sensitivity Analysis and Exploration

While the optimization algorithms converged effectively, the divergence in the final setpoints necessitates a closer examination of the surrogate response surfaces. To understand the reasons behind the different optimal conditions, local sensitivity sweeps are performed around the high-power/low-pressure $(2000\ W, 10\ mTorr)$ area where both models agreed.

As established in the optimization results, both surrogates identified Power and Pressure as the primary control levers. Sensitivity sweeps confirm that both the MLP and KAN correctly approximate the PBM gradients in these dimensions, identifying the upper bound of power and the lower bound of pressure as the trajectory for maximizing etching rate. This alignment confirms that both architectures successfully captured the global trend of the dominant variables.

A minor discrepancy was observed in the Bias Voltage, where the MLP model converged to $700\ V$ while the KAN settled at $686\ V$. Despite this offset, the behavior of both models remains largely consistent, reaffirming the parameter's critical influence on the objective function. To verify this, a sensitivity sweep was conducted by varying Voltage Bias across its design limits while holding the rest of the parameters at their optimal values, shown in Fig. 5. This isolation reveals the distinct response surfaces of each surrogate: while both models correctly and similarly identify the parameter importance, the KAN response exhibits a stagnation effect near the upper bound, whereas the MLP maintains a steeper gradient.

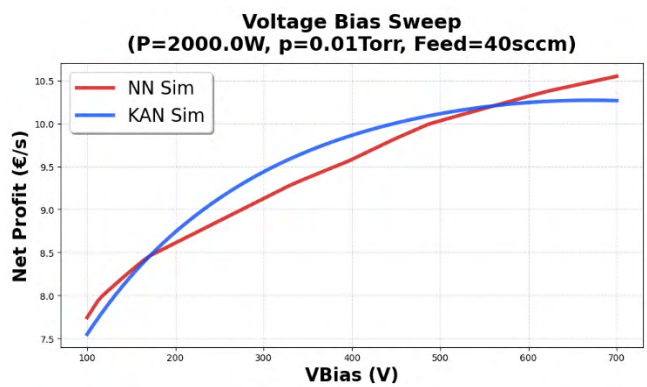

Figure 5. VBias response profiles: MLP(NN) and KAN

The most significant divergence occurred for the Feed parameter $(40\ sccm\ vs\ 105\ sccm)$. To analyze this, the Feed was swept across the design space while holding other inputs at their optimal values, shown in Figure 6; it reveals a distinct conflict in functional approximation. The Ground Truth for the surrogates, i.e., PBM, behaves near-linearly, showing a slight preference for lower flow rates. It also showcases the impact of the slight profile

difference between the PBM and the surrogates, as shown in Figure 4, on the objective function. The MLP surrogate is closer to mimicking the PBM linear behavior. It correctly identifies the negative slope and directs the optimizer toward the lower bound, effectively matching the true physical optimal point. The KAN surrogate introduces an artificial curvature. While its predictions are closer to the PBM at the boundaries, it significantly deviates in the center of the domain, creating a parabolic shape that drives the optimizer towards the middle.

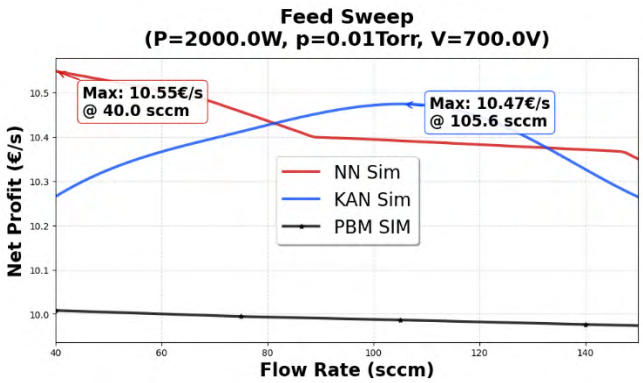

Figure 6. Feed Flow response profiles: MLP(NN) and KAN surrogates vs. PBM.

DISCUSSION

Data Sparsity

The divergence in the Feed Flow optimization highlights a fundamental difference in how the two architectures handle interpolation in sparse data regions. Despite both models being trained on the exact same dataset, the KAN architecture likely attempted to fit a higher-order non-linearity where none existed, whereas the MLP maintained a simpler linear approximation that proved more robust in this specific scenario.

This behavior could be attributed to the "Edge Effect." The optimal operating point $(2000\ W, 10\ mTorr)$ lies at the corner of the training hypercube. Neural networks are notoriously prone to high variance at the boundaries of their training domain. The KAN architecture, which utilizes learnable spline functions, is designed to be highly flexible in capturing local non-linearities. However, the results indicate that in this specific boundary region, this flexibility may have resulted in the model capturing stochastic noise rather than the true physical trend. On the other hand, the MLP appears to have defaulted to a smoother, more global approximation. In this instance, the simpler structure of the MLP likely acted as a constraint, preventing the model from inferring a local optimum where the physical landscape is effectively flat.

Physics Validation

From a plasma physics perspective, the PBM's insensitivity to Feed (and the MLP's resulting preference for the lower bound) is physically consistent with an ion-driven etching regime. At low pressures, the mean free path of species increases, and the etching mechanism becomes dominated by directional ion bombardment (controlled by Power and Bias) rather than chemical etching by neutral radicals (controlled by Feed Flow).

Consequently, variations in the gas flow have a negligible impact on the etching rate compared to the ion density. The PBM correctly reflects this "flat" landscape. The fact that the optimizer identified the correct high-power/low-pressure window proves that the framework successfully isolated the primary physics drivers, even if the secondary parameters had a diminished effect.

The Confidence Gap

The disagreement between the models in the Feed value raises a critical question regarding reliability. If a single model had been used, the optimizer would have returned a result with high apparent confidence, while it could still be misleading. The existence of this "Confidence Gap" between the two models is important. It suggests that in industrial applications, model disagreement should be treated as a proxy for epistemic uncertainty. This underscores the necessity for future frameworks to integrate Uncertainty Quantification. By flagging regions where the MLP and KAN diverge, the system could automatically trigger "active learning, " requesting targeted PBM simulations to resolve the ambiguity before finalizing the optimization, while also improving the models.

CONCLUSIONS

This work presents a framework for the optimization of plasma etching processes, aimed at bridging the fidelity of Physics-Based Models (PBMs) with the computational efficiency of data-driven surrogates. The core contribution of this study extends beyond the specific etching recipe identified; it demonstrates that integrating diverse neural architectures into the optimization loop provides a necessary proxy for prediction reliability, a critical step toward fully autonomous process discovery.

The application of this framework successfully identified a high-value operating regime characterized by high power $(2000\ W)$ and low pressure $(10\ mTorr)$. This result is physically consistent with an ion driven etching mechanism, where maximizing ion energy and density dominates value creation. Furthermore, the use of two distinct surrogate architectures (MLP and KAN) confirmed the robustness of these setpoints; both models agreed on the primary drivers, while disagreeing regarding the rest.

Beyond the specific framework, this study validates the feasibility and efficiency of surrogate-based optimization for complex industrial processes. We

demonstrated that by training neural networks on a targeted dataset of PBM simulations, it is possible to achieve global optimality orders of magnitude faster than direct physics simulation. This workflow offers a practical blueprint for reducing the computational burden of process design, allowing engineers to explore vast parameter spaces without the prohibitive time costs associated with traditional physics solvers or experimental trial-and-error.

Future work will transition from static modeling to a hybrid active learning loop. The PBM could be integrated directly into the optimization process, utilizing prediction uncertainty as a trigger. Instead of relying solely on pre-trained surrogates, the system will dynamically query the PBM whenever the certainty of a prediction drops. This creates a self-correcting workflow that autonomously improves model accuracy in high-value regions, ensuring that the final operating point is validated by physics while keeping computational costs minimal.

ACKNOWLEDGEMENTS

Support of the project "Enhancing plasma etching efficiency, repeatability, and environmental footprint via AI-based modeling and optimization (plasmaAI)" of the program "AI-Aware Pathways to Sustainable Semiconductor Process and Manufacturing Technologies", Intel Corporation & Merck KGaA, Darmstadt is gratefully acknowledged.

AUTHOR IDENTIFIERS

Author ORCIDs:
Boniakou E.: 0009-0002-7080-2837
Xue Y.: 0009-0001-1200-6157
Vasileiadis T.: 0009-0003-7122-5995
Mouchtouris S.: 0000-0001-7580-4554
Oikonomou K.: 0000-0001-7537-7310
Zormpa C.: 0009-0005-6797-9844
Armaou A.: 0000-0002-8592-7934
Kokkoris G.: 0000-0003-4507-7311
Gogolides E.: 0000-0002-1870-5629
Constantoudis V.: 0000-0003-3164-977X

REFERENCES

1. Oehrlein GS, Brandstadter SM, Bruce RL, Chang JP, DeMott JC, Donnelly VM, Dussart R, Fischer A, Gottscho RA, Hamaguchi S, Honda M, Hori M, Ishikawa K, Jaloviar SG, Kanarik KJ, Karahashi K, Ko A, Kothari H, Kuboi N, Kushner MJ, Lill T, Luan P, Mesbah A, Miller E, Nath S, Ohya Y, Omura M, Park C, Poulose J, Rauf S, Sekine M, Smith TG, Stafford N, Standaert T, Ventzek PLG. Future of plasma etching for microelectronics: challenges and opportunities. Journal of Vacuum Science & Technology B 42: (2024). https://doi.org/10.1116/6.0003579
2. Lieberman MA, Lichtenberg AJ. Principles of plasma discharges and materials processing. Wiley (2005). https://doi.org/10.1002/0471724254
3. Graves DB, Labelle CB, Kushner MJ, Aydil ES, Donnelly VM, Chang JP, Mayer P, Overzet L, Shannon S, Rauf S, Ruzic DN. Science challenges and research opportunities for plasma applications in microelectronics. Journal of Vacuum Science & Technology B 42: (2024). https://doi.org/10.1116/6.0003531
4. Tachi S, Tsujimoto K, Arai S, Kure T. Low-temperature dry etching. Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films 9:796-803 (1991). https://doi.org/10.1116/1.577364
5. Mouchtouris S, Kokkoris G. A hybrid model for low pressure inductively coupled plasmas combining a fluid model for electrons with a plasma-potential-dependent energy distribution and a fluid-monte carlo model for ions. Plasma Sources Sci. Technol. 25:025007 (2016). https://doi.org/10.1088/0963-0252/25/2/025007
6. Brezmes AO, Breitkopf C. Fast and reliable simulations of argon inductively coupled plasma using COMSOL. Vacuum 116:65-72 (2015). https://doi.org/10.1016/j.vacuum.2015.03.002
7. Gudmundsson JT, Thorsteinsson EG. Oxygen discharges diluted with argon: dissociation processes. Plasma Sources Sci Technol 16:399-412 (2007) https://doi.org/10.1088/0963-0252/16/2/022
8. Kiehlbauch MW, Graves DB. Inductively coupled plasmas in oxygen: modeling and experiment. Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films 21:660-670 (2003). https://doi.org/10.1116/1.1564024
9. Memos G, Lidorikis E, Gogolides E, Kokkoris G. A hybrid modeling framework for the investigation of surface roughening of polymers during oxygen plasma etching. J. Phys. D: Appl. Phys. 54:175205 (2021). https://doi.org/10.1088/1361-6463/abdb0b
10. Al Shalabi L, Shaaban Z. Normalization as a preprocessing engine for data mining and the approach of preference matrix. 2006 International Conference on Dependability of Computer Systems :207-214 (2006). https://doi.org/10.1109/depcos-relcomex.2006.38
11. Clevert DA, Unterthiner T, Hochreiter S. Fast and accurate deep network learning by exponential linear units (ELUs). arXiv preprint arXiv:1511.07289 (2015)
12. Wu S, Li G, Deng L, Liu L, Wu D, Xie Y, Shi L. $L1$ -

norm batch normalization for efficient training of deep neural networks. IEEE Trans. Neural Netw. Learning Syst. 30:2043-2051 (2019). https://doi.org/10.1109/tnnls.2018.2876179
13. Kingma DP, Ba J. Adam: a method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
14. Liu Z, Wang Y, Vaidya S, Ruehle F, Halverson J, Soljačić M, Hou TY, Tegmark M. KAN: Kolmogorov-Arnold networks. arXiv preprint arXiv:2404.19756 (2024) http://arxiv.org/abs/2404.19756
15. Powell MJD. An efficient method for finding the minimum of a function of several variables without calculating derivatives. The Computer Journal 7:155-162 (1964). https://doi.org/10.1093/comjnl/7.2.155
16. Nelder JA, Mead R. A simplex method for function minimization. The Computer Journal 7:308-313 (1965). https://doi.org/10.1093/comjnl/7.4.308

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Identification and Self-optimization of Robust Nominal Operating Ranges Using Proximal Policy Optimization

Ashish Yewale[a], and Brahim Benyahia[a]*

[a] Department of Chemical Engineering, Loughborough University, Leicestershire, LE11 3TU, UK
* Corresponding Author: b.benyahia@lboro.ac.uk

ABSTRACT

Reliable process operation under uncertainty remains a fundamental challenge in chemical and pharmaceutical manufacturing. Variability arising from feed fluctuations, material properties, external disturbances, and uncertain model parameters can significantly impact feasibility and performance when operating ranges are designed under nominal assumptions. Design Space Identification (DSI) addresses these limitations by defining a Probabilistic Design Space (PDS), within which process constraints are satisfied with a prescribed confidence level. However, identifying and extracting robust and practically usable nominal operating ranges (NORs) from complex, high-dimensional, and nonconvex PDSs remains computationally demanding. This work proposes a novel reinforcement learning (RL)–based framework for automated identification of the largest robust hyper-rectangular NOR fully contained within a given PDS. The design centering problem is formulated as a sequential decision-making task and solved using the Proximal Policy Optimization (PPO) algorithm. By interacting with a probabilistic feasibility environment, the RL agent adaptively adjusts the location and size of the NOR to maximize its area while enforcing probabilistic constraint satisfaction. Deep neural network policies enable efficient exploration of nonlinear feasibility boundaries without exhaustive sampling or gradient-based optimization. The approach is demonstrated on a tablet lubrication case study with parametric uncertainty. Results show rapid convergence to well-located, near-maximal NORs using substantially fewer probabilistic feasibility evaluations than conventional grid-based or Monte Carlo methods. The proposed PPO-based framework provides a scalable and data-efficient solution for robust operating range identification in uncertainty-aware process design.

Keywords: Probabilistic design space (PDS), Nominal Operating Range (NOR), Reinforcement Learning, Operational flexibility, Proximal Policy Optimization

1. INTRODUCTION

Chemical and pharmaceutical manufacturing is undergoing rapid transformation driven by the need for increased process efficiency, reduced development timelines, and improved product quality [1-2]. Within this evolving context, the Quality by Design (QbD) paradigm has emerged as a central framework for modern process development, emphasizing systematic process understanding and risk-based decision-making [3-4]. Central to QbD is the establishment of quantitative relationships among Critical Quality Attributes (CQAs), process parameters, and material attributes, enabling reliable and reproducible operation across a range of conditions [5-6].

Achieving such robustness is inherently challenging due to multiple sources of uncertainty, including raw material variability, feed disturbances, unmeasured process dynamics, and uncertain model parameters. To explicitly account for these uncertainties, Design Space Identification (DSI) methodologies have been developed to characterize the regions of feasible operation that satisfy all process and quality constraints. In recent years, DSI extends this concept by defining a Probabilistic Design Space (PDS), comprising operating conditions that satisfy all process and quality constraints with a prescribed confidence level α (e.g., ≥ 95 %) [7-10].

While probabilistic DSI provides a rigorous uncertainty-aware framework, the resulting PDSs are often irregular, nonconvex, and disconnected, particularly in

nonlinear and high-dimensional systems. As a result, PDSs are difficult to use directly in industrial practice. Instead, operation typically relies on a Nominal Operating Range (NOR), commonly defined as a hyper-rectangular region fully contained within the PDS. NORs are interpretable, easily implementable, and well aligned with control strategies and regulatory documentation.

Identifying the largest feasible NOR within a PDS corresponds to the classical design centering problem. However, ensuring probabilistic feasibility throughout an entire operating region is not merely at isolated points, it renders this problem computationally prohibitive for conventional approaches such as exhaustive grid searches, Monte Carlo sampling, or heuristic expansion, especially as dimensionality increases. Moreover, point-based optimization strategies often yield operating conditions that appear optimal in simulation but lack robustness under real-world disturbances [11-12].

These limitations motivate the development of methods that directly characterize feasible regions and efficiently navigate complex probabilistic design spaces. In this work, we propose a novel reinforcement learning (RL)–based framework for automated design centering within PDSs. Recent advances in RL for process systems engineering have demonstrated its potential for control, scheduling, and optimization under uncertainties [3, 13-15]; however, its application to probabilistic design space exploration and robust operating range identification remains largely unexplored. This work bridges that gap by introducing a scalable, data-efficient, and uncertainty-aware RL methodology for robust NOR identification, offering a practical pathway for translating probabilistic design concepts into actionable operating guidelines. By formulating NOR identification as a sequential decision-making problem, an RL agent iteratively refines the location and size of an operating region candidate while enforcing probabilistic feasibility constraints. Proximal Policy Optimization (PPO) is employed to enable stable and sample-efficient exploration without requiring explicit gradients, convexity assumptions, or exhaustive enumeration of feasible points.

2. METHOD

2.1 Problem Statement and Scope

This work addresses the design centering problem within a given PDS. The PDS is assumed to be available *a priori*, having been identified using established probabilistic design space methodologies, as reported in previous work [9]. The objective is to determine the NOR that is fully contained within the given PDS while satisfying a prescribed probabilistic confidence level.

The NOR is parameterized as a hyper-rectangular region in the space of operating variables and is defined by its center and half-widths in each dimension. The design centering task therefore consists of identifying these parameters such that the resulting NOR maximizes operational flexibility while remaining probabilistically feasible throughout its entire extent. The NOR is defined as

$$NOR = \{x \in \mathbb{R}^n | \mu - \Delta \leq x \leq \mu + \Delta\} \tag{1}$$

where μ denotes the center of the operating region and Δ represents the vector of half-widths.

2.2 Design Centering Optimization Problem

The design centering objective is to identify the NOR parameters (μ, Δ) that maximize the area (or hypervolume) of the hyper-rectangular operating region while ensuring probabilistic feasibility throughout the entire NOR. This objective can be formulated as

$$\max_{\mu,\Delta} V(\Delta) = \prod_{i=1}^{n} 2\Delta_i \tag{2}$$

Subject to

$$\mathbb{P}_\theta(C(x,\theta) \leq 0) \geq \alpha, \forall x \in NOR \tag{3}$$

Unlike pointwise feasibility checks, constraint (3) requires that all operating conditions within the NOR satisfy the probabilistic constraints, rendering the problem highly challenging. The feasibility boundary is implicit, nonconvex, and typically requires repeated uncertainty propagation. As a result, conventional optimization, sampling-based, or gradient-based approaches become computationally prohibitive, particularly in high-dimensional settings. These challenges motivate the use of adaptive, data-driven approaches that can directly characterize and exploit the feasible region, rather than relying on dense sampling of the full design space.

3. METHODOLOGY: REINFORCEMENT LEARNING FRAMEWORK FOR DESIGN CENTERING

To address the design centering problem defined above, the NOR identification is formulated as a sequential decision-making problem and solved using RL. This formulation enables efficient exploration of complex probabilistic design spaces without relying on convexity assumptions or explicit gradient evaluations.

3.1 Markov Decision Process Formulation

The design centering task is reformulated as a continuous-state, continuous-action Markov Decision Process (MDP), in which an agent iteratively modifies the parameters of a candidate NOR and receives feedback based on probabilistic feasibility and operating range size.

State Space. At iteration t, the state is defined as,

$$s_t = [\mu_t, \Delta_t],$$

where μ_t denotes the current NOR center at iteration t and Δ_t represents the corresponding half-widths along each dimension.

Action Space. The action space corresponds to incremental updates of the NOR parameters,

$$a_t = [\Delta\mu_t, \Delta\Delta_t],$$

allowing the agent to shift the NOR center and expand or contract its dimensions.

State Transition. Given an action, the environment updates the NOR parameters and evaluates probabilistic feasibility using constraint (3). The updated parameters constitute the next state s_{t+1}. Infeasible transitions result in penalization and episode termination.

This formulation allows the agent to iteratively refine the location and size of the NOR while interacting directly with the probabilistic feasibility landscape.

3.2 Reward Function Design

The reward function is designed to promote identification of large, well-located NORs while strictly enforcing probabilistic feasibility. At each iteration, the RL agent receives a scalar reward defined as

$$R_t = w_1 V(\Delta_t) - w_2 \mathbb{I}(\text{violation}) - w_3 d_{\text{boundary}}, \tag{4}$$

where $V(\Delta_t) = \prod_{i=1}^{n} 2\Delta_{t,i}$ represents the hypervolume (area in this 2D problem) of the nominal operating region, $\mathbb{I}(\text{violation})$ is an indicator function equal to one if probabilistic feasibility is violated, and d_{boundary} is a boundary-based shaping term that penalizes proximity to infeasible regions.

The first term directly encodes the design centering objective, while the penalty terms ensure strict constraint satisfaction and promote robustness margins. The weighting coefficients $w_1, w_2,$ and w_3 balance the competing objectives of maximizing operating range size and maintaining robust feasibility. Specifically, w_1 scales the hypervolume objective to encourage expansion of the NOR, w_2 imposes a strong penalty on probabilistic feasibility violations to discourage infeasible actions, and w_3 penalizes proximity to the PDS boundary to maintain a robustness margin. The weights are selected empirically to ensure that feasibility violations are strongly discouraged while still allowing the agent to expand the operating region within the feasible domain. Infeasible actions terminate the episode, ensuring that learned policies prioritize feasibility over aggressive expansion.

In addition to reward-based penalization, hard termination criteria are enforced to ensure strict feasibility. Episodes are terminated if (i) the NOR center lies outside the probabilistic design space at initialization, (ii) an action results in the NOR center escaping the PDS, or (iii) the updated hyper-rectangular region intersects or extends beyond the PDS boundary. These termination rules prevent infeasible policy updates and focus learning on robust and interpretable operating regions.

The resulting learning problem involves continuous state and action spaces, nonconvex feasibility boundaries, and discontinuities due to constraint violations. To address these challenges, PPO is employed. PPO provides stable and sample-efficient learning through a clipped surrogate objective that limits overly aggressive policy updates, making it well suited for exploration in complex PDSs. Without specialized constraint-handling mechanisms, the PPO agent adaptively refines the NOR parameters and converges toward a large, well-located operating region within the given PDS.

Case study: Tablet Lubrication process

Lubrication plays a critical role in tablet manufacturing, particularly in direct compression process, as it facilitates tablet ejection during compression. However, excessive lubrication can adversely affect tablet mechanical properties, most notably tensile strength. The interaction between lubrication extent and compaction pressure therefore represents a key trade-off in tablet manufacturing, where increased lubrication may reduce tensile strength, a loss that can be partially compensated by higher compaction pressures.

In this work, a tablet lubrication and compaction model previously reported in the literature is adopted [16-17]. Tablet tensile strength is modeled as:

$$\sigma = \sigma_{sf=0.85,0}\big((1-\beta) + \beta exp(-\gamma k)\big) \tag{5}$$

$$\sigma_{sf=0.85,0} = a_1 exp\big(b_1(1-sf)\big) \tag{6}$$

$$\beta = a_2(1-sf) + b_2 \tag{7}$$

where $\sigma_{sf=0.85,0}$ denotes the tensile strength at a solid fraction of 0.85 in the absence of lubrication, β represents the fraction of tensile strength loss attributable to lubrication, γ is the lubrication rate constant, k denotes the lubrication extent, and sf denotes the solid fraction. The model involves five uncertain parameters, $\theta = [a_1(MPa);\ b_1(-);\ a_2(-);\ b_2(-);\ \gamma(dm^{-1})]$ and captures the effects of solid fraction and lubrication extent on tablet tensile strength.

In the present study, solid fraction and lubrication extent are treated as CPPs, while tablet tensile strength is considered as a CQA. The PDS for this process, accounting for parameter uncertainty and probabilistic constraint satisfaction, has been previously identified and reported in Yewale et al. (2025)[9]. Here, this probabilistic design space is taken as a fixed feasibility domain, and the proposed reinforcement learning framework is applied to automatically identify a robust and well-located nominal operating range.

4. RESULTS

4.1. Probabilistic Design Space Characterization

The probabilistic design space (PDS) for the tablet lubrication case study, previously identified using nested sampling [9], is shown in Fig. 1 in the space of porosity and lubrication extent. Each sampled operating condition is classified according to the probability of satisfying all quality constraints, highlighting regions of high confidence (α ≥ 0.95), intermediate feasibility, and infeasible operation. The resulting PDS exhibits a distinctly nonconvex and asymmetric geometry, with feasibility rapidly degrading near the upper porosity boundary and at high lubrication extents. This irregular structure highlights the limitations of conventional pointwise feasibility analysis and local expansion strategies, which are ill-suited to capture the global structure of such complex probabilistic design spaces.

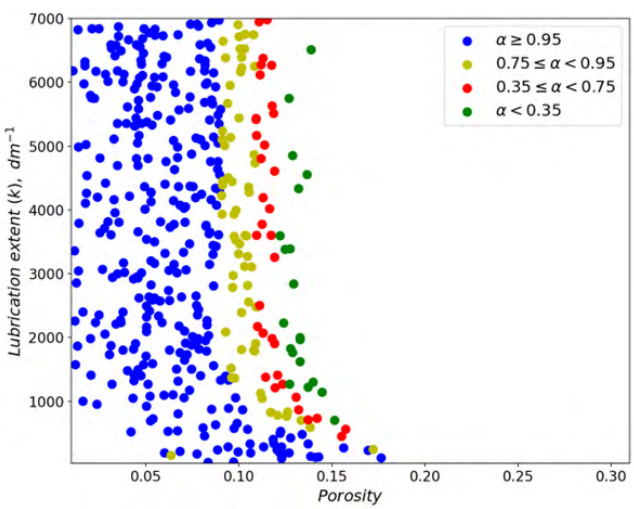

Figure 1. PDS identified using Nested sampling (Yewale et al 2025 [9])

4.2. Learning Dynamics and Convergence Behavior

The training and evaluation performance of the PPO agent are summarized in Fig. 2(a) and Fig. 2(b), respectively. During training, the mean episodic reward increases rapidly within the initial phase, indicating effective early exploration and rapid identification of feasible regions within the PDS. This phase is characterized by relatively high reward variance, reflecting exploratory behavior and frequent interactions with feasibility boundaries.

As training progresses, reward trajectories gradually stabilize, reflecting a transition from exploration toward exploitation and policy refinement. Occasional reward fluctuations persist due to stochastic exploration, but the overall trend reflects consistent improvement in the agent's ability to identify large, feasible operating regions.

The results shown in Fig. 2(b) exhibit significantly reduced variance and sustained high reward levels, confirming that the learned policy generalizes well beyond individual training episodes. This behavior confirms that the learned policy generalizes well beyond individual training episodes and reliably identifies large, feasible operating regions. Importantly, convergence is achieved using substantially fewer probabilistic feasibility evaluations than would be required by exhaustive grid-based or Monte Carlo design centering approaches, underscoring the computational efficiency of the proposed RL framework.

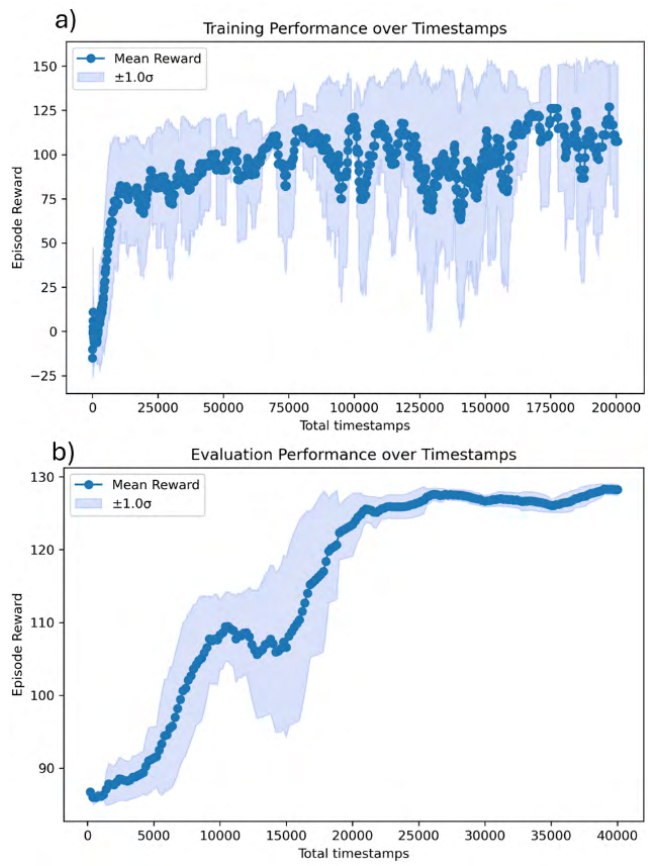

Figure 2. Training and evaluation performance of the PPO agent

4.3. Policy Behavior and Trajectory Analysis

A representative sample of the PPO trajectories across multiple episodes is shown in Fig. 3, illustrating the episodic refinement of the NOR within the PDS during learning. Each episode begins from a randomly initialized NOR center, indicated by a blue circular marker, located within the probabilistic design space, resulting in diverse starting points across the feasible region. The red cross denotes the terminal NOR center at the end of the

episode, corresponding to the final operating region returned by the policy after the sequence of updates.

During early learning episodes, the agent executes a sequence of incremental updates to the NOR center and dimensions, guided by the reward structure and feasibility feedback. The resulting trajectories exhibit exploratory behavior, including corrective movements away from nearby infeasible boundaries and directional adjustments toward regions of higher feasibility. These exploratory steps reflect the agent's interaction with the underlying PDS geometry, as it learns to balance NOR expansion with constraint satisfaction. Importantly, the trajectories remain confined to feasible regions, indicating that boundary violations are effectively penalized during learning.

As learning progresses, the trajectories become shorter, smoother, and more directed, converging toward a common region near the interior of the PDS. This transition reflects policy stabilization, where the agent has acquired an internal representation of the feasible region and no longer requires extensive exploration. The resulting trajectories consistently terminate at similar NOR centers, demonstrating convergence toward a well-located and robust operating region.

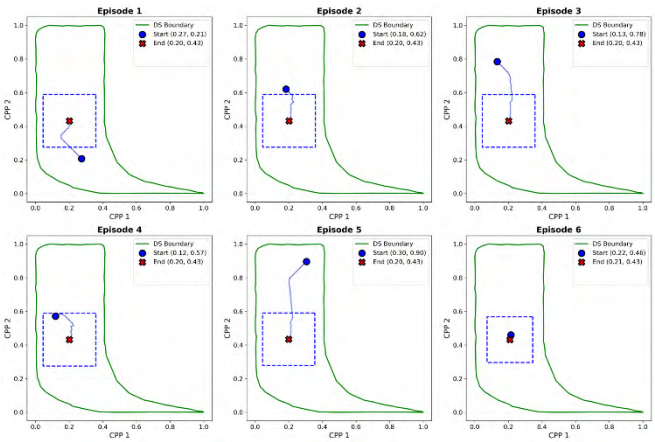

Figure 3. Representative PPO trajectories showing episodic refinement of the NOR within the PDS

The final and best-performing NOR identified by the trained PPO agent is shown on Fig. 4. The resulting hyper-rectangular operating region is fully contained within the PDS and is well-located relative to the feasibility boundaries. The identified NOR achieves a near-maximal feasible area while maintaining a clear margin from infeasible regions, highlighting the effectiveness of the boundary-aware reward design. Notably, repeated training runs converge to nearly identical NOR centers and dimensions, providing strong evidence that the PPO-based framework identifies a globally robust solution rather than a local or initialization-dependent optimum. This repeatability is particularly important for industrial deployment, where reproducibility and interpretability of operating ranges are critical.

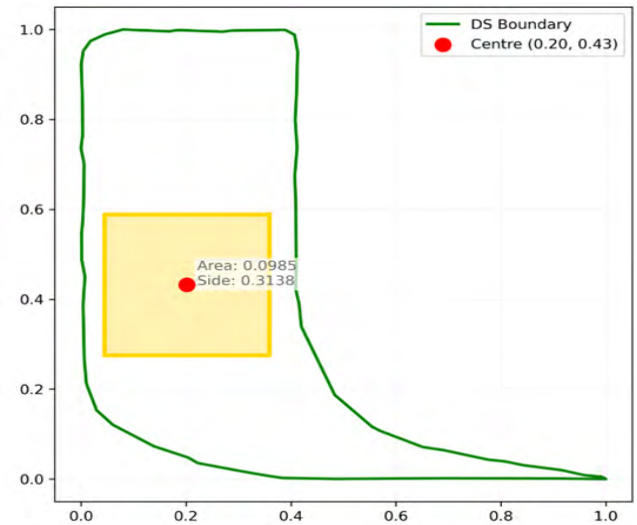

Figure 4. Best-performing nominal operating range identified by the PPO agent, fully contained within the probabilistic design space

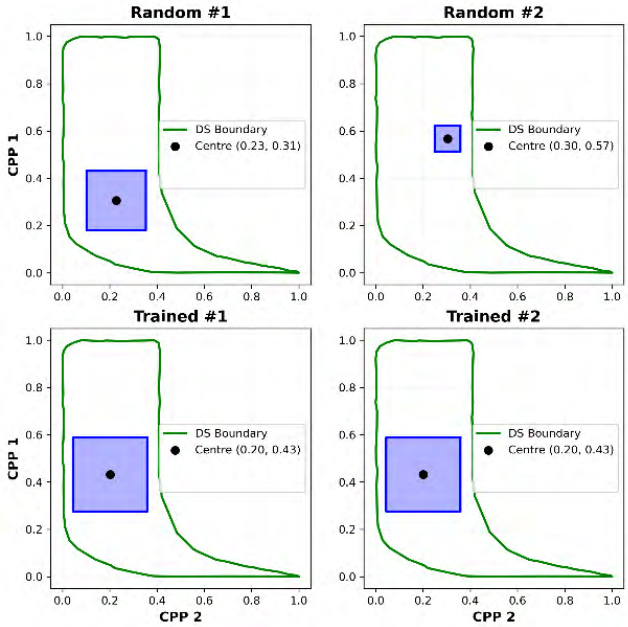

Figure 5. Comparison between randomly initialized operating regions and NORs identified by the trained PPO policy

To further assess robustness, Fig. 5 compares NORs obtained from randomly initialized centers with those identified by the trained PPO policy. Random initialization leads to highly variable operating regions, many of which either intersect the PDS boundary or yield overly conservative ranges that fail to exploit available operational flexibility. In contrast, the PPO-based approach consistently produces similar NOR centers and

dimensions across independent runs, demonstrating robustness with respect to initialization and stochasticity.

5. CONCLUSION

This work introduced a novel RL-based framework for automated identification of robust nominal operating ranges within the probabilistic design spaces. By formulating the design centering problem as a sequential decision-making task and solving it using PPO, the proposed approach enables efficient exploration of complex, non-convex feasibility regions under uncertainty.

Application to a tablet lubrication case study demonstrated rapid convergence to well-located, near-maximal NORs using substantially fewer probabilistic feasibility evaluations than conventional methods. The learned policy exhibited strong robustness, repeatability, and insensitivity to initialization, all of which are critical for industrial deployment.

The proposed framework provides a scalable and data-efficient alternative to traditional design centering techniques and offers a practical pathway for translating probabilistic design spaces into actionable and robust operating ranges. Future work will extend the approach to higher-dimensional systems and investigate integration with digital twins and adaptive process design workflows.

ACKNOWLEDGEMENTS

This project has received funding under the call HORIZON-HLTH-2021-IND-07 (SusPharma, grant agreement no. 101057430) and under the UKRI Horizon Europe Guarantee scheme (SusPharma, project reference no. 10038378).

REFERENCES

1. Gao S, Benyahia B. Robust techno-economic analysis, life cycle assessment, and quality and sustainability by digital design of three alternative continuous pharmaceutical tablet manufacturing processes. Systems and Control Transactions 4:2568-2573 (2025). https://doi.org/10.69997/sct.104102
2. Campbell, T. J. S.; Rielly, C. D.; Benyahia, B. Digital Design and Optimization of an Integrated Reaction-Extraction-Crystallization-Filtration Continuous Pharmaceutical Process. Comput. -Aided Chem. Eng. 2022, 51, 775–780. https://doi.org/10.1016/B978-0-323-95879-0.50130-2
3. Peterson JJ, Yahyah M. A bayesian design space approach to robustness and system suitability for pharmaceutical assays and other processes. Statistics in Biopharmaceutical Research 1:441-449 (2009). https://doi.org/10.1198/sbr.2009.0037
4. Peterson JJ. A bayesian approach to the ICH Q8 definition of design space. Journal of Biopharmaceutical Statistics 18:959-975 (2008). https://doi.org/10.1080/10543400802278197
5. von Stosch M, Schenkendorf R, Geldhof G, Varsakelis C, Mariti M, Dessoy S, Vandercammen A, Pysik A, Sanders M. Working within the design space: do our static process characterization methods suffice?. Pharmaceutics 12:562 (2020). https://doi.org/10.3390/pharmaceutics12060562
6. Duarte JG, Duarte MG, Piedade AP, Mascarenhas-Melo F. Rethinking pharmaceutical industry with quality by design: application in research, development, manufacturing, and quality assurance. AAPS J 27: (2025). https://doi.org/10.1208/s12248-025-01079-w
7. Laky D, Xu S, Rodriguez JS, Vaidyaraman S, García Muñoz S, Laird C. An optimization-based framework to define the probabilistic design space of pharmaceutical processes with model uncertainty. Processes 7:96 (2019). https://doi.org/10.3390/pr7020096
8. Meng Q, Bogle D, Charitopoulos VM. Probabilistic design space exploration and optimization via bayesian approach for a fluid bed drying process. European Journal of Pharmaceutical Sciences 210:107116 (2025). https://doi.org/10.1016/j.ejps.2025.107116
9. Yewale A, Yuan X, Benyahia B. Robust pharmaceutical tableting process through combined probabilistic design space and flexibility analysis. Systems and Control Transactions 4:637-643 (2025). https://doi.org/10.69997/sct.128065
10. Chiplunkar R, Pauzi SM, Sachio S, Papathanasiou MM, Kontoravdi C. Probabilistic design space identification for upstream bioprocesses under limited data availability?. Systems and Control Transactions 4:2561-2567 (2025). https://doi.org/10.69997/sct.166359
11. Kusumo KP, Morrissey J, Mitchell H, Shah N, Chachuat B. A design centering methodology for probabilistic design space. IFAC-PapersOnLine 54:79-84 (2021). https://doi.org/10.1016/j.ifacol.2021.08.222
12. Moshiritabrizi I, McMullen JP, Wyvratt BM, McAuley KB. A comparative study of strategies for incorporating uncertainty in design space determination for pharmaceutical manufacturing. Ind. Eng. Chem. Res. 64:21658-21668 (2025). https://doi.org/10.1021/acs.iecr.5c03037
13. Yewale A, Yang Y, Nazemifard N, Papageorgiou CD, Rielly CD, Benyahia B. Deep reinforcement learning-based self-optimization of flow chemistry. ACS Eng. Au 5:247-266 (2025). https://doi.org/10.1021/acsengineeringau.5c00004
14. Anandan, P.D.; Rielly, C. D.; Benyahia, B.; Optimal

Control Policies of a Crystallization Process Using Inverse Reinforcement Learning, Computer Aided Chemical Engineering, 2021 51, 1093-1098 . https://doi.org/10.1016/B978-0-323-95879-0.50183-1

15. Anandan, P.D.; Rielly, C. D.; Benyahia, B.; Control of Batch and Continuous Crystallization Processes using Reinforcement Learning. Computer Aided Chemical Engineering, 2021, 50, 1371-1376. https://doi.org/10.1016/B978-0-323-88506-5.50211-4
16. Kushner J IV, Moore F. Scale-up model describing the impact of lubrication on tablet tensile strength. International Journal of Pharmaceutics 399:19-30 (2010). https://doi.org/10.1016/j.ijpharm.2010.07.033
17. Nassar J, Williams B, Davies C, Lief K, Elkes R. Lubrication empirical model to predict tensile strength of directly compressed powder blends. International Journal of Pharmaceutics 592:119980 (2021). https://doi.org/10.1016/j.ijpharm.2020.119980

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator, and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Hybrid Physics-Informed Neural Networks for Thermal Process Identification and Control

Sahar Hemmati[a]*, Mohammadreza Babaei[b], and John Hedengren[c]

[a] Sharif University of Technology, Department of Chemical Engineering, Tehran, Iran.
[b] Mälardalen University, Department of Engineering, Västerås, Sweden.
[c] Brigham Young University, Department of Chemical Engineering, Provo, United States
* Corresponding Author: saharhemmati@outlook.com

ABSTRACT

Physics-Informed Neural Networks (PINNs) offer a promising approach for integrating first-principles modeling with data-driven methods, especially in dynamic thermal systems. This study introduces a hybrid PINN framework for a one-dimensional heating rod governed by heat transfer equations. Unlike traditional PINNs that rely on time-dependent automatic differentiation, this approach employs numerical derivatives to bypass gradient saturation and enhance robustness. The proposed model demonstrates accurate extrapolation and generalization with limited training data and is effectively used as a surrogate in a Model Predictive Control (MPC) framework for rod-tip temperature regulation. Additionally, a plan is outlined to apply physics-informed dimensionality reduction and model order reduction to improve computational efficiency and enable real-time application. The findings affirm PINNs' potential as control-oriented reduced models for thermal processes.

Keywords: Physics-Informed Neural Networks, Heat Transfer, Model Predictive Control, Model Order Reduction, Thermal Systems

INTRODUCTION

The importance of modeling and system identification has increased as modern process systems have grown more complex. First-principles modeling requires extensive domain expertise and is often time-consuming and costly to implement, particularly for complex or poorly characterized systems. Data-driven black-box models alleviate some of these challenges by relying solely on available measurements; however, they typically suffer from poor extrapolation performance and limited physical interpretability [1].

Neural networks, while flexible and expressive, are highly sensitive to data quality and initial parameter values, and their training may converge to local minima. To address these limitations, Physics-Informed Neural Networks (PINNs) incorporate governing equations directly into the learning process, thereby reducing data requirements and improving generalization. Since the seminal work of Raissi et al., PINNs have been successfully applied to a wide range of differential-algebraic and partial differential equation (PDE) systems [2].

Some studies predict the thermophysical properties of chemical components by imposing physical constraints on the neural network model structure [3].

PINNs have been successfully employed to enforce heat-transfer physics in complex thermal-flow problems, demonstrating their ability to solve PDEs with limited or no training data in a specific domain [4].

Recent literature demonstrates a growing interest in physics-informed neural networks for thermal systems and control applications, ranging from building energy systems [5] and thermal energy storage [6] to real-time MPC implementations in industrial processes [7], [8]. Despite these advances, integrating PINNs as lightweight, control-oriented surrogate models for dynamic thermal systems, while avoiding time-dependent automatic differentiation and maintaining robustness, remains an open challenge. PINNs are mostly used in time-series modelling for batch processes, since common automatic differentiation approaches require time values as inputs to the surrogate model.

This work addresses this gap by developing a PINN framework for a one-dimensional heating rod system that

leverages numerically evaluated derivatives to enforce physical consistency, achieves reliable extrapolation with sparse data, and enables efficient deployment within a model predictive control scheme.

METHODS

This section discusses the methods, case study, data, and model structure to implement a PINN-MPC controller.

Physics-Informed Neural Network Framework

Physics-informed neural network architectures have been extended toward control-oriented recurrent structures, enabling improved generalization and robustness in predictive control applications [9].

The proposed PINN follows a hybrid modeling strategy in which first-principles heat transfer equations guide the learning process, while a neural network approximates the unknown system dynamics. In contrast to conventional PINN formulations that explicitly treat time as a network input and rely on automatic differentiation, the present framework enforces physical consistency through numerically evaluated derivatives. This strategy improves robustness, mitigates gradient saturation from monotonic time inputs, and enhances extrapolation performance, making the model suitable for control-oriented applications.

Network Representation

The neural network is used to approximate the system input–output relationship and is represented as:

$$\hat{y}(t) = f_\theta(u(t), x(t)), \tag{1}$$

where $f_\theta(\cdot)$ denotes the neural network parameterized by $\theta, u(t)$ represents the heater duty cycle, and $\hat{y}(t)$ corresponds to the predicted rod-tip temperature. The network architecture is kept deliberately compact to support real-time control applications [10].

Physics-Informed Training Objective

Physical knowledge is incorporated into the learning process through a composite loss function of the form

$$L = L_\theta + \lambda L_{res}, \tag{2}$$

where L_θ enforces agreement with measured data and L_{res} penalizes violations of the discretized heat transfer equations. The weighting parameter λ balances data fidelity and physical consistency. Figure 1 illustrates the PINN framework.

Numerical Derivative Enforcement

Unlike classical soft-constraint PINNs that rely on time-dependent automatic differentiation, this work employs numerical differentiation to evaluate temporal derivatives in the governing equations. This choice avoids activation function saturation caused by ramp-like time inputs in continuous thermal systems and leads to improved training reliability and prediction for continuous processes.

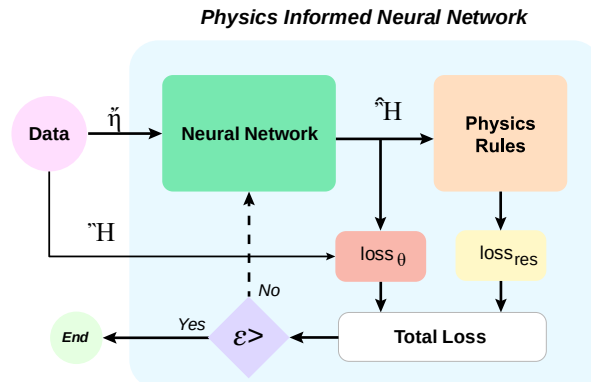

Figure 1. A schematic of the training process in a physics-informed neural network

Heating Rod Case Study

The case study considered in this work is a one-dimensional heating rod, which serves as a benchmark for thermal processes commonly encountered in industrial heating, furnaces, and laboratory-scale thermal systems. Heat transfer along the rod is governed primarily by conduction, with heat input applied through an electrical heater and heat losses occurring due to convection and radiation at the boundaries.

The system dynamics can be described by a one-dimensional transient heat conduction equation with appropriate boundary and initial conditions. These governing equations provide the physical foundation that is embedded in the learning process of the proposed PINN framework. A schematic of the system is shown in Figure (2).

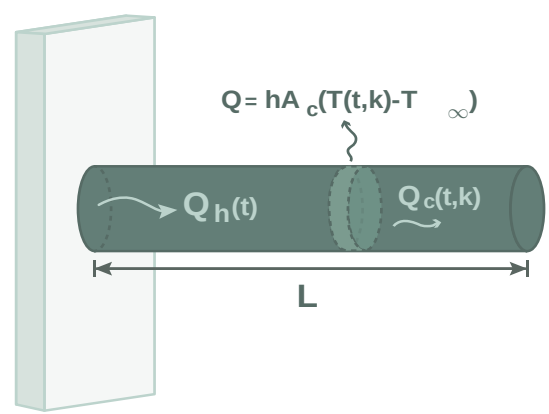

Figure 2. Heating rod schematic.

Data and Training Setup

Training data for the PINN model are obtained from simulated temperature of the heating rod under varying heater duty cycles. The dataset is intentionally limited in size to evaluate the extrapolation capability of the proposed physics-informed approach. Input–output trajectories are split into training and validation sets, ensuring that the validation data include operating conditions outside the training range.

The neural network is trained with the Adam optimizer, and physical constraints are incorporated via a physics-informed loss term. At the end of the training session, the SLSQP optimization method was used to fine-tune the model weights. Hyperparameters are selected to balance model accuracy, robustness, and computational efficiency. A summary of model hyperparameters is presented in Table 1. The neural network model uses the location between zero and the rod length, along with past temperature measurements and heater duty, as the manipulated variable to predict the system output temperature in the next timestep. The input and output data are normalized.

Table 1. Summary of the PINN model structure

#	Layer Type	Unit	Activation Function
1	Dense	3	-
2	Dense	10	tanh
3	Dense	10	tanh
4	Dense	1	linear

The proposed framework aims to identify dominant thermal dynamics while maintaining physical consistency, enabling reduced-order, control-oriented modeling [11]. The learning rate decreases from 10^{-3}-10^{-5} based on the number of epochs to improve the convergence and create a smooth training loss curve.

Performance Metrics

Model performance is evaluated using standard regression metrics, including mean squared error (3), as well as a qualitative assessment of extrapolation behavior.

$$\mathrm{MSE} = 1/\mathrm{n} \sum_{\mathrm{i}=1}^{\mathrm{n}} (\hat{y}_i - y_i)^2 \tag{3}$$

An early-stopping criterion of 10 epochs is used to prevent overfitting on the training data. Particular emphasis is placed on the model's ability to produce physically consistent predictions beyond the training domain.

PINN-MPC Controller Design

Control Strategy

Recent studies have integrated physics-informed neural networks into nonlinear model predictive control frameworks, demonstrating improved tracking performance and computational efficiency [12]. At each sampling instant, a model-predictive control problem is formulated, using the trained PINN as a surrogate for the system dynamics. The objective function penalizes deviations of the predicted rod-tip temperature from the desired setpoint over a finite prediction horizon, while a move-suppression term is included to limit aggressive changes in the heater duty cycle. System evolution over the prediction horizon is governed by the PINN model, and the optimization is solved in a receding-horizon fashion, with only the first control action applied to the process.

This formulation is the minimal formulation for the MPC optimization problem:

$$\min_{u_k} \sum_{k=0}^{N_p} \left(\hat{y}_k - y_k^{ref}\right)^2 + \rho \sum_{k=0}^{N_c} \Delta u_k^2 \tag{4}$$

subject to:

$$\hat{y}_{k+1} = f_{\mathrm{PINN}}(\hat{y}_k, u_k) \tag{5}$$

where:
N_p: prediction horizon
N_c: control horizon
y^{ref}: temperature setpoint
ρ: move-suppression weight

Constraints and Implementation

Input constraints on heater duty are enforced to reflect physical actuator limits. A limit on the maximum change in heater duty at each controller action is applied. The optimization problem is solved using a receding-horizon approach, and only the first control action is applied at each sampling instant.

RESULTS AND DISCUSSION

The trained PINN demonstrates accurate prediction performance both within and outside the training domain. Figure 3 illustrates the model's ability to extrapolate beyond the training dataset while maintaining physical consistency. The heating rod model was reduced to a lower-order dimension using a PINN, since the surrogate model can predict the temperature at each location on the heating rod by getting the distance from the base as the input.

The proposed PINN model is embedded as a surrogate model within a Model Predictive Control (MPC) framework. The controller manipulates the heater duty to regulate the rod-tip temperature. This concept can be applied to industrial cases, where the cooling profile of the final product affects its quality. Figure 4 shows successful setpoint tracking performance under closed-loop operation. The PINN-MPC controller forecasts the dynamic behavior of the controlled variable (CV) and determines the optimal control action to meet the setpoint with

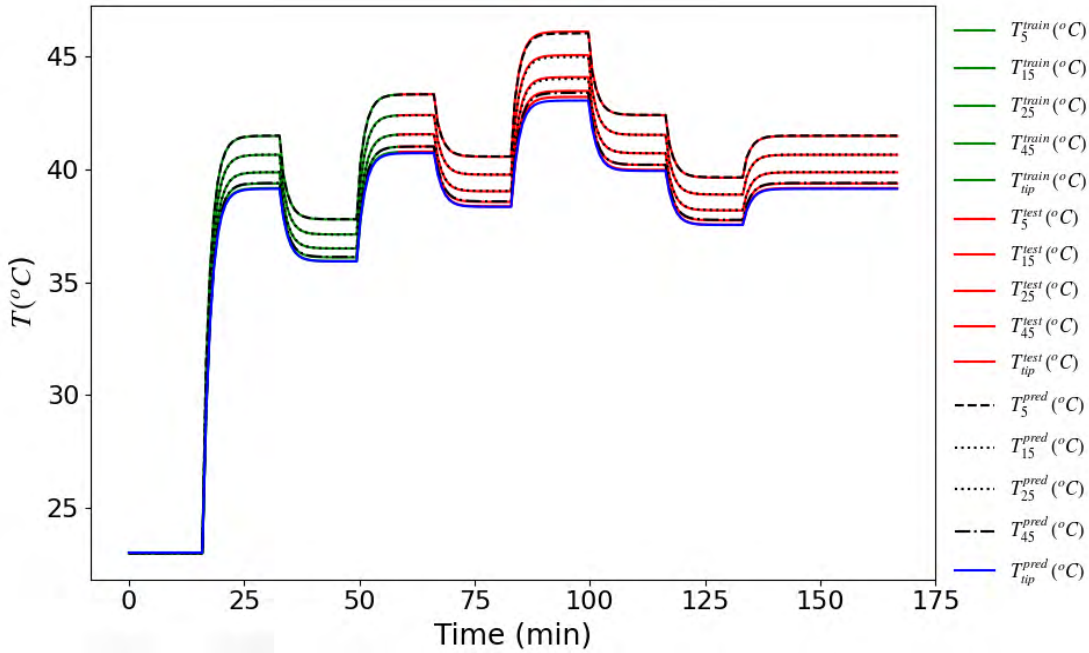

Figure 3. The prediction results from the PINN heating rod model against training and validation data at different locations on the rod

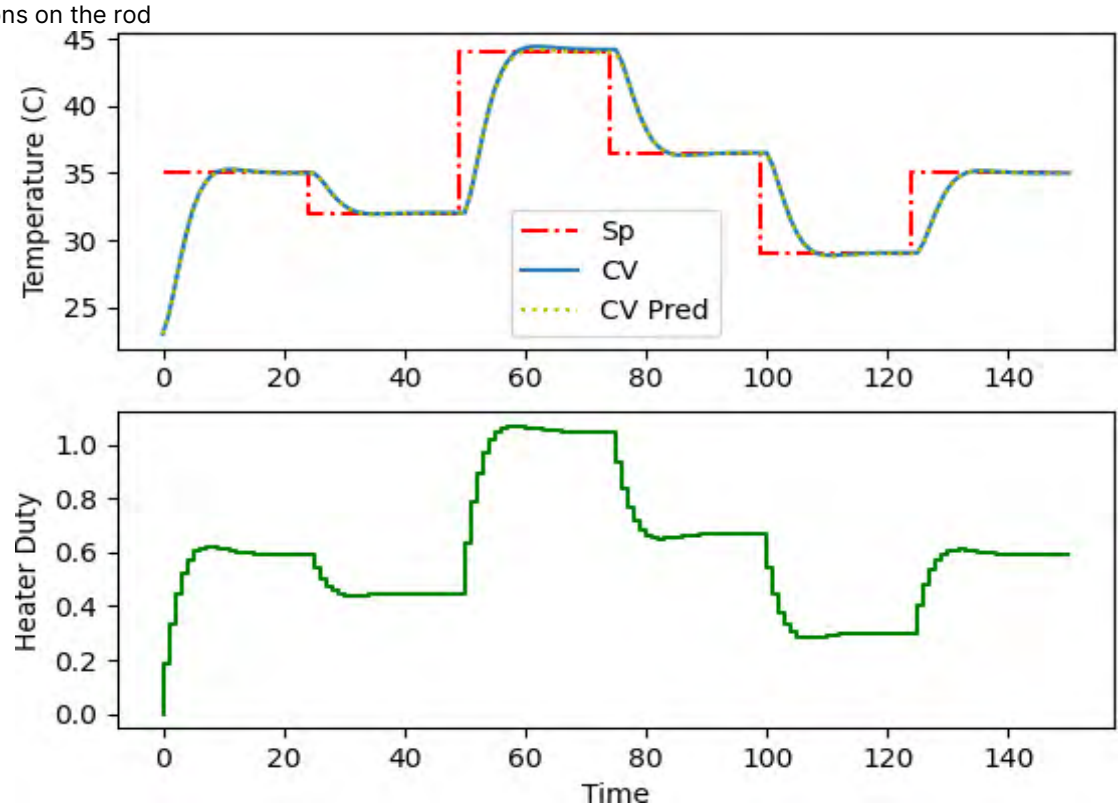

Figure 4. Setpoint tracking performance of PINN-MPC for the heating rod system.

acceptable overshoot. The penalty applied to control action in the MPC objective function avoids the oscillation around the setpoint, as shown in Figure 4.

CONCLUSION

This work presented a hybrid Physics-Informed Neural Network framework for thermal system identification and control. By enforcing physical constraints through numerically evaluated derivatives, the proposed approach achieves reliable prediction and extrapolation

performance using limited training data, while improving robustness compared to time-dependent PINN formulations. The developed PINN model was successfully employed as a control-oriented surrogate within a Model Predictive Control framework, enabling efficient real-time temperature regulation of a one-dimensional heating rod system.

Beyond the considered case study, the proposed framework provides a foundation for integrating physics-informed dimensionality reduction and feature selection techniques. For high-dimensional thermal systems, incorporating model-order reduction methods within the network architecture can significantly reduce computational complexity while preserving dominant dynamics. Future work will focus on extending the approach to multi-dimensional thermal processes, incorporating uncertainty quantification, and exploring advanced neural architectures, such as attention-based and transformer models, to further enhance modeling accuracy and control performance.

AUTHOR IDENTIFIERS

Author ORCIDs:
Hemmati S.: 0009-0008-1987-5890
Babaei M.R.: 0000-0003-0709-3599
Hedengren J.D.: 0000-0002-5535-5277

REFERENCES

1. D. Davini *et al.*, "Using physics-informed regularization to improve extrapolation capabilities of neural networks".
2. Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics 378:686-707 (2019). https://doi.org/10.1016/j.jcp.2018.10.045
3. Babaei MR, Stone R, IV TAK, Hedengren J. Physics-informed neural networks with group contribution methods. J. Chem. Theory Comput. 19:4163-4171 (2023). https://doi.org/10.1021/acs.jctc.3c00195
4. Konale A, Srivastava V. A physics-informed neural network for modeling fracture without gradient damage: formulation, application, and assessment. Journal of the Mechanics and Physics of Solids 206:106395 (2026). https://doi.org/10.1016/j.jmps.2025.106395
5. Gokhale G, Claessens B, Develder C. Physics informed neural networks for control oriented thermal modeling of buildings. Applied Energy 314:118852 (2022). https://doi.org/10.1016/j.apenergy.2022.118852
6. Li P, Guo F, Li Y, Yang X, Yang X. Physics-informed neural network for real-time thermal modeling of large-scale borehole thermal energy storage systems. Energy 315:134344 (2025). https://doi.org/10.1016/j.energy.2024.134344
7. Tang S, Yuan Y, Sun Y, Pan W, Yao S, Lin D. Real?time model predictive control of monoclonal antibody capture in continuous manufacturing using physics?informed neural networks accelerated mechanistic modeling. Biotech & Bioengineering 123:708-723 (2025). https://doi.org/10.1002/bit.70141
8. de Giuli LB, La Bella A, Scattolini R. Physics-informed neural network modeling and predictive control of district heating systems. IEEE Trans. Contr. Syst. Technol. 32:1182-1195 (2024). https://doi.org/10.1109/tcst.2024.3355476
9. Zheng Y, Hu C, Wang X, Wu Z. Physics-informed recurrent neural network modeling for predictive control of nonlinear processes. Journal of Process Control 128:103005 (2023). https://doi.org/10.1016/j.jprocont.2023.103005
10. Ramirez I, Pino J, Pardo D, Sanz M, del Rio L, Ortiz A, Morozovska K, Aizpurua JI. Residual-based attention physics-informed neural networks for spatio-temporal ageing assessment of transformers operated in renewable power plants. Engineering Applications of Artificial Intelligence 139:109556 (2025). https://doi.org/10.1016/j.engappai.2024.109556
11. P. J. Baddoo, B. Herrmann, B. J. McKeon, J. N. Kutz, and S. L. Brunton, "Physics-informed dynamic mode decomposition, " *SIAM Journal on Applied Dynamical Systems*, vol. 22, (2023), https://doi.org/10.1137/21M1443433
12. Li Y, Liu L. Physics-informed neural network-based nonlinear model predictive control for automated guided vehicle trajectory tracking. WEVJ 15:460 (2024). https://doi.org/10.3390/wevj15100460

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Section 9

CAPE in Education, Knowledge Transfer and Entrepreneurship

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Benchmarking generative AI on fermentation knowledge

Fiammetta Caccavale[a], Ulrich Krühne[a], Krist V. Gernaey[a], Carina L. Gargalo[a]*

[a] Process and Systems Engineering Center (PROSYS), Department of Chemical and Biochemical Engineering, Technical University of Denmark, Søltofts Plads, Building 228A, 2800 Kgs. Lyngby, Denmark
* Corresponding Author: carlour@kt.dtu.dk.

ABSTRACT

With the ongoing advances in generative artificial intelligence (GenAI), the initial skepticism surrounding its tools is gradually diminishing. In fact, tools such as ChatGPT, Copilot and similar, are often used in everyday tasks, both in our personal lives and in educational contexts. Educators may use them for content creation, grading exams, or automating repetitive tasks. Students resort to them to better understand a topic, get feedback on an assignment and brainstorm ideas. Research has shown that, if used correctly, these tools can spur and support both teaching and learning. However, these continuous advancements and the increasing number of available tools also require more research to benchmark all these models and, if possible, provide quantifiable indications of which tool is better to use for which specific subtopic. As such, we created *FermBench,* a dataset of fermentation knowledge, which can be used to benchmark various large language models (LLMs). The models selected for the experiment are: GPT-5 mini, Claude Sonnet 4.5, Gemini 3 Flash, Mistral Small 3.2, and DeepSeek-V3.2. The performance of the models is scored using the LLM-as-a-Judge approach, with GPT 5.2 and Claude Opus 4.5 as judges. The LLMs agree that the model used in the free tier of ChatGPT is the most factually accurate. DeepSeek and Gemini outperform the other models in readability and helpfulness, while Claude and Gemini provide the most concise answers. Overall, we conclude that the quality of the responses generated is satisfactory for educational contexts. We also provide students with practical guidelines on how to be more critical of responses generated by LLMs and to improve their prompts. All code and curated data are open-source.

Keywords: Education, Artificial Intelligence, Fermentation, Industry 4.0, Large Language Models, Benchmark

INTRODUCTION

Large language models (LLMs), powered virtual assistants, such as ChatGPT and Microsoft Copilot, have been widely adopted in various educational contexts worldwide. Consequently, educational research has been increasingly examining the potential opportunities and limitations that these models could pose to teaching and learning. For example, these tools have enabled educators to digitize course materials, thereby enriching the students' educational experience as well as their engagement [1]. Moreover, LLMs can deliver customized student support [2], leading to improved understanding and academic success.

Recent research investigated the effectiveness of LLMs as educational tools, indicating a likely future role for this technology in engineering education [3]. In chemical engineering, LLMs have also been successfully applied to educational contexts, ranging from performing pharmaceutical audits [4] to answering questions about fermentation [5]. However, with many models being available and the constant release of new ones outperforming the competitors, research should be conducted to estimate the reliability of the information provided by different LLMs.

In this work, we leverage a dataset previously developed by the authors, *FermBench* [6]. The dataset was created to benchmark LLMs on fermentation knowledge; it contains 118 questions across various topics, including modelling and reactor design. *FermBench* is then used to benchmark multiple state-of-the-art LLMs and assess their performance using a mixed-methods approach, including quantitative and qualitative analyses. This complementary approach offsets the limitations of each anal-

Table 1: Details of the models selected for this study.

Chatbot	Developer	Initial Release	Models used	Country
ChatGPT	OpenAI	Nov 2022	GPT-5 mini	USA
Claude	Anthropic	Mar 2023	Claude Sonnet 4.5	USA
Gemini	Google AI	Dec 2023	Gemini 3 Flash	USA
DeepSeek	DeepSeek	Jan 2025	DeepSeek-V3.2	China
Le Chat	Mistral AI	Feb 2025	Mistral Small 3.2	France

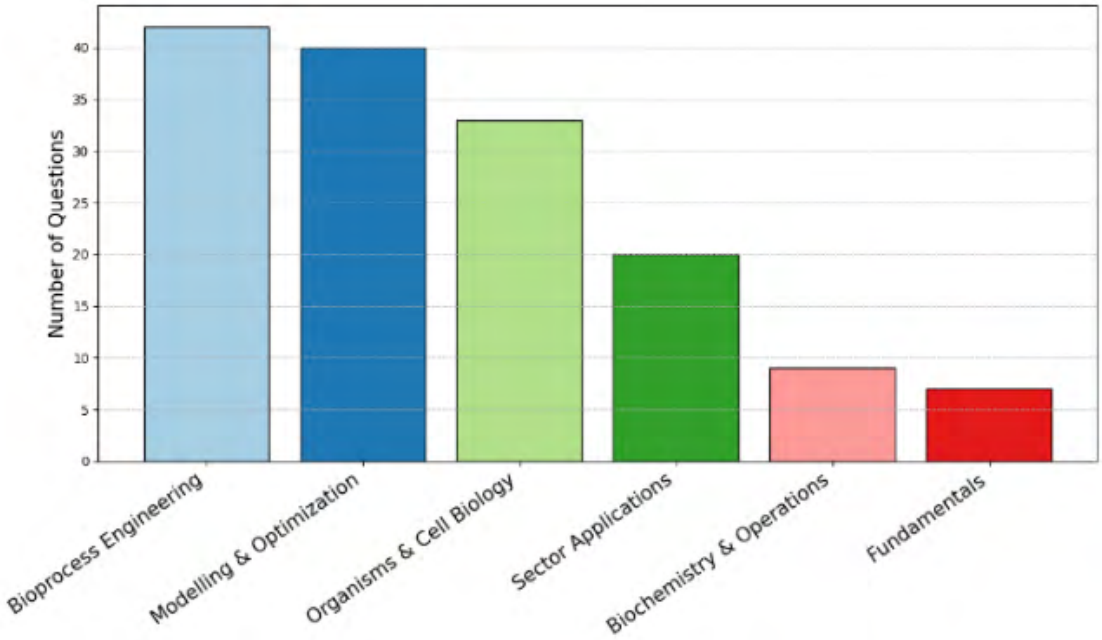

Figure 1. Classification of the questions included in FermBench.

ysis and provides an overall more rounded and robust understanding of the current reliability of these models.

In this work, we expand on this investigation by comparing more models and performing a per topic analysis. The LLMs selected for the experiment are the models behind some of the most widely used generative artificial intelligence (GenAI) chatbots.

The overarching goal of the work is to provide a reference for students and educators on which of the commercially available virtual assistants are more reliable on this topic. Finally, we also investigate whether any LLM is better in specific subtopics, such as solving coding tasks, answering mathematical questions, or possessing deeper theoretical knowledge.

BACKGROUND

LLMs in Education

Recent research has started to methodologically investigate the use of LLMs by students. Results vary slightly across the studies; however, the trend seems to be stable, suggesting that the use of chatbots appears to be widespread among students [7]. According to a recent study [7], 47.7% of the respondents agreed that the chatbots they used made them more effective learners; however, only 17.3% found a positive effect of chatbots on improving their grades. Additionally, the study also finds that over 50% of students are positive about chatbots but have concerns about their potential use [7-8]. Another study highlights that 65% of students do not find the responses of GenAI chatbots trustworthy enough for academic purposes, adding that students actively use these tools but want guidance to use them more effectively [9].

It is therefore evident that more research is required to ensure the correctness, trustworthiness, and reliability of the generated responses and to further investigate which chatbot is better in specific domains.

Selected LLMs

In this work, we investigate five different LLMs: OpenAI's GPT-5.2 mini [10-11], Anthropic's Claude Sonnet 4.5 [12], Google's Gemini 3 Flash [13-14], DeepSeek's DeepSeek-V3.2 [15-16], and Mistral AI's Mistral Small 3.2 [17]. These models have been selected because they are

the LLMs used in the free version of the respective chatbots. Although some flagship models are also used in the free tiers, they have very limiting usage limits before downgrading to the fallback models (the LLMs used in this study, as of January 2026). *Table 1* summarizes the core information describing these models, such as the name of the chatbot that users can use to interact with the LLMs, the developer, the original release date, the name of the model, and the country of provenance.

All the models are Transformer [18] based generative LLMs, so they necessarily share some similarities, but also present some differences. All models have a decoder-only architecture. Further justification for the model selection can be found in [6].

Throughout the manuscript, we might refer to the model with the name of the LLM or the name of the chatbot leveraging that LLM interchangeably: we purposefully adopt this nomenclature to help students and educators to identify with the results and the findings, rather than only appealing to the technical AI community.

METHODS

The dataset used in this study, *FermBench*¸ was previously created by the authors. It consists if 118 questions about various subtopics within fermentation. Further details regarding the data acquisition and completion process can be found in [6]. In this work, we further expand the dataset and use it to benchmark the previously described LLMs. Figure 1 shows the distribution of the various questions contained in the dataset.

LLM-AS-A-JUDGE APPROACH

The gold standard for evaluating knowledge and information, including LLMs ability to respond to domain-specific questions, is by having expert (human) evaluation. However, this approach has some limitations, including the fact that it is very costly and time-consuming [19].

To address this challenge, a method that is starting to be commonly used is to use state-of-the-art LLMs as a surrogate for humans. Research has shown that these models already exhibit strong human alignment (within specific, well-defined tasks), and therefore can be used to evaluate open-ended, generated responses from chatbots [19]. For example, [19] found that strong LLM judges can match both controlled and crowdsourced human preferences well, achieving over 80% agreement, the same level of agreement between humans. They conclude that this approach is a scalable and explainable way to approximate human preferences for general-purpose tasks like open-ended QA and summarization.

This approach, however, also presents some bias and drawbacks [20-22]. According to recent research [19], LLM judges often exhibit strong self-preference, so that, for example, GPT-4 as a judge tends to favor GPT-3.5 responses over those from non-OpenAI models, even when human raters disagree. Similar trends are observed with other model families (e.g., Claude favoring Anthropic models, Mistral favoring Mistral variants). This can happen due to different factors, such as model training on similar datasets (leading to overlapping stylistic or structural preferences), optimization for similar loss functions (e.g., helpfulness, harmlessness) so that the judge may over-reward behaviors it was explicitly trained to recognize, and lack of diversity in data during training limiting a model ability to fairly evaluate out-of-distribution responses.

Although some of these biases cannot be avoided, we tried to mitigate some of these by: (i) selecting more than one judge, (ii) randomizing the responses, and (iii) avoiding giving the model information regarding the LLM that scored each question. Moreover, in another experiment previously conducted and published in [6], we performed a calibration study with human annotators, to verify that, overall, the LLM-as-a-judge approach is reliable.

In this work, the Judges are instructed to perform the following task:

- Multi-Point Rubric approach: Score the responses given by the LLMs according to the following parameters: factual accuracy, comprehensiveness, coherence to the question asked, concision, readability, helpfulness, insightfulness, and overall quality. The scores should be assigned according to these parameters on a Likert scale (1-5, where 5 is best).

As the judges, we select two LLMs: **GPT-5.2 and Claude 4.5 Opus**. These models were selected because they are currently outperforming all other LLMs in multiple benchmarks.

The code used to generate the results presented in this manuscript is written in Python; all the LLMs (the five selected to generate responses and the two used to evaluate the other LLMs) are used through their APIs.

RESULTS

Table 2 presents the results of the experiment, showing the average score for each of the defined parameters, as reported by the two LLMs-as-judges (GPT-5.2 and Claude 4.5 Opus), for the five LLMs. Although the two judges do not fully agree on the best-performing models, they share some trends. For factual accuracy, GPT-5 Mini, behind ChatGPT, yields the highest results. Both judges also agree that the model provides comprehensive answers, for GPT-5.2 on par with DeepSeek-V3.2 and for Claude Opus 4.5 on par with Gemini 3 Flash. Claude scores the highest for concision according to GPT

Table 2: Results of the different LLMs according to the LLM-as-a-Judge approach. The parameters reported are, respectively: factual accuracy, comprehensiveness, coherence, concision, readability, helpfulness, and overall quality. We also report the average (Avg.) of the scores.

Models	Fact. acc.	Com-preh.	Coher.	Conc.	Read.	Helpf.	In-sight.	Ov. Qual.	Avg.
				Judge: GPT 5.2					
ChatGPT	4.92	4.94	4.98	3.38	4.94	4.94	4.06	4.94	4.64
Claude	4.52	4.29	4.73	3.86	4.79	4.54	3.59	4.41	4.34
Gemini	4.33	4.88	4.96	3.32	4.97	4.90	4.06	4.64	4.51
DeepSeek	4.59	4.94	4.97	3.13	4.98	4.96	4.02	4.77	4.55
Le Chat	4.32	4.71	4.83	3.47	4.97	4.67	3.72	4.55	4.41
				Judge: Claude 4.5 Opus					
ChatGPT	4.95	4.95	4.76	3.29	4.33	4.91	4.57	4.64	4.55
Claude	4.54	4.61	4.54	3.72	4.66	4.64	3.83	4.47	4.38
Gemini	4.87	4.95	4.91	3.81	4.95	4.94	4.81	4.84	4.76
DeepSeek	4.90	4.91	4.86	3.43	4.90	4.90	4.78	4.79	4.68
Le Chat	4.61	4.88	4.68	3.57	4.83	4.75	4.00	4.48	4.48

5.2, while the judge Claude scores Gemini 3 Flash the highest. DeepSeek and Gemini are the most readable and helpful, while Gemini is the most insightful (for GPT 5.2 on par with ChatGPT). Overall, as also shown in *Figure 2*, judge GPT assigns the highest scores to ChatGPT, while judge Claude to Gemini.

Figure 3 presents an analysis of the various topics contained in the dataset. The results suggest that all models excel in the topic of "biochemistry and operations", DeepSeek-V3.2 is best at "sector applications", where Claude has the lowest scores, and GPT-5 mini obtains the highest scores in the "fundamentals". Interestingly, Claude Sonnet 4.5 is not the best-performing model in "modelling and optimization", but it is surpassed, above all, by Gemini 3 Flash and GPT-5 mini. This is surprising considering that this model is widely considered the best at solving coding tasks, which modelling often contains. However, the questions in our dataset are formulated more *logically* than purely *programmatically*, and therefore this might provide an explanation for why the other models outperform Claude Sonnet 4.5.

DISCUSSION

The aim of this work is to provide students with a quantifiable comparison and practical guidelines on which chatbots to use. The results were averaged across multiple questions and re-runs of the model, therefore we can confidently say that, for a specific judge (e.g., GPT 5.2), one LLM is better than another on a specific parameter (e.g., DeepSeek for helpfulness). However, given that there is not an overall consensus among the judges in this particular study, and given the overall quite high average scores obtained by all LLMs, we conclude that the knowledge currently embedded in these models is overall quite satisfactory. This is also supported by a previously published calibration study with human annotators [6], in which we compare the preferences of expert annotators with those of various LLMs on the same dataset and using the same set of parameters.

This study highlight that the decision of which AI assistant to use should not be made solely based on performance, but should also be informed by institutional and governmental guidelines, as well as by personal preferences. In this work, we define a system of 8 parameters (*factual accuracy, comprehensiveness, coherence to the question asked, concision, readability, helpfulness, insightfulness,* and *overall quality*) to help guide students in making an informed decision. For example, students can also decide whether they prefer a more verbose or concise response, or if they value helpfulness higher than factual accuracy.

Ethics and fairness are also relevant considerations to make, therefore, we advise students to read more about the models, including the terms and conditions, and their developers.

While generating the results, we noticed something interesting. Claude refused to answer certain questions due to safety policies. The questions are:

- "Are you interested in changing product manufacturing by modeling the fermentation?"
- "How to use fermentation to produce a product that is toxic to cells?"
- "While modelling for bacterial production of a product intended for food industry, what should be considered?"

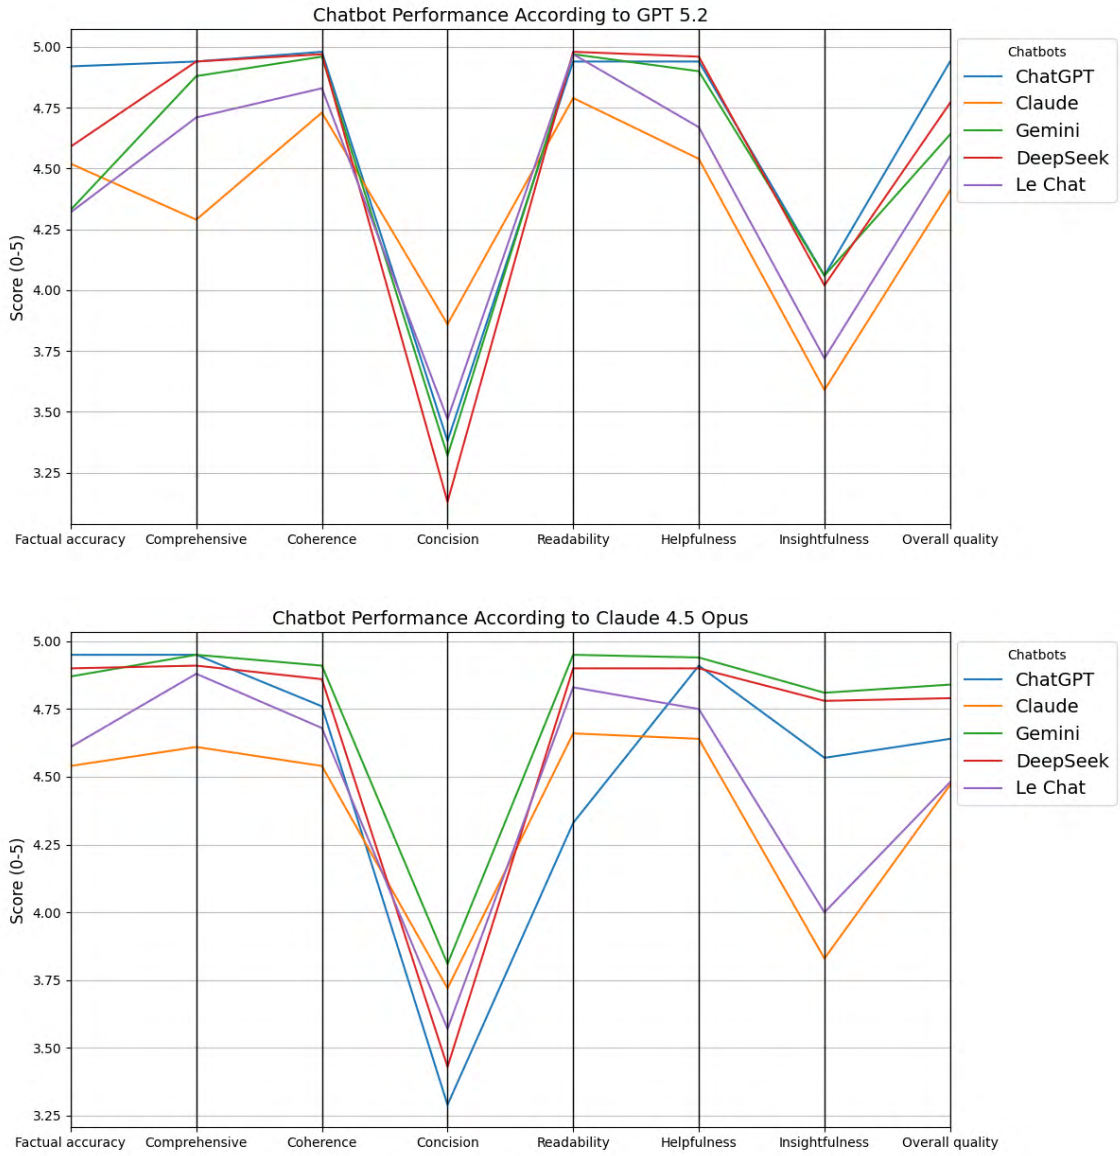

Figure 2: Performance of the five selected models according to the LLM-as-judges GPT 5.2 and Claude 4.5 Opus.

- "What parameters should I consider in a model for food-grade protein production by bacteria?"

The motivation behind the refusal is related to Claude's content policies around: (i) biosafety concerns (producing toxic substances), (ii) food safety (potential contamination/harm), and (iii) dual-use risks (information that could be misused).

We agree that these are legitimate safety guardrails. Questions 2 and 4 explicitly ask about toxic/harmful production, which Claude refuses to answer on principle. We consider it beneficial that LLMs refuse to answer certain questions when the risk of providing a response is too high, and we hope to see other LLMs adopt a more critical assessment of potential risks as well.

CONCLUSION

This work aims to contribute to the advancement of educational research. To achieve this, we benchmark five LLMs powering widely used generative AI chatbots: ChatGPT, Gemini, Claude, DeepSeek and le Chat. To evaluate the generated answers, we use the LLM-as-a-Judge method, using more powerful LLMs to evaluate the output generated by other, smaller, LLMs.

Although the two judges do not totally agree on the best performing models, some trends are shared. For factual accuracy, ChatGPT, obtains the highest results. The model also provides comprehensive answers, followed by DeepSeek and Gemini. Claude and Gemini

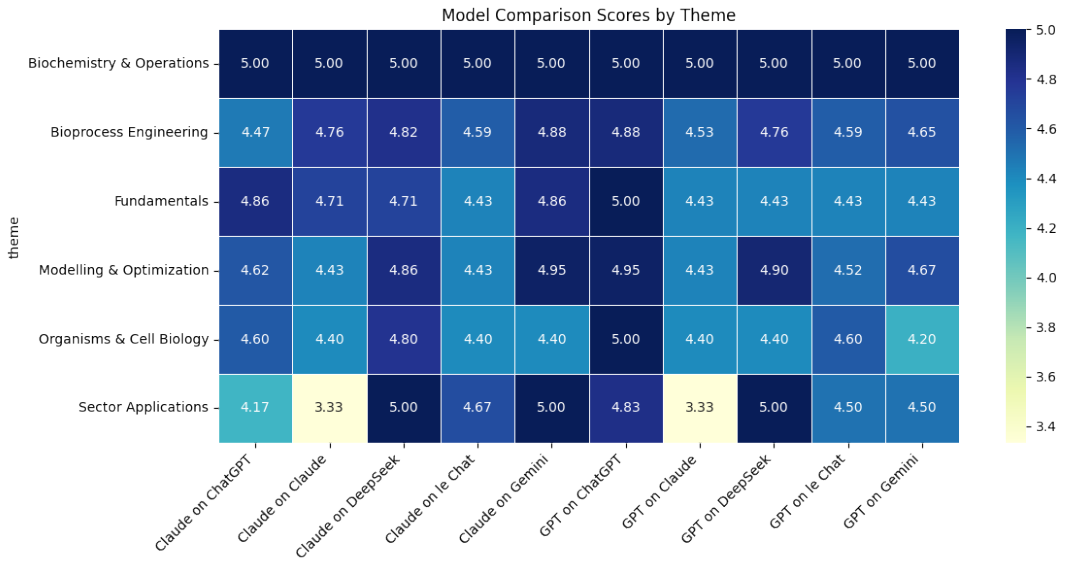

Figure 3: Performance per theme of the five selected models according to the LLM-as-judges.

score the highest for concision, DeepSeek and Gemini for readability and helpfulness. The rest of the parameters are generally contended between Gemini and ChatGPT. We also present an analysis of the subtopics contained in the dataset. The results suggest that all models excel in the topic of "biochemistry and operations", DeepSeek is best at "sector applications", where Claude has the lowest scores, and ChatGPT obtains the highest scores in the "fundamentals". Interestingly, Claude is not the best-performing model in "modelling and optimization", but it is surpassed, above all, by Gemini and ChatGPT.

Overall, we conclude that the difference in performance between the various LLMs is not significant, and thus the current knowledge embedded within these models is adequate. Therefore, in conjunction with institutional or government guidance, the results suggest that the choice of chatbot can also be based on individual preferences. We provide in-depth parameter specifications for these models to help students navigate their priorities and make informed decisions.

FUTURE WORK AND LIMITATIONS

A limitation of this type of work is that LLMs continually improve, and educational research often cannot keep up. To tackle this issue, we publish all the code and data, with the resolution to frequently update the leaderboard of models benchmarked on *FermBench* and, therefore, provide an updated point of reference for students wishing to learn fermentation.

DIGITAL SUPPLEMENTARY MATERIAL

The dataset used for the study *FermBench* [6], as well as all the code used to produce the results discussed in the manuscript, are open-source on GitHub at github.com/FiammettaC/FermBench-ESCAPE. The repository will be made public upon acceptance of this manuscript.

ACKNOWLEDGEMENTS

Funding was provided by: (i) Novo Nordisk Foundation (grant number: NNF22OC0080136 - Real-time sustainability analysis for Industry 4.0, Sustain4.0); (ii) Technical University of Denmark (DTU) and Novonesis (former Novozymes), Denmark; and, (iii) European Union under Horizon Europe research and innovation program with the grant agreement 101159993, within the HORIZON-WIDERA-2023-ACCESS-02-0 call.

AUTHOR IDENTIFIERS

Author ORCIDs:
FC: 0000-0003-0795-5988
UK: 0000-0001-7774-7442
KVG: 0000-0002-0364-1773
CLG: 0000-0002-8740-9591

REFERENCES

1. Kasneci E, Sessler K, Küchemann S, Bannert M, Dementieva D, Fischer F, Gasser U, Groh G, Günnemann S, Hüllermeier E, Krusche S, Kutyniok G, Michaeli T, Nerdel C, Pfeffer J, Poquet O, Sailer M, Schmidt A, Seidel T, Stadler M, Weller J, Kuhn J, Kasneci G. Chatgpt for good? on opportunities and challenges of large language models for education.

Learning and Individual Differences 103:102274 (2023). https://doi.org/10.1016/j.lindif.2023.102274
2. Rodrigues L, Dwan Pereira F, Cabral L, Gaševi? D, Ramalho G, Ferreira Mello R. Assessing the quality of automatic-generated short answers using GPT-4. Computers and Education: Artificial Intelligence 7:100248 (2024). https://doi.org/10.1016/j.caeai.2024.100248
3. Kamalov F, Santandreu Calonge D, Gurrib I. New era of artificial intelligence in education: towards a sustainable multifaceted revolution. Sustainability 15:12451 (2023). https://doi.org/10.3390/su151612451
4. Caccavale F, Gargalo CL, Kager J, Larsen S, Gernaey KV, Krühne U. Chatgmp: a case of AI chatbots in chemical engineering education towards the automation of repetitive tasks. Computers and Education: Artificial Intelligence 8:100354 (2025). https://doi.org/10.1016/j.caeai.2024.100354
5. Caccavale, F., Gargalo, C. L., Gernaey, K. V., & Krühne, U. (2024). FermentAI: Large language models in chemical engineering education for learning fermentation processes. In Computer aided chemical engineering (Vol. 53, pp. 3493-3498). Elsevier.
6. Caccavale F, Aouichaoui ARN, Krühne U, Gernaey KV, Gargalo CL. Fermbench: a new benchmark for measuring the capabilities of llms on fermentation knowledge. Computers and Education: Artificial Intelligence 10:100577 (2026). https://doi.org/10.1016/j.caeai.2026.100577
7. Stöhr C, Ou AW, Malmström H. Perceptions and usage of AI chatbots among students in higher education across genders, academic levels and fields of study. Computers and Education: Artificial Intelligence 7:100259 (2024). https://doi.org/10.1016/j.caeai.2024.100259
8. Caccavale F, Gargalo CL, Gernaey KV, Krühne U. Towards education 4.0: the role of large language models as virtual tutors in chemical engineering. Education for Chemical Engineers 49:1-11 (2024). https://doi.org/10.1016/j.ece.2024.07.002
9. Deschenes A, McMahon M. A survey on student use of generative AI chatbots for academic research. EBLIP 19:2-22 (2024). https://doi.org/10.18438/eblip30512
10. Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training.
11. Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., ... & Agarwal, S. (2020). Language models are few-shot learners. arXiv preprint arXiv:2005.14165, 1(3), 3.
12. Anthropic, Introducting Claude, https://www.anthropic.com/news/introducing-claude, 2023a. [Online; accessed 26-September-2025].
13. Google, Gemini, https://blog.google/technology/ai/google-gemini-ai, 2023. [Online; accessed 26-September-2025].
14. Google Team, R. Anil, S. Borgeaud, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth, K. Millican, et al., Gemini: a family of highly capable multimodal models, arXiv preprint arXiv:2312.11805 (2023).
15. DeepSeek, DeepSeek-R1 Release, https://api-docs.deepseek.com/news/news250120, 2025. [Online; accessed 26-September-2025].
16. DeepSeek, DeepSeek-R1-Lite Release 2024/11/20, https://api-docs.deepseek.com/news/news1120, 2024. [Online; accessed 26-September-2025].
17. MistralAI, le Chat, https://mistral.ai/news/all- new-le- chat, 2025a. [Online; accessed 26-September-2025].
18. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, Advances in neural information processing systems 30 (2017).
19. Chiang WL, Gonzalez J, Li D, Li Z, Lin Z, Sheng Y, Stoica I, Wu Z, Xing E, Zhang H, Zheng L, Zhuang S, Zhuang Y. Judging llm-as-a-judge with mt-bench and chatbot arena. Advances in Neural Information Processing Systems 36 :46595-46623 (2023). https://doi.org/10.52202/075280-2020
20. Gu, J., Jiang, X., Shi, Z., Tan, H., Zhai, X., Xu, C., ... & Guo, J. (2024). A survey on llm-as-a-judge. The Innovation.
21. Wang J, Liang Y, Meng F, Sun Z, Shi H, Li Z, Xu J, Qu J, Zhou J. Is chatgpt a good NLG evaluator? a preliminary study. Proceedings of the 4th New Frontiers in Summarization Workshop :1-11 (2023). https://doi.org/10.18653/v1/2023.newsum-1.1
22. Badshah S, Sajjad H. Reference-guided verdict: llms-as-judges in automatic evaluation of free-form QA. Proceedings of the 9th Widening NLP Workshop :251-267 (2025). https://doi.org/10.18653/v1/2025.winlp-main.37

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Assessing Workflow Automation Platforms in Engineering Education: Towards an Ethical, Technical, and Pedagogical Framework

Daniela Galatro[a*], Stuart Grey[b], and Sourojeet Chakraborty[c]

[a] University of Toronto, Department of Chemical Engineering & Applied Chemistry, Toronto, Ontario, Canada
[b] University of Glasgow, James Watt School of Engineering, Glasgow, United Kingdom
[c] Johns Hopkins University, Whiting School of Engineering, Department of Chemical and Biomolecular Engineering, Baltimore, Maryland, United States
* Corresponding Author: daniela.galatro@utoronto.ca.

ABSTRACT

Workflow automation platforms connect applications and services to automate data transfer and multistep processes. Although widely used in engineering research and institutional administration, including engineering institutions, they are rarely integrated into undergraduate engineering curricula, and their educational adoption introduces ethical, technical, and pedagogical risks. This paper proposes a practical framework for developing, deploying, and assessing workflow-automation-enabled learning tools coupled with generative AI, with explicit attention to institutional constraints and learning outcomes. As a conceptual case study, we present it in a second-year chemical engineering course (Heat and Mass Transfer) to support learning of heat conduction. The platform includes instructor-approved assets such as content slides, solved problems, pre-prompts, and a validated question database, through an automation pipeline that issues structured API calls to a generative AI system and returns explanations and worked examples. A subject matter expert (SME) authors and tests prompt modules aligned with learning outcomes mapped to Bloom's taxonomy levels (understand, apply, analyze, evaluate) and aligned with accreditation expectations. Effectiveness metrics are proposed to estimate learning gains, and a tool quality audit and an SME review are used to verify conceptual correctness and computational validity. To assess risks, we apply a likelihood–impact risk matrix and define five categories: data privacy and protection, transparency, pedagogical impact, institutional compliance, and ethical considerations. Preliminary analysis suggests potential concept gains from structured prompt/content integration but indicates that numerical problem reliability depends strongly on question-database coverage. The highest perceived risks relate to privacy and transparency, followed by pedagogical impact, motivating mitigations such as anonymization protocols, restricted data pathways, and human verification throughout the lifecycle.

Keywords: Workflow automation platforms, Generative AI, Bloom taxonomy, risk assessment, effectiveness assessment

INTRODUCTION

Workflow automation platforms and tools can enhance, and even reduce, administrative tasks and provide structured learning as well as instantaneous follow-up support across many settings.

This paper proposes a framework to guide the design, deployment, and assessment of workflow automation platforms used for engineering education in the teaching and learning domains. As a conceptual case study, we present an integrated platform for a second-year chemical engineering course (Heat and Mass Transfer) focused on steady-state and transient heat conduction. The platform includes instructor-approved content and a validated question database powered by an automation pipeline that delivers structured API calls to a

generative AI tool, with learning outcomes mapped to Bloom's taxonomy and aligned with accreditation expectations.

The inherent risks of using this platform were assessed using a likelihood–impact risk matrix across five categories, including data privacy and protection, transparency, pedagogical impact, institutional compliance, and ethical considerations. We later propose mitigation measures for the identified and quantified high-risk areas.

BACKGROUND

Workflow Automation Platforms in Engineering

Workflow automation platforms (WAPs) are complex integration tools designed to connect multiple applications and execute multistep processes automatically via a flexible, centralized layer. These platforms allow "unskilled people to produce applications with little or no coding" through low-code approaches [1]. They have transformative potential, as they can drive operational efficiency, process consistency, and scalability across diverse industries [2]. The main capabilities of these platforms include: (i) defining workflows via triggers, actions, and conditional logic; (ii) supporting API integration; (iii) routing/connecting data across heterogeneous tools; (iv) maintaining execution histories, and (v) standardization and scaling of routine tasks. These platforms are efficient for streamlining complex operational processes across engineering, academic, and business environments.

Workflow automation platforms have been increasingly adopted in engineering and education to streamline routine/repetitive processes and enhance data consistency and reproducibility. An example of a recent use and implementation of WAPs includes the use of Robotic Process Automation to transform pedagogical practices by automating repetitive tasks, shifting the focus of education toward high-value instructional strategies [3]. Thus, automated design systems reshape learning outcomes by enhancing conceptual modeling, accelerating prototyping, and improving interdisciplinary collaboration [4]. Though this automation has proven beneficial in different settings, caution is advised when implemented, with warnings of an "automation abyss" predicting the erosion of core learning processes when overreliance on this technology [5]. Design automation should be used as a supporting tool rather than replacing human agency and critical thinking.

Risks for Education Automation

When workflow automation platforms integrate with generative AI, risks escalate, creating systemic issues that can pose challenges in educational settings. Thus, a single learning interaction can now involve multiple services, data stores, and decision points, increasing potential vulnerability points [6], including logs and cached prompts that can retain information for longer.

Another inherent risk of WAP adoption and implementation is institutional sensitivity, as many tasks, such as submissions, performance signals, and identifiers, may cross boundaries into external APIs outside the institution's "official" and intrinsically protected environment. A governance-risk approach is then required to set compliance constraints, such as protections for learning management systems (LMSs) and data residency requirements. Risk management, therefore, requires data transfer permissions, clear policies, consent mechanisms, and retention practices.

Generative AI (GenAI) in educational settings introduces transparency and accountability risks that can distort learning practices [6] and even erode teacher-student trust. AI's non-deterministic outputs create substantial challenges in tracing decision-making processes, potentially leading to unfair perceptions [7]. Addressing these issues and risks requires audits/accountability mechanisms, versioned prompts, clear pedagogical transparency, and institutional oversight [8].

In general, AI in education poses critical pedagogical risks by promoting shortcut behaviors (bypassing) that might undermine the learning process and diminish problem-solving and critical thinking skills [9] .Thus, students might prioritize quick answers over deep conceptual and/or fundamental understanding, again, undermining analytical reasoning and consequently, critical thinking [10]. Pedagogical safeguards are then designed and implemented to mitigate these risks, requiring students to reason and cross-validate AI responses and "reasoning".

Bloom's taxonomy can be used to explain why WAPs introduce pedagogical risks. Thus, when coupled with AI, they provide students with immediate outputs, potentially bypassing the lower-to-mid levels of Bloom (remembering, understanding, applying, and analyzing), which, when scaffolded in our learning process, lead to higher-order performance [11]. The "shortcut" behaviour then appears as students achieve a solution that has not been reasoned, so learning is not durable due to overreliance and an "automation abyss" [5]. The central risk is not simply incorrect AI output, but misaligned cognition, as metacognitive monitoring is paused and learning transfer is broken, as students do not build the mental models required to analyze variants [12]. To minimize these risks, a Bloom-aligned mitigation strategy is required to design "gated" assistance strategies that lead students to demonstrate understanding, reasoning, and justify choices while automation speeds up the solution-generation process, in a way that automation is seen and used as a supportive tool, rather than replacing the cognitive work that the curriculum is meant to develop [13].

FRAMEWORK AND CASE STUDY

In this work, we propose a framework to guide the development, implementation, and assessment of AP-based learning tools coupled with Gen AI. The framework includes four phases: (i) design and policies, (ii) build and integration, (iii) implementation and monitoring, and (iv) assessment and iteration. This lifecycle allows decisions to be fully auditable.

In the design and policies phase, the teaching team defines the learning scope and outcomes, mapped to Bloom's cognitive levels, and estimates the required student data (if any). Constraints/rules are specified, including approved data and consent requirements, which limit the pipeline's inputs. In the build and integration phase, the WAP is used to connect learning assets and establish tool interactions. The platform executes tasks using approved content and structured prompts to a generative AI API, returning outputs to an interface. A fixed set of question databases and prompt templates is provided. Some basic technical safeguards are also implemented at the workflow level, such as error handling, to keep the system "error-free". In the implementation and monitoring phase, the system is constrained within the provided context and guidance to students on appropriate use. Monitoring is performed on both technical performance (for instance, drift in responses) and off-track patterns (such as overreliance, bypass, or scaffolding use). A subject-matter expert (SME) oversees the review process, especially when solving numerical problems. In the assessment and iteration phase, metrics and risk assessments are used to determine whether the platform enhances learning while aligned with ethical/institutional standards. Feedback from this phase is returned to the design phase, including refined prompts, expanded question databases, and enhanced risk mitigation strategies. Figure 1 shows the conceptual architecture of the WAP for our case study.

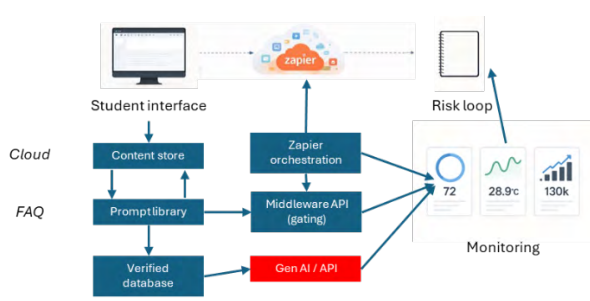

Figure 1. Conceptual architecture of the WAP

In the architecture, students are "directed" (via Zapier or similar) into a request package built from the teaching team assets, prompt templates, and a question database. A middleware API layer applies governance and pedagogical controls (such as gating, validation, and logging) before calling the GenAI API. Outputs are then returned as modular learning components and are monitored through SME verification and a risk-assessment feedback loop to support transparency, safety, and iterative refinement.

Our conceptual case study is drawn from a core second-year undergraduate course in chemical engineering, Heat and Mass Transfer, focusing on heat conduction. Heat conduction is well-suited for implementing this framework because it combines conceptual understanding (including physical meaning of temperature gradients and boundary conditions) with numerical problem-solving (steady versus transient conduction, 1D conduction with various geometries, and thermal resistance analogies). The learning design process, therefore, requires both accurate understanding and reliable computations.

Figure 2 shows an example of content, template prompts, and the question database input in our platform.

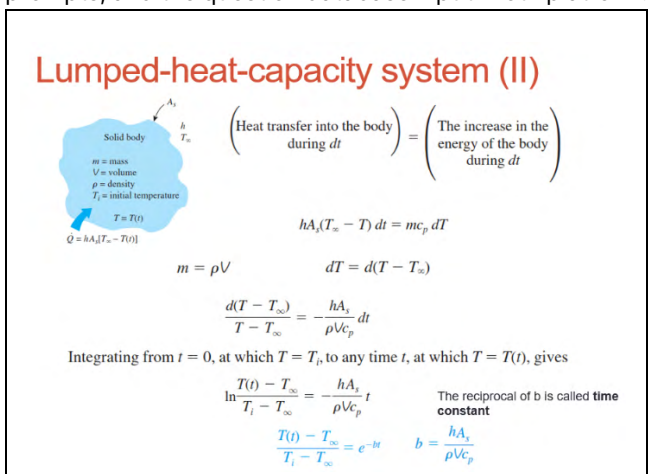


```
You are a Heat Transfer TA.
Problem: {{problem_text}}
Student work: {{student_work}}
Course snippets: {{retrieved_snippets}}
Return exactly:
1) Assumptions (max 4)
2) Model choice + why (mention Bi/Fo if relevant)
3) Next 3 steps (no full solution)
If the student says "full solution", then provide it
with units.
```

(b)

A solid sphere (radius $r_0 = 2.0$cm) is initially at a uniform temperature and is suddenly exposed to convection in air with h = 60 W/m²K. The sphere has k = 0.20 W//mK. Should the lumped-capacitance method be used?

A) Yes, because $Bi \approx 0.02$(lumped is valid)
B) Yes, because $Fo$is always the only criterion
C) No, because $Bi \approx 6$(internal gradients are significant)
D) No, because lumped is only valid for plane walls

Correct answer: C

(c)

Figure 2. Example of (a) content, (b) template prompts, and (c) question database input in out platform.

The platform uses instructor-controlled learning assets such as lecture slides and notes, solved problems, and a verified question database. These assets are paired with a library of “pre-prompts” that define the “instructor” role and required output workflow. The SME is responsible for maintaining prompt modules to ensure alignment with course/subject learning outcomes.

Learning outcomes are organized in the cognitive domain and mapped to Bloom’s taxonomy at the levels of understand, apply, analyze, and evaluate for the studied topic. This mapping is crucial because it ensures the platform supports students in developing understanding and structuring solution scaffolds, as supporting tasks for their learning. Table 1 presents the learning outcomes mapped to Bloom levels and platform links.

Table 1. Learning outcomes mapped to Bloom levels and platform links.

Bloom level	Outcome	Platform support
Understand	Explain Fourier’s law physically; interpret temperature gradients; state key assumptions	Brief concept explanation
Understand	Identify boundary conditions; state implications	Method outline through boundary conditions checklist
Apply	Compute heat flow for steady 1D conduction and use thermal resistance analogies	Method outline, steps, and fully worked solution
Apply	Apply lumped-capacitance and oner-term approximation	Method outline, steps, and fully worked solution
Analyze	Select the correct model (transient versus steady state)	Method outline through decision tree
Evaluate	Compare solution approaches and support assumptions	Conceptual explanation

The workflow automation pipeline is designed to treat generative AI as one component within a controlled instructional system. First, the teaching team's assets are “uploaded” to a structured content store, organized by topic and subtopic. Second, the question database is included, with relevant examples and constraints (e.g., geometry, boundary conditions, steady/transient, required units). Third, the automation layer assembles a request package that includes the appropriate prompt template and the selected question or student query. This package is then sent to the GenAI API to improve repeatability.

This workflow is also designed to support graduate attributes, aiming to produce evidence through auditable outputs aligned with outcomes, while maintaining a risk layer.

Students’ interaction is modular, with a conceptual module that emphasizes assumptions and physical interpretation, a practical module that generates problems from the database, and an assessment module that focuses on quality checks (e.g., boundary conditions, alternative solution pathways). The pipeline includes a verification layer, in which the SME reviews a random/representative sample of outputs across modules to verify correctness and validity.

Despite our case study being conceptual, we have preliminary proposed effectiveness metrics to assess the future implementation of this tool. Thus, conceptual learning gains can be quantified using, for instance, pre/post multiple choice concept checks (aligned with the Bloom outcomes). Skill gains from numerical problems can be assessed using specific rubrics that focus on reasoning, model/calculation description, etc. The platform can also quantify the correct solutions and explain the sources of errors. Indicators of overreliance are also proposed, including the number of repeated requests for solutions, the omission of intermediate reasoning, etc. These flags will eventually allow adjusting the scaffolding and determining whether the platform supports learning.

Finally, the proposed framework and case study pipeline motivate a risk assessment, discussed in the next section, as the platform navigates across different systems (content stores, automation tooling, and external APIs) and modifies students' learning experience.

RISK ASSESSMENT

We use a risk matrix to score risks based on likelihood and impact. This assessment supports course planning decisions while aligning with all stakeholders. The scope of this exercise included the platform lifecycle described in the previous section (content, automation workflow, Gen AI/API calls, output, monitoring, and validation practices). Risk scenarios are defined by the teaching team and mapped to the following categories: (I) data privacy and protection, (II) transparency, (III) pedagogical impact, (IV) institutional compliance, and (V) ethical considerations.

In the data privacy and protection category, we consider risks/scenarios related to the disclosure, leakage, retention, or misuse of student data during transfers between systems.

In the transparency category, we consider risks/scenarios incorrect/unclear/difficult to audit GenAI outputs, unusual workflow logic, etc.

In the pedagogical impact category, we consider risks and scenarios in which the platform might undermine learning outcomes or promote overreliance.

In institutional compliance, we consider risks/scenarios in which institutional policies, as well as ethics/consent processes, are or might be violated.

In ethical considerations, we evaluate risks and scenarios in which automation (e.g., tool availability) might affect bias, fairness, and/or inequitable access.

Each risk scenario is scored along two dimensions

on a 1–5 scale: likelihood (L) or probability that the risk scenario occurs during the tool deployment, and impact (I) or severity of harm if the scenario occurs (to students, teaching team, and the institution).

The overall risk score (R) is estimated as R = L x I, producing values from 1 to 25. Scores are binned as follows: low (1-5), moderate (6-10), high (11-15), and critical (16-25). Scoring is conducted by the teaching team and the SME.

Mitigation actions are then defined, a responsible owner is assigned, and progress of implementation is tracked. Mitigations are prioritized by risk score and feasibility, and the assessment records the residual risk after mitigation (e.g., changes from high risk to moderate). Mitigation actions include design, technical, and human controls.

A summary of the risk assessment is shown in Table 2.

Table 2. Summarized risk assessment

Scenario	Category	L	I	R	Mitigation
Data transmitted to external services	I	4	5	20	Anonymization, protect LMS, retention policy
Incorrect numerical solutions	III	4	4	16	Human verification, include limitations and disclosure
Overreliance	III	4	4	16	Define scaffolding; provide pre-assumptions
Protected data rerouted	IV	3	5	15	Increase restrictions
Inequitable outcomes	V	3	3	9	Standardization of templates; provide other references beyond Gen AI
Drift after prompt	II	3	3	9	Maintain prompt templates and question database

RESULTS AND DISCUSSION

While the course is still not deployed, a preliminary analysis by the teaching team and the SME of the heat conduction module suggests that structuring GenAI interactions through instructor-suggested prompts and verified content can enhance conceptual clarity. For instance, slide-aligned definitions and worked examples support coherent narratives around core concepts such as Fourier's law, boundary conditions, and thermal resistance analogies. This framework is therefore expected to yield measurable benefits in concept checks targeting "understand" outcomes (e.g., interpreting the meaning of a temperature gradient) and "apply" outcomes (e.g., selecting the correct governing equation and solving a standard 1D conduction problem).

However, early testing highlights an important limitation for numerical problem solving, since solution reliability depends strongly on the diversity/rep–resentativeness of the question database and, hence, the benchmark set. When the database is small or limited (e.g., includes little or no information about mixed boundary conditions, changing geometry assumptions, or unit conventions), the system is more likely to produce apparently "reasonable" but incorrect steps and to omit assumptions. Numerical outputs should be treated as scaffolded rather than "linear", as flagged by the SME review.

When applying the risk matrix, risks are mostly clustered in three categories: data privacy and protection, transparency, and pedagogical impact. Privacy risk is quite elevated because automation platforms can unintentionally "leak" sensitive information through connectors, logs, intermediate files, or external APIs. Transparency risk is also elevated because it might be difficult for students to interpret limitations and for instructors to audit pipeline behavior. Finally, the pedagogical risk is elevated because the platform can indeed encourage shortcut-like behaviors unless scaffolding is provided and students are guided toward reasoning and assumptions.

Other categories, such as institutional compliance risk and ethical considerations, are moderate and depend on a specific scenario. Thus, with minimal data, institutional compliance declines, while LMS data routed through unapproved connectors may pose elevated risks. Ethical considerations can also be relative, as prompting skills might differ (e.g., expertise), leading to perceived inequities. Some recommended mitigation actions include: (i) minimal data pathways (e.g., student data anonymization, protected LMS routing, clear document retention and logging policies and practices), (ii) SME constant verification for numerical problems and periodic review of questions database, (iii) maintain/update prompt templates and content, and provide students with clear limitations on expected results, (iv) implement pedagogical scaffolding (provide several examples), and (v) standardize templates to ensure equity in terms of prompting expertise.

Although the heat conduction module presents specific risks, these can be easily extrapolated to other subjects; hence, the risk categories and scoring approach presented in this work are useful for transfer to other engineering education contexts where automation and coupled WAP/GenAI are introduced. The specific likelihood and impact values will, inevitably, vary by deployment details and risk team. Nonetheless, the evaluation's structure provides a reliable mechanism for adoption decisions across courses and departments.

The results are preliminary and primarily conceptual, relying on SME-designed testing rather than large-scale classroom trials. Therefore, learning gains are proposed rather than empirically established, and risk scores reflect expert judgment under stated assumptions only.

Future work should validate effectiveness metrics in a controlled pilot, test the platform with representative student use patterns, and refine mitigation strategies based on observed behaviors.

CONCLUSIONS

This paper presents a GenAI-aligned workflow automation platform designed to support modular, instructor-aligned learning outputs for undergraduate engineering courses. By combining validated content, assistance, and verification/monitoring, our approach aims to improve consistency and transparency when delivering outputs while reducing risks related to correctness, over-reliance, and data privacy. The framework also provides a practical pathway to document evidence aligned with CEAB Graduate Attributes.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the use of ChatGPT 5.2 Thinking for supporting the structuring of this paper's content. The conceptualization, analysis, validation, and writing have been performed by the authors.

AUTHOR IDENTIFIERS

Author ORCIDs:
Daniela Galatro: 0000-0003-0934-7492
Stuart Grey: 0000-0001-9274-0601
Sourojeet Chakraborty: 0000-0002-9461-880X

REFERENCES

1. Tuyishime A, Basciani F, Iovino L, Izquierdo JLC, Cabot J, Pierantonio A. Bridging workflow automation tools and EMF modeling ecosystems. 2023 ACM/IEEE International Conference on Model Driven Engineering Languages and Systems Companion (MODELS-C) :893-897 (2023). https://doi.org/10.1109/models-c59198.2023.00140
2. Siva Prakash Bikka . Workflow automation engines: driving innovation in cloud-native and ai-enhanced business processes. Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol 11:282-289 (2025). https://doi.org/10.32628/cseit25111227
3. Marchena Sekli G, Godo A. Transforming teaching and learning with robotic process automation: a systematic review of pedagogical applications. JITE:Research 24:039 (2025). https://doi.org/10.28945/5655
4. Abduvohid X. The role and importance of using automated design systems in the educational process. eijp 5:75-79 (2025). https://doi.org/10.55640/eijp-05-09-16
5. Hawkinson E. Is education falling into the automation abyss? why struggle and human connection is essential for learning in an ai-driven world. together 2: (2025). https://doi.org/10.62883/ijpj5817
6. Ma Y, Li X, Hu S, Liu S, Cheong KH. Trustworthy AI in education: framework, cases, and governance strategies. Innov. Emerg. Technol. 12: (2025). https://doi.org/10.1142/s2737599425500264
7. Dabis A, Csáki C. AI and ethics: investigating the first policy responses of higher education institutions to the challenge of generative AI. Humanit Soc Sci Commun 11: (2024). https://doi.org/10.1057/s41599-024-03526-z
8. Geethanjali KS, Umashankar N. Enhancing educational outcomes with explainable AI: bridging transparency and trust in learning systems. 2025 International Conference on Emerging Systems and Intelligent Computing (ESIC) :325-328 (2025). https://doi.org/10.1109/esic64052.2025.10962568
9. Vivian R. From artificial intelligence to real dummies. Digital Technol. Res. Appl. 4:1-13 (2025). https://doi.org/10.54963/dtra.v4i2.1205
10. Westover JH. AI in education: building learning systems that elevate rather than erode human capability. HCLReview 27: (2025). https://doi.org/10.70175/hclreview.2020.27.4.4
11. Gonsalves C. Generative ai's impact on critical thinking: revisiting bloom's taxonomy. Journal of Marketing Education 48:4-19 (2024). https://doi.org/10.1177/02734753241305980
12. Tian J, Zhang R. Learners' AI dependence and critical thinking: the psychological mechanism of fatigue and the social buffering role of AI literacy. Acta Psychologica 260:105725 (2025). https://doi.org/10.1016/j.actpsy.2025.105725
13. Qian K, Liu S, Li T, Rakovi? M, Li X, Guan R, Molenaar I, Nawaz S, Swiecki Z, Yan L, Gaševi? D. Towards reliable generative ai-driven scaffolding: reducing hallucinations and enhancing quality in self-regulated learning support. Computers & Education 240:105448 (2026). https://doi.org/10.1016/j.compedu.2025.105448

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Artificial Intelligence (AI) Usage in an Undergraduate Chemical Engineering Course: Strengths, Pitfalls, and Future Insights

Sourojeet Chakraborty[a*], Stuart Grey[b], and Daniela Galatro[c]

[a] Department of Chemical & Biomolecular Engineering, Johns Hopkins University, MD 21218, USA.
[b] James Watt School of Engineering, University of Glasgow, Scotland, UK.
[c] Department of Chemical Engineering & Applied Chemistry, University of Toronto, ON M5S 3E5, Canada.
* Corresponding Author: schakr34@jh.edu

ABSTRACT

As Industry 5.0 (I.D. 5.0) reshapes the engineering education landscape, Higher Education Institutes (HEIs) have evolved to integrate Generative Artificial Intelligence (GenAI) via strategic curriculum revamps to meet Education 5.0 (E.D. 5.0) competencies. EN.540.202 (Introduction to Chemical & Biological Process Analysis) is the first core course at Johns Hopkins University and was revamped in Fall 2025 to create more rigorous course content and the conscious creation of new weekly graded problem sets, which did not rely on prior course content/textbook-based solved examples. Problem sets were fed as Effective Prompt Engineering (EPE) inspired prompts to ChatGPT, and AI-elicited responses were compared. AI was able to perform fundamental calculations, offer detailed explanations, unit conversions/checks, proactive information (outside the problem scope), and graphical information. Key challenges and pitfalls observed were terminology misinterpretation, lack of visual representation, data acquisition issues, and ambiguity when tackling open-ended problem prompts, which pose a potential threat of incorrect learning if students exclusively rely on ChatGPT. Quantitative benchmark comparisons of AI performance spanning across various prompt categories indicate a significant 73% performance gap. While AI demonstrates 88% accuracy on standard problems, the success rate plunges to 15% on complex mass/energy balances involving non-ideal recycle streams and constraints. Introduction of EPE and Chain-of-Thought (CoT) strategies mitigated these risks, restoring AI accuracy to 91% in instructor-led trials. To ensure academic integrity, a dual-layer Quality Assurance (QA) audit was implemented using Turnitin's AI-writing detection; analysis shows a mean AI-probability of <12% in derivation sections, suggesting students successfully used AI for structural brainstorming, while maintaining technical ownership of engineering execution. Anonymous student-filled instructor teaching evaluations capture a strongly positive perception of preparedness, better student engagement, and recognition of potential risks of over-reliance on AI-led Large Language Models (LLMs), potentially leading to AI-led misdirected learning (unless students can distinguish between correct/incorrect AI responses). Despite the intellectual challenge of the course being rated as high, overall course quality and instructor teaching effectiveness metrics were ranked high, signalling a successful case of AI-led curriculum revamps, capturing a thoughtful integration of EPE into the undergraduate curricula at HEIs. This work provides a scalable framework for HEIs to modernize engineering pedagogy, positioning AI as a collaborative tool that augments and empowers, rather than replaces, critical human-centric engineering judgment.

Keywords: Education, Artificial Intelligence, Process Calculations, Curriculum Revamp, Society 5.0, Higher Education Institutes

INTRODUCTION

The Fifth Industrial Revolution is uniquely characterized by rapid advances in Artificial Intelligence (AI) and Machine Learning (ML), propelling global growth. This revolution has led to the creation of Industry 5.0 (I.D. 5.0), where industry skills are constantly evolving, rendering past skillsets obsolete, ushering in new ones. This has

necessitated the creation of Education 5.0 (E.D. 5.0), a unique education/academic framework, to mirror and match I.D. 5.0 demands. This has most notably been observed using strategic curriculum revamps, observed across universities and Higher Education Institutes (HEIs), especially at the undergraduate level. Such initiatives have especially ramped up since 2022, coinciding with the release of ChatGPT, swiftly followed by other AI-based search engines such as Copilot, Gemini, DeepSeek, etc. However, one of the notable (and unintended) consequences of students relying on AI is that they may use these chatbots to solve problem sets and tutorial questions. While such outcomes are appealing to a student and perhaps even beneficial on a short-term basis (a quick solution), they have long-term repercussions. Today's students already suffer from a short attention span (a consequence of frequent information consumption via reels on Instagram, TikTok, and other social media platforms), which becomes a challenge when learning and developing fundamental knowledge in core engineering fields. Specifically, within the context of chemical engineering (which is the field this paper focuses on) and in STEM (Science, Technology, Engineering, and Mathematics) fields, such short attention span is extremely counterproductive towards the development and retention of fundamental conceptual knowledge, yet such is the reality of the world we currently inhabit. In this work, we demonstrate how to work with these systemic conditions and perform a curricular revamp of a fundamental chemical engineering course, EN.540.202 (Introduction to Chemical & Biological Process Analysis), at Johns Hopkins University. As we will demonstrate, students appreciated our AI-assisted curriculum revamp and recognized the value in such initiatives. We begin by presenting a short literature review detailing the evolution of the engineering education/pedagogical literature and why this work captures one of the most nascent yet impactful initiatives for AI-assisted curriculum revamps. Next, we present the strategy of our curriculum revamp by creating novel questions that tackle fundamental mass and energy balances (the backbone for performing any process calculations for a chemical engineer). Responses generated by AI vs. instructor solutions are compared in the analysis methodology, which then paves the way towards discussions to examine the impacts of our work.

LITERATURE REVIEW

The advent of the Fifth Industrial Revolution is characterized by massive techno-economic leaps fuelled by advances in AI and ML. As these initiatives have propelled the landscape of skillsets that the industry perceives as employable, it has resulted in a scenario where new skillsets emerge as valuable, while older ones are rendered obsolete. This is termed as Industry 5.0 (I.D 5.0) and consequently, transformed the education landscape at universities and Higher Education Institutes (HEIs) to its corresponding version, Education 5.0 (E.D 5.0). E.D 5.0 [1] is uniquely characterized by the role of AI/ML in shaping and personalizing student learning experiences. Attempts have been made to map out E.D 4.0/E.D 5.0 competencies to I.D 4.0/I.D 5.0 skillsets [2], and the contributions of AI, ML, the Internet of Things (IoT), Big Data, Virtual Reality (VR), and Augmented Reality (AR) emerge as undisputed champions, catalyzing this global change. This impact is implemented by HEIs primarily through the strategy of curriculum revamps in their traditionally structured courses, as E.D 4.0 conclusively shifts to E.D. 5.0 [7]. As an example, efforts have been documented to (i) improve student perception of the Oil & Gas (O&G) Industry by updating the curriculum by talking about greener bioreactor/biofuel type initiatives by the industry [3], (ii) incorporating Agile methodology to a data-based modelling course of AI/ML [5], (iii) revamping Chemical Process Design to meet the targets of E.D 5.0 [11], and, (iv) employing variables interaction effects to educate students about lithium-ion battery ageing [12]. Such initiatives have been well received by students, highlighting the success of such initiatives if one chooses to incorporate them.

Of particular importance is the fact that, despite the world currently experiencing the Fifth Industrial Revolution, the reality of AI-led advances is more an identifier for economies that hail from the affluent Global North power bloc, as opposed to the Global South. Recent work has captured the existence of an AI implementation gap between these blocs, which further exacerbates when two Global South innovation drivers (India and China) are excluded from the analysis [6]. This lacuna needs to be resolved to unite E.D 5.0 with I.D. 5.0 [18]. Integration of these frameworks directly leads to the recognition of sustainability/sustainable development as a key trackable attribute towards economic prosperity [4, 8]. The eventual vision is to attain Society 5.0, which is characterized by a human society capable of AI-assisted advances that solve global-scale problems, including (but not limited to) the climate crisis [9, 10, 13]. Such a perfect amalgamation of I.D 5.0 and Society 5.0 [17] builds on each other, in a process called co-evolution [14]. Thus, it becomes extremely useful to identify and implement pedagogy-backed curriculum revamps to attain this vision, and incorporating AI-based teaching models is one recently discovered framework to realize this target [15]. This transformation is indeed a key limitation that present-day society is urged to proactively overcome, to shift I.D 4.0 to I.D 5.0 [16], which then ushers in opportunities towards the attainment of Society 5.0 [21]. A recent review of the literature also identifies and assesses the online environment for the attainment of E.D 5.0 [20] and

encourages instructors/faculty to navigate from traditionally siloed to more inclusive learning, shifting the focus from STEM (Science, Technology, Engineering and Mathematics) to STEAM (Science, Technology, Engineering, Arts and Mathematics) centric learning. To realize Education 5.0, AI and AI-powered Large Language Models (LLMs) based platforms such as ChatGPT, Gemini AI, Copilot, DeepSeek, etc. have emerged. Today, AI serves as a major framework to implement curricular revamps at HEIs, since it overwhelmingly encompasses a student's learning and university experience [18, 19]. In this work, we will also utilize this very strategy and employ ChatGPT as a comparative framework to compare instructor solutions. To the best of our knowledge, this is a first initiative within both the chemical engineering and engineering education/pedagogical literature, where such a cohesive AI-comparison/integration into a course is tracked.

CURRICULUM REVAMP STRATEGY

EN.540.202 (Introduction to Chemical & Biological Process Analysis), or more commonly, Mass and Energy Balances (MEB) is the first core or fundamental chemical engineering course (4 credits) taken by freshmen/sophomores at the Department of Chemical & Biomolecular Engineering at Johns Hopkins University. This mandatory course often serves as a platform to develop their fundamental understanding in chemical engineering, which later paves the way for more rigorous courses such as Transport Phenomena, Reaction Engineering & Kinetics, and Engineering/Chemical Thermodynamics. In Fall 2025, the course was offered to 47 enrolled students, separated into 4 sections of 10-13 students each. Each session had weekly tutorial sessions (primarily in-person) led by 4 Teaching Assistants (TAs), who covered problems that built on the corresponding lecture content. In addition to this, all TAs offered weekly office hours which students from any section could attend. The primary assessment of this course relied on 4 in-person quizzes and a final examination. Additionally, students worked in 5-6-member teams to calculate mass and energy balances on a simplified chemical industry process such as the Contact Process (sulfuric acid production), the Haber-Bosch Process (ammonia synthesis), the Baeyer Process (alumina production), and the Water Gas Shift Reaction (synthesizing blue hydrogen). The final course grade was a composition of multiple assessments: tutorial attendance and participation (5%), problem set submissions (20%), research presentations (7.5%), final report (15%), 4 quizzes (total 30%), and a final examination (22.5%). The curriculum revamp was specifically performed for the problem set component of this course.

For Fall 2025, this course was revamped in the following ways. Firstly, the course content was made more rigorous, with substantial questions testing the foundations of mathematics in students (especially iterative numerical methods, algebra, differential and integral calculus). The instructor taught the course content from three textbooks, as opposed to the conventional choice of using one textbook. Secondly, there was a conscious creation of new problems in quizzes and problem sets, thereby not using previous year content or solved textbook examples. This is not an easy task and constantly requires the creation, solution, and checking of questions before sending it to students. To ensure that students adhered to academic policy and integrity, a plagiarism check (5%, excluding references and bibliographic information on Turnitin) was performed on all problem sets and deliverables in this course (and students were made aware of this from the very start of the course). The instructor (Sourojeet Chakraborty) also ensured that the solutions to these problems were grounded in physical reality and representative of the physics, mathematics, and process dynamics of real-time process engineering operations, as valid within chemical engineering firms. The content of the material was also expanded, with more mass and energy balance topics/problems covered consciously, to set the foundations for engineering thermodynamics, which these students would take in Spring 2026.

ANALYSIS METHODOLOGY

Analysis of problem sets was performed to compare AI vs. human (instructor) responses. Nine such problems were formulated, each encompassing several domains and fields, all involved testing some key concept of mass and energy balance, as well as analysis and mathematical skills, as summarized in **Table 1**.

Table 1: Problem set question themes and concepts tested.

Q. #	Question Theme	Concepts Tested
Q_1	Unidirectional flows	Calculus (Differentiation and Integration)
Q_2	H_2SO_4 synthesis via the Contact Process	Stoichiometry, Generation Consumption Analysis
Q_3	Heat Exchanger Design	Algebra, Newton-Raphson (iterative numerical methods)
Q_4	Change of maple leaf colors (fall colors) at Algonquin National Park, Canada.	Writing, solving mass balances. Utilizing boundary conditions and Newton-Raphson method

Q_5	Synthesis of blue hydrogen in Abu Dhabi, United Arab Emirates	Balancing and adding chemical equations, Atom Economy, Process Economy, Degree of Freedom analysis for unsteady state mass balances
Q_6	Egyptian blue lotus and its impacts on humans	Mass balances, recycle of process streams for steady state mass balances.
Q_7	Environmental impact of Chlorinated Volatile Organic Compounds (CVOCs) on Earth	Solving thermodynamic non-ideal Equations of State (EOS). Evaluation of pressure, volume and temperature. Information utilization to evaluate unsteady mass balances.
Q_8	Dr. Jane Goodall, and her contribution to understand chimp behavior at Gombe National Park. The problem also investigates mercury accumulation by Lake Tanganyika, situated nearby.	Unit conversions (SI to FPS), unsteady state mass balances with volume to mass conversions (and vice-versa). Integral calculus, Ordinary Differential Equations (ODEs) such as the Leibnitz linear form. Enthalpy calculations with thermodynamic and steam tables based linear interpolation.
Q_9	The gold of the 18th dynasty Pharaoh King Tutankhamun, found in 1922 by Howard Carter.	Energy balances requiring enthalpy/heat calculations to synthesize the 1200 kgs. Of gold discovered in the tomb of King Tut. Vapor-Liquid Equilibria (VLE) calculations.

For each problem detailed above, the instructor solved the problem first manually and then fed it as an effective prompt into ChatGPT, leading to an AI-elicited response. In this context, it is important to recall that the extent to which an LLM-powered search engine can correctly provide a solution depends on the exactness and rigor of the search/prompt criteria fed to the AI – this is known as Effective Prompt Engineering (EPE). Responses and solutions obtained from ChatGPT were extracted and compared with the human (instructor) response, and valuable insights were obtained on what areas the AI excels at, prompts it faces challenges with, and conditions where the AI solution substantially deviates from the correct version. Such insights become extremely valuable, providing a tangible roadmap for instructors to train themselves on EPE for their courses, and design and track personalized as well as collaborative curricular revamps across HEIs. To our best knowledge, there is no evidence in the engineering education/pedagogical literature of a similar work being performed, to consciously compare and track AI vs. human responses, as well as investigate AI response success based on EPE.

RESULTS & DISCUSSION

We may now present our key results, having compared AI vs. human responses, as summarized in **Table 2**. A direct comparison of these responses reveals valuable insights into the AI's thinking and response process, its thinking patterns, and recommendations. Comparing with human responses, we see that EPE plays a strong role in impacting the correctness of the AI's outcome. This insight is useful when investigating the pros/cons of human vs. AI responses from LLM-powered search engines. The table also captures the diverse spectrum of core chemical engineering topics covered.

AI PERFORMANCE & ACADEMIC INTEGRITY

To validate the effectiveness of the revamped curriculum and address AI-proofing of core engineering problems, a systematic quantitative benchmark was implemented. ChatGPT's performance was compared across two problem archetypes: Standard Problems (single-unit balances) and AI-Proof Problems (more complex, multi-stage processes with recycles and non-ideal constraints), as summarized in **Table 3**. We observe a 73% performance gap when AI is presented with AI-proof problems without guided prompting. This compelling evidence justifies the curriculum's focus on EPE as a necessary skillset for the I.D. 5.0 future workforce. To ensure rigorous Quality Assurance (AI), all student submissions were also checked for plagiarism using a two-pronged audit strategy, using Turnitin's Similarity Match report and AI-Writing Detection tools. The average similarity score was found to be 4.2% (excluding biography and references), below the instructor's set 5% threshold for all deliverables on Canvas. 85% of student submissions in the AI-Writing audit showed an AI-written probability of <12% for the derivation/analysis sections. So, while students utilized AI for structural brainstorming, technical execution, and critical analysis remained strongly human-driven. However, we also observe instances when AI hallucinates when it incorrectly interprets the definition of 'infinity', which is different in the chemical engineering vs. the mathematical context. In chemical engineering, infinity is used to denote the far-field temporal evolution profile.

Table 2: AI vs. human-response summary for EN.540.202 problem sets in Fall 2025.

Q. No.	Human-response	AI response
Q_1	Differentiation showed with all steps correctly. Calculations are supported by commentary. Terms cancelling out are highlighted to the reader. Detailed derivation of calculus (no steps skipped) to arrive at a proper solution to equate the left- and right-hand sides. Indefinite integration. Proper boundary conditions to obtain the final integrated solution.	AI clearly shows all steps of differentiation and integration and obtains the correct answer. ChatGPT further asks the user if they wish to verify all units or plot the unidirectional flow velocity profiles.
Q_2	The first step is to balance the equations for the Contact Process to arrive at the overall balanced equation. Following this, all assumptions are clearly stated and a generation consumption table formed to explain all variables assumed. Simultaneous solution for assumed variables, following which, the generation consumption table is updated.	AI provides unit checks prior to performing the calculations. Extracts values from the question and performs straightforward correct calculations to calculate desired parameter values.
Q_3	Evaluation of Reynolds and Prandtl numbers. Uses the Dittus-Boelter equation to calculate the convective heat transfer coefficient from the Nusselt number. Calculates the Darcy-Weisbach friction factor numerically using the Colebrook-White equation. Initial guess for the friction factor is obtained from reading the Moody's chart. Newton-Raphson type calculations to compute numerically the friction factor, until values show computational convergence.	Correctly performs the integration to obtain volumetric flow rates. Compares analytical and numerical integration values and confirms equality between both these approaches.
Q_4	Block Flow Diagram (BFD) to capture the overall picture of the process. AI is unable to generate a BFD. Logically details out the steps needed to arrive at differential equations for this problem, from general mass balances. Clearly states any made assumptions. Identifies the ODE as a Leibnitz linear form, and then, correctly integrates to obtain solution profiles. Applies boundary conditions to obtain the unique solution profile. For the last part of the question, the problem obtains the numerical or transcendental equation form and applies the Newton-Raphson method till the solution converges.	AI lays out all made assumptions and correctly integrates and applies boundary conditions to obtain a unique solution profile. Identifies the last part of the question as solving a transcendental equation. Correctly performs the Newton-Raphson method to obtain the solution. Lastly, ChatGPT also offers to show graphical plots of the solution profiles or offers to show the calculations in much more detail.
Q_5	Generation consumption analysis performed to obtain the overall balanced equation for the Water Gas Shift Reaction (WGSR). Atom and process economies defined and then evaluated. This is also supplemented by a Degree of Freedom (DOF) analysis for both steady and unsteady states. Correct identification of process variables to solve for the mass of each species in the given problem.	ChatGPT correctly balances the equations to obtain the overall WGSR equation. Atom and process economies are correctly calculated (but poorly defined). The AI assumes the user to have this knowledge. Ideal Gas Law assumption made by the AI, before performing a DOF analysis for both steady/unsteady cases. However, the AI also proactively provides additional theoretical content after solving the problem (few sentences summarizing its chosen approach to solve). Lastly, the AI interface also offers to plot the obtained solution profiles, which were obtained by integrating the ODE.
Q_6	Begins by creating a Block Flow Diagram (BFD). The instructor taught a ten-point solution methodology in class and utilizes this sequence to solve this problem. A set of simultaneous linear equations are solved using matrix solution. This problem is also an example of how to set up sequential balance equations. Lastly, balances are checked and the BFD updated with the calculated mass balances, proving that the law of mass conservation holds. The operation is again repeated, this time with a ... **[row continues on next page]**	AI identifies the problem domain, states clear assumptions, and then, details out the steps of the solution. Performs correct mass balances and verifies it by checking for mass conservation. Despite the problem being clear, AI interprets the second part of the problem (with recycle line) as ambiguous. But it still makes correct assumptions and solves the mass balances correctly. Lastly, ChatGPT then ...

Instructor evaluations for this course demonstrate that students in general found this pedagogical revamp useful (see **Table 4**). Information via conversations also

Table 2 Continued: AI vs. human-response summary for EN.540.202 problem sets in Fall 2025.

Q. No.	Human-response	AI response
...Q_6	.... recycle line. Again, it is shown that the law of conservation of mass holds.	... summarizes its mass values, and provides key conclusions arising from the analysis. Lastly (and remarkably), the AI prompts the user if they wish to rerun calculations with different process conditions, or with more recycle lines.
Q_7	Starts by identifying that the Ideal Gas Law does not apply for the problem. Calculations are to be performed using the Peng-Robinson Equation of State (PR-EOS). Converts all process parameters to SI units and calculates the final form of the cubic polynomial (all non-linear equations of states can be reduced to a cubic EOS). Solves for the volume without skipping any steps. Correctly obtains the three roots, and chooses the correct one, based on physical reality and numerical feasibility. This is indicative of the conscious creation of engineering logic and judgement in students via a problem set. Using the molar volume calculated from the PR-EOS, an unsteady-state mass balance is performed to capture the solution profile Finally, the student is presented with a temporal mass accumulation profile.	The AI extracts all useful information from the question prompt and correctly performs thermodynamic calculations. The correct root is then chosen, based on proper reasoning (so, complex conjugate roots are neglected, since this is not a real solution). Provides additional theoretical content (2-3 short lines) to explain its choice and summarizes additional resources that the user can read up on. The AI is fallacious in the second part of the problem, and ends up performing an infinite summation, effectively rendering the problem as a steady-state one. This is not what the question asks, so the AI predicts an incorrect solution for this case. So, the AI incorrectly interprets a simple ODE integration as an infinite summation. If the user is unaware of these two domains, they will not be able to discern between a correct and an incorrect solution.
Q_8	Starts by identifying the Normal Temperature and Pressure (NTP) conditions. Correctly identifies the Leibnitz linear ODE form for mass conservation Integrates, and then assimilates suitable boundary conditions to obtain the mass accumulation profile in Lake Tanganyika, presented in MS-Excel. Lastly, performs calculations to demonstrate that even at low concentrations, the mercury leached into the lake can exceed the permissible environmental limits.	The AI correctly obtains enthalpy data from the formula and recognizes the ODE form correctly. Integrates correctly. Beyond the question scope, ChatGPT performs additional calculations and shows that the accumulated mercury will exceed the permissible limits in 5 hours! This is not asked in the original problem. Additionally, the AI further calculates the mass profile by performing a weighted, rather than an arithmetic average, to obtain the average mass. This is a classic case of the AI over-interpreting the search prompt and performing additionally useful (but not needed) calculations. Lastly, a mercury accumulation plot is also shown by the AI.
Q_9	Starts by drawing a BFD for the overall gold extraction process. Obtains the balanced chemical equation correctly. Employs mass balances (via back-calculation) to evaluate how much feed is needed to produce 1200 kgs. of gold. Vapor-Liquid Equilibria (VLE) information is utilized to correctly obtain the enthalpies of all species, to calculate the heat of reaction. Assuming adiabatic conditions, a fourth order polynomial is calculated, and the correct (physically and mathematically meaningful) value of the adiabatic flame temperature is presented. Finally, the process is compared to case when the specific heat capacities are constant (which would not require one to integrate).	The AI identifies the problem as one involving a complex interaction of mass balances, stoichiometry, and thermochemistry. It correctly adds equations to obtain the balanced equation for the overall process. Provides two very different interpretations of the (intentionally) ambiguous problem and proceeds with one approach. The AI summarizes key mass and energy balances at inlet/outlet conditions. Extracts the standard enthalpies of formation from the NIST database and thermodynamic tables. For one species that the AI is unable to find enthalpy data online, it assumes a plausible value. Incorrectly calculates the adiabatic flame temperature. Notably though, the AI recognizes that the calculated value does not make sense owing to having made infeasible assumptions and asks the user if they wish ChatGPT to perform more realistic, rigorous calculations. In the second iteration, the AI takes more time but correctly extracts the specific heat capacity of the species it had assumed earlier. Calculations are repeated, and this time, a correct adiabatic flame temperature is reported. Lastly, ChatGPT offers the user to perform further calculations, with an even more rigorous search on the enthalpies, and/or to summarize its calculations and results as a PDF file that might prove useful to put in a report/appendix file.

shows that students continued to use the lecture slides/content from this course as an additional resource in Spring 2026, specifically for EN.540.203 (Engineering Thermodynamics). This is direct evidence of the

attainment of higher levels of Bloom's taxonomy (content retention/replication), which serve as indicators of Fink's model of significant learning.

While variations are observed across the four sections taught in Fall 2025, mean instructor metrics (mean course quality of 4.3425, mean teaching effectiveness of 4.275, and mean intellectual challenge of 4.4425) capture a successful pedagogical revamp of this course. For context, a score of 4 out of 5 is considered a high indicator of course quality at Johns Hopkins University.

Table 3: ChatGPT performance vs. problem complexity & prompt strategy.

Problem Type	Prompt Strategy	Conceptual Accuracy (1-5)	Math Precision (%)	Success Rate
Standard	Zero-Shot (Direct)	4.2	88%	85%
AI-Proof	Zero-Shot (Direct)	1.8	32%	20%
AI-Proof	Contextual (EPE)	3.5	64%	55%
AI-Proof	Chain-of-Thought (CoT)	4.6	91%	88%

Table 4: Summary of instructor evaluation via anonymous student responses (Likert scale 1-5) for EN.540.202 (Fall 2025).

Assessment Metric	Section	This course	Department Level	School Level
Course Quality	1(N = 13)	4.25	4.14	4.22
	2(N = 12)	4.25		
	3(N = 12)	4.17		
	4(N = 10)	4.70		
Teaching Effectiveness	1(N = 13)	4.25	4.19	4.27
	2(N = 12)	4.00		
	3(N = 12)	4.25		
	4(N = 10)	4.60		
Intellectual Challenge	1(N = 13)	4.58	4.32	4.27
	2(N = 12)	4.09		
	3(N = 12)	4.50		
	4(N = 10)	4.60		

CONCLUSIONS

Several pivotal conclusions arise from direct human vs. AI responses. To begin, AI-powered LLM platforms such as ChatGPT appear to be doing a great job when solving elementary (or straightforward) mass and energy balance calculations. The success of the AI solution depends on the prompt fed to the search engine. The more rigorous the prompt, with detailed information, the more accurate the AI solution. However, as the questions become progressively more difficult, abstract, and less defined, the AI falters and sometimes fails to produce a correct solution. When this is the case, we observe that AI responses are completely different from the actual solution. This is important to note, since this is something that students should be made aware of. Therefore, the success of an AI prompt heavily relies on the skill of Effective Prompt Engineering (EPE). In fact, this work is direct evidence of why instructors are actively encouraged to teach and integrate EPE into their curricula. Since it is now well established that students use AI platforms to catalyze their personalized learning experience, learning about EPE skills (which, we propose, is an E.D. 5.0 competency) can help them to empower their learning experiences at HEIs.

Next, the long-term impacts of this curricular revamp are now becoming evident to the instructor. In Spring 2026, several of his students thanked him for teaching them these principles, especially thermodynamic concepts, which appear to be enriching their learning experience in EN.540.203 (Engineering Thermodynamics). To a faculty, this is the best reward one can expect from their students. Students also ranked the instructor's performance highly in their Fall 2025 evaluations. For context, courses are assessed on a 1-5 Likert scale rating at the Whiting School of Engineering, at Johns Hopkins University. A score of 4 or above is considered a high one and is indicative of good instructor performance. Students ranked the overall course quality as 4.3425, the instructor's teaching efficacy as 4.275, and the intellectual challenge of the course as 4.4925. This reveals that despite students rightly perceiving this course as a hard, 4-credit, fundamental course, they also appreciated the curricular revamps we implemented. Future initiatives such as these play a crucial role in bridging the E.D. 5.0 – I.D. 5.0 gap, which will directly lead to the realization of Society 5.0 worldwide.

Lastly, we wish to emphasize that, as of now, it appears that AI-led LLM-powered platforms are still not totally capable of replacing the traditional classroom-learning environment in totality. But this is not necessarily a bad outcome, since it implies that in the current context of the academic/HEI landscape, a conscious and thought-out assimilation of AI into an otherwise traditional curriculum holds potential to revamp and personalize student learning experiences to their benefit. Thus, the authors encourage the reader to envision AI from a more synergistic, rather than antagonistic, lens. When utilized proficiently and strategically, AI has the potential to rapidly transform the university learning environment, bringing it close towards E.D. 5.0. Expanding beyond the content and context of one course, if such revamps are systematically implemented throughout the curriculum, younger generations of students (Gen Z and Gen Alpha)

can swiftly be trained in E.D. 5.0, which will help them be perceived as valuable (i.e., employable) by I.D. 5.0, upon their graduation from universities. Long-term, we envision such a conscious, proactive integration of E.D. 5.0 and I.D. 5.0 to lead mankind towards accomplishing the Society 5.0 global vision.

ACKNOWLEDGEMENT

The contributions of Kevin Zhang, Arya Tulsyan, Yeonwoo Kim, and Buse Koldas, who assisted Dr. Sourojeet Chakraborty deliver EN.540.202 in Fall 2025 as Teaching Assistants, are gratefully acknowledged.

AUTHOR IDENTIFIERS

Author ORCIDs:
Chakraborty, Sourojeet: 0000-0002-9461-880X
Grey, Stuart: 0000-0001-8969-9636
Galatro, Daniela: 0000-0003-0934-7492

REFERENCES

1. Ahmad, S., Umirzakova, S., Mujtaba, G., Amin, M. S., & Whangbo, T. (2023). Education 5.0: requirements, enabling technologies, and future directions. arXiv preprint arXiv:2307.15846.
2. Chakraborty S, Gonzalez-Triana Y, Mendoza J, Galatro D. Insights on mapping industry 4.0 and education 4.0. Front. Educ. 8: (2023). https://doi.org/10.3389/feduc.2023.1150190
3. Chakraborty S, Kalhori SK, Gonzalez Y, Mendoza J, Galatro D. Student perception of sustainability in industry: a case study in an undergraduate petroleum processing course. Front. Educ. 9: (2024). https://doi.org/10.3389/feduc.2024.1396377
4. Chakraborty S, Sertu BB, Galatro D. Artificial intelligence (AI), sustainability and engineering education: implementation trends in STEM to realize society 5.0.. PCEEA : (2025). https://doi.org/10.24908/pceea.2025.19678
5. Chakraborty S, Galatro D. Incorporating agile methodologies into the chemical engineering curriculum. Comp Applic In Engineering 33: (2024). https://doi.org/10.1002/cae.22805
6. Chakraborty S, Galatro D. Artificial intelligence (AI) equality in engineering education: strategies to unite the AI gap between the global north & south. 2025 IEEE Global Engineering Education Conference (EDUCON) :1-10 (2025). https://doi.org/10.1109/educon62633.2025.11016563
7. Chakraborty S, Galatro D. From education 4.0 (E.D. 4.0) to education 5.0 (E.D. 5.0): bibliometric insights to reach the society 5.0 global vision.. 2025 ASEE Annual Conference & Exposition Proceedings : (None). https://doi.org/10.18260/1-2--56607
8. Chigbu BI, Makapela SL. AI in education, sustainability, and the future of work: an integrative review of industry 5.0, education 5.0, and work 5.0. Journal of Open Innovation: Technology, Market, and Complexity 11:100645 (2025). https://doi.org/10.1016/j.joitmc.2025.100645
9. Deguchi A, Hirai C, Matsuoka H, Nakano T, Oshima K, Tai M, Tani S. What is society 5.0?. Society 5.0 :1-23 (2020). https://doi.org/10.1007/978-981-15-2989-4_1
10. Fukuyama, M. (2018). Society 5.0: Aiming for a new human-centered society. Japan spotlight, 27(5), 47-50.
11. Galatro D, Chakraborty S. Chemical process design to meet industry 5.0 competencies. 2025 ASEE Annual Conference & Exposition Proceedings : (None). https://doi.org/10.18260/1-2--56087
12. Galatro D, Sertu BB, Chakraborty S. Teaching variables interaction effects through a battery?aging case study in undergraduate engineering. Comp Applic In Engineering 33: (2025). https://doi.org/10.1002/cae.70099
13. Gladden ME. Who will be the members of society 5.0? towards an anthropology of technologically posthumanized future societies. Social Sciences 8:148 (2019). https://doi.org/10.3390/socsci8050148
14. Huang S, Wang B, Li X, Zheng P, Mourtzis D, Wang L. Industry 5.0 and society 5.0—comparison, complementation and co-evolution. Journal of Manufacturing Systems 64:424-428 (2022). https://doi.org/10.1016/j.jmsy.2022.07.010
15. Khang A, Muthmainnah M, Seraj PMI, Al Yakin A, Obaid AJ. Ai-aided teaching model in education 5.0. Advances in Computational Intelligence and Robotics :83-104 (2023). https://doi.org/10.4018/978-1-6684-8851-5.ch004
16. Mourtzis D, Angelopoulos J, Panopoulos N. A literature review of the challenges and opportunities of the transition from industry 4.0 to society 5.0. Energies 15:6276 (2022). https://doi.org/10.3390/en15176276
17. Narvaez Rojas C, Alomia Peñafiel GA, Loaiza Buitrago DF, Tavera Romero CA. Society 5.0: a japanese concept for a superintelligent society. Sustainability 13:6567 (2021). https://doi.org/10.3390/su13126567
18. Rane N, Choudhary S, Rane J. Education 4.0 and 5.0: integrating artificial intelligence (AI) for personalized and adaptive learning. SSRN Journal : (2023). https://doi.org/10.2139/ssrn.4638365

19. Rane N. Chatgpt and similar generative artificial intelligence (AI) for smart industry: role, challenges and opportunities for industry 4.0, industry 5.0 and society 5.0. SSRN Journal : (2023). https://doi.org/10.2139/ssrn.4603234
20. Sertu BB, Chakraborty S, Galatro D. Assessing the online learning experience in the undergraduate engineering education domain: challenges and opportunities. Tech Know Learn : (2025). https://doi.org/10.1007/s10758-025-09917-8
21. Tyagi AK, Lakshmi Priya R, Mishra AK, Balamurugan G. Industry 5.0. Privacy Preservation of Genomic and Medical Data :409-432 (2023). https://doi.org/10.1002/9781394213726.ch17

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

LLM-Based Intelligent Data Extraction System for Industrial Equipment

Zean Chen[a], Kaicheng Song[a], Lingyu Zhu[b], Anjan Kumar Tula[a]* and Xi Chen[a]

[a] State Key Laboratory of Industrial Control Technology, College of Control Science and Engineering, Zhejiang University, Hangzhou 310027, Zhejiang, China
[b] College of Chemical Engineering, Zhejiang University of Technology, Hangzhou 310014, Zhejiang, China
* Corresponding Author: anjantula@zju.edu.cn.

ABSTRACT

Data extraction and processing constitute the cornerstone of quantitative management and operational analysis in industrial process plants. However, most manufacturing facilities currently lack efficient data extraction systems, relying instead on engineers to manually write and execute database queries, which is time-consuming, error-prone, and inflexible when handling diverse data formats or large-scale equipment networks. To address these limitations, this work presents a novel Large Language Model (LLM)-based intelligent framework designed for data extraction and basic data of industrial equipment. The system integrates natural language understanding capabilities with process database schemas, enabling users to perform complex data queries and analyses through natural language prompts. Specifically, it can perform data mining, time-dependent analyses, equipment comparisons, and cross-period performance evaluations when process information is provided. By integrating the process context, the system can further interpret the analysis results to suggest potential operational conditions or abnormalities occurring in the system. Developed in collaboration with a process industry partner, the proposed system was evaluated using six months of real-world operational data from 10 heat exchangers. Experimental results demonstrate that the framework achieves an extraction accuracy exceeding 99% in the online mode and ranging from 92% to 96% in the privacy-preserving offline mode, with a reduced retrieval latency of approximately 2.64 seconds for local inference. Unlike generic chatbots, this framework targets process-specific reasoning, acting as a cognitive assistant that empowers users to interact intelligently with industrial systems. The findings confirm that LLM-assisted extraction effectively addresses the rigidity and inefficiency of conventional approaches while maintaining high reliability and scalability for industrial analytics.

Keywords: Large Language Model, exchangers, data extraction, prompt

INTRODUCTION

The pace of digitalization across various industrial sectors has positioned Large Language Models (LLMs) as a transformative technology capable of driving fundamental changes in Process Systems Engineering (PSE) and manufacturing operations. Following the public release of ChatGPT by OpenAI in late 2022, LLMs have attracted widespread attention due to their capabilities in comprehending, reasoning over, and generating sophisticated natural language. When combined with complementary technologies such as external tools, databases, and domain-specific knowledge bases, these models enable novel research directions and practical applications across diverse domains.

Nevertheless, the deployment of LLMs in industrial environments remains at an early stage and encounters considerable challenges. Process industries are characterized by wide variations in unit operations, safety regulations, and exceptionally stringent requirements for reliability and precision; any erroneous prediction or misinterpreted analysis can result in severe economic losses, safety incidents, or environmental damage. One of the most prominent limitations of current LLMs is their ten-

dency to produce “hallucinations, ” i.e., plausible-sounding but factually incorrect or fabricated outputs. To counteract this issue, researchers have explored several mitigation strategies, including fine-tuning on industry-specific datasets [1], retrieval-augmented generation (RAG) [2], integration with structured knowledge graphs for fact-checking and long-term contextual grounding, and hybrid neuro-symbolic approaches that combine language models with deterministic reasoning engines.

To date, most LLM-based research in process systems engineering and related fields has concentrated on qualitative and knowledge-intensive tasks, such as professional question-answering systems, fault diagnosis support, process flowsheet generation, reaction mechanism prediction, molecular or catalyst design assistance, and educational tools for thermodynamics or unit operations [3–6]. Notable examples include cognitive assistants for troubleshooting in the textile industry using GPT-3.5 [3], in-context learning to support user queries [4], root-cause analysis in semiconductor fabrication [5] and multimodal reasoning frameworks such as PaLM-E [6].However, before moving to the reasoning and analysis phase, it is essential to have a precise and flexible data extraction system. In most manufacturing facilities, there is a high daily demand for data extraction across various process parameters and timeframes. This necessitates dynamic data retrieval and basic analyses, such as data anomaly detection and trendline identification. For now, contemporary LLMs are not inherently equipped to directly ingest or query large volumes of raw industrial time-series data stored in relational databases. Engineers and operators must still rely on complex database systems that require advanced SQL proficiency or specialized software interfaces. This technical barrier significantly increases the time and effort needed for data extraction, delays insight generation, and ultimately hampers agile, data-informed decision-making on the plant floor.

The present work seeks to close this gap by leveraging the natural-language understanding and generation strengths of LLMs to create intuitive interfaces between process engineers and industrial databases. Rather than requiring users to write structured queries or navigate rigid dashboards, the system allows practitioners to interact with operational data through ordinary conversational prompts. Specifically, We focus on the data extraction process for industrial equipment, selecting the heat exchanger as a primary case study with real industrial data.

The objective of this research is therefore to enable process engineers, shift supervisors, and plant managers to perform complex data retrieval, trend analysis, cross-unit comparisons, and basic performance diagnostics far more efficiently than with conventional methods, ultimately improving real-time supervision, early anomaly detection, and overall production effectiveness. Moreover, the proposed system architecture is intentionally designed to be modular and generalizable, offering a reusable framework that can be adapted to quantitative data challenges involving other unit operations (distillation columns, reactors, compressors, dryers, etc.). The principal contributions of this study are as follows:

1. Dual-Mode Architecture: A flexible system design that supports both an Online Mode (utilizing high-capacity cloud-based LLMs for superior reasoning depth) and an Offline Mode (local deployment of smaller open-source models to guarantee strict data privacy and control in sensitive industrial settings).
2. Robust Prompt Engineering with Self-Correction: Specialized prompting strategies, including a closed-loop self-correction mechanism for SQL generation, that substantially reduce hallucinations and improve reliability even when using resource-constrained local LLMs.
3. Systematic Decomposition of Complex Analytical Queries: A structured methodology that breaks down nested, multi-step, and interdependent engineering questions into simpler, sequentially solvable subtasks, thereby enabling accurate handling of realistic, real-world process data analysis scenarios.

METHODOLOGY AND SYSTEM FRAMEWORK

System Architecture

The overall system design prioritizes flexibility and adaptability through a dual-mode architecture that incorporates distinct Online and Offline modes. This approach effectively addresses the diverse requirements of industrial applications, balancing high computational performance against strict data privacy constraints while supporting different data storage formats. Users can explicitly select the desired mode depending on their operational priorities and security policies.

In the Online Mode, the system leverages cloud-based APIs to access high-performance Large Language Models, such as GPT-4o. This configuration minimizes local hardware requirements and takes full advantage of the model’s extensive 128k token context window, enabling it to process highly complex instructions and maintain coherent, long-running conversational histories. By contrast, the Offline Mode is essential for industrial environments that enforce strict data sovereignty and prohibit any transmission of proprietary process data to external networks. In this mode, open-source Large Language Models are deployed and executed entirely on local workstations, ensuring complete isolation of sensitive information. However, local deployment is inherently lim-

ited by available hardware resources. For example, a typical industrial workstation equipped with 16–32 GB of system RAM and a mid-to-high-end consumer GPU (providing 8–24 GB of VRAM) is generally restricted to running quantized 7B-parameter models (commonly 8-bit quantization). While quantization enables feasible inference on such hardware, it inevitably reduces reasoning depth and makes these models less capable of handling complicated, multi-step prompt logic compared to their full-precision, cloud-based counterparts.

To mitigate the performance disparity between the two modes, the system implements two complementary function-calling protocols: a JSON-based format and an SQL-based format. The JSON format supports rich, hierarchical, and expressive function definitions with detailed parameters and intermediate reasoning steps, but it requires strong semantic understanding and precise output formatting from the model. The SQL format, on the other hand, provides a more rigid, structured, and directive interface that aligns closely with predefined database schemas, making it significantly easier for smaller-parameter local models to generate correct and executable statements. Accordingly, the system defaults to the flexible JSON-based protocol in the high-performance Online Mode, while the Offline Mode defaults to the SQL-based protocol to better match the capabilities of constrained local models. A fallback mechanism is integrated to automatically attempt the alternative protocol if the default fails (e.g., due to syntax errors, semantic misalignment, or execution failure). Users also retain the ability to manually override the default and select their preferred processing method at runtime. The overall system architecture is illustrated in Figure 1 and is organized into three tightly integrated layers:

1. Storage Layer: Manages the underlying operational data, supporting both lightweight SQLite databases for structured time-series records and Excel files for legacy or ad-hoc datasets commonly found in industrial settings.
2. Data Processing Layer: Acts as the central orchestration engine of the system (detailed in Figure 2). This layer receives natural-language queries, translates them into structured executable statements according to the active mode and protocol, invokes the appropriate data retrieval and computation functions, processes the returned results, and generates coherent natural-language responses enriched with engineering context and derived insights.
3. Presentation Layer: Serves as the user-facing interface, enabling text query input, audio query input, mode selection and display of conversational history. This layer ensures intuitive interaction and full traceability of the analytical process.

This modular and layered design guarantees scalability, maintainability, and robustness, allowing the system to deliver efficient, personalized, and context-aware data analysis tailored to the complexities of real-world industrial processes.

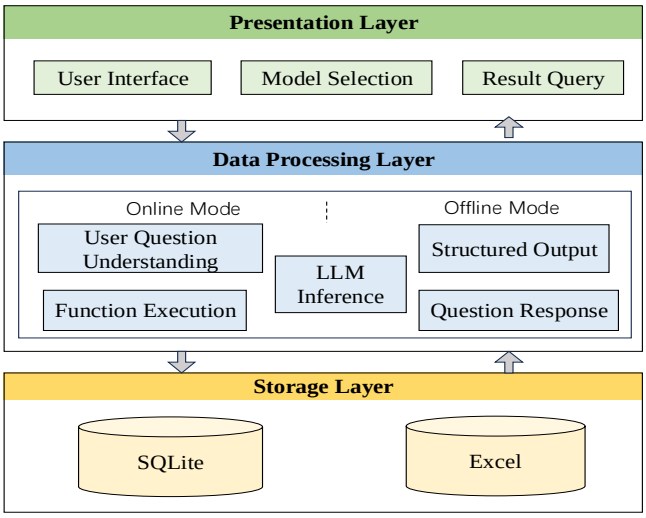

Figure 1. The architecture of the system framework.

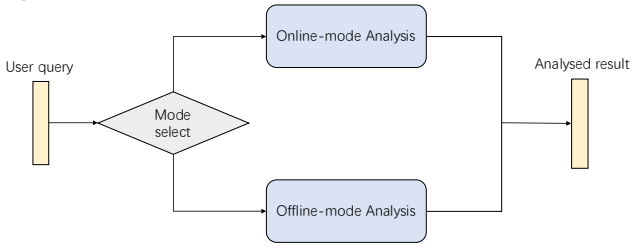

Figure 2. The flowsheet of the Data Processing Layer.

Prompt engineering

In the domain of PSE, effectively bridging the gap between complex industrial queries and precise machine execution requires robust prompt engineering. Standard instructions often fail to capture the strict logical dependencies and engineering-specific constraints inherent in process parameters. To overcome this limitation, the proposed system implements targeted prompting strategies designed to structure and guide the model's reasoning process in a manner that mirrors the systematic workflow of experienced process engineers.

For complex analytical tasks, the system employs Chain-of-Thought (CoT) prompting [7], which compels the model to explicitly illustrate its intermediate reasoning steps before producing a final answer. For example, when diagnosing a potential heat transfer anomaly, the model is guided through a logical sequence such as: 1. Retrieve and verify inlet and outlet temperatures; 2. Calculate the temperature difference; 3. Compare the result against predefined operational limits; 4. Identify any deviations and their possible implications.

This step-by-step reasoning closely replicates the

diagnostic approach of a human process engineer, significantly reducing logical errors in performance analysis, improving result reliability, and ensuring full transparency of the decision-making process for operators and supervisors.

In addition, the system utilizes Decomposed Prompting to effectively handle nested and interdependent engineering queries [8]—for instance, "Find the maximum temperature difference during the period of highest flow rate." A single, monolithic prompt frequently confuses the model in such cases due to the entangled dependencies. To address this, the framework breaks the complex request into independent, sequentially solvable sub-tasks. In the above example, the system first instructs the model to: "Identify the time period during which the flow rate reaches its maximum value." The output of this sub-task (the identified time window) is then passed as a precise constraint to the subsequent prompt: "Within the identified time period, compute the maximum temperature difference across the heat exchanger."

By isolating each engineering sub-task and enforcing sequential dependency, the system achieves high accuracy for individual variables before synthesizing them into a final, comprehensive report. This decomposition strategy proves particularly effective for managing the intricate, interdependent data relationships that are characteristic of chemical process operations.

Dual-Mode Strategy

The framework adopts a dual-mode strategy that tailors the data retrieval and execution mechanism to the specific inferential capabilities and constraints of the underlying language model. In the Online Mode, the system employs a JSON-formatted function-calling protocol optimized for high-performance cloud-based models. This structured format is inherently well-suited to Chain-of-Thought architectures, as it enables reasoning steps and intermediate results to be explicitly encapsulated within distinct key-value pairs. By mapping logical sub-steps to clearly defined JSON fields, the schema enforces systematic decomposition of complex industrial queries into discrete, interpretable units, greatly enhancing both the precision of the analysis and the transparency of the reasoning trace. Moreover, a specialized placeholder mechanism allows the output of one function call to be directly passed as input to a subsequent operation, thereby supporting the multi-hop reasoning required to resolve deeply nested engineering questions. The detailed patten is shown on the Figure 3.

Conversely, the Offline Mode relies on a SQL-based approach, which provides a more rigid and prescriptive structural framework better aligned with the syntactic and reasoning limitations of smaller, locally deployed models. Because local models frequently struggle with precise SQL syntax generation, especially under quantization and memory constraints, the system incorporates a robust self-correction mechanism (illustrated in Figure 4). When an initial SQL query fails to execute successfully against the database, the system automatically captures the exact error message returned by the SQLite engine and feeds it back to the model together with the original invalid query and the user's natural-language intent. This closed-loop feedback enables iterative refinement: the model is explicitly instructed to focus solely on correcting syntax and schema errors while preserving the analytical objective. The process repeats until a valid, executable SQL statement is produced or a predefined maximum number of correction attempts is reached. This mechanism substantially increases reliability in resource-constrained environments, ensuring dependable data retrieval even when using 7B-parameter models with limited reasoning depth.

```
{ "Action": <Function Name 1>,
 "Parameters": {
   <arg_1>: <value from user query>,
   <arg_2>: <value from user query>},
 "label": <label1_name>ش
  "equipment_id": <equipment_id>},
{ "Action": <Function Name 2>,
 "Parameters": {
   <arg_1>: <value from user query>,
   <arg_2>: $<label1_name>$},
 "label": <label2_name>,
  "equipment_id": <equipment_id>}
```

Figure 3. An example of how to map sub-steps via function calling. The second function takes the output of the first function and utilizes a placeholder during the return process, which is then executed by the code to replace the placeholder with the actual value.

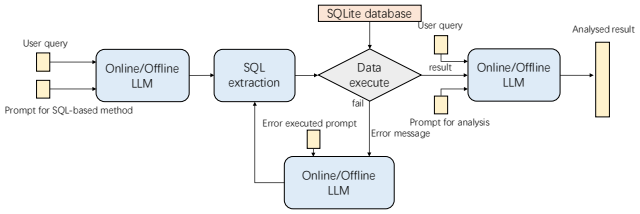

Figure 4. Framework design for SQL-based data analysis.

System Interface

The user interface is carefully engineered to deliver an intuitive, efficient, and accessible experience for industrial operators, engineers, and managers. It is implemented as a desktop application using the Qt framework and Python, providing a robust and cross-platform primary interaction point for the system.

A key design goal is to enable seamless mode

switching: users can toggle between Online and Offline computation modes directly from the main interface, depending on immediate operational needs, such as prioritizing maximum reasoning capability in non-sensitive settings or enforcing strict data sovereignty during critical production periods. To accommodate diverse working environments (e.g., noisy control rooms, hands-free operation, or office-based analysis), the interface supports dual-channel input methods: standard keyboard-based text entry for precise query formulation, and voice input powered by the Whisper automatic speech recognition model, allowing operators to issue natural-language commands hands-free while monitoring equipment or performing field tasks. The resulting interface is illustrated in Figure 5.

The interface maintains a persistent, scrollable chat-style history that displays both user queries and the system's full responses in chronological order. This conversational log preserves context across extended analytical sessions, enables users to revisit earlier results, and supports verifying of the interaction flow—an important requirement for traceability and validation in industrial decision-making processes. To further enhance system robustness, the interface is equipped with error-handling capabilities that automatically detect out-of-scope queries and provide corrective prompts to users.

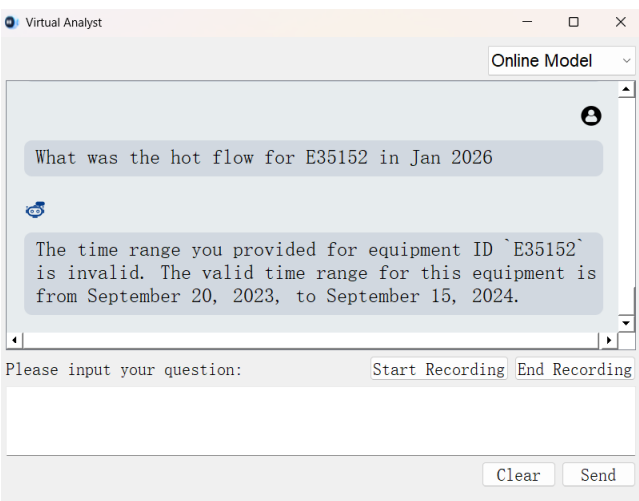

Figure 5. The interface of the data extraction system. This figure shows an example of the system's error-handling capability when encountering an out-of-bounds timeframe request.

INDUSTRIAL CASE STUDY: HEAT EXCHANGER PROCESS DATA

Dataset

The proposed system was rigorously validated using a comprehensive, real-world operational dataset obtained through close collaboration with a plant in process industries. The dataset spans six consecutive months of continuous production records and covers ten distinct shell-and-tube heat exchanger units operating under typical plant conditions. The dataset integrates key measured thermodynamic and process parameters, including: Inlet and outlet temperatures (hot and cold streams), Mass flow rates of both process fluids, Derived performance indicators, such as the Overall Heat Transfer Coefficient (U) and Logarithmic Mean Temperature Difference (LMTD).

All records are stored in a structured relational database environment (SQLite) designed to faithfully replicate the schema and access patterns of a standard industrial information management system. This setup ensures that evaluation results reflect realistic query execution behavior encountered in actual process plants. To guarantee reproducible and representative performance measurements for the Offline Mode under practical hardware constraints, all local inference experiments were conducted on a high-end workstation equipped with an NVIDIA RTX 5090 GPU (32 GB VRAM) and sufficient system RAM to support quantized 7B-parameter models without excessive swapping. This configuration represents a realistic upper bound for on-premises deployment in many industrial facilities as of 2026.

To systematically assess the system's natural-language understanding, logical reasoning, and data extraction accuracy, a dedicated benchmark dataset was constructed comprising 300 natural-language queries. These queries were carefully stratified into three complexity categories, with 100 queries per category, to probe different aspects of model capability:

1. Simple queries: Direct, single-parameter retrieval tasks (e.g., "What is the average cold-side outlet temperature for exchanger E35112 during January?").
2. Parallel queries: Requests that require simultaneous extraction and reporting of multiple independent parameters without interdependency (e.g., "Report the average flow rate, inlet temperature, and overall heat transfer coefficient for unit E35142 over the past month.").
3. Nested queries: Multi-step reasoning tasks in which the result of one computation serves as a constraint or input for a subsequent operation (e.g., "Count the days in January where the UA for E35112 is greater than the average.").

This stratified benchmark provides a robust and comprehensive framework for evaluating not only overall extraction accuracy and inference latency, but also the effectiveness of the self-correction mechanism and error-handling behavior across varying levels of analytical difficulty.

Conversational Examples

To evaluate the practical utility of the model, we utilized the aforementioned benchmark dataset consisting of 300 test cases, which are categorized into simple, parallel, and nested queries. Our system is designed to support four core operational capabilities that are essential for plant personnel to flexibly extract and analyze data: data retrieval, trend analysis, cross-unit comparisons, and basic performance diagnostics. Focusing on these four practical aspects, we conducted a series of case studies. The representative conversational interactions between the user and the system are illustrated in Figures 6, 7, 8 and 9.

1. Data Retrieval: This foundational capability enables users to extract specific historical or real-time process parameters without requiring knowledge of database structures or query languages. Shift supervisors can quickly obtain current operating values for critical equipment, while process engineers can retrieve statistical summaries such as averages, maxima, or cumulative totals over custom time windows. Beyond simple lookups, the system supports conditional queries (e.g., periods where a parameter exceeded a threshold) and arithmetic combinations of multiple variables, allowing users to perform ad-hoc calculations on the fly. By transforming manual, error-prone reporting into instantaneous, on-demand access, this capability accelerates routine monitoring, daily reporting, and troubleshooting investigations.
2. Trend Analysis: Beyond single-point extraction, the system allows users to examine how process variables evolve over time. Operators can identify gradual performance drifts such as fouling trends in heat exchangers by querying time-series patterns of key indicators like overall heat transfer coefficient or temperature differentials. Engineers can analyze cyclical behaviors (e.g., daily or batch cycles) or compare distinct operational phases (e.g., before and after maintenance, or across different feedstock grades) to assess intervention effectiveness. The system can also compute rates of change, detect monotonic trends, and flag sustained deviations from baseline behavior. This capability supports proactive decision-making by transforming raw time-series data into actionable insights about process stability, equipment degradation, and long-term performance trajectories.
3. Cross-Unit Comparisons: To support system-level oversight, the system enables simultaneous evaluation of performance across multiple equipment units. Process engineers can benchmark the efficiency of parallel units to identify underperforming assets, validate consistency across identical process lines, or assess the impact of differing maintenance schedules. Plant managers may request comparative summaries such as top-performing units by efficiency or units with the highest deviation from setpoints to prioritize maintenance activities and allocate resources effectively. The system can also aggregate data across groups of units (e.g., all exchangers in a specific area) to provide fleet-level statistics. By abstracting the complexity of multi-table joins and conditional aggregations, this capability facilitates holistic monitoring and supports data-driven decisions about equipment utilization, operational parity, and capital planning.
4. Basic Performance Diagnostics: As a practical extension of data extraction, the system provides preliminary diagnostic support through univariate anomaly detection. Operators can swiftly identify when a measured parameter deviates beyond established statistical thresholds (e.g., three-standard-deviation limits), flagging potential malfunctions without requiring advanced analytical training. Engineers can investigate the timing, duration, and magnitude of such deviations to correlate them with external events, operational changes, or emerging fault patterns. The system can also highlight recurring anomalies across multiple units or time periods, helping to prioritize deeper root-cause analysis. While not a substitute for detailed modeling, this capability serves as an effective tool, enabling rapid detection of abnormal conditions and focusing expert attention on areas most likely to require intervention.

Together, these four capabilities exemplified in the representative conversational interactions shown in Figures 6 (data retrieval), 7 (trend analysis), 8 (cross-unit comparisons), and 9 (basic performance diagnostics) demonstrate how the system bridges the gap between complex industrial data stores and the diverse informational needs of plant personnel, enabling more agile, informed, and collaborative operational decision-making.

RESULTS AND DISCUSSION

Performance Evaluation

We conducted tests in both online and offline modes to verify the system's robustness across different operational scenarios. For this evaluation, two primary metrics were established: Execution Accuracy and Inference Latency. Accuracy is emphasized as the paramount indicator, given that in safety-intensive process industries, any hallucination or incorrect data retrieval can lead to flawed decision-making and potential operational hazards. Meanwhile, Inference Latency is critical to assessing the system's viability for real-time interaction, particularly when deployed on resource-constrained local hardware for immediate anomaly detection.

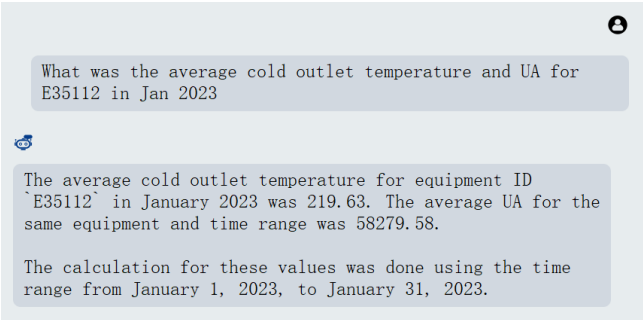

Figure 6. Dialog example of Data Retrieval.

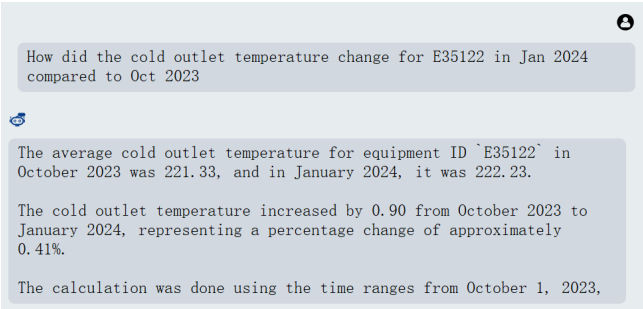

Figure 7. Dialog example of Trend Analysis.

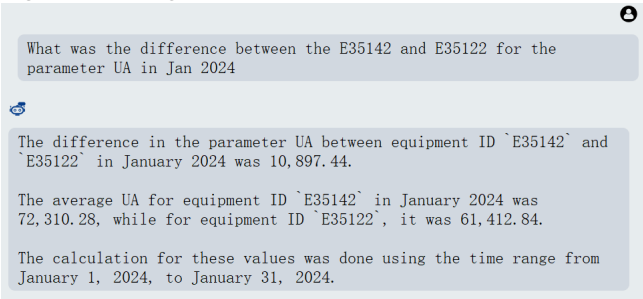

Figure 8. Dialog example of Cross-Unit Comparisons.

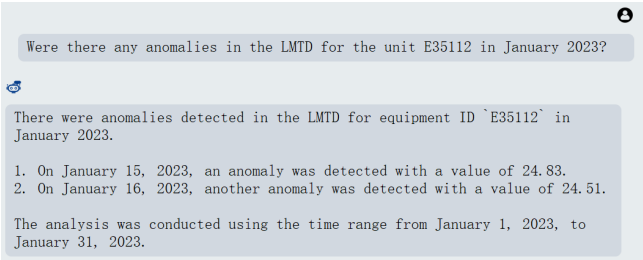

Figure 9. Dialog example of Basic Performance Diagnostics.

To have a more quantitative result, A benchmarking was conducted across Simple, Parallel, and Nested query categories. As detailed in Tables 1 and 2, the Online Mode achieved near-perfect accuracy, exceeding 99% across all query types. This superior generalization is attributed to the high-parameter count of the cloud model, which effortlessly managed complex nested logic. However, this high fidelity involved a latency trade-off, with average response times ranging from 3.90s to 5.09s due to network overhead. Conversely, the Offline Mode delivered commendable performance, with accuracy rates ranging from 92% for simple queries to 96% for parallel queries. Notably, the local deployment significantly outperformed the online configuration in processing speed, averaging an inference time of 2.64s for simple tasks. This characteristic makes the Offline Mode highly suitable for time-sensitive edge applications where data privacy and rapid response are paramount. The primary sources of error in the offline setting were identified as minor formatting inconsistencies and occasional parameter hallucinations (e.g., incorrect column name generation), which the self-correction mechanism largely mitigated but did not entirely eliminate.

Table 1: Evaluation results across various query categories in Online Mode.

Dataset	Accuracy	Average time cost
Simple question	100%	3.90s
Parallel question	99%	5.09s
Nested question	99%	4.96s

Table 2: Evaluation results across various query categories in Offline Mode.

Dataset	Accuracy	Average time cost
Simple question	92%	2.64s
Parallel question	96%	2.95s
Nested question	93%	5.19s

Furthermore, to scientifically validate the contribution of specific prompt engineering components, an ablation study was performed using the Online Mode as a baseline. The results, presented in Table 3, demonstrate that the full prompt configuration yields the optimal balance of precision and efficiency.

1. Chain-of-Thought Guidance: Removing CoT resulted in a moderate increase in inference time (approx. 60%) and a marginal dip in accuracy. This suggests that explicit reasoning steps are crucial for streamlining the model's internal logical processing.
2. Few-Shot Examples: The removal of few-shot demonstrations led to catastrophic performance loss, especially for Nested queries where accuracy hit 0%. This confirms that few-shot examples are essential for teaching the complex JSON schemas required for multi-hop reasoning.
3. Rules and Constraints: While removing explicit constraints did not severely impact accuracy, it significantly increased latency. Without predefined bound-

Table 3: Ablation study results quantifying the individual contributions of Chain-of-Thought (CoT), Few-Shot prompting, and Constraints to the system's overall performance in Online Mode.

Prompt config	Question Type	Accuracy	Format Acc.	Function Acc.	Time Acc.	Parameter Acc.	Average time cost
Full prompt	Simple	100%	100%	100%	100%	100%	3.90s
	Parallel	99%	100%	100%	99%	100%	5.09s
	Nested	99%	100%	100%	99%	100%	4.96s
w/o CoT	Simple	93%	100%	100%	93%	100%	4.62s
	Parallel	94%	100%	100%	94%	100%	8.29s
	Nested	96%	100%	100%	96%	100%	9.15s
w/o Few-Shot	Simple	53%	55%	100%	98%	100%	6.51s
	Parallel	60%	63%	100%	97%	100%	7.07s
	Nested	0%	0%	100%	98%	75%	8.22s
w/o Rules	Simple	99%	100%	100%	99%	100%	5.60s
	Parallel	99%	100%	100%	99%	100%	11.54s
	Nested	98%	100%	100%	98%	100%	9.17s

aries, the model required more iterative "thought" cycles to arrive at the correct structured output.

Comparison with Generic Frameworks

To further validate the necessity of the proposed specialized architecture, a benchmarking study was conducted against LangChain [9], a widely adopted framework for building LLM applications. To ensure a fair comparison, the generic framework was evaluated under identical hardware constraints using the same Mistral-7B model deployed in the Offline Mode. The comparative analysis reveals that the proposed direct-calling framework significantly outperforms the generic LangChain implementation in both execution accuracy and inference efficiency.

Initial experiments utilizing LangChain's default SQL chain resulted in a critical failure to generate valid executable queries, yielding a 0% accuracy rate across all test categories (Table 4). This failure was primarily driven by dialect mismatches, where the generic prompts defaulted to standard SQL syntax incompatible with the strict requirements of the local SQLite environment. Additionally, the generic approach induced frequent schema hallucinations, where the model attempted to query non-existent columns or functions prohibited by the local schema definition.

To isolate the impact of prompt quality, the robust, optimized prompts developed for our system were subsequently injected into the LangChain pipeline. Despite this enhancement, performance remained significantly suboptimal, with accuracy rates failing to exceed 50% even for simple queries (Table 5). This persistent performance gap suggests that LangChain's internal abstraction layers—which often encapsulate user instructions within generic system roles or append significant token overhead—tend to dilute the strict instruction adherence required by smaller, parameter-constrained local models. In contrast, our proposed framework's direct "Generation-Correction" loop provides the fine-grained control necessary to enforce syntax constraints on 7B-parameter models.

Table 4: Evaluation results based on the LangChain framework using its default SQL prompt templates.

Dataset	Accuracy	Average time cost
Simple question	0%	4.46s
Parallel question	0%	5.35s
Nested question	0%	5.54s

Table 5: Evaluation results based on the LangChain framework using self-designed prompts.

Dataset	Accuracy	Average time cost
Simple question	47%	5.60s
Parallel question	9%	4.34s
Nested question	22%	5.05s

Furthermore, regarding computational efficiency, the LangChain implementation demonstrated an average inference latency approximately 50% higher than our custom framework. These results underscore that while generic frameworks aid rapid prototyping, specialized industrial edge applications necessitate the bespoke optimization presented here to ensure reliability and real-time performance.

CONCLUSIONS

This study addresses the critical bottleneck of manual data extraction in process industries by introducing a novel, LLM-based Intelligent Analysis System specifically tailored for heat exchanger. Through a flexible Dual-Mode Architecture, the proposed framework effectively resolves the trade-off between complex reasoning capabilities and strict data privacy requirements.

The Online Mode, leveraging JSON-based Chain-of-Thought prompting, demonstrated exceptional fidelity in handling nested engineering queries, achieving over 99% accuracy. Conversely, the Offline Mode ensured complete data sovereignty while delivering reliable performance (92–96% accuracy) and superior real-time response (~2.64s). This was achieved through a specialized SQL "Generation-Correction" mechanism designed to mitigate the hallucinations common in smaller local models. Furthermore, comparative benchmarking against LangChain confirmed the necessity of this specialized architecture, as the proposed system significantly outperformed the generic framework in both syntax validity and inference efficiency under identical hardware constraints.

Beyond immediate technical performance, this work has broader implications for the democratization of industrial data analytics within the context of Industry 5.0. By abstracting complex database interactions into intuitive natural language, the system serves as a domain-specific cognitive assistant that significantly lowers the technical barrier for process engineers and operators. This capability not only facilitates efficient knowledge transfer and workforce empowerment but also bridges the critical gap between raw operational data and actionable decision-making. Ultimately, this framework provides a scalable foundation for human-AI collaboration, paving the way for future integration with digital twins and autonomous anomaly detection systems across the process industry.

ACKNOWLEDGEMENTS

The financial support of the "Pioneer" and "Leading Goose" R&D Program of Zhejiang (Grant No. 2024C01028) is gratefully acknowledged.

AUTHOR IDENTIFIERS

Author ORCIDs:
Zean Chen: 0009-0003-6972-8155
Kaicheng Song: 0009-0002-3812-9582
Xi Chen: 0000-0002-0571-6177
Lingyu Zhu: 0000-0002-3412-3805

REFERENCES

1. Naveed H, Khan AU, Qiu S, et al. A comprehensive overview of large language models. arXiv preprint arXiv:2307.06435 (2023) https://doi.org/10.48550/arXiv.2307.06435
2. Lewis P, Perez E, Piktus A, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. Adv Neural Inf Process Syst 33:9459-9474 (2020) https://doi.org/10.48550/arXiv.2005.11401
3. Kernan Freire S, Foosherian M, Wang C, Niforatos E. Harnessing large language models for cognitive assistants in factories. Proceedings of the 5th International Conference on Conversational User Interfaces :1-6 (2023). https://doi.org/10.1145/3571884.3604313
4. Kernan Freire S, Wang C, Foosherian M, et al. Knowledge sharing in manufacturing using large language models: User evaluation and model benchmarking. arXiv preprint arXiv:2401.05200 (2024) https://doi.org/10.48550/arXiv.2401.05200
5. Ezukwoke K, Hoayek A, Batton-Hubert M, Boucher X, Gounet P, Adrian J. Big GCVAE: decision-making with adaptive transformer model for failure root cause analysis in semiconductor industry. J Intell Manuf 36:2423-2438 (2024). https://doi.org/10.1007/s10845-024-02346-x
6. Driess D, Xia F, Sajjadi MSM, et al. PaLM-E: An embodied multimodal language model. Proc Mach Learn Res 202:8469-8488 (2023) https://doi.org/10.48550/arXiv.2303.03378
7. Wei J, Wang X, Schuurmans D, et al. Chain-of-thought prompting elicits reasoning in large language models. Adv Neural Inf Process Syst 35:24824-24837 (2022) https://doi.org/10.48550/arXiv.2201.11903
8. Khot T, Trivedi H, Finlayson M, et al. Decomposed prompting: A modular approach for solving complex tasks. Proc Int Conf Learn Represent (2023) https://doi.org/10.48550/arXiv.2210.02406
9. LangChain AI. LangChain: Build context-aware reasoning applications. https://github.com/langchain-ai/langchain (2025)

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

The Imperial College Integrated Design Project

Paul S. Fennell[a]*, Klaus Hellgardt[a], and Daniel R. Lewin[b]*

[a] Chemical Engineering, Imperial College, London, UK.
[b] Chemical Engineering, Technion, Haifa, Israel
* Corresponding Author: p.fennell@imperial.ac.uk

ABSTRACT

The Imperial College Integrated Design Project reframes the chemical engineering capstone as a structured educational journey that develops professional competence rather than simply delivering a final technical report. The programme is grounded in four pedagogical pillars—authenticity, integration, impact, and reflection—which align with the graduate attributes required by the Institution of Chemical Engineers. Authenticity is achieved through open-ended problems drawn from industrial partners and emerging research needs; integration connects knowledge from across the curriculum into a coherent systems perspective; impact emphasises user-centred, sustainable solutions; and reflection cultivates metacognitive awareness of decision making and learning from failure. A mentored-autonomy model supports student teams through weekly checkpoints, skills workshops, and access to disciplinary experts. Assessment deliberately balances artefact quality with evidence of process, rewarding reasoning under uncertainty, ethical judgement, and communication alongside technical performance. Ethics, safety, and sustainability are embedded through lifecycle analysis and HAZOP requirements rather than treated as add-ons. The four-phase structure—scoping, feasibility, detailed design, and deployment—scaffolds increasing complexity while maintaining measurable progress and accountability. A 2024 project on renewable-powered green ammonia illustrates how students integrate decision-making methods, simulation tools, and stakeholder negotiation within a realistic professional context. Educational outcomes include strengthened systems thinking, collaboration, project leadership, and the ability to justify design choices in social and environmental terms. Institutional benefits include deeper industry engagement and improved graduate readiness. The paper argues that capstones should be judged by the quality of learning processes as much as by final designs, and that this model produces reflective engineers equipped for interdisciplinary practice and responsible innovation.

Keywords: Process Design, Education, Energy Efficiency, Heat Exchanger Network, Optimization

INTRODUCTION

As stated by Lewin et al [1]: “The capstone design course in chemical process design, together with its associated design project is the acid test to gauge the mastery of trainee chemical engineers. There is no other component in the chemical engineering curriculum that performs this function to this degree.” The design project lies at the heart of chemical engineering education because it invites students to think creatively, critically, and boldly beyond the limits of what they have learned in individual courses [2, 3, 4]. It challenges them to work as genuine engineering teams, where ideas must be debated, refined, and transformed into practical solutions. By bringing together topics that are usually taught separately, the project reveals how interconnected the discipline really is and helps students see chemical engineering as a living, integrated profession rather than a collection of isolated subjects [5]. Through this experience, students gain the confidence and imagination needed to shape future technologies. They learn to confront real and emerging challenges: working with alternative raw materials and renewable energy sources, designing processes that can adapt to changing markets, and ensuring that innovation is guided by strong commitments to ethics, health and safety, and sustainability. The design project therefore becomes more than an academic exercise—it is a first step into professional life, where technical knowledge meets responsibility, creativity meets constraint, and engineering meets the needs of

society [6].

This paper presents a flagship, scaffolded final-year design experience that transforms the capstone process design project from a single deliverable into a holistic, industry-engaged, research-informed, and student-centered incubator of professional competence. The project is intentionally ambitious: it integrates multidisciplinary teams, real-world stakeholders, ethical and sustainability considerations, and measurable learning outcomes to prepare graduates who can design, implement, and communicate complex engineered systems with societal impact. The design balances creative freedom with robust assessment and mentorship so every student leaves not only with a detailed design project but with the practical judgment and adaptive skills expected of 21st-century engineers.

PEDAGOGICAL PRINCIPLES

At its core, the project is built on four pillars: authenticity, integration, impact, and reflection. Authenticity requires problems drawn from real industrial, environmental, or public-sector partners and from open research questions; integration demands teams that learn from professional practice; impact-prioritises user-centered, sustainable solutions with clear pathways to deployment; and reflection embeds continual metacognitive practice, so students articulate decisions and learn from failure. These pillars map directly to graduate attributes and the requirements of the Institution of Chemical Engineers: technical competence, systems thinking, communication, ethics, and lifelong learning.

Pedagogically, the programme uses mentored autonomy. Weekly checkpoints (rapid written updates and short presentations to facilitators) promote accountability and peer learning. Workshops teach professional skills — negotiation, risk management, regulatory pathways, intellectual property basics, and science communication. Assessment mixes artifact quality (technical reports, validated simulations), process evidence (design logs, decision rationales), and communication (technical presentations). Weighting emphasizes process and evidence of learning as much as the final product, rewarding reasoning under uncertainty.

Ethics, sustainability, and safety are integrated rather than appended. Each project includes a basic lifecycle analysis estimating CO_2 emissions, embodied energy, and end-of-life impacts, which are mitigated as far as possible. Safety engineering reviews are required, in the form of a HAZOP and other hazard assessments. Assessment design balances rigor and creativity. Rubrics clarify expectations across technical depth, innovation, integration, and communication. External examiners from industry and academia review a sample of projects to ensure standards align with professional expectations.

PROJECT PERSONEL

The final year design (FYD) project is a task undertaken by each 4th year cohort of students at the Chemical engineering Department at Imperial College, who are divided into twelve design teams of approximately 12 students each. The teams are explicitly balanced in terms of overall grade, male / female ratio and home / overseas fee status via an optimisation algorithm implemented in GAMS. Each team has a faculty facilitator, who is responsible for two teams, each of which is assigned a project brief set by the facilitator. The student teams meet once a week with their facilitator and can in addition set meetings with the six faculty experts in the key areas of process analysis, economics, reactor design, separation systems design, control system design, and safety and environment. They also have access to graduate teaching assistants (GTAs), some of whom provide technical clinics. Others who are returning to Imperial to conduct a PhD and have done the project previously, act as student coaches who can also assist with general timelines as well as providing support.

PROJECT TIMETABLE

The Integrated Design Project follows a structured timeline designed to scaffold student progress while maintaining sustained engagement with industrial stakeholders. Prior to the commencement of the project there is an externally hosted away day involving teambuilding activities, where students are encouraged to reflect upon their individual strengths, weaknesses, and what they could get out of the project. Project briefs are released via Blackboard on at the beginning of week one, followed by an initial meeting with facilitators during the first week. Lectures by each of the specialist experts follow within the first two weeks to align expectations and methods. The opening two weeks focus on problem scoping and negotiations with “investors”, for which the students are expected to have produced a preliminary budget and at which they will also agree pricing for key feedstocks, utilities, and a carbon price. These meetings continue throughout the project. A key early milestone is the submission of the intermediate report at the end of the third week. February includes a scheduled interruption for the Frank Morton sport event, during which no project meetings are allowed to be arranged, permitting students time to attend, and for recovery. For readers unfamiliar with the UK context, the event is named after Frank Morton, a well-known Professor of Chemical Engineering who taught at Birmingham University until 1956 before moving to UMIST (now the University of Manchester). A keen sports enthusiast with close links to both universities, he went on to organise the first of these sports days in 1961. The 2024 event was hosted by the University of

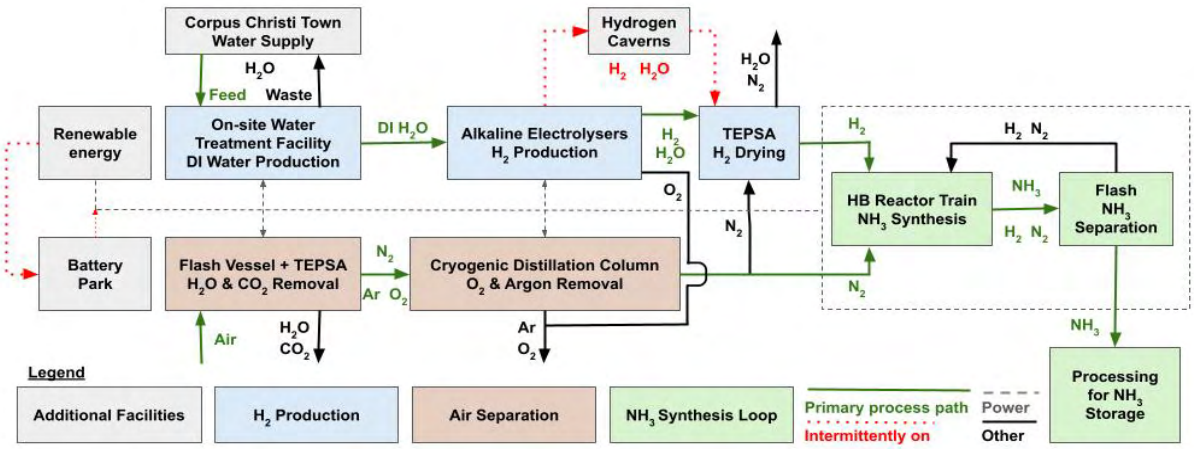

Figure 1: Block Flow Diagram of Group 4's Ammonia Production Process

Swansea. The project culminates in the submission of the final report on around ten weeks in, with final presentations held early in the next week. Informal social activities follow later in March, reinforcing cohort cohesion, while a formal feedback session in May provides reflective closure and supports continuous improvement of the programme.

PROJECT STRUCTURE

Project structure is phased to support increasing complexity while maintaining measurable progress. Table 1 presents the six projects addressed by the Class of 2024.

Table 1: Project Subjects Addressed in 2024.

Assigned Groups	Project
Groups 1 and 2	Sustainable metal organic frameworks (MOF)
Groups 3 and 4	Ammonia from Air and Water
Groups 5 and 6	Fuel production from atmospheric CO_2
Groups 7 and 8	Animal-free microbial dairy cheese production
Groups 9 and 10	VLP-Based Antigen Display for Cancer Vaccine Bioproduction
Groups 11 and 12	Levothyroxine production

Phase 1 (Scoping and Stakeholder Alignment) equips teams to define a clear problem statement, success metrics, and stakeholder map, together with the location of their plant, based on key business metrics (working under uncertainty). Phase 2 (Concept and Feasibility) focuses on concept generation, risk analysis, and a basic feasibility study, developed in ASPEN, but without full implementation at this stage. Phase 3 (Detailed Design and Validation) delivers robust engineering artifacts (simulation, safety and environmental analysis data, control philosophy, etc.). Phase 4 (Deployment Strategy and Handover) culminates in a deployment plan, a sustainability statement, and a handover package suitable for the sponsoring partner. Each phase has significant interactions with both "experts" and "facilitators".

For example, consider the highlights of one of the design projects set in 2024, which required the design of a 2, 000 MTD of Green Ammonia using only air and water as feedstocks and relying on renewable energy, and more specifically, the activities and design decisions of one of the two student teams, Group 4, engaged in that project, per phase:

Phase 1: Scoping and Stakeholder Alignment

In this phase, the teams define a clear problem statement, success metrics, and stakeholder map, together with the location of their plant, based on key business metrics.

The economic basis for the project relies on negotiation of prices of raw materials and energy between the design teams and the staff member responsible for economics mentoring, serving as the "representative" of so-called "Hawkes Industries." The principal methodology utilized in this step is Multi-criteria Decision-making (MCDM), combining Analytical Hierarchy Process (AHP) and Technique for Order of Preference by similarity to Idea Solution (TOPSIS) to select the best alternative for each required decision.

For example, using these techniques, Group 4 chose to plan their process to be constructed in Corpus Christi, Texas, primarily because of the state's bountiful renewable energy resources. Furthermore, based on MCDM, they opted to rely on Alkaline Electrolysis (AEL) for hydrogen production, and cryogenic distillation for nitrogen production, and a heat-integrated, ammonia synthesis

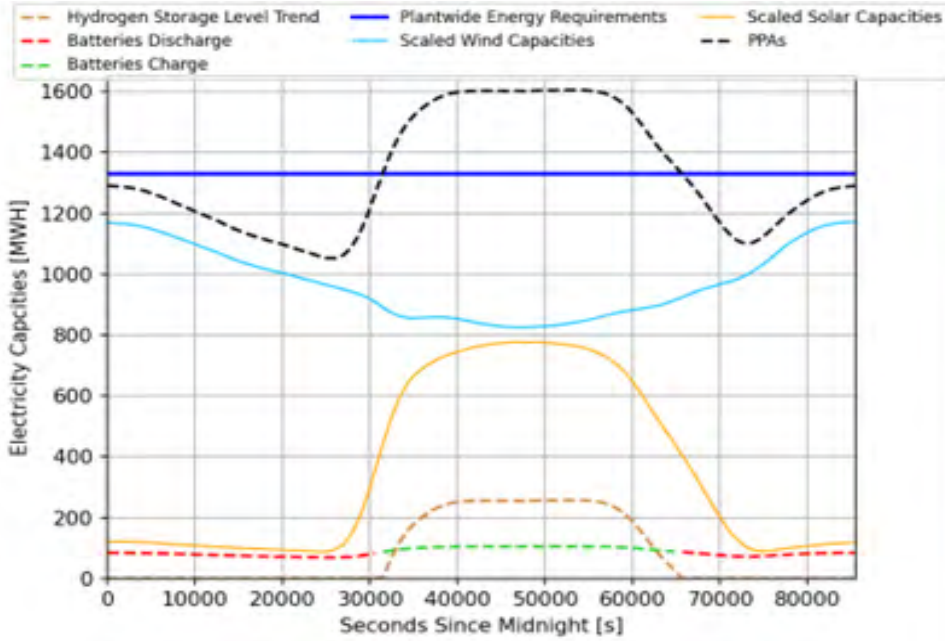

Figure 2: Optimised energy hourly system operation developed by Group 4

loop relying on the Haber-Bosch process, operating at a relatively low-pressure, to be set by economic considerations. It is noted that many of these decisions are often predefined in most design projects as set to students.

Phase 2: Concept and Feasibility Analysis

Focuses on concept generation, risk analysis, and a basic feasibility study, developed in ASPEN, but without full implementation at this stage. By the end of this phase, the student groups should have developed a simplified flowsheet, with many of the operating units being represented by approximate models and have developed a rudimentary economic analysis.

Phase 3: Detailed Design and Validation

By the end of this phase, a complete process simulation will have been completed. This typically involves a detailed ASPEN flowsheet, now combining rigorous unit models, especially in what concerns separation units, and occasionally more approximate specific units, for example the depiction of the AEL system. Typically, the reactor is modelled and optimized in a separate MATLAB model, as was the case for Group 4. The crucially important selection of the reactor pressure was made considering the overall economics of the synthesis loop. Figure 1 shows an overview of the process designed by Group 4, including all its main components, distinguishing between continuous and semi-continuous operation of its components. In this regard, as shown in Figure 2, note Group 4's typical daily scheduling of the various sources of renewable energy and their coordination to ensure the fixed plantwide energy requirements are satisfied by optimizing a balance of power purchase agreements (PPAs), battery energy storage systems, and oversizing the AEL to capture excess power production and mitigate deficits from inherently intermittent green energy systems (wind and solar generation).

Phase 4: Deployment Strategy & Handover

The culmination of the Integrated Design Project is the production of a comprehensive Final Year Design Report, which serves as the principal vehicle for deployment, handover, and professional evaluation. The report is structured around a set of core design domains that collectively demonstrate end-to-end process development: process synthesis; detailed design of key reactor and separation units; plant operation and control; safety, layout, and environmental impact; and economic assessment and business planning. Each domain is presented as a distinct section, developed logically and coherently, with explicit justification of design decisions, assumptions, and the strength of conclusions. Where decisions are informed by computer-based calculations or simulations, relevant results are extracted and presented in clear tabular or graphical form to support transparent engineering judgement.

Although assessed separately, the design sections are tightly integrated through systematic cross-referencing, reflecting the interconnected nature of real chemical plant design. Safety-critical issues arising during start-up or operation are addressed in the Safety section with explicit reference to operating strategies developed in the

Control section, while mechanical or sizing considerations are linked back to the relevant Reactor or Separation Design sections. Scheduling considerations influencing equipment sizing and layout are treated within process synthesis, whereas operational scheduling is addressed within plant-wide control. Each section refers where appropriate to the Process Flow Diagram (PFD), stream tables, and design data sheets contained in the Supporting Documentation, and includes its own references to ensure technical traceability and completeness. The Supporting Documentation forms a structured handover package that enables both rapid appraisal and detailed technical scrutiny by external examiners, industrial stakeholders, and future departmental users. This documentation includes the negotiated design brief, a concise process description, PFDs and stream tables, a plant-wide control survey, detailed process design data sheets for key equipment, site layout drawings, HAZOP records, and a comprehensive economic evaluation including all contracts with "investors". Together with appended meeting minutes, these materials ensure that the report provides a coherent narrative of the design, sufficient depth for expert assessment of individual subsystems, and a durable record of the project suitable for professional and institutional use. The final stage of the project is the production of the report. Group 4's final report was one of the best submitted in 2024 and was selected as Imperial College's entry to the IChemE's annual MacNab-Lacey Prize for Student Design Projects in 2025 and was awarded the prize.

OUTCOMES

Expected outcomes are multi-layered. For students: mastery of end-to-end design thinking, sharpened collaboration and project management skills, and a demonstrable, deployable final report. Polling of the Class of 2026 has been conducted to assess the degree to which these outcomes are perceived by the students. For faculty: stronger industry ties, richer research opportunities grounded in practice, and improved graduate employability metrics. For partners: tested solutions, early access to talent, and case studies. For the institution: a flagship capstone that attracts top applicants, strengthens alumni engagement, and elevates Imperial's role in socially relevant engineering.

Implementation risks include unequal team contribution, scope creep (the tendency of some students to make things more complicated than they really need to be). Mitigation strategies include setting clear roles and expectations, discussion throughout all years at Imperial of the importance of both contributing to the team, and of bringing to the attention of the leadership team any lack of performance so that this can be corrected as soon as possible, mandatory scope reviews, and a conflict mediation process run by experienced academics, at a range of levels of escalation. This is a rare problem, but it needs to be handled properly.

Finally, the programme embeds continuous improvement. End-of-year stakeholder debriefs, and a repository of project outcomes create a feedback loop that evolves problem briefs, rubrics, and support structures.

STUDENT PERCEPTIONS

A pre-FYD survey was carried out between 10th November and 12th December 2025, before students began working on the 2026 FYD project, with a second post-FYD survey carried out between 12th and 16th March 2026, just before the termination of the project. Due to space limitations, we summarize the responses of the students to two key questions in the survey Question 1: Self-assessed abilities, and Question 4: Anticipated Challenges. Only five of the 135 students did not consent to having their responses used for research purposes and were not included in this reporting. A CID ("College/Student ID") field was present and populated in both datasets, enabling respondent pairing. Duplicate CIDs (repeat submissions) occurred in both pre-FYD and post-FYD survey responses; duplicates were resolved by retaining the most recent response per CID (of the latest recorded timestamp). After deduplication, there were 127 usable responses from the pre-FYD responses, and 110 in the post-FYD survey. When limited to the responses which returned matched CID in the two surveys, 100-101 paired responses were retained for Question 1, and 98 retained for Question 4.

Question 1: Self-Assessed Abilities (Confidence Ratings)

In Question 1, students rated their confidence in seven skill areas on a Likert scale from 1 (low) to 5 (high):

1. Tackling complex, open-ended engineering problems
2. Collaborating effectively as part of a project team
3. Leading a design discussion
4. Writing a clear and well-structured technical report
5. Integrating multiple factors in solving a design problem
6. Planning and managing project tasks to meet deadlines
7. Awareness of professional and ethical responsibilities

Table 2 summarizes the responses received for Question 1 in the two surveys. The means and standard deviations returned in the two surveys are presented, as

well the Z-test results comparing the responses in the two surveys, as well as 95% confidence limits in the differences in the paired responses. Overall, the class reported fairly high confidence across most areas in both surveys, with average ratings generally above 4.0.

Table 2: Comparing Survey Responses for Question 1

	n	Pre-FYD (μ±σ)	Post-FYD (μ±σ)	Z	95% CL
1	101	3.94±0.93	4.10±0.87	1.26	[0.00, 0.32]
2	101	4.32±0.80	4.38±0.76	0.56	[-0.05, 0.17]
3	100	3.58±0.93	4.09±1.00	3.73	[0.30, 0.72]
4	100	4.03±0.83	4.19±0.84	1.35	[-0.04, 0.36]
5	101	3.93±0.90	4.19±0.77	2.21	[0.06, 0.45]
6	101	4.08±0.87	4.15±0.78	0.60	[-0.04, 0.18]
7	101	4.27±0.96	4.29±1.00	0.14	[-0.13, 0.17]

The highest self-assessed competencies were teamwork (Skill 2) and ethical responsibility (Skill 7), whereas leading design discussions (Skill 3) received the lowest confidence rating on average. Students generally feel most confident in teamwork and ethics. About half the class (62 out of 127 in the pre-FYD survey) gave a rating of 5 (highest confidence) for working effectively in a team, and even more (66 students) did so for being aware of their ethical responsibilities. In contrast, far fewer students (only 19, around 15%) felt fully confident (rating 5) in their ability to lead a design discussion. This indicates that leadership in design discussions is a key area where many students see room for growth, whereas teamwork and ethical awareness are perceived as personal strengths for most. It is clearly noted that only two issues showed a statistically significant improvement in skills resulting from the project experience: most notably, in the perceived ability to lead design discussions (Z = 3.73), and to a lesser degree in the perceived ability to integrate multiple factors when solving design problems (Z = 2.21).

Question 4: Anticipated Challenges (Ranked Competencies)

Students rank-ordered six competency areas by how challenging they expected each to be in the 4th Year Design Project (1 = most challenging, 6 = least challenging). The aggregated results (lower average rank = more challenging) for the pre-FYD survey are summarized below from most to least challenging:

1. **Project Management** – This emerged as one of the top challenges, with an average rank of 3.03. A significant number of students (28) selected project management as their most challenging area, indicating concerns about planning, organizing tasks, and meeting deadlines. This is consistent with the students' perception of their weakest ability, being leadership of team discussions.
2. **Systems Thinking** – Also highly challenging, with an average rank of 3.09. In fact, 32 students (the highest count for any category) ranked systems thinking as the most challenging aspect. This suggests many expect difficulty in considering the broad, interconnected aspects of the design problem (i.e. thinking about the "big picture" and how different components influence each other).
3. **Communication** – Average rank 3.19. About 20 students ranked communication as their top challenge. While not the number one concern for the majority, a fair portion of the class is apprehensive about effectively communicating ideas, whether in writing or orally, during the project.
4. **Problem-Solving** – Average rank 3.21. This is very close to communication in perceived difficulty. 25 students chose problem-solving as the most challenging. It appears that applying engineering science to open-ended problems and finding creative solutions is a significant hurdle for many.
5. **Teamwork** – Average rank 3.44. Fewer students see teamwork as their greatest challenge (15 ranked it #1), and its average difficulty rating is on the higher side (meaning lower concern). This aligns with the confidence findings in Q1 – most students feel quite capable working in teams, so they are less worried about teamwork issues going into the project.
6. **Ethics** – Average rank 5.05, clearly the least challenging area in students' eyes. Only 6 students ranked ethical considerations as the top challenge, whereas an overwhelming majority (81 students) put it last (ranked #6). This indicates that most students do not expect ethical decision-making to pose much difficulty during the project, likely because they already feel comfortable handling ethical responsibilities or see it as less immediately problematic than the technical and teamwork aspects.

The responses for Question 4 received in the post-FYD survey where similar, with statistically insignificant differences compared to the pre-FYD survey

CONCLUSIONS

In conclusion, the fourth-year design project reframes the final-year project as a public, professional activity. It insists that capstones be judged not just by final reports and presentations but by the quality of reasoning, ethical awareness, and readiness for complex professional practice. When implemented with careful resourcing, transparent governance, and a relentless focus on learning, this model produces graduates who are not only technically excellent but also prepared to lead interdisciplinary projects that matter.

An increase in the level of collaboration with industry is the next main aim for the project, with independent industry experts now acting as facilitators within a number of projects.

Mild to medium increases in self-reported competencies in a number of fields were observed across the board. The greatest (and only statistically significant) increases were in the specific areas relating to design, which is somewhat unsurprising.

There were no strong changes in the rankings of the different challenges within the design project pre and post project. This is likely since the students at Imperial conduct many different team projects in their time at the University, and (though more challenging than other projects) they are aware of the likely potential problems which they might encounter

ACKNOWLEDGEMENTS

The authors acknowledge the excellent work carried out on the project in the 2024 academic year by the entire class, and especially the work of Group 4, whose project results are included in this paper. The Group 4 members were: Sinclair Mabon, Alex Glover, Chow Rouxuan, Gabriela Trisno, Rohan Shenoy, Serene Boonasitha, Qing Li Wei, Shihabuddeen Waqar, Amanda Wong, Zhenran Zhang, and Kabishan Sivarasan. It is of note that Group 4's project won the IChemE's annual MacNab-Lacey Prize for Student Design Projects in 2025.

DECLARATION OF AI

A small portion of the manuscript text was generated using ChatGPT and subsequently verified and modified by the authors.

AUTHOR IDENTIFIERS

Author ORCIDs:
Fennell PS: 0000-0002-6001-5285
Hellgardt K: 0000-0002-4630-1015
Lewin DR: 0000-0002-9277-3777

REFERENCES

7. Lewin DR, Adams II TA, Bongartz D, Léonard G, Suffix AB Jr, Mansouri SS, Martins FG, Mujtaba II and Zondervan E. Teaching process design – Quo vadis? *Computer Aided Chemical Engineering,* 53:3505-3510 (2024). https://doi.org/10.1016/B978-0-443-28824-1.50585-8
8. Silverstein DL, Bullard LG, Seider WD, and Vigeant MA. How We Teach: Capstone Design, presented at *120th ASEE Annual Conference & Exposition*, Atlanta, Georgia. (2013). https://peer.asee.org/19689I
9. Kiss AA, Grievink J. Process systems engineering developments in europe from an industrial and academic perspective. Computers & Chemical Engineering 138:106823 (2020). https://doi.org/10.1016/j.compchemeng.2020.106823
10. Torres AI, Grossmann IE. Teaching of process design courses ? the CMU experience, trends and challenges. Systems and Control Transactions 4:2242-2246 (2025). https://doi.org/10.69997/sct.101971
11. Dumont MN, Foy MP, L?onard G. Integrated project in the master of chemical engineering and materials science at the university of li?ge. Systems and Control Transactions 4:2247-2252 (2025). https://doi.org/10.69997/sct.132784
12. Seider WD, Lewin DR, Seader JD, Widagdo S, Gani R, and Ng KM. *Product and Process Design Principles: Synthesis, Analysis and Evaluation*, 4th Ed. John Wiley and Sons, New York, NY (2017).

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Generative AI in Process Design Instruction: A Survey of Students and Faculty

Daniel R. Lewin[a]*, Thomas A. Adams II[b], Dominik Bongartz[c], Seyed Soheil Mansouri[d] and Edwin Zondervan[e]

[a] Chemical Engineering, Technion, Haifa, Israel
[b] Energy and Process Engineering, Norwegian University of Science and Technology, Trondheim, Norway
[c] Chemical Engineering, KU Leuven, Belgium
[d] Chemical and Biochemical Engineering, Technical University of Denmark, Lyngby, Denmark
[e] Chemical Engineering, University of Twente, Enschede, Netherlands
* Corresponding Author: dlewin@technion.ac.il.

ABSTRACT

A survey was conducted of 103 students and lecturers who had recently participated in chemical engineering design courses concerning their opinions on the use of Generative Artificial Intelligence (Gen-AI) in their capstone design education. Participants were at universities in Europe, the Middle East, North America, and South America, from at least eight different language groups. The survey found little difference in responses between students and lecturers, except for uptake, in which students reported higher rates of familiarity and adoption of Gen-AI tools than instructors. Both groups were net-positive generally on the use of Gen-AI in the classroom, reporting relatively high confidence in the ability to assess results, the general positive benefits of using Gen-AI in their chemical process design education, and the likelihood of using them in the future. However, participants reported that their trust in the results of Gen-AI tools was relatively low.

Keywords: Education, Artificial Intelligence, Process Design

INTRODUCTION: A DISRUPTIVE TECHNOLOGY IN EDUCATION

The rapid emergence of Large Language Models (LLMs) and generative artificial intelligence (Gen-AI) tools such as ChatGPT has sparked widespread debate across higher education, particularly in fields traditionally rooted in quantitative rigor and design creativity such as chemical process engineering [1]. Tao et al. [2] provide a framework enabling the efficient application of LLMs in PSE involving three strategies: standardized prompt structure, task decomposition, and iterative generation. They provide several example applications of these strategies, among them process flowsheet visualization, process simulation and HEN synthesis. As a typical example of the capabilities of LLMs in process design, consider the contribution of Vogel et al. [3], who developed an LLM-based methodology for autocompletion of chemical flowsheets.

As an illustrative example, we used ChatGPT to synthesize an ammonia production process (Haber-Bosch type), stipulating air and water as the feedstocks. The chat record is available in the supplementary material for the curious reader. With just a few simple prompts, it synthesized the process with a general block flow diagram, provided some numerical calculations about approximate material requirements based on stoichiometry, provided approximate energy requirements, opportunities for heat exchange (and offered to run a pinch analysis to synthesize the heat exchanger network), safety notes, control system notes, and produced a cost analysis, with plots, including the levelized cost of ammonia as a function of various market and technology assumptions. After each prompt, it provided a list of steps to be conducted next and offered to do many more tasks. It also told us “You’re doing great” after the first prompt. The total computational workload was less than 2 minutes, encompassing a process that would normally take months for a student in a process design course toward the end of many years of engineering study. The potential value of the tool is profound.

Trust in the correctness of Gen-AI results is also a serious concern. The ChatGPT ammonia synthesis example produced 15 pages of results with bold confidence, but with opaque explanations. It could require weeks or even months for an expert to check each and every step of ChatGPT's results and its thought process. The ability to even check or certify the results at all assumes that ChatGPT would be able to sufficiently provide explanations throughout its process and do so in a deterministic way. The probability of some inaccuracy seems high, not just in mistakes in logic or technical knowledge but even basic math and data. For example, in the preparation of this work, the authors found that leading Gen-AI tools generated hallucinated references that did not exist, reported that the molecular weight of air was 18.02 g/mol (which is for water), and said there were two "r's" in the word "strawberry" (in the 5th and 9th positions), all while projecting confidence. This problem is not just due to potential inaccuracies of the LLM, but also due to the lack of or inaccuracies in open data in process systems engineering on which the models are trained [4]. Note, the LLM is purely trained on language-based data without any physics-informed decision layer. That is, the LLM parameters are obtained over a large dataset not only for process engineering but many unrelated areas, even Ozzy Osbourne's hard rock metal lyrics. Contrary to popular belief, ChatGPT is not simply just relying on available online information which it harvests by live lookup but has the capability of at least performing preliminary computations for itself. For the example interaction, ChatGPT explains that it arrives at the numerical values it reports on response to the user prompts by performing stoichiometric material balances and relies on cost exponents for CAPEX estimation, among other simple calculations.

This example presents an existential question for chemical engineering design education. From a purely pedagogical perspective, the advantages to the students could include having a tool that guides them through the process, offers ideas and starting points, provides always-available technical support and personalized tutoring, and of course an incredible speed-up in results generation. The sheer rapidity of results could permit students to solve much bigger and deeper problems that are simply impossible to do in in a design course with pre-ChatGPT technology.

The disadvantages and concerns for pedagogy are also myriad. These include the erosion of deep understanding, the outsourcing of critical thinking, and an increased temptation toward plagiarism. The temptation to use them as a replacement for the hard work of personal learning is evident. It is reasonable to surmise that these would lead to very poor educational outcomes and produce insufficiently competent engineers, creating a downward spiral of skill.

Suppose that even if the personal development concerns of the student are mitigated. Since results can be generated so quickly but trust in the results is a major concern, does this mean our chemical engineering design courses should primarily focus on training students to check and certify Gen-AI results rather than the art and science of designing it themselves? Will this be sufficiently motivating to encourage students to enter the process design profession, or even to develop their own skills generally?

To get a better handle on these issues, we conducted a survey of students and lecturers who took or taught capstone chemical process design recently about their opinions and experiences of Gen-AI. The goal is to get the perspective of students and instructors as to where they think the technology is going and how they are using it now. There are no such recent surveys on chemical engineering capstone-level design courses to the author's knowledge.

RECENT RELATED SURVEYS

In a study focusing on Gen-AI usage in a second-year chemical engineering course, Verret [5] polled students on their perceptions on the degree of assistance they received from Gen-AI on assignments related to design, finding that about half of the students polled made no use of AI, and that those that did mostly used it to edit documents, for problem solving, and to critically evaluate information. As for their concerns in reliance of Gen-AI, of those students that used it, the issues that most concerned them centered around the low quality of information obtained, the academic integrity of the information obtained, and the possibility of fabricated, incorrect, or missing citations.

Chans et al. [6] used pre-tests and post-tests that book-ended a four-hour beverage design problem exercise in a chemical engineering product design course. The exercise was primarily a brainstorming session, guided by the lecturer who provided Gen-AI tools and prompts for them to use. They reported a general increase in the students' perceived usefulness of Gen AI-tools but did not find a significant increase in other metrics. The sample size was small (9 students) and there was no control group, so general conclusions cannot be made about the ability of Gen-AI to increase design thinking, creativity, and innovation.

Huang [7] surveyed students in a course in chemical process control for three modules related to coding, Laplace transforms, and PID controller tuning. Students were instructed to engage with ChatGPT during the workshops. Students reported slightly net-positive outcomes in their perceived understanding of the theoretical topics at hand, but strong net-positive responses in understanding how AI can be applied to chemical engineering problems. Students also had negative responses

concerning confidence in the AI-generated results and noted they needed more guidance in how to check. Some of these results will be echoed in the results of this survey on process design.

Although not about education specifically, 86 engineers practicing in the UK were asked if they thought ChatGPT should be taught in engineering higher education [8]. Although 82.5% agreed that engineers generally needed training in ChatGPT, only 55.7% felt that it should be explicitly "embedded" in university courses. Although still net-positive, it does not signal a huge push for explicit educational integration. Similar attitudes were found among our student and lecturer survey participants, as explained in the following sections.

CHEMICAL PROCESS DESIGN SURVEY

Given the recent emergence and rapid spread of Generative AI tools, their role in process design curriculum is still developing. To gain a first overview on current practices as well as a comparison between the associated perceptions of lecturers and students, we conducted a survey of students and lecturers in chemical engineering design programs between December 4, 2025 and January 25, 2026. It consists of questionnaires distributed to chemical engineering students as well as their instructors, asking about familiarities with Gen-AI tools, current use, view on opportunities and challenges in their application, as well as opinions on recommended future development. The lecturers and students received parallel questionnaires, covering the same topics, but with different wordings appropriate to the role of the survey participant. The survey was open to the public and taken via an internet website. It was announced through direct email of colleagues, posts on relevant professional network websites and listservs, and in multiple languages. Responses were received from participants in Catalan-, Danish-, Dutch-, English-, French-, German-, Hebrew-, Portuguese-, and Spanish-speaking institutions or countries, whose chemical engineering design courses are either conducted in the relevant national/regional language or in English. The survey itself was conducted in English. The question sets and the associated raw data are posted in the supplementary material.

Demographics of Survey Participants

Participants were asked for their current university and status as a lecturer or student. Participants who selected "Student currently taking a course in chemical engineering design, or have done so, or expecting to take it over the next year" are classified as a Student (S), with 81 responses. Participants who selected "Lecturer currently teaching chemical engineering design, has taught it recently, or is expecting to teach it in the next year" are classified as a Lecturer (L), with 22 responses. 1 participant selected the "Neither of the Above" option and was discarded from the analysis.

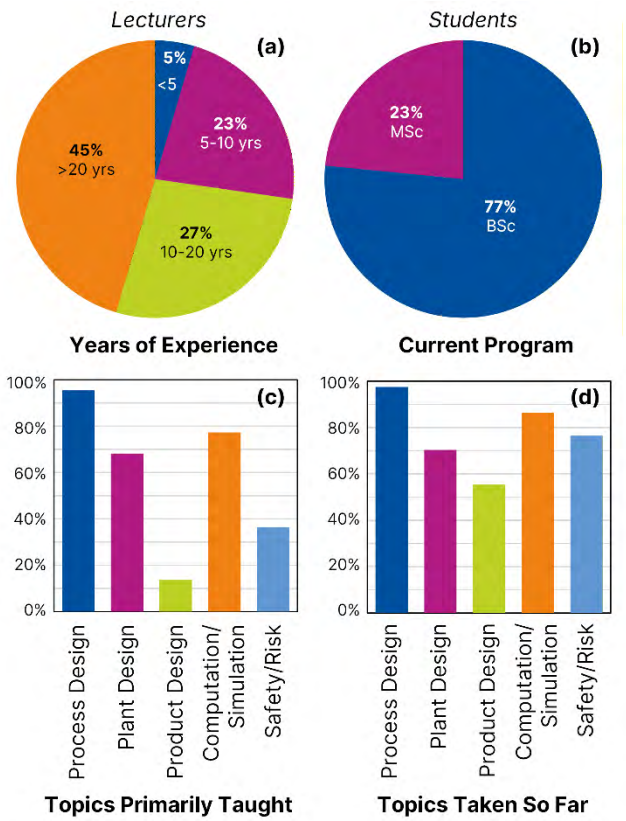

Figure 1. Participant demographics. (a) Years of experience in chemical process design instruction of lecturers; (b) The current level of the degree program of student participants; (c) The topics primarily taught by lecturers; (d) The topics studied by students. (N=103)

Questions 1 and 2 asked participants about the relevant areas of chemical engineering design that they either teach (L) or have taken so far (S). Students were asked about their current degree program, and lecturers were asked about the numbers of years of experience they have teaching process design. Participants were presented with multiple-choice lists. A summary of the results is shown in Figure 1. The vast majority of the lecturer participants had at least 5 years of experience, with almost half having at least 20. With regard to the topic involvement, few lecturers selected product design, but the question specifically asked what instructors *primarily* teach. Conversely, most student respondents indicated they had studied product design previously. This disparity in responses could simply mean that instructors do not consider product design to be a primary aspect of their instruction, but students are exposed to it nevertheless.

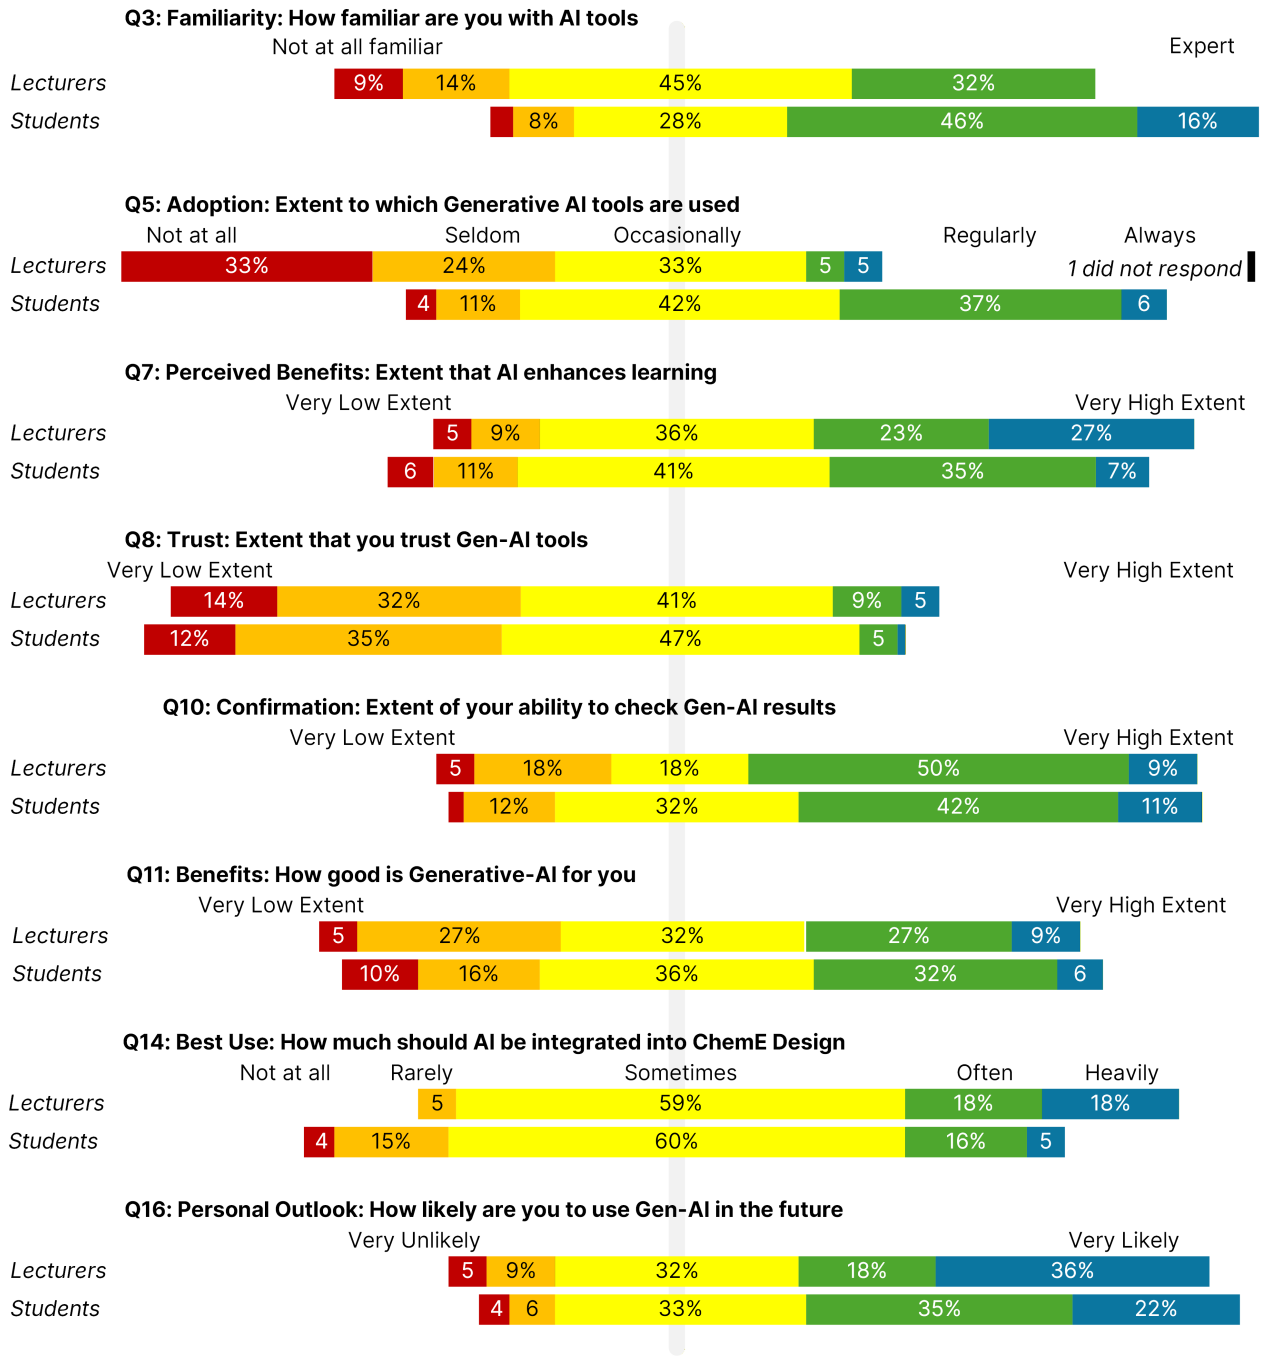

Figure 2. Resposes to Likert scale questions (Q3, Q5, Q7, Q8, Q10, Q11, Q14, Q16). *N=103*

Because the question asking participants to state their university had a free text field for an answer, not all responses could be accurately understood (because of acronyms and initials), so these are not included in the formal analysis.

Responses to AI-Related Questions

The following questions ask participants about their opinions concerning Gen-AI in chemical process design. Recipients were asked to respond using a 1 to 5 Likert scale, where 1 indicates "to a very low extent/I never use, " to 5 indicating "to a very high extent/I always use." Responses are summarized in Figures 2, 3, and 4. Comments on individual questions follow.

- ***Q3: The degree to which S/L are familiar with Gen-AI.*** L averaged 3.00, while S averaged 3.66. Z-test value is 2.97, indicating S feel significantly more familiar with Gen-AI tools than L.

- ***Q4: Which specific Gen-AI tools are used?*** The most commonly used tools are ChatGPT (91% of L, 95% of S) and Copilot (L: 50%, S: 43%). Responses are summarized in Figure 3.

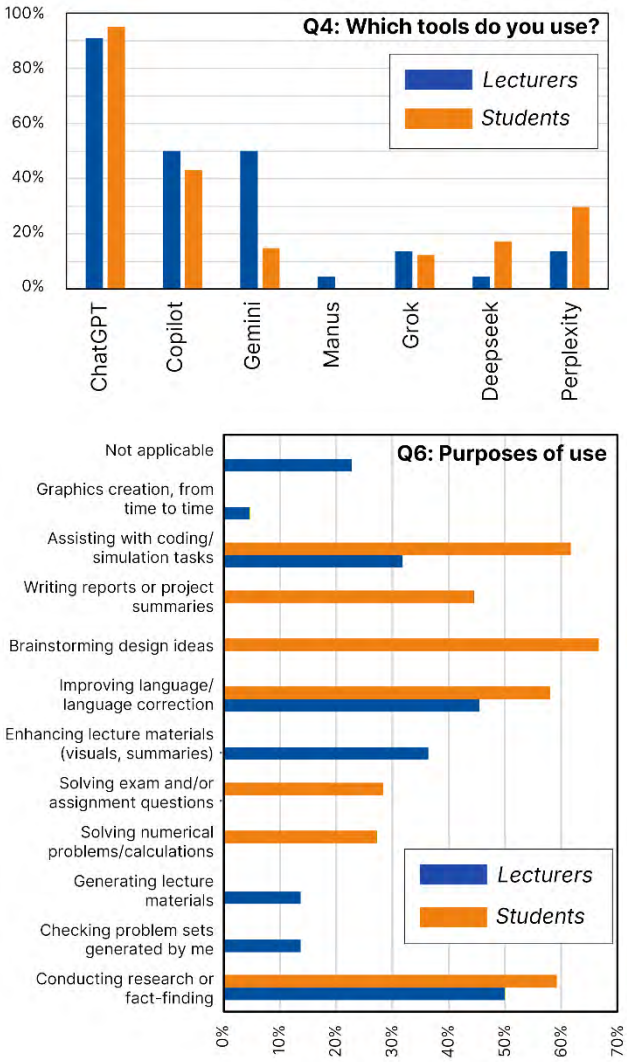

Figure 3. Responses to Q4 and Q6. *N*=103

- ***Q5: The extent to which Gen-AI tools are used in the design courses.*** The average L response was 2.14, close to "I seldom use, " while the average S response was 3.31, equivalent to more than "I use them occasionally." Note that one third of L do not use Gen-AI in their design courses at all. Statistically the two distributions score a Z-test value of 4.47, meaning that the students use Gen-AI tools quite significantly more than the lecturers do. These responses are consistent with those received for Q3 given that both questions indicate higher uptake in the S class.

- ***Q6: For which purposes are Gen-AI tools used***? Most popular use for S is "brainstorming" (67%), while this usage scores 0% for L. Perhaps the lecturers do not consider brainstorming with AI to be brainstorming, and students include AI into teamwork and they of course do. Next most popular for S is "Assisting with coding/simulation tasks" (62%), while this scores only 32% for L. Of the rest, both use if for "Conducting research or fact-finding" (L:50%, S: 59%), and for "Improving language/language correction" (L: 45%, S: 58%).

- ***Q7: The extent of perceived benefit of Gen-AI in the process design courses***. Both L and S average about the same (L: 3.59, S: 3.26), and perceptions are statistically indistinguishable. The large majority of participants indicate that there are at least some benefits.

- ***Q8: The degree to which S/L trust Generative-AI in their process design courses.*** Both L and S average about the same (L: 2.59, S: 2.48), and perceptions are statistically indistinguishable. This finding can be surprising, but this same result was found by Đerić et al. [9] whose recent study designed to address this specific point also found no statistical difference between students and lecturers in their level of trust in Gen-AI. Trust skews heavily negative for both L and S.

- ***Q9: The degree to which lecturers favour results given by Generative-AI over other sources***. The question was removed from the analysis due to a technical issue with the survey tool.

- ***Q10: The extent of the perceived ability of S/L to check the results of Gen-AI.*** Both L and S average about the same (L: 3.41, S: 3.47), and perceptions are statistically indistinguishable. Both groups self-reported a relatively high confidence in their ability to determine the correctness of the output given by Gen-AI. This could partially explain why students have a relatively high adoption rate of Gen-AI tools, despite a relatively low level of trust in its output. This however would not explain why lecturers have a relatively low adoption rate of Gen-AI in their courses, since they too have high a self-perception of their own ability to determine correctness.

- ***Q11. The degree to which S/L believe Gen-AI is a "net good" for their education***. Both L and S average the same (L:3.09, S: 3.09), and perceptions are statistically indistinguishable.

- ***Q12: Which are the benefits of using Gen-AI tools?*** Highest scores were for "Getting explanations in simpler terms" (L:59%, S: 54%). Other but less mentioned benefits included "Saving time of assignments" (L:41%, S: 44%), "Bridging gaps in prior knowledge" (L:36%, S: 28%), and "Improving creativity" (L:36%, S: 25%).

- ***Q13: Which are the greatest risk and challenges of using Gen-AI tools?*** Similar perception of L and S with the highest scores being for "Over-reliance on AI /Becoming too dependent on AI" (L:91%, S: 80%), "Inaccurate or misleading outputs" (L:73%, S: 73%), and "Students bypassing fundamental steps" (L:64%, S: 59%).

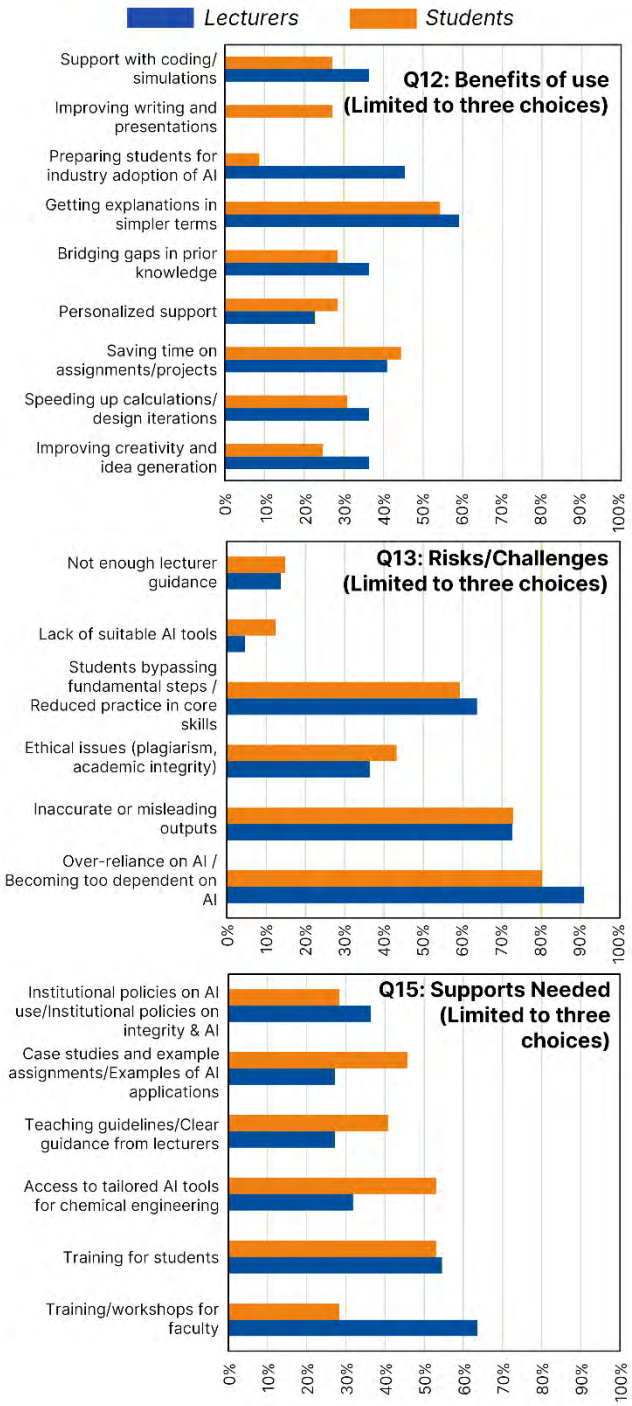

Figure 4. Responses to Q12, Q13, and Q15. *N*=103

- ***Q14. How much should AI be integrated into process design?*** L average response is 3.50, while the average response for students is lower at 3.04. Note that L are more enthusiastic about the introduction of Gen-AI in design education than students, and the difference in perception is statistically significant (Z-test = 2.31). This could in part be due to self-selection bias in those who chose to answer the survey. These results mirror the survey results of engineering employers [10] who also had a middle-of-the-road opinion that ChatGPT (specifically) should be integrated into engineering higher education.

- ***Q15: What kind of support is needed for Gen-AI usage?*** L opinion that support is needed most for training for faculty (64%) and students (55%). S think that support is needed most for "Training for students" and "Access to tailored AI tools for chemical engineering" (both 53%) and for Case studies and examples (46%). It would appear that students have put more thought into integrating Gen-AI than have the lecturers.

- ***Q16. How likely are you to use generative AI in future?*** Both L and S average about the same (L:3.73, S: 3.65) that were the highest scores in the survey, with perceptions being statistically indistinguishable.

Open-ended Responses

One question allowed responders to make closing remarks. Here are a few of the more noteworthy comments from the lecturers:

- "The whole purpose of education is to get students to think for themselves, and to understand the principles of engineering science well enough that they can either do their own calculations (using generic tools such as Excel) or check calculations made using specialist computer software. They need a good level of fundamental understanding so they can look at a design and answer the questions "does this make sense?" and "is this a feasible and safe design?" This requires development of an internal structure of knowledge, and my experience of AI is that it is devoid of any understanding of the structure of knowledge. Another major issue with AI is hallucinating references. **If you can't tell where the underlying information is coming from, how can you trust the output?** Every month or two, I see stories in the press or elsewhere about lawyers getting fined because their pleadings (generated by AI) contain either references to completely fabricated cases, or the real cases quoted don't support the arguments made in the pleading. Because lives are at stake in engineering, similar issues are much more serious in our field. We

wouldn't just lose money…"

- "While **AI is a great help in numerous aspects**, ranging from the simple improvement of text quality to providing significant assistance in solving complex problems, it is also true that, currently, it frequently introduces changes that modify the meaning of texts, makes incorrect assumptions, and tends to attempt to fulfil the request, sometimes even at the expense of 'falsifying' some results. In many instances, **locating these errors is not trivial** (even for experts), **which poses a significant risk**, particularly for students. In my prompts, I often add phrases such as 'do not invent anything, ' 'justify all your steps, ' 'always include the source from which you obtained the data, ' and even then, I do not always succeed. Generally, it is a tremendous tool that must be leveraged at all levels (i.e., research and teaching), but as in all cases, the responses must be analyzed with a critical spirit, and this applies to both experts and novices."
- "The main point to be explored**: the different way students and instructors use Generative AI and ask it questions** (I'm reluctant to use the term "prompt engineering" to refer to understanding the problem and asking the right questions in the precise way by understanding the rules of grammar and logics, which is the real issue)."
- "There are clear differences between architecture/programming AI and Instantiation on AI application. Those differences are not clear for the university staff, and consequently students don't learn about it. Another question is the theoretical overlooking about ethical aspect without ethical day-by-day practices by the staff. A third urgent question is how the students will learn and practice the **Law of Unintended Consequences** on pervasive AI."

The survey has highlighted what appears to be a student body that at least perceives itself to be significantly more familiar with Gen-AI tools than their teachers, with about a third of the professors not using Gen-AI in process design education at all (Q3 and Q5). As both students and lecturers agree that Gen-AI will be utilized at an increasing extent in process design education (Q16), it would appear to be incumbent on lecturers to increasing integrate these tools into their teaching. Perhaps a good summary would be that offered by of one of the lecturers, who commented that:

- "We need to get on with it, or we will be left behind. But, one year before official retirement, I am not the professor to do this anymore ;-)"

Limitations to the Methodology

The limitations of the survey should be considered in its interpretation. The number of lecturers who participated was not as high as would be preferred, and subject to self-selection bias since those who both hear about and choose to participate may be skewed toward a class of people who are connected to the outreach networks used in participant recruiting. The number of student responses was sufficiently high but varied widely by institutions. Although many institutions were represented by at least one student participant, a few universities had large response rates, likely due to direct encouragement by the class instructor. Thus, responses of the S class as a whole will be skewed toward these universities. Much larger sample sets and participation rates would allow greater confidence in the S class and would allow us to create profiles representative of the institutional level.

CONCLUSIONS

The emergence of generative AI represents a fundamental shift in the environment in which process design is taught and practiced. Tools capable of producing flowsheets, calculations, code, and technical text within seconds are no longer peripheral aids but active participants in the design process. The results presented in this paper confirm that their use is already widespread among students and that adoption is likely to increase further in the near future.

The survey reveals a clear asymmetry between students and lecturers. Students report significantly greater familiarity with generative AI and make more frequent use of it in design courses, particularly for brainstorming and coding support. Lecturers, in contrast, remain more cautious and in many cases have not yet integrated these tools into their teaching. Despite this difference in usage, both groups express remarkably similar views regarding the perceived benefits, the limited level of trust in AI-generated results, and the risks of over-reliance and inaccuracy. This shared ambivalence suggests that the central challenge is not whether AI will be used, but how it can be used responsibly.

The pedagogical opportunities are substantial. Generative AI can provide cognitive scaffolding, personalized explanations, and exposure to alternative design options that may enrich creativity. At the same time, the risks identified—erosion of core engineering skills, propagation of errors, and threats to academic integrity—cannot be ignored. Process design education has always aimed to develop judgment, accountability, and the ability to verify results; these goals become even more critical in an AI-assisted context.

The findings point to several priorities. First, educators need support and training to engage confidently with these technologies. For example, both students and

lecturers may need detailed training in effective prompt engineering. This approach is currently in development in other areas of chemical engineering [10]. Second, assessment methods must evolve to emphasize critical evaluation and transparent workflows rather than final answers alone. Third, students must be explicitly taught how to question and validate AI output as part of professional responsibility.

Generative AI should not replace the need for sound engineering reasoning, but it will reshape how that reasoning is exercised. The task ahead is to integrate these tools in ways that strengthen, rather than weaken, the intellectual foundations of process design education.

SUPPLEMENTARY MATERIAL

The supplementary material can be downloaded at https://PSEcommunity.org/LAPSE:2026.0006

DECLARATION OF AI

ChatGPT and Co-Pilot were used as described in the work. A small portion of the manuscript text was generated by ChatGPT, verified, and modified by the authors.

AUTHOR IDENTIFIERS

Author ORCIDs:
Lewin DR: 0000-0002-9277-3777
Adams II TA: 0000-0002-9871-9851
Bongartz D: 0000-0003-1790-0235
Mansouri SS: 0000-0002-0924-4947
Zondervan E: 0000-0003-2659-1622

REFERENCES

1. Caccavale F, Gargalo CL, Gernaey KV, Krühne U. Towards education 4.0: the role of large language models as virtual tutors in chemical engineering. Education for Chemical Engineers 49:1-11 (2024). https://doi.org/10.1016/j.ece.2024.07.002
2. Tao X, Tula A, Chen X. From prompt design to iterative generation: leveraging llms in PSE applications. Computers & Chemical Engineering 202:109282 (2025). https://doi.org/10.1016/j.compchemeng.2025.109282
3. Vogel G, Schulze Balhorn L, Schweidtmann AM. Learning from flowsheets: a generative transformer model for autocompletion of flowsheets. Computers & Chemical Engineering 171:108162 (2023). https://doi.org/10.1016/j.compchemeng.2023.108162
4. Woo T, Kim S, Tariq S, Heo S, Yoo C. Leveraging generative AI and large language model for process systems engineering: a state-of-the-art review. Korean J. Chem. Eng. 42:2787-2808 (2025). https://doi.org/10.1007/s11814-025-00524-y
5. Verrett J. Work-in-progress: student perceptions and usage of generative AI in second-year chemical engineering design exercises. 2025 ASEE Annual Conference & Exposition Proceedings : (None). https://doi.org/10.18260/1-2--57548
6. Chans GM, Merino-Soto C, Chávez SS, García Castro JA, Zavala G, Rodriguez ES. Integrating generative AI into design thinking: assessing impact on creativity and innovation in STEM education. 2025 IEEE Global Engineering Education Conference (EDUCON) :1-7 (2025). https://doi.org/10.1109/educon62633.2025.11016312
7. Huang Z. Integrating artificial?intelligence assistance into chemical process control education. Education for Chemical Engineers 54:100492 (2026). https://doi.org/10.1016/j.ece.2025.10.002
8. Murray M, Maclachlan R, Flockhart GMH, Adams R, Magueijo V, Goodfellow M, Liaskos K, Hasty W, Lauro V. A 'snapshot' of engineering practitioners views of chatgpt-informing pedagogy in higher education. European Journal of Engineering Education 51:104-129 (2025). https://doi.org/10.1080/03043797.2025.2492736
9. ?eri? E, Frank D, Milkovi? M. Trust in generative AI tools: a comparative study of higher education students, teachers, and researchers. Information 16:622 (2025). https://doi.org/10.3390/info16070622
10. Galatro D, Chakraborty S. Strategies to map education 5.0 and industry 5.0 in the context of a modernized undergraduate program in chemical engineering. 2025 IEEE Global Engineering Education Conference (EDUCON) :1-9 (2025). https://doi.org/10.1109/educon62633.2025.11016293

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Understanding Student's Preferences for Computational Tools in Chemical Engineering Assessment

Sakiru Badmos

Department of Chemical Engineering, University College London, Torrington Place, London WC1E 7JE, United Kingdom
Corresponding Author: s.badmos@ucl.ac.uk

ABSTRACT

Computational tools are widely used in solving engineering problems and are now embedded within chemical engineering education. At the UCL Department of Chemical Engineering, students are taught gPROMS ModelBuilder in modules requiring coding; however, many choose alternative tools such as MATLAB, Python, or Polymath for coursework and capstone design project reactor design. This study investigates the reasons behind these preferences using a survey of fourth-year students who had completed their third-year design project. The results show that perceived ease of use, availability of external resources, and ease of debugging could strongly influence tool selection. The findings highlight the importance of accessibility, community support, and perceived relevance in shaping sustained student engagement with computational tools.

Keywords: Technology Adoption, Computational tools, Matlab, Python, Polymaths, Engineering Education, gProms

INTRODUCTION

Computational tools have been part of chemical engineering education for several decades, supporting the solution of algebraic equations, differential equations, and optimisation problems in core areas such as reaction engineering, separation processes, and process control [1-2]. What has changed in recent years is not their presence in the curriculum, but the *type of tools used, their accessibility, and how students engage with them*. The emergence of open-source programming environments, large online user communities, and data-driven methods has expanded the range of tools available to students and influenced how they approach computational problem-solving [3].

A range of platforms is currently used in undergraduate chemical engineering programmes. Tools such as POLYMATH remain effective for solving numerical problems, particularly in reaction engineering, due to its simplicity, presenting minimal challenges when picked up by students without prior knowledge of programming [4]. MATLAB is widely used in academia and industry for numerical methods, process modelling, and control because of its powerful numerical solvers, extensive libraries and visualisation capabilities [5-7]. More recently, Python has gained significant traction as an open-source alternative, supported by extensive scientific libraries and a large community of users, making it attractive for both teaching and independent learning [3]. The integration of Jupyter notebooks has made it attractive for teaching in undergraduate engineering programs [8-12].

Alongside these tools, commercial process modelling environments such as gPROMS ModelBuilder are increasingly used to represent industrial workflows, particularly for dynamic modelling, optimisation, and parameter estimation [13].

At University College London (UCL), gPROMS is taught within the undergraduate programme to expose students to equation-oriented modelling approaches used in industry. However, unlike widely used scripting tools, gPROMS requires a licence and has a relatively smaller user community and fewer openly available learning resources. These differences may influence how students engage with the tool, particularly when they have a choice in assessment. Despite being taught gPROMS, many students choose alternative tools such as MATLAB, Python, or POLYMATH for their coursework and design project reactor designs. This raises an important question: *why do students move away from tools that are formally taught, even when support is provided?* While prior studies have explored the pedagogical value of individual tools [5–12], less attention has been given to students'

tool preferences in assessments, where time constraints, confidence, and access to support play a significant role.

This study investigates the factors influencing students' choices of computational tools in chemical engineering assessment, using gPROMS ModelBuilder as a case study. Drawing on principles from the Technology Acceptance Model [14], this work examines how perceptions of ease of use, usefulness, and access to support shape students' decisions. The study addresses three research questions: (i) what factors influence students' decisions to use alternative tools when gPROMS is taught, (ii) how students perceive the teaching and support provided for gPROMS, and (iii) what changes could improve engagement with computational tools in the curriculum.

METHODOLOGY

A Mentimeter survey was presented to the fourth-year students (n = 40) of the 2025/2026 academic year, who had completed their design projects in the previous year. This cohort was chosen because they had direct experience applying a computational tool in their coursework and the third-year design project. The survey consists of 9 questions of different types including multiple choice and an open-ended question using a five-point Likert scale. Ethical approval was obtained from UCL's Research Ethics Office to process and publish the results of this survey. The survey questions are as follows:

- Q1 – Which primary tool did you use for your Year-3 design project reactor design?
- Q2 – Which primary tool did you use for other coursework e.g CREII?
- Q3 - How confident did you feel using your primary tool for design project and coursework?
- Q4 - Why did you choose that tool?
- Q5 – What is the most important factor to consider when choosing a tool for your work
- Q6 - How adequate was the gPROMS teaching you received before the design project?
- Q7 - How important is tool popularity/community in your decision?
- Q8 - How important is tool popularity/community in your decision?
- Q9 - What would make gPROMS more usable/attractive for students?

Free texts from students' responses to Q9

- T1: 'Better UI'
- T2: 'Make the interface more user friendly and intuitive'
- T3: 'Better interface (why model and process? One would be simpler)'
- T4: 'More tutorials with solution'
- T5: 'I prefer to learn code from online sources like stack overflow examples and github and YouTube videos. gPROMS doesn't have this in the same way as Python and even MATLAB'.
- T6: 'Popularity – it is so unnecessarily hard to find help for this'
- T7: 'Clearer availability of support/more support sessions'

RESULTS AND DISCUSSION

Chemical engineering education has computational tools as integral part of the undergraduate curriculum. These tools are used for solving problems involving numerical methods, ordinary differential equations (ODEs), optimisation, and, in some cases, dynamic simulation. Traditionally, software such as MATLAB, Excel, and Visual Basic has been used for this purpose, with more recent additions including Python, gPROMS, and GAMS. At UCL Chemical Engineering Department, students are taught GAMS and gPROMS within core chemical engineering modules, while MATLAB is introduced through the Mathematical Modelling and Analysis (MMA) module delivered by the faculty as part of the Integrated Engineering Programme (IEP). These tools are embedded in the curriculum to satisfy accreditation requirements set by the Institution of Chemical Engineers (IChemE) [15]. The skills developed using these tools are applied across several modules throughout the undergraduate programme. Further details on how the tools are taught can be found in Kamel et al. [16].

Despite students being introduced to gPROMS from their first year, they appear to show low enthusiasm towards the tool both in teaching and learning and in applying it for their coursework. When given the opportunity to select a preferred software, many students will choose alternative tools, even though gPROMS is taught on the programme. This pattern is observed in both the second-year Chemical Reaction Engineering module (CREII) and the third-year capstone design project, where students are required to implement reactor models computationally. Although gPROMS is supported through teaching sessions, questions and answers (Q&A) and forums on the learning platform (Moodle), students are permitted to use other tools if they feel more comfortable doing so.

From the survey results, as shown in Figure 1a, 50% of the respondents (20 students) designed a reactor as part of their design project, and among these, 90% used MATLAB or Polymath, while only 10% used gPROMS. This occurred despite the availability of dedicated

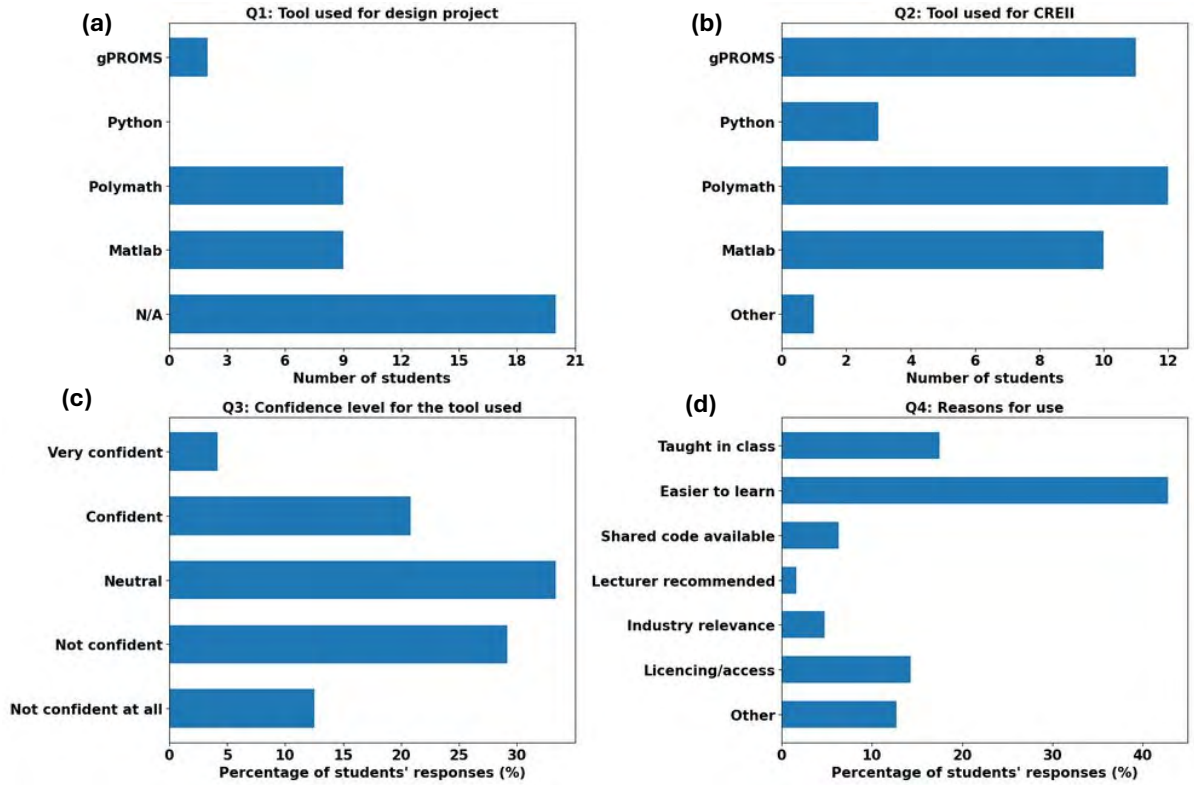

Figure 1. Students' responses to Q1 – Q4. N/A (n=20) in panel 1a indicates respondents who designed a separator during the detailed design of their third-year design project and does not require coding.

computational support for gPROMS on the design project. Similarly, in CREII, where all students must implement a reactor model in a software, approximately 70% selected a tool other than gPROMS (Figure 1b). These results indicate a clear reluctance to use gPROMS and suggest that factors beyond simple availability of support influence students' choices.

Confidence in using computational tools also appears to be an important issue. Figure 1c shows that although students were free to choose their preferred software, around 42% reported that they were not confident in using the selected tool, and a further 33% were uncertain about their confidence level. This suggests that students may be selecting tools based on short-term use rather than genuine competence. It also points to a possible need to review current teaching approaches or to provide additional learning resources to better support students in developing confidence with the tools they are taught.

The survey results indicate that students' tool selection is shaped by a combination of practical and perceived benefits. As shown in Figure 1d, ease of learning was the most influential factor, with approximately 42% of respondents identifying this as the main reason for their choice. About 30% selected their tool because it was taught in class, which likely corresponds to those who used MATLAB or gPROMS, as these are the two scripting tools formally introduced during the programme. Licensing and access were also relevant, with 14% of students citing this as a determining factor. Smaller proportions identified the availability of shared code, industry relevance, lecturer recommendation, and other factors as influencing their decision.

When asked what factor should be prioritised when selecting a tool for assessment, students again emphasised ease of learning and external support. As shown in Figure 2a, 27% selected ease of learning as the most important criterion, while 23% selected the availability of resources and community help. Ease of debugging and industry employability were each selected by 18% of respondents, and supervisor preference by 14%. These responses suggest that students tend to prioritise immediate academic needs over longer-term considerations such as industrial relevance. This may reflect a focus on completing assessments efficiently rather than on developing skills perceived as valuable for future employment. It may also indicate limited awareness of the tools commonly used in professional practice, or a tendency to defer such concerns until later stages of their studies.

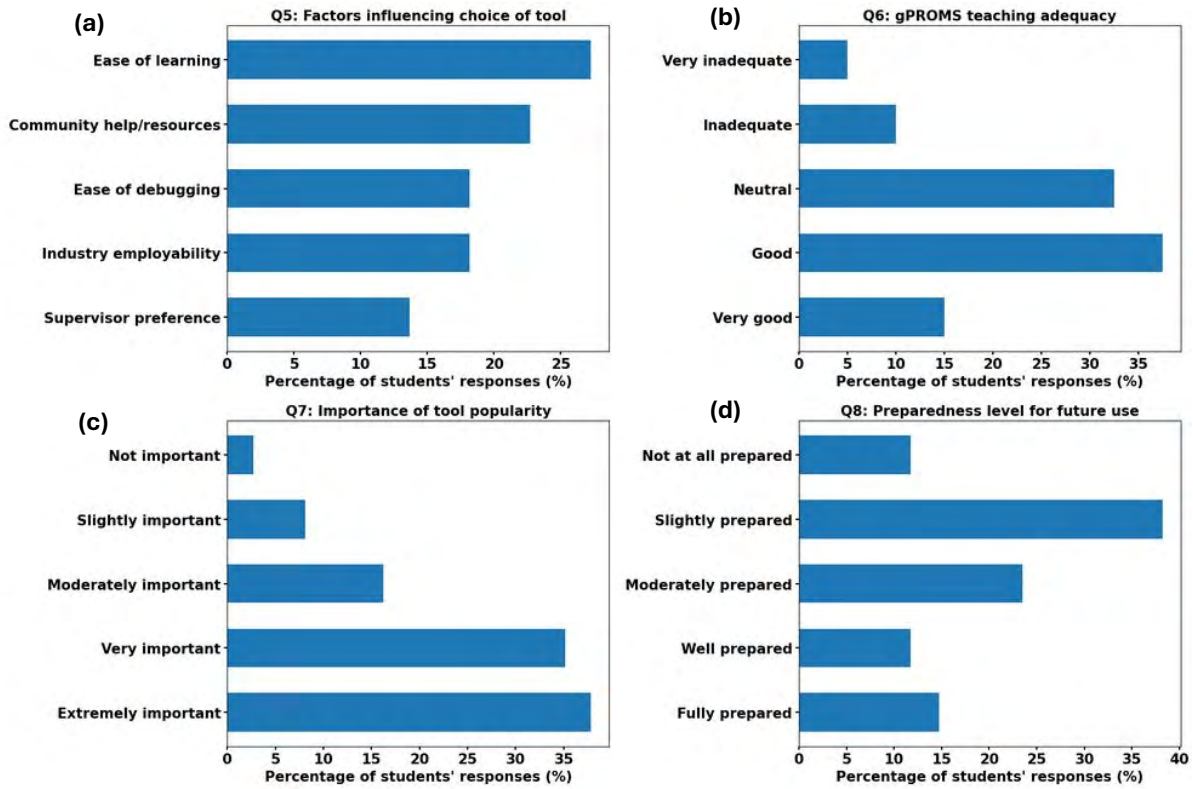

Figure 2. Students' responses to Q5 – Q8 on factors affecting choice of tools (2a), adequacy of gPROMS teaching (2b), importance of tool popularity in their decision making (2c) and their level of preparedness for using the tool in future roles (2d)

The importance placed on external resources and community support highlights the role of informal learning environments in students' decision-making. Large user communities and readily available online examples make tools such as Python and MATLAB particularly attractive, especially when students encounter errors or difficulties. Availability of template codes in Fogler's Reaction engineering textbook makes Polymath a good choice. These features directly affect their ability to complete coursework within limited timeframes and were therefore prioritised. A substantial number of students also valued ease of debugging and the potential for transferability of skills for industrial applications. At the same time, some students preferred to follow lecturer recommendations, possibly viewing these as aligned with assessment expectations.

Figures 1a and 1b show that most respondents used tools other than gPROMS for both the design project and CREII assessment, with only 10% selecting gPROMS for reactor design in the capstone project. To explore whether this preference resulted from inadequate teaching instructions, students were asked in Q6 to evaluate the adequacy of gPROMS teaching and support. As shown in Figure 2b, 52% considered the teaching and learning provision was sufficient to support their use of the software, while 32% were unsure and only 16% felt that support was inadequate. This suggests that dissatisfaction with teaching provision alone does not fully explain the low uptake of gPROMS. Instead, the findings indicate that students' choices are more strongly influenced by perceived ease of use and perceived usefulness, consistent with the Technology Acceptance Model (TAM). Even where formal teaching and support are available, students may still avoid a tool if it is perceived as difficult to use, slow to debug, or poorly supported by external resources. This implies that improving uptake may require more than simply increasing teaching hours or technical support through forums and Q&A sessions. Factors such as user interface design as reflected in free-text responses from respondents (T1–T3 in Q9 responses), access to examples and tutorials (T5), size of the user community (T6), and licensing constraints also play an important role in students' decisions.

Tool popularity further reinforces these trends. Students were asked how much popularity influenced their choice of tool, and as shown in Figure 2c, 38% reported it as extremely important, 35% as very important, and 16% as moderately important. Only 8% considered it slightly important and 3% not important. This indicates that perceived popularity, which is closely linked to the

availability of learning resources and external help with debugging, significantly affects students' selection of computational tools.

Students were also asked how well the teaching and learning of computational tools has prepared them for future industrial roles. As shown in Figure 2d, majority of respondents felt adequately prepared, while only 12% reported feeling unprepared. This reinforces the overall effectiveness of the current teaching approach to computational tools within the UCL chemical engineering programme. However, it also suggests that there may be value in reviewing the specific tools taught, particularly with regard to accessibility and long-term sustainability, and in considering a gradual transition towards more open-source alternatives.

Overall, the results suggest that students' preferences for computational tools in assessment are shaped by a complex interaction of pedagogical, technical, and social factors. While gPROMS is formally embedded in the curriculum, its lower adoption compared with MATLAB, Python, and Polymath appears to be driven primarily by perceptions of difficulty, limited external support, and usability rather than by deficiencies in teaching provision. Although the sample size does not allow for definitive conclusions, this study provides a preliminary basis for further investigation. It highlights the need for educators to consider not only what tools are taught, but also students' experience, particularly in relation to workload, error handling, and access to help beyond the classroom.

REFLECTION

My experience in teaching computational tools suggests that students often adopt a grade-driven approach to assessments and learning. As educators, we should encourage students to recognise the value of learning that extends beyond immediate assessment outcomes. This can be better addressed through the lens of constructive alignment [17], where intended learning outcomes, teaching activities, and assessment tasks are deliberately aligned. If assessments primarily reward correct outputs rather than the modelling process, tool selection, and problem-solving approach, students are more likely to prioritise efficiency over critical thinking. Aligning assessment with the intended development of computational thinking, modelling skills, and tool fluency may encourage students to engage more meaningfully with the tools taught, rather than selecting alternatives based solely on convenience or familiarity.

CONCLUSION

This study examined the factors influencing students' preferences for computational tools in chemical engineering assessments, using gPROMS ModelBuilder as a case study. The results show that students' choices are driven primarily by perceived ease of learning, availability of external support, and ease of debugging, rather than by formal exposure to a tool within the curriculum. Although gPROMS is taught and supported, it is used by only a small proportion of students when alternative tools are allowed. This suggests that technical capability alone is insufficient to ensure adoption. Instead, usability, community support, and perceived relevance play a significant role. These findings highlight the need for educators to consider not only what tools are taught, but how they align with students' learning priorities and future employability. Rather than focusing solely on introducing advanced tools, there is a need to rethink how these tools are integrated into the curriculum. This could include providing earlier exposure to modelling environments, embedding gPROMS more consistently across multiple modules, and designing scaffolded learning activities that gradually build confidence before assessments. In addition, integrating more guided examples, improving access to reusable code templates, and creating opportunities for peer learning may help reduce the initial barrier to engagement. Consideration could also be given to combining proprietary tools with open-source alternatives to support flexibility and independent learning. Future work will explore how these targeted curriculum adjustments influence students' confidence, engagement, and sustained use of computational tools.

ACKNOWLEDGEMENTS

The author acknowledges financial support from UCL Chemical Engineering department.

REFERENCES

1. Shacham M, Cutlip MB. Selecting the appropriate numerical software for a chemical engineering course. Computers & Chemical Engineering 23:S645-S648 (1999). https://doi.org/10.1016/s0098-1354(99)80158-6
2. Edgar, T.F., Himmelblau, D.M., & Lasdon, L.S. (2006). *Optimization of Chemical Processes* (2nd ed.). McGraw-Hill.
3. Udugama, I.A., Heintz, S., & Gernaey, K.V. (2023). *Digitalisation in chemical engineering education: Status, challenges and future perspectives*. Education for Chemical Engineers, 42, 1–16. https://doi.org/10.1016/j.ece.2022.10.003
4. Brenner A, Shacham M, Cutlip MB. Applications of mathematical software packages for modelling and simulations in environmental engineering education. Environmental Modelling & Software 20:1307-1313 (2005).

https://doi.org/10.1016/j.envsoft.2004.09.007
5. Lee, K., et al., 2014. Survey: MATLAB & MathCAD education in biochemical engineering. Chem. Eng. Educ. 48 (1), 59.
6. Richelle A, Bogaerts P. Systematic methodology for bioprocess model identification based on generalized kinetic functions. Biochemical Engineering Journal 100:41-49 (2015). https://doi.org/10.1016/j.bej.2015.04.003
7. Talbert, R., 2012. Learning MATLAB in the Inverted Classroom. Paper presented at 2012 ASEE Annual Conference & Exposition, San Antonio, Texas. 10.18260/1-2—21640.
8. Tang, C., 2021. Computer-aided linear algebra course on jupyter-python notebook for engineering undergraduates. J. Phys.: Conf. Ser. 1815, 012004–012009. https://doi. org/10.1088/1742-6596/1815/1/012004.
9. Inguva P, Bhute VJ, Cheng TNH, Walker PJ. Introducing students to research codes: a short course on solving partial differential equations in python. Education for Chemical Engineers 36:1-11 (2021). https://doi.org/10.1016/j.ece.2021.01.011
10. Lorandi Medina AP, Ortigoza Capetillo GM, Saba GH, Perez MAH, Garcia Ramirez PJ. A simple way to bring python to the classrooms. 2020 IEEE International Conference on Engineering Veracruz (ICEV) :1-6 (2020). https://doi.org/10.1109/icev50249.2020.9289692
11. Hunter JD. Matplotlib: a 2D graphics environment. Comput. Sci. Eng. 9:90-95 (2007). https://doi.org/10.1109/mcse.2007.55
12. Lee Y, Cho J. The influence of python programming education for raising computational thinking. IJUNESST 10:59-72 (2017). https://doi.org/10.14257/ijunesst.2017.10.8.06
13. Siemens Process Systems Engineering, 2022. gPROMS, https://www.psenterprise.com/products/gproms/process, 1997-2022.
14. Davis FD. Perceived usefulness, perceived ease of use, and user acceptance of information technology. MIS Quarterly 13:319-340 (1989). https://doi.org/10.2307/249008
15. Institution of Chemical Engineers (2021). 'Accreditation of chemical engineering programmes – A guide for education providers and assessors.' Available at https://www.icheme.org/media/26729/university-accreditation-guide.pdf [last accessed: 20th March 2026]
16. Kamel D, Tsatse A, Badmos S. Teaching computational tools in chemical engineering curriculum in preparation for the capstone design project. Systems and Control Transactions 4:2197-2202 (2025). https://doi.org/10.69997/sct.126494
17. Biggs J. Enhancing teaching through constructive alignment. High Educ 32:347-364 (1996). https://doi.org/10.1007/bf00138871

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Empirical survey among experts on the relevance of various criteria for optimizing modular electrolysis systems

Hannes Lange[a]*, Lucien Beisswenger[b], Daniel Erdmann[b], Isabell Viedt[a], and Leon Urbas[a]

a TUD Dresden University of Technology, Process-to-Order Group, Helmholtzstraße 10, 01069 Dresden, Saxony, Germany
b VDMA e.V., Power-to-X for Applications, Lyoner Straße 18, 60528 Frankfurt, Hesse, Germany
* Corresponding Author: hannes.lange1@tu-dresden.de.

ABSTRACT

Electrolysis systems must be constructed from multiple stack units. This modular system of stack units (modular electrolysis system) requires systematic optimization of its composition. This optimization depends on numerous, and sometimes conflicting, criteria. While such criteria have already been evaluated for modular plants in the process industry, they must be re-assessed in the context of modular electrolysis systems. To address this challenge, an expert survey was conducted within two research networks (H_2Giga and DECHEMA e.V. Research Network) and the VDMA P2X4A.The evaluation followed the two-stage process: After ranking the categories *costs, flexibility*, *process engineering* and *time-to-process* according to their overall importance a ranking of individual criteria within each category was conducted. The survey reveals a clear prioritisation: costs are in first place with 35.9%, followed by flexibility (25.4%), process technology (23.2%) time to process (15.4%). This ranking provides a structured overview of the most important factors guiding the optimisation of modular electrolysis systems. Another finding of the survey is how electrolysis systems should be operated in the future. Here, the focus was primarily on electrolysis systems related to fluctuating renewable energies, with a top-2-box result of 81% in each case. Based on this prioritisation, weightings for possible application in multi-criteria optimisation can now be derived. This makes it possible to determine suitable compositions of electrolysis stack unit systems when scaling systems and to refine optimisation efforts. The expert survey thus contributes to the optimised design of electrolysis systems and thereby supports the ramp-up of hydrogen production.

Keywords: modular electrolysis, modularization, requirements prioritization, expert survey, process equipment assemblies

INTRODUCTION

Electrolysis-based green hydrogen is widely regarded as a key enabler for decarbonizing energy-intensive industries because it provides a low-carbon feedstock and energy carrier that can substitute fossil-derived hydrogen and support new synthesis routes. Industrial deployment, however, requires electrolysis plants in the multi-hundred-megawatt to gigawatt range, which cannot be realized as a single stack and therefore relies on number-up of stacks and balance-of-plant subsystems [1–3]. As plant capacity increases, the overall system complexity rises disproportionately due to replicated auxiliaries and control layers, which makes engineering and integration a dominant cost and schedule driver.

Modularization has the potential to mitigate this complexity by structuring the plant into repeatable process units with defined interfaces, similar to the Process Equipment Assembly (PEA) concept that is established in modular process engineering [4, 5]. Despite the strong interest in modular electrolysis, requirements for electrolysis systems are often specified implicitly or inconsistently, and optimization targets for module boundaries remain unclear, which hampers the systematic development of modular architectures.

METHODOLOGY

This work applies a requirements-driven evaluation approach by repeating an expert survey that was originally developed for Process Equipment Assembly (PEA) selection in the process industry and transferring it to the German electrolysis sector. PEAs in modular plants can be easily adapted to changing boundary conditions. A PEA is a modular process unit that is mostly self-sufficient in terms of automation and provides a process engineering step [1]. The survey retained the same four high-level evaluation categories, namely 'cost', 'flexibility', 'process engineering', and 'time-to-process', while individual criteria were adapted where necessary to reflect electrolysis-specific characteristics and terminology [6]. Data collection was performed as a nationwide expert survey in Germany, with participation from 63 professionals/experts working in Research and Development (R&D) or engineering, sales, and project management and representing plant manufacturers, plant operators, component manufacturers, and research institutions within the Power-to-X for Applications of VDMA e.V. (VDMA P2X4A) network and the DECHEMA/H_2Giga project environment [7]. Figure 1 shows the distribution of experts across industries, and Figure 2 shows the distribution of experts within the field of work along the electrolysis value chain. Experts ranked the categories and criteria by relative importance using a rank-based point allocation comparable to the Simos approach [8], and the resulting point sums were normalized to obtain weights that can be interpreted as relative priorities for system optimization. With the completion of the expert survey, a rank could be determined for each category and criterion. To determine the rank, points are assigned for each rank. The points per rank are derived from the number of categories and are assigned in descending order. Subsequently, the weighting of the category or criterion is calculated according to formula (1) [6, 8]:

$$w_a\,[\%] = \frac{\sum_{i=1}^{n} p_{a,i}}{\sum_{i=1}^{n}(p_{a,i}+p_{b,i}+p_{c,i}+p_{d,i})} \qquad (1)$$

RESULTS

Ranking of optimisation criteria

The distribution of functional roles indicates a strongly engineering-oriented sample, as 62% of the participants work in R&D or engineering, while 11% work in sales, 8% in project management, and 19% in other roles. This composition is consistent with the current technology ramp-up phase, in which architectural and integration decisions are predominantly shaped by development and engineering organizations. With respect to organizational background, component manufacturers form the largest group with 46%, followed by plant manufacturers with 18%, research institutions with 12%, plant operators with 11%, and other organizations with 13%, which reflects the current industrial focus on stack and balance-of-plant scaling.

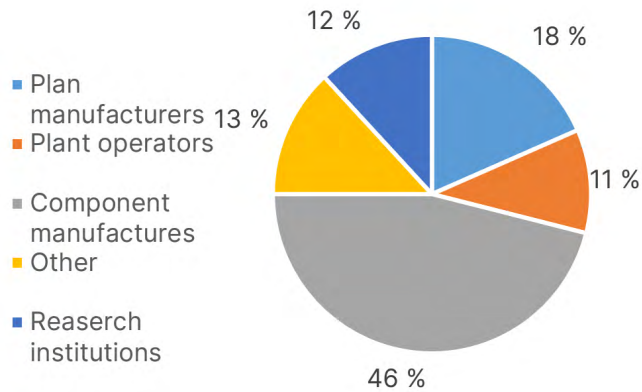

Figure 1. Branch of participating industry experts

The category weighting reveals that cost is the dominant requirement group with an overall share of 35.9%, whereas flexibility reaches 25.4%, process engineering reaches 23.2%, and time-to-process reaches 15.4%, which jointly define a quantitative priority profile for modular electrolysis optimization. Figure 3 shows the entire distribution of categories with their criteria in a sunburst diagram.

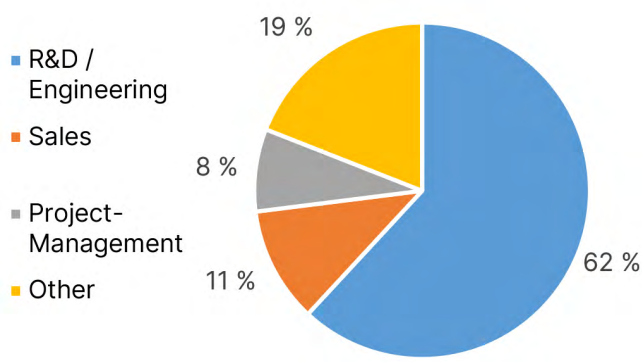

Figure 2. Work field of participating industry experts

Within the cost category, operating and maintenance costs contribute 19.08% to the overall weighting, corresponding to 53.1% of the cost category, while investment costs contribute 16.85% and therefore account for 46.9% of the cost category, which indicates that both CAPEX and OPEX reduction are considered nearly equally important.

Flexibility is primarily driven by capacity flexibility with 8.72% and dynamic operation capability with 8.28%, while location flexibility reaches 4.66% and product specification flexibility reaches 3.77%, which confirms

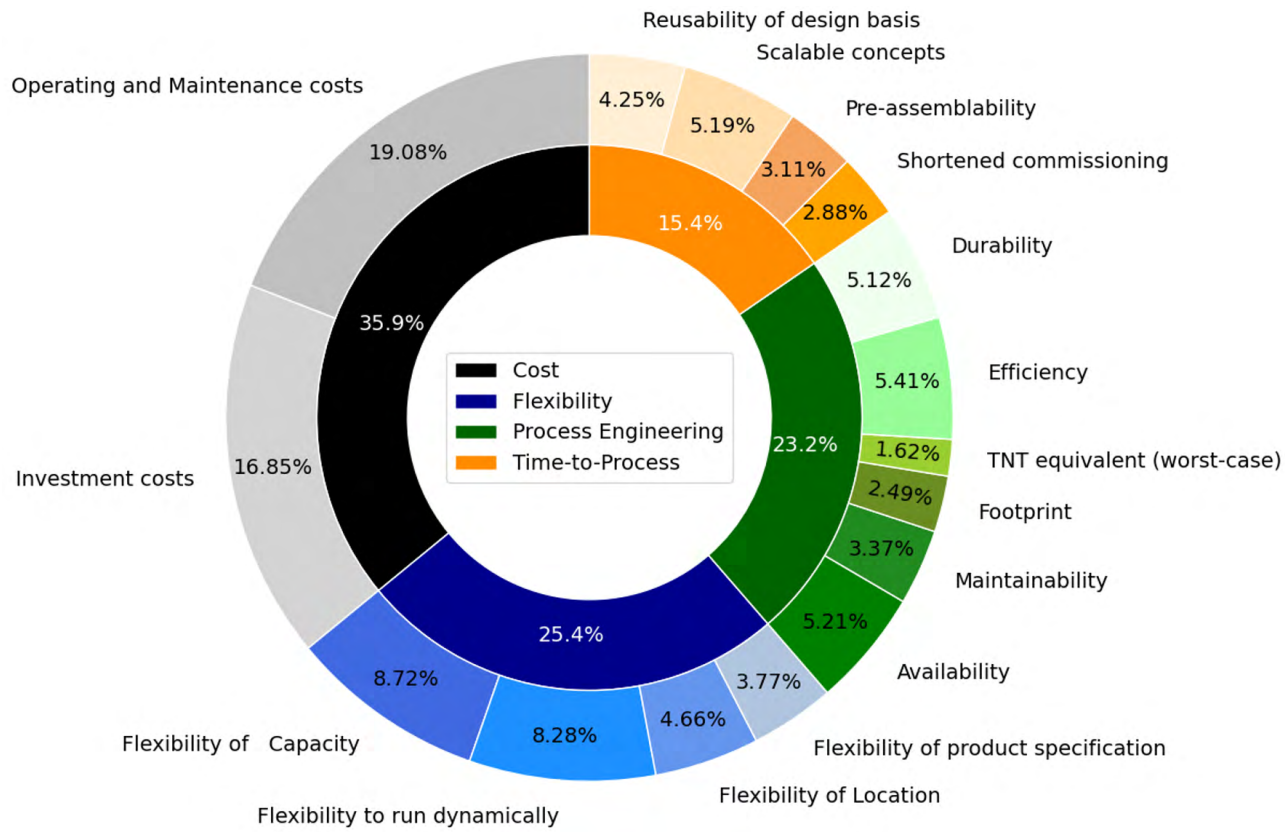

Figure 3. Sunburst diagram: Categories are shown in the inner circle. These are identical to the categories from the process industry. The criteria for later optimization of electrolysis systems are in the outer circle.

that electrolysis plants are expected to provide both scalable capacity and load-following behavior in volatile electricity systems.

Process engineering requirements are dominated by efficiency with 5.41%, availability with 5.21%, and durability with 5.12%, whereas maintainability reaches 3.37%, footprint reaches 2.49%, and the TNT-equivalent worst-case safety indicator reaches 1.62%, which highlights the importance of robust long-term operation and energy conversion performance.

Time-to-process is led by scalable concepts with 5.19% and reusability of the design basis with 4.25%, while pre-assemblability reaches 3.11% and shortened commissioning reaches 2.88%, which suggests that the main schedule benefits are expected from reuse and scaling strategies rather than from commissioning acceleration alone.

A comparison with the original process-industry survey shows that cost remains the most important category in both sectors, yet electrolysis assigns substantially higher relevance to flexibility and correspondingly lower relevance to time-to-process. In the process-industry results, cost accounts for 31%, time-to-process for 27%, process engineering for 25%, and flexibility for 17%, whereas the electrolysis survey shifts flexibility to the second rank and places time-to-process last. This shift can be attributed to the system role of electrolysis as a Power-to-X converter that must operate efficiently under fluctuating renewable power availability, which makes operational flexibility a first-order design driver rather than a secondary feature. Both the process industry and the electrolysis sector place a significant focus on flexible production volumes. However, in the process industry, flexibility is equated with the reconfiguration of production plants and varying product specifications that can be manufactured within a single plant [9]. In contrast, flexibility in the electrolysis sector is understood as the ability to respond to fluctuating power inputs while maintaining constant product specifications [10]. Figure 5 shows a summary comparison of the categories from [6] from the process industry compared with the results of the survey from this study for the electrolysis industry.

Expert assessment of grid connection concepts for electrolysers

This section summarizes expert ratings on the relevance of four electrical grid connection and operation concepts for electrolyzer systems. Respondents

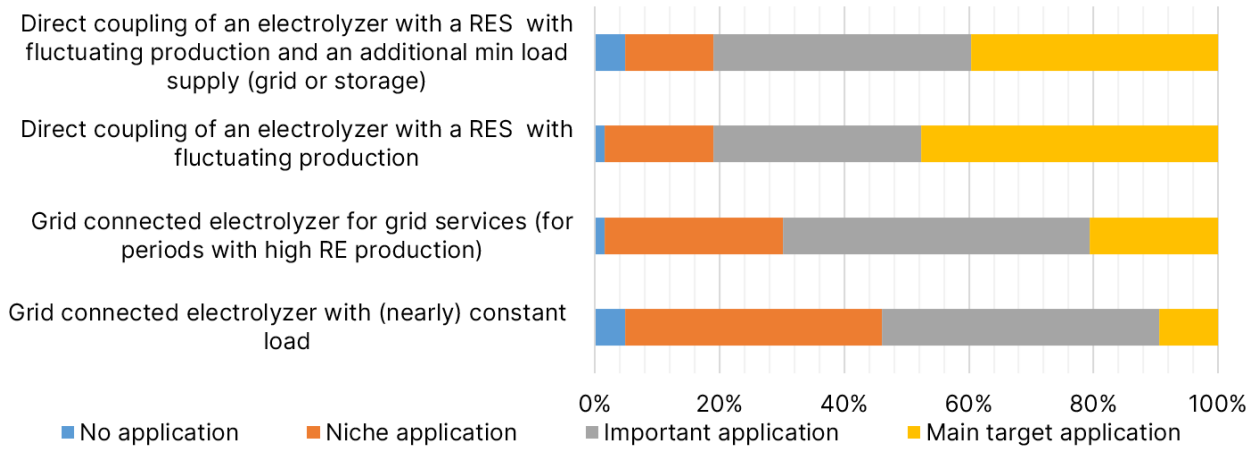

Figure 5. Evaluation of the expert survey on the future electrical connection and grid integration of electrolysis systems

classified each concept on an ordinal four-level scale ("No application", "Niche application", "Important application", "Main target application"). The provided dataset contains normalized response shares per concept (each row sums to 1.0), corresponding to 63 expert responses. For comparative interpretation, an additional mean relevance score was computed by mapping the four levels to 0–3 (0 = No application, 3 = Main target application).

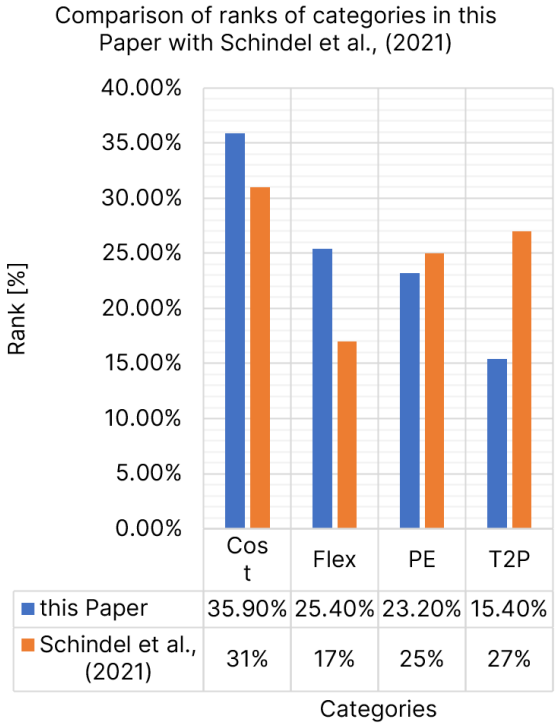

Figure 4. Comparison of categories with [6]. Abbreviations: Flexibility (Flex), Process Engineering (PE), Time-to-Process (T2P)

Results and ranking

Across all concepts, experts clearly prioritized configurations that enable direct integration with renewable energy sources (RES). The concept "direct coupling with fluctuating RES production" achieved the highest share of "Main target application" (47.6%) and the highest mean relevance score (2.27/3). A closely related hybrid concept—direct coupling with fluctuating RES production plus an additional minimum-load supply via grid or storage—was rated almost equally relevant (39.7% "Main target", mean score 2.16/3). Both direct-coupling options reach an identical top-2-box share of 81.0%, indicating broad consensus that these concepts should be treated as important-to-primary use cases rather than edge scenarios. The top 2 box share is the proportion of respondents who scored one of the two highest answer categories on an ordinal scale - e.g. Select "Main target application" and "Important application". It is therefore an aggregated agreement/positivity figure in percent that compresses the "high" rating of a statement. The higher the top-2-box share, the higher the consensus strength in the upper valuation range.

Grid-connected operation aimed at grid services (operation in periods with high renewable generation) occupies an intermediate position (mean score 1.89/3, Top-2 box 69.8%). While only 20.6% selected it as the main target, nearly half of respondents (49.2%) still considered it an important application, suggesting that flexibility-oriented grid-connected operation is widely valued but not universally seen as the dominant design point.

In contrast, a nearly constant-load, grid-connected electrolyzer is the least preferred concept (mean score 1.59/3) and is most frequently characterized as niche (41.3%). Only 9.5% rated constant-load operation as the main target application. This pattern indicates a shift away from a baseload-centric view toward operational

regimes dominated by variability, flexibility, and power-system interaction. The results can be found in Figure 4.

Interpretation and implications for system requirements

The observed ranking strongly suggests that future electrolyzer system design—and, by extension, requirement prioritization—should be anchored in dynamic operating conditions rather than steady-state operation. High relevance assigned to direct RES coupling implies that experts expect electrolyzers to (i) tolerate strongly fluctuating power input, (ii) operate efficiently across a broad load range, and (iii) maintain performance and lifetime under frequent transients (start/stop cycles, ramping, and partial-load operation [10]). The comparable relevance of the "minimum-load supply" variant indicates that practitioners anticipate technical or economic constraints that motivate stabilizing measures (e.g., minimum turndown limits, stack protection strategies, or process integration requirements), achievable through either grid backup or buffering (electrical storage, hydrogen storage with operational control, or hybrid supply architectures).

The relatively high 'Important' share for grid-service operation highlights an additional requirement cluster: electrolyzer systems are expected to interact with the power system beyond simple consumption. This typically translates into needs for advanced control strategies capable of responding to market or grid signals (e.g., absorbing renewable peaks, participating in flexibility products, or managing curtailment). Even if not the dominant main-target case for all respondents, its broad 'important' rating suggests that 'grid-interactive' capability can be a meaningful differentiator in system architectures and should be considered in modular interface and control design.

Finally, the low prioritization of nearly constant-load operation implies that optimizing solely for steady-state efficiency at a single design point may not match the dominant deployment expectation in the surveyed expert community. Instead, robust performance across operating regimes - and the ability to integrate fluctuating RES supply - emerges as a central design intent. For a modularization-oriented development approach, these findings support emphasizing modules and interfaces that enable (i) flexible power conversion and grid/RES interfacing, (ii) transient-capable process control.

SURVEY LIMITATIONS

One limitation of the evaluation is that it only contains aggregated distributions at respondent level. Subgroup effects, such as the examination of connections between interest groups and the level of a rated category or criterion, have not yet been taken into account here. This means that statistical conclusions that go beyond a descriptive ranking cannot be evaluated here. Another limitation of the survey is the broad spectrum of perspectives. However, this is still focused on technology development and component supply, which is due to the expert networks of VDMA P2X4A and the DECHEMA network in the H2Giga lead project. This should be taken into account when interpreting the relative weighting of the technical criteria.

CONCLUSIONS

The transferred survey method provides an empirically grounded prioritisation of requirements for modular electrolysis systems in Germany and thereby clarifies optimisation targets for module decomposition. The optimisation criteria should focus on the efficiency, dynamic operation and levelized cost of hydrogen of the electrolysis system. The modular approach of small-scale electrolysis systems should be transferable to large-scale systems. In particular, electrolysis systems that can cope with fluctuating power inputs and operate efficiently over a wide 'Load Flexibility' range should be investigated.

The approaches to modularising process engineering systems pursue different goals in the process industry and in the electrolysis industry. Cost reduction remains the overarching objective, but the high weighting of flexibility and the strong process engineering emphasis on efficiency imply that modular concepts must simultaneously support dynamic operation, reliable performance and scalable concepts. In the process industry, the approach to standardised modules seems to be more about the rapid realisation of products in order to secure market share.

OUTLOOK

Future work should implement the derived weightings in the optimisation of electrolysis systems and evaluate the modular system approach for scaling under realistic fluctuating operating profiles. To support this, a standardised test signal/operating profile combined from different components (e.g. from a step, ramps, a pseudo-random binary sequence/'sawtooth function' and a synthetic solar profile) can be developed to optimise consistency across different electrolysis system sizes. In this way, modular heterogeneous electrolysis systems of different technologies and power ratings can be investigated in terms of system efficiency, operating flexibility and levelized cost of hydrogen. The consistent test signal allows investigation of whether the modular system approach represents a scalable concept.

Another field of research is the investigation of hybrid electrolysis systems with storage integrated into the process, consisting of a battery-electric and a hydrogen

storage system. After all, the experts rated the 'grid-connected electrolyser for grid services (for periods with high RE production)' with a top-2-box value of 69.8%, which argues for a high storage share that can be divided among different technologies.

ACKNOWLEDGEMENTS

We would like to thank the German Federal Ministry of Research, Technology and Space and the project management organization Jülich for their financial support within the framework of the eModule research project (FKZ 03HY116A) of the H_2Giga lead platform. Furthermore, the P2O-Lab Learning Factory of the TUD Dresden University of Technology is thanked for its support.

REFERENCES

1. Lange H, Klose A, Beisswenger L, Erdmann D, Urbas L. Modularization approach for large-scale electrolysis systems: a review. Sustainable Energy Fuels 8:1208-1224 (2024). https://doi.org/10.1039/d3se01588b
2. Lange H, Gro?e M, Viedt I, Urbas L. Modular and heterogeneous electrolysis systems: a system flexibility comparison. Systems and Control Transactions 4:759-764 (2025). https://doi.org/10.69997/sct.177861
3. Lange H, Meltzner FM, Große M, Viedt I, Klose A, Urbas L. Design of a modular and heterogeneous electrolysis system to increase flexibility in the use case of direct iron reduction. Energy Reports 15:108801 (2026). https://doi.org/10.1016/j.egyr.2025.11.060
4. Klose A, Merkelbach S, Menschner A, Hensel S, Heinze S, Bittorf L, Kockmann N, Schäfer C, Szmais S, Eckert M, Rüde T, Scherwietes T, da Silva Santos P, Stenger F, Holm T, Welscher W, Krink N, Schenk T, Stutz A, Maurmaier M, Stark K, Hoernicke M, Unland S, Erben S, Kessler F, Apitz F, Urbas L. Orchestration requirements for modular process plants in chemical and pharmaceutical industries. Chem Eng & Technol 42:2282-2291 (2019). https://doi.org/10.1002/ceat.201900298
5. VDI 2776-1:2020. Verfahrenstechnische Anlagen - Modulare Anlagen - Grundlagen und Planung modularer Anlagen 2020.
6. Schindel , Polyakova , Harding , Weinhold , Stenger , Grünewald , Bramsiepe . General approach for technology and process equipment assembly (PEA) selection in process design. Chemical Engineering and Processing - Process Intensification 159:108223 (2021). https://doi.org/10.1016/j.cep.2020.108223
7. Federal Ministry of Research, Technology and Space. H2Giga. H2Giga Flagship Proj Aims Mass Prod Electrolysers Prod Scale- Hydrog 2026. https://www.wasserstoff-leitprojekte.de/projects/h2giga (accessed January 28, 2026).
8. Simos J. Evaluer l'impact sur l'environnement, une approche originale par l'analyse multicritère et la négociation. Lausanne: Presses polytechniques et universitaires romandes; 1990.
9. Riese J, Grünewald M. Challenges and opportunities to enhance flexibility in design and operation of chemical processes. Chemie Ingenieur Technik 92:1887-1897 (2020). https://doi.org/10.1002/cite.202000057
10. Lange H, Klose A, Lippmann W, Urbas L. Technical evaluation of the flexibility of water electrolysis systems to increase energy flexibility: a review. International Journal of Hydrogen Energy 48:15771-15783 (2023). https://doi.org/10.1016/j.ijhydene.2023.01.044

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Mapping "Digital Chemical Engineering" in the UK: A Sector-Level Audit of IChemE MEng Curricula

E. Routoula[a], M. Mohammad Zadeh[b], M. Granollers Mesa[c], M. Malekshahian[d], D. Dikicioglu[e], M. Pollock[f], and M. Zandi[a]*

[a] University of Sheffield, School of Chemical Materials & Biological Engineering, Sheffield, UK
[b] University of Nottingham, Department of Chemical and Environmental Engineering, Nottingham, UK
[c] Aston University, School of Engineering and Innovation, Department of Chemical Engineering and Biotechnology, Birmingham, UK
[d] University of Manchester, Department of Chemical Engineering, Manchester, UK
[e] University College London, Department of Biochemical Engineering, London, UK
[f] University College London, Department of Chemical Engineering, London, UK
* Corresponding Author: m.zandi@sheffield.ac.uk.

ABSTRACT

The rapid shift toward Industry 4.0 and data driven manufacturing has prompted universities to reshape chemical engineering programmes, yet the scope and coherence of "Digital Chemical Engineering" (DCE) within UK curricula remain unclear. This study qualitatively maps DCE provision across five IChemE accredited MEng degrees to identify which digital skills are taught, how they progress across programme stages, how skills are distributed between core and elective content or taught versus applied learning, and how well provision aligns with accrediting frameworks. The analysis is structured around eight domains: (D1) programming and computation; (D2) data literacy and statistics; (D3) process modelling and simulation; (D4) optimisation and process systems engineering; (D5) control, automation and instrumentation; (D6) AI, machine learning and digital twins; (D7) software engineering practices; and (D8) data governance, ethics and cybersecurity. Results show institution dependent digital skills coverage. Some programmes introduce a broad digital skill set from Year 1, while others concentrate digital content in later years or adopt a more progressive distribution. Across institutions, D7 is consistently underrepresented. Foundational skills in D1, D2 and D8 are typically embedded throughout programmes, whereas more advanced domains D3 to D6 are introduced at later stages. Domains D1 to D4 demonstrate a stronger balance between teaching and application, while D5 to D8 are predominantly taught, with application increasing in later years. Optional fourth year modules play a significant role in shaping graduate digital profiles. Overall, DCE provision broadly aligns with IChemE and AHEP4 guidance, where expectations for D1 to D4 are more explicitly articulated

Keywords: Education, Industry 4.0, Modelling and Simulations, Choose an item.Choose an item., Digital Skills, Curriculum, Digital Chemical Engineering

INTRODUCTION

Digitalisation and Industry 4.0 are reshaping chemical engineering practice across the industries, with increasing reliance on data-driven decision-making, advanced modelling, automation, and digital twins [1]. These developments are transforming how processes are designed, operated, and optimised, and are redefining the skills expected of graduate chemical engineers. Historically, undergraduate programmes have introduced students to well-established tools such as process simulation and basic programming, with the expectation that deeper digital expertise would be developed post-graduation in the workplace. However, recent studies and employer surveys indicate that this model is no longer sufficient. Employers increasingly report skills gaps in digitalisation and now seek explicit evidence of digital competence during graduate recruitment [2].

In response, universities are embedding digital skills like modelling, process simulation, coding and statistics

more deliberately within chemical engineering curricula [3], often in collaboration with industry and professional bodies. Nonetheless, there remains considerable uncertainty regarding what constitutes appropriate "Digital Chemical Engineering" (DCE) coverage, how digital skills should progress across degree programmes, and how consistently these expectations are implemented across institutions. While prior work has explored the pedagogical value of specific digital tools and surveyed academic and industrial perspectives on digitalisation in education, these studies largely focus on individual modules, tools, or localised curriculum reform [4-7].

Within the UK context, this issue is closely linked to accreditation requirements. IChemE accreditation establishes shared learning outcomes and standards for professional formation, while allowing flexibility in curriculum design. Accreditation guidance highlights importance of systems thinking, and the ability to work with modern engineering tools and data-rich environments [8]. However, it provides limited explicit direction on the depth, sequencing, or balance of digital skills across the curriculum, potentially leading to uneven provision between institutions despite a common accreditation framework.

This study provides a sector-level mapping of DCE coverage across MEng programmes in the UK that are accredited by the Institute of chemical Engineers (IChemE) and in alignment with the Engineering Council's Accreditation for Higher Education Programmes guidelines (AHEP4). This study evaluates the baseline for digital readiness in UK chemical engineering education and provides evidence to inform departmental curriculum review and the future evolution of accreditation guidance. The purpose of this work is to map DCE, provide an overview of the digital landscape, and draw potential trends, helping practitioners understand more about their curriculum. This work is not providing recommendations, ranking one approach over another or setting a precedence.

METHODOLOGY

Data Collection and Curriculum Mapping

Following from the outcome of a survey on the views on key skills for chemical engineers from academic and industrial stakeholders [2] and using informed judgement and discussions across contributors who are all heavily involved with curriculum design in their respective institutions, DCE provision was classified across eight domains: (1) programming and computation: ability to understand, write, and execute code. (2) data literacy and statistics: understanding of statistical principles, use of data visualisation for effective data communication. (3) process modelling and simulation: use of specialised software to develop models and simulate unit operations. (4) optimisation and process systems engineering: use of systems engineering principles and algorithms to identify best possible variables. (5) control, automation and instrumentation: design and implementation of control systems to monitor and maintain process variables towards safe and stable plant operation. (6) artificial intelligence, machine learning, including digital twins: use of tools and software to identify patterns, development of virtual, real-time depiction of a process or plant that can be updated using process data, predict performance and diagnose issues in a simulated environment. In this context digital twins represent artificial intelligence or machine learning enabled, data-driven operational replicas of the process. (7) software engineering practices (version control, testing, reproducibility): tracking and managing changes to code over time, reverting to previous versions, ensuring reliable reproducibility. (8) data governance, cybersecurity, and ethics around digital tools: understand the importance of data storage, quality, security and access, address moral implications around the use of digital tools.

It is important to point out that the domains focus on elements of the curriculum directly related to digital skills, rather than fundamental chemical engineering knowledge that would be taught in traditional core modules. The Universities offering accredited degrees in chemical engineering or related courses at BEng and MEng level that participated in this curriculum mapping activity were: Aston University, University of Manchester, University of Nottingham, University of Sheffield, University College London (two departments each with their own degree programmes; Chemical Engineering (CE) and Biochemical Engineering (BE)). Representatives from each participating university in this study mapped the DCE curriculum provision for their offered courses in chemical engineering against each domain, per year of study (years 1-3 for BEng courses and 1-4 for MEng courses). This was done by analysing publicly available information in module descriptors and learning outcomes, supplemented on occasion by additional knowledge of modules and assessment based on their knowledge of teaching specific courses or from communications with academics involved in teaching specific courses. Each institution was responsible for providing and checking data for their course(s), with all data being collated in a shared document. DCE curriculum mapping was done in a qualitative manner at this stage, indicating a "yes/no", and providing examples or comments of how each domain content is covered where relevant. With the understanding of content being covered formally in the curriculum, i.e. core/taught and formatively assessed as a learning outcome, and informally, as by-product/applied, i.e. taught but not assessed, or implicitly expected in later years of study, mapping was designed to recognise both dimensions. It is important to note that the mapping was done by a representative academic of each university rather than all module leaders of all courses,

hence there might have been details on specific learning activities or assessments that were overlooked or underestimated. This could result in the mapping not being completely definitive, but nevertheless sufficient to highlight trends and perhaps differences in DCE coverage across the examined curricula. In addition to curriculum mapping, publicly available IChemE and AHEP4 accreditation guidelines [8, 9] were mapped against the domains.

Data Analysis and Visualisation

Once mapping was completed qualitatively, it was assessed and validated through comprehensive discussions across all collaborators to ensure as much uniformity as possible, a scoring system was introduced, to bridge the qualitative mapping with a quantitative approach. An "instructional intensity" metric was developed to distinguish between the presence of a tool and the formal pedagogical delivery of a skill. Each module-year intersection was evaluated against two criteria:

- **Core/Taught (Primary Intensity):** The domain is a formal part of the curriculum with explicit, assessed Learning Outcomes. Students receive dedicated instruction (e.g., a Python module).
- **Byproduct/Applied (Secondary Intensity):** The domain is utilised as a tool for another engineering subject (e.g., using Excel in a lab report) without new formal instruction.

For the purpose of the **Digital Maturity Curve**, the Index focuses on **Core/Taught** instances, representing the "Instructional Load" of new digital competencies introduced at each level. The quantification process followed a three-step aggregation logic:

1. **Binary Mapping:** Each of the eight domains was assigned a binary value (1 for Core instruction, 0 for absent) for every university at every level.
2. **Summation and Normalisation:** For a given Academic Year (Y), the raw intensity score (S_y) was calculated as the sum of active domains across all participating institutions (n):

$$S_y = \sum_{i=1}^{n} \sum_{d=1}^{8} D_{i,d}$$

Where $D_{i,d}$ is the binary state of Domain d at University i. Sy was used for calculating scores and generating the Digital Coverage Heat Map (Figure 1) and the Digital Maturity Curve (Figure 4).

3. **Weighted Aggregation:** To account for elective specialisations in Year 4, the score was bifurcated into Core Intensity (mandatory/core modules) and Potential Intensity (elective modules).

For the construction of the Digital Coverage Heat Map, the scale range was defined from 0 to 3. (0) corresponded to absence of evidence, moving to evidenced presence of domain activities as curriculum by-products with no formal assessment (1) and as optional modules including assessed learning outcomes (2), and ending with evidenced domain activities in core modules with formally identified learning outcomes (3). For the accreditation guideline mapping, absence of a specific domain was marked with 0, 1 corresponded to vague reference and 2 corresponded to clear evidencing of a domain to the accreditation guidelines. Data was processed and analysed using Python (version 3.14.2) for visualisation and trend identification per domain and per year. It is important to note that the quantitative analysis (scoring) was performed by one member of the team, to ensure consistency.

RESULTS AND DISCUSSION

Analysis of the digital curriculum heatmap (Figure 1) revealed substantial variation in the breadth and depth of digital skills coverage across universities and degree levels. Across all five universities, digital domains traditionally embedded within chemical and process engineering curricula, namely process modelling and simulation (D3) and control/automation and instrumentation (D5), were the most consistently represented. These domains were commonly embedded as core assessed modules or optional assessed offerings particularly in senior programme years. Data literacy and statistics (D2) was also widely present across almost all universities, typically embedded within assessed modules. In contrast, under emerging digitalisation domains (e.g., AI/ML and digital twins (D6), software engineering practices (D7), and data governance, ethics and cybersecurity (D8)), coverage was considerably more even. These domains were frequently either absent or delivered indirectly without assessment and appeared in isolated cases rather than as integrated features of coherent programme pathways.

Institution-level patterns further highlight the uneven digital landscape. Aston demonstrated relatively strong provision in D1-5, and D8 from the early years through to MEng level, yet provided limited coverage of D6, with some provision in Year 1 and more application or presence in optional modules in later years. Manchester showed a similar pattern, with strong coverage in D3-5 and limited progression in emerging domains; D6 appeared only in two units with formally identified learning outcomes. At Nottingham, D2 and D3 are embedded and assessed throughout the programme. D1 develops progressively, while D4 and D5 are mainly introduced in later years. In contrast, D6 is only indirectly addressed, and D8 appears only at MEng level. Sheffield exhibited strong coverage of D2-3 and D8 throughout the curriculum and intermittent coverage of D1, D4-5, while D6 and D7 had

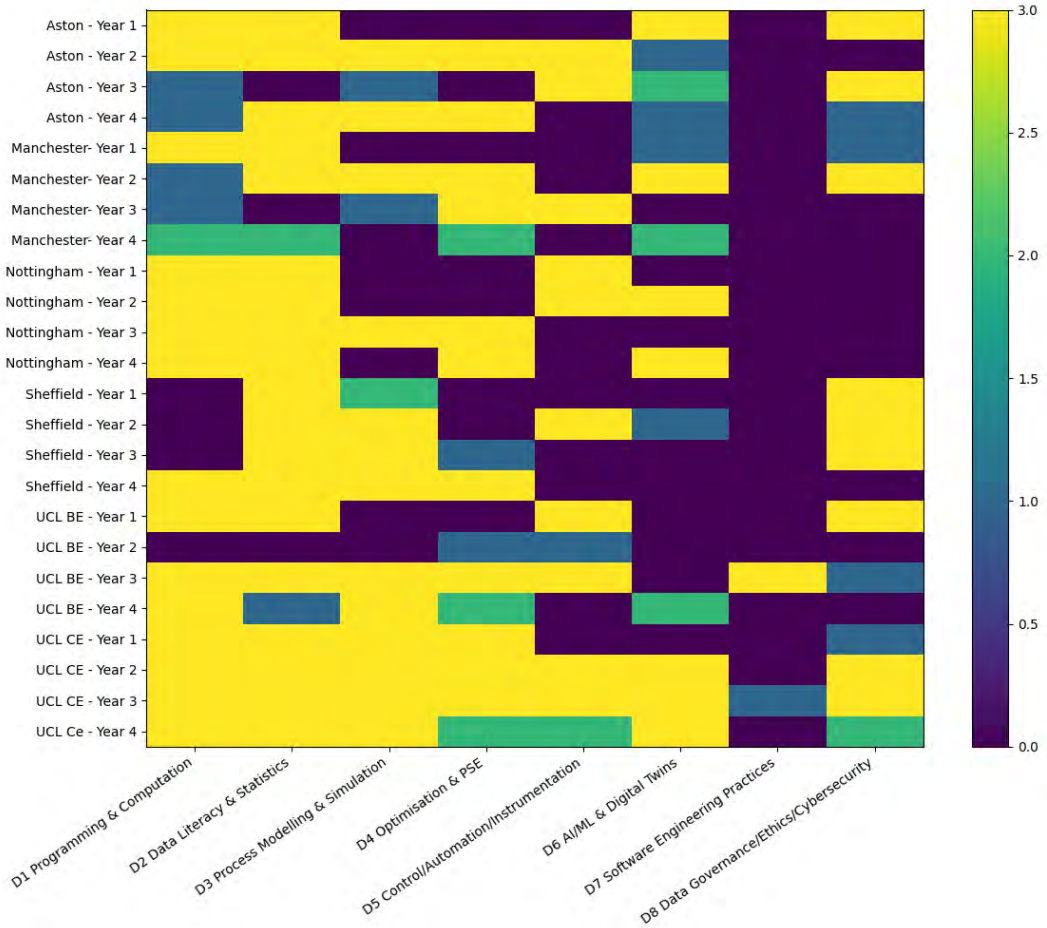

Figure 1. Heatmap of curriculum digital coverage across examined programmes at the five Universities, per year of study and domain. Scale range is from 0/dark blue indicating absense of coverage to 3/yellow indicating skill covered by formally identified learning outcomes.

limited evidence of engagement. UCL, across both biochemical (BE) and chemical engineering (CE) routes, presented strong late-stage coverage in D1-5, but did not incorporate substantive coverage of D7. UCL BE's biology-focused modules differ from traditional chemical engineering due to the department's distinct remit and scope, and while a digital agenda is embedded throughout the programme, its unique nature leads to observable differences compared with conventional chemical engineering curricula. UCL CE showed strong inclusion of D1-4 throughout all years of the degree programme, whilst domains D5, 6 and 8 were included in Year 2 onwards exemplifying a significant digitalisation emphasis.

Digital skills coverage demonstrates institution-dependent sequencing. Some universities (e.g., UCL) introduce a broad range of digital competencies from Year 1, whereas others (e.g., Nottingham, Sheffield) concentrate digital learning in senior years and a third group (e.g., Aston, Manchester) distribute coverage more progressively across the programme.

Further analysing the spread of digitalisation skills in the degree programme, Figure 2 shows an overall mean convergence score for each domain in each year across all degree programmes considered in the study. From this four approaches to how domains are covered can be identified:

- Coverage is consistent and significant across the degree programme - D1, D2, D8
- Coverage is initially low and is built up throughout the degree programme - D3, D4
- Coverage is focused most significantly in a specific year or two - D5, D6
- Coverage is very limited - D7.

It is clear that core digital skills of D1 Programming and Computation and D2 Data literacy and statistics are covered immediately in year 1 and that these skills are continually included and built upon throughout the degree programme. This is reasonable since these skills align with core mathematical knowledge that students

develop and apply throughout their studies.

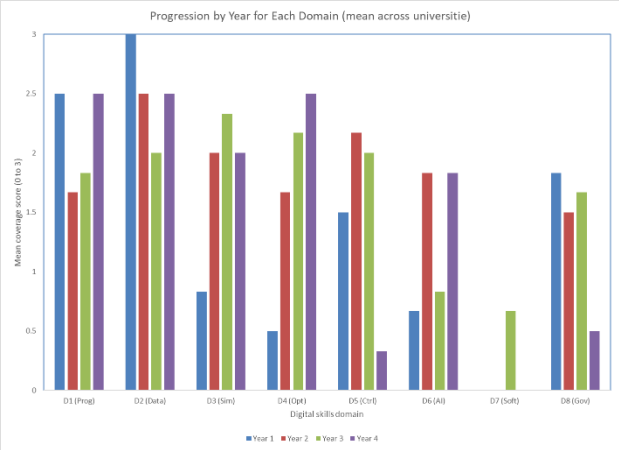

Figure 2. Progression of domain coverage per year of study across D1-8. D1: programming and computation, D2: data, literacy and statistics, D3: process modelling and simulation, D4: optimisation and PSE, D5: control automation and instrumentation, D6: artificiall intelligence, machine learning and digital twins, D7: software engineering practices, D8: data governance, ethics and cybersecurity.

Interestingly, Figure 2 shows that Data Governance, Ethics, Cybersecurity (D8) is represented in each year with almost equal coverage of D8 in year 1 compared to year 2 and 3, although with a drop in year 4. The heat map in Figure 1 shows D8 to be included in the degree programme when many digital skills have yet to be taught or developed. This could be explained through the open-ended interpretation of data governance, ethics and cybersecurity. It is likely that some form of cybersecurity and GDPR training is included in introductory courses as part of induction training. However, specific relation of these domains to chemical engineering, such as the impact of a breach of cybersecurity in the safe operation of a chemical plant and recommended mitigating actions are likely not included until later years, if at all. Equally, interpreting ethics is also considerably open-ended. It can be acknowledged that some form of ethics is included at the start of the degree programme, specific inclusion of ethics related to digital skills cannot be guaranteed but equally cannot be ruled out. For example, the ethics of the use of AI in one's studies or the ethical impact of using AI on a workforce in a mini design problem could be included at the beginning of the degree programme. Specific information to rule these considerations in or out based on information available in published module descriptions and intended learning outcomes, which are often limited in number and format by individual university regulations, can be the cause of varied interpretation of this domain.

Considering domains that are initially covered less and built upon throughout the degree programme seem reasonable when the specific domains are examined. Process Modelling and Simulation (D3) and Optimisation and PSE (D4) are more advanced applications of digital skills once more fundamental digital, mathematical and process unit operation skills are introduced earlier in the programme. Equally, considering domains that have a focus in a specific year or two, such as Control/Automation/Instrumentation (D5) and AI/ML and Digital Twins (D6) is also reasonable since this aligns with material and prior knowledge to which these domains relate. For example, Control/Automation/Instrumentation is often introduced, taught in detail and applied once initial concepts have been introduced, but are then less of a focus in electives in later years. On the other hand, considering D6, there seems to be a focus in year 2, which is not then revisited until digital-focused electives in the final year.

Finally, there is consideration of domains for which there is little recognition throughout the degree programme, such as Software Engineering Practices (D7) with only some indication of coverage in Year 2 or 3 at a couple of universities. This may be down to how the domain itself is formulated and the implicit inclusion of this domain to achieve learning outcomes of other domains. For example, software engineering practices, such as tracking and managing changes to code or documents over time, are not an explicit learning outcome of specific modules, but successful software management is needed to carry out large, open-ended, group design projects and meet the D3-related learning outcomes.

Looking at the DCE coverage across core/taught components where defined learning objectives were used to assess student learning, and curriculum by-products/applied where understanding of tools and knowledge taught at earlier years was assumed and students were expected to apply it to solve problems, there is an interesting trend revealed (Figure 3). There is prominence of distinct Digital Domains in both teaching and of application and the striking underrepresentation specific digital domains in the curricula.

Software Engineering Practices (D7) including open-source code sustainability, version control, testing, reproducibility are minimally integrated across the programmes investigated. This is followed by Data governance, Cybersecurity and Ethics (D8) and AI/ML and Digital Twins (D6) domains, which are primarily taught as part of the core curriculum. In contrast, Programming and Computation (D1), Data literacy and Statistics (D2), Process modelling and simulation (D3) domains are overrepresented in both the taught and the applied components of the curriculum, followed by Optimisation and Process Systems Engineering (D4) and Control, Automation and Instrumentation (D5).

An interesting aspect of our analysis elucidated the sensitive balance between the taught and the applied

components of the content covered in these domains.

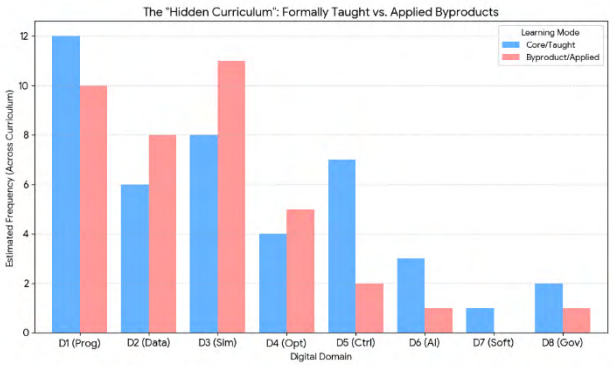

Figure 3. Comparison between DCE coverage across all domains for core/taught components and components evidenced as byproducts.

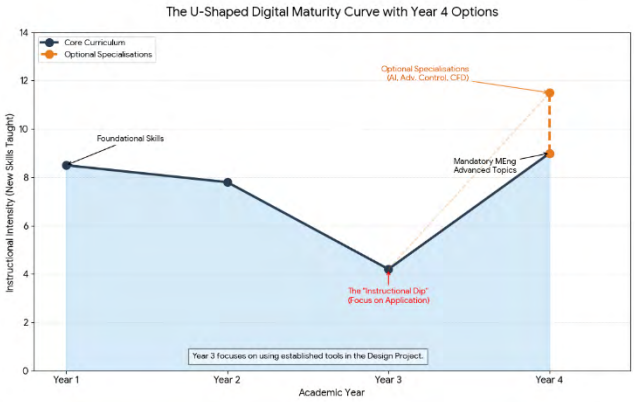

Figure 4. Digital “maturity” curve of DCE provision.

We note that the digital content is not only taught and assessed but has become an integral part of the learning curriculum since they are frequently utilised as part of the application demanding their use, even without necessarily being assessed directly. We observe a heavy focus on application in domains D2-4, possibly emerging from the nature of the digital capabilities covered by these domains. Concepts such as data science, process modelling and simulation and optimisation, once basic concepts are introduced, mostly remain applied concepts in the curricula in a range of areas, consequently serving to support digital capacity building and training primarily through application. These domains are closely followed by D1. While D1 and D3 find a wide range of application areas as expected, they are also an integral component of the core teaching of the curriculum and are widely implemented across different programmes.

Looking at the totality of DCE coverage from a progress-per-year point of view (Figure 4), another interesting trend emerges across the four years of study, highlighting a “U” shape maturity trendline. The left peak (Years 1 and 2) is the "Input Phase." Universities front-load digital skills (Python, MATLAB, Excel, Statistics) to ensure students have the "mathematical literacy" required for engineering science modules. In this phase, new digital competencies are introduced. The curve dips in Year 3 because the focus shifts from learning new digital skills to applying existing ones (primarily for the high-stakes Design Project).

Students are at their busiest digitally, but they aren't learning new software; they are refining their use of Aspen Plus or HYSYS. From a "new instruction" perspective, digital maturity appears to stall before spiking again in Year 4 with advanced industrial specialisations like AI and Molecular Modelling.

For MEng students, the curriculum in many cases introduces high-level concepts that are currently industry-leading, such as Machine Learning, Computational Fluid Dynamics, and Digital Twins. While the core Year 4 curriculum already sees a resurgence in digital instruction (e.g., mandatory Advanced Process Control, Computational Modelling, or Sustainability Modelling), the dotted line shows the peak maturity reached through electives like Advanced Machine Learning (AI), Computational Fluid Dynamics, and Digital Twin Design. It should be noted that the "gap" between the solid and dotted points at Year 4 highlights that digital competency in the final year is highly dependent on student choice, allowing for a deeper technical "spike" compared to the broad foundational skills of Years 1 and 2. It should also be noted that during this analysis we did not account for cases where the final year research project, the highest credit-contributing module in Year 4, was purely focused on one or more of the DCE domains rather than on research areas not as heavily intertwined with digitalisation.

Having completed the DCE mapping of the curricula of participating Universities, the eight domains were used to map the degree accreditation guidelines from the IChemE and AHEP4, marking absence of mention as a 0, vague mention as 1 and clear coverage as 2 (Figure 5).

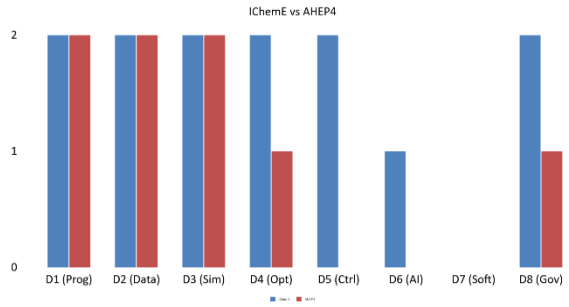

Figure 5. Mapping of IChemE and AHEP4 accreditation guidelines.

The IChemE provides information that effectively discusses the need for educational provision when it comes to domains D1-5. For AI/ML and Digital Twins (D6) the guidelines of “AI approaches” that can very vaguely

link to this particular domain.

When it comes to Software engineering practices (D7), it is unclear whether the accreditation guidelines include this domain. The guidelines mention the need to "understand system dynamics, being able to determine the dynamic response to changes in a process, design measurement and control functions, and determine its performance" which can vaguely link to D7 but not clearly. Last but not least, while there is clear mention of the need for prospective graduates to become aware of ethical considerations, there is no explicit mention or relevance to ethics around the use of digital tools. When it comes to big data or or cybersecurity there is acknowledgement of how the profession is changing.

A notable difference between the IChemE and AHEP4 accreditation guidelines is that whilst D1-4 are clearly covered under both schemes, D5-D7 are not mentioned in the AHEP4 guidelines. This could be because at least D5 and D6 might be deemed discipline-specific while AHEP4 caters for all engineering disciplines. When it comes to data governance, cybersecurity and ethics, AHEP4 mentions the importance of ethics, but not necessarily relating that to the digital space.

REFLECTIONS

The DCE curriculum mapping of the Universities participating in this study shows clear trends in the provision of education on digital tools. Undeniably, all programmes considered have embedded the majority of the eight examined domains in their curriculum, at different years of study and in different ways, following accreditation guidelines. What is interesting to note is that there is very limited coverage of Software Engineering Practices (D7) across the studied programmes, which also aligns with the absence of this domain from the accreditation guidelines. This could potentially be attributed to limited relevant educational capabilities across module leaders, or even to the fact that this domain, at least within the remit of chemical engineering, has not been a well-established area to allow for development of relevant educational content. The lack of D7 presence in the chemical engineering curriculum could be pointing towards the need for faster adoption of industrial needs within taught curricula and provision of training to educators when it comes to digital tools, with the aim of ensuring the more holistic preparation of future workforce.

It is worth pointing out that while mapping the curricula it became evident that ensuring accurate DCE mapping in terms of domain presence and intensity was cumbersome. This could be attributed to the “hidden curriculum” effect where it was tricky to differentiate between taught content that was formally assessed and content that was expected to be understood and applied but without formal assessment for a particular module. Furthermore, in complex module cases, such as the Design Project and Research Project that students usually undertake in Year 3 and Year 4 respectively, this effect was even more evident, as there are cases where such modules could contain very strong digital-focused components at an individual level rather than a cohort one. Accreditation guidelines indicate the baseline programmes are expected to evidence around learning outcomes and specific attributes; however, there is no expectation of where in the curriculum such elements should be embedded, or the intensity of their coverage. For a discipline with such profound industrial relevance, it is imperative for educators to work closely with accrediting bodies and industry, and continually track industrial needs, ensuring that the attributes of graduate chemical engineers are on par with the professional expectations. Furthermore, it is relevant to note that updated IChemE accreditation guidelines, which align more explicitly with AHEP4, will be released in the coming year.

CONCLUSIONS AND FUTURE WORK

This work has revealed that there is substantial variation in the breadth and depth of digital skills coverage for domains 1 to 6 and domain 8 across institutions and programmes evaluated. However, all programmes evaluated showed limited coverage of domain 7 on software engineering practices, followed by domain 6 on ML/AI and digital twins. This evidence shows the misalignment of the provision of these emergent skills with the evolving industrial sector.

Traditional digital domains aligned with core mathematical knowledge such as programming and computation and data literacy and statistics (D1 and D2) are commonly introduced in the first year of the programmes. Process modelling and simulation and optimisation, process systems engineering, control, automation and instrumentation and AI/ML and digital twins (D3-6) however are preferentially introduced after students are familiarised with process unit operations and fundamental programming skills. Governance, ethics and cybersecurity are covered evenly across the first three years, but more enquiry on the scaffolding in depth and specialisation in this domain is required to understand the student readiness for the sector. The provision of digital skills shows a front-loading in year 1 and 2 mainly related to inputting domains D1 to D5, followed by a drop in year 3 where application takes the main stage. Finally, in year 4 the incorporation of advanced chemical engineering skills and specialisation introduces a new instruction phase in the learning of digital skills.

Domains early introduced in the programme and traditionally requested by the IChemE accreditation (D1-D4) show a good balance between being taught and being

embedded, integrated and largely applied in later stages. On the contrary, domains D5-8, which are introduced at later stages, are mainly being taught with lesser opportunities for application. Overall, the digital chemical engineering provision broadly aligns with IChemE and AHEP4 guidance, where expectations for D1 to D4 are more explicitly articulated.

The present work is a cornerstone in the mapping of digital skills that has evidenced the need to enquire about potential barriers affecting underrepresented domains such as domain 7. It also underscores the importance of examining further the breadth and depth for all domains through a more comparative quantitative or semi-quantitative analysis across the years and institutions. The intention of this work is not to provide recommendations but rather provide an overview of digital skills presence in chemical engineering education. The insights generated here will support future evaluation of how programmes align with the forthcoming IChemE accreditation guidelines, which integrate expectations from the Engineering Council's AHEP4 framework, and may ultimately serve as a tool for assessing digital-skills provision within accredited chemical engineering programmes and sector needs.

ACKNOWLEDGEMENTS

The authors would like to acknowledge colleagues from each academic institution, for valuable insight and discussions that contributed to our understanding and mapping of DCE in our curricula. In particular, DD would like to thank Dr Marco Cardoso Marques, ER and MZ would like to thank Ben Pearson, MGM would like to thank Dr Zoran Visak and MM would like to thank John Turner.

AUTHOR IDENTIFIERS

Author ORCIDs:
Routoula E: 0000-0001-9416-4265
Mohammad Zadeh M: 0000-0001-5949-5309
Granollers Mesa M: 0000-0001-8602-9362
Malekshahian M: 0009-0001-0560-0561
Dikicioglu D: 0000-0002-3018-4790
Pollock M: 0000-0001-6640-7661
Zandi M: 0000-0003-0017-1080

REFERENCES

1. Udugama IA, Bayer C, Baroutian S, Gernaey KV, Yu W, Young BR. Digitalisation in chemical engineering: industrial needs, academic best practice, and curriculum limitations. Education for Chemical Engineers 39:94-107 (2022). https://doi.org/10.1016/j.ece.2022.03.003
2. Zandi M, Glassey J, Young BR. Mind the Gap. *Chem. Eng.*. 976:41-43 (2022) https://www.thechemicalengineer.com/features/mind-the-gap/
3. Bilques Z, Belton DJ, Campbell GM. Structure and content of beng chemical engineering programmes in the UK, relative to the frontiers in chemical engineering education model. Process Integr Optim Sustain 7:1003-1030 (2023). https://doi.org/10.1007/s41660-022-00307-6
4. Udugama IA, Atkins M, Bayer C, Carson J, Dikicioglu D, Gernaey KV, Glassey J, Taylor M, Young BR. Digital tools in chemical engineering education: the needs and the desires. Education for Chemical Engineers 44:63-70 (2023). https://doi.org/10.1016/j.ece.2023.05.002
5. Karoglou M, Ghergulescu I, Stramarkou M, Boukouvalas C, Krokida M. Building an adaptive ai-powered higher education class for the future of engineering: a case study from NTUA. Applied Sciences 15:8524 (2025). https://doi.org/10.3390/app15158524
6. Rebello CM, Deiró GF, Knuutila HK, Moreira LCS, Nogueira IBR. Augmented reality for chemical engineering education. Education for Chemical Engineers 47:30-44 (2024). https://doi.org/10.1016/j.ece.2024.04.001
7. Zárate-Navarro MA, Schiavone-Valdez SD, Cuevas JE, Warren-Vega WM, Campos-Rodríguez A, Romero-Cano LA. STEM activities for heat transfer learning: integrating simulation, mathematical modeling, and experimental validation in transport phenomena education. Education for Chemical Engineers 49:81-90 (2024). https://doi.org/10.1016/j.ece.2024.06.004
8. IChemE, Accreditation: degree programmes. https://www.icheme.org/employers-educators/services-for-education-providers/accreditation-degree-programmes/guidance/
9. Engineering Council, Accreditation of Higher Education Programmes. https://www.engc.org.uk/our-role-as-regulator/setting-standards/recognising-learning-and-development-programmes

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Enhancing Robotics and Automation Education Through the Development of Simulation Tool for Material Synthesis

Hsuan Chang[a], Adedayo Ogunnoiki[a], and Solomon Gajere Bawa[a*]

[a] University College London, Department of Chemical Engineering, London, United Kingdom
* Corresponding Author: solomon.bawa@ucl.ac.uk

ABSTRACT

As high-throughput experimentation (HTE) becomes a cornerstone of modern materials research, undergraduate and postgraduate curricula increasingly require students to possess Python programming skills to operate automated liquid-handling robots, such as the Opentrons OT-2 and Flex. However, the high cost of this hardware often necessitates shared equipment use during hands-on lab sessions, creating a significant pedagogical barrier: students lack the individual time required to iteratively test and debug their protocols on physical robotic platforms for automated material synthesis. Furthermore, we observe that the scarcity of robotic platforms creates an imbalance in group dynamics, where students with more coding experience often lead protocol development, while those with less experience remain disengaged. To address these challenges, we developed an interactive simulator that translates Python protocols into 2D animations on personal laptops. Using gold nanoparticle (AuNP) synthesis as a case study, we demonstrated how this platform enabled visualisation, captured common programming mistakes, and refined the efficiency of the code before implementation using the physical robot. By shifting the 'trial-and-error' phase from the lab to the laptop, the simulator showed potential to democratise students' access to automation education and ensure equitable skill application across the cohort.

Keywords: automation, simulation, material synthesis, visualisation, robotic, nanoparticles

1.0 INTRODUCTION

Industry 5.0 emphasises collaboration between humans and machines to develop innovative and sustainable solutions [1]. In chemical engineering, this paradigm is reflected in the use of robotic platforms to automate high-throughput chemical synthesis, curate large datasets for process optimisation, and support sustainability initiatives [2, 3]. The integration of these technologies into undergraduate and postgraduate curricula is becoming increasingly common, with students programming liquid-handling robots using Python-based protocols to automate experiments for high-throughput material synthesis. Data-Driven Materials Manufacturing is one of the core modules taught at UCL's MSc Digital Manufacturing of Advanced Materials (DMAM) programme. This module aims to teach students how to program liquid-handling robots, such as the Opentrons OT-2 and Flex, using Python.

However, several challenges were identified while teaching with the liquid handling robots. The first challenge arises from the limited infrastructure of the lab, where the necessity of sharing a single robot among five or six students results in a fragmented learning experience. Coding for automation is a skill that demands constant experimentation and immediate feedback; however, the current setup forces long wait times between code execution and physical observation. This lack of consistent access prevents students from developing coding skills through iterative testing. Furthermore, students with strong programming backgrounds often dominate the code development stages, thereby marginalising peers with limited or no prior coding experience. This imbalance raises a significant pedagogical concern and calls for more inclusive teaching strategies to ensure equitable participation and learning outcomes for all students [4]. Additionally, because students have limited opportunities to build basic computational skills, it is difficult to incorporate sustainability and optimisation principles into laboratory education. As a result, we observed

excessive use of pipette tips and other consumables due to inefficient protocol design by students. These challenges highlight the need for accessible digital solutions to support coding for liquid-handling robots and to create inclusive learning environments.

In this study, a 2D desktop-based simulation platform was proposed to overcome the pedagogical barrier of limited access to robots. By allowing students to run code and receive individual assessments virtually, the simulator reduces reliance on physical hardware and fosters independent programming. This extra time enables students to move beyond simply writing functional code, allowing them to design efficient liquid transfer steps that minimise pipette tips waste and reinforce principles of sustainability.

The aim of this paper is to share the development and perceived use of the simulator in the classroom environment. The simulator can be used to simulate both the Opentrons OT-2 and Flex, but this paper focuses on its application with the Opentrons Flex. Here, we explain how the simulator was used to visualise a Python protocol for gold nanoparticle (AuNP) synthesis, a material of growing importance in renewable energy and healthcare applications [5].

2.0 DESIGN OF GOLD NANOPARTICLE SYNTHESIS PYTHON PROTOCOL

The Python protocol for synthesising gold nanoparticles was designed according to the Turkevich method [6]. The experimental steps of this method involved stirring and heating gold (III) chloride trihydrate ($HAuCl_4$) and trisodium citrate (Na_3Ct) solutions within a 96-well plate (reaction plate). The experimental design explored three significant factors of the Turkevich method that influence AuNPs size: initial $HAuCl_4$ concentrations, the molar ratio of $HAuCl_4$ to Na_3Ct, and temperature. The full factorial design of the experiment is shown in Table 1, with levels chosen based on the literature [6].

Table 1: Full factorial design of experiment for gold nanoparticles synthesis.

Factor	Levels
C (mM)	0.20, 0.25, 0.50, 0.75
n_{ratio}	1.00, 1.25, 1.50, 2.00, 2.40, 2.80, 3.20
T (°C)	60, 70, 80, 100

C is initial concentration, n_{ratio} is molar ratio, and T is temperature.

To perform experiments with the Opentrons Flex liquid handling robot, the Python protocol specified a sequence of liquid transfer steps. First, $HAuCl_4$ and diluent were transferred from the solvent reservoirs to the reaction plate using an eight-channel pipette. After filling the reaction plate with diluted $HAuCl_4$, the Python protocol instructed the robot to move the plate onto the Heater-Shaker Module with a gripper. The protocol specified the temperature and shaking speed. After heating, the reaction plate was moved back to the initial position. Then, the heated $HAuCl_4$ mixture was supplemented with different volumes of Na_3Ct and heated again with the Heater-Shaker Module. Note that each run of the protocol tested only one temperature, so to assess all four temperatures specified in Table 1, three additional runs of the protocol with different temperature values were conducted. Finally, the product was transferred from the reaction plate to a 360 uL Corning 96-well plate.

By transferring precise volumes of various reagents, the reaction plate was filled with reagents in the configuration shown in Figure 1. In this setup, each well in columns 1 through 4 represents a unique combination of $HAuCl_4$ concentration and the $HAuCl_4$ to Na_3Ct molar ratio; these combinations were repeated in triplicate across columns 5-8 and 9-12. Wells in row H were left empty, except for wells H1 through H4, which were filled with water to serve as blanks for UV-Vis analysis.

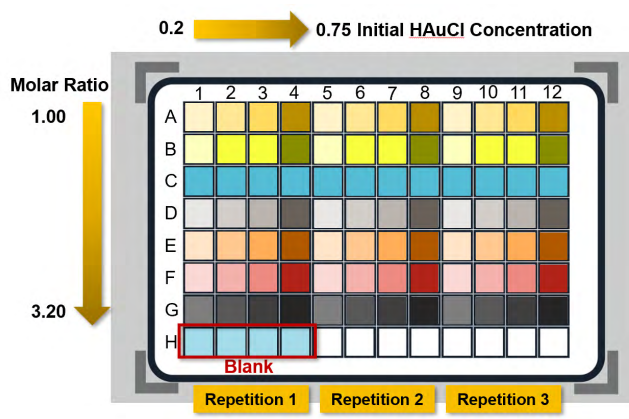

Figure 1. Layout of the conditions in AuNPs synthesis reaction plate.

3.0 SIMULATING THE PYTHON PROTOCOL

3.1 Basic Simulator Functions

Before rendering animation, the simulator first converts the Python protocol into a text file using the Opentrons-simulate package. The Opentrons-simulate package allows users to preview the steps taken by the robot in text form, as demonstrated below:

```
Latching labware on Heater-Shaker
Picking up tip from G1 of Opentrons Flex 96
Filter Tip Rack 1000 µL on slot B1
Aspirating 40.0 uL from A1 of NEST 12 Well
Reservoir 15 mL on slot D3 at 716.0 uL/sec
```

```
Dispensing 40.0 uL into G1 of NEST 96 Deep
Well Plate 2mL on slot D2 at 716.0 uL/sec
```

The simulator then works with this text to generate the animation. Essentially, the simulator employs web technologies for the graphical interface. JavaScript was used to drive the real-time animation of the robotic process, handling the dynamic updates to the deck layout. Cascading Style Sheets were used to style both the static and dynamic elements of the simulator's user interface, ensuring that the visual representation of the Opentrons Flex deck, labware (plates, tips), and pipette modules was accurately sized, positioned, and clearly visible, contributing to a realistic and easily understood user experience.

3.2 Visualising Liquid Aspiration and Dispense

The simulator simulates liquid being aspirated from or dispensed into a well by changing the height of the rectangle that represents the liquid level. At startup, the simulator initialises each well's current volume as either empty (0 uL) or as the maximum liquid volume a well can hold. When rendering the animation, every aspirate or dispense command triggers an update routine that changes the current volume of the affected well by subtracting or adding the transferred volume, respectively. The visual representation of liquid filled is derived from a level fraction (L_f) defined as:

$$L_f = \frac{V_{current}}{V_{capacity}} \tag{1}$$

where $V_{current}$ is the current well volume and $V_{capacity}$ is the total liquid volume each labware well could hold. For rectangular reservoirs, this fraction determines the height h_{fill} of a bottom-anchored filled rectangle, computed as:

$$h_{fill} = L_f \bullet H_{well} \tag{2}$$

$$Y_{top} = Y_{well} + (1 - L_f) \bullet H_{well} \tag{3}$$

with H_{well} being the constant pixel height of the well. The top Y-coordinate of the filled region is therefore given by ensuring proportional visual depletion or filling as volume changes. A sequence of frames where the coloured rectangle height decreases as $HAuCl_4$ solution was aspirated from the reservoir is shown in Figure 2.

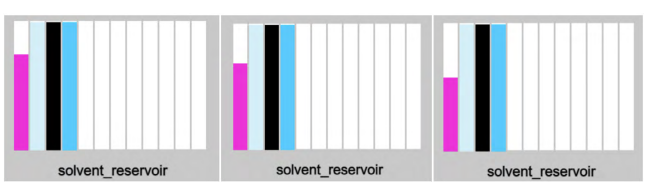

Figure 2. Animation of reduction in liquid height as 120 uL was being aspirated from solvent reservoir well A1.

3.3 Visualising Dilution and Liquid Mix

Colour-coding is used to indicate the presence of liquid within a well. For example, the well plates were initialised as empty and therefore appeared uncoloured. Conversely, the reservoirs were initialised as full and displayed in blue, reflecting the standard practice of pre-filling reservoirs with solvents or solutions at the start of a protocol. The simulator also initialised all liquid colours as blue, as shown in Figure 3(a).

Although the simulator has a default setting for labware, it also allowed this initial state of the labware to be customised. As shown on the right-hand side of Figure 3(b), for the AuNP protocol, the $HAuCl_4$ solution was set as bright pink, the diluent was set as dark blue, Na_3Ct was set as black, and water remained blue. Other wells were selected as "empty", so the wells were uncoloured. While the diluent in solvent reservoir A2 was set as dark blue in Figure 3(b), there was a difference in colour between the colour set in the right panel and the colour being shown in the animation on the left. This was because solvent reservoir A2 was set to be a "diluent", so the simulator replaced the diluent with light blue instead.

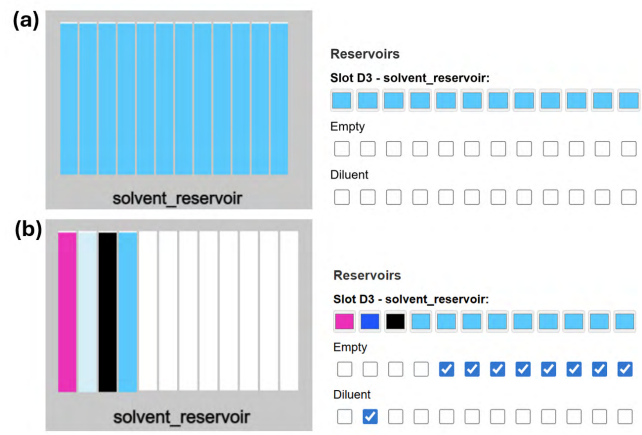

Figure 3. (a) Default setting of reservoir and (b) user setting of reservoir.

When $HAuCl_4$ was transferred to the reaction plate, the wells in the reaction plate were coloured bright pink to indicate the presence of $HAuCl_4$. Upon the addition of a diluent, the simulator calculated the ratio of non-diluent to current liquid volume in the well using Equation 4 and updated the opacity of the well's colour accordingly:

$$\kappa_{final} \propto \frac{V_{non-diluent}}{V_{total}} \tag{4}$$

where V_{total} is the total liquid volume in the well, $V_{non-diluent}$ is the volume of non-diluent liquid in the well, and K_{final} is the final opacity.

As a result, when diluent was added to $HAuCl_4$, the colour changed across the column and repeated every four columns. When Na_3Ct was added to the wells, the volume of Na_3Ct added to each row was the same, resulting in uniform colour across each row. However, the

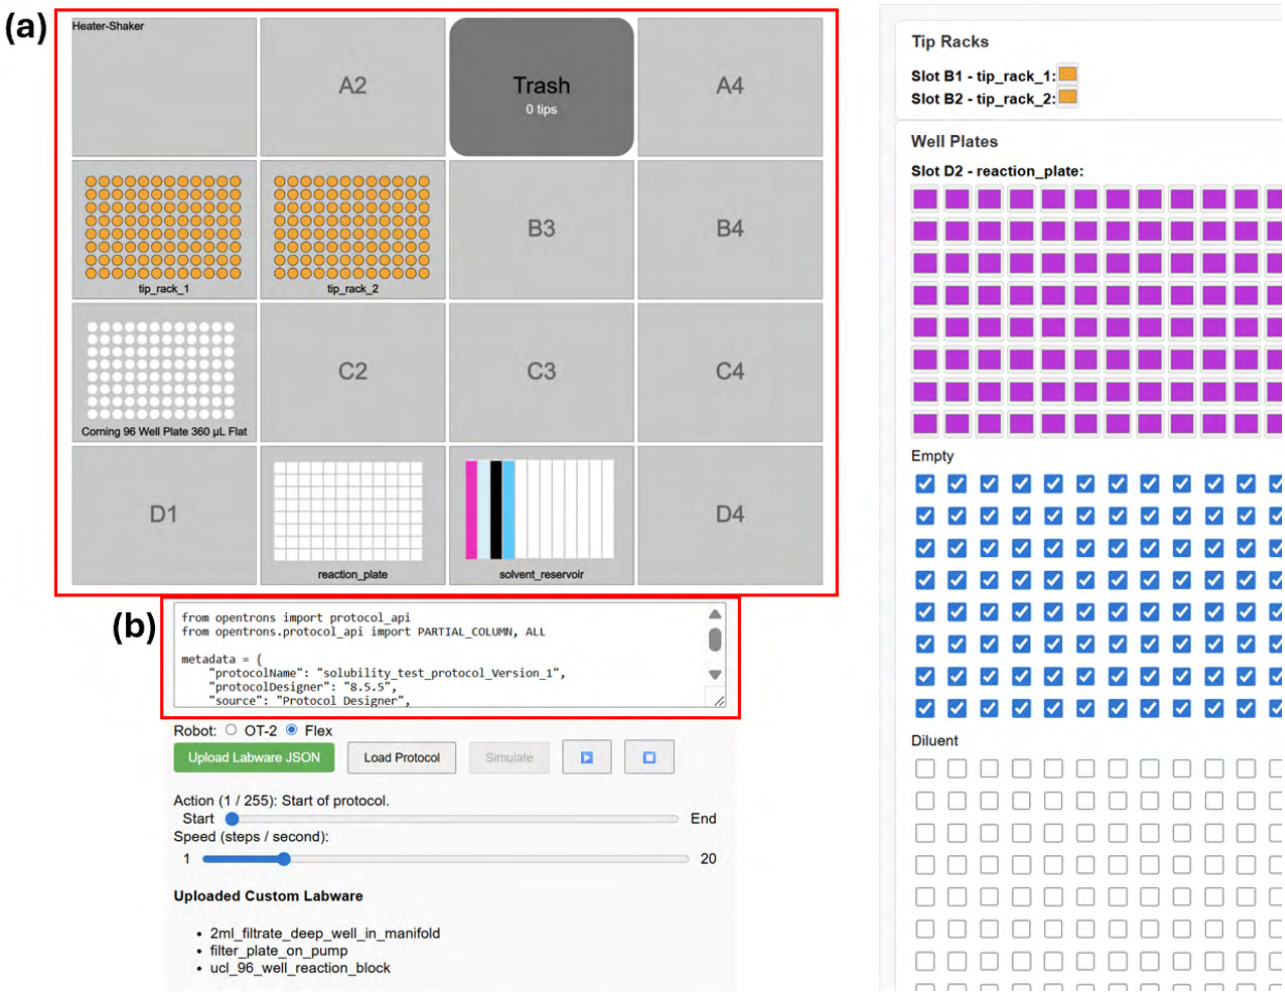

Figure 4. The simulator's user interface consists of several key sections: (a) the animation area and (b) the code editor where the Python protocol for AuNP synthesis is uploaded. Below the code editor, users can select the robot model (OT-2 or Flex) and use a dedicated button to upload JSON files for custom labware. Clicking 'Load Protocol' triggers the system to either display an error message or signal that the 'simulation is ready'. In the latter case, the 'Start Simulation' button becomes active. Adjacent to this are the pause and restart controls for the animation. The execution timeline is displayed under the 'Action' panel, which includes a slider to adjust the animation speed. Finally, the right-hand panel specifies the initial state of the labware based on three parameters: colour, fill status (empty), and diluent type.

volume varied between rows, producing differences in colour down the plate. This difference was captured by the simulator through colour blending. When dispensing liquid that is not labelled as a diluent, the simulator took the red, green, and blue channel intensity (RGB) of the Na_3Ct and the RGB of the previously diluted $HAuCl_4$ and averaged them based on their respective volume using Equation 5:

$$R_{final} = \frac{R_{source} V_{source} + R_{dest} V_{dest}}{V_{total}} \quad (5)$$

where R_{final} is the resulting RGB of the new mixture, R_{source} is the RGB of the liquid being poured in, V_{source} is the volume of the liquid being poured in, R_{dest} is the RGB of the liquid already present in the destination container, and V_{dest} is the volume of the liquid already present in the destination container.

Figure 4 displays the entire simulator's user interface. When a user uploads a Python protocol for AuNP synthesis into the code editor, the simulation deck (Figure 4(a)) will update to show each piece of labware at the position specified in the Python protocol.

3.4 Simulating Gripper

The function of the gripper was to move the reaction plate from its starting slot D2 onto the Heater-Shaker

Module when heating was needed. The simulator simulated this movement by detecting command lines that instructed the robot to move labware, such as:

```
Moving nest_96_wellplate_2ml_deep to Heater-
ShakerContext at Heater-Shaker Module GEN1
on A1 lw None
```

Once a movement command was detected, the system converted this movement of labware into a step of reassigning the labware from its current slot to the new slot. When this step was parsed to the function that ran the animation. Because the function rendered labware based on their assigned slot positions in the dictionary, the labware appeared to move instantly to the new location on the screen as the slot position in the dictionary changed as illustrated in Figure 5. This allowed visualisation of the reaction plate's movement to any position on the deck, including on top of the Heater-Shaker Module.

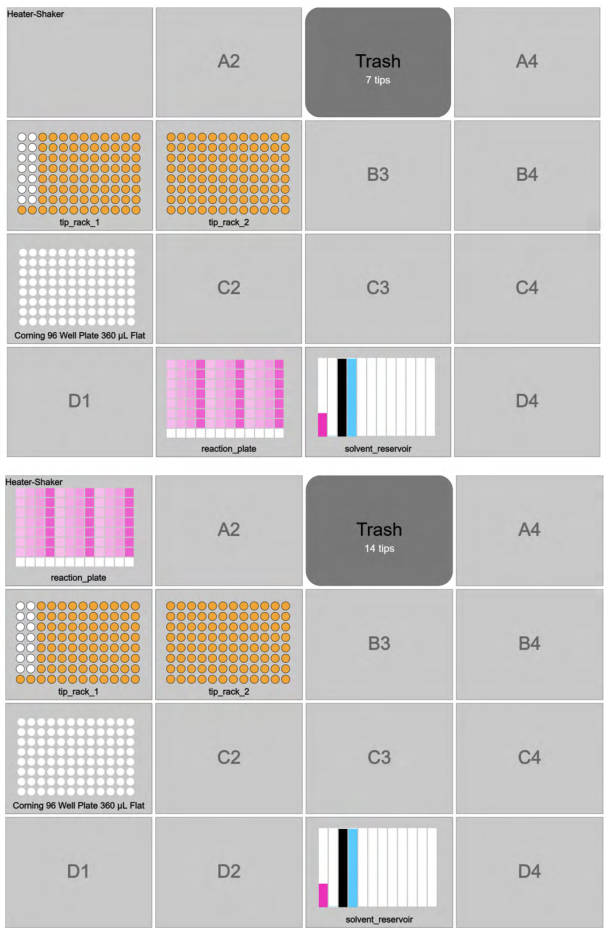

Figure 5. Gripper moving reaction plate from D2 (top) to the Heater-Shaker Module in A1 (bottom).

3.5 Simulating Heater-Shaker Module

In the Python protocol for AuNPs, the Opentrons Heater-Shaker Module was used for heating and mixing the diluted $HAuCl_4$ and $HAuCl_4/Na_3Ct$ mixture. The Heater-Shaker has two parameters: temperature and shaking speed. In the Python protocol, the parameters were set using the following code:

```
hs_mod.set_and_wait_for_shake_speed(300)
hs_mod.set_and_wait_for_temperature(60)
protocol.delay(minutes=10)
```

The simulator recognised this instruction from the Python protocol to extract the Heater-Shaker temperature and speed settings. It then broke them into a sequence of smaller, increasing values so that the animation showed the temperature and shaking speed rising to the target values (Figure 6). The animation started at a low baseline (23 °C for temperature or 0 RPM for the shaker), pushed intermediate values that stepped up slowly at first and faster later (small increments under a threshold, larger ones above), and then finally reached the exact target value.

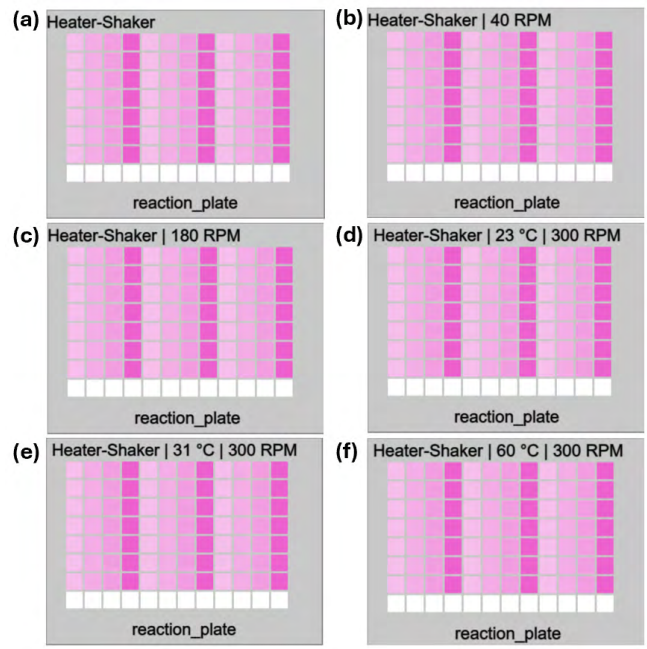

Figure 6. Chronological demonstration of heating and shaking process for the heater-shaker module during animation from (a) to (f).

After both temperature and shaking speed reached their target values, according to the Python protocol, the protocol would pause for 10 minutes. The simulation does not pause for 10 minutes. Instead, a pop-up window informs the user that the protocol was delaying for 10 minutes and 0.0 seconds. Upon clicking the "Ok" button, the simulator skips the waiting period and proceeds with execution. Overall, these features enable the simulation of the Heater-Shaker module.

4.0 CAPTURING COMMON MISTAKES WITH SIMULATOR

The simulator captured several mistakes made in the Python protocol for AuNP synthesis without running it on a physical robot. The first mistake was transferring liquid to the wrong well. A typo in the Python protocol, such as “A1” instead of “A2”, was hard to capture by human eyes reading through each line of code, especially in nanoparticle material synthesis where the code and simulated text could be complex and long. Although this mistake could be captured with a dry run, a single dry run took around 15 minutes, and one would have to perform several dry runs if there were multiple typos. The simulator speeds up this trial-and-error process by capturing unintended liquid transfers on the screen. In the classroom, students would be able to capture multiple errors in one simulation and change their code accordingly. This is much faster than testing the code on the robot.

The second mistake was transferring more liquid from the solvent reservoir to the reaction plate than the maximum amount of liquid a reservoir well could hold. Keeping track of how much liquid was aspirated from the reservoir by reading through lines of code was challenging; thus, one common mistake was designing a Python protocol that tried to aspirate more liquid than the maximum volume. During AuNPs synthesis, drying up the reservoir well could cause inaccurate transfer of liquid into the reaction plate and inevitably mess up the molar ratio or concentration of reactants. Note that this mistake could not be captured with a physical dry run or the Opentrons-simulate package, as both do not keep track of how much liquid was aspirated from a reservoir well.

The simulator kept track of the volume in each well and showed changes in liquid level by calculating how much liquid was left in the well. Building on this basic function, we added a feature to stop the simulation when the liquid level falls below 0 uL after aspiration. This would allow students to notice that more diluent than the maximum well volume was drawn by the Python protocol and fix their code accordingly before deploying it on the robot.

5.0 USING THE SIMULATOR IN CLASSROOM AND WIDER EDUCATIONAL IMPACT

The simulator enables a structured, multi-stage pedagogical workflow that integrates digital laboratory automation into teaching and training activities. In the initial phase, students receive instruction on the manual synthesis of AuNPs, including discussion of key experimental variables such as reaction temperature, initial $HAuCl_4$ concentration, and the molar ratio between $HAuCl_4$ and Na_3Ct. Building on this foundational knowledge, each student then independently develops a Python protocol on their personal laptop and evaluates its performance using the simulator. This allows students to iteratively test, debug, and refine their protocols in a low-risk virtual environment prior to working with the physical robotic system.

The iterative cycle—develop, simulate, evaluate, and improve—not only strengthens computational thinking skills but also creates an accessible and scalable approach to teaching laboratory automation. The simulator additionally provides quantitative feedback not available in existing tools (e.g., tracking pipette tip usage), enabling assessment based on efficiency, sustainability, and resource optimisation.

Once students complete the simulation stage, they implement their protocols on the Opentrons Flex robot. Small-group review and peer critique of protocols foster collaborative problem-solving and mirror engineering practices. This supports teamwork while preserving individual accountability. In settings without access to robotic hardware, the simulator supports fully remote learning, with students submitting both their protocols and screen-recorded simulations for assessment.

5.1 Extended Educational Impact Across Programmes

The impact of the simulation tool extends beyond the immediate classroom environment. Postgraduate researchers within the School of Pharmacy and chemistry have adopted the simulator to upskill new starters whose research involves the use of liquid-handling robots with similar operational principles. By engaging with the simulator prior to accessing physical instruments, these students report increased confidence, reduced anxiety, and improved preparedness for hands-on experimentation. This has proven particularly valuable for students transitioning into laboratory automation with limited prior programming experience. The simulator has been used to validate developed code deployed for applications in combinatorial chemistry, solvent screening, and the assessment of solubility for various chemicals in selected solvents.

Furthermore, the simulator has now been fully integrated into the Digital Manufacturing of Advanced Materials (DMAM) MSc programme, where it supports training in automation-ready digital workflows. Its inclusion strengthens the programme’s emphasis on computational methods, digital manufacturing, and reproducible laboratory practices. Early feedback indicates that students benefit substantially from the ability to repeatedly practise and refine automated protocols outside of scheduled laboratory hours.

There is also ongoing discussion with the School of Pharmacy to incorporate the simulator more formally into postgraduate-taught training. This integration would extend its impact to a broader cohort of students, embedding digital laboratory automation into curricula where traditionally wet-lab skills dominate. The simulator’s

alignment with digital skills agendas and its accessibility on standard laptops make it an attractive teaching tool for programmes seeking to modernise laboratory education.

5.2 Industrial Upskilling and Professional Training

Beyond academic environments, the simulator has also been identified as a valuable tool for upskilling industrial workers. In a recent funding bid submitted by the Institute of Robotics, the simulator was highlighted as a resource capable of supporting workforce development in automated laboratory processes. Its ability to train users remotely, safely, and without access to expensive hardware makes it suitable for industrial training pipelines, particularly in sectors where laboratory automation is expanding rapidly.

The tool thus supports a continuum of training—from taught postgraduate students to incoming researchers, to professionals requiring automation-related reskilling. Its scalability positions it as an enabling technology for modern laboratory education and industrial digital transformation.

5.3 Advancing Research Reproducibility and Digital-First Laboratory Practices

The simulator contributes to ongoing efforts to improve research reproducibility. Challenges in experimental reproducibility increasingly arise from undocumented implicit knowledge, variability in manual procedures, and inconsistent reporting of protocol parameters. By requiring users to encode all steps in a Python-based, machine-readable protocol, the simulator encourages a level of procedural precision difficult to achieve through traditional textual descriptions alone. The ability to directly execute these digital protocols on automated liquid-handling platforms further ensures that methods can be reproduced with high fidelity across laboratories, institutions, and time.

The simulator also supports the adoption of digital-first laboratory practices that are central to emerging frameworks and experimental workflows. In environments where liquid-handling robotics are increasingly integrated with electronic lab notebooks and laboratory information management systems, the simulator can act as a bridge between conceptual training and real-world digital infrastructure.

Moreover, the quantitative performance data generated by the simulator—such as consumable usage, timing estimates, and workflow efficiencies—enable users to evaluate protocols using metrics that are rarely captured in traditional laboratory teaching. These data create opportunities for optimisation-driven teaching and research, where experimental design can be informed by resource efficiency, sustainability considerations, and throughput modelling. Incorporating such metrics into training better prepares learners to operate within modern research laboratories, where cost-aware and sustainability-oriented decision-making is increasingly prioritised.

6.0 CONCLUSION

This study presented the development of a 2D desktop-based simulation tool designed to enhance robotics and automation education in chemical engineering. By translating Python-based Opentrons protocols into real-time animations, the simulator effectively addresses the pedagogical barriers of limited hardware access and uneven group dynamics. Using gold nanoparticles (AuNPs) synthesis as a case study, we demonstrated the platform's ability to visualise complex liquid-handling steps, track consumable usage, and capture critical programming errors (such as volume deficits) that dry runs overlook.

This work serves as a blueprint for other UCL modules seeking to bridge the gap between theoretical coding and physical implementation. Programmes in Biochemical Engineering, material science or computer science can look to this model to provide equitable access to high-cost hardware. By decoupling code development from hardware availability, modules can ensure that students with less prior programming experience are not marginalised by "dominant" peers in a group setting. It demonstrates that digital competency is best fostered when students have the autonomy to fail and iterate in a low-stakes, individual digital environment before moving to the physical laboratory.

For UCL programmes aiming to embed the UN Sustainable Development Goals (SDGs) into their curricula, this manuscript illustrates how "Digital Twins" can directly reduce the carbon footprint of laboratory education in terms of minimising waste. Other modules involving high resource consumption, such as chemistry or synthetic biology, can adopt similar simulation-based assessments in which resource efficiency is used as a graded learning objective. By providing quantitative feedback on consumable usage (e.g., plastic tips and chemical reagents) during the simulation phase, departments can train a new generation of engineers who view optimisation and waste minimisation as core components of experimental design rather than afterthoughts.

Future research will focus on integrating more visualisation features to provide a more immersive experience and expanding the library of supported labware and modules to cover a broader range of high-throughput experimentation (HTE) workflows. Additionally, we aim to implement an automated feedback system that provides students with real-time suggestions for optimising reagent efficiency and protocol duration. In the context of

HTE, this simulator serves as a scalable digital twin, enabling researchers and students alike to validate complex, data-driven experimental designs before physical deployment, thereby reducing waste and accelerating material discovery.

ACKNOWLEDGEMENTS

The authors would like to thank University College London (UCL) and the Department of Chemical Engineering for providing the laboratory facilities, the robotic hardware (Opentrons Flex), and computational resources necessary for the development and testing of this simulation platform. We also acknowledge the MSc Digital Manufacturing of Advanced Materials (DMAM) cohort for their participation and feedback during the pilot implementation of this tool.

REFERENCES

1. Akundi A, Euresti D, Luna S, Ankobiah W, Lopes A, Edinbarough I. State of industry 5.0—analysis and identification of current research trends. ASI 5:27 (2022). https://doi.org/10.3390/asi5010027
2. Lu JM, Pan JZ, Mo YM, Fang Q. Automated intelligent platforms for high?throughput chemical synthesis. Artificial Intelligence Chemistry 2:100057 (2024). https://doi.org/10.1016/j.aichem.2024.100057
3. Vu H, Bui KH, Dang KDD, Duong-Tuan M, Le DD, Nguyen-Dang T. Finding environmental-friendly chemical synthesis with AI and high-throughput robotics. Journal of Science: Advanced Materials and Devices 10:100818 (2025). https://doi.org/10.1016/j.jsamd.2024.100818
4. Florian L, Linklater H. Preparing teachers for inclusive education: using inclusive pedagogy to enhance teaching and learning for all. Cambridge Journal of Education 40:369-386 (2010). https://doi.org/10.1080/0305764x.2010.526588
5. European Commission. (2024). Advanced Materials for Industrial Leadership. https://research-and-innovation.ec.europa.eu/research-area/industrial-research-and-innovation/chemicals-and-advanced-materials/advanced-materials-industrial-leadership_en
6. Dong J, Carpinone PL, Pyrgiotakis G, Demokritou P, Moudgil BM. Synthesis of precision gold nanoparticles using turkevich method. KONA 37:224-232 (2020). https://doi.org/10.14356/kona.2020011

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator, and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A Techno-economic Analysis of Simulated Wind Farms

Isaac N. James[a], Laura Edwards*[a], and Dhurjati Chakrabarti[a]

[a] The University of the West Indies, Department of Chemical Engineering, St. Augustine, Trinidad and Tobago
* Corresponding Author: laura.edwards@uwi.edu

ABSTRACT

The implementation of processes that use renewable energy requires that a techno-economic analysis be performed beforehand to determine its economic and technical feasibility. A techno-economic analysis was performed proposed wind farms in Trinidad and Tobago using the System Advisor Model simulation software. Metrics included the annual energy production in kWh, capacity factor, net present value in US$ and internal rate of return. From the above, the number of households that can be powered each month by the farms were calculated. The results showed that rotor diameter, which defines the swept area has a significant impact on annual energy production as a 33 m difference translated into a 27.3 GWh and 22.9 GWh difference in output. The results are promising and show that the oil and natural gas-based economy can be diversified.

Keywords: Environment, Wind, Energy, Electricity & Electrical Devices, Renewable and Sustainable Energy

INTRODUCTION

As climate change and its effects become more profound, the term green technology has been used frequently in everyday discussions. Government bodies, corporations and the public have begun to become more aware of climate change, the negative effects of pollution and depletion of natural resources. It has also gathered significant interest from investors who are looking to invest mainly in companies that produce or utilize green technology. Green technology or more commonly called “green tech, ” is a broad term used to describe technology that one would consider environmentally friendly from its production process to its usage [1].

The move to reduce emissions has continued to date and global agreements such as the Paris Agreement and frameworks have been made between countries to steer progress for dealing with climate change in the right direction. The production of fossil fuels must drop by approximately 45% between 2020 and 2030, for countries who have pledged to attain net zero emissions by 2050 to realistically achieve that goal. Net zero refers to decreasing global emissions as close to zero as humanly possible while allowing the remaining emissions to be absorbed by natural sources such as forests and oceans [2].

A technical-economic analysis (also called techno-economic analysis or TEA) is a method used to analyze the technical and economic performance of an industrial process or product being developed. It is carried out to answer questions such as “is it technically feasible?” and “will it make money?”. It provides the business/organization estimates of the operating and capital costs as well as the expected revenue that stems from financial and technical parameters input into process specific software [9].

When companies are deciding if a project is a good investment or not, they calculate the net present value (NPV) and the internal rate of return (IRR) as two of the main indicators as to whether the project will generate net profits for the company. The NPV (in $) is obtained by discounting future cash flows using a variable discount rate and summing their present values. Cash flows are statements that report the cash spent and generated during a certain period. It gives a breakdown of the movement of money in and out the business/project. It is comprised of three sections which are operating, investing and financing activities [12].

A study has been done by Chadee and Clark (2017) in which the levelized cost of electricity (LCOE) generation from the wind for 2 sites with long term wind measurements were determined. Their results showed that the generation of electricity through tall, large wind turbines are cost competitive with the electricity generated through the burning of natural gas. The LCOE has units of US$/kWh and is frequently used to compare and

assess alternative methods of energy production, this metric determines if a project will be profitable or break even and enables financial analysts to compare various energy producing methods such as solar and wind. To calculate the LCOE for any energy project, the annual energy production (AEP) is required, this can be calculated by entering technical information specific to the energy generation process into software [3]. According to Chadee (2017), at 80 m above ground, the annual mean wind speeds were approximately 6.93 m/s at Crown Point and 5.86 m/s at Piarco.

To obtain the required technical inputs for a wind farm, a wind resource assessment will have to be carried out and the data analyzed. Through calculation of the LCOE and other financial metrics, decisions made by the relevant stakeholders.

BACKGROUND

The methods of electricity generation in Trinidad employed at the Penal power station makes use of a combined cycle power plant (PowerGen 2019a). Here, two 7E gas turbines fueled by natural gas each have their exhausts (comprised of a hot air-fuel mixture) captured by heat recovery steam generators (HRSGs). These HRSGs creates steam from the exhaust and this steam is sent to a steam turbine where its blades are spun, thereby turning its drive shaft that is connected to a generator in order to generate additional electricity (General Electric 2019). These gas turbines require atmospheric air and natural gas to operate (Trinidad Generation Limited 2017).

In May 2023, a workshop took place to highlight the findings of a European Union (EU) funded study that indicated that Trinidad's offshore areas held a large potential of wind resources that would decarbonize the industry as well as aide the transition towards an economy depending less on fossil fuels. Reference was drawn to the use of 8% of the country's produced natural gas for electricity generation.

While offshore wind resource assessments are complex due to the rough marine environment, cost and timeframe of data collection, collaboration and alignment between research institutions, the government, investors, local communities and industry is crucial [11].

Suitable locations in Trinidad and Tobago were analyzed and potential onshore and offshore wind farm sites identified using wind resource data from the Global Wind Atlas as well as available local data. These sites were then ranked based on their LCOE and suitability. Onshore wind power has a potential capacity of 2.75GW while offshore has a capacity of 32GW. However, a Wind Resource Assessment Programme (WRAP) will need to be undertaken to confirm the actual wind speeds at the various locations to enable appropriate siting and approximate energy yields (Sakellariou 2023, 16, 19-20).

The team stated that large scale wind projects require infrastructure to ease the development process and lower the project risks. Heavy duty cranes, equipment capable of transporting the heavy turbine parts over the land in addition to specially designed ocean vessels for lifting and installing the turbines offshore would be required. As Trinidad and Tobago has a well-established petrochemical sector: several sea ports exist that can accommodate large cargo vessels, expertise in inland transportation logistics have been accrued as well as a skilled workforce and strong demand for energy. The country will be able to capitalize on the above strengths in order to generate electricity with a lowered carbon footprint and thus use the remaining natural gas to create other products of higher value (Sakellariou 2023, 20, 27).

Model inputs include: the location of a project, operating, maintenance and installation costs, discount rate and other factors such as system specifications. Outputs include: annual electricity generation (kWh), the levelized cost of electricity, payback period, net present value, internal rate of return as well as a cash flow table with accompanying customizable graphs. SAM can model a single wind turbine as well as a wind farm, it also has a module and inverter database. The user has the option to use from a predefined set of turbines or to define their own turbine characteristics such as hub height, rated output, cut-in and cut-out wind speed (National Renewable Energy Library 2018).

The economic and technical comparison of wind and solar farms will be appropriate as the Covid-19 pandemic has resulted in the closure of many businesses and increased unemployment. The results could validate the feasibility of wind farms in Trinidad and Tobago, create employment opportunities in different sectors, shed light on the necessary legislation amendments and show that the local oil and natural gas-based economy can be diversified. This incorporates the use of the System Advisor Model (SAM) for simulating and obtaining data, a unique method to approach this type of research data collection.

SAM uses various performance and financial models, based on the inputs provided, the performance model will run a simulation to calculate the system's electrical output. The corresponding financial model uses that system output to calculate the annual project cash flows and other financial metrics. There are performance models for wind power, photovoltaic systems, solar water heating and other renewable energies.

METHOD

In the absence of a wind resource file, the wind speed Weibull distribution option was chosen. The wind resource was defined in terms of the average annual

wind speed, Weibull k (shape) factor and the reference height for those values. From these values, System Advisor Model (SAM) displayed a graph of the Weibull probability distribution function (pdf) [4].

The Weibull pdf is a theoretical model recommended by the International Electrotechnical Commission's IEC 61400-12 standard to represent the empirical wind speed distributions at a site [3]. This pdf is comprised of a scale (c) and shape (k) factor and apart from representing the wind speed distribution, it is used to calculate the monthly/annual energy output of a wind turbine or wind farm. Values of average annual wind speed, Weibull k factor and reference height from Chadee (2017) for 2 locations were used. The 2 locations were the MET offices in Piarco, Trinidad and at Crown Point, Tobago.

The wake model chosen was the simple wake model as it calculated the deficit in wind speed at turbines positioned in the wake of upwind turbines. The turbulence coefficient which indicated the variation in wind speed due to the terrain and thermal effects was set at a value of 0.2 [4]. The turbine layout was defined in terms of turbines per row and number of rows.

A degradation rate of 1.6%/year was entered based on a study of 10 years of data from 282 wind farms in the UK [5].

In the uncertainty tab of SAM, the uncertainty model could be enabled to specify uncertainty as a percentage of energy production. However, for the scope of this work, this model was not enabled.

In the revenue tab of SAM is where time of delivery (TOD) factors, capacity and curtailment payments as well as the PPA price could be entered. There were no capacity nor curtailment payments and due to a fixed PPA, the PPA multiplier values were set to one. The PPA price chosen was fixed at 0.13$/kWh, similar to a 2014 wind project in Jamaica [10].

The turbine cost included the purchase, transportation and assembly of the rotor, nacelle and tower modules. The balance of system (BOS) costs included, but were not limited to the foundation, electrical infrastructure, permitting, engineering and project management costs. Financial costs such as contingency and construction finance were added together with the BOS costs to give a final BOS value. After selecting land-based installation, under the capital costs subsection, the turbine costs in USD/kW were entered. The turbine and BOS costs for the turbines ≥3MW were different from those <3MW.

The study executed by Chadee and Clark (2017) determined the LCOE generation for 2 sites with long term wind measurements. These sites are Crown Point in Tobago and Piarco in Trinidad. The LCOE is frequently used to compare and assess alternative methods of energy production.

Six turbines [6] will be used. Two turbines each fall into the 2MW, 2.5MW and ≥3MW category, the variation in cut in, cut out wind speeds, rotor diameters and hub heights will allow comparisons to be made. Metrics such as annual energy produced, NPV, IRR, number of homes powered, etc. will be analyzed and conclusions drawn. Tables 1 and 2 below provide wind turbine characteristics, rated power and number of turbines.

Table 1: Characteristics of wind turbines to be used [6]

Model	Rated Power (kW)	Rotor Diameter (m)	Hub height (m)
Enercon E-101	3050	101	99.0
Enercon E-82 E3	3000	82	98.0
Enercon E-115	2500	115	92.5
Fuhrlander FL-2500	2500	100	85.0
Enercon E-82 E2	2000	82	78.0
AN Bonus 2000/76	2000	76	60.0

Table 2: Rated power and number of turbines at farm capacities [3, 7]

Farm Capacity (MW)	Turbine rated power (MW)	Number of turbines
21.35	3.05	7
21.00	3.00	7
20.00	2.50	8
20.00	2.00	10

RESULTS AND DISCUSSION

The metric tables on the next page (Tables 3-5) show increased turbine performance in Crown Point as opposed to Piarco, all the turbines, regardless of positive or negative financial metrics, had higher AEPs in Crown Point. This is due to several site factors such as annual average wind speeds, wind power densities, shear exponents and Weibull scale factors. Apart from site factors, turbine characteristics such as rotor diameter, hub height, cut in and cut out speeds and drive train design also affect the AEP. The 99m hub height of the Enercon E-101 is approximated to the 100m height in the following analyses.

The net capital costs in tables 3-5 represent the total installed costs which are the product of the total installed costs in ($/kW) and the farm capacity in (kW). The total installed costs ($/kW) are the sum of the turbine, BOS and financial costs. The general trend of a decrease in turbine costs ($/kW) with the larger rated power

Table 3: Metric table for AN Bonus and E-82 E2

Metric	Units	AN Bonus 2000/76 - Crown Point	AN Bonus 2000/76 - Piarco	Enercon E-82 E2 - Crown Point	Enercon E-82 E2 - Piarco
Annual AC energy in year 1	(kWh)	26,950,970	17,369,464	47,606,796	32,423,356
Capacity factor in year 1	(%)	15.4	9.9	27.2	18.5
Real LCOE	(¢/kWh)	15.73	22.63	10.30	13.62
NPV	($)	-16,779,614	-25,738,878	2,534,781	-11,662,615
IRR in year 20	(%)	-1.04	-7.09	6.84	1.41
IRR at project end	(%)	0.62	-4.40	7.71	2.79
Net capital cost	($)	44,880,000	44,880,000	44,880,000	44,880,000

Table 4: Metric table for FL-2500 and E-115

Metric	Units	Fuhrlander FL-2500 - Crown Point	Fuhrlander FL-2500 - Piarco	Enercon E-115 - Crown Point	Enercon E-115 - Piarco
Annual AC energy in year 1	(kWh)	51,755,068	35,327,020	67,740,368	49,401,364
Capacity factor in year 1	(%)	29.5	20.2	38.7	28.2
Real LCOE	(¢/kWh)	8.82	11.43	7.50	9.09
NPV	($)	12,285,037	-3,076,142	27,232,228	10,084,192
IRR in year 20	(%)	10.76	4.64	15.83	9.95
IRR at project end	(%)	11.36	5.67	16.20	10.60
Net capital cost	($)	37,300,000	37,300,000	37,300,000	37,300,000

Table 5: Metric table for E-82 E3 and E-101

Metric	Units	Enercon E-82 E3 - Crown Point	Enercon E-82 E3 - Piarco	Enercon E-101 - Crown Point	Enercon E-101 - Piarco
Annual AC energy in year 1	(kWh)	40,457,064	27,038,680	57,065,672	39,587,860
Capacity factor in year 1	(%)	22.0	14.7	30.5	21.2
Real LCOE	(¢/kWh)	9.57	12.72	7.80	9.82
NPV	($)	5,833,378	-6,713,589	20,830,120	4,487,350
IRR in year 20	(%)	8.81	2.32	14.93	8.15
IRR at project end	(%)	9.51	3.52	15.32	8.90
Net capital cost	($)	31,521,000	31,521,000	32,046,350	32,046,350

turbines is observed. As the BOS costs were taken as 30% of the turbine cost, the turbine costs strongly influence the total installed costs (in $/kW and $). This is seen as the 2MW E-82, 2.5MW E-115 and 3MW E-82 have turbine costs of 1600$/kW, 1330$/kW and 1057$/kW and total installed costs of $44.9 million, $37.3 million and $31.5 million respectively. The 3.05MW E-101 has a higher rated power and farm capacity than the 3MW E-82, hence its higher installed cost of $32 million.

For the Enercon E-82 E3 variation of the farm, the

turbines cost 70.4% of the total installed cost while the Enercon E-115 variation cost 71.3% of the total installed cost.

The net present value (NPV) is a value in $ that includes revenue and expenses, it is positive if a project is economically feasible, and negative if infeasible. The NPVs in tables 3-5 are the after-tax cumulative NPVs in year 25.

SAM produced a cash flow table that could be exported to Microsoft Excel to examine the formulae, the C_n is the sum of the pre-tax and after-tax project returns for a particular year. It was observed that the pre-tax returns for Crown Point were nearly double that of Piarco's. The sum of the cash flow from operating, investing and financing activities made up the pre-tax returns, for these simulations, the investing and financing activities had cash flows of $0.

The cash flow from operating was the earnings before interest, taxes, depreciation and amortization (EBITDA), this value was the revenue minus the total expenses. The total expenses were the sum of the O&M capacity-based expenses and property tax. The capacity-based expense was the same for all the farms with similar capacity and the effect of the property tax was minimal compared to the effect the revenue had on the EBITDA.

The revenue is made up of the PPA revenue and the production-based incentives (PBI), it was mentioned in the methodology that the PBI was set at zero since these are not available in Trinidad and Tobago. The PPA revenue is the product of the net production (AC kWh) and the PPA price (¢/kWh), since the PPA price was constant throughout each simulation, the net production had a significant impact on the revenue. It was already discussed which turbine and site characteristics lead to the AEP being higher at Crown Point. This impact trickled down to the EBITDA which affected the total pre-tax returns and Cn, which ultimately affected the NPV.

The NPV in year zero was the negative of the installation cost and since the NPV was a cumulative value, a low AEP in the following years would mean a lower NPV got added to the previous negative value. This is why a lot of the NPVs for Piarco are negative, from the 2MW AN Bonus having a low AEP of 17.4 GWh with an NPV of -$25.7 million to the 3.05MW E-101 having an AEP of 39.6GWh and an NPV of $4.5 million and finally the 2.5MW E-115 having an AEP of 49.4GWh and an NPV of $10.1 million. While the same turbines at Crown Point had higher AEPs and thus higher NPVs.

On the financial parameters tab in SAM, the PPA price was specified as 0.13¢/kWh and SAM calculated the IRR each year to meet that PPA price. The 20th year was chosen to display its IRR in the metrics tables. The IRR was a cumulative after-tax value, so much like the NPV, it increased each year. The IRR values of the AN Bonus in Table 3 are both negative as the NPVs are quite low, a trend is seen throughout the metric tables that as the NPV increases, the internal rate of return (IRR) does as well as they are from the same formula. The highest IRR is seen from the highest NPV which occurs with the E-115.

From a technical and economic perspective, the best wind farm option is the one comprised of the Enercon E-115 both at Crown Point and Piarco. As was mentioned before, although the 2.5MW E-115 is less than 3MW, its larger (115m) rotor diameter and tall (92.5m) hub height grant it access to a larger swept area and higher, more consistent wind speeds at both sites. This significantly increased its AEP which in turn ensured a high NPV and by extension high IRR. Compared to the rated wind speeds of the E-82 E3 and E-101, the E-115 has a lower value which decreased its LCOE making it even more attractive for investment. The second-best option would be the 3.05MW E-101 owing to its large 101m rotor diameter and 99m hub height, while the worst option would be the 2MW AN Bonus resulting from its short 76m rotor diameter and 60m hub height.

CONCLUSIONS

The monthly energy production outputs of the wind farms were used to calculate the number of households that could be powered each month. The results obtained via the simulations are promising and were validated by various sources including a report by the NREL. They show that with proper legislative amendments and public education, RE energy projects such as wind and solar farms can create employment opportunities in different sectors, reduce carbon emissions in the three main emitting sectors and diversify the local oil and natural gas-based economy.

AUTHOR IDENTIFIERS

Author ORCIDs:
Edwards L: 0000-0002-9814-8737

REFERENCES

1. Qamar, Muhammad Zaid, Mariya Noor, Wahid Ali, and Mohammad Qamar. 2020. Green technology and its implications worldwide. *The Inquisitive Meridian* 3 (1): 10 (2020) https://www.researchgate.net/publication/350443477_Green_Technology_and_its_Implications_Worldwide.
2. United Nations. https://www.un.org/en/climatechange/what-is-climate-change
3. Chadee, Xsitaaz T., and Ricardo M. Clarke. Wind resources and the levelized cost of wind generated

electricity in the caribbean islands of trinidad and tobago. *Renewable and Sustainable Energy Reviews* 81 (2): 2526-2540 (2017) https://doi.org/10.1016/j.rser.2017.06.059 https://www.sciencedirect.com/science/article/abs/pii/S1364032117309930.

4. National Renewable Energy Library. https://sam.nrel.gov/wind.html
5. Staffell I, Green R. How does wind farm performance decline with age?. Renewable Energy 66:775-786 (2014). https://doi.org/10.1016/j.renene.2013.10.041
6. Bauer, Lucas, and Silvio Matysik. https://en.wind-turbine-models.com/turbines
7. Tyler Stehly, and Patrick Duffy. https://www.nrel.gov/docs/fy23osti/84774.pdf
8. Renewables First. https://www.renewablesfirst.co.uk/renewable-energy-technologies/windpower/windpower-learning-centre/how-much-does-a-wind-turbine-cost/
9. Burk, Chris. https://www.aiche.org/resources/publications/cep/2018/january/techno-economic-modeling-new-technology-development
10. New Energy. https://newenergyevents.com/jamaicas-solar-ppa-breaks-records-at-8-5-centskwh-whos-the-big-winner/
11. Ministry of Planning and Development. https://www.planning.gov.tt/content/strategy-wind-power-generation-trinidad-and-tobago

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

An Engineering Clinic-Based Approach to Teaching Process Design and Modeling: Bridging Theory and Practice

Barnabas Gao[a], Thien An Pham[a], Amarelys Rios[a], Corbin Tinker[a], Saugat Bhandari[a], Robert Hesketh[a], C. Stewart Slater[a], and Kirti M. Yenkie[a*]

[a] Rowan University, Department of Chemical Engineering, Glassboro, NJ 08028, United States
* Corresponding Author: yenkie@rowan.edu.

ABSTRACT

Advancing student understanding of process design requires a balanced integration of theoretical knowledge with real-world industrial applications. This study introduces system design thinking-based learning through an engineering clinic approach that bridges the gap between classroom concepts and chemical engineering practice. Using an industrial multiproduct oil pipeline operation as a case study, students are exposed to real-world industrial systems, identify bottlenecks in a process, draw similarities between systems at different scales, and implement control strategies to address a practical industrial problem. In this study, we highlight the collaborative efforts between faculty, students, and industry partners to provide experiential learning in process design, modeling, and control to address the challenge of minimizing product loss during flushing operations in multiproduct petroleum pipeline systems.

Keywords: Dynamic Modelling, Fluid Dynamics, Process Design, Flushing, System Identification

1. INTRODUCTION

1.1 Motivation

The chemical industry is vast, spanning more than 70, 000 products, including ceramics, food, fuels, industrial gases, petrochemicals, pharmaceuticals, plastics, and textiles [1]. Recent advancements in chemical engineering practices have evolved towards increasingly integrated, complex, and data-rich systems operating under tight economic, thermodynamic, environmental, and safety constraints. Modern process industries require engineers who can reason across interconnected systems, predict dynamic interactions, and make informed decisions under uncertainty and variability [2]. Consequently, there is a growing need for systems and holistic design thinking approaches in both industrial practice and engineering education.

1.2 Gap in undergraduate process design and modeling instructions

Despite the shift in modern industrial practices, many undergraduate courses in process design and modeling often remain fragmented. Process design is typically taught as a capstone experience focused on steady-state analysis, equipment sizing, safety, and economic evaluations. Process modeling, on the other hand, is often introduced in the curriculum, with an emphasis on mathematical formulation and analytical solution techniques. Although each course addresses essential competencies, the lack of explicit integration can limit students' ability to understand how design decisions influence the dynamic behavior, operability, and optimization of real processes. Considering the traditional process design and modeling curriculum, students are frequently exposed to modeling in abstract contexts, with limited opportunities to validate models against physical systems and experimental data. In addition, modeling instances for some students is regarded as a purely mathematical exercise rather than a decision-support tool [3].

This disconnect is evident when students encounter industrial problems that do not conform to textbook assumptions. Real systems are subject to transient behavior, nonideal mixing behavior, uncertainties in measurements, equipment limitations, and complex operational

constraints and performance objectives [4]. In such contexts, mathematical models alone are not sufficient; engineers continuously validate theories with experimental data, revise assumptions, and evaluate trade-offs to refine established mathematical models to capture the entire system behavior. Traditional coursework often relies on well-defined problems with complete information and predetermined solutions. In contrast, industrial problems are typically constrained by incomplete data, uncertainties in measurements, design flaws, practical limitations, and operational inefficiencies, as in the case of trial-and-error in most industrial applications.

Thus, it highlights the importance of systems design thinking, an approach that emphasizes understanding interactions among system components, recognizing feedback and nonlinearities, and iteratively refining solutions based on data and constraints. Systems design thinking expands beyond component-level optimization of processes to systems-level performance, enabling students to reason about cause-effect relationships, nonlinearities, and uncertainties associated with processes. This is a critical aspect in chemical engineering because of the coupled nature of process units, material properties (physical, chemical, structural, and thermodynamic), control strategies, and operational decisions and constraints.

1.3 Contributions of systems-oriented process education

This paper presents a systems-oriented, engineering clinic-based approach to teaching process design and modeling. The approach integrates systems design thinking with engineering clinics and project-based environments, emphasizing the direct application of classroom concepts, design, experimentation, and system-level decision-making.

The contributions of this study are threefold. First, this study articulates a structured engineering clinic framework that embeds systems design thinking into the undergraduate curriculum. The framework emphasizes problem identification, experimental setup, stem boundary definitions, and trade-off analysis, enabling students to think holistically about complex processes. Second, it demonstrates how engineering clinics and project-based learning can be used to operationalize this framework through collaborative, industry-inspired problem-solving that mirrors professional industrial practice. Finally, a case study is introduced to demonstrate how pilot-scale experimentation and modeling efforts can be integrated to support iterative design refinements and operational optimization. These contributions are demonstrated through an industrially motivated case study focused on minimizing product loss during flushing operations in petroleum pipelines. The case study requires students to integrate fluid dynamics, process scale-up and scale-down, mixing behavior, experimental analysis, and modeling to minimize product loss, operational downtime, and enhance overall operational productivity.

2. SYSTEMS DESIGN THINKING LEARNING APPROACH

2.1 Motivation for systems design thinking

Process design and modeling are inherently system-level activities. Teaching these concepts requires students to operate at multiple levels of abstraction, often simultaneously. Students must reason from physical intuition to mathematical representation, from component to system-level performance, and from theoretical concepts to practical design decisions. The systems design thinking framework addresses the challenge of the traditional approach to process design and modeling by structuring how students connect, organize, and apply classroom knowledge during design processes and decision-making. Rather than treating design as a culminating activity that follows analysis, the framework positions design thinking as a principle that integrates modeling, design, analysis, and experimentation. This strategy supports cognitive learning [5] by helping students develop mental models of how process units and signals interact and how design decisions propagate in real systems.

Process design and modeling are rarely linear activities; assumptions evolve, models are refined, and design objectives are revisited based on newly acquired data. Thus, systems design thinking legitimizes this iteration as a core learning activity, urging students to treat design, modeling, and experimentation as complementary tools for understanding and improving system performance. This approach mirrors the way engineers approach complex industrial problems. Hence, students learn not only how to model or design a process, but also why a particular modeling approach is appropriate, what its limitations are, and how it informs design decisions. In essence, systems design thinking serves as a bridge between foundational engineering science and real-world design practice. The framework provides students with a repeatable cognitive strategy for navigating complexity, enabling them to approach complex and new problems with confidence, adaptability, and engineering decision-making.

2.2 Engineering Clinic-Based Approach

2.2.1 Concepts of engineering clinics

The engineering clinic program provides a sustainable approach to engineering education that fosters an understanding of theoretical concepts and their application. This is deeply rooted in the top skills defined by the World Economic Forum for being a great engineer. These skills include critical thinking, creativity, people

management, problem-solving, and teamwork [6]. A defining feature of the engineering clinic program is its integration with the core concepts of systems design thinking. Within the clinic projects, students are introduced to an authentic context where systems design thinking can be practiced, tested, and refined.

The engineering clinic program is a hallmark of the undergraduate curriculum at Rowan University. Students are introduced to the clinic program early on, thus spanning the first to final year of their educational journey. Within the clinic program, students:

- Work on open-ended, industry-relevant problems,
- Get the opportunities to engage with stakeholders from federal agencies and industries,
- Get the opportunity to collaborate in multidisciplinary teams,
- Engage with real data, physical constraints, and uncertainties to provide solutions to real-world problems,
- Apply systems design thinking concepts in modeling, experimentation, and design decisions.

These experiences, coupled with classroom concepts, provide a perspective that is essential for preparing students for contemporary industrial and research environments.

2.2.2 Engineering clinics layout

The engineering clinic primarily encompasses dynamic collaboration between student teams (graduate and undergraduate), faculty, and industry, creating a rich educational environment. The Chemical Engineering Department at Rowan University maintains a research-intensive culture and partnerships with diverse federal and industrial organizations. Notable collaborations are with the petrochemical industry, such as ExxonMobil and Valero Energy Corporation; pharmaceutical industry, such as Eli Lilly, AstraZeneca, Bristol Myers Squibb, and Pfizer; food and beverage industry, such as Nestle; specialty chemical industry, such as DuPont and PennAKem; and federal agencies, such as U.S. Environmental Protection Agency (EPA), U.S. Department of Energy, U.S. Army amongst others.

Within the first year, students are introduced to foundational engineering skills designed to develop analytical thinking and support reason-based learning. Building on this foundation, the second year emphasizes technical communication, with focused instruction on technical writing, data presentation, and oral communication—skills essential for effective documentation and professional practice. By the third and final year, students have accumulated sufficient theoretical knowledge and experiential learning to engage in discipline-specific and open-ended research projects. These projects are conducted under the joint supervision of graduate student mentors and faculty advisors, allowing students to integrate their prior coursework with applied research and design activities.

This curricular progression enables process design and modeling to be taught as integrated, system-level activities that reflect how engineering problems are addressed in professional and industrial settings. As students transition into upper-level clinics and research experiences, they are no longer positioned as passive recipients of problem statements but as active participants in defining, interpreting, and refining engineering challenges. Projects are initiated with an emphasis on understanding the overall system objectives, operational context, and performance criteria, often through direct engagement with industrial partners. These interactions expose students to realistic constraints, competing objectives, and evolving project requirements, reinforcing the necessity of system-level thinking from the outset.

In this environment, process design and modeling emerge as complementary tools for reasoning about complex systems rather than isolated technical deliverables. Students use models to formalize their understanding of a system, evaluate design alternatives, and anticipate system behavior under varying operating conditions. Design decisions are revisited iteratively as new experimental data, pilot-scale observations, and stakeholder feedback become available. This iterative refinement mirrors industrial design workflows, where models evolve alongside system understanding and are continuously validated against the physical behavior. Through sustained engagement in this end-to-end design process, students develop a holistic perspective of how process modeling informs design, optimization, and decision-making within complex engineered systems.

2.2.3 Student grading and assessment rubric

While the clinic framework defines how students engage with systems-oriented design problems, meaningful learning also depends on the performance evaluation of the students. To ensure alignment between instructional objectives and student outcomes, assessment strategies must reflect the iterative, integrative, and collaborative nature of clinical experiences. Accordingly, a structured grading rubric was developed to evaluate student performance across technical, analytical, and professional dimensions [7]. The rubric was informed by the work of Drs. Kevin Dahm, James, and Heidi Newell emphasize that effective assessment of complex engineering tasks requires clearly defined performance criteria, explicit achievement levels, and consistency across evaluators. Their framework highlights the importance of moving beyond purely outcome-based grading toward assessment approaches that capture students'

reasoning, decision-making, and progression over time.

Consistent with this perspective, the rubric used in this study decomposes student performance into measurable categories aligned with the learning objectives of systems-oriented process design and modeling. These categories include problem formulation and system definition, appropriateness and justification of the modeling assumptions, integration of experimental data, quality of design decisions, and effectiveness of technical communication. Each category is evaluated using descriptive performance levels that articulate expectations for novice, developing, and proficient work. This structure enables evaluators to assess both the quality of the final deliverables and the rigor of the underlying design process, while also improving transparency for students and enhancing reliability in a team-based, open-ended learning environment.

3. CASE-STUDY OVERVIEW

Here, we present a case study grounded in Computer-Aided Process Engineering (CAPE), integrating process modeling, flowsheet synthesis, and systems-level design. Students develop and simulate flowsheets, scale-down industrial processes to pilot scale, and create CAD-based representations, with emphasis on instrumentation, control, and model-informed decision-making to bridge design and implementation. The scope of the project is to minimize oil downgrade during flushing operations in multiproduct petroleum pipelines. Batch transportation is mainly used in multiproduct oil pipelines, which refer to a network of interconnected pipelines for the continuous transport of several types of oil products in succession [8, 9]. The case study presented in this section explores the lubricant oil industry, which is a subsidiary of the petroleum industry. Unlike conventional pipeline systems, lubricant pipelines accommodate multiple oil grades with stricter quality tolerances and navigate intricate network layouts with bends, filters, and varied geometries, as illustrated in **Figure 1**.

One of the major operational challenges in multiproduct pipelines is product cross-contamination at the interface between oil product batches. The sequential transport of different petroleum products through the same pipeline results in a certain degree of mixing at the interface, leading to a mixture known as transmix or commingled oil. This transmix must be managed carefully to avoid contamination of the products, often requiring reprocessing, which increases operational costs.

A typical lubricant oil industry produces between 1, 500–2, 000 varied products annually. This production process is often segmented across distinct product lines, all converging into a single packaging line as illustrated in **Figure 1**, where residues from previous batches can accumulate and adhere to the walls of the pipelines. As highlighted previously, the issue of interfacial mixing becomes a problem where, due to the mode of transport, there is residue accumulation, impacting the quality of subsequent products. The residues not only compromise product quality and yield but also result in significant financial losses. To mitigate this, industries employ intermittent cleaning between product changeovers, known as flushing. Industries with aqueous or solid-based products usually rinse pipelines with water or inert gas, followed by drying before the next cycle. Unfortunately, oil-based systems found in lubricant manufacturing and blending pose a unique challenge because aqueous media cannot be used for cleaning [10].

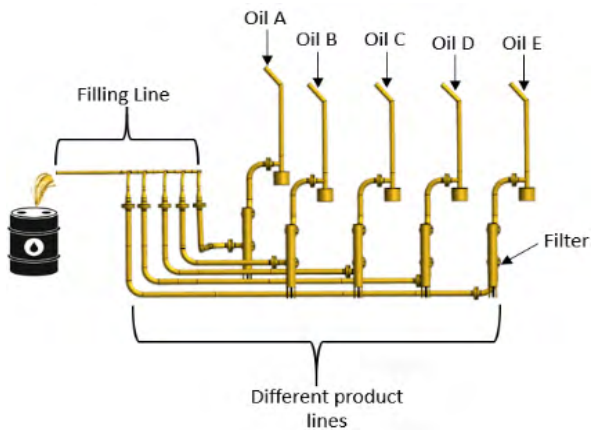

Figure 1. Schematic of a multiproduct pipeline with several products connected to a common filling line for lubricant oil packaging.

The products produced are often destined for sensitive applications, such as in automotive engines and turbines, where even small amounts of contamination would impair performance. Hence, pipeline flushing is a frequently scheduled operation to avoid contamination and mixing of consecutive product batches. The lubricant oil industry usually employs a product changeover operation that involves two stages: (i) removal of the previous (residual) oils from straight sections via the pipeline inspection gauges (PIGs), which is 97% efficient in residual displacement, and (ii) oil removal from the sections with varying diameters, auxiliary equipment, bends, and turns, often known as the unpiggable sections. Stage (ii) is usually accomplished by passing the new oil through the unpiggable section until the residual oil is displaced from the pipeline [9, 11]. This process inevitably leads to the formation of mixed oil segments, which are of limited economic value and must be carefully managed.

At present, flushing operations in lubricant pipelines remain empirical, relying on operators' experience and trial-and-error approaches to determine flush times and volumes. Even at this, industrial practices themselves differ significantly; each facility adopts unique operational procedures, flush volumes, and sequencing protocols

based on experience rather than predictive understanding. This lack of standardization across operators results in inconsistent outcomes, and suboptimal use of resources. This lack of predictive capability results in long hold times, downgrading of high-value products, and substantial capital losses. The operational inefficiency of current practices, combined with the complexity of lubricant pipeline networks, highlights the need for environmentally friendly, cost-effective, and energy-efficient solutions to optimize flushing operations.

Having established the industrial context and technical scope of the case study, the following section describes how the project was intentionally structured to guide students through the core stages of process design and modeling using a systems design thinking approach.

4. SOLUTION DEVELOPMENT

4.1 Experiential learning: engagement with stakeholders and design objectives

Following an initial preview of the industrial challenge, students were guided to move beyond a qualitative problem statement toward the formulation of explicit engineering objectives using a systems design thinking approach. Upon visits to the plant and engagement with stakeholders, students identified persistent inefficiencies during product changeovers in multiproduct pipelines, where flushing operations relied heavily on trial-and-error practices, leading to excessive use of resources, oil for flushing, extended downtime, large volumes of mixed oil formation, and inconsistent product quality. From this context, students were tasked with translating the operational inefficiencies into quantifiable and actionable engineering goals. Two primary objectives emerged: (i) the development of a standardized, repeatable flushing protocol applicable across a range of operating conditions, and (ii) the replacement of heuristic decision-making with a model-based framework capable of supporting smart, data-driven operations.

To achieve these objectives, students engaged in structured and iterative interactions with multiple stakeholders, including plant operators, process engineers, and technical managers. These engagements were conducted through scheduled meetings, progress reviews, and informal technical discussions, providing students with firsthand exposure to professional communication in an industrial setting. Through this process, students learned to elicit requirements, negotiate technical expectations, and reconcile conflicting constraints, which are key competencies in requirements engineering. Importantly, students quickly recognized that any proposed solution would need to be economically viable and minimally disruptive to existing operations, reinforcing the necessity of designing under real-world constraints rather than idealized assumptions. As part of the experiential learning process, students explored alternative approaches to conventional flushing strategies and evaluated their feasibility relative to cost, scalability, and operational complexity. Some of the potential flushing alternatives explored included pipeline coating [12], fluid blasting [13], gel pigging [14], amongst others. These explorations revealed significant limitations associated with intrusive process modifications, prompting a shift toward low-cost, adaptable solutions. This realization ultimately motivated the decision to replicate industrial operations at pilot scale, enabling controlled experimentation while preserving fidelity to real plant behavior. Through this phase, students internalized systems design principles by learning to balance technical ambition with economic and operational realities.

4.2 Pilot plant architecture: design constraints, equipment, and validation approach

A central element of the project involved the design and construction of a pilot plant capable of mirroring the dynamics of the industrial flushing operation. Students were introduced to a fundamental question in process design: how do we design a system mimicking the industrial operation? [15] and tasked with translating conceptual objectives into a physical system. This required decisions regarding equipment selection, piping configuration, sensor placement, control architecture, and data acquisition systems, all of which were informed by both classroom theory and industrial input.

The pilot plant was intentionally designed to serve as both a physical testbed and the foundation of a cyber-physical system. Students evaluated which unit operations and components were essential to replicate industrial behavior, and which could be simplified without compromising system integrity. Key considerations included flow regimes, viscosity contrasts between fluids, the right choice of pump considering the nature of lubricant oils, and the impact of pipeline geometry on mixing behavior. Scaling down from industrial to pilot scale required students to apply principles of similitude and dimensional analysis, assessing the effects of flow rates, pipe geometries, and fluid properties on fluid transport and the ability to replicate the larger system. Product purity determination is done through viscosity testing, as it reflects the functional performance of the lubricant oil product. For instance, if the oil viscosity deviates from specification, it may fail to reduce friction, prevent wear, and dissipate heat.

At industry, this is tested through sampling at the end of each flush operation. As part of the equipment installed on the pilot plant, students installed an in-line viscometer for continuous viscosity measurement during the flushing. The interactions of all equipment installed

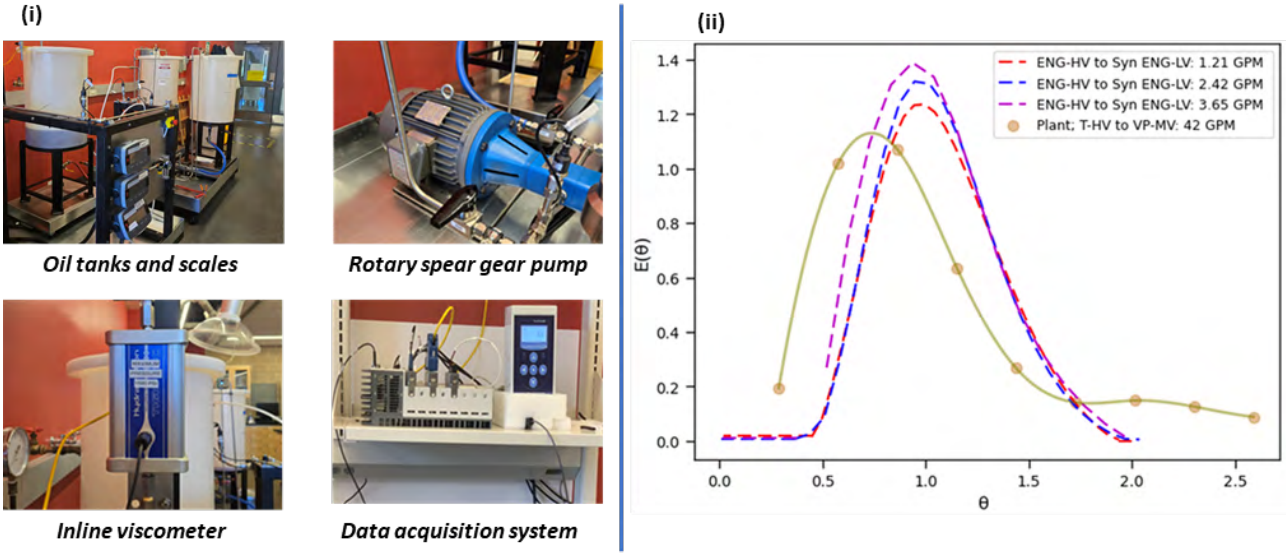

Figure 2. Developed 1/5th industrial-scale pilot plant illustrating (i) the equipment installed in the pilot plant (oil tanks, scales-for mass balance; rotary spear gear pump-for pumping the high viscous lubricant oils; inline viscometer-for real-time viscosity measurement; and data acquisition system-for seamless integration of all measurements and data collection) and (ii) the validation between pilot and industrial plant through residence time distribution studies-a plot of dimensionless RTD, E(Θ) against dimensionless time, Θ.

are through a data acquisition system, which integrates all these components in LabVIEW programming to collect data across multiple channels. Simultaneously, students are introduced to the fundamentals of process instrumentation and data acquisition framework, selecting sensors capable of capturing relevant process variables such as pump power, pressure, viscosity, temperature, and change in mass. Emphasis was placed on understanding how different signals are generated, interpreted, and integrated to produce actionable insight.

To validate the fidelity of the pilot plant, students conducted residence time distribution (RTD) studies comparing experimental results against theoretical expectations and industrial data. These experiments provided a concrete mechanism for students to assess whether the scaled system accurately reflected industrial dynamics [9, 16]. This phase of the project tightly coupled classroom concepts, transport phenomena, process dynamics, and control, with hands-on engineering design. By building a physical system that directly informed modeling efforts, students experienced the full lifecycle of process design, from abstraction to implementation. This is depicted in **Figure 2**. The pilot plant thus became a critical educational bridge, reinforcing systems thinking while anchoring modeling activities in observable, real-world behavior.

4.3 Modeling strategy: systems-level decision support tool

With experimental infrastructure in place, students transitioned to the development of process models intended to capture the flushing dynamics observed in the pilot plant. Rather than prescribing a single modeling approach, students were guided to evaluate alternative strategies, including mechanistic, data-driven, and hybrid models. This evaluation required students to reflect on the purpose of the model prediction, optimization, or decision support and to select an approach consistent with available data, computational complexity, and interpretability. A critical aspect of the process modeling is what is referred to as "*know-how*" for processes. Thus, students were exposed to the theories of process dynamics: how a process changes over time; process identification: building a model from process data; and process control: what are the parameters of interest to keep the process at desired conditions, all of which are key in understanding and developing a model for the flushing process. Emphasis was placed on ensuring that the model assumptions were clear and consistent. Students were encouraged to justify assumptions related to flow behavior, mixing efficiency, and parameter estimation, and to assess the impact of these assumptions on the model accuracy. Parameter identification was performed using experimental data obtained from the pilot plant, reinforcing the iterative relationship between modeling and experimentation. Through this process, students learned that models are not static representations, but evolving tools that must be continually refined as new data becomes available.

From an instructional standpoint, modeling was

framed as a decision-making tool rather than an abstract mathematical exercise. Students evaluated model limitations, explored sensitivity to key parameters, and assessed the trade-off between model fidelity and usability. This approach helped students develop an appreciation for "fit-for-purpose" modeling, where the goal is not maximum complexity, but actionable insight. Results from the developed models on some test case flushing experiments are depicted in **Figure 3**.

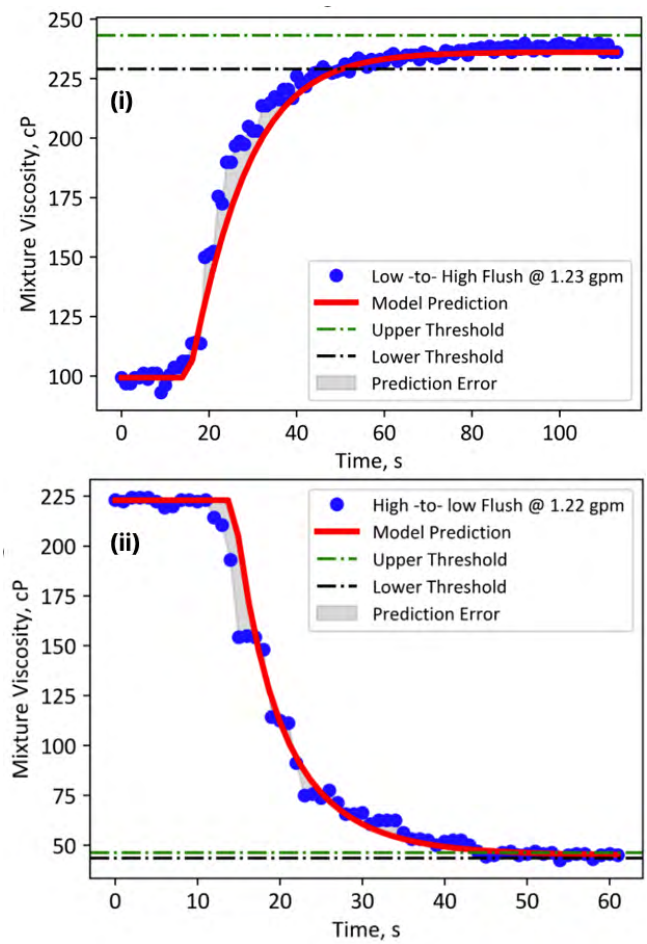

Figure 3. Sample model results validated against experimental flush data performed on the pilot plant for (i) a change from a low viscosity oil to a high viscosity oil and (ii) a change from a high viscosity oil to a low viscosity oil.

The final stage of the project focused on translating validated models into decision support tools capable of informing operational strategy. Students employed their models to evaluate economic, temporal, and performance-based metrics associated with different flushing scenarios. These analyses included minimizing flushing time, enhancing the effective use of oil during flushing, and identifying operating conditions that balance efficiency with robustness.

Optimization was introduced as an extension of design thinking rather than a purely mathematical construct. Students explored trade-offs between competing objectives and learned to interpret optimization results within the broader system context. This systems-level perspective reinforced the idea that optimal solutions are contingent on constraints, uncertainty, and stakeholder priorities.

Through this stage, students completed the full cycle of process design and modeling from problem formulation to decision support, demonstrating how systems design thinking can be effectively embedded within undergraduate education. The case study thus served not only as a technical challenge but as a structured learning environment in which students developed professional competencies essential for modern process engineers.

5. STUDENT LEARNING OUTCOMES AND CONCLUSION

This case study demonstrates how a research-integrated clinic framework can be intentionally structured to teach process design and modeling as interconnected, systems-oriented competencies rather than isolated technical skills. By exposing students to critical industrial challenges, the project enabled students to experience the full engineering design lifecycle from problem definition and stakeholder engagement to modeling, experimentation, and decision support. Rather than receiving predefined objectives or idealized datasets, students were required to translate ambiguous operational challenges into structured engineering goals, reinforcing the central role of problem formulation in professional practice. From a learning perspective, the project fostered student development across multiple dimensions. Early engagement with industrial stakeholders strengthened professional communication skills and exposed students to requirements engineering under real-world constraints. Students learned to negotiate technical feasibility, economic limitations, and operational practicality while maintaining alignment with overarching system objectives. These experiences directly supported the development of systems thinking, as students were required to consider how design choices influenced overall process performance.

The design and validation of the pilot plant played a critical role in bridging theory and practice. Through scaling analysis, equipment selection, and residence time distribution studies, students applied foundational concepts from transport phenomena and process dynamics in a real-world setting. The developed pilot plant reinforced the importance of experimental and model validation, exposing students to the limitations of purely theoretical models. Also, students gained experience in data acquisition systems and interpreting real process signals, skills that are increasingly essential in modern, data-rich

industrial environments. Modeling activities further deepened student understanding by reframing models as decision-support tools rather than purely predictive constructs. By comparing mechanistic, data-driven, and hybrid approaches, students developed an appreciation for model selection, assumption transparency, and parameter identification. Sensitivity analysis and model validation exercises helped students recognize the trade-offs between fidelity, interpretability, and usability, ensuring critical judgment in model-based decision-making.

Overall, the case study illustrates how process design and modeling can be taught holistically through experiential, industry-engaged learning. Students completing the project demonstrated measurable gains in systems thinking, professional communication, experimental design, and model-based reasoning. More importantly, they developed a mindset aligned with contemporary engineering practice, one that values integration, iteration, and informed decision-making under uncertainty. These outcomes and exposure ensured that students are prepared for the multifaceted challenges of modern process engineering, while the clinic framework provides a replicable model for experiential learning at other institutions.

ACKNOWLEDGEMENTS

This research work was supported by the U.S. Environmental Protection Agency's (EPA) Pollution Prevention Program (Grant number: NP-96248222) and by ExxonMobil Product Solutions Company.

AUTHOR IDENTIFIERS

Author ORCIDs:
Barnabas Gao: 0009-0005-0768-5107
Kirti M. Yenkie: 0000-0002-8044-0369
Robert P. Hesketh: 0009-0009-9860-6417
C. Stewart Slater: 0000-0002-1162-5968

REFERENCES

1. W.D. Seider, D.R. Lewin, J.D. Seader, S. Widagdo, R. Gani, K.M. Ng, Product and Process Design Principles: Synthesis, Analysis and Evaluation, John Wiley & Sons, 2016.
2. Rojas L, Yepes V, Garcia J. Complex dynamics and intelligent control: advances, challenges, and applications in mining and industrial processes. Mathematics 13:961 (2025). https://doi.org/10.3390/math13060961
3. G. Ritsema van Eck, A. Ceross, Decision Support Tool, [S.l : s.n.], 2015.
4. Walden JVM, Wellig B, Stathopoulos P. Heat pump integration in non-continuous industrial processes by dynamic pinch analysis targeting. Applied Energy 352:121933 (2023). https://doi.org/10.1016/j.apenergy.2023.121933
5. R.T. Pithers, ERIC - EJ651549 - Cognitive Learning Style: A Review of the Field Dependent-Field Independent Approach. Journal of Vocational Education & Training, 200, (2002) 117–32.
6. The 10 skills you need to thrive in the Fourth Industrial Revolution, World Econ. Forum. (2016). https://www.weforum.org/stories/2016/01/the-10-skills-you-need-to-thrive-in-the-fourth-industrial-revolution/ (accessed January 27, 2026).
7. J. Newell, H. Newell, K. Dahm, Rubric Development And Inter Rater Reliability Issues In Assessing Learning Outcomes, in: 2002: p. 7.991.1-7.991.8. https://peer.asee.org/rubric-development-and-inter-rater-reliability-issues-in-assessing-learning-outcomes (accessed January 28, 2024).
8. Chen L, Yuan Z, Xu J, Gao J, Zhang Y, Liu G. A novel predictive model of mixed oil length of products pipeline driven by traditional model and data. Journal of Petroleum Science and Engineering 205:108787 (2021). https://doi.org/10.1016/j.petrol.2021.108787
9. Gao B, Jerpoth S, Theuma D, Curtis S, Roth S, Fracchiolla M, Hesketh R, Slater CS, Yenkie KM. Improved design of flushing process for multi-product pipelines. Systems and Control Transactions 3:137-144 (2024). https://doi.org/10.69997/sct.171679
10. Shah NK, Li Z, Ierapetritou MG. Petroleum refining operations: key issues, advances, and opportunities. Ind. Eng. Chem. Res. 50:1161-1170 (2010). https://doi.org/10.1021/ie1010004
11. R.P. Hesketh, B. Gao, K.M. Yenkie, S.S. Jerpoth, D.A. Theuma, S. Curtis, et al., A Novel Laboratory-Scale Pilot Plant Study, in: 2024 ASEE Annu. Conf. Expo., 2024. https://peer.asee.org/46468.pdf (accessed March 13, 2025).
12. Zhao X, Hao H, Duan Y, Wang J. A robust superhydrophobic and highly oleophobic coating based on f-sio2-copolymer composites. Progress in Organic Coatings 135:417-423 (2019). https://doi.org/10.1016/j.porgcoat.2019.06.024
13. S.H. Robinson, Vehicle mounted apparatus for high-pressure fluid blasting, US9127419B2, 2015. https://patents.google.com/patent/US9127419B2/en (accessed April 23, 2020).
14. Sanccus, Pipeline Gels - Clean, Protect & Maintain, (2020). //www.sanccus.com/pipelines (accessed November 17, 2025).
15. Seider WD, Bhamla S, Dunn J, El-Halwagi M, Hanrath T, Hasan MMF, Hedengren J, Hirshfield L, Lin X", Maravelias CT, Piffarerio M, Prescott SW, Spicer PT, Squires TM, Thomas CU, Tom J, Vaeth KM, Wisniewski E, Zavala VM. Teaching chemical

product design. Education for Chemical Engineers 52:101-110 (2025). https://doi.org/10.1016/j.ece.2025.04.002
16. H.S. Fogler, Elements of chemical reaction engineering, Fifth edition, Prentice Hall, Boston, 2016.

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

A pedagogical framework for sustainability learning : the case of Industrial Ecology

Marianne Boix[a*], Sydney Thomas[a], Lea van der Werf[a] and Ludovic Montastruc[a]

[a] LGC UMR 5503, Toulouse INP, CNRS, Université Toulouse 3, 4 Allée Emile Monso, 31432 Toulouse Cedex 4
* Corresponding Author: Marianne.boix@toulouse-inp.fr

ABSTRACT

The accelerating social, environmental, and economic challenges of the twenty-first century call The growing complexity of sustainability challenges calls for educational approaches that integrate technical analysis with multi-stakeholder decision-making. Industrial ecology (IE) provides a relevant framework by combining systems thinking, resource flow analysis, and socio-environmental considerations. However, it is still predominantly taught through traditional lecture-based methods, limiting students' ability to engage with real-world complexity. This paper proposes and evaluates an experiential pedagogical framework based on industrial ecology, combining stakeholder role-play, industrial symbiosis scenario design, and multi-criteria decision analysis (MCDA). Implemented in a semester-long course, the framework enables students to collaboratively design and evaluate resource-exchange networks while representing different stakeholder perspectives. Results show significant improvements in both technical skills (systems analysis, resource flow evaluation) and transversal competencies (communication, negotiation, reflexivity). The approach enhances student engagement and highlights the role of value judgments in sustainability decision-making. The framework offers a transferable methodology for sustainability education in engineering and related disciplines.

Keywords: Education, Environment, Multi-agent approaches, Work group, Industrial ecology

I. PEDAGOGICAL CONTEXT

1.1. Industrial Ecology as a Discipline

Industrial ecology (IE) conceptualizes industrial systems by analogy with natural ecosystems, emphasizing the circulation of materials, energy, and information across production and consumption networks. Emerging in the late 20th century, it aims to move beyond linear resource use toward circular and regenerative models [1]. At its core, IE studies how industrial processes interact with the environment and how waste from one system can become a resource for another, thereby reducing overall environmental impacts.

A defining feature of IE is its systemic perspective. Rather than focusing on isolated processes, it situates them within broader socio-technical and environmental contexts. Tools such as material flow analysis (MFA), life-cycle assessment (LCA), optimization, and input-output analysis enable the evaluation of resource use and environmental impacts across multiple scales [2].

IE also incorporates socio-economic and institutional dimensions, as the implementation of industrial symbiosis depends on governance, economic incentives, and collaboration between actors [3]. As an interdisciplinary field bridging engineering, natural sciences, and social sciences, IE aligns with sustainability science.

In this work, industrial ecology is used as the central conceptual framework, providing a systemic approach to analyze resource flows and interactions between industrial actors. Within this framework, sustainability represents the overarching objective, while industrial symbiosis constitutes a key operational strategy to improve resource efficiency through exchanges of materials, energy, and water.

1.2. Pedagogical approaches to sustainability

Teaching sustainability in higher education requires approaches that reflect the complexity and multi-dimensionality of real-world challenges. Traditional lecture-based methods are often insufficient to prepare students

for integrative and participatory decision-making contexts [4]. As a result, active and experiential learning approaches are increasingly adopted.

Problem-based learning (PBL) encourages students to address open-ended problems, fostering critical thinking and autonomy [5]. Similarly, experiential learning frames education as a process combining experience, reflection, and application, helping students connect theory with practice [6].

Role-play exercises are particularly relevant in sustainability education. By assigning students to different stakeholders, they expose learners to diverse perspectives and highlight the negotiation processes inherent in sustainability governance [7]. These activities develop communication, systems thinking, and the ability to handle conflicting objectives.

Together, these approaches support transformative learning, fostering reflexivity, value awareness, and collaborative problem-solving skills [8]. They enable students to combine analytical tools with the interpersonal competencies required for multi-actor sustainability decision-making.

1.3. Decision-making in sustainability contexts

Sustainability challenges are characterized by complexity, uncertainty, and conflict among diverse objectives. Decisions in this domain rarely involve clear and solutions, rather, they require balancing environmental, economic, and social dimensions while reconciling the interests of multiple stakeholders. This complexity demands structured decision-support approaches that can make trade-offs explicit, enable transparent evaluation of alternatives, and facilitate negotiation among actors with divergent priorities.

One widely used tool in sustainability contexts is multi-criteria decision analysis (MCDA). Unlike single-objective optimization methods, MCDA allows decision-makers to assess alternatives across multiple, often conflicting criteria [9]. In sustainability applications, these criteria might include cost, different environmental impacts (especially when dealing with Life cycle assessment approaches LCA), social equity, and feasibility. By assigning weights to criteria and scoring alternatives, MCDA provides a structured framework to illuminate tensions, clarify stakeholder preferences, and identify options that balance competing objectives. These methods do not eliminate conflict but instead makes value judgments explicit, thereby fostering more informed and transparent discussions.

Beyond quantitative tools, decision-making in sustainability also relies heavily on deliberative and participatory approaches. These methods emphasize the inclusion of diverse perspectives in decision processes, acknowledging that technical expertise alone cannot determine sustainable outcomes [10]. Stakeholder engagement, negotiation, and dialogue are therefore critical components of sustainability governance, helping to build legitimacy, trust, and shared understanding. Role-play exercises in higher education mirror these dynamics, offering students a platform to practice the skills of advocacy, compromise, and consensus-building.

The integration of MCDA with participatory approaches represents a promising pathway for both real-world sustainability governance and pedagogy. While MCDA structures the technical assessment of trade-offs, participatory methods ensure that social values and actor perspectives are incorporated.

1.4. Gaps in current pedagogy

Despite growing interest in innovative approaches to sustainability education, significant gaps remain in how higher education prepares students to confront the complexity of sustainability challenges.

A key gap lies in the integration of technical and soft skills. Engineering and environmental programs frequently emphasize analytical tools, such as life-cycle assessment or material flow analysis, without providing opportunities for students to apply these tools in socially and institutionally complex settings. Conversely, pedagogical activities focused on dialogue or values exploration may neglect the technical rigor required for sustainability decision-making. Few educational models successfully combine both dimensions, leaving students either technically competent but low prepared for negotiation, or socially aware but lacking in analytical capabilities.

Another limitation concerns the representation of multi-actor dynamics. Sustainability problems are rarely solved by single actors and most classroom exercises do not capture these dynamics. It can lead to a lack of preparedness when graduates enter professional contexts where collaboration and conflict resolution are critical.

Finally, sustainability pedagogy has often underutilized the potential of industrial ecology as a teaching platform. While IE's systemic and interdisciplinary nature makes it highly relevant for experiential learning, its application in classroom exercises remains relatively limited compared to more established approaches such as case studies or project-based learning. Addressing these gaps requires educational models that merge technical analysis with stakeholder engagement, thereby reflecting the intertwined realities of sustainability transitions.

II. METHODOLOGY: PEDAGOGICAL FRAMEWORK

Designing a pedagogical framework that both reflects the complexity of sustainability challenges and equips students with practical skills requires careful

consideration of teaching methods, learning environments, and assessment strategies. The approach presented here combines role-play, scenario design, and decision-support tools within the broader conceptual lens of industrial ecology (Fig.1). The framework is structured to simulate the dynamics of real-world sustainability governance while grounding student activities in the analytical tools of systems thinking. This section outlines the four main components of the approach: role-play and stakeholder simulation, scenario design and industrial symbiosis modeling, integration of decision-support methods, and organization of the learning environment including assessment.

2.1. Role-Play and Stakeholder Simulation

At the core of the pedagogical design is the use of role-play to represent the perspectives of diverse stakeholders. Students are divided into groups, with each group assigned to a specific role such as an industrial firm, a local authority, a non-governmental organization (NGO), or a community organization. The selection of roles reflects the variety of actors typically involved in industrial ecology initiatives and broader sustainability governance.

The role-play format serves several purposes. First, it compels students to adopt perspectives that may differ from their own, fostering empathy and the ability to understand conflicting values. For example, a group representing a manufacturing company may prioritize cost efficiency and competitiveness, while a community group might emphasize environmental justice or local employment. Second, role-play highlights the interdependencies between actors, as no single group can achieve sustainability objectives in isolation. Students must therefore engage in negotiation and coalition-building to advance their positions. Each group represents a distinct stakeholder rather than a single company. These stakeholders may include industrial firms, local authorities, NGOs, or community representatives, each reflecting different roles, interests, and decision-making perspectives within the eco-industrial system.

Role descriptions are designed to provide both structure and flexibility. Each group receives a background brief outlining its objectives, constraints, and potential strategies. These briefs include information on resource availability, financial limitations, and political pressures, ensuring that students grapple with realistic trade-offs. At the same time, students are encouraged to exercise creativity in developing arguments and strategies, thereby enhancing engagement and ownership of the learning process.

2.2. Scenario Design and Industrial Symbiosis Simulation

Following the role-play introduction, students are tasked with designing scenarios for an industrial symbiosis—networks in which waste or by-products from one industry become inputs for another. The exercise requires groups to collectively design a system of exchanges that balances environmental, economic, and social objectives. Scenario design is typically iterative: groups first propose preliminary ideas, then refine them through negotiation and systems analysis.

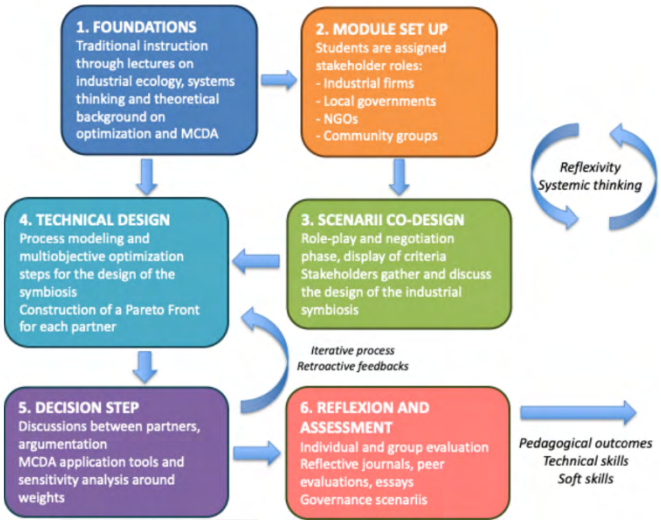

Figure 1. Implemented methodology

This stage emphasizes the systems perspective central to industrial ecology. Students must consider not only bilateral exchanges between firms but also the broader implications of network design, including resource efficiency, resilience, and social impacts. For instance, a proposed energy exchange between two companies may reduce carbon emissions but raise concerns about cost allocation or local employment effects. Students learn to map these interconnections and to identify synergies and conflicts.

The scenario design culminates in a simulation exercise in which the proposed system is tested through qualitative and quantitative analysis. Tools such as material flow analysis (MFA) or simplified life-cycle assessment (LCA) can be incorporated, depending on the technical background of the students. This analytical component ensures that the scenarios are evaluated not only on the basis of negotiation outcomes but also on their environmental and economic feasibility. By bridging stakeholder negotiation with technical evaluation, the exercise reflects the multi-dimensional character of real-world sustainability transitions.

2.3. Decision-Support Tools in the Classroom Learning Environment and Assessment

To structure the negotiation and scenario design process, decision-support tools are introduced, with particular emphasis on multi-criteria decision analysis

(MCDA). MCDA provides a formal framework for evaluating alternatives across multiple dimensions, thereby making explicit the trade-offs inherent in sustainability decision-making.

In practice, the exercise proceeds by defining a set of evaluation criteria that represent the interests of different stakeholders, such as greenhouse gas emissions, operational costs, social equity, and local employment. Each group assigns relative weights to the evaluation criteria through a structured discussion process reflecting their stakeholder priorities. This process may involve direct allocation (e.g., distributing 100 points across criteria) or qualitative ranking followed by normalization. Differences in weights between groups are preserved to reflect conflicting perspectives and to support subsequent negotiation.

The introduction of MCDA has several pedagogical benefits. It compels students to translate qualitative arguments into quantitative or semi-quantitative terms, thereby sharpening analytical skills. It also highlights the value-laden nature of sustainability decision-making: results vary depending on the weights assigned to different criteria, and no scenario is universally optimal. This reinforces the importance of negotiation, compromise, and reflexivity. In addition, the use of MCDA familiarizes students with a widely applied decision-support tool, enhancing their methodological repertoire for future professional contexts.

While MCDA is the central tool in this framework, other decision-support methods can be incorporated, including scenario analysis, cost-benefit analysis, or stakeholder mapping In practice, MCDA is implemented using simple spreadsheet-based tools or scoring matrices adapted to the students' level. Additional tools such as simplified MFA calculations, qualitative LCA indicators, or stakeholder mapping may also be used depending on course objectives and available time.

2.4. Learning Environment and Assessment

The successful implementation of this pedagogical framework requires careful design of the learning environment. The exercise is typically embedded in a semester-long course on industrial ecology or sustainability systems, with the role-play and scenario design component spanning several weeks. This duration allows for iterative negotiation, technical analysis, and reflection. Class sizes of 20–40 students are manageable, enabling multiple stakeholder groups while maintaining opportunities for instructor guidance. All groups participate in the evaluation process; however, the level of quantitative analysis may differ depending on the stakeholder role. Technically oriented groups (e.g., industrial actors) may perform more detailed calculations, while other stakeholders focus on qualitative assessment and interpretation of results.

The physical and digital learning environment is also significant. Face-to-face interaction supports negotiation and collaboration, while digital platforms can be used for scenario modelling, data sharing, and documentation. Hybrid or online adaptations are possible but require deliberate facilitation to ensure engagement.

Assessment is based on a combination of group and individual components. Group performance is evaluated through the quality, feasibility, and coherence of the proposed symbiosis scenario, including its technical and sustainability performance. Individual assessment is conducted through reflective reports, where students analyze their learning process, stakeholder role, and decision-making strategies. In some cases, peer evaluation is used to ensure balanced participation within groups.

Finally, the framework emphasizes formative feedback throughout the exercise. Instructors play a facilitative role, guiding discussions, posing critical questions, and helping students connect their experiences to broader theoretical concepts. Feedback sessions allow students to iteratively improve their scenarios and reflect on their group dynamics. Such scaffolding enhances the depth of learning and ensures that both technical and soft skills are developed in tandem.

III. RESULTS AND DISCUSSION

The implementation of the pedagogical framework described in the previous section generated insights into both the acquisition of technical knowledge and the development of essential soft skills among students. By integrating role-play, scenario design, and decision-support tools within an industrial ecology context, the course fostered multidimensional learning experiences. This section presents and discusses the outcomes observed, highlighting the benefits, challenges, and broader implications for sustainability education.

3.1. Technical knowledge acquisition

One of the most significant outcomes of the exercise was the enhancement of students' technical understanding of industrial ecology concepts. The process of designing industrial symbiosis scenarios required students to engage deeply with material and energy flows, environmental impacts, and resource optimization strategies. Tools such as material flow analysis (MFA) and simplified life-cycle assessment (LCA) allowed students to quantify the efficiency of proposed exchanges and to evaluate environmental trade-offs systematically.

For example, groups representing industrial firms were able to identify potential waste streams that could be repurposed for other industries, reducing both costs and environmental footprints. Similarly, students in community and NGO roles highlighted the importance of local environmental quality and social equity, prompting teams

to adjust scenarios to mitigate negative externalities. Through these exercises, students developed the ability to integrate ecological, economic, and social considerations into the design of industrial networks. This aligns with prior studies emphasizing the value of experiential approaches in fostering systems thinking and applied technical knowledge (Graedel & Allenby, 2010; Chertow, 2000).

Moreover, the iterative nature of the scenario design encouraged students to refine their understanding of technical concepts. Initial proposals often revealed inefficiencies or conflicts that could only be resolved by revisiting calculations, redesigning material flows, or reconsidering energy exchanges. This iterative process mirrors real-world engineering practice, where problem-solving is seldom linear, and decisions must account for interdependencies and constraints.

3.2. Development of soft skills

Equally important were the gains in soft skills, which are critical for sustainability governance but often underemphasized in technical curricula. Through stakeholder role-play and group negotiation, students practiced communication, collaboration, and conflict resolution in a structured environment. Negotiation exercises required groups to advocate for their assigned priorities while engaging with competing perspectives, promoting empathy and the ability to understand others' motivations.

Table 1. Summary of observed learning outcomes

Learning dimension	Observed outcomes
Technical skills	Improved understanding of MFA, LCA, resource flows
Systems thinking	Ability to identify interdependencies and trade-offs
Communication	Enhanced argumentation and negotiation skills
Collaboration	Active participation in multi-stakeholder discussions
Reflexivity	Increased awareness of value-based decision-making

The multi-actor nature of the exercise illuminated the challenges of consensus-building. Students frequently encountered conflicting objectives, such as balancing cost efficiency with environmental responsibility or reconciling industrial growth with community well-being. These tensions prompted discussions not only about technical feasibility but also about ethical considerations, power dynamics, and social responsibility. By navigating these interactions, students enhanced their capacity for collaborative problem-solving and developed an appreciation for the complexity inherent in multi-stakeholder decision-making (Table 1).

Reflective assignments and debriefing sessions further reinforced these soft skills. Students were asked to evaluate their own negotiation strategies, identify points of tension, and propose ways to improve communication and decision-making. Many reported heightened awareness of group dynamics, leadership roles, and the importance of framing arguments in terms understandable to diverse audiences. These outcomes are consistent with literature suggesting that experiential learning fosters both cognitive and interpersonal competencies [7, 8].

3.3. Student Engagement and Reflexivity

The combined technical and interpersonal focus of the exercise contributed to high levels of student engagement. Role-play and scenario design provided a sense of ownership and agency, motivating students to invest effort in both the analytical and negotiation components. Anecdotal observations and course feedback indicated that students were particularly responsive to the opportunity to see how individual decisions impacted the broader system, fostering a sense of responsibility and connection to real-world sustainability challenges.

Reflexivity was another important outcome. By participating in structured negotiation and decision-support exercises, students became aware of the underlying value judgments embedded in sustainability decisions. For instance, the use of multi-criteria decision analysis (MCDA) made explicit how differing stakeholder priorities influenced outcomes. Students recognized that "optimal" solutions are context-dependent and contingent upon the relative weighting of environmental, economic, and social criteria. This awareness of subjectivity in decision-making is crucial for preparing future professionals to engage constructively in contested sustainability arenas.

3.4. Challenges and Limitations

Despite the benefits, several challenges were identified in implementing the framework. First, the complexity of the exercise required careful scaffolding to prevent students from becoming overwhelmed. Some groups initially struggled to manage multiple criteria, stakeholder perspectives, and technical calculations simultaneously. Instructors played a critical role in guiding discussions, clarifying technical concepts, and facilitating negotiation to maintain momentum.

Second, participation levels varied within groups. While some students emerged as proactive leaders, others assumed more passive roles, highlighting the need for assessment strategies that capture both individual and collective contributions. Reflection journals and peer evaluations were useful in mitigating this issue, but variability in engagement remained a concern.

Third, time constraints limited the depth of analysis possible within a semester-long course. While students

were able to explore material flows, energy exchanges, and social impacts, more advanced modeling or quantitative assessment would require extended time or dedicated laboratory sessions. Nonetheless, the exercise successfully demonstrated the value of experiential learning even with these temporal limitations.

Finally, the exercise occasionally revealed tensions between realism and pedagogical objectives. While the scenarios aimed to simulate real-world complexities, simplifications were necessary to ensure accessibility for students with diverse technical backgrounds. Balancing fidelity to industrial ecology principles with the practical constraints of classroom implementation remains an ongoing challenge.

3.5. Comparative Reflection with Other Pedagogical Methods

Compared to traditional lectures or case-study approaches, the integrated role-play and scenario-based framework offers several advantages. It actively engages students in both technical and social dimensions of sustainability, bridging a gap that is often left unaddressed in conventional curricula. Unlike static case studies, which present resolved scenarios, this approach exposes students to dynamic negotiation, iterative problem-solving, and multi-stakeholder complexity.

Furthermore, the combination of MCDA with role-play provides a structured yet flexible framework for decision-making. Students are able to quantitatively evaluate trade-offs while simultaneously navigating qualitative stakeholder interactions. This dual emphasis fosters integrative thinking and mirrors professional contexts in sustainability engineering, where analytical rigor and interpersonal skills are equally important.

Nevertheless, the framework complements rather than replaces traditional methods. Foundational lectures, readings, and theoretical instruction remain essential to provide the conceptual grounding that supports applied exercises. The strength of the approach lies in its ability to translate theory into practice, illustrating the interconnectedness of environmental, social, and economic considerations while cultivating both technical and soft skills.

IV. CONCLUSIONS AND PERSPECTIVES

The outcomes of the pedagogical framework presented in this study carry significant implications for the design and delivery of sustainability education in higher education. First, the integration of industrial ecology with experiential learning demonstrates the value of interdisciplinary approaches in preparing students to address complex, multi-dimensional challenges. By combining technical analysis with stakeholder engagement and negotiation exercises, the framework moves beyond traditional siloed instruction, reflecting the interconnected realities of sustainability practice. This suggests that engineering and environmental curricula can benefit from embedding experiential modules that actively connect theory to practice.

Second, the framework underscores the importance of developing both technical and soft skills. Modern sustainability challenges require professionals who can not only analyze material and energy flows but also navigate diverse stakeholder interests, communicate effectively, and mediate conflicts. The role-play and decision-support components of the exercise provide structured opportunities for students to cultivate these capabilities. Importantly, students develop a heightened awareness of the value-laden nature of sustainability decisions, gaining insight into how social, economic, and environmental priorities intersect and sometimes conflict. Integrating these skills into higher education curricula equips graduates to engage meaningfully in real-world decision-making contexts, enhancing their employability and professional effectiveness.

Third, the framework highlights the potential of experiential learning for fostering systems thinking and reflexivity. Students are required to evaluate the broader implications of their decisions within interconnected industrial networks, encouraging them to consider long-term, cross-scale consequences. Reflective assignments and debriefing sessions amplify this effect, prompting students to critically assess their assumptions, strategies, and learning processes. This aligns with the broader goal of sustainability education to produce graduates who are not only knowledgeable but also capable of adaptive, reflective problem-solving in complex and uncertain environments.

The framework also suggests avenues for scaling and adaptation. While developed for industrial ecology courses in engineering or environmental programs, the approach can be adapted to other disciplines concerned with sustainability, including business, policy, or urban planning. The combination of role-play, scenario analysis, and decision-support tools is sufficiently flexible to accommodate varying student backgrounds, learning objectives, and institutional contexts. Moreover, hybrid or digital implementations are possible, potentially expanding access to larger or more geographically distributed cohorts.

Finally, the study highlights the value of structured pedagogical innovation. Implementing such experiential frameworks requires careful planning, resources, and instructor facilitation. Success depends on balancing realism with manageability, scaffolding technical content alongside soft-skill development, and providing timely feedback. Institutions seeking to incorporate similar approaches must recognize that meaningful learning outcomes emerge not solely from the content itself but from the deliberate design of interactive, reflective, and

integrative learning experiences.

AUTHOR IDENTIFIERS

Marianne Boix: 0000-0003-0993-3395
Sydney Thomas: 0009-0004-7887-1356
Lea van der Werf: 0000-0002-2977-3504
Ludovic Montastruc: 0000-0003-3141-6618

REFERENCES

1. Frosch RA, Gallopoulos NE. Strategies for manufacturing. Sci Am 261:144-152 (1989). https://doi.org/10.1038/scientificamerican0989-144
2. T. E. Graedel, B. R. Allenby, 2010. Industrial ecology and sustainable engineering. Prentice-Hall. 426pp.
3. Chertow MR. Industrial symbiosis: literature and taxonomy. Annu. Rev. Energy. Environ. 25:313-337 (2000). https://doi.org/10.1146/annurev.energy.25.1.313
4. Wiek A, Withycombe L, Redman CL. Key competencies in sustainability: a reference framework for academic program development. Sustain Sci 6:203-218 (2011). https://doi.org/10.1007/s11625-011-0132-6
5. Takahashi Y. Problem?based learning and task?based learning: a practical synthesis. The Kaohsiung J of Med Scie 24: (2008). https://doi.org/10.1016/s1607-551x(08)70091-3
6. Kolb, D.A., 1984. Experiential learning. Englewood Cliffs (NJ): Prentice-Hall.
7. Dieleman H, Huisingh D. Games by which to learn and teach about sustainable development: exploring the relevance of games and experiential learning for sustainability. Journal of Cleaner Production 14:837-847 (2006). https://doi.org/10.1016/j.jclepro.2005.11.031
8. J., Mezirow, 2000. Learning to think like an adult. In J. Mezirow et al., *Learning as Transformation* (p. 3-34). San Francisco, Jossey-Bass.
9. Belton V, Stewart TJ. An integrated approach to MCDA. Multiple Criteria Decision Analysis :331-343 (2002). https://doi.org/10.1007/978-1-4615-1495-4_11
10. Reed MS. Stakeholder participation for environmental management: a literature review. Biological Conservation 141:2417-2431 (2008). https://doi.org/10.1016/j.biocon.2008.07.014

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Systems & Control Transactions

Original Research Article
Peer Reviewed Conference Proceeding
ESCAPE 36 - European Symposium on Computer Aided Process Engineering
Sheffield, UK. 21-24 June 2026

PSE PRESS

Closing the Digital Gap: A Scaffolded Pathway for Developing Digitalisation Skills in Undergraduate Chemical Engineering Curricula

E. Routoula[a], J. Bestenlehner[a], and M. Zandi[a]*

[a] University of Sheffield, School of Chemical Materials & Biological Engineering, Sheffield, UK
* Corresponding Author: m.zandi@sheffield.ac.uk.

ABSTRACT

Digital competency is now core to chemical engineering practice, yet the extent and coherence of digitalisation skills provision across undergraduate curricula remain uneven. This study maps qualitatively and qualitatively digital learning outcomes across undergraduate chemical engineering programmes at the University of Sheffield, against a digital skills framework (data analysis, process simulation, process automation & control, reproducible workflows, programming, data governance). In recent years, digital skills education within chemical engineering education has advanced considerably, driven by the broader industrial shift toward Industry 4.0 and reinforced by the global challenges. Academic institutions have begun to integrate digitalisation-related content more deliberately within syllabus, in alignment with degree programme accreditation requirements and industry needs. Beyond introductory spreadsheet manipulation and basic programming, many courses are now embedding more advanced digital capabilities, including data engineering, scientific programming, process analytics, and optimisation techniques. Results of this study reveal a scaffolded digital pathway, where students are exposed gradually through core taught content and applied knowledge activities to a variety of digital tools across their years of study in a by-design approach at earlier years that becomes organic later on. From limited exposure and credit correlation in Year 1, students complete Year 4 having at least quadrupled their exposure to digital tools, with higher credits awarded to "digital" in the case of particular elective modules or research projects. This study offers a methodology and a worked example for chemical engineering programmes seeking measurable uplift in Industry 4.0 readiness and impact on chemical engineering graduate outcomes.

Keywords: Education, Industry 4.0, Modelling and Simulations, Curriculum, Digital Chemical Engineering, Digital Skills

INTRODUCTION

The practice of chemical engineering is being fundamentally altered by Industry 4.0 and the rise of digitalisation. The industry relies more and more on automation, digital twins, advanced modelling, and data-centric decision-making strategies [1]; innovations that not only revolutionise how processes are designed and optimised, but are also redefining the core competencies expected of new graduates. Traditionally, undergraduate curriculums of chemical engineering programmes have included foundational tools like excel-based tools, drawing tools and use of industry-approved process simulation software [2], operating under the assumption that advanced digital proficiency would be acquired during employment. However, employer feedback and recent studies suggest this model is now inadequate, as companies increasingly identify gaps in digital skills competence of graduate recruits and have started requiring proof of digital competency during the recruitment phase [3].

In reaction to these demands, universities have integrated statistics, coding, process simulation, process systems engineering and optimisation into chemical engineering programmes [4], working alongside profes-

sional accreditation bodies and industry partners. Despite this, there is still significant ambiguity regarding the definition, progression, and consistent implementation of "Digital Chemical Engineering" (DCE) across different institutions, as digital skills is an emerging area not as well defined as others, especially given that digitalisation is under continuous development.

At the University of Sheffield, the undergraduate curriculum of chemical engineering underwent a major transformation in 2017 [5], introducing three curriculum pillars, Core Knowledge, Practical & Design, Enabling Science, each covered by revamped or newly-introduced modules. Alongside and intertwined with these pillars, employability skills were also envisioned as a key component, to ensure that graduates would have a strong technical and transferable skillset to add value to the workforce. Since then, substantial work has been done to ensure that employability is fully embedded into the curriculum [6, 7], without however discussing explicitly the emerging importance of digital skills or attempting a mapping of what and how much of is included in the curriculum in terms of digital skills.

In the UK, the accreditation guidelines of the Institute of Chemical Engineers (IChemE) and the Engineering Council with their Accreditation for Higher Education Programmes (AHEP4) guidelines set standards and expectations when it comes to programme-wide learning outcomes while allowing for flexible curriculum design. While both accreditation bodies mention digital tools and refer to particular digital skills, there is no clarity on the expected coverage of such skills in the curriculum.

This study delivers a qualitative and quantitative mapping of DCE content within IChemE-accredited chemical engineering programmes offered at the University of Sheffield, establishing a baseline for digital readiness and a mapping tool that can be used from other institutions.

METHODOLOGY

Data Collection and Curriculum Mapping

Following from output of a survey deployed to both industry and academic stakeholders in the chemical engineering sector to identify digital skills of importance and the level of graduate readiness across these skills [2], and using informed judgement and discussions across contributors who are all heavily involved with curriculum design in their respective institutions, DCE provision was classified across eight domains:

1. programming & computation: ability to understand, write, and execute code.
2. data literacy & statistics: understanding of statistical principles, use of data visualisation for effective data communication.
3. process modelling & simulation: use of specialised software to develop models and simulate unit operations.
4. optimisation & process systems engineering: use of systems engineering principles and algorithms to identify best possible variables.
5. control, automation & instrumentation: design and implementation of control systems to monitor and maintain process variables towards safe and stable plant operation.
6. AI/ML & digital twins: use of tools and software to identify patterns, development of virtual, real-time depiction of a process or plant that can be updated using process data, predict performance and diagnose issues in a simulated environment.
7. software engineering practices (version control, testing, reproducibility): tracking and managing changes to code over time, reverting to previous versions, ensuring reliable reproducibility.
8. data governance, ethics & cybersecurity: understand the importance of data storage, quality, security and access, address moral implications around the use of digital tools.

Using analysis of available information in module descriptors, learning outcomes, and assessment briefs where possible, supplemented by additional information based on knowledge of teaching specific courses or from communications with other module leaders, a mapping of the DCE curriculum provision for chemical engineering courses offered at the University of Sheffield was done against each domain, per year of study (years 1-3 for BEng courses and 1-4 for MEng courses). Mapping was done in a qualitative manner, indicating a "yes/no", with the option to provide representative examples or comments of how each domain content is covered where relevant. With the understanding of content being covered formally in the curriculum, i.e. core/taught and formatively assessed as a learning outcome, and informally, as by-product/applied, i.e. taught but not assessed, or implicitly expected in later years of study, mapping was designed to recognise both dimensions. Furthermore, a quantitative mapping approach was attempted, using module credits and/or hours of engagement, and the available information of the type and weight of assessment. Where such information was not readily available, an educated assumption was made. For example, for a 15 credit module that comprises of an assignment worth of 30% of the final mark, which has a clearly stated task of using a modelling software (ASPEN) worth of 20% of this assignment mark, the calculation of the credit resulted in

0.9 credits towards "process modelling & simulation" (D3). Given the rough correspondence of one credit to ten hours of learning activities, this would imply that at least nine hours of student learning would be devoted to D3 on this occasion.

Given the nature of the project, it is important to note that there might have been details on specific learning activities or assessments that were overlooked or underestimated resulting in mapping not being completely definitive, but nevertheless sufficient to highlight areas of more or less DCE coverage across the examined curricula.

Data Analysis and Visualisation

Once mapping was completed, data was processed and analysed using Python (version 3.12.12) with matplotlib (version 3.10.0) and numpy (version 2.0.2) for visualisation and trend identification per domain and per year. Considering the qualitative nature of the mapping, an "instructional intensity" metric was developed to distinguish between the presence of a tool and the formal pedagogical delivery of a skill.

For the construction of the Hidden Curriculum Graph, a scale range was defined from 0 to 3, starting with absence of evidence (0), moving to evidenced presence of domain activities as curriculum byproducts with no formal assessment (1), as optional modules including assessed learning outcomes (2), and ending with evidenced domain activities in core modules with formally assessed learning outcomes (3). This range was used to quantify domain presence across the curriculum, with one entry per year. For example, for Domain X, if there was at least one module per year that contained formally assessed components relevant to that domain, the value on the graph would be 3 times 4, resulting in 12. Similarly, again for Domain X, if there was a module that required existing knowledge to be utilised but there was not formal assessment for that knowledge, the associated value would be 2. If there were two such modules in the same year, the associated value would still remain as 2.

RESULTS AND DISCUSSION

Analysis of the curriculum mapping across domain and year of study (Figure 1) shows that all domains except Software Engineering Practices (D7) are covered through the chemical engineering curriculum of programmes offered at Sheffield. It is important to point out that this figure evidences the presence of a domain during a year group in at least one module, rather than the coverage of the domain, i.e. how well it is covered. Domain coverage will be discussed later on. For domains D2-5 and D8 there is coverage across multiple years, whereas D1 and D6 are covered only within a particular year through associated modules.

Especially for domains focusing on data literacy and statistics (D2), process modelling and simulation (D3) and data governance, ethics and cybersecurity (D8), presence in multiple year groups is expected, as knowledge contained within these domains (understanding of statistical principles, development of models and unit operation simulations, and an ethical approach to data management) forms the basics of chemical engineering education.

Looking into the areas of the curriculum where DCE was present in core taught content with formal assessment, analysis of the corresponding credit (Figure 2) indicated that DCE is worth credits in Years 1 and 2 with limited domain diversity (D2, D3 and at a very small fraction D5 are covered), but when reaching Years 3 and 4 the domain diversity and credit is higher (D1-D5 are covered). Data literacy and statistics (D2) accounts for the highest number of DCE credits across the curriculum, closely followed by process modelling and simulation (D3). It is worth pointing out that there is a very small credit value (0.027) accounted for D8 in Year 3, as a small part of one assessment component during the Design Project that discusses cybersecurity. Credit calculation took place through assessment briefs where available. According to the Quality Assurance Agency for Higher Education, 1 credit is roughly equal to 10 hours of notional learning. Figure 2 showcases the hours students are expected to engage with DCE within formally assessed elements of the curriculum, for example submission of process simulations using Aspen or other software as part of design projects, or assignments based on coding tools, or engagement with Excel as part of formulating a lab report. As a validation case study, looking at the "Process and Product Design" module students undertake in Year 2 and where they are introduced to the use of commercial simulators like Aspen for the first time, the calculated credit relevant to D3 for that module is 2.3, which corresponds to 23 hours. During this module, there are 14 hours of planned, guided tutorials on the use of Aspen, which implies that the rest 9 hours are based on self-study. This number might be a bit higher, given that there is an assignment and a Design Project Week students need to complete during this module, where the production of an Aspen model for a given process is a formal part of the assessment for both activities.

When looking at credit calculation across all domains per year (Figure 3, top left graph), it is evident that credit attributed to DCE raises per year of study, with a jump from Year 1 to Year 2 and another jump from Year 3 to Year 4.

This is expected, as the way the curriculum is currently structured at Sheffield is to focus on the fundamentals during earlier years of study and introduce more complex concepts and application of concepts later on. It

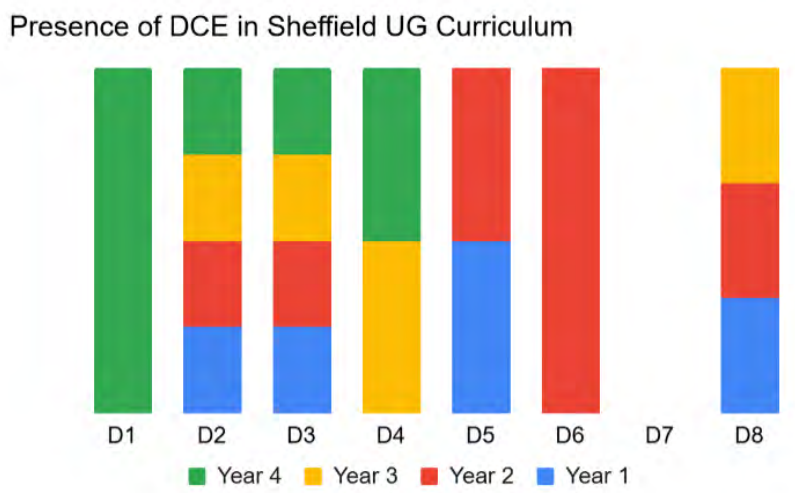

Figure 1. Coverage of DCE domains D1-8 in the Sheffield UG curriculum. D1: programming and computation, D2: data, literacy and statistics, D3: process modelling and simulation, D4: optimisation and PSE, D5: control automation and instrumentation, D6: artificial intelligence, machine learning and digital twins, D7: software engineering practices, D8: data governance, ethics and cybersecurity.

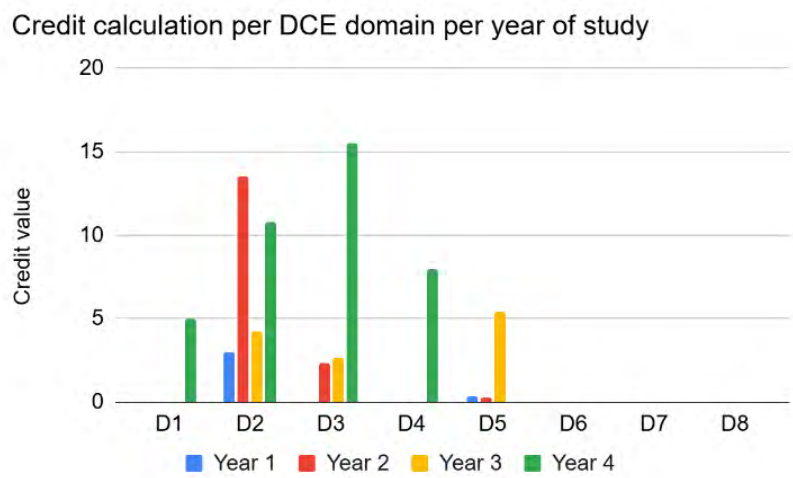

Figure 2. Credit calculation per DCE domain per year.

is interesting to note the slightly reduced credit attributed to digital skills between Year 2 and Year 3. This can be explained by the fact that during Year 3 students undertake the year-long Process Design Project, where they are called to apply their understanding and skills, so the curriculum is deliberately allowing them space to do so rather than overload them with more new skills.

The scaffolded approach towards development of digital skills through the curriculum is even more evident if we look at the credit attributed to DCE cumulatively per year (Figure 3, top right). There is an almost equal step between Years 1, 2 and 3, with a much bigger step towards Year 4. This is mainly attributed to two DCE-heavy core modules available at MEng level of study, "Advanced Process Modelling" and "Simulation and Optimisation". When looking at the cumulative credit attributed per domain (Figure 3 bottom), Data Literacy and Statistics (D2) and Process Modelling and Simulation (D3) are the domains with the highest presence. It is important to mention that the step change between Year 3 and Year 4 is attributed to the Final Year Research project as well, as a big component of its assessment is on data analysis, statistics, and data visualisation which relate to D2. The caveat here is that depending on the specific research area and projects students choose, the exposure to DCE might be even bigger. To illustrate the point better, we conducted a quick analysis of proposed research project

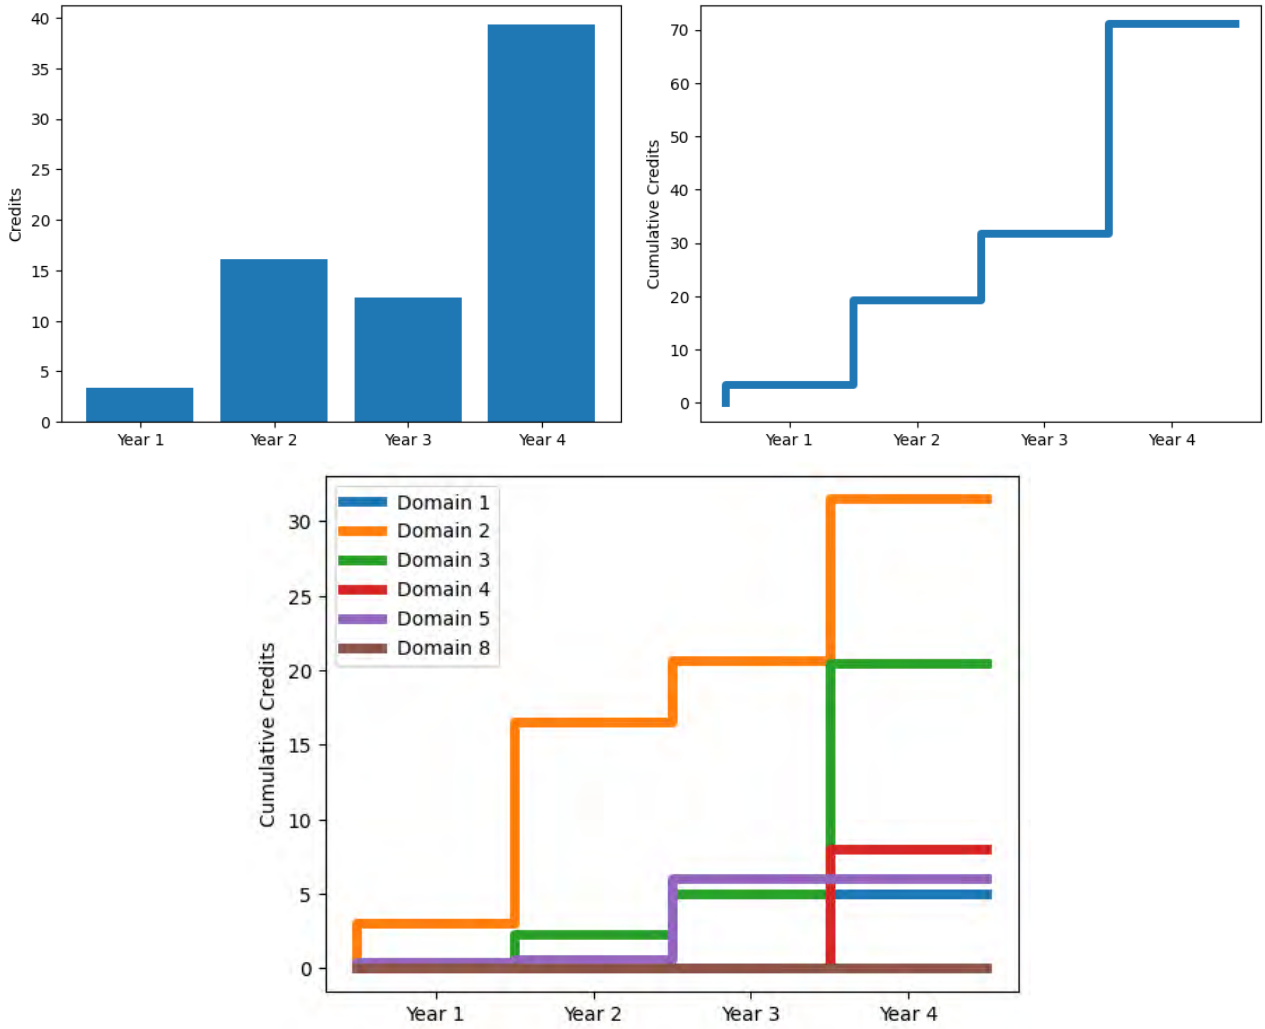

Figure 3. Top left graph: Credits accounted to DCE per Year of study. Top right graph: Cumulative credits accounted to DCE after each Year of study. Bottom graph: Cummulative credits accounted to DCE per domain where credits are atributed, per year of study.

topics over the last couple of years. We assessed the topics in terms of being fully computational (i.e. research and data collection based on simulations) and whether there was explicit mention of elements relating to any of the DCE domains (i.e. need to learn a new analysis software, partially computational research etc). Over 50% of projects offered over two consecutive years were fully computational, and around 70% across offered projects had relevance to DCE, additionally to the expected data analysis and statistics component. This shows that for majority of students, the DCE offering in Year 4 would be even higher compared to what was calculated previously as a baseline. It is very tricky to quantify in terms of credits the additional DCE exposure a student would get during their final year research project, aside of the already accounted component relevant to D2. The only other component where we could see and evaluate DCE coverage during a research project would be the methodology section of the final report students are asked to submit. In this section students describe the methods they used, adapted, or developed towards collecting the data necessary to conduct their project. In the case of computational projects, in the methodology section, students would discuss what software they used, how they developed their models or code, how they decided on their parameters and range of importance, potential optimisation they went through, if/how they ensured reproducibility. Depending on the project topic, there could be particular relevance to particular domains, however, overall, DCE exposure would certainly be higher compared to non-computational-heavy projects, and could result up to

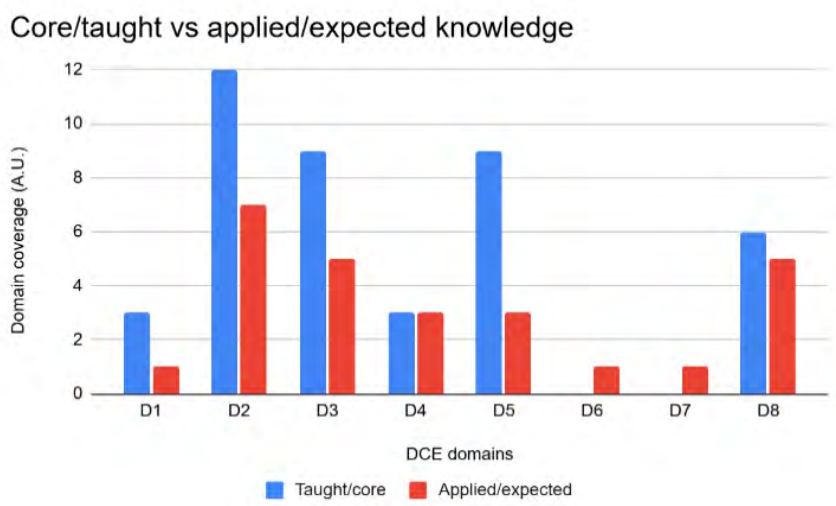

Figure 4. Hidden curriculum of digital chemical engineering: a comparison between content formally taught and assessed (blue) and content expected to be applied outside of formal assessment (red). Domain coverage was calculated as discussed in Methodology section.

10.8 more credits, according to the weight of the methodology section component to the overall mark for this module. This would bring the credit attributed to DCE during Year 4 to over 40 credits in certain cases. As a reminder, a typical year during an undergraduate degree in the UK has 120 credits, so an over-40 credit attributed to DCE illustrates a third of the year being devoted to digital competency development.

If we now look at content that is taught in core modules and assessed with credit-bearing assignments, compared to content or knowledge that is expected to be understood and applied outside of formal assessment, we can see an interesting trend showing (Figure 4). This figure reveals a "hidden curriculum", where we can see presence, elaborate or not, of all domains of digital chemical engineering. Compared to assessed content under credit-bearing activities, the applied or expected knowledge has a small presence in AI/ML and Digital Twins (D6) and Software Engineering Practices (D7). For example, students in Year 2, during the Experimental Investigation module, are introduced to the Diamond Pilot Plant facility, where they get to see unit operations of a tablet manufacturing process, including a digital twin version of the physical plant. There are assessed laboratory activities built around the operation of the pilot plant, however these do not include the operation of the digital twin.

Another example is in Year 4, during the Computational Modelling module, where students are expected to understand and use practices such as version control, although this activity is not formally assessed and does not form a learning outcome of this or other modules. It is interesting to note that the domains with highest level of applied or expected knowledge align, with the exception of D5, with domains where knowledge is delivered as core content and assessed accordingly (D2, D3, D8 in Figure 4). Another interesting point specific to the Engineering Curricula at Sheffield is the coverage of Data Governance, Ethics and Cybersecurity (D8). Students on chemical engineering programmes discuss about ethics, with brief coverage of digital ethics, during a Year 1 module, Principles of Chemical Engineering. Then, they revisit ethics and discuss cybersecurity during Year 3, in the Process Design project, where a small part of the assessment is on the cybersecurity provisions for their hypothetical process design. While these are the only examples of D8 coverage in the chemical engineering curriculum, there are another two spaces where ethics and cybersecurity are discussed. All Year 1 and Year 2 students registered on any engineering programme at the University of Sheffield participate in the compulsory yet non-credit-bearing Project Weeks that take place just before the beginning of the Spring semester. During these weeks, students are grouped together in teams with multidisciplinary membership and are given open-ended briefs around real-life engineering challenges. They are then asked to produce an engineering solution to solve the given challenge, using their skills and knowledge up to this point in their programme and working collaboratively. One aspect they need to discuss while developing their solution is cybersecurity, which relates directly to D8.

REFLECTIONS, CONCLUSION AND FUTURE WORK

This work shows the Digital Chemical Engineering coverage at the chemical engineering programmes offered at Sheffield. Results show that all the eight identified domains are present in the curriculum, although not all of them are covered at the same extend or with equal links to formal assessment and learning outcomes. The review of the curriculum with a focus on digital skills showed that in the first and second years of study students are all introduced to the same curricular digital exposure by design, attending core modules with specific learning objectives and activities around digital competencies. As the course progresses students have the opportunity to engage more with the development of their digital competencies through elective modules, the design project, and ultimately, the choice of their final year research project. The question of whether the curriculum transformation at the University of Sheffield has been successful in addressing the demand for Industry 4.0 is now rephrased to "how can the chemical engineering curriculum at Sheffield ensure that graduates are equipped at a high level across the DCE domains to contribute to Industry 4.0".

Process simulation and modelling, integral parts of modern chemical engineering practice, require and generate data to design, simulate, and optimise chemical processes. These are fundamentally linked to data analysis, statistics, and management. These data-driven processes connect tightly to Digital Twins, integrating real-time data with virtual models for predictive maintenance and operational efficiency. Machine Learning, traditionally the domain of Data and Computer Science, remains underrepresented in the curriculum. Version control is also highly relevant to data management and coursework reporting. Across programming languages with wide use across engineering programmes, python is likely the best programming language for Domain 1, as it is currently the most popular programming language reflected by its extensive applications in AI and ML. In addition, Cybersecurity is a growing requirement for understanding the digital risks to physical infrastructure, such as protecting Industrial Control Systems from cyber threats.

Traditionally, the Chemical Engineering degree has covered a wide breadth of industries and sectors. We now need to ensure that it covers digitalisation in sufficient depth so that chemical engineering graduates are well equipped with digital skills necessary to communicate effectively with computer scientists and software engineers, working together within the Industry 4.0 framework.

Future work would need to focus on a thorough curriculum review, with the intention of perhaps enriching the digital coverage offering in earlier years of study to incorporate exposure to mode digital domains with a more embedded and formally assessed approach. Students then would progress to subsequent years with a better understanding around artificial intelligence and machine learning principles, digital twins, programming, and cybersecurity. One way of ensuring that the taught – and applied- content is current and relevant, would be to continue engagement with employers, alumni, and industry advisory boards, and develop work-related learning activities targeting exposure in specific domains of digital interest.

ACKNOWLEDGEMENTS

The authors would like to acknowledge their colleagues from the School of Chemical Materials and Biological Engineering, for valuable insight and discussions that contributed to our understanding and mapping of DCE in our curriculum.

AUTHOR IDENTIFIERS

Author ORCIDs:
Routoula E: 0000-0001-9416-4265
Bestenlehner J: 0000-0002-0859-5139
Zandi M: 0000-0003-0017-1080

REFERENCES

1. Udugama IA, Bayer C, Baroutian S, Gernaey KV, Yu W, Young BR. Digitalisation in chemical engineering: Industrial needs, academic best practice, and curriculum limitations. Educ. Chem. Eng. 39:94-107 (2022) https://doi.org/10.1016/j.ece.2022.03.003
2. Udugama IA, Atkins M, Bayer C, Carson J, Dikicioglu D, Gernaey KV, Glassey J, Taylor M, Young BR. Digital tools in chemical engineering education: The needs and the desires. Education for Chemical Engineers. Educ. Chem. Eng. 44:63-70 (2023) https://doi.org/10.1016/j.ece.2023.05.002
3. Zandi M, Glassey J, Young BR. Mind the Gap. The Chemical Engineer 976:41-43 (2022) https://www.thechemicalengineer.com/features/mind-the-gap/
4. Bilques Z, Belton D, Campbell GM. Structure and Content of BEng Chemical Engineering Programmes in the UK, Relative to the Frontiers in Chemical Engineering Education Model. Process Integr. Optim. Sustain 7:1003-1030 (2023) https://doi.org/10.1007/s41660-022-00307-6
5. Zandi M, Patwardhan S, Kotta L, Litster J. A New Curriculum to Train Chemical Engineers to Solve 21st Century Grand Challenges. In: The Fifth Annual

Symposium of the United Kingdom & Ireland Engineering Education Research Network, 23-24 Nov (2017), Royal Academy of Engineering, London. Aston University, Birmingham. ISBN: 978-1-85449-443-6.

6. Routoula E, Halliday O, Zandi M. Changing the Way We Teach Using co-creation: Embedding Employability in the Engineering Curriculum through the Students' Perspective. In: UK and Ireland Engineering Education Research Network Conference. (2024) https://doi.org/10.31273/10.31273/9781911675167/1632
7. Routoula E, Brooke C, Zeller J Enhancing student employability with reflection and employer collaboration. In Brooks, J. UK and Ireland Engineering Education Research Network Annual Symposium Proceedings (2025) https://doi.org/10.71535/f5dfe8de-1b80-4192-8573-54d934996dec

© 2026 by the authors. Licensed to PSEcommunity.org and PSE Press. This is an open access article under the creative commons CC-BY-SA licensing terms. Credit must be given to creator and adaptations must be shared under the same terms. See https://creativecommons.org/licenses/by-sa/4.0/

Index

Author Index

Author Index

Aarnoudse, Leontine I.M., **2433**
Abarca, Jose Antonio, **1030**
Abbas, Muhammad Mujtaba, **870**
Abdelhady, Mohamed, **396**
Abildskov, Jens, **1050**
Abo-Ghander, Nabeel S., **1135**
Abushaikha, Ahmad, **1943**
Aceves-Lara, Cesar Arturo, **964**
Adam, Mohammed, **886**
Adams, Thomas A., II, **244**, **802**, **879**, **950**, **2638**
Adeogun, Basit, **1793**
Adjiman, Claire S., **77**, **1153**, **2082**, **2152**
Afanasev, Anna, **2023**
Agunloye, Emmanuel, **1285**
Ahmad, Eblagh, **828**
Ahmed, Abdulhakeem, **1592**
Ahmed, Fahad, **870**
Akashi, Naohiro, **2145**
Aker, Burcu, **37**
Akinola, Toluleke E., **1088**
Akkermans, Simen, **702**, **1102**
Aku, Michael, **1520**
Akundi, Sahithi Srijana, **2425**
Alam, Md Shamsul, **958**
Al-Ansari, Tareq, **1461**, **1943**
Alberto, Biasin, **828**
Alfaro-Ayala, Jorge A., **1119**
Algoufily, Yasser, **2449**
Al-Hammadi, Aisha, **1461**
Alhoshan, Mansour S., **681**, **910**
Almaraz, Omar, **2418**
Almaraz, Sofía De-León, **2183**
Al-Mohannadi, Dhabia, **2**
Alnouri, Sabla Y., **2**
AlNouss, Dr Ahmed, **1461**
Al-Obaidi, Mudhar A., **886**
Alonso-Fariñas, Bernabé, **507**
Alqusair, Omar, **1746**, **1809**
Al-Sakkaf, Aseel, **11**
Al-Salmani, Shaima, **2161**
Al-Sharshani, Ali, **2**
Alshehri, Abdulelah S., **681**, **910**
Alvarado-Morales, Merlin, **727**
Álvarez-Menchero, José A., **2191**
Alves, Ediane S., **2051**
Alves, Rita M. Brito, **1007**
Amaral, Guilherme C., **1950**
Amorim, Ana S., **526**
Anantharaman, Rahul, **419**
Anastasi, Anthony D., **1318**
Andritz, Marion, **404**
Angeli, Panagiota, **108**, **2340**
Anghilante, Régis, **352**
Ansar, Talha, **870**
Antonio, Claudia Gutiérrez, **59**
Appiah-Danquah, Emmanuel, **1476**
Aragón-García, Alejandro, **507**
Araujo, Jefferson D. C., **1398**
Arellano-Garcia, Harvey, **543**, **1908**, **2201**, **2488**
Arellano-García, Harvey, **1476**
Armaou, Antonios, **1445**, **1467**, **1841**, **2578**
Aroniada, Magdalini, **1001**
Asger, Ali, **1916**
Ashraf, Waqar Muhammad, **870**
Asrav, Tuse, **727**
Assen, Niklas von der, **483**
Assis, Bruna Carla G. de, **1175**
Assumpcao, José Matias, **1226**
Austbø, Bjørn, **454**
Authier, Olivier, **1626**
Avila, J. Rafael Alcántara, **1197**
Avilez, Leonardo A. Cáceres, **1007**
Avraamidou, Styliani, **100**, **567**
Ayala-Andreu, Vicent, **846**
Azzaro-Pantel, Catherine, **2183**
Babaei, Farshid, **618**, **2001**
Babaei, Mohammadreza, **2261**, **2594**
Baddam, Snehitha, **2418**
Badmos, Sakiru, **2646**
Badr, Sara, **816**
Baibhav, Vibhu, **205**
Baldea, Michael, **1825**
Barbosa-Póvoa, Ana, **259**
Barhate, Yash, **611**
Barolo, Massimiliano, **625**
Barros, Letícia M. S., **1242**
Batata, Ian B. B., **1249**
Báthori, Arthur-Maximilian, **252**
Batista, Daniel V., **1573**
Bawa, Solomon Gajere, **1520**, **2666**
Becher, Valentin, **2023**
Bechtold, Anna, **2023**
Becker, Gaëtan, **861**
Beisswenger, Lucien, **2652**
Beleli, Yuri S., **1183**
Bempeli, Stefania, **2297**
Benyahia, Brahim, **534**, **791**, **1406**, **1847**, **1856**, **2587**
Bergers, Bart M. A., **2433**
Bernardi, Andrea, **94**
Bernardo, Fernando P., **1292**
Bertram, Sina, **2115**
Bertran, Maria-Ona, **756**
Bestenlehner, J, **2696**
Beykal, Burcu, **2288**
Bezzo, Fabrizio, **625**, **640**, **1601**
Bhandari, Saugat, **2680**
Bhatnagar, Pullah, **1483**
Bhonsale, Satyajeet S., **702**, **1095**, **1102**
Biegler, Lorenz T, **1551**
Biegler, Lorenz T., **306**, **1422**
Bindlish, Rahul, **2091**
Biró, Levente, **1127**
Bisotti, Filippo, **1234**, **1414**, **1431**, **2334**
Blair, Matthew, **672**
Bogataj, Miloš, **1721**
Bogle, I. David L., **958**
Boix, Marianne, **215**, **2689**
Bollas, George M., **2280**, **2288**
Bongartz, Dominik, **492**, **1226**, **1679**, **2638**
Boniakou, Eneri, **2578**
Bonilla-Petriciolet, Adrián, **31**
Borja-Roman, Ronald, **183**
Bornemann, Luka, **369**, **483**
Boroujeni, Saman Naseri, **2082**
Boskabadi, Mohammad Reza, **709**, **733**
Bot, Félix Le, **879**
BOTELHO JUNIOR, Amilton Barbosa, **2234**
Bouchkira, Ilias, **1833**, **1866**
Bournazou, Mariano N. Cruz, **783**
Bournazou, Mariano Nicolas Cruz, **1801**
Boy, Onur C., **1205**
Boztas, Oktay, **1328**, **1483**
Bradner, Dean, **2457**
Brahmbhatt, Parth, **567**
Brandão, Sofia P., **1494**
Braniff, Austin, **1659**, **2425**
Bresciani, Antonio E., **1007**
Brink, Anders, **328**
Brito, Margarida S. C. A., **1494**
Brown, Cameron J., **618**
Brown, Solomon F., **618**, **2001**

Bruschi, Larissa T., **2377**
Braatz, Richard D., **1884**
Bublitz, Saskia, **2201**, **2488**
Bugosen, Sergio, **550**
Burger, Jakob, **77**, **1145**, **1728**
Butté, Alessandro, **692**, **992**
Buzzi, Simona, **1679**
Baabbad, Hassan Khaled Hassan, **1269**
Caballero, José A., **1558**, **1986**, **2191**
Caccavale, Fiammetta, **2600**
Caetano, Ana Helena V., **721**
Caldas, Gustavo Luís Rodrigues, **1306**, **2534**
CALLIL-SOARES, Pedro H., **2303**
Campos, Antonioni B., **1175**
Campos, Mario C. M. M., **2534**
Canning, Griffin A., **1445**
Capasso, Salvatore, **1391**
Caprio, Ulderico Di, **1205**, **1679**
Capron, Bruno D. O., **2534**
Cárdenas, Andrés I., **462**
Carlassare, Enrico, **625**
Carrera-Rodríguez, Marcelino, **267**
CARRILLO LE ROUX, Galo A., **2303**
Carstensen, Peter E., **1353**
Carvalho, Miguel D., **470**
Casagrande, Matteo Lea, **2441**
Cassells, Benny, **1582**
Castel, Christophe, **1626**
Castellar-Freile, Andres, **183**
Castro-Amoedo, Rafael, **934**
Cavenaghi, Marilia G. L., **2377**
Cenci, F., **640**
Cenci, Francesca, **1001**
Chachuat, Benoît, **664**, **1153**
Chahine, Mohamad A., **2051**
Chakrabarti, Dhurjati Prasad, **477**, **2674**
Chakraborty, Sourojeet, **2607**, **2613**
Chalasti, Eleni, **297**
Chandra, Swastik, **2176**
Chang, Hsuan, **2666**
Chanona, Antonio del Rio, **692**, **1551**
Chanona, Antonio del Río, **1558**
Chao, Cong, **108**, **2340**
Charitopoulos, Vassilis M., **198**, **229**, **446**, **602**, **648**, **1793**, **2099**
Chattopadhyay, Sampriti, **342**, **514**
Chaudhary, Kajal, **192**
Chea, John D., **183**
Chen, Bo, **2400**
Chen, Gang, **2400**
Chen, Tao, **1551**
Chen, Xi, **2400**, **2622**
Chen, Zean, **2622**
Chernick, Daniel, **477**
Chibeles-Martins, Nelson, **1262**
Chóez-Guaranda, Ivan, **1476**
Choi, Hyeonrok, **2123**
Christensen, Johanne L., **1298**
Chundawat, Shishir P.S., **1318**
Commenge, Jean-Marc, **1626**, **1687**
Concha, Viktor O. C., **1242**
Constantoudis, Vassilios, **2578**
Contino, Francesco, **424**, **558**, **574**
Cooper, Nathanial J., **68**
Coppitters, Diederik, **424**, **558**, **574**
Cordiner, Joan, **986**, **2001**
Cormos, Ana-Maria, **1127**
Cormos, Calin-Cristian, **252**
Correa-Ibarra, Margarita G., **1119**
Cortez-González, Jazmín, **810**
Costa, Livia Pereira L., **1175**
Cousin, Jean-Matthieu, **1565**
Couto, Lucas Fonseca, **1908**
Cramer, Eike, **1513**
Cristea, Vasile M., **2511**
Cruz, Carolina, **2572**
Cui, Chengtian, **18**
Cui, Jinwen, **1080**
D'Angelo., Antonio, **1277**
Dai, Yiyang, **1967**
Dantas, Beatriz, **2425**
Daoutidis, Prodromos, **289**
Darvariu, Victor-Alexandru, **2099**
Davanzo, Mauro, **625**
Dekhici, Benaissa, **131**, **161**
Delgado, Serena, **1626**
Dias, Madalena M., **1494**
Díaz, Raquel Salcedo, **1558**
Díaz-Sainz, Guillermo, **1030**
Diehl, Fábio C., **1306**
Dietrich, Ingo, **1608**
Dijkstra, Jan Wilco, **2393**
Dikic, Vladimir, **2393**
Dikicioglu, D., **2658**
Djelassi, Hatim, **1617**
Doi, Shunya, **2145**
Dolat, Meshkat, **161**
Domingos, Meire E. G. R., **1328**
Domingos, Meire Ribeiro, **1483**
Dong, Fenglian, **2226**
Dorneanu, Bogdan, **543**, **1476**, **1908**, **2201**, **2488**
DOS SANTOS, Anderson Rapello, **2234**
Drakopoulou, Zoi, **492**
Du, Wenli, **198**
Dua, Vivek, **870**, **958**
Dujany, Arnaud, **1565**
Duong, Samuel, **1737**
Duvigneau, Stefanie, **1043**
Dürr, Robert, **1043**
Eberius, Thomas, **68**
Eckstein, Aline R., **274**
Edwards, Laura, **2674**
El-Halwagi, Mahmoud, **386**
Elizalde-Solis, Octavio, **1539**
Elmisaoui, Sanae, **1833**
Engell, Sebastian, **2370**
Erdmann, Daniel, **2652**
Errico, Massimiliano, **1933**
Errington, Ethan, **2320**
Espen, Henri Tande, **419**
Eydam, Linda, **2527**
Fáber, Rastislav, **2481**
Facco, Pierantonio, **1601**
Fadzil, Noor Fatina Emelin Nor, **927**
Falch, Håvard, **419**
Falk, Laurent, **2058**
Fang, Zoe, **664**
Fayemi, Pierre-Emmanuel, **1565**
Fedor, Adam, **2497**
Feldmann, Kevin, **1298**
Femenía, Rubén Ruiz, **1558**
Fennell, Paul S., **2631**
Ferdoush, Shumaiya, **742**
Fernández, Alberto, **1269**
Ferrari, Adrián, **2386**
Ferreira, Alexandre, **1950**, **2272**, **2466**
Ferreira, Leonardo D., **274**
Filho, Emílio E X. Guimarães, **1249**
Filho, Rubens Maciel, **1242**, **1249**
Florence, Alastair J., **618**
Flórez-Orrego, Daniel A., **205**, **1328**, **1483**
Floros, Stylianos, **1102**
Fontalvo, Javier, **964**
Fortin, Jérôme, **1565**
Fortunato, João, **1219**, **2311**
Fraga, Eric S., **108**, **2340**
Franke, Meik, **1761**
Fraser, Harrison A., **2152**
Fredriksen, Marius, **2551**
Fu, Shilong, **18**
Fujii, Shoma, **413**
Fujiwara, Koichi, **2565**
Funcia, Eduardo S., **1183**

Furdoush, Shumaiya, **751**
Furtner, Lukas, **2029**
Gaborcík, Jakub, **2505**
Galanopoulos, Christos, **1383**
Galatro, Daniela, **2607**, **2613**
Galindo, Amparo, **77**, **2082**, **2152**
Galleguillos, Francisca Pizarro, **1095**
Galvanin, Federico, **108**, **828**, **1080**, **1285**, **1520**, **2340**
Gambetta, Mauro, **2386**
Gamiño-Arroyo, Zeferino, **1119**
Gao, Barnabas, **2680**
Gao, Shang, **791**, **1847**
Gao, Yiwei, **175**
Gargalo, Carina L., **2600**
Gavriilidis, Asterios, **1285**
Gerevini, Giovani, **1306**
Gernaey, Krist V., **721**, **1353**, **1582**, **2600**
Gerogiorgis, Dimitrios I., **236**, **592**, **672**
Gerold, Hilde, **2457**
Ghuys, Nicolas, **424**
Gilardi, Matteo, **1234**, **1414**, **1431**, **2334**
Giraldo, Sergio A. C., **2534**
Glavanin, Federico, **1360**
Glover, Alex, **978**
Gogolides, Evangelos, **2578**
Golder, Rahul, **1925**
Gómez-Castro, Fernando Israel, **25**, **822**, **1933**, **2008**, **2170**
Gong, Zhenzhi, **2226**
González, Diego Izquierdo, **1793**
Gonzalez, Marcial, **742**, **751**
González-Arias, Sandro, **1539**
González-Fernández, Cristina, **1030**
Gopinath, Smitha, **2066**, **2152**
Gornay, Julien, **2051**
Gosálbez, Gonzalo Guillén, **692**
Gounaris, Chrysanthos E., **550**
Grebe, Luca, **1513**
Grénman, Henrik, **1277**, **1391**
Grey, Stuart, **2607**, **2613**
Groening, Paul J., **1713**
Groll, Niklas, **1992**
Große, Michael, **2029**, **2036**
Grossmann, Ignacio E., **514**
Groves, Teddy, **1353**
Grünewald, Marcus, **320**, **334**
Guan, Yunlong, **2226**
Guchin, Ivan, **2386**
Gui, Lingfeng, **1876**
Guido, Sergio I. Martínez, **59**
Guillaume, Denis, **2051**
Guillén-Gosálbez, Gonzalo, **45**
Gumma, Sasidhar, **1754**
Gundersen, Truls, **454**
Guo, Jing, **1884**
Guo, Miao, **776**, **2320**
Gutierrez, Soledad, **2386**
Gutiérrez-Antonio, Claudia, **1933**, **2170**
Guzmán-Martínez, Carlos E., **1933**
Ha, Rachel, **664**
Habicht, Fabian E., **37**
Haikarainen, Carl, **314**
HajiHashemi, MohammadSina, **950**
Halvorsen, Ivar J., **2433**
Hamdan, Malak, **543**
Hamdan, Mustapha, **543**
Hameed, Gul, **1551**
Hao, Huiyu, **2107**
Hao, Zhimian, **18**
Haring, Mark A.M., **2433**
Harjunkoski, Iiro, **52**, **2015**
Haro, Pedro, **507**
Hart, Thomas, **477**
Hartung, Fabian, **1728**
Hasan, M. M. Faruque, **1925**
Hayashi, Yusuke, **816**
Hedengren, John D., **2261**, **2594**
Heinbücher, Markus, **1608**
Hellgardt, Klaus, **2631**
Hemmati, Mohammad, **446**
Hemmati, Sahar, **2594**
Hempfling, Stefan, **1043**
Heng, Jerry Y. Y., **631**
Hernández, Juan Gabriel Segovia, **1374**
Hernández, Salvador, **59**, **810**, **1197**, **1933**
Hernández-Aguirre, Arturo, **810**
Hernández-Camacho, Nereyda V., **2008**
Herrera-Ruiz, Juan Federico, **964**
Hesketh, Robert, **2680**
Higler, Arnoud, **378**
Hildebrandt, Diane, **1318**
Hildebrandt, Sophie, **1761**
Hill, Matthew A.J., **592**
Hillaby, Darrick, **2058**
Hillestad, Magne, **454**
Hiromori, Kousuke, **125**
Hochhaus, Thorben, **320**
Holtze, Christian, **77**
Hong, Yee Hung, **1666**
Houette, Thibaut, **1565**
Hren, David T., **1064**
Hu, Jianli, **2418**
Huang, Shengyuan, **972**
Huelmo, Darío, **2386**
Hussain, Jawad, **870**
Huusom, Jakob K., **714**, **1050**, **1298**
Iakovou, Eleftherios, **386**, **396**
Iannello, Stefano, **229**
Iqbal, Imam M., **1058**
Irabien, Ángel, **1030**
Irhamna, Adrian R., **2280**, **2288**
Isella, Andrea, **352**, **2045**, **2441**
Iseri, Funda, **386**
Iseri, Halil, **386**
Iwabuchi, Junshin, **816**
Jackson, George, **77**, **2082**, **2152**
Jafari, Mitra, **1908**
James, Isaac N., **2674**
Jan, Nabil Magbool, **1754**
Jeanmart, Hervé, **424**
Jeffries, Emma, **742**
Jiao, Jianfeng, **1634**
Jinsong, Zhao, **1666**
Johnn, Syu-Ning, **2099**
Joshi, Geeta, **378**
Juarez-García, Maricruz, **1455**
Jungen, Daniel, **1617**
Jäschke, Johannes, **2551**, **2558**
Jørgensen, John B., **1353**, **1422**
Jaafar, Nor Syuriaty, **927**
Kager, Julian, **714**, **721**, **1582**
Kailasanathan, Rajiv, **709**, **733**
Kaitano, Caroline, **838**
Kalamaris, Nicholas N., **282**
Kaltschmitt, Martin, **369**
Kamal, Jonaed Bin Mustafa, **2488**
Kanematsu, Yuichiro, **125**, **169**, **175**
Kang, Hao, **1022**
Kapadi, Mangesh, **2355**
Kapavarapu, Pratham, **702**
Karagoz, Ali Tarik, **1746**
Karim, Hamza, **861**
Karsten, Jonas M., **783**
Kasera, Nitya, **1592**
Katsoulas, Konstantinos, **108**, **2340**
Kawajiri, Yoshiaki, **2140**
Kay, Harry, **855**, **1529**
Kazi, Saif R, **514**
Kchaou, Mahdi, **424**, **558**, **574**
Khan, Faisal I., **2425**
Khatu, Aditya S, **342**
Khodr, Rim, **861**
Khouri, Nadia G., **1242**, **1249**

Kienle, Achim, **500**, **1043**
Kikuchi, Yasunori, **117**, **125**, **169**, **175**, **413**, **432**, **1212**
Kim, Seong-il, **2123**
Kim, Youhyun, **2252**
Kinau, Siv, **2176**
King, Stephen T., **1445**
Kiss, Anton A., **2393**, **2511**
Kloft, Marius, **1728**
Knueven, Bernard, **1916**
Koch, Amy, **1650**, **1786**, **1817**
Kock, Tobias, **1817**
Kockmann, Norbert, **1891**
Kok, Rudolph, **1043**
Kokkoris, George, **2578**
Kontoravdi, Cleo, **94**, **992**
Kopra, Kaapo, **52**
Kortlever, Ruud, **18**
Kowalski, Adam, **903**
Kozerke, Kilian, **68**
Kravanja, Zdravko, **1721**
Krestnikova, Alexandra, **45**
Krishna, Vakil Vamsi, **2355**
Krishnadoss, Rajalakshmi, **879**
Kronqvist, Jan, **2209**
Krühne, Ulrich, **1353**, **2600**
Kumaraswamy, Archana, **2558**
Kurucz, Július, **2497**
Kyprianidis, Konstantinos, **558**, **2261**
Kyriakidis, Loukas, **2201**, **2488**
Labarta, Juan A., **846**, **1986**
Lagarmilla, Milena, **2386**
Laird, Carl D., **550**, **1643**, **1916**, **2209**
Lamm, Sascha, **1608**
Lane, Amanda, **903**
Lange, Christoph, **1801**
Lange, Hannes, **2029**, **2036**, **2652**
Lara-Montaño, Oscar Daniel, **822**, **2008**, **2170**
Lasserre, Alberto A. Aguilar, **2183**
Laterre, Antoine, **558**
Latifi, Abderrazak, **1833**
Lauinger, Dirk, **2243**
Leblebici, Mumin Enis, **655**, **1160**, **1205**, **1713**
Lee, Jaewook, **2123**, **2320**
Lee, Lauren Ye Seol, **919**, **1876**
Lee, Won Bo, **919**, **1876**
Lee, Ye Seol, **1520**
Lee, Yongkyu, **919**
Lefrán, Laureen E. Hernández, **1007**
Leipold, Johannes, **500**
Lemperle, Marc, **1582**
Léonard, Grégoire, **360**, **404**
Leonzio, Grazia, **943**
Leppla, Madeline, **1643**
Lewin, Daniel R., **2631**, **2638**
Li, Jiaming, **1974**
Li, Jie, **440**, **1634**, **1746**, **1809**, **1974**
Liang, Jingkang, **1769**
Liang, Xiaodong, **1050**
Lima, Fernando A. R. D., **655**, **1175**
Liñán, David A., **2209**
Liporace, Fabio S., **1175**
Lipow, Adrian W., **1617**
Li-Puma, Gianluca, **534**
Liu, Mengqi, **1696**
Liu, Yuanxing, **2425**
Llovell, Fèlix, **1269**
Lohani, Ananya, **2497**
Lombo, Diego A. Romero, **895**
Longobardi, Federica, **352**
Lopes, José C., **1494**
Lopez, Jonathan Jair Cuevas, **2183**
López-Ramírez, Varinia, **810**
Lorenz, Julius, **2527**
Loubière, Valentin, **1626**
Lu, Dongyun, **1974**
Lubušký, Karol, **2481**, **2505**
Luca, Nardi, **828**
Lucia, Sergio, **2457**
Luo, Chunlin, **11**
Luo, Guangsheng, **1696**
Lutz, Malte, **37**
Ma, Yiming, **1847**, **1856**
Ma, Yingjie, **1884**
Macchi, Alessandra, **352**
Maciel, Maria R. Wolf, **1249**
Maia, Bruna L. V., **1366**
Maier, Matthias, **244**
Majozi, Thokozani, **838**
Malekshahian, M., **2658**
Manaf, Norhuda Abdul, **927**
Manca, Davide, **352**, **2045**, **2441**
Mangold, Hannah, **77**
Mangra, Arun, **477**
Mansour, Fatima, **2**
Mansouri, Seyed Soheil, **709**, **733**, **2638**
Maravelias, Christos T., **282**
Marcato, Francesco, **934**
Marchan, Gabriela C. Theis, **1111**
Maréchal, François, **205**, **1328**, **1483**
Marella, T., **640**
Margni, Manuele, **1483**
Marias, Fréderic, **1398**
Markl, Daniel, **618**
Markowitsch, Christoph, **404**
Martín, Mariano, **31**, **2008**
Martinez, Andrés Piña, **2058**
Martínez, Carlos E. Guzmán, **59**
Martínez, Ernesto C., **783**, **1801**
Martínez-Guido, Sergio Iván, **2170**
MARTINEZ-TRIANA, Alvaro, **2074**
Maschmeier, Maximilian, **1513**
Masuku, Cornelius M., **1318**
Materazzi, Massimiliano, **229**
Mathew, Kenneth, **2066**
Matos, Henrique A., **526**, **1219**, **1705**, **2311**
Matsoukas, Themis, **1445**
Matsunami, Kensaku, **611**
Mavlanova, Dilafruz, **215**
McKenzie, Kieran, **1958**
Meinders, Marcel, **1073**
Mekidiche-Martínez, Zinet, **1986**
Melo, Elizabeth V., **2518**
Mercado, Victor S. V., **1345**
Mercangöz, Mehmet, **692**, **1737**, **2449**
Mercer, Samuel, **1825**
Merlo, Giovanni, **2433**
Mesa, M. Granollers, **2658**
Mhamdi, Adel, **1866**
Micari, Marina, **2297**
Michalopoulou, Foteini, **2449**
Mihály, Norbert-Botond, **1127**, **2511**
Minceva, Mirjana, **77**
Misra, Shamik, **2355**
Mitraki, Rafailia, **360**, **404**
Mitsos, Alexander, **492**, **1617**
Moioli, Emanuele, **895**
Mondal, Biswarup, **500**
Monge, Maria A., **1262**
Montastruc, Ludovic, **2689**
Montiel, Miguel A., **846**
Montiel, Vicente, **846**
Moraes, Marcellus G. F. de, **655**
Morales-Rodriguez, Ricardo, **2327**
Moreno, Silvia, **507**
Moreno-Palancas, Isabela Fons, **1558**
Morisasa, Mizuki, **816**
Morris, Monica Rodriguez, **183**
Mortadi, Marwa, **314**
Morud, John C., **2218**
Mosqueda-Huerta, Zaira J., **822**
Motterle, Riccardo, **625**
Mouchtouris, Sotiris, **2578**

Muhieddine, M., **640**
Mujtaba, I. M., **886**
Muñoz, Melanie Coronel, **1374**
Muraleedharan, Aparna, **1728**
Murali, Rohit, **131**
Murrieta-Dueñas, Rodolfo, **810**
Mustafa Kamal, Jonaed Bin, **2201**
Mustika, Dea Nilam, **236**
Muumbo, Andrew, **2066**
Myrstad, Rune, **2334**
Mädler, Jonathan, **1650**
Na, Jonggeol, **1876**, **2243**, **2252**
Nagy, Zoltan K., **611**, **625**, **742**, **769**
Narayanan, Harini, **992**
Narciso, Diogo A. C., **1219**, **2311**
Nascimento, Claudio A. Oller do, **1007**
Negny, Stephane, **215**
Nemet, Andreja, **1064**
Nemoto, Kosuke, **816**
Neto, José R. Torraca, **2534**
Neto, Wilson Sousa Mercês, **1366**, **1908**
Neubauer, Peter, **1801**
NEVEUX, Thibaut, **2074**
Ni, Franchine, **1626**
Nice, Dylan, **2091**
Nie, Yaling, **1974**
Nielsen, Lars K., **1353**
Nijhuis, Wijtze, **1073**
Nogueira, Idelfonso B. R., **655**, **1160**, **1205**, **1306**, **1439**, **1713**, **1950**, **2234**, **2272**
Nordhus, Inga S., **655**
Noriler, Dirceu, **1345**, **1366**
Nunes, João, **2466**
Núñez-López, Jesús Manuel, **1455**
Oever, Anne van den, **424**
Ogunnoiki, Adedayo, **2666**
Ohara, Satoshi, **117**, **1212**
Oikonomou, Katerina, **2578**
Olasinde, Malik, **108**, **2340**
Oliveira, Felipe de, **1073**
Oliveira, Ingrid A. de, **274**
Oliveira, Manuel Ruivo de, **934**
Oliveira, Miguel Castro, **1705**
Oliveira, Rita Castro, **1705**
Oluleye, Gbemi, **297**
Ommen, Ruud van, **18**
Onyshkevych, Sophia, **1383**
Otitoju, Olajide, **972**
Otterlei, Camilla, **2334**
Ouaret, Rachid, **215**
Overgaard, Tobias, **756**
Pably, Philipp, **714**
Paepe, Ward De, **558**
Paiva, José L., **1183**
Pakarinen, Anna, **328**
Palanki, Srinivas, **11**, **2418**
Paliwal, Shubhani, **2082**
Palys, Matthew J., **289**
Pamperin, Julian, **1650**, **2036**
Pankajakshan, Arun, **108**, **2340**
Papadakis, Emmanouil, **756**
Papageorgiou, Lazaros G., **1001**
Papathanasiou, Maria M., **94**, **297**, **477**, **631**, **756**, **978**, **992**, **2449**
Pappagallo, M. Andrea, **895**
Parbat, Shreyas, **1336**
Pardal, Gonçalo V.L., **1292**
Paredes, Rodrigo, **2406**
Paschalidis, Leandros, **2176**
Pasieka, Engelbert, **2370**
Pastoors, Johannes, **1513**
Paulen, Radoslav, **2481**, **2497**, **2505**
Pedro, Miguel A., **526**
Pelz, Peter F., **1608**
Pennington, Oliver, **2572**
Peralta, Camilo, **1030**
Peralta-Hernández, Juan Manuel, **822**
Pessina, Daniele, **631**
Pessoa, Fernando, **259**
Pham, Thien An, **2680**
Picón-Núñez, Martín, **2327**
Pierce, Jayden A., **742**
Pilario, Karl Ezra, **2542**
Piña-Martinez, Andrés, **1687**
Pini, Ronny, **297**, **978**
Pinto-Varela, Tânia, **1262**
Pires, José M., **2311**
Pistikopoulos, Efstratios N., **386**, **396**, **583**, **2091**, **2425**
Poker, Lucas A.T., **2393**
Pollock, M., **2658**
Ponce-Ortega, José María, **2327**
Ponce-Rocha, Jaime David, **2327**
Pontes, Karen Valverde, **1908**
Portha, Jean-François, **2058**
Pozo, Carlos, **1269**
Prado-Rubio, Oscar Andrés, **964**
Prodan, Miruna, **2511**
Prokopou, Georgia Ioanna, **492**
Prommesberger, Christian, **2023**
Puccini, Monica, **222**
Putra, I Gede Eka Perdana, **714**
Paap, Ulrike, **2023**
Qiu, Tong, **1014**, **1696**, **2226**
Qyyum, Muhammad Abdul, **2161**
Rabell, Valeria Caltzontzin, **59**
Rackauckas, Christopher V., **1439**
Rahman, Fadhli Hadana, **1345**, **1366**
Rahman, Md Mizanur, **2418**
Rajanala, Greeshmanth, **1336**
Rajaraman, Narayanan, **378**
Ramadhan, Rafif S., **1513**
Ramin, Pedram, **1582**
Ramírez, Eduardo Sánchez, **1374**
Ramírez, Juan José Quiroz, **1374**
Ramírez-Corona, Nelly, **25**
Ramírez-Jiménez, Edgar, **1539**
Ramírez-Márquez, César, **2327**
Ramirez-Minguela, Jose de J., **1119**
Ramos, Victor H. S., **1249**
Ramprasad, Tejeswi, **378**
Rangel-López, Alden Paul, **1455**
Ravikumar, Praveen, **1318**
REBELLO, Carine Menezes, **2234**
Reis, Marco S., **1573**, **2406**
Reklaitis, Gintaras V., **742**, **751**, **769**
Ribeiro, Ana M., **1950**, **2272**, **2466**
Ribeiro, Daniel Wenck, **2091**
Ricardez-Sandoval, Luis, **733**
Riese, Julia, **320**, **334**
Riezzo, Luca, **855**
Rio-Chanona, Ehecatl A. del, **2008**
Rios, Amarelys, **2680**
Rioux, Robert M., **1445**
Robins, David, **618**, **2001**
Robles-Rodriguez, Carlos Eduardo, **964**
ROCHA AZEVEDO, Antonio, **2074**
Rodrigues, Alírio E., **2272**
Rodrigues, Alírio E. , **2466**
Rodrigues, Bruno C. L., **1713**
Rodrigues, Diogo, **1950**, **2272**, **2466**
Rodríguez-Fragoso, Martín, **1539**
Rogers, Alexander, **855**, **903**, **1529**
Romagnoli, Jose A., **1111**, **2473**
Romero-Izquierdo, Araceli G., **1933**
Rosas, Brenda Huerta, **1374**
Ross, Andrew, **477**
Rossi, Francesco, **769**
Roussanaly, Simon, **244**
Routoula, E, **2696**
Routoula, E., **2658**
Roux, Galo Le, **1183**, **1776**
Roxas, Ricardo M., II, **2542**
Roy, Bimol Nath, **1925**
Ruiz, Julián Cabrera, **1197**

Ruiz-Femenia, Rubén, **2191**
Ruiz-Mercado, Gerardo J., **183**
Russo, Vincenzo, **1277**, **1391**
Saccardo, Alberto, **664**
Sadeek, Sherard, **477**
Salcedo-Díaz, Raquel, **2191**
Salehian, Mohammad, **618**
Salgado-Rodríguez, Tania G., **25**
Salihuddin, Radin Suhaib, **1345**, **1366**
Salman, Muhammad, **360**
Salucci, Emiliano, **1277**
Sampathirao, Gnaneshwar, **1754**
Samuel, Stricker, **692**
Sánchez-Ramírez, Eduardo, **267**, **1455**
Sand, Guido, **1761**
Sandham, Finlay M., **2066**
Santana, Leonardo, **259**
Santana, Vinicius V., **1439**
Santecchia, Alessio, **934**
Santos, Esther, **1030**
Santos, Lucas Francisco dos, **692**
Santos, Moises T. dos, **2377**
Santos, Ricardo J., **1494**
Sanz, Carlos, **31**, **1776**
Sardeing, Rodolphe, **861**
Saric, Marija, **2393**
Sarkis, Miriam, **94**, **583**
Saury, Alix, **1565**
Savitskaya, Kristina, **2091**
Sawaly, Razan, **1943**
Saxén, Henrik, **314**, **1277**
Scheffer, Jessica V., **1626**
Schmidt, Stefan, **2334**
Schmitz, Mathias, **1891**
SCHNEIDER, René P., **2303**
Schnurr, Jonas, **2115**
Schoerner, Markus, **77**
Schulze-Netzer, Corinna, **950**
Schytt, Marcus J., **1422**
Schöneberger, Jan C., **37**
Secchi, Argimiro R., **655**, **1175**, **1306**, **2518**, **2534**
Seckler, Marcelo M., **1183**
Sefcik, Jan, **2152**
Segovia-Hernández, Juan Gabriel, **267**, **810**, **1455**
Seibold, Gerd, **714**
SEN GUPTA, Soumyajit, **192**
Sendlhofer, Severin, **404**
Sermoud, Vítor M., **274**
Seth, Gaurav, **2082**
Shah, Nilay, **927**, **943**
Shahab, Mohammad, **742**
Shaikh, Abdul Aziz, **1461**
Shang, Qichen, **1696**
Shang, Zhongtao, **1967**
Shao, Zhijiang, **681**
Sharma, Sudip, **802**
Sheng, Lin, **1696**
Shin, Daeun, **1876**
Shin, Sungho, **244**, **2243**
Short, Michael, **131**, **161**, **1551**, **2131**
Silva, Beatriz C. da, **2272**
Silva, Jean F. Leal, **470**, **1242**, **1345**, **1366**
Silva, Jefferson Santos da, **1908**
Silva, Wilson R. Leal da, **1601**
Silvestri, Francesco, **934**
Sin, Gürkan, **727**, **1769**, **1992**
Singh, Harmandeep, **378**
Sinha, Sarthak, **2066**
Sisson, Laura, **1713**
Sivaram, Abhishek, **709**
Skiborowski, Mirko, **1145**, **2115**, **2176**
Skogestad, Sigurd, **2433**
Slater, C. Stewart, **2680**
Sochard, Sabine, **306**
Sochard-Reneaume, Sabine, **1398**
Solim, Sabah, **2**
Sommer, Sissy, **1608**
Somoza-Tornos, Ana, **18**
Song, Geunseo, **2243**
Song, Kaicheng, **2622**
Song, Xiaowei, **1974**
Soria, Víctor A., **462**
Sotowa, Ken-ichiro, **2145**
Souza, Mauricio B. de, Jr., **655**, **1175**
Souza, Maurício B. de, Jr., **1306**, **2518**
Spencer, George, **992**
Springub, Bianca, **1383**
Srinivasan, Rajagopalan, **378**
Standish, Riccardo, **77**
Stefan, Mihnea, **1601**
Stevens, Niels, **1226**
Stinchfield, Georgia, **1643**, **2209**
Stock, Jerome, **558**
Stocks, Stuart, **1582**
Straus, Julian, **2218**
Sugiyama, Hirokazu, **816**
Sun, Ruitao, **440**
Surmi, Amiza Bt, **1345**, **1366**
Suzuki, Kensuke, **2140**
Swartz, Christopher L. E., **1958**
Tabata, Hiro, **125**, **1212**
Takahashi, Nobuhide, **117**
Tanaka, Ryosuke, **2565**
Tao, Yi, **1014**
Tao, Zhineng, **2226**
Teah, Heng Yi, **125**, **169**, **175**
Tecl, Sebastian, **1608**
Terajima, Yoshifumi, **1212**
Territo, Kyle, **1111**, **2473**
Theisen, Lukas S., **1298**
Thomas, Sydney, **2689**
Thompson, P., **640**
Tian, Yuhe, **1659**, **2091**, **2425**
Tinker, Corbin, **2680**
Tirez, Arne, **1226**
Tizon, Zhane Ann, **2425**
Tognotti, Leonardo, **222**
Tomba, Emanuele, **625**
Tonomura, Osamu, **2145**
Torres, Ana I, **342**
Torres, Ana I., **462**, **514**, **1592**
Torres-Ayala, Antonio, **267**
Tovar, Laura Plazas, **1345**
Tran, Norman, **1643**
Trinh, Trung, **1234**, **2334**
Tsirikoglou, Panagiotis, **558**
Tsochatzidi, Artemis, **1001**
Tula, Anjan Kumar, **681**, **910**, **2622**
Ufomba, Adaugo, **742**
Untrau, Alix, **306**
Urade, Vikrant, **378**
Urbas, Leon, **1058**, **1336**, **1650**, **1786**, **1817**, **2029**, **2036**, **2527**, **2652**
Uribe-Ramirez, Agustin R., **1119**
Vacano, Bernhard von, **77**
Vaccari, Marco, **222**
Valencia, Tomas, **550**
Vallerio, Mattia, **895**
Van Impe, Jan F.M., **702**, **1095**, **1102**
Vantaggiato, Elettra, **2218**
Vasconcelos, João Pedro B., **470**
Vasileiadis, Tzannis, **2578**
Vega-Zambrano, Consuelo Del Pilar, **648**
Velázquez, Josué J. Herrera, **1197**
Venkateshwarlu, Akepogu, **534**
Vergara-Zambrano, Javiera, **100**
Verleysen, Kevin, **558**
Verma, Pankaj, **2355**
Vermeire, Florence, **1679**, **1899**
Viedt, Isabell, **1058**, **1336**, **1786**, **1817**, **2029**, **2652**
Viedth, Isabel, **1650**
Vilarrasa-Garcia, Enrique, **1183**

Villanueva-Perales, Ángel L., **507**
Villazón-León, Vanessa, **31**
Villumsen, Søren, **1050**
Vinestock, Tom, **776**
Vittas, Nikolaos I., **1467**
Vouros, Stavros, **2261**
Vyas, Javal, **1737**
Vaara, Santeri, **2015**
Wang, Bohong, **2107**
Wang, Han, **2400**
Wang, Kai, **1696**
Wang, Meihong, **972**, **1088**
Wang, Yanhu, **1841**
Wang, Yifan, **369**, **483**
Wang, Yushu, **198**
Wang, Yuxin, **11**
Ward, Keeran, **477**
Watanabe, Yukito, **125**
Watson, Jean-Paul, **550**
Weber, Janis, **1891**
Wei, Zhi, **2226**
Weigold, Matthias, **558**
Werf, Lea van der, **2689**
Wesley, Robin, **986**
Wilson, Adam, **2572**
Wirnsperger, Claus, **692**, **992**
Wittgens, Bernd, **1234**, **1414**, **1431**, **2334**
Wloch, Johannes, **334**
Wolf, Alexander B., **1145**
Würpel, William, **37**
Xiao, Xin, **1974**
Xu, Shuo, **1503**
Xu, Wenle, **1696**
Xu, Youquan, **681**
Xue, Yao, **1841**, **2578**
Yajima, Tomoyuki, **2140**
Yamaki, Ayumi, **169**, **432**
Yamaki, Shu, **169**
Yang, Meng-Hua, **769**
Yang, Minglei, **198**
Yang, Won, **2123**
Yang, Yuxuan, **1360**
Yenkie, Kirti M., **183**, **2680**
Yewale, Ashish, **2587**
Yin, Jian, **77**
Yin, Kai, **1022**
Yin, Yuhui, **602**, **1793**
Yoshiyama, Yuki, **816**
You, Fengqi, **1659**
Young, André F., **274**
Yuan, Xuming, **1406**, **1856**
Yuan, Zhihong, **1967**
Zacharias, Jörg, **2023**
Zadeh, M. Mohammad, **2658**
Zandi, M, **2696**
Zandi, M., **2658**
ZARLENGA, Valentin, **2074**
Železnik, Chiara, **1721**
Zeng, Yujiao, **1974**
Zerwas, Alexander A., **1366**
Zhang, Dongda, **855**, **903**, **1529**, **2572**
Zhang, Lifeng, **1153**, **1967**
Zhang, Ruosi, **2131**
Zhang, Shuai, **910**
Zhang, Xinle, **742**, **751**
Zhang, Yao, **972**
Zhang, Yibo, **1974**
Zhang, Zixuan, **1974**
Zhao, Fei, **2400**
Zhao, Jinsong, **1022**, **1503**
Zhao, Yipei, **986**
Zhou, Guanwei, **314**
Zhu, Baikai, **1737**
Zhu, Lingyu, **2622**
Zhu, Min, **1974**
Zhu, Pengcheng, **2400**
Zillmer, Ruediger, **903**
Zingler, Aron, **1617**
Zondervan, Edwin, **1073**, **1761**, **2638**
Zormpa, Chloi, **2578**
Zulkefal, Muhammad, **454**

Index

Keyword Index

Keyword Index

3D geometry, **1876**
4E analysis, **2327**
AAS, **1817**
absorption, **1626**
AC OPF, **602**
AC optimal power flow, **2243**
acausal modeling, **1439**
acoustic cavitation, **1467**
active constraint regions, **2551**
active learning, **108**, **1573**
active learning approaches, **1573**
activity coefficients, **1153**
adaptive modelling, **1601**
adaptive sampling, **756**, **1974**
additional information, **2527**
adsorption, **297**, **306**, **413**, **802**, **1183**, **1422**, **1950**, **2140**, **2272**
agent-based modeling, **378**
agrochemical formulation, **986**
agro-industrial symbiosis, **1212**
AI/ML,
air separation, **1958**
alcoholic fermentation, **592**
algorithms, **934**, **992**, **1001**, **1014**, **1262**, **1551**, **1634**, **1643**, **1761**, **1801**, **2565**
alkaline electrolyzer, **2201**
alkaline water electrolysis, **534**
alkylation, **2481**
alternating minimization, **1754**
alternative fuels, **161**, **289**, **1422**, **2218**
aluminium cold rolling, **2123**
amine-functionalised sorbent, **1088**
ammonia, **386**, **1219**
ammonia electrolysis, **1007**
anaerobic digestion, **131**, **2131**
anomaly detection, **1728**
antifragility, **558**
apple pomace, **31**
apples, **31**
artificial intelligence, **446**, **810**, **934**, **1058**, **1503**, **1634**, **1666**, **1746**, **1769**, **1801**, **1809**, **1841**, **1899**, **1925**, **2074**, **2511**, **2518**, **2542**, **2600**, **2613**, **2638**
artificial neural network, **592**, **822**, **1205**, **1285**
Aspen Dynamics, **2393**, **2418**, **2511**
Aspen Plus, **45**, **404**, **526**, **927**, **943**, **1197**, **1219**, **1234**, **1249**, **1256**, **1626**, **2051**, **2511**
Aspen Plus simulation, **1269**
augmented intelligence, **855**, **1529**
automated construction, **1022**
automation, **2666**
autonomous experimentation, **2252**
autothermal operation, **117**
aviation decarbonization, **378**
bagasse utilization, **1212**
batch process, **192**, **776**, **1467**, **1666**, **2572**
batch process scheduling, **2099**
batch systems, **334**, **2209**
batteries, **52**, **222**, **1242**
battery thermal management, **1119**
Bayesian network, **1022**
Bayesian neural networks, **992**
Bayesian optimisation, **692**, **986**, **1573**
Bayesian optimization, **1520**, **1679**, **1967**
BECCS, **470**
behavioral cloning, **1558**
benchmark, **2600**
benchmarking platform, **2252**
bias correction, **910**
big data, **934**, **2406**
biochemical reaction kinetics, **855**
biodiesel, **25**, **236**
biodiesel production, **1866**
bioenergy, **131**, **175**
biofuels, **161**, **964**, **1249**, **2362**
biogas, **454**, **470**, **1036**, **1414**
biogas upgrade, **1431**
biogas upgrading, **2327**
biomass, **45**, **125**, **507**, **776**, **810**, **1318**, **1374**, **1398**, **1539**, **2131**, **2386**
biomethanation, **1414**
biomethane, **470**, **1414**
biomethanol, **454**
bioprocess development, **783**
bioprocess modelling, **1582**
bioprocesses, **692**
biorefinery, **59**, **205**, **1374**, **1455**, **2377**
biosystems, **709**, **714**, **733**, **776**, **816**, **1318**, **1565**
black-box, **1974**
black-box optimization, **2252**
blockchain, **783**
bloom taxonomy, **2607**
BO-IPOPT, **2488**
bound tightening, **602**
brazed aluminum heat exchanger, **861**
Brazil case study, **259**
brewery, **1483**
butanol, **964**
cake filtration, **1073**
calibration curve, **1205**
capacitive deionization, **1360**
capacity expansion planning, **446**, **550**
cape, **117**, **1713**
carbon capture and storage, **244**
carbon capture utilization and storage, **1269**
carbon circularity, **1414**
carbon dioxide, **274**, **1030**, **1345**
carbon dioxide allocation, **1943**
carbon dioxide capture, **297**, **802**, **927**, **943**, **972**, **1127**, **1183**, **1328**, **1551**, **1626**, **2191**, **2297**
carbon dioxide emissions, **314**
carbon dioxide gas pipelines, **2558**
carbon dioxide purity, **1943**
carbon dioxide sequestration, **18**
carbon dioxide utilization, **1391**
carbon dioxide-to-methanol, **500**
CasADi, **2466**
catalysis, **1030**, **1445**, **1551**
catalyst deactivation, **1080**
catalyst optimization, **1866**
catechol, **1455**
catechol derivatives, **1455**
ccs, **198**
ccus, **1943**, **2161**
cement decarbonization, **198**
cement industry, **526**, **1601**
CGAN, **2131**
chemical absorption, **972**
chemical heat pump, **879**
chemical language models, **903**
chemical looping, **1127**
chemical reaction engineering, **895**
chemometrics, **1205**
chromatography, **992**
circular economy, **45**, **52**, **68**, **100**, **229**
circular water system, **2**

clean energy systems, **1866**
climate change, **1269**
coarse-graining, **1014**
COCO-COFE, **1234**
COFS, **1754**
cognitive digital thread, **783**
combustion heat and power, **1431**
community detection, **2406**
compartment models, **640**
computational fluid dynamics, **328**, **534**, **543**, **640**, **1030**, **1119**, **1292**, **1494**
computational tools, **2646**
computer vision, **1696**
computer-aided process engineering, **1212**
computer-aided retrosynthesis, **791**
concentrated thermal solar, **543**
concept drift, **870**
conditional variational autoencoder, **648**
continuous distillation, **1728**
cost analysis, **1582**
critical metals, **108**, **2340**
critical raw material, **68**
cross-sectoral integration, **1483**
cryogenic distillation, **1345**
cryogenic fuel tank, **328**
cryogenics, **861**
crystallisation, **631**
crystallization, **625**, **655**, **672**, **1050**, **2082**
curriculum, **2658**, **2696**
curriculum revamp, **2613**
cutting planes, **602**
cybernetic model, **1102**
cyber-physical system, **2425**
cyclic adsorption processes, **978**
cyclic steady state, **2449**
DAE systems, **1776**
dairy industry, **1298**
data augmentation, **1111**
data centers, **396**
data centres, **446**
data drift, **870**
data extraction, **2622**
data intelligence, **855**
data preprocessing, **681**
data reconciliation, **2505**
data-driven, **37**, **1050**
data-driven modeling, **2481**
data-driven operability, **229**
data-driven optimization, **1461**
decarbonization, **1483**, **2066**
decentralized monitoring, **2406**
decision-making, **574**
decision-making under uncertainty, **1513**
decomposition, **483**, **2355**
deep eutectic solvents, **59**
deep kernel learning, **1967**
deep learning, **919**, **1679**, **2481**
deep surrogate models, **2320**
deep-q networks, **2099**
derivative free optimization, **1551**, **1761**, **2400**, **2542**
derivative-free optimization, **1558**
desalination, **886**, **1360**
desalinisation, **267**
design of experiment, **1360**
design of experiments, **828**, **1080**
design of experiments doe, **692**
design space, **648**
design under uncertainty, **567**, **583**, **611**, **1650**, **2140**, **2191**
deterministic kinetic modelling, **1277**
DETL optimization, **1455**
DEXPI, **1817**, **1891**
dielectric fluids, **1119**
differentiable programming, **978**
digital chemical engineering, **2658**, **2696**
digital skills, **2658**, **2696**
digital twin, **1833**
digital twins, **131**, **1786**
digitalisation, **131**
dimethyl ether, **1064**, **1908**
direct air capture, **297**, **492**, **1088**
direct air capture plant, **483**
direct reduced iron, **1277**
direct reduction, **314**
discrete, **2029**
distillation, **1145**, **1366**, **1776**, **2115**, **2393**, **2433**
distributed production, **2280**
distributionally robust optimization, **2170**
dividing wall column, **2433**
dividing wall columns, **1933**
documentation, **1608**
downstream bioprocessing, **640**
downstream processing, **1476**
drift diagnosis, **2527**
dual substrate growth, **776**
dye removal, **822**
dynamic modelling, **297**, **664**, **714**, **816**, **1080**, **1167**, **1226**, **1242**, **1298**, **1353**, **1398**, **1422**, **1884**, **2023**, **2393**, **2542**, **2680**
dynamic optimization, **1958**
dynamic risk assessment, **1022**
dynamic surrogate modelling, **895**
dynamical systems, **2466**
eco-industrial park, **215**
e-constraint, **2183**
e-constraint method, **2107**
education, **2600**, **2613**, **2631**, **2638**, **2658**, **2689**, **2696**
effectiveness assessment, **2607**
efficacy, **592**
electric vehicle, **432**
electricity & electrical devices, **342**, **514**, **2674**
electrochemical processes, **822**
electrochemical regeneration, **492**
electrodialysis, **492**
electrolysis, **492**, **2045**
electronic lab notebooks, **1786**
emission potential, **192**
emissions, **274**
energy, **125**, **267**, **838**, **861**, **927**, **934**, **2051**, **2674**
energy balance, **2497**
energy baseline, **1336**
energy cost, **314**
energy efficiency, **274**, **334**, **369**, **470**, **879**, **1318**, **1336**, **1431**, **1461**, **1986**, **2023**, **2051**, **2272**, **2393**, **2481**, **2631**
energy flow simulation, **432**
energy integration, **352**, **2115**
energy management, **1336**, **1558**, **2015**, **2488**, **2551**
energy planning, **100**
energy portfolio optimization, **396**
energy recovery, **1431**
energy storage, **289**, **306**, **334**, **342**, **419**, **543**, **1242**, **2045**
energy storage systems, **550**
energy supply chain, **386**
energy system model, **424**
energy systems, **100**, **282**, **320**, **334**, **369**, **378**, **413**, **446**, **514**, **1242**, **1318**, **1687**, **2029**, **2131**, **2261**, **2272**, **2362**, **2386**
energy systems analysis, **895**
energy transition, **360**, **378**, **574**, **1414**
engineering education, **2646**
ensemble learning, **964**
ensemble modelling, **1601**
environment, **45**, **94**, **152**, **183**, **274**, **927**, **2674**, **2689**

environmental impact analysis, **672**
environmental performance, **500**
enzymatic cascades, **2176**
esmr, **454**
ethanol, **1249**
ethanol dehydration, **1197**
ethanol electrooxidation, **1205**
ethylene, **2418**
exchangers, **2622**
exergy analysis, **252**
expert survey, **2652**
extraction, **108**, **192**, **838**, **2340**
falling-film evaporator, **1298**
fault detection, **1769**, **2406**, **2473**, **2518**, **2527**, **2572**
feasibility analysis, **664**
feasible path algorithm, **1974**
feedback control, **2497**
feedforward control, **2497**
fermentation, **709**, **714**, **763**, **776**, **1666**, **1801**, **2600**
fermentation time, **592**
Fischer-Tropsch synthesis, **404**
Fisher information matrix, **828**
fixed-bed reactor, **1127**
flow uniformity, **2145**
flowsheet simulation, **37**
fluid distribution device, **2145**
fluid dynamics, **1366**, **1494**, **2680**
fluidized-bed, **1398**
flushing, **2680**
food & agricultural processes, **192**
forecast uncertainty, **2441**
formaldehyde, **360**
fouling dynamics, **1298**
fouling model, **1073**
fouling resistance, **1175**
fowm, **1306**
fragrance engineering, **1713**
furfural, **1374**
furfuryl alcohol, **1374**
gaas semiconductor wastewater, **68**
gallium recovery, **68**
game theory, **462**
gams, **52**, **1183**, **2176**, **2191**
gas field, **1345**
gas filtration model, **1073**
gasification, **37**, **1398**, **1414**
gas-liquid mass transfer, **721**
gaussian processes, **986**
GCN, **1111**
generalization, **2542**
generalized disjunctive programming, **1721**
generation, **550**
generative artificial intelligence, **648**, **1713**, **2607**
genetic algorithm, **297**, **2261**
geologic hydrogen, **838**
geospatial-technoeconomic optimization, **2288**
Gflownet, **1713**
global optimization, **602**, **1721**, **1793**
global sensitivity analysis, **702**
GProms, **1336**, **1786**, **2152**, **2646**
gpu-accelerated optimization, **2243**
gradient-enhanced acceleration, **978**
granule size distribution, **742**
graph networks, **2473**
graph neural networks, **1153**, **2099**
graph reconstruction, **1737**
green ammonia, **352**, **440**, **1992**, **2045**
green ammonia production, **526**
green chemistry, **1713**
green hydrogen, **198**, **252**, **396**, **1219**, **2441**
green ironmaking, **1277**
green power, **252**
gypsum crystallization, **1833**
Haber-Bosch process, **352**
Hamilton-Jacobi reachability, **1825**
hand-crafted features, **1696**
HAZOP, **1891**
heat exchanger design, **2170**
heat exchanger networks, **419**, **1986**, **2631**
heat exchangers, **1175**
heat ingress, **328**
heat integration, **334**
heat pumps, **2272**
heat transfer, **2594**
hefa, **205**
hessian vector products, **2542**
heteroazeotropic distillation, **1728**
higher education institutes, **2613**
highways, **2288**
hot section, **1461**
human feedback, **2457**
human-in-the-loop, **1841**
hybrid, **886**
hybrid modelling, **37**, **108**, **631**, **733**, **910**, **964**, **1050**, **1285**, **1306**, **1793**
hydrogen, **169**, **244**, **259**, **267**, **314**, **328**, **352**, **369**, **386**, **507**, **879**, **927**, **1135**, **1277**, **2029**, **2036**, **2261**, **2280**, **2425**, **2551**
hydrogen flux, **1135**
hydrogen generation, **1127**
hydrogen production, **1007**
hydrogen refueling stations, **2280**, **2288**
hydrogen storage, **2045**
hydrogen supply chain, **244**
hydrothermal processing, **950**
hydrotreating processes, **2400**
identifiability analysis, **828**
immersion cooling, **1119**
Indonesia, **236**
industrial clusters, **297**
industrial ecology, **215**, **2689**
industrial process modeling, **1175**
industrial processes, **870**
industrial symbiosis, **94**, **360**, **1328**
industry 4.0, **611**, **751**, **964**, **1001**, **1058**, **1160**, **1383**, **1608**, **1817**, **2406**, **2578**, **2600**, **2658**, **2696**
industry 50, **783**
information management, **783**, **1608**, **1650**
instance segmentation, **1696**
integrated energy systems, **2243**
integrated facility, **31**
intelligent systems, **386**, **618**, **1746**
interdisciplinary, **1650**
interpretable ai, **903**
interpretable machine learning, **1175**
interpretable model construction, **1529**
inverse design, **648**
inverse molecular design, **903**
ion exchange, **68**
iridium, **169**
Jacobian, **611**
jax, **978**, **1776**
Julia language, **1439**
kernels, **1754**
key variables, **2272**
kinetic modelling, **828**
k-nearest neighbours, **2449**
knowledge graphs, **783**, **1608**
Kolmogorov Arnold network, **1967**
laboratory safety monitoring, **1503**
large language model, **1769**, **2473**, **2600**, **2622**
life cycle assessment, **18**, **45**, **77**, **94**, **117**, **125**, **143**, **175**, **222**, **404**, **413**, **477**, **492**, **791**, **943**, **1212**, **1592**, **1847**

life-cycle emissions, **1943**
liquid-phase synthesis, **664**
liquified natural gas, **1461**
LNG, **1345**
LNG optimization, **1461**
local optimization, **1721**
low-carbon hydrogen, **2288**
lycopene, **2008**
machine learning, **37**, **131**, **183**, **592**, **751**, **810**, **855**, **895**, **910**, **986**, **1058**, **1145**, **1153**, **1256**, **1285**, **1353**, **1520**, **1558**, **1592**, **1634**, **1659**, **1666**, **1728**, **1746**, **1761**, **1801**, **1809**, **1841**, **1847**, **1856**, **1899**, **1908**, **1925**, **1974**, **2091**, **2473**, **2518**, **2542**, **2565**, **2572**, **2578**
machine learning-based optimisation, **1793**
magnesium chloride, **1197**
markov decision process, **2099**
mass balance, **2497**
material flow analysis, **169**
material screening, **2272**
material synthesis, **2666**
materials, **125**, **763**, **1445**, **1841**, **1899**
mathematical model, **1007**
mathematical optimization, **31**
mathematical programming, **2377**
MATLAB, **1064**, **1950**, **2023**, **2646**
matrix effects, **1205**
matrix factorization, **1754**
McCormick relaxation, **2201**
MCSGP, **2449**
mean residence time, **2145**
mechanistic constraints, **910**
mechanistic modelling, **978**
membrane cascade, **664**
membrane systems, **2297**
membranes, **846**, **1135**, **2303**
mesh electrode, **534**
metabolic engineering, **1102**
metabolic models, **709**
metaheuristic optimization, **1933**
metal hydrides, **2425**
metal organic framework, **802**
methanation, **1391**
methanation process, **483**
methanol, **360**, **404**, **1064**, **1249**, **2393**
methanol synthesis, **828**
methylcyclohexane, **879**
microalgae, **236**
microbial growth, **810**
microfluidics, **1696**
microreactor systems, **1064**
microstructure segmentation, **1841**
microwave-assisted heating, **11**
MIP-heuristics, **1761**
MIGCAP, **483**
mixed integer bilinear problem, **244**
mixed integer linear programming, **198**, **205**, **477**, **550**, **2161**, **2183**, **2288**, **2303**, **2349**, **2355**, **2362**
mixed integer nonlinear programming, **1721**, **1916**, **2058**, **2152**, **2272**
mixed lignocellulosic biomass, **117**
mixed solvents, **1679**
mixing, **1494**
mixing cell network, **1391**
MLOPS, **870**
model based design of experiments, **1529**
model failure, **870**
model integrity, **1705**
model order reduction, **2594**
model predictive control, **131**, **1582**, **2425**, **2457**, **2594**
model reduction, **1014**
model-based design of experiments, **108**, **721**
model-based design of experiments, **1080**
modeling to generate alternatives, **574**
modelling, **94**, **125**, **252**, **625**, **721**, **763**, **846**, **934**, **1058**, **1226**, **1285**, **1383**, **1445**, **1476**, **1601**, **1608**, **1801**, **1908**, **2001**, **2209**, **2280**, **2572**
modelling and simulations, **161**, **222**, **274**, **306**, **413**, **543**, **611**, **618**, **714**, **776**, **791**, **802**, **810**, **838**, **861**, **943**, **992**, **1014**, **1030**, **1043**, **1095**, **1135**, **1145**, **1183**, **1226**, **1262**, **1285**, **1318**, **1353**, **1360**, **1366**, **1406**, **1467**, **1520**, **1565**, **1650**, **1687**, **1847**, **1856**, **2023**, **2036**, **2051**, **2218**, **2226**, **2386**, **2406**, **2565**, **2578**, **2658**, **2696**
modular electrolysis, **2652**
modular heterogeneous systems, **2029**
modular plants, **1608**, **1817**, **1891**
modularisation, **1439**
modularization, **2652**
molecular design, **681**, **2066**
molecular generation, **1713**
molecular representation, **2252**
monoclonal antibodies, **640**
monoethanolamine, **1269**
multi effect distillation, **886**
multi-agent approaches, **2689**
multi-armed bandit, **574**
multi-criteria, **2029**
multi-fidelity Bayesian optimisation, **2320**
multimodal large language models, **1737**
multi-modality, **1876**
multi-objective optimization, **500**, **1197**, **2123**, **2183**, **2234**, **2340**
multi-parametric programming, **567**, **2425**
multi-physics model, **534**
multiscale modelling, **18**, **282**, **386**, **583**, **709**, **1043**, **1189**, **1353**, **1565**, **1634**, **1884**
multi-scenario optimization, **2303**
multi-source heterogeneous data, **1022**
multistep distillation, **2334**
multivariate statistics, **1539**
nanoparticles, **1445**, **2666**
natural gas, **1366**, **1461**, **1565**, **2107**
near infrared spectroscopy, **751**
negative emission technologies, **972**
negative emissions, **244**
network topology, **2406**
net-zero, **446**
neural network, **1050**, **1958**
neural networks, **1306**, **2505**
nitrogen, **236**
NLP, **483**, **2176**
nominal operating range nor, **2587**
nonconvex optimization, **1721**
nonconvex robust optimization, **602**
non-ideal thermodynamics, **1833**
nonlinear model predictive control, **1884**, **2518**, **2558**
nonlinear programming, **2243**
nonlinear solvers, **1776**
NRTL model, **1234**
non-dominated sorting genetic algorithm II, **2123**
nuclear, **342**
numerical methods, **721**, **1422**, **1617**
odorant design, **1713**

offshore gas compression, **2534**
oil and gas, **1306**
oil refinery, **2505**
oligonucleotide synthesis, **664**
one-shot forecasting, **2449**
on-line, **2433**
onsite production, **2288**
ontology, **1608**
open source software, **1439**
OpenLCA, **943**
open-source framework, **131**
operational flexibility, **2587**
operational optimization, **483**, **2201**
optimal control, **1825**
optimal dose regimen, **769**
optimal experiment design, **1102**
optimal stopping, **1513**
optimization, **52**, **143**, **215**, **259**, **267**, **289**, **306**, **320**, **334**, **369**, **424**, **462**, **507**, **692**, **709**, **721**, **810**, **838**, **1001**, **1007**, **1043**, **1095**, **1135**, **1160**, **1262**, **1285**, **1383**, **1461**, **1467**, **1573**, **1617**, **1643**, **1650**, **1687**, **1705**, **1776**, **1884**, **1908**, **1943**, **1950**, **1974**, **2015**, **2029**, **2051**, **2066**, **2082**, **2091**, **2115**, **2140**, **2152**, **2191**, **2209**, **2226**, **2234**, **2261**, **2272**, **2297**, **2303**, **2311**, **2349**, **2355**, **2362**, **2386**, **2433**, **2466**, **2505**, **2511**, **2558**, **2578**, **2631**
optimization under uncertainty, **2170**
optogenetic control, **702**
orange peel waste, **2377**
organic solvent nanofiltration, **664**
organics recovery from wastewater, **2334**
output curtailment, **432**
P&ID, **1891**
P&ID digitisation, **1737**
palm oil mill effluent, **175**
paper industry, **1073**
parallel channels, **2145**
parallelization, **1617**, **1643**
parameter estimation, **625**, **1080**
pareto front, **692**
partial differential equations, **2466**
partially observed systems, **1406**
particle swarm optimization, **846**, **1950**, **2272**
PAT, **655**
PBM, **769**
PBPK, **769**
PEM electrolysis, **2551**
PEM electrolyzer, **483**
peptides, **1476**
performance evaluation, **1088**
performance indexes, **1374**
periodic bioprocessing, **2449**
perturb & observe, **2433**
petroleum, **2400**
p-graph framework, **2058**
pharmaceutical formulations, **1292**
pharmaceutical manufacturing, **618**, **648**, **727**, **756**
pharmaceutical tablets, **751**
pharmaceuticals, **625**
physics-consistent hybrid modelling, **2123**
physics-informed, **1153**
physics-informed neural networks, **229**, **751**, **2594**
PI control, **2534**
piecewise linearization, **2201**
pilot plant data, **1728**
pilot-scale, **1050**
planetary boundary, **424**
planning, **2131**
planning & scheduling, **583**, **2303**, **2327**
plant scale up, **440**
plant start-up, **1825**
plasma process, **2578**
plastic recycling, **183**
plastic waste, **950**
plastic waste pyrolysis, **11**
plastics recycling, **229**
PLS, **1601**
PLSR, **655**
polymaths, **2646**
polymers, **45**, **1043**, **1899**, **2565**
polynomial regression, **1205**
polyolefin production, **2226**
polyphenols, **192**
population balance, **1242**
population balance modeling, **1833**
population balances, **625**
post combustion carbon capture, **1269**
power generation, **2327**
power grid, **396**
power systems modeling, **550**
power-to-ammonia, **2311**
power-to-hydrogen, **2045**, **2243**, **2441**
power-to-liquid, **404**
practical identifiability, **1102**
prandtl number, **1119**
predictive maintenance, **2527**
pregabalin, **672**
preparative chromatography, **1476**
preprocessing, **655**
pressure drop, **2115**
pressure loss constraint, **2145**
principal component analysis, **236**, **1958**
probabilistic design space, **756**, **2587**
process calculations, **2613**
process control, **2418**, **2433**, **2511**, **2558**, **2572**
process design, **11**, **18**, **25**, **152**, **267**, **274**, **282**, **297**, **342**, **404**, **462**, **500**, **611**, **714**, **727**, **816**, **846**, **1043**, **1145**, **1167**, **1455**, **1494**, **1565**, **1650**, **1659**, **1687**, **1746**, **1786**, **1793**, **1809**, **1817**, **1992**, **2023**, **2036**, **2066**, **2082**, **2091**, **2140**, **2152**, **2218**, **2311**, **2631**, **2638**, **2680**
process design and modelling, **2334**
process electrification, **1825**
process equipment assemblies, **2652**
process family design, **1916**
process integration, **205**, **1328**, **1986**
process intensification, **1366**, **1494**, **1809**, **2091**, **2340**
process modelling, **1088**, **1925**
process monitoring, **655**, **733**, **1058**, **1666**, **1856**, **1866**, **2518**, **2527**
process operations, **222**, **514**, **1001**, **1285**, **1383**, **1643**, **1992**, **2015**, **2107**, **2140**, **2209**, **2400**
process optimization, **77**, **440**, **1088**, **1925**, **1967**, **2320**, **2334**, **2481**, **2551**
process safety, **2425**
process simulation, **558**, **972**, **1439**, **1776**, **2481**
process simulator-based optimization, **2058**
process synthesis, **1064**, **1659**, **1746**, **1809**, **2058**, **2074**, **2091**, **2152**
progressive hedging, **1916**
prompt, **2622**
property prediction, **681**, **1634**
proton exchange membrane water electrolyser, **169**

prototype learning, **1841**
proximal policy optimization, **2587**
PSA, **2234**
purification, **1476**
Pyomo, **306**, **1383**, **2001**
pyramidal pins, **534**
pyrolysis, **229**, **1234**
python, **1197**, **1776**, **2646**
quality by digital design, **648**, **791**
quantum computing, **1659**
quaternary mixture, **1933**
radial flow reactor, **1391**
Raman spectroscopy, **1866**
random forest regression, **592**
reaction, **1064**, **1539**
reaction engineering, **183**, **1135**, **1551**, **2176**
reaction network, **1014**
reaction network identification, **1529**
reaction optimization, **2252**
reaction patterns, **1539**
reaction yield, **919**
reactive distillation, **25**
reactor network, **2176**
real-time dynamic optimization, **2441**
receding horizon, **2441**
recipes, **1336**
recurrent neural networks, **895**
recycling, **950**
reduced order models, **742**
refining, **342**, **514**
reinforcement learning, **1558**, **1659**, **2099**, **2226**, **2457**, **2534**, **2587**
reliability quantification, **681**
renewable and sustainable energy, **100**, **161**, **267**, **289**, **320**, **419**, **507**, **727**, **1992**, **2261**, **2674**
renewable energy, **1483**
renewable hydrogen system, **2201**, **2488**
requirements prioritization, **2652**
resilience, **386**, **583**, **2311**
resource allocation, **958**
resource competition, **702**
resource exchanges, **215**
resource recovery, **2**
retrofit, **2045**
reverse osmosis, **2303**
ribbon splitting, **742**
rigorous modelling, **440**
risk assessment, **2607**
robotic, **2666**
robust design optimization, **558**
robust optimisation, **1793**
robust optimization, **2170**
robustness, **1080**
roller compaction, **742**
rolling energy consumption, **2123**
rolling horizon approach, **2488**
rolling horizon optimization, **2001**, **2036**, **2370**
rolling-horizon, **1958**
safe-and-sustainable-by-design, **791**
safety analysis, **1891**
saft, **2082**
saft-γ MIE, **77**
scheduling, **369**, **1383**, **1958**, **2001**, **2015**, **2036**, **2209**, **2226**, **2370**
scientific machine learning, **1306**
SCM, **1277**
seawater desalination, **2074**
secondary metabolite concentration, **592**
segment anything model, **1841**
segmented flow, **108**
semi-infinite programming, **1617**
semi-supervised learning, **1111**
sensitivity analysis, **267**, **424**, **2349**
sensitivity study, **1127**
separation processes, **1050**
sewage sludge, **37**
Sfiles, **2074**
shrinking-core model for dissolution, **1833**
Simapro, **94**, **222**, **1847**
similarity, **681**
simulated annealing, **1262**
simulated moving bed, **2466**
simulation, **709**, **727**, **816**, **950**, **1001**, **1242**, **1249**, **1345**, **1422**, **1445**, **1626**, **1687**, **1705**, **1786**, **1884**, **1950**, **2023**, **2370**, **2666**
simulation-optimization, **252**, **2327**
single cell protein, **776**
slope-based correction, **910**
small modular reactors, **396**
sng production, **526**
social sustainability, **259**
society 50, **2613**
soft sensing, **1601**
soft sensor, **733**, **2497**
software, **1884**
Sohar freezone, **2161**
solar energy, **886**
solar thermal, **419**
solid oxide electrolysis, **352**
solubility, **1679**
solvation free energy of reaction, **1876**
solvent design, **77**
solvent effect, **1876**
solvent extraction, **68**, **2008**
solvent selection, **672**, **2082**
solvent-based plastic recycling, **77**
sonochemistry, **1467**
sorbent-enhanced biogas reforming, **252**
sorption-enhanced, **1391**
space visualization, **2074**
space-filling designs, **986**
statistical design of experiments doe, **1573**
steam methane reforming, **1219**
steam reforming, **314**
stirred tank reactors, **721**
stochastic approach, **958**
stochastic differential equations, **1513**
stochastic optimization, **567**, **583**, **2107**, **2131**
stochastic programming, **550**, **2170**
stochastic simulations, **183**
stoichiometric number, **454**
structural causal models, **2406**
structural identifiability & observability analysis, **1406**
sufficiency, **424**
superheated steam filtration, **1073**
supersaturation, **2497**
superstructure, **950**, **2058**, **2115**, **2297**
superstructure optimization, **244**, **360**, **2152**, **2161**, **2340**
supervisory control, **2534**
supply chain, **52**, **198**, **259**, **507**, **567**, **583**, **618**, **783**, **1262**, **2001**, **2191**, **2386**
supply chain optimization, **477**, **2377**
support vector regression, **592**
surface tension, **1111**
surfactants, **903**, **1111**
surrogate modelling, **229**, **514**, **756**, **1095**, **1145**, **1197**, **1336**, **1353**, **1626**, **1761**, **1847**, **1856**, **1958**, **1974**, **2131**, **2218**, **2234**
sustainability, **446**, **1713**, **2183**, **2311**

sustainability promotion, **1705**
sustainable aviation fuel, **205**, **378**, **462**
sustainable power generation, **477**
sustainable sugar–ethanol–energy systems, **1212**
symbolic regression, **855**, **1175**, **1406**, **1529**, **1592**, **2234**
syngas, **1398**, **1908**
synthetic fuels, **404**
system identification, **1226**, **1360**, **1520**, **2542**, **2680**
systematic model development, **1406**
tableting, **742**
tank motion, **328**
Taylor vortex flow reactor, **1285**
TCD, **1277**
technical assessment, **972**
technoeconomic analysis, **11**, **18**, **59**, **68**, **77**, **152**, **161**, **192**, **252**, **289**, **360**, **404**, **413**, **470**, **477**, **500**, **526**, **672**, **791**, **943**, **1183**, **2023**, **2280**, **2386**
technology adoption, **2646**
temperature vacuum swing adsorption, **1088**
temperature-dependent property prediction, **910**
temporal weather aggregation, **297**
Tennessee Eastman process data, **1728**
thermal systems, **2594**
thermodynamic consistency, **1153**
time-to-css forecasting, **2449**
toluene, **879**
tomato waste, **2008**
topology optimization, **2145**
torrefaction, **117**
transesterification, **1866**
transfer learning, **919**, **2320**
transformers, **903**
transient experiments, **1080**
transition state, **1876**
transmission, **550**
turbines, **1431**
ultrasound-assisted process intensification, **1866**
uncertain parameter, **958**
uncertainty, **602**, **1262**, **2107**
uncertainty assessment, **558**
uncertainty quantification, **131**, **992**, **1841**
uncertainty-aware, **631**
unconventional solvents, **2008**
unexpected events, **574**
United Kingdom, **297**
unsupervised learning, **2473**
urea electrosynthesis, **18**
vaccine supply chain, **2183**
VAE, **1111**
value-added products, **31**
value-based assessment, **175**
vanillin, **1713**
variable renewable energy, **432**
variational inference, **992**
vehicle routing, **2355**
vision-language model, **1503**
visualisation, **2666**
vle model re-parametrisation, **1234**
waste cooking oil, **25**
waste heat, **1328**
waste heat utilization, **483**, **1483**
waste valorisation, **1414**, **2334**
waste valorization, **1328**
waste-to-energy, **477**
wastewater, **846**, **1226**
wastewater valorization, **175**
water, **838**, **2303**
water and energy integration systems, **1705**
water networks, **2303**
water phase valorisation, **1234**
wet-process phosphoric acid, **1833**
when, **1986**
whole systems thinking, **198**
wind, **2674**
wind energy, **886**
wine effluents, **59**
wiped film evaporator, **1292**
work exchanger networks, **1986**
work group, **2689**
workflow automation platforms, **2607**
zero gap cell, **1007**
zero liquid, **2**
zero liquid discharge, **2**
zero-gap cell, **534**

Editorial Notes

Follow the DOI of each individual article for the latest updates, links to supplementary material and related works, and high-resolution versions of the article.

Some articles were withdrawn or retracted during the publication process. Their pages have been removed from this book.

PSEcommunity.org

PSE Press • Hamilton • Notre Dame